
Distributed Robotic Radiation Mapping

R. Andres Cortez, Herbert G. Tanner, and Ron Lumia

Mechanical Engineering Department, University of New Mexico
{acortez,tanner,lumia}@unm.edu

Abstract

We develop a control scheme to coordinate a group of mobile sensors for radiation mapping of a given planar polygon region. The control algorithm is based on the concept of information surfing, where navigation is done by means of following information gradients, taking into account sensing performance as well as dynamics of the observed process. The control scheme steers mobile sensors to locations at which they maximize the information content of their measurement data. The resulting controllers depend on both time, sensor measurements and local robot configurations, and the approach exploits the asymptotic properties of mutual information to offer an invariance-type stability result for a time-varying distributed cooperative system.

1 Motivation

This work is motivated by the emerging threat of contamination from a malicious attack or accidental release of radioactive material, as well as by problems related to nuclear forensics and non proliferation. In the latter cases, robots could investigate the possibility that fissile material has been processed, by searching for small specks of nuclear material. Currently, searching for radiation sources is usually done manually, by operators waving radiation counters in front of them.

2 Problem Statement

The goal is to estimate the distribution of radiation levels over a given area of interest, to a certain confidence level, using radiation measurements collected in a distributed fashion by multiple mobile robots. This radiation map should

be available as it is being constructed in real-time, and should be offered at a confidence level that gradually improves with time. We need to design and implement distributed control algorithms that steer the mobile robots into completing this map.

3 Related Work

Existing control design methods for multiple mobile robot coordination apply to problems ranging from topological mapping to active sensing and source seeking, as well as reconfigurable sensor networks. Within these existing control designs, entropy-based metrics have become more commonplace as a way to express estimate uncertainty or to aid in localizing target features.

3.1 Active sensing and source seeking approaches

The idea of active sensing has become an area of interest recently because of its applications in environmental monitoring, search and rescue, and source seeking problems. The approach is to use a mobile sensor platform that reacts “intelligently” to the sensed environment based on the measurements. An early approach to this technique is found in [1], where Chung et. al use a gradient decent algorithm to cooperatively estimate the state of dynamic targets in an optimal fashion. In [2] Yang et. al extend the work of [1] by making the control algorithm implemented in a distributed fashion. This is achieved by constructing a dynamic average consensus estimator and using a one-hop neighbor for communication.

Active sensing is not limited to target estimation. In [3], Zhang et. al address a source seeking problem for an autonomous vehicle modeled as a nonholonomic unicycle. In [4, 5, 6] the problem of source seeking is addressed when direct gradient information of the measured quantity is unavailable. Pang and Farrell [4] address chemical plume source localization by constructing a source likelihood map based on Bayesian inference methods. Mesquita et. al [5] induces source seeking behavior without direct gradient information by mimicking *E. Coli* bacteria. Lastly Mayhew et. al [6] propose a hybrid control strategy to locate a radiation source utilizing only radiation intensity measurements, however they do not consider the statistics of nuclear measurement.

3.2 Entropy-based metrics in robotic control

Entropy-based metrics quantifying uncertainty or information gain have been utilized in robotic exploration, localization and mapping. Moorehead [7] uses entropy, among other utility measures, to evaluate the benefit of visiting different locations in the context of exploration. However, that entropy describes

the uncertainty over a certain location being reachable; it is not directly associated with the quality of the model nor is it linked with the statistics of measurement. Mutual information is used as a metric of significance of different discrete locations containing features or targets of interest, in an application where mobile robots and unmanned aerial vehicles (UAVs) localize target features in their environment [8]. Each vehicle takes local actions that reduce estimate uncertainty the most by following mutual information gradients. The Fisher information is another metric used to describe control actions for general information acquisition tasks [9] as well as motion coordination tasks [10].

