

Distance Optimal Target Assignment in Robotic Networks under Communication and Sensing Constraints

Jingjin Yu¹ Soon-Jo Chung² Petros G. Voulgaris²

Abstract—We study the problem of minimizing the total distance incurred in assigning a group of mobile robots to an equal number of static targets. Assuming that the robots have limited, range-based communication and target-sensing capabilities, we present a necessary and sufficient condition for ensuring distance optimality when robots and targets are uniformly randomly distributed. We then provide an explicit, non-asymptotic formula for computing the number of robots needed for guaranteeing optimality in terms of the robots' sensing and communication capabilities with arbitrarily high probabilities. The bound given in the formula is also asymptotically tight. Due to the large number of robots needed for high-probability optimality guarantee, we continue to investigate strategies for cases in which the number of robots cannot be freely chosen. We show that a properly designed strategy can be asymptotically optimal or suboptimal with constant approximation ratios.

I. INTRODUCTION

In this paper, we study permutation-invariant assignments of a set of networked robots to a set of targets of equal cardinality, with a primary focus on minimizing the total path distance. Both robots and targets are assumed to be uniformly randomly distributed in a two-dimensional unit square. Under communication and target-sensing limitations, we seek optimality guarantees, in terms of necessary and sufficient conditions, as well as asymptotically optimal or suboptimal strategies when the conditions for optimality cannot be satisfied. In characterizing the performance of suboptimal strategies, we show that these strategies can often provide constant ratio approximations with respect to distance optimality.

Our problem considers the problem of *target assignment in robotic networks*. This problem is studied by, among others, Smith and Bullo in [25], in which the performance of several classes of algorithms for achieving time optimality (i.e., minimizing the time until every target is occupied) were established. In contrast, we focus on minimizing the total distance traveled by all robots. The total distance serves as a proper proxy to quantities such as the energy consumption of all robots. Simple examples show that a distance-optimal

solution for the target assignment problem generally does not imply time optimality and vice versa [33].

The problem of target assignment in robotic networks requires the resolution of an *assignment* (or *matching*) problem. Assignment problems are extensively studied in the area of combinatorial optimization, with efficient algorithms available for solving many of its variations [1], [4], [5], [7], [11], [18], [35]. If we instead put more emphasis on multi-robot systems, the problems of robotic task allocation [16], [27], [28], [34], swarm reconfiguration [9], multi-robot path planning [17], [24], [29], and multi-agent consensus [10], [15], [19], [20] come up. For a review on some of these topics, see [6].

Our work is also related to the study on the connectivity of wireless networks. If n robots are uniformly randomly distributed in a unit square, then each robot needs to have $k = \Theta(\log n)$ nearest neighbors for the entire network to be asymptotically connected [31]. In particular, the authors of [31] 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 [3]. These nearest neighbor based connectivity models were further studied in [13], [14], [21], to list a few. In these work, a *geometric graph* structure is often used [23].

Our contribution is twofold. First, for robots with arbitrarily limited range-based 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 r_{comm}). In contrast to related connectivity results [22], [31], we give n as an explicit function of r_{comm} without asymptotic assumptions. 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, adopting results on one-dimensional random walk, we 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 the approximation ratio 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) can be achieved.

The rest of the paper is organized as follows. In Section

¹Jingjin Yu is with the Computer Science and Artificial Intelligence Lab at the Massachusetts Institute of Technology and the Mechanical Engineering Department at Boston University. E-mail: jingjin@csail.mit.edu.

²Soon-Jo Chung and Petros G. Voulgaris are with the Coordinated Science Lab and the Department of Aerospace Engineering, University of Illinois at Urbana-Champaign. E-mail: {sjchung, vougari}@illinois.edu.

This work was supported in part by AFOSR grant FA95501210193 and NSF grant IIS-1253758. This paper is intended as an early dissemination of results of an extended draft [32] which contains more complete proofs and significant generalizations. We thank the reviewers for their constructive comments.

II, we introduce notations and formally define the problem that we study. Sections III and IV then elaborate on the two stated contributions, one contribution per section. We conduct simulations in Section V to confirm our theoretical findings and conclude in Section VI. Due to limited space, some proofs are sketched or omitted; for complete proofs, see [32].

