
Local-Global Phenomena in Graphs

NATHAN LINIAL

Institute of Computer Science, Hebrew University, Jerusalem, Israel

Received 26 March 1993

For Paul Erdős on his 80th birthday

This is a survey of a number of recent papers dealing with graphs from a geometric perspective. The main theme of these studies is the relationship between graph properties that are local in nature, and global graph parameters. Connections with the theory of distributed computing are pointed out and many open problems are presented.

1. Introduction

How well can global properties of a graph be inferred from observations that are purely local? This general question gives rise to numerous interesting problems that we want to discuss here. Such a *local-global* approach is often taken in geometry, where it has a long and successful history, but a systematic study of graphs from this perspective has not begun until recently. Nevertheless, a number of older results in graph theory do fit very nicely into this framework, as we later point out. Most of the specific problems fall in two categories. In the first, local structural information on the graph is collected and then used to derive certain consequences for the graph as a whole. The other class of problems concerns *consistency* of local data. Namely, one asks to characterize those sets of local data that may come from some graphs.

As the reader will soon see, the local-global paradigm leads to many questions in which graphs are viewed as geometric objects, a point of view that we believe can greatly benefit graph theory. Besides the geometric connection, ties also exist with the theory of combinatorial algorithms. We suggest a specific test case for the heuristic notion that polynomial-time algorithms are capable of examining only local phenomena. In distributed computing, locality of computation is an already recognized and studied notion, and some connections with this discipline are pointed out as well.

2. Packing and covering with spheres and local-global averaging

Let $W \subseteq V(G)$ be a set of vertices in a graph G . If the vertices in W form a majority in every ball of radius between 1 and r in G , does this imply that W has a large cardinality?

. At the moment, we
, and we believe that
fixed $r \geq 3$. As to the
t $c_3 = 1/3$.

seven chromatic number,
§ 205–218.

488 pp.
e Texts in Mathematics

l graph theory. *Discrete*

Kugel, Wiss Z. Martin-

graphs. In: Bollobás, B.

162.

Lapok 48 436–452 (in

As an illustration, consider the following example with $r = 1$. In this graph, W is a clique of \sqrt{n} vertices. Each vertex in W has a set of $\sqrt{n} - 1$ neighbors not in W , each of which has degree 1. It is a routine matter to check that this graph satisfies the assumption for $r = 1$. It is also not hard to modify the construction for any fixed $\alpha < 1$ so that W occupies a fraction $\geq \alpha$ of any 1-neighborhood, while $|W| = O(\sqrt{n})$ (here α was $1/2$).

Let us introduce some notation. The ball[†] of radius k centered at x , denoted $B_k(x)$ consists of all vertices y whose distance from x does not exceed k , and its cardinality $|B_k(x)|$ is denoted $\beta_k(x)$. Our question is how small $|W|$ may be in terms of r and n , the order of G .

If we represent W by its characteristic function, we are led to consider a more general problem. Namely, let f be a nonnegative function defined on the vertices of an n -vertex graph G . Suppose that we have a lower bound on the average of f on every ball in G of radius between 1 and r . What can we conclude for the overall average of f ?

This subject has been recently taken up by Linial, Peleg, Rabinovich and Saks [23] who show the following.

Theorem 2.1. (Local Averages) *Let f be a nonnegative function defined on the vertices of an n -vertex graph G . Suppose that the average of f over every ball of radius $r \geq t \geq 1$ in G is at least μ . Then, the average of f over all of V is at least $\mu \cdot n^{-O(1/\log r)}$. The bound is tight.*

Consequently, if we let r be n^c for some positive constant c , local averages do reflect the true global behavior of f . Examples are given in [23] showing that smaller r 's will not do. It is natural to ask at this point what happens if we only know a lower bound for the average of f over balls of radius r (and not for every $r \geq t \geq 1$). Examples are given showing that only very weak conclusions can be drawn about the overall average of f , however big r may be. Namely, it may be that the average of f is only $O(n^{-1/3})$. It is also worthwhile noting that the conclusions of the theorems remain unchanged even if we make the assumption only for balls whose radius $r \geq t \geq 1$ is a power of 2.

The result for local averages is proved as a consequence of tight theorems about sphere packing and about covering by spheres in general graphs. Either 0-1 or fractional packing and covering results will do for this purpose.

Theorem 2.2. (Covering by Spheres) *For integers $n > r$, the vertices of an n -vertex graph can be covered by a collection of balls with radii in the range $[1, \dots, r]$, that cover no vertex more than $n^{O(1/\log r)}$ times. The bound is tight.*

Theorem 2.3. (Sphere Packing) *In any n -vertex graph, there is a collection of disjoint balls whose radii are in the range $[1, \dots, r]$, which together cover at least $n^{1-O(1/\log r)}$ vertices. The bound is tight.*

[†] The words ball and sphere are used interchangeably here.

