RADIATION MAPPING USING MULTIPLE ROBOTS

R. A. Cortez and H. G. Tanner

Mechanical Engineering Department, University of New Mexico, Albuquerque NM

We present a motion planning methodology to allow for a group of mobile sensors to scan a given polygonal area and map the radiation levels over it. We assume low radiation levels, which implies that detectors obtain meaningful measurements at small distances, and that radiation detectors are carried around mounted on mobile robots. The area is discretized at a resolution suggested by the footprint of the detector, and then portions of it are allocated to different robots for mapping. Robots cover their allocated areas in an efficient manner, collecting enough measurements at each location to allow a certain confidence level about the measured intensity to be reached. The algorithm is complete, in the sense that all locations within the area of interest are visited and mapped, and efficient since locations are being visited only once.

I. INTRODUCTION

Both in nuclear forensics and in emergency management after an accident or a deliberate attack, it is important to confirm the distribution of radiation levels over an area of interest as soon as possible, so that confident decisions can be made on the scene and rescue crews can be deployed safely and efficiently. In situations like these, autonomous mobile robots can be of use as a platform for portable radiation sensors.

Bringing detectors close to radiation sources brings important benefits to radiation mapping because of the inverse square relationship between the SNR and the distance to the radiating source. Robots also offer a straightforward way to map measurements to sensing locations. When mapping spatial distributions of physical quantities, measurement data are useless unless correlated to spatial coordinates. Finally, robots as platforms for radiation sensors allow for a systematic way of scanning and mapping areas, without exposing human operators to hazardous environments.

Deploying several robotic devices to search different portions of the area of interest accelerates the completion of the search and mapping task significantly. In addition, the use of multiple autonomous sensor platforms in a sensor network configuration offers robustness to individual subsystem failures: if one robot or sensor fails, others can pick up its allocated task and still finish the task cooperatively.

In previous work,[3] we developed a technique for robotic radiation search and mapping which has proved very promising in terms of map completion time. It is inspired by sequential testing theory for radiation source detection as far as sensor motion is concerned, but the underlying statistics and its implementation is quite different. The method uses the variance of an assumed probability distribution of source intensity at a given location as a measure of estimate accuracy, and updates this metric by taking measurements for different time intervals at each location. The main reason why it has been the fastest over alternative methods in our single-robot tests, is the fact that the robot visits each location only once. This highlights the importance of robot motion planning in such search and mapping tasks.

In this paper we extend this method to a multi-robot setting, we consider a group of N homogeneous robots with sensors. Assuming that the robot group is initially positioned randomly, robots set out to distribute themselves over the area, and this distribution induces a decomposition the area to be mapped into N polygonal regions. Each robot is assigned to map a different region, balancing the subsequent mapping effort (motion) of the robot to the effort spent during the initial deployment. In this way, we are able to minimize overhead time for the robots to transition to their assigned regions, and we run N parallel "sequential" mapping algorithms. Another novelty compared to our earlier work is related to the mapping motion strategy: our search areas are now polytopes as opposed to rectangles, and thus if robots are required to visit each location only once, more advanced planning strategies need to be employed to ensure that robots do not get trapped away from unexplored regions.

The rest of the paper is organized as follows. In Section II, we review reported work in literature that is related to the problem addressed here. In Section III, we formalize the problem to be solved and we state the assumptions that we have made. Section IV outlines our approach to solving the problem and is followed by Section V, which provides numerical results from simulation tests. Section VI concludes the paper by summarizing our results and contributions.

II. RELATED WORK

The problem addressed in this paper is conceptually similar to robotic exploration and mapping. However, there is an important caveat. Robot exploration typically involves creating a map of the known workspace which depicts the location of obstacles and landmarks. This is not what we do here. Our goal is to map the spatial distribution of a physical quantity over a given area. The workspace geometry is assumed known, and the robot localization problem considered solved in some way.

