1 Hadwiger's conjecture

The four-color theorem, conjectured in 1852 and only proved in 1976, states that every planar graph has a proper four-coloring. This theorem is one of the landmark results in graph theory, but every known proof requires a substantial amount of computer-assisted case analysis. This naturally raises the question (which we will shortly return to) of whether there is a simple, human-readable proof.

But an even more basic question about the four-color theorem is: why four colors? The answer, of course, is that four is best possible: K_4 is a planar graph with $\chi(K_4) = 4$, hence we certainly can't obtain a stronger bound than the four-color theorem. On the other hand, we might hope that the four-color theorem is true, just because K_5 is non-planar: it is the simplest example of a graph of chromatic number 5, and its non-planarity rules out the simplest possible reason why the four-color theorem might fail to be true.

This reasoning is extremely naive, but (a) it gives the right answer of four, and (b) many other instances of the same reasoning also give the right answer. A prominent example of this is the so-called Heawood conjecture, also known as the Ringel–Youngs theorem, which states that the maximum chromatic number of a graph embeddable on a genus-g (orientable) surface is

$$\left\lfloor \frac{7+\sqrt{1+48g}}{2} \right\rfloor.$$

Here, it turns out to be quite easy to prove the upper bound in all cases $g \geqslant 1$ (note that $g = 0$ is precisely the four-color theorem), and this was done by Heawood in 1890. The hard part turns out to be the lower bound, which splits into a rather complex casework depending on the residue of g modulo 12, and which was slowly completed over the course of the 1960s. But in all cases, the tightness of Heawood's upper bound is again exhibited by a complete graph: that is, the reason why $\lfloor \frac{7+\sqrt{1+48g}}{2} \rfloor$ is the right answer is "because" a clique of that order *can* be embedded on the genus-g surface, whereas a clique of order one larger *cannot*.

All of this leads to Hadwiger's conjecture.

Conjecture 1.1 (Hadwiger, 1943). *Let $\mathcal{F}$ be a minor-closed family of graphs. Then*

$$\max\{\chi(G) : G \in \mathcal{F}\} = \max\{k : K_k \in \mathcal{F}\}.$$

Here, we recall that a graph H is a *minor* of a graph G if H can be obtained from G by deleting vertices, deleting edges, and contracting edges, and a class $\mathcal{F}$ is *minor-closed* if $H \in \mathcal{F}$ whenever H is a minor of some $G \in \mathcal{F}$. In particular, as deleting or contracting the edges of a graph embedded on some surface do not change the embedding, we see that the class of graphs embeddable on any fixed surface is a minor-closed family, hence Hadwiger's conjecture is essentially a formalization of the naive intuition discussed above. Moreover, as there are many other natural examples of minor-closed classes (e.g. more or less any class of graphs defined by some sort of topological condition), it would immediately determine the maximum chromatic number of any such family.

Hadwiger's conjecture remains one of the central open problems in graph theory. Among other things, one could hope that it is the "right" strengthening of the four-color theorem;

perhaps, by virtue of being a more general and abstract statement, one could obtain a human-readable proof of Hadwiger's conjecture, and hence of the four-color theorem.

An equivalent formulation of Hadwiger's conjecture, which is a little bit more convenient because it deals with a single graph rather than a family, is the following.

Conjecture 1.2 (Hadwiger, 1943). *If $\chi(G) \geqslant t$, then G contains K_t as a minor.*

It is not hard to see that this is equivalent to the previous formulation. Another advantage of this formulation is that it suggests another natural way of viewing Hadwiger's conjecture. Namely, we know that every graph containing K_t as a *subgraph* has chromatic number at least t, and it is well-known that the converse is false: there are even triangle-free graphs of arbitrarily large chromatic number. However, Hadwiger's conjecture suggests a natural remedy to the falseness of this converse, namely, if we weaken the subgraph relation to the minor relation, then this converse becomes true.

As indicated above, Hadwiger's conjecture is wide open. The cases $t \leqslant 4$ of the conjecture are elementary to prove, and it is not hard to show that the case $t = 5$ is equivalent to the four-color theorem. In a major work, Robertson, Seymour, and Thomas proved that the $t = 6$ case is *also* equivalent to the four-color theorem (and in particular is true). However, nothing is known about the veracity of the conjecture for any $t \geqslant 7$.