The most closely related work to the one presented here is probably the one of Sujan and Dubowksy [11], who derive a metric for determining optimal sensing poses for mobile robot agents that create a detailed three dimensional model of their environment. The metric used is based on Shannon's entropy. The main difference between [11] and our approach is that the underlying statistics are completely different. Here, radiation measurement statistics results in closed loop controllers which are both state and time dependent, making the stability analysis considerably more difficult.

3.3 Sensor network deployment

In [12], Popa et. al employ potential fields to reconfigure a mobile sensor network. Through these artificial potentials, the spatial distribution of the network can be regulated. Cortes, Martinez, and Bullo [13, 14], use gradient climbing algorithms to distribute sensor platforms in a geometrically optimal fashion over a given polygonal area. The area is divided up among the agents using Voronoi partitions. Hussein and Stipanovic [15] use another gradient climbing method for control of the sensor network, which does not partition the area among the team members, thus reducing computational overhead.

Scalable approaches to flocking and formation control rely on decentralized control schemes. Lindhe et. al [16] use a navigation function with Voronoi partitions to move the group of agents from one region to another *goal* region in the area while keeping formation. The navigation function allows for obstacle avoidance, and only neighboring agents position information is needed to compute the control actions. Yamaguchi [17] addresses a type of pursuit evasion game, where the group of robots try to keep a formation and enclose the evader. Coordination is achieved with no communication between robots and the control law uses position rules to coordinate the group around the target. Fredslund and Mataric [18] allow robots to fall into formation by selecting a leader of the group and having each agent maintain a certain position and orientation with respect to its neighbor. Minimal communication is required by each agent to verify that its neighbors are participating in the formation task.

In hazardous environments mobile robots can maneuver, take measurements, and build a map without human intervention. Utilizing large numbers

of mobile robots speeds up the map building task and makes the map available faster. Mobile robots can be equipped with a variety of sensors, so both topological (obstacles and freespace) and spatial distribution maps can be created. Spatial distributions of interest include temperature, gas, as well as radiation.

3.4 Spatial distributions and topological maps

Most of the recent work in the area of mapping focuses on obtaining a detailed topological or metric map of the environment in which the robots are deployed. Accurate two-dimensional and three-dimensional maps are constructed for static [19] and dynamic environments [20]. To build detailed metric maps, laser-range finders [21] or sonar sensors [22, 23] are typically used. These sensors report objects in line of sight at relatively long ranges. Thus establishing line of sight between sensors and environmental features is a critical aspect of topological mapping.

Mapping spatial distributions can provide insight to the short or long term effects of variation in temperature, pressure, water current, and may warn against threats from pollutants. In [24], Lilienthal and Duckett create maps of gas concentrations by maneuvering a robot over a predefined path that covers the entire area. An approach to search for ocean features is found in [25] where Orgen et. al use virtual bodies and artificial potentials to coordinate multiple robots into gradient climbing, in order to locate and track ocean features such as fronts and eddies. Singh et. al [26], use active learning techniques to map water currents in a freshwater lake to understand the hydrodynamics present in Lake Wingra.

Some attempts to automate nuclear search appear in [27, 28], using a strategy based on the classical sequential testing theory [29]. This strategy, however, is a detection strategy, not a mapping technique. It only gives a positive or negative response to the question of whether a source of a particular strength is present at that location. It gives no information on background radiation levels of the area. Using traditional sequential testing theory for radiation mapping would involve setting a number of different thresholds, one set for each anticipated radiation level. The area has to be scanned each time a certain level needs to be confirmed. This approach is far from time-optimal, and the required completion time increases very fast with the desired map resolution.

Even as a detection tool, for which the technique is optimal in terms of completion time, there is an important caveat. The algorithm is sensitive to the assumption on the strength of the source that is to be detected.

To overcome the shortcomings of [27, 28] an algorithm is developed in [30] to calculate and update, in real-time, the belief about local radiation levels, thus creating a radiation map of the area in question. The variance of the local radiation distribution level is used as a metric to assess the accuracy of the radiation map.

