





Towards a Better Understanding of (Partial Weighted) MaxSAT Proof Systems

Javier Larrosa  and Emma Rollon 

Universitat Politècnica de Catalunya, Barcelona Tech, Barcelona, Spain
{larrosa,erollon}@cs.upc.edu
<http://www.cs.upc.edu>

Abstract. MaxSAT is a very popular language for discrete optimization with many domains of application. While there has been a lot of progress in MaxSAT solvers during the last decade, the theoretical analysis of MaxSAT inference has not followed the pace. Aiming at compensating that lack of balance, in this paper we do a proof complexity approach to MaxSAT resolution-based proof systems. First, we give some basic definitions on completeness and show that refutational completeness makes completeness redundant, as it happens in SAT. Then we take three inference rules such that adding them sequentially allows us to navigate from the weakest to the strongest resolution-based MaxSAT system available (i.e., from standalone MaxSAT resolution to the recently proposed ResE), each rule making the system stronger. Finally, we show that the strongest system captures the recently proposed concept of Circular Proof while being conceptually simpler, since weights, which are intrinsic in MaxSAT, naturally guarantee the flow condition required for the SAT case.

Keywords: MaxSAT · Rule-based proof systems · Circular proofs

1 Introduction

Proof Complexity is the field aiming to understand the computational cost required to prove or refute statements. Different proof systems may provide different proofs for the same formula and some proof systems are provably more efficient than others. When that happens, proof complexity cares about which elements of the more powerful proof system really make the difference.

In propositional logic, proof systems that work with CNF formulas have attracted the interest of researchers for several decades. One of the reasons is that CNF is the working language of the extremely successful SAT solvers and the search spaces that they traverse can be understood and analyzed as proofs [5].

Projects TIN2015-69175-C4-3-R and RTI2018-094403-B-C33, funded by: FEDER/Ministerio de Ciencia e Innovación – Agencia Estatal de Investigación, Spain.

© Springer Nature Switzerland AG 2020

L. Pulina and M. Seidl (Eds.): SAT 2020, LNCS 12178, pp. 218–232, 2020.

https://doi.org/10.1007/978-3-030-51825-7_16

(Partial Weighted) *MaxSAT* is the optimization version of SAT. Although discrete optimization problems can be modeled and solved with SAT solvers, many of these problems are more naturally treated as MaxSAT. For this reason the design of MaxSAT solvers has attracted the interest of researchers in the last decade. Interestingly, while some of the first efficient MaxSAT solvers were strongly influenced by MaxSAT inference [9], this influence has diminished along time. The currently most efficient algorithms solve MaxSAT by sophisticated sequences of calls to SAT solvers [1, 4, 11].

We think it is important to understand this scientific trend with a more formal approach and such understanding must go through an analysis of the possibilities and limitations of MaxSAT proof systems (how MaxSAT inference compares with obtaining the same result with a sequence of SAT inferences?). The purpose of this paper is to start contributing in that direction by improving the understanding of MaxSAT proof systems. With that aim we extend some classic proof complexity concepts (i.e, entailment, completeness, etc) to MaxSAT and analyze three proof systems of increasing complexity: from stand-alone MaxSAT resolution (**Res**) [9] to the recently proposed resolution with extension (**ResE**) [10]. For the sake of clarity, we break the extension rule of **ResE** into two atomic rules: *split* and *virtual*; and analyze their incremental power. Our results show that each add-on makes a provable stronger system. More precisely, we have observed that: **Res** is sound and refutationally complete. Adding the split rule (**ResS**) we get completeness and (unlike what happens in SAT) some exponential speed-up in certain refutations. Further adding the virtual rule (**ResSV**), which allows to keep negative weights during proofs, we get further speed-up by capturing the concept of *circular proofs* [3]. We also give the interesting and somehow unexpected result that in some cases rephrasing a MaxSAT refutation as a MaxSAT entailment may transform the problem from exponentially hard to polynomial when using **ResSV**.

The structure of the paper is as follows. In Sects. 2 and 3 we provide preliminaries on SAT and MaxSAT, respectively. In Sect. 4 we define some variations of the Pigeon Hole Problem that we need for the proofs of the theorems. In Sect. 5 we provide basic definition and properties on MaxSAT proof systems and introduce and analyze the different systems addressed in the paper. In Sect. 6 we show how the strongest proof system **ResSV** captures the notion of Circular Proof. Finally, in Sect. 7, we give some conclusions.

2 SAT Preliminaries

A *boolean variable* x takes values on the set $\{0, 1\}$. A *literal* is a variable x (positive literal) or its negation \bar{x} (negative literal). A *clause* is a disjunction of literals. A clause C is satisfied by a truth assignment X if X contains at least one of the literals in C . The *empty clause* is denoted \square and cannot be satisfied. The negation of a clause $C = l_1 \vee l_2 \vee \dots \vee l_p$ is satisfied if all its literals are falsified and this can be trivially expressed in CNF as the set of unit clause $\{\bar{l}_1, \bar{l}_2, \dots, \bar{l}_p\}$.

A *CNF formula* \mathcal{F} is a set of clauses (understood as a conjunction). A truth assignment satisfies a formula if it satisfies all its clauses. If such an assignment exists, we say that the assignment is a model and the formula is satisfiable, noted $SAT(\mathcal{F})$. Determining whether a formula is satisfiable constitutes the well-known SAT Problem.

We say that formula \mathcal{F} *entails* formula \mathcal{G} , noted $\mathcal{F} \models \mathcal{G}$, if every model of \mathcal{F} is also a model of \mathcal{G} . Two formulas \mathcal{F} and \mathcal{G} are *equivalent*, noted $\mathcal{F} \equiv \mathcal{G}$, if they entail each other.

An *inference rule* is given by a set of antecedent clauses and a set of consequent clauses. In SAT, the intended meaning of an inference rule is that if some clauses of a formula match the antecedents, the consequents can be *added*. The rule is *sound* if every truth assignment that satisfies the antecedents also satisfies the consequents. The process of applying an inference rule to a formula \mathcal{F} is noted $\mathcal{F} \vdash \mathcal{F}'$.

