



HAL
open science

Treasure Hunt with Volatile Pheromones

Evangelos Bampas, Joffroy Beauquier, Janna Burman, William Guy–Obé

► **To cite this version:**

Evangelos Bampas, Joffroy Beauquier, Janna Burman, William Guy–Obé. Treasure Hunt with Volatile Pheromones. 37th International Symposium on Distributed Computing (DISC 2023), Oct 2023, L'Aquila, Italy. pp.8:1-8:21, 10.4230/LIPIcs.DISC.2023.8 . hal-04470589

HAL Id: hal-04470589

<https://hal.science/hal-04470589>

Submitted on 21 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Treasure Hunt with Volatile Pheromones

Evangelos Bampas ✉ 🏠 

Université Paris-Saclay, CNRS, Laboratoire Interdisciplinaire des Sciences du Numérique, 91400 Orsay, France

Joffroy Beauquier ✉

Université Paris-Saclay, CNRS, Laboratoire Interdisciplinaire des Sciences du Numérique, 91400 Orsay, France

Janna Burman ✉

Université Paris-Saclay, CNRS, Laboratoire Interdisciplinaire des Sciences du Numérique, 91400 Orsay, France

William Guy-Obé ✉

Université Paris-Saclay, CNRS, Laboratoire Interdisciplinaire des Sciences du Numérique, 91400 Orsay, France

Abstract

In the treasure hunt problem, a team of mobile agents need to locate a single treasure that is hidden in their environment. We consider the problem in the discrete setting of an oriented infinite rectangular grid, where agents are modeled as synchronous identical deterministic time-limited finite-state automata, originating at a rate of one agent per round from the origin. Agents perish τ rounds after their creation, where $\tau \geq 1$ is a parameter of the model. An algorithm solves the treasure hunt problem if every grid position at distance τ or less from the origin is visited by at least one agent. Agents may communicate only by leaving indistinguishable traces (pheromone) on the nodes of the grid, which can be sensed by agents in adjacent nodes and thus modify their behavior. The novelty of our approach is that, in contrast to existing literature that uses permanent pheromone markers, we assume that pheromone traces evaporate over μ rounds from the moment they were placed on a node, where $\mu \geq 1$ is another parameter of the model. We look for uniform algorithms that solve the problem without knowledge of the parameter values, and we investigate the implications of this very weak communication mechanism to the treasure hunt problem. We show that, if pheromone persists for at least two rounds ($\mu \geq 2$), then there exists a treasure hunt algorithm for all values of agent lifetime. We also develop a more sophisticated algorithm that works for all values of μ , hence also for the fastest possible pheromone evaporation of $\mu = 1$, but only if agent lifetime is at least 16.

2012 ACM Subject Classification Theory of computation \rightarrow Distributed algorithms; Theory of computation \rightarrow Design and analysis of algorithms

Keywords and phrases Mobile Agents, Exploration, Search, Treasure Hunt, Pheromone, Evaporation

Digital Object Identifier 10.4230/LIPIcs.DISC.2023.8

Related Version *Full Version:* <https://hal.science/hal-04177364v1>

Acknowledgements We thank the DISC 2023 anonymous reviewers for their careful reading and comments.

1 Introduction

Treasure hunt is the fundamental problem of employing a team of searchers to locate a “treasure” that is hidden somewhere in their environment. It is one of the fundamental primitives in swarm robotics and a natural abstraction of foraging behavior of animals. Although various formulations of the problem exist at least since the 1960s, when Beck



© Evangelos Bampas, Joffroy Beauquier, Janna Burman, and William Guy-Obé; licensed under Creative Commons License CC-BY 4.0

37th International Symposium on Distributed Computing (DISC 2023).

Editor: Rotem Oshman; Article No. 8; pp. 8:1–8:21



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

introduced the linear search problem [14], treasure hunt as a group search problem was first investigated from a distributed algorithms perspective by Feinerman et al. [41, 42, 43], under the name ANTS (Ants Nearby Treasure Search). In the ANTS problem, the search is performed by a team of randomized searchers, starting at the origin of an infinite 2-dimensional rectangular grid and having no means of communication once they start moving. Subsequent works considered stronger communication models, such as local communication by exchanging constant-size messages when two agents are located on the same node [40, 39, 25, 21, 54, 53], or communication by leaving permanent markers on grid nodes [56, 1, 2], that can be detected by other agents.

In this paper, we introduce a new model in which not only agents communicate indirectly, by dropping and sensing markers on nodes, but also these markers gradually evaporate and eventually disappear. This is directly inspired by the behavior of actual pheromone trails in nature. A common feature of the papers that we mentioned above is that the team of searchers is of constant size. However, with evaporating pheromones, we can no longer expect a constant-size team of constant-memory agents to explore all the grid nodes up to arbitrary distances.¹ Therefore, we propose a new model taking into account pheromone evaporation, in which a potentially infinite number of identical, synchronous, deterministic, time-limited finite-state automata are created at a rate of one agent per round at the origin of a 2-dimensional grid. Agents have a finite lifetime represented by the parameter τ , and the treasure is guaranteed to be within reach, i.e., at distance $\leq \tau$ from the origin. Pheromone evaporation is controlled by a parameter μ , which determines the number of rounds it takes for a pheromone marker to disappear from the system, assuming it is not refreshed in the meantime by a new pheromone drop on the same node. Agents can sense the presence or absence of pheromone in their neighboring nodes, and they can compare pheromone values, i.e., they know, for any pair of directions, which neighbor has the freshest pheromone. Agent memory cannot depend on the parameters τ, μ .

1.1 Related work

Searching is a well-studied family of problems in which a group consisting of one or multiple searchers (mobile agents) need to find a target placed at some unknown location. The search is typically concluded when the first searcher finds the target. Numerous books and research papers have been written on this subject, studying diverse models involving stationary or mobile targets, graphs or geometric terrains, different types of knowledge about the environment, one or many searchers, etc. [5, 6, 17, 23, 45, 47, 58].

Deterministic search on a line with a single robot was introduced in [14, 15]. In the original formulation of [14], a probability distribution of treasure placements is known to the agent. An optimal algorithm with competitive ratio 9, for an unknown probability distribution, is proposed in [15]. The problem is further generalized in [8, 35], by introducing a cost for turning, as well as a more general star topology. Further variants include searching for multiple targets [9], maximizing the searched area with a given time budget [10], and providing a hint to the searcher before it starts exploring [7].

¹ Indeed, intuitively, if they find themselves sufficiently far from each other, then they can no longer communicate because pheromone will evaporate before it can be sensed by another agent, whereas if they never find themselves more than a constant distance from each other, then their overall behavior can be described by a single finite automaton, which fails to explore a sufficiently large grid due to state repetition that forces it to explore at most a constant-width half-line (see, e.g., [39, Lemma 5]).

Various maintenance and patrolling problems have also been formulated as linear group search problems, under requirements and assumptions such as perpetual exploration [27, 50, 26] or distinct searcher speeds [29, 50, 13]. The closely related evacuation problem on the line, in which the search is concluded when *all* searchers reach the target, has also been studied in a series of papers [11, 22, 12, 20]. See also [30] for a survey of group search and evacuation in different domains.

Searching with advice (hints) is studied under various assumptions in several papers. The size of advice that must be provided to a lone deterministic searcher in a polygonal terrain with polygonal obstacles, in order to locate the treasure at a cost linear in the length of a shortest path from the initial position of the agent to the treasure, is investigated in [59]. An algorithm that enables a deterministic agent to find an inert treasure in the Euclidean plane, taking advantage of hints about the general direction of the treasure, is given in [18]. In [51, 57], they explore the tradeoff between advice size and search cost in graphs. In trees, [19] explores the impact of different kinds of initial knowledge given to a lone searcher on the time cost of treasure hunt, and [16] considers treasure hunt with faulty hints.

The speedup in search time obtained by multiple independent random walkers has been studied for various graph families, such as expanders and random graphs [37, 4, 38, 48, 28]. Multiple searchers following Lévy walk processes, a type of random walk in which jump lengths are drawn from a power-law distribution, and for which there is significant empirical evidence that it models the movement patterns of various animal species [31], are investigated in [24]. A further abstraction of multiple independent randomized searchers is studied in [46], where a group of non-communicating agents need to find an adversarially placed treasure, hidden in one of an infinite set of boxes indexed by the natural numbers. In this Bayesian search setting, searchers have random access to the boxes. A game-theoretic perspective to the Bayesian search framework of [46] is given in [52].

The ANTS (Ants Nearby Treasure Search) problem was introduced in [41, 42, 43] as a natural abstraction of foraging behavior of ants around their nest. They explore the tradeoff between searcher memory and the speedup obtained by using multiple probabilistic searchers vs using a single searcher. Searchers may not communicate once they leave the nest. A variant of the ANTS problem in the geometric plane, with searchers that are susceptible to crash faults, is investigated in [3]. A notion of selection complexity, which measures how likely a given ANTS algorithm is to arise in nature, is introduced in [55], where they study the tradeoff between selection complexity and speedup in search time.

