

Synchronization Helps Robots to Detect Black Holes in Directed Graphs^{*}

Adrian Kosowski^{1,2}, Alfredo Navarra³, and Cristina M. Pinotti³

¹ LaBRI - Université Bordeaux 1 - CNRS,

351 cours de la Libération, 33405 Talence cedex, France.

² Department of Algorithms and System Modeling, Gdańsk University of Technology,
Narutowicza 11/12, 80233 Gdańsk, Poland. Email: adrian@kaims.pl

³ Dipartimento di Matematica e Informatica, Università degli Studi di Perugia,
Via Vanvitelli 1, 06123 Perugia, Italy. Email: {navarra, pinotti}@dmi.unipg.it

Abstract. The paper considers a team of robots which has to explore a graph G where some nodes can be harmful. Robots are initially located at the so called *home base* node. The dangerous nodes are the so called *black hole* nodes, and once a robot enters in one of them, it is destroyed. The goal is to find a strategy in order to explore G in such a way that the minimum number of robots is wasted. The exploration ends if there is at least one surviving robot which knows all the edges leading to the black holes. As many variations of the problem have been considered so far, the solution and its measure heavily depend on the initial knowledge and the capabilities of the robots. In this paper, we assume that G is a directed graph, the robots are associated with unique identifiers, they know the number of nodes n of G (or at least an upper bound on n), and they know the number of edges Δ leading to the black holes. Each node is associated with a white board where robots can read and write information in a mutual exclusive way.

A recently posed question [Czyzowicz et al., *Proc. SIROCCO'09*] is whether some number of robots, expressed as a function of parameter Δ only, is sufficient to detect black holes in directed graphs of arbitrarily large order n . We give a positive answer to this question for the synchronous case, i.e., when the robots share a common clock, showing that $O(\Delta \cdot 2^\Delta)$ robots are sufficient to solve the problem. This bound is nearly tight, since it is known that at least 2^Δ robots are required for some instances. Quite surprisingly, we also show that unlike in the case of undirected graphs, for the directed version of the problem, synchronization can sometimes make a difference: for $\Delta = 1, 2$ robots are always sufficient and sometimes required to explore the graph regardless of whether synchronization is present; however, for $\Delta = 2$, in the synchronous case 4 robots are always sufficient, whereas in the asynchronous case at least 5 robots are sometimes required.

^{*} This work was done during a research collaboration supported by the Italian CNR Short-Term Mobility Program.

1 Introduction

The subject of exploring an unknown graph by means of mobile entities has been widely considered during the last years. The increasing interest to the problem comes from the variety of applications that it meets. In robotics, it might be very useful to let a robot or a team of robots exploring dangerous or impervious zones. In networking, software agents might automatically discover nodes of a network and perform updates and/or refuse their connections. In this paper we are interested in the exploration of a graph with faulty nodes, i.e. nodes that destroy any entering entity. Such nodes are called *black holes*, and the exploration of a graph in such kind of networks is usually referred as *black hole search*. In what follows, we refer to the mobile entities as robots. According to the assumed initial settings of the network, and the knowledge and the capabilities of the robots, many results have been provided. Pure exploration strategies, without dealing with black holes, has been widely addressed, see for instance [8, 11] and references therein. In this case, the requirement is usually to perform the exploration as fast as possible. When black holes are considered, along with the time (or equivalently the number of edge traversals) required for a full exploration, the main goal resides in minimizing the number of robots that may fall into some black hole. A full exploration means that, at the end, all the edges which do not lead to any black hole must be marked as safe edges.

1.1 Related work

In [2], the exploration subject to black holes has been restricted to tree topology in a synchronous setting. The authors show that 2 robots are enough for computing the full exploration when only one black hole is admitted. They provide a $\frac{5}{3}$ -approximation algorithm with respect to the required edge traversals. In [6], the attention is devoted to rings. The authors show that 2 robots are enough in order to detect one black hole and they provide an algorithm which requires $O(n \log n)$ moves, where n is the dimension of the ring. Other restrictions have been considered like the knowledge of a map in [7] where the authors provide a strategy for locating one black hole by means of 2 robots in $O(n + d \log d)$ moves, where n is the size of the network and d its diameter. Also in [10], the map is known and the authors show that the problem of locating one black hole by means of 2 synchronous robots is not polynomial time approximable within any constant factor less than $\frac{388}{389}$, and provide a 6-approximation algorithm. A slightly different problem has been considered in [1] where the network admits many black holes and if a robot gets destroyed into one, then such a black hole disappears but not the underlying node. In [4], one black hole is optimally located by 2 robots in $O(n)$ moves in the case the input graph has size n and its topology is among hypercubes, cube-connected cycles, star graphs, wrapped butterflies, chordal rings, multidimensional meshes or tori of restricted diameter. Most general results can be found in [5] where the asynchronous case is considered in presence of a single black hole without any limitation to the network topology. The authors show that $\Delta + 1$ robots are sufficient and require $\Theta(n)$

