A Mobile Robot System for Automatic Floor Marking

Patric Jensfelt, Gunnar Gullstrand
CAS,NADA,KTH
patric@kth.se, gunnar@gullstrand.nu

Erik Förell
Stockholm International Fair
erik.forell@stofair.se

April 10, 2006

Abstract

This paper describes a patent awarded system for automatically marking the positions of stands for a trade fair or exhibition. The system has been in operation since August 2003 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. What used to be a job for two men over eight hours now takes one robot monitored by one man four hours to complete. The operators of the robot are from the same group of people that previously performed the marking task manually.

Environmental features are much further away than in most other indoor applications and even many outdoor applications. 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. Reaching the required level of accuracy was only possible by proper modelling of the laser scanner. The system has been evaluated by hand measuring 680 marked points.

To make the integration of the robot system into the overall system as smooth as possible the robot uses information from the existing CAD model of the environment in combination with a SICK LMS 291 laser scanner to localize the robot. This allows the robot to make use of the same information about changes in the environment as the people administrating the CAD system.

1 Introduction

The Stockholm International Fairs (StoFair)[1] is the leading exhibition, fair and congress organizer in Scandinavia. It has 56,500 m^2 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. There are three main exhibition halls that are about 70m wide and between 220 and 270m long, i.e., extremely large single room indoor environments.

[1]Stockholm International Fairs, SE-125 80 Stockholm, SWEDEN. http://www.stofair.se/

The work with a fair or exhibition starts with the planning phase where the area in the hall is sold and the layout is worked out. This phase can be very long. It is not uncommon that the planning starts many years before the actual event. Then, 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 place 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 and typically under stress.

Traditionally the marking has been carried out manually with tape and a long 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. Furthermore, after all the tapes have been put on the floor information that give the identity of the stand is added separately. A large exhibition can have several hundred stands and for each stand several coordinates are marked (see Figure 1). This in combination with the odd hours motivates automation of the process.

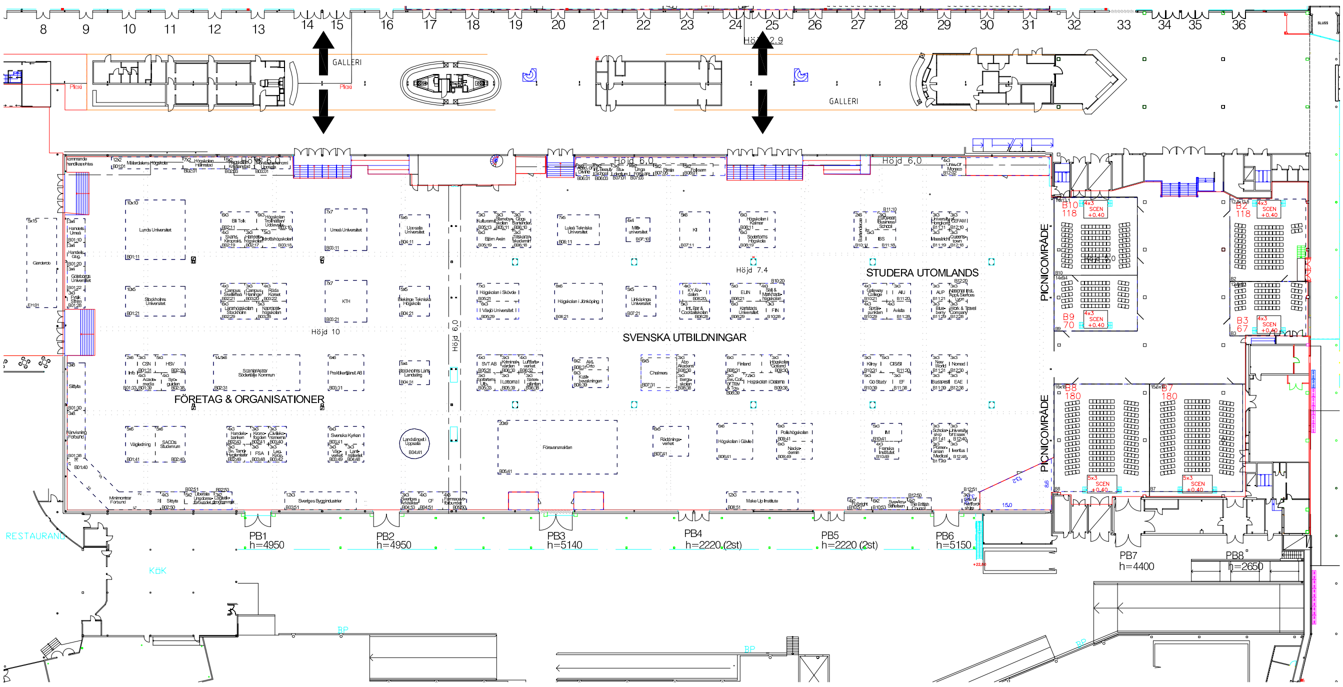

Figure 1: Example of a typical fair layout in one of the halls at the Stockholm International Fair. The stands are marked with dashed lines and are located in the interior of the hall. On the right hand side of the hall, sliding walls are used to create seminar room areas.

The automation of the marking process can be realized in several ways and with different levels of autonomy. It is clear that most of the time is spent on finding the location to mark and not the marking itself. One set of solutions thus provides only this functionality, i.e., help find where to mark. An example of this type could be to have one or more actuated laser pointers mounted in the ceiling pointing out the location of the coordinates. One problem here is the sensitivity to small errors in the angle that would lead to very large errors in position. Another option would be to mount laser pointers on sliders along two walls and let the intersection of the laser beams mark the point. This would however require that the hall was completely empty and had no pillars blocking the line of sight.

Increasing the autonomy one step further requires a moving device, a mobile robot, which can do the marking on its own. 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.

To realize an autonomous marking system there are several subproblems that need to be addressed. The robot needs to be able to localize accurately, navigate safely through the environment and be able to mark locations on the floor. In addition all the subsystems need to be integrated in a stable and reliable system that is easy enough for normal people to use. There are no systems available off the shelf yet for this task to the best knowledge of the authors.

1.1 Contributions

Some of the contributions of this paper that are worth highlighting are:

- The paper demonstrates accurate localization in an environment that is truly large scale. What really sets it aside from earlier work is the large distance to the landmarks.
- The fact that features are at long range has brought forth the need for modelling of the laser scanner in a way typically not needed in a normal indoor environment.
- The system has been thoroughly evaluated by hand measuring the marked points to verify the accuracy of the system. Two such experiments are presented in the paper where a total of 680 points were compared to ground truth, an effort not matched by many.
- The system has been operational for more than two years, i.e., it has been tested over an extensive period of time. It is now a natural part of the marking process at the Stockholm International Fair.

Another contribution of the paper that is not so much a scientific one is that it provides a very good example of what localization can be used for in robotics. It is often argued, and in most cases correctly, that it is of little use to strive for a localization system that has very high absolute precision. Most acts demand only that the relative precision is high such as when passing through a narrow door.

1.2 Outline

The rest of the paper is outlined as follows. Section 2 discusses related work and Section 3 lists some of the requirements on the system. Section 4 presents the sensor choice, sensor modelling and design of the localization and navigation subsystems. Section 5 discusses implementation issues. The results of an experimental evaluation of the system are given in Section 6. A summary and some conclusions can be found in Section 7. A hands on overview of the system and how it is used at StoFair is presented in Appendix A.

2 Related Work

In (Simmons & Koenig (1995)) it is said that "the major difficulty in achieving reliable navigation is in maintaining a good estimate of the robot's whereabouts". In the current application this is clearly true as the main objective is to accurately mark the location of a set of coordinates on the floor. Most of the robot localization research has been carried out in indoor environments. Much of the earlier work used ultra sonic sensors (Leonard & Durrant-Whyte (1991); Burgard et al. (1996)). Lately the laser scanner has become more or less a standard sensor for mobile robot systems (Fox et al. (1999); Arras & Tomatis (1999)). Most systems have been operating in relatively small office/domestic environments and in cases where the environment is large (Bosse et al. (2003)) it typically consists of long corridors and many small rooms. The StoFair environment is comparably very large.

