Accurate and automatic surveying of beacon positions for a laser guided vehicle

Magnus Oskarsson, Kalle Åström
Dept of Mathematics, Lund University
Box 118, S-221 00 Lund, Sweden
E-mail: {magnuso,kalle}@maths.lth.se
phn: +46 46 222 41 45, +46 46 222 45 48
fax: +46 46 222 40 10

Abstract

In this paper we outline a system for automatic surveying of beacon positions for an autonomous guided vehicle. The vehicle measures the bearings to retroreflective beacons using a laser scanner. As the vehicle moves, asynchronous measurements to unidentified beacons are obtained. The goal of the system is to automatically group those measurements that belong to the same beacons, to eliminate all erroneous measurements and to estimate the two-dimensional map of the beacon positions. Grouping of the beacons is achieved in a RANSAC type of algorithm. The map of the beacons is estimated using maximum likelihood estimation using bearing measurements and odometry. Recursive approximations of the estimation problem is used for increased speed.

1. Introduction

Autonomous guided vehicles are important in factory automation. One way, and the way that has been used traditionally, to guide the vehicle is by wires in the ground. This works, but is very inflexible. To change the path of the vehicle one has to rewire the wires. Another way that is much more flexible is to use a vehicle that, from the information where it is, can calculate where it is supposed to move next. An example of such a vehicle is a laser guided vehicle, see [1, 4, 6]. This vehicle is equipped with a laser that rotates and scans a horizontal plane of the premises. The surroundings are equipped with beacons in the form of reflective tapes. These give rise to reflections when hit by the laser. The bearing to the beacon relative to the heading of the vehicle is measured by a sensor that is mounted on the laser.

Introduce an object coordinate system which will be held fixed. The bearing α defined in fig 1, depends on the position of the beacon (U_x, U_y) and of the position (P_x, P_y) and orientation P_θ of the scanner, cf. figure 1, according to

$$\alpha(P,U) = \operatorname{atan2}(U_y - P_y, U_x - P_x) - P_\theta \ , \tag{1}$$

where atan2 is the four quadrant inverse tangent as explained in [5]. The vector (P_x, P_y, P_θ) is called the **meter state**. The actual measurement $\bar{\alpha}$ deviates from the true angle, due to noise and quantization effects in the laser scanner:

$$\bar{\alpha} = \alpha(P,U) + \Delta\alpha \ . \tag{2}$$

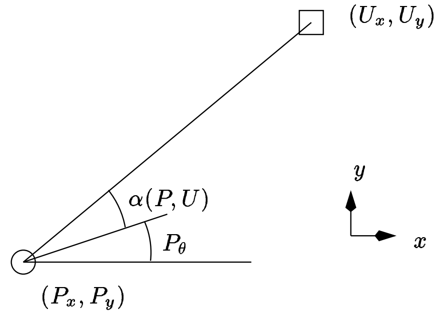

Figure 1. The figure illustrates the measured angle α as a function of scanner position (P_x, P_y), scanner orientation P_θ and beacon position (U_x, U_y).

The position and the heading of the vehicle is calculated using these bearing measurements and the positions of the beacons, [6]. As the positions of the beacons have to be known, one has to develop a method by which you can calculate the map of beacons the first time that a new locale is encountered. Previously this was done in a semi-automatic way, see [2, 3]. The association of synchronous bearings was done manually, and then optimized by a computer, or by hand using surveying instruments. To do the association manually is not completely satisfactory, since it is both time-consuming and difficult, and therefore prone to errors. Traditional surveying using theodolites is very time consuming. It may take a whole day to create the map of a relatively large locale. This is not very convenient and it is also very expensive. It would be much better to have a completely automatic procedure, using the autonomous vehicle itself with the laser scanner. This is the problem that is discussed in this paper.

The time factor is not that important, since the map of beacons is only constructed one or a few times for a new area. The most important thing is the precision of the map, that all beacon positions are accurate. Another important factor to be considered with a fully automatic procedure, is robustness.

The vehicle measures, apart from the bearings where reflections occur, information about the vehicle's speed and the direction of the steering wheel. This information is used to calculate the path that the vehicle has moved in.

The resolution of the scanner is about 1 mrad. Apart from random errors in the measurements of reflections there may be systematic errors due to misalignment of the laser and the mirror. Errors in the odometry may depend on loss of traction of the wheels of the vehicle, which causes the speed to be measured wrongly.