steps, with Δ being the maximum degree of the input network and n its size. The variant in which the robot entering the node cannot learn the link by which it has arrived was studied in [9]. For this scenario, a bound of $\frac{\Delta^2 + \Delta}{2} + 1$ on the number of required robots was established and shown to be tight for some instances of the problem.

So far, the input graph has been always considered as undirected. The additional property of having directed edges can only increase the difficulty of the problem as once traversed an edge, the robot cannot directly come back in general. The first results concerning directed graphs have been recently published in [3]. The authors have considered a setting where robots are associated with unique identifiers (IDs), they know the number of nodes n of the input directed graph $G = (V, A)$, and they know the number Δ of edges belonging to G leading to the black holes. If $BH \subset V$ is the set of black holes in G , the induced subgraph $G[V \setminus BH]$ is assumed to be strongly connected. Each node $v \in V$ is associated with a so called *white board* which is an available bounded memory (with size polynomial in n) accessible in a mutually exclusive way by the robots located at node v . The obtained results show a general case lower bound on the number of robots required in order to accomplish the black hole search.

Theorem 1 ([3]). *In both the asynchronous and synchronous models, 2^Δ robots are sometimes required to solve the directed black hole search problem, for all $\Delta \geq 1$.*

Moreover, in [3] it is shown that $2\Delta + 1$ robots are enough for solving the problem on planar graphs with the planar embedding known by the robots, and 2Δ robots are needed in the worst case.

1.2 Our results

In this paper, we extend the results on directed graphs by also separately considering the synchronous and the asynchronous cases, i.e. when the robots share or not a common clock. Under the same settings of [3], for the synchronous case we provide a general strategy which requires $O(\Delta \cdot 2^\Delta)$ robots. This answers the main question posed in [3], i.e. whether some number of robots, expressed as a function of parameter Δ only, is sufficient to detect black holes in directed graphs of arbitrarily large order n . We also provide a strategy for $\Delta = 1$ and $\Delta = 2$ which requires 2 and 4 robots, respectively. For the asynchronous case, we show that 2 robots are still sufficient when $\Delta = 1$, but for $\Delta = 2$ at least 5 robots are sometimes required.

1.3 Outline

The paper is organized as follows. In Section 2 we recall the most important assumptions regarding the directed black hole search problem, and introduce some further notation. In Section 3 we state the main positive results of the paper for synchronous robots, proving that $O(\Delta \cdot 2^\Delta)$ robots suffice to explore

any graph, and 4 robots suffice when $\Delta = 2$. The results for asynchronous robots for $\Delta = 1$ and $\Delta = 2$ are given in Section 4. Section 5 contains some concluding remarks.

2 The model

We apply essentially the same scenario as that introduced in [3]. The explored digraph $G = (V, A)$ contains a distinguished node called the *home base* (hb) from which all the robots start the exploration, and a distinguished set of black holes $BH \subseteq V \setminus \{hb\}$. All the black holes are sinks of the digraph (i.e., with an out-degree of 0); we remark that from the perspective of this work, setting the number of black holes equal to $|BH| = 1$ does not affect any of the stated results. It is assumed that the induced subgraph $G[V \setminus BH]$ is strongly connected. The set of arcs leading into some black hole, $BA = A \cap ((V \setminus BH) \times BH)$, is known as the set of *black hole arcs*. All robots know two parameters of the graph: its order $n = |V|$, and the total number of black hole arcs, $\Delta = |BA|$. A robot located at a node v has access to its own built-in memory, the white board associated with the node, a global timer counting the steps of the algorithm, and a local labeling of the arcs leaving node v (called an assignment of port numbers), which can be used to select the next arc in the robot's traversal. Note that nodes do not have unique identifiers, although such identifiers can potentially be created by the robots exploring the graph.

