DISTRIBUTED HETEROGENEOUS SENSING FOR OUTDOOR MULTI-ROBOT LOCALIZATION, MAPPING, AND PATH PLANNING

Lynne E. Parker, Kingsley Fregene*, Yi Guo, and Raj Madhavan
Computer Science and Mathematics Division
Oak Ridge National Laboratory, P.O. Box 2008, Oak Ridge, TN 37831-6355 †
parkerle@ornl.gov,guoy@ornl.gov,kocfrege@ieee.org,raj_madhavan@yahoo.com

Abstract Our objective is to develop a team of autonomous mobile robots that are able to operate in previously unfamiliar outdoor environments. In these environments, the robot teams should be able to cooperatively localize even when DGPS is not consistently available, to autonomously generate rough elevation maps of their terrain, and to use these generated maps to plan multi-robot paths that enable them to accomplish their mission objective, such as reconnaissance and surveillance or perimeter security. This paper briefly outlines our approaches to achieving this objective, along with some of our implementation results on our team of four ATRV-mini mobile robots.

Keywords: Distributed sensing, heterogeneous robots, outdoor navigation.

1. Introduction

In practical applications of teams of mobile robots in outdoor terrains, a serious consideration is the navigation of the robots across previously unfamiliar terrain. For nearly all applications, these robots must be

*Research performed while author visited ORNL. Author's current address is University of Waterloo, Ontario, Canada.

†This research is sponsored in part by the Engineering Research Program of the Office of Basic Energy Sciences, U. S. Department of Energy. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U. S. Government purposes. Oak Ridge National Laboratory is managed by UT-Battelle, LLC for the U.S. Dept. of Energy under contract DE-AC05-00OR22725.

able to move safely to avoid navigation hazards. However, for many applications *safe* navigation alone is not sufficient; the robots are also required to find *efficient* paths through their terrain based upon their mission requirements. These robots may need to operate for a period of time in an outdoor area, and may need to develop knowledge about the outdoor terrain. For example, reconnaissance and surveillance tasks may require the robots to set up security patrols based upon terrain visibility. Exploration tasks may require robots to search an outdoor area for an object or feature of interest. Military applications may require a robot team to move from point to point within given boundaries along the most efficient route possible.

All of these practical applications require the robot teams to be able to 1) localize within the outdoor environment, 2) map their terrain sufficiently to enable efficient path planning, and 3) plan their paths according to the mission goals. A significant amount of research has addressed these individual and related problems, including localization, mapping for indoor planar environments, cross-country and road-following navigation, following trajectories roughly specified by a human operator, and path planning for both indoor and outdoor environments given a terrain map. In particular, the cooperative localization and mapping issue has been very extensively studied. However, most of this prior research has addressed the indoor environment. Very little prior research has addressed the complete problem of developing approaches that enable a team of robots to be immediately placed in a previously unfamiliar outdoor environment, to generate sufficient knowledge of the terrain for safe and efficient navigation, and to derive efficient multi-robot path plans.

A key challenge in this research is enabling the robot team to autonomously develop a terrain map of their outdoor working environment. Commonly available Digital Elevation Maps (DEMs) are not provided at the terrain resolution needed for safe and efficient robot navigation. The motion of a robot across a terrain using the Differential Global Positioning System (DGPS) to make position and elevation measurements will not operate in environments that include trees, buildings, steep hills, and so forth that generate a multi-pathing problem for DGPS. Even if continual DGPS could be guaranteed, robots would still need an additional mechanism for recognizing obstacles and unsafe navigation regions that should not be entered. In many applications, human operators could mark unsafe regions on an image roughly correlated to the DGPS positions, but this approach does not address the need to have a map model with sufficient elevation detail to enable efficient, repeated navigation across the working area.