II. PROBLEM STATEMENT

We introduce some notations before formally defining the target assignment problem. 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 . $|\cdot|$ denotes the cardinality for sets and the absolute values for real numbers. 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 \mathcal{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 $f, g: \mathbb{R}^+ \rightarrow \mathbb{R}^+$, $f(x) = O(g(x))$ (resp. $\Omega(g(x))$) if $\lim_{x \rightarrow \infty} f(x)/g(x) < \infty$ (resp. $\lim_{x \rightarrow \infty} f(x)/g(x) \geq c > 0$). Note that under this context, “=” behaves as “ \in ”. 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 $\neg(f(x) = \Theta(x))$.

We now move to stating the problem. Let $X^0 = \{x_1^0, \dots, x_n^0\}, Y^0 = \{y_1^0, \dots, y_n^0\} \subset \mathcal{Q}$ be two uniformly randomly selected point sets of cardinality n .¹ The superscript emphasizes that these points are obtained at the start time $t = 0$. Place n 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 control space for a robot a_i is $\dot{x}_i = u_i$ 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: A robot a_i may communicate with other robots within a disc of radius r_{comm} centered at $x_i(t)$.² At any given time $t \geq 0$, we define the (undirected) *communication graph* $G(t)$ as follows, which is a geometric graph [23] 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 $\|x_i(t) - x_j(t)\|_2 \leq r_{\text{comm}}$.

Since the communication overhead is often negligible with respect to the time it takes for 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 as far as decision making is concerned.

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 is defined as follows.

Problem 1 (Target Assignment in Robotic Networks)

Given $X^0, Y^0, r_{\text{comm}}$, and r_{sense} , 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 1, we are interested in minimizing the total distance traveled by all robots, which can be expressed as

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

As a proper proxy to measures such as the energy consumption of the entire system, the cost defined in (1) is an appropriate objective in practice. Unless otherwise specified, *optimality* refers to minimizing D_n in this paper. 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 1 given a fixed set of capabilities for the robots, $\inf_{\mathcal{U}} D_n$ is then the greatest lower bound achievable under these capabilities.³

III. GUARANTEEING OPTIMALITY FOR ARBITRARY r_{comm} AND r_{sense}

Intuitively, without global communication, optimality can be hard to guarantee (i.e., $\inf_{\mathcal{U}} D_n = D_n^*$), because global assignment is not possible in general at $t = 0$. For example, as $r_{\text{sense}} \rightarrow 0$, the robots must search for the targets before assignment can be made; it is 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: For arbitrary fixed r_{comm} and r_{sense} , under what conditions can we ensure optimality? This question is answered in the following Theorem.

Theorem 1 *Under sensing and communication constraints, $\inf_{\mathcal{U}} D_n = D_n^*$ 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

³Here \inf is used instead of \min because it is not immediately clear that the minimum can always be reached.

¹Via scaling, our result readily extends to arbitrary square environments.

²While discs may be too simple for modeling the communication or sensing range of a robot precisely, they remain valid for purposes such as establishing performance bounds (the subject of this paper).

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 $x \in X^0$.

For establishing the first claim, note that the robots must decide at $t = 0$ a pairing between elements of X^0 and Y^0 that minimizes D_n . We now show that this is impossible in general for $n = 2$ and $r_{\text{comm}} < \sqrt{2}$. A possible configuration of two robots and two targets (a_1, a_2, y_1, y_2 , respectively) under these assumptions is given in Fig. 1 (solid blue and red dots). Because they are more than r_{comm} apart, the robots cannot communicate. 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 must have a_1 go to y_1 and a_2 go to y_2 . However, 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 may expand the locations of the robots and targets to include neighborhoods around them (the dotted circles in Fig. 1) to establish that there is a non-zero probability with which an optimal assignment cannot be made at $t = 0$. Thus, for distance optimality, $G(0)$ cannot have more than one connected component and must be connected.




Fig. 1. A general setup in which two robots cannot communicate 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 the robot to follow a path that is not a straight line with probability one, implying that $D_n = D_n^*$ with zero probability. It straightforward to see that the two necessary conditions from the two claims are also sufficient for optimality. \square

Theorem 1 provides a simple way of ensuring optimality by either increasing the number of robots or increasing r_{comm} and/or r_{sense} , which leads to a centralized 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. An optimal assignment in the unit square can be computed in $O(n^3)$ using the strongly polynomial⁴ Hungarian algorithm [11], [18] or other asymptotically faster algorithms [1], [30].

