Return to Session Directory

DYNAMIC POSITIONING CONFERENCE
October 15-16, 2013

Sensors I SESSION

Performance modelling for local position reference sensors applied to RadaScan

By Russ Miles

Guidance Navigation Ltd

Abstract

When first presented with a measurement device of any sort, one very natural question to ask is “How accurate is it?” It’s often easy enough to give a simple, order of magnitude answer, but a fully specified, testable answer is necessarily more detailed.

In this paper we discuss the selection of performance metrics for a position reference sensor with a particular focus on the specification of accuracy.

The use and benefits of a quantitative model of measurement errors is examined, in particular data fusion and measurement validation. This is in the context of the Kalman filter and alternative estimation techniques.

The experimental technique used to obtain a detailed error model is described and illustrated by results obtained from the RadaScan microwave position reference sensor.

The most important environmental factors which affect the accuracy of RadaScan are described, particularly at the effect of sea reflection multi-path.

Data from both land and sea trials are presented along with recordings from routine operation at sea. These show how RadaScan performs across a range of representative operating circumstances.

Finally a summary performance model is specified which gives a usable guide to the accuracy the RadaScan can be relied upon to provide.

Performance metrics

The performance of a GPS navigation system can be characterised in terms of:

- Accuracy

How close are the GPS fixes to the true position of the vessel?

- Availability

How often is the GPS ready when I want to start using it?

- Continuity

How often does the GPS continue to operate throughout an operation?

- Integrity

How often does the GPS give me a fix which is completely wrong?

More fully specified definitions of these terms can be found in [1] and [2]. These parameters are defined with respect to an operating area (i.e. coverage).

In the case of a local position reference sensor, availability, continuity and integrity depend on local operational factors, e.g. how well maintained are the reflectors? What obstructs the view from sensor to reflector? We aren't in a position to provide statistics for these parameters across all installations, and those statistics probably wouldn't be very useful anyway. In this paper we will concentrate on accuracy.

For Dynamic Positioning a constant offset error in a local position reference is of little concern. For excellent station keeping performance we require:

- Local accuracy

The measurement of displacement over several metres is accurate.

- Repeatability

If we return to the same actual position, we get the same reported position as before.

So when modelling position measurement errors, we don't need to worry about offset errors which vary slowly as a function of position.

Kalman filters

The Kalman filter has been used for dynamic positioning since the 1970s [3], [4]. Cadet [5] gives a very readable introduction to Kalman filters as applied to DP.

In the DP application, the main attraction of the Kalman filter approach is that it combines the sequence of measurements from a position reference sensor with a model of likely motions of the vessel. In this way it is usually possible to filter out high frequency measurement noise. More importantly, it is also possible to rule out measurement blunders in an intelligent way, i.e. rogue measurements where the error is unusually large. In a Kalman Filter framework a validation gate can be set to best balance the risk of accepting a bad measurement against the risk of rejecting a good measurement, see for example Bar-Shalom and Fortmann [6].

The Kalman filter is also a popular mechanism for multi-sensor data fusion where the measurements from each of the position sensors are combined to give a best overall estimate. This is the so-called “blended” mode. This is very effective in applications where different sensors have complementary performance profiles across the operating range – where at all positions there is always at least one sensor which gives accurate measurements, but no sensor gives accurate readings everywhere. In the DP application, this is not so important. Multiple sensors are used in DP to provide redundancy and a voting arrangement is perhaps more appropriate than a blended mode.

In the Kalman filter the error in a sensor is modelled as white Gaussian noise. In configuring the Kalman filter it is necessary to specify a measurement covariance for each sensor. Experience has shown that it is not necessary to fine-tune the noise covariance terms in a Kalman filter to the nth degree. A reasonable approximation is good enough. And of course the measurement noise process is seldom either Gaussian or white. In practice very usable results can be obtained by increasing the measurement covariance in the Kalman filter to compensate.

Over the past decade or so new estimation techniques have become available which allow for non-Gaussian distributions for the measurement errors. A recent paper by Rigatos [7] compares the Particle Filter against the Kalman filter for dynamic positioning.

This sort of Kalman filter validation gate compares the measurement prediction error against the a priori measurement prediction variance. That is, it compares how much the measurements jump around compared to how much they are expected to jump around according to the model. This can be contrasted with an approach which relies on comparing the latest measurement prediction error against the sample measurement prediction variance. That approach compares how much each measurement jumps around compared to how much measurements have been jumping around recently.

Other statistical tests (w- and F- tests) can also be applied to a sequence of position measurement innovations to check that they are indeed distributed in an approximately Gaussian fashion with a stable variance. See, for example, Russell [8].

Elliptical errors

The error in a position fix obtained from a GPS receiver is said to be circular. The easting and the northing have approximately the same accuracy, making the confidence region around a point estimate circular.

With a range-bearing local position reference sensor that is not the case. When the range-bearing measurement is interpreted as a position fix, the error in the radial direction may be of quite a different size to the error in the tangential direction. Radial error is more or less constant over a wide operating range. For a given bearing error, tangential position error is proportional to the distance from the sensor to the responder.

To illustrate this, imagine that across the normal range of operating circumstances that the typical range error is 0.25m and that the typical bearing error is 2.5mrad. At 50m the tangential position will be accurate to 0.125m. At 200m the tangential position will be accurate to 0.5m.

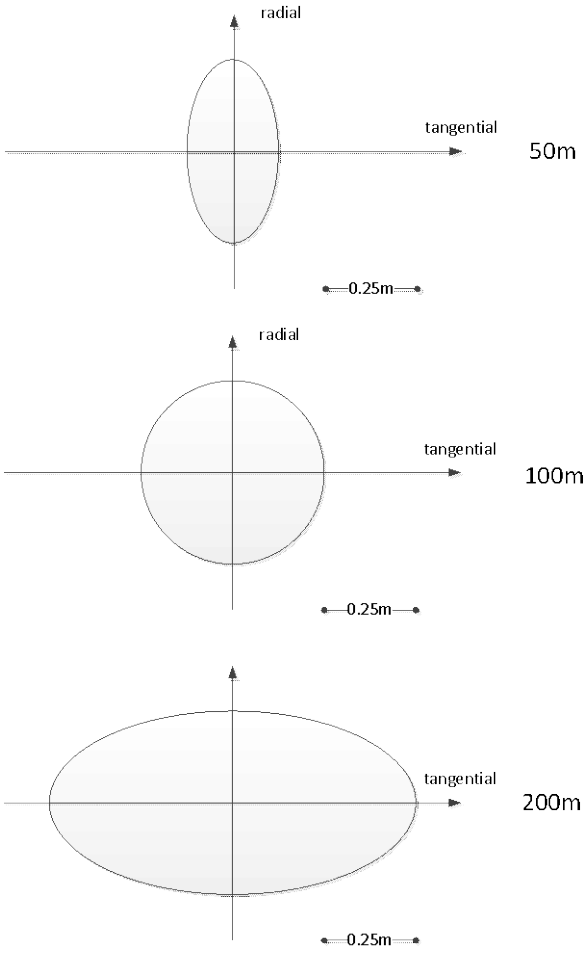

When blending measurements from a local position reference sensor with those from GPS, it is best to calculate the weight for the range and bearing components of the measurement separately. This applies whether weights are assigned on the strength of an a priori measurement noise model or in response to the observed variation of measurements.

Modelling the errors

If we are to use a position reference sensor in a Kalman filter, then we need to be able to calculate a measurement noise covariance for each measurement. If we use an alternative estimation technique to the Kalman filter, we need some other way to represent the probability density function of the measurement prediction error. Even if we use a simpler technique, we need to be able to set an objective standard for what counts as a good measurement even if the performance of the device is validated manually. We choose to do this by specifying a covariance matrix.

We are concerned only for local accuracy and repeatability, so we can treat the mean error as if it is zero.

With a scanning local position reference sensor such as CyScan or RadaScan, we obtain one range-bearing measurement to each reflector from each scan. The device scans once per second. The error on one scan is independent of the error on the previous scan. That is to say, the measurement noise is "white".

We accept that underlying random process may not be perfectly Gaussian. We aim to compose a covariance matrix such that no more than 5% of range measurement or of bearing measurements lie more than 2 standard deviations away from the true value.

The noise process is likely to be heteroskedastic, that is, the errors may be larger at some times than at others. We model most of the most important factors. But even accounting for these, the typical size of the errors varies due to influences that we don't model.

We expect range errors and bearing errors to be uncorrelated.

The modelling task is to identify the most important factors influencing the size of the errors and to quantify those effects so that we can supply a measure of the confidence we have in each measurement in the form of a covariance matrix. This requires extensive experimentation and observation of performance in a realistic working environment. Let's now see how this unfolds in the case of the RadaScan system.

You can find an overview of the RadaScan system in IMCA M209 [9].

RadaScan accuracy on land

We've done extensive tests with RadaScan at test track at a former airfield. In the typical experimental arrangement a responder is placed at a fixed position at a height above the ground of about 1.5m. The sensor mounted on a truck at about the same height above the ground. The truck then moves away from the responder in a straight line across the airfield. For the most part the truck maintains a constant speed. It pauses for a few minutes every so often so that we can obtain measurements at a fixed position. We can go to distances well in excess of 1km. The airfield is substantially flat over these distances.

Bear in mind that these are very favourable circumstances for RadaScan. Operation at sea is rather more challenging. Still, in these friendly conditions RadaScan accuracy performance is very impressive.

At 30m

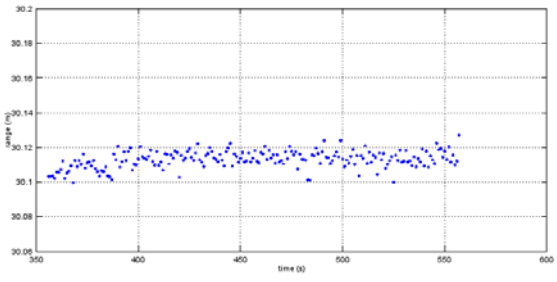

standard deviation: 0.005m

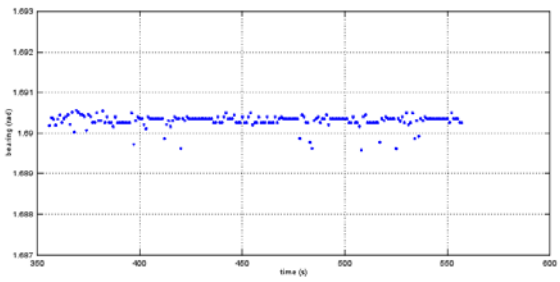

standard deviation: 0.16mrad

At 100m

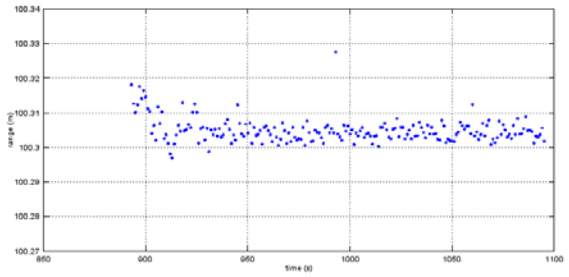

standard deviation: 0.004m

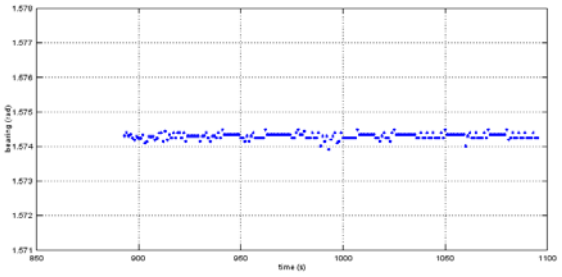

standard deviation: 0.09mrad

At 300m

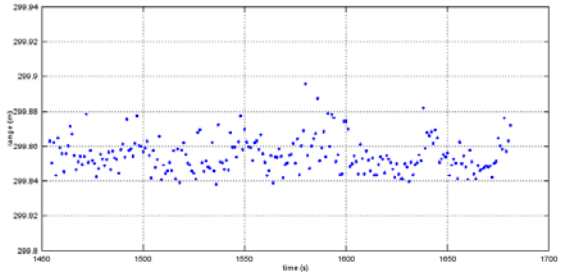

standard deviation: 0.010m

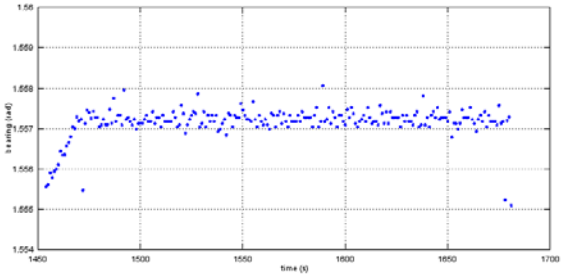

standard deviation: 0.39mrad

At 600m

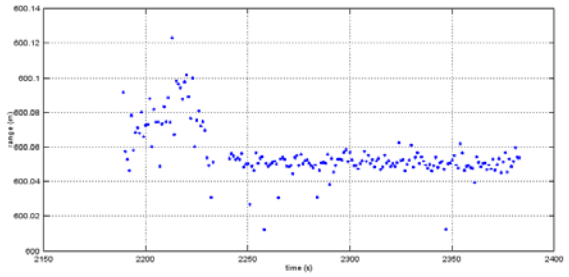

standard deviation: 0.014m

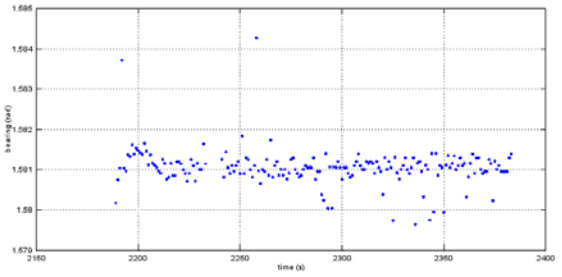

standard deviation: 0.46mrad

Bruntingthorpe Mini RadaScan 2.12.3.11 21/11/12

In these conditions we see centimetre level precision in position at ranges typical of operating circumstances, and comfortably sub-metre precision all the way out to 600m.

Responder angle of incidence

We get the maximum signal power at the sensor when the responder is facing directly towards the sensor. The signal level falls off as the responder is pointed away from the sensor. This is illustrated in the following plot:

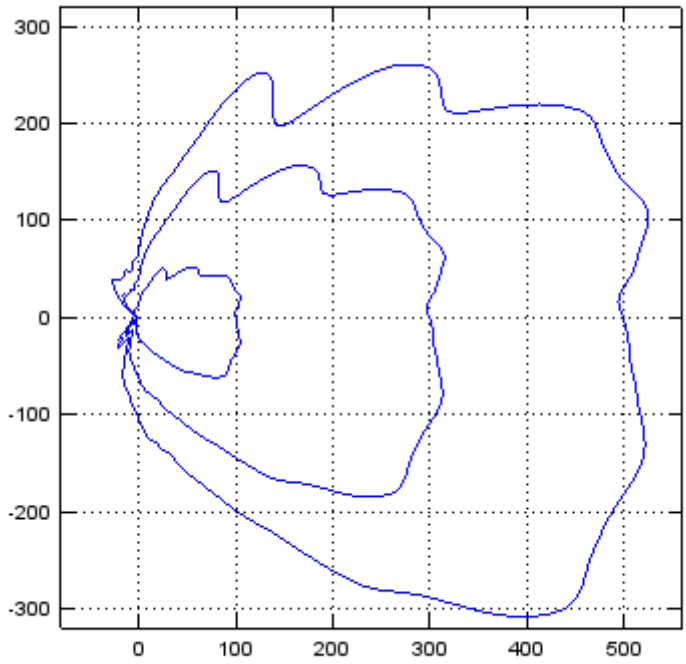

Narrow beam responder antenna pattern May 2013

The responder is at the origin. The curve shows the positions at which the signal power received by the sensor is equal to that received at boresight at 500m.

The natural tendency for any measurement system is for the errors to increase as the signal power decreases and RadaScan is no exception. The RadaScan sensor monitors the signal to noise ratio of the reflection. When this falls below a threshold, the measurement is suppressed and the fix is not supplied to the DP system. This threshold is set so that measurements cut out before the errors grow much above the level seen at high signal power.

The signal to noise ratio threshold can be reduced to make these reduced accuracy fixes available to the DP system. The DP system needs to handle these fixes carefully, allowing for relatively large measurement errors.

Radar position reference sensors and multi-path

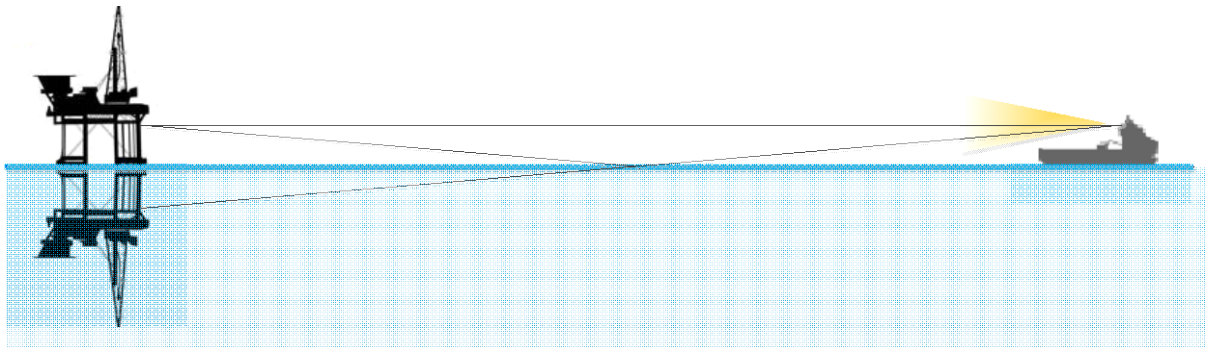

It is well known in the field of microwave radar that ground or sea reflection at very low elevation angles produces deep nulls where the radar target disappears at particular combinations of height and distance. See Kingsley [10]. The direct signal from the target interferes with the signal reflected from the water to produce "Lloyd's mirror" interference fringes. This phenomenon affects all microwave radars. For example, see [11]. An arrangement of antennas at different heights can mitigate the effect and give a signal which is adequate throughout the operating range [12].

Multi-path nulls are seldom a practical problem for DP vessels in any case. The effect is most prominent at long distance at low height. A radar position reference sensor is typically mounted at least 15m above the surface of the water. At this height we can rely on the frequency diversity of the FMCW radar to deliver a usable signal at least as far as 1000m. Height diversity measures are not required for ordinary work boat operations, at least, not with RadaScan.

At more modest ranges (200m to 300m), sea reflection does have a significant effect on accuracy. Here the direct image and the reflected image are not so close that they cancel each other out in destructive interference. But they are closer the range resolution capability of the radar. That means that the two images interfere with each other and the range measured to either of them is distorted. Resulting range errors are a few decimetres. Again, this phenomenon comes from the underlying physics which affects all microwave radar systems. The effect is a function of bandwidth rather than of radio frequency. Various mitigating measures are taken within RadaScan, but it remains the fact that operation in this range is the most challenging in terms of range accuracy. And in this range, sea reflection is by far the most important source of measurement error, so tests done on land at about 2m above the ground are not a good indication of the accuracy which is achievable in typical operation at sea.

At shorter ranges, sea reflection is less of a problem. First of all the difference in distance between the direct image and the reflected image becomes great enough for the two images to be separated, allowing accurate measurement of both ranges. Secondly the reflected image becomes much fainter. As the grazing angle increases, the reflectivity of the water decreases. Also the reflection moves to the edge of the main lobe of the radar antennas.

To get a measure of the effect of sea reflection on RadaScan accuracy, we took some measurements with a vessel moored in harbour. Sea reflection is thought to be at its most troublesome when the sea is calm. Waves tend to break up the reflection. So conditions in the harbour approach our worst case. The sensor and responder were relatively low (average height above the water of 15m), which brings the onset of measurement errors due to sea reflection closer.

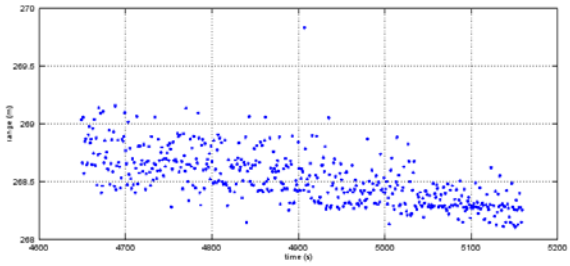

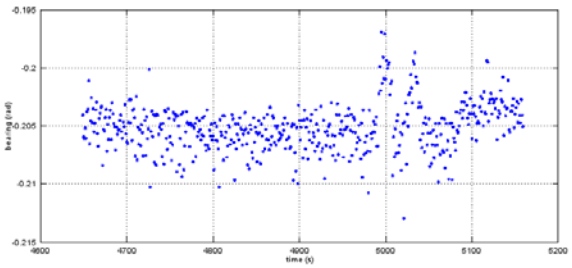

standard deviation: 0.230m

standard deviation: 2.06mrad

mini RadaScan Aberdeen 11/06/13

Both range and bearing errors are quite a lot larger than in the most comparable case at the airfield. This is nearly all due to sea reflection.

The impact of sea reflection on accuracy depends on the region of operation:

short range	**medium range**	**long range**
less than 125m	125m to 250m	250m to 1000m
Very little radiation reaches the sea surface.		
Reflected image is a few metres further away than the direct image and can easily be rejected.	Reflected image is a metre or two further away than the direct image. It can usually be rejected, but it can distort the direct image.	Reflected image is only a fraction of a metre further away than the direct image. The images can't be separated.
No errors from sea reflection	Occasional errors	Frequent small errors

Accuracy in a typical operational setting

Here are some range and bearing measurements typical of station-keeping operation in the North Sea:

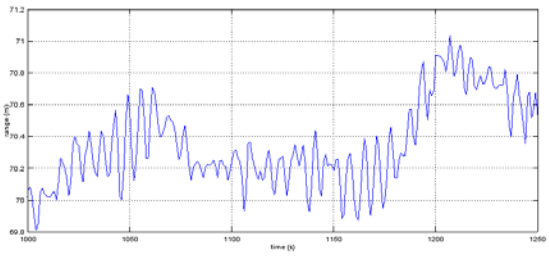

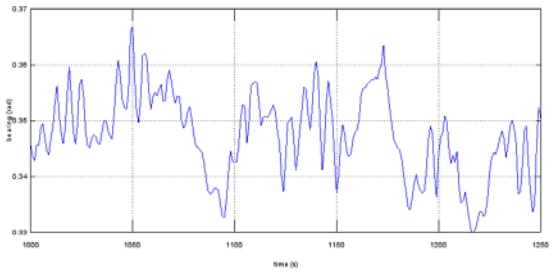

standard deviation ≤ 0.085m

standard deviation ≤ 2.41mrad

Edda Frende 01/08/13

The wave motion is fairly strong, making it more difficult to quantify the accuracy performance. In a comparison with an alternative sensor (e.g. GPS), timing and other registration errors dominate over RadaScan measurement errors. We can apply a high-pass filter which has unity gain for white noise while attenuating the wave motion. This gives us an upper bound on the standard deviation of the range and bearing measurements.

Here are some further examples of RadaScan performance in North Sea operating conditions:

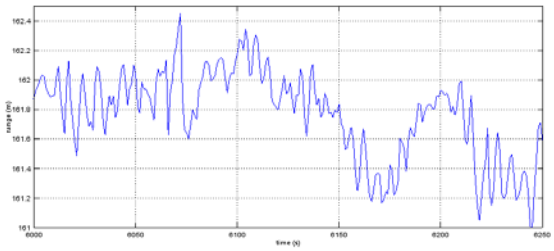

standard deviation ≤ 0.095m

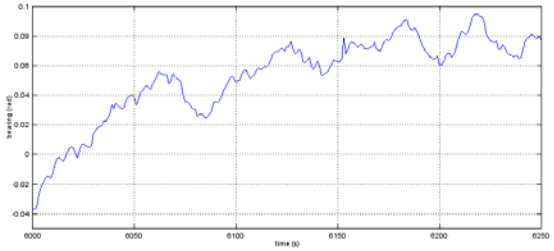

standard deviation ≤ 2.09mrad

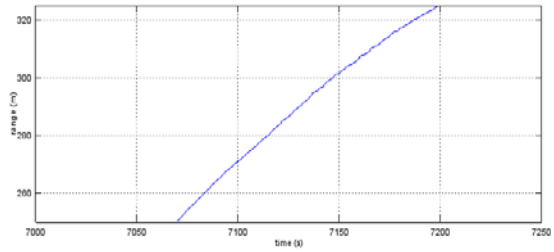

standard deviation ≤ 0.159m

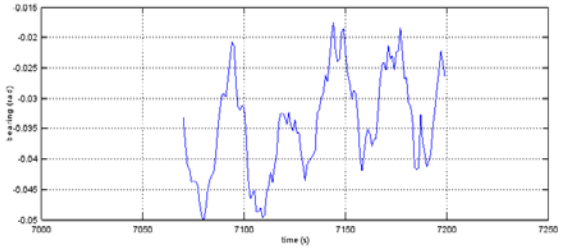

standard deviation ≤ 2.05mrad

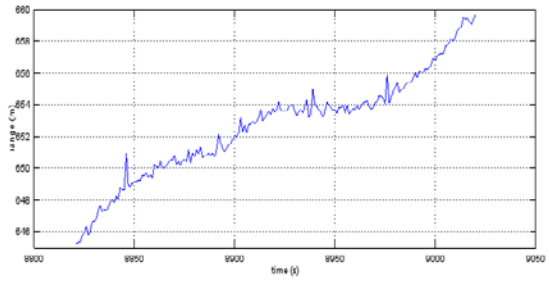

standard deviation ≤ 0.318m

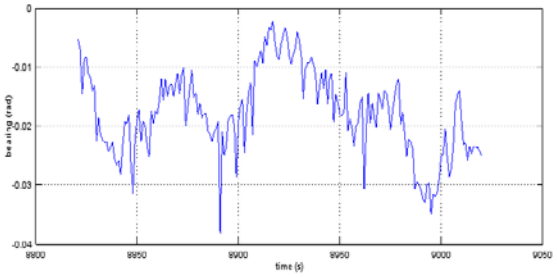

standard deviation ≤ 3.04mrad

In the second of these examples the vessel is on the move which makes it difficult to pick out the detail in the range plot. In that case it is helpful to remove the trend from the range plot:

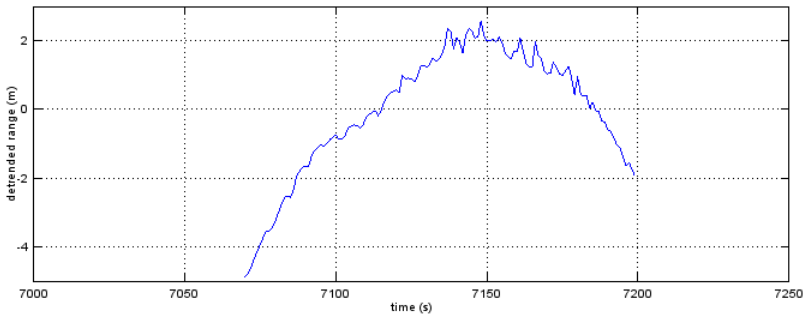

The effect of sea reflection on range accuracy is variable. The following data set shows the effect very prominently:

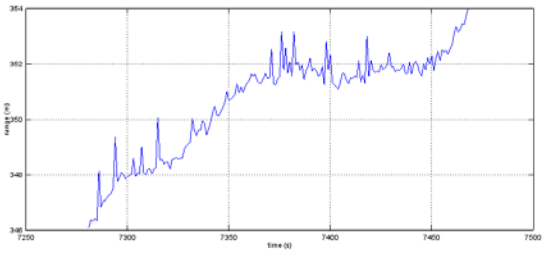

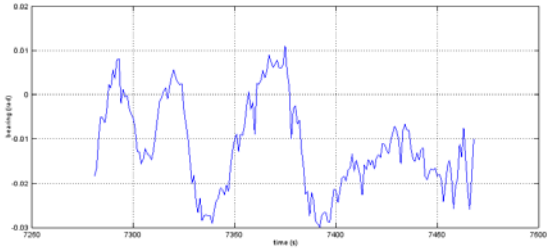

standard deviation ≤ 0.362m

standard deviation ≤ 2.51mrad

The range measurement can jump by a metre, although it nearly always jumps straight back again.

RadaScan accuracy model

Under typical operating conditions mini RadaScan can be expected to provide 1-σ accuracy of 0.25m in range and 2.5mrad in bearing out to a range of 500m. At greater distances errors grow roughly in proportion to distance.

As the distance decreases below 250m, the effect of sea reflection diminishes and accuracy improves. Sea reflection has no significant effect at less than 125m and 1-σ accuracy of 0.1m in range and 1mrad in bearing are commonly achieved.

While RadaScan is configured to supress measurements where the signal level is low, no model of the increase in range-bearing noise with low signal level is required.

We can express our model in the form of a measurement noise covariance suitable for a Kalman filter observer.

$$R = \begin{bmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_\theta^2 \end{bmatrix}$$

$$\sigma_r = \begin{cases} 0.1 & r < 125m \\ 0.001r - 0.025 & 125m \leq r < 375m \\ -0.0008r + 0.55 & 375m \leq r < 500m \\ 0.0005r & 500m \leq r \end{cases}$$

$$\sigma_\theta = \begin{cases} 1 & r < 125m \\ 0.012r - 0.5 & 125m \leq r < 250m \\ 2.5 & 250m \leq r < 500m \\ 0.005r & 500m \leq r \end{cases}$$

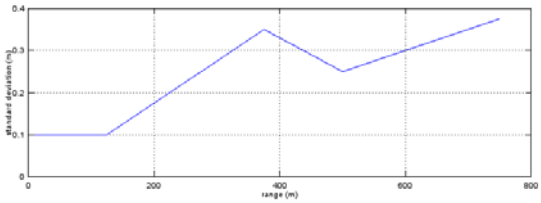

σ_r

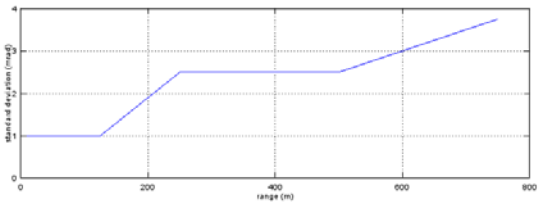

σ_θ

References

1. GNSS Performances

ESA Navipedia, 2011, http://www.navipedia.net/index.php/GNSS_Performances

2. 2008 Federal Radionavigation Plan

National Technical Information Service, Springfield, Virginia 22161, DOT-VNTSC-RITA-08-02/DoD-4650.5

3. Dynamic positioning using Kalman filtering and optimal control theory

J G Balchen, N A Jenssen, S Sælid

IFAC/IFIP Symposium on Automation in Offshore Oil Filed Operation, 1976

4. Use of Kalman filtering techniques in dynamic ship positioning systems

M J Grimble, R J Patton, D A Wise

IEE Proceedings D (Control Theory and Applications), Volume 127, issue3, May 1980, p. 93 – 102

5. Introduction to Kalman Filter and its Use in Dynamic Positioning Systems

Olivier Cadet

Dynamic Positioning Conference, September 16-17, 2003

6. Validation of Measurements

Yaakov Bar-Shalom and Thomas Fortmann

Tracking and Data Association, section 5.1, 1988, Academic Press, ISBN 0-12-089760-7

7. Sensor fusion-based dynamic positioning of ships using Extended Kalman and Particle Filtering

Gerasimos G. Rigatos

Robotica Volume 31 Issue 03 May 2013, pp 389-403

8. DGNSS Position Quality Information for DP Applications

David Russell

Dynamic Positioning Conference, October 9-10, 2007

9. RadaScan Microwave Radar Sensor for Dynamic Positioning Operations

IMCA M209, April 2011.

10. Multipath.

Simon Kingsley and Shaun Quegan

Understanding radar systems, section 7.4, Feb 1999, SciTech Publishing.

11. Guidance on RADius Relative Positioning System

IMCA M 224 – February 2013

12. Active Target With Height Diversity

Patent application EP 2369366, David Patrick, Guidance IP Ltd, published Sept 2011.