1 Introduction to graph theory

Definition 1. A *graph* G is defined to be a pair of sets (V, E) satisfying the following properties. V, called the *vertex set*, is some finite set. E, the *edge set*, records which pairs of vertices are adjacent; formally, we require that $E \subseteq V \times V$ (which means that elements of E are ordered pairs of vertices) and that E is *symmetric*, meaning that if $(u, v) \in E$, then $(v, u) \in E$ as well. This means that every edge can be thought of as an unordered pair of vertices. Note that we allow (v, v) to be an edge of our graph for any $v \in V$; we think of this as a *loop* connecting v to itself. If a graph G has no loops (i.e. $(v, v) \notin E$ for all $v \in V$), then we say G is *simple*.

Often, intsead of writing $(u, v) \in E$, we will write $u \sim v$, which we read as "u is adjacent to v".

Example. Let's draw the graph $G = (V, E)$, where $V = \{a, b, c, d\}$ and

$$E = \{(a, b), (b, c), (c, d), (a, d), (b, d), (a, a)\}.$$

Note that for brevity, I only included one copy of each unordered pair of vertices in E, though technically E should include (b, a) since it includes (a, b).

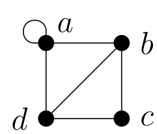

Example. Often, and especially in this class, we will need to deal with graphs that we can't draw directly, but which we can nevertheless understand because their vertices and edges are defined in some concrete way. One important class of graphs like this is the class of *Kneser graphs*. Given two positive integers $n \geq k$, the Kneser graph $\mathrm{KG}(n, k)$ is defined as follows. The vertex set V consists of all k-element subsets of $[n]$, where the notation $[n]$ means the set $\{1, 2, \ldots, n\}$ (we will frequently use this notation in this class). Thus, $|V| = \binom{n}{k}$. Additionally, given two vertices $S, T \in V$ (which are both k-element subsets of $[n]$), we will declare (S, T) to be an edge of $\mathrm{KG}(n, k)$ if and only if $S \cap T = \varnothing$.

Here's a picture of $\mathrm{KG}(5, 2)$, where the pair of integers next to each vertex is the subset of $[5]$ that this vertex represents.

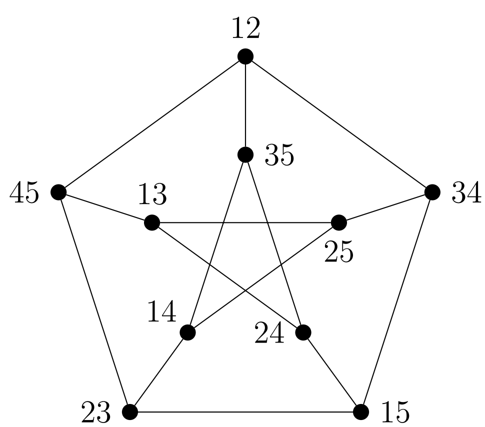

Note that if we double the parameters and look at $\mathrm{KG}(10,4)$, then we will already need to draw $\binom{10}{4} = 210$ vertices, which is quite a lot, and doubling again gets us to $\binom{20}{8} = 125970$, which is way too much to draw on a board. So to understand Kneser graphs in general, we have to reason about their structure, rather than drawing them and looking at what we get. Throughout this class, the graphs we'll be dealing with will generally be too big to draw, so we will need to be able to reason about graphs defined purely abstractly.

1.1 Colorings

One of the central concepts in graph theory, and one which we'll be dealing with throughout this class, is that of a proper coloring.

Definition 2. Given a graph $G = (V, E)$ and a positive integer c, a *proper c-coloring* is a function $\varphi : V \to [c]$ with the property that if $u \sim v$, then $\varphi(u) \neq \varphi(v)$. We think of the numbers $1, \ldots, c$ as colors and as φ assigning a color to each vertex, and then this condition says that adjacent vertices must get different colors.

If a graph G has a proper c-coloring, then we say that G is *c-colorable*. The *chromatic number* of G, $\chi(G)$, is defined to be the smallest c such that G is c-colorable.

Example. It's not too hard to convince yourself that $\chi(\mathrm{KG}(5,2)) = 3$. For instance, we can construct a proper 3-coloring by assigning color 1 to every vertex that has the number 1 in its label, color 2 to all remaining vertices with the number 2 in their label, and color 3 to the rest of the vertices. This proves that $\chi(\mathrm{KG}(5,2)) \leq 3$. On the other hand, even the outside pentagon can't be properly 2-colored, which proves that $\chi(\mathrm{KG}(5,2)) > 2$.

Note that if G is any graph with a loop, then G cannot be properly c-colored for any c. Indeed, suppose v has a loop in G. Then given any map $\varphi : V \to [c]$, the vertex v is assigned some color $\varphi(v)$. But since $v \sim v$ in G, the condition of a coloring requires $\varphi(v) \neq \varphi(v)$, which is impossible. Thus, graphs with loops do not have a well-defined chromatic number. Because of this, from now on, we will never refer to the chromatic number of a non-simple graph.

1.2 The tensor product of graphs

As you probably know very well if you've seen some abstract algebra (e.g. group theory or ring theory), it is often the case in math that once we define a class of objects, we also want to define the product of two such objects. It turns out that, unlike in abstract algebra, there are several interesting ways of defining the product of two graphs. In this class, one specific product, called the tensor product, will be the most important to us, though we will eventually encounter also the strong product and (on the homework) the Cartesian product.

Definition 3. Let $G = (V(G), E(G)), H = (V(H), E(H))$ be graphs. Their *tensor product* $G \times H$ is the graph defined as follows. Its vertex set $V(G \times H)$ is $V(G) \times V(H)$; i.e. a vertex of $G \times H$ is an ordered pair of vertices from G, H, respectively. Edges are defined by

$$((g_1, h_1), (g_2, h_2)) \in E(G \times H) \qquad \Longleftrightarrow \qquad (g_1, g_2) \in E(G) \text{ and } (h_1, h_2) \in E(H).$$

In other words, edges in $G \times H$ consist of coordinate-wise pairs of edges.

Example. Suppose $G =$ 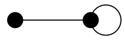 and $H =$ 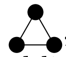, and let's draw $G \times H$. To do so, it's helpful to place G and H as a vertical and horizontal line of vertices, and then to place the vertices of $G \times H$ in the coordinate positions these define. That way, it's easy to see where we have a pair of vertices that are adjacent in each coordinate.