⁴A 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.

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 1 can be met. Although connectivity results on Random Geometric Graphs [22] can be used for this purpose, these results only yield implicit formulas of an asymptotic nature. We take a different approach and produce the number of robots n as an explicit function of r_{comm} , without the asymptotic assumption.

A. Guaranteeing a Connected $G(0)$

When the robots can be anywhere in the unit square \mathcal{Q} , given a communication radius of $r_{\text{comm}} < \sqrt{2}$, at least $\Theta(1/r_{\text{comm}}^2)$ robots are needed for a connected $G(0)$, which requires the robots 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 robots are needed to ensure a connected communication graph $G(0)$.

Lemma 2 *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*

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

PROOF. We divide the unit square \mathcal{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. 2) contains at least a robot, the robot can communicate with any other robot in the four squares sharing a side with q_i . Therefore, $G(0)$ is connected if




Fig. 2. 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.

each q_i contains a robot. Let n_i denote the number of robots in q_i . Then $\mathbf{P}(n_i = 0) = (1 - 1/m)^n < e^{-\frac{n}{m}}$. The inequality holds because $(1 - x)^n < e^{-nx}$ for $0 < x < 1$. 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

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

Setting $me^{-n/m} = \varepsilon$ and replacing $m = \lceil \sqrt{5}/r_{\text{comm}} \rceil^2$ shows that (2) guarantees that each small square contains at least one robot with probability $1 - \varepsilon$. \square

Remark. In contrast to asymptotic results (see, e.g., [22]), Lemma 2 provides n as an explicit function of r_{comm} . The sufficient condition on n given in (2) is non-asymptotic and 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 3 *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*

$$n \geq (2 \log \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil + c) \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2. \quad (3)$$

PROOF. Given the division scheme used in the proof of Lemma 2, distributing robots into the unit square Q is equivalent to tossing the robots (balls) into the m small squares (bins), uniformly randomly. By classical results on *balls and bins* (e.g., Inequality (2) from [12]), 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 3 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. This many robots turns out to be also necessary for the high probability guarantee, which we prove next.

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 arbitrarily small for n up to $m \log m$.

Lemma 4 *Suppose that $1 \leq n < m \log m$ balls are tossed uniformly randomly into m bins. As $m \rightarrow \infty$, $\mathbf{P}_{n,m}(E_1) > 0.34$.*

PROOF SKETCH. We sketch the general idea behind the proof due to limited space. Our proof partitions all $n \in [1, m \log m]$ into two pieces: $n \in [1, m]$ and $n \in (m, m \log m)$. When $1 \leq n \leq m$, since on average there is no more than one ball per bin, intuitively some bin must have exactly one ball in it. We can show that $\mathbf{P}_{n,m}(E_1) > e^{-1}$. For $m < n < m \log m$, we first establish that $\mathbf{P}_{n',m}(E_0)$ is large for n' up to $m \log m - m$ using again Inequality (2) from [12]. Using $\mathbf{P}_{n'+k,m}(E_1) \geq \mathbf{P}_{n',m}(E_0) \mathbf{P}_{k,m}(\text{exactly one ball falls in the empty bin})$, we can show that $\mathbf{P}_{n,m}(E_1) \geq 0.34$ for $m < n < m \log m$. \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 5 *For uniformly randomly distributed robots in a unit square with a communication radius r_{comm} ,*

$$n = \Theta\left(\frac{1}{r_{\text{comm}}^2} \log \frac{1}{r_{\text{comm}}}\right) \quad (4)$$

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 SKETCH. Lemma 3 covers sufficiency; we are to show that there is some non-trivial probability that $G(0)$ is disconnected if the number of robots satisfies $n = o((1/r_{\text{comm}}^2) \log(1/r_{\text{comm}}))$. To prove the claim, we partition the unit square Q into $m = b^2 = \lfloor 1.1/r_{\text{comm}} \rfloor^2$ small squares and group them into 3×3 blocks. We set the size of the square so that a robot in the center square of a 3×3 block cannot communicate with others if the other eight squares of the same block are unoccupied. By Lemma 4, the probability that one of the 3×3 blocks gets exactly one robot is non-trivial even when the number of robots is of order $m \log m$. Since there is $1/9$ chance the robot is in the center of the block, we are done. \square

B. Ensuring Target Observability

With a connected communication graph $G(0)$, 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 2 lead to a similar lower bound.

