

Escaping Local Optima in a Class of Multi-Agent Distributed Optimization Problems: A Boosting Function Approach

Xinmiao Sun
Division of Systems Engineering
Boston University

Christos G. Cassandras
Division of Systems Engineering
Boston University

Kagan Gokbayrak
Industrial Engineering Department
Bilkent University

Abstract—We address the problem of multiple local optima commonly arising in optimization problems for multi-agent systems, where objective functions are nonlinear and nonconvex. For the class of coverage control problems, we propose a systematic approach for escaping a local optimum, rather than randomly perturbing controllable variables away from it. We show that the objective function for these problems can be decomposed to facilitate the evaluation of the local partial derivative of each node in the system and to provide insights into its structure. This structure is exploited by defining “boosting functions” applied to the aforementioned local partial derivative at an equilibrium point where its value is zero so as to transform it in a way that induces nodes to explore poorly covered areas of the mission space until a new equilibrium point is reached. The proposed boosting process ensures that, at its conclusion, the objective function is no worse than its pre-boosting value. However, the global optima cannot be guaranteed. We define three families of boosting functions with different properties and provide simulation results illustrating how this approach improves the solutions obtained for this class of distributed optimization problems.

I. INTRODUCTION

Multi-agent systems involve a team of agents (e.g., vehicles, robots, sensor nodes) that cooperatively perform one or more tasks in a mission space which may contain uncertainties such as unexpected obstacles or random event occurrences. The agents communicate, usually wirelessly and over limited ranges, so there are constraints on the information they can exchange. Optimization problems are often formulated in the context of such multi-agent systems and, more often than not, they involve nonlinear, nonconvex objective functions resulting in solutions where global optimality cannot be easily guaranteed. The structure of the objective function can sometimes be exploited, as in cases where it is additive over functions associated with individual agents; for example, in [1], a sum of local nonconvex objective functions is minimized over nonconvex constraints using an approximate dual sub-gradient algorithm. In many problems of interest, however, such an additive structure is not appropriate, as in coverage control or active sensing [2]–[5] where a set of agents must be positioned so as to cooperatively maximize a given objective function. Communication costs and constraints imposed on multi-agent systems, as well as the need to avoid single-point-of-failure

issues, motivate *distributed* optimization schemes allowing agents to achieve optimality, each acting autonomously and with as little information as possible.

Nonconvex environments for coverage control are treated in [6]–[9]. In [3], [8], [10], algorithms concentrate on Voronoi partitions of the mission space and the use of Lloyd’s algorithm. We point out that partition-based algorithms do not take into account the fact that the coverage performance can be improved by sharing observations made by several nodes. This is illustrated by a simple example in Figures. 1–2 where a common objective function when a Voronoi partition is used is outperformed by a distributed gradient-based approach which optimally positions nodes with overlapping sensor ranges (darker-colored areas indicate better coverage).

The nonconvexity of objective functions motivates us to seek systematic methods to overcome the presence of multiple local optima in multi-agent optimization problems. For off-line centralized solutions, one can resort to global optimization algorithms that are typically computationally burdensome and time-consuming. However, for on-line distributed algorithms, this is infeasible; thus, one normally seeks methods through which controllable variables escape from local optima and explore the search space of the problem aiming at better equilibrium points and, ultimately, a globally optimal solution. In gradient-based algorithms, this is usually done by randomly perturbing controllable variables away from a local optimum, as in, for example, simulated annealing [11], [12]. However, in practice, it is infeasible for agents to perform such a random search which is notoriously slow and computationally inefficient. In [13], a “ladybug exploration” strategy is applied to an adaptive controller which aims at balancing coverage and exploration. This approach allows only two movement directions, thus limiting the ability of agents to explore a larger fraction of the mission space. In [9], a gradient-based algorithm was developed to maximize the joint detection probability in a mission space with obstacles. Recognizing the problem of multiple local optima, a method was proposed to balance coverage and exploration by modifying the objective function and assigning a higher reward to points with lower values of the joint event detection probability metric.

In this paper, we propose a systematic approach for coverage optimization problems that moves nodes to locations with potentially better performance, rather than randomly perturbing them away from their current equilibrium. This is accomplished by exploiting the structure of the problem

The authors’ work is supported in part by NSF under grant CNS-1239021, by AFOSR under grant FA9550-12-1-0113, by ONR under grant N00014-09-1-1051, and by ARO under Grant W911NF-11-1-0227. xmsun@bu.edu, cgc@bu.edu, kgokbayr@bilkent.edu.tr

considered. In particular, we focus on the class of optimal coverage control problems. Our first contribution is to show that each node can decompose the objective function into a local objective function dependent on this node's controllable position and a function independent of it. This facilitates the evaluation of the local partial derivative and provides insights into its structure which we subsequently exploit. The second contribution is the development of a systematic method to escape local optima through "boosting functions" applied to the aforementioned local partial derivative. The main idea is to alter the local objective function whenever an equilibrium is reached. A boosting function is a transformation of the associated local partial derivative which takes place at an equilibrium point, where its value is zero; the result of the transformation is a non-zero derivative, which, therefore, forces a node to move in a direction determined by the boosting function and explore the mission space. When a new equilibrium point is reached, we revert to the original objective function and the gradient-based algorithm converges to a new (potentially better and never worse) equilibrium point. We define three families of boosting functions and discuss their properties.

In Sec. II, we formulate the optimization problem and review the distributed gradient-based method developed in [9]. In Sec. III, we derive the local objective function associated with a node and its derivative. In Sec. IV, we introduce the boosting function approach and three families of boosting functions with different properties. Sec. V provides simulation results illustrating how this approach improves the objective function value and we conclude with Sec. VI.

