NP-Completeness of Refutability by Literal-Once Resolution*

Stefan Szeider

Institute of Discrete Mathematics
Austrian Academy of Sciences
Sonnenfelsgasse 19, 1010 Vienna, Austria
stefan.szeider@oeaw.ac.at

Abstract. A boolean formula in conjunctive normal form (CNF) F is refuted by *literal–once resolution* if the empty clause is inferred from F by resolving on each literal of F at most once. Literal–once resolution refutations can be found nondeterministically in polynomial time, though this restricted system is not complete. We show that despite of the weakness of literal–once resolution, the recognition of CNF-formulas which are refutable by literal–once resolution is NP-complete. We study the relationship between literal–once resolution and *read-once resolution* (introduced by Iwama and Miyano). Further we answer a question posed by Kullmann related to minimal unsatisfiability.

1 Introduction

Resolution is a method for establishing the unsatisfiability of formulas in conjunctive normal form (CNF), based on the *resolution rule*: if $C_1 \cup \{\ell\}$ and $C_2 \cup \{\overline{\ell}\}$ are clauses, then the clause $C_1 \cup C_2$ may be inferred, *resolving on* the literal ℓ. A *resolution refutation* of a CNF-formula F is a derivation of the empty clause $\square$ from F, using the resolution rule. It is well-known that resolution is *sound* and *complete*, i.e., a CNF-formula is unsatisfiable if and only if there is a resolution refutation of it ([14]). Resolution refutations can be represented as binary trees, where the leaves are labeled by clauses of F (see Figure 1 for an example). Unfortunately, the size of a shortest resolution refutation of a CNF-formula F

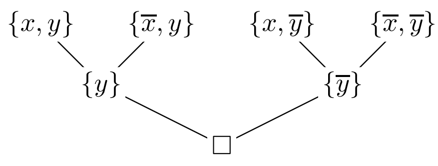

Fig. 1. A resolution refutation of $F = \{\{x, y\}, \{x, \overline{y}\}, \{\overline{x}, y\}, \{\overline{x}, \overline{y}\}\}$.

* This work has been supported by the Austrian Science Fund, P13417-MAT.

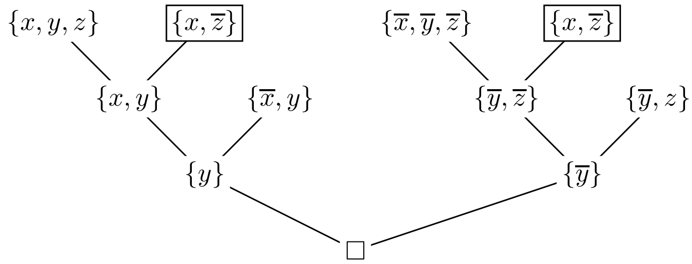

Fig. 2. A resolution refutation which is not read–once.

can be exponential in the number of clauses of F ([6, 7]). Therefore, considerable effort has been made to identify restricted (and incomplete) classes of resolution refutations where the size of refutations is polynomially bounded by the size of input formulas (see [10] for a survey). One of the best known examples is *unit resolution*, where the resolution rule is only applied to pairs of clauses C_1, C_2 if C_1 or C_2 is a unit clause (i.e., a singleton). Unit resolution is not complete any more, but the class of formulas which can be refuted by unit resolution can be recognized in linear time (see, eg., [10]).

Iwama and Miyano ([8]) considered *read–once resolution*, where each clause of the input formula must be used at most once in a refutation; i.e., two leaves of the resolution tree may not be labeled by the same clause. (In [8] also resolution refutations are considered, where clauses of the input formula may used more than once, but the number of repetitions is restricted.) For example, the refutation exhibited in Figure 2 is not read-once, since the clause $\{x, \overline{z}\}$ occurs at two leaves (in fact, it can be shown that for $F = \{\{x, y, z\}, \{x, \overline{z}\}, \{\overline{x}, y\}, \{\overline{x}, \overline{y}, \overline{z}\}, \{\overline{y}, z\}\}$ no read–once resolution exists, despite F being unsatisfiable; see [8] or Proposition 1 below). It is easy to see that the size of a read–once resolution refutation is polynomially bounded by the size of the input formula. However, in [8] it is shown that—in spite of the shortness of read–once resolution refutations—it is NP-complete to recognize formulas which can be refuted by read–once resolution.

If we modify the above example by adding two clauses $\{w, x, \overline{z}\}$ and $\{\overline{w}, x, \overline{z}\}$ to F, then we get a read–once resolution refutation (exhibited in Figure 3). There are still two occurrence of $\{x, \overline{z}\}$, but one occurrence became an interior vertex of the tree, and so the refutation became read-once. Thus, it is natural to consider resolution trees where no clause appears more than once *at any position* in the resolution tree. We call such refutations *strict read-once*. It can be shown that there are CNF-formulas which are refutable by read–once resolution, but not by strict read-once resolution (see Proposition 1 below). Since strict read–once resolution is therefore weaker than read–once–resolution, it is conceivable that refutability by strict read-once resolution can be decided in polynomial time. We will show, however, that recognition of formulas refutable by strict read-once resolution is NP-complete.

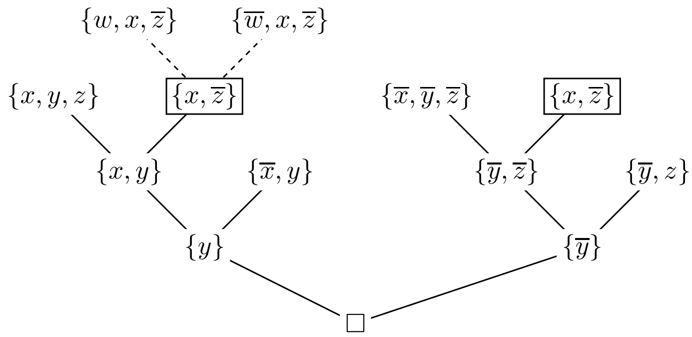

Fig. 3. A resolution refutation obtained from Figure 3; it is read–once, but not strict read–once.

