THE JOURNAL OF NAVIGATION (2006), **59**, 71–89.
Published by the Royal Institute of Navigation with permission from QinetiQ Ltd.
doi:10.1017/S0373463305003462 Printed in the United Kingdom

Optimising the Integration of Terrain Referenced Navigation with INS and GPS

Paul D Groves and Robin J Handley

(*QinetiQ ltd, Farnborough*)
(Email: pdgroves@qinetiq.com)

Andrew R Runnalls

(*Data Fusion Research Limited and University of Kent*)

The benefits of integrated INS/GPS systems are well known. However, the knowledge required to jam GPS is becoming public and the hardware to achieve this is basic. When GPS data are unavailable and a low grade INS is used, navigation accuracy quickly degrades to an unacceptable level. The addition of one or more terrain referenced navigation (TRN) systems to an integrated INS/GPS navigation system enables the INS to be calibrated during GPS outages, increasing the robustness of the overall navigation solution. TRN techniques are compared and integration architectures are reviewed. For the initial studies of INS/GPS/TRN integration, radar altimeter based terrain contour navigation (TCN) with a batch processing algorithm is used in conjunction with a centralised integration filter. Four different approaches for using these TCN fixes to calibrate the INS are compared. These are a best fix method, a weighted fix method using a probabilistic data association filter (PDAF) and single and multi-hypothesis versions of the Iterative Gaussian Mixture Approximation of the Posterior (IGMAP) method. Simulation results are presented showing that the single hypothesis IGMAP technique offers the best balance between accuracy, robustness and processing efficiency.

KEY WORDS

1. TRN. 2. GPS. 3. INS

1. INTRODUCTION. An inertial navigation system (INS) operates continuously (bar hardware faults) and provides a high bandwidth (>50 Hz) output with low short term noise. It also provides effective attitude, angular rate and acceleration measurements as well as position and velocity. However, its navigation accuracy degrades with time as the noise and biases on its inertial instrument outputs are mathematically integrated through the navigation equations that generate the final output.

The Global Positioning System (GPS), and other satellite navigation systems, provide a high accuracy ($\sim$5 m) position solution that does not degrade with time.

The GPS navigation solution is noisier than that of an INS, has a lower bandwidth (~1 Hz) and does not normally include attitude. GPS and INS are thus complementary. Consequently, many aircraft and guided weapons use an integrated INS/GPS navigation system. The INS provides the core navigation solution, whilst the GPS measurements are used to correct and calibrate the INS via an integration algorithm.

However, satellite navigation signals are extremely vulnerable to interference. Unintentional interference sources include broadcast television, mobile satellite services, ultra-wide-band communications, over-the-horizon radar and cellular telephones (Carroll *et al.*, 2001). In military applications, deliberate jamming is highly likely, and must be planned for. Interference can be mitigated using a controlled reception pattern antenna (CRPA) system (e.g. Owen and Wells, 2001; Wells and Owen, 2002; Boasman and Briggs, 2002), together with advanced INS/GPS integration techniques such as adaptive tightly-coupled (ATC) (Groves and Long, 2003; 2005) and deep integration (e.g. Sennott and Senffner, 1997; Gustafson *et al*, 2000; Soloviev *et al*, 2004). These techniques enable satellite navigation signals to be tracked under higher levels of interference. However, they do not eliminate the effects of jamming and interference completely. The cost and complexity of jamming technology that can defeat them is significantly less than that of the CRPA systems themselves and is being communicated across the internet!

As soon as GPS measurements are lost, the INS begins to drift out of calibration. Aircraft grade INS can maintain a horizontal position accuracy within 100 m through GPS outages of more than 10 minutes. However, the lower cost INS common in guided weapons, unmanned air vehicles and general aviation (private) aircraft can only maintain this accuracy for 2 to 3 minutes. To attain robust navigation in a GPS jamming environment, reversionary navigation systems are required. Terrain referenced navigation (TRN) techniques offer a solution.

The most established TRN technique, terrain contour navigation (TCN), uses measurements from a radio altimeter (radalt) and requires undulating terrain. Performance may be enhanced by using a laser range-finder as the sensor. A second, and complementary, technique is scene/line feature matching, which uses a dedicated imaging sensor. The current state-of-the-art in line feature matching systems is represented by the Continuous Visual Navigation (CVN) system (Handley *et al.* 2001). TRN techniques are reviewed in Section 2.

The focus of this paper is the integration of TRN with INS and GPS. As with INS/GPS integration, TRN is usually integrated with INS and GPS using techniques based on the Kalman filter. Selecting the integration architecture to make the best use of information from the different sensors is discussed in Section 3. How best to handle potential false or ambiguous fixes from a TRN system within the integration filter is covered in the succeeding sections.

For the initial studies of INS/GPS/TRN integration, radalt based TCN was selected because a relatively simple simulation can be used and the research is readily applicable to current systems. A batch processing TCN algorithm is used, whereby a series of radalt measurements is compared with a terrain height database to produce one or more position fix hypotheses, each with an associated covariance and probability. To determine how best to use these TCN fixes to calibrate the INS whilst obtaining the optimum balance of accuracy, robustness and processing efficiency, four different approaches are compared. Section 4 describes a best fix method and a weighted fix method using a probabilistic data association filter (PDAF). The single and multi-hypothesis versions of the Iterative Gaussian Mixture Approximation of the Posterior (IGMAP) method are described in Section 5.

The simulation environment used for comparing the different INS/TCN integration architectures is described in Section 6, with emphasis on the radar altimeter simulation. Simulation results comparing the performance of the best fix, PDAF and two IGMAP methods are presented in Section 7. Section 8 discusses the effect of radalt and terrain height database errors on performance. Where strong GPS signals are available, GPS will always be more accurate than TRN. However, under moderate jamming or where the number of GPS signals available is limited, there can be benefits in using TRN alongside GPS. Section 9 presents the results of simulations to determine the benefit of combining TCN with adaptive tightly-coupled anti-jam INS/GPS integration.

2. REVIEW OF TRN TECHNIQUES.

Development of terrain contour navigation (TCN) started in the 1970s and a number of systems have been commercially available since before GPS was fully operational. Conventionally, such systems estimate the height of the terrain below the air vehicle by subtracting radio altimeter height from [Barometric/]INS altitude. Measurements are typically taken around once a second. These are then compared with a terrain height database, such as Digital Terrain Elevation Data (DTED) (Defense Mapping Agency, 1983). A range of different processing techniques have been developed to obtain position fixes from the comparisons of measured and database terrain heights (Metzger *et al*, 2002). These may be divided into two broad categories: sequential and batch.