The process of traditional robot exploration is often based on an occupancy grid, which discretizes the area of interest into a large number of cells. The notion of grid maps or occupancy maps was introduced early on in the area of mobile robotics[4] and subsequently for topological mapping and exploration.[9,8] In most cases the cells of a grid map contain a probability value of whether that cell is *occupied by an obstacle*. Yamauchi[9] uses occupancy grids to define *frontiers*, which mobile robots push to expand their knowledge of the environment. Romero *et al.*[8] uses occupancy grids to minimize the cost to travel to an unoccupied cell for further investigation of the area. In this paper, we use each grid cell (henceforth called *location*) to hold a metric of the uncertainty regarding the radiation levels in that location.

Some work on mapping the distribution of physical quantities over a region (e.g. gas concentration, temperature) using robots has been reported in literature. Maps of gas concentrations have been constructed,[5] by maneuvering a robot using a predefined path that covers the entire area. In an approach to search for ocean features[7] multiple robots follow gradients to locate and track ocean features such as fronts and eddies. For the radiation mapping problem addressed here, sensor measurements at given locations are in theory random samples drawn from a Poisson distribution, and therefore vary widely making gradient calculations meaningless. Our approach is closer in principle to that of gas concentration mapping[5] in the sense that we follow open-loop motion plans, but these plans are not entirely predetermined. They are designed on-line, as the task progresses and the allocation of portions of the area of interest to robots is finalized.

Several methods have been reported in robotics literature for the (complete) coverage of known spaces[1,6,10]. Algorithms based on trapezoidal cell decomposition[1] can be complete but they are naturally suited to trapezoidal environments. Such algorithms guarantee that points will be visited *at least* once, where in our case, for efficiency reasons we require each cell to be visited *exactly once*. We therefore design our motion plans inspired by the algorithms that are based on the notion of a "distance transform,"[10], where one assigns a value to each cell in a discretized environment, and uses this assignment to guide the robot from high values to low values until the entire space is covered. This algorithm is also complete, and in our modified implementation, also ensures that cells will be visited only once.

III. PROBLEM STATEMENT

We are given a polygonal region and we are asked to create a map of radiation intensity over this area. Available are N mobile robots equipped with radiation sensors (Fig. 1–2) that can be autonomously controlled to wander and collect measurements. The problem is to design an algorithm to coordinate the motion of the robots, and plan sensor data collection so that measurements are taken from all locations in the given area, without a location being visited twice by any robot.

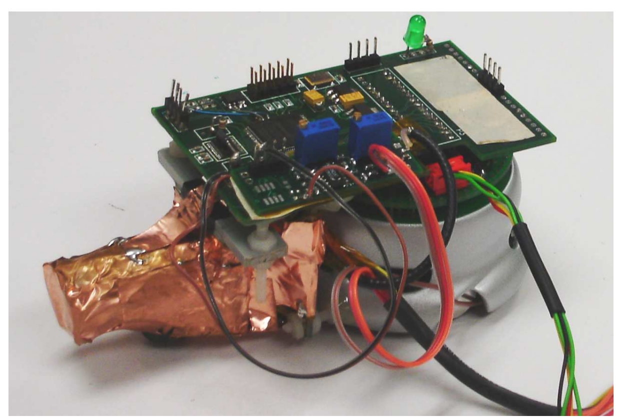

Fig. 1. The experimental platform used in our earlier work[3]. It consists of a Khepera II mobile robot on which a La_2Br scintillator is mounted.

We assume that robots can move in any direction – implying that we ignore the nonhonolomic motion constraints imposed on any wheeled vehicle. This assumption is justified from the fact that differential drive robots such as the one shown in Fig. 1 as well as others available in the market can turn in place. This allows us to use first order linear equations of the form