Going one step further, we also consider a type of resolution which is even weaker than strict read–once resolution: a resolution tree is *literal–once* if it does not contain two or more vertices whose clauses are inferred by resolving on the same literal. For example, the resolution refutation depicted in Figure 1 is strict read–once, but it is not literal–once, since clauses at two positions are inferred by resolving on the same literal x. However, it is easy to see that every literal–once resolution refutation is a (strict) read-once resolution. The main result of this paper is the intractability of literal–once resolution; i.e., it is NP-complete to recognize CNF-formulas which are refutable by literal–once resolution.

Furthermore, we show that intractability of read–once resolution can be obtained as corollary of our main result. This fact may be of interest, since Iwama and Miyano obtain the quoted result solely by presenting a single example without giving an accurate proof.

In [11] Kullmann asked for the computational complexity of finding a subset F' of a given formula F such that

(i) F' is minimal unsatisfiable (F' is unsatisfiable, but every proper subset of F' is satisfiable), and

(ii) F' has exactly one more clause than variables.

We denote by MU(1) the class of formulas F' satisfying (i) and (ii). This class is of special interest; for example, every minimal unsatisfiable Horn formula belongs to MU(1) ([4]). We show that F has a subset $F' \in$ MU(1) if and only if F is refutable by literal–once resolution. Whence the intractability of Kullmann's problem follows from the NP–completeness of refutability by literal–once resolution.

2 Notation

2.1 Digraphs

We denote a digraph D by an ordered pair (V, A) consisting on a finite nonempty set V of *vertices* and a set A of *arcs*; an arc is an ordered pair (u, v) of distinct vertices $u, v \in V$. Let $D = (V, A)$ be a digraph and $v \in V$. We denote the sets of *incoming* and *outgoing* arcs of v by $\mathsf{out}(v) = \{ (u, w) \in A \mid u = v \}$ and $\mathsf{in}(v) = \{ (u, w) \in A \mid w = v \}$, respectively. For $(u, v) \in A$ we say that u is a *predecessor of* v and that v *is a successor of* u.

A digraph $T = (V, A)$ is an *in–tree* if there is exactly one vertex v without successors (the *root* of T), and for every vertex $w \in V$ there is exactly one (directed) path P_w from w to v. Consequently, every vertex which is different from the root has exactly one successor. A vertex without predecessors is a *leaf.* An in–tree T is *binary* if every non-leaf has exactly two predecessors. Note that a binary in–tree with k leaves has $2k - 1$ vertices. For graph theoretic terminology not defined here, the reader is referred to [2].

2.2 CNF-Formulas

Let var be a set of boolean variables. A *literal* ℓ is an object of the form x or $\overline{x}$ for $x \in \mathsf{var}$; in the first case we call ℓ *positive*, in the second case *negative*; for a negative literal $\ell = \overline{x}$, $x \in \mathsf{var}$, we put $\overline{\ell} = x$. Literals ℓ and $\overline{\ell}$ are *complements* of each other. If x is a variable and $\ell \in \{x, \overline{x}\}$, then we call x the *variable of* ℓ and write $\mathsf{var}(\ell) = x$. A *clause* is a finite set of literals without complements. The empty clause is denoted by $\Box$. For a clause C we put $\mathsf{var}(C) = \{ \mathsf{var}(\ell) \mid \ell \in C \}$. A *CNF-formula* (or *formula*, for short) is a finite set of clauses. For a formula F we put $\mathsf{var}(F) = \bigcup_{C \in F} \mathsf{var}(C)$. A literal ℓ is a *pure literal* of F if $\ell \in \bigcup_{C \in F} C \not\ni \overline{\ell}$. A formula F is *Horn* if every clause in F contains at most one positive literal.

A *truth assignment* t to a formula F is a map $t : \mathsf{var}(F) \to \{0, 1\}$. Let t be a truth assignment to F; we put $t(\overline{x}) = 1 - t(x)$ for $x \in \mathsf{var}(F)$, and we say that t *satisfies* a clause $C \in F$ if $t(\ell) = 1$ for at least one literal $\ell \in C$. Furthermore, we say that t *satisfies* F if t satisfies all clauses of F. A formula F is *satisfiable* if there is a truth assignment which satisfies F; otherwise F is called *unsatisfiable.* We denote the set of all unsatisfiable formulas by UNSAT.

2.3 Resolution

Let C_1, C_2 be two clauses. If there is *exactly one* literal ℓ such that $\ell \in C_1$ and $\overline{\ell} \in C_2$ then we call the clause $C = (C_1 \setminus \{\ell\}) \cup (C_2 \setminus \{\overline{\ell}\})$ the *resolvent of* C_1 *and* C_2; in this case we also say that C is obtained from C_1, C_2 by *resolving on* ℓ.

Let $T_0 = (V, A)$ be an in–tree and λ a labeling of its vertices such that $\lambda(v)$ is a clause for every $v \in V$. We call $T = (V, A, \lambda)$ a *resolution tree* if for every vertex $v \in V$ with predecessors v_1, v_2 it holds that $\lambda(v)$ is the resolvent of $\lambda(v_1)$ and $\lambda(v_2)$. Let $T = (V, A, \lambda)$ be a resolution tree and $v \in V$. If v is a leaf, then we put $\mathsf{rlit}(v) = \emptyset$; otherwise v has two predecessors, say v_1 and v_2; we put

$\mathsf{rlit}(v) = (\lambda(v_1) \cup \lambda(v_2)) \setminus \lambda(v)$. We call the elements of $\mathsf{rlit}(v)$ *resolution literals of* v. A clause C is a *premise* of a resolution tree T if $\lambda(v) = C$ for some leaf v of T. We write $\mathsf{pre}(T)$ for the set of all premises of T. A clause C is the *conclusion* of T if $\lambda(v) = C$ for the root v of T; in this case we write $\mathsf{con}(T) = C$. A resolution tree T is a *resolution refutation* if $\mathsf{con}(T) = \Box$. Let F be a formula and T a resolution refutation. If $\mathsf{pre}(T) \subseteq F$ then we say that F is *refuted* by T, or that T is a resolution refutation *of* F. A resolution tree $T = (V, A, \lambda)$ is *trivial* if $|V| = 1$. Clearly, a formula F is refuted by the trivial resolution tree $T = (\{v\}, \emptyset, \lambda)$ if and only if $\lambda(v) = \Box \in F$.

