



Are hitting formulas hard for resolution?

Tomáš Peitl^{*}, Stefan Szeider

TU Wien, Favoritenstraße 9-11, 1040, Vienna, Austria



ARTICLE INFO

Article history:

Received 1 July 2022

Received in revised form 21 April 2023

Accepted 3 May 2023

Available online 20 May 2023

Keywords:

Hitting formulas

Resolution

Satisfiability

SAT

Symmetry breaking

ABSTRACT

Hitting formulas, introduced by Iwama, are an unusual class of propositional CNF formulas. Not only is their satisfiability decidable in polynomial time, but even their models can be counted in closed form. This stands in stark contrast with other polynomial-time decidable classes, which usually have algorithms based on backtracking and resolution and for which model counting remains hard, like 2-SAT and Horn-SAT. However, those resolution-based algorithms usually easily imply an upper bound on resolution complexity, which is missing for hitting formulas. Are hitting formulas hard for resolution?

In this paper we take the first steps towards answering this question. We show that the resolution complexity of hitting formulas is dominated by so-called irreducible hitting formulas, first studied by Kullmann and Zhao, that cannot be composed of smaller hitting formulas. However, by definition, large irreducible unsatisfiable hitting formulas are difficult to construct; it is not even known whether infinitely many exist. Building upon our theoretical results, we implement an efficient algorithm on top of the Nauty software package to enumerate all irreducible unsatisfiable hitting formulas with up to 14 clauses. We also determine the exact resolution complexity of the generated hitting formulas with up to 13 clauses by extending a known SAT encoding for our purposes. Our experimental results suggest that hitting formulas are indeed hard for resolution.

© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Hitting formulas (also known as hitting clause-sets) are propositional formulas in conjunctive normal form (CNF) with the property that each pair of clauses clashes, meaning there is a variable that occurs in one clause and its negation in the other (see Fig. 1). Introduced by Iwama [10] in 1989 and given the name ‘hitting’ by Kleine Büning and Zhao [12], they are notorious for Iwama’s stunningly elegant argument (reproduced here as Theorem 1), which shows that their satisfiability can be decided in polynomial time—in fact, the number of models can be obtained in closed form. This is a rare encounter in the bay of tractability: most other polynomial-time decidable subclasses of SAT (‘islands of tractability’) have algorithms with more typical ingredients like unit propagation, (bounded-depth) backtracking, and resolution [6]. Such algorithms usually do not scale to the case of model counting, which remains #P-hard for 2-SAT and Horn-SAT, two of the most prominent tractable cases of SAT [26]. Due to their remarkable properties, hitting formulas have been the subject of many theoretical investigations [7,12,14–17,19,22,28,31] and served as a building block of larger classes of tractable formulas via backdoors [21] and treewidth [8].

But one fundamental question remains unsettled: while polynomial-time resolution-like algorithms usually provide a corresponding polynomial upper bound on the resolution complexity (length of a shortest resolution refutation) of formulas

^{*} Corresponding author.

E-mail addresses: peitl@ac.tuwien.ac.at (T. Peitl), sz@ac.tuwien.ac.at (S. Szeider).

$$\{x, \bar{y}\}, \{y, \bar{z}\}, \{z, \bar{x}\}, \{x, y, z\}, \{\bar{x}, \bar{y}, \bar{z}\}$$

Fig. 1. An example hitting formula written as a set of clauses, each itself a set of literals.

Table 1

Values of $h(\text{IUH}(n, m))$, i.e., the lengths of the longest shortest refutation required by a (regular) irreducible unsatisfiable hitting formula with n variables and m clauses. The subscript gives the number of formulas that attain maximum hardness, the superscript is the total number of formulas in each category. All counts are modulo isomorphisms. Empty areas contain no IUHs.

$n \setminus m$	1	2	3	4	5	6	7	8	9	10	11	12	13	14
0	1_1^1													
1														
2														
3					10_1^1									
4							14_2^2	19_1^2	19_1^1					
5									19_7^{15}	22_9^{47}	25_{15}^{138}	28_2^{245}	30_{11}^{511}	$\geq 33_{\leq 3}^{804}$
6											24_6^{112}	27_{10}^{618}	30_{34}^{5178}	$\geq 34_{\leq 29}^{25235}$
7													29_1^{1019}	31_{608}^{7765}
8														
9														
10														

they target, it is far from clear how one could get such a bound from the counting argument for hitting formulas. Resolution complexity has established itself as an important measure due to its relation to the running time of conflict-driven clause-learning (CDCL [27]) SAT solvers—resolution complexity gives a lower bound for CDCL running time and, under some assumptions also an upper bound (although these assumptions include non-determinism) [2,3,25]. Even though the resolution complexity of hitting formulas might not have immediate practical consequences, it is still a natural question to ask for a famous class of formulas and a fundamental complexity measure. So, are hitting formulas hard for resolution? On the one hand, they are easy even for model counting; on the other hand, the reason why they are easy to solve is a global counting argument, the kind which is notoriously hard for resolution (as witnessed, for instance, by the Pigeonhole Principle [9] or Tseitin Formulas [30]). With this paper, we attempt to shed some light on this question.

One reason why we know so little about the resolution complexity of hitting formulas is that we have no good methods of constructing large hard instances. We can construct large hitting formulas by carefully gluing together smaller hitting formulas, but that only produces *reducible* formulas, a notion recently proposed by Kullmann and Zhao [19]. In contrast, we observe (Corollary 5) that any sequence of hitting formulas must implicitly contain arbitrarily large *irreducible* formulas to have a chance at being hard for resolution. However, large irreducible formulas are non-trivial to construct because they cannot be assembled from smaller ones by definition, at least not by known constructions.

That leads us to another straightforward yet non-trivial question: do there even exist infinitely many irreducible unsatisfiable hitting formulas? We hint that this is indeed the case by enumerating all irreducible unsatisfiable hitting formulas on up to 14 clauses (and some larger) modulo isomorphisms and observing that their numbers grow fast. Kullmann and Zhao proved [19, Lemma 39], perhaps surprisingly, that all irreducible unsatisfiable hitting formulas other than $x \wedge \neg x$ must be *regular*—meaning each literal must occur at least twice (also known as *nonsingular*, though we prefer the more positive term). Another surprise we discovered is that there is no irreducible unsatisfiable hitting formula with 4 or 6 clauses (and none with 2 or 3, but that is an artifact of regularity; see Table 1).

Now that we have the formulas, we might as well compute their shortest refutations. Armed with recent developments in both the theory of hitting formulas [19] as well as in practical computation of shortest resolution refutations [24], we can compute shortest refutations of all irreducible hitting formulas on up to 13 clauses (and some larger) and all regular hitting formulas on up to 11 clauses. The picture that this unveils indicates that hitting formulas are perhaps not as hard as the hardest formulas in general (by comparing with the *resolution hardness numbers* [24]), but not significantly easier either.

We need some theoretical advances to narrow our search space and compute shortest refutations faster. Following the lead of Kullmann and Zhao [18,19], we focus on the aforementioned concept of *irreducibility*, and we observe that a new property we call *strong irreducibility* is also essential. We prove irreducibility and strong irreducibility coincide for hitting formulas, allowing us to take advantage of the stronger variant while only checking the weaker one.

Our contributions are thus twofold. First, we contribute theoretical knowledge about irreducibility and regularity in (not only) hitting formulas, with a focus on connections to resolution complexity. Second, we extensively evaluate the resolution complexity of hitting formulas using the new theoretical results. In the experimental part, we rely on two key tools: the ability to generate formulas modulo isomorphisms, provided by a customized version of Nauty [20], and an efficient way to compute shortest refutations of generated formulas. For the latter, we use a recently proposed SAT encoding [24], but we enhance it with further symmetry breaking and constraints that exploit irreducibility.