$$\dot{x} = u,$$

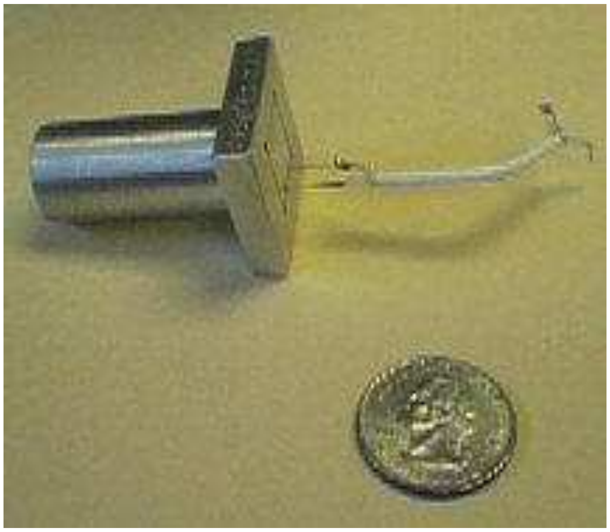

Fig. 2. A closer view of the scintillator mounted on the mobile robot.

with x being the robot's position and u being the control input, as a first approximation of the robots' kinematics.

Initially, the robots are assumed to be positioned at random configurations, but clustered around an initial location. No particular initial configuration for the robot team is to be assumed. We consider our measurements to be local (from within each cell), motivated by the fact that we are searching for weak radiation sources and the SNR of nuclear measurements drops with the square of the distance between sensor and source –in fact, if the sensor's solid angle is also taken into account (given that the sensor's geometry does not change) we have a combined effect of R^4. The sensor is therefore assumed to be picking radiation only from the cell it is in.

Finally, we assume that the topological structure of the workspace (boundaries etc.) is known exactly, and that robots are able to localize themselves within this environment accurately (through the use of some absolute positioning system, such as an overhead camera).

IV. APPROACH TO SOLUTION

We approach the multi-agent radiation mapping problem by first discretizing the area to be mapped by applying a grid, which can be as refined as the footprint of our radiation sensor, assumed $\frac{R}{2}$ allows. To capture the effect of the attenuating SNR, we model the ability of the radiation sensor to obtain (reliable) measurements from a given location q as a function of the distance from the sensor at point p, $s = \|q-p\|$

$$f(s) = \begin{cases} \frac{\exp(\frac{-1}{R^2-s})}{\exp(\frac{-1}{R^2-s})+\exp(\frac{-1}{s-R^2})} & s \in (\frac{R}{2}, R) \\ 1 & s \leq \frac{R}{2} \\ 0 & s > R \end{cases} \qquad (1)$$

Our map is an array corresponding to this grid, with each element holding a measure of uncertainty about the level of the distributed quantity in the associated location.

IV.A. Area Decomposition

We partition the area between the available mobile robotic platforms using Voronoi partitions in which the centroid of each Voronoi cell is taken to be the position of a single mobile robot. Thus, a certain region within this area (namely the corresponding Voronoi cell) is allocated to each robot for mapping.

Due, however, to the randomness of initial robot positions, the resulting Voronoi partition does not reflect fairness in the distribution of work, nor does it facilitate efficient task execution. We thus let this partition evolve, by controlling the robotic agents using the following scheme[2]:

$$u_i = -\int_{V_i} \frac{\partial f(\|q-x_i\|)}{\partial x_i} \phi(q)dq \qquad (2)$$

where V_i is the Voronoi cell allocated to agent i, f is the sensor performance function defined in (1), and $\phi(q)$ is the a density function distributed over the area, which we define as

$$\phi(q) = \|q-q_0^*\|,$$

with $q_0^* = \frac{1}{N}\sum_{i=1}^{N} x_i(0)$ being the centroid of the mobile robot group at initial time.

The effect of control law (2), is that agents distribute themselves in such away so that they get a bigger "chunk" of the density function within their cells. In the process of competing for higher density function values, the integrals of this density function over their Voronoi cells, become equal asymptotically. Each one ends up with either more space but smaller density values, or less space around higher values of the density function.