Whenever you have a hard problem, it is natural to try both to weaken it and to strengthen it. In terms of weakening, the most well-studied question is to find a function $f(t)$ with the property that if $\chi(G) \geqslant f(t)$, then G contains K_t as a minor. Hadwiger's conjecture assserts that $f(t) = t$ works, and the goal is to make progress towards this by finding smaller and smaller functions with the same property. For a long time, the best known bound on this problem was due independently to Kostochka and Thomason from the 1980s, who proved that $f(t) = Ct\sqrt{\log t}$ works, for an appropriate constant C. In fact, they proved a much stronger statement, namely that every graph without a K_t minor contains a vertex of degree at most $Ct\sqrt{\log t}$ (which readily yields the statement about the chromatic number by induction). Moreover, for this stronger statement, this bound is actually best possible up to the constant C. This state of affairs remained unchanged until a breakthrough of Norin, Postle, and Song, who broke this "degeneracy barrier" and proved that $f(t) = t(\log t)^{1/4+o(1)}$ works. Their result led to a flurry of follow-up papers improving the bounds, culminating in the current record, due to Delcourt and Postle, that $f(t) = Ct \log\log t$ works. It remains a major open problem to even prove the *linear Hadwiger conjecture*, which asserts that $f(t) = Ct$ works, for some appropriate constant C.

Another natural way to weaken the problem is to study it for a certain restricted class of graphs. Of these, a very important and well-studied class consists of those graphs of independence number 2. Note that if G is an n-vertex graph with $\alpha(G) = 2$, then $\chi(G) \geqslant n/2$, and hence Hadwiger's conjecture predicts that G contains $K_{n/2}$ as a minor. Thus, many people have tried to prove this statement, or weakenings of it, as another potential approach to resolving the full conjecture. A short and elegant argument of Duchet and Meyniel proves that every such graph contains $K_{n/3}$ as a minor, and unfortunately this remains essentially the state of the art: the best result is due to Fox, who found a clique minor of order $\frac{n}{3} + \Omega(n^{4/5}(\log n)^{1/5})$, improving the Duchet–Meyniel bound by a sublinear factor. It remains a major open problem to prove the existence of a clique number of order $(\frac{1}{3} + \delta)n$, for any fixed $\delta > 0$.

2 The odd Hadwiger conjecture

Having discussed weakenings of Hadwiger's conjecture, let us now turn to strengthenings. Many natural strengthenings turn out to be false (e.g. if one replaces chromatic number with list chromatic number, or if one replaces minor by topological minor). However, one of the most prominent strengthenings remained open for over 30 years, and concerns the notion of *odd minors*. In order to motivate this definition, it is helpful to think about the statement of Hadwiger's conjecture for small values of t.

For $t = 3$, the conjecture asserts that if $\chi(G) \geqslant 3$, then G contains K_3 as a minor. It is not hard to see that G contains K_3 as a minor if and only if G contains a cycle, hence this statement says that if G is not bipartite, then G contains a cycle. This statement is true, but it's clearly not the “right” statement: the “right” statement is that if G is not bipartite, then G contains an *odd* cycle.

In the case $t = 4$, one can again strengthen Hadwiger's conjecture to obtain a true statement involving odd cycles. Namely, Catlin proved that if $\chi(G) \geqslant 4$, then G contains a subdivision of K_4 in which all the cycles between three branch vertices have odd length.

These results motivated Gerards and Seymour to define the notion of an odd minor. Roughly speaking, when defining an odd minor, we do not allow arbitrary edge contractions, but only contractions of sets of edges which preserve the parities of cycle lengths. In order to make this definition formal, note that G contains K_t as a minor if and only if we can find vertex-disjoint trees $T_1, \ldots, T_t \subseteq G$ such that for all $i \neq j$, there is some edge joining T_i to T_j. Indeed, given such a configuration, we can obtain a K_t minor by contracting all edges within each tree; conversely, given a K_t minor, every vertex must come from some connected subgraph of G, and one can obtain this structure by picking a spanning tree for each such subgraph.

Definition 2.1. We say that G has K_t as an *odd minor* if there exist vertex-disjoint trees $T_1, \ldots, T_t \subseteq G$, as well as functions $f_i : V(T_i) \to \{\text{black}, \text{white}\}$ with the following two properties.

- Each f_i is a proper coloring of T_i, that is, $f_i(u) \neq f_i(v)$ for every $uv \in E(T_i)$.
- For every $i \neq j$, there exists some $v_i \in V(T_i), v_j \in V(T_j)$ such that $v_iv_j \in E(G)$ and $f_i(v_i) = f_j(v_j)$.

