

Recolorable Graph Exploration by an Oblivious Agent with Fewer Colors

Shota Takahashi ✉ 

Hosei University, Tokyo, Japan

Haruki Kanaya ✉ 

Nara Institute of Science and Technology, Ikoma, Japan

Shoma Hiraoka ✉ 

The University of Osaka, Suita, Japan

Ryota Eguchi ✉ 

Nara Institute of Science and Technology, Ikoma, Japan

Yuichi Sudo ✉ 

Hosei University, Tokyo, Japan

Abstract

Recently, Böckenhauer, Frei, Unger, and Wehner (SIROCCO 2023) introduced a novel variant of the graph exploration problem in which a single memoryless agent must visit all nodes of an unknown, undirected, and connected graph before returning to its starting node. Unlike the standard model for mobile agents, edges are not labeled with port numbers. Instead, the agent can color its current node and observe the color of each neighboring node. To move, it specifies a target color and then moves to an adversarially chosen neighbor of that color. They analyzed the minimum number of colors required for successful exploration and proposed an elegant algorithm that enables the agent to explore an arbitrary graph using only eight colors. In this paper, we present a novel graph exploration algorithm that requires only six colors. Furthermore, we prove that five colors are sufficient if we consider only a restricted class of graphs, which we call the φ -free graphs, a class that includes every graph with maximum degree at most three and every cactus.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases mobile agents, recolorable graphs, graph exploration

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2025.32

Related Version *Previous Version*: <https://arxiv.org/abs/2505.02789>

Funding *Yuichi Sudo*: This work is supported by JST FOREST Program JPMJFR226U and JSPS KAKENHI Grant Numbers JP20KK0232, JP25K03078, and JP25K03079.

1 Introduction

Autonomous mobile entities that migrate from node to node on a graph – commonly referred to as *mobile agents* or simply *agents* – have been extensively studied in the theoretical distributed computing community. Numerous fundamental problems involving these agents have been investigated in the literature, including graph exploration [13, 12, 14, 18, 19, 3], gathering [11, 16] (called rendezvous if exactly two agents are considered), dispersion [8, 9, 10, 1, 20], uniform deployment [15, 17], black hole search [5, 6, 7], and others.

Among these, graph exploration (performed by a single agent) is arguably one of the most fundamental problems, since exploration algorithms often serve as building blocks for solving many other problems. The objective is to enable the agent to visit every node in a graph, return to its starting node, and then terminate. This problem has been studied under various settings: the nodes may be *labeled* or *anonymous*; the agent may have access to the



© Shota Takahashi, Haruki Kanaya, Shoma Hiraoka, Ryota Eguchi, and Yuichi Sudo; licensed under Creative Commons License CC-BY 4.0

29th International Conference on Principles of Distributed Systems (OPODIS 2025).

Editors: Andrei Arusoaie, Emanuel Onica, Michael Spear, and Sara Tucci-Piergiovanni; Article No. 32; pp. 32:1–32:15



Leibniz International Proceedings in Informatics

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

local memories of the nodes (often referred to as *whiteboards*); the agent may have its own memory or be *oblivious*; and the initial configuration may be fixed or arbitrary (in which case the solution is self-stabilizing), among others.

All these settings for graph exploration share a common characteristic: the edges are locally labeled with *port numbers*. Typically, in a graph $G = (V, E)$, the edges incident to a node $v \in V$ are assigned distinct port numbers $0, 1, \dots, \delta_v - 1$, where δ_v is the degree of v . The agent uses these port numbers to distinguish between the edges and decide which neighbor to move to. Moreover, most studies in the literature assume that the agent can access the *incoming port* information: when the agent moves from a node u to v , it learns the port number assigned to the edge $\{u, v\}$ at the destination v , whereas the well-known rotor-router algorithm [13] and some other algorithms [3] do not use this information.

In 2023, Böckenhauer, Frei, Unger, and Wehner [2] introduced a surprisingly novel model for mobile agents in which *port numbers are unavailable*. Moreover, they assume that the nodes are anonymous and that the agent is oblivious, so additional assumptions are required to make graph exploration feasible. To overcome these limitations, they introduced *node coloring*. Specifically, at each time step, the agent at a node v can observe the color of v as well as the colors of its neighboring nodes. Based on this information, the agent updates the color of v and computes a *destination color* c , after which it moves to a neighbor with color c . If two or more neighbors of v are colored c , the destination is selected adversarially among them. At the start of exploration, all nodes are assigned a common initial color.

The authors of [2] thoroughly investigated the minimum number of colors required to solve this new variant of the graph exploration problem. It is worth mentioning that they measure color complexity by excluding the initial color. In contrast, we include the initial color in our count – a convention we believe improves clarity. Consequently, whenever we refer to results from [2], our reported color counts are always one greater.

If the agent is restricted to changing the color of each node at most once, they prove that n colors are both necessary and sufficient for arbitrary simple, undirected, and connected graphs, while four colors are necessary and sufficient for trees. (Throughout the remainder of this paper, we assume that the graph is simple, undirected, and connected, and thus omit this assumption from subsequent descriptions.) They further show that $n - 1$ colors remain necessary even for bipartite graphs of size $n = |V|$. Moreover, they prove that if the exact network size $n = |V|$ is known, then $n - 2$ and $n - 1$ colors are necessary and sufficient, respectively, for exploring arbitrary graphs.

In the model where the agent is allowed to recolor each node an unlimited number of times, which we refer to as the recolorable model, the color complexity decreases significantly. Specifically, they prove that eight colors are sufficient to explore arbitrary graphs, although no lower bound is provided.

To be candid, we could find no immediate practical application for this exploration setting. However, we believe that this new variant holds significant theoretical interest. In the absence of port numbers, most of the commonly used techniques for graph exploration and related problems cannot be applied. Consequently, entirely novel methods are required to minimize the color complexity. In this regard, the new model introduced by [2] may pave the way for exciting new avenues of research in the field of mobile agents.

Our Contribution

In this paper, we address the aforementioned variant of graph exploration introduced by [2]. In particular, within the recolorable model, we present an exploration algorithm for arbitrary graphs that uses only six colors, improving on the eight-color bound of [2] by two. Moreover,

for triangle-free graphs – those containing no cycle of length 3 – this algorithm can be adapted to use only five colors. We also prove that five colors suffice for another restricted class of graphs, which we call φ -free graphs (see Section 4); this class includes all graphs of maximum degree 3 and all cactus graphs.

Similarly to [2], we do not provide any non-trivial lower bound for exploration on arbitrary graphs in the recolorable model. Therefore, the main open question is whether six colors are indeed necessary or if this number can be further reduced. Another promising direction for future research is to allow the agent to have its own internal memory. Even a constant number of bits – or even just one bit – might significantly help to reduce the required number of colors.

