Combinatorics, Probability and Computing (1993) 2, 491–503

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?

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 $[1, \ldots, r]$, 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 $[1, \ldots, r]$, 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 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,

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, $V(G) = \dot{\bigcup} S_i$. 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 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 $\|\mathbf{x} - \mathbf{y}\|_\infty = \mathbf{1}$.

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

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 $\Pi : V(G) = \dot{\bigcup} S_i$ 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:

Problem 3.5. Is there a deterministic, distributed, polylog-time algorithm to find a maximal independent set (MIS) in a graph?

It has been observed in [4] that low-diameter decompositions with a low-chromatic decomposition graph may help provide such an algorithm. Assuming such (an already colored) decomposition is available, we construct, in parallel, an MIS within each part colored 1. Since each part has diameter at most r, an MIS for it can be constructed by an elected leader, where both election and construction take time $O(r)$. Also, there are no edges between different parts of color 1, so these activities in different parts can be performed in parallel without affecting each other. Vertices selected so far for the MIS eliminate their neighbors (also in other parts), and we move on to parts colored 2 etc. Using the terminology of Theorem 3.2, a time bound of $O(r\chi)$ can be achieved, which by proper choice of parameters may be made $O(\log^2 n)$. The difficulty is, of course, that this argument assumes a partition to be already available. Currently, however, only *randomized* distributed algorithms are known that find such decompositions in polylogarithmic time (Theorem 3.2 and [3]). Problem 3.5 thus remains open.

Another problem for which low-diameter decompositions help is *distributed job scheduling* [5]. In this problem, processors try to efficiently share their workloads. Initially, each processor is assigned a number of (unit-cost) jobs to perform. In each step, a processor can perform one of its assigned jobs, as well as send some of its assigned jobs to neighbors. It can also communicate messages to its neighbors. A processor knows only its own history and the contents of the messages it receives. An algorithm is sought where the completion time is early as possible. Moreover, the following strong ('competitive') criterion is applied: the time for completion should compare favorably with the best that can be achieved by an optimal central controller having a complete view of the situation at all times (and not just local views at a certain processor). Using low-diameter decompositions, [5] manages to guarantee a completion time that is only $O(\log n)$ longer than can be attained by a knowledgable central controller. This result is shown to be almost optimal for certain families of graphs in [2]. For further applications see [6, 7].

4. Distributed coloring and related problems

The systematic study of locality in distributed processing was begun in [22]. Our first result concerned the time required to 3-color an n-cycle of processors. A clever algorithm by Cole and Vishkin [11] does this in time $O(\log^* n)$. (Recall that $\log^* n$ is the number of times one has to iterate the log function to come down from n to 1).

The first result in [22] says that this algorithm is optimal.

Theorem 4.1. *A distributed algorithm that properly colors the n-cycle with only 3 colors requires time $\Omega(\log^* n)$. The bound is tight.*

It was later shown by Naor [30] that the same statement also holds for *randomized* algorithms.

Assuming n is even, how long does it take to 2-color an n-cycle? A huge difference shows up, in comparison with the time complexity of 3-coloring.

Theorem 4.2. *A distributed algorithm that properly colors the n-cycle (n even) with only 2 colors requires time $\Omega(n)$. The bound is tight.*

This example captures a big difference in locality of 2 and 3-coloring of cycles. Of course, a 2-coloring requires perfect *global* coordination, which results in an excessive time complexity.

Other results from [22] are:

Theorem 4.3. *Let T be the d-regular tree of radius r. Any algorithm that properly colors T and runs for time $< 2r/3$ requires at least $\Omega(\sqrt{d})$ colors.*

This result can probably be improved to $\Omega(d/\log d)$. An intriguing open question in this area is:

Problem 4.4. Consider distributed algorithms that properly color n-vertex graphs and take time $\log^{O(1)} n$. What can be said about the least number of colors required by such an algorithm? Specifically, is it possible that $\Delta + 1$ colors suffice, where Δ is the largest vertex degree of the graph?

This question is closely related to Problem 3.5. Some partial results have been provided in [22].

Theorem 4.5. *An $O(\log^* n)$-time algorithm exists to color any n-vertex graph G with $O(\Delta^2)$ colors, where $\Delta = \Delta(G)$ is the largest vertex degree in G.*

See also [34] for some recent progress in this area. Naor and Stockmeyer [31] have recently investigated the limits of what can be computed with a *constant* diameter of locality.

