



HAL
open science

Planning for Connected Agents in a Partially Known Environment

Arthur Queffelec, Ocan Sankur, François Schwarzenruber

► **To cite this version:**

Arthur Queffelec, Ocan Sankur, François Schwarzenruber. Planning for Connected Agents in a Partially Known Environment. AI 2021 - 34th Canadian Conference on Artificial Intelligence, May 2021, Vancouver / Virtual, Canada. pp.1-23. hal-03205744

HAL Id: hal-03205744

<https://hal.science/hal-03205744>

Submitted on 22 Apr 2021

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.

Planning for Connected Agents in a Partially Known Environment

Arthur Queffelec[†], Ocan Sankur[‡], François Schwarzentruber[†]

[†] Univ Rennes, CNRS, IRISA

[‡] Univ Rennes, Inria, CNRS, IRISA

Abstract

The Connected Multi-Agent Path Finding (CMAPF) problem asks for a plan to move a group of agents in a graph while respecting a connectivity constraint. We study a generalization of CMAPF in which the graph is not entirely known in advance, but is discovered by the agents during their mission. We present a framework introducing this notion and study the problem of searching for a strategy to reach a configuration in this setting. We prove the problem to be PSPACE-complete when requiring all agents to be connected at all times, and NEXPTIME-complete in the decentralized case, regardless of whether we consider a bound on the length of the execution.

Keywords: multi-agent path finding, multi-agent planning, connectivity, imperfect knowledge, decentralized planning

1. Introduction

The coordination of mobile agents is at the heart of many real world problems such as traffic control [1], robotics [2, 3], aviation [4] and more [5, 6]. Some of these problems have multiple aspects which make them complex: (1) Some systems are *multi-agent*, that is, the behaviors of agents influence others' and these influences must be taken into consideration when computing missions; this can be due for instance, to collisions [7], sensor interferences [8, 9] etc.; (2) Some missions must ensure *connectivity*, that is, ensure periodic or constant connection to a station/agent to share acquired information [10]; (3) The environment may be only *partially known*, and the agents may discover it during the mission [11, 12]. Several works have considered problems containing these three aspects. For instance, several algorithmic approaches have been investigated to solve the coordination of multi-robot exploration [13–15]. Our objective in this paper is to present a framework to study the theoretical complexity of planning problems with respect to these three aspects.

The theoretical complexity of some related problems have been studied in the literature. Multi-Agent Path Finding (MAPF) is an important framework introduced to study collision-free navigation of agents in warehouses (see [7, 16]). This problem was intensively studied and gave rise to a popular algorithm known as Conflict-Based Search (CBS) [17]. An extension of MAPF with connectivity constraints, called *Connected MAPF* (CMAPF), was studied as well [18]. The complexity of CMAPF and algorithmic solutions were studied in [19, 20]. However, CMAPF only addresses the multi-agent and connectivity aspects, and not the partial knowledge of the environment. The latter aspect is considered in the Canadian Traveler Problem (CTP), which is a well-known problem to study the navigation of an agent in a partially known graph [21]. While the initial framework was for a single agent, CTP has been extended to multiple agents in the settings of packet routing [22], multi-robot exploration [23] and more [24]. While a notion of communication was considered in [25, 26], it is limited to settings where all agents can receive information at all times or only designated agents can send information. In contrast, we are interested in studying the setting where agents' ability to communicate depends on their positions in the graph, and in establishing theoretical complexity results of resulting problems.

In this paper, we study the theoretical complexity of generating plans for a group of agents to reach a given target configuration. More precisely, we analyze the impact of enforcing or

^{†‡}firstname.lastname@irisa.fr

ignoring: (A) connectivity; (B) collision; (C) a bound on the length of the execution. For (A), we consider either *fully-connected* strategies, requiring that the agents remain connected at all times during the mission, or a *decentralized* strategy, allowing agents to disconnect and reconnect. (B) In some applications, collisions can be handled by a local collision avoidance system [18], and one can thus abstract away and ignore collisions in graph-based planning algorithms. (C) By providing a bound on the execution length, we can study the complexity of the decision problem associated to the optimization problem. Our results are summarized in Table 1. Interestingly, The PSPACE algorithm for the connected problem with a bounded execution is subtle and relies on a variant of the Savitch’s theorem [27] we present here. Additionally, the PSPACE-completeness holds even in the case in which agents can always communicate, thus the hardness of the problem already comes from the incomplete knowledge of the movement graph. For the decentralized case, we prove the NEXPTIME-hardness in the bounded and unbounded cases by two separate reductions from the True Dependency Quantified Boolean Formula problem (TDQBF) [28], thus showing that the problem becomes significantly harder in this case.

Let us compare our results with known complexity results. In the fully known environment, the CMAPF problem is PSPACE-complete in the connected and unbounded case [19], while it is NP-complete in the bounded case (with the bound given in unary) [18]. Thus, the partial knowledge of the environment does not render the problem harder in terms of complexity. In contrast, recall that in MAPF (without connectivity), one can check the existence of a solution in polynomial time [29], while the bounded problem is NP-hard [30], so the hardness is due to connectivity constraints. Some algorithms were presented for CMAPF in [19] that can scale up to about ten agents. Since both problems are similar and belong to PSPACE, one can hope that these approaches can be extended to the partial knowledge case. On the other hand, our results show that the complexity of the decentralized case is significantly higher. While some tools and algorithms are available for Decentralized POMDPs, which are in the same complexity class, the scalability is limited and the development of efficient algorithms for this case seems more challenging (see [31] for a survey).

	Decentralized	Connected
bounded	NEXP-complete (Th. 7)	PSPACE-complete (Th. 4, 5)
unbounded	NEXP-complete (Th. 6)	PSPACE-complete (Th. 3)

Table 1. Complexity Results.




Figure 1. Example of a topological graph.

2. Preliminaries

Connected MAPF. We consider a setting where agents move on a graph, with the *vertices/nodes* being their possible locations and the *movement edges* defining how they can move. In addition, *communication edges* define how the agents can communicate. The graphs we consider thus have two types of edges and are called *topological graphs*.

Definition 1. A topological graph is a tuple $G = \langle V, E^m, E^c \rangle$, with V a finite non-empty set of vertices, $E^m \subseteq V \times V$ a set of movement edges and $E^c \subseteq V \times V$ a set of undirected communication edges.

A movement edge (u, v) is said to be undirected if $(u, v), (v, u) \in E^m$. Figure 1 gives an example of a topological graph. For instance, an agent can go from 1 to 3 in one step. Two agents at vertices 1 and 5 can communicate, but two agents at 1 and 4 cannot.

We suppose that each vertex has a self-loop movement edge and a self-loop communication edge. This respectively represents the ability of an agent to stay at their location and to communicate with a nearby agent (i.e. at the same location/vertex).

Definition 2. A configuration is a tuple $c = \langle c_1, \dots, c_n \rangle$ where c_i is the vertex of agent i .

We write $c \rightarrow c'$ when $(c_i, c'_i) \in E^m$ for all $1 \leq i \leq n$. This means that each agent i can move from their vertex c_i in c to their vertex c'_i in c' in one step.

Definition 3 (Execution). An execution π is a sequence of configurations $\langle \pi[0], \pi[1], \dots, \pi[\ell] \rangle$ such that $\pi[t] \rightarrow \pi[t+1]$ for all $t \in \{0, \dots, \ell-1\}$.

We denote the length of π by $|\pi|$, and write $\pi[0..t-1]$ to denote the sub-execution $\langle \pi[0], \pi[1], \dots, \pi[t-1] \rangle$ of π . Moreover, $\text{last}(\pi)$ denotes the last configuration of π .

We say that a configuration c is *connected* iff the subgraph of the vertices occupied by the agents form a connected graph for relation E^c , i.e. the graph $\langle V_a, E^c \cap (V_a \times V_a) \rangle$ is connected with $V_a = \{c_1, \dots, c_n\}$. An execution is said to be *connected* iff all its configurations are connected. The Connected Multi-Agent Path Finding problem consists in finding a connected execution for the agents from a given initial configuration to a target configuration. We summarize below the known complexity results for the CMAPF problem.

Theorem 1 ([18]). *The problem of deciding whether, in a given instance (G, c^s, c^t, k) where k is unary, there is a connected execution from c^s to c^t of length at most k is NP-complete.*

Theorem 2 ([19]). *The problem of deciding whether for a given instance (G, c^s, c^t, ∞) , there exists a connected execution from c^s to c^t is PSPACE-complete.*

We do not consider collision constraints which would forbid agents from sharing the same vertex. For CMAPF with perfect knowledge, executions are not harder to compute with or without collision constraints [19, 20]. The same holds in our case, and we discuss, in the Section 6, how to incorporate these constraints in our setting.

Dependency Quantified Boolean Formula. A Dependency QBF (DQBF) is a formula in which dependencies of existential variables over universal variables are explicitly specified. A DQBF is of the form $\forall y_1, \dots, y_n \exists x_1(O_{x_1}) \dots \exists x_n(O_{x_n}) \psi$, where each O_{x_i} is, the *dependency set*, a subset of universally quantified variables, and ψ is a Boolean formula in CNF over $x_1, \dots, x_n, y_1, \dots, y_n$. It is worth noting that a QBF can be seen as a DQBF with $O_{x_1} \subseteq O_{x_2} \subseteq \dots \subseteq O_{x_n}$.

The *True DQBF* (TDQBF) is the problem of deciding whether a given DQBF holds. Formally, a DQBF φ holds iff there exists a collection of Skolem functions $A = (A_{x_i} : \{0, 1\}^{O_{x_i}} \rightarrow \{0, 1\})_{i=1..n}$ such that replacing each existential variable x_i by a Boolean formula representing A_{x_i} , turns ψ into a tautology. TDQBF is NEXPTIME-complete [28], and will be used to prove NEXPTIME lower bounds in Section 5.

3. Our framework

3.1. Modeling Imperfect Information

To formalize CMAPF in the imperfect information setting, let us show how to represent the initial knowledge of the agents, and how the information they have evolves during the execution. Agents initially know the exact set of vertices, but only have a lower and an upper approximation of the actual graph: they know that some (communication or movement) edges are *certain* (they must be present), while some are *uncertain* (they may be absent).

Definition 4 (Initial Knowledge). *The initial knowledge is modeled by a pair of topological graphs (G_1, G_2) , with the graph $G_1 = \langle V, E_1^m, E_1^c \rangle$ a lower bound, and $G_2 = \langle V, E_2^m, E_2^c \rangle$ an upper bound on the knowledge about the actual graph with $E_1^m \subseteq E_2^m$ and $E_1^c \subseteq E_2^c$.*

The agents initially know G_1 and G_2 while the actual graph $G = \langle V, E^m, E^c \rangle$ is initially unknown to them. They only know that $E_1^m \subseteq E^m \subseteq E_2^m$ and $E_1^c \subseteq E^c \subseteq E_2^c$, written as $G_1 \subseteq G \subseteq G_2$. The perfect information case is captured by $G_1 = G_2 (= G)$.




Figure 2. An initial knowledge of (G_1, G_2) .




Figure 3. An initial knowledge of a topological graph. Vertex s_7 has communication edges with all other vertices.