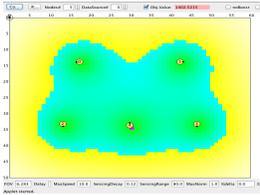


Fig. 1: Gradient-based algorithm; optimal obj.function = 1388.1

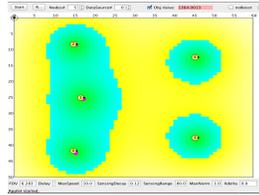


Fig. 2: Voronoi partition; optimal obj. function = 1346.5

II. PROBLEM FORMULATION AND DISTRIBUTED OPTIMIZATION SOLUTION

We concentrate on the optimal coverage control problem. A *mission space* $\Omega \subset \mathbb{R}^2$ is modeled as a non-self-intersecting polygon, i.e., a polygon such that any two non-consecutive edges do not intersect. For any $x \in \Omega$, the function $R(x) : \Omega \rightarrow \mathbb{R}$ describes some a priori information associated with Ω . When the problem is to detect random events that may take place in Ω , this function captures an a priori estimate of the frequency of such event occurrences and is referred to as an *event density* satisfying $R(x) \geq 0$ for all $x \in \Omega$ and $\int_{\Omega} R(x) dx < \infty$. The mission space may contain obstacles modeled as m non-self-intersecting polygons denoted by M_j , $j = 1, \dots, m$ which block the movement of agents. The interior of M_j is denoted by $\overset{\circ}{M}_j$ and the overall *feasible space* is $F = \Omega \setminus (\overset{\circ}{M}_1 \cup \dots \cup \overset{\circ}{M}_m)$, i.e., the space Ω excluding all interior points of the obstacles. There

are N agents in the mission space and their positions at time t are defined by $s_i(t)$, $i = 1, \dots, N$ with an overall position vector $\mathbf{s}(t) = (s_1(t), \dots, s_N(t))$. Figure 3 shows a mission space with two obstacles and an agent located at s_i . The agents may communicate with each other, but there is generally a limited communication range so that it is customary to represent such a system as a network of nodes with a link (i, j) defined so that nodes i, j can communicate directly with each other. This limited communication and the overall cost associated with it are major motivating factors for developing distributed schemes to allow agents to operate so as to optimally achieve a given objective with each acting as autonomously as possible.

In a coverage control problem (e.g., [9], [7], [3]), the agents are *sensor nodes*. We assume that each such node has a bounded sensing range captured by the *sensing radius* δ_i . Thus, the sensing region of node i is $\Omega_i = \{x : d_i(x) \leq \delta_i\}$ where $d_i(x) = \|x - s_i(t)\|$. The presence of obstacles inhibits the sensing ability of a node, which motivates the definition of a *visibility set* $V(s_i) \subset F$ (we omit the explicit dependence of s_i on t for notational simplicity). A point $x \in F$ is *visible* from $s_i \in F$ if the line segment defined by x and s_i is contained in F , i.e., $[\lambda x + (1 - \lambda)s_i] \in F$ for all $\lambda \in [0, 1]$, and x can be sensed, i.e. $x \in \Omega_i$. Then, $V(s_i) = \Omega_i \cap \{x : [\lambda x + (1 - \lambda)s_i] \in F\}$ is a set of points in F which are visible from s_i . We also define $\bar{V}(s_i) = F \setminus V(s_i)$ to be the *invisibility set* (e.g., the grey area in Fig. 3).

A sensing model for any node i is given by the probability that i detects an event occurring at $x \in V(s_i)$, denoted by $p_i(x, s_i)$. We assume that $p_i(x, s_i)$ is expressed as a function of $d_i(x) = \|x - s_i\|$ and is monotonically decreasing and differentiable in $d_i(x)$. An example of such a function is $p_i(x, s_i) = p_{0i} e^{-\lambda_i \|x - s_i\|}$. For points that are invisible by node i , the detection probability is zero. Thus, the overall *sensing detection probability* is denoted as $\hat{p}_i(x, s_i)$ and defined as

$$\hat{p}_i(x, s_i) = \begin{cases} p_i(x, s_i) & \text{if } x \in V(s_i) \\ 0 & \text{if } x \in \bar{V}(s_i) \end{cases} \quad (1)$$

Note that $\hat{p}_i(x, s_i)$ is not a continuous function of s_i . We may now define the *joint detection probability* that an event at $x \in \Omega$ is detected by at least one of the N nodes:

$$P(x, \mathbf{s}) = 1 - \prod_{i=1}^N [1 - \hat{p}_i(x, s_i)] \quad (2)$$

where we assume that detection events by nodes are independent. Finally, assuming that $R(x) = 0$ for $x \notin F$, we define the optimal coverage control problem to maximize $H(\mathbf{s})$, where

$$H(\mathbf{s}) = \int_F R(x) P(x, \mathbf{s}) dx \quad (3)$$

such that $s_i \in F, i = 1, \dots, N$. Thus, we seek to control the node position vector $\mathbf{s} = (s_1, \dots, s_N)$ to maximize the overall joint detection probability of events taking place in the environment. Note that this is a nonlinear, generally nonconvex, optimization problem and the objective function $H(\mathbf{s})$ cannot be expressed in an additive form such as $\sum_{i=1}^N H_i(\mathbf{s})$.