In *sequential* processing, each measurement is processed separately. The difference between the radalt generated and database indicated terrain height is input as a measurement to a Kalman filter, noting that the database height is obtained at the current best estimate of position. The terrain gradient is then used to attribute this to a linear combination of the latitude, longitude and height components of the aided INS position error. To obtain independent estimates of the three components of the position error, at least three radalt measurements are required, usually more in practice. Sequential processing is well established in commercial TCN systems such as BAE Systems' TERPROM (Robins, 1998) and the American SITAN (Hollowell, 1990). The principal advantage of the sequential approach is relative simplicity and comparatively low processor load. However, it relies on accurate knowledge of the terrain gradient below the aircraft, which is a demanding requirement on existing low resolution, low accuracy databases like DTED. To a certain extent, the limitations of terrain height databases may be overcome by using sophisticated linearisation algorithms (Yu *et al*, 1991). However, a fundamental problem remains in that the gradient is calculated below the navigation system indicated aircraft position, not its true position. Thus, if the horizontal position error exceeds about 250 m, main-stream sequential processing does not work and a 'recovery' mode must be instigated, for example batch processing or a parallel solutions approach such as Multiple Model Adaptive Estimation (Hollowell, 1990). An alternative approach is to process the radalt measurements with a non-linear filter, such as a particle filter (Metzger and Trommer, 2003).

In *batch* processing, a series of typically 5–16 terrain height measurements, known as a transect, are processed together (Runnalls, 1985). The transect is fitted to the terrain height database at different offsets in latitude and longitude from the current estimated position. The residuals of each fit are used to calculate a likelihood at each point in the grid, producing a likelihood surface as the output of the matching process. The simplest way of obtaining a position fix from the likelihood surface is to take the highest point. However, the likelihood surface tends to be noisy, so this does not provide a good position estimate (Metzger *et al*, 2002). Another straightforward approach is to fit a bivariate Gaussian distribution to the likelihood surface. However, the likelihood surface is often multi-modal, i.e. it has more than one peak. Another method is to multiply successive likelihood surfaces until a single dominant peak is attained (Carlbom and Johansson, 2002). An optimal approach is to use Monte-Carlo Markov Chain methods with a Bayesian network based data fusion algorithm (Runnalls and Handley, 1998); however this imposes an extreme processing load. A suitable compromise is to fit multiple Gaussian distributions (typically up to five) to the likelihood surface, providing a multiple hypothesis position fix for processing by the data fusion algorithm. This approach forms the basis of the TCN algorithm used for the studies presented here.

Conventional, radalt based, TCN systems are significantly less accurate than civil GPS. Although there is scope for improving TCN performance through using more sophisticated processing, accuracy can be enhanced by using a more precise height sensor. Commercially available Laser Line Scanners (LLS), such as those from Optech, provide many range measurements from the air vehicle to the terrain, to an accuracy of centimetres, in a fraction of a second. Unlike a radalt, the LLS is able to acquire a sequence of terrain elevation profiles by rapidly scanning the terrain surface. By combining these profiles with the forward motion of the air vehicle, as indicated by the INS, the LLS can construct elevation measurements over a two-dimensional surface, where a radalt can only provide a one-dimensional elevation strip. This increased coverage, coupled with the greater inherent accuracy of the LLS, holds the promise of improved accuracy and greater operational flexibility. LLS are already in use for obstacle avoidance and target detection. Consequently they will be increasingly common on future airborne platforms.

A Laser Line scan Navigation (LLN) system has been developed by Hi-Q Systems and QinetiQ (Handley *et al*, 2003; McNeil *et al*, 2002). The LLN algorithm takes a series of range measurements from the LLS and performs co-ordinate transformations to obtain a terrain surface measurement. This is then compared with the region of the DTED terrain height database corresponding to the estimated position using an optimised surface correlation algorithm. This process outputs the difference between the estimated and measured terrain surfaces which may then be used to update the navigation filter. Using recorded LLS data, navigation accuracies in the 30 to 45 m range, depending on the grade of the associated INS, have been demonstrated using LLN. A similar laser navigation system has been developed by Ohio University (Campbell *et al*, 2003; 2004).

Terrain contour matching techniques rely on the flight profile overflying undulating terrain in order to provide useful navigation data. In order to accommodate this requirement, mission planners need to work hard to find acceptable routes, with a significant probability of not succeeding. By contrast, scene and linear feature matching systems work optimally over flat terrain.

Linear feature matching techniques usually work by linking observed linear features with their representation within a linear feature database, and can provide comparable accuracy to civil GPS over favourable terrain. Linear feature matching is complementary to TCN because linear features:

- commonly appear wherever there is development by man (which predominantly occurs, for practical reasons, on flat terrain where TCN is ineffective and ground-based GPS jammers are most effective at jamming low flying aircraft!);
- are usable even when flying very low;
- can be detected by simple and well established image processing techniques (although complex techniques can produce better results);
- are easily extracted from satellite and photographic imagery (and consequently database production is fairly quick and inexpensive and can incorporate valuable target information).

Continuous Visual Navigation (CVN) (Handley *et al*, 1998; 2001) was jointly conceived by QinetiQ and Hi-Q Systems. The goal was a linear feature matching technique that is not reliant upon linear feature intersections, requires minimal mission planning, imposes a relatively low computational load, is cheap to implement and can cope with ambiguous linear features. The CVN system comprises a video or infra-red camera, pointing downwards, to provide images of the terrain below the host vehicle, a database of linear features covering the area overflown and a processor implementing the CVN algorithm. It also uses the host vehicle's radalt to determine the height above ground for scaling the image and the INS to determine the spatial separation of successive images and the area of the database to search.

CVN grabs a digitised image from the imaging sensor whenever line features are predicted using the database and current navigation solution, typically around once a second. Standard imaging processing techniques are then used to extract linear features from the image. These features are then compared with both the predicted line features from the database and features over a wider search area to produce measurements to update the navigation filter. In the event of ambiguous matches, parallel Kalman filters are maintained, each with a likelihood score, until the ambiguity is resolved from subsequent images.

Over favourably-featured terrain, CVN exhibits a navigation accuracy of around 10 m with an aircraft grade (SNU84) INS and around 20 m with a lower quality tactical grade INS, such as an LN200. CVN has been successfully demonstrated in real-time flight trials (Handley *et al*, 2001). Employing separate sensors for CVN and LLN imposes a significant burden on the host vehicle, in terms of the financial cost of acquiring separate sensors and their increased mass, power consumption, supporting electronics requirements and impact on the design of the host airframe. However, with some modifications to the processing algorithm, the LLS can be used to capture images for CVN. Furthermore, use of a common sensor enables the CVN and LLN algorithms to be integrated, reducing the occurrence of false and ambiguous matches (Handley *et al*, 2003).

3. INS/GPS/TRN INTEGRATION ARCHITECTURE. As with INS/GPS integration, TRN is usually integrated with INS and GPS using Kalman filter

based techniques. Selection of the integration architecture is a trade-off between a number of sometimes conflicting requirements:

- maximisation of the navigation accuracy for most of the time, which is met by cross-calibration between the navigation sensors;
- preventing cross-contamination of errors between subsystems;
- maximising the availability of information for integrity monitoring;
- minimising algorithm complexity and processor load.

