



DYNAMIC POSITIONING CONFERENCE
November 15-16. 2--5

Sensors II

CyScan: The Benefits of Multiple Hypothesis Tracking
for Laser-Based DP Reference Sensors

J Grothusen

Guidance Navigation Ltd (Leicester, UK)

Introduction

CyScan is a high performance local position reference sensor specifically engineered for marine vessel control applications at offshore installations, oil rigs or other offshore platforms. To achieve this high accuracy the system employs a number of data correction methods and tracking algorithms as detailed below.

It comprises a continuously rotating 360 °/s fanned laser measuring the range and bearing to retro-reflective targets. The range is determined by measuring the 'time-of-flight' of laser pulses while the bearing is referenced to an internal optical encoder. The retro-reflective targets are industry standard low-cost passive 3M reflective tube or flat targets or clusters of nitrogen filled glass prisms.

This data is automatically fused with attitude information from an internal vertical reference unit which provides automatic wave motion compensation to enable the unit to remain horizontal or at a given upwards or downwards viewing angle irrespective of the vessel roll and pitch. The vertical reference sensors module contains solid-state tri-axial accelerometers and gyros as well as a gravity referenced inclinometer. These sensor signals are processed to provide a rate stabilised vertical reference, a yaw rate, a lateral acceleration and a forward acceleration. Based on the calculated system inclination and the mechanical drive ball-screw displacement the system can calculate the actual inclination of the vessel. This paper will not provide details on the vertical reference measurements and attitude calculation.

This attitude-corrected basic range and bearing data can then be used for navigation purposes such as position-only basic single target tracking, i.e. range and bearing to a single reflector. However, when two or more reflections are combined in a multiple hypothesis tracking algorithm the system can provide local position and heading information. This paper is mainly concerned about the latter tracking mode.

Basic Sensing and Target Configuration

A number (m) of retro-reflective beacons are positioned around a platform, which the vessel is approaching. Ideally the relative positions of these reflectors are known a priori but this is practically rarely possible, so instead they are automatically measured via an auto-survey, typically when the vessel approaches a platform for the first time¹.

The position of a reflector is specified as the position of the centre of the reflector and a reflector can be either a flat reflector or a cylindrical reflector. The orientation of the reflector is also recorded. In the case of a cylindrical reflector this is the orientation of the surface behind the reflector. The laser scanner primarily measures the bearing to a reflector, the range to a reflector and the angle subtended by the reflector at the scanner (other information such as brightness etc. are less relevant for the actual target tracking).

An important constraint is that the reflector configuration is required to be convex. This means that each of the reflectors face out from the platform so that it would not be possible for a laser scanner placed at the position of one of the reflectors to see any of the other reflectors. A consequence of this is that reflectors are always seen in the same order within a particular revolution of the scanner. This means that the system instantly rejects a hypothesis for which reflectors appear in the wrong order in one revolution. It can also reject a pair of associations

¹ This survey result can then be stored and used again for navigation against this platform at subsequent approaches.

where the field of view for the first reflector does not intersect the field of view of the second. Even when the fields of view do intersect, it can obtain an upper bound on the difference between the bearing to the first reflector and the bearing to the second reflector.

The use of cylindrical reflectors has special implications on the convexity assumption. It still holds if there is a barrier some way behind the reflector which blocks the view and the scanner will never be between the reflector and that barrier. This is typically the case in standard offshore rig installations.

Range and Bearing Correction for Tilt

In addition to the standard pulse-by-pulse signal calibration for receiver gain, output power, measured signal strength and temperature which are not covered in more detail here, the system has to correct for the vessel attitude.