For a resolution tree $T = (V, A, \lambda)$ and $v \in V$ we define T_v to be the resolution tree (V', A', λ') where (V', A') is the maximal subtree of (V, A) with root v and λ' is the restriction of λ to V'.

It is well–known that a formula F is unsatisfiable if and only if it can be refuted by some resolution refutation T.

3 Restricted Types of Resolution

Read–Once Resolution. A resolution tree $T = (V, A, \lambda)$ is *read–once* if $\lambda(v) \neq \lambda(w)$ for any two distinct leaves v, w of T. We denote by ROR the class of all formulas refutable by read–once resolution refutations. (ROR corresponds to the class which is denoted by $R(0)$ in [8].)

Strict Read–Once Resolution. A resolution tree $T = (V, A, \lambda)$ is *strict read–once* if $\lambda(v) \neq \lambda(w)$ for any two distinct vertices v, w of T. We denote by SROR the class of all formulas refutable by strict read–once resolution refutations.

Literal–Once Resolution. A resolution tree $T = (V, A, \lambda)$ is *literal–once* if $\mathsf{rlit}(v) \neq \mathsf{rlit}(w)$ for any two distinct non–leaves v, w of T. We denote by LOR the class of all formulas refutable by literal–once resolution refutations.

Proposition 1 LOR $\subsetneq$ SROR $\subsetneq$ ROR $\subsetneq$ UNSAT.

Proof. If a resolution refutation is literal–once, then it is obviously strict read–once; thus LOR $\subseteq$ SROR. Consider the formula $F = \{\{x, y\}, \{x, \overline{y}\}, \{\overline{x}, y\}, \{\overline{x}, \overline{y}\}\}$. Figure 1 shows a strict read-once resolution refutation T of F, hence $F \in$ SROR. (We note in passing that F belongs to a subclass of minimal unsatisfiable formulas characterized in [9].) However, T is not literal-once. It is easy to see that there is no literal–once resolution refutation of F at all. Whence LOR $\subsetneq$ SROR follows.

We have SROR $\subseteq$ ROR by definition. Consider the formula $F = \{C_1, \ldots, C_5\}$ with

$$C_1 = \{x, \overline{z}\}, \quad C_4 = \{x, y, z\},$$
$$C_2 = \{\overline{x}, y\}, \quad C_5 = \{\overline{x}, \overline{y}, \overline{z}\},$$
$$C_3 = \{\overline{y}, z\}.$$

Figure 2 exhibits a resolution refutation of F, hence $F \in$ UNSAT. We show that $F \notin$ ROR. Consider a resolution refutation T of F with root v, and let v_1, v_2 the predecessors of v. Clearly $|\mathsf{con}(T_{v_1})| = |\mathsf{con}(T_{v_2})| = 1$. However, no pair of clauses $C', C'' \in F$ have a resolvent C with $|C| = 1$. Thus $|\mathsf{pre}(T_{v_1})|, |\mathsf{pre}(T_{v_2})| \geq 3$. Since $|F| = 5$ it follows that $\mathsf{pre}(T_{v_1}) \cap \mathsf{pre}(T_{v_2}) \neq \emptyset$. Consequently, T is not read–once. Hence $F \notin$ ROR and so ROR $\neq$ UNSAT.

Let $W_1 = \{w, x, \overline{z}\}$, $W_2 = \{\overline{w}, x, \overline{z}\}$, and consider $F^* = F \cup \{W_1, W_2\}$. Observe that C_1 is the resolvent of W_1 and W_2. The resolution tree exhibited in Figure 3 shows that $F^* \in$ ROR. Consider a read–once resolution refutation T of F^*. We show that T is not strict read–once. Again, let v_1, v_2 be the predecessors of the root of T. W.l.o.g., we assume $|\mathsf{pre}(T_{v_1})| \leq |\mathsf{pre}(T_{v_2})|$. Similarly as above, $|\mathsf{pre}(T_{v_1})|, |\mathsf{pre}(T_{v_2})| \geq 3$ follows. Since T is assumed to be read–once, $|\mathsf{pre}(T_{v_1})| + |\mathsf{pre}(T_{v_2})| \leq |F^*|$; thus $|\mathsf{pre}(T_{v_1})| = 3$. It can be verified that there is no resolution tree T' with $\mathsf{pre}(T') \subseteq F^*$, $|\mathsf{pre}(T')| = 3$ and $|\mathsf{con}(T')| = 1$, such that $W_1 \in \mathsf{pre}(T')$ or $W_2 \in \mathsf{pre}(T')$. However, $W_1, W_2 \in \mathsf{pre}(T)$ since $F \notin$ ROR. It follows that $W_1, W_2 \in \mathsf{pre}(T_{v_2})$ and $|\mathsf{pre}(T_{v_2})| = 4$. Hence we have $\mathsf{pre}(T_{v_2}) = \{W_1, W_2, D_1, D_2\}$ for some $D_1, D_2 \in \{C_2, \ldots, C_5\}$. Checking all possibilities for D_1, D_2 shows that either $\{D_1, D_2\} = \{C_2, C_4\}$ or $\{D_1, D_2\} = \{C_3, C_5\}$. In both cases, the two vertices u_1, u_2 of T_{v_2} which are labeled by W_1 and W_2, respectively, have a common successor u. Evidently u is labeled by C_1. Since $C_1 \in \mathsf{pre}(T)$, it follows that T is not strict read–once. Whence SROR $\neq$ ROR. $\square$

4 NP-Completeness Results

Let F be a formula with m clauses and $T = (V, A, \lambda)$ a read–once (strict read–once, literal–once, respectively) resolution refutation of F. Clearly T has at most m leaves, and so $|V| \leq 2m - 1$. Thus one can guess such resolution refutation T of F and verify in deterministic polynomial time whether T is indeed read–once (strict read–once, literal–once, respectively). Hence the following holds.

Lemma 1 *The recognition problems for* LOR, SROR, *and* ROR *are in* NP.

Next we state our main result whose proof we present in Section 6.

Theorem 1 *Recognition of* LOR *is* NP*-complete.*

We are going to show that recognition of ROR and recognition of SROR are both NP-complete problems as well. We proceed by reducing recognition of LOR to recognition of SROR and ROR, respectively. For these reductions, the following construction is crucial.

