On Regularity Lemmas and their Algorithmic Applications

JACOB FOX[1][†],   LÁSZLÓ MIKLÓS LOVÁSZ[2]   and   YUFEI ZHAO[3][‡]

[1] Department of Mathematics, Stanford University, Stanford, CA 94305, USA
(e-mail: `jacobfox@stanford.edu`)

[2] Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
(e-mail: `lmlovasz@math.mit.edu`)

[3] Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK
(e-mail: `yufei.zhao@maths.ox.ac.uk`)

Szemerédi's regularity lemma and its variants are some of the most powerful tools in combinatorics. In this paper, we establish several results around the regularity lemma. First, we prove that whether or not we include the condition that the desired vertex partition in the regularity lemma is equitable has a minimal effect on the number of parts of the partition. Second, we use an algorithmic version of the (weak) Frieze–Kannan regularity lemma to give a substantially faster deterministic approximation algorithm for counting subgraphs in a graph. Previously, only an exponential dependence for the running time on the error parameter was known, and we improve it to a polynomial dependence. Third, we revisit the problem of finding an algorithmic regularity lemma, giving approximation algorithms for several co-NP-complete problems. We show how to use the weak Frieze–Kannan regularity lemma to approximate the regularity of a pair of vertex subsets. We also show how to quickly find, for each $\varepsilon' > \varepsilon$, an ε'-regular partition with k parts if there exists an ε-regular partition with k parts. Finally, we give a simple proof of the permutation regularity lemma which improves the tower-type bound on the number of parts in the previous proofs to a single exponential bound.

2010 *Mathematics subject classification*: Primary 05C85
Secondary 05C50, 05D99

1. Introduction

Szemerédi's regularity lemma [25] is one of the most powerful tools in graph theory. Szemerédi [24] used an early version in the proof of his celebrated theorem on long arithmetic progressions

in dense subsets of the integers. Roughly speaking, the regularity lemma says that every large graph can be partitioned into a small number of parts such that the bipartite subgraph between almost every pair of parts is random-like.

To state Szemerédi's regularity lemma requires some terminology. Let G be a graph, and let X and Y be (not necessarily disjoint) vertex subsets. Let $e(X,Y)$ denote the number of pairs of vertices $(x,y) \in X \times Y$ that are edges of G. The *edge density* $d(X,Y) = e(X,Y)/(|X||Y|)$ between X and Y is the fraction of pairs in $X \times Y$ that are edges. The pair (X,Y) is ε-*regular* if, for all $X' \subseteq X$ and $Y' \subseteq Y$ with $|X'| \geqslant \varepsilon|X|$ and $|Y'| \geqslant \varepsilon|Y|$, we have $|d(X',Y') - d(X,Y)| < \varepsilon$. Qualitatively, a pair of parts is ε-regular with small ε if the edge densities between pairs of large subsets are all roughly the same. A vertex partition $V = V_1 \cup \cdots \cup V_k$ is *equitable* if the parts have size as equal as possible, that is, we have $||V_i| - |V_j|| \leqslant 1$ for all i, j. An equitable vertex partition with k parts is ε-*regular* if all but εk^2 pairs of parts (V_i, V_j) are ε-regular. The regularity lemma states that for every $\varepsilon > 0$ there is a (least) integer $K(\varepsilon)$ such that every graph has an ε-regular equitable vertex partition into at most $K(\varepsilon)$ parts.

Arguably the main drawback of Szemerédi's regularity lemma is that the proof gives an enormous upper bound $K(\varepsilon)$ on the number of parts, namely an exponential tower of twos of height $O(\varepsilon^{-5})$. That such a huge bound is indeed necessary was an open problem for many years, until Gowers [15] proved a lower bound on $K(\varepsilon)$ which is an exponential tower of twos of height $\Omega(\varepsilon^{-1/16})$. Further results by Conlon and Fox [5] determine the dependence on the number of irregular pairs, and a simpler proof of Gowers' result was obtained by Moshkovitz and Shapira [23]. The first two authors [12] determine the tower height up to a constant factor in a version of the regularity lemma (see Section 2 for details). In this version, we show in Section 2 that the requirement that the partition is equitable has a negligible effect on the number of parts (up to changing the regularity parameter a little bit).

Due to the many applications of the regularity lemma, there has been a great deal of research on developing algorithmic versions of the regularity lemma and its applications (see the survey by Komlós and Simonovits [19]). We would like to be able to find an ε-regular partition of a graph on n vertices in time polynomial in n. Szemerédi's original proof of the regularity lemma was not algorithmic. The reason for this is that it needs to be able to check if a pair of parts is ε-regular, and if not, to use subsets of the parts that realize this. This is problematic because it is shown in [1] that determining whether a given pair of parts is ε-regular is co-NP-complete. They use this to show that checking whether a given partition is ε-regular is co-NP-complete.

However, Alon, Duke, Lefmann, Rödl and Yuster [1] show how to find, if a given pair of vertex subsets of size n are not ε-regular, a pair of subsets which realize that the pair is not $\varepsilon^4/16$-regular. The running time is $O_\varepsilon(n^{\omega+o(1)})$, where $\omega < 2.373$ is the matrix multiplication exponent (multiplying two $n \times n$ matrices in $n^{\omega+o(1)}$ time) [7, 20]. Here we use the subscript ε to mean that the hidden constants depend on ε. Finding a pair of subsets of vertices that detect irregularity is the key bottleneck for the algorithmic proof of the regularity lemma. It was shown [1] that one can find an ε-regular partition with the number of parts at most an exponential tower of height $O(\varepsilon^{-20})$ in an n-vertex graph in time $O_\varepsilon(n^{\omega+o(1)})$. Thus, the following surprising fact holds: while checking whether a given partition is ε-regular is co-NP-complete, finding an ε-regular partition can be done in polynomial time.

Frieze and Kannan [14] later found a simple algorithmic proof of the regularity lemma based on a spectral approach. Using expander graphs, Kohayakawa, Rödl and Thoma [18] gave a faster

algorithmic regularity lemma with optimal running time of $O_\varepsilon(n^2)$. Alon and Naor [2] develop an algorithm which approximates the cut norm of a graph within a factor 0.56 using Grothendieck's inequality and apply this to find a polynomial-time algorithm which finds, for a given pair of vertex subsets of order n which is not ε-regular, a pair of subsets which realize that the pair is not $\varepsilon^3/2$-regular. They further observe that their approach gives an improvement on the tower height in the algorithmic regularity lemma to $O(\varepsilon^{-7})$.

However, due to the tower-type dependence for the number of parts on the regularity parameter, these are not practical algorithms. While most graphs have a small regularity partition, the previous algorithmic proofs would not necessarily find it, and would only guarantee finding a regular partition with a tower-type number of parts. Addressing this issue, Fischer, Matsliah and Shapira [11] give a probabilistic algorithm which runs in constant time (depending on ε and k) and which finds, in a graph with an $\varepsilon/2$-regular partition with k parts, an ε-regular partition with at most k parts (implicitly defined). Tao [26] gives a probabilistic algorithm which, with high probability in constant time (depending on ε), produces an ε-regular partition. The algorithm takes a random sample of vertices (the exact number of which is also random) and outputs the common refinement of the neighbourhoods of these vertices.

Still, it is desirable to have a fast *deterministic* algorithm for finding a regularity partition, which we obtain here. We give several deterministic approximation algorithms for these co-NP-complete problems.

Theorem 1.1. *There exists an $O_{\varepsilon,\alpha,k}(n^2)$-time algorithm which, given $0 < \varepsilon, \alpha < 1$ and k, and a graph G on n vertices that admits an equitable ε-regular partition with k parts, outputs an equitable $(1+\alpha)\varepsilon$-regular partition of G into k parts.*

In other words, if a graph has a regular partition with few parts, then we can quickly find a regular partition (losing very slightly on the regularity) with the same number of parts. In particular, we obtain an algorithmic regularity lemma which is optimal in terms of the number of parts as it is exactly the same as in the non-algorithmic version (with a very slight loss on the regularity parameter).

We also give an approximation algorithm for checking whether a given pair of vertex subsets is ε-regular, in the sense that if the pair is not ε-regular, then we can algorithmically find a pair of vertex subsets that witness that its failure to be $(1-\alpha)\varepsilon$-regular. We will formulate this in terms of regularity of bipartite graphs. We say that bipartite graph G with bipartition (X,Y) is ε-regular if the pair (X,Y) is ε-regular.

Theorem 1.2. *There exists an $O_{\varepsilon,\alpha}(n^2)$-time algorithm which, given $\varepsilon, \alpha > 0$, and a bipartite graph G between vertex sets X and Y, each of size at most n, outputs one of the following.*

(1) *It correctly states that G is ε-regular.*
(2) *It finds a pair of vertex subsets $U \subseteq X$ and $W \subseteq Y$ which realize that G is not $(1-\alpha)\varepsilon$-regular, that is, $|U| \geqslant (1-\alpha)\varepsilon|X|$, $|W| \geqslant (1-\alpha)\varepsilon|Y|$, and $|d(U,W) - d(X,Y)| > (1-\alpha)\varepsilon$.*

Using this result, by checking the regularity of each pair of parts in a partition, we have the following corollary.

Corollary 1.3. *Given $\varepsilon, \alpha > 0$, we can distinguish in time $O_{\varepsilon,\alpha}(n^2)$ between an ε-regular partition and a partition which is not $(1-\alpha)\varepsilon$-regular.*

Remark. In Theorems 1.1 and 1.2 and Corollary 1.3, the dependence of the running time on the parameters ε, α, k may be improved at the cost of worsening the dependence on n from n^2 to $n^{\omega+o(1)}$. This is because we use the recent algorithmic version of the Frieze–Kannan weak regularity lemma due to Dellamonica, Kalyanasundaram, Martin, Rödl and Shapira [8, 9]. In the more recent paper [9], they develop an $O_\varepsilon(n^2)$ algorithm for finding a weak ε-regular partition, but it has a double exponential in $1/\varepsilon$ constant factor dependence. The older paper [8] has the advantage of not having this poor dependence on the regularity parameter. See Section 3 for more information.

