
A System for Automatic Marking of Floors in Very Large Spaces

Patric Jensfelt¹, Gunnar Gullstrand¹, and Erik Förell²

¹ Centre for Autonomous Systems, Royal Institute of Technology, SE-100 44 Stockholm, patric@nada.kth.se, gunnar@gullstrand.nu

² Stockholm International Fairs, SE-125 80 Stockholm, Erik.Forell@stofair.se

Summary. This paper describes a system for automatic marking of floors. Such systems can be used for example when marking the positions of stands for a trade fair or exhibition. Achieving a high enough accuracy in such an environment, characterized by very large open spaces, is a major challenge. Environmental features will be much further away than in most other indoor applications and even many outdoor applications.

A SICK LMS 291 laser scanner is used for localization purposes. Experiments show that many of the problems that are typically associated with the large beam width of ultra sonic sensors in normal indoor environments manifest themselves here for the laser because of the long range.

The system that is presented has been in operation for almost two years to date and has been used for every exhibition in the three main exhibition halls at the Stockholm International Fair since then. The system has speeded up the marking process significantly. For example, what used to be a job for two men over eight hours now takes one robot monitored by one man four hours to complete.

Key words: floor marking, localization, large space, long range, laser

1 Background

The Stockholm International Fairs (StoFair) [?] is the leading exhibition, fair and congress organizer in Scandinavia. It has 56,500 m² of indoor exhibition area. There are about 70 fairs and over 1000 congresses, conferences and seminars per year. The StoFair uses a completely free layout for the exhibitions, that is, each exhibition can have its own unique layout and there are no restrictions on the shape of the individual stands.

Shortly before each event, the production phase is initiated by marking on the floor where each stand is located. Then the stands are built, the customers move in, the fair runs, the customer move out and the stands are torn down. The next marking takes places and the cycle continues. To maximize the

utility of the space, the time between fairs should be as small as possible. The marking of the stand positions are therefore often performed during the night or at other odd hours.

Traditionally the marking has been carried out manually with tape and a tape measure. It is a very tedious and boring job. For each tape that is placed the person has to bend down all the way to the floor. A large exhibition can have several hundred stands and for each stand several coordinates are marked. This in combination with the odd hours motivates automation of the process.

The automation of the marking process can be realized in several ways and with different levels of autonomy. Most of the time is spent on finding the location to mark and not the marking itself. Clearly, a fully autonomous system provides the largest potential gain. One person could then supervise one or more marking robots and do the job faster than before. It would also relieve the person from the hard labor that the manual process involves.

2 Introduction

To realize an autonomous marking system there are several subproblems that need to be addressed. The robot need to be able to localize accurately, navigate safely through the environment and be able to mark locations on the floor. There are no systems available off the shelf yet for this task to the best knowledge of the authors.

The problem of localization has been thoroughly studied in the robotics literature, see for example [?, ?, ?]. Most of the localization research have been carried out in indoor environments. The main sensor modalities have been the ultra sonic sensors and now lately the laser scanner. Both of these provide range information. In outdoor applications other types of sensors such as GPS and inertial sensors are also common (see e.g. [?]). GPS however, is not available indoor. Different representations have been used where the two main directions are to use features [?, ?] and occupancy grids [?, ?].

Navigation and obstacle avoidance have also attracted a lot of attention over the years. In most indoor environments this is a key component as the distance to obstacle at all times is relatively small. Several methods have been proposed such as the Vector Field Histogram [?], the Dynamic Window Approach [?] and the Nearness Diagram [?]. Also here the sonar and laser scanners have been the most commonly used sensors. In the current application obstacle avoidance is easier then under regular indoor conditions as the environment is very large with few obstacles.

2.1 Outline

The rest of the paper is outlined as follows. Sections ?? lists some of the requirements on the system and presents the overall design. Section ?? discusses

the implementation except for the localization part which is described in Section ???. The results of an experimental evaluation of the system are given in Section ??? and a summary and some conclusions can be found in Section ???.

3 Requirements and Design

3.1 Requirements

This section lists some of the requirements that was put on the system by StoFair.

- The environment is constantly undergoing changes. A system that that relies entirely on artificial landmarks, such as retro reflective tapes, being placed throughout the environment would be costly to maintain. It is therefore desirable that the system uses the natural environment as much as possible.
- Since the marks on the floor are the basis for building the walls and placing the carpets they need to have a certain level of accuracy. The StoFair specified 3 cm as an acceptable level.
- It must be easy to maintain the map so that it can be adapted to changes in the environment.
- The system must avoid collisions if there are objects in the way and instead report points as unreachable if blocked.
- The marking is sometimes done several weeks ahead of time. The floor is cleaned with machines between the fairs and each mark that is lost in this process has to be re-measured. This means that the markings have to withstand significant wear for quite some time.
- The system must have the means to notify an operator if there is a problem, such as the batteries being low, the robot is stuck, etc.
- The robot must be able to report what coordinates where marked and which failed. Based on such a report the system should also be able to continue a mission if it was interrupted.
- It would be desirable to add information besides the location of coordinates to the markings.

3.2 Design Decisions

It is clear that the sensor used for localization must have a long range since the environment is very large and the features are sparsely spaced. With GPS being ruled out, radio localization systems not yet being accurate enough, sonar sensor not having the necessary range, the remaining candidate sensors are cameras and laser scanners. Vision was also ruled out since the lighting conditions are often quite bad and the uncertainty was too large whether or not the necessary accuracy could be reached. The choice was therefore to use a laser scanner.

As mentioned in Section ?? there are several ways to represent the environment in a map. Depending on the application one will be better suited than the other. Here, the requirement that a user should be able to maintain the map tipped the scale in favor of the feature based representations. Another reason for this was that a CAD model of the environment already existed which would be easier to exploit in a feature map setting.

To get a head start in the project it was decided to use a commercially available platform. Building a custom made platform would take too much time and cost too much money in the initial phase of the project.

A four wheel platform would have the advantage of handling uneven floors well and be able to traverse small obstacles such pieces of wood that had been left on the floor. The downside was that it would have significantly worse odometric performance and might not even be able to turn on a carpet for example. A conventional two wheeled, differential drive system was therefore chosen.

4 Implementation

4.1 Platform

The platform of choice for the unit was a Pioneer2-DXE from ActivMedia [?] with an on-board 800MHz computer running Linux. Figure ?? shows the platform equipped with everything that is needed to carry out a marking mission. The standard batteries of the Pioneer robots were later replaced by packs of Ni-MH cells to boost the autonomy from about 2h up to 6h. The laser scanner is the main sensor for localization and navigation. In addition, sonar and bumpers are used as a compliment for close range obstacle avoidance.

Regarding the laser scanner a more thorough investigation should have been made. Now the choice was biased by already being familiar with the SICK family of laser scanners. A SICK LMS 291 was chosen, as it has better range than the indoor SICK LMS 200.

4.2 Map

A map of the environment already existed in the CAD system that is used for planning the layout of the fairs and many other tasks. The CAD system is already familiar to the staff and offers an easy to use interface for editing a map. It was therefore decided that the map format for the robot should be CAD compatible. The DXF format was chosen because it is text based and easy to parse. Changes to the environment made in the CAD system can thus be carried over to the robot effortlessly. Figure ?? shows the features available to the robot for localization in HallC at StoFair. This hall is about 80m by 270m. Looking at the figure it seems to be dense with features, but pillars in the middle are 25m apart.




Fig. 1. The platform base is a Pioneer-2DXE. A SICK LMS 291 is the main sensor. For the marking it has been equipped with a industrial ink-jet printing system consisting of a printer head, an ink bottle, hoses and a printer control unit.




Fig. 2. The wall and pillar features in HallC which is about 80m by 270m large.

4.3 Marking Device

Clearly the system would not be functional without any means to mark locations on the floor. The requirement that the markings must be sturdy enough to withstand running a cleaning machine over them and to allow addition information to be added to the marking led us to use an industrial ink-jet printer from EBS [?]. Such systems are typically used to print on for example cardboard boxes. In its normal application the printer head is fixed and the boxes move past the head on a conveyer belt. Here the printer head is moved by the robot over the floor to create the equivalent of a gigantic plotter. This has given the first platform its name, Harry Plotter.

4.4 Navigation

The navigation system have three main parts that are responsible for i) local planning to make sure that obstacles are avoided, ii) marking of points and iii) global path planning to select which order to traverse the coordinates. The navigation system operates at 20Hz, driven by the rate of the odometric data.

Obstacle Avoidance

Obstacle avoidance in exhibition hall environments is often not as challenging as in a typically indoor environment as there is so much space to maneuver on. However, the system still needs to be able to handle situations with heavy clutter not to get stuck when an area has not been properly cleaned. The system switches between two modes depending on the amount of clutter. In the normal mode a controller similar to the one used in the High Safety situations of the Nearness Diagram [?] is used. In situations with heavy clutter the Global Dynamic Window Approach [?] is used. The speed had to be limited to 0.6m as the platform was unable to maintain a straight course at higher speed with all the extra equipment that added to the weight. Clearly, a custom made platform could have performed better here.

