Grid-Based Navigation for Autonomous, Mobile Robots

Carsten BUSCHMANN, Florian MÜLLER and Stefan FISCHER

Institute of Operating Systems and Networks, Technical University of Braunschweig
Braunschweig, Germany, e-mail: (buschmann, fmueller, fischer)@ibr.cs.tu-bs.de

Abstract - **Navigation is a major challenge for autonomous, mobile robots. The problem can basically be divided into positioning and path planning. In this paper we present an approach which we call grid-based navigation. Though we also propose a scheme for path finding, we focus on positioning. Our approach uses minimal environmental infrastructure and only two light sensors on the mobile device itself. Starting out from a predefined tile and orientation in the grid, the mobile robot can autonomously head for destination tiles in the grid. On its way it determines the current location in the grid using a finite state machine by picking up line-crossing events with its sensors.**

1   Introduction

A key ability needed by an autonomous, mobile robot is the possibility to navigate through the space. The problem can basically be decomposed into positioning and path planning. Though we also propose a scheme for the latter, we clearly focus on the detection of the current position by monitoring grid crossing events by using two light sensors.

Especially if the robot is severely resource-constrained, simple schemes are favourable to elaborated algorithms. Rather simple sensors and actuators as well as a limited computing platform also demand simple, robust techniques due to inaccuracy and the lack of resources.

The paper is structured as follows: In section 2, we give an overview over the environment that the robot which is described in the following section needs for navigation. In section 4 we present the underlying algorithm and discuss various details. After that the software structure of the implementation is presented. The paper is concluded by a summary and an outlook to future work.

2   Environment

The robot navigates on a grid (see Figure 1) which regularly divides the ground into square tiles that are identified with Cartesian coordinates. As mentioned earlier, the robot starts out from a predefined position (e.g. the centre of tile 0,0) and orientation (e.g. North, which means looking up the y-axis). It is sufficient that the full grid is rectangular and all tiles are the same size. It is not required that the robot knows the dimension (n,m) of the grid. It is satisfactory to only request navigation to tiles that really exist. Heading from tile to tile, the robot distinguishes eight driving directions: north, north-east, east, south-east, south and so on. Depending on the size of the robot and the sensitivity of the


Figure 1: The grid

luminance sensors, different tile sizes and contrasts between plane and grid are possible. We used a grid made out of black self-adhesive tape of about 2 cm width on white ground with tiles of about 45 by 45 cm.

3   The Robot

The robot requires only very basic sensing and actuation capabilities. All that is needed are two ground-observing light sensors attached in a line orthogonal to the driving direction at the bottom of the robot, and two individually controllable wheels. For precise positioning (see Figure 3) the light sensors (A) are attached between the wheels (B) to make sure the sensor readings represent the wheels' location. As shown in section 4, the distance between wheels and sensors should be as small as possible to facilitate the alignment when crossing the grid orthogonally. At the back, the robots rests on a ball (C) that is spherical seated, allowing the robot to turn on the spot. The distance between the two light sensors should be no smaller than double the grid line width. This ensures that orthogonal grid line crossings where both light sensors get dark simultaneously (see Figure 2a) can be securely distinguished from diagonal crossings at the intersection of grid lines (see Figure 2b). If the sensors are located too close to each other, misinterpretations may appear (see Figure 2c).


Figure 2: Consequences of different distances between the light sensors

Basically only binary sensor readings are required; in order to cope with a wider spectrum of readings, an initial calibration on light and dark ground can be used to determine a threshold between the two by calculating the average of both readings. Figure 3 to Figure 5 show the robot we used for experimenting built out of the LEGO Mindstorms Robotics Invention System 2.0 [1].

It does not only offer a cheap possibility to experiment with various robots, it also allows for focusing our work on algorithmic aspects without putting much effort into mechanical construction. The entire navigation and positioning code executes on the yellow RCX control unit (D in Figure 5). The PDA shown on the robot in Figure 4 serves solely as a communication gateway to a Wireless LAN Network [2].

Nevertheless, the Robotics Invention Systems also yields some disadvantages. The RCX is extremely resource constrained. It offers only 32kB of memory firmware, program code and data. In addition, engines and sensors are rather imprecise, e.g. the driving speed of the robot heavily varies with the remaining battery capacity. All this demands simple, robust algorithms for navigation.