The three main architectures for integrating multiple sensor navigation systems are federated, cascaded and centralised. In a federated filters architecture, the INS is integrated with each of the other navigation systems using a separate Kalman filter, known as a local filter, each of which produces a navigation solution (Carlson, 1990). The integrated navigation solution is obtained simply by fusing the solutions from the Kalman filters, weighted according to their uncertainties. The advantages of this architecture are ease of combining equipment from different manufacturers and integrity. The integrity benefits arise because the contamination of one local filter by false data does not affect the others and, with three or more filters, they can be used to monitor each other. The main disadvantages are that there is no cross-calibration between GPS and TRN and, because the correlations between the Kalman filters are ignored, the uncertainty of the fused navigation solution tends to be over-optimistic.

A cascaded architecture is similar to a federated architecture, except that a master Kalman filter is used to combine the outputs of the local filters. This enables some cross-calibration between the local filters. However, in any cascaded architecture, low gains must be used in the master Kalman filter to prevent destabilisation due to correlated noise on the local filter outputs. Such an architecture can be hazardous where information about the local filters is not available due to commercial confidentiality.

In a centralised architecture, shown in Figure 1, measurements from all sub-systems are processed by a single Kalman filter. For example the GPS integration is tightly-coupled, inputting pseudo-range and pseudo-range rate to the integration filter. With all sensor errors modelled centrally, full cross calibration between the GPS and TRN systems can be obtained. There are also no potential stability problems arising from un-modelled cross-coupling or cascading Kalman filters. Thus, a centralised Kalman filter should provide the most accurate integrated navigation solution. However, the effectiveness of this architecture is contingent upon representative models of the sensor errors, making co-operation with the sensor manufacturers essential. Integrity monitoring can be implemented using parallel filter techniques (Moore *et al*, 2001). A centralised integration architecture has been selected for the work presented here.

Because TRN systems operate by comparing a series of measurements or an image of the terrain below the host vehicle with a database, they are liable to produce occasional false position fixes, either as a result of errors in the database, or because the measurements match the database in more than one location. A lot of TRN systems, such as CVN and many batch-processed TCN systems, will also output multiple hypothesis position fixes where there is not a unique match between measurements and database. Each candidate is accompanied by an associated

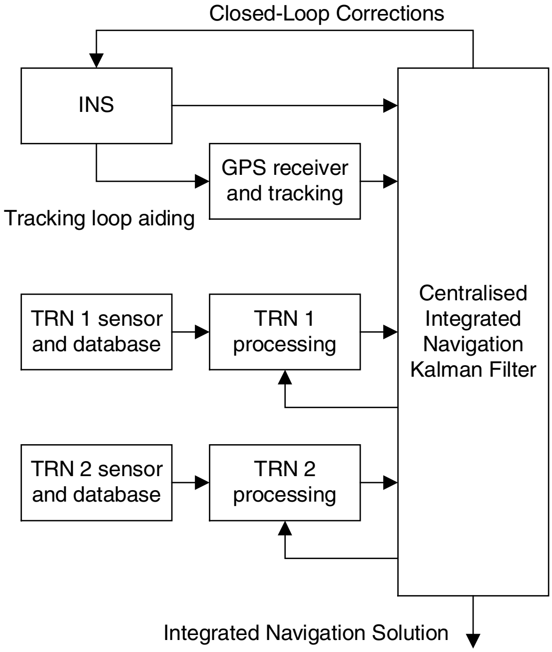

Figure 1. Centralised integrated navigation architecture.

covariance and an estimated probability. There are three main ways in which such ambiguous fixes may be handled by a navigation filter: best fix, weighted fix and parallel filters. The following sections describe and compare a best fix method, PDAF and IGMAP weighted fix methods and the parallel filters IGMAP method for INS/TCN integration.

4. BEST FIX AND PDAF INS/TCN INTEGRATION.

The best fix and PDAF INS/TCN integration strategies have both been tested in conjunction with a Gaussian clumping TCN algorithm. This is a batch processing TCN technique, as described above in Section 2. To obtain position fixes from the likelihood surface, the likelihoods are first converted to probabilities by dividing them by the total likelihood summed across the surface. Grid points above a noise floor are then grouped into contiguous clumps. The probabilities of the points in each clump are summed and the lowest probability clumps are merged with their nearest neighbours where necessary. From each clump, a mean position fix is obtained by calculating the weighted average of the points within the clump. Each point in the likelihood surface has a height error associated with it, so the fix is 3-dimensional. To indicate the uncertainty of the position fixes, a 3×3 covariance matrix is computed for each clump from the distribution of points within the clump and the estimated DTED error standard deviations.

The best fix method is the simplest method of handling uncertain TCN measurements in a Kalman filter. The integration algorithm processes the fix with the highest probability score and rejects the others. To minimise contamination of the

https://doi.org/10.1017/S0373463305003462 Published online by Cambridge University Press

integration filter by false fixes, measurements should only be accepted where there is reasonable confidence in the best fix. This is done by setting two thresholds: a minimum probability score for the best fix and a maximum probability score for the second best fix. Low thresholds enable more TCN information to be processed in the integration filter, whereas high thresholds provide more protection against false fixes. The choice of threshold depends on the accuracy and robustness requirements, INS grade, integrity monitoring system and the probability of false TCN fixes being produced with high scores.

The probabilistic data association filter inputs all the candidate measurements to a single integration filter, weighted according to their probabilities. An extra term is applied in the covariance update to account for the spread of fix hypotheses. The system update phase of the PDAF is the same as for a standard Kalman filter. An advantage of the PDAF approach is that all true fixes are input to the integration filter. The main disadvantage is that all of the false fixes are also input, degrading the integrated navigation solution. However, provided the probability scores determined by the TRN system are reasonably representative, the false fixes will not destabilise the filter as the covariance of the PDAF estimates will remain representative.

For both the best fix and PDAF integration algorithms, closed loop correction of the INS errors was implemented to minimise linearisation errors in the integration filter, which was based on the ATC INS/GPS integration algorithm set (Groves and Long, 2003; 2005). Accelerometer and gyro biases were estimated in addition to position, velocity and attitude. In addition, measurement pre-filtering was implemented to reject those TCN fix hypotheses that were widely inconsistent with the integration filter's state estimates.

5. IGMAP INS/TCN INTEGRATION.

The IGMAP (Iterative Gaussian Mixture Approximation of the Posterior) method, developed by Data Fusion Research, combines the fitting of a set of Gaussian distributions to the likelihood surface and the TCN measurement input to the integration filter into a single iterative process. Measurement inputs to the integration filter need not be a linearised function of the state estimates nor possess a Gaussian error distribution.

