2012 IEEE International Conference on Robotics and Automation
RiverCentre, Saint Paul, Minnesota, USA
May 14-18, 2012

Robust Optimal Deployment of Mobile Sensor Networks

Seth Hutchinson and Timothy Bretl

Abstract— **A common algorithm for deployment of a mobile sensor network in a bounded domain moves each sensor toward the centroid of its Voronoi cell. This algorithm is optimal for a particular cost function that is expressed as a sum over Voronoi cells, where the placement of a sensor in its own cell has no effect on cost in other cells. We provide a probabilistic interpretation of this "partitioned" cost function in the context of a target detection task, where each sensor has a chance of seeing the target that decreases monotonically with distance and where the goal is to minimize the total probability of missed detection. We show that the partitioned cost function is exactly the probability of missed detection given that a sensor can only see a target in its own Voronoi cell. We derive the probability of missed detection in the general case—where each sensor might see the target anywhere—and show that optimal sensor placement changes. Finally, we derive the probability of missed detection given the possibility of sensor failure, producing a robust measure of cost with respect to which optimal sensor placement is different yet again. Our results are illustrated by several examples in simulation.**

I. INTRODUCTION

In this paper, we consider the deployment of a mobile sensor network with the goal of target detection. The target location is drawn from a given probability distribution on a bounded domain. Each sensor can take a single measurement and has a chance of detecting the target that decreases monotonically with the distance away—olfactory receptors and audio microphones, for example, are consistent with this model. We would like to derive a deployment algorithm resulting in sensor locations that minimize the total *probability of missed detection*, i.e., the probability that no sensor detects the target. We are particularly interested in deployment algorithms that are decentralized, for example where the motion of any given sensor depends only on the location of its nearest neighbors.

Our problem can be viewed as an instance of distributed coverage, for which a natural solution approach might be the deployment algorithm of Cortes et al [1]. With this algorithm, each sensor moves toward the centroid of its Voronoi cell. This algorithm is decentralized in that the descent direction for each sensor depends only on its own location and on the location of its Voronoi neighbors. This algorithm also produces sensor locations that are locally optimal for a particular choice of cost function. This cost function can be expressed—and computed in a distributed way—as a sum over Voronoi cells, where the placement of a sensor in its own cell has no effect on cost in other cells. The resulting sensor locations form a centroidal Voronoi tessellation of the domain, a nice property from the standpoint of analysis [2].

However, it is not clear how the "partitioned" cost function of Cortes et al [1] corresponds, if at all, to the probability of missed detection. Without understanding this relationship, we cannot say exactly why the algorithm of Cortes et al [1] should be considered a "good" or "bad" approach to deployment of a mobile sensor network for target detection or any other task. Furthermore, when asking questions about the performance of this algorithm under perturbation or with the possibility of sensor failure (when a sensor with certainty does not see the target), we are tempted to draw conclusions that are obviously incorrect. For example, to measure the effect of failure at a given sensor i, one might simply ignore cost accumulated over the Voronoi cell $\mathcal{W}_i$. Doing so is nonsense, since it means that sensor failure decreases rather than increases total cost. What should be done instead?

We will show that the partitioned cost function of Cortes et al [1] is—for the right choice of parameters—exactly the probability of missed detection, given that a sensor can only see a target in its own Voronoi cell. We will proceed to derive the probability of missed detection in the general case, where each sensor might see the target anywhere, regardless of its location. We will show that the partitioned cost function is, in fact, an upper bound to the probability of missed detection in this general case. In other words, the decentralized deployment algorithm of Cortes et al [1] is based on a sensible approximation to the cost associated with our problem, but as a consequence results in sensor locations that are sub-optimal in the general case. Finally, we will return to the question of robustness and derive the probability of missed detection given the possibility of sensor failure. When we say that a sensor "fails" we mean that it will not detect the target, regardless of the target's location. In our analysis, we will assume that each sensor fails independently with equal probability. Again, we will show that optimal sensor placement with respect to this robust cost function is different from what would result from application of the partitioned cost function.

The deployment algorithms we produce—for the "total" cost that minimizes the probability of missed detection, and for the "robust" cost that minimizes this probability given the possibility of sensor failure—are not decentralized. They are descent methods, but computing the gradient is not something that we do in a distributed way. However, by understanding the way in which the partitioned cost of Cortes et al [1] is a relaxation of the total cost, we hope to have provided the foundation for the design of other decentralized algorithms in the future, particularly those

S. Hutchinson is with the Department of Electrical and Computer Engineering and T. Bretl is with the Department of Aerospace Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA {tbretl,seth}@illinois.edu

resulting in deployment that is robust to sensor failure.