An ideal solution would be to automate the challenging aspects of this problem so that the robot team can indeed be placed in a new, outdoor environment and operate successfully according to the mission requirements. Our research is aimed at developing the algorithms and the overall system that will enable this type of application to be solved with robot teams. Our approach takes advantage of the heterogeneous distributed sensing capabilities afforded by a team of multiple robots. Robots should be able to assist each other as needed to provide collaborative sensing capabilities that enable them to accomplish their mission.

The problem statement for the robot team that we are addressing is as follows: given an unknown outdoor environment with incomplete DGPS availability and unsafe navigation regions, develop an elevation map of the terrain marked with the unnavigable areas and use this map to plan multi-robot paths that satisfy the mission objectives (such as patrol paths). For the purposes of this research, the unsafe navigation regions are considered to be positive obstacles (e.g., trees, large rocks, etc.) and areas whose slope or local roughness exceeds a pre-specified limit. For now, we are not addressing the recognition of negative obstacles (e.g., holes in the ground or pools of water) or hidden obstacles (e.g., in grassy areas), since much recent work is addressing this issue and it is expected that these approaches can easily be inserted into our system.

The organization of this paper is as follows: Section 2 gives an overview of the experimental setup. The next three sections then outline the approaches to the three key issues in this research – multi-robot localization in Section 3, multi-robot mapping in Section 4, and multi-robot path planning in Section 5. Examples of the results of our implementation to date are given in Section 6. We conclude with summary remarks in Section 7.

2. Robot Team and Experimental Setup

The experimental platform (see Figure refsensorsuite) is a team of four ATRV-Mini wheeled mobile robots with 4-wheel differential-drive skid-steering. The experimental setup consists of a wireless mini-LAN, a Local Area DGPS (LADGPS), a software platform (*Mobility* from RWI) and codes developed in-house under Linux to read and log the data for the sensors on each robot. The wireless LAN is set up outdoors between an Operator Console Unit (OCU) and the robots. The OCU consists of a rugged notebook equipped with a BreezeCOM access point and antennas. Each robot has a BreezeCOM station adapter and an antenna. The LADGPS is formed by the base station/antenna hardware connected to the OCU and remote stations/antennas directly mounted

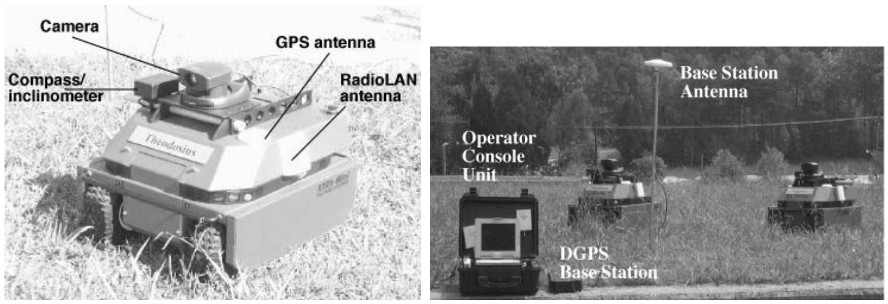

Figure 1. The ATRV-Mini sensor suite and experimental setup. The sensor suite consists of encoders, DGPS, a compass and a PTZ camera. The experimental setup depicted in the second photo consists of an operator console unit, a DGPS base station and a base station antenna. See text for further details.

on each robot. Each robot's station receives differential corrections from the base station such that LADGPS accuracy of up to 10 centimeters is obtainable. The distributed CORBA-based interface offered by *Mobility* ensures that querying the sensor slots of particular robots is done in a transparent decentralized manner by simply appending the robot's ID to all such queries.

The sensor suite is comprised of encoders that measure the wheel speeds and heading, DGPS, and a magnetic compass. Two of the robots are equipped with a pan-tilt-zoom (PTZ) capable camera for visual perception and the remaining two robots are equipped with a SICK scanning laser rangefinder.

