

Stigmergy-based, Dual-Layer Coverage of Unknown Regions

Ori Rappel

Technion - Israel Institute of
Technology, Haifa, Israel
rappelori@campus.technion.ac.il

Michael Amir

Technion - Israel Institute of
Technology, Haifa, Israel
ammicha3@cs.technion.ac.il

Alfred M. Bruckstein

Technion - Israel Institute of
Technology, Haifa, Israel
freddy@cs.technion.ac.il

ABSTRACT

We present algorithms for uniformly covering an unknown indoor region with a swarm of simple, anonymous and autonomous mobile agents. The exploration of such regions is made difficult by the lack of a common global reference frame, severe degradation of radio-frequency communication, and ground obstacles. We propose addressing these challenges by using airborne agents, such as Micro Air Vehicles, in dual capacity, both as mobile explorers and, once they land, as beacons that help other agents navigate the region.

The algorithms we propose are designed for a swarm of identical ant-like agents with local sensing capabilities. The agents enter the region, which is represented as a graph, over time from one or more entry points and are required to occupy all of its vertices. Unlike many works in this area, we consider the task of informing an outside operator with limited information that the coverage mission is complete. Even with this additional requirement we show, both through simulations and mathematical proofs, that the dual role concept results in linear-time termination, while also improving many well-known algorithms in the literature in terms of energy use.

KEYWORDS

swarm robotics, mapping and exploration, uniform dispersion, unknown environment, dual-layer coverage, ant-like robots

ACM Reference Format:

Ori Rappel, Michael Amir, and Alfred M. Bruckstein. 2023. Stigmergy-based, Dual-Layer Coverage of Unknown Regions. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

Rescue workers coming to a rescue scene have neither an up-to-date map of the region nor a way of monitoring hazards such as harmful chemical leakages or structural instability. For the workers to safely and efficiently navigate the scene, it is desirable to continuously monitor it before and during the workers' operations. Several works have proposed addressing this problem by flooding the region with a swarm of simple, airborne agents prior to workers' entrance [7, 16, 17, 19]. The robots' purpose is to explore and monitor the scene in a distributed fashion, providing rescue workers with real time, potentially life-saving information.

Swarm-robotic algorithms whose purpose is to flood an unknown region with robots are broadly referred to as "uniform dispersion" algorithms in the literature [19]. Uniform dispersion algorithms are designed with an indoor, unknown rescue scene in

mind. Such a setting poses four unique challenges for multi-robot systems: (i) robots' movements may be hampered by rubble on the ground or other obstacles, (ii) lack of a common reference frame makes coordination based on common landmarks difficult or even impossible, (iii) robots have limited energy stores and can only operate at full capacity for a limited time, and (iv) in difficult indoor rescue regions, severe multi-path and signal attenuation errors limit radio frequency communication between the robots. Although numerous works have been written about each of these challenges as separate topics (see below the "Related work" section), we are not aware of any prior work about uniform dispersion that attempts to address all four challenges.

In this work, we propose overcoming these challenges by using Micro Air Vehicles (e.g., small quadcopters), which we will refer to as agents. We envision agents small enough to navigate difficult indoor environments (such as collapsed buildings or underground tunnels) and to settle inside. The agents enter the region one after another via predefined entry points and explore it in a distributed manner according to a local sensing-based algorithm. Agents gradually create *beacons* throughout the region. These beacons help agents navigate through local visual signals, and enable real time monitoring and pathfinding inside the region.

This kind of airborne swarm approach is limited by several factors: first, using airborne agents comes at the cost of limited flight time due to their high power consumption and limited battery capacity. Next, although rescue scenes such as collapsed buildings benefit small agents, their limited payload capability severely limits their ability to carry and dispense the aforementioned beacons. To overcome both energy and payload limitations, rather than have beacons be something that the agents dispense throughout the region, our idea is to let the agents settle down as soon as possible and act as beacons themselves, so as both to conserve energy and guide the remaining flying agents. This idea draws upon the bio-inspired concept of *stigmergy*—communication via the environment [20, 29, 31]. In swarm robotics, stigmergy is usually implemented through agents *deploying* beacons, "virtual pheromones", or other kinds of marks inside the environment. In contrast, our implementation of stigmergy utilizes *agents themselves* as a medium to change the state of the environment by having agents *become* beacons.

The algorithms we propose successfully cover regions that can be represented as a graph $\mathbb{G}(V, E)$, where vertices represent locations and edges between them represent the ability of an agent to move from one location to the next. For the sake of concreteness, we model the rescue scene as a discrete 2D grid graph composed of equally sized cells (Figure 3). We assume that each cell can contain at most two agents: one agent on the ground serving as a beacon, and one agent flying above-ground. Once an agent enters the region it moves autonomously according to a local sensing-based algorithm. As is common in swarm robotics, we assume the agents to be

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

anonymous and identical, hence executing the same algorithm. The algorithm results in an environment uniformly covered with beacons, such that each cell of the grid graph contains a beacon. We assume agents are capable of reading local visual signals from nearby beacons. Visual signals (i.e., lights) can encode information using color and blinking frequency, as shown in e.g. [21]. Since the agents’ direct (e.g., radio) communication capabilities are severely degraded, they rely on these signals to navigate the environment.

Results. We present several stigmergy-based “gradient ascent” algorithms for uniformly filling the region with beacons and study their time to completion and energy use via simulations and formal analysis. The algorithms work by maintaining and growing a sub-region of beacons that generate a virtual “gradient” which flying agents ascend to efficiently find new locations to explore. We prove that the algorithms’ time to completion is linear in the size of the environment in both a synchronous and a bounded asynchronous time setting (in the sense of [12]). The best-performing algorithm that we study is the “Single-Layer Unlimited Gradient” (SLUG) algorithm. We compare SLUG to recent works [3] as well as to the well-known, non-stigmergy based algorithms of Hsiang et al. [19] and find that SLUG’s energy use is significantly lower, suggesting that dual layer-based approaches to coverage have significant energy savings over traditional approaches. Desirably, we also show that SLUG can recover from several kinds of crashes and errors.

