

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 $\leq d$. 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.

processors of the system. This problem is particularly significant in networks with large diameter and non-negligible message transmission time.

For these reasons, considerable amounts of research in distributed algorithms has focused on finding ways to decentralize distributed computation. For many problems it is possible to design algorithms in which each node acts with knowledge only of the activity of nearby nodes, and these "local" activities combine together to produce a global solution to the problem. The extent to which this is possible is referred to informally as the *locality* of the problem. Exploiting locality for specific problems leads to algorithms that are among the most novel and interesting in the area, for example, the beautiful symmetry breaking techniques of Cole and Vishkin ([8]) which have been used for graph coloring ([11]). Limitations on locality were addressed in [13].

This leads to the following general problem: find techniques for the design of distributed network algorithms that can be used to exploit locality. One class of techniques that has been proposed involves partitioning the network into regions of small diameter, coordinating action in each region through a local coordinator and combining the partial solutions together. This natural methodology has been considered by a number of authors; its earliest explicit statement known to us is in Awerbuch's [3] analysis of time and communication trade-offs required to achieve network synchronization. In [4] it is further used to improve the time complexity of distributed deterministic maximal independent set (MIS) algorithms, for graph coloring and distributed breadth first search. More recently, the approach has been applied to distributed routing [5] for a distributed all-pairs-shortest-distance algorithm and other similar problems [1].

The above discussion motivates the following graph definitions:

Let G be a graph. A subset W of vertices will be called a *block* of G . The *strong diameter* of a block W , $SD(W)$ is the maximum diameter of any connected component of the graph G_W induced on W . The *weak diameter* $WD(W)$ is the maximum distance in G between two vertices of W that belong to the same connected component of G_W . (The difference between strong and weak diameter is that when computing weak diameter we are allowed to shortcut through vertices not in W and thus $WD(W) \leq SD(W)$.) A partition Π of the vertex set of a graph G into λ disjoint blocks is called a λ -*decomposition* of G . The *strong diameter* $SD(\Pi)$ (*weak diameter* $WD(\Pi)$) of Π is the maximum strong diameter (weak diameter) of any of its blocks.

For a given graph we are interested in finding decompositions into a small number of (possibly disconnected) blocks each of small (strong or weak) diameter. This problem was introduced (with somewhat different terminology) in [4] as an attempt to solve a major outstanding problem in the theory of distributed algorithms: is there a deterministic algorithm for finding a maximal independent set in a distributed network that runs in polylog time? There are various partial results known for this tantalizing problem: a randomized distributed algorithm that runs in expected polylog time (Luby [14] and Alon *et al.* [2]), and a deterministic polylog time algorithm for bounded degree graphs (Goldberg *et al.* [11]). It was noted in [4] that given a λ -decomposition of the graph, an MIS can be constructed in a sequence of λ rounds each taking time $O(D \text{ polylog}(n))$ where D is the strong diameter, through the following iterative procedure: after i iterations there will be

an MIS for the subgraph of vertices in the first i blocks of the decomposition. During round $i+1$, the vertices in block $i+1$ that have no neighbors in the current independent set try to find an MIS among themselves. This can be done separately by each connected component of this subgraph, and the information about the connected component can be collected at one node in time $O(D)$ where D is the strong diameter of the decomposition. In fact, it is easy to see that D can be taken to be the weak diameter of the decomposition as well.

In [4], it was shown that for some $k = n^{O(\sqrt{\log \log n / \log n})}$, every graph has a k -decomposition with strong diameter at most $O(k)$. Furthermore, they gave a distributed algorithm running in time $O(k)$ for finding this decomposition. This enabled them to construct the fastest known deterministic distributed algorithm for the maximal independent set problem, which runs in time $n^{O(\sqrt{\log \log n / \log n})}$, (which is $o(n^\epsilon)$ for any positive ϵ , but bigger than any polylog(n)).

The obvious questions that come out of their work are:

1. What are the trade-offs between the number of blocks λ of a decomposition and its strong (or weak) diameter? In particular, what is the smallest $k(n)$ such that every graph on n vertices has a $k(n)$ -decomposition with weak (or strong) diameter $k(n)$?
2. What is the smallest $j(n)$ such that there is a deterministic distributed algorithm that runs on any n -vertex graph in polylog(n) time and produces a $j(n)$ -decomposition with weak diameter $j(n)$.

Of course, the functions $j(n)$ and $k(n)$ answering these questions satisfy $j(n) \geq k(n)$. What we'd like is that $j(n)$ is $O(\text{polylog}(n))$, in which case we could use the approach of [4] to get an MIS algorithm with the desired time complexity.

As it happens, it is not hard to show that $k(n) = O(\log n)$, using a simple but powerful constructive technique that was introduced by Awerbuch ([3]) and modified by Awerbuch and Peleg ([5]) to solve some graph covering theorems closely related to the decomposition problem. This technique provides a sequential way to build the decomposition one block at a time and enables us to give a deterministic sequential algorithm that, for any n -vertex graph G and any $p \in (0, 1)$ produces a λ -decomposition with $\lambda < \frac{\log n}{\log(1/(1-p))}$ and strong (and hence also weak) diameter at most $\frac{2 \log n}{\log(1/p)}$. For $\lambda \leq \log n$, this implies that the decomposition has diameter at most $2n^{1/\lambda} \log n$. In particular, the function $k(n)$ referred to in the first question above is $O(\log n)$.

