

Des réfutations SAT aux réfutations Max-SAT*

Matthieu Py[†] Mohamed Sami Cherif Djamal Habet

Aix-Marseille Université, Université de Toulon, CNRS, LIS, Marseille, France
 {matthieu.py, mohamed-sami.cherif, djamal.habet}@univ-amu.fr

Résumé

Adapter une preuve SAT par résolution en une preuve valide pour Max-SAT sans augmenter considérablement sa taille est une question longtemps restée ouverte. Cet article, qui résume les travaux publiés dans [3], contribue à y répondre en présentant des adaptations linéaires de réfutations SAT en réfutations Max-SAT, dans les cas *tree-like regular*, *tree-like* et *semi-tree-like*. On étend également ces résultats en proposant une adaptation complète pour toute réfutation SAT qui est exponentielle dans le pire des cas.

1 Introduction

Étant donnée une formule sous Forme Normale Conjonctive, le problème Max-SAT consiste à déterminer le nombre maximum de clauses qu'il est possible de satisfaire par une affectation des variables alors que le problème SAT consiste simplement à déterminer si la formule est satisfiable. Dans le contexte du problème SAT, une formule peut être démontrée insatisfiable à l'aide d'une séquence de résolutions [4], appelée réfutation par résolution, qui déduit de la formule initiale de nouvelles clauses jusqu'à en déduire la clause vide qui, par définition, est impossible à satisfaire. De même, pour Max-SAT, on utilise un système de preuve bien connu basé sur la règle de la max-résolution [2], qui étend la règle de résolution utilisée dans SAT. Les séquences de max-résolutions sont plus contraintes que les séquences de résolutions car la max-résolution remplace les prémisses par les conclusions, ce qui consomme les prémisses et empêche de les utiliser à nouveau. Adapter une réfutation SAT (par résolution) en une réfutation valide pour Max-SAT est par conséquent simple si la formule est *read-once*, c'est à dire si chaque clause est utilisée une seule fois comme prémisses d'une résolution. Dans ce cas, il suffit de remplacer chaque étape de

résolution par une max-résolution pour obtenir une réfutation valide pour Max-SAT dont la taille (exprimée en nombres d'étapes de transformation) est linéaire par rapport à la taille de la réfutation SAT [1]. En revanche, adapter une réfutation SAT en une réfutation Max-SAT dans le cas général et sans augmenter considérablement sa taille reste une question ouverte.

Dans cet article, on propose d'apporter une contribution pour répondre à cette question. Pour cela, on augmente la règle de la max-résolution avec la règle du *split*, permettant de remplacer une clause par deux clauses en y ajoutant un littéral supplémentaire. Ainsi, on est désormais capable d'adapter n'importe quelle réfutation SAT *tree-like regular* en une réfutation Max-SAT de taille linéaire. On étend ensuite ce résultat aux cas *tree-like* et *semi-tree-like*. Enfin, on propose une adaptation des réfutations SAT dans le cas général mais dont la taille de la réfutation Max-SAT obtenue peut être exponentielle dans le pire des cas.

2 Des réfutations *tree-like regular* aux réfutations Max-SAT

Considérons tout d'abord le cas des réfutations par résolution *tree-like regular*. Une réfutation est dite *tree-like* si aucune clause intermédiaire (déduite par une résolution) n'est utilisée plusieurs fois en tant que prémisses d'une étape de résolution. Une réfutation est dite *regular* si chacune de ses branches (séquences de résolutions commençant depuis des clauses de la formule et se terminant par la déduction de la clause vide) contient au plus une résolution par variable.

Exemple 1. $\phi = (\bar{x}_1 \vee x_3) \wedge (x_1) \wedge (\bar{x}_1 \vee x_2) \wedge (\bar{x}_2 \vee \bar{x}_3)$. La réfutation par résolution de ϕ représentée dans la figure 1 n'est pas *read-once* mais elle est *tree-like regular*.

*Cet article est un résumé de [3].

[†]Papier doctorant : Matthieu Py est auteur principal.