An often overlooked topic in swarm robotics which is especially relevant in non-synchronous time settings (where we cannot infer much from the passage of time) is the ability of an operator to *detect* mission completion. In this work, we take the perspective that an external observer’s (e.g., a rescue worker’s) ability to tell whether the agents have completed the coverage mission is just as important as the agents’ ability to do so. Hence, in our view, the agents’ labor is not complete until they have managed to somehow deliver information about the job’s completion to such an observer. The primary metric we study and minimize is thus termination time – the earliest point in time in which robots can inform an external observer monitoring only the region’s entry points that the region is completely covered with beacons. We show that the observer can quickly and correctly be informed of mission termination through back-propagation of the beacons’ states.

Related work. Many works have been written on the topic of multi-agent deployment under constraints [1, 10, 24, 25, 30]. The problem of completely filling an environment with robots in a decentralized fashion is often called “uniform dispersion.” The version of uniform dispersion we consider in this work can be traced back to Hsiang et al. [19] or to Howard et al. [17, 18]. The literature on uniform dispersion and related problems can be organized according to assumptions regarding the capabilities of the agents, the time scheme used (synchronous or asynchronous), the initial location of the agents, and the properties of the region. As an example, in [19], agents are assumed capable of direct (radio-based) inter-agent communication, time is synchronous, agents are initially located outside the region, and the region is grid-like. Conversely, in some works the initial location of the agents may be inside the region [9, 17]. The region to be covered may be a continuous part of the two-dimensional plane [17, 22, 23, 29] or modelled as a discrete graph (usually as a grid graph) [7–9, 15, 16, 19].

In this work we broadly follow Hsiang et al. [19]’s setting and assume that our agents are initially located outside the region, gradually entering it via predefined *entry points*. However, whereas Hsiang et al.’s algorithms require synchronous time and direct inter-agent messaging, our algorithms only require indirect visual communication and can operate in a bounded asynchronous time scheme. Additionally, unlike [19] and most works in this area, we focus on external detection of coverage completion rather than just achieving coverage. Hsiang et al.’s problem setting and model are extended in numerous works including [2, 3, 7–9, 15, 16]. The closest work to ours among these is [3], in which the concept of covering an environment with two layers of agents is similarly explored. The coverage algorithm presented in [3] (“DLTT”) is used as a benchmark in our simulation results. We show that SLUG beats it significantly in terms of termination time and energy use.

Stigmergy-based graph exploration strategies in which agents place beacons in the environment to aid the exploration process are widely studied [11, 14]. In contrast to the majority of these works, our robots do not deploy beacons in the environment, but *become* them. We show that using *agents themselves* as beacons results in highly attractive properties, including significant energy savings and ant-like recovery from crashes and errors. The idea of using settled airborne agents as beacons was first studied by Stirling et al. [28, 29] and later by Aznar et al. [4, 5] with the objective of exploring (not coverage) and in [3, 26] with the objective of coverage. In [26] the authors present a uniform coverage algorithm based on stigmergy and depth-first search, however the number of agents in each cell is not limited and agents possess unique IDs. The assumption of unique IDs is made in [5, 6, 28, 29] as well. Conversely, in this work and in [3] (whose DLTT algorithm we improve upon, as per the previous paragraph) it is assumed all robots in the swarm are anonymous and identical.

Energy is the primary constraint on the performance of swarms composed of flying agents, and energy consumption is often measured as a performance metric [2, 5, 26, 29]. In [2], a time and energy-optimal uniform dispersion algorithm for simply connected environments is presented alongside theoretical lower bounds on energy use in general environments. Aznar et al. [5] propose predicting energy consumption via recurrence equations. We study our algorithms’ energy use in Section 4, measuring both the maximum individual energy use of an agent and the cumulative energy use of the swarm. Our thesis, which is corroborated by simulation results (Section 4), is that a stigmergy-based approach where agents can land and become beacons leads to significant energy savings.

2 MODEL AND PROBLEM FORMULATION

In this section we define the *beacon coverage problem*, a type of uniform dispersion task wherein a swarm of agents seeks to completely cover a discretized, a priori unknown environment \mathbf{R} with beacons and to detect termination, such that in the final configuration every location in \mathbf{R} contains a beacon and an external observer monitoring only the state of the agents at the region’s entry points is capable of detecting that the coverage mission is complete.

We model \mathbf{R} as a graph $\mathbb{G}(V, E)$ whose vertices, V , represent locations in \mathbf{R} and whose edges, E , represent the connections between locations. We define $n = |V|$ and $m = |E|$. For simplicity, we assume here that \mathbb{G} is an arbitrary grid graph composed of equally

sized cells. However, this assumption is not required by our algorithms and formal analysis, both of which can readily be extended to arbitrary graph environments.

We assume a sufficiently large swarm of identical and anonymous robots a_1, a_2, \dots is ready to enter the region one by one via predefined entry points located somewhere in \mathbf{R} . a_i is defined to be the i th robot that enters \mathbf{R} .

2.1 The Agents

At all times, an agent is either in flight and searching for a place to settle, or has already settled somewhere in \mathbf{R} . The former type of agent is called *mobile* whereas the latter is called a *beacon*. Beacons have two sub-states: they can be *open* or *closed*. Beacons remain in place forever, conserving energy, and their function is to provide simple information to nearby mobile agents via visual signals.

Agents are assumed to be identical (they all act according to the same algorithm), anonymous (have no identifiers) and autonomous. Agents wake up once per time step and perform a Look, Compute and Move action cycle. An agent can sense other agents' presence at any vertex neighboring its location. As we assume \mathbf{R} is a grid graph, this means agents sense adjacent cells at Manhattan distance 1 from themselves. Mobile agents are allowed to move to an adjacent cell as long as that cell contains no other mobile agent. Mobile agents can settle and become beacons at any adjacent cell that contains no other beacons. Hence, each cell of \mathbf{R} may contain at most one mobile agent and one beacon, reflecting the idea that beacons have landed on the ground and mobile agents can fly above them.

Definition 2.1. The set of agents and cells sensed by agent a_i at time t , denoted $\xi(i, t)$, is called "the neighborhood of a_i at time t ."

All agents are capable of detecting the state of agents in their neighborhood (mobile, open beacon, closed beacon). Each agent i further keeps track of an integer called a step count s_i . Beacons are capable of projecting their step counts to neighboring agents.