2. General framework

As the scanner moves around time stamped bearings are obtained. These bearings occur at asynchronous times. We have no a priori information at which times we will get a reflection. This means that if we only have information about the direction of a reflection and no information about distance to beacons and no information about how the vehicle has moved, the problem of reconstructing the map of beacons is impossible to solve. So some information about the movement of the vehicle is necessary. On the other hand if we had full and accurate knowledge of the vehicle's movement it would be an easy task to reconstruct the map of beacons in an accurate way. As we have access to both odometry data and bearing measurements it would be desirable to use the combination of information in an optimal way. The odometry is most accurate on a small time scale; errors in odometry tend to accumulate. If we were to rely on the odometry on a large time scale we would soon have a very poor estimate of the vehicle's position relative to the origin. The solution is to use the odometry on a small time scale and then use the bearings to connect beacons that are the same.

The entire data log is divided into smaller segments. We have used two second long segments. The path of the vehicle within the segment is then calculated using the odometry data. From this path and the reflections, the beacons are identified and incorporated into the map. Old beacons in the map are first tried to be associated with bearings in the active segment. After this step the remaining bearings are grouped and then associated with beacons that are incorporated into the map of beacons. The last step is an optimization, which is done over all the acquired data up to this point in time. We use both a full optimization and a recursive optimization. The recursive optimization is done when we believe that we are relatively close to a minimum of the values of our parameters. The whole procedure is then repeated with the next segment from the data log.

3. Odometry

The vehicle used has three wheels, one front steering wheel, and two wheels on a back axis. It is assumed that the vehicle can only move in a direction normal to the back axis, i e the back wheels do not move side ways.

The vehicle measures both its speed and its direction. The direction is registered as the angle of the steer wheel of the vehicle. From this information the path that the vehicle has moved in from an initial position can be calculated. This is simply done as adding small distances from the initial position.

$$x = x_0 + \sum_i v_i \Delta t_i \tag{3}$$

This is how an origo on the vehicle has moved. The important information is how the position of the angle meter has moved. To calculate this position one has to know its position relative to the origo of the vehicle. Apart from this information the wheel base of the vehicle, i e the distance between the steering wheel and the back axis, is known. The position of the angle meter is fixed on the vehicle, and hence the position of the angle meter moves in the same way as the origo. It is not enough to know how the position of the angle meter has moved, it is also necessary to know how the angle meter has rotated in the coordinate system. This is done by calculating the angular velocity of the vehicle, Ω. From this the angle of the angle meter is calculated.

$$\theta = \theta_0 + \sum_i \Omega_i \Delta t_i \tag{4}$$

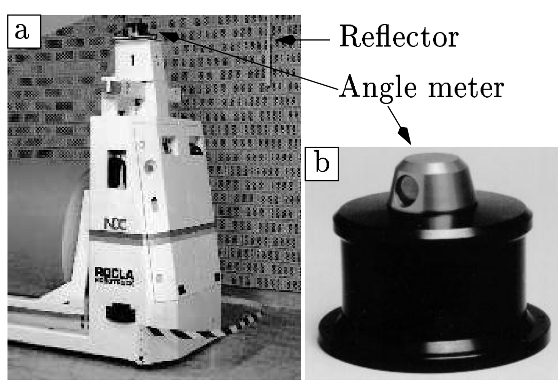

Figure 2. a: A laser guided vehicle. b: A laser scanner or angle meter.

4. Automatic association of measurements

Before the reflections are grouped, the odometry for the current segment of data is calculated. The grouping of bearing measurements is done by the following method. First a bearing measurement is randomly chosen. From this angle one can predict projected angles for other times in the segment. If there are bearing measurements that correspond to this angle within some degree of accuracy these bearings are grouped. To check that these bearings correspond to an actual beacon two things have to be checked. First there has to be a minimum number of associated measurements, e g 4. Secondly the root mean square of the difference between projected angles and measured angles must not be too large. If this is OK then a new reflector is created in the map. If it's not OK the reflector position has to be established more accurately. This is done by choosing two bearings from the previously investigated. From these two measurements a new reconstruction is done. This is tried a few times or until a reflector position is established sufficiently correct. The whole procedure is repeated until all the bearings in the segment have been associated or tried to be associated.

When a map of beacons has been established one can first of all, in a new segment, try to associate bearings with existing reflectors in the map. This is done by calculating the difference between the measured bearing and the projected bearing. The projected bearing is calculating using the odometry data and the earlier established map of beacons. Not all the beacons in the map are tried; only those that are active, i e the beacons that were used in the data segment prior to the current one.