► **Note 1.** This paper adheres strictly to the original model of Böckenhauer, Frei, Unger, and Wehner [2], but there is one ambiguity – apparently due to a typo in their definition – concerning whether the agent may recolor a node from a non-initial color back to the initial color. In this paper, we assume that it can, and we conjecture that the authors of [2] intended the same (see Section B in the appendix for details). If this ability were disallowed, the color complexity of our first two algorithms – for arbitrary graphs and for triangle-free graphs – would each increase by one (to seven and six, respectively). In contrast, our third algorithm for φ -free graphs would remain unchanged, since it never relies on recoloring back to the initial color. We further conjecture that even a substantial modification of the eight-color algorithm of [2] would not benefit from this capability.

2 Preliminaries

Throughout this paper, the term *graph* refers exclusively to a simple, undirected, and connected graph.

An algorithm is defined as a triple $\mathcal{A} = (C, c_0, \xi)$, where C is a set of colors, $c_0 \in C$ is the initial color, and

$$\xi : C \times \mathcal{M}(C) \rightarrow C \times (C \cup \{\text{stay}, \text{stop}\})$$

is a rule function. Here, $\mathcal{M}(C)$ denotes the set of all multisets over C , that is, all functions $m : C \rightarrow \mathbb{N}_0$, where \mathbb{N}_0 is the set of nonnegative integers. When an agent executes \mathcal{A} on a graph $G = (V, E)$, the execution proceeds as follows. Let v_t be the node where the agent is located at time $t = 0, 1, 2, \dots$, and let $c(v, t)$ denote the color of the node $v \in V$ at time t . Initially, every node is assigned the color c_0 , i.e., $c(v, 0) = c_0$ for all $v \in V$. At each time step t , the agent observes the color of its current node v_t and the colors of all its neighbors. More precisely, it obtains the current color $c(v_t, t)$ and the multiset $M = \bigsqcup_{u \in N(v_t)} \{c(u, t)\}$. Here, the symbol \bigsqcup denotes the sum (union) of multisets, and $N(v)$ represents the set of neighbors of a node $v \in V$. The agent then recolors v_t with a new color x and moves to a neighbor with color y , where


$$(x, y) = \xi(c(v_t, t), M).$$

Note that if more than one neighbor in $N(v_t)$ has color y , the agent moves to one of those nodes, with the choice made *adversarially*. We have two exceptions:

- The agent stays at the current node v_t if $y = \text{stay}$ ¹, and
- The agent terminates its execution at v_t if $y = \text{stop}$ or $y \notin M$ holds.

Even in these exceptional cases, the agent recolors v_t with x .

¹ In [2], the stay option `stay` is not included. Note that introducing the stay option does not change or extend the model: if $\xi(c, M) = (x, \text{stay})$, then we can equivalently define the rule function recursively by setting $\xi(c, M) = \xi(x, M)$, as in the original model.



■ **Figure 1** An example configuration illustrating why simulating DFS is difficult. The agent must backtrack from u_5 to u_4 , but it cannot determine which neighbor of u_5 corresponds to u_4 .

A global state, or *configuration*, is a triple $\gamma = (v, \psi, b)$, where $v \in V$ is the current node where the agent is located, $\psi : V \rightarrow C$ assigns a color to every node, and $b \in \{\mathbf{false}, \mathbf{true}\}$ indicates whether the algorithm has terminated. We write $\gamma \rightarrow \gamma'$ to denote that configuration γ transitions to γ' by one action of the agent (i.e., by updating the color and executing a movement or termination). We denote by $\gamma_{\text{init}}(v)$ the configuration in which all nodes are colored c_0 , the agent is located at node v , and the agent has not terminated. A finite or infinite sequence $\mathcal{E} = \gamma_0, \gamma_1, \dots$ of configurations is called an *execution* of an algorithm \mathcal{A} starting at node v if $\gamma_0 = \gamma_{\text{init}}(v)$ and $\gamma_t \rightarrow \gamma_{t+1}$ holds for every consecutive pair of configurations. We say that an execution \mathcal{E} is *maximal* if it is infinite or if the agent has terminated in its last configuration.

► **Definition 2** (Graph Exploration). *An algorithm \mathcal{A} is said to solve exploration on a graph $G = (V, E)$ if, for any node $s \in V$, the agent visits every node in V and terminates at s in every maximal execution of \mathcal{A} starting at s .*

Our goal is to minimize the number of colors used in graph exploration. For an algorithm $\mathcal{A} = (C, c_0, \xi)$, we define the number of colors to be $|C|$. Note that in the original work by Böckenhauer, Frei, Unger, and Wehner [2], the number of colors is defined as $|C| - 1$, since the initial color c_0 is not counted. In this paper, we count the initial color c_0 , as this definition may be clearer for most readers.

3 Exploration on Arbitrary Graphs

In this section, we show that six colors suffice for exploring arbitrary graphs. Böckenhauer, Frei, Unger, and Wehner [2] presented a simple and elegant algorithm for graph exploration that employs eight colors. In their algorithm, the agent visits all nodes in a breadth-first search (BFS) manner and terminates at the starting node s . To reduce the number of colors, we adopt a different approach in which the agent explores the graph in a manner reminiscent of depth-first search (DFS), though its movement deviates slightly. We refer to this approach as the *semi-DFS*. In the remainder of this section, we first explain the semi-DFS and then show how the agent can simulate it using only six colors.

3.1 Semi-DFS

A depth-first search (DFS) can be described in terms of a maintained path P and a set F of *finished* nodes, as follows:

1. Initially, set $P = (s)$ and $F = \emptyset$, where s is the starting node.
2. Let $P = (u_1, u_2, \dots, u_k)$. If there exists a neighbor $u \in N(u_k) \setminus (F \cup \{u_1, u_2, \dots, u_k\})$, choose any such node and append it to P . Otherwise, remove the last node (or *head*) u_k from P and add it to F .
3. If P is empty (i.e., $P = ()$), terminate; otherwise, return to step (2).