The rest of our paper proceeds as follows. In Section II, we provide a brief overview of related work. In Section III, we define formally the problem of target detection for a mobile sensor network. In Section IV, we derive the probability of missed detection and show how it relates to the partitioned cost of Cortes et al [1]. In Section V, we derive an expression for the cost of sensor failure, and extend this result in Section VI to derive the probability of missed detection with the possibility of sensor failure. In Section VII, we compute gradients of each measure of cost (total, partitioned, robust), which can be used as the basis for a simple deployment algorithm, gradient descent. Finally, in Section VIII, we show example results in simulation and discuss their implications. We conclude with opportunities for future work (Section IX).

II. RELATED WORK

The problem that we consider here is an instance of the general search problem [3], [4]. Most research in this area considers active search, in which sensors patrol the workspace in search of a target [5]–[9]. We restrict ourselves in this paper to a static deployment of sensors. This restriction allows us to compute cost gradients analytically, which is often not possible with standard formulations of active search. Our problem may also be considered an instance of robot coverage [10], but again most research considers the case in which robots incrementally cover the environment as they move through it (sometimes called "sweep coverage"), whereas we focus on static deployment.

Coverage in sensor networks is a problem that has recently received considerable interest (e.g., [11]). These problems are often concerned with connectivity within the network [12]–[14] or with using sensor data to estimate properties of the environment [15]. Often, these approaches search for purely decentralized solutions, which would be suboptimal for the cost function that we present in the present paper. Nonetheless, we note that [16] considers a cost function that is directly related to our own, and that [17] considers a variant of our "robust" cost that is generated not by sensor fail but by communication dropouts.

III. PROBLEM DESCRIPTION

We consider the case of n sensors deployed in a workspace $\mathcal{W} \subseteq \mathbb{R}^2$, with the purpose of detecting a target whose position is unknown. We denote by p_i the position of the i^{th} sensor, and define a sensor configuration as $\mathcal{P} = \{p_1, \ldots, p_n\}$. The configuration space for the sensor network is thus $\mathcal{C} = \mathbb{R}^{2n}$. The target location is a random vector $\mathbf{X} \in \mathcal{W}$, with prior probability distribution ϕ, which we choose as the uniform density when no prior knowledge of target position is available. We define the cost for a given sensor configuration as

$$\mathcal{L}(\mathcal{P}) = P\{\text{missed detection} \mid \mathcal{P}\} \tag{1}$$

in which the probability of missed detection (to be derived in Section IV) is a function of both the sensor deployment $\mathcal{P}$ and the prior distribution for the target position, ϕ.

We assume that the vehicles obey simple first-order dynamics, such that $u_i = \dot{p}_i$, where u_i is the control input for the i^{th} vehicle. The optimal distributed control law satisfies

$$\mathcal{P}^* = \arg\min_{\mathcal{P} \in \mathcal{C}} \mathcal{L}(\mathcal{P}) \tag{2}$$

$$u_i(\mathcal{P}^*) = 0, \quad \text{for all } i \tag{3}$$

If the control law is truly decentralized, each u_i will be a function of only some subset of $\mathcal{P}$ (i.e., u_i will not depend on full configuration information for the entire network).

Now, suppose that one or more sensors fail to function (we will assume that they are unable to detect the target). Let $\mathcal{F}$ denote the set of sensors that fail. We will denote by $\mathcal{L}_{\mathcal{F}}(\mathcal{P})$ the probability that the network will fail to detect the target in the case when the sensors in $\mathcal{F}$ fail to operate. We can now define optimality as a min-max problem, in which the goal is to minimize the worst case performance of the network under prescribed failure conditions. For example, if we wish the network to be robust to failures by individual sensors, we would choose $\mathcal{P}^*$ as

$$\mathcal{P}^* = \arg\min_{\mathcal{P} \in \mathcal{C}} \max_{|\mathcal{F}|=1} \mathcal{L}_{\mathcal{F}}(\mathcal{P}) \tag{4}$$

Note that under this formulation, we have assumed that the autonomous vehicles are able to successfully execute their control algorithms, even though their ability to detect the target has been compromised.

