

Hardness and Algorithms for Rainbow Connection

Sourav Chakraborty ^{*} Eldar Fischer [†] Arie Matsliah [‡] Raphael Yuster [§]

Abstract

An edge-colored graph G is *rainbow connected* if any two vertices are connected by a path whose edges have distinct colors. The *rainbow connection* of a connected graph G , denoted $rc(G)$, is the smallest number of colors that are needed in order to make G rainbow connected. In the first result of this paper we prove that computing $rc(G)$ is NP-Hard solving an open problem from [6]. In fact, we prove that it is already NP-Complete to decide if $rc(G) = 2$, and also that it is NP-Complete to decide whether a given edge-colored (with an unbounded number of colors) graph is rainbow connected. On the positive side, we prove that for every $\epsilon > 0$, a connected graph with minimum degree at least ϵn has *bounded* rainbow connection, where the bound depends only on ϵ , and a corresponding coloring can be constructed in polynomial time. Additional non-trivial upper bounds, as well as open problems and conjectures are also presented.

1 Introduction

Connectivity is perhaps the most fundamental graph-theoretic property, both in the combinatorial sense and the algorithmic sense. There are many ways to strengthen the connectivity property, such as requiring hamiltonicity, k -connectivity, imposing bounds on the diameter, requiring the existence of edge-disjoint spanning trees, and so on.

An interesting way to quantitatively strengthen the connectivity requirement was recently introduced by Chartrand et al. in [5]. An edge-colored graph G is *rainbow connected* if any two vertices are connected by a path whose edges have distinct colors. Clearly, if a graph is rainbow connected, then it is also connected. Conversely, any connected graph has a trivial edge coloring that makes it rainbow connected; just color each edge with a distinct color. Thus, one can properly define the *rainbow connection* of a connected graph G , denoted $rc(G)$, as the smallest number of colors that are needed in order to make G rainbow connected. An easy observation is that if G is connected and has n vertices then $rc(G) \leq n - 1$, since one may color the edges of a given spanning tree with distinct colors. We note also the trivial fact that $rc(G) = 1$ if and only if G is a clique, the (almost) trivial fact that $rc(G) = n - 1$ if and only if G is a tree, and the easy observation that a cycle with

^{*}Department of Computer Science, University of Chicago, Chicago, IL-60637 USA. Email: sourav@cs.uchicago.edu

[†]Department of Computer Science, Technion, Haifa 32000, Israel. Email: eldar@cs.technion.ac.il

[‡]Centrum Wiskunde & Informatica (CWI), Amsterdam, Netherlands. Email: ariem@cwi.nl

[§]Department of Mathematics, University of Haifa, Haifa 31905, Israel. Email: raphy@math.haifa.ac.il

$k > 3$ vertices has rainbow connection $\lceil k/2 \rceil$. Also notice that, clearly, $rc(G) \geq \text{diam}(G)$ where $\text{diam}(G)$ denotes the diameter of G .

In addition to being a natural combinatorial measure, rainbow connectivity can be motivated by its interesting interpretation in the area of networking. Suppose that G represents a network (e.g., a cellular network). We wish to route messages between any two vertices in a pipeline, and require that each link on the route between the vertices (namely, each edge on the path) is assigned a distinct channel (e.g. a distinct frequency). Clearly, we want to minimize the number of distinct channels that we use in our network. This number is precisely $rc(G)$.

Chartrand et al. computed the rainbow connection of several graph classes including complete multipartite graphs [5]. Caro et al. [6] considered the extremal graph-theoretic aspects of rainbow connection. They proved that if G is a connected graph with n vertices and with minimum degree 3 then $rc(G) < 5n/6$, and if the minimum degree is δ then $rc(G) \leq \frac{\ln \delta}{\delta} n(1 + f(\delta))$ where $f(\delta)$ tends to zero as δ increases. They also determine the threshold function for a random graph $G(n, p(n))$ to have $rc(G) = 2$. In their paper, they conjecture that computing $rc(G)$ is an NP-Hard problem, as well as conjecture that even deciding whether a graph has $rc(G) = 2$ is NP-Complete.

In this paper we address the computational aspects of rainbow connection. Our first set of results solve, and extend, the complexity conjectures from [6]. Indeed, it turns out that deciding whether $rc(G) = 2$ is an NP-Complete problem. Our proof is by a series of reductions, where on the way it is shown that 2-rainbow-colorability is polynomial-time equivalent to the seemingly harder question of deciding the existence of a 2-edge-coloring that is required to rainbow-connect only vertex pairs from a prescribed set.

Theorem 1.1 *Given a graph G , deciding if $rc(G) = 2$ is NP-Complete. In particular, computing $rc(G)$ is NP-Hard.*

Suppose we are given an edge coloring of the graph. Is it then easier to verify whether the colored graph is rainbow connected? Clearly, if the number of colors is constant then this problem becomes easy. However, if the coloring is arbitrary, the problem becomes NP-Complete:

Theorem 1.2 *The following problem is NP-Complete: Given an edge-colored graph G , check whether the given coloring makes G rainbow connected.*

For the proof of Theorem 1.2, we first show that the $s - t$ version of the problem is NP-Complete. That is, given two vertices s and t of an edge-colored graph, decide whether there is a rainbow path connecting them.

We now turn to positive algorithmic results. Our main positive result is that connected n -vertex graphs with minimum degree $\Theta(n)$ have *bounded* rainbow connection. More formally, we prove:

Theorem 1.3 *For every $\epsilon > 0$ there is a constant $C = C(\epsilon)$ such that if G is a connected graph with n vertices and minimum degree at least ϵn , then $rc(G) \leq C$. Furthermore, there is a polynomial time algorithm that constructs a corresponding coloring for a fixed ϵ .*

The proof of Theorem 1.3 is based upon a modified degree-form version of Szemerédi's Regularity Lemma that we prove and that may be useful in other applications. From our algorithm it is

also not hard to find a probabilistic polynomial time algorithm for finding this coloring with high probability (using on the way the algorithmic version of the Regularity Lemma from [1] or [7]).

We note that connected graphs with minimum degree ϵn have bounded diameter, but the latter property by itself does *not* guarantee bounded rainbow connection. As an extreme example, a star with n vertices has diameter 2 but its rainbow connection is $n - 1$. The following theorem asserts however that having diameter 2 and only logarithmic minimum degree suffices to guarantee rainbow connection 3.

Theorem 1.4 *If G is an n -vertex graph with diameter 2 and minimum degree at least $8 \log n$ then $rc(G) \leq 3$. Furthermore, such a coloring is given with high probability by a uniformly random 3-edge-coloring of the graph G , and can also be found by a polynomial time deterministic algorithm.*

Since a graph with minimum degree $n/2$ is connected and has diameter 2, we have as an immediate corollary:

Corollary 1.5 *If G is an n -vertex graph with minimum degree at least $n/2$ then $rc(G) \leq 3$.*

The rest of this paper is organized as follows. The next section contains the hardness results, including the proofs of Theorem 1.1 and Theorem 1.2. Section 3 contains the proof of Theorem 1.3 and the proof of Theorem 1.4. At the end of the proof of each of the above theorems we explain how the algorithm can be derived – this mostly consists of using the conditional expectation method to derandomize the probabilistic parts of the proofs. The final Section 4 contains some open problems and conjectures.