5. Optimizing the beacon map

After a segment has been associated, the whole map of beacons is optimized and updated. The optimization is done over all the old segments and all the beacons in a nonlinear least-squares sense. All the relative positions of the vehicle in each segment are fixated so that only stiff body motion of an entire segment is allowed. Then the sum of the squares of the differences between the projected angles a_p and the measured angles a_m divided by the standard deviation σ in estimated angles, for all bearing measurements, is calculated,

$$f = \sum_i \left(\frac{a_p(i) - a_m(i)}{\sigma} \right)^2 \tag{5}$$

This sum is minimized, with respect to the positions of the beacons in the map and the positions of the segments. Every beacon gives rise to two variables, the x and y position, and every segment gives rise to three variables, x and y position of the midpoint of the segment, and the orientation of the segment. In total there are n*2+m*3 variables, where n is the number of beacons and m is the number of measurement segments. The size of the residual vector is the same as the total number of bearing measurements that are associated. This grows with about 140 elements with every segment that is investigated. This means that the optimization becomes significantly slower the more segments that have been associated. In the trial runs that have been investigated, there have been about ten thousand bearing measurements in every run. This means that the vector that is to be minimized is quite large.

6. Recursive optimization of beacon map

As the number of segments increase, the number of measurements becomes very large. Each step of the optimization as described above takes longer and longer time. It would be preferable to use more of the information of the old map since it doesn't change that much after awhile.

Let x denote linearized deviations of the unknown parameters (beacon map, vehicle motion etc) from the minimum. Let Y be the vector of weighted residuals $Y(i) = \frac{a_p(i) - a_m(i)}{\sigma}$. The derivative $A = \frac{\partial Y}{\partial x}$ gives valuable information about the quality of the estimate. The covariance matrix of the estimated parameters is $C[x] = (A^T A)^{-1}$ and the function f can be estimated as

$$f \approx f_0 + x^t (A^T A) x \tag{6}$$

close to the minima.

Let M be the number of measurements and N the number of estimated parameters. Using the cholesky factorization it is possible to write $A^T A$ as $L^T L$, where L is an upper triangular matrix of size $N \times N$ as opposed to A which is of size $M \times N$. Since the number of measurements M is often much larger than N, L is much smaller than A.

So it would make sense to use the following approximation

$$f_{old} \approx f_0 + x^t (L^T L) x = f_0 + (Lx)^T (Lx) \tag{7}$$

instead of all the old measurements.

Optimizing the map using both new and old measurements

$$\min f = f_{old} + f_{new} = (Y_{old})^T (Y_{old}) + (Y_{new})^T (Y_{new})$$

is equivalent to minimizing

$$(Lx)^T (Lx) + (Y_{new})^T (Y_{new}) \ ,$$

i e minimizing the square of

$$Z = \begin{bmatrix} Lx \\ Y_{new} \end{bmatrix} \tag{8}$$

The derivative of Z with respect to the motion parameters is

$$A = \begin{bmatrix} L & 0 \\ \frac{\partial Y_{new}}{\partial x} & \frac{\partial Y_{new}}{\partial y} \end{bmatrix} \tag{9}$$

Motion parameters can be updated by solving

$$A \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = - \begin{bmatrix} Lx \\ Y_{new} \end{bmatrix}$$

in a least-squares sense. The covariance matrix of estimated motion parameters x and y can now be estimated as

$$C[x, y]^{-1} = A^T A \ .$$

This method of updating the motion parameters is substantially faster since only the new residuals and their derivatives have to be calculated in each step. As an example in a run after 80 seconds and 14 beacons in the map the size of the approximation Lx is 28×1 where as the exact Y_{old} has size 5146×1. The size of Y_{new} which is used in both the recursive and the full optimization is 139×1. So the advantage of the recursive optimization over the full optimization is very large in terms of computing time. When using the recursive optimization it should be possible to run the program in real-time. This is not possible if a full optimization has to be done after every segment. The program uses the full optimization in the beginning of the run and changes to the recursive optimization after a few segments. If the errors become too large a full optimization is done. This does not happen too frequently, once every four or five minutes, depending on how many disturbances that are present in the bearing log.

The information in the old measurements is encoded in L.

7. Optimizing vehicle parameters

As described in section 3 the wheel base of the vehicle is a parameter that plays a role in the calculation of motion of the vehicle. Apart from this parameter the position of the laser scanner relative to the vehicle has also to be known, as well as the zero angle of the angle-meter relative to the vehicle.