IV. PROBABILITY OF MISSED DETECTION

We denote by D_i the event that the i^{th} sensor detects the target, and by $\overline{D}_i$ the event that the i^{th} sensor fails to detect the target. In general, the probability that the i^{th} sensor fails to detect the target increases with the distance between sensor and target. Let the conditional probability that the i^{th} sensor fails to detect the target when the target is located at q be

$$f_i(\|q - p_i\|) = P\{\overline{D}_i \mid \mathbf{X} = q\} \tag{5}$$

Assuming that the sensors are independent, the events D_i are conditionally independent given the position of the target (i.e., if the target position is known, detection by sensor i tells us nothing new about the probability of detection by other sensors[1]). The joint conditional probability that no sensor will detect the target located at position q is given by

$$P\{\overline{D} \mid \mathbf{X} = q\} = \prod_{i=1}^{n} P\{\overline{D}_i \mid \mathbf{X} = q\} = \prod_{i=1}^{n} f_i(\|q - p_i\|).$$

We apply the law of total probability to find the probability that a target will not be detected by any of the n sensors

$$\begin{aligned} P\{\overline{D}\} &= \int_{\mathcal{W}} P\{\overline{D} \mid \mathbf{X} = q\}\phi(q)dq \\ &= \int_{\mathcal{W}} \left(\prod_{i=1}^{n} f_i(\|q - p_i\|) \right) \phi(q)dq \end{aligned} \tag{6}$$

[1] If target position is unknown, the D_i are *not* independent events, since detection by sensor i gives information about target position, which can be used to infer something about the probability of detection by other sensors.

Note that $P\{\overline{D}\}$ depends on the sensor placement $\mathcal{P}$, and therefore should technically be written as $P\{\overline{D} \mid \mathcal{P}\}$. We have suppressed the dependence on $\mathcal{P}$ here merely to simplify notation. The algorithms in Section VII explicitly consider this dependence, and solve the problem of choosing $\mathcal{P}$ to minimize the probability of missed detection.

A. Partitioned Cost Functions

We consider now the special case in which the cost function (here, the probability of missed detection) can be "partitioned." In particular, we consider the case in which the cost can be expressed in terms of n independent cost functions, f_i, such that f_i depends only on p_i, and f_i is constant except for the so-called "dominant region" of the i^{th} sensor.

Define a partition of $\mathcal{W}$ by $\{W_j\},\ j = 1 \dots n$, such that $\cup W_j = \mathcal{W}$ and $W_i \cap W_j = \emptyset$ for $i \neq j$. When the indices are chosen so that $p_i \in W_i$, we say that W_i is the *dominant region* for the i^{th} sensor. We can write Equation 6 as a sum of the failure probabilities of the regions in the partition

$$P\{\overline{D}\} = \sum_j \int_{W_j} \left(\prod_{i=1}^{n} P\{\overline{D}_i \mid \mathbf{X} = q\} \right) \phi(q) dq \tag{7}$$

Equation 7 is again a cost function associated to the particular deployment of sensors $P = \{p_1, \dots, p_n\}$.

Consider now the special case for which a sensor can only detect the target if the target lies within its dominant region. In this case, we have

$$P\{\overline{D}_i \mid \mathbf{X} = q\} = \begin{cases} f_i(\|q - p_i\|) & q \in W_i \\ 1 & q \notin W_i \end{cases} \tag{8}$$

and Equation 7 reduces to

$$P_{\text{cvt}}\{\overline{D}\} = \sum_j \int_{W_j} f_j(\|q - p_j\|) \phi(q) dq. \tag{9}$$

This result is of interest, because it corresponds exactly to the cost function given in [1]. Note that P_{cvt} provides an upper bound on the probability that the network will fail to detect the target. We use the subscript *cvt* because in this case, the choice of $\mathcal{P}$ that minimizes cost corresponds to a centroidal Voronoi tesselation of $\mathcal{W}$ [2].

V. COST OF SENSOR FAILURE

We now consider the case in which one or more sensors fail. If the i^{th} sensor fails, its probability of missed detection will be one, without regard to the location of the target:

$$P\{\overline{D}_i \mid \mathbf{X} = q\} = \begin{cases} 1 & i^{th} \text{ sensor fails} \\ f_i(\|q - p_i\|) & \text{else} \end{cases} \tag{10}$$

Let $\mathcal{F}$ denote the set of sensors that fail to function properly (thus, only the sensors in the set $\overline{\mathcal{F}}$ function properly). We denote by $\Delta P_{\mathcal{F}}$ the increase in the probability of missed detection when those sensors in the set $\mathcal{F}$ fail to function properly. We can determine $\Delta P_{\mathcal{F}}$ as follows:

$$\begin{aligned} \Delta P_{\mathcal{F}} &= P\{\overline{D} \mid \text{sensors in } \mathcal{F} \text{ fail}\} - P\{\overline{D}\} \\ &= \sum_j \int_{W_j} \left(\prod_{i \in \overline{\mathcal{F}}} f_i(\|q - p_i\|) \right) \phi(q) dq \\ &\quad - \sum_j \int_{W_j} \left(\prod_{i=1}^{n} f_i(\|q - p_i\|) \right) \phi(q) dq \\ &= \sum_j \int_{W_j} \left(\prod_{i \in \overline{\mathcal{F}}} f_i(\|q - p_i\|) \right. \\ &\qquad \left. - \prod_{i=1}^{n} f_i(\|q - p_i\|) \right) \phi(q) dq \\ &= \sum_j \int_{W_j} M_{\overline{\mathcal{F}}}(q) \left(1 - M_{\mathcal{F}}\right) \phi(q) dq \end{aligned} \tag{11}$$

in which

$$M_{\mathcal{F}}(q) = \left(\prod_{i \in \mathcal{F}} f_i(\|q - p_i\|) \right)$$

$$M_{\overline{\mathcal{F}}}(q) = \left(\prod_{i \notin \mathcal{F}} f_i(\|q - p_i\|) \right).$$

Equation 11 provides a quantitative assessment of the effects of sensor failure on the probability of missed detection.

A. Application to Partitioned Cost Functions

Consider again the special case when sensors can only detect targets within their dominant region (as given by Equation 8). For dominant region j, there are two cases: either sensor j functions properly, and thus $j \in \overline{\mathcal{F}}$, or sensor j fails, and thus $j \in \mathcal{F}$. For a functioning sensor i.e., when $j \in \overline{\mathcal{F}}$, when $q \in W_j$ we have

$$\begin{aligned} &\prod_{i \in \overline{\mathcal{F}}} f_i(\|q - p_i\|) - \prod_{i=1}^{n} f_i(\|q - p_i\|) \\ &\qquad = f_j(\|q - p_j\|) - f_j(\|q - p_j\|) = 0 \end{aligned}$$

For a non-functioning sensor (i.e., when $j \in \mathcal{F}$), when $q \in W_j$ we have

$$\prod_{i \in \overline{\mathcal{F}}} f_i(\|q - p_i\|) - \prod_{i=1}^{n} f_i(\|q - p_i\|) = 1 - f_j(\|q - p_j\|)$$

Thus, for this model,

$$\Delta P_{\mathcal{F}} = \sum_{j \in \mathcal{F}} \int_{W_j} \left(1 - f_j(\|q - p_j\|)\right) \phi(q) dq$$

For the special case of failure by a single sensor, say sensor n, we have

$$\Delta P_{\{n\}} = \int_{W_n} \left(1 - f_n(\|q - p_n\|)\right) \phi(q) dq$$

In this case, the cost of sensor failure clearly increases as the size of the sensor's dominant region increases, the probability mass within the sensor's dominant region increases, or the effectiveness of the sensor increases.

VI. ROBUSTNESS TO SENSOR FAILURE

In this section, we derive the probability of missed detection in the case when sensors might fail. We assume that the sensors fail independently, each with probability p_{fail}. We first derive the general probability for missed detection, and then examine the specific case when exactly one sensor fails. The latter is useful if we wish to design a network configuration that is robust to individual sensor failures. We could, if desired, extend this analysis to the case of robustness to two sensor failures, three, etc., analogously to the derivation of error correcting codes (e.g., Hamming codes) that are robust to k-bit errors.

As a shorthand, denote the conditional probability of missed detection given failure by sensors in the set $\mathcal{F}$ as $P\{\overline{D} \mid \mathcal{F}\}$, and the probability that exactly the sensors in $\mathcal{F}$ fail as $P\{\mathcal{F}\}$. We can again use the law of total probality to determine the probability of missed detection for a sensor configuration $\mathcal{P}$:

$$\begin{aligned} P\{\overline{D}\} &= \sum_{\mathcal{F}\subseteq\{1\ldots n\}} P\{\overline{D} \mid \mathcal{F}\} P\{\mathcal{F}\} \\ &= \sum_{\mathcal{F}\subseteq\{1\ldots n\}} P\{\overline{D} \mid \mathcal{F}\} p_{\text{fail}}^{|\mathcal{F}|}(1-p_{\text{fail}})^{n-|\mathcal{F}|} \\ &= \sum_{\mathcal{F}\subseteq\{1\ldots n\}} \int_{\mathcal{W}} \left(\prod_{i\in\overline{\mathcal{F}}} f_i(\|q-p_i\|) \right) \phi(q)dq \\ &\qquad \times p_{\text{fail}}^{|\mathcal{F}|}(1-p_{\text{fail}})^{n-|\mathcal{F}|} \end{aligned} \tag{12}$$

In particular, for the probability of missed detection when exactly one sensor fails (i.e., of the joint event comprised of missed detection and single sensor failure) we have

$$\begin{aligned} P_1 &= \sum_{j=1}^{n} P\{\overline{D} \mid \{j\}\} P\{j\} \\ &= \sum_{j=1}^{n} \left(\int_{\mathcal{W}} \left(\prod_{i\neq j} f_i(\|q-p_i\|) \right) \phi(q)dq \right) \\ &\qquad \times p_{\text{fail}}(1-p_{\text{fail}})^{n-1} \end{aligned}$$

where $p_{\text{fail}}(1-p_{\text{fail}})^{n-1}$ is the probability that exactly one sensor fails (and therefore $n-1$ sensors do not fail).