For a node v of a directed (multi)graph H , the number of arcs entering v is denoted by $\text{indeg}_H(v)$, and the number of arcs exiting v is denoted by $\text{outdeg}_H(v)$.

3 Upper bound in the synchronous model

Theorem 2. *There exists a strategy for solving all instances of the directed black hole search problem using $O(\Delta \cdot 2^\Delta)$ synchronous robots.*

Proof. The strategy works as follows. The robots are released from hb one-by-one, at regular intervals of a time steps (where $a \in \Theta(n^4)$ steps is a value which results from the subsequent formulation of the strategy), in such a way as to guarantee that within a steps, i.e. before the next robot is released, exactly one of the following events will occur:

- The currently active robot successfully completes the entire exploration, returns to hb , and terminates the algorithm.
- The currently active robot is destroyed in some black hole.

The procedure is divided into a certain number of phases, and the p -th phase consists of $(\Delta + 2) \cdot S_p$ robot releases, where S_p is a value which will be fixed further on in the analysis.

- In the first part of the p -th phase, S_p *searching robots* are released using a set of rules specified below.

- The second part of the phase consists of Δ iterations. In each iteration, S_p *cycle detection robots* are released which follow along a slight modification of the paths used by the searching robots from the first part of the phase.
- In the third part of the phase, S_p *cycle contraction robots* are released in a similar way as in one iteration of the second part of the phase.

Throughout the algorithm, each node v maintains a label $f(v)$ stored in its white board. Let the subgraph $H_{f(v)}$ of G consist of all nodes $u \in V \setminus BH$ having $f(u) = f(v)$, and all arcs $(u_1, u_2) \in A$ such that $f(u_1) = f(u_2) = f(v)$ and (u_1, u_2) has been visited by at least one robot. The procedure is defined in such a way that initially $f(v)$ is unique for all visited nodes (the definition of $f(v)$ for unvisited nodes is irrelevant), whereas each of the graphs $H_{f(v)}$ is strongly connected; moreover, for each node v , the structure of the graph $H_{f(v)}$ and its embedding in the ports of G is encoded on the white board of v . To avoid confusion, we will sometimes write $f^{(2)}$ and $H^{(2)}$ to denote the value of labels f and graphs H throughout the second part of the phase (the cycle detection robots do not update the labels), and likewise $f^{(3)}$ and $H^{(3)}$ to denote these graphs at the end of the whole phase (after releasing all the cycle contraction robots).

Procedure for the first part of the phase. The procedure for each searching robot r is defined as follows. Directly before and directly after traversing an arc, the robot writes this information on the white boards of the starting node and of the end node of the arc, respectively, together with its unique ID and number of the move in the robot's move sequence. The robot proceeds to make its next move as follows:

1. Let the robot be located at an arbitrary node of some graph H_i . Then, the robot explores all nodes of H_i , identifying all arcs of G which leave or enter these nodes and do not belong to H_i . If there exists an outgoing arc (u_1, v_1) which was used by some robot r' at time t_1 , and an incoming arc (v_2, u_2) which was used by the same robot r' at some later time t_2 , where $u_1, u_2 \in V(H_i)$ and $v_1, v_2 \notin V(H_i)$, then robot r proceeds to extend component H_i by including in it all the nodes on the route (v_1, \dots, v_2) used by the robot r' between times t_1 and t_2 . (Robot r' was either a previously released robot, or $r' = r$.) More precisely, robot r visits all the nodes of all the graphs $H_{f(v)}$, with $v \in (v_1, \dots, v_2)$, setting their labels to $f(v) := i$, and updating the stored graphs $H_{f(v)}$ accordingly.
2. When the extension from Step 1 can no longer be applied, the robot performs the next move as follows.
 - If there exists an arc leaving H_i which has never been visited by any robot, then the robot proceeds along this arc.
 - Otherwise, the robot proceeds along the arc leaving H_i which has not been used by any robot for the longest time. (Observe that since the extension from Step 1 can no longer be applied, each arc leaving H_i has been used by each robot at most once, so the robot simply chooses the arc leaving H_i used by the robot which was released the earliest.)

Note that the actions performed in Step 1 will never lead the robot into a black hole, and do not lead it out of the current explored strongly connected component, whereas either of the actions performed in Step 2 can potentially destroy the robot or lead it into a node which does not belong to component H_i .