4 Technical Approach

Low-rate counting of radiation from nuclear decay is described by the Poisson statistics. On the other hand, the expected number of source counts μ to be registered by a moving sensor drops dramatically as a function of the distance $\mu = \chi \cdot \alpha \int_0^t \frac{1}{r^2(\tau)} d\tau$, where χ is the cross sectional area of the sensor, α is the activity of the source measured, and $r(t)$ is the instantaneous distance of the source to the sensor. The probability density function (PDF) associated with the random variable expressing the total number of counts recorded by the moving sensor c as $f(c) = f(c|\alpha, \chi, r(t))$, and can be updated recursively using Bayes rule. The most conservative assumption regarding $f(\alpha)$ prior is to take it as being uniformly distributed between two values α_1 (background) and α_2 (a strong source). For this problem, the mutual information of a measurement, I when the latter is viewed as a channel from the environment to the robot is given as $I(A; C) = h(A) - h(A|C)$, where A and C are random variables expressing radiation intensity and sensor counts, respectively, and $h(\cdot)$ denotes entropy. Following [14], we define an objective function

$$\mathcal{W}(P) = \int_{V_i(P)} f(\|q - p_i\|) I(q, p_i, t) dq, \quad (1)$$

where $f(\|q - p_i\|)$ is a smooth (sensor performance) sigmoid function, identically equal to one over the footprint of the radiation sensor (Fig. 1) and decreasing by (at least) the inverse square of the distance from the sensor [31].¹

5 Results

Our approach then traces the steps of [14], with one significant difference: in (1), the “density function” $I(q, p_i, t)$ is both configuration and time dependent, and invariance-type arguments cannot be made. In this work, we exploit the asymptotic properties of the mutual information associated with nuclear measurement:

Lemma 1 ([32]). *If $I(q, p_i, t)$ be the mutual information of nuclear measurement, then $\lim_{t \rightarrow \infty} \frac{\partial I(q, p_i, t)}{\partial t} = 0$.*

Using this lemma we can recover local stability (integrals are evaluated over Voronoi cells following [14]):

Proposition 1. *Consider the gradient field defined by:*

¹ Unless the solid angle around the source remains constant, SNR actually scales proportionally to R^4 !




Fig. 1. A small radiation sensor (NaI scintillator), covered with copper tape to reduce noise, and interfaced with a Khepera II robot for low-count radiation detection.

$$\dot{p}_i = \int_{V_i(P)} \frac{\partial f(\|q - p_i\|)}{\partial p_i} I(q, p_i, t) dq + \int_{V_i(P)} f(\|q - p_i\|) \frac{\partial I(q, p_i, t)}{\partial p_i} dq. \quad (2)$$

This system stabilizes at configurations that (locally) maximizes the information flow from each sensor, as expressed by the product $I(q, p_i, t)f(\|q - p_i\|)$, for $i = 1, \dots, n$.

Proof. Take $\frac{1}{\mathcal{W}(P)}$ to be a Lyapunov-like function candidate $\mathcal{V}(P)$. Then,

$$\begin{aligned} \dot{\mathcal{V}}(P) = & - \left[\int_{V_i(P)} \frac{\partial f(\|q - p_i\|)}{\partial p_i} I(q, p_i, t) dq \right. \\ & \left. + \int_{V_i(P)} f(\|q - p_i\|) \frac{\partial I(q, p_i, t)}{\partial p_i} dq \right]^2 \\ & + \int_{V_i(P)} f(\|q - p_i\|) \frac{\partial I(q, p_i, t)}{\partial t} dq. \end{aligned}$$

Now assume that the system does not stabilize at the configurations mentioned. Then Lemma 1 suggests there exists a finite time T when one can find an $\epsilon > 0$ such that $\left| \frac{\partial (I(q, p_i, t)f(\|q - p_i\|))}{\partial p_i} \right| > \epsilon \quad \forall t > T$. It follows that

$$\left[\int_{V_i(P)} \left| \frac{\partial f(\|q - p_i\|)}{\partial p_i} I(q, p_i, t) \right| dq + \int_{V_i(P)} \left| f(\|q - p_i\|) \frac{\partial I(q, p_i, t)}{\partial p_i} \right| dq \right]^2 > \int_{V_i(P)} \epsilon^2 dq.$$

