

Exploration via Structured Triangulation by a Multi-Robot System with Bearing-Only Low-Resolution Sensors

SeoungKyou Lee, Aaron Becker, Sándor P. Fekete, Alexander Krölller, and James McLurkin

Abstract—This paper presents a distributed approach for exploring and triangulating an unknown region using a multi-robot system. The objective is to produce a covering of an unknown workspace by a fixed number of robots such that the covered region is maximized, solving the *Maximum Area Triangulation Problem* (MATP). The resulting triangulation is a *physical data structure* that is a compact representation of the workspace; it contains distributed knowledge of each triangle, adjacent triangles, and the dual graph of the workspace. Algorithms can store information in this physical data structure, such as a routing table for robot navigation

Our algorithm builds a triangulation in a closed environment, starting from a single location. It provides coverage with a breadth-first search pattern and completeness guarantees. We show the computational and communication requirements to build and maintain the triangulation and its dual graph are small. Finally, we present a physical navigation algorithm that uses the dual graph, and show that the resulting path lengths are within a constant factor of the shortest-path Euclidean distance. We validate our theoretical results with experiments on triangulating a region with a system of low-cost robots. Analysis of the resulting quality of the triangulation shows that most of the triangles are of high quality, and cover a large area. Implementation of the triangulation, dual graph, and navigation all use communication messages of fixed size, and are a practical solution for large populations of low-cost robots.

I. INTRODUCTION AND RELATED WORK

Many practical applications of multi-robot systems, such as search-and-rescue, exploration, mapping and surveillance require robots to disperse across a large geographic area. Large populations of robots offer two large advantages: they can search the environment rapidly using a breadth-first approach, and can maintain coverage of the environment after the dispersion is complete.

In this paper, we demonstrate that triangulating the workspace with a multi-robot system is a useful approach to dispersion and monitoring. Triangulations are used in a large variety of applications because of their useful properties. In our application they provide complete coverage, they can be built with only basic local geometry, and they allow proofs of properties for coverage, navigation, and distributed data storage. The underlying topological structure of a triangulation allows us to exploit its dual graph for mapping and routing, with performance guarantees for these purposes. Fig. 1 shows an example output demonstrating a triangulated network, its dual graph, and a navigating robot.

S. Lee, A. Becker, and J. McLurkin are with the Computer Science Department, Rice University, Houston, TX, 77005 USA e-mail: sl28@rice.edu.

A. Krölller and S. Fekete are with the Computer Science Department, TU Braunschweig, Braunschweig, Germany.




Fig. 1: A sample triangulation and navigation experiment result with 17 r-one robots. The planar network of triangulation is a subset of the full network (dashed gray lines). Each robot creates a new triangle (dark green) by expanding toward the frontier or discovers new triangles (light green triangles) by examining local network geometry. Small black arcs indicate which robot creates and stores which triangle. This is a distributed *physical data structure*; there is no centralized storage of triangulation information. The network between adjacent triangles forms a *dual graph* of the triangulation. In this example, a navigating robot uses a tree rooted at the red triangle to guide it from its start location to its current location. It follows the red path, the numbers indicate hops in the dual graph from the goal triangle.

We are interested in solutions for large populations of robots, and focus our attention on approaches applicable on small, low-cost robots with limited sensors and capabilities. In this work, we assume that robots do not have a map of the environment, nor the ability to localize itself relative to the environment geometry, *i.e.* SLAM-style mapping is beyond the capabilities of our platform. We exclude solutions that use centralized control, as the communication and processing constraints do not allow these approaches to scale to large populations. We also do not assume that GPS localization or external communication infrastructure is available, which are limitations present in an unknown indoor environment. Finally, we assume that the communication range is much smaller than the size of the environment, so a multi-hop network is required for communication, and the *local network geometry* provides each robot with geometric information about its neighboring robots.

The basic problem requires exploring an unknown region by triangulation from a given starting position. The maximum edge length of a triangle is bounded by the communications range of the robots. If the number of available robots is not bounded a priori, the problem of minimizing their number for covering all of the region is known as

the *Minimum Relay Triangulation Problem* (MRTP); if their number is fixed, the objective is to maximize the covered area, which is known as the *Maximum Area Triangulation Problem* (MATP). Both problems have been studied both for the *offline* scenario, in which the region is fully known, and the *online* scenario, where the region is not known in advance [1]. Online MRTP admits a 3-competitive strategy, while the online MATP does not allow a bounded competitive factor: If the region consists of many narrow corridors, we may run out of robots exploring them, and thereby miss a large room that could permit large triangles. In this work, we focus on the online MATP problem. Our algorithm proceeds by extending the covered region by adding new triangles to the frontier of the exploration. The motion controllers we present use local geometric information. In particular, we focus on a simple platform that can only measure angles between neighbors and detect nearby obstacles. We provide a number of results:

- We develop simple and efficient exploration methods based on triangulation.
- We show that these methods only require local information and geometry.
- We demonstrate that well-known abstract concepts (such as the dual graph) can be implemented in a distributed network of robots.
- We provide provable performance guarantees for online triangulation and routing.
- We demonstrate the practicality of our method by implementing it with simple, low-cost robots. (See our video [2] for an overview of [1].)

Related Work

Our work combines ideas of self-organization and routing in stationary sensor networks with approaches to dynamic robot swarms. For the former, e.g., see [3], [4]. The new challenges arise from considering a large number of mobile nodes with limited capabilities, and real-life platforms and constraints. Classical triangulation problems seek a triangulation of all vertices of a polygon, but allow arbitrary length of the edges in the triangulation [1]. This differs from our problem, in which edge lengths are bounded by communication length. Triangulations with shape constraints for the triangles and the use of Steiner points are considered in mesh generation, see for example the survey by Bern and Eppstein [5].