Procedure for the second and third parts of the phase. Looking at the labels $f^{(2)}$ and graphs $H^{(2)}$, which are defined at the time of the destruction of the last searching robot, observe that each searching robot r can be seen as performing a traversal of some sequence of graphs $(H_{f_1}^{(2)}, \dots, H_{f_k}^{(2)})$, where a graph $H_{f_i}^{(2)}$ appears in the sequence if at least one of its nodes has been visited by robot r . This implies a natural precedence relation (\prec) between graphs $H_{f_i}^{(2)}$ and $H_{f_j}^{(2)}$, that is, we will write $H_{f_i}^{(2)} \prec H_{f_j}^{(2)}$ if for some searching robot r , graph $H_{f_i}^{(2)}$ appears before graph $H_{f_j}^{(2)}$ in the robot's sequence of visits; see Figure 1 for an illustration.




Fig. 1. (a) A schematic representation of the paths followed by three consecutive searching robots r_1 , r_2 , r_3 . The strongly connected components H_α , H_β , H_γ associated with respective black hole arcs are assumed to remain unchanged during the phase. (b) The corresponding graph of the precedence relation (\prec) on H_α , H_β , H_γ .

The goal of the cycle detection robots is to detect cycles of length at most Δ in the graph of relation \prec . To do this, we release the robots in such a way that after the l -th of Δ iterations, each node of graph $H_{f_j}^{(2)}$ stores a list of all f_i such that $H_{f_i}^{(2)} \prec^l H_{f_j}^{(2)}$, where (\prec^l) is the l -th power of relation (\prec) . This is achieved using a standard approach. With each graph $H_{f_i}^{(2)}$, we associate a precedence set P_{f_i} whose copies are stored in the white boards of all of the nodes of $H_{f_i}^{(2)}$, and initially $P_{f_i} = \emptyset$. In the l -th iteration, for each searching robot r , exactly

one cycle detection robot r' will retrace the route of r . Robot r' is defined so as to visit all nodes of all graphs in the same order $(H_{f_1}^{(2)}, \dots, H_{f_k}^{(2)})$ as robot r , with transitions between adjacent graphs $H_{f_i}^{(2)}$ performed using the same arcs as those visited by robot r (this can be achieved since robot r may be assumed to store the information about its next move on the white boards of the nodes it is passing through). Additionally, for each $H_{f_j}^{(2)}$, robot r' now appends to the contents of set P_{f_j} the contents of all sets P_{f_i} , for all $i < j$, and also adds to P_{f_j} the element f_i , for all $i < j$; all the copies of set P_{f_j} stored in the nodes of $H_{f_j}^{(2)}$ are updated accordingly. Clearly, we have $H_{f_i}^{(2)} \prec^l H_{f_j}^{(2)}$ if and only if $f_i \in P_{f_j}$. For easier manipulation of data in the next part of the phase, the precedence sets should store additional information about the ID of the searching robots r inducing the respective precedence relations.

Finally, the goal of the cycle contraction robots is to contract into a single connected component all graphs $H_{f_i}^{(2)}$ forming cycles of length at most Δ in the precedence graph induced by relation (\prec) . The cycle contraction once again retraces the route of corresponding searching robots from r . Upon detecting a cycle of length Δ in the precedence relation, i.e., $f_i \in P_{f_i}$ for some graph $H_{f_i}^{(2)}$, they follow the arcs of the directed cycle of G which induces this relation (the identifiers of these arcs should be stored in set P_{f_i} as auxiliary information). For all the nodes v of all the graphs $H_{f_j}^{(2)}$ encountered in this cycle, the robots then update their labels, setting the new value $f^{(3)}(v) = f_i$ and constructing the new contracted component $H_{f_i}^{(3)}$ accordingly. These labels and components are then used by the searching robots in the next phase of the algorithm.

Termination condition. The algorithm terminates once the strongly connected component containing hb , $H_{f(hb)}$ has exactly Δ arcs exiting it and no arcs entering it. Any searching robot which is released can of course detect this immediately, and return with this information to hb .

Analysis. To complete the proof, we will show that, when the number of searching robots in the p -th phase is set as $S_p = \lceil 2^\Delta / \sqrt{3}^{p-1} \rceil$, the algorithm will always terminate within Δ phases. Let $\{v_1, \dots, v_\Delta\}$ be the set of vertices from which it is possible to exit by a black hole arc. At any given time step, consider the set of labels $\{f(v_1), \dots, f(v_\Delta)\} = \{f_1, f_2, \dots, f_d\}$, with $1 \leq d \leq \Delta$, and let δ_i denote the number of black hole arcs exiting from nodes of H_{f_i} , with $\sum_{1 \leq i \leq d} \delta_i = \Delta$. As in the case of the other parameters, we will use the notation $d^{(2)}$, $d^{(3)}$, $f_i^{(2)}$, $f_i^{(3)}$, and $\delta_i^{(2)}$, $\delta_i^{(3)}$ to denote the respective values during the second part and at the end of the third part of the phase.