3. Distributed EKF Localization

In outdoor environments, errors introduced due to distance traveled can be significant and unpredictable. This is a direct consequence of the undulatory nature of the terrain of travel and the uncertainties introduced into sensor data. These challenges make it comparatively difficult to realize successful navigation in unstructured outdoor environments. Motivated by these factors, our approach is an Extended Kalman Filter (EKF) based multi-robot heterogeneous localization framework similar to that developed in Roumeliotis and Bekey, 2000, but differing in the following ways: 1) the kinematic model of the robots is nonlinear, 2) no absolute positioning system capable of providing relative pose information is assumed to be available, and 3) the robots traverse on uneven and unstructured outdoor terrain. In the first case, a kinematic model that sufficiently captures the nonlinear vehicle motion is key to efficient use of sensor data and is central to successful autonomous navigation. A nonholonomic robot with a nonlinear kinematic model performs significantly better as the model efficiently captures the maneuvers of the robot. In the second case, even though we consider systems including DGPS, it only provides absolute position information for a single robot subject to the number of satellites in view at any given time. DGPS is not guaranteed to be continually available.

When some robots of the team do not have absolute positioning capabilities or when the quality of the observations from the absolute positioning sensors deteriorate, another robot in the team with better positioning capability can assist in the localization of the robots whose sensors have deteriorated or failed. In such cases, if relative pose information is obtained, casting the EKF-based localization algorithm in a form such that the update stage of the EKF utilizes this relative pose thereby provides reliable pose estimates for all the members of the team. Under this approach, at least one of the robots must maintain global positioning. (In future work, we will be examining how to maintain this constraint as the robots perform their primary mission.) We obtain relative pose information through one of two ways – a scanning laser range finder-based, and a vision-based cooperative localization approach.

In the case of cooperative localization via laser, consider the two-robot cases where robot #2 has a scanning laser range finder. The localization process proceeds as follows. First, robot #2 identifies robot #1 and acquires a range and bearing laser scan. Then, after the necessary preprocessing to discard readings that are greater than a predefined threshold, the range and bearing to the minima identified in the laser profile of robot #1 are determined. Finally, from the range and bearing pertaining to the minima, the pose of robot #2 is inferred and relative pose information is available for use.

In the case of vision-based cooperative localization, the robot's camera is used to provide relative position information. In the case where two robots are performing cooperative localization with the camera-equipped robot #1 lacking in absolute positioning capability, relative position information is obtained as follows. First, robot #1 searches the vicinity for another robot (say, robot #2) whose pose is known (this is determined via communication). Robot #1 then visually acquires robot #2 using an object recognition algorithm. The algorithm identifies the centroid of the robot within the image frame using a color segmentation scheme and marks its pixel coordinates on that frame. An incremental depth-from-motion algorithm (see Fregene *et al.*, 2002 for more details) computes the depth for a window within the frame that encloses these coordinates. The required relative position is inferred from the computed depth and the bearing of robot #2 relative to robot #1 is approximately determined from the lateral displacement between the enclosed pixel coordinates and the coordinates of the frame's optical center. The robot states are then updated. More details on these approaches are available in Madhavan *et al.*, 2002.

4. Multi-Robot Mapping

Incremental terrain mapping takes place via four main processes. An incremental dense depth-from-camera-motion algorithm (which is an adaptation of the work reported in L. Matthies, *et al.*, 1989) is used to obtain depth ranges to various features in the environment. The relative pose of the robots at these locations are associated with particular depth information. An elevation gradient of the terrain is determined by fusing GPS altitude information and vertical displacements obtained from inclinometer pitch angles. The depth and elevation information are then registered with their associated covariances. The terrain map is updated to incorporate the registered values at their proper coordinates. The covariances associated with each measurement provides the confidence the algorithm has in that measurement. In the case of overlapping areas, this confidence determines whether or not the map is updated. The overall schematic diagram of the algorithm is shown in Figure 2. More details on our approach are available in Fregene *et al.*, 2002.