The problem of placing a minimum number of relays with limited communication range in order to achieve a connected network (a generalization of the classical Steiner tree problem) has been considered by Efrat et al. [6], who gave a number of approximation results for the offline problem (a 3.11-approximation for the one-tier version and a PTAS for the two-tier version of this problem); see the survey [7] for related problems. A similar question was considered by Bredin et al. [8], who asked for the minimum number of relays to be placed to assure a k -connected network. They presented approximation results for the offline problem. For swarms, Hsiang et al. [9] consider the problem of dispersing




Fig. 2: (a) Robot u can measure the bearing to neighbor u_L , $B_u(u_L)$, and the orientation of neighbor u_L , $Ori(u_L)$. (b), triangle angles (black arrows) are measured from neighbors of robot u , and shared with u using a local broadcast message. We define left (right) inner angle for the triangle angle of u 's left (right).

a swarm of simple robots in a cellular environment, minimizing the time until every cell is occupied by a robot. For workspaces with a single entrance door, Hsiang et al. present algorithms with time optimal *makespan* and $\Theta(\log(k+1))$ -competitive algorithms for k doors.

There are many types of exploration in the multi-robot literature. McLurkin and Smith [10] present a breadth-first distribution from a more practical view, using a swarm of 100 robots. Durham et al. [11] present an algorithm for a team of robots to cover entire region using pursuit-evasion problem, using the robot's state to store intermediate results of the search. Spears and Spears describe algorithms for to produce a triangle lattice, but this is not a *triangulation*: there is no knowledge of triangles, the dual graph, or distributed data structures for computation [12].

II. COMPUTATIONAL MODEL

We have a system of n robots. The communication network is an undirected graph $G = (V, E)$. Each robot is modeled as a vertex, $u \in V$, where V is the set of all robots and E is the set of all robot-to-robot communication links. The neighbors of each vertex u are the set of robots within line-of-sight communication range r_{max} of robot u , denoted $N(u) = \{v \in V \mid \{u, v\} \in E\}$. Robot u sits at the origin of its local coordinate system, with the \hat{x} -axis aligned with its current heading. Each robot can measure the angles of the geometry of its local network, as shown in Fig. 2a. Robot u cannot measure distance to its neighbors, but can only measure the *bearing* and *orientation*. We assume that these angular measurements have limited resolution.

Robots share their angle measurements with their neighbors. In this way, robot u can learn of all angles in its 2-hop neighborhood. Fig. 2b shows the relevant *inner angles* of a triangle around u . Each neighbor of u computes these angles from local bearing measurements, then announced them. The communication used by these messages is $O(\max(\delta(u \in V)^2))$, where $\delta(u)$ is the degree of vertex u .

Each robot has contact sensors that detect collisions with the environment. There is an obstacle avoidance behavior that can effectively maneuver the robot away from these collisions. The robots also have a short-range obstacle sensor

that can detect walls closer than ≈ 50 cm. The obstacle sensor does not detect neighboring robots.

Algorithm execution occurs in a series of synchronous rounds, t_r . This greatly simplifies analysis and is straightforward to implement in a physical system [13]. At the end of each round, every robot u broadcasts a message to all of its neighbors. The robots randomly offset their initial transmission to minimize collisions. During the duration of each round, robot u receives a message from each neighbor $v \in N(u)$. Each message contains a set of public variables, including the sending robot's unique ID number $u.id$. The remaining variables will be defined later, but we note that the number of bits needed for each variable is bounded by $\log_2 n$, i.e. the number of bits required to identify each robot. This produces a total message of constant size.

III. MAX-AREA TRIANGULATION ALGORITHM

Fig. 3 illustrates the execution of the Max-Area Triangulation (MAT) algorithm. Initially, two *base robots* mark the *base edge* — such as a door to an unexplored building. The algorithm starts with this base edge and proceeds by constructing a triangulation in a breadth-first manner. The triangulation is extended as robots construct triangles along the current *frontier* of exploration. The frontier is shown as blue lines in Fig. 3, and it delineates the boundary between triangulated space and untriangulated space. All the area between the base edge and the frontier is triangulated. Each mobile robot extends the frontier by moving into unexplored space and forming a triangle with itself and at least two other adjacent robots from the frontier. The algorithm terminates when either all of the workspace has been explored, or the maximum number of robots has been exhausted.

Each robot tries to build a *high-quality* triangle—one that does not have edges that are too short or angles that are too small. Equilateral triangles are ideal, but cannot always be constructed due to errors or environmental constraints.

During algorithm execution, we distinguish the following types of edges in the robot network G : 1) *Frontier edges* (Blue lines in Fig. 3), $\{u, v\} \in E_F$, which belong to only one triangle and have at least one vertex that is not in contact with the wall. 2) *Internal edges*, $\{u, v\} \in E_I$ which belong to two adjacent triangles. 3) *Wall edges*, $\{u, v\} \in E_W$, which also belong to only one triangle, but both vertices of the edge are in contact with a wall. The yellow lines indicate the dual graph, D , which connects adjacent triangles. We will address the detail of the dual graph in Section. III-B.