Let F be a formula. For each $x \in \mathsf{var}(F)$ we take two new variables $x[1]$, $x[2]$, and for every clause $C \in F$ we define

$$C^\circ = \{\, \overline{x[1]} \mid \overline{x} \in C \,\} \cup \{\, \overline{x[2]} \mid x \in C \,\}.$$

We put

$$F^\circ = \{\, C^\circ \mid C \in F \,\} \cup \{\, \{x[1], x[2]\} \mid x \in \mathsf{var}(F) \,\}.$$

Observe that F° is satisfiable if and only if F is satisfiable; furthermore, for every $x[i] \in \mathsf{var}(F^\circ)$ there is exactly one clause $C \in F^\circ$ with $x[i] \in C$.

The following result is a direct consequence of Lemmas 4, 5, and 7, which are more technical and will be presented in the Appendix.

Proposition 2 *For every formula F the following statements are equivalent.*

$$F \in \mathrm{LOR}; \quad F^{\circ} \in \mathrm{ROR}; \quad F^{\circ} \in \mathrm{SROR}.$$

The next two results follow from Theorem 1 and Proposition 2.

Theorem 2 *Recognition of* SROR *is* NP*-complete.*

Theorem 3 (Iwama and Miyano [8]) *Recognition of* ROR *is* NP*-complete.*

5 Literal–Once resolution and Minimal Unsatisfiable Formulas

In this section we apply Theorem 1 to answer a question posed by Kullmann ([11]). A formula F is *minimal unsatisfiable* if F is unsatisfiable but $F \setminus \{C\}$ is satisfiable for every $C \in F$. The *deficiency* $\delta(F)$ of a formula F is defined by

$$\delta(F) = |F| - |\mathsf{var}(F)|\,.$$

Let k be an integer; we write $\mathrm{MU}(k)$ for the class of minimal unsatisfiable formulas F with $\delta(F) = k$. By a result due to Tarsi ([1]), $\mathrm{MU}(k) = \emptyset$ for $k \leq 0$. Recognition of minimal unsatisfiable formulas is D^P-complete ([12]); however, for every fixed k, the class $\mathrm{MU}(k)$ can be recognized in polynomial time ([11, 5]). In [11], Kullmann asked whether recognizing

$$\mathcal{C} = \{\ F \mid \text{there is some } F' \subseteq F \text{ with } F' \in \mathrm{MU}(1)\ \}$$

is NP-complete. We answer this question positively: in the next lemma we show $\mathcal{C} = \mathrm{LOR}$; hence NP-completeness of $\mathcal{C}$ follows from Theorem 1.

Proposition 3 *Let F be a formula. Then $F \in \mathrm{MU}(1)$ if and only if there is a literal–once resolution refutation T with $\mathsf{pre}(T) = F$. Consequently $\mathrm{LOR} = \mathcal{C}$.*

Proof. We apply the following results from [4].

(i) If $F \in \mathrm{MU}(1)$ and $F \neq \square$ then there is a literal ℓ and clauses $C_1, C_2 \in F$ such that C_1 is the only clause of F containing ℓ; C_2 is the only clause of F containing $\overline{\ell}$.

(ii) Let F be a formula and ℓ a literal such that there are unique clauses $C_1, C_2 \in F$ with $\ell \in C_1$ and $\overline{\ell} \in C_2$; let $C_{1,2}$ be the resolvent of C_1 and C_2. Then $F \in \mathrm{MU}(1)$ if and only if $(F \setminus \{C_1, C_2\}) \cup \{C_{1,2}\} \in \mathrm{MU}(1)$.

We proceed by induction on $|F|$. The proposition evidently holds if $|F| = 1$; hence consider $|F| > 1$. Assume $F \in \mathrm{MU}(1)$ and choose ℓ, C_1, and C_2 according to (i). It follows now from (ii) that $F^* = (F \setminus \{C_1, C_2\}) \cup \{C_{1,2}\} \in \mathrm{MU}(1)$.

By induction hypothesis, there is a literal–once resolution refutation T^* with $C_{1,2} \in \mathsf{pre}(T^*) = F^*$. We extend T^* to a a literal–once resolution refutation T with $\mathsf{pre}(T) = F$ by adding leaves v_1, v_2 (labeled by C_1 and C_2, respectively) to T^*.

Conversely, assume that there is a literal–once resolution refutation $T = (V, A, \lambda)$ with $\mathsf{pre}(T) = F$. Choose two leaves v_1, v_2 of T which have a common successor v. Put $C_i = \lambda(v_i)$, $i = 1, 2$ and $C_{1,2} = \lambda(v)$. Consequently, there is a literal ℓ such that $\ell \in C_1$ and $\overline{\ell} \in C_2$. Hence removing v_1 and v_2 from T yields a literal–once resolution refutation T^* with $\mathsf{pre}(T^*) = (F \setminus \{C_1, C_2\}) \cup \{C_{1,2}\}$; $\mathsf{pre}(T^*) \in \mathrm{MU}(1)$ by induction hypothesis. It follows now from (ii) that $F \in \mathrm{MU}(1)$. □

In [4] it is shown that every minimal unsatisfiable Horn formula belongs to MU(1). Since every unsatisfiable Horn formula contains a minimal unsatisfiable Horn formula, Proposition 3 implies the following.

Proposition 4 *Every unsatisfiable Horn formula is refutable by literal–once resolution.*

6 Proof of Theorem 1

This section is devoted to a proof of Theorem 1. We reduce 3-SAT to recognition of LOR (in fact we could reduce SAT as well, but we choose 3-SAT to keep notation simpler). In a first step we reduce 3-SAT to the problem of finding a "satisfying path" in a digraph D, i.e., a path which does not run through prescribed pairs of vertices. In a second step we mimic this path problem by constructing a formula F such that literal–once resolution refutations of F and satisfying paths of D correspond to each other.

First we prove two short lemmas which we will need below.