2.2 Time Model

The algorithms we describe work in many kinds of synchronous, semi-synchronous and asynchronous time settings. For the purposes of formal modelling and time bound analysis, in this paper we focus on a type of fair 2-bounded asynchronous scheduler (in the sense of [12]). Time, t , is divided into integer time steps $t = 0, 1, 2, 3, \dots$, such that each time step is further sub-divided into M equal sub-time steps of duration $dt = 1/M$. Let $U(0, M-1)$ denote the uniform distribution over the integers $\{1, 2, \dots, M-1\}$, and let $t_{i,k}$ denote the k th time that agent a_i wakes up. We define:

$$t_{i,k} = k + \frac{U(0, M-1)}{M} \quad (1)$$

In other words, at every time step, each agent wakes up during a randomly chosen sub-time step (see Figure 1). An agent that wakes up at time t senses the configuration of agents at the previous sub time step $t - dt$. If two agents attempt to move to the same location at the same sub time-step, we assume one of them stays put arbitrarily. This kind of time scheme is applicable to scenarios in which agents have an internal timer that wakes them up at some fixed nominal rate, but are not controlled by a centralized scheduler that guarantees precise synchronization between the agents' clocks.

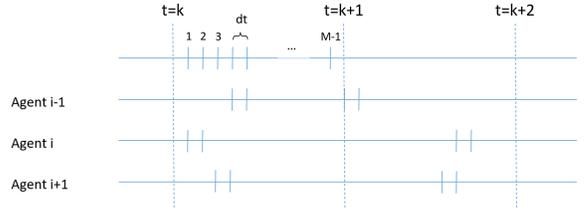


Figure 1: An illustration of the time scheme. Large vertical bars indicate the beginning of time steps; the short bars indicate sub-time steps. The randomly chosen wake-up times of agents $i-1$, i and $i+1$ are depicted.

For the rest of this work, unless stated otherwise, the phrase "at time step t " refers to the time interval $[t, t+1)$, and the phrase "at time t " refers to the configuration of the agents at sub-time step t , after any agent wake-ups.

As a special case, our time bound result (Theorem 3.1) holds also for synchronous time schemes where all agents always wake up at the same sub-time step.

2.3 Agent Entry Model

Agents are initially outside the region and enter the region one by one via predefined *entry points*. We assume a constant rate entry model in which the time between successive entry attempts of agents to the region through an entry point is bounded by an integer parameter $\Delta T \geq 1$. Specifically, let p be an entry point in \mathbf{R} . For every integer $k \geq 0$, an agent will enter p at the earliest sub-time step $t_0 \in [k\Delta T, (k+1)\Delta T)$ at which p contains no mobile agent. Since t_0 is unique by definition (there is just one "earliest" time $t \in [k\Delta T, (k+1)\Delta T)$ where p is empty), at most k agents can enter the region before time $k\Delta T$. An agent entering the environment at a given time step remains there without waking for the rest of the current time step and activates only at the subsequent time step.

Although many works in the literature of robotic coverage assume similar entry models, the effect of the entry process on mission completion time is rarely explored in depth. The interaction between the entry delay ΔT and the coverage algorithm is especially important when trying to develop a tight upper bound on the time to mission completion. We study this topic in Sections 3.1 (analytically) and 4 (empirically).

2.4 The Agents' Mission

It is common in works on multi-robot coverage to place primary importance on the amount of time a coverage algorithm takes to complete. One limitation of this approach is that an external observer cannot necessarily tell that coverage has been completed, and might not be able to proceed to the next stage of their operation without such information. In this work, we take the perspective that an external observer's ability to tell whether the agents have arrived at the desired configuration is just as important as the agents' ability to do so, and so the agents' mission is not complete until they have managed to somehow deliver this information to such an observer. In light of this, our agents' mission is both (a) to completely cover \mathbf{R} with beacons and (b) to deliver this information back to external observers. Since \mathbf{R} is an a priori unknown region, we assume external observers can only see and read the states of

the agents at each of \mathbf{R} 's entry points. We then follow the convention that, to announce mission completion, each entry point must contain a closed beacon. Formally:

Definition 2.2. The **termination time**, $T_C(\mathbf{R})$, of the multi-agent system with respect to an environment \mathbf{R} is the first time t at which every entry point contains a closed beacon. The algorithm is said to have executed correctly if every cell contains an (open or closed) beacon at termination time. The **termination time step** is defined as the first time step after termination occurs, i.e., $\lceil T_C(\mathbf{R}) \rceil$.

We are also interested in the *total energy use* of a multi-agent system as a measure of its efficiency. We assume here that the power consumption of a mobile agent is constant and that once an agent becomes a beacon its energy consumption is 0 (since it no longer needs to fly). Hence energy consumption is linear with respect to the time an agent is mobile, and the total energy used is the sum of the flight times of all the agents:

Definition 2.3. The energy consumption of an agent a_i , denoted E_i , is the duration of time a_i spends as a mobile agent after entering \mathbf{R} up to time $T_C(\mathbf{R})$.

Definition 2.4. Assuming k agents entered \mathbf{R} before time $T_C(\mathbf{R})$, the swarm's total energy consumption is equal to $\sum_{i=0}^k E_i$. Total energy consumption is denoted by $E_{total}(\mathbf{R})$.

We may also consider the *maximal energy use* of an agent in the system. Whereas total energy measures the swarm's cost-effectiveness (e.g., how much fuel or electricity it consumes), maximum energy tells us what energy storage capabilities our robots require to complete the mission:

Definition 2.5. The maximum energy consumption of a multi-agent system with respect to an environment \mathbf{R} is equal to $\max E_i$, the maximum energy used among all agents at termination time.

We note that even after the termination time, redundant mobile agents may remain in \mathbf{R} and continue consuming energy. Although post-termination energy use is not one of the main metrics we explore in this paper, we will later show that beacons can be used to send redundant agents back to entry points where they can be recovered and stop consuming energy.

3 ALGORITHMS

In this section we describe two beacon coverage algorithms: "Dual-Layer Limited Gradient," abbreviated *DLLG*, and "Single-Layer Unlimited Gradient," abbreviated *SLUG*. *DLLG* is a relatively straightforward algorithm that completely fills the region with two layers of agents: beacons and mobile agents. We study its properties and prove that, when there is one entry point, its termination time step obeys the inequality $(2n - 1)\Delta T + 1 \leq \lceil T_C(\mathbf{R}) \rceil \leq (2n - 1)\Delta T + 2$ for $\Delta T \geq 2$. We then discuss *SLUG*, which is a strict improvement over *DLLG* and uses *DLLG* as a building block. *SLUG* improves *DLLG* in two ways: by making better use of closed beacons, and by enabling greater freedom of movement for the robots. These improvements can be implemented separately. In Section 4, we study (numerically) the contribution of each type of improvement to the algorithm's performance in terms of termination time and energy use.

