

A Wowzer-Type Lower Bound for the Strong Regularity Lemma

Subrahmanyam Kalyanasundaram* Asaf Shapira†

May 17, 2012

Abstract

The regularity lemma of Szemerédi asserts that one can partition every graph into a bounded number of quasi-random bipartite graphs. In some applications however, one would like to have a strong control on how quasi-random these bipartite graphs are. Alon, Fischer, Krivelevich and Szegedy obtained a powerful variant of the regularity lemma, which allows one to have an *arbitrary* control on this measure of quasi-randomness. However, their proof only guaranteed to produce a partition where the number of parts is given by the Wowzer function, which is the iterated version of the Tower function. We show here that a bound of this type is unavoidable by constructing a graph H , with the property that even if one wants a very mild control on the quasi-randomness of a regular partition, then the number of parts in any such partition of H must be given by a Wowzer-type function.

1 Introduction

The regularity lemma of Szemerédi [23] is one of the most widely used tools in extremal combinatorics. The lemma was originally devised as part of Szemerédi’s proof of his (eponymous) theorem [22] on arithmetic progressions in dense sets of integers. Since then it has turned into a fundamental tool in extremal combinatorics, with applications in diverse areas such as theoretical computer science, additive number theory, discrete geometry and of course graph theory. We refer the reader to [14] and its references for more details on the rich history and applications of the regularity lemma.

Let us turn to formally state the regularity lemma. For a graph $G = (V, E)$ and two disjoint vertex sets A and B , we denote by $e_G(A, B)$ the number of edges of G with one vertex in A and one in B . The *density* $d_G(A, B)$ of the pair (A, B) in the graph G is

*Department of Computer Science and Engineering, IIT Hyderabad, India. Email: subruk@iith.ac.in. This work was done while being a student in School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332.

†School of Mathematics, Tel-Aviv University, Tel-Aviv, Israel 69978, and Schools of Mathematics and Computer Science, Georgia Institute of Technology, Atlanta, GA 30332. Email: asafico@tau.ac.il. Supported in part by NSF Grant DMS-0901355, ISF Grant 224/11 and a Marie-Curie CIG Grant 303320.

$$d_G(A, B) = e_G(A, B)/|A||B|, \tag{1}$$

that is, $d_G(A, B)$ is the fraction of pairs $(x, y) \in A \times B$ such that (x, y) is an edge of G . For $\gamma > 0$, we say that the pair (A, B) in a graph G is γ -regular if for any choice of $A' \subseteq A$ of size at least $\gamma|A|$ and $B' \subseteq B$ of size at least $\gamma|B|$, we have $|d_G(A', B') - d_G(A, B)| \leq \gamma$. Note that a large random bipartite graph is γ -regular for all $\gamma > 0$. Thus we can think of γ as measuring the quasi-randomness of the bipartite graph connecting A and B ; the smaller γ is the more quasi-random the graph is. We will sometimes drop the subscript G in the above notations when the graph G we are referring to is clear from context.

Let $\mathcal{Z} = \{Z_1, \dots, Z_k\}$ be a partition of $V(G)$ into k sets. Throughout the paper, we will only consider partitions into sets Z_i of equal size¹. We will refer to each $Z \in \mathcal{Z}$ as a *cluster* of the partition \mathcal{Z} . The *order* of a partition is the number of clusters it has (k above). We will sometimes use $|\mathcal{Z}|$ to denote the order of \mathcal{Z} . We say that a partition $\mathcal{Z} = \{Z_1, \dots, Z_k\}$ *refines* another partition $\mathcal{Z}' = \{Z'_1, \dots, Z'_{k'}\}$ if each cluster of \mathcal{Z} is contained in one of the clusters of \mathcal{Z}' .

A partition $\mathcal{Z} = \{Z_1, \dots, Z_k\}$ of $V(G)$ is said to be γ -regular if all but γk^2 of the pairs (Z_i, Z_j) are γ -regular. Szemerédi’s regularity lemma can be formulated as follows:

Theorem 1. (Szemerédi [23]) *For any $\gamma > 0$ and t there is an integer $K = K(t, \gamma)$ with the following property; given a graph G and a partition \mathcal{A} of $V(G)$ of order t , one can find a γ -regular partition \mathcal{B} of $V(G)$ which refines \mathcal{A} and satisfies $|\mathcal{B}| \leq K$.*

Let $T(x)$ be the function satisfying $T(0) = 1$ and $T(x) = 2^{T(x-1)}$ for $x \geq 1$. So $T(x)$ is a tower of 2’s of height x . Szemerédi’s proof of the regularity lemma [23] showed that the function $K(t, \gamma)$ can be bounded from above² by $T(1/\gamma^5)$. For a long time it was not clear if one could obtain better upper bounds for $K(t, \gamma)$. Besides being a natural problem, further motivation came from the fact that some fundamental results, such as Roth’s Theorem [19, 20], could be proved using the regularity lemma. Hence improved upper bounds for $K(t, \gamma)$ might have resulted in improved bounds for several other fundamental problems. In a major breakthrough, Gowers [9] proved that the Tower-type dependence is indeed necessary. He showed that for any $\gamma > 0$ there is a graph where any γ -regular partition must have size at least $T(1/\gamma^{1/16})$.

Gowers’ lower bound [9] can be stated as saying that if one wants a regular partition of order k , then there are graphs in which the best quasi-randomness measure one can hope to obtain is merely $1/(\log^*(k))^{16}$. Suppose however that for some $f : \mathbb{N} \mapsto (0, 1)$, we would like to find a partition of a graph of order k that will be “close” to being $f(k)$ -regular. Alon, Fischer, Krivelevich and Szegedy [1] formulated the following notion of being close to $f(k)$ -regular.

¹In some papers partitions of this type are called *equipartitions*.

²We note that in essentially any application of Theorem 1, one takes t to be (at least) $1/\gamma$ so some papers simply consider the function $K'(\gamma) = K(1/\gamma, \gamma)$. The reason is that one wants to avoid “degenerate” regular partitions into a very small number of parts, where most of the graph’s edges will belong to the sets V_i where one has no control on the edge distribution.

Definition 1.1. ((ϵ, f) -regular partition) Let f be a function $f : \mathbb{N} \mapsto (0, 1)$. An (ϵ, f) -regular partition of a graph G is a pair of partitions $\mathcal{A} = \{V_i : 1 \leq i \leq k\}$ and $\mathcal{B} = \{U_{i,i'} : 1 \leq i \leq k, 1 \leq i' \leq \ell\}$ of $V(G)$, where \mathcal{B} is a refinement of \mathcal{A} (with $U_{i,i'} \subseteq V_i$), satisfying the following:

1. \mathcal{B} is $f(k)$ -regular.
2. Say that a pair (V_i, V_j) of clusters of \mathcal{A} is good if all but at most $\epsilon \ell^2$ of pairs $1 \leq i', j' \leq \ell$ satisfy $|d(U_{i,i'}, U_{j,j'}) - d(V_i, V_j)| < \epsilon$. Then, at least $(1 - \epsilon) \binom{k}{2}$ of the pairs (V_i, V_j) are good.

One useful way of thinking about the above notion is to “forget” for a moment about the partition \mathcal{B} and just treat partition \mathcal{A} as an $f(k)$ -regular partition. One then tries to extract some useful information from the assumption that \mathcal{A} itself is $f(k)$ -regular. Finally, one uses the second property of Definition 1.1, which says that the two partitions are *similar*, in order to show that the information deduced from the *assumption* that \mathcal{A} is $f(k)$ -regular can actually be deduced from the *fact* that \mathcal{B} is $f(k)$ -regular.

One of the main results of [1] was that given a graph G and *any* function f , one can construct an (ϵ, f) -regular partition of G of bounded size. This version of the regularity lemma is sometimes referred to as the *strong regularity lemma*. As we have mentioned above, in order to avoid degenerate partitions we will assume henceforth that an (ϵ, f) -regular partition has order at least $1/\epsilon$.

Theorem 2. (Alon et al. [1]) For every $\epsilon > 0$ and $f : \mathbb{N} \mapsto (0, 1)$, there is an integer $S = S(\epsilon, f)$ such that any graph $G = (V, E)$ has an (ϵ, f) -regular partition $(\mathcal{A}, \mathcal{B})$ where $1/\epsilon \leq |\mathcal{A}|, |\mathcal{B}| \leq S$.

Let us describe two cases where one needs to have a better control of the measure of quasi-randomness of a regular partition. A first example is when proving certain variants of the graph removal lemma [20]. In such a scenario we are given a regular partition and would like to be able to say that since the partition behaves in a quasi-random way, then we can find “small” subgraphs that we expect to find in a truly random graph. The only problem is that as the “small” structure we are trying to find becomes larger, we need the measure of quasi-randomness to decrease with it. Some examples where Theorem 2 was used to overcome such difficulties can be found in [1, 2, 4, 5, 13, 21]. We note that in some of these papers, Theorem 2 was used with functions f that go to zero extremely fast, so the ability to apply the theorem with arbitrary functions was crucial.

Another example when one wants a better control of the measure of quasi-randomness is when the graph we are trying to partition is very sparse. It is not hard to see that for the notion of γ -regularity to make sense, the graph should have density at least γ . A well known case where one is faced with increasingly sparse graphs is in the proofs of the hypergraph regularity lemma, that were obtained independently by Gowers [11], by Rödl et al. [8, 17, 18] and later also by Tao [24]. In those proofs, one is partitioning not only the vertices of the hypergraph (as in Theorem 1) but also the pairs of vertices into quasi-random bipartite graphs. However, in the process these bipartite graphs become sparser so one needs to control their quasi-randomness as a function of their density. Out of the aforementioned proofs of the hypergraph regularity lemma, Tao’s proof actually uses

Theorem 2 in order to address this issue. See the survey of Gowers [11] for an excellent account of this.

We finally note that the strong regularity lemma is also related to the notion of the limit of convergent graph sequences defined and studied in [7]. Without defining these notions explicitly, we just mention that many of the results mentioned above which were proved using Theorem 2, were later reproved using graph limits, see e.g. Lovász and Szegedy [16]. Furthermore, some of the important properties of the limit of a convergent graph sequence, such as its uniqueness [15], also hold for (ϵ, f) -regular partitions, see [4]. Hence, one can view an (ϵ, f) -regular partition as the discrete analogue of the (analytic) limit of a convergent graph sequence.

Let $W(x)$ be the function satisfying $W(0) = 1$ and $W(x) = T(W(x - 1))$ for $x \geq 1$. So the function W is an iterated version of the Tower function $T(x)$. The function W is sometimes referred to as the Wowzer³ function (for obvious reasons). The proof of Theorem 2 in [1] gave a W -type upper bound for the function $S(\epsilon, f)$ in Theorem 2. As we have mentioned above, in some applications of this lemma one uses functions f that go to zero extremely fast. But in some cases, as was the case in [1], one uses moderate functions like $f(x) = 1/x^2$. However, even when the function f is $f(x) = 1/x$, the upper bound given in [1] for the function $S(\epsilon, f)$ is (roughly) $W(1/\epsilon)$. Hence it is natural to ask if better bounds can be obtained for such versions of Theorem 2. Our main result here is that a W -type dependence is unavoidable even in this case.

Theorem 3. *Set $f(x) = 1/x$. For every small enough $\epsilon \leq c_0$ there is a graph H with the following property: If $(\mathcal{A}, \mathcal{B})$ is an (ϵ, f) -regular partition of H , and⁴ $|\mathcal{A}| \geq 1/\epsilon$, then $|\mathcal{A}| \geq W(\sqrt{\log(1/\epsilon)}/100)$.*

An interesting aspect of our proof is that it gives the same lower bound even if one considers a much weaker condition than the second condition in Definition 1.1. What we show is that the lower bound of Theorem 2 holds even if one wants only $\epsilon^{1/10}k^2$ of the pairs (V_i, V_j) to be good. Observe that Definition 1.1 asks⁵ for $(1 - \epsilon)\binom{k}{2}$ good pairs! In other words, the lower bound holds even if one is interested in having a very weak similarity⁶ between the partitions \mathcal{A} and \mathcal{B} .

Another interesting aspect of the proof of Theorem 3 is that by resetting the parameters appropriately, one can get W -type lower bounds for (ϵ, f) -regularity for any function $f : \mathbb{N} \mapsto (0, 1)$ going to zero faster than $1/\log^*(x)$. Observe that this is not a caveat of the proof; when $f(x) = 1/\log^*(x)$, Theorem 1 can be formulated as saying that any graph has an (ϵ, f) -regular partition of order

³This name was coined by Graham, Rothschild and Spencer [12].

⁴As we have mentioned before, in order to rule out degenerate partitions (such as taking a partition into 1 set) we assume that $|\mathcal{A}| \geq 1/\epsilon$. A similar assumption was used in [1], where they assume that $f(x) \leq \epsilon$. These two assumptions are basically equivalent (recall that $f(x) = 1/x$), but the one we use makes the notation somewhat simpler.

⁵We note that the application of Theorem 2 in [1] (as well as in most other papers) critically relied on the partition having $(1 - \epsilon)\binom{k}{2}$ good pairs.

⁶Recall the discussion following Definition 1.1.

$T(1/\epsilon^5)$. Hence, one cannot obtain a W -type lower bound for f of this type. So we see that even if one wants to have a very weak relation between the order of \mathcal{A} and the regularity measure of \mathcal{B} (say, $1/\log\log(k)$) one would still have to use a partition of size given by a W -type function⁷.

The ideas we use here in order to prove Theorem 3 appear to be useful also for proving W -type lower bounds for the hypergraph regularity lemma [8, 10, 11, 17, 18, 24]. As we explained above, in this case also one is faced with the need to control a measure of quasi-randomness approaching 0, and this seems to be the main reason why the current bounds for this lemma are of W -type. This investigation is part of a joint work of the second author with Dellamonica and Rödl.

Organization: The rest of the paper is organized as follows. In the following section we describe the graph H which we use in proving Theorem 3. In Section 3 we give an overview of the proof, state the two key lemmas that are needed to prove Theorem 3 and then derive Theorem 3 from them. In Section 4 we prove several preliminary lemmas that we would later use in the proofs of the two key Lemmas. In Sections 5 and 6 we prove the key lemmas stated in Section 3.

2 A Hard Graph for the Strong Regularity Lemma

In this section we describe the graph H which will have the properties asserted in Theorem 3. The description will be somewhat terse; the reader can find in Section 3 an overview of the proof of Theorem 3, which includes some intuition/motivation for the way we define H .

2.1 A weighted reformulation of Theorem 3

Suppose P is a weighted complete graph, where each edge (x, y) is assigned a weight $d_P(x, y) \in [0, 1]$. For two sets A, B define the weighted density between A, B

$$d_P(A, B) = \sum_{x \in A, y \in B} d_P(x, y) / |A||B|. \quad (2)$$

Note that if we think of a graph as a weighted complete graph with 0/1 weights then the above definition coincides with the definition of $d_G(A, B)$ given in (1). Also note that when $A = \{x\}$, $B = \{y\}$ are just two vertices then $d_P(A, B)$ is just the weight $d_P(x, y)$ assigned to (x, y) as above. The following simple claim follows immediately from a standard application of Chernoff's inequality.

Claim 2.1. *Let $\zeta > 0$. Suppose P is a weighted complete graph with weights in $[0, 1]$, and H is a random graph, where each edge (x, y) is chosen independently to be included in H with probability $d_P(x, y)$. Then with probability at least $1/2$ we have*

$$|d_H(A, B) - d_P(A, B)| \leq \zeta,$$

for all sets A, B of size at least $20\zeta^{-2} \log(n)$.

⁷But in such cases the bound might become $W(\log\log(1/\epsilon))$ or some other W -type function.

It is clear that we can prove Theorem 3 by constructing an arbitrarily large graph, such that the number of vertices n will be much larger than all the constants involved. Hence, by the above claim, we see that in order to prove Theorem 3 it is enough to construct a *weighted* graph H satisfying the condition of the theorem with respect to the notion of $d(A, B)$ defined in (2). The reason is that by Claim 2.1, if we have a weighted graph H satisfying Theorem 3, then a random graph generated as in Claim 2.1 will satisfy the assertion of Theorem 3 with high probability. Therefore, from this point and *throughout the paper* we will focus on the construction of a *weighted* graph H satisfying the condition of Theorem 3. Hence, from now on, whenever we talk about $d(A, B)$ we will be referring to the weighted density between A, B as in (2).

2.2 A preliminary construction

In this subsection we describe the first step in defining the graph H of Theorem 3. This graph will be a variant of the graph used by Gowers in [9]. We start with the following definition.

Definition 2.2. (Balanced Partitions) *Let M be an integer and suppose we have a sequence $(A_i, B_i)_{i=1}^m$ of (not necessarily distinct) partitions of $[M]$. We call this sequence of partitions balanced if for any distinct $j, j' \in [M]$, the number of $1 \leq i \leq m$ for which j and j' lie in the same set of the partition (A_i, B_i) is at most $3m/4$.*

The following claim appears in [9]. For completeness, we will reproduce a simple proof later on in the paper (see Section 4).

Claim 2.3. *Let $\phi(m) = 2^{\lceil m/16 \rceil}$. Then for every $m \geq 1$ there exists a sequence of m balanced partitions of $\phi(m)$.*

Let $T^\phi(x)$ be the function satisfying $T^\phi(0) = 1$ and $T^\phi(x) = T^\phi(x-1)\phi(T^\phi(x-1))$ for $x \geq 1$, where $\phi(x) = 2^{\lceil x/16 \rceil}$ is the function defined in Claim 2.3. It is not hard to see that T^ϕ is a Tower-type function, and that in fact $T^\phi(x) \geq T(\lfloor x/2 \rfloor)$ (see Section 4).

Let us define a sequence of integers as follows. We set

$$w(1) = \lfloor \log \log(1/\epsilon) \rfloor, \tag{3}$$

and define inductively

$$w(x+1) = \lfloor \log \log(T^\phi(w(x))) \rfloor. \tag{4}$$

It is also not hard to see that $w(x)$ has a W -type dependence on x . Specifically we will later (see Section 4) observe that:

Claim 2.4. *For every integer $x \geq 1$, we have $w(x) \geq W(\lfloor x/2 \rfloor)$.*

We now turn to define a graph G , which we will later modify in order to get the actual graph H that will satisfy the assertion of Theorem 3. In order to define G we will first define a sequence of

partitions of the vertex set of G . For simplicity we will identify the n vertices of G with the integers $[n]$. So let $n \in \mathbb{N}$ and set $s = w(\frac{1}{48}\sqrt{\log(1/\epsilon)})$, where $w(x)$ is the function defined in (4). We set $m_0 = 1$ and for $1 \leq r \leq s$, let $m_r = m_{r-1}\phi(m_{r-1})$. For each $0 \leq r \leq s$, let $X_1^{(r)}, X_2^{(r)}, \dots, X_{m_r}^{(r)}$ be a partition of $[n]$ into m_r intervals of integers of equal size⁸. We will later refer to this partition as *canonical partition* \mathcal{P}_r . Thus at level r , we have a canonical partition \mathcal{P}_r consisting of m_r clusters. So \mathcal{P}_0 is just the entire vertex set of G . Note that using the notation we introduced above we have

$$|\mathcal{P}_r| = m_r = T^\phi(r). \quad (5)$$

A crucial observation that will be used repeatedly in the paper is that for every $r < r'$, partition $\mathcal{P}_{r'}$ refines partition \mathcal{P}_r .

We finally arrive at the actual definition of G . We will start with the graph G where each pair of vertices (x, y) has weight 0. We will then go over the partitions $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_s$ one after the other, and in each case increase the weight between some of the pairs (x, y) .

Consider some $r \geq 1$ and focus on \mathcal{P}_r and \mathcal{P}_{r-1} . Let us simplify the notation a bit and set $m = m_{r-1}$, $M = \phi(m)$ and $m_r = Mm$. So m is the number of clusters of \mathcal{P}_{r-1} , M is the number of clusters of \mathcal{P}_r inside each cluster of \mathcal{P}_{r-1} , and mM is the number of clusters of \mathcal{P}_r . Let us use X_1, \dots, X_m to denote the m clusters of \mathcal{P}_{r-1} . Also, for each $1 \leq i \leq m$ we use $X_{i,1}, \dots, X_{i,M}$ to denote the M clusters of \mathcal{P}_r inside X_i . Now, for each $1 \leq i \leq m$, let $(A'_{i,j}, B'_{i,j})_{j=1}^M$ be a sequence of balanced partitions of $[M]$. Such a collection exists since $M = \phi(m)$ so Claim 2.3 can be used here. One can think of each of these partitions as partitioning the clusters of \mathcal{P}_r within cluster X_i . Let $A_{i,j} = \cup_{t \in A'_{i,j}} X_{i,t}$ and $B_{i,j} = \cup_{t \in B'_{i,j}} X_{i,t} = X_i \setminus A_{i,j}$. We now update the weights of G as follows: If $(x, y) \in X_i \times X_j$, then we increase $d_G(x, y)$ by $4^{-r}/4\sqrt{\log(1/\epsilon)}$ if and only if $(x, y) \in A_{i,j} \times A_{j,i}$ or $(x, y) \in B_{i,j} \times B_{j,i}$. We will later refer several times to the following observation.

Fact 2.5. *For any $x, y \in V(G)$ we have $d_G(x, y) \leq 4^{-\sqrt{\log(1/\epsilon)}}$.*

2.3 Adding Traps to G

We will now need to modify the graph G defined above in order to obtain the graph H from Theorem 3. To this end we will need to define certain quasi-random graphs. Let $b' < b$ and consider two of the canonical partitions $\mathcal{P}_{b'}$ and \mathcal{P}_b defined in the previous subsection. Suppose \mathcal{P}_b has order m_b and let V be a set of m_b vertices, where we identify vertex $i \in V$ with cluster $X_i \in \mathcal{P}_b$. Note that with this interpretation in mind, one can think of a cluster $U \in \mathcal{P}_{b'}$ as a subset of vertices $U' \subseteq V$, where vertex j belongs to U' if and only if cluster $X_j \in \mathcal{P}_b$ is a subset of U . It follows that for every $b' < b$, partition $\mathcal{P}_{b'}$ defines a natural partition of V into $m_{b'}$ subsets $U_1^{b'}, \dots, U_{m_{b'}}^{b'}$ corresponding to its $m_{b'}$ clusters.

⁸We assume that n is such that it can be divided into equal sized parts of size m_r for all $0 \leq r \leq s$.

We now arrive at a critical definition. We will use $e(R, R')$ to denote the number of edges in a graph with one vertex in R and another in R' , where edges in $R \cap R'$ are counted twice⁹.

Definition 2.6. (Trap) Let \mathcal{P}_b , m_b , V and the partitions $U_1^{b'}, \dots, U_{m_b}^{b'}$ be as above. Let $\mathcal{O} = (V, E)$ be an m_b -vertex graph on V . Then \mathcal{O} is said to be a trap if it satisfies the following two conditions:

- For every pair of sets $R, R' \subseteq V(\mathcal{O})$ of size $\lceil \sqrt{m_b}/4 \rceil$ we have

$$\left| e(R, R') - \frac{1}{2}|R||R'| \right| \leq \frac{1}{4}|R||R'|.$$

- For every $b' < b$, for every $1 \leq i, j \leq m_{b'}$, every choice of $200 \leq k \leq \log(m_b)$, every choice of $R \subseteq U_i^{b'}$ of size k^6 and every choice of $R' \subseteq U_j^{b'}$ of size $\lceil |U_j^{b'}|/k \rceil$, we have

$$\left| e(R, R') - \frac{1}{2}|R||R'| \right| \leq \frac{1}{k^2}|R||R'|.$$

We will later prove the following (see Section 4).

Claim 2.7. *There is a constant C , such that for every $m > C$, there exists a trap on m vertices.*

We are now ready to describe the modifications needed to turn G into the graph H . We do the following for every integer $1 \leq g \leq \frac{1}{48}\sqrt{\log(1/\epsilon)}$; let $b = w(g)$ be the integer defined in (4), let m_b be the order of \mathcal{P}_b and let $\mathcal{O}_b = (V, E)$ be¹⁰ a trap on a vertex set V of size m_b . Recall that we identify vertex $i \in V$ with cluster $X_i \in \mathcal{P}_b$. We now modify G as follows; for every pair of clusters (X_i, X_j) , if $(i, j) \in E(\mathcal{O}_b)$ we increase by 4^{-g} the weight of every pair of vertices $(x, y) \in X_i \times X_j$. If $(i, j) \notin E(\mathcal{O}_b)$ we do not increase the weight of (x, y) . Let us state the following fact to which we will later refer.

Fact 2.8. *The smallest weight used when placing a trap in H is $4^{-\frac{1}{48}\sqrt{\log(1/\epsilon)}}$.*

Later on in the paper we will say that we have *placed* a trap on partition \mathcal{P}_b if b is one of the integers $w(1), \dots, w(\frac{1}{48}\sqrt{\log(1/\epsilon)})$. If a trap was placed on \mathcal{P}_b and (i, j) is an edge of the graph \mathcal{O}_b that was used in the previous paragraph, then we will say that the pair (X_i, X_j) *belongs* to the trap placed on \mathcal{P}_b . Also, if $b = w(g)$, then we will refer to the trap placed on \mathcal{P}_b as the g^{th} trap placed in H . Finally, if $(x, y) \in X_i \times X_j$ and (X_i, X_j) belongs to the trap placed on $\mathcal{P}_{w(g)}$ then we will say that (x, y) *received* an extra weight of 4^{-g} from the g^{th} trap placed in H .

⁹Note that this definition is compatible with the definition of $e(A, B)$ we used earlier, where we assumed that the sets A, B are disjoint.

¹⁰Note that since we only ask Theorem 3 to hold for small enough ϵ , we can assume that ϵ is small enough so that already $m_{w(1)} = T^\phi(w(1))$ would be larger than C , thus allowing us to pick a trap via Claim 2.7 (where $w(1)$ is defined in (3)).

Using the above jargon, we can thus say that in order to obtain the graph H from the graph G we do the following for every $1 \leq g \leq \frac{1}{48}\sqrt{\log(1/\epsilon)}$; setting $b = w(g)$, we place the g^{th} trap on partition \mathcal{P}_b , by increasing the weight of (x, y) by 4^{-g} if and only if $(x, y) \in X_i \times X_j$ and (X_i, X_j) belongs to the trap.

Let us draw some distinction between the way we assigned weights to edges in G and the way we have done so when modifying G to obtain H . When defining G we looked at *each* of the partitions \mathcal{P}_r , and for *every* $X_i, X_j \in \mathcal{P}_{r-1}$ added weight $4^{-r}/4\sqrt{\log(1/\epsilon)}$ only to *some* of the pairs $(x, y) \in X_i \times X_j$. More specifically, we considered the partitions of $X_i = A_{i,j} \cup B_{i,j}$ and $X_j = A_{j,i} \cup B_{j,i}$ and only added the weight $4^{-r}/4\sqrt{\log(1/\epsilon)}$ when either $(x, y) \in A_{i,j} \times A_{j,i}$ or $(x, y) \in B_{i,j} \times B_{j,i}$. When adding the traps, we have only added weights to *some* of the partitions \mathcal{P}_b , that is, those for which $b = w(g)$ for some $1 \leq g \leq \frac{1}{48}\sqrt{\log(1/\epsilon)}$. Moreover, when placing a trap on \mathcal{P}_b we added weight 4^{-g} only to pairs (x, y) connecting *some* of the pairs (X_i, X_j) (those that belong to the trap). Finally, for each such pair (X_i, X_j) we either added more weights to all the pairs $(x, y) \in X_i \times X_j$ or to none of them.

Another important distinction is the following; suppose $b = w(g)$. Then in G , the weight that was added to \mathcal{P}_b was $4^{-b}/4\sqrt{\log(1/\epsilon)}$ while the weight we added when placing a trap on \mathcal{P}_b is 4^{-g} . Since w is a W -type function we see that the weights assigned in G to a specific partition \mathcal{P}_b are extremely small compared to those assigned to \mathcal{P}_b when placing a trap on it (assuming a trap was placed on \mathcal{P}_b).

We also observe that for every pair of vertices (x, y) of H , the total weight it can receive from all the traps we placed is bounded by $1/4 + 1/16 + \dots < 1/3$. We also recall Fact 2.5 stating that the total weight assigned to a pair (x, y) in G is bounded by $1/4\sqrt{\log(1/\epsilon)}$. This means that $d_H(x, y) \leq 1$, as needed for the application of Claim 2.1.

3 Proof Overview, Key Lemmas and Proof of Theorem 3

Our goal in this section is fourfold; give an overview of the proof of Theorem 3, describe the main intuition behind the construction of H , state the two key lemmas that will be used to prove Theorem 3 and finally derive Theorem 3 from these two lemmas.

Perhaps the best way to approach our construction of H is to first consider the proof of Theorem 2 in [1]. For simplicity, let us consider the case $f(x) = 1/x$; we start by taking \mathcal{A}_1 to be an arbitrary partition of G of order $1/\epsilon$, and then apply Theorem 1 in order to find a $1/|\mathcal{A}_1|$ -regular partition, \mathcal{B}_1 , of G which refines \mathcal{A}_1 . Note that by definition, \mathcal{A}_1 and \mathcal{B}_1 satisfy the first condition of Definition 1.1, so if they also satisfy the second, then we are done. If they do not, then we set \mathcal{A}_2 to be \mathcal{B}_1 and use Theorem 1 to find a $1/|\mathcal{A}_2|$ -regular partition, \mathcal{B}_2 , of G which refines \mathcal{A}_2 . Note that \mathcal{A}_2 and \mathcal{B}_2 satisfy the first property, so if they satisfy the second we are done. The process thus goes on till we end up with a pair of partitions $\mathcal{A}_i, \mathcal{B}_i$ which satisfy the second condition. The main argument in [1] shows that this process must stop after (about) $1/\epsilon$ steps with a pair $\mathcal{A}_i, \mathcal{B}_i$ which satisfies

the second condition, and also (by definition) the first condition. Since the above proof applies Theorem 1 repeatedly, where each time we take $1/\gamma$ to be the order of the previous partition, the bound we obtain is of W -type.

Of course, if we want to have any chance of proving Theorem 3, we need to come up with a graph for which the *proof* of Theorem 2 will produce a partition of W -size. Given the overview of this proof described above, the graph H needs to have two properties: (1) For every $\gamma > 0$, any γ -regular partition of H has size given by a Tower-type function; (2) one needs to iteratively apply Theorem 1 a super-constant¹¹ number of times in order to get two partitions \mathcal{A} and \mathcal{B} satisfying the second condition of Definition 1.1. The first property will guarantee that each time we apply Theorem 1 we get a Tower-type increase in the size of \mathcal{A}_i while the second condition will guarantee that we will have to repeat this sufficiently many times.

Let us describe how to get a graph satisfying property (1) mentioned above. Recall that Gowers showed [9] that for every γ there exists a graph with the property that any γ -regular partition has a size $T(1/\gamma^{1/16})$. It is not hard to see that by a minor “tweak” of his construction¹² one can get a *single* graph that works for all γ bounded away from 0. This is basically¹³ the graph G we defined in Subsection 2.2. For completeness let us describe the intuition behind Gowers’ construction. Let us explain why the partitions \mathcal{P}_r used in the construction of G cannot be used as γ -regular partitions of G . Recall that at each iteration, we take every pair of sets $X_i, X_j \in \mathcal{P}_{r-1}$ split them as $X_i = A_{i,j} \cup B_{i,j}$ and $X_j = A_{j,i} \cup B_{j,i}$ and increase the weight between $A_{i,j}, A_{j,i}$ and $B_{i,j}, B_{j,i}$. So, in some sense, each partition \mathcal{P}_r is used in order to rule out the possibility of using the previous partition \mathcal{P}_{r-1} as a γ -regular partition. We note that when one comes about to actually prove that no other (small) partition can be γ -regular one relies critically on the fact that the weights assigned to the partitions \mathcal{P}_r in G decrease exponentially (as a function of r). This makes sure that any irregularity found in level r cannot be canceled by weights assigned to levels $r' > r$.

Let us describe how to get a graph satisfying property (2) mentioned above. Recall that G was defined over a sequence of partitions \mathcal{P}_r . Suppose we want to make sure that two specific partitions in this sequence \mathcal{P}_r and $\mathcal{P}_{r'}$, with $\mathcal{P}_{r'}$ refining \mathcal{P}_r , will not satisfy the second property of Definition 1.1. Then we can do the following; we take a random graph \mathcal{O} whose vertices are the clusters of $\mathcal{P}_{r'}$, and for every edge $(i', j') \in E(\mathcal{O})$ increase the weight of all pairs $(x, y) \in U_{i'} \times U_{j'}$, where $U_{i'}, U_{j'} \in \mathcal{P}_{r'}$. This is just the trap we used in Subsection 2.3. Since we use a random graph, we expect *all* pairs of clusters (X_i, X_j) of \mathcal{P}_r to not be good (in the sense of Definition 1.1) since close to half of the clusters $(U_{i'}, U_{j'})$ with $U_{i'} \subseteq X_i, U_{j'} \subseteq X_j$, will get an extra weight while the other half will not. Now it is not hard to see that for this to work we do not actually have to put the

¹¹To be precise, in order to get a W -type lower bound the number of iterations needs to be larger than $W^{-1}(1/\epsilon)$.

¹²In fact, we will be tweaking the construction of Gowers [9] which gives a slightly weaker lower bound of $T(\log(1/\gamma))$, and is much simpler to analyze. Since we are trying to prove W -type lower bounds it makes little difference if we are iterating the function $T(x)$ or $T(\log(x))$.

¹³If we were only interested in getting a graph that for all $\gamma > 0$ had only γ -regular partitions of Tower-size, then we could have used the weights 4^{-r} instead of $4^{-r}/4^{\sqrt{\log(1/\epsilon)}}$ like we do.

trap on $\mathcal{P}_{r'}$; it is enough to do that on some partition \mathcal{P}_b with $r \leq b \leq r'$. Since we will make sure that a γ -regular partition must be huge, in order to satisfy the first condition of Definition 1.1 one would have to pick two partitions $\mathcal{P}_{r'}, \mathcal{P}_r$ with r' being much larger than r . Therefore, in order to make sure that all pairs $\mathcal{P}_{r'}, \mathcal{P}_r$ will fail the second condition, it is enough to place the traps only on very few partitions \mathcal{P}_b , where by few we mean that there will be a Tower-type jump between their indices.

So with one serious caveat, if one wants to construct an (ϵ, f) -regular partition by taking \mathcal{A} and \mathcal{B} to be two of the canonical partitions $\mathcal{P}_r, \mathcal{P}_{r'}$, then one is forced to take two partitions that refine the last trap we have placed in H . The reason is that by property (1) the integers r and r' must be very far apart, and the way we have placed the traps will guarantee that there will be a trap in between them which will then make sure that they do not satisfy the second property of Definition 1.1. The caveat we are referring to is the fact that once we have added the traps to G , we have destroyed the critical feature of the graph G , which is that the weights decrease exponentially (recall the observation we made above and the discussion at the end of Subsection 2.3). Hence, it is no longer true that once we find a discrepancy in some partition \mathcal{P}_r , this discrepancy cannot be canceled by lower levels. In terms of analyzing Gowers' example, it might be the case that some pairs which were not γ -regular in G , might become γ -regular in H . Actually, there *will* be such pairs. This might completely ruin our ability to prove the H has only γ -regular partitions of Tower-size.

We overcome the above problem by proving that it cannot happen *very often*. Namely, since the trap we have added originates from a random graph, then at least on average we expect it to contribute the same density to all pairs of vertex sets. So *on average*, we do not expect a trap to cancel a discrepancy caused by partitions that are refined by it. This is of course only true on average. To turn this into a deterministic statement, we formulate a condition that holds in random graphs, and show that if too many pairs that were supposed to be not γ -regular somehow turn out to be γ -regular, then we get a violation of the property we assume the trap to satisfy. Turning this intuition into formality is probably the most challenging part of this paper. One of the main reasons is that we cannot run this argument over all the pairs; instead we need to somehow “pack” them together and then argue about each of these packaged pairs. See Lemmas 6.4 and 6.5.

We now turn to the key lemmas of the paper. To state them we will need to define the notion of β -refinement. We briefly mention that this notion is crucial in overcoming another assumption we have used in the above discussion, that one is trying to construct an (ϵ, f) -regular partition by using only the canonical partitions \mathcal{P}_r . Using the notion of β -refinement we will show that one actually has to *approximately* use only such partitions.

Let $0 \leq \beta < 1/2$. Given two sets Z and X , we write $Z \subset_\beta X$, to denote the fact that $|Z \cap X| \geq (1 - \beta)|Z|$. We will sometimes also say that X β -contains Z or that Z is β -contained in X to refer to the fact that $Z \subset_\beta X$. Note that since we assume that $\beta < 1/2$, there can be at most one set X which β -contains a set Z . Given two partitions $\mathcal{P} = \{X_1, \dots, X_m\}$ and $\mathcal{Z} = \{Z_1, \dots, Z_k\}$

of $V(H)$ and $0 \leq \beta < 1/2$, we shall say that \mathcal{Z} is a β -refinement of \mathcal{P} if for at least $(1 - \beta)k$ values of t , there exists i such that $Z_t \subset_\beta X_i$. Observe that if $\beta = 0$, then β -refinement coincides with the standard notion of one partition refining another one, that we discussed earlier.

In what follows, when we refer to the graph H we mean the graph H defined in the previous section. We now state the two key lemmas we will prove later on in the paper. Getting back to the intuitive discussion above, one can think of the first lemma as formalizing condition (1) mentioned above, which we wanted H to satisfy.

Lemma 3.1. *Let $f(x) = 1/x$. Suppose \mathcal{A} and \mathcal{B} form an (ϵ, f) -regular partition of H . If $|\mathcal{A}| = k \geq 1/\epsilon$ then \mathcal{B} is an $\epsilon^{1/5}$ -refinement of $\mathcal{P}_{2 \log \log k}$.*

Note that if $\beta < 1/2$ and partition \mathcal{A} is a β -refinement of \mathcal{P}_r then the order of \mathcal{A} is at least half the order of \mathcal{P}_r . Hence the above lemma (implicitly) says that partition \mathcal{B} , which must be $1/k$ -regular, must have order at least half times the order of $\mathcal{P}_{2 \log \log k}$. Recalling (5), this means that $|\mathcal{B}| \geq (1/2) \cdot T^\phi(2 \log \log k)$. We note however, that knowing that \mathcal{B} must have Tower size is not enough for our proof to work. We actually need to know that \mathcal{B} is a good refinement of partition $\mathcal{P}_{2 \log \log k}$. This is needed in order to show that if a trap was placed between \mathcal{A} and \mathcal{B} then they will indeed fail to satisfy the second property of Definition 1.1. This is exactly where the notion of β -refinement becomes useful, as we state in the second key lemma, which formalizes property (2) mentioned above that we wanted H to satisfy.

Lemma 3.2. *Suppose \mathcal{A}, \mathcal{B} are two partitions of H with the following properties*

- \mathcal{B} is a refinement of \mathcal{A} .
- $|\mathcal{A}| = k$ and H has a trap on a canonical partition \mathcal{P}_b whose order is at least k^2 .
- \mathcal{B} is an $\epsilon^{1/5}$ -refinement of \mathcal{P}_b .

Then \mathcal{A} and \mathcal{B} do not satisfy the second condition of Definition 1.1. In particular they do not form an (ϵ, f) -regular partition of H .

We end this section with the derivation of Theorem 3 from Lemma 3.1 and Lemma 3.2.

Proof of Theorem 3: Suppose \mathcal{A} and \mathcal{B} form an (ϵ, f) -regular partition of H , where $|\mathcal{A}| = k \geq 1/\epsilon$. Let m_s denote the order of \mathcal{P}_s , which is the largest partition on which we have placed a trap. Recall that $s = w(\frac{1}{48} \sqrt{\log(1/\epsilon)})$ and that $m_s \geq s$ (in fact, $m_s = T^\phi(s)$). Hence, by Claim 2.4 we have $m_s \geq W(\frac{1}{96} \sqrt{\log(1/\epsilon)})$. Therefore, if $k \geq \sqrt{m_s}$ we are done since $\sqrt{W(\frac{1}{96} \sqrt{\log(1/\epsilon)})} > W(\frac{1}{100} \sqrt{\log(1/\epsilon)})$ (with a lot of room to spare).

We can thus assume that $|\mathcal{A}| = k \leq \sqrt{m_s}$, and choose b to be the *smallest* index of a partition \mathcal{P}_b , on which we have placed a trap satisfying $|\mathcal{P}_b| \geq k^2$. If we could show that \mathcal{B} forms an $\epsilon^{1/5}$ -refinement of \mathcal{P}_b , then an application of Lemma 3.2 would give that \mathcal{A} and \mathcal{B} do not form an

(ϵ, f) -regular partition of H , which would be a contradiction. Now, Lemma 3.1 tells us that \mathcal{B} is an $\epsilon^{1/5}$ -refinement of $\mathcal{P}_{2\log\log k}$. Note that if \mathcal{B} is an $\epsilon^{1/5}$ -refinement on $\mathcal{P}_{2\log\log k}$ then it is also an $\epsilon^{1/5}$ -refinement of any partition that is refined by $\mathcal{P}_{2\log\log k}$. In other words, it is enough¹⁴ that we show that $b \leq 2\log\log(k)$.

Suppose first that $b = w(1)$, that is, the first trap of size at least k^2 is the first trap placed in H . Then recalling (3) and the fact that $k \geq 1/\epsilon$, we have

$$b = w(1) = \lfloor \log\log(1/\epsilon) \rfloor \leq 2\log\log(k) ,$$

as needed. Suppose now that $b = w(g+1)$ for some $g \geq 1$ and that the trap with largest order smaller than k^2 was placed on $\mathcal{P}_{b'}$ where $b' = w(g)$. Then recalling (4) we see that $b = \lfloor \log\log(T^\phi(b')) \rfloor$. We also recall (5) stating that $|\mathcal{P}_{b'}| = T^\phi(b')$. We thus infer that

$$T^\phi(b') = |\mathcal{P}_{b'}| \leq k^2 ,$$

implying that

$$b = \lfloor \log\log(T^\phi(b')) \rfloor \leq \log\log(k^2) \leq 2\log\log(k) ,$$

thus completing the proof. ■

As one can see from our proof of Theorem 3, what we show is not only that an (ϵ, f) -regular partition must be large, but that the only way to get such a partition is to basically take \mathcal{A} and \mathcal{B} to be refinements of partition \mathcal{P}_s in H . Recall that we started this section by saying that one should design H in a way that will make sure that at least the proof of Theorem 2 will produce a large partition. The fact that the only way to get an (ϵ, f) -regular partition is to take partition \mathcal{P}_s , can be interpreted as saying that the only way to *prove* Theorem 2 is to go through the process described at the beginning of this section.

4 Some Preliminary Lemmas

In this section we prove some simple lemmas that will be used later on in the paper. But we start with proving the claims that were stated without proof in the previous sections. From this point on, when we write something like $x \leq_{(3)} y$, we mean that the fact that $x \leq y$ follows from the facts stated in equation (3). As the reader will inevitably notice, we will be very loose in many of the proofs. The main reason is that as we are dealing with W -type and Tower-type functions, many “improvements” make absolutely no difference even on the quantitative bounds one obtains. Hence, we opted for statements that are simpler to state and apply.

Proof of Claim 2.3: First, notice that for any $m \geq 1$, we can simply repeat the partition $A_i = \{1\}, B_i = \{2\}$, a total of m times to get m partitions of the set $\{1, 2\}$ such that there is no i

¹⁴Recall that each partition \mathcal{P}_r is a refinement of all the partitions $\mathcal{P}_{r'}$ with $r' \leq r$.

for which (distinct) j, j' appear in the same part. Since for $1 \leq m \leq 16$, we have $\phi(m) = 2^{\lceil m/16 \rceil} = 2$, the claim holds for these values of m .

Suppose now that $m \geq 17$, set $M = \phi(m) = 2^{\lceil m/16 \rceil}$ and consider a randomly generated sequence $(A_i, B_i)_{i=1}^m$ of partitions of $[M]$ obtained as follows; for each $1 \leq i \leq m$ and each $1 \leq j \leq M$ we assign element j to A_i with probability $1/2$ (all mM choices being independent). Fix a pair of distinct elements $j, j' \in [M]$. Clearly the number of i such that j, j' belong to the same class in (A_i, B_i) is distributed as the binomial random variable $B(m, 1/2)$. Hence, we get from a standard application of Chernoff's inequality that the probability that the number of these i is larger than $3m/4$ is bounded by $e^{-m/6}$. Hence, the probability that some pair of distinct $j, j' \in [M]$ belong to the same part in more than $3m/4$ of the partitions is bounded by $\binom{M}{2}e^{-m/6} < 1$ so the required sequence of partitions exists. \blacksquare

Proof of Claim 2.4: Let us start by proving that

$$T^\phi(x) \geq T(\lfloor x/2 \rfloor), \quad (6)$$

as we have previously claimed. We first notice that when $x \geq 256$ we have $2^{x/16} \geq 16x$, implying that in this case we have

$$\phi(\phi(t)) \geq 2^{2^{t/16}/16} \geq 2^t. \quad (7)$$

Now, one can verify that (6) holds when $1 \leq x \leq 10$ and that $T(x) \geq 256$ when $x \geq 4$. Thus, when $x \geq 11$, we have

$$T^\phi(x) \geq \phi(\phi(T^\phi(x-2))) \geq_{(6)} \phi(\phi(T(\lfloor x/2 \rfloor - 1))) \geq_{(7)} 2^{T(\lfloor x/2 \rfloor - 1)} = T(\lfloor x/2 \rfloor).$$

We now recall (3) which implies that since we can assume that ϵ is small enough, we can also assume that $w(1)$ is large enough. In particular we have $w(1) \gg W(1) = T(1) = 2$. Let us denote $\hat{T}(t) = \lfloor \log \log(T^\phi(t)) \rfloor$. So $w(i)$ is just \hat{T} iterated i times with $w(1) = \lfloor \log \log(1/\epsilon) \rfloor$. Now we shall show that for any large enough t , $\hat{T}(\hat{T}(t)) > T(t)$. Using induction, it would follow that for all $i \geq 1$, $w(i) > W(\lfloor i/2 \rfloor)$, thus completing the proof. Now

$$\begin{aligned} \hat{T}(\hat{T}(t)) &= \lfloor \log \log(T^\phi(\lfloor \log \log(T^\phi(t)) \rfloor)) \rfloor \\ &\geq \frac{1}{4} \log \log \left(T \left(\frac{1}{4} \log \log (T(t/4)) \right) \right) \\ &\geq \frac{1}{4} T \left(\frac{1}{4} T(t/4 - 2) - 2 \right) \\ &\geq \frac{1}{4} T \left(\frac{1}{5} T \left(\frac{t}{5} \right) \right) \\ &\geq T(t), \end{aligned}$$

where in the first inequality we apply (6), in the second we use the fact that $\log \log(T(x)) = T(x-2)$, and the last holds for all large enough t . \blacksquare

We now turn to the proof of Claim 2.7. Recall that given two sets of vertices R, R' , which are not necessarily disjoint, we used $e(R, R')$ to denote the number of edges connecting a vertex in R to a vertex in R' , where an edge belonging to $R \cap R'$ is counted twice.

Claim 4.1. *There is a constant C , such that if $m = m_b \geq C$ and \mathcal{O} is a random graph from $G(m, 1/2)$, then with probability at least $3/4$ it satisfies the first condition of a trap (as stated in Definition 2.6).*

Proof: Fix two sets R, R' of size $r = \lceil \sqrt{m}/4 \rceil$. Given distinct ℓ, ℓ' let $z_{\ell, \ell'}$ be the indicator for the event that $(\ell, \ell') \in E(\mathcal{O})$, and $z_{R, R'} = \sum_{\ell \in R, \ell' \in R'} z_{\ell, \ell'}$. Then,

$$\frac{3r^2}{8} \leq \binom{r}{2} \leq \mathbb{E}[z_{R, R'}] = \mathbb{E}[e(R, R')] = \frac{1}{2} (r^2 - |R \cap R'|) \leq \frac{r^2}{2},$$

for all large enough m . Now observe that $z_{R, R'}$ is a sum of at least $\binom{r}{2}$ indicators $z_{\ell, \ell'}$ and each $z_{\ell, \ell'}$ can change the value of $z_{R, R'}$ by at most 2. We thus get from a standard application of Chernoff's inequality that

$$\mathbb{P} \left[\left| e(R, R') - \frac{1}{2}r^2 \right| \geq \frac{1}{4}r^2 \right] \leq \mathbb{P} \left[|z_{R, R'} - \mathbb{E}[z_{R, R'}]| \geq \frac{1}{8}r^2 \right] \leq e^{-\frac{r^2}{100}}.$$

Hence the probability that there is any pair of sets R, R' satisfying $|e(R, R') - \frac{1}{2}r^2| > \frac{1}{4}r^2$ is at most

$$\binom{m}{r}^2 2^{-\frac{1}{100}r^2} \leq m^{\sqrt{m}} e^{-m/1600} \ll 1/4,$$

for all large enough m . \blacksquare

Claim 4.2. *There is a constant C , such that if $m = m_b \geq C$ and \mathcal{O} is a random graph from $G(m, 1/2)$, then with probability at least $3/4$, it satisfies the second condition of a trap (as stated in Definition 2.6).*

Proof: Let us start by considering the case $b' = b - 1$. Suppose $U_1, \dots, U_{m_{b-1}}$ is the partition of $V(\mathcal{O})$ induced by the partition \mathcal{P}_{b-1} (as discussed prior to Definition 2.6). Now recall (see Subsection 2.2) that the integers m_b satisfy the relation

$$m = m_b = m_{b-1} \phi(m_{b-1}) = m_{b-1} 2^{\lceil m_{b-1}/16 \rceil}.$$

This means that

$$\log(m) \leq m_{b-1} \leq 17 \log(m), \tag{8}$$

so the size of the sets U_i , which we will denote by h_{b-1} , satisfies

$$m/17 \log(m) \leq h_{b-1} = m/m_{b-1} \leq m/\log(m). \quad (9)$$

Fix now two sets U_i, U_j , an integer $200 \leq k \leq \log(m)$, a subset $R \subseteq U_i$ of size k^6 and a subset $R' \subseteq U_j$ of size $\lceil h_{b-1}/k \rceil$. Given distinct ℓ, ℓ' with $\ell \in R$ and $\ell' \in R'$ let $z_{\ell, \ell'}$ be the indicator for the event that $(\ell, \ell') \in E(\mathcal{O})$, and $z_{R, R'} = \sum_{\ell \in R, \ell' \in R'} z_{\ell, \ell'}$. Then

$$\begin{aligned} \frac{|R||R'|}{2} \geq \mathbb{E}[z_{R, R'}] = \mathbb{E}[e(R, R')] &= \frac{1}{2} (|R||R'| - |R \cap R'|) \\ &\geq \frac{1}{2} |R||R'| - \frac{1}{2} |R| \\ &\geq \left(\frac{1}{2} - \frac{1}{2k^2} \right) |R||R'|. \end{aligned}$$

where in the last inequality we use the facts that $k \leq \log(m)$, that $|R'| = \lceil h_{b-1}/k \rceil \stackrel{(9)}{\geq} m/17k \log(m) \geq m/17 \log^2(m)$ and that we can pick m to be large enough so that $|R'| \geq k^2$.

Note that $z_{R, R'}$ is a sum of at least $|R|(|R'| - |R|) \geq |R||R'|/2$ indicators $z_{\ell, \ell'}$ (we are using the fact that $|R| \ll |R'|$). Since each of them can change $z_{R, R'}$ by at most 2, we get from Chernoff's inequality, the fact that $k \geq 200$, and the estimate for $\mathbb{E}[z_{R, R'}]$ from the previous paragraph that

$$\begin{aligned} \mathbb{P} \left[\left| e(R, R') - \frac{1}{2} |R||R'| \right| \geq \frac{1}{k^2} |R||R'| \right] &\leq \mathbb{P} \left[|z_{R, R'} - \mathbb{E}[z_{R, R'}]| \geq \frac{1}{2k^2} |R||R'| \right] \\ &\leq e^{-\frac{|R||R'|}{64k^4}} \\ &\leq e^{-kh_{b-1}/64} \\ &\leq e^{-2h_{b-1}}. \end{aligned}$$

Now, there are $m_{b-1}^2 = O(\log^2(m))$ ways to pick the sets U_i, U_j , $O(\log(m))$ ways to choose k , $\binom{h_{b-1}}{k^6}$ ways to pick R and $\binom{h_{b-1}}{h_{b-1}/k}$ ways to pick R' . Overall, we get from a union bound that the probability that some choice of U_i, U_j, k, R and R' will violate the second condition of Definition 2.6 is bounded by

$$O(\log^3 m) \binom{h_{b-1}}{k^6} \binom{h_{b-1}}{h_{b-1}/k} e^{-2h_{b-1}} \leq m^{2k^6} (ek)^{h_{b-1}/k} e^{-2h_{b-1}} \leq m^{2 \log^6(m)} e^{-h_{b-1}}, \quad (10)$$

where in the first inequality we use the inequality $\binom{n}{k} \leq (en/k)^k$ and in the second the fact that $k \leq \log(m)$.

Let us now consider an arbitrary $b' < b$. Note that since $m_{b'} \leq m_{b-1}$, we still have $m_{b'} \leq 17 \log(m)$. Hence there are still only $O(\log^2(m))$ many ways to choose the sets $U_i^{b'}, U_j^{b'}$. This means that the upper bound obtained in (10) for the probability of partition \mathcal{P}_{b-1} violating the condition applies to any given partition $\mathcal{P}_{b'}$, with h_{b-1} replaced by $h_{b'}$. But since $h_{b'} \geq h_{b-1}$ the right hand side of the bound in (10) still holds.

We finally recall (5) stating that $m_b = T^\phi(b)$. As we noted in (6) we have $T^\phi(b) > T(\lfloor b/2 \rfloor)$. Hence the number of $b' < b$ we need to consider is only $O(\log^*(m))$. So combining this fact with the discussion in the previous paragraph we get that the probability of any partition $\mathcal{P}_{b'}$ violating the second condition of Definition 2.6 is bounded by

$$m^{3 \log^6(m)} e^{-h_{b-1}} \ll 1/4$$

where we apply the fact that $h_{b-1} \geq m/17 \log(m)$, stated in (9). ■

Proof of Claim 2.7: Follows immediately from Claims 4.1 and 4.2. ■

We will now prove two lemmas which will somewhat streamline the application of the properties of traps later on in the paper. Both lemmas will rely on the observation stated in Lemma 4.3 below. In what follows, we use $v_S \in \mathbb{R}^n$, with $S \subseteq [n]$ to denote the vector whose i^{th} entry is $1/|S|$ when $i \in S$ and 0 otherwise. Let $\mathcal{V}_k = \{v_S : S \subseteq [n], |S| = k\}$.

Lemma 4.3. *If $x \in [0, 1/k]^n$ and $\sum x_i = 1$, then x is a convex combination of the vectors of \mathcal{V}_k .*

Before we prove this lemma, we need a standard theorem from linear programming theory, which we state without proof. A polyhedron $P \subseteq \mathbb{R}^n$ is the set of points satisfying a finite number of linear inequalities. P is *bounded* if there is a constant C such that $\|x\| \leq C$ for all $x \in P$. Finally, a point $x \in P$ is said to be a *vertex* of P if it cannot be represented as a proper convex combination of two distinct points $x', x'' \in P$.

Theorem 4 ([6]). *For every bounded polyhedron $P \subseteq \mathbb{R}^n$ and $x \in P$, the point x can be written as a convex combination of the vertices of P .*

Proof of Lemma 4.3: Consider the polyhedron

$$P = \left\{ x : \sum_i x_i = 1, \text{ and } 0 \leq x_1, \dots, x_n \leq 1/k \right\}.$$

Notice that for all $x \in P$, we have $\|x\| \leq 1$. Let \mathcal{V} be the set of vertices of P . By Theorem 4, we have that any $x \in P$ is a convex combination of \mathcal{V} . So we need to show that¹⁵ $\mathcal{V} \subseteq \mathcal{V}_k$.

Suppose $u \in \mathcal{V}$. If all its entries are either 0 or $1/k$ it obviously belongs to \mathcal{V}_k . So suppose that u has an entry $u_i \in (0, 1/k)$. Then there exists at least one more entry $u_j \in (0, 1/k)$, because otherwise the entries cannot sum to 1. Let $\varepsilon_u = \frac{1}{2} \min\{u_i, u_j, 1/k - u_i, 1/k - u_j\}$. Let e_i denote the canonical basis vector where the i th entry is 1 and all the other entries are 0. Similarly define e_j . Let $u' = u + \varepsilon_u e_i - \varepsilon_u e_j$ and $u'' = u - \varepsilon_u e_i + \varepsilon_u e_j$. It can be checked that both $u', u'' \in P$ and that $u' + u'' = 2u$. So u can be written as the convex combination of two other vectors in P , which means that u is not a vertex of P . ■

¹⁵We clearly have $\mathcal{V}_k \subseteq \mathcal{V}$ but this direction is not needed.

We now turn to prove two lemmas. The first one will help us in applying the first property of traps in proving Lemma 3.2, while the second one will help us in applying the second property of traps in proving Lemma 3.1.

Lemma 4.4. *Suppose \mathcal{O} is the graph that was used when defining the trap on partition \mathcal{P}_b (so $|V(\mathcal{O})| = m_b$ and we can assume that \mathcal{O} satisfies the first condition of Definition 2.6). Let Q be the adjacency matrix of \mathcal{O} , and suppose $x, y \in [0, 1]^{m_b}$ satisfy $\sum x_i = \sum y_i = g \geq \sqrt{m_b}/2$. Then we have*

$$\left| x^T Q y - \frac{1}{2} g^2 \right| \leq \frac{1}{4} g^2.$$

Proof: The vectors x/g and y/g satisfy the condition of Lemma 4.3 with $k = \lceil \sqrt{m_b}/4 \rceil$. Hence we can express x/g and y/g as convex combinations of the vectors of \mathcal{V}_k as $x/g = \sum_R a_R v_R$ and $y/g = \sum_{R'} b_{R'} v_{R'}$. Observe further that $(v_R)^T Q v_{R'} = e(R, R')/|R||R'|$. Since $|R| = |R'| = k = \lceil \sqrt{m_b}/4 \rceil$ and we assume that \mathcal{O} satisfies the first condition of being a trap, we can infer that for any R and R' we have

$$1/4 \leq (v_R)^T Q v_{R'} \leq 3/4. \quad (11)$$

We can thus infer from (11) and the fact that $\sum_R a_R v_R$ and $\sum_{R'} b_{R'} v_{R'}$ are convex combinations that

$$\begin{aligned} (x/g)^T Q (y/g) &= \left(\sum_R a_R v_R \right)^T Q \left(\sum_{R'} b_{R'} v_{R'} \right) \\ &= \sum_{R, R'} a_R b_{R'} (v_R)^T Q v_{R'} \\ &\leq \frac{3}{4} \sum_{R, R'} a_R b_{R'} \\ &= \frac{3}{4}, \end{aligned}$$

implying that $x^T Q y \leq \frac{3}{4} g^2$. An identical argument gives $x^T Q y \geq \frac{1}{4} g^2$, which completes the proof. \blacksquare

Lemma 4.5. *Suppose \mathcal{O} is the graph that was used when defining the trap placed on partition \mathcal{P}_b (so $|V(\mathcal{O})| = m_b$ and we can assume that \mathcal{O} satisfies the second condition of Definition 2.6). Let Q be the adjacency matrix of \mathcal{O} . Let $b' < b$, set $m = m_{b'}$ and let X_1, \dots, X_m be the partition of $V(\mathcal{O})$ induced¹⁶ by $\mathcal{P}_{b'}$. Suppose each of the sets X_i has size h and let X_i, X_j be two of these sets. Suppose δ and $x, y \in [0, 1]^{m_b}$ satisfy the following conditions:*

¹⁶This was defined explicitly just before Definition 2.6. Since we are identifying the clusters of \mathcal{P}_b with the vertices of \mathcal{O} we can also identify these clusters with the indices of the adjacency matrix Q . Hence, since we think of X_i as a subset of vertices of \mathcal{O} , we can say (as we will in item 2) that an index of a vector $x \in [0, 1]^{m_b}$ belongs to X_i .

1. $1/\log(m_b) < \delta < 1/200$.
2. The vector x has non-zero entries only in X_i and y has non-zero entries only in X_j .
3. For each $1 \leq p' \leq m_b$ we have $x_{p'}/(\sum_p x_p) < \delta^6$.
4. $\sum_{p=1}^{m_b} y_p > 2\delta h$.

Then, setting $g_1 = \sum_p x_p$ and $g_2 = \sum_p y_p$, we have

$$\left| x^T Q y - \frac{1}{2} g_1 g_2 \right| \leq 2\delta^2 g_1 g_2. \quad (12)$$

Proof: Put $k = \lfloor 1/\delta \rfloor$. Then item (1) of the lemma guarantees that $200 \leq k \leq \log(m_b)$. Item (3) of the lemma guarantees that the vector x/g_1 satisfies the condition of Lemma 4.3 with respect to k^6 . Hence we can write $x/g_1 = \sum_R a_R v_R$ using the vectors of \mathcal{V}_{k^6} . Moreover, since item (2) guarantees that x has non-zero entries only in X_i we know that in the convex combination $\sum_R a_R v_R$ all the sets R satisfy $R \subseteq X_i$. Observe now that item (2) guarantees that y has non-zero entries only in X_j . Item (4) of the lemma guarantees that the vector y/g_2 satisfies the condition of Lemma 4.3 with respect to $\lceil h/k \rceil$. Hence we can write $y/g_2 = \sum_{R'} b_{R'} v_{R'}$ using the vectors of $\mathcal{V}_{\lceil h/k \rceil}$. Again, we know that in this convex combination we are only using sets $R' \subseteq X_j$.

Now, $(v_R)^T Q v_{R'} = e(R, R')/|R||R'|$. Hence, if $|R| = k^6$ and $|R'| = \lceil h/k \rceil$ and $R \subseteq X_i$, $R' \subseteq X_j$, then we can use the assumption that \mathcal{O} satisfies the second condition of being a trap, to conclude that

$$\left| (v_R)^T Q v_{R'} - \frac{1}{2} \right| \leq 1/k^2 \leq 2\delta^2. \quad (13)$$

We can thus infer from (13) and the facts that $\sum_R a_R v_R$ and $\sum_{R'} b_{R'} v_{R'}$ are convex combinations that

$$\begin{aligned} (x/g_1)^T Q (y/g_2) &= \left(\sum_R a_R v_R \right)^T Q \left(\sum_{R'} b_{R'} v_{R'} \right) \\ &= \sum_{R, R'} a_R b_{R'} (v_R)^T Q v_{R'} \\ &\leq (1/2 + 2\delta^2) \sum_{R, R'} a_R b_{R'} \\ &= (1/2 + 2\delta^2) \end{aligned}$$

implying that $x^T Q y \leq (1/2 + 2\delta^2) g_1 g_2$. An identical argument gives $x^T Q y \geq (1/2 - 2\delta^2) g_1 g_2$, which completes the proof. ■

5 Proof of Lemma 3.2

Suppose $\mathcal{A} = \{V_i : 1 \leq i \leq k\}$ and $\mathcal{B} = \{U_{i,i'} : 1 \leq i \leq k, 1 \leq i' \leq \ell\}$ (so $|\mathcal{B}| = k\ell$). We will say that a pair of sets (V_i, V_j) is *bad* if there are two sets $C_1, C_2 \subseteq [\ell] \times [\ell]$, each of size at least $\epsilon\ell^2$ such that $|d(U_{i,i_1}, U_{j,j_1}) - d(U_{i,i_2}, U_{j,j_2})| \geq 2\epsilon$ for every $(i_1, j_1) \in C_1$ and $(i_2, j_2) \in C_2$. Note that if (V_i, V_j) is bad then it cannot be good in the sense of Definition 1.1. Hence, to show that \mathcal{A} and \mathcal{B} fail to satisfy the second condition of Definition 1.1 it is enough to show that there are at least $\epsilon \binom{k}{2}$ bad pairs (V_i, V_j) . As we mentioned after the statement of Theorem 3, we will actually show that there are at least $(1 - 2\epsilon^{1/10}) \binom{k}{2}$ bad pairs.

A set $U_{i,i'}$ is called *useful* if there is an $X \in \mathcal{P}_b$ such that $U_{i,i'} \subset_{\epsilon^{1/5}} X$. If $U_{i,i'}$ is not useful, we call it *useless*. A set V_i is called *useful* if it contains¹⁷ less than $\epsilon^{1/10}\ell$ useless sets $U_{i,i'}$. If V_i is not useful, we call it *useless*. Observe that there can be at most $\epsilon^{1/10}k$ useless sets V_i , as otherwise \mathcal{B} would not be an $\epsilon^{1/5}$ -refinement of \mathcal{P}_b , which would contradict the third assumption of the lemma. Hence, there are at least $(1 - 2\epsilon^{1/10}) \binom{k}{2}$ pairs of useful sets (V_i, V_j) . By the previous paragraph, it is enough to show that every such pair is bad.

So for the rest of the proof, let us fix a pair of useful sets (V_i, V_j) . Let us assume that ϵ is small enough so that $\epsilon^{1/5} < 1/2$. Given a useful set $U_{i,i'} \subset_{\epsilon^{1/5}} X \in \mathcal{P}_b$, we let $X_{\mathcal{P}_b}(U_{i,i'})$ denote this (unique) cluster in \mathcal{P}_b that $\epsilon^{1/5}$ -contains $U_{i,i'}$. We will later prove the following claim:

Claim 5.1. *If V_i and V_j are both useful, then there are $D_1, D_2 \subseteq [\ell] \times [\ell]$ satisfying the following:*

- D_1 and D_2 have size at least $\frac{1}{32}\ell^2$.
- For every $(i_1, j_1) \in D_1$ both U_{i,i_1} and U_{j,j_1} are useful and the pair $(X_{\mathcal{P}_b}(U_{i,i_1}), X_{\mathcal{P}_b}(U_{j,j_1}))$ belongs to the trap placed on \mathcal{P}_b .
- For every $(i_2, j_2) \in D_2$ both U_{i,i_2} and U_{j,j_2} are useful and the pair $(X_{\mathcal{P}_b}(U_{i,i_2}), X_{\mathcal{P}_b}(U_{j,j_2}))$ does not belong to the trap placed on \mathcal{P}_b .

In the next subsection we prove the lemma assuming Claim 5.1, in the subsection following it we will prove this claim.

5.1 Proof of Lemma 3.2 via Claim 5.1

Let α be the weight added to H by the trap that was placed on \mathcal{P}_b . Let D_1, D_2 be the subsets of $[\ell] \times [\ell]$ guaranteed by Claim 5.1. Take any pair $(i_1, j_1) \in D_1$ and let $X_1 = X_{\mathcal{P}_b}(U_{i,i_1})$ and $X_2 = X_{\mathcal{P}_b}(U_{j,j_1})$. Since $(i_1, j_1) \in D_1$ we know that the pair (X_1, X_2) was assigned an extra weight of α by the trap placed on \mathcal{P}_b . Now consider the traps with weight larger than α , that is, the traps that were placed on partitions \mathcal{P}' which are refined by \mathcal{P}_b . Note that (X_1, X_2) might get an extra

¹⁷Recall that each V_i is the union of ℓ sets $U_{i,i'}$.

weight from a subset of these traps¹⁸. But since H contains only $\frac{1}{48}\sqrt{\log(1/\epsilon)}$ many traps, the number of ways to choose the subset of the traps with weight larger than α from which (X_1, X_2) get an extra weight is bounded by $2^{\frac{1}{48}\sqrt{\log(1/\epsilon)}} \ll \frac{1}{32\epsilon}$. Hence D_1 must have a subset of pairs of size at least $\epsilon\ell^2$, denoted D'_1 , and set of weights W_1 (all larger than α) with the following property; if $\alpha' > \alpha$ and \mathcal{P}' is the partition on which the trap with weight α' was placed then for any $(i_1, j_1) \in D'_1$ the pair $(X_{\mathcal{P}'}(U_{i_1, i_1}), X_{\mathcal{P}'}(U_{j_1, j_1}))$ belongs to the trap on \mathcal{P}' if and only if $\alpha' \in W_1$. We can also define D'_2 and W_2 in the same manner.

We now claim that we can take C_1 and C_2 (the sets showing that (V_i, V_j) is bad) to be the sets D'_1 and D'_2 . First, as noted above, both D'_1 and D'_2 have size at least $\epsilon\ell^2$. So to finish the proof we will have to show that for every $(i_1, j_1) \in D'_1$ and $(i_2, j_2) \in D'_2$ we have

$$|d(U_{i_1, i_1}, U_{j_1, j_1}) - d(U_{i_2, i_2}, U_{j_2, j_2})| \geq 2\epsilon. \quad (14)$$

Let α' be the largest weight that belongs to exactly one of the sets W_1 and W_2 . Assume without loss of generality that $\alpha' \in W_1$ and $\alpha' \notin W_2$. If there is no such weight (that is, $W_1 = W_2$) then set $\alpha' = \alpha$. We now recall Fact 2.8 which tells us that

$$\alpha' \geq 4^{-\frac{1}{48}\sqrt{\log(1/\epsilon)}}. \quad (15)$$

Let \mathcal{P}' be the partition on which the trap with weight α' was placed. Since traps with weight at least α are placed on partitions that are refined by \mathcal{P}_b , we see that if a set $U_{i, i'}$ is useful with respect to \mathcal{P}_b it must also be useful with respect to \mathcal{P}' . This means that for each pair $(i_1, j_1) \in D'_1$ the trap at \mathcal{P}' increases $d(U_{i_1, i_1}, U_{j_1, j_1})$ by at least

$$\alpha' \left(1 - \epsilon^{1/5}\right)^2 \geq \alpha'(1 - 2\epsilon^{1/5}) \geq 0.99\alpha'.$$

Similarly, for each pair $(i_2, j_2) \in D'_2$ the trap at \mathcal{P}' increases $d(U_{i_2, i_2}, U_{j_2, j_2})$ by at most

$$2\alpha'\epsilon^{1/5} \leq 0.01\alpha'.$$

Hence, disregarding for a moment all the other weights that can be assigned to these sets in H , we see that all the pairs in $(i_1, j_1) \in D'_1$ are such that $d(U_{i_1, i_1}, U_{j_1, j_1}) \geq 0.99\alpha'$ while all $(i_2, j_2) \in D'_2$ are such that $d(U_{i_2, i_2}, U_{j_2, j_2}) \leq 0.01\alpha'$. We will now show that this discrepancy is (essentially) maintained even when considering the entire graph H .

First, recall that by Fact 2.5 the total weight assigned to any pair of vertices of H in the graph G is bounded by $1/4\sqrt{\log(1/\epsilon)}$. Hence, recalling (15), we see that even after taking into account these weights, we have $d(U_{i_2, i_2}, U_{j_2, j_2}) \leq 0.02\alpha'$ for any $(i_2, j_2) \in D'_2$. Let us now consider the contribution of the weights coming from traps that were assigned a weight smaller than α' . Since these weights

¹⁸More precisely, if X_1 and X_2 are subsets of the same cluster $X' \in \mathcal{P}'$, then they will never get an extra weight from the trap placed on \mathcal{P}' . If they belong to different clusters $X'_1, X'_2 \in \mathcal{P}'$, then they will receive an extra weight only if (X'_1, X'_2) belong to the trap placed on \mathcal{P}' .

are $\alpha'/4, \alpha'/16, \dots$ their sum is bounded by $\alpha'/3$, so after taking these weights into account we still have $d(U_{i,i_2}, U_{j,j_2}) \leq 0.36\alpha'$ for any $(i_2, j_2) \in D'_2$. Let us now consider the contribution coming from traps with weight more than α' . Consider any trap with weight $\alpha'' > \alpha'$ that was placed on a partition \mathcal{P}'' . Recall that by definition of W_1, W_2 and by our choice of α' , either the extra weight α'' was added to all pairs $(X_{\mathcal{P}''}(U_{i,i'}), X_{\mathcal{P}''}(U_{j,j'}))$ with $(i', j') \in D'_1 \cup D'_2$ or to none of them. Since all the sets U_{i,i_1} and U_{j,j_1} are useful we see that for each pair $(i_1, j_1) \in D'_1$ the pair (U_{i,i_1}, U_{j,j_1}) gets from the trap at \mathcal{P}'' a total weight at least

$$\alpha'' \left(1 - \epsilon^{1/5}\right)^2 \geq \alpha''(1 - 2\epsilon^{1/5}).$$

Set w to be the sum of the weights in W_1 that are larger than α' . Then the above discussion implies that for each $(i_1, j_1) \in D'_1$ we have

$$d(U_{i,i_1}, U_{j,j_1}) \geq (1 - 2\epsilon^{1/5})w + 0.99\alpha' \geq w + 0.99\alpha' - 2\epsilon^{1/5}. \quad (16)$$

Consider now a pair $(i_2, j_2) \in D'_2$; If a weight $\alpha'' \geq \alpha'$ belongs to W_2 then it can contribute to $d(U_{i,i_2}, U_{j,j_2})$ a weight of at most α'' , hence such weights contribute to $d(U_{i,i_2}, U_{j,j_2})$ a total weight of at most¹⁹ w . As to weights $\alpha'' > \alpha'$ that do not belong to W_2 , we see that since U_{i,i_2} and U_{j,j_2} are useful, they can increase $d(U_{i,i_2}, U_{j,j_2})$ by at most $2\alpha''\epsilon^{1/5}$. As the total sum of weights of all traps is at most 1, this extra contribution is bounded by $2\epsilon^{1/5}$. All together, we see that for every $(i_2, j_2) \in D'_2$,

$$d(U_{i,i_2}, U_{j,j_2}) \leq w + 0.36\alpha' + 2\epsilon^{1/5}. \quad (17)$$

Recalling (15), we see that $4\epsilon^{1/5} < 0.1\alpha'$. Hence, (16) and (17) imply that

$$d(U_{i,i_1}, U_{j,j_1}) - d(U_{i,i_2}, U_{j,j_2}) > 0.5\alpha' >_{(15)} 2\epsilon$$

for every choice of $(i_1, j_1) \in D'_1$ and $(i_2, j_2) \in D'_2$. This establishes (14), thus completing the proof.

5.2 Proof of Claim 5.1

Let us start with observing that since V_i is assumed to be useful, it contains (more than) $\frac{1}{2}\ell$ useful sets $U_{i,i'}$. Let V'_i be the union of $\frac{1}{2}\ell$ such sets, and define V'_j in a similar way. From now on we will focus on V'_i and V'_j and their subsets $U_{i,i'}$ and $U_{j,j'}$ so we will only be talking about sets $U_{i,i'}$ and $U_{j,j'}$ that are useful. Recall that for any useful set $U_{i,i'}$ there is a (unique) set $X_{\mathcal{P}_b}(U_{i,i'}) \in \mathcal{P}_b$ such that $U_{i,i'} \subset_{\epsilon^{1/5}} X_{\mathcal{P}_b}(U_{i,i'})$.

Suppose \mathcal{P}_b has m clusters and recall that we defined the trap on \mathcal{P}_b using an m -vertex graph \mathcal{O} satisfying the first condition of Definition 2.6. That is (u, v) is an edge of \mathcal{O} if and only if (X_u, X_v) belongs to the trap on \mathcal{P}_b . Define a vector $x \in [0, 1]^m$ by setting $x_u = |V'_i \cap X_u|/|X_u|$. Define $y \in [0, 1]^m$ similarly by setting $y_u = |V'_j \cap X_u|/|X_u|$. Recall that each of the sets V_i contains a

¹⁹Recall that by choice of α' the sets W_1 and W_2 contain the same weights larger than α' .

$1/k$ -fraction of the vertices H (since $|\mathcal{A}| = k$) so $|V'_i|$ contains a $1/2k$ -fraction of the vertices of H . Since \mathcal{P}_b has order m (so there are m sets X_u) and we assume that $m \geq k^2$ (in the second item of Lemma 3.2) we infer that

$$\sum_u x_u = \sum_u y_u = \frac{m}{2k} \geq \sqrt{m}/2. \quad (18)$$

If we take Q to be the adjacency matrix of \mathcal{O} , then by (18) we can apply Lemma 4.4 (with $g = m/2k$) to infer that

$$\frac{1}{4}(m/2k)^2 \leq x^T Q y \leq \frac{3}{4}(m/2k)^2. \quad (19)$$

Given a set $U_{i,i'}$ we define a vector $x^{i'}$ by setting $x_u^{i'} = |U_{i,i'} \cap X_u|/|X_u|$. Similarly given a set $U_{j,j'}$ we define a vector $y^{j'}$ by setting $y_u^{j'} = |U_{j,j'} \cap X_u|/|X_u|$. Observe that since V'_i is the union of the sets $U_{i,i'}$ we have $x = \sum_{i'} x^{i'}$ where the sum ranges over all the $\ell/2$ indices i' for which $U_{i,i'} \subseteq V'_i$. Similarly $y = \sum_{j'} y^{j'}$ where the sum ranges over all the $\ell/2$ indices j' for which $U_{j,j'} \subseteq V'_j$. Hence, we get from (19) that

$$\frac{1}{4}(m/2k)^2 \leq \sum_{i',j'} (x^{i'})^T Q y^{j'} \leq \frac{3}{4}(m/2k)^2. \quad (20)$$

Consider now any pair i', j' in the above sum. Let $X_{u'} = X_{\mathcal{P}_b}(U_{i,i'})$ and $X_{v'} = X_{\mathcal{P}_b}(U_{j,j'})$. Recall that $U_{i,i'}$ contains a $1/k\ell$ fraction of $V(H)$ while the sets X_u contains a $1/m$ fraction of $V(H)$. This means that

$$\sum_u x_u^{i'} = m/k\ell,$$

and similarly we have

$$\sum_u y_u^{j'} = m/k\ell.$$

Hence

$$0 \leq (x^{i'})^T Q y^{j'} \leq m^2/k^2\ell^2. \quad (21)$$

More importantly, since $|U_{i,i'} \cap X_{u'}| \geq (1 - \epsilon^{1/5}) |U_{i,i'}|$ we have

$$x_{u'}^{i'} = |U_{i,i'} \cap X_{u'}|/|X_{u'}| \geq (1 - \epsilon^{1/5}) m/k\ell, \quad (22)$$

and since $|U_{j,j'} \cap X_{v'}| \geq (1 - \epsilon^{1/5}) |U_{j,j'}|$ we have

$$y_{v'}^{j'} = |U_{j,j'} \cap X_{v'}|/|X_{v'}| \geq (1 - \epsilon^{1/5}) m/k\ell. \quad (23)$$

Suppose now that $(X_{u'}, X_{v'})$ belong to the trap placed on \mathcal{P}_b , that is, that $Q_{u',v'} = 1$. We then get from (21), (22) and (23) that

$$0.99m^2/k^2\ell^2 \leq (1 - \epsilon^{1/5})^2 m^2/k^2\ell^2 \leq (x^{i'})^T Q y^{j'} \leq m^2/k^2\ell^2. \quad (24)$$

Suppose now that $(X_{u'}, X_{v'})$ does not belong to the trap placed on \mathcal{P}_b , that is, that $Q_{u',v'} = 0$. We then get from (21), (22) and (23) that

$$0 \leq (x^{i'})^T Q y^{j'} \leq 2\epsilon^{1/5}m^2/k^2\ell^2 \leq 0.01m^2/k^2\ell^2. \quad (25)$$

We thus see from (25) that the total contribution to (20) of pairs (i', j') for which $(X_{u'}, X_{v'})$ does not belong to the trap is bounded by $(\ell/2)^2 \cdot 0.01m^2/k^2\ell^2 = 0.01(m/2k)^2$. Combining (20), (24) and (25) it thus must be the case that there are at least

$$\frac{\frac{1}{4}(m/2k)^2 - 0.01(m/2k)^2}{m^2/k^2\ell^2} \geq \frac{1}{32}\ell^2,$$

pairs (i', j') for which $(X_{u'}, X_{v'})$ belongs to the trap placed on \mathcal{P}_b . Hence we can take D_1 to be the collection of these pairs. Finally, we see from (20), (24) and (25) that the number of pairs (i', j') for which $(X_{u'}, X_{v'})$ belongs to the trap on \mathcal{P}_b cannot be larger than

$$\frac{\frac{3}{4}(m/2k)^2}{0.99m^2/k^2\ell^2} \leq \frac{31}{32}\ell^2,$$

so we can take D_2 to be the collection of pairs (i', j') that do not belong to D_1 . We thus complete the proof of Claim 5.1.

6 Proof of Lemma 3.1

We will prove Lemma 3.1 by first performing a series of reductions that will culminate in Lemma 6.5. We will then spend most of this section proving Lemma 6.5. Let us first derive Lemma 3.1 from the following lemma:

Lemma 6.1. *Suppose $\gamma \leq \epsilon$ and $\mathcal{Z} = \{Z_1, \dots, Z_k\}$ is a γ -regular partition of H . Assume*

- $r < \frac{\log(1/\gamma)}{10\sqrt{\log(1/\epsilon)}}$
- $\gamma^{1/4} \leq \beta \leq 1/100$

Then, if \mathcal{Z} is a β -refinement of \mathcal{P}_{r-1} it is also an 8β -refinement of \mathcal{P}_r .

Proof that Lemma 6.1 implies Lemma 3.1: By the definition of (ϵ, f) -regularity, we get that if $|\mathcal{A}| = k$ then \mathcal{B} must be $\frac{1}{k}$ -regular. Since $k \geq 1/\epsilon$ we have $1/k \leq \epsilon$. Since \mathcal{B} is a refinement of \mathcal{P}_0 (recall that \mathcal{P}_0 is just the entire vertex set of H), it is in particular a $(1/k)^{1/4}$ -refinement of \mathcal{P}_0 . Hence, starting with $\beta = (1/k)^{1/4}$ we can repeatedly apply Lemma 6.1 (with $\gamma = 1/k$) as long as

$$r \leq \frac{\sqrt{\log(k)}}{10} \leq \frac{\log k}{10\sqrt{\log(1/\epsilon)}} \tag{26}$$

and

$$8^r/k^{1/4} \leq 1/100. \tag{27}$$

Taking $r = 2 \log \log(k)$, we thus make sure that both (26) and (27) hold²⁰ with a lot of room to spare. Hence, after these $r = 2 \log \log k$ applications of Lemma 6.1 we get that \mathcal{B} must be an $8^{2 \log \log k} / k^{1/4}$ -refinement of $\mathcal{P}_{2 \log \log k}$. Since

$$8^{2 \log \log k} / k^{1/4} \leq 1/k^{1/5} \leq \epsilon^{1/5},$$

we get that \mathcal{B} is indeed an $\epsilon^{1/5}$ -refinement of $\mathcal{P}_{2 \log \log k}$. ■

Let us now continue with the proof of Lemma 6.1. So throughout the rest of this section we assume all the conditions that are stated in the lemma. Suppose $\mathcal{P}_{r-1} = \{X_i : 1 \leq i \leq m\}$ and $\mathcal{P}_r = \{X_{i,i'} : 1 \leq i \leq m, 1 \leq i' \leq M\}$. Recall the sets $A_{i,j}, B_{i,j}$ that were used in the construction of the graph G in Subsection 2.2. With respect to these, we make the following definition:

Definition 6.2. *A pair of sets (Z_t, Z_u) is said to be β -helpful if*

1. *There are²¹ $1 \leq i, j \leq m$ such that $Z_t \subset_\beta X_i$ and $Z_u \subset_\beta X_j$ (we are not requiring $i \neq j$).*
2. *We have $\min(|Z_t \cap A_{i,j}|, |Z_t \cap B_{i,j}|) \geq \beta^2 |Z_t|$.*

We will need the following lemma, restated from [9].

Lemma 6.3. ([9]) *Let M be an integer and let $(A_j, B_j)_{j=1}^m$ be a sequence of balanced partitions of $[M]$. Let $0 < \zeta \leq 1/2$ and let $\eta, \xi > 0$ be such that*

$$(1 - \eta)(1 - 4\xi) > 1 - \zeta + \zeta^2. \quad (28)$$

Then for every sequence $\lambda = (\lambda_1, \dots, \lambda_M)$ such that $\lambda_{i'} \geq 0$ for every i' , $\|\lambda\|_1 = 1$ and $\|\lambda\|_\infty < 1 - \zeta$, there are at least ηm values of j for which $\min(\sum_{i' \in A_j} \lambda_{i'}, \sum_{i' \in B_j} \lambda_{i'}) > \xi$.

Lemma 6.4. *Suppose \mathcal{Z} is a β -refinement of \mathcal{P}_{r-1} . Then, if $Z_t \subset_\beta X_i$ for some i , but there is no i' for which $Z_t \subset_{8\beta} X_{i,i'}$, then there are at least $2\beta m$ sets X_j such that each of these sets X_j β -contains at least $\frac{k}{2m}$ sets Z_u such that (Z_t, Z_u) are β -helpful.*

Proof: Let $Z_t \subset_\beta X_i$ and suppose that there is no $1 \leq i' \leq M$ for which $Z_t \subset_{8\beta} X_{i,i'}$. Write $\lambda_{i'}$ for $|Z_t \cap X_{i,i'}|/|Z_t|$. Then $\lambda_{i'} \geq 0$ for all i' , $\|\lambda\|_1 \geq 1 - \beta$ (since $Z_t \subset_\beta X_i$) and $\|\lambda\|_\infty \leq 1 - 8\beta$ (since we assume that there is no i' for which $Z_t \subset_{8\beta} X_{i,i'}$). Set $\zeta = 7\beta/(1 - \beta) < 1/2$ and note that we have

$$(1 - 6\beta)(1 - 8\beta^2) > 1 - 6\beta - 8\beta^2 > 1 - \zeta + \zeta^2, \quad (29)$$

where in the second inequality we use the fact that $\beta < 1/100$. Define the vector $\lambda' = \lambda/\|\lambda\|_1$. Then $\|\lambda'\|_1 = 1$ and

$$\|\lambda'\|_\infty \leq (1 - 8\beta)/\|\lambda\|_1 \leq (1 - 8\beta)/(1 - \beta) = 1 - \zeta. \quad (30)$$

²⁰Recall that $k \geq 1/\epsilon$. Since Theorem 3 allows us to assume that ϵ is sufficiently small, we can assume that k is large enough so that $2 \log \log k < \frac{\sqrt{10 \log(k)}}{10}$ and that $8^{2 \log \log k} / k^{1/4} \leq 1/100$.

²¹Note that since $\beta < 1/2$ there is (at most) one choice of X_i and X_j such that $Z_t \subset_\beta X_i$ and $Z_u \subset_\beta X_j$.

Since $(A'_{i,j}, B'_{i,j})_{j=1}^m$ are balanced partitions of $[M]$, we can apply Lemma 6.3 to the vector λ' (with $\eta = 6\beta$ and $\xi = 2\beta^2$), and conclude that there are at least $6\beta m$ values of j , for which $\min(\sum_{i' \in A'_{i,j}} \lambda'_{i'}, \sum_{i' \in B'_{i,j}} \lambda'_{i'}) > 2\beta^2$. Recalling that $\lambda' = \lambda/\|\lambda\|_1$ and that $\|\lambda\|_1 \geq 1 - \beta$ this means that for each such j we have $\min(\sum_{i' \in A'_{i,j}} \lambda_{i'}, \sum_{i' \in B'_{i,j}} \lambda_{i'}) > 2\beta^2(1 - \beta) > \beta^2$. Notice that by the construction of the sets $A_{i,j}, B_{i,j}$, (that is $A_{i,j} = \cup_{i' \in A'_{i,j}} X_{i,i'}$ and $B_{i,j} = \cup_{i' \in B'_{i,j}} X_{i,i'}$) and by the definition of λ , these j 's satisfy

$$\min(|Z_t \cap A_{i,j}|, |Z_t \cap B_{i,j}|) \geq \beta^2 |Z_t|, \quad (31)$$

that is, they satisfy the second condition of being β -helpful. This means that if a set Z_u is β -contained in X_j then (Z_t, Z_u) is β -helpful. So to finish the proof, we need to show that out of the $6\beta m$ values of j that satisfy (31), at least $2\beta m$ are such that X_j β -contains at least $k/2m$ sets Z_u . Hence, it is enough to show that \mathcal{P}_{r-1} has at most $4\beta m$ sets X that β -contain less than $k/2m$ sets $Z \in \mathcal{Z}$.

Call a vertex $v \in V(H)$ *bad* if it either belongs to a set $Z \in \mathcal{Z}$ that is not β -contained in any $X \in \mathcal{P}_{r-1}$ or if it belongs to $Z \setminus X$ where $Z \subset_\beta X$. Note that since we assume that \mathcal{Z} is a β -refinement of \mathcal{P}_{r-1} then the fraction of H 's vertices that are bad is bounded by 2β . Suppose now that there are more than $4\beta m$ sets X that β -contain less than $k/2m$ sets Z . Recall that each set X contains a $1/m$ -fraction of vertices of H , while each Z contains a $1/k$ -fraction. Therefore, if X has less than $k/2m$ sets Z that are β -contained in it, then half of its vertices belong to sets Z that are either β -contained in another set X' or that are not β -contained in any set. Hence, if \mathcal{P}_{r-1} has more than $4\beta m$ such sets X , then more than 2β -fraction of H 's vertices would be bad which is impossible. \blacksquare

The main part of the proof of Lemma 6.1 will be the proof of the following lemma

Lemma 6.5. *Suppose $Z \in \mathcal{Z}$ and $X_i, X_j \in \mathcal{P}_{r-1}$. Suppose $Z \subset_\beta X_i$ and there are $\frac{k}{2m}$ sets $Z_u \subset_\beta X_j$ such that (Z, Z_u) is β -helpful. Then at least $\frac{k}{4m}$ of the sets Z_u are such that (Z, Z_u) is not γ -regular.*

We first derive Lemma 6.1 from Lemmas 6.4 and 6.5.

Proof of Lemma 6.1: By Lemma 6.4 we know that if $Z_t \subset_\beta X_i$ for some i , but there is no i' for which $Z_t \subset_{8\beta} X_{i,i'}$, then there is $S_t \subseteq [m]$ of size at least $2\beta m$ such that for any $j \in S_t$, the set X_j β -contains at least $k/2m$ sets Z_u for which (Z_t, Z_u) is β -helpful. By Lemma 6.5, each of these sets X_j β -contains at least $k/4m$ sets Z_u such that (Z_t, Z_u) is not γ -regular. Hence, all together (that is, when considering all the sets X_j where $j \in S_t$) there are at least $\beta k/2$ sets Z_u such that (Z_t, Z_u) is not γ -regular. Hence, since $\beta^2 > \gamma$ and we assume that \mathcal{Z} is γ -regular, there cannot be more than $2\beta k$ sets Z_t as above.

Since we assume that for at least $(1 - \beta)k$ of the sets Z_t there is a set X_i such that $Z_t \subset_\beta X_i$, it follows that for at least $(1 - 3\beta)k > (1 - 8\beta)k$ of the sets Z_t there exists an X_i and i' such that $Z_t \subset_{8\beta} X_{i,i'}$, which means that \mathcal{Z} is an 8β -refinement of \mathcal{P}_r . \blacksquare

In the next subsections we complete the proof of Lemma 6.1 by proving Lemma 6.5.

6.1 Setting the stage for the proof of Lemma 6.5

We start by setting some notation and observing some relations between the parameters involved. We remind the reader again that we will be assuming the conditions of Lemma 6.1. Also, hereafter we focus only on the $k/2m$ sets $Z_u \subset_\beta X_j$ such that (Z, Z_u) are β -helpful, namely the sets in the statement of Lemma 6.5.

Let us set $A = Z \cap A_{i,j}$ and $B = Z \cap B_{i,j}$. Also for each of the sets $Z_u \subset_\beta X_j$, if $|Z_u \cap A_{j,i}| \geq |Z_u \cap B_{j,i}|$ we set $W_u = Z_u \cap A_{j,i}$, otherwise we set $W_u = Z_u \cap B_{j,i}$. Since we assume that all the pairs (Z, Z_u) are β -helpful and that $\beta \geq \gamma^{1/4}$ we can deduce that

$$\min(|A|, |B|) \geq \beta^2 |Z| \geq \gamma^{1/2} |Z|, \quad (32)$$

and for all u we have

$$|W_u| \geq (1 - \beta) |Z_u| / 2 \geq |Z_u| / 4. \quad (33)$$

Let $\mathcal{P}_{r_1}, \dots, \mathcal{P}_{r_f}$ be the canonical partitions which refine \mathcal{P}_{r-1} and on which we have placed a trap. For each $1 \leq \ell \leq f$, let α_ℓ be the weight²² that was added to H when placing a trap on partition \mathcal{P}_{r_ℓ} . Recall that H contains $\frac{1}{48} \sqrt{\log(1/\epsilon)}$ many traps so

$$f \leq \frac{1}{48} \sqrt{\log(1/\epsilon)}. \quad (34)$$

Also recall that by Fact 2.8 we have that all weights $\alpha_1, \dots, \alpha_f$ satisfy

$$\alpha_1, \dots, \alpha_f \geq 4^{-\frac{1}{48} \sqrt{\log(1/\epsilon)}}. \quad (35)$$

Set

$$\delta = \frac{4^{-r}}{4 \sqrt{\log(1/\epsilon)}}, \quad (36)$$

and recall that δ is the extra weight we have added to some of the pairs (x, y) in G when considering partition \mathcal{P}_{r-1} . Since in Theorem 3 we can assume that ϵ is sufficiently small, we get from (34), (35) and (36) that

$$\delta \ll \frac{1}{f}, \alpha_1, \dots, \alpha_f. \quad (37)$$

We also observe that since $\gamma \leq \epsilon$, and Lemma 6.1 assumes that $r \leq \frac{\log(1/\gamma)}{10 \sqrt{\log(1/\epsilon)}}$ we get from (36) that

$$\gamma^{1/3} \ll \delta. \quad (38)$$

We now define a set $A' \subseteq A$ using the following iterative process. We first set $A_0 = A$. If each of the clusters $X \in \mathcal{P}_{r_1}$ is such that $|A_0 \cap X| < \delta^6 |A_0|$, then the process ends with $A' = A_0$. If there is a cluster $X \in \mathcal{P}_{r_1}$ such that $|A_0 \cap X| \geq \delta^6 |A_0|$ then we set $A_1 = |A_0 \cap X|$, and continue to the next phase. If each of the clusters $X \in \mathcal{P}_{r_2}$ is such that $|A_1 \cap X| < \delta^6 |A_1|$, then the process ends

²²So recalling the way we have defined H in Subsection 2.3, we get that if $r_\ell = b = w(g)$ then $\alpha_\ell = 4^{-g}$.

with $A' = A_1$. If there is a cluster $X \in \mathcal{P}_{r_2}$ such that $|A_1 \cap X| \leq \delta^6 |A_1|$ then we set $A_2 = |A_2 \cap X|$ and move to the next phase. So the process either stops at some level \mathcal{P}_{r_t} in which none of the clusters of \mathcal{P}_{r_t} contains more than a δ^6 -fraction of A_{t-1} , or it goes all the way to \mathcal{P}_{r_f} .

Let us make two important observations about A' . First, if the process stops at level \mathcal{P}_{r_t} (where $t \leq f$) then for any $t' > t$ we have $|A' \cap X| < \delta^6 |A'|$ for all $X \in \mathcal{P}_{r_{t'}}$. This follows from the fact that $\mathcal{P}_{r_{t'}}$ refines \mathcal{P}_{r_t} . Therefore, A' has the property, that for each partition \mathcal{P}_{r_t} the set A' is either contained in a single cluster $X \in \mathcal{P}_{r_t}$ or none of the clusters contains more than a δ^6 -fraction of A' .

The second observation is that at each iteration the process picks a subset A_i satisfying $|A_i| \geq \delta^6 |A_{i-1}|$. Since we have at most f iterations, we get that the final set A' we end up with satisfies

$$|A'| \geq \delta^{6f} |A| = \left(\frac{4^{-r}}{4\sqrt{\log(1/\epsilon)}} \right)^{6f} |A| \stackrel{(34)}{\geq} \left(\frac{4^{-r}}{4\sqrt{\log(1/\epsilon)}} \right)^{\frac{6}{48} \sqrt{\log(1/\epsilon)}} |A| \geq \epsilon^{1/4} \gamma^{1/4} |A| \geq \gamma |Z|, \quad (39)$$

where the third inequality relies on the assumption of Lemma 6.1 that $r \leq \frac{\log(1/\gamma)}{10\sqrt{\log(1/\epsilon)}}$ and the last uses (32) and the fact that $\gamma \leq \epsilon$. We now use the same process to pick a set $B' \subseteq B$ satisfying the same properties discussed above, and whose size also satisfies

$$|B'| \geq \gamma |Z|. \quad (40)$$

Take one of the sets $W = W_u$ and assume without loss of generality that $W \subseteq A_{j,i}$. Recall that by $d_G(A', W)$ and $d_G(B', W)$ we denote the densities between these sets in the graph G , that is, before adding the traps to obtain the final graph H . First note that since A', B' both belong to $X_i \in \mathcal{P}_{r-1}$ and $W \subseteq X_j$, we can infer that exactly the same weight was added in G to $d(A', W)$ and $d(B', W)$ by the partitions \mathcal{P} that are refined by \mathcal{P}_{r-1} . Now recall that we put weight δ between all the edges connecting a vertex in $A_{i,j}$ and a vertex in $A_{j,i}$ and that we did not do so for edges connecting a vertex in $B_{i,j}$ and a vertex in $A_{j,i}$. Since $A' \subseteq A_{i,j}$, $B' \subseteq B_{i,j}$ and $W \subseteq A_{j,i}$ this means that \mathcal{P}_{r-1} creates a discrepancy of δ between $d_G(A', W)$ and $d_G(B', W)$. Now recall that the weights assigned by G to the partitions \mathcal{P} which refine \mathcal{P}_{r-1} are $\delta/4, \delta/4^2, \delta/4^3, \dots$. Since the sum of these weights is at most $\delta/3$ we get that

$$|d_G(A', W) - d_G(B', W)| \geq \frac{2}{3} \delta \stackrel{(38)}{\geq} \gamma. \quad (41)$$

It thus follows from (33) (39), (40) and (41) that if we had not added the traps to G , we would have thus concluded that *every* β -helpful pair (Z, Z_u) is not γ -regular. So to finish the proof we need to show that a large number of these β -helpful pairs are not γ -regular in H as well.

Recall that $\mathcal{P}_{r_1}, \dots, \mathcal{P}_{r_f}$ are the partitions which refine \mathcal{P}_{r-1} and on which we have placed a trap. For $1 \leq \ell \leq f$ we let $d_\ell(A, B)$ be the weight added to $d(A, B)$ by the trap placed on \mathcal{P}_{r_ℓ} . We thus have the following claim:

Claim 6.6. *If (Z, Z_u) is γ -regular, then there is $1 \leq \ell \leq f$ for which*

$$|d_\ell(A', W_u) - d_\ell(B', W_u)| > 4\delta^2. \quad (42)$$

Proof: Recall that since both $A', B' \subseteq X_i \in \mathcal{P}_{r-1}$ and $W_u \subseteq X_j \in \mathcal{P}_{r-1}$ we get that $d_H(A', W_u)$ and $d_H(B', W_u)$ get the same weight from each of the traps placed on partitions $\mathcal{P}_{r'}$ that are refined by \mathcal{P}_{r-1} (that includes the case that a trap was placed on \mathcal{P}_{r-1}). This means that a discrepancy between $d_H(A', W_u)$ and $d_H(B', W_u)$ can come either from $d_G(A', W_u)$ and $d_G(B', W_u)$ or from traps placed on partitions $\mathcal{P}_{r_1}, \dots, \mathcal{P}_{r_f}$. Thus, if (42) does not hold for all $1 \leq \ell \leq f$ then we would have

$$\begin{aligned} |d_H(A', W_u) - d_H(B', W_u)| &= \left| d_G(A', W_u) - d_G(B', W_u) + \sum_{\ell=1}^f (d_\ell(A', W_u) - d_\ell(B', W_u)) \right| \\ &\geq |d_G(A', W_u) - d_G(B', W_u)| - \sum_{\ell=1}^f |(d_\ell(A', W_u) - d_\ell(B', W_u))| \\ &\geq \frac{2}{3}\delta - 4f\delta^2 \stackrel{(37)}{\geq} \frac{1}{3}\delta \stackrel{(38)}{\geq} \gamma, \end{aligned}$$

where in the second inequality we use (41). Recalling (33), (39) and (40) we thus infer that (Z, Z_u) is not γ -regular which is a contradiction. \blacksquare

Assume that for each u for which (Z, Z_u) is γ -regular, we set ℓ_u to be the *smallest* integer for which (42) holds. Recall that α_{ℓ_u} is the weight added by the trap placed on the partition $\mathcal{P}_{r_{\ell_u}}$. In the following subsection we prove Lemma 6.5 via Claim 6.7 (stated below) and in the subsection following it we prove this claim thus completing the proof of Lemma 6.5.

Claim 6.7. *If (Z, Z_u) is γ -regular, then either A' or B' satisfies the following two conditions (we write the condition with respect to A'):*

- *There is no $X \in \mathcal{P}_{r_{\ell_u}}$ such that $A' \subseteq X$.*
- *$|d_{\ell_u}(A', W_u) - \frac{1}{2}\alpha_{\ell_u}| > 2\delta^2$.*

6.2 Proof of Lemma 6.5 via Claim 6.7

Once again, let us recall that given $Z \subset_\beta X_i$ and X_j we are focusing only on the $k/2m$ sets $Z_u \subset_\beta X_j$ such that (Z, Z_u) are β -helpful. We need to show that at least $k/4m$ of the sets Z_u are such that (Z, Z_u) is not γ -regular.

Suppose to the contrary that there are $k/4m$ sets Z_u for which (Z, Z_u) is γ -regular. Then by Claim 6.7, for such Z_u either A' or B' satisfies the two conditions of Claim 6.7. Suppose without loss of generality that in at least $k/8m$ of these cases the set is A' . Also, suppose without loss of generality that out of these $k/8m$ cases, in at least $k/16m$ we have $d_{\ell_u}(A', W_u) > \alpha_{\ell_u}/2 + 2\delta^2$. Finally, since there are only f traps in the canonical partitions that refine \mathcal{P}_{r-1} , we get that there must be an integer $1 \leq \ell \leq f$ for which there are at least $k/16mf$ sets W_u for which the above holds such that $\ell_u = \ell$. So for each of these sets we have

$$d_\ell(A', W_u) > \frac{1}{2}\alpha_\ell + 2\delta^2. \tag{43}$$

For what follows we set S to be the collection of $k/16mf$ values of u for which (43) holds and such that $\ell_u = \ell$.

We now make a simple observation which relates $d_\ell(A', W_u)$, the graph \mathcal{O}_{r_ℓ} that was used to define the trap which was placed on level \mathcal{P}_{r_ℓ} and the way in which A' and W are “spread” over the clusters of \mathcal{P}_{r_ℓ} . Let m_{r_ℓ} denote the number of clusters of \mathcal{P}_{r_ℓ} (which is also the number of vertices of \mathcal{O}_{r_ℓ}). Let us use $Y_1, \dots, Y_{m_{r_\ell}}$ to denote the clusters of \mathcal{P}_{r_ℓ} . Suppose X_i and X_j each contain h clusters of \mathcal{P}_{r_ℓ} .

Let $x^a \in [0, 1]^{m_{r_\ell}}$ be the vector satisfying $x_p^a = |A' \cap Y_p|/|Y_p|$ for every $1 \leq p \leq m_{r_\ell}$. Similarly, let $x^u \in [0, 1]^{m_{r_\ell}}$ be the vector satisfying $x_p^u = |W_u \cap Y_p|/|Y_p|$ for every $1 \leq p \leq m_{r_\ell}$. If we take Q to be the adjacency matrix of \mathcal{O}_{r_ℓ} then

$$d_\ell(A', W_u) = \frac{(x^a)^T (\alpha_\ell Q) x^u}{(\sum_p x_p^a)(\sum_p x_p^u)}. \quad (44)$$

Our plan now is to show that the information we have gathered thus far contradicts Lemma 4.5. Let us start setting the stage for applying this lemma. First, as partition \mathcal{P}_b in Lemma 4.5 we will take partition \mathcal{P}_{r_ℓ} . So we are using m_{r_ℓ} as m_b in Lemma 4.5.

Second, as partition $\mathcal{P}_{b'}$ in Lemma 4.5 we will take partition \mathcal{P}_{r-1} . Note that here and in Lemma 4.5 we use m to denote the number of clusters in partitions \mathcal{P}_{r-1} and $\mathcal{P}_{b'}$ and that we use X_1, \dots, X_m to name the m clusters of both partitions. As δ in Lemma 4.5 we use the same δ used here, that is $\delta = 4^{-r}/4\sqrt{\log(1/\epsilon)}$ as defined in (36). We clearly have $\delta < 1/200$. Also, to satisfy the first condition of Lemma 4.5 we need to make sure that $\delta > 1/\log(m_{r_\ell})$, or equivalently that

$$m_{r_\ell} \stackrel{(5)}{=} T^\phi(r_\ell) \stackrel{(6)}{\geq} T(\lfloor r_\ell/2 \rfloor) > 2^{4r + \sqrt{\log(1/\epsilon)}} \stackrel{(36)}{=} 2^{1/\delta}, \quad (45)$$

We need to verify the second inequality. Recall that $r_\ell \geq r$ since we are only considering traps that were placed on partitions refining \mathcal{P}_{r-1} . Recalling (3) we also have $r_\ell \geq \log \log(1/\epsilon)$ since the first trap was placed on the partition with this index. It is easy to see that these two facts imply that the second inequality in (45) indeed holds.

As the vector x in Lemma 4.5 we will take the vector x^a defined above, and as the vector y we take $\sum_{u \in S} x^u$ with S the set defined just after equation (43). Note that since $A' \subseteq X_i$ and for all u we have $W_u \subseteq X_j$, these vectors satisfy the second condition of Lemma 4.5.

Now, by Claim 6.7 there is no cluster²³ $X \in \mathcal{P}_{r_\ell}$ such that $A' \subseteq X$. By the process we have used to define A' , this means that each of the clusters of $X \in \mathcal{P}_{r_\ell}$ contains no more than a δ^6 -fraction of the vertices of A' . This means that the vector x^a defined above satisfies the third item of Lemma 4.5.

Finally, observe that each of the sets Y_p contains a $1/mh$ -fraction of H 's vertices²⁴ while each set Z_u takes a $1/k$ -fraction. We thus get from (33) that the sum of entries of each of the vectors x^u

²³Recall that we assume that $\ell_u = \ell$ for the set W_u with $u \in S$. See the discussion at the beginning of this subsection.

²⁴Since each X_i contains a $1/m$ fraction of H 's vertices and we assumed that X_i is partitioned into h sets Y_p .

is at least $mh/4k$. Since we assume that there are at least $k/16mf$ sets W_u , we infer that the sum of entries of y is at least $h/64f \stackrel{(37)}{\geq} 2\delta h$. Hence y satisfies the fourth condition of Lemma 4.5.

Since we assume that each of the sets W_u satisfies (43), we can use the formulation of (44) to infer that

$$(x^a)^T Q x^u > (1/2 + 2\delta^2) \left(\sum_p x_p^a \right) \left(\sum_p x_p^u \right) = (1/2 + 2\delta^2) g_1 g_2^u, \quad (46)$$

where we set $g_1 = \sum_p x_p^a$ and $g_2^u = \sum_p x_p^u$. Now set $g_2 = \sum_p y_p = \sum_u g_2^u$. Summing over all vectors x^u , and applying (46) we have

$$(x^a)^T Q y = (x^a)^T Q \left(\sum_u x^u \right) > (1/2 + 2\delta^2) g_1 \sum_u g_2^u = (1/2 + 2\delta^2) g_1 g_2,$$

which contradicts (12) in Lemma 4.5.

6.3 Proof of Claim 6.7

We recall that we use α_ℓ to denote the weight added to H when placing a trap on partition \mathcal{P}_{r_ℓ} , and that for a set W_u we defined ℓ_u just before Claim 6.7.

