Low Diameter Graph Decompositions

Josias Thöny (thoenyj@student.ethz.ch)

Seminar of Distributed Computing
ETH Zürich
December 2003

1 Introduction

Graph decompositions are used in distributed computing as a tool for fast distributed algorithms. The goal is to decentralize computations (symmetry-breaking) and to exploit locality, which can be achieved by decomposing a graph into blocks of nearby vertices. More formally, a decomposition of a graph $G = (V, E)$ is a partition of the vertex set into disjoint subsets, called blocks. We're interested in decompositions into a small number of blocks, each having small diameter (maximum shortest path between any two vertices). A decomposition with a good trade-off between these two qualities can be used to improve the performance of several distributed algorithms, like MIS (maximal independent set), graph coloring, or distributed breadth first search.

2 Summary

The paper [1] presents some interesting results about graph decompositions. The authors show that every n-vertex graph has a decomposition into $O(\log n)$ blocks of diameter $O(\log n)$. This trade-off is proven to be optimal by constructing families of graphs where the trade-off cannot be improved. Furthermore, the paper shows two algorithms for decomposing graphs, a sequential one and a distributed randomized one. Both yield a decomposition of diameter $O(\log n)$ into $O(\log n)$ blocks. In the case of the sequential algorithm, these bounds are guaranteed, and the running time is polynomial. The randomized distributed algorithm achieves the logarithmic bounds with high probability, using a time complexity of $O(\log^2 n)$. Both algorithms create blocks which can be internally disconnected. This implies use in applications where the purpose of a decomposition is more symmetry-breaking than clustering.

The two presented algorithms use different definitions for the diameter of a block B in the graph G. In the case of the sequential algorithm, the so-called *strong diameter* is defined as the length of the maximum shortest path in B, using only vertices of B. The distributed algorithm uses the *weak diameter*, which is not bound to this restriction. That means we are allowed to shortcut through vertices not in B.

It may be desireable to have a bound on the strong diameter for the distributed algorithm also, but the construction of the decomposition as done by the algorithm doesn't imply such a bound.

3 Related Work

The idea of network decomposition was introduced in [2], where an efficient distributed algorithm was presented. However, the quality of the decomposition (in terms of both diameter and number of blocks) was not optimal. In [1], these results were clearly improved. Yet, there are several things about the distributed algorithm presented in [1] which are still not quite satisfying.

First, there is the before-mentioned issue with weak and strong diameter. One may be interested in finding a decomposition with bounded strong diameter.

Secondly, the distributed algorithm uses randomization. This might not be acceptable in some cases, since we cannot give a guarantee about the quality of the decomposition. For some problems, a deterministic algorithm will be a preferred solution.

Another aspect is about local neighborhood. The blocks as described in [1] use only a weak notation of neighborhood, and they can even be internally disconnected. There are applications where this approach is not suitable and a stronger concept of neighborhood is needed, which has been introduced later under the name of sparse neighborhood cover. A neighborhood cover is a set of overlapping clusters of a graph, such that for every vertex there exists a cluster which contains its neighborhood (=set of nearby vertices).

In [3] , further research on network decomposition and neighborhood covers has been done, and efficient deterministic distributed algorithms have been presented for both problems. The new algorithm for network decomposition works with the weak diameter, but it's claimed that a neighborhood cover algorithm can be used to construct a decomposition with good strong diameter.

But still, the results presented in [1] can be used for specific applications, with the advantage of the algorithm being simple and efficient.

References

[1] N. Linial and M. Saks: *Low Diameter Graph Decompositions*, Combinatorica 13 (4) (1993), pages 441-454

[2] B. Awerbuch, A. Goldberg, M. Luby, S. Plotkin: *Network Decomposition and Locality in Distributed Computation*, Proc. 30th IEEE Symp. on Foundations of Comp. Sci. (1989), pages 364-369

[3] B.Awerbuch, B.Berger, L.Cowen, D.Peleg: *Fast Distributed Network Decompositions and Covers*, Journal of Parallel and Distributed Computing, Volume 39, Number 2, December 1996, pages 105-114