ARTICLE IN PRESS YJCSS:3041

Journal of Computer and System Sciences ••• (••••) •••–•••

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Collision-free network exploration ☆

Jurek Czyzowicz [a], Dariusz Dereniowski [b,*], Leszek Gąsieniec [c,*], Ralf Klasing [d,*], Adrian Kosowski [e,*], Dominik Pająk [f]

[a] Université du Québec en Outaouais, Gatineau, Québec J8X 3X7, Canada
[b] Faculty of Electronics, Telecommunications and Informatics, Gdańsk University of Technology, 80-233 Gdańsk, Poland
[c] University of Liverpool, Liverpool L69 3BX, UK
[d] CNRS, LaBRI, Université de Bordeaux, 33405 Talence, France
[e] LIAFA, Inria Paris-Rocquencourt, 75013 Paris, France
[f] Department of Computer Science at the Faculty of Fundamental Problems of Technology, Wrocław University of Science and Technology, 50-370 Wrocław, Poland

ARTICLE INFO

Article history:
Received 8 June 2014
Received in revised form 25 November 2016
Accepted 26 November 2016
Available online xxxx

Keywords:
Mobile agents
Network exploration
Synchronous agents

ABSTRACT

Mobile agents start at different nodes of an n-node network. The agents synchronously move along the network edges in a *collision-free* way, i.e., in no round two agents may occupy the same node. An agent has no knowledge of the number and initial positions of other agents. We are looking for the shortest time required to reach a configuration in which each agent has visited all nodes and returned to its starting location. In the scenario when each mobile agent knows the map of the network, we provide tight (up to a constant factor) lower and upper bounds on the collision-free exploration time in arbitrary graphs, and the exact bound for the trees. In the second scenario, where the network is unknown to the agents, we propose collision-free exploration strategies running in $O(n^2)$ rounds in tree networks and in $O(n^5 \log n)$ rounds in networks with an arbitrary topology.

1. Introduction

The graph searching problem is a task of central importance in many contexts, including network maintenance, terrain patrolling, and robotics. Its different aspects have been thoroughly investigated, cf. [12]. The *rendezvous search* problem has been often presented as a game with two mobile players walking within the *search space* and having the common goal of arriving at the same time at the same location (see [4]). On the other hand, the *exploration problem* consists in examining all elements of the search space by a mobile agent (e.g. visiting all graph nodes or traversing all its edges), e.g., in order to find a hidden target (see [1,21]).

In this paper we consider a graph searching problem in which each of a set of mobile agents must explore a given undirected graph, in such a way that two agents may never visit the same node of the graph at the same time. This

☆ Research partially supported by the ANR project DISPLEXITY (ANR-11-BS02-014) and by NCN under contract DEC-2011/02/A/ST6/00201. This study has been carried out in the frame of "the Investments for the future" Programme IdEx Bordeaux – CPU (ANR-10-IDEX-03-02). Dariusz Dereniowski has been partially supported by a scholarship for outstanding young researchers funded by the Polish Ministry of Science and Higher Education. An extended abstract of this work appeared in Proc. 11th Latin American Theoretical INformatics Symposium (LATIN 2014), LNCS 8392, pp. 342–354 [9].

* Corresponding authors.
E-mail addresses: deren@eti.pg.gda.pl (D. Dereniowski), l.a.gasieniec@liverpool.ac.uk (L. Gąsieniec), ralf.klasing@labri.fr (R. Klasing), adrian.kosowski@inria.fr (A. Kosowski).

http://dx.doi.org/10.1016/j.jcss.2016.11.008
0022-0000/

Please cite this article in press as: J. Czyzowicz et al., Collision-free network exploration, J. Comput. Syst. Sci. (2016), http://dx.doi.org/10.1016/j.jcss.2016.11.008

Table 1
The time of optimal collision-free graph exploration. $\Delta(G)$ denotes the maximum degree of a node in graph G, and $\Delta^*(G) = \Delta(T)$, where T is a minimum-degree spanning tree of G.