Algorithm 1 Semi-DFS.

```

1  $P \leftarrow (s), F \leftarrow \emptyset$  //  $s$  is the starting node
2 while  $P$  is not empty (i.e.,  $P \neq ()$ ) do
3   Let  $(u_1, u_2, \dots, u_k) = P$ , where  $k = |P|$ .
4   Let  $U(P, F) = N(u_k) \setminus (F \cup \{u_1, u_2, \dots, u_k\} \cup \bigcup_{i=1}^{k-1} N(u_i))$ .
5   if  $U \neq \emptyset$  then
6     Choose an arbitrary node  $u \in U$ .
7      $P \leftarrow (u_1, u_2, \dots, u_k, u)$  // expand path  $P$ 
8   else
9      $P \leftarrow (u_1, u_2, \dots, u_{k-1})$  // remove  $u_k$  from path  $P$ 
10     $F \leftarrow F \cup \{u_k\}$ 

```

Generally speaking, during a DFS, two nodes u_i and u_j in the path $P = (u_1, u_2, \dots, u_k)$ may be adjacent even if $|i - j| \neq 1$. In particular, the head u_k may have two or more neighbors that belong to P . When the agent attempts to simulate DFS using a constant number of colors, this fact poses a significant challenge (Figure 1). To backtrack from u_k to u_{k-1} , the agent must distinguish u_{k-1} from the other neighbors of u_k . If we were allowed to assign different colors to the $k - 1$ nodes u_1, u_2, \dots, u_{k-1} , the agent could easily identify u_{k-1} . However, this approach would require $\Omega(|V|)$ colors generally. Thus, there appears to be no straightforward way to implement DFS with a constant number of colors.

Therefore, we modify DFS so that any two nodes u_i, u_j in the path $P = (u_1, u_2, \dots, u_k)$ are adjacent only when $|i - j| = 1$. We refer to this modified algorithm as the semi-DFS. As we shall see in Section 3.2, this property greatly facilitates graph exploration. The pseudocode for the semi-DFS is provided in Algorithm 1. The sole difference from the original DFS occurs during path expansion: we extend the path $P = (u_1, u_2, \dots, u_k)$ only if the property is preserved. Specifically, let U denote the set of neighbors of u_k that have not been visited and are not adjacent to any node in the earlier part of the path, i.e., $U(P, F) = N(u_k) \setminus (F \cup \{u_1, u_2, \dots, u_k\} \cup \bigcup_{i=1}^{k-1} N(u_i))$. If $U \neq \emptyset$, we expand the path; otherwise, we backtrack from u_k to u_{k-1} .

► **Lemma 3.** *For any graph $G = (V, E)$ and for any starting node $s \in V$, the semi-DFS algorithm eventually terminates, and by the time it terminates, every node in G has been visited (i.e., $F = V$ eventually holds).*

Proof of Lemma 3. We prove that Algorithm 1 terminates and that $F = V$ holds upon termination.

Termination. At each iteration of the while loop, either (i) a node is appended to the path P , or (ii) a node is removed from P and added to F . Since each node can be appended to P at most once and removed from P at most once – and once a node is added to F , it is never removed – the total number of iterations is bounded by $2|V|$. Hence, the algorithm terminates after a finite number of iterations.


Full Visitation. Let $X \subseteq V$ be the set F upon termination. Assume, for the sake of contradiction, that $X \subsetneq V$. Since G is connected, the boundary set $B(X) = \{v \in X \mid \exists u \in V \setminus X : \{u, v\} \in E\}$ is not empty. Let v_{in} be the node in $B(X)$ that was removed from the return path P most recently, and choose any $v_{\text{out}} \in N(v_{\text{in}}) \setminus X \neq \emptyset$. Suppose the return

■ **Algorithm 2** Graph Exploration Algorithm for General Graphs $\mathcal{A}_{\text{GEN6}} = (C_{\text{GEN6}}, \text{init}, \xi_{\text{GEN6}})$.

Colors: $C_{\text{GEN6}} = \{\text{init}, \text{path}, \text{fin}, \text{head1}, \text{head2}, \text{neigh}\}$

Rule Function:

$$\begin{aligned} \xi_{\text{GEN6}}(\text{init}, M) &= \begin{cases} (\text{neigh}, \text{head1}) & \text{if } \text{path} \in M & \text{(Rule 1)} \\ (\text{head1}, \text{head1}) & \text{else if } \text{head1} \in M & \text{(Rule 2)} \\ (\text{head1}, \text{stay}) & \text{otherwise} & \text{(Rule 3)} \end{cases} \\ \xi_{\text{GEN6}}(\text{head1}, M) &= \begin{cases} (\text{head1}, \text{init}) & \text{if } \text{init} \in M \wedge \text{head1} \notin M & \text{(Rule 4)} \\ (\text{head2}, \text{stay}) & \text{otherwise} & \text{(Rule 5)} \end{cases} \\ \xi_{\text{GEN6}}(\text{head2}, M) &= \begin{cases} (\text{head2}, \text{neigh}) & \text{if } \text{neigh} \in M & \text{(Rule 6)} \\ (\text{path}, \text{head1}) & \text{else if } \text{head1} \in M & \text{(Rule 7)} \\ (\text{fin}, \text{path}) & \text{else if } \text{path} \in M & \text{(Rule 8)} \\ (\text{fin}, \text{stop}) & \text{otherwise} & \text{(Rule 9)} \end{cases} \\ \xi_{\text{GEN6}}(\text{path}, M) &= (\text{head1}, \text{stay}) & \text{(Rule 10)} \\ \xi_{\text{GEN6}}(\text{neigh}, M) &= (\text{init}, \text{head2}) & \text{(Rule 11)} \end{aligned}$$



■ **Figure 2** Lifecycle of node colors in algorithm $\mathcal{A}_{\text{GEN6}}$.


path was $P = (u_1, u_2, \dots, u_k)$ with $v_{\text{in}} = u_k$ at the moment when v_{in} was removed, at which $U(P, F) = N(u_k) \setminus (F \cup \{u_1, u_2, \dots, u_k\} \cup \bigcup_{i=1}^{k-1} N(u_i))$ must be empty. Since $v_{\text{out}} \in N(u_k)$, at least one of the following must hold:

1. $v_{\text{out}} \in F \subseteq X$,
2. $v_{\text{out}} \in \{u_1, u_2, \dots, u_k\} \subseteq X$, or
3. v_{out} is adjacent to some u_i with $1 \leq i \leq k-1$.


Cases 1 and 2 contradict $v_{\text{out}} \notin X$. Case 3 yields $u_i \in B(X)$ for some $i < k$, contradicting the choice of v_{in} as the most recently removed boundary node. Hence $X = V$, as required. ◀

3.2 How to simulate Semi-DFS with Six Colors

In this section, we present an algorithm $\mathcal{A}_{\text{GEN6}} = (C_{\text{GEN6}}, \text{init}, \xi_{\text{GEN6}})$ that enables the agent to explore a graph using six colors, namely `init`, `path`, `fin`, `head1`, `head2`, and `neigh`, by simulating the movement of semi-DFS. The definition of $\mathcal{A}_{\text{GEN6}}$ is given in Algorithm 2. Some readers might be concerned about the absence of a definition for $\xi_{\text{GEN6}}(\text{fin}, M)$ in Algorithm 2; however, this poses no issue since, as we shall see later, the agent is located at a node colored `fin` only after the algorithm terminates. Figure 2, which describes how a node changes its color from `init` to `fin`, may help the reader quickly grasp the algorithm.