FIGURE 1 – Réfutation par résolution

Théorème 1. *Étant donnée une formule insatisfiable ϕ et une réfutation par résolution tree-like regular P de ϕ , il existe une réfutation Max-SAT de ϕ contenant $O(|P|)$ étapes d'inférence.*

Comme énoncé dans le théorème 1, il est possible d'adapter une réfutation *tree-like regular* en une réfutation Max-SAT de taille linéaire par rapport à la taille de la réfutation SAT. Pour cela, si une clause est utilisée plusieurs fois, on cherche le point de jonction de toutes les branches partant de cette clause. Ce point de jonction est une étape de résolution sur une variable telle que l'application de la règle du *split* sur la clause initiale et cette variable génère deux clauses qui permettent de la remplacer sans impacter la validité de la preuve. En répétant cette opération autant de fois que nécessaire, on retombe sur une réfutation que l'on peut qualifier de *read-once* et on peut remplacer chaque résolution par une max-résolution pour obtenir une réfutation Max-SAT de la formule de départ.

Exemple 2. *Dans la réfutation de l'exemple 1, la clause (x_1) est utilisée deux fois et le point de jonction des deux branches qui partent de (x_1) est une résolution sur la variable x_3 . On applique donc une étape de *split* sur la clause (x_1) et la variable x_3 pour obtenir deux nouvelles clauses permettant de remplacer (x_1) . On remplace enfin chaque résolution par une max-résolution et on obtient la réfutation Max-SAT de la figure 2.*

3 Extension aux cas tree-like, semi-tree-like et général

Le résultat linéaire sur les réfutations *tree-like regular* s'étend aux réfutations *tree-like* en utilisant un résultat connu permettant de transformer n'importe quelle réfutation *tree-like* en une réfutation *tree-like regular* plus petite [5]. Pour cela, à chaque fois que l'on détecte deux résolutions sur la même variable dans une branche, on supprime la première résolution ainsi




FIGURE 2 – Application de la règle *split* pour générer une réfutation Max-SAT

que toutes les autres résolutions devenues inapplicables. Ce résultat s'étend également au cas *semi-tree-like*, où chaque branche de la réfutation contient au plus une clause utilisée plusieurs fois. Dans ce cas-là, on découpe la réfutation en deux morceaux indépendants, le premier *read-once* et le second *tree-like* et on les traite de manière indépendante avant de les refusionner.

Théorème 2. *Étant donnée une formule insatisfiable ϕ et une réfutation par résolution semi-tree-like P de ϕ , il existe une réfutation Max-SAT de ϕ contenant $O(|P|)$ étapes d'inférence.*

Enfin, dans le cas général, il est possible de rendre la réfutation *semi-tree-like* moyennant une augmentation au pire exponentielle de sa taille. Pour cela, on duplique chaque portion de la preuve conduisant à une clause intermédiaire utilisée plusieurs fois afin d'en créer autant de copies que nécessaire, ce qui permet de rendre la preuve *semi-tree-like* et de revenir au cas précédent.

Théorème 3. *Étant donnée une formule insatisfiable ϕ et une réfutation par résolution P de ϕ , il existe une réfutation Max-SAT de ϕ contenant $O(2^{\mu(P)} \times |P|)$ étapes d'inférence, où $\mu(P)$ est le nombre de multi-utilisation des clauses intermédiaires de P .*

Références

- [1] Federico HERAS et Joao MARQUES-SILVA : Read-Once resolution for Unsatisfiability-Based Max-SAT Algorithms. *In IJCAI*, 2011.
- [2] María LUISA BONET, Jordi LEVY et Felip MANYÀ : Resolution for Max-SAT. *Artificial Intelligence*, 2007.
- [3] M. PY, M. S. CHERIF et D. HABET : Towards Bridging the Gap Between SAT and Max-SAT Refutations. *In 2020 IEEE 32nd International Conference on Tools with Artificial Intelligence (ICTAI)*, 2020.
- [4] J. A. ROBINSON : A machine-oriented logic based on the resolution principle. *In Journal of the Association for Computing Machinery*, 1965.
- [5] Alasdair URQUHART : The Complexity of Propositional Proof. *Bulletin of Symbolic Logic*, 1995.

Towards Bridging the Gap Between SAT and Max-SAT Refutations

Matthieu Py, Mohamed Sami Cherif and Djamel Habet

Aix-Marseille Université, Université de Toulon, CNRS, LIS, Marseille, France
{matthieu.py, mohamed-sami.cherif, djamel.habet}@univ-amu.fr

Abstract—Adapting a resolution proof for SAT to a Max-SAT resolution proof without increasing considerably the size of the proof is an open question. This paper contributes to this topic by exhibiting linear adaptations, in terms of the input SAT proof size, in restricted cases which are regular tree resolution refutations, tree resolution refutations and a new introduced class of refutations that we refer to as semi-tree resolution refutations. We also extend these results by proposing a complete adaptation for any unrestricted SAT refutation to a Max-SAT refutation, which is exponential in the worst case.

