



KATEDRA INFORMATIKY
FAKULTA MATEMATIKY, FYZIKY A INFORMATIKY
UNIVERZITA KOMENSKÉHO, BRATISLAVA

MOBILE AGENTS IN SEARCH FOR A FAULTY NODE

(Master Thesis)

Peter Glaus

Program: Informatika 9.2.1

Advisor: doc. RNDr. Rastislav Královič, PhD.

Bratislava, 2009

I hereby declare that I wrote this thesis by myself,
only with the help of the referenced literature, under
the careful supervision of my thesis advisor.

.....

Acknowledgements

For useful advice, valuable consultations and linguistic support, I would like to thank my thesis advisor doc. Rastislav Královič.

I would also like to thank my parents for their support and my girlfriend Baška for the help and motivation she provided.

Abstract

Author: Peter Glaus
Title: Mobile Agents in Search for a Faulty Node
Advisor: doc. RNDr. Rastislav Kráľovič, PhD.
University: Univerzita Komenského
Faculty: Fakulta Matematiky Fyziky a Informatiky
Department: Katedra informatiky
Number of pages: 47
Year: 2009

We study a group of mobile agents operating on an arbitrary unknown distributed system. One of the nodes of the distributed system is extremely harmful and destroys any incoming agent without notification. The task of exploring the distributed system and locating the harmful node, *Black Hole Search*, has been studied with various modifications.

We are studying the effects of the *knowledge of incoming link* on the size of the optimal solution. When an agent enters a node, the information which port leads back can be given to it. We refer to this as to the *knowledge of incoming link*. In previous research, it was always assumed that the agent is given this information.

In this paper we study arbitrary, unknown distributed systems without the knowledge of incoming link. Agents are asynchronous and they communicate via whiteboards. We present a lower bound on the size of the optimal solution, proving that at least $\frac{\Delta^2 + \Delta}{2} + 1$ agents are necessary to locate the black hole. We provide an algorithm for the Black Hole Search without the knowledge of incoming link as well. We prove that this algorithm is correct, and that it uses $\frac{\Delta^2 + \Delta}{2} + 1$ agents, thus providing optimal solution.

KEYWORDS: distributed system, mobile agents, faulty node, black hole search, knowledge of incoming link

Abstrakt

Autor:	Peter Glaus
Názov práce:	Mobile Agents in Search for a Faulty Node
Školiteľ:	doc. RNDr. Rastislav Kráľovič, PhD.
Univerzita:	Univerzita Komenského
Fakulta:	Fakulta Matematiky Fyziky a Informatiky
Katedra:	Katedra informatiky
Počet strán:	47
Rok:	2009

V práci sa zaoberáme algoritmami pre agentov prehľadávajúcich distribuovaný systém. Jeden z uzlov distribuovaného systému je chybný, bez varovania zničí každého agenta, ktorý do neho vstúpi. Úloha prehľadávania distribuovaného systému a lokalizovania chybného uzla, *Hľadanie čiernej diery*, bola preštudovaná vo viacerých modifikáciách.

V tejto práci skúmame dôsledky *znalosti príchodzej hrany* na počet agentov potrebných na vyriešenie úlohy. Keď agent vstúpi do uzla, informácia o tom, ktorá hrana vedie do predchádzajúceho uzla, mu môže byť poskytnutá. Túto informáciu nazývame *znalosť príchodzej hrany*. Vo všetkých predošlých štúdiách bola táto znalosť agentom automaticky poskytnutá.

Zameriavame sa na problém hľadania čiernej diery v náhodných, neznámych distribuovaných systémoch bez znalosti príchodzej hrany. Uvažujeme asynchrónnych agentov komunikujúcich pomocou tabúľ. Dokázali sme spodnú hranicu počtu agentov potrebných na nájdenie čiernej diery. Ľubovoľný korektný algoritmus musí v najhoršom prípade použiť aspoň $\frac{\Delta^2 + \Delta}{2} + 1$ agentov.

Taktiež sa nám podarilo nájsť algoritmus na hľadanie čiernej diery v takomto distribuovanom systéme. O tomto algoritme sme ukázali, že používa $\frac{\Delta^2 + \Delta}{2} + 1$ agentov v najhoršom prípade, čo znamená, že je optimálnym riešením.

Kľúčové slová: distribuovaný systém, mobilní agenti, chybný uzol, hľadanie čiernej diery, znalosť príchodzej hrany

Contents

1	Introduction	2
1.1	Black Hole Search	2
1.2	Aim of this Thesis	3
2	Definitions	4
2.1	Computational Model	4
2.2	The Problem	6
2.3	Variations of the Black Hole Search Problem	7
3	Previous Results	11
3.1	General Statements and Ideas	11
3.2	The Effects of Prior Knowledge of the Distributed System	13
3.3	Other Results	18
4	The Knowledge of Incoming Link	21
4.1	Proposed Model	21
4.2	Lower Bound on the Size of the Optimal Solution	23
5	Size-Optimal Algorithm	26
5.1	Algorithm	27
5.2	Correctness	33
5.3	Upper Bound on the Size of the Optimal Solution	38
	Conclusion	40
	Bibliography	42

Chapter 1

Introduction

1.1 Black Hole Search

Special part of theoretical computer science concerning multi-processor or network environments is the field of Distributed algorithms. Distributed algorithms can be studied through a very intriguing agent-based model. In this model, the distributed system can be regarded as a kind of maze consisting of nodes connected with passages and in this maze is a group of autonomous mobile agents trying to fulfill a given task. In real world, the nodes would be computers in a computer network and the agents would be represented as processes executed on these computers, the task might be finding the best routing path.

Agent-based distributed algorithms have been widely studied. Problems such as graph exploration, finding the shortest path, finding a spanning tree, gathering of agents, leader election and other were already tackled. There are many modifications of these tasks based on agents' knowledge and properties of the distributed system. In most cases though, the distributed system is considered non-faulty. We focus our study on distributed systems containing a *black hole*.

Black hole is a harmful node which destroys all agents that enter it. It is indistinguishable from the outside and once an agent enters, it is unable to send a message or warn other agents. *Black Hole Search* is a task in which the agents have to create a map of the distributed system with the location of the black hole pointed out.

1.2 Aim of this Thesis

The *Black Hole Search Problem* (BHS) is the problem of finding an algorithm for the agents to fulfill the Black Hole Search. The solution to this problem and its efficiency can vary a lot and are dependent on the distributed system and the agent model properties. In unoriented arbitrary unknown graphs the optimal algorithm needs $\Delta + 1$ agents in the worst case[7], whereas in oriented graph scenario, the number of necessary agents can be up to 2^Δ [3]. In this work we want to study this gap by proposing a new model which has some restrictions similar to the oriented graph scenario.

In all previous studies of BHS, it was anticipated that when an agent enters a node, it automatically knows from which link it entered the node. We call this information the *knowledge of the incoming link*. This information is not inherently present in all distributed systems. For instance, one can think of a software agent that starts its execution on a host without any information about the incoming port.

We study the effects of the knowledge of incoming link on unoriented arbitrary unknown distributed systems. We show that without this knowledge, the problem solution needs more agents to work correctly. We show that there are $\Theta(n^2)$ agents needed for an algorithm that solves BHS on arbitrary unknown graphs, more precisely $\frac{n^2+n}{2} + 1$ agents are necessary and sufficient.

Outline of this Thesis

The basic definitions of BHS are presented in the chapter 2. We also note possible modifications of the distributed systems and mention the influence on the problem solution. We summarize the most important results provided by previous studies in the chapter 3. We focus on studies of BHS on arbitrary graphs with various degrees of agents' prior knowledge.

In the chapter 4 we present the model of a distributed system in which the agents lack the knowledge of incoming link. The lower bound on the number of agents needed by any general problem solution is provided as well. An algorithm for BHS without the knowledge of incoming link with a proof of its size-optimality is presented in the chapter 5.

Summarization of our work, possible future work and conclusion follow in the last chapter.

Chapter 2

Definitions

2.1 Computational Model

We start by defining properties of the computational model that are similar in most of BHS studies and will be used in this work as well. The model consists of a distributed system, which can be simply described as a graph, a group of agents acting according to a predefined algorithm and one or more black holes, which are harmful nodes. Roughly said, our goal is to design such algorithm for the agents, that the agents together fulfill a predefined task, exact definition is stated below.

Distributed System

Distributed system is a graph $G = (V, E)$ with its port labeling. We mostly denote vertices as *nodes* and edges as *links*. The size of the distributed system is the number of the nodes of the graph $n = |V|$. The maximal degree of a node is denoted Δ .

