

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

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

Abstract—We study the problem of multi-robot target assignment to minimize the total distance traveled by the robots until they all reach an equal number of static targets. In the first half of the paper, we present a necessary and sufficient condition under which true distance optimality can be achieved for robots with limited communication and target-sensing ranges. Moreover, we provide an explicit, non-asymptotic formula for computing the number of robots needed to achieve distance optimality in terms of the robots' communication and target-sensing ranges with arbitrary guaranteed probabilities. The same bounds are also shown to be asymptotically tight. In the second half of the paper, we present suboptimal strategies for use when the number of robots cannot be chosen freely. Assuming first that all targets are known to all robots, we employ a hierarchical communication model in which robots communicate only with other robots in the same partitioned region. This hierarchical communication model leads to constant approximations of true distance-optimal solutions under mild assumptions. We then revisit the limited communication and sensing models. By combining simple rendezvous-based strategies with a hierarchical communication model, we obtain decentralized hierarchical strategies that achieve constant approximation ratios with respect to true distance optimality. Results of simulation show that the approximation ratio is as low as 1.4.

Index Terms—Networked robots, optimality.

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. Focusing on minimizing the total distance traveled by the robots in a planar setting, we seek optimality guarantees and near-optimal strategies. For robot-to-robot communication, we investigate both a simple circular range-based model and a region-based model in which all robots within the same region can communicate with each other. When we consider the limited target-sensing capability of the robots, a circular range sensing model is used.

Manuscript received July 17, 2013; revised January 28, 2014; accepted July 5, 2014. Date of publication July 30, 2014; date of current version January 21, 2015. This work was supported in part by AFOSR Grant FA95501210193, NSF Grant IIS-1253758, and NSF Grant ECCS 10-27437. Recommended by Associate Editor E. Frazzoli.

J. Yu was with the University of Illinois at Urbana-Champaign, Champaign, IL 61820-5711 USA. He is now with the Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: jingjin@csail.mit.edu).

S.-J. Chung and P. G. Voulgaris are with the Coordinated Science Lab and the Department of Aerospace Engineering, University of Illinois at Urbana-Champaign, Champaign, IL 61820-5711 USA (e-mail: sjchung@illinois.edu; vougari@illinois.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TAC.2014.2344291

The class of problems that we study is denoted as *target assignment in robotic networks* as it shares many similarities with the problems studied in [1]. In [1], 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 significantly different assumptions on the communication and sensing models of the robots. The total distance serves as a proper proxy to quantities such as the total energy consumption of all the robots. Note that a distance-optimal solution for the target assignment problem generally does not imply time optimality and vice versa [2].

As its name implies, the problem of *target assignment in robotic networks* requires solving an *assignment* (or *matching*) problem. The assignment problem is extensively studied in the area of combinatorial optimization, with polynomial time algorithms available for solving many of its variations [3]–[8]. If we instead put more emphasis on multi-robot systems, the problems of robotic task allocation [9]–[12], swarm reconfiguration [13], multi-robot path planning [14]–[16], and multi-agent consensus [17]–[20] are relevant. For a more comprehensive review on these topics, see [21].

Our work is also closely related to the study of the connectivity of wireless networks. An interesting result [22] 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 [22] 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 [23]. These nearest neighbor based connectivity models were further studied in [24]–[26], to list a few. In many of these settings, a *geometric graph* structure is used [27].

This research effort brings forth three contributions. First, for robots with limited range-based target-sensing and communication capabilities (the ranges are 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 by n) sufficient for all robots to form a connected network for a fixed communication radius r_{comm} . In contrast to the asymptotic connectivity results from [22], [28], we give

¹ETSP stands for the *Euclidean traveling salesman problem*.

n as an explicit function of r_{comm} that is also non-asymptotic. Therefore, our bounds hold without requiring $n \rightarrow \infty$. We further show that the same bounds are 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 show that an infinite family of hierarchical strategies can lead to decentralized target assignments while ensuring that the total expected distance traveled by the robots is asymptotically within a constant multiple 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 savings on computation time (in certain cases, a speedup of 1000 times or more) are 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 for the lack of global communication) are combined to obtain decentralized suboptimal algorithms. These hybrid strategies, under mild assumptions, preserve the constant approximation ratios on distance optimality achieved by the “pure” hierarchical strategies. We further show that the global target-sensing assumption can be removed without affecting asymptotic optimality.

Portions of this work were presented in [29] and [30] for the early dissemination of results. Compared with [29] and [30], this paper provides a more comprehensive view of the results along with complete proofs for all theorems. Many of the proofs have been significantly improved to illustrate more clearly proof techniques that may be of interest on their own. In addition, the current paper discusses extensively generalizations of the stochastic target assignment problem to mismatching number of robots and targets, and to higher dimensions.

The rest of the paper is organized as follows. In Section II, we present notations and well-known results from other branches of research needed for the development of our results. After stating the problem formally in Section III, we then elaborate on the three stated contributions in Sections IV–VI. We present results of simulation studies in Section VII to validate our theoretical results and conclude in Section VIII.

II. PRELIMINARIES

In this section, we review results on the balls and bins problem, 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 $\lceil x \rceil$ (resp., $\lfloor x \rfloor$) denotes the smallest (resp., largest) integer that is larger (resp., smaller) than or equal to x . $|\cdot|$ denotes the cardinality for a set and the absolute value for a real number. We use $\|v\|_2$ to denote the Euclidean 2-norm of a vector v . The unit square $[0, 1]^2 \subset \mathbb{R}^2$ is denoted as Q . The expectation of a random variable X is denoted as $\mathbf{E}[X]$. We use $E(\cdot)$ to represent 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

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

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

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

If $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$, then we say $f(x) = \Theta(g(x))$. Finally, $f(x) = o(g(x))$ (resp., $f(x) = \omega(g(x))$) if and only if $f(x) = O(g(x))$ (resp., $f(x) = \Omega(g(x))$) and $f(x) = \Theta(g(x))$ does not hold.

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 [31]): 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 $k = 1$ follows readily.

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), letting $k = 1$ yields

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

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 $T_1 = m \log m$ 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 at least $m \log m + 5m$ balls are tossed. On the other hand, the same probability is no more than 0.001 when no more than $m \log m - 2m$ balls are tossed.

B. Linear Assignment Problem

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

Problem 1 (Linear Assignment): 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

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

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 represent physical locations. The *Hungarian method* for the assignment problem, proposed by Kuhn [7], has an $O(n^4)$ running time, which was subsequently improved to $O(n^3)$ by Edmonds and Karp [6]. Many other algorithms for the assignment problem exist, including other primal-dual (linear programming) methods [5], auction based methods [3], and parallel algorithms [4], [8]. 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 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 [32]. Alternatively, near linear time approximation algorithms can yield near optimal solutions with high probability [33].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 and connecting any two vertices whose underlying points x_1 and x_2 satisfy $\|x_1 - x_2\|_2 \leq r_{\text{comm}}$. When 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 [27] 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 [28]): Let G be a 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 $n \rightarrow \infty$

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

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 objective.

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 (the superscript emphasizes that these points are

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 [33], [32] 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 not strongly polynomial time algorithms like the Hungarian method.

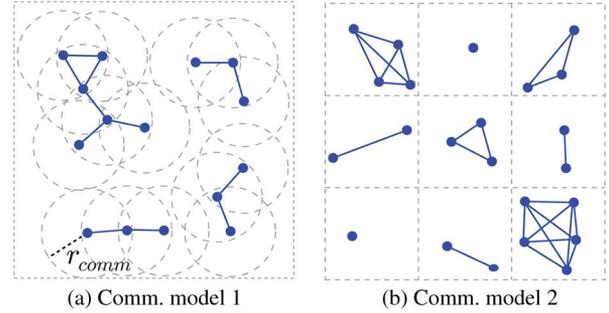


Fig. 1. (a) The communication graph (solid blue nodes and edges) for a set of robots under Communication Model 1 with a communication radius of r_{comm} . Robots (blue dots) in the same connected component of a communication graph can freely communicate with each other. (b) The communication graph for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

obtained at the start time $t = 0$) in the unit square Q ,⁴ selected uniformly randomly. 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: The robots are single integrators, i.e., $\dot{x}_i(t) = u_i(t)$ with $u_i(t)$ being piece-wise smooth and $\|u_i(t)\|_2 \in \{0, 1\}$. We assume the size of the robots is 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* $G(t)$, which is a geometric graph that changes over time, as follows. $G(t)$ 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 $\|x_i(t) - x_j(t)\|_2 \leq r_{\text{comm}}$. Fig. 1(a) provides an example of a (disconnected) communication graph.