Once we define the constant $W_3(\epsilon) \triangleq \int_{V_i(P)} \epsilon^2 dq$ and the function $\beta(t) \triangleq \int_{V_i(P)} f(\|q - p_i\|) \frac{\partial I(q, p_i, t)}{\partial t} dq$ then the derivative of \mathcal{V} can be bounded as follows $\dot{\mathcal{V}}(P) \leq -W_3(\epsilon) + \beta(t) \leq -(1-\theta)W_3(\epsilon) - \theta W_3(\epsilon) + \beta(t)$, where $0 < \theta < 1$. Knowing that $\lim_{t \rightarrow \infty} \frac{\partial I(q, p_i, t)}{\partial t} = 0$, and noticing that the integration in $\beta(t)$ is over position, we conclude that $\lim_{t \rightarrow \infty} \int_{V_i(P)} f(\|q - p_i\|) \frac{\partial I(q, p_i, t)}{\partial t} dq = 0$. Since W_3 is a strictly positive number there exists a time τ such that $-\theta W_3(\epsilon) + \beta(t) \leq 0 \quad \forall t > \tau$. After time τ , $\dot{\mathcal{V}}(P) \leq -(1-\theta)W_3(\epsilon) \triangleq \dot{\gamma}(t)$. Noting that $\dot{\gamma}(t)$ only depends on ϵ , we can get $\gamma(t) = \gamma(0) - (1-\theta)W_3(\epsilon)t$. The Comparison Lemma suggests that $\mathcal{V}(P) \leq \gamma(t) = \gamma(0) - (1-\theta)W_3(\epsilon)t$. This in turn implies that there exists a finite time when $\mathcal{V}(P) < 0$; contradiction, since by construction $\mathcal{V}(P) \geq 0$.

6 Experiments

Rather than experimenting with actual radioactive material² we use a light source to emulate radiation emission intensity (Fig. 2). Each robot takes its light intensity measurement and passes it through a Poisson distribution filter as the mean count rate at that particular cell. This filter returns a randomly distributed number from a Poisson distribution taken with a mean of the measured light intensity.

6.1 Algorithm Implementation Issues

The radiation map is supposed to be obtained by applying a Bayesian update for t ranging from the time of algorithm initialization until current time. The Bayesian update then gives the distribution of radiation activity mean value at each location $r(t)$, conditioned upon the measurements and the robots motion. There are, however, several reasons why the direct application of the Bayesian update may be impractical.

One issue is the numerical instability that may be observed when calculating the derivatives of mutual information, I , for large t . This phenomenon is due to the expression of this derivative involving incomplete gamma functions, which are evaluated over long time periods (implying large values for

² Preliminary experiments with a 10 nCi Na-22 radiation source, using the Khepera platform of Fig. 1 have been conducted by the first author at Los Alamos National Laboratory facilities.




Fig. 2. The experimental test bed with two mobile robots, each carrying a cricket beacon. The origin of the coordinate system is at the near right corner of the platform, and the light (simulating radiation) source in the far left corner at position (6,6). A red light was utilized because IR sensors are more sensitive to the red light wavelength.

c). In such cases, the incomplete gamma functions evaluate to very large numbers, sometimes causing numerical overflow. For this reason, we resort to a “receding horizon” type of approach, where we calculate these derivatives for an appropriately short time interval, update the radiation prior, reinitialize the algorithm and repeat with the new initial conditions.

