A SAT Solver's Opinion on the Erdős-Faber-Lovász Conjecture

Markus Kirchweger
Algorithms and Complexity Group, TU Wien, Austria

Tomáš Peitl
Algorithms and Complexity Group, TU Wien, Austria

Stefan Szeider
Algorithms and Complexity Group, TU Wien, Austria

Abstract

In 1972, Paul Erdős, Vance Faber, and Lászlo Lovász asked whether every linear hypergraph with n vertices can be edge-colored with n colors, a statement that has come to be known as the *EFL conjecture*. Erdős himself considered the conjecture as one of his three favorite open problems, and offered increasing money prizes for its solution on several occasions. A proof of the conjecture was recently announced, for all but a finite number of hypergraphs. In this paper we look at some of the cases not covered by this proof.

We use SAT solvers, and in particular the SAT Modulo Symmetries (SMS) framework, to generate non-colorable linear hypergraphs with a fixed number of vertices and hyperedges modulo isomorphisms. Since hypergraph colorability is NP-hard, we cannot directly express in a propositional formula that we want only non-colorable hypergraphs. Instead, we use one SAT (SMS) solver to generate candidate hypergraphs modulo isomorphisms, and another to reject them by finding a coloring. Each successive candidate is required to defeat all previous colorings, whereby we avoid having to generate and test all linear hypergraphs.

Computational methods have previously been used to verify the EFL conjecture for small hypergraphs. We verify and extend these results to larger values and discuss challenges and directions. Ours is the first computational approach to the EFL conjecture that allows producing independently verifiable, DRAT proofs.

2012 ACM Subject Classification Theory of computation → Automated reasoning; Mathematics of computing → Hypergraphs; Mathematics of computing → Graph coloring; Mathematics of computing → Graph enumeration; Mathematics of computing → Matchings and factors

Keywords and phrases hypergraphs, graph coloring, SAT modulo symmetries

Digital Object Identifier 10.4230/LIPIcs.SAT.2023.13

Supplementary Material *Software (Source Code)*: `https://github.com/markirch/sat-modulo-symmetries/`, archived at `swh:1:dir:c50f1ce7d22717e113b915ae9976da0cb0e67c1d`
Software (Documentation): `https://sat-modulo-symmetries.readthedocs.io/`

Funding This work was supported by the Austrian Science Fund (FWF, project P32441), and the Vienna Science and Technology Fund (WWTF, project ICT19-065).

1 Introduction

In 1972, Paul Erdős and László Lovász gathered at Vance Faber's apartment for some tea, and to prove what was supposed to be an easy theorem about hypergraph edge colorings: every linear hypergraph can be edge-colored with no more colors than it has vertices. A hypergraph is a collection of subsets, called (hyper)edges, of a finite set, and it is linear if any two hyperedges intersect in at most one element (vertex). An edge coloring is a coloring of the edges such that intersecting edges have different colors. The conjecture, now known as

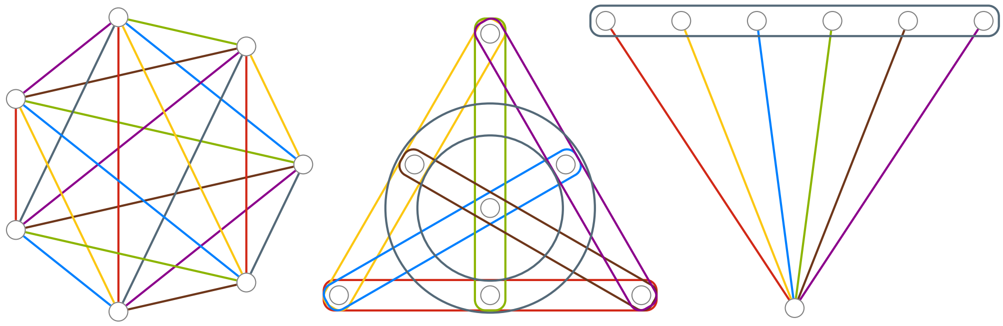

Figure 1 The famous extremal examples to the EFL conjecture for $n = 7$. Left to right: the complete graph, the projective plane of order 2 (also known as the Fano plane), and the degenerate plane. Hyperedges are drawn as bounding boxes around vertices, 2-edges are drawn simply as lines. Each hypergraph is colored with its optimal number of colors; in the case of the planes this is trivial as the number of colors equals the number of hyperedges.

the Erdős-Faber-Lovász conjecture (EFL) conjecture, is inspired by its prominent extremal examples, the complete graph, the degenerate plane, and the projective plane (see Figure 1), each of which requires the full number of colors.

Needless to say, they did not find an easy proof, and the conjecture has remained open for fifty years. Erdős wrote about it as one of his favorite combinatorial problems [7] and offered money prizes for a solution [6]. Throughout its lifetime, the conjecture attracted a lot of attention, and many partial results have been proved. The newest in the series is the recently announced proof of the conjecture for all but a finite number of cases [17].

In this paper, we look at the conjecture by computational means. We investigate classes of small hypergraphs, and verify that they can indeed be colored with the conjectured number of colors. In this respect we follow in the footsteps of Hindman [12] and Romero and Alonso-Pecina [22], who also investigated small classes of hypergraphs and partially verified the EFL conjecture. The novel aspects of our work are:

1. we verify the EFL conjecture for more classes of small hypergraphs;
2. where the previous two works had to enumerate and afterwards color all linear hypergraphs, we interleave these two processes: we generate hypergraphs and colorings alternately, and thus our method is not constrained by the number of linear hypergraphs, but only by the number of colorings needed to color them all. In doing this we build on the recently-proposed *co-certificate learning (CCL)* [18].
3. We use SAT encodings and SAT solvers, and hence our method enables us to produce independently verifiable proofs, including of the previous computational results. We are able to use them effectively due to the recently introduced *SAT modulo symmetries (SMS)* framework [20], in which a CDCL SAT solver [9] is enhanced with a custom symmetry propagator which prunes non-canonical search paths. In this work, we extend SMS, which was previously used for graphs, and later matroids, to hypergraphs. We achieve this by modeling hypergraphs by their bipartite incidence graphs, which in turn can be fit into the SMS framework with some care.

We report on the results of three experiments.

In the first, we tackle the EFL conjecture itself and analyze the nature of the hardness of verifying the conjecture experimentally in cases we processed. Naturally, we have not found any counterexamples (the title of the paper would otherwise have been different).

In the second experiment, we search for extremal examples to the EFL conjecture, namely linear hypergraphs with n vertices that require n colors. In addition to the known extremal examples, we were able to enumerate all other small extremal hypergraphs. Inspired by our findings, we also proved extremality of some infinite families of linear hypergraphs by hand.

In the last experiment, we tackle a generalization of the EFL conjecture (Conjecture 3), which at the same time also generalizes Vizing's Theorem about edge colorings in graphs. This experiment proceeds in a very similar way as the first, and we did not find any counterexamples in this case either.

In the next sections, we first review the necessary background on hypergraphs and colorings, followed by an explanation of our SAT-based process, and a discussion of results.

2 The EFL Conjecture

For positive integers n, m, let $[n] = \{1, \ldots, n\}$ and $[n, m] = \{n, n+1, \ldots, m\}$. A hypergraph $H = (V, E)$ consists of a finite vertex set V (we assume $V = [n]$), and a set E of subsets of V called *hyperedges*, or sometimes simply *edges*. A hyperedge of size k is also called a *k-edge*. A hypergraph where each hyperedge has size 2 is called a *graph*. The graph, denoted by K_n, which contains all possible 2-edges between n vertices is called the *complete graph*, or the *clique*. The *degree* of a vertex is the number of hyperedges that contain v. A hypergraph is called *linear* if any two of its hyperedges intersect in at most one vertex. Every graph is trivially a linear hypergraph.

A *χ-edge coloring* of a hypergraph $H = (V, E)$ is a mapping $c : E \to [\chi]$ which colors intersecting hyperedges with different colors: $e \cap f \neq \emptyset \implies c(e) \neq c(f)$. The *chromatic index* of a hypergraph is the smallest integer χ for which a χ-edge coloring exists. The *Erdős-Faber-Lovász (EFL) conjecture* postulates the existence of certain edge colorings.

▶ **Conjecture 1** (EFL conjecture). *Every linear hypergraph with n vertices is n-edge-colorable.*

While this paper is primarily about edge colorings of hypergraphs, we will also need to work with vertex colorings of graphs. A *χ-coloring* of a graph $G = (V, E)$ is a mapping $c : V \to [\chi]$ that maps adjacent vertices to different colors: $\{u, v\} \in E \implies c(u) \neq c(v)$. The *chromatic number* of a graph is the smallest χ for which a χ-coloring exists.

A hypergraph H is *covered* if every pair of vertices $\{u, v\} \subseteq V$ is contained in some hyperedge $e \in E$. Linear hypergraphs which are covered and whose hyperedges have size at least two are known as *linear spaces*.

▶ **Observation 2.** *If there is a counterexample to the EFL Conjecture with n vertices, then there is a counterexample with n vertices which is a linear space.*

Indeed, one can simply delete singleton hyperedges without decreasing the chromatic index, due to the fact that singleton hyperedges intersect at most $n-1$ other hyperedges in a linear hypergraph, and any n-edge coloring can always be extended to them. Similarly, adding a 2-edge to cover a pair $\{u, v\}$ cannot decrease the chromatic index either. By Observation 2, we thus restrict our search for counterexamples to the EFL conjecture to linear spaces.

The EFL conjecture, if true, is tight, as witnessed by the examples of the (odd) complete graphs, the degenerate planes, and the projective planes, each of which requires n colors.

When restricted to graphs, the EFL conjecture is easy to prove. The only linear space on n vertices is the complete graph, and it is an instructive exercise to verify that it is indeed n-colorable. In fact, a stronger result known as Vizing's Theorem holds for graphs: a graph of maximum degree Δ can be $(\Delta+1)$-edge-colored (for these graphs we can no longer assume they are covered). Together with the trivial bound arising from the fact that the edges around a vertex of degree Δ require Δ different colors, Vizing's Theorem implies that the chromatic index of any graph is either Δ or $\Delta + 1$. Again, it is easy to verify that Vizing's Theorem in terms of the maximum degree does not hold for linear hypergraphs in general, but a different, beautiful generalization has been conjectured. This generalization arises from a reformulation of Vizing's Theorem in terms of the *closed neighborhood* of a vertex, which is the union of the hyperedges containing a vertex, $\bigcup_{v\in e} e$. In the case of graphs, the size of the closed neighborhood of v is, by definition, the degree of v + 1, and the maximum size of any closed neighborhood is $\Delta + 1$. Vizing's Theorem then says that a graph can be edge-colored with a number of colors equal to the largest size of a closed neighborhood. Now, *this* version has been conjectured for linear hypergraphs in general.

▶ **Conjecture 3** (Füredi [10], Berge [1]). *Let $H = (V, E)$ be a linear hypergraph, and let $\zeta = \max_{v\in V} \left|\bigcup_{v\in e} e\right|$. Then H is ζ-edge-colorable.*

We call this conjecture the FB Conjecture.

In the context of this conjecture, we can no longer assume that hypergraphs are covered, because arbitrary edge addition could increase neighborhood size beyond the admissible bound. Instead, we say a hypergraph is *weakly covered* if each vertex pair $\{u, v\} \subseteq V$ either is contained in some hyperedge $e \in E$ or one of the two vertices has maximum closed neighborhood. We state a similar observation as before:

▶ **Observation 4.** *If there is a counterexample to the FB Conjecture with n vertices, then there is a weakly covered counterexample with n vertices without hyperedges of size 1.*

Note how the two formulations of Vizing's Theorem coincide for graphs, but in presence of larger hyperedges the difference between counting and unifying the neighborhood materializes.

In our experiments, we will validate both conjectures, and we will also search for extremal examples in addition to those listed in Figure 1.

2.1 Previous Work

A large body of work on the EFL conjecture (Conjecture 1) is available, roughly categorizable into two groups. In the first group, there are the asymptotic results [3, 13–16, 23], of which the culmination is the recently announced proof of Conjecture 1 for all hypergraphs with a sufficiently large number of vertices [17]. If correct, this result, at least in theory, leaves the complete verification of the conjecture down to a finite computation, the verification of the colorability of a finite number of linear spaces.

In the second group, there are the results which verify the conjecture for a finite number of small hypergraphs. The first such result known to us is due to Hindman [12], who verified Conjecture 1 for all hypergraphs with at most 10 vertices. In fact, Hindman proved a stronger result, namely that every hypergraph whose hyperedges of size at least 3 span at most 10 vertices is colorable; and thus in particular those with at most 10 vertices in total.

▶ **Theorem 5** (Hindman [12]). *Let $H = (V, E)$ be a linear hypergraph with n vertices, let $E' = \{e \in E : |e| \geq 3\}$. If $|\bigcup E'| \leq 10$, then H is n-edge colorable.*

Later, Romero and Alonso-Pecina, building on an existing program to enumerate linear spaces [2], verified Conjecture 1 up to $n \leq 12$ [22]; their contribution lied mainly in a meta-heuristic coloring algorithm. In our work, we replace ad-hoc hypergraph generation and coloring tools with a SAT-based workflow, which is more flexible, reliable, and easier to use. We verify the existing results and explore some generalizations and larger cases.

3 Reduction to SAT

In the sequel, we shall implicitly generate linear hypergraphs with n vertices and check that they are colorable with the right number of colors χ in order to validate Conjectures 1 and 3. For the former conjecture, we set $\chi = n$, for the latter, χ is the size of the largest closed neighborhood. Since in each run we will generate hypergraphs with a fixed number of hyperedges, we first need to determine bounds on the number m of hyperedges to be considered for a given number of vertices n.

We can assume the number of hyperedges is at least $\chi + 1$, as $\leq \chi$ hyperedges are trivially colorable by χ colors, though for generating extremal examples, we also need to consider $m = \chi$, so we take $m \geq \chi$. On the upper side, since each hyperedge covers at least one pair of vertices and no pair is covered twice (which is the same as linearity), the largest possible number of hyperedges is $\binom{n}{2}$, when the hypergraph is the complete graph on n vertices.

Additionally, we impose the restrictions given by Observation 2 and 4 respectively.

To enumerate linear hypergraphs with n vertices and m hyperedges, we produce a propositional formula $\mathcal{L}(n, m)$, whose models are exactly these linear hypergraphs. Before we describe the encoding itself, let us discuss the entire process including coloring and other parts on a high level. The two hallmarks of our approach are the use of *Satisfiability modulo symmetries (SMS)* [20] to enumerate hypergraphs without isomorphic copies, and an χ-edge-coloring learning technique thanks to which we do not need to enumerate all hypergraphs in order to show that they are all colorable.

▶ **Definition 6.** *Two hypergraphs $H_1 = (V, E_1)$ and $H_2 = (V, E_2)$ are* isomorphic*, written $H_1 \cong H_2$, if there is a bijection $\pi : V \to V$, extended to hyperedges and sets of hyperedges in the obvious way, so that $H_2 = \pi(H_1)$. For any total order $\preccurlyeq$ on the set of hypergraphs, we say that a hypergraph $H = (V, E)$ is* canonical *if $H \preccurlyeq \pi(H)$ for every $\pi : V \to V$.*

SMS is a general framework which augments a SAT solver with a custom propagator to check whether the currently constructed hypergraph is canonical. If not, a *symmetry clause* is learned and the solver backtracks. SMS performs *dynamic symmetry breaking*, and thus prevents the solver from redundantly exploring isomorphic (symmetric) copies of the search space. Other methods of symmetry exploitation include *static* symmetry breaking, where a single symmetry-breaking formula is added at the beginning, but those do not break all symmetries [4]. We expand on our use of SMS in Section 3.2.