Here and elsewhere, all logarithms are in base 2.

On the other hand, we can show that the trade-off given by this construction is nearly tight; precisely we show that for each $\lambda = O(\log n)$ there is a graph S with at most n vertices such that any λ -decomposition has weak diameter (and hence strong diameter) at least $\Omega(n^{1/\lambda})$. The graph S is simple: Choose m to be the greatest integer such that $\binom{m+\lambda}{\lambda} \leq n$. The vertex set of S consists of all vectors $\bar{x} = (x_0, \dots, x_\lambda)$ where $\sum x_i = m$ and all x_i are nonnegative integers. Vertices \bar{x}, \bar{y} are adjacent iff $\bar{x} - \bar{y} \in \{-1, 0, 1\}^{\lambda+1}$. (This graph contains as a spanning subgraph the one-skeleton of the standard triangulation of the λ -dimensional simplex). The

lower bound on the diameter of any λ -decomposition is deduced from a well known lemma in combinatorial topology due to Sperner. For λ substantially greater than $\log n$, it is easy to see that the upper bound is tight for expander graphs with the appropriate parameters. Another class of graphs where the tradeoff cannot be improved has vertices all integer vectors $\bar{x} = (x_1, \dots, x_\lambda)$ where \bar{x}, \bar{y} are adjacent iff their difference $\bar{x} - \bar{y}$ is in either $\{0, 1\}^\lambda$ or $\{0, -1\}^\lambda$. The proof depends this time on Tucker's lemma in combinatorial topology. While this paper was being refereed, we discovered that the relevant property of these graphs was proved previously by Gale [10] (with a different proof) in the analysis of certain combinatorial games.

Most interesting from the point of view of distributed computation, we give a randomized distributed algorithm that runs in expected time $O(n^{1/\lambda} \text{polylog}(n))$ and produces a λ -decomposition with weak diameter $O(n^{1/\lambda})$ and degree $\leq \lambda$. We note that we do not know how to make a similar guarantee on the strong diameter.

Our results fall short of the original goal; for this we would need to replace the randomized algorithm by a deterministic one. We believe that careful study of the randomized algorithm may lead to a deterministic one, or alternatively, suggest why no deterministic algorithm is possible. Furthermore, as discussed in [4] and [5], low diameter decompositions have a variety of potential applications and a fast randomized distributed algorithm is itself a potentially useful tool for replacing serialization by randomness in distributed computing. For instance, in [3], a graph decomposition is used as a data structure for providing a way to force synchrony in an asynchronous network. The construction given there is serial, and our randomized algorithm can be used to give a fast distributed construction. (Awerbuch [personal communication] has an alternative randomized algorithm with similar properties).

2. The existence of low diameter λ -decompositions

The first result of this section is:

Theorem 2.1. *Let p be a real number between 0 and 1, G an n -vertex graph and $\lambda = \frac{\log n}{\log(1/(1-p))}$. Then there is a λ -decomposition of G with strong diameter at most $\frac{2 \log n}{\log(1/p)}$.*

By analyzing the above theorem for the two ranges of λ depending on whether it is less than or greater than $\log n$ we obtain the following corollaries:

Corollary 2.1. *For $\lambda \leq \log n$, every n vertex graph has a λ -decomposition with strong diameter at most $2n^{1/\lambda} \log n$.*

Corollary 2.2. *For $\lambda \geq \log n$, every n vertex graph has a λ -decomposition with strong diameter at most $\frac{2 \log n}{1 + \log \lambda - \log \log n}$.*

In particular, when $p = 1/2$ we get $\lambda \leq \log n$ and strong diameter at most $2 \log n$.

The proof of theorem 2.1 is based on a technique of Awerbuch and is implicit in the work of Awerbuch and Peleg [5].

Proof. For two vertices x, y in a graph H let $d(x, y) = d_H(x, y)$ be their distance in H . For an integer r let $B_r(x)$ be the ball of radius r around x in H , i.e., $\{y \in V(H) : d_H(x, y) \leq r\}$. Whenever the need arises the graph H will be mentioned explicitly.

Fix p between 0 and 1 and call an integer r a *safe radius* for a vertex x if $p|B_r(x)| < |B_{r-1}(x)|$. Note that if $1, \dots, r$ are all unsafe for x , then necessarily $|B_j(x)| \geq (1/p)^j$ for all $1 \leq j \leq r$, and in particular $n \geq (1/p)^r$. In other words, every vertex x has a safe radius not exceeding $\frac{\log n}{\log(1/p)}$.