5. Multi-Robot Path Planning

Our multi-robot path planning approach operates as follows. First, each robot plans its own path independently using D^*. The path is broadcast to all other robots, so every robot knows all path information. Under our approach, the paths that are planned for each robot are fixed, i.e., the following steps will not alter the (x, y) sequences of the paths. Instead, we define velocity profiles so that, while robots follow their paths, they insert delays as required to avoid collisions. Once the paths are planned, the collision check is then executed. If the collision is a time-space collision, that is, two or more robots reach the same point at the same time, an N-dimensional coordination diagram (CD) is constructed with collision regions marked as obstacles in the diagram. D^* searches for a free trajectory in the coordination diagram. The trajectory is then interpreted into a velocity profile for each robot, and the performance index of the current trajectory solution is calculated. Since the searching in CD is distributed across the robots, each search can take a different cost function to minimize based upon differences in priorities

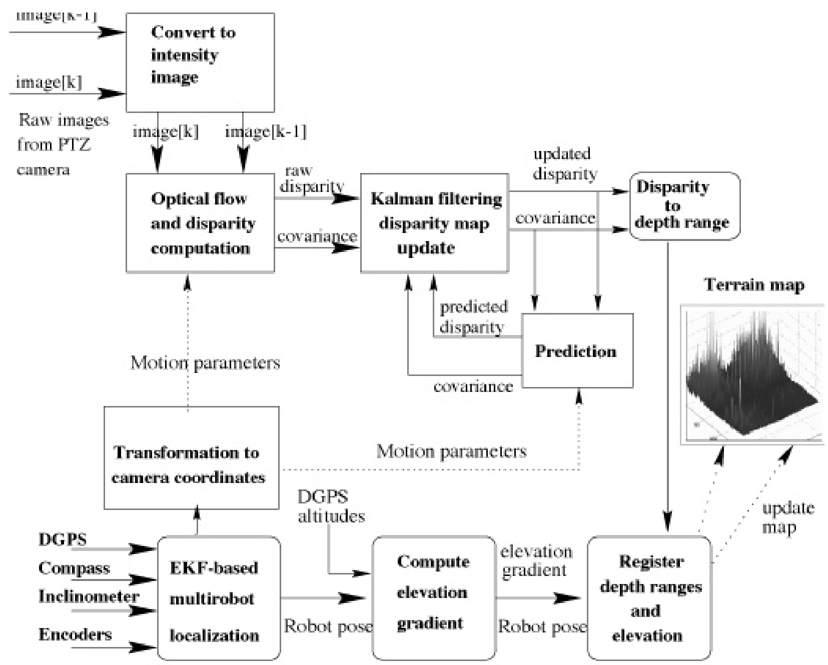

Figure 2. The overall terrain mapping scheme

between robots at intersections. Then the performance index and velocity profile are broadcast to all other robots. An evaluation is done to get a minimum value of the performance index, and the corresponding velocity profile is chosen. More details on our approach are available in Guo and Parker, 2002.

6. Experimental Results

We have implemented portions of this approach for multi-robot localization, mapping, and path planning in outdoor environments using distributed sensing. We briefly mention some of these results here, referring the reader to Fregene *et al.*, 2002; Madhavan *et al.*, 2002; Guo and Parker, 2002 for more details.