In follow-up work [40, 39, 25, 21, 54, 53] to the original ANTS formulation, searchers are modeled as finite state machines and can communicate outside the nest, when they are sufficiently close to each other, by exchanging messages of constant size. Under these assumptions, it is shown in [40] that the optimal search time can still be achieved by probabilistic finite state machines, matching the lower bound of [42]. The minimum number of searchers that can solve the ANTS problem, when they are controlled by randomized/deterministic finite/push-down automata, is investigated in [39, 25, 21]. A probabilistic fault-tolerant constant-memory algorithm is presented in [54], for the synchronous case. An algorithm that tolerates obstacles is presented in [53].

A different communication mechanism is considered in [56, 1, 2], where it is assumed that agents may communicate only by leaving permanent markers (pheromones) on nodes, which can be sensed later by other agents. Note that, although these papers use the word “pheromone” to describe the traces that agents leave on nodes, these are assumed permanent and do not evaporate. The usual term in the mobile agent literature to describe this type of movable or immovable marker that agents may choose to leave on nodes, and which can be detected later by other agents, is “token” or “pebble” [34].

1.2 Our contributions

We study the treasure hunt problem in the model that we outlined above, and which is developed in detail in Section 2. Thematically, our work is closest to the literature descending from the original ANTS problem formulation, and in particular to these papers that use pheromone (or tokens) as a means of communication [56, 1, 2]. The novelty of our approach is that we use *evaporating* pheromones as an agent communication mechanism. Indeed, in our model, a pheromone trace disappears μ rounds after it was dropped, unless it is refreshed by a new pheromone drop on the same node. Tokens that may disappear instantly from the system have actually been considered before in the literature, but only in the context of faults [44, 33, 32, 36].

To our knowledge, evaporating pheromone markers have never been considered before as a communication mechanism, from an algorithmic point of view. We study the impact of this weak agent communication model on the treasure hunt problem.

Our first result is a treasure hunt algorithm that works for all $\tau \geq 1$, assuming that the pheromone markers persist for at least two rounds ($\mu \geq 2$). This algorithm is optimal in terms of search time, number of pheromone drops, and number of agents used. Intuitively, the algorithm dispatches agents to the North and to the South of the origin by means of pheromone patterns around the nest. An agent knows when to leave the vertical axis in order to explore a horizontal half-line by detecting pheromone markers that were dropped by previous agents when they left the vertical axis. Because of the North-South dispatching at the origin, successive agents on the same side of the origin are at distance 2 from each other, therefore it is crucial that $\mu \geq 2$ for an agent to be able to detect pheromone that was dropped by the previous agent.

Our second result is a more complex algorithm that works for all $\mu \geq 1$. This algorithm is also based on a North-South dispatching of agents at the origin. The challenge here is that, since pheromone may be detectable for only one round after being dropped, we can no longer use the same approach as in the first algorithm. We resolve this by introducing *differentiation* of agent roles as a result of observing different pheromone patterns as they walk along the vertical axis. Now, some agents become *signaling* agents that stop moving at key positions and start dropping pheromone according to a predetermined pattern, whereas other agents become *explorers* that are dispatched to different horizontal half-lines according to these signals. This algorithm works for all $\tau \geq 16$.

Both algorithms are deterministic and uniform, i.e., they do not assume knowledge of the values of the parameters τ, μ . They solve the problem for *all* parameter values in the specified ranges, and the required memory per agent is constant.

Some proofs have been omitted due to lack of space, and they can be found in the full version of the paper.

2 Model and problem setting

The environment in which the agents operate is an infinite two-dimensional rectangular grid graph, equipped with a Cartesian coordinate system. Each node of the grid is identified by a pair of integer coordinates $(x, y) \in \mathbb{Z}^2$. The node $(0, 0)$ is called *nest*, as newly created agents appear at $(0, 0)$. We assume that the grid is oriented, with the four outgoing edges from each node receiving globally consistent distinct local port labels from $\{N, E, S, W\}$. Each node stores a nonnegative integer that represents the amount of pheromone present on that node. This value is decremented by 1 at each round and a value of zero represents the absence of pheromone.

Given natural numbers a, b , we use the notation $a \dot{-} b$ for proper subtraction: $a \dot{-} b = \max(a - b, 0)$. Moreover, if x is a nonnegative integer, we use \mathbb{B}_x for the set of nodes at distance at most x from the nest, and \mathbb{L}_x for the set of nodes at distance exactly x from the nest. We have $|\mathbb{B}_x| = 2x^2 + 2x + 1$ and $|\mathbb{L}_x| = 4x$.

2.1 Agent model

Agents are modeled as identical copies of a deterministic finite-state machine (FSM). An agent can move from node to node along the edges of the grid graph, and it may decide to drop pheromone before or after each move (but not both on the origin and on the destination node). It computes its next move based on the relative pheromone values of the neighboring nodes. More precisely, the agent does not have access to the actual stored pheromone values, but it can detect the presence or absence of pheromone in any direction, as well as whether one direction has equal, less, or more pheromone than another.

- **Definition 1 (Agents).** *An agent is a finite-state machine $\mathcal{A} = (Q, q_0, \text{In}, \text{Out}, \delta)$ where:*
- Q is a finite set of states and $q_0 \in Q$ is the initial state.
 - In is the input alphabet. A symbol of In encodes the presence or absence of pheromone in the four cardinal directions, as well as the result of the comparison of pheromone levels for any pair of directions. This is clearly a finite amount of information, hence In is a finite set.
 - $\text{Out} = \{N, S, E, W, \perp\} \times \{\text{before}, \text{after}, \perp\}$ is the output alphabet, where the first element of an output symbol is the local port label through which the agent will exit the current node (\perp for no movement), and the second element indicates whether pheromone will be dropped before or after the move (\perp for no pheromone drop).
 - $\delta : Q \times \text{In} \rightarrow Q \times \text{Out}$ is the transition function.

► **Note 2.** By definition, an agent does not perceive other agents that may be present on the same node or on neighboring nodes. Moreover, an agent does not perceive and therefore cannot compare the pheromone level of its current node to those of neighboring nodes.

2.2 Model parameters

Agents have limited life, which is a parameter of the model and is represented by a positive integer τ . An agent “dies” upon having performed τ state transitions, meaning that it essentially disappears from the system.² We will call this parameter *lifetime*.

We also assume that every node has a maximum amount of pheromone that it can store, which is a second parameter of the model and is represented by a positive integer μ , which we will call *pheromone duration*. Whenever any number of agents decide to drop pheromone on a node at the same time, the pheromone value of that node is updated to μ . If an agent drops pheromone on some node, the pheromone value of that node will decrease from μ to 0 over the following μ rounds (assuming it is not refreshed in the meantime).

2.3 Execution

Given a protocol \mathcal{A} (FSM) and an assignment of values to the parameters (τ, μ) , the execution of the system proceeds deterministically in synchronous rounds.

² Perhaps less fatally, we may assume that after τ transitions, an agent is so tired that it cannot continue executing the protocol before returning to the nest for a brief nap. It may re-emerge from the nest at a later round without retaining its state.

► **Definition 3** (Execution). *The execution of an FSM \mathcal{A} for parameter values (τ, μ) is an infinite sequence of system configurations defined as follows: In the initial configuration, there are no agents and no pheromone present on the grid. In each round i ($i \geq 1$), the next configuration is obtained from the current configuration by synchronously executing the following steps in the given order:*

1. A new agent (copy of \mathcal{A}) is created on node $(0,0)$, in the initial state q_0 .
2. All agents read their inputs.
3. At each node, the quantity of pheromone is decreased by 1, if not already zero (pheromone evaporation).
4. All agents compute their transition function based on the input from step 2 and change their state accordingly.
5. All agents that computed in step 4 a pheromone drop action “before” drop pheromone on their current nodes. For each node on which at least one agent drops pheromone, the pheromone quantity of that node is updated to μ .
6. All agent moves (as computed in step 4) are executed.
7. All agents that computed in step 4 a pheromone drop action “after” drop pheromone on their current nodes. For each node on which at least one agent drops pheromone, the pheromone quantity of that node is updated to μ .
8. If this is round $i \geq \tau$, the agent that was created at the beginning of round $i - \tau + 1$ “dies” as it has now performed τ state transitions.

► **Note 4.** As agents are anonymous and deterministic, and because pheromone does not accumulate higher than μ on a single node, if two (or more) agents ever find themselves at the same node and in the same state, then they will effectively continue moving as one agent from that point on. In particular, agents do not appear simultaneously at the nest, but they are created at a rate of one agent per round.

► **Definition 5.** For $i \geq 1$, we denote by A_i the agent that is created at the beginning of round i .

2.4 The treasure hunt problem


In the *treasure hunt* problem, a treasure is placed at an unknown location in the grid and the goal is for at least one agent to visit that node. In that case, we say that the agent *locates* the treasure. Locating the treasure for any (unknown) treasure location up to distance d from the nest is trivially equivalent to exploring all nodes up to distance d from the nest. We recast, then, the treasure hunt problem as an exploration problem:

► **Definition 6** (Treasure hunt problem). *A given FSM \mathcal{A} solves the treasure hunt problem for the pair of parameters (τ, μ) if, with lifetime τ and pheromone duration μ , every node at distance τ or less is visited by at least one agent. In this case, we will also say that the FSM is correct for (τ, μ) .*

We will seek a uniform algorithm that solves the problem without knowledge of the model parameters, i.e., a single FSM that is correct for arbitrarily large values of τ and μ (ideally, for all $\tau \geq \tau_0$ and $\mu \geq \mu_0$, for some constants τ_0, μ_0).

