Ramsey numbers of sparse digraphs

Yuval Wigderson (Stanford)
Joint with Jacob Fox and Xiaoyu He

March 4, 2021

Warmup: Hamiltonian paths in tournaments

Theorem (Rédei 1934)

Every tournament contains a Hamiltonian path.

Tournament = complete directed graph (every pair of vertices connected by a directed edge)

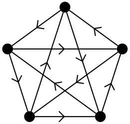

Questions and results about Hamiltonian paths in tournaments abound!

What structures must appear in every N-vertex tournament?

Definition

The *Ramsey number* $\vec{r}(H)$ of a digraph H is the minimum N such that every N-vertex tournament contains a copy of H.

Rédei's theorem $\iff \vec{r}(P_n) = n$, where $P_n =$ directed n-vertex path.

Directed and undirected Ramsey numbers

Definition

The *Ramsey number* $r(H)$ of a graph H is the minimum N such that every two-edge-coloring of K_N contains a monochromatic copy of H.

For a complete graph K_n,

$$2^{n/2} \le r(K_n) \le 2^{2n}.$$

The upper bound implies that $r(H)$ exists for all H.
If H has εn^2 edges, then

$$r(H) \ge 2^{\varepsilon n}.$$

Definition

The *Ramsey number* $\vec{r}(H)$ of a digraph H is the minimum N such that every N-vertex tournament contains a copy of H.

For a transitive tournament $\overrightarrow{T_n}$,

$$2^{n/2} \le \vec{r}(\overrightarrow{T_n}) \le 2^n.$$

The upper bound implies that $\vec{r}(H)$ exists for all acyclic H.
If H has εn^2 edges, then

$$\vec{r}(H) \ge 2^{\varepsilon n}.$$

So the Ramsey number is exponential if H is dense.

For the rest of the talk, we'll focus on sparse (di)graphs.

Ramsey numbers of sparse undirected graphs

If H is a tree or cycle, then $r(H) = O(n)$.

Burr–Erdős (1975): Does $r(H) = O(n)$ for all sparse H?

Theorem (Chvátal–Rödl–Szemerédi–Trotter 1983)

If H has n vertices and maximum degree Δ, then $r(H) = O_\Delta(n)$.

A more refined notion of sparsity is degeneracy, defined by

$$\max_{H' \subseteq H} (\text{minimum degree of } H').$$

If H is d-degenerate, then $r(H) \geq 2^{d/2}$. So graphs of unbounded degeneracy have "large" Ramsey numbers.

Conjecture (Burr–Erdős 1975), Theorem (Lee 2017)

If H is d-degenerate, then $r(H) = O_d(n)$.

Upshots: H has linear Ramsey number "if and only if" H is sparse. Qualitatively, n and d control $r(H)$.

Ramsey numbers of sparse digraphs

Conjecture (Sumner 1971)

If H is any orientation of an n-vertex tree, then $\vec{r}(H) \leq 2n - 2$.

Häggkvist–Thomason (1991): $\vec{r}(H) \leq 12n$.

Kühn–Mycroft–Osthus (2011): $\vec{r}(H) \leq 2n - 2$ for $n \geq n_0$.

Theorem (Thomason 1986)

If H is any acyclic orientation of C_n, then $\vec{r}(H) = n$ for $n \geq n_0$.

Bucić–Letzter–Sudakov: Is $\vec{r}(H)$ linear for all bounded-degree H?

Theorem (Yuster 2020, Girão 2020, DDFGHKLMSS 2020)

If H has bandwidth k, (i.e. there is an edge $v_i \to v_j$ only if $1 \leq j - i \leq k$) then $\vec{r}(H) = O_k(n)$.

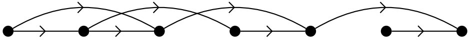

Main results

Bucić–Letzter–Sudakov: Is $\vec{r}(H)$ linear for all bounded-degree H?
No!

Theorem (Fox–He–W. 2021)

For all $C > 0$ and $n \geq n_0$, there is a bounded-degree ($\Delta \leq C^{3/2+o(1)}$) n-vertex acyclic digraph H with

$$\vec{r}(H) > n^C.$$

Theorem (Fox–He–W. 2021)

Let H be an n-vertex acyclic digraph with maximum degree Δ.

- $\vec{r}(H) \leq n^{O_\Delta(\log n)}$.
- *If H has height h, then* $\vec{r}(H) \leq n \cdot h^{O_\Delta(\log h)} = O_{\Delta,h}(n)$.
- *If H is chosen randomly, then* $\vec{r}(H) \leq n \cdot (\log n)^{O_\Delta(1)}$ *w.h.p.*

Height (aka depth) = length of longest directed path

What determines if $\vec{r}(H)$ is large?

Recall: In the undirected setting, number of vertices and degeneracy determine how large $r(H)$ is.

What additional parameters are relevant in the directed setting?

If H is an acyclic digraph, we can order its vertices as $v_1, \ldots, v_n$ such that all edges go to the right ($v_i \to v_j$ implies $i < j$).

Given such an ordering, the *length* of an edge $v_i \to v_j$ is $j - i$.

"Definition"

If H has "many" edges of length in $[2^t, 2^{t+1})$ for "most" $0 \le t \le \log n$, then H has *high multiscale complexity*.
If not, H has *low multiscale complexity*.

"Theorem"

Let H be a bounded-degree acyclic digraph. Then $\vec{r}(H)$ is large "if and only if" H has high multiscale complexity.

Multiscale complexity affects $\vec{r}(H)$

Multiscale complexity: Many edges in many dyadic length scales.

"Theorem"

Let H be a bounded-degree acyclic digraph. Then $\vec{r}(H)$ is large "if and only if" H has high multiscale complexity.