As already mentioned, it is highly desirable to develop distributed algorithms to solve (3) so as to (i) limit costly

communication among nodes and (ii) impart robustness to the system as a whole by avoiding single-point-of-failure issues. Towards this goal, a distributed gradient-based algorithm was developed in [9] based on the iterative scheme:

$$s_i^{k+1} = s_i^k + \zeta_k \frac{\partial H(\mathbf{s})}{\partial s_i^k}, \quad k = 0, 1, \dots \quad (4)$$

where the step size sequence $\{\zeta_k\}$ is appropriately selected (see [14]) to ensure convergence of the resulting node trajectories. If nodes are mobile, then (4) can be interpreted as a motion control scheme for the i th node. In general, a solution through (4) can only lead to a local maximum and it is easy to observe that many such local maxima result in poor performance [9] (we will show such examples in Section V).

Our approach in what follows is to first show that $H(\mathbf{s})$ can be decomposed into a “local objective function” $H_i(\mathbf{s})$ and a function independent of s_i so that node i can locally evaluate its partial derivative with respect to its own controllable position through $H_i(\mathbf{s})$ alone. Our idea then is to alter $H_i(\mathbf{s})$ after a local optimum is attained when $\frac{\partial H_i(\mathbf{s})}{\partial s_i} = 0$, and to define a new objective function $\hat{H}_i(\mathbf{s})$. By doing so, we force $\frac{\partial \hat{H}_i(\mathbf{s})}{\partial s_i} \neq 0$, therefore, node i can “escape” the local optimum and explore the rest of the mission space in search of a potentially better equilibrium point. Because of the structure of $\frac{\partial H_i(\mathbf{s})}{\partial s_i}$ and the insights it provides, however, rather than explicitly altering $H_i(\mathbf{s})$ we instead alter $\frac{\partial H_i(\mathbf{s})}{\partial s_i}$ through what we refer to in Section IV as a *boosting function*.

III. LOCAL OBJECTIVE FUNCTIONS FOR DISTRIBUTED GRADIENT-BASED ALGORITHMS

We begin by defining B_i to be a set of neighbor nodes with respect to i : $B_i = \{k : \|s_i - s_k\| < 2\delta_i, k = 1, \dots, N, k \neq i\}$. Clearly, this set includes all nodes k whose sensing region Ω_k has a nonempty intersection with Ω_i . Accordingly, given that there is a total number of N nodes, we define a complementary set $C_i = \{k : k \notin B_i, k = 1, \dots, N, k \neq i\}$. In addition, let $\Phi_i(x)$ denote the joint probability that a point $x \in \Omega$ is not detected by any neighbor node of i , defined as

$$\Phi_i(x) = \prod_{k \in B_i} [1 - \hat{p}_k(x, s_k)] \quad (5)$$

Similarly, let $\bar{\Phi}_i(x)$ denote the probability that a point $x \in \Omega$ is not covered by nodes in C_i : $\bar{\Phi}_i(x) = \prod_{j \in C_i} [1 - \hat{p}_j(x, s_j)]$. The following theorem establishes the decomposition of $H(\mathbf{s})$ into a function dependent on s_i , for any $i = 1, \dots, N$, and one dependent on all other node positions except s_i .

Theorem 1: The objective function $H(\mathbf{s})$ can be written as: $H(\mathbf{s}) = H_i(\mathbf{s}) + \bar{H}(\bar{\mathbf{s}}_i)$ for any $i = 1, \dots, N$, where $\bar{\mathbf{s}}_i = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_N)$ and $H_i(\mathbf{s}) = \int_{V(s_i)} R(x) \Phi_i(x) p_i(x, s_i) dx$, $\bar{H}(\bar{\mathbf{s}}_i) = \int_F R(x) \{1 - \prod_{k=1, k \neq i}^N [1 - \hat{p}_k(x, s_k)]\} dx$. The proof is omitted here but can be found in [15].

We refer to $H_i(\mathbf{s})$ as the *local objective function* of node i and observe that it depends on $V(s_i)$, $p_i(x, s_i)$, and $\Phi_i(x)$ which are all available to node i (the latter through some communication with nodes in B_i). This result enables a distributed gradient-based optimization solution approach with

each node evaluating $\frac{\partial H_i(\mathbf{s})}{\partial s_i} = (\frac{\partial H_i(\mathbf{s})}{\partial s_{ix}}, \frac{\partial H_i(\mathbf{s})}{\partial s_{iy}})$. We proceed to derive this derivative using the same method as in [16]. Based on the extension of the Leibnitz rule [17], we get

$$\begin{aligned} \frac{\partial H_i(\mathbf{s})}{\partial s_{ix}} &= \frac{\partial}{\partial s_{ix}} \int_{V(s_i)} R(x) \Phi_i(x) p_i(x, s_i) dx \\ &= \int_{V(s_i)} R(x) \Phi_i(x) \frac{\partial p_i(x, s_i)}{\partial s_{ix}} dx \\ &\quad + \int_{\partial V(s_i)} R(x) \Phi_i(x) p_i(x, s_i) (u_x dx_y - u_y dx_x) \end{aligned} \quad (6)$$

where (u_x, u_y) illustrates the “velocity” vector at a boundary point $x = (x_x, x_y)$ of $V(s_i)$. The first term, denoted by E_{ix} , is

$$\begin{aligned} E_{ix} &= \int_{V(s_i)} R(x) \Phi_i(x) \frac{\partial p_i(x, s_i)}{\partial s_{ix}} dx \\ &= \int_{V(s_i)} R(x) \Phi_i(x) \left[-\frac{dp_i(x, s_i)}{dd_i(x)} \right] \frac{(x - s_i)_x}{d_i(x)} dx \end{aligned} \quad (7)$$

where $(x - s_i)_x$ is the x component of the vector $(x - s_i)$. Similarly, we can obtain an integral E_{iy} . Let $E_i = (E_{ix}, E_{iy})$. The integrand of E_i can be viewed as a weighted normalized direction vector $\frac{(x - s_i)}{d_i(x)}$ connecting s_i to $x \in F$ where x is visible by the i th node. The weight is

$$w_1(x, s_i) = -R(x) \Phi_i(x) \frac{dp_i(x, s_i)}{dd_i(x)} \quad (8)$$

Observe that $w_1(x, s_i) \geq 0$ because $\frac{dp_i(x, s_i)}{dd_i(x)} < 0$ since $p_i(x, s_i)$ is a decreasing function of d_i .

