Information and Computation 247 (2016) 254–265

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Topology recognition with advice ☆

Emanuele G. Fusco [a,*], Andrzej Pelc [b,1], Rossella Petreschi [c]

[a] *Department of Computer, Control, and Management Engineering Antonio Ruberti, Sapienza, University of Rome, 00185 Rome, Italy*
[b] *Département d'informatique, Université du Québec en Outaouais, Gatineau, Québec J8X 3X7, Canada*
[c] *Computer Science Department, Sapienza, University of Rome, 00198 Rome, Italy*

A R T I C L E   I N F O

Article history:
Received 24 October 2013
Received in revised form 4 June 2015
Available online 21 January 2016

Keywords:
Topology recognition
Network
Advice
Time
Tradeoff

A B S T R A C T

In topology recognition, each node of an anonymous network has to deterministically produce an isomorphic copy of the underlying graph, with all ports correctly marked. This task is usually unfeasible without any a priori information. Such information can be provided to nodes as advice. An oracle knowing the network can give a (possibly different) string of bits to each node, and all nodes must reconstruct the network using this advice, after a given number of rounds of communication. During each round each node can exchange arbitrary messages with all its neighbors and perform arbitrary local computations. The time of completing topology recognition is the number of rounds it takes, and the size of advice is the maximum length of a string given to nodes.

We investigate tradeoffs between the time in which topology recognition is accomplished and the minimum size of advice that has to be given to nodes. We provide upper and lower bounds on the minimum size of advice that is sufficient to perform topology recognition in a given time, in the class of all graphs of size n and diameter $D \leq \alpha n$, for any constant $\alpha < 1$. In most cases, our bounds are asymptotically tight.

1. Introduction

1.1. The model and the problem

Learning an unknown network by its nodes is one of the fundamental distributed tasks in networks. Once nodes acquire a faithful labeled map of the network, any other distributed task, such as leader election [21,28], minimum weight spanning tree construction [5], renaming [4], etc. can be performed by nodes using only local computations. Knowing topology also simplifies the task of routing. More generally, constructing a labeled map converts all distributed problems to centralized ones, in the sense that nodes can solve them simulating a central monitor.

If nodes are a priori equipped with unique identifiers, they can deterministically construct a labeled map of the network, by exchanging messages, without any additional information about the network. However, even if nodes have unique identities, relying on them for the task of learning the network is not always possible. Indeed, nodes may be reluctant to reveal their identities for security or privacy reasons. Hence it is important to design algorithms reconstructing the topology of the

☆ This research was done during the visit of Andrzej Pelc at Sapienza, University of Rome, partially supported by a visiting fellowship from this university. A preliminary version of this paper appeared in the Proceedings of the 27th International Symposium on Distributed Computing (DISC 2013).

* Corresponding author.
E-mail addresses: fusco@diag.uniroma1.it (E.G. Fusco), pelc@uqo.ca (A. Pelc), petreschi@di.uniroma1.it (R. Petreschi).

[1] Partially supported by NSERC discovery grant 8136–2013 and by the Research Chair in Distributed Computing at the Université du Québec en Outaouais.

http://dx.doi.org/10.1016/j.ic.2016.01.005
0890-5401/

network without assuming any node labels, i.e., for anonymous networks. In this paper we are interested in deterministic solutions.

Ports at each node of degree d are arbitrarily numbered $0, \dots, d-1$, and there is no assumed coherence between port numbers at different nodes. A node is aware of its degree, and it knows on which port it sends or receives a message. The goal is, for each node, to get an isomorphic copy of the graph underlying the network, with all port numbers correctly marked. (Port numbers are very useful, e.g., to perform the routing task.) There are two variants of this task: a weaker version, that we call *anonymous topology recognition*, in which the nodes of the reconstructed map are unlabeled, and a stronger version, that we call *labeled topology recognition*, in which all nodes construct a map of the network assigning distinct labels to all nodes in the same way, and know their position in this map. Even anonymous topology recognition is not always feasible without any a priori information given to nodes, as witnessed, e.g., by the class of oriented rings in which ports at each node are numbered 0,1 in clockwise order. No amount of information exchange can help nodes to recognize the size of the oriented ring and hence to reconstruct correctly its topology. Thus, in order to accomplish (even anonymous) topology recognition for arbitrary networks, some information must be provided to nodes. This can be done in the form of *advice*. An oracle knowing the network gives a (possibly different) string of bits to each node. Then nodes execute a deterministic distributed algorithm that does not assume knowledge of the network, but uses message exchange and the advice provided by the oracle to nodes, in order to reconstruct the topology of the network by each of its nodes.