Our findings indicate that small hitting formulas are almost as hard for resolution as the hardest formulas with the same number of clauses, suggesting that hitting formulas in general might not admit polynomial-size resolution refutations.

2. Preliminaries

We assume familiarity with standard notions of graph theory, including those of (un)directed graphs, acyclicity, and in- and out-degree of a vertex—we refer to a standard graph theory handbook [4]. All graphs considered, directed or undirected, do not contain any self-loops or parallel edges.

2.1. Formulas

We consider propositional formulas in conjunctive normal form (CNF) represented as sets of clauses. We assume an infinite set \mathcal{V} of (propositional) variables. A literal ℓ is a variable x or a negated variable $\neg x$; we write $\text{lit} := \{x, \neg x \mid x \in \mathcal{V}\}$. For a literal ℓ we write $\bar{\ell} := \neg x$ if $\ell = x$, and $\bar{\ell} := x$ if $\ell = \neg x$. For a set C of literals, we define $\bar{C} := \{\bar{\ell} \mid \ell \in C\}$; we say C is tautological if $C \cap \bar{C} \neq \emptyset$. A finite non-tautological set of literals is a clause; a finite set of clauses is a (CNF) formula. The empty clause is denoted by \perp . We say two clauses C and D clash if $C \cap \bar{D} \neq \emptyset$; the literals in $C \cap \bar{D}$ and $\bar{C} \cap D$ are the clashing literals, and their variables are the clashing variables. For a clause C , we put $\text{var}(C) = \{\text{var}(\ell) \mid \ell \in C\}$, and for a formula F , $\text{var}(F) = \bigcup_{C \in F} \text{var}(C)$. Similarly, we put $\text{lit}(F) := \text{var}(F) \cup \bar{\text{var}}(F)$.

An (partial) assignment is a mapping $V \rightarrow \{0, 1\}$ for a subset of variables $V \subseteq \text{var}(F)$, extended to literals as $\tau(\bar{x}) = 1 - \tau(x)$. An assignment is called total if its domain is $\text{var}(F)$. We identify an assignment with the set of literals $\tau^{-1}(1) \cup \tau^{-1}(0)$, i.e., with the set of all literals set to 1 by τ . For a clause C and a set of literals τ (which could be an assignment) we define the restriction of C by τ , written $C[\tau]$, as \top if $C \cap \tau \neq \emptyset$ (in which case we say τ satisfies C), and $C \setminus \tau$ otherwise. If $C[\tau] = \perp$ for a clause C and total assignment τ , we say C covers τ . The restriction of a formula F by an assignment τ , written $F[\tau]$, is defined as $\{C[\tau] \mid C \in F\} \setminus \{\top\}$. A formula F is satisfiable if there is a satisfying assignment (also model), i.e., a total assignment τ with $F[\tau] = \emptyset$; otherwise it is unsatisfiable. For two formulas F and G , we say F entails G , denoted by $F \models G$, if every model of F is a model of G . We say F and G are (logically) equivalent, denoted $F \equiv G$, if $F \models G$ and $G \models F$. We write $\#F$ for the number of models of F .

2.2. Resolution refutations

If $C_1 \cap \bar{C}_2 = \{\ell\}$ for clauses C_1, C_2 and a literal ℓ , then the resolution rule allows the derivation of the clause $D = (C_1 \cup C_2) \setminus \{\ell, \bar{\ell}\}$; D is the resolvent of the premises C_1 and C_2 , and we say that D is obtained by resolving on ℓ . For a formula F , a sequence $P = C_1, \dots, C_s$ of clauses is a resolution derivation of C_s from F if for each $i \in \{1, \dots, s\}$ at least one of the following holds.

1. $C_i \in F$ (“ L_i is an axiom”);
2. C_i is the resolvent of L_j and $L_{j'}$ for some $1 \leq j < j' < i$ (“ C_i is obtained by resolution”).

We write $|P| := s$ and call s the length of P . If L_s is the empty clause, then P is a resolution refutation of F . A clause C_i in a resolution derivation may have different possible ‘histories.’ C_i may be the resolvent of more than one pair of clauses preceding C_i , or C_i may be both an axiom and obtained from preceding clauses by resolution. In the sequel, however, we assume that an arbitrary but fixed single history is associated with each considered resolution derivation. Thus, with a refutation P we can associate the directed acyclic graph (refutation DAG) $G(P)$ whose vertices are the clauses of P , and which has an arc from C_i to C_j if there is $C_k, i, k < j$, such that C_j is the resolvent of C_i and C_k . In any refutation DAG $G(P)$, the in-degree of each axiom vertex is 0, while each resolvent has in-degree 2.

It is well known that resolution is a complete proof system for unsatisfiable formulas: a formula is unsatisfiable if and only if there exists a resolution refutation of it [5]. The resolution complexity or resolution hardness $h(F)$ of an unsatisfiable formula F is the length of a shortest resolution refutation of F . For a nonempty set C of formulas, we define $h(C) = \sup_{F \in C} h(F)$.

2.3. Unsatisfiability minimal and saturated

A formula F is minimally unsatisfiable if it is unsatisfiable and removing any clause $C \in F$ leaves $F \setminus \{C\}$ satisfiable. A formula F is saturated minimally unsatisfiable if it is unsatisfiable and adding any literal (including on a fresh variable) to any clause makes it satisfiable. The central notion of the paper is, of course, that of a hitting formula.

Definition 1. A formula F is hitting if for every $C, D \in F, C \neq D, C$ and D clash.

Hitting formulas are notable for the following useful property.

Theorem 1 (Iwama [10]). The number of satisfying assignments of a hitting formula $F = \{C_1, \dots, C_m\}$ with n variables is given by $2^n(1 - \sum_{i=1}^m 2^{-|C_i|})$.

Proof. Because each pair of clauses of F clashes, every assignment can be covered by at most one clause. Since for each clause we can explicitly calculate the number of assignments it covers, we get the overall number of covered assignments—models are simply assignments that are not covered. \square

Corollary 1. A hitting formula F is satisfiable if and only if $\sum_{C \in F} 2^{-|C|} < 1$.

Corollary 2. All unsatisfiable hitting formulas are saturated minimally unsatisfiable.

Proof. Consider what happens to the model-counting sum after adding a literal or deleting a clause (both of which are operations that trivially preserve hittingness). \square

Two important notions in the context of minimally unsatisfiable formulas are *singularity* and *deficiency*. A literal $\ell \in \text{lit}(F)$ is *singular* if it is only contained in one clause. A variable $v \in \text{var}(F)$ is singular if v or \bar{v} is a singular literal. A formula is singular if it has a singular variable, and is *regular* otherwise. Singular formulas can be reduced by *DP-reduction*—the elimination of a variable v by taking all possible resolvents on v and then removing all clauses that contained v or \bar{v} . While DP-reduction can be applied to any variable of any formula preserving satisfiability, *singular DP-reduction*—when the eliminated variable v is singular—always decreases formula size, and thus can be applied exhaustively until a regular formula is obtained in overall polynomial time. Thus arises the importance of regular formulas—they are a kind of ‘core’ left over once singular variables are eliminated.