A movement (resp. communication) edge is said to be *certain* (i.e. sure to be present) if it is in E_1^m (resp. E_1^c); it is said *uncertain* (i.e. can be absent) if it is in $E_2^m \setminus E_1^m$ (resp. $E_2^c \setminus E_1^c$). We assume that the communication edges of the actual graph are undirected, so for all $(u, v) \in E_2^c \setminus E_1^c$, either $(u, v), (v, u) \in E^c$, or $(u, v), (v, u) \notin E^c$.

We say that an edge (u, v) is an *uncertain undirected movement edge* when $(u, v), (v, u) \in E_2^m \setminus E_1^m$. The environment can leave both edges $(u, v), (v, u)$, remove both edges $(u, v), (v, u)$, but also remove (u, v) and leave (v, u) , or remove (v, u) and leave (u, v) . That is, the movement edges of the actual graph are not necessarily undirected.

Example 1. Figure 2 depicts an initial knowledge (G_1, G_2) . The area is divided in two zones connected by two bridges, represented by the edges (s_2, s_4) and (s_3, s_5) , with an uncertainty on which bridge is open and on the communication between s_1 and s_4 .

A strategy σ_i for agent i tells where to go next after a given execution π . Formally:

Definition 5. A strategy σ_i for agent i maps any execution π to a vertex such that $(v, \sigma_i(\pi)) \in E^m$ where v is the vertex at which agent i is in the last configuration of π .

A *joint strategy* σ is a tuple $\langle \sigma_1, \dots, \sigma_n \rangle$ where σ_i is a strategy for agent i . The *outcome* of a joint strategy σ starting from configuration c^s is the execution π defined by induction as follows: $\pi[0]$ is c^s , and for $t \geq 1$, $\pi[t]$ is the configuration in which agent i is at vertex $\sigma_i(\pi[0..t-1])$.

In the context of imperfect information, the behaviors of the agents only depend on their observations, as in imperfect information games as in [32]. The strategies, as defined above, do not necessarily take observations of the agents into account. We will now formalize observations and *uniform* strategies, that is, those respecting the observations of the agents.

In our setting, at any time, an agent observes all movement edges adjacent (both in- and out-coming) to the vertex v it occupies. Moreover, they observe the presence or absence of a communication edge between v and v' if v' is occupied by another agent with which there is a direct or indirect communication (via other agents). Intuitively, during an execution, at each step, each agent updates their knowledge about the graph with these observations they receive. Moreover, they share all their knowledge with all agents with which they are connected at each step.

The observation of adjacent movement edges has been a recurrent practice in theoretical works [21, 33] as well as robotics [34, 35], and our formalism is inspired from these works.

The knowledge of an agent at any time corresponds intuitively to a pair of graphs as in Definition 4. For agent i and execution π , let us denote by $k_i(\pi)^G$ the knowledge of agent i about the graph after observing the execution π in actual graph G . Given such knowledge $K = k_i(\pi)^G$, the agent can deduce an under-approximation and an over-approximation of the actual graph; let us denote these by \underline{G}^K and \overline{G}^K respectively. In particular, if K is the knowledge the agents have initially, then $\underline{G}^K = G_1$ and $\overline{G}^K = G_2$. We present a representation of the knowledge using predicates in the appendix where a detailed formalization of $k_i(\pi)^G$ can be found.

In the rest of the paper, we assume all considered strategies to be *uniform*, that is, they comply with the knowledge of the agents: the strategies prescribe the same move to all executions that are indistinguishable with the agent's observations. Formally, if $|\pi| = |\pi'|$ and $k_i(\pi)^G = k_i(\pi')^G$, then $\sigma_i(\pi) = \sigma_i(\pi')$.

3.2. Decision Problems

We consider the decision problem of reaching a configuration c^t from a configuration c^s in less than k steps, using uniform strategies. For two configurations c^s, c^t , let us call a topological graph G (c^s, c^t) -*admissible* if there is an execution from c^s to c^t , which is not necessarily connected.

Definition 6. *We say that an instance (G_1, G_2, c^s, c^t, k) is positive if there exists a joint strategy σ such that in all (c^s, c^t) -admissible graphs G satisfying $G_1 \subseteq G \subseteq G_2$, the outcome of σ starting in c^s ends in c^t in less than k steps.*

Observe that the above problem requires that a strategy ensures the reachability of the target configuration only for graphs that are (c^s, c^t) -admissible and compatible with the initial knowledge. In fact, intuitively, we would like the strategy to work under all possible graphs G with $G_1 \subseteq G \subseteq G_2$. However, requiring a strategy to ensure reachability in a non-admissible graph does not make sense, since even a strategy with full information would fail. We thus require the strategies to make their best efforts, that is, to ensure the objective unless it is physically impossible.

Example 2. *Consider the example of Figure 2. If both bridges (i.e. movement edges (s_2, s_4) and (s_3, s_5)) are absent in the actual graph, the graph is not (s_1, s_6) -admissible and there cannot be a strategy ensuring reachability. The admissible graphs contain either (s_2, s_4) , or (s_3, s_5) , or possibly both. Note that, this instance is negative for $k < 6$ (that is, it does not admit a solution). Indeed, consider a strategy that moves the agent, for instance, to s_2 . In the graph where (s_2, s_4) is absent, the agent would need to come back to s_1 and take the alternative path, which requires an execution of total length 6; and the situation is symmetric if the first move is towards s_3 . The instance is nonetheless positive for $k \geq 6$ with the described strategy. However, if the edges (s_2, s_1) and (s_3, s_1) were not present, then the instance would be negative. In fact, once the agent moves to s_2 or s_3 , they get stuck if the graph only contains the other bridge.*

We now define the *connected* version of Definition 6. For two configurations c^s, c^t , we say that a topological graph G is (c^s, c^t) -*c-admissible* if there is a connected execution from c^s to c^t . We will often omit the pair of configurations which will be clear from the context, and write simply admissible or c-admissible.

Definition 7. *We say that an instance (G_1, G_2, c^s, c^t, k) is c-positive if there exists a joint strategy σ such that in all (c^s, c^t) -c-admissible graphs G satisfying $G_1 \subseteq G \subseteq G_2$, the outcome of σ starting in c^s is connected and ends in c^t in less than k steps.*

In the connected case, agents cannot visit a disconnected configuration. Hence, the considered strategies only visit configurations that are certainly connected. Observe that agents can make observations about the presence or absence of communication edges while being connected and use this information later.

Example 3. *Let us illustrate the above property on the example of Figure 3. Assume there are two agents, the starting and goal configurations are $c^s = \langle s_1, s_6 \rangle$ and $c^t = \langle s_5, s_7 \rangle$, and the only uncertainty is about the movement edges (s_3, s_5) and (s_4, s_5) . Here, Agent 2 could immediately move to her target s_7 ; however, she could also cooperate with Agent 1 and lower the total completion time. Indeed, from their start configuration $\langle s_1, s_6 \rangle$, the agents*

first move to $\langle s_2, s_4 \rangle$ where Agent 2 observes whether (s_4, s_5) is present. Assume the edge is present. Then, they follow the sequence $\langle s_4, s_6 \rangle \cdot \langle s_5, s_7 \rangle$; and otherwise $\langle s_3, s_6 \rangle \cdot \langle s_5, s_7 \rangle$. Thus, in order to minimize the length of the execution, the agents do not always take their shortest paths but might help other agents by obtaining information about the graph.

Consider now the same example in which the communication edge (s_3, s_6) is uncertain. If this edge is absent then Agent 2 cannot help Agent 1 achieve the target faster since if the former moves to s_4 and (s_4, s_5) is absent, then, in order to maintain connectivity, the next configurations should be $\langle s_4, s_6 \rangle \cdot \langle s_2, s_7 \rangle \cdot \langle s_3, s_7 \rangle \cdot \langle s_5, s_7 \rangle$. An execution of the same size is obtained when Agent 2 moves to s_7 in the first step.

For both Definitions 6 and 7 above, let us call a joint strategy a *witness* if it witnesses the fact that the given instance is positive, and respectively, c-positive.

We instantiate the Connected MAPF problem in four different settings. The four following decision problems are defined depending on whether we consider the connectivity requirement and whether the bound is finite. Note that the bounded problems are the decision problems associated to the optimization problems.

<i>Bounded Decentralized Reachability.</i> Is given (G_1, G_2, c^s, c^t, k) , with $k < \infty$, positive?	<i>Unbounded Decentralized Reachability.</i> Is given $(G_1, G_2, c^s, c^t, \infty)$ positive?
<i>Bounded Connected Reachability.</i> Is given (G_1, G_2, c^s, c^t, k) , with $k < \infty$, c-positive?	<i>Unbounded Connected Reachability.</i> Given $(G_1, G_2, c^s, c^t, \infty)$ c-positive?

As we will see, the encoding of the integer k does not change the overall complexity. Lower bounds are all obtained directly for the unary encoding; the lower bounds for the binary encoding follows. Concerning the upper bounds, we explain for each case how to design an algorithm with k encoded in binary.

4. Connected Reachability

We first address the case where agents must be connected at each step of the execution. In this case, agents share their knowledge at all times and thus the group of agents can be considered as a single agent playing against the environment.

4.1. Unbounded Case

We first focus on the existence of an unbounded connected strategy. Interestingly, we show that verifying the existence of a connected strategy in a partially known environment is not harder than in a perfectly known environment.

Theorem 3. *The unbounded connected reachability problem is PSPACE-complete.*

4.2. Bounded Case

We now study the existence of a bounded connected strategy. We show that this problem is PSPACE-complete even when the communication graph is complete. Please find in supplementary material detailed proofs of the Lemmas 1, 2 and Theorem 5.

Theorem 4. *The bounded connected reachability problem is in PSPACE when the bound is given in binary.*

Proof. Let us first prove the upper bound when k is given in unary. As $\text{APTITUDE} = \text{PSPACE}$ [36], we give an alternating algorithm that runs in polynomial time, as follows. At each step, the existential player chooses the next connected configuration to move the agents; and the universal player chooses the information about the newly discovered edges. After k steps




Figure 4. Gadgets for the reduction from QBF into the bounded reachability problem in the complete connectivity case.

the algorithm accepts if the target configuration is reached, or the revealed edges mean that the graph is not c -admissible. The number of steps is bounded by k , which is polynomial, thus the algorithm runs in polynomial time.

There is one subtlety to prove the correctness. The alternating algorithm actually corresponds to a slight variant of our setting which can be seen as a game. In our setting, the environment chooses a graph G with $G_1 \subseteq G \subseteq G_2$ at the beginning, and the agents discover the graph G as they move. In contrast, in the alternating algorithm, the universal player reveals the graph step by step; therefore the environment might adapt the graph to the moves of the existential player.

Lemma 1. *The alternating algorithm is correct.*

When k is binary, the previous algorithm does not run in polynomial time. However, observe that the number of alternations can be bounded by a polynomial because there is only a polynomial number of steps in which the universal players reveal *new* information to the coalition of agents. In fact, the universal player is only useful when some agent is at a vertex that has not been seen before, and this can only happen a linear number of times. Furthermore, the previous algorithm runs in polynomial space.

When k is binary, our problem is in $\text{STA}(\text{poly}(n), *, \text{poly}(n))$ where $\text{STA}(s(n), t(n), a(n))$ is the set of problems decided in space $O(s(n))$, time $O(t(n))$ with $O(a(n))$ alternations. Our problem is in PSPACE thanks to the generalization of Savitch's theorem we prove:

Lemma 2. $\text{STA}(\text{poly}(n), *, \text{poly}(n)) \subseteq \text{PSPACE}$.

