Rolling backwards can move you forward: on embedding problems in sparse expanders

Nemanja Draganić * Michael Krivelevich † Rajko Nenadov ‡

Abstract

We develop a general embedding method based on the Friedman-Pippenger tree embedding technique (1987) and its algorithmic version, essentially due to Aggarwal et al. (1996), enhanced with a roll-back idea allowing to sequentially retrace previously performed embedding steps. This proves to be a powerful tool for embedding graphs of large girth into expander graphs. As an application of this method, we settle two problems:

- For a graph H, we denote by H^q the graph obtained from H by subdividing its edges with $q-1$ vertices each. We show that the k-size-Ramsey number $\hat{R}_k(H^q)$ satisfies $\hat{R}_k(H^q) = O(qn)$ for every bounded degree graph H on n vertices and for $q = \Omega(\log n)$, which is optimal up to a constant factor. This settles a conjecture of Pak (2002).
- We give a deterministic, polynomial time algorithm for finding vertex-disjoint paths between given pairs of vertices in a strong expander graph. More precisely, let G be an (n, d, λ)-graph with $\lambda = O(d^{1-\varepsilon})$, and let $\mathcal{P}$ be any collection of at most $c\frac{n \log d}{\log n}$ disjoint pairs of vertices in G for some small constant c, such that in the neighborhood of every vertex in G there are at most $d/4$ vertices from $\mathcal{P}$. Then there exists a polynomial time algorithm which finds vertex-disjoint paths between every pair in $\mathcal{P}$, and each path is of the same length $\ell = O\left(\frac{\log n}{\log d}\right)$. Both the number of pairs and the length of the paths are optimal up to a constant factor; the result answers the offline version of a question of Alon and Capalbo (2007).

1 Introduction

Given a graph H from some class of graphs, and a graph G with specific properties, is there a copy of H in G? In other words, does there exist an embedding of H into G? This general question is one of the central settings of combinatorics. Embedding questions lie at the heart of many classical problems, in particular problems in graph Ramsey theory and Turán-type extremal theory.

We will consider embedding problems where the host graph G is sparse, i.e. the number of edges in G is linear in its number of vertices. This is a natural and important setup both for theoretical and practical reasons, and its potential applicability ranges from problems in extremal combinatorics like Ramsey-type problems, to construction of lean but resilient networks in computer networking.

In particular, we will work with sparse expanders — those are sparse graphs in which all sets of vertices S of (up to) a certain size have a relatively large neighborhood. For a comprehensive source of information about expanders, see the survey of Hoory, Linial and Wigderson [34].

*Department of Mathematics, ETH, 8092 Zürich, Switzerland. Email: nemanja.draganic@math.ethz.ch.

†School of Mathematical Sciences, Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv 6997801, Israel. Email: krivelev@tauex.tau.ac.il. Supported in part by USA-Israel BSF grant 2018267, and by ISF grant 1261/17.

‡Department of Mathematics, ETH, 8092 Zürich, Switzerland. Email: rajko.nenadov@math.ethz.ch.

Closely related to expander graphs is the notion of pseudo-random graphs. Informally, a graph is pseudo-random if it behaves similarly to a random graph when it comes to edge distribution. A very popular class of examples of such graphs are (n, d, λ)-graphs, introduced by Alon. A d-regular graph G on n vertices is an (n, d, λ)-graph if all of the eigenvalues of its adjacency matrix, except the largest one, are at most λ in absolute value. One can show that the smaller λ is, the closer the graph resembles a random graph in terms of edge distribution (see Section 1.1.1 for some details). A small λ also means that the graph has good expansion properties and we will use a few such results throughout the paper. For a survey on pseudo-random graphs, see the paper of Krivelevich and Sudakov [42].

For our embedding problems, usually the host graph G will be an (n, d, λ)-graph, for sufficiently small λ and a constant d. What kind of subgraphs can we hope to find in such graphs? One natural restriction will be that the girth of the graph we embed is $\Omega(\log n)$, as there exist (n, d, λ)-graphs with small *spectral ratio* λ/d and of logarithmic girth, as shown in the seminal paper of Lubotzky, Phillips and Sarnak [44]. Thus we are normally confined to embedding trees and other graphs with large girth.

There is a large body of research devoted to finding (almost-spanning and spanning) bounded degree trees in sparse expanders and in sparse random graphs. Beck [10] used results about long paths in expanding graphs to argue that one can find monochromatic linear sized paths in 2-colored sparse random graphs. Friedman and Pippenger [28] proved an analogous statement for arbitrary bounded degree trees in sparse expanders, which was improved upon by Haxell, who showed that under similar assumptions one can embed even larger trees into (sparse) expanders. Alon, Krivelevich and Sudakov [5] proved the existence of every almost spanning tree of bounded degree in both sparse random graphs and in appropriate (n, d, λ)-graphs, later improved by Balogh, Csaba and Samotij and Pei [8], and for a resilience version of this result see [9]. Finally, for random graphs $G \sim G(n, p)$ with $p = \frac{C \log n}{n}$ and for a fixed d, Montgomery [47] recently proved that for large enough C, G typically contains all spanning trees of maximum degree at most d, resolving an old conjecture of Kahn. For results about finding small minors of logarithmic girth in sparse expanders, see, e.g. [46, 54].

In this paper, we will show two different results related to embedding into sparse expanders — the first one deals with size-Ramsey numbers of logarithmic subdivisions of bounded degree graphs and resolves a conjecture of Pak from 2002 [48], while the second is concerned with the classical problem of finding vertex-disjoint paths in graphs, and solves the offline version of the problem of Alon and Capalbo from 2007 [3]. For each of those problems we develop a separate variation of our embedding technique. Both are based on the result of Friedman and Pippenger [28] about embedding trees in expander graphs vertex by vertex and an idea by Daniel Johannsen [36], which allows us to successively remove vertices from the list of already embedded vertices. This *roll-back* result turns out to be very powerful for tackling problems of this sort. One of the variants which we show is algorithmic, and uses ideas by Dellamonica and Kohayakawa [20], who showed an algorithmic version of the original Friedman-Pippenger embedding result, by reducing it to a certain online matching problem solved by Aggarwal et al. [1].

1.1 Size-Ramsey numbers of subdivided graphs

Given a graph H and an integer $k \geq 2$, a graph G is said to be *k-Ramsey* for H if every coloring of the edges of G with k colors contains a monochromatic copy of H. This notion was introduced by Ramsey [50], who proved that for every graph H there exists $N \in \mathbb{N}$ such that K_N is k-Ramsey for H. The smallest such N, denoted by $R_k(H)$, is called the *Ramsey number*. Determining the asymptotic order of $R_2(K_\ell)$ is one of the most important open problems in this area [16, 55]. We will be concerned with the related notion of *size-Ramsey numbers*, introduced by Erdős, Faudree, Rousseau and Schelp [25]. Given a graph H and an integer $k \geq 2$, the *size-Ramsey* *number* $\hat{R}_k(H)$ is the smallest integer m such that there exists a graph G with m edges which is k-Ramsey for H. The existence of the Ramsey number immediately implies the upper bound $\hat{R}_k(H) \le \binom{R_k(H)}{2}$. Other related notions include Folkman numbers, chromatic-Ramsey numbers, degree-Ramsey numbers, etc. We refer the reader to a recent survey by Conlon, Fox and Sudakov [17] for a thorough treatment of the topic.