Index Terms—Resolution, Max-SAT Resolution, Refutation, Regular Resolution, Tree Resolution

I. INTRODUCTION

Given a Boolean formula in Conjunctive Normal Form (CNF), the Max-SAT problem consists in determining the maximum number of clauses that it is possible to satisfy by an assignment of the variables, while the SAT problem asks for the existence of an assignment which satisfies all the clauses. A well-known proof system for Max-SAT is Max-SAT resolution [20] which extends the resolution rule [24] used in the context of SAT. Max-SAT resolution plays a prominent role in Max-SAT as it is the most studied inference rule, both in theory and practice [1], [2], [7], [18], [20], [22].

In the context of SAT, an unsatisfiable formula can be refuted with a sequence of resolution steps which leads to the empty clause. Sequences of Max-SAT resolution steps are more constrained than sequences of resolution steps. Indeed, resolution adds the conclusion to the premises whereas the premise clauses are replaced by the conclusions when applying Max-SAT resolution. Switching from a read-once resolution proof, where each clause is used once, to a Max-SAT resolution proof is possible and well-known [11]. However, the adaptation of any resolution proof to a Max-SAT resolution one, especially in the context of a refutation, is an established problem. Bonet et al. state that "it seems difficult to adapt a classical resolution proof to get a Max-SAT resolution proof, and it is an open question if this is possible without increasing substantially the size of the proof" [20].

This paper attempts to contribute to this open question on refutation proofs by proposing a way to deal with non-read-once clauses, i.e. clauses used several times as a premise of a resolution step. Indeed, to adapt any resolution refutation, it is necessary to duplicate the non-read-once clauses in some form while also preserving Max-SAT equivalence. To this

end, we augment Max-SAT resolution with a simple split rule which allows to generate two clauses subsumed by the original clause. Intuitively, applying the split rule on a non-read-once clause will duplicate it since only literals that will not affect the rest of the proof will be added.

Accordingly, we deal first with regular tree resolution refutations [4], [26], showing that a linear adaptation of a SAT refutation to a Max-SAT one is possible in this case. Then, we extend this result to tree resolution refutations using a known result in [25] which stipulates that a minimal tree resolution refutation is regular. Furthermore, we introduce a new class of refutations that we refer to as semi-tree-like, which is a generalization of tree resolution refutations, and we extend our linear result to this class of refutations. Finally, we propose a complete adaptation of any (or unrestricted) resolution refutation to a Max-SAT refutation, although with an exponential factor, using the fact that any resolution refutation can be made tree-like with an exponential cost.

This paper is organized as follows. Section II gives some necessary definitions and notations. Sections III to VI describe the above contributions in detail. Finally, we conclude and discuss future work in Section VII.

II. PRELIMINARIES

A. Definitions and Notations

Let X be the set of propositional variables. A literal l is a variable $x \in X$ or its negation \bar{x} . A clause c is a disjunction (or a set) of literals $(l_1 \vee l_2 \vee \dots \vee l_k)$. A formula in Conjunctive Normal Form (CNF) ϕ is a conjunction (or a multiset) of clauses $\phi = c_1 \wedge c_2 \wedge \dots \wedge c_m$. An assignment $I : X \rightarrow \{true, false\}$ maps each variable to a boolean value and can be represented as a set of literals. A literal l is satisfied (resp. falsified) by an assignment I if $l \in I$ (resp. $\bar{l} \in I$). A clause c is satisfied by an assignment I if at least one of its literals is satisfied by I , otherwise it is falsified by I . The empty clause \square contains zero literals and is always falsified. A clause c opposes a clause c' if c contains a literal whose negation is in c' , i.e. $\exists l \in c, \bar{l} \in c'$. A clause c subsumes a clause c' if each literal of c is a literal of c' , i.e. $\forall l \in c, l \in c'$. We denote $var(c)$ the variables appearing in the clause c . A CNF formula ϕ is satisfied by an assignment I , that we call model of ϕ , if each clause $c \in \phi$ is satisfied by I , otherwise it is falsified by I . Solving the Satisfiability (SAT) problem consists in determining whether there exists an

assignment I that satisfies a given CNF formula ϕ . In the case where such an assignment exists, we say that ϕ is satisfiable, otherwise we say that ϕ is unsatisfiable or inconsistent. The cost of an assignment I , denoted $cost_I(\phi)$, is the number of clauses falsified by I . The Maximum Satisfiability (Max-SAT) problem is an optimization extension of SAT which, for a given CNF formula ϕ , consists in determining the maximum number of clauses that can be satisfied by an assignment of ϕ . Equivalently, it consists in determining the minimum number of clauses that each assignment must falsify, i.e. $\min_I cost_I(\phi)$.