Intuitively, this lemma is proved by guessing the computations between each universal choice by a PSPACE oracle, which yields an overall PSPACE algorithm. \square

Theorem 5. *The bounded connected reachability problem with complete connectivity is PSPACE-hard.*

Proof. The lower bound is proven by reduction from TQBF.

Consider a QBF φ of the form $\forall z_1 \exists z_2 \dots Q_n z_n \psi$ where ψ is a Boolean formula in conjunctive normal form with n variables and m clauses.

In the reduction, we call a *movement path*, from node v to node u , a chain of nodes linking v to u by movement edges. In addition, we denote an occurrence of a positive (resp. negative) literal of a variable z , by z (resp. $\neg z$)

We create the graphs G_1 and G_2 as described in Figure 6. More precisely, for each variable z , we create a gadget shown in Figure 6a. For each clause, we create a gadget as depicted in Figure 6b. In addition, we create a movement edge between a node $v_{\underline{z}}$ (resp. $v_{\neg\underline{z}}$) and a node t_γ if the literal z (resp. $\neg z$) is present in the clause γ .

We define the initial and target configurations c^s, c^t as follows. There is a single agent at each vertex of the form s_γ (resp. s_z) whose target is t_γ (resp. t_z). Furthermore, at each $s_{\underline{z}}$ (resp. $s_{\neg\underline{z}}$) there is one agent for each clause that contains z (resp. $\neg z$) and her target is $t_{\underline{z}}$ (resp. $t_{\neg\underline{z}}$).

We show that φ is true iff (G_1, G_2, c^s, c^t, k) is c-positive with $k = 3(n-1) + 5$.

For simplicity, we classify the agents as follows.

- A *Variable agent* is an agent starting at a node s_z ;
- a *Clause agent* is an agent starting at a node s_γ ;
- a *Positive (resp. negative) occurrence agent* is an agent starting at $s_{\underline{z}}$ (resp. $s_{\neg\underline{z}}$).

(\Rightarrow) Assume that the QBF φ is true. There exists a collection of Skolem functions A such that for each existential variable z_i (where i is even), and an assignment ν to universally quantified variables in z_1, z_3, \dots, z_{i-1} , $A_{z_i}(\nu) \in \{\top, \perp\}$ is the value assigned to z_i such that φ is true under assignment ν augmented with the values of A . We construct the following strategy σ , which guarantees that c^t is reached in k steps from c^s .

Intuitively, the lengths of the initial movement paths are designed so that agents starting at s_{z_i} arrive at v_{z_i} in the order of their indices. For an existential variable agent, the choice of the successor from v_{z_i} determines the value of z_i ; while for a universal variable agent, the choice is made by the environment. More precisely, if the agent moves to $\mathbf{r}_{\underline{z_i}}$, then z_i is set to true, if she moves to $\mathbf{l}_{\neg\underline{z_i}}$ it is set to false.

Formally, all agents start by moving to their respective v nodes (e.g. A clause agent at s_γ moves to v_γ). They arrive at these nodes at different moments due to the sizes of their movement paths. An existential variable agent arrives at node v_{z_i} at time $3(i-1) + 1$. At this point, all universal variable agents among z_1, z_3, \dots, z_{i-1} have arrived to their respective nodes v_{z_j} , thus revealing the values of these variables. Let ν be this assignment. If $A_{z_i}(\nu) = \top$, the agent moves to $\mathbf{r}_{\underline{z_i}}$, and otherwise, to $\mathbf{l}_{\neg\underline{z_i}}$.

When a universal variable agent arrives to v_{z_i} at time $3(i-1) + 1$, she follows the only edge dictated by the environment, either to $\mathbf{r}_{\underline{z_i}}$ or to $\mathbf{l}_{\neg\underline{z_i}}$. This assigns \top to the variable z_i in the former case, and \perp in the latter case. Observe that if the graph is admissible, then there must exist a path from source to target for each agent; this means that one of these edges must be present.

Consider a positive (resp. negative) occurrence agent associated to a clause γ . This agent arrives to $v_{\underline{z_i}}$ (resp. $v_{\neg\underline{z_i}}$) at time $3(i-1) + 2$. Observe that the variable agent has already determined the value of z_i in the previous step.

- if z_i is assigned \top (resp. \perp) then the agent moves to t_γ , observe which edges are available at t_γ , and immediately comes back to $v_{\underline{z_i}}$ (resp. $v_{\neg\underline{z_i}}$). Now, the edge between $\mathbf{r}_{\underline{z_i}}$ (resp. $\mathbf{l}_{\neg\underline{z_i}}$) and $u_{\underline{z_i}}$ (resp. $u_{\neg\underline{z_i}}$) has been observed by agent z_i ; so if this edge is present, she moves to $\mathbf{r}_{\underline{z_i}}$ and $u_{\underline{z_i}}$; and if not, then the edge from $\mathbf{l}_{\underline{z_i}}$ to $u_{\underline{z_i}}$ must be available, and she arrives to $t_{\underline{z_i}}$ (resp. $t_{\neg\underline{z_i}}$) at time k .

- if z_i is assigned \perp (resp. \top) then she does not visit t_γ , but moves to $\mathbf{l}_{\underline{z_i}}$ (resp. $\mathbf{l}_{\neg\underline{z_i}}$). If the edge between $\mathbf{l}_{\underline{z_i}}$ (resp. $\mathbf{l}_{\neg\underline{z_i}}$) and $u_{\underline{z_i}}$ (resp. $u_{\neg\underline{z_i}}$) is available, she moves to $t_{\underline{z_i}}$ (resp. $t_{\neg\underline{z_i}}$), otherwise she moves back and reach $t_{\underline{z_i}}$ (resp. $t_{\neg\underline{z_i}}$) at time k , through $\mathbf{r}_{\underline{z_i}}$ (resp. $\mathbf{r}_{\neg\underline{z_i}}$).

It remains to argue that clause agents can reach their target nodes within k steps. Since φ is true, by the definition of A , whatever the choice for the universal variables, some literal ℓ of each clause γ is assigned to true. Therefore, the positive or negative occurrence agent corresponding to this literal visits t_γ , thus revealing the edges available from t_γ to \mathbf{r}_γ and

1_γ . Note that at least one of these edges must be available for the graph to be admissible. Thus, a clause agent arriving to v_γ at time $3(n-1)+3$ can follow the available path to reach t_γ exactly at time k .

(\Leftarrow) Let σ be a witness joint strategy. Following σ , each clause agent c must know the available edges in the rest of their paths at time $3(n-1)+3$ since otherwise they cannot ensure reaching t_γ at time k . Thus, for each clause γ , the node t_γ is visited by some occurrence agent under strategy σ . Furthermore, an occurrence agent z (resp. $\neg z$) can visit node t_γ and still make it to t_z (resp. $t_{\neg z}$) in time iff the associated variable agent has observed the presence of the edge between u_z (resp. $u_{\neg z}$) and r_z (resp. $r_{\neg z}$) beforehand. In fact, otherwise, if the occurrence agent makes a wrong guess between l_z and r_z , they will not arrive to t_z (resp. $t_{\neg z}$) at time k . Hence, the joint strategy of the variable agents determines an assignment function which satisfies φ . □

Our reduction actually builds an undirected movement graph. Thus, PSPACE-hardness holds already for undirected movement graphs. Note that in our current setting, pairs of uncertain edges of the form (u, v) and (v, u) are treated separately, but the lower bound proof still holds when they are seen as one.

5. Decentralized Reachability

We now tackle the case where agents are allowed to be disconnected; at each configuration, they share their knowledge with all agents to which they are connected. This case is harder because agents no longer follow a centralized strategy and they must cooperate to exchange information at the right moment to reach their targets.

5.1. Unbounded Case

Theorem 6. *The unbounded decentralized reachability problem is NEXPTIME-complete.*

Proof Sketch. For the upper bound, an NEXPTIME algorithm consists in guessing uniform strategies for all agents and checking whether the joint strategy is a witness. Such a strategy has exponential size since it is a function of the sets of knowledge of the agents and the current vertex. One can enumerate all graphs G between G_1 and G_2 , and execute the joint strategy on G to check that it ensures the reachability of the target. Moreover, the executions to be checked have at most exponential length. In fact, executions can be seen as paths in a meta-graph where vertices are configurations augmented with the sets of knowledge of the agents. This meta-graph is of exponential size, so it is sufficient to consider executions of exponential length. The overall non-deterministic algorithm is thus in exponential time.

The lower bound is shown by reduction from TDQBF. Given a DQBF $\varphi = \forall y_1, \dots, y_n \exists x_1(O_{x_1}) \dots \exists x_n(O_{x_n}) \psi$, we build an instance (G_1, G_2, c^s, c^t, k) of unbounded decentralized reachability. We denote by $\gamma_1, \dots, \gamma_m$ the clauses in ψ .

The graph G_1 and G_2 as follows. For each variable z , we create the gadget depicted in Figure 5a.

We create the observation gadget for all existential variables x and for all (universal) variables $y \in O_x$, depicted in Figure 5b. For convenience, we write O for the pair (x, y) corresponding to observation of y by x .

Finally, we create the clause gadget, depicted in Figure 5c. A vertex γ_i certainly communicates with \top_z iff $z \in \gamma_i$, and with \perp_z iff $\neg z \in \gamma_i$. Moreover, the vertex v_γ communicates with all \top_z and \perp_z for all variables z .

We define the initial and target configurations as $c^s = \langle s_\gamma, s_{\gamma_1}, \dots, s_{\gamma_m}, s_{x_1}, \dots, s_{x_n}, s_{y_1}, \dots, s_{y_n}, s_{O_1}, \dots, s_{O_k} \rangle$, and $c^t = \langle t_\gamma, t_{\gamma_1}, \dots, t_{\gamma_m}, t_{x_1}, \dots, t_{x_n}, t_{y_1}, \dots, t_{y_n}, t_{O_1}, \dots, t_{O_k} \rangle$.

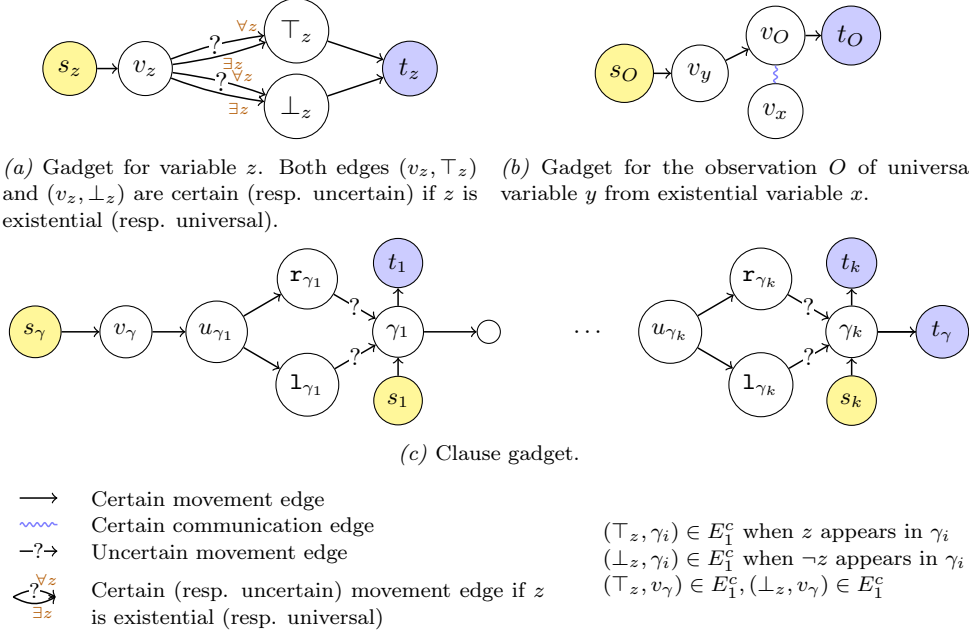


Figure 5. Gadgets in the reduction from DQBF to unbounded decentralized reachability.