Note that in this structure, for every three trees T_i, T_j, T_k, the unique cycle passing through them and the edges v_iv_j, v_jv_k, v_kv_i has odd length, because the colors of vertices alternate in each step within a tree, but stay the same on the three crossing edges. In particular, it is not hard to show that G contains K_3 as an odd minor if and only if it contains an odd cycle, and that it contains K_4 as an odd minor if and only if it contains a subdivision of K_4 in which all cycles between three branch vertices are odd. That is, the results discussed above are precisely the cases $t = 3, 4$ of the following conjecture.

Conjecture 2.2 (Gerards–Seymour, 1993)**.** *If $\chi(G) \geqslant t$, then G contains K_t as an odd minor.*

Other than the cases of $t = 3, 4$ discussed above, the only case for which the odd Hadwiger conjecture is known is the $t = 5$ case; this is a substantial strengthening of the four-color theorem, and a proof was announced in the early 2000s by Guenin. Moreover, proving asymptotic weakenings of this conjecture turns out to be quite a bit more difficult than for Hadwiger's conjecture. For example, while Kostochka and Thomason proved that every graph without a K_t minor has a vertex of bounded degree (in terms of t), no such statement is true for graphs avoiding odd minors. For example, any bipartite graph avoids K_3 as an odd minor, but certainly there exist bipartite graphs of arbitrarily large minimum degree. This means that many of the density-based techniques used to make progress on Hadwiger's conjecture are simply unavailable in the setting of odd minors. On the other hand, it also suggests that perhaps the notion of odd minor is more accurately picking up the key properties of chromatic number, and hence that it may be the "better" statement to try to prove.

Despite the difficulties, a series of works was able to prove various weakenings of the odd Hadwiger conjecture, just as in the case of the usual Hadwiger conjecture. For example, Kawarabyashi and Song proved an analogue of the Duchet–Meyniel result for odd minors, proving that every n-vertex graph of independence number two contains $K_{n/3}$ as an odd minor. Additionally, there have been a series of works finding smaller and smaller functions f such that $\chi(G) \geqslant f(t)$ implies that G contains K_t as an odd minor. Moreover, in many instances when the bound for Hadwiger's conjecture was improved, researchers were able to adapt some of those techniques and obtain a similar improvement in the setting of odd minors. Nevertheless, the two problems seemed to have some pretty significant differences, until the following remarkable result of Steiner showed that they are in fact very closely related.

Theorem 2.3 (Steiner, 2022)**.** *Suppose that f is any function such that $\chi(G) \geqslant f(t)$ implies that G contains K_t as a minor. Then any graph G satisfying $\chi(G) \geqslant 2 \cdot f(t)$ contains K_t as an odd minor.*

In particular, if Hadwiger's conjecture is true, then the odd Hadwiger conjecture is true "up to a factor of 2": every graph with $\chi(G) \geqslant 2t$ contains K_t as an odd minor.

And yet, despite all the similarities and partial evidence, the odd Hadwiger conjecture is false.

Theorem 2.4 (Kühn–Sauermann–Steiner–W., 2025+)**.** *There exists a graph G with $\chi(G) \geqslant (\frac{3}{2} - o(1))t$ and not containing K_t as an odd minor.*

That is, not only is the odd Hadwiger conjecture false, it is false by an asymptotic factor of $\frac{3}{2}$. As discussed above, Steiner's theorem states that if Hadwiger's conjecture is true, then this "failure factor" is at most 2. Thus, in some very vague sense, our result is "halfway" to disproving Hadwiger's conjecture. More precisely, the fact that the odd Hadwiger conjecture is wrong by a factor of $\frac{3}{2}$ casts at least some doubt on the veracity of Hadwiger's conjecture.

Our construction does not give just any graph G, but one of independence number 2. So we really prove the following stronger result.

Theorem 2.5. *There exists an n-vertex graph with independence number 2 whose largest odd clique minor is of order $(\frac{1}{3} + o(1))n$.*

Note that the odd Hadwiger conjecture would predict an odd clique minor of order $\frac{n}{2}$, which explains where the failure factor of $\frac{3}{2}$ comes from. Moreover, as discussed above, Theorem 2.5 is asymptotically best possible, as Kawarabyashi–Song proved that every such graph *does* have $K_{n/3}$ as an odd minor.

