
Preface

Overview

Sensing and planning are at the core of robot motion. Traditionally, mobile robots have been used for performing various tasks with a general-purpose processor on-board. This book grew out of our research enquiry into alternate architectures for sensor-based robot motion. It describes our research starting early 2002 with the objectives of obtaining a time, space and energy-efficient solution for processing sensor data for various robotic tasks.

New algorithms and architectures have been developed for exploration and other aspects of robot motion. The research has also resulted in design and fabrication of an FPGA-based mobile robot equipped with ultrasonic sensors. Numerous experiments with the FPGA-based mobile robot have also been performed and they confirm the efficacy of the alternate architecture.

Organization and Features

Chapter 1 presents the motivation for the work described in this book. Chapter 2 surveys prior work on sensors and processors on mobile robots. Prior work on exploration and navigation algorithms is also described and the algorithms are examined from a VLSI point of view. Chapter 3 presents the details of design and fabrication of the FPGA-based mobile robot used for experiments reported in this book. Chapter 4 presents a hardware-efficient grid-based exploration strategy for dynamic planar environments. Chapter 5 studies the problem of constructing landmarks, in a hardware-efficient manner, using data gathered by the mobile

robot during exploration. Chapter 6 presents a summary of the work described in this book and discusses extensions.

Programs have been developed in Verilog for exploration and other tasks. To give the reader a better feel for the approach, central modules in the Verilog code for exploration have been included in this book as an appendix. A list of suggested mini-projects is included as another appendix.

Audience

This book presents material that is appropriate for courses at the senior undergraduate level and graduate level in the areas of robotics and embedded systems. It is also suitable for researchers in the areas of VLSI-efficient algorithms and architectures. Practising engineers in the area of FPGA-based system design would also find the book to be immensely useful.

Basic familiarity with logic design and hardware description languages would be adequate to follow the material presented in this book.

Acknowledgments

The authors owe a word of thanks to many people who helped in various ways. The authors thank their families and friends for their support. Professor S. Srinivasan provided support and encouragement. Xilinx, Inc. and Digilent Inc. provided valuable support in the form of development boards for the experiments. Special thanks go to Dr. Thomas Ditzinger, Springer editor, for obtaining reviews for various chapters in this book. Thanks also to Ms. Heather King of Springer for assistance with the cover text. The authors thank the anonymous reviewers for their comments. The authors also acknowledge the support of Indian Institute of Technology Madras and Nanyang Technological University, Singapore.

*K. Sridharan
Panakala Rajesh Kumar*

Literature Survey

The survey is organized into sections covering prior work on various aspects presented in the subsequent chapters. Since a core component of the material presented in this book is description of an FPGA-based mobile robot developed for experiments, review of prior work on sensors and processors for mobile robots is first taken up. This is followed by a review of the literature on robotic exploration, landmark determination and navigation. Finally, we touch upon applications where FPGAs have been hitherto employed.

2.1 Sensors and Processors for Mobile Robots

Autonomous ground vehicles and mobile robots have been developed with different configurations and sensing capabilities since approximately the 1970s. One of the earliest wheeled robots was the SHAKEY [10]. The SHAKEY was based on a combination of on-board computers and a radio link to larger computers elsewhere. SHAKEY's primary sensor was a scanning camera. Around the early 1980's, the CMU Rover represented the state of the art in robotic mobility. The CMU Rover also had multiple general-purpose processors on-board and had provisions for connecting to a large remote computer.

While the robots described had sensing and planning capabilities, others which lacked elaborate planning capabilities were also simultaneously developed. Example are the Robart-I and Robart-II mobile robots [11]. Robart-I had a range of sensors, including ultrasonic sensors, infrared proximity detectors and tactile

feelers. Robart-II employed a microwave motion detector and a smoke detector besides infrared sensors. A good account of these robots is provided by [12].