Given our focus on distance optimality, we make the simplifying assumption that all robots corresponding to vertices in a connected component of the communication graph may exchange information instantaneously. In other words, robots in a connected component of $G(t)$ can be treated as a single robot insofar as decision making is concerned.

Communication Model 2: The unit square Q is divided into $m = b^2$ equal-sized smaller squares (regions).⁵ Robots within each region can communicate with one another but

⁴Our results are scale-invariant because all the theorems hold for squares of any size with proper scaling. Hence, a unit square environment is used throughout the paper.

⁵In this paper, m is frequently used to denote the number of small squares in a division of the unit square Q and $b = \sqrt{m}$ is the number of resulting partitions on an edge of the unit square. The value of m and b may vary.

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

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 Communication Model 1 with r_{comm} or Communication Model 2, find a control strategy $\mathbf{u}(t) = [u_1(t), \dots, u_n(t)]$, 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

$$D_n := \sum_{i=1}^n \int_0^{t_i^f} \|\dot{x}_i(t)\|_2 dt. \quad (5)$$

As an accurate proxy to 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 . Over all permutations σ of the numbers $1, \dots, n$, and for fixed X^0 and Y^0 , the minimum total distance for robots moving along continuous paths is

$$D_n^* := \min_{\sigma} \sum_{i=1}^n \|x_i^0 - y_{\sigma(i)}^0\|_2 \quad (6)$$

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, we say that distance optimality is achieved if $\min_{\mathcal{U}} D_n = D_n^*$. 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 by T_n . If all robots start moving toward targets and do not stop in the middle, then $T_n = D_n$. In particular, we define $T_n^* := D_n^*$.

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, 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 assignments 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.

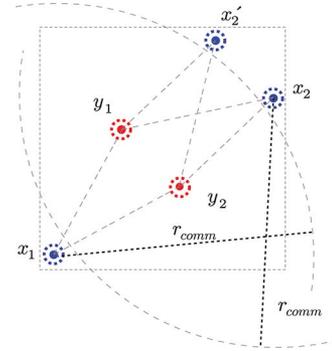


Fig. 2. General setup in which the two robots cannot communicate with each other at $t = 0$ and therefore cannot always decide an optimal assignment at $t = 0$.

Theorem 4: In a unit square, under sensing and communication constraints (i.e., $r_{\text{comm}}, r_{\text{sense}} < \sqrt{2}$), distance optimality can be achieved with probability one if and only if at $t = 0$:

- i) the communication graph is connected;
- ii) every target is within a distance of r_{sense} to some robot.

Proof: We first prove that the conditions are necessary with two claims: 1) an optimal assignment that minimizes D_n is possible 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 distance-optimal assignments forbid robots from moving unnecessarily, requiring 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 when $r_{\text{comm}} < \sqrt{2}$. For $n = 2$, assume that the two targets are located at y_1 and y_2 as given in Fig. 2 (solid red dots). Assume the first robot a_1 is located at x_1 (the blue solid dot at the lower left of Fig. 2) and a_1 is of equal distance to y_1 and y_2 . Let the second robot a_2 take two possible locations x_2 and x'_2 as shown, which are symmetric along a diagonal of Q . If a_2 is at x_2 (resp. x'_2), then a_2 should go to y_2 (resp. y_1), forcing a_1 to go to y_1 (resp. y_2). Not knowing a_2 's location because a_1 is out of a_2 's communication radius, a_1 has at most 50% chance of picking the distance minimizing choice at $t = 0$. 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 a non-zero probability that 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. The example can be extended to work for arbitrary n by adding additional robots and targets to close vicinities of x_1 and y_1 , respectively.

For the second claim, suppose that at $t = 0$, some $y \in Y^0$ is not within a distance of r_{sense} to any $x \in X^0$. A robot must move to search for that y . This will cause the 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 for ensuring distance optimality by either increasing the number of robots or increasing one or both of r_{comm} and r_{sense} . This 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) handle the entire assignment process.

Strategy 1: CENTRALIZED ASSIGNMENT

Initial condition: X^0, Y^0

Outcome: permutation σ that assigns a 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 compute over $\{d_{i,j}\}$ an assignment that minimizes D_n
 - 3 communicate the assignment to all the 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 yields an asymptotic relationship between the required number of robots for $G(0)$ to be connected and r_{comm} . We take a different approach and produce the required number of robots 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 robots to take a lattice-like formation such as a grid. It turns out that when the robots are uniformly randomly distributed, only a logarithmic factor more robots are needed to ensure a connected $G(0)$.

Lemma 5: Suppose that n robots are uniformly randomly distributed in the unit square. For fixed $r_{\text{comm}} < \sqrt{2}$ and $0 < \varepsilon < 1$, at $t = 0$, the communication graph is connected with probability at least $1 - \varepsilon$ if

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

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 $\{q_1, \dots, q_m\}$. Under this division scheme, a robot residing in a small square q_i can communicate with any other robot in the four squares sharing a side with q_i (see Fig. 3). Therefore, $G(0)$ is connected if each q_i contains a robot. Let n_i denote the number of robots in q_i . Then

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

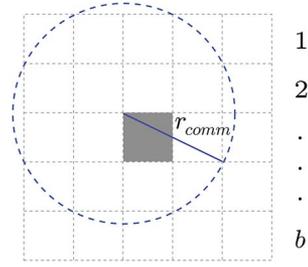


Fig. 3. If the small squares have a side length of $\lceil \sqrt{5}/r_{\text{comm}} \rceil$ or smaller, then a robot in such a square (e.g., the gray square) can communicate with any robot in the four neighboring small squares sharing a side with the gray square.

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 (i.e., the union bound), the probability that at least one of 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

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

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

The bound in Lemma 5 can be further tightened; Corollary 6 (below) illustrates one way to achieve this. It produces n smaller than that given by (7) when $r_{\text{comm}} < \sqrt{5}/2$.

Corollary 6: Suppose that n robots are uniformly randomly distributed in the unit square. For fixed $r_{\text{comm}} < \sqrt{2}$ and $0 < \varepsilon < 1$, at $t = 0$, the communication graph is connected with probability at least $1 - \varepsilon$ if

$$n \geq \left\lceil \frac{\sqrt{5}}{r_{\text{comm}}} \right\rceil^2 \log \left[\frac{1}{\varepsilon} \left(\frac{1}{2} \left\lceil \frac{\sqrt{5}}{r_{\text{comm}}} \right\rceil^2 + \left\lceil \frac{\sqrt{5}}{r_{\text{comm}}} \right\rceil \right) \right]. \quad (8)$$

Proof: If each of the shaded small squares in Fig. 4 has at least one robot, then $G(0)$ must be connected: any robot falling in a small white square must be connected to some robot in a shaded small square. This shows that (8) 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 (7) [and (8)], 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.

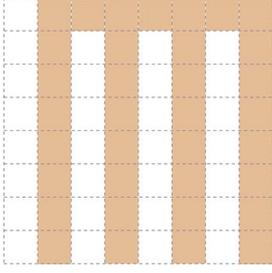


Fig. 4. As long as each of the shaded small squares contains an robot, $G(0)$ must be connected. Therefore, only $b^2/2 + b$ small squares need to have robots in them.

