

Target Assignment in Robotic Networks: Distance Optimality Guarantees and Hierarchical Strategies

Jingjin Yu, *Student Member, IEEE*, Soon-Jo Chung, *Senior Member, IEEE*, Petros G. Voulgaris, *Fellow, IEEE*

Abstract—We study the problem of assigning a group of mobile robots to an equal number of distinct static targets, seeking to minimize the total distance traveled by all robots until each target is occupied by a robot. In the first half of our paper, the robots assume limited communication and target-sensing range; otherwise, the robots have no prior knowledge of target locations. Under these assumptions, we present a necessary and sufficient condition under which true distance optimality can be achieved. Moreover, we provide an explicit, non-asymptotic formula for computing the number of robots needed for achieving distance optimality in terms of the robots’ communication and target-sensing ranges with arbitrary guaranteed probabilities. We also show that the same bound is asymptotically tight.

Because a large number of robots is required for guaranteeing distance optimality with high probability, in the second half of our study, we present suboptimal strategies when the number of robots cannot be freely chosen. Assuming that each robot is aware of all target locations, we first work under a hierarchical communication model such that at each hierarchy level, the workspace is partitioned into disjoint regions; robots can communicate with one another if and only if they belong to the same region. This communication model leads naturally to hierarchical strategies, which, under mild assumptions, yield constant approximations of true distance-optimal solutions. We then revisit the range-based communication model and show that combining hierarchical strategies with simple rendezvous-based strategies results in decentralized strategies which again achieve constant approximation ratios on distance optimality. Results from simulation show that the approximation ratio is as low as 1.4.

I. INTRODUCTION

In this paper, we study the permutation-invariant assignment of a set of networked robots to a set of targets of equal cardinality in a planar setting. Focusing on minimizing the total distance traveled by all robots, we seek true optimality guarantees (i.e., sufficient conditions) and suboptimal strategies under various communication and target-sensing capabilities for the robots. For robot-to-robot communication, we investigate a simple circular range-based model as well as a region-based model (in which all robots within the same region can communicate with each other) that leads naturally to hierarchical decision-making processes. For detecting targets, a circular range sensing model is used. In the study of suboptimal strategies, we further establish their parametrized performance characteristics.

This work was supported by AFOSR grant FA95501210193.

Jingjin Yu is with the MIT Computer Science and Artificial Intelligence Lab, Cambridge, MA 02142 USA. E-mail: jingjin@csail.mit.edu. Soon-Jo Chung and Petros G. Voulgaris are with the Coordinated Science Lab and the Department of Aerospace Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA. E-mail: sjchung, voulgaris@illinois.edu

The class of problems that we study is denoted as *target assignment in robotic networks* as it shares many similarities with the problems studied by Smith and Bullo in [25]. In [25], the authors characterized the performance of ETSP¹ ASSGMT and GRID ASSGMT algorithms (strategies) in achieving time optimality (i.e., minimizing the time until every target is occupied). In contrast, we focus on minimizing the total distance traveled by all robots with significant different assumptions on the robots’ models. The total distance serves as a proper proxy to quantities such as the total energy consumption of all robots. Note that a distance-optimal solution for the target assignment problem generally does not imply time optimality and vice versa [31].

As the name implies, an *assignment* (or *matching*) problem is embedded in the problem of target assignment in robotic networks. The assignment problem is extensively studied in the area of combinatorial optimization, with polynomial time algorithms available for solving many of its variations [3], [4], [6], [9], [16], [33]. If we instead put more emphasis on multi-robot systems, the problems of robotic task allocation [14], [26], [27], [32], swarm reconfiguration [7], multi-robot path planning [15], [24], [28], and multi-agent consensus [8], [13], [17], [18] emerge. For a more comprehensive review on these topics, see [5].

Our work is also closely related to the study of connectivity of wireless networks. An interesting result [30] showed that, if n robots are uniformly randomly scattered in a unit square, then each robot needs to communicate with $k = \Theta(\log n)$ nearest neighbors for the entire robotic network to be asymptotically connected as n approaches infinity. In particular, the authors of [30] showed that $k < 0.074 \log n$ leads to an asymptotically disconnected network whereas $k > 5.1774 \log n$ guarantees asymptotic connectivity. This pair of bounds was subsequently improved and extended in [2]. These nearest neighbor based connectivity models were further studied in [11], [12], [19], to list a few. In many of these settings, a *geometric graph* structure is used [22].

In this research effort, we bring forth three contributions. First, for robots with arbitrarily limited range-based target-sensing and communication capabilities (with ranges captured by radii r_{sense} and r_{comm} , respectively), we derive necessary and sufficient conditions for ensuring a distance-optimal solution. In particular, we provide a probabilistic estimate of the number of robots (denoted n) sufficient for all robots to form a connected network given a communication range (some radius

¹ETSP stands for the *Euclidean traveling salesman problem*, in which the distance (weight) between two graph nodes is determined by the physical distance of their embeddings in the plane.

Report Documentation Page		Form Approved OMB No. 0704-0188
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.		
1. REPORT DATE 2013	2. REPORT TYPE	3. DATES COVERED 00-00-2013 to 00-00-2013
4. TITLE AND SUBTITLE Target Assignment in Robotic Networks: Distance Optimality Guarantees and Hierarchical Strategies		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER	
	5e. TASK NUMBER	
	5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Illinois at Urbana-Champaign, Department of Aerospace Engineering, Urbana, IL, 61801		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited		
13. SUPPLEMENTARY NOTES IEEE Transactions on Automatic Control, under review, 2013.		
14. ABSTRACT We study the problem of assigning a group of mobile robots to an equal number of distinct static targets, seeking to minimize the total distance traveled by all robots until each target is occupied by a robot. In the first half of our paper, the robots assume limited communication and target-sensing range otherwise, the robots have no prior knowledge of target locations. Under these assumptions, we present a necessary and sufficient condition under which true distance optimality can be achieved. Moreover, we provide an explicit, non-asymptotic formula for computing the number of robots needed for achieving distance optimality in terms of the robots' communication and target-sensing ranges with arbitrary guaranteed probabilities. We also show that the same bound is asymptotically tight. Because a large number of robots is required for guaranteeing distance optimality with high probability, in the second half of our study, we present suboptimal strategies when the number of robots cannot be freely chosen. Assuming that each robot is aware of all target locations, we first work under a hierarchical communication model such that at each hierarchy level, the workspace is partitioned into disjoint regions; robots can communicate with one another if and only if they belong to the same region. This communication model leads naturally to hierarchical strategies, which, under mild assumptions, yield constant approximations of true distance-optimal solutions. We then revisit the range-based communication model and show that combining hierarchical strategies with simple rendezvous-based strategies results in decentralized strategies which again achieve constant approximation ratios on distance optimality. Results from simulation show that the approximation ratio is as low as 1.4.		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 14	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

r_{comm}). In contrast to connectivity results from [21], [30], we give n as an explicit function of r_{comm} without the asymptotic assumption employed in [21], [30]. Therefore, our bounds do not depend on n being large. We further show that our bound is also asymptotically tight when a high probability guarantee is required.

Second, allowing the robots to have global target-sensing capabilities coupled with a region-based communication model, we adapt results on one-dimensional random walk to show that an infinite family of hierarchical strategies can produce assignments in a decentralized way while simultaneously ensuring that the total distance traveled by the robots is within a constant (asymptotic) bound of the optimal distance. Our simulation results show that this bound can often be smaller than two. Moreover, because hierarchical strategies avoid running a centralized assignment algorithm, significant saving on computation time (in certain cases, a speedup of $1000\times$ or more) is achieved.

Third, for robots with global target-sensing capabilities and a range-based communication model, hierarchical strategies (for assignment) and rendezvous-based strategies (for compensating the lack of global communication) are combined to obtain decentralized suboptimal algorithms. These hybrid strategies, under some mild assumptions, preserve the (asymptotic) constant approximation ratios on distance optimality achieved by the “pure” hierarchical strategies. A result of simulation shows that the cost introduced by the rendezvous requirement becomes negligible as the number of robots increases.

The rest of the paper is organized as follows. In Section II, we introduce notations and well-known results from other branches of research needed for the development of our results. Sections IV–VI then elaborate on the three stated contributions, one contribution per section, in the order as they are given. We conduct simulations in Section VII to confirm our theoretical findings and conclude in Section VIII.

II. PRELIMINARIES

In this section, we review results on balls and bins, linear assignment, and random geometric graphs. The symbols $\mathbb{R}, \mathbb{R}^+, \mathbb{N}$ denote the set of real numbers, the set of positive reals, and the set of positive integers, respectively. For a positive real number x , $\log x$ denotes the natural logarithm of x ; the function $\lceil x \rceil$ (respectively, $\lfloor x \rfloor$) denotes the smallest (respectively, largest) integer that is larger (respectively, smaller) than x . Let A be a set, $|A|$ denotes its cardinality. We use $\|v\|_2$ to denote the Euclidean 2-norm of a vector v . The unit square $[0, 1] \times [0, 1] \subset \mathbb{R}^2$ is denoted as Q . The expectation of a random variable X is denoted as $\mathbf{E}[X]$. We use $E(\cdot)$ to denote a probabilistic event and the probability with which an event e occurs is denoted as $\mathbf{P}(e)$.