By designing the density function to capture the distance from the initial centroid of the team, we are penalizing the length of motion paths. Thus, agents who travel longer (which translates to more time and energy required to get into position for mapping) are allocated smaller portions of the area to be mapped. This does not have so big an effect on task completion time, since mapping will only start once every robot stabilizes and the partition is finalized, but it represents a fairer distribution of work in view of limited on-board power resources on behalf of the robots (typically powered by rechargeable batteries).

IV.B. Sweeping Strategy

The last component of our methodology relates to the way robots sweep their Voronoi cells, visiting all locations and lowering the variance there below the given threshold. The problem is not-trivial since we want each cell to be visited only once, yet we want to cover the whole Voronoi cell, which can have an arbitrary polygonal shape. Here, a robot runs the risk of disconnecting the unvisited regions within the cell, and getting trapped in a scanned portion when another portion is unexplored.

This is a motion planning (coverage) problem in a discrete environment, and the solution we implement is inspired by distance transforms.[10] We label each location inside a Voronoi cell with an integer that relates to the Manhattan distance from the cell's centroid, on an underlying lattice with diagonal links. The strategy then is to run an algorithm similar earlier algorithms,[10] in the sense that the robot seeks to transition to the location that is the highest within its immediate neighborhood. As a location is visited, its label is reset to a low value (zero). As the algorithm progresses, the robot first transitions to the location with the cell with the highest label and then, starting from the boundary works its way in. No cell is visited twice, and all the area inside the Voronoi cell is covered.

As the robot moves from location to location it collects measurements and maps the radiation levels in each location in a way that is described in the following section.

IV.C. Regional Mapping

Natural gamma ray background radiation has a cosmic ray component, and a component from naturally occurring radioactive isotopes. Small detectors (such as a one cubic inch La_2Br scintillator of Fig. 2) typically record low count rates, and the probability of observing k counts, given a mean expected count rate λ, is well described by the Poisson distribution

$$P(X=k|\lambda)=\frac{\lambda^k e^{-\lambda}}{k!}.$$

In each location within a Voronoi cell, assume a Gamma distribution for λ,

$$\pi(\lambda)=\beta^\gamma\lambda^{\gamma-1}e^{-\beta\lambda}\times\frac{1}{\Gamma(\gamma)},$$

where γ is the shape parameter, β is scale parameter and $\Gamma(\gamma)=\int_0^\infty t^{\gamma-1}e^{-t}dt$. The mean and variance of the gamma distribution can be expressed in terms of its shape and scale parameters

$$\mathbb{E}(\Gamma)=\frac{\gamma}{\beta}, \qquad V(\Gamma)=\frac{\gamma}{\beta^2}. \tag{3}$$

If no a priori information is available, the mean and variance for the prior distribution in each cell can be set to be equal.

The goal for the robots is now to lower the variance $V(\Gamma)$ at the cell they are currently visiting below a certain, pre-defined threshold. This is achieved by extending the integration time for the sensor since the application of Bayes rule updates the shape and scale parameters of the Γ distribution of a cell indexed by i and j between two successive time steps as follows[3]

$$\gamma_{ij}^{+}=\gamma_{ij}+c, \qquad \beta_{ij}^{+}=\beta_{ij}+1, \tag{4}$$

where c is the number of counts registered at that cell within the time step.

V. SIMULATION RESULTS

In our simulation tests, we have assumed a rectangular area that is to be mapped by four mobile robots.

V.A. Initial Deployment

The robots start clustered randomly around an arbitrary initial position. Fig. 3 shows this initial configuration and the Voronoi partition of the area of interest that is associated with it.

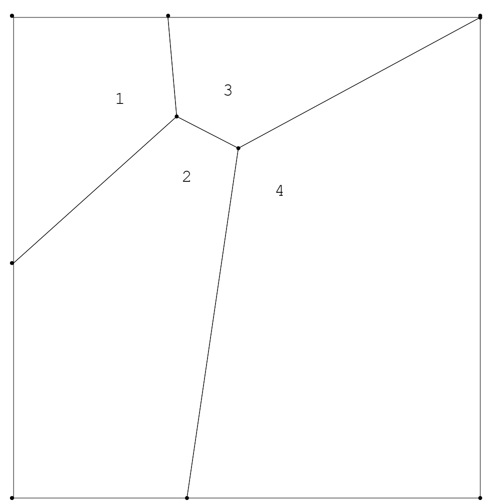

Fig. 3. The initial configuration of the robot team and the associated Voronoi partition.