Lemma 7: Suppose that n robots, each with a communication radius of r_{comm} , are uniformly randomly distributed in the unit square. At $t = 0$, the communication graph is asymptotically connected with arbitrarily high probability $e^{-e^{-c}}$ (for some $c > 0$) if

$$n \geq \left(2 \log \left\lceil \frac{\sqrt{5}}{r_{\text{comm}}} \right\rceil + c \right) \left\lceil \frac{\sqrt{5}}{r_{\text{comm}}} \right\rceil^2. \quad (9)$$

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, as $m \rightarrow \infty$, 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 these many robots are also necessary for the high probability guarantee.

Let $\mathbf{P}_{n,m}(E)$ denote the probability of an event E happening after tossing n balls into m bins. We work with two events: E_0 , the event that “at least one bin is empty”, 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 in the next two lemmas.

Lemma 8: Suppose that $1 \leq n \leq m$ balls are tossed uniformly randomly into m bins. Then

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

Proof: First we prove a useful inequality: for $m \in \mathbb{N}$

$$\left(1 - \frac{1}{m} \right)^{m-1} > e^{-1}. \quad (10)$$

To see this, note that the function $\log(1-x)^{(1/x)-1}$ has a Taylor expansion of $-1 + x/2 + o(x^2) > -1$ for small $x > 0$, yielding that $(1-x)^{(1/x)-1} > e^{-1}$ for small $x > 0$. Since the derivative of $(1-x)^{(1/x)-1}$ is positive for $x \in (0, 1)$,

(10) holds for all $m > 0$ (we use the definition $0^0 = 1$ here).

To prove Lemma 8, 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 $1 \leq k \leq m$, we have

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

\square

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

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

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 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

$$\begin{aligned} \mathbf{P}_{n'+k,m}(E_1) \\ \geq \mathbf{P}_{n',m}(E_0) \mathbf{P}_{k,m}(\text{exactly one ball falls in bin 1}) \\ = \frac{k}{m} \left(1 - \frac{1}{m} \right)^{k-1} \mathbf{P}_{n',m}(E_0). \end{aligned} \quad (11)$$

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

$$\lim_{m \rightarrow \infty} \mathbf{P}(T_1 \geq m \log m - m) = 1 - e^{-e}. \quad (12)$$

That is, as $m \rightarrow \infty$, for $0 < n' < m \log m - m$, $\mathbf{P}_{n',m}(E_0) \geq 1 - e^{-e}$. Plugging this into (11) and letting $k = m$, we have that for $m < n < m \log m$, as $m \rightarrow \infty$

$$\mathbf{P}_{n,m}(E_1) \geq (1 - e^{-e}) \frac{m}{m} \left(1 - \frac{1}{m} \right)^{m-1} > (1 - e^{-e})e^{-1}$$

in which the last inequality is by (10). \square

Under the assumptions of Lemmas 8 and 9, we always have that as $m \rightarrow \infty$, $\mathbf{P}_{n,m}(E_1) > \min\{e^{-1}, (1 - e^{-e})e^{-1}\} > 0.34$. 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.

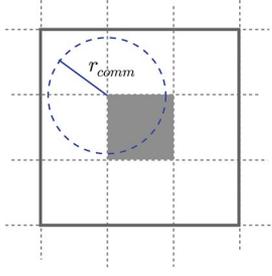


Fig. 5. A 3×3 block as defined in the proof Theorem 10.

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

$$n = \Theta \left(\frac{1}{r_{\text{comm}}^2} \log \frac{1}{r_{\text{comm}}} \right) \quad (13)$$

is necessary and sufficient to ensure that at $t = 0$, the communication graph is asymptotically connected with arbitrarily high probability.

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 in which $b = \lfloor 1.1/r_{\text{comm}} \rfloor$. 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 $(m/9) \times \log(m/9) = 2 \lfloor 1.1/r_{\text{comm}} \rfloor^2 \log(\lfloor 1.1/r_{\text{comm}} \rfloor / 3) / 9$ robots, the probability of having at least one of these 3×3 blocks 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

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: Suppose that n robots and n targets are uniformly randomly distributed in the unit square. For fixed r_{sense}

and $0 < \varepsilon < 1$, every target is observable by some robot at $t = 0$ with probability at least $1 - \varepsilon$ if

$$n \geq \left\lceil \frac{\sqrt{2}}{r_{\text{sense}}} \right\rceil^2 \log \left(\frac{1}{\varepsilon} \left\lceil \frac{\sqrt{2}}{r_{\text{sense}}} \right\rceil^2 \right). \quad (14)$$

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 (14) 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: Suppose that n robots and n targets are uniformly randomly distributed in the unit square. Fixing $0 < \varepsilon < 1$, at $t = 0$, the communication graph is connected and every target is observable by some robot with probability at least $1 - \varepsilon$ if

$$n \geq \left\lceil \frac{\sqrt{10}}{\theta} \right\rceil^2 \log \left(\frac{1}{\varepsilon} \left\lceil \frac{\sqrt{10}}{\theta} \right\rceil^2 \right) \quad (15)$$

in which $\theta := \min\{\sqrt{5}r_{\text{sense}}, \sqrt{2}r_{\text{comm}}\}$.

Proof: When $\theta = \sqrt{5}r_{\text{sense}}$, (15) becomes (14), which implies (9). Therefore, $G(0)$ is connected with probability $1 - \varepsilon$.

When $\theta = \sqrt{2}r_{\text{comm}}$, i.e., $r_{\text{sense}} \geq \sqrt{10}r_{\text{comm}}/5$, by Lemma 5, (9) 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 robots in a small square observes all targets within the same small square. Therefore, every $y \in Y^0$ is within a distance of r_{sense} to some $x \in X^0$. \square

Remark: Theorem 12 is not an asymptotic result and applies to all r_{comm} and r_{sense} . If a high probability asymptotic result is desirable, Lemma 11 can be readily turned into a version similar to Theorem 10, by following the same proof techniques. 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 Section VI as we revisit Communication Model 1.

A region-based communication model naturally leads to a hierarchical strategy for solving Problem 2 under the optimality criterion of minimizing the cost defined by (5). Let $h \geq 1$ be the number of hierarchies and m_i , $1 \leq i \leq h$, be the number of equal-sized regions at hierarchy i . We make the following assumptions that are mainly used in Theorem 16: *i*) $m_1 \equiv 1$, *ii*) $m_{i+1} \geq m_i$, and *iii*) a region at a higher numbered hierarchy is contained in a single region at a lower numbered hierarchy.

For example, dividing Q into 4^{i-1} squares at hierarchy i satisfies these requirements.

Strategy 2: HIERARCHICAL-DIVIDE-AND-CONQUER

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

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

```

1 for each hierarchy  $i \in \{1, \dots, h\}$  in decreasing order do
2   for each region  $j \in \{1, \dots, m_i\}$  do
3     let  $n_a$  and  $n_g$  be the number of unmatched robots
4     and targets in region  $j$ , respectively
5     if  $n_a \geq n_g > 0$  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  $n_g \geq n_a > 0$  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  $n_a$  targets in region  $j$ 
15    else
16      continue

```

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.

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 an optimal assignment σ that yields D_n^* as defined by (6). This assignment problem can be solved using a bipartite matching algorithm such as the Hungarian method. Ajtai, Komlós, and Tusnády proved the following about D_n^* .

Theorem 13 (Optimal Matching [34]): Assuming that n points are i.i.d. following the uniform distribution over a unit square, then, with probability $1 - o(1)$

$$C_1 \sqrt{n \log n} \leq D_n^* \leq C_2 \sqrt{n \log n} \quad (16)$$

in which C_1 and C_2 are positive constants.

Remark: The second inequality in (16) remains true in expectation and also for arbitrary probability measures on $[0, 1]^2$, albeit with a different universal constant than C_2 , by a result of Talagrand [35]. Therefore, $D_n^* = \Theta(\sqrt{n \log n})$ in expectation. Although no formulas for C_1 and C_2 from (16) were given in [34], a simulation study suggests 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 Fig. 6).

