The Intrinsic Dimensionality of Graphs

Robert Krauthgamer*
International Computer Science Institute and
Computer Science Division, U.C. Berkeley
Berkeley, CA 94720

James R. Lee†
Computer Science Division, U.C. Berkeley
Berkeley, CA 94720

ABSTRACT

We resolve the following conjecture raised by Levin together with Linial, London, and Rabinovich [16]. Let $\mathbb{Z}_\infty^d$ be the infinite graph whose vertex set is $\mathbb{Z}^d$ and which has an edge (u,v) whenever $||u-v||_\infty = 1$. Let $\dim(G)$ be the smallest d such that G occurs as a (not necessarily induced) subgraph of $\mathbb{Z}_\infty^d$. The *growth rate* of G, denoted ρ_G, is the minimum ρ such that every ball of radius $r > 1$ in G contains at most r^ρ vertices. By simple volume arguments, $\dim(G) = \Omega(\rho_G)$. Levin conjectured that this lower bound is tight, i.e., that $\dim(G) = O(\rho_G)$ for every graph G.

Previously, it was not known whether $\dim(G)$ could be upper bounded by *any function* of ρ_G, even in the special case of trees. We show that a weaker form of Levin's conjecture holds by proving that, for every graph G, $\dim(G) = O(\rho_G \log \rho_G)$. We disprove, however, the specific bound of the conjecture and show that our upper bound is tight by exhibiting graphs for which $\dim(G) = \Omega(\rho_G \log \rho_G)$. For families of graphs which exclude a fixed minor, we salvage the strong form, showing that $\dim(G) = O(\rho_G)$. This holds also for graphs without long induced simple cycles. Our results extend to a variant of the conjecture for finite-dimensional Euclidean spaces due to Linial [15].

1. INTRODUCTION

The geometry of graphs, a fascinating area of combinatorics dealing with the geometric representation of graphs, has found many algorithmic applications in recent years. Embedding the metric of a weighted graph into some finite-dimensional real-normed space (see, for instance, the surveys [12, 19]) is a very fruitful and actively studied line of research. In their breakthrough paper [16], Linial, London, and Rabinovich were the first to realize the algorithmic importance of low-distortion metric embeddings. However, their primary motivation was to understand the relationship between the *dimensionality* of a graph and its combinatorial properties.

A notion of dimensionality is usually based on a particular way of embedding a graph into some space that possesses an intrinsic dimension (e.g., a finite-dimensional Euclidean space). One then defines the dimension of a graph to be the minimum dimension into which the graph can be embedded. Several such notions have been extensively studied, see [18]. In [16], the authors wished to express the concept that graphs of large diameter should have low dimensionality. With the help of Leonid Levin, this concept was was formalized as follows.

Let $\mathbb{Z}_\infty^d$ be the infinite graph whose vertices are the elements of $\mathbb{Z}^d$ and such that two elements u and v are adjacent whenever $||u-v||_\infty = 1$. For a graph $G = (V,E)$, define $\dim(G)$ to be the smallest d such that G occurs as a (not necessarily induced) subgraph of $\mathbb{Z}_\infty^d$. For a vertex $v \in V$, let $B(v,r)$ be the ball of radius r (in G) centered at v, and define the *growth rate* of G to be

$$\rho_G := \min\{\rho : |B(v,r)| \le r^\rho \text{ for all } v \in V \text{ and } r > 1\}.$$

Stated differently, $\rho_G = \max\{\frac{\log |B(v,r)|}{\log r} : v \in V, r > 1\}$. Notice that $\rho_{\mathbb{Z}_\infty^d} = \Theta(d)$, so by a simple counting argument, we must have $\dim(G) = \Omega(\rho_G)$. Levin, together with Linial, London, and Rabinovich [16], conjectured that $O(\rho_G)$ dimensions suffice.

Conjecture 1 ([16]). *Every graph G with growth rate ρ_G occurs as a (not necessarily induced) subgraph of $\mathbb{Z}_\infty^{O(\rho_G)}$. In other words,* $\dim(G) = \Theta(\rho_G)$.

It is shown in [16] that Conjecture 1 holds for the k-dimensional hypercube and the complete binary tree, but nothing beyond these two special cases was known. Indeed, previously it was not known whether $\dim(G)$ could be upper bounded by any function of ρ_G, even in the special case of trees. Linial [15] asked about a Euclidean analogue to this notion of dimensionality.

Question 2 ([15]). *For a graph G, what is the minimum dimension $d = \dim_2(G)$ such that there exists a mapping $\gamma : G \to \ell_2^d$ with the following properties?*

1. $||\gamma(u) - \gamma(v)||_2 \ge 1$ *for all* $u \ne v$, *and*
2. $||\gamma(u) - \gamma(v)||_2 \le 2$ *for* $(u,v) \in E(G)$.

*Supported in part by NSF grants CCR-9820951 and CCR-0121555 and by a US-Israeli BSF grant 19999325.
Email: robi@cs.berkeley.edu

†Supported by NSF grant CCR-0121555.
Email: jrl@cs.berkeley.edu

STOC'03, June 9–11, 2003, San Diego, California, USA.

It is easy to verify that the same lower bound holds for $\dim_2(G)$, i.e. $\dim_2(G) = \Omega(\rho_G)$. As we will see later, $\dim(G)$ and $\dim_2(G)$ are closely related.

Dimensionality is a highly important issue in various contexts. For instance, dimensionality reduction is the main tool in solving many problems efficiently, see [12]. Another example is the algorithms designed in [13, 11] to achieve superior performance on restricted growth metrics (a related, but different, notion of dimensionality for metrics spaces). Low-dimensional representations also have a variety of combinatorial consequences, like the existence of good graph decompositions [16]. Indeed, if $\dim(G)$ is small, then one can efficiently find small vertex separators in G. (A precise statement and proof of this is deferred to the full version).

1.1 Results and techniques

In Section 2, we provide a self-contained proof of Levin's conjecture for trees. We first recursively decompose a tree into many hierarchically nested partitions called *levels*. Each level is responsible for pairs of vertices whose distance (in G) falls into a certain range. Fixing a single level, we show that a map drawn from a particular distribution is a good embedding of the level with high probability. If each level is embedded independently, the dimension of the host lattice becomes too large. The key technique, which completes the proof, is to find a way of handling all the levels simultaneously (using the same coordinates). We then extend this result to graphs whose induced simple cycles are of bounded length by relying on known low-distortion embeddings of such graphs into trees, due to [5, 6].

In Section 3, using different techniques, but many ideas from our proof for trees, we give a general upper bound on $\dim(G)$ in terms of certain graph decompositions. In this setting, choosing a good embedding for a level with high probability is more difficult. For this purpose, we modify a technique of [21] (which was used there to embed planar graphs into Euclidean space with low distortion). Again, we must discover a delicate way of handling all the levels simultaneously.

In Section 4, we employ the decomposition of [14], combined with the results of Section 3, to prove the conjecture for any family of graphs which excludes a fixed minor (this includes planar graphs, for instance).

In Section 5, we modify a probabilistic decompositions of Linial and Saks [17] and of Bartal [3] for use with growth-restricted graphs. Our modifications are two-fold. First, the parameters of our decomposition depend on ρ (and not on $n = |V|$ as in [3]). This is essential to our application. Secondly, our decomposition is local in the sense that events which are far apart (in G) are mutually independent (a similar idea was used for a different purpose in [17]). As a result, we are able to apply the Lovász Local Lemma, yielding decompositions which, when combined with the results of Section 3, give $\dim(G) = O(\rho_G^3)$ for any graph G. By delicately combining the steps of Section 3 together via some Chernoff-type tail bounds, and applying the local lemma to their composition, we are able to show an improved upper bound: For any graph G, $\dim(G) = O(\rho_G \log \rho_G)$. In addition, it is shown that our use of the local lemma can be made algorithmic.

In Section 6, we give a lower bound of $\Omega(\rho_G \log \rho_G)$ on the dimension necessary to embed low-degree expander graphs. This shows that Levin's conjecture does not hold in general and that our upper bound is tight. Finally, we show in Section 7 that all our results for $\dim(G)$ hold also for Linial's variant, $\dim_2(G)$.

1.2 Related work

Notions of dimensionality were perhaps first considered by Erdös, Harary, and Tutte [8]. The geometric representations of graphs have been extensively studied in other settings; see, for instance, the survey of Lovász and Vesztergombi [18]. If one defines the cubical dimension of a graph G to be the least k for which G occurs as a subgraph of the k-dimensional hypercube, then characterizing the graphs which have cubical dimension k is one of the most famous open problems in graph theory [4]. Also, as mentioned before, low-distortion metric embeddings have found many applications (see, e.g., the surveys [12, 19]).

A seemingly related notion was considered by Assouad [2] (see also [10]). He conjectured that every metric space (X, d) with restricted growth (suitably defined) is bi-Lipschitz embeddable into $\mathbb{R}^d$ with the Euclidean norm, where d and the bi-Lipschitz constant depend only on the growth rate. He proved that this holds for (X, d^ϵ), for any fixed $0 < \epsilon < 1$, but the conjecture itself (the case $\epsilon = 1$) was disproved by Semmes [22].

Conjecture 1 is actually a dual of the bandwidth problem for graphs. The bandwidth problem asks for the minimum stretch of any edge in an embedding of the graph into $\mathbb{Z} = \mathbb{Z}^1_\infty$. Conjecture 1 asks for the minimum dimension needed to achieve a stretch of one (no stretch). Interestingly, the density bound $D = \max_{v,r}\{\frac{|B(v,r)|}{2r}\}$ (the one-dimensional analogue of the growth rate), which is a straightforward lower bound on the bandwidth, is conjectured to be within an $O(\log n)$ factor of the bandwidth. We know from [9] that the bandwidth and density differ by only a polylog(n) factor. However, even when $D = O(1)$, there are graphs for which the bandwidth is $\Omega(\log n)$ [7]. Thus for $d = 1$, the minimum achievable stretch into $\mathbb{Z}^d_\infty$ cannot be bounded by a function of D. In contrast, our results imply that when $d = \Omega(\rho_G)$, the stretch *can* be bounded by a function which depends only on ρ_G.

1.3 Preliminaries.

DEFINITION 1.1. *For a graph $G = (V, E)$, we will say that a map $\varphi : V \to \mathbb{Z}^d$ is a* contraction *(or* contractive*) if $(u, v) \in E$ implies $||\varphi(u) - \varphi(v)||_\infty \le 1$. Furthermore, if $\{\varphi_i\}$ is a finite set of mappings, we define the* direct sum*, $\varphi = \oplus_i \varphi_i$ to be the mapping $\varphi(u) = (\varphi_1(u), \varphi_2(u), \ldots)$ (where coordinates are concatenated).*

Notice that if a map $\varphi : V \to \mathbb{Z}^d$ is both contractive and injective, then G occurs as a subgraph of $\mathbb{Z}^d_\infty$, and, in particular, $\dim(G) \le d$. We can think of any such embedding φ as consisting of d separate one-dimensional maps $\varphi_1, \ldots, \varphi_d$ such that $\varphi = \oplus_{i=1}^n \varphi_i$. The following trivial lemma will serve as our guide.

LEMMA 1.1. *Let $G = (V, E)$ and $\varphi = \oplus_{i=1}^d \varphi_i$ where $\varphi_i : V \to \mathbb{Z}$, then the following are true.*

1. *φ is a contraction $\iff$ every φ_i is a contraction.*
2. *φ is injective $\iff$ for every pair $u, v \in V$, there exists some φ_i such that $\varphi_i(u) \neq \varphi_i(v)$.*

In what follows, $||\cdot|| = ||\cdot||_\infty$.

2. TREES

In this section we show that every tree T with growth rate $\rho = \rho_T$ occurs as a subgraph of $\mathbb{Z}^d_\infty$, with $d = O(\rho)$. Relying on Lemma 1.1, we will exhibit a map $\varphi : T \to \mathbb{Z}^d$ that is both contractive and injective.

2.1 Embedding trees by random walks

In light of Lemma 1.1, it is natural to define a distribution over random contractions and then argue that some such map must be injective.

Let $T = (V, E)$ be a tree whose growth rate is at most ρ. Fix an arbitrary root r of T, and suppose that the height of T is $O(h)$. Let d_T be the shortest path metric on T. For some constant c to be chosen later, let $T_1, T_2, \ldots, T_{c\rho}$ be $c\rho$ weighted copies of T, where the weight of every edge in T_i is chosen independently and uniformly at random from the set $\{-1, +1\}$. Let d_i be the weighted shortest path metric on T_i. For a vertex v, define its image in $\mathbb{Z}^{c\rho}$ as $\varphi(v) = (v_1, v_2, \ldots, v_{c\rho})$ where $v_i = d_i(r, v)$, i.e. the distance from the root to v in T_i.

Clearly φ is a contraction. Now consider any two vertices $u, v \in V$ for which $d_T(u, v) \geq \sqrt{h}$. The probability that the images of u and v agree in any single coordinate, i.e. that $u_i = v_i$, is the probability that a random walk of length $\sqrt{h}$ ends at its starting point, namely $O(h)^{-1/4}$. Hence the probability that $\varphi(u) = \varphi(v)$ is at most $O(h)^{-c\rho/4}$. Observe that since T is contained in a ball of radius $2h$ centered at r, it contains at most $O(h)^\rho$ vertices. Taking a union bound over at most $O(h)^{2\rho}$ pairs u, v, the probability that any such pair collides is, for any $c > 9$,

$$O(h)^{-c\rho/4} O(h)^{2\rho} = o(1).$$

It follows that for some universal constant h_0 and all $h \geq h_0$, there exists a map $\varphi : V \to \mathbb{Z}^{c\rho}$ such that $d_T(u, v) \geq \sqrt{h} \implies \varphi(u) \neq \varphi(v)$. In what follows, we carefully utilize this simple but powerful idea to show that $\dim(T) = O(\rho_T)$ for any tree T, thus proving Conjecture 1 for the special case of trees.

2.2 Relative embeddings

Consider a tree $\mathcal{T} = (V, E)$ with $\rho = \rho_{\mathcal{T}}$ and fix a root $r_{\mathcal{T}}$ of $\mathcal{T}$. Define a *rooted subtree* of $\mathcal{T}$ to be a connected vertex-induced subgraph T with a distinguished root r_T. Let $W = \{-1, 0, +1\}$ be the set of edge weights. For a rooted subtree $T = (V_T, E_T)$, we define a *d-dimensional relative embedding* of T to be a map $\mu_T : E_T \to W^d$. Finally, we will denote by $\mu^*_T : V_T \to \mathbb{Z}^d$ the *absolute embedding* induced by the relative embedding μ_T (with respect to the root r_T), which is the map obtained as follows: For a vertex $v \in V_T$, define its image $\mu^*_T(v) \in \mathbb{Z}^d$ as the sum of the edge weights along the unique path from r_T to v in T, where an edge $e \in E_T$ has weight $\mu_T(e)$. Note, in particular, that induced embeddings are always contractions. Furthermore, given any contraction $\varphi : V_T \to \mathbb{Z}^d$, there exists a unique d-dimensional relative embedding μ such that $\varphi = \mu^*$ (with respect to r_T). Let us define $\mathbf{0} = (0, 0, \ldots, 0)$ to be the all-zero vector. Notice that the construction of Section 2.1, when applied to a subtree T, essentially yields a relative embedding of T with the following nice property.

LEMMA 2.1 (RELATIVE EMBEDDINGS). *For any rooted subtree $T = (V_T, E_T)$ of $\mathcal{T}$ with height $O(h)$ (where h is larger than some constant), there exists a relative embedding $\mu_T : E \to W^{O(\rho)}$ such that $\mu^*_T(u) \neq \mu^*_T(v)$ whenever $d_{\mathcal{T}}(u, v) \geq \sqrt{h}$.*

In essence μ^*_T "separates" points in T which are far apart relative to the height of T. Notice that the above lemma only works for h sufficiently large. When $h = O(1)$, we can do even better.

LEMMA 2.2 (SMALL SUBTREES). *For any rooted subtree $T = (V_T, E_T)$ of $\mathcal{T}$ with height $h = O(1)$, there exists a relative embedding $\mu_T : E \to W^{O(\rho)}$ such that μ^*_T is injective.*

PROOF. Let $V_T = \{v_1, v_2, \ldots, v_m\}$ and define $\varphi(v_i) = B(i)$ where $B(i)$ is the binary representation of i represented as a $\log m$-dimensional vector. Now φ is clearly injective and also a contraction since $||B(i) - B(j)||_\infty \leq 1$ for all i, j. Finally, notice that since T is of height $h = O(1)$, we have $m \leq O(1)^\rho$ and hence $\log m = O(\rho)$. Let μ be the unique relative embedding such that $\varphi = \mu^*$. It follows that $\mu : E_T \to W^{O(\rho)}$ is a relative embedding with μ^* injective. □

Suppose $\{T_1, T_2, \ldots, T_k\}$ is a collection of vertex-disjoint rooted subtrees of $\mathcal{T}$, and let each $T_i = (V_i, E_i)$ have root r_i. Furthermore, suppose that for each T_i, we have a relative embedding $\mu_i : E_i \to W^d$. Then we can define a relative embedding μ on all of V by

$$\mu(e) = \begin{cases} \mu_i(e) & \text{if } e \in E_i \text{ for some } i \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

Note that μ has the desirable property that $||\mu^*(u) - \mu^*(v)|| = ||\mu^*_i(u) - \mu^*_i(v)||$ whenever $u, v \in V_i$. We will say that μ is obtained by *gluing* the relative embeddings $\{\mu_i\}$ together.

In the sequel, we will construct, for a given h, an embedding $\varphi : V \to \mathbb{Z}^{O(\rho)}$ that satisfies $\varphi(u) \neq \varphi(v)$ whenever $h^{1/2} \leq d(u, v) \leq h$. To do this, we will partition $\mathcal{T}$ into subtrees of height $O(h)$, find for each subtree a relative embedding that satisfies the desired property, and then glue all these relative embeddings into an embedding for $\mathcal{T}$. There is the slight problem that for pairs u, v with $h^{1/2} \leq d(u, v) \leq h$ that end up in different subtrees, we have no guarantee that their images (under μ^*) will be distinct. To handle this, we will actually use two sets of disjoint subtrees which are "staggered" so that every pair u, v with $d(u, v) \leq h$ ends up in the same subtree in at least one of the sets. A more challenging problem is that this embedding is guaranteed to handle only one value of h.