2 Hardness results

We first give an outline of our proof of Theorem 1.1. We begin by showing the computational equivalence of the problem of rainbow connection 2, that asks for a red-blue edge coloring in which *all* vertex pairs have a rainbow path connecting them, to the problem of *subset rainbow connection* 2, asking for a red-blue coloring in which every pair of vertices in a *given subset* of pairs has a rainbow path connecting them. This is proved in Lemma 2.1 below.

In the second step, we reduce the problem of *extending to 2-rainbow connectivity*, asking whether a given partial red-blue coloring can be completed to obtain a rainbow connected graph, to the subset 2-rainbow connectivity problem. This is proved in Lemma 2.2 below.

Finally, the proof of Theorem 1.1 is completed by reducing 3-SAT to the problem of *extending to rainbow connection* 2.

Lemma 2.1 *The following problems are polynomial-time equivalent:*

1. *Given a graph G decide whether $rc(G) = 2$.*
2. *Given a graph G and a set of pairs $P \subseteq V(G) \times V(G)$, decide whether there is an edge coloring of G with 2 colors such that all pairs $(u, v) \in P$ are rainbow connected.*

Proof. It is enough to describe a reduction from Problem 2 to Problem 1. Given a graph $G = (V, E)$ and a set of pairs $P \subseteq V \times V$, we construct a graph $G' = (V', E')$ as follows.

For every vertex $v \in V$ we introduce a new vertex x_v , and for every pair $(u, v) \in (V \times V) \setminus P$ we introduce a new vertex $x_{(u,v)}$. We set

$$V' = V \cup \{x_v : v \in V\} \cup \{x_{(u,v)} : (u, v) \in (V \times V) \setminus P\}$$

and

$$E' = E \cup \left\{ \{v, x_v\} : v \in V \right\} \cup \left\{ \{u, x_{(u,v)}\}, \{v, x_{(u,v)}\} : (u, v) \in (V \times V) \setminus P \right\} \cup \left\{ \{x, x'\} : x, x' \in V' \setminus V \right\}.$$

It remains to verify that G' is 2-rainbow connected if and only if there is an edge coloring of G with 2 colors such that all pairs $(u, v) \in P$ are rainbow connected. In one direction, notice that when G is considered as a subgraph of G' , no pair of vertices of G that appear in P has a path of length two in G' that is not fully contained in G . Hence, if G is not colorable in a way connecting the pairs in P , the graph G' is not 2-rainbow-connected.

In the other direction, assume that $\chi : E \rightarrow \{\text{red}, \text{blue}\}$ is a coloring of G that rainbow-connects the pairs in P . To extend it to a rainbow-coloring $\chi' : E' \rightarrow \{\text{red}, \text{blue}\}$, define $\chi'(\{v, x_v\}) = \text{blue}$ for all $v \in V$, $\chi'(\{u, x_{(u,v)}\}) = \text{blue}$ and $\chi'(\{v, x_{(u,v)}\}) = \text{red}$ for all $(u, v) \notin P$ (note that we treat P as a set of ordered pairs – an unordered pair can be represented by putting both orderings in P), and finally $\chi'(\{x, x'\}) = \text{red}$ for all $x, x' \in V' \setminus V$. One can see that χ' is indeed a valid 2-rainbow-coloring of G' , concluding the proof. ■

Lemma 2.2 *The first problem defined below is polynomially reducible to the second one:*

1. *Given a graph $G = (V, E)$ and a partial 2-edge-coloring $\hat{\chi} : \hat{E} \rightarrow \{0, 1\}$ for $\hat{E} \subset E$, decide whether $\hat{\chi}$ can be extended to a complete 2 edge-coloring $\chi : E \rightarrow \{0, 1\}$ that makes G rainbow connected.*
2. *Given a graph G and a set of pairs $P \subseteq V(G) \times V(G)$ decide whether there is an edge coloring of G with 2 colors such that all pairs $(u, v) \in P$ are rainbow connected.*

Proof. Since the identity of the colors does not matter, it is more convenient that instead of a coloring $\chi : E \rightarrow \{0, 1\}$ we consider the corresponding partition $\pi_\chi = (E_1, E_2)$ of E . Similarly, in the case of a partial coloring $\hat{\chi}$, the pair $\pi_{\hat{\chi}} = (\hat{E}_1, \hat{E}_2)$ will contain the corresponding disjoint subsets of E (which may not cover E).

Now, given such a partial coloring $\hat{\chi}$ we extend the original graph $G = (V, E)$ to a graph $G' = (V', E')$, and define a set P of pairs of vertices such that for the resulting graph the answer for Problem 2 is “yes” if and only if the answer for Problem 1 for the original graph is “yes”.

Let $\ell : V \rightarrow [|V|]$ be arbitrary linear ordering of the vertices, and let **high** : $E \rightarrow V$ be a mapping that maps an edge $e = \{u, v\}$ to u if $\ell(u) > \ell(v)$, and to v otherwise. Similarly, let **low** : $E \rightarrow V$ be a mapping that maps an edge $e = \{u, v\}$ to u if $\ell(u) < \ell(v)$, and to v otherwise.

We construct G' as follows. We add $3 + |\hat{E}_1| + |\hat{E}_2|$ new vertices

$$\{b_1, c, b_2\} \cup \{c_e : e \in (\hat{E}_1 \cup \hat{E}_2)\}$$

and add the edges

$$\{\{b_1, c\}, \{c, b_2\}\} \cup \{\{b_i, c_e\} : i \in \{1, 2\}, e \in \hat{E}_i\} \cup \{\{c_e, \mathbf{low}(e)\} : e \in (\hat{E}_1 \cup \hat{E}_2)\}.$$

Now we define the set P of pairs of vertices that have to be 2-rainbow connected:

$$P = \{\{b_1, b_2\}\} \cup \{\{u, v\} : u, v \in V\} \cup \{\{c, c_e\} : e \in (\hat{E}_1 \cup \hat{E}_2)\} \\ \cup \{\{b_i, \mathbf{low}(e)\} : i \in \{1, 2\}, e \in \hat{E}_i\} \cup \{\{c_e, \mathbf{high}(e)\} : e \in (\hat{E}_1 \cup \hat{E}_2)\}.$$

Now, if there is a 2-rainbow-coloring $\pi_\chi = (E_1, E_2)$ of G which extends $\pi_{\hat{\chi}} = (\hat{E}_1, \hat{E}_2)$, then we color G' as follows. E is colored as mandated by χ , that is, an edge is colored red if it is in E_1 and otherwise it is colored blue. Now $\{b_1, c\}, \{b_2, c_e\}$ for $e \in \hat{E}_2$ and $\{c_e, \mathbf{low}(e)\}$ for $e \in \hat{E}_1$ are all colored blue, and $\{b_2, c\}, \{b_1, c_e\}$ for $e \in \hat{E}_1$ and $\{c_e, \mathbf{low}(e)\}$ for $e \in \hat{E}_2$ are all colored red. One can see that this coloring indeed rainbow-connects all the pairs in P .

On the other hand, any 2-edge-coloring of G' that connects the pairs in P clearly contains a 2-rainbow-coloring of G , because P contains all vertex pairs of G when considered as an induced subgraph of G' , and also G' contains no path of length 2 between vertices of G that is not contained in G . Also, such a coloring would have to color $\{c, b_1\}$ and $\{c, b_2\}$ differently. It would also have to color every $\{b_i, c_e\}$ in a color different from that of $\{c, b_i\}$, and would hence color $\{c_e, \mathbf{low}(e)\}$ in a color identical to that of $\{c, b_i\}$, because it has to be the color different from that of $\{b_i, c_e\}$. Finally, every $e \in \hat{E}_i$ would have to be colored with the color different from that of $\{c, b_i\}$ so as to rainbow-connect $\mathbf{high}(e)$ and c_e . This means that the coloring of G' not only provides a 2-rainbow-coloring of G , but that it also conforms to the original partial coloring $\hat{\chi}$. ■

Proof of Theorem 1.1. We show that Problem 1 of Lemma 2.2 is NP-hard, and then deduce that 2-rainbow-colorability is NP-Complete by applying Lemma 2.1 and Lemma 2.2 while observing that it clearly belongs to NP.