Mobile robotic platforms with other types of processors are relatively recent. One example is the work reported in [13]. The authors in [13] report development of a mobile robot with two processors on-board: a TMS320F2812 DSP for implementing the control system for the robot and a general-purpose processor (as part of a laptop computer) for handling algorithms pertaining to navigation. Earlier work [14] has focussed on use of a digital signal processor for acquiring data from multiple ultrasonic sensors and processing. No experiment with an actual robot is reported in [14].

Many different sensors and instruments have been used on mobile robots. These include ultrasonic sensors, infrared sensors, vision sensors, tactile sensors, and encoders.

Ultrasonic sensors [15], [16] have been one of the simplest and widely used sensors for measuring distance to the nearest obstacle from a mobile robot. Design and development of circuits for distance measurement between a transmitter and receiver ultrasonic transducer have been extensively researched [17]. Approaches to correct various types of errors during distance measurement have also been developed [18]. Use of multiple receivers (along with one transmitter) has also been studied [19]. The exact point attaining minimum distance cannot be easily obtained using an ultrasonic sensor. This, however, is not a severe limitation for many applications in robotics.

Ultrasonic sensors have a range of approximately 3 metres and constitute the primary sensor on the mobile robot developed as a part of this book.

While sensors based on the ultrasonic, pulse-echo technique described in the previous paragraph have been the mainstay for collision avoidance systems in mobile robots, other alternatives that are not based on sound have also been examined by researchers [20]. In particular, where it is difficult to provide a medium for propagation of sound waves, light-based sensors have been explored. Some of these sensors are based on time-of-flight and triangulation techniques. In the mobile robot prototype developed

in this research, three infrared sensors are used for detection of objects within a very small distance (less than 20 cm) from the robot and to cut off power (via a relay) to the motors driving the robot wheels. It is worth noting that light-based sensors are typically used to determine the presence of a “target” (or an object) rather than to measure distance to them.

Studies on tactile sensors and whiskers have also been made by researchers in the context of path planning for mobile robots. One example is the work described in [21]. Work on vision sensors for mobile robots also exists [22]. The authors in [22] discuss various aspects of mobile robot navigation using vision sensors. The work described in this research does not explicitly use tactile or vision sensors.

2.2 Robotic Operation in Known and Unknown Environments

A brief review of the literature on robotic operation in environments where the geometry of the objects is known beforehand (or a map is available) is presented. We then move on to study prior work on exploration and navigation in unknown environments.

2.2.1 Environment with Prior Knowledge of Object Geometries and Locations

It is worth noting that prior exploration (combined perhaps with other methods such as probing) would yield information on location and shape of objects (that function typically as obstacles for the robot during the course of navigation) in the environment. An important task then is to find a collision-free path for the robot from some initial position to a predefined destination. Several algorithms have been proposed for navigation in an environment when the geometry of the objects is completely known. These are based on notions such as the *visibility graph* [2] and Voronoi diagrams [1]. Much work has been done in the context of the *piano mover’s problem* [1], [23]. The domain of computational geometry [3] has dealt with computational aspects of geometric problems dealing with polygons representing objects in the environment. While the solutions proposed for robotic navigation for different models of

objects in the environment are interesting from a theoretical point of view, it is often difficult to obtain accurate *a priori* information on shapes and positions of the objects. The research reported in this book assumes that the geometry of the objects is not available.

2.2.2 Unknown Environments

Since the core of this book is on various tasks in unknown environments, we discuss work in different directions for this case.

Exploration and Mapping

One of the early efforts on *exploration* in unknown environments is [24]. The authors in [24] develop a method that is particularly appropriate for exploration and mapping in large-scale spatial environments. The authors examine the problem of errors in calculation of position for a mobile robot based on its encoders. The authors present a topology-based approach that relies on sensor signatures rather than on coordinates of a reference point on the moving robot with respect to a global coordinate frame. Several aspects pertaining to the difficulty in building a map using topological information are brought out. The authors present simulation results for map building.

Another piece of work on exploration is reported in [25]. The authors in [25] consider the notion of an agent that locates itself on a graph by matching nodes and the adjacency relationships between them. The approach assumes that the agent can label each node by depositing a marker at the nodes. The work reported is of an algorithmic nature. Since the focus of the work reported in [25] is not on implementation aspects, it is not clear how easily the ideas can be realized in hardware.