Another related problem is that of finding *happy partitions*. A partition of the vertex set of a graph $V = A \dot{\cup} B$ is called happy if every $x \in A$ has most of its neighbors in B and vice versa. That such partitions always exist is easy to show, and a sequential algorithm to construct such partitions is easy to find. The distributed time complexity of this is still unknown: Linial and Saks conjecture (unpublished) as follows.

Conjecture 4.6. *There are n-vertex graphs where a distributed algorithm to find a happy partition requires time $\Omega(\sqrt{n})$.*

5. Coloring

The chromatic number of a graph is a good example of a global parameter where the behavior of small induced subgraphs seems to be a weak indicator of global properties.

(But notice Remark 3.3.) Up to this point 'local' has always been taken in the sense of distance. It is also interesting to examine assumptions about the behavior of (cardinality) small sets of vertices. In this section we consider sets that are small in either diameter or cardinality. One easy consequence of Theorem 3.2 is the following.

Theorem 5.1. *If the subgraph spanned by every k vertices in G is 2-colorable, $\chi(G) = O(n^{O(1/k)})$. This bound is tight. Moreover, it is possible to find a proper coloring with this number of colors in polynomial time.*

Dealing with 2-colorabilty is usually much easier than with any larger coloring number. Is there, perhaps, a similar result for graphs that are, say, locally 3-colorable? To simplify our notation, we will only consider 3-colorability, leaving out the obvious extension to more colors.

Problem 5.2. Let $\chi(n, k)$ be the largest chromatic number of an n-vertex graph G if the subgraph spanned by every k vertices in G is 3-colorable. Determine the behavior of $\chi(n, k)$.

It is not hard to see that

$$n^{1/2+o(1)} \geq \chi(n, k) \geq n^{7/20+o(1)},$$

where the $o(1)$ terms tend to zero as k grows. The upper bound follows, *e.g.*, from Wigderson's argument [35] mentioned below. The lower bound combines an argument from [24] with a lower bound due to Gallai [15] on the least number of edges in minimally non-3-colorable graphs. Note the difference compared with locally 2-colorable graphs (Theorem 5.1).

Besides the interest in the problem *per se*, it is related to approximating chromatic numbers in polynomial time. That it is NP-hard to determine the chromatic number has been known for a long time [16]. How well this quantity may be approximated is still unknown, although considerable progress has been made. An early positive result on approximating chromatic numbers is a polynomial-time algorithm by Wigderson [35], which colors any n-vertex 3-colorable graph with $O(\sqrt{n})$-colors. Here is the argument: as long as you can find a vertex x of degree $\geq \sqrt{n}$, allot two fresh colors for the neighbors of x and discard them (they are two-colorable, since $\chi(G) = 3$). When the remaining graph has all degrees $< \sqrt{n}$, it can be $\sqrt{n}$-colored by a greedy algorithm. Altogether, only $O(\sqrt{n})$ colors are utilized.

Observe that the algorithm actually applies not only to 3-colorable graphs, but, in fact, to every graph in which the neighborhood of every vertex is 2-colorable. Now, bounds on Ramsey numbers naturally fit into the local-global framework. For example, the fact that

$$R(3, k) = k^{2-o(1)}$$

(see [17] for the sharpest known bounds) answers the following question: given that the neighbors of any vertex in G form an anti-clique, what is the best lower bound on the largest anti-clique in G (answer: $n^{1/2-o(1)}$). In particular, triangle-free graphs exist of

chromatic number $\Omega(n^{1/2-o(1)})$. But in a triangle-free graph, the neighborhood of every vertex is, in fact an independent set, so under the more general assumptions, Wigderson's algorithm is in fact optimal.

These arguments were further improved by A. Blum [9], who showed how to color a 3-colorable graph with $n^{3/8}$ colors in polynomial time. Interestingly, Blum's algorithm (which is much more involved than that of [35]) also exploits only local (neighborhoods of radius 2) properties of 3-colorable graphs. This leads us to ask some questions to capture the heuristic claim that *polynomial-time graph-coloring algorithms can only check local properties*. We first observe that the answer to Problem 5.2 yields a completely trivial algorithm to tell 3-chromatic graphs from those not colorable in n^c colors. We expect this algorithm to be better than Blum's in this respect.