Next, we look at the general case with $h > 1$ hierarchies. To bound D_n , at each hierarchy i , we need to know the number of robots that cannot be matched locally. We derive this number in

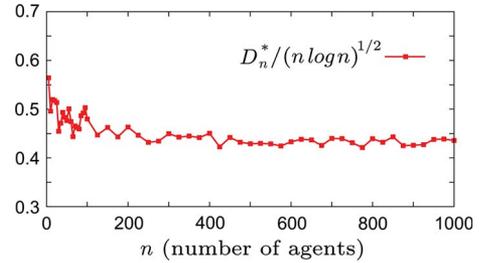


Fig. 6. The ratio of $D_n^*/\sqrt{n \log n}$. Each data point is an average of 25 runs.

Lemma 14. Note that Lemma 14 does not depend on m and n being large.

Lemma 14: Suppose that n robots and n targets are uniformly randomly distributed in the unit square Q , and Q is divided into m equal-sized regions. Within each of these m regions, the robots are matched one-to-one with the targets until no more matchings can be made. The total number of robots that are left unmatched is no more than $\sqrt{n(m-1)}/2$ in expectation.

Proof: Restricting to one of the m equal-sized regions, say q_i , we know for $x_j^0 \in X^0$ and $y_j^0 \in Y^0$

$$\mathbf{P}(x_j^0 \in q_i) = \mathbf{P}(y_j^0 \in q_i) = \frac{1}{m}$$

and

$$\mathbf{P}(x_j^0 \in q_i, y_j^0 \notin q_i) = \mathbf{P}(x_j^0 \notin q_i, y_j^0 \in q_i) = \frac{m-1}{m^2}$$

in which the event $(x_j^0 \in q_i, y_j^0 \notin q_i)$ represents a surplus of one robot in q_i and the event $(x_j^0 \notin q_i, y_j^0 \in q_i)$ a deficit in q_i . Thus, we may view the experiment of picking x_j^0 and y_j^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.

Under this random walk analogy, we may use a random variable $Z_j \in \{0, \pm 1\}$ to represent the outcome of picking (x_j^0, y_j^0) . We immediately have that $\mathbf{E}[Z_j^2] = 2(m-1)/m^2$. Letting $S_n = Z_1 + \dots + Z_n$, we can compute the variance of S_n as

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

Applying Jensen's inequality to the concave function \sqrt{x} with $x = |S_n|^2 = S_n^2$, we have

$$\begin{aligned} \mathbf{E}[|S_n|] &= \mathbf{E}[\sqrt{S_n^2}] \leq \sqrt{\mathbf{E}[S_n^2]} \\ &\Rightarrow \mathbf{E}[|S_n|] \leq \sqrt{\frac{2n(m-1)}{m^2}}. \end{aligned}$$

Because, in expectation, an equal number of the m regions have surpluses (more robots than targets) and deficits (fewer robots than targets), and some of the m regions may have neither, no more than half of the m regions should have a surplus of robots on average. The total number of unmatched robots in expectation is then no more than $(m/2) * \mathbf{E}[|S_n|] \leq \sqrt{n(m-1)}/2$. \square

The distance traveled by the matched robots at the bottom hierarchy with m regions can be bounded easily. For simplicity, we now assume that these regions are equal-sized squares.

Lemma 15: Suppose that n robots and n targets are uniformly randomly distributed in the unit square Q , and Q is divided into m equal-sized small squares. Within each of these m small squares, the robots are matched one-to-one with the targets until no more matchings can be made. The minimum total distance of matchings made between the robots and the targets within the small squares is no more than $C\sqrt{n \log n}$ in expectation, for some positive constant C .

Proof: Since Q is divided into m squares, these squares all have a side length of $1/\sqrt{m}$. Let one such square be q_i with n_i robots (note that $\sum_{i=1}^m n_i = n$). Since a uniform distribution restricted to q_i is again uniform, we can apply Theorem 13 to q_i . If we let these n_i robots match only with targets inside q_i , then the total distance incurred locally will not exceed $C\sqrt{n_i \log n_i/m}$ in expectation. Here C is some positive constant.

Note that it is not necessarily the case that all n_i robots will be matched locally in q_i . This does not affect the current proof. For some $1 \leq i \leq m$, it may be the case that no local matchings can be made because either $n_i = 0$ or there is no target in q_i . Let $m' \leq m$ denote the number of these m squares in which local matchings can be made. The total distance incurred by local matchings is then upper bounded by (note that n_i is now indexed with respect to these m' squares)

$$\sum_{i=1}^{m'} C \sqrt{\frac{n_i \log n_i}{m}} = C \frac{m'}{\sqrt{m}} \sum_{i=1}^{m'} \frac{1}{m'} \sqrt{n_i \log n_i}.$$

Here we assume that $m' > 0$, otherwise the local matchings would have a distance cost of zero. Since the function $\varphi(x) = \sqrt{x \log x}$ is concave, by Jensen's inequality, $\mathbf{E}[\sqrt{x \log x}] \leq \sqrt{\mathbf{E}[x] \log(\mathbf{E}[x])}$. Letting $x = n_i$ and the expectation be carried out over the discrete uniform distribution with $1/m'$ probability each, we have

$$\begin{aligned} & C \frac{m'}{\sqrt{m}} \sum_{i=1}^{m'} \frac{1}{m'} \sqrt{n_i \log n_i} \\ & \leq C \frac{m'}{\sqrt{m}} \sqrt{\left(\sum_{i=1}^{m'} \frac{n_i}{m'} \right) \log \left(\sum_{i=1}^{m'} \frac{n_i}{m'} \right)} \\ & = C \sqrt{\frac{m'}{m}} \sqrt{\left(\sum_{i=1}^{m'} n_i \right) \left(\log \left(\sum_{i=1}^{m'} n_i \right) - \log(m') \right)} \\ & \leq C \sqrt{n \log n}. \end{aligned}$$

□

Remark: With minor modifications, Lemma 15 can be applied to regions with shapes other than squares. Defining the diameter of a two-dimensional region as the diameter of the region's smallest enclosing circle, the main requirement for the modification to work is that the maximum diameter of these regions is $O(1/\sqrt{m})$.

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

Theorem 16: Suppose that n robots and n targets are uniformly randomly distributed in the unit square Q , and Q is divided into m_i equal-sized small squares at hierarchy i with a total of $h \geq 2$ hierarchies. For all applicable $i \geq 1$, assume that $m_{i+1} \geq m_i$ and any small square at hierarchy $i + 1$ falls within a single square at hierarchy i . Then Strategy 2 yields

$$\mathbf{E}[D_n] \leq C\sqrt{n \log n} + \sum_{i=1}^{h-1} \sqrt{\frac{nm_{i+1}}{m_i}}. \quad (17)$$

Proof: The $C\sqrt{n \log n}$ term on the RHS of (17) is due to Lemma 15. Then at each hierarchy i with $1 \leq i < h$, each of the m_i squares contains m_{i+1}/m_i smaller squares from hierarchy $i + 1$. Here we use the assumption that a region at a higher numbered hierarchy falls completely within a single region at a lower numbered hierarchy. This means that a robot that gets matched at hierarchy i needs to travel at most a distance of $\sqrt{2/m_i}$. Since there are no more than $\sqrt{n(m_{i+1} - 1)/2} < \sqrt{m_{i+1}n/2}$ unmatched robots at hierarchy i in expectation by Lemma 14, the distance incurred at hierarchy i is no more than $\sqrt{nm_{i+1}/m_i}$ for $1 \leq i < h$. Summing up all the distances then gives us the inequality (17). □

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 $\{m_i\}$ independent of n , the first term $C\sqrt{n \log n}$ dominates the other terms in (17) as $n \rightarrow \infty$. This implies that Strategy 2 yields assignments whose total distance is at most a constant multiple of the optimal distance. This observation is summarized in Corollary 17. Recall that D_n^* is the minimum possible distance defined by (6).