Next, we evaluate the second term in (6), referred to as E_b . This evaluation is more elaborate and requires some additional notation (see Fig. 3). Let v be a reflex vertex (definition can be found in [9]) of an obstacle and let $x \in F$ be a point visible from v . A set of points $I(v, x)$, which is a ray starting from v and extending in the direction of $v - x$, is defined by $I(v, x) = \{q \in V(v) : q = \lambda v + (1 - \lambda)x, \lambda > 1\}$. The ray intersects the boundary of F at an *impact* point.

An *anchor* of s_i is a reflex vertex v such that it is visible from s_i and $I(v, s_i)$ is not empty. Denote the anchors of s_i by v_{ij} , $j = 1, \dots, Q(s_i)$, where $Q(s_i)$ is the number of anchors of s_i . An *impact* point of v_{ij} , denoted by V_{ij} , is the intersection of $I(v_{ij}, s_i)$ and ∂F . As an example, in Fig. 3, v_{i1} , v_{i2} , v_{i3} are anchors of s_i , and V_{i1} , V_{i2} , V_{i3} are the corresponding impact points. Let $D_{ij} = \|s_i - v_{ij}\|$ and $d_{ij} = \|V_{ij} - v_{ij}\|$. Define θ_{ij} to be the angle formed by $s_i - v_{ij}$ and the x -axis, which satisfies $\theta_{ij} \in [0, \pi/2]$, that is, $\theta_{ij} = \arctan \frac{|s_i - v_{ij}|_y}{|s_i - v_{ij}|_x}$. Using this notation, a detailed derivation of the second term in (6) may be found in [16] with the final result being:

$$\begin{aligned} E_{bx} &= \sum_{j \in \Gamma_i} \text{sgn}(n_{jx}) \frac{\sin \theta_{ij}}{D_{ij}} \int_0^{z_{ij}} R(\rho_{ij}(r)) \Phi_i(\rho_{ij}(r)) p_i(\rho_{ij}(r), s_i) r dr \end{aligned} \quad (9)$$

where $\Gamma_i = \{j : D_{ij} < \delta_i, j = 1, \dots, Q(s_i)\}$; $z_{ij} = \min(d_{ij}, \delta_i - D_{ij})$ and $\rho_{ij}(r)$ is the Cartesian coordinate of a point on I_j which is a distance r from v_{ij} : $\rho_{ij}(r) = (V_{ij} - v_{ij}) \frac{r}{d_{ij}} + v_{ij}$. In the same way, we can also obtain E_{by} . Note that $E_b = (E_{bx}, E_{by})$ is the gradient component in (6) due to points on

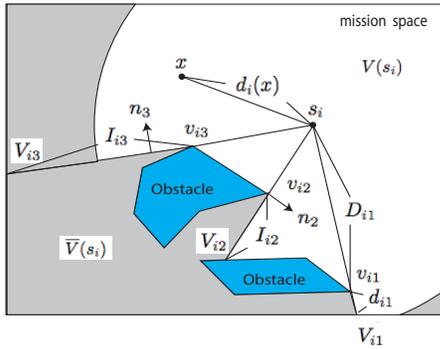


Fig. 3: Mission space with two polygonal obstacles

the boundary $\partial V(s_i)$. In particular, for each boundary, this component attracts node i to move in a direction perpendicular to the boundary and pointing towards $V(s_i)$. We can see in (9) that every point x written as $\rho_{ij}(r)$ in the integrand has an associated weight which we define as $w_2(x, s_i)$:

$$w_2(x, s_i) = R(x)\Phi_i(x)p_i(x, s_i) \quad (10)$$

and observe that $w_2(x, s_i) \geq 0$, as was the case for $w_1(x, s_i)$. Combining (7), (9), (8) and (10), we finally obtain the derivative of $H_i(\mathbf{s})$ with respect to s_i :

$$\begin{aligned} \frac{\partial H_i(\mathbf{s})}{\partial s_{ix}} &= \int_{V(s_i)} w_1(x, s_i) \frac{(x-s_i)_x}{d_i(x)} dx \\ &+ \sum_{j \in \Gamma_i} \text{sgn}(n_{jx}) \frac{\sin \theta_{ij}}{D_{ij}} \int_0^{z_{ij}} w_2(\rho_{ij}(r), s_i) r dr \end{aligned} \quad (11)$$

$\frac{\partial H_i(\mathbf{s})}{\partial s_{iy}}$ can be obtained similarly. We can see that the essence of each derivative is captured in the weights $w_1(x, s_i)$, $w_2(x, s_i)$. In the first integral, $w_1(x, s_i)$ controls the mechanism through which node i is attracted to different points $x \in V(s_i)$. If obstacles are present, then $w_2(x, s_i)$ in the second integral controls the attraction that boundary points exert on node i with the geometrical features of the mission space contributing through n_{jx} , n_{jy} , θ_{ij} , and D_{ij} . This viewpoint motivates the boosting function approach described next.