However, once such an approach is adopted a second issue arises: the expressions for the marginal distributions and the mutual information derivatives are based on the assumption that the initial radiation prior is uniform. This assumption enables one to obtain closed form expressions for these derivatives, and subsequently for the control law. If $f(\alpha)$ is updated in real time, the radiation prior is no longer uniform when initializing the algorithm at the next step of the “receding horizon” method. One approach is to evaluate the derivatives numerically; another, which we opt for, is to discretize the radiation and mutual information maps in the form of a three-dimensional bar chart, where each cell is assigned to a uniform distribution, but the distribution is different among different cells. The map can then be updated cell by cell, since within each cell a uniform prior is assumed.

6.2 Radiation Map Construction

The area over which radiation is mapped is discretized into a $m \times n$ cell grid. Without knowing the type of radiation source in the given area, we assume that each cell in the area could contain a radioactive source of activity α , which could vary between α_1 and α_2 .

At time step $t = 0$, each robot stays for five time periods in the cell that it occupies and calculates the total number of the radiation counts collected during this period. This is done because too small sensor integration times may yield very large variations in count measurements, and temporarily increase uncertainty.

With each new set of count measurements c_{ij} over a cell (i, j) , the average of the total number of counts collected over that cell over the collective integration time, is updated as follows

$$\bar{c}_{ij}^+ = \frac{\bar{c}_{ij} \cdot (M_{ij} - 1) + c_{ij}}{M_{ij}} ,$$

where M_{ij} is the number of time periods a robot has spent in cell (i, j) , and \bar{c}_{ij}^+ is the updated average. Over time, the average number of counts collected in each cell approaches the true average according to the Law of Large Numbers.

This average \bar{c}_{ij} corresponds to the mean of the distribution that would have been obtained by the Bayesian update. To enable the next iteration of the algorithm in a “receding horizon” fashion, we update $f(\alpha)$ in each cell by approximating with a uniform one, over the interval $[\alpha_1, \alpha_2]$ as follows

$$\alpha_1^+ = \bar{c}_{ij} - \delta_{ij} \quad (3a)$$

$$\alpha_2^+ = \bar{c}_{ij} + \delta_{ij} \quad (3b)$$

$$\delta_{ij} \triangleq \min \{ |\bar{c}_{ij} - \alpha_1|, |\bar{c}_{ij} - \alpha_2| \} .$$

To speed up the map construction at the initial phases, we consider the (remote, for the case of weak radioactivity) possibility that counts collected by robots over particular cells may have been a result of radiation from neighboring cells. In this case, the radiation intensity of these neighboring cells is adjusted according to

$$c_{ij} = \sum_{k=1}^n \bar{c}_k \frac{f_k(\|q - p_k\|)}{\chi d_{ij}(k)} \quad (4)$$

where c_k is number of counts collected (remotely) by robot k , χ the radiation sensor cross-section, $f_k(\|q - p_k\|)$ is robot’s k sensing performance function, and the function $d_{i,j}(k) \triangleq \int_0^t \frac{1}{\|q_{i,j} - p_k(t)\|^2} dt$ reflects the fact that the likelihood of registering counts from remote radiation sources drops with the square of the distance to the source. Since (4) suggests that registering c counts from remote locations implies that the assumed source at these locations radiates

intensively, the resulting c_{ij} estimate is compared with the current belief α_{ij} about the radiation intensity at (i, j) , using the Bayesian update

$$F_{\alpha_{ij}} = \frac{f(\alpha) \cdot f(\bar{c}_{ij}^+ | \alpha_{ij}, \chi, r(t))}{f_c(\bar{c}_{ij}^+)} \cdot \sigma, \quad (5)$$

where $F_{\alpha_{ij}}$ plays the role of our confidence on this new projected radiation count \bar{c}_{ij}^+ initiated at location (i, j) . The process of updating the intensity in cell (i, j) is now conditioned upon whether this number of counts is more likely to have been produced by a source of the assumed intensity, or from a source with a candidate intensity of

$$\alpha'_{ij} = \frac{\frac{1}{\alpha_2 - \alpha_1} \cdot \alpha_{ij} + F_{\alpha_{ij}} \cdot c_{ij}}{\frac{1}{\alpha_2 - \alpha_1} + F_{\alpha_{ij}}}.$$

If (5), when evaluated for α'_{ij} instead of α_{ij} , yields a value that is larger than $F_{\alpha_{ij}}$, then the distribution in (i, j) is updated through (3) using $\bar{c}_{ij} = \alpha'_{ij}$; otherwise, the existing estimate is kept.

