

A Sonar-Based Mapping and Navigation System

Alberto Elfes

The Robotics Institute
Carnegie-Mellon University
Pittsburgh, PA 15213

Abstract

This paper describes a sonar-based mapping and navigation system for autonomous mobile robots operating in unknown and unstructured surroundings. The system uses sonar range data to build a multi-levelled description of the robot's environment. Sonar maps are represented in the system along several dimensions: the Abstraction axis, the Geographical axis, and the Resolution axis. Various kinds of problem-solving activities can be performed and different levels of performance can be achieved by operating with these multiple representations of maps. The major modules of the Dolphin system are described and related to the various mapping representations used. Results from actual runs are presented, and further research is mentioned. The system is also situated within the wider context of developing an advanced software architecture for autonomous mobile robots.

1. Introduction

The Dolphin system is intended to provide sonar-based mapping and navigation for an autonomous mobile robot operating in unknown and unstructured environments. The system is completely autonomous in the sense that it has no *a priori* model or knowledge of its surroundings and also carries no user-provided map. It acquires data from the real world through a set of sonar sensors and uses the interpreted data to build a multi-levelled and multi-faceted description of the robot's operating environment. This description is used to plan safe paths and navigate the vehicle towards a given goal.

The system is intended for indoor as well as outdoor use; it may be coupled to other systems, such as vision, to locate landmarks that would serve as intermediate or final destinations.

In the course of this paper, we will briefly identify some of the conceptual processing levels needed for mobile robot software, relate the present system to this framework, discuss the multiple representations developed for sonar maps as well as their use in different kinds of problem-solving activities, describe the overall system architecture and show some results from actual runs. We finish with an outline of further research.

2. Conceptual Processing Levels for an Autonomous Mobile Robot

The sonar mapping and navigation system discussed here is part of a research effort that investigates various issues involved in the development of the software structure of an autonomous mobile robot. To situate the Dolphin system within this wider context, we characterize

in this section some of the conceptual processing levels required for an autonomous vehicle (see Fig. 2-1). Each is briefly discussed below:

- **Robot Control:** This level takes care of the physical control of the different sensors and actuators available to the robot. It provides a set of primitives for locomotion, actuator and sensor control, data acquisition, etc., that serve as the robot interface, freeing the higher levels of the system from low-level details. This would include dead-reckoning motion estimation and monitoring of internal sensors. *Internal Sensors* provide information on the status of the different physical subsystems of the robot, while *External Sensors* are used to acquire data from the robot's environment.
- **Sensor Interpretation:** On this level the acquisition of sensor data and its interpretation by Sensor Modules is done. Each Sensor Module is specialized in one type of sensor or even in extracting a specific kind of information from the sensor data. They provide information to the higher levels using a common representation and a common frame of reference.
- **Sensor Integration:** Due to the intrinsic limitations of any sensory device, it is essential to integrate information coming from qualitatively different sensors. Specific assertions provided by the Sensor Modules are correlated to each other on this level. For example, geometric boundaries of an obstacle extracted by sonar can be projected onto an image provided by the vision subsystem and can help in identifying a certain object. On this level, information is aggregated and assertions about specific portions of the environment can be made.
- **Real-World Modelling:** To achieve any substantial degree of autonomy, a robot system must have an understanding of its surroundings, by acquiring and manipulating a rich model of its environment of operation. This model is based on assertions integrated from the various sensors, and reflects the data acquired and the hypotheses proposed so far. On this level, local pieces of information are used in the incremental construction of a coherent global Real-World Model; this Model can then be used for several other activities, such as landmark recognition, matching of newly acquired information against already stored maps, and generation of expectancies and goals.
- **Navigation:** For autonomous locomotion, a variety of problem-solving activities are necessary, such as short-term and long-term path-planning, obstacle-avoidance, detection of emergencies, etc. These different activities are performed

VII. Global Control

VI. Global Planning

V. Navigation

IV. Real-World Modelling

III. Sensor Integration

II. Sensor Interpretation

I. Robot Control

Figure 2-1: Conceptual Activity Levels in a Mobile Robot Software Architecture.

by modules that provide specific services.