Given two functions $f, g: \mathbb{R}^+ \rightarrow \mathbb{R}^+$, $f(x) = O(g(x))$ if and only if there exist $M_O, x_O \in \mathbb{R}^+$ such that

$$\forall x > x_O, |f(x)| \leq M_O |g(x)|.$$

Similarly, $f(x) = \Omega(g(x))$ if and only if there exist $M_\Omega, x_\Omega \in \mathbb{R}^+$ such that

$$\forall x > x_\Omega, |f(x)| \geq M_\Omega |g(x)|.$$

If $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$, then we say $f(x) = \Theta(g(x))$. Finally, $f(x) = o(g(x))$ (respectively, $f(x) = \omega(g(x))$) if and only if $f(x) = O(g(x))$ (respectively, $f(x) = \Omega(g(x))$) and $\neg(f(x) = \Theta(x))$.

A. Balls and bins

The well-studied problem in probability theory known as the *urns-problem*, or the problem of *balls and bins*, considers the distribution generated as a number of balls are randomly tossed into a set of bins. The following classical result on the ball and bins problem is due to Erdős and Rényi.

Theorem 1 (Balls and bins [10]) *Suppose that a number of balls are tossed uniformly randomly into m bins, one ball per time step. Let T_k denote the first time such that k balls are collected in every bin. Then for any real number c ,*

$$\lim_{m \rightarrow \infty} \mathbf{P}(T_k < m \log m + (k-1)m \log \log m + cm) = e^{-e^{-\frac{c}{(k-1)!}}}. \quad (1)$$

It is worth noting that the proof of Theorem 1 is fairly short and elegant, employing only basic tools from analysis and combinatorics. A useful corollary for $k = 1$ readily follows.

Corollary 2 *For an arbitrary real number c , suppose that no fewer than $(m \log m + cm)$ balls are tossed uniformly randomly into m bins. As $m \rightarrow \infty$, every bin contains at least one ball with probability $e^{-e^{-c}}$.*

PROOF. In (1), let $k = 1$ yields

$$\mathbf{P}(T_1 < m \log m + cm) = e^{-e^{-c}}. \quad (2)$$

The corollary directly follows (2) (recall that T_1 is the number of tosses needed so that every bin has at least one ball). $\square$

Corollary 2 says that $T_1 \approx m \log m$ is a “sharp” threshold: Letting $c = 5$ in (2) yields that the probability of every bin being occupied by at least one ball is greater than 0.99 when there are at least $m \log m + 5m$ balls. On the other hand, the same probability is no more than 0.001 when there are $m \log m - 2m$ balls.

B. The assignment problem

The (linear) assignment problem, as a fundamental combinatorial optimization problem, can be defined as follows.

Problem 1 (Linear Assignment Problem) *Given two finite sets X and Y with $|X| = |Y|$, together with a weight function $C: X \times Y \rightarrow \mathbb{R}$, find a bijection $f: X \rightarrow Y$ that minimizes the cost function*

$$\sum_{x \in X} C(x, f(x)). \quad (3)$$

Problem 1 is also called the *perfect weighted bipartite matching* problem. Here, the mapping C is essentially a square matrix that can be used to represent a variety of weights, such as the Euclidean distance when X and Y represents physical locations. The *Hungarian method* for the assignment problem,

proposed by Kuhn [16], has an $O(n^4)$ running time, which was subsequently improved to $O(n^3)$ by Edmonds and Karp [9]. Many other algorithms for the assignment problem exist, including other primal-dual (linear programming) methods [6], auction based methods [3], and parallel algorithms [4], [33]. Nevertheless, the strongly polynomial² $O(n^3)$ Hungarian method remains as the fastest exact (sequential) algorithm, which we use in our simulations.

When X and Y are restricted to be points on the plane with the weight function C being the Euclidean distances between the points, the special linear assignment problem is known as the *Euclidean bipartite matching* problem, which can be solved exactly using an $O(n^{2.5} \log n)$ primal-dual algorithm [29]. Alternatively, near linear time approximation algorithms can yield near optimal solutions with high probability [23].³

C. Random geometric graphs

Let $X = \{x_1, \dots, x_n\}$ be a set of n points in the unit square Q . For a fixed *communication radius* r_{comm} , the *geometric graph* G over this set of points is formed by taking each point as a vertex of the graph and connect any two vertices if the underlying points x_1 and x_2 satisfies $\|x_1 - x_2\|_2 \leq r_{\text{comm}}$. When each $x \in X$ is selected randomly following some distribution, the resulting graph is called a *random geometric graph*.

Properties of random geometric graphs have been studied extensively; see [22] for a thorough coverage. One such property is the connectivity of these graphs, which is of particular interest to wireless communication and robotic networks.

Theorem 3 (Random Geometric Graphs [21]) *Let G be the random geometric graph obtained following the uniform distribution over the unit square for some n and r_{comm} . Then for any real number c , as $n \rightarrow \infty$ ($r_{\text{comm}} \rightarrow 0$),*

$$\mathbb{P}(G \text{ is connected} \mid \pi n r_{\text{comm}}^2 - \log n \leq c) = e^{-e^{-c}}. \quad (4)$$

From (4), it is possible to estimate the number of robots required to guarantee a connected geometric graph G .

III. TARGET ASSIGNMENT IN ROBOTIC NETWORKS

In this section, we formally define the problem of *target assignment in robotic networks* and the optimality objectives.

A. Problem statement


Let $X^0 = \{x_1^0, \dots, x_n^0\}$ and $Y^0 = \{y_1^0, \dots, y_n^0\}$ be two sets of points in the unit square Q (the superscript emphasizes that these points are obtained at the start time $t = 0$), selected uniformly randomly.⁴ Place $n = |X^0| = |Y^0|$ point robots on the points in X^0 , with robot a_i occupying x_i^0 . Each robot has a

unique integer label (e.g., i). In general, we denote robot a_i 's location (coordinates) at time $t \geq 0$ as $x_i(t)$. The basic task (to be formally defined) is to move the robots so that at some *final time* $t^f \geq 0$, every $y \in Y^0$ is occupied by a robot (we may assume that there is a final time t_i^f for each robot a_i , such that $x_i(t) \equiv x_i(t_i^f)$ for $t \geq t_i^f$). For convenience, we also refer to X^0 and Y^0 as the set of initial locations and the set of target locations, respectively.

Motion model: A robot a_i may start and stop moving with unit speed in any direction. The control policy for a_i is then some $\dot{x}_i = u_i$ with $\|u_i\|_2 \in \{0, 1\}$. We assume that robots' sizes are negligible with respect to the distance they travel and ignore collisions between robots.

Communication Model 1: We study two communication models in this paper. In the first communication model, a robot a_i may communicate with other robots within a disc of radius r_{comm} centered at $x_i(t)$. At any given time $t \geq 0$, we define the (undirected) *communication graph* $G(t)$ as follows, which is a geometric graph that changes over time. $G(t)$ has n vertices $v_1, \dots, v_n$, corresponding to robots $a_1, \dots, a_n$, respectively. There is an edge between two vertices v_i and v_j if the corresponding robot locations $x_i(t)$ and $x_j(t)$, respectively, satisfy $\|x_i(t) - x_j(t)\|_2 \leq r_{\text{comm}}$. Figure 1(a) provides an example of a (disconnected) communication graph.

Since communication overhead is often negligible with respect to the time it takes the robots to move, we assume that all robots corresponding to vertices in a connected component of the communication graph may exchange information as needed instantaneously. In other words, robots in a connected component of $G(t)$ can be effectively treated as a single robot insofar as decision making is concerned.



(a) Comm. model 1

(b) Comm. model 2

Fig. 1. (a) The communication graph (solid blue nodes and edges) for a given set of robots under Communication Model 1 with a communication radius of r_{comm} . Robots (blue dots) in the same component can freely communicate. (b) The communication graph for a set of robots under Communication Model 2 with $m = b^2 = 9$.

Communication Model 2: In the second communication model, the unit square Q is divided into some $m = b^2$ equal sized smaller squares (regions).⁵ Robots within each region can communicate with one another but robots from different regions cannot exchange information (see, e.g., Fig. 1(b)). This model mimics the natural (geometrical) resource allocation process in which supplies and demands are first matched

²A polynomial time algorithm runs in *strongly polynomial time* only if its running time does not depend on the size of the input parameters. Note that n is the *number* of input parameters in this case.

³Although algorithms from [23], [29] have theoretically faster running times than the Hungarian method and apply to the problem that we study, they are more difficult to implement and slower in practice unless $|X|$ is very large because they are no longer strongly polynomial time algorithms like the Hungarian method.

⁴Our results readily extend to a square environment with arbitrary side length through scaling.

⁵In this paper, m is frequently used to denote the number of small squares in a division of the unit square Q and b denotes the number of segments a side of the unit square is divided into; the value of m and b may vary depending on the context.

locally; the surpluses and deficits within the each region then get balanced out at larger regions, giving rise to a hierarchical decision process.

Target-sensing model: We assume that a robot is aware of a point $y \in Y^0$ if $\|y - x_i(t)\|_2 \leq r_{\text{sense}}$, the *target sensing radius*.

The problem we consider in this paper is defined as follows.