Counting the number of copies of a graph H in another graph G is a famous problem in algorithmic graph theory. For example, a special case of this problem is to determine the clique number, the size of the largest clique, in a graph. This is a well-known NP-complete problem. In fact, Håstad [16] and Zuckerman [27] proved that it is NP-hard to approximate the clique number of an n-vertex graph within a factor $n^{1-\varepsilon}$ for any $\varepsilon > 0$.

There is a fast *probabilistic* algorithm for approximating up to ε the fraction of k-tuples which make a copy of H. The algorithm takes $s = 10\varepsilon^{-2}$ samples of k-tuples of vertices uniformly at random from G and outputs the fraction of them that make a copy of H. The number of copies of H is a binomial random variable with standard deviation at most $s^{1/2}/2$, and hence the fraction of k-tuples which make a copy of H in this random sample is likely within ε of the fraction of k-tuples which makes copies of H. However, this algorithm has no guarantee of success. It is therefore desirable to have a *deterministic* algorithm for counting copies which gives an approximation for the subgraph count with complete certainty.

The algorithmic regularity lemma is useful for deterministically approximating the number of copies of any fixed graph in a graph. Indeed, the counting lemma shows that if k parts $V_1, \ldots, V_k$ are pairwise regular, then the number of copies of a graph H with k vertices with the copy of the ith vertex in V_i is close to what is expected in a random graph with the same edge densities between the pairs of parts. Adding up over all k-tuples of parts in an ε-regular partition, and noting that almost all k-tuples of parts have all pairs ε-regular, we get an algorithm which runs in time $O_{\varepsilon,k}(n^2)$ and which computes the number of copies of a graph H on k vertices in a graph on n vertices, up to an additive error of εn^k. The major drawback with this result is the tower-type dependence on ε and k, which comes from the number of parts in the regularity lemma.

Duke, Lefmann and Rödl [10] gave a faster approximation algorithm for the number of copies of H in a graph G. They first develop a weak regularity lemma with an exponential dependence instead of a tower-type dependence. This gives an algorithm which runs in time $2^{(k/\varepsilon)^{O(1)}} n^{\omega+o(1)}$ and which computes the number of copies of a graph H on k vertices in a graph on n vertices, up to an additive error of εn^k.

In Section 4, we will use the algorithmic version of the Frieze–Kannan weak regularity lemma [8] to get the following even faster approximation algorithm for the subgraph counting problem. It improves the previous exponential dependence on the error parameter to a polynomial dependence. Here $v(H)$ and $e(H)$ denote the number of vertices and edges in H, respectively.

Theorem 1.4. *Let H be a graph, and let $\varepsilon > 0$ be given. There is a deterministic algorithm that runs in time $O_H(\varepsilon^{-O(1)} n^{\omega + o(1)} + \varepsilon^{-O(e(H))} n)$, and finds the number of copies of H in G up to an error of at most $\varepsilon n^{v(H)}$.*

For example, we can count the number of cliques of order 1000 in an n-vertex graph up to an additive error $n^{1000 - 10^{-6}}$ in time $O(n^{2.4})$.

In the final section of the paper, Section 7, we turn our attention to a regularity lemma for permutations. Cooper [6] proved a permutation regularity lemma which was later refined by Hoppen, Kohayakawa and Sampaio [17]. We give a new short proof of the permutation regularity lemma, improving the number of parts from tower-type to single exponential, and further extend it to an interval regularity lemma for graphs and matrices.

2. Equitable partitions with small irregularity

Let G be a graph, and let X and Y be (not necessarily disjoint) vertex subsets. The *irregularity* of the pair X, Y is defined as

$$\text{irreg}(X, Y) = \max_{U \subseteq X, W \subseteq Y} \left| e(U, W) - |U||W| d(X, Y) \right|.$$

The *irregularity* of a partition $\mathcal{P}$ of the vertex set of G is defined to be

$$\text{irreg}(\mathcal{P}) = \sum_{X, Y \in \mathcal{P}} \text{irreg}(X, Y).$$

One version of Szemerédi's regularity lemma [22, Lemma 2.2] states that given any ε, one can find an $M(\varepsilon)$ such that any graph G has a partition into at most $M(\varepsilon)$ parts with irregularity at most $\varepsilon |V|^2$. The proof of the regularity lemma gives $M(\varepsilon) \leqslant \text{tower}(O(\varepsilon^{-2}))$, and we now know that this is essentially tight [5, 12, 15], in the sense that $M(\varepsilon) = \text{tower}(\Theta(\varepsilon^{-2}))$ [12]. Here the tower function is defined by $\text{tower}(1) = 2$ and $\text{tower}(k+1) = 2^{\text{tower}(k)}$.

We say that a partition is *equitable* if any two parts differ in size by at most one. It is a convenient property to have in a regularity partition. The main result of this section shows that for any vertex partition, one can refine it a bit further to obtain a partition which is close to an equitable partition whose irregularity is not substantially larger.

Theorem 2.1. *Let $0 < \alpha < 1/2$, let m be a positive integer, and let G be a graph on $n \geqslant 10^8 m \alpha^{-5}$ vertices. If $\mathcal{P}$ is a vertex partition of G into m parts, then there is an equitable vertex partition $\mathcal{Q}$ of G into at most $4m/\alpha$ parts such that $\text{irreg}(\mathcal{Q}) \leqslant \text{irreg}(\mathcal{P}) + \alpha n^2$.*

Let $M_{eq}(\varepsilon)$ be the smallest M such that, for any graph $G = (V, E)$, there is an equitable partition into at most M parts with total irregularity at most $\varepsilon |V|^2$. We have $M_{eq}(\varepsilon) \geqslant M(\varepsilon)$ trivially. As a consequence of Theorem 2.1, we show directly that adding the condition that the partition is equitable has a very small effect on the size of the smallest partition with small irregularity.

Theorem 2.2. *Let $0 < \varepsilon < 1$ and $0 < \alpha < 1/2$. We have $M_{eq}(\varepsilon + \alpha) \leqslant \alpha^{-O(1)} M(\varepsilon)$.*

In particular, taking α small but not too small, such as $\alpha = 2^{-1/\varepsilon}$, we see that the tower height in Szemerédi's regularity lemma is not significantly affected by adding the equitability requirement.

Note that Theorem 2.1 also applies to graphs G whose number of parts in the regularity partition is not as large as the worst case $M(\varepsilon)$. To prove Theorem 2.1, we randomly divide each part of the partition $\mathcal{P}$ into parts of (essentially) equal size (apart from a small remaining subset), and then arbitrarily partition the relatively few remaining vertices into parts of equal size to obtain an equitable partition. We show that this works with high probability.

As a first step, the following lemma shows that with high probability, a pair of random subsets X', Y' of a pair of parts X, Y induces roughly the same subgraph density as X and Y.

Lemma 2.3. *Let X and Y be vertex subsets of a graph G. Let $X' \subseteq X$ and $Y' \subseteq Y$ be picked uniformly at random with $|X'| = |Y'| = k$. Then*

$$\mathbb{P}\big(|d(X',Y') - d(X,Y)| < \delta\big) \geqslant 1 - 2e^{-\delta^2 k/4}.$$

Proof. Consider picking the vertices of X' and Y' one at a time, starting with the vertices in X'. Let $Z_0, \ldots, Z_{2k}$ be the martingale where Z_i is the expected value of $e(X', Y')$ conditioned on knowing the first i vertices already chosen (this is sometimes referred to as the vertex-exposure martingale). We have $|Z_i - Z_{i-1}| \leqslant k$ as the choice of each vertex in X' and Y' changes the final $e(X', Y')$ by at most k. By the Azuma–Hoeffding inequality (see Alon and Spencer [3, Chapter 7]),

$$\mathbb{P}\big(|Z_{2k} - Z_0| \geqslant t\big) \leqslant 2e^{-t^2/(4k^3)}.$$

We have $Z_0 = k^2 d(X,Y)$ and $Z_{2k} = e(X',Y')$. Set $t = \delta k^2$. We then obtain

$$\mathbb{P}\big(|e(X',Y') - k^2 d(X,Y)| \geqslant \delta k^2\big) \leqslant 2e^{-\delta^2 k/4}.$$

Noting that $e(X',Y') = k^2 d(X',Y')$, the lemma follows. $\qquad\square$

The next lemma shows that the irregularity parameter remains roughly the same when restricted to a random, much smaller, subset of vertices. Recall that the cut metric $\mathsf{d}_\square$ between two graphs G and H on the same vertex set $V = V(G) = V(H)$ is defined by

$$\mathsf{d}_\square(G,H) := \max_{U,W \subseteq V} \frac{|e_G(U,W) - e_H(U,W)|}{|V|^2}.$$

When G and H are bipartite graphs on $V = X \cup Y$, we define the cut metric as

$$\mathsf{d}_\square(G,H) := \max_{U \subseteq X, W \subseteq Y} \frac{|e_G(U,W) - e_H(U,W)|}{|X||Y|}.$$

We also use the same notation for edge-weighted graphs, where $e(U,W)$ denotes the sum of weights of all edges in $U \times W$.