■ **Figure 3** Path expansion.



■ **Figure 4** Removal of the head.

Algorithm $\mathcal{A}_{\text{GEN6}}$ simulates the Semi-DFS specified by Algorithm 1. In $\mathcal{A}_{\text{GEN6}}$, nodes colored **fin** represent finished nodes (i.e., the nodes in F), while nodes colored **path** and **head1** form the path P of the semi-DFS. Nodes colored **init** indicate that they are neither finished nor on P . The colors **head2** and **neigh** are employed to find a node in $U(P, F) = N(u_k) \setminus (F \cup \{u_1, u_2, \dots, u_k\} \cup \bigcup_{i=1}^{k-1} N(u_i))$. A detailed explanation of $\mathcal{A}_{\text{GEN6}}$ is provided in the proof of the following theorem.

► **Theorem 4.** *Algorithm $\mathcal{A}_{\text{GEN6}} = (C_{\text{GEN6}}, \text{init}, \xi_{\text{GEN6}})$, which uses six colors, solves the exploration problem on any graph $G = (V, E)$.*

Proof. We fix a starting node $s \in V$ and consider an arbitrary execution $\mathcal{E} = \gamma_0, \gamma_1, \dots$ of $\mathcal{A}_{\text{GEN6}}$ starting at s . To prove correctness, it suffices to show that \mathcal{E} is finite and that every node is colored **fin** in the final configuration since a node is recolored from the initial color **init** only after it has been visited.


We say that a configuration $\gamma = (v, \psi, b)$ of $\mathcal{A}_{\text{GEN6}}$ is *regular* if either

1. $b = \text{true}$, $v = s$, and every node u is colored **fin** (i.e., $\psi(u) = \text{fin}$), or
2. $b = \text{false}$ and there exists a path $P = (u_1, u_2, \dots, u_k)$ with $k \geq 1$ such that $u_1 = s$, $u_k = v$, u_i and u_j are adjacent only when $|i - j| = 1$ for all $1 \leq i, j \leq k$, $\psi(u_k) = \text{head1}$, $\psi(u_i) = \text{path}$ for all $1 \leq i < k$, and every node not on P is colored either **init** or **fin**.


Note that in any regular configuration, no node is colored **neigh** or **head2**. We refer to the regular configuration satisfying the first condition as the *final configuration*. If a configuration γ is regular but not final, we denote the unique path defined above by $P(\gamma)$ and call it the *return path* of γ . Moreover, we denote $F(\gamma) = \{u \in V : \psi(u) = \text{fin}\}$.

In the initial configuration γ_0 , each node has color **init**, and the agent is located at s . Therefore, by **Rule 3**, only s is colored **head1**, while all other nodes remain colored **init**, and the agent remains at s in γ_1 . This configuration γ_1 is regular and corresponds to the initial setting in the semi-DFS, i.e., $P(\gamma_0) = (s)$ and $F(\gamma_0) = \emptyset$. In the remainder of this proof, we show that from any regular but non-final configuration γ , $\mathcal{A}_{\text{GEN6}}$ simulates one iteration of the semi-DFS while loop, yielding a new regular configuration. By Lemma 3, it follows that \mathcal{E} eventually reaches a final configuration, thereby proving correctness.

Let $\gamma = (v, \psi, b)$ be any regular, non-final configuration. In γ , the agent is located at the head of the return path $P(\gamma) = (u_1, u_2, \dots, u_k)$; that is, the agent is at node $v = u_k$, which is colored **head1**. Our goal is to verify that the agent simulates one iteration of the while



■ **Figure 5** φ -shaped graphs.



■ **Figure 6** φ -shaped graph in the proof of Lemma 7.

loop of the semi-DFS from this configuration: if there exists a node in $U(P(\gamma), F(\gamma)) = N(u_k) \setminus (F(\gamma) \cup \{u_1, u_2, \dots, u_k\} \cup \bigcup_{i=1}^{k-1} N(u_i))$, the agent appends an arbitrary such node to P ; if $U(P(\gamma), F(\gamma)) = \emptyset$ and $k \geq 2$, the agent removes the head u_k from P and backtracks to u_{k-1} , recoloring u_k with `fin`; otherwise, the agent terminates.

Let $I(\gamma) = \{u \in N(u_k) \mid \psi(u) = \text{init}\}$. By **Rule 4**, the agent visits each node $u \in I(\gamma)$ one by one. Checking whether $u \in U(P(\gamma), F(\gamma))$ is straightforward: since u is colored `init`, it belongs to $U(P(\gamma), F(\gamma))$ if and only if none of its neighbors is colored `path`. The agent colors u with `head1` if $u \in U(P(\gamma), F(\gamma))$, and with `neigh` otherwise, after which the agent returns to u_k (**Rules 1 and 2**). We now consider two cases (Figures 3 and 4 may help the reader quickly grasp the two scenarios).

Case 1: $U(P, F) \neq \emptyset$. In this case, the agent visits the nodes in $I(\gamma)$, recoloring them with `neigh`, until it encounters a node $u \in U(P, F)$. Then, the agent colors u with `head1` and returns to u_k . At u_k , upon observing that u is colored `head1`, the agent recolors u_k with `head2` (**Rule 5**). By **Rules 6 and 11**, the agent then visits all neighbors colored `neigh` and reverts their colors back to `init`. Finally, the agent expands P by recoloring u_k with `path` and moving to u (**Rule 7**).

Case 2: $U(P, F) = \emptyset$. In this case, the agent recolors all nodes in $I(\gamma)$ with `neigh`, recolors u_k with `head2` (**Rule 5**), and reverts their colors back to `init` (**Rules 6 and 11**). If $k \geq 2$, the agent moves to and recolors u_{k-1} with `head1` (**Rules 8 and 10**); otherwise, it terminates (**Rule 9**). In any case, u_k is recolors with `fin`.

In both cases, we have shown that the agent correctly simulates one complete iteration of the semi-DFS while loop, resulting in a new regular configuration. ◀

3.3 Triangle-free Case

When we restrict our attention to triangle-free graphs, i.e., graphs that contain no cycle of length 3, we can reduce the number of colors by one. Specifically, by merging the two colors `head1` and `head2`, our algorithm requires only five colors. This yields the following theorem. A detailed description of the algorithm and its correctness is provided in the appendix (Section A).