IV. THE BOOSTING FUNCTION APPROACH

As defined in (3), this nonlinear, generally nonconvex, optimization problem may have multiple local optima to which a gradient-based algorithm may converge. When we apply a distributed optimization algorithm based on $\frac{\partial H_i(\mathbf{s})}{\partial s_i}$ as described above, any equilibrium point is characterized by $\frac{\partial H_i(\mathbf{s})}{\partial s_i} = 0$. Since node i controls its position based on its local objective function $H_i(\mathbf{s})$, a simple way to “escape” a local optimum \mathbf{s}^1 is to alter $H_i(\mathbf{s})$ by replacing it with some $\hat{H}_i(\mathbf{s}) \neq H_i(\mathbf{s})$ thus forcing $\left. \frac{\partial \hat{H}_i(\mathbf{s})}{\partial s_i} \right|_{\mathbf{s}^1} \neq 0$ and inducing the node to explore the rest of the mission space for potentially better equilibria. Subsequently, when a new equilibrium is reached with node i at $\mathbf{s}_i^1 \neq \mathbf{s}_i^1$ and $\left. \frac{\partial \hat{H}_i(\mathbf{s})}{\partial s_i} \right|_{\mathbf{s}_i^1} = 0$, we can revert to $H_i(\mathbf{s})$, which, in turn will force $\left. \frac{\partial H_i(\mathbf{s})}{\partial s_i} \right|_{\mathbf{s}_i^1} \neq 0$ and the node will seek a new equilibrium at \mathbf{s}_i^2 .

Selecting the proper $\hat{H}_i(\mathbf{s})$ to temporarily replace $H_i(\mathbf{s})$ is not a simple process. However, focusing on $\frac{\partial H_i(\mathbf{s})}{\partial s_i}$ instead of $H_i(\mathbf{s})$ is much simpler due to the nature of the derivatives we derived in (11). In particular, the effect of altering $H_i(\mathbf{s})$ can be accomplished by transforming the weights $w_1(x, s_i)$, $w_2(x, s_i)$ in (11) by “boosting” them in a way that forces $\frac{\partial H_i(\mathbf{s})}{\partial s_i} = 0$ at a local optimum to become nonzero. The net effect is that the attraction exerted by some points $x \in F$ on s_i is “boosted” so as to promote exploration of the mission space by node i in search of better optima.

In contrast to various techniques which aim at randomly perturbing controllable variables away from a local optimum, this approach provides a systematic mechanism for accomplishing this goal by exploiting the structure of the specific optimization problem reflected through the form of the derivatives (11). Specifically, it is clear from these expressions that this can be done by assigning a higher weight (i.e., boosting) to directions in the mission space that provide greater opportunity for exploration and, ultimately “better coverage”. To develop such a systematic approach, we define transformations of the weights $w_1(x, s_i)$, $w_2(x, s_i)$ for interior points and boundary points respectively as follows:

$$\hat{w}_1(x, s_i) = g_i(w_1(x, s_i)) \quad (12)$$

$$\hat{w}_2(x, s_i) = h_i(w_2(x, s_i)) \quad (13)$$

where $g_i(\cdot)$ and $h_i(\cdot)$ are functions of the original weights $w_1(x, s_i)$ and $w_2(x, s_i)$ respectively. We refer to $g_i(\cdot)$ and $h_i(\cdot)$ as *boosting functions* for node $i = 1, \dots, N$. Note that these may be node-dependent and that each node may select the time at which this boosting is done, independent from other nodes. In other words, the boosting operation may also be implemented in distributed fashion, in which case we refer to this process at node i as *self-boosting*.

In the remainder of this paper, we concentrate on functions $g_i(\cdot)$ and $h_i(\cdot)$ which have the form

$$\hat{w}_1(x, s_i) = \alpha_{i1}(x, \mathbf{s})w_1(x, s_i) + \beta_{i1}(x, \mathbf{s}) \quad (14)$$

$$\hat{w}_2(x, s_i) = \alpha_{i2}(x, \mathbf{s})w_2(x, s_i) + \beta_{i2}(x, \mathbf{s}) \quad (15)$$

where $\alpha_{i1}(x, \mathbf{s})$, $\beta_{i1}(x, \mathbf{s})$, $\alpha_{i2}(x, \mathbf{s})$, and $\beta_{i2}(x, \mathbf{s})$ are functions dependent on the point x and the node position vector \mathbf{s} in general. We point out that although the form of (14)-(15) is linear, the functions $\alpha_{ij}(x, \mathbf{s})$, $\beta_{ij}(x, \mathbf{s})$, $j = 1, 2$, $i = 1, \dots, N$ are generally nonlinear in their arguments.