A. Triangulation

Construction of a new triangle begins with the addition of a new *navigating robot*, u . To build the triangulation in a breadth-first fashion, a *frontier triangle* is selected that is the minimum distance in the dual graph from the base triangle. This triangle will have at least one frontier edge, we select it to be the *goal frontier edge*, $\{l, r\}$. The robot uses the dual graph to navigate to the frontier triangle, these algorithms are described in Secs. III-B and III-C.




Fig. 3: Constructing a triangulation in a BFS manner. The frontier edges are blue and interior edges are green. The edges of the BFS tree (the dual graph) are yellow. The blue tick marks on each robot show the direction of the *frontier normal*. This points into unexplored space, in the direction perpendicular to the frontier edges incident at each robot.

A new triangle can be formed in two ways, *expansion* or *discovery*. Fig. 4a illustrates the construction of a triangle by expansion. When navigating robot u is within the frontier triangle, it switches to the *expanding state*, and moves towards the equilateral point for the new triangle. When u crosses the frontier edge $\{l, r\}$, it creates a new expansion triangle Δulr ($l = l_0$ and $r = r_0$ in Fig. 4a). Once robot u arrives at the equilateral point, it switches to the *expanded state*, and adds Δulr to its list of triangles, becoming its owner. Edge $\{l, r\}$ becomes an internal edge, and robot u broadcasts a message to neighbors l and r , so that they update their right and left frontier neighbors to u . Because the edge $\{l, r\}$ is now internal, it is not used for expansion again, which prevents creating overlapping triangles.

When u enters the expanded state, it needs to discover all of the unexpanded high-quality triangles adjacent to Δulr . Fig. 4b shows an example of triangle discovery. We describe the process for the left frontier neighbor (l), it is analogous for the right. We label the left neighbors $\{l_0, l_1, \dots\}$ where $l_0 \equiv l$. Robot u first considers neighbor l_1 , then proceeds through each neighbor on its left side in counter-clockwise order. For each neighbor l_i , $i \geq 1$, robot u checks for edge $\{l_i, l_{i-1}\} \in E_F$. If this edge exists, then u forms a candidate triangle, $\Delta ul_i l_{i-1}$ (light green in Fig. 4b), and evaluates its quality using definition 3.1. The search terminates if the triangle is not high-quality, or there are no further neighbors to consider. If the candidate triangle is high-quality, robot u becomes its owner, and switches its left frontier neighbor from l_{i-1} to l_i . Robot u then broadcasts a message to l_i to update its right frontier neighbor from l_{i-1} to u .

B. Dual Graph Construction

The dual graph of our a triangulation, D , describes the adjacencies between adjacent triangles. The dual graph can




Fig. 4: (a) An example of *Expansion* triangle: Red arrow is the path of u in *Expanding* state. After arriving at the equilateral point, u switches to *Expanded* state. (b) An example of *Discovery* triangle (light green). u checks θ_F to evaluate the quality of candidate triangle, $\Delta u l_1 l_0$.

be used for realizing global objectives, such as routing. However, one difficulty for a distributed swarm of robots is the absence of a centralized authority that can explicitly keep track of a dual graph, as there are only “primal” vertices, i.e., robots. Our solution is to establish and maintain the dual graph implicitly, by assigning each triangle Δ to a unique robot “owner”, $o(\Delta)$, and then mapping edges between triangles in the dual graph to edges between robots in the primal graph.

We first observe that all owners are connected because all robots in our network, with the exception of the base robots, are owners by construction; every time a new navigating robot is added to the network, it becomes the owner of at least one constructed triangle. A robot can own multiple discovered triangles, and must maintain multiple vertices in the dual graph.

We must ensure that two triangle owners connected by an edge in the dual graph can communicate with each other through the primal graph. This is trivial for two triangles Δ_1 and Δ_2 owned by the same robot, so we must show that for two different triangle owners $o(\Delta_1) \neq o(\Delta_2)$ with a dual graph edge, $\{o(\Delta_1), o(\Delta_2)\}_D \in D$, $\{o(\Delta_1), o(\Delta_2)\}$ is an edge in the primal graph.

Lemma 3.1: Consider edge $\{a, b\} \in E_F$ and Δabo , where o is the triangle owner. Then $o = a$ or $o = b$.

Proof: By contradiction: assume $o \neq a$ and $o \neq b$. Then consider the expanding state for Δoab . Since o is the owner in the expanded state, o must have been the navigation robot in the expanding state. Therefore $\{a, b\}$ was the frontier edge in the expanding state and $\{a, b\}$ is now the internal edge in the expanded state, a contradiction. ■

Theorem 3.2: The owners of two adjacent triangles must also be connected.

Proof: Let Δabc and Δabd be the two adjacent triangles, and $\{a, b\}$ be the edge they share. These two triangles can be formed in the following two ways (in the expanding state): 1) Robot a was the navigation robot. Then a is the owner for both Δabc and Δabd . a is connected to itself. 2) Robot d was the navigation robot. This makes Δabc an existing triangle and $\{a, b\}$ a frontier edge in the expanding state. d is also the owner robot for Δabd in the expanded state. Either a or b is the owner of Δabc by Lemma 3.1, so d , the owner of Δabd , must be connected to the owner of Δabc through either edge $\{a, d\}$ or edge $\{b, d\}$. By symmetry, b is equivalent to a and c is equivalent to d . ■

C. Dual Graph Navigation

We use the dual graph as a navigation guide for robots in our triangulation. If the destination triangle is known, such as a frontier triangle, then a broadcast message can be used to build a BFS tree suitable for navigation [14]. Our previous work shows there is no lower bound on the competitive factor of the stretch of a path in the online MATP problem [1], but this requires narrow corridors of infinitesimal width. In the following, we show that more realistic assumptions do allow constant-factor performance.