Answering a 100\$–question of Erdős [24], Beck [10] showed that paths have linear size-Ramsey number, that is $\hat{R}_2(P_n) \le Cn$ for an absolute constant C. He also raised the question [11] of whether $\hat{R}_2(H)$ grows linearly for graphs with bounded maximum degree. This was proven for trees by Friedman and Pippenger [28] and for cycles by Haxell, Kohayakawa and Łuczak [32]. However, the general case was settled in the negative by Rödl and Szemerédi [51], who showed that there exists a constant $c > 0$ such that for every sufficiently large n there is a graph H with n vertices and maximum degree 3 for which $\hat{R}_2(H) \ge n \log^c n$. In the same paper, they conjectured that $\log^c n$ can be improved to n^ε for some constant $\varepsilon > 0$, but this remains open. For further results about size-Ramsey numbers, see for example [7, 12, 15, 19, 23, 31, 35, 37, 41, 43].

1.1.1 Subdivisions of graphs

Since we are far from understanding size-Ramsey numbers of bounded degree graphs in general, one natural step in this direction is to consider subdivisions of those graphs. Given a graph H and a function $\sigma\colon E(H) \to \mathbb{N}$, the *$\sigma$-subdivision* H^σ of H is the graph obtained from H by replacing each edge $e \in E(H)$ with a path of length $\sigma(e)$ joining the endpoints of e, such that all these paths are mutually vertex-disjoint (except possibly at the endpoints). In other words, we *subdivide* each edge $\sigma(e) - 1$ times.

Size-Ramsey numbers of 'short' subdivisions were first studied by Kohayakawa, Retter and Rödl [38]. In a recent paper [22] we improved their bounds by showing that $\hat{R}_k(H^q) \le O(n^{1+1/q})$, for constant q, k and for all bounded degree graphs H, thus removing a polylogarithmic factor from their bound and answering their question. In general, these graphs were considered in the context of Ramsey theory by Burr and Erdős [14] and by Alon [2].

In Section 3.1 (Theorem 1.2) we show that bounded degree graphs with n vertices such that every two vertices of degree ≥ 3 are at distance $q = \Omega(\log n)$, have linear size-Ramsey numbers (in their order). In fact we prove a stronger result on arbitrary long *subdivisions* of bounded degree graphs, answering a conjecture of Pak [48] along the way. He conjectured that long subdivisions of bounded degree graphs have linear size-Ramsey number.

Conjecture 1.1 ([48])**.** *For every $k, D \in \mathbb{N}$ there exist $C, L > 0$ such that if H is a graph with $\Delta(H) \le D$ and $\sigma(e) = \ell \ge L \log(v(H^\sigma))$ for all $e \in E(H)$ then $\hat{R}_k(H^\sigma) \le Cv(H^\sigma)$.*

Pak [48] showed that $\hat{R}_k(H^\sigma) = O(v(H^\sigma)\log^3(v(H^\sigma)))$ and the special case where H is a fixed (small) graph and $\sigma(e)$ grows was resolved by Donadelli, Haxell and Kohayakawa [21].

We show that every *η-uniform* graph on n vertices is k-Ramsey for H^σ with $v(H^\sigma) \le \alpha n$ and $\sigma(e) \ge \log n$, for some small $\alpha > 0$. As a typical random graph with n vertices and $m = Cn$ edges is η-uniform, for sufficiently large C, there are an abundance of η-uniform graphs with $O(n)$ edges, thus confirming Conjecture 1.1.

Definition. Given $0 < \eta \le 1$, we say that a graph G with n vertices and density $p = e(G)/\binom{n}{2}$ is *η-uniform* if for every disjoint subsets $U, W \subseteq V(G)$ of size $|U|, |W| \ge \eta n$, we have

$$\big|e(U,W) - |U||W|p\big| \le \eta|U||W|p.$$

Now we state our result.

Theorem 1.2. *For every $k, D \in \mathbb{N}$ and for every $\delta > 0$, there exist $\eta, \alpha, C > 0$, such that the following holds for every η-uniform graph G with n vertices and $m \geq Cn$ edges: every k-edge-coloring of G contains a monochromatic copy of every graph H^σ, where H is a graph with maximum degree at most D, $v(H^\sigma) \leq \alpha n$ and $\sigma(e) \geq \delta \log n$ for every $e \in E(H)$.*

Besides random graphs, explicit constructions of η-uniform graphs of constant average degree are also known. One class of examples of such graphs are (n, d, λ)-graphs, for suitably chosen parameters. Indeed, the well known Expander Mixing Lemma [4] states that for every (n, d, λ)-graph G and for every $U, V \subseteq V(G)$ it holds:

$$\left| e_G(U, V) - \frac{d|U||V|}{n} \right| \leq \lambda \sqrt{|U||V|}. \tag{1}$$

From this, one can see that every (n, d, λ)-graph is η-uniform for $\eta = 2\sqrt{\frac{\lambda}{d}}$. Hence, for a fixed d, the parameter λ is accountable for the uniformity of the distribution of the edges of a d-regular graph. But how small can λ be in terms of d, so that there exists a (n, d, λ)-graph? One can show that $\lambda = \Omega(\sqrt{d})$ for every such graph whenever $d < 0.99n$, and there are known constructions of d-regular graphs for which λ attains this bound, and n is arbitrarily large. This provides us examples of bounded degree graphs, which are η-uniform (for $\eta \sim d^{-1/4}$). For several constructions of such graphs, see, e.g., [42]. As discussed above, there are known constructions of such graphs which have logarithmic girth, showing that our result is asymptotically tight with respect to the bound on σ. Indeed, if H is a triangle, and G is a graph with girth strictly larger than $c \log n$, then G does not contain H^σ for any σ bounded from above by $c \log n/3$.

Note that Theorem 1.2 is in fact a *universality* result, meaning that the η-uniform graph in question is k-Ramsey for *all* graphs in the class we are interested in, hence our theorem confirms Pak's conjecture in a strong way. Furthermore, from our proof it can be seen that we actually find a monochromatic subgraph of the graph we color, which contains all described subdivided graphs. Extending the definition in [39], we say that a graph G is *k-partition universal* for a class of graphs $\mathcal{F}$ if for every k-coloring of the edges of G, there exists a monochromatic subgraph of G which contains a copy of every graph in $\mathcal{F}$. Under this framework, we actually prove that the graph we color is up to a constant factor the optimal k-partition universal graph for the class of all described subdivisions of graphs. For further universality-type results in Ramsey theory see for example [18, 22, 38, 39].

1.2 The vertex-disjoint paths problem

For a given graph G and a collection of k disjoint pairs of vertices (a_i, b_i) from G, can we find for each i a path from a_i to b_i, such that the found paths are all vertex-disjoint? This decision problem is $\mathcal{NP}$-complete [30] when G is allowed to be an arbitrary graph. Furthermore, it remains $\mathcal{NP}$-complete, even when G is restricted to be in the class of planar graphs. For fixed k, it is shown to be in $\mathcal{P}$ [52]. A variant of this problem in random graphs was studied indepedently by Hochbaum [33], and by Shamir and Upfal [53]. Both papers proved that for a fixed set of at most $O(\sqrt{n})$ disjoint pairs of vertices in the random graph $G(n, m)$, with high probability (whp) there exist vertex-disjoint paths between every pair if $m > Cn \log n$, for a constant $C > 1$. Subsequently, Broder, Frieze, Suen and Upfal [13] improved this result:

Theorem ([13])**.** *There exist $\alpha, \beta > 0$, such that whp the following holds. Let $G = G(n, m)$ for $m = (\log n + \omega(1))\frac{n}{2}$, and let $d = 2m/n$. For every collection $\mathcal{F}$ of at most $\alpha \frac{n \log d}{\log n}$ disjoint pairs of vertices (a_i, b_i) in G, there exists a path for every i connecting a_i to b_i, such that all paths are vertex-disjoint, if the following condition is satisfied:*

$$|N_G(v) \cap (A \cup B)| < \beta d_G(v), \quad \textit{where } A = \cup_i \{a_i\} \textit{ and } B = \cup_i \{b_i\}.$$

This is an improvement over the mentioned previous papers in many aspects. The number of pairs is optimal up to a constant factor — most pairs in $G(n,m)$ are at distance $\Omega(\log n/\log d)$, so in general one can hope to connect at most $O(\frac{n\log d}{\log n})$ pairs (a_i, b_i) in a graph on n vertices. Furthermore, the pairs (a_i, b_i) are not fixed before generating $G(n,m)$, but are rather chosen adversarily after having exposed a random graph $G \sim G(n,m)$. The last constraint is also optimal up to a constant factor — if the adversary chooses a_1 and b_1 to be at distance 2, and then chooses all neighbours of a_1 to be in other pairs from $\mathcal{F}$, then obviously one cannot find the requested disjoint paths. The bound on the number of edges m is also asymptotically optimal, as this many edges are needed for G to be connected whp.