As the search progresses the system is able to narrow down the possible source activities in each cell: the difference between α_1 and α_2 becomes smaller. Our metric for the completion of the task is based on the range of possible radiation source activities in each cell. We take an average difference between α_1 and α_2 over the entire area to measure how well our estimated radiation map represents the actual radiation distribution in the area. When the average difference between α_1 and α_2 is below a threshold of 1.5 counts/sec the algorithm terminates.

6.3 Experimental Snapshots

The area over which radiation is mapped is discretized into a 6×6 cell grid, and a uniform distribution between $\alpha_1 = 1$ and $\alpha_2 = 10$ counts/sec is assumed each cell's radiation intensity α . Localization of the robots is done through odometry and triangulation using distance measurements from Crossbow Crickets attached on each one of the robots shown in Fig. 2.

Fig. 3 shows the change in the information map and the updated radiation intensity map after two updates (each update takes five seconds). The completed radiation intensity map is shown in Fig. 4 and is obtained after 17 updates.

In the setting described here, the algorithm is implemented in a centralized fashion, so issues related to keeping, communicating, and updating the mutual information map that needs to be shared among robots are not addressed. The control architecture allows decentralization, up to map maintenance and sharing. The latter task depends on available communication bandwidth and local memory storage.

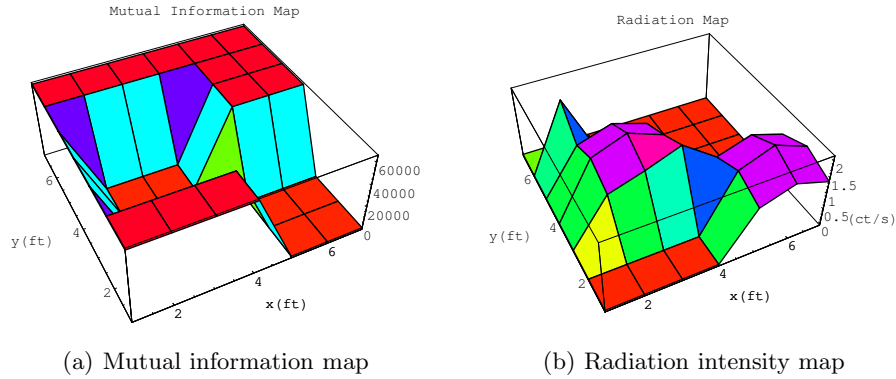


Fig. 3. Map estimate after two updates.

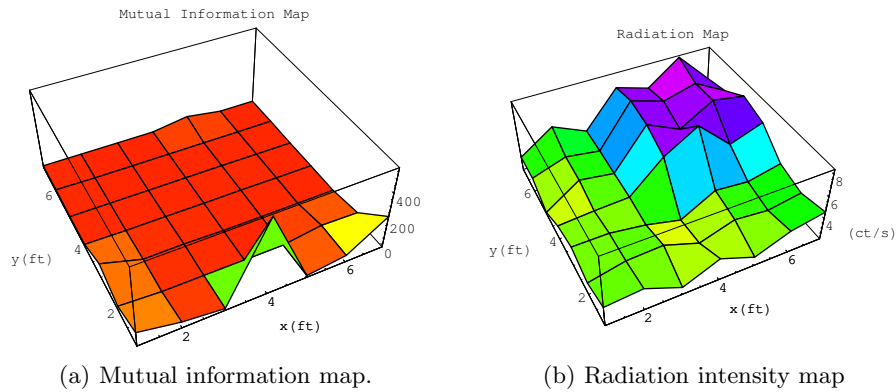


Fig. 4. Map estimate after 17 updates.