First, we will prove by induction that at the end of the p -th phase, $1 \leq p < \Delta$, we have $d^{(3)} \leq \Delta - p$. The claim is clearly true for $p = 0$; we will show that in each subsequent phase, the value of $d^{(3)}$ decreases by at least 1, or equivalently, that some two components H_{f_i} and H_{f_j} are contracted into each other during each phase. Consider for a moment the situation during the second part of the

phase, and define the directed multigraph (\mathcal{H}, R) , where the set of nodes is $\mathcal{H} = \{s, H_{f_1}^{(2)}, H_{f_2}^{(2)}, \dots, H_{f_d}^{(2)}\}$, with s being a special source node, while arcs are defined as follows. For each searching robot r released in the current phase, we add one arc from node $H_{f_i}^{(2)}$ to $H_{f_j}^{(2)}$ if and only if robot r visited at least one node of $H_{f_i}^{(2)}$ before visiting a node of $H_{f_j}^{(2)}$, and did not visit any node of any other graph from $\mathcal{H} \setminus \{H_{f_i}^{(2)}, H_{f_j}^{(2)}\}$ in between these two visits. For each robot r , we also add an arc from the special source node s to the first of the graphs $H_{f_i}^{(2)}$ which the robot visits. Note that $\text{indeg}_{(\mathcal{H}, R)}(s) = 0$, and $\text{outdeg}_{(\mathcal{H}, R)}(s) = S_p$.

We consider the following cases.

(a) Multigraph (\mathcal{H}, R) contains a directed cycle. Then, since the relation R restricted to nodes from $\mathcal{H} \setminus \{s\}$ is a sub-relation of precedence relation (\prec) , and $|F| \leq \Delta$, the graph of relation (\prec) also contains a cycle of length at most Δ on the connected components from \mathcal{H} . Since all such components are contracted into one in the cycle contraction phase, we immediately obtain that the value of $d^{(3)}$ is smaller than that of $d^{(2)}$ for the current phase, hence $d^{(3)}$ decreases with respect to the previous phase.

(b) Multigraph (\mathcal{H}, R) is acyclic, and no contractions are performed ($d^{(2)} = d^{(3)}$); we can then extend R to a linear order, simply writing $f_1 < f_2 < \dots < f_d$; each searching robot of the current phase visits some subsequence of graphs $H_{f_i}^{(2)}$ in ascending orders of labels. For each f_i , consider the set of black hole arcs exiting the component $H_{f_i}^{(2)}$. If at the start of the current phase some of these arcs exited different strongly connected components H_a, H_b (which were contracted into one by some searching robot), then the inductive claim holds, since for the current phase the value of $d^{(2)}$ is less than that of $d^{(3)}$ at the end of the previous phase. Thus, w.l.o.g. we can assume that $H_{f_i}^{(2)}$ is obtained from some component H_{f_i} , which is perhaps enlarged by searching robots, but has the same set of exiting black hole arcs as $H_{f_i}^{(2)}$, throughout the phase. With this assumption, we now proceed to show the following auxiliary lemma.

Lemma 1. *We either have $d = 1$ and $H_{f_1}^{(2)} = G[V \setminus BH]$, or for all i , $1 \leq i \leq d^{(2)}$, the following condition holds:*

$$\text{outdeg}_{(\mathcal{H}, R)} H_{f_i}^{(2)} \geq \left\lfloor \frac{\text{indeg}_{(\mathcal{H}, R)} H_{f_i}^{(2)}}{\delta_i + 1} \right\rfloor.$$