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

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

$$\begin{aligned} \mathbf{E}[D_n] & \leq C\sqrt{n \log n} + \sum_{i=1}^{h-1} \sqrt{4n} \\ & = C\sqrt{n \log n} + 2(h-1)\sqrt{n}. \end{aligned} \quad (18)$$

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

$$\mathbf{E}[D_n] \leq C\sqrt{n \log n} + \sum_{i=1}^2 \sqrt{n \log n} = (C+2)\sqrt{n \log n}. \quad (19)$$

Since hierarchical strategies need not 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.

Remark: Before concluding this section, it is worth mentioning that the results of this section continue to hold in only slightly weaker forms when the point sets X^0, Y^0 are drawn i.i.d. from the same arbitrary distribution over $[0, 1]^2$ (based

on Talagrand [35]). Since the topic of arbitrary probability measures diverges from the main focus of this paper, we only briefly discuss extending the results of this section to deal with arbitrary probability measures on $[0, 1]^2$.

To adapt Lemma 14 for arbitrary probability measures, assume that each region q_i (see the proof of Lemma 14) has an overall probability of p_i of receiving a robot or target. Note that $\sum_{i=1}^m p_i = 1$. This changes the upper bound of $\mathbf{E}[|S_n|]$ for the region q_i to $\sqrt{2np_i(1-p_i)}$. Then, over all m regions, the total number of unmatched robots is bounded by

$$\begin{aligned} \sum_{i=1}^m \sqrt{2np_i(1-p_i)} &= m\sqrt{2n} \sum_{i=1}^m \frac{1}{m} \sqrt{p_i(1-p_i)} \\ &\leq m\sqrt{2n} \sqrt{\sum_{i=1}^m \frac{p_i}{m} \left(1 - \sum_{i=1}^m \frac{p_i}{m}\right)} \\ &= m\sqrt{2n} \sqrt{\frac{1}{m} \left(1 - \frac{1}{m}\right)} \\ &= \sqrt{2n(m-1)} \end{aligned}$$

in which the inequality is obtained by applying Jensen's inequality to the concave function $\sqrt{x(1-x)}$.

Besides updating the uniform distribution of X^0 and Y^0 to an arbitrary probability measure, the statement and proof of Lemma 15 remain largely unchanged. This is because the second inequality in (16) does not change asymptotically as the underlying robot and target distribution changes. Then, the inequality (17) from Theorem 16 merely adds a multiplicative constant of 2 to its second term on the RHS. Because the first inequality in (16) is not known to hold for arbitrary probability measures, we do not have a parallel of Corollary 17 for arbitrary probability measures. Nevertheless, these bounds for arbitrary probability measures suggest that the uniform distribution is among the *worst* distributions for Problem 2 under the optimality constraint of minimizing (5). This is because the uniform distribution leads to an optimal assignment distance of $\Omega(\sqrt{n \log n})$, and an arbitrary distribution leads to an optimal assignment distance of $O(\sqrt{n \log n})$. Note that these updates also apply to the results in the next section with appropriate modifications.

VI. NEAR OPTIMAL STRATEGIES

After exploring hierarchical strategies for the region-based Communication Model 2, we now return to the range-based Communication Model 1. If r_{comm} is arbitrary and the conditions specified in Theorem 4 are not known to hold, the best we can do is obtain near distance-optimal strategies. In this section, we show that constant ratio approximation of distant optimality is possible for arbitrary r_{sense} and r_{comm} . The basic idea behind our strategies is to move the robots to pass around information about the locations of other robots. The assumption $r_{\text{sense}} \geq \sqrt{2}$ is made temporarily. At the end of this section, we show how to remove this assumption without affecting asymptotic optimality.

A. Near Distance-Optimal Rendezvous Strategy

Our first suboptimal strategy uses moving robots for information aggregation until some 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. Although some robots will move and change their locations during this process, the moved robots nevertheless are aware of their initial locations in X^0 . 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 Fig. 7).

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 does the same. The procedure continues until a robot reaches the middle of Q (row $\lceil b/2 \rceil$). Then, the robots in the squares on 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 moved robots then reverse their travel directions to deliver the assignment information to all robots. The outline of the strategy is given in Strategy 3.

Strategy 3: RENDEZVOUS

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

Outcome: produces a permutation σ that assigns robots to targets and communicates σ to all the 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 $q_{i,j}$
 - 2 **for** each $q_{i,j}$, $1 \leq i, j \leq b = \lceil \sqrt{2}/r_{\text{comm}} \rceil$ **do**
 - 3 **if** $i \neq \lceil b/2 \rceil$ **then**
 - 4 $waitTime \leftarrow \lceil \lceil b/2 \rceil - i \rceil / b$
 - 5 **else**
 - 6 $waitTime \leftarrow 1/2 + \lceil \lceil b/2 \rceil - j \rceil / b$
 - 7 $a_{i,j}$ waits for up to $waitTime$ units of time for information from a robot coming from the previous square. After the information is received or after $waitTime$ passes, $a_{i,j}$ starts moving to the next squares and delivers its information once it can communicate with another robot in these squares. It then stops
 - 8 robot $a_{\lceil b/2 \rceil, \lceil b/2 \rceil}$ computes σ ; the earlier communication process is then reversed to deliver σ to all the robots.
-

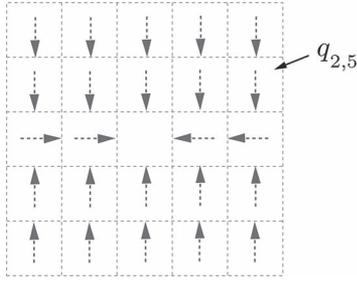


Fig. 7. Directions for robots to move in the rendezvous strategy.

The correctness of Strategy 3 as an algorithm is proven by construction. Besides the distance cost 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, in expectation, $D_n < D_n^* + 2b + 2$. Since $D_n^* = \Theta(\sqrt{n \log n})$, D_n^* 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 (13) 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 $T_n \approx 2n + T_n^*$ in expectation, which is undesirable 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

We first look at a decentralized hierarchical strategy that combines 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 $q_{i,j}$ 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), as $n \rightarrow \infty$

$$\mathbf{E}[D_n] \leq C\sqrt{n \log n} + \sqrt{mn} + 2\sqrt{m} + 2 \quad (20)$$

and

$$\mathbf{E}[T_n] = \Theta(\sqrt{n \log n} + \sqrt{mn}). \quad (21)$$

Proof: The bound on $\mathbf{E}[D_n]$, given by (20), is straightforward to compute using Theorem 16, in which the first two terms on the right side of (20) correspond to the first and second terms of the right side of (17), respectively, and the last two terms are due to communication overhead. For total completion time, all but $\Theta(\sqrt{mn})$ robots can start moving to their targets at $t = 0$. For the $\Theta(\sqrt{mn})$ robots, they need to wait no more than two units of time each before moving to their targets. This gives us (21). \square

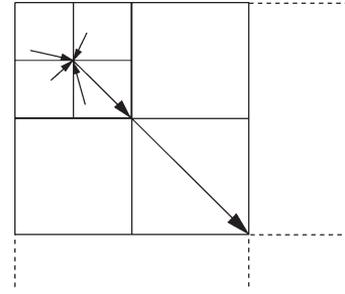


Fig. 8. Illustration of robot movements in a potential hierarchical strategy.

Remark: Similar to Strategy 3, for any fixed m , in expectation, $D_n/D_n^* = O(1)$ as $n \rightarrow \infty$. Moreover, in contrast to Strategy 3, for any fixed m , $T_n/T_n^* = O(1)$ in expectation. Suppose that a centralized algorithm requires $t(n)$ running time. Using the same centralized algorithm, Strategy 4 has a running time of $O(mt(n/m) + t(\sqrt{mn}))$. 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 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 a 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 Theorem 16. For example, we look at an extension to Strategy 4 with three hierarchies; let 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 the following.