For a given FSM, we will consider the following measures of efficiency as functions of τ and μ :

- *Completion time:* the number of rounds until the treasure is located.
- *Pheromone utilization:* the total number of times that any agent decides to drop pheromone at its destination node until the treasure is located.



■ **Figure 1** Illustration of the sequence of states for a typical agent with signature $[v; h]$ executing Algorithm 1.

- *Agent utilization*: the number of agents *effectively* used by the algorithm, i.e., the smallest r such that the algorithm remains correct even if the system stops creating new agents after round r .

3 A treasure hunt algorithm for $\tau \geq 1$ and $\mu \geq 2$

We propose a deterministic and uniform algorithm that solves the treasure hunt problem for all combinations of parameters (τ, μ) with $\tau \geq 1$ and $\mu \geq 2$. We give a compact representation of the algorithm as a hybrid state transition diagram in Appendix A, and the full pseudocode in Section 3.1.

Before we give an informal description of the algorithm, we define the notion of *agent signature*:

► **Definition 7 (Agent signature)**. Let $v, h \in \mathbb{Z}$ with $|v| + |h| = \tau$. We say that an agent has signature $[v; h]$ if it starts (from the nest) by moving $|v|$ steps to the North (resp. South), up to node $(0, v)$, if $v \geq 0$ (resp. $v < 0$), followed by $|h|$ steps to the East (resp. West), up to node (h, v) , if $h \geq 0$ (resp. $h < 0$).

The algorithm creates agents of all possible signatures $[v; h]$, thus ensuring correctness by visiting all nodes at distance $\leq \tau$ from the nest. Each agent drops pheromone once upon leaving the nest on its first move, and once more if and when it leaves the vertical axis. Figure 1 shows the sequence of states of a typical agent executing the algorithm.

The first two agents use pheromone information to the East and to the West of the nest to take signatures $[0, -\tau]$ and $[0, \tau]$ (state *Init*, lines 2-5). Subsequently created agents use pheromone information to the North and to the South of the nest to alternate between the two vertical directions: If there is more pheromone to the North of the nest then they start moving South, otherwise they start moving North (state *Init*, lines 7-9).

A northbound agent (southbound agents behave symmetrically) starts moving to the North in state *Vert-seek*. In this state, it checks horizontally adjacent nodes for the presence of pheromone previously dropped by agents leaving the vertical axis. Once it finds such pheromone traces, it switches to state *Vert-bypass* and keeps moving to the North until it reaches the first node $(0, v)$ whose East and West neighbors do not *both* have pheromone.

At that point, if no horizontal neighbor has pheromone then it turns East, taking signature $[v, \tau - v]$, whereas if only the East neighbor has pheromone then it turns West, taking signature $[v, v - \tau]$ (state *Vert-bypass*). Once it leaves the vertical axis, an agent keeps moving horizontally until the end of its lifetime in state *Horiz*.

3.1 Pseudocode

We give the transition function executed by each agent during step 4 of each round (cf. Definition 3) in Algorithm 1. We denote by φ_x , for $x \in \{N, E, S, W\}$, the pheromone value of the neighboring node in direction x . These represent the input to the FSM. In accordance with Definition 1, the pheromone values are never used directly but only as part of comparisons to each other and to the value 0. The output of the FSM is composed of the pair of values $(\text{dir}, \text{drop})$ at the end of the transition function computation.

3.2 Correctness

► **Theorem 8.** *Algorithm 1 correctly solves the treasure hunt problem for all combinations of parameters (τ, μ) with $\tau \geq 1$ and $\mu \geq 2$.*

The complete proof of Theorem 8 is available in the full version. The proof is based on the following simple properties of Algorithm 1:

- Whenever an agent switches to state **Horiz** it moves horizontally (East or West) and drops pheromone. Subsequently, it keeps moving in the same direction in the same state without dropping pheromone until the end of its lifetime.
- Whenever an agent switches to state **Vert-seek** it moves vertically (North or South) and drops pheromone. Subsequently, it keeps moving in the same direction in the same state without dropping pheromone until one of the following happens: it reaches the end of its lifetime, or it switches to state **Vert-bypass** moving in the same direction as before, or it switches to state **Horiz** moving West.
- Whenever an agent switches to state **Vert-bypass** it moves vertically, and it drops pheromone only if it switches from state **Init** to **Vert-bypass**. Subsequently, it keeps moving in the same direction in the same state without dropping pheromone until one of the following happens: it reaches the end of its lifetime, or it switches to state **Horiz**.
- During its first transition, every agent switches to one of the states **Horiz**, **Vert-seek**, or **Vert-bypass**.

Based on these, we conclude that every agent has a signature as per Definition 7. The rest of the proof is devoted to showing that the first 4τ agents pick up distinct signatures, and thus they explore all nodes at distance τ or less from the nest. This is accomplished by a series of lemmas, whose proofs are omitted. We show first that agents A_1 and A_2 have signatures $[0; \tau]$ and $[0; -\tau]$, respectively, and that subsequent agents are alternately dispatched to the North and to the South half-planes. Then, the following two technical lemmas, whose proofs are omitted, describe completely the state transitions of agents on the vertical axis:

► **Lemma 9.** *For all odd i with $3 \leq i \leq 4\tau - 1$, and for all y with $1 \leq y \leq \lceil \frac{i-1}{4} \rceil$, A_i is at node $(0, y)$ at the beginning of round $i + y$ and:*

1. *It senses pheromone $\mu \div (i - 4y)$ to the East and $\mu \div (i - 2 - 4y)$ to the West.*
2. *If $y = 1$, it is in state **Vert-seek** if $i - 3 \geq \mu$, otherwise it is in state **Vert-bypass**.*
3. *If $y \geq 2$, it is in state **Vert-seek** if $i - 4y + 2 \geq \mu$, otherwise it is in state **Vert-bypass**.*

► **Lemma 10.** *For all even i with $4 \leq i \leq 4\tau$, and for all y with $1 \leq y \leq \lceil \frac{i-2}{4} \rceil$, A_i is at node $(0, -y)$ at the beginning of round $i + y$ and:*

1. *It senses pheromone $\mu \div (i - 4y - 1)$ to the East and $\mu \div (i - 3 - 4y)$ to the West.*
2. *If $y = 1$, it is in state **Vert-seek** if $i - 4 \geq \mu$, otherwise it is in state **Vert-bypass**.*
3. *If $y \geq 2$, it is in state **Vert-seek** if $i - 4y + 1 \geq \mu$, otherwise it is in state **Vert-bypass**.*

From Lemmas 9 and 10, we deduce the signatures of the first 4τ agents and conclude the proof of Theorem 8.

■ **Algorithm 1** A treasure hunt algorithm for $\tau \geq 1, \mu \geq 2$.

Variables

state $\in \{\text{Init}, \text{Vert-seek}, \text{Vert-bypass}, \text{Horiz}\}$	▷ Initial value: Init
dir $\in \{N, E, S, W, \perp\}$	▷ Initial value: \perp
drop $\in \{\text{before}, \text{after}, \perp\}$	▷ Initial value: \perp

Transition function

```

1: if state = Init then
2:   if  $\varphi_N = \varphi_W = \varphi_S = \varphi_E = 0$  then
3:     state  $\leftarrow$  Horiz; dir  $\leftarrow$  E; drop  $\leftarrow$  after
4:   else if  $\varphi_E > \varphi_W$  then
5:     state  $\leftarrow$  Horiz; dir  $\leftarrow$  W; drop  $\leftarrow$  after
6:   else
7:     state  $\leftarrow$  Vert-bypass if  $\varphi_W > 0$  else Vert-seek
8:     dir  $\leftarrow$  S if  $\varphi_N > \varphi_S$  else N
9:     drop  $\leftarrow$  after
10:  end if
11: else if state = Vert-seek then
12:   if  $\varphi_W = \varphi_E = 0$  then
13:     state  $\leftarrow$  Vert-seek; drop  $\leftarrow$   $\perp$ 
14:   else
15:     INTERPRET-SIGNALS
16:   end if
17: else if state = Vert-bypass then
18:   if  $\varphi_W = \varphi_E = 0$  then
19:     state  $\leftarrow$  Horiz; dir  $\leftarrow$  E; drop  $\leftarrow$  after
20:   else
21:     INTERPRET-SIGNALS
22:   end if
23: else if state = Horiz then
24:   drop  $\leftarrow$   $\perp$ 
25: end if

26: procedure INTERPRET-SIGNALS
27:   if  $\varphi_W = 0$  and  $\varphi_E > 0$  then
28:     state  $\leftarrow$  Horiz; dir  $\leftarrow$  W; drop  $\leftarrow$  after
29:   else if  $\varphi_W > 0$  then
30:     state  $\leftarrow$  Vert-bypass; drop  $\leftarrow$   $\perp$ 
31:   end if
32: end procedure

```

3.3 Complexity

Recall the definitions of \mathbb{B}_x (ball of radius x around the nest) and \mathbb{L}_x (layer of nodes at distance x from the nest) from Section 2.