Claim 6.8. *Set $\alpha = \alpha_{\ell_u}$. If $|d_{\ell_u}(A', W_u) - d_{\ell_u}(B', W_u)| \geq 0.4\alpha$ then (Z, Z_u) is not γ -regular.*

Proof: Recall that ℓ_u was chosen to be the smallest integer for which (42) holds. Hence

$$\left| \sum_{\ell=1}^{\ell_u-1} d_\ell(A', W_u) - d_\ell(B', W_u) \right| \leq 4f\delta^2 \stackrel{(37)}{\leq} \frac{1}{100}\alpha.$$

The assumption of this claim thus gives

$$\left| \sum_{\ell=1}^{\ell_u} d_\ell(A', W_u) - d_\ell(B', W_u) \right| \geq 0.39\alpha.$$

Since the weights assigned to traps with weight smaller than α are given by $\alpha/4, \alpha/16, \dots$, after taking into account all the traps placed on $\mathcal{P}_{r_1}, \dots, \mathcal{P}_{r_f}$ we still have

$$\left| \sum_{\ell=1}^f d_\ell(A', W_u) - d_\ell(B', W_u) \right| \geq 0.05\alpha. \quad (47)$$

As we have noted in the proof of Claim 6.6, the only traps that can create a discrepancy between $d_H(A', W_u)$ and $d_H(B', W_u)$ are those placed on $\mathcal{P}_{r_1}, \dots, \mathcal{P}_{r_f}$. Hence we can disregard the traps that were placed on partitions refined by \mathcal{P}_{r-1} , that is partitions other than $\mathcal{P}_{r_1}, \dots, \mathcal{P}_{r_f}$. Thus, (47) holds even when considering *all* the traps placed in H . Finally, by Fact 2.5 the total weight assigned to edges in G is at most $1/4\sqrt{\log(1/\epsilon)} \stackrel{(35)}{\leq} 0.01\alpha$. We thus conclude that

$$|d_H(A', W_u) - d_H(B', W_u)| \geq 0.04\alpha >_{(35)} \epsilon \geq \gamma.$$

Recalling (33), (39) and (40) we can deduce that (Z, Z_u) is not γ -regular. ■

Claim 6.9. *If there is a cluster $X \in \mathcal{P}_{r_\ell}$ such that $A' \subseteq X$ and*

$$\delta^2 \leq d_\ell(A', W_u) \leq \alpha_\ell - \delta^2, \quad (48)$$

then (Z, Z_u) is not γ -regular²⁵.

Proof: Let us define the vectors x^a and x^u as we have done just before equation (44). Let us also use the terminology used when defining these vectors. So $X = Y_q$ for some $Y_q \subseteq X_i$ implying that $x_q^a = |A'|/|Y_q|$ and all the other entries of x^a are 0. Suppose Y_1, \dots, Y_h are the clusters of \mathcal{P}_{r_ℓ} within X_j . Let \mathcal{O}_{r_ℓ} be the graph used when placing the trap on \mathcal{P}_{r_ℓ} , let $v_q \in V(\mathcal{O})$ be the vertex corresponding to cluster Y_q and let u_1, \dots, u_h be the vertices corresponding to Y_1, \dots, Y_h . Finally set $N = \{p : (v_q, u_p) \in E(\mathcal{O})\}$ to be the indices of the vertices u_1, \dots, u_h which are neighbors of v_q in \mathcal{O} . Then by (44) and (48) we have

$$\delta^2 \leq \frac{\alpha_\ell \sum_{p \in N} x_p^u}{\sum_{p=1}^h x_p^u} \leq \alpha_\ell - \delta^2,$$

implying that

$$\delta^2 \leq \frac{\sum_{p \in N} x_p^u}{\sum_{p=1}^h x_p^u} \leq 1 - \delta^2.$$

This means that if we take $W^1 = W_u \cap (\bigcup_{p \in N} Y_p)$ then

$$\delta^2 |W_u| \leq |W^1| \leq (1 - \delta^2) |W_u|. \quad (49)$$

Let $W^2 = W_u \setminus W^1$ and note that it satisfies (49) as well. A critical observation now is that our choice of N implies that for all $p \in N$ the pair (Y_q, Y_p) belongs to the trap placed on \mathcal{P}_{r_ℓ} and for all $p \notin N$ the pair (Y_q, Y_p) does not belong to this trap. This means that $d_\ell(A', W^1) = \alpha_\ell$ while $d_\ell(A', W^2) = 0$.

We will now show that we can find $W' \subseteq W^1$ and $W'' \subseteq W^2$, satisfying $|W'| \geq \epsilon^{1/10} |W^1|$, $|W''| \geq \epsilon^{1/10} |W^2|$ and

$$|d_H(A', W') - d_H(A', W'')| \geq \gamma. \quad (50)$$

Recalling (39), this will imply that (Z, Z_u) is not γ -regular as the fact that $|W'| \geq \epsilon^{1/10} |W^1|$ means that

$$|W'| \geq \epsilon^{1/10} |W^1| \stackrel{(49)}{\geq} \epsilon^{1/10} \delta^2 |W_u| \stackrel{(33)}{\geq} \frac{1}{4} \epsilon^{1/10} \delta^2 |Z_u| \geq \frac{1}{4} \gamma^{1/10} \delta^2 |Z_u| \stackrel{(38)}{\geq} \gamma |Z_u|,$$

where in the fourth inequality we use the fact that $\gamma \leq \epsilon$. A similar derivation would show that $|W''| \geq \gamma |Z_u|$.

²⁵Note that in this claim we are not assuming that $\ell = \ell_u$. That is, the claim is true for all $1 \leq \ell \leq f$. However, we will only apply it with $\ell = \ell_u$.

So we are left with picking the sets W' and W'' . Let us focus on W' . Consider some $1 \leq \ell' < \ell$. Since we assume that A' is contained in one of the clusters of \mathcal{P}_{r_ℓ} there must be a cluster $Y'_q \in \mathcal{P}_{r_{\ell'}}$ such that $A' \subseteq Y'_q$. Take some $p \in N$ and let $Y'_p \in \mathcal{P}_{r_{\ell'}}$ be the cluster containing Y_p . So we see that for each pair (Y_q, Y_p) , either all the vertices $(x, y) \in Y_q \times Y_p$ get an extra weight of $\alpha_{\ell'}$ from that trap or none of them do (depending on whether (Y'_q, Y'_p) belongs to the trap placed on $\mathcal{P}_{r_{\ell'}}$). So for each pair (Y_q, Y_p) there is a subset $S_p \subseteq [\ell - 1]$ representing those traps from which (Y_q, Y_p) got an extra weight. Recall now that H contains only $\frac{1}{48}\sqrt{\log(1/\epsilon)}$ many traps, so there are (much) less than $1/\epsilon^{1/10}$ ways to pick a set $S_p \subseteq [\ell - 1]$. So there must be a subset $N' \subseteq N$ such that $S_p = S_{p'}$ for all $p, p' \in N'$ and such that $|W^1 \cap \bigcup_{p \in N'} Y_p| \geq \epsilon^{1/10}|W^1|$. We now take $W' = W^1 \cap \bigcup_{p \in N'} Y_p$ and take S' to be the subset of $[\ell - 1]$ which is common to all $p \in N'$. Recapping the above, we see that if $\ell' \in S'$ then $d_{\ell'}(A', W') = \alpha_{\ell'}$ and if $\ell' \notin S'$ then $d_{\ell'}(A', W') = 0$. We can now define W'' and S'' in a similar way, such that if $\ell' \in S''$ then $d_{\ell'}(A', W'') = \alpha_{\ell'}$ and if $\ell' \notin S''$ then $d_{\ell'}(A', W'') = 0$.

If $S' = S''$ set $\ell' = \ell$, otherwise, let ℓ' be the smallest index that appears in exactly one of the sets S' and S'' . Also, set $\alpha = \alpha_{\ell'}$. Let us now compare $d_H(A', W')$ and $d_H(A', W'')$. By our choice of α , the traps with weight larger than α have the same contribution to both $d_H(A', W')$ and $d_H(A', W'')$. Using again the way we chose α we get that

$$\left| \sum_{\ell=1}^{\ell'} d_{\ell}(A', W') - d_{\ell}(A', W'') \right| = \alpha.$$

Now observe that the total weight added by traps with weight smaller than α is bounded by $\alpha/4 + \alpha/16 \dots < \alpha/3$. So after taking into account all traps $\mathcal{P}_{r_1}, \dots, \mathcal{P}_{r_f}$ there is still a discrepancy of at least

$$\left| \sum_{\ell=1}^f d_{\ell}(A', W') - d_{\ell}(A', W'') \right| \geq \alpha/2.$$

As in previous proofs, we do not need to consider the weight coming from traps not placed on $\mathcal{P}_{r_1}, \dots, \mathcal{P}_{r_f}$ (that is, traps placed on partitions refined by \mathcal{P}_{r-1}) since $A' \subseteq X_i \in \mathcal{P}_{r-1}$ and $W_u \subseteq X_j \in \mathcal{P}_{r-1}$. Finally, by Fact 2.5 the total weight assigned to edges in G is bounded by $1/4\sqrt{\log(1/\epsilon)} \stackrel{(35)}{\leq} \alpha/4$, so after taking into account all the weights assigned to (A', W') and (A', W'') in H we still have

$$|d_H(A', W') - d_H(A', W'')| \geq \alpha/4 \stackrel{(35)}{\geq} \epsilon \geq \gamma.$$

This proves (50) thus completing the proof. ■

Claim 6.10. *If there is a cluster $X \in \mathcal{P}_{r_{\ell_u}}$ such that $A' \subseteq X$ and a cluster $Y \in \mathcal{P}_{r_{\ell_u}}$ such that $B' \subseteq Y$ then (Z, Z_u) is not γ -regular.*

Proof: If either A' or B' satisfies (48) then Claim 6.9 implies that (Z, Z_u) is not γ -regular. So suppose both do not satisfy (48). Now note $d_{\ell_u}(A', W_u), d_{\ell_u}(B', W_u) \leq \alpha_{\ell_u}$ since α_{ℓ} is the maximum

weight a pair of sets can get from the trap placed on \mathcal{P}_{r_ℓ} . Recall that ℓ_u is an integer for which (42) holds hence one of the sets (say A') satisfies $0 \leq d_{\ell_u}(A', W_u) \leq \delta^2$ while the other satisfies $\alpha_{\ell_u} - \delta^2 \leq d_{\ell_u}(B', W_u) \leq \alpha_{\ell_u}$. But this means that

$$|d_{\ell_u}(A', W_u) - d_{\ell_u}(B', W_u)| \geq \alpha_{\ell_u} - 2\delta^2 \stackrel{(37)}{\geq} \alpha_{\ell_u}/2 ,$$

so (Z, Z_u) is not γ -regular by Claim 6.8. ■

We are now ready to complete the proof of Claim 6.7. We know from Claim 6.10 that one of the sets A' or B' must satisfy the first requirement of the claim. Suppose it is A' . If A' also satisfies the second item then we are done, so suppose it does not.

If B' also satisfies the first requirement of the claim, then since ℓ_u is chosen to satisfy (42) and since we assume that A' does not satisfy the second requirement of the lemma, we get that B' must satisfy the second requirement and we are done.

So suppose now that the B' does not satisfy the first item. If $\delta^2 \leq d_{\ell_u}(B', W_u) \leq \alpha_{\ell_u} - \delta^2$ then by Claim 6.9 (Z, Z_u) is not γ -regular, which contradicts the assumption of Claim 6.7 that (Z, Z_u) is γ -regular. Finally, if either $d_{\ell_u}(B', W_u) \geq \alpha_{\ell_u} - \delta^2$ or $d_{\ell_u}(B', W_u) \leq \delta^2$ we can combine this with the assumption that A' does not satisfy the second requirement of the claim to get that

$$|d_{\ell_u}(A', W_u) - d_{\ell_u}(B', W_u)| \geq \frac{1}{2}\alpha_{\ell_u} - 3\delta^2 \stackrel{(37)}{>} 0.4\alpha_{\ell_u} .$$

Claim 6.8 then implies that (Z, Z_u) is not γ -regular which again contradicts the assumption of Claim 6.7.

Note added. After completing this paper, we learned that D. Conlon and J. Fox have independently (and simultaneously) obtained a result similar to the one stated in Theorem 3. Their proof gives a lower bound of $W(1/\epsilon^c)$ for some $c > 0$ to the strong regularity lemma.

References

- [1] N. Alon, E. Fischer, M. Krivelevich and M. Szegedy, Efficient testing of large graphs, *Combinatorica* 20 (2000), 451-476.
- [2] N. Alon and A. Shapira, A characterization of the (natural) graph properties testable with one-sided error, *SIAM J. Comput.*, 37 (2008), 1703-1727.
- [3] N. Alon, A. Shapira and U. Stav, Can a graph have distinct regular partitions? *SIAM J. Discrete Math.*, 23 (2009), 278-287.
- [4] N. Alon and U. Stav, What is the furthest graph from a hereditary property?, *Random Structures and Algorithms* 33 (2008), 87-104.

- [5] C. Avart, V. Rödl and M. Schacht, Every monotone 3-graph property is testable, *SIAM J. Discrete Math.* 21 (2007), 73-92
- [6] D. Bertsimas and J. Tsitsiklis, **Introduction to Linear Optimization**, Athena Scientific, 1st edition, 1997.
- [7] C. Borgs, J. Chayes, L. Lovász, V.T. Sós and K. Vesztegombi, Convergent Sequences of dense graphs I: subgraph frequencies, metric properties and testing, *Adv. Math.* 219 (2008), 1801-1851.
- [8] P. Frankl and V. Rödl, Extremal problems on set systems, *Random Structures and Algorithms* 20 (2002), 131-164.
- [9] T. Gowers, Lower bounds of tower type for Szemerédi's uniformity lemma, *Geom. Funct. Anal.* 7 (1997), 322-337.
- [10] T. Gowers, Hypergraph regularity and the multidimensional Szemerédi theorem, *Ann. of Math.* 166 (2007), 897-946.
- [11] T. Gowers, Quasirandomness, counting and regularity for 3-uniform hypergraphs, *Combin. Probab. Comput.*, 15 (2006), 143-184.
- [12] R. Graham, B. Rothschild and J. Spencer, **Ramsey Theory**, 2nd Edition, Wiley, New York, 1990.
- [13] Y. Kohayakawa, B. Nagle and V. Rödl, Efficient testing of hypergraphs, *Proc. of ICALP 2002*, 1017-1028.
- [14] J. Komlós and M. Simonovits, Szemerédi's Regularity Lemma and its applications in graph theory. In: *Combinatorics, Paul Erdős is Eighty*, Vol II (D. Miklós, V. T. Sós, T. Szönyi eds.), János Bolyai Math. Soc., Budapest (1996), 295–352.
- [15] L. Lovász and B. Szegedy, Limits of dense graph sequences, *J. Combin. Theory Ser. B* 96 (2006), 933-957.
- [16] L. Lovász and B. Szegedy, Testing properties of graphs and functions, *Israel J. Math.* 178 (2010), 113-156.
- [17] B. Nagle, V. Rödl and M. Schacht, The counting lemma for regular k -uniform hypergraphs, *Random Structures and Algorithms* 28 (2006), 113-179.
- [18] V. Rödl and J. Skokan, Regularity lemma for k -uniform hypergraphs, *Random Structures and Algorithms* 25 (2004), 1-42.
- [19] K. F. Roth, On certain sets of integers, *J. Lond. Math. Soc.* 28, 104-109, 1953.

- [20] I. Z. Ruzsa and E. Szemerédi, Triple systems with no six points carrying three triangles, in *Combinatorics (Keszthely, 1976)*, Coll. Math. Soc. J. Bolyai 18, Volume II, 939-945.
- [21] V. Rödl and M. Schacht, Generalizations of the removal lemma, *Combinatorica* 29(4), 467-501
- [22] E. Szemerédi, Integer sets containing no k elements in arithmetic progression, *Acta Arith.* 27 (1975), 299-345.
- [23] E. Szemerédi, Regular partitions of graphs, In: *Proc. Colloque Inter. CNRS* (J. C. Bermond, J. C. Fournier, M. Las Vergnas and D. Sotteau, eds.), 1978, 399–401.
- [24] T. Tao, A variant of the hypergraph removal lemma, *J. Combin. Theory Ser. A* 113 (2006), 1257-1280.