Lemma 3. *DQBF φ holds if and only if $(G_1, G_2, c^s, c^t, \infty)$ is positive.*

Proof. We denote the agents as such: (1) an agent that starts at s_z as the *existential agent* z if z is an existential variable; the agent is called *universal* if z is universal; (2) the *verification agent* is the one starting at s_γ ; (3) the *clause agents* γ_i start at s_{γ_i} ; (4) the *observation agents* O start at s_O .

(\Rightarrow) Suppose the DQBF φ holds, and let A be the collection of Skolem functions. We build the following joint strategy. The environment chooses the truth values of universal variables z by deleting some edges v_z to \top_z or v_z to \perp_z . If the environment deletes the edge v_z to \top_z , the agent is forced to pass in \perp_z , thus the variable z is considered to be false. If the environment deletes the edge v_z to \perp_z , the agent is forced to pass in \top_z , thus variable z is considered to be true. If the environment deletes neither edge, then we define the strategy for agent a_y to choose to pass in y , making y true by default.

The rest of the strategy is defined as follows. At the first step, each variable agent for variable z moves to v_z , and each observation agent for the pair (x, y) moves to v_y , and thus observes the value of universal variable y . At the second step, existential agents remain in place, while observation agents move to v_O , thus sharing their observations with the corresponding existential agents. Thus, at this point, each existential agent corresponding to variable x knows the values of all universal variables $y \in O_x$. Then, agent z moves to \top_z if $A_z(\nu) = 1$ and to \perp_z otherwise, where ν is the valuation of the variables in O_z .

All clause agents move from s_i to γ_i and remain at γ_i for two steps so that all existential and universal variable agents z are at \top_z or \perp_z . The verification moves to v_γ and also waits for two steps. Since each clause is satisfied by the currently read valuation, each clause agent γ_i communicates at least with one existential or universal agent. Thus, the verification agent communicates with *all* clause agents via these variable agents. Since the clause agents communicate with the verification agent at this moment, the latter can see which edges are present in the clause gadget, and can continue go to t_γ without getting stuck.

(\Leftarrow) Conversely, suppose there is a witness joint strategy, in particular ensuring that the verification agent goes to t_z . This means that the agent must have received all the information about the topology around vertices $\gamma_1, \dots, \gamma_k$. But this is only possible if the agents have occupied a configuration in which the verification agent is at v_γ , all clause agents are at γ_i such that for each clause γ_i , there is at least one variable agent z at \top_z if $z \in \gamma_i$ and at \perp_z if $\neg z \in \gamma_i$. Thus, φ is a positive instance of TDQBF. □

□

□

5.2. Bounded Case

In the bounded case, the problem is NEXPTIME-complete independently of the encoding of the bound. Moreover, the hardness holds even for undirected graphs.

Theorem 7. *The bounded decentralized reachability problem is NEXPTIME-complete. NEXPTIME-hardness holds for undirected graphs.*

The NEXPTIME algorithm is similar to that of Theorem 6. We just add a counter in the algorithm to count the number of steps when the joint strategy is checked. The lower bound is by reduction from TDQBF and follows the ideas of the reduction in Theorem 6.

6. Discussion

Additional Results. We present results obtained by a simple observation/modification.

Unbounded Reachability and Undirected Graphs. Both the unbounded connected and unbounded decentralized reachability become *trivial* on undirected graphs. This is because we only require reachability for (c-)admissible graphs. In the decentralized case, each agent can run a DFS independently, and eventually reach their targets in at most $2|V|$ steps, and a similar search can be done by the set of agents in the connected case.

Base Station. Several works consider a designated *base* vertex to which all agents must stay connected during the execution [19, 20, 37]. This concept is only relevant in the connected case. Our results also hold with this additional constraint. In fact, the lower bound of Theorem 3 follows from [19], which proves the bound also with a base. In Theorem 5, we can add the base vertex as an isolated vertex so that the reduction is still valid.

Collisions. We did not require the paths to be collision-free in the results presented in this paper. However, this property is already ensured by our proofs or can be obtained by simple modifications. The lower bound proof of Theorem 3 relies on Theorem 2 from [19] which holds with collision constraints as well, so this is also true for our case. The proof of Theorem 5 does not generate collision-free paths as the groups of occurrence agents start and finish at the same location and follow almost the same path. This proof can be adapted to prevent collisions by delaying each occurrence agent by 3 steps behind one another. This can be achieved easily by extending the movement paths and shifting the starting location and target location of an agent up by 3 vertices behind the previous agent. The proof of lower bound of Theorem 6 features a construction ensuring a collision-free strategy. Indeed, the observations agents only need to take turns to visit the universal variables. Thus, the result holds with collision constraints as well. The algorithms of Theorem 3, 5, and 7 can be adapted by restricting all considered configurations to collision-free ones; while c-admissibility of a graph with collision constraints can be checked using Theorem 2.

Graph Classes. The MAPF and CMAPF problems have been studied for different classes of graphs (planar, grid, ...). The proof of lower bound in Theorem 3 relies on the proof of unbounded reachability done in [20], thus the result of PSPACE-hardness on planar graphs also carries over to our problem. Planar QBF is known to be PSPACE-complete [38], and

the construction of Theorem 5 is such that when applied to a planar QBF, the resulting graph is planar. Our PSPACE-completeness result thus holds on planar graphs.

Related Work. Different definitions of robust plans [39–41] have been studied. A k -robust plan guarantees the reachability of the target in the events of at most k delays. A p -robust plan executes without a conflict with probability at least p . Our framework does not consider delayed agents but focus on synchronous executions with imperfect knowledge of the area.

The problem of MAPF with a dynamic environment has multiple formulations. The Adversarial Cooperative Path-Finding [42] considers that the obstacles are agents which reason to prevent the cooperation to reach its goal. [43] considered the problem where the dynamics of the environment is predictable. Additionally, when obstacles have unknown dynamics, one can estimate their movements and plan to minimize the probability of a collision [44], or predict their movements [45] and plan online the movement of the agents [46]. In our setting, the environment is static, thus, all observations are fixed.

MAPF with Uncertainty (MAPFU) asks for a plan which guarantees that mishaps, localization and sensing errors do not impact the proper execution of the plan. This problem can be solved by temporal logic [47], POMDPs [48], replanning [49–51], interaction regions [1, 52], and belief space planning [53–55]. Nebel et al. [56] studied the MAPF problem with an uncertainty on the destination of the agents and lack of communication. The asynchronous movement of the agents, studied in those papers, cannot be expressed in our framework as we require the agent to follow some universal clock to execute their plan.

Perspectives. We proposed a setting for CMAPF in the imperfect case and studied the theoretical complexity of the reachability problem. The first natural question is to find classes of graphs (e.g. grid graphs) on which the reachability problem is easier to solve, as it was done for MAPF in [57, 58], and CMAPF in [20]. Another possible direction is to study the coverage of all vertices [20]. An alternative way to handle non-admissible graphs is to require that agents return to their starting configuration if the graph is discovered not to be admissible. We believe that such variants should be as hard as reachability. Furthermore, there are several possible generalizations that could be considered by introducing dynamic environments (instead of static), faulty sensing of agents, robustness, uncertainty, etc.

References

- [1] K. Dresner and P. Stone. “A Multiagent Approach to Autonomous Intersection Management”. In: *JAIR* 31.1 (2008), 591–656. doi: [10.1613/jair.2502](https://doi.org/10.1613/jair.2502).
- [2] E. Erdem, D. G. Kisa, U. Oztok, and P. Schüller. “A General Formal Framework for Pathfinding Problems with Multiple Agents”. In: *Proc. of AAAI*. 2013, 290–296.
- [3] J. Yu and S. LaValle. “Planning Optimal Paths for Multiple Robots on Graphs”. In: 2012, pp. 3612–3617. doi: [10.1109/ICRA.2013.6631084](https://doi.org/10.1109/ICRA.2013.6631084).
- [4] L. Pallottino, V. G. Scordio, A. Bicchi, and E. Frazzoli. “Decentralized Cooperative Policy for Conflict Resolution in Multivehicle Systems”. In: *IEEE Trans, on Rob.* 23.6 (2007), 1170–1183. doi: [10.1109/TR0.2007.909810](https://doi.org/10.1109/TR0.2007.909810).
- [5] D. Silver. “Cooperative Pathfinding”. In: *Proc. of AIIDE*. 2005, 117–122.
- [6] J. Yu and S. M. LaValle. “Optimal Multirobot Path Planning on Graphs: Complete Algorithms and Effective Heuristics”. In: *IEEE Trans. on Rob.* 32.5 (2016), 1163–1177. doi: [10.1109/TR0.2016.2593448](https://doi.org/10.1109/TR0.2016.2593448).
- [7] H. Ma and S. Koenig. “AI Buzzwords Explained: Multi-Agent Path Finding (MAPF)”. In: *AI Matters* 3 (2017). doi: [10.1145/3137574.3137579](https://doi.org/10.1145/3137574.3137579).
- [8] D. Goldberg and M. J. Mataric. “Interference as a Tool for Designing and Evaluating Multi-Robot Controllers”. In: *Proc. of AAAI*. 1997, 637–642.
- [9] M. Schneider-Fontán and M. Mataric. “Territorial multi-robot task division”. In: *IEEE Trans. on Rob. and Autom.* 14 (1998), pp. 815–822.