The laser measures the bearing of the beam in the plane of the laser scanner. The position calculations require the bearing of the horizontal projection of the laser beam relative to the orientation of the sensor unit:

$$\theta_{corrected} = \arctan 2(\cos \theta_{roll} \sin \theta_{uncorrected}, \cos \theta_{pitch} \cos \theta_{uncorrected} + \sin \theta_{pitch} \sin \theta_{roll} \sin \theta_{uncorrected})$$

Equally, the laser scanner measures the range to the reflector in the plane of the laser scanner, but the position calculations require the length of the horizontal projection of the ray to the target:

$$r_{corrected} = r_{uncorrected} \sqrt{\cos^2 \theta_{roll} \sin^2 \theta_{uncorrected} + (\cos \theta_{roll} \cos \theta_{uncorrected} + \sin \theta_{pitch} \sin \theta_{roll} \sin \theta_{uncorrected})^2}$$

There are a number of other corrections that should be applied in order to accurately translate from the sensor frame to the vessel frame. For example, the sensor heading is not exactly the same as the heading of the vessel. The latter is understood to be the horizontal projection of the stern-bow axis of the vessel frame, because when both the sensor pitch and roll angles are non-zero, then the heading of the vessel in the sensor frame is also non-zero, but the details of this analysis reach beyond the scope of this paper.

Multiple Hypothesis Tracking

The basic principle behind the position and heading estimation algorithm is to use the difference of two bearing measurements and the two corresponding range measurements to calculate the position of the sensor. This equates to three constraints on two unknowns. The extra constraint is disposed of by minimizing the bearing and range prediction errors².

Each time the sensor receives a reflection, whether from a true reflector or a rogue reflection from other environmental features, it has to make an assumption about where that reflection came from. It may be a false reflection or it may come from any of the m expected reflectors³. This gives rise to a total of $(m+1)$ hypotheses after one reflection⁴. After n reflections a total of $(m+1)^n$ hypotheses arise so implausible hypotheses must be rejected as soon as possible.

² It is also possible to dispose of the extra constraint by calculating both the position of the scanner and the range measurement offset but experimental and site measurements show that the range offset in this hardware implementation is very small.

³ The operator either selects on the user console interface which reflections to use for an auto-survey or directly picks a stored defined set of reflectors.

⁴ A hypothesis is a set of associations, one for each reflection received since start up. The term association refers to the decision made about the source of a particular reflection. If a reflection has come from a bad reflector (rogue) this is referred to as the “null association” which results in the $(m+1)$ -th hypothesis.

Secondly, it is also possible for two or more hypotheses to emerge which are identical except for the decision made about a single reflection in the distant past. That decision makes very little difference to the position estimates that these hypotheses produce, so a further means of detecting that two hypotheses are very close and deleting one of them is needed.

Last, whilst the system needs to track multiple hypotheses, it must be able to identify the most plausible of the current hypotheses so that it can report a position to the DP system.

The requirements above can be satisfied by calculating a plausibility score for each hypothesis and a measure of the statistical distance between two hypotheses. The plausibility is based on the normalised squared innovations for that hypothesis. This score is always positive and a very small plausibility score indicates that a hypothesis is very good and a large score indicates that it is implausible.

The First Reflection and Forming a New Hypothesis

When the first reflection is obtained, a total of $m+1$ hypotheses can be generated. The first of these makes the null association for the first reflection which is awarded a score of ζ_{null} . For each of the other hypotheses, there is only one reflection from a known reflector without any further assessment of the plausibility⁵.

Each time a reflection is processed a new set of hypotheses is produced consisting of an old hypothesis about the sources of all previous reflections combined with an association for the current reflection. The plausibility score for a new hypothesis is calculated according to

$$\zeta_{\text{new}} = e^{\lambda_{\zeta}(t_{\text{new}} - t_{\text{old}})} \zeta_{\text{old}} + \zeta_{\text{assoc}}$$

where ζ_{assoc} is a measure of how well the new association fits in with all the previous associations that have gone to make up this hypothesis and λ_{ζ} is the plausibility score forgetting factor⁶.