Figure 3: Bottom view of the robot with light sensors (A), wheels (B) and ball (C)


Figure 4: Robot with PDA for WLAN communication

Figure 5: Top view of the robot including the yellow RCX control unit (D)

4   Algorithm

The process of driving to a destination tile always follows the same 5-step-algorithm: (1) calculate the next tile to go to, (2) turn towards the direction of the next tile, (3) start driving and wait for line-crossing events to happen, (4) thereby decide which tile the robot arrived in, and (5) finally determine whether the robot has reached the destination tile. The 5-step-algorithm is also visualized in Figure 6. Determining the next tile is done via a breadth-first search: starting from the current tile, the coordinates of all adjacent tiles that have not been visited on the way to the current destination are

stored into a vector together with information on its predecessor. If the destination tile is not among the vector entries, again each entry's all adjacent and not yet considered tiles are stored together with their predecessor. The algorithm terminates when the destination tile is in one of the vectors. The path to that destination tile can then be derived using the chain of predecessors stored together with the coordinates. This rather naive approach has potential to be optimized. An example would be to employ branch-and-bound-techniques. Nevertheless, in large grids the applicability of such tree-based techniques is limited. In this case greedy algorithms might be an alternative. However, such path finding issues were not the focus of our research.

Once a path has been derived, the robot turns towards the direction of the next tile (which can be calculated from the current and next coordinates) by turning both wheels in different directions for a predefined time corresponding to 45 degrees (repeatedly if necessary). Because wheels and sensors are attached in a line, the light sensors will not "move" but only rotate.

The robot then starts driving, waiting for events to occur. We distinguish light sensor events and timeout events. A series of sensor events corresponding to the crossing of the grid) is always concluded by a timeout event. If no sensor events have occurred for a certain time period, it is assumed that the robot has reached the next tile. The length of the interval depends on the size of the tiles, the robot dimensions,
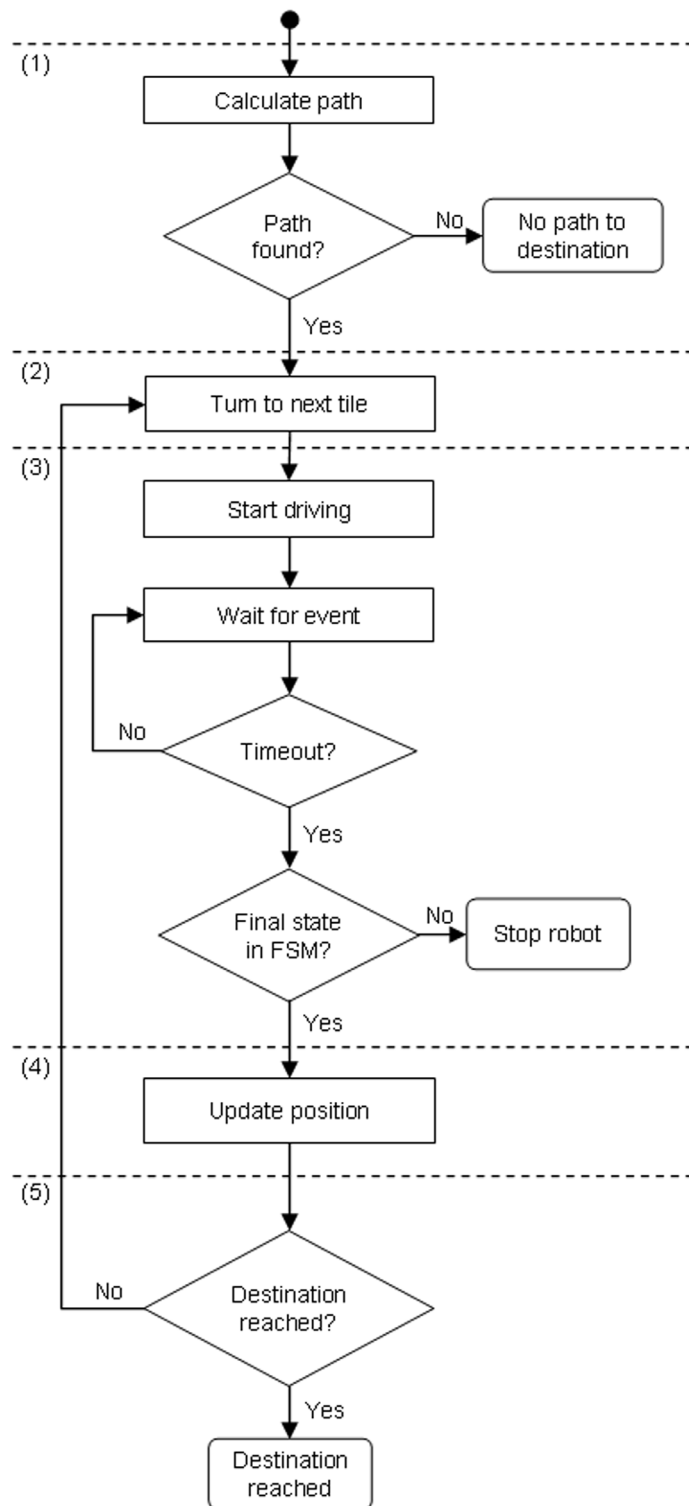

Figure 6: The 5-step-algorithm

path deviation and speed. If it is chosen too long (and tiles are small), the robot might already have started to cross the next lines at the opposite side of the new tile. If it is chosen too short, the timeout might occur during the crossing of the grid. For our implementation, we chose a timeout interval of four seconds.