In this paper we study tradeoffs between the *size of advice* provided to nodes and the *time* of topology recognition. The size of advice is defined as the length of the longest string of bits given by the oracle to nodes. For communication, we use the extensively studied $\mathcal{LOCAL}$ model [27]. In this model, communication proceeds in synchronous rounds and all nodes start simultaneously. In each round each node can exchange arbitrary messages with all its neighbors and perform arbitrary local computations. The time of completing a task is the number of rounds it takes. The central question of the paper is:

> What is the minimum size of advice that enables (anonymous or labeled) topology recognition in a given time T, in the class of n-node networks of diameter D?

It should be stressed that the reason why we adopt the $\mathcal{LOCAL}$ model is to focus on how deeply nodes should probe the network in order to discover the topology. Since nodes are anonymous, this depth of probing may often go beyond the diameter of the network, enabling the nodes to "see" other nodes several times, along different paths, and deduce some information from these deeper views. At the same time, as it is always done when using the $\mathcal{LOCAL}$ model [27], the size of messages is ignored.

1.2. Our results

We provide upper and lower bounds on the minimum size of advice sufficient to perform topology recognition in a given time, in the class $\mathcal{C}(n, D)$ of all graphs of size n and diameter $D \le \alpha n$, for any constant $\alpha < 1$. All our upper bounds are valid even for the harder task of labeled topology recognition, while our lower bounds also apply to the easier task of anonymous topology recognition. Hence we will only use the term *topology recognition* for all our results. We prove upper bounds $f(n, D, T)$ on the minimum size of advice sufficient to perform topology recognition in a given time T, for the class $\mathcal{C}(n, D)$, by providing an assignment of advice of size $f(n, D, T)$ and an algorithm, using this advice, that accomplishes this task, within time T, for any network in $\mathcal{C}(n, D)$. We prove lower bounds on the minimum size of advice, sufficient for a given time T, by constructing graphs in $\mathcal{C}(n, D)$ for which topology recognition within this time is impossible with advice of a smaller size. (Notice that, while our upper bounds would hold even if $D \approx n$, our lower bound constructions rely on the constraint $D \le \alpha n$.)

The meaningful span of possible times for topology recognition is between 0 and $2D + 1$. Indeed, while advice of size $O(n^2 \log n)$ permits topology recognition in time 0 (i.e., without communication), we show that topology recognition in time $2D + 1$ can be done with advice of size 1, which is optimal.

For most values of the allotted time, our bounds are asymptotically tight. This should be compared to many results from the literature on the advice paradigm (see, e.g., [9,14,15,17,26]), which often either consider the size of advice needed for feasibility of a given task, or only give isolated points in the curve of tradeoffs between resources (such as time) and the size of advice.

We show that, if the allotted time is $D-k$, where $0 < k \le D$, then the optimal size of advice is $\Theta((n^2 \log n)/(D-k+1))$. If the allotted time is D, then this optimal size is $\Theta(n \log n)$. If the allotted time is $D+k$, where $0 < k \le D/2$, then the optimal size of advice is $\Theta(1 + (\log n)/k)$. The only remaining gap between our bounds is for time $D+k$, where $D/2 < k \le D$. In this time interval our upper bound remains $O(1 + (\log n)/k)$, while the lower bound (that holds for any time) is 1. This leaves a gap if $D \in o(\log n)$. See Table 1 for a summary of our results.

Our results show how sensitive is the minimum size of advice to the time allowed for topology recognition: allowing just one round more, from D to $D+1$, decreases exponentially the advice needed to accomplish this task. Our tight bounds on the minimum size of advice also show a somewhat surprising fact that the amount of information that nodes need to reconstruct a labeled map of the network, in a given time, and that needed to reconstruct an anonymous map of the network in this time, are asymptotically the same in most cases.