Initialising a Hypothesis

Once a second non-null association has been made under a hypothesis, it is possible to make a first estimate of the pose of the system. For example, by calculating the distance between the two reflectors and the difference in range measurements, the hypothesis can be filtered for plausibility because as a consequence of the convexity assumption the magnitude of the difference in range cannot be greater than separation of the reflectors. An iterative least-squares estimation can be used to obtain the first position and heading estimate for the hypothesis.

Given a state vector $[\hat{\psi} \ \hat{x} \ \hat{y} \ \hat{\rho}]^T$ and a measurement vector $[\theta_1 \ r_1 \ \theta_2 \ r_2]^T$ an information matrix for this 4-state estimate can be constructed $Y_4 = H^T R^{-1} H$ where R is the measurement covariance matrix and

$$H = \begin{bmatrix} \frac{\partial \theta_1}{\partial \psi} & \frac{\partial \theta_1}{\partial x} & \frac{\partial \theta_1}{\partial y} & \frac{\partial \theta_1}{\partial \rho} \\ \frac{\partial r_1}{\partial \psi} & \frac{\partial r_1}{\partial x} & \frac{\partial r_1}{\partial y} & \frac{\partial r_1}{\partial \rho} \\ \frac{\partial \theta_2}{\partial \psi} & \frac{\partial \theta_2}{\partial x} & \frac{\partial \theta_2}{\partial y} & \frac{\partial \theta_2}{\partial \rho} \\ \frac{\partial r_2}{\partial \psi} & \frac{\partial r_2}{\partial x} & \frac{\partial r_2}{\partial y} & \frac{\partial r_2}{\partial \rho} \end{bmatrix} = \begin{bmatrix} -1 & \frac{dy_{i1}}{r_{i1}^2} & -\frac{dx_{i1}}{r_{i1}^2} & 0 \\ 0 & -\frac{dx_{i1}}{r_{i1}} & -\frac{dy_{i1}}{r_{i1}} & 1 \\ -1 & \frac{dy_{i2}}{r_{i2}^2} & -\frac{dx_{i2}}{r_{i2}^2} & 0 \\ 0 & -\frac{dx_{i2}}{r_{i2}} & -\frac{dy_{i2}}{r_{i2}} & 1 \end{bmatrix}$$

⁵ It is possible to specify a restricted range for the heading with respect to the world frame allowing a test for the maximum angle of incidence.

⁶ Hypotheses are assessed mostly on their performance over the past 10 seconds.

The 4-state information matrix is augmented to produce a state information matrix for a 9 state Kalman filter where the initial estimates of vehicle speed, accelerometer offsets and the gyro offset are zero. At this stage, the first non-null association has already been awarded a score of ζ_{null} , and this score has been incorporated into the score for the hypothesis. The score for the second non-null association can be computed from the measurement prediction and measurement vector by:

$$\zeta_{assoc} = \sum_{i=1}^2 (\underline{\theta}_{mi} - \hat{\theta}_i)^T (H_i P H_i^T + R)^{-1} (\underline{\theta}_{mi} - \hat{\theta}_i) + k_{\zeta\omega} \frac{(\omega_{m,i} - \hat{\omega}_i)^2}{(\omega_{max,i} - \omega_{min,i})^2}$$

where $k_{\zeta\omega}$ is a subtended angle innovation weighting factor.

Having initialised the pose estimate for a hypothesis, the system maintains and refines its estimate of the sensor unit pose by means of its 9 state Kalman filter incorporating not only the orientation and X,Y position but also forward and lateral velocities and accelerations as well as gyro and range offsets. The speed is modelled as a coloured noise process⁷.

The gyro measurement of angular speed is subject to a slow drift. By an appropriate choice of gains in the Kalman filter, it is possible to eliminate the effect of an offset in the gyro reading in less than 1 minute. From one scanner reading to the next, the rate gyro reading is integrated to produce a measurement of the change in the heading of the vessel. Likewise the accelerometer readings are integrated to produce the change in speed. Each time a scanner reading is obtained, the state and covariance are extrapolated. The state transition matrix relies on a speed estimator forgetting factor which must be chosen to be the reciprocal of the time over which the speed of the vessel changes.