Then the robots move to distribute themselves over the area to be mapped steered by control laws (2). The final (steady state) configuration for the closed loop control system under (2) is shown in Fig. 4. This depicts the finalized Voronoi partition. Note that robots 2 and 4 which have moved over longer paths are assigned slightly smaller Voronoi cells.

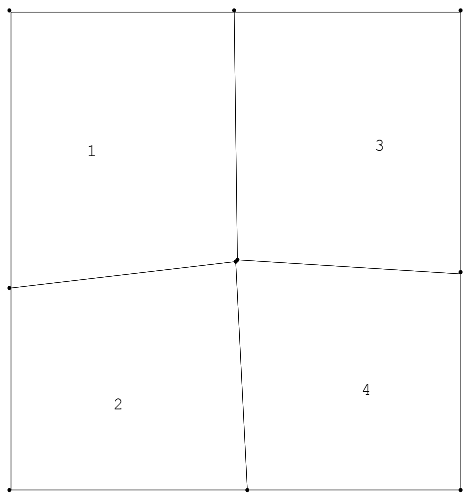

Fig. 4. The final configuration of Voronoi partition, as the robot team has reached a steady state driven under control law (2).

Our implementation of the “distance transform" for the cell allocated to each mobile platform is shown in Fig. 5. It is the squared Euclidean distance from the centroid of the cell –where each robotic platform stabilizes to under (2).

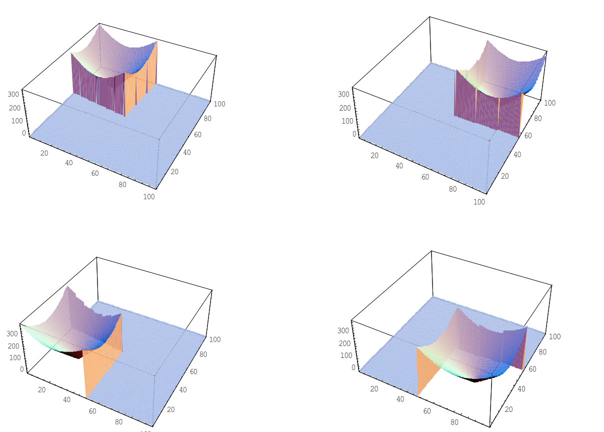

Fig. 5. The distance transforms used to generate labeling for the locations in each cell, before the motion plan is generated. Note that the view point for these plots is at the top left corner when looking at Fig. 4.

We assume a distribution of radiation over the area that is shown in Fig. 6. The goal of the simulation test is for the robots to reconstruct this plot using (simulated) measurements that they make as they sweep the area according to the motion planning strategy.

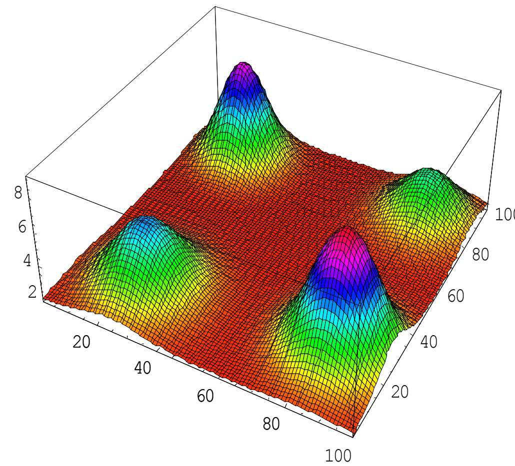

Fig. 6. The true radiation intensity distribution assumed for the area of interest.