Proof. Let i be arbitrarily chosen, and consider an arbitrary sequence of $\delta_i + 1$ robots $r_1, \dots, r_{\delta_i + 1}$ which consecutively went through $H_{f_i}^{(2)}$, ordered by increasing release times. We will show that at least one of the robots did not use a black hole arc to leave $H_{f_i}^{(2)}$. Let $H_{f_i}(r_j)$ denote component H_{f_i} directly after the destruction of robot r_j ; recall that we have $H_{f_i}(r_1) \subseteq H_{f_i}(r_2) \subseteq \dots \subseteq H_{f_i}(r_{\delta_i + 1}) \subseteq H_{f_i}^{(2)}$, and each $H_{f_i}(r_j)$ has the same set of black hole arcs as $H_{f_i}^{(2)}$. Suppose now, to the contrary, that each of the robots r_j left $H_{f_i}^{(2)}$ by a

black hole arc; this means that each of these robots also left $H_{f_i}(r_j)$ by a black hole arc; moreover, since preference is given to arcs which were not used for a longer time, each of the robots r_1, \dots, r_{δ_i} will be destroyed by a different black hole arc. Consider now the robot r_{δ_i+1} . By the strong connectivity property of graph $G[V \setminus BH]$, unless $H_{f_i}(r_{\delta_i+1}) = G[V \setminus BH]$, there exists at least one arc e outgoing from $H_{f_i}(r_{\delta_i+1})$ which is not a black hole arc. This arc could not have been visited by any robot r released between r_1 (inclusive) and r_{δ_i+1} (exclusive). Indeed, r never returns to component $H_{f_i}(r_{\delta_i+1})$ after leaving it by arc e (since otherwise arc e would have been included into $H_{f_i}(r_{\delta_i+1})$ by robot r_{δ_i+1}), and since we have $r \in \{r_1, \dots, r_{\delta_i}\}$ (because r visits $H_{f_i}(r_{\delta_i+1}) \subseteq H_{f_i}^{(2)}$), one of the robots $\{r_1, \dots, r_{\delta_i}\}$ would have not been destroyed by any black hole arc leading out of $H_{f_i}(r_{\delta_i+1})$, a contradiction. Hence, robot r_{δ_i+1} finally leaves component $H_{f_i}(r_{\delta_i+1})$ by an arc which is not a black hole arc, so it must contribute to the out-degree of node $H_{f_i}^{(2)}$ in multigraph (\mathcal{H}, R) . Since the sequence of $\delta_i + 1$ consecutive robots going through $H_{f_i}^{(2)}$ was arbitrarily chosen, the claim follows directly.

Now, we make a simple structural claim which relies only on properties of directed acyclic multigraphs.

Lemma 2. *Let M be any directed acyclic multigraph having set of nodes $\{s, v_1, v_2, \dots, v_d\}$, such that $\text{indeg}(s) = 0$, and for all v_i , $1 \leq i \leq d$, we have: $\text{outdeg}_M(v_i) \geq \left\lfloor \frac{\text{indeg}_M(v_i)}{\delta_i + 1} \right\rfloor$. Then: $\text{outdeg}(s) < \prod_{i=1}^d (\delta_i + 1)$.*

The proof of the lemma proceeds by induction with respect to d ; we omit the details. We also make one more simple arithmetic observation, which we leave without proof.

Lemma 3. *Let $\{\delta_1, \dots, \delta_d\}$ be positive integers such that $\sum_{i=1}^d \delta_i = \Delta$. Then $\prod_{i=1}^d (\delta_i + 1) \leq 2^\Delta / \sqrt{3}^{\Delta-d}$.*

Now, set the number of searching robots released in the current p -th phase as $S_p = \lceil 2^\Delta / \sqrt{3}^{p-1} \rceil$. Since we have that $d \leq \Delta - p + 1$ (by the inductive claim from phase $p - 1$), we obtain from the preceding lemma: $\text{outdeg}_{(\mathcal{H}, R)}(s) = S_p \geq 2^\Delta / \sqrt{3}^{\Delta-d} \geq \prod_{i=1}^d (\delta_i + 1)$. This means that the claim of Lemma 2 does not hold, and so the assumption $\text{outdeg}_{(\mathcal{H}, R)} H_{f_i}^{(2)} \geq \left\lfloor \frac{\text{indeg}_{(\mathcal{H}, R)} H_{f_i}^{(2)}}{\delta_i + 1} \right\rfloor$ is violated for some component $H_{f_i}^{(2)}$. From Lemma 1 we thus immediately conclude that $d = 1$ and $H_{f_1}^{(2)} = G[V \setminus BH]$, hence in this case one of the robots successfully solves the black hole search problem.