We construct a λ -decomposition one block at a time. The construction of V_1 is performed iteratively. Pick any vertex x_1 in $G_1 = G$ and let r_1 be the smallest safe radius for x_1 . Add all of the vertices in $B_{r_1-1}(x_1)$ to V_1 and define G_2 to be the graph obtained by deleting $B_{r_1}(x_1)$ from G_1 . Select x_2 in G_2 and let r_2 be the smallest safe radius for x_2 (in the graph G_2). Add the vertices of $B_{r_2-1}(x_2)$ (restricted to G_2) to V_1 and define G_3 to be $G_2 \setminus B_{r_2}(x_2)$. Continue this process constructing sequences $\{x_i\}$ of vertices $\{r_i\}$ of radii and $\{G_i\}$ of graphs and enlarging V_1 at each stage. The construction of V_1 is complete when G_i is empty.

The blocks V_i , for $i > 1$ are constructed inductively. Having constructed V_1, \dots, V_{i-1} define $G^i = G \setminus (V_1 \cup V_2 \cup \dots \cup V_{i-1})$ and apply the above process to G^i to obtain V_i . Continue until all vertices belong to some V_i .

The construction of the first (and each subsequent) block guarantees that its strong diameter is at most twice the largest radius of any of the selected balls. Since these radii are bounded by $\frac{\log n}{\log(1/p)}$ we obtain:

$$SD(\Pi) \leq 2 \frac{\log n}{\log(1/p)}.$$

Since the ratio of $|B_{r_i-1}(x_i)|$ to $|B_{r_i}|$ is always at least p for each selected ball, the fraction of vertices of G^i that are not assigned to V_i is at most $1-p$, and the number of vertices in G^{i+1} is at most $1-p$ times the number in G^i . Thus the number of blocks of the decomposition is bounded as follows:

$$\lambda \leq \frac{\log n}{\log(1/(1-p))}.$$

Note that the above proof provides a polynomial time sequential algorithm for constructing the decomposition.

3. Tightness of the existence theorem

Our existence theorem cannot be significantly improved for $\lambda > \log n$. A graph of chromatic number k and girth g cannot have a $((k-1)/2)$ -decomposition of diameter less than $\lceil \frac{g}{2} \rceil$, since in such a decomposition each block would be a forest and thus, by two coloring each block we would have a $k-1$ coloring of the graph. Since graphs of chromatic number k and girth $g = \Omega(\log n / \log k)$ exist for $k = O(\log n)$ (see [7]) this shows that Corollary 2.2 is essentially tight.

lower bound on the diameter of any λ -decomposition is deduced from a well known lemma in combinatorial topology due to Sperner. For λ substantially greater than $\log n$, it is easy to see that the upper bound is tight for expander graphs with the appropriate parameters. Another class of graphs where the tradeoff cannot be improved has vertices all integer vectors $\bar{x} = (x_1, \dots, x_\lambda)$ where \bar{x}, \bar{y} are adjacent iff their difference $\bar{x} - \bar{y}$ is in either $\{0, 1\}^\lambda$ or $\{0, -1\}^\lambda$. The proof depends this time on Tucker's lemma in combinatorial topology. While this paper was being refereed, we discovered that the relevant property of these graphs was proved previously by Gale [10] (with a different proof) in the analysis of certain combinatorial games.

Most interesting from the point of view of distributed computation, we give a randomized distributed algorithm that runs in expected time $O(n^{1/\lambda} \text{polylog}(n))$ and produces a λ -decomposition with weak diameter $O(n^{1/\lambda})$ and degree $\leq \lambda$. We note that we do not know how to make a similar guarantee on the strong diameter.

Our results fall short of the original goal; for this we would need to replace the randomized algorithm by a deterministic one. We believe that careful study of the randomized algorithm may lead to a deterministic one, or alternatively, suggest why no deterministic algorithm is possible. Furthermore, as discussed in [4] and [5], low diameter decompositions have a variety of potential applications and a fast randomized distributed algorithm is itself a potentially useful tool for replacing serialization by randomness in distributed computing. For instance, in [3], a graph decomposition is used as a data structure for providing a way to force synchrony in an asynchronous network. The construction given there is serial, and our randomized algorithm can be used to give a fast distributed construction. (Awerbuch [personal communication] has an alternative randomized algorithm with similar properties).

2. The existence of low diameter λ -decompositions

The first result of this section is:

Theorem 2.1. *Let p be a real number between 0 and 1, G an n -vertex graph and $\lambda = \frac{\log n}{\log(1/(1-p))}$. Then there is a λ -decomposition of G with strong diameter at most $\frac{2 \log n}{\log(1/p)}$.*

By analyzing the above theorem for the two ranges of λ depending on whether it is less than or greater than $\log n$ we obtain the following corollaries:

Corollary 2.1. *For $\lambda \leq \log n$, every n vertex graph has a λ -decomposition with strong diameter at most $2n^{1/\lambda} \log n$.*

Corollary 2.2. *For $\lambda \geq \log n$, every n vertex graph has a λ -decomposition with strong diameter at most $\frac{2 \log n}{1 + \log \lambda - \log \log n}$.*

In particular, when $p = 1/2$ we get $\lambda \leq \log n$ and strong diameter at most $2 \log n$.

The proof of theorem 2.1 is based on a technique of Awerbuch and is implicit in the work of Awerbuch and Peleg [5].