There are commercially available off-the-shelf-solutions for localizing a mobile robot. The old standard for AGV:s (Automated Guided Vehicles) was to use wire guidance in which an electric wire was buried in the ground and the magnetic field generated by an AC current flowing in the wire was utilized. Today, laser technology is taking over as the standard solution. It offers a more flexible solution where the robot is not forced to travel along a predefined path. The LazerWay system from the company NDC[2] uses a scanning laser measuring the angle to reflective tapes for localization. The laser scanner from SICK has also been used in several other systems (see e.g. (Arras et al. (2003))). One of the advantages with using laser scanner such as the SICK sensor is that it can be used both for localization and for detecting obstacles. SICK AG also provides a product, the NAV200 sensor, which integrates the ability to navigate with respect to reflectors. The use of vision as the main sensory modality for localization and SLAM (Simultaneous Localization and Mapping) in particular has become the focus of much of the recent research. Evolution Robotics offers a system (Goncavles et al. (2005)) that uses a single low end web camera and targets consumer type applications where cost is a major factor. This system and many other vision based systems (Se et al. (2002); Sim et al. (2005); Jensfelt et al. (2006)) use the so called SIFT features (Lowe (1999)).

Another group of localization methods are based on radio technology. Many indoor environments are now equipped with wireless communication facilities. With such an infrastructure in place it is possible to implement a localization system based in measuring for example the signal strength (Krumm & Horvitz (2004)). The accuracy is typically in the order of meters. The use of RFID technology for localization has also been investigated (Hähnel et al. (2004)). RFID tags are becoming more and more frequent in the environment, a development driven for example by the idea to put a tag on every shopping item. The accuracy with a realistic density of RFID tags is currently also in the order of meters. Ultra-Wideband radio technology currently seems like one of the most promising radio solutions and provides means for an accuracy approaching centimeter level (Gezici et al. (2005)).

In outdoor applications GPS in combinations with inertial sensors are common (see e.g. (Nebot et al. (1997))). The advent of the RTK (Real-Time Kinematic) GPS (see for example (Lenain et al. (2003))) has increased the accuracy of GPS system to a few centimeters. The system is commonly used in for example surveying applications. However, the GPS signals are not available indoors and building a dedicated indoor GPS system with high enough accuracy is too expensive.

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 the obstacle at all times is

[2]Netzler & Dahlgren Co AB, http://www.ndc.se/

relatively small. Several methods have been proposed. Some of the more popular are the Vector Field Histogram (Borenstein & Koren (1991)), the Dynamic Window Approach (Fox et al. (1998)) and the Nearness Diagram (Minguez & Montano (2000)). The sonar and laser scanners have been the most commonly used sensors. In the current application obstacle avoidance is easier than under regular indoor conditions as the environment is very large, typically with few obstacles.

3 Requirements

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

- The environment is constantly undergoing changes. A system 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.
- It must be easy to maintain the map so that it can be adapted to changes in the environment.
- 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. It was clear from the beginning that an automatic solution was unlikely to compete in terms of relative accuracy of point positions within a stand. The points in the stand are manually marked relative to each other which gives millimeter accuracy. However, the absolute location of the stands can be off in the order of a meter in some cases. This sometimes resulted in problems when a pillar ended up inside a stand. The StoFair specified 3 cm as a guideline for the accuracy of an automatically marked point.
- The system should avoid collisions if there are objects in the way and report points as unreachable if blocked.
- 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. This would remove the need for the second step in the marking process where, historically, the information about the identity of the stand was added manually.
- 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.

4 System Design and Modelling

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. Sonar sensors do not have long enough range and the speed of sound would also lead to low update rates. At the time when the system was designed, radio technology did not present any solution that was accurate enough and was not too expensive (such as a local GPS system). The commercially available laser based localization systems were also rejected because of cost. The remaining candidate sensors were cameras and laser scanners. Vision was also ruled out since the lighting conditions are often quite bad, there is little texture and the uncertainty was too large whether or not the necessary accuracy could be reached. The choice was therefore to use a laser scanner that has proven its value in other systems. This choice also had the added benefit that it could serve as main sensor both for localization and navigation.

4.1 Sensor Modelling

At typical indoor distances (< 10m) and with an indoor laser scanner the foot print of the laser beams is often small enough to neglect. However, in the large halls this effects causes large errors if not dealt with. The beam width of the SICK LMS 291 is in the order of $0.01 rad \approx 0.60°$ (SICK AG (n.d.)). The footprint can be in the order of a meter in some cases when the distance is large.

Figure 2a) shows how the compensated point can be calculated assuming that the reflecting surface is rough enough to reflect the laser. Figure 2b) shows how the estimate of the direction of a line will be biased because of the beam width. The solid line is the true direction of the line whereas the dashed line is the resulting estimate if the beam width is not taken into account. In an indoor environment there are often walls on either side of the robot and the bias will be in different directions that results in a less apparent effect. However, when travelling along one long wall that is often the case outdoor and in the StoFair environment this effect can lead to large position errors.

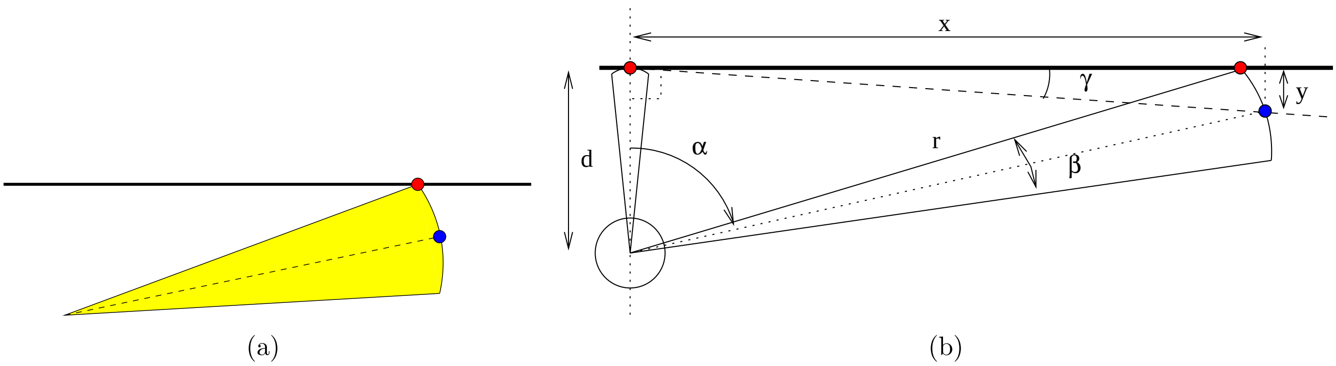

Figure 2: a) The center point (blue) is often assumed to be the point of reflection of a laser beam. The point at the end of the arc (red) is an approximation of the true reflection point. b) The thick solid line marks the true line and the dashed line is the estimated line resulting from using the center point model.

The bias in the line direction illustrated by Figure 2b) is γ. Using the notation from the figure we get

$$
\begin{aligned}
r &= \frac{d}{\cos(\alpha - \frac{\beta}{2})} \\
x &= r \sin\alpha \\
y &= d - r\cos\alpha \\
\tan(\gamma) &= \frac{y}{x}
\end{aligned}
$$

Using the fact that γ is small we get $\tan\gamma \approx \gamma$ and thus

$$
\gamma \approx \frac{y}{x} = \frac{d - r\cos\alpha}{r\sin\alpha} = \frac{\cos(\alpha - \frac{\beta}{2}) - \cos\alpha}{\sin\alpha} = \frac{(\cos\frac{\beta}{2} - 1)\cos\alpha + \sin\alpha\sin\frac{\beta}{2}}{\sin\alpha} \tag{1}
$$

Now using the fact that β is small we get $\cos\frac{\beta}{2} \approx 1$ and $\sin\frac{\beta}{2} \approx \frac{\beta}{2}$ which leads to

$$
\gamma \approx \frac{\beta}{2}, \tag{2}
$$

that is the bias in the line direction is half the beam width of the laser scanner.