Number, kind and order of sensor events determine which tile the robot drove to. Due to path deviations due to wheel slip, different engine powers etc. the robot does not necessarily arrive at the tile it was actually heading for.

Basically two ways of crossing the grid can be distinguished: the change to a tile in horizontal (East, West) or vertical (North, South) direction which we will call orthogonal crossing and the change to a tile at north-east, north-west south-east or south-west which we will call diagonal crossing for the rest of this paper.

Determining orthogonal crossings is relatively easy: after both sensors have gone dark and light again once and nothing happened afterwards for a while, the robot has reached the next tile. From the order in which the two events occurred a first indication can be drawn on the robot's deviation from the wanted direction (orthogonal to the grid line).


(a) (b) (c)

Figure 7: Robot alignment at orthogonal crossings

We use this fact to adjust the driving direction when crossing the line (see Figure 7). When the robot is heading for an orthogonal line crossing and the first light sensor event occurs, we stop the robot. Let us assume that it was the right sensor that got dark first (a). We then turn the corresponding (right) wheel back to make the robot turn until the according light sensor gets off the line again. The robot then again approaches the line until one light sensor recognizes it (b). This procedure is repeated until both light sensors are on the line simultaneously. By doing so the robot almost perfectly aligns with the grid line (c). Please note that this procedure is only invoked if the robot intentionally performs an orthogonal crossing.

When heading for a tile in diagonal direction (e.g. north-east), things are more complicated because five different cases can occur (see Figure 8). Depending on the position in the current tile, driving diagonally can lead to arriving in one out of three tiles that can be distinguished by the order in which the events occur. Please note that the upper left arrow and the lower right arrow denote special cases of an orthogonal crossing: though intending to cross horizontally, the robot effectively performs an orthogonal crossing. Unfortunately, this can only be detected after the robot arrived in tile 1 or 3 respectively.
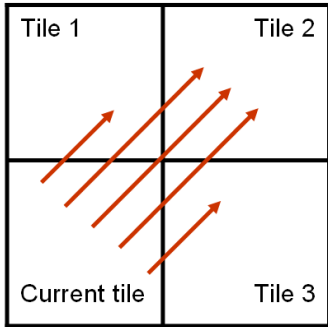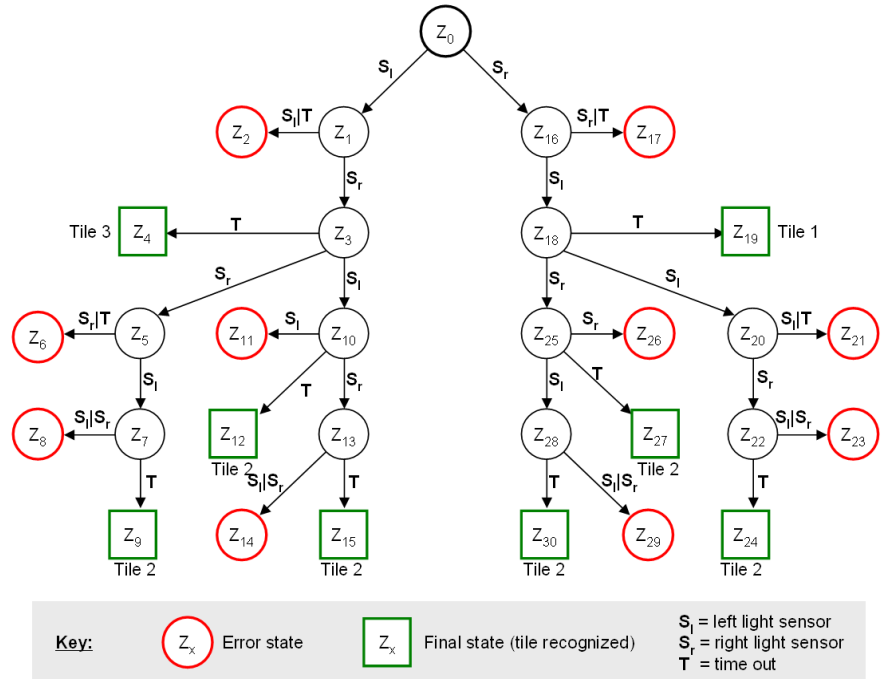
Figure 8: Possible cases when crossing diagonally

Figure 9: Finite state machine for determining which tile the robot arrived in

Key: Z_x Error state Z_x Final state (tile recognized)

S_l = left light sensor
S_r = right light sensor
T = time out

We use a finite state machine (FSM, see Figure 9) to figure out which of the five possible cases occurred: Events indicating that a line was crossed by a light sensor as well as timeouts induce state changes. The order in which these occur determines the destination tile. Non-valid event orders lead to error states which cause the robot to stop at the current position. Since unintended orthogonal crossing can only be detected after the robot arrived in the new tile, no alignment is carried out.