► **Theorem 5.** *There exists an algorithm that uses five colors and solves the exploration problem on any graph $G = (V, E)$ containing no cycle of length 3.*

■ **Algorithm 3** Graph Exploration Algorithm for φ -free Graphs $\mathcal{A}_{\text{PF5}} = (C_{\text{PF5}}, \text{init}, \xi_{\text{PF5}})$.
 (Note: All arithmetic on indices in $\{0, 1, 2\}$ is performed modulo 3; the modulo notation is omitted.)

Colors: $C_{\text{PF5}} = \{\text{init}, d_0, d_1, d_2, \text{fin}\}$

Notation: $f(M) = \min\{i \in \{0, 1, 2\} \mid d_{i-1} \in M \wedge d_{i+1} \notin M\}$

Rule Function:

$$\xi_{\text{PF5}}(\text{init}, M) = \begin{cases} (\text{fin}, d_0) & \text{if } \{d_0, d_1, d_2\} \subseteq M & \text{(Rule 1)} \\ (d_{f(M)}, \text{stay}) & \text{else if } \exists j \in \{0, 1, 2\} : d_j \in M & \text{(Rule 2)} \\ (d_0, \text{stay}) & \text{otherwise} & \text{(Rule 3)} \end{cases}$$

$$\xi_{\text{PF5}}(d_i, M) = \begin{cases} (d_i, \text{init}) & \text{if } \text{init} \in M & \text{(Rule 4)} \\ (d_i, d_{i+1}) & \text{else if } d_{i+1} \in M & \text{(Rule 5)} \\ (\text{fin}, d_{i-1}) & \text{else if } d_{i-1} \in M & \text{(Rule 6)} \\ (\text{fin}, \text{stop}) & \text{otherwise} & \text{(Rule 7)} \end{cases}$$

for each $i \in \{0, 1, 2\}$.

4 Exploration on φ -free Graphs

In this section, we show that five colors suffice for exploring φ -free graphs, defined below.

► **Definition 6** (φ -free graph). A graph $G = (V, E)$ is said to be φ -shaped if there exist three distinct nodes $s, t, u \in V$ such that:

- there exist three node-disjoint paths $p_1, p_2,$ and p_3 between s and t that do not contain u
- every node in $V \setminus \{u\}$ lies on at least one of $p_1, p_2,$ or $p_3,$ and
- $\{t, u\} \in E,$

A graph is said to be φ -free if it does not contain any φ -shaped subgraph.

Typical examples of φ -shaped graphs are shown in Figure 5. The term φ -free is derived from the shape of the φ -shaped graphs. Every graph with maximum degree at most three is φ -free, since in any φ -shaped graph the node t has degree four. Moreover, every cactus is also φ -free because every φ -shaped graph violates the cactus condition that no two cycles share more than one node. The following useful lemma holds.

► **Lemma 7.** A simple, undirected graph $G = (V, E)$ is φ -free if and only if, for every node $v \in V$ with degree at least four, no connected component of $G - v$ contains three or more neighbors of v .

Proof.

Sufficiency. Every φ -shaped graph contains a node t with degree at least four such that three neighbors of t lie in the same connected component of $G - t$ (see Figure 5).

Necessity. Assume that there exists a node $t \in V$ with degree at least four for which three of its neighbors, say $u_1, u_2,$ and $u_3,$ lie in the same connected component of $G - t$. Then, in that connected component of $G - t$, there is a spanning tree containing $u_1, u_2,$ and u_3 . Consequently, there exists a node s in this component such that there are three node-disjoint paths $q_1, q_2,$ and q_3 in $G - t$, where each q_i is a path from s to u_i .

Now, by adding the edge $\{t, u_i\}$ to each q_i , we obtain three node-disjoint paths p_1, p_2 , and p_3 from s to t in G . Let u_4 be another neighbor of t different from u_1, u_2 , and u_3 . If none of p_1, p_2 , or p_3 contains u_4 , then G contains a φ -shaped subgraph, as required. Otherwise, assume without loss of generality that p_1 contains u_4 (see Figure 6). Let q_4 be the subpath of p_1 connecting s and u_4 . Then, the three node-disjoint paths $p_4 = q_4 + \{u_4, t\}$, p_2 , and p_3 from s to t (with the edge $\{t, u_i\}$ appended as before) and edge $\{t, u_1\}$ form a φ -shaped subgraph. Note that this discussion holds even in $s = u_2$ or $s = u_3$. ◀

As mentioned at the beginning of Section 3, The authors of [2] use eight colors to explore arbitrary graphs. Their approach employs a breadth-first search (BFS) strategy, encoding the distance between the starting node s and each node v modulo 3 as the *color* of v . Specifically, they use the color set

$$\{\text{init}, \text{fin}, g_0, g_1, g_2, r_0, r_1, r_2\},$$

where we have renamed these colors to be consistent with our notation. As you can see, exactly two colors are used to encode each distance (modulo 3) from the starting node, which they refer to as the green and red colors. This multiplicity (i.e., having both a green and a red color for each distance) plays a critical role in computing the correct distance (modulo 3) from the starting node in the BFS manner. Given these distances, the agent has a *sense of direction*: when the agent is located at a node v with distance $i \in \{0, 1, 2\}$, the neighbors in $N(v)$ with distance $i - 1 \pmod 3$ are the nodes closer to s , while those with distance $i + 1 \pmod 3$ are the nodes further away. This sense of direction greatly facilitates the agent's task of visiting all nodes in the graph.

In this section, we show that the multiplicity of colors (i.e., the separate green and red colors) can be eliminated when we consider only φ -free graphs. Our algorithm $\mathcal{A}_{\text{PF5}} = (C_{\text{PF5}}, \text{init}, \xi_{\text{PF5}})$ uses a color set of size five:

$$C_{\text{PF5}} = \{\text{init}, \text{fin}, d_0, d_1, d_2\}.$$

The key idea is to abandon computing the exact distance between each node v and the starting node s . Instead, we construct a directed acyclic graph (DAG) on the graph $G = (V, E)$ using the three colors d_0, d_1, d_2 , which we refer to collectively as *distance colors*. Specifically, given a configuration $\gamma = (v, \psi, b)$, we define a digraph $D(\gamma) = (V, A)$ as follows:

$$A = \{(u, v) \in V \times V \mid \{u, v\} \in E, \exists i \in \{0, 1, 2\} \text{ such that } \psi(u) = d_i \text{ and } \psi(v) = d_{i+1}\}.$$

Intuitively, a node u colored d_i has an outgoing arc to a neighbor $v \in N(u)$ if and only if v is colored d_{i+1} . Here and throughout, all arithmetic on indices in $\{0, 1, 2\}$ is performed modulo 3 (the modulo notation is omitted for brevity); for example, $d_{2+1} = d_0$ and $d_{0-1} = d_2$. When no ambiguity arises, we omit the argument γ and simply denote $D(\gamma)$ by D .