Problem 2 (Target Assignment in Robotic Networks)

Given X^0 , Y^0 , r_{sense} , and a communication model, find a control strategy $\mathbf{u} = [u_1, \dots, u_n]$, such that for some $0 \leq t_i^f < \infty$ and some permutation σ of the numbers $1, \dots, n$, $x_i(t_i^f) = y_{\sigma(i)}^0$ for all $1 \leq i \leq n$.

Over all feasible solutions to an instance of Problem 2, we are interested in minimizing the total distance traveled by all robots, which can be expressed as

$$D_n = \sum_{i=1}^n \int_0^{t_i^f} \dot{x}_i(t) dt. \quad (5)$$

As an accurate proxy to measures such as the energy consumption of the entire system, the cost defined in (5) is an appropriate objective in practice. Unless otherwise specified, *distance optimality* refers to minimizing D_n . Assuming that robots must follow continuous paths, we let D_n^* denote the best possible D_n , which may or may not be achievable depending on the capabilities of the robots (e.g., if the robots cannot follow straight line paths, then $D_n > D_n^*$). Let $\mathcal{U}$ denote the set of all possible control strategies that solve Problem 2 given a fixed set of capabilities for the robots, $\inf_{\mathcal{U}} D_n$ is then the greatest lower bound achievable under these capabilities.⁶ Besides distance optimality, we also briefly discuss the total task completion time (i.e., the sum of the individual task completion times as targets are occupied), denoted T_n . If all robots start moving toward targets and do not stop in the middle, $T_n = D_n$.

IV. GUARANTEEING DISTANCE OPTIMALITY FOR ARBITRARY r_{comm} AND r_{sense}

In this section, we use Communication Model 1. In general, when $r_{\text{sense}} < \sqrt{2}$ or $r_{\text{comm}} < \sqrt{2}$, it is impossible to guarantee distance optimality (i.e., $\inf_{\mathcal{U}} D_n = D_n^*$), since global assignment is no longer possible in general. For example, as $r_{\text{sense}} \rightarrow 0$, the robots must search for the targets before assignment can be made; it is very unlikely that the paths taken by the robots toward the targets will be straight lines, which is required to obtain D_n^* . This raises the following question: Given a pair of r_{comm} and r_{sense} , under what conditions can we ensure distance optimality? Theorem 4 answers this question.

Theorem 4 *Under sensing and communication constraints (i.e., $r_{\text{comm}}, r_{\text{sense}} < \sqrt{2}$), $\inf_{\mathcal{U}} D_n = D_n^*$ if and only if $G(0)$ is connected and every target $y \in Y^0$ is within a distance of r_{sense} to some $x \in X^0$.*

PROOF. We first prove the necessary conditions with two claims: 1) an optimal assignment that minimizes D_n is possible

⁶Here $\inf$ is used instead of $\min$ because it is not immediately clear that the minimum can always be reached.

in general only if $G(0)$ is connected, and 2) an optimal assignment that minimizes D_n is possible only if for all $y \in Y^0$, y is within a distance of r_{sense} to some $x \in X^0$.

To see that the first claim is true, we note that to get a distance-optimal assignment, the robots cannot make any unnecessary moves. This means that they must be able to decide at $t = 0$ a pairing between elements of X^0 and Y^0 that minimizes D_n . We now show that this is not possible in general for $n = 2$ and $r_{\text{comm}} < \sqrt{2}$. For $n = 2$, assume that the two robots and two targets (a_1, a_2, y_1, y_2 , respectively) are located as given in Fig. 2 (solid blue and red dots). Because they are out of each other's communication radius, the robots cannot communicate with each other. Robot a_1 is of equal distance to y_1 and y_2 whereas robot a_2 is closer to y_2 than it is to y_1 . An optimal assignment requires that a_1 goes to y_1 and a_2 goes to y_2 . Given this setting, it is impossible for a_1 to decide at $t = 0$ to go to y_1 or y_2 without knowing where a_2 is. We can readily extend the locations of the robots and targets to include neighborhoods around them (the dotted circles in Fig. 2) to show that there is non-zero probability with which an optimal assignment cannot be made at $t = 0$. This proves that that $G(0)$ cannot have more than one connected component and must be connected.




Fig. 2. A general setup in which the two robots cannot communicate with each other at $t = 0$ and therefore, cannot decide an optimal assignment at $t = 0$.

For the second claim, suppose that at $t = 0$, some $y \in Y^0$ is not within a distance of r_{sense} to any $x \in X^0$. Then some robot must move to search for that y . This will cause some robot to follow a path that is not a straight line with probability one, implying that $D_n = D_n^*$ with probability zero.

It is not hard to see that the necessary conditions from the two claims are also sufficient: when $G(0)$ is connected and each target is observable by some robot a_i , the robots can decide at $t = 0$ a global assignment that minimizes D_n . $\square$

Theorem 4 suggests a simple way of ensuring distance optimality by either increasing the number of robots or increasing one or both of r_{comm} and r_{sense} , which essentially leads to a centralized communication and control strategy (Strategy 1). Note that given the assignment permutation σ , each robot a_i can easily compute its straight-line path between x_i^0 and $y_{\sigma(i)}^0$. Since every robot can carry out the computation in Strategy 1, to resolve conflicting decisions and avoid unnecessary

computation, we may let the highest labeled robot (e.g., a_n) dictate the assignment process.

Strategy 1: CENTRALIZED ASSIGNMENT

Initial condition: X^0, Y^0

Outcome: permutation σ that assigns robot a_i to $y_{\sigma(i)}^0$

- 1 compute $d_{i,j} = \|x_i - y_j\|_2$ between each pair of (x_i, y_j) in which $x_i \in X^0$ and $y_j \in Y^0$
 - 2 based on $\{d_{i,j}\}$, compute an optimal assignment for the robots that minimizes D_n
 - 3 communicate the assignment to all robots
-

The rest of this section establishes how the conditions from Theorem 4 can be met. We point out that similar conclusions can also be reached by exploring Theorem 3, which indirectly yields an asymptotic estimate of the required number of robots for $G(0)$ to be connected, given an r_{comm} . We take a different approach and produce the number as an explicit function of r_{comm} , without the asymptotic assumption.

A. Guaranteeing a connected $G(0)$

Since the robots can be anywhere in the unit square Q , given a communication radius of $r_{\text{comm}} < \sqrt{2}$, intuitively, at least $\Theta(1/r_{\text{comm}}^2)$ robots are needed for a connected $G(0)$, which requires the robot to take a roughly “regular” formation such as a grid. It turns out that when the robots are randomly distributed, not a great many more than $\Theta(1/r_{\text{comm}}^2)$ robots are needed to ensure a connected communication graph $G(0)$.

Lemma 5 *Given a fixed $r_{\text{comm}} < \sqrt{2}$ and $0 < \varepsilon < 1$, the communication graph $G(0)$ is connected with probability at least $1 - \varepsilon$ if the number of robots n satisfies*

$$n = \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2 \log\left(\frac{1}{\varepsilon} \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2\right). \quad (6)$$

PROOF. We divide the unit square Q into $m = b^2$ equal sized small squares with $b = \lceil \sqrt{5}/r_{\text{comm}} \rceil$. Label these small squares as $\{q_1, \dots, q_m\}$. Under this division scheme, if a small square q_i (see, e.g., the gray one in Fig. 3) contains at least a robot, the robot can communicate with any other robot in the four squares sharing a side with q_i .




Fig. 3. If the small squares have a side length of $\lceil \sqrt{5}/r_{\text{comm}} \rceil$ or smaller, then a robot in such a square (e.g., the gray square) can communicate with any robot in the four neighboring small squares.

If a small square q_i contains a robot a , then a robot in a neighboring square (those that share a side with q_i) is

within r_{comm} of robot a by the assumption. Therefore, $G(0)$ is connected if each q_i contains a robot. Let n_i denote the number of robots in q_i . Then

$$\mathbf{P}(n_i = 0) = (1 - \frac{1}{m})^n < e^{-\frac{n}{m}}.$$

The inequality holds because $(1 - x)^n < e^{-nx}$ for $0 < x < 1$. To see this, let $f(x) = \log(1 - x)/x$. The Taylor expansion of $f(x)$ at $x = 0$ is $-1 - x/2 - x^2/3 + o(x^3) < -1$ for $0 < x < 1$. This shows that $\log(1 - x) < -x$ for $0 < x < 1 \Rightarrow n \log(1 - x) < -nx \Rightarrow (1 - x)^n < e^{-nx}$. By Boole's inequality, the probability that at least one of the squares $q_1, \dots, q_m$ is empty can be upper bounded as

$$\mathbf{P}\left(\bigcup_{i=1}^m E(n_i = 0)\right) \leq \sum_{i=1}^m \mathbf{P}(n_i = 0) < m e^{-\frac{n}{m}}.$$

Setting $m e^{-n/m} = \varepsilon$ and replacing $m = \lceil \sqrt{5}/r_{\text{comm}} \rceil^2$ yields

$$\begin{aligned} \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2 \exp\left(-n \frac{1}{\lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2}\right) &= \varepsilon \\ \Rightarrow n &= \left(\lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2\right) \log\left(\frac{1}{\varepsilon} \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2\right), \end{aligned}$$

which guarantees that each small square contains at least one robot with probability $1 - \varepsilon$. $\square$

We derive (6) in Lemma 5 so that the expression is more uniform (see Theorem 12); it is possible to make n even smaller. The following corollary illustrates one way to tighten this bound.