In multi-hypothesis IGMAP, the system state estimate is represented as a set of n components, each comprising a weight, an estimate of the system state vector, and a corresponding covariance matrix, with the weights summing to unity. (These components can be thought of as representing n hypotheses about the value of the state vector.) In single-hypothesis IGMAP, there is just one component, with unit weight. For the present work, a 15-element state vector has been used, comprising position, velocity and attitude errors and the accelerometer and gyro biases. Each TCN transect is processed as follows:

- Each of the n components of the system state estimate is processed as follows:
  - (a) The component is projected into the three-dimensional space of the likelihood function by discarding elements of the state vector estimate and covariance matrix other than those representing components of position error. The component, thus projected, defines a Gaussian probability distribution over these three dimensions.

(b) This Gaussian distribution is multiplied by the likelihood function arising from the data in the transect, to produce the (un-normalised) Bayesian posterior distribution corresponding to this component of the state estimate.
(c) The resulting posterior distribution is approximated by a mixture of m Gaussian distributions. This is done iteratively using an expectation-maximisation algorithm.
(d) Each of these m Gaussian distributions is back-projected into the original state space to yield 15-element state estimate components.

- The result of Stage 1 is an updated state estimate comprising nm weighted Gaussian components. To prevent the number of components growing exponentially over time, a fusing algorithm is now applied to reduce the number of components back down to n. Fusing is done by merging one pair of components at a time, with Kullback-Leibler divergence used to select which pair to fuse at each step: this tends to select components with low weights and similar means and covariance matrices.
- Finally, each of the n components of the state estimate is propagated forward to the time of the next TCN transect using the Kalman filter time update equations.

A more detailed description of the IGMAP process can be found in (Runnalls *et al*, 2005). Comparing single hypothesis IGMAP with the PDAF method, IGMAP will generally produce a much closer fit to the probability distribution than the clumping algorithm, but at a higher cost in terms of processor load and developmental complexity. However, IGMAP can run in real-time on a 1 GHz processor with multiple state vector hypotheses. A multi-hypothesis integration filter conveys the advantage over its single hypothesis counterpart that less information is discarded between update cycles, particularly where a TCN measurement produces an ambiguous state vector update. Thus, more information is retained from which the current navigation solution may be estimated. The downside of a multi-hypothesis approach is increased processor load.

6. SIMULATION ENVIRONMENT. Simulations using the PDAF and best fix methods were conducted on the QinetiQ INS/GPS/TCN Integration Simulation (QIGTIS). This uses the same INS and GPS models as navigation simulations described in previous papers (Groves and Long, 2003; Gouldsworthy *et al*, 2002; Groves *et al*, 2002). The parameters for the INS error model were loosely based on a low grade (10°/hr specified drift) INS, the Boeing Digital Quartz Inertial Measurement Unit (DQI) (Boeing, 1997). Table 1 lists the rms values of the inertial instrument errors. However, for most runs, it was assumed that a transfer alignment (Groves, 2003) had been performed prior to the start of each simulation, reducing the accelerometer static biases to around 150 μg and the gyro static biases to around 1°/hr.

Simulation of radar altimeter measurements comprises a terrain simulation and a radalt model. A 'truth' terrain is generated from a level 1 DTED terrain height database based on user-defined error statistics. This ensures that the DTED data used by the TCN algorithm differs from the 'truth' terrain used to generate the radalt measurements by a representative amount. For most of the simulations in this

Table 1. 10°/hr INS inertial instrument errors (root mean square values are averaged over each component and expressed to 2 significant figures).

INS error	rms value
Accelerometer static biases	1500 μg
Accelerometer dynamic biases	150 μg
Correlation time of accel. dynamic biases	60 s
Accelerometer scale factor errors	200 ppm
Accelerometer cross-coupling errors	270 ppm
Accelerometer random walk	60 $\mu g/\sqrt{Hz}$
Accelerometer incremental quantisation	0·001 m/s
Gyro. static biases	10°/hr
Gyro. dynamic biases	1°/hr
Correlation time of gyro. dynamic biases	60 s
Gyro. scale factor errors	350 ppm
Gyro. cross-coupling errors	350 ppm
Gyro. g-dependent errors	1·0°/hr/g
Gyro. random walk	1·8°$/hr/\sqrt{Hz}$
Gyro. incremental quantisation	0·001°

study, a random height error of 5 m standard deviation was simulated with a 15 m standard deviation per axis horizontal displacement, correlated over 10 km. This reflects the relatively high accuracy of DTED over the UK. Four different radalt models can be simulated, each representing a different design: averaging, first return, peak return and high integrity. The averaging radalt model returns the average distance from the host vehicle to the terrain across the radalt footprint. The first return model returns the nearest distance across the radalt footprint above a certain threshold. The peak return model returns the distance corresponding to the peak in the spectrum of signal return against range. These models account for the distribution of the radalt footprint over a wide area of terrain, rather than the measurement errors of the sensor itself, being the main source of error in a height above terrain measurement.

For most of the simulations in this study, the high integrity radalt model was used. This is based on the design aims of advanced radalts, such as the BAE Systems Covert Radar Altimeter. These measure the whole spectrum of radar signal return against range and then process the data to derive a measurement of height above terrain. The high integrity radalt model outputs either the truth, plus or minus a height dependent random error or gives no measurement. For a measurement to be given, the maximum height above terrain was set at 5000 m, the maximum aircraft attitude (combined pitch and roll) at 30° and the maximum terrain gradient at 45°. The random error standard deviation was set at 5 m, height independent, corresponding to the typical performance of the other three radalt models 500 to 1000 m above the terrain.

A separate simulation software suite was used for the IGMAP runs. However, this implements the same INS model as QIGTIS and made use of QIGTIS simulated radalt measurements and the same flight profiles. Therefore direct comparisons can be made between INS/TRN results generated by the two simulations. Four trajectories were selected, overflying terrain of different levels of roughness as listed in Table 2, where the roughness is computed from the terrain gradient. Each trajectory

https://doi.org/10.1017/S0373463305003462 Published online by Cambridge University Press

Table 2. Locations and roughness of simulation trajectories.

Trajectory: roughness level	Location	Mean roughness
Low	Chilterns, UK	3·1%
Medium	South Pennines, UK	10·1%
High	Scottish Highlands, UK	21·7%
Very high	Swiss & Italian Alps	48·2%

comprised 100 km of straight and level flight at 300 ms^{-1} with a 5° 0·25 g turn to port starting at 100 s and a corresponding turn to starboard starting at 138 s. The flight duration was 333 s. Radalt measurements were taken at a rate of 2 Hz, noting that this gave significantly better TCN performance than a 1 Hz radalt measurement rate at this speed.