Finally, we remark that if the algorithm reaches phase Δ , we have $d = 1$, and so only one strongly connected component H_{f_1} . Since the algorithm cannot perform any further contractions, by the same argument, $S_\Delta = \Delta + 1$ searching robots in this phase are sufficient to complete the exploration of the graph.

To finish the proof, observe that the total number of robots used during all the phases of the algorithm is at most: $\sum_{p=1}^{\Delta} (\Delta+2)S_p = (\Delta+2) \sum_{p=1}^{\Delta} \lceil 2^{\Delta}/\sqrt{3}^{p-1} \rceil \in O(\Delta \cdot 2^{\Delta})$.

Theorem 3. *In the synchronous model, 4 robots are always sufficient and sometimes required to solve the directed black hole search problem with $\Delta = 2$.*

Proof. The lower bound on the number of required robots is a direct consequence of Theorem 1. The upper bound is obtained by a similar strategy as that used in the proof of Theorem 2, simply releasing 4 searching robots r_1, r_2, r_3, r_4 in a single phase (without any cycle detection or cycle contraction robots). Let $\{e_1, e_2\} = BA$ be the two black hole arcs, originating from nodes v_1 and v_2 respectively. Robot r_1 may either successfully solve the black hole search instance, or may be destroyed on one of the black hole arcs, say, e_1 . In the latter case, robot r_2 will either successfully solve the instance or be destroyed in the other black hole arc e_2 (recall that robots choose to use unvisited arcs, whenever possible, and the graph is strongly connected). Suppose that robot r_3 is also destroyed in some black hole. Now, taking into account Lemma 2, we observe that at some point during the traversal performed by robot r_4 , nodes v_1 and v_2 will necessarily belong to the same strongly connected component H_f ($f(v_1) = f(v_2) = f$). Once robot r_4 has reached this component, it will always exit it by unvisited arcs. Since both arc e_1 and e_2 have already been used (by robots r_1 and r_2 , respectively), robot r_4 will never enter a black hole and will proceed to enlarge component H_f , until the whole of the graph has been explored, $H_f = G[V \setminus BH]$, and will thus successfully complete the task.

4 Results for the asynchronous model

Theorem 4. *In both the asynchronous and synchronous models, 2 robots are always sufficient and sometimes required to solve the directed black hole search problem with $\Delta = 1$.*

Proof. The lower bound on the number of required robots is a direct consequence of Theorem 1. The upper bound (which we obviously need to show for the asynchronous model, only) is obtained by simultaneously releasing two robots, each of which applies a strategy identical to that used by searching robots in the proof of Theorem 2. It is easy to observe that, throughout the process, each robot is either traversing arcs of an already identified strongly connected subgraph H_f of $G[V \setminus BH]$, or is exiting such a subgraph by some outgoing arc e . Arc e has either never been visited before by any robot, or is the unique arc exiting H_f , which, due to the strong connectivity property of $G[V \setminus BH]$, has to lead to some node of $V \setminus BH$. Thus, a robot may potentially enter a black hole only when using an arc which has been never visited before by any robot. Since $\Delta = 1$, this means that the black hole may destroy at most one robot, since the remaining robot will not traverse the unique black hole arc at any later time. The surviving robot(s) perform a search of the graph, exploring unvisited arcs, until the whole

of the strongly connected component $G[V \setminus BH]$ has been discovered and the search is complete.

Theorem 5. *In the asynchronous model, at least 5 robots are sometimes required to solve the directed black hole search problem with $\Delta = 2$.*

Proof. The proof considers some possible cases that must be considered by any strategy \mathcal{A} in order to solve the directed black hole search problem with $\Delta = 2$. We first assume that 4 robots are sufficient for solving the problem and then we obtain the claim by showing how an adversary can easily defeat all the robots. As usual, all the 4 robots start from hb which has two outgoing arcs. According to \mathcal{A} , there are three possible cases for the first move of each robot (see Figure 2 for a visualization of the corresponding strategies of the adversary): (i) all of them follow one direction; (ii) one goes through one direction and the remaining three to the other; (iii) the four robots are equally distributed among the two directions.




Fig. 2. The three possible ways that an adversary has to defeat any strategy which aims to solve the directed black hole search problem when using less than 5 robots. Dashed arcs will be never traversed because of the adversary.