Let r_{\max} be the maximum length of a triangulation edge. We also consider a lower bound of r_{\min} on the length of the shortest edge in the triangulation; in particular, we assume that the local construction ensures that any non-boundary edge is long enough to let a robot pass between the two robots marking the vertices of the edge, so $r_{\min} \geq 2\delta$, where δ is the diameter of a robot. (The practical validity of these assumptions for a real-world robot platform will be shown in the experimental Section V.) Finally, angular measurements of neighbor positions let us guarantee a minimum angle of α in all triangles. These constraints give rise to the following:

Definition 3.1: Let \mathcal{T} be a triangulation of a planar region \mathcal{R} , with vertex set V . \mathcal{T} is (ρ, α) -fat, if it satisfies the following properties:

- The ratio r_{\max}/r_{\min} of longest to shortest edge in \mathcal{T} is bounded by some positive ρ .
- All angles in \mathcal{T} have size at least α .

This definition is used to prove properties of triangulations.

1) Covered Area:

Theorem 3.3: Consider a (ρ, α) -fat triangulation of a set V with n vertices, with maximum edge length r_{\max} and minimum edge length r_{\min} . Then the total triangulated area is within $\sqrt{3}\rho^2/2\sin(\alpha)$ of the optimum.

Proof:

Each edge has length at least r_{\min} , and any angle is bounded from below by α . The claim follows by trigonometry. ■

Note that in a practical setting, ρ will be much smaller than the theoretically possible worst case; see Fig. 11b for a real-world evaluation.

2) *Path Stretch:* Now we establish that the dual graph of our triangulations can be exploited for provably good routing. We make use of the following terminology.

Definition 3.2: Consider a triangulation \mathcal{T} of a planar region \mathcal{R} , with vertex set V . Let s, g be points in \mathcal{R} and let $p(s, g)$ be a polygonal path in \mathcal{R} that connects s to g ; let $d_p(s, g)$ be its length. Let Δ_s and Δ_g be the triangles containing s and g , respectively, and let $D(s, g) := \Delta_s, \Delta_1, \dots, \Delta_\ell, \Delta_g$ be a shortest path in the dual graph of \mathcal{T} . Then a \mathcal{T} -greedy path between s and g is a path s, q_1, \dots, q_ℓ, g , such that $q_i \in \Delta_i$, and consecutive vertices of the path are connected by a straight line.

In other words, a \mathcal{T} -greedy path between s and g builds a short connection in the dual graph of the triangulation, and then goes from triangle to triangle along straight segments. Note that we do not make any assumptions whatsoever concerning where we visit each of the triangles.




Fig. 5: (a) A shortest s, g -path (shown in red) in a region covered by a triangulation \mathcal{T} . The resulting \mathcal{T} -greedy path is depicted in yellow; a shortest dual path is indicated by colored triangles. Note that each point q_i may be anywhere in the respective triangle Δ_i . (b) A triangle Δ is intersected by a straight line L . If L passes the triangle not too close to one of the endpoints, the length of the intersection is long. If the line passes the triangle close to one of the endpoints (indicated by the dashed line L'), then the intersection with a circle of radius $r_{min}/2$ must be long.

Lemma 3.4: Consider a (ρ, α) -fat triangle Δ with minimum edge length at least r_{min} ; let Δ be intersected by a straight line L . Then the total length of the intersection of L and Δ is at least $\frac{r_{min}}{2 \sin(\alpha/2)}$, or the length of the intersection of L with the $r_{min}/2$ -disk around one of Δ 's vertices is at least $\frac{r_{min}}{2 \sin(\alpha/2)}$.

Proof: Refer to Fig. 5b. Consider the closest distance between L and one of the vertices of Δ . If this is larger than $2r_{min} \cos(\alpha/2)$, then we see from Pythagoras' theorem that the intersection of L and Δ must have length at least $2r_{min} \sin(\alpha/2)$. Otherwise the distance is at most $2r_{min} \cos(\alpha/2)$, and the intersection of L with the $r_{min}/2$ -disk around the closest vertex of Δ must have length at least $2r_{min} \sin(\alpha/2)$. ■

With this, we can proceed to the proof of the theorem.

Theorem 3.5: Consider a (ρ, α) -fat triangulation \mathcal{T} of a planar region \mathcal{R} , with vertex set V , maximum and minimum edge length r_{max} and r_{min} , respectively. Let s, g be points in \mathcal{R} that are separated by at least one triangle, i.e., the triangles Δ_s, Δ_g in \mathcal{T} that contain s and g do not share a vertex. Let $p(s, g)$ be a shortest polygonal path in \mathcal{R} that connects s with g , and let $d_p(s, g)$ be its length. Let $p_{\mathcal{T}}(s, g)$ be a \mathcal{T} -greedy path between s and g , of length $d_{p_{\mathcal{T}}}(s, g)$. Then $d_{p_{\mathcal{T}}}(s, g) \leq c \cdot d_p(s, g) + 2$, for $c = \lfloor \frac{2\pi}{\alpha} \rfloor \frac{\rho}{\sin(\alpha/2)}$, and $d_{p_{\mathcal{T}}}(s, g) \leq c' \cdot d_p(s, g)$, for $c' = \lfloor \frac{6\pi}{\alpha} \rfloor \frac{\rho}{\sin(\alpha/2)}$.