Exploration of environment for building specific types of geometric structures has also been performed [26], [27]. In [26], the authors present an incremental approach to construct a Voronoi diagram-like structure. In particular, the authors consider the problem of learned navigation of a circular robot through a two-dimensional terrain whose model is not *a priori* known. The authors present algorithmic details for (a) the visit problem where the robot is required to visit a sequence of destination points and

(b) the terrain model acquisition problem where the robot is required to acquire the complete model of the terrain. Experimental results are not available.

Construction of hierarchical generalized Voronoi graph for an unknown environment in two and higher dimensions is presented in [27], [28]. The authors in [27] present a scheme based on prediction and correction for incremental construction of the hierarchical generalized Voronoi graph. Planar and three-dimensional simulations are presented. Some experiments on a mobile robot equipped with a general-purpose processor and sonar are briefly described.

Construction of generalized local Voronoi diagram using laser range scanner data has been studied in [29]. The scheme developed in [29] is based on clusterization of scan points based on some property. In particular, clusterizations involving interdistance of successive scan points and distance to the nearest neighbor have been studied. It is worth noting that the above scheme is for a static environment and simulations are only presented.

Construction of other geometric structures based on sensor data has also been researched. In [30], the authors study the construction of a *Local Tangent Graph (LTG)* using range sensor data. The local tangent graph is then used for navigation. In particular, the local tangent graph helps to choose the locally optimal direction for a robot as it moves towards the destination point. Also, the robot uses the LTG for following the boundary of obstacles. The authors present simulation results for their scheme.

An approach to exploration based on visiting landmarks is presented in [31]. In [31], the work is based on a robot that maintains a list of list of all unvisited landmarks in the environment.

Considerable work has also been done on planning robot motion strategies for efficient model construction [32]. Work in this direction has been on finding a function that reflects intuitively how the robot should explore the space so that we have a compromise between possible elimination of unexplored space and distance travelled.

While most of the work in the area of robotic exploration has concentrated on static environments, there has been some work on semi-dynamic and dynamic environments. Mapping semi-dynamic environments where objects (such as chairs, tables) move

periodically from one location to another has been studied in [33]. Some work on detecting changes in a dynamic environment for updating an existing map supported by computer simulations is described in [34]. Experiments with a mobile robot have been reported by some authors [35], [33].

More recently, an approach for planning exploration strategies for simultaneous localization and mapping in a *static environment* has been proposed [36]. The authors in [36] give a method to select the next robot position for exploration based on a novel utility function. The utility function defined in [36] combines geometric information with intensive usage of results obtained from perceptual algorithms. The outcome of the exploration is a multi-representational map made up of polygons, landmarks and a roadmap. Experiments with a real robot and simulations are presented in [36] but the focus is on using a general-purpose processor. Also, this work does not deal with *dynamic* environments.

Landmark Determination

The problem of selecting landmarks for path execution is addressed in [37]. In particular, the authors in [37] determine which landmarks a robot should detect and track at different parts of a given path so as to minimize the total cost of detecting and tracking landmarks. Experimental results on a mobile robot are not available. The authors present results of simulation based on implementations of a shortest path algorithm and a min-cost flow algorithm.

Localization of a free-navigating mobile robot called MACROBE is described in [38]. The authors in [38] use a 3D-laser range camera as the key sensor. Natural landmarks are extracted from the 3D laser range image and these are matched with landmarks predicted from an environmental model.

A method to learn dynamic environments for spatial navigation is proposed in [39]. The authors in [39] develop what are known as adaptive place networks and a relocalization technique. They then develop an integrated system known as ELDEN that combines a controller, an adaptive place network and the relocalization technique to provide a robust exploration and navigation method in a dynamic environment.