2.3 The leveled decomposition

Let $\Delta = \operatorname{diam}(\mathcal{T})$ and $k = \log \log \Delta$. We define k *levels* $L_0, L_1, \ldots, L_{k-1}$ as follows. Level i consists of two partitions of $\mathcal{T}$ into rooted subtrees, Denote these two partitions A_i and B_i and let $L_i = A_i \cup B_i$. The subtrees in L_i will cover $\mathcal{T}$ (in a sense that will be defined soon) and will each have height at most $3h(i)$, where $h(i) = \Delta^{1/2^i}$. (For convenience, define $h(k) = 1$.) To form A_i, let O^A_i be the set of edges in $\mathcal{T}$ whose *depth* (i.e., distance from the root $r_{\mathcal{T}}$) is a multiple of $3h(i)$. Removing O^A_i from $\mathcal{T}$ results in a collection of disjoint subtrees; let A_i consist of these subtrees, each rooted at its (unique) closest vertex to $r_{\mathcal{T}}$. B_i is formed similarly, except that O^B_i is defined as the set of edges in $\mathcal{T}$ whose depth modulo $3h(i)$ is equal to $h(i)$ (rather than 0). The edges in O^A_i and O^B_i are called the *open edges* of level L_i. The next lemma is easy to verify. In particular, property (3) follows from the "staggering" of the two sets of subtrees A_i and B_i. Property (4) follows from the specifics of the decomposition; it provides a nesting that will turn out to be useful in Section 2.5.

LEMMA 2.3 (THE LEVELED DECOMPOSITION). *For every tree $\mathcal{T} = (V, E)$, the above construction satisfies the following properties.*

1. *Each A_i and each B_i is a partition of V.*
2. *The height of any subtree $T \in L_i$ is at most $3h(i)$.*
3. *For any pair $u, v \in V$ with $d(u, v) \le h(i)$, there is some $T \in L_i$ containing both u and v.*
4. *Each subtree in A_i is entirely contained in some subtree in A_{i+1}; thus, $O^A_{i+1} \subseteq O^A_i$. The same holds for the subtrees in B_i. In this sense, each level is a refinement of the previous level.*

2.4 A first attempt

Let us say that a map φ *separates* A_i (and similarly B_i) if, for every $T \in A_i$ and for every pair $u, v \in V(T)$ with $h(i+1) \le d(u, v) \le h(i)$, we have $\varphi(u) \ne \varphi(v)$. Notice that if $\varphi : V \to \mathbb{Z}^d$ separates A_i and B_i for all $i \in \{0, 1, \ldots, k-1\}$, then φ is injective (by the properties of the decomposition in Lemma 2.3).

Now consider the partition A_i of $\mathcal{T}$. Applying the embedding technique of Lemma 2.1 to each $T \in A_i$ and gluing the relative embeddings together yields an induced embedding φ^A_i which separates A_i. Let φ^B_i be a similar embedding obtained from the partition B_i, and let $\varphi_i : V \to \mathbb{Z}^{O(\rho)}$ be defined as $\varphi_i = \varphi^A_i \oplus \varphi^B_i$. Then φ_i separates A_i and B_i. Finally, as noted before, the map $\varphi = \varphi_0 \oplus \cdots \oplus \varphi_{k-1}$ separates every A_i and B_i, and is hence injective. Since φ is also a contraction, it yields $\dim(\mathcal{T}) = O(\rho k) = O(\rho \log\log \Delta)$. Unfortunately, this bound is not good enough for our purposes.

2.5 Conserving randomness or "Not using all your ammo at once"

Why did we use so many dimensions? Because we needed a distinct set of coordinates for every level. In essence, after fixing an embedding for level L_i, there is no randomness left for "higher" levels $L_{i-1}, L_{i-2}, \ldots, L_0$ (since all the edge weights in the relative embedding for L_i are fixed).

Now consider the open edges of level L_i, namely, O^A_i and O^B_i, which run between disjoint subtrees. In Section 2.2, when the relative embeddings for subtrees are glued together, the open edges are assigned a weight of $\mathbf{0}$. But they might as well have been assigned any other weight in W^d. Clearly the resulting embedding would still be a contraction. Thus even after fixing a relative embedding for L_i, there is still some freedom left to us in deciding how to choose weights for the edges in O^A_i and O^B_i. And as it turns out, this is enough.

We will now show that, after finding a relative embedding for the subtrees in L_{i+2}, there is still enough randomness left to embed the subtrees in L_i simply by assigning random weights to the open edges of L_{i+2}. Notice that this process goes up two levels at a time, from L_{i+2} to L_i, so we will need to do it twice, once for "even" levels and once for "odd" levels. This will increase the number of coordinates used by only a factor of 2.

THEOREM 2.1. *Any tree $\mathcal{T}$ with growth rate ρ occurs as a (not necessarily induced) subgraph of $\mathbb{Z}^{O(\rho)}_\infty$, thus* $\dim(\mathcal{T}) = O(\rho_{\mathcal{T}})$.

PROOF. We will construct four contractions,

$$\varphi^A_{\text{even}}, \varphi^A_{\text{odd}}, \varphi^B_{\text{even}}, \varphi^B_{\text{odd}} : V \to \mathbb{Z}^{O(\rho)}.$$

Let $L_0, L_1, \ldots, L_{k-1}$ be the levels of the decomposition for $\mathcal{T}$, and assume for simplicity that k is odd. Then φ^A_{even} will separate $A_{k-1}, A_{k-3}, \ldots, A_0$, φ^A_{odd} will separate $A_{k-2}, A_{k-4}, \ldots, A_1$, and φ^B_{even} and φ^B_{odd} will satisfy similar properties for the B_i. It will follow from Lemma 2.3, Lemma 1.1, and the discussion in Section 2.4, that $\varphi = \varphi^A_{\text{even}} \oplus \varphi^A_{\text{odd}} \oplus \varphi^B_{\text{even}} \oplus \varphi^B_{\text{odd}}$ is a contractive injection and thus gives an embedding of $\mathcal{T}$ into $\mathbb{Z}^{O(\rho)}_\infty$.

We will construct the map φ^A_{even} inductively. The other maps are constructed similarly. Let $\mu_{k-1} : E \to W^{c\rho}$ be a relative embedding for which μ^*_{k-1} separates A_{k-1}. Since the trees in A_{k-1} have constant height, we can use the construction of Lemma 2.2.

The inductive step. Now assume that we have a relative embedding $\mu_{i+2} : E \to W^{c\rho}$ for which μ^*_{i+2} separates $A_{k-1}, A_{k-3}, \ldots, A_{i+2}$. We will show the existence of a relative embedding $\mu_i : E \to W^{c\rho}$ which satisfies

1. For $T \in A_{i+2}$ and every $u, v \in V(T)$, $||\mu^*_{i+2}(u) - \mu^*_{i+2}(v)|| = ||\mu^*_i(u) - \mu^*_i(v)||$,
2. μ_i separates A_i.

Since the subtrees of A_j for $j \ge i+2$ are all completely nested within the subtrees of A_{i+2} (recall property (4) of Lemma 2.3), the first condition guarantees that μ^*_i separates $A_{k-1}, A_{k-3}, \ldots, A_{i+2}$, since μ^*_{i+2} does.

To obtain μ_i from μ_{i+2}, we will only change the edge weights in O^A_{i+2}, i.e. those running between disjoint subtrees of A_{i+2}. Condition (1) then follows immediately. We now define the relative embedding μ_i randomly and show that it satisfies (2) with positive probability. For $e \in O^A_{i+2}$, choose $\mu_i(e)$ uniformly at random from $\{-1, +1\}^{c\rho}$ and define $\mu_i(e) = \mu_{i+2}(e)$ otherwise.

Let us now show that with positive probability, μ^*_i separates A_i. Fix some $T \in A_i$ and consider two points $u, v \in V(T)$ such that $d(u, v) \ge h(i+1) = h(i)^{1/2}$. Let P_{uv} be the unique path from u to v in T. Since P_{uv} has length at least $h(i)^{1/2}$ and each subtree of A_{i+2} has height at most $3h(i+2) = 3h(i)^{1/4}$, P_{uv} must pass through at least $\frac{1}{3}h(i)^{1/4}$ such subtrees. In particular, the path includes at least $\frac{1}{3}h(i)^{1/4}$ edges from O^A_{i+2}.

Now consider the part of P_{uv} which is composed of edges whose weights are already fixed (i.e., edges not in O^A_{i+2}). The sum of their weights is fixed, and the probability that a random walk of length $\frac{1}{3}h(i)^{1/4}$ (along open edges) is equal to the negation of any fixed sum is at most $O(h(i))^{-1/8}$. This also upper bounds the probability that the images of u and v (under μ^*_i) agree in any single coordinate. So the probability of this occurring in $c\rho$ independent coordinates is $\Pr[\mu^*_i(u) = \mu^*_i(v)] = O(h(i))^{-c\rho/8}$. Finally, notice that T has height at most $3h(i)$, and thus contains at most $(3h(i))^{2\rho}$ pairs of vertices. Choosing c to be a large enough constant (which is *independent of* i), the union bound $O(h(i))^{-c\rho/8} O(h(i))^{2\rho} = o(1)$ shows the existence of the desired map μ_i on T. Continuing in this way for each disjoint subtree $T \in A_i$, we see that with positive probability, μ_i satisfies condition (2).

By induction, μ_0^* separates each of $A_{k-1}, \ldots, A_0$. Setting $\varphi_{\text{even}}^A = \mu_0^*$ completes the proof. $\square$

For a graph G, let $\lambda(G)$ be the length of the longest induced simple cycle in G. The proof of the following theorem, which uses Theorem 2.1 and the low-distortion embeddings of [5, 6] is omitted from this version.

THEOREM 2.2. *Conjecture 1 holds for for any class of graphs in which $\lambda(G)$ is bounded, For such graphs,* $\dim(G) = O(\rho_G)$. *This includes trees and chordal graphs.*

3. DIMENSION UPPER BOUNDS VIA GRAPH DECOMPOSITION

In this section, we use the ideas of the previous section to prove a result on general graphs in terms of their decompositions (Theorem 3.1).