- [10] F. Amigoni, J. Banfi, and N. Basilico. “Multirobot Exploration of Communication-Restricted Environments: A Survey”. In: *IEEE Intell. Sys.* 32.6 (2017), pp. 48–57. DOI: [10.1109/MIS.2017.4531226](https://doi.org/10.1109/MIS.2017.4531226).
- [11] N. Rao, S. Kareti, W. Shi, and S. Iyengar. “Robot navigation in unknown terrains: Introductory survey of non-heuristic algorithms”. In: 1993.
- [12] S. Thrun. “Robotic Mapping: A Survey”. In: *Exploring Artificial Intelligence in the New Millennium*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2003, 1–35.
- [13] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider. “Coordinated multi-robot exploration”. In: *IEEE Transactions on Robotics* 21.3 (2005), pp. 376–386. DOI: [10.1109/TR0.2004.839232](https://doi.org/10.1109/TR0.2004.839232).
- [14] K. M. Wurm, C. Stachniss, and W. Burgard. “Coordinated multi-robot exploration using a segmentation of the environment”. In: *IROS*. IEEE. 2008, pp. 1160–1165.
- [15] L. Matignon, L. Jeanpierre, and A.-I. Mouaddib. “Coordinated multi-robot exploration under communication constraints using decentralized markov decision processes”. In: *Proc. of AAAI*. Vol. 26. 1. 2012.
- [16] A. Felner, R. Stern, S. E. Shimony, E. Boyarski, M. Goldenberg, G. Sharon, N. Sturtevant, G. Wagner, and P. Surynek. “Search-based optimal solvers for the multi-agent pathfinding problem: Summary and challenges”. In: *SoCS*. 2017, pp. 28–37.
- [17] G. Sharon, R. Stern, A. Felner, and N. Sturtevant. “Conflict-based search for optimal multi-agent pathfinding”. In: *Artificial Intelligence* 219 (Feb. 2015), pp. 40–66. DOI: [10.1016/j.artint.2014.11.006](https://doi.org/10.1016/j.artint.2014.11.006).
- [18] G. A. Hollinger and S. Singh. “Multirobot Coordination With Periodic Connectivity: Theory and Experiments”. In: *IEEE Trans. on Rob.* 28.4 (2012), pp. 967–973. DOI: [10.1109/TR0.2012.2190178](https://doi.org/10.1109/TR0.2012.2190178).
- [19] D. Tateo, J. Banfi, A. Riva, F. Amigoni, and A. Bonarini. “Multiagent Connected Path Planning: PSPACE-Completeness and How to Deal With It”. In: *Thirty-Second Conference on Artificial Intelligence*, 2018.
- [20] T. Charrier, A. Queffelec, O. Sankur, and F. Schwarzenrüber. “Complexity of planning for connected agents”. In: *JAAMAS* 34.2 (2020), p. 44. DOI: [10.1007/s10458-020-09468-5](https://doi.org/10.1007/s10458-020-09468-5).
- [21] C. H. Papadimitriou and M. Yannakakis. “Shortest paths without a map”. In: *Theoretical Computer Science* 84.1 (1991), pp. 127–150. DOI: [10.1016/0304-3975\(91\)90263-2](https://doi.org/10.1016/0304-3975(91)90263-2).
- [22] A. Itai and H. Shachnai. “Adaptive source routing in high-speed networks”. In: *ISTCS*. 1993, pp. 212–221. DOI: [10.1109/ISTCS.1993.253468](https://doi.org/10.1109/ISTCS.1993.253468).
- [23] W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun. “Collaborative multi-robot exploration”. In: *Proc. of ICRA*. Vol. 1. 2000, pp. 476–481. DOI: [10.1109/ROBOT.2000.844100](https://doi.org/10.1109/ROBOT.2000.844100).
- [24] L. V. Lita, J. Schulte, and S. Thrun. “A System for Multi-Agent Coordination in Uncertain Environments”. In: *Proc. of AGENTS*. 2001, pp. 21–22. DOI: [10.1145/375735.375806](https://doi.org/10.1145/375735.375806).
- [25] H. Zhang and Y. Xu. “The k-Canadian Travelers Problem with Communication”. In: *FAW-AAIM*. 2011, pp. 17–28. DOI: [10.1007/s10878-012-9503-x](https://doi.org/10.1007/s10878-012-9503-x).
- [26] D. Shiri and F. S. Salman. “On the Online Multi-Agent O—D k-Canadian Traveler Problem”. In: *J. of Comb. Opt.* 34.2 (2017), 453–461. DOI: [10.1007/s10878-016-0079-8](https://doi.org/10.1007/s10878-016-0079-8).
- [27] W. J. Savitch. “Relationships Between Nondeterministic and Deterministic Tape Complexities”. In: *J. Comput. Syst. Sci.* (1970). DOI: [10.1016/S0022-0000\(70\)80006-X](https://doi.org/10.1016/S0022-0000(70)80006-X).
- [28] G. Peterson, J. Reif, and S. Azhar. “Lower Bounds for Multiplayer Noncooperative Games of Incomplete Information”. In: *Comput & Math. Appl.* 41.7-8 (2001), pp. 957–992. DOI: [10.1016/S0898-1221\(00\)00333-3](https://doi.org/10.1016/S0898-1221(00)00333-3).
- [29] J. Yu and S. M. LaValle. “Multi-agent Path Planning and Network Flow”. In: *Algorithmic Foundations of Robotics X*. 2013, pp. 157–173.
- [30] P. Surynek. “An Optimization Variant of Multi-Robot Path Planning is Intractable”. In: *Proc. of AAAI*. 2010, 1261–1263.
- [31] C. Amato, G. Chowdhary, A. Geramifard, N. K. Üre, and M. J. Kochenderfer. “Decentralized Control of Partially Observable Markov Decision Processes”. In: *CDC*. 2013, pp. 2398–2405. DOI: [10.1109/CDC.2013.6760239](https://doi.org/10.1109/CDC.2013.6760239).

- [32] R. Berthou, B. Maubert, A. Murano, S. Rubin, and M. Y. Vardi. “Strategy logic with imperfect information”. In: *LICS*. 2017, pp. 1–12. DOI: [10.1109/LICS.2017.8005136](https://doi.org/10.1109/LICS.2017.8005136). eprint: [1805.12592](https://arxiv.org/abs/1805.12592).
- [33] A. Bar-Noy and B. Schieber. “The Canadian Traveller Problem.” In: *SODA*. Vol. 91. 1991, pp. 261–270.
- [34] G. Dudek, M. Jenkin, E. Miliotis, and D. Wilkes. “Robotic exploration as graph construction”. In: *IEEE Transactions on Robotics and Automation* 7.6 (1991), pp. 859–865. DOI: [10.1109/70.105395](https://doi.org/10.1109/70.105395).
- [35] S. Koenig, C. Tovey, and W. Halliburton. “Greedy Mapping of Terrain”. In: *Proc. of ICRA*. Vol. 4. 2001, 3594–3599 vol.4. DOI: [10.1109/ROBOT.2001.933175](https://doi.org/10.1109/ROBOT.2001.933175).
- [36] A. K. Chandra, D. Kozen, and L. J. Stockmeyer. “Alternation”. In: *J. of ACM* 28.1 (1981), pp. 114–133. DOI: [10.1145/322234.322243](https://doi.org/10.1145/322234.322243).
- [37] T. Charrier, A. Queffelec, O. Sankur, and F. Schwarzentruber. “Reachability and Coverage Planning for Connected Agents”. In: *Proc. of IJCAI*. 2019, pp. 144–150. DOI: [10.24963/ijcai.2019/21](https://doi.org/10.24963/ijcai.2019/21).
- [38] D. Lichtenstein. “Planar Formulae and Their Uses”. In: *SIAM Journal on Computing (SICOMP)* 11 (1982), pp. 329–343. DOI: [10.1137/0211025](https://doi.org/10.1137/0211025).
- [39] H. Ma, T. K. S. Kumar, and S. Koenig. “Multi-Agent Path Finding with Delay Probabilities”. In: *Proc. of AAAI*. 2017, 3605–3612.
- [40] D. Atzmon, A. Felner, R. Stern, G. Wagner, R. Barták, and N.-F. Zhou. “k-Robust Multi-Agent Path Finding”. In: *International Symposium on Combinatorial Search (SoCS)*. 2017, pp. 157–158.
- [41] D. Atzmon, R. Stern, A. Felner, N. R. Sturtevant, and S. Koenig. “Probabilistic Robust Multi-Agent Path Finding”. In: *Proc. of ICAPS*. 2020, pp. 29–37.
- [42] M. Ivanová and P. Surynek. “Adversarial Cooperative Path-Finding: Complexity and Algorithms”. In: *ICTAI*. 2014, pp. 75–82. DOI: [10.1109/ICTAI.2014.22](https://doi.org/10.1109/ICTAI.2014.22).
- [43] A. Murano, G. Perelli, and S. Rubin. “Multi-agent Path Planning in Known Dynamic Environments”. In: *PRIMA*. 2015. DOI: [10.1007/978-3-319-25524-8_14](https://doi.org/10.1007/978-3-319-25524-8_14).
- [44] J. Miura and Y. Shirai. “Probabilistic Uncertainty Modeling of Obstacle Motion for Robot Motion Planning”. In: *Journal of Robotics and Mechatronics* 14 (2002), pp. 349–356.
- [45] N. C. Griswold and J. Eem. “Control for mobile robots in the presence of moving objects”. In: *IEEE Trans. on Rob. and Autom.* 6.2 (1990), pp. 263–268. DOI: [10.1109/70.54744](https://doi.org/10.1109/70.54744).
- [46] Yun Seok Nam, Bum Hee Lee, and Nak Yong Ko. “A View-Time Based Potential Field Method for Moving Obstacle Avoidance”. In: *Proc. of SICE*. 1995, pp. 1463–1468. DOI: [10.1109/SICE.1995.526730](https://doi.org/10.1109/SICE.1995.526730).
- [47] A. Ulusoy, S. L. Smith, X. C. Ding, and C. Belta. “Robust multi-robot optimal path planning with temporal logic constraints”. In: *Proc. of ICRA*. 2012, pp. 4693–4698. DOI: [10.1109/ICRA.2012.6224792](https://doi.org/10.1109/ICRA.2012.6224792).
- [48] S. A. Miller, Z. A. Harris, and E. K. P. Chong. “Coordinated Guidance of Autonomous UAVs via Nominal Belief-State Optimization”. In: *ACC*. 2009, pp. 2811–2818. DOI: [10.1109/ACC.2009.5159963](https://doi.org/10.1109/ACC.2009.5159963).
- [49] A. Stentz. “Optimal and Efficient Path Planning for Unknown and Dynamic Environments”. In: *IJRA* 10 (Feb. 1993).
- [50] D. Ferguson, N. Kalra, and A. Stentz. “Replanning with RRTs”. In: *Proc. of ICRA*. 2006, pp. 1243–1248. DOI: [10.1109/ROBOT.2006.1641879](https://doi.org/10.1109/ROBOT.2006.1641879).
- [51] M. Likhachev, D. Ferguson, G. Gordon, A. Stentz, and S. Thrun. “Anytime Dynamic A*: An Anytime, Replanning Algorithm”. In: *Proc. of ICAPS*. 2005, 262–271.
- [52] C. Ferrari, E. Pagello, J. Ota, and T. Arai. “Multirobot motion coordination in space and time”. In: *Robotics and Autonomous Systems* 25.3-4 (1998), pp. 219–229. DOI: [10.1016/S0921-8890\(98\)00051-7](https://doi.org/10.1016/S0921-8890(98)00051-7).
- [53] A. Bry and N. Roy. “Rapidly-exploring Random Belief Trees for motion planning under uncertainty”. In: *Proc. of ICRA* (2011), pp. 723–730.
- [54] J. P. Gonzalez and A. Stentz. “Planning with uncertainty in position an optimal and efficient planner”. In: *IROS*. 2005, pp. 2435–2442. DOI: [10.1109/IROS.2005.1545048](https://doi.org/10.1109/IROS.2005.1545048).

- [55] S. Prentice and N. Roy. “The Belief Roadmap: Efficient Planning in Belief Space by Factoring the Covariance”. In: *IJRR* 28 (Oct. 2009), pp. 1448–1465. DOI: [10.1177/0278364909341659](https://doi.org/10.1177/0278364909341659).
- [56] B. Nebel, T. Bolander, T. Engesser, and R. Mattmüller. “Implicitly Coordinated Multi-Agent Path Finding under Destination Uncertainty: Success Guarantees and Computational Complexity”. In: *JAIR* 64.1 (2019), 497–527. ISSN: 1076-9757. DOI: [10.1613/jair.1.11376](https://doi.org/10.1613/jair.1.11376).
- [57] K.-H. C. Wang and A. Botea. “Tractable Multi-Agent Path Planning on Grid Maps”. In: *Proc. of IJCAI*. 2009, pp. 1870–1875.
- [58] J. Banfi, N. Basilico, and F. Amigoni. “Intractability of Time-Optimal Multirobot Path Planning on 2D Grid Graphs with Holes”. In: *RA-L* 2.4 (2017), pp. 1941–1947. DOI: [10.1109/LRA.2017.2715406](https://doi.org/10.1109/LRA.2017.2715406).