One of the common problems in landmark-based localization approaches is that some landmarks may not be visible or may be confused with other landmarks. The authors in [40] address this problem by using a set of good landmarks. In particular, the authors present an algorithm to learn a function that can select the appropriate subset of landmarks. The authors present empirical results to support their method. The authors in [41] develop a technique for on-line selection of stable visual landmarks under uncertainty of vision and motion.

A landmark-based approach for topological representation of indoor environments is proposed in [42]. Geometric features detected by exteroceptive sensors of a robotic vehicle are grouped into landmarks. A complete representation of the environment is constructed incrementally as the robot relocalizes along its trajectory. Some robotic experiments using a HILARE-2bis and a laser range finder are also presented.

A technique based on training a neural network to learn landmarks that optimize the localization uncertainty is presented in [43]. Techniques to optimally select landmarks for mobile robot localization by matching terrain maps are presented in [44].

An approach to construct and tracking abstract landmark chunks from a video sequence is presented in [45]. Potential landmarks are acquired in point mode but aggregations of them are utilized to represent interesting objects that can then be maintained throughout the path. The approach has been tested on a video sequence of 1200 frames.

Navigation in Unknown Environments

We have so far studied literature on the *exploration* problem. We now briefly discuss literature on the *navigation* problem. With regard to *navigation* in unknown environments, one of the early approaches to sensor-based mapping (for navigation) is based on the notion of *occupancy grids* [15]. It is done by imposing a grid on the workspace. Each of the grid cells can have one of the following values: occupied, empty and unknown. The method proposed in [15] is restricted to a static unknown environment.

The occupancy grid technique has been applied to collision avoidance [46] and to learn shape models of objects in unknown

environments [47]. In particular, [47] assumes a semi-dynamic environment where the objects are static over short time periods but move over longer periods (such as chairs moved in and out of a room). Some experiments with a *Pioneer* robot equipped with a general-purpose processor are described in [47]. Recently, variations to the occupancy grid approach have been proposed. One example is the *coverage map* defined in [48]. The coverage map is intended to overcome some of the difficulties encountered in the occupancy grid technique with regard to handling of walls and other obstacles.

Some work on navigation in busy dynamic environments also exists. The authors in [49] present the development of a robotic wheelchair and experiments in autonomous navigation in a railway station during busy periods. The authors employ a SICK two-dimensional laser range finder PLS200 and indicate that the real-time application calls for use of this expensive sensor. The authors employ a personal computer in their setup.

Hardware-directed Schemes

Since this book concerns digital hardware-friendly algorithms and FPGA implementations, we focus here only on reviewing prior digital hardware-based schemes for robotic planning and related problems. In general, not much is known on hardware-directed algorithms for robotic exploration and navigation.

A parallel algorithm and architecture for robot path planning are presented in [50]. The algorithm is based on wave propagations and the algorithm is mapped on to a systolic array architecture with N processing elements assuming the work area (for the robot) is divided into $N \times N$ discrete regions. The authors in [50] also indicate development of a prototype VLSI chip with five processors using 2-micron CMOS technology.

A parallel algorithm for collision-free path planning of a diamond-shaped robot among arbitrarily shaped obstacles, which are represented as a discrete image is presented in [51]. A cellular architecture and its implementation in VLSI using a 1.2 μm double-layer metal CMOS technology are also presented. No experimental results on an actual robot are available.

There exists some prior work on construction of certain geometric structures in hardware. The authors in [51] have proposed a cellular architecture for constructing a discrete Voronoi diagram on a binary image based on a d_4 metric (the d_4 distance between two pixels at coordinates (x_1, y_1) and (x_2, y_2) in an image is given by $\|x_1 - x_2\| + \|y_1 - y_2\|$). An extension of this work to construct the discrete Voronoi diagram based on the Euclidean distance metric in a binary image is presented in [52]. The approach in [52] is based on dilating each object iteratively and establishing connectivity between neighboring pixels belonging to the same dilated object at each iteration.