VII. DEPLOYMENT ALGORITHMS

We have now derived three different cost functions governing variants of our target detection problem. We have the "total cost," $P\{\overline{D}\}$ given in Equation 6, which measures the probability of missed detection. We have the "partitioned cost," P_{cvt} given in Equation 9, which measures the probability of missed detection given that sensors see only their own Voronoi partition. And, we have the "robust cost," $P\{\overline{D}\}$ given in Equation 12, which measures the probability of missed detection when sensors fail independently with probability p_{fail}. In this section, we will compute the gradient of each cost with respect to the location p_j of each sensor. The result is a simple deployment algorithm—gradient descent—that leads to a locally optimal placement of sensors in each case.

A. Gradient of Total Cost

Let $g(p_1, \ldots, p_n)$ denote the probability of missed detection given in Equation 6. In this case

$$g(p_1, \ldots, p_n) = \int_{\mathcal{W}} \left(\prod_{i=1}^{n} f_i(\|q-p_i\|) \right) \phi(q)dq \tag{13}$$

and the gradient of g with respect to p_j is

$$\begin{aligned} &\nabla_{p_j} g(p_1, \ldots, p_n) \\ &= \nabla_{p_j} \int_{\mathcal{W}} \left(\prod_{i=1}^{n} f_i(\|q-p_i\|) \right) \phi(q)dq \\ &= \int_{\mathcal{W}} \nabla_{p_j} \left(\prod_{i=1}^{n} f_i(\|q-p_i\|) \right) \phi(q)dq \\ &= \int_{\mathcal{W}} \nabla_{p_j} f_j(\|q-p_j\|) \left(\prod_{\substack{i=1\\ i\neq j}}^{n} f_i(\|q-p_i\|) \right) \phi(q)dq. \end{aligned} \tag{14}$$

B. Gradient of Partitioned Cost

Now let g denote the partitioned cost function given in Equation 9. The gradient of

$$g(p_1, \ldots, p_n) = \sum_{i=1}^{n} \int_{W_i} f_i(\|q-p_i\|) \phi(q)dq \tag{15}$$

with respect to p_j, holding W_i constant for all i, is

$$\begin{aligned} &\nabla_{p_j} g(p_1, \ldots, p_n) \\ &= \nabla_{p_j} \sum_{i=1}^{n} \int_{W_i} f_i(\|q-p_i\|) \phi(q)dq \\ &= \int_{W_j} \nabla_{p_j} f_j(\|q-p_j\|) \phi(q)dq \end{aligned} \tag{16}$$

Remarkably, (16) still holds if each W_i is a Voronoi cell, hence depends on the location of p_i and its Voronoi neighbors [1], [2]. In other words, we need not consider the effect of p_j on the partition $\mathcal{W}_1, \ldots, \mathcal{W}_n$ when computing the gradient. We note that there are only two differences between (14) and (16). In particular, (16) lacks the product term

$$\prod_{\substack{i=1\\ i\neq j}}^{n} f_i(\|q-p_i\|)$$

and integrates only over W_j and not over all of $\mathcal{W}$. Also note that—in contrast to (14)—the gradient (16) depends only on the location of p_j and its Voronoi neighbors, hence can be computed in a distributed way.

C. Gradient of Robust Cost

Now denote by g the robust cost given in Equation 12. The gradient of

$$\begin{aligned} g(p_1, \ldots, p_n) = \sum_{\mathcal{F}\subseteq\{1\ldots n\}} \int_{\mathcal{W}} \left(\prod_{i\in\overline{\mathcal{F}}} f_i(\|q-p_i\|) \right) \phi(q)dq \\ \times p_{\text{fail}}^{|\mathcal{F}|}(1-p_{\text{fail}})^{n-|\mathcal{F}|} \end{aligned} \tag{17}$$