Appendix A. Complements on the Imperfect Information Setting

We present here a formal modeling of the imperfect information in our setting.

In our setting, at any time, an agent observes all movement edges adjacent (in- and out-coming edges) to the vertex v it occupies. Moreover, they observe the presence or absence of a communication edge between (v, v') if v' is occupied by another agent with which there is a direct or indirect communication (via other agents). Intuitively, during an execution, at each step, each agent updates their knowledge about the graph with these observations they receive. Moreover, they share all their knowledge with all agents with which they are connected at each step.

Given graph $G = \langle V, E^m, E^c \rangle$, let us define the direct observation $obs_i(c)$ of agent i at a configuration c to be the set:

$$\{o_{c_i, v'}^m \mid (c_i, v') \in E^m\} \cup \{o_{c_i, v'}^{-m} \mid (c_i, v') \notin E^m\} \quad (\text{A.1})$$

$$\cup \{o_{v', c_i}^m \mid (v', c_i) \in E^m\} \cup \{o_{v', c_i}^{-m} \mid (v', c_i) \notin E^m\} \quad (\text{A.2})$$

$$\cup \{o_{c_i, c_j}^c \mid j \text{ is an agent and } (c_i, c_j) \in E^c\} \quad (\text{A.3})$$

$$\cup \{o_{c_i, c_j}^{-c} \mid j \text{ is an agent connected to } i \text{ in } c, (c_i, c_j) \notin E^c\} \quad (\text{A.4})$$

where $o_{u, v}^m$, $o_{u, v}^{-m}$, $o_{u, v}^c$ and $o_{u, v}^{-c}$ are abstract terms that represent the *observations*. In the definition of $obs_i(c)$, points (1) and (2) mean that agent i directly observes the set of movement edges adjacent to her current position. Point (3) means that agent i observes a communication edge when she can communicate with another agent j . Point (4) means that agent i observes the absence of a communication edge when she sees that she can not communicate directly with another agent j but can communicate with j via multi-hop. That is, we say that j is an agent connected to i in c when there is a communication path $c_{i_1}, c_{i_2}, \dots, c_{i_k}$ with $i_1 = i$ and $i_k = j$.

Let O denote the set of all observations. We define the knowledge of an agent as a subset of O . We define the *initial knowledge for the pair* (G_1, G_2) as follows:

$$K^0(G_1, G_2) = \begin{aligned} & \{o_{u, v}^m \mid (u, v) \in E_1^m\} \cup \{o_{u, v}^c \mid (u, v) \in E_1^c\} \cup \\ & \{o_{u, v}^{-m} \mid (u, v) \notin E_2^m\} \cup \{o_{u, v}^{-c} \mid (u, v) \notin E_2^c\} \\ & \text{where } G_i = \langle V, E_i^m, E_i^c \rangle. \end{aligned}$$

This corresponds to the *a priori* knowledge on the graph all agents have before making any observation.

During the execution, agents update their knowledge at each step, upon visiting a new configuration. Formally, the knowledge $k_i((K_j)_{1 \leq j \leq n}, \pi)^G$ of agent i after observing execution π in actual graph G , where each agent j starts with initial knowledge K_j , is defined by induction on π :

- $k_i((K_j)_{1 \leq j \leq n}, \epsilon)^G = K_i$
- $k_i((K_j)_{1 \leq j \leq n}, \pi c)^G$ is the union of:

$$k_i((K_j)_{1 \leq j \leq n}, \pi)^G; \quad (\text{a})$$

$$obs_i(c); \quad (\text{b})$$

$$\bigcup_{j \text{ connected to } i \text{ in } c} (k_j((K_j)_{1 \leq j \leq n}, \pi)^G \cup obs_j(c)). \quad (\text{c})$$

Intuitively, the knowledge of an agent is composed of (a) her knowledge collected until now, (b) her current observation, and (c) the knowledge of the agents she is connected to and their current observation.

When the agents start with an initial knowledge (G_1, G_2) that is clear from the context, we will omit the tuple $(K_j)_{1 \leq j \leq n}$ and simply write $k_i(\pi)^G$.

Note that during a connected execution, all agents have an identical knowledge at all times. We thus omit the subscript i , and replace the tuple of initial knowledge sets by a single set K and write $k(K, \pi)^G$ rather than $k_i((K_j)_{1 \leq j \leq n}, \pi)^G$.

Appendix B. Proofs of Section 4

Unbounded Case. We formalize the intuition presented in the proof sketch of Theorem 3. We define the following recursive property: $P(\underline{G}, \overline{G}, c, c^t)$ holds for graphs $\underline{G}, \overline{G}$, and configurations c, c^t if either \overline{G} is not (c^s, c^t) -c-admissible, or there exists a connected execution π from c to c^t in \overline{G} such that for all G with $\underline{G} \subseteq G \subseteq \overline{G}$, writing K_0 for the initial knowledge for the pair $(\underline{G}, \overline{G})$, there exists $0 \leq i_0 \leq |\pi|$ with $k(K_0, c)^G = k(K_0, \pi[0..i_0 - 1])^G$, and either $\pi[i_0] = c^t$ or $(k(K_0, c)^G \subsetneq K = k(K_0, \pi[0..i_0])^G$ and $P(\underline{G}^K, \overline{G}^K, \pi[i_0], c^t)$).

The following two lemmas prove Theorem 3.

Lemma 4. *An instance $I = (G_1, G_2, c^s, c^t, \infty)$ is c-positive if, and only if $P(G_1, G_2, c^s, c^t)$.*

Proof. We prove the following more general property by induction on the number of edges present in G_2 and absent in G_1 , that is, $|E_2^m| - |E_1^m| + |E_2^c| - |E_1^c|$: For all graphs $G_1 \subseteq \overline{G} \subseteq G_2$, and configurations c , if the instance $(\underline{G}, \overline{G}, c, c^t, \infty)$ is c-positive then $P(\underline{G}, \overline{G}, c, c^t)$.

If $\underline{G} = \overline{G}$ then either the instance is not c-admissible and, then, $P(\underline{G}, \overline{G}, c, c^t)$ holds, or there is a connected execution in \overline{G} from c to c^t in which case the property holds as well.

Assume $\underline{G} \subsetneq \overline{G}$, and consider σ a witness strategy for the instance $(\underline{G}, \overline{G}, c, c^t, \infty)$. Consider a graph $\underline{G} \subseteq G \subseteq \overline{G}$. If the execution π induced by σ does not reveal any new observation in G , then it must end in c^t and $P(\underline{G}, \overline{G}, c, c^t)$ holds. Otherwise, let i_0 be the first step where a new observation is made. Since σ is a witness strategy, it is a witness for the instance $(\underline{G}^K, \overline{G}^K, \pi[i_0], c^t, \infty)$ as well where $K = k(\pi[0..i_0])^G$. Since \overline{G}^K and \underline{G}^K have a smaller number of differences than \overline{G} and \underline{G} , we conclude by induction that $P(\underline{G}, \overline{G}, \pi[i_0], c^t)$. Thus, $P(\underline{G}, \overline{G}, c, c^t)$ holds.

Let us now show that all instances that satisfy the property are c-positive. Assume $P(G_1, G_2, c^s, c^t)$ holds. We define the joint strategy σ on all graphs \overline{G}^K and executions π in \overline{G}^K , such that $P(\underline{G}^K, \overline{G}^K, \text{last}(\pi), c^t)$, by induction on the length of π .

Assume σ is constructed for an execution π and knowledge K . If \overline{G}^K is not c-admissible, then any strategy is a witness strategy so σ can be defined arbitrarily. Otherwise, consider the connected execution π' given by $P(\underline{G}^K, \overline{G}^K, \text{last}(\pi), c^t)$. We define σ so that agents follow π' until index i_0 , in which case either $\pi'[i_0] = c^t$ or $P(\underline{G}^{K'}, \overline{G}^{K'}, \pi\pi'[1..i_0], c^t)$. \square

Lemma 5. *$P(G_1, G_2, c^s, c^t)$ can be checked in polynomial space.*

Proof. The existence of a connected execution can be checked in polynomial space by Theorem 2. However, the size of such an execution can be exponential, and checking $P(G_1, G_2, c^s, c^t)$ requires iterating over the step of the execution. We thus need to combine the enumeration of the connected execution as we check P recursively.

The procedure to check $P(\underline{G}, \overline{G}^K, c, c^t)$ works as follows. We non-deterministically guess a connected execution step by step, from c to c^t using the PSPACE algorithm of Theorem 2. We thus only keep the last configuration in memory, a binary integer counter to bound the length of the execution (bounded by the number of configurations, thus an exponential), and the current graph \overline{G}^K . If there is no such execution, we accept. Otherwise, at each step, say, after having visited execution π and generated next configuration c' we enumerate all possible sets $K' = k(K_0, \pi c')^G$, where K_0 is the initial knowledge for the pair (G_1, G_2) . This can be done by enumerating all subsets of movement edges adjacent to c' , present in \overline{G}^K but not in \underline{G}^K , and similarly communication edges revealed by c' . Note that $k(K_0, \pi c')^G$

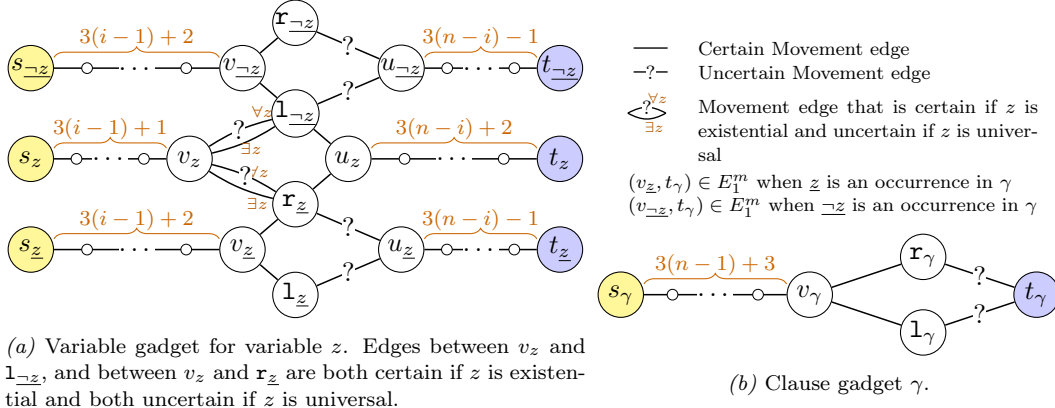


Figure 6. Gadgets for the reduction from QBF to the bounded reachability problem in the complete connectivity case.

only depends on $k(K_0, \pi)^G$ and c' which the algorithm already has. There is an exponential number of possibilities, and these can be enumerated in polynomial space. For each case, we check recursively whether $P(G^{K'}, \overline{G}^{K'}, c', c^t)$. Since the knowledge can increase only a polynomial number of times (since the knowledge can only increase when an edge is added or removed), the depth of the recursive calls is polynomial. Thus, overall, the procedure uses polynomial space. \square

Bounded Case.