We also use the term *port*, which denotes the places where a link is connected to a node. Link (u, v) has two ports, one is in the node u and the other is in the node v .

The graphs we use are simple, without double edges or loops of type (u, u) . As we consider only undirected graphs the link (u, v) is the same as (v, u) . We expect the graph to be biconnected, otherwise finding the black hole might be impossible[7].

The use of the port labeling is necessary for the purpose of clarity. Link (u, v) has a port in node u , its label is denoted as $\lambda_u(u, v)$. The set of

all port labels in a node v is denoted as λ_v and the whole port labeling is $\lambda = \{\lambda_u | \forall u \in V\}$.

Agents

A group of agents is initially placed in a random node of the distributed system. The node where the agents are at the beginning is called *home base*. The agents are autonomous and proceed according to the same algorithm. We focus on the asynchronous agents only. In general, the agents are unaware of each other, although they can communicate by means of messages or by changing state of a node either using whiteboards or tokens.

An agent's behavior can be divided into basic steps. In one step agent checks the state of a current node, chooses some action, modifies the state of the node or communicates and, at last, it performs the action. The action can be either entering a specific port of the node, staying in the current node and waiting for some event to happen or finishing.

Agents have to fulfill a specific task, in our case it is to find a black hole. By finding a black hole we mean that agents that decide to take the action *finish* should have a full map of the graph with the position of the black hole marked down, we will call this task *Black Hole Search*.

Black Hole

There is a harmful node in the distributed system. Every agent that enters a port of a link leading to such node can be considered dead or not functional any more. Once an agent enters it, there are no means of communication with other agents. The distributed system itself carries no information whether link leads into such harmful node or not. We denote such node as *black hole*. Without loss of generality, we assume the degree of the black hole to be the maximal degree of a node — Δ . Throughout this work we focus on distributed systems with one black hole only, even though models with multiple black holes can be considered as well.

For agents that did not enter the black hole we say that they are alive. On the other hand, we will use the term dead or we say that agent died when it entered the black hole.

2.2 The Problem

Agents' task is to explore the graph and locate the black hole. All agents act according to the predefined algorithm. We say that the algorithm ends when all the remaining agents take action *finish*. We use the following definition when considering the correctness of the algorithm:

Definition 1 (Correctness). *The algorithm is correct if for any distributed system the algorithm ends with at least one agent alive and all surviving agents have a correct map of the distributed system with correctly marked location of the black hole.*

The problem of finding such algorithm is denoted as *Black Hole Search Problem*.

Solution Evaluation

To allow us comparison of BHS solutions we have to be able to measure the efficiency of the algorithms. The first measure of the efficiency is the number of agents needed by the algorithm to work correctly, it is denoted as the *size* of the solution.

The second measure, the *cost* of the solution, is the number of steps that agents make in total before finishing the task. The size is much more important than the cost and thus the focus is on using the optimal number of agents first and then improving the number of steps made by the agents.

When considering the size of an algorithm on a specific distributed system, we can simply look at the number of agents that enter the black hole in the worst case for that particular distributed system. We can expect the number of agents needed by the algorithm as the number of agents that die in the worst case plus one, otherwise it could be improved.

Adversary

When evaluating solutions, we always consider the worst case scenario. Instead of referring to the worst case we use the concept of adversary. Adversary is a kind of powerful opponent, his goal is to either make the algorithm fail the task or to work least efficiently as possible.

Firstly, when the agents are not able to distinguish between two ports, then the adversary can decide where these ports lead to, possibly make them

lead to the black hole. Secondly, adversary can slow down agents traveling through some link so an agent waiting for them is not able to decide whether they fell into the black hole or not.

For more details we advice the reader to [7, 6].

2.3 Variations of the Black Hole Search Problem

Agents Behavior

Most important property of the whole computational model is *synchronous-ness* or the fact whether agents' steps are somehow synchronized or they are totally asynchronous. We say that the agents are synchronized when every step can maximally last a given amount of time. The synchronization of the agents provides great advantage. The algorithm can be designed in such manner that the agents can easily prove whether node is dangerous or not. Every agent has to return from a newly discovered node immediately, thus when an agent does not return from an unknown node after two steps, the other agents know that the node is in fact the black hole. As we will show in the Lemma 3 provided by [7], it is not possible in asynchronous model.

In publications related to the synchronous model [13, 2, 14] the Black Hole Search Problem is shifted from studying what is the asymptotically optimal solution, to the study of how hard it is to find exact optimal solution for a particular distributed system. In such scenario the problem of finding optimal traversing algorithm with consideration of the black hole is shown to be NP-Hard [13]. Our work is focused on asynchronous models only and thus we will be coping with uncertainty and looking for the best universal solution.

Agents Communication

Without some kind of communication between the agents, the problem would be unsolvable. If they could not send some notice to each other, all agents could make the same first move and end up in the black hole. We mention two types of communication, both based on the same principle. Agents communicate only by changing states of the nodes.

Whiteboards

First type of communication uses *whiteboards*. In every node there is a whiteboard and all agents present in the node can read it and write on it. We can easily imagine this in the following manner: when an agent comes into the node it looks on the whiteboard and before taking the next action it can leave any message on the whiteboard. The important fact is that we expect atomicity of agents behavior and exclusive access to the whiteboards. We assume that from the point when an agent starts reading it to the point when it finishes writing or decides not to write on it, no other agent can use the whiteboard.

Whiteboards are being used for various purposes. As we will show later, the most common purpose is to let the other agents know which port did an agent enter. Agents can also "draw" a map of the distributed system, or at least a part of it that they know. The whiteboards are the only way of communication, thus when an agent decides to wait in a node for some event to happen, the only event it can notice is the change of the whiteboard.

Tokens

When using whiteboards, one could also study how much space is needed on the whiteboards for the agents to communicate. With the use of whiteboards with constant amount of space we get very close to the *token* model. In this type of communication every agent can carry a specific amount of tokens. Each node has one or more places where the tokens can be laid. Usually, the tokens can be laid in front of each port and also in the middle of the node. An agent can see all the tokens that are laid down in the node, it can lay down a token and of course pick up tokens and carry them. Once again, we expect atomicity and exclusive access. Between the time when an agent reads the state of the node and the time the agent changes it, no other agent can "access" the state of the node.

The token model of communication is much more restricting than using whiteboards. As an example, in case of anonymous nodes it is hard to distinguish between different nodes, with the use of whiteboards, agents can easily give every node unique identifier.

Prior Knowledge of the Distributed System

As will be shown later, in the asynchronous model, agents need to know the size of the distributed system. Without this knowledge, the algorithm would either never stop or could work incorrectly. Apart from this information, agents know the degree of the black hole as well. Number of agents used by an algorithm is in some cases dependent on Δ . With the awareness of Δ , the agents can make sure that only as many agents operate as is necessary. These two information are sufficient for the agents to locate the black hole[7], nevertheless adding more information about the distributed system can improve the efficiency of BHS solution as we will see in the chapter 3.

Sense of Direction

Sense of direction gives agents some kind of orientation based on port labeling. A simple example would be a mesh graph with ports labeled **left**, **right**, **up**, **down**. The exact definition is quite straightforward: we say that a distributed system has the sense of direction if for any node v and two paths starting in v , it is always possible to tell whether they end up in the same node based on the sequences of port labels of the two paths. As we will show in the section 3.2, the sense of direction can improve the size of the solution but does not help much with the cost of the optimal solution.

Topological Knowledge

We say that agents have *topological knowledge* if they possess a map of the distributed system a-priori. Agents with a given map of the graph do not need to explore the whole graph, they just need to locate the black hole. The problem in the asynchronous model is that sometimes the only way to decide whether somewhere is the black hole is to make sure that it is not anywhere else. It can be observed that with the map of the graph agents obviously have a very good sense of direction as well. We provide results showing that with the map of the graph the cost of the solution can be improved in the next chapter.

Properties of the Distributed System

Algorithms solving general BHS have to work correctly on any arbitrary graph provided and the size and cost of the solution are based on the worst case scenario. Probably simplest modification of BHS is narrowing it by restricting the graph to be of some type. In such cases agents might take the advantages of patterns or special topological properties of the graphs.

The focus is usually on rings, as in [6, 11], or on the other hand, on graphs with small diameter as hypercubes and tori [4]. The rings proved to be an useful instrument for showing the lower bounds. We mention the ring networks in the section 3.3.

Chapter 3

Previous Results

In this chapter, we would like to present a short survey of previous studies, methods used and functional solutions for the Black Hole Search Problem.