In what follows, let $G = (V, E)$ be a simple graph with growth rate $\rho = \rho_G$. For any set S, define $\mathcal{P}(S)$ to be the power set of S. For a subset $S \subseteq V$, define $\partial S = \{u \in S : \exists v \notin S \text{ s.t. } (u, v) \in E\}$ to be the *boundary of* S. For a collection $\mathcal{C} \subseteq \mathcal{P}(V)$, define $\partial\mathcal{C} = \cup_{S \in \mathcal{C}} \partial S$ as the the boundary of $\mathcal{C}$. For a subset $S \subseteq V$, let $G[S]$ be the subgraph induced on the vertices of S. When $G[S]$ is connected, we refer to such a subset as a *cluster*, and define the *weak diameter of* S to be $\operatorname{diam}(S) = \max_{u,v \in S} d_G(u, v)$. Finally, set $\Delta = \operatorname{diam}(G)$.

3.1 Relative embeddings and the padded decomposition

Suppose we are given a cluster $S \subseteq V$. Define a *d-dimensional relative embedding* of S to be a contraction $\varphi : S \to \mathbb{Z}^d$ such that $\varphi(\partial S) = \mathbf{0}$, i.e. the boundary is mapped to $\mathbf{0}$. Suppose further that we would like to find a relative embedding of S with the following property: For every $u, v \in S$ with $d(u, v) > r$ and such that $B(u, 3r^{1/2}) \subseteq S$, we have $\varphi(u) \neq \varphi(v)$. In other words, since we are imposing the rather stringent condition that $\varphi(\partial S) = \mathbf{0}$, we only make requirements on vertices which are far enough away from the boundary.

We will produce such an embedding using a technique similar to that of Rao [21]. First we will partition S into clusters of diameter at most r. It follows that if $d(u, v) > r$, then u and v will end up in different clusters. We will define the image of a vertex in a single coordinate to be simply the distance from that vertex to the boundary of its cluster, but to achieve injectiveness with high probability, we will first randomly contract the boundary inward. We first discuss our method of choosing clusters. Let us call a collection of disjoint subsets $P \subseteq \mathcal{P}(S)$ a *semi-partition* of S.

DEFINITION 3.1 (THE PADDED DECOMPOSITION). *Let $S \subseteq V$ be a cluster in G. A set $\{P_1, P_2, \ldots, P_m\}$ of m semi-partitions of S is called an* r-padded decomposition of S with m layers *if the following properties are satisfied.*

1. *Every $C \in \cup_{i=1}^m P_i$ is a cluster with $\operatorname{diam}(C) \leq r^\alpha$.*
2. *For every $u \in S$ with $B(u, 3r) \subseteq S$, there exists some $C \in \cup_{i=1}^m P_i$ such that $B(u, 3r) \subseteq C$.*

Any $\alpha = O(1)$ suffices in the above definition. For ease of notation, define an r-inner decomposition *to be an $r^{1/\alpha}$-padded decomposition.*

Note that in an r-inner decomposition, clusters have $\operatorname{diam}(C) \leq r$ and vertices have "padding" of the form $B(u, 3r^{1/\alpha})$. We now show how to use an r-inner decomposition to produce a "good" relative embedding.

LEMMA 3.1 (RELATIVE EMBEDDINGS). *Let $S \subseteq V$ be a cluster with $|S| \leq r^{O(\rho)}$. If S has an r-inner decomposition with m layers, then there exists a relative embedding $\varphi : S \to \mathbb{Z}^{O(m\rho)}$ such that for every $u, v \in S$ with $d(u, v) > r$ and $B(u, 3r^{1/\alpha}) \subseteq S$, we have $\varphi(u) \neq \varphi(v)$.*

PROOF. Let $P_1, \ldots, P_m$ be the semi-partitions of the inner decomposition. For each P_i ($i = 1, \ldots, m$) we will construct a map $\varphi_i : S \to \mathbb{Z}^{c\rho}$; eventually we shall take the direct sum of all these maps.

Fix some semi-partition P_i and form a single coordinate $\varphi_i^0 : S \to \mathbb{Z}$ as follows: For every cluster $C \in P_i$, choose uniformly at random $r_C \in \{0, 1, \ldots, r^{1/\alpha}\}$, and let $\partial_C^* = \{v \in C : d(v, \partial C) \leq r_C\}$ be the boundary of C randomly contracted inward. Now for every $u \in C$, define $\varphi_i^0(u) = d(u, \partial_C^*)$.

Clearly $\varphi_i^0(u)$ is a contraction, since whenever $(u, v) \in E$, their distances to ∂_C^* can differ by at most 1. It is also clear that $u \in \partial S$ implies $u \in \partial C$ for some C and hence $\varphi_i^0(u) = 0$. It follows that $\varphi_i^0(\partial S) = 0$.

Now independently form $c\rho$ such coordinates (each time picking fresh values for the r_C) and let φ_i be the direct sum of the resulting maps. Finally, set $\varphi = \varphi_1 \oplus \cdots \oplus \varphi_m$. From the properties of φ_i, we conclude that $\varphi : S \to \mathbb{Z}^{cm\rho}$ is a contraction which maps ∂S to $\mathbf{0}$.

Consider a pair u, v with $d(u, v) > r$ and such that $B(u, 3r^{1/\alpha}) \subseteq S$. It follows from from property (2) of Definition 3.1 that there exists a semi-partition P_i and a cluster $C \in P_i$ for which $B(u, 3r^{1/\alpha}) \subseteq C$. Since $d(u, v) > r$, u and v must lie in different clusters of P_i. It follows that, in any single coordinate φ_i^0 of the map φ_i, the value of $\varphi_i^0(u)$ is distributed uniformly over an interval of size $r^{1/\alpha}$ independently of the value $\varphi_i^0(v)$, hence $\Pr[\varphi_i^0(u) = \varphi_i^0(v)] \leq r^{-1/\alpha}$. Thus the probability that u and v collide under φ_i is $\Pr[\varphi_i(u) = \varphi_i(v)] \leq r^{-c\rho/\alpha}$. Since $|S| \leq r^{O(\rho)}$, there are at most $r^{O(2\rho)}$ such pairs u, v, and hence the probability that some pair collides is at most $r^{O(2\rho)} r^{-c\rho/\alpha} < 1/2$ for sufficiently large constant c (remembering that $\alpha = O(1)$). The existence of a map φ satisfying the lemma follows. $\square$

3.2 A simple approach

Now that we can find relative embeddings for clusters, we will use the padded decomposition to decompose G into levels of disjoint clusters that cover G (by setting $S = V$ in Definition 3.1). Finding a relative embedding for each cluster, and then gluing all these embedding together, we will arrive at a good embedding for G. Note that the padded decomposition is serving two purposes here. First, it is being used *inside* clusters to compute a good relative embedding (this is like choosing random edge weights to produce a relative embedding for trees). Second, it is being used to decompose the graph into the clusters which will be separately embedded (this is similar to the leveled decomposition of Section 2).

Here is a simple approach which will fail in the end, but will give some intuition as to how the padded decomposition will be used. Let $k = \log\log\Delta$. Set $r_i = \Delta^{1/2^i}$ for $i \in \{0, 1, \ldots, k-1\}$, and $r_k = 0$ (notice that $r_{k-1} =$

$O(1)$). Suppose we use the padded decomposition with $r = r_0, r_1, \ldots, r_{k-1}$. For each value of r, the r-padded decomposition will be used to break the graph into clusters of diameter at most r^α such that every two vertices within a distance r are contained in some such piece (and are "far" from the boundary of the piece—this is needed to ensure that they are "separated" by the relative embedding for that cluster).

An embedding for one level. Assume that $j < k-1$. Let $\{P_1, P_2, \ldots, P_m\}$ be the semi-partitions produced by the r_j-padded decomposition. We will show how to construct a contraction φ_j that satisfies: For every pair $u, v \in V$ with $r_{j+1} < d(u,v) \le r_j$, $\varphi_j(u) \neq \varphi_j(v)$.

Fix some single semi-partition P_i of V. For every cluster $C \in P_i$, compute a relative embedding ψ_C by applying Lemma 3.1 with parameter $r_{j+1} = r_j^{1/2}$. Note that the lemma is applicable since $\mathrm{diam}(C) \le r_j^\alpha$ implies $|C| \le r_{j+1}^{2\alpha\rho} = r_{j+1}^{O(\rho)}$. Now set $\psi_i(u) = \psi_C(u)$ if u belongs to some cluster C and $\psi_i(u) = \mathbf{0}$ otherwise. Notice that this map is well-defined since the clusters $C \in P_i$ are disjoint. Also, notice that it is a contraction, for suppose $(u,v) \in E$. If u and v are in the same cluster C, then $||\psi_i(u) - \psi_i(v)|| = ||\psi_C(u) - \psi_C(v)|| \le 1$ since ψ_C is a relative embedding, and hence a contraction. If u and v are in different clusters, or one or both of them are contained in no cluster, then $\psi_i(u) = \psi_i(v) = \mathbf{0}$ since both of u and v are either on the boundary of their cluster (which is mapped to $\mathbf{0}$), or in no cluster at all. Finally, set $\varphi_j = \psi_1 \oplus \cdots \oplus \psi_m$.

Now consider some $u, v \in V$ with $r_{j+1} < d(u,v) \le r_j$. By property (2) of Definition 3.1, there exists some partition P_i and a cluster $C \in P_i$ such that $B(u, 3r_j) \subseteq C$. Thus $u, v \in C$, and certainly $B(u, 3r_j^{1/2\alpha}) \subseteq C$, so by Lemma 3.1, $\psi_C(u) \neq \psi_C(v)$. It follows that $\psi_i(u) \neq \psi_i(v)$, and hence $\varphi_j(u) \neq \varphi_j(v)$.