In addition to its performance improvements over *DLLG*, *SLUG* can recover from various errors and crash faults - see Section 3.2.

3.1 Dual-Layer Limited Gradient

The idea of *DLLG* is to have mobile agents read the step counts of neighboring beacons and treat them as a gradient which they must ascend. Specifically, at every time step, mobile agents with step count s_i attempt to move to neighboring locations that contain a beacon with step count $s_i + 1$. Hence, the step count of every mobile agent necessarily equals the step count of the beacon it flies over, and agents are always "climbing" the gradient defined by the beacons in single-step increments. Mobile agents themselves become beacons over time: as a first priority, whenever a mobile agent sees a neighboring empty cell, it attempts to settle in that cell, changing its state to "open beacon" and thereafter signaling its step count to any agent that sees it. Beacons become "closed" in a backpropagating fashion until each entry point contains a closed beacon. Specifically, open beacons become "closed" when they contain a mobile agent, all their neighboring cells contain beacons, and every neighboring beacon with higher step count (if any such neighbor exists) is a closed beacon. This condition guarantees that the algorithm terminates when precisely $2n$ agents are present in the region: n mobile agents and n closed beacons.

DLLG is a distributed algorithm that is implemented via a local action rule that all agents in \mathbf{R} execute at every time step. The local rule for agent a_i at time t is defined by Algorithm 1. Lines 1-9 of Algorithm 1 govern mobile agents' behavior and we shall refer to them as the *mobile agent action rules*. Lines 10-11 govern the behavior of beacons and shall be referred to as *beacon action rules*. In *DLLG*, beacon action rules are only necessary for detecting termination; the mobile agent action rules are entirely sufficient for the coverage process itself.

Algorithm 1 Local rule for agent a_i waking up at time $t + dt$ and executing Dual-Layer Limited Gradient.

Require: a_i has entered \mathbf{R}

```

1: if  $a_i$  is a mobile agent then
2:   if  $\xi(a_i, t)$  contains an empty cell  $u$  then
3:      $a_i$  moves to  $u$ 
4:      $a_i$  sets state to "open beacon"
5:      $s_i \leftarrow s_i + 1$ 
6:   else if  $\xi(a_i, t)$  contains a cell  $u$  with exactly one agent  $a_j$ 
   and  $a_j$  is an open or closed beacon and  $s_j = s_i + 1$  then
7:      $a_i$  moves to  $u$ 
8:      $s_i \leftarrow s_i + 1$ 
9:   end if
10: else if  $a_i$  is an open beacon and there is a mobile agent at  $a_i$ 's
   location and every beacon  $a_j \in \xi(a_i, t)$  with  $s_j > s_i$  is closed
   and  $\xi(a_i, t)$  does not contain an empty cell then
11:    $a_i$  sets state to "closed beacon"
12: end if

```

We can think of *DLLG* as creating a directed acyclic graph \mathcal{D} whose vertices are beacons and where there is an edge from a beacon a_i to any neighboring beacon whose step count is $s_i + 1$ (Definition 3.2). Mobile agents are only allowed to move via directed edges of this DAG, and they follow the paths induced by its edges until the region is completely filled.

Analysis. The rest of this section concerns the mathematical analysis of DLLG. We prove that when $\Delta T \geq 2$, the termination time step $\lceil T_C(\mathbf{R}) \rceil$ of DLLG over any environment with one entry point obeys the inequality $(2n - 1)\Delta T + 1 \leq \lceil T_C(\mathbf{R}) \rceil \leq (2n - 1)\Delta T + 2$ (Theorem 3.1). This bound on termination is deterministically true regardless of the randomly determined wake-up times of the agents in a given run of DLLG, as long as we agents obey our bounded asynchronous time scheme; in particular, it remains true if we assume all agents always wake up at the same sub-time step, i.e., remains true in a synchronous time scheme.

THEOREM 3.1. *If $\Delta T \geq 2$, the termination time step $\lceil T_C(\mathbf{R}) \rceil$ of DLLG over any environment \mathbf{R} with n cells and one entry point obeys the inequality $(2n - 1)\Delta T + 1 \leq \lceil T_C(\mathbf{R}) \rceil \leq (2n - 1)\Delta T + 2$.*

When there are multiple entry points, or when $\Delta T = 1$, intuition and countless numerical simulations suggest that the termination time of any execution of DLLG is bounded by, but can be lower than $(2n - 1)\Delta T + 2$. The formal time bound analysis of these two situations, and especially the $\Delta T = 1$ case, is (surprisingly) much more complex and is left as an open problem.

For the rest of this section, let p be the lone entry point of \mathbf{R} . As mentioned earlier in this section, DLLG implicitly generates a directed acyclic graph \mathcal{D} which spans \mathbf{R} . Let us formally define \mathcal{D} .

Definition 3.2. In a given execution of DLLG, define $\mathcal{D}(t)$ to be the subgraph of \mathbf{R} rooted at p whose vertices are all the vertices of \mathbf{R} containing beacons at time t , and where there is an edge (u, v) if u contains a beacon a_i and v contains a beacon a_j such that $s_j = s_i + 1$. We define $\mathcal{D} = \mathcal{D}(\infty)$.

We note that \mathcal{D} might not be the same between different runs of DLLG—its precise structure depends on the stochastic activation times of the agents, as well as their sometimes arbitrary movement decisions. Since it is simpler to think in deterministic terms, for the rest of this section we will be studying a fixed, arbitrary execution of DLLG over \mathbf{R} , thus we may assume \mathcal{D} is fixed.

If $\mathcal{D}(t)$ does not contain all cells in \mathbf{R} , then there is at time t a beacon a_i which is adjacent to an empty cell u . There is necessarily also a path in $\mathcal{D}(t)$ from p to a_i . DLLG constantly sends agents across the edges of $\mathcal{D}(t)$, thus eventually an agent will arrive at u . This argument shows that $\mathcal{D}(t)$ grows with t until there is a beacon at every cell of \mathbf{R} . Hence \mathcal{D} contains every vertex of \mathbf{R} .

