Big Ramsey degrees and Galvin-Prikry theorems for binary free-amalgamation classes

Andy Zucker
Department of Pure Mathematics
University of Waterloo

October 26th, 2022
Logic Seminar
Caltech (Zoom)

Recall Ramsey's theorem:

Recall Ramsey's theorem:

Theorem (Ramsey 1930)

Let $n, r < \omega$. Then

$$\aleph_0 \to (\aleph_0)^n_r$$

meaning that for any coloring γ of $[\aleph_0]^n$ into r colors, there is an infinite $X \subseteq \aleph_0$ with $|\gamma[[X]^n]| = 1$.

Recall Ramsey's theorem:

Theorem (Ramsey 1930)

Let $n, r < \omega$. Then

$$\aleph_0 \to (\aleph_0)^n_r$$

meaning that for any coloring γ of $[\aleph_0]^n$ into r colors, there is an infinite $X \subseteq \aleph_0$ with $|\gamma[[X]^n]| = 1$.

How to generalize? Can try to color infinite subsets.

Recall Ramsey's theorem:

Theorem (Ramsey 1930)

Let $n, r < \omega$. Then

$$\aleph_0 \to (\aleph_0)^n_r$$

meaning that for any coloring γ of $[\aleph_0]^n$ into r colors, there is an infinite $X \subseteq \aleph_0$ with $|\gamma[[X]^n]| = 1$.

How to generalize? Can try to color infinite subsets.

Theorem (Galvin-Prikry 1973)

For any Borel coloring of $[\aleph_0]^{\aleph_0}$ into finitely many colors, there is an infinte $X \subseteq \aleph_0$ with $[X]^{\aleph_0}$ monochromatic.

Another generalization: can demand more structure in the "monochromatic" set X.

Another generalization: can demand more structure in the "monochromatic" set X.

For instance, suppose $r < \omega$ and $\gamma\colon [\mathbb{Q}]^2 \to r$ is a coloring. Can we find $X \subseteq \mathbb{Q}$ order isomorphic to $\mathbb{Q}$ with $|\gamma[[X]^2]| = 1$?

Another generalization: can demand more structure in the "monochromatic" set X.

For instance, suppose $r < \omega$ and $\gamma\colon [\mathbb{Q}]^2 \to r$ is a coloring. Can we find $X \subseteq \mathbb{Q}$ order isomorphic to $\mathbb{Q}$ with $|\gamma[[X]^2]| = 1$?

NO! Enumerate $\mathbb{Q} = \{q_n : n < \omega\}$. We define a coloring $\gamma\colon [\mathbb{Q}]^2 \to 2$, where given $m < n < \omega$, we set

$$\gamma(\{q_m, q_n\}) = 0 \Leftrightarrow q_m < q_n.$$

Another generalization: can demand more structure in the "monochromatic" set X.

For instance, suppose $r < \omega$ and $\gamma\colon [\mathbb{Q}]^2 \to r$ is a coloring. Can we find $X \subseteq \mathbb{Q}$ order isomorphic to $\mathbb{Q}$ with $|\gamma[[X]^2]| = 1$?

NO! Enumerate $\mathbb{Q} = \{q_n : n < \omega\}$. We define a coloring $\gamma\colon [\mathbb{Q}]^2 \to 2$, where given $m < n < \omega$, we set

$$\gamma(\{q_m, q_n\}) = 0 \Leftrightarrow q_m < q_n.$$

If $X \subseteq \mathbb{Q}$ is order isomorphic to $\mathbb{Q}$, we must have $|\gamma[[X]^2]| = 2$.

Remarkably, 2 colors is the worst possible:

Remarkably, 2 colors is the worst possible:

Theorem (Galvin 1968)

For any $r < \omega$, we have

$$\mathbb{Q} \to (\mathbb{Q})^2_{r,2}$$

meaning that for any coloring $\gamma \colon [\mathbb{Q}]^2 \to r$, there is $X \subseteq \mathbb{Q}$ order isomorphic to $\mathbb{Q}$ so that $|\gamma[[X]^2]| \leq 2$.

Remarkably, 2 colors is the worst possible:

Theorem (Galvin 1968)

For any $r < \omega$, we have

$$\mathbb{Q} \to (\mathbb{Q})^2_{r,2}$$

meaning that for any coloring $\gamma\colon [\mathbb{Q}]^2 \to r$, there is $X \subseteq \mathbb{Q}$ order isomorphic to $\mathbb{Q}$ so that $|\gamma[[X]^2]| \leq 2$.

Theorem (Laver (unpublished), D. Devlin 1979)

For every $n < \omega$, there is $T_n < \omega$ so that for every $r < \omega$, we have

$$\mathbb{Q} \to (\mathbb{Q})^n_{r,T_n}$$

Devlin gives precise characterization of least T_n that works.

A few remarks about Devlin's characterization:

A few remarks about Devlin's characterization:

Identify $\mathbb{Q}$ with $2^{<\omega}$ ordered lexicographically. Associate to each finite $F \subseteq \mathbb{Q}$ its envelope, the finite subtree generated by F. Different possible envelopes correspond to different "bad" colors. Some envelopes can be avoided. Devlin trees cannot.

A few remarks about Devlin's characterization:

Identify $\mathbb{Q}$ with $2^{<\omega}$ ordered lexicographically. Associate to each finite $F \subseteq \mathbb{Q}$ its envelope, the finite subtree generated by F. Different possible envelopes correspond to different "bad" colors. Some envelopes can be avoided. Devlin trees cannot.

Finite Devlin trees: rooted binary trees where each level either has exactly one splitting node or exactly one terminal node. We don't really care about passing numbers (or just demand they are all 0).

A few remarks about Devlin's characterization:

Identify $\mathbb{Q}$ with $2^{<\omega}$ ordered lexicographically. Associate to each finite $F \subseteq \mathbb{Q}$ its envelope, the finite subtree generated by F. Different possible envelopes correspond to different "bad" colors. Some envelopes can be avoided. Devlin trees cannot.

Finite Devlin trees: rooted binary trees where each level either has exactly one splitting node or exactly one terminal node. We don't really care about passing numbers (or just demand they are all 0).

The terminal nodes code the points of the finite linear order. Call these coding nodes

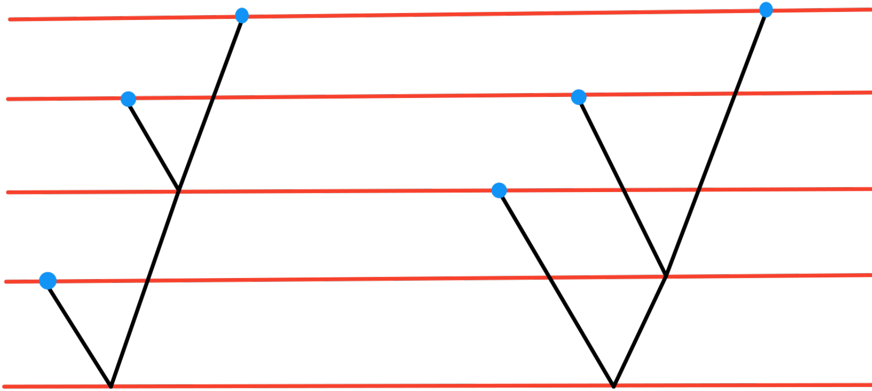

Some examples of Devlin trees.

Any subset of coding nodes induces a Devlin subtree of the original Devlin tree by closing under meets. Gives us a notion of embedding of one Devlin tree into another.

One can also consider infinite Devlin trees. We add the demand that terminal nodes are cofinal. Thus infinite Devlin trees code infinite linear orders.

One can also consider infinite Devlin trees. We add the demand that terminal nodes are cofinal. Thus infinite Devlin trees code infinite linear orders.

We pay special attention to those infinite Devlin trees that code the rational order. To prove the exact characterization of T_n, Devlin proves the following stronger result.

One can also consider infinite Devlin trees. We add the demand that terminal nodes are cofinal. Thus infinite Devlin trees code infinite linear orders.

We pay special attention to those infinite Devlin trees that code the rational order. To prove the exact characterization of T_n, Devlin proves the following stronger result.

Theorem (Devlin 1979)

Any two Devlin trees which code the rational order are bi-embeddable.

One can also consider infinite Devlin trees. We add the demand that terminal nodes are cofinal. Thus infinite Devlin trees code infinite linear orders.

We pay special attention to those infinite Devlin trees that code the rational order. To prove the exact characterization of T_n, Devlin proves the following stronger result.

Theorem (Devlin 1979)

Any two Devlin trees which code the rational order are bi-embeddable.

Question: Is this extra strength a fluke or a feature?

Definition

Let K be a countably infinite first-order structure, and let A be a finite structure with $\mathrm{Emb}(\mathsf{A}, \mathsf{K}) \neq \emptyset$. Let $\ell < r < \omega$. We write

$$\mathsf{K} \to (\mathsf{K})^{\mathsf{A}}_{r,\ell}$$

if for any coloring $\gamma\colon \mathrm{Emb}(\mathsf{A}, \mathsf{K}) \to r$, there is $\eta \in \mathrm{Emb}(\mathsf{K}, \mathsf{K})$ with $|\gamma[\eta \cdot \mathrm{Emb}(\mathsf{A}, \mathsf{K})]| = |\mathrm{Im}(\gamma \cdot \eta)| \leq \ell$.

Definition

Let K be a countably infinite first-order structure, and let A be a finite structure with $\mathrm{Emb}(\mathsf{A}, \mathsf{K}) \neq \emptyset$. Let $\ell < r < \omega$. We write

$$\mathsf{K} \to (\mathsf{K})^{\mathsf{A}}_{r,\ell}$$

if for any coloring $\gamma \colon \mathrm{Emb}(\mathsf{A}, \mathsf{K}) \to r$, there is $\eta \in \mathrm{Emb}(\mathsf{K}, \mathsf{K})$ with $|\gamma[\eta \cdot \mathrm{Emb}(\mathsf{A}, \mathsf{K})]| = |\mathrm{Im}(\gamma \cdot \eta)| \leq \ell$.

The Ramsey degree of A in K is the least $\ell < \omega$, if it exists, with $\mathsf{K} \to (\mathsf{K})^{\mathsf{A}}_{r,\ell}$ for every $r > \ell$.

Definition

Let K be a countably infinite first-order structure, and let A be a finite structure with $\mathrm{Emb}(\mathsf{A}, \mathsf{K}) \neq \emptyset$. Let $\ell < r < \omega$. We write

$$\mathsf{K} \to (\mathsf{K})^{\mathsf{A}}_{r,\ell}$$

if for any coloring $\gamma\colon \mathrm{Emb}(\mathsf{A}, \mathsf{K}) \to r$, there is $\eta \in \mathrm{Emb}(\mathsf{K}, \mathsf{K})$ with $|\gamma[\eta \cdot \mathrm{Emb}(\mathsf{A}, \mathsf{K})]| = |\mathrm{Im}(\gamma \cdot \eta)| \leq \ell$.

The Ramsey degree of A in K is the least $\ell < \omega$, if it exists, with $\mathsf{K} \to (\mathsf{K})^{\mathsf{A}}_{r,\ell}$ for every $r > \ell$.

If $\mathcal{K}$ is a Fraïssé class with limit K, we say that $\mathsf{A} \in \mathcal{K}$ has big Ramsey degree $\ell < \omega$ if A has Ramsey degree ℓ in K.

Definition

Let K be a countably infinite first-order structure, and let A be a finite structure with $\mathrm{Emb}(\mathsf{A},\mathsf{K}) \neq \emptyset$. Let $\ell < r < \omega$. We write

$$\mathsf{K} \to (\mathsf{K})^{\mathsf{A}}_{r,\ell}$$

if for any coloring $\gamma\colon \mathrm{Emb}(\mathsf{A},\mathsf{K}) \to r$, there is $\eta \in \mathrm{Emb}(\mathsf{K},\mathsf{K})$ with $|\gamma[\eta \cdot \mathrm{Emb}(\mathsf{A},\mathsf{K})]| = |\mathrm{Im}(\gamma \cdot \eta)| \leq \ell$.

The Ramsey degree of A in K is the least $\ell < \omega$, if it exists, with $\mathsf{K} \to (\mathsf{K})^{\mathsf{A}}_{r,\ell}$ for every $r > \ell$.

If $\mathcal{K}$ is a Fraïssé class with limit K, we say that $\mathsf{A} \in \mathcal{K}$ has big Ramsey degree $\ell < \omega$ if A has Ramsey degree ℓ in K.

We say that $\mathcal{K}$ has finite big Ramsey degrees if every $\mathsf{A} \in \mathcal{K}$ has some finite big Ramsey degree.

If $\mathsf{A} \in \mathcal{K}$ has big Ramsey degree ℓ, it is interesting to consider unavoidable ℓ-colorings of $\mathrm{Emb}(\mathsf{A}, \mathsf{K})$, i.e. a coloring witnessing that the big Ramsey degree is at least ℓ.

If $\mathrm{A} \in \mathcal{K}$ has big Ramsey degree ℓ, it is interesting to consider unavoidable ℓ-colorings of $\mathrm{Emb}(\mathrm{A}, \mathrm{K})$, i.e. a coloring witnessing that the big Ramsey degree is at least ℓ.

Easy: if $\mathrm{A} \leq \mathrm{B} \in \mathcal{K}$ have finite big Ramsey degrees ℓ_A and ℓ_B, respectively, then there are unavoidable colorings $\gamma_X \colon \mathrm{Emb}(\mathrm{X}, \mathrm{K}) \to \ell_X$ ($X \in \{A, B\}$), so that whenever $f \in \mathrm{Emb}(\mathrm{A}, \mathrm{B})$ and $x, y \in \mathrm{Emb}(\mathrm{B}, \mathrm{K})$, then $\gamma_B(x) = \gamma_B(y)$ implies $\gamma_A(x \circ f) = \gamma_A(y \circ f)$.

If $\mathrm{A} \in \mathcal{K}$ has big Ramsey degree ℓ, it is interesting to consider unavoidable ℓ-colorings of $\mathrm{Emb}(\mathrm{A}, \mathrm{K})$, i.e. a coloring witnessing that the big Ramsey degree is at least ℓ.

Easy: if $\mathrm{A} \leq \mathrm{B} \in \mathcal{K}$ have finite big Ramsey degrees ℓ_A and ℓ_B, respectively, then there are unavoidable colorings $\gamma_X \colon \mathrm{Emb}(\mathrm{X}, \mathrm{K}) \to \ell_X$ ($X \in \{A, B\}$), so that whenever $f \in \mathrm{Emb}(\mathrm{A}, \mathrm{B})$ and $x, y \in \mathrm{Emb}(\mathrm{B}, \mathrm{K})$, then $\gamma_B(x) = \gamma_B(y)$ implies $\gamma_A(x \circ f) = \gamma_A(y \circ f)$.

In other words, if we know $\gamma_B(x)$, then we automatically know $\gamma_A(x \circ f)$ for every $f \in \mathrm{Emb}(\mathrm{A}, \mathrm{B})$.

What about $A_0 \leq A_1 \leq \cdots$? If each A_i has finite big Ramsey degree ℓ_i, it is no longer clear that we can find unavoidable colorings $\gamma_i \colon \mathrm{Emb}(A_i, K) \to \ell_i$ with this coherence property.

What about $A_0 \leq A_1 \leq \cdots$? If each A_i has finite big Ramsey degree ℓ_i, it is no longer clear that we can find unavoidable colorings $\gamma_i\colon \mathrm{Emb}(A_i, K) \to \ell_i$ with this coherence property.

Definition (Z. 2019)

Assume $\mathcal{K}$ has finite big Ramsey degrees, and write $K = \mathrm{Flim}(\mathcal{K})$. We say that $\mathcal{K}$ admits a big Ramsey structure if there is an expansion K' of K so that for every $A \in \mathcal{K}$ with big Ramsey degree ℓ, the map sending $f \in \mathrm{Emb}(A, K)$ to the expansion on $f[A]$ is an unavoidable ℓ-coloring of $\mathrm{Emb}(A, K)$.

What about $\mathrm{A}_0 \leq \mathrm{A}_1 \leq \cdots$? If each A_i has finite big Ramsey degree ℓ_i, it is no longer clear that we can find unavoidable colorings $\gamma_i\colon \mathrm{Emb}(\mathrm{A}_i, \mathrm{K}) \to \ell_i$ with this coherence property.

Definition (Z. 2019)

Assume $\mathcal{K}$ has finite big Ramsey degrees, and write $\mathrm{K} = \mathrm{Flim}(\mathcal{K})$. We say that $\mathcal{K}$ admits a big Ramsey structure if there is an expansion K' of K so that for every $\mathrm{A} \in \mathcal{K}$ with big Ramsey degree ℓ, the map sending $f \in \mathrm{Emb}(\mathrm{A}, \mathrm{K})$ to the expansion on $f[\mathrm{A}]$ is an unavoidable ℓ-coloring of $\mathrm{Emb}(\mathrm{A}, \mathrm{K})$.

Example: Devlin trees coding the rational order. However, the statement that any two Devlin trees are bi-embeddable is even stronger.

Big Ramsey degrees are often discussed in terms of copies rather than embeddings. Embeddings are better for dynamical applications, while copies are more intuitive combinatorially.

Big Ramsey degrees are often discussed in terms of copies rather than embeddings. Embeddings are better for dynamical applications, while copies are more intuitive combinatorially.

Enumerated structures: A is enumerated if its underlying set is $|\mathsf{A}|$. Given enumerated structures A and B, write $\mathrm{OEmb}(\mathsf{A}, \mathsf{B}) := \{f \in \mathrm{Emb}(\mathsf{A}, \mathsf{B}) : f \text{ is monotone}\}$.

Big Ramsey degrees are often discussed in terms of copies rather than embeddings. Embeddings are better for dynamical applications, while copies are more intuitive combinatorially.

Enumerated structures: A is enumerated if its underlying set is $|A|$. Given enumerated structures A and B, write $\mathrm{OEmb}(A, B) := \{f \in \mathrm{Emb}(A, B) : f \text{ is monotone}\}$.

If $\mathcal{K}$ is SAP with limit K, then any two enumerations of K will be bi-embeddable. So we can define the ordered big Ramsey degree of an enumerated $A \in \mathcal{K}$. Ordinary BRD can be recovered from this.

Another example: the class $\mathcal{K}$ of finite graphs. Sauer (2006) shows that $\mathcal{K}$ has finite BRD, and Laflamme-Sauer-Vuksanovic (2007) give the precise characterization.

Another example: the class $\mathcal{K}$ of finite graphs. Sauer (2006) shows that $\mathcal{K}$ has finite BRD, and Laflamme-Sauer-Vuksanovic (2007) give the precise characterization.

As in Devlin's characterization, the (ordered) BRD of an (enumerated) finite graph is characterized by the number of certain tree-like objects which code the enumerated graph.

Another example: the class $\mathcal{K}$ of finite graphs. Sauer (2006) shows that $\mathcal{K}$ has finite BRD, and Laflamme-Sauer-Vuksanovic (2007) give the precise characterization.

As in Devlin's characterization, the (ordered) BRD of an (enumerated) finite graph is characterized by the number of certain tree-like objects which code the enumerated graph.

The main difference between Devlin trees and LSV-trees is that we now need to encode the graph relation via passing numbers.

Enumerated graph A

Two LSV-trees coding A

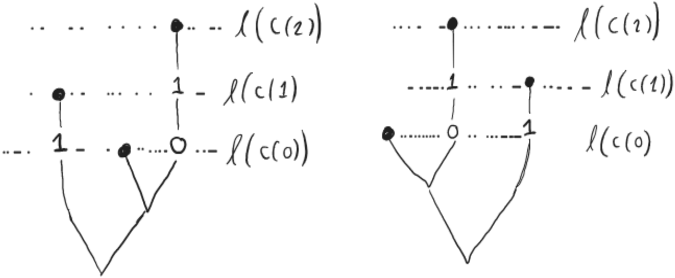

Similar to Devlin trees, any subset of coding nodes of an LSV-tree induces an LSV-subtree by closing under meets. Note that we don't record passing numbers unless the coding node on that level is in the subtree.

Similar to Devlin trees, any subset of coding nodes of an LSV-tree induces an LSV-subtree by closing under meets. Note that we don't record passing numbers unless the coding node on that level is in the subtree.

We can also consider those infinite LSV-trees which code the Rado graph. As before, lower bounds on big Ramsey degrees are obtained by proving something stronger.

Similar to Devlin trees, any subset of coding nodes of an LSV-tree induces an LSV-subtree by closing under meets. Note that we don't record passing numbers unless the coding node on that level is in the subtree.

We can also consider those infinite LSV-trees which code the Rado graph. As before, lower bounds on big Ramsey degrees are obtained by proving something stronger.

Theorem (Laflamme, Sauer, Vuksanovic 2007)

Any two LSV-trees coding the Rado graph are bi-embeddable.

Similar to Devlin trees, any subset of coding nodes of an LSV-tree induces an LSV-subtree by closing under meets. Note that we don't record passing numbers unless the coding node on that level is in the subtree.

We can also consider those infinite LSV-trees which code the Rado graph. As before, lower bounds on big Ramsey degrees are obtained by proving something stronger.

Theorem (Laflamme, Sauer, Vuksanovic 2007)

Any two LSV-trees coding the Rado graph are bi-embeddable.

In particular, any LSV-tree coding the Rado graph is a big Ramsey structure for the Rado graph.

In recent years, there has been rapid progress in understanding big Ramsey degrees, especially for binary free-amalgamation classes. The following classes all have finite big Ramsey degrees.

In recent years, there has been rapid progress in understanding big Ramsey degrees, especially for binary free-amalgamation classes. The following classes all have finite big Ramsey degrees.

- Dobrinen (2020): finite triangle-free graphs.

In recent years, there has been rapid progress in understanding big Ramsey degrees, especially for binary free-amalgamation classes. The following classes all have finite big Ramsey degrees.

- Dobrinen (2020): finite triangle-free graphs.
- Dobrinen (2022+): finite k-clique-free graphs.

In recent years, there has been rapid progress in understanding big Ramsey degrees, especially for binary free-amalgamation classes. The following classes all have finite big Ramsey degrees.

- Dobrinen (2020): finite triangle-free graphs.
- Dobrinen (2022+): finite k-clique-free graphs.
- Z. (2022): Any class in a finite binary relational language of the form $\mathrm{Forb}(\mathcal{F})$ for $\mathcal{F}$ a finite set of finite irreducible structures.

In recent years, there has been rapid progress in understanding big Ramsey degrees, especially for binary free-amalgamation classes. The following classes all have finite big Ramsey degrees.

- Dobrinen (2020): finite triangle-free graphs.
- Dobrinen (2022+): finite k-clique-free graphs.
- Z. (2022): Any class in a finite binary relational language of the form $\mathrm{Forb}(\mathcal{F})$ for $\mathcal{F}$ a finite set of finite irreducible structures.

Here irreducible means that every pair of points participates in a non-trivial relation.

In recent joint work with several authors, we precisely characterize the big Ramsey degrees for classes in a finite binary relational language of the form $\mathrm{Forb}(\mathcal{F})$ for $\mathcal{F}$ a finite set of finite irreducible structures.

In recent joint work with several authors, we precisely characterize the big Ramsey degrees for classes in a finite binary relational language of the form $\mathrm{Forb}(\mathcal{F})$ for $\mathcal{F}$ a finite set of finite irreducible structures.

Similar to the rational order and the Rado graph, the big Ramsey degree of an enumerated $\mathrm{A} \in \mathrm{Forb}(\mathcal{F})$ is the number of tree-like objects, which we call diagonal diaries, that code the structure.

In recent joint work with several authors, we precisely characterize the big Ramsey degrees for classes in a finite binary relational language of the form $\mathrm{Forb}(\mathcal{F})$ for $\mathcal{F}$ a finite set of finite irreducible structures.