B. Resolution Refutations in SAT

To certify that a CNF formula is satisfiable, it is sufficient to simply exhibit a model of the formula. On the other hand, to prove that a CNF formula is unsatisfiable, we need to refute the existence of a model. To this end, we can exhibit a SAT refutation which consists of a sequence of equivalence-preserving transformations (in the sense of SAT as defined below) starting from the formula and ultimately deducing an empty clause.

Definition 1 (SAT Equivalence). *Let ϕ and ϕ' be two CNF formulas. We say that ϕ is equivalent (in the sense of SAT) to ϕ' if for any assignment $I : var(\phi) \cup var(\phi') \rightarrow \{true, false\}$, I is a model of ϕ if and only if I is a model of ϕ' .*

A well-known SAT refutation system is based on an inference rule for SAT called resolution [24]. Refutations in this system are referred to as resolution refutations. The resolution rule, defined below, deduces a clause called resolvent from two opposed clauses which can be added to the formula while preserving SAT equivalence. Resolution plays an important role in the context of Conflict Driven Clause Learning (CDCL) [21]. Furthermore, it was shown that CDCL can polynomially simulate general resolution [23]. As showcased in Example 1, a resolution proof can be represented as a Directed Acyclic Graph (DAG) whose nodes are clauses in the proof either having two or zero incoming arcs (resp. if they are resolvents or clauses of the initial formula).

Definition 2 (Resolution [24]). *Given two clauses $c_1 = (x \vee A)$ and $c_2 = (\bar{x} \vee B)$, the resolution rule is defined as follows:*

$$\frac{c_1 = (x \vee A) \quad c_2 = (\bar{x} \vee B)}{c_3 = (A \vee B)}$$

Example 1. *We consider the CNF formula $\phi = (\bar{x}_1 \vee x_3) \wedge (x_1) \wedge (\bar{x}_1 \vee x_2) \wedge (\bar{x}_2 \vee \bar{x}_3)$. A resolution refutation of ϕ is represented as a DAG in Fig. 1.*

Many restricted classes of resolution refutations have been studied in the literature namely linear resolution [19], unit resolution [12], input resolution [12], regular resolution [26], read-once resolution [14] and tree (or tree-like) resolution [4] refutations among others. In particular, a resolution refutation is tree-like if every intermediate clause, i.e. resolvent, is used at most once in the proof. It is known that DPLL algorithms [8] on unsatisfiable instances correspond to tree resolution




Fig. 1. Resolution refutation

refutations [9]. Similarly, a resolution refutation is read-once if each clause is used at most once in the proof. Clearly, read-once resolution refutations are also tree-like since they form a restricted class of tree resolution refutations. It was shown in [14] that there exists unsatisfiable CNF formulas which cannot be refuted using read-once resolution. Finally, a resolution is regular if every variable is resolved on at most once in each branch of the DAG, i.e. path from a clause of the initial formula to the empty clause. It was shown that CDCL without restarts can polynomially simulate regular resolution [6]. We say that an irregularity is a sequence of clauses (each clause must be deduced using the previous one as premise) such that the first clause and the last one contain a literal l but at least one of the intermediate clauses does not contain this literal l . In other words, an irregularity is a certificate that a resolution refutation is not regular.

Example 2. *We consider the refutation of ϕ in Example 1. The refutation is clearly tree-like but it is not read-once since clause (x_1) is used two times as a premise of a resolution step. The refutation is also regular as every variable is resolved on at most once in every branch of the DAG in Fig. 1.*

C. Max-SAT Refutations

Several complete proof systems for Max-SAT were introduced in the literature, namely the Max-SAT resolution Calculus in [20] and the Clause Tableau Calculus in [17]. In particular, Max-SAT resolution, one of the first known complete systems for Max-SAT, was inspired from Resolution. The aim of complete Max-SAT systems is not to refute the formula per se but to compute the Max-SAT optimum of a given CNF formula, i.e. the maximum number of falsified clauses. The formula is thus refuted as many times as its optimum through equivalence-preserving transformations in the sense of Max-SAT as defined below.

Definition 3 (Max-SAT Equivalence). *Let ϕ and ϕ' be two CNF formulas. We say that ϕ is equivalent (in the sense of Max-SAT) to ϕ' if for any assignment $I : var(\phi) \cup var(\phi') \rightarrow \{true, false\}$, we have $cost_I(\phi) = cost_I(\phi')$.*

The Max-SAT resolution proof system relies on an inference rule that extends resolution for Max-SAT. Other than the resol-

non empty since if there exists an empty subset v can't be the first junction point of the branches. Let x be the variable eliminated at this resolution step and suppose w.l.o.g that the first premise contains literal x while the second contains literal \bar{x} . As P is regular, x is not a variable of c and the subset of branches starting from c leading to the first premise accepts the substitution of c by $c \vee x$ while the subset of branches leading to the second premise accepts the substitution of c by $c \vee \bar{x}$. ■

The result established in Lemma 1 ensures the possibility to fix any non-read-once clause used $k > 1$ times by using the split rule. Indeed, we can apply this rule to replace a non-read-once clause used $k > 1$ times by two clauses used respectively $1 \leq k_1 < k$ and $1 \leq k_2 < k$ such that $k = k_1 + k_2$. By iterating this method, we can fix every non-read-once clause. Then, we only need to replace the resolution rule by the Max-SAT resolution rule to obtain an adaptation from any regular tree resolution refutation to a max-refutation in linear size.

Theorem 1. *Given an unsatisfiable formula ϕ and a regular tree resolution refutation P of ϕ , there exists a max-refutation of ϕ containing $O(|P|)$ inference steps.*

Proof. Let P be a regular tree resolution refutation of ϕ . We set $T_1 = \emptyset$ and $T_2 = MR(P)$, where $MR(P)$ is obtained from P after replacing each resolution by Max-SAT resolution. If P is read-once, T_2 is a max-refutation of ϕ containing $|P|$ inference steps (which is obviously in $O(|P|)$). Now, let c be a non-read-once clause of P . Using Lemma 1, there exists a variable $x \notin \text{var}(c)$ and a partition of the branches starting from c into two non-empty subsets, the first accepting $c \vee x$ and the second accepting $c \vee \bar{x}$. We apply the Max-SAT split rule on c to obtain $c \vee x$ and $c \vee \bar{x}$ and we replace c as premise by $c \vee x$ on the first subset of branches and c by $c \vee \bar{x}$ on the second. Doing this, we augment T_1 by adding one split and we change T_2 by replacing the premise clause c as described above. As T_2 is a tree-like regular resolution refutation of $(\phi \setminus c) \wedge (c \vee x) \wedge (c \vee \bar{x})$, it is possible to iteratively apply this operation on T_2 until we obtain a read-once regular tree resolution refutation. Therefore, after the last iteration, we have a couple (T_1, T_2) such that T_1 is a sequence of applications of the split rule transforming ϕ into a Max-SAT equivalent ϕ' and T_2 is a read-once regular max-refutation of ϕ' . Therefore, these transformations form a max-refutation of ϕ .

To prove that the size of the max-refutation is in $O(|P|)$, we first consider how to fix a leaf clause of P (i.e. how to replace it by read-once clauses). If c is a leaf clause of P used k times, we prove by induction on k that it is possible to fix this clause using at most $k - 1$ splits:

- If $k = 1$, we clearly need 0 splits to fix the read-once clause c .
- Suppose that the assertion is true for any $k' < k$ and let c be a clause used k times. Using Lemma 1, it is possible to use 1 split to replace c by two clauses c_1 and

c_2 respectively used k_1 and k_2 times with $k_1, k_2 > 0$ and $k_1 + k_2 = k$. Using our assertion for k_1 and k_2 , it is possible to fix c_1 with at most $k_1 - 1$ splits and c_2 with at most $k_2 - 1$ splits. Therefore, it is possible to fix c with at most $1 + (k_1 - 1) + (k_2 - 1) = k - 1$ splits.

Let c_1, \dots, c_p be the leaf clauses of P used respectively k^1, k^2, \dots, k^p times. Notice that $k^1 + k^2 + \dots + k^p = |P| + 1$ since P has exactly $2|P|$ premises, i.e. uses of clauses, and $|P| - 1$ intermediate clauses (the empty clause is not used and we neglect the trivial cases where a non-empty intermediate clause is not used and where the proof produces several empty clauses). Using the previous induction, we need at most $k^1 - 1 + k^2 - 1 + \dots + k^p - 1 \leq |P|$ splits to fix every non-read-once leaf clause of P . Consequently, $|T_1| \leq |P|$. On the other hand, the number of Max-SAT resolutions in T_2 is by construction equal to the number of resolution steps in P and, therefore, $|T_2| = |P|$. We conclude that the complete max-refutation contains at most $2|P|$ inference steps, which is in $O(|P|)$. ■

Example 4. *We consider the regular tree resolution refutation from Example 1 represented by the DAG in Fig. 1. We observe that the original clause (x_1) is used two times as a premise of a resolution step. The junction point of the left and right branches eliminates variable x_3 such that the branch on the left leads to the premise containing literal x_3 and the branch on right leads to the premise containing literal \bar{x}_3 . We apply the split rule on clause (x_1) to get $(x_1 \vee x_3)$ and $(x_1 \vee \bar{x}_3)$ and we replace (x_1) by $(x_1 \vee x_3)$ and $(x_1 \vee \bar{x}_3)$ respectively on the left and right branches. Finally, we replace all resolutions by Max-SAT resolutions to obtain the complete max-refutation.*




Fig. 3. Applying the split rule to deal with a non-read once clause

IV. FROM TREE RESOLUTION REFUTATIONS TO MAX-REFUTATIONS

In the previous section, we proposed a linear adaptation from regular tree resolution refutations to max-refutations. We propose in this section to extend the case where this adaptation guarantees linear size of the obtained max-refutation to tree resolution refutations. To this end, we simply exhibit a known transformation from any tree resolution refutation to a regular

resolution refutation is bounded by $O(2^{\mu(P)} \times |P|)$. ■

Now that we can transform any resolution refutation to a tree-like resolution refutation, we just have to adapt the obtained tree-like resolution refutation with the method described in section IV as shown in the following Theorem.

Theorem 3. *Given an unsatisfiable formula ϕ and an unrestricted resolution refutation P of ϕ , there exists a max-refutation of ϕ with $O(2^{\mu(P)} \times |P|)$ inference steps.*

Proof. Let P be an unrestricted resolution refutation of ϕ . We adapt P to obtain a tree resolution refutation P_t of size $O(2^{\mu(P)} \times |P|)$ using Lemma 3. Then, using Theorem 1, we obtain a max-refutation of size $O(2^{\mu(P)} \times |P|)$. ■

Example 8. *We consider the resolution refutation represented in Fig. 8. This refutation is not semi-tree-like since the clauses (x_1) and (x_4) are two non read-once clauses in the same branch. First, we duplicate the resolutions leading to (x_1) and we obtain the tree-like resolution refutation represented in Fig. 9. Then, we apply the transformations described in Section IV to get the max-refutation represented in Fig. 10.*




Fig. 8. Unrestricted resolution refutation




Fig. 9. Adapting a resolution refutation to a tree-like resolution refutation

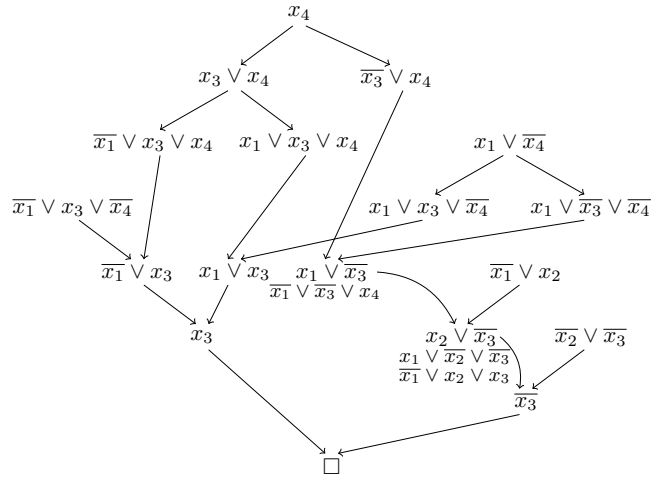


Fig. 10. Adapting an unrestricted resolution refutation to a max-refutation

We finish this section by exhibiting resolution refutations whose adaptations as in Theorem 3 is exponential. To this end, we introduce in the following definition a new pattern which we will use to build such refutations.

Definition 8 (Diamond pattern). *Let A be a disjunction of literals and let $x \notin \text{var}(A)$ and $y \notin \text{var}(A)$ two distinct variables. We define the diamond pattern (x, y, A) as the sequence of resolutions represented in Fig.11.*




Fig. 11. Diamond pattern (x, y, A)



Fig. 12. Simplified representation of a diamond pattern

We can represent this pattern by a diamond as in Fig. 12. Notice that in particular, the diamond pattern (x, y, \square) is a resolution refutation. Now, imagine that the topmost clause of (x, y, \square) is derived through another diamond pattern. We iterate the same reasoning to define a k -stacked diamonds pattern as follows:

Definition 9 (k -stacked diamond pattern). *Let $k \geq 1$ be a natural number and let x_i and y_i where $1 \leq i \leq k$ be distinct variables. A k -stacked diamond pattern is formed by k diamond patterns (x_i, y_i, A_i) where $1 \leq i \leq k$ such that $A_1 = \square$ and $A_i = (x_1 \vee \dots \vee x_{i-1})$ for $1 < i \leq k$. Each diamond (x_i, y_i, A_i) is stacked on top of $(x_{i-1}, y_{i-1}, A_{i-1})$ such that the last conclusion of the former is the topmost central premise of the latter.*

A k -stacked diamond pattern is represented as a stack of diamonds as shown in Fig.13 for $k = 3$. Clearly, k -stacked diamond are resolution refutations as they deduce the empty clause \square . In particular, when $k > 2$, a k -stacked diamond is not semi-tree-like. The size of a k -stacked diamond P is $|P| = 3k$. Furthermore, we have $\mu(P) = k - 1$. Therefore, after the application of the adaptation described in Theorem 3, we obtain a max-refutation whose size is at least 2^{k-1} showing that the proposed adaptation can be exponential in the worst case.




Fig. 13. Simplified representation of a 3-stacked diamond pattern

VII. CONCLUSIONS AND FUTURE WORK

The contributions of this work are related to adapting resolution refutations to Max-SAT refutations. In particular, we have proposed linear adaptations regarding the size of the resolution refutations in the following cases: regular tree resolution, tree resolution and semi-tree resolution. These results are achieved by augmenting Max-SAT resolution with the split rule which enabled us to duplicate clauses by adding literals when necessary. We have also generalised our adaptation to unrestricted resolution refutations, even though the proposed transformation can produce a max-refutation whose size is exponential in the worst case. Notice that our results remain valid for weighted Max-SAT formulas as we simply need to augment the previous rules with another split rule for weights. Indeed, the overhead of this rule is linear in terms of the size of the refutation since we need to apply it once on the clauses in the weighted MAX-SAT formula to produce clauses with same weight, i.e. the minimum of all the weights as done in the context of SAT-based (weighted) Max-SAT algorithms [3].

These results may help to exhibit proofs based on Max-SAT resolution for Max-SAT algorithms which remains an unexplored topic whereas, in SAT, practically all modern solvers are able to compute a resolution proof of unsatisfiability (in different formats [10], [13]). Indeed, it would be interesting to include the proposed adaptations in SAT-based algorithms for Max-SAT. Such extended algorithms would thus iteratively call a SAT oracle to get a resolution refutation, adapt this resolution refutation to get a Max-SAT refutation based on our results and transform the formula accordingly. This treatment is repeated until reaching a satisfiable formula and a set of empty clauses whose size is the optimum value of the formula. Such implementation would require an efficient SAT oracle which returns a resolution refutation. Finally, the existence of an adaptation that does not increase substantially the size of

an unrestricted resolution proof remains an open question. We will continue to investigate this topic either by exhibiting a polynomial adaptation or refuting its existence.