If (i), then the adversary can easily locate one black hole on the reached node, and all the robots get destroyed. If (ii), then the adversary can locate one black hole on the node reached by the majority of the robots, hence destroying three of them. The surviving one reaches a safe node admitting another branching off. One of the two options leads to the black hole, while the other one to hb . As the robot has no further knowledge to distinguish among the two options, the branch it chooses will be set by the adversary as the one leading to the black hole, and again all the robots get destroyed. If (iii), by referring to Figure 2 we have that after the first move, two robots are on node u and two robots are on node v . Note that since from the point of view of the robots located at u (respectively, v), this situation is indistinguishable from that in graph (ii), where all but one

of the robots are destroyed; hence, no robot can wait at u (resp. v), since this could lead to an infinite deadlock. Now, the pair of robots on u (resp. v) cannot choose to move along the same arc because they may both be destroyed by the adversary. So, they have to move on distinct arcs. This implies that one robot from u and one from v get destroyed. The snapshot of the network now gives a robot on v and another on z . From v , both outgoing arcs have been already traversed and hence the robot has no further knowledge to distinguish among the two options apart from the IDs of the robots that have been passing through such arcs. However, these IDs can be exchanged by the adversary by swapping the labels of ports exiting v , thus making this information useless. The robot is thus forced either to wait at v , or enters the black hole. The other robot, which is on z , has again two unexplored options, one leading to hb , and the other to u . The adversary forces the robot to reach u , from there it is analogously forced either to wait or to enter the black hole. Thus, all the robots are either destroyed or deadlocked, and the claim holds.

5 Conclusions

In this paper, we have considered the black hole search in directed graphs under both the synchronous and the asynchronous settings. The obtained results reflects the difficulty of the problem on directed graphs as well as the difference between the synchronous and the asynchronous settings. It is worth noting how the behavior of the robots must change even for small values of Δ like 2. This was not the case for undirected graphs. This makes the study of the problem on directed graphs even more intriguing. One major remaining open problem concerns the possibility of bounding the number of required robots in terms of Δ in the asynchronous case.

Many variations of the problem still deserve investigation. Changing the assumptions to the initial knowledge of the robots, for instance, could lead to completely different results.

References

1. C. Cooper, R. Klasing, and T. Radzik. Locating and Repairing Faults in a Network with Mobile Agents. In Proc. of 15th International Colloquium on Structural Information and Communication Complexity (SIROCCO), pages 20–32, 2008.
2. J. Czyzowicz, D. Kowalski, E. Markou, and A. Pelc. Searching for a black hole in tree networks. In Proc. 8th International Conference on Principles of Distributed Systems (OPODIS), pages 35–45, 2004.
3. J. Czyzowicz, S. Dobrev, R. Kralovic, S. Miklik, and D. Pardubska. Black Hole Search in Directed Graphs. In Proc. of 16th International Colloquium on Structural Information and Communication Complexity (SIROCCO), to appear, 2009.
4. S. Dobrev, P. Flocchini, R. Kralovic, G. Prencipe, P. Ruzicka, and N. Santoro. Black hole Search in Common Interconnection Networks. *Networks*, 47(2):61–71, 2006.

5. S. Dobrev, P. Flocchini, G. Prencipe, and N. Santoro. Finding a black hole in an arbitrary network: optimal mobile agents protocols. In Proc. of 21st ACM Symposium on Principles of Distributed Computing (PODC), pages 153–162, 2002.
6. S. Dobrev, P. Flocchini, G. Prencipe, and N. Santoro. Mobile search for a black hole in an anonymous ring. *Algorithmica*, 48(1):67–90, 2007.
7. S. Dobrev, P. Flocchini, and N. Santoro. Improved bounds for optimal black hole search in a network with a map. In Proc. of 11th International Colloquium on Structural Information and Communication Complexity (SIROCCO), pages 111–122, 2004.
8. L. Gaśieniec, R. Klasing, R.A. Martin, A. Navarra and X. Zhang. Fast periodic graph exploration with constant memory. *Journal of Computer and System Sciences (JCSS)*, 74(5):802–822, 2008.
9. P. Glaus. Locating a Black Hole without the Knowledge of Incoming Link. In Proc. of 5th International Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS), to appear, 2009.
10. R. Klasing, E. Markou, T. Radzik, and F. Sarracco. Approximation bounds for Black Hole Search problems. *Networks*, 52(4):216–226, 2008.
11. A. Kosowski and A. Navarra. Graph Decomposition for Improving Memoryless Periodic Exploration. In Proc. of 34th International Symposium on Mathematical Foundations of Computer Science (MFCS), pages 501–512, 2009.