Corollary 6 *Given a fixed $r_{\text{comm}} < \sqrt{2}$ and $0 < \varepsilon < 1$, the communication graph $G(0)$ is connected with probability at least $1 - \varepsilon$ if the number of robots n satisfies*

$$n = \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2 \log \left[\frac{1}{\varepsilon} \left(\frac{1}{2} \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2 + \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil \right) \right] \quad (7)$$




Fig. 4. As long as each of the colored squares contains a robot, $G(0)$ must be connected. Therefore, only $b^2/2 + b$ squares need to have robots in them.

PROOF. If each of the colored squares in Fig. 4 has at least one robot, then $G(0)$ must be connected: any robot falling in a white square must be connected to some robot in a colored square. This shows that (7) is sufficient. $\square$

Remark. In comparison to Theorem 3, Lemma 5 provides n as an explicit function of r_{comm} . Moreover, our sufficient condition on n given in (6) (and (7)), unlike (4), is not an

asymptotic bound. Therefore, our bound applies to an arbitrary r_{comm} . On the other hand, if we let $r_{\text{comm}} \rightarrow 0$, then an asymptotic statement can also be made.

Lemma 7 *As $r_{\text{comm}} \rightarrow 0$, the communication graph $G(0)$ is connected with arbitrarily high probability $e^{-e^{-c}}$ (for some $c > 0$) if the number of robots n satisfies*

$$n \geq (2 \log \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil + c) \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2. \quad (8)$$

PROOF. Given the division scheme used in the proof of Lemma 5, distributing robots into the unit square Q is equivalent to tossing the robots (balls) into the m small squares (bins), uniformly randomly. By Corollary 2, having $n \geq m \log m + cm = (2 \log \lceil \sqrt{5}/r_{\text{comm}} \rceil + c) \lceil \sqrt{5}/r_{\text{comm}} \rceil^2$ robots guarantees that all m small squares must have at least one robot each with probability $e^{-e^{-c}}$. $\square$

Since $f(x) = cx$ grows slower than $g(x) = x \log x$ as $x \rightarrow \infty$, Lemma 7 says that $n = \Theta((1/r_{\text{comm}})^2 \log(1/r_{\text{comm}}))$ robots can ensure that $G(0)$ is connected with probability arbitrarily close to one asymptotically. Next, we show that this many robots are also necessary for the high probability guarantee.

Let $\mathbf{P}_{n,m}(E)$ denote the probability of event E happening after tossing n balls into m bins. We work with two events: E_0 , the event that “at least one bin has zero balls in it”, and E_1 , the event that “at least one bin contains exactly one ball”. We want to show that $\mathbf{P}_{n,m}(E_1)$ is not too small for n up to $m \log m$, which is proven over the next two lemmas.

Lemma 8 *Suppose that $1 \leq n \leq m$ balls are tossed uniformly randomly into m bins. Then for large m ,*

$$\mathbf{P}_{n,m}(E_1) \geq e^{-1}.$$

PROOF. Because all bins are initially empty, after tossing the first ball, some bin contains exactly one ball. That is, $\mathbf{P}_{1,m}(E_1) = 1$. Let the bin occupied by the first ball be bin 1. As $k-1$ additional balls are tossed into the m bins, the probability that none of these $k-1$ balls occupy bin 1 is $(1-1/m)^{k-1}$. Therefore, for $1 \leq k \leq m$ tosses, we have

$$\begin{aligned} \mathbf{P}_{k,m}(E_1) &\geq \mathbf{P}_{1,m}(E_1) \left(1 - \frac{1}{m}\right)^{k-1} \\ &> \mathbf{P}_{1,m}(E_1) \left(1 - \frac{1}{m}\right)^m \\ &\approx \mathbf{P}_{1,m}(E_1) e^{-1} = e^{-1} \approx 0.37. \end{aligned}$$

The first approximation in the above derivation is due to the function $\log(1-x)^{1/x}$ having a Taylor expansion of $-1 + o(1)$ for small positive x . This is where the assumption of “large m ” is used in this lemma. $\square$

Lemma 9 *Suppose that $m < n < m \log m$ balls are tossed uniformly randomly into m bins. As $m \rightarrow \infty$,*

$$\mathbf{P}_{n,m}(E_1) \geq 0.34.$$

PROOF. Suppose that after an experiment of n' tosses into m bins, E_0 holds; i.e., at least one bin is empty. Without loss

of generality, we assume that the empty bin is bin 1. Now consider tossing an additional k balls into the m bins. The probability of exactly one of these k balls falling in bin 1 is

$$\begin{aligned} &\mathbf{P}_{k,m}(\text{exactly one ball falls in bin 1}) \\ &= \binom{k}{1} \frac{1}{m} \left(1 - \frac{1}{m}\right)^{k-1} = \frac{k}{m} \left(1 - \frac{1}{m}\right)^{k-1}. \end{aligned}$$

Therefore,

$$\begin{aligned} &\mathbf{P}_{n'+k,m}(E_1) \\ &\geq \mathbf{P}_{n',m}(E_0) \mathbf{P}_{k,m}(\text{exactly one ball falls in bin 1}) \\ &\approx \frac{k}{m} \left(1 - \frac{1}{m}\right)^{k-1} \mathbf{P}_{n',m}(E_0). \end{aligned} \quad (9)$$

Let $c = -1$ in Corollary 2, we have

$$\mathbf{P}(T_1 \geq m \log m - m) = 1 - e^{-e^1} > 0.93. \quad (10)$$

That is, as $m \rightarrow \infty$, for $0 < n' < m \log m - m$, $\mathbf{P}_{n',m}(E_0) \geq 0.93$. Plugging this into (9) and let $k = m$, then for $m < n < m \log m$, as $m \rightarrow \infty$,

$$\mathbf{P}_{n,m}(E_1) \geq 0.93 \frac{m}{m} \left(1 - \frac{1}{m}\right)^{m-1} \approx 0.93e^{-1} \approx 0.34. \quad \square$$

We now show that $n = \Theta((1/r_{\text{comm}})^2 \log(1/r_{\text{comm}}))$ is a tight bound on the number of robots for guaranteeing the connectivity of $G(0)$ with high probability.

Theorem 10 *For uniformly randomly distributed robots in a unit square with a communication radius r_{comm} ,*

$$n = \Theta\left(\frac{1}{r_{\text{comm}}^2} \log \frac{1}{r_{\text{comm}}}\right) \quad (11)$$

robots are necessary and sufficient to ensure a connected communication graph at $t = 0$ with arbitrarily high probability as $r_{\text{comm}} \rightarrow 0$.

PROOF. Lemma 7 covers sufficiency; we are to show that there is some non-trivial probability that $G(0)$ is disconnected if the number of robots satisfies

$$n = o\left(\frac{1}{r_{\text{comm}}^2} \log \frac{1}{r_{\text{comm}}}\right).$$

To prove the claim, we partition the unit square Q into $m = b^2$ equal sized small squares where $b = \lfloor 1.1/r_{\text{comm}} \rfloor$. The factor of 1.1 in the expression makes the side of the small square larger than r_{comm} . Assuming that m is divisible by 3 (it is always possible to truncate some small squares to satisfy this), we may group the small squares into $m/9$ groups of 3×3 blocks (see, e.g., Fig. 5).

If there is a single robot in a 3×3 block, the robot cannot communicate with the rest of the robots if it falls inside the small square in the center of the block (e.g., the solid gray square in Fig. 5). By Lemmas 8 and 9, for less than $(m/9) \log(m/9) = 2 \lfloor 1.1/r_{\text{comm}} \rfloor^2 \log(\lfloor 1.1/r_{\text{comm}} \rfloor / 3) / 9$ robots, the probability of having at least one of these 3×3 block containing exactly one robot is at least 0.34 as $m \rightarrow \infty$ (i.e., $r_{\text{comm}} \rightarrow 0$). If a 3×3 block has exactly one robot in it, with probability of $1/9$ the robot is in the middle square. Therefore, with probability at least $0.34/9 \approx 0.04$, $G(0)$ is disconnected. $\square$




Fig. 5. A 3×3 block as defined in the proof Theorem 10.

B. Ensuring target observability

With a connected communication graph $G(0)$ guaranteed by Lemma 5, we can solve a single assignment problem if for each $y \in Y^0$, $\|y - x\|_2 \leq r_{\text{sense}}$ for some $x \in X^0$. Similar techniques used in the proof of Lemma 5 lead to a similar lower bound on n .

Lemma 11 *For fixed r_{sense} and $0 < \varepsilon < 1$, every target $y \in Y^0$ is observable by some robot at $t = 0$ with probability at least $1 - \varepsilon$ when*

$$n \geq \left\lceil \frac{\sqrt{2}}{r_{\text{sense}}} \right\rceil^2 \log\left(\frac{1}{\varepsilon} \left\lceil \frac{\sqrt{2}}{r_{\text{sense}}} \right\rceil^2\right). \quad (12)$$

PROOF. If we partition the unit square Q into $\lceil \sqrt{2}/r_{\text{sense}} \rceil^2$ equal sized small squares and there is at least one robot in each small square, then any point of Q is within r_{sense} distance to some robot. Following the same argument used in the proof of Lemma 5, the inequality from (12) ensures that this happens with probability at least $1 - \varepsilon$. $\square$

Putting together Lemmas 5 and 11, we obtain a lower bound on n that makes a distance-optimal assignment possible.

