THE BULLETIN OF SYMBOLIC LOGIC
Volume 27, Number 4, December 2021

NOTES ON THE STABLE REGULARITY LEMMA

M. MALLIARIS AND S. SHELAH

Abstract. This is a short expository account of the regularity lemma for stable graphs proved by the authors, with some comments on the model theoretic context, written for a general logical audience.

Some years ago, we proved a "stable regularity lemma" showing essentially that Szemerédi's celebrated regularity lemma is much stronger for graphs which do not contain large half-graphs [12, Theorem 5.18], thus characterizing the existence of irregular pairs in Szemerédi's lemma by instability in the sense of model theory. Since that time, it has been a pleasure to see the work which has grown out from this theorem, with various interesting extensions, further developments, and new directions worked out by many different colleagues. Nonetheless, it seems the clear 'picture' of the original proof has not necessarily been widely communicated. Perhaps having a short exposition available may help inspire further interactions and applications.

So in these brief expository notes we give a short overview of the proof itself and the model-theoretic ideas behind the proof. Recall that our story begins with:

THEOREM A (Szemerédi's regularity lemma, 1978). *For every $\varepsilon > 0$ there is $N(\varepsilon)$ s.t. every finite graph G may be partitioned into m classes $V_1 \cup \cdots \cup V_m$ where $m \leq N$ and*:

- *all of the pairs V_i, V_j satisfy $||V_i| - |V_j|| \leq 1$.*
- *all but at most εm^2 of the pairs (V_i, V_j) are ε-regular.*

Szemerédi's regularity lemma can be thought of as saying that huge finite graphs can be well approximated by much smaller random graphs. It was proved by Szemerédi in the course of establishing that sets of natural numbers of positive upper density contain arbitrarily long arithmetic progressions [16] and has since been extremely useful, especially in combinatorics and theoretical computer science.

It was known by work of Gowers that the bound N on the number of pieces is very large as a function of ε [6]. Regarding whether the irregular pairs can be eliminated, in [8], Section 1.8, Komlós and Simonovits write:

Received December 2, 2020.

© The Author(s), 2021. Published by Cambridge University Press on behalf of Association for Symbolic Logic
1079-8986/21/2704-0004
DOI :10.1017/bsl.2021.69

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

"*The Regularity Lemma does not assert that all pairs of clusters are regular. In fact, it allows* εk^2 *pairs to be irregular. For a long time it was not known if there must be irregular pairs at all. It turned out that there must be at least* ck *irregular pairs.*" They continue: "*Alon, Duke, Leffman, Rödl and Yuster* [2, 3] *write: 'In* [17] *the author raises the question if the assertion of the lemma holds when we do not allow any irregular pairs in the definition of a regular partition. This, however, is not true, as observed by several researchers, including Lovász, Seymour, Trotter and ourselves. A simple example showing irregular pairs are necessary is a bipartite graph with vertex classes* $A = \{a_1, \dots, a_n\}$ *and* $B = \{b_1, \dots, b_n\}$ *in which* $a_i b_j$ *is an edge iff* $i \leq j$*.'*"[1]

The stable regularity lemma will show that half-graphs *characterize* existence of irregular pairs by proving that in the absence of long half-graphs, there is a much stronger regularity lemma, which among other things, has no irregular pairs. At the time, the idea that there might be better regularity lemmas on certain sub-classes of graphs was not new: it was already known [2, 3, 9] that assuming bounded VC dimension, the number of pieces could be taken to be polynomial in $\frac{1}{\varepsilon}$ (though necessarily with irregular pairs). However, to our knowledge, the order property was not suspected by the combinatorial community be an indicator of a sea-change in structure.

To a model theorist, half-graphs are an instance of the *order property* for the graph edge relation. In the case of infinite structures, we know from the second author's *Classification Theory*, Theorem II.2.2 that the presence or absence of the order property, here the presence or absence of infinite half-graphs, is a very strong indicator of a change in structural properties, the *dividing line* at stability. One of the contributions of [12] was the idea that one might try to finitize some of the structure familiar from stability to prove that when half-graphs are small relative to the size of the finite graph, one may look for suitable finite approximations to stable behavior and thus build a much stronger regularity lemma. Note that an interesting line of work starting with Terry and Wolf [18] has since carried this idea further and into an arithmetic setting.

§1. Proof of Stable Regularity. In this section we review the Stable Regularity Lemma as it was proved in [12, Section 5]. All graphs in the paper are finite.

Given $k \in \mathbb{N}$, we say the graph G is *k-edge stable* if it contains no half-graph of length k. That is, there do not exist distinct vertices $a_1, \dots, a_k, b_1, \dots, b_k$ of G such that $R(a_i, b_j)$ holds iff $i < j$. (Note that this forbids a family of graphs, not just a bipartite half-graph.)

THEOREM 1.1 (Stable regularity lemma, [12]). *For each* $\varepsilon > 0$ *and* $k \in \mathbb{N}$ *there is* $N = N(\varepsilon, k)$ *such that for any sufficiently large finite* k*-edge stable graph* G*, for some* ℓ *with* $\ell \leq N$*,* G *can be partitioned into disjoint pieces* $A_1, \dots, A_\ell$ *and*:

(1) *the partition is equitable, i.e., the sizes of the pieces differ by at most* 1,

[1] This important graph is called the half-graph.

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

(2) *all pairs of pieces* (A_i, A_j) *are* ε*-regular, and moreover have density either* $> 1 - \varepsilon$ *or* $< \varepsilon$, *and*
(3) $N < (\frac{4}{\varepsilon})^{2^{k+3}-7}$.

For completeness, we recall the definition of ε*-regular*. However, one of the key points in the stable setting will be that we can mostly avoid working with this definition directly, by means of excellence.

DEFINITION 1.2. Given a pair of finite vertex sets A, B the density $d(A, B)$ is $e(A, B)/|A||B|$, where $e(A, B)$ is the number of edges between A and B. We say the pair (A, B) is ε *-regular* when[2] for every $A' \subseteq A$, $B' \subseteq B$ with $|A'| \geq \varepsilon|A|$, $|B'| \geq \varepsilon|B|$ we have that $|d(A, B) - d(A', B')| < \varepsilon$.

More precisely, the strategy of [12, Section 5] is to use the hypothesis of k-edge stability to divide a given graph into pieces which have an atomicity property called ε-excellence. It will follow from the definition that the distribution of edges between any two ε-excellent sets is highly uniform, in fact δ-regular for a related δ. So if we can get an equipartition into excellent pieces, we will have regularity for *all* pairs, just by construction. This is exactly what the paper does: proves a regularity lemma with excellent pieces (Theorem 1.17), from which Theorem 1.1 is easily deduced.

After giving the proof, we will make some comments on the model theoretic ideas behind it.

Comment on notation: we consider graphs model-theoretically, that is, as a set of vertices on which there is a certain symmetric irreflexive binary relation, the edge relation. This is reflected in writing things like (a) "$A \subseteq G$" to mean A is an induced subgraph of G corresponding to this set of vertices, (b) "$b \in G$" to mean b is a vertex of G, and (c) writing "size of" a graph G or of some $A \subseteq G$ to mean the number of vertices.

DEFINITION 1.3. Call $A \subseteq G$ ε*-good* if for any $b \in G$ either $|\{a \in A : R(b, a)\}| < \varepsilon|A|$ or $|\{a \in A : \neg R(b, a)\}| < \varepsilon|A|$.

In other words, any $b \in G$ (not necessarily outside A) induces a partition of A into two pieces, and ε-good means any such partition is strongly imbalanced: there is a *majority opinion* in A regarding whether to connect to b. We may express this by saying that there is a truth value $\mathbf{t} = \mathbf{t}(b, A) \in \{0, 1\}$ and that for all but $< \varepsilon|A|$ of the elements of A, $R(b, a)^{\mathbf{t}}$ holds. (In logical shorthand, for a formula φ, $\varphi^1 = \varphi$ and $\varphi^0 = \neg\varphi$.)

Note that for any $a \in G$ and any $\frac{1}{2} > \varepsilon > 0$, $\{a\}$ is ε-good. So a posteriori the definition "A is good" can be described by saying: elements of A have a majority opinion with respect to *certain specific* good sets, namely the singletons. Of course, we can ask for majority opinions with respect to larger good sets. This leads naturally to the definition of *excellent*, informally, that

[2]Informally, (A, B) is called ε*-regular* if the density of edges between these sets doesn't change too much when we replace A, B by subsets which are not too small. Some version of "not too small" is necessary to rule out choosing a single vertex on each side. One could keep track of the different uses of epsilon separately, as in [16], but this isn't necessary here.

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

elements of A have coherent, majority opinions with respect to *all* other good sets.

DEFINITION 1.4. Call $A \subseteq G$ ε *-excellent* if for any $B \subseteq G$, if B is ε-good then either $|\{a \in A : \mathbf{t}(a, B) = 1\}| < \varepsilon|A|$ or $|\{a \in A : \mathbf{t}(a, B) = 0\}| < \varepsilon|A|$.

First note that if A is ε-excellent it is ε-good. This is because single points are always trivially ε-good (and ε-excellent). In more detail, when $B = \{b\}$, $\mathbf{t}(a, \{b\})$ is simply 1 if $R(a, b)$ and 0 if $\neg R(a, b)$, so if A is excellent and $b \in G$, then either $|\{a \in A : R(a, b)\}| < \varepsilon|A|$ or $|\{a \in A : \neg R(a, b)\}| < \varepsilon|A|$. Thus A is ε-good. Thus the main existence result for ε-excellent sets, Claim 1.10, gives existence of ε-good sets by the same proof (taking the B's there to be singletons).

Second, in the notation of Definition 1.4, when A is excellent and B is good (or even excellent), any element $a \in A$ will reveal a majority opinion among elements of B regarding whether to connect to a. If A is an excellent set, most of the time the revealed opinion is the same. We may express this by saying that when A, B are excellent, $\mathbf{t} = \mathbf{t}(A, B) \in \{0, 1\}$ is well defined, where this expression means that for all but at most $< \varepsilon|A|$ of the elements $a \in A$, we have that $\mathbf{t}(a, B) = \mathbf{t}$.

Looking now towards the existence result, it isn't obvious that larger ε-excellent subsets of an arbitrary graph should exist (e.g., a random graph tends not to have nontrivial ε-good subsets); our proof will use k-edge stability in a direct way. We will use the following definition and fact from model theory, which specialized to our case says that from edge stability we may infer a specific finite bound on the height of a certain tree, which locally in this paper let us call a special tree. (So for us, a special tree is always full.)

DEFINITION 1.5. A *special tree of height n* in a graph is a configuration consisting of two sequences of vertices, $\langle b_\rho : \rho \in 2^{<n}\rangle$, called nodes, and $\langle a_\eta : \eta \in 2^n\rangle$, called leaves, with edges satisfying the following constraint: given $\eta \in 2^n$ and $\rho \in 2^{<n}$, *if* $\rho^\frown\langle \ell\rangle \trianglelefteq \eta$ *then* $R(a_\eta, b_\rho)^\ell$.

Other than the edges and non-edges which we have explicitly mentioned, anything is allowed; so as with half-graphs, asserting that there is no special tree of a certain height forbids a family of configurations.

EXAMPLE 1.6. Consider a special tree of height 2 with nodes $b_\emptyset, b_0, b_1$ and leaves $a_{00}, a_{01}, a_{10}, a_{11}$. Then the following edges and non-edges must occur: $R(b_\emptyset, a_{11})$, $R(b_\emptyset, a_{10})$, $\neg R(b_\emptyset, a_{01})$, $\neg R(b_\emptyset, a_{00})$, $R(b_1, a_{11})$, $\neg R(b_1, a_{10})$, $R(b_0, a_{01})$, $\neg R(b_0, a_{00})$.

The relevant fact relating edge stability to special trees is the following, which can be seen as a finite case of the Unstable Formula Theorem from [15].

FACT 1.7 (see e.g., Hodges [7, Lemma 6.7.9, p. 313]).
(1) *If the graph G is k-edge stable, then there is no special tree in G of height $2^{k+2} - 2$.*
(2) *If G contains no special tree of height n, then G is 2^{n+1}-edge stable.*

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

CONVENTION 1.8. *Given $k \in \mathbb{N}$, we may define $t = t(k)$, the tree bound, to be a strict upper bound on the height of a special tree in a k-edge stable graph. By the Fact above, $t \leq 2^{k+2} - 2$.*

DISCUSSION 1.9. For even finer control, we could define $t(G)$, the strict upper bound on the height of a special tree in a given graph G. The above says that if G is k -edge stable, $t(G) \leq t = t(k) \leq 2^{k+2} - 2$. Both inequalities deserve mention. First, the bounds in Fact 1.7 are not known to be tight, so t may a priori be much smaller than $2^{k+2} - 2$. Secondly, even once this is worked out, a priori some k-edge stable graphs may have shorter trees than others. So in the rest of this section, we keep track of t in addition to k, and the reader can obtain addition information by reading it as $t(G)$ when the graph is known. For other information on bounds see [5], [4].

When G is k-edge stable, we can use Fact 1.7 to partition into ε-excellent sets as follows. (Excellent sets are of course good, but notice this proof extracts ε-good sets directly if we take the B_η's to be singletons.)

CLAIM 1.10. *Suppose G is a k-edge stable graph, $t = t(k)$ is the tree bound, and $\varepsilon < \frac{1}{2^t}$. Then for every $A \subseteq G$ with $|A| \geq \frac{1}{\varepsilon^t}$ there exists an ε-excellent subset $A' \subseteq A$ with $|A'| \geq \varepsilon^{t-1}|A|$.*

PROOF. By induction on $m \leq t$ let us try to choose $\langle A_\eta : \eta \in 2^m \rangle$ and $\langle B_\eta : \eta \in 2^m \rangle$ such that:

(i) in general, for $m \geq 0$ and ρ of length m, B_ρ is an ε-good set witnessing that A_ρ is not ε-excellent;
(ii) for $m = 0$, $A_\emptyset = A$ and $B_\emptyset$ is an ε-good set witnessing that $A_\emptyset$ is not ε-excellent;
(iii) for $m > 0$ and η of length $m - 1$,
- $A_{\eta^\frown\langle 0\rangle} = \{a \in A_\eta : \mathbf{t}(a, B_\eta) = 0\}$ and
- $A_{\eta^\frown\langle 1\rangle} = \{a \in A_\eta : \mathbf{t}(a, B_\eta) = 1\}$

noting that both are nonempty, and in fact of size at least $\varepsilon|A_\eta|$, since B_η witnesses that A_η is not excellent. (Why is $\mathbf{t}(a, B_\eta)$ always defined? Because B_η is good.)

If we can indeed choose the sets B_η at each stage up to and including $m = t$, we arrive at a contradiction by building a special tree as follows. First we choose the leaves: for each $\eta \in 2^t$, let a_η be any element of the set A_η, which is nonempty as it is of size $\geq \varepsilon^t|A|$ and $|A| \geq \frac{1}{\varepsilon^t}$. Next we choose the nodes. For each $m < t$, each $\rho \in 2^m$, and each $\eta \in 2^t$ such that $\rho \trianglelefteq \eta$, the set $U_\eta = \{b \in B_\rho : R(a_\eta, b)^{1-\mathbf{t}(a_\eta, B_\rho)}\}$ is small, of size $< \varepsilon|B_\rho|$ because B_ρ is ε-good. Letting $U_\rho = \bigcup\{U_\eta : \rho \trianglelefteq \eta \in 2^t\}$, we have that $|U_\rho| < 2^t\varepsilon|B_\rho| < |B_\rho|$, with the last inequality using our assumption that $\varepsilon < \frac{1}{2^t}$. Choose b_ρ to be any element of $B_\rho \setminus U_\rho$. This constructs a special tree, which gives a contradiction.

Therefore the construction must stop before stage t, so one of the A_η must have been ε-excellent, and it will be a subset of A of size at least $\varepsilon^m|A|$ where $m =$length $(\eta) \geq t - 1$. ⊣

Our plan is now to partition the graph G into ε-excellent sets by induction: running Claim 1.10 on the graph, setting aside the excellent subset A_0,

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

running Claim 1.10 on the remainder $A = G \setminus A_0$ to obtain A_1, ... and iterating as far as we can until the leftover vertices are few enough to distribute among the excellent sets already obtained without causing much trouble. The issue to be solved is that we would like to end up with an equipartition, but as we've stated it, Claim 1.10 gives us little control on the size of the excellent sets it returns. If the construction stops at the zero-th level the excellent set will have size $|A|$; at the first level, size $\geq \varepsilon|A|$; at the second level, $\geq \varepsilon^2|A|$; and so on, with a lower bound of $\varepsilon^{t-1}|A|$. A first simple modification is to choose a short list of possible sizes in advance, as we now do.[3]

DEFINITION 1.11. Call the sequence $s_0, \ldots, s_{t-1}$ of natural numbers a *size sequence for* ε when:

(a) $\varepsilon s_\ell \geq s_{\ell+1}$ for $\ell < t-2$.
(b) s_{t-1} divides all other elements of the sequence.
(c) $s_{t-1} > t$.

CLAIM 1.12. *Suppose G is a k-edge stable graph, t is the tree bound, and $\varepsilon < \frac{1}{2^t}$. Let $s_0, \ldots, s_{t-1}$ be a size sequence for ε. Then for every $A \subseteq G$ with $|A| \geq \max\{s_0, \frac{1}{\varepsilon^n}\}$ there exists an ε-excellent subset $A' \subseteq A$ with $|A'| = s_\ell$ for some $\ell = 0, \ldots, t-1$.*

PROOF. Just as in the proof of Claim 1.10, adding to (ii) of the inductive hypothesis the condition that at stage m, both $A_{\eta^\frown\langle 0\rangle}$ and $A_{\eta^\frown\langle 1\rangle}$ have size exactly s_m. This is handled by a simple modification to the construction. At stage 0, let $A_\emptyset$ be any subset of A of size s_0. At stage m, by inductive hypothesis, the set A_η will have size s_{m-1}, so $A_{\eta^\frown\langle 0\rangle}$ and $A_{\eta^\frown\langle 1\rangle}$ will have size at least $\varepsilon s_{m-1} \geq s_m$; if necessary throw away vertices to obtain size exactly s_m *before* continuing. ⊣

At this point, two points about excellent sets come in to play. First, that any two interact uniformly:

CLAIM 1.13. *Suppose $A \subseteq G$, $B \subseteq G$ and $\varepsilon < \frac{1}{2}$ and both A, B are ε-excellent. Then for $\delta = \sqrt{2\varepsilon}$, the pair (A, B) is δ-regular, with density $d(A, B) > 1 - \delta$ or $d(A, B) < \delta$.*

PROOF. This is a straightforward calculation made in [12, Claim 5.17]. (The hypothesis of that claim says that $\mathbf{t}(A, B)$ is well defined.) ⊣

Second, for $\varepsilon < \frac{1}{2}$, observe that ε-excellence need not be preserved under subset, nor indeed is ε-goodness. (Suppose the partition induced on A by b is imbalanced, with the smaller side of size ℓ: A has a subset of size 2ℓ whose partition induced by b is exactly balanced, so this subset clearly isn't ε-good. To preserve excellence, a subset should retain roughly proportional intersection with every partition induced on A.) However, we can ensure this

[3] 1.11(b) suggests that having once fixed a size sequence in 1.11 and obtained a partition of G into excellent sets whose sizes are all elements of this sequence, we will aim for an equipartition into pieces of size s_{t-1}.

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

if we do it randomly, because k-edge stability implies the VC dimension [19] is also small. This is explained by the following, an instance of the Sauer–Shelah lemma.

FACT 1.14 (see [15, Theorem II.4.10(4), p. 72]). *If G is a k-edge stable graph, then for any finite $A \subseteq G$,*

$$|\{\{a \in A : R(a,b)\} : b \in G\}| \leq |A|^k$$

or more precisely, we can replace $|A|^k$ on the right by $\Sigma_{i \leq k} \binom{|A|}{i}$.

COROLLARY 1.15. *Whenever k is fixed, for any graph G which is k-edge stable, any $A \subseteq G$ and any $\varepsilon > 0$ the family*

$$\{\{a \in A : R(a,b)\} : b \in G\} \cup \{\{a \in A : \neg R(a,b)\} : b \in G\} \quad \subseteq \mathcal{P}(A)$$

has VC dimension $\leq k+1$.

Again, informally, call a subset of A a *trace* if it is the intersection of A with a neighborhood of some $b \in G$, or if it is the complement (in A) of the neighborhood of some $b \in G$. By the previous corollary there are relatively few traces. In order that the partition retain excellence for a related ε, it suffices that every piece of the partition intersect all of the traces of A in approximately the expected proportion.

COROLLARY 1.16. *For all $\zeta > \varepsilon > 0$ and $r, k \geq 1$, there exists $M_{1.16} = M(\zeta, \varepsilon, r, k)$ such that if*:

(a) *A is a subset of a k-edge stable graph G, $|A| \geq M_{1.16}$,*
(b) *A is ε-excellent in G, and*
(c) *the size of A is divisible by r*

then there exists a partition of A into r disjoint pieces of equal size each of which is ζ-excellent.

About the bounds: Recently Ackerman et al. [1] have computed a bound on this quantity, in the context of proving a stable regularity lemma for hypergraphs, indeed for finite structures in arbitrary finite relational languages. To read [1, Proposition 4.5] specialized to the case of graphs, the number $|\mathcal{L}|$ of non-equality symbols is 1, the maximal arity of a relation $q_{\mathcal{L}}$ is 2, and $\hat{\tau}$ is their notation for the tree bound.

Returning to the main line of [12, Section 5], here is the core result of that section, essentially [12, Theorem 5.18] (with some more information about bounds).

THEOREM 1.17 (Stable regularity lemma—version with excellence). *For each $k \geq 1$ and $\varepsilon > 0$, there are $N = N(\varepsilon, k)$ such that if G is any sufficiently large k-edge stable graph there is a partition of G into disjoint pieces $A_0, \ldots, A_{\ell-1}$ with $\ell \leq N$ such that*:

(a) *the partition is equitable, i.e., for $i, j < N$, $|A_i|, |A_j|$ differ by at most* 1,

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

(b) *for each* $i < N$, A_i *is* ε*-excellent*,
(c) *thus every pair* (A_i, A_j) *is* ε*-uniform meaning that* $\mathbf{t}(A, B) \in \{0, 1\}$ *is well defined, and*
(d) *if* $\varepsilon < \frac{1}{2^t}$, *then* $N \leq 4\left(\frac{8}{\varepsilon}\right)^{t-2}$, *where* $t \leq 2^{k+2} - 2$ *is the tree bound from* 1.8.

PROOF. Let G be a graph on n vertices. We're given k and thus $t = t(k)$, recalling that from Fact 1.7, $t \leq 2^{k+2} - 2$. Without loss of generality, to simplify notation, suppose $\varepsilon < \frac{1}{2^t}$, and ε is a fraction whose numerator is 1.

For the proof, we'll need $0 < \alpha < \beta < \varepsilon$; for definiteness, we will use $\alpha = \frac{\varepsilon}{4}$, $\beta = \frac{\varepsilon}{3}$.

Let $q = \lceil \frac{1}{\alpha} \rceil \in \mathbb{N}$, so $\frac{2}{\alpha} \geq q \geq \frac{1}{\alpha}$. Let c be a natural number which is maximal such that $q^{t-1}c \in \left(\frac{\alpha n}{2} - q^{t-1}, \frac{\alpha n}{2}\right]$. Since G is sufficiently large (see step 4), we can assume $c > t$ and define a sequence by $s_0 = q^{t-1}c, s_1 = q^{t-2}c, \ldots, s_{t-1} = c > t$. This is a size sequence (1.11): it is integer valued and satisfies the conditions on divisibility.

Step 1. By induction, construct a partition $\{B_j : j < j_*\} \cup \{B\}$ of G into α-excellent pieces each with size s_ℓ for some $\ell \in \{0, \ldots, t-1\}$, *plus* a remainder B of size $< s_0$. That is, apply Claim 1.12 to G to obtain B_0; if the remainder $G \setminus B_0$ has size $\geq s_0$, apply Claim 1.12 again to obtain B_1; continue until fewer than s_0 elements remain.

Step 2. As the elements B_j of the partition are sufficiently large, *see step* 4, we may randomly partition each of them into pieces of size s_{t-1}, i.e. c, which are β-excellent. Call this new partition $\{B'_i : i < i_*\} \cup \{B\}$, since we still have the remainder.

Step 3. Distribute the remainder among the pieces $\{B'_i : i < i_*\}$ to obtain $\{A_i : i < i_*\}$ where for all $i, j < i_*$, $||A_i| - |A_j|| \leq 1$. Since our choice of c implies $s_0 \leq \frac{\alpha n}{2}$, a short calculation shows that this new partition remains 3β-excellent. (The calculation is given in [12] Claim 5.14(3), with β here for ε' there, and noting that our assumption implies $c > \frac{1}{\beta}$.)

Step 4. As for a lower bound on n, looking at what we have used so far, it is sufficient that

(i) $\frac{\alpha^2 n}{4} - 1 > \max\{t, \frac{3}{\varepsilon}\}$ and also

(ii) $\frac{\alpha^2 n}{4} - 1 > \frac{1}{\alpha} M_{1.16}(\beta, \alpha, r, k)$ for all integer values of r in $\{\frac{\alpha n}{2q^{t-1}} - 1, \ldots, \frac{\alpha n}{2q^{t-1}}\}$, so it suffices to check for the largest, say $r = \lceil \frac{\alpha n}{2q^{t-1}} \rceil$.

Note that by our construction, $c \geq (\frac{\alpha n}{2} - q^{t-1})(q^{t-1})^{-1} = \frac{\alpha n}{2q^{t-1}} - 1 \geq \frac{\alpha^2 n}{4} - 1$. So the assumption (i) says that c, really s_{t-1}, is sufficiently large to get a size sequence and to run steps 1 and 3, and condition (ii) says that all the pieces in the partition are large enough to admit a random partition in step 2.

Step 5. To bound the number of pieces: first note that by our choice of c, $\frac{\alpha n}{2} - q^{t-1} < q^{t-1}c$, so $\frac{\alpha n}{2q^{t-1}} - 1 < c$, so $\frac{\alpha n}{4q^{t-1}} < c$. Note also that by choice of q, $q \leq \frac{2}{\alpha}$. Then we can bound $\frac{n}{c}$ by:

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

$$\frac{n}{c} \leq (n)\left(\frac{\alpha n}{4q^{t-1}}\right)^{-1} = \frac{4q^{t-1}}{\alpha} \leq \frac{4(\frac{2}{\alpha})^{t-1}}{\alpha} = 4\left(\frac{2}{\alpha}\right)^{t-2}.$$

In terms of ε, the right-hand quantity is $4\left(\frac{8}{\varepsilon}\right)^{t-2}$. ⊣

Discussion 1.18. Often such lemmas are stated with an input m corresponding to a lower bound on ℓ. It should be clear that by shrinking ε, one can ensure the minimum number of pieces is as large as desired.

Discussion 1.19. As may now be clear from the proof, the bound in 1.17(d), thus in 1.1(3) may be read aloud as "polynomial in $\frac{1}{\varepsilon}$." The specific constants there may just reflect the current proof.

Discussion 1.20. We will shortly replace t by its upper bound in terms of k, but if 1.7 is later improved, or if we are dealing with some specific G whose $t(G)$ is smaller, then the present form gives more information.

We've done the work needed for Theorem 1.1 (i.e., Conclusion 5.19 of [12]).

Proof of Theorem 1.1. Apply Theorem 1.17 with k and $\frac{\varepsilon^2}{2}$, and since the statement does not mention t, replace t by $2^{k+2} - 2$. ⊣

§2. A spectrum of regularity lemmas. Section 5 is the final section of [12]. What were the aims of the rest of the paper?

As the plural in the title suggests, the main thread of the paper investigates the structure of stable graphs by proving a spectrum of Regularity Lemmas, capturing different aspects of stability. For the first, inspired by the fact that infinite models of stable theories contain large indiscernible sets, we prove that for finite stable graphs or hypergraphs (indeed, in any finite stable relational structure, suitably defined) one can extract much larger indiscernible sets than expected from Ramsey's theorem, of size n^c rather than $\log n$ for c depending on the set of relations and their 'stability', as measured by rank. (This was well exposited, in the case of graphs, in [14], where it found a nice application.) By iteratively using this theorem, one then can, with some additional care, build a first regularity lemma for stable graphs in which all pieces are cliques or independent sets of size n^c, plus a remainder, though necessarily the number of pieces grows with the size of the graph. The second and third regularity lemmas in some sense progressively relax the 'uniformity' conditions on the pieces until arriving at the fourth, the stable regularity lemma described above, in which the pieces are now 'only' approximately uniform, i.e., ε-excellent, at the gain of the number of pieces no longer growing with the size of the graph.

We may note that for a structure M and infinite $A \subseteq M$ and ultrafilter $\mathcal{D}$ on A, there is an average type $\mathrm{Av}(A, \mathcal{D})$. For stable T we know that for suitable sets (so-called indiscernible) the filter is degenerated: the co-finite sets are enough. The notions of ε-good and ε-excellent are finitary analogues.

The reader may wonder: the importance of both stability theory and the regularity lemma were independently well understood by the early 80s.

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

Why did they come together some 30-odd years later? As we have written elsewhere, to our knowledge the first connections of Szemerédi regularity to model theory came in the context of thinking about the relation of finite and infinite combinatorics necessary to understand Keisler's order [10, 11]. It was not by mistake that [12], the first joint paper of the authors, came at the beginning of our joint work on Keisler's order. Indeed, ideas from regularity play a certain, perhaps more hidden role in our recent discovery [13] that Keisler's order has the maximum number of classes, continuum many (we may refer the interested reader to the open problems section in that paper). As our knowledge of this order develops, our understanding of the finite is also changing.

Acknowledgments. Malliaris was partially supported by NSF DMS-1553653 and by a Minerva Research Foundation membership at the IAS. Shelah was partially supported by ERC grant 338821. Both authors thank NSF grant 1362974 (Rutgers) and ERC 338821. This is paper E98 in Shelah's list. These notes benefitted from lectures in the Chicago REU over several summers. We are grateful to E. Bajo, C. Terry, and the anonymous referee for very helpful comments on the manuscript.

REFERENCES

[1] N. Ackerman, C. Freer, and R. Patel, *Stable regularity for relational structures*, preprint, 2014, arXiv:1712.09305.

[2] N. Alon, R. Duke, H. Leffman, V. Rödl, and R. Yuster, *The algorithmic aspects of the regularity lemma*. ***Foundations of Computer Science***, vol. 33 (1992), pp. 479–481.

[3] ———, *The algorithmic aspects of the regularity lemma*. ***Journal of Algorithms***, vol. 16 (1994), pp. 80–109.

[4] N. Alon, J. Fox, and Y. Zhao, *Efficient arithmetic regularity and removal lemmas for induced bipartite patterns*. ***Discrete Analysis***, vol. 3 (2019), 14 pp.

[5] D. Conlon and J. Fox, *Bounds for graph regularity and removal lemmas*. ***Geometric and Functional Analysis***, vol. 22 (2012), pp. 1192–1256.

[6] W. T. Gowers, *Lower bounds of tower type for Szemerédi's uniformity lemma*. ***Geometric and Functional Analysis***, vol. 7 (1997), pp. 322–337.

[7] W. Hodges, ***Model Theory***, Encyclopedia of Mathematics, Cambridge University Press, Cambridge, 1993.

[8] J. Komlós and M. Simonovits, *Szemerédi's Regularity Lemma and its applications in graph theory*, ***Combinatorics: Paul Erdös is Eighty***, (D. Miklós, V. T. Sós and T. Szönyi, editors), Bolyai Society Mathematical Studies, vol. 2, Keszthely, Hungary, 1996, pp. 295–352.

[9] L. Lovász and B. Szegedy, *Regularity partitions and the topology of graphons*, ***An Irregular Mind (Szemerédi is 70)***, Bolyai Society Mathematical Studies, vol. 21, Keszthely, Hungary, 2010, pp. 415–446.

[10] M. Malliaris, ***Persistence and regularity in unstable model theory***, Ph.D. thesis, Berkeley, 2009.

[11] ———, *Edge distribution and density in the characteristic sequence*. ***Annals of Pure and Applied Logic***, vol. 162 (2010), no. 1, pp. 1–19.

[12] M. Malliaris and S. Shelah, *Regularity lemmas for stable graphs*. ***Transactions of the American Mathematical Society***, vol. 366 (2014), pp. 1551–1585. First public version: arXiv:1102.3904 (Feb 2011).

[13] ———, *Keisler's order is not simple* (*and simple theories may not be either*). ***Advances in Mathematics*** 2021, Paper No. 108036, 94 pp.

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press

[14] M. Malliaris and C. Terry, *On unavoidable induced subgraphs in large prime graphs*. ***Journal of Graph Theory***, vol. 88 (2018), no. 2, pp. 255–270.

[15] S. Shelah, ***Classification Theory and the Number of Non-isomorphic Models***, first ed., North-Holland, Amsterdam, 1978, rev. ed. 1990.

[16] E. Szemerédi, *On sets of integers containing no k elements in arithmetic progression*. ***Acta Arithmetica***, vol. 27 (1975), pp. 199–245.

[17] E. Szemerédi, *Regular partitions of graphs*, ***Colloques Internationaux C.N.R.S.*** No. 260, Problèmes Combinatoires et Théorie Des Graphes, Univ. Orsay, Orsay, 1976, pp. 399–401.

[18] C. Terry and J. Wolf, *Stable arithmetic regularity in the finite-field model*. ***The Bulletin of the London Mathematical Society***, vol. 51 (2019), no. 1, pp. 70–88.

[19] V. Vapnik and A. Chervonenkis, *On the uniform convergence of relative frequencies of events to their probabilities*. ***Theory of Probability and its Applications***, vol. XVI (1971), no. 2, pp. 264–280.

DEPARTMENT OF MATHEMATICS
UNIVERSITY OF CHICAGO
5734 S. UNIVERSITY AVENUE
CHICAGO, IL 60637, USA
E-mail: mem@math.uchicago.edu

EINSTEIN INSTITUTE OF MATHEMATICS
EDMOND J. SAFRA CAMPUS, GIVAT RAM
THE HEBREW UNIVERSITY OF JERUSALEM
JERUSALEM, 91904, ISRAEL
and
DEPARTMENT OF MATHEMATICS
HILL CENTER – BUSCH CAMPUS, RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY
110 FRELINGHUYSEN ROAD, PISCATAWAY, NJ 08854, USA
E-mail: shelah@math.huji.ac.il
URL: http://shelah.logic.at

https://doi.org/10.1017/bsl.2021.69 Published online by Cambridge University Press