One of the challenges with using this model is that the effective beam width depends (at least) both on distance and on the reflecting material.

4.2 Localization

Because of the nature of the marking task the localization component in the system is a key to success. Different map representations have been used in the literature, the two main directions are to use features (Leonard & Durrant-Whyte (1991); Arras & Tomatis (1999)) and occupancy grids (Elfes (1986); Burgard et al. (1996)). The laser scanner has been used in combination with both of these representation. Depending on the application one will be better suited then the other. Here, the feature based representation was used as it best accommodated the requirement that a user should be able to maintain the map easily. Furthermore, a CAD model of the environment already existed which would be easier to exploit in a feature map setting.

It has been shown by many that indoor localization using a feature representation can be made highly accurate (Arras & Tomatis (1999)). The challenge in this application is that the environment is much larger, the floor is not as even and the map might not be perfectly correct and complete. As can be seen in Figure 3 the distance between features is large. The map of the 7th floor of the Centre for Autonomous Systems (CAS) in Stockholm is overlayed to illustrate that the environment is orders of magnitude larger than most indoor environments. In essence it is three truly gigantic rooms. The area between the halls hosts restaurants, rest rooms, etc. Note that some parts of the halls do not have any walls on the side. These parts consist of glass or other materials that cannot be reliably detected by the laser. The distances between the pillars inside the halls are about 25m. The surface of the floor varies between relatively smooth concrete to uneven asphalt where the robot can almost get stuck in places.

It would not be possible to rely only on the walls, as they are not reliably detected from some poses in the large hall. Notice also that the uneven floor limits the maximum usable range as well. Even a small angle will make the robot look down into the ground or far up on a wall when the distances are long.

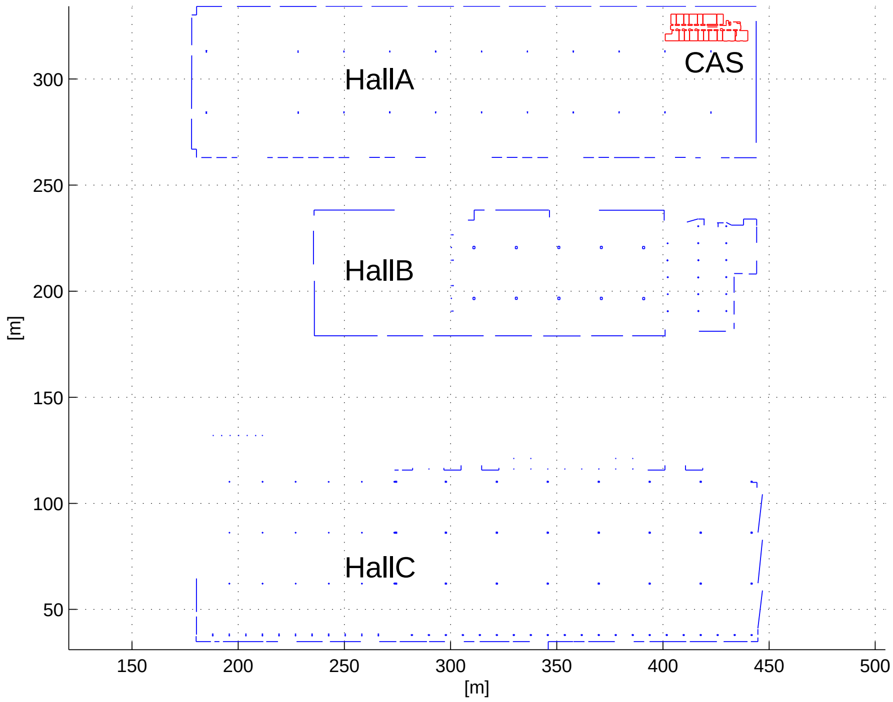

Figure 3: The wall and pillar features in StoFair environment. The map of the Centre for Autonomous Systems (CAS) at KTH, Stockholm, Sweden, has been overlayed for comparison.

4.2.1 Walls

There are many examples of using line segments for localization in the literature. What make this case special are the large distances. It was shown in Section 4.1 how the large beam width of the LMS 291 can lead to a bias in the estimated line direction. 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. At StoFair the surface of the walls are often changed. Accurately estimating and updating the reflectivity information for the walls is therefore difficult.

The effect of the line direction bias is particularly bad when a line segment is observed at an oblique angle far away from the robot. That is, when the projection of the sensor position onto the line direction vector is far the line segment (compare distance D in Figure 4). The biased angle will then through the "lever arm effect" lead to a large and incorrect compensation in the position.

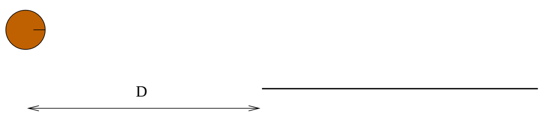

Figure 4: Line observed far away from robot at an oblique angle.

The geometry of the environment with four long walls in each hall provides for a solution. Only wall segments that can be observed within 15° from perpendicular are used for localization when the robot is marking. The measurement from a line is given by the distance to and the orientation of the line segment.

4.2.2 Pillars

The pillars come in many different shapes at StoFair. They are represented as sets of lines and arcs. Three examples of pillars are shown in Figure 5.

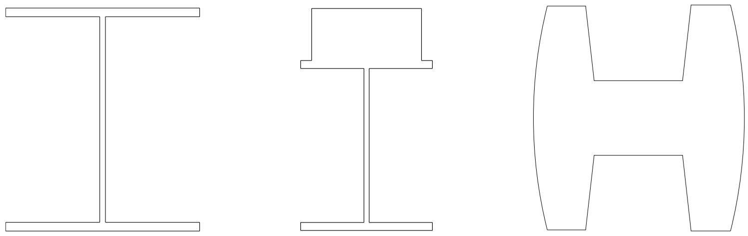

Figure 5: Some of the pillar shapes present in the environment. The depicted pillars are about 0.7m in diameter. The complicated shapes combined with long-range makes the data association problem very hard.

When treating the measurements from each pillar we divide between two cases, i) many points associated with a pillar ii) few points associated with a pillar. These two cases typically correspond to close range and long-range observations. Figure 6 illustrates the two measurement cases.

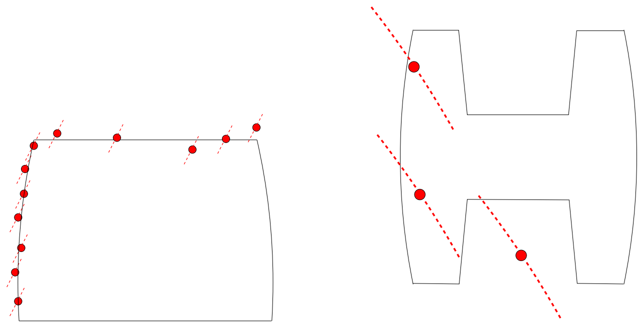

Figure 6: The amount of information provided by the laser scanner about the location of a pillar depends on the distance between sensor and objects. At long range few and uncertain measurements makes the estimation process hard. The dashed arcs corresponds to the foot print of the laser beams. Left: Close range observation. Right: Long range observation.

With many associated scan points, i.e. at close range, we can estimate the relative position of the pillar well. We do this by matching the scan points to the contour of the pillars using the Iterative Closest Point (ICP) algorithm (Besl & McKay (1992)) in the following steps.

1. Select the arc and line segments from each pillar that is visible from the current sensor pose. This way we remove the chance to get associations with parts of the pillars that are not visible.