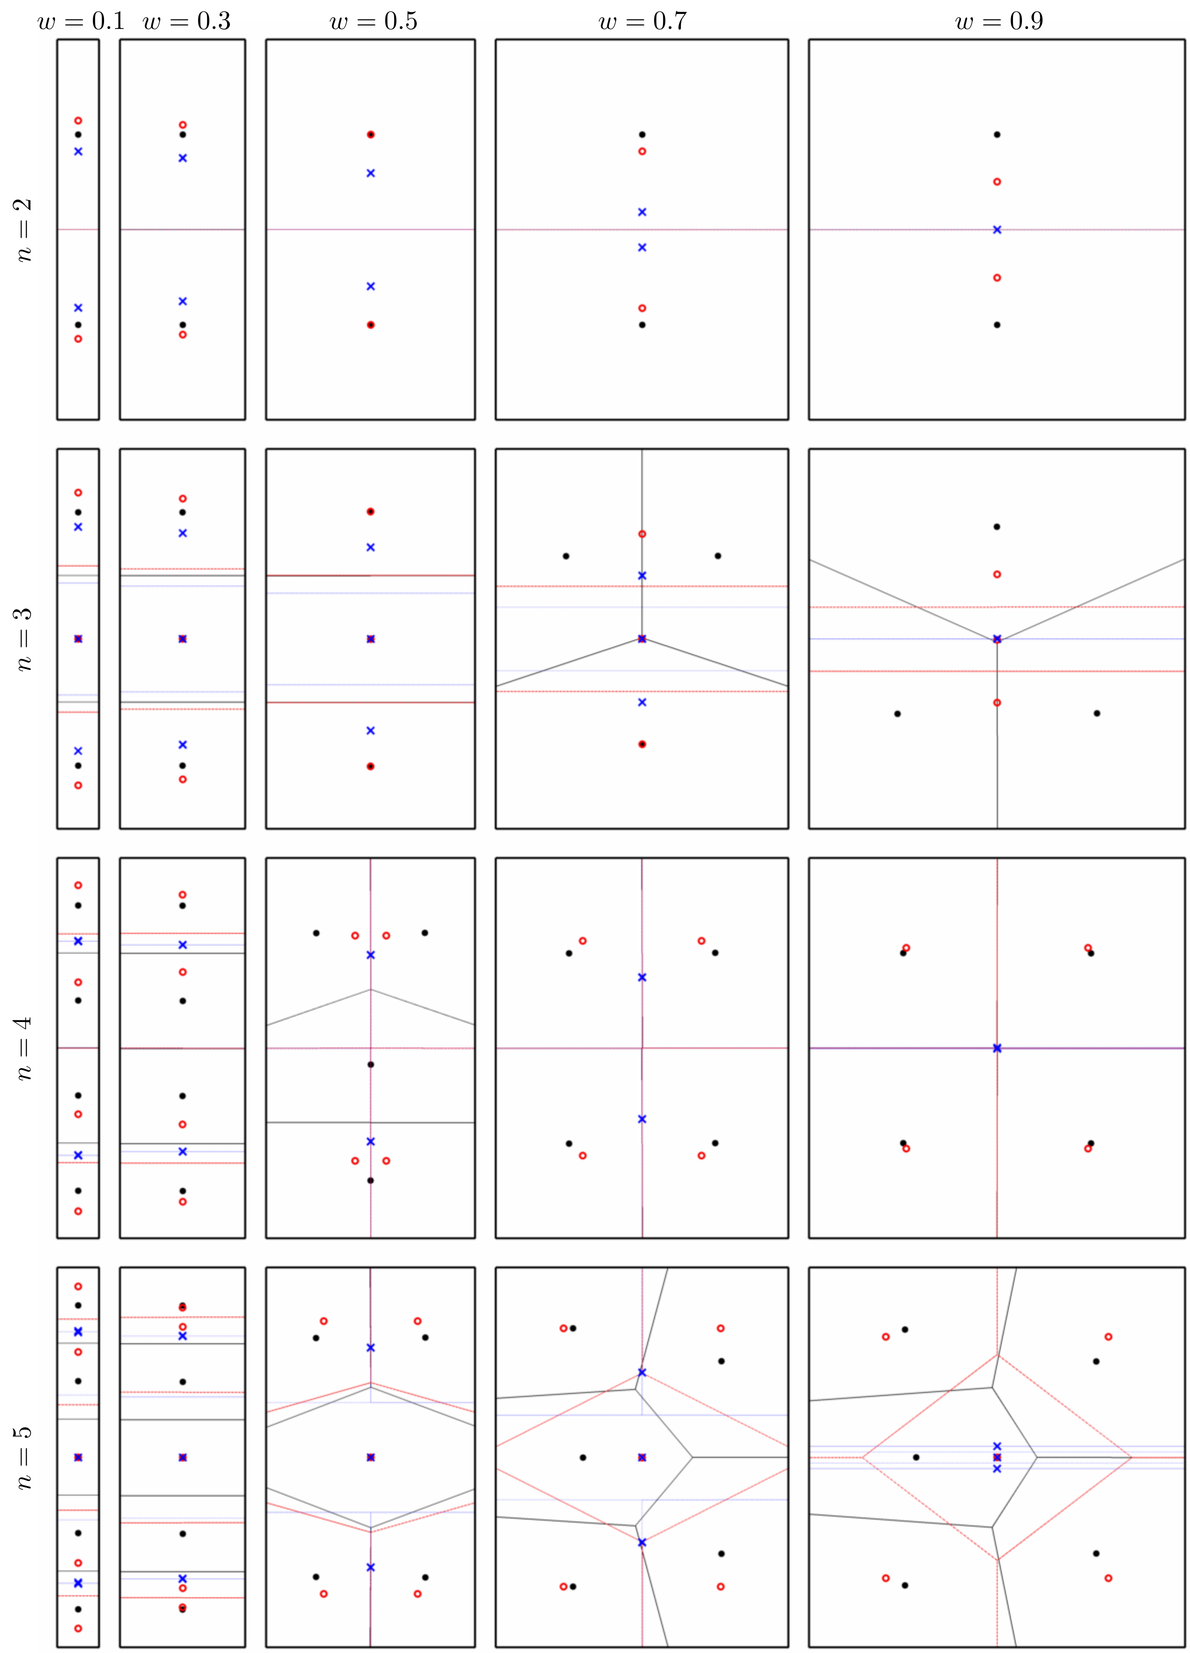

Fig. 1. Deployment of sensors for target detection in a rectangular workspace as the number n of sensors and the width w of the workspace are varied. The target is drawn from a uniform distribution over the workspace, i.e., $\phi(q) = 1/w$ for all q. The deployment algorithm is gradient descent. The resulting deployment is locally optimal with respect to the total cost (red circles) given by Eq. (6), the partitioned cost (black dots) given by Eq. (9), and the robust total cost (blue crosses) given by Eq. (12). Recall that these costs correspond to the probability of missed detection, the probability of missed detection given that sensors see only their own Voronoi partition, and the probability of missed detection when sensors fail independently with probability p_{fail}. It is clear from this example that what is locally optimal for each model of cost can be quite different.

with respect to p_j is

$$\nabla_{p_j} g(p_1, \ldots, p_n) = \sum_{\mathcal{F} \subseteq \{1 \ldots n\} \setminus \{j\}} \int_{\mathcal{W}} \nabla_{p_j} f_j(\|q - p_j\|) \times \left(\prod_{\substack{i \in \overline{\mathcal{F}} \\ i \neq j}} f_i(\|q - p_i\|) \right) \phi(q) dq \times p_{\text{fail}}^{|\mathcal{F}|} (1 - p_{\text{fail}})^{n - |\mathcal{F}|}. \quad (18)$$

The summation in (18) has a combinatorial number of terms, so it is not practical to compute this gradient for large n.

VIII. EXAMPLES

In this section, we show results in simulation for deployment based on gradient descent (Fig. 1). In our examples, the workspace is simply $\mathcal{W} = [0, 1] \times [0, w] \subset \mathbb{R}^2$. The target is drawn from a uniform distribution over $\mathcal{W}$, so $\phi(q) = 1/w$ for all $q \in [0, 1] \times [0, w]$. We assume