3 Heuristics: random graphs

Whenever confronted with a conjecture about graphs, a natural first thing to try is to verify whether one can disprove it by considering a random graph. Since it is conceptually easier (at least for me) to work with sparse random graphs, let us pass to the complement: our goal will be to construct a graph H, such that its complement $G = \overline{H}$ has independence number 2 and no large odd clique minor. The first condition is equivalent to demanding that H be triangle-free. Concretely, we will sample H as a binomial random graph $G(n, p)$, and hope to show that with high (or at least positive) probability, we have both that H is triangle-free and that $G = \overline{H}$ has does not have $K_{n/2}$ as an odd minor. If we can find some p for which both of these events happen with high probability, then we will have disproved the odd Hadwiger conjecture. Note that, for the moment, we are only focusing on finding G which avoids $K_{n/2}$ as an odd minor, which is a weaker statement than Theorem 2.5, but still strong enough if our goal is just to disprove odd Hadwiger.

We understand pretty well when random graphs are or are not triangle-free, so let's begin by understanding the odd minors. By definition, an odd $K_{n/2}$ minor in G consists of vertex-disjoint trees $T_1, \ldots, T_{n/2}$, as well as special proper colorings of them. Since we have $n/2$ vertex-disjoint trees in a universe of size n, the average size of one of the trees is at most 2, so let us for simplicity assume that each T_i has exactly two vertices. Then each T_i is a single edge, and there is only one way to properly 2-color it. Thus, we may assume that $V(T_i) = \{u_i, v_i\}$, and that u_i is colored black and v_i white. For this to form an odd $K_{n/2}$ minor in G, we must have that for all $i \neq j$, either $u_i u_j$ or $v_i v_j$ is an edge of G (or both). Moreover, each T_i must actually be a tree, i.e. we also need that $u_i v_i \in E(G)$ for all i.

This last condition sounds annoying, so let's just drop it. Thus, let us define an *odd connected pairing* in G to be a sequence of pairs $\{u_1, v_1\}, \ldots, \{u_{n/2}, v_{n/2}\}$ such that for all $i \neq j$, at least one of $u_i u_j, v_i v_j$ is an edge of G. Recalling that we are choosing $G = \overline{H}$ for $H \sim G(n, p)$, we see that for a fixed collection of pairs, the probability that it forms an odd connected pairing in G is precisely $(1 - p^2)^{\binom{n/2}{2}}$. Indeed, for at least one of $u_i u_j, v_i v_j$ to be an edge of G, we cannot have that both are edges of H, and this happens with probability

$1-p^2$. We then need this to happen for all $\binom{n/2}{2}$ choices of i, j, and these events are all independent.

Thus, for a fixed pairing $\{u_1, v_1\}, \ldots, \{u_{n/2}, v_{n/2}\}$, the probability that it forms an odd connected pairing in G is precisely $(1-p^2)^{\binom{n/2}{2}} \approx e^{-p^2n^2/8}$. On the other hand, we have many potential pairings to consider: the number of them is precisely the number of perfect matchings in K_n, which is known to be $(n/e + o(n))^{n/2} \approx e^{n \log n/2}$. In particular, by the union bound we know that

$$\Pr(G \text{ contains an odd connected pairing}) \lesssim e^{n \log n/2} \cdot e^{-p^2n^2/8}.$$

As a consequence, if we take $p > 2\sqrt{\log n/n}$, then this probability will tend to 0 as $n \to \infty$, which is precisely what we want to happen.

Now, all that remains is to verify that we can find some $p > 2\sqrt{\log n/n}$ for which $G(n,p)$ is triangle-free with high probability. Unfortunately, we start running into issues here. Indeed, it is well-known that $G(n,p)$ almost surely has triangles once $p \gg 1/n$, which is far lower than the constraint we just derived.

Still, not all hope is lost: we are OK if H has some triangles, so long as there are "not too many of them", so that we can delete these triangles without significantly affecting the structure of H. In particular, it is not hard to show that if $p < \sqrt{n}/1000$, then one can destroy all triangles in $G(n,p)$ by deleting a very small number of edges, and there are a large number of results whose proof uses this fact: concretely, they show that if one samples $G(n,p)$ at this density and then deletes these few edges, then the resulting triangle-free graph is still "highly random". However, this is still not good enough for us: we need $p > 2\sqrt{\log n/n}$, and it's well-known that once $p \gg 1/\sqrt{n}$, then a huge number of edges of $G(n,p)$ must be deleted to destroy all triangles.