Proof: Consider $p(s, g)$, triangles Δ_s, Δ_g and the sequence $\Delta_1, \dots, \Delta_{\ell'}$ of ℓ' other triangles intersected by it; by assumption, $\ell' \geq \ell \geq 1$, where ℓ is the number of triangles contained in $p_{\mathcal{T}}(s, g)$. Furthermore, note that the disjointness of Δ_s, Δ_g implies $d_p(s, g) \geq r_{min}$.

We first show that $d_p(s, g) \geq \ell' \lfloor \frac{2\pi}{\alpha} \rfloor \frac{\rho}{2 \sin(\alpha/2)}$. For this purpose, charge the intersection of $p(s, g)$ with Δ_i to Δ_i , if its length is at least $\frac{\rho}{2 \sin(\alpha/2)}$; if it is shorter, we charge the length of the intersection of $p(s, g)$ with the $r_{min}/2$ -disk around one of Δ 's vertices p_j evenly to all of the triangles Δ_i that are incident to p_j . Because the minimum angle in a triangle is bounded from below by α , the preceding lemma implies the lower bound on the length of $d_p(s, g)$.




Fig. 6: Finite state machine for MAT algorithm.

On the other hand, it is straightforward to see that no edge in a \mathcal{T} -greedy s, g -path can be longer than $2r_{max}$. Therefore, $d_{p_{\mathcal{T}}}(s, g) \leq 2(\ell + 2)r_{max}$. Comparing the lower bound on $d_p(s, g)$ and the upper bound on $d_{p_{\mathcal{T}}}(s, g)$ yields the claim $d_{p_{\mathcal{T}}}(s, g) \leq c \cdot d_p(s, g) + 2$ with c as stated. The additive term of 2 results from the s and g possibly being close to the boundaries of Δ_s and Δ_g , respectively; it can be removed by noting that $\ell' \geq \ell \geq 1$ implies $(\ell + 2) \leq 3\ell'$, as indicated by the second comparison and the choice of c' . ■

This provides constant stretch factors even under minimal, purely theoretical and highly pessimistic assumptions. The practical performance in real-world settings (where the greedy paths do not visit worst-case points in the visited triangles) is considerably better, as we demonstrate in Section V.

IV. IMPLEMENTATION

A high-level finite-state machine of our implementation of triangulation construction is shown in Fig. 6. Two robots are initialized in the Frontier-Wall state and placed at the base-edge. All other robots begin behind the base edge in the Navigation state. Table I lists helper functions for all algorithms below.

A. Navigation State

The navigation contains three states; *Nav-Internal*, *Expand-Triangle*, and *Wall-Follow*. A new robot, u , enters the network in the *Nav-Internal* state, and runs algorithm 1 to navigate to a frontier triangle. Line 2 runs an occupancy test function, shown in Fig. 7a, that returns the current triangle, T_c , that contains robot u , and its owner, o . If T_c is a non-frontier triangle, then u moves to an adjacent triangle that is closer to (fewer hops from) the frontier (line 10 to 11). Theorem 3.2 ensures that the owner of T_c is connected to owners of adjacent triangles, so u learns the hops of all adjacent triangles with a 2-hop message similar to the geometry message from Fig. 2b. If T_c is a frontier triangle (line 3) or null (only true if u has just crossed the base edge, line 6), then u will create a new triangle. The variables $u.L$ and $u.R$ are set to the left and right neighbors of the frontier edge (line 4 and 7), and the robot changes its state to *Expand-Triangle* (line 5 and 8).

Once in the *Expand-Triangle* state, u runs algorithm 2. Line 2 computes the left and right inner angles to the

Algorithm 1 NAV-INTERNAL

```

1: while  $u.state = \text{Navigate-Internal}$  do
2:    $\mathcal{T}_c \leftarrow \text{GETCURRENTTRIANGLE}()$ 
3:   if  $\text{ISFRONTIERTRIANGLE}(\mathcal{T}_c)$  then
4:      $(u.L, u.R) \leftarrow \text{GETFRONTIEREDGENBR}(\mathcal{T}_c)$ 
5:      $u.state \leftarrow \text{Expand-Triangle}$ 
6:   else if  $\text{ISONLYBASEEDGE}(N(u))$  then
7:      $(u.L, u.R) \leftarrow \text{GETBASEEDGENBR}()$ 
8:      $u.state \leftarrow \text{Expand-Triangle}$ 
9:   else
10:     $\mathcal{T}_{next} \leftarrow \text{GETMINHOPADJTRI}(\mathcal{T}_c)$ 
11:     $\text{MOVETONEXTTRIANGLE}(\mathcal{T}_{next})$ 
12:   end if
13: end while

```

TABLE I: Table of Helper Functions

$\text{GetCurrentTriangle}()$	Runs occupancy test and returns current triangle, \mathcal{T}_c .
$\text{GetMinHopAdjTri}(\mathcal{T}_c)$	Get \mathcal{T}_c 's min-hop adjacent triangle.
$\text{DiscoverTriangle}(u.L, u.R)$	Runs discovery procedure and gets discovery triangles, \mathcal{T}_D , and list of u 's old and new frontier neighbors.
$\text{IsIsoscelesTriangle}(u.L, u.R)$	Checks if $\theta_L = \theta_R$ in an expand triangle.
$\text{GetFrontierWallNbr}(u.L, u.R)$	Returns $u.L$ or $u.R$ in frontier-wall state.
$\text{BCastFMsg}(u^{old}, u^{new})$	Broadcast new frontier msg to nbrs $\in u^{new}$.
$\text{RecvFMsg}(u^{old}, u^{new})$	Receive new frontier nbrs.
$\text{UpdateFNbr}(u^{old}, u^{new})$	Change frontier nbr from u^{old} to u^{new} .
$\text{BCastDisconnectMsg}(u^{old})$	Broadcast disconnect msg to nbrs $\in u^{old}$.
$\text{RecvDisconnectMsg}()$	Return u_{sender} if u_{sender} disconnects u .
$\text{IsContainFrontierEdge}(T_i)$	Checks if T_i has a frontier edge.
$\text{UpdateTriangleHop}()$	For each triangle u owns, sets its hop to 1 + minimum among all adjacent triangles' hops.
$\text{BCastTriangleHop}(N(o))$	Broadcast all hops of all triangles u owns.