We reduce 3-SAT to Problem 1 of Lemma 2.2. Given a 3CNF formula $\phi = \bigwedge_{i=1}^m c_i$ over variables x_1, x_2, \dots, x_n , we construct a graph G_ϕ and a partial 2-edge coloring $\chi' : E(G_\phi) \rightarrow \{0, 1\}$ such that there is an extension χ of χ' that makes G_ϕ rainbow connected if and only if ϕ is satisfiable.

We define G_ϕ as follows:

$$V(G_\phi) = \{c_i : i \in [m]\} \cup \{x_i : i \in [n]\} \cup \{a\}$$

$$E(G_\phi) = \{\{c_i, x_j\} : x_j \in c_i \text{ in } \phi\} \cup \{\{x_i, a\} : i \in [n]\} \cup \{\{c_i, c_j\} : i, j \in [m]\} \cup \{\{x_i, x_j\} : i, j \in [n]\}$$

and we define the partial coloring χ' as follows:

$$\forall_{i, j \in [m]} \chi'(\{c_i, c_j\}) = 0$$

$$\forall_{i,j \in [n]} \chi'(\{x_i, x_j\}) = 0$$

$$\forall_{\{x_i, c_j\} \in E(G_\phi)} \chi'(\{x_i, c_j\}) = 0 \text{ if } x_i \text{ is positive in } c_j, 1 \text{ otherwise}$$

while all the edges in $\{\{x_i, a\} : i \in [n]\}$ (and only they) are left uncolored.

Assuming without loss of generality that all variables in ϕ appear both as positive and as negative, one can verify that a 2-rainbow-coloring of the uncolored edges corresponds to a satisfying assignment of ϕ and vice versa. ■

The proof of Theorem 1.2 is based upon the proof of the following theorem.

Theorem 2.3 *The following problem is NP-complete: Given an edge colored graph G and two vertices s, t of G , decide whether there is a rainbow path connecting s and t .*

Proof. Clearly the problem is in NP. We prove that it is NP-Complete by reducing 3-SAT to it. Given a 3CNF formula $\phi = \bigwedge_{i=1}^m c_i$ over variables x_1, x_2, \dots, x_n , we construct a graph G_ϕ with two special vertices s, t and a coloring $\chi : E(G_\phi) \rightarrow [|E(G_\phi)|]$ such that there is a rainbow path connecting s and t in G_ϕ if and only if ϕ is satisfiable.

We start by constructing an auxiliary graph G' from ϕ . The graph G' has $3m + 2$ vertices, that are partitioned into $m + 2$ layers $V_0, V_1, \dots, V_m, V_{m+1}$, where $V_0 = \{s\}$, $V_{m+1} = \{t\}$ and for each $i \in [m]$, the layer V_i contains the three vertices corresponding to the literals of c_i (a clause in ϕ). The edges of G' connect between all pairs of vertices residing in consecutive layers. Formally,

$$E(G') = \left\{ \{u, v\} : \exists i \in [m + 1] \text{ s.t. } u \in V_{i-1} \text{ and } v \in V_i \right\}.$$

Intuitively, in our final colored graph G_ϕ , every rainbow path from s to t will define a satisfying assignment of ϕ in a way that for every $i \in [m]$, if the rainbow path contains a vertex $v \in V_i$ then the literal of c_i that corresponds to v is satisfied, and hence c_i is satisfied. Since any path from s to t must contain at least one vertex from every layer V_i , this will yield a satisfying assignment for the whole formula ϕ . But we need to make sure that there are no contradictions in this assignment, that is, no opposite literals are satisfied together. For this we modify G' by replacing each literal-vertex with a gadget, and we define an edge coloring for which rainbow paths yield only consistent assignments.

For every variable x_j , $j \in [n]$, let $v_{j_1}, v_{j_2}, \dots, v_{j_k}$ be the vertices of G' corresponding to the positive literal x_j , and let $\bar{v}_{j_1}, \bar{v}_{j_2}, \dots, \bar{v}_{j_\ell}$ be the vertices corresponding to the negative literal \bar{x}_j . We can assume without loss of generality that both $k \geq 1$ and $\ell \geq 1$, since otherwise the formula ϕ can be simplified. For every such variable x_j we also introduce $k \times \ell$ distinct colors $\alpha_{1,1}^j, \dots, \alpha_{k,\ell}^j$. Next, we transform the auxiliary graph G' into the final graph G_ϕ .

For every $a \in [k]$ we replace the vertex v_{j_a} that resides in layer (say) V_i with $\ell + 1$ new vertices $v_1, v_2, \dots, v_{\ell+1}$ that form a path in that order. We also connect all vertices in V_{i-1} to v_1 and connect all vertices in V_{i+1} to $v_{\ell+1}$. For every $b \in [\ell]$, we color the edge $\{v_b, v_{b+1}\}$ in the new path with the color $\alpha_{a,b}^j$. Similarly, for every $b \in [\ell]$ we replace the vertex \bar{v}_{j_b} from layer (say) $V_{i'}$ with $k + 1$ new vertices $\bar{v}_1, \bar{v}_2, \dots, \bar{v}_{k+1}$ that form a path, and connect all vertices in $V_{i'-1}$ to \bar{v}_1 and all vertices in $V_{i'+1}$ to \bar{v}_{k+1} . For every $a \in [k]$, we color the edge $\{v_a, v_{a+1}\}$ with $\alpha_{a,b}^j$. All other edges of G_ϕ (which were the original edges of G') are colored with fresh distinct colors.

Clearly, any path from s to t in G_ϕ must contain at least one of the newly built paths in each layer. On the other hand, it is not hard to verify that any two paths of opposite literals of the same variable have edges sharing the same color. ■

Proof of Theorem 1.2. We reduce from the problem in Theorem 2.3. Given an edge colored graph $G = (V, E)$ with two special vertices s and t , we construct a graph $G' = (V', E')$ and define a coloring $\chi' : E' \rightarrow [|E'|]$ of its edges such that s and t are rainbow connected in G if and only if the coloring of G' makes G' rainbow connected.

Let $V = \{v_1 = s, v_2, \dots, v_n = t\}$ be the vertices of the original graph G . We set

$$V' = V \cup \{s', t', b\} \cup \{v_1^1 = s^1, v_2^1, v_2^2, \dots, v_{n-1}^1, v_{n-1}^2, v_n^2 = t^2\}$$

and

$$E' = E \cup \left\{ \{s', s\}, \{t', t\}, \{s, s^1\}, \{t, t^2\} \right\} \cup \left\{ \{b, v_i\} : i \in [n] \right\} \cup \left\{ \{v_i, v_i^\alpha\} : i \in [n], \alpha \in \{1, 2\} \right\} \cup \left\{ \{v_i^\alpha, v_j^\beta\} : i, j \in [n], \alpha, \beta \in \{1, 2\} \right\}.$$

The coloring χ' is defined as follows:

- all edges $e \in E$ retain the original color, that is $\chi'(e) = \chi(e)$;
- the edges $\{s, s'\}, \{t, b\}$ and $\left\{ \{v_i, v_i^2\} : i \in [2, n] \right\}$ are colored with a special color c_1 ;
- the edges $\{t, t'\}, \{s, b\}$ and $\left\{ \{v_i, v_i^1\} : i \in [n-1] \right\}$ are colored with a special color c_2 ;
- the edges in $\left\{ \{v_i, b\} : i \in [2, n-1] \right\}$ are colored with a special color c_3 ;
- the edges in $\left\{ \{v_i^\alpha, v_j^\beta\} : i, j \in [n], \alpha, \beta \in \{1, 2\} \right\}$ are colored with a special color c_4 .

Now we can verify that χ' makes G' rainbow connected if and only if there was a rainbow path from s to t in G under χ .

Assume that there is a rainbow path from s to t in G under χ , and let us prove that χ' makes G' rainbow connected. First, let us fix $\hat{t} \in V'$ and argue that b is rainbow connected with \hat{t} . If $\hat{t} \in V$, then b and \hat{t} are neighbors. If $\hat{t} \in \{s', t'\}$ then there is a path from b to \hat{t} that goes thru s or t and uses only colors c_1, c_2 . If $\hat{t} = v_i^\alpha$ for some $i \in [2, n-1]$ and $\alpha \in \{1, 2\}$, then a path between b and \hat{t} can be formed by going thru (say) v_2 , then v_2^1 and then \hat{t} (using colors c_3, c_2, c_4 respectively).

Next notice that there is a rainbow path from s to any other vertex $\hat{t} \in V' \setminus \{s', b\}$, even without using the color c_1 . If $\hat{t} = t$ then by our assumption such a path exists, and since it uses only the original colors, this path can be extended to t' as well (using color c_2). If $\hat{t} = v_i$ for some $i \in [2, n-1]$, then this path can be formed by first going to b (using color c_2) and then going from b to \hat{t} (using color c_3). If $\hat{t} = v_i^\alpha$ for some $i \in [2, n-1]$ and $\alpha \in \{1, 2\}$, then such a path can be formed by first going to s^1 (using color c_2) and then directly to \hat{t} (using color c_4). Since all these options did not use the color c_1 , we deduce that s' is rainbow connected with all $\hat{t} \in V'$ as well.

Using similar (symmetric) arguments, it is also clear that there is a rainbow path from both t and t' to any other vertex in V' .

At last, let us fix v_i^α and $\hat{t} \in V' \setminus \{s, t, s', t', b\}$, and argue that v_i^α is rainbow connected with \hat{t} . There are only two options. Either $\hat{t} = v_j^\beta$ for some $j \in [2, n-1]$ and $\beta \in \{1, 2\}$, and in this case v_i^α and \hat{t} are neighbors. Otherwise, $\hat{t} = v_j$ for some $j \in [2, n-1]$. Then the rainbow path can be formed by going thru (say) v_j^1 to \hat{t} (using colors c_4, c_2).

Now assume that there is *no* rainbow path from s to t in G under χ , and let us prove that G' is not rainbow connected under χ' . In particular, we show that there is *no* rainbow path from s' to t' in G' under χ' . To see that, notice that in any path from s' to t' , the first edge must be $e_1 = \{s', s\}$ and the last edge must be $e_2 = \{t, t'\}$. Therefore, the special colors $\chi'(e_1) = c_1$ and $\chi'(e_2) = c_2$ cannot be used anywhere else along the path. This means that none of the vertices $\{v_i^\alpha : i \in [n], \alpha \in \{1, 2\}\}$ can participate in this path. In addition, the vertex b cannot participate either, since we cannot use an edge from b to s or t , or use any two edges from b outside those to s or t . So, any rainbow path from s' to t' must contain a rainbow path from s to t in the original graph G , therefore no such path exists.

■

3 Upper bounds and algorithms

The proof of our main Theorem 1.3 is based upon a modified degree-form version of Szemerédi's Regularity Lemma, that we prove here and that may be useful in other applications. We begin by introducing the Regularity Lemma and the already known degree-form version of it.

3.1 Regularity Lemma

The Regularity Lemma of Szemerédi [9] is one of the most important results in graph theory and combinatorics, as it guarantees that every graph has an ϵ -approximation of constant descriptive size, namely a size that depends only on ϵ and not on the size of the graph. This approximation “breaks” the graph into a constant number of pseudo-random bipartite graphs. This is very useful in many applications since dealing with random-like graphs is much easier than dealing with arbitrary graphs. In particular, as we shall see, the Regularity Lemma allows us to prove that graphs with linear minimum degree have bounded rainbow connection.

We first state the lemma. For two nonempty disjoint vertex sets A and B of a graph G , we define $E(A, B)$ to be the set of edges of G between A and B . The *edge density* of the pair is defined by $d(A, B) = |E(A, B)| / (|A||B|)$.

Definition 3.1 (ϵ -regular pair) *A pair (A, B) is ϵ -regular if for every $A' \subseteq A$ and $B' \subseteq B$ satisfying $|A'| \geq \epsilon|A|$ and $|B'| \geq \epsilon|B|$, we have $|d(A', B') - d(A, B)| \leq \epsilon$.*

An ϵ -regular pair can be thought of as a pseudo-random bipartite graph in the sense that it behaves almost as we would expect from a random bipartite graph of the same density. Intuitively,

in a random bipartite graph with edge density d , all large enough sub-pairs should have similar densities.

A partition V_1, \dots, V_k of the vertex set of a graph is called an *equipartition* if $|V_i|$ and $|V_j|$ differ by no more than 1 for all $1 \leq i < j \leq k$ (so in particular every V_i has one of two possible sizes). The *order* of an equipartition denotes the number of partition classes (k above). An equipartition V_1, \dots, V_k of the vertex set of a graph is called ϵ -regular if all but at most $\epsilon \binom{k}{2}$ of the pairs (V_i, V_j) are ϵ -regular. Szemerédi's Regularity Lemma can be formulated as follows.

Lemma 3.2 (Regularity Lemma [9]) *For every $\epsilon > 0$ and positive integer K , there exists $N = N_{3.2}(\epsilon, K)$, such that any graph with $n \geq N$ vertices has an ϵ -regular equipartition of order k , where $K \leq k \leq N$.*

As mentioned earlier, the following variation of the lemma comes useful in our context.

Lemma 3.3 (Regularity Lemma - degree form [8]) *For every $\epsilon > 0$ and positive integer K there is $N = N_{3.3}(\epsilon, K)$ such that given any graph $G = (V, E)$ with $n > N$ vertices, there is a partition of the vertex-set V into $k + 1$ sets V_0', V_1', \dots, V_k' , and there is a subgraph G' of G with the following properties:*

1. $K \leq k \leq N$,
2. $s \triangleq |V_0'| \leq \epsilon^{10}n$ and all other components V_i' , $i \in [k]$ are of size $\ell \triangleq \frac{n-s}{k}$,
3. for all $i \in [k]$, V_i' induces an independent set in G' ,
4. for all $i, j \in [k]$, the pair (V_i', V_j') is ϵ^{10} -regular in G' , with density either 0 or at least $\frac{\epsilon}{4}$,
5. for all $v \in V$, $\deg_{G'}(v) > \deg_G(v) - \frac{\epsilon}{3}n$.

This form of the lemma (see e.g. [8]) can be obtained by applying the original Regularity Lemma (with a smaller value of ϵ), and then “cleaning” the resulting partition. Namely, adding to the exceptional set V_0' all components V_i incident to many irregular pairs, deleting all edges between any other pairs of clusters that either do not form an ϵ -regular pair or they do but with density less than ϵ , and finally adding to V_0 also vertices whose degree decreased too much by this deletion of edges.

3.2 A modified degree form version of the Regularity Lemma

In order to prove that graphs with minimum degree ϵn have bounded rainbow connection, we need a special version of the Regularity Lemma, which is stated next.

Lemma 3.4 (Regularity Lemma - new version) *For every $\epsilon > 0$ and positive integer K there is $N = N_{3.4}(\epsilon, K)$ so that the following holds: If $G = (V, E)$ is a graph with $n > N$ vertices and minimum degree at least ϵn then there is a subgraph G'' of G , and a partition of V into V_1'', \dots, V_k'' with the following properties:*

1. $K \leq k \leq N$,
2. for all $i \in [k]$, $(1 - \epsilon)\frac{n}{k} \leq |V_i''| \leq (1 + \epsilon^3)\frac{n}{k}$,
3. for all $i \in [k]$, V_i'' induces an independent set in G'' ,
4. for all $i, j \in [k]$, (V_i'', V_j'') is an ϵ^3 -regular pair in G'' , with density either 0 or at least $\frac{\epsilon}{16}$,
5. for all $i \in [k]$ and every $v \in V_i''$ there is at least one other class V_j'' so that the number of neighbors of v in G'' belonging to V_j'' is at least $\frac{\epsilon}{2}|V_j''|$.

In the following proof we assume (without mentioning explicitly) that ϵ is a small enough constant (this assumption is not limiting, since adding edges can only improve the rainbow connection of a graph).

Proof. Given $\epsilon > 0$ and K , let us apply the degree-form Regularity Lemma (Lemma 3.3), with the parameters ϵ and K , and let $N = N_{3.3}(\epsilon, K)$. Let V'_0, V'_1, \dots, V'_k be the partition promised by Lemma 3.3, and let G' be the corresponding subgraph of G . Recall the parameters $s \leq \epsilon^{10}n$ and $\ell = \frac{n-s}{k}$ defined in Lemma 3.3.

Consider a mapping $f : V'_0 \rightarrow [k]$ where $f(v) = i$ implies that v has more than $\frac{1}{2}\epsilon\ell$ neighbors in G' that belong to V'_i . Such a mapping clearly exists since otherwise the degree of v in G' would have been at most $(s - 1 + k\frac{1}{2}\epsilon\ell) + \frac{\epsilon}{3}n < \epsilon n$.

Next consider the unique mapping $g : [k] \rightarrow \mathcal{P}([k])$ that maps an index $i \in [k]$ to the subset of indices $g(i) \subseteq [k]$ such that $j \in g(i)$ if and only if $d(V'_i, V'_j) \geq \epsilon/4$ in G' . We notice that the cardinality of $g(i)$ is not too small. Indeed, if v is any vertex of V'_i then the cardinality of $g(i)$ is at least

$$\frac{(\deg_{G'}(v) - s)}{\ell} > \frac{\epsilon}{2}k.$$

Now consider the following process that creates G'' from G' by placing the vertices of V'_0 one-by-one in the other vertex classes. Assume that $V'_0 = \{v_1, \dots, v_s\}$ and denote by V_i^t the extension of V_i after placing v_t somewhere. Initially, $V_i^0 = V'_i$ for $i = 1, \dots, k$, and eventually $V_i^s = V_i''$ for $i = 1, \dots, k$ forming the desired partition.

At stage t , we want to place $v_t \in V'_0$ somewhere. Consider the subset of indices $g(f(v_t))$. We will choose $j \in g(f(v_t))$ so that V_j^{t-1} has the smallest cardinality (if there are several candidates, we may choose any of them). We define $V_j^t = V_j^{t-1} \cup \{v_t\}$ and define $V_i^t = V_i^{t-1}$ for $i \neq j$. The neighbors of v_t that are kept in G'' are all the neighbors of v_t in G' that belong to V_x , where $x \in g(j)$. Notice that in particular $f(v_t) \in g(j)$.

Having created G'' it remains to prove its claimed properties. Clearly, properties 1 and 3 hold. For property 2, notice that the final sets $V_i'' = V_i^s$ have grown from the initial V'_i by no more than

$$\frac{s}{\left(\frac{\epsilon}{2}k\right)} < \frac{2}{\epsilon k}\epsilon^{10}n = 2\epsilon^9\frac{n}{k}$$

and hence

$$(1 - \epsilon)\frac{n}{k} \leq \ell = |V'_i| \leq |V_i''| = |V_i^s| \leq \ell + 2\epsilon^9\frac{n}{k} \leq (1 + \epsilon^8)\ell \leq (1 + \epsilon^3)\frac{n}{k}$$

and property 2 holds. (Recall that $\ell = (1 - \epsilon^{10})\frac{n}{k}$, and that we assume ϵ to be a small enough constant.)

For property 4, Let us first prove the requirement on the density. Notice that if the original density of (V'_i, V'_j) in G' was 0, it remained so also in G'' . Otherwise, if the density of (V'_i, V'_j) in G' was at least $\epsilon/4$, the fact that $|V''_i| \leq |V'_i| + 2\epsilon^4\frac{n}{k} < 2|V'_i|$ implies that the density of (V''_i, V''_j) is larger than $\frac{1}{16}\epsilon$, no matter where the neighbors of the newly added vertices reside.

As for the regularity condition in property 4, notice that if (V'_i, V'_j) had density 0 in G' , then the same holds in G'' and the pair (V'_i, V'_j) is trivially ϵ^3 -regular. Otherwise, (V'_i, V'_j) was ϵ^{10} -regular with density $\delta' \geq \epsilon/4$ in G' (i.e., there were precisely $\delta'\ell^2$ edges between V'_i and V'_j). Let us bound the density δ'' of (V''_i, V''_j) in G'' . The fact $\ell \leq |V''_i|, |V''_j| \leq (1 + \epsilon^8)\ell$ implies the following inequalities:

1. $\ell^2 = |V'_i||V'_j| \leq |V''_i||V''_j| \leq (1 + \epsilon^8)^2|V'_i||V'_j| < (1 + \epsilon^7)\ell^2$
2. $|E_{G'}(V'_i, V'_j)| \leq |E_{G''}(V''_i, V''_j)| \leq |E_{G'}(V'_i, V'_j)| + 2\epsilon^8(1 + \epsilon^8)\ell^2 \leq |E_{G'}(V'_i, V'_j)| + \epsilon^7\ell^2$

Hence, $\delta'' = \frac{|E_{G''}(V''_i, V''_j)|}{|V''_i||V''_j|}$ can be bounded as follows

$$\delta' - \epsilon^7 \leq \frac{\delta'}{(1 + \epsilon^7)} = \frac{|E_{G'}(V'_i, V'_j)|}{(1 + \epsilon^7)\ell^2} \leq \delta'' \leq \frac{|E_{G'}(V'_i, V'_j)| + \epsilon^7\ell^2}{\ell^2} = \delta' + \epsilon^7.$$