Marking

The printing on the floor is done at constant speed (0.15m/s). The printer is “loaded” with a text string and the printing is started by a signal that is normally generated when the object to print on passes a photodetector. Here this signal is generated by the robot.

Besides marking the location of a certain coordinate the system can add extra information that makes it easier to identify the coordinate and make it easier to see how the carpets or walls should be placed. This extra information is given by two optional text strings.

At a speed of 0.15m/s and a sampling rate of 20Hz the error just from the discrete time control could be in the order of 1cm. In addition, there are no real-time guarantees in standard Linux which can also contribute to errors. Another source of errors is the delay between when the signal is given and the printing starts. This delay is measured and compensated for. To reduce the influence of the first type errors the distance to the mark position is calculated in each step. Instead of signalling to print when the mark has been passed, printing starts when it is predicted to be shorter to the mark location in the current iteration than it will be when scheduled the next time.

Global Path Planning

The global path planning here is similar to the traveling salesman problem, each coordinate has to be visited to mark it. Currently a simplistic approach

is used where the stands are completed sequentially. The advantage of this approach, besides being computationally attractive, is that construction of the stands in principal could begin immediately after the first stand has been marked. Planning is performed once before each mission and results in an XML-file with the information needed for each coordinate.

The direction of the text information is defined by the direction of the line that connects the current coordinate with the previous. In cases where there is no previous point, special coordinates which are not marked can be added to the file.

The mission specification file in XML-format can be used directly as a report file. A tag for each coordinate tells if it has been marked or not. A comment regarding the cause of failures can also be added and easily be picked out because of the XML-format.

The robot is unable to mark a certain coordinate in mainly two cases i) the space required by the robot to mark is occupied or ii) it is not sure enough about its position.

4.5 User Interface

The user interface is graphical and has been written in JAVA. It allows the operator to monitor the progress of the robot continuously. The planned path of the robot is displayed and coordinates that have been successfully marked are faded out. Points that the robot was unable to mark are highlighted. In cases where the operator knows that certain parts of a hall are blocked these areas can be de-selected graphically. The graphical user interface also provides information about the estimate time left to complete the mission and the status of the batteries.

As an extra security measure the robot is equipped with a GSM module so that it can send text messages (SMS) to the operator. Such messages are sent when the batteries need to be changed, if the robot is stuck or lost somewhere.

5 Localization

Because of the nature of the task the localization component in the system was one of the keys to success. It has been shown by many that indoor localization using a feature representation can be made highly accurate [?]. The challenge in this application is that the environment is much larger, the floor is not as even and the map might not be completely correct and complete. As was seen in Figure ?? the distance between features is large, in the order of 25m. The surface of the floor varies between relatively smooth concrete to uneven asphalt where the robot can almost get stuck in places.

Two types of features are used in the map, i) lines which correspond to walls and ii) pillars. A wall provides accurate information about the distance and orientation relative to it. However, it would not be possible to rely only

on the walls as they are not visible at all from some poses in the large hall. Notice also that the uneven floor limits the maximum usable range as well. The different features are kept in different layers in the CAD map so that they can be distinguished easily. The pillars come in many different shapes and are defined as sets of lines and arcs. An Extended Kalman Filter (EKF) is used to fuse the information from the measurements of the features and the odometry.

5.1 Lines

There are many examples of using lines for localization in the literature. What makes this case special is the large distances. At typical indoor distance ($< 10\text{m}$) the foot print of the laser scanner is often small enough to neglect. However, in the large halls this will lead to large errors. The beam width of the SICK LMS 291 is in the order of $0.01\text{rad} \approx 0.60^\circ$ [?]. The footprint can easily be in the order of a meter in some cases. When performing line extraction for the localization this must be taken into account. Figure ?? shows how the compensated point can be calculated assuming that the reflecting surface is rough.




Fig. 3. The center point (blue) is reported by the laser scanner and the one at the end of the arc (red) is an approximation of the true reflection point

The problem with this model is that the effective beam width, depends (at least) both on distance and on the reflecting material. In the environment under consideration some walls are very smooth, whereas some are rough. Just as ignoring the beam width is bad, erroneously compensating for it is equally bad. Therefore, to avoid having to augment the map with reflectivity information for the walls another approach is taken. Only wall segments that can be observed close to perpendicular are used for localization. Because of the layout of the halls this is generally not a limitation.