3.1 General Statements and Ideas

First we summarize limitations of the Black Hole Search that are mostly due to agents' asynchronous behavior. Most of the statements are straightforward, for exact proofs we advice the reader to [7].

A trivial technique for safer graph exploration is presented afterwards.

BHS Solution Limitations

It is quite clear that with the use of only one agent, the Black Hole Search is impossible. The agent has to move from the home base and thus the first move it makes can lead into the black hole.

Lemma 1 ([6]). *At least two agents are needed to locate the black hole*

As was noted before, the adversary has great powers over the agents and the distributed system. Especially when the agents are slowed down when entering a link, it might be indistinguishable from the scenario when they enter the black hole. The two following statements are based on this fact.

Lemma 2 ([7]). *If there are two or more unexplored nodes, the Black Hole Search is not solved.*

The Lemma 2 tells us that the agents always have to explore the whole distributed system to locate the black hole correctly.

Next statement addresses the problem of uncertainty of the black hole existence. If the problem is defined in such way that the agents know there may be one black hole, then all they can do is locate its possible position. At the end, the surviving agents are not able to distinguish whether the last unexplored node is the black hole or not. The last unexplored node may be the black hole as well as a safe node with all links slowed down by the adversary.

Lemma 3 ([7]). *It is not possible to verify whether there is a black hole in the distributed system.*

Without loss of generality we assume throughout the whole work that there is exactly one black hole in the distributed system.

Lemma 4 ([7]). *If G has a cut vertex different from the home base, then it is impossible to determine the location of the black hole.*

As was stated in the chapter 2, we will always use biconnected graphs. The Lemma 4 provides reason behind this restriction. For graphs that are not biconnected, the black hole as a cut vertex would make the search impossible.

Last statement tells us what agents need to know about the distributed system, so the Black Hole Search will be possible. Even if the agents do not have a map of the distributed system or any other topological knowledge, they still need to know the size of the distributed system.

Lemma 5 ([7]). *It is impossible to determine the location of the black hole if the size of G is not known.*

Cautious Walk

The technique that allows algorithms to use as few agents as possible and was used in the majority of previous studies is called *cautious walk*. There are two possible variations based on the knowledge that agents have. As first, we explain variant in which the agents have no topological information about the distributed system.

The main idea is that through every unknown edge which might possibly lead into the black hole, only one agent is allowed to pass. Ports are classified into three categories: unknown, dangerous and safe. When an agent is

entering an unknown port, it marks the port as dangerous. No agent is allowed to enter a dangerous port. If the dangerous port leads into the black hole, then the agent that entered it died in the black hole, but no other agent will make the same "mistake". On the other hand, if the port led into a safe node, the agent marks the other port of that link as safe, returns to the previous node and marks the port it entered as safe and then proceeds with the exploration. This technique ensures that maximally one agent passes each link leading to the black hole, thus maximally Δ agents die in total.

When agents have the sense of direction or even the topological knowledge of the distributed system, then the cautious walk can be improved even more. With such knowledge they can always tell whether two ports lead towards the same node. In this case, not the ports will be considered dangerous but the nodes to which they lead. When an agent is entering a port that leads into an unknown node, it needs to know which ports did other agents enter and make sure that it is not entering the same node. As it will be shown later, in such scenario only one agent enters the black hole thus two agents are sufficient for the Black Hole Search.

3.2 The Effects of Prior Knowledge of the Distributed System

In this section, we focus on the dependency of the size and cost based on different levels of agents' prior knowledge of the distributed system. Firstly we present tight bounds on the size and cost of the solution when agents explore arbitrary unknown distributed systems. The size of the solution can be improved if the agents are given the sense of direction. It is not possible to improve the cost of such solution as well. The last observation is of agents with full topological knowledge of the distributed system. In this case, two agents are sufficient and the cost of the solution can be improved using various methods.

Exploring an Unknown Graph

Let us start with the lower bound on the number of agents necessary to fulfill the Black Hole Search. For graphs with $\Delta = 2$ — rings, two agents

are necessary and sufficient for BHS [6]. As can be seen from the following theorem, for the maximal degree of a node $\Delta \leq n-4$ adversary can make sure that Δ agents enter the black hole. When almost every node is connected to the black hole — $\Delta > n-4$ the theorem shows that the adversary can make $\Delta-1$ agents die in the black hole.

Theorem 1 (Lower Bound on the Size of the Solution without any Topological Information [7]). *There is an n node graph G with maximum degree $3 \leq \Delta \leq n-4$ such that, without topological information, any algorithm for locating the black hole in arbitrary networks needs at least $\Delta+1$ agents in G . In addition, if $n-3 \leq \Delta < n$ then any such algorithm needs at least Δ agents.*

For both cases $\Delta \leq n-4$ and $\Delta > n-4$ it is possible to construct algorithms which use the optimal number of agents according to the previous theorem [7]. Thus for both cases, tight bound on the size of the optimal solution is provided.

As for the cost of the algorithm which uses the optimal number of agents, tight bound is provided as well. Theorem stating the lower bound follows.

Theorem 2 (Lower Bound on the Cost of the Size-Optimal Solution [7]). *Let A be an algorithm that, with topological ignorance, correctly locates the black hole in all graphs in $\mathcal{G}(\Delta, n)$ using $\Delta+1$ agents, where $n < 7$ and $3 \leq \Delta \leq n-4$. Then there exists a $G \in \mathcal{G}(\Delta, n)$ such that A incurs $\Omega(n^2)$ steps on G in the worst case.*

We provide a sketch of a proof for the Theorem 2. This proof serves as a very good example on how is the concept of adversary used in such proofs.

Proof. The proof uses a graph which is composed of special cells that are connected into a ring. Nodes in the cell are ordered into Δ consecutive pairs which might resemble a ladder with an extra node on each end, an example of such cell for $\Delta = 5$ can be seen on the Fig. 3.1(a). An example of such graph can be seen on the Fig. 3.1(b).

The adversary will play a following game. He will always let the agents explore one cell at a time and slow down an agent entering the cell in the opposite direction of the ring. By exploring a cell, we mean that he ensures that agents enter all Δ links connecting the pairs first, and then he allows

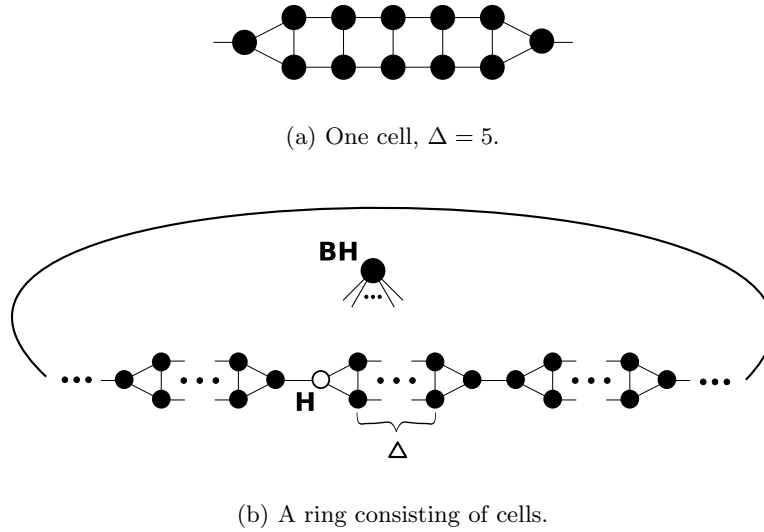


Figure 3.1: Worst case for the cost of a correct algorithm.

them to explore the whole cell. An example of the situation can be seen on the Fig. 3.2.

It can be proven that the agents have to enter all Δ connecting links. If they would not, then it means that m agents, $m \leq \Delta$, had entered some link and the remaining agents had to wait somewhere for the m agents to return. The problem is that adversary can make sure that all m links that were entered by an agent and also the blocked link in the opposite direction lead into the black hole and thus no agent returns.

Once all Δ agents entered some link in the cell, the adversary reveals the whole cell and blocks only the link leading to the next cell. The adversary now also lets the agent on the opposite side of the ring to enter the next cell. By this, he forces all $\Delta - 1$ agents to traverse the already explored part to get to the cell on the opposite side.

To cross one cell, agent needs $\Delta + 3$ steps. When the adversary unblocks i -th cell, agents make $(\Delta - 1)(\Delta + 3)(i - 1)$ steps in total. There are $2\Delta + 2$ nodes in one cell which means that there are $\lfloor n/(2\Delta + 2) \rfloor$ cells in the graph.




Figure 3.2: Agents exploring a cell. Arrow marks port entered by an agent.

