The Max-Atom Problem and its Relevance

Marc Bezem*, Robert Nieuwenhuis**, and Enric Rodríguez-Carbonell**

Abstract. Let F be a conjunction of atoms of the form $\max(x, y)+k \geq z$, where x, y, z are variables and k is a constant value. Here we consider the satisfiability problem of such formulas (e.g., over the integers or rationals). This problem, which appears in unexpected forms in many applications, is easily shown to be in NP. However, decades of efforts (in several research communities, see below) have not produced any polynomial decision procedure nor an NP-hardness result for this -apparently so simple- problem.
Here we develop several ingredients (small-model property and lattice structure of the model class, a polynomially tractable subclass and an inference system) which altogether allow us to prove the existence of small unsatisfiability certificates, and hence membership in NP intersection co-NP. As a by-product, we also obtain a weakly polynomial decision procedure.
We show that the Max-atom problem is PTIME-equivalent to several other well-known -and at first sight unrelated- problems on hypergraphs and on Discrete Event Systems, problems for which the existence of PTIME algorithms is also open. Since there are few interesting problems in NP intersection co-NP that are not known to be polynomial, the Max-atom problem appears to be relevant.

Keywords: constraints, max-plus algebra, hypergraphs.

1 Introduction

Difference Logic (DL) is a well-known fragment of linear arithmetic in which atoms are constraints of the form $x+k \geq y$, where x, y are variables and the *offset* k is a constant value. Due to its many applications to verification (e.g., timed automata), it is one of the most ubiquitous theories in the context of Satisfiability Modulo Theories (SMT). In SMT systems, a *theory solver* is essentially a decision procedure for the satisfiability of *conjunctions* of theory atoms. For DL satisfiability is equivalent to the absence of negative cycles in the digraph having one edge $x \xrightarrow{k} y$ for each atom $x+k \geq y$, and can be decided in polynomial time (e.g., by the Bellman-Ford algorithm; cf. [NOT06] for background on SMT and algorithms for DL, among other theories).

Motivated by the need of SMT techniques for reasoning about delays in digital circuits, it is natural to extend the atoms of DL to *max-atoms* of the form $\max(x, y) + k \geq z$. The satisfiability of conjunctions of such constraints appears to be a new problem, hereafter referred to as the *Max-atom problem.*

* Department of Informatics, University of Bergen, Norway.
** Technical University of Catalonia (UPC), Barcelona, Spain. Partially supported by Spanish Ministry Educ. and Science LogicTools-2 project (TIN2007-68093-C02-01).

The Max-atom problem is easily seen to belong to NP, since after guessing in each atom $\max(x, y)+k \geq z$ which one of x and y is the maximal variable, the problem reduces to DL. As in DL, there is no essential difference here between interpretations over integers or rationals[1]: Given a conjunction of n atoms with rational offsets $\max(x_i, y_i)+p_i/q_i \geq z_i$, for i in $1 \ldots n$, if lcm is the least common multiple of the $q_i's$, one can express each atom as $\max(x_i, y_i)+r_i/lcm \geq z_i$ for certain r_i's and solve the equisatisfiable conjunction of atoms $\max(x_i, y_i)+r_i \geq z_i$ over the integers. Therefore, unless explicitly stated otherwise, here we will only consider integer models and offsets.

The language of conjunctions of max-atoms of the form $\max(x, y)+k \geq z$ is quite expressive, and many interesting problems can be modeled by polynomially many such max-atoms. Some simple examples follow. DL literals $x+k \geq y$ can of course be expressed as $\max(x, x)+k \geq y$. Equalities $\max(x, y)+k = z$ can be written as $\max(x, y)+k \geq z \;\wedge\; z-k \geq x \;\wedge\; z-k \geq y$. Strict inequalities $\max(x, y)+k > z$ can be expressed as $\max(x, y)+k-1 \geq z$. One can express max on both sides, as in $\max(x, y)+k = \max(x', y')+k'$ by introducing a fresh variable z and writing $\max(x, y)+k = z \;\wedge\; \max(x', y')+k' = z$. One can also express different offsets on different arguments of max; for instance $\max(x+5, y-3) \geq z$ can be written as $\max(x, y')+5 \geq z \;\wedge\; y'+8 = y$, where y' is fresh. Furthermore, since $\max(e_1, e_2, e_3)$ is equivalent to $\max(e_1, \max(e_2, e_3))$, one can express nested or larger-arity max-atoms such as $\max(e_1, e_2, e_3) \geq z$ by writing $\max(e_1, x) \geq z \;\wedge\; \max(e_2, e_3) = x$, where x is fresh.

A less simple equivalence (see Section 5) exists with a problem used in Control Theory for modeling Discrete Event Systems. It amounts to solving *two-sided linear max-plus systems:* sets of equations of the form

$\max(\ x_1+k_1,\ \ldots,\ x_n+k_n\ ) = \max(\ x_1+k'_1,\ \ldots,\ x_n+k'_n\ )$

where *all* n variables of the system occur on both sides of every equation, which makes it non-trivial to show that max-atoms can be equivalently expressed in this form. Finding a polynomial algorithm for this problem has been open for more than 30 years in the area of max-plus algebras [BZ06]. An elegant algorithm was given and claimed to be polynomial in [BZ06], but unfortunately in [BNRC08] we have given an example on which it behaves exponentially. Currently still no polynomial algorithm is known.

Yet another equivalent problem (see again Section 5) concerns shortest paths in directed weighted hypergraphs. In such hypergraphs, an edge goes from a set of vertices to another vertex. Hence a natural notion of a hyperpath (from a set of vertices to a vertex) is a tree, and a natural notion of length of the hyperpath is the maximal length (the sum of the weights) of a path from a leaf to the root of this tree. For arbitrary directed hypergraphs with positive or negative weights, no polynomial algorithm for determining (the existence of) such shortest hyperpaths has been found.

Slight increases in expressive power lead to NP-hardness. For instance, having both max and min it is easy to express any 3-SAT problem with variables $x_1 \ldots x_n$, by: (i) an atom $T > F$ (T, F are variables); (ii) for all x_i the atoms

[1] Except, possibly, for the weakly polynomial algorithm that will be given in Section 3.

$\min(x_i, x_i') = F$ and $\max(x_i, x_i') = T$; (iii) for each clause like $x_p \vee \overline{x_q} \vee x_r$, an atom $\max(x_p, x_q', x_r) \geq T$.

Altogether, decades of efforts in the hypergraph and the max-plus communities have not produced any polynomial decision procedure nor an NP-hardness result for the different versions of the –apparently so simple– Max-atom problem. In this paper we give several interesting new insights.

In Section 2 we first prove some relevant results on the models of sets (conjunctions) of max-atoms: we give a small-model property, and show that the model class is a (join semi-) lattice. These properties allow us to prove that a set of max-atoms is unsatisfiable if, and only if, it has an unsatisfiable subset which is *right-distinct*, i.e., where each variable occurs at most once as a right-hand side of a max-atom.

In Section 3 we define *max-derivations* as transformation systems on states (assignments to the variables) as a formalism for searching models, and use the properties of the previous section to obtain a weakly polynomial algorithm for the integers, which is also a strongly polynomial one for a relevant subclass of problems.

In Section 4 we define a *chaining* inference system for max-atoms of the form $\max(x_1 + k_1, \ldots, x_m + k_m) \geq z$, and building upon the previous results we show that it is sound and refutation complete. Moreover, we prove that for right-distinct sets chaining can be turned into a polynomial-time decision procedure, thus showing that the Max-atom problem is in co-NP (one only needs to guess the small unsatisfiability certificate: the right-distinct unsatisfiable subset). Since there are few interesting problems in NP $\cap$ co-NP that are not known to be polynomial, this one appears to be relevant. Moreover, given the history of problems in this class, such as deciding primality [AKS04], there is hope for a polynomial-time algorithm.

The paper ends with the proofs of equivalence with solving two-sided linear max-plus systems and shortest paths in hypergraphs (Section 5) and the conclusions (Section 6).

2 Models of Conjunctions of Max-atoms

The following lemma ensures that models of a set of max-atoms are invariant under "uniform" translations:

Lemma 1. *Given a set of max-atoms S defined over the variables V and an assignment $\alpha : V \to \mathbb{Z}$ which is a model of S, for any $d \in \mathbb{Z}$ the assignment α' defined by $\alpha'(x) = \alpha(x) + d$ is a model of S.*

Definition 1. *Given a set of variables V, the* size *of an assignment $\alpha : V \to \mathbb{Z}$ is the difference between the largest and the smallest value assigned to the variables, i.e.,* $\mathrm{size}(\alpha) = \max_{x,y \in V}(\alpha(x) - \alpha(y))$.

Lemma 2 (Small Model Property). *If a set of max-atoms S is satisfiable, then it has a model of size at most the sum of the absolute values of the offsets,* *i.e., at most*

$$K_S = \sum_{\max(x,y)+k \geq z \ \in \ S} |k|.$$

Proof. We may assume that all constraints in the set are equations: replace each max-atom $\max(x, y) + k \geq z$ by the constraints $\max(x, y) + k = z'$ and $\max(z, z') = z'$. The class of models does not change essentially by adding these auxiliary constraints and variables, as one just has to add/omit interpretations for the fresh variables. Furthermore, the sum of the absolute values of the offsets does not change. Therefore, we may assume that S is a set of constraints of the form $\max(x, y) + k = z$ (where possibly x and y are the same variable).

Let α be a model of S. Based on α we define a weighted graph whose vertices are the variables. For every constraint $\max(x, y) + k = z$, if $\alpha(x) \geq \alpha(y)$ then we add a red edge (x, z) with weight k and a green edge (y, x) without a weight; and otherwise, if $\alpha(y) > \alpha(x)$ then we add a red edge (y, z) with weight k and a green edge (x, y) without a weight. While changing the model, the graph will remain all the time the same.

A red (weakly) connected component is a subgraph such that there are red paths between any two variables in the subgraph, where the red edges may be used in any direction. The *segment* of a red connected component is the range of integers from the lowest value to the highest one assigned to the variables in the component. The size of such a segment is at most the sum of the absolute values of the weights of the edges in the component.

Red connected components partition the set of variables. If their segments overlap, then already $\mathrm{size}(\alpha) \leq K_S$. If there is a gap somewhere, say of size p, then the gap is closed by a suitable translation, e.g., by decreasing by p all values assigned to variables above the gap. This respects all red edges and their weights since the gap is between segments of red connected components and components are translated as a whole. Green edges are also respected since we only close gaps and never a variable x with initially a higher value than another variable y ends up with a value strictly lower than y. Since all edges are respected we keep a model, all the time closing gaps until there are no gaps left. We end up with a model α' without gaps and hence $\mathrm{size}(\alpha') \leq K_S$. $\square$

The previous lemma gives an alternative proof of membership in NP of the Max-atom problem: it suffices to guess a "small" assignment; checking that it is indeed a model is trivially in P.

Lemma 3 below proves that the model class of a set of max-atoms is a (join semi-) lattice, where the partial ordering is $\geq$ (pointwise $\geq$):

Definition 2. *Given a set of variables V and assignments $\alpha_1, \alpha_2 : V \to \mathbb{Z}$, we write $\alpha_1 \geq \alpha_2$ if for all $x \in V$, $\alpha_1(x) \geq \alpha_2(x)$.*

Definition 3. *Given a set of variables V and two assignments $\alpha_1, \alpha_2 : V \to \mathbb{Z}$, the* supremum *of α_1 and α_2, denoted by $\sup(\alpha_1, \alpha_2)$, is the assignment defined by $\sup(\alpha_1, \alpha_2)(x) = \max(\alpha_1(x), \alpha_2(x))$ for all $x \in V$.*

Lemma 3. *Given a set of max-atoms S defined over the variables V and two assignments $\alpha_1, \alpha_2 : V \to \mathbb{Z}$, if $\alpha_1 \models S$ and $\alpha_2 \models S$ then $\sup(\alpha_1, \alpha_2) \models S$.*

Proof. Let us denote $\sup(\alpha_1, \alpha_2)$ by α^*. Assume $\alpha_1 \models S$ and $\alpha_2 \models S$ and let $\max(x, y)+k \geq z$ be an atom in S. By assumption, $\max(\alpha_i(x), \alpha_i(y))+k \geq \alpha_i(z)$ for $i = 1, 2$. Also by definition for $i = 1, 2$ we have $\alpha^*(x) \geq \alpha_i(x)$ and $\alpha^*(y) \geq \alpha_i(y)$, so $\max(\alpha^*(x),\ \alpha^*(y)) + k \geq \alpha_i(z)$. Thus $\max(\alpha^*(x),\ \alpha^*(y)) + k \geq \alpha^*(z)$, that is, the atom $\max(x, y) + k \geq z$ is satisfied by α^*. Hence $\alpha^* \models S$. □

Using the previous lemmas, we have the following result:

Lemma 4. *Let S be a set of max-atoms, and let z be a variable such that for some $r > 1$ all max-atoms with z as right-hand side are L_1, ..., L_r. The set S is satisfiable if and only if all $S - \{L_i\}$ (i in $1 \ldots r$) are satisfiable.*

Proof. The "only if" implication is trivial, since $S - \{L_i\} \subseteq S$ for all i in $1 \ldots r$. Now, for the "if" implication, let α_i be a model of $S - \{L_i\}$. By Lemma 1, for every i in $2 \ldots r$ we can assume w.l.o.g. that $\alpha_i(z) = \alpha_1(z)$. Let us define $\alpha^* = \sup(\alpha_1, ..., \sup(\alpha_{r-1}, \alpha_r)...)$. Then $\alpha^*(z) = \alpha_i(z)$ for all i in $1 \ldots r$. Moreover, since for all i in $1 \ldots r$ we have in particular $\alpha_i \models S - \{L_1, ..., L_r\}$, by iterating Lemma 3, $\alpha^* \models S - \{L_1, ..., L_r\}$. It remains to be seen that $\alpha^* \models L_i$ for any i in $1 \ldots r$. Let thus $\max(x, y) + k \geq z$ be L_i, for a given i in $1 \ldots r$. Since $r > 1$, there is j in $1 \ldots r$ such that $i \neq j$. Since $\alpha_j \models S - \{L_j\}$ and $i \neq j$, $\alpha_j \models L_i$. So $\max(\alpha^*(x),\ \alpha^*(y)) + k \geq \max(\alpha_j(x),\ \alpha_j(y)) + k \geq \alpha_j(z) = \alpha^*(z)$. Hence $\alpha^* \models L_i$. □

The next definition and lemma will be paramount for building short certificates of unsatisfiability:

Definition 4. *A set of max-atoms S is said to be* right-distinct *if variables occur at most once as right-hand sides, i.e., for every two distinct max-atoms $\max(x, y) + k \geq z$ and $\max(x', y') + k \geq z'$ in S we have $z \neq z'$.*

Lemma 5. *Let S be a set of max-atoms. If S is unsatisfiable, then there exists an unsatisfiable right-distinct subset $S' \subseteq S$.*

Proof. Let V be the set of variables over which S is defined. Let us prove the result by induction on $N = |S| - |\{z \in V \mid z \text{ appears as a right-hand side in } S\}|$:

- Base step: $N = 0$. Then all variables appearing as right-hand sides are different. So S is right-distinct, and we can take $S' = S$.
- Inductive step: $N > 0$. Then there is a variable which appears at least twice as a right-hand side. Let z be such a variable. By Lemma 4, since S is unsatisfiable, there exists an atom $L \in S$ with right-hand side z such that $S - \{L\}$ is unsatisfiable. Now, by induction hypothesis on $S - \{L\}$ there is an unsatisfiable right-distinct set $S' \subseteq S - \{L\} \subset S$. □

3 Max-derivations

W.l.o.g. in this section max-atoms are of the form $\max(x, y) + k \geq z$ with $x \neq z$, $y \neq z$. This can be assumed by removing trivial contradictions $\max(x, x) + k \geq x$ $(k < 0)$, trivial tautologies $\max(x, y) + k \geq x$ $(k \geq 0)$, and by replacing $\max(x, y) + k \geq x$ by $\max(y, y) + k \geq x$ if $k < 0$ and $x \neq y$.

Definition 5. *Given a set of max-atoms S defined over the variables V and two assignments α, α', we write $\alpha \rightarrow_S \alpha'$ (or simply $\alpha \rightarrow \alpha'$, if S is understood from the context) if there is a max-atom* $\max(x, y) + k \geq z \in S$ *such that:*

1. $\alpha'(z) = \max(\alpha(x), \alpha(y)) + k$
2. $\alpha'(z) < \alpha(z)$ *(hence we say that z* decreases *in this step)*
3. $\alpha'(u) = \alpha(u)$ *for all $u \in V$, $u \neq z$.*

Any sequence of steps $\alpha_0 \rightarrow \alpha_1 \rightarrow \cdots$ is called a max-derivation *for S.*

Lemma 6. *Let S be a set of max-atoms defined over the variables V. An assignment $\alpha : V \rightarrow \mathbb{Z}$ is a model for S if and only if α is* final*, i.e., there is no α' such that $\alpha \rightarrow \alpha'$.*

The following lemma expresses that max-derivations, while decreasing variables, never "break through" any model:

Lemma 7. *Let S be a set of max-atoms and let α be a model of S. If $\alpha_0 