Theorem 12 *Fixing $0 < \varepsilon < 1$, the communication graph is connected and every target $y \in Y^0$ is observable by some robot at $t = 0$ with probability at least $1 - \varepsilon$ when*

$$n \geq \begin{cases} \left\lceil \frac{\sqrt{2}}{r_{\text{sense}}} \right\rceil^2 \log\left(\frac{1}{\varepsilon} \left\lceil \frac{\sqrt{2}}{r_{\text{sense}}} \right\rceil^2\right), & r_{\text{sense}} < \frac{\sqrt{10}r_{\text{comm}}}{5} \\ \left\lceil \frac{\sqrt{5}}{r_{\text{comm}}} \right\rceil^2 \log\left(\frac{1}{\varepsilon} \left\lceil \frac{\sqrt{5}}{r_{\text{comm}}} \right\rceil^2\right), & r_{\text{sense}} \geq \frac{\sqrt{10}r_{\text{comm}}}{5} \end{cases} \quad (13)$$

PROOF. For the case of $r_{\text{sense}} < \sqrt{10}r_{\text{comm}}/5$, it is straightforward to check that (12) implies (8). Therefore, $G(0)$ is connected with probability $1 - \varepsilon$.

When $r_{\text{sense}} \geq \sqrt{10}r_{\text{comm}}/5$, by Lemma 5, (8) implies that $G(0)$ is connected with probability $1 - \varepsilon$. Moreover, there is at least one robot in each of the small squares with a side length of at most $r_{\text{comm}}/\sqrt{5}$ (as specified in the proof of Lemma 5). Having $r_{\text{sense}} \geq \sqrt{10}r_{\text{comm}}/5$ guarantees that a robot in a small square observes all targets within the same small square. Therefore, every $y \in Y^0$ is within a distance of r_{comm} to some $x \in X^0$. $\square$

Remark. Theorem 12 is not an asymptotic result and works for all r_{comm} and r_{sense} . If high-probability asymptotic result is desirable, Lemma 11 can be easily turned into a version similar to Theorem 10, following essentially the same proof techniques; we omit the details. In view of this fact, the bounds from Theorem 12 are asymptotically tight.

V. HIERARCHICAL STRATEGIES FOR $r_{\text{sense}} \geq \sqrt{2}$: OPTIMAL DISTANCE AND PERFORMANCE GUARANTEES

In this section, we work with the (region-based) Communication Model 2 and assume that $r_{\text{sense}} \geq \sqrt{2}$ (that is, every robot is aware of the entire Y^0). The study of Communication Model 2, besides leading to interesting conclusions on hierarchical strategies, also facilitates the analysis in the next section as we revisit Communication Model 1. As mentioned, a region-based communication model naturally leads to a hierarchical strategy. Let $h \geq 1$ be the number of hierarchies and $m_i, 1 \leq i \leq h$, be the number of regions at hierarchy i , we require that: 1. $m_{i+1} > m_i$, 2. $\sqrt{m_i}$ divides $\sqrt{m_{i+1}}$, 3. $m_1 \equiv 1$, and 4. a region at a higher numbered hierarchy does not span multiple regions at a lower numbered hierarchy. We call the associated strategy under these assumptions the hierarchical divide-and-conquer strategy, the details of which are described in Strategy 2. Note that for each region in Strategy 2, the robots can again let the highest labeled robot within the region carry out the strategy locally.

Strategy 2: HIERARCHICAL-DIVIDE-AND-CONQUER

Initial condition: $X^0, Y^0, h, m_1, \dots, m_h$

Outcome: permutation σ that assigns robot a_i to $y_{\sigma(i)}^0$

```

1 for each hierarchy  $i$  in decreasing order do
2   for each region  $j, 1 \leq j \leq m_i$  do
3     let  $n_a$  and  $n_g$  be the number of unmatched robots
4     and targets in region  $j$ , respectively
5     if  $n_a > n_g > 0$  then
6       pick the first  $n_g$  robots from the  $n_a$ 
7       unmatched robots and run an assignment
8       algorithm to match them with the  $n_g$ 
9       unmatched targets in region  $j$ 
10    else if  $n_g > n_a > 0$  then
11      pick the first  $n_a$  targets from the  $n_g$ 
12      unmatched targets and run an assignment
13      algorithm to match the  $n_a$  unmatched robots
14      with these targets in region  $j$ 
15    else
16      continue

```

It is clear that Strategy 2 is correct by construction because $|X^0| = |Y^0|$. The rest of this section is devoted to analyzing the strategy. We begin with a single hierarchy ($h = 1$). Since $r_{\text{sense}} \geq \sqrt{2}$ implies that all robots are aware of the entire set Y^0 , the robots may form a consensus of which robot should go to which target at $t = 0$ by finding the quantity

$$D_n^* = \min_{\sigma} \sum_{i=1}^n \|x_{\sigma(i)}^0 - y_i^0\|_2, \quad (14)$$

in which $\min_\sigma$ is taken over all permutations σ of the integers $1, \dots, n$. This assignment problem can be solved using the Hungarian method or any one of the Euclidean bipartite matching algorithms. Ajtai, Komlós, and Tusnády proved the following.

Theorem 13 (Optimal Matching [1]) *With high probability,*

$$C_1 \sqrt{n \log n} \leq D_n^* \leq C_2 \sqrt{n \log n}, \quad (15)$$

in which C_1, C_2 are positive constants.

That is, $D_n^* = \Theta(\sqrt{n \log n})$. Here, high probability means that the given event happens with a probability of $1 - O(n^{-c})$ for some positive constant c . Although the authors did not provide formulas for C_1 and C_2 in [1], simulation seems to suggest that $C_1 < C_2 < 1$ and $C_2/C_1 \rightarrow 1$ as $n \rightarrow \infty$. As an example, for $200 \leq n \leq 10000$, $0.4 \sqrt{n \log n} \leq D_n^* \leq 0.5 \sqrt{n \log n}$ on average (see, e.g., Fig. 6). Since all robots can start moving directly to their targets at $t = 0$, $T_n^* = D_n^* = \Theta(\sqrt{n \log n})$ as well. Note that this implies a total completion time of $\Theta(\sqrt{\log n / n})$, which goes to zero as $n \rightarrow \infty$.




Fig. 6. The ratio of $D_n^* / \sqrt{n \log n}$. Each data point is an average of 25 runs.

Next we look at the general case of having $h > 1$ hierarchies. To bound D_n , at each hierarchy i , we need to know the number of robots that can be matched locally, given by Lemma 14.

Lemma 14 *Suppose that the unit square Q is divided into m equal sized small squares. The number of robots that are not matched locally is $\sqrt{mn/\pi}$ for large n .*

PROOF. Focusing on a small square, say q_i , the probabilities of $x_1^0 \in X^0$ and $y_1^0 \in Y^0$ falling into q_i are both $1/m$. The probability of having x_1^0 but not y_1^0 in q_i is $(m-1)/m^2$; same is true for the event of having y_1^0 but not x_1^0 in q_i . The former event represents a surplus of a robot in q_i and the later a deficit in q_i . Thus, we may view the experiment of picking x_1^0 and y_1^0 as a one step walk on the real line starting at the origin, with $(m-1)/m^2$ probability of moving ± 1 . The entire process of picking X^0 and Y^0 can then be treated as a random walk of n such steps.

Denoting the move on the real line from the i -th step of the random walk as a random variable Z_j and let $S_n = Z_1 + \dots + Z_n$, we know that [20]

$$\lim_{n \rightarrow \infty} \mathbb{E}[|S_n|] = \frac{2\sqrt{n(m-1)}}{\sqrt{\pi m}}.$$

To see why this is intuitively true, note $\mathbb{E}[Z_j^2] = 2(m-1)/m^2$ and observe that

$$\begin{aligned} \mathbb{E}[S_n^2] &= \mathbb{E}[(Z_1 + \dots + Z_n)^2] = \mathbb{E}[Z_1^2 + \dots + Z_n^2] \\ &= n\mathbb{E}[Z_j^2] = \frac{2n(m-1)}{m^2}. \end{aligned}$$

Since at most half of the m small squares should have a surplus of robots on average, the total number of unmatched robots in expectation, for large n , is no more than

$$\frac{m\mathbb{E}[|S_n|]}{2} = \frac{m}{2} \frac{2\sqrt{n(m-1)}}{\sqrt{\pi m}} \approx \sqrt{\frac{mn}{\pi}}.$$

□

Lemma 15 *The total distance of matchings made at the bottom hierarchy is no more than $C_2 \sqrt{n \log n}$ as $n \rightarrow \infty$.*

PROOF. For the robots that are matched locally, by (15), the distance traveled by the robots in a square q_i (which contains n_i robots) is no more than $C_2 \sqrt{n_i \log n_i / m}$. Applying Jensen's inequality to the concave function $\sqrt{x \log x}$ and letting $x = n_i$ (we may simply ignore an n_i if it is zero) yields

$$\sum_{i=1}^m C_2 \sqrt{\frac{n_i \log n_i}{m}} \leq C_2 \sqrt{m} \sqrt{\frac{\sum_i n_i}{m} \log \frac{\sum_i n_i}{m}} \leq C_2 \sqrt{n \log n}.$$

□

We now give an upper bound on D_n for general hierarchical strategies.