Consider the following two rules [3,12],

$$\frac{x \vee A \quad \bar{x} \vee B}{A \vee B} \qquad \frac{A}{A \vee x \quad A \vee \bar{x}}$$

(resolution) (split)

where A and B are arbitrary (possibly empty) disjunctions of literals and x is an arbitrary variable. In propositional logic it is customary to define rules with just one consequent because one rule with s consequents can be obtained from s one-consequent rules. As we will see, this is not the case in MaxSAT. For this reason, here we prefer to introduce the two-consequents split rule instead of the equivalent weakening rule [3] to keep the parallelism with MaxSAT more evident.

A *proof system* Δ is a set of inference rules. A *proof* of length e under a proof system Δ is a finite sequence $\mathcal{F}_0 \vdash \mathcal{F}_1 \vdash \dots \vdash \mathcal{F}_e$ where \mathcal{F}_0 is the original formula and each \mathcal{F}_i is obtained by applying an inference rule from Δ . We will use \vdash^* to denote an arbitrary number of inference steps. A *short* proof is a proof whose length can be bounded by a polynomial on $|\mathcal{F}|$. A *refutation* is a proof such that $\square \in \mathcal{F}_e$. Refutations are important because they prove unsatisfiability.

A proof system is *sound* if all its rules are sound. All the SAT rules and proof systems considered in this paper are sound. A proof system is *complete* if for every \mathcal{F}, \mathcal{G} such that $\mathcal{F} \models \mathcal{G}$, there is a proof $\mathcal{F} \vdash^* \mathcal{H}$ with $\mathcal{G} \subseteq \mathcal{H}$. Although completeness is a natural and elegant property, it has limited practical interest. For that reason a weaker version of completeness has been defined. A proof system is *refutationally complete* if for every unsatisfiable formula \mathcal{F} there is a refutation starting in \mathcal{F} (i.e, completeness is required only for refutations). It is usually believed that refutational completeness is enough for practical purposes. The reason is that $\mathcal{F} \models \mathcal{G}$ if and only if $\mathcal{F} \cup \bar{\mathcal{G}}$ is unsatisfiable, so any implicationally complete proof system can prove the entailment by deriving a refutation from a CNF formula equivalent to $\mathcal{F} \cup \bar{\mathcal{G}}$.

It is well-known that the proof system made exclusively of resolution is refutationally complete and adding the split rule makes the system complete. The following property says that adding the split rule does not give any advantage to resolution in terms of refutational power.

Property 1 [(see Lemma 7 in [2]). A proof system with resolution and split as inference rules cannot make shorter refutations than a proof system with only resolution.

3 MaxSAT Preliminaries

A *weight* w is a positive number or ∞ (i.e, $w \in \mathbb{R}^+ \cup \{\infty\}$). We extend sum and subtraction to weights defining $\infty + w = \infty$ and $\infty - w = \infty$ for all w . Note that $v - w$ is only defined when $w \leq v$.

A *weighted clause* is a pair (C, w) where C is a clause and w is a weight associated to its falsification. If $w = \infty$ we say that the clause is *hard*, else it is *soft*. A weighted *MaxSAT CNF formula* is a set of weighted clauses $\mathcal{F} = \{(C_1, w_1), (C_2, w_2), \dots, (C_p, w_p)\}$. If all the clauses are hard, we say that the formula is hard. We say that $\mathcal{G} \subseteq \mathcal{F}$ if for all $(C, w) \in \mathcal{G}$ there is a $(C, w') \in \mathcal{F}$ with $w \leq w'$.

Given a formula \mathcal{F} , we define the cost of a truth assignment X , noted $\mathcal{F}(X)$, as the sum of weights over the clauses that are falsified by X . The MaxSAT problem is to find the minimum cost over the set of all truth assignments,

$$\text{MaxSAT}(\mathcal{F}) = \min_X \mathcal{F}(X)$$

This definition of MaxSAT including weights and hard clauses is sometimes referred to as Partial Weighted MaxSAT [11]. Note that any clause (C, w) can be broken into two clauses $(C, u), (C, v)$ as long as $u + v = w$. In the following we will assume that clauses are separated and merged as needed.

We say that formula \mathcal{F} *entails* formula \mathcal{G} , noted $\mathcal{F} \models \mathcal{G}$, if $\mathcal{G}(X)$ is a lower bound of $\mathcal{F}(X)$. That is, if for all X , $\mathcal{F}(X) \geq \mathcal{G}(X)$. We say that two formulas \mathcal{F} and \mathcal{G} are *equivalent*, noted $\mathcal{F} \equiv \mathcal{G}$, if they entail each other. That is, if for all X , $\mathcal{F}(X) = \mathcal{G}(X)$.

In the following Sections we will find useful to deal with negated clauses. Hence, the corresponding definitions and useful properties. Let A and B be arbitrary disjunctions of literals. Let $(A \vee \bar{B}, w)$ mean that falsifying $A \vee \bar{B}$ incurs a cost of w . Although $A \vee \bar{B}$ is not a clause, the following property shows that it can be efficiently transformed into a CNF equivalent,

Property 2. $\{(A \vee \overline{l_1 \vee l_2 \vee \dots \vee l_p}, w)\} \equiv \{(A \vee \bar{l}_1, w), (A \vee l_1 \vee \bar{l}_2, w), \dots, (A \vee l_1 \vee \dots \vee l_{p-1} \vee \bar{l}_p, w)\}$.

Observe that if we restrict the MaxSAT language to hard formulas we have standard SAT CNF formulas where ∞ corresponds to *false* and 0 corresponds to *true*. Note that all the previous definitions naturally instantiate to their SAT analogous.

4 Pigeon Hole Problem and Variations

We define the well-known Pigeon Hole Problem *PHP* and three MaxSAT versions *SPHP*, *SPHP*⁰ and *SPHP*¹, that we will be using in the proof of our results.