Thus it can be proven that agents have to make $O(n^2)$ steps. \square

An algorithm which uses only $\Delta + 1$ agents that make $O(n^2)$ steps in total is quite simple. Main idea is that from each node only one link might lead into the black hole. Thus two agents can together explore one node — determine destination of all its links.

A node is considered unexplored when it is safe but the destination of the links is unknown. Agents keep a list of unexplored nodes in the home base. Using this list, maximally two agents assign themselves for exploration of one node. When the node is explored — every port is either safe or dangerous, agent returns to the home base, and assigns itself for the next unexplored node.

Every time an agent finds a new node, it returns to the home base, adds the new node to the list of unexplored nodes and returns to the node which it was exploring. There are n nodes to be explored in the distributed system. An agent returning to the home base and back to the last node makes $O(n)$ steps, thus the algorithm uses $O(n^2)$ steps in total.

Theorem 3 (The Cost-Optimal Solution Using Optimal Number of Agents [7]). *Above described algorithm correctly locates the black hole in $O(n^2)$ steps using $\Delta + 1$ agents.*

Using the Sense of Direction

With the sense of direction, it is possible to determine whether two ports lead into the same node. As was mentioned before, with the use of cautious walk not more than one agent enters an unknown node. There is an efficient algorithm which uses only two agents that have to make $O(n^2)$ steps in total.

The two agents explore the nodes together. When one of the agents, A , does not return from a dangerous node, the other agent, B , notes where does the port lead and proceeds with the exploration while avoiding the dangerous node. If the agent A returns from the dangerous node, it follows the exploration of B until it reaches the node which is being explored. Agent A notifies B that it has survived and joins B in the exploration process.

Theorem 4 (Upper Bound on the Cost of the BHS with the Sense of Direction Using Two Agents [7]). *In an arbitrary network with sense of direction, the black hole can be located by two agents with cost $O(n^2)$.*

The cost optimality of this algorithm can be proved as well. A very simple sense of direction is used, ports are labeled based on their destination node. This sense of direction gives only information about the node on the other side of the link but no other knowledge is provided.




Figure 3.3: A ring consisting of triangles.

A ring graph with nodes replaced by triangles will be used, see Fig. 3.3. If the adversary makes the links between the triangles very slow, then agents have to explore in opposite direction. Once an agent gets stuck on such link, the other agent has to proceed with the exploration and only unknown link is on the opposite side of the explored part. When both agents enter these links, they eventually get into the node with two unknown ports. The only way of deciding whether any of the new ports lead into the same node is by comparing the port labels, which means that the agents have to meet or visit the other agent's node.

It is quite clear now that to explore next triangles on each side, at least one agent has to cross the explored part, which implies that the agents make $O(n^2)$ steps in total.

Theorem 5 (Lower Bound on the Cost of the BHS with the Sense of Direction Using Two Agents [7]). *The worst case cost for locating a black hole in arbitrary networks, independently of the type of sense of direction available, is $\Omega(n^2)$.*

Knowledge of the Whole Graph

Agents that have topological knowledge of the whole graph automatically gain the sense of direction as well. With the use of previously mentioned algorithm, two agents can perform BHS in $O(n^2)$ steps. Taking the advantage of full topological knowledge can easily improve the cost of the solution to $O(n \log n)$ steps [7]. This is the optimal solution for arbitrary graphs in

terms of n in deed. It is due to the fact that for a ring networks, the optimal cost is $\Theta(n \log n)$ [6].

Theorem 6 (The Optimal Cost of the Solution Using Two Agents with Topological Knowledge [7]). *The black hole can be located by two agents with full topological knowledge in arbitrary networks of vertex connectivity 2 with cost $O(n \log n)$, and this is optimal.*

The idea is that the two agents first divide the whole distributed system into two separate equal-size parts. The black hole can be only in one of the parts. Each agent explores its part and one of the agents eventually finishes. The agent that finishes follows the other agent until it reaches the dangerous link where the other agent got stuck. It divides the remaining unexplored part of the graph into two parts, leaves a message for the other agent and starts exploring one of the smaller parts. This process repeats until there is only one node in the remaining unexplored part. As the graph can be halved $O(\log n)$ times and finding the other agent and entering the newly assigned part might take up to $O(n)$ steps, the cost of the algorithm is $O(n \log n)$.

In [4] similar algorithm is applied on graphs with small diameter such as hypercubes, cube connected cycles, wrapped butterflies, tori and other. For these graphs, the cost of the solution is $O(n)$ steps only.

With the consideration of the graph diameter, even more interesting results can be obtained. For arbitrary graph, two agents with full topological knowledge can perform the Black Hole Search in $O(n + d \log d)$ steps in total [9].

Different approach was used in [10] where authors use open vertex cover by cycles and provide an efficient algorithm using different principle.

3.3 Other Results

This section provides a brief overview of more results obtained by different studies of the Black Hole Search Problem.

Tokens

Tokens as a type of communication are much more restricted than whiteboards. If agents are allowed to put one token into the middle of each node

and one in front of each port, then the model is very similar to a whiteboard model with $\log(d)$ bits per whiteboard, where d is the degree of the node. On one hand, tokens provide more places where it is possible to put the "bit", on the other hand, the usual setting is that every agent carries only one token. Thus there are only as many tokens as agents in the graph.

For arbitrary unknown graphs the size of the optimal solution is similar to the whiteboard model, $\Delta + 1$ agents are able to locate the black hole. Nevertheless, the cost of the provided solution is $O(\Delta^2 m^2 n^7)$ [5].

When the agents are given full topological knowledge of the distributed system, the token model is proved to be as efficient as the whiteboards model and the black hole can be located in $\Theta(n \log n)$ steps [12].

Scattered Agents

When assuming the absence of the home base, agents are initially scattered around the distributed system. The agents are thus unable to cooperate at beginning of their exploration. They have to avoid entering the black hole and gather with other agents as well.

First studies were accomplished on ring networks. Conditions under which are the agents able to gather and tight bounds on the maximal number of agents that are able to gather are provided in [8].

Study of scattered agents exploring a faulty distributed system can be found in [1] as well. Instead of assuming presence of the black hole, the distributed system contains faulty links. Clearly, it is only more general problem to our setting, in which all the faulty links lead into the black hole.




Figure 3.4: Oriented graph — 2^Δ agents necessary.

Oriented Graphs

All previously mentioned results use the cautious walk technique to reduce the number of agents that enter the black hole. When considering oriented graphs, this technique is not available. If an agent marks a port as dangerous, it is not able to return straight away and mark it safe. The agent might even enter a black hole and thus more agents have to enter the dangerous port. For the graph shown on the Fig. 3.4, 2^{Δ} agents are needed in the worst case [3].

Black Hole Search Problem on oriented graphs and the size of the optimal solution served as one of the motivations for the model of distributed systems without the knowledge of incoming link and its further studies.

Chapter 4

The Knowledge of Incoming Link

All previously stated results and studies assumed that when an agent enters a new node, the agent automatically knows what port did it enter through. We refer to this as to *the knowledge of incoming link*.

Main focus of this work is on the effects of the knowledge of incoming link, more precisely on how lack of this knowledge affects BHS problem solution.

4.1 Proposed Model

We study BHS problem in the following scenario:

Definition 2 (Black Hole Search Problem without the Knowledge of Incoming Link). *Asynchronous autonomous agents operate on a distributed system consisting of arbitrary graph and its port labeling. One node of the graph is very harmful — a black hole. Agents work according to a predefined algorithm. At the beginning, they are all stationed in the home base. Their task is to perform the Black Hole Search.*

Agents' only prior knowledge of the distributed system is the number of nodes n and the degree of the black hole Δ . They communicate by means of whiteboards, which are located in every node. Moreover, the agents have no knowledge of incoming link: when they enter a new node they do not have the information which port leads to the previous node.

All agents and nodes have their identifiers. Asynchronous agents with the use of whiteboards can easily assign identifiers to themselves and to the nodes as well, thus this assumption can be done without loss of generality.

It can be observed that our modification makes agents less powerful and that the size of the optimal solution will be at least $\Delta + 1$ and the cost will be $O(n^2)$. We start by examining the basic properties of this model and the effect of our modification on the size of the solution.

Number of Agents Needed

First observation of the newly proposed model is that the cautious walk is no longer possible in its original form. If the agent is not able to decide which port it should enter to get back, then it is not necessarily able to get back and withdraw the warning. The cautious walk is a mechanism that ensures that maximally one agent travels through an edge leading into the black hole and thus only Δ agents die.