Theorem 16 *Suppose that the unit square Q is divided into m_i equal sized small squares at hierarchy i with a total of $h \geq 2$ hierarchies. Then*

$$D_n \leq C_2 \sqrt{n \log n} + \sum_{i=h-1}^2 \left[\sqrt{\frac{2n}{\pi}} \left(1 - \sqrt{\frac{m_i}{m_{i+1}}} \right) \right] + \sqrt{\frac{2m_2 n}{\pi}}. \quad (16)$$

PROOF. The $C_2 \sqrt{n \log n}$ term on the right side of (16) is due to Lemma 15. Then at each hierarchy i with $2 \leq i < h$, the number of matched robots in total at this hierarchy is given by $\sqrt{m_{i+1}n/\pi} - \sqrt{m_i n/\pi}$. Since each of these robots needs to travel at most a distance of $\sqrt{2/m_{i+1}}$, the total distance incurred at hierarchy i is $\sqrt{2n/\pi} (1 - \sqrt{m_i/m_{i+1}})$. The last term on the right side of (16) is the distance incurred at the top hierarchy (i.e., hierarchy 1). □

Remark. Theorem 16 allows us to upper bound the performances of different hierarchical strategies depending on the choices of h and $\{m_i\}$. We observe that for fixed h and $\{m_i\}$ that do not depend on n , the first term $C_2 \sqrt{n \log n}$ dominates the other terms in (16) as $n \rightarrow \infty$. This implies that Strategy 2 yields assignments of which the total distance is at most a multiple of the true optimal distance.

Corollary 17 *For fixed h and $m_1, \dots, m_h$ that do not depend on n , Strategy 2 yields target assignment with $D_n/D_n^* = O(1)$.*

For example, letting $h \geq 2$ and $m_i = 4^{i-1}$ at hierarchy i , we have:

$$\begin{aligned} D_n &\leq C_2 \sqrt{n \log n} + \sum_{i=h-1}^2 \sqrt{\frac{2n}{\pi}} (1 - \sqrt{\frac{1}{4}}) + 4 \sqrt{\frac{n}{2\pi}} \\ &= C_2 \sqrt{n \log n} + (h+2) \sqrt{\frac{n}{2\pi}}. \end{aligned} \quad (17)$$

For any fixed h , as $n \rightarrow \infty$, $D_n/D_n^* \leq C_2/C_1 = O(1)$. Constant approximation ratio can also be achieved when h and/or $\{m_i\}$ depend on n . For example, letting $h = 3$, $m_3 = \log^2 n$, and $m_2 = \log n$, we have

$$\begin{aligned} D_n &\leq C_2 \sqrt{n \log n} + \sqrt{\frac{2n}{\pi}} (1 - \sqrt{\frac{\log n}{\log^2 n}}) + \sqrt{\frac{2n \log n}{\pi}} \\ &\stackrel{n \rightarrow \infty}{=} (C_2 + \sqrt{\frac{2}{\pi}}) \sqrt{n \log n}. \end{aligned} \quad (18)$$

Finally, since hierarchical strategies need not to run centralized assignment algorithms for all robots, the computational part of these strategies can be significantly faster. We will come back to this point in the next section.

VI. NEAR OPTIMAL STRATEGIES FOR ARBITRARY r_{comm}

After exploring hierarchical strategies for the region-based communication model, we now return to Communication Model 1 (Fig. 1(a)). If r_{comm} is arbitrary and the conditions specified in Theorem 4 may not hold, the best we can do is near distance-optimal strategies. In this section, we show that constant ratio approximation of distant optimality is again possible for the case of $r_{\text{sense}} \geq \sqrt{2}$ with arbitrary r_{comm} . In particular, we study a simple rendezvous strategy followed by a hierarchical rendezvous strategy that combines the rendezvous strategy with the hierarchical strategies from Section V. The basic idea behind these strategies is to move the robots to pass around information about other robots' locations.

A. A near distance-optimal rendezvous strategy

Our first suboptimal strategy uses moving robots for information aggregation until a robot is aware of the locations of all robots (i.e., the set X^0), at which point a centralized optimal assignment can be made. To carry out the strategy, the unit square Q is divided into $m = b^2$ disjoint, equal-sized small squares, with $b = \lceil \sqrt{2}/r_{\text{comm}} \rceil$. These small squares are labeled as $q_{i,j}$'s, in which i and j are the row number and column number of the square, respectively (see, e.g., Fig. 7).




Fig. 7. Directions for robots to move in the rendezvous strategy.

Based on its initial location, each robot can identify the small square $q_{i,j}$ it lies in. At $t = 0$, the robots in the squares on row 1 and row b start moving in the direction as indicated in Fig. 7. We want to use these robot to pass the information of where all robots are. At most one robot per square is required to move since all robots in a small square can communicate with each other by the assumption $b = \lceil \sqrt{2}/r_{\text{comm}} \rceil$.

Assuming that a robot in a square $q_{i,j}$ is moving downwards, it keeps moving until it is within the communication radius of a robot in a square with label $q_{i+k,j}$, $k \geq 1$, at which point it passes over the information it has and stops. The robot in $q_{i+k,j}$ then starts doing the same. The procedure continues until a robot reaches the middle of Q (row $\lceil b/2 \rceil$). It takes at most $1/2$ time unit (recall that we assume robots travel at unit speed) for this to happen. Once this happens, the robots in the squares belonging to row $\lceil b/2 \rceil$ repeat the same process horizontally until a robot in the center of Q knows the locations of all other robots. At this point, the robot in the center of Q that knows the location of all other robots makes a global assignment so that each robot is matched with a target. The moving process is then reversed to deliver the assignment information to all robots. The pseudo code of the strategy is given in Strategy 3.

Strategy 3: RENDEZVOUS

Initial condition: $X^0, Y^0, r_{\text{comm}}$

Outcome: produces permutation σ that assigns robots to targets and communicate σ to all robots

- 1 each robot computes its square $q_{i,j}$ based on r_{comm} , let the highest labeled robot within each $q_{i,j}$ be $a_{i,j}$, which represents $q_{i,j}$ **for each** $q_{i,j}$, $1 \leq i, j \leq b = \lceil \sqrt{2}/r_{\text{comm}} \rceil$ **do**
 - 2 **if** $i \neq \lceil b/2 \rceil$ **then**
 - 3 $a_{i,j}$ waits for up to $|\lceil b/2 \rceil - i|/b$ units of time for information from the previous square; after receiving information or after the wait time passes, it starts moving to the next squares and delivers its information once it can communicate with another robot in these squares; it then stops
 - 4 **else**
 - 5 $a_{i,j}$ waits for up to $1/2 + |\lceil b/2 \rceil - j|/b$ units of time for information from the previous square; after receiving information or after the wait time passes, it starts moving to the next squares and delivers its information once it can communicate with another robot in these squares; it then stops
 - 6 robot $a_{\lceil b/2 \rceil, \lceil b/2 \rceil}$ computes σ ; the earlier communication process is then reversed to deliver σ to all robots.
-

Strategy 3 is correct by construction. Besides the distance from the assignment, the robots in each column travel at most a total distance of two. The middle row incurs an extra distance of at most two. Thus, $D_n < D_n^* + 2b + 2$. Since $D_n^* \approx \Theta(\sqrt{n \log n})$, it dominates $2b + 2$ when $b = o(\sqrt{n \log n})$. In particular, $n = \Theta(1/r_{\text{comm}}^2)$ satisfies this requirement. Therefore,

Strategy 3 can yield near distance-optimal solution without requiring an n as large as (11) with respect to $1/r_{\text{comm}}$.

A drawback of Strategy 3 is that no robot can move to the targets until the assignment phase is complete. This yields a total task completion time of $T_n \approx 2n + T_n^*$. Such a T_n is not desirable since $T_n^* = O(\sqrt{n \log n})$ asymptotically. Furthermore, Strategy 3 requires running a centralized assignment algorithm for all robots. This might be impractical for large n . We address these issues with decentralized hierarchical strategies.

B. Decentralized hierarchical strategies

Among possible decentralized hierarchical strategies under Communication Model 1, we first look at one by combining Strategies 2 and 3. Instead of waiting for a centralized assignment to be made, in each of the small square $q_{i,j}$ as specified in Strategy 3, we let the robots in the square be assigned to targets that belong to the same square (we refer to these as *local assignments*). The robots that are not matched to targets then carry out Strategy 3. We denote this hierarchical rendezvous strategy as Strategy 4 and omit the pseudo code.

Corollary 18 For strategy 4 (2-level Hierarchical Rendezvous),

$$D_n \leq C_2 \sqrt{n \log n} + \sqrt{\frac{2mn}{\pi}} + 2\sqrt{m} + 2, \quad (19)$$

and

$$T_n = \Theta(\sqrt{n \log n} + \sqrt{mn}). \quad (20)$$

PROOF. The bound on D_n , given by (19), is straightforward to compute using Theorem 16, in which the first two terms on the right side of (19) correspond to the first and third terms of the right side of (16), respectively, and the last two terms are due to communication overhead. For total completion time, all but $\sqrt{mn}/\pi$ robots can start moving to their targets at $t = 0$. For the $\sqrt{mn}/\pi$ robots, they need to wait no more than two units of time each before moving to their targets. This gives us (20).

Remark. Similar to Strategy 3, for any fixed m , $D_n/D_n^* = O(1)$ (as $n \rightarrow \infty$). Moreover, in contrast to Strategy 3, for any fixed m , $T_n/T_n^* = O(1)$. Suppose that a centralized algorithm requires time $t(n)$, using the same centralized algorithm, Strategy 4 has a computational time complexity $O(mt(n/m) + t(\sqrt{mn}))$. If $t(n) = O(n^3)$ as given by the Hungarian method, then Strategy 4 has a running time of $O(n^3/m^2 + (mn)^{3/2})$. Taking $n = 10000, m = 10$, for example, we get about a 1000-time speedup.

By introducing additional hierarchies, Strategy 4 can be easily extended to a multi-hierarchy decentralized strategy. Depending on how the subdivisions are made, many such strategies are possible. For example, using $h \geq 2$ hierarchies with each hierarchy i having 4^{i-1} small squares, we get a “quad-merging” strategy as illustrated in Fig. 8, in which up to four representatives in four adjacent squares meet to decide local assignment of the robots in these squares at a given hierarchy level.

Although these suboptimal strategies vary in detail, they can be easily analyzed with the help of Theorem 16. For example, we look at an extension to Strategy 4 with three hierarchies; let




Fig. 8. Illustration of robot movements in a potential hierarchical strategy.

us call this strategy Strategy 5. After partitioning the bottom (or third) hierarchy to m squares, the middle (or second) hierarchy is partitioned into $k = \sqrt{m}$ small squares. At either the third or the second hierarchy, local assignments are made, followed by applying the rendezvous strategy as given in Strategy 3. It is again straightforward to derive

Corollary 19 For Strategy 5 (3-level Hierarchical Rendezvous),

$$D_n \leq C_2 \sqrt{n \log n} + 2\sqrt{n\sqrt{m}} + 4\sqrt{m} + 2. \quad (21)$$

Remark. Again, $D_n/D_n^* = O(1)$ as $n \rightarrow \infty$ for a fixed m . Introducing more hierarchy levels extends the total completion time T_n , which is approximately $2\sqrt{m}$ more; thus, the total completion time of Strategy 5 is also given by (20). Following similar analysis, the overall computation time required by Strategy 5 is $O(mt(n/m) + \sqrt{mt}(\sqrt{n}) + t(\sqrt{n\sqrt{m}}))$ given a centralized assignment algorithm that runs in $t(n)$ time.

VII. SIMULATION STUDIES

A. Number of required robots for a connected $G(0)$




Fig. 9. Effects of n on the connectivity of $G(0)$ for different values of r_{comm} .

In this subsection, we show a result of simulation to verify our theoretical findings in Section IV. Since the bounds over r_{comm} and r_{sense} are similar, we focus on r_{comm} and confirm the requirement for the connectivity of $G(0)$ for several r_{comm} 's ranging from 0.01 to 0.2. For each fixed r_{comm} , varying numbers of robots are used starting from $n = \log(1/r_{\text{comm}})/r_{\text{comm}}^2 = -\log r_{\text{comm}}/r_{\text{comm}}^2$ (the number of robots goes as high as 3×10^5 for the case of $r_{\text{comm}} = 0.01$). 1000 trials were run for each fixed combination of r_{comm}

and n ; the percentages of the runs with a connected $G(0)$ were reported in the simulation result shown in Fig. 9. The simulation suggests that the bounds on n from Theorem 10 are fairly tight.

TABLE I
COMPARISON BETWEEN (4) AND (6)

prob.	r_{comm}				
	0.2	0.1	0.05	0.02	0.01
0.1	0.001, 0.82	0.001, 0.96	0.001, 0.99	0.001, 1	0.003, 1
0.5	0.007, 0.92	0.006, 0.98	0.027, 0.99	0.064, 1	0.081, 1
0.9	0.2, 0.99	0.31, 1	0.381, 1	0.477, 1	0.502, 1
0.99	0.702, 1	0.742, 1	0.794, 1	0.834, 1	0.855, 1

To compare to (4), which also allows for estimation of n in terms of r_{comm} with a specified probability for obtaining a connected $G(0)$, for r_{comm} from 0.01 to 0.2, we computed n based on (4) and (6) for several probabilities (from 0.1 to 0.99). We then use these n 's to estimate the actual probability of having a connected $G(0)$. We list the result in Table I. Each main entry of the table has two probabilities separated by a comma, obtained using (4) and (6), respectively. As we can see, (4) gives underestimates (due to its asymptotic nature) and cannot be used to provide probabilistic guarantees. On the other hand, (6) provides overestimates that guarantees the desired probability.

B. Performance of near optimal strategies

Next, we simulate Strategies 2-5 and evaluate D_n , T_n , and computational time for these strategies over varying values of n and r_{comm} . Due to our choice of $k = \sqrt{m}$ in Strategy 5, for uniformity, we pick specific r_{comm} 's so that $m = \lceil \sqrt{2}/r_{\text{comm}} \rceil$ is a always perfect square. These values are $r_{\text{comm}} = 0.16, 0.09, 0.057$, and 0.04 , which correspond to $m = 81, 256, 625$, and 1296 , respectively. The number of robots used in each simulation ranges from 100 to 10000. For each n , 10 problems are randomly generated and used across all strategies. We test Strategy 2 using the same (two-hierarchy and three-hierarchy) partitions from Strategies 4 and 5.




Fig. 10. Distance optimality of Strategy 3 over varying n and r_{comm} .

Distance optimality: The ratios D_n/D_n^* for Strategy 3 over different n and r_{comm} are plotted in Fig. 10. We observe that the overhead for establishing global communication among the




Fig. 11. Distance optimality of Strategy 4 over varying n and r_{comm} .




Fig. 12. The effect of varying n on the distance optimality of Strategy 4 with $r_{\text{comm}} = 0.16$ ($m = 81$).

robots becomes insignificant as n increases, driving D_n/D_n^* to close to one.

For Strategy 4, the ratios were plotted similarly in Fig. 11 but with (small) errorbars. The errorbars display the standard deviation over the 10 runs (we omitted these from a figure, such as Fig. 10, when they are too small to see). They can be better seen in Fig. 12, which is a zoomed-in version of the $r_{\text{comm}} = 0.16$ line from Fig. 11.




Fig. 13. Distance optimality of Strategy 5 over varying n and r_{comm} .

As expected, for a fixed r_{comm} , D_n/D_n^* decreases as n increases. For $n = 10000$, the approximation ratios for our choices of r_{comm} are around 1.4 (due to the slow growing nature of $D_n^* \sim \sqrt{n \log n}$; fixing any r_{comm} , this ratio should be close to one for large n). On the other hand, for a fixed n , as the division of the unit square Q gets finer, D_n/D_n^* increases,

implying that decreasing communication radius has a negative effect on distance optimality. We observe similar results on distance optimality of Strategy 5 (see Fig. 13).

If we remove the rendezvous part from Strategies 4 and 5, they become versions of Strategy 2. The distance optimality performance of these two particular versions of Strategy 2 are briefly evaluated and given in Fig. 14 and Fig. 15. For all subdivisions made ($m = 81, 256, 625, 1296$), D_n/D_n^* ratios of less than two are achieved and can go as low as 1.06, showing that hierarchical decision making processes can provide very good optimality guarantees.




Fig. 14. The assignment cost of a two-level “pure” hierarchical strategy.




Fig. 15. The assignment cost of a three-level “pure” hierarchical strategy.

Computational time: We list the computational time, in seconds, for Strategies 3-5 in Table II (a table instead of a figure is used due to the number of data points and the large differences between the strategies). The standard $O(n^3)$ Hungarian method is used as the baseline assignment algorithm. Each main entry of the table lists three numbers corresponding to the computational time of Strategies 3, 4, and 5, respectively, for the given r_{comm} and n combination (note that any version of Strategy 2 has the same amount of computation as a corresponding rendezvous-based strategy). As expected, hierarchical assignment greatly reduces the computational time, often by a factor over 10^3 .

Time optimality: Since Strategies 3-5 sacrifice distance (and therefore, time) to compensate for limited communication, we do not expect the total completion time T_n of these strategies to compete with T_n^* . For example, in (20), although $T_n \rightarrow T_n^*$ as $n \rightarrow \infty$ for fixed $m = \lceil \sqrt{2}/r_{\text{comm}} \rceil^2$, it requires very large n for

TABLE II
COMPUTATIONAL TIME FOR STRATEGIES 3-5

# of robots, n	$r_{\text{comm}}(m)$			
	0.16 (81)	0.09 (256)	0.057 (625)	0.04 (1296)
100	0.007	0.007	0.007	0.007
	0.001	0.002	0.002	0.003
	0.001	0.0001	0.0004	0.0004
200	0.02	0.02	0.02	0.02
	0.001	0.005	0.01	0.02
	0.0001	0.0003	0.0004	0.0006
500	0.34	0.34	0.34	0.34
	0.005	0.02	0.07	0.14
	0.0006	0.001	0.002	0.003
1000	2.76	2.76	2.76	2.76
	0.015	0.07	0.22	0.54
	0.002	0.003	0.003	0.006
2000	22.3	22.3	22.3	22.3
	0.05	0.20	0.70	1.90
	0.009	0.006	0.011	0.015
5000	345	345	345	345
	0.02	0.78	2.84	8.28
	0.069	0.032	0.043	0.058
10000	2756	2756	2756	2756
	0.83	2.32	8.35	24.4
	0.43	0.11	0.11	0.14

$\sqrt{\log n}$ to dominate $\sqrt{m}$. Thus, we only compare T_n among Strategies 3-5. Using $T_n(i)$ to denote the T_n for Strategy i , $T_n(4)/T_n(3)$ and $T_n(5)/T_n(3)$ are plotted in Fig. 16 and 17. As n increases, Strategies 4 and 5 both take much less total completion time on average.