$$f_i(\|q - p_i\|) = \eta \|q - p_i\|^2,$$

for $i \in \{1, \ldots, n\}$, where η is chosen so that $f_i(\|q - p_i\|) \leq 1$ for all $q, p_i \in \mathcal{W}$. As a consequence, we have

$$\nabla_{p_i} f_i(\|q - p_i\|) = 2\eta (q - p_i)^T$$

for $i \in \{1, \ldots, n\}$. We chose $p_{\text{fail}} = 0.01$. For various choices of n and w, we sampled initial locations for $p_1, \ldots, p_n$ uniformly at random, applied gradient descent with respect to each measure of cost (total, partitioned, or robust), and show results in Fig. 1. Note that there is nothing special about our choice of sensor model—other models (e.g., exponential detection) produce similar results.

These results make clear that optimal sensor locations are different for each measure of cost. For example, the decentralized deployment algorithm that results from minimizing the partitioned cost does not in general produce sensor locations that minimize the total cost. Similarly, the possibility of sensor failure strongly influences what is optimal. In particular, even a small possibility of failure causes the optimal placement of sensors for $w = 0.9$ to be much different from what results from consideration of the total and partitioned cost. In the examples shown here, for which ϕ is the uniform distribution, the sensors tend to move to the middle of the workspace for the robust detection solution. Intuitively, this corresponds to a more conservative sensor configuration, in which the sensors move independently to positions from which they can more effectively observe a larger portion of the environment.