To prove Theorem 1 we first need to study the structure of \mathcal{D} . We make use of the following definitions:

Definition 3.3. The **children** of a vertex $v \in \mathcal{D}$ are all vertices reachable from v via a directed path of \mathcal{D} . A **leaf** of \mathcal{D} is a vertex $v \in \mathcal{D}$ with no children.

Definition 3.4. Let v be a vertex of \mathcal{D} and let a_i be the beacon at v . Let t be the sub-time step in which a_i becomes a beacon. The depth of v , denoted $depth(v)$, is the step count s_i of a_i at time t (e.g., $depth(p) = 0$, since the beacon at p has step count 0).

Definition 3.5. $dist(p, v)$ is the number of edges of the shortest directed path in \mathcal{D} from p to v .

LEMMA 3.6. $dist(p, v) = depth(v)$

PROOF. By definition, the directed edges of \mathcal{D} determine which vertices a mobile agent can move to from a given location. In the

DLLG algorithm, a mobile agent whose step count is s can only move to vertices of depth $s + 1$. Hence, an agent can only move from p to v in exactly $depth(v)$ steps, which means $dist(p, v)$ must equal $depth(v)$. \square

We define next the “depth” of an agent (Definition 3.7), a time-dependent value associated with the agent. At every time step t , we keep track of the depth of both agents inside \mathbf{R} and agents outside \mathbf{R} that will enter \mathbf{R} in the future. When a_i is inside \mathbf{R} , we shall show in Lemma 3.9 that its depth is related to the depth of the vertex v it is located at (Definition 3.4).

Definition 3.7. Let $W_i(t) := \{t' \mid 0 \leq t' \leq t \wedge a_i \text{ wakes up at time } t'\}$ be the set of wake-up times of agent a_i before time t (inclusive), and define $W_i^*(t) := \{t' \in \mathbb{Z} \mid 0 \leq t' \leq t \wedge a_i \text{ has not entered } \mathbf{R} \text{ at time } t' - dt\}$ to be the set of integer times that a_i begins outside \mathbf{R} .

The depth of an agent a_i at time t is defined as follows:

$$depth(a_i, t) := |W_i(t)| + |W_i^*(t)| - (i - 1)\Delta T$$

The expression $|W_i(t)| + |W_i^*(t)|$ in Definition 3.7 counts the number of times a_i has woken up by time t , treating the integer times where a_i is outside \mathbf{R} as “virtual” wake-ups. Every agent increases its depth exactly once per time step during some sub-time step (note that this is not the same as moving to a vertex of greater vertex depth). Example: suppose that $\Delta T = 2$, that the agent a_3 enters \mathbf{R} at time $t = 4$, and that its first wake-up time occurs at time $5 + 3dt$. Then $depth(a_3, 2) = -2$, $depth(a_3, 3) = -1$, $depth(a_3, 4) = 0$, $depth(a_3, 5 + dt) = 0$, $depth(a_3, 5 + 3dt) = 1$ (see Section 2.2 for the definition of dt). In this example a_3 enters \mathbf{R} exactly when its depth is 0. In Lemma 3.9 we prove this must always be the case.

Definition 3.8. The vertex location of the agent a_i at time t is denoted $v(a_i, t)$.

LEMMA 3.9. *Assuming $\Delta T \geq 2$, at any time $0 \leq t \leq T_C(\mathbf{R})$ the following holds for any vertex $v \in \mathcal{D}$ and any mobile agent a_i :*

- (1) *If $(i - 1)\Delta T \leq t$, a_i has entered \mathbf{R} at time $(i - 1)\Delta T$ (which is the time a_i 's depth equals 0).*
- (2) *If $v(a_i, t)$ contains an open beacon then $depth(a_i, t) = depth(v(a_i, t))$ or the beacon at $v(a_i, t)$ becomes a closed beacon by time $\lceil t \rceil - dt$.*
- (3) *If v contains a closed beacon, it contains a mobile agent and all its children are closed beacons.*

PROOF. (Sketch.) We prove this lemma by induction on t . The base case is straightforward: at any time in $[0, \Delta T)$, there is just one agent at p with depth 0, thus statements (1)-(3) hold. Next, we assume for induction that at times $0, 1, \dots, t$, statements (1)-(3) are true, and show that that (1)-(3) are true at time $t + dt$ as well. Due to space constraints, for the full details of the proof please see the full version of this paper [27]. \square

COROLLARY 3.10. *DLLG fills the graph with $2n$ agents, such that the last agent always enters at time $(2n - 1)\Delta T$.*

Corollary 3.10 follows immediately from Lemma 3.9, (1). We now prove Theorem 3.1.

PROOF. Let us first establish the lower bound. Lemma 3.9, (3) implies that termination can only occur when \mathbf{R} contains $2n$ agents (n mobile and n closed beacons). Corollary 3.10 says that agent a_{2n}

enters at time $(2n - 1)\Delta T$. The beacon at p can only become closed once it wakes up and sees agent a_{2n} , hence $\lceil T_C \mathbf{R} \rceil \geq (2n - 1)\Delta T + 1$.

It remains to establish the upper bound. When a_{2n} enters \mathbf{R} there are n mobile agents in \mathbf{R} . By time $(2n - 1)\Delta T + 2 - dt$, a_{2n} has woken up at least once, hence its depth is 1, but it could not have moved because every vertex in the graph contains a mobile agent. Hence by Lemma 3.9, (2), p is closed at time $(2n - 1)\Delta T + 2 - dt$, establishing $\lceil T_C \mathbf{R} \rceil \leq (2n - 1)\Delta T + 2$. \square

The bound of Theorem 3.1 is precise. Whether the termination time step is $(2n - 1)\Delta T + 1$ or $(2n - 1)\Delta T + 2$ depends on the wake-up times of the agents. The upper bound $(2n - 1)\Delta T + 2$ can for example be attained on a region that is a linear path of n vertices (Figure 3, (a)), by always waking up the beacons before the mobile agents at each time step; the lower bound $(2n - 1)\Delta T + 1$ is attained on the same region by always waking the beacons after the mobile agents.

3.2 Single-Layer Unlimited Gradient

In the previous section we studied the DLLG algorithm, which terminates with $2n$ agents in \mathbf{R} - n agents and n beacons. DLLG is limited by two factors:

- (1) Agents can traverse only the edges of the directed acyclic graph \mathcal{D} , but oftentimes there are edges in \mathbf{R} that are not in \mathcal{D} which could shorten the agents' paths.
- (2) The algorithm requires $2n$ agents to be present in \mathbf{R} before termination can be detected, but only n agents are necessary to cover \mathbf{R} .