The base case. For $r = r_{k-1}$, we will construct φ_{k-1} in a special way so that for $u, v \in V$ with $d(u,v) \le r_{k-1}$, $\varphi_{k-1}(u) \neq \varphi_{k-1}(v)$. We do this as above, but using a much simpler relative embedding technique: Given a cluster C with $C = \{v_1, v_2, \ldots, v_s\}$ and $s = r_{k-1}^{O(\rho)}$, define the relative embedding $\psi_C(v_i) = B(i)$ if $v_i \notin \partial C$ and $\psi_C(v_i) = \mathbf{0}$ otherwise, where $B(i)$ is the binary representation of i as a $\log s$-dimensional vector. Note that this map is a contraction and satisfies $\psi_C(u) \neq \psi_C(v)$ whenever $u, v \notin \partial C$. Using this technique in the above argument (instead of Lemma 3.1) yields the desired map φ_{k-1}. Also notice that φ_{k-1} uses only $m \log s = O(m\rho)$ coordinates (since $r_{k-1} = O(1)$).

Putting it all together. If we let $\varphi = \varphi_0 \oplus \varphi_1 \oplus \cdots \oplus \varphi_{k-1}$, we see that φ is in fact a contractive, injective embedding because for any pair u, v, the distance $d(u,v)$ falls into some range $r_{j+1} < d(u,v) \le r_j$ and thus $\varphi_j(u) \neq \varphi_j(v)$.

If we assume that, for every sufficiently large value of r, we can construct r-padded decompositions with m layers, then each ψ_i (From Lemma 3.1) uses $O(m\rho)$ coordinates and thus each φ_j uses $O(m^2\rho)$ coordinates. It follows that φ uses $O(m^2\rho \log\log\Delta)$ coordinates in all. It turns out that this bound is of the right form, except for the appearance of Δ. And indeed, this dependence on Δ will be removed in the next subsection.

In the case of trees, we removed the dependence of the dimension on Δ by exploiting some "untapped randomness." We saw that even after fixing a relative embedding for a single level, we were still free to assign arbitrary weights to the open edges of that level. In the next section, we exploit a similar observation, namely that the boundary of a cluster need not be mapped to $\mathbf{0}$.

3.3 Forced nesting, contractions, and untapped randomness

We can now prove the main result of this section.

THEOREM 3.1. *Let G be a graph with $\rho = \rho_G$. If for every cluster $S \subseteq V$ and every value $r \le diam(G)$ that is larger than a suitable constant, there exists an r-padded decomposition of S with m layers, then* $\dim(G) = O(m^2\rho)$.

First, consider a single layer P_i of an r-padded decomposition, consisting of many disjoint clusters. For each $C \in P_i$, let ψ_C be the relative embedding for C. Previously, we set $\psi_i(u) = \psi_C(u)$ for $u \in C$ and $\psi_i(u) = \mathbf{0}$ otherwise. This yielded a contraction defined on all of V with the property that for $u, v \in C$, $||\psi_i(u) - \psi_i(v)|| = ||\psi_C(u) - \psi_C(v)||$. The following lemma gives a simple way of maintaining this property, while allowing some freedom in choosing ψ_i.

LEMMA 3.2. *Let P be a partition of V into disjoint clusters. Suppose that, for each cluster $C \in P$, we have a contraction $\psi_C : C \to \mathbb{Z}^d$. Let $G_P = (V_P, E_P)$ be the graph obtained from G by contracting (in the graph-theoretic sense) each cluster C to a single vertex, and let $\psi_P : V_P \to \mathbb{Z}^d$ be a contraction of G_P. We may identify V_P with P and think of ψ_P as a map defined on P. Let C_u be the cluster containing u and define the map $\psi : V \to \mathbb{Z}^d$ as follows: $\psi(u) = \psi_P(C_u) + \psi_C(u)$. Then ψ is a contraction and for $u, v \in C$, $||\psi(u) - \psi(v)|| = ||\psi_C(u) - \psi_C(v)||$.*

PROOF. It is clear that for $u, v \in C$, $||\psi(u) - \psi(v)|| = ||\psi_C(u) - \psi_C(v)||$. It follows that ψ contracts vertices which are in the same cluster. For u, v in different clusters with $(u,v) \in E$, $||\psi(u) - \psi(v)|| = ||\psi_P(C_u) - \psi_P(C_v)|| \le 1$ because ψ_P is a contraction and $(u,v) \in E$ implies $(C_u, C_v) \in E_P$. $\square$

The above lemma tells us that even after fixing a relative embedding for each cluster $C \in P_i$, we still have some freedom in choosing the map ψ_i (every contraction ψ_P in Lemma 3.2 gives rise to a valid ψ_i). The following proof has a structure which is similar to that of Theorem 2.1 for trees, but is technically more difficult. Basically, we want to find an embedding φ_j which is "good" for level j and then modify it so that it remains good for j, and at the same time it is also good for level $j - t$ for some $t = O(1)$ (in Theorem 2.1 we had $t = 2$, giving rise to the even and odd levels). Continuing inductively, we will handle a constant fraction of the levels with only one set of $O(m^2\rho)$ coordinates. Doing this t times and taking the direct sum of the results, we obtain an embedding which uses only $O(tm^2\rho) = O(m^2\rho)$ coordinates.

PROOF OF THEOREM 3.1 (SKETCH). Let $\Delta = \mathrm{diam}(G)$, and suppose that for some constant R_0, every r with $R_0 \le r \le \Delta$, and every cluster $S \subseteq V$, there exists an r-padded decomposition of S with m layers. Define $k = \log\log\Delta$ and set $r_i = \Delta^{1/2^i}$ for $i \in \{0, 1, \ldots, k-1\}$, $r_k = 0$. Let $\mathcal{P}_i = \{P_1^i, \ldots, P_m^i\}$ be the r-padded decompositions corresponding to $S = V$ and $r = r_i$ for $i \in \{0, 1, \ldots, k-1\}$.

In Section 3.2, we used one $O(m^2\rho)$-dimensional embedding for every level $\mathcal{P}_i$. This resulted in the undesirable $\log\log\Delta$ factor in the dimension. To overcome this, we shall use the same $O(m^2\rho)$-dimensional embedding for say the first layer of every level $\mathcal{P}_i$, and so forth. Consider then the sequence of layers $P_1^1, P_1^2, \ldots, P_1^{k-1}$. Let us say that a map φ *separates* a partition P_1^i if, for any $u, v \in V$ with $r_{i+1} < d(u,v) \le r_i$, such that there exists a cluster $C \in P_1^i$ with $B(u, 3r_i) \subseteq C$, we have $\varphi(u) \neq \varphi(v)$.

We will show how to construct a map $\varphi_1 : V \to \mathbb{Z}^{O(m\rho)}$ which separates the first layer of every level simultaneously, i.e., φ_1 separates $P_1^0, P_1^1, \ldots, P_1^{k-1}$. Constructing similar maps φ_i for each sequence $P_i^0, P_i^1, \ldots, P_i^{k-1}$, and taking the sum $\hat{\varphi} = \oplus_{i=1}^m \varphi_i$, we arrive at an embedding of G into the $O(m^2\rho)$-dimensional lattice, and hence $\dim(G) = O(m^2\rho)$, proving the theorem.

Let $t = \log 4\alpha$ (where $\alpha = O(1)$ appears in Definition 3.1) and set $b = \lfloor (k-1)/t \rfloor$. We will construct, for each $i = 0, 1, \ldots, t-1$, a map φ_1^i which separates $P_1^i, P_1^{t+i}, P_1^{2t+i}, \ldots, P_1^{bt+i}$. Letting $\varphi_1 = \oplus_{i=0}^{t-1} \varphi_1^i$ will yield the desired map φ_1. In what follows, we will construct only $\varphi = \varphi_1^0 : V \to \mathbb{Z}^{O(m\rho)}$, which will separate $P_1^0, P_1^t, \ldots, P_1^{bt}$. The other maps are constructed similarly. We now sketch the remaining details. A more technical proof is deferred to the full version.

Forced nesting. First, with some negligible loss, we can force the sequence of semi-partitions $P_1^0, P_1^t, \ldots, P_1^{bt}$ to be nested so that P_1^{it} is a refinement of $P_1^{(i+1)t}$ for all $0 \le i < b$. In other words, for every cluster $S \in P_1^{(i+1)t}$, there is a cluster $S' \in P_1^{it}$ which contains S. This can be achieved inductively as follows: Suppose that $P_1^{(i+1)t}, P_1^{(i+2)t}, \ldots, P_1^{bt}$ are nested. If there is any cluster $S \in P_1^{(i+1)t}$ that is not contained completely in a cluster $S' \in P_1^{it}$ (for instance, if it is split between two or more clusters of P_1^{it}), then simply make every vertex of S its own cluster in P_1^{it}. That is, clusters of P_1^{it} are contracted inwards to avoid breaking any clusters of $P_1^{(i+1)t}$. The nesting property is now clearly satisfied. Also, note that if $B(u, 3r_{it})$ is contained in some cluster S' of P_1^{it}, then $B(u, 3r_{it} - r_{(i+1)t}^\alpha) \subseteq S'$ holds after the modification, since the diameter of any removed cluster $S \in P_1^{(i+1)t}$ is at most $r_{(i+1)t}^\alpha$ by Definition 3.1. Note that $r_{(i+1)t}^\alpha \le r_{it}^{1/4}$, so this loss is asymptotically negligible.