These four parameters, the wheel base of the vehicle the, x and y position of the laser scanner and the angle of the scanner, are measured when the laser scanner is mounted on the vehicle. Especially the position of the scanner can change slightly during transportation, and of course it need not be necessary that they are correctly measured. In this case it would be preferable to be able optimize these parameters so that the errors in the map are minimized. This is done in a non-linear least squares sense. The optimization of the vehicle parameters is not done automatically at the moment but can be used in the post-processing of the beacon map to increase the accuracy of the map.

8. Robustness

Since this is to be used in an industrial application it is important to have a robust system. It is desirable, but not entirely possible, to have a system with a go button that can handle every thing that can possibly happen. There are several things that can go wrong, ranging from things that are easily remedied to things that are impossible to anticipate or handle.

One thing that happens is the registration of false reflections. False in the sense that they arise from objects that are not to be considered as beacons in the environment. These may arise from metal parts and windows. The scanner registers with every reflex a status of how good the reflex is. This is used to help distinguish real from false reflections. Only reflections that have status over a specific level, that can be set, are used in the association procedure.

A thing that frequently happens when the map is reconstructed, is that instead of associating bearings that should be associated with an old beacon in the map, a new beacon in the map is created. This means that there will exist two or more beacons in the reconstructed map that in reality correspond to one actual beacon. This is undone by an algorithm that merges beacons. The routine uses a measurement :

$$f = (q - p)^T C^{-1} (q - p) \tag{10}$$

to determine whether two beacons are the same. In this equation p denotes the coordinates of the current map and q is the map under the hypothesis that two beacons in the map are really the same. C is the covariance matrix, so f measures the difference between the current map and the hypothesis of merging two beacons weighted with the covariances for the different beacons. If f is small there is a good chance that the two beacons in question are the same. In order to merge the beacons a variety of things are checked, so that there are no conflicts in merging the two beacons, and the beacons can't be too badly estimated since two very badly estimated beacons could then be merged without being near each other.

Although a lot of things have been done to increase the robustness of the system, there are still many problems that must be addressed. One big problem is to have a robust association. Things that may occur that have to be dealt with are crossing tracks in the bearing logs. These crossings have to detected to ensure that false associations don't occur. It is better to not associate bearings to an old beacon in dubious cases, since we have the ability to merge beacons. It is harder to remedy the case when bearings are falsely associated to a beacon in the map.

9. Experimental Validation

The procedure seems to work adequately in most cases. Although it has only been tested on a moderate amount of data logs it seems that it finds all beacons and positions them correctly.

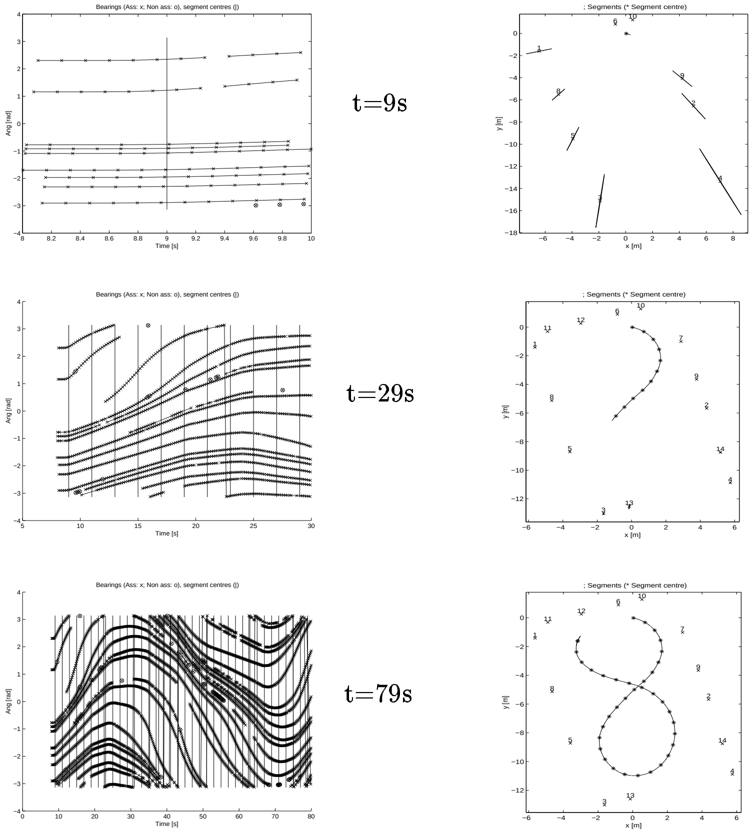

a b

Figure 3. a) associated bearings after respectively 9, 29 and 79 seconds and b) reconstructed maps from the above bearings after respectively 9, 29 and 79 seconds