Lemma 6 *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*

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

Putting together Lemmas 2 and 6, we obtain a lower bound on n that makes a distance-optimal assignment possible; we omit the straightforward proof.

Theorem 7 *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*

$$n \geq \begin{cases} \lceil \frac{\sqrt{2}}{r_{\text{sense}}} \rceil^2 \log\left(\frac{1}{\varepsilon} \lceil \frac{\sqrt{2}}{r_{\text{sense}}} \rceil^2\right), & r_{\text{sense}} < \frac{\sqrt{10}r_{\text{comm}}}{5} \\ \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2 \log\left(\frac{1}{\varepsilon} \lceil \frac{\sqrt{5}}{r_{\text{comm}}} \rceil^2\right), & r_{\text{sense}} \geq \frac{\sqrt{10}r_{\text{comm}}}{5} \end{cases} \quad (6)$$

Remark. Theorem 7 is not an asymptotic result and works for all r_{comm} and r_{sense} . If high-probability asymptotic result is desirable, Lemma 6 can be easily turned into a version similar to Theorem 5, following essentially the same proof techniques; we omit the details. In view of this fact, the bounds from Theorem 7 are asymptotically tight.

IV. HIERARCHICAL ASYMPTOTICALLY OPTIMAL AND SUBOPTIMAL STRATEGIES

In view of Theorem 7, a large number robots may be necessary to guarantee an optimal solution with high probability. When such a guarantee is unavailable due to the lack of resources, alternative strategies are needed. In this section, we seek asymptotically optimal and suboptimal strategies for arbitrary r_{comm} and r_{sense} that do not require as many robots.

To make our strategies more modular, sensing and communication are treated orthogonally; we now handle the sensing part. Because there can be targets anywhere in Q , the robots' joint sensing area must cover all of Q to obtain all target locations. For this to happen for arbitrary r_{sense} , Q must be swept through. To achieve this, we partition Q into $\lceil 1/(2r_{\text{sense}}) \rceil^2$ small squares and let a robot in the top-left square "zig-zag" through Q (i.e., follows a Boustrophedon path [8]) until it covers the bottom side of Q . If there is no robot in the top-left 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 then repeats a similar path (the unit square is now divided into $\lceil 1/(2r_{\text{comm}}) \rceil^2$ small squares) 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 above procedure is at most $2\lceil 1/(2r_{\text{sense}}) \rceil + \lceil 1/(2r_{\text{comm}}) \rceil$. Taking this penalty, we assume that all robots are aware of all target locations.

A. Generic Abstract Hierarchical Strategy

Because the strategies to be proposed share the general feature of being *hierarchical*, we first characterize the performance of an abstract hierarchical strategy. Let $h \geq 1$ be the number of hierarchies and $m_i, 1 \leq i \leq h$, be the number of regions (disjoint squares within Q) at hierarchy i , we require that: 1. $m_1 \equiv 1$, 2. $m_{i+1} > m_i$, and 3. a region at a higher numbered hierarchy does not span multiple regions at a lower numbered hierarchy. Given such a setup, a straightforward strategy is to assign the robots to targets locally when possible and balance the surplus or deficit of robots at higher (lower numbered) hierarchies. The pseudo code of such a strategy is outlined in Strategy 2.