Lemma 2 *Let T be a literal–once resolution tree and $C_1, C_2 \in \mathsf{pre}(T)$ with $\ell \in C_1$ and $\overline{\ell} \in C_2$ such that $\mathsf{rlit}(v) = \{\ell, \overline{\ell}\}$ for the root of T. Then $C_1 \cap C_2 \subseteq \mathsf{con}(T)$.*

Proof. Let v be the root of T and v_1, v_2 the predecessors of v. Consider $\ell' \in C_1 \cap C_2$. Since T is literal–once, it follows that ℓ' cannot be an element of both $\mathsf{rlit}(T_{v_1})$ and $\mathsf{rlit}(T_{v_2})$. Hence $\ell' \in \lambda(v) = \mathsf{con}(T)$. □

Lemma 3 *Let $T = (V, A, \lambda)$ be a literal–once resolution refutation and $C_1, C_2 \in \mathsf{pre}(T)$. Then there cannot be distinct literals $\ell, \ell' \in C_1$ such that $\overline{\ell}, \overline{\ell'} \in C_2$.*

Proof. We observe that there are vertices $v, v_1, v_2 \in V$ such that v_1, v_2 are predecessors of v and $\ell \in \mathsf{rlit}(v)$. W.l.o.g., assume $\ell \in \lambda(v_1)$ and $\overline{\ell} \in \lambda(v_2)$. It follows that $C_1 \in \mathsf{pre}(T_{v_1})$ and $C_2 \in \mathsf{pre}(T_{v_2})$. Since $\mathsf{rlit}(T_{v_1}) \cap \mathsf{rlit}(T_{v_2}) = \emptyset$, ℓ is the only literal with $\ell \in C_1$ and $\overline{\ell} \in C_2$. □

Construction I Let $F_3 = \{C_1, \ldots, C_n\}$ be a formula with $C_i = \{\ell_{i,1}, \ell_{i,2}, \ell_{i,3}\}$ for $1 \leq i \leq n$. We write L for the set of literals ℓ such that $\mathsf{var}(\ell) \in \mathsf{var}(F_3)$. Further, for $\ell \in L$ we put

$$q(\ell) = \{\ i \mid \ell \in C_i,\ 1 \leq i \leq n\ \}.$$

Observe that $i \notin q(\overline{\ell})$ for every $\ell \in C_i$, $1 \leq i \leq n$, since clauses do not contain complementary pairs of literals. We assume w.l.o.g. that F_3 has no pure literals; i.e., $|q(\ell)| \geq 1$ for every $\ell \in L$.

We construct a digraph $D = (V, A)$ as follows. We take a set of $n+1$ vertices $\{u_0, \ldots, u_n\}$, and for $i = 1, \ldots, n$ we join u_{i-1} and u_i by three (directed) paths $P_{i,1}$, $P_{i,2}$, $P_{i,3}$ of length $\left|q(\overline{\ell_{i,1}})\right| + 1$, $\left|q(\overline{\ell_{i,2}})\right| + 1$, $\left|q(\overline{\ell_{i,3}})\right| + 1$, respectively. We denote the set of inner vertices of $P_{i,j}$ by $V_{i,j}$ $(1 \leq i \leq n,\ 1 \leq j \leq 3)$. Hence we have $|V_{i,j}| = \left|q(\overline{\ell_{i,j}})\right|$ for $0 \leq i \leq n$, $1 \leq j \leq 3$. Now we form a set S of pairs (v, v') of vertices $v, v' \in V \setminus \{u_0, \ldots, u_n\}$ such that

- there is a pair $(v, v') \in S$ with $v \in V_{i,j}$ and $v' \in V_{i',j'}$ $(1 \leq i < i' \leq n$, $1 \leq j, j' \leq 3)$ if and only $\ell_{i,j} = \overline{\ell_{i',j'}}$, and
- every vertex in $V \setminus \{u_0, \ldots, u_n\}$ is contained in exactly one pair of S.

Note that such set S exists and can be obtained efficiently. We call a directed path in D *satisfying* if it runs from u_0 to u_n and contains at most one vertex of each pair in S. Observe that each satisfying path has to pass through all of the vertices $u_0, \ldots, u_n$ in increasing order.

Claim 1 *F_3 is satisfiable if and only if D has a satisfying path.*

Proof. If F_3 is satisfied by some truth assignment t, then we can choose $\sigma(i) \in \{1, 2, 3\}$ for $0 \leq i \leq n$ such that $t(\ell_{i,\sigma(i)}) = 1$. We observe that

$$P = P_{0,\sigma(0)} \ldots P_{n,\sigma(n)} \qquad (1)$$

is a satisfying path. Conversely, by definition, every satisfying path P is of the form (1) for some $\sigma : \{0, \ldots, n\} \to \{1, 2, 3\}$. Thus, if P is a satisfying path, then putting $t(\ell_{i,\sigma(i)}) = 1$ for $0 \leq i \leq n$ induces a truth assignment t which satisfies F_3. □

Note that the above construction is closely related to the *connection method* (see, e.g., [13, 3, 10]).

Construction II Let $D = (V, A)$ be the digraph obtained from a given 3-CNF formula F_3 according to Construction I. We consider a portion of distinct boolean variables: for $0 \leq i \leq n$ we take a new variable ν_i; for each arc $a \in A$ we take a new variable α_a; for each pair $p \in S$ we take three distinct new variables $\beta_p, \gamma_p, \delta_p$. We define a formula F with

$$\mathsf{var}(F) = \{\nu_0, \ldots, \nu_n\} \cup \{\ \alpha_a \mid a \in A\ \} \cup \{\ \beta_p, \gamma_p, \delta_p \mid p \in S\ \}$$

by

$$F = \{\{\nu_0\}, \{\overline{\nu_n}\}\} \cup \bigcup_{v \in V} F(v)$$

and the following definitions (recall that $\mathsf{in}(v)$ and $\mathsf{out}(v)$ denote the sets of arcs incoming to and outgoing from v, respectively). For $0 \leq i \leq n$ let

$$\begin{aligned} F(u_i) = \{\ \{\overline{\alpha_a}, \nu_i\} \mid a \in \mathsf{in}(u_i)\ \} \cup \\ \{\ \{\alpha_b, \overline{\nu_i}\} \mid b \in \mathsf{out}(u_i)\ \}. \end{aligned} \tag{2}$$

For $p = (v, v') \in S$ with

$$\mathsf{in}(v) = \{a\}, \quad \mathsf{out}(v) = \{b\}, \quad \mathsf{in}(v') = \{a'\}, \quad \mathsf{out}(v') = \{b'\} \tag{3}$$

we put

$$\begin{aligned} F(v) = \{&\{\overline{\alpha_a}, \beta_p, \gamma_p\}, \\ &\{\overline{\beta_p}, \gamma_p\}, \\ &\{\overline{\gamma_p}, \overline{\delta_p}\}, \\ &\{\alpha_b, \overline{\gamma_p}, \delta_p\}\} \end{aligned} \qquad \text{and} \qquad \begin{aligned} F(v') = \{&\{\overline{\alpha_{a'}}, \beta_p, \overline{\gamma_p}\}, \\ &\{\overline{\beta_p}, \overline{\gamma_p}\}, \\ &\{\gamma_p, \overline{\delta_p}\}, \\ &\{\alpha_{b'}, \gamma_p, \delta_p\}\}, \end{aligned}$$

and write $F(p) = F(v) \cup F(v')$.

Claim 2 *Let $T = (V, A, \lambda)$ be a literal–once resolution refutation of F and $p = (v, v') \in S$. If $F(p) \cap \mathsf{pre}(T) \neq \emptyset$ then either $F(p) \cap \mathsf{pre}(T) = F(v)$ or $F(p) \cap \mathsf{pre}(T) = F(v')$.*

Proof. Let $a, a', b, b' \in A$ according to (3). We use the shorthands

$$\begin{aligned} C_1 &= \{\overline{\alpha_a}, \beta_p, \gamma_p\}, & C_1' &= \{\overline{\alpha_{a'}}, \beta_b, \overline{\gamma_p}\}, \\ C_2 &= \{\overline{\beta_p}, \gamma_p\}, & C_2' &= \{\overline{\beta_p}, \overline{\gamma_p}\}, \\ C_3 &= \{\overline{\delta_p}, \overline{\gamma_p}\}, & C_3' &= \{\overline{\delta_p}, \gamma_p\}, \\ C_4 &= \{\alpha_b, \overline{\gamma_p}, \delta_p\}, & C_4' &= \{\alpha_{b'}, \gamma_p, \delta_p\} \end{aligned}$$

so that $F(v) = \{C_1, \ldots, C_4\}$ and $F(v') = \{C_1', \ldots, C_4'\}$. First we show

$$\{C_1, C_1'\} \not\subseteq \mathsf{pre}(T). \tag{4}$$

Suppose to the contrary that $\{C_1, C_1'\} \subseteq \mathsf{pre}(T)$. Consequently, there is some $v \in V$ such that $\gamma_p \in \mathsf{rlit}(v)$. Thus $C_1, C_1' \in \mathsf{pre}(T_v)$. By Lemma 2 it follows that $\beta_p \in \lambda(v)$. Hence there must be a clause $C \in \mathsf{pre}(T) \setminus \mathsf{pre}(T_v)$ with $\overline{\beta_p} \in C$. By construction of F, C_2 and C_2' are the only clauses of F which contain $\overline{\beta_p}$. Observe that $\gamma_p \in C_2$ and $\overline{\gamma_p} \in C_2'$. Thus $\gamma_p \in \mathsf{rlit}(T)$, since $\mathsf{con}(T) = \Box$. It follows that $\gamma_p \in \mathsf{rlit}(T) \setminus \mathsf{rlit}(T_v)$. However, $\gamma_p \in \mathsf{rlit}(T_v)$, and therefore we have a contradiction to the assumption T being literal–once. Whence (4) holds. By analogous arguments one can show

$$\begin{aligned} \{C_4, C_4'\} &\not\subseteq \mathsf{pre}(T), \\ \{C_2, C_2'\} &\not\subseteq \mathsf{pre}(T), \\ \{C_3, C_3'\} &\not\subseteq \mathsf{pre}(T). \end{aligned} \tag{5}$$

We show that

$$C_1 \in \mathsf{pre}(T) \quad \Leftrightarrow \quad C_2 \in \mathsf{pre}(T). \tag{6}$$

Assume $C_1 \in \mathsf{pre}(T)$. Since $\beta_p \in C_1$, there must be a clause $C \in \mathsf{pre}(T)$ with $\overline{\beta_p} \in C$; C_2 and C_2' are the only clauses of F which contain $\overline{\beta_p}$. By Lemma 3 we conclude that $C_2' \notin \mathsf{pre}(T)$; thus $C_2 \in \mathsf{pre}(T)$. Whence we have shown one direction of (6). The converse can be shown similarly applying Lemma 3. Moreover, one can show by analogous arguments that

$$\begin{array}{lcl} C_1' \in \mathsf{pre}(T) & \Leftrightarrow & C_2' \in \mathsf{pre}(T), \\ C_3 \in \mathsf{pre}(T) & \Leftrightarrow & C_4 \in \mathsf{pre}(T), \\ C_3' \in \mathsf{pre}(T) & \Leftrightarrow & C_4' \in \mathsf{pre}(T). \end{array} \tag{7}$$

Finally we observe that

$$\mathsf{pre}(T) \cap \{C_1, C_2, C_3', C_4'\} \neq \emptyset \quad \Leftrightarrow \quad \mathsf{pre}(T) \cap \{C_1', C_2', C_3, C_4\} \neq \emptyset. \tag{8}$$

Claim 2 now follows from (4)–(8). ☐

Claim 3 *D has a satisfying path if and only if $F \in$ LOR.*

Proof. Assume that D has a satisfying path P. We denote by $V(P)$ and $A(P)$ the vertices and arcs of P, respectively. For $0 \leq i \leq n$ we put

$$\begin{array}{rl} F_P(u_i) = & \{\ \{\overline{\alpha_a}, \nu_i\} \mid a \in \mathsf{in}(u_i) \cap A(P)\ \} \cup \\ & \{\ \{\alpha_b, \overline{\nu_i}\} \mid b \in \mathsf{out}(u_i) \cap A(P)\ \} \end{array}$$

and for $v \in V(P) \setminus \{u_0, \ldots, u_n\}$ we put $F_P(v) = F(v)$. We show that