To keep notation simple, let us concentrate on a single node i and omit the subscript i in $\alpha_{ij}(x, \mathbf{s})$, $\beta_{ij}(x, \mathbf{s})$ above. By replacing $w_1(x, s_i)$, $w_2(x, s_i)$ with $\hat{w}_1(x, s_i)$, $\hat{w}_2(x, s_i)$ respectively, we obtain the *boosted derivative* $\frac{\partial \hat{H}(\mathbf{s})}{\partial s_i}$ as follows

$$\begin{aligned} \frac{\partial \hat{H}(\mathbf{s})}{\partial s_{ix}} &= \int_{V(s_i)} \alpha_1(x, \mathbf{s})w_1(x, s_i) \frac{(x-s_i)_x}{d_i(x)} dx \\ &+ \int_{V(s_i)} \beta_1(x, \mathbf{s}) \frac{(x-s_i)_x}{d_i(x)} dx \\ &+ \sum_{j \in \Gamma_i} \text{sgn}(n_{jx}) \frac{\sin \theta_{ij}}{D_{ij}} \int_0^{z_{ij}} \alpha_2(x, \mathbf{s})w_2(x, s_i) r dr \\ &+ \sum_{j \in \Gamma_i} \text{sgn}(n_{jx}) \frac{\sin \theta_{ij}}{D_{ij}} \int_0^{z_{ij}} \beta_2(x, \mathbf{s}) r dr \end{aligned} \quad (16)$$

Obviously, the boosting process (14)-(15) actually changes the objective function $H(\mathbf{s})$. Thus, when a new equilibrium is reached in the boosted derivative phase of system operation, it is necessary to revert to the original objective function by setting $\alpha_1(x, \mathbf{s}) = \alpha_2(x, \mathbf{s}) = 1$ and $\beta_1(x, \mathbf{s}) = \beta_2(x, \mathbf{s}) = 0$.

We summarize the boosting process as follows. Initially, node i uses (11) until an equilibrium \mathbf{s}^1 is reached at τ^1 .

- 1) At $t = \tau^1$, evaluate $H(\mathbf{s}(\tau^1))$ and set $\mathbf{s}^* = \mathbf{s}^1$ and $H^* = H(\mathbf{s}(\tau^1))$. Then, evaluate (16) by using (14)-(15), and iterate on the controllable node position using (4).
- 2) Wait until $\frac{\partial \hat{H}(\mathbf{s})}{\partial s_{ix}} = \frac{\partial \hat{H}(\mathbf{s})}{\partial s_{iy}} = 0$ at time $\hat{\tau}^1 > \tau^1$.
- 3) At $t = \hat{\tau}^1$, set $\alpha_1(x, \mathbf{s}) = \alpha_2(x, \mathbf{s}) = 1$ and $\beta_1(x, \mathbf{s}) = \beta_2(x, \mathbf{s}) = 0$ and revert to $\frac{\partial H(\mathbf{s})}{\partial s_i}$.
- 4) Wait until $\frac{\partial H(\mathbf{s})}{\partial s_{ix}} = \frac{\partial H(\mathbf{s})}{\partial s_{iy}} = 0$ at time $\tau^2 > \hat{\tau}^1$ and evaluate $H(\mathbf{s}(\tau^2))$. If $H(\mathbf{s}(\tau^2)) \geq H^*$, then set $\mathbf{s}^* = \mathbf{s}(\tau^2)$ and $H^* = H(\mathbf{s}(\tau^2))$. Otherwise, \mathbf{s}^* , H^* remain unchanged (if nodes are mobile and have already been moved to $\mathbf{s}(\tau^2)$, then return them to \mathbf{s}^*).
- 5) Either STOP, or repeat the process from the current \mathbf{s}^* with a new or the same boosting function to further explore the mission space for better equilibrium points.

Note that if \mathbf{s}^1 is a global optimum, the boosting process simply perturbs node locations until Step 4 returns them to \mathbf{s}^1 . The process stops if no solution is better than \mathbf{s}^1 after some fixed iterations. It is also possible that there are multiple global optima, in which $H(\mathbf{s}(\tau^2)) = H(\mathbf{s}(\tau^1))$ and the new equilibrium point is equivalent to the original one.

The process above assumes that all nodes wait until they have all reached an equilibrium point \mathbf{s}^1 before each initiates its boosting process. However, this may also be done in a distributed function through a self-boosting process: node i may apply (14)-(15) as soon as it observes $\frac{\partial \hat{H}(\mathbf{s})}{\partial s_{ix}} = \frac{\partial H(\mathbf{s})}{\partial s_{iy}} = 0$.

Boosting function selection: The selection of boosting functions generally depends on the mission space topology. In what follows, we present three families of boosting functions that we have investigated; each has different properties and has provided promising results.

Before proceeding, we make a few observations which guide the selection of boosting functions. First, we exclude cases such that $\alpha_1(x, s_i) = \alpha_2(x, s_i) = C$ independent of x , and $\beta_1(x, s_i) = \beta_2(x, s_i) = 0$. In such cases, the boosting effect is null, since it implies that $\frac{\partial \hat{H}(\mathbf{s})}{\partial s_i} = C \frac{\partial H(\mathbf{s})}{\partial s_i}$, which has no effect on $\frac{\partial H(\mathbf{s})}{\partial s_i} = 0$. Second, we observe that if $|\beta_1(x, s_i)| \gg \alpha_1(x, s_i)w_1(x, s_i)$, then the first integral in (16) is dominated by the second one, and the net effect is that nodes tend to be attracted to a single point instead of exploring the mission space. The third observation is more subtle. The first term of (11) contains information on points of the visible set $V(s_i)$, which is generally more valuable (i.e., more points in $V(s_i)$) than the information in the second term related to the boundary points in Γ_i . Thus, a boosting function should ensure that the first integral in (11) dominates the second when $\frac{\partial H_i(\mathbf{s})}{\partial s_{ix}} \neq 0$. In order to avoid such issues, we will boost $w_1(x, s_i)$ only and set $\alpha_2(x, s_i) = 1$, $\beta_2(x, s_i) = 0$.