As each robot visits a location, a sample is drawn from a Poisson distribution having a mean corresponding to the level of intensity of the distribution shown in Fig. 6. With each sample, the parameters of the local Gamma distribution are updated according to (4). Samples are continuously collected until the variance given in (3) drops below a certain threshold. At the end of the mapping procedure, the estimated radiation map (the distribution of the expected value as calculated in (3)) has the form seen in Fig. 7.

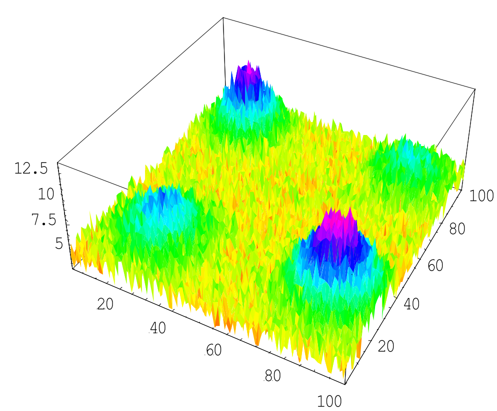

Fig. 7. The estimated radiation intensity distribution, reconstructed from local measurements.

VI. CONCLUSIONS

In this paper we extend one of the radiation mapping strategies that we developed in our earlier work[3] to a robot-enabled sensor network setting. We select a method for mapping that is inspired by the sequential testing theory, since our preliminary studies have shown it to be efficient in terms of completion time. The availability of multiple robots, serving as mobile sensor platforms for the radiation detectors, enable the mapping task to be divided among the robots and the mapping process to be accelerated. We present a methodology for area decomposition that takes into account the energy/time spent by each robot in order to reach its allocated area. We adapt an established robotic coverage algorithm to the case of polygonal workspace and we use it for navigation during mapping. We verify the algorithmic completeness of our algorithm through numerical simulations.

ACKNOWLEDGMENTS

Portions of this work were supported in part by Los Alamos National Laboratory Award No. STB-UC:06-36, and DoE URPR grant DE-FG52-04NA25590.

REFERENCES

1. H. Choset. Coverage of known spaces: The boustrophedon cellular decomposition. *Autonomous Robots*, 9(3):247–253, December 2000.

2. J. Cortés, S. Martinez, and F. Bullo. Spatially-distributed coverage optimization and control with limited-range interactions. *ESAIM: Control, Optimisation and Calculus of Variations*, 11(4):691–719, 2005.

3. R.A. Cortez, X. Papageorgiou, H.G. Tanner, A.V. Klimenko, K.N. Borozdin, R. Lumia, and W.C. Priedhorsky. Smart radiation sensor management; radiation search and mapping using mobile robots. *IEEE Robotics & Automation Magazine*, 2007. (in print).

4. Alberto Elfes and Hans P. Moravec. High resolution maps from wide angle sonar. In *IEEE International Conference on Robotics and Automation*, pages 116–121, 1985.

5. Achim Lilienthal and Tom Duckett. Building gas concentration gridmaps with a mobile robot. *Robotics and Autonomous Systems*, 48(1):3–16, 2004.

6. H.H. González-Ba nos and J.-C. Latombe. Navigation strategies for exploring indoor environments. *Journal of Robotic Systems*, 21(10-11):829–848, Oct-Nov 2002.

7. Petter Orgen, Edward Fiorelli, and Naomi Ehrich Leonard. Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment. *IEEE Transactions on Automatic Control*, 49(8):1292–1302, August 2004.

8. L. Romero, E. Morales, and E. Sucar. An exploration and navigation approach for indoor mobile robots considering sensor's perceptual limitations. In *IEEE International Conference on Robotics and Automation*, pages 3092–3097, May 2001.

9. B. Yamauchi. Frontier-based exploration using multiple robots. In *Second International Conference on Autonomous Agents*, pages 47–53, 1998.

10. A. Zelinsky, R.A. Jarvis, J.C. Byrne, and S. Yuta. Planning paths of complete coverage of an unstructured environment by a mobile robot. In *Proceedings of the International Conference on Advanced Robotics*, pages 533–538, 1993.