2. Run the ICP algorithm to associate the points with the different segments on the pillars.

3. The position of each scan point is then compensated for (see Figure 2), according to the associated pillar segment[3].

4. Run a second round of the ICP algorithm to make adjustments caused by the beam width compensation.

The measurement uncertainty for each pillar is estimated based on the number of matches and the direction of the matched segments. The more points that are associated to a segment the lower the uncertain will be in the direction orthogonal to that segment's direction. The uncertainty model is important in order to be able to reason about the uncertainty of the robot pose and thus determine if the accuracy is high enough to perform the marking or not.

The case illustrated on the right side in Figure 6 represents a very difficult data association problem with a low number of data points, each with a large foot print. Because of the complicated shapes of the pillars the range is too uncertain especially since some of the measurements can be mixed pixel measurements (Ye & Borenstein (2004)). In this case we use only a bearing-only measurement. The angle is estimated as the average angle of the scan points associated with the pillar. The uncertainty is based on the field of view that the pillar occupies and the number of associated scan points.

Currently, the threshold between the two cases is 5 associated scan points. That is, the close range model is used when there are at least 5 associated points. This corresponds to a distance of about 20m.

4.2.3 Temporary landmarks

If enough of the features in the map are not visible the system will not reach the required level of accuracy. This is a problem sometimes when there is still material in the halls blocking the view to the map features. Such a situation is shown in Figure 7.

To allow the system to operate under such conditions a method based on temporary landmarks is devised. The idea is to let the system automatically detect which of these temporary landmarks are available. However, to do this in a robust way the robot needs to be able to identify the landmarks. The SICK laser scanner has the ability to measure intensity as well as range simultaneous. However, early experiments showed that to get reliable intensity readings at long range the targets had to be in the order of several tens of centimeter. This would result in very bulky landmarks.

The halls at the StoFair all have a grid with 8 m cells permanently marked on the floor to aid the manual marking process. During manual marking the location of the points are sometimes measured relative to this grid. This grid allowed for an alternative solution for identifying temporary landmarks based on their position. Thin vertical sticks are used as temporary landmarks. When a task is started, all temporary landmarks are treated as unverified and are not used for localization. Whenever a grid point, i.e. the location of a potential temporary landmarks, comes into view the system tries to determine if a stick is present or not. If a stick is found it is flagged in the map as existing and used for localization.

[3]The pillars are more similar in reflectivity than the walls and therefore it was possible to use the sensor model compensation here.

Figure 7: In some cases there is lots of building material blocking the view of the features in the map which makes localization difficult.

4.3 Navigation

The navigation system has 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. This section describes the design of the obstacle avoidance and the global path planning.

4.3.1 Obstacle Avoidance

Obstacle avoidance in exhibition hall environments is often not as challenging as in a typical 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 basis for the obstacle avoidance is a local occupancy grid where fusion of sensor data is performed over time. The system switches mainly between two modes depending on the amount of clutter when traveling between two locations.

Mostly Open Space In the normal mode a controller similar to the one used in the High Safety situations of the Nearness Diagram (Minguez & Montano (2000)) and in the Vector Field Histogram (Borenstein & Koren (1991)) is used. The idea is to make small adjustments to the course so that the robot steers clear of for example pillars. The following steps illustrate the algorithm.

1. Calculate the closest distance d_{min} to an obstacle
2. Create a binary histogram where anything closer than d_{min} plus some margin accounting for robot size and desired clearance is considered blocked.
3. Select an unblocked direction of motion closest to goal direction and a minimum distance m from any blocked sector.

Figure 8 shows how the robot adjusts the direction of motion to keep a minimum distance m from a pillar when heading for a goal location.

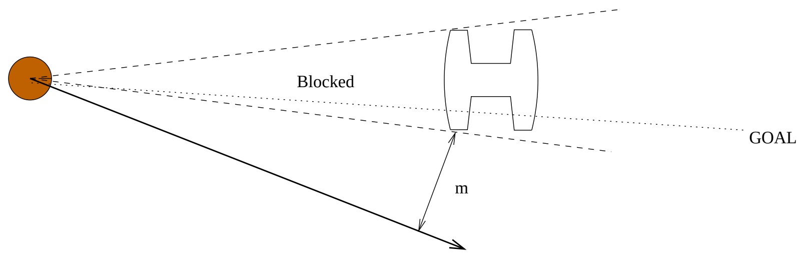

Figure 8: When navigating in open spaces the direction of motion is adjusted in a way similar to the High Safety situations in the Nearness Diagram (Minguez & Montano (2000)).

Clutter In situations with clutter the Global Dynamic Window Approach (GDWA) (Brock & Khatib (1999)) is used. The GDWA is an extension to the dynamic window approach (DWA) (Fox et al. (1998)) and adds the ability to plan the way out of more elaborate situations that DWA otherwise gets stuck in. This path planning uses the occupancy grid to find a path from the current robot position to the goal location. The DWA is then used when following the path. The switch to the GDWA from open space navigation is carried out if there is an obstacles closer than 2m in front of the robot and less than 0.5m to the side.

Collisions There are situations, although very rare, when the sensors of the robot are unable to detect obstacles and the robot collides. One such example is when pieces of wood or metal are left lying on the ground. In these cases the bumpers trigger a recover maneuver. In this maneuver the robot tries to drive away from the obstacle guided by which bumper element was activated by the collision (see Figure 9). An obstacle is added to the occupancy grid at the position of the bumper that detected the collision to avoid repeating the mistake.

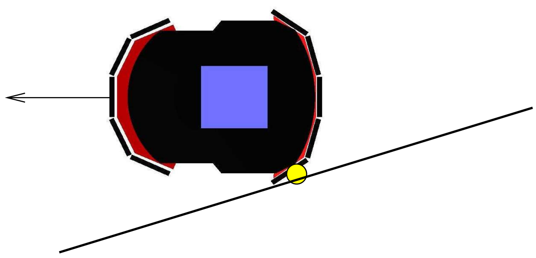

Figure 9: If a collision is detected a recovery maneuver is initiated where the robot drives away from the point of collision. An obstacle is added to the occupancy grid at the position of the bumper.

Timeout In some cases the robot is unable to get close to the point to mark, typically because there is building material blocking the path. The navigation system will make an effort to find a path around but to avoid this process from taking too much time a simple timeout mechanism is used. The straight-line distance to the point is calculated in every iteration and compared to the shortest distance so far. Every time the distance is the shortest so far, the timer is reset, i.e., whenever “progress” has been made. When the timer exceeds a threshold the point is deemed unreachable and the next point in the list is selected instead. Points are also skipped when it can be established that they are occupied.

4.3.2 Global Path Planning

The global path planning here is similar to the traveling salesman problem; each coordinate has to be visited once to mark it. To get a more flexible system that does not depend on any special structure for the fair layout the high-level path planning is performed off-board. The robot system bases the printing mission on the content of an XML-file that specifies where to mark and in what order. This way, the robot acts more like a plotter accepting print jobs and it is easy to change what to print without having to change the system on the robot.

The path planning is integrated into the CAD system. Currently a simplistic approach is used where the stands are completed sequentially. The advantage of this approach, besides being computationally attractive, is that the building of the stands in principal can begin immediately after the first stand or block of stands has been marked.

5 Implementation

5.1 Platform

To get a head start in the project it was decided to build upon a commercially available robot base. A custom made platform would allow for better performance but would take too much time to build 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 as 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.

The platform of choice was the Pioneer base from ActivMedia[4]. It comes with an on-board 800MHz computer running standard Linux. The left of Figure 10 shows the first prototype platform equipped with everything that it needs to carry out a marking mission. The standard batteries of the Pioneer robots are 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. A SICK LMS 291 was chosen, as it has better range than the more frequently used indoor SICK LMS 200.

