A Survey on the Stable Arithmetic Regularity Lemma

Course Report

Zhou Yangrui

June 2026

Abstract

This report surveys the stable arithmetic regularity lemma for subsets of finite vector spaces. The focus is on the proof and on the meaning of the main definitions. We first recall graph regularity and the stable graph regularity theorem. We then explain three related notions in the finite-field setting: Cayley-pair regularity, Fourier uniformity on cosets, and coset goodness. After that, we show how stability turns uniformity into goodness. Finally, we prove the Terry–Wolf stable arithmetic regularity lemma and compare it with Green's arithmetic regularity lemma.

Contents

Introduction

This report studies a stable version of arithmetic regularity in finite vector spaces. In the whole report, p is a fixed prime, $V = \mathbb{F}_p^n$, and $A \subseteq V$. The theorem says the following. If A is stable in the model-theoretic sense, then A can be described well by cosets of a subspace of small codimension. More precisely, for every $\mu > 0$, there is a subspace $H \leq V$ of polynomially bounded codimension such that each coset $g + H$ is almost empty or almost full inside A.

This result is part of the general idea of regularity. A regularity lemma replaces a large complicated object by a simpler object with bounded size. In graph theory, Szemerédi's regularity lemma partitions the vertex set into a bounded number of parts. Most pairs of parts then have almost constant edge density. In additive combinatorics over finite fields, Green's arithmetic regularity lemma uses cosets of a subspace instead of vertex parts [1]. For a set $A \subseteq \mathbb{F}_p^n$, one looks at the restrictions of A to cosets $g + H$. The main test is whether the balanced function

$$\mathbf{1}_A(g + h) - \mathbb{E}_{h \in H} \mathbf{1}_A(g + h)$$

has small correlation with every nontrivial character of H.

These results are useful because they reduce many questions to a quotient space. In Green's theorem, after passing to a suitable H, the set A is Fourier-uniform on most cosets of H. On such cosets, A behaves randomly enough for many counting arguments. This is similar to graph regularity. The quotient keeps the coarse information, and the uniform part is small for the patterns one wants to count. This idea is used in counting linear patterns, removal lemmas, and affine-invariant property testing [1, 2, 3].

There are two problems with ordinary regularity lemmas. First, the quantitative bounds are usually very large. Even in the finite-field setting, Green's lemma has tower-type bounds in the error parameter. Second, the conclusion often allows some exceptional parts. These two issues make later arguments harder, because one has to keep track of huge parameters and ignore exceptional pieces.

Stability gives a much better conclusion. In model theory, stability means that the relation has no long order pattern. For finite graphs, this means that the graph has no large half-graph. Malliaris and Shelah proved that stable graphs have a stronger regularity lemma: the number of parts is polynomial in the error parameter, and the pairs of parts are almost homogeneous [6]. Terry and Wolf proved the corresponding result for subsets of finite vector spaces [8]. For $A \subseteq V$, the relevant relation is

$$R_A(x, y) \iff x + y \in A.$$

We say that A is k-stable if there are no elements $a_1, \dots, a_k, b_1, \dots, b_k \in V$ such that

$$a_i + b_j \in A \iff i \le j \qquad (1 \le i, j \le k).$$

This is the additive version of forbidding a long order pattern.

We now fix the basic notation. A subspace $H \le V$ gives a partition of V into cosets $g + H$. On one coset, we translate back to H and define

$$F_{A,g,H}(h) := \mathbf{1}_A(g + h), \qquad h \in H.$$

The coset is called Fourier-uniform if every nontrivial Fourier coefficient of this function is small. The density of A on $g + H$ is

$$\alpha_{g+H} := \frac{|A \cap (g + H)|}{|H|}.$$

The coset is called good if this density is at most μ or at least $1 - \mu$. Thus a good coset is almost empty or almost full. Goodness is stronger than Fourier uniformity, because an almost constant function has small nontrivial Fourier coefficients.

The main theorem can be stated in simple terms:

> For fixed p and fixed k, every k-stable set $A \subseteq \mathbb{F}_p^n$ is, up to error μ on each coset, a union of cosets of some subspace $H \le \mathbb{F}_p^n$. Moreover, $\operatorname{codim}(H)$ is polynomially bounded in μ^{-1}.

So compared with Green's arithmetic regularity lemma, the stable theorem improves the result in two ways. The codimension bound becomes polynomial instead of tower-type. Also, every coset is controlled, not only most cosets.

The proof has two parts. The first part is Fourier analytic. If a coset $g + H$ is not uniform, then some nontrivial character detects this. The kernel of that character cuts H by one dimension. On one of the smaller cosets, the density of A changes by a definite amount. This is the density-increment step. It does not use stability.

The second part uses stability. Suppose no bounded-codimension quotient is good. Then one can keep applying the density-increment step on both the A side and the complement side. If this branching continued for too long, it would build a full binary tree for the relation $R_A(x, y) \iff x + y \in A$. Stability forbids such a tree after bounded height. Therefore the process must stop. This gives a bounded-codimension quotient in which every coset is good.

This structure is useful in applications. Once every coset is almost empty or almost full, the set A can be replaced by a subset of the small quotient V/H. Counting problems then become counting problems in V/H, with a controlled error on each coset. For removal problems and property testing, this avoids two common costs of ordinary regularity: huge bounds and exceptional parts. In this report, these are only guiding consequences. The main goal is to explain the theorem and its proof.

The report is organized as follows. Chapter 1 recalls graph regularity and stable graph regularity. Chapter 2 explains Cayley-pair regularity, Fourier uniformity, and goodness. Chapter 3 introduces stable subsets of $\mathbb{F}_p^n$ and proves that, for stable sets, uniform cosets are good. Chapter 4 reviews Green's arithmetic regularity lemma. Chapter 5 proves the stable arithmetic regularity lemma. The report ends with a comparison between the stable theorem and Green's theorem.

Chapter 1

Graph Regularity and Stable Graph Regularity

1.1 Szemerédi regularity

Let $\Gamma = (V, E)$ be a finite graph. For nonempty subsets $X, Y \subseteq V$, define $d(X, Y) := |E \cap (X \times Y)|/(|X||Y|)$.

Definition 1.1. *Let $\varepsilon > 0$. A pair (X, Y) of nonempty vertex sets is* ε-regular *if for every $X' \subseteq X$ and $Y' \subseteq Y$ with $|X'| \geq \varepsilon|X|$ and $|Y'| \geq \varepsilon|Y|$, one has $|d(X', Y') - d(X, Y)| \leq \varepsilon$.*

Regularity means the following simple fact. If we take large enough subsets of the two parts, then the edge density does not change much.

Theorem 1.2 (Szemerédi's regularity lemma)**.** *For every $\varepsilon > 0$ and every $m_0 \in \mathbb{N}$ there are integers $M = M(\varepsilon, m_0)$ and $N = N(\varepsilon, m_0)$ such that every finite graph $\Gamma = (V, E)$ with $n := |V| \geq N$ admits a partition $V = V_0 \cup V_1 \cup \cdots \cup V_k$ satisfying*

(i) $m_0 \leq k \leq M$*;*

(ii) $|V_0| < \varepsilon n$*;*

(iii) $V_1, \ldots, V_k$ *are nonempty and* $|V_1| = \cdots = |V_k|$*;*

(iv) all but at most εk^2 of the pairs (V_i, V_j), $1 \leq i, j \leq k$, are ε-regular.

Theorem 1.2 says that every large enough finite graph can be approximated by a bounded number of vertex classes. The number $N(\varepsilon, m_0)$ is needed because we want at least m_0 nonempty classes of the same size. The set V_0 is small, and only a small number of pairs of classes are allowed to be bad. For most pairs (V_i, V_j), the density between large subsets of V_i and V_j is close to the density between V_i and V_j itself. The bound $M(\varepsilon, m_0)$ does not depend on $n = |V|$. In general, however, it grows like a tower in ε^{-1}.

1.2 Stability and tree bounds

The model-theoretic condition we need is the absence of long order patterns. In the graph setting of Malliaris–Shelah, this is written as the following finite condition.

Definition 1.3 (Non-k-order property)**.** *Let $\Gamma = (V, E)$ be a finite graph with edge relation R. We say that Γ has the* k-order property *if there are $a_0, \dots, a_{k-1}, b_0, \dots, b_{k-1} \in V$ such that, for all $i < j < k$, $a_i R b_j$ and $\neg(a_j R b_i)$. If no such configuration exists, then Γ has the* non-k-order property*; equivalently, we also say that the edge relation R is* k-stable*. Up to harmless conventions about diagonal entries, this means that Γ contains no half-graph of length k.*

Stability improves regularity because it limits branching. It does not directly rule out every local witness to irregularity. Instead, it says that these witnesses cannot be arranged into a very tall binary tree. We record this tree bound first, because the same bound will be used later in the arithmetic proof.

Definition 1.4 (Binary tree configuration)**.** *Let $R \subseteq X \times Y$ be a binary relation and let $d \geq 1$. We use one-based indexing for bits: if $\eta \in 2^t$, then $\eta = (\eta_1, \dots, \eta_t)$ with $\eta_r \in \{0, 1\}$ for $1 \leq r \leq t$. A* full binary tree configuration *of height d for R consists of elements $a_\eta \in X$ for $\eta \in 2^d$ and $b_\rho \in Y$ for $\rho \in 2^{<d}$ such that, whenever $\rho \in 2^{<d}$ and $\eta \in 2^d$ extends ρ, one has*

$$R(a_\eta, b_\rho) \iff \eta_{|\rho|+1} = 1.$$

Thus $b_\varnothing$ reads the first bit η_1. If $|\rho| = s$, then b_ρ reads the next bit η_{s+1}. Along the path to a leaf η, the element a_η records exactly the bits of η.

Theorem 1.5 (Quantitative stable tree bound)**.** *Let R be a k-stable relation and set $t(k) := 2^{k+2} - 2$. Then R admits no full binary tree configuration of height $t(k)$.*

This is the standard explicit tree bound used for finite stable regularity. It comes from the tree characterization of stability in model theory [5, 6].

So if a graph has the non-k^*-order property, we may take $k^{**} := t(k^*) = 2^{k^*+2} - 2$. This number is not meant to be optimal. What matters is that it depends only on k^*. It does not depend on the size of the graph or on any ambient dimension.

The stable graph regularity lemma first gives a stronger conclusion than ordinary ε-regularity. It produces parts satisfying the notions of goodness, excellence, and uniformity below.

Definition 1.6 (Goodness, excellence, and uniformity)**.** *Let $0 < \varepsilon, \zeta < 1/2$ and let $A, B \subseteq V$ be nonempty.*

(i) A is ε-good *if for every $b \in V$ there exists a truth value $t(b, A) \in \{0, 1\}$ such that*

$$|\{a \in A : (aRb) \not\equiv t(b, A)\}| < \varepsilon |A|.$$

(ii) A is (ε, ζ)-excellent *if, whenever $B \subseteq V$ is ζ-good, there exists a truth value $t(B, A) \in \{0, 1\}$ such that*

$$|\{a \in A : t(a, B) \neq t(B, A)\}| < \varepsilon |A|.$$

*When $\varepsilon = \zeta$, we simply say that A is ε-*excellent.

*(iii) The pair (A, B) is (ε, ζ)-*uniform *if there exists $t(A, B) \in \{0, 1\}$ such that, for all but fewer than $\varepsilon|A|$ elements $a \in A$, for all but fewer than $\zeta|B|$ elements $b \in B$,*

$$aRb \equiv t(A, B).$$

*When $\varepsilon = \zeta$, we say that (A, B) is ε-*uniform.

Theorem 1.7 (Stable graph regularity, Malliaris–Shelah)**.** *Fix $k^* \geq 1$ and set $k^{**} := 2^{k^*+2} - 2$. Let $G = (V, E)$ be a finite graph with the non-k^*-order property. Then for every $\varepsilon > 0$ there are integers $N = N(k^*, \varepsilon)$ and $m = m(k^*, \varepsilon)$ such that every subset $A \subseteq V$ with $|A| \geq N$ admits a partition $A = A_1 \cup \cdots \cup A_\ell$ into nonempty pieces with $\ell \leq m$ satisfying*

(i) the partition is equitable: $||A_i| - |A_j|| \leq 1$ for all i, j;

(ii) each piece A_i is ε-excellent;

(iii) every pair (A_i, A_j) is $(\varepsilon, \varepsilon)$-uniform.

Moreover, if $\varepsilon < 2^{-k^{**}}$, one may take $m \leq (3 + \varepsilon)(8/\varepsilon)^{k^{**}}$.

The usual graph-regularity statement follows from this stronger form. After changing the parameter, Malliaris–Shelah get the following corollary.

Corollary 1.8 (Regular-form consequence)**.** *Let G have the non-k^*-order property. For every $0 < \varepsilon < 1/2$ there exist N and m such that, whenever $A \subseteq V$ and $|A| \geq N$, there is an equitable partition $A = A_1 \cup \cdots \cup A_\ell$ with $\ell \leq m$, where each A_i is ε-excellent and, for every $i < j$, the pair (A_i, A_j) is ε-regular. Moreover, if $B_i \subseteq A_i$, $B_j \subseteq A_j$, $|B_i| \geq \varepsilon|A_i|$, and $|B_j| \geq \varepsilon|A_j|$, then $d(B_i, B_j) < \varepsilon$ or $d(B_i, B_j) \geq 1 - \varepsilon$.*

Thus the stable theorem has no exceptional pairs. Still, the clean version is best stated using ε-excellence and ε-uniformity. If we only say that every pair is regular and almost empty or almost complete, then we lose some information. We also hide the condition that the subsets must be large enough. The polynomial bound on the number of parts comes from the tree parameter $k^{**} = 2^{k^*+2} - 2$.

The arithmetic theorem below is the finite-field version of this story. A set $A \subseteq \mathbb{F}_p^n$ defines a bipartite graph by the rule $x + y \in A$. In the arithmetic setting, vertex classes are replaced by cosets of subspaces. The next section explains how graph regularity, Fourier uniformity, and the stronger condition called goodness are related.

Chapter 2

Regularity, Uniformity and Goodness

This chapter separates three notions that are often mixed together. We work in $V := \mathbb{F}_p^n$. Let $A \subseteq V$, and let $H \leq V$ be a linear subspace.

2.1 Cayley graph regularity

Definition 2.1 (Bipartite Cayley graph)**.** *The* bipartite Cayley graph *of A is the bipartite graph $\Gamma_A \subseteq V \times V$ whose edge relation is $(x, y) \in \Gamma_A \iff x + y \in A$.*

Definition 2.2 (Cayley-pair regularity)**.** *Let $\varepsilon > 0$ and let $u, v \in V$. The pair of cosets $(u + H, v + H)$ is* ε-regular *for A if it is an ε-regular pair in the bipartite graph Γ_A, namely if for all $X \subseteq u + H$ and $Y \subseteq v + H$ with $|X| \geq \varepsilon|H|$ and $|Y| \geq \varepsilon|H|$, one has $|d_{\Gamma_A}(X, Y) - d_{\Gamma_A}(u + H, v + H)| \leq \varepsilon$. The quotient partition V/H is called* ε-regular in the Cayley sense *if every pair of cosets $(u + H, v + H)$ is ε-regular for A.*

This is just graph regularity applied to the bipartite graph coming from the rule $x + y \in A$.

2.2 Fourier uniformity

We use only the standard finite-field Fourier transform. Put $\omega := e^{2\pi i/p}$ and, for $x, \xi \in V = \mathbb{F}_p^n$, write $\langle x, \xi \rangle := x_1\xi_1 + \cdots + x_n\xi_n \in \mathbb{F}_p$. For $f : V \to \mathbb{C}$, define $\widehat{f}(\xi) := \mathbb{E}_{x \in V} f(x)\omega^{-\langle x, \xi \rangle}$ for $\xi \in V$. Then

$$f(x) = \sum_{\xi \in V} \widehat{f}(\xi)\omega^{\langle x, \xi \rangle} \qquad \text{and} \qquad \sum_{\xi \in V} |\widehat{f}(\xi)|^2 = \mathbb{E}_{x \in V}|f(x)|^2.$$

For a subspace $H \leq V$, define its annihilator by $H^\perp := \{\xi \in V : \langle h, \xi \rangle = 0 \text{ for all } h \in H\}$. If $F : H \to \mathbb{C}$, we write its Fourier coefficient in the direction $\xi \in V$ as $\widehat{F}_H(\xi) := \mathbb{E}_{h \in H} F(h)\omega^{-\langle h, \xi \rangle}$. This coefficient depends only on the coset $\xi + H^\perp$. The directions in

$H^\perp$ are exactly the trivial frequency on H; the nontrivial frequencies on H are represented by $\xi \notin H^\perp$.

For a coset $g + H$, set $\alpha_{g+H} := |A \cap (g + H)|/|H|$. The most direct way to look at A on $g + H$ is to translate the coset back to H and define $F_{A,g,H}(h) := \mathbf{1}_A(g + h)$ for $h \in H$; thus $\widehat{F_{A,g,H}}_H(\xi) = \mathbb{E}_{h \in H} \mathbf{1}_A(g + h)\omega^{-\langle h,\xi\rangle}$. The balanced version is $\phi_{A,g,H}(h) := F_{A,g,H}(h) - \alpha_{g+H} = \mathbf{1}_A(g+h) - \alpha_{g+H}$. Subtracting α_{g+H} only changes the trivial frequency. Indeed,

$$\widehat{F_{A,g,H}}_H(\xi) = \widehat{\phi_{A,g,H}}_H(\xi) \qquad \text{for every } \xi \notin H^\perp.$$

So the density subtraction is optional once we agree to test only the frequencies nontrivial on H.

Definition 2.3 (Fourier uniformity). *Let $\varepsilon > 0$ and let $C = g + H$ be a coset of the subspace H. We say that the coset C is* ε-uniform for A *if*

$$\sup_{\xi \notin H^\perp} \left| \widehat{F_{A,g,H}}_H(\xi) \right| \le \varepsilon.$$

Equivalently, one may replace $F_{A,g,H}$ by the balanced function $\phi_{A,g,H}$ in the same formula. The phrase "nontrivial frequency" means $\xi \notin H^\perp$, not merely $\xi \neq 0$.

When $g = 0$, this also says that the linear subspace $H = 0 + H$ is ε-uniform for A, viewed as one particular coset.

The quotient partition V/H is called ε-uniform for A *if every coset $g + H \in V/H$ is ε-uniform for A.*

Thus uniformity is a Fourier condition on one coset. After translating the coset back to H, we test A against every nontrivial character of H. If all these correlations are small, the coset is uniform. Saying that V/H is uniform means that every coset has this property.

2.3 Goodness

Definition 2.4 (Goodness). *Let $\mu > 0$ and let $C = g + H$ be a coset of H. We say that C is* μ-good for A *if $|A \cap (g + H)|/|H| \le \mu$ or $|A \cap (g + H)|/|H| \ge 1 - \mu$. When $g = 0$, this also says that the linear subspace $H = 0 + H$ is μ-good for A, viewed as one particular coset.*

The quotient partition V/H is called μ-good for A *if every coset $g + H \in V/H$ is μ-good for A.*

Goodness is also a condition on one coset. It says that the coset is almost empty or almost full. At the level of the partition V/H, it says that every coset has this property. Then, up to error μ on each coset, the set A is simply a union of cosets of H.

2.4 General implications

Proposition 2.5 (Coset goodness implies coset uniformity). *Let $A \subseteq V$, let $H \le V$, let $g \in V$, and let $\mu > 0$. If the coset $g + H$ is μ-good for A, then the coset $g + H$ is*

μ-*uniform for* A.

Proof. Write $\alpha := |A \cap (g+H)|/|H|$. If $\alpha \le \mu$, then for every $\xi \notin H^\perp$,

$$|\widehat{F_{A,g,H}}_H(\xi)| = \left|\mathbb{E}_{h\in H}\mathbf{1}_A(g+h)\omega^{-\langle h,\xi\rangle}\right| \le \mathbb{E}_{h\in H}\mathbf{1}_A(g+h) = \alpha \le \mu.$$

If $\alpha \ge 1-\mu$, apply the same argument to the complement of A on the coset $g+H$. Since the constant function has zero Fourier coefficient at every nontrivial frequency,

$$\widehat{F_{A,g,H}}_H(\xi) = -\widehat{F_{V\setminus A,g,H}}_H(\xi) \qquad (\xi \notin H^\perp),$$

and therefore $|\widehat{F_{A,g,H}}_H(\xi)| \le 1-\alpha \le \mu$. □

Corollary 2.6. *If the quotient partition* V/H *is* μ*-good for* A*, then* V/H *is* μ*-uniform for* A.

Proposition 2.7 (Uniformity implies Cayley regularity)**.** *Let* $\varepsilon > 0$*, let* $H \le V$*, and let* $u, v \in V$*. If the coset* $u+v+H$ *is* ε^2*-uniform for* A*, then the pair* $(u+H, v+H)$ *is* ε*-regular for* A *in the Cayley graph* Γ_A.

Proof. Set $B := (A-(u+v))\cap H \subseteq H$ and let $\alpha := |B|/|H|$. Then $d_{\Gamma_A}(u+H, v+H) = \alpha$. Indeed, write $x = u+h, \qquad y = v+k \qquad (h,k \in H)$. By the definition of the bipartite Cayley graph,

$$(x,y) \in \Gamma_A \iff x+y \in A \iff u+v+h+k \in A \iff h+k \in B.$$

For each fixed $h \in H$, there are exactly $|B|$ choices of $k \in H$ such that $h+k \in B$. Hence $e_{\Gamma_A}(u+H, v+H) = |H|\,|B|$, and therefore

$$d_{\Gamma_A}(u+H, v+H) = \frac{|H|\,|B|}{|H|^2} = \frac{|B|}{|H|} = \alpha.$$

Let $X \subseteq u+H$ and $Y \subseteq v+H$ satisfy

$$|X| \ge \varepsilon|H|, \qquad |Y| \ge \varepsilon|H|.$$

Write $X' := X-u \subseteq H$ and $Y' := Y-v \subseteq H$. Then

$$d_{\Gamma_A}(X,Y) = \frac{1}{|X'||Y'|}\sum_{x\in X'}\sum_{y\in Y'}\mathbf{1}_B(x+y).$$

Subtracting α and writing $f := \mathbf{1}_B - \alpha$, we get

$$d_{\Gamma_A}(X,Y) - \alpha = \frac{|H|^2}{|X'||Y'|}\mathbb{E}_{x,y\in H} f(x+y)\mathbf{1}_{X'}(x)\mathbf{1}_{Y'}(y).$$

Using Fourier inversion on H in finite-field notation,

$$f(x+y) = \sum_{\xi\in V/H^\perp}\widehat{f}_H(\xi)\omega^{\langle x+y,\xi\rangle},$$

where the sum is over one representative of each coset of $H^\perp$. Therefore

$$\mathbb{E}_{x,y\in H} f(x+y)\mathbf{1}_{X'}(x)\mathbf{1}_{Y'}(y) = \sum_{\xi\in V/H^\perp} \widehat{f}_H(\xi)\widehat{\mathbf{1}_{X'}}_H(-\xi)\widehat{\mathbf{1}_{Y'}}_H(-\xi).$$

Since f is balanced, the trivial coefficient is zero. Since $u+v+H$ is ε^2-uniform for A, and since $f = \mathbf{1}_B - \alpha$ has the same nontrivial Fourier coefficients as $\mathbf{1}_B$, every nontrivial coefficient of f has magnitude at most ε^2. Hence

$$|\mathbb{E}_{x,y\in H} f(x+y)\mathbf{1}_{X'}(x)\mathbf{1}_{Y'}(y)| \le \varepsilon^2 \sum_{\xi\in V/H^\perp} \left|\widehat{\mathbf{1}_{X'}}_H(\xi)\widehat{\mathbf{1}_{Y'}}_H(\xi)\right|.$$

By Cauchy–Schwarz and Parseval on H,

$$\begin{aligned}
&\sum_{\xi\in V/H^\perp} \left|\widehat{\mathbf{1}_{X'}}_H(\xi)\widehat{\mathbf{1}_{Y'}}_H(\xi)\right| \\
&\le \left(\sum_{\xi\in V/H^\perp} |\widehat{\mathbf{1}_{X'}}_H(\xi)|^2\right)^{1/2} \left(\sum_{\xi\in V/H^\perp} |\widehat{\mathbf{1}_{Y'}}_H(\xi)|^2\right)^{1/2} \\
&= \sqrt{\frac{|X'|}{|H|}}\sqrt{\frac{|Y'|}{|H|}}.
\end{aligned}$$

Therefore

$$|d_{\Gamma_A}(X,Y) - \alpha| \le \varepsilon^2 \frac{|H|}{\sqrt{|X'||Y'|}} \le \varepsilon.$$

Since $\alpha = d_{\Gamma_A}(u+H, v+H)$, the pair $(u+H, v+H)$ is ε-regular. □

Corollary 2.8. *If the quotient partition V/H is ε^2-uniform for A, then V/H is ε-regular in the Cayley sense.*

Green's arithmetic regularity lemma uses a related but weaker condition. It asks for Fourier uniformity on most cosets. It does not ask for Cayley regularity on every pair of cosets.

Definition 2.9 (Green regularity). *Let $\varepsilon > 0$. The quotient partition V/H is called* ε-regular for A in Green's arithmetic sense *if all but at most an ε-fraction of the cosets $g+H \in V/H$ are ε-uniform for A.*

Chapter 3

Stable Sets in the Arithmetic Setting

3.1 The associated relation

Let $A \subseteq G$. The correct way to import stability into additive combinatorics is to regard A as defining the bipartite relation $R_A(x, y) \iff x + y \in A$.

Definition 3.1. *Let $A \subseteq G$ and let $k \geq 1$. We say that A is* k-stable *if the relation R_A is k-stable. Equivalently, there do not exist $a_1, \dots, a_k, b_1, \dots, b_k \in G$ such that $a_i + b_j \in A \iff i \leq j$ for all $1 \leq i, j \leq k$.*

This is the definition used by Terry and Wolf. It is unchanged by translation. It is also stable, up to changing the parameter by a bounded amount, under the basic operations used in the proof.

Lemma 3.2. *Let $A \subseteq G$ be k-stable.*

(1) For every $g \in G$, the translate $A + g$ is k-stable.

(2) The complement $G \setminus A$ is $(k + 1)$-stable.

(3) If $B \subseteq G$ is definable from A by fixing one coordinate in the relation R_A, then B has bounded stability depending only on k.

(4) Any fixed finite Boolean combination of translates of A has bounded stability depending only on k and on the Boolean complexity.

Proof. Part (1) is immediate from the definition. Part (2) is the standard finite stability fact that reversing membership can increase the half-graph parameter by at most one. Part (3) follows from monotonicity: passing to induced subrelations or fibres cannot create order patterns of unbounded size. Part (4) is the corresponding closure property of stable relations under fixed finite Boolean combinations; quantitatively, the new stability bound depends only on the old bound and the number of formulas being combined. □

Lemma 3.3 (Quantitative stability closure used in the proof)**.** *The following closure facts hold.*

(1) Every subgroup of an abelian group is 2*-stable. Moreover, if* $L \leq G$ *is a subgroup and* $B \subseteq G$ *is* r*-stable, then* $B \cap L$ *is* r*-stable.*

(2) There is a function $\mathfrak{b}(r, s)$ *such that whenever* $B \subseteq G$ *is* r*-stable and* $C \subseteq G$ *is* s*-stable, both* $B \cap C$ *and* $B \setminus C$ *are* $\mathfrak{b}(r, s)$*-stable.*

Proof. For (1), a subgroup is 2-stable by the elementary cancellation argument: if three corners of a 2-order pattern lie in the subgroup, then the fourth corner lies there as well. The same argument shows that intersecting an r-stable set with a subgroup does not create an r-order pattern: an order pattern in $B \cap L$ would force the missing subgroup-membership entries into L, and hence would already give an order pattern in B.

Part (2) is the standard Ramsey closure of stability under finite Boolean combinations. If, for example, $B \cap C$ had arbitrarily large order patterns, then colouring the relevant positive and negative incidences according to membership in B and C and passing to a large homogeneous subsequence would yield a large order pattern in either B or C (or in a complement, which only changes the stability parameter by a bounded amount). The exact numerical form of $\mathfrak{b}$ is not important here; only that it depends on r and s alone. $\square$

3.2 The relation of regularity, uniformity and goodness in stable subsets

Lemma 3.4 (Large stable regular pairs are homogeneous)**.** *Fix* $k \geq 2$ *and* $\mu \in (0, 1/2)$*, and set* $\rho_0 := (\mu/4)^k$*. For every* $0 < \rho \leq \rho_0$ *and every integer* $m > 2(4/\mu)^k$*, the following holds. Let* $R \subseteq X \times Y$ *be a* k*-stable bipartite relation with* $|X| = |Y| \geq m$*. If* (X, Y) *is* ρ*-regular with respect to* R*, then* $d_R(X, Y) \leq \mu$ *or* $d_R(X, Y) \geq 1 - \mu$*.*

Proof. Set $\beta := \mu/4$ and $\rho_0 := \beta^k$. Since $\rho \leq \rho_0$, every ρ-regular pair is also ρ_0-regular: the size condition becomes weaker and the allowed error becomes larger. It is therefore enough to prove the result for the parameter ρ_0.

The assumption on m says precisely that $\beta^k m > 2$. In particular, whenever a set has size at least $\beta^k|X|$, it has size greater than 2.

Assume, towards a contradiction, that (X, Y) is ρ_0-regular and that $\mu < \alpha := d_R(X, Y) < 1 - \mu$. We will construct elements $x_1, \ldots, x_k \in X$, $y_1, \ldots, y_k \in Y$ such that $R(x_i, y_j) \iff i \leq j$, contradicting k-stability.

We construct them inductively together with sets $X_i \subseteq X$ and $Y_i \subseteq Y$. The intended invariants after stage i are

$$|X_i| \geq \beta^i|X|, \qquad |Y_i| \geq \beta^i|Y|,$$

$$Y_i \subseteq N(x_1) \cap \cdots \cap N(x_i), \qquad X_i \subseteq X \setminus \big(N(y_1) \cup \cdots \cup N(y_i)\big),$$

where $N(x) = \{y \in Y : R(x, y)\}$ and $N(y) = \{x \in X : R(x, y)\}$.

Suppose that X_i, Y_i have been constructed for some $0 \leq i < k$, where $X_0 = X$ and $Y_0 = Y$. Since $i < k$ and $0 < \beta < 1$, we have

$$|X_i|, |Y_i| \geq \beta^i|X| \geq \beta^k|X| = \rho_0|X|.$$

Thus ρ_0-regularity gives

$$d_R(X_i, Y_i) \geq \alpha - \rho_0 > \mu - \beta^k = 4\beta - \beta^k \geq 3\beta > 2\beta.$$

Hence there exists $x_{i+1} \in X_i$ such that $|N(x_{i+1}) \cap Y_i| \geq 2\beta|Y_i|$. Put $Y_{i+1} := N(x_{i+1}) \cap Y_i$. Then

$$|Y_{i+1}| \geq 2\beta|Y_i| \geq 2\beta^{i+1}|Y| \geq \beta^{i+1}|Y|.$$

In particular, since $i+1 \leq k$, we also have $|Y_{i+1}| \geq \rho_0|Y|$.

Now set $Z := X_i \setminus \{x_{i+1}\}$. Since $|X_i| \geq \beta^i|X| \geq \beta^k m > 2$, we have $|Z| = |X_i| - 1 \geq \frac{1}{2}|X_i|$. As $i < k$ and $\beta < 1/2$, this implies

$$|Z| \geq \frac{1}{2}\beta^i|X| \geq \beta^k|X| = \rho_0|X|.$$

Together with $|Y_{i+1}| \geq \rho_0|Y|$, this allows us to apply ρ_0-regularity to the pair (Z, Y_{i+1}). We get

$$d_R(Z, Y_{i+1}) \leq \alpha + \rho_0 < 1 - \mu + \beta^k = 1 - 4\beta + \beta^k \leq 1 - 3\beta < 1 - 2\beta.$$

Therefore some $y_{i+1} \in Y_{i+1}$ satisfies $|Z \setminus N(y_{i+1})| \geq 2\beta|Z|$. Define $X_{i+1} := Z \setminus N(y_{i+1})$. Since $|Z| \geq |X_i|/2$, we have

$$|X_{i+1}| \geq 2\beta|Z| \geq \beta|X_i| \geq \beta^{i+1}|X|.$$

This completes the induction.

At the end, for $i \leq j$ we have $y_j \in Y_{j-1} \subseteq N(x_i)$ if $i < j$, and $y_i \in Y_i \subseteq N(x_i)$ if $i = j$. Hence $R(x_i, y_j)$ holds. If $i > j$, then $x_i \in X_{i-1} \subseteq X_j \subseteq X \setminus N(y_j)$, so $R(x_i, y_j)$ fails. Thus the sequence $(x_i, y_j)_{1 \leq i,j \leq k}$ is a half-graph of length k, contradicting k-stability. □

Proposition 3.5 (Stable uniform cosets are good)**.** *Fix $k \geq 2$, $\mu \in (0, 1/2)$, and a prime p. There exists $\varepsilon = \varepsilon(k, \mu, p) > 0$ such that the following holds.*

Let $A \subseteq \mathbb{F}_p^n$ be k-stable, let $H \leq \mathbb{F}_p^n$ be a subspace, and let $g \in \mathbb{F}_p^n$. If the coset $g + H$ is ε-uniform for A, then the coset $g + H$ is μ-good for A. Consequently, if the quotient partition $\mathbb{F}_p^n/H$ is ε-uniform for A, then $\mathbb{F}_p^n/H$ is μ-good for A.

Proof. Set

$$\rho := \left(\frac{\mu}{4}\right)^k \qquad \text{and} \qquad m := \left\lfloor 2\left(\frac{4}{\mu}\right)^k \right\rfloor + 1.$$

Then $m > 2(4/\mu)^k$, so Lemma 3.4 applies with this choice of ρ and m. Choose $\varepsilon < \min\{\rho^2, \sqrt{\mu(1-\mu)/m}\}$. Assume that the coset $g + H$ is ε-uniform for A. We show that its density $\alpha := \frac{|A \cap (g+H)|}{|H|}$ cannot lie strictly between μ and $1 - \mu$.

First suppose $|H| \geq m$. Since $g + H$ is ε-uniform and $\varepsilon < \rho^2$, Proposition 2.7, applied with $u = 0$ and $v = g$, implies that the Cayley pair $(H,\ g + H)$ is ρ-regular in the bipartite graph Γ_A. The relation $R_A(x, y) \iff x + y \in A$ is k-stable, so Lemma 3.4 applies to this pair. Its edge density is

$$d_{\Gamma_A}(H, g + H) = \frac{|A \cap (g + H)|}{|H|} = \alpha.$$

Therefore $\alpha \le \mu$ or $\alpha \ge 1-\mu$.

It remains to handle the case $|H| < m$. This is the only point where the proof cannot use Lemma 3.4, because that lemma requires the two parts of the regular pair to have size at least m. The small case is instead forced directly by Parseval.

If $|H| = 1$, then the coset $g+H$ consists of a single point, so α is either 0 or 1. Hence the coset is automatically μ-good. We may therefore assume that $|H| > 1$ and argue by contradiction. Suppose $\mu < \alpha < 1-\mu$. Define the balanced function on H by $\phi(x) := \mathbf{1}_A(g+x) - \alpha \qquad (x \in H)$. Its average on H is zero:

$$\mathbb{E}_{x\in H}\phi(x) = \frac{|A\cap(g+H)|}{|H|} - \alpha = 0.$$

Thus the trivial Fourier coefficient of ϕ is zero. Equivalently, if $\xi \in H^\perp$, then the character $x \mapsto \omega^{-\langle x,\xi\rangle}$ is constant equal to 1 on H, and hence $\widehat{\phi}_H(\xi) = \mathbb{E}_{x\in H}\phi(x) = 0$. On the other hand, since $\mathbf{1}_A(g+x)$ is a Bernoulli function of mean α on the coset $g+H$, we have

$$\mathbb{E}_{x\in H}|\phi(x)|^2 = \mathbb{E}_{x\in H}\big(\mathbf{1}_A(g+x)-\alpha\big)^2 = \alpha(1-\alpha).$$

Indeed, this is just the variance of a $\{0,1\}$-valued function with mean α. By Parseval on H,

$$\sum_{\xi\in\mathbb{F}_p^n/H^\perp} |\widehat{\phi}_H(\xi)|^2 = \alpha(1-\alpha),$$

where the sum is taken over one representative of each coset of $H^\perp$. Since the trivial coefficient vanishes, this becomes

$$\sum_{\substack{\xi\in\mathbb{F}_p^n/H^\perp\\ \xi\neq H^\perp}} |\widehat{\phi}_H(\xi)|^2 = \alpha(1-\alpha) > \mu(1-\mu).$$

Here $\xi \neq H^\perp$ means that ξ represents a nontrivial character on H.

There are exactly $|H|-1$ nontrivial characters on H, and in the present case $|H|-1 < m$. If every nontrivial Fourier coefficient satisfied

$$|\widehat{\phi}_H(\xi)| \le \sqrt{\frac{\mu(1-\mu)}{m}},$$

then the previous sum would be bounded by $(|H|-1)\frac{\mu(1-\mu)}{m} < \mu(1-\mu)$, contradicting Parseval. Therefore there must be a nontrivial frequency $\xi \notin H^\perp$ such that

$$|\widehat{\phi}_H(\xi)| > \sqrt{\frac{\mu(1-\mu)}{m}}.$$

For such a nontrivial frequency, subtracting the constant α does not change the Fourier coefficient, so

$$\widehat{\phi}_H(\xi) = \mathbb{E}_{x\in H}\mathbf{1}_A(g+x)\omega^{-\langle x,\xi\rangle}.$$

This contradicts the assumed ε-uniformity of the coset $g+H$, because $\varepsilon < \sqrt{\frac{\mu(1-\mu)}{m}}$. Hence the intermediate-density alternative is impossible also when $|H| < m$.

Thus $g+H$ is μ-good. The partition-level statement follows by applying the coset-level statement to every coset of H. $\square$

Remark 3.6. Fourier uniformity gives graph regularity of large Cayley pairs. Stability then turns this regular pair into an almost empty or almost complete pair. Since the density of the pair $(H, g+H)$ is exactly the density of A on the coset $g+H$, this graph homogeneity is exactly coset goodness.

Chapter 4

Green's Arithmetic Regularity Lemma

Theorem 4.1 (Green). *Fix a prime p and $0 < \varepsilon < 1$. There exists $M = M(p, \varepsilon)$ such that, for every n and every set $A \subseteq V := \mathbb{F}_p^n$, there is a subspace $H \leq V$ with $\operatorname{codim}_V(H) \leq M$ such that the quotient partition V/H is ε-regular for A in Green's arithmetic sense.*

Green's theorem is useful here for two reasons. First, it gives the background for the stable theorem. Second, it contains the Fourier idea that we still use later. If many cosets are not Fourier-uniform, then their Fourier witnesses can be used to refine the quotient. This refinement makes an averaged energy increase by a fixed amount.

4.1 Energy increment sketch

For a subspace $H \leq V$, write

$$\alpha_{g,H} := \mathbb{E}_{h \in H} \mathbf{1}_A(g + h) = \frac{|A \cap (g + H)|}{|H|}.$$

Define the energy of the quotient partition V/H by

$$\mathcal{E}(H) := \mathbb{E}_{g+H \in V/H} \alpha_{g,H}^2.$$

One always has $0 \leq \mathcal{E}(H) \leq 1$.

Suppose that V/H is not ε-regular in Green's arithmetic sense. Then more than an ε-fraction of the cosets $g + H$ are not ε-uniform. For each such bad coset, choose a non-trivial character $\xi_g \in H^*$ witnessing this failure, so that

$$\left| \mathbb{E}_{h \in H} \big(\mathbf{1}_A(g + h) - \alpha_{g,H} \big) \omega^{-\xi_g(h)} \right| > \varepsilon,$$

where $\omega = e^{2\pi i/p}$. Let Γ be the collection of all these chosen witnesses, and refine H by taking their common kernel:

$$H' := \bigcap_{\xi \in \Gamma} \ker \xi.$$

Then $H' \leq H$, so the quotient partition V/H' refines V/H. More explicitly, each old coset decomposes as

$$g + H = \bigsqcup_{t+H' \in H/H'} (g + t + H').$$

This is the precise sense in which all bad cosets are refined simultaneously: we do not choose a single favourable subcoset inside each bad coset; instead, we add all bad Fourier witnesses to the quotient structure at once.

The energy increment is measured by the variance of the new subcoset densities inside each old coset. Namely,

$$\mathcal{E}(H') - \mathcal{E}(H) = \mathbb{E}_{g+H \in V/H} \mathbb{E}_{t+H' \in H/H'} \big(\alpha_{g+t,H'} - \alpha_{g,H}\big)^2.$$

For a bad coset $g + H$, the witness ξ_g is trivial on H' and hence descends to a non-trivial character of the quotient H/H'. By Parseval's identity on H/H', the corresponding local variance satisfies

$$\mathbb{E}_{t+H' \in H/H'} \big(\alpha_{g+t,H'} - \alpha_{g,H}\big)^2 \geq \big|\mathbb{E}_{h \in H} \big(\mathbf{1}_A(g + h) - \alpha_{g,H}\big)\omega^{-\xi_g(h)}\big|^2 > \varepsilon^2.$$

Since bad cosets have total proportion greater than ε, we obtain the global increment

$$\mathcal{E}(H') \geq \mathcal{E}(H) + \varepsilon^3.$$

As the energy is bounded above by 1, this iteration must terminate after at most ε^{-3} steps. At termination, the current quotient partition is ε-regular in Green's arithmetic sense.

The cost appears in the codimension. In one refinement step, $\mathrm{codim}_H(H')$ can be as large as the number of bad cosets. This number is at most $|V/H| = p^{\mathrm{codim}_V(H)}$. Thus the codimension can grow roughly like

$$d_{i+1} \lesssim d_i + p^{d_i}.$$

After several steps, this gives tower-type bounds. This is the arithmetic version of the usual energy-increment proof of Szemerédi's regularity lemma.

This argument gives a Green-regular quotient partition, but it does not make every coset uniform. After refinement, new bad cosets may appear at the next scale, even inside old cosets that were uniform before. Green's lemma controls this only on average, using the energy. In the stable setting, the tree bound stops this bad branching from continuing forever. This is why the stable theorem gives a stronger almost-homogeneous partition.

Chapter 5

The Stable Arithmetic Regularity Lemma

5.1 The theorem

We now state the theorem in the form used in this report.

Theorem 5.1 (Terry–Wolf)**.** *Fix a prime p, an integer $k \geq 1$, and $\mu \in (0,1)$. There is a constant $C = C(k,p,\mu)$, polynomial in μ^{-1}, such that for every $n \geq 1$ and every k-stable set $A \subseteq \mathbb{F}_p^n$, there exists a subspace $H \leq \mathbb{F}_p^n$ with $\operatorname{codim}(H) \leq C$ and $|A \cap (g+H)|/|H| \leq \mu$ or $|A \cap (g+H)|/|H| \geq 1-\mu$ for every $g \in \mathbb{F}_p^n$. In other words, G/H is μ-good for A.*

5.2 A Fourier density-increment step

Before proving the dense affine piece statement, we record the Fourier fact we need. It says that if an affine subspace is not uniform, then it contains a codimension-one affine subspace where the density changes. This part does not use stability.

Proposition 5.2 (Fourier witness to density deviation)**.** *Let $A \subseteq V$, $H \leq V$, $g \in V$, and $\varepsilon > 0$. If $g+H$ is not ε-uniform for A, then there are a codimension-one subspace $H' \leq H$ and a coset $g+u+H' \subseteq g+H$ such that $\big||A \cap (g+u+H')|/|H'| - |A \cap (g+H)|/|H|\big| \geq \varepsilon$.*

Proof. Write $\alpha := |A \cap (g+H)|/|H|$. Since $g+H$ is not ε-uniform, there exists $\xi \notin H^\perp$ such that $|\widehat{F_{A,g,H}}_H(\xi)| > \varepsilon$. At a nontrivial frequency on H, subtracting a constant has no effect. Hence

$$|\widehat{\phi_{A,g,H}}_H(\xi)| > \varepsilon, \qquad \phi_{A,g,H}(h) = \mathbf{1}_A(g+h) - \alpha.$$

Define the linear map $L : H \to \mathbb{F}_p, \qquad L(h) := \langle h, \xi \rangle$. Because $\xi \notin H^\perp$, the map L is nonzero. Since the target $\mathbb{F}_p$ is one-dimensional over itself, L is surjective. Therefore

$H' := \ker L = \{h \in H : \langle h, \xi\rangle = 0\}$ is a codimension-one subspace of H, and for every $\lambda \in \mathbb{F}_p$ we may choose $x_\lambda \in H$ with $\langle x_\lambda, \xi\rangle = \lambda$. The fibres of L are exactly the cosets

$$x_\lambda + H' = \{h \in H : \langle h, \xi\rangle = \lambda\}, \qquad \lambda \in \mathbb{F}_p,$$

and hence

$$H = \bigsqcup_{\lambda \in \mathbb{F}_p} (x_\lambda + H').$$

For each $\lambda \in \mathbb{F}_p$, define the density on the corresponding subcoset by

$$\alpha_\lambda := \frac{|A \cap (g + x_\lambda + H')|}{|H'|} = \mathbb{E}_{h \in x_\lambda + H'} \mathbf{1}_A(g + h),$$

and define its deviation from the original density by

$$a_\lambda := \alpha_\lambda - \alpha = \mathbb{E}_{h \in x_\lambda + H'} \phi_{A,g,H}(h).$$

On the coset $x_\lambda + H'$, the value of $\langle h, \xi\rangle$ is constantly λ. Therefore

$$\begin{aligned}
\widehat{\phi_{A,g,H}}_H(\xi) &= \mathbb{E}_{h \in H} \phi_{A,g,H}(h) \omega^{-\langle h, \xi\rangle} \\
&= \frac{1}{p} \sum_{\lambda \in \mathbb{F}_p} \omega^{-\lambda} \mathbb{E}_{h \in x_\lambda + H'} \phi_{A,g,H}(h) \\
&= \frac{1}{p} \sum_{\lambda \in \mathbb{F}_p} a_\lambda \omega^{-\lambda}.
\end{aligned}$$

Thus

$$\varepsilon < \left| \widehat{\phi_{A,g,H}}_H(\xi) \right| = \left| \frac{1}{p} \sum_{\lambda \in \mathbb{F}_p} a_\lambda \omega^{-\lambda} \right| \le \frac{1}{p} \sum_{\lambda \in \mathbb{F}_p} |a_\lambda| \le \max_{\lambda \in \mathbb{F}_p} |a_\lambda|.$$

Hence some λ satisfies $|a_\lambda| > \varepsilon$. Taking $u = x_\lambda$ gives

$$\left| \frac{|A \cap (g + u + H')|}{|H'|} - \frac{|A \cap (g + H)|}{|H|} \right| = |a_\lambda| > \varepsilon,$$

which proves the claim. □

Corollary 5.3 (Non-uniformity gives a density increment)**.** *Let $A \subseteq V$, $H \le V$, $g \in V$, and $\varepsilon > 0$. If $g + H$ is not ε-uniform for A, then there are a codimension-one subspace $H' \le H$ and a subcoset $g + u + H' \subseteq g + H$ such that $|A \cap (g + u + H')|/|H'| \ge |A \cap (g + H)|/|H| + \varepsilon/2$.*

Proof. Use the notation from the proof of Proposition 5.2. We have real numbers

$$a_\lambda = \alpha_\lambda - \alpha, \qquad \lambda \in \mathbb{F}_p,$$

whose average is zero, since

$$\frac{1}{p}\sum_{\lambda\in\mathbb{F}_p}\alpha_\lambda = \alpha.$$

Moreover

$$\left|\frac{1}{p}\sum_{\lambda\in\mathbb{F}_p} a_\lambda \omega^{-\lambda}\right| > \varepsilon.$$

It follows from the triangle inequality that

$$\frac{1}{p}\sum_{\lambda\in\mathbb{F}_p}|a_\lambda| > \varepsilon.$$

Let $M := \max_{\lambda\in\mathbb{F}_p} a_\lambda$. Since the a_λ are real and have average zero, the total positive mass equals the total negative mass in absolute value. Therefore

$$\frac{1}{p}\sum_{\lambda\in\mathbb{F}_p}|a_\lambda| = 2\cdot\frac{1}{p}\sum_{a_\lambda>0} a_\lambda \le 2M.$$

Thus $M > \varepsilon/2$. Choose λ with $a_\lambda = M$ and put $u = x_\lambda$. Then

$$\frac{|A\cap(g+u+H')|}{|H'|} = \alpha_\lambda = \alpha + a_\lambda \ge \alpha + \frac{\varepsilon}{2},$$

which is the desired density increment. □

5.3 Dense affine pieces

Proposition 5.4 (Dense affine piece)**.** *Fix a prime p, an integer $k \ge 1$, and $\alpha \in (0,1]$. Then there are constants $\nu = \nu(k,p,\alpha) \in (0,\alpha/2)$ and $M = M(k,p,\alpha)$ such that the following holds. If V is a finite-dimensional vector space over $\mathbb{F}_p$ and $B \subseteq V$ is a k-stable set with $|B| \ge \alpha|V|$, then there are a subspace $H \le V$ and an element $x \in V$ such that* $\operatorname{codim}(H) \le M$ *and $|B\cap(x+H)| \ge (1-\nu)|H|$. Equivalently, the translated slice $X := (B-x)\cap H$ is k-stable and has density at least $1-\nu$ in H.*

Proof. Choose $\nu := \alpha/4$. By Proposition 3.5, there exists a number $\varepsilon_0 = \varepsilon_0(k,\nu,p) > 0$ such that every ε_0-uniform coset of a k-stable set is ν-good. Shrinking ε_0 if necessary, assume $\varepsilon_0 < \frac{\alpha}{2}$. We shall show that one may take $M := \lceil 2/\varepsilon_0\rceil$. The quantitative proof of Proposition 3.5 gives ε_0 polynomially small in α for fixed k and p. Hence M is polynomial in α^{-1}.

Start with the affine subspace $y_0 + H_0 := 0 + V$ and write $\alpha_i := \frac{|B\cap(y_i+H_i)|}{|H_i|}$ for the density of B on the current affine subspace. Initially $\alpha_0 \ge \alpha$.

If $y_i + H_i$ is ε_0-uniform for B, then Proposition 3.5 says that this coset is ν-good. Since $\alpha_i \ge \alpha > \nu$, the sparse alternative is impossible, and so $\alpha_i \ge 1-\nu$. Taking $H = H_i$ and $x = y_i$ gives the required affine piece.

Suppose instead that $y_i + H_i$ is not ε_0-uniform for B. By Corollary 5.3, there are a codimension-one subspace $H_{i+1} \le H_i$ and a subcoset $y_{i+1} + H_{i+1} \subseteq y_i + H_i$ such that

$$\alpha_{i+1} = \frac{|B \cap (y_{i+1} + H_{i+1})|}{|H_{i+1}|} \ge \alpha_i + \frac{\varepsilon_0}{2}.$$

Thus each non-uniform step increases the density by at least $\varepsilon_0/2$ while increasing the codimension by exactly one.

This process cannot continue for more than $\lceil 2/\varepsilon_0 \rceil$ steps, since the density is always at most 1. Therefore it stops at some $i \le M$ with an ε_0-uniform coset $y_i + H_i$. The previous paragraph then shows that this coset has density at least $1-\nu$. Hence $\mathrm{codim}_V(H_i) = i \le M$ and $|B \cap (y_i + H_i)| \ge (1-\nu)|H_i|$. Set $H := H_i$ and $x := y_i$.

Finally, the translated slice $(B - x) \cap H$ is stable because it is obtained from the translate of the stable set B and then restricted to the subgroup H. This proves the proposition. □

5.4 Proof of the stable arithmetic regularity lemma

We now combine the previous ingredients to prove Theorem 5.1. This is the main proof of the report.

5.4.1 The inductive data

Fix a k-stable set $A \subseteq G$ and a number $\mu \in (0, 1/2)$. The case $k = 1$ is trivial, because then $A = \varnothing$. So we assume $k \ge 2$.

In the recursive construction, we often intersect the current set X_η with a translate of A, or with the complement of such a translate. Because of this, the stability parameter of X_η may change. We keep track of this change explicitly.

By Lemma 3.3, there is a function $\mathfrak{b}(r, s)$ with the following property: if B is r-stable and C is s-stable, then both $B \cap C$ and $B \setminus C$ are $\mathfrak{b}(r, s)$-stable. Define a sequence of stability parameters by

$$K_0 := 2, \qquad K_{t+1} := \mathfrak{b}(k+1, K_t) \qquad (t \ge 0).$$

Here $K_0 = 2$ is used because the full ambient group G is a 2-stable set, and the factor $k + 1$ accounts for the possibility of taking complements of translates of A.

Let $d := t(k) = 2^{k+2} - 2$ be the explicit stable tree bound from Theorem 1.5. For the depth-d construction set $M_t := M(K_{t+1}, p, \mu/2) \qquad (0 \le t < d)$, where $M(\cdot)$ is the codimension loss in Proposition 5.4. Define

$$C_0 := 0, \qquad C_t := \sum_{s=0}^{t-1} M_s \quad (1 \le t \le d), \qquad C := C_d.$$

Thus C is the explicit codimension bound allowed in the contradiction argument. Also set $\nu := \frac{\mu}{8}$. This is the value produced by Proposition 5.4 when the density parameter is $\alpha = \mu/2$.

Suppose, towards a contradiction, that there is no subspace $H \le G$ with $\operatorname{codim}(H) \le C$ such that the quotient partition G/H is μ-good for A. We shall construct data indexed by the full binary tree up to height d:

$$H_\eta \le G, \qquad X_\eta \subseteq H_\eta \qquad (\eta \in 2^{\le d}),$$

together with translations

$$x_\eta \in G \qquad (\eta \in 2^{\le d},\ \eta \ne \varnothing)$$

and branching parameters $g_\eta \in G \qquad (\eta \in 2^{<d})$. For a binary string η of length t, write $\eta = (\eta_1, \dots, \eta_t)$ and use one-based bit indexing. For an integer $0 \le r \le |\eta|$, write $\eta \restriction r$ for the initial segment of η of length r. Thus, if $\sigma = \eta \restriction s$, then the next bit after the prefix σ is η_{s+1}.

It is convenient to write $N^1(g) := A - g, \qquad N^0(g) := G \setminus (A - g)$. Thus $x \in N^i(g)$ means exactly that $x + g \in A$ when $i = 1$, and $x + g \notin A$ when $i = 0$.

The construction will satisfy the following invariants. If $|\eta| = t \le d$, then

(i) $\operatorname{codim}(H_\eta) \le C_t$;
(ii) $X_\eta \subseteq H_\eta$;
(iii) X_η is K_t-stable;
(iv) $|X_\eta| \ge (1 - \nu)|H_\eta|$.

If $\eta = \sigma^\frown i$ with $|\sigma| < d$ and $i \in \{0, 1\}$, then in addition

$$X_\eta = \big(N^i(g_\sigma + x_\eta) \cap (X_\sigma - x_\eta)\big) \cap H_\eta. \tag{5.1}$$

Equivalently,

$$X_\eta + x_\eta \subseteq N^i(g_\sigma) \cap X_\sigma.$$

Finally, we maintain the cumulative membership invariant: whenever $|\eta| = t \le d$, $0 \le s < t$, and $\sigma = \eta \restriction s$, then for every $x \in X_\eta$,

$$x + g_\sigma + \sum_{r=s+1}^{t} x_{\eta \restriction r} \in A \quad \Longleftrightarrow \quad \eta_{s+1} = 1. \tag{5.2}$$

This last invariant is exactly what we need later to get the forbidden tree.

5.4.2 The initial node

Set $H_\varnothing := G, \qquad X_\varnothing := G$. Then $\operatorname{codim}(H_\varnothing) = 0 = C_0$, the set $X_\varnothing$ is 2-stable, and $|X_\varnothing| = |H_\varnothing|$. Since $C \ge 0$, the contradiction hypothesis implies that G/G is not μ-good for A. Hence there is a choice of $g_\varnothing \in G$ such that both

$$|N^0(g_\varnothing) \cap H_\varnothing| > \mu |H_\varnothing| \qquad \text{and} \qquad |N^1(g_\varnothing) \cap H_\varnothing| > \mu |H_\varnothing|.$$

For the one-coset partition G/G one may take $g_\varnothing = 0$, but it is useful to keep the notation uniform.

5.4.3 The branching step

Assume that all data have been constructed up to depth $t < d$, and fix a node $\eta \in 2^t$. Since $\operatorname{codim}(H_\eta) \le C_t \le C$, the contradiction hypothesis says that the quotient partition G/H_η is not μ-good for A. Therefore we may choose $g_\eta \in G$ such that

$$|N^i(g_\eta) \cap H_\eta| > \mu |H_\eta| \qquad (i = 0, 1). \tag{5.3}$$

For $i \in \{0, 1\}$ set

$$S_{\eta,i} := N^i(g_\eta) \cap X_\eta \subseteq H_\eta.$$

By the density of X_η in H_η and (5.3),

$$|S_{\eta,i}| \ge |X_\eta| + |N^i(g_\eta) \cap H_\eta| - |H_\eta| > (\mu - \nu)|H_\eta|.$$

Since $\nu = \mu/8$, each $S_{\eta,i}$ has density at least $\mu/2$ in H_η.

The set $N^1(g_\eta) = A - g_\eta$ is k-stable and $N^0(g_\eta) = G \setminus (A - g_\eta)$ is $(k+1)$-stable. Since X_η is K_t-stable, the definition of K_{t+1} implies that each $S_{\eta,i}$ is K_{t+1}-stable.

Apply Proposition 5.4 inside the ambient vector space H_η to the set $S_{\eta,i}$, with stability parameter K_{t+1} and density parameter $\mu/2$. We obtain a subspace $H_{\eta^\frown i} \le H_\eta$ and an element $x_{\eta^\frown i} \in H_\eta$ such that $\operatorname{codim}_{H_\eta}(H_{\eta^\frown i}) \le M_t$ and

$$|S_{\eta,i} \cap (x_{\eta^\frown i} + H_{\eta^\frown i})| \ge (1 - \nu)|H_{\eta^\frown i}|.$$

Define

$$X_{\eta^\frown i} := \big(S_{\eta,i} - x_{\eta^\frown i}\big) \cap H_{\eta^\frown i}.$$

Then $X_{\eta^\frown i} \subseteq H_{\eta^\frown i}$, its density in $H_{\eta^\frown i}$ is at least $1 - \nu$, and its stability parameter is at most K_{t+1}, because it is a translate of $S_{\eta,i}$ followed by intersection with the subgroup $H_{\eta^\frown i}$, which does not increase the stability parameter by Lemma 3.3. Moreover $\operatorname{codim}(H_{\eta^\frown i}) \le C_t + M_t = C_{t+1}$. The identity (5.1) follows from

$$(S_{\eta,i} - x_{\eta^\frown i}) = N^i(g_\eta + x_{\eta^\frown i}) \cap (X_\eta - x_{\eta^\frown i}).$$

It remains to verify the cumulative invariant for the new child. Let $\tau = \eta^\frown i$ and let $x \in X_\tau$. If $s = t$, then $\tau \restriction s = \eta$ and, by (5.1),

$$x + g_\eta + x_\tau \in A \iff i = 1,$$

which is precisely (5.2) for the last edge of the path.

If $s < t$, then $\tau \restriction s = \eta \restriction s$ and $x + x_\tau \in X_\eta$. Applying the cumulative invariant at the parent node η to $x + x_\tau$ gives

$$(x + x_\tau) + g_{\eta\restriction s} + \sum_{r=s+1}^{t} x_{\eta\restriction r} \in A \iff \eta_{s+1} = 1.$$

Since $\tau \restriction r = \eta \restriction r$ for $r \le t$ and $x_\tau = x_{\tau\restriction(t+1)}$, this is exactly (5.2) for τ. Thus the induction continues.

By induction, the data exist for all nodes in $2^{\le d}$.

5.4.4 Building the forbidden tree

For each leaf $\eta \in 2^d$, choose an arbitrary element $c_\eta \in X_\eta$ and define

$$a_\eta := c_\eta + \sum_{r=1}^{d} x_{\eta \restriction r}.$$

For each internal node $\sigma \in 2^{<d}$, with $|\sigma| = s$, define

$$b_\sigma := g_\sigma - \sum_{r=1}^{s} x_{\sigma \restriction r}.$$

Now let $\sigma \in 2^{<d}$ be an initial segment of a leaf $\eta \in 2^d$, and write $s := |\sigma|$. Then

$$a_\eta + b_\sigma = c_\eta + g_\sigma + \sum_{r=s+1}^{d} x_{\eta \restriction r}.$$

Applying (5.2) to the leaf η and the prefix $\sigma = \eta \restriction s$, with $x = c_\eta$, gives

$$a_\eta + b_\sigma \in A \quad \iff \quad \eta_{s+1} = 1.$$

Thus the elements $(a_\eta)_{\eta \in 2^d}$ and $(b_\sigma)_{\sigma \in 2^{<d}}$ form a full binary tree configuration of height d for the relation $R_A(x, y) \iff x + y \in A$. This contradicts Theorem 1.5, because R_A is k-stable.

5.4.5 Termination and conclusion

We have shown that the contradiction hypothesis forces the existence of a forbidden tree of height $d = t(k) = 2^{k+2} - 2$ for R_A. Therefore the hypothesis was false. Hence there exists a subspace $H \leq G$ with $\operatorname{codim}(H) \leq C$ such that the quotient partition G/H is μ-good for A.

The explicit bound produced above is

$$C = \sum_{t=0}^{d-1} M(K_{t+1}, p, \mu/2) \leq d \max_{0 \leq t < d} M(K_{t+1}, p, \mu/2).$$

For fixed k and p, the explicit number $d = 2^{k+2} - 2$ and the finitely many stability parameters $K_1, \ldots, K_d$ depend only on k. Since each dense-piece bound $M(K_{t+1}, p, \mu/2)$ is polynomial in μ^{-1} for fixed K_{t+1} and p, the final codimension bound is polynomial in μ^{-1}. This proves Theorem 5.1.

5.5 Comparison with Green's theorem

The stable theorem is a stronger version of Green's theorem under an extra assumption. Both proofs use the same local Fourier fact: a large nontrivial Fourier coefficient gives a direction where the density changes. The difference is how this local information is used globally.

In Green's theorem, $A \subseteq \mathbb{F}_p^n$ is arbitrary. If too many cosets of the current subspace are not uniform, we add all their Fourier witnesses to the quotient at once. The averaged energy

$$\mathcal{E}(H) = \mathbb{E}_{g+H \in V/H} \alpha_{g,H}^2$$

then increases by a fixed amount. Since this energy is at most 1, the process stops. The final statement is an average statement: all but an ε-fraction of cosets are Fourier-uniform. This method does not force every coset to be uniform. Also, the repeated common-kernel refinements lead to tower-type bounds.

In the stable theorem, we do not run an energy increment over all bad cosets. Instead, we assume that every bounded-codimension quotient still has a bad coset, and then we try to branch from it. A bad coset has noticeable density from both A and A^c. The dense-piece proposition turns these two sides into two large affine children. If this could be repeated many times, the choices would form a binary tree for R_A. Stability forbids such a tree. So the process stops, and no bad coset remains. This is why the stable theorem controls every coset. It also explains why the bound is polynomial for fixed k and p.

In short, Green's theorem gives a quotient that is uniform on most cosets for arbitrary sets. The stable theorem gives a quotient that is almost empty or almost full on every coset, with polynomial complexity, for sets with no long order pattern. Stability turns arithmetic regularity into a simple description by cosets.

Bibliography

[1] Ben Green, *A Szemerédi-type regularity lemma in abelian groups, with applications*, Geometric and Functional Analysis 15 (2005), no. 2, 340–376.

[2] Ben Green and Terence Tao, *An arithmetic regularity lemma, an associated counting lemma, and applications*, In *An irregular mind*, Bolyai Society Mathematical Studies 21, Springer, 2010, 261–334.

[3] Arnab Bhattacharyya, Eldar Fischer, Hamed Hatami, Pooya Hatami, Shachar Lovett, and Oded Lovett, *Testing low complexity affine-invariant properties*, ACM Transactions on Computation Theory 5 (2013), no. 2, Article 9.

[4] Jacob Fox, Shachar Lovász, and Lisa Sauermann, *A polynomial bound for the arithmetic k-cycle removal lemma in vector spaces*, Journal of the European Mathematical Society 24 (2022), no. 1, 293–317.

[5] Wilfrid Hodges, *Model Theory*, Cambridge University Press, 1993.

[6] Maryanthe Malliaris and Saharon Shelah, *Regularity lemmas for stable graphs*, Transactions of the American Mathematical Society 366 (2014), no. 3, 1551–1585.

[7] Maryanthe Malliaris and Anand Pillay, *The stable regularity lemma revisited*, Proceedings of the American Mathematical Society 143 (2015), no. 4, 1551–1558.

[8] Caroline Terry and Julia Wolf, *Stable arithmetic regularity in the finite-field model*, Bulletin of the London Mathematical Society 51 (2019), no. 1, 65–77.