REFERENCES

- [1] A. Abramé and D. Habet. On the resiliency of unit propagation to max-resolution. In *IJCAI*, 2015.
- [2] A. Abramé and D. Habet. ahmaxsat: Description and evaluation of a branch and bound max-sat solver. In *Journal on Satisfiability, Boolean Modeling and Computation. Vol. 9*, pp. 89–128, 2015.
- [3] C. Ansótegui, M. L. Bonet, and J. Levy. Solving (weighted) partial maxsat through satisfiability testing. In O. Kullmann, editor, *SAT 2009*, volume 5584 of *Lecture Notes in Computer Science*, pages 427–440. Springer, 2009.
- [4] E. Ben-sasson, R. Impagliazzo, and A. Wigderson. Near optimal separation of tree-like and general resolution. *Combinatorica*, 24:585–603, 09 2004.
- [5] M. Bonet, J. Levy, and F. Manyà. A complete calculus for max-sat. In *SAT 2006*, volume 4121, pages 240–251, 08 2006.
- [6] S. Buss, J. Hoffmann, and J. Johannsen. Resolution Trees with Lemmas: Resolution Refinements that Characterize DLL Algorithms with Clause Learning. *Logical Methods in Computer Science*, 4, 11 2008.
- [7] M. S. Cherif and D. Habet. Towards the characterization of max-resolution transformations of ucsp by up-resilience. In T. Schiex and S. de Givry, editors, *CP*, pages 91–107, Cham, 2019. Springer International Publishing.
- [8] M. Davis, G. Logemann, and D. Loveland. A machine program for theorem-proving. *Communications of the ACM*, 5(7):394–397, 1962.
- [9] A. V. Gelder. Extracting (easily) checkable proofs from a satisfiability solver that employs both preorder and postorder resolution. In *ISAIM*, 2002.
- [10] A. V. Gelder. Verifying rup proofs of propositional unsatisfiability. In *ISAIM*, 2008.
- [11] F. Heras and J. Marques-Silva. Read-once resolution for unsatisfiability-based max-sat algorithms. In *IJCAI*, 2011.
- [12] A. Hertel and A. Urquhart. Algorithms and complexity results for input and unit resolution. *Journal of Satisfiability, Boolean Modeling and Computation*, 6, 05 2009.
- [13] M. J. Heule, W. A. Hunt, and N. Wetzler. Trimming while checking clausal proofs. In *2013 Formal Methods in Computer-Aided Design*, pages 181–188. IEEE, 2013.
- [14] K. Iwama and E. Miyano. Intractability of read-once resolution. In *Proceedings of Structure in Complexity Theory. Tenth Annual IEEE Conference*, 1995.
- [15] A. Küegel. Improved exact solver for the weighted max-sat problem. In *Daniel Le Berre (editor). POS-10. Pragmatics of SAT, vol 8*, pages 15–27, 2012.
- [16] J. Larrosa and F. Heras. Resolution in max-sat and its relation to local consistency in weighted csps. In *IJCAI*, pages 193–198, 01 2005.
- [17] C.-M. Li, F. Manyà, and J. R. Soler. A clause tableau calculus for maxsat. In *IJCAI, IJCAI'16*, page 766–772. AAAI Press, 2016.
- [18] C.-M. Li, F. Manyà, and J. Planes. New inference rules for max-sat. *J. Artif. Intell. Res. (JAIR)*, 30:321–359, 09 2007.
- [19] D. Loveland. A linear format for resolution. *Symposium on Automatic Demonstration*, pages 147–162, 01 1970.
- [20] M. Luisa Bonet, J. Levy, and F. Manyà. Resolution for max-sat. *Artificial Intelligence Volume 171, Issues 8–9, June 2007, Pages 606-618*, 2007.
- [21] J. P. Marques Silva and K. A. Sakallah. Grasp-a new search algorithm for satisfiability. In *Proceedings of International Conference on Computer Aided Design*, pages 220–227, 1996.
- [22] N. Narodytska and F. Bacchus. Maximum satisfiability using core-guided maxsat resolution. In *AAAI*, 2014.
- [23] K. Pipatsrisawat and A. Darwiche. On the power of clause-learning sat solvers as resolution engines. *Artificial Intelligence*, 175(2):512 – 525, 2011.
- [24] J. A. Robinson. A machine-oriented logic based on the resolution principle. In *Journal of the Association for Computing Machinery, vol. 12 (1965)*, pp. 23–41., 1965.
- [25] A. Urquhart. The complexity of propositional proofs. *Bull. Symbolic Logic*, 1(4):425–467, 12 1995.
- [26] A. Urquhart. A near-optimal separation of regular and general resolution. *SIAM J. Comput.*, 40:107–121, 01 2011.