In fig 3b one can see how the map of beacons is built as time progresses. The reconstructed path of the vehicle is also shown. In fig 3a the corresponding time stamped bearing measurements are shown. Associated bearings are connected with lines, and one can see that most of the bearings are associated. In the plot of the beacon-map the covariance ellipses around every beacon are also plotted. The size of the ellipses is 5 standard deviations. As can be seen the ellipses shrink quite fast and in the last reconstructed map they are hardly visible. They have in the end shrunk to a few centimeters. A further indication of the accuracy is seen in fig 4. Here the result of several runs in the same locale are shown in the same plot. For every run a map of the beacons is reconstructed and these maps are then fitted so that the distances in the map are minimized. As can be seen they match each other closely. In the close-up in fig 4 one can see that the difference between the different maps is a few millimeters. It's the same for all the beacons in the maps that have been examined.

The residuals between measured angles and the projected angles are in the final map of the size mrads, which is the same magnitude as the resolution of the laser scanner. These errors in residuals and the implied errors in the maps mentioned above are no larger than those that are used in today's maps, that are measured manually. These accuracies are adequate for safe navigation in the places where the vehicles are used. There are still errors in the residuals between the measured and the projected angles that seem to be of a systematic nature. We have not been able to determine where these errors occur or whether they indeed are systematic.

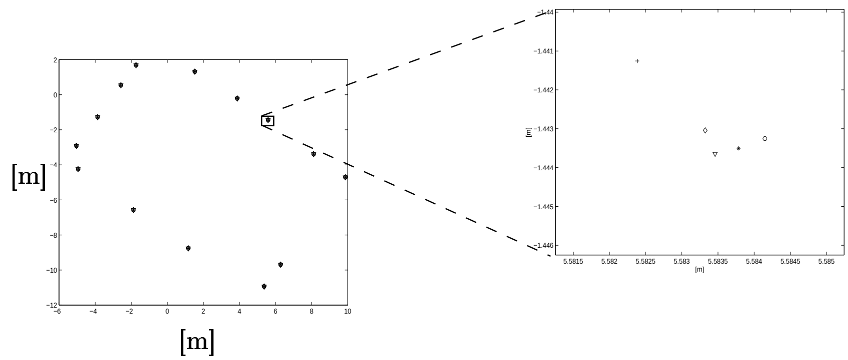

Figure 4. Five maps super-imposed on each other

10. Conclusions and Future Research

The procedure works almost automatically. The errors in the final maps are of the order of millimeters and the residuals between the measured and projected angles are of the order mrads, which is the same as the resolution of the laser scanner. There are still some systematic errors that we do not have an explanation for.

The routine works fast when the recursive optimization is used, and this seems to give quite good results. When the errors become too large a full optimization is done, which takes some time. The full optimizations though, are done with several minutes apart.

The hardest problem at the moment is to get a robust routine. Several measures have been taken to ensure a safe and automatic process, including the merging of beacons, and the merging of maps. There are still many problems in getting a robust association. Problems occur in the association when track crosses in the bearing logs.

Acknowledgments

The authors thank Netzler & Dahlgren Co AB.

References

[1] U. Andersson. Trajectory estimation and control of autonomous guided vehicles. Licentiate thesis, Luleå University of Technology, Luleå, Sweden, 1989.

[2] K. Åström. Where am I and what am I seeing? Algorithms for a laser guided vehicle. Master's thesis, Dept. of Mathematics, Lund Institute of Technology, Sweden, 1991.

[3] K. Åström. *Invariancy Methods for Points, Curves and Surfaces in Computational Vision.* PhD thesis, 1996.

[4] K. Hyyppä. Optical navigation system using passive identical beacons. In *Proceedings intelligent autonomous systems, Amsterdam*, 1987.

[5] The Math Works Inc. *MATLAB Reference Guide.* The Math Works, Inc., 1992.

[6] U. Wiklund. Algorithms for navigation of autonomous guided vehicles. Licentiate thesis, Luleå University of Technology, Luleå, Sweden, 1988.