Changing the focus to the sparse(r) setting, Alon and Capalbo [3] studied graphs with constant average degree with good expansion properties. In particular, they proved that for any graph G which is a d-blowup of a (n,d,λ)-graph with a small spectral ratio, and any collection of $O\left(\frac{nd\log d}{\log n}\right)$ pairs of vertices in G which satisfy a similar local condition like in [13], one can connect those pairs with vertex-disjoint paths. The number of pairs is optimal up to a constant factor, and they provide a polynomial time algorithm for finding them.

The argument of Alon and Capalbo does not allow to control the length of the paths found by the algorithm. Accordingly, they ask for a similar result where the length of the paths between each pair is at most $O(\log n)$. In Section 3.2, we prove such a result (Theorem 1.3). Furthermore, we do it not only for blowups, but directly for (n,d,λ)-graphs for $\lambda < d^{0.99}$. We get the optimal dependency on n and d, both for the number of pairs and for the upper bound on the length of the paths.

Theorem 1.3. *Let $\varepsilon > 0$, and let G be an (N,D,λ)-graph, with $\lambda < D^{1-\varepsilon}/150$ and $D^\varepsilon > 25$. Let $P = \{a_i, b_i\}$ be a collection of at most $\frac{\varepsilon N\log D}{160\log N}$ disjoint pairs of vertices in G, such that $|N_G(x) \cap (A \cup B)| \leq \frac{D}{4}$ for every $x \in V(G)$, where $A = \cup_i\{a_i\}, B = \cup_i\{b_i\}$. There exists a polynomial time algorithm to find vertex-disjoint paths in G between every pair of vertices $\{a_i, b_i\}$, such that the paths are of equal length which is less than $5\frac{\log N}{\varepsilon\log D}$.*

These results are closely related to the study of non-blocking networks, which arise in a variety of applications, including construction of communication networks and distributed-memory architectures. For some results see, e.g., [26, 27, 49]. In contrast to our results, the graphs which are usually considered here have pre-determined sets of vertices ("inputs" and "outputs") from which the pairs are chosen, while the pairs in our result can be chosen by an adversary, but in such a way that they satisfy an essentially minimal local property. Besides that, the path lengths in some constructions of non-blocking networks are also of optimal $O(\log n)$ size [6]. Hence, in some sense our results are a common generalization of [6] and [3], as we both allow the adversary to choose the pairs, and our paths are logarithmic in size, although our algorithm is less efficient than the one in [6], and is not online in the same sense like in [3].

A lot of attention has also been paid to the edge-disjoint paths problem. For a short survey, see [29], and for a more recent result on edge-disjoint paths in sufficiently strong expander graphs see [3].

1.3 Outline of the paper and notation

In Section 2 we show two versions of our main embedding technique — in Section 2.1 we show the non-algorithmic version of it, and in Section 2.2 we give an algorithmic version of the technique. In Section 3, we prove our main results — Theorem 1.2 (the resolution of Pak's conjecture) in Section 3.1, and Theorem 1.3 (vertex-disjoint paths in (N,D,λ)-graphs) in Section 3.2. In section 4 we give some concluding remarks.

Notation. We follow standard graph theoretic notation. In particular, given a graph G and a vertex $x \in V(G)$, we denote by $N_G(x)$ the neighborhood of x in G. Similarly, for a subset of vertices $X \subseteq V(G)$ we denote by $N_G(X)$ the external neighborhood of X, that is $N_G(X) = (\bigcup_{x \in X} N_G(x)) \setminus X$. By $\partial_G(x)$ we denote the set of edges incident with vertex x in G. Given disjoint subsets of vertices $A, B \subseteq V(G)$, we denote by $e_G(A, B)$ the number of edges with one endpoint in A and the other in B, and with $d_G(A, B) = e_G(A, B)/|A||B|$ the density of such a induced bipartite graph. We denote by $v(G)$ the number of vertices of G, and by $e(G)$ the number of edges of G. Given graphs G and H, we say that a mapping $\phi\colon V(H) \to V(G)$ is an *embedding*, with the notation $\phi\colon H \hookrightarrow G$, if it is injective and preserves edges of H (i.e. if $\{v, w\} \in E(H)$ then $\{\phi(v), \phi(w)\} \in E(G)$). For an embedding $\phi\colon H \hookrightarrow G$ and subsets $S_1 \subseteq V(H), S_2 \subseteq V(G)$ we denote by $\phi(S_1)$ the image of S_1 under ϕ, and by $\phi^{-1}(S_2)$ the preimage of S_2 under ϕ, i.e. $\phi(S_1) = \{y \in V(G) \mid \exists x \in S_1\colon \phi(x) = y\}$, and $\phi^{-1}(S_2) = \{x \in V(H) \mid \phi(x) \in S_2\}$. We omit floors and ceilings whenever it is not crucial. Given two constant ε and α, we use somewhat informal notation $\varepsilon \ll \alpha$ to denote that ε is sufficiently small compared to α. We denote by $\log n$ the natural logarithm of n.

2 Friedman-Pippenger type embedding theorems

Now we describe the main embedding machinery behind our proofs. It relies on the idea of Friedman and Pippenger, used for embedding trees in expanders vertex by vertex, by maintaining a certain invariant. An algorithmic version of this technique was presented by Dellamonica and Kohayakawa, based on a result about an online matching game by Aggarwal et al. [1]. In the following two subsections, we give two Friedman-Pippenger type embedding theorems, non-algorithmic and algorithmic, enhanced with a roll-back idea, which allows us to sequentially retrace previously performed embedding steps. While the algorithmic result requires the host graph to have stronger expansion properties, it also enables us to embed larger graphs than with the technique described in Section 2.1.

2.1 The original Friedman-Pippenger theorem with rollbacks

We start with a standard definition of expansion.

Definition 2.1. Let $s \in \mathbb{N}$ and $K > 0$. We say that a graph G is (s, K)-*expanding* if for every subset $X \subseteq V(G)$ of size $|X| \leq s$ we have $|N_G(X)| \geq K|X|$.

In order to develop our machinery, we define the notion of an (s, D)-*good* embedding.

Definition 2.2. Let G be a graph and let $s, D \in \mathbb{N}$. Given a graph F with maximum degree at most D, we say that an embedding $\phi\colon F \hookrightarrow G$ is (s, D)-*good* if

$$|N_G(X) \setminus \phi(F)| \geq \sum_{v \in X} \left[D - \deg_F(\phi^{-1}(v))\right] \tag{2}$$

for every $X \subseteq V(G)$ of size $|X| \leq s$. Here we slightly abuse the notation by setting $\deg_F(\emptyset) := 0$, i.e. if a vertex $v \in V(G)$ is not used by ϕ to embed F, then we set $deg_F(\phi^{-1}(v)) = 0$.