The definition of $\mathcal{A}_{\text{PF5}} = (C_{\text{PF5}}, \text{init}, \xi_{\text{PF5}})$ is provided in Algorithm 3. At time step 0, the agent colors the starting node s with d_0 (**Rule 3**). Thereafter, the agent moves to a node with the initial color *init* only from a node with a distance color by **Rule 4**. Upon arriving at a node colored *init*, say $v \in V$, the agent observes the multiset $M = \bigsqcup_{u \in N(v)} \{c(u)\}$, where $c(u)$ denotes the color of a neighbor $u \in N(v)$. The agent then colors v based on this multiset M and the following two principles:

1. The digraph D must remain acyclic (i.e., no cycles are ever created), and
2. Any node v that is colored d_i for some $i \in \{0, 1, 2\}$ must be reachable from s in D .

Specifically, **Rule 2** enforces these principles: whenever $\{d_0, d_1, d_2\} \not\subseteq M$, the agent assigns v the color $d_{f(M)}$, where

$$f(M) = \min\{i \in \{0, 1, 2\} \mid d_{i-1} \in M \text{ and } d_{i+1} \notin M\}.$$

The first principle is preserved because no neighbor of v is colored $d_{f(M)+1}$, and the second principle is maintained because there is at least one neighbor colored $d_{f(M)-1}$. Note that $f(M)$ is uniquely determined here since M contains at least one distance color (recall that the agent has moved to v by **Rule 4**) and $\{d_0, d_1, d_2\} \not\subseteq M$. If, on the other hand, $\{d_0, d_1, d_2\} \subseteq M$, then the agent colors v with **fin** and then moves to an arbitrary neighbor with color d_0 (**Rule 1**).

As mentioned above, the starting node s is colored d_0 . Consequently, in the digraph D , any node v with color d_i is at distance i from s (modulo 3). This provides the agent with a sense of direction, as in the eight-color algorithm of [2]. Suppose that the agent is at a node v with distance color d_i and none of its neighbors is colored **init**. Then, if there exists a neighbor with color d_{i+1} , the agent moves forward to that neighbor (**Rule 5**); otherwise, it backtracks to a neighbor with color d_{i-1} (**Rule 6**), at which the agent colors v with **fin**. The only exception occurs when the agent is at the starting node s (with color d_0) and no neighbor is colored d_1 ; in that case, the agent terminates the execution (**Rule 7**).

The following lemma guarantees that the two principles mentioned above are always preserved.

► **Lemma 8.** *During any execution of \mathcal{A}_{PF5} starting at any node $s \in V$, the following two invariants are maintained: (i) the digraph D is acyclic, and (ii) every node assigned a distance color is reachable from the starting node s in D .*

Proof. At the start of exploration, these invariants clearly hold: no node has been assigned a distance color, and consequently D is acyclic. In what follows, we show that none of the rules (i.e., **Rules 1, 2, ..., 7**) violate these invariants.

The first invariant could potentially be broken only when the agent colors a node v with a distance color via **Rules 2** or **Rule 3**. By the definition of the function f , **Rule 2** never creates a cycle in D . Similarly, **Rule 3** preserves acyclicity because it is applicable only when v has no neighbor with a distance color, ensuring that no cycle is introduced.

The second invariant could be compromised only when the agent removes a distance color from a node v , as occurs in **Rules 6** or **Rule 7**. However, these rules are applied only when no neighbor of v is colored d_{i+1} (if v is colored d_i), thus v has no out-going arc in D . Hence, the removal does not disconnect any node from s in D . Therefore, both invariants are maintained throughout the execution. ◀

► **Lemma 9.** *During any execution of \mathcal{A}_{PF5} on any φ -free graph, **Rule 1** is applicable only when the agent is at a node of degree exactly 3.*

Proof. Let v be a node that the agent colors with **fin** by applying **Rule 1** during an execution of \mathcal{A}_{PF5} . For **Rule 1** to be applicable at v , the agent must observe all three distance colors d_0 , d_1 , and d_2 among the neighbors of v . Hence, v must have at least three neighbors.

Assume for contradiction that v has degree at least four. By **Rules 1, 2, and 3**, once the agent visits v , it assigns v either a distance color or **fin** and never reverts it back to **init**. Thus, **Rule 1** is applicable at v only during the first visit to v . By Lemma 7, before this first visit, the agent can have visited at most two neighbors of v , so that at most two neighbors have been assigned a distance color. This contradicts the requirement that all three distance colors must be observed at v . ◀

► **Lemma 10.** *Let F denote the set of nodes colored fin. During any execution $\gamma_0, \gamma_1, \dots$ of \mathcal{A}_{PF5} starting at any node $s \in V$ on any φ -free graph, the induced subgraph $G[V \setminus F]$ is always connected.*

Proof. Let $F(i)$ be the set F in configuration γ_i . Assume for contradiction that there exists an integer $i \geq 0$ such that $G[V \setminus F(i)]$ is connected, but $G[V \setminus F(i+1)] = G[V \setminus F(i)] - v$ is disconnected, where v is the node at which the agent is located in γ_i . We consider two cases, depending on the rule that causes v to be colored fin.

Case 1: The agent applied Rule 1 in γ_i . By Lemma 9, v then has exactly three neighbors, and all three must already have been assigned a distance color.

Case 2: The agent applied Rule 6 or Rule 7 in γ_i . In this case, **Rule 4** was not applicable in γ_i , implying that every neighbor of v in $N(v) \setminus F(i)$ already has a distance color.

In either case, every neighbor of v in $G[V \setminus F(i)]$ has a distance color. By Lemma 8, all these neighbors remain reachable from the starting node s in the digraph D , even after the removal of v . Thus, $G[V \setminus F(i)] - v$ must be connected, which contradicts our assumption. ◀

► **Theorem 11.** *Algorithm $\mathcal{A}_{PF5} = (C_{PF5}, \text{init}, \xi_{PF5})$, which uses five colors, solves the exploration problem on any φ -free graph.*

Proof. We first show that the agent eventually terminates. Notice that the agent changes the color of each node at most twice. Consequently, the rules that alter a node's color (**Rule 1,2,3,6,7**) are applied only finitely many times. Moreover, **Rule 4** is applied only finitely many times since each application of **Rule 4** eventually triggers an application of one of **Rule 1,2,3** in the subsequent step. The only remaining rule is **Rule 5**; however, by Lemma 8, the digraph D remains acyclic throughout the execution, so the agent cannot apply **Rule 5** indefinitely. Thus, the agent must eventually terminate.