Lemma 2.4. *Let X,Y be vertex subsets of a graph G. Let $X' \subseteq X$ and $Y' \subseteq Y$ be picked uniformly at random with $|X'| = |Y'| = k$. Then, with probability at least $1 - 6e^{-\sqrt{k}/10}$,*

$$\left| \frac{\text{irreg}(X',Y')}{k^2} - \frac{\text{irreg}(X,Y)}{|X||Y|} \right| \leqslant \frac{9}{k^{1/4}}, \tag{2.1}$$

Proof. We use the so-called First Sampling Lemma [4, Theorem 2.10] (we quote the statement from [21, Lemma 10.5]): if G and H are weighted graphs with $V(G) = V(H)$ and edge weights in $[0,1]$, and $S \subseteq V(G)$ is chosen uniformly at random with $|S| = k$, then with probability at least $1 - 4e^{-\sqrt{k}/10}$,

$$|\mathsf{d}_\square(G[S], H[S]) - \mathsf{d}_\square(G, H)| \leqslant \frac{8}{k^{1/4}}.$$

Let $G[X,Y]$ denote the bipartite (weighted) graph with vertex sets X and Y, and whose edges are induced from G. A bipartite version of this sampling lemma holds true, that with probability at least $1 - 4e^{-\sqrt{k}/10}$,

$$|\mathsf{d}_\square(G[X',Y'], H[X',Y']) - \mathsf{d}_\square(G[X,Y], H[X,Y])| \leqslant \frac{8}{k^{1/4}}, \tag{2.2}$$

and its proof is nearly identical to the first version stated above. Note that

$$\text{irreg}(X,Y) = |X||Y|\mathsf{d}_\square(G[X,Y], d(X,Y)),$$

where the second argument denotes the complete graph with loops with all edge weights equal to $d(X,Y)$. Similarly,

$$\text{irreg}(X',Y') = k^2\mathsf{d}_\square(G[X',Y'], d(X',Y')).$$

By letting H in (2.2) be the complete graph with loops and all edge weights $d(X,Y)$, we obtain that with probability at least $1 - 4e^{-\sqrt{k}/10}$,

$$\left| \frac{\text{irreg}(X',Y')}{k^2} - \frac{\text{irreg}(X,Y)}{|X||Y|} \right| \leqslant \frac{8}{k^{1/4}} + |d(X,Y) - d(X',Y')|.$$

We then apply Lemma 2.3 with $\delta = 1/k^{1/4}$ to reach the desired conclusion. $\square$

As a corollary of Lemma 2.3, and noting that the left-hand side of (2.1) is at most 1, we have

$$\mathbb{E}\left(\frac{\text{irreg}(X',Y')}{k^2} \right) \leqslant \frac{\text{irreg}(X,Y)}{|X||Y|} + \frac{9}{k^{1/4}} + 6e^{-\sqrt{k}/10} \leqslant \frac{\text{irreg}(X,Y)}{|X||Y|} + \frac{20}{k^{1/4}}. \tag{2.3}$$

Proof of Theorem 2.1. We shall omit floors and ceilings for the sake of clarity of presentation. Let $k = \alpha n/(4m)$. Let $V_1, \ldots, V_m$ be the parts of $\mathcal{P}$. Uniformly at random partition each V_i into parts of size k, with possibly one remainder part of size less than k. Call the resulting partition $\mathcal{P}'$.

By (2.3),[†]

$$\mathbb{E}\left(\sum_{\substack{X',Y'\in\mathcal{P}' \\ |X'|=|Y'|=k}} \mathrm{irreg}(X',Y')\right) \leqslant \sum_{X,Y\in\mathcal{P}}\left(\mathrm{irreg}(X,Y)+\frac{20}{k^{1/4}}|X||Y|\right)$$

$$\leqslant \mathrm{irreg}(\mathcal{P})+\frac{20n^2}{k^{1/4}}.$$

So there exists some such partition $\mathcal{P}'$ such that

$$\sum_{\substack{X',Y'\in\mathcal{P}' \\ |X'|=|Y'|=k}} \mathrm{irreg}(X',Y') \leqslant \mathrm{irreg}(\mathcal{P})+\frac{20n^2}{k^{1/4}}. \tag{2.4}$$

Fix $\mathcal{P}'$ to be this partition.

Let S be the union of the parts of $\mathcal{P}'$ of size less than k. There is at most one such part of $\mathcal{P}'$ for each V_i, so $|S| < mk$. Arbitrarily partition S into sets of size k, and let $\mathcal{Q}$ be the equitable vertex partition consisting of these parts of S along with the parts of $\mathcal{P}'$ of size k.

Parts arising from S contribute at most $2|S|n < 2mkn$ to $\mathrm{irreg}(\mathcal{Q})$, whereas the other contributions to $\mathrm{irreg}(\mathcal{Q})$ are bounded by (2.4). Thus

$$\mathrm{irreg}(\mathcal{Q}) \leqslant \mathrm{irreg}(\mathcal{P})+\frac{20n^2}{k^{1/4}}+2mkn \leqslant \mathrm{irreg}(\mathcal{P})+\alpha n^2,$$

where the last step follows from

$$\frac{20}{k^{1/4}} = \frac{20(4m)^{1/4}}{(\alpha n)^{1/4}} \leqslant \frac{20(4m)^{1/4}}{(\alpha\cdot 10^8 m\alpha^{-5})^{1/4}} < \frac{\alpha}{2}$$

and

$$\frac{2mk}{n} = \frac{2m}{n}\frac{\alpha n}{4m} \leqslant \frac{\alpha}{2}.$$

Therefore $\mathcal{Q}$ is the required equipartition.   □

As a consequence of Theorem 2.1, we can prove Theorem 2.2.

Proof of Theorem 2.2.   Let G be a graph with a partition $\mathcal{P}$ of the vertex set into $m \leqslant M(\varepsilon)$ parts with irregularity at most εn^2. If $n < 10^8 m\alpha^{-5}$, then we just partition the vertices into singleton sets, which have zero irregularity, using at most $10^8 m\alpha^{-5} \leqslant 10^8\alpha^{-5}M(\varepsilon)$ parts. Otherwise apply Theorem 2.1 to obtain a partition with irregularity at most $(\varepsilon + \alpha)n^2$, and at most $4m/\alpha \leqslant 4\alpha^{-1}M(\varepsilon)$ parts.   □

3. Algorithmic weak regularity

In this section we review some results concerning algorithmic versions of the Frieze–Kannan weak regularity lemma. We will be applying these results in subsequent sections.

[†] It is easy to modify the proof to address the case when X' and Y' are within the same part of $\mathcal{P}$.

Given any edge-weighted graph G and any partition $\mathcal{P}\colon V = V_1 \cup V_2 \cup \cdots \cup V_t$ of the vertex set of G into t parts, let $G_{\mathcal{P}}$ denote the weighted graph with vertex set V obtained by giving weight $d_{ij} := d(V_i, V_j) = e(V_i, V_j)/(|V_i||V_j|)$ to all pairs of vertices in $V_i \times V_j$, for every $1 \leqslant i \leqslant j \leqslant t$. We say $\mathcal{P}$ is an ε-*regular Frieze–Kannan* (or ε-*FK-regular*) partition if $\mathsf{d}_{\square}(G, G_{\mathcal{P}}) \leqslant \varepsilon$. In other words, $\mathcal{P}$ is an ε-regular Frieze–Kannan partition if

$$\left| e(S,T) - \sum_{i,j=1}^{t} d_{ij}|S \cap V_i||T \cap V_j| \right| \leqslant \varepsilon |V|^2. \tag{3.1}$$

for all $S, T \subseteq V$. We say that sets S and T *witness* that $\mathcal{P}$ is not ε-FK-regular if the above inequality is violated.

Frieze and Kannan [13] proved the following regularity lemma.

Theorem 3.1 (Frieze–Kannan). *Let $\varepsilon > 0$. Every graph has an ε-regular Frieze–Kannan partition with at most $2^{2/\varepsilon^2}$ parts.*

We apply efficient deterministic algorithms for generating Frieze–Kannan regular partitions. Such algorithms were recently given in [8, 9]. Specifically, Dellamonica, Kalyanasundaram, Martin, Rödl and Shapira [8] gave an $\varepsilon^{-6} n^{\omega + o(1)}$-time algorithm to generate an equitable ε-regular Frieze–Kannan partition of a graph on n vertices into at most $2^{O(\varepsilon^{-7})}$ parts. Recall that $\omega < 2.373$ is the matrix multiplication exponent. In [9] they gave a different algorithm which improved the dependence of the running time on n from $O_{\varepsilon}(n^{\omega + o(1)})$ to $O_{\varepsilon}(n^2)$, while sacrificing the dependence of ε. Namely, it was shown that there is a deterministic algorithm that finds, in $O(2^{2^{\varepsilon^{-O(1)}}} n^2)$ time, an ε-regular Frieze–Kannan partition into at most $2^{\varepsilon^{-O(1)}}$ parts. It remains an open problem to improve the dependence on ε in the running time.

The proof of the Frieze–Kannan regularity lemma and its algorithmic versions, roughly speaking, run as follows.

- Given a partition (starting with the trivial partition with one part), either it is ε-FK-regular (in which case we are done), or we can exhibit some pair of subsets S, T of vertices that witness the irregularity by violating (3.1) (in the algorithmic versions, one may only be guaranteed to find S and T that violate (3.1) for some smaller value of ε).
- Refine the partition by using S and T to split each part into at most four parts, thereby increasing the total number of parts by a factor of at most 4.
- Repeat. Use a mean-square-density increment argument to upper-bound the number of possible iterations.

Remark. As in the case of the usual regularity lemma, it is possible to obtain an equitable partition in the Frieze–Kannan regularity lemma, increasing the number of parts by a small factor. We will not need this for our algorithm, however.

For example, in the algorithmic version [8], the first step (also the key step) is given as the following result [8, Corollary 3.1].

Theorem 3.2. *There is an $n^{\omega + o(1)}$-time algorithm which, given $\varepsilon > 0$, an n-vertex graph G and a partition $\mathcal{P}$ of $V(G)$, does one of the following.*

(1) *It correctly states that $\mathcal{P}$ is ε-FK-regular.*
(2) *It finds sets S, T which witness the fact that $\mathcal{P}$ is not $\varepsilon^3/1000$-FK-regular.*

In [9], an alternative theorem is given for finding an irregular pair. The statement below is a consequence of [9, Theorem 8.1].

Theorem 3.3. *There is an $O(2^k n^2)$-time algorithm which, given $\varepsilon > 0$, an n-vertex graph G and a partition $\mathcal{P}$ of $V(G)$ into k parts, does one of the following.*
 (1) *It correctly states that $\mathcal{P}$ is ε-FK-regular.*
 (2) *It finds sets S, T which witness the fact that $\mathcal{P}$ is not $\varepsilon^{O(1)}$-FK-regular.*

There is a variant of the weak regularity lemma, where the final output is not a partition of V into $2^{\varepsilon^{-O(1)}}$ parts, but rather an approximation of the graphs as a sum of $\varepsilon^{-O(1)}$ complete bipartite graphs each assigned some weight. Below we use *weighted graph* to mean a graph with edge weights. For $S, T \subseteq V$, by $K_{S,T}$ we mean the weighted graph where an edge $\{s,t\}$ has weight 1 if $s \in S$ and $t \in T$ (and weight 2 if $s, t \in S \cap T$) and weight zero otherwise. For any $c \in \mathbb{R}$, by cG we mean the weighted graph obtained from G by multiplying every edge weight by c. For a pair of weighted graphs G_1, G_2 on the same set of vertices, we will use the notation $G_1 + G_2$ to denote the graph on the same vertex set with edge weights summed (and weight 0 corresponding to not having an edge). Additionally, we write c to mean the constant graph with all edge weights equal to c.

Theorem 3.4 (Frieze–Kannan). *Let $\varepsilon > 0$. Let G be any weighted graph with $[-1,1]$-valued edge weights. There exists some $k \leqslant O(\varepsilon^{-2})$, subsets $S_1, \ldots, S_k, T_1, \ldots, T_k \subseteq V$, and $c_1, \ldots, c_k \in [-1,1]$, so that*

$$\mathsf{d}_\square(G, d(G) + c_1 K_{S_1, T_1} + \cdots + c_k K_{S_k, T_k}) \leqslant \varepsilon.$$

See [22, Lemma 4.1] for a proof (given there in the more general setting of arbitrary Hilbert spaces). Roughly speaking, to find the appropriate S_i, T_i, c_i, in the second step of the above outline of the proof of the weak regularity lemma, instead of using S and T to refine the existing partition, we subtract $cK_{S,T}$ from the remaining weighted graph (starting with G), where c is the density between S and T in the remaining weighted graph. We record the corresponding S_i, T_i, c_i in step i of this iteration. We can bound the number of iterations by observing that the L^2-norm of $G - d(G) - c_1 K_{S_1, T_1} - \cdots - c_i K_{S_i, T_i}$ must decrease by a certain amount at each step.

As for the algorithmic version, using Theorems 3.2 and 3.3 (or minor modifications thereof), we can efficiently approximate G as a weighted sum of $\varepsilon^{-O(1)}$ complete bipartite graphs.

Corollary 3.5. *There exists a $\min\{\varepsilon^{-O(1)} n^{\omega + o(1)}, O(2^{2^{\varepsilon^{-O(1)}}} n^2)\}$-time algorithm which, given $\varepsilon > 0$ and an n-vertex graph G, outputs subsets $S_1, \ldots, S_k, T_1, \ldots, T_k \subseteq V(G)$ and real numbers $c_1, \ldots, c_k$, for some $k \leqslant \varepsilon^{-O(1)}$, such that*

$$\mathsf{d}_\square(G, d(G) + c_1 K_{S_1, T_1} + \cdots + c_k K_{S_k, T_k}) \leqslant \varepsilon.$$

We will use a variation for bipartite graphs.

Corollary 3.6. *There exists a* $\min\{\varepsilon^{-O(1)}n^{\omega+o(1)}, O(2^{2^{\varepsilon^{-O(1)}}}n^2)\}$*-time algorithm which, given* $\varepsilon > 0$ *and a bipartite graph G between vertex sets X and Y each with at most n vertices, outputs subsets* $S_1, \ldots, S_k \subseteq X$ *and* $T_1, \ldots, T_k \subseteq Y$ *and real numbers* $c_1, \ldots, c_k$, *for some* $k \leqslant \varepsilon^{-O(1)}$, *such that*

$$\mathsf{d}_\square(G, d(X,Y) + c_1 K_{S_1,T_1} + \cdots + c_k K_{S_k,T_k}) \leqslant \varepsilon,$$

where the constant $d(X,Y)$ *here denotes the weighted complete bipartite graph* $d(X,Y)K_{X,Y}$.

4. Approximation algorithm for subgraph counts

Suppose that we are given a graph G on n vertices, and we would like to find the number of copies of a small graph H on k vertices that are contained in G. We would like to count them up to an error at most εn^k. In this section we will provide a deterministic algorithm that can do so. Specifically, we prove Theorem 1.4, reproduced below for convenience.

Theorem 1.4. *Let H be a graph, and let* $\varepsilon > 0$ *be given. There is a deterministic algorithm that runs in time* $O_H(\varepsilon^{-O(1)}n^{\omega+o(1)} + \varepsilon^{-O(e(H))}n)$, *and finds the number of copies of H in G up to an error of at most* $\varepsilon n^{v(H)}$.

It will be cleaner to work instead with $\hom(H,G)$, the number of graph homomorphisms from H to G. This quantity differs from the number of (labelled) copies of H in G by a negligible $O(n^{v(H)-1})$ additive error (the hidden constants here and onward may depend on H). We extend the definition of $\hom(H,G)$ to edge-weighted graphs G: if the edge xy in G has weight $G(x,y)$, then we define

$$\hom(H,G) = \sum_{f:V(H)\to V(G)} \prod_{\{u,v\}\in E(H)} G(f(u),f(v)).$$

Note that here $G(x,y)$ is defined on all pairs, with $G(x,y) = 0$ if there is no edge between x and y.

The idea is to apply a weak regularity lemma in the form of Corollary 3.5. A weakly regular approximation also gives an approximate H-count, via a standard counting lemma (see [21, Lemma 10.22]).

Lemma 4.1 (counting lemma). *Given any graph H and any two weighted graphs* G_1 *and* G_2 *on the same set V of n vertices, we have*

$$|\hom(H,G_1) - \hom(H,G_2)| \leqslant e(H)\mathsf{d}_\square(G_1,G_2)n^{v(H)}.$$

Here is the algorithm. Apply Corollary 3.5 to find any approximation

$$G' = d(G) + c_1 K_{S_1,T_1} + \cdots + c_k K_{S_k,T_k}$$

of G with $\mathsf{d}_\square(G,G') \leqslant \varepsilon/e(H)$ and $k \leqslant \varepsilon^{-O(1)}$. By the counting lemma, it suffices to compute

$$\hom(H,G') = \hom(H, d(G) + c_1 K_{S_1,T_1} + \cdots + c_k K_{S_k,T_k}), \tag{4.1}$$

which can be done in $O_H(k^{e(H)}n)$ time, as follows. We can expand the right-hand side of (4.1) via the distributive property, writing

$$\hom(H, G') = \sum_{\varphi\colon E(H)\to\{0,\dots,k\}} \hom^\varphi(H, (d(G), c_1 K_{S_1, T_1}, \dots, c_k K_{S_k, T_k})), \qquad (4.2)$$

where for each assignment $\varphi\colon E(H) \to \{0, \dots, k\}$ of edges of H to the components of G' we write $\hom^\varphi(H, (G_0, G_1, \dots, G_k))$ to mean the homomorphism count obtained where the image of each edge $e \in E(H)$ is restricted to $G_{\varphi(e)}$, that is,

$$\hom^\varphi(H, (G_0, G_1, \dots, G_k)) = \sum_{f\colon V(H)\to[n]} \prod_{uv\in E(H)} G_{\varphi(uv)}(f(u), f(v)).$$

Here by $G_i(x, y)$ we mean the edge weight of (x, y) in G_i.

There are $(k+1)^{e(H)}$ possible maps φ. We claim that each term on the right-hand side of (4.2), corresponding to some φ, can be exactly computed in $O_H(n)$ time. Taking out a constant factor, it remains to compute the value of

$$\hom^\varphi(H, (1, K_{S_1, T_1}, \dots, K_{S_k, T_k})).$$

We further decompose each K_{S_i, T_i} (viewed as an adjacency matrix) as a sum $1_{S_i \times T_i} + 1_{T_i \times S_i}$ and apply the distributive property once again to expand the quantity as a sum of $2^{e(H)}$-terms. Each term counts the number of maps $f\colon V(G) \to [n]$ such that, for every $v \in V(H)$, $f(v) \in \bigcap_{e \in E(H)} R_{\varphi(e)}$ for some choice of $R_{\varphi(e)} = S_{\varphi(e)}$ or $T_{\phi(e)}$. The size of such an intersection can be computed in $O_H(n)$ time, and we can compute this term (one of $2^{e(H)}$-terms) by multiplying over all $v \in V(H)$. There are $2^{e(H)}$ choices for which summand in $1_{S_i \times T_i} + 1_{T_i \times S_i}$ to take in the expansion over all i, and by summing over all $2^{e(H)}$ choices, we can evaluate $\hom^\varphi(H, (1, K_{S_1, T_1}, \dots, K_{S_k, T_k}))$. By summing over all φ in (4.2), we see that $\hom(H, G')$ can be exactly computed in $O_H(k^{e(H)}n)$ time, thereby providing the desired approximation to $\hom(H, G)$.

As for the running time, it took $\varepsilon^{-O(1)} n^{\omega + o(1)}$ time to find the approximation G', and it took $O_H(k^{e(H)}n) = O_H(\varepsilon^{-O(e(H))}n)$ time to compute $\hom(H, G')$, giving the claimed total running time.

5. Finding an irregular pair

In this section we prove Theorem 1.2, reproduced below for convenience.

Theorem 1.2.   *There exists an $O_{\varepsilon, \alpha}(n^2)$-time algorithm which, given $\varepsilon, \alpha > 0$, and a bipartite graph G between vertex sets X and Y, each of size at most n, outputs one of the following.*

(1) *It correctly states that G is ε-regular.*
(2) *It finds a pair of vertex subsets $U \subseteq X$ and $W \subseteq Y$ which realize that G is not $(1-\alpha)\varepsilon$-regular, that is, $|U| \geqslant (1-\alpha)\varepsilon|X|$, $|W| \geqslant (1-\alpha)\varepsilon|Y|$, and $|d(U, W) - d(X, Y)| > (1-\alpha)\varepsilon$.*

We can assume that $\alpha < 1/2$, since for larger α we can just apply the algorithm with a lower value of α. We shall give an $O(2^{2^{(\alpha\varepsilon)^{-O(1)}}} n^2)$-time algorithm. Using Corollary 3.6, we approximate G by $G' = d(G) + c_1 K_{S_1, T_1} + \cdots + c_k K_{S_k, T_k}$ so that $k = (\alpha\varepsilon)^{-O(1)}$ and $\mathsf{d}_\square(G, G') \leqslant \alpha\varepsilon^3/4$. Here

$S_1, \ldots, S_k \subseteq X$ and $T_1, \ldots, T_k \subseteq Y$. We shall assume that k is small compared to $|X|$ and $|Y|$, namely,

$$100 \cdot k 2^k \leqslant \alpha \varepsilon^3 \min\{|X|, |Y|\}, \tag{5.1}$$

for otherwise we can accomplish the task by a complete search (say when $|X| \leqslant |Y|$) over all subsets of X in $2^{O(|X|)} = 2^{O(\alpha^{-1} \varepsilon^{-3} k 2^k)} = 2^{2^{(\alpha \varepsilon)^{-O(1)}}}$ time, which is enough.