P-Boosting function: In this function, $\beta_1(x, \mathbf{s}) = 0$ and we

only concentrate on $\alpha_1(x, \mathbf{s})$ which we set:

$$\alpha_1(x, \mathbf{s}) = kP(x, \mathbf{s})^{-\gamma} \quad (17)$$

where $P(x, \mathbf{s})$ is the joint detection probability defined in (2), γ is a positive integer parameter and k is a gain parameter. Thus, the boosted derivative with this P -boosting function is

$$\begin{aligned} \frac{\partial \hat{H}(\mathbf{s})}{\partial s_{ix}} &= \int_{V(s_i)} kP(x, \mathbf{s})^{-\gamma} w_1(x, s_i) \frac{(x-s_i)_x}{d_i(x)} dx \\ &+ \sum_{j \in \Gamma_i} \text{sgn}(n_{jx}) \frac{\sin \theta_{ij}}{D_{ij}} \int_0^{z_{ij}} w_2(x, s_i) r dr \end{aligned} \quad (18)$$

The motivation for this function is similar to a method used in [9] to assign higher weights for low-coverage interior points in $V(s_i)$, in order for nodes to explore such low coverage areas. This is consistent with the following properties of this boosting function: $(P(x, \mathbf{s}))^{-\gamma} \rightarrow \infty$ as $P(x, \mathbf{s}) \rightarrow 0$, and $(P(x, \mathbf{s}))^{-\gamma} \rightarrow 1$ as $P(x, \mathbf{s}) \rightarrow 1$.

Neighbor-Boosting function: We set $\alpha_1(x, \mathbf{s}) = 1$ and focus on $\beta_1(x, \mathbf{s})$. Every node applies a repelling force on each of its neighbors with the effect being monotonically decreasing with their relative distance. We define:

$$\beta_1(x, \mathbf{s}) = \sum_{j \in B_i} \delta(x - s_j) \frac{k_j}{\|s_i - x\|^\gamma} \quad (19)$$

where $k_j \geq 0$ is a gain parameter which may vary with j , γ is a positive integer, and $\delta(x - s_j)$ is the delta function. The boosted derivative with the neighbor-boosting function is

$$\begin{aligned} \frac{\partial \hat{H}(\mathbf{s})}{\partial s_{ix}} &= \int_{V(s_i)} w_1(x, s_i) \frac{(x-s_i)_x}{d_i(x)} dx \\ &+ \sum_{j \in \Gamma_i} \text{sgn}(n_{jx}) \frac{\sin \theta_{ij}}{D_{ij}} \int_0^{z_{ij}} w_2(x, s_i) r dr \\ &+ \sum_{j \in B_i} \frac{k_j}{\|s_j - s_i\|^{\gamma+1}} (s_j - s_i)_x \end{aligned} \quad (20)$$

Φ -boosting function: This function aims at varying $\alpha_1(x, \mathbf{s})$ by $\Phi_i(x)$ defined in (5), which is the probability that point x is not detected by neighboring nodes of i . $\beta_1(x, \mathbf{s}) = 0$. Large $\Phi_i(x)$ values imply a lower coverage by neighbors, therefore higher weights are set. In particular, we define

$$\alpha_1(x, \mathbf{s}) = k\Phi_i(x)^\gamma \quad (21)$$

where k is a gain parameter and γ is a positive integer parameter. The boosted derivative here is

$$\begin{aligned} \frac{\partial \hat{H}(\mathbf{s})}{\partial s_{ix}} &= \int_{V(s_i)} k\Phi_i(x)^\gamma w_1(x, s_i) \frac{(x-s_i)_x}{d_i(x)} dx \\ &+ \sum_{j \in \Gamma_i} \text{sgn}(n_{jx}) \frac{\sin \theta_{ij}}{D_{ij}} \int_0^{z_{ij}} w_2(x, s_i) r dr \end{aligned} \quad (22)$$

Observe that $\Phi_i(x) = 0$ means that x is well-covered by neighbors of i , so, node i has no incentive to move closer to x . On the other hand, $\Phi_i(x) = 1$ means that no neighbor covers x , so the boosted weight is the value of k .

In this section, we provide a simulation example illustrating how the objective function value in (3) is improved by using the boosting function process and how boosting parameters affect performance.

Figure 4 presents a mission space with a “Room Obstacle” configuration (obstacles shown as blue polygons). There are 10 nodes shown as numbered circles and the event density function is uniform, i.e., $R(x) = 1$. The mission space is colored from dark to lighter as the joint detection probability decreases (the joint detection probability is ≥ 0.97 for purple areas, ≥ 0.50 for green areas, and near zero for white areas). Nodes start from the upper left corner and reach equilibrium configurations with associated objective function value shown in the caption. It is easy to see that the deployment is sub-optimal due to the obvious imbalanced coverage. The upper and lower rightmost “rooms” are poorly covered while there are 4 nodes clustered together near the first obstacle on the left side. We expect that boosting functions can guide nodes towards exploration of poorly covered areas in the mission space, thus leading to a more balanced, possibly globally optimal, equilibrium. Figures.