It would be efficiency of these results happens if the low-dimension

2.1. Connect

There is an maximal function Metanya Ben-

Briefly, the $x \in \mathbf{R}^d$, the d - be the average The function f over the years findings is that

Our proof for analysis to combine technique involving argument under of the form

where $C_{d,p}$ grows in this method, concerning maximum that these methods It would also be graphs.

The theory of distributed communication networks or processors residing in neighboring vertices. Let us restrict our attention among this class of graphs in t time units and within distance at most t . questions of the local to symmetry breaking processors, then all

It would be very interesting to understand how various properties of a graph affect the efficiency of sphere packing and of covering by spheres. Also, it is not hard to extend these results to general finite metric spaces. We still do not know, for example, what happens if the metric space is embedded in a d -dimensional Euclidean space or other low-dimensional normed space. These questions lead us to our next subject.

2.1. Connections with the theory of maximal functions

There is an appealing connection between this class of problems and the theory of maximal functions in analysis (e.g. [32]). This observation came up in discussions with Metanya Ben-Artzi.

Briefly, the connection is this: again let $B_r(x)$ denote the ball of radius r centered at $x \in \mathbf{R}^d$, the d -dimensional Euclidean space. Let f be a real function on \mathbf{R}^d , and let $a_r(x)$ be the average of f over $B_r(x)$. Define $f^*(x)$ as the supremum of $a_r(x)$ over all $r > 0$. The function f^* is called the *maximal function* of f . Numerous results have been derived over the years concerning maximal functions. Informally speaking, among the most basic findings is that ' f^* is not much larger than f '.

Our proof for Theorem 2.1 shows a significant similarity with the methods used in analysis to compare the p -norms of f^* and f . Specifically, the most traditional proof technique involves some geometric covering arguments (Vitali's Lemma), and a similar argument underlies some of our proofs as well. In analysis, such arguments lead to results of the form

$$\|f^*\|_p < C_{d,p} \|f\|_p$$

where $C_{d,p}$ grows exponentially with the dimension d . This bad dependency is unavoidable in this method, since the bounds in Vitali's lemma do grow this way. More modern results concerning maximal functions (e.g. [33]) manage to bypass this difficulty. It is conceivable that these methods may help settle our questions on low-dimensional finite metric spaces. It would also be interesting to see if similar ideas can be developed for other classes of graphs.

3. Locality in distributed systems

The theory of distributed computing concerns a set of processors connected through a communication network. The network is depicted as a graph in whose vertices computers or processors reside. Communication takes place as messages are exchanged between neighboring vertices. The processors' goal is to perform some computational task together. Let us restrict our attention to deterministic and synchronized networks – the simplest among this class of computational models. In such an environment it is easy to see that in t time units a processor can only learn about the situation at processors that are within distance at most t from itself in the graph. This observation gives rise to numerous questions of the local-global type. In studying such questions, some care has to be given to symmetry breaking. If processors 'have no identity' and cannot be told apart by other processors, then almost nothing interesting can be done. We do not elaborate on this,

In this graph, W is a
sors not in W , each of
atisfies the assumption
fixed $\alpha < 1$ so that W
) (here α was $1/2$).

d at x , denoted $B_k(x)$
 k , and its cardinality
terms of r and n , the

onsider a more general
vertices of an n -vertex
on every ball in G of
age of f ?

ich and Saks [23] who

ined on the vertices of
of radius $r \geq t \geq 1$ in
 $-O(1/\log r)$. The bound is

al averages do reflect
g that smaller r 's will
know a lower bound
 $t \geq 1$). Examples are
ut the overall average
f f is only $O(n^{-1/3})$. It
ain unchanged even if
power of 2.

heorems about sphere
l or fractional packing

s of an n -vertex graph
, that cover no vertex

lection of disjoint balls
 $-O(1/\log r)$ vertices. The

but rather say that the common practice in this area is to assume that processors are equipped with individual (distinct) ID-numbers, and so symmetry causes no problems.

3.1. Low-diameter decompositions of graphs

Perhaps the most fundamental difficulty in distributed processing, as compared with more traditional computational models, is the absence of central control. It is very difficult to have many processors perform in concert when there is no conductor around. Indeed, much research effort in distributed processing concerns efficient and reliable methods for electing a leader. We will not pursue this fascinating subject, and only point out some of the shortcomings of this approach. It creates a communication bottleneck around the elected leader. It is also very sensitive to failures, or latency of the leader and its neighbors. Moreover, if the graph underlying the communication network has a large diameter, this method is also very wasteful in terms of communication.