Figures 3 and 4 show the results for the laser-based cooperative localization described in Section 3. Figure 3 shows the estimated paths of robots #1 and #2. The pose standard deviations of robot #2 in Figure 4 demonstrate the utility of the relative pose information in accomplishing cooperative localization. At *time* = 21 seconds, DGPS becomes unavailable as indicated by the rise in the x standard deviation. It can be seen that as a result of the laser-based relative position information,

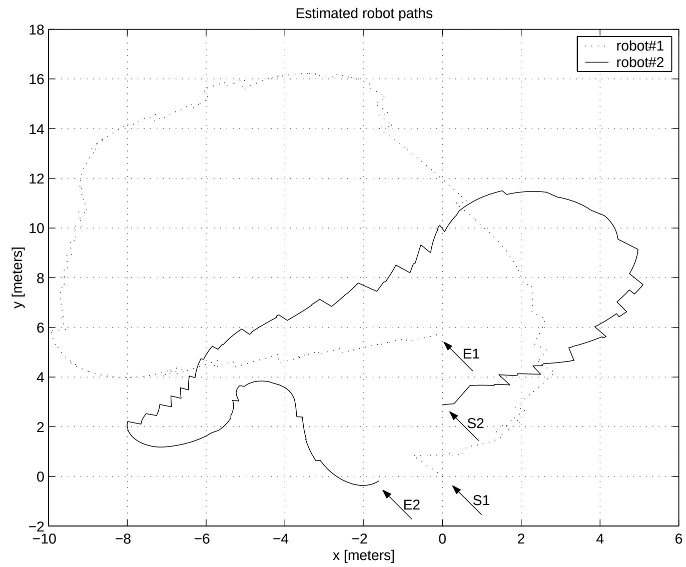

Figure 3. EKF estimated robot paths. The solid line denotes the estimated path of robot #2 and the dotted line that of robot #1. (S1,E1) and (S2,E2) denote the start and end positions for robots #1 and #2, respectively.

there is a sharp decrease in the position standard deviations of robot #2 (marked by arrows). As the motion of the robot is primarily in the x direction when the corrections are provided, the resulting decrease in the x standard deviation is noticeable compared to those in y and ϕ.

Figure 5 shows a partially updated terrain map that was developed by two robots, `Augustus` and `Theodosius`, using the mapping procedure outlined in Section 4. Although this update is still performed offline for now, it shows the elevation profile across the area traversed by each robot, with prominent features within the robot's field of view during the motion segment being marked on the map.

Our multi-robot motion planning algorithm has been implemented in a 3D vehicle planner and control simulation environment. For typical multi-robot paths, collisions will occur if the paths are planned separately. Therefore velocity planning is necessary to resolve potential collisions. The velocity planning (D^* search in coordination diagram) on a typical $117 \times 95 \times 99$ grid took about 4 minutes. No consideration was given to reduce computation time in the software implementation. The velocity profiles for a typical example will give several solutions. For example, for a three-robot situation, one solution would be to insert three unit time delays for robot 2 at the beginning of its movement, a second solution is to insert four unit time delays for robot 3 at the beginning of

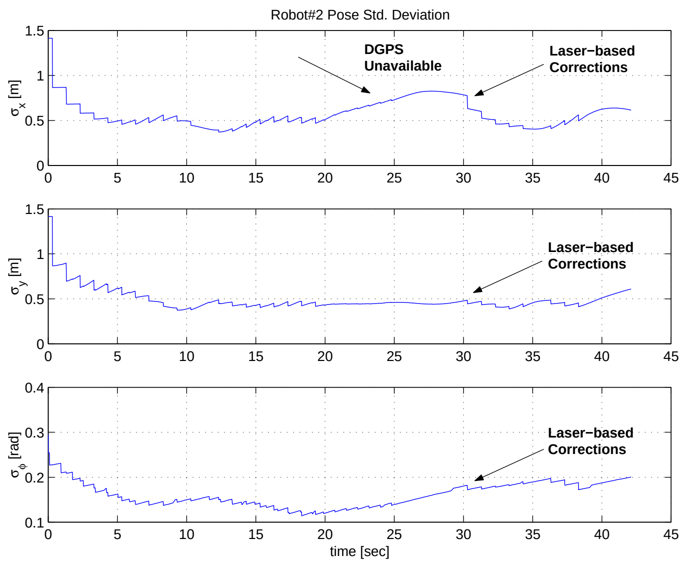

Figure 4. Laser-based cooperative localization, showing the standard deviation of the pose of robot #2. The external corrections offered by the laser-based localization scheme are marked by arrows.