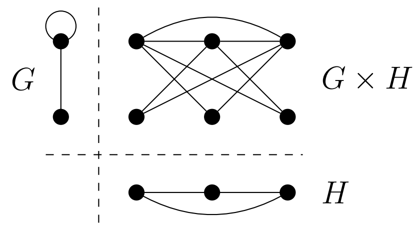

Proposition 1. *For any simple graphs G, H,*

$$\chi(G \times H) \leq \min\{\chi(G), \chi(H)\}.$$

Proof. First, suppose that $\varphi : V(G) \to [c]$ is a proper c-coloring of G. Then consider the map $\widehat{\varphi} : V(G \times H) \to [c]$ defined by

$$\widehat{\varphi}(g, h) = \varphi(g) \in [c].$$

We claim that $\widehat{\varphi}$ is a proper c-coloring of $G \times H$. To see this, suppose that $(g_1, h_1) \sim (g_2, h_2)$ in $G \times H$. Then by the definition of adjacency in the tensor product, this implies that $g_1 \sim g_2$ and $h_1 \sim h_2$. In particular, since φ was assumed to be a proper coloring of G, this implies that $\varphi(g_1) \neq \varphi(g_2)$, and therefore

$$\widehat{\varphi}(g_1, h_1) = \varphi(g_1) \neq \varphi(g_2) = \widehat{\varphi}(g_2, h_2),$$

so $\widehat{\varphi}$ is a proper c-coloring. Thus, if G is c-colorable, then so is $G \times H$, which implies that $\chi(G \times H) \leq \chi(G)$. The exact same argument shows that $\chi(G \times H) \leq \chi(H)$, and putting these together gives the desired result. □

This inequality is very tantalizing because, if you play around with small examples, you'll always find that $\chi(G \times H) = \min\{\chi(G), \chi(H)\}$. This led Hedetniemi to make the following conjecture in 1966.

Conjecture (Hedetniemi, 1966)**.** *For all simple graphs G, H,*

$$\chi(G \times H) = \min\{\chi(G), \chi(H)\}.$$

A huge amount of work went into proving this conjecture over the past fifty years, and there's been a lot of partial progress which proves it in various special cases. Nevertheless, in this class we will disprove it.

Theorem 1 (Shitov, 2019)**.** *There exist simple graphs G, H and a positive integer c such that $\chi(G) > c, \chi(H) > c$, but $\chi(G \times H) \leq c$.*

In fact, we will even prove the stronger result that for all sufficiently large c, there exist graphs G, H with this property.

1.3 Exponential graphs

The graph H that Shitov uses in his counterexample to Hedetniemi's conjecture is a very special type of graph called an exponential graph. To define it, we'll need one more coloring concept: that of co-proper colorings. The terminology is meant to be reminiscent of coprime (aka relatively prime) integers; these integers aren't necessarily prime, but they are, in some sense, prime when considered relative to one another.

Definition 4. Given a graph $G = (V, E)$ and a positive integer c, two maps $\varphi_1, \varphi_2 : V \to [c]$ are called *co-proper colorings* if, whenever $u \sim v$, we have that $\varphi_1(u) \neq \varphi_2(v)$.

Note that in this definition, we **do not** require φ_1, φ_2 to be proper colorings. In fact, we can characterize proper colorings through the language of co-proper colorings.

Proposition 2. *A map $\varphi : V \to [c]$ is a proper coloring if and only if φ is co-proper with itself.*

Proof. First, suppose that φ is a proper coloring. Then if $u \sim v$, we know that $\varphi(u) \neq \varphi(v)$, which is precisely the condition guaranteeing that φ is co-proper with itself. Conversely, if φ is co-proper with itself, then for every $u \sim v$, we have that $\varphi(u) \neq \varphi(v)$, which is the condition for φ being a proper coloring. $\square$

Definition 5. Given a graph G and a positive integer c, the *exponential graph $\mathcal{E}_c(G)$* is the graph defined as follows. Its vertex set $V(\mathcal{E}_c(G))$ consists of all functions $V(G) \to [c]$, so that

$$|V(\mathcal{E}_c(G))| = c^{|V(G)|},$$

hence the name "exponential". Two vertices (i.e. two maps $\varphi_1, \varphi_2 : V(G) \to [c]$) are adjacent in $\mathcal{E}_c(G)$ if and only if φ_1 and φ_2 are co-proper colorings.

Recall that we proved that a map φ is a proper coloring if and only if it's co-proper with itself. Thus, $\mathcal{E}_c(G)$ contains loops if and only if G has a proper c-coloring, i.e. if and only if $c \leq \chi(G)$.

The reason exponential graphs are useful for disproving Hedetniemi's conjecture is the following result.

Lemma 1. *Let G be any graph and c any positive integer. Then*

$$\chi(G \times \mathcal{E}_c(G)) \leq c.$$

Proof. To prove this, we need to construct a proper c-coloring Φ of $G \times \mathcal{E}_c(G)$. Such a Φ will be a map $V(G \times \mathcal{E}_c(G)) \to [c]$. So the input of Φ is a pair (v, φ), where v is a vertex of G and φ is a vertex of $\mathcal{E}_c(G)$, i.e. a map $\varphi : V(G) \to [c]$. Given this input, there's a natural choice for $\Phi(v, \varphi) \in [c]$, namely

$$\Phi(v, \varphi) := \varphi(v) \in [c].$$

Thus, our map Φ is defined by evaluation, and we claim that it's a proper c-coloring. To check this, suppose that (v_1, φ_1) and (v_2, φ_2) are adjacent in $G \times \mathcal{E}_c(G)$. Then by the definition of the tensor product, this implies that $v_1 \sim v_2$ in G and that $\varphi_1 \sim \varphi_2$ in $\mathcal{E}_c(G)$. By the definition of the exponential graph, this latter condition implies that φ_1, φ_2 are co-proper colorings. But then since $v_1 \sim v_2$, the definition of co-proper colorings says that $\varphi_1(v_1) \neq \varphi_2(v_2)$. Therefore,

$$\Phi(v_1, \varphi_1) = \varphi_1(v_1) \neq \varphi_2(v_2) = \Phi(v_2, \varphi_2),$$

which shows that Φ is indeed a proper c-coloring. $\square$

2 A bit more on colorings

Definition 6. A set of vertices $I \subseteq V$ is called an *independent set* if I contains no edges, i.e. if for every $u, v \in I$, we have that $(u, v) \notin E$. The *independence number* $\alpha(G)$ of a graph G is defined to be the size of the largest independent set in G.