The authors in [53] consider the computation of generalized Voronoi diagrams on polygon rasterization hardware. The Voronoi diagram in [53] is obtained via a polygonal mesh approximation to the distance function of a Voronoi site over a planar rectangular grid of point samples. The design in [53] is focussed on a single processor solution and is supported by C++ simulations. More recently, the work in [53] has been extended in [54] for real-time path planning for virtual agents in dynamic environments. However, the focus of the work in [54] is on a PC-based solution with a special graphics card. The objective in [54] is to primarily obtain high speed and their focus is not on energy or space-efficiency.

It is also worth noting that in all of the above articles, there is no explicit sensor-based Voronoi construction. Other structures which have been explored from a hardware perspective include visibility graphs [55], [56], [57] and tangent graphs [58]. The authors in [56] present the notion of a virtual rectangle for path calculations. Work on definition of T -vectors in the context of mobile robot motion planning is presented in [55]. The approach presented in [55] is in the context of construction of reduced visibility graphs and is digital-hardware friendly. It is based on bit operations involving vertices of objects. However, the implementation of the approach in [55] is directed to a general-purpose processor.

In [57], the authors present a hardware-efficient scheme for construction of the complete visibility graph for (i) an environment with purely convex polygonal objects and (ii) an environment with non-convex polygonal objects. Results of FPGA-based implementation are also available.

In [58], the authors present a VLSI-efficient scheme for construction of the tangent graph. Results of synthesis using Synopsys Design Compiler 2001.08-SP1 with Compass Passport CB35OS142 standard cell library in 0.35 μm CMOS process are also presented.

In [55], [56], [57] and [58], the environment is assumed to be known.

2.3 FPGA-based Design

To the best of our knowledge, there is no prior work on FPGAs for robotic exploration in dynamic environments. This section is primarily intended to give a feel for diversity of applications of FPGAs. The survey is not exhaustive. We have just attempted to cover some representative work.

FPGAs have been identified for use in space robotics [59] in view of the ability to get high performance solutions on a small device that may otherwise take up an entire board of parts.

A high-performance FPGA-based implementation of Kalman filter is reported in [60]. This is valuable in the context of probabilistic robotics with the focus on hardware-directed solutions. Codesign systems and coprocessors using FPGAs have been developed and employed in many applications in the last decade [61], [62], [63].

The authors in [64] study FPGA implementation of a fuzzy wall-following control scheme for a mobile robot. Work on handling other types of sensors on a mobile robot by an FPGA has also been reported recently [65]. The authors in [65] report an FPGA-based scheme for colour image classification for mobile robot navigation. No experiments on an actual report, however, are available in [65].

A proposal for a dynamic self-reconfigurable robotic navigation system is presented in [66]. The author in [66] presents an embedded system platform which is based on the partial reconfiguration capabilities of Xilinx Virtex-II Pro FPGA. The reconfiguration is suggested as a valuable feature in the context of image processing required on the mobile robot (in particular, to allow image loading and image matching to execute at the same time). The author presents some results pertaining to usage of the FPGA. Experimental results on navigation are not available.

An approach and architecture for dynamic reconfiguration on autonomous mini-robots is presented in [67]. The authors in [67] consider two application examples, one dealing with reconfigurable digital controllers and the second dealing with image processing to discuss their ideas. Two robotic platforms are also described.

An approach to incorporate partial reconfiguration in an environment with multiple mobile robots is presented in [68]. The authors in [68] present a case study involving autonomous fault handling and reconfiguration to show the behavior reconfiguration in a team of two robots (denoted by R_1 and R_2). In particular, R_1 and R_2 perform wall following initially. In the middle of this experiment, R_1 's IR sensor fails so it notifies R_2 . R_2 , upon obtaining a resource list from R_1 , formulates that Leader/Follower behaviour is appropriate and configures R_1 to follower behavior. The authors present snaps of actual point where reconfiguration takes place. They also present results of FPGA usage.

2.4 Summary

In this chapter, we have discussed the literature on sensors and processors used on mobile robots. The literature on exploration and navigation in unknown environments has then been discussed. Finally, applications of FPGAs have been mentioned.