7. TCN INTEGRATION PERFORMANCE COMPARISON. To compare the performance of the different INS/TCN integration techniques, simulation runs were conducted without GPS. For most of the runs, a 100 m initial position uncertainty was assumed, representing marginal GPS performance prior to loss. However, the velocity, attitude and accelerometer and gyro biases were assumed to be well calibrated. The first runs with the best fix and PDAF INS/TCN algorithms were conducted to tune the TCN and integration algorithms. The navigation performance was found to be more sensitive to tuning over the low roughness terrain than over the other trajectories. This is because, with less terrain height variation, a lower roughness terrain provides less information for the TCN algorithm to work with.

Optimisation of the transect length was found to be a trade-off. For low roughness terrain, longer transects give better performance as they increase the chance of a good match between the measurements and the DTED database. For higher roughness terrain, there is generally sufficient data in a shorter transect for a good match, so shorter transects can give better performance as there is then a greater frequency of position fixes. A transect length of 8 s was found to be a good compromise. Where low roughness terrain was overflown and the search area was wide, weighting the likelihood surface according to the distance from the filter indicated position divided by the filter position uncertainty was found to improve performance. Otherwise, it had little effect. Varying the weighting coefficients had little effect on performance. However, positive feedback problems were found during QIGTIS testing when the filter weighting coefficients were too high, so the lowest value assessed, 0·05, was selected for the main performance comparison runs.

Using the maximum database search area when the integrated navigation solution was relatively accurate was found to degrade performance. This is because widening the search area increases the chance of matching the measurements with the wrong features in the database. Thus, it was confirmed that it is better to scale the search area with the navigation filter position uncertainty.

For the best fix integration algorithm, the optimum minimum probability of the best fix was found to be 0·5 and the maximum probability for the second best fix 0·25. The remaining algorithm tuning parameters were the same for both best

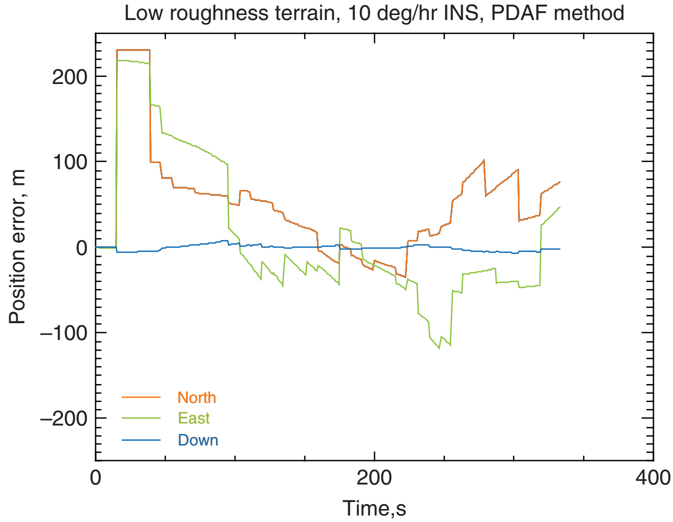

Figure 2. INS/TCN position error – transfer aligned 10°/hr INS; 100 m per axis initial position uncertainty; PDAF method; low terrain roughness trajectory.

fix and PDAF. The measurement covariances output by the TCN algorithm were found to be optimal, not needing re-weighting by the integration algorithm. However, it was found to be necessary to estimate the correlated DTED database errors as Kalman filter states in order to attain a reasonable match between errors and uncertainties. The optimum threshold for the spike filter, which rejects TCN measurements inconsistent with the filter estimates, was found to be 3 standard deviations.

The best fix and PDAF integration methods were compared over all four trajectories under a range of conditions. The initial position uncertainty was varied between 10 m, 100 m and 1 km per axis; the INS grade varied between 1, 10 and 100°/hr gyro drift, and some runs were conducted with a non transfer aligned INS. Performance over the medium, high and very high roughness trajectories was similar. However, performance over the low roughness terrain was notably poorer. Figures 2 and 3 illustrate this with the PDAF method, 10°/hr transfer aligned INS and 100 m initial uncertainty. Each figure shows the north, east and vertical components of position error for a single run.

Comparing the results obtained with the best fix and PDAF integration methods, performance was similar in most cases. However, in most of the runs with the 1 km initial position uncertainty, the navigation solution initially locked onto a false TCN fix. The PDAF algorithm recovered the correct navigation solution in all cases, but the best fix algorithm failed to recover in the high roughness run. Thus, the PDAF approach is more robust. Figures 4 and 5 show the position errors using the two TCN integration methods over the low roughness terrain with a 1 km initial uncertainty. Over the first 40 s, the maximum north position error was ~1500 m and the maximum east error ~200 m with both methods. In general, the two methods performed similarly, but with the PDAF north position converging more quickly during the 50–150 s period. In general the parameters of the IGMAP algorithm were set either

https://doi.org/10.1017/S0373463305003462 Published online by Cambridge University Press

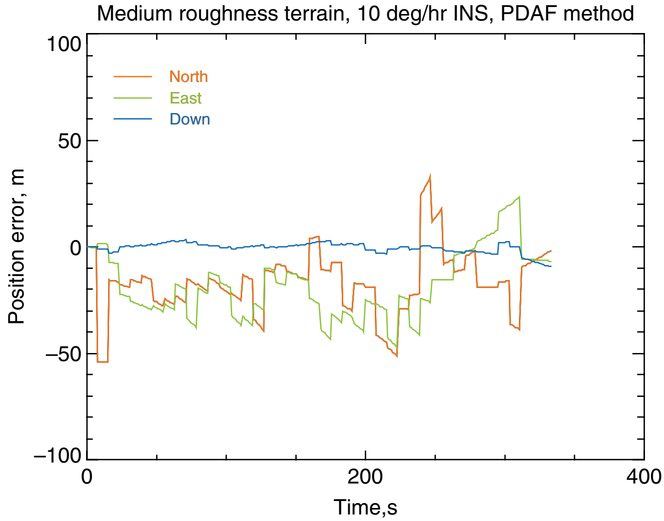

Figure 3. INS/TCN position error – transfer aligned 10°/hr INS; 100 m per axis initial position uncertainty; PDAF method; medium terrain roughness trajectory.

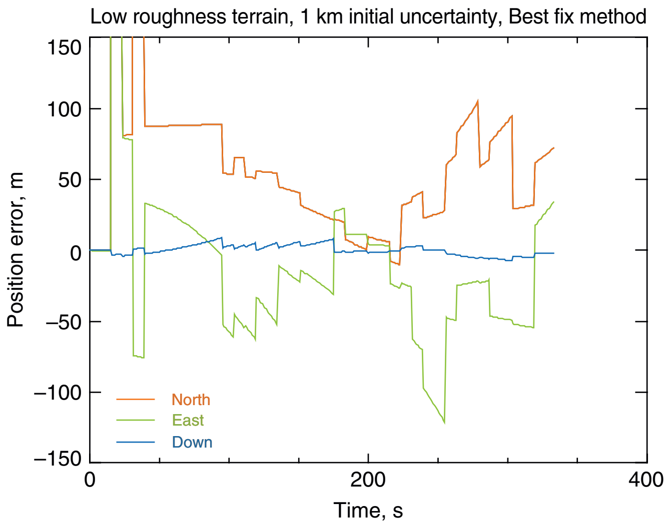

Figure 4. INS/TCN position error – transfer aligned 10°/hr INS; 1 km per axis initial position uncertainty; best fix method; low terrain roughness trajectory.

to correspond to the values obtained by tuning the best fix and PDAF algorithms, or on the basis of previous experience with IGMAP. The number of Gaussian components used to approximate the posterior distribution, *m*, was set to four, and in multiple hypothesis IGMAP the number of hypotheses, *n*, was also set to four.

IGMAP works well with short transects, and can indeed work in a sequential mode (i.e. with a single radalt sample per transect) though performance is better with somewhat longer transects. For the purposes of the comparisons in this paper, a transect of duration 3·5 s has been used, comprising seven terrain height measurements. This is slightly less than half of the transect length used for the best fix and PDAF algorithms. For the IGMAP simulations, four of the more demanding scenarios were selected for simulation: the low and high roughness trajectories with a 1 km initial uncertainty and 10°/hr transfer aligned INS; and the low roughness trajectory with a 100 m initial uncertainty, 100°/hr aligned INS and 10°/hr non-aligned INS. Each of these was run with both the single-hypothesis and multi-hypothesis (4 filter) versions of the algorithm.

Over the low roughness terrain runs, IGMAP position performance was up to a factor of two better than the best fix and PDAF algorithms' performance, particularly where the poorer INS were used. Figure 6 illustrates this for the 1 km initial uncertainty run. The IGMAP algorithm was able to obtain a first fix earlier than the PDAF and best fix algorithms due to differences in simulation implementation. IGMAP also obtained a relatively poor initial fix, in error by ~400 m north and ~250 m east, but this was much less than obtained with the PDAF method and was consistent with the filter uncertainties. IGMAP performance over the high roughness terrain was very close to that of the PDAF method, however. This confirms earlier observations that TCN performance is much more sensitive to the algorithm design and tuning over low roughness terrain where it is more difficult to obtain good matches between the terrain height measurements and the database.

IGMAP offers two advances over the PDAF integration method with the Gaussian clumping TCN algorithm. Firstly, a more sophisticated and more accurate process is used to fit a set of Gaussian distributions to the TCN probability distribution and, secondly, the fitting process and filter update are integrated. Further research is needed to determine which of these advances bears greater responsibility for the improved performance.

Comparing the performance of the single and multi-hypothesis versions of IGMAP, very little difference in position error was observed. Figure 7 illustrates this for the 1 km initial uncertainty run. This suggests that there is little benefit in moving to a multi-hypothesis integration filter for TCN, though it should be noted that these results were obtained with a simulated radalt and terrain.

8. RADALT AND DATABASE ERRORS.

To determine the effect of radalt and database errors on TCN performance, a series of simulations was run using the PDAF INS/TCN algorithm with a 10°/hr transfer aligned INS, 100 m initial position uncertainty and PDAF measurement model. With a 12° full width at half maximum (FWHM) intensity radar beam simulated, performance with the four different radalt models (see Section 6) over the low, medium and high roughness terrain was similar, with the high integrity model giving the best results on average, followed by the peak return, first return and averaging models. When the FWHM was increased to 25° (not applicable to the high integrity model), the TCN performance with the averaging radalt model was significantly degraded, whilst that with the peak return and first return models was affected only marginally. Thus, the latter two designs are much more suited to TCN applications.

Over the very high roughness terrain simulations, TCN performance was poor with the averaging, peak return and first return radalt models. A major reason for this was that the average host vehicle height above terrain was much greater for these runs, spreading the radalt footprint over a wider area of terrain. To get reasonable navigation performance, it was necessary to reduce the radalt beamwidth to 3° FWHM and re-tune the TCN algorithm. Thus, the design of a TCN system is contingent on the range of host vehicle heights above terrain it is intended to operate over. To gauge the effect of terrain height database errors on TCN performance, a set of runs was conducted with the simulated database errors doubled. This led to an approximate doubling of the maximum position errors observed. Thus database accuracy is an important factor in determining TCN performance.

Lastly, a set of runs was conducted with a perfect radalt and DTED database simulated. The position performance in these runs was about a factor of two better than that obtained with the standard models. This shows that there is considerable scope to improve TCN performance by improving the sensor and the terrain height database.

9. WEIGHTING OF TCN AND GPS. The simulations described so far considered TCN operation when GPS is unavailable. Where strong GPS signals are available, GPS will always be more accurate than TCN. However, at the margins of GPS operation, adding TCN measurements can potentially improve the integrated navigation performance. To test this, a number of simulations were conducted using QIGTIS with PDAF TCN integrated with ATC INS/GPS (Groves and Long, 2005) in a centralised filter. These scenarios were simulated with the 10°/hr transfer aligned INS model over the medium terrain roughness trajectory.

Three scenarios were considered: jamming at the maximum tolerable level for robust ATC integration; three strong GPS signals tracked and a good GPS environment with 5 satellites tracked. In each case, the TCN/GPS measurement weighting ratio was varied between 0 and 5. The weights are applied to the Kalman gain matrices and the TCN and GPS weights sum to unity. This is done to prevent over-correction of the INS due to two measurement sources being used to calibrate the same errors.

Considering first the jamming scenario, GPS measurements become very noisy as the maximum tolerable jamming level is approached. Figure 8 shows the position error with INS and GPS only, whilst Figure 9 shows the INS/GPS/TRN position error with both measurements weighted equally. The performance with all three navigation systems was significantly better than that obtained with either INS and GPS or INS and TCN (Figure 3). TCN/GPS weightings of 0·5 and 2 gave similar performance, whilst a weighting of 5 in favour of TCN gave slightly poorer performance than that shown in Figure 9.

With only three satellites tracked, the INS/GPS navigation solution drifted, with position errors growing with time, whereas with TCN added, the position errors were bounded within 11 m per axis for all of the TCN/GPS weighting ratios considered. With five good GPS signals, adding TCN brought no improvements to the navigation accuracy. However, with a weighting ratio of 0·5 or 1·0, adding TCN had no detrimental effect on navigation accuracy either. These results suggest that, in an integrated INS/GPS/TCN navigation system, optimum performance can be obtained by weighting GPS and TCN measurements equally in all GPS environments.

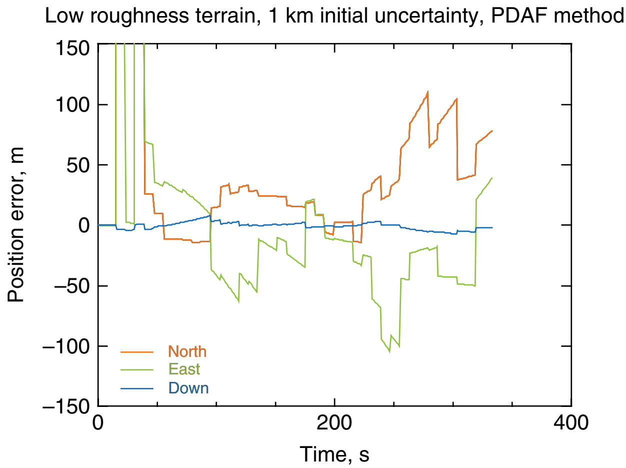

Figure 5. INS/TCN position error – transfer aligned 10°/hr INS; 1 km per axis initial position uncertainty; PDAF method; low terrain roughness trajectory.

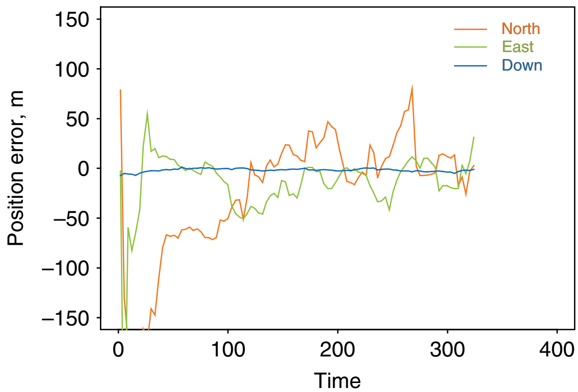

Figure 6. INS/TCN position error – transfer aligned 10°/hr INS; 1 km per axis initial position uncertainty; single hypothesis IGMAP method; low terrain roughness.

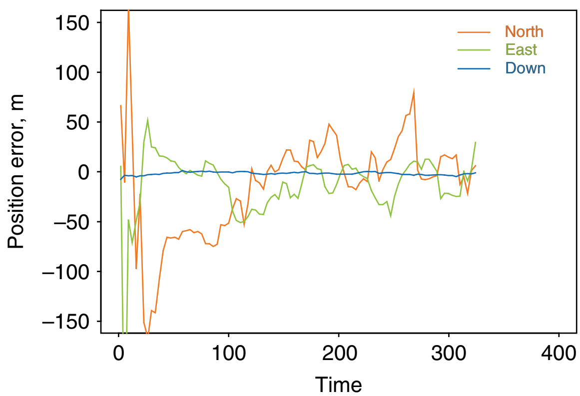

Figure 7. INS/TCN position error – transfer aligned 10°/hr INS; 1 km per axis initial position uncertainty; multiple hypothesis IGMAP method; low terrain roughness.

https://doi.org/10.1017/S0373463305003462 Published online by Cambridge University Press

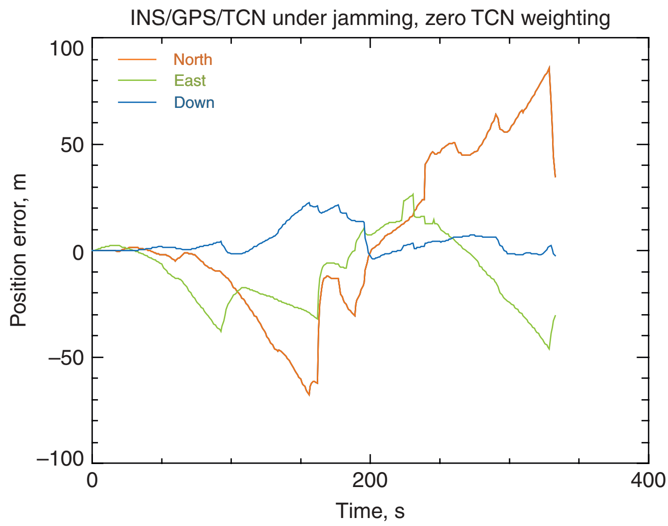

Figure 8. INS/GPS position error – transfer aligned 10°/hr INS under jamming.

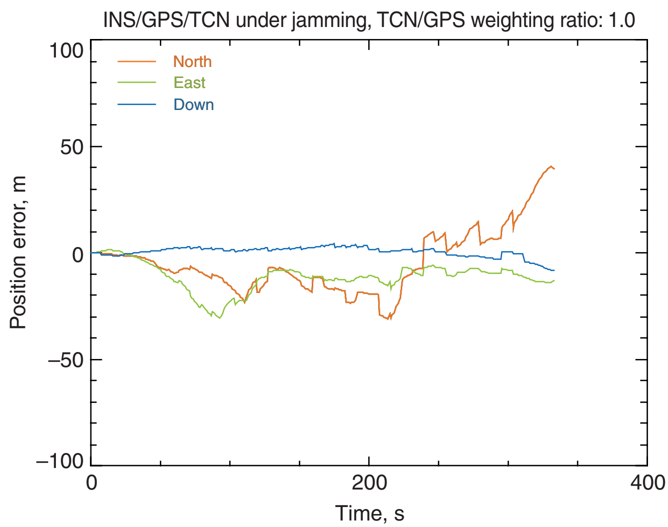

Figure 9. INS/GPS/TCN position error – transfer aligned 10°/hr INS under jamming, GPS & TCN weighted equal.

10. CONCLUSIONS. A number of different techniques for integrating TCN with INS and GPS have been developed and assessed by simulation. Navigation performance is much more sensitive to the design and tuning of the TCN and INS/TCN integration algorithms when the terrain overflown has relatively little height variation, making it more difficult to obtain unique matches between the terrain height measurements and the region of the database they are compared with. The

https://doi.org/10.1017/S0373463305003462 Published online by Cambridge University Press

design of a TCN system is also contingent on the range of host vehicle heights above terrain it is intended to operate over.

Using a weighted fix integration technique, such as PDAF, makes the navigation solution more robust against false TCN fixes than a simple best fix integration. Of the weighted fix TCN integration techniques, IGMAP performs better than the PDAF algorithm in association with Gaussian clumping TCN, justifying the greater complexity and processor load. However, the simulations did not demonstrate any benefit in implementing a multi-hypothesis integration filter over the corresponding single hypothesis approach.

Where TCN is combined with INS and GPS, optimum performance is obtained by weighting the TCN and GPS measurements equally in the integration filter. Where only a limited number of GPS signals can be tracked or the GPS measurements are noisy due to jamming, the TCN measurements improve the overall navigation solution. The next step is to validate the work presented here with flight trials results.

ACKNOWLEDGEMENTS

This work was funded by the Applied Research Programme of the UK Ministry of Defence. The authors would also like to thank Martin Heath, Adam Kirby, Daniel Long and Jane Mercer who contributed to the simulation software development; Martin Brown who suggested the use of a clumping algorithm and a PDAF; and Paul McNeil and Lawrence Dack of Hi-Q Systems who developed the CVN and LLN algorithms. This paper was first presented at ION GNSS 2004 in Long Beach, CA, USA, 21–24 September 2004.

REFERENCES

Boasman, N. J. and Briggs, P. (2002). The Development of an Anechoic GPS Test Facility, *Proceedings of the ION 58th Annual Meeting and CIGTF 21st Guidance Test Symposium*, Albuquerque, NM.

Boeing Defence UK (1997). DQI Digital Quartz Inertial Measurement Unit (Specification); Integrated GPS/INS – C-Migits II Users Guide.

Campbell, J., Uijt de Haag, M. and Van Graas, F. (2003). Light Detection and Ranging-Based Terrain Navigation – A Concept Exploration. *Proceedings of ION GPS/GNSS 2003*, Portland, OR.

Campbell, J., Uijt de Haag, M. and Van Graas, F. (2004). Terrain Reference Navigation using Airborne Laser SCAnner (ALASCA) – Preliminary Flight Test Results. *Proceedings of ION 60th Annual Meeting*, Dayton, OH.

Carlbom, P. and Johansson, N. (2002). Terrain Navigation and IntegrityMonitoring for Avionic Platforms. *Proceedings of NATO SET Symposium* '*Emerging Military Capabilities Enabled by Advances in Navigation Sensors*', Istanbul, Turkey.

Carlson, N. A. (1990). Federated Square Root Filter for Decentralized Parallel Processes. *IEEE Transactions on Aerospace and Electronic Systems* **26** 517–525.

Carroll, J. *et al* (2001). Vulnerability Assessment of the Transportation Infrastructure Relying on the Global Positioning System. *John A Volpe National Transportation Systems Center report for US Department of Transportation*.

Defense Mapping Agency (1983). Product Specifications for Digital Landmass System (DLMS) Data Base.

Gouldsworthy, S. N., Groves, P. D. and Wells, M. (2002). High-fidelity model development for navigation warfare simulation studies. *Proceedings of ION GPS 2002*, Portland, OR.

Groves, P. D. (2003). Optimising the Transfer Alignment of Weapon INS. *The Journal of Navigation* **56** 323–335.

Groves, P. D. and Long, D. C. (2003). Adaptive tightly-coupled, a low cost alternative anti-jam INS/GPS integration technique. *Proceedings of the 2003 ION National Technical Meeting*, Anaheim, CA.

Groves, P. D. and Long, D. C. (2005). Combating GNSS interference with advanced inertial integration INS/GNSS. *The Journal of Navigation* **58**, 419–432.

Groves, P. D., Wilson, G. G. and Mather, C. J. (2002). Robust rapid transfer alignment with an INS/GPS reference. *Proceedings of the 2002 ION National Technical Meeting*, San Diego, CA.

Gustafson, D., Dowdle, J. and Flueckiger, K. (2000). A High Anti-Jam GPS-Based Navigator. *Proceedings of the 2000 ION National Technical Meeting*, Anaheim, CA.

Handley, R. J., Abbott, J. P. and Surawy, C. R. (1998). Continuous Visual Navigation – An Evolution of Scene Matching. *Proceedings of the 1998 ION National Technical Meeting*, Long Beach, CA.

Handley, R. J., Dack, L. and McNeil, P. (2001). Flight Trials of the Continuous Visual Navigation System. *Proceedings of the 2001 ION National Technical Meeting*, Long Beach, CA.

Handley, R. J., Groves, P. D., McNeil, P. and Dack, L. (2003). Future Terrain Referenced Navigation Techniques exploiting Sensor Synergy. *Proceedings of GNSS 2003, The European Navigation Conference*, Graz, Austria.

Hollowell, J. (1990), HELI/SITAN: A Terrain Referenced Navigation Algorithm, for Helicopters. *Proceedings of the IEEE Position Location And Navigation Symposium* (*PLANS*) *1990*.

McNeil, P., Dack, L., Handley, R. (2002). Terrain Referenced Reversionary Navigation. *Proceedings of Royal Institute of Navigation NAV02* '*GNSS Vulnerability*', London, UK.

Metzger, J., Wendel, J. and Trommer, G. F. (2002). Comparison of Different Terrain Referenced Navigation Techniques. *Proceedings of ION 58th Annual Meeting*, Albuquerque, NM.

Metzger, J and Trommer, G. F. (2003). Improvement of Modular Terrain Navigation Systems by Measurement Decorrelation. *Proceedings of ION 59th Annual Meeting, pp 353–362*, Albuquerque, NM.

Moore, S., Myers, G. and Hunt, R. (2001). Future Integrated Navigation Guidance System. *Proceedings of the 2001 ION National Technical Meeting*, Long Beach, CA.

Robins, A. (1998). Recent Developments in the 'TERPROM' integrated navigation system. *Proceedings of ION 44th Annual Meeting*.

Runnalls, A. R. (1985). A Bayesian Approach to Terrain Contour Navigation. *Proceedings of* '*Guidance-Control-Navigation Automation for Night All-Weather Tactical Operation*' *NATO AGARD-CP-387*.

Runnalls, A. R., Groves, P. D. and Handley, R. J. (2005). Terrain Referenced Navigation Using the IGMAP Data Fusion Algorithm. *Proceedings of ION 61st Annual Meeting*, Boston, MA.

Runnalls, A. R. and Handley, R. J. (1998). The 'Gold Standard' Navigator, *Proceedings of IEE Eurofusion 98*, Great Malvern, UK.

Sennott, J. and Senffner, D. (1997). A GPS Carrier Phase Processor for Real-Time High Dynamics Tracking. *Proceedings of ION 53rd Annual Meeting*.

Soloviev, A., Van Graas, F. and Gunawardena, Sanjeev. (2004). Implementation of Deeply Integrated GPS/Low-Cost IMU for Reacquisition and Tracking of Low CNR GPS Signals. *Proceedings of the 2004 ION National Technical Meeting*, San Diego, CA.

Owen, J. I. R. and Wells, M. (2001). An Advanced Digital Antenna Control Unit for GPS. *Proceedings of the 2001 ION National Technical Meeting*, Long Beach, CA.

Wells, M. and Owen, J. I. R. (2002). A Digital Adaptive Array Demonstrator for GPS. *Proceedings of NATO SET Symposium* '*Emerging Military Capabilities Enabled by Advances in Navigation Sensors*', Istanbul, Turkey.

Yu, P., Chen, Z. and Hung, J. C. (1991). Performance Evaluation of Six Terrain Stochastic Linearisation Techniques for TAN. *Proceedings of the IEEE National Aerospace and Electronics Conference* (*NAECON*) *1991*.

https://doi.org/10.1017/S0373463305003462 Published online by Cambridge University Press