Now, consider an execution $\gamma_0, \gamma_1, \dots, \gamma_t$ of \mathcal{A}_{PF5} starting at some node $s \in V$ on a φ -free graph, where at step $t-1$ the agent applies **Rule 7** at a node v with distance color d_i , thereby coloring v with fin and terminating the execution. We claim that v must be the starting node s . Otherwise, by Lemma 8, v would have a neighbor with color d_{i-1} , contradicting the condition required to apply **Rule 7**. Hence, the agent terminates at s .

Finally, we show that in the final configuration γ_t , every node is colored fin. In γ_t , no neighbor of s is colored init, because the agent does not apply **Rule 4** in the step preceding termination. Moreover, no node retains a distance color in γ_t , since any node with a distance color would be reachable from s in D , yet s has no outgoing arcs in D in the final configuration. Therefore, all nodes must be colored fin. By Lemma 10, the induced subgraph $G[V \setminus F]$ remains connected, which implies that $F = V$ in γ_t . This completes the proof. ◀

5 Conclusion

In this paper, we have advanced the study of graph exploration by a single oblivious agent in the recolorable model, originally introduced by Böckenhauer et al. [2]. Our main contributions are twofold:

- We present a novel exploration algorithm that reduces the color complexity on arbitrary undirected connected graphs from eight to six colors. Moreover, if the graph is triangle-free (i.e., contains no 3-cycles), the color complexity further reduces to five colors.
- We identify a broad graph class – φ -free graphs, which includes all graphs of maximum degree at most three and all cacti – and show that five colors suffice for exploring every graph in this class.

Several intriguing open problems remain. On arbitrary graphs, no nontrivial lower bound on the number of colors is known, so determining whether six is optimal (or can be further decreased) is a compelling challenge. Moreover, it would be interesting to study how a small amount of internal agent memory – or relaxations such as randomized moves – might impact the color complexity.

References

- 1 John Augustine and William K. Moses Jr. Dispersion of mobile robots. *Proceedings of the 19th International Conference on Distributed Computing and Networking*, January 2018.
- 2 Hans-Joachim Böckenhauer, Fabian Frei, Walter Unger, and David Wehner. Zero-memory graph exploration with unknown inports. In *International Colloquium on Structural Information and Communication Complexity*, pages 246–261. Springer, 2023. doi:10.1007/978-3-031-32733-9_11.
- 3 Dominik Bojko, Karol Gotfryd, Dariusz R. Kowalski, and Dominik Pająk. Tree Exploration in Dual-Memory Model. In *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, pages 22:1–22:16, 2022. doi:10.4230/LIPIcs.MFCS.2022.22.
- 4 Hans-Joachim Böckenhauer, Fabian Frei, Walter Unger, and David Wehner. Zero-memory graph exploration with unknown inports, 2023. doi:10.48550/arXiv.2301.13860.
- 5 Stefan Dobrev, Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. Searching for a black hole in arbitrary networks: Optimal mobile agents protocols. *Distributed Computing*, 19:1–99999, 2006.
- 6 Stefan Dobrev, Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. Mobile search for a black hole in an anonymous ring. *Algorithmica*, 48:67–90, 2007. doi:10.1007/S00453-006-1232-Z.
- 7 Stefan Dobrev, Nicola Santoro, and Wei Shi. Using scattered mobile agents to locate a black hole in an un-oriented ring with tokens. *International Journal of Foundations of Computer Science*, 19(06):1355–1372, 2008. doi:10.1142/S0129054108006327.
- 8 Ajay D Kshemkalyani and Faizan Ali. Efficient dispersion of mobile robots on graphs. In *Proceedings of the 20th International Conference on Distributed Computing and Networking*, pages 218–227, 2019. doi:10.1145/3288599.3288610.
- 9 Ajay D Kshemkalyani, Anisur Rahaman Molla, and Gokarna Sharma. Efficient dispersion of mobile robots on arbitrary graphs and grids. *arXiv preprint arXiv:1812.05352*, 2018.
- 10 Ajay D Kshemkalyani, Anisur Rahaman Molla, and Gokarna Sharma. Fast dispersion of mobile robots on arbitrary graphs. In *International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics*, pages 23–40. Springer, 2019. doi:10.1007/978-3-030-34405-4_2.
- 11 Fukuhito Ooshita, Shinji Kawai, Hirotugu Kakugawa, and Toshimitsu Masuzawa. Randomized gathering of mobile agents in anonymous unidirectional ring networks. *IEEE Transactions on Parallel and Distributed Systems*, 25(5):1289–1296, 2013. doi:10.1109/TPDS.2013.259.
- 12 P. Panaite and A. Pelc. Exploring unknown undirected graphs. *Journal of Algorithms*, 33(2):281–295, 1999. doi:10.1006/JAGM.1999.1043.
- 13 Vyatcheslav B Priezzhev, Deepak Dhar, Abhishek Dhar, and Supriya Krishnamurthy. Eulerian walkers as a model of self-organized criticality. *Physical Review Letters*, 77(25):5079, 1996.
- 14 O. Reingold. Undirected connectivity in log-space. *Journal of the ACM (JACM)*, 55(4):1–24, 2008. doi:10.1145/1391289.1391291.
- 15 Masahiro Shibata, Toshiya Mega, Fukuhito Ooshita, Hirotugu Kakugawa, and Toshimitsu Masuzawa. Uniform deployment of mobile agents in asynchronous rings. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, pages 415–424, 2016. doi:10.1145/2933057.2933093.

- 16 Masahiro Shibata, Daisuke Nakamura, Fukuhito Ooshita, Hirotsugu Kakugawa, and Toshimitsu Masuzawa. Partial gathering of mobile agents in arbitrary networks. *IEICE TRANSACTIONS on Information and Systems*, 102(3):444–453, 2019. doi:10.1587/TRANSINF.2018FCP0008.
- 17 Masahiro Shibata, Yuichi Sudo, Junya Nakamura, and Yonghwan Kim. Uniform deployment of mobile agents in dynamic rings. In *International Symposium on Stabilizing, Safety, and Security of Distributed Systems*, pages 248–263. Springer, 2020. doi:10.1007/978-3-030-64348-5_20.
- 18 Yuichi Sudo, Daisuke Baba, Junya Nakamura, Fukuhito Ooshita, Hirotsugu Kakugawa, and Toshimitsu Masuzawa. A single agent exploration in unknown undirected graphs with whiteboards. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 98(10):2117–2128, 2015. doi:10.1587/TRANSFUN.E98.A.2117.
- 19 Yuichi Sudo, Fukuhito Ooshita, and Sayaka Kamei. Self-stabilizing graph exploration by a single agent. In *32nd International Colloquium on Structural Information and Communication Complexity (SIROCCO 2025)*, pages 384–400, 2025.
- 20 Yuichi Sudo, Masahiro Shibata, Junya Nakamura, Yonghwan Kim, and Toshimitsu Masuzawa. Near-Linear Time Dispersion of Mobile Agents. In *38th International Symposium on Distributed Computing (DISC 2024)*, pages 38:1–38:22, 2024. doi:10.4230/LIPIcs.DISC.2024.38.