Corollary 19: For Strategy 5 (3-level Hierarchical Rendezvous), as $n \rightarrow \infty$

$$\mathbf{E}[D_n] \leq C\sqrt{n \log n} + 2\sqrt{n\sqrt{m}} + 4\sqrt{m} + 2. \quad (22)$$

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 increased by approximately $2\sqrt{m}$. Thus, the total completion time of Strategy 5 is also given by (21). Following similar analysis, the overall running 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.

C. Handling Arbitrary r_{sense}

Because there can be targets anywhere in Q , to carry out the algorithms stated in this section, each robot must be aware of all target locations. For this to happen for arbitrary r_{sense} , Q must be swept through in a worst scenario. To do this, we partition Q

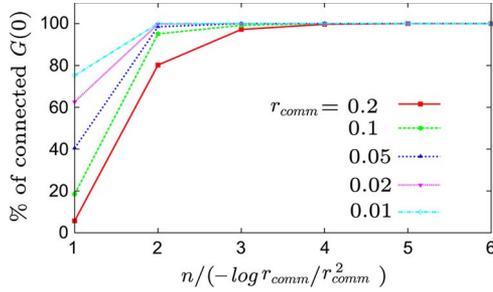


Fig. 9. Effects of n on the connectivity of $G(0)$ for different values of r_{comm} .

into $\lceil 1/(2r_{\text{sense}}) \rceil^2$ small squares and let a robot in the top-left small square “zig-zag” through Q (i.e., following a Boustrophedon path [36]) until it covers the bottom side of Q . If there is no robot in the top-left small square, then a robot in a square along the Boustrophedon path is used; implicit timing can be used to determine this. Once the end of the path is reached, the robot then reverses its course until it gets back to the top-left small square. At this point, this robot is aware of all target locations. It can then repeat a similar path to communicate that information to all other robots. This procedure ensures that all robots are aware of all target locations. The total distance cost of the procedure is about $2\lceil 1/(2r_{\text{sense}}) \rceil + \lceil 1/(2r_{\text{comm}}) \rceil$. Taking this penalty, which does not depend on n and therefore has no impact on the asymptotic optimality, we can then effectively assume $r_{\text{sense}} \geq \sqrt{2}$.

VII. SIMULATION STUDIES

A. Number of Required Robots for a Connected $G(0)$

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 verify 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} , various numbers of robots are used starting from $n = -\log r_{\text{comm}}/r_{\text{comm}}^2$ (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 percentage of the runs with a connected $G(0)$ is reported in Fig. 9. The simulation suggests that the bounds on n from Theorem 10 are fairly tight.

To compare to (4), which also allows for estimation of n in terms of r_{comm} with a specified probability for obtaining a connected $G(0)$, we computed n based on (4) and (7) for a range of r_{comm} -probability pairs. 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 probability numbers separated by a comma, obtained using (4) and (7), 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, (7) provides overestimates that guarantee the desired probability.

B. Performance of Near Optimal Strategies

Next, we simulate Strategies 2–5 and evaluate D_n , T_n , and running time for these strategies over various values of

TABLE I
COMPARISON BETWEEN (4) AND (7)

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

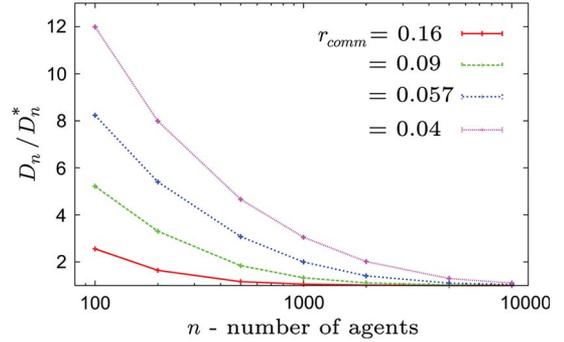


Fig. 10. Distance optimality of Strategy 3 over varying n and r_{comm} .

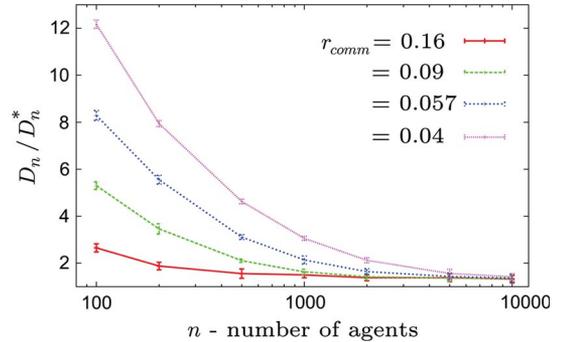


Fig. 11. Distance optimality of Strategy 4 over varying n and r_{comm} .

n and r_{comm} , assuming $r_{\text{sense}} \geq \sqrt{2}$. Due to our choice of $k = \sqrt{m}$ in Strategy 5, we pick specific r_{comm} 's so that $m = \lceil \sqrt{2}/r_{\text{comm}} \rceil$ is always a 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 assignment problem instances are randomly generated. These problem instances are then used to test all strategies. We test Strategy 2 using the same (two-hierarchy and three-hierarchy) partitions that are used with Strategies 4 and 5.

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 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) error bars. The error bars 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. The similarities between Fig. 10 and Fig. 11 for small n are not surprising since both strategies spend most of their effort (distance traveled) in establishing

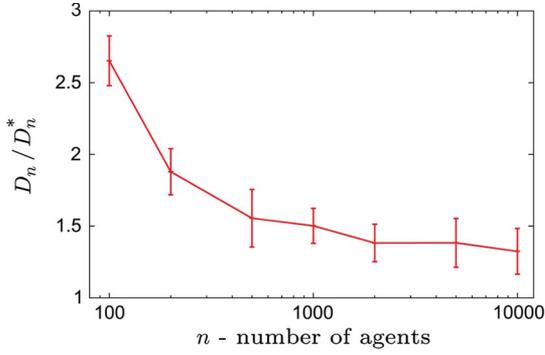


Fig. 12. The effect of varying n on the distance optimality of Strategy 4 with $r_{\text{comm}} = 0.16$ ($m = 81$).

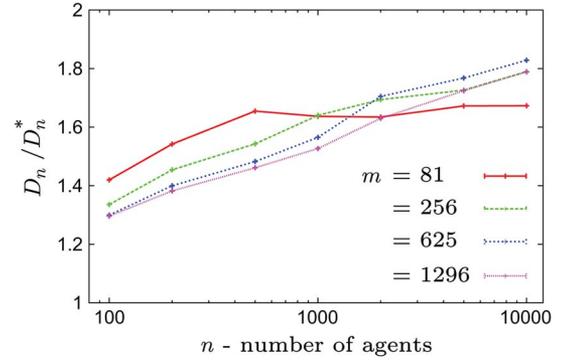


Fig. 15. The assignment cost of a three-level “pure” hierarchical strategy.

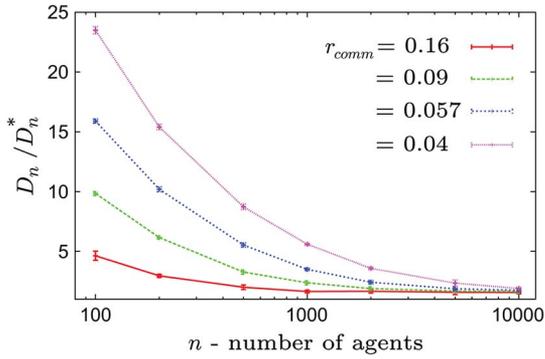


Fig. 13. Distance optimality of Strategy 5 over varying n and r_{comm} .