Proof of Lemma 1. First, observe that if the existential player has a strategy σ in the alternating algorithm, then the instance is c -positive. In fact, for any graph G with $G_1 \subseteq G \subseteq G_2$, consider strategy τ of the universal player which makes choices according to G . Since this σ wins against τ , either the graph is not admissible or the agents successfully arrive to the target configuration. Conversely, assume that the instance is c -positive, that is, for all choices of an admissible graph G , the agents arrive at a target configuration under some joint strategy σ . We apply σ in the alternating algorithm. Consider any strategy τ of the universal player, and observe the execution induced by (σ, τ) . If the graph induced by τ is revealed not to be admissible, then the existential player wins. Otherwise, consider any admissible graph G with $G_1 \subseteq G \subseteq G_2$ compatible with the edges revealed during the execution of (σ, τ) . Since σ is winning in the original game when the underlying graph is G , the existential player also wins. \square

Proof of Lemma 2. To build a PSPACE algorithm, we perform a DFS of the computation tree T in a succinct manner. Consider the tree T' , built from T , where we only keep vertices in which the universal player makes a decision, while paths along which only the existential player moves are shortcut into single edges. The depth of this tree is polynomial by definition.

The idea is to run a DFS on T' to check whether the machine accepts. This can be done in polynomial space provided that the children of all vertices can be computed in polynomial space. Successors of an existential configuration c in T' are computed as follows: we generate on-the-fly all possible configurations c' and test whether c' is reachable from c in the original alternating machine by using a PSPACE oracle. The DFS that runs in PSPACE augmented with this PSPACE oracle gives a polynomial space procedure. \square

Proof of Theorem 5. The lower bound is proven by reduction from TQBF.

Consider a QBF φ of the form $\forall z_1 \exists z_2 \dots Q_n z_n \psi$ where ψ is a Boolean formula in conjunctive normal form with n variables and m clauses.

In the reduction, we call a *movement path*, from vertex v to vertex u , a chain of vertices linking v to u by movement edges. In addition, we denote an occurrence of a positive (resp. negative) literal of a variable z , by \underline{z} (resp. $\neg \underline{z}$)

Let us describe the construction of the graph G_1 depicted in Figure 6. For each variable z , we create a gadget depicted in Figure 6a. Formally, for all variables z , the i -th quantified variable, we create the following vertices: $s_z, s_{\underline{z}}, s_{\neg \underline{z}}, v_z, v_{\underline{z}}, v_{\neg \underline{z}}, u_z, u_{\underline{z}}, u_{\neg \underline{z}}, t_z, t_{\underline{z}}, t_{\neg \underline{z}}, \mathbf{l}_{\underline{z}}, \mathbf{l}_{\neg \underline{z}}, \mathbf{r}_{\underline{z}}$ and $\mathbf{r}_{\neg \underline{z}}$. We create the following movement paths: between s_z and v_z of length $3(i-1)+1$, between $s_{\underline{z}}$ (resp. $s_{\neg \underline{z}}$) and $v_{\underline{z}}$ (resp. $v_{\neg \underline{z}}$) of length $3(i-1)+2$, between u_z and t_z of length $3(n-i)+2$, and between $u_{\underline{z}}$ (resp. $u_{\neg \underline{z}}$) and $t_{\underline{z}}$ (resp. $t_{\neg \underline{z}}$) of length $3(n-i)-1$. Then, we add the following movement edges: from $u_{\underline{z}}$ (resp. $u_{\neg \underline{z}}$) to $\mathbf{l}_{\underline{z}}$ (resp. $\mathbf{l}_{\neg \underline{z}}$) and $\mathbf{r}_{\underline{z}}$ (resp. $\mathbf{r}_{\neg \underline{z}}$), from u_z to $\mathbf{r}_{\underline{z}}$ and $\mathbf{l}_{\neg \underline{z}}$. Finally, if the variable z is an existential variable then we add movement edges from v_z to $\mathbf{r}_{\underline{z}}$ and $\mathbf{l}_{\neg \underline{z}}$.

We create clause gadgets as depicted in Figure 6b. Formally, for all clauses γ , we create the following vertices: $s_\gamma, v_\gamma, \mathbf{l}_\gamma, \mathbf{r}_\gamma$ and t_γ . We create a movement path of length $3(n-1)+3$ from s_γ to v_γ . Finally, we create movement edges from v_γ to \mathbf{l}_γ and \mathbf{r}_γ .

Note that G_2 contains G_1 and for all universal variables, it contains additionally the following movement edges: from v_z to $\mathbf{r}_{\underline{z}}$ and $\mathbf{l}_{\neg \underline{z}}$, from u_z to $\mathbf{r}_{\underline{z}}$ and $\mathbf{l}_{\underline{z}}$, from $u_{\neg \underline{z}}$ to $\mathbf{r}_{\neg \underline{z}}$ and $\mathbf{l}_{\neg \underline{z}}$, and from t_γ to \mathbf{l}_γ and \mathbf{r}_γ .

We define the initial configuration and target configuration as follows:

$$\begin{aligned} c^s &= \langle s_{\gamma_1}, \dots, s_{\gamma_m}, s_{z_1}, \dots, s_{z_n}, s_{\underline{z}_1}, \dots, s_{\underline{z}_1}, \dots, s_{\neg \underline{z}_n} \rangle; \\ c^t &= \langle t_{\gamma_1}, \dots, t_{\gamma_m}, t_{z_1}, \dots, t_{z_n}, t_{\underline{z}_1}, \dots, t_{\underline{z}_1}, \dots, t_{\neg \underline{z}_n} \rangle. \end{aligned}$$

We show that φ is true iff (G_1, G_2, c^s, c^t, k) is c-positive with $k = 3(n-1) + 5$.

For simplicity, we classify the agents as follows.

- A *Variable agent* is an agent starting at a vertex s_z ;
- a *Clause agent* is an agent starting at a vertex s_γ ;
- a *Positive (resp. negative) occurrence agent* is an agent starting at $s_{\underline{z}}$ (resp. $s_{\neg \underline{z}}$).

(\Rightarrow) Assume that the QBF φ is true. There exists a collection of Skolem functions A such that for each existential variable z_i (where i is even), and an assignment ν to universally quantified variables in z_1, z_3, \dots, z_{i-1} , $A_{z_i}(\nu) \in \{\top, \perp\}$ is the value assigned to z_i such that φ is true under assignment ν augmented with the values of A . We construct the following strategy σ , which guarantees that c^t is reached in k steps from c^s .

Intuitively, the lengths of the initial movement paths are designed so that agents starting at s_{z_i} arrive at v_{z_i} in the order of their indices. For an existential variable agent, the choice of the successor from v_{z_i} determines the value of z_i ; while for a universal variable agent, the choice is made by the environment. More precisely, if the agent moves to $\mathbf{r}_{\underline{z}_i}$, then z_i is set to true, if she moves to $\mathbf{l}_{\neg \underline{z}_i}$ it is set to false.

Formally, all agents start by moving to their respective v vertices (e.g. A clause agent at s_γ moves to v_γ). They arrive at these vertices at different moments due to the sizes of their movement paths. An existential variable agent arrives at vertex v_{z_i} at time $3(i-1)+1$. At this point, all universal variable agents among z_1, z_3, \dots, z_{i-1} have arrived to their respective vertices v_{z_j} , thus revealing the values of these variables. Let ν be this assignment. If $A_{z_i}(\nu) = \top$, the agent moves to $\mathbf{r}_{\underline{z}_i}$, and otherwise, to $\mathbf{l}_{\neg \underline{z}_i}$.

When a universal variable agent arrives to v_{z_i} at time $3(i-1)+1$, she follows the only edge dictated by the environment, either to $\mathbf{r}_{\underline{z}_i}$ or to $\mathbf{l}_{\neg \underline{z}_i}$. This assigns \top to the variable z_i in the former case, and \perp in the latter case. Observe that if the graph is admissible, then there must exist a path from source to target for each agent; this means that one of these edges must be present.

Consider a positive (resp. negative) occurrence agent associated to a clause γ . This agent arrives to v_{z_i} (resp. $v_{\neg z_i}$) at time $3(i-1) + 2$. Observe that the variable agent has already determined the value of z_i in the previous step.

- if z_i is assigned \top (resp. \perp) then the agent moves to t_γ , observe which edges are available at t_γ , and immediately comes back to v_{z_i} (resp. $v_{\neg z_i}$). Now, the edge between r_{z_i} (resp. $l_{\neg z_i}$) and u_{z_i} (resp. $u_{\neg z_i}$) has been observed by agent z_i ; so if this edge is present, she moves to r_{z_i} and u_{z_i} ; and if not, then the edge from l_{z_i} to u_{z_i} must be available, and she arrives to t_{z_i} (resp. $t_{\neg z_i}$) at time k .

- if z_i is assigned \perp (resp. \top) then she does not visit t_γ , but moves to l_{z_i} (resp. $l_{\neg z_i}$). If the edge between l_{z_i} (resp. $l_{\neg z_i}$) and u_{z_i} (resp. $u_{\neg z_i}$) is available, she moves to t_{z_i} (resp. $t_{\neg z_i}$), otherwise she moves back and reach t_{z_i} (resp. $t_{\neg z_i}$) at time k , through r_{z_i} (resp. $r_{\neg z_i}$).

It remains to argue that clause agents can reach their target vertices within k steps. Since φ is true, by the definition of A , whatever the choice for the universal variables, some literal ℓ of each clause γ is assigned to true. Therefore, the positive or negative occurrence agent corresponding to this literal visits t_γ , thus revealing the edges available from t_γ to r_γ and l_γ . Note that at least one of these edges must be available for the graph to be admissible. Thus, a clause agent arriving to v_γ at time $3(n-1) + 3$ can follow the available path to reach t_γ exactly at time k .

(\Leftarrow) Let σ be a witness joint strategy. Following σ , each clause agent c must know the available edges in the rest of their paths at time $3(n-1) + 3$ since otherwise they cannot ensure reaching t_γ at time k . Thus, for each clause γ , the vertex t_γ is visited by some occurrence agent under strategy σ . Furthermore, an occurrence agent z (resp. $\neg z$) can visit vertex t_γ and still make it to t_z (resp. $t_{\neg z}$) in time iff the associated variable agent has observed the presence of the edge between u_z (resp. $u_{\neg z}$) and r_z (resp. $l_{\neg z}$) beforehand. In fact, otherwise, if the occurrence agent makes a wrong guess between l_z and r_z , they will not arrive to t_z (resp. $t_{\neg z}$) at time k . Hence, the joint strategy of the variable agents determines an assignment function which satisfies φ . □

Appendix C. Proofs of Section 5

Unbounded Case. Given φ , let $(G_1, G_2, c^s, c^t, \infty)$ denote the graph built in the proof sketch of Theorem 6.

Lemma 6. *DQBF φ holds if and only if $(G_1, G_2, c^s, c^t, \infty)$ is positive.*

Proof. We denote the agents as (1) an agent that starts at s_z as the *existential agent* z if z is an existential variable; the agent is called *universal* if z is universal; (2) the *verification agent* is the one starting at s_γ ; (3) the *clause agents* γ_i start at s_{γ_i} ; (4) the *observation agents* O start at s_O .