Conjecture 5.3. *Let G be an n-vertex graph in which every induced subgraph of order k is 3-colorable. Then, $\chi(G) < n^{\gamma+o(1)}$ for some $3/8 > \gamma \geq 7/20$, and where the $o(1)$ term tends to zero with $k \to \infty$.*

An exhaustive algorithm running in time $O(n^k)$ can obviously test this condition. It is an interesting possibility that this procedure may be transformed into an algorithm that actually provides a $n^{\gamma+o(1)}$ coloring.

If this conjecture fails, it may be possible to save it by adding an assumption such as that the neighborhood of every vertex is 2-colorable, a condition that is again polynomial-time verifiable.

A more daring conjecture is:

Conjecture 5.4. *If $P \neq NP$, then no polynomial time algorithm can color every 3-chromatic n-vertex graph with fewer than n^θ colors for some $\theta > 0$.*

There have been many new and exciting results on the difficulty of approximating NP-hard problems. The first step in establishing such a result for coloring has been taken by Lund and Yanakakis [28], who establish a separation between coloring numbers n^{c_1} and n^{c_2} for some fixed $1 > c_1 > c_2 > 0$. A simpler proof has been provided recently by Khanna, Linial and Safra [21], who also show that it is NP-hard to 5-color 3-colorable graphs. All this is, obviously, still a far cry from Conjecture 5.4, but some progress in this direction is likely to occur in the foreseeable future.

6. Sizes of neighborhoods

Perhaps the most obvious 'local' information about a graph is the degree sequence, classically characterized by Erdős and Gallai [12]. Briefly, $d_1 \geq \ldots d_n \geq 0$ is such a sequence iff (i) $\sum d_i$ is even and (ii) for all $1 \leq k \leq n$ it is the case that $\sum_1^k d_j \leq k(k-1) + \sum_{j>k} \min\{k, d_j\}$. The necessity of these conditions is easy to establish and the thrust of the theorem is that they are also sufficient. Pursuing our local-global approach we ask: what else can be said about the possible *rate of growth* of (balls in) a graph?

Recall that $\beta_k(x)$ is the number of vertices y whose distance from x does not exceed

k. In a connected n-vertex graph, one obtains n integer sequences, one for each vertex, $1 = \beta_0(x) \leq \beta_1(x) \ldots \leq \beta_n(x) = n$. Following Erdős–Gallai's result, it is appealing to ask:

Problem 6.1. Characterize those sets of n integer sequences of the type

$$1 = \beta_0(x) \leq \beta_1(x) \ldots \leq \beta_n(x) = n$$

that are obtained from connected graphs.

This question, in full generality, is presently too difficult, and at this time one should settle for less. Here are some illustrative special cases of this problem:

— Is it possible to characterize sets of n pairs $\beta_1(x) \leq \beta_2(x)$ that come from graphs? One possible approach would be to get sufficient information on squares of graphs and then resort to Erdős–Gallai. Note, however, that Motwani and Sudan [29] have shown that it is NP-complete to decide whether a given graph is a square.

— In the context of the previous question, it is not hard to derive some necessary conditions, *e.g.*, that

$$\sum_x (\beta_2(x) - 1) \leq \sum (\beta_1(x) - 1)^2,$$

with equality iff $girth(G) \geq 5$. This inequality suggests that there might exist some comparison theorems between norms of the various vectors $\bar{\beta}_i = (\beta_i(x) | x \in V)$.

— Obviously, for any fixed x, the sequence $1 = \beta_0(x) \leq \beta_1(x) \ldots \leq \beta_n(x) = n$ is unrestricted. It may be possible to characterize pairs of such sequences, one for vertex x and one for y. Such an analysis could start by considering for any i, j the number of vertices z that are at distance i from x and j from y.

— For which parameters is it possible that all x satisfy $\beta_1(x) = d + 1$, while for every $i \leq c \log n$ it is the case that $\beta_{i+1}(x) \geq (1 + \delta)\beta_i(x)$? This question is clearly related to the existence of constant-degree expanders. Methods developed in that area may prove helpful in studying growth rates of graphs in general.

— What is the largest girth of a d-regular n-vertex graph? Specifically, is it

$$(2 - o(1)) \log n / \log(d - 1)?$$

This is also an instance of the general problem. We conjecture the answer to be negative. The best current lower bound [26, 8] gives 4/3 instead of the 2.

A problem related to the last item in this list concerns the ratio between girth and diameter. Consider the distance between two vertices that are antipodal in a shortest cycle. This consideration shows that $2 \cdot diameter(G) \geq girth(G)$. Equality holds for even cycles, but what if all degrees are ≥ 3? Examples are known with $girth(G) = 2 \cdot diameter(G)$, where the numbers are small, *e.g.*, the points-lines graph of a projective plane. We are not aware of similar constructions with large girth, so we ask:

Problem 6.2. Consider graphs G with all degrees ≥ 3. What are possible values for the pair $(girth(G), diameter(G))$? In particular, can their ratio be kept as close to 2 as we wish?

See [14] for a related classical work.

Together with S. Hoori [18], we have recently obtained some results concerning the existence of a 'center of mass' in both graphs and sets in Euclidean spaces. Namely, we are looking for a vertex x where we can establish a tight lower bound on the numbers $\beta_k(x)$ $(k = 1, 2, \ldots)$.

7. Cliques

A number of people have investigated how well the clique number of a graph can be inferred from local behavior. The earliest work we are aware of is by Erdős and Rogers [13]. Recent work on the subject can be found in [10, 24] and the references therein. The main question studied here can be stated in terms of computing, or estimating, the quantities related to the following arrow relation. Say that a graph G has property (p, q) if every set of p vertices in G contains a q-clique. We say that $(p, q) \to (f, n)$, if every G of order $\geq p$ having the (p, q)-property must satisfy (f, n) as well. The question then is, for given p, q, n estimate the least $f = f(p, q, n)$ for which $(p, q) \to (f, n)$. The exact determination of f includes, as a special case, the exact evaluation of Ramsey numbers, so it is more realistic to ask for estimates, or to settle for special cases. We use both the arrow notation and the function f to describe the results.

Bollobas and Hind [10] concentrate on the case of large p and small q, n. Among their results are:

Theorem 7.1. *For $s > r \geq 3$,*

$$(n, r) \to (cn^{s-r+1}, s)$$

for some constant $1 > c > 0$. Also, for $r \geq 3$ and n large enough

$$(n, r) \not\to (n^{1+r/(r^2-2)-\epsilon}, r + 1).$$

Linial and Rabinovich [24] consider fixed p, q and n tending to infinity. Their main result breaks down into three cases, roughly according to whether p/q is smaller than, equal to or bigger than 2.

Theorem 7.2.

— *For $p \leq 2q - 2$ and all n,*

$$f(p, q, n) = n + p - q.$$

— *For $p = 2q - 1$ and all $n \geq p$,*

$$n^{1+2/(q-3)+o(1)} \geq f(p, q, n) \geq n^{1+1/(8q-5)}.$$

— *For all $n \geq p \geq q$,*

$$(p, q) \to (R(r, n) + p - 1, n),$$

where $r = \lceil \frac{p}{q-1} \rceil$, $R(r, n)$ is the Ramsey number and c is an absolute constant. On the other hand,

$$(p, q) \not\to (n^{(T-1)/(p-2)+o(1)}, n),$$

where T is a Turan number: the least number of edges in a p-vertex graph without a q-anticlique.

All $o(1)$ terms are for fixed p, q and growing n.

References

[1] Alon, N., Babai, L. and Itai, A. (1986) A fast and simple randomized algorithm for the maximal independent set problem. *J. of Algorithms* **7** 567–583.

[2] Alon, N., Kalai, G., Ricklin, M. and Stockmeyer, L. (1992) Lower bounds on the competitive ratio for mobile user tracking and distributed job scheduling. *FOCS* **33** 334–343.

[3] Awerbuch, B., Berger, B., Cowen, L. and Peleg, D. (1992) Fast distributed network decomposition. *PODC* **11** 169–178.

[4] Awerbuch, B., Goldberg, A. V., Luby, M. and Plotkin, S. A. (1989) Network decomposition and locality in distributed computation. *FOCS* **30** 364–369.

[5] Awerbuch, B., Kutten, S. and Peleg, D. (1992) Competitive distributed job scheduling. *STOC* **24** 571–580.

[6] Awerbuch, B. and Peleg, D. (1990) Sparse partitions. *FOCS* **31** 503–513.

[7] Awerbuch, B. and Peleg, D. (1990) Network synchronization with polylogarithmic overhead. *FOCS* **31** 514–522.

[8] Biggs, N. L. and Boshier, A. G. (1990) Note on the girth of Ramanujan graphs. *J. Combin. Th.* ser. B **49** 190–194.

[9] Blum, A. (1990) Some tools for approximate 3-coloring. *FOCS* **31** 554–562.

[10] Bollobás, B. and Hind, H. R. (1991) Graphs without large triangle free subgraphs. *Discrete Math.* **87** 119–131.

[11] Cole, R. and Vishkin, U. (1986) Deterministic coin tossing and accelerating cascades: micro and macro techniques for designing parallel algorithms. *STOC* **18** 206–219.

[12] Erdös, P. and Gallai, T. (1960) Graphen mit Punkten vorgeschriebenen Grades. *Mat. Lapok* **11** 264–274.

[13] Erdös, P. and Rogers, C. A. (1962) The construction of certain graphs. *Canad. J. Math.* **14** 702–707.

[14] Feit, W. and Higman, G. (1964) The nonexistence of certain generalized polygons. *J. Algebra* **1** 114–131.

[15] Gallai, T. (1963) Kritische Graphen I. *Publ. Math. Inst. Hungar. Acad. Sci.* **8** 165–192.

[16] Garey, M. R. and Johnson, D. S. (1979) *Computers and Intractability: A Guide to NP-completeness*, W. H. Freeman.

[17] Graham, R. L., Rothschild, B. L. and Spencer, J. (1980) *Ramsey Theory*, Wiley, New York.

[18] Hoori, S. and Linial, N. (March 1993) Work in progress.

[19] Hurewicz, W. and Wallman, H. (1948) *Dimension Theory*, Princeton University Press.

[20] Karp, R. and Wigderson, A. (1985) A fast parallel algorithm for the maximal independent set problem. *J. ACM* **32** 762-773.

[21] Khanna, S., Linial, N. and Safra, S. (to appear) On the hardness of approximating the chromatic number. *ISTCS*.

[22] Linial, N. (1992) Locality in distributed graph algorithms. *SIAM J. Comp.* **21** 193–201. (Preliminary version: Linial, N, (1987) Distributive graph algorithms – global solutions from local data. *FOCS* 331–335.

[23] Linial, N., Peleg, D., Rabinovich, Yu. and Saks, M. (to appear) Sphere packing and local majorities in graphs. *ISTCS*.

[24] Linial, N. and Rabinovich, Yu. (in press) Local and global clique numbers. *J. Combin. Th.* ser. B.

[25] Linial, N. and Saks, M. (to appear) Low diameter graph decompositions. *Combinatorica*. (Preliminary version (1991) published in *SODA* **2** 320–330.)

[26] Lubotsky, A.

[27] Luby, M. (19 *J. Comp.* **15**

[28] Lund, C. and manuscript.

[29] Motwani, R.

[30] Naor, M. (1 *SIAM J. Dis*

[31] Naor, M. an

[32] Stein, E. M. (Press.

[33] Stein, E. M. Rubio de Fra Studies **111**,

[34] Szegedy, M.

[35] Wigderson, A 325–329.

[26] Lubotsky, A., Phillips, R. and Sarnak, P. (1988) Ramanujan graphs. *Combinatorica 8* 261–278.
[27] Luby, M. (1986) A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comp.* **15** 1036–1053.
[28] Lund, C. and Yanakakis, M. (1992) On the Hardness of Approximating Minimization Problems, manuscript.
[29] Motwani, R. and Sudan, M. (1991) Computing roots of graphs is hard, manuscript, Stanford.
[30] Naor, M. (1991) A lower bound on probabilistic algorithms for distributive ring coloring. *SIAM J. Disc. Math.* **4** 409–412.
[31] Naor, M. and Stockmeyer, L. (1993, to appear) What can be Computed Locally? *STOC* **25**.
[32] Stein, E. M. (1970) Topics in Harmonic Analysis. *Annals of Math. Study* **63**, Princeton University Press.
[33] Stein, E. M. (1985) Three variations on the theme of maximal functions. In: Peral, I. and Rubio de Francia, J.-L. (eds.) *Recent Progress in Fourier Analysis*, North-Holland Mathematics Studies **111**, North-Holland, 229–244.
[34] Szegedy, M. and Vishwanatan, S. (1993, to appear) Locality-based graph coloring. *STOC* **25**.
[35] Wigderson, A. (1983) Improving the performance for approximate graph coloring. *J. ACM* **30** 325–329.