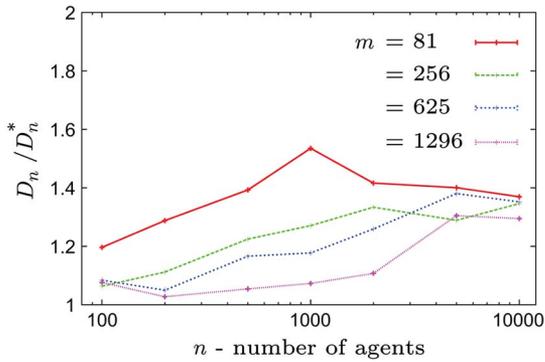


Fig. 14. Assignment cost of a two-level “pure” hierarchical strategy.

communication. As this extra communication cost diminishes as n grows, the actual assignment cost dominates. Strategy 3, with assignment being done in a centralized manner, becomes better than the decentralized Strategy 4.

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 partition of the unit square Q gets finer, D_n/D_n^* increases, implying that decreasing the communication radius has a negative effect on distance optimality. We observe similar results on the distance optimality of Strategy 5 (see Fig. 13).

If we remove the rendezvous part from Strategies 4 and 5, they become similar to Strategy 2. The distance optimality

TABLE II
RUNNING 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

performance of these two particular versions of Strategy 2 is shown in Figs. 14 and 15, respectively. For all partitions 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 strategies can provide very good optimality guarantees.

Computational Performance: We list the running time, in seconds, for Strategies 3–5 in Table II. 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 running time of Strategies 3, 4, and 5, respectively, for the given combination of r_{comm} and n . Note that any version of Strategy 2 has the same amount of computation as a corresponding rendezvous-based strategy. As expected, a hierarchical assignment greatly reduces the computation time, often by a factor over 10^3 . The computation was performed on an Intel Core-i7 3970K PC under a 8 GB Java virtual machine.

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 match T_n^* closely. For example, in (21), although $T_n \rightarrow T_n^*$ as $n \rightarrow \infty$ for fixed $m = \lceil \sqrt{2}/r_{\text{comm}} \rceil^2$, it requires a very large n for $\sqrt{\log n}$ to dominate \sqrt{m} . Thus, we only compare T_n among

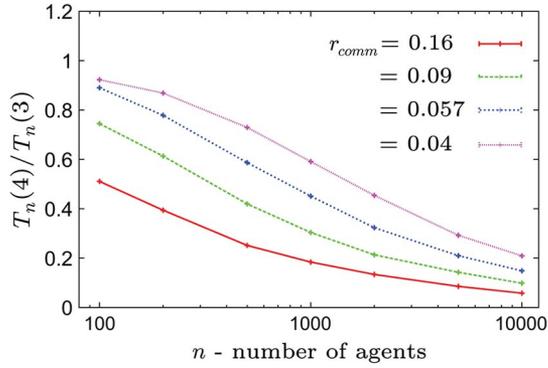


Fig. 16. Ratio of total completion time between Strategies 3 and 4.

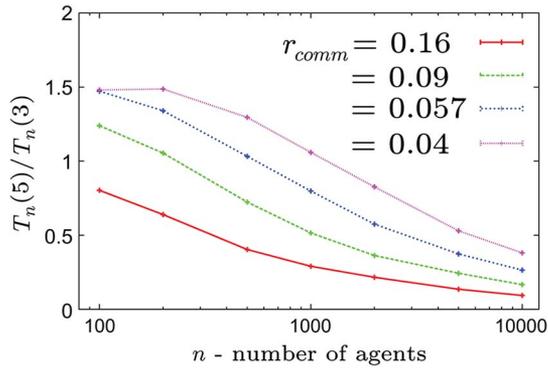


Fig. 17. Ratio of total completion time between Strategies 3 and 5.

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 Figs. 16 and 17. As n increases, Strategies 4 and 5 both take much less total completion time on average.

VIII. CONCLUSION

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 optimality can be achieved. We also provided a direct formula for computing the number of robots sufficient for probabilistically guaranteeing such an optimal solution. Then, we took a different angle; we looked at suboptimal strategies and their asymptotic performance as the number of robots goes to infinity. We showed that these strategies yield a constant approximation ratio when compared with the true distance optimal solution. Many of these decentralized strategies also provide computational advantages over a centralized one.

We conclude the paper by discussing our choice on certain elements that can be generalized in a future work.

Equal Number of Initial and Target Locations: In the problem statement we assume that $|X^0| = |Y^0|$. If $|X^0| > |Y^0|$, some robots do not need to move and if $|X^0| < |Y^0|$, some robots may need to reach multiple targets, assuming that the main goal is to serve the targets. Our result readily generalizes to the case in which $|X^0|/|Y^0|$ is close to 1. 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 [1]. 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 $C \in [0.1, 10]$).

Distribution of Initial and Target Locations: Although it is beyond the scope of this paper, it would be interesting to establish a lower bound on the optimal assignment distance for arbitrary probability measures. Also, it would be interesting to investigate the case in which the robots and the targets assume different distributions. Another important aspect not covered in this paper is the issue of targets distributed somewhat randomly over time.

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 [34]. That is, for

$$D_{n,p}^* := \min_{\sigma} \sum_{i=1}^n \|x_i^0 - y_{\sigma(i)}^0\|_2^p \quad (23)$$

it holds true that

$$D_{n,p}^* \sim n(\log n/n)^{\frac{p}{2}}. \quad (24)$$

Theorem 13 corresponds to the special case of $p = 1$. As $p \rightarrow \infty$, (23) minimizes the longest distance traveled by any robot. This is true because for fixed X^0, Y^0 , and a sufficiently large p , the largest $\|x_i^0 - y_{\sigma(i)}^0\|_2^p$ becomes the dominating term in the sum $\sum_{i=1}^n \|x_i^0 - y_{\sigma(i)}^0\|_2^p$. Although we restrict our attention to $p = 1$ in this paper, our results readily extend to other values of p (i.e., other optimality criteria) with (24). Note that this means the D_n definition given by (5) needs to be updated accordingly to an appropriately defined $D_{n,p}$.

ACKNOWLEDGMENT

The authors wish to thank the anonymous reviewers and S. Har-Peled and A. Nayeri for their constructive comments.

REFERENCES

- [1] S. L. Smith and F. Bullo, "Monotonic target assignment for robotic networks," *IEEE Trans. Autom. Control*, vol. 54, no. 9, pp. 2042–2057, Sep. 2009.
- [2] J. Yu and S. M. LaValle, "Distance optimal formation control on graphs with a tight convergence time guarantee," in *Proc. IEEE Conf. Decision Control*, 2012, pp. 4023–4028.
- [3] D. P. Bertsekas, "The auction algorithm: A distributed relaxation method for the assignment problem," *Annals Oper. Res.*, vol. 14, pp. 105–123, 1988.
- [4] D. P. Bertsekas and D. A. Castañón, "Parallel synchronous and asynchronous implementations of the auction algorithm," *Parallel Comp.*, vol. 17, pp. 707–732, 1991.
- [5] R. Burkard, M. Dell'Amico, and S. Martello, *Assignment Problems*. Philadelphia, PA: SIAM, 2012.
- [6] J. Edmonds and R. M. Karp, "Theoretical improvements in algorithmic efficiency for network flow problems," *J. ACM*, vol. 19, no. 2, pp. 248–264, 1972.
- [7] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Res. Logistics Quart.*, vol. 2, pp. 83–97, 1955.

[8] M. M. Zavlanos, L. Spesivtsev, and G. J. Pappas, "A distributed auction algorithm for the assignment problem," in *Proc. IEEE Conf. Decision Control*, 2008, pp. 1212–1217.

[9] M. Ji, S. Azuma, and M. Egerstedt, "Role-assignment in multi-agent coordination," *Int. J. Assitive Robot. Mechatron.*, vol. 7, no. 1, pp. 32–40, Aug. 2006.

[10] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Flocking in fixed and switching networks," *IEEE Trans. Autom. Control*, vol. 52, no. 5, pp. 863–868, 2007.

[11] K. Treleaven, M. Pavone, and E. Frazzoli, "Asymptotically optimal algorithms for one-to-one pickup and delivery problems with applications to transportation systems," *IEEE Trans. Autom. Control*, vol. 59, no. 9, pp. 2261–2276, 2013.