Correcting the Estimate

When a new association is made under a hypothesis, the state estimate and the state covariance must be corrected. From the observation equations a linearised observation matrix can be derived

$$H = \begin{bmatrix} -1 & \frac{y_T - \hat{y}_v}{\tilde{r}^2} & -\frac{x_T - \hat{x}_v}{\tilde{r}^2} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -\frac{x_T - \hat{x}_v}{\tilde{r}} & -\frac{y_T - \hat{y}_v}{\tilde{r}} & 1 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

where \tilde{r} is the estimated range from the scanner to the particular target. Given the measurement prediction error covariance $\Sigma = H P H^T + R$ and Kalman gain $K = P H^T \Sigma^{-1}$ the state correction is

$$\hat{x}(+) = \hat{x}(-) + K (\underline{\theta}_m - \hat{\theta})$$

Crucially, at this stage a number of tests such as a range test, bearing test, incidence angle test and subtended angle test can be applied to reject implausible associations. Finally a score for the new association can be calculated

$$\zeta_{assoc} = (\underline{\theta}_m - \hat{\theta})^T (H P H^T + R)^{-1} (\underline{\theta}_m - \hat{\theta}) + k_{\zeta\omega} \frac{(\omega_m - \hat{\omega})^2}{(\omega_{max} - \omega_{min})^2}$$

The new hypothesis is immediately rejected if $\zeta_{assoc} > \zeta_{max}$ otherwise the score for the hypothesis is given by

$$\zeta_{new} = \exp(\lambda_{\zeta} (t_{new} - t_{old})) \zeta_{old} + \zeta_{assoc}$$

If this score is very much greater than the current best score then the new hypothesis is rejected

⁷ See chapter 4.5, "Tracking and Data Association.", Bar-Shalom and Fortmann

by $\zeta_{new} > \zeta_{best} + \zeta_{trim}$ where ζ_{trim} is a trimming margin.

Trimming Hypotheses and the Null Hypothesis

As described above, on each rotor revolution the system extrapolates the state and covariance for each of its hypotheses and deletes all uninitialised ones. The remaining hypotheses are sorted by ascending score and starting from the tail all hypotheses which are very much less plausible than the best hypothesis are deleted, i.e. where the score is at least ζ_{trim} greater than the best score.

Hypotheses which are very similar, i.e. non-distinct, can also be deleted, because if all the associations made under one hypothesis during the rev just completed are the same as the associations made under the other hypothesis, then the one with the larger score can be removed. This is an efficient means of keeping the number of hypotheses tracked at any point in time to a manageable group of only the most likely solution candidates.

Finally, a null hypothesis needs to be added to the solution pool which ensures that there is a mechanism to reinitialize a new navigation track if the current solutions deteriorate. This hypothesis is awarded a score a little greater than that for the best hypothesis in the list, i.e.

$$\zeta_{new} = \min\{\zeta_i\} + \zeta_{pen}$$

where ζ_{pen} is the null hypothesis penalty. In order for the system to continue to hold on to a good hypothesis if it gets a couple of bad readings in a row, ζ_{pen} should be chosen to be $> \zeta_{null}$.

DP Position Calculation

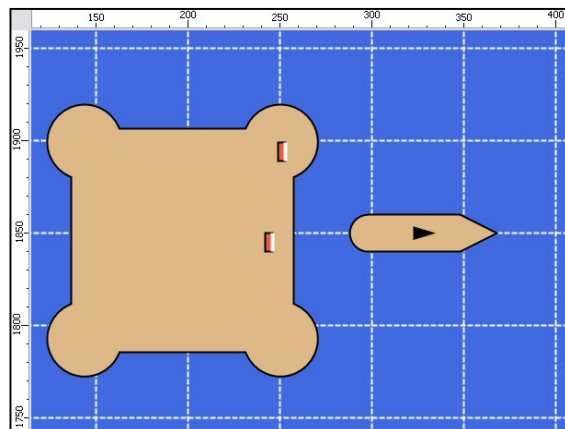
For each hypothesis the position and heading of the sensor unit is calculated based on the Kalman filter output of all the associated reflections since the hypothesis started up. Typically, the sensor unit supplies position and heading values to the DP system which then uses these as pseudo observations in its own Kalman filter. In these circumstances it is not appropriate to supply a Kalman filter output for use as an input to the DP system's Kalman filter. Instead, on each rev, a position and heading based entirely on measurements made during that rev are supplied.