Strategy 2: GENERIC HIERARCHICAL STRATEGY

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
       and targets in region  $j$ , respectively; assign any
        $n_a$  (or  $n_g$  if  $n_a > n_g$ ) unmatched robots to any
       (equal number of) unmatched targets in the
       region

```

Note that in Strategy 2, we did not mention how the robots reach consensus under limited communication; this will be specified in each concrete strategy. The total distance D_n incurred by Strategy 2 consists of two parts: 1. D_n^a , the distance between robots and their assigned targets, and 2. D_n^c , the distance the robots must travel to compensate for the lack of global communication and sensing (i.e., the extra distance traveled for reaching consensus). We note that regardless of the strategy, the best possible D_n^a cannot be smaller than

$$D_n^* = \min_{\sigma} \sum_{i=1}^n \|x_{\sigma(i)}^0 - y_i^0\|_2, \quad (7)$$

in which \min_{σ} is taken over all permutations σ of the integers $1, \dots, n$. Ajtai, Komlós, and Tusnády proved the following.

Theorem 8 (Optimal Matching [2]) *With high probability,*

$$C_1 \sqrt{n \log n} \leq D_n^* \leq C_2 \sqrt{n \log n}, \quad (8)$$

in which C_1, C_2 are positive constants.

Although the authors did not provide formulas for C_1 and C_2 in [2], simulation seems to suggest that $C_1 < C_2 < 1$ and $C_2/C_1 \rightarrow 1$ as $n \rightarrow \infty$. Moreover, the second inequality in (8) also hold in expectation [26] with a universal constant C .

With $h \geq 2$ hierarchies, to bound D_n^a , at each hierarchy i , we need to know the number of robots that can be matched locally (Lemma 9) and the assignment distance incurred by these robots (Lemma 10).

Lemma 9 *Suppose that the unit square Q is divided into m equal-sized small squares. The number of robots that are not matched locally is $\sqrt{mn}/2$ in expectation.*

PROOF SKETCH. The process of picking X^0 and Y^0 is equivalent to picking (x_i^0, y_i^0) pairs for n times. For each small square q , we are interested in the events $x_i^0 \in q, y_i^0 \notin q$ and $x_i^0 \notin q, y_i^0 \in q$ (each having probability $(m-1)/m^2$). Combining these two aspects, in each small square we end up with a random walk of n steps on the integer line with each step having a probability of $(m-1)/m^2$ moving ± 1 . Applying Jensen's inequality to the concave function \sqrt{x} with x being the total variance summed over all small squares gives us the bound $\sqrt{mn}/2$. \square

Lemma 10 *Dividing the unit square into m equal sized squares and matching robots and targets within the boundaries of each small square, the total distance of matchings made this way is no more than $C\sqrt{n \log n}$ in expectation for some constant C .*

PROOF SKETCH. In a square q_i with n_i robots, local matching distance is bounded by $C\sqrt{n_i \log n_i/m}$ (see [26]). Applying Jensen's inequality to the concave function $\sqrt{x \log x}$ and letting $x = n_i$ yields the result. \square

We now give an upper bound on D_n^a .

Theorem 11 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. in expectation,

$$D_n^a \leq C\sqrt{n \log n} + \sum_{i=1}^{h-1} \sqrt{\frac{nm_{i+1}}{m_i}}. \quad (9)$$

PROOF. The $C\sqrt{n \log n}$ term is due to Lemma 10. Then at each hierarchy i with $2 \leq i < h$, the number of matched robots in total at this hierarchy is bounded by $\sqrt{m_{i+1}n/2}$. Since each of these robots needs to travel at most a distance of $\sqrt{2/m_i}$, we get the second term on the RHS side of (9). \square

Remark. We observe that for fixed h and $\{m_i\}$ that do not depend on n , the first term $C\sqrt{n \log n}$ dominates the other terms in (9) as $n \rightarrow \infty$. This implies that Strategy 2 yields assignments of which D_n^a is at most a multiple of the true optimal distance. As long as D_n^a is not dominated by D_n^c , an hierarchical strategy based on Strategy 2 achieves constant approximation ratio on distance optimality.

B. A Near-Optimal Rendezvous Strategy

Our first concrete strategy uses moving robots for communication until a robot is aware of the locations of all robots and targets, 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. 3).




Fig. 3. 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. 3. 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 $b = \lceil \sqrt{2}/r_{\text{comm}} \rceil$. The pseudo code of the strategy is given in Strategy 3.