We remark that the notion of a good embedding is the same as the one used by Friedman and Pippenger [28]. The following is implicit in [28].

Theorem 2.3. *Let F be a graph with $\Delta(F) \leq D$ and $v(F) < s$, for some $D, s \in \mathbb{N}$. Suppose we are given a $(2s, D{+}1)$-expanding graph G and a $(2s, D)$-good embedding $\phi\colon F \hookrightarrow G$. Then for every graph F' with $v(F') \leq s$ and $\Delta(F') \leq D$ which can be obtained from F by successively adding a new vertex of degree 1, there exists a $(2s, D)$-good embedding $\phi'\colon F' \hookrightarrow G$ which extends ϕ.*

The second result we need is a simple corollary of the definition of (m, D)-goodness. While easy to prove, this observation [36] turns out to yield a powerful method for connecting vertices in expanding graphs. It has also been utilized in the recent paper by Montgomery [47] for embedding spanning trees in random graphs.

Lemma 2.4. *Suppose we are given graphs G and F and an (s, D)-good embedding $\phi\colon F \hookrightarrow G$, for some $s, D \in \mathbb{N}$. Then for every graph F' obtained from F by successively removing a vertex of degree 1, the restriction ϕ' of ϕ to F' is also (s, D)-good.*

Proof. We show that the statement holds for the case where F' is obtained from F by removing a single vertex $v \in V(F)$ of degree 1. The lemma then follows by iterating it.

Let ϕ' be a restriction of F to such F', and let $w \in F'$ denote the unique neighbor of v. If $X \subseteq V(G)$ does not contain $\phi(w)$ then the right hand side of (2) does not change. Otherwise (if $\phi(w) \in X$) the right hand side of (2) increases by 1 (as the degree of w in F' is one less than it was in F). However, as $\phi(v)$ is no longer occupied (i.e. $\phi(v) \notin \phi'(F')$) and $\phi(v) \in N_G(\phi(w))$, the left hand side also increases by one, hence the inequality again holds. □

2.2 Algorithmic Friedman-Pippenger with roll-backs

In this section we prove an algorithmic version of the embedding technique provided by Theorem 2.3 and Lemma 2.4 from Section 2.1. We start with a description of an online matching game, to which we reduce our embedding problem.

Let $m \geq 0$ be an integer. The game is played on a bipartite graph $H = (U \cup V, E)$. In the beginning we set M (*the current matching*) to be empty. At each step an adversary chooses a vertex $x \in U$ which is not covered by M, and we match it to some free vertex in V to extend M. After each step the adversary is allowed to remove any number of edges from the current matching M, but at most m times in total during the game. In [1, Lemma 2.2.7], Aggarwal et al. describe a polynomial time algorithm which finds a matching of size n, against any adversary, if H satisfies the property that for each $X \subset U$ of size $|X| \leq n$, even if we remove at most half of the edges incident to every vertex in X, there are still at least $2|X|$ neighbors of X in the obtained graph.

Theorem 2.5 ([1], Aggarwal et al.). *Let $H = (U \cup V, E)$ be a bipartite graph and let $n, m \in \mathbb{N}$, such that for every $X \subseteq U$ of size $|X| \leq n$ and for every $F \subseteq E$ such that $|F \cap \partial_H(x)| \leq d_H(x)/2$ for every $x \in X$, we have that $|N_{H-F}(X)| \geq 2|X|$. Then there is an algorithm which finds a matching of size n against any adversary, if the adversary is allowed to remove edges from the matching at most m times in total during the game. Furthermore, the number of operations which the algorithm performs is polynomial in $m + |V(H)|$.*

Definition. We say that a graph $G = (V, E)$ has *property $P_\alpha(n, d)$* if for every $X \subseteq V$ of size $|X| \leq n$ and every $F \subseteq E$ such that $|F \cap \partial_G(x)| \leq \alpha \cdot d_G(x)$ for every $x \in X$, we have $|N_{G-F}(X)| \geq 2d|X|$.

Definition 2.6. Given a graph G, a subset of vertices $P \subseteq V(G)$, and natural numbers $n, m, d \in \mathbb{N}$, we define the following online game, which we call *the (G, P, n, m, d)-forest building game.* At each step there is a forest $T \subseteq G$ (initially $T := (P, \emptyset)$) with less than n edges in G, and the adversary *requests* a vertex $v \in T$ such that $d_T(v) < d$ and we are supposed to find a neighbor of v in $V(G) - V(T)$, hence extending T by a new leaf. The adversary is is allowed to successively remove any number of vertices of degree 1 in T after every step, but he is allowed to do so at most m times in total, and none of the removed vertices are allowed to be in P. We win if at some point T has n edges.

The next theorem gives a handy tool for embedding forests algorithmically in a robust way. In comparison to the technique presented in Section 2.1, here we require a stronger notion of expansion (the $P_\alpha(n,d)$-property) for the host graph, but the graphs we are embedding can have more vertices than before. The idea of the proof is similar to the one in [20].

Theorem 2.7. *Let $\alpha, \beta > 0$ with $\alpha - \beta \geq 1/2$ and let G be a graph with property $P_\alpha(n,d)$. Let P be a non-empty subset of vertices $P \subseteq V(G)$, such that for every vertex $x \in V(G)$ it holds that $|N_G(x) \cap P| \leq \beta \cdot d_G(x)$. Then there is an algorithm which wins the (G, P, dn, m, d)-forest building game after performing a number of operations polynomial in $m + |V(G)|$.*

Proof. In order to use Theorem 2.5, we construct the following auxiliary graph. Let H be a bipartite graph with classes $U = V(G) \times [d]$ and $V = \{\bar{v} \mid v \in V(G) - P\}$. In other words, U consists of d copies of $V(G)$, and V is a copy of $V(G) - P$. Two vertices $(u, j) \in U$ and $\bar{v} \in V$ are adjacent iff $\{u, v\}$ is an edge in G. Now we show that H satisfies the condition of Theorem 2.5 (with dn instead of n).

Let $X \subseteq U$ be of size $|X| \leq dn$, and $F \subseteq E(H)$ be such that $|F \cap \partial_H(x)| \leq d_H(x)/2$ for every $x \in X$. We want to show that $|N_{H-F}(X)| \geq 2|X|$. By the pigeonhole principle, one of the d copies of $V(G)$ in U contains at least $|X|/d$ elements from X, or in other words, there is an $i \in [d]$ such that the set $X_i := \{(u,i) \mid (u,i) \in X\}$ is of size $|X_i| \geq |X|/d$. Let Y be an arbitrary subset of X_i of size exactly $\lceil |X|/d \rceil$, and let $Y' = \{u \mid (u,i) \in Y\} \subseteq V(G)$.

We also define $F' \subseteq E(G)$ as follows:

$$F' = \Big\{\{u, v\} \in E(G) \mid u \in Y',\, v \notin Y', \text{ and } \{(u,i), \bar{v}\} \in F\Big\}.$$

Let G' be the graph obtained from G by deleting all edges in F' and by deleting all edges which have one vertex in Y' and the other in $P \setminus Y'$. Note the following facts:

(i) $|N_{H-F}(Y)| \geq |N_{G'}(Y')|$,

(ii) $d_{G'}(x) \geq (1-\alpha) d_G(x)$ for all $x \in Y'$.

The first claim is true as for every vertex $v \in N_{G'}(Y')$ there is a vertex $(u,i) \in Y$ such that $\{(u,i), \bar{v}\}$ is an edge in $H - F$. For the second claim, notice that after deleting F' from G each vertex $x \in Y'$ loses at most half of its edges, and by deleting the edges incident to $P \setminus Y'$, it loses at most another $\beta \cdot d_G(x)$ edges, which is in total at most $(1/2 + \beta) d_G(x) \leq \alpha \cdot d_G(x)$ edges.

It follows from the second claim and from $|Y'| = \lceil |X|/d \rceil \leq \lceil nd/d \rceil = n$ (and from the assumption that G has the $P_\alpha(n,d)$-property), that $|N_{G'}(Y')| \geq 2d|Y'| \geq 2|X|$. Together with ($i$) this implies $|N_{H-F}(X)| \geq |N_{H-F}(Y)| \geq 2|X|$.

Now we reduce our forest building game on the graph G to the matching game on the graph H. At the beginning our initial forest is set to be the empty graph on P, i.e. $T := (P, \emptyset) \subseteq G$. We also set our auxiliary matching M in H to be empty in the beginning. During the game M and T will have the same number of edges. In each step the adversary requests a vertex $u \in T$ such that $d_T(u) < d$, and we want to find a vertex v in $N_G(u) \setminus V(T)$ which extends T, such that $\{u, v\}$ is a new leaf in T. In order to do this, we find a vertex (u, j) in H for some $j \leq d$, which is not covered by M (in the next paragraph we show that such a vertex exists), and we extend M by finding a match $\bar{v} \in V$ for (u, j), using the algorithm from Theorem 2.5. Now we add the edge (u, v) (which is in G by the definition of H) to T. Note also that v was not in T before, as v certainly is not in P (by the definition of $\bar{v}$), and for every other vertex $x \in T$, the vertex $\bar{x}$ is covered by M, as x has been added to T by the same procedure, so $\bar{x} \neq \bar{v}$.

When the adversary wants to delete an edge (u, v) (where v is of degree 1 in T) from T, then we also delete the corresponding edge $\{(u, j), \bar{v}\}$ from M. Note that if at any step the adversary requests a vertex u such that $d_T(u) < d$, then a vertex of the form (u, i) (for some $i \in [d]$) has been used only at most $d-1$ times by the current matching M, so it is valid to assume that in each step we can find such a vertex which is not covered by M. Since the algorithm finds a forest T with dn edges at the same point when M contains dn edges, and we remove edges from the matching only at most m times in total, thanks to Theorem 2.5, we are done. □

3 Applications

3.1 Size-Ramsey number of long subdivisions

Before we start with the proof of Theorem 1.2, we state a few preliminary results which will help us find a subgraph with good expansion properties in the edge colored graph in question.

3.1.1 Preliminaries

The following lemma tells us that if in a graph all sets of a specified size expand well, we can delete relatively few vertices, so that in the remaining graph all smaller sets also expand well. For related results see for example [40]. A similar statement also appeared in [45].

Lemma 3.1. *Let G be a graph such that $|N_G(X)| \geq 2Ks$ for every subset $X \subseteq V(G)$ of size $|X| = s$, for some $s, K \in \mathbb{N}$. Then there exists a subset $B \subseteq V(G)$ of size $|B| < s$ such that $G - B$ is (s, K)-expanding.*

Proof. Let $B \subset V(G)$ be a largest set such that $|N_G(B)| < K|B|$ and $|B| < s$ (or $B = \emptyset$ if no such set exists). We show that $H = G - B$ is (s, K) expanding. Let $X \subseteq V(H)$ be an arbitrary non-empty set of size $|X| \leq s$ and suppose $|N_H(X)| < K|X|$. Then $|N_G(X \cup B)| < K|X| + K|B| = K|X \cup B|$, so by assumption we have $|X \cup B| \geq s$. Therefore, we conclude:

$$|N_H(X)| \geq |N_G(X \cup B)| - |N_G(B)| \geq 2Ks - Ks \geq Ks \geq K|X|$$

which contradicts the assumption on X, so we are done. □

Regular pairs

The proof of Theorem 1.2 combines results from Section 2.1 with a sparse version of Szemerédi's regularity lemma for multicolored graphs (or rather its corollary given shortly).

Definition. Given a graph G and disjoint subsets $U, W \subseteq V(G)$, we say that the pair (U, W) is *(G, ε, p)-regular* for some $\varepsilon, p \in (0, 1)$ if

$$|d_G(U', W') - d_G(U, W)| \leq \varepsilon p$$

for every $U' \subseteq U$ of size $|U'| \geq \varepsilon|U|$, and $W' \subseteq W$ of size $|W'| \geq \varepsilon|W|$.

Remark 3.2. *If $U' \subseteq U$ and $W' \subseteq W$ are as above and $d_G(U, W) > \varepsilon p$, then there exists at least one edge between U' and W' in G, as otherwise $d_G(U', W') = 0$, which contradicts $|d_G(U', W') - d_G(U, W)| \leq \varepsilon p$. It follows that $|N_G(U')| > (1-\varepsilon)|W|$.*

The following corollary of Szemerédi's regularity lemma was proven in [32, Lemma 3.4].

Lemma 3.3. *For every $k \geq 2$ and $0 < \varepsilon < 1$, there exist $\mu, \eta > 0$ such that the following holds: Suppose $G = (V, E)$ is an η-uniform graph with n vertices and density $p = e(G)/\binom{n}{2} > 0$, and let $E = E_1 \cup E_2 \cup \ldots \cup E_k$ be an k-edge-coloring of G. Then, for some $1 \leq z \leq k$, there exist pairwise disjoint subsets $V_1, V_2, V_3 \subseteq V$ of size $|V_i| = \mu n$ such that*

(a) (V_i, V_j) is (G_z, ε, p)-regular, where $G_z = (V, E_z)$, and

(b) $d_{G_z}(V_i, V_j) \geq p|V_i||V_j|/2k$,

for every $1 \leq i < j \leq 3$.

We are ready to prove Theorem 1.2, which we restate here.

Theorem 1.2. *For every* $k, D \in \mathbb{N}$ *and for every* $\delta > 0$, *there exist* $\eta, \alpha, C > 0$, *such that the following holds for every* η*-uniform graph* G *with* n *vertices and* $m \geq Cn$ *edges: every* k*-edge-coloring of* G *contains a monochromatic copy of every graph* H^σ, *where* H *is a graph with maximum degree at most* D, $v(H^\sigma) \leq \alpha n$ *and* $\sigma(e) \geq \delta \log n$ *for every* $e \in E(H)$.

3.1.2 Proof of Theorem 1.2 — resolution of Pak's conjecture

Proof of Theorem 1.2. Let $\mu = \mu(k, \varepsilon)$ and $\eta = \eta(k, \varepsilon) > 0$ be given by Lemma 3.3 for a sufficiently small constant $\varepsilon \ll D^{-1}, k^{-1}$. Also assume w.l.o.g. that $D \gg 1/\delta$. Suppose we are given an η-uniform graph G with n vertices and a k-edge-coloring $E(G) = E_1 \uplus E_2 \uplus \ldots \uplus E_k$, and let $1 \leq z \leq k$ and $V_1, V_2, V_3 \subseteq V(G)$ be obtained by applying Lemma 3.3. In the rest of the proof we show that $\Gamma = (V(G), E_z)$ contains H^σ for every H satisfying conditions of the theorem with $\alpha = \varepsilon\mu$.

Prepare Γ. Let $t = |V_i| = \mu n$. Let $\Gamma' = \Gamma[V_1, V_2]$ be a bipartite subgraph of Γ induced by V_1 and V_2. From (Γ, ε, p)-regularity of (V_1, V_2) and from the assumption $\varepsilon \ll 1/k, 1/D$, we conclude (Remark 3.2) that for every subset $X \subseteq V(\Gamma')$ of size $|X| = 2s$, where

$$s = 2D^2 \varepsilon t,$$

we have

$$|N_{\Gamma'}(X)| \geq t - \varepsilon t - |X| \geq t/2 \geq 2(D+3)|X|.$$

Therefore, by Lemma 3.1 there exists a subset $B \subseteq V(\Gamma')$ of size $|B| = s$ such that $\Gamma_B = \Gamma' \setminus B$ is $(2s, D+3)$-expanding. Let $V_1' = V_1 \setminus B$ and $V_2' = V_2 \setminus B$, so that $\Gamma_B = \Gamma_B[V_1', V_2']$. Most of H^σ will be embedded using Γ_B and the machinery from Section 2.1, with occasional help from set V_3.

Embed H. Consider a graph H with maximum degree D and let $\sigma\colon E(H) \to \mathbb{N}$ be a function such that $v(H^\sigma) = v(H) + \sum_{e \in E(H)} \sigma(e) < \varepsilon t$ and $\sigma(e) \geq \delta \log n$ for every $e \in E(H)$. Let $(e_1, \ldots, e_m)$ be an arbitrary ordering of the edges of H, and for each $0 \leq i \leq m$ set $H_i = (V(H), \{e_1, \ldots, e_i\})$. Note that H_0 is just an empty graph on the vertex set $V(H)$. We inductively show that for each $0 \leq i \leq m$ there exists an embedding $\phi_i\colon H_i^\sigma \hookrightarrow \Gamma$ such that the following holds:

(1) $\phi_i(V(H)) \subseteq V_1'$, and

(2) the restriction of ϕ_i to $F_i = \phi_i^{-1}(V(\Gamma_B))$, denoted by $f_i\colon F_i \hookrightarrow \Gamma_B$, is $(2s, D)$-good.

Let us first prove the base case $i = 0$. Note that $H_0^\sigma = H_0$ is an empty graph on the vertex set $V(H)$. Let a be a vertex (some new auxiliary vertex not used before) and $v \in V_1$, and set $\phi_0'(a) = v$. As Γ_B is $(2s, D+3)$-expanding, it is easy to see that ϕ_0' is a $(2s, D+2)$-good embedding of a graph consisting of a single vertex. Let us extend such a one-vertex graph to a path P of length $2\varepsilon t$. By Theorem 2.3, there exists an $(2s, D+2)$-good embedding $\phi_0'\colon P \hookrightarrow \Gamma_B$. Consider an arbitrary bijection between $V(H)$ and the set of *odd* vertices in P (i.e. the first vertex, third vertex, etc.). As $\phi_0'(a)$ is mapped into V_1', all these vertices are also necessarily mapped into V_1'. Together with ϕ_0', such a bijection gives an embedding $\phi\colon H_0 \hookrightarrow \Gamma_B$ with $\phi(V(H)) \subseteq V_1'$. As ϕ_0' was a $(2s, D+2)$-good embedding, it is easy to verify that ϕ_0 is a $(2s, D)$-good embedding.

Suppose the induction holds for some $i < m$ and let $e_{i+1} = \{a, b\}$. In short, we need to find a path from $\phi_i(a)$ to $\phi_i(b)$ of length $\sigma(e_{i+1})$, such that the part of it that goes through Γ_B maintains $(2s, D)$-goodness. In the proof we use auxiliary parameters $\ell_1, \ell_2, h \in \mathbb{N}$, defined as follows: choose $h \in \mathbb{N}$ to be the smallest integer such that $(D-1)^h \geq \varepsilon t$, and set $\ell_1 = \lfloor \sigma(e_{i+1})/2 \rfloor - h - 1$ and $\ell_2 = \lceil \sigma(e_{i+1})/2 \rceil - h - 1$. Note that $\ell_1, \ell_2 > 1$ since $\lfloor \sigma(e_{i+1})/2 \rfloor \geq \lfloor \delta \log n/2 \rfloor \geq \log_{D-1} n > h+2$, where we used $1/\varepsilon \gg D \gg 1/\delta$.

Let $F_i = \phi_i^{-1}(\Gamma_B)$ be the part of H_i^σ embedded into Γ_B, and f_i be the restriction of ϕ_i to F_i. We start by constructing the graph F_i' in two steps: First attach to F_i two paths of lengths ℓ_1 and ℓ_2, one rooted in a and the other in b, and let a' and b' denote other ends of such paths. Then attach two complete $(D-1)$-ary trees of depth h, one rooted in a' and the other in b'. Let us denote the set of leaves of these trees by L_a and L_b, respectively, and note that $|L_a| = |L_b| = (D-1)^h \geq \varepsilon t$ by the choice of h. Such trees have less than $(D-1)^{h+1} \leq (D-1)^2 \varepsilon t$ vertices each, which together with a trivial bound $\ell_1 \leq \ell_2 < \sigma(e_{i+1}) < v(H^\sigma)$ implies

$$v(F_i') \leq v(F_i) + 2(\ell_2 - 1) + 2 \cdot ((D-1)^2 \varepsilon t - 1) \leq v(H^\sigma) + 2v(H^\sigma) + 2(D-1)^2 \varepsilon t < s.$$

Assuming $D \geq 3$ each vertex has degree at most D in F_i' and, by its definition, F_i' can be constructed from F_i by successively adding a vertex of degree 1. Therefore, we can apply Theorem 2.3 to obtain a $(2s, D)$-good embedding $f_i' \colon F_i' \hookrightarrow \Gamma_B$ which extends f_i.

Every vertex in L_a is at distance exactly $\ell_1 + h$ from a, and every vertex in L_b is at distance exactly $\ell_2 + h$ from b. Thus $f_i'(L_a) \subseteq V_{j_1}'$ and $f_i'(L_b) \subseteq V_{j_2}'$ for some $j_1, j_2 \in \{1, 2\}$. Next, we find a path of length 2 from $f_i'(L_a)$ to $f_i'(L_b)$ with the internal vertex lying in V_3. From (Γ, ε, p)-regularity of the pairs (V_1, V_3) and (V_2, V_3), and $|f_i'(L_a)|, |f_i'(L_b)| \geq \varepsilon t$, we know that all but at most $2\varepsilon t$ vertices in $V_3 \setminus \phi_i(H_i^\sigma)$ are adjacent to both $f_i'(L_a)$ and $f_i'(L_b)$. As $|V_3| = t$ and $v(H^\sigma) < \varepsilon t$, this implies that there exists a free vertex in V_3 adjacent both to $f_i'(L_a)$ and $f_i'(L_b)$, which gives a desired path of length 2.

To summarize, we have found a path $P(x, y)$ of length 2 from $f_i'(x)$ to $f_i'(y)$, for some $x \in L_a$ and $y \in L_b$, with the internal vertex avoiding $V_1 \cup V_2$ and $\phi_i(H_i^\sigma)$. By Lemma 2.4, the restriction of f_i' to the graph obtained by removing all newly added vertices to F_i which do not lie either on the path from x to a or from y to b is $(2s, D)$-good. Together with the path $P(x, y)$, this defines an embedding ϕ_{i+1} of H_{i+1}^σ into Γ. □

3.2 Vertex-disjoint paths in expanding graphs

Theorem 2.7 provides a framework for embedding forests (in polynomial time) into graphs with certain expansion properties, while allowing arbitrary leaf deletions along the way. We present an application of this result to the classical problem of finding vertex-disjoint paths between given pairs of vertices in graphs.

Now we state the key result of this subsection. Theorem 1.3 (stated in the introduction) will then follow directly from the properties of (N, D, λ)-graphs.

Theorem 3.4. *Let G be a graph with the $P_\alpha(n, d)$ property for $3 \leq d < n$, and such that for every two disjoint $U, V \subseteq V(G)$ of sizes $|U|, |V| \geq n/8d$ there exists an edge between U and V. Let $P = \{a_i, b_i\}$ be a collection of at most $\frac{dn \log d}{5 \log n}$ disjoint pairs of vertices in G, such that $|N_G(x) \cap (A \cup B)| \leq \beta d_G(x)$ for every $x \in V(G)$, where $A = \cup_i \{a_i\}, B = \cup_i \{b_i\}$. If $\alpha - \beta \geq 1/2$ then there exists a polynomial time algorithm to find vertex-disjoint paths in G between every pair of vertices $\{a_i, b_i\}$, such that the length of each path is $2\lceil \frac{\log(n/8d)}{\log(d-1)} \rceil + 1$.*

Proof. By Theorem 2.7, there is an algorithm which works in time polynomial in $V(G)$, and wins the (G, P, nd, n^3, d)-forest building game. We construct the required disjoint paths one by one as follows. Let h be the smallest integer such that $(d-1)^h > \frac{n}{8d}$.

For the first pair $\{a_1, b_1\}$ we find two disjoint complete $(d-1)$-ary trees of depth h in G, rooted at a_1 and b_1, using the algorithm for winning the forest building game. Since both trees have more than $\frac{n}{8d}$ leaves, there is an edge connecting these sets of leaves, thus creating a path (between a_1 and b_1) of length $2h+1$. Remove from our current forest all other edges which do not lie on this path. We continue in the same fashion, by finding two complete $(d-1)$-ary trees rooted at a_2 and b_2 (disjoint from the path connecting a_1 and b_1), then finding a connecting edge between the sets of leaves, and removing all edges from the $(d-1)$-ary trees, which do not lie on the found path. We delete the edges successively, by always removing the edges which are incident with vertices of degree 1, just like in the forest building game.

We do this procedure for every pair of vertices, and note that we can do this as at any given point the current forest which we use for our argument has at most

$$\frac{dn \log d}{5 \log n} \cdot (2h+1) + 2 \cdot 2 \cdot \frac{n}{8} < \frac{dn}{2} + \frac{n}{2} < dn$$

edges, where the first term is a bound on the total number of edges used in previous paths, and the second one bounds the number of edges in the current $(d-1)$-ary trees we use. Furthermore we delete vertices of degree 1 at most $|A \cup B| \cdot n < n^3$ times. This completes the proof. □

The following result can be derived from the Expander Mixing Lemma through rather routine calculations.

Lemma 3.5 ([20], Lemma 2.7)**.** *Let G be an (N, D, λ)-graph and let d, n be positive integers. G has property $P_\alpha(n, d)$ for $\alpha > 0$ if the following holds:*

$$1 - \alpha > \frac{n(1+4d)}{2N} + \frac{\lambda}{D}(1 + \sqrt{2}d).$$

We are ready to give the promised proof of Theorem 1.3.

Theorem 1.3. *Let $\varepsilon > 0$, and let G be an (N, D, λ)-graph, with $\lambda < D^{1-\varepsilon}/150$ and $D^\varepsilon > 25$. Let $P = \{a_i, b_i\}$ be a collection of at most $\frac{\varepsilon N \log D}{160 \log N}$ disjoint pairs of vertices in G, such that $|N_G(x) \cap (A \cup B)| \le \frac{D}{4}$ for every $x \in V(G)$, where $A = \cup_i \{a_i\}, B = \cup_i \{b_i\}$. There exists a polynomial time algorithm to find vertex-disjoint paths in G between every pair of vertices $\{a_i, b_i\}$, such that the paths are of equal length which is less than $5 \frac{\log N}{\varepsilon \log D}$.*

Proof. Let $n = N/16d$ and $d = D^{\varepsilon/2}$. From Lemma 3.5 we see that G has the $P_{3/4}(n, d)$-property. Furthermore, by the Expander Mixing Lemma (eq. (1)), we have that for sets $U, V \subseteq V(G)$ of size at least $\frac{n}{8d}$ it holds:

$$\left| e_G(U, V) - \frac{Dn^2}{64Nd^2} \right| \le \lambda \frac{n}{8d}.$$

which, together with $\lambda < D^{1-\varepsilon}/150$ gives $e_G(U, V) > 0$. Applying Theorem 3.4 to G completes the proof; here are the final calculations.

- Number of pairs:

$$\frac{dn \log d}{5 \log n} = \frac{N \log D^{\varepsilon/2}}{16 \cdot 5 \log(N/16d)} = \frac{\varepsilon N \log D}{160 \log(N/16d)} > \frac{\varepsilon N \log D}{160 \log N};$$

- Length of paths:

$$2 \left\lceil \frac{\log(n/8d)}{\log(d-1)} \right\rceil + 1 = 2 \left\lceil \frac{\log(N/128d^2)}{\log(D^{\varepsilon/2} - 1)} \right\rceil + 1 \le 5 \frac{\log N}{\varepsilon \log D}.$$

□

4 Concluding remarks

We presented two embedding techniques (algorithmic and non-algorithmic) for embedding graphs of large girth into sparse expanders. Both are based on the embedding result by Friedman and Pippenger, enhanced with a roll-back idea which allows retracing previous embedding steps. We showed two applications of these techniques:

- We proved that the size-Ramsey number of logarithmic subdivisions of bounded degree graphs is linear in their order;
- For a given (n, d, λ)-graph with relatively small spectral ratio and any collection of $c\frac{n \log d}{\log n}$ disjoint pairs of vertices which satisfy a natural local condition, we gave a polynomial time algorithm which finds vertex disjoint paths of (the same) logarithmic length between each pair.

The first result answers a question of Pak [48], and the second one answers an offline version of a question of Alon and Capalbo [3]. With regards to the latter result, our offline algorithm can be made online in the following sense: instead of all the pairs being given in advance, the adversary can choose a set S of vertices of size $|S| = c\frac{n \log d}{\log n}$ (which satisfies the same local condition as before, and is given after G is exposed). Then he chooses pairs from S one by one, and we connect each pair as soon as it is given by the adversary. The same proof works for this stronger version of our algorithm. Finding an online polynomial time algorithm which does not require S to be chosen in advance (i.e. the pairs are chosen from $V(G)$ one by one) and where the lengths of the paths are logarithmic, remains open.

It can also be seen from our proof that we can allow the adversary to terminate arbitrary already established connections between pairs (and thus freeing the used vertices in the corresponding paths) a finite number of times during the mentioned online algorithm. This feature is related to the study of permutation networks [1].

References

[1] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswami, B. Schieber, and M. Sudan. Efficient routing in optical networks. *Journal of the ACM (JACM)*, 43(6):973–1001, 1996.

[2] N. Alon. Subdivided graphs have linear Ramsey numbers. *Journal of Graph Theory*, 18(4):343–347, 1994.

[3] N. Alon and M. Capalbo. Finding disjoint paths in expanders deterministically and online. *48th annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, 518–524, 2007.

[4] N. Alon and F. R. Chung. Explicit construction of linear sized tolerant networks. *Discrete Mathematics*, 72(1-3):15–19, 1988.

[5] N. Alon, M. Krivelevich, and B. Sudakov. Embedding nearly-spanning bounded degree trees. *Combinatorica*, 27(6):629–644, 2007.

[6] S. Arora, T. Leighton, and B. Maggs. On-line algorithms for path selection in a nonblocking network. *Proceedings of the twenty-second annual ACM Symposium on Theory of Computing (STOC'90)*, 149–158, 1990.

[7] D. Bal and L. DeBiasio. New lower bounds on the size-Ramsey number of a path. *arXiv preprint arXiv:1909.06354*, 2019.

[8] J. Balogh, B. Csaba, M. Pei, and W. Samotij. Large bounded degree trees in expanding graphs. *Electronic Journal of Combinatorics* 17, Research paper 6, 2010.

[9] J. Balogh, B. Csaba, and W. Samotij. Local resilience of almost spanning trees in random graphs. *Random Structures & Algorithms*, 38(1–2):121–139, 2011.