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. 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. For the second

[4]ActivMedia Robotics, http://www.activmedia.com/

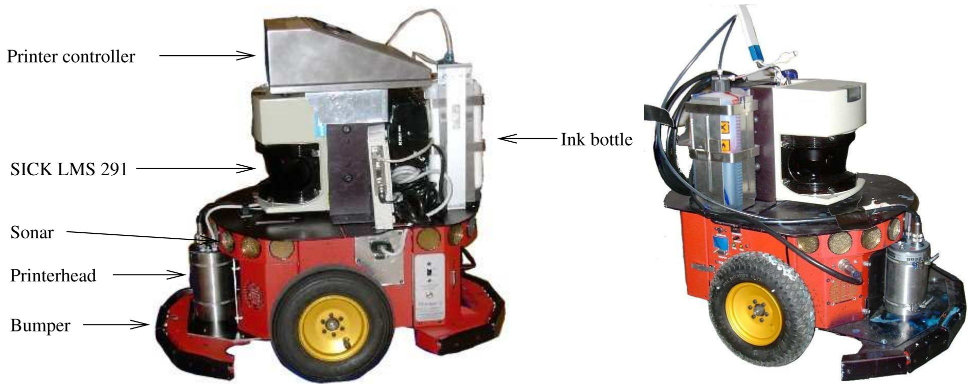

Figure 10: The first (left) and second (right) generation marking robots. The base is a standard Pioneer platform from ActivMedia. A SICK LMS 291 acts as main sensor. An industrial ink-jet printing system provides the ability to mark on the floor.

generation platform, right image in Figure 10, the bulky printer control unit has been replaced by a custom made circuit inside the robot.

5.2 Localization

An Extended Kalman Filter (EKF) (Jazwinski (1970)) is used to fuse the information from the measurements of the features and the odometry. The SICK laser scanner is configured to provide scans with 0.5° angular spacing over a field of view of 180°. The SICK laser scanner, produces these scans by interlacing two scans each with 1° resolution. At the long ranges in this application this separation is clearly visible when moving fast and then especially when rotating fast. A more rigorous treatment would account for this but in our application the speeds are relatively low.

The uncertainty of the localization system is monitored closely when the robot is about to mark a location. Unless it is accurate enough the robot will skip the point and head for the next one.

5.3 Map representation

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. The CAD system contains a lot of information that is not of interest for the robot such as electrical wiring, pipes, etc. These objects are difficult to automatically tell apart from information about the location of for example walls and pillars. To aid in this the robot map objects are kept in special layers in the CAD map, which allows the robot to easily identify what are walls, pillars and temporary landmarks.

5.4 Navigation

The vehicle control subsystem of the navigation system operates at 20Hz, driven by the rate of the odometric data from the robot base. The maximum speed had to be limited to 0.6m/s 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.

5.4.1 Marking

The printing on the floor is done at constant speed (0.15m/s). Besides marking the location of a certain coordinate the system can add extra information that makes it easier to identify the coordinate and to see how the carpets or walls should be placed. This extra information is given by two optional text strings (see Figure 11). The printing process consists of three main steps, marking the actual point and the two optional texts. The tolerances for the two optional texts are not as tight as for the marking point.

The signal to start the printing is controlled by the computer and is monitored at 20Hz. 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. To reduce this source of errors the distance to the mark position is calculated in each iteration. 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. This is currently one of the limiting factors in terms of marking speed. The higher the speed is, the larger the error. Further sources of errors are for example: i) there are no real-time guarantees in standard Linux, which can also contribute to errors and ii) the delay between when the signal is given and the printing starts. Experiments were conducted to measure this delay and compensate for it.

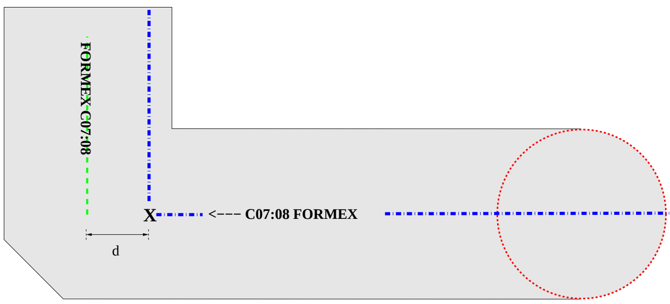

Figure 11: Each coordinate is marked with a 'X' and two optional text strings. Notice how the second text can be displaced w.r.t. the stand for better visibility. The location of the walls or the borders of the carpet are shown with the dash-dotted lines intersecting at the 'X'.

5.4.2 Path planning

The global path planing is integrated into the CAD system as a plug-in. Given the layout of the exhibition the path is generated by connecting the stands with a polygon line in the CAD application program. The plug-in automatically detects what stands are under the polygon line and in what order that are connected and automatically generates the XML-file. An example of a description of a point in the XML-file is shown in Figure 12, where the coordinates are given in mm in the CAD coordinate system.

```
<point id="1"  fairid="FORMEX" marked="FALSE" visible="TRUE">
  <text1>C07:08</text1>
  <text2>C07:08</text2>
  <x>298977</x>
  <y>41784</y>
  <Type>COORD</Type>
  <Symbol>X</Symbol>
</point>
```

Figure 12: Example of description of a point to mark in the XML-file.

The directions of the text strings (see Figure 11) are defined by the directions of the lines that connect the current coordinate with the previous and next one respectively. To handle the cases where the next or previous point is on another stand or otherwise unfit to define the direction special types of points ("VOID" instead of "COORD") that are not marked can be added.

The mission specification file in XML-format can be used directly as a report file that can be fed in again to continue a marking task that was stopped. An attribute ("marked) for each point element in the list tells if it has been marked or not.

5.4.3 Reversing marking direction

In many cases where the robot is uncertain of its pose for a mark, the pose can be determined more accurately if the robot is reversing the mark direction. This is the case for example when marking a point facing a long wall that only provides information about the orientation and the distance to the wall, but not the position along the wall. 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.

5.5 User Interface

The user interface is graphical and web-based. A screen shot from the applet is shown in Figure 13. It allows the operator to for example select a job and 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 will be blocked during the marking session these areas can be de-selected graphically which saves time as the robot will not even attempt to mark them. The graphical user interface also provides information about the estimated time left to complete the mission and the status of the batteries for example.

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 for example when the batteries need to be changed or if the robot is stuck or lost somewhere. This is an important feature as the operator

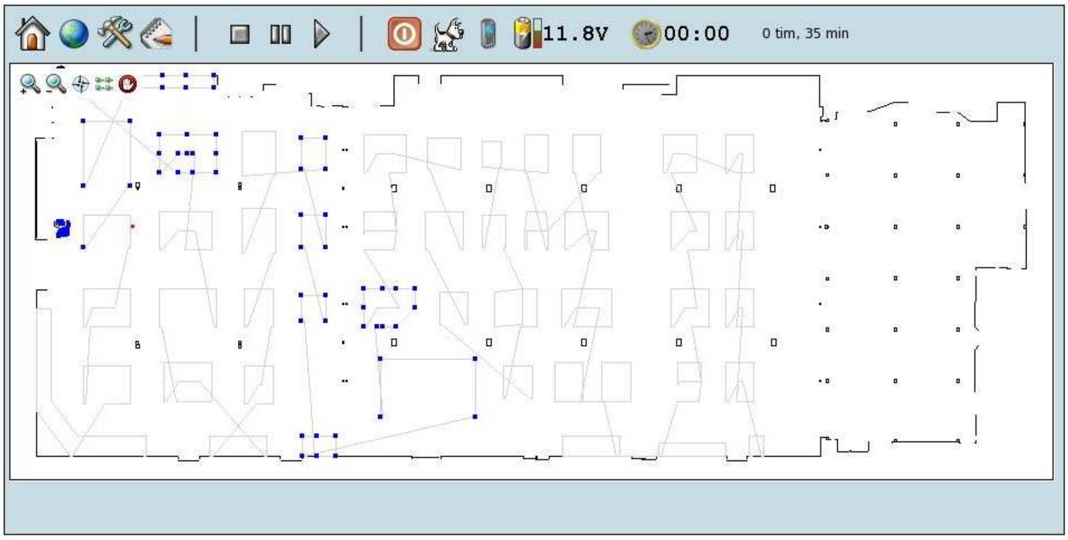

Figure 13: A screen shot from the user interface applet. The buttons at the top allows the operator to for example start, pause and stop a printing task, have the robot follow and look at messages from the robot. The planned robot path is shown. The blue dots in the corners of the stands are points left to mark. When a point is marked successfully it disappears and if it fails it is highlighted in red. In the current scene a report file (created at an earlier time) has been loaded into the system to mark the remaining points that were previously blocked by building material.

need to be informed as soon as possible in case there is a problem so that it can be dealt with. A delay could be expensive especially when the schedule is tight.