- If H has bandwidth k, then every edge in H has length $\leq k$.
- If H has height h, then "most" edges have length in $[n/h, n]$.
- Suppose H is chosen randomly by connecting $v_i \to v_j$ with probability $p = c/n$. Then

$$\mathbb{E}[\#(\text{edges of length} \leq \ell)] \leq p(n\ell) = c\ell.$$

  So a $o(1)$ fraction of H's edges have length $o(n)$.
- Our construction of a bounded-degree H with $\vec{r}(H) > n^C$ has many edges at every dyadic scale ("interval mesh").

Lower bound proof sketch

Theorem

There exists an n-vertex acyclic digraph H with maximum degree ≤ 1000 *and* $\vec{r}(H) > n^{\log_2(3)-\varepsilon}$.

We need (1) a construction of H, (2) a tournament T on $n^{\log_2(3)-\varepsilon}$ vertices, and (3) a proof that there is no embedding $H \hookrightarrow T$.

For (2): We let T be an iterated blowup of a cyclic triangle.

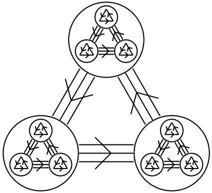

For (3): Construct H so that in any embedding $H \hookrightarrow T$, some subinterval of $[n]$ of length $\geq 0.49n$ is mapped into a single part.

Ensure that the induced subgraph on this subinterval has the same property, so we can iterate. At each step, $|T|$ drops by a factor of 3, but $|H|$ drops by a factor of 2.01.

Lower bound proof sketch: interval meshes

Want: In any embedding $H \hookrightarrow T$, some subinterval of $[n]$ of length $\geq 0.49n$ is mapped into a single part, and this is hereditary.

Definition

H is an *interval mesh* if

- H has a Hamiltonian path $1 \to 2 \to \cdots \to n$.
- For all $1 \leq a < b \leq c < d \leq n$ with $c - b \leq 100 \min(b - a, d - c)$, there is an edge between $[a, b]$ and $[c, d]$.

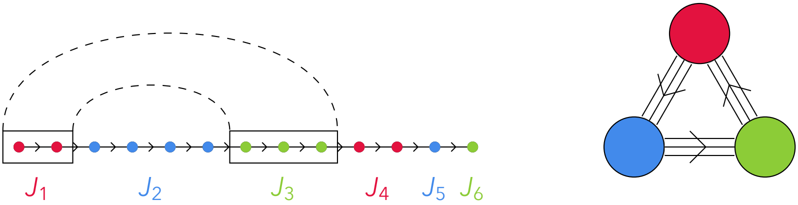

Thus, $|J_i| > 100 \min(|J_{i-1}|, |J_{i+1}|)$. So $|J_i| \geq 0.49n$ for some i.

Greedy algorithm yields an interval mesh with max degree ≤ 1000.

Upper bound proof sketch: greedy embedding

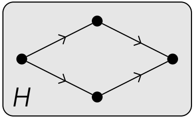

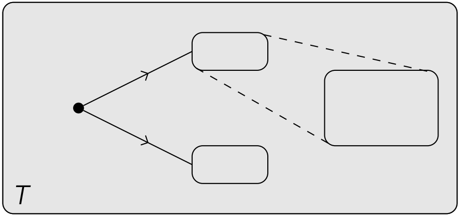

Lemma

If T is H-free, then T contains two large vertex sets with most edges between them oriented the same way.

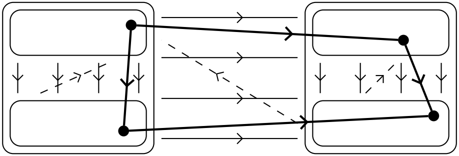

The multiscale complexity of *H* controls the number of iterations.

More colors and ordered Ramsey numbers

Summary: If H has n vertices and maximum degree Δ, then $\vec{r}(H) \le n^{O_\Delta(\log n)}$, but $\vec{r}(H) \ge n^C$ is possible.

With more colors, the upper bound is closer to the truth.

$$\vec{r_k}(H) = \min \left\{ N \;\middle|\; \begin{array}{c} \text{any } k\text{-edge-colored } N\text{-vertex tournament} \\ \text{contains a monochromatic copy of } H \end{array} \right\}.$$

Theorem (Fox–He–W. 2021)

If H has n vertices and maximum degree Δ, then

$$\vec{r_k}(H) \le n^{O_\Delta(\log^{O_k(1)} n)}.$$

For $k \ge 2$, there exists H of maximum degree 3 and

$$\vec{r_k}(H) \ge n^{\Omega(\log n / \log\log n)}.$$

Proof uses a connection to ordered Ramsey numbers.
Conlon–Fox–Lee–Sudakov and Balko–Cibulka–Král–Kynčl proved that random ordered matchings have super-polynomial ordered Ramsey numbers.

Conclusion and open questions

Let H have n vertices and maximum degree Δ.

- There is a gap between the n^C lower bound and $n^{O_\Delta(\log n)}$ upper bound on $\vec{r}(H)$.
  We conjecture that the upper bound is closer to the truth.
  Perhaps the same iterated blowup construction for T works?
- If H is random, we conjecture $\vec{r}(H) = O_\Delta(n)$ w.h.p., but can only prove $\vec{r}(H) \leq n(\log n)^{O_\Delta(1)}$.
  This boils down to improving one technical lemma.
- Some notion of multiscale complexity affects whether $\vec{r}(H)$ is small or large.
  - Can one formalize this?
  - Which other digraph parameters are relevant?
- Can one combine greedy embedding with existing techniques (e.g. median ordering)?

Thank you!