Proposition 3. *In any proper coloring $\varphi : V \to [c]$, every color class is an independent set, where the ith color class is $\varphi^{-1}(i)$, the set of all vertices receiving color i.*

Proof. Let $I = \varphi^{-1}(i) \subseteq V$ be the ith color class, and suppose I were not an independent set. Then this means that there exist $u, v \in I$ with $u \sim v$. But then by the definition of a proper coloring, this implies that $\varphi(u) \neq \varphi(v)$, so u and v can't both get color i, a contradiction. $\square$

Because of this proposition, we see that a proper c-coloring is the same thing as a partition of V into c disjoint independent sets, and being c-colorable means that such a partition exists. We will occasionally use this interpretation of c-colorability, so it's important to keep in mind that this is an equivalent condition.

3 Colorings of the exponential graph

3.1 Suited colorings

With the final lemma from last time, our task in proving Theorem 1 becomes "merely" to find a graph G and an integer c such that $\chi(G) > c$ and $\chi(\mathcal{E}_c(G)) > c$. Since we have the freedom to choose G, the first condition is not so hard to satisfy, so most of our work will be focused on the second condition: how can we understand the exponential graph well enough to lower-bound its chromatic number, and how can we choose G and c to make this lower bound good?

Specifically, since we want to prove that $\chi(\mathcal{E}_c(G)) > c$ for a certain choice of G and c, we wish to prove that $\mathcal{E}_c(G)$ has no proper c-colorings. To do so, we will collect a large amount of information about the structure that a proper c-coloring of $\mathcal{E}_c(G)$ must satisfy, and then eventually use these properties to derive a contradiction. Our first result in this direction tells us that we can always assume that a proper c-coloring of $\mathcal{E}_c(G)$ plays nicely with the maps that are the vertices of $\mathcal{E}_c(G)$.

Definition 7. Suppose Ψ is a proper c-coloring of $\mathcal{E}_c(G)$. We say that Ψ is *suited* if, for every $\varphi \in V(\mathcal{E}_c(G))$, we have that

$$\Psi(\varphi) \in \mathrm{im}(\varphi).$$

In other words, if Ψ colors some vertex with color i, then that vertex itself uses color i when we think of it as a map.

Lemma 2. *If $\chi(\mathcal{E}_c(G)) \leq c$, then $\mathcal{E}_c(G)$ has a suited c-coloring.*

Proof. Since $\chi(\mathcal{E}_c(G)) \leq c$, there is some proper c-coloring $\Psi : V(\mathcal{E}_c(G)) \to [c]$. Now, for every $i \in [c]$, consider the constant map $\varphi_i : V(G) \to [c]$ that maps every vertex to i, i.e.

$$\varphi_i(v) = i \text{ for every } v \in V(G).$$

Then φ_i is a vertex of $\mathcal{E}_c(G)$. Moreover, observe that if $i \neq i'$, then φ_i and $\varphi_{i'}$ are co-proper colorings, since φ_i and $\varphi_{i'}$ never color any vertices by the same color, so in particular they never color adjacent vertices by the same color. Therefore, $\varphi_1, \ldots, \varphi_c$ give us c vertices in $\mathcal{E}_c(G)$ that are all pairwise adjacent. Since Ψ must assign adjacent vertices different colors, this implies that $\Psi(\varphi_1), \ldots, \Psi(\varphi_c)$ is just some permutation of the colors $1, \ldots, c$. Now, we define a new map $\widetilde{\Psi} : V(\mathcal{E}_c(G)) \to [c]$ by renaming the colors in Ψ according to this permutation. In other words, Ψ and $\widetilde{\Psi}$ have the same color classes, the only difference is that we've permuted the actual colors that they use so that $\widetilde{\Psi}(\varphi_i) = i$ for all i. Since Ψ and $\widetilde{\Psi}$ have the same color classes and since Ψ was a proper c-coloring, then $\widetilde{\Psi}$ is a proper c-coloring as well, and we claim that it's suited.

To see this, let $\varphi : V(G) \to [c]$ be an arbitrary map. Note that if $i \notin \operatorname{im}(\varphi)$, then φ and φ_i are co-proper, since they never assign any vertices the same color, so in particular never assign adjacent vertices the same color. Therefore, since $\widetilde{\Psi}(\varphi_i) = i$ and $\widetilde{\Psi}$ is a proper coloring, we see that $\widetilde{\Psi}(\varphi) \neq i$. Since this holds for all $i \notin \operatorname{im}(\varphi)$, we find that $\widetilde{\Psi}(\varphi) \in \operatorname{im}(\varphi)$, as desired. □

3.2 Large independent sets

Recall that our goal is to show that for a good choice of G and c, the graph $\mathcal{E}_c(G)$ cannot be properly c-colored. Suppose that G has n vertices, so that $\mathcal{E}_c(G)$ has c^n vertices. A proper c-coloring yields a partition of the vertices of $\mathcal{E}_c(G)$ into c independent sets, so the average size of a color class is c^{n-1}. So our first task is to understand what a rather large (i.e. of size roughly c^{n-1}) independent set in $\mathcal{E}_c(G)$ can look like.

There is one natural way to come up with c^{n-1} maps $V(G) \to [c]$. Namely, suppose we fix some vertex $v \in V(G)$ and some color $b \in [c]$. Let $S_{v,b}$ denote the set of all maps $\varphi : V(G) \to [c]$ with the property that $\varphi(v) = b$. Then $|S_{v,b}| = c^{n-1}$, since we have c choices for the color of each vertex in G other than v itself.

Is this $S_{v,b}$ an independent set in $\mathcal{E}_c(G)$? Recall that an independent set is a set that contains no edges, and that a pair of vertices of $\mathcal{E}_c(G)$ form an edge if and only if they are co-proper colorings. So $S_{v,b}$ will be independent if and only if all $\varphi_1, \varphi_2 \in S_{v,b}$ are *not* co-proper. In general, this might not be the case; simply because φ_1 and φ_2 both map v to b doesn't mean they can't be co-proper. So S is not an independent set in general.