6 Experimental Evaluation

This section presents part of the experimental evaluation performed with the system. The staff at StoFair carried out the bulk of the evaluation, which resulted in changes in the way an operator can interact with the system and how the user interface looks.

The evaluation aims at establishing how accurately the system is able to mark the desired locations. Ideally the contribution from the positioning system and the marking system should be decoupled. However, this is difficult. The experiments performed to calibrate the marking device in Section 5.4.1 showed that at the speed of 0.15m/s the error under normal circumstances are less than 1cm. This together with the results below indicates that the errors in marking accuracy are mostly from the position uncertainty. This position uncertainty can only be indirectly observed by studying the accuracy in the final markings. There are no means to get the ground truth of the robot pose when the robot is moving and the position estimate when the robot is stationary includes a convergence phase that is not available when the robot is moving.

6.1 Experiment 1

At the end of the project, in February 2003, a few months before the system was put into operation, a full-scale test of the system was performed. The tests were carried out in an almost clean Hall C (see figure 3). The total number of points to mark was 518. The trajectory contained an additional 204 points for helping to define the direction (VOID points). The length of the trajectory was 4374m and contained 246 stands corresponding to a total stand area of 10,150m^2. Note that the stands sometimes share walls and form blocks of stands. The total time for the task was about 4.5h and the robot successfully marked 492 out of the 518 points.

To evaluate the marking accuracy, the 492 marked coordinates were hand measured[5]. Figure 14 shows the error distribution from the experiment both in terms of the absolute error and separated into the x- and y-components. The average absolute error was 28mm and the standard deviation was 18mm.

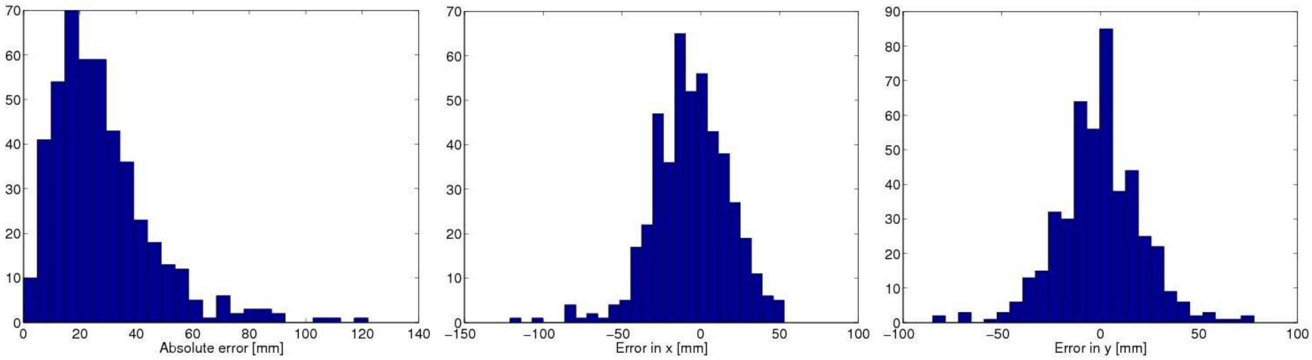

Figure 14: Error distribution for the marked point for the first full-scale test. From left: absolute error, error in x-direction and error in y-direction.

There are many sources of errors. First of all it is important to note that the positions of the marked points include errors that are not caused by the localization system such as those mentioned before, in Section 5.4.1, 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 (SICK AG (n.d.)). Finally, the hand measured positions also have an error associated with them. Put together these errors can explain most of the fluctuations in the accuracy.

The really large errors ($\geq$ 75mm), which represent less than 3% of the total number of points, are the result of data association errors. These errors typically occur when foreign objects are placed close to the pillars. It is for example common to lean material against the pillars as this is a natural way to get things out of the way for trucks, cranes, etc driving in the environment. At a long distance it is difficult to tell these objects apart from the pillar and this cause errors.

Figure 15 illustrates and example where a paper tube has been places inside one of the pillars (same shape as rightmost pillar in Figure 5). The tube is shown to the left in Figure 15. To the right in the same figure the resulting mark with and without the tube are shown. The mark denoted by 2 is the one without the tube and this mark is less than a centimeter from the correct absolute position. Notice how the mark when the tube was in place, marked 1, has been shifted away from the pillar which is a result of erroneous ICP matching with the foreign object. Also shown to the right in Figure 15 is the manually marked position (the right angle marked in tape above 1 and 2). This mark is more wrong then both of the automatically marked positions in global coordinates.

[5]Hand measuring the points gave the authors a valuable lesson in how tedious the manual marking is and took 4 people a full day to complete.

Figure 15: An example of how a foreign obstacle can cause marking errors. Left: images shows how a paper tube has been leaned against a pillar. Right: The resulting marks with and without the tube present are circled and denoted 1 and 2 respectively. Mark 2 is very close to correct and the error for 1 corresponds to the size of the tube. The manual mark is also shown next to the two circles. This mark is more incorrect in absolute sensor then both of the automatic marks.

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. This resulted in the addition described in Section 5.4.3.
- The robot is unable to mark coordinates that are too close to obstacles. Furthermore, the robot needs to drive straight for a distance before the mark is made to stabilize everything and reach the right printing speed. This corresponds to the stretch between the circle and the beginning of the horizontal text in Figure 11.

6.2 Experiment 2

The same fair that was evaluated in 2003 was evaluated again in 2005. This time some temporary sticks where used by the operators. Note that the fairs do not have the same layout every year. This time a selection of 194 points spread over the hall were investigated. Out of the 194 investigated points, 188 points were marked and their position were hand measured for the evaluation. Six points out of the 194 had been skipped because they were too close to pillars.

Figure 16 shows the results for the global accuracy of the points. The average absolute error is here 27mm and the standard deviation is 17mm, i.e. very close to the result in the first experiment.