► **Theorem 11.** *If the treasure is located at distance at most d , where $1 \leq d \leq \tau$, then Algorithm 1 locates the treasure in time at most $5d - 1$.*

8:10 Treasure Hunt with Volatile Pheromones

Proof. From Lemmas 9 and 10, it follows that agents A_1, \dots, A_{4d} have all possible signatures with vertical component at most d (in absolute value). Moreover, since the distance of every agent from the nest strictly increases in each round, each agent A_i reaches distance d from the nest in round $i + d - 1$. It follows that, by the time agent A_{4d} reaches distance d from the nest, hence by round $5d - 1$, all nodes at distance d or less from the nest have been explored. ◀

► **Theorem 12.** *If the treasure is located at distance $d = \tau$, then any treasure hunt algorithm needs at least $5\tau - 1$ rounds to locate the treasure in the worst case.*

Proof. A given agent can explore at most one node at distance τ within its lifetime. Since \mathbb{L}_τ contains 4τ nodes, a correct algorithm must create at least 4τ agents, the last of which reaches distance τ in round $4\tau + \tau - 1 = 5\tau - 1$. It follows that, in the worst case, the treasure cannot be located before round $5\tau - 1$. ◀

► **Theorem 13.** *Let \mathcal{A} be any treasure hunt algorithm that is correct for a pair of parameters (τ, μ) . For every $d \leq \tau$, \mathcal{A} needs at least $\sqrt{5}d$ rounds to explore all nodes up to distance d .*

Proof. Fix a $d \leq \tau$ and let T be the first round at the end of which \mathcal{A} explores all nodes up to distance d . Clearly, $T \geq d$ because otherwise no agent can reach any node at distance d . We also assume that $T < \sqrt{5}d$, and we will show a contradiction.

Consider \mathbb{B}_x , for $x \leq d$ to be determined below. Among the agents A_1, \dots, A_T , those with $i \geq T - x + 1$ have moved at most x times by the end of round T , therefore they are unable to explore any node outside of \mathbb{B}_x . For every $i \leq T - x$, agent A_i moves at most $T - i + 1$ times by the end of round T , and it needs at least x moves before it can exit \mathbb{B}_x . Therefore, A_i explores at most $T - i + 1 - x$ nodes outside of \mathbb{B}_x . Summing over all agents with $i \leq T - x$ and taking also into account \mathbb{B}_x itself, we conclude that, by the end of round T , algorithm \mathcal{A} can explore at most

$$|\mathbb{B}_x| + \sum_{i=1}^{T-x} (T - i + 1 - x) = 2x^2 + 2x + 1 + \frac{(T-x)(T-x+1)}{2}$$

nodes. The above expression is minimized for $x = \frac{T}{5} - \frac{3}{10} < d$, whence we obtain that \mathcal{A} explores at most

$$\frac{2T^2}{5} + \frac{4T}{5} + \frac{31}{40}$$

nodes. By definition of T , at that round \mathcal{A} has explored at least \mathbb{B}_d , therefore:

$$\frac{2T^2}{5} + \frac{4T}{5} + \frac{31}{40} \geq 2d^2 + 2d + 1$$

whence it follows that $T > \sqrt{5}d$, a contradiction. ◀

► **Theorem 14.** *Algorithm 1 effectively uses 4τ agents, and that is optimal.*

Proof (sketch). By Lemmas 9 and 10, agents $A_1, \dots, A_{4\tau}$ have all possible signatures with vertical component at most τ (in absolute value). Moreover, it can be shown that every agent A_i only senses pheromone left by some earlier agent A_j , $j < i$ (details omitted). It follows that, even if no agents are generated after round 4τ , the above agents $A_1, \dots, A_{4\tau}$ will still perform the same trajectories and explore all nodes up to distance τ . Therefore,

the effective agent utilization of Algorithm 1 is 4τ . This is optimal because there exist 4τ nodes at distance τ , and an agent can only visit at most one node at distance τ during its lifetime. ◀

► **Theorem 15.** *The pheromone utilization of Algorithm 1 is at most $O(d)$, and this is asymptotically optimal.*

Proof. By Theorem 11, the treasure is located in time $O(d)$, and each of the $O(d)$ agents that are created until then drops pheromone at most 2 times: once when it leaves the nest, and once if and when it leaves the vertical axis. Hence, the pheromone utilization of Algorithm 1 is $O(d)$.

To prove optimality, consider a treasure hunt algorithm \mathcal{A} that uses asymptotically less than d pheromone, i.e., its pheromone utilization is bounded by some function $f(d)$ such that $\lim_{d \rightarrow \infty} \frac{f(d)}{d} = 0$. Let N be the number of states of the FSM \mathcal{A} .

By our assumption on $f(d)$, for every $\varepsilon > 0$ there exists a d_ε such that for all $d > d_\varepsilon$, $f(d) < \varepsilon \cdot d$. Let us fix, then, a $d_0 > N + 1$ such that $f(d_0) < \frac{d_0}{N+1}$. Moreover, it is well known and has been observed several times in the literature (see, e.g., [39, Lemma 5]) that a deterministic FSM that moves in a grid and does not interact with its environment can explore at most a constant-width band, infinite in one direction. Let W be the constant that bounds the number of nodes of any particular layer that are visited by such an agent.

Now, consider the execution of \mathcal{A} in a system with parameters (τ, μ) , where $\tau \geq WNd_0 + 1$ and the treasure is located at distance d_0 . The number of layers on which at least one agent drops pheromone is clearly bounded by the pheromone utilization of \mathcal{A} , and hence by $f(d_0) < \frac{d_0}{N+1}$. It follows that there exists at least one layer $d_1 \leq d_0 - (N + 1)$, such that no agent drops pheromone on any of the layers $d_1, d_1 + 1, \dots, d_1 + N$. Therefore, any agent that arrives at layer d_0 is already repeating a sequence of states during which it drops no pheromone.

Consider, now, the execution of \mathcal{A} in the same system but with the treasure placed at distance $d^* = WNd_0 + 1$. As agents do not perceive the presence of treasure, they will behave as in the previous case. In particular, even though there is an infinite number of agents coming out of layer d_0 , their trajectories are contained in at most most $4d_0 \cdot N$ distinct bands of constant width W , infinite in one direction. This is because the trajectory of an agent that is coming out of d_0 is completely determined by the node from which it exits layer d_0 and the state in which it leaves the layer.

It follows that algorithm \mathcal{A} explores at most $4WNd_0$ nodes of layer d^* , but layer d^* contains $4d^* > 4WNd_0$ nodes. Therefore, the adversary can place the treasure at a node that will not be explored by \mathcal{A} . ◀

4 A treasure hunt algorithm for $\mu \geq 1$ and $\tau \geq 16$

Similarly to Algorithm 1, the algorithm that we present in this section creates agents of all possible signatures $[v; h]$, as per Definition 7.

The main difficulty here is that the dropped pheromone can evaporate in one round only, in the case of $\mu = 1$. To explore a grid up to an unknown distance τ , where τ is also the lifetime of an agent, every node at distance τ has to be visited by at least one agent (Definition 6). This agent cannot stop even for one round and has to follow a shortest path to the node at distance τ . At the same time, agents have to be sent alternatively exploring each half of the grid (north and south, in our case), and so the shortest time interval between two

following agents (moving to the positions to explore) is two rounds. This makes it difficult to solve the problem with pheromone evaporating in one round. It disappears too fast to provide any information to the next arriving agent.

In order to overcome this challenge, we use in Algorithm 2 two types of agents: *signaling* agents, that stop moving at key positions and start dropping pheromone according to some predetermined pattern, and *explorer* agents, that read these patterns on their way to the extreme grid positions without stopping even for a single round. Since signaling consumes rounds from the lifetime of signaling agents, these agents must stop at a sufficient distance away from the extreme positions, to still have enough lifetime to signal the required pattern. This distance is expressed by the parameter s of our algorithm. This also has an impact on the minimum agent lifetime that is required for the algorithm to operate correctly, as the furthest signaling agents must have enough lifetime to reach their signaling positions and complete the required pattern. Algorithm 2 works for all values of μ , but only for $\tau \geq 16$.

Binary word notations. Let us define some finite binary word notations that we will use in order to present the algorithm. The empty word is denoted by ϵ and the length of a word w by $|w|$. For any word w and integer $j \in \{1, \dots, |w|\}$, $w[j]$ denotes the j^{th} most significant bit of w . Let $\text{shiftright}(w)$ return a word obtained by removing the most significant bit ($w[1]$) from w .

We now present Algorithm 2 by referring to the pseudocode that we give in this section. All proofs are omitted. Algorithm 2 uses a constant number of special states, as follows: Given a binary word w of length at most 9, $\text{Pattern}(w)$ is the first of a sequence of $|w|$ states, during which the agent stays on the same node and drops (or not) pheromone according to the bit pattern w . $\text{Forward}(s)\text{-Explore}(E)$ (resp. $\text{Forward}(s)\text{-Explore}(W)$) is the first of a sequence of states during which the agent moves s steps forward (north or south, in the same direction as it was moving before entering this state) and then turns east (resp. west) and keeps moving in that direction until the end of its lifetime.

In the main part of the algorithm, agents leave the nest alternatively moving either north or south, on the vertical axis, until arriving s steps away from a non-explored yet line where they either stop for signaling (moving to state $\text{Pattern}(w)$) or continue moving to reach this non-explored yet line to turn there either east (state $\text{Forward}(s)\text{-Explore}(E)$) or west (state $\text{Forward}(s)\text{-Explore}(W)$) for exploring each half of the line, s steps away.