Once the robot has reached the next tile, it checks whether it has reached the destination tile. If not, it continues its travel towards the next tile on the way to the destination; otherwise it restarts the 5-step-algorithm for a new destination.

Videos showing a demonstration of the algorithm can be watched at http://www.ibr.cs.tu-bs.de/arbeiten/cbuschma/lego-pos/videos.html.

5 The Software

The software for the robot was developed using the BrickOS operating and programming system for the Robotic Invention System [3]. It is derived from LegOS [4] and offers numerous advantages such as an elaborate firmware, native coding in C/C++, multi threading including semaphores, direct hardware access as well as event based sensor readings and comfortable network routines for the IR interface [5]. In addition, the full code is available as open-source.

Figure 10: Software and hardware components

Figure 10 depicts the major components involved in implementing the algorithm: Ellipses represent hardware whereas boxes stand for software components. The Javastorms interface [6] and the infrared port are not part of the navigation itself but are used for reporting current locations and receiving future destinations from a wireless LAN via the Gateway as described in [2].

6 Conclusion and Future Work

In this paper we presented an approach for mobile robot positioning and navigation which we call grid-based navigation. It uses only minimal environmental infrastructure and two light sensors on the robot. Starting out from a predefined location and orientation in the grid, the mobile robot can autonomously head for destination tiles. On the way it determines its location in the grid using a finite state machine by picking up line-crossing events with its sensors. In addition, we demonstrated how we implemented the underlying algorithm in software and robot hardware. Nevertheless, the required grid limits this approach to scenarios where either a grid can be set up or a natural grid exists (e.g. due to a tiled floor). Nevertheless, it is easy to adjust both sensing and algorithm to different scenarios like a chessboard-like floor.

In this paper obstacles are not considered at all. A touch sensor could enable the robot to detect obstacles in tiles which then could be marked as blocked.

The path finding algorithm described in this paper is far from optimal. Greedy algorithms would allow for quite some improvement. A possible approach is to calculate the next tile to go to whenever the robot reaches a new tile. A possible metric is the line of sight distance to the destination tile. However in combination with obstacles blocking certain tiles this may lead to locally optimal but globally unfavourable paths if obstacles are arranged in a maleficent way. Nevertheless, strategies to escape such problems exist e.g. in the fields of linear optimization or geographic routing.

A major cause of errors is the robot's inaccuracy. Due to its simple construction, the robot shows a hardly predictable behaviour in terms of driving and turning speed as well as path deviation when going straight. Thus we found that it does not make much sense to develop more sophisticated algorithmic approaches. Hence we stuck to simple, qualitative approaches that are robust to such errors. Given a robot that is able to move more exactly and at defined speeds, new possibilities would

arise. By measuring the time intervals between light sensor events, the angle in which the robot crosses the grid line could be derived.

In addition to the navigation software on the robot we also developed a gateway application that resides on a PDA which is mounted on the robot and that allows for giving driving commands in the form of destination coordinates to the robot. Furthermore, the robot can send back its current location in the grid whenever it reaches a new tile. As future work we plan to replace the PDA by the wireless sensor mote ESB 430/2 developed at the Freie Universität Berlin [7]. We are currently implementing an infrared API to enable communication between the RCX and these devices. The motes will then be able to use the robot for autonomous wandering, e.g. broadcasting their current location periodically and thus providing a location service to a sensor network consisting of additional motes. When the robot comes into the communication range of a sensor mote, the mote gets a first indication on its location. Hearing the robot multiple times from different locations would enable the mote the estimate its position by averaging the locations transmitted by the robot.

7   Acknowledgements

References

[1]   Website of the Lego Mindstorms Robotics Inventions System 2.0: http://mindstorms.lego.com/eng/default.asp

[2]   Axel Wegener: WLAN-Gateway for the RCX Robot on a PDA, Semester project, Technical University of Braunschweig, 2003

[3]   Website of the BrickOS project: http://brickos.sourceforge.net

[4]   Website of the LegOS project: http://sourceforge.net/projects/legos

[5]   Jonathan B. Knudsen: The Unofficial Guide to Lego Mindstorms Robots, O'Reilly, 1999

[6]   Website of the Javastorms project: http://www.javastorms.org

[7]   Website of the Embedded Sensor Board ESB 430/2: http://www.scatterweb.com

[8]   Carsten Buschmann, Stefan Fischer, Norbert Luttenberger and Florian Reuter: Middleware for Swarm-Like Collections of Devices, in IEEE Pervasive Computing Magazine, Vol. 2, No. 4, 2003