The Single-Layer Unlimited Gradient algorithm (Algorithm 2) addresses these weaknesses through three straightforward improvements to DLLG. The first improvement SLUG makes is using “looser” (i.e., less constrained) agent traversal rules. In DLLG, a mobile agent a_i can only move to a beacon whose step count is $s_i + 1$. In SLUG, a_i can move to any beacon whose step count is greater than s_i (if a_i 's step count is 5, it can move to a beacon with step count 6, or 10, or 100, etc.). Under this “*unlimited gradient*” traversal rule, the agents still implicitly generate and traverse a DAG \mathcal{D} that spans the cells of \mathbf{R} , but \mathcal{D} has many more edges for the agents to traverse.

The second improvement of SLUG is its better use of stigmergy and beacons. DLLG is heavily redundant, requiring both a mobile agent and a beacon at every location, and closed beacons are used only to signal the algorithm's termination to an external observer. In SLUG, closed beacons back-propagate without requiring a mobile agent to be on top of them, meaning that we only require a beacon at every vertex to attain termination (hence “Single-Layer” Unlimited Gradient). Moreover, in SLUG, closed beacons are used as a signal for mobile agents not to enter already-explored regions.

The third improvement is giving mobile agents flying above closed beacons the ability to *descend* (rather than ascend) the gradient of step counts, enabling them to move from fully explored areas to either unexplored areas or back to the entry point, if they cannot find any. Specifically, whereas mobile agents located at open beacons attempt to *ascend* the gradient (Algorithm 2, lines 6-8), mobile agents located at closed beacons attempt to *descend* the gradient (lines 9-11). This modification speeds up the exploration process. It also naturally results in redundant mobile agents going back toward an entry point at the end of the algorithm's execution, should the operator wish to recover them. We find this feature

important enough to point out, despite the matter of mobile agent recovery not being a main focus of this work.

Algorithm 2 Local rule for agent a_i waking up at time $t + dt$ and executing Single-Layer Limited Gradient.

Require: a_i has entered \mathbf{R}

- 1: **if** a_i is a “mobile” agent **then**
 - 2: **if** $\xi(a_i, t)$ contains an empty cell **u then**
 - 3: a_i moves to u
 - 4: a_i sets state to “open beacon”
 - 5: $s_i \leftarrow s_i + 1$
 - 6: **else if** $\xi(a_i, t)$ contains a cell u with just one agent a_j **and** a_j is an open beacon **and** $s_j > s_i$ **then**
 - 7: a_i moves to u
 - 8: $s_i \leftarrow s_j$
 - 9: **else if** a_i 's current location contains a closed beacon **and** $\xi(a_i, t)$ contains a cell u with just one agent a_j **and** a_j is an open or closed beacon **and** $s_j < s_i$ **then**
 - 10: a_i moves to u
 - 11: $s_i \leftarrow s_j$
 - 12: **end if**
 - 13: **else if** a_i is an open beacon **and** every beacon $a_j \in \xi(a_i, t)$ with $s_j > s_i$ is closed **and** $\xi(a_i, t)$ does not contain an empty cell **then**
 - 14: a_i sets state to “closed beacon”
 - 15: **end if**
-

An example run of SLUG in a square grid region is depicted in Figure 2. Simulation results suggest that SLUG terminates faster than DLLG, and that its total energy use tends to be an order of magnitude smaller (see Section 4).

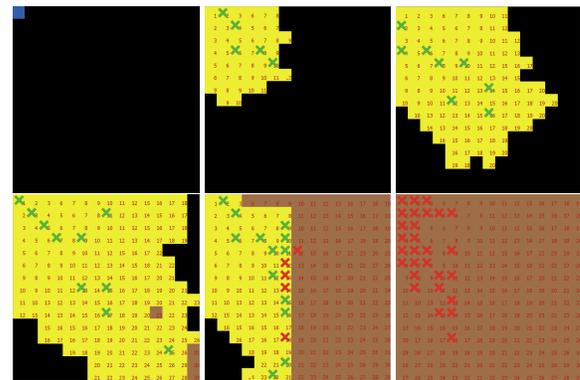


Figure 2: Snapshots from a run of SLUG ordered left to right, top to bottom. Empty cells are black, obstacles are white, cells containing an open beacon are yellow and cells containing a closed beacon are brown. Beacons' step count is denoted by the red numbers. Green “X”s are mobile agents that have never entered a closed beacon's location, and red “X”s are mobile agents that have. Red agents descend the gradient until they find new open beacons. $\lceil T_C(\mathbf{R}) \rceil = 508$.

“Ant-like” recovery from crash faults and errors. An often-sought property of multi-robot systems is their ability to recover from crash faults and memory errors. The mobile agent layer of SLUG

is provably resilient to crash faults: a mobile robot can crash and disappear from the environment at any time without affecting the algorithms' correctness (and often barely affecting termination time). Furthermore, three straight-forward changes to SLUG can be implemented to make the mobile agents provably resilient to memory corruption and positional errors, without any additional assumptions. Due to space constraints, please see the full version of this paper: [27].

4 SIMULATION RESULTS AND BENCHMARKS

In this section we employ numerical simulations to study the performance of SLUG and DLLG in terms of energy, termination time, and other related parameters, and compare these algorithms to well-known algorithms from the literature.

Recall that SLUG makes two independent conceptual improvements over DLLG: (i) it has looser traversal rules, and (ii) it makes better use of closed beacons. It is interesting to ask what each of these improvements provides: what is the impact of looser, or stricter traversal rules on DLLG's termination time and energy use? What is the impact of the closed beacon strategy?