[10] J. Beck. On size Ramsey number of paths, trees, and circuits. I. *Journal of Graph Theory*, 7(1):115–129, 1983.

[11] J. Beck. On size Ramsey number of paths, trees and circuits. II. *Mathematics of Ramsey theory, Algorithms Combin.*, 5: 34–45. Springer, Berlin, 1990.

[12] S. Berger, Y. Kohayakawa, G. S. Maesaka, T. Martins, W. Mendonça, G. O. Mota, and O. Parczyk. The size-Ramsey number of powers of bounded degree trees. *Acta Mathematica Universitatis Comenianae*, 88(3):451–456, 2019.

[13] A. Z. Broder, A. M. Frieze, S. Suen, and E. Upfal. An efficient algorithm for the vertex-disjoint paths problem in random graphs. *Proceedings of the seventh annual ACM-SIAM Symposium on Discrete Algorithms (SODA'96)*, 261–268, 1996.

[14] S. A. Burr and P. Erdős. On the magnitude of generalized Ramsey numbers for graphs. Colloq. Math. Soc. János Bolyai, 10: 215–240, 1975.

[15] D. Clemens, M. Jenssen, Y. Kohayakawa, N. Morrison, G. O. Mota, D. Reding, and B. Roberts. The size-Ramsey number of powers of paths. *Journal of Graph Theory*, 91(3):290–299, 2019.

[16] D. Conlon. A new upper bound for diagonal Ramsey numbers. *Annals of Mathematics*, 941–960, 2009.

[17] D. Conlon, J. Fox, and B. Sudakov. Recent developments in graph Ramsey theory. *Surveys in combinatorics*, London Mathematical Society Lecture Notes 424: 49–118, 2015.

[18] D. Conlon and R. Nenadov. Size Ramsey numbers of triangle-free graphs with bounded degree. *Preprint.*

[19] D. Dellamonica Jr. The size-Ramsey number of trees. *Random Structures & Algorithms*, 40(1):49–73, 2012.

[20] D. Dellamonica Jr. and Y. Kohayakawa. An algorithmic Friedman–Pippenger theorem on tree embeddings and applications. *Electronic Journal of Combinatorics* 15, Research paper 127, 2008.

[21] J. Donadelli, P. E. Haxell, and Y. Kohayakawa. A note on the size-Ramsey number of long subdivisions of graphs. *RAIRO – Theoretical Informatics and Applications*, 39(1):191–206, 2005.

[22] N. Draganić, M. Krivelevich, and R. Nenadov. The size-Ramsey number of short subdivisions. *arXiv preprint arXiv:2004.14139*, 2020.

[23] A. Dudek and P. Prałat. On some multicolor Ramsey properties of random graphs. *SIAM Journal on Discrete Mathematics*, 31(3):2079–2092, 2017.

[24] P. Erdős. On the combinatorial problems which I would most like to see solved. *Combinatorica*, 1(1):25–42, 1981.

[25] P. Erdős, R. J. Faudree, C. C. Rousseau, and R. H. Schelp. The size Ramsey number. *Periodica Mathematica Hungarica*, 9(1-2):145–161, 1978.

[26] P. Feldman, J. Friedman, and N. Pippenger. Non-blocking networks. *Proceedings of the eighteenth annual ACM Symposium on Theory of Computing (STOC'86)*, 247–254, 1986.

[27] P. Feldman, J. Friedman, and N. Pippenger. Wide-sense nonblocking networks. *SIAM Journal on Discrete Mathematics*, 1(2):158–173, 1988.

[28] J. Friedman and N. Pippenger. Expanding graphs contain all small trees. *Combinatorica*, 7(1):71–76, 1987.

[29] A. M. Frieze. Disjoint paths in expander graphs via random walks: A short survey. *International Workshop on Randomization and Approximation Techniques in Computer Science*, 1–14. Springer, 1998.

[30] M. R. Garey and D. S. Johnson. *Computers and intractability*, vol. 174, Freeman, 1979.

[31] J. Han, M. Jenssen, Y. Kohayakawa, G. O. Mota, and B. Roberts. The multicolour size-Ramsey number of powers of paths. *Journal of Combinatorial Theory, Ser. B*, 145: 359–375, 2020.

[32] P. E. Haxell, Y. Kohayakawa, and T. Łuczak. The induced size-Ramsey number of cycles. *Combinatorics, Probability and Computing*, 4(3):217–239, 1995.

[33] D. S. Hochbaum. An exact sublinear algorithm for the max-flow, vertex disjoint paths and communication problems on random graphs. *Operations research*, 40(5):923–935, 1992.

[34] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

[35] R. Javadi, F. Khoeini, G. R. Omidi, and A. Pokrovskiy. On the size-Ramsey number of cycles. *Combinatorics, Probability and Computing*, 28(6):871–880, 2019.

[36] D. Johannsen. Personal communication.

[37] N. Kamčev, A. Liebenau, D. Wood, and L. Yepremyan. The size Ramsey number of graphs with bounded treewidth. *arXiv preprint arXiv:1906.09185*, 2019.

[38] Y. Kohayakawa, T. Retter, and V. Rödl. The size-Ramsey number of short subdivisions of bounded degree graphs. *Random Structures & Algorithms*, 54(2):304–339, 2016.

[39] Y. Kohayakawa, V. Rödl, M. Schacht, and E. Szemerédi. Sparse partition universal graphs for graphs of bounded degree. *Advances in Mathematics*, 226(6):5041–5065, 2011.

[40] M. Krivelevich. Expanders — how to find them, and what to find in them. *Surveys in Combinatorics*, A. Lo et al., Eds., London Mathematical Society Lecture Notes 456: 115–142, 2019.

[41] M. Krivelevich. Long cycles in locally expanding graphs, with applications. *Combinatorica*, 39(1):135–151, 2019.

[42] M. Krivelevich and B. Sudakov. Pseudo-random graphs. *More sets, graphs and numbers*, Bolyai Soc. Math. Stud., Springer, 15: 199–262, 2006.

[43] S. Letzter. Path Ramsey number for random graphs. *Combinatorics, Probability and Computing*, 25(4):612–622, 2016.

[44] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[45] R. Montgomery. Sharp threshold for embedding combs and other spanning trees in random graphs. *arXiv preprint arXiv:1405.6560*, 2014.

[46] R. Montgomery. Logarithmically small minors and topological minors. *Journal of the London Mathematical Society*, 91(1):71–88, 2015.

[47] R. Montgomery. Spanning trees in random graphs. *Advances in Mathematics*, 356:106793, 2019.

[48] I. Pak. Mixing time and long paths in graphs. *Proceedings of the thirteenth annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*, 321–328, 2002.

[49] N. Pippenger. Telephone switching networks. *Proceedings of Symposia in Applied Mathematics*, 26: 101–133, 1982.

[50] F. P. Ramsey. On a problem of formal logic. *Proc. London Math. Soc.*, 2(1):264–286, 1930.

[51] V. Rödl and E. Szemerédi. On size Ramsey numbers of graphs with bounded degree. *Combinatorica*, 20(2):257–262, 2000.

[52] P. Seymour and N. Robertson. Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Ser. B*, 63(1): 65–110, 1995.

[53] E. Shamir and E. Upfal. A fast construction of disjoint paths in networks. *Annals of Discrete Mathematics*, 24: 141–154, 1985.

[54] A. Shapira and B. Sudakov. Small complete minors above the extremal edge density. *Combinatorica*, 35(1):75–94, 2015.

[55] J. Spencer. Ramsey's theorem — a new lower bound. *Journal of Combinatorial Theory, Ser. A*, 18(1):108–115, 1975.