Strategy 3 is correct by construction. Besides the distance from the assignment and sensing penalty, 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, $D_n \leq D_n^* + 2\lceil 1/(2r_{\text{sense}}) \rceil + \lceil 1/(2r_{\text{comm}}) \rceil + 2\lceil \sqrt{2}/r_{\text{comm}} \rceil + 2$. Since $D_n^* = \Theta(\sqrt{n \log n})$, it dominates the other terms when, for example, $n = \Theta(1/r_{\text{comm}}^2)$ and $n = \Theta(1/r_{\text{sense}}^2)$. Therefore, Strategy 3 yields asymptotically optimal solution without requiring an n as large as (6) with respect to $1/r_{\text{comm}}$ and $1/r_{\text{sense}}$.

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** $q_{i,j}$, $1 \leq i, j \leq b = \lceil \sqrt{2}/r_{\text{comm}} \rceil$
 - 2 **do**
 - 3 **if** $i \neq \lceil b/2 \rceil$ **then**
 - 4 $a_{i,j}$ waits for up to $\lceil \lceil b/2 \rceil - i \rceil / b$ 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
 - 5 **else**
 - 6 $a_{i,j}$ waits for up to $1/2 + \lceil \lceil b/2 \rceil - j \rceil / b$ 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.
-

Among other shortcomings, Strategy 3 requires running a centralized assignment algorithm for all robots, which can be computationally intensive for large n . Decentralized hierarchical strategies, to be discussed next, can address such issues.

C. Decentralized Hierarchical Strategies

By playing with h and $\{m_i\}$, many decentralized strategies are possible; we first look at one 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 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 12 For strategy 4 (2-level Hierarchical Rendezvous), in expectation,

$$D_n \leq C_2 \sqrt{n \log n} + 2\lceil 1/(2r_{\text{sense}}) \rceil + \lceil 1/(2r_{\text{comm}}) \rceil + \left\lceil \frac{\sqrt{2}}{r_{\text{comm}}} \right\rceil \sqrt{n} + 2\left\lceil \frac{\sqrt{2}}{r_{\text{comm}}} \right\rceil + 2. \quad (10)$$

PROOF. Simple application of Theorem 11. \square

Similar to Strategy 3, for any fixed r_{comm} and r_{sense} , $D_n/D_n^* = O(1)$ (as $n \rightarrow \infty$). Suppose that a centralized assignment algorithm requires time $t(n)$, using the same algorithm, Strategy 4 has a computational time complexity $O(mt(n/m) + t(\sqrt{mn}))$ (recall that $m = b^2 = \lceil \sqrt{2}/r_{\text{comm}} \rceil^2$). If

$t(n) = O(n^3)$ (e.g., the Hungarian method), then Strategy 4 has a running time of $O(n^3/m^2 + (mn)^{3/2})$. For $n = 10000, m = 10$, we get a roughly 1000-time speedup.

The second decentralized strategy we look at is an extension to Strategy 4 with three hierarchies; let us call this strategy Strategy 5. After partitioning the bottom (third) hierarchy to m squares, the middle (second) hierarchy is partitioned into \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. This yields following corollary.

Corollary 13 *For Strategy 5 (3-level Hierarchical Rendezvous), in expectation,*

$$D_n \leq C_2 \sqrt{n \log n} + 2 \lceil 1/(2r_{\text{sense}}) \rceil + \lceil 1/(2r_{\text{comm}}) \rceil + 2 \sqrt{n \lceil \frac{\sqrt{2}}{r_{\text{comm}}} \rceil + 4 \lceil \frac{\sqrt{2}}{r_{\text{comm}}} \rceil} + 2. \quad (11)$$

Again, $D_n/D_n^* = O(1)$ as $n \rightarrow \infty$, fixing other parameters. 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.

V. SIMULATION STUDIES

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




Fig. 4. 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 III. 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. 4. The simulation suggests that the bounds on n from Theorem 5 are fairly tight.

B. Performance of Near-Optimal Strategies

Next, we simulate Strategies 3-5 and evaluate D_n and computational time for these strategies over varying values of n and r_{comm} . Since the effect of r_{sense} on optimality is not as important, we assume $r_{\text{sense}} \geq \sqrt{2}$ so that all robots are aware of all target locations. Due to our choice of subdivisions 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.




Fig. 5. 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. 5. 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.




Fig. 6. Distance optimality of Strategy 4 over varying n and r_{comm} .




Fig. 7. Distance optimality of Strategy 5 over varying n and r_{comm} .

For Strategy 4, the ratios were plotted similarly in Fig. 6. 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. 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 optimality. We observe similar results on distance optimality of Strategy 5 (see Fig. 7).