However, suppose that the graph G has a loop at v. Then this means that v is adjacent to itself, and $\varphi_1(v) = \varphi_2(v) = b$ for all $\varphi_1, \varphi_2 \in S_{v,b}$. Therefore, if v has a loop, then S is an independent set in $\mathcal{E}_c(G)$, and it has size c^{n-1}.

As it turns out, a useful heuristic is that all large independent sets in $\mathcal{E}_c(G)$ "look like" such an $S_{v,b}$. Of course, this heuristic is fairly problematic, since not all such $S_{v,b}$ even are independent sets, and I didn't tell you what "large" or "look like" even mean.

4 Robust colors

Last time, we came up with a (somewhat problematic) heuristic, which suggests that every large independent set in $\mathcal{E}_c(G)$ looks like a set $S_{v,b} = \{\varphi : V(G) \to [c] \mid \varphi(v) = b\}$. The key lemma in Shitov's proof spells out the extent to which this heuristic is true. First, we will need a definition.

Definition 8. Fix a graph G and a positive integer c. Suppose $\Psi : V(\mathcal{E}_c(G)) \to [c]$ is a suited c-coloring of $\mathcal{E}_c(G)$. Then given a vertex $v \in V(G)$ and a color $b \in [c]$, we say that b *is v-robust* if, for every $\varphi \in \Psi^{-1}(b)$, we have that either $\varphi(v) = b$ or $\varphi(w) = b$ for some w adjacent to v (or both).

This definition is a bit of a mouthful, so let's unpack it a bit. First, let's let $\overline{N}(v)$ denote the *closed neighborhood* of v, namely $\overline{N}(v)$ consists of v and all vertices adjacent to v; in symbols,

$$\overline{N}(v) = \{v\} \cup \{w \in V(G) : v \sim w\}.$$

Next, recall that in any suited coloring Ψ of $\mathcal{E}_c(G)$, if a vertex φ is colored with color b, then φ must map some vertex of G to b (this is the definition of suited colorings). So if $\varphi \in \Psi^{-1}(b)$, then there is some "witness" vertex w with $\varphi(w) = b$; this is a witness that $b \in \operatorname{im}(\varphi)$, and therefore a witness that Ψ can indeed assign the color b to φ. With all of this, we say that the color b is v-robust if this witness w can actually be chosen in $\overline{N}(v)$.

Note that if the $S_{v,b}$ we defined previously were a color class, say $S_{v,b} = \Psi^{-1}(b)$, then the witness of every $\varphi \in \Psi^{-1}(b)$ could in fact be taken to be v itself. v-robustness is a weaker condition: the witness might not be v itself, but it's either v or a neighbor of v. In this sense, a v-robust color class "looks like" one of these sets $S_{v,b}$ we defined before.

With this definition, we can state Shitov's main lemma. The proof in Shitov's paper is three paragraphs, but this is somewhat deceptive, since the argument is fairly subtle and the result is quite strong. It says that in any suited c-coloring of $\mathcal{E}_c(G)$, there is a single vertex v such that almost all the colors are v-robust. In other words, not only do most of the colors "look like" the $S_{v,b}$ from before, but they all do so with respect to the *same* vertex v.

Lemma 3 (Key lemma)**.** *Let G be a graph on n vertices, let c be a positive integer, and suppose that $\mathcal{E}_c(G)$ has a suited c-coloring $\Psi : V(\mathcal{E}_c(G)) \to [c]$. Let $s = \sqrt[n]{n^3 c^{n-1}}$. Then there exists a vertex $v \in V(G)$ such that at least $c - s$ of the colors in $[c]$ are v-robust.*

Proof. For every vertex $u \in V(G)$ and every $b \in [c]$, let

$$I(u, b) = \{\varphi \in \Psi^{-1}(b) : \varphi(u) = b\}.$$

Note that this is similar to the set S from before, except that every map in $I(u, b)$ is assigned color b by Ψ. In other words, $I(u, b)$ consists of all maps that are assigned color b by Ψ and who have u as a witness vertex. By the definition of a suited coloring, every $\varphi \in V(\mathcal{E}_c(G))$ has some witness vertex, which means that every $\varphi \in V(\mathcal{E}_c(G))$ is in at least one $I(u, b)$, though it may be in several of them.

Now, let's say that $I(u,b)$ is *large* if $|I(u,b)| \geq n^2c^{n-2}$. Suppose $I(u,b)$ is large, and fix maps $\varphi \in I(u,b)$ and $\psi \in \Psi^{-1}(b) \setminus I(u,b)$. Note that since $I(u,b) \subseteq \Psi^{-1}(b)$, we know that φ and ψ are both assigned color b by Ψ, which means that they must not be adjacent in $\mathcal{E}_c(G)$. Therefore, φ and ψ are not co-proper colorings. By definition, this means that there exist some vertices $x, y \in V(G)$ such that $x \sim y$ and $\varphi(x) = \psi(y)$. In particular, we find that there exists some vertex $x \in V(G)$ such that $\varphi(x) \in \operatorname{im}(\psi)$.

We've found that for any fixed $\psi \in \Psi^{-1}(b) \setminus I(u,b)$, we have that every $\varphi \in I(u,b)$ has a vertex $x \in V(G)$ such that $\varphi(x) \in \operatorname{im}(\psi)$. Let $I' \subseteq I(u,b)$ denote the set of $\varphi \in I(u,b)$ for which this x can be chosen to be distinct from u, i.e.

$$I' = \{\varphi \in I(u,b) : \exists x \in V(G) \setminus \{u\} \text{ with } \varphi(x) \in \operatorname{im}(\psi)\}.$$