[12] M. M. Zavlanos and G. J. Pappas, "Dynamic assignment in distributed motion planning with local coordination," *IEEE Trans. Robotics*, vol. 24, no. 1, pp. 232–242, 2008.

[13] S.-J. Chung, S. Bandyopadhyay, I. Chang, and F. Y. Hadaegh, "Phase synchronization control of complex networks of Lagrangian systems on adaptive digraphs," *Automatica*, vol. 49, no. 5, pp. 1148–1161, 2013.

[14] S. Kloder and S. Hutchinson, "Path planning for permutation-invariant multirobot formations," *IEEE Trans. Robotics*, vol. 22, no. 4, pp. 650–665, 2006.

[15] V. Sharma, M. A. Savchenko, E. Frazzoli, and P. G. Voulgaris, "Transfer time complexity of conflict-free vehicle routing with no communications," *Int. J. Robot. Res.*, vol. 26, no. 3, pp. 255–271, 2007.

[16] M. Turpin, K. Mohta, N. Michael, and V. Kumar, "Goal assignment and trajectory planning for large teams of aerial robots," in *Proc. Robotics: Sci. Syst.*, 2013.

[17] J. Cortés, S. Martínez, and F. Bullo, "Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions," *IEEE Trans. Autom. Control*, vol. 51, no. 8, pp. 1289–1298, Aug. 2006.

[18] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, 2003.

[19] J. Lin, A. S. Morse, and B. D. O. Anderson, "The multi-agent rendezvous problem. Part 1: The synchronous case," *SIAM J. Control Optim.*, vol. 46, no. 6, pp. 2096–2119, Nov. 2007.

[20] J. Lin, A. S. Morse, and B. D. O. Anderson, "The multi-agent rendezvous problem. Part 2: The asynchronous case," *SIAM J. Control Optim.*, vol. 46, no. 6, pp. 2120–2147, Nov. 2007.

[21] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*. Princeton, NJ: Princeton Univ. Press, ser. Applied Mathematics Series, 2009. [Online]. Available: <http://coordinationbook.info>

[22] F. Xue and P. R. Kumar, "The number of neighbors needed for connectivity of wireless networks," *Wireless Networks*, vol. 10, no. 2, pp. 169–181, Mar. 2004.

[23] P. Balister, B. Bollobás, A. Sarkar, and M. Walters, "Connectivity of random k-nearest-neighbour graphs," *Adv. Appl. Prob.*, vol. 37, pp. 1–24, 2005.

[24] N. M. Freris, H. Kowshik, and P. R. Kumar, "Fundamentals of large sensor networks: Connectivity, capacity, clocks and computation," *Proc. IEEE*, vol. 98, no. 11, pp. 1828–1846, Nov. 2010.

[25] A. Ganesh and F. Xue, "On the connectivity and diameter of small-world networks," *Adv. Appl. Prob.*, vol. 39, pp. 853–863, 2007.

[26] G. Mao and B. D. O. Anderson, "Connectivity of large wireless networks under a general connection model," *IEEE Trans. Inform. Theory*, vol. 59, no. 3, pp. 1761–1772, 2013.

[27] M. Penrose, *Random Geometric Graphs*. London, U.K.: Oxford Univ. Press, 2003, ser. Oxford Studies in Probability.

[28] M. Penrose, "The longest edge of the random minimal spanning tree," *Annals Appl. Prob.*, vol. 7, pp. 340–361, 1997.

[29] J. Yu, S.-J. Chung, and P. G. Voulgaris, "Distance optimal target assignment in robotic networks under communication and sensing constraints," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2014, pp. 1098–1105.

[30] J. Yu, S.-J. Chung, and P. G. Voulgaris, "Traveled distance minimization and hierarchical strategies for robotic networks," in *Proc. Int. Symp. Commun., Control, Signal Processing, Special Session Modeling Control Complex Netw.*, May 2014, pp. 491–496.

[31] P. Erdős and A. Rényi, "On a classical problem of probability theory," *Publ. Math. Inst. Hung. Acad. Sci.*, vol. Ser. A 6, pp. 215–220, 1961.

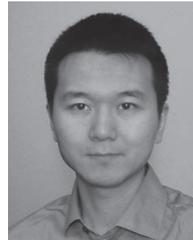
[32] P. M. Vaidya, "Geometry helps in matchings," *SIAM J. Comp.*, vol. 18, no. 6, pp. 1201–1225, 1989.

[33] R. Sharathkumar and P. K. Agarwal, "A near-linear time ϵ -approximation algorithm for geometric bipartite matching," in *Proc. 44th Symp. Theory Comp. Conf.*, 2012, pp. 385–394.

[34] M. Ajtai, J. Komlós, and G. Tusnády, "On optimal matchings," *Combinatorica*, vol. 4, no. 4, pp. 259–264, 1984.

[35] M. Talagrand, "The Ajtai-Komlós-Tusnády matching theorem for general measures," *Prob. Banach Spaces*, vol. 30, pp. 39–54, 1992.

[36] H. Choset, "Coverage of known spaces: The boustrophedon cellular decomposition," *Autonomous Robots*, vol. 9, pp. 247–253, 2000.



Jingjin Yu (S'11–M'13) received the B.S. degree in materials science from the University of Science and Technology of China, Hefei, China, in 1998, the M.S. degree in chemistry from the University of Chicago, Chicago, IL, USA, in 2000, the M.S. degree in mathematics from the University of Illinois at Chicago, Chicago, IL, USA, in 2001, and the M.S. degree in computer science and Ph.D. degree in electrical engineering from the University of Illinois at Urbana-Champaign, Urbana, IL, USA, in 2010 and 2013, respectively.

He is currently with the Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, MA, USA.



Soon-Jo Chung (M'06–SM'12) received the B.S. degree (with highest honors) from the Korea Advanced Institute of Science and Technology, Daejeon, Korea, in 1998 and the S.M. degree in aeronautics and astronautics and the Sc.D. degree in estimation and control from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2002 and 2007, respectively.

He is currently an Assistant Professor in the Department of Aerospace Engineering and the Coordinated Science Laboratory, the University of Illinois at Urbana-Champaign (UIUC), Urbana, IL, USA. He was a Summer Faculty Fellow with the Jet Propulsion Laboratory, working on guidance and control of spacecraft swarms during the summers of 2010–2014. His research areas include nonlinear control and estimation theory and flight guidance and control, with application to aerial robotics, distributed spacecraft systems, and vision-based navigation.

Dr. Chung received the UIUC Engineering Dean's Award for Excellence in Research, the Beckman Fellowship of the UIUC Center for Advanced Study, the U.S. Air Force Office of Scientific Research Young Investigator Award, the National Science Foundation Faculty Early Career Development Award, and two best conference paper awards from the IEEE and The American Institute of Aeronautics and Astronautics.



Petros G. Voulgaris (F'11) received the Diploma in mechanical engineering from the National Technical University, Athens, Greece, in 1986, and the S.M. and Ph.D. degrees in aeronautics and astronautics from the Massachusetts Institute of Technology, Cambridge, in 1988 and 1991, respectively.

Since August 1991, he has been with the Department of Aerospace Engineering, University of Illinois at Urbana-Champaign, where he is currently a Professor. He also holds joint appointments with the Coordinated Science Laboratory, and the department of Electrical and Computer Engineering at the same university. His research interests include robust and optimal control and estimation, communications and control, networks and control, and applications of advanced control methods to engineering practice including flight control, nano-scale control, robotics, and structural control systems.

Dr. Voulgaris received the National Science Foundation Research Initiation Award (1993), the Office of Naval Research Young Investigator Award (1995) and the UIUC Xerox Award for research. He has been an Associate Editor for the IEEE TRANSACTIONS ON AUTOMATIC CONTROL and the ASME *Journal of Dynamic Systems, Measurement and Control*.