Similar to the rational order and the Rado graph, the big Ramsey degree of an enumerated $\mathbf{A} \in \mathrm{Forb}(\mathcal{F})$ is the number of tree-like objects, which we call diagonal diaries, that code the structure.

Theorem (Balko-Chodounský-Dobrinen-Hubička-Konečný-Vena-Z.)

For any such $\mathrm{Forb}(\mathcal{F})$, any two diagonal diaries coding the Fraïssé limit are bi-embeddable. In particular, these classes all admit big Ramsey structures.

The precise definition of a diagonal diary is extremely technical. The most important new feature is that in addition to splitting and coding levels, there is a new type of critical level called an age-change level.

The precise definition of a diagonal diary is extremely technical. The most important new feature is that in addition to splitting and coding levels, there is a new type of critical level called an age-change level.

As an example, in an enumerated Fraïssé limit of the class of triangle-free graphs, given a non-edge $\{a, b\}$, it is interesting to record the least point adjacent to both a and b.

The precise definition of a diagonal diary is extremely technical. The most important new feature is that in addition to splitting and coding levels, there is a new type of critical level called an age-change level.

As an example, in an enumerated Fraïssé limit of the class of triangle-free graphs, given a non-edge $\{a, b\}$, it is interesting to record the least point adjacent to both a and b.

In diagonal diaries which code triangle-free graphs, we can record this information by putting a graph structure on each level of the tree.

Diagonal diaries coding triangle-free graphs are in one-one correspondence with runs of the following procedure which produces a sequence of finite, linearly ordered graphs:

Diagonal diaries coding triangle-free graphs are in one-one correspondence with runs of the following procedure which produces a sequence of finite, linearly ordered graphs:

- At step 0, G_0 is the empty graph.

Diagonal diaries coding triangle-free graphs are in one-one correspondence with runs of the following procedure which produces a sequence of finite, linearly ordered graphs:

- At step 0, G_0 is the empty graph.
- If G_k is empty, we can choose to stop.

Diagonal diaries coding triangle-free graphs are in one-one correspondence with runs of the following procedure which produces a sequence of finite, linearly ordered graphs:

- At step 0, G_0 is the empty graph.
- If G_k is empty, we can choose to stop.
- Given G_k, we can form G_{k+1} by doing exactly one of the following:

Diagonal diaries coding triangle-free graphs are in one-one correspondence with runs of the following procedure which produces a sequence of finite, linearly ordered graphs:

- At step 0, G_0 is the empty graph.
- If G_k is empty, we can choose to stop.
- Given G_k, we can form G_{k+1} by doing exactly one of the following:
  - Add a new minimum vertex adjacent to all previous vertices.

Diagonal diaries coding triangle-free graphs are in one-one correspondence with runs of the following procedure which produces a sequence of finite, linearly ordered graphs:

- At step 0, G_0 is the empty graph.
- If G_k is empty, we can choose to stop.
- Given G_k, we can form G_{k+1} by doing exactly one of the following:
  - Add a new minimum vertex adjacent to all previous vertices.
  - Duplicate a vertex; the two children are order-consecutive and non-adjacent.

Diagonal diaries coding triangle-free graphs are in one-one correspondence with runs of the following procedure which produces a sequence of finite, linearly ordered graphs:

- At step 0, G_0 is the empty graph.
- If G_k is empty, we can choose to stop.
- Given G_k, we can form G_{k+1} by doing exactly one of the following:
  - Add a new minimum vertex adjacent to all previous vertices.
  - Duplicate a vertex; the two children are order-consecutive and non-adjacent.
  - Delete an edge.

Diagonal diaries coding triangle-free graphs are in one-one correspondence with runs of the following procedure which produces a sequence of finite, linearly ordered graphs:

- At step 0, G_0 is the empty graph.
- If G_k is empty, we can choose to stop.
- Given G_k, we can form G_{k+1} by doing exactly one of the following:
  - Add a new minimum vertex adjacent to all previous vertices.
  - Duplicate a vertex; the two children are order-consecutive and non-adjacent.
  - Delete an edge.
  - Delete a vertex not belonging to any triangle.

Diagonal diaries coding triangle-free graphs are in one-one correspondence with runs of the following procedure which produces a sequence of finite, linearly ordered graphs:

- At step 0, G_0 is the empty graph.
- If G_k is empty, we can choose to stop.
- Given G_k, we can form G_{k+1} by doing exactly one of the following:
  - Add a new minimum vertex adjacent to all previous vertices.
  - Duplicate a vertex; the two children are order-consecutive and non-adjacent.
  - Delete an edge.
  - Delete a vertex not belonging to any triangle.
- When considering infinite runs of this procedure, demand that every vertex in every G_k has a descendant which gets deleted.

The 5 diagonal diaries coding an (enumerated) non-edge

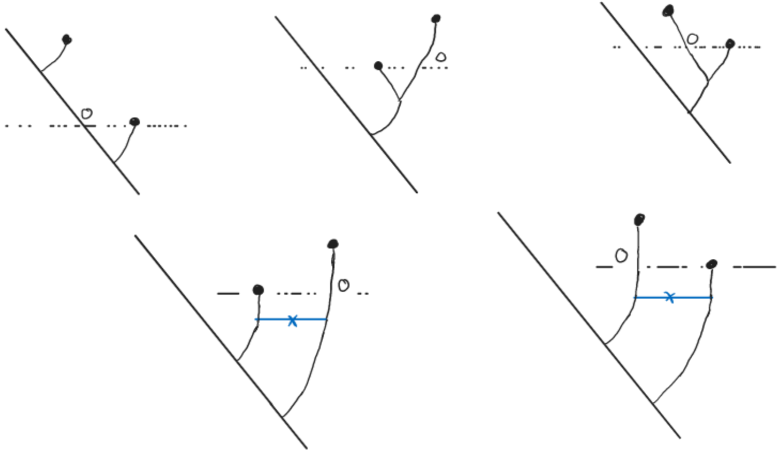

Theorem (Galvin-Prikry 1973)

For any finite Borel coloring of $[\omega]^\omega$, there is $X \in [\omega]^\omega$ with $[X]^\omega$ monochromatic.

Theorem (Galvin-Prikry 1973)

For any finite Borel coloring of $[\omega]^\omega$, there is $X \in [\omega]^\omega$ with $[X]^\omega$ monochromatic.

This can be rephrased using the semigroup S of increasing injections of ω into itself: for any finite Borel coloring of S, there is a monochromatic right ideal.

Theorem (Galvin-Prikry 1973)

For any finite Borel coloring of $[\omega]^\omega$, there is $X \in [\omega]^\omega$ with $[X]^\omega$ monochromatic.

This can be rephrased using the semigroup S of increasing injections of ω into itself: for any finite Borel coloring of S, there is a monochromatic right ideal.

Definition

We say that a Borel semigroup S is Galvin-Prikry if S satisfies the above statement.

For certain S, Todorčević has given sufficient structural criteria A1 through A4 which imply that S is Galvin-Prikry.

For certain S, Todorčević has given sufficient structural criteria A1 through A4 which imply that S is Galvin-Prikry.

We consider $S = \mathrm{Emb}(\mathrm{M}, \mathrm{M})$ for some countable relational structure M, which we write as a union of finite substructures in a distinguished way, $\mathrm{M} = \bigcup_n \mathrm{A}_n$.

For certain S, Todorčević has given sufficient structural criteria A1 through A4 which imply that S is Galvin-Prikry.

We consider $S = \mathrm{Emb}(\mathsf{M}, \mathsf{M})$ for some countable relational structure M, which we write as a union of finite substructures in a distinguished way, $\mathsf{M} = \bigcup_n \mathsf{A}_n$.

Write $\mathcal{A}_k = \{f \in \mathrm{Emb}(\mathsf{A}_k, \mathsf{M}) : \exists\, \phi \in \mathrm{Emb}(\mathsf{M}, \mathsf{M}) \text{ with } \phi|_{\mathsf{A}_k} = f\}$ and $\mathcal{A} = \bigcup_n \mathcal{A}_n$. Put $|f| = k$ iff $f \in \mathcal{A}_k$.

For certain S, Todorčević has given sufficient structural criteria A1 through A4 which imply that S is Galvin-Prikry.

We consider $S = \mathrm{Emb}(\mathsf{M}, \mathsf{M})$ for some countable relational structure M, which we write as a union of finite substructures in a distinguished way, $\mathsf{M} = \bigcup_n \mathsf{A}_n$.

Write $\mathcal{A}_k = \{f \in \mathrm{Emb}(\mathsf{A}_k, \mathsf{M}) : \exists\, \phi \in \mathrm{Emb}(\mathsf{M}, \mathsf{M}) \text{ with } \phi|_{\mathsf{A}_k} = f\}$ and $\mathcal{A} = \bigcup_n \mathcal{A}_n$. Put $|f| = k$ iff $f \in \mathcal{A}_k$.

This is enough to give A1.

A2 asks for a quasi-order $\leq_{fin}$ on $\mathcal{A}$ satisfying certain properties. For today, we take

$$f \leq_{fin} g \Leftrightarrow \mathrm{Im}(f) \subseteq \mathrm{Im}(g) \text{ and } g^{-1} \circ f \in \mathcal{A}_{|f|}.$$

A2 asks for a quasi-order $\leq_{fin}$ on $\mathcal{A}$ satisfying certain properties. For today, we take

$$f \leq_{fin} g \Leftrightarrow \mathrm{Im}(f) \subseteq \mathrm{Im}(g) \text{ and } g^{-1} \circ f \in \mathcal{A}_{|f|}.$$

This is also enough to yield one part of A3 called A3(1). We will almost never be in a situation where A3(2) holds.