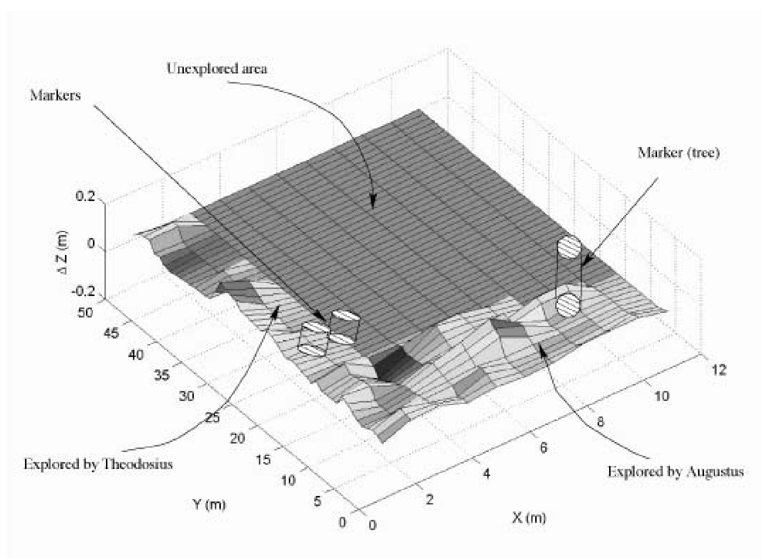

Figure 5. Partially updated terrain map.

its movement, and a third solution is identical to the first. The differences in schedules are caused by assigning different priorities to robots. Since the first index is the smallest, the corresponding set of velocity profiles are chosen for each robot. It should be noted that although more complicated velocity profiles (with many stop-move schedules in the middle of the velocity profile) can be generated by the described algorithm, from a practical concern, based on the same or a comparable performance index value, it is preferred to have delays at the beginning of the velocity profiles, or to be consolidated, instead of requiring a lot of move-stop-move procedures during the robot movement. This can be achieved by smoothing zig-zag paths in the searching algorithm. Experimental work is underway to implement this algorithm on our group of ATRV-mini all-terrain mobile robots.

7. Summary

In this paper, we have briefly outlined our approach toward using distributed heterogeneous sensing to achieve cooperative localization, mapping, and path planning in outdoor terrains using teams of mobile robots. We have sketched our algorithms toward achieving this goal and have given some initial results from implementation on our team of ATRV-mini robots. Our ultimate objective is to generate teams of mobile robots that can be placed in a previously unfamiliar outdoor environment, and use their distributed sensing capabilities to localize themselves, generate an approximate elevation map, and generate multi-robot paths that enable them to accomplish their intended objectives, such as perimeter security, reconnaissance and surveillance, and exploration.

References

Fregene, K., Madhavan, R., and Parker, L. E. (2002). Incremental multiagent robotic mapping of outdoor terrains. In *Proceedings of IEEE International Conference on Robotics and Automation.*

Guo, Y. and Parker, L. E. (2002). A distributed and optimal motion planning approach for multiple mobile robots. In *Proceedings of IEEE International Conference on Robotics and Automation.*

Matthies, L., Kanade, T., and Szelinski, R. (1989). Kalman Filter-based Algorithms for Estimating Depth from Image Sequences. *International Journal of Computer Vision*, 3:209–236.

Madhavan, R., Fregene, K., and Parker, L. E. (2002). Distributed heterogenous outdoor multirobot localization. In *Proceedings of IEEE International Conference on Robotics and Automation.*

Roumeliotis, S. and Bekey, G. (2000). Distributed Multi-Robot Localization. In *Distributed Autonomous Robotic Systems - Chapter 6: Localization*, pages 179–188. Springer-Verlag.