However during experimental testing, issues arose with wireless localization information from the Crossbow Crickets attached to the robots interfering with wireless transmissions of radiation information from the robots to the central computer. This was due to the close frequency ranges of the wireless devices. To overcome these interference issues, an overhead camera system is being developed to localize the robots during experimentation which will replace the Crossbow Cricket sensors.

7 Conclusion

Radiation map construction is demonstrated on an experimental testbed, simulating radiation intensity with light intensity and sampling from Poisson

distributions generally accepted to describe the statistics of nuclear measurement. The algorithm is shown to be a useful tool in assessing nuclear contaminations and situations of nuclear forensics.

Acknowledgment

R. Andres Cortez was supported in part by the Los Alamos National Laboratory Award No. STB-UC:06-36 and in part by DoE URPR grant DE-FG52-04NA25590. The latter grant also supported Herbert Tanner and Ron Lumia.

References

1. T. H. Chung, V. Gupta, J. W. Burdick, and R. M. Murray, "On a decentralized active sensing strategy using mobile sensor platforms in a network," in *IEEE Conference on Decision and Control*, December 2004, pp. 1914–1919.
2. P. Yang, R. A. Freeman, and K. M. Lynch, "Distributed cooperative active sensing using consensus filters," in *IEEE International Conference on Robotics and Automation*, April 2007, pp. 404–410.
3. C. Zhang, D. Arnold, N. Ghods, A. Siranosian, and M. Krstic, "Source seeking with nonholonomic unicycle without position measurement part 1: Tuning of forward velocity," in *IEEE Conference on Decision and Control*, December 2006, pp. 3040–3045.
4. S. Pang and J. A. Farrell, "Chemical plume source localization," *IEEE Transactions On Systems, Man, And Cybernetics*, vol. 36, no. 5, pp. 1068–1080, October 2006.
5. A. R. Mesquita, J. P. Hespanha, and K. Astrom, "Optimotaxis: A stochastic multi-agent on site optimization procedure," in *11th International Conference on Hybrid Systems: Computation and Control*, April 2008.
6. C. G. Mayhew, R. G. Sanfelice, and A. R. Teel, "Robust source-seeking hybrid controllers for autonomous vehicles," in *American Control Conference*, July 2007, pp. 1185–1190.
7. S. J. Moorehead, "Autonomous surface exploration for mobile robots," Ph.D. dissertation, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, Pennsylvania 15213, August 2001.
8. B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, "Cooperative air and ground surveillance," *IEEE Robotics and Automation Magazine*, pp. 16–26, September 2006.
9. B. Grocholsky, A. Makarenko, and H. Durrant-Whyte, "Information-theoretic coordinated control of multiple sensor platforms," in *IEEE International Conference on Robotics and Automation*, September 2003, pp. 1521–1526.
10. R. Olfati-Saber, "Distributed tracking for mobile sensor networks with information-driven mobility," in *American Control Conference*, July 2007, pp. 4606–4612.
11. V. A. Sujan and S. Dubowsky, "Visually guided cooperative robot actions based on information quality," *Autonomous Robots*, vol. 19, pp. 89–110, 2005.
12. D. O. Popa, H. E. Stephanou, C. Helm, and A. C. Sanderson, "Robotic deployment of sensor networks using potential fields," in *IEEE International Conference on Robotics and Automation*, April 2004, pp. 642–647.