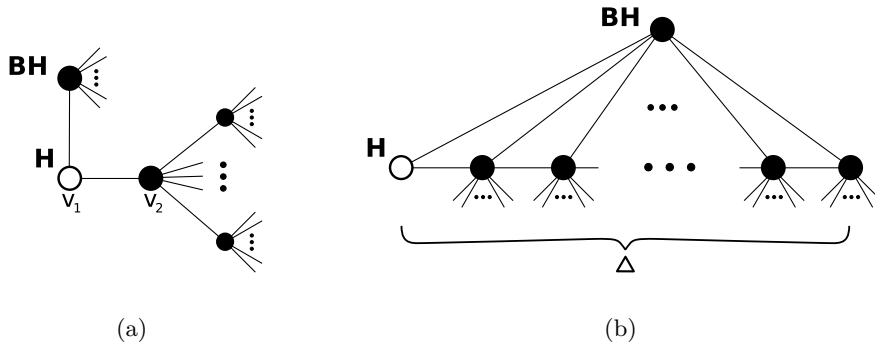


Figure 4.1: Two example graphs.

In the Fig. 4.1(a) the node v_1 is the home base. The question is, how many agents have to leave the node v_2 so at least one gets back into the home base. We know that the node v_2 can have Δ outgoing ports and thus if we send one agent through each port then at least one will get back into the home base. This leads us to a rough estimation that if we send Δ agents into the node v_2 then at least one will return.

The agents do not have a map and cannot distinguish a port leading to v_2 from a port leading to the black hole and thus we need to send Δ agents

through both ports, to make sure that Δ agents get to v_2 and at least one returns.

It might seem that the model is very similar to the oriented graph scenario, where edges have orientation and it is never possible to get back using the same edge. Luckily it is not quite so. If at least one agent returns from v_2 , we might be able to distinguish which port leads into still unknown node and which one into the safe node v_2 . It is simply accomplished by writing on the whiteboard which port did agents use when they left.

Let us consider example on the Fig. 4.1(b). We see that by using the previous reasoning, in order to decide which port is leading into the next node on the right, we need to send Δ agents through each. Thus we send approximately Δ^2 agents into the black hole.

We proposed very rough estimates for the number of agents that can die in the studied model. We can see that designing a graph where more than Δ agents die is quite easy. In the second example, it is noticeable that Δ^2 agents is more than necessary. With this in mind we present the lower bound on the number of agents needed for finding the black hole.

4.2 Lower Bound on the Size of the Optimal Solution




Figure 4.2: Worst case scenario: G_{lb}

We show that any algorithm solving BHS for an arbitrary unknown graph needs at least $\frac{\Delta^2 + \Delta}{2} + 1$ agents to solve BHS on the graph G_{lb} shown on the Fig. 4.2.

Let us consider the Fig. 4.3(b). Agents start the exploration in the home base v_1 . They see two indistinguishable edges. The algorithm has to send a number of agents through each edge and then other agents wait

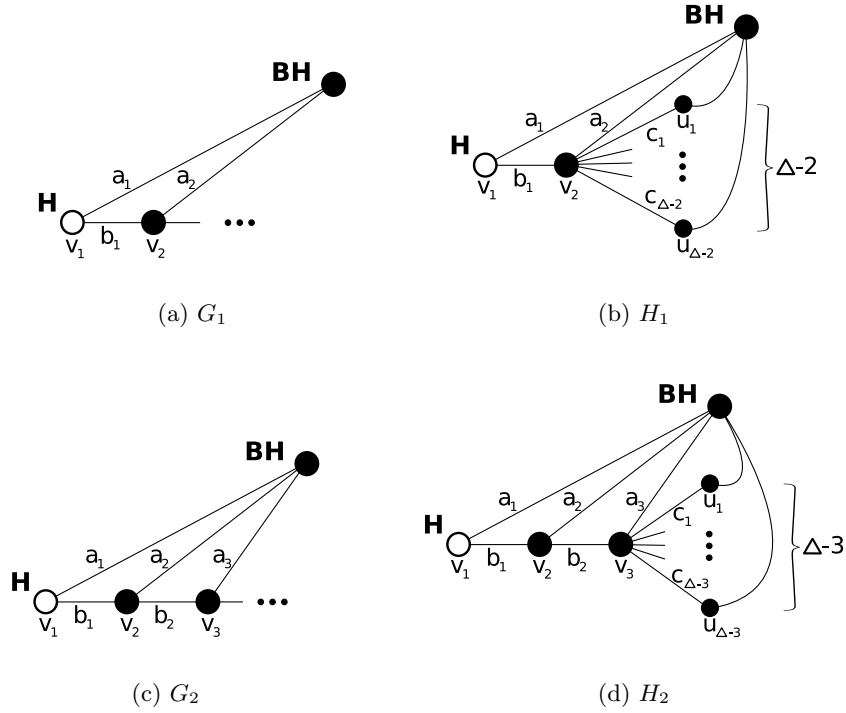


Figure 4.3: Indistinguishable graphs.

until some agents come back. If the algorithm sends a different number of agents through each edge, the adversary can make sure that more agents pass through the edge a_1 and thus end up in the black hole. This means that an efficient algorithm has to send the same number of agents through the both edges a_1 and b_1 .

It can be seen that there have to be at least Δ agents walking through the edge b_1 . If the algorithm sends only $k < \Delta$ agents, the adversary can make sure that none of the agents returns to the home base and thus the algorithm fails. There are Δ edges leading out of v_2 . Without loss of generality it can be assumed that no two agents use the same edge, one agent uses the edge a_2 and the other $k - 1$ agents use the edges $c_1 \dots c_{k-1}$, they get into the nodes $u_1 \dots u_{k-1}$ and from these they use edges leading into the black hole. All k agents die in the black hole and no agent returns.

For the graph H_1 , there have to be at least Δ agents using the edge b_1 and thus also Δ agents using the edge a_1 . When considering the Fig. 4.3(a) and the graph G_1 , it can be observed that for agents waiting in the

home base, the graphs G_1 and H_1 are indistinguishable. This means that the algorithm has to act in the same way in both cases, thus also in the graph G_1 there have to be at least Δ agents passing through the edges a_1 and b_1 .

A very similar situation can be seen in the Fig. 4.3(d). As stated before, the adversary can make sure that there are at least Δ agents traveling through the edge a_1 . Now the edge b_1 is known to be safe and agents can freely move into the node v_2 . Let us assume that the algorithm now knows that a_1 leads into the black hole so we do not have to take it into account. Once again there are two indistinguishable edges a_2, b_2 and there is no point for sending more agents into one of them.

Similarly as in the graph H_1 , to make sure that at least one agent gets back into the node v_2 we need at least $\Delta - 1$ agents in the node v_3 . Otherwise all agents that went into v_3 can end up in the black hole. With the knowledge of v_1 and v_2 , the graph G_2 from the Fig. 4.3(c) is indistinguishable from the graph H_2 . This implies that the algorithm, which solves BHS for an arbitrary unknown graph, will in the worst case send Δ agents through the edge a_1 and $\Delta - 1$ agents through the edge a_2 in both, G_2 and H_2 .

Extending the previous argument on graphs $G_3 \dots G_\Delta$, where G_Δ is similar to G_{lb} we can prove that for $i \in \{1, \dots, \Delta\}$ there are $\Delta - i + 1$ agents traveling through an edge a_i and thus dying in the black hole. Following theorem is thus implied.

Theorem 7 (Lower Bound on the Size of the Optimal Solution). *Any algorithm solving the Black Hole Search Problem for an arbitrary unknown graph without the knowledge of incoming link needs at least $\frac{\Delta^2 + \Delta}{2} + 1$ agents.*

Chapter 5

Size-Optimal Algorithm

Before we describe the algorithm that uses the optimal number of agents, let us look again at the example 4.1(a). There are some ideas left undiscussed and further investigation leads us to the optimal solution.

In the previous chapter, we discussed that if we send Δ agents into the node v_2 then one will come back, if we design the algorithm in such way that no more than one agent will enter two ports of v_2 . The question is whether we always need Δ agents. The fact that some agents do not leave through the returning port does not mean that they are dead and useless. Some of them might come back. If some of the ports eventually lead into the black hole, we could expect that all agents leaving through these ports would die, but the rest could come back.

The second idea that should be considered are the conditions under which we actually need to send more than one agent through an unknown port in the first place. The thing is that we need some agents to come back, thus they cannot fall into the black hole. If the home base would have $\Delta + k$ ports then it might be enough to send one agent through each. Maybe Δ of them would fall into the black hole but the rest of them would be perfectly safe because the black hole has only Δ links leading into it.

We clarify these conjectures while describing the algorithm and proving its correctness.