$$F(P) = \{\{\nu_0\}, \{\overline{\nu_n}\}\} \cup \bigcup_{v \in V(P)} F_P(v)$$

can be refuted by literal–once resolution (observe that $F(P) \subseteq F$). Consider a vertex $v \in V(P)$ with $p = (v, v') \in S$. Using the same notation as in the proof of Claim 2, we have $F(v) = \{C_1, C_2, C_3, C_4\} \subseteq F(P)$. Now $C_{1,2} = \{\overline{\alpha_a}, \gamma_p\}$ is a resolvent of C_1 and C_2; $C_{3,4} = \{\alpha_b, \overline{\gamma_p}\}$ is a resolvent of C_3 and C_4. Further, $C_v = \{\overline{\alpha_a}, \alpha_b\}$ is a resolvent of $C_{1,2}$ and $C_{3,4}$. Hence finding a literal–once resolution refutation of $F(P)$ reduces to finding a literal–once resolution refutation of $(F(P) \setminus F(v)) \cup \{\{\overline{\alpha_a}, \alpha_b\}\}$. Similarly, if $v' \in V(P)$ with $p = (v, v') \in S$, then it suffices to find a literal–once resolution refutation of $(F(P) \setminus F(v')) \cup \{\{\overline{\alpha_{a'}}, \alpha_{b'}\}\}$. By multiple applications of this argument, $F(P)$ can be reduced to a formula of the form

$$F_{\text{lin}} = \{\{\ell_1\}, \{\overline{\ell_1}, \ell_2\}, \{\overline{\ell_2}, \ell_3\}, \ldots, \{\overline{\ell_{r-1}}, \ell_r\}, \{\overline{\ell_r}\}\}.$$

It is easy to construct a literal–once resolution refutation T_{lin} for F_{lin}. Now T_{lin} can be extended by the above considerations to a literal–once resolution refutation of $F(P)$. Whence $F \in$ LOR follows.

Conversely, assume that $F \in$ LOR. We show that D has a satisfying path. Let T be a literal–once resolution refutation of F and put $F' = \mathsf{pre}(T)$. Let W

be the set of vertices $w \in W$ such that there is at least one arc $a \in \mathsf{in}(w) \cup \mathsf{out}(w)$ with $\alpha_a \in \mathsf{var}(F')$. Clearly $W \neq \emptyset$. Since F' has no pure literals, it follows that for every $w \in W \setminus \{u_0, u_n\}$ there are arcs $a \in \mathsf{in}(w)$, $b \in \mathsf{out}(w)$ such that $\alpha_a, \alpha_b \in \mathsf{var}(F')$ (if $w = u_i$ for some $1 \leq i \leq n-1$ this is obvious; on the other hand, if w belongs to some pair in S, then it follows by Claim 2). Thus, for every $w \in W$, at least one predecessor and at least one successor belongs to W. Consider the subdigraph D_W of D induced by W. Clearly D_W is acyclic, since D is acyclic by construction. Every nonempty acyclic digraph has at least one vertex s without incoming arcs and at least one vertex t without outgoing arcs. For D_W the only possibility is $s = u_0$ and $t = u_n$. We conclude that D_W contains a path from u_0 to u_n. By Claim 2 it follows that for every $(v, v') \in S$ at most one of v, v' belongs to W. Thus P must be a satisfying path necessarily. This completes the proof of the claim. □

In view of Lemma 1, Theorem 1 now follows from Claims 1, 3, and the NP-completeness of the 3-SAT problem.

Appendix: Technical Lemmas

Lemma 4 $F \in \mathrm{LOR}$ *implies* $F^\circ \in \mathrm{LOR}$ *for every formula* F.

Proof. We show by induction on $|V|$ that for every literal–once resolution tree $T = (V, A, \lambda)$ there is a literal–once resolution tree T' with $\mathsf{pre}(T') = \mathsf{pre}(T)^\circ$, $\mathsf{con}(T') = \mathsf{con}(T)^\circ$, and $\mathsf{rlit}(T') = \{\ x[i], \overline{x[i]} \mid x \in \mathsf{var}(\mathsf{rlit}(T)),\ i = 1, 2\ \}$. If $|V| = 1$, then there is nothing to show. Assume $|V| > 1$ and let v be the root of T and x the variable in $\mathsf{rlit}(v)$. Moreover, let v_1, v_2 the predecessors of v such that $\overline{x} \in \lambda(v_1)$ and $x \in \lambda(v_2)$. For $i = 1, 2$ let T_i' be a literal–once resolution tree obtained from T_{v_i} as supplied by the induction hypothesis. Since $\mathsf{rlit}(T_{v_1}) \cap \mathsf{rlit}(T_{v_2}) = \emptyset$, it follows that $\mathsf{rlit}(T_1') \cap \mathsf{rlit}(T_2') = \emptyset$. Now $\overline{x[i]} \in \mathsf{con}(T_i') = \mathsf{con}(T_{v_i})^\circ$. It is obvious how T_1' and T_2' can be assembled to literal–once resolution tree T' with the desired properties by adding two non–leaves and a leaf w with $\lambda(w) = \{x[1], x[2]\}$. □

The following Lemma is due to an observation by Kullmann.

Lemma 5 $F^\circ \in \mathrm{ROR}$ *implies* $F^\circ \in \mathrm{LOR}$ *for every formula* F.

Proof. Observe that for every resolution tree $T = (V, A, \lambda)$ and two distinct vertices $v, v' \in V$ with $\mathsf{rlit}(v) = \mathsf{rlit}(v') = \{x, \overline{x}\}$, there must be at least four distinct leaves $u_1, u_2, u_1', u_2' \in V$ such that $x \in \lambda(u_1) \cap \lambda(u_1')$ and $\overline{x} \in \lambda(u_2) \cap \lambda(u_2')$. (Every vertex $v \in V$ with $\mathsf{rlit}(v) = \{x, \overline{x}\}$ "consumes" at least one leaf u_1 with $x \in \lambda(u_1)$ and one leaf u_2 with $x \in \lambda(u_2)$.) However, for every variable $x[i]$ of F° there is exactly one clause $C \in F^\circ$ such that $x[i] \in C$. Hence every read–once resolution refutation of F° is literal–once. □

Lemma 6 *Let F be a formula and T a resolution refutation with* $\mathsf{pre}(T) \subseteq F^\circ$. *Then* $\mathsf{pre}(T) = F_1^\circ$ *for some* $F_1 \subseteq F$.

Proof. Follows from the fact that $\mathsf{pre}(T)$ has no pure literals. □

Lemma 7 $F^\circ \in \mathrm{LOR}$ *implies* $F \in \mathrm{LOR}$ *for every formula* F.

