Information and Computation ••• (••••) •••–•••

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Randomised distributed MIS and colouring algorithms for rings with oriented edges in $O(\sqrt{\log n})$ bit rounds

Y. Métivier *, J.M. Robson, A. Zemmari

Université de Bordeaux - Bordeaux INP, LaBRI, CNRS UMR 5800, 351 cours de la Libération, 33405 Talence, France

ARTICLE INFO

Article history:
Received 27 May 2014
Received in revised form 6 May 2016
Available online xxxx

ABSTRACT

We present and analyse Las Vegas distributed algorithms which compute a MIS or a colouring for anonymous rings with an arbitrary orientation of the edges; their bit complexity and time complexity are $O(\sqrt{\log n})$ with high probability. These algorithms are optimal modulo a multiplicative constant.

1. Introduction

1.1. The problem

Let $G = (V, E)$ be a simple connected undirected graph. An independent set is a subset I of V such that no two members of I are adjacent. An independent set I is said to be maximal (MIS for short) if any vertex of G is in I or adjacent to a vertex of I. A vertex *colouring* of G assigns colours to each vertex in such a way that neighbours have different colours. If we need at most 3 colours then we call it a 3-colouring.

In this paper we discuss how greedy selection for the computation of a maximal independent set or of a vertex colouring in a ring of processors such that edges have an arbitrary orientation can be accomplished by exchange of messages of size $O(1)$ between adjacent processors.

Usually, the topology of a distributed system is modelled by a graph and paradigms of distributed systems are encoded by classical problems in graph theory; among these classical problems one may cite the problems of vertex colouring, computing a maximal independent set, finding a vertex cover, finding a maximal matching or finding a graph decomposition. Each solution to one of these problems is a building block for many distributed algorithms: symmetry breaking, topology control, routing, resource allocation, network synchronisation.

Even if ring graphs are simple, they are used as a case study in many problems, as explained by Attiya and Welch in [3], page 31: *"rings are a convenient structure for message-passing systems and correspond to physical communication systems, for example, token rings."*

1.2. The model

A general presentation may be found in [26] (Chapter 9) or in [19] (Chapter 8).

* Corresponding author.
E-mail addresses: yves.metivier@labri.fr (Y. Métivier), mike.robson@labri.fr (J.M. Robson), akka.zemmari@labri.fr (A. Zemmari).

http://dx.doi.org/10.1016/j.ic.2016.09.006
0890-5401/

The network We consider the standard message passing model for distributed computing. The communication model consists of a point-to-point communication network described by a simple ring graph $G = (V, E)$ where the vertices V represent network processors and the edges represent bidirectional communication channels. Processors communicate by message passing: a processor sends a message to another by depositing the message in the corresponding channel. Note that we consider only reliable systems: no fault can occur on processors or communication links.

We assume that the system is synchronous with simultaneous wakeup of processors: processors have access to a global clock and all processors start the algorithm at the same time.

Time complexity A round (cycle) of each processor is composed of the following three steps: 1. Send messages to (some of) the neighbours, 2. Receive messages from (some of) the neighbours, 3. Perform some local computation. As usual (see for example Peleg [23]) the time complexity is the maximum possible number of rounds needed until every node has completed its computation.

Bit complexity A bit round is a round such that each processor can send (and receive) at most 1 bit to (and from) each neighbour. As in [13], the bit complexity of an algorithm $\mathcal{A}$ is the number of bit rounds to complete algorithm $\mathcal{A}$. One round of the algorithm contains 1 or more bit rounds.

Remark 1. If we consider a distributed algorithm using messages of size $O(1)$ (and this is the case in this paper) then the bit complexity and the time complexity are equivalent modulo a multiplicative constant.

Network and processor knowledge The network is anonymous: unique identities are not available to distinguish the processors. We do not assume any knowledge on the size of the ring or on an upper bound on the size of the ring, or any position or distance information. Each processor knows from which channel it receives a message. An important fact due to the initial symmetry is: *there is no deterministic distributed algorithm for anonymous ring graphs for solving the MIS problem or the maximal matching problem assuming all vertices wake up simultaneously,* see [23].

Las Vegas distributed algorithms A probabilistic algorithm is an algorithm which makes some random choices based on some given probability distributions.

A distributed probabilistic algorithm is a collection of local probabilistic algorithms. Since our network is an anonymous ring, two processes have the same degree thus their local probabilistic algorithms are identical and have the same probability distribution.

A Las Vegas algorithm is a probabilistic algorithm which terminates with probability 1 and always produces a correct result.

1.3. Our contribution

We present and analyse Las Vegas distributed algorithms which compute a MIS or a colouring for anonymous rings having oriented edges. Their bit complexity and time complexity are $O(\sqrt{\log n})$ with high probability[1] (w.h.p. for short). From [13] (Theorem 2.2), we deduce that the colouring algorithm is optimal (modulo a multiplicative constant). As it is easy to deduce from a MIS a 3-colouring in a ring with oriented edges (see Section 3), we deduce that the MIS algorithm is also optimal (modulo a multiplicative constant).

The descriptions and the analyses of these algorithms follow the scheme introduced in [9] and use rewriting rules. We think that the algorithm description in the framework of rewriting rules is simpler than the standard algorithmic description for this kind of algorithm.

This result shows the impact of the orientation of edges on the complexity of the colouring problem in rings: without orientation of the edges, colouring can be computed only in $\Omega(\log n)$ rounds on rings with high probability (see [13]). Our result improves the result presented in [13]. To obtain a $O(\sqrt{\log n})$ colouring algorithm for rings Kothapalli et al. assume in [13] that edges have the same orientation: the ring is a directed cycle.

1.4. Related works: comparison and comments

1.4.1. Bit complexity and single bit messages

Classically, there are two models for the size of the messages: the LOCAL model and the CONGEST model (see [23] p. 27). The first allows message of unlimited size while the second allows messages with a size bounded by $O(\log n)$ (n is the size of the network). In both models vertices have unique identifiers.

The model of this paper is anonymous (no uniqueness of identifiers for the nodes) and nodes have no global knowledge on the network such as its size. Thus, in this context, when a processor builds a message its size cannot depend on the size of the network and it is natural to consider single bit messages or more generally messages with bounded sizes.

[1] With high probability means with probability $1 - o(n^{-1})$.

Please cite this article in press as: Y. Métivier et al., Randomised distributed MIS and colouring algorithms for rings with oriented edges in $O(\sqrt{\log n})$ bit rounds, Inf. Comput. (2016), http://dx.doi.org/10.1016/j.ic.2016.09.006