5.1 Algorithm

We propose the following algorithm as a solution for the Black Hole Search Problem on a distributed system without the knowledge of incoming link.

General Ideas of the Algorithm

All ports are classified into four categories, similar to those used in the *cautious walk* with one extra. Unmarked ports are called *unknown*. When one or more agents have entered a port, it is marked as *dangerous*. If a port is dangerous and it is known that at least one of the agents that entered such port has been alive afterwards, then the port can be marked as *safe*. When a link has both ports safe, then the link is safe as well and both ports of the link are marked as *explored*. Furthermore, for both explored ports the nodes to which they lead are known and are written on the whiteboard.

One or more nodes of the graph that are connected by safe links are called *safe component*. Safe component which contains the home base is called *main safe component* (MSC).

Every agent searching the graph has a counter of steps it has made. Every time an agent is going to use a port, it leaves a message on the whiteboard stating the port's label, agent's *ID* and the current state of its counter, it increases its counter as well. Each agent keeps track of how many steps have other agents done. When an agent *A* enters a node, it looks on the whiteboard and for every agent *B* that left a message, it updates its information about *B*'s number of steps. An agent always remembers the complete list of agents that visited the current node. The list of counters, known map of the distributed system and a list of agents that entered a dangerous port are kept in the home base as well.

The lists are used as follows. An agent that enters a node *u* first compares its list of agents in the previous node with the list of agents that left the node *u*. By comparing the state of their counters, the agent can immediately find out, whether any of the agents went through the same link and it knows by which port as well. If there is such agent, it marks the port explored, goes back to mark the other port as well and returns to the node *u*.

Second thing an agent has to do when it enters a new node *u* is to compare its list of all agents and maxima of their counters to the list of agents that left *u*. If there is a port marked as dangerous and it knows that

one of the agents that entered it stayed alive afterwards then it marks this port safe and enters it.

After these two checks, the agent proceeds with its own exploration.

Now we can state rules for the agents, which they obey at all circumstances. We will prove the correctness of this algorithm in the next section.

1. When an agent is in a node that is not part of MSC, it will never enter a dangerous port.
2. When an agent is in a node, it will always make a move, unless the problem is solved or all ports are dangerous. If all ports are dangerous the agent will wait until one of the ports is made safe and then proceeds with the exploration.
3. When agent is in MSC it can enter a dangerous port under following conditions:
 - There are no unknown or safe ports in MSC.
 - The size of MSC is less than $n - 1$.
 - There are k dangerous ports leading out of MSC and $k \leq \Delta$.
 - There are less than $\Delta - k + 2$ agents, that entered the port it wants to enter.

This condition ensures that if there are $k \leq \Delta$ dangerous ports at a certain time, then there will be no more than $\Delta - k + 2$ agents entering any of the dangerous ports. Of course, if in some previous state there were k' dangerous ports where $k' < k$ then there might be some ports that were used by $\Delta - k' + 2$ agents which is more than $\Delta - k + 2$. In this case, however, no more agents will enter these ports.

4. An agent has to enter a port when it wants to mark it as safe.
5. An agent moves through an unknown port from node u to node v , upon arrival to the node v it checks its list of agents that visited the node u and compares it with the whiteboard. If it finds out that some agent previously traveled from v through port p and came to u then it has to mark the port p and the whole link as safe¹.

¹By marking a link safe we always mean marking its ports as explored as well.

From these rules it follows that when an agent is not able to reach the home base by a safe path, it does not enter any dangerous ports. The agent does not stop exploring. If all ports are dangerous it waits until some port is marked safe. When the agent discovers that some port/link can be marked as safe, then it is its duty to do so.

More Detailed Description

Agents collect information about MSC in the home base. There is a map of MSC with all outgoing ports being updated by the agents. For every dangerous port, there is a list of agents that entered this port. List of agents with last known values of their counters is kept in the home base as well. Except from this, all nodes of MSC are labeled and will have marked port that leads to the home base.

The algorithm can be divided into three main parts based on where the agent is and what has it done. The first part is for agents without any mission stationed in MSC. The second part is used when agent is outside of MSC. The last part is designed for an agent that has just returned to MSC.

In the procedure `EXPLORE`, procedure `SEARCH` is used to search for a next node with an unknown port. All nodes connected by safe and explored ports can be considered as a subgraph in which agents move as in a directed graph. Dangerous ports are being ignored. Ports marked as safe are directed edges, because on the other side of the link it is not clear which port leads back through that link. Explored ports, with their safe links can be considered as two symmetric directed edges.

For such graph a basic algorithm for exploration of directed graphs can be used, which forms the main part of the procedure `SEARCH`. Apart from that, agent always looks for ports that could be made explored or safe, similarly to the initial checks of the procedure `EXPLORE`.

Procedure 1 LEAVE MSC

```

1: return to home base
   update lists and map
2: if size(MSC) =  $n - 1$  then
3:   if  $\#_{DANGEROUS} \leq \Delta$  then
4:     FINISH
5:   else
6:     add  $MY(ID)$  to waiting list
       sleep until wake up
       goto 1
7:   end if
8: end if
9: if  $(\exists p)$  AND ( $p$  is dangerous) AND
    $(\exists A)$  AND ( $A$  entered  $p$ ) AND  $(\text{counterAtPort}(A,p) < \text{counterMax}(A))$ 
   then
10:  mark  $p$  as safe
     EXPLORE( $p$ )
11: end if
12: if  $(\exists p)$  AND ( $p$  is unknown) then
13:  mark  $p$  as dangerous
     EXPLORE( $p$ )
14: else if  $(\exists p)$  AND ( $p$  is safe) then
15:  EXPLORE( $p$ )
16: else
17:  if  $(\exists p)$  AND ( $p$  is dangerous) AND  $(\#_{DANGEROUS} < \Delta)$  AND
      $(\#_{ENTERED\_AGENTS}(p) < \Delta - \#_{DANGEROUS} + 2)$  then
18:    EXPLORE( $p$ )
19:  end if
20: end if
21: add  $MY(ID)$  to waiting list
     sleep until wake up
22: goto 1

```

Procedure 2 EXPLORE

```

1:  $u \leftarrow$  current node
2:  $u' \leftarrow$  previous node
3: if  $(\exists A)$  AND  $(A \text{ in } u)$  AND  $(A \text{ in } u')$  AND
    $(\text{counterAtPort}(A, \text{lastPort}(A, u)) + 1 = \text{counterAtPort}(A, \text{lastPort}(A, u')))$ 
   AND
    $(\text{lastPort}(A, u) \text{ is dangerous OR safe})$  then
4:    $q \leftarrow \text{lastPort}(A, u)$ 
     mark  $q$  as explored
     enter  $q$ 
     mark  $\text{lastPort}(MY(ID), u')$  as explored
     enter  $\text{lastPort}(MY(ID), u')$ 
5: end if
6: if  $(\exists A)$  AND  $(A \text{ in } u)$  AND  $(\text{lastPort}(A, u) \text{ is dangerous})$ 
    $(\text{counterAtPort}(A, \text{lastPort}(A, u)) < \text{maxCounter}(A))$  then
7:    $p \leftarrow \text{lastPort}(A, u)$ 
     mark  $p$  as safe
     EXPLORE( $p$ )
8: end if
9: if  $(\exists p)$  AND  $(p \text{ is unknown})$  then
10:  mark  $p$  as dangerous
     EXPLORE( $p$ )
11: else if  $(\exists p)$  AND  $((p \text{ is safe}) \text{ OR } (p \text{ is explored}))$  then
12:  SEARCH()
13: else
14:  add  $MY(ID)$  to waiting list
     sleep until wake up
15: end if

```

Procedure 3 RETURN

```

1:  $u \leftarrow$  current node
2:  $u' \leftarrow$  previous node
3: if  $(\exists A)$  AND  $(A \text{ in } u)$  AND  $(A \text{ in } u')$  AND
   (counterAtPort( $A, \text{lastPort}(A, u)$ )+1 = counterAtPort( $A, \text{lastPort}(A, u')$ ))
   AND
   (lastPort( $A, u$ ) is dangerous OR safe) then
4:    $q \leftarrow$  lastPort( $A, u$ )
   mark  $q$  as explored
   enter  $q$ 
   mark lastPort( $MY(ID), u'$ ) as explored
   enter lastPort( $MY(ID), u'$ )
5: end if
6: find Home Base
   update lists and map
   write down Wake up
7: if size(MSC) =  $n - 1$  then
8:   if  $\#_{DANGEROUS} < \Delta$  then
9:     FINISH
10:  else
11:    add  $MY(ID)$  to waiting list
    sleep until wake up
    goto 7
12:  end if
13: end if
14: determine outgoing port  $p$  previously used
15: if  $p$  is dangerous then
16:   mark  $p$  as safe
17: end if
18: locate  $p$ 
   write down destination of  $p$ 
   EXPLORE( $p$ )

```

The procedure RETURN is automatically initiated once an agent returns to MSC.

5.2 Correctness

First we state some supporting lemmas and then prove that MSC will keep expanding until the whole graph except the black hole is explored.

Lemma 6. *When an agent A is not in MSC and A is in a node v_D with all ports marked as dangerous, one of the ports will be eventually marked as explored.*

Proof. The node v_D is not the home base, thus there is an agent A_0 that entered this node as first using the link l_0 . That means that it read an empty whiteboard, wrote down its ID , port by which it left and the state of its counter on the whiteboard. All the agents that enter after the agent A_0 will read its message on the whiteboard.

Now all ports are marked as dangerous thus there exists an agent A_X that left through the link l_0 . When the agent A_X arrives to the other end of the link l_0 , it finds out that the agent A_0 was there just before the node v_D and thus it can identify a port of the link l_0 . He marks the port as explored and goes back to v_D and marks the other port as explored as well, thus one of the dangerous ports has been marked explored and the agent A can leave. \square

The use of the first agent A_0 in the previous proof has its reason. By picking the first agent it is clear that the agent A_X read the both messages from A_0 . Important fact is that before A_X performs any other action and thus risks falling into the black hole it has to make l_0 safe. This lemma can be expanded for more than one node.

Corollary 1. *When an agent is not in MSC and it is in a safe component with all outgoing ports marked as dangerous, one of the ports will be eventually marked as explored.*

Proof. Once again, there was an agent A_0 entering one of the nodes of the safe component from outside of the safe component as first by the link l_0 . Thus there is an agent that left the safe component through the link l_0 and will mark it safe. \square

Lemma 7. *A link that has both ports safe will be marked as safe, and the ports as explored. Furthermore, it will be marked safe by one of the agents that marked its ports safe.*

Proof. Link l has ports p and p' , one of the ports had to be marked safe first, let us assume, that it was the port p . If the port p is safe, it means that an agent A_1 traveled already through it. Then second agent A_2 with the knowledge that A_1 survived, marked it safe and traveled through it as well.

Agent B_2 that is going to mark the port p' as safe can do it only after A_2 does. This means that A_1 has already emerged from the link l and thus B_2 reads its messages on the whiteboard. When B_2 emerges on the other side of the link l it finds previous message of A_1 and knows that it can mark the whole link safe and its ports as explored. \square

In the following lemma we focus on the agents exploring outside of MSC. These agents are not allowed to enter any dangerous ports. We show that there are always enough agents outside, so that at least one does not die in the black hole. Right now we will expect that there are enough agents in the main safe component and that it is always possible to send out more agents if needed. The actual number of agents that are really needed is shown in the section 5.3.

Lemma 8. *For MSC of size smaller than $n - 1$ the following statement holds. When there are no more agents that could leave MSC according to the third rule, then at least one of the agents outside of MSC does not enter the black hole.*

Proof. The fact that no more agents could leave MSC can only mean that all k outgoing ports of MSC are dangerous and through each one went at least $\Delta - k + 2$ agents if $k \leq \Delta$, and at least one agent if $k > \Delta$, thus no more agents are allowed.

Let l be the number of ports leading out of MSC and into the black hole, we know that $l < k$ unless the size of MSC is $n - 1$.

When $k > \Delta$, we do not need all k ports being entered by at least one agent. We expect that we have enough agents, so that there are at least $m \geq \Delta + 1$ outgoing ports that were used by at least one agent. This means that at least $m - l$ agents survived after leaving MSC. There are $\Delta - l$ ports leading into the black hole from outside of MSC and thus maximally $\Delta - l$ agents can enter these ports and die. Thus we get that at least $(m - l) - (\Delta - l)$ agents in total will survive outside MSC. As we mentioned, $m \geq \Delta + 1$ and thus $m - \Delta \geq 1$, which means that at least one agent survives.

Let us consider $k \leq \Delta$. We can expect that all allowed agents entered a port meaning that at least $\Delta - k + 2$ agents entered each port. If l ports were dangerous then $(k - l) \times (\Delta - k + 2)$ is the number of agents that got out of MSC and not into the black hole. From these agents, only $\Delta - l$ will die in the black hole outside MSC. Thus the number of surviving agents can be expressed:

$$\begin{aligned} \#_s &= (k - l) \times (\Delta - k + 2) - (\Delta - l) \\ &= k\Delta - k^2 + 2k - l\Delta + lk - 2l - \Delta + 2 \\ &= -k^2 + k\Delta + 2k - \Delta + l \times (k - \Delta - 1) \end{aligned}$$

We know that Δ and k are fixed and that coefficient of l : $(k - \Delta - 1)$ is negative, so $\#_s$ drops with raising l . Number of ports leading from MSC into the black hole is bounded by the number of all outgoing ports, $l < k$. Thus the smallest number of agents survive when $l = k - 1$:

$$\begin{aligned} \#_s &= (k_s - l) \times (\Delta - k + 2) - (\Delta - l) \\ &= (k - (k - 1)) \times (\Delta - k + 2) - (\Delta - (k - 1)) \\ &= 1 \end{aligned}$$

Which again means that at least one agent outside MSC survives. \square

Lemma 9. *For MSC of size smaller than $n - 1$ the following statement holds. When there are no more agents that could leave MSC according to the third rule, then at least one of the agents that left MSC returns back into MSC.*

Proof. What we are going to show is that one of the surviving agents will eventually enter a port leading back to MSC and thus it returns.

From the Lemma 8 we know that at least one of the agents that left MSC survives. Agents work in such way that they always enter some port, unless all are dangerous in which case they wait. The Lemma 6 further tells us that even though an agent does not enter a dangerous port and waits, it does not get stuck somewhere either. From these three statements it is quite clear that the agent never really stops moving, once it leaves MSC.

There are only n nodes and thus less than n^2 links between them. Every link has two ports and every port can have four states. Thus there is a finite

number of ports outside MSC and they have only finite number of states together.

We will show that agents not only never stop exploration, but when they move through the graph, the states of the ports are always being changed. This means that either an agent itself changes a state of a port or it gets into a situation in which the state of a port will be inevitably changed by some other agent.

Agents can encounter the following situations:

1. When an agent is in a node with an unknown port then it enters it and changes its state to dangerous.
2. If there are only dangerous ports, then the agent waits until one of the ports is made explored, we proved in the Lemma 6 that it will happen.
3. If there are only dangerous, safe and explored ports in a node the agent searches the subgraph of such nodes using only safe and explored ports. When the agent finds an unknown port, it stops the search protocol and enters the port. If there are no reachable unknown ports, then the agent can wait, because one of the dangerous ports will be changed by other agent. To prove this, we could use the same reasoning as in the proof of the Lemma 6.

If surviving agents never stop changing states of the ports outside MSC and there is only finite number of ports, eventually some agent has to change a state of a port leading back to MSC, put in different words, the agent will enter it. \square

We prove that MSC will be expanding until it contains $n - 1$ nodes. At the beginning, only the home base is a part of MSC and it has all ports marked as unknown. Agents start the exploration by entering these ports. We now proceed with a general case, there is the main safe component, with k outgoing ports.

Let us assume that MSC does not expand. Once again we use the fact that there is only finite number of ports leading out of MSC and they can be of only 3 different states. If a port would be marked as explored, it means that MSC is being expanded. The fact that MSC does not expand implies that there must be a state of MSC that does not change any more.

By state of MSC we will understand state of all outgoing ports, with additionally, for each safe outgoing port a list of agents that entered it. When an agent returns to MSC, it leaves it through the safe port, which it has already used before, in such case the agent will be on the list twice. Moreover, if an agent is entering outgoing port second time, it knows where does it lead and it writes the information on the whiteboard in the node of that port.

It is important to note that if at least n agents enter a safe port, then there is certainly an agent that entered the link from the other side, because unknown links are entered first. Thus when the agent comes out of such port into the node of MSC it does not know by which port it arrived but it either reads information left by an agent that entered the port twice and is able to mark the port explored. Or the agent leaves information about itself on the whiteboard so that the next agent which travels through that port is able to mark it explored afterwards.

If the number of agents that entered the safe port would be increasing, MSC would eventually expand, so it has to be stable as well.

There are only two options for the stable state. First one is that there are agents in MSC but they are all waiting in the home base. This further means that all outgoing ports are dangerous and no other agents are allowed to leave through these ports, thus they are waiting. The second option is that there is no agent in MSC otherwise it would exit and thus change the stable state.

Both these states would change by any agent returning to MSC. In the first case, an agent that enters MSC had to leave the MSC before by a port p . Now that the agent is back in MSC it can, and will, mark the port p as safe, thus changing the state of MSC, allowing more agents to continue the exploration.

In the second case, the fact of an agent returning to MSC changes the stable state. Even if the state would not be changed directly, the agent had previously left MSC and traveled through outside of MSC by some path. Part of this path might have become a part of MSC but there must be one port which now leads out. This port is either dangerous and the agent can mark it as safe, or it was marked as safe in the meanwhile. In such case it will do as mentioned above, leave information about the port's destination, add itself into the list and enter it again.

The Lemma 9 proves that when the number of nodes of MSC is less than $n - 1$ then every time there are no more agents able to leave MSC at least one agent returns to MSC. We formulate our statement in the following theorem.

Theorem 8 (Correctness). *Agents with their exploration, will expand the main safe component until it contains $n-1$ nodes, thus exploring whole graph except the edges possibly leading into the black hole. By this they finish the Black Hole Search.*

5.3 Upper Bound on the Size of the Optimal Solution

In the proof of the lower bound value we used the graph G_{lb} from the Fig. 4.2. Note that our algorithm uses optimal number of agents to fulfill BHS on such graph — $\frac{\Delta^2 + \Delta}{2} + 1$. We do not need to prove this statement for the purposes of the upper bound proof.

Even more, let us order all ports leading into the black hole based on the number of agents that used them in descending order, these numbers make a following sequence: $p_1 = \Delta, p_2 = \Delta - 1, \dots, p_{\Delta-1} = 2, p_{\Delta} = 1$. This corresponds with the claim from the lower bound proof that for each edge a_i leading into the black hole, there had to be $\Delta - i + 1$ agents that passed through this edge.

Theorem 9 (Upper Bound on the Size of the Optimal Solution). *Algorithm for the Black Hole Search Problem obeying conditions stated before will succeed with $\frac{\Delta^2 + \Delta}{2} + 1$ agents in the worst case.*

Proof. Consider a sequence of Δ numbers similar to the one mentioned before, for $i \in \{1, \dots, \Delta\}; p_i = \Delta - i + 1$.

Let us consider graph G_{ub} , with the degree of the black hole Δ , where our algorithm needs more than $\frac{\Delta^2 + \Delta}{2} + 1$ agents to solve BHS. In other words, there are at least $\frac{\Delta^2 + \Delta}{2} + 1$ agents that die in the black hole. Now again we take the numbers of agents that used a port leading into the black hole and order them in descending order. These numbers form a non-increasing sequence of length Δ : $q_1, q_2, \dots, q_{\Delta}$. We stated previously that there are

more agents ending up in the black hole in the graph G_{ub} and thus:

$$\sum_{i=1}^{\Delta} q_i > \sum_{i=1}^{\Delta} p_i = \frac{\Delta^2 + \Delta}{2}$$

Both sequences have the same lengths and thus there must exist index j where $q_j > p_j$. We know that $q_j > p_j = \Delta - j + 1$, and thus $q_j \geq \Delta - j + 2$. Now we have to consider what does this mean in terms of the rules for our algorithm.

In the last rule about the number of agents entering a dangerous port at certain conditions, we state that for $k \leq \Delta$ maximally $\Delta - k + 2$ agents enter. When at least $\Delta - j + 2$ agents enter dangerous port it means that:

$$\begin{aligned} \Delta - j + 2 &\leq \Delta - k + 2 \\ j &\geq k \end{aligned}$$

For all ports corresponding to the numbers q_1, \dots, q_j , there is an agent that entered these ports last. For this last agent, some conditions must have been fulfilled. There were no unknown ports in MSC. Number of the nodes in MSC was less than $n - 1$, put in other words, there were still some nodes that were not part of MSC. And at last, there were maximally k dangerous ports at that moment, where $k \leq j$ from above stated equation.

Numbers q_1, \dots, q_j correspond to j different ports leading **into** the black hole. If the graph is biconnected then we are missing at least one dangerous port, that **did not lead** into the black hole. So the last agent would obviously have to break one of the conditions by entering the dangerous port, which leads us to contradiction. \square

Corollary 2 (Size Optimal Solution). *Proposed algorithm provides size optimal solution for the Black Hole Search Problem without the knowledge of incoming link.*

Conclusion

In this work, we have studied the Black Hole Search Problem without the knowledge of incoming link. We have shown that this modification has effects on the size of the solution.

We provided lower bound on the number of agents that are necessary to locate the black hole. Any correct algorithm solving the Black Hole Search Problem without the knowledge of incoming link needs at least $\frac{\Delta^2 + \Delta}{2} + 1$ agents. The algorithm is presented with the proof of correctness, it uses the optimal number of agents in the worst case. The cost of the algorithm and bounds on the optimal cost of the solution are left for further investigation.

Bibliography

- [1] Jérémie Chalopin, Shantanu Das, and Nicola Santoro. *Distributed Computing*, volume 4731 of *LNCS*, chapter Rendezvous of Mobile Agents in Unknown Graphs with Faulty Links, pages 108–122. Springer, 2007.
- [2] Colin Cooper, Ralf Klasing, and Tomasz Radzik. Searching for black-hole faults in a network using multiple agents. In *Proceedings of the 10th International Conference on Principles of Distributed Systems (OPODIS '06)*, volume 4305 of *LNCS*, pages 320–332. Springer, 2006.
- [3] J. Czyzowicz, S. Dobrev, R. Královič, S. Miklík, and D. Pardubská. Black hole search in directed graphs. In *SIROCCO 2009*, LNCS. Springer, 2009. to appear.
- [4] Stefan Dobrev, Paola Flocchini, Rastislav Královič, Giuseppe Prencipe, Peter Ružička, and Nicola Santoro. Black hole search in common interconnection networks. *Networks*, 47(2):61–71, 2006.
- [5] Stefan Dobrev, Paola Flocchini, Rastislav Královič, and Nicola Santoro. Exploring an unknown graph to locate a black hole using tokens. In *IFIP TCS*, volume 209 of *IFIP*, pages 131–150. Springer, 2006.
- [6] Stefan Dobrev, Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. Mobile search for a black hole in an anonymous ring. In *Proceedings of the 15th International Conference on Distributed Computing (DISC '01)*, volume 2180 of *LNCS*, pages 166–179. Springer, 2001.
- [7] Stefan Dobrev, Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. Searching for a black hole in arbitrary networks: optimal mobile agent protocols. In *Proceedings of the twenty-first annual symposium on Principles of distributed computing (PODC '02)*, pages 153–162. ACM, 2002.
- [8] Stefan Dobrev, Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. Multiple agents rendezvous in a ring in spite of a black hole. In *Proceedings of the 7th International Conference on Principles of Distributed Systems (OPODIS '03)*, volume 3144 of *LNCS*, pages 34–46. Springer, 2003.

- [9] Stefan Dobrev, Paola Flocchini, and Nicola Santoro. Improved bounds for optimal black hole search with a network map. In *Structural Information and Communication Complexity (SIROCCO '04)*, volume 3104 of *LNCS*, pages 111–122. Springer, 2004.
- [10] Stefan Dobrev, Paola Flocchini, and Nicola Santoro. Cycling through a dangerous network: A simple efficient strategy for black hole search. In *26th IEEE International Conference on Distributed Computing Systems (ICDCS '06)*, page 57. IEEE Computer Society, 2006.
- [11] Stefan Dobrev, Rastislav Kráľovič, Nicola Santoro, and Wei Shi. Black hole search in asynchronous rings using tokens. In *CIAC*, volume 3998 of *LNCS*, pages 139–150. Springer, 2006.
- [12] Paola Flocchini, David Ilcinkas, and Nicola Santoro. Ping pong in dangerous graphs: Optimal black hole search with pure tokens. In *Proceedings of the 22nd international symposium on Distributed Computing (DISC '08)*, volume 5218 of *LNCS*, pages 227–241. Springer-Verlag, 2008.
- [13] Ralf Klasing, Euripides Markou, Tomasz Radzik, and Fabiano Sarracco. Hardness and approximation results for black hole search in arbitrary networks. *Theoretical Computer Science*, 384(2–3):201–221, 2007.
- [14] Ralf Klasing, Euripides Markou, Tomasz Radzik, and Fabiano Sarracco. Approximation bounds for black hole search problems. *Networks*, 52(4):216–226, 2008.