frontier neighbors, θ_L and θ_R . Line 3 then runs the triangle-expansion controller illustrated in Fig. 7b until u is in region 3.

We lack the space here for a complete description of the controller, we sketch its operation here. When robot u enters region 3, if $\theta_L > \theta_R$, u first moves toward $B_u(u.L) + \pi$ until $\theta_R \geq \frac{\pi}{3}$. It then changes its heading toward $B_u(u.R) + \pi$, and moves until it reaches the goal region (region 4). The opposite control happens when $\theta_L < \theta_R$.

Robot u stores the triangle on its list (line 5), runs the Discover Triangle procedure to discover all adjacent triangles as described in section III-A (line 6). The *Frontier angle*, θ_F , provides a simple way to evaluate the quality of candidate triangles; we define a triangle to be high-quality if $\theta_F < k$, with k manually tuned to reduce errors. After adding triangles, u updates its frontier neighbors, those of $u.L$ and $u.R$ (line 7), adds new frontier neighbors in u^{new} , and disconnects frontier neighbors in u^{old} .

If u detects a wall while expanding a triangle (line 12), it changes its state to Wall-Follow (line 13). This controller moves u along the wall until it forms an isosceles triangle. Then u stores the triangle, and broadcasts disconnect message to $u.L$ or $u.R$, and changes its state to Frontier-Wall.

B. Frontier and Frontier-Wall State

When robot u enters the frontier or frontier-wall state, it becomes stationary and runs algorithm 3. In lines 3-7,

Algorithm 2 EXPAND-TRIANGLE

```

1: while  $u.state = \text{Expand-Triangle}$  do
2:    $(\theta_L, \theta_R) \leftarrow \text{GETINNERANGLE}(u.L, u.R)$ 
3:    $\text{TRIANGLEEXPANSIONCONTROLLER}(\theta_L, \theta_R)$ 
4:   if  $\text{ISINGOALREGION}(\theta_L, \theta_R)$  then
5:      $\text{STORETRIANGLESToList}(\Delta u u.L u.R)$ 
6:    $(u^{old}, u^{new}, \mathcal{T}_D) \leftarrow \text{DISCOVERTRIANGLE}(u.L, u.R)$ 
7:    $\text{UPDATEFNBR}(u^{old}, u^{new})$ 
8:    $\text{BCASTFMMSG}(u^{old}, u^{new})$ 
9:    $\text{BCASTDISCONNECTMSG}(u^{old})$ 
10:   $\text{STORETRIANGLESToList}(\mathcal{T}_D)$ 
11:   $u.state \leftarrow \text{Frontier}$ 
12:  else if  $\text{ISWALLDETECTED}()$  then
13:     $u.state \leftarrow \text{Wall-Follow}$ 
14:  end if
15: end while

```




Fig. 7: (a) The occupancy test algorithm determines if a robot is inside a given triangle. If any angle between neighbors of u is greater than π , then u is outside of the triangle. (b) Diagram of triangle expansion controller regions between robots u_L and u_R , each with $\frac{\pi}{8}$ bearing resolution. A robot 1 in region 1 rotates around u_L or u_R , so that it always converges in region 2. A robot in region 2 always moving toward the direction where its two inner angles are getting smaller. By doing so, The controller always guides a robot from lower number region to higher number adjacent region by only considering θ_L and θ_R . All sample trajectories in red lines converge to the goal region.

u labels all of its triangles which include a frontier edge to frontier triangles. These triangles become sources for the frontier message that guides navigating robots to them. (line 8-9). The frontier robots compute the *Frontier angle*, θ_F , between adjacent frontier neighbors, in the direction of the frontier normal, shown in Figs. 4b and 3. This is done in line 11-12. To maintain the simply-connected frontier subnetwork, robot u will need to update its frontier edges when new triangles are added. After a new navigation robot, v , expands and discovers new triangles, lines 14-16 ensure the frontier edges adjacent to robot u are updated when messages from v are received. If robot u receives a disconnect message from v , it has no more incident frontier edges, and transitions to the Internal state in lines 17-19. This is illustrated in Fig. 4b) by robot l_0 . If u and v are both frontier-wall robots, v 's disconnect message to u will cause it to transition to internal state and create a wall edge (line 20-22).