We will now proceed as follows. For $i = b, b-1, \ldots, 0$, we will exhibit a map $\gamma_i : V \to \mathbb{Z}^{O(m\rho)}$ which separates $P_1^{bt}, \ldots, P_1^{(i+1)t}, P_1^{it}$. Setting $\varphi = \gamma_0$ will complete the proof. γ_i will be constructed by modifying γ_{i+1} using Lemma 3.2.

Suppose we have a map γ_{i+1} separating $P_1^{bt}, \ldots, P_1^{(i+1)t}$. We would like to modify γ_{i+1} to a map γ_i so that it also separates P_1^{it}. This is accomplished by mapping the boundaries of the clusters in $P_1^{(i+1)t}$ to values other than $\mathbf{0}$ using Lemma 3.2.

To this end, consider a cluster $C \in P_1^{it}$. Let us see how to modify the proof of Lemma 3.1 to construct a relative embedding of C which only changes the images of the boundaries of the clusters in $P_1^{(i+1)t}$. We construct an r_{it}-inner decomposition of C as before, but force it to be nested with the clusters $U = \{C' \in P_1^{(i+1)t} : C' \subset C\}$ using the technique discussed above. Then, each $C' \in U$ is contracted to a single vertex. To the resulting contracted graph, we apply the embedding technique of Lemma 3.1. The only concern is that a padded vertex, i.e. one for which $B(u, 3r_{it}) \subseteq C$ has significantly less padding in the contracted graph. But it is not too difficult to see that, during contraction, the padding decreases by only a polynomial factor (in r_{it}). Thus by increasing the constant of Lemma 3.1, the union bound can still be applied.

After a relative embedding for the contracted version of C is found, we apply Lemma 3.2 to arrive at a relative embedding of C which is still good for the clusters $C' \in U$. Finally, we construct γ_i as in Section 3.2, using the relative embeddings exhibited for the clusters $C \in P_1^{it}$. $\square$

4. GRAPHS EXCLUDING A FIXED MINOR

Let G be a graph which excludes a $K_{s,s}$ minor for some fixed s. Then we can use a decomposition technique of Klein, Plotkin, and Rao [14], to construct, for any value of r, an r-padded decomposition of G with only $O(2^s)$ layers. Applying Theorem 3.1, we arrive at the main result of this section. Due to space considerations, the proof is deferred to the full version of this paper.

THEOREM 4.1. *Conjecture 1 is true for any family of graphs that excludes a fixed minor, i.e.,* $\dim(G) = O(\rho_G)$ *for any such graph* G.

5. A GENERAL DIMENSION UPPER BOUND

In this section we give a tight upper bound on the dimension of general graphs: $\dim(G) = O(\rho_G \log \rho_G)$ for any graph G. (We show that this upper bound is met by expanders in Section 6). First, we devise a decomposition for growth-restricted metrics (Section 5.1) and use Theorem 3.1 to obtain a weaker upper bound of $O(\rho_G^3)$ (Section 5.2). Then, by combining the previous arguments more carefully and utilizing some Chernoff-type tail bounds, we obtain the aforementioned tight upper bound (Section 5.3).

5.1 Decomposition of growth-restricted graphs

Linial and Saks [17] and Bartal [3] show that for any graph $G = (V, E)$ and $1 \le r \le \mathrm{diam}(G)$, there exists a probabilistic partition of G into disjoint clusters of diameter at most $O(r \ln n)$, such that for any pair of vertices $u, v \in V$, the probability that u and v end up in different clusters is at most $O(d(u,v)/r)$. Let $\rho = \rho_G$. In this section, we give a similar decomposition, but we replace the diameter bound of $O(r \ln n)$ with a bound which is independent of n, namely $O(\rho\, r \ln r)$, for $r \ge \rho$. Our partitioning method is similar to those of [17] and [3], but different in a subtle and crucial way: It is local. Events which are sufficiently far apart are mutually independent.

First, take the continuous exponential distribution with mean r, truncate it at M and rescale the remaining density function. The resulting distribution, which we denote $\mathrm{Texp}(r, M)$, has density function $p(z) = \frac{e^{M/r}}{r(e^{M/r}-1)} e^{-z/r}$ for $z \in (0, M)$.

The algorithm. Let $V = \{v_1, v_2, \ldots, v_n\}$. For each v_t, choose a radius r_t according to the distribution $\mathrm{Texp}(r, 8\rho r \ln r)$. For a vertex v, define $B_v = \min\{t : v \in B(v_t, r_t)\}$. Finally, define $S_t = \{v : B_v = t\}$ as the set of vertices v for which $B(v_t, r_t)$ is the first ball containing v. Notice that $G[S_t]$ may be disconnected, so define the set of clusters to be $\mathcal{C} = \cup_t \{\text{connected components of } G[S_t]\}$.

It is easy to bound the (weak) diameter of every cluster $C \in \mathcal{C}$ by $\mathrm{diam}(C) \le 16\rho r \ln r$. Further analysis will require the following simple facts, which show that the truncated exponential distribution is "almost" memoryless.

FACT 5.1. *Consider a random variable $R \sim \mathrm{Texp}(r, M)$ for $M \geq 2r$. Then,*

1. $\Pr[R \geq \alpha] \leq 2e^{-\alpha/r}$.
2. $\Pr[R \leq \alpha] \leq 2(1 - e^{-\alpha/r}) \leq 2\alpha/r$.
3. *if* $R_0 \leq M/2$ *then* $\Pr[R \leq R_0 + \alpha \,|\, R \geq R_0] \leq 2\alpha/r$.

For a vertex $u \in V$, let $\mathcal{E}_u^x$ be the event that $B(u, x)$ is split between multiple clusters, i.e., the event that no cluster $C \in \mathcal{C}$ fully contains $B(u, x)$.

THEOREM 5.1. *Let $u \in V$ and $r \geq \rho$. Then $\Pr[\mathcal{E}_u^x] \leq 10x/r$.*

PROOF. Assume $1 \leq x \leq r$ (the theorem says nothing for larger x) and let $\mathcal{B} = B(u, x)$, $B_t = B(v_t, r_t)$. Let us say that the ball $\mathcal{B}$ is cut by the ball B_t if $S_t \cap \mathcal{B}$ is non-empty but $\mathcal{B} \nsubseteq S_t$. Then $\mathcal{E}_u^x$ is precisely the event that some ball B_t cuts $\mathcal{B}$.

Let us separate the cuts into two classes, depending on the distance from v_t to u. Define $\mathcal{E}_{\mathrm{far}}$ to be the event that B_t cuts $\mathcal{B}$ and $d(v_t, u) \geq 4\rho r \ln r$. Define $\mathcal{E}_{\mathrm{near}}$ to be the event that B_t cuts $\mathcal{B}$ and $d(v_t, u) \leq 4\rho r \ln r$.

Fix v_t with $d(u, v_t) \geq 4\rho r \ln r$ and notice that by Fact 5.1,

$$\Pr[B_t \text{ cuts } \mathcal{B}] \leq \Pr[r_t \geq 4\rho r \ln r - x] \leq 2r^{-4\rho} e^{x/r} \leq 6r^{-4\rho}.$$

But the number of such v_t for which B_t can possibly cut $\mathcal{B}$ is at most the number of points in a ball of radius $(8\rho r \ln r + x) \leq r^3$ which is at most $r^{3\rho}$. Taking a union bound over all such possible v_t, we see that $\Pr[\mathcal{E}_{\mathrm{far}}] \leq 6r^{-4\rho} r^{3\rho} \leq 6/r^{\rho} \leq 6/r$. Thus we are left only to bound the probability of $\mathcal{E}_{\mathrm{near}}$.

Let the random variable T be the minimum t such that $B_T \cap \mathcal{B} \neq \emptyset$ (note that possibly $v_T \in \mathcal{B}$). The ball B_T can either cut $\mathcal{B}$ (in which case $\mathcal{E}_u^x$ occurs) or contain $\mathcal{B}$ (and then $\mathcal{B} \subseteq S_t$ is not cut by any ball B_t). By the principle of deferred decision it suffices to upper bound the conditional probability $\Pr[\mathcal{E}_{\mathrm{near}} | T = t]$. To this end, we may assume that $d(v_t, u) \leq 4\rho r \ln r$ (as otherwise this conditional probability is 0) and then $\mathcal{E}_{\mathrm{near}}$ happens if and only if B_t cuts $\mathcal{B}$, which in turn happens only if $r_t \leq d(v_t, u) + x$. Hence,

$$\Pr[\mathcal{E}_{\mathrm{near}} \,|\, T = t] \leq \Pr\left[r_t \leq d(v_t, u) + x \,|\, r_t \geq d(v_t, u) - x\right] \leq \frac{4x}{r},$$

where we have used Fact 5.1 in conjunction with $d(v_t, u) \leq 4\rho r \ln r$. Thus, $\Pr[\mathcal{E}_{\mathrm{near}}] = \sum_t \Pr[T = t] \cdot \Pr[\mathcal{E}_{\mathrm{near}} | T = t] \leq 4x/r$ and $\Pr[\mathcal{E}_u^x] \leq \Pr[\mathcal{E}_{\mathrm{near}}] + \Pr[\mathcal{E}_{\mathrm{far}}] \leq 10x/r$. $\square$

5.2 General layered decompositions

Now we describe how to obtain an r-padded decomposition with $O(\rho_G)$ layers for general graphs G. Plugging these values into Theorem 3.1 yields an embedding into $O(\rho_G^3)$ dimensions. We will only be able to show the existence of such decompositions under the assumption that $r \geq \rho$. In the case where $r \leq \rho$, clusters of diameter $r^{O(1)}$ have at most $\rho^{O(\rho)}$ points, so we will be able to embed these by brute force using only $O(\rho \log \rho)$ dimensions (similar to the base case of Section 3.2). The final result appears in Theorem 5.3.

THEOREM 5.2. *Let r_0 be a sufficiently large constant. Then for any graph $G = (V, E)$ with $\rho = \rho_G$ and any $r \geq \max\{r_0, \rho\}$, there exists an r-padded decomposition with $O(\rho)$ layers.*

PROOF. Assume $r \geq \max\{r_0, \rho\}$. To produce a single layer of the decomposition (a partition of V into clusters), we will use the algorithm of Section 5.1, with the parameter r (in the algorithm and in Theorem 5.1) set to r^2. Notice that the clusters produced have diameter at most $32\rho r^2 \ln r \leq r^4$ (for r_0 sufficiently large). For a vertex $v \in V$, let $\mathcal{E}_v$ be the event that the ball of radius $3r$ about v is cut (i.e., split amongst two or more clusters). From Theorem 5.1, we know that $\Pr[\mathcal{E}_v] \leq O(1/r)$.

Now produce ℓ layers independently (with fresh random coins each time) and let $\mathcal{E}_v^\ell$ be the event that the ball of radius $3r$ about v is cut in *every* layer. Clearly $\Pr[\mathcal{E}_v^\ell] \leq O(1/r)^\ell$. We would like to say that $\Pr\left[\bigwedge_{v \in V} \overline{\mathcal{E}_v^\ell}\right] > 0$. If we could show this with $\ell = O(\rho)$, the theorem would follow. And indeed, this is our goal. We will employ the following symmetric form of the Lovász Local Lemma, see e.g. [1].

LEMMA 5.2 (LOVÁSZ LOCAL LEMMA). *Let $A_1, \ldots, A_n$ be events in an arbitrary probability space. Suppose that for each A_i there is a set that contains all the other events A_j but at most d, such that A_i is mutually independent of this set of events, and suppose that $\Pr[A_i] \leq p$ for all $1 \leq i \leq n$. If $ep(d + 1) \leq 1$ then $\Pr[\wedge_{i=1}^n \overline{A_i}] > 0$.*

Let $r_1 = 32r^3 \ln r + 6r$. An event $\mathcal{E}_u^\ell$ is mutually independent of all events $\mathcal{E}_v^\ell$ for which $d(u, v) > r_1$ because every ball of the decomposition of Section 5.1 has radius at most $16r^3 \ln r$ and thus cannot intersect both $B(u, 3r)$ and $B(v, 3r)$. It follows that $\mathcal{E}_u^\ell$ is mutually independent of the set of all events $\mathcal{E}_v^\ell$ except those for which $v \in B(u, r_1)$, and there are at most $d = r_1^\rho$ such events. Thus if $\Pr[\mathcal{E}_u^\ell] \leq e/(d+1)$, we can apply the local lemma and the theorem is proved. But this is easily accomplished by choosing $\ell = 4\rho$ (for r_0 a sufficiently large constant). By applying Lemma 5.2 we conclude that there exists an r-padded decomposition for V (with $\alpha = 4$). To get an r-padded decomposition for any cluster $S \subseteq V$, just apply the theorem to $G[S]$. $\square$

THEOREM 5.3. *For every graph G with growth rate ρ_G, $\dim(G) = O(\rho_G^3)$.*

PROOF (SKETCH). Instead of starting the proof of Theorem 3.1 at $r = O(1)$, we start with the level corresponding to $r = \max\{\rho, r_0\}$. In this case, the clusters have diameter $r^{O(1)} = \rho^{O(1)}$, so we can easily give a relative embedding for each cluster using only $O(\rho \log \rho)$ coordinates, similar to the base case of 3.2. The proof then proceeds unchanged. $\square$

5.3 A tight upper bound

As mentioned previously, Theorem 5.2, combined with Theorem 3.1, shows that $\dim(G) = O(\rho_G^3)$ for every graph G. Thus the dimension of a graph is indeed bounded above by a function which depends only on its growth rate. We can do better, though; by carefully combining the previous arguments and utilizing some Chernoff-type tail bounds, we are able to find a tight upper bound, $\dim(G) = O(\rho_G \log \rho_G)$.

THEOREM 5.4. *For every graph $G = (V, E)$ with growth rate ρ_G, $\dim(G) = O(\rho_G \log \rho_G)$.*

PROOF (SKETCH). Consider again the $O(\rho)$-layer r-padded decomposition who existence is guaranteed by Theorem 5.2. Suppose we have $\ell = c_0\rho$ layers $P_1, P_2, \ldots, P_\ell$. Let us say that a vertex u is padded in a layer i if there exists a cluster

$C \in P_i$ such that $B(u, 3r) \subseteq C$ (otherwise, we will say that it is unpadded in layer i).

First let us show that with positive probability, every vertex $u \in V$ is padded in a constant fraction of the ℓ layers. The probability that u is unpadded in layer P_i is at most c_1/r (by the analysis of Theorem 5.1) for some constant c_1, so the expected number of layers in which u is bad is at most $c_1\ell/r$. We now need the following Chernoff-type tail bound (see, e.g., [20]).

LEMMA 5.3 (A TAIL BOUND). *Let $X_1, X_2, \ldots, X_n$ be independent Poisson trials such that, for $1 \le i \le n$, $\Pr[X_i = 1] = p_i$, where $0 < p_i < 1$. Then for $X = \sum_i X_i$, $\mu = E[X]$, and any $\delta > 0$,*

$$\Pr[X > (1+\delta)\mu] < \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^\mu < \left(\frac{e}{1+\delta}\right)^{(1+\delta)\mu}.$$

Let Y_u be the expected number of layers in which u is unpadded, and let $\mathcal{E}_u$ be the event that u is unpadded in more than $\ell/2$ layers. Then, applying the above lemma,

$$\Pr[\mathcal{E}_u] = \Pr\left[Y_u > \frac{r}{2c_1} \cdot \frac{c_1\ell}{r}\right] = \frac{1}{\Omega(r)^{\ell/2}}.$$

We will now show that each layer of an r-padded decomposition $P_1, P_2, \ldots, P_\ell$ can be embedded using only $O(\rho \log \rho)$ coordinates, i.e., we can find, for each $C \in P_i$, a relative embedding $\varphi_C : C \to \mathbb{Z}^{O(\log \rho)}$, such that, letting $\varphi_i(u) = \varphi_C(u)$ for $u \in C$ and $\varphi_i(u) = \mathbf{0}$ otherwise, and setting $\varphi = \oplus_{i=1}^{\ell}\varphi_i$, we have $\varphi(u) \ne \varphi(v)$ whenever $\sqrt{r} < d(u,v) \le r$. Using the nesting techniques of Theorem 3.1, the existence of an embedding into only $O(\rho \log \rho)$ dimensions will follow.

We first handle the base case, where $r = \rho$ (notice that this is where our decomposition breaks down). We will show how to produce a contraction $\varphi : V \to \mathbb{Z}^{O(\rho \log \rho)}$ which satisfies $\varphi(u) \ne \varphi(v)$ whenever $d(u,v) \le \rho$. The idea is simple: First, produce an ℓ-layer r-padded decomposition $P_1, \ldots, P_\ell$. Now for every cluster $C \in P_i$, form a relative embedding of C as follows: If $u \notin \partial C$, then let $\varphi_C(u)$ be a $\log \rho$-dimensional vector chosen uniformly at random from $\{0,1\}^{\log \rho}$ and let $\varphi_C(u) = \mathbf{0}$ otherwise. This is clearly a contraction. Now set $\varphi_i(u) = \varphi_C(u)$ where $C \in P_i$ is the cluster containing u, and then set $\varphi = \oplus_{i=1}^{\ell}\varphi_i$.

Consider two points u, v with $d(u,v) \le r$ and let P_i be a layer in which u is padded. In this layer, u and v occur in the same cluster C and $u \notin \partial C$, so $\Pr[\varphi_C(u) = \varphi_C(v)] \le (\frac{1}{2})^{\log \rho} = \frac{1}{\rho}$. Let $\mathcal{E}_{u,v}$ be the event that $\varphi(u) = \varphi(v)$, then it follows that $\Pr[\mathcal{E}_{u,v}] \le \Pr[\mathcal{E}_u] + (1/\rho)^{\ell/2} \le O(1/\rho)^{\ell/2}$. There are $\Omega(n)$ events $\mathcal{E}_{u,v}$ and we would like to argue that with positive probability, none of them occur. Again, the local lemma comes to our rescue. It is not difficult to see that $\mathcal{E}_{u,v}$ is independent of all events $\mathcal{E}_{u',v'}$ for which $d(u,u'), \ge 3\rho^4$. This is because the r-padded decomposition is a local probabilistic procedure and every ball appearing in its formation has radius at most ρ^4 (assuming without loss that ρ is not smaller than a suitable constant as otherwise we could increase it by a constant factor). Since the relative embedding technique is also local, we get independence between distant events. It follows that $\mathcal{E}_{u,v}$ is mutually independent of all but at most $d = O(\rho^{8\rho})$ other events (close enough pairs u', v'). Choosing c_0 to be a large enough constant, we see that $\Pr[\mathcal{E}_{u,v}] \le O(1/\rho)^{c\rho/2} \le e/(d+1)$. Thus, applying Lemma 5.2 yields an embedding for which none of the events $\mathcal{E}_{u,v}$ occur.