The other important technique that we use is the *learning of coloring clauses*, recently proposed under the name *co-certificate learning* [18]. A naive approach to our problem, even with symmetry breaking, would consist of enumerating, modulo isomorphisms, all hypergraphs with n vertices and m hyperedges, and testing that they are all χ-edge-colorable.[1] The coloring process would come wholly after the enumeration process. We, on the other hand, interleave these two processes. As soon as we generate a candidate hypergraph, we color it. If it is not colorable, we have found a counterexample to the conjecture. If it is colorable, we learn a new clause, which we feed back to the hypergraph-generating SAT solver. From this point onwards, any generated hypergraph will not be colorable by this coloring whose clause has been learned. Thus, as soon as we discover a set of colorings that colors all canonical hypergraphs (thanks to SMS, we only generate those), the solver's active formula immediately becomes unsatisfiable. See Figure 2 for an overview of the process.

[1] Unless NP = coNP, non-colorability cannot be encoded in a polynomial-size propositional formula, i.e., we cannot simultaneously encode graph search and non-colorability.

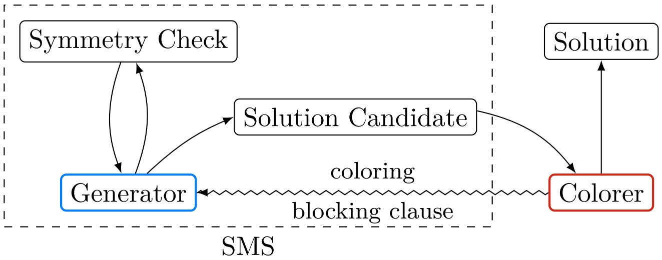

Figure 2 SMS-powered enumeration interleaved with color filtering. One SAT solver generates candidates, the other colors them, and either reports a solution, or returns a coloring clause to be learned by the generator.

We will now explain the details of each of these components: the encoding, the details of hypergraph SMS, and the coloring.

3.1 The Encoding

Our encoding of the existence of a linear hypergraphs is relatively straightforward; the magic happens in SMS and the learning of coloring clauses. We represent a hypergraph with n vertices and m hyperedges by its *incidence graph* and *incidence matrix*. The incidence graph $\mathcal{I}(H)$ of a hypergraph $H = (V, E)$ is the bipartite graph between V and E, with those edges $\{v, e\}$ where $v \in e$. The incidence matrix $\mathcal{M}(H)$ of H is the $n \times m$ binary matrix where $\mathcal{M}(H)_{i,j} = 1 \iff v_i \in e_j$, for some fixed ordering of the vertices and edges $V = \{v_1, \ldots, v_n\}$, $E = \{e_1, \ldots, e_m\}$.

Clearly, any of H, $\mathcal{I}(H)$, $\mathcal{M}(H)$ uniquely determines the other two, and $H \cong H' \iff \mathcal{I}(H) \cong \mathcal{I}(H')$.[2] We will use propositional variables $I_{i,j}$ to represent the incidence matrix, and we will construct incidence graphs modulo isomorphisms. As the canonical incidence graph, we pick the one with the lexicographically least incidence matrix, seen as a string of zeros and ones obtained by concatenating the rows.

We describe the encoding $\mathcal{L}(n, m)$ of the linear hypergraphs with n vertices and m edges in terms of a conjunction of clauses together with cardinality constraints. To encode cardinality constraints, we use sequential counters [24].

- To ensure all hyperedges have size at least 2, we use $\bigwedge_{j=1}^{m} \sum_{i=1}^{n} I_{i,j} \geq 2$.
- Similarly, we can request that the degree of each vertex be at least 2 (vertices of degree 1 contribute nothing to the intersection structure): $\bigwedge_{i=1}^{n} \sum_{j=1}^{m} I_{i,j} \geq 2$.
- To ensure linearity, we will use the auxiliary variables $S_{i,j,k}$ to denote that $v_k \in e_i \cap e_j$:

$$\bigwedge_{1 \leq i < j \leq m} \bigwedge_{k=1}^{n} (\overline{I_{k,i}} \vee \overline{I_{k,j}} \vee S_{i,j,k}) \wedge (\overline{S_{i,j,k}} \vee I_{k,i}) \wedge (\overline{S_{i,j,k}} \vee I_{k,j}) \wedge \bigwedge_{1 \leq k < k' \leq n} \overline{S_{i,j,k}} \vee \overline{S_{i,j,k'}}.$$

Now let us look at the constraints added to $\mathcal{L}(n, m)$ for the EFL Conjecture. We start with restricting the search to covered hypergraphs. To ensure that every pair of vertices is covered, we use the auxiliary variables $O_{k,k',i}$ to denote that $\{v_k, v_{k'}\} \subseteq e_i$:

$$\bigwedge_{1 \leq k < k' \leq n} \left(\bigwedge_{i=1}^{m} (\overline{I_{k,i}} \vee \overline{I_{k',i}} \vee O_{k,k',i}) \wedge (\overline{O_{k,k',i}} \vee I_{k,i}) \wedge (\overline{O_{k,k',i}} \vee I_{k',i}) \right) \wedge \left(\bigvee_{i=1}^{m} O_{k,k',i} \right).$$

[2] For isomorphisms between incidence graphs we require that vertices be mapped to vertices and hyperedges to hyperedges or, what is the same, that the incidence matrix is not transposed by an isomorphism.

For the FB Conjecture we additionally parametrize the formula over the maximum size of the closed neighborhood ζ. Then $\mathcal{L}(n, m)$ is enhanced by the following two constraints:

- We use $O_{i,j}$ for $v_i, v_j \in e$ for some edge e and $\mathcal{N}(k) = \sum_{k' \in [n], k' \neq k} O_{\min(k,k'),\max(k,k')}$. To ensure that the size of the closed neighborhood is limited to ζ we add

$$\bigwedge_{1 \leq k < k' \leq n} \left(\bigwedge_{i \in [m]} (\overline{O_{k,k',i}} \vee O_{k,k'}) \wedge (\overline{O_{k,k'}} \vee \bigvee_{i \in [m]} O_{k,k',i}) \right) \wedge \bigwedge_{k \in [n]} \mathcal{N}(k) < \zeta;$$

- we restrict the search to weakly covered hypergraphs:

$$\bigwedge_{1 \leq k < k' \leq n} (O_{k,k'} \vee \mathcal{N}(k) = \zeta - 1 \vee \mathcal{N}(k') = \zeta - 1)\,.$$

  Note that the complicated-looking expression $\mathcal{N}(k) = \zeta - 1$ is a propositional variable appearing in the encoding of the cardinality constraint.

Last but not least, the variables $S_{i,j}$ to denote that $e_i \cap e_j \neq \emptyset$ are needed to describe the *intersection graph* of the hypergraph whose existence we are encoding. The intersection graph of a hypergraph $H = (V, E)$ is the graph $\mathcal{S}(H) = (E, E^\cap)$, where $E^\cap = \{\{e_1, e_2\} : e_1 \neq e_2 \text{ and } e_1 \cap e_2 \neq \emptyset\}$. We will need this explicit description of the intersection graph to encode a necessary condition for a high chromatic index, and also to learn coloring clauses, as explained in Section 3.3. The variables $S_{i,j}$ can be encoded as follows: $\bigwedge_{1 \leq i < j \leq m} \left(\overline{S_{i,j}} \vee \bigvee_{k=1}^{n} S_{i,j,k}\right) \wedge \left(\bigwedge_{k=1}^{n} S_{i,j} \vee \overline{S_{i,j,k}}\right)$.