To do this, the best hypothesis (with the lowest score) is selected from which two associations are selected which maximize the distance between the reflector positions. These two associations can then be used to obtain a position and heading for the CyScan unit using the same algorithm as the hypothesis initialization while using the current state estimate for the chosen hypothesis as the starting solution. Alternatively, all the associated observations in a rev can be used to minimize the measurement prediction errors over an arbitrary number of measurements.

Results

Typical DP operations for platform supply vessels involve reversing stern towards a rig which might have two or more laser targets installed as shown to the right and then holding the position while load transfers are completed.

The following screen shots and graphs show the inherent robustness of the multiple hypothesis tracking approach against persistent rogue reflections without resorting to naive range or brightness gating. In this example the system is

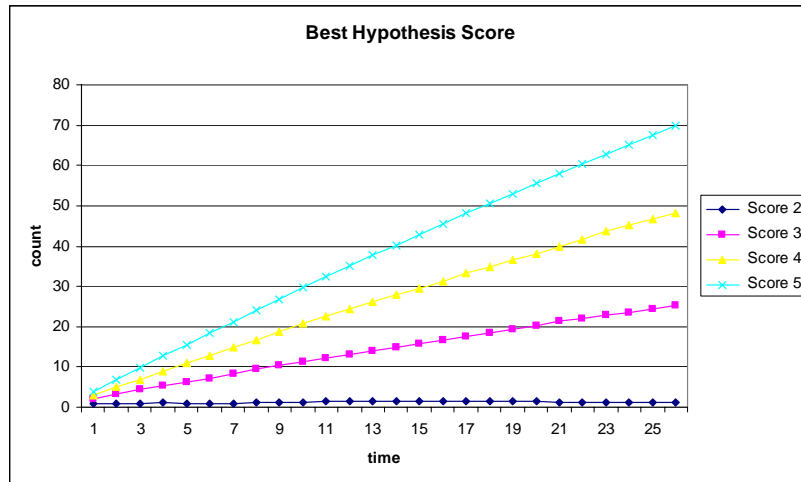


holding position at ca. 80m from the closest target and maintains this position in the presence of other persistent rogue reflections without degradation of position and heading fix.



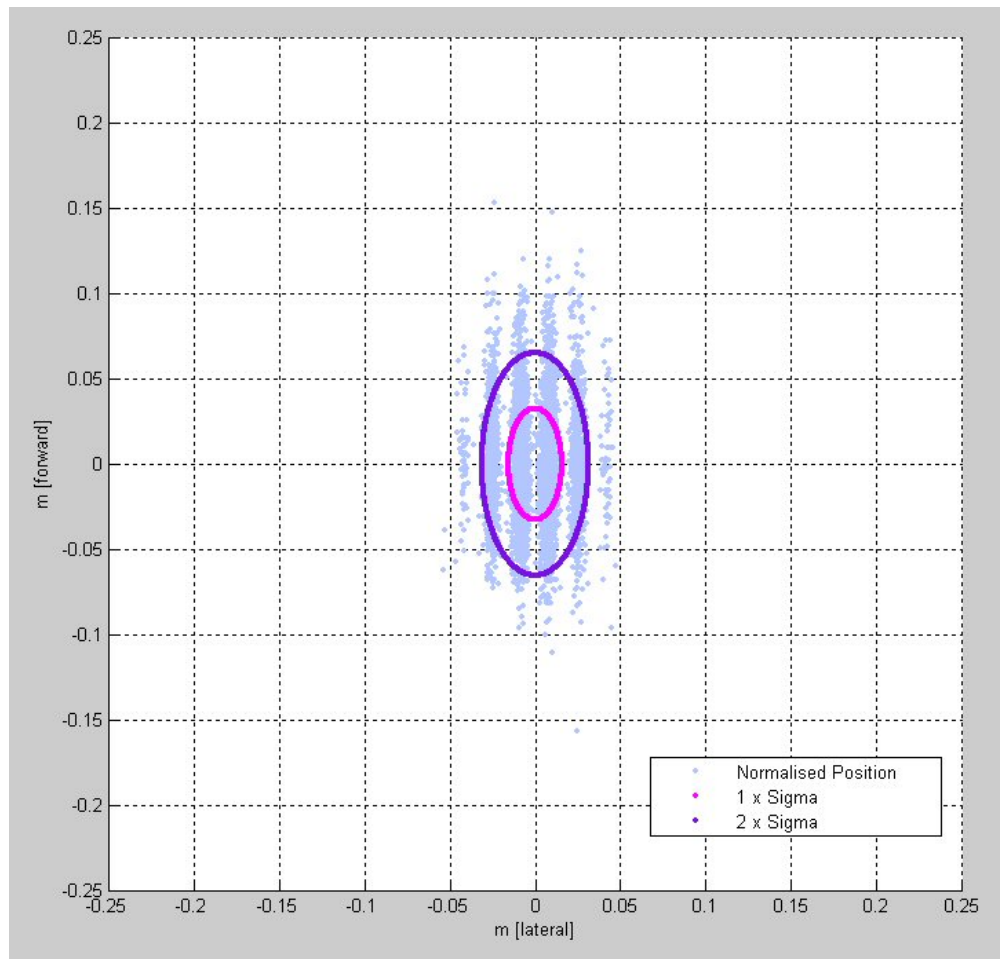
This example is particularly strong as one of the persistent rogue reflections forms a very unfortunate ambiguous non-convex constellation which could theoretically throw the system into confusion.

However, the multi-hypothesis algorithm with its careful scoring of the possible solutions in the solution space ensures that the correct, i.e. best fitting, solution is presented as the position fix to the DP system. It is important to highlight that the rogue reflections are actually generated not from weak clutter but by real targets generating equally strong reflections in both brightness and subtended angle. Simple gating on these parameters would not allow the correct solution to be determined. The presence of these false reflections results in a significant increase of the hypothesis score as shown to the right where in the 2 target case the score remains well in the single digit realm, but in the 5 target case quickly approaches a 3 digit value.



This inherent robustness in multi-target tracking ensures that coupled with the raw measurement accuracy of the laser head itself a very high degree of positional accuracy and repeatability can be achieved in the most demanding circumstances.

A typical normalised position accuracy plot for a PSV navigating at a few hundred meters target range with a 2σ value less than $\pm 7\text{cm}$ is shown below. The grating effect is due to the different measurement characteristics in range and bearing which become cross-correlated when translated into a Cartesian X,Y position. The data has been aligned so that the bearing measurement corresponds to the lateral direction and the range measurement to the forward direction.



Appendix: Short Specification

Positional accuracy:	better than 0.5% of range depending on target layout
Range resolution:	10 cm
Range repeatability:	20 cm
Angular resolution:	0.006° (0.1 mrad)
Angular repeatability:	0.03° (0.5 mrad)
Scan rate:	uni-directional $360^\circ/\text{s}$ (6.28rad/s)
Beam characteristics:	16° vertical $\pm 20^\circ$ mechanical tilt, 0.23° horizontal
Tilt compensation:	$\pm 20^\circ$ roll and pitch combined, $5^\circ/\text{s}$ max speed $\pm 5^\circ$ for 15s wave period $\pm 2.5^\circ$ for 5 second wave period