To prove that the pair (V''_i, V''_j) is ϵ^3 -regular in G'' , let $A \subseteq V''_i$, $|A| \geq \epsilon^3|V''_i|$ and $B \subseteq V''_j$, $|B| \geq \epsilon^3|V''_j|$ be arbitrary, and let us show that $d_{G''}(A, B) = \delta'' \pm \epsilon^3$. Let $A' \triangleq A \cap V'_i$ and $B' \triangleq B \cap V'_j$, and let $A'' \triangleq A \setminus A'$ and $B'' \triangleq B \setminus B'$. Notice that $|E_{G''}(A, B)| = |E_{G''}(A', B')| + |E_{G''}(A'', B)| + |E_{G''}(A, B'')|$. Clearly $|A''|, |B''| \leq \epsilon^8\ell$, and hence $|A''| \leq \epsilon^5|A|$ and $|B''| \leq \epsilon^5|B|$. Therefore,

$$0 \leq |E_{G''}(A'', B)|, |E_{G''}(A, B'')| \leq \epsilon^5|A||B|.$$

Since (V'_i, V'_j) was ϵ^5 -regular in G' ,

$$(\delta' - \epsilon^5)|A'||B'| \leq |E_{G''}(A', B')| \leq (\delta' + \epsilon^5)|A'||B'|.$$

Combining the two, we have

$$(\delta' - \epsilon^5)|A'||B'| \leq |E_{G''}(A, B)| \leq (\delta' + \epsilon^5)|A'||B'| + 2\epsilon^5|A||B| \leq (\delta' + 3\epsilon^5)|A||B|$$

and using $|A'| \geq (1 - \epsilon^5)|A|$, $|B'| \geq (1 - \epsilon^5)|B|$ we get

$$(\delta' - \epsilon^4)|A||B| \leq (\delta' - \epsilon^5)(1 - \epsilon^5)^2|A||B| \leq |E_{G''}(A, B)| \leq (\delta' + 3\epsilon^5)|A||B| \leq (\delta' + \epsilon^4)|A||B|.$$

Finally, since $\delta'' = \delta' \pm \epsilon^7$ we get

$$\frac{|E_{G''}(A, B)|}{|A||B|} = \delta'' \pm \epsilon^3$$

as required.

For property 5, this is easily checked to hold for vertices not originally from V'_0 because of property 5 of G' . Property 5 also holds for vertices of V'_0 since when v_t is placed in some V'_j , we know that it has at least $\frac{1}{2}\epsilon\ell$ neighbors in V'_i where $i = f(v_t)$. ■

We note that the partition guaranteed by our modified version of the Regularity Lemma can be found in polynomial time for a fixed ϵ (with somewhat worse constants), by using the exact same methods that were used in [1] for constructing an algorithmic version of the original Regularity Lemma.

3.3 Proof of Theorem 1.3

In this section we use our version of the Regularity Lemma to prove Theorem 1.3. First we need some definitions. Given a graph $G = (V, E)$ and two subsets $V_1, V_2 \subseteq V$, let $E(V_1, V_2)$ denote the set of edges having one endpoint in V_1 and another endpoint in V_2 . Given a vertex v , let $\Gamma(v)$ denote the set of v 's neighbors, and for $W \subseteq V$, let $\Gamma_W(v)$ denote the set $W \cap \Gamma(v)$.

For an edge coloring $\chi : E \rightarrow \mathcal{K}$, let π_χ denote the corresponding partition of E into (at most) $|\mathcal{K}|$ components. For two edge colorings χ and χ' , we say that χ' is a *refinement* of χ if $\pi_{\chi'}$ is a refinement of π_χ , which is equivalent to saying that $\chi'(e) = \chi'(e')$ always implies $\chi(e) = \chi(e')$.

Observation 3.5 *Let χ and χ' be two edge-colorings of a graph G , such that χ' is a refinement of χ . For any path P in G , if P is a rainbow path under χ , then P is a rainbow path under χ' . In particular, if χ makes G rainbow connected, then so does χ' . ■*

We define a set of eight distinct colors $\mathcal{K} = \{a_1, a_2, a_3, a_4, b_1, b_2, b_3, b_4\}$. Given a coloring $\chi : E \rightarrow \mathcal{K}$ we say that $u, v \in V$ are *a-rainbow connected* if there is a rainbow path from u to v using only the colors a_1, a_2, a_3, a_4 . We similarly define *b-rainbow connected* pairs. The following is a central lemma in the proof of Theorem 1.3.

Lemma 3.6 *For any $\epsilon > 0$, there is $N = N_{3.6}(\epsilon)$ such that any connected graph $G = (V, E)$ with $n > N$ vertices and minimum degree at least ϵn satisfies the following. There is a partition Π of V into $k \leq N$ components V_1, V_2, \dots, V_k , and a coloring $\chi : E \rightarrow \mathcal{K}$ such that for every $i \in [k]$ and every $u, v \in V_i$, the pair u, v is both a-rainbow connected and b-rainbow connected under χ .*

Proof. Proof given in Section 3.4. ■

Using Lemma 3.6 we derive the proof of Theorem 1.3. For a given $\epsilon > 0$, set $N = N_{3.6}(\epsilon)$ and set $C = \frac{3}{\epsilon}N + 8$. Clearly, any connected graph $G = (V, E)$ with $n \leq C$ vertices satisfies $rc(G) \leq C$. So we assume that $n > C \geq N$, and let $\Pi = V_1, \dots, V_k$ be the partition of V from Lemma 3.6, while we know that $k \leq N$.

First observe that since the minimal degree of G is ϵn , the diameter of G is bounded by $3/\epsilon$. This can be verified, e.g., by taking an arbitrary vertex $r \in V$ and executing a *BFS* algorithm from it. Let L_1, \dots, L_t be the layers of vertices in this execution, where L_i are all vertices at distance i from r . Observe that since the minimal degree is at least ϵn , the total number of vertices in every

three consecutive layers must be at least ϵn , thus $t \leq 3/\epsilon$. Since the same claim holds for any $r \in V$, this implies that $\text{diam}(G) \leq t \leq 3/\epsilon$.

Now let $T = (V_T, E_T)$ be a connected subtree of G on at most $k \cdot \text{diam}(G) \leq \frac{3}{\epsilon}N$ vertices such that for every $i \in [k]$, $V_T \cap V_i \neq \emptyset$. Such a subtree must exist in G since as observed earlier, $\text{diam}(G) \leq 3/\epsilon$. Let $\chi : E \rightarrow \mathcal{K}$ be the coloring from Lemma 3.6, and let $\mathcal{H} = \{h_1, h_2, \dots, h_{|E_T|}\}$ be a set of $|E_T| \leq \frac{3}{\epsilon}N$ fresh colors. We refine χ by recoloring every $e_i \in E(T)$ with color $h_i \in \mathcal{H}$. Let $\chi' : E \rightarrow (\mathcal{K} \cup \mathcal{H})$ be the resulting coloring of G . The following lemma completes the proof of Theorem 1.3.

Lemma 3.7 *The coloring χ' makes G rainbow connected. Consequently, $rc(G) \leq |E_T| + 8 \leq C$.*

Proof. Let $u, v \in V$ be any pair of G 's vertices. If u and v reside in the same component V_i of the partition Π , then (by Lemma 3.6) they are connected by a path P of length at most four, which is a rainbow path under the original coloring χ . Since χ' is a refinement of χ , the path P remains a rainbow path under χ' as well (see Observation 3.5).