Such a signaling, for exploring each next line at distance h , is achieved by using three agents. One is placed s lines before, and at one cell east from the vertical axis, i.e. at $(1, h - s)$ if heading north (resp. $(1, -(h - s))$, if heading south). The second one is also s lines before, but at one cell west from the vertical axis, i.e. at $(-1, h - s)$ (resp. $(-1, -(h - s))$). The third agent, is at $(0, h - s + 1)$ (resp. $(0, -(h - s + 1))$). A newly arrived agent (at $(0, h - s)$ (resp. $(0, -(h - s))$)) senses the pheromone dropped by these three agents and performs actions according to the parsing of the sensed pattern. This part of the algorithm is controlled mainly by the $\text{INTERPRET-SIGNALS-PHASE2}()$ procedure (see Alg. 2).

Operating in this way, with agents “jumping” each time s steps vertically, for exploring a line there, leaves at least the s first horizontal lines of each half of the grid unexplored. Hence, we need a special procedure for exploring these lines. For that, up to horizontal lines at distance s , the signaling agents stay longer for guiding some of the incoming agents to explore these lines (some other agents are still guided to “jump” for exploring the lines s steps further). This part of the algorithm is controlled mainly by the $\text{INTERPRET-SIGNALS-PHASE1}()$ procedure (see Alg. 2).

■ **Algorithm 2** A treasure hunt algorithm for $\tau \geq 16, \mu \geq 1$ and $s = 6$.

Variables

state \in {Init, Vert-seek, Vert-bypass, Horiz, Pattern(w), Forward(k)-Explore(E)
 Forward(k)-Explore(W)}, $w \in \{0, 1\}^9, k \in [0, s]$ ▷ Initial value: Init
 dir \in {N, E, S, W, \perp } ▷ Initial value: \perp
 drop \in {before, after, \perp } ▷ Initial value: \perp
 moves \in [0, $s + 1$] ▷ Initial value: 0

Transition function

```

1: if state = Init then
2:   if  $\varphi_N = \varphi_W = \varphi_S = \varphi_E = 0$  then
3:     state  $\leftarrow$  Pattern(11001); dir  $\leftarrow$  E; drop  $\leftarrow$  after ▷ Start signaling E
4:   else if  $\varphi_N = \varphi_W = \varphi_S = 0$  and  $\varphi_E > 0$  then
5:     state  $\leftarrow$  Pattern(111); dir  $\leftarrow$  W; drop  $\leftarrow$  after ▷ Start signaling W
6:   else if  $\varphi_N = \varphi_S = 0$  and  $\varphi_W = \varphi_E > 0$  then
7:     state  $\leftarrow$  Pattern(01); dir  $\leftarrow$  N; drop  $\leftarrow$  after ▷ Start signaling N
8:   else if  $\varphi_S = 0$  and  $\varphi_N = \varphi_W = \varphi_E > 0$  then
9:     state  $\leftarrow$  Horiz; dir  $\leftarrow$  E; drop  $\leftarrow$   $\perp$  ▷ Explore E
10:  else if  $\varphi_S = 0$  and  $\varphi_N = \varphi_E < \varphi_W$  then
11:    state  $\leftarrow$  Horiz; dir  $\leftarrow$  W; drop  $\leftarrow$   $\perp$  ▷ Explore W
12:  else if  $\varphi_S = 0$  and  $\varphi_N = \varphi_W > \varphi_E$  then
13:    state  $\leftarrow$  Vert-bypass; dir  $\leftarrow$  N; drop  $\leftarrow$  after ▷ Go signaling E on line (0, 1)
14:  else if  $\varphi_S = 0$  and  $\varphi_N = \varphi_E > \varphi_W$  then
15:    state  $\leftarrow$  Vert-bypass; dir  $\leftarrow$  S; drop  $\leftarrow$  after ▷ Go signaling E on line (0, -1)
16:  else if  $\varphi_S > \varphi_N$  and  $\varphi_S > \varphi_E$  and  $\varphi_S > \varphi_W$  then
17:    state  $\leftarrow$  Vert-seek; dir  $\leftarrow$  N; drop  $\leftarrow$  after ▷ Go signaling W on line (0, 1)
18:  else if  $\varphi_N > \varphi_S$  and  $\varphi_N > \varphi_E$  and  $\varphi_N > \varphi_W$  then
19:    state  $\leftarrow$  Vert-seek; dir  $\leftarrow$  S; drop  $\leftarrow$  after ▷ Go signaling W on line (0, -1)
20:  end if
21: else if state = Vert-seek then
22:   if  $\varphi_{dir} = \varphi_E = \varphi_W = 0$  then
23:     drop  $\leftarrow$   $\perp$  ▷ keep searching
24:   else if moves  $< s + 1$  then
25:     state  $\leftarrow$  Vert-bypass; INTERPRET-SIGNALS-PHASE1
26:   else
27:     state  $\leftarrow$  Vert-bypass; INTERPRET-SIGNALS-PHASE2
28:   end if
29: else if state = Vert-bypass then
30:   if moves  $< s + 1$  then
31:     INTERPRET-SIGNALS-PHASE1
32:   else
33:     INTERPRET-SIGNALS-PHASE2
34:   end if
35: else if state = Horiz then
36:   drop  $\leftarrow$   $\perp$ 
37: else if state = Forward( $k$ )-Explore( $d$ ) then
38:   if  $k > 1$  then
39:     state  $\leftarrow$  Forward( $k - 1$ )-Explore( $d$ )
40:   else
41:     state  $\leftarrow$  Horiz; dir  $\leftarrow$   $d$ 
42:   end if
43: drop  $\leftarrow$   $\perp$ 

```

8:14 Treasure Hunt with Volatile Pheromones

► Algorithm 2 (continued)

```

44: else if state = Pattern( $w$ ) then
45:   dir  $\leftarrow \perp$ 
46:   if  $|w| > 1 \wedge w[1] = 1$  then                                ▷  $w[1]$  returns the first bit of the binary word  $w$ 
47:     drop  $\leftarrow$  after
48:   else
49:     drop  $\leftarrow \perp$ 
50:   end if
51:   if  $w \neq \epsilon$  then
52:     state  $\leftarrow$  Pattern(shiftleft( $w$ ))                            ▷ shiftleft( $w$ ) removes the first bit of  $w$ 
53:   end if
54: end if
55: if dir  $\neq \perp \wedge$  moves  $< s + 1$  then
56:   moves  $\leftarrow$  moves + 1
57: end if
58: procedure INTERPRET-SIGNALS-PHASE1
59:   if  $\varphi_{\text{dir}} = \varphi_E = \varphi_W = 0$  then
60:     state  $\leftarrow$  Pattern(1 01 01 01); dir  $\leftarrow E$ ; drop  $\leftarrow \perp$                                 ▷ Start signaling E
61:   else if  $\varphi_{\text{dir}} = \varphi_W = 0$  and  $\varphi_E > 0$  then
62:     state  $\leftarrow$  Pattern(1 01 00 01 01); dir  $\leftarrow W$ ; drop  $\leftarrow \perp$                                 ▷ Start signaling W
63:   else if  $\varphi_{\text{dir}} = 0$  and  $\varphi_E = \varphi_W > 0$  then
64:     state  $\leftarrow$  Pattern(1 01 00 01 01); drop  $\leftarrow \perp$                                 ▷ Start signaling N or S
65:   else if  $\varphi_{\text{dir}} = \varphi_E = \varphi_W > 0$  then
66:     state  $\leftarrow$  Horiz; dir  $\leftarrow E$ ; drop  $\leftarrow \perp$                                 ▷ Explore E
67:   else if  $\varphi_{\text{dir}} = \varphi_E > \varphi_W$  then
68:     state  $\leftarrow$  Horiz; dir  $\leftarrow W$ ; drop  $\leftarrow \perp$                                 ▷ Explore W
69:   else if  $\varphi_{\text{dir}} = \varphi_E < \varphi_W$  then
70:     state  $\leftarrow$  Forward( $s$ )-Explore( $E$ ); drop  $\leftarrow \perp$                                 ▷ Move forward  $s$  steps and explore E
71:   else if  $\varphi_{\text{dir}} = \varphi_W > \varphi_E$  then
72:     state  $\leftarrow$  Forward( $s$ )-Explore( $W$ ); drop  $\leftarrow \perp$                                 ▷ Move forward  $s$  steps and explore W
73:   else if  $\varphi_{\text{dir}} > \varphi_E$  and  $\varphi_{\text{dir}} > \varphi_W$  then
74:     drop  $\leftarrow \perp$                                 ▷ Continue to bypass pheromone traces
75:   end if
76: end procedure
77: procedure INTERPRET-SIGNALS-PHASE2
78:   if  $\varphi_{\text{dir}} = \varphi_E = \varphi_W = 0$  then
79:     state  $\leftarrow$  Pattern(1 01 01); dir  $\leftarrow E$ ; drop  $\leftarrow \perp$                                 ▷ Start signaling E
80:   else if  $\varphi_{\text{dir}} = \varphi_W = 0$  and  $\varphi_E > 0$  then
81:     state  $\leftarrow$  Pattern(1 01 01); dir  $\leftarrow W$ ; drop  $\leftarrow \perp$                                 ▷ Start signaling W
82:   else if  $\varphi_{\text{dir}} = 0$  and  $\varphi_E = \varphi_W > 0$  then
83:     state  $\leftarrow$  Pattern(1 01 01); drop  $\leftarrow \perp$                                 ▷ Start signaling N or S
84:   else if  $\varphi_{\text{dir}} = \varphi_E = \varphi_W > 0$  then
85:     state  $\leftarrow$  Forward( $s$ )-Explore( $E$ ); drop  $\leftarrow \perp$                                 ▷ Move forward  $s$  steps and explore E
86:   else if  $\varphi_{\text{dir}} = \varphi_W > \varphi_E$  then
87:     state  $\leftarrow$  Forward( $s$ )-Explore( $W$ ); drop  $\leftarrow \perp$                                 ▷ Move forward  $s$  steps and explore W
88:   else if  $\varphi_{\text{dir}} > \varphi_E$  and  $\varphi_{\text{dir}} > \varphi_W$  then
89:     drop  $\leftarrow \perp$                                 ▷ Continue to bypass pheromone traces
90:   end if
91: end procedure