- *Global Planning*: To achieve a global goal proposed to the robot, this level provides task-level planning for autonomous generation of sequences of actuator, sensor and processing actions. Other necessary activities include simulation, error detection, diagnosis and recovery, and replanning in the case of unexpected situations or failures.
- *Global Control*: Finally, on this level Supervisory Modules are responsible for the scheduling of different activities and for combining Plan-driven with Data-driven activities in an integrated manner so as to achieve coherent behaviour.

This conceptual structure provides a paradigm within which several of our research efforts are situated [6, 11, 12]. It has influenced, in particular, the architecture of the *Dolphin* system for sonar-based mapping and navigation, as mentioned in Section 5.

3. Sonar Mapping

3.1. Introduction

The *Dolphin* sonar system is able to build dense maps of the robot's environment and use them for autonomous navigation. The central representation of sonar mapping information is the *Probabilistic* or *Sensor-Level Local Map*, which uses a medium-resolution grid (with a typical accuracy of 0.5 ft). The cells of a two-dimensional array spanning the area of interest are used to store occupancy information (EMPTY, OCCUPIED or UNKNOWN), as well as the associated confidence factors.

Currently, the cycle of operation of the sonar system is as follows: from its current position, the robot acquires a set of range measurements provided by the sonar sensor array; these readings are then interpreted as assertions concerning *empty* and *occupied* areas, and serve to update the sonar map. The map is now used to plan a safe path around obstacles, and the robot moves a certain distance along the path. It updates its position and orientation estimate and repeats the cycle.

3.2. Building Maps

The Local Map building process is discussed in detail in [11], and is reviewed here only briefly. We proceed to describe how other representations are derived from it.

The sonar sensor array is composed of 24 Polaroid laboratory grade ultrasonic transducers. These devices are arranged in a ring and controlled by a microprocessor that also interfaces to a VAX mainframe. For experimental runs, the array was mounted on two different robots (*Neptune* [13] for indoor runs, and the *Terragator* [12] for outdoors).

The mapping system processes range measurements obtained from the sonar transducers, annotated with the positions of the corresponding sensors, which are derived from the position and orientation of the robot. Each measurement provides information about *probably empty* and *possibly occupied* volumes in the space subtended by the beam (a 30° cone for the present sensors). This occupancy information is projected onto a rasterized two-dimensional horizontal map. Sets of readings taken both from different sensors and from different positions of the robot are incrementally integrated into the sonar map, using a probabilistic approach. In this way, errors and uncertainties are reduced and the map becomes gradually more detailed.

The sonar beam is modelled by probability distribution functions. Informally, these functions describe our confidence that the points inside the cone of the beam are empty and our uncertainty about the location of the point that caused the echo. The functions are based on the range value and on the spatial sensitivity pattern of the sonar device.

These sonar maps are very useful for motion planning. They are much denser than those made by typical stereo vision programs, and computationally at least one order of magnitude faster to produce.

3.3. Related Work

In the Robotics area, ultrasonic range transducers have recently attracted increasing attention. This is due in part to their simplicity, low cost and the fact that distance measurements are provided directly. Some research has focused specifically on the development of more elaborate beam-forming and detection devices (see, for example, [8]), or on the application of highly sophisticated signal processing techniques [1] to complex sonar signals.

Specific applications of sonar sensors in robot navigation include determining the position of a robot given a known map of the environment [9, 10, 5] and some *ad hoc* navigation schemes [2]. An independent CMU sonar mapping and navigation effort [3, 4] uses a narrow beam, formed by a parabolic reflector, to build a line-based description of the environment.

4. Multiple Axis of Representation of Sonar Mapping Information

From the Probabilistic Local Maps described in the previous section, several other data structures are derived. We use the following dimensions of representation (Fig. 4-1):

- THE ABSTRACTION AXIS: Along this axis we move from a sensor-based, data-intensive representation to increasingly higher levels of interpretation and abstraction. Three levels are defined: the *Sensor Level*, the *Geometric Level* and the *Symbolic Level*.
- THE GEOGRAPHICAL AXIS: Along this axis we define *Views*, *Local Maps* and *Global Maps*, depending on the extent and characteristics of the area covered.
- THE RESOLUTION AXIS: Sonar Maps are generated at different values of grid resolution for different applications. Some computations can be performed satisfactorily at low levels of detail, while others need higher or even multiple degrees of resolution.




Figure 4-1: Multiple Axis of Representation of Sonar Maps.

4.1. The Abstraction Axis

The first kind of sonar map built from the sonar range data uses the *Probabilistic* representation described earlier. A two-dimensional grid covering a limited area of interest is used. This map is derived directly from the interpretation of the sensor readings and is, in a sense, the description closest to the real world. It serves as the basis from which other kinds of representations are derived. Along the Abstraction Axis, this data-intensive description is also defined as the *Sensor Level Map*.

The next level is called the *Geometric Level*. It is built by scanning the Sensor Level Map and identifying blobs of cells with high OCCUPIED confidence factors. These are merged into uniquely labeled objects with explicitly represented polygonal boundaries. If needed, the same can be done with EMPTY areas.

The third is the *Symbolic Level*, where maps of larger areas (typically Global Maps) are described using a graph-like representation. This description bears only a topological equivalence to the real world. Nodes represent "interesting" areas, where more detailed mapping information is necessary or available, while edges correspond to simpler or "uninteresting" areas (navigationally speaking), such as corridors.

Different kinds of problem-solving activities are better performed on different levels of abstraction. For example, global path-planning (such as how to get from one building wing to another) would be done on the symbolic level, while navigation through a specific office or lab uses the sensor-level map, where all the detailed information about objects and free space, as well as the associated certainty factors, is stored.

4.2. The Geographical Axis

In order to be able to focus on specific geographical areas and to handle portions of as well as complete maps, we define a hierarchy of maps with increasing degrees of coverage. Progressing along the Geographical Axis, we start with *Views*, which are maps generated from scans taken from the current position, and that describe the area visible to the robot from that place. As the vehicle moves, several *Views* are acquired and integrated into a *Local Map*. The latter corresponds to physically delimited spaces such as labs or offices, which define a connected region of visibility. *Global Maps* are sets of several *Local Maps*, and cover wider spaces such as a whole wing of a building, with labs, offices, open areas, corridors, etc.

4.3. The Resolution Axis

Finally, along the Resolution Axis, we again start with the Sensor-Level Local Map and generate a progression of maps with increasingly less detail. This allows certain kinds of computations to be performed either at lower levels of resolution with correspondingly less computational expense, or else enables operations at coarser levels to guide the problem-solving activities at finer levels of resolution.

The most detailed sonar maps that can be obtained from the method outlined in Section 3 (considering the intrinsic limitations of the sensors) have a cell size of 0.1×0.1 ft. For navigation purposes, we have typically been using a 0.5 ft grid for indoors and a 1.0 ft grid for outdoors. Nevertheless, several operations on the maps are expensive and are done more quickly at even lower levels of resolution. For these cases we reduce higher resolution maps by an averaging process that produces a coarser description. One example of an application of this technique is the Map Matching procedure described in [11], where two Local Maps being compared with each other are first matched at a low level of detail. The result then constrains the search for a match at the next higher level of resolution.

5. Overall System Architecture

To provide a context for these multiple descriptions, we present in this Section the overall architecture of the **Dolphin Sonar-Based Mapping and Navigation** system (Fig. 5-1). The function of the major modules and their interaction with the various sonar map representations[7] is described below:

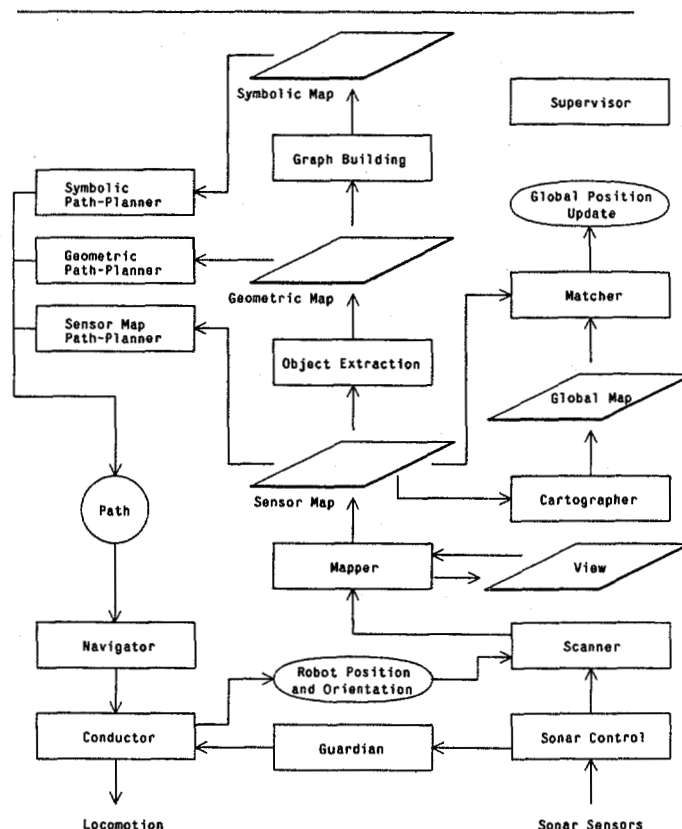


Figure 5-1: Architecture of the Sonar Mapping and Navigation System.

Sonar Control: Interfaces to and runs the sonar sensor array, providing range readings.

Scanner: Preprocesses and filters the sonar data. Annotates it with the position and orientation of the corresponding sensor, based on the robot's motion estimate.

Mapper: Using the information provided by the Scanner, generates a View obtained from the current position of the robot. This View is then integrated into a Local Map.

Cartographer: Aggregates sets of Local Maps into Global Maps. Provides map handling and bookkeeping functions.

Matcher: Matches a newly acquired Local Map against portions of Global Maps for operations such as landmark identification or update of the absolute position estimate.

Object Extraction: Obtains geometric information about obstacles. Objects are extracted by merging blobs of OCCUPIED cells and determining the corresponding polygonal boundaries. A region-coloring approach is used for unique labeling.

Graph Building: Searches for larger regions that are either empty or else have complex patterns of obstacles, labeling them as "free" or "interesting" spaces.

Path-Planning: Three levels of path-planning are possible: *Symbolic Path-Planning* is done over wider areas (Global Maps) and at a higher level of abstraction (Symbolic Maps); *Geometric Path-Planning* is done as an intermediary stage, when the uncertainty in Local Maps is low; and *Sensor Map Path-Planning* is used to generate detailed safe paths. The latter performs an A* search over the map cells, with the cost function taking into account the obstacle certainty factors and the distance to the goal. The planned path is provided to the Navigator.

Navigator: Takes care of the overall navigation issues for the vehicle. This includes examining already planned paths to determine whether they are still usable, invoking the path-planner to provide new paths, setting intermediary goals, overseeing the actual locomotion, etc.

Conductor: Controls the physical locomotion of the robot along the planned path. The latter is currently approximated by sequences of line segments, using a line-fitting approach. Provides an estimate of the new position and orientation of the robot.

Guardian: During actual locomotion, this module checks the incoming sonar readings and signals a stop if the robot is coming too close to a (possibly moving) obstacle not detected previously. It serves as a "sonar bumper".

Supervisor: Oversees the operation of the various modules and takes care of the overall control of the system. It also provides a user interface.

Comparing this architecture with the activities outlined in Section 2, we see that the Sonar Control and Conductor modules belong to the Robot Control level; the Scanning and Mapping modules operate on the Sensor Interpretation level; the Object Extraction, Graph Building, Cartographer and Matcher modules provide functions on the Real-World Modelling level; Path-Planning, the Guardian and Navigation are situated on the Navigation level; and the Supervisor belongs to the Control level.

6. Tests of the System

The *Dolphin* system described here was tested in several indoor runs in cluttered environments using the *Neptune* mobile robot [13], developed at the Mobile Robot Laboratory of the Robotics Institute, CMU. It was

also tested in outdoor environments, operating among trees, using the *Terragator* robot in the context of the CMU ALV project. The system operated successfully in both kinds of environments, navigating the robot towards a given destination.

In Fig. 6-1, an example run is given. The sequence of maps presented shows how the sonar map becomes gradually more detailed and how the path is improved, as more information is gathered. The example corresponds to an indoor run, done in our laboratory. A distance of approximately 25 ft was covered; the grid size is 0.5 ft. Objects present in the lab included chairs, tables, boxes, workstations, filing cabinets, etc. Empty spaces with high certainty factors are represented by white areas; lower certainty factors by "." symbols of increasing thickness. Occupied areas are shown using "x" symbols, and Unknown areas using "-". The planned path is shown as a dotted line, and the route actually followed by the robot as solid line segments. The starting point is a solid + and the goal a solid x.

In Fig. 6-2, an outdoor run is shown, together with an example of the Object Extraction algorithm. The objects are uniquely identified and the polygonal boundaries are shown. The map corresponds to a run done among trees. A distance of approximately 50 ft was traversed. The grid size was 1.0 ft, which proved adequate for navigation, but did not allow a more precise description of the real boundaries of the detected objects.

7. Further Research

We conclude our discussion by outlining in this Section some research lines to be further pursued.

7.1. Handling Position Uncertainty


Our current system presupposes that the position and orientation of the robot (and by that, of the sonar sensors) as it acquires sonar data is known with reasonable precision. This is crucial for integrating readings taken over shorter distances, which are combined as previously outlined. Drifts over longer distances are inevitable, but lead only to a topological distortion of the map.

To update the current position of the robot, we presently rely on dead-reckoning estimates based on wheel encoders and an onboard inertial navigation system. These drift with travelling time and distance. As a result, ground truth (the real-world environment) and the sonar map drift apart. This problem is characteristic of navigation without access to absolute position information. In stereo vision navigation, it has traditionally been addressed by estimating motion based on image matching.


We are currently investigating two complementary approaches to this problem: incorporating the uncertainty in the position of the robot into the map-making process and do motion solving by matching new sets of readings against the map being incrementally built.

7.2. Extending the Architecture


The architecture described above embodies a sequential control-flow organization. This, however, does not reflect the problem-solving characteristics inherent to mobile robot software. The various modules involved in the problem-solving effort are frequently quasi-independent and have a low degree of coupling; therefore, they should conceptually proceed in parallel, interacting with each other as needed. We have recently started the implementation of a distributed version of *Dolphin* [12] along the lines discussed in [6], where multiple agents work on concurrent activities.




(a)




(d)



(b)



(e)



(c)

Figure 6-1: An Example Run. This run was performed indoors, in the Mobile Robot Lab. Distances are in ft. Grid size is 0.5 ft.




Figure 6-2: Objects Extracted from a Sonar Map. The objects are numbered and their polygonal boundaries are shown. This map describes an outdoor run, and the objects are trees. Distances are in ft. Grid size is 1.0 ft.

Another issue we are currently investigating is the development of a task-level Global Planner that would automatically generate a Control Plan, establishing sequences of parallel and sequential actions. We are considering a hierarchical approach similar to NOAH [14], using a graph to represent the plan and explicitly storing alternatives and sensor-dependent conditions as part of it. The elementary operations of sensor information gathering, interpretation, actuator control and specific problem-solving activities are the primitives used by the planner.

8. Conclusions

We have described a system that uses a Sensor Level, probability-based sonar map representation of medium resolution to build several kinds of maps. Three different dimensions of representation are defined: the Abstraction Axis, the Geographical Axis and the Resolution Axis. These maps are used by a sonar mapping and navigation system that performed successfully in indoor and outdoor environments. We are now investigating motion recovery techniques and expanding the system to test distributed control and global planning mechanisms.