IX. CONCLUSION

In this paper, we have presented a general approach to optimal sensor deployment for the problem of distributed target detection. We have shown that deployment algorithms based on centroidal Voronoi partitions optimize a specialized version of our cost function. Further, we have developed expressions for the probability of missed detection, conditioned on the possibility of failure by multiple sensors in the network. Our results show that the solutions for these three problems can be very different, thus demonstrating that approaches that rely on centroidal Voronoi partitions are not robust with respect to sensor failure.

Our analysis represents the first steps toward a generalized theory of optimal robust sensor deployment. Future research will extend our approach to problems in which non-uniform priors for object location are used, and to problems in which one or more agents in the network may act maliciously.

ACKNOWLEDGMENTS

This material is based on work supported by the National Science Foundation under Grant Nos. 0931871 and 0955088.

REFERENCES

[1] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Trans. Robot. Autom.*, vol. 20, no. 2, pp. 243 – 255, Apr. 2004.

[2] Q. Du, V. Faber, and M. Gunzburger, "Centroidal voronoi tessellations: Applications and algorithms," *SIAM Review*, vol. 41, no. 4, pp. 637–676, 1999.

[3] L. D. Stone, "What's happened in search theory since the 1975 Lanchester Prize?" *Operations Research*, vol. 37, no. 3, pp. pp. 501–506, 1989.

[4] S. J. Benkoski, M. G. Monticino, and J. R. Weisinger, "A survey of the search theory literature," *Naval Research Logistics (NRL)*, vol. 38, no. 4, pp. 469–494, 1991.

[5] D. W. Gage, "Randomized search strategies with imperfect sensors," in *Proceedings of SPIE Mobile Robots VIII*, 1993, pp. 270–279.

[6] T. H. Chung and J. W. Burdick, "Multi-agent probabilistic search in a sequential decision-theoretic framework," in *IEEE Int. Conf. Rob. Aut.*, May 2008, pp. 146–151.

[7] T. Chung, "On probabilistic search decisions under searcher motion constraints," in *Algorithmic Foundation of Robotics VIII*, ser. Springer Tracts in Advanced Robotics. Berlin, Germany: Springer, 2009, vol. 57, pp. 501–516.

[8] E.-M. Wong, F. Bourgault, and T. Furukawa, "Multi-vehicle Bayesian search for multiple lost targets," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, April 2005, pp. 3169–3174.

[9] F. Bourgault, T. Furukawa, and H. Durrant-Whyte, "Optimal search for a lost target in a Bayesian world," in *Field and Service Robotics*, ser. Springer Tracts in Advanced Robotics, S. Yuta, H. Asama, E. Prassler, T. Tsubouchi, and S. Thrun, Eds. Berlin, Germany: Springer, 2006, vol. 24, pp. 209–222.

[10] H. Choset, "Coverage for robotics-A survey of recent results," *Annals of Mathematics and Artificial Intelligence*, vol. 31, pp. 113–126, 2001.

[11] J. Yick, B. Mukherjee, and D. Ghosal, "Wireless sensor network survey," *Computer Networks*, pp. 2292–2330, Aug. 2008.

[12] E. Stump, N. Michael, V. Kumar, and V. Isler, "Visibility-based deployment of robot formations for communication maintenance." in *Proc. IEEE Int. Conf. on Robotics and Automation*, 2011, to appear.

[13] O. Tekdas, W. Yang, and V. Isler, "Robotic routers: Algorithms and implementation," *Int. J. of Robotics Research*, vol. 29, no. 1, 2010.

[14] S. Poduri, S. Pattem, B. Krishnamachari, and G. S. Sukhatme, "Using local geometry for tunable topology control in sensor networks," *IEEE Trans. Mobile Comput.*, vol. 8, no. 2, pp. 218–230, Feb 2009.

[15] A. Kansal, W. Kaiser, G. Pottie, M. Srivastava, and G. S. Sukhatme, "Reconfiguration methods for mobile sensor networks," *ACM Transactions on Sensor Networks*, vol. 3, no. 4, pp. 22–23, 2007.

[16] M. Zhong and C. Cassandras, "Distributed coverage control and data collection with mobile sensor networks," *IEEE Trans. Autom. Control*, vol. 56, no. 10, pp. 2445–2455, Oct. 2011.

[17] J. Cortés, "Deployment of an unreliable robotic sensor network for spatial estimation," in *IEEE Conf. Dec. Cont.*, Dec. 2010, pp. 376–381.