We investigate these questions numerically by comparing the termination time and energy use of DLLG to SLUG and to two variants of DLLG: a variant with stricter traversal rules called *Dual-Layer Tree Traversal* (DLTT), and a variant with looser traversal rules called *Dual-Layer Unlimited Gradient* (DLUG). All three Dual-Layer algorithms operate under the same underlying principle of gradient climbing and dual-layer coverage, filling \mathbf{R} with n mobile agents and n closed beacons—the only difference is the underlying directed acyclic graph which the agents construct and traverse, \mathcal{D} . In DLTT, an algorithm first described in [3] and which can directly be compared to our algorithms as a benchmark, \mathcal{D} is restricted to be a tree. This results in more limited traversal than DLLG, where \mathcal{D} is not necessarily a tree. In DLUG, similar to SLUG, the agents are not restricted in terms of gradient steepness: an agent a_i can move to any neighboring beacon whose step count is at least $s_i + 1$, thus \mathcal{D} has more edges compared to DLLG (where a_i can only move to a neighboring beacon whose step count is exactly $s_i + 1$), and mobile agents' traversal rules are consequently less restricted. Comparing DLUG, DLLG and DLTT gives us insight into the impact of traversal rules on the energy and time costs of a dual-layer coverage algorithm. We shall show that DLUG, which has the least restrictive gradient ascent rules, outperforms DLGG and DLTT. Comparing these algorithms to SLUG gives us insight into the cumulative impact of using looser traversal rules *and* the closed beacon strategy. We find that SLUG significantly outperforms DLLG, DLTT and DLUG.

Our numerical simulations were run on the regions depicted in Figure 3. The top three regions are simple regions that can be generalized to different sizes, and through which the effect of region size on termination time and energy use can be investigated. The three regions depicted in the bottom row were replicated from [13] as examples of complex regions that more closely resemble real-life scenarios. Each data point in the figures of this section is the average of data from 50 runs of an algorithm on one of these regions (when simulating the time model we set $M = 100$).

The impact of ΔT on DLLG and SLUG's termination time step $[T_C(\mathbf{R})]$ is investigated in Figure 4, (a) for the 6 types of regions

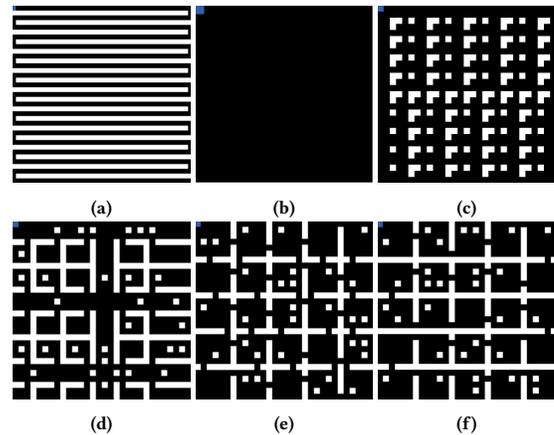


Figure 3: Regions used in simulation experiments. Obstacles are marked in white and empty cells are marked in black. The entry point is at the upper-left corner, marked in blue. Region labels are as follows: (a) linear (b) square (c) sawtooth (d) complex 1 (e) complex 2 (f) complex 3. Regions (a), (b) and (c) can be generalized to different sizes - e.g., we run simulations on linear regions of different sizes.

depicted in Figure 3. We see that for every value of ΔT and every region, SLUG's termination time is significantly lower than DLLG's. Figure 4, (a) also confirms the termination time bounds of Theorem 3.1 for DLLG: as claimed, we clearly see that $[T_C(\mathbf{R})]/n \approx 2\Delta T$ in the plot of Figure 4, (a) corresponding to DLLG. Figure 4, (b) investigates the number of agents present in a region at termination time when using SLUG and DLLG. By construction, every region always contains $2n$ agents at DLLG's termination, but when it comes to SLUG, the ratio of agents to region size quickly converges to 1 as ΔT goes to infinity. In other words, as ΔT grows, there are less redundant agents in the environment. Hence, by increasing ΔT in SLUG, we can reduce the number of redundant agents at the cost of larger makespan.

Figure 5 depicts the effect of region size on the termination time and total energy use of SLUG, DLTT, DLLG and DLUG when $\Delta T = 1$ (Figure 5, (a) and (c)) and when $\Delta T = 2$ (Figure 5, (b) and (d)). We see that SLUG significantly outperforms DLTT, DLGG and DLUG in terms of energy use and termination time. Additionally, several facts about the impact of traversal rules can be inferred from the figure: first, looser traversal rules confer a benefit on energy use: E_{total} is reduced in DLUG compared to DLTT and DLLG. Next, when $\Delta T \geq 2$, termination time is unaffected by the looseness or strictness of traversal rules. However, when $\Delta T = 1$, looser traversal rules improve termination time: of the three dual-layer algorithms, DLUG's termination time is lowest, followed by DLLG.

Benchmarks from the literature. We compared SLUG to the well-known Depth-First and Breadth-First Leader-Follower coverage algorithms of Hsiang et al. [19] in terms of termination time and energy use. Hsiang's algorithms assume synchronous time, a greater sensing range than ours, and direct radio-based inter-agent communication. Due to these differing assumptions, an exact one-to-one comparison of our work and [19] is not possible—but we can nevertheless draw some conclusions. Figure 6, (a) compares the termination time of SLUG to Hsiang et al.'s algorithms on square grid regions. Assuming $\Delta T = 1$, SLUG's termination time is significantly lower than BFLF or DFLF's, whereas when $\Delta T = 2$, its

REFERENCES

- [1] Yaniv Altshuler, Alex Pentland, and Alfred M Bruckstein. *Introduction to Swarm Search*. Springer, 2018.
- [2] Michael Amir and Alfred M. Bruckstein. Minimizing Travel in the Uniform Dispersal Problem for Robotic Sensors. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, pages 113–121, Richland, SC, 2019. International Foundation for Autonomous Agents and Multiagent Systems.
- [3] Michael Amir and Alfred M. Bruckstein. Fast uniform dispersion of a crash-prone swarm. *Robotics: Science and Systems*, 2020.
- [4] Fidel Aznar, Mar Pujol, Ramon Rizo, Francisco A. Pujol, and Carlos Rizo. Energy-Efficient Swarm Behavior for Indoor UAV Ad-Hoc Network Deployment. *Symmetry* (20738994), 10(11):632, November 2018. Publisher: MDPI Publishing.
- [5] Fidel Aznar, Mar Pujol, Ramón Rizo, and Carlos Rizo. Modelling multi-rotor UAVs swarm deployment using virtual pheromones. *PLOS ONE*, 13(1):e0190692, January 2018. Publisher: Public Library of Science.
- [6] Fidel Aznar Gregori, Mar Pujol, and Ramón Rizo. UAV Deployment Using Two Levels of Stigmergy for Unstructured Environments. October 2020. Accepted: 2020-11-05T14:46:58Z Publisher: MDPI.
- [7] Eduardo Mesa Barrameda, Shantanu Das, and Nicola Santoro. Deployment of Asynchronous Robotic Sensors in Unknown Orthogonal Environments. In Sándor P. Fekete, editor, *Algorithmic Aspects of Wireless Sensor Networks*, volume 5389, pages 125–140. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [8] Eduardo Mesa Barrameda, Shantanu Das, and Nicola Santoro. Uniform Dispersal of Asynchronous Finite-State Mobile Robots in Presence of Holes. In Paola Flocchini, Jie Gao, Evangelos Kranakis, and Friedhelm Meyer auf der Heide, editors, *Algorithms for Sensor Systems*, volume 8243, pages 228–243. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
- [9] Lali Barrière, Paola Flocchini, Eduardo Mesa-Barrameda, and Nicola Santoro. UNIFORM SCATTERING OF AUTONOMOUS MOBILE ROBOTS IN A GRID. *International Journal of Foundations of Computer Science*, 22(03):679–697, April 2011.
- [10] J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, April 2004.
- [11] Yann Disser, Jan Hackfeld, and Max Klimm. Tight bounds for undirected graph exploration with pebbles and multiple agents. *Journal of the ACM (JACM)*, 66(6):1–41, 2019.
- [12] Xavier Défago, Maria Gradinariu, Stéphane Messika, and Philippe Raipin-Parvédy. Fault-Tolerant and Self-stabilizing Mobile Robots Gathering. In Shlomi Dolev, editor, *Distributed Computing*, Lecture Notes in Computer Science, pages 46–60, Berlin, Heidelberg, 2006. Springer.
- [13] Ettore Ferranti, Niki Trigoni, and Mark Levene. Brick Mortar: an on-line multi-agent exploration algorithm. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 761–767, April 2007. ISSN: 1050-4729.
- [14] Nobuhiro Funabiki, Benjamin Morrell, Jeremy Nash, and Ali-akbar Aghamohammadi. Range-aided pose-graph-based slam: Applications of deployable ranging beacons for unknown environment exploration. *IEEE Robotics and Automation Letters*, 6(1):48–55, 2020.
- [15] Attila Hideg and Tamás Lukovszki. Asynchronous filling by myopic luminous robots. In *International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics*, pages 108–123. Springer, 2020.
- [16] Attila Hideg, Tamás Lukovszki, and Bertalan Forstner. Improved runtime for the synchronous multi-door filling. *Periodica Polytechnica Electrical Engineering and Computer Science*, 66(1):12–19, 2022.
- [17] Andrew Howard, Maja J. Mataric, and Gaurav S. Sukhatme. An Incremental Self-Deployment Algorithm for Mobile Sensor Networks. *Autonomous Robots*, 13(2):113–126, September 2002.
- [18] Andrew Howard, Maja J. Mataric, and Gaurav S. Sukhatme. Mobile Sensor Network Deployment using Potential Fields: A Distributed, Scalable Solution to the Area Coverage Problem. *SpringerLink*, pages 299–308, 2002.
- [19] Tien-Ruey Hsiang, Esther M. Arkin, Michael A. Bender, Sándor P. Fekete, and Joseph S. B. Mitchell. Algorithms for Rapidly Dispersing Robot Swarms in Unknown Environments. *SpringerLink*, pages 77–93, 2004.
- [20] James F. Kennedy, Russell C. Eberhart, and Yuhui Shi. *Swarm intelligence*. The Morgan Kaufmann series in evolutionary computation. Morgan Kaufmann Publishers, San Francisco, 2001.
- [21] Alexander B Maxseiner, Daniel M Lofaro, and Donald A Sofge. Visible light communications with inherent agent localization and simultaneous message receiving capabilities for robotic swarms. In *2021 18th International Conference on Ubiquitous Robots (UR)*, pages 633–639. IEEE, 2021.
- [22] Yash Mulgaonkar, Anurag Makineni, Luis Guerrero-Bonilla, and Vijay Kumar. Robust Aerial Robot Swarms Without Collision Avoidance. *IEEE Robotics and Automation Letters*, 3(1):596–603, January 2018.
- [23] Anil Özdemir, Melvin Gauci, Andreas Kolling, Matthew D Hall, and Roderich Groß. Spatial coverage without computation. In *2021 18th International Conference on Robotics and Automation (ICRA)*, pages 9674–9680. IEEE, 2019.
- [24] David Peleg. Distributed Coordination Algorithms for Mobile Robot Swarms: New Directions and Challenges. In *Distributed Computing – IWDC 2005*, pages 1–12. Springer, Berlin, Heidelberg, December 2005.
- [25] Dmitry Rabinovich, Michael Amir, and Alfred M. Bruckstein. Optimal physical sorting of mobile agents. <https://arxiv.org/abs/2111.06284>, 2021.
- [26] O. Rappel and J.Z. Ben-Asher. Area coverage – A swarm based approach. In *59th Israel Annual Conference on Aerospace Sciences, IACAS 2019*, pages 35–55. Israel Annual Conference on Aerospace Sciences, 2019.
- [27] Ori Rappel, Michael Amir, and Alfred M. Bruckstein. Stigmergy-based, dual-layer coverage of unknown indoor regions. [arXiv 2209.08573](https://arxiv.org/abs/2209.08573) (<https://arxiv.org/abs/2209.08573>), 2022.
- [28] Timothy Stirling and Dario Floreano. Energy-Time Efficiency in Aerial Swarm Deployment. In Alcherio Martinoli, Francesco Mondada, Nikolaus Correll, Grégory Mermoud, Magnus Egerstedt, M. Ani Hsieh, Lynne E. Parker, and Kasper Støy, editors, *Distributed Autonomous Robotic Systems*, volume 83, pages 5–18. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [29] Timothy Stirling, Steffen Wischmann, and Dario Floreano. Energy-efficient indoor search by swarms of simulated flying robots without global information. *Swarm Intelligence*, 4(2):117–143, 2010.
- [30] Vu Phi Tran, Matthew A Garratt, Kathryn Kasmarik, and Sreenatha G Anavatti. Robust multi-robot coverage of unknown environments using a distributed robot swarm. *arXiv preprint arXiv:2111.14295*, 2021.
- [31] Ouarda Zedadra, Nicolas Jouandeau, Hamid Seridi, and Giancarlo Fortino. Multi-agent foraging: state-of-the-art and research challenges. *Complex Adaptive Systems Modeling*, 5(1):1–24, 2017.