(\Rightarrow) Suppose the DQBF φ holds, and let A be the collection of Skolem functions. We build the following joint strategy. The environment chooses the truth values of universal variables z by deleting some edges v_z to \top_z or v_z to \perp_z . If the environment deletes the edge v_z to \top_z , the agent is forced to pass in \perp_z , thus the variable z is considered to be false. If the environment deletes the edge v_z to \perp_z , the agent is forced to pass in \top_z , thus variable z is considered to be true. If the environment deletes neither edge, then we define the strategy for agent a_y to choose to pass in y , making y true by default.

The rest of the strategy is defined as follows. At the first step, each variable agent for variable z moves to v_z , and each observation agent for the pair (x, y) moves to v_y , and thus observes the value of universal variable y . At the second step, existential agents remain in place, while observation agents move to v_O , thus sharing their observations with the

corresponding existential agents. Thus, at this point, each existential agent corresponding to variable x knows the values of all universal variables $y \in O_x$. Then, agent z moves to \top_z if $A_z(\nu) = 1$ and to \perp_z otherwise, where ν is the valuation of the variables in O_z .

All clause agents move from s_i to γ_i and remain at γ_i for two steps so that all existential and universal variable agents z are at \top_z or \perp_z . The verification moves to v_γ and also waits for two steps. Since each clause is satisfied by the currently read valuation, each clause agent γ_i communicates at least with one existential or universal agent. Thus, the verification agent communicates with *all* clause agents via these variable agents. Since the clause agents communicate with the verification agent at this moment, the latter can see which edges are present in the clause gadget, and can continue go to t_γ without getting stuck.

(\Leftarrow) Conversely, suppose there is a witness joint strategy, in particular ensuring that the verification agent goes to t_z . This means that the agent must have received all the information about the topology around vertices $\gamma_1, \dots, \gamma_k$. But this is only possible if the agents have occupied a configuration in which the verification agent is at v_γ , all clause agents are at γ_i such that for each clause γ_i , there is at least one variable agent z at \top_z if $z \in \gamma_i$ and at \perp_z if $\neg z \in \gamma_i$. Thus, φ is a positive instance of TDQBF. \square

Bounded Case.

Proof of Theorem 7. The upper bound when the bound k is given in unary is obtained by the following non-deterministic algorithm:

- (1) Guess a strategy σ_i for each agent i , up to executions of length $\leq k$. Such a strategy can be represented as a tree of depth k , and thus has size exponential in k .
- (2) Check that σ_i is uniform for agent i .
- (3) For all admissible graphs G such that $G_1 \subseteq G \subseteq G_2$, execute the joint strategy σ and check that the outcome execution from the initial configuration leads to the target configuration.

The obtained algorithm is non-deterministic and runs in exponential time. Note that the encoding of k is not relevant since for $k \geq 2|V|$ there is always a solution following the unbounded case.

We now prove the NEXPTIME-hardness result by reduction from TDQBF. Given an instance of TDQBF $\forall y_1, \dots, y_n \exists x_1(O_{x_1}) \dots \exists x_n(O_{x_n}) \psi$, we build an instance of bounded decentralized reachability (G_1, G_2, c^s, c^t, k) . We denote the number of clauses by m .

We construct the graph G_1 as follows:

For each variable z , we create a gadget, as depicted in Figure 7a. We create the vertices $s_z, v_z, \top_z, \perp_z$ and t_z , and we create a movement path of length $3m + 1$ from t_z to \top_z and \perp_z . Then, if the variable z is universal, we create a movement path of length $3 \times n^2 + 1$ from s_z to v_z . If the variable z is x_i , that is the i -th existential variable, we build a movement path of length $3 \times n^2 - 2|O_z| + 1$ from s_z to v_z as follows. We first build a movement path of length $3 \times (in - |O_z|)$. We then extend this movement by a vertex ρ_y for each $y \in O_z$. We extend our path from the last such vertex ρ_y to v_z by a movement path of length $3 \times (n^2 - in)$. Furthermore, we add a bidirectional edge between ρ_y and v_y .

We create the clause gadget, depicted in Figure 7b, composed of the vertices $s_\gamma, v_\gamma, t_\gamma$ and for all clauses γ_i , the vertices $u_{\gamma_i}, \mathbf{l}_{\gamma_i}, \mathbf{r}_{\gamma_i}, s_i, c_i$ and t_i . We create a movement path of length $3 \times n^2 + 2$ between s_γ and v_γ and between all s_i and γ_i . For all γ_i , we create movement edges from u_{γ_i} to \mathbf{l}_{γ_i} and to \mathbf{r}_{γ_i} , from the vertex γ_i to $u_{\gamma_{i+1}}$, or t_γ if $i = m$. In addition, we add a movement edge between v_γ and u_{γ_1} . For all γ_i , we add a movement path of length $3m + 1$ between vertices γ_i and t_i . Vertex γ_i communicates with \top_z iff $z \in \gamma_i$, and with \perp_z iff $\neg z \in \gamma_i$. Moreover, vertex v_γ communicates with all \top_z and \perp_z for all variables z .

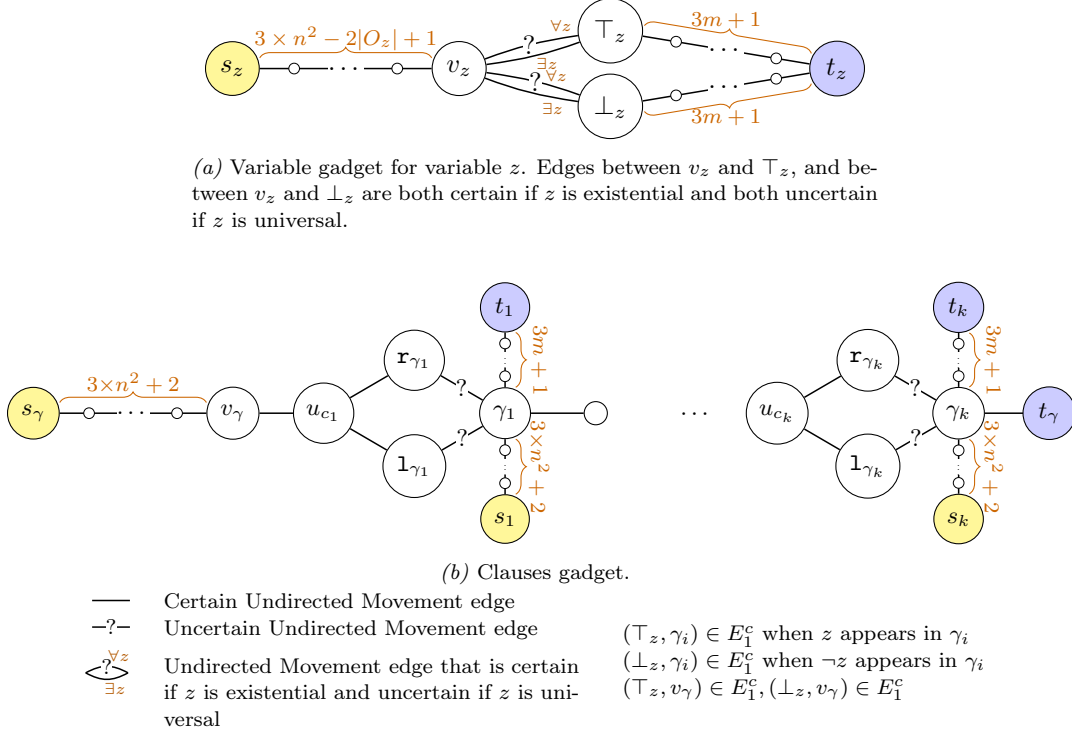


Figure 7. Gadgets for the reduction from DQBF to the bounded decentralized reachability problem.

Graph G_2 contains G_1 and for all universal variables z , we add movement edges: from v_z to \top_z and to \perp_z , and for all clauses γ_i , we add movement edges from the vertex γ_i to \mathbf{l}_{γ_i} and to \mathbf{r}_{γ_i} .

We define the initial and target configurations as $c^s = \langle s_\gamma, s_{\gamma_1}, \dots, s_{\gamma_m}, s_{z_1}, \dots, s_{z_n} \rangle$ and $c^t = \langle t_\gamma, t_{\gamma_1}, \dots, t_{\gamma_m}, t_{z_1}, \dots, t_{z_n} \rangle$.

We show that φ is a positive instance of TDQBF iff (G_1, G_2, c^s, c^t, k) is positive with $k = 3 \times n^2 + 2 + 3m + 1$. We refer to an agent that starts at s_z as the *existential agent* z if z is an existential variable; the agent is called *universal* if z is universal. The *verification agent* is the agent that starts at s_γ and the *clause agents* γ_i start at s_{γ_i} .

(\Rightarrow) Suppose the DQBF holds, and let A be an assignment function. We build a joint strategy. The environment chooses the truth values of variables z by deleting some edges v_z to \top_z or v_z to \perp_z . If the environment deletes the edge v_z to \top_z , it enforces the agent to pass in \perp_z , thus the variable z is considered to be false. If the environment deletes the edge v_z to \perp_z , it enforces the agent to pass in \top_z , thus variable z is considered to be true. If the environment deletes neither edge, then we define the strategy for agent a_y to choose to pass in y , making y true by default.

The strategy is defined as follows. Each existential agent z follows the movement path of length $3 \times n^2 - 2|O_z| + 1$ to v_z , but whenever they have vertex v_y as a neighbor, they visit v_y , come back, and continue their paths. This happens exactly O_z times, so at time $3 \times n^2 + 1$, the existential agent is at u_z . Note that along this path the agent has visited all vertices $v_{z'}$ with $z' \in O_z$, thus knows which edge among $(v_{z'}, \top_{z'})$ and $(v_{z'}, \perp_{z'})$ is present. Thus, upon arriving to v_z , the agent has the knowledge of the valuation for all variables in O_z . Observe also that by construction of the movement paths between s_z and v_z , the

agents never meet in this phase of the execution. Agent z moves to \top_z if $A_z(\nu) = 1$ and to \perp_z otherwise, where ν is the valuation of the variables in O_z .

All clause agents reach γ_i at time $3 \times n^2 + 2$. Moreover, the verification agent is at v_γ at this point, and all existential and universal variable agents z are at \top_z or \perp_z . Recall that all these vertices communicate. Since each clause is satisfied by the currently read valuation, each clause agent γ_i communicates at least with one existential or universal agent. Thus, the verification agent communicates with *all* clause agents via these variable agents. Since the clause agents can see which edges are present in the clause gadget, the verification agent has now full information about this gadget, and can continue their path until t_γ without backtracking, thus in total time $3 \times n^2 + 2 + 3m + 1$.

(\Leftarrow) Conversely, suppose there is a joint strategy enforcing that the verification agent goes to t_γ in k steps. Thus, it means that she must have received all the information about the surroundings of the vertices $\gamma_1, \dots, \gamma_k$, as she has no time to backtrack from a wrong choice. This information can only be sent to the verification agent from the clause agents through the variable agents after $3 \times n^2 + 2$ steps. The variable agents representing the assignments are connecting (i.e. satisfying) all clauses. Indeed, for all γ_i , there is one variable agent z at \top_z if $z \in \gamma_i$ and at \perp_z if $\neg z \in \gamma_i$. Thus, the DQBF is a positive instance of TDQBF. \square