Albert-Ludwigs-Universität Freiburg
Dept. of Computer Science
Algorithms and Complexity

Scheduling Algorithms in Radio Networks

Bachelor or Master Thesis (or Project)

Project Description

Assume that we are given a network that is represented by a graph $G = (V, E)$ and assume that we are given a collection of distributed algorithms $\mathcal{A}_1, \dots, \mathcal{A}_k$ that we want to run on the network G. Each of the distributed algorithms $\mathcal{A}_i$ consists of a sequence of synchronous time slots (or rounds) and in each time slot, algorithm $\mathcal{A}_i$ sends a message consisting of B bits over each edge of G. Typically, we assume that $B = O(\log n)$, where n is the number of nodes of G. Finally, assume that the running time of algorithm $\mathcal{A}_i$ is T_i (the running time of an algorithm is the total number of time slots required).

The algorithms $\mathcal{A}_1, \dots, \mathcal{A}_k$ now need to be run at the same time in parallel on G. When doing this, we assume that over each edge, we can still send at most B bits per time slots. If different algorithms want to use the same edge, they therefore have to use it in different time slots. Assume that for each edge $e \in E$, C_e is the total number of messages sent over edge e, summed over all algorithms and all time slots. Let $C := \max_{e \in E} C_e$ and let $T := \max_{i \in \{1,\dots,k\}} T_i$. It is clear that when running the k algorithms in parallel, the total running time has to be at least C and it also has to be at least T. In a paper a few years ago, Ghaffari showed that it is always possible to run the k algorithms in parallel in time $O((C + D) \cdot \log^c n)$ for some constant $c \geq 0$ [Gha15]. That is, the trivial lower bound of $\Omega(C + D)$ can be achieved up to a polylogarithmic factor.

The objective of this project will be to generalize the result of [Gha15] to different distributed communication models. In particular, we are interested in generalizing it to a standard radio network model. Here, it is still possible to send messages of the edges of a graph. However, in each time slot, a node always locally broadcasts the same message to all its neighbors and a receiving node can only receive a message if there is exactly one neighbor that sends a message.

Requirements

- mathematical maturity, interest in mathematical questions
- algorithm theory and/or distributed systems (network algorithms) lecture (or comparable lectures) are an advantage

Contact

- **Fabian Kuhn**, kuhn@cs.uni-freiburg.de

References

[Gha15] Mohsen Ghaffari. Near-optimal scheduling of distributed algorithms. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*, pages 3–12, 2015.