The other important notion in the context of minimally unsatisfiable formulas is that of the *deficiency*, defined as the difference between the number of clauses and the number of variables of a formula, and known to be positive for minimally unsatisfiable formulas [1]. It is known that minimally unsatisfiable formulas of deficiency 1 with more than 1 clause are always singular, and, up to isomorphism, there is only one regular minimally unsatisfiable formula of deficiency 2 for each number of clauses $m \geq 4$ [11] (otherwise there are none); we will refer to it as \mathcal{F}_m^2 . These formulas are actually saturated minimally unsatisfiable, and for $m = 4, 5$ also hitting. Kullmann and Zhao refer to \mathcal{F}_4^2 and \mathcal{F}_5^2 as \mathcal{F}_2 and \mathcal{F}_3 respectively [19].

It is known that singular DP-reduction of minimally unsatisfiable formulas preserves both minimal unsatisfiability as well as deficiency, and it also preserves saturated minimal unsatisfiability [18, Corollaries 10 and 13]. While singular DP-reduction is not necessarily confluent in minimally unsatisfiable formulas (different orders of elimination can result in different formulas), it is confluent in saturated minimally unsatisfiable formulas [18, Thm 23].

We will use the following acronyms to describe various classes of formulas: UH (unsatisfiable hitting), RUH (regular UH), IUH (irreducible UH, necessarily regular), SMU (saturated minimally unsatisfiable), RSMU (regular SMU), SSMU (singular SMU), RISMU (irreducible RSMU), SISMU (irreducible SSMU). For each of these classes we may append one or two integral parameters in parentheses: for example IUH(n, m) are IUHs with exactly n variables and m clauses, and IUH(m) are IUHs with exactly m clauses (without the restriction on the number of variables).

2.4. Isomorphisms and symmetries

An *isomorphism* between two formulas F and G is a bijection $\varphi : \text{lit}(F) \rightarrow \text{lit}(G)$ such that for each literal $\ell \in \text{lit}(F)$ we have $\varphi(\bar{\ell}) = \overline{\varphi(\ell)}$ and for each $C \subseteq \text{lit}(F)$ we have $C \in F$ if and only if $\varphi(C) := \{\varphi(\ell) \mid \ell \in C\} \in G$. If there is an isomorphism between F and G , we say F and G are *isomorphic*. For instance the formulas $F = \{\{x, y\}, \{\bar{x}, y\}, \{\bar{y}\}\}$, and $G = \{\{\bar{z}, w\}, \{z, w\}, \{\bar{w}\}\}$ are isomorphic. Obviously, two isomorphic formulas have the same properties concerning satisfiability, minimal unsatisfiability, and resolution complexity.

An *automorphism* (or symmetry) of a CNF formula F is an isomorphism to itself. $\text{Aut}(F)$ denotes the set of automorphisms of F . For an automorphism $\varphi \in \text{Aut}(F)$ we denote by φ^* the bijection $\text{var}(F) \rightarrow \text{var}(F)$ defined by $\varphi^*(x) = \text{var}(\varphi(x))$; φ^* is the *variable action* of φ on x . Let $\text{Aut}^*(F)$ denote the set of all variable actions of $\text{Aut}(F)$.

The *orbit* of a literal ℓ in F is the set $\text{Orb}_F(\ell) = \{\varphi(\ell) \mid \varphi \in \text{Aut}(F)\}$. The *variable orbit* of a variable x in F is the set $\text{Orb}_F^*(x) = \{\varphi^*(x) \mid \varphi \in \text{Aut}^*(F)\}$.

For sets S, S', T , we write $T = S \uplus S'$ if $T = S \cup S'$ and $S \cap S' = \emptyset$. A *2-graph* is an undirected graph $G = (V, E)$ together with a partition of its vertex set into two disjoint subsets $V_1 \uplus V_2 = V$. Two 2-graphs $G = (V_1 \uplus V_2, E)$ and $G' = (V'_1 \uplus V'_2, E')$ are *isomorphic* if there exists a bijection $\varphi : V_1 \uplus V_2 \rightarrow V'_1 \uplus V'_2$ such that $v \in V_i$ if and only if $\varphi(v) \in V'_i$, $i = 1, 2$, and $\{u, v\} \in E$ if and only if $\{\varphi(u), \varphi(v)\} \in E'$.

The *clause-literal graph* of a formula F is the 2-graph $G(F) = (V_1 \uplus V_2, E)$ with $V_1 = \text{lit}(F)$, $V_2 = F$, and $E = \{\{x, \bar{x}\} \mid x \in \text{var}(F)\} \cup \{\{C, \ell\} \mid C \in F, \ell \in C\}$. We refer to the edges $\{x, \bar{x}\}$ as *variable edges*. It is easy to verify that any two formulas are isomorphic if and only if their clause-literal graphs are isomorphic (as 2-graphs).

3. Irreducibility

The central notion of this section is that of a formula *factor* (called ‘clause-factor’, and later ‘clause-irreducible’ by Kullmann and Zhao [19]).

Definition 2 (Kullmann and Zhao [19]). A *factor* of a formula G is a subset $F \subseteq G$ logically equivalent to a single clause C . The clause C is called the *basis* of the factor F and is uniquely determined.

Lemma 1. Let $F \subseteq G$. Then F is a factor if and only if there is a clause C such that $C \subseteq D$ for every $D \in F$ and $\{D \setminus C \mid D \in F\}$ is unsatisfiable.

Proof. If F is a factor, then its basis satisfies the conditions on C . Conversely, if there is such a C , then F must be equivalent to it. \square

Non-trivial formula factors, like non-trivial factors for integers and other kinds of algebraic objects, can be used to decompose or reduce formulas. That leads to the central notion of an *irreducible* formula—one which cannot further be reduced by collapsing factors to their bases.

Definition 3 (Kullmann and Zhao [19]). A formula G is called *irreducible* if for each factor $F \subseteq G$ either $|F| = 1$ or $F = G$. Otherwise, it is reducible.

There is an important connection between factors and resolution hardness. In a nutshell, in presence of non-trivial factors, one can always pursue the strategy of reducing any remaining non-trivial factors to their bases, gradually shrinking the formula. These reductions can be performed in resolution, giving a way of constructing a canonical ‘decomposition’ refutation. An upper bound on resolution complexity follows.

Lemma 2. Let G be an unsatisfiable formula with m clauses, $F \subseteq G$ a factor of G , C the basis of F . Then G has a proof of length at most $h(F[\bar{C}]) + h(\{C\} \cup G \setminus F) - 1$.

Proof. Because F is a factor, $F[\bar{C}]$ is unsatisfiable and can be refuted in $h(F[\bar{C}])$ resolution steps. By re-inserting C , we can derive C from F in $h(F[\bar{C}])$ resolution steps. Because $F \equiv C$, $G \setminus F \cup \{C\}$ is still unsatisfiable, and hence can be refuted in $h(\{C\} \cup G \setminus F)$ steps. Combining the two refutations minus 1 for the shared clause C we get the result. \square

In the previous proof, we do not actually need that $F \equiv C$. All we need is that $F \models C$ and $G \setminus F \cup C$ is unsatisfiable. In other words, that C is a kind of ‘interpolant’ of F and $G \setminus F$. Such generalized factors will turn out quite useful in the context of resolution.

Definition 4. A clause C is a *clause interpolant* of (F, G) if $F \models C$ and $C \models \neg G$. We say that $F \subseteq H$ is a *pseudo-factor* of H if $(F, H \setminus F)$ has a clause interpolant. A formula H with m clauses is *strongly irreducible* if it has no pseudo-factors of size other than 1 or m , otherwise it is *weakly reducible*.

Corollary 3. Let G be an unsatisfiable formula with m clauses, $F \subseteq G$ a pseudo-factor of G , C a clause interpolant of $(F, G \setminus F)$. Then G has a refutation of length at most $h(F[\bar{C}]) + h(\{C\} \cup G \setminus F) - 1$.

It is easy to see that strong irreducibility is indeed a strengthening of irreducibility (Lemma 3), but the converse is not true even for MU formulas (Example 1).

Lemma 3. Every unsatisfiable strongly irreducible formula is irreducible.

Proof. A non-trivial factor is by definition also a non-trivial pseudo-factor. \square

Example 1. The formula $\{x, y, \{\bar{x}\}, \{\bar{y}\}\}$ is minimally unsatisfiable and irreducible (no pair of clauses is equivalent to a single clause). But it is not strongly irreducible: the first two clauses form a pseudo-factor with the clause interpolant $\{y\}$.

The formula in the previous example is minimally unsatisfiable but not saturated. As it turns out, it must be; for saturated minimally unsatisfiable formulas (which includes hitting formulas), irreducibility implies strong irreducibility.

Theorem 2. Any pseudo-factor of a saturated minimally unsatisfiable formula is also a factor. Consequently, an SMU formula is strongly irreducible if and only if it is irreducible.

Proof. For contradiction, let $F \subseteq G$ be a non-trivial pseudo-factor of the weakly reducible SMU formula G , let C be the clause interpolant. Let $F' = \{D \cup C \mid D \in F\}$. Then $F' \equiv F \vee C$, and hence $F' \models C$ and so $F' \cup G \setminus F$ is unsatisfiable. Because G is saturated, $F' = F$. Clearly $C \models F' = F$, and so F is in fact a factor and G is reducible. \square

Theorem 2 is an important background link that allows us to benefit from strong irreducibility while only having to check irreducibility, which is simpler (cf. Section 5, where we describe how we generate irreducible hitting formulas, and Section 6.1, where we use strong irreducibility in computing shortest resolution refutations).

Although saturated minimal unsatisfiability is sufficient to elevate irreducibility onto strong irreducibility, it is not necessary.

Example 2. The minimally unsatisfiable formula $\{x_1, \bar{x}_2, \{x_2, \bar{x}_3\}, \{x_3, \bar{x}_4\}, \{x_4, \bar{x}_1\}, \{x_1, x_2, x_3\}, \{\bar{x}_1, \bar{x}_2, \bar{x}_3, \bar{x}_4\}\}$ is not saturated (it is obtained by deleting the literal x_4 from the 5th clause of the unique deficiency-2 RSMU formula \mathcal{F}_6^2), yet it is strongly irreducible. Admittedly, why it is strongly irreducible is not obvious: we verified this by computer, trying every possible split and every possible clause interpolant.




Fig. 2. A shortest resolution refutation of the formula \mathcal{F}_5^2 from Example 3 (axioms highlighted).

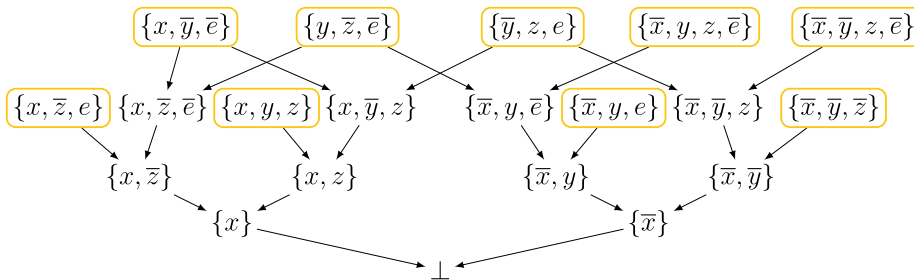


Fig. 3. A shortest resolution refutation of the formula G from Example 4 (axioms highlighted).

Finally, we arrive at the statement that we will actually use in our encoding to prune the search spaces by asserting certain kinds of refutations do not exist for strongly irreducible minimally unsatisfiable formulas.

Lemma 4. *Let F be a strongly irreducible minimally unsatisfiable formula with more than 2 clauses, let P be a resolution refutation of F . Let $C, D \in F$ be two axioms that are resolved together in P . Then at least one of C, D is used at least once more in P .*

Proof. $\{C, D\} \subseteq F$ is not unsatisfiable, because F is minimally unsatisfiable. Hence P continues beyond the resolution of C and D . If neither C nor D is used another time, then the resolvent of C and D is a clause interpolant for $(\{C, D\}, F \setminus \{C, D\})$, contradicting F 's strong irreducibility. \square

Kleine Büning and Zhao [13] studied *read-once* refutations for minimally unsatisfiable formulas—ones in which every clause is only used once. Lemma 4 has implications for the existence of read-once refutations for strongly irreducible formulas.

Corollary 4. *Strongly irreducible minimally unsatisfiable formulas with more than 2 clauses do not have read-once refutations.*

Example 3. Corollary 4 is in a sense tight, because there are strongly irreducible formulas with refutations where only a single axiom is read twice and everything else only once. An example is the formula

$$\mathcal{F}_5^2 = \{\{x_1, \bar{x}_2\}, \{x_2, \bar{x}_3\}, \{x_3, \bar{x}_1\}, \{x_1, x_2, x_3\}, \{\bar{x}_1, \bar{x}_2, \bar{x}_3\}\},$$

known to have hardness 10 by [24], with a shortest refutation shown in Fig. 2.

One could ask whether the decomposition refutation from Lemma 2 is optimal. This is not the case, as witnessed by the following example.

Example 4. Let $F = \{\{x, y, z\}, \{\bar{x}, \bar{y}, \bar{z}\}, \{\bar{x}, y, e\}, \{x, \bar{y}, \bar{e}\}, \{x, \bar{z}, e\}, \{\bar{x}, z, \bar{e}\}, \{y, \bar{z}, \bar{e}\}, \{\bar{y}, z, e\}\}$. F is the formula $F_{4,8,52}$ from [24] and so is known to have hardness 19. Now, consider G obtained from F by replacing the clause $\{\bar{x}, z, \bar{e}\}$ with the clauses $\{\bar{x}, y, z, \bar{e}\}$ and $\{\bar{x}, \bar{y}, z, \bar{e}\}$. Clearly, G is reducible: the two new clauses are a factor, whose basis is the replaced clause. If the decomposition refutation were optimal, the hardness of G would have been $3 + h(F) - 1 = 21$. But G has a refutation of length 20 (see Fig. 3).

Thus concludes our tour of irreducibility and its strong cousin. In the following subsection, we show that asymptotic resolution hardness of hitting formulas is decided on irreducible formulas—if they have polynomial-size resolution refutations, so do all unsatisfiable hitting formulas. Then, in Section 6.1, we employ strong irreducibility to improve the computation of shortest refutations.

3.1. Irreducible formulas and asymptotic hardness

The goal of this subsection is to show that irreducible hitting formulas are no easier (asymptotically up to a polynomial factor) for resolution than non-irreducible ones, thus showing that the question of whether hitting formulas have polynomial refutations can be decided solely by looking at irreducible formulas. Even though our focus is on hitting formulas here, there is nothing special about hittingness that makes (strongly) irreducible formulas asymptotically hardest. We can state and prove the theorem generally for any class of unsatisfiable formulas that is closed under factorization.

Definition 5. Let \mathcal{C} be a set of unsatisfiable formulas. We say \mathcal{C} is *closed under factorization* if for every weakly reducible formula $G \in \mathcal{C}$ and pseudo-factor F of G with clause interpolant C , both $F[\bar{C}]$ and $G \setminus F \cup \{C\}$ are in \mathcal{C} .

Theorem 3. Let \mathcal{C} be a set of unsatisfiable formulas closed under factorization, let $\mathcal{C}^* \subseteq \mathcal{C}$ be the set of strongly irreducible formulas of \mathcal{C} , and let $\mu : \mathbb{N} \rightarrow \mathbb{N}$ be a non-decreasing function such that every formula $F \in \mathcal{C}^*$ has a resolution refutation of length at most $\mu(|F|)$. Then every formula $F \in \mathcal{C}$ has a resolution refutation of length at most $(|F| - 1)\mu(|F|) + 1$.

Proof. By induction on the number of clauses.

Base case. Every unsatisfiable formula has at least one clause. The induction hypothesis holds for the only 1-clause unsatisfiable formula, which contains only the empty clause and has a refutation of length $(1 - 1)\mu(1) + 1 = 1$. The induction hypothesis also holds for all $G \in \mathcal{C}^*$ because when $|G| > 1$, we have $(|G| - 1)\mu(|G|) + 1 > \mu(|G|)$.

Inductive step. Consider a weakly reducible $G \in \mathcal{C} \setminus \mathcal{C}^*$. Pick a non-trivial pseudo-factor $F \subset G$, let $C = \bigcap F$ be the clause interpolant of F and $G \setminus F$, let $F' = F[\bar{C}]$ and $H = G \setminus F \cup \{C\}$. By [Corollary 3](#), G has a refutation of size $h(F') + h(H) - 1$. Because \mathcal{C} is closed under factorization, $F', H \in \mathcal{C}$. Both F' and H have fewer clauses than G , and so by the induction hypothesis $h(F') \leq (|F'| - 1)\mu(|F'|) + 1$ and $h(H) \leq (|H| - 1)\mu(|H|) + 1$. Summing up, we get that $h(G) \leq h(F') + h(H) - 1 \leq (|F'| + |H| - 2)\mu(|G|) + 2 - 1 = (|G| - 1)\mu(|G|) + 1$. \square

When $\mathcal{C} = \text{UH}$, [Theorem 3](#) translates to the following corollary.

Corollary 5. If irreducible unsatisfiable hitting formulas have polynomial-size resolution refutations, then so do all unsatisfiable hitting formulas.

Proof. In order to apply [Theorem 3](#) to hitting formulas and obtain the statement, all that needs to be shown is that the set of hitting formulas is closed under factorization. This follows from Lemma 33 of [\[19\]](#) (and from the fact that any pseudo-factor of a hitting formula is a factor by [Theorem 2](#)), but we also repeat the argument below for sake of completeness.

It is easy to see that if G is hitting, then $F \subseteq G$ is also hitting, and hittingness is further preserved under restriction, hence $F[\bar{C}]$ is hitting. Now, consider the other formula $G \setminus F \cup \{C\}$, where F is a factor of G and C is the basis of F . Clearly $G \setminus F \subseteq G$ is hitting. It remains to show that C clashes with all clauses in $G \setminus F$. By contradiction, assume it does not clash with $D \in G \setminus F$. Then, because G is hitting, D clashes with every clause in F , but because D does not clash with C , in fact D clashes with every clause in $F[\bar{C}] = \{E \setminus C \mid E \in F\}$. This is the same as saying that \bar{D} is a satisfying assignment for $F[\bar{C}]$, contradicting the assumption that $F \models C$. \square

In fact, a slightly stronger statement still than [Corollary 5](#) can be derived from [Theorem 3](#). Let (G_1, \dots) be an infinite sequence of hitting formulas of super-polynomial resolution hardness, assuming one exists. Let \mathcal{C} be the closure of $\{G_1, \dots\}$ under factorization, i.e., the intersection of all sets closed under factorization containing $\{G_1, \dots\}$. Then \mathcal{C} contains an infinite sequence of irreducible formulas of superpolynomial hardness—for if it did not, then [Theorem 3](#) would imply that the sequence (G_1, \dots) also has polynomial-size refutations. In other words, in order to construct a family of hard hitting formulas, one must at least implicitly construct a family of hard irreducible hitting formulas.

4. Symmetries

In this section we will see how formula symmetries can be used to manipulate resolution refutations and to single out certain canonical ones.

Lemma 5. Let G be an unsatisfiable formula, $\varphi \in \text{Aut}(G)$ a symmetry with $\varphi^*(x) = y$ for some $x, y \in \text{var}(G)$. If there is a shortest refutation of G that ends with resolving x , then there is one that ends with resolving y . Moreover, the refutation can be arranged such that it ends with either the sequence $\{y\}, \{\bar{y}\}, \perp$ or $\{\bar{y}\}, \{y\}, \perp$ (both are always possible).

Proof. The first part follows by applying φ to the refutation (literal by literal). For the rest we need to show that neither of $\{y\}, \{\bar{y}\}$ is needed to derive the other. Suppose $\{y\}$ is used in the derivation of $\{\bar{y}\}$; surely by removing the use of $\{y\}$ we will end up with a derivation of a clause that at most contains an additional \bar{y} —but it already did to begin with. \square

Lemma 6. Let G be an unsatisfiable formula, $\text{Aut}(G)$ its symmetry group, \mathcal{O} the set of orbits under the variable action. Let $R \subseteq \text{var}(G)$ be a set of representatives of each orbit from \mathcal{O} , i.e., for all $O \in \mathcal{O}$ $|R \cap O| = 1$. Then there is a variable $v \in R$ and a shortest refutation of G which ends in the sequence $\{v\}, \{\bar{v}\}, \perp$, and one which ends in the sequence $\{\bar{v}\}, \{v\}, \perp$.

Proof. There is some shortest refutation of G and it ends by resolving some variable x . Let $v \in R \cap \mathcal{O}(x)$, and apply Lemma 5. \square

5. Generating hitting formulas

We generate hitting formulas with a tailor-made adaptation of Nauty [20], which generates clause-literal graphs modulo isomorphisms. In order to generate formulas efficiently, we hook into Nauty’s generation process and prune partially constructed formulas on the fly when we determine they cannot be extended to full formulas.

Suppose we are generating formulas with n variables and m clauses. Nauty generates the formulas by progressively adding vertices corresponding to clauses in all possible ways modulo isomorphisms. Thus, it effectively exploring a tree of ‘partial’ formulas, and the leaf nodes are ‘full’ formulas we want to generate. Suppose we have a partially constructed candidate formula $F = \{C_1, \dots, C_{m'}\}$, $m' \leq m$ (including the possibility $m' = m$, i.e., the formula is a leaf node). Thanks to the way Nauty generates graphs, we can assume that $C_{m'}$ is the largest clause of F and that all further clauses added in any extension of F will be at least as large as $C_{m'}$. As we are interested in generating only IUHs and not arbitrary formulas, we can afford to prune as follows (here *prune* means discard the current partial formula and immediately backtrack in the tree explored by Nauty):

- If F is not hitting, prune.
- If $\#F > (m - m')2^{n - |C_{m'}|}$, prune (otherwise future clauses cannot cover all satisfying assignments).
- If $\#F < m - m'$, prune (every clause must uniquely cover at least one assignment, and here we have too few satisfying assignments to be uniquely covered by future clauses).
- If F has a non-trivial factor, or if F itself is a factor and $1 < m' < m$, prune.

The first three checks can be implemented in linear time: models can be counted easily thanks to Theorem 1, for hittingness it is sufficient to check whether the last clause hits every other, because the same check has been performed on $\{C_1, \dots, C_{m'-1}\}$ already.

We test the fourth condition by enumerating all non-empty subsets $F' \subseteq F$ and checking whether $F' \equiv I := \{\bigcap F'\}$. Since obviously $I \models F'$, it is sufficient to compare the number of models in order to see whether I and F' are equivalent—and this can be done in polynomial time as before, because F' is still hitting. Moreover, thanks to the incremental nature of the check, we can focus only on subsets F' with $C_{m'} \in F'$.

We note that it is possible to test irreducibility faster than by traversing all subsets of clauses. The key observation is that any factor in a minimally unsatisfiable formula must be *intersection-maximal*, meaning that any its strict superset has a strictly smaller intersection (we leave this as exercise for the reader, it follows straightforwardly from the definitions). Thus, one only needs to enumerate intersection-maximal subsets of clauses, which can be done with polynomial delay (enumerate subsets recursively and after each branching, include all clauses that contain the current intersection; again we leave the details to the reader). However, we did not implement this in Nauty, as the full traversal has an advantage in technical simplicity that outweighs algorithmic gains for small numbers of clauses.

We also generated some reducible regular unsatisfiable hitting formulas (RUHs). In that case, we can re-use the pruning described above, and simply skip the fourth test. However, focusing on irreducible formulas allows us to prune significantly more and generate formulas faster and up to bigger size; and the resulting number of formulas is smaller making subsequent processing more manageable (see Section 7).

6. Computing shortest refutations

In this section, we will see how to use (strong) irreducibility and formula symmetries in order to rule out certain refutations, thereby pruning the search space and making it easier for the solver to show no refutations of a given length exist.

The problem we are solving asks, given a formula F and integer s , whether F has a resolution refutation of length at most s , i.e., whether there is a sequence C_1, \dots, C_s such that each C_i is either a clause of F , or is derived by resolution from two previous clauses. We will solve the problem by reduction to SAT, building on a previous encoding of ours [24]; for a full presentation of the encoding we refer to that paper. Here we will only recall that the encoding is centered around the following three sets of propositional variables:

- $\text{pos}[i, v]$ and $\text{neg}[i, v]$ denote that the v and \bar{v} occur in the i th clause of the refutation;
- $\text{arc}[i, j]$ denotes that the i th clause is used to obtain the j th clause via resolution. Every non-axiom clause has exactly two incoming arcs, i.e., for all $j > m$ we have that $\sum_{i=1}^{j-1} \text{arc}[i, j] = 2$ (where m is the number of axiom clauses; without loss of generality we assume those appear at the beginning of the proof).

An assignment to these variables together fully determines a candidate resolution derivation: it is then left to constraints to validate that the derivation is a valid refutation.

We generally assume that F is minimally unsatisfiable, as all formulas we work with are. In the following two subsections, we first show how we can add a further redundant constraint when we additionally assume strong irreducibility, and then how we can reason about refutation symmetries via formula symmetries.

6.1. Strong irreducibility and clause reuse

Recall [Lemma 4](#), which tells us that we must resolve some axioms more than once. We will simply translate it into a SAT encoding. We define a set of fresh variables $\text{active}[i, j]$ equal to the clause $\{\text{arc}[i, j], \dots, \text{arc}[i, s]\}$ to denote that the i th clause is still ‘active’ at position j in the refutation, i.e., it is used for resolution to obtain the j th or later clause.¹ With this, we can rewrite [Lemma 4](#) right into the set of clauses

$$\{\overline{\text{arc}[i, k]}, \overline{\text{arc}[j, k]}, \text{active}[i, k + 1], \text{active}[j, k + 1]\},$$

for $1 \leq i < j \leq m$ and $m < k < s$. In practice, because this constraint can get quite large, we decided to use only a limited version, where k is fixed to $m + 1$ (i.e., we only apply [Lemma 4](#) to the first resolvent of the refutation).

We point out that while the constraint just described is a translation of [Lemma 4](#), successful application also relies on [Theorem 2](#). That is because while we require strong irreducibility, we generate formulas that are only irreducible. But because they are hitting (saturated minimally unsatisfiable), [Theorem 2](#) kicks in and warrants strong irreducibility.

6.2. Symmetries

Recall [Lemma 6](#), which says that formula symmetries can be applied in order to obtain symmetric images of refutations. If we fix an arbitrary set R of variable representatives, thanks to [Lemma 6](#) we need only look for refutations which end in a variable from R . Further, we can place the two clashing unit clauses at the end of the refutation in the order we pick (say the penultimate clause has a negative literal). This is quite straightforward to put into constraints: for all $v \in \text{var}(F)$ we include the unit clauses $\{\text{pos}[s - 1, v]\}$ and $\{\text{neg}[s - 2, v]\}$, and for all $v \in \text{var}(F) \setminus R$, we include the unit clauses $\{\text{neg}[s - 1, v]\}$ and $\{\text{pos}[s - 2, v]\}$. Additionally, we stipulate that the literals appearing in C_{s-1} and C_{s-2} are on the same variable, i.e., for all $v \in R$, $\text{neg}[s - 1, v] = \text{pos}[s - 2, v]$, and the arc structure at the end of the refutation: $\{\text{arc}[s - 2, s]\}$, $\{\text{arc}[s - 1, s]\}$, and $\{\text{arc}[s - 2, s - 1]\}$.

The encoding presented in our previous work [24] uses symmetry breaking to allow only one sequence for any given refutation DAG—the canonical topological sort. Care must be taken to ensure that our new formula symmetry breaking is compatible with canonical topological sorting. Fortunately, there is a simple way out: disable canonical topological sorting on the clauses C_{s-1}, C_{s-2} .

In practice, we obtain the symmetry group of a formula and the variable orbits with Sage [29].

7. Results

We will now present our experimental results and connections to the work of others.

The main results are summarized in [Tables 1](#) and [2](#). For each pair n, m , they show the maximum hardness of IUHs ([Table 1](#)) and RUHs ([Table 2](#)) with n variables and m clauses, together with the number of formulas which attain maximum hardness (subscript) and total number of formulas with n variables and m clauses (superscript; both numbers modulo isomorphisms). We computed these numbers (the shortest refutations mainly) on a cluster of heterogeneous machines with the SAT solver CaDiCaL,² which we already observed earlier to be most effective for finding shortest refutations [24].

[Fig. 4](#) compares IUH hardness to hardness of RUHs and other related classes like SMU, depicting both maximum as well as average hardness.³ The data on SMU hardness is taken from our previous work [24], where we looked for *resolution hardness numbers* (maximum resolution complexity with a fixed number of clauses) and showed that the hardest formulas with a given number of clauses are saturated minimally unsatisfiable, and consequently analyzed hardness in SMU formulas with up to 10 clauses in detail.

We can draw several conclusions.

Firstly, it is not true that IUHs are always the hardest of hitting formulas for a given number of clauses. On the other hand, IUHs exhibit greater hardness *on average* than all hitting formulas, suggesting that irreducibility is indeed positively correlated with hardness.

Secondly, with the caveat that our limited study cannot provide answers for larger formulas, it seems that hitting formulas, irreducible or not, are not substantially easier than formulas in general. Interestingly, while for most values of m (number of clauses) the hardest formulas are not hitting, for 5 and 8 clauses there is an IUH that is as hard as any formula (for $m = 5$ it is \mathcal{F}_5^2 , for $m = 8$ it is the formula F from [Example 4](#)).

¹ We use the Tseitin encoding in order to translate this set of equations into CNF.

² <http://fmv.jku.at/cadical>

³ We calculate average hardness as weighted average by the number of isomorphic copies of each formula.

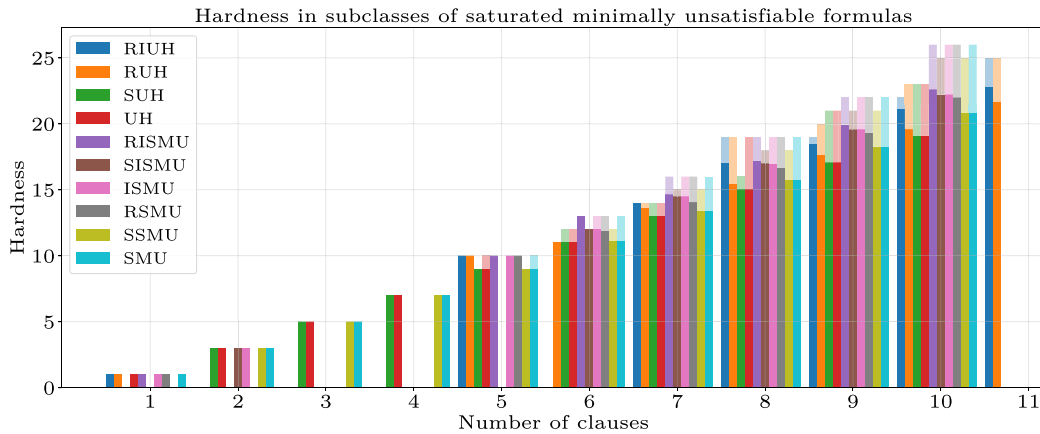


Fig. 4. Hardness in subclasses of SMU. The opaque bar shows average hardness for a given number of clauses, the transparent bar extends to maximum hardness. While irreducible hitting formulas are not as hard as general hitting formulas in the worst case, they are harder on average.

Table 2

Values of $h(\text{RUH}(n, m))$, i.e., the lengths of the longest shortest refutation required by a regular unsatisfiable hitting formula with n variables and m clauses. The subscript gives the number of formulas that attain maximum hardness, the superscript is the total number of formulas in each category. All formula counts are modulo isomorphisms. Empty areas contain no RUHs.

$n \setminus m$	1	2	3	4	5	6	7	8	9	10	11
0	1_1^1										
1											
2											
3				7_1^1							
4					10_1^1	11_3^3	13_1^1	15_1^1			
5							14_{10}^6	19_{11}^{49}	20_1^{79}	21_{21}^{94}	22_{21}^{70}
6								16_5^9	19_7^{207}	23_3^{772}	25_{21}^{8203}
7									19_2^4	21_7^{281}	25_3^{7449}
8										21_1^1	23_{50}^{261}
9											

Of course, limited conclusions should be drawn from our limited computational evaluation. We note that it is probably not possible to advance to much higher numbers of clauses with this approach: we are looking at over 200K IUHs with 15 clauses, a number that is likely going to be very difficult to process for shortest refutations. The situation is even worse for RUHs, with almost 140K with 12 clauses and over 1M formulas with 13 clauses. A further complication is that the computation of shortest refutations grows hard quickly with increasing formula size [24]. The longest successfully solved case took more than 86 days for an IUH with 6 variables and 14 clauses, and that still leaves many formulas with 14 clauses left. The formula with the highest found hardness, also with 6 variables and 14 clauses, with hardness 34, took two weeks to solve. As we noted in our previous work [24], this workflow is difficult to parallelize as most of the effort is concentrated in one hardest SAT query.

On the flip side though, the pace of growth suggests there might indeed be infinitely many IUHs. A refined existence question could ask whether IUHs with each number of clauses exist. Table 1 shows that is not the case: IUHs with 2, 3, 4, and 6 clauses are missing. For 2 and 3, regularity is the culprit: there are no regular minimally unsatisfiable formulas with 2 or 3 clauses either [18,24]. The case of 4 and 6 is a different story: there are both RUHs and RISMUs with 6 clauses, and there is a RUH with 4 clauses, but interestingly no IUH with 6 clauses, and not even a RISMU with 4 clauses. This reminds us of other algebraic objects like orthogonal Latin squares, which also exist for almost all sizes (cf. OEIS sequences A160368, A305570, A305571, and A287761), but it is unclear whether there is any connection.

Kullmann and Zhao [19] were interested in the existence of RUHs, and in particular they investigated the ‘Finiteness Conjecture’ that for any fixed deficiency only finitely many RUHs exist. An exact, stronger variant of their conjecture [19, Conjecture 5] states explicitly that a RUH of deficiency $k \geq 2$ can have at most $4k - 5$ variables (and thus $5k - 5$ clauses), which they prove for $k = 3$ (and which was already known for $k = 2$ [11]). This exactness allows our work to intersect with theirs, and indeed, wherever we touch the bound of $4k - 5$, Tables 1 and 2 confirm that no RUHs (or IUHs) of deficiency k and with more than $4k - 5$ variables exist. In fact, all three extremal cases that we cover (for $k = 1, 2, 3, n = 0, 3, 7$) have a unique RUH formula modulo isomorphism (the unique formula for $k = 3$ is the formula K_2 described by Kullmann and Zhao [19, Lemma 4]). This was already well known for $k = 1, 2$, and followed for $k = 3$ from our previous work [24], where we generated all RSMU formulas with 7 variables and 10 clauses, but we did not see the connection to Kullmann

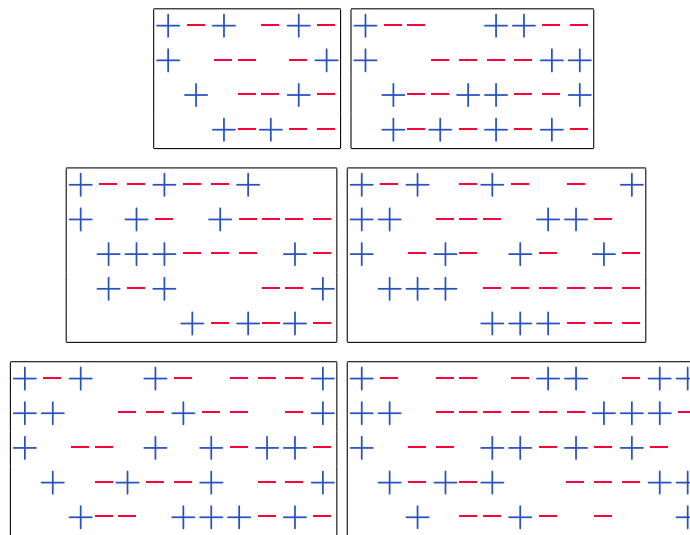


Fig. 5. Selected hardest IUHs with 7 and 9 – 13 clauses written in ‘incidence matrix’ form (rows are variables, columns are clauses, ‘+’ and ‘-’ mean positive and negative occurrence). The hardest IUH with 1 clause is $\{\perp\}$, with 5 clauses \mathcal{F}_5^2 , already shown in Example 3, and with 8 clauses the formula F from Example 4.

and Zhao’s work then. There is another intriguing pattern of similar nature apparent in Table 1: Formulas with n variables show up starting with $m = 2n - 1$ clauses. We leave it to future work to find out whether these patterns are coincidental or not.

Our catalog of RUHs and IUHs will be added to our existing catalog of SMU formulas [23], and will hopefully help with the testing of hypotheses and as inspiration for constructing infinite classes of irreducible hitting formulas. In Fig. 5, we show hand-picked examples of notable formulas. For each number of clauses, we picked among the IUHs with highest hardness one with the smallest length (number of literal occurrences, i.e. $\sum_{C \in F} |C|$), and in case there was more than one, the one with the largest symmetry group.