In view of the difficulties involved with such a 'central government' the next thing to try is a set of cooperative 'local authorities'. Namely, in the previous section we were covering vertices by balls; now we consider *decomposing* the vertices, subject to a certain upper bound on the diameter of each part. Let us introduce some notation: if Π is a decomposition of the vertices of graph G into subsets, $V(G) = \bigcup S_i$. The *diameter* of this decomposition is defined as the maximum over all $diam(S_i)$.

Remark 3.1. In defining the diameter, we may consider the graphs induced by the parts, and compute distances within these graphs. Alternatively, we may consider distances as inherited from the whole graph. Our statements, slightly modified, hold for either definition.

The graph *induced* by Π has one vertex per S_i , with vertices i, j adjacent iff there is a vertex in S_i and one in S_j that are adjacent in G . The goal is to find partitions Π with small diameter and favorable properties for the induced graph. Linial and Saks [25] show (see also [6, 7]):

Theorem 3.2. *An n -vertex graph has a decomposition of diameter r , where the induced graph has chromatic number $\leq \chi$, if both*

$$\chi = \Omega\left(\frac{\log n}{\log r}\right) \text{ and } r = \Omega\left(\frac{\log n}{\log \chi}\right)$$

hold. Examples exist showing these bounds are tight. A randomized distributed algorithm of $\log^{O(1)} n$ run time is provided to obtain such decompositions.

We briefly discuss some extreme examples for Theorem 3.2. It is easily seen that there are two interesting ranges to this theorem:

$$r \geq \frac{\log n}{\log \log n} \geq \chi.$$

In this r:

The know
Euclidean
dimension:

where the

Here trees

Remark 3.
if every ba
factor, the

More or
So far v
decomposi
the analog
The follow
Let $\Pi : V(G)$
 x , let $\gamma(x)$ b

Problem 3.4
sition Π of

Possibly,
Theorem 3.2

3.2. Applic

Low diamet
computing.
(MIS) probl
the search fo
of course, a
the MIS and
algorithm sc
this question
Karp and W
[1, 27]). In fa
randomizatio

that processors are
solves no problems.

In this range, the tradeoff between r and χ is given by:

$$\chi = \Omega\left(\frac{\log n}{\log r}\right).$$

The known extreme examples in this range are graphs corresponding to triangulations of Euclidean spaces. For example, the graph whose vertices are all lattice points in $\log n / \log r$ dimensions, with adjacency between \mathbf{x}, \mathbf{y} iff $\|\mathbf{x} - \mathbf{y}\|_\infty = 1$.

$$\chi \geq \frac{\log n}{\log \log n} \geq r,$$

where the condition is

$$r = \Omega\left(\frac{\log n}{\log \chi}\right).$$

Here *trees* and *expander graphs* provide extreme examples.

Remark 3.3. Notice that radius $\log n$ along with $\chi = O(\log n)$ are possible. Consequently, if every ball or radius $\log n$ in G is k -colorable, $\chi(G) = O(k \log n)$. So, up to a logarithmic factor, the coloring number can be inferred from radius $\log n$ views of G .

More on coloring from the local-global perspective will be said later.

So far we have considered only the chromatic number of the graph induced by a decomposition. Other properties of this graph are of interest as well. Let us point out the analogy between these questions and notions from dimension theory in topology [19]. The following question is inspired by the notion of covering dimension of metric spaces. Let $\Pi : V(G) = \bigcup S_i$ be, again, a decomposition of the vertices of graph G . For a vertex x , let $\gamma(x)$ be the number of S_i in which x has a neighbor. $\Delta(\Pi)$ is defined as $\max_x(\gamma(x))$.

Problem 3.4. What is the least $D = D(r, n)$, such that any n -vertex graph has a decomposition Π of diameter $\leq r$ with $\Delta(\Pi) \leq D$?

Possibly, the tradeoff between D, r and n is the same as the one for χ, r and n in Theorem 3.2.

3.2. Applications of low-diameter decompositions

Low diameter graph decompositions have found numerous applications in distributed computing. We briefly sketch some of these. We begin with the Maximal Independent Set (MIS) problem. (We mean inclusion-maximal. This problem is not to be confused with the search for an independent set of largest cardinality, which is NP-complete.) There is, of course, a most simple sequential algorithm, which at each step adds a new vertex to the MIS and eliminates all its neighbors from the graph. While such a naive sequential algorithm solves the problem in optimal time, finding efficient *parallel* algorithms for this question is not nearly as obvious. An efficient parallel algorithm was first found by Karp and Wigderson [20] with numerous improvements and ramifications by others (e.g. [1, 27]). In fact, Luby's algorithm [27] works also in the distributed model, but it does use randomization, however. One of the tantalizing questions that remain in this area is:

