Theoretical Computer Science 610 (2016) 169–181

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Exploring an unknown dangerous graph with a constant number of tokens

B. Balamohan [b], S. Dobrev [a,*], P. Flocchini [b], N. Santoro [c]

[a] Institute of Mathematics, Slovak Academy of Sciences, Bratislava, Slovak Republic
[b] School of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, Canada
[c] School of Computer Science, Carleton University, Ottawa, Canada

ARTICLE INFO

Article history:
Received 19 November 2012
Received in revised form 22 May 2014
Accepted 9 July 2014
Available online 16 July 2014

Keywords:
Black hole
Dangerous networks
Unknown graph exploration
Tokens
Asynchronous mobile agents
Distributed algorithms

ABSTRACT

Consider a team of asynchronous agents that must explore an unknown graph G in presence of a black hole, a node which destroys all incoming agents without leaving any observable trace. Inter-agent communication and coordination is achieved using tokens that agents can pick up, carry, and drop on the nodes. It is known that, if the agents have a map of G, the problem can be solved with a constant number of tokens. The question we study is under what conditions it is possible to explore with $O(1)$ tokens if the agents have no map of G. The contribution of this paper is to provide a definite answer to this question.

We first prove the unexpected negative result that if only a constant number of tokens are available, then $\Delta+1$ agents are *not* sufficient, where Δ is the maximum node degree. We also prove that, regardless of the team size, two tokens are *not* sufficient. In other words, any solution protocol using a constant number of tokens, must use at least $\Delta+2$ agents and at least three tokens.

We then show that these bounds are *tight* by presenting a protocol that allows the exploration of an unknown anonymous dangerous graph using only three tokens and $\Delta+2$ agents. At the core of our solution is a novel token-based asynchronous communication protocol, which is of independent interest.

Our algorithm assumes that the number of agents is known; in case the number of agents is sufficient but unknown, we show that the problem has a simple solution that uses only five tokens.

1. Introduction

1.1. The problem

The classic problem of exploring an unknown graph using a team of one or more mobile agents has been extensively studied since its initial formulation in 1951 by Shannon [31]. It requires the agents, starting from the same node, to visit within finite time all the nodes of a graph whose topology is unknown to them.

Different instances of the problem exist depending on a variety of factors, including the (a)synchrony of the agents, the presence of distinct agent identifiers, the amount of memory, the coordination and communication tools available to the agents, etc. (e.g., see [1,3,12–14,23,24,30]).

---

[*] Corresponding author.
E-mail addresses: stefan.dobrev@savba.sk (S. Dobrev), flocchin@site.uottawa.ca (P. Flocchini), santoro@scs.carleton.ca (N. Santoro).

http://dx.doi.org/10.1016/j.tcs.2014.07.013
0304-3975/

Notice that, except for trees, the exploration of anonymous graphs is possible only if the agents are allowed to mark the nodes in some way; various methods of marking nodes have been used by different authors ranging from the weak model of *tokens* to the most powerful model of *whiteboards*. The majority of the solutions proposed in the literature succeed in their task assuming that the network is *safe* for the agents.

The exploration problem has been examined also when the network is *unsafe*; the danger considered is the presence in the network of a *black hole* (BH), a node that disposes of any incoming agent without leaving any observable trace of this destruction (e.g., [2,6–11,15,16,18–22,25–29,32]). Note that such a dangerous presence is not uncommon. For example, a network site that contains a local program (virus) that destroys incoming mobile code is a black hole; similar effect can have just the malfunctioning of the site's communication system. From a computational point of view, the undetectable crash failure of a node in an asynchronous network transforms that node into a black hole.

The problem of exploring the graph in spite of this harmful presence is called *black hole search* (BHS). It requires the team of agents to explore the network and, within finite time, discover the location of the BH. More precisely, at least one agent must survive, and any surviving agent must have constructed a map of the network indicating the edges leading to the BH.

It is known that, when the graph is unknown (i.e., the agents have no map of the network) some information must however be available for BHS to be solvable; in particular, some metric information such as the number n of nodes and the number m of links, or the number of safe links (i.e., not leading to the black hole) is necessary for termination [18]. Furthermore, if the graph is unknown, at least $\Delta + 1$ agents are needed, where Δ is the maximum node degree in the graph [17].

The problem of asynchronous agents exploring a dangerous graph has been investigated assuming powerful inter-agent communication mechanisms: *whiteboards* at all nodes. In the whiteboard model, each node has available a local storage area (the whiteboard) accessible in fair mutual exclusion to all incoming agents; upon gaining access, the agent can write messages on the whiteboard and can read all previously written messages. This mechanism can be used by the agents to communicate and mark nodes or/and edges, and has been employed e.g. in [12,15,17,18,23]. Using whiteboards, BHS can be solved with a minimal team size and performing a polynomial number of moves (e.g., [15,17,18]). Indeed, the whiteboard model is very powerful, providing not only direct and explicit communication but also a mechanisms for leader election (and assigning distinct ids) when agents are co-located, and FIFO capabilities even when the network is not so.

An alternative coordination and communication mechanism is provided by the use of identical *tokens* (or pebbles) that an agent can hold, carry while moving, place on nodes, and pick up from a node. The *token* model is generally viewed as less taxing of the system resources than whiteboards, and has been employed in the exploration of safe graphs (e.g., [3,13]); it gives rise to additional research questions, in particular: How many tokens are needed? Can tokens be placed only on nodes (pure token model) or also on edges (hybrid token model)? etc.

These questions focus on how much *space* is required from the system by an asynchronous token-based solution. These questions are also relevant because of the relationship between solutions in the two models. In fact, any protocol which uses at most t tokens in the pure token model, can be directly implemented using whiteboards of size at most $\lceil \log t \rceil$ bits; note that h tokens in the hybrid model correspond to $t = h\Delta$ tokens in the pure token model. This measure t, which we call the *token load*, is clearly important in that it determines the usability of token-protocols in the whiteboard model, and provides a simple mechanism to transform[1] complexity results from the token setting to the whiteboard one. Importantly, the transformation preserves also other costs (such as total number of moves by agents and time).

The recent token-based solution for graphs of *known* topology with minimum number of agents and $O(1)$ pure tokens opens immediately the question of whether a similar result holds also if no map is available to the agents. In other words, is it possible to locate the black hole in a graph of *unknown* topology using a constant number of tokens? and if so, under what conditions?

The current results for unknown graphs are not very useful and provide no hints. The existing whiteboard solution uses $O(\log n)$ bits whiteboards [17], implying a solution that uses $t = O(n^2)$ tokens; the existing token-based solution uses $t = O(\Delta^2)$ tokens [16]. Similar questions have been recently raised, in the case of synchronous rings and synchronous tori [5,6].

In this paper we provide a definite answer to those questions.

1.2. Main contributions

We first examine how many agents are needed to locate a black hole without a map of the graph. We prove the unexpected negative result that if $t = O(1)$, then $\Delta + 1$ agents are *not* sufficient. That is, the minimum team size to solve the problem with a map $(\Delta + 1)$ is not sufficient to solve the problem without a map if the number of tokens is a fixed constant. This must be contrasted with the case when the graph is known: in that case, constraining the number of tokens to be constant does not affect the minimum team size.

We next examine how many tokens are really needed, i.e. how small t can be. We prove that, regardless of the team size, $t = 2$ tokens are *not* sufficient.

---

[1] Note that an automatic transformation exists also in the other direction: any solution that uses whiteboards of size s bits can be implemented with a token load at most $t = n2^s$.