Computational time: We list the computational time, in seconds, for Strategies 3-5 in Table I, for $n = 1000, 5000$, and 10000 . 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. As expected, hierarchical assignment greatly reduces the computational time, often by a factor over 10^3 . The computation was performed on a Intel Core-i7 3970K cpu under a 8GB Java virtual machine.

TABLE I
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)
1000	2.76 s	2.76 s	2.76 s	2.76 s
	0.015 s	0.07 s	0.22 s	0.54 s
	0.002 s	0.003 s	0.003 s	0.006 s
5000	345 s	345 s	345 s	345 s
	0.02 s	0.78 s	2.84 s	8.28 s
	0.069 s	0.032 s	0.043 s	0.058 s
10000	2756 s	2756 s	2756 s	2756 s
	0.83 s	2.32 s	8.35 s	24.4 s
	0.43 s	0.11 s	0.11 s	0.14 s

VI. 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 further provided an explicit formula for computing the number of robots sufficient for probabilistically guaranteeing such an optimal solution. Then, we took a different angle and looked at strategies with good 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. Some of these decentralized strategies also provide computational advantages over a centralized one.

REFERENCES

- [1] P. K. Agarwal, A. Efrat, and M. Sharir. Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications. In *The Eleventh annual Symposium on Computational Geometry*, pages 39–50, 1995.
- [2] M. Ajtai, J. Komlós, and G. Tusnády. On optimal matchings. *Combinatorica*, 4(4):259–264, 1984.
- [3] P. Balister, B. Bollobás, A. Sarkar, and M. Walters. Connectivity of random k -nearest-neighbour graphs. *Advances in Applied Probability*, 37:1–24, 2005.
- [4] D. P. Bertsekas. The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of Operations Research*, 14:105–123, 1988.
- [5] D. P. Bertsekas and D. A. Castañón. Parallel synchronous and asynchronous implementations of the auction algorithm. *Parallel Computing*, 17:707–732, 1991.
- [6] 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>.
- [7] R. Burkard, M. Dell’Amico, and S. Martello. *Assignment Problems*. Siam Society for Industrial and Applied Mathematics, Philadelphia, 2012.
- [8] H. Choset. Coverage of known spaces: The boustrophedon cellular decomposition. *Autonomous Robots*, 9:247–253, 2000.
- [9] 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.

- [10] J. Cortés, S. Martínez, 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.
- [11] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM*, 19(2):248–264, 1972.
- [12] P. Erdős and A. Rényi. On a classical problem of probability theory. *Publ. Math. Inst. Hung. Acad. Sci., Ser. A* 6:215–220, 1961.
- [13] 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.
- [14] A. Ganesh and F. Xue. On the connectivity and diameter of small-world networks. *Advances in Applied Probability*, 39:853–863, 2007.
- [15] 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.
- [16] 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.
- [17] S. Kloder and S. Hutchinson. Path planning for permutation-invariant multirobot formations. *IEEE Transactions on Robotics*, 22(4):650–665, 2006.
- [18] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] M. Penrose. The longest edge of the random minimal spanning tree. *Annals of Applied Probability*, 7:340–361, 1997.
- [23] M. Penrose. *Random Geometric Graphs*. Oxford Studies in Probability. Oxford University Press, London, 2003.
- [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] M. Talagrand. The ajtai-komlós-tusnády matching theorem for general measures. In *Probability in Banach spaces*, 30, 1992.
- [27] 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.
- [28] 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*, 59(9):2261–2276, 2013.
- [29] 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.
- [30] P. M. Vaidya. Geometry helps in matchings. *SIAM Journal on Computing*, 18(6):1201–1225, 1989.
- [31] F. Xue and P. R. Kumar. The number of neighbors needed for connectivity of wireless networks. *Wireless Networks*, 10(2):169–181, March 2004.
- [32] J. Yu, S.-J. Chung, and P. G. Voulgaris. Target assignment in robotic networks: Distance optimality guarantees and hierarchical strategies. *IEEE Transactions on Automatic Control*, 2013. conditionally accepted, draft available at <http://arxiv.org/abs/1308.3956>.
- [33] 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.
- [34] 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.
- [35] 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.