C. Internal State

Eventually, robot u is likely to become an Internal robot. It remains stationary, and relays broadcast messages. Every robot processes broadcast messages by updating the hops

Algorithm 3 FRONTIER/FRONTIER-WALL()

```

1: while  $u.state = \text{Frontier}$  OR  $u.state = \text{Frontier-Wall}$  do
2:   for all  $T_i \in \text{TriangleList}$  do
3:     if ISCONTAINFRONTIEREDGE( $T_i$ ) then
4:       SETFRONTIERTRIANGLE( $T_i$ )
5:     else
6:       CLEARFRONTIERTRIANGLE( $T_i$ )
7:     end if
8:     UPDATETRIANGLEHOP( $T_i, N(o)$ )
9:     BCASTTRIANGLEHOP( $N(o)$ )
10:  end for
11:  COMPUTE/ROTATETONORMALVEC( $B(u.L), B(u.R)$ )
12:   $\theta_F \leftarrow \text{COMPUTEFRONTIERANGLE}(B(u.L), B(u.R))$ 
13:  BROADCASTFRONTIERANGLE( $\theta_F$ )
14:  if RECVFMSG( $u^{old}, u^{new}$ ) then
15:    UpdateFNbr( $u^{old}, u^{new}$ )
16:  end if
17:   $v \leftarrow \text{RECVDISCONNECTMSG}()$ 
18:  if  $v.state = \text{Frontier}$  then
19:     $u.state \leftarrow \text{Internal}$ 
20:  else if  $v.state = \text{Frontier-Wall} \wedge u.state = \text{Frontier-Wall}$  then
21:     $u.state \leftarrow \text{Internal}$ 
22:  end if
23: end while

```

of each triangle they own by considering the hops to adjacent triangles, finding the minimum, and adding one. This procedure propagates the broadcast message, and the hops updated, and ensures that any new robot crossing the base edge will move to the frontier triangle that is nearest in the dual graph, providing a breadth-first construction.

V. EXPERIMENTAL RESULTS

We have performed several real world experiments, using the r-one robots shown in [15]. The capabilities of this platform supports the assumptions in our problem statement.; each robot can measure the bearings to its nearby robots, despite of a limited resolution of only $\frac{\pi}{8}$, and exchange messages including those bearings and necessary information to run an implemented algorithm in Section IV using inter-robot communication. Each robot also has 8 bump sensors that provides wall detection. To evaluate a resulting triangulation quality or trace the trajectory of a navigation robot, we use the April-Tag system by APRIL group [16]. This measures the ground truth position, $P_u = \{x_u, y_u, \theta_u\}$, of each robot u . The u , however, cannot measure or use the ground-truth position while executing our algorithms. All robots only know the two-hop local network geometry shown in Fig. 2b.

A. Maximum Area Triangulation using MATalgorithm

Fig. 8 shows snapshots of triangulation. Over 8 trials using 9-16 robots, the average triangulated area is $1.5 \pm 0.29 m^2$. It takes 7.8 ± 2.1 robots to cover a unit area ($1 m^2$). The resulting triangulations are ($\rho = 3.6, \alpha = 0.36 rad$)-fat. Fig. 9a shows that our triangulations cover about 91% of the region behind the frontier edges. The uncovered region is because the top-left and bottom-left corner in Fig. 8 are wall edges (incident on two wall robots), and are not expanded by navigating robots. Fig. 9b shows the distribution of area covered by individual triangles. The initial length of the base edge predicts the area of an ideal equilateral triangle should be $0.088 m^2$, our triangles have a mean area of $0.13 m^2$,






Fig. 9: (a) Pie chart for covered area by all triangles. (b) Histogram of covered area by each triangle.



(a) Histogram for minimum angle of each triangle.



(b) Histogram for MaxMin ratio of each triangle.

Fig. 10: (a) Distribution of minimum angle of each triangle, not for a global ρ . All triangles not constrained by a wall satisfy the lower bound, $\frac{\pi}{8}$. (b) Distribution of MaxMin ratio of each triangle. All triangles whose minimum angle is larger than the $\frac{\pi}{8}$ (blue colored) also satisfy corresponding upperbound MaxMin ratio.

with a std. dev. of $0.065 m^2$. This discrepancy caused by the angle-based sensors; the robots cannot measure range, and therefore cannot control the area of the triangle they produce. We show this by studying individual triangle quality.

Figs. 10a and 10b show our measurements of individual triangle quality; the distribution of minimum angle and maximum/minimum edge length ratio (MaxMin ratio) for each triangle. The individual data shows triangle quality in a way that overall ρ cannot. An ideal equilateral triangle has a minimum angle of $\frac{\pi}{3}$ rad and MaxMin ratio of 1. Triangles satisfying the lower bound for minimum angle and the upper bound for MaxMin ratio are 95% and 96.7% of overall triangles, respectively. We note that all triangles not constrained by a wall satisfy these bounds, meaning they are approximately the correct shape, but not always the correct size. Knowing range would let us address this, but it is unclear how robots expanding the triangulation should choose between making a triangle of the correct shape, or the correct size. We leave this for future work.




Fig. 8: Screenshots of constructing triangulation with 12 robots.




Fig. 11: (a) Sample navigation results with 12 robots. (b) Stretch factor histogram with 34 trials of navigation experiment.

B. Dual Graph Navigation

We start each navigation experiment with a constructed triangulation. The triangulation is ($\rho = 1.36, \alpha = 0.88rad$)-fat. For each trial, we randomly select one triangle as a goal, and the robots build a tree on the dual graph. Fig. 11a, shows five trials of the 34 we conducted. The numbers inside the triangles indicate the hops in the dual graph from the goal triangle. (The trial from the 8-hop triangle is also shown in Fig. 1) The thick blue lines show connectivity between owners of adjacent triangles. Note that this graph is not complete, but it is a spanning graph of all triangle owners in G , which is implied by Theorem. 3.2.

