Collecting information by power-aware mobile agents (abstract)*

Julian Anaya[1], Jérémie Chalopin[2], Jurek Czyzowicz[1],
Arnaud Labourel[2], Andrzej Pelc[1] , Yann Vaxès[2]
[1]Université du Québec en Outaouais, C.P. 1250, succ. Hull, Gatineau, Qc. J8X 3X7 Canada.
[2] LIF, CNRS & Aix-Marseille University, 13453 Marseille, France.

The model and the problem. A set of identical, mobile agents is deployed in a weighted network. Every agent has a battery: a power source allowing it to move along the network edges. The movements of an agent use its battery proportionally to the distance traveled. Each agent starts with the same amount of power noted P, allowing all agents to travel the same distance P. Initially, each agent has an individual piece of information. When two (or more) agents are at the same point of the network at the same time, they exchange their information. The purpose of a *convergecast* algorithm is to schedule the movements of the agents, so that the exchanges of the currently possessed information between the meeting agents eventually result in some agent holding the union of individual information of all the agents. This task is important, e.g., when agents have partial information about the topology of the network and the aggregate information can be used to construct a map of it, or when individual agents hold measurements performed by sensors located at their initial positions and collected information serves to make some global decision based on all measurements. Agents try to cooperate so that the convergecast is achieved with the agent's smallest possible initial battery power P_{OPT}, i.e., minimizing the maximum distance traveled by an agent. We investigate the problem in two possible settings, centralized and distributed.

In the centralized setting, the problem must be solved by a centralized authority knowing the network and the initial positions of all the agents. We consider two different versions of the problem: the decision problem, i.e., deciding if convergecast can be achieved with power P (where P is the input of the problem) and the optimization problem, i.e., computing the smallest amount of power that is sufficient to achieve convergecast.

In the distributed setting, the problem must be approached individually by each agent. Each agent is unaware of the network, of its position in the network and without the knowledge of positions of any other agents. The agents are anonymous, i.e., they must execute the same algorithm. The agent has a very simple sensing device allowing it to detect the presence of other agents at its current location in the network. The agent is also aware of the degree d of the node at which it is located and it can identify all ports represented by integers $1, 2, \ldots d$. The agent is aware of the directions from which are coming all agents currently present at its location. Each agent has memory sufficient to store all information initially belonging to all agents as well as a small (constant) number of real values. The efficiency of a distributed solution is expressed by the *competitive ratio*, which is the worst-case ratio of the amount of power necessary to solve the convergecast by the distributed algorithm with respect to the amount of power computed by the optimal centralized algorithm, which is executed for the same agents' initial positions.

Our results. In the case of centralized setting we give a linear-time deterministic algorithm finding an optimal convergecast strategy for line networks. We show that, already for the case of tree networks, the centralized problem is strongly NP-complete. We give a 2-approximation centralized algorithm for general graphs. For the distributed setting, we show that the convergecast is possible for tree networks if all agents have the amount of initial power equal to twice the power necessary to achieve centralized convergecast. The competitive ratio of 2 is proved to be the best possible for this problem, even if we only consider line networks.

---

*An extended version of this paper appeared in DISC 2012.