We say that a sequence of numbers $u, u_1, \ldots, u_k, w, w_1, \ldots, w_k$ is *feasible* if there exists a function $\mu : X \cup Y \to [0, 1]$ (we write $\mu(S) = \sum_{x \in S} \mu(x)$ from now on) such that the quantities

$$\frac{|\mu(X) - u|}{|X|}, \quad \frac{|\mu(Y) - w|}{|Y|}, \quad \frac{|\mu(S_i) - u_i|}{|X|}, \quad \frac{|\mu(T_i) - t_i|}{|Y|}, \quad \text{for all } 1 \leqslant i \leqslant k,$$

are each at most $\alpha \varepsilon^3 / (100k)$. One can think of μ as representing subsets $U \subseteq X$ and $W \subseteq Y$ with $[0, 1]$-valued weights attached to its elements. One can determine via a linear program if a given sequence is feasible (see Lemma 5.1 below).

Here is the algorithm. We perform a complete search through all sequences

$$u, u_1, \ldots, u_k, w, w_1, \ldots, w_k$$

of non-negative integers at most n, where u and each u_i are divisible by $\lfloor \alpha \varepsilon^3 |X| / (100k) \rfloor$, and w and each w_i are divisible by $\lfloor \alpha \varepsilon^3 |Y| / (100k) \rfloor$. For each such sequence, we check whether it is feasible, and if so we then check whether the inequalities

$$\left| \sum_{i=1}^{k} c_i u_i w_i \right| > (1 - \alpha/2) \varepsilon u w, \quad u \geqslant (1 - \alpha/2) \varepsilon |X| \quad \text{and} \quad w \geqslant (1 - \alpha/2) \varepsilon |Y| \tag{5.2}$$

hold. If they never hold for any feasible sequence, then we state that G is ε-regular. On the other hand, if they hold for some feasible sequence, then we can convert f into actual sets U and W (as we shall explain) that witness that G is not $(1 - \alpha)\varepsilon$-regular.

Next we prove the correctness of the algorithm if the output is that G is ε-regular.

Consider the partition of X given by the common refinement by $S_1, \ldots, S_k$. For any index set $I \subseteq [k]$, let

$$S_I = \left(\bigcap_{i \in I} S_i \right) \cap \left(\bigcap_{i \notin I} (X \setminus S_i) \right)$$

denote the part in the common refinement indexed by I. We can compute the sizes $|S_I|$ for all $I \subseteq [k]$ in $O(2^k n)$ time. With this information at hand, we obtain the following result.

Lemma 5.1. *There exists a $2^{O(k)}$-time algorithm that determines whether a given sequence u, $u_1, \ldots, u_k$, w, $w_1, \ldots, w_k$ is feasible.*

Proof. It suffices to show that one can determine in the required time whether there exists $\mu : X \to [0, 1]$ such that $|\mu(X) - u| \leqslant a$ and $|\mu(S_i) - u_i| \leqslant a_i$, for each i. Here

$$a = a_i = \lfloor \alpha \varepsilon^3 n / (100k) \rfloor$$

is the required bound (though it could be chosen arbitrarily for the purpose of this lemma). The situation for Y is analogous.

For the purpose of satisfying the inequalities $|\mu(X) - u| \leqslant a$ and $|\mu(S_i) - u_i| \leqslant a_i$, one only needs to know the sum of values of μ on parts in the partition of X induced by the common refinement of $S_1, \ldots, S_k$.

For each $I \subseteq [k]$, the variable x_I is supposed to correspond to the value of $\mu(S_I)$. Then μ exists if and only if there exists $(x_I)_{I \subseteq [k]} \in \mathbb{R}^{2^k}$ satisfying the following inequalities:

$$-a \leqslant \left(\sum_{I \subseteq [k]} x_I \right) - u \leqslant a,$$

$$-a_i \leqslant \left(\sum_{I \ni i} x_I \right) - u_i \leqslant a_i \quad \text{for all } i \in [k],$$

and $\quad 0 \leqslant x_I \leqslant |S_I| \quad$ for all $I \subseteq [k]$.

This is a linear program in $2^k + 1$ variables, which can be solved in $2^{O(k)}$ time. The original sequence is feasible if and only if the above system of linear inequalities has some solution in (x_I). □

Suppose the algorithm does not find any feasible sequence satisfying (5.2). We claim that G is ε-regular. Assume otherwise. Then there exist $U \subseteq X$ and $W \subseteq Y$ such that $|U| \geqslant \varepsilon|X|$, $|W| \geqslant \varepsilon|Y|$, and $|d(U,W) - d(X,Y)| > \varepsilon$. Since $\mathsf{d}_{\square}(G,G') \leqslant \alpha\varepsilon^3/4$, we have $|e_G(U,W) - e_{G'}(U,W)| \leqslant (\alpha\varepsilon^3/4)|X||Y|$. Thus

$$
\begin{aligned}
|e_{G'}(U,W) - d_G(X,Y)|U||W|| &\geqslant |e_G(U,W) - d_G(X,Y)|U||W|| - |e_G(U,W) - e_{G'}(U,W)| \\
&\geqslant |d_G(U,W) - d_G(X,Y)||U||W| - \frac{1}{4}\alpha\varepsilon^3|X||Y| \\
&\geqslant \varepsilon|U||W| - \frac{1}{4}\alpha\varepsilon|U||W| \\
&\geqslant \left(1 - \frac{1}{4}\alpha\right)\varepsilon|U||W|.
\end{aligned}
$$

On the other hand, since $G' = d_G(X,Y) + c_1 K_{S_1,T_1} + \cdots + c_k K_{S_k,T_k}$, we have

$$e_{G'}(U,W) - d_G(X,Y)|U||W| = \sum_{i=1}^{k} c_i |U \cap S_i||W \cap T_i|.$$

So

$$\left| \sum_{i=1}^{k} c_i |U \cap S_i||W \cap T_i| \right| \geqslant \left(1 - \frac{1}{4}\alpha\right)\varepsilon|U||W|.$$

Let u and u_i be $|U|$ and $|U \cap S_i|$, each respectively rounded to the nearest integer multiple of $\lfloor \alpha\varepsilon^3|X|/(100k) \rfloor$, for all $1 \leqslant i \leqslant k$. Similarly let w, w_i be $|W|$ and $|W \cap S_i|$, each respectively rounded to the nearest integer multiple of $\lfloor \alpha\varepsilon^3|Y|/(100k) \rfloor$, for all $1 \leqslant i \leqslant k$. The sequence $u, u_1, \ldots, u_k, w, w_1, \ldots, w_k$ is feasible as witnessed by $\mu = 1_{U \cup W}$. We claim that (5.2) holds. Indeed, we have

$$u \geqslant |U| - \frac{1}{100}\alpha\varepsilon^3|X| \geqslant \left(1 - \frac{1}{100}\alpha\varepsilon^2\right)\varepsilon|X|,$$

and

$$w \geqslant |W| - \frac{1}{100}\alpha\varepsilon^3|Y| \geqslant \left(1 - \frac{1}{100}\alpha\varepsilon^2\right)\varepsilon|Y|.$$

Furthermore, we have

$$
\begin{aligned}
\left|\sum_{i=1}^{k} c_i u_i w_i\right| &\geqslant \left|\sum_{i=1}^{k} c_i |U \cap S_i||W \cap T_i|\right| - \frac{3}{100}\alpha\varepsilon^3|X||Y| \\
&\geqslant \left(1 - \frac{1}{4}\alpha\right)\varepsilon|U||W| - \frac{3}{100}\alpha\varepsilon^3|X||Y| \\
&\geqslant \left(1 - \frac{1}{4}\alpha - \frac{3}{100}\alpha\right)\varepsilon|U||W| \\
&\geqslant \left(1 - \frac{1}{4}\alpha - \frac{3}{100}\alpha\right)\varepsilon\left(1 + \frac{1}{100}\alpha\varepsilon^2\right)^{-2} uw \\
&> \left(1 - \frac{1}{2}\alpha\right)\varepsilon uw.
\end{aligned}
$$

The first inequality above follows from the fact that for each i,

$$|u_i - |U \cap S_i|| \leqslant \frac{\alpha\varepsilon^3|X|}{100k},$$

$$|u_i - |U \cap S_i|| \leqslant \frac{\alpha\varepsilon^3|Y|}{100k},$$

and thus

$$|u_i w_i - |U \cap S_i||W \cap T_i|| \leqslant \frac{3\alpha\varepsilon^3|X||Y|}{100k}.$$

The penultimate inequality follows from

$$u \leqslant |U| + \frac{1}{100}\alpha\varepsilon^3|X| \leqslant \left(1 + \frac{1}{100}\alpha\varepsilon^2\right)|U|$$

and similarly with w. So we have a feasible sequence satisfying (5.2), which is a contradiction.

Now suppose instead that the algorithm does find some feasible sequence that satisfies (5.2). By adjusting μ, we may assume that μ takes $\{0, 1\}$-value on all but at most one element in each part in the common refinement partition of X by $S_1, \ldots, S_k$, and likewise in Y by $T_1, \ldots, T_k$. Let $U \subseteq X$ and $W \subseteq Y$ denote the elements where μ is positive. Then we have

$$||U| - u| \leqslant \frac{\alpha\varepsilon^3}{100k}|X| + 2^k \leqslant \frac{\alpha\varepsilon^3}{50k}|X|.$$

Here the extra 2^k-term accounts for rounding up non-integral values of μ. We used assumption (5.1) to bound 2^k. It thus follows from the above and (5.2) that

$$|U| \geqslant \left(1 - \frac{1}{2}\alpha - \frac{1}{50k}\alpha\varepsilon^2\right)\varepsilon|X| \geqslant (1 - \alpha)\varepsilon|X|.$$

In particular, this means that

$$||U| - u| \leqslant \frac{\alpha\varepsilon^3}{50k}|X| \leqslant \frac{\alpha\varepsilon^2}{(1-\alpha)50k}|U|.$$

Similarly, we have

$$|W| \geqslant \left(1 - \frac{1}{2}\alpha - \frac{1}{50k}\alpha\varepsilon^2\right)\varepsilon|Y| \geqslant (1-\alpha)\varepsilon|Y|,$$

along with

$$||U \cap S_i| - u_i| \leqslant \frac{\alpha\varepsilon^2}{(1-\alpha)50k}|U|, \quad \text{for all } 1 \leqslant i \leqslant k,$$

and

$$||W| - w| \leqslant \frac{\alpha\varepsilon^2}{(1-\alpha)50k}|W|, \quad \text{and} \quad ||W \cap T_i| - w_i| \leqslant \frac{\alpha\varepsilon^2}{(1-\alpha)50k}|W|, \quad \text{for all } 1 \leqslant i \leqslant k.$$

Further,

$$\begin{aligned}
|d_G(U,W) - d_G(X,Y)| &\geqslant |d_{G'}(U,W) - d_G(X,Y)| - |d_G(U,W) - d_{G'}(U,W)| \\
&\geqslant \frac{1}{|U||W|}\left|\sum_{i=1}^{k} c_i|U \cap S_i||W \cap T_i|\right| - \frac{|X||Y|}{|U||W|}\mathsf{d}_\square(G,G') \\
&\geqslant \frac{1}{|U||W|}\left(\left|\sum_{i=1}^{k} c_i u_i w_i\right| - \frac{3}{(1-\alpha)50}\alpha\varepsilon^2|U||W|\right) - \frac{1}{4}\alpha\varepsilon^3\frac{|X||Y|}{|U||W|} \\
&\geqslant \frac{1}{|U||W|}\left(\left(1 - \frac{\alpha}{2}\right)\varepsilon u w - \frac{3}{(1-\alpha)50}\alpha\varepsilon^2|U||W|\right) - \frac{1}{4}\alpha\varepsilon \\
&\geqslant (1-\alpha)\varepsilon.
\end{aligned}$$

Hence the pair (U,W) witnesses that G is not $(1-\alpha)\varepsilon$-regular.

We will need the following easy corollary of Theorem 1.2 for the next section.

Corollary 5.2. *There exists an $O_{\varepsilon,\alpha,k}(n^2)$-time algorithm which, given $\varepsilon, \alpha, k > 0$, a graph G on n vertices, and a partition $\mathcal{P}$ of the vertex set of G into k parts, does one of the following.*

(1) It correctly states that $\mathcal{P}$ is $(1+\alpha)\varepsilon$-regular.
(2) It correctly states that $\mathcal{P}$ is not ε-regular.

Sometimes both options are correct. The algorithm that we give runs in $O(k^2 2^{2^{(\alpha\varepsilon)^{-O(1)}}}n^2)$ time.

Proof. Let $\mathcal{P}$ be the partition of V into $V_1, \ldots, V_k$. Apply the algorithm in Theorem 1.2 to each pair V_i, V_j so that either it correctly states that (V_i, V_j) is $(1+\alpha)\varepsilon$-regular or it is not ε-regular. If at least a $(1-\varepsilon)$-fraction of pairs is seen to be $(1+\alpha)\varepsilon$-regular, then we know that $\mathcal{P}$ is $(1+\alpha)\varepsilon$-regular. Otherwise, more than an ε-fraction of pairs fails to be ε-regular, so $\mathcal{P}$ is not ε-regular. $\square$

6. Approximating regularity

In this section we prove Theorem 1.1, reproduced below for convenience.

Theorem 1.1. *There exists an $O_{\varepsilon,\alpha,k}(n^2)$-time algorithm which, given $0 < \varepsilon, \alpha < 1$ and k, and a graph G on n vertices that admits an equitable ε-regular partition with k parts, outputs an equitable $(1+\alpha)\varepsilon$-regular partition of G into k parts.*

Here is the algorithm, which runs in $O(2^{2^{(k/(\alpha\varepsilon))^{O(1)}}} n^2)$ time. Using Corollary 3.5, we find $S_1, \ldots, S_s, T_1, \ldots, T_s \subseteq V$, with $s \leqslant (k/\alpha\varepsilon)^{O(1)}$, such that $\mathsf{d}_\square(G,G') \leqslant \alpha\varepsilon/(10k^2)$, where

$$G' = d(G) + c_1 K_{S_1,T_1} + \cdots + c_k K_{S_s,T_s}.$$

Let $\mathcal{Q}$ denote the partition of $V(G)$ given by the common refinement of the sets

$$S_1, \ldots, S_s, T_1, \ldots, T_s.$$

Let $\mathcal{Q}$ have $r \leqslant 4^s$ parts, with sizes $q_1, \ldots, q_r$. We shall search over all tuples $(q_{i,j})_{1 \leqslant i \leqslant r, 1 \leqslant j \leqslant k}$ of non-negative integers satisfying all of the following requirements:

- $q_i = q_{i,1} + \cdots + q_{i,k}$ for each $1 \leqslant i \leqslant r$,
- each $q_{i,j}$ with $j < k$ is divisible by $\lfloor \alpha\varepsilon n/(25rk) \rfloor$ (no divisibility requirements for $q_{i,k}$), and
- the sums $\sum_{i=1}^{r} q_{i,j}$ for different values of j differ from n/k by at most $\alpha\varepsilon n/(50k)$.

For each eligible tuple $(q_{i,j})$, consider a partition $\mathcal{P} : V = V_1 \cup \cdots \cup V_k$ where $Q_i \cap V_j = q_{i,j}$ (there are many such partitions: pick an arbitrary one). Apply Corollary 5.2 to certify that $\mathcal{P}$ is either $(1+3\alpha/4)\varepsilon$-regular or not $(1+\alpha/2)\varepsilon$-regular. It turns out that the latter option cannot always be true for all $\mathcal{P}$ searched, as we assume that G admits some ε-regular partition with k parts (we will justify this claim). From this search, we find a $(1+3\alpha/4)\varepsilon$-regular partition $\mathcal{P}$ which is almost equitable in the sense that its parts have sizes differing from n/k by at most $\alpha\varepsilon n/(50k)$. We modify $\mathcal{P}$ by moving a minimum number of vertices to make it equitable. We claim that the resulting partition is $(1+\alpha)\varepsilon$-regular.

We next analyse the running time of this algorithm. Corollary 3.5 takes $O(2^{2^{(k/(\alpha\varepsilon))^{O(1)}}} n^2)$ time to find the cut norm decomposition. The number of tuples $(q_{i,j})$ is at most

$$(25rk\alpha^{-1}\varepsilon^{-1})^{kr} \leqslant 2^{2^{(k/(\alpha\varepsilon))^{O(1)}}}.$$

For each $(q_{i,j})$, the algorithm in Corollary 5.2 takes $O(k^2 2^{2^{(\alpha\varepsilon)^{-O(1)}}} n^2)$ time. Therefore, the entire algorithm takes $O(2^{2^{(k/(\alpha\varepsilon))^{O(1)}}} n^2) = O_{\alpha,\varepsilon,k}(n^2)$ time.

Now we verify correctness. We shall prove the following claims, which together imply the result. Indeed, (1) shows that the algorithm always finds some $(1+3\alpha/4)\varepsilon$-regular partition $\mathcal{P}$, and (2) shows that making $\mathcal{P}$ equitable by moving a minimum number of vertices between parts results in a $(1+\alpha)\varepsilon$-regular partition.

(1) If a partition $\mathcal{P} = \{V_1, V_2, \ldots, V_k\}$ of V is ε-regular, then we can modify it slightly to obtain $\mathcal{P}' = \{V_1', V_2', \ldots, V_k'\}$ such that $q_{i,j} = |Q_i \cap V_j'|$ form an eligible tuple, and $\mathcal{P}'$ is $(1+\alpha/2)\varepsilon$-regular for G (so the search would not pass over this $(q_{i,j})$).

(2) If a partition $\mathcal{P}$ of V is $(1+3\alpha/4)\varepsilon$-regular for G, then by modifying $\mathcal{P}$ by adding or deleting at most $\alpha\varepsilon n/(50k)$ vertices from each part, the resulting partition is $(1+\alpha)\varepsilon$-regular.

In order to show these claims, we first establish a few simple lemmas.

Lemma 6.1. *Let X, X', Y be vertex subsets of a graph with $X \subset X'$ and $|X| \geqslant (1-\delta)|X'|$. Then $|d(X',Y) - d(X,Y)| \leqslant \delta$.*

Proof. We have the identity

$$d(X',Y) - d(X,Y) = \frac{e(X' \setminus X, Y) + e(X,Y)}{|X' \setminus X||Y| + |X||Y|} - \frac{e(X,Y)}{|X||Y|}$$
$$= \big(d(X' \setminus X, Y) - d(X,Y)\big) \frac{|X' \setminus X|}{|X'|}.$$

The lemma follows from noting that densities are between 0 and 1 and $|X' \setminus X| \leqslant \delta|X'|$. □

Recall that $A \Delta B := (A \setminus B) \cup (B \setminus A)$ denotes the symmetric difference between A and B.

Lemma 6.2. *If U, U', W, W' are vertex subsets of a graph with $|U \Delta U'| \leqslant \delta|U \cup U'|$ and $|W \Delta W'| \leqslant \delta|W \cup W'|$, then $|d(U,W) - d(U',W')| \leqslant 2\delta$.*

Proof. It suffices to prove the lemma in the case $W = W'$ and with the bound 2δ replaced by δ. Indeed, the lemma would then follow by applying this case twice and the triangle inequality. By the triangle inequality and applying Lemma 6.1 twice with $X' = U \cup U'$, first with

$$\delta_1 = \frac{|U \cup U'| - |U|}{|U \cup U'|}$$

and then with

$$\delta_2 = \frac{|U \cup U'| - |U'|}{|U \cup U'|},$$

and finally using

$$\delta_1 + \delta_2 = \frac{|U \Delta U'|}{|U \cup U'|} \leqslant \delta,$$

we have

$$|d(U,W) - d(U',W)| \leqslant |d(U,W) - d(U \cup U', W)| + |d(U \cup U', W) - d(U',W)|$$
$$\leqslant \delta_1 + \delta_2 \leqslant \delta. \qquad \square$$

Lemma 6.3. *Suppose (V_1, V_2) is an ε-regular pair of vertex subsets of a graph. Suppose we modify them slightly to V_1' and V_2', with $|V_i \Delta V_i'| \leqslant \delta\varepsilon|V_i|$ for $i = 1, 2$. Then V_1' and V_2' are $\varepsilon + 4\delta$-regular.*

Proof. Clearly we may assume that $\varepsilon + 4\delta \leqslant 1$. Let $U' \subseteq V_1'$ and $W' \subseteq V_2'$ with $|U'| \geqslant (\varepsilon + 4\delta)|V_1'|$ and $|W'| \geqslant (\varepsilon + 4\delta)|V_2'|$. Let $U = U' \cap V_1$ and $W = W' \cap V_2$. Then we have

$$|U| = |U'| - |U' \setminus V_1| \geqslant |U'| - |V_1' \setminus V_1| \geqslant (\varepsilon + 4\delta)|V_1'| - \delta\varepsilon|V_1|$$
$$\geqslant (\varepsilon + 4\delta)(|V_1| - \delta\varepsilon|V_1|) - \delta\varepsilon|V_1|$$
$$= \varepsilon|V_1| + 4\delta|V_1| - (1 + \varepsilon + 4\delta)\delta\varepsilon|V_1| \geqslant \varepsilon|V_1|.$$

Similarly $|W| \geqslant \varepsilon|V_2|$. Thus, by the regularity of the pair (V_1, V_2), we have

$$|d(U,W) - d(V_1, V_2)| \leqslant \varepsilon.$$

Now, we have that

$$|U \Delta U'| \leqslant \delta\varepsilon|V_1| \leqslant \delta|U| \leqslant \delta|U \cup U'|,$$

and similarly $|W \Delta W'| \leqslant \delta|W \cup W'|$, and thus $|d(U,W) - d(U',W')| \leqslant 2\delta$ by Lemma 6.2. Similarly

$$|d(V_1', V_2') - d(V_1, V_2)| \leqslant 2\delta\varepsilon \leqslant 2\delta.$$

By the triangle inequality, we have $|d(U',W') - d(V_i', V_j')| \leqslant \varepsilon + 4\delta$, showing that (V_i', V_j') is $(\varepsilon + 4\delta)$-regular. $\qquad\square$

As a corollary, we have the following.

Corollary 6.4. *Let $0 < \varepsilon, \delta < 1$. Let G be a graph with n vertices. Let $\mathcal{P}$ be a partition of $V(G)$ into k parts, with each part having size at least $n/(2k)$. Suppose that $\mathcal{P}$ is ε-regular for G. If we modify $\mathcal{P}$ by adding or deleting at most $\delta\varepsilon|V|/k$ vertices from each part of $\mathcal{P}$, then the resulting partition is $(\varepsilon + 8\delta)$-regular for G.*

Proof. Indeed, for any part V_i of $\mathcal{P}$, if we let V_i' be its modification, then $|V_i \Delta V_i'| \leqslant \delta\varepsilon|V|/k \leqslant 2\delta\varepsilon|V_i|$. This means that if a pair (V_i, V_j) was ε-regular, then after the modification it is $(\varepsilon + 8\delta)$-regular, and so the proportion of pairs that are not $\varepsilon + 8\delta$-regular is at most $\varepsilon \leqslant \varepsilon + 8\delta$. $\qquad\square$

Now we prove claim (1). Let $\mathcal{P}$ be an equitable ε-regular partition of G. Since $\mathsf{d}_{\square}(G, G') \leqslant \alpha\varepsilon/(10k^2)$, $\mathcal{P}$ is $(1 + \alpha/10)\varepsilon$-regular for G'. In G', edges between the same parts of $\mathcal{Q}$ have equal weights, and we can take $\mathcal{P}'$ such that $|Q_i \cap V_j|$ differs from $|Q_i \cap V_j'|$ by at most $\alpha\varepsilon n/(50rk)$ for each i, j. This means that $\mathcal{P}'$ can be taken so that V_j and V_j' differ by at most $\alpha\varepsilon n/(50k)$ for each j, so it follows from the lemma above that $\mathcal{P}'$ is $(1 + 3\alpha/10)\varepsilon$-regular for G'. Therefore, $\mathcal{P}'$ must be $(1 + \alpha/2)\varepsilon$-regular for G.

Claim (2) follows immediately from the corollary above.

7. Permutation regularity lemma

In this section we give a new proof of a regularity lemma for permutations that requires fewer parts than previous results in literature. To define regular partitions for permutations, it is natural to state it as a special case in a more general setting for matrices.

Let $Y = (y_{ij})$ be an $n \times n$ matrix. We use *interval* to mean a subset of $[n]$ of consecutive integers. For any intervals I, J of $[n]$, we write

$$d_Y(I,J) := \frac{1}{|I||J|} \sum_{i \in I, j \in J} y_{ij}.$$

Let $f\colon [0,1]^2 \to [0,1]$ be a measurable function. For any intervals $I,J \subseteq [0,1]$, we write

$$d_f(I,J) := \frac{1}{\lambda(I)\lambda(J)} \int_{I\times J} f(x,y)\,dx\,dy.$$

Here λ denotes the Lebesgue measure.

Definition. Let $f\colon [0,1]^2 \to [0,1]$ be a measurable function. Let $I,J \subseteq [0,1]$ be intervals. We say that (I,J) is *interval ε-regular* for f if, for all subintervals $A \subseteq I$ and $B \subseteq J$ with $\lambda(A) \geqslant \varepsilon\lambda(I)$ and $\lambda(B) \geqslant \varepsilon\lambda(J)$, we have

$$|d_f(A,B) - d_f(I,J)| \leqslant \varepsilon.$$

Let $\mathcal{P}$ be a partition of $[0,1]$ into k intervals. We say that $\mathcal{P}$ is *interval ε-regular* for f if all except at most εk^2 pairs of intervals (I,J) of $\mathcal{P}$ are interval ε-regular for f.

Theorem 7.3. *For every $\varepsilon > 0$ and positive integer m there is some $M = m\varepsilon^{-O(\varepsilon^{-5})}$ with the following property. For every measurable function $f\colon [0,1]^2 \to [0,1]$, there is some integer $k \in [m,M]$ such that the partition of $[0,1]$ into k equal-length intervals*

$$[0,1/k) \cup [1/k,2/k) \cup \cdots \cup [(k-1)/k,1]$$

is interval ε-regular for f.

Remark. In Theorem 7.3, it is possible to take $k = mq^i$, where $q = \lceil 16\varepsilon^{-3}\rceil$ and $0 \leqslant i < \lceil 4\varepsilon^{-5}\rceil$ is some integer.

Before proving Theorem 7.3, we first prove a lemma showing that the density $d_f(A,B)$ does not change very much if A and B are changed only slightly.

Lemma 7.4. *Let $f\colon [0,1]^2 \to [0,1]$ be a measurable function. For any intervals*

$$A,A',B,B' \subseteq [0,1],$$

we have

$$|d_f(A,B) - d_f(A',B')| \leqslant \frac{2\lambda((A\times B)\Delta(A'\times B'))}{\lambda(A)\lambda(B)}.$$

Proof. By the triangle inequality,

$$\begin{aligned}
\lambda(A)&\lambda(B)|d_f(A,B) - d_f(A',B')| \\
&\leqslant |\lambda(A)\lambda(B)d_f(A,B) - \lambda(A')\lambda(B')d_f(A',B')| + d_f(A',B')|\lambda(A)\lambda(B) - \lambda(A')\lambda(B')| \\
&\leqslant \left| \int_{A\times B} f\,d\lambda - \int_{A'\times B'} f\,d\lambda \right| + |\lambda(A)\lambda(B) - \lambda(A')\lambda(B')| \\
&\leqslant 2\lambda((A\times B)\Delta(A'\times B')). \qquad \square
\end{aligned}$$

The bound in Lemma 7.4 can be improved by a factor of 2 via the proof of Lemma 6.2.

Proof of Theorem 7.3.   Let f_k denote the function obtained from f by replacing its value inside each box,

$$B_{ij} = \left[\frac{i}{k}, \frac{i+1}{k}\right) \times \left[\frac{j}{k}, \frac{j+1}{k}\right),$$

with its average inside that box, that is,

$$f_k(x,y) := k^2 \int_{B_{ij}} f\, d\lambda \quad \text{if } (x,y) \in B_{ij}, \quad \text{for } i,j = 0,1,\dots,k-1$$

(when i or j equals $k-1$, the corresponding interval should be modified to be closed on the right). Write

$$\|f\|_2 := \left(\int_{[0,1]^2} |f|^2\, d\lambda\right)^{1/2}$$

for the L^2-norm.

Let $q = \lceil 16\varepsilon^{-3}\rceil$. Consider the sequence $f_m, f_{mq}, f_{mq^2}, \dots$. Since $0 \leqslant \|f_k\|_2 \leqslant 1$ for all k, there exists some $k = mq^i$ for $0 \leqslant i < \lceil 4\varepsilon^{-5}\rceil$ such that

$$\|f_{kq}\|_2^2 \leqslant \|f_k\|_2^2 + \frac{\varepsilon^5}{4}. \tag{7.1}$$

We will show that the partition of $[0,1]$ into k equal-length intervals is interval ε-regular. Indeed, if this were not the case, then there would exist more than εk^2 irregular pairs of intervals (I,J), where $I = [i/k, (i+1)/k)$ and $J = [j/k, (j+1)/k)$ for some integers i and j. Due to the irregularity, there exist subintervals $A \subseteq I$ and $B \subseteq J$ such that $\lambda(A) \geqslant \varepsilon\lambda(I)$, $\lambda(B) \geqslant \varepsilon\lambda(J)$, and

$$|d_f(I,J) - d_f(A,B)| > \varepsilon. \tag{7.2}$$

Let A' be the smallest interval containing A with both ends being multiples of $1/(kq)$. Note that $A' \subseteq I$. Define B' similarly. We see that $A' \times B'$ contains $A \times B$, and the difference in area is at most $4/(k^2 q)$. By Lemma 7.4,

$$|d_f(A,B) - d_f(A',B')| \leqslant \frac{2(4/(k^2 q))}{(\varepsilon/k)^2} = \frac{8}{q\varepsilon^2} = \frac{8}{\lceil 16\varepsilon^{-3}\rceil \varepsilon^2} \leqslant \frac{\varepsilon}{2}.$$

By (7.2) we have

$$|d_f(I,J) - d_f(A',B')| > \frac{\varepsilon}{2}.$$

Since the endpoints of I and J are multiples of $1/k$ and those of A' and B' are multiples of $1/(kq)$, the function $f_k - f_{kq}$ has average value $d_f(I,J) - d_f(A',B')$ over the box $A' \times B'$. So the contribution to $\|f_k - f_{kq}\|_2^2$ from $A' \times B'$ is at least $\lambda(A')\lambda(B')(\varepsilon/2)^2 \geqslant \varepsilon^4/(4k^2)$. As there are more than εk^2 irregular pairs (I,J), and all the rectangles $I \times J$ are disjoint, we have

$$\|f_k - f_{kq}\|_2^2 > \frac{\varepsilon^5}{4}.$$

Note that

$$\int_{[0,1]^2} (f_k - f_{kq}) f_k\, d\lambda = 0$$

since f_k is constant over each box $[i/k,(i+1)/k) \times [j/k,(j+1)/k)$, and f_{kq} averages to f_k on this box. Thus f_k and $f_k - f_{kq}$ are orthogonal, so by the Pythagorean theorem,

$$\|f_{kq}\|_2^2 = \|f_k - (f_k - f_{kq})\|_2^2 = \|f_k\|_2^2 + \|f_k - f_{kq}\|_2^2 > \|f_k\|_2^2 + \frac{\varepsilon^5}{4},$$

which contradicts (7.1). It follows that the partition of $[0,1]$ into k equal-length intervals is interval ε-regular for f. $\qquad\square$

7.2. Dealing with equitable partitions

Here is a lemma that will be useful for the proof of Theorem 7.1. It says that (I,J) being interval regular is robust under changing I and J by a small amount.

Lemma 7.5. *Let $f \colon [0,1]^2 \to [0,1]$ be a measurable function. Let $I,I',J,J' \subseteq [0,1]$. Let $0 < \varepsilon \leqslant 1$. Let $\varepsilon' > 0$ be a quantity less than each of*

$$\varepsilon - \frac{4\lambda((I \times J)\Delta(I' \times J'))}{\varepsilon^2 \lambda(I)\lambda(J)}, \quad \frac{\lambda(I)\varepsilon - \lambda(I \setminus I')}{\lambda(I')}, \quad \frac{\lambda(J)\varepsilon - \lambda(J \setminus J')}{\lambda(J')}.$$

If (I',J') is interval ε'-regular for f, then (I,J) is interval ε-regular for f.

Proof. Let $A \subseteq I$ and $B \subseteq J$ be subintervals such that $\lambda(A) \geqslant \varepsilon\lambda(I)$ and $\lambda(B) \geqslant \varepsilon\lambda(J)$. Let $A' = A \cap I'$ and $B' = B \cap J'$. The second and third hypotheses about ε' above imply that $\lambda(A') \geqslant \varepsilon'\lambda(I')$ and $\lambda(B') \geqslant \varepsilon'\lambda(J')$. Since (I',J') is ε'-regular for f, we have

$$|d_f(A',B') - d_f(I',J')| \leqslant \varepsilon'.$$

By Lemma 7.4, we have

$$|d_f(I,J) - d_f(I',J')| \leqslant \frac{2\lambda((I \times J)\Delta(I' \times J'))}{\lambda(I)\lambda(J)}$$

and

$$|d_f(A,B) - d_f(A',B')| \leqslant \frac{2\lambda((A \times B)\Delta(A' \times B'))}{\lambda(A)\lambda(B)} \leqslant \frac{2\lambda((I \times J)\Delta(I' \times J'))}{\varepsilon^2\lambda(I)\lambda(J)}.$$

It follows by the triangle inequality and the first hypotheses on ε' that

$$|d_f(A,B) - d_f(I,J)| \leqslant \varepsilon,$$

which proves that (I,J) is ε-regular for f. $\qquad\square$

Proof of Theorem 7.1. Let $f \colon [0,1]^2 \to [0,1]$ be the function that takes constant value y_{ij} on the rectangle $[(i-1)/n,i/n) \times [(j-1)/n,j/n)$, for each $1 \leqslant i,j \leqslant n$. By Theorem 7.3, there is some $k \in [m, m\varepsilon^{-O(\varepsilon^{-5})}]$ such that the partition of $[0,1]$ into k equal-length intervals is interval $(\varepsilon/2)$-regular.

Any equitable partition $\mathcal{P}$ of $[n]$ into sets of sizes $c_1,\dots,c_k$, gives rise to a partition $\mathcal{Q}$ of $[0,1]$ into intervals of length $c_1/n,\dots,c_k/n$. Since $\mathcal{P}$ is an equitable partition, the ith interval I_i of $\mathcal{Q}$ differs, in terms of symmetric difference, from $[(i-1)/k,i/k)$ by at most k/n in measure. It

follows from Lemma 7.5 that if n is large enough, say, $n \geqslant 100k^3\varepsilon^{-3}$, then (I_i, J_i) is interval ε-regular for f whenever $[(i-1)/k, i/k) \times [(j-1)/k, j/k)$ is interval $(\varepsilon/2)$-regular for f. It follows that $\mathcal{Q}$ is interval ε-regular for f.

When $n < 100k^3\varepsilon^{-3}$, we can take the partition of $[n]$ into singletons, which is trivially interval ε-regular. $\qquad\square$

References

[1] Alon, N., Duke, R. A., Lefmann, H., Rödl, V. and Yuster, R. (1994) The algorithmic aspects of the regularity lemma. *J. Algorithms* **16** 80–109.

[2] Alon, N. and Naor, A. (2006) Approximating the cut-norm via Grothendieck's inequality. *SIAM J. Comput.* **35** 787–803.

[3] Alon, N. and Spencer, J. H. (2008) *The Probabilistic Method*, third edition, Wiley.

[4] Borgs, C., Chayes, J. T., Lovász, L., Sós, V. T. and Vesztergombi, K. (2008) Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing. *Adv. Math.* **219** 1801–1851.

[5] Conlon, D. and Fox, J. (2012) Bounds for graph regularity and removal lemmas. *Geom. Funct. Anal.* **22** 1191–1256.

[6] Cooper, J. N. (2006) A permutation regularity lemma. *Electron. J. Combin.* **13** 22.

[7] Coppersmith, D. and Winograd, S. (1990) Matrix multiplication via arithmetic progressions. *J. Symbol. Comput.* **9** 251–280.

[8] Dellamonica, D., Kalyanasundaram, S., Martin, D., Rödl, V. and Shapira, A. (2012) A deterministic algorithm for the Frieze–Kannan regularity lemma. *SIAM J. Discrete Math.* **26** 15–29.

[9] Dellamonica Jr, D., Kalyanasundaram, S., Martin, D. M., Rödl, V. and Shapira, A. (2015) An optimal algorithm for finding Frieze–Kannan regular partitions. *Combin. Probab. Comput.* **24** 407–437.

[10] Duke, R. A., Lefmann, H. and Rödl, V. (1995) A fast approximation algorithm for computing the frequencies of subgraphs in a given graph. *SIAM J. Comput.* **24** 598–620.

[11] Fischer, E., Matsliah, A. and Shapira, A. (2010) Approximate hypergraph partitioning and applications. *SIAM J. Comput.* **39** 3155–3185.

[12] Fox, J. and Lovász, L. M. A tight lower bound for Szemerédi's regularity lemma. *Combinatorica*, to appear.

[13] Frieze, A. and Kannan, R. (1999) Quick approximation to matrices and applications. *Combinatorica* **19** 175–220.

[14] Frieze, A. and Kannan, R. (1999) A simple algorithm for constructing Szemerédi's regularity partition. *Electron. J. Combin.* **6** 17.

[15] Gowers, W. T. (1997) Lower bounds of tower type for Szemerédi's uniformity lemma. *Geom. Funct. Anal.* **7** 322–337.

[16] Håstad, J. (1999) Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica* **182** 105–142.

[17] Hoppen, C., Kohayakawa, Y. and Sampaio, R. M. (2012) A note on permutation regularity. *Discrete Appl. Math.* **160** 2716–2727.

[18] Kohayakawa, Y., Rödl, V. and Thoma, L. (2003) An optimal algorithm for checking regularity. *SIAM J. Comput.* **32** 1210–1235.

[19] Komlós, J. and Simonovits, M. (1996) Szemerédi's regularity lemma and its applications in graph theory. In *Combinatorics: Paul Erdős is Eighty*, Vol. 2, János Bolyai Mathematical Society, pp. 295–352.

[20] Le Gall, F. (2014) Powers of tensors and fast matrix multiplication. In *ISSAC '14: Proc. 39th International Symposium on Symbolic and Algebraic Computation*, ACM, pp. 296–303.

[21] Lovász, L. (2012) *Large Networks and Graph Limits*, Vol. 60 of *American Mathematical Society Colloquium Publications*, AMS.

[22] Lovász, L. and Szegedy, B. (2007) Szemerédi's lemma for the analyst. *Geom. Funct. Anal.* **17** 252–270.

[23] Moshkovitz, G. and Shapira, A. (2016) A short proof of Gowers' lower bound for the regularity lemma. *Combinatorica* **36** 187–194.

[24] Szemerédi, E. (1975) On sets of integers containing no *k* elements in arithmetic progression. In *Proc. International Congress of Mathematicians 1974*, Vol. 2, Canadian Mathematical Congress, pp. 503–505.

[25] Szemerédi, E. (1978) Regular partitions of graphs. In *Problèmes Combinatoires et Théorie des Graphes*, Vol. 260 of *Colloq. Internat. CNRS*, CNRS, pp. 399–401.

[26] Tao, T. (2010) *An Epsilon of Room, II*, AMS.

[27] Zuckerman, D. (2007) Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory Comput.* **3** 103–128.