```

► **Remark 16.** The technical analysis shows that the algorithm works with $s = 6$. We give a short intuition for this value. As briefly explained above, the point is that signaling patterns cannot be established too far from the nest, because agents do not have enough remaining lifetime to complete the pattern. As such, the exploration of horizontal lines that are far from the nest must be signaled by patterns that are set up closer to the nest. In fact, s depends on the longest such signaling pattern, dropped by a signaling agent (at a distance further than s from the nest). This in turn establishes the closest position of such agent to the grid extremity (at distance τ), where it can complete the signaling before it dies. In our algorithm, to explore lines after distance s , only 6 rounds are used by a signaling agent, which explains why $s = 6$. We actually need to encode 6 actions (3 for signaling agents and 3 for exploring). This requires 12 rounds of signaling due to the N/S dispatching at the nest, but we can get away with s being only 6 because the agents arrive at different times. Still, signaling agents have to stay alive there only for 6 rounds each.

Regarding the minimal τ which is 16, it is due to the transition from operation in the s first lines to the next ones. During this transition, signaling agents should have enough remaining lifetime to reach line s and to signal the required pattern (in these lines, the signaling pattern of each agent requires 10 rounds; there are $3 + 5$ actions to signal here). So 10 rounds for signaling and 6 rounds to reach the line at distance 6 gives $\tau \geq 16$ rounds.

Let us detail now the operation of the algorithm during the first rounds intended to explore the x -axis (this differs from the exploration of other lines). Agents start at the nest in state `Init`. Each of the first three agents are placed respectively east, west and north to the nest and start signaling according to the predetermined pattern (lines 3, 5 and 7, Alg. 2). This signaling instructs the 4th agent (A_4) to explore the east half of the x -axis (line 9) and the 5th agent (A_5), to explore the remaining (west) half of the x -axis (line 11). The next four agents are instructed to move to lines $(0, 1)$ and $(0, -1)$ (lines 13 - 19), two agents on each line, to stop on the East and West from the vertical axis (cells $(1, 1)$, $(-1, 1)$, $(1, -1)$, $(-1, -1)$). This is for instructing to explore lines $(0, 1)$, $(0, -1)$, $(0, s + 1)$ and $(0, -s - 1)$ (as explained in the previous paragraph).

Notice that starting from round 8, every even round, an agent in `Vert-see` leaves the nest to the North, and every odd round, an agent in `Vert-see` leaves the nest to the South. This alternation allows to explore both the north and the south halves of the grid, without knowing its size.

States `Vert-see` and `Vert-bypass` are used in a similar way as in the previous algorithm, to overcome the difficulty caused by the pheromone traces left from previous drops in case of $\mu > 1$. An agent has to bypass (in state `Vert-bypass`) these traces (lines 74 and 89) until arriving to a line with either no pheromone or with “fresh” pheromones, just dropped in the previous round (treated in all other lines of the `INTERPRET-SIGNALS-PHASE1()` and `INTERPRET-SIGNALS-PHASE2()` procedures). Starting with the 8th agent, agents leave the nest in state `Vert-see` and move vertically in this state until sensing some dropped pheromone, moving then to `Vert-bypass` (lines 22 - 27).

► **Theorem 17.** *Algorithm 2 solves the treasure hunt problem for $\mu \geq 1$, $\tau \geq 16$ and $s = 6$ in $11\tau - 6s + 2$ rounds, using $10\tau - 6s + 3$ agents and $28\tau + O(s) + 8$ pheromone drops.*

5 Concluding remarks

We have presented the first algorithms for the treasure hunt problem under the weak communication mechanism of evaporating pheromone markers. In Algorithm 1, the assumption that pheromone lasts for at least two rounds ($\mu \geq 2$) leads to a fairly simple algorithm design with very few states. By contrast, Algorithm 2 is significantly more complicated, as it needs to be able to handle both an extremely fast evaporation rate ($\mu = 1$) and larger values of μ .

Algorithm 2 covers all values of the evaporation parameter $\mu \geq 1$, but it requires a lifetime of $\tau \geq 16$. It would be interesting to determine the smallest τ_0 such that there exists a treasure hunt algorithm that works for all $\mu \geq 1$ and for all $\tau \geq \tau_0$. With ad-hoc arguments, it can be seen that $\tau_0 > 2$. However, it is far from obvious how to generalize these arguments to larger values of τ_0 . On the other hand, there may be room to improve the upper bound of $\tau_0 \leq 16$, with some fine-tuning of the signaling patterns.

Another interesting direction for future work is improving on the complexities of Algorithm 2, or studying tradeoffs between completion time, pheromone utilization, and agent utilization. Since both of our algorithms use only a constant number of pheromone drops per agent, one idea would be to increase the frequency of pheromone drops. It seems that this would not help to reduce agent utilization or the completion time. Indeed, the limiting factor in Algorithm 2 seems to be not the amount of pheromone that is dropped or that might be dropped, but indeed the number of grid positions that are available in order to set up an efficient pattern, i.e., a pattern that resides in the neighborhood of the main axis so that it can be immediately sensed by agents.

The assumption of detecting pheromones only in adjacent nodes to the agent, although natural, could be relaxed. However, if the sensing range is increased even to 2 while maintaining the principle that the agent can pinpoint exactly the position of the pheromone and compare pheromone levels between all nodes in its 2-neighborhood, then Algorithm 1 resolves the problem for all values of the parameters. Indeed, the only reason why Algorithm 1 fails for $\mu = 1$ is that, due to the North-South dispatching at the nest, agents are dispatched into the same half-plane every two rounds, and therefore any pheromone dropped by an agent evaporates before the next agent can sense it. Consequently, in order to study a meaningful problem with an increased sensing range, some loss of information would have to be introduced at distance 2 or more.

Our algorithms are quite far from modeling natural ant foraging patterns. Indeed, depending on species, ants in nature tend to employ a wide range of communication methods, including multiple types of pheromone of various degrees of volatility, repellent pheromones, contact, or sounds [49]. However, our proposed solutions are more appropriate for artificial agent systems, where the parameter τ might correspond to agents with limited energy, and evaporating markers could be useful to prevent area pollution. The appropriate parameter values will depend on the specific application.

As a general remark, we believe that the communication model of evaporating pheromone markers is inherently interesting and we would like to study other agent coordination problems in this model. Orthogonally, one may consider less predictable evaporation mechanisms, such as evaporation governed by a random process, or controlled by an adversary.