We want to upper-bound $|I'|$. In counting the number of $\varphi \in I'$, we have $n-1$ choices for the vertex x, then $|\operatorname{im}(\psi)|$ choices for $\varphi(x)$, and finally c^{n-2} choices for the image of every vertex other than u and x. Thus,

$$|I'| \leq (n-1)|\operatorname{im}(\psi)|c^{n-2}.$$

Finally, observe that since G has n vertices, $|\operatorname{im}(\psi)| \leq n$. Thus,

$$|I'| \leq (n-1)nc^{n-2} < n^2c^{n-2}.$$

On the other hand, we assumed that $I(u,b)$ is large, meaning that $|I(u,b)| \geq n^2c^{n-2}$. Therefore, since $I' \subseteq I(u,b)$ but $|I'|$ is strictly smaller than $|I(u,b)|$, we may find some $\varphi_0 \in I(u,b)$ that is not in I'. Recall that by the above, there is some vertex x so that $\varphi_0(x) \in \operatorname{im}(\psi)$, and since $\varphi_0 \notin I'$, we must in fact have that $x = u$. Thus, we find that u is the only vertex of G that is mapped by φ_0 into $\operatorname{im}(\psi)$. On the other hand, since $\varphi_0 \in I(u,b)$, we know that $\varphi_0(u) = b$. So we find that $\operatorname{im}(\varphi_0) \cap \operatorname{im}(\psi) = \{b\}$, and that u is the unique vertex in $\varphi_0^{-1}(b)$.

Now, recall that ψ and φ_0 must not be co-proper, since they are both assigned color b by Ψ. So ψ and φ_0 must assign some pair of adjacent vertices the same color, say $x \sim y$ and $\varphi_0(x) = \psi(y)$. However, b is the only color in $\operatorname{im}(\varphi_0) \cap \operatorname{im}(\psi)$, so this same color must be b, i.e. $\varphi_0(x) = \psi(y) = b$. Additionally, since u is the unique vertex in $\varphi_0^{-1}(b)$, we know that $x = u$. Thus, we conclude that $\psi(y) = b$ for some vertex y adjacent to u. In other words, ψ has a witness adjacent to u.

This argument worked for every $\psi \in \Psi^{-1}(b) \setminus I(u,b)$. On the other hand, if $\psi \in I(u,b)$, then $\psi(u) = b$ by definition. So we find that for every $\psi \in \Psi^{-1}(b)$, we can find a witness for ψ that's either u itself or adjacent to u, meaning that the witness for ψ is in $\overline{N}(u)$. Our only assumption in this argument was that $I(u,b)$ is large, namely that $|I(u,b)| \geq n^2c^{n-2}$. So we conclude that if $I(u,b)$ is large, then the color b is u-robust, by definition.

For every vertex $u \in V(G)$, let $B_u \subseteq [c]$ denote the set of colors $b \in [c]$ for which $I(u,b)$ is small. If $|B_v| \leq s$ for some $v \in V(G)$, then we find that $I(v,b)$ is large for at least $c-s$ choices of b, so at least $c-s$ of the colors are v-robust, as desired. So we may assume that $|B_u| > s$ for all $u \in V(G)$. Finally, consider the set of maps

$$T = \{\varphi \in V(\mathcal{E}_c(G)) : \varphi(u) \in B_u \text{ for all } u \in V(G)\}.$$

Then since $|B_u| > s$ for all u, we find that $|T| > s^n$, since we have more than s choices for each vertex u. By the definition of s, we know that $s^n = n^3c^{n-1}$. Moreover, note that every $\varphi \in T$ does not lie in any large $I(u,b)$, by the definition of B_u. So we conclude that

$$|T| \leq \sum_{\substack{u \in V(G), b \in [c] \\ I(u,b) \text{ is small}}} |I(u,b)| \leq \sum_{\substack{u \in V(G), b \in [c] \\ I(u,b) \text{ is small}}} n^2c^{n-2} \leq \sum_{\substack{u \in V(G) \\ b \in [c]}} n^2c^{n-2} = (nc)(n^2c^{n-2}) = n^3c^{n-1}.$$

We've found that $|T| \leq n^3c^{n-1}$ and that $|T| > n^3c^{n-1}$, which is a contradiction. Therefore, there is some vertex v for which $I(v,b)$ is large for at least $c - s$ choices of b, which finishes the proof. $\square$

5 Constructing the counterexamples

The key lemma in the previous section shows gives us a lot of structural information about what suited colorings of $\mathcal{E}_c(G)$ can look like. In this section, we'll see how to use this structural information to cleverly pick G and c so that $\mathcal{E}_c(G)$ cannot be c-colored.

5.1 The strong product

Definition 9. Given two graphs G, H, their strong product $G \boxtimes H$ is the graph with vertex set $V(G \boxtimes H) = V(G) \times V(H)$ and (g_1, h_1) adjacent to (g_2, h_2) if one of the following conditions hold:

$$g_1 \sim g_2, h_1 \sim h_2 \qquad \text{or} \qquad g_1 \sim g_2, h_1 = h_2 \qquad \text{or} \qquad h_1 \sim h_2, g_1 = g_2.$$

Note that $G \times H$ is a subgraph of $G \boxtimes H$.

Example. Suppose G is any graph, and let K_q denote the complete graph on q vertices. Then $G \boxtimes K_q$ looks like a copy of G, where we've replaced every vertex of G by a copy of K_q, and every edge of G by all q^2 possible edges between the corresponding K_qs.

If G is a simple graph, let G° be the graph we get by adding a loop to every vertex of G. Our next lemma shows that the exponential graph of $G \boxtimes K_q$ is closely related to the exponential graph of G°.

Lemma 4. *Let G be any simple graph and c, q any positive integers. There is an injective map $\iota : V(\mathcal{E}_c(G^\circ)) \hookrightarrow V(\mathcal{E}_c(G \boxtimes K_q))$ that maps every edge of $\mathcal{E}_c(G^\circ)$ to an edge of $\mathcal{E}_c(G \boxtimes K_q)$; in other words, ι realizes $\mathcal{E}_c(G^\circ)$ as a subgraph of $\mathcal{E}_c(G \boxtimes K_q)$. In particular, a suited coloring of $\mathcal{E}_c(G \boxtimes K_q)$ gives rise to a suited coloring of $\mathcal{E}_c(G^\circ)$.*

Proof. To define ι, we need to decide how to take a map $\varphi : V(G^\circ) \to [c]$ and output a map $\iota(\varphi) : V(G \boxtimes K_q) \to [c]$. There is a natural guess, namely the map $\varphi^* : V(G \boxtimes K_q) \to [c]$ defined by

$$\varphi^*(g, k) = \varphi(g),$$

i.e. the map that ignores the second coordinate and applies φ to the first. Indeed, we define $\iota(\varphi) = \varphi^*$. Then ι is certainly injective, since φ^* uniquely determines φ. To see that ι maps edges to edges, suppose that $\varphi_1 \sim \varphi_2$ in $\mathcal{E}_c(G^\circ)$, meaning that φ_1, φ_2 are co-proper colorings of G°. We need to check that φ_1^* and φ_2^* are co-proper colorings of $G \boxtimes K_q$, so let (g_1, k_1) be adjacent to (g_2, k_2) in $G \boxtimes K_q$. If $g_1 \sim g_2$, then

$$\varphi_1^*(g_1, k_1) = \varphi_1(g_1) \neq \varphi_2(g_2) = \varphi_2^*(g_2, k_2),$$

since φ_1, φ_2 are co-proper. On the other hand, if $g_1 \nsim g_2$, then since G° has loops on every vertex, this also implies that $g_1 \neq g_2$. But then (g_1, k_1) and (g_2, k_2) cannot be adjacent, by the definition of the strong product, so ι indeed maps edges to edges.

For the final statement of the lemma, note that if $\Psi : V(\mathcal{E}_c(G \boxtimes K_q)) \to [c]$ is a suited coloring, then $\iota \circ \Psi$ is a proper coloring of $\mathcal{E}_c(G^\circ)$ (since $\mathcal{E}_c(G^\circ)$ is a subgraph of $\mathcal{E}_c(G \boxtimes K_q)$), and it's suited since φ and φ^* have the same images in $[c]$. $\square$

5.2 Constructing uncolorable maps

Now, we'll show that if we pick a graph G' and an integer c appropriately, then $\chi(\mathcal{E}_c(G')) > c$. Below, we'll take G' to be the strong product of a fixed graph G with a large complete graph K_q. First, we need two definitions.

Definition 10. The *girth* of a graph G, denoted $\operatorname{girth}(G)$, is the length of the shortest cycle in G. If G has no cycles, we declare $\operatorname{girth}(G) = \infty$.

Definition 11. Given two vertices x, y in a graph G, their *distance* $\operatorname{dist}(x, y)$ is the number of edges in the shortest path connecting x and y. Thus, $\operatorname{dist}(x, y) = 1$ if and only if $x \sim y$, and $\operatorname{dist}(x, y) \leq 2$ if x and y have a common neighbor, and so on.

Theorem 2. *Suppose G is a graph with girth at least 6 on n vertices. Let $q \geq 2^n n^3/6$ be a large integer, and let $c = 6q$. Then $\chi(\mathcal{E}_c(G \boxtimes K_q)) > c$.*

Proof. Suppose for contradiction that $\chi(\mathcal{E}_c(G \boxtimes K_q)) \leq c$, meaning that $\mathcal{E}_c(G \boxtimes K_q)$ has a proper c-coloring. By Lemma 2, this implies that there is a suited c-coloring $\Psi : V(\mathcal{E}_c(G \boxtimes K_q)) \to [c]$. Moreover, by Lemma 4, we know that $\mathcal{E}_c(G^\circ)$ is a subgraph of $\mathcal{E}_c(G \boxtimes K_q)$, meaning that we also get a suited c-coloring Ψ° of $\mathcal{E}_c(G^\circ)$.

Now, recall our key lemma, Lemma 3. It tells us that if we let $s = \sqrt[n]{n^3 c^{n-1}}$, then there is some vertex $v \in V(G^\circ)$ such that at least $c - s$ of the colors in the coloring Ψ° are v-robust. Fix this vertex v.

Now, for every $i \in \{2q + 1, 2q + 2, \ldots, c\}$, we define a map $\mu_i : V(G \boxtimes K_q) \to [c]$ by

$$\mu_i(g, k) = \begin{cases} k & \text{if } \operatorname{dist}(v, g) \in \{0, 2\}, \\ q + k & \text{if } \operatorname{dist}(v, g) = 1, \\ i & \text{otherwise.} \end{cases}$$

In this definition, g is a vertex of G and k is a vertex of K_q, which we think of as a number between 1 and q. We claim that if $i \neq i'$, then μ_i and $\mu_{i'}$ are co-proper colorings of $G \boxtimes K_q$. To see this, suppose that (g_1, k_1) and (g_2, k_2) are adjacent vertices in $G \boxtimes K_q$. Notice that $\mu_i(g_1, k_1) \in \{k_1, q + k_1, i\}$, while $\mu_{i'}(g_2, k_2) \in \{k_2, q + k_2, i'\}$. Moreover, we have that

$$k_1, k_2 \in \{1, \ldots, q\} \qquad \text{and} \qquad q+k_1, q+k_2 \in \{q+1, \ldots, 2q\} \qquad \text{and} \qquad i, i' \in \{2q+1, \ldots, c\}.$$

Thus, since $i \neq i'$, the only way for $\mu_i(g_1, k_1)$ to equal $\mu_{i'}(g_2, k_2)$ is for k_1 to equal k_2, for otherwise μ_i and $\mu_{i'}$ will use disjoint sets of colors. Thus, we now assume $k_1 = k_2$. But by the definition of the strong product, the fact that $(g_1, k_1) \sim (g_2, k_2)$ now implies that $g_1 \sim g_2$. As you saw on the homework, the fact that G has girth at least 6 means that if g_1, g_2 are adjacent and both have distance at most 2 from v, then $\operatorname{dist}(v, g_1)$ and $\operatorname{dist}(v, g_2)$ must have different parities. But in that case μ_i and $\mu_{i'}$ will assign these vertices colors $q + k_1$ and k_1 (in some order), and in particular they won't get the same color. On the other hand, if one of g_1, g_2 has distance at least 3 from v, then it will be assigned color i or i', and again $\mu_i(g_1, k_1)$ will not equal $\mu_{i'}(g_2, k_2)$. Thus, we indeed find that $\mu_i, \mu_{i'}$ are co-proper colorings.

Last time, we found a set $\{\mu_{2q+1}, \ldots, \mu_c\}$ of $c - 2q$ vertices of $\mathcal{E}_c(G \boxtimes K_q)$ that are all pairwise adjacent. This means they must all receive distinct colors under the proper coloring Ψ. Since $c - 2q = (6-2)q > 2q$, the pigeonhole principle implies that one of the colors $\Psi(\mu_{2q+1}), \ldots, \Psi(\mu_c)$ must not lie in the set $\{1, \ldots, 2q\}$. Let $t \in \{2q+1, \ldots, c\}$ be some index with this property, i.e. one for which $\Psi(\mu_t) \notin \{1, \ldots, 2q\}$. On the other hand, we know that Ψ is a suited coloring, which means that $\Psi(\mu_t) \in \operatorname{im}(\mu_t)$, and by definition, we have that $\operatorname{im}(\mu_t) = \{1, \ldots, 2q, t\}$. Thus, we conclude that $\Psi(\mu_t) = t$.

Now, recall that at least $c - s$ of the colors in Ψ° are v-robust. By our assumption that $q \geq 2^n n^3/6$, we see that $c \geq 2^n n^3$, which implies that

$$s = \sqrt[n]{n^3 c^{n-1}} = c\sqrt[n]{\frac{n^3}{c}} \leq c\sqrt[n]{\frac{n^3}{2^n n^3}} = \frac{c}{2}.$$

Therefore, $c - s \geq c/2 = 3q$. So at least $3q > 2q + 1$ colors of Ψ° are v-robust, so we can pick some color $r \in [c]$ which is v-robust and with $r \notin \{1, \ldots, 2q, t\}$, again by the pigeonhole principle. We define another map $\rho : V(\mathcal{E}_c(G \boxtimes K_q)) \to [c]$ by

$$\rho(g, k) = \begin{cases} t & \text{if } \operatorname{dist}(v, g) \leq 1 \\ r & \text{otherwise.} \end{cases}$$

Observe that ρ ignores the K_q coordinate, which means that it lies in the image of the embedding ι from Lemma 4. This means we can also think of ρ as a vertex of $\mathcal{E}_c(G^\circ)$. Consider $\Psi^\circ(\rho)$. Since Ψ° is a suited coloring, we know that $\Psi^\circ(\rho) \in \operatorname{im}(\rho) = \{t, r\}$. If $\Psi^\circ(\rho) = r$, then since r is a v-robust color in the color Ψ°, there must be a witness vertex w satisfying $\rho(w) = r$ and $w \in \overline{N}(v)$. However, by the definition of ρ, we see that this is impossible: the only vertices that ρ maps to r are those outside of $\overline{N}(v)$. Therefore, we must have that $\Psi^\circ(\rho) = t$, and since Ψ° is just the restriction of Ψ onto the subgraph $\mathcal{E}_c(G^\circ)$, we also conclude that $\Psi(\rho) = t$.

On the other hand, observe that ρ and μ_t are co-proper. Indeed, the only color in $\operatorname{im}(\rho) \cap \operatorname{im}(\mu_t)$ is t, so the only way ρ and μ_t could not be co-proper is if they assigned two adjacent vertices color t. However, ρ only assigns color t to those vertices whose first coordinate has distance at most 1 from v, while μ_t only assigns color t to those vertices whose first coordinate has distance at least 3 from v. No vertex of distance ≤ 1 and distance ≥ 3 can be adjacent, so ρ and μ_t are co-proper. But $\Psi(\rho) = t = \Psi(\mu_t)$, which means that Ψ is not a proper coloring, giving us our contradiction. □

5.3 Lower bounding the chromatic number of a strong product

Recall from the homework that for any graph G,

$$\chi(G) \geq \frac{|V(G)|}{\alpha(G)},$$

since each color class is an independent size, and thus has size at most $\alpha(G)$. On the homework, you saw that

$$\chi(G \boxtimes H) \leq \chi(G)\chi(H). \tag{1}$$

If this inequality were always an equality, we could conclude that

$$\chi(G \boxtimes K_q) = \chi(G)\chi(K_q) \geq q\frac{|V(G)|}{\alpha(G)}. \tag{2}$$

However, the inequality (1) is not an equality in general. Nevertheless, the bound in (2) is always true.

Proposition 4. *For any graph G and any positive integer q,*

$$\chi(G \boxtimes K_q) \geq \frac{q|V(G)|}{\alpha(G)}.$$

Proof. Let I be an independent set of $G \boxtimes K_q$. Then if (g_1, k_1) and (g_2, k_2) are distinct vertices in I, then they must not be adjacent. We split our analysis into three cases.

- If $g_1 = g_2$, then since our vertices are distinct, we must have $k_1 \neq k_2$. But then since K_q is the complete graph, this implies that $k_1 \sim k_2$, so by the definition of the strong product, we have that $(g_1, k_1) \sim (g_2, k_2)$, a contradiction.
- Next, suppose $g_1 \sim g_2$. If $k_1 = k_2$, then by the definition of the strong product, $(g_1, k_1) \sim (g_2, k_2)$. On the other hand, if $k_1 \neq k_2$, then $k_1 \sim k_2$, so again $(g_1, k_1) \sim (g_2, k_2)$. In either case we get a contradiction.
- Therefore, we must have $g_1 \nsim g_2$ and $g_1 \neq g_2$.

This implies that for any pair of vertices in I, they have distinct first coordinates and the first coordinates are non-adjacent in G. This means that if we consider the projection map $\pi : I \to V(G)$ that forgets the second coordinate, then we find that π is an injection and its image is an independent set in G. Therefore, $|I| \leq \alpha(G)$. Since this holds for any independent set I in $G \boxtimes K_q$, we conclude that $\alpha(G \boxtimes K_q) \leq \alpha(G)$. Therefore,

$$\chi(G \boxtimes K_q) \geq \frac{|V(G \boxtimes K_q)|}{\alpha(G \boxtimes K_q)} \geq \frac{q|V(G)|}{\alpha(G)},$$

as desired. □

5.4 Putting it all together

The main theorem we proved yesterday and today shows that if G has girth at least 6 and if q is large enough, then $\chi(\mathcal{E}_c(G')) > c$, where $c = 6q$ and $G' = G \boxtimes K_q$. Additionally, by Proposition 4, we know that

$$\chi(G') = \chi(G \boxtimes K_q) \geq \frac{|V(G)|}{\alpha(G)}q.$$

Thus, if we choose G so that $|V(G)| > 6\alpha(G)$, we will conclude that $\chi(G') > 6q = c$. However, as we saw in Lemma 1, we always have that

$$\chi(G' \times \mathcal{E}_c(G')) \leq c.$$

So if we can find a graph G with $\mathrm{girth}(G) \geq 6$ and $|V(G)| > 6\alpha(G)$, we can disprove Hedetniemi's conjecture.