In the manual marking process the relative measures are typically very accurate thanks to the tape measure. However, as illustrated in Figure 15 the global precision is not as high. As the relative measure is more important for the construction of the stands it is interesting to analyze it for the automatic system. To do so the relative position of consecutive hand measured points were compared to the distance given by the print job specification. That is, if points 1,2 and 4 were

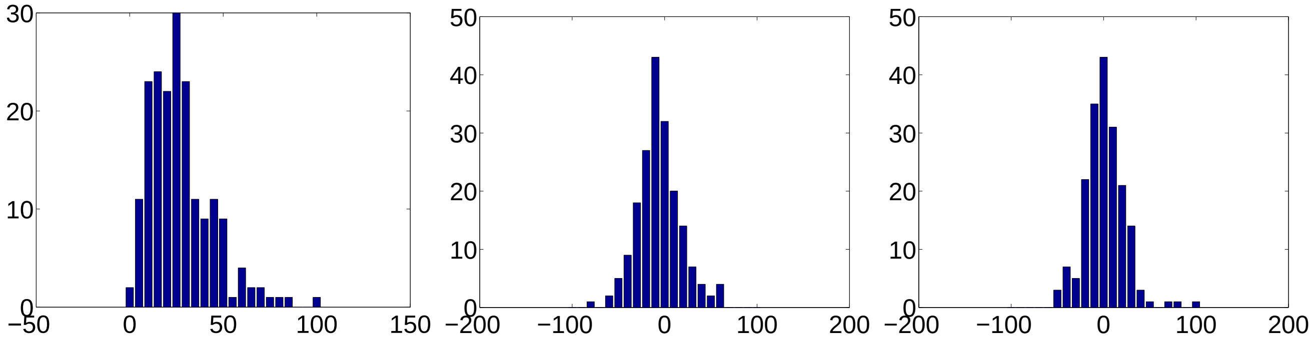

(a) Global absolute error (mm) (b) Global x error (mm) (c) Global y error (mm)

Figure 16: Histogram showing the error in the global position of the 188 hand measured points.

hand measured the relative position of 1 and 2 as well as 2 and 4 were analyzed. Notice that the relative accuracy is estimated indirectly by using the absolute position of both coordinates. The results are presented in Figure 17. The average absolute error was 30mm for the relative distance and the standard deviation 21 mm. The guideline for the point accuracy specified that each point should have an accuracy of 30mm which allowed the relative distance (the interesting measure) to be up to 60mm wrong.

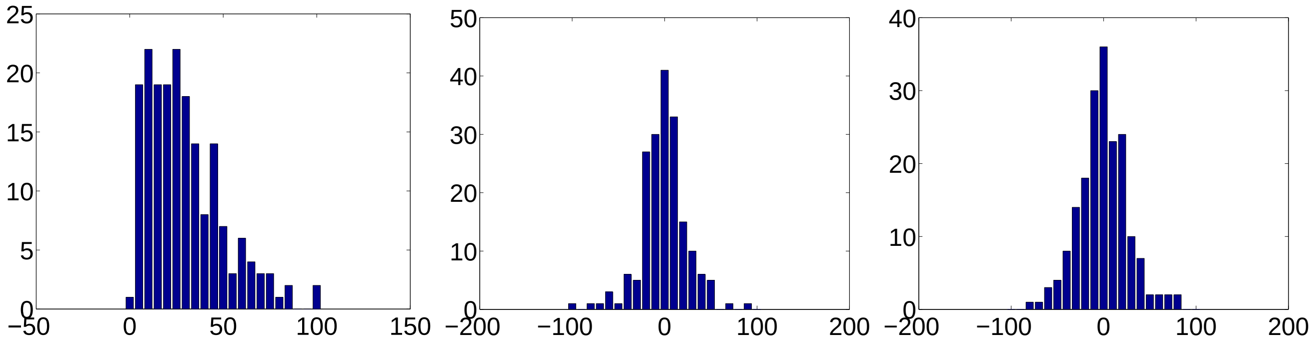

(a) Relative absolute error (mm) (b) Relative x error (mm) (c) Relative y error (mm)

Figure 17: The error in the relative measure between two consecutive marked and measured points. Error both given as absolute value and split into the x- and y-coordinates.

Note that the robot will mark many of the consecutive points with different direction of motion (compare the four corners of a rectangle that all will have). This means that the robot looked at different environmental features for the two marks.

6.3 End result

Figure 18 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 walls.

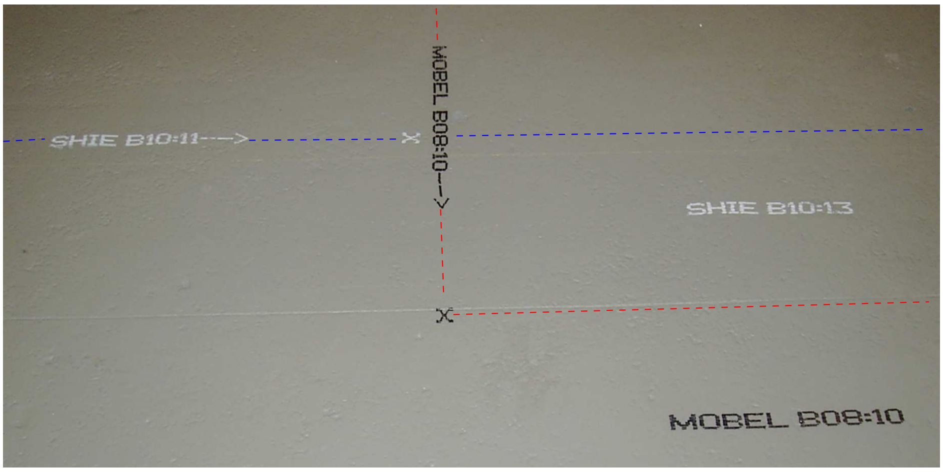

Figure 18: An example of how it looks on the floor when two fairs have been marked. The blue and red dashed lines have been overlayed to indicate where the walls would be placed. Note how one of texts is outside the wall and thus still readable when the walls have been built.

As a final remark a few words about the gain in productivity that the system provides. A typical fair that took two men working 8h each to mark can now be completed by the robot and an operator in about 4h. The operator can do other things while the robot is marking and could run more than one robot at a time to even further increase efficiency. This shows that there could be a business for such a robot marking system in countries where labor is expensive. This is typically the case in the countries where most of the fairs are located. The system is covered in patents SE527498, WO2004107073 and EP1627264.

7 Summary and Conclusion

To summarize we have presented a system for automatic marking of floors. The environment is truly large scale with long distances between features. This type of environment is close to an outdoor environment in many aspects. The laser scanner has been modelled in more detail than what is customary in typical indoor environments.

The system has been thoroughly evaluated. The paper presented two evaluations where the location of a total of 680 marked points were hand measured and compared to the desired position. For global position accuracy the average absolute error was found to be around 28 mm with a standard deviation of 18 mm. The relative position error between two points was estimated to an average of 30mm with a standard deviation of 21mm. As mentioned in Section 6 there are many sources contributing to these errors with the map accuracy of about 1-2 cm being one of the main contributions.

The system has been used to mark every fair in the three main halls at StoFair since August 2003. This together with the fact that the system has become a natural part of the marking process is the most important indicator of success in this project. 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.

A The New Marking Process Step-by-Step

This section provides a hands on overview of the system at StoFair and the way it is used. It has two main components, the CAD system and the robots. The CAD system provides the map information in DXF format and the trajectory/marking specification in XML format. The original idea was to transfer this information from the CAD system to the robots using wireless LAN. Practical reasons, such as the robot not being on when not marking, gave way for a solution with a USB stick instead. Figure 19 gives an overview of the system. In principle any number of robots could be operating in parallel. Currently the system does not support communication between robots to perform the same task together. At the StoFair the robots are used for different fairs when operating in parallel.

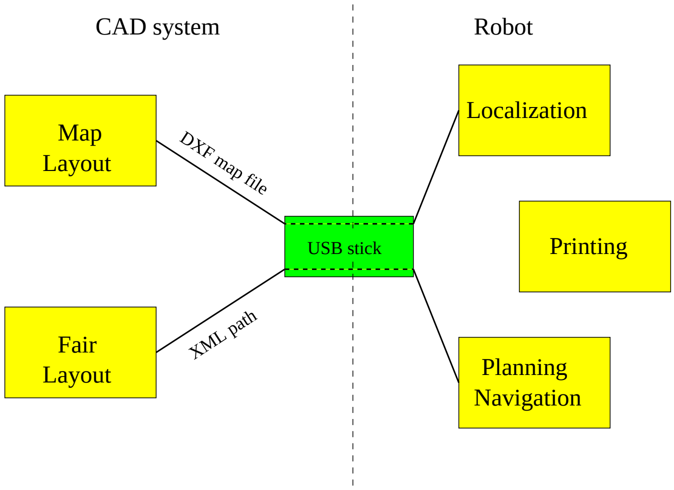

Figure 19: The system at StoFair consists of two parts. The CAD system provides the map information and is used to plan that layout of the fairs. This is also where the trajectory planning is performed. The trajectory in XML format and the map in DXF format is given to the robot using a USB stick.