9. Acknowledgments

The author would like to thank Hans P. Moravec for his support and for providing several important insights into the topics discussed in this paper, Gregg W. Podnar for his help with the Neptune robot, Chuck Thorpe for the initial code for the path-planner, and Richard Redpath for providing more reliable communication to the sonar ring, as well as for assistance during outdoor runs.

This research is being supported in part by the Office of Naval Research under Contract N00014-81-K-0503, in part by Denning Mobile Robotics, Inc., and in part by the Western Pennsylvania Advanced Technology Center. The author is supported in part by the Conselho Nacional de Desenvolvimento Científico e Tecnológico - CNPq, Brazil, under Grant 200.986-80, in part by the Instituto Tecnológico de Aeronáutica - ITA, Brazil, and in part by The Robotics Institute, Carnegie-Mellon University.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the funding agencies.

References

- [1] Baggeeroc, A.B.
Sonar Signal Processing.
Signal Processing Series. Applications of Digital Signal Processing.
Prentice-Hall, Englewood Cliffs, N.J., 1978.
- [2] Chattergy, A.
Some Heuristics for the Navigation of a Robot.
Technical Report, Robotics Research Laboratory, Department of Electrical Engineering, University of Hawaii, 1984.
- [3] Crowley, J.L.
Position Estimation for an Intelligent Mobile Robot.
In *1983 Annual Research Review*. The Robotics Institute, Carnegie-Mellon University, Pittsburgh, PA, 1984.
- [4] Crowley, J.L.
Dynamic World Modelling for an Intelligent Mobile Robot Using a Rotating Ultra-Sonic Ranging Device.
In *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*. IEEE, St. Louis, Missouri, March, 1985.
- [5] Drumheller, M.
Mobile Robot Localization Using Sonar.
Technical Report AI-M-826, Artificial Intelligence Lab, Massachusetts Institute of Technology, January, 1985.
- [6] Elfes, A. and Talukdar, S.N.
A Distributed Control System for the CMU Rover.
In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence - IJCAI-83*. IJCAI, Karlsruhe, Germany, August, 1983.
- [7] Elfes, A.
Multiple Levels of Representation and Problem-Solving Using Maps From Sonar Data.
In Weisbin, C.R. (editor), *Proceedings of the DOE/CESAR Workshop on Planning and Sensing for Autonomous Navigation*. Oak Ridge National Laboratory, UCLA, Los Angeles, August 18-19, 1985.
Invited Paper.
- [8] Miller, G.L., Boie, R.A. and Sibilia, M.J.
Active Damping of Ultrasonic Transducers for Robotic Applications.
In *Proceedings of the International Conference on Robotics*. IEEE, Atlanta, Georgia, March, 1984.
- [9] Miller, D.
Two Dimensional Mobile Robot Positioning Using Onboard Sonar.
In *Pecora IX Remote Sensing Symposium Proceedings*. IEEE, Sioux Falls, SD, October, 1984.
- [10] Miller, D.
A Spatial Representation System for Mobile Robots.
In *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*. IEEE, St. Louis, Missouri, March, 1985.
- [11] Moravec, H.P. and Elfes, A.
High-Resolution Maps from Wide-Angle Sonar.
In *IEEE International Conference on Robotics and Automation*. IEEE, March, 1985.
- [12] Moravec, H.P. et al.
Towards Autonomous Vehicles.
1985 Robotics Research Review.
The Robotics Institute, Carnegie-Mellon University, Pittsburgh, PA, 1985.
- [13] Podnar, G.W., Blackwell, M.K. and Dowling, K.
A Functional Vehicle for Autonomous Mobile Robot Research.
Technical Report CMU-RI-TR-84-28, The Robotics Institute, Carnegie-Mellon University, April, 1984.
- [14] Sacerdoti, E.D.
A Structure for Plans and Behavior.
Elsevier, New York, N.Y., 1977.