COMBINATORICA
Akadémiai Kiadó – Springer-Verlag

COMBINATORICA 13 (4) (1993) 441–454

LOW DIAMETER GRAPH DECOMPOSITIONS*

NATHAN LINIAL** and MICHAEL SAKS†

Received October 6, 1990

A *decomposition* of a graph $G=(V,E)$ is a partition of the vertex set into subsets (called *blocks*). The *diameter* of a decomposition is the least d such that any two vertices belonging to the same connected component of a block are at distance $\le d$. In this paper we prove (nearly best possible) statements of the form: Any n-vertex graph has a decomposition into a small number of blocks each having small diameter. Such decompositions provide a tool for efficiently decentralizing distributed computations. In [4] it was shown that every graph has a decomposition into at most $s(n)$ blocks of diameter at most $s(n)$ for $s(n)=n^{O(\sqrt{\log\log n/\log n})}$. Using a technique of Awerbuch [3] and Awerbuch and Peleg [5], we improve this result by showing that every graph has a decomposition of diameter $O(\log n)$ into $O(\log n)$ blocks. In addition, we give a randomized distributed algorithm that produces such a decomposition and runs in time $O(\log^2 n)$. The construction can be parameterized to provide decompositions that trade-off between the number of blocks and the diameter. We show that this trade-off is nearly best possible for two families of graphs: the first consists of skeletons of certain triangulations of a simplex and the second consists of grid graphs with added diagonals. The proofs in both cases rely on basic results in combinatorial topology, Sperner's lemma for the first class and Tucker's lemma for the second.

1. Introduction

In this paper, we investigate a problem in algorithmic graph theory that originated in the theory of distributed computing. The systems we are concerned with can be modeled as graphs whose nodes correspond to processors and whose links correspond to communication channels between certain processors. One of the basic difficulties in designing algorithms for such systems is determining the extent to which the actions of the processors must be coordinated, and accomplishing this coordination as efficiently as possible. The most naive approach is to centralize the network operation by appointing one of the processors as a coordinator for the whole network and having all processes act under the direction of the coordinator. Centralization has several advantages; it often simplifies the problem considerably and facilitates the development of distributed algorithms based on known serial algorithms. On the other hand, rigid centralization often degrades system performance because of delays in communication between the coordinator and the other

AMS subject classification codes (1991): 05 C 12, 05 C 15, 05 C 35, 05 C 70, 05 C 85, 68 Q 22, 68 R 10

* A preliminary version of this paper appeared as "Decomposing Graphs into Regions of Small Diameter" in Proc. 2nd ACM–SIAM Symposium on Discrete Algorithms (1991) 321–330.

** This work was supported in part by NSF grant DMS87-03541 and by a grant from the Israel Academy of Science.

† This work was supported in part by NSF grant DMS87-03541 and CCR89-11388.