The following gives a short step-by-step description of how to work with the system.

1. Produce the map in DXF format. This is mostly a one time effort, even though the map needs to be kept up to date if there are major changes. In places where a CAD drawing exists this is straightforward. An obvious extension to the system would be to use SLAM techniques to build the map automatically. Using an existing CAD drawing has the advantage though that you only keep one representation of the space, the same for robot and humans.

2. Produce the XML trajectory specification file. At StoFair this is done inside the CAD program that also handles the layout of the fair. The file is generated by drawing a polygon connecting the blocks of stands that should be marked in the desired order. A plug-in to CAD automatically detects which stands are under the polygon and in which order the corresponding points should be marked.

3. Make sure that the area to mark is free from obstacles. This is not because obstacle avoidance does not work on the robot but simply that every point to mark that is occupied by something

means one point that has to be marked again which takes time. The localization performance is also improved if the features in the map are visible. The floor was typically clean before as well when performing the manual marking, but obviously the manual process was not as sensitive to obstacles and debris on the floor as the automatic robot system.

4. If there is a lot of material around the marking region that blocks the view of the features in the map place some sticks around it to assist the robot in the localization task.
5. Start the robot
6. Plug in the USB stick with the map (if changed) and the XML trajectory file.
7. Open up the web-based interface and select what job to run.
8. Indicate on the screen where the robot is[6].
9. Press the start button
10. Watch out for SMS messages about the need for battery changes or the robot being stuck somewhere.
11. Repeat from 2 for next fair.

References

Arras, K., Philippsen, R., Battista, N. T. M. de, Schilt, M., & Siegwart, R. (2003). A navigation framework for multiple mobile robots and its application at the expo.02 exhibition. In *Proc. of the ieee international conference on robotics and automation (icra'03).* Taipei, Taiwan.

Arras, K., & Tomatis, N. (1999, September 6-9,). Improving robustness and precision in mobile robot localization by using laser range finding and monocular vision. In *Proc. of the 3rd european workshop on advanced mobile robots (eurobot'99)* (pp. 177–185). Zürich, Switzerland.

Besl, P. J., & McKay, N. D. (1992). A method for registration of 3-D shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, *14*(2), 239–256.

Borenstein, J., & Koren, Y. (1991, June). The vector field histogram - fast mobile obstacle avoidance for mobile robots. *IEEE Transactions on Robotics and Automation*, *7*(3), 278–288.

Bosse, M., Newman, P., Leonard, J., Soika, M., Feiten, W., & Teller, S. (2003). An Atlas framework for scalable mapping. In *Proc. of the ieee international conference on robotics and automation (icra'03).* Taipei, Taiwan.

Brock, O., & Khatib, O. (1999, May). High-speed navigation using the global dynamic window approach. In *Proc. of the ieee international conference on robotics and automation (icra'99)* (Vol. 1, pp. 341–346). Detroit, Michigan, USA.

[6]An automatic process is possible. However, the highly symmetric and huge environment would sometimes force the robot to move far to unambiguously determine the position and it only takes a click on the screen to do it manually

Burgard, W., Fox, D., Henning, D., & Schmidt, T. (1996, August). Estimating the absolute position of a mobile robot using position probability grids. In *Proc. of the national conference on artificial intelligence (aaai-96)* (pp. 896–901). Portland, Oregon, USA.

Elfes, A. (1986, April). A sonar-based mapping and navigation system. In *Proc. of the ieee international conference on robotics and automation (icra'86)* (Vol. 3, pp. 1151–1156). San Francisco, CA.

Fox, D., Burgard, W., Dellaert, F., & Thrun, S. (1999, July 18-22,). Monte carlo localization - efficient position estimation for mobile robots. In *Proc. of the national conference on artificial intelligence (aaai-99).* Orlando, Florida, USA.

Fox, D., Burgard, W., Thruna, S., & Cremers, A. B. (1998, May). A hybrid collision avoidance method for mobile robots. In *Proc. of the ieee international conference on robotics and automation (icra'98)* (Vol. 2, pp. 1238–1243). Leuven, Belgium.

Gezici, S., Tian, Z., Biannakis, G. B., Kobayashi, H., Molisch, A. F., Poor, H. V., et al. (2005, July). Localization via ultra-wideband radios. In *Ieee signal processing magazine* (pp. 70–84).

Goncavles, L., Bernardo, E. di, Benson, D., Svedman, M., Ostrovski, J., Karlsson, N., et al. (2005, apr). A visual front-end for simultaneous localization and mapping. In *Proc. of the ieee international conference on robotics and automation (icra'05)* (p. 44-49). Barcelona, Spain.

Hähnel, D., Burgard, W., Fox, D., Fishkin, K., & Philipose, M. (2004). Mapping and localization with rfid technology. In *Proc. of the ieee international conference on robotics and automation (icra'04).* New Orleans, LA, USA.

Jazwinski, A. H. (1970). *Stochastic processes and filtering theory* (Vol. 64). Academic Press.

Jensfelt, P., Folkesson, J., Kragic, D., & Christensen, H. I. (2006, March). Exploiting distinguishable image features in robotic mapping and localization. In H. I. Christensen (Ed.), *1st european robotics symposium (euros-06).* Palermo, Italy.

Krumm, J., & Horvitz, E. (2004, August). Locadio: Inferring motion and location from wi-fi signal strengths. In *Proc. of mobiquitous 2004, first annual international conference on mobile and ubiquitous systems: Networking and services* (pp. 4–14). Boston, MA, USA.

Lenain, R., Thuilot, B., Cariou, C., & Martinet, P. (2003). Adaptive control for car like vehicles guidance relying on RTK GPS: Rejection of sliding effects in agricultural applications. In *Proc. of the ieee international conference on robotics and automation (icra'03)* (p. 115-120). Taipei, Taiwan.

Leonard, J., & Durrant-Whyte, H. (1991). Mobile robot localization by tracking geometric beacons. *IEEE Transactions on Robotics and Automation*, *7*(3), 376–382.

Lowe, D. G. (1999, September). Object recognition from local scale-invariant features. In *Proc. of the international conference on computer vision (iccv 1999)* (pp. 1150–57). Corfu, Greece.

Minguez, J., & Montano, L. (2000). Nearness diagram navigation (nd): A new real time collision avoidance approach. In *Proc. of the ieee/rsj international conference on intelligent robots and systems (iros'00)* (pp. 2094–2100). Takamatsu, Japan.

Nebot, E., Sukkarieh, S., & Durrant-Whyte, H. (1997, September). Inertial navigation aided with GPS information. In *Proc. of fourth annual conference on mechatronics and machine vision in practice* (pp. 169–174). Toowoomba, Qld., Australia.

Se, S., Lowe, D. G., & Little, J. (2002). Mobile robot localization and mapping with uncertainty using scale-invariant visual landmarks. *International Journal of Robotics Research*, *21*(8), 735–58.

SICK AG. (n.d.). *LMS200 / LMS211 / LMS220 / LMS221 / LMS291 laser measurement systems, technical description.* (8 008 970/06-2003)

Sim, R., Elinas, P., Griffin, M., & Little, J. J. (2005, July). Vision-based slam using the rao-blackwellised particle filter. In *Proc. ijcai workshop on reasoning with uncertainty in robotics.* Edinburgh, Scotland.

Simmons, R., & Koenig, S. (1995). Probabilistic navigation in partially observable environments. In *Proc. of the international joint conference on artificial intelligence (ijcai'95)* (pp. 1080–1087). Montréal, Québec, Canada.

Ye, C., & Borenstein, J. (2004, October). A novel filter for terrain mapping with laser rangefinders. *IEEE Transactions on Robotics*, *20*(5), 913–921.