5.2 Pillars

The description of the pillars gives a flexible representation that can account for the many types of pillars that are present in the environment. When matching scan points to the contour of the pillars the ICP algorithm [?] is used. The contour is sampled to generate the reference point to match the scans against

(see Figure ??). The pillars all have similar painted surfaces and thus the compensation illustrated in Figure ?? can be used. Only the sample points that correspond to model segments that are visible from the current pose are used for the different pillars. After associating scan points to pillars based on their position the ICP algorithm is run to associate the points with the different segments on the pillars. The position of each scan point is then compensated for (see Figure ??), according to the associated pillar segment. After this a second turn of the ICP algorithm is run to make adjustments caused by the beam width compensation.




Fig. 4. Left: Sampled contour as reference points for the ICP algorithm. Mid: Close range observation. Right: Long range observation.

6 Experimental Evaluation and Modifications

At the end of the project, in February 2003, a few month before the system was put into operation, a full scale test of the system was performed. The total number of coordinates was 722, out of which 518 should be marked³. The length of the trajectory was 4374m and contained 246 stands corresponding to a total stand area of 10150m². The total time for the task was 4.5h and the robot succeeded to mark 492 out of the 518 points.

To evaluate the marking accuracy, the 492 marked coordinates were hand measured. Figure ?? shows the error distribution both in terms of the absolute error and separated into the x - and y -components. The standard deviation for the absolute error was 18mm and the average error was 28mm.

There are many sources of errors. One was mentioned before dealing with the printer. Furthermore, the CAD map that is the basis for the localization is accurate to about 1-2cm. The laser sensor also has a systematic error in the order of centimeters [?]. Finally, the hand measured positions also have an error associated with them. Put together these errors can explain most

³ The rest defined directions

of the fluctuations in the accuracy. The really large errors ($\geq 75\text{mm}$), which represents less than 3% of the total number points, are the result of data association errors. Many of these are caused by unmodeled objects being in the scene. For example, the pillars are often used to lean all sorts of things against to get them out of the way. At large distances it is very difficult to detect such errors.

Fig. 5. Error distribution for the marked point for the first full scale test.

There were two causes for the 26 coordinates that could not be marked.

- In some places the trajectory was planned such that only a single wall was visible, i.e. the robot was driving close to and toward the wall. In these cases the pose of the robot could not be determined with enough accuracy and these coordinates were skipped.
- The robot is unable to mark coordinates that are too close to obstacles. Furthermore, in the current implementation the robot also wants to be able to drive straight for a distance before the mark is made to stabilize everything and reach the right printing speed.

6.1 Modifications

Some modifications were made to the system before it was put into operation in August 2003.

A method based on using temporary localization stick landmarks was developed. The halls at the StoFair all have some coordinates permanently marked on the floor to assist in the manual marking process. Whenever such a point comes into view the system searches for a stick there. If it is found very close to the predicted position it is added to the map and henceforth used in the localization process. The sticks have their own layer in the CAD map.

In many cases where the robot is uncertain of its pose for a mark, the pose can be determined better if the robot is reversing the mark direction. Therefore, the robot attempts to mark a point in the opposite direction before reporting a failure due to insufficient accuracy. The result of this modification is that the robot is able to mark all positions free from obstacles under normal conditions.

6.2 Re-evaluation

The same fair that was evaluated in 2003 was evaluated again in 2005. This time a total of 194 coordinates spread out over the hall were examined manually. Out of these, 188 points were marked and 6 had been skipped because

they were too close to obstacles (pillars). The ability to turn around and mark in the other direction reduced the number of failures. The marked coordinates had the same average error and standard deviation as in the evaluation two years earlier.

The system has now become a natural part of the marking process. It has been used for marking every fair in the three main halls since it was taken into service in 2003. The robot operators are from the same group of people that used to perform the manual marking. They do not have any special computer or robotics training which indicates that the system is easy enough to use by people that have no robotics background.

7 Summary and Conclusion

To summarize we have presented a system for automatic markings of floors. It has been used to mark every fair in the three main halls at StoFair since August 2003. The system has shown that high accuracy can be achieved in very large areas, but that the laser scanner has to be modelled better than in typical indoor applications.

Figure ?? shows an example of how it might look when two fairs have been marked in different colors. The dashed lines have been overlayed graphically to highlight the location of the walls for the two stands. Notice how one of the two texts will be visible even after the construction as it ends up outside the wall.




Fig. 6. An example of how it looks on the floor when two fairs have been marked

The most important indicators of success in this project is the fact that the system has been used now for almost two year continuously and is greatly appreciated by the staff and that a second platform has been built. At first many considered it to be mysterious and some were sceptical but today it is taken for granted. As a final remark the time to mark a typical fair has been

cut from 8h with two men to about 4h with one robot and one man that potentially can run more than one robot at a time.