3.2 SMS

The encoding presented above is rich in symmetries. Any permutation of the rows and columns of the incidence matrix $\mathcal{M}(H)$ of a hypergraph H leads to an isomorphic hypergraph. Since linearity and non-colorability of the intersection graph are invariants, it is sufficient to keep only one hypergraph per isomorphism class in the search space.

Kirchweger and Szeider [20] introduced SAT modulo Symmetries (SMS) for graphs, a method that supports the search for lexicographically minimal graphs with any additional property encoded with a propositional formula. Later, it was extended to matroids of fixed rank [19]. SMS for graphs produces only lexicographically minimal graphs given by the concatenation of the rows of the adjacency matrix, i.e., those for which no permutation of the set of vertices produces a lexicographically smaller adjacency matrix. The framework checks during the CDCL procedure whether there is a permutation leading to a lexicographically smaller graph; if so, a symmetry-breaking clause is added. This is a problem with two pillars of hardness: during the search, the presence or absence of some edges might be unknown, and for those an exponential number of possibilities must be accounted for; and even if all edges are known, finding a suitable permutation is NP-hard [5]. A procedure called the *minimality check* verifies some necessary criteria for a partially defined graph to be extensible to a minimal graph. The procedure tries to construct a permutation leading to a smaller graph for all extensions of the partially defined graph. If such a permutation exists, we can add a symmetry-breaking clause to trigger a backtrack. The minimality check builds the permutation gradually using a branching algorithm for different choices. Most of the time in practice, the algorithm is fast enough, although sometimes it degrades to exponential behavior. One can limit the total number of branching steps per call of the minimality check. This potentially makes the symmetry breaking incomplete, but does not have an impact on the satisfiability of the formula. Further, it is easy to filter these copies during postprocessing using tools like Nauty [21].

Let us see how we can use SMS for the conjectures on hypergraphs. One possibility to use SMS as is, is by breaking the symmetries over the intersection graph. Preliminary tests showed that this approach does not perform well, chiefly due to the fact that for a candidate intersection graph, the solver must also find the underlying linear hypergraph, and assuming it breaks symmetries on the intersection graph, it cannot break symmetries of the hypergraph anymore. However, we can use this framework to break symmetries also for hypergraphs without (significant) changes.

Let $V = \{v_1, \ldots, v_n\}$ and the number of hyperedges of a hypergraph be fixed by m. A hypergraph H_1 is lexicographically smaller than H_2 (short $H_1 \prec H_2$) if the concatenation of the rows of the incidence matrix $\mathcal{M}(H_1)$ is lexicographically smaller than the concatenation of the rows of $\mathcal{M}(H_2)$.

Since SMS is originally designed for graphs, we apply the symmetry breaking on the incidence graph. Note that the incidence matrix $\mathcal{M}(H)$ coincides with the first n rows and last m columns of the adjacency matrix of the incidence graph $\mathcal{I}(H)$. Let us have a look at an example: let H be a hypergraph with the edges $e_1 = \{v_0, v_1, v_2\}, e_2 = \{v_0, v_3\}, e_3 = \{v_1, v_3\}$, and $e_4 = \{v_2, v_3\}$. We have

$$\mathcal{M}(H) = \begin{array}{c|cccc} & e_1 & e_2 & e_3 & e_4 \\ \hline v_0 & 1 & 1 & 0 & 0 \\ v_1 & 1 & 0 & 1 & 0 \\ v_2 & 1 & 0 & 0 & 1 \\ v_3 & 0 & 1 & 1 & 1 \end{array}, \quad \mathcal{I}(H) = \begin{array}{c|cccccccc} & v_0 & v_1 & v_2 & v_3 & e_1 & e_2 & e_3 & e_4 \\ \hline v_0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ v_1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ v_2 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ v_3 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ e_1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ e_2 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ e_3 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ e_4 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \end{array}.$$

The red box in $\mathcal{I}(H)$ marks the part which is in common with the incidence matrix $\mathcal{M}(H)$.

▶ **Observation 7.** *$\mathcal{M}(H)$ is lexicographically minimal if and only if there is no permutation $\pi : [n+m] \to [n+m]$ with $\pi([n]) = [n]$ such that $\pi(\mathcal{I}(H)) \prec \mathcal{I}(H)$.*

SMS supports not only complete symmetry breaking but also symmetry breaking on a set of given symmetries which can be described by a *generalized ordered partition (GOP)*. Informally, a GOP gives for each vertex a range to which it can be mapped in the permutation. For a formal definition we refer to the original SMS paper [20]. In our case the set of permutations can be described as $[([n], 1, n), ([n+1, n+m], n+1, n+m)]$ which means that all elements in $[n]$ must be mapped by the permutation between 1 and n; all elements in $[n+1, n+m]$ must be mapped between $n+1$ and $n+m$.

This allows us to use SMS as follows. Let $e_{u,v}$ denote the variables representing whether the edge $\{u, v\}$ is present in the bipartite incidence graph. By definition of bipartite graph, we can add the unit clauses $\overline{e_{v,u}}$ for $v, u \in [n]$ and also $\overline{e_{v,u}}$ for $v, u \in \{n+1, \ldots, n+m\}$. Additionally, we set $e_{v,n+i} = I_{v,i}$. All clauses of $\mathcal{L}(n, m)$ can be added unmodified.

Proofs

SAT solvers, in contrast to many other tools for enumeration of combinatorial objects, can produce formal proofs, typically in the DRAT format. These proofs can be checked by external (formally verified) tools [25]. In our workflow, on top of ordinary DRAT-emitting CDCL, we add additional symmetry-breaking and coloring clauses during the search with a custom propagator altering the formula. In previous work [19] the authors suggested a multi-step approach for producing the proof. First, SMS is run until it concludes unsatisfiability.

All learned symmetry-breaking clauses are stored. Then, a second solver without SMS is run on the initial formula augmented with the symmetry-breaking clauses, and produces a proof in the ordinary fashion.

In this work, we use an adaption of the CaDiCaL solver [8] with support for custom propagators and proof logging, including in combination with propagators. The clauses added by the propagator are assumed to be part of the initial formula for the DRAT proof. Proof production in the current CaDiCaL version is only experimental.

The DRAT proof is a certificate for the correctness of the reasoning of the solver, but not for the correctness of the symmetry-breaking clauses. One can either see the symmetry-breaking clauses as part of the initial formula and apply no additional check, or one can check whether the symmetry breaking clauses follow a certain structure, and therefore preserve lexicographically minimal objects, and by extension also satisfiability. Given the permutation used for generating the symmetry-breaking clause, it is easy to check the correctness of the clause [19]. For each symmetry-breaking clause we additionally store the permutation and use a second script for checking the correctness of these clauses.

3.3 Coloring

Suppose the SMS solver has produced a hypergraph $H = (V, E)$ with $V = [n]$. We check whether the hypergraph is χ-edge-colorable with $\chi = n$ for the EFL Conjecture and $\chi = \zeta$ for the FB Conjecture. Observe that whether H is χ-edge-colorable is the same as whether its intersection graph is χ-(vertex-)colorable. We produce another propositional formula, $\mathcal{C}_{H,\chi}$, with the variables $c_{e,i}$ for each hyperedge $e \in E$ and $i \in [\chi]$ to express that e is colored with the color i, and with the following constraints:

- each edge should have some color: $\bigwedge_{e\in E} \bigvee_{i\in[\chi]} c_{e,i}$;
- intersecting edges should have different colors:

$$\bigwedge_{\substack{e_1,e_2\in E\\ e_1\neq e_2\wedge e_1\cap e_2\neq\emptyset}} \bigwedge_{i\in[\chi]} \overline{c_{e_1,i}} \wedge \overline{c_{e_2,i}}.$$

It is not a problem if an edge receives more than one color, as extra colors can always be dropped without violating the constraints.

We enhance this basic encoding by finding and fixing the colors of a clique in the intersection graph as follows. We find the vertex v contained in the largest number of hyperedges $\{e_1, \ldots, e_r\}$ and give these hyperedges the colors $1, \ldots, r$ by adding the unit clauses $c_{e_1,1}, \ldots, c_{e_r,r}$. We further find any other hyperedges $\{e_{r+1}, \ldots, e_{r'}\}$ that intersect with all of these, for a subset-maximal clique in the intersection graph. Since these hyperedges must all have different colors, we can break the symmetries that arise from permutations of color names by fixing the colors.

A second SAT solver then finds a coloring $c : E \to [\chi]$ (a *co-certificate*, because it certifies that the hypergraph is *not* non-χ-edge-colorable [18]), and from this we learn the following *coloring clause*, which we add to the SMS SAT solver:

$$\bigvee_{k=1}^{\chi} \bigvee_{e_i,e_j\in c^{-1}(k)} S_{i,j}.$$

The coloring clause says that at least one pair of hyperedges colored the same by c should intersect, thereby invalidating the coloring c for future hypergraphs.

3.4 Restrictions on the Intersection Graph

We can impose further restrictions on the intersection graph for both conjectures based on the fact that the intersection graph should not be χ-colorable. Any graph that is not χ-colorable must contain some vertex-critical subgraph, i.e., one which is rendered χ-colorable by the removal of any vertex, but which is not χ-colorable itself. It is easy to see that a vertex-critical graph with chromatic number $\chi + 1$ has minimum degree $\geq \chi$, because adding a vertex of degree $< \chi$ to a χ-colorable graph does not break χ-colorability.

We cannot require that the whole intersection graph itself be critical because that is incompatible with (weak) coverage, but we can select a critical subgraph of the intersection graph, using additional variables s_e indicating whether a vertex is part of the subgraph, and restrict the minimum degree for this subgraph accordingly. Without loss of generality, we can additionally request that each vertex (in the intersection graph) representing a hyperedge of size ≥ 3 is part of the critical subgraph and therefore must have degree $\geq \chi$. To see that this last restriction is sound, consider the following process. Given a non-χ-colorable intersection graph $\mathcal{S}(H)$ of some (weakly) covered linear hypergraph H:

1. find a χ-vertex-critical subgraph $G \leq \mathcal{S}(H)$;
2. remove all hyperedges of size ≥ 3 from H which are not contained (as vertices) in G to obtain H';
3. add however many 2-edges are needed to make H' (weakly) covered;
4. observe that $G \leq \mathcal{S}(H')$, and so H' is not χ-edge-colorable.

A hypergraph H whose intersection graph $\mathcal{S}(H)$ contains a subgraph containing all hyperedges of size at least 3 and having minimum degree at least δ is called *δ-reduced*. It follows that if a counterexample with n vertices and m hyperedges to either conjecture exists, a $(\chi - 1)$-reduced (weakly) covered counterexample also exists, with the same number of vertices n, and a number of hyperedges $m' \geq m$.

3.5 Conjectures

Let us summarize the encoding and how the side constraints relate to the original Conjectures 1 and 3. We will formulate modified versions of the conjectures that match our encoding and precisely state their relationship to the original conjectures.

▶ **Conjecture 8** (EFL' Conjecture). *Every $(n-1)$-reduced linear space with n vertices and m hyperedges is n-edge-colorable.*

▶ **Proposition 9.** *If there is a counterexample to the EFL Conjecture with n vertices and m hyperedges of size at least 2, then there is a counterexample to EFL' Conjecture with n vertices and $m' \geq m$ hyperedges.*

▶ **Corollary 10.** *The EFL Conjecture and the EFL' Conjecture are equivalent.*

▶ **Conjecture 11** (FB' Conjecture). *Every $(\zeta - 1)$-reduced weakly covered linear hypergraph with n vertices, m hyperedges of size at least 2, and maximum closed neighborhood size ζ is ζ-edge-colorable.*

▶ **Proposition 12.** *If there is a counterexample to FB Conjecture with n vertices and m hyperedges of size at least 2, then there is a counterexample to the FB' Conjecture with n vertices and $m' \geq m$ hyperedges.*

▶ **Corollary 13.** *The FB Conjecture and the FB' Conjecture are equivalent.*

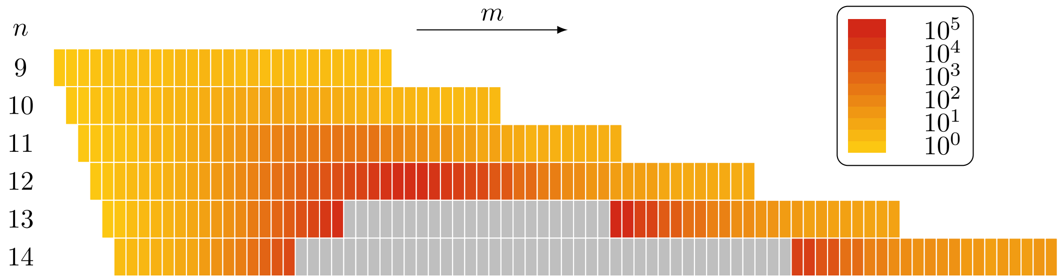

Figure 3 For selected n, the regions of m for which we verified the EFL' conjecture, within 3 days of CPU time. For each n, m runs from n to $\binom{n}{2}$. The color transitions from yellow to red in proportion to the logarithm of the running time in seconds, red means higher. White regions on the left are trivial, on the right impossible for linear spaces. Regions we could not solve are in gray.

For given n and m, the individual cases of the EFL' and FB' conjectures are, in some sense, the canonical or essential cases of EFL and FB, respectively, even though the correspondence between the (n, m)-buckets between the two versions of each conjecture is not one-to-one (and is made precise by Propositions 9 and 12).

4 Computational Results

In this section we will present the results of our computations. We performed three different experiments, each on a cluster of machines with different processors[3], running Ubuntu 18.04 on Linux 4.15. The source code and scripts for reproducing the results are available as part of the SMS package at `https://github.com/markirch/sat-modulo-symmetries`.

The first experiment was targeted at verifying Conjecture 8, and by extension the EFL conjecture itself. We managed to verify, this time with formal methods capable of proof logging, the previous known results, and we verified the conjecture in additional cases. Selected results of this experiment are summarized in Figure 3. The colored cells indicate which values of n and m we could solve, within a time limit of three days. Our results in every case agree with the prediction of Erdős, Faber, and Lovász: all linear hypergraphs are n-colorable. In addition to those shown in Figure 3, we verified all cases listed in the following theorem.

▶ **Theorem 14.** *The EFL' Conjecture (Conjecture 8) holds for n, m if $n \leq 12$; or*

- $n = 13$ *and* $m \in [13, 32] \cup [55, 78]$*; or*
- $n = 14$ *and* $m \in [14, 28] \cup [70, 91]$*; or*
- $n = 15$ *and* $m \in [15, 29] \cup [84, 105]$*; or*
- $n = 16$ *and* $m \in [16, 30] \cup \{99\} \cup [101, 120]$*; or*
- $n = 17$ *and* $m \in [17, 30] \cup [117, 136]$*; or*
- $n = 18$ *and* $m \in [18, 31] \cup [134, 153]$.

The longest successfully terminated case that we encountered was that of $n = 15$ and $m = 85$, which took two days, 15 hours, and 26 minutes to solve, of which one day and 15 and a half hours were spent in the coloring solver.

[3] Intel Xeon {E5540, E5649, E5-2630 v2, E5-2640 v4}@ at most 2.60 GHz, AMD EPYC 7402@2.80GHz

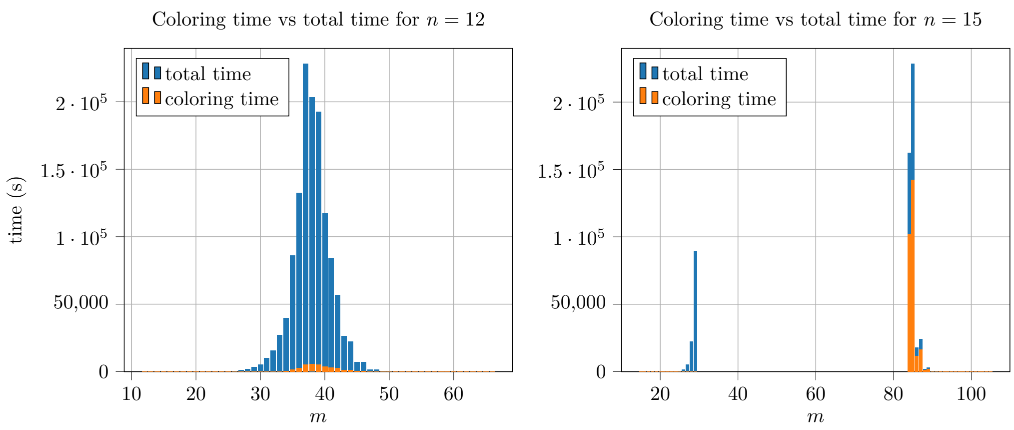

Figure 4 The distribution of the running times for $n = 12$ and $n = 15$.

In general, we observed a pattern where hardness of the individual instances culminated around the middle of the region of the available values $\left[n, \binom{n}{2}\right]$, and fell off roughly symmetrically on both sides. However, the nature of the hardness differed on the two sides: for smaller values of m, coloring was negligibly easy (less than 0.1% of the total time) and most of the time was spent in hypergraph search, while for higher values of m, coloring took up a much more significant portion of the total time; this proportion peaked for $n = 15$ and $m = 87$, at just over 67%. Selected cases are depicted in Figure 4.

For smaller values of n we can exhaustively enumerate all $(n-1)$-reduced linear spaces and compare their number with the number of learned colorings, to get an idea of the speedup. We tested this for $10 \leq n \leq 12$ and found that the total number of colorings is about 10% smaller than the number of hypergraphs. This is somewhat in contrast with much higher speedups reported in related work [18], and it would be interesting future work to investigate why our case behaves differently. One hypothesis that we have is that in our case, the intersection graphs are not canonical, and hence the colorings do not work as well. In fact, we could even see isomorphic copies of intersection graphs, which require different colorings. This is testable by setting up an appropriate experiment, but is somewhat beyond the scope of this paper.

The aim of the second experiment was to gain insights about the extremal examples to the EFL conjecture. The second experiment is similar to the first, except that

- instead of asking for hypergraphs that cannot be edge-colored with n colors (and getting none), we asked for hypergraphs that cannot be edge-colored with $n-1$ colors;
- and we dropped the constraint that the hypergraphs should be $(n-1)$-reduced.

Thus, we have enumerated all extremal linear *spaces* with respect to the EFL conjecture.

The known extremal examples are the degenerate plane (for any $n \geq 3$), the complete graph (for odd $n \geq 3$), and the projective plane, where $n = k^2 + k + 1$ for many k including all prime powers. According to Kang et al. [17], in addition to the odd clique, also "minor modifications thereof" work. Kahn [15] is similarly cryptic: "*[The EFL Conjecture] is sharp when H is a projective plane or complete graph K_n with n odd, and also in a few related cases, but there ought to be some slack in the bound away from these extremes.*" We enumerated all extremal examples with $n \leq 12$ and in some cases of $n = 13$, and can confirm that these "minor modifications," listed in Theorem 15, are the only other extremal examples.

Let $\mathcal{H}_{n,k}$ be the linear space on n vertices that consists solely of 2-edges and one k-edge, $\mathcal{H}_{n,3\not\cap 3}$ the linear space with 2-edges and two disjoint 3-edges, and $\mathcal{H}_{n,3\cap 3}$ the space with 2-edges and two intersecting 3-edges.

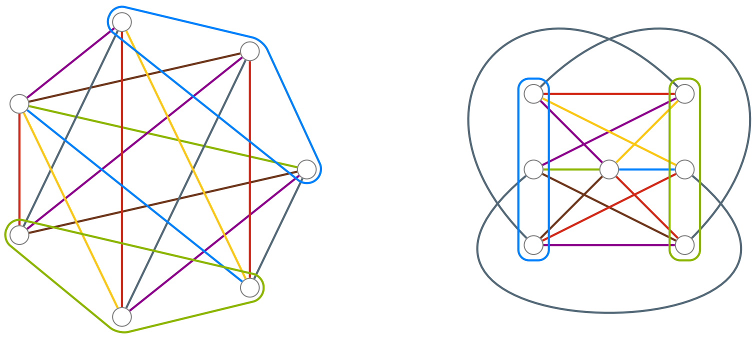

Figure 5 The extremal hypergraph $\mathcal{H}_{7,3\not\cap 3}$, on the left drawn analogously to K_7 from Figure 1, on the right in a way that highlights symmetries.

▶ **Theorem 15.** *The linear spaces with $n \leq 12$ vertices and chromatic index n are precisely $\mathcal{H}_{n,k}$ for all $k \not\equiv n \mod 2$, and additionally $\mathcal{H}_{7,3}, \mathcal{H}_{9,3}, \mathcal{H}_{11,3}$, $\mathcal{H}_{7,3\not\cap 3}$, $\mathcal{H}_{11,3\cap 3}$, $\mathcal{H}_{11,3\not\cap 3}$, and the Fano plane.*

We proved Theorem 15 almost entirely computationally, with the only exception of $\mathcal{H}_{11,3\cap 3}$ and $\mathcal{H}_{11,3\not\cap 3}$, which were too hard to prove non-colorable. In Section 4.1, we provide the missing manual proofs of extremality and generalize some patterns from Theorem 15. Theorem 15 remains necessary to show that the enumerated extremal examples are complete.

The hypergraphs $\mathcal{H}_{n,k}$ can be seen as a chain of steps of size 2 linking the complete graph to the degenerate plane. Even though the complete graph is not extremal for even n, this chain still exists, only shifted to odd k. In this case the chain can even be considered somewhat "purer," as there is neither an "intermediate link" of the wrong parity, as there is $\mathcal{H}_{n,3}$ for odd n, nor other exceptions as listed in Theorem 15; and moreover there are no projective planes for even n ($k^2 + k + 1$ is always odd). One is almost tempted to conjecture that these are the only extremal examples when n is even.

The enumeration of extremal examples is computationally harder than the verification of the EFL conjecture, for two reasons. The first is that the coloring clauses are weaker, and as such more hypergraphs must be searched. The second is that while in the case of the EFL conjecture all graphs are colorable (all colorability queries satisfiable), in this case non-colorability sometimes has to be proved. Even though we break color symmetries by assigning arbitrary colors to one clique in the intersection graph, frequently a clique of the maximum possible size $n - 1$, a large part of the intersection graph remains to be colored, often with other large cliques. Since the formula encoding the existence of a (vertex) coloring of a clique is nothing else than the famous pigeonhole principle formula [11], known to be hard for resolution and CDCL solvers, we can expect proving non-colorability to be hard. This is indeed what we observe: for $n = 11$ we could not completely solve the case $m = 51$ ($\mathcal{H}_{11,3\cap 3}$ and $\mathcal{H}_{11,3\not\cap 3}$), and for $n = 13$ there were many cases that were easy in the first experiment, but which we could not solve now; in fact, we could not even prove (with SAT) that K_{13} is not 12-edge-colorable. Since in these cases we can easily enumerate all linear spaces, the reason for hardness must be unsatisfiable colorability queries. An interesting avenue for future research would be to improve performance on these hard coloring instances.

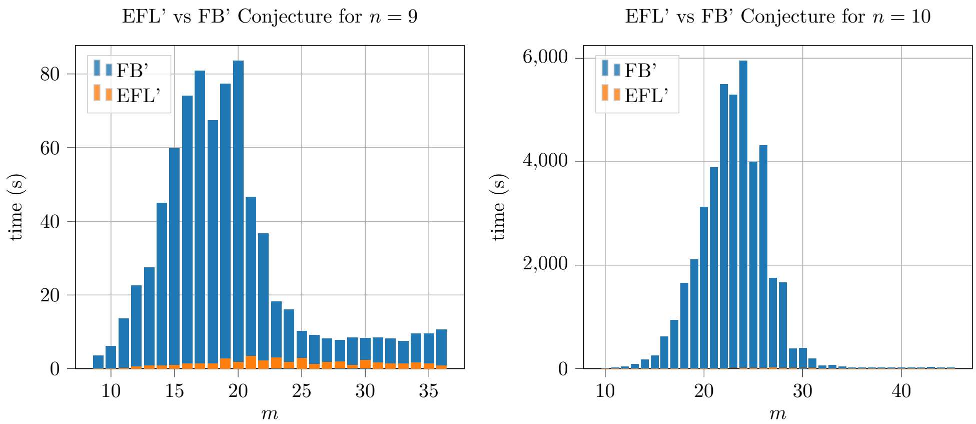

Figure 6 The distribution of the running times for the EFL' and FB' conjectures for $n \in [9, 10]$.

This experiment also uncovered that the minimality check sometimes struggles and runs out of the limit (see Section 3.2), and duplicate solutions are produced. In this case, this is not a problem in and of itself, mainly because the copies are very obviously isomorphic, but it shows that these highly symmetric hypergraphs pose a challenge for current SMS, and could serve as benchmarks for future development.

In the third experiment we tackled Conjecture 11, and through it Conjecture 3. The setup is almost identical to the first experiment, except that we use the other version of the encoding described in Section 3.1. We summarize our results in another theorem.

▶ **Theorem 16.** *The FB' Conjecture holds for n and m if $n \leq 10$; or*
- $n = 11$ *and* $m \in [11, 20] \cup [37, 55]$*; or*
- $n = 12$ *and* $m \in [12, 17] \cup [49, 66]$*.*

▶ **Corollary 17.** *The FB Conjecture holds for $n \leq 10$.*

Verifying the FB Conjecture is much harder than verifying the EFL conjecture; for example for $n = 11$ and $m = 35$, Conjecture 8 was solved in 12 minutes, while the same case took over a day for Conjecture 11. The running times for $n = 9, 10$ are shown in Figure 6.

4.1 Extremality Proofs

The main results of this section are Theorems 18, 22, 24, and 27. Theorem 18 says that a linear space with an odd number of vertices that does not have enough or large enough hyperedges of size ≥ 3 is extremal. Theorem 22 establishes extremality of $\mathcal{H}_{n,k}$ when n and k have opposite parity. The other two theorems fill some of the gaps left by the first two. Theorem 24 gives a sufficient condition for a space of odd size to be non-extremal, and Theorem 27 shows that spaces with an even number of vertices and only even hyperedges are non-extremal provided that large hyperedges do not intersect.

We say a color c *sees* a vertex v (and vice versa) if a c-colored hyperedge contains v.

▶ **Theorem 18.** *Let H be a linear space with n vertices, n odd, whose hyperedges of size ≥ 3 have sizes $a_1, \ldots, a_\mu$. If $n \geq \sum_{i=1}^{\mu} a_i(a_i - 2) + \mu + 2$, then H is not $(n-1)$-edge-colorable.*

Proof. For a vertex of degree d, there are $n-1-d$ colors that do not see it. We have $\deg(v) = n - 1 - \sum_{v\in e,|e|\geq 3}(|e|-2)$, so at most

$$\sum_{v\in V}\left(\sum_{v\in e,|e|\geq 3}(|e|-2)\right) = \sum_{i=1}^{\mu} a_i(a_i-2)$$

colors do not see every vertex. There are $\leq \mu$ colors used for the hyperedges of size ≥ 3, so at least one color sees every vertex and is used only on 2-edges; a contradiction for odd n. ◀

▶ **Corollary 19.** *$\mathcal{H}_{n,3\not\cap 3}$ and $\mathcal{H}_{n,3\cap 3}$ are not $(n-1)$-colorable when $n \geq 11$ is odd.*

▶ **Theorem 20.** *$\mathcal{H}_{7,3\not\cap 3}$ is not 6-edge-colorable.*

Proof. Let A and B be the 3-edges, and v the other vertex. A and B cannot be colored the same because v needs to see all colors, so the nine 2-edges between A and B use 4 colors, and some c is used ≥ 3 times. But then each vertex of A and B sees c, and so v cannot see c. ◀

▶ **Theorem 21.** *$\mathcal{H}_{9,3\not\cap 3}$ is 8-edge-colorable.*

Proof. $\{u,v\} \mapsto u+v \mod 8$ if $u,v < 8$, $\{u,8\} \mapsto 2u$ if $u \leq 3$, else $\{u,8\} \mapsto 2u-1 \mod 8$, $\{3,4,5\} \mapsto 0$, $\{6,7,8\} \mapsto 5$. ◀

▶ **Theorem 22.** *Let $n \not\equiv k \mod 2$, $k \in [2,n]$. Then $\mathcal{H}_{n,k}$ is not $(n-1)$-edge-colorable.*

Proof. Let the k-edge have the color c. All other vertices must see every color, in particular c. The color c forms a perfect matching with $1+(n-k)/2$ edges; but $n-k$ is odd. ◀

▶ **Corollary 23.** *$\mathcal{H}_{n,3}$ is not $(n-1)$-edge-colorable when $n = 4$ or $n \geq 6$.*

Proof. For even n this follows from Theorem 22, for odd n from Theorem 18. ◀

▶ **Theorem 24.** *Let H be a linear space with n vertices, n odd, and assume that H contains a vertex v^* that belongs to every hyperedge of size greater than two. If the number of even hyperedges (including 2-edges) containing v^* is $\leq \frac{n-1}{2}$, then H is $(n-1)$-edge-colorable.*

Proof. Rename vertices so that $v^* = n-1$, the even hyperedges containing v^* are $A_0,\dots,A_r$, with $i \in A_i$, and the odd hyperedges are $B_1,\dots,B_s$, and they are ordered as follows:

$$0,1,\dots,r,(A_0\setminus\{0,n-1\}),\dots,(A_r\setminus\{r,n-1\}),(B_1\setminus\{n-1\}),\dots,(B_s\setminus\{n-1\}),n-1$$

Color as follows: $u,v < n-1 \implies \{u,v\} \mapsto u+v \mod n-1$, if $u \leq \frac{n-1}{2}$ and $\{u,n-1\} \subseteq E$, then $E \mapsto 2u$; all thus unassigned hyperedges must be B_i, and we color them $B_i \mapsto \min(B_i) + \max(B_i\setminus\{n-1\}) \mod n-1$. The coloring is proper because (1) the doubling color is applied to at most $\frac{n-1}{2}$ edges; (2) $\min(B_i)$ and $\max(B_i\setminus\{n-1\})$ always have opposite parity, and hence do not conflict with the doubled colors, which are even; and (3) two colors assigned by the last rule cannot be the same because $\min(B_i) \geq \frac{n+1}{2}$ $\left(\frac{n-1}{2} < a < b < c < d < n-1 \implies a+b \not\equiv c+d \mod n-1\right)$. ◀

▶ **Corollary 25.** *Let $n \leq 2k-1$, $n \equiv k \equiv 1 \mod 2$. Then $\mathcal{H}_{n,k}$ is $(n-1)$-edge-colorable.*

▶ **Corollary 26.** *$\mathcal{H}_{5,3}$ is 4-, $\mathcal{H}_{7,3\cap 3}$ is 6-, and $\mathcal{H}_{9,3\cap 3}$ is 8-edge-colorable.*

▶ **Theorem 27.** *Let H be a hypergraph with n vertices, n even. If all hyperedges of H are even, and all those of size greater than 2 are pairwise disjoint, then H is $(n-1)$-edge-colorable.*

Proof. Let the hyperedges of size greater than two be $L_1,\dots,L_r$, let $L := \bigcup_{i=1}^{r} L_i$. Rename the vertices so that all L_i are closed under additive inverse: $x \in L_i \implies n-1-x \in L_i$. Color $L_i \mapsto 0$, $u,v < n-1 \implies \{u,v\} \mapsto u+v \mod n-1$, $\{u,n-1\} \mapsto 2u \mod n-1$. ◀

▶ **Corollary 28.** *Let $n \equiv k \equiv 0 \mod 2$, $k \in [2,n]$. Then $\mathcal{H}_{n,k}$ is $(n-1)$-edge-colorable.*

5 Conclusion

In this paper, we used SAT solvers to partially confirm two conjectures about hypergraph edge colorings for small instances. We achieved this by interleaving the two SAT solvers: one for symmetry-breaking search for hypergraphs, the other for finding colorings of candidate hypergraphs, to prune the search space early.

Our experiments show that the hypergraph search and edge-coloring problem is a challenging benchmark for all components of our framework. To improve the minimality check in SMS, we plan to combine dynamic symmetry breaking with (partial) static symmetry breaking, for example by way of sorting the vertices by degree (by a static, poly-size encoding), and limit the dynamic symmetry breaking to vertices of the same degree. Similarly, it would be very interesting to improve CDCL performance on these hard coloring problems. Here, it seems, symmetry-breaking could help, but we have a chicken-and-egg problem: if symmetry-breaking is sometimes hard for the minimality check, will it be feasible for the coloring part? In any case, hypergraphs and colorings are very general and expressive combinatorial objects, and as such are worthy of the attention to improve the performance of general-purpose SAT solvers.

References

1 Claude Berge. On the chromatic index of a linear hypergraph and the Chvátal conjecture. *Annals of the New York Academy of Sciences*, 555(1):40–44, 1989.

2 Anton Betten and Dieter Betten. Linear spaces with at most 12 points. *Journal of Combinatorial Designs*, 7(2):119–145, 1999.

3 William I. Chang and Eugene L. Lawler. Edge coloring of hypergraphs and a conjecture of Erdős, Faber, Lovász. *Combinatorica*, 8(3):293–295, 1988.

4 Michael Codish, Alice Miller, Patrick Prosser, and Peter J. Stuckey. Constraints for symmetry breaking in graph representation. *Constraints*, 24(1):1–24, 2019. `doi:10.1007/s10601-018-9294-5`.

5 James M. Crawford, Matthew L. Ginsberg, Eugene M. Luks, and Amitabha Roy. Symmetry-breaking predicates for search problems. In Stuart C. Shapiro Luigia Carlucci Aiello, Jon Doyle, editor, *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR'96), Cambridge, Massachusetts, USA, November 5-8, 1996*, pages 148–159. Morgan Kaufmann, 1996.

6 Paul Erdős. Problems and results in graph theory and combinatorial analysis. *Proc. British Combinatorial Conj., 5th*, pages 169–192, 1975.

7 Pául Erdős. On the combinatorial problems which I would most like to see solved. *Combinatorica*, 1(1):25–42, 1981.

8 Katalin Fazekas, Aina Niemetz, Mathias Preiner, Markus Kirchweger, Stefan Szeider, and Armin Biere. IPASIR-UP: User propagators for CDCL. In Meena Mahajan and Friedrich Slivovsky, editors, *The 26th International Conference on Theory and Applications of Satisfiability Testing (SAT 2023), July 04–08, 2023, Alghero, Italy*, LIPIcs, pages 8:1–8:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. `doi:10.4230/LIPIcs.SAT.2023.8`.

9 Johannes K. Fichte, Daniel Le Berre, Markus Hecher, and Stefan Szeider. The silent (r)evolution of SAT. *Communications of the ACM*, 66(6):64–72, June 2023. `doi:10.1145/3560469`.

10 Zoltán Füredi. The chromatic index of simple hypergraphs. *Graphs and Combinatorics*, 2(1):89–92, 1986.

11 Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–308, 1985. `doi:10.1016/0304-3975(85)90144-6`.

12 Neil Hindman. On a conjecture of Erdős, Faber, and Lovász about n-colorings. *Canadian Journal of Mathematics*, 33(3):563–570, 1981. `doi:10.4153/CJM-1981-046-9`.

13 Jeff Kahn. Coloring nearly-disjoint hypergraphs with $n + o(n)$ colors. *Journal of combinatorial theory, Series A*, 59(1):31–39, 1992.

14 Jeff Kahn. Asymptotics of hypergraph matching, covering and coloring problems. In *Proceedings of the International Congress of Mathematicians: August 3–11, 1994 Zürich, Switzerland*, pages 1353–1362. Springer, 1995.

15 Jeff Kahn. On some hypergraph problems of Paul Erdős and the asymptotics of matchings, covers and colorings. *The Mathematics of Paul Erdős I*, pages 345–371, 1997.

16 Jeff Kahn and Paul D. Seymour. A fractional version of the Erdős-Faber-Lovász conjecture. *Combinatorica*, 12(2):155–160, 1992.

17 Dong Yeap Kang, Tom Kelly, Daniela Kühn, Abhishek Methuku, and Deryk Osthus. A proof of the Erdős-Faber-Lovász conjecture, 2021. doi:10.48550/arXiv.2101.04698.

18 Markus Kirchweger, Tomáš Peitl, and Stefan Szeider. Co-certificate learning with SAT modulo symmetries. In *Proceedings of the 34th International Joint Conference on Artificial Intelligence, IJCAI 2023*. AAAI Press/IJCAI, 2023. To appear.

19 Markus Kirchweger, Manfred Scheucher, and Stefan Szeider. A SAT attack on Rota's Basis Conjecture. In *Theory and Applications of Satisfiability Testing – SAT 2022 – 25th International Conference, Haifa, Israel, August 2-5, 2022, Proceedings*, 2022. doi:10.4230/LIPIcs.SAT.2022.4.

20 Markus Kirchweger and Stefan Szeider. SAT modulo symmetries for graph generation. In *27th International Conference on Principles and Practice of Constraint Programming (CP 2021)*, LIPIcs, pages 39:1–39:17. Dagstuhl, 2021. doi:10.4230/LIPIcs.CP.2021.34.

21 Brendan D. McKay and Adolfo Piperno. Practical graph isomorphism, II. *J. Symbolic Comput.*, 60:94–112, 2014. doi:10.1016/j.jsc.2013.09.003.

22 David Romero and Federico Alonso-Pecina. The Erdős–Faber–Lovász conjecture is true for $n \leq 12$. *Discrete Mathematics, Algorithms and Applications*, 06(03):1450039, 2014. doi:10.1142/S1793830914500396.

23 Paul D. Seymour. Packing nearly-disjoint sets. *Combinatorica*, 2:91–97, 1982.

24 Carsten Sinz. Towards an optimal CNF encoding of Boolean cardinality constraints. In Peter van Beek, editor, *Principles and Practice of Constraint Programming – CP 2005, 11th International Conference, CP 2005, Sitges, Spain, October 1-5, 2005, Proceedings*, volume 3709 of *Lecture Notes in Computer Science*, pages 827–831. Springer Verlag, 2005. doi:10.1007/11564751_73.

25 Nathan Wetzler, Marijn J. H. Heule, and Warren A. Hunt. DRAT-trim: Efficient checking and trimming using expressive clausal proofs. In *Theory and Applications of Satisfiability Testing – SAT 2014*, volume 8561 of *Lecture Notes in Computer Science*, pages 422–429. Springer Verlag, 2014. doi:10.1007/978-3-319-09284-3_31.