Otherwise, let $u \in V_i$ and $v \in V_j$ for $i \neq j$. Let t_i and t_j be vertices of the subtree T , residing in V_i and V_j respectively. By definition of χ' , there is a rainbow path from t_i to t_j using colors from \mathcal{H} . Let P_t denote this path. In addition, by Lemma 3.6 we know that for the original coloring χ , there is a rainbow path P_a from u to t_i using colors a_1, \dots, a_4 and there is a rainbow path P_b from v to t_j using colors b_1, \dots, b_4 . Based on the fact that χ' is a refinement of χ , it is now easy to verify that P_t, P_a and P_b can be combined to form a rainbow path from u to v under χ' . ■

This concludes the proof of Theorem 1.3, apart from the existence of a polynomial time algorithm for finding this coloring. We note that all arguments above apart from Lemma 3.6 admit polynomial algorithms for finding the corresponding structures. The algorithm for Lemma 3.6 will be given with its proof.

3.4 Proof of Lemma 3.6

First we state another auxiliary lemma, which is proved in the next section.

Lemma 3.8 *For every $\epsilon > 0$ there exists $N = N_{3.8}(\epsilon)$ such that any graph $G = (V, E)$ with $n > N$ vertices and minimum degree at least ϵn satisfies the following: There exists a partition $\Pi = V_1, \dots, V_k$ of V such that for every $i \in [k]$ and every $u, v \in V_i$, the number of edge disjoint paths of length at most four from u to v is larger than $8^5 \log n$. Moreover, these sets can be found using a polynomial time algorithm for a fixed ϵ .*

Proof. (of Lemma 3.6) First we apply Lemma 3.8 to get the partition Π . Now the proof follows by a simple probabilistic argument. Namely, we color every edge $e \in E$ by choosing one of the colors in $\mathcal{K} = \{a_1, \dots, a_4, b_1, \dots, b_4\}$ uniformly and independently at random. Observe that a fixed path P of length at most four is an a -rainbow path with probability at least 8^{-4} . Similarly, P is

a b -rainbow path with probability at least 8^{-4} . So any fixed pair $u, v \in V_i$ is not both a -rainbow-connected and b -rainbow-connected with probability at most $2(1 - 8^{-4})^{8^5 \log n} < n^{-2}$, and therefore the probability that all such pairs are both a -rainbow connected and b -rainbow connected is strictly positive. Hence the desired coloring must exist.

To find the coloring algorithmically, we note that for every *partial* coloring of the edges of the graph it is easy to calculate the *conditional* probability that the fixed pair of vertices u, v is not both a -rainbow-connected and b -rainbow-connected. Therefore we can calculate the conditional expectation of the number of pairs that are not so connected for any partial coloring. Now we can derandomize the random selection of the coloring above by using the conditional expectation method (cf. [2]): In every stage we color one of the remaining edges in a way that does not increase the conditional expectation of the number of unconnected pairs. Since this expectation is smaller than 1 in the beginning, in the end we will have less than 1 unconnected pair, and so all pairs will be connected. ■

3.5 Proof of Lemma 3.8

Given $\epsilon > 0$ let $L = N_{3.4}(\epsilon, 1)$ and set N to be the smallest number that satisfies $\epsilon^4 \frac{N}{L} > 8^5 \log N$. Now, given any graph $G = (V, E)$ with $n > N$ vertices and minimum degree at least ϵn , we apply Lemma 3.4 with parameters ϵ and 1. Let $\Pi = V_1, V_2, \dots, V_k$ be the partition of V obtained from Lemma 3.4, while as promised, $k \leq L = N_{3.4}(\epsilon, 1)$.

Fix $i \in [k]$ and $u, v \in V_i$. From Lemma 3.4 we know that there is a component V_a such that u has at least $\frac{\epsilon}{3k}n$ neighbors in V_a . Similarly, there is a component V_b such that v has at least $\frac{\epsilon}{3k}n$ neighbors in V_b . Let $\Gamma_{u,a}$ denote the set of u 's neighbors in V_a , and similarly, let $\Gamma_{v,b}$ denote v 's neighbors in V_b . We assume in this proof that $V_a \neq V_b$, and later it will be clear that the case $V_a = V_b$ can only benefit.

We say that a set $W_u = \{w_1, \dots, w_t\} \subseteq V_i$ is *distinctly reachable from u* if there are distinct vertices $w'_1, \dots, w'_t \in \Gamma_{u,a}$ such that for every $j \in [t]$, $\{w_j, w'_j\} \in E$. Notice that the collection of pairs $\{w_j, w'_j\}$ corresponds to a matching in the graph G , where all edges of the matching have one endpoint in V_i and the other endpoint in $\Gamma_{u,a}$. Similarly, we say that $W_v \subseteq V_i$ is *distinctly reachable from v* if there are distinct vertices $w'_1, \dots, w'_t \in \Gamma_{v,b}$ such that for every $j \in [t]$, $\{w_j, w'_j\} \in E$. Observe that it is enough to prove that there exists a set $W \subseteq V_i$ of size $\epsilon^4 \frac{N}{L} > 8^5 \log N$ which is distinctly reachable from both u and v . This will imply the existence of $8^5 \log N$ edge disjoint paths of length four from u to v . Before we proceed, notice that if the same holds when $V_a = V_b$, then the number of disjoint paths from u to v remains the same while their lengths can shrink to 2 (if the corresponding subsets of $\Gamma_{u,a}$ and $\Gamma_{v,b}$ intersect), which is even better for our analysis.

Our first goal is to bound from below the size of the maximal set W_u as above. Since (by Lemma 3.4) V_a and V_i are ϵ^3 -regular pairs with density $\geq \frac{\epsilon}{16}$ and since $\epsilon^3 < \epsilon/3$, the number of edges between $\Gamma_{u,a}$ and V_i is at least $(\frac{\epsilon}{16} - \epsilon^3) |\Gamma_{u,a}| \cdot |V_i|$. Before proceeding, we make the following useful observation.

Observation 3.9 Let $H = (A, B)$ be a bipartite graph with $\gamma|A||B|$ edges. Then H contains a matching M of size $\gamma \frac{|A||B|}{|A|+|B|}$.

Proof. Consider the following process that creates M . Initially $M_0 = \emptyset$. Then in step i , we pick an arbitrary edge $\{a, b\} \in E(H)$, set $M_{i+1} = M_i \cup \{a, b\}$ and remove from $E(H)$ all the edges incident with either a or b . Clearly, in each step the number of removed edges is bounded by $|A| + |B|$, so the process continues for at least $\frac{E(H)}{|A|+|B|} = \gamma \frac{|A||B|}{|A|+|B|}$ steps. Hence $|M| = |\bigcup_i M_i| \geq \gamma \frac{|A||B|}{|A|+|B|}$. ■

Returning to the proof of Lemma 3.8, by Observation 3.9 the size of a maximal set W_u as above is at least

$$\left(\frac{\epsilon}{16} - \epsilon^3\right) \frac{|\Gamma_{u,a}||V_i|}{|\Gamma_{u,a}| + |V_i|} \geq \left(\frac{\epsilon}{16} - \epsilon^3\right) \frac{(\epsilon n/(3k))(n/k)}{\epsilon n/(3k) + n/k} \geq \frac{\epsilon^2}{100k}n.$$

To prove that $W = W_u \cap W_a$ is large, we similarly use the regularity condition, but now on the pair $(\Gamma_{v,b}, W_u)$. We get,

$$|E(\Gamma_{v,b}, W_u)| \geq \left(\frac{\epsilon}{16} - \epsilon^3\right) |\Gamma_{v,b}||W_u|.$$

