Ramsey numbers of trees

Jun Yan

University of Warwick

Joint work with Richard Montgomery and Matías Pavez-Signé

LSE PhD Combinatorics Seminar

24th January, 2025

Outline

Ramsey numbers of trees

Definition

The Ramsey number of a graph G, denoted as $R(G)$, is the smallest integer n such that any red/blue colouring of K_n contains a monochromatic copy of G.

Theorem

- [Gerencsér, Gyárfás, 1967]

$$R(P_k) = \begin{cases} \frac{3k}{2} & \text{if } k \text{ is even,} \\ \frac{3k+1}{2} & \text{if } k \text{ is odd.} \end{cases}$$

- [Harary, 1972]

$$R(K_{1,k}) = \begin{cases} 2k-1 & \text{if } k \text{ is even,} \\ 2k & \text{if } k \text{ is odd.} \end{cases}$$

Burr and Erdős' Conjecture

Conjecture (Burr, Erdős, 1976)

Let T be a tree on n vertices, then

$$R(T) \leq \begin{cases} 2n-3 & \text{if } n \text{ is odd,} \\ 2n-2 & \text{if } n \text{ is even.} \end{cases}$$

Theorem (Zhao, 2016)

Let T be a tree on n vertices with n large, then $R(T) \leq 2n-2$.

This confirms the conjecture for large even n, but the odd case is still open.

Burr's conjecture

Lemma

Let T be a tree with bipartition classes of sizes $t_1 \geq t_2 \geq 2$, then

$$R(T) \geq R_B(T) := \max\{2t_1, t_1 + 2t_2\} - 1.$$

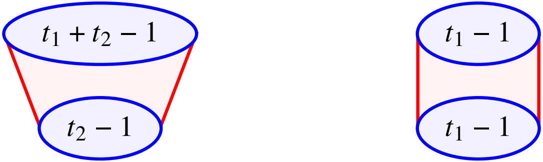

Conjecture (Burr, 1974)

Let T be a tree with bipartition classes of sizes $t_1 \geq t_2 \geq 2$, then

$$R(T) = R_B(T) = \max\{2t_1, t_1 + 2t_2\} - 1.$$

Counterexamples

Definition

The double star S_{m_1,m_2} is the tree obtained by joining the central vertices of the two stars K_{1,m_1} and K_{1,m_2} with an edge.

Note that if $m_1 \geq m_2$, then

$$R_B(S_{m_1,m_2}) = \max\{2m_1 + 2, m_1 + 2m_2 + 3\} - 1.$$

- [Grossman, Harary, Klawe, 1979]

$$R(S_{3m,m}) = 6m + 2 = R_B(S_{3m,m}) + 1.$$

- [Norin, Sun, Zhao, 2016]

$$R(S_{2m,m}) \geq (4.2 + o(1))m \geq (1.1 + o(1))R_B(S_{2m,m}).$$

Note that double stars have large maximum degrees.

Approximate version

Conjecture (Burr, 1974)

Let T be a tree with bipartition classes of sizes $t_1 \geq t_2 \geq 2$, then

$$R(T) = R_B(T) = \max\{2t_1, t_1 + 2t_2\} - 1.$$

Theorem (Haxell, Łuczak, Tingley, 2002)

For every $\mu > 0$, there exists $c > 0$ such that for every large n and every n-vertex tree T with $\Delta(T) \leq cn$, we have

$$R(T) \leq (1 + \mu) R_B(T).$$

Our Result

Theorem (Montgomery, Pavez-Signé, Y., 2025++)

There exists $c > 0$ such that for every n-vertex tree T with $\Delta(T) \leq cn$ and bipartition classes of sizes $t_1 \geq t_2$, we have

$$R(T) = R_B(T) = \max\{2t_1, t_1 + 2t_2\} - 1.$$

This confirms Burr's conjecture for all trees maximum degree at most cn.

The Haxell, Łuczak, Tingley proof sketch

Theorem (Haxell, Łuczak, Tingley, 2002)

For every $\mu > 0$, there exists $c \in (0,1)$ such that for every large n and every n-vertex tree T with $\Delta(T) \leq cn$ and bipartition classes of sizes $t_1 \geq t_2$, we have

$$R(T) \leq (1+2\mu)R_B(T) = (1+2\mu)\max\{2t_1, t_1+2t_2\}.$$

We may assume $t_1 \leq 2t_2$ by adding additional leaves.

Step 1: In any red/blue coloured graph G on $(1+2\mu)(t_1+2t_2)$ vertices, find a monochromatic "HŁT structure" in the reduced graph.

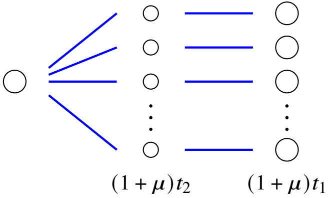

Step 2: Show that T can be embedded into the HŁT structure using regularity.

Our proof sketch

Definition

A red/blue coloured graph/reduced graph is μ-extremal if we can find one of the following approximate extremal constructions as a subgraph.

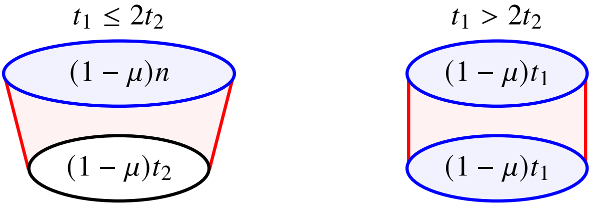

Stability part: starting with a scaled down HŁT structure in the reduced graph

- either we can find a structure to embed monochromatic T using regularity,
- or the reduced graph, and thus G must be μ-extremal.

Extremal part: embed a monochromatic T into a μ-extremal G.

Outline

Stability part

Starting with a scaled down HŁT structure in the reduced graph, we move through three stages.

In each stage,

- either we can find a structure to embed monochromatic T using regularity,
- or we find the structure that represents the beginning of the next stage.

At the end of Stage 3, we conclude that the reduced graph is extremal, and hence so is the original graph G.

Stage 1

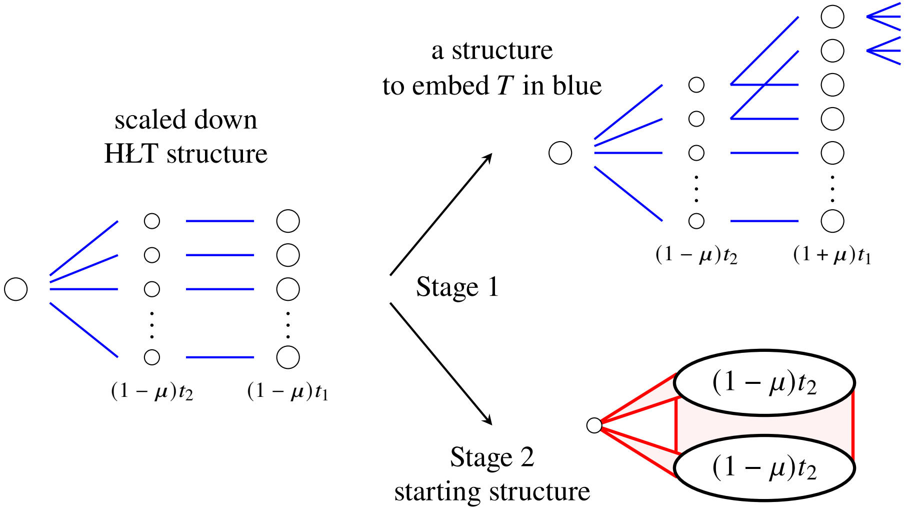
scaled down
HŁT structure
$(1-\mu)t_2$
$(1-\mu)t_1$
a structure
to embed T in blue
$(1-\mu)t_2$
$(1+\mu)t_1$
Stage 1
Stage 2
starting structure
$(1-\mu)t_2$
$(1-\mu)t_2$

Stage 2

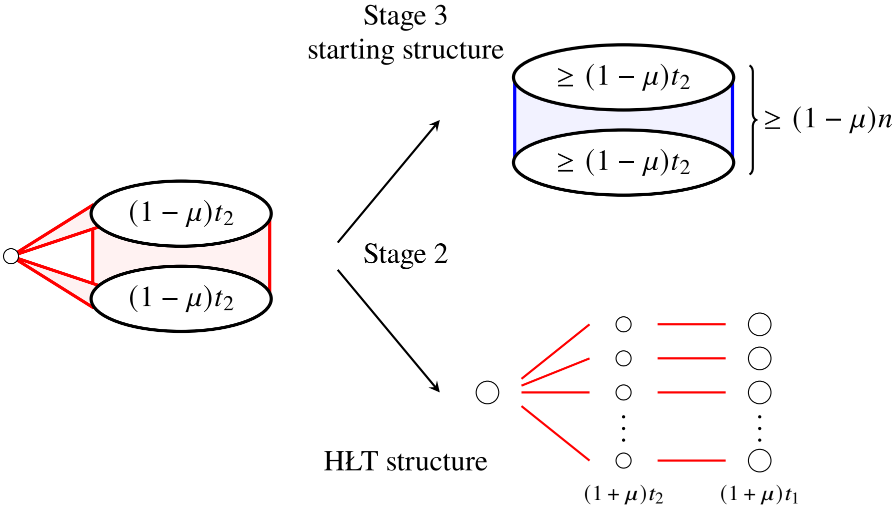
Stage 3
starting structure
$\geq (1-\mu)t_2$
$\geq (1-\mu)t_2$
$\geq (1-\mu)n$
$(1-\mu)t_2$
$(1-\mu)t_2$
Stage 2
HŁT structure
$(1+\mu)t_2$
$(1+\mu)t_1$

Stage 3

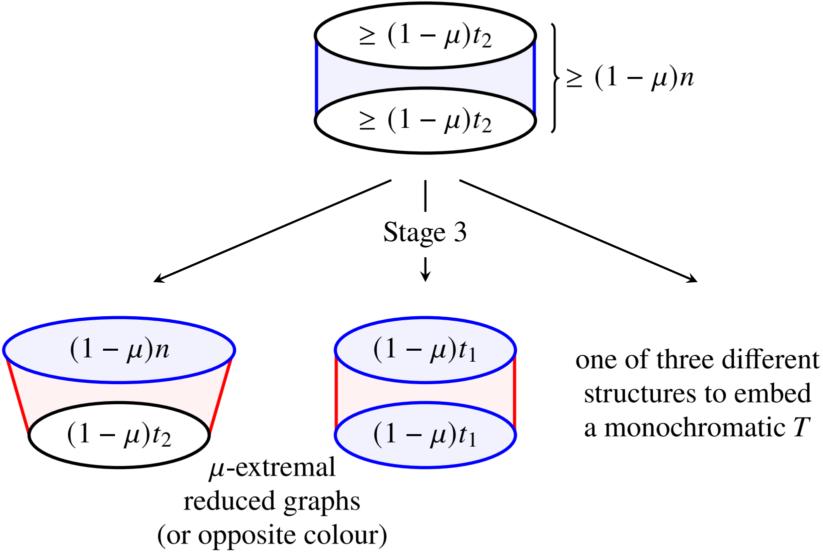
$\geq (1-\mu)t_2$
$\geq (1-\mu)t_2$
$\geq (1-\mu)n$
Stage 3
$(1-\mu)n$
$(1-\mu)t_2$
$(1-\mu)t_1$
$(1-\mu)t_1$
one of three different structures to embed a monochromatic T
μ-extremal reduced graphs (or opposite colour)

Outline

Extremal part

It follows from the Stability part that if a graph G on $R_B(T)$ vertices does not contain a monochromatic T, then G must be μ-extremal, with each vertex having at most μn neighbours in the wrong colour.

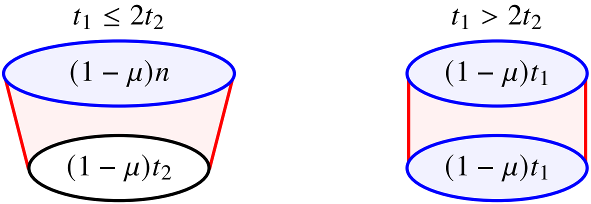

We show that we can find still a monochromatic T in any μ-extremal graph. Absorption type arguments relying on a random embedding technique.

Dichotomy between leaves and bare paths

Definition

A path P in a tree T is a bare path if every vertex in P has degree exactly 2.

Lemma (Krivelevich, 2010)

Let T be a tree on n vertices, then

- either T contains at least ℓ leaves,
- or T contains at least $\frac{n}{s+1} - 2\ell$ vertex-disjoint bare paths of length s.

We will assume from now on that T has many leaves.

Generalised Hall's matching condition

Lemma

Let G be a bipartite graph with bipartition $A \cup B$ and let $(f_a)_{a \in A}$ be a tuple of non-negative integers indexed by elements of A.
Suppose that $|N(S)| \geq \sum_{a \in S} f_a$ for all $S \subset A$. Then, there exists a disjoint collection of f_a neighbours in B for all $a \in A$.

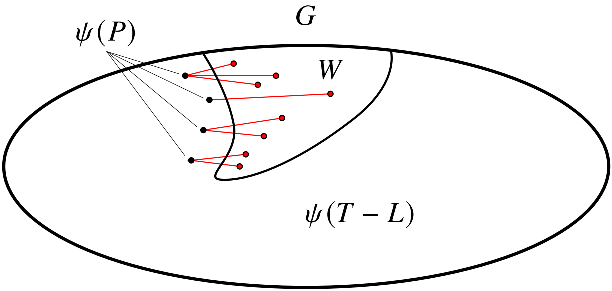

Random embedding method: Version 1

Lemma (Montgomery, Pavez-Signé, Y., 2025++, Version 1)

- $H = U_1 \cup U_2$ is a μn-almost-complete bipartite graph with $|U_1| = t_1$, $|U_2| = t_2 + 10\mu n$.
- T is a tree with n vertices and bipartition class sizes $t_1 \geq t_2$.
- T contains a set L of $\lambda n \gg \mu n$ leaves in the t_1 side, such that every parent of leaves in L in has at most $d \ll n/\log n$ children in L.

Then H contains a copy of T.

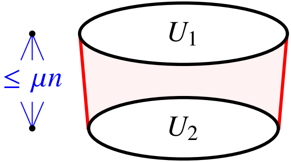

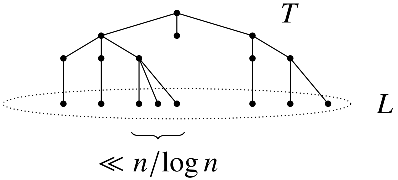

Random embedding method: Sketch

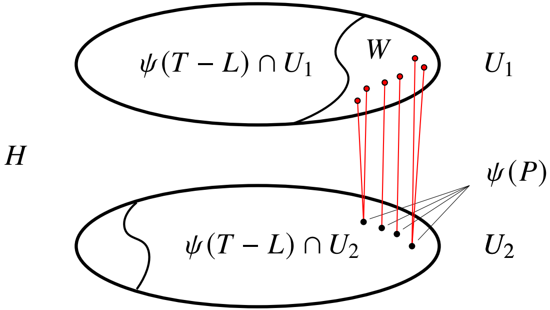

- Embed $T - L$ between U_1 and U_2 randomly.
- Verify generalised Hall's condition to embed L into W.
  - An unused vertex $w \in W$ is bad if all the parents in $\psi(P)$ that it is adjacent to together has few leaves in L.
  - Number of bad vertices is $n\exp(-\Theta(n/d)) \ll 1$ when $d \ll n/\log n$.

Random embedding method: Version 2

Lemma (Montgomery, Pavez-Signé, Y., 2025++, Version 2)

- $H = U_1 \cup U_2$ is a $\sqrt{n}$-almost-complete bipartite graph with $|U_1| \geq t_1$, $|U_2| = t_2 + 10\mu n$.
- T is a tree with n vertices and bipartition class sizes $t_1 > 2t_2$.
- T contains a set L of $\lambda n \gg \mu n$ leaves in the t_1 side, such that every parent of leaves in L in has at most $d \leq cn$ children in L.

Then H contains a copy of T.

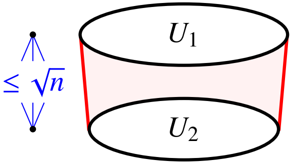

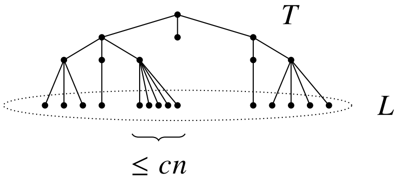

Extremal part: $t_1 > 2t_2$ cleaning

As $t_1 > 2t_2$, G is μ-extremal with $2t_1 - 1$ vertices.

Cleaning:
Assign every remaining vertex appropriately to either U_1 or U_2.

Maintain similar degree conditions in the resulting partition $V(G) = U_1' \cup U_2'$.

Extremal part: $t_1 > 2t_2$ brief sketch

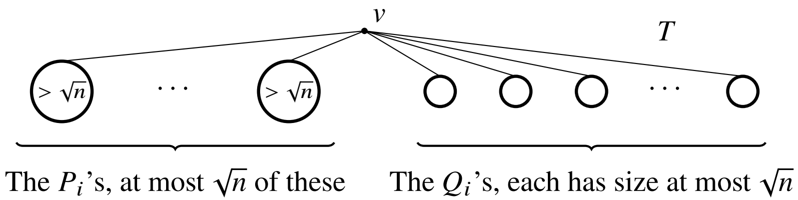

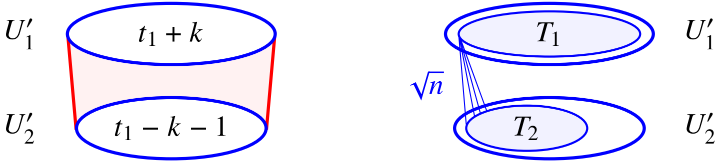

- $|\bigcup Q_i| > n/2$, can embed in red with Version 1 as $d \leq \sqrt{n} \ll n/\log n$.
- $|\bigcup P_i| > n/2$ and $\sqrt{n}$-almost-complete in red, embed with Version 2.
- $|\bigcup P_i| > n/2$ and there exists $u \in U_1$ with $\sqrt{n}$ blue neighbours in U_2, greedily embed in blue by sending appropriate amount over.

Random embedding method: Version 3

Lemma (Montgomery, Pavez-Signé, Y., 2025++, Version 3)

- G is a graph with $|G| \geq n$ and $\delta(G) \geq |G| - \mu n$.
- T is a tree with $|T| = n$.
- T contains a set L of $\lambda n \gg \mu n$ leaves, such that every parent of leaves in L in has at most $d \ll n/\log n$ children in L.

Then G contains a copy of T.

- Still works with some modifications if a set of at most $10\mu n$ vertices only have degree $\beta n \gg \lambda n$.
- Similarly, there is a variant that allows $d \leq cn$ at the cost of stronger degree conditions.
- Very flexible. Can use randomness to guarantee many other conditions satisfied by the embedding, e.g. certain vertices go into a certain set.

Extremal part: $t_1 \le 2t_2$ cleaning

As $t_1 \le 2t_2$, G contains $t_1 + 2t_2 - 1 = n + t_2 - 1$ vertices.

Cleaning:
For each remaining vertex, add it to U_1 if it has at least βn blue neighbours in U_1, and add it to U_2 otherwise.

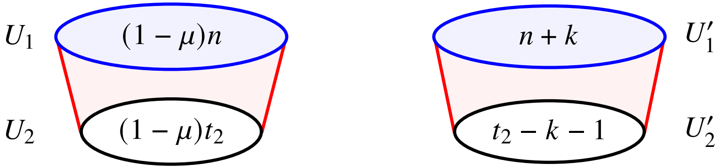

Extremal part: $t_1 \leq 2t_2$ and $d \ll \frac{n}{\log n}$

If $k \geq 0$, we can find a blue T in U_1' with Version 3.

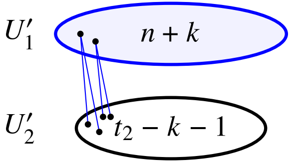

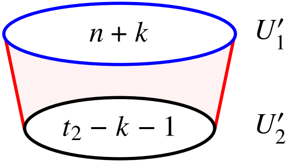

If $k \leq -1$, then

- either several vertices in U_1 have at least $-k$ blue neighbours in U_2': Embed T in blue in U_1' with Version 3, except that $-k$ leaves are embedded into U_2'.
- or most vertices in U_1 have at least t_2 red neighbours in U_2': Embed T in red between U_1 and U_2' greedily.

Extremal part: $t_1 \leq 2t_2$ and $d \gg \frac{n}{\log n}$

More complicated...
Tradeoff between conditions required to embed in blue and red.

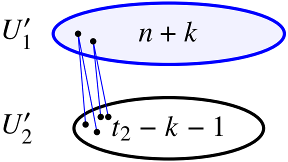

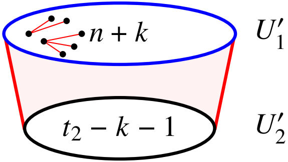

- If several vertices in U_1' have higher blue degrees:
  Use a variant of the random embedding method to embed T in blue.
- Otherwise, there are enough red edges in U_1':
  Embed T in red mostly greedily, but "flip" enough vertices to create space in U_2'.