Fig. 16. Ratio of total completion time between Strategies 3 and 4.




Fig. 17. Ratio of total completion time between Strategies 3 and 5.

VIII. CONCLUSION AND DISCUSSIONS

Focusing on the distance optimality for the target assignment problem in a robotic network setting, we have characterized a necessary and sufficient condition under which true optimality can be achieved. We further provided a direct formula for computing the number of robots sufficient for probabilistically guaranteeing such an optimal solution. Then, we took a difference angle and looked at suboptimal strategies and their asymptotic performances as the number of robots goes to infinity. We showed that these strategies generally yield a constant approximation ratio when it comes to minimizing the total distance traveled by all robots. Many of these decentralized strategies also provide computational advantages over a centralized one.

We conclude the paper with a discussion on our choice on certain elements in the problem formulation and how they may be generalized.

Equal number of initial and target locations: In the problem statement we assume that $|X^0| = |Y^0|$; our result based on this assumption generally holds when $|X^0|$ and $|Y^0|$ are roughly the same because the locations in X^0 and Y^0 are randomly and independently selected. If this is not the case, when $|X^0| \gg |Y^0|$, it is likely that for a $y_i \in Y^0$, there is a unique $x_i \in X^0$ that is closest to y_i [25]. Moreover, for two different y_i, y_j , $x_i \neq x_j$. The spatial assignment problem then degenerates to finding the nearest robot for each $y \in Y^0$. When $|X^0| \ll |Y^0|$, the problem becomes a multiple salesmen version of the traveling salesman problem (we have a standard traveling salesman problem when $|X^0| = 1$), which is an NP-hard problem. It remains an interesting open question to investigate the middle ground, i.e., $|X^0| = C|Y^0|$ for some constant C (for example $C \in [0.1, 10]$).

Suboptimal strategies for arbitrary r_{sense} : In this paper, we did not provide suboptimal strategies in the case of arbitrary r_{sense} . However, we note that by modifying the rendezvous strategy (Strategy 3), it is possible to again get asymptotically constant approximation ratio on distance optimality for fixed r_{sense} as $n \rightarrow \infty$. It could be interesting to investigate whether there are better strategies when n is not very large. The same could be said for suboptimal strategies when r_{comm} is arbitrary.

Placement of initial and target locations: The assumption that all initial and target locations are independently picked uniformly randomly at $t = 0$ captures spatial randomness quite well. Although it is beyond the scope of this paper, it could be interesting to address the issue of randomness in time and the case in which the spatial distribution is not uniformly random (i.e., other spatial-temporal distributions).

Minimizing over other powers of the 2-norm: On the side of optimality measures, we note that Theorem 13 generalizes to arbitrary powers of the Euclidean 2-norm [1]. That is, for

$$D_{n,p}^* = \min_{\sigma} \sum_{i=1}^n \|x_{\sigma(i)}^0 - y_i^0\|_2^p, \quad (22)$$

it holds true that

$$D_{n,p}^* \sim n(\log n/n)^{p/2}. \quad (23)$$

As $p \rightarrow \infty$, (22) minimizes the longest distance traveled by any robot. Although we restrict our attention to the case of

$p = 1$ in this paper, our results readily extend to other values of p (i.e., other optimality criteria) with (23).

REFERENCES

- [1] M. Ajtai, J. Komlós, and G. Tusnády. On optimal matchings. *Combinatorica*, 4(4):259–264, 1984.
- [2] P. Balister, B. Bollobás, A. Sarkar, and M. Walters. Connectivity of random k -nearest-neighbour graphs. *Advances in Applied Probability*, 37:1–24, 2005.
- [3] D. P. Bertsekas. The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of Operations Research*, 14:105–123, 1988.
- [4] D. P. Bertsekas and D. A. Castañón. Parallel synchronous and asynchronous implementations of the auction algorithm. *Parallel Computing*, 17:707–732, 1991.
- [5] F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009. Electronically available at <http://coordinationbook.info>.
- [6] R. Burkard, M. Dell’Amico, and S. Martello. *Assignment Problems*. Siam Society for Industrial and Applied Mathematics, Philadelphia, 2012.
- [7] S.-J. Chung, S. Bandyopadhyay, I. Chang, and F. Y. Hadaegh. Phase synchronization control of complex networks of lagrangian systems on adaptive digraphs. *Automatica*, 49(5):1148–1161, 2013.
- [8] J. Cortés, S. Martínéz, and F. Bullo. Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions. *IEEE Transactions on Automatic Control*, 51(8):1289–1298, August 2006.
- [9] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM*, 19(2):248–264, 1972.
- [10] P. Erdős and A. Rényi. On a classical problem of probability theory. *Publ. Math. Inst. Hung. Acad. Sci.*, Ser. A 6:215–220, 1961.
- [11] N. M. Freris, H. Kowshik, and P. R. Kumar. Fundamentals of large sensor networks: Connectivity, capacity, clocks and computation. *Proceedings of the IEEE*, 98(11):1828–1846, November 2010.
- [12] A. Ganesh and F. Xue. On the connectivity and diameter of small-world networks. *Advances in Applied Probability*, 39:853–863, 2007.
- [13] A. Jadbabaie, J. Lin, and A. S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, 2003.
- [14] M. Ji, S. Azuma, and M. Egerstedt. Role-assignment in multi-agent coordination. *International Journal of Assitive Robotics and Mechatronics*, 7(1):32–40, August 2006.
- [15] S. Kloder and S. Hutchinson. Path planning for permutation-invariant multirobot formations. *IEEE Transactions on Robotics*, 22(4):650–665, 2006.
- [16] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
- [17] J. Lin, A. S. Morse, and B. D. O. Anderson. The multi-agent rendezvous problem. part 1: The synchronous case. *SIAM Journal on Control and Optimization*, 46(6):2096–2119, November 2007.
- [18] J. Lin, A. S. Morse, and B. D. O. Anderson. The multi-agent rendezvous problem. part 2: The asynchronous case. *SIAM Journal on Control and Optimization*, 46(6):2120–2147, November 2007.
- [19] G. Mao and B. D. O. Anderson. Connectivity of large wireless networks under a general connection model. *IEEE Transactions on Information Theory*, 59(3):1761–1772, 2013.
- [20] F. Mosteller, R. E. K. Rourke, and G. B. Thomas. *Probability and Statistics*. Addison-Wesley, Reading, MA, 1961.
- [21] M. Penrose. The longest edge of the random minimal spanning tree. *Annals of Applied Probability*, 7:340–361, 1997.
- [22] M. Penrose. *Random Geometric Graphs*. Oxford Studies in Probability. Oxford University Press, London, 2003.
- [23] R. Sharathkumar and P. K. Agarwal. A near-linear time ϵ -approximation algorithm for geometric bipartite matching. In *Proceedings of the 44th Symposium on Theory of Computing Conference*, pages 385–394, 2012.
- [24] V. Sharma, M. A. Savchenko, E. Frazzoli, and P. G. Voulgaris. Transfer time complexity of conflict-free vehicle routing with no communications. *International Journal of Robotics Research*, 26(3):255–271, 2007.
- [25] S. L. Smith and F. Bullo. Monotonic target assignment for robotic networks. *IEEE Transactions on Automatic Control*, 54(9):2042–2057, September 2009.
- [26] H. G. Tanner, A. Jadbabaie, and G. J. Pappas. Flocking in fixed and switching networks. *IEEE Transactions on Automatic Control*, 52(5):863–868, 2007.

- [27] K. Treleaven, M. Pavone, and E. Frazzoli. Asymptotically optimal algorithms for one-to-one pickup and delivery problems with applications to transportation systems. *IEEE Transactions on Automatic Control*. To appear.
- [28] M. Turpin, K. Mohta, N. Michael, and V. Kumar. Goal assignment and trajectory planning for large teams of aerial robots. In *Proc. Robotics: Science and Systems*, 2013.
- [29] P. M. Vaidya. Geometry helps in matchings. *SIAM Journal on Computing*, 18(6):1201–1225, 1989.
- [30] F. Xue and P. R. Kumar. The number of neighbors needed for connectivity of wireless networks. *Wireless Networks*, 10(2):169–181, March 2004.
- [31] J. Yu and S. M. LaValle. Distance optimal formation control on graphs with a tight convergence time guarantee. In *Proceedings IEEE Conference on Decision & Control*, pages 4023–4028, 2012.
- [32] M. M. Zavlanos and G. J. Pappas. Dynamic assignment in distributed motion planning with local coordination. *IEEE Transactions on Robotics*, 24(1):232–242, 2008.
- [33] M. M. Zavlanos, L. Spesivtsev, and G. J. Pappas. A distributed auction algorithm for the assignment problem. In *Proceedings IEEE Conference Decision & Control*, pages 1212–1217, 2008.