In the *Pigeon Hole Problem PHP* the goal is to assign $m + 1$ pigeons to m holes without any pair of pigeons sharing their hole. In the usual SAT encoding there is a boolean variable x_{ij} (with $1 \leq i \leq m + 1$, and $1 \leq j \leq m$) which is true if pigeon i is in hole j . There are two groups of clauses. For each pigeon i , we have the clause,

$$\mathcal{P}_i = \{x_{i1} \vee x_{i2} \vee \dots \vee x_{im}\}$$

indicating that pigeon i must be assigned to a hole. For each hole j we have the set of clauses,

$$\mathcal{H}_j = \{\bar{x}_{ij} \vee \bar{x}_{i'j} \mid 1 \leq i < i' \leq m + 1\}$$

indicating that hole j is occupied by at most one pigeon. Let \mathcal{K} be the union of all these sets of clauses $\mathcal{K} = \cup_{1 \leq i \leq m+1} \mathcal{P}_i \cup \cup_{1 \leq j \leq m} \mathcal{H}_j$. It is obvious that \mathcal{K} is an unsatisfiable CNF formula. In MaxSAT notation the pigeon hole problem is,

$$PHP = \{(C, \infty) \mid C \in \mathcal{K}\}$$

and clearly $MaxSAT(PHP) = \infty$.

In the *soft Pigeon Hole Problem SPHP* the goal is to find the assignment that falsifies the minimum number of clauses. In MaxSAT language it is encoded as,

$$SPHP = \{(C, 1) \mid C \in \mathcal{K}\}$$

and it is obvious that $MaxSAT(SPHP) = 1$.

The $SPHP^0$ problem is like the soft pigeon hole problem but augmented with one more clause $(\square, m^2 + m)$ where m is the number of holes. Note that $MaxSAT(SPHP^0) = m^2 + m + 1$.

Finally, the $SPHP^1$ problem is like the soft pigeon hole problem but augmented with a set of unit clauses $\{(x_{ij}, 1), (\bar{x}_{ij}, 1) \mid 1 \leq i \leq m + 1, 1 \leq j \leq m\}$. Note that $MaxSAT(SPHP^1) = m^2 + m + 1$.

5 MaxSAT Proof Systems

A *MaxSAT inference rule* is given by a set of antecedent clauses and a set of consequent clause. In MaxSAT, the application of an inference rule is to *replace* the antecedents by the consequents. The process of applying an inference rule to a formula \mathcal{F} is also noted $\mathcal{F} \vdash \mathcal{F}'$. The rule is *sound* if it preserves the equivalence of the formula.

As in the SAT case, given a proof system Δ (namely, a set of rules) a *proof* of length e is a sequence $\mathcal{F}_0 \vdash \mathcal{F}_1 \vdash \dots \vdash \mathcal{F}_e$ where \mathcal{F}_0 is the original formula and each \mathcal{F}_i is obtained by applying an inference rule from Δ . If $\mathcal{G} \subseteq \mathcal{F}_e$, we say that the proof is a proof of \mathcal{G} from \mathcal{F} .

A proof system is *sound* if all its rules are sound. In this paper all MaxSAT rules and proof systems are sound. A proof system is *complete* if for every \mathcal{F}, \mathcal{G}

such that $\mathcal{F} \models \mathcal{G}$, there is a proof of \mathcal{G} from \mathcal{F} . A *refutation* of \mathcal{F} is a proof of (\square, k) from \mathcal{F} with $k = \text{MaxSAT}(\mathcal{F})$. A proof system is *refutationally complete* if it can derive a refutation of every formula \mathcal{F} .

Next we show that, similarly to what happens in SAT, refutationally completeness is sufficient for practical purposes. The reason is that it can also be used to proof or disproof general entailment, making completeness somehow redundant. We need first to define the maximum soft cost of a formula as $\top(\mathcal{F}) = \sum_{(C,w) \in \mathcal{F} | w \neq \infty} w$ and the negation of a MaxSAT formula as the negation of all its clauses $\overline{\mathcal{F}} = \{(\overline{C}, w) \mid (C, w) \in \mathcal{F}\}$. The following property tells the effect of negating a formula without hard clauses,

Property 3. If \mathcal{F} is a CNF MaxSAT formula without hard clauses, then

$$\overline{\mathcal{F}}(X) = \top(\mathcal{F}) - \mathcal{F}(X)$$

Proof. Let X be a truth assignment, \mathcal{S} be the set of clauses satisfied by X and \mathcal{U} be the set of clauses falsified by X . It is clear that $\mathcal{F}(X) = \sum_{(C_i, w_i) \in \mathcal{U}} w_i$ while $\overline{\mathcal{F}}(X) = \sum_{(C_i, w_i) \in \mathcal{S}} w_i$. Since $\mathcal{S} \cap \mathcal{U} = \emptyset$ and $\mathcal{S} \cup \mathcal{U} = \mathcal{F}$, then $\sum_{(C_i, w_i) \in \mathcal{U}} w_i + \sum_{(C_i, w_i) \in \mathcal{S}} w_i = \top(\mathcal{F})$. Therefore, $\mathcal{F}(X) + \overline{\mathcal{F}}(X) = \top(\mathcal{F})$ and, as a consequence, $\overline{\mathcal{F}}(X) = \top(\mathcal{F}) - \mathcal{F}(X)$.

We can now show that an entailment $\mathcal{F} \models \mathcal{G}$ can be rephrased as a MaxSAT problem,

Theorem 1. *Let \mathcal{F} and \mathcal{G} be two MaxSAT formulas, possibly with hard clauses. Then,*

$$\mathcal{F} \models \mathcal{G} \text{ iff } \text{MaxSAT}(\mathcal{F} \cup \overline{\mathcal{G}'}) \geq \top(\mathcal{G}')$$

where \mathcal{G}' is a softened version of \mathcal{G} in which infinity weights are replaced by $\max\{\top(\mathcal{G}), \top(\mathcal{F})\} + 1$.

Proof. Let us prove the if direction. $\mathcal{F} \models \mathcal{G}$ means that $\forall X, \mathcal{F}(X) \geq \mathcal{G}(X)$. Also, by construction $\forall X, \mathcal{G}(X) \geq \mathcal{G}'(X)$. Therefore, $\forall X, \mathcal{F}(X) \geq \mathcal{G}(X)$. Because \mathcal{G}' does not contain hard clauses, $\mathcal{G}'(X) \neq \infty$, which means that, $\forall X, \mathcal{F}(X) - \mathcal{G}'(X) \geq 0$. Adding $\top(\mathcal{G}')$ to both sides of the disequality we get, $\forall X, \mathcal{F}(X) + \top(\mathcal{G}') - \mathcal{G}'(X) \geq \top(\mathcal{G}')$. By Property 3, we have, $\forall X, \mathcal{F}(X) + \overline{\mathcal{G}'}(X) \geq \top(\mathcal{G}')$ which clearly means that, $\text{MaxSAT}(\mathcal{F} \cup \overline{\mathcal{G}'}) \geq \top(\mathcal{G}')$.

Let us proof the else if direction. $\text{maxSAT}(\mathcal{F} \cup \overline{\mathcal{G}'}) \geq \top(\mathcal{G}')$ implies that $\forall X, \mathcal{F}(X) + \overline{\mathcal{G}'}(X) \geq \top(\mathcal{G}')$. Moreover, since $\overline{\mathcal{G}'}$ does not have hard clauses, from Property 3 we know that, $\forall X, \mathcal{F}(X) + \top(\mathcal{G}') - \mathcal{G}'(X) \geq \top(\mathcal{G}')$ so we have, $\forall X, \mathcal{F}(X) \geq \mathcal{G}'(X)$ and we need to have, $\forall X, \mathcal{F}(X) \geq \mathcal{G}(X)$. We reason on cases for truth assignment X :

1. If $\mathcal{G}'(X) \leq \top(\mathcal{G})$, by definition of \mathcal{G}' , $\mathcal{G}'(X) = \mathcal{G}(X) \neq \infty$. Therefore, $\mathcal{F}(X) \geq \mathcal{G}'(X) = \mathcal{G}(X)$, which proofs this case.

2. If $\mathcal{G}'(X) > \top(\mathcal{G})$, by definition of \mathcal{G}' , $\mathcal{G}(X) = \infty$. We show that in this case, $\mathcal{F}(X) = \infty$.
- if $\top(\mathcal{F}) \leq \top(\mathcal{G})$ then $\mathcal{F}(X) \geq \mathcal{G}'(X) > \top(\mathcal{G}) \geq \top(\mathcal{F})$. We show that $\mathcal{F}(X) > \top(\mathcal{F})$ implies that $\mathcal{F}(X) = \infty$. We proceed by contradiction. Let us suppose that $\mathcal{F}(X) > \top(\mathcal{F})$ and $\mathcal{F}(X) \neq \infty$. The latter means that X satisfies all hard clauses. As a consequence, $\mathcal{F}(X) \leq \top(\mathcal{F})$, which contradicts the hypothesis.
 - if $\top(\mathcal{F}) > \top(\mathcal{G})$, then there are no X such that $\top(\mathcal{G}) < \mathcal{G}'(X) \leq \top(\mathcal{F})$. By definition of \mathcal{G}' , for all $(C_i, \infty) \in \mathcal{G}$, $(C_i, \top(\mathcal{F}) + 1) \in \mathcal{G}'$. Therefore, either X satisfies all hard clauses in \mathcal{G} and then $\mathcal{G}'(X) \leq \top(\mathcal{G})$ or X falsifies at least one hard clause in \mathcal{G} and then $\mathcal{G}'(X) > \top(\mathcal{F})$.

which proofs the theorem.

The application of the previous theorem to single clause entailment yields the following corollary.

Corollary 1. *Let \mathcal{F} be a formula and (C, w) be a weighted clause. Then,*

$$\mathcal{F} \models (C, w) \text{ iff } \text{MaxSAT}(\mathcal{F} \cup \{(\overline{C}, u)\}) \geq u$$

$$\text{where } u = \begin{cases} w, & \text{if } w \neq \infty \\ \top(\mathcal{F}) + 1, & \text{if } w = \infty \end{cases}$$

A useful application of this corollary will be shown in Sect. 5.3.

In the rest of the section we introduce and analyze the incremental impact of the three inference rules.

5.1 Resolution

The MaxSAT *resolution* rule [8] is,

$$\frac{(x \vee A, v) \quad (\bar{x} \vee B, w)}{(A \vee B, m)} \\ (x \vee A, v - m) \quad (\bar{x} \vee B, w - m) \\ (x \vee A \vee \bar{B}, m) \quad (\bar{x} \vee B \vee \bar{A}, m)$$

where A and B are arbitrary (possibly empty) disjunctions of literals and $m = \min\{v, w\}$. When A (resp. B) is empty, \bar{A} (resp. \bar{B}) is constant true, so $x \vee \bar{A} \vee B$ (resp. $x \vee A \vee \bar{B}$) is tautological. Note that MaxSAT resolution, when applied to two hard clauses, corresponds to SAT resolution.

It is known that the proof system **Res** made exclusively of the resolution rule is refutationally complete,

Theorem 2 [6,9]. *Res is refutationally complete.*

However, as we show next, it is not complete.

Theorem 3. *Res is not complete.*

Proof. Consider formula $\mathcal{F} = \{(x, 1), (y, 1)\}$. It is clear that $\mathcal{F} \models (x \vee y, 1)$ which cannot be derived with **Res**.

It is well-known that **Res** cannot compute short refutations for *PHP* [12] or *SPHP* [6]. However, it can efficiently refute *SPHP*¹. We write it as a property and sketch the proof (which is a direct adaptation of what was proved in [7] and [10]) because it will be instrumental in the proof of several results in the rest of this section,

Property 4. There is a short *Res* refutation of *SPHP*¹.

Proof. The proof is based on the fact that for each one of the $m + 1$ pigeons there is a short refutation

$$\{(x_{i1} \vee x_{i2} \vee \dots \vee x_{im}, 1)\} \cup \{(\bar{x}_{ij}, 1) \mid 1 \leq j \leq m\} \vdash^* \mathcal{G} \cup (\square, 1)$$

and for each one of the m holes there is a short refutation

$$\{(\bar{x}_{ij} \vee \bar{x}_{i'j}, 1) \mid 1 \leq i < i' \leq m + 1\} \cup \{(x_{ij}, 1) \mid 1 \leq i \leq m + 1\} \vdash^* \mathcal{G} \cup \{(\square, m)\}$$

Because each derivation is independent of the other we can concatenate them into,

$$SPHP^1 \vdash^* \mathcal{G} \cup \{(\square, m^2 + m + 1)\}$$

which is a refutation of *SPHP*¹.

5.2 Split

The *split* rule,

$$\frac{(A, w)}{(A \vee x, w) \quad (A \vee \bar{x}, w)}$$

is the natural extension of its SAT counterpart. Consider the proof system **ResS**, made of resolution and split. We show that, as it happens in the SAT case, the split rule brings completeness,

Theorem 4. *ResS is complete.*

Proof. The proof is based on the following facts:

1. For every formula \mathcal{F} there is a proof $\mathcal{F} \vdash^* \mathcal{F}^e$ where \mathcal{F}^e is made exclusively of splits in which all the clauses of \mathcal{F}^e contain all the variables in the formula and there are no repeated clauses. Each clause $(C, w) \in \mathcal{F}$ can be expanded to a new variable not in C using the split rule. This process can be repeated until all clauses in the current formula contain all the variables in the formula. Note that all clauses (C', u) , (C', v) can be merged and, as a result, \mathcal{F}^e does not contain repeated clauses.

2. If $\mathcal{F} \vdash^* \mathcal{F}^e$ then $\mathcal{F}^e \vdash^* \mathcal{F}$. Let $\mathcal{F} = \mathcal{F}_0 \vdash \mathcal{F}_1 \vdash \mathcal{F}_2 \vdash \dots \vdash \mathcal{F}_p = \mathcal{F}^e$ be the proof from \mathcal{F} to \mathcal{F}^e . Then, $\mathcal{F}^e = \mathcal{F}_p \vdash \dots \vdash \mathcal{F}_2 \vdash \mathcal{F}_1 \vdash \mathcal{F}_0 = \mathcal{F}$ is done resolving the pairs of clauses in \mathcal{F}_i that were splitted in the $\mathcal{F}_{i-1} \vdash \mathcal{F}_i$ step.
3. If $\mathcal{F}^e(X) = w$ then there exists a unique clause $(C, w) \in \mathcal{F}^e$ which is falsified by X

By fact (1), $\mathcal{F} \vdash^* \mathcal{F}^e$ and $\mathcal{G} \vdash^* \mathcal{G}^e$. Because of soundness, $\vdash, \forall X, \mathcal{F}(X) = \mathcal{F}^e(X)$ and $\forall X, \mathcal{G}(X) = \mathcal{G}^e(X)$. Since $\mathcal{F} \models \mathcal{G}, \forall X, F(X) \geq G(X)$. Therefore, $\forall X, F^e(X) \geq G^e(X)$ which, by fact (3), means that for each X there exists a unique $(C, F^e(X)) \in \mathcal{F}^e$ and $(C, G^e(X))$ which is falsified by X . Separating all $(C, \mathcal{F}^e(X))$ into $(C, \mathcal{G}^e(X)), (C, \mathcal{F}^e(X) - \mathcal{G}^e(X))$ we have $\mathcal{F}^e = \mathcal{G}^e \cup \mathcal{H}^e$. Therefore, $\mathcal{F} \vdash \mathcal{F}^e = \mathcal{G}^e \cup \mathcal{H}^e$. By fact (2), $\mathcal{G}^e \cup \mathcal{H}^e \vdash^* \mathcal{G} \cup \mathcal{H}^e$.

However, unlike what happens in the SAT case (see Property 1), **ResS** is stronger than **Res**,

Theorem 5. *ResS is stronger than Res.*

Proof. On the one hand, it is clear that **ResS** can simulate any proof of **Res** since it is a superset of **Res**. On the other hand, unlike **Res**, **ResS** can produce short refutations for $SPHP^0$, as shown below.

First, let us proof that **Res** cannot produce short refutations for $SPHP^0$. Since the resolution rule does not apply to the empty clause (\square, w) , if **Res** could refute $SPHP^0$ in polynomial time it would also refute $SPHP$ in polynomial time, which is impossible [6].

ResS can produce short refutations for $SPHP^0$ because it can transform $SPHP^0$ into $SPHP^1$ and then apply Property 4. The transformation is done by a sequence of splits,

$$\frac{(\square, 1)}{(x_{ij}, 1) \quad (\bar{x}_{ij}, 1)}$$

that move one unit of weight from the empty clause to every variable in the formula and its negation.

5.3 Virtual

In a recent paper [10] we proposed a proof system in which clauses with negative weights can appear during the proof. This is equivalent to adding to **ResS** the *virtual* rule,

$$\frac{}{(A, w) \quad (A, -w)}$$

which allows to introduce a fresh clause (A, w) into the formula. To preserve soundness (i.e, cancel out the effect of the addition) it also adds $(A, -w)$.

Let **ResSV** be the proof system made of resolution, split and virtual (note that resolution and split are only defined for antecedents with positive weights).

It has been shown that if $\mathcal{F}_0 \vdash^* \mathcal{F}_e$ is a **ResSV** proof and \mathcal{F}_e does not contain any negative weight, then for every $\mathcal{G} \subseteq \mathcal{F}_e$ we have that $\mathcal{F} \models \mathcal{G}$.

The following theorem shows that the virtual rule adds further strength to the proof system,

Theorem 6. *ResSV is stronger than ResS.*

Proof. On the one hand, it is clear that **ResSV** can simulate any proof of **ResS** since it is a superset of **ResS**. On the other hand, **ResSV** can produce a short refutation of *SPHP* and **ResS** cannot.

The short refutation of **ResSV**, as shown in [10], is obtained by first virtually transforming *SPHP* into *SPHP*¹. Then, it uses Property 4 to derive $(\square, m^2 + m + 1)$. Finally, it splits one unit of the empty clause cost to each pair x_{ij}, \bar{x}_{ij} to cancel out negative weights. At the end of the process all clauses have positive weight while still having $(\square, 1)$.

It is clear that **ResS** cannot polynomially refute *SPHP* because otherwise a SAT proof system with resolution and split rules would produce shorter refutations than a SAT proof system with only resolution, which contradicts Property 1.

We will finish this section showing that Theorem 1 has an unexpected application in the context of **ResSV**. Consider the problem of proving *PHP* $\models (\square, \infty)$. This can be done with a refutation of *PHP*. Namely $PHP \vdash^* (\square, \infty) \cup \mathcal{F}$ or using Corollary 1, which tells that $\mathcal{F} \models (\square, \infty)$ if and only if $MaxSAT(\mathcal{F}) \geq 1$. The following two theorems shows that **ResSV** cannot do efficiently the first approach, but can do efficiently the second.

Theorem 7. *There is no short ResSV refutation of PHP.*

Proof. Virtual rule cannot introduce hard clauses and resolution and split rules only produce a hard consequence if they have hard antecedents. As a consequence, (\square, ∞) can only be obtained by resolving or splitting hard clauses in *PHP*. If **ResSV** produce a short refutation for *PHP*, **ResS** and, as a consequence **Res**, also produce the same short refutation for *PHP*, which contradicts Property 1.

Theorem 8. *There is a polynomial ResSV proof of $(\square, 1)$ from PHP.*

Proof. We only need to apply the virtual rule,

$$\frac{}{(\square, m^2 + m) \quad (\square, -m^2 - m)}$$

and then split,

$$\frac{(\square, 1)}{(x_{ij}, 1) \quad (\bar{x}_{ij}, 1)}$$

for each i, j . The resulting problem is similar to *SPHP*¹ but with hard clauses. At this point and adapting the proof of 4 we can derive $(\square, m^2 + m + 1)$ cancel out the negative weight while still retaining $(\square, 1)$.

6 MaxSAT Circular Proofs

In this section we study the relation between **ResSV** and the recently proposed concept of circular proofs [3]. Circular proofs allows the addition of an arbitrary set of clauses to the original formula. It can be seen that conclusions are sound as long as the added clauses are *re-derived* as many times as they are used. In the original paper this condition is characterized as the existence of a flow in a graphical representation of the proof. Here we show that the **ResSV** proof system naturally captures the same idea and extends it from SAT to MaxSAT with an arguably simpler notation. In particular, the virtual rule guarantees the existence of the flow.

6.1 SAT Circular Proofs

We start reviewing the SAT case, as defined in [3]. Given a CNF formula \mathcal{F} a *circular pre-proof* of C_r from \mathcal{F} is a sequence,

$$\Pi = C_1, C_2, \dots, C_p, C_{p+1}, C_{p+2}, \dots, C_{p+q}, C_{p+q+1}, C_{p+q+2}, \dots, C_r$$

such that $\mathcal{F} = \{C_1, C_2, \dots, C_p\}$, $\mathcal{B} = \{C_{p+1}, C_{p+2}, \dots, C_{p+q}\}$ is an arbitrary set of clauses and each C_i ($i > p + q$) is obtained from previous clauses by applying an inference rule in the proof system. Note that the same clause can be both derived and used several times during the proof.

A *circular pre-proof* Π can be associated with a directed bi-partite graph $G(\Pi) = (I \cup J, E)$ such that there is one node in J for each element of the sequence (called clause nodes) and one node in I for each inference step (called inference nodes). There is an arc from $u \in J$ to $v \in I$ if u is an antecedent clause in the inference step of v . There is an arc from $u \in I$ to $v \in J$ if v is a consequent clause in the inference step of u . The graph is compacted by merging nodes whose associated clause is identical to one in \mathcal{B} . Note that before the compactation the graph is acyclic, but the compactation may introduce cycles. The set of in-neighbors and out-neighbors of node $C \in J$ are denoted $N^-(C)$ and $N^+(C)$, respectively.

A *flow assignment* for a circular pre-proof is an assignment $f : I \rightarrow \mathbb{R}^+$ of positive reals to inference nodes. The *balance* of node $C \in J$ is the inflow minus the outflow,

$$b(C) = \sum_{R \in N^-(C)} f(R) - \sum_{R \in N^+(C)} f(R)$$

Definition 1. A SAT circular proof of clause A from CNF formula \mathcal{F} is a pre-proof Π whose proof-graph $G(\Pi)$ admits a flow in which all clauses not in \mathcal{F} have non-negative balance and $A \in J$ has a strictly positive balance.

Theorem 9 (soundness). Assuming a sound SAT proof system, if there is a SAT circular proof of A from formula \mathcal{F} then $\mathcal{F} \models A$.

Property 5. Using the proof system with the following two rules,

$$\frac{x \vee D \quad \bar{x} \vee D}{D} \quad \frac{D}{D \vee x \quad D \vee \bar{x}}$$

(symmetric resolution) (split)

there is a short circular refutation of *PHP*.

6.2 ResSV and MaxSAT Circular Proofs

Now we show that the MaxSAT **ResSV** proof system is a true extension of circular proofs from SAT to MaxSAT. The following two theorems show that, when restricted to hard formulas, **ResSV** and SAT circular proofs can simulate each other. Recall that specializing Corollary 1 to hard formulas, $\mathcal{F} \models (A, \infty)$ and $\text{MaxSAT}(\mathcal{F} \cup \{(\bar{A}, 1)\}) \geq 1$ is equivalent. Therefore, one can show $\mathcal{F} \models (A, \infty)$ with a proof $\mathcal{F} \cup \{(\bar{A}, 1)\} \vdash^* (\square, 1) \cup \mathcal{G}$.

Theorem 10. *Let Π be a SAT circular proof of clause A from formula $\mathcal{F} = \{C_1, \dots, C_p\}$ using the proof system symmetric resolution and split. There is a MaxSAT **ResSV** proof of $(\square, 1)$ from $\mathcal{F}' = \{(C_1, \infty), \dots, (C_p, \infty)\} \cup \{(\bar{A}, 1)\}$. The length of the proof is $O(|\Pi|)$.*

Proof. Let $G(\Pi) = (J \cup I, A)$ be the proof graph and $f(\cdot)$ the flow of Π . By definition of SAT circular proof, $A \in J$ and $b(A) > 0$.

The ResSV proof starts with $\mathcal{F}' = \{(C_1, \infty), \dots, (C_p, \infty)\} \cup \{(\bar{A}, 1)\}$ and consists in 3 phases. In the first phase, the virtual rule is applied for each $C \in J$ not in \mathcal{F} , introducing $\{(C, o), (C, -o)\}$ with $o = \sum_{R \in N^-(C)} f(R)$. In the second phase, there is an inference step for each $u \in I$. If u is a SAT split, the inference step is a MaxSAT split generating two clauses with weight $f(u)$. If u is a SAT symmetric resolution, the inference step is a MaxSAT resolution generating one clause with weight $f(u)$. Note that this phase never creates new clauses because all of them have been virtually added at the first phase. It only moves weights around the existing ones. Note as well that we guarantee by construction that at each step of the proof the antecedents are available no matter in which order the proof is done because the first phase has given enough weight to each added clause to guarantee it and original clauses are hard, so their weight never decreases. At the end of the second phase we have $\mathcal{F} \cup \{(\bar{A}, 1)\} \cup \mathcal{C}$ with $\mathcal{C} = \{(C, b(C)) \mid C \in J, b(C) > 0\}$ with $b(C)$ being the balance of C . Therefore $(A, b(A))$ is in \mathcal{C} . The third phase is a final sequence of q steps in which $(\square, 1)$ is derived from $\{(\bar{A}, 1), (A, b(A))\}$ which completes the proof. Note that the size of the proof is $O(|J + I|) = O(|\Pi|)$.

Theorem 11. *Consider a hard formula $\mathcal{F} = \{(C_1, \infty), \dots, (C_p, \infty)\}$ and a **ResSV** proof $\mathcal{F} \cup \{(\bar{A}, 1)\} \vdash^* \mathcal{F} \cup \mathcal{G} \cup \{(\square, 1)\}$ with e inference steps. There is a SAT circular proof Π of A from $\mathcal{F}' = \{C_1, \dots, C_p\}$ with proof system symmetric resolution and split. Besides, $|\Pi| = O(e)$.*

Proof. We need to build a graph $G(\Pi) = (J \cup I, E)$ with $\mathcal{F}' \subset J$ and $A \in J$, and a flow $f(\cdot)$ that satisfies the balance conditions and with which A has strictly positive balance.

Because the virtual rule does not have antecedents all its applications can be done at the beginning of the proof and all the cancellation of all the virtual clauses can be done at the end. Therefore, we can omit all those inference steps and assume without loss of generality that the proof is a ResS (that is, without virtual),

$$\mathcal{F} \cup \{(\bar{A}, 1)\} \cup \mathcal{B} \vdash \mathcal{F}_1 \vdash \mathcal{F}_2 \vdash \dots \vdash \mathcal{F}_e = \mathcal{F} \cup \mathcal{G} \cup \{(\square, 1)\} \cup \mathcal{B}$$

where \mathcal{B} is the set of virtually added clauses. Note as well that any application of MaxSAT resolution between $x \vee A$ and $\bar{x} \vee B$ can be simulated by a short sequence of splits to both clauses until their scope is the same and then one resolution step between $x \vee A \vee B$ and $\bar{x} \vee A \vee B$. So, again without loss of generality we can assume that the proof only contains splits and symmetric resolutions.

Our proof contains three phases. First, we are going to build an acyclic graph $G'(\Pi)$ which is an unfolded version of $G(\Pi)$ and a flow function $f'(\cdot)$ that may have ∞ flows. Second we will compute $f(\cdot)$ traversing the graph $G'(\Pi)$ bottom-up and replacing any infinite flow in $f'(\cdot)$ by a finite one that still guarantees the flow condition. In the third and final phase, we will compact the graph which will constitute the circular proof.

Phase 1:

We build $G'(\Pi) = (J' \cup I', E')$ by following the proof step by step. Let $G'_i = (J'_i \cup I'_i, E'_i)$ be the graph associated to proof step i . We define the *front* of G_i as the set of clause nodes in J'_i with strictly positive balance. By construction of G'_i we will guarantee a connection between the current formula \mathcal{F}_i and the front of the current graph G'_i

$$\mathcal{F}_i = \{(C, b(C)) \mid C \in J'_i, b(C) = \sum_{R \in N^-(C)} f'(R) - \sum_{R \in N^+(C)} f'(R)\}$$

where we define $\infty - \infty = \infty$.

G'_1 contains one clause node for each clause in \mathcal{F} , $\{(\bar{A}, 1)\}$ and \mathcal{B} , respectively. For each clause node there is one dummy inference node pointing to it. The flow $f'(\cdot)$ of the inference node is the weight of the clause it points to. This set of dummy inference nodes will be removed at step three. Then we proceed through the proof. At inference step i , we add a new inference node i to I . Its in-neighbors will be nodes from the front (that must exist because of the invariant) and its out-neighbors will be newly added clause nodes. Its flow $f'(i)$ is the weight moved by the inference rule (which may be infinite). If the inference rule is split we add two clause nodes, one for each consequent and add the corresponding arcs. If the inference rule is a resolution we add one clause node for its consequent and add the corresponding arcs. Note that, the out-neighbors of node i have a positive balance and in-neighbors of i have their out-flow decreased, but cannot turn negative. Finally, we merge any pair of nodes in the front of G'_i whose associated clause is the same (which preserves the property of balances being

non-negative). Graph G' is obtained after processing the last inference step. Note that the invariant guarantees that \square is in G' and its balance is 1.

Phase 2:

Now we traverse the inference nodes of G' in the reverse order of how they were added transforming infinite flows into finite. When considering node i , because of the traversing order, we know that every $C \in N^+(i)$ has finite outgoing flow. We compute the flow value $f(i)$ as follows: if $f'(i)$ is finite, then $f(i) = f'(i)$, else $f(i)$ is the minimum value that guarantees that the balance of every $C \in N^+(i)$ is non-negative.

Phase 3:

We obtain G by doing some final arrangements to G' . First, we remove dummy inference nodes pointing to clauses in \mathcal{F} , $(\overline{A}, 1)$ and \mathcal{B} added in Phase 1. As a result, the balance of these nodes is negative. In particular, the balance of nodes representing \overline{A} and \mathcal{B} is its negative weight.

Since $\mathcal{B} \subseteq \mathcal{F}_e$, we know that all nodes representing \mathcal{B} are included in the front of G' with balance greater than or equal to its weight. We compact these nodes with the ones in G'_1 and, as a result, its balance is positive.

Finally, we add some split nodes with flow 1 from node \square (recall that $b(\square) = 1$) in order to generate A and \overline{A} , and we compact the latter ones with the ones in G'_1 . As a result, the balance of A is 1 and the balance of \overline{A} nodes is positive.

7 Conclusions

This paper constitutes a first attempt towards MaxSAT resolution-based proof complexity analysis. We have provided some basic definitions and results emphasizing the similarities and differences with respect to SAT. In particular, we have shown that MaxSAT entailment can be rephrased as a MaxSAT refutation problem and, as a consequence, refutation completeness is sufficient for practical purposes. Interestingly, when such rephrasing is applied to hard formulas it transforms a SAT query into a MaxSAT one, and such transformation turns out to be relevant in our analysis of SAT circular proofs.

We have also provided three basic inference MaxSAT rules used in resolution-based proof systems (e.g. resolution, split and virtual) and have analysed their incremental effect in terms of refutation power. Finally, we have related **ResSV**, the strongest of the proof systems considered, with the recently proposed concept of circular proofs. We have shown that **ResSV** generalizes SAT circular proofs as defined in [3].

An additional contribution of the paper is to put together under a formal framework and common notation some ideas spread around in different recent papers such as [3, 7, 10].

References

1. Ansótegui, C., Bonet, M.L., Levy, J.: SAT-based maxsat algorithms. *Artif. Intell.* **196**, 77–105 (2013). <https://doi.org/10.1016/j.artint.2013.01.002>
2. Atserias, A.: On sufficient conditions for unsatisfiability of random formulas. *J. ACM* **51**(2), 281–311 (2004). <https://doi.org/10.1145/972639.972645>
3. Atserias, A., Lauria, M.: Circular (yet sound) proofs. In: Janota, M., Lynce, I. (eds.) *SAT 2019*. LNCS, vol. 11628, pp. 1–18. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-24258-9_1
4. Bacchus, F., Hyttinen, A., Järvisalo, M., Saikko, P.: Reduced cost fixing for maximum satisfiability. In: Lang, J. (ed.) *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI 2018)*, 13–19 July (2018), Stockholm, Sweden, pp. 5209–5213 (2018). [ijcai.org. https://doi.org/10.24963/ijcai.2018/723](https://doi.org/10.24963/ijcai.2018/723)
5. Biere, A., Heule, M., van Maaren, H., Walsh, T. (eds.): *Handbook of Satisfiability, Frontiers in Artificial Intelligence and Applications*, vol. 185. IOS Press, Amsterdam (2009)
6. Bonet, M.L., Levy, J., Manyà, F.: Resolution for Max-Sat. *Artif. Intell.* **171**(8–9), 606–618 (2007). <https://doi.org/10.1016/j.artint.2007.03.001>
7. Ignatiev, A., Morgado, A., Marques-Silva, J.: On tackling the limits of resolution in SAT solving. In: Gaspers, S., Walsh, T. (eds.) *SAT 2017*. LNCS, vol. 10491, pp. 164–183. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66263-3_11
8. Larrosa, J., Heras, F.: Resolution in Max-Sat and its relation to local consistency in weighted CSPs. In: Kaelbling, L.P., Saffiotti, A. (eds.) *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, Edinburgh, Scotland, UK, 30 July–5 August 2005, pp. 193–198. Professional Book Center (2005). <http://ijcai.org/Proceedings/05/Papers/0360.pdf>
9. Larrosa, J., Heras, F., de Givry, S.: A logical approach to efficient Max-SAT solving. *Artif. Intell.* **172**(2–3), 204–233 (2008). <https://doi.org/10.1016/j.artint.2007.05.006>
10. Larrosa, J., Rollon, E.: Augmenting the power of (partial) maxsat resolution with extension. In: *Proceedings of the Thirty-third AAAI Conference on Artificial Intelligence (AAAI 2020)*, New York, New York, USA, 7–12 February 2020
11. Morgado, A., Heras, F., Liffiton, M.H., Planes, J., Marques-Silva, J.: Iterative and core-guided MaxSAT solving: a survey and assessment. *Constraints* **18**(4), 478–534 (2013). <https://doi.org/10.1007/s10601-013-9146-2>
12. Robinson, J.A.: A machine-oriented logic based on the resolution principle. *J. ACM* **12**(1), 23–41 (1965). <https://doi.org/10.1145/321250.321253>