How do we go about finding such a graph? If you try playing around with small examples, you're unlikely to succeed: you'll find that the condition that $\mathrm{girth}(G)$ be large means that it's very easy to produce large independent sets. Because of this, for a long time no one was really sure if graphs like this existed. Nevertheless, Erdős proved the following amazing result.

Theorem 3 (Erdős, 1959)**.** *For every $a, b > 0$, there exists some graph G with* $\mathrm{girth}(G) > a$ *and $|V(G)| > b\alpha(G)$.*

This theorem is one of the famous examples of the *probabilistic method*: rather than explicitly describing how to construct such a graph, Erdős merely guarantees that a randomly chosen graph (according to some specific probability distribution) will have this property with strictly positive probability. Since the probability would be zero if no such graph existed, then one must exist, even though we have absolutely no idea how to actually go about finding it. Since Erdős's work, people have since developed techniques to explicitly construct such graphs (probably the most important such graphs are the Ramanujan graphs of Lubotzky, Phillips, and Sarnak), but the probabilistic approach is still the simplest and most versatile. Unfortunately, it is still somewhat complicated, and uses a lot of techniques that we don't have time to cover, so we'll have to simply accept Erdős's theorem as a black-box result.

However, with this theorem in hand, we are done. Let's fix a graph G with $\mathrm{girth}(G) \geq 6$ and $|V(G)| > 6\alpha(G)$, and define $n = |V(G)|$. We now pick an integer q larger than $2^n n^3/6$, let $c = 6q$, and define $G' = G \boxtimes K_q$. By the argument above, we know that $\chi(G') > c, \chi(\mathcal{E}_c(G')) > c$, but $\chi(G' \times \mathcal{E}_c(G')) \leq c$, which disproves Hedetniemi's conjecture.

Just how big is our counterexample? As far as I know, the smallest known graph with girth at least 6 and $n > 6\alpha(G)$ has about $n \approx 10^{10}$ vertices. This means that we need to take

$$q \gtrsim 2^n n^3 \approx 2^{10^{10}} 10^{30} \approx 2^{10^{10}},$$

and c is of basically the same size. This means that our first graph $G' = G \boxtimes K_q$ has $qn \approx 2^{10^{10}}$ vertices, which is way way more than the number of atoms in the universe, approximately 10^{80}. On the other hand, our second graph $\mathcal{E}_c(G')$ has

$$c^{qn} \approx \left(2^{10^{10}}\right)^{2^{10^{10}}} = 2^{2^{10^{10}} \cdot 10^{10}} \approx 2^{2^{10^{10}}}$$

vertices. This is an incomprehensibly big number.

The argument I presented is not optimal in terms of these constants, and one can do a bit better by being more careful: one can take n to only be around 2 million, and q to only be of order roughly n^3, so that G' has around 10^{25} vertices. This is much smaller! It's only about as many grains of sand as are on earth. However, with $\mathcal{E}_c(G')$, we're again in big trouble, since it will have about $10^{10^{25}}$ vertices.

6 What's next?

Hedetniemi's conjecture is now known to be false, but there's still a lot to do, and I suspect that there will be many developments in this area over the next few years. To state the questions one can ask as follow-ups, we need the following definition.

Definition 12. The *Poljak–Rödl function* is a function $f : \mathbb{N} \to \mathbb{N}$ defined by

$$f(c) = \min_{\substack{\text{graphs } G,H \\ \chi(G),\chi(H)\geq c}} \chi(G \times H).$$

In other words, the Poljak–Rödl function measures how close Hedetniemi's conjecture is to being true; our simple observation that $\chi(G \times H) \leq \min\{\chi(G), \chi(H)\}$ implies that $f(c) \leq c$ for all c, and Hedetniemi's conjecture asserts that $f(c) = c$ for all c. What we showed is that this is false: there is some c for which $f(c) < c$.

The natural next question is how quickly f grows as a function of c; we know it grows strictly slower than the identity function $c \mapsto c$, but how much slower? The simplest version of this question is the following conjecture.

Conjecture ("Weak Hedetniemi conjecture")**.**

$$\lim_{c\to\infty} f(c) = \infty.$$

In other words, the weak Hedetniemi conjecture asserts that f is not bounded: if we pick c large enough, then we can make $f(c)$ take an arbitrarily large value. This conjecture is still open, though the following very surprising result of Poljak and Rödl suggests that it's likely true.

Theorem 4 (Poljak–Rödl)**.** *Either* $\lim_{c\to\infty} f(c) = \infty$ *or else* $f(c) \leq 9$ *for all* c*.*

In other words, if one could find some c such that $f(c) \geq 10$, then the weak Hedetniemi conjecture is true.

The weak Hedetniemi conjecture asserts that Hedetniemi's conjecture can't be "too" false—maybe $\chi(G \times H)$ doesn't always equal $\min\{\chi(G), \chi(H)\}$, but at least it can't lag too far behind, since it needs to tend to infinity. The other natural question is to ask for upper bounds on $f(c)$, which corresponds to proving that Hedetniemi's conjecture is more than just false, but in fact quite false. The following result was proved by me and another grad student, Xioayu He.

Theorem 5. *There is some constant* $\varepsilon > 0$ *such that for all sufficiently large* c*,*

$$f(c) \leq (1 - \varepsilon)c.$$

In other words, this theorem says that not only does $f(c)$ not equal c, but in fact the ratio $f(c)/c$ is bounded away from 1 when c is large. The value of ε we get is pretty small (roughly 10^{-9}), but this is after several improvements to the paper—our original proof gave $\varepsilon \approx 2^{-10^{20}}$!

This result doesn't rule out that f grows linearly in c; for instance, maybe $f(c) \geq .99c$ for all c. Nevertheless, the following conjecture says this doesn't happen, and seems plausible, though I think some pretty substantial new ideas would need to be introduced to prove it.

Conjecture (Tardif–Zhu)**.** *f grows sub-linearly, i.e.*

$$\lim_{c \to \infty} \frac{f(c)}{c} = 0.$$

In other words, for every $\delta > 0$, there is some sufficiently large c for which $f(c) \leq \delta c$.