13. S. Martinez, J. Cortes, and F. Bullo, "Motion coordination with distributed information," *IEEE Control Systems Magazine*, vol. 27, no. 4, pp. 75–88, August 2007.
14. J. Cortes, S. Martinez, and F. Bullo, "Spatially-distributed coverage optimization and control with limited-range interactions," *ESIAM: Control, Optimization and Calculus of Variations*, vol. 11, no. 4, pp. 691–719, 2005.
15. I. I. Hussein and D. M. Stipanovic, "Effective coverage control for mobile sensor networks," in *IEEE Conference on Decision and Control*, 2006, pp. 2747–2752.
16. M. Lindhe, P. Orgen, and K. H. Johansson, "Flocking with obstacle avoidance: A new distributed coordination algorithm based on voronoi partitions," in *IEEE International Conference on Robotics and Automation*, April 2005, pp. 1785–1790.
17. H. Yamaguchi, "A cooperative hunting behavior by mobile robot," in *IEEE International Conference on Robotics and Automation*, May 1998, pp. 3204–3209.
18. J. Fredslund and M. J. Mataric, "A general algorithm for robot formations using local sensing," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 837–846, October 2002.
19. S. Thrun, "A probabilistic on-line mapping algorithm for teams of mobile robots," *The International Journal of Robotics Research*, vol. 20, no. 5, pp. 335–363, May 2001.
20. D. Hahnel, R. Triebel, W. Burgard, and S. Thrun, "Map building with mobile robots in dynamic environments," in *IEEE International Conference on Robotics and Automation*, September 2003, pp. 1557–1563.
21. W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun, "Collaborative multi-robot exploration," in *IEEE International Conference on Robotics and Automation*, April 2000, pp. 476–481.
22. R. Grabowski, L. E. Navarro-Serment, C. J. Paredis, and P. K. Khosla, "Heterogeneous teams of modular robot for mapping and exploration," *Autonomous Robots*, vol. 8, no. 3, pp. 293–308, June 2000.
23. D. Pagac, E. Nebot, and H. Durrant-Whyte, "An evidential approach to map-building for autonomous vehicles," *IEEE Transactions On Robotics and Automation*, vol. 14, no. 4, pp. 623–629, August 1998.
24. A. Lilienthal and T. Duckett, "Building gas concentration gridmaps with a mobile robot," *Robotics and Autonomous Systems*, vol. 48, no. 1, pp. 3–16, 2004.
25. P. Orgen, E. Fiorelli, and N. E. Leonard, "Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment," *IEEE Transactions on Automatic Control*, vol. 49, no. 8, pp. 1292–1302, August 2004.
26. A. Singh, R. Nowak, and P. Ramanathan, "Active learning for adaptive sensing networks," in *5th International Conference on Information Processing in Sensor Networks*, 2006, pp. 60–68.
27. A. Kumar, H. G. Tanner, A. V. Klimenko, K. Borozdin, and W. C. Priedhorsky, "Automated sequential search for weak radiation sources," in *IEEE Mediterranean Conference on Control and Automation*, June 2006, pp. 1–6.
28. R. Cortez, X. Papageorgiou, H. G. Tanner, A. V. Klimenko, K. N. Borozdin, and W. C. Priedhorsky, "Experimental implementation of robotic sequential nuclear search," in *Proceedings of the IEEE 15th Mediterranean Conference on Control and Automation*, June 2007.

29. A. Wald, "Sequential tests of statistical hypotheses," *Annals of Mathematical Statistics*, vol. 16, pp. 117–186, 1945.
30. R. Cortez, X. Papageorgiou, H. G. Tanner, A. V. Klimenko, K. N. Borozdin, R. Lumia, and W. C. Priedhorsky, "Smart radiation sensor management: Nuclear search and mapping using mobile robots," (to appear) *IEEE Robotics and Automation Magazine*, 2007.
31. A. Klimenko, W. C. Priedhorsky, N. Hegartner, and K. N. Borozdin, "Efficient strategies for low-statistics nuclear searches," *IEEE Transactions on Nuclear Science*, vol. 53, no. 3, pp. 1435–1442, June 2006.
32. R. A. Cortez and H. G. Tanner, "The entropy of cooperative radiation sensing by distributed sensors," in *ANS 2nd International Joint Topical Meeting on Emergency Preparedness and Response and Robotic and Remote Systems*, March 2008.