A The Omitted Proof for Triangle-free Case

■ **Algorithm 4** Graph Exploration Algorithm For Triangle-free graphs $\mathcal{A}_{\text{TF5}} = (C_{\text{TF5}}, \text{init}, \xi_{\text{TF5}})$.

Colors: $C = \{\text{init}, \text{path}, \text{neigh}, \text{fin}, \text{head}\}$ and $c_0 = \text{init}$

Rule Function:

$$\begin{aligned}
 \xi_{\text{TF5}}(\text{init}, M) &= \begin{cases} (\text{path}, \text{init}) & \text{if } |M| > 0 \wedge \forall x \in M, x = \text{init} & \text{(Rule 1)} \\ (\text{neigh}, \text{head}) & \text{else if } \text{head} \in M \wedge \text{path} \in M & \text{(Rule 2)} \\ (\text{head}, \text{head}) & \text{else if } \text{head} \in M \wedge \text{path} \notin M & \text{(Rule 3)} \\ (\text{head}, \text{stay}) & \text{otherwise} & \text{(Rule 4)} \end{cases} \\
 \xi_{\text{TF5}}(\text{head}, M) &= \begin{cases} (\text{head}, \text{neigh}) & \text{if } \text{head} \in M \wedge \text{neigh} \in M & \text{(Rule 5)} \\ (\text{path}, \text{head}) & \text{else if } \text{head} \in M \wedge \text{path} \in M & \text{(Rule 6)} \\ (\text{fin}, \text{head}) & \text{else if } \text{head} \in M & \text{(Rule 7)} \\ (\text{head}, \text{init}) & \text{else if } \text{path} \in M \wedge \text{init} \in M & \text{(Rule 8)} \\ (\text{head}, \text{path}) & \text{else if } \text{path} \in M & \text{(Rule 9)} \\ (\text{path}, \text{init}) & \text{else if } \text{init} \in M & \text{(Rule 10)} \\ (\text{fin}, \text{stop}) & \text{otherwise} & \text{(Rule 11)} \end{cases} \\
 \xi_{\text{TF5}}(\text{path}, M) &= (\text{head}, \text{head}) & \text{(Rule 12)} \\
 \xi_{\text{TF5}}(\text{neigh}, M) &= (\text{init}, \text{head}) & \text{(Rule 13)}
 \end{aligned}$$

Proof of Theorem 5. In a manner similar to Section 3.2, we define $\mathcal{A}_{\text{TF5}} = (C_{\text{TF5}}, \text{init}, \xi_{\text{TF5}})$ and show that \mathcal{A}_{TF5} simulates semi-DFS on triangle-free graphs. (The definition of \mathcal{A}_{TF5} is given in Algorithm 4.)

We follow an outline similar to that of Theorem 3.2. First, we redefine the notion of *regular*, with the modification that we use the term *head* instead of *head1*. Let the return path be $P = (u_1, u_2, \dots, u_k)$. Suppose there exists a node $u \in N(u_k)$ that is colored *init*. We distinguish between two scenarios based on the value of k .

If $k = 1$, then the node u necessarily belongs to $U(P, F)$. Hence, the agent visits the init-colored node in accordance with **Rule 10** and subsequently recolors it with **head**, as prescribed by **Rule 4**.

If $k > 1$, the agent visits a init-colored node according to **Rule 8**. Moreover, if the visited node u belongs to $U(P, F)$, it is recolored with **head**; otherwise, it is colored **neigh**. In either case, the agent returns to u_k , following **Rule 2** or **Rule 3**, respectively.

Next, we consider the following two cases:

Case 1: $U(P, F) \neq \emptyset$. Similarly to Theorem 3.2, once a node $u \in U(P, F)$ is reached, that node is recolored with **head** and the agent returns to u_k . Then, following **Rules 5 and 13**, every node in $N(u_k)$ that is colored **neigh** is reset to **init**, with the agent returning to u_k after each reset. Here, the triangle-free property of the graph is crucial: since any node colored **neigh** has exactly one adjacent node colored **head**, the return destination is uniquely determined. Finally, once all vertices of **neigh** have been restored to **init**, **Rule 6** directs the agent to change the color of u_k to **path** and, by moving to u , adds u to P , thereby extending the path.

Case 2: $U(P, F) = \emptyset$. If $k = 1$, then every node in $N(u_k)$ is colored **fin**, and according to **Rule 11**, the process terminates. In the remaining case where $k \geq 2$, every node in $N(u_k)$ that is not colored **path** is recolored with **neigh**. Moreover, since $k \geq 2$, there is exactly one node adjacent to u_k that is colored **path**. The agent moves to that node, recolors it with **head**, and returns to u_k . Afterwards, every node adjacent to u_k that is colored **neigh** is reset to **init** (again following **Rules 5 and 13**), and the agent returns to u_k . Owing to the triangle-free property of the graph, the unique **head** node to which the agent must return is unambiguously determined. Finally, according to **Rule 7**, the agent changes the color of u_k to **fin** and returns to u_{k-1} .

Therefore, it can be shown that Case 1 simulates the expand operation of the semi-DFS, while Case 2 simulates the remove operation. ◀

B Model Clarification

As mentioned in Section 1, there is an ambiguity in the original model of [2], specifically in Section 1.2 of that paper: it is unclear whether the agent may recolor a node from a non-initial color back to the initial color. In their description, the rule function – given the current node’s color and the multiset of neighbor colors – outputs a pair (x, y) in the range

$$\{1, 2, \dots, n_c\} \times (\{1, 2, \dots, n_c\} \cup \{\text{stop}\}),$$

where n_c is the number of colors *excluding* the initial color 0. Here, x specifies the new color for the current node, and y specifies which color to move to (or **stop** to terminate). However, this range contains a typo: the second component must include 0 so that the agent can move to a node still colored with the initial color, and the first component should also include 0, since Lemma 7 of the arXiv version [4] considers the agent leaving a node colored 0 without changing its color. We therefore conclude that the intended output range in the model of Böckenbauer, Frei, Unger, and Wehner [2] is

$$\{0, 1, \dots, n_c\} \times (\{0, 1, \dots, n_c\} \cup \{\text{stop}\}),$$

which coincides with the rule-function range used in this paper. (As noted in the footnote of Section 2, the **stay** option does not extend the model.) Under this corrected interpretation, the agent can indeed recolor a node from a non-initial color back to the initial color.