Fig. 11b shows the distribution of the stretch factor over all trials of navigation tests with various start-goal pairs. The mean stretch factor is 1.38 ± 0.19 . This is much less than the theoretical bound implied by Theorem 3.5, which is based on worst-case assumptions. Our occupancy algorithm produces 91% correctness in returning the triangle that actually includes u . We define navigation correctness as the ratio of times u moves to the correct adjacent triangle. This result is 99%, with incorrect navigation caused by occupancy errors.

VI. CONCLUSION

We have presented a distributed algorithm to triangulate a workspace, produce a physical data structure, and use this structure for communications and robot navigation. There are many exciting new challenges that lie ahead. The next step is to extend this approach with a self-stabilizing algorithm that can construct and repair the triangulation and dual graph dynamically, starting from an arbitrary distribution of robots. While existing controllers can already form triangulated graphs [17], what is needed is construction and maintenance of the physical data structures, i.e., the primal and dual graphs. Future work can extend these ideas to very large populations and dynamic environments. Another objectives will be to improve the routing algorithm to replace the simple dual graph paths by more sophisticated geodesic trajectories.

We are currently working on multi-robot patrolling using the physical data structure to store visitation frequencies and implement a geodesic Lloyds controller to provide periodic coverage of the triangulation with multi-patrolling robots.

REFERENCES

- [1] S. P. Fekete, T. Kamphans, A. Krölller, J. Mitchell, and C. Schmidt, "Exploring and triangulating a region by a swarm of robots," in *Proc. APPROX 2011*, ser. LNCS, vol. 6845. Springer, 2011, pp. 206–217.
- [2] A. Becker, S. P. Fekete, A. Krölller, S. K. Lee, J. McLurkin, and C. Schmidt, "Triangulating unknown environments using robot swarms," in *Proc. 29th Annu. ACM Sympos. Comput. Geom.*, 2013, pp. 345–346, video available at <http://imaginary.org/film/triangulating-unknown-environments-using-robot-swarms>.
- [3] S. P. Fekete and A. Krölller, "Geometry-based reasoning for a large sensor network," in *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, 2006, pp. 475–476.
- [4] —, "Topology and routing in sensor networks," in *Proc. 3rd Intern. Workshop Algorithmic Aspects Wireless Sensor Networks (ALGOSENSORS)*, 2007, pp. 6–15.
- [5] M. Bern and D. Eppstein, "Mesh generation and optimal triangulation," in *Computing in Euclidean Geometry*, ser. Lecture Notes Series on Computing, D.-Z. Du and F. K. Hwang, Eds. Singapore: World Scientific, 1992, vol. 1, pp. 23–90.
- [6] A. Efrat, S. P. Fekete, P. R. Gaddehosur, J. S. Mitchell, V. Polishchuk, and J. Suomela, "Improved approximation algorithms for relay placement," in *Proc. 16th Annu. Europ. Sympos. Algor.* Springer, 2008, pp. 356–367.
- [7] B. Degener, S. P. Fekete, B. Kempkes, and F. M. auf der Heide, "A survey on relay placement with runtime and approximation guarantees," *Computer Science Review*, vol. 5, no. 1, pp. 57–68, 2011.
- [8] J. Bredin, E. Demaine, M. Hajiaghayi, and D. Rus, "Deploying sensor networks with guaranteed fault tolerance," *Networking, IEEE/ACM Transactions on*, vol. 18, no. 1, pp. 216–228, feb. 2010.
- [9] T.-R. Hsiang, E. M. Arkin, M. A. Bender, S. P. Fekete, and J. S. B. Mitchell, "Algorithms for rapidly dispersing robot swarms in unknown environments," in *Algorithmic Foundations of Robotics V*, ser. STAR, J. B. et al., Ed., vol. 7. Springer, 2004, pp. 77–93.
- [10] J. McLurkin and J. Smith, "Distributed algorithms for dispersion in indoor environments using a swarm of autonomous mobile robots," in *Proc. 7th Internat. Sympos. Distr. Auton. Robot. Syst.*, 2004.
- [11] J. W. Durham, A. Franchi, and F. Bullo, "Distributed pursuit-evasion with limited-visibility sensors via frontier-based exploration," in *IEEE International Conference on Robotics and Automation*, 2010, pp. 3562–3568.
- [12] W. M. Spears, D. F. Spears, R. Heil, W. Kerr, and S. Hettiarachchi, "An overview of physicomimetics," *Lecture Notes in Computer Science-State of the Art Series*, vol. 3342, 2005.
- [13] J. McLurkin, "Analysis and implementation of distributed algorithms for Multi-Robot systems," Ph.D. thesis, Massachusetts Institute of Technology, 2008.
- [14] Q. Li and D. Rus, "Navigation protocols in sensor networks," *ACM Trans. Sen. Netw.*, vol. 1, no. 1, pp. 3–35, 2005.
- [15] J. McLurkin, A. J. Lynch, S. Rixner, T. W. Barr, A. Chou, K. Foster, and S. Bilstein, "A low-cost multi-robot system for research, teaching, and outreach," in *Distributed Autonomous Robotic Systems*, vol. 83. Springer Tracts in Advanced Robotics, 2013, pp. 597–609.
- [16] E. Olson, "Apriltag: A robust and flexible visual fiducial system," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 3400–3407.
- [17] W. M. Spears, D. F. Spears, J. C. Hamann, and R. Heil, "Distributed, Physics-Based control of swarms of vehicles," *Autonomous Robots*, vol. 17, no. 2, pp. 137–162, 2004.