It is instructive to pause to understand why $1/\sqrt{n}$ is the critical density here. Recall that in $G(n,p)$, with high probability all vertices have degree roughly pn. Therefore, the number of cherries (copies of $K_{1,2}$) in $G(n,p)$ is equal to $\sum_v \binom{\deg(v)}{2} \approx \sum_v \binom{pn}{2} \approx \frac{1}{2}p^2n^3$. So long as $p \ll 1/\sqrt{n}$, this quantity is much smaller than $\binom{n}{2}$. As a consequence, the number of pairs of vertices that are *closed by a cherry* (that is, are the endpoints of a copy of $K_{1,2}$) is negligible. But such pairs are precisely the ones we need to worry about: every such pair which is itself an edge of $G(n,p)$ contributes to a triangle, and hence such pairs are the ones we need to worry about deleting. As long as their number is negligible compared to $\binom{n}{2}$, the total number of vertex pairs, we have a hope of destroying the triangles in $G(n,p)$ without significantly affecting the structure. However, once $p \gg 1/\sqrt{n}$, then the number of cherries is actually much larger than $\binom{n}{2}$, and the number of pairs closed by cherries is no longer insignificant. This explains why we cannot hope to push $G(n,p)$ beyond $p \asymp 1/\sqrt{n}$ while being able to effectively destroy the triangles.

Nevertheless, all hope is *still* not lost. The past few decades have seen remarkable advancements in the study of various constrained random graph processes, of which the most notable is the *triangle-free process*, which is a random graph process formed by starting with an empty graph, and repeatedly picking adding a uniformly random edge among all edges whose addition would not create a triangle. Building on breakthrough work of Bohman, it was proved independently by Fiz Pontiveros–Griffiths–Morris and Bohman–Keevash that, when this process terminates, it produces a highly pseudorandom triangle-free graph of edge density $p \approx \sqrt{\log n/2n}$. In fact, the triangle-free process is so random-like that in many precise senses, it is indistinguishable from $G(n,p)$ at the same density, other than the fact that it contains no triangles.

However, this is still not good enough for us: we want to end up with $p > 2\sqrt{\log n/n}$, and the triangle-free process gets us to $p \approx \sqrt{\log n/2n}$, which is worse by a factor of $\sqrt{8}$. And even worse, there is no saving grace here: the quantity $\sqrt{\log n/2n}$ has a real meaning, and in a very precise sense, it is impossible to construct a denser triangle-free graph without giving up on some pseudorandomness.

This means we are stuck, but it also means that we shouldn't give up hope. Thinking about the odd Hadwiger conjecture, it is not obvious that we need extremely strong pseudorandomness: we would be OK with trading some pseudorandomness for density, so long as we do this in some controlled way that allows us to still analyze the presence of odd connected pairings. The fact that we got *so* close to disproving the odd Hadwiger conjecture by just using off-the-shelf random graphs suggests that any new idea on constructing denser pseudorandom triangle-free graphs could translate to a disproof of the odd Hadwiger conjecture.

And this is indeed what happened. The triangle-free process is instrumental to the study of the off-diagonal Ramsey number $r(3,k)$, and previously gave the best known lower bound for this quantity. However, just as in our problem, there is an apparent tradeoff there between density and pseudorandomness. A remarkable recent breakthrough of Hefty, Horn, King, and Pfender (building on an earlier breakthrough of Campos, Jenssen, Michelen, and Sahasrabudhe, who were the first to beat the triangle-free process) yielded an elegant construction, which is much simpler to analyze than the triangle-free process, and which can produce a pseudorandom triangle-free graph of *much* higher density. We used (essentially) the same construction (with a different analysis) to prove Theorem 2.5.