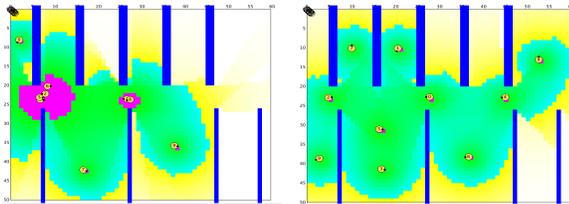


Fig. 4: Room obstacle with $H(s_0^*) = 1183.5$

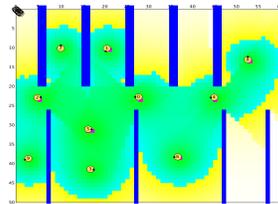


Fig. 5: P-boost, $\gamma = 4$, $k = 100$; $H(s^*) = 1419.5$

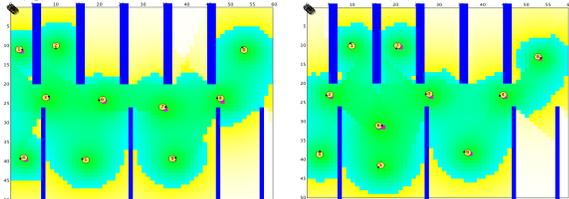


Fig. 6: Neighbor-boost, $\gamma = 1$, $g = 500$; $H(s^*) = 1415.1$

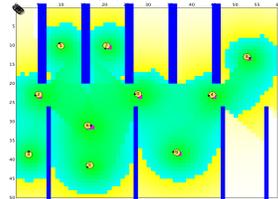


Fig. 7: Φ -boost, $\gamma = 1$, $k = 1000$; $H(s^*) = 1419.1$

5-7 show the effect of the boosting function methods. All boosting functions do increase the objective function values and reach better optima. Comparing with Fig. 4, the clustered nodes in the former have spread apart and the upper and lower rightmost “rooms” are now covered by two nodes. As a result, $H(s)$ increases from 1183.5 to 1419.5. Note that all boosting functions converge to approximately the same results. But this may not be true for different scenarios (See additional examples in [15]). Focusing on the neighbor-boosting function, we select the gain parameters k_j so that $k_j = 0$ for all neighboring nodes except for the closest neighbor of s_i . Figure 6 presents the effect of the neighbor-boosting function with the best parameter pair we have found.

Finally, the boosting function approach converges faster than a random perturbation. Details can be found in [15].

We have shown that the objective function $H(s)$ for the class of optimal coverage control problems in multi-agent system environments can be decomposed into a local objective function $H_i(s)$ for each node i and a function independent of node i 's controllable position s_i . This leads to the definition of boosting functions to systematically (as opposed to randomly) allow nodes to escape from a local optimum so that the attraction exerted by some points on a node i is “boosted” to promote exploration of the mission space by i in search of better optima. We have defined three families of boosting functions, and provided simulation results illustrating their effects and relative performance. Ongoing research aims at combining different boosting functions to create a “hybrid” approach and at studying self-boosting processes whereby individual nodes autonomously control their boosting in a distributed manner.

REFERENCES

- [1] M. Zhu and S. Martinez, “An approximate dual subgradient algorithm for multi-agent non-convex optimization,” *IEEE Transactions on Autom. Control*, vol. 58, no. 6, pp. 1534–1539, 2013.
- [2] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. Srivastava, “Coverage problems in wireless ad-hoc sensor networks,” in *Proc. of 20th Annual Joint Conf. of the IEEE Computer and Commun. Societies*, vol. 3, pp. 1380–1387, 2001.
- [3] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, “Coverage control for mobile sensing networks,” *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.
- [4] L. Mihaylova, T. Lefebvre, H. Bruyninckx, K. Gadeyne, and J. D. Schutter, “Active sensing for robotics - a survey,” in *Proc. 5th Intl Conf. On Numerical Methods and Applications*, pp. 316–324, 2002.
- [5] C. G. Cassandras and W. Li, “Sensor networks and cooperative control,” *European Journal of Control*, vol. 11, no. 4, 2005.
- [6] C. Caicedo-Nuez and M. Zefran, “A coverage algorithm for a class of non-convex regions,” in *Proc. of the 47th IEEE Conf. on Decision and Control*, pp. 4244–4249, 2008.
- [7] C. H. Caicedo-Nunez and M. Zefran, “Performing coverage on non-convex domains,” in *Proc. of the 2008 IEEE Conf. on Control Applic.*, pp. 1019–1024, 2008.
- [8] A. Breitenmoser, M. Schwager, J.-C. Metzger, R. Siegwart, and D. Rus, “Voronoi coverage of non-convex environments with a group of networked robots,” in *Proc. of the 2010 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4982–4989, 2010.
- [9] M. Zhong and C. Cassandras, “Distributed coverage control and data collection with mobile sensor networks,” *IEEE Transactions on Automatic Control*, vol. 56, no. 10, pp. 2445–2455, 2011.
- [10] A. Gusrialdi, S. Hirche, T. Hatanaka, and M. Fujita, “Voronoi based coverage control with anisotropic sensors,” in *Proc. of the 2008 American Control Conf.*, pp. 736–741, 2008.
- [11] P. J. Van Laarhoven and E. H. Aarts, *Simulated annealing*. Springer, 1987.
- [12] D. Bertsimas and J. Tsitsiklis, “Simulated annealing,” *Statistical Science*, pp. 10–15, 1993.
- [13] M. Schwager, F. Bullo, D. Skelly, and D. Rus, “A ladybug exploration strategy for distributed adaptive coverage control,” in *Proc. of the IEEE International Conference on Robotics and Automation*, pp. 2346–2353, 2008.
- [14] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific, 1995.
- [15] X. Sun, C. G. Cassandras, and K. Gokbayrak, “Escaping local optima in a class of multi-agent distributed optimization problems: A boosting function approach,” arXiv:1409.0776v3[cs.SY].
- [16] C. G. Cassandras and M. Zhong, “Distributed coverage control in sensor network environments with polygonal obstacles,” in *Proc. of the 17th IFAC World Congress*, vol. 17, pp. 4162–4167, 2008.
- [17] H. Flanders, “Differentiation under the integral sign,” *The American Mathematical Monthly*, vol. 80, no. 6, pp. 615–627, 1973.