Proof. We show by induction on $|F|$ that if T is a literal–once resolution refutation with $\mathsf{pre}(T) = F^\circ$, then there is a literal–once resolution refutation T' with $\mathsf{pre}(T') = F$; the lemma will follow by Lemma 6. If $|F| = 1$, then $F = F^\circ = \{\Box\}$, and the result follows by taking $T' = T$. Now assume $|F| > 1$ and let T be a literal–once resolution refutation with $\mathsf{pre}(T) = F^\circ$. We call a vertex v' of T *mistimed* if there is a predecessor v_1 of v' with $\lambda(v_1) = \{x[1], x[2]\}$, $x \in \mathsf{var}(F)$, and a successor v of v' such that $\mathsf{rlit}(v) \cap \{x[1], x[2]\} = \emptyset$. Mistimed vertices can be successively eliminated as follows (roughly speaking, we shift leaves labeled by clauses of the form $\{x[1], x[2]\}$ towards the root). Consider a mistimed vertex v' of T with predecessors v_1 and v_2 such that $\lambda(v_2) = \{x[1], x[2]\}$, $x \in \mathsf{var}(F)$. Let v be the successor of v' such that v' and v'' are the predecessors of v. We remove the arcs (v_1, v') and (v'', v) from T and add instead the arcs (v_1, v) and (v'', v'). Clearly $\lambda(v')$ and $\lambda(v)$ can be modified appropriately such that the result is still a read–once resolution refutation with same set of premises. Hence we can assume, w.l.o.g., that T has no mistimed vertices.

We write L_1 for the set of leaves v of T with $\lambda(v) = C^\circ$ for some $C \in F$, and we write L_2 for the set of leaves of T not in L_1 (i.e., if $v \in L_2$, then $\lambda(v) = \{x[1], x[2]\}$ for some $x \in \mathsf{var}(F)$). Observe that for any two leaves v_1, v_2 of T which have the same successor, either $v_1 \in L_1$ and $v_2 \in L_2$, or vice versa. Therefore, if T is nontrivial, then the height of T (i.e., the length of a longest path in T) is at least 2.

We choose a vertex v of T such that T_v has height 2. Since T has no mistimed vertices by assumption, we conclude that exactly one leaf of T_v is in L_1. Hence v has two predecessors v' and v'' such that v' has two predecessors $v_1 \in L_1$ and $v_2 \in L_2$, and $v'' \in L_1$. Let $Q, R \in F$ and $x \in \mathsf{var}(F)$ such that $\lambda(v_2) = \{x[1], x[2]\}$, $\lambda(v_1) = Q^\circ$, and $\lambda(v'') = R^\circ$. It follows for $\{i, j\} = \{1, 2\}$ that $x[i] \in \mathsf{rlit}(v')$ and $x[j] \in \mathsf{rlit}(v)$. Observe that $x[i] \notin \mathsf{var}(R^\circ)$; otherwise there would be a leaf $v_2^* \neq v_2$ with $\lambda(v_2^*) = \lambda(v_2)$. We conclude that $x[1], x[2] \notin \mathsf{var}(\lambda(v))$. Thus Q and R have a resolvent C with $\lambda(v) = C^\circ$. Let T_0 be the resolution tree obtained from T by removing v_1, v_2, v', v''. We have

$$\mathsf{pre}(T_0) = (\mathsf{pre}(T) \setminus \{\{x[1], x[2]\}, Q^\circ, R^\circ\}) \cup \{C^\circ\}.$$

Clearly T_0 is literal–once, hence the induction hypothesis applies. Thus, there is a literal–once resolution refutation T_0' with $\mathsf{pre}(T_0')^\circ = \mathsf{pre}(T_0)$; in particular, $C \in \mathsf{pre}(T_0')$. Let w be the leaf of T_0' labeled by C. It is now obvious how a literal–once resolution refutation $T' = (V', A', \lambda')$ can be obtained from T_0': we add two vertices w_1, w_2, the arcs (w_1, w), (w_2, w) to T_0', and we put $\lambda'(w_1) = Q$ and $\lambda'(w_2) = R$. Hence the lemma follows. □

References

1. R. Aharoni and N. Linial. Minimal non-two-colorable hypergraphs and minimal unsatisfiable formulas. *J. Combin. Theory Ser. A*, 43:196–204, 1986.
2. J. Bang-Jensen and G. Gutin. *Digraphs: theory, algorithms, applications*. Springer Monographs in Mathematics. Springer Verlag, 2000.
3. W. Bibel. On matrices with connections. *Journal of the ACM*, 28(4):633–645, 1981.
4. G. Davidov, I. Davydova, and H. Kleine Büning. An efficient algorithm for the minimal unsatisfiability problem for a subclass of CNF. *Annals of Mathematics and Artificial Intelligence*, 23:229–245, 1998.
5. H. Fleischner, O. Kullmann, and S. Szeider. Polynomial-time recognition of minimal unsatisfiable formulas with fixed clause-variable difference. Submitted, 2000.
6. Z. Galil. On the complexity of regular resolution and the Davis and Putnam procedure. *Theoretical Computer Science*, 4:23–46, 1977.
7. A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–308, 1985.
8. K. Iwama and E. Miyano. Intractability of read-once resolution. In *Proceedings of the Tenth Annual Structure in Complexity Theory Conference*, pages 29–36, Minneapolis, Minnesota, 1995. IEEE Computer Society Press.
9. H. Kleine Büning. On subclasses of minimal unsatisfiable formulas. To appear in Discr. Appl. Math., 2001.
10. H. Kleine Büning and T. Lettman. *Propositional logic: deduction and algorithms*. Cambridge University Press, Cambridge, 1999.
11. O. Kullmann. An application of matroid theory to the SAT problem. In *Fifteenth Annual IEEE Conference of Computational Complexity*, pages 116–124, 2000.
12. C. H. Papadimitriou and D. Wolfe. The complexity of facets resolved. *J. Comput. Syst. Sci.*, 37(1):2–13, 1988.
13. D. Prawitz. Advances and problems in mechanical proof procedures. In *Machine Intelligence*, volume 4, pages 59–71. American Elsevier, New York, 1969.
14. J. A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, 1965.