Here too, by Observation 3.9 we can bound from below the size of a maximal matching in the pair $(\Gamma_{v,b}, W_u)$ with

$$\left(\frac{\epsilon}{16} - \epsilon^3\right) \frac{|\Gamma_{v,b}||W_u|}{|\Gamma_{v,b}| + |W_u|} \geq \left(\frac{\epsilon}{16} - \epsilon^3\right) \frac{\left(\frac{\epsilon}{3k}n\right)\left(\frac{\epsilon^2}{100k}n\right)}{\frac{\epsilon}{3k}n + \frac{\epsilon^2}{100k}n} \geq \epsilon^4 \frac{n}{k} \geq \epsilon^4 \frac{n}{L} > 8^5 \log N,$$

where the inequalities follow from our choice of N and the assumption that ϵ is small. Recall that the matching that we found defines the desired set W , concluding the proof. An algorithmic version of this lemma can be derived by simply using an algorithmic version of Lemma 3.4 in the selection of V_1, \dots, V_k above. ■

3.6 Graphs with diameter 2

Proof of Theorem 1.4. Consider a random 3-coloring of E , where every edge is colored with one of three possible colors uniformly and independently at random. It is enough to prove that for all pairs $u, v \in V$ the probability that they are not rainbow connected is at most $1/n^2$. Then the proof follows by the union bound (cf. [2]).

Let us fix a pair $u, v \in V$, and bound from above the probability that this pair is not rainbow connected. We know that both $\Gamma(u)$ and $\Gamma(v)$ (the neighborhoods of u and v) contain at least $8 \log n$ vertices.

1. If $\{u, v\} \in E$ then we are done.
2. If $|\Gamma(u) \cap \Gamma(v)| \geq 2 \log n$ then there are at least $2 \log n$ edge-disjoint paths of length two from u to v . In this case, the probability that none of these paths is a rainbow path is bounded by $(1/3)^{2 \log n} < 1/n^2$, and we are done.




Figure 1: Rainbow diameter is not monotone

3. Otherwise, let $A = \Gamma(u) \setminus \Gamma(v)$ and $B = \Gamma(v) \setminus \Gamma(u)$. We know that $|A|, |B| \geq 6 \log n$, and in addition, since the first two cases do not hold and the diameter of G is two, all the (length two) shortest paths from A 's vertices to v go through the vertices in B . This implies that every vertex $x \in A$ has a neighbor $b(x) \in B$ ($b(x)$ need not be a one-one function). Let us consider the set of at least $6 \log n$ edge-disjoint paths $P = \{u, x, b(x) : x \in A\}$. For each $x \in A$, the probability that $u, x, b(x), v$ is a rainbow path (given the color of the edge $(b(x), v)$) is $2/9$. Moreover, this event is independent of the corresponding events for all other members of A , because this probability does not change even with full knowledge of the colors of all edges incident with v . Therefore, the probability that none of the paths in P extends to a rainbow path from u to v is at most $(7/9)^{6 \log n} \leq 1/n^2$, as required.

The above proof immediately implies a polynomial expected time randomized algorithm with zero error probability (since we can also efficiently check if the coloring indeed makes G 3-rainbow connected). The algorithm can be derandomized and converted to a polynomial time probabilistic algorithm using the method of conditional expectations (cf. [2]) similarly to the proof of Lemma 3.6: For every partial coloring of the edges we can efficiently bound the conditional probability that a fixed pair u, v is not rainbow-connected, using the relevant one of the three cases concerning u and v that were analyzed above. Now we can color the edges one by one, at each time taking care not to increase the bound on the conditional expectation of unconnected pairs that results from the above probability bound for every u and v . Since the bound on the expectation was smaller than 1 before the beginning of the process, in the end we would get a valid 3-rainbow-coloring of G . ■

4 Concluding remarks and open problems

- Theorem 1.3 asserts that a connected graph with minimum degree at least ϵn has bounded rainbow connectivity. However, the bound obtained is huge as it follows from the Regularity Lemma. It would be interesting to find the “correct” bound. It is even possible that $rc(G) \leq C/\epsilon$ for some absolute constant C .

- The proof of Theorem 1.1 shows that deciding whether $rc(G) = 2$ is NP-Complete. Although this suffices to deduce that computing $rc(G)$ is NP-Hard, we still do not have a proof that deciding whether $rc(G) \leq k$ is NP-Complete for every fixed k . It is tempting to conjecture that for every k it is NP-Hard even to distinguish between 2-rainbow-colorable graphs and graphs that are not even k -rainbow-colorable.
- A parameter related to rainbow connectivity is the *rainbow diameter*. In this case we ask for an edge coloring so that for any two vertices, there is a rainbow *shortest* path connecting them. The rainbow diameter number, denoted $rd(G)$ is the smallest number of colors used in such a coloring. Clearly, $rd(G) \geq rc(G)$ and obviously every connected graph with n vertices has $rd(G) < \binom{n}{2}$. Unlike rainbow connectivity, which is a monotone graph property (adding edges never increases the rainbow connectivity number) this is not the case for the rainbow diameter (although we note that constructing an example that proves non-monotonicity is not straightforward). Clearly, computing $rd(G)$ is NP-Hard since $rc(G) = 2$ if and only if $rd(G) = 2$. It would be interesting to prove a version of Theorem 1.3 for rainbow diameter. We conjecture that, indeed, if G is a connected graph with minimum degree at least ϵn then it has a bounded rainbow diameter.
- Suppose that we are given a graph G for which we are *told* that $rc(G) = 2$. Can we rainbow-color it in polynomial time with $o(n)$ colors? For the usual coloring problem, this version has been well studied. It is known that if a graph is 3-colorable (in the usual sense), then there is a polynomial time algorithm that colors it with $\tilde{O}(n^{3/14})$ colors [3].

Acknowledgment We thank Van Bang Le and Zsolt Tuza for pointing out an error in one of our concluding remarks in an earlier version of this paper.

References

- [1] N. Alon, R.A. Duke, H. Lefmann, V. Rödl, and R. Yuster, *The algorithmic aspects of the Regularity Lemma*, Journal of Algorithms 16 (1994), 80–109.
- [2] N. Alon and J. H. Spencer, **The Probabilistic Method**, *Second Edition*, Wiley, New York, 2000.
- [3] A. Blum and D. Karger, *An $\tilde{O}(n^{3/14})$ -coloring algorithm for 3-colorable graphs*, Inform. Process. Lett., 61(1) : 49-53, 1997.
- [4] B. Bollobás, **Modern Graph Theory**, Graduate Texts in Mathematics 184, Springer-Verlag 1998.
- [5] G. Chartrand, G. L. Johns, K. A. McKeon, and P. Zhang, *Rainbow connection in graphs*, Math. Bohem., 133 (2008), no. 1, 85-98.

- [6] Y. Caro, A. Lev, Y. Roditty, Z. Tuza, and R. Yuster, *On rainbow connection*, The Electronic Journal of Combinatorics 15(2008), Paper R57.
- [7] E. Fischer, A. Matsliah, and A. Shapira, *Approximate Hypergraph Partitioning and Applications*, Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS) 579–589 (2007).
- [8] J. Komlós and M. Simonovits, *Szemerédi Regularity Lemma and its applications in graph theory*, In: Combinatorics, Paul Erdős is Eighty (D. Miklós, V. T. Sós, and T. Szőnyi, Eds.), Bolyai Society Mathematical Studies, Vol. 2, Budapest (1996), 295-352.
- [9] E. Szemerédi Regular partitions of graphs, *Proc. Colloque Inter. CNRS 260*, (CNRS, Paris) 399–401 (1978).