Weizmann Institute of Science
הפקולטה למתמטיקה ולמדעי המחשב
Faculty of Mathematics and
Computer Science

דוד פלג
מכון ויצמן למדע
טלפון 08 934 3478
פקס 08 934 6598

david.peleg@weizmann.ac.il

David Peleg
Norman D. Cohen Professorial Chair
of Computer Science
Weizmann Institute of Science
76100 Rehovot, Israel
Phone: 972 8 934 3478
Fax: 972 8 934 6598

Nov. 13, 2020

Review of the Ph.D. thesis
"On Spanning Trees and Small Cuts in Congested Clique and MPC"
By Krzysztof Nowicki

Summary

The thesis studies a number of graph-algorithmic problems in the distributed setting, and specifically in two models of distributed computing known as the *Congested Clique (CC)* model and the *Massively Parallel Computation (MPC)* model. The problems studied are related to connectivity, cuts, and (minimum weight) spanning trees and forests, and the focus is on the development of (a) graph theoretic understanding and techniques, and (b) computational tools and procedures in the CC and MPC models, that will enable more efficient solutions for the problems under consideration.

Overall, the thesis consists of five main chapters, each dealing with a different problem, plus a chapter of summary and open problems.

Chapter 1 of the thesis introduces the area, the models and the problems, and also provides some of the computational tools in the CC and MPC models to be used in the algorithms presented later on. In particular, it presents some basic procedures, e.g., for sorting and grouping, and also compares the two models and discusses reductions between them, establishing the extent to which algorithms devised in one model can be translated to the other model and vice versa.

Chapter 2 studies the problem of distributedly constructing a minimum-weight spanning tree (MST) for a given graph. The main result presented is an efficient deterministic distributed algorithm for solving the MST problem on a given graph G with n vertices and m edges. The algorithm works in the CC model in O(1) rounds using O(m) communication, and in the MPC model in O(1) rounds using O(n) memory per machine and O(m) global memory. This is a strong result that concludes a long sequence of previous improvements on the time complexity of the problem (most of which were obtained for randomized algorithms), from O(logn) via O(loglogn), O(logloglogn) and O(log*n), to O(1). The result is obtained by means of some insightful observations on the nature of the MST

problem, that enable more efficient processing of the input graph. In particular, the given graph is sparsified in a clever way (via several steps) while preserving the edges of the original MST.
As a by-product of the developed techniques, the chapter presents, along with the solution for MST, also efficient algorithmic solutions for the related problems of constructing a spanning forest, deciding graph connectivity, and identifying the connected components of the input graph.

Chapter 3 studies the edge connectivity problem. It presents a novel method for contracting non-essential edges (i.e., ones that do not belong to a minimum cut) in order to sparsify the graph. This method is used in order to develop improved randomized (high probability) distributed algorithms for edge connectivity in the CC and MPC models, again with time complexity $O(1)$, using $O(m)+\tilde{O}(n)$ bits of communication.

Chapter 4 studies the *Minimum Cut* problem, mainly in the MPC model. Two algorithms are presented. The first is an exact algorithm that requires $O(1)$ time, $\tilde{O}(n)$ memory per machine and $\tilde{O}(m)$ global memory. The second algorithm is an approximation algorithm with approximation ratio $2+\varepsilon$, $O(\log n \cdot \log\log n)$ time, $O(n^{\alpha})$ memory per machine and $\tilde{O}(m)$ global memory (for any constant $\alpha<1$ and $\varepsilon>0$). Both algorithms rely on existing methods, skillfully adapted to the distributed setting of the MPC model. The chapter presents also some partial results for the CC model.

In Chapter 5, it is shown how to use the contraction method introduced in Chapter 3 in order to obtained improved algorithms for some of the above problems in other models, such as the distributed CONGEST model (which does not assume that the underlying communication model is complete), the Broadcast CC model (where the messages sent by a machine to the other machines on a given round are identical), and the standard sequential model.

Evaluation

The research contributions presented in this thesis are interesting and novel, and constitute a significant progress over the best-known results to date. This refers both to the background development of efficient procedures in the CC and MPC models and graph-theoretic insights on the problem, and to the final algorithms obtained for the problems under study.

Most impressive in my opinion, in terms of novelty and ingenuity, are the new methods developed in Chapters 2 and 3. It is reasonable to expect that these techniques will attract attention and find additional applications in the future.

The distributed algorithms and their analysis are also interesting technically, and demonstrate creativity, strong analytic abilities and familiarity with the state of the art ideas and methods in this area.

The write-up is clear and readable and the presentation is easy to follow. Some minor comments are included next.

Minor comments

While the general writeup of the thesis is very good, section 1.2 of the introduction, which presents the models and some of their basic properties, appears to be written too tersely, perhaps aiming at readers already familiar with the models, so some issues are discussed sketchily and leave unclear points. This is unfortunate since understanding the thesis relies on familiarity with these models. I will mention briefly just a few of these issues, but generally, it would be a good idea to explain the models in more detail.

The definition of the CC model is informal and leaves some aspects unspecified. This is significant because the literature has more than one variant of the CC model. In particular, in one "pure" variant of CC, the underlying communication network is explicit, i.e., each machine has n-1 physical ports connecting it via communication channels to the other machines. In contrast, in another (more realistic) variant of the CC model the network is implicit, i.e., a machine has a single port giving it access to a communication mechanism (or routing protocol), and can use it to send out a message to another machine only if it initially knows, or has already learned, the ID of that machine. In the latter model, it is important to define whether each vertex initially knows the ID's of all other vertices (this does not really make a difference in the former model).

In the description of the MPC model, the notion of "global memory" is not defined explicitly. It may be understood as referring to the availability of some external memory accessible by all machines, like in the PRAM model, but I assume the intention is just to the "sum" of the local memories. This makes a difference, and should be defined explicitly.

The 3rd paragraph in subsection 1.2.2 is confusing. The description of the model seems to imply that the local memory at each machine may be unbounded, and there is a bound only on the amount of information that can be communicated to / from each machine in each round (namely, each machine may be a source and a destination of O(S) messages). Then, however, it is stated that

> "the parameter S is a bound on communication of a single machine, but one can also think of it as a bound on the amount of information (i.e., O(S log n) bits of information or O(S) words) that can be stored in the local memory of a machine between the phases of local computation."

This informal claim, stating that the two types of bounds are somehow "equivalent" in some sense, is not clear, since it could be possible that only bounded information is communicated in each round, yet nevertheless the machines may accumulate the new information they learn in each round in their (unbounded size) local memories. This issue should be clarified, and the claim should be either justified or removed.
Later in the text, it is clarified that

> "In this thesis, S is also called a bound on the local memory of a single machine."

So perhaps it may be clearer if you just define it in this way from the beginning.

Conclusion

The thesis is suitable, in its current form, for fulfillment of the requirements for the degree of Doctor of Philosophy.

In my opinion the contributions are outstanding, and the thesis is deserving of an excellence prize.

Sincerely,

David Peleg