4 The construction

The first basic idea in the constructions of both Campos et al. and Hefty et al. is the following simple observation. Let us fix some integers $n > m$, and let H_m be some triangle-free graph on vertex set $[m]$ and with edge density p. Now, let $\pi : [n] \to [m]$ be a uniformly random function, and define a random graph H_n on vertex set $[n]$ by "pulling back" the graph H_m along the function π: more precisely, we set $uv \in E(H_n)$ if and only if $\{\pi(u), \pi(v)\} \in E(H_m)$. This operation has several appealing properties. First, it preserves triangle-freeness: for all distinct $u, v, w \in [n]$, at least one of $\{\pi(u), \pi(v)\}, \{\pi(u), \pi(w)\}, \{\pi(v), \pi(w)\}$ must be a non-edge of H_m, as H_m is triangle-free, and hence H_n is triangle-free as well. Second, this operation preserves edge density (at least in expectation): for any given pair $u, v \in [n]$, the probability that they are joined by an edge of H_n is (essentially) equal to p. But this is extremely good for us: if we choose $n \gg m$, then we could take H_m to have edge density roughly $1/\sqrt{m}$, which is much larger than $1/\sqrt{n}$, and hence this operation gives us access to n-vertex triangle-free graphs that are much denser than what we can obtain by only using $G(n, p)$ and its variants. Finally, and arguably most importantly, many natural pseudorandomness properties of H_m are preserved by this operation with high probability, e.g. if the edges of H_m are well-distributed among large vertex sets, then the same holds for H_n.

Another way of viewing this operation, at least when $n \gg m$, is simply taking a random blowup: all vertices of $[n]$ which map to the same vertex in $[m]$ are simply clones of that vertex, and hence we are doing nothing but blowing up each vertex $x \in [m]$ to $|\pi^{-1}(x)|$ copies of it. This also demonstrates that the randomness inherent in π is sort of irrelevant, as all that actually matters are the sizes of the fibers. Nevertheless, the perspective of a random function will be very useful in what follows.

The problem with this construction is that, while it preserves certain global pseudorandomness properties, its local structure is highly non-random: we have these blowup sets corresponding to fibers of π, and between any pair of them we either see no edges or a complete bipartite graph. In particular, it is not hard to check that odd connected pairings in H_m can be used to construct much larger odd connected pairings in H_n, and hence this operation will be completely useless if our goal is to disprove the odd Hadwiger conjecture.

The key insight of Hefty, Horn, King, and Pfender is to not give up, and instead to do this operation twice. More precisely, we will now start with *two* graphs H_R, H_B, which we think of as colored red and blue, both on the vertex set $[m]$. We now pick two independent, uniformly random functions $\pi_R, \pi_B : [n] \to [m]$, and construct H_n by overlaying the two graphs we get via the discussion above: we set $uv \in E(H_n)$ if $\{\pi_R(u), \pi_R(v)\} \in E(H_R)$ or $\{\pi_B(u), \pi_B(v)\} \in E(H_B)$ (or both). It will be convenient to record which of the two "causes" led to the edge uv being present in H_n, so we color the edges of H_n red or blue depending on whether $\{\pi_R(u), \pi_R(v)\} \in E(H_R)$ or $\{\pi_B(u), \pi_B(v)\} \in E(H_B)$; if both happen, we give the edge both colors.

Here, the fact that the functions π_R, π_B are independent is crucial. While the red and blue graphs are each simply blowups of an m-vertex graph, they interact in complicated ways, and each one "hides" the blowup structure of the other one. This means that we can now gain back some of the "local" pseudorandomness that we lost when only using a single random function.

Unfortunately, we've lost something as well, which is the triangle-freeness: even if H_R and H_B are each individually triangle-free, there is no reason why H_n should be. Instead, while this implies that there is no *monochromatic* triangle in H_n, there may be plenty of red-red-blue and blue-blue-red triangles in H_n. So the final step in the Hefty–Horn–King–Pfender construction is to destroy all these triangles by deleting, for each one, the edge of the minority color, thus forming an n-vertex triangle-free graph.

This is a natural thing to try, but is there any reason to expect it to be useful? A priori, as there may be very many non-monochromatic triangles in H_n, it is not obvious why this operation won't simply delete essentially all edges of H_n. Indeed, the number of red cherries is still roughly $\frac{1}{2}p^2n^3$, since every vertex of H_n has red degree roughly pn. But the whole idea here was to work at an edge density p that is much larger than $1/\sqrt{n}$, and hence the number of red cherries is much larger than $\binom{n}{2}$, and it seems that we are back with the same issue we encountered when trying to use $G(n,p)$.

However, the blowup structure of the red graph now comes to save us: while there are very many red cherries in H_n, they overlap strongly, and they actually close very few pairs of vertices in H_n. Indeed, for a given pair $u, v \in [n]$, the probability that they are closed by a red cherry is precisely the probability that $\pi_R(u), \pi_R(v)$ have a common neighbor in H_R. But the number of cherries in H_R is roughly $\frac{1}{2}p^2m^3$, which is much less than $\binom{m}{2}$ if we choose $p \ll 1/\sqrt{m}$; in this case the probability that u, v are closed by a red cherry is only at most $(\frac{1}{2}p^2m^3)/m^2 = o(1)$. Another way of saying this is that while there are many red cherries in H_n, they close certain pairs with high multiplicity: if u, v are closed by a red cherry, then they are actually closed by roughly n/m different red cherries, one for each vertex in the blowup set corresponding to the common neighbor of $\pi_R(u), \pi_R(v)$. Thus, the number of pairs closed by a red cherry is roughly $(\frac{1}{2}p^2n^3)/(n/m) \approx (p^2m)\binom{n}{2} = o(1)\binom{n}{2}$. Hence, we can be hopeful that deleting those blue edges that are closed by red cherries will not be a very significant change.

With all this in mind, all that remains is to choose the parameters, recalling that we need $p = o(1/\sqrt{m})$ and $p = \omega(1/\sqrt{n})$. There are actually a wide array of different choices that will all work with essentially the same analysis; in our paper we choose, essentially for convenience,

$$m = \frac{n}{(\log n)^8} \qquad \text{and} \qquad p = \frac{1}{\sqrt{m} \cdot (\log m)^2} \approx \frac{(\log n)^2}{\sqrt{n}}.$$

The rest of the construction proceeds essentially as indicated above:

1. Sample two independent random graphs $H_R^\circ, H_B^\circ \sim G(m,p)$.
2. Obtain H_R, H_B from H_R°, H_B°, respectively, by deleting every edge that lies in a triangle.
3. Pick two independent random functions $\pi_R, \pi_B : [n] \to [m]$.
4. Define the graph H_n on vertex set n by setting $uv \in E(H_n)$ if $\{\pi_R(u), \pi_R(v)\} \in E(H_R)$ or $\{\pi_B(u), \pi_B(v)\} \in E(H_B)$ (or both).
5. Define $H \subseteq H_n$ by deleting every blue edge that lies on a red-red-blue triangle, and every red edge that lies on a blue-blue-red triangle.

In this way, we construct an n-vertex graph H which is triangle-free with probability 1. Finally, we set $G = \overline{H}$, so that G has independence number two. The rest of the proof boils down to proving that

$$\Pr(G \text{ contains an odd connected pairing}) \overset{n\to\infty}{\longrightarrow} 0. \tag{1}$$

5 Proof sketch

As in our analysis of odd connected pairings in $G(n,p)$, the most natural way to try to prove (1) is via a union bound. Recall that there are at most roughly $e^{n \log n/2}$ different choices of potential pairings $\{u_1, v_1\}, \dots, \{u_{n/2}, v_{n/2}\}$, hence to prove (1) it suffices to prove that for each such fixed pairing, we have

$$\Pr(\{u_1, v_1\}, \dots, \{u_{n/2}, v_{n/2}\} \text{ form an odd connected pairing in } G) \leqslant e^{-10n \log n}. \quad (2)$$

For $\{u_1, v_1\}, \dots, \{u_{n/2}, v_{n/2}\}$ to form an odd connected pairing in G, we must have that for all $i \neq j$, neither $u_i u_j$ nor $v_i v_j$ is an edge of H. However, the intuition discussed above states that since $p = o(1/\sqrt{m})$, both steps (2) and (5) of the construction above (in which we delete triangles) do not seriously affect anything. Therefore, for what follows, we will pretend that in fact $H_R = H_R^\circ, H_B = H_B^\circ$, and $H = H_n$. Obviously, making such things rigorous is highly non-trivial (and comprises a large fraction of the proof), but to get intuition we can safely make these assumptions.

Therefore, in order to obtain (2), we need to prove

$$\Pr(\text{for all } i \neq j, \text{ we have } u_i u_j \notin E(G) \text{ or } v_i v_j \notin E(G)) \leqslant e^{-10n \log n}$$

or equivalently

$$\Pr(\text{there exists } i \neq j \text{ for which } u_i u_j \in E(H) \text{ and } v_i v_j \in E(H)) \geqslant 1 - e^{-10n \log n}. \quad (3)$$

In order to prove (3), we introduce the following auxiliary objects. First, let Γ_R be the graph defined as follows. Its vertex set is $\binom{[m]}{2}$, i.e. a vertex of Γ_R is an *edge* of K_m. Then, for every $i \neq j$, we add an edge to Γ_R between $\{\pi_R(u_i), \pi_R(u_j)\}$ and $\{\pi_R(v_i), \pi_R(v_j)\}$. Note that, under this definition, Γ_R is a multigraph with exactly $\binom{n/2}{2}$ edges, but we delete all loops and parallel edges to obtain a simple graph, which may have anywhere between 0 and $\binom{n/2}{2} \approx n^2/8$ edges. We define the graph Γ_B analogously. Note that both of these graphs depend only on the choice of pairing and on the random functions π_R, π_B; in particular, the structure of these graphs is independent from the randomness inherent in H_R, H_B. The key lemma we now prove is the following.