Now let us consider the layers $P_1, \ldots, P_\ell$ corresponding to an r-padded decomposition for some value $r \ge \rho$, and assume that we already found a relative embedding for the layers corresponding to say $r^{1/64}$. In Theorem 3.1, for each cluster $C \in P_i$, we applied an m-layer $r^{1/2}$-inner decomposition to C to produce the relative embedding. But now we know that with high probability, a vertex u is padded in at least half of the ℓ the layers, so we will only decompose C into one (randomly chosen) layer. Furthermore, the relative embedding we design will not use $O(\rho)$ coordinates but rather only one coordinate. Let us see that this works, i.e., with positive probability $\varphi(u \ne \varphi(v)$ whenever $\sqrt{r} < d(u,v) \le r$.

First, assume that u is padded in at least $\ell/2$ layers. (Recall that his happens with probability $1 - O(1/r)^{l/2}$.) Consider a layer P_i in which u is padded and let $C \in P_i$ be such that $u \in C$. To produce the one layer to which we apply the relative embedding technique of Lemma 3.1, use the randomized decomposition from the proof of Theorem 5.2 with parameter $r^{1/8} = (\sqrt{r})^{1/4}$. The probability that u is not padded in this layer, i.e. that $B(u, 3r^{1/8})$ is not contained in any cluster, is at most $O(r^{1/8})$ by the same proof as in Theorem 5.1. But there are at least $\ell/2$ of these events since u is padded in at least $\ell/2$ layers, so, applying a Chernoff bound, we see that u is $r^{1/8}$-padded in at least $\ell/4$ one-layer cluster decompositions with probability $1 - O(1/r^{\ell/4})$. So assume that this event happens, and then the probability that u and v collide in every one of them is $O(1/r^{1/8})^{\ell/4}$. By a union bound on the three bad events mentioned above we see that $\Pr[\varphi(u) = \varphi(v)] \le O(1/r^{\ell/32})$. Finally, we would like to apply Lemma 5.2 on the events $\mathcal{E}_{u,v} = \{\varphi(u) = \varphi(v)\}$ where $\sqrt{r} < d(u,v) \le r$. It can be seen that every event $\mathcal{E}_{u,v}$ is mutually independent of all the other events $\mathcal{E}_{u',v'}$ but the $r^{O(1)}$ events for which $d(u,u') \ge 3r^4$. Hence, for $c > 0$ a sufficiently large constant we can apply Lemma 5.2, which completes the proof of the theorem. $\square$

Algorithmic aspects. It is possible to make the applications of the local lemma in this section algorithmic. A proof is deferred to the full version.

6. EXPANDERS

THEOREM 6.1. *Let G be a $\log n$-degree expander, then $\dim(G) = \Omega(\rho_G \log \rho_G)$. In particular, Conjecture 1 is not true (for general graphs).*

PROOF. Let $G = (V, E)$ be an expander graph on n vertices with all vertices having roughly the same degree $\Theta(k)$, for $1 \le k \le \log n$. Specifically, the expansion properties that we need are:

(i) The diameter of G is $O(\log_k n)$.

(ii) Every two sets of $n/\log n$ vertices in G are connected by a path of length $O(\log_k \log n)$.

Observe that these two properties follow from standard vertex expansion. It follows from Property (i) that $\rho_G = \Theta(\frac{\log n}{\log \log_k n})$. We shall show that if G occurs as a subgraph of $\mathbb{Z}_\infty^d$ then $d = \Omega(\frac{\log n}{\log \log_k \log n})$. Note that for $k = \log n$, this implies that $\dim(G) = \Omega(\log n) = \Omega(\rho \log \rho)$ and this lower bound is tight, up to constant factors, since the trivial upper bound $d = O(\log n)$ holds for any n-vertex graph (by a bijection into $\{0,1\}^{\log n}$).

Assume for contradiction that G occurs as a subgraph of $\mathbb{Z}_\infty^d$ with $d = o(\frac{\log n}{\log \log_k \log n})$. Let φ be the corresponding embedding of G into $\mathbb{Z}_\infty^d$, and let φ_i be the projection of φ on the coordinate $i = 1, \ldots, d$. Let the set S_i consist of the $n/\log n$ vertices $v \in V$ with smallest $\varphi_i(v)$,and let the set L_i consist of the $n/\log n$ vertices $v \in V$ with largest $\varphi_i(v)$.

We claim that $\varphi_i(V \setminus (S_i \cup L_i))$ is an interval of size $O(\log_k \log n)$. Indeed, by property (ii) above G contains a path of length $O(\log_k \log n)$ that connects some vertex $s \in S_i$ with some vertex $l \in L_i$. Since φ is contractive, $\varphi_i(l) - \varphi_i(s) \leq O(\log_k \log n)$. By the definition of S_i and L_i, for every $v \in V \setminus (S_i \cup L_i)$ we have $\varphi_i(s) \leq \varphi_i(v) \leq \varphi_i(l)$, which proves the claim.

Finally, the set of vertices $V' = V \setminus (\cup_{i=1}^d (S_i \cup L_i))$ contains at least $n - dn/\log n > n/2$ vertices. By the above claim, $\varphi(V')$ is contained in a subset of the lattice $\mathbb{Z}^d$ which is the cartesian product of d intervals of size $O(\log_k \log n)$. However, this subset of $\mathbb{Z}^d$ contains at most $(O(\log_k \log n))^d \ll n/2$ points, which contradicts the assumption that φ is injective. $\square$

7. RELATED NOTIONS OF DIMENSIONALITY

THEOREM 7.1. *All the upper bounds for* $\dim(G)$ *hold also for* $\dim_2(G)$.

PROOF. Consider a contractive, injective embedding of $G = (V, E)$ into $\mathbb{Z}_\infty^d$ such that for every two distinct vertices u, v, their images $\varphi(u)$ and $\varphi(v)$ differ in at least $\frac{d}{4}$ coordinates. Our constructions can be easily modified to yield such an embedding by applying appropriate Chernoff bounds when the coordinates are formed (see the application of Lemma 5.3 in Section 5.3, for instance).

This embedding satisfies

1. $||\varphi(u) - \varphi(v)||_2 \geq \frac{1}{2}\sqrt{d}$ for $u \neq v \in V$ and
2. $||\varphi(u) - \varphi(v)||_2 \leq \sqrt{d}$ for $(u, v) \in E$.

Now scaling the value of every coordinate by $\frac{2}{\sqrt{d}}$ yields the desired mapping. $\square$

THEOREM 7.2. *For a* $\Theta(k)$*-degree expander* G *with* $1 \leq k \leq \log n$, $\dim_2(G) = \Omega(\frac{\log n}{\log \log_k \log n})$. *For a* $\log n$*-degree expander,* $\dim_2(G) = \Omega(\rho_G \log \rho_G)$.

PROOF. Similar to that of Theorem 6.1. $\square$

Acknowledgments. We are indebted to Nati Linial for his part in initiating this research. We are particularly grateful to him for valuable discussions in an early stage of the work and for his suggestion to focus on trees.

8. REFERENCES

[1] N. Alon and J. H. Spencer. *The probabilistic method.* Wiley-Interscience [John Wiley & Sons], New York, second edition, 2000.

[2] P. Assouad. Plongements lipschitziens dans $\mathbf{R}^n$. *Bull. Soc. Math. France*, 111(4):429–448, 1983.

[3] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *37th Annual Symposium on Foundations of Computer Science*, pages 184–193. IEEE, 1996.

[4] J. A. Bondy and U. S. R. Murty. *Graph theory with applications.* American Elsevier Publishing Co., Inc., New York, 1976.

[5] A. Brandstädt, V. Chepoi, and F. Dragan. Distance approximating trees for chordal and dually chordal graphs. *J. Algorithms*, 30(1):166–184, 1999.

[6] V. Chepoi and F. Dragan. A note on distance approximating trees in graphs. *European J. Combin.*, 21(6):761–766, 2000. Discrete metric spaces (Marseille, 1998).

[7] F. R. K. Chung. Labelings of graphs. In *Selected topics in graph theory, 3*, pages 151–168. Academic Press, San Diego, CA, 1988.

[8] P. Erdős, F. Harary, and W. T. Tutte. On the dimension of a graph. *Mathematika*, 12:118–122, 1965.

[9] U. Feige. Approximating the bandwidth via volume respecting embeddings. *J. Comput. System Sci.*, 60(3):510–539, 2000.

[10] J. Heinonen. *Lectures on analysis on metric spaces.* Universitext. Springer-Verlag, New York, 2001.

[11] K. Hildrum, J. D. Kubiatowicz, S. Rao, and B. Y. Zhao. Distributed object location in a dynamic network. In *Proceedings of the Fourteenth ACM Symposium on Parallel Algorithms and Architectures*, pages 41–52. August 2002.

[12] P. Indyk. Algorithmic applications of low-distortion geometric embeddings. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 10–33. October 2001.

[13] D. Karger and M. Ruhl. Finding nearest neighbors in growth-restricted metrics. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, pages 63–66, 2002.

[14] P. Klein, S. A. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *25th Annual ACM Symposium on Theory of Computing*, pages 682–690, May 1993.

[15] N. Linial. Variation on a theme of Levin. In J. Matoušek, editor, *Open Problems, Workshop on Discrete Metric Spaces and their Algorithmic Applications.* Haifa, March 2002.

[16] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.

[17] N. Linial and M. Saks. Low diameter graph decompositions. *Combinatorica*, 13(4):441–454, 1993.

[18] L. Lovász and K. Vesztergombi. Geometric representations of graphs. In *Paul Erdös, Proc. Conf.*, Budapest, 1999.

[19] J. Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, 2002.

[20] R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, 1995.

[21] S. Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proceedings of the 15th Annual Symposium on Computational Geometry*, pages 300–306. ACM, 1999.

[22] S. Semmes. Bilipschitz embeddings of metric spaces into Euclidean spaces. *Publ. Mat.*, 43(2):571–653, 1999.