References

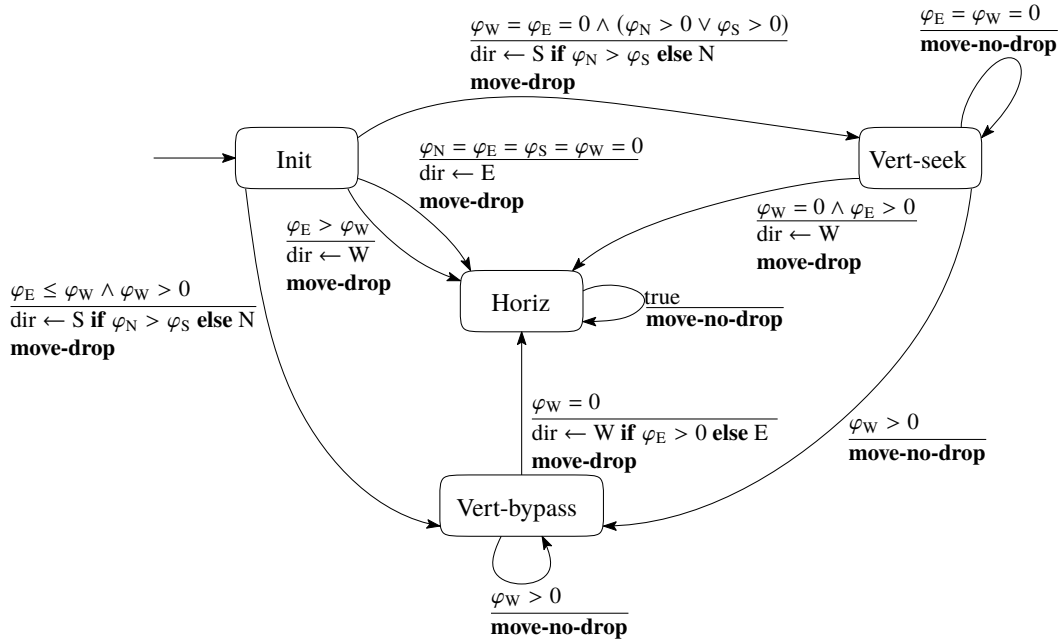
- 1 Yehuda Afek, Roman Kecher, and Moshe Sulamy. Optimal pheromone utilization. In *2nd Workshop on Biological Distributed Algorithms (BDA)*, 2014.
- 2 Yehuda Afek, Roman Kecher, and Moshe Sulamy. Optimal and resilient pheromone utilization in ant foraging. *CoRR*, abs/1507.00772, 2015. [arXiv:1507.00772](https://arxiv.org/abs/1507.00772).
- 3 Abhinav Aggarwal and Jared Saia. ANTS on a plane. In Andrea Werneck Richa and Christian Scheideler, editors, *Structural Information and Communication Complexity - 27th International Colloquium, SIROCCO 2020, Paderborn, Germany, June 29 - July 1, 2020, Proceedings*, volume 12156 of *Lecture Notes in Computer Science*, pages 47–62. Springer, 2020. doi:10.1007/978-3-030-54921-3_3.
- 4 Noga Alon, Chen Avin, Michal Koucký, Gady Kozma, Zvi Lotker, and Mark R. Tuttle. Many random walks are faster than one. *Combinatorics, Probability and Computing*, 20(4):481–502, 2011. doi:10.1017/S0963548311000125.
- 5 Steve Alpern, Robbert Fokkink, Leszek Gąsieniec, Roy Lindelauf, and V.S. Subrahmanian, editors. *Search Theory: A Game Theoretic Perspective*. Springer New York, NY, 2013. doi:10.1007/978-1-4614-6825-7.
- 6 Steve Alpern and Shmuel Gal. *The Theory of Search Games and Rendezvous*. International Series in Operations Research & Management Science. Springer New York, NY, 2003. doi:10.1007/b100809.
- 7 Spyros Angelopoulos. Online search with a hint. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPICs*, pages 51:1–51:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ITCS.2021.51.
- 8 Spyros Angelopoulos, Diogo Arsénio, and Christoph Dürr. Infinite linear programming and online searching with turn cost. *Theoretical Computer Science*, 670:11–22, 2017. doi:10.1016/j.tcs.2017.01.013.
- 9 Spyros Angelopoulos, Alejandro López-Ortiz, and Konstantinos Panagiotou. Multi-target ray searching problems. *Theoretical Computer Science*, 540:2–12, 2014. doi:10.1016/j.tcs.2014.03.028.
- 10 Spyros Angelopoulos and Malachi Voss. Online search with maximum clearance. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 3642–3650. AAAI Press, 2021. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/16480>.
- 11 Ricardo A. Baeza-Yates and René Schott. Parallel searching in the plane. *Computational Geometry*, 5:143–154, 1995. doi:10.1016/0925-7721(95)00003-R.
- 12 Evangelos Bampas, Jurek Czyzowicz, Leszek Gąsieniec, David Ilcinkas, Ralf Klasing, Tomasz Kociumaka, and Dominik Pajak. Linear search by a pair of distinct-speed robots. *Algorithmica*, 81(1):317–342, 2019. doi:10.1007/s00453-018-0447-0.
- 13 Evangelos Bampas, Jurek Czyzowicz, David Ilcinkas, and Ralf Klasing. Beachcombing on strips and islands. *Theoretical Computer Science*, 806:236–255, 2020. doi:10.1016/j.tcs.2019.04.001.
- 14 Anatole Beck. On the linear search problem. *Israel Journal of Mathematics*, 2:221–228, 1964. doi:10.1007/BF02759737.
- 15 Anatole Beck and Donald J. Newman. Yet more on the linear search problem. *Israel Journal of Mathematics*, 8:419–429, 1970. doi:10.1007/BF02798690.
- 16 Lucas Boczkowski, Uriel Feige, Amos Korman, and Yoav Rodeh. Navigating in trees with permanently noisy advice. *ACM Transactions on Algorithms*, 17(2):15:1–15:27, 2021. doi:10.1145/3448305.

- 17 Anthony Bonato and Richard J. Nowakowski. *The Game of Cops and Robbers on Graphs*. American Mathematical Society, 2011.
- 18 Sébastien Bouchard, Yoann Dieudonné, Andrzej Pelc, and Franck Petit. Deterministic treasure hunt in the plane with angular hints. *Algorithmica*, 82(11):3250–3281, 2020. doi:10.1007/s00453-020-00724-4.
- 19 Sébastien Bouchard, Arnaud Labourel, and Andrzej Pelc. Impact of knowledge on the cost of treasure hunt in trees. *Networks*, 80(1):51–62, 2022. doi:10.1002/net.22075.
- 20 Sebastian Brandt, Klaus-Tycho Foerster, Benjamin Richner, and Roger Wattenhofer. Wireless evacuation on m rays with k searchers. *Theoretical Computer Science*, 811:56–69, 2020. doi:10.1016/j.tcs.2018.10.032.
- 21 Sebastian Brandt, Jara Uitto, and Roger Wattenhofer. A tight lower bound for semi-synchronous collaborative grid exploration. *Distributed Computing*, 33(6):471–484, 2020. doi:10.1007/s00446-020-00369-0.
- 22 Marek Chrobak, Leszek Gasieniec, Thomas Gorry, and Russell Martin. Group search on the line. In Giuseppe F. Italiano, Tiziana Margaria-Steffen, Jaroslav Pokorný, Jean-Jacques Quisquater, and Roger Wattenhofer, editors, *SOFSEM 2015: Theory and Practice of Computer Science - 41st International Conference on Current Trends in Theory and Practice of Computer Science, Pec pod Sněžkou, Czech Republic, January 24-29, 2015. Proceedings*, volume 8939 of *Lecture Notes in Computer Science*, pages 164–176. Springer, 2015. doi:10.1007/978-3-662-46078-8_14.
- 23 Timothy H. Chung, Geoffrey A. Hollinger, and Volkan Isler. Search and pursuit-evasion in mobile robotics - A survey. *Autonomous Robots*, 31(4):299–316, 2011. doi:10.1007/s10514-011-9241-4.
- 24 Andrea E. F. Clementi, Francesco D’Amore, George Giakkoupis, and Emanuele Natale. Search via parallel Lévy walks on \mathbb{Z}^2 . In Avery Miller, Keren Censor-Hillel, and Janne H. Korhonen, editors, *PODC ’21: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, July 26-30, 2021*, pages 81–91. ACM, 2021. doi:10.1145/3465084.3467921.
- 25 Lihi Cohen, Yuval Emek, Oren Louidor, and Jara Uitto. Exploring an infinite space with finite memory scouts. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 207–224. SIAM, 2017. doi:10.1137/1.9781611974782.14.
- 26 Jared Ray Coleman and Oscar Morales-Ponce. The snow plow problem: Perpetual maintenance by mobile agents on the line. In *24th International Conference on Distributed Computing and Networking, ICDCN 2023, Kharagpur, India, January 4-7, 2023*, pages 110–114. ACM, 2023. doi:10.1145/3571306.3571396.
- 27 Andrew Collins, Jurek Czyzowicz, Leszek Gasieniec, Adrian Kosowski, Evangelos Kranakis, Danny Krizanc, Russell Martin, and Oscar Morales-Ponce. Optimal patrolling of fragmented boundaries. In Guy E. Blelloch and Berthold Vöcking, editors, *25th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA ’13, Montreal, QC, Canada - July 23 - 25, 2013*, pages 241–250. ACM, 2013. doi:10.1145/2486159.2486176.
- 28 Colin Cooper, Alan M. Frieze, and Tomasz Radzik. Multiple random walks in random regular graphs. *SIAM Journal on Discrete Mathematics*, 23(4):1738–1761, 2009. doi:10.1137/080729542.
- 29 Jurek Czyzowicz, Leszek Gasieniec, Konstantinos Georgiou, Evangelos Kranakis, and Fraser MacQuarrie. The beachcombers’ problem: Walking and searching with mobile robots. *Theoretical Computer Science*, 608:201–218, 2015. doi:10.1016/j.tcs.2015.09.011.
- 30 Jurek Czyzowicz, Kostantinos Georgiou, and Evangelos Kranakis. Group search and evacuation. In Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro, editors, *Distributed Computing by Mobile Entities, Current Research in Moving and Computing*, volume 11340 of *Lecture Notes in Computer Science*, pages 335–370. Springer, 2019. doi:10.1007/978-3-030-11072-7_14.

- 31 Francesco D'Amore. *On the collective behaviors of bio-inspired distributed systems. (Sur les comportements collectifs de systèmes distribués bio-inspirés)*. PhD thesis, Côte d'Azur University, Nice, France, 2022. URL: <https://tel.archives-ouvertes.fr/tel-03906167>.
- 32 Shantanu Das. Mobile agent rendezvous in a ring using faulty tokens. In Shrisha Rao, Mainak Chatterjee, Prasad Jayanti, C. Siva Ram Murthy, and Sanjoy Kumar Saha, editors, *Distributed Computing and Networking, 9th International Conference, ICDCN 2008, Kolkata, India, January 5-8, 2008*, volume 4904 of *Lecture Notes in Computer Science*, pages 292–297. Springer, 2008. doi:10.1007/978-3-540-77444-0_29.
- 33 Shantanu Das, Matús Mihalák, Rastislav Srámek, Elias Vicari, and Peter Widmayer. Rendezvous of mobile agents when tokens fail anytime. In Theodore P. Baker, Alain Bui, and Sébastien Tixeuil, editors, *Principles of Distributed Systems, 12th International Conference, OPODIS 2008, Luxor, Egypt, December 15-18, 2008. Proceedings*, volume 5401 of *Lecture Notes in Computer Science*, pages 463–480. Springer, 2008. doi:10.1007/978-3-540-92221-6_29.
- 34 Shantanu Das and Nicola Santoro. Moving and computing models: Agents. In Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro, editors, *Distributed Computing by Mobile Entities, Current Research in Moving and Computing*, volume 11340 of *Lecture Notes in Computer Science*, pages 15–34. Springer, 2019. doi:10.1007/978-3-030-11072-7_2.
- 35 Erik D. Demaine, Sándor P. Fekete, and Shmuel Gal. Online searching with turn cost. *Theoretical Computer Science*, 361(2-3):342–355, 2006. doi:10.1016/j.tcs.2006.05.018.
- 36 Yoann Dieudonné and Andrzej Pelc. Deterministic network exploration by a single agent with Byzantine tokens. *Information Processing Letters*, 112(12):467–470, 2012. doi:10.1016/j.ipl.2012.03.017.
- 37 Klim Efremenko and Omer Reingold. How well do random walks parallelize? In Irit Dinur, Klaus Jansen, Joseph Naor, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21-23, 2009. Proceedings*, volume 5687 of *Lecture Notes in Computer Science*, pages 476–489. Springer, 2009. doi:10.1007/978-3-642-03685-9_36.
- 38 Robert Elsässer and Thomas Sauerwald. Tight bounds for the cover time of multiple random walks. *Theoretical Computer Science*, 412(24):2623–2641, 2011. doi:10.1016/j.tcs.2010.08.010.
- 39 Yuval Emek, Tobias Langner, David Stolz, Jara Uitto, and Roger Wattenhofer. How many ants does it take to find the food? *Theoretical Computer Science*, 608:255–267, 2015. doi:10.1016/j.tcs.2015.05.054.
- 40 Yuval Emek, Tobias Langner, Jara Uitto, and Roger Wattenhofer. Solving the ANTS problem with asynchronous finite state machines. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II*, volume 8573 of *Lecture Notes in Computer Science*, pages 471–482. Springer, 2014. doi:10.1007/978-3-662-43951-7_40.
- 41 Ofer Feinerman and Amos Korman. Memory lower bounds for randomized collaborative search and implications for biology. In Marcos K. Aguilera, editor, *Distributed Computing - 26th International Symposium, DISC 2012, Salvador, Brazil, October 16-18, 2012. Proceedings*, volume 7611 of *Lecture Notes in Computer Science*, pages 61–75. Springer, 2012. doi:10.1007/978-3-642-33651-5_5.
- 42 Ofer Feinerman and Amos Korman. The ANTS problem. *Distributed Computing*, 30(3):149–168, 2017. doi:10.1007/s00446-016-0285-8.
- 43 Ofer Feinerman, Amos Korman, Zvi Lotker, and Jean-Sébastien Sereni. Collaborative search on the plane without communication. In Darek Kowalski and Alessandro Panconesi, editors, *ACM Symposium on Principles of Distributed Computing, PODC '12, Funchal, Madeira, Portugal, July 16-18, 2012*, pages 77–86. ACM, 2012. doi:10.1145/2332432.2332444.

- 44 Paola Flocchini, Evangelos Kranakis, Danny Krizanc, Flaminia L. Luccio, Nicola Santoro, and Cindy Sawchuk. Mobile agents rendezvous when tokens fail. In Rastislav Kralovic and Ondrej Sýkora, editors, *Structural Information and Communication Complexity, 11th International Colloquium, SIROCCO 2004, Smolenice Castle, Slovakia, June 21-23, 2004, Proceedings*, volume 3104 of *Lecture Notes in Computer Science*, pages 161–172. Springer, 2004. doi:10.1007/978-3-540-27796-5_15.
- 45 Fedor V. Fomin and Dimitrios M. Thilikos. An annotated bibliography on guaranteed graph searching. *Theoretical Computer Science*, 399(3):236–245, 2008. doi:10.1016/j.tcs.2008.02.040.
- 46 Pierre Fraigniaud, Amos Korman, and Yoav Rodeh. Parallel Bayesian search with no coordination. *Journal of the ACM*, 66(3):17:1–17:28, 2019. doi:10.1145/3304111.
- 47 Subir Kumar Ghosh and Rolf Klein. Online algorithms for searching and exploration in the plane. *Computer Science Review*, 4(4):189–201, 2010. doi:10.1016/j.cosrev.2010.05.001.
- 48 Andrej Ivaskovic, Adrian Kosowski, Dominik Pajak, and Thomas Sauerwald. Multiple random walks on paths and grids. In Heribert Vollmer and Brigitte Vallée, editors, *34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, 2017, Hannover, Germany*, volume 66 of *LIPICs*, pages 44:1–44:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.STACS.2017.44.
- 49 Duncan E. Jackson and Francis L.W. Ratnieks. Communication in ants. *Current Biology*, 16(15):R570–R574, 2006. doi:10.1016/j.cub.2006.07.015.
- 50 Akitoshi Kawamura and Yusuke Kobayashi. Fence patrolling by mobile agents with distinct speeds. *Distributed Computing*, 28(2):147–154, 2015. doi:10.1007/s00446-014-0226-3.
- 51 Dennis Komm, Rastislav Královic, Richard Královic, and Jasmin Smula. Treasure hunt with advice. In Christian Scheideler, editor, *Structural Information and Communication Complexity - 22nd International Colloquium, SIROCCO 2015, Montserrat, Spain, July 14-16, 2015, Post-Proceedings*, volume 9439 of *Lecture Notes in Computer Science*, pages 328–341. Springer, 2015. doi:10.1007/978-3-319-25258-2_23.
- 52 Amos Korman and Yoav Rodeh. Multi-round cooperative search games with multiple players. *Journal of Computer and System Sciences*, 113:125–149, 2020. doi:10.1016/j.jcss.2020.05.003.
- 53 Tobias Langner, Barbara Keller, Jara Uitto, and Roger Wattenhofer. Overcoming obstacles with ants. In Emmanuelle Anceaume, Christian Cachin, and Maria Gradinariu Potop-Butucaru, editors, *19th International Conference on Principles of Distributed Systems, OPODIS 2015, December 14-17, 2015, Rennes, France*, volume 46 of *LIPICs*, pages 9:1–9:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICs.OPODIS.2015.9.
- 54 Tobias Langner, Jara Uitto, David Stolz, and Roger Wattenhofer. Fault-tolerant ANTS. In Fabian Kuhn, editor, *Distributed Computing - 28th International Symposium, DISC 2014, Austin, TX, USA, October 12-15, 2014. Proceedings*, volume 8784 of *Lecture Notes in Computer Science*, pages 31–45. Springer, 2014. doi:10.1007/978-3-662-45174-8_3.
- 55 Christoph Lenzen, Nancy A. Lynch, Calvin Newport, and Tsvetomira Radeva. Searching without communicating: tradeoffs between performance and selection complexity. *Distributed Computing*, 30(3):169–191, 2017. doi:10.1007/s00446-016-0283-x.
- 56 Christoph Lenzen and Tsvetomira Radeva. The power of pheromones in ant foraging. In *1st Workshop on Biological Distributed Algorithms (BDA)*, 2013.
- 57 Avery Miller and Andrzej Pelc. Tradeoffs between cost and information for rendezvous and treasure hunt. *Journal of Parallel and Distributed Computing*, 83:159–167, 2015. doi:10.1016/j.jpdc.2015.06.004.
- 58 Paul J. Nahin. *Chases and Escapes: The Mathematics of Pursuit and Evasion*. Princeton University Press, 2007.
- 59 Andrzej Pelc and Ram Narayan Yadav. Advice complexity of treasure hunt in geometric terrains. *Information and Computation*, 281:104705, 2021. doi:10.1016/j.ic.2021.104705.

A State transition diagram of Algorithm 1



■ **Figure 2** A hybrid state transition diagram representing Algorithm 1. On each transition, the guard condition is given above the horizontal line. The actions that are executed if the transition is triggered are given below the horizontal line. The values φ_x , for $x \in \{N, E, S, W\}$, represent the pheromone values in neighboring nodes at the beginning of the round. $\text{dir} \in \{N, E, S, W\}$ is a variable whose value persists between transitions. The statement **move-drop** instructs the agent to move in the direction indicated by the variable dir , dropping pheromone on the destination node. The statement **move-no-drop** instructs the agent to move in the direction indicated by the variable dir , without dropping pheromone on the destination node. Exactly one guarded transition is enabled from each state.