Scenario	Tree	General graph
With complete map:	$n\Delta(G)$ Theorem 2.2	$\Theta(n\Delta^*(G))$ Theorem 2.4
With local knowledge:	$O(n^2)$ Theorem 3.1	$O(n^5 \log n)$ Theorem 3.2

property of the model, which we call *collision avoidance*, is motivated by the fact that the processes executed by mobile agents (software agents or physical robots) sometimes require exclusive access to network resources. Our problem may have practical applications. For example, mobile software agents may need exclusive access to a node's resources when updating its data. Robots (or nano-robots) distributing interacting chemical or pharmacological agents within a battlefield or a human body must avoid to be simultaneously present at a small distance apart. Individuals, one of which is highly infectious or socially conflicting, should avoid a meeting. This problem has been studied in [20] for the case of tree networks. A question related to the offline version for our problem has been given some attention in the context of routing (cf. [3]).

In our considerations, time is divided into synchronous rounds. Initially, each agent is placed at a different node and in each round it may choose to move to a neighboring node or to stay motionless. The agents are independent in the sense that they cannot communicate and none of them knows the number of other agents, their initial placement in the graph, and is unaware of the current location of the other agents. The agents move independently, and each of them executes the same algorithm. The effectiveness of the algorithm is measured in terms of the *collision-free exploration time*, i.e., the number of rounds until all potentially existing agents are certain to have completed the exploration and returned to their initial location. Details of our model are discussed at the end of this section.

1.1. Our results

We consider two scenarios, differing in the amount of global information about the network topology which is available to each agent. Our results are summarized in Table 1.

For the first scenario, considered in Section 2, we assume that a map of the network is *a priori* known to the agents. We show that a collision-free exploration strategy exists for any graph, and provide efficient solutions for trees and general graphs. We start by considering the case of trees, proposing a strategy which involves the simultaneous activation of agents located at the endpoints forming a matching in some optimal edge-coloring of the tree. This strategy is shown to yield optimal exploration time. We then extend this approach from the case of trees to the case of general graphs, by requiring that the agents perform exploration using only the edges of a well-chosen spanning tree of the graph. Somewhat surprisingly, it turns out that this approach is asymptotically the best possible, i.e., within a constant factor of the optimum. To prove the corresponding lower bound on the collision-free exploration time in graphs, we establish a tight connection between our problem and the fractional relaxation of the LP formulation of the minimum-degree spanning tree problem.

In the second scenario, discussed in Section 3, we deal with synchronous agents possessing only local knowledge about the graph to explore. In particular, no knowledge of the size of the graph is assumed. We suppose that each agent executes a local, distributed algorithm, in every round making a decision based on the information concerning the currently occupied node and the identifiers of the neighboring nodes. For this scenario, we show that a collision-free exploration is always feasible in finite time and we give algorithms for trees and general graphs. Our collision-free exploration strategies are of length $O(n^2)$ for trees and $O(n^5 \log n)$ for arbitrary graphs, and make use of the application of universal exploration sequences.

Throughout the paper, we assume that the strategies for collision-free exploration are required to return the agent to their initial location. This assumption allows us to see our strategies as an analogue of the classical Traveling Salesman Problem with mutually-exclusive salesmen on an unweighted graph, and also allows the agents to engage in perpetual (periodic) exploration of the graph. After minor modification of the proofs, all the results presented in Table 1 also hold up to constant factors for the variant of the problem in which agents may end exploration at an arbitrary node of the graph.

1.2. Related work

We remark that a similar process was studied in a different context by Nakaminami et al. [20], who obtained a $\Theta(\Delta n)$ bound for trees. Herein, we provide our own analysis of this case for the sake of completeness. We then generalize this analysis to the case of arbitrary graphs, and also to the case where agents do not have a map of a graph. Our analysis for trees shows in a simple way that the number of steps required in a tree is precisely Δn.

The offline setting of our question is related to the following problem (cf. [3]), which was studied in the context of routing. Each vertex of a given graph is initially occupied by a "pebble", which has to be moved to a destination, so that the destinations of different pebbles are different. In every synchronous round a set of edges is selected and the pebbles

Please cite this article in press as: J. Czyzowicz et al., Collision-free network exploration, J. Comput. Syst. Sci. (2016), http://dx.doi.org/10.1016/j.jcss.2016.11.008