Theoretical Computer Science 509 (2013) 86–96

Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Gathering asynchronous oblivious agents with local vision in regular bipartite graphs

Samuel Guilbault, Andrzej Pelc *

Département d'informatique, Université du Québec en Outaouais, Gatineau, Québec J8X 3X7, Canada

ARTICLE INFO

Keywords:
Asynchronous
Mobile agent
Gathering
Regular graph
Bipartite graph

ABSTRACT

We consider the problem of gathering identical, memoryless, mobile agents in one node of an anonymous graph. Agents start from different nodes of the graph. They operate in Look-Compute-Move cycles and have to end up in the same node. In one cycle, an agent takes a snapshot of its immediate neighborhood (Look), makes a decision to stay idle or to move to one of its adjacent nodes (Compute), and in the latter case makes an instantaneous move to this neighbor (Move). Cycles are performed asynchronously for each agent. The novelty of our model with respect to the existing literature on gathering asynchronous oblivious agents in graphs is that the agents have very restricted perception capabilities: they can only see their immediate neighborhood.

An initial configuration of agents is called *gatherable* if there exists an algorithm that gathers all the agents of the configuration in one node and keeps them idle from then on, regardless of the actions of the asynchronous adversary. (The algorithm can be even tailored to gather this specific configuration.) The gathering problem is to determine which configurations are gatherable and find a (universal) algorithm which gathers all gatherable configurations. We give a complete solution of the gathering problem for regular bipartite graphs. Our main contribution is the proof that the class of gatherable initial configurations is very small: it consists only of "stars" (an agent *A* with all other agents adjacent to it) of size at least 3. On the positive side we give an algorithm accomplishing gathering for every gatherable configuration.

1. Introduction

The aim of gathering is to bring mobile entities (agents), initially situated at different locations of some environment, to the same location (not determined in advance) and stop. Agents may operate either in the plane, in which case they usually represent mobile robots, or in a network, modeled by a simple undirected graph, in which case they may model software agents. Gathering is a basic task which permits, e.g, to exchange information between agents or coordinate their further actions.

A lot of effort has been devoted to studying gathering in very weak scenarios, where agents represent simple devices that could be cheaply mass-produced. One of these scenarios is the CORDA model, initially formulated for agents operating in the plane [4–6,11,16–18] and then adapted to the network environment [12–14]. In this paper we study gathering in networks, in a scenario even weaker than the above. Below we describe our model and point out its differences with respect to that from [12–14].

Consider a simple undirected graph. Neither nodes nor links of the graph have any labels. Initially, some nodes of the graph are occupied by identical agents and there is at most one agent in each node. The goal is to gather all agents in one node

* Corresponding author. Tel.: +1 819 595 3900.
E-mail addresses: samuel.guilbault@gmail.com (S. Guilbault), andrzej.pelc@uqo.ca, pelc@uqo.ca (A. Pelc).

0304-3975/$ – see front matter
doi:10.1016/j.tcs.2012.07.004

of the graph and stop. Agents operate in Look-Compute-Move cycles. In one cycle, an agent takes a snapshot of its immediate neighborhood (Look), then, based on it, makes a decision to stay idle or to move to one of its adjacent nodes (Compute), and in the latter case makes an instantaneous move to this neighbor (Move). Cycles are performed asynchronously for each agent. This means that the time between Look, Compute, and Move operations is finite but unbounded, and is decided by the adversary for each agent. The only constraint is that moves are instantaneous, and hence any agent performing a Look operation can see other agents at its own or adjacent nodes and not on edges, while performing a move. However an agent A may perform a Look operation at some time t, perceiving agents at some adjacent nodes, then Compute a target neighbor at some time $t' > t$, and Move to this neighbor at some later time $t'' > t'$ in which some agents are in different nodes from those previously perceived by A, because in the meantime they performed their Move operations. Hence agents may move based on significantly outdated perceptions, which adds to the difficulty of achieving the goal of gathering. It should be stressed that agents are memoryless (oblivious), i.e., they do not have any memory of past observations. Thus the target node (which is either the current position of the agent or one of its neighbors) is decided by the agent during a Compute operation solely on the basis of the location of other agents perceived in the previous Look operation. Agents are anonymous and execute the same deterministic algorithm. They cannot leave any marks at visited nodes, nor send any messages to other agents.

The only difference between our scenario and that from [12–14] is in what an agent perceives during the Look operation. While in the above papers the agent was assumed to see the entire configuration of all agents in the graph, we assume that it only sees its immediate neighborhood, i.e., agents located at its own and at adjacent nodes. The reason for this change of model is applicability. It is hard to imagine how otherwise weak and simple agents could implement a global snapshot of the network, whereas local perception can be easily performed by exchanging signals between adjacent nodes.

An important and well studied capability in the literature on agent gathering is the *multiplicity detection* [11–14,17]. This is the ability of the agents to perceive, during the Look operation, if there is one or more agents at a given location. It has been shown in [14] that without this capability gathering in networks is usually impossible (even for the ring and even if agents can take global snapshots). On the other hand, it was proved in [12] that multiplicity detection can be weakened to apply only to the node on which the agent is currently located. In this case, multiplicity detection is called *local*. In this paper we consider both the local and the global version of multiplicity detection. Our negative result holds even with *global* multiplicity detection (which means, in our case, that an agent can detect if there is one or more agents at its own and at adjacent nodes) and our positive result holds even for *local* multiplicity detection, when the agent can only distinguish if it is alone or not at its node. It should be stressed that, during a Look operation, an agent can only tell if at some node there are no agents, there is one agent, or there are more than one agents: an agent does not see a difference between a node occupied by a or b agents, for distinct $a, b > 1$.

In this paper we study the gathering problem in regular bipartite graphs. A configuration of agents is called *gatherable* if there exists an algorithm that gathers all the agents of the configuration in one node and keeps them idle from then on, regardless of the actions of the asynchronous adversary. (The algorithm can be even tailored to gather this specific configuration.) The gathering problem is to determine which initial configurations are gatherable and find a (universal) algorithm which gathers all initial gatherable configurations.

1.1. Our results

We give a complete solution of the gathering problem for regular bipartite graphs. Our main contribution is the proof that the class of gatherable initial configurations is very small: it consists only of "stars" (an agent A with all other agents adjacent to it) of size at least 3. On the positive side we give a (universal) algorithm accomplishing gathering for every gatherable configuration.

1.2. Discussion of assumptions

We would like to argue that our model is likely to be the weakest under which the gathering problem can be meaningfully studied in networks. Already the CORDA model adapted to networks in [12–14] is extremely weak, as agents do not have memory of previous snapshots and the asynchronous adversary has the full power of arbitrarily scheduling Look-Compute-Move cycles for all agents. Also agents are identical (anonymous), hence gathering protocols cannot rely on labels to break symmetry. However, the perception capabilities of agents were previously assumed very strong, thus giving rise to agents with contrasting features: no memory but powerful view. In our model the latter power is also weakened, as we allow the agents to perceive only their immediate neighborhood. It is easy to see that further restrictions are impossible: agents that see only their own node cannot gather. As for the assumption on multiplicity detection, we solve the gathering problem both for its local and global version.

It remains to discuss our assumptions concerning the class of networks in which agents operate. We look at regular bipartite graphs. This is a large class of networks including e.g., even cycles, multidimensional tori with even sides, hypercubes and many other networks. It should be noted that our result remains valid without change in arbitrary cycles of size different from 3 and arbitrary multidimensional tori with every dimension of size different from 3. For multidimensional tori with at least one dimension of size 3, the result changes as follows. For global multiplicity detection, gatherable