8. Conclusion

Inspired by the observation that hitting formulas are remarkable among polynomial-time decidable classes of propositional formulas in that they admit polynomial-time model counting and at the same time lack obvious resolution-complexity upper bounds, we set out to answer whether hitting formulas are hard for resolution. This quest led us into the land of (strong) irreducibility and to discover a number of interesting related questions and phenomena. With our theoretical and experimental results we now understand that IUHs are a key subset of hitting formulas with respect to both resolution complexity, as well as just plain existence. Because the number of IUHs (and even more so RUHs) grows so fast, it seems safe to conjecture that there are infinitely many—but at the same time there are intriguing similarities with other algebraic objects like orthogonal Latin squares, and the rules governing existence of IUHs (and RUHs) with fixed parameters n and m also exhibit simple-looking patterns that deserve further investigation. Based on the resolution complexity we could observe, we see no signs that hitting formulas should be significantly easier for resolution than formulas in general, even if they are a bit easier than the absolutely hardest formulas; although we cannot draw definitive conclusions beyond the formulas whose proofs we computed.

While we focused on resolution here, it would also be interesting to look at hitting formulas’ complexity in other proof systems.

Data availability

Data will be made available on request.

Acknowledgments

We would like to acknowledge the grants J-4361N and P32441 from the Austrian Science Fund (FWF) and the grant ICT19-065 from the Vienna Science and Technology Fund (WWTF).

References

- [1] Ron Aharoni, Nathan Linial, Minimal non-two-colorable hypergraphs and minimal unsatisfiable formulas, *J. Combin. Theory Ser. A* 43 (1986) 196–204.
- [2] Albert Atserias, Johannes Klaus Fichte, Marc Thurley, Clause-learning algorithms with many restarts and bounded-width resolution, *J. Artificial Intelligence Res.* 40 (2011) 353–373.
- [3] Paul Beame, Henry A. Kautz, Ashish Sabharwal, Towards understanding and harnessing the potential of clause learning, *J. Artificial Intelligence Res.* 22 (2004) 319–351.
- [4] J.A. Bondy, U.S.R. Murty, *Graph Theory*, in: Graduate Texts in Mathematics, vol. 244, Springer Verlag, New York, 2008.
- [5] M. Davis, H. Putnam, A computing procedure for quantification theory, *J. ACM* 7 (3) (1960) 201–215.
- [6] John Franco, John Martin, A history of satisfiability, in: Armin Biere, Marijn Heule, Hans van Maaren, Toby Walsh (Eds.), *Handbook of Satisfiability*, 2nd Ed., IOS Press, 2021, pp. 3–74.
- [7] Nicola Galesi, Oliver Kullmann, Polynomial time SAT decision, hypergraph transversals and the Hermitian rank, in: SAT 2004 - the Seventh International Conference on Theory and Applications of Satisfiability Testing, 10–13 May 2004, Vancouver, BC, Canada, Online Proceedings, 2004.
- [8] Robert Ganian, Stefan Szeider, New width parameters for SAT and #SAT, *Artificial Intelligence* 295 (2021) 103460.
- [9] Armin Haken, The intractability of resolution, *Theoret. Comput. Sci.* 39 (1985) 297–308.
- [10] Kazuo Iwama, CNF-satisfiability test by counting and polynomial average time, *SIAM J. Comput.* 18 (2) (1989) 385–391.
- [11] Hans Kleine Büning, On subclasses of minimal unsatisfiable formulas, *Discrete Appl. Math.* 107 (1–3) (2000) 83–98.
- [12] Hans Kleine Büning, Xishun Zhao, Satisfiable formulas closed under replacement, in: Henry Kautz, Bart Selman (Eds.), *Proceedings for the Workshop on Theory and Applications of Satisfiability*, in: Electronic Notes in Discrete Mathematics, vol. 9, Elsevier Science Publishers, North-Holland, 2001.
- [13] Hans Kleine Büning, Xishun Zhao, The complexity of read-once resolution, *Ann. Math. Artif. Intell.* 36 (4) (2002) 419–435.
- [14] Hans Kleine Büning, Xishun Zhao, On the structure of some classes of minimal unsatisfiable formulas, *Discrete Appl. Math.* 130 (2) (2003) 185–207.
- [15] Oliver Kullmann, The combinatorics of conflicts between clauses, in: Enrico Giunchiglia, Armando Tacchella (Eds.), *Sixth International Conference on Theory and Applications of Satisfiability Testing*, S. Margherita Ligure - Portofino, Italy, May 5–8, 2003 (SAT 2003), in: *Lecture Notes in Computer Science*, vol. 2919, Springer Verlag, 2004.
- [16] Oliver Kullmann, Green-Tao numbers and SAT, in: Ofer Strichman, Stefan Szeider (Eds.), *Theory and Applications of Satisfiability Testing - SAT 2010*, 13th International Conference, SAT 2010, Edinburgh, UK, July 11–14, 2010. Proceedings, in: *Lecture Notes in Computer Science*, vol. 6175, Springer, 2010, pp. 352–362.
- [17] Oliver Kullmann, Constraint satisfaction problems in clausal form II: Minimal unsatisfiability and conflict structure, *Fund. Inform.* 109 (1) (2011) 83–119.
- [18] Oliver Kullmann, Xishun Zhao, On davis-putnam reductions for minimally unsatisfiable clause-sets, *Theoret. Comput. Sci.* 492 (2013) 70–87.
- [19] Oliver Kullmann, Xishun Zhao, Unsatisfiable hitting clause-sets with three more clauses than variables, *CoRR* (2016) arXiv:1604.01288, arXiv:1604.01288, URL <http://arxiv.org/abs/1604.01288>.
- [20] Brendan D. McKay, Adolfo Piperno, Practical graph isomorphism, {ii}, *J. Symbolic Comput.* 60 (2014) 94–112.
- [21] Naomi Nishimura, Prabhakar Ragde, Stefan Szeider, Solving #sat using vertex covers, *Acta Inform.* 44 (7–8) (2007) 509–523.
- [22] Sebastian Ordyniak, Daniël Paulusma, Stefan Szeider, Satisfiability of acyclic and almost acyclic CNF formulas, *Theoret. Comput. Sci.* 481 (2013) 85–99.
- [23] Tomáš Peitl, Stefan Szeider, Saturated Minimally Unsatisfiable Formulas on Up to Ten Clauses, 10.5281/zenodo.3951545, Zenodo, 2021, DOI.
- [24] Tomáš Peitl, Stefan Szeider, Finding the hardest formulas for resolution, *J. Artificial Intelligence Res.* 72 (2021) 69–97, Conference Award Track, best paper CP 2020.
- [25] Knot Pipatsrisawat, Adnan Darwiche, On the power of clause-learning SAT solvers with restarts, in: Ian P. Gent (Ed.), *Principles and Practice of Constraint Programming - CP 2009*, 15th International Conference, CP 2009, Lisbon, Portugal, September 20–24, 2009, Proceedings, in: *Lecture Notes in Computer Science*, vol. 5732, Springer Verlag, 2009, pp. 654–668.
- [26] Dan Roth, On the hardness of approximate reasoning, *Artificial Intelligence* 82 (1–2) (1996) 273–302.
- [27] João P. Marques Silva, Kareem A. Sakallah, GRASP - A new search algorithm for satisfiability, in: *International Conference on Computer-Aided Design (ICCAD '96)*, November 10–14, 1996, San Jose, CA, USA, ACM and IEEE, 1996, pp. 220–227.
- [28] Balázs Szörényi, Disjoint DNF tautologies with conflict bound two, *J. Satisf. Boolean Model. Comput.* 4 (1) (2008) 1–14.
- [29] The Sage Developers, SageMath, the sage mathematics software system, 2022, <http://dx.doi.org/10.5281/zenodo.6259615>.
- [30] Alasdair Urquhart, Hard examples for resolution, *J. ACM* 34 (1) (1987) 209–219.
- [31] Xishun Zhao, Complexity results on minimal unsatisfiable formulas, in: *Mathematical Logic in Asia*, World Sci. Publ., Hackensack, NJ, 2006, pp. 302–319.