Lemma 5.1. *With high probability over the choice of π_R, π_B, the following holds for every pairing $\{u_1, v_1\}, \dots, \{u_{n/2}, v_{n/2}\}$: at least one of the graphs Γ_R, Γ_B has at least $n^2/100$ edges.*

Note that Γ_R has very few edges precisely when the pairing $\{u_1, v_1\}, \dots, \{u_{n/2}, v_{n/2}\}$ is "highly correlated" with the blowup structure of H_R. Thus, this lemma roughly states that it is impossible for this to happen simultaneously with respect to both H_R and H_B: the two blowup structures are "orthogonal", so at least one of the two graphs Γ_R, Γ_B contains a constant fraction of all the possible edges it might have. The proof of Lemma 5.1 is non-trivial but completely elementary, boiling down to a bit of careful combinatorics and a simple union bound.

Now, we may actually fix some outcome of π_R, π_B such that the statement of Lemma 5.1 holds, and in so doing we still have all the randomness in H_R, H_B available to us. Without loss of generality, let us assume that for our fixed pairing, we have $e(\Gamma_R) \geqslant n^2/100$. Since our goal is to prove (3), note that

$$\begin{aligned}\Pr(&\text{there exists } i \neq j \text{ for which } u_iu_j \in E(H) \text{ and } v_iv_j \in E(H))\\ &\geqslant \Pr(\text{there exists } i \neq j \text{ for which } u_iu_j \text{ and } v_iv_j \text{ are both \textit{red} edges of } H)\\ &= \Pr(\text{there exists } i \neq j \text{ for which } \{\pi_R(u_i), \pi_R(u_j)\}, \{\pi_R(v_i), \pi_R(v_j)\} \in E(H_R))\\ &= \Pr(\text{some edge of } \Gamma_R \text{ has both of its endpoints included in } H_R).\end{aligned}$$

In the final step, we recall that the vertices of Γ_R are *edges* of K_m, and hence we can view the sampling $H_R \sim G(m, p)$ as nothing more than keeping a p-random set of vertices of Γ_R.

Having done all of these reductions, we have reduced our problem to the following. We have a graph Γ_R with at least $n^2/100$ edges, we keep a p-random sample of its vertices, and wish to prove that with high probability, the random induced subgraph we obtain has at least one edge. Note that the expected number of edges in the sample is precisely $p^2e(\Gamma_R)$. Moreover, there are a wide array of tools in probabilistic combinatorics which are precisely designed to handle such settings; the most notable of these is Janson's inequality, which gives exponential bounds on the probability that the sample has far fewer edges than its expectation. Applying Janson's inequality, we find that

$$\Pr(\text{some edge of } \Gamma_R \text{ has both of its endpoints included in } H_R) \geqslant 1 - e^{-p^2e(\Gamma_R)/10}.$$

Finally, we recall that $e(\Gamma_R) \geqslant n^2/100$, and hence this final probability is at least $1 - e^{-p^2n^2/1000}$. Finally, recalling our choice of p, we see that this is asymptotically greater than $1 - e^{-10n \log n}$ (in fact better by a polylogarithmic factor in the exponent), which, in sum, proves (3).

In reality, we cannot apply Janson's inequality, because it cannot handle the fact that I cheated and assumed $H_R = H_R^\circ, H_B = H_B^\circ$, and $H = H_n$. Additionally, even if I weren't cheating here, the conditions of Janson's inequality are actually not met in this setup. We overcome both problems in essentially the same way: we reveal the randomness of H_R°, H_B° in stages, first sampling much denser random graphs and then further sparsifying them. By doing this appropriately, and relying on other probabilitic tools (we use a version of Talagrand's inequality to make the proof short, but it could all be done with just the Azuma and Chernoff inequalities), we can still prove the desired exponential tail bound.