The key pigeon-hole principle is A4:

The key pigeon-hole principle is A4:

Given $f \in \mathcal{A}_k$, suppose n is least with $\mathrm{Im}(f) \subseteq \mathsf{A}_n$. Then for every finite coloring γ of $\{g \in \mathcal{A}_{k+1} : g|_{\mathsf{A}_k} = f\}$, there is $\phi \in \mathrm{Emb}(\mathsf{M}, \mathsf{M})$ with $\phi|_{\mathsf{A}_n} = \mathrm{id}|_{\mathsf{A}_n}$ and with $\gamma \circ \phi$ constant.

The key pigeon-hole principle is A4:

Given $f \in \mathcal{A}_k$, suppose n is least with $\mathrm{Im}(f) \subseteq \mathsf{A}_n$. Then for every finite coloring γ of $\{g \in \mathcal{A}_{k+1} : g|_{\mathsf{A}_k} = f\}$, there is $\phi \in \mathrm{Emb}(\mathsf{M}, \mathsf{M})$ with $\phi|_{\mathsf{A}_n} = \mathrm{id}|_{\mathsf{A}_n}$ and with $\gamma \circ \phi$ constant.

Dobrinen isolates a mild strengthening of A4 which implies that $\mathrm{Emb}(\mathsf{M}, \mathsf{M})$ is Galvin-Prikry even without A3(2).

Theorem (Dobrinen-Z. (2022+))

Fix a class in a finite binary relational language of the form $\mathrm{Forb}(\mathcal{F})$ *for* $\mathcal{F}$ *a finite set of finite irreducible structures. Let* Δ *be any diagonal diary which codes the Fraïssé limit. Then* $\mathrm{Emb}(\Delta, \Delta)$ *is Galvin-Prikry.*

Theorem (Dobrinen-Z. (2022+))

Fix a class in a finite binary relational language of the form $\mathrm{Forb}(\mathcal{F})$ for $\mathcal{F}$ a finite set of finite irreducible structures. Let Δ be any diagonal diary which codes the Fraïssé limit. Then $\mathrm{Emb}(\Delta, \Delta)$ is Galvin-Prikry.

This is new even for the class of finite graphs.

The idea is to first show that certain strong diaries, along with a restricted notion of embedding between them, satisfies the strengthened version of A4.

The idea is to first show that certain strong diaries, along with a restricted notion of embedding between them, satisfies the strengthened version of A4.

This shows that the semigroup of strong self-embeddings of the strong diary is Galvin-Prikry.

The idea is to first show that certain strong diaries, along with a restricted notion of embedding between them, satisfies the strengthened version of A4.

This shows that the semigroup of strong self-embeddings of the strong diary is Galvin-Prikry.

Then use the bi-embeddability of diagonal diaries and properties of strong embeddings to transfer the result to any diagonal diary with ordinary embeddings.

Remarkably, this gives us examples of Galvin-Prikry semigroup where even A4 fails.

Remarkably, this gives us examples of Galvin-Prikry semigroup where even A4 fails.

We build Δ an LSV-tree coding the Rado graph with the property that in $\mathrm{Emb}(\Delta, \Delta)$, the identity is metrically isolated. We can arrange that if $\phi \in \mathrm{Emb}(\Delta, \Delta)$ satisfies $\phi|_{\mathsf{A}_1} = \mathrm{id}|_{\mathsf{A}_1}$, then $\phi = \mathrm{id}$.

Remarkably, this gives us examples of Galvin-Prikry semigroup where even A4 fails.

We build Δ an LSV-tree coding the Rado graph with the property that in $\mathrm{Emb}(\Delta, \Delta)$, the identity is metrically isolated. We can arrange that if $\phi \in \mathrm{Emb}(\Delta, \Delta)$ satisfies $\phi|_{\mathsf{A}_1} = \mathrm{id}|_{\mathsf{A}_1}$, then $\phi = \mathrm{id}$.

However, there will be $f \in \mathcal{A}_1$ with the property that $\{g \in \mathcal{A}_2 : g|_{\mathsf{A}_1} = f\} \geq 2$ (in fact infinite). For this f, A4 must fail.

An LSV-tree with isolated identity embedding

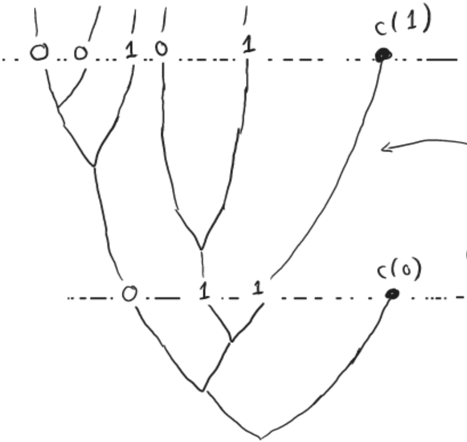

Key idea: $c(n+1)|_m$ does not split for any $m \in [\ell(c(n)), \ell(c(n+1))]$

Thanks!