

Terrain-Referenced Navigation Using the IGMAP Data Fusion Algorithm

Andrew R. Runnalls, *Data Fusion Research Ltd, UK*
Paul D. Groves, *QinetiQ Ltd, UK*, Robin J. Handley, *QinetiQ Ltd, UK*

BIOGRAPHY

Dr Andrew Runnalls received his M.A. and Ph.D. degrees from Cambridge University. For ten years he was with the guidance systems division of GEC Avionics at Rochester, UK, where he played a key part in the development of the SPARTAN terrain-referenced navigation system, which was flight-trialled in F16, Tornado, A6 and other aircraft. He joined the Computing Laboratory of the University of Kent in 1988, where he has continued to work on statistical data fusion techniques, and established Data Fusion Research Ltd in 1999. He is a Fellow of the Royal Institute of Navigation. Email: arr@datafusion.co.uk

Dr Paul Groves has been with the Navigation Systems team of QinetiQ (and DERA before that) for eight years. He specialises in integration, alignment and modelling and is an author of about 20 publications on transfer alignment, INS/GPS and multi-sensor integration, GPS modelling and GPS signal monitoring. He also teaches a course on the principles of navigation technology. He holds an M.A. and a D.Phil. in physics from Oxford University. He is an Associate Fellow of the Royal Institute of Navigation, secretary of the R&D special interest group and a member of the NAV05 conference committee.

Robin Handley received his B.Sc. in Mathematics from Leeds University in 1986 and has worked as a navigation scientist at QinetiQ (and its predecessors) since then. During this period he has worked extensively on airborne navigation trials, analytical software, and data analysis, through which he has developed considerable experience in the behaviour of inertial, GPS, and terrain-based navigation systems. His specialisation is terrain-based navigation techniques, and their integration with inertial sensors. He has written and presented technical papers on this subject at several international conferences, and had two papers published in the *Journal of Defence Science*. He is currently the Technical Capability Leader of the Integrated Naviga-

tion Group, in QinetiQ, and is an Associate Fellow of the Royal Institute of Navigation.

ABSTRACT

Accuracy and robustness are vital to present and future air navigation, both in the military and civil spheres. The strengths of INS and GPS are well known; integrated INS/GPS systems combine their advantages. However, the knowledge required to jam GPS is becoming public, and it can be carried out with basic hardware. When GPS data are unavailable, and a low grade INS is used, navigation accuracy quickly degrades to an unacceptable level.

Terrain-referenced navigation (TRN) techniques—for example, terrain contour navigation based on radio altimeter measurements over undulating terrain—provide a complementary technology: when integrated with INS and GPS, TRN can allow the system to establish and maintain high accuracy even in sustained GPS outages. The authors' previous paper [1] explored different techniques for performing the triple integration of TRN, INS and GPS, and found that there were performance advantages if TRN data were processed using a novel data fusion algorithm known as IGMAP; in particular, IGMAP was found to provide more accurate and robust performance over low roughness terrain, which can prove challenging to conventional TRN algorithms. The present paper explores the IGMAP algorithm and its performance in more detail.

IGMAP (Iterative Gaussian Mixture Approximation of the Posterior) is an advanced data fusion algorithm for handling non-linear measurements, particularly ambiguous measurements (i.e. measurements for which the likelihood function may be multimodal), in conjunction with a linear or linearisable system model. It is particularly well suited to system models of high dimensionality, and applications where it is desired to interoperate with existing approaches using a Kalman Filter or multi-hypothesis

Kalman Filter. Although devised with integrated TRN/INS or TRN/INS/GPS systems in mind, the algorithm has potential applications to other data fusion problems, for example in target tracking.

The paper outlines the mathematical foundations of the algorithm, and illustrates its operation using recorded flight data based on the use of an inertial system aided by terrain height information from a radio altimeter.

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 position solution—of the order of 10 m (2σ) in each axis—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 [2]. 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. [3, 4, 5]), together with advanced INS/GPS integration techniques, such as adaptive tightly-coupled (ATC) [6, 7] and deep integration (e.g. [8, 9, 10]). 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 this technology 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 [11, 12]. 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 [13], developed and tested by QinetiQ and Hi-Q Systems.

Our previous paper [1] used simulation to compare a number of different techniques for integrating TCN with INS, and for the three-way integration of TCN with INS and GPS. The techniques considered for integrating TCN data included a best-fix method, a probabilistic data association filter (PDAF), and a new algorithm called IGMAP. It concluded that a weighted fix integration technique—of which PDAF and IGMAP are examples—makes the navigation solution more robust against false TCN fixes than a simple best-fix integration. The simulation results obtained indicated that IGMAP performs sufficiently better than the PDAF algorithm to justify the greater complexity and processor load that it entails.

In the present paper we take a closer look at the IGMAP algorithm. Sec. 2 reviews the various approaches that have been taken to TCN, to place the IGMAP method in context. In Sec. 3 we set out the design objectives that led to the development of this algorithm. In Sec. 4 we describe how the algorithm works, and in Sec. 5 we describe the behaviour of the IGMAP algorithm as applied to data recorded during a complete sortie of a Tornado aircraft. Finally in Sec. 6 we summarise conclusions and identify possible areas of future work.

2 REVIEW OF TCN TECHNIQUES

Development of terrain contour navigation (TCN) started in the 1970s and a number of systems have been produced commercially over the years. Conventionally, such systems estimate the height of the terrain below the air vehicle by subtracting radio altimeter height from INS or 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) [14]. A range of different processing techniques have been developed to obtain position fixes from the comparisons of measured and database terrain heights

[15]. 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 data fusion algorithm. Typically this data fusion algorithm is an extended Kalman filter (EKF): the terrain gradient at the current best estimate of position is used to attribute the observed height difference to a linear combination of the latitude, longitude and height components of the aided INS position error. Sequential processing is well established in commercial TCN systems such as BAE Systems' TERPROM [16] and the American SIRTAN [17, 18, 19]. Particle-filtering approaches [20, 15] (also known as bootstrap filtering) also take a sequential approach, but without the linear/Gaussian assumptions of the extended Kalman filter

The principal advantage of the EKF 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 [21]. However, a fundamental problem remains in that the gradient is calculated below the aircraft's position *as indicated by the navigation system*, 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 [19]. An alternative approach is to process the radalt measurements with a non-linear filter [22].

In batch processing, a series of terrain height measurements, known as a transect, are processed together; this was the approach taken in the original TERCOM system [23], and in early work of one of the present authors [24]. 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 [15]. Another straightforward approach is to fit a Gaussian distribution to the likelihood surface: this is the approach taken in [24]. However, as we shall see in Sec. 4, the likelihood surface is often decidedly non-Gaussian in form, and may well be multimodal, i.e. it may have more than one peak. The approach used in the SPARTAN system [25, 26] was to make a pro-

visional Gaussian approximation to the likelihood surface, but to carry forward residuals for consideration alongside the data from subsequent transects. Another approach is to fit multiple Gaussian distributions to the likelihood surface, providing a multiple hypothesis position fix for processing by the data fusion algorithm. This approach was the basis of the Gaussian clumping algorithm used in two of the three algorithms considered in our previous paper [1], namely the best-fix and PDAF algorithms.

The IGMAP algorithm is essentially a batch processing algorithm, although it can be configured so that the 'batches' comprise a single radalt measurement, in which case it will effectively operate as a sequential method. However, unlike the EKF sequential methods considered above, it does not require any linear approximation to the terrain surface. Like the best-fix and PDAF algorithms of [1], it uses a process of fitting multiple Gaussian distributions, but instead of being fitted to the likelihood function itself, they are fitted to the resulting posterior distribution; the fitting method is also completely different. The following section amplifies the design objectives that led to the development of IGMAP, and the algorithm itself is described in Sec. 4.

3 DESIGN OBJECTIVES

The IGMAP algorithm described in this paper arose from a need for a data fusion algorithm suited to multiway integrated navigation with terrain-referenced navigation as one (or more) of the inputs. In this section we consider the design objectives that led to the development of the IGMAP algorithm. There were three primary objectives:

1. **The algorithm should be capable of updating the system's navigation solution quickly—e.g. within a few seconds—in response to newly-gathered terrain data.**

This is in contrast, for example, to early TCN systems [23, 24] which would gather radio altimeter over a distance of several kilometres before comparing the terrain profile thus measured with a digital elevation map (DEM). This approach has the advantage that the terrain profile over long distances is likely to be effectively unique within the area of navigational uncertainty, and the data can therefore be processed into a precise and unambiguous position fix, which is easy to process with conventional integrated navigation schemes.

The disadvantage with this 'long transect' approach is there there is a considerable delay before the navigational estimate is updated. If this delay could be reduced, there would be less time for residual navigational drifts to accumulate before the TRN update becomes available. This in turn may allow the air vehi-

cle to fly at lower altitude, since the position of terrain features and obstacles in relation to the aircraft will be better known. This consideration is particularly strong for applications such as unmanned air vehicles and guided weapons, where the use of lower grade INS will result in high rates of navigational drift.

2. The algorithm should be capable of handling the position ambiguities that often arise using terrain-referenced navigation, in such a way as to make maximum and timely use of the terrain data.

As has just been remarked, if TCN data are gathered over several kilometres, an unambiguous position fix can be obtained. Conversely, if radio altimeter data are gathered only over a short period—perhaps a single radio altimeter sample, or perhaps a short ‘transect’ comprising radio altimeter data over a few hundred metres—then often the terrain profile thus measured will be a *good* match to the DEM at several distinct places within the area of navigational uncertainty, and a *satisfactory* match over an extended (and usually irregularly-shaped) area. We shall see vivid examples of this later in the paper. Consequently there is an inherent ambiguity in the navigational information provided by the terrain data.

Likewise, in line-feature matching systems such as CVN [13], the lines in a particular scene captured by the imaging sensor may match the line feature database in more than one place, thus again leading to a navigational ambiguity.

One simple approach to such cases (applicable to both TCN and line-feature matching) is simply to discard batches of data (or images) that lead to appreciable ambiguities. A rather better approach, less wasteful of terrain data, is to defer processing of the data: instead of being processed straight away, the data batch is aggregated with subsequent data until the ambiguity is resolved. This is the approach taken by [27]; the SPARTAN algorithm [25, 26] also utilises this idea. However, with either approach, the underlying navigational drift will continue to accumulate inexorably until the ambiguity is resolved.

It would arguably be better if the data could be fed immediately into the overall navigation solution, but in a way which takes explicit account of possible ambiguities: this is the second design objective.

3. The algorithm should be directly compatible with Kalman filter approaches to multiway integrated navigation, including in particular the multi-hypothesis Kalman filter (MHKF).

Conventional Kalman filter approaches to integrated navigation are based on linear (or linearised) statistical error models of the navigation system components, with Gaussian error statistics. The second

design objective above implies that we must move away from such linear/Gaussian statistics towards approaches that can handle more general statistical models of the system components.

Data Fusion Research Ltd (DFRL) and QinetiQ Ltd have devoted considerable effort to exploring the potential of Monte-Carlo Markov Chain (MCMC) methods in application to terrain-referenced navigation—see, in particular, [28]. Such approaches have the strong advantage that they can incorporate arbitrary statistical models of the navigation system components, and—using these models—can process the input data into an evolving integrated navigation estimate using strict Bayesian statistical reasoning. These methods take ambiguities, non-linearities, and non-Gaussian error statistics all in their stride. The related particle-filtering approaches to TRN [20] have similar advantages.

There are, however, some disadvantages with MCMC (including particle filter) approaches to integrated navigation:

- For relatively straightforward integrated navigation tasks, such as INS/GPS integration, MCMC methods are an overkill: as compared with a Kalman filter, they require much greater computational loads, but offer little or no performance advantage.
- Consequently, for multiway integration, e.g. INS/GPS/TRN, it would seem to be desirable for the overall integration to be carried out using Kalman filtering techniques, leaving MCMC approaches to deal with the non-linearities and ambiguities of TRN, and somehow feeding the MCMC output into the overall Kalman filter. Unfortunately, the output of MCMC methods is not in a form which lends itself to use as a Kalman filter input.
- Typical INS models often have 15 or more elements in the state vector. Such high dimensionality may cause problems for MCMC approaches.

The design requirements outlined above led QinetiQ and DFRL to search for techniques that would offer some of the versatility of MCMC methods of data fusion, but in a way that would fit into a Kalman filter framework for multiway integrated navigation. The IGMAP algorithm was the upshot of this search.

4 OPERATION OF THE IGMAP ALGORITHM

In this section we explain how the IGMAP algorithm works, as applied to an integrated TCN/INS system us-

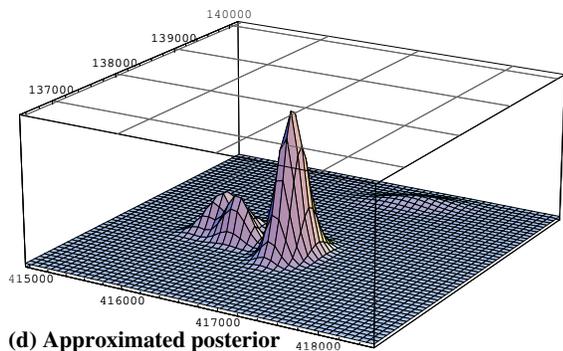
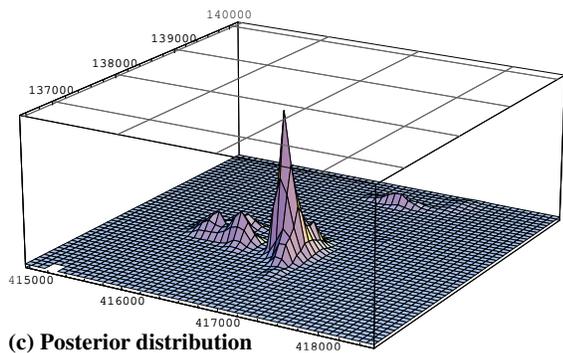
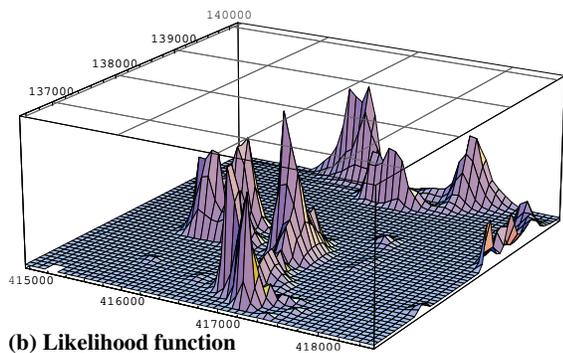
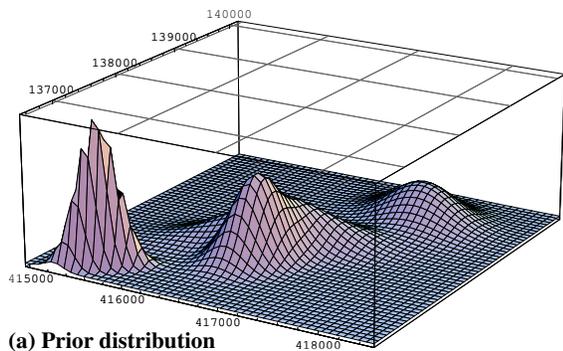


Figure 1 Overview of the IGMAP algorithm. The horizontal axes are labelled with UK National Grid coordinates, and the graticule on the top surface of each box consists of 1 km squares.

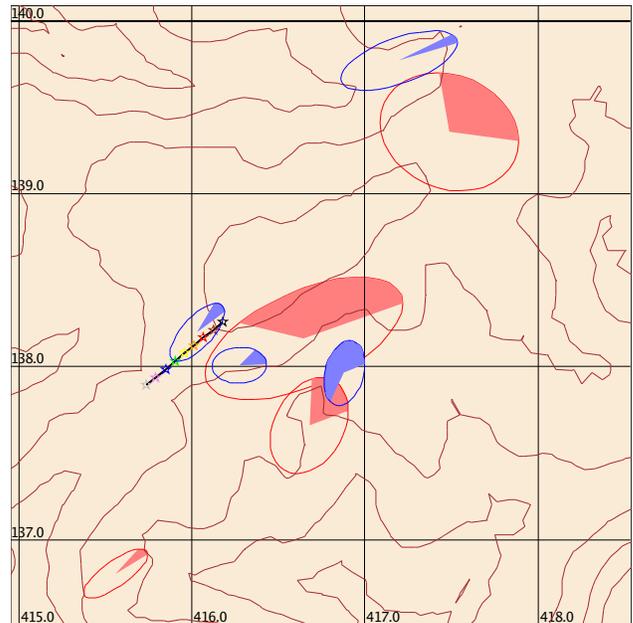


Figure 2 Plan view of transect processing. Coordinates as in Fig. 1. The brown lines are terrain contours at 10 m vertical spacing. See text for further explanation. The stars mark the true (south-westerly) path of the aircraft during the transect, as given by INS/GPS data.

ing a radio altimeter as the terrain sensor. An overview of its operation is presented in Fig. 1. At any time, the system's navigational estimate is represented as a mixture (i.e. a weighted sum) of Gaussian distributions, each defined over the d dimensions of the state vector defining the error dynamics of the INS (or combined INS/GPS). For example, for an INS the state vector will typically comprise three elements representing the components of the current position error, three elements representing the components of velocity error, three platform misalignment angles, plus further elements representing accelerometer and gyro calibration errors.

It is difficult to represent distributions over such a large number of dimensions in visual form, so in this overview we shall simplify matters by considering just two dimensions, representing horizontal position. With this simplification, consider an example run in which the system starts with a very high uncertainty about position. In this example run, the system's positional estimate shortly after the IGMAP algorithm has started operating is as shown in the probability density function in Fig. 1(a). This distribution is a mixture of four Gaussian components, which are shown in plan view as the red ellipses in Fig. 2. Each ellipse encloses 50% of the probability volume of the Gaussian component, and the proportion of the area of the ellipse shown as a shaded sector represents the weight of the component within the 4-component mixture: these weights sum to unity. In the case shown, the probability distribu-

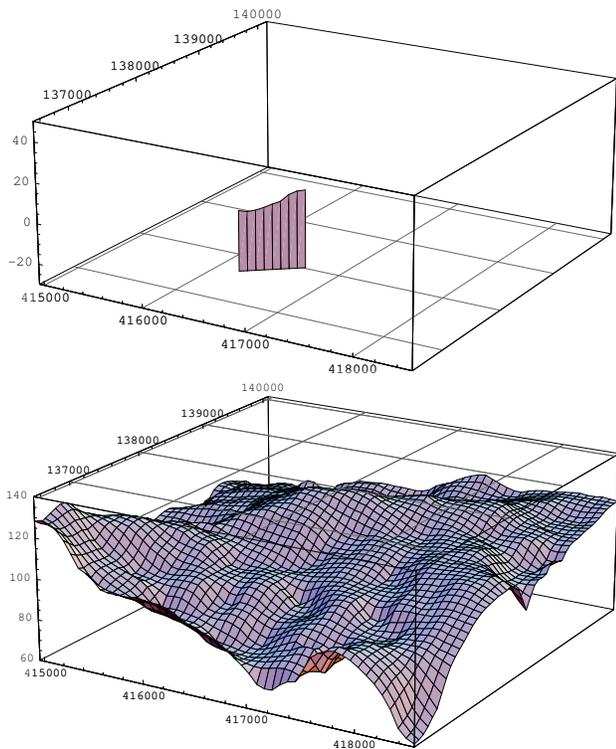


Figure 3 (a) The sensed terrain profile observed during a transect; the aircraft was flying approximately south-west while the transect data were gathered. (b) The terrain topography as given by the Digital Elevation Map (DEM). In both panels, heights are in metres. Note that the vertical scale is exaggerated by a factor of about 17.5 in comparison with the horizontal scale.

tion comprises a sharp peak towards the south-west corner of the region shown, contributing about 12% of the total probability. (The jaggedness of this peak is a plotting artefact.) There is a broader, rounded peak towards the north-east corner, with about 27% of the total probability. The remaining two Gaussian components are close together in the middle of the area shown, and form an irregular, roughly T-shaped peak in the perspective view. (Note that despite our use of topographical metaphors such as ‘peak’ and ‘ridge’, we are here talking about the shape of the probability density function, not about the geography of the terrain.)

Now consider what happens when we process a batch of radio altimeter data, known as a **transect**. It is permissible for a transect to comprise a single radio altimeter sample; more typically, however, we have used transects spanning a few seconds of input data, with the radio altimeter sampled at 1–2 Hz, so as to achieve a horizontal separation of about 100 m between the samples. (Closer sampling than this would increase the processing load but yield little accuracy benefit, owing to the limited resolution both of the radio altimeter and the DEM.) The transect data are processed by subtracting the height above ground measured by

the radio altimeter from the aircraft height measured by the INS (as corrected by the integrated navigation system) to yield a **sensed terrain profile**. Fig. 3(a) shows the sensed terrain profile measured during a transect of 4 seconds’ duration, with the radio altimeter sampled at 2 Hz, yielding 9 samples in all. Notice that sensed terrain profile plotted here will be offset from the true terrain profile—the section of terrain profile that was actually being overflown while the transect data were gathered—because of residual position errors in the INS, both horizontal and vertical. In the initial stages of the operation of the algorithm, this *absolute* offset may be of the order of many hundreds of metres horizontally, and many tens of metres vertically. However, the *relative* positions of the points along the sensed terrain profile will normally be in much better agreement with the relative positions of the points along the true terrain profile, although they will still be affected by residual velocity errors in the INS, radio altimeter errors, and errors in the digital elevation map (DEM).

The next stage is therefore to search for horizontal and vertical position offsets that will bring the sensed profile into good agreement with the terrain profile given by the DEM. Fig. 3(b) shows the form of the terrain surface within the area of navigational uncertainty, as given by the DEM. This search proceeds by working through a series of hypotheses about the true position of the aircraft, both vertically and horizontally, at the time of the **mid-transect point**: the time when the fifth of the nine radio altimeter readings forming the transect was sampled. (Working from the middle of the transect helps to minimise the effects of residual INS velocity errors.) For each hypothesis about the mid-transect point, we examine the DEM to determine the terrain profile that would have been overflown during the transect if that hypothesis were true, and compare that DEM terrain profile with the (appropriately offset) sensed terrain profile. This comparison is performed using a statistical model characterising the errors arising from the radio altimeter, and from inaccuracies in the DEM itself, and yields a statistical quantity known as the **likelihood** of the mid-transect point hypothesis.

Different hypotheses about the position of the aircraft at mid-transect yield different values of the likelihood, thus yielding a three-dimensional **likelihood function**. High values of the likelihood function indicate hypotheses where there is good agreement between the shapes of the sensed terrain profile and the DEM terrain profile, and conversely. For the transect shown in Fig. 3, the likelihood function (reduced for presentational clarity to two dimensions) is shown in Fig. 1(b). The reader will observe that there are three areas of relatively high likelihood towards the north of the area shown, and further such regions to the south and to the west of the centre of the region. There is generally rather poor agreement elsewhere, for example towards the south-west corner.

The next step is to use Bayes' Theorem to use the data from the transect to update the state estimate. Stated roughly, the theorem asserts:

$$\begin{array}{ccc} \text{Probability} & & \text{Likelihood} \\ \text{density} & \times & \text{function} \\ \text{function} & & \text{arising from} \\ \text{conditional on} & & \text{data} \\ \text{data} & & \end{array} \times \begin{array}{c} \text{Prior} \\ \text{probability} \\ \text{density} \\ \text{function} \end{array}$$

In other words, to update the probability density function in Fig. 1(a) to take account of the transect data, we simply need to multiply it by the likelihood function in Fig. 1(b), and renormalise it so that it integrates to unity. The result, known as the **posterior** probability density function, is shown in Fig. 1(c). Like the likelihood function, it is very irregular in shape, with a sharp peak towards the centre, an irregular raised region to the west of that, south-west corner, another peak near the centre of the area, a more diffuse raised area towards the north-east, and a slight east-west ridge towards the north-centre.

Unfortunately, the great irregularity of the posterior distribution means that even in two dimensions it requires a lot of data to represent it: within the nine or more dimensions of the system state vector, manipulating this posterior distribution directly would be computationally intractable. The final stage of processing transect data is therefore to approximate the posterior distribution, so that it can be represented once again as a mixture of Gaussian distributions. This is achieved by an iterative procedure, thus leading to the acronym IGMAP: Iterative Gaussian Mixture Approximation of the Posterior.

The result of this approximation is shown in perspective form in Fig. 1(d), and the individual components are shown as the blue ellipses in Fig. 2. As compared with the prior distribution, there is now no peak in the south-west corner, and the peak in the north-east has shifted north-westwards and reduced in volume: this Gaussian component now accounts for only 7% of the total probability. The remaining three Gaussian components are all located near the centre of the region shown: one of them contains 63% of the total probability and corresponds to the sharp peak in Fig. 1(c), and the other two form a compound peak to the west of that. (As it happens, the true aircraft position at mid-transect lies within this latter peak.) Despite the remaining ambiguities, overall uncertainty about the aircraft's horizontal position has been much reduced by this transect.

This completes the processing of the data from the transect in Fig. 3(a). Each of the Gaussian components can now be propagated forward in time using the familiar Kalman filter time update equations, leaving the component weights unchanged. Time updates continue until another batch of transect data is available, whereupon the procedure shown in Fig. 1 is repeated.

If it is desired to combine data from GPS into the mix, this can be accomplished using the standard measurement

update equations for the multi-hypothesis Kalman filter [29, 30], and the same applies for data from many other 'well-behaved' sources of navigation fixes. The full procedure of Fig. 1 is required only for non-linear/ambiguous fixing systems like TRN.

As has already been stated, the overview of the algorithm given in Fig. 1 is a simplification, reduced as it is to two dimensions. More detail of the operation of the algorithm is given in Fig. 4. Between TRN updates, the navigational state is represented as a mixture of n Gaussian distributions over d dimensions, where d is the dimensionality of the system state vector. In Fig. 1 n was 4, but in Fig. 4 n is 2 to keep the diagram simple. The state estimate just before a TRN update (corresponding to Fig. 1(a)) is represented in Fig 4 by the rectangle on the left-hand side containing two shaded ellipses, which represent these n Gaussian components.

Processing of a batch of TRN data (e.g. a TCN transect) now proceeds through the following stages:

Stage 1(a) Each Gaussian component is projected down into the d' -dimensional subspace of the state space on which the TRN measurement directly depends. For example, the sensed terrain profile observed during a TCN transect depends primarily on the *position* errors in the navigation system, so each Gaussian distribution is projected down into the $d' = 3$ dimensions corresponding to the horizontal and vertical position errors, by discarding the remaining elements of the distribution's mean vector and covariance matrix. (The terrain data may also have a second-order dependency on velocity and attitude errors, but we ignore this.)

In the remaining substages of Stage 1, each of the n reduced-dimension Gaussian components is processed separately.

Stage 1(b) Each of the n d' -dimensional Gaussian components is multiplied by the likelihood function generated by the TRN measurements, which is itself defined over d' -dimensions. This yields a per-component true posterior distribution, which will typically be non-Gaussian, and irregular in form.

(In practice, this is not carried out as a distinct stage prior to Stage 1(c): instead, evaluation of the likelihood function is driven by the numerical integration algorithm used in Stage 1(c).)

Stage 1(c) Each per-component true posterior is approximated as a mixture of some number m of d' -dimensional Gaussian components. In Fig. 4, $m = 3$, and the result of Stage 1(c) is shown as the rectangles containing three hollow ellipses.

This approximation is carried out by an iterative procedure involving numerical integration. However,

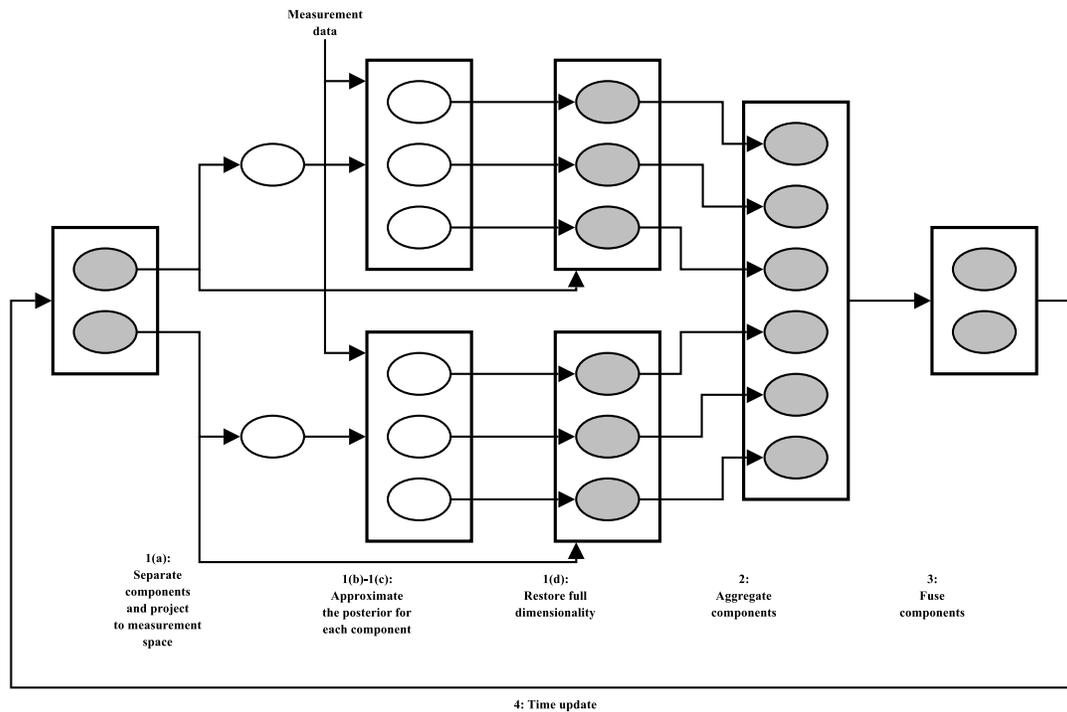


Figure 4 Schematic diagram of the IGMAP algorithm. Gaussian components of the full dimension of the state vector are shown as shaded ellipses, while Gaussian components projected to the measurement space are shown as hollow ellipses.

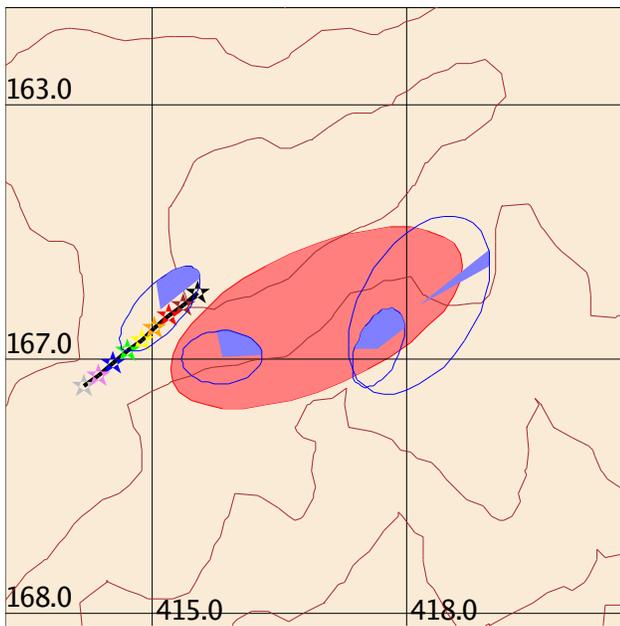


Figure 5 Illustration of transect processing at Stage 1(c). The red ellipse shows the most heavily weighted component of the prior distribution in Fig. 2, and the blue ellipses show an approximation to the corresponding posterior distribution as a mixture of $m = 4$ Gaussian components. The other three components of the prior distribution are processed in a similar way.

since the integration needs to be carried out only over d' dimensions this is numerically very tractable. A special adaptive integration procedure has been developed to support this stage of the algorithm.

Fig. 5 shows the processing of the transect in Fig. 1 up to this stage.

Stage 1(d) Each d' -dimensional component is now converted back to the full dimensionality of the state vector, by reference to the prior component from which it originated.

Stage 2 The results of Stage 1(d) are now assembled together into a weighted mixture of nm d -dimensional Gaussian components. This represents the algorithm's first-cut approximation to the posterior distribution resulting from the TRN measurement.

Fig. 6 shows the processing of the transect in Fig. 1 up to this stage.

Stage 3 Unfortunately, the result of Stage 2 contains nm components rather than the n we started with. Obviously, if this algorithm is to be used recursively, we cannot allow the number of components to increase on each iteration. At Stage 3, therefore, the algorithm chooses the two components of this mixture which are most similar to each other and **fuses** them together, replacing them by a single component whose weight is

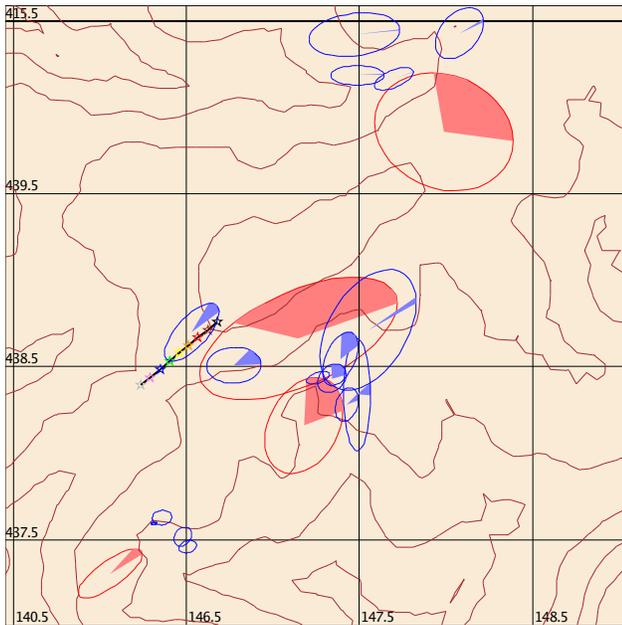


Figure 6 Illustration of transect processing at Stage 2. The red ellipses shows the prior distribution as in Fig. 2, and the blue ellipses show a first-cut approximation to the corresponding posterior distribution as a mixture of $nm = 16$ Gaussian components. Among these blue ellipses the four from Fig. 5 are visible (but with scaled-down weights); the remaining twelve come from the other prior components. The four ellipses towards the south-west corner have very low weights, and the shaded sector indicating the weight is visible barely if at all.

equal to the sum of the weights of the fused components. This pairwise fusing (or merging) is repeated until the mixture is brought back down to n components.

The result of this stage is represented in Fig 4 by the rectangle on the right-hand side containing two shaded ellipses: this corresponds to Fig. 1(d).

For the transect in Fig. 1, the result of this stage of processing is shown as the blue ellipses in Fig. 2: it will be seen that the 16 ellipses in Fig. 6 have been reduced back to four.

This completes the processing of the batch of TRN data.

Stage 4 Each of the n Gaussian components undergoes time updates (using the standard Kalman filter time update equations). This continues until more TRN data are available.

5 THE IGMAP ALGORITHM IN ACTION

As the introduction has mentioned, [1, Sec. 7] reported results obtained using the IGMAP algorithm using simulated data. It included results for $10^\circ/\text{hr}$ and $100^\circ/\text{hr}$ inertial systems with various INS alignment scenarios and various levels of terrain roughness, and compared the IGMAP results with those obtained using a best-fix and a PDAF TCN algorithm. In these results, IGMAP consistently performed as well as or better than either of the other algorithms.

In this section, we describe the results of applying IGMAP to real flight data recorded during a sortie of a QinetiQ Tornado GR1 aircraft over southern Britain. The equipment for the trial included a Honeywell H764G Embedded GPS/INS (EGI) and a BAE Systems AD1990 radio altimeter, both mounted in a pod fitted under the fuselage of the aircraft. The H764G incorporates a GPS receiver and a ring-laser gyro INS, and provides both a blended GPS/INS navigation output and a pure inertial output (with barometric damping of the vertical channel). GPS signals were obtained via a typical fixed reception pattern antenna (FRPA). The IGMAP algorithm was applied to combine data from the AD1990 with the pure inertial output from the H764G, using DTED Level 1 [14] (compiled *circa* 1986) as the digital elevation map. The blended INS/GPS output from the H764G was used as the ‘truth’ measure with which the IGMAP output is compared.

Although the INS in this trial was of aircraft grade, the model of the INS incorporated in the Kalman filter system model was very pessimistic, particularly as regards the initial conditions: the initial position was assumed to be accurate only to within ± 2 km (2σ) in each horizontal axis, and to within ± 200 m (2σ) in height. The initial velocity was assumed to be accurate only to within ± 20 m/s (2σ) in each horizontal axis, and ± 2 m/s (2σ) vertically. Although the H764G was barometrically aided, the height channel in the Kalman filter model does not assume this.

The data from this sortie, of duration just over $1\frac{1}{2}$ hours, were analysed using the IGMAP algorithm, with the system state estimated as a mixture of $n = 4$ Gaussian components. (In fact, Fig. 1 is based on data from the third transect of this sortie.)

First consider the initial capture phase, immediately after terrain data becomes available. The position errors during the first 45 seconds are shown in Fig. 7. In the figure the black lines represent the components of position error, based on comparing the overall mean of the 4-component Gaussian mixture with the ‘true’ position given by the INS/GPS blended data. The green band represents a 2σ tolerance band based on the overall standard deviation of the Gaussian mixture.

It will be noted that already after the first 4 second transect has been analysed, there has been a substantial reduction

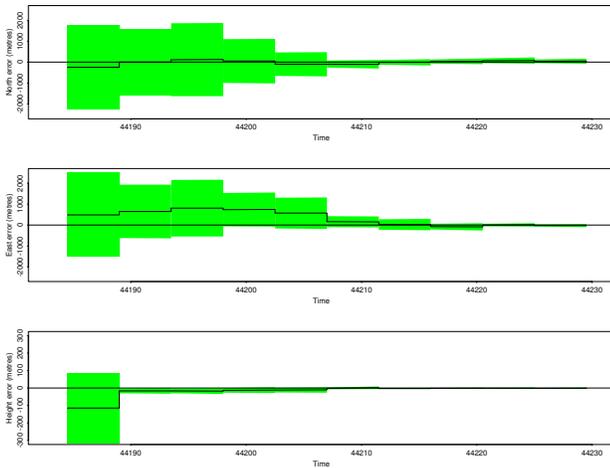


Figure 7 Position errors during the first 45 seconds after terrain-aided navigation begins. (Time is measured as seconds after midnight.)

the horizontal position uncertainty, particularly in the east-west axis, as well as in the vertical position uncertainty. After just 22.5 seconds (five transects), the position uncertainty is of the order of ± 220 m (2σ) in each horizontal axis, ± 6 m (2σ) vertically. After 45 seconds these uncertainties are further reduced to ± 48 m and ± 2.6 m respectively.

Now let us consider the sortie as a whole. A characteristic of the sortie was the high level of manoeuvre, with the aircraft making numerous sharp turns and climbs and dives. Consequently, for a large proportion of the flight, the radio altimeter could not provide usable data, either because the aircraft's height above ground was too great (i.e. above about 500 m), or its bank angle too great (i.e. greater than about 30°). Fig. 8, which shows the components of position error throughout the sortie, indicates these parts of the sortie by colouring the 2σ tolerance bands red. Yellow coloration indicates that the radio altimeter was usable, but that the terrain roughness was no greater than 1%. (Terrain roughness is here measured as the RMS terrain gradient along the true path of the aircraft during a transect, with the true path measured by INS/GPS.) The remaining periods are when the radio altimeter was usable and the terrain roughness exceeded 1%; these periods together accounted for about a quarter of the sortie duration, and are indicated by green colour.

Obviously, during the 'red' periods, the horizontal and vertical position uncertainties increase continuously as a result of INS drift, though the rate of drift decreases as the sortie progresses because the INS becomes better calibrated, particularly the velocity errors. During the yellow periods, horizontal uncertainties normally continue to grow, but vertical errors are kept in check. Only during the green periods is the horizontal navigation materially assisted by the TCN

data.

Fig. 8 clearly indicates the ability of the IGMAP algorithm quickly to recover accurate navigation once 'green' data comes along following a period of drift. Taking all the green periods together, the radial horizontal position error had a median value of just under 28 m.

The RMS height error during these periods works out as 5.4 m; however, it is noticeable in Fig. 8 that the height error appears to have a slowly varying bias: this is believed to be due to GPS errors influencing the measurement of 'true' height. Consequently the height accuracy is probably better than this.

6 CONCLUSIONS AND PROSPECTS

In this paper we have described the design objectives that led to the development of the IGMAP algorithm: in a nutshell, this was to have a data fusion algorithm that could cope with the ambiguities inherent in TRN data, but at the same time make efficient use of the available data, in a manner compatible with the use of a Kalman filter (or MHKF) architecture for multiway integrated navigation. The paper has given an overview of the operation of the IGMAP algorithm, and described its performance when applied to recorded data from a fast-jet sortie with some challenging characteristics. The results illustrated the algorithm's ability rapidly to acquire and reacquire accurate navigation from high initial position and velocity uncertainties.

The following are some areas for possible exploration in the future:

- The studies of the use of IGMAP with medium- and low-grade IN systems reported in [1] were based on simulated data. It would be desirable to corroborate its conclusions using real data, particularly real radio altimeter data, or—better still—laser rangefinder data.
- All studies of IGMAP applied to TRN have so far been based on terrain-contour navigation (TCN): it would be interesting to apply it to an imaging-based technique such as CVN [13].
- IGMAP may also be applicable to some target-tracking problems.

ACKNOWLEDGEMENTS

This work was funded by the Applied Research Programme of the UK Ministry of Defence.

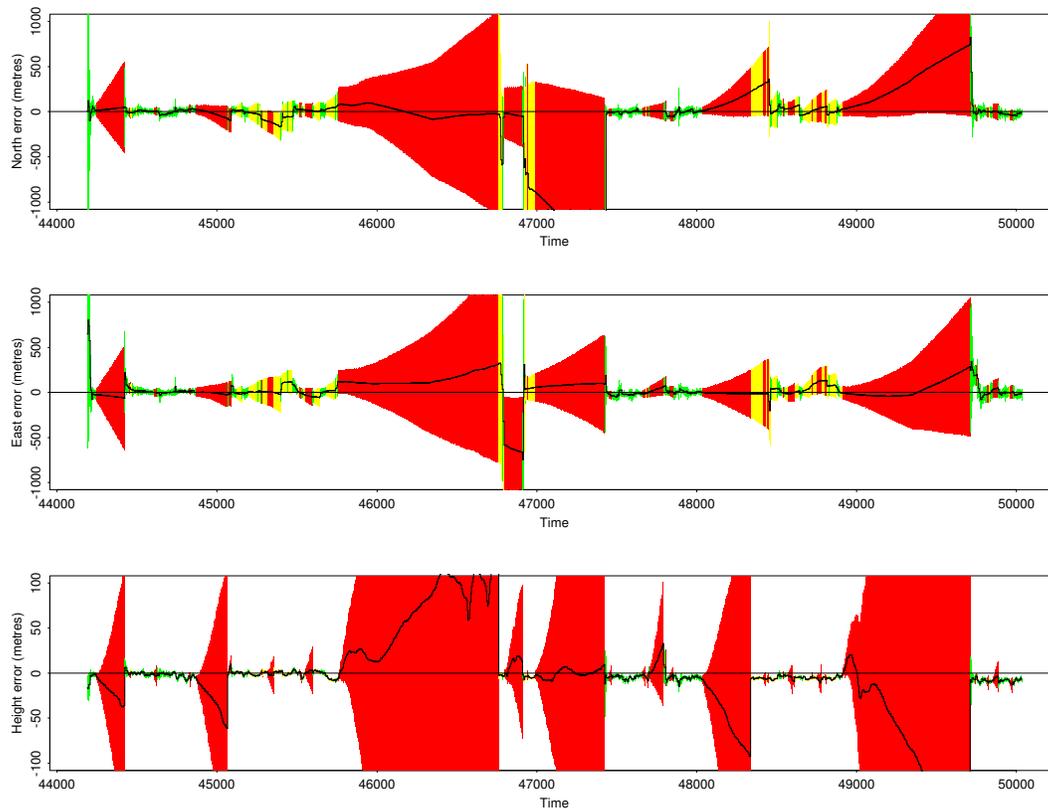


Figure 8 Position errors throughout the sortie. Timescale is in seconds. Red colour indicates periods in which the radio altimeter was unusable. Yellow colour indicates terrain roughness $\leq 1\%$; green indicates terrain roughness $> 1\%$.

REFERENCES

- [1] Paul D. Groves, Robin J. Handley, and Andrew R. Runnalls. Optimising the integration of terrain-referenced navigation with INS and GPS. In *Proceedings of Institute of Navigation National GNSS2004*, pages 1048–59, September 2004. Long Beach, California, USA.
- [2] James Carroll et al. Vulnerability assessment of the transportation infrastructure relying on the global positioning system. Technical report, Volpe National Transportation Systems Center, August 2001. Report for US Department of Transportation.
- [3] John I. R. Owen and Michael Wells. An advanced digital antenna control unit for GPS. In *Proceedings of the 2001 ION National Technical Meeting*, pages 402–7, 2001.
- [4] Michael Wells and John Owen. A digital adaptive array demonstrator for GPS. In *Proceedings of NATO SET Symposium 'Emerging Military Capabilities Enabled by Advances in Navigation Sensors'*, October 2002.
- [5] Neil J. Boasman and Peter Briggs. The development of an anechoic GPS test facility. In *Proceedings of the ION 58th Annual Meeting and CIGTF 21st Guidance Test Symposium, 2002*.
- [6] Paul D. Groves and Daniel Long. Adaptive tightly-coupled, a low-cost alternative anti-jam INS/GPS integration technique. In *Proceedings of ION National Technical Meeting, 2003*.
- [7] Paul D. Groves and Daniel Long. Combatting GNSS interference with advanced inertial navigation. *Journal of Navigation*, 58(3), September 2005.
- [8] Jim Sennott and Dave Senffner. A GPS carrier phase processor for real-time high dynamics tracking. In *Proceedings of ION 53rd Annual Meeting*, pages 299–308, 1997.
- [9] Donald Gustafson, John Dowdle, and Karl Flueckinger. A high anti-jam GPS-based navigator. In *Proceedings of ION National Technical Meeting*, pages 438–46, 2000.
- [10] Andrey Soloviev, Frank van Graas, and Sanjeev Gunawardena. Implementation of deeply integrated GPS/low-cost IMU for reacquisition and tracking of

- low CNR GPS signals. In *Proceedings of ION National Technical Meeting*, pages 923–35, 2004.
- [11] Robin J. Handley, Paul D. Groves, Paul McNeil, and Lawrence Dack. Future terrain referenced navigation techniques exploiting sensor synergy. In *Proceedings of GNSS 2003, The European Navigation Conference*, April 2003.
- [12] Jacob Campbell, Maarten Uijt de Haag, and Frank van Graas. Terrain reference navigation using Airborne LAsER SCanner (ALASCA)—preliminary flight test results. In *ION 60th Annual Meeting*, pages 671–8, June 2004.
- [13] Robin J. Handley, Lawrence Dack, and Paul McNeil. Flight trials of the continuous visual navigation system. In *Proceedings of the ION National Technical Meeting*, pages 185–92, 2001.
- [14] Defense Mapping Agency. *Product Specifications for the Digital Landmass System (DLMS) Data Base*, April 1983.
- [15] Jürgen Metzger, Jan Wendel, and Gert F. Trommer. Comparison of different terrain referenced navigation techniques. In *Proceedings of ION 58th Annual Meeting*, pages 132–41, 2002.
- [16] Alan Robins. Recent developments in the ‘TERPROM’ integrated navigation system. In *Proceeding of the ION 44th Annual Meeting*, June 1998.
- [17] Larry D. Hostetler and Richard C. Beckman. Expanding the region of convergence for SITAN through improved modelling of terrain nonlinearities. In *NAECON*, pages 1023–30, 1979.
- [18] L. D. Hostetler and R. D. Andreas. Non-linear Kalman filtering techniques for terrain-aided navigation. *IEEE Transactions on Automatic Control*, AC-28(3):315–23, March 1983.
- [19] Jeff Hollowell. HELI/SITAN: a terrain referenced navigation algorithm, for helicopters. In *Proceedings of the IEEE Position Location and Navigation Symposium (PLANS)*, 1990.
- [20] Niclas Bergman. *Recursive Bayesian Estimation: Navigation and Tracking Applications*. PhD thesis, Linköping University, Sweden, 1999.
- [21] Pei-jun Yu, Zhe Chen, and James C. Hung. Performance evaluation of six terrain stochastic linearisation techniques for TAN. In *Proceedings of the IEEE National Aerospace and Electronics Conference (NAECON)*, 1991.
- [22] Jürgen Metzger and Gert F. Trommer. Improvement of modular terrain navigation systems by measurement decorrelation. In *Proceedings of ION 59th Annual Meeting*, pages 353–62, 2003.
- [23] J. P. Golden. Terrain contour matching (TERCOM): A cruise missile guidance aid. *Proc. Int’l Soc. for Optical Eng. (SPIE)*, 238, July/Aug 1980.
- [24] A. R. Runnalls. A Bayesian approach to terrain contour navigation. In *Guidance and Control Panel, 40th Symposium*. NATO Advisory Group for Aerospace Research and Development, May 1985.
- [25] Andrew R. Runnalls. Apparatus for measuring the state of a dynamic system. UK Patent Application GB 2178571A, February 1987.
- [26] Andrew R. Runnalls. Apparatus incorporating recursive estimators. United States Patent 4786908, November 1988.
- [27] Pelle Carlbom and Nicklas Johansson. Terrain navigation and integrity monitoring for avionic platforms. In *Proceedings of NATO SET Symposium ‘Emerging Military Capabilities Enabled by Advances in Navigation Sensors’*, October 2002.
- [28] A. R. Runnalls and R. J. Handley. The ‘gold standard’ navigator. In Mark Bedworth and Jane O’Brien, editors, *Proceedings of Eurofusion98*, pages 77–82, 1998.
- [29] Samuel Blackman and Robert Popoli. *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.
- [30] Lawrence D. Stone, Carl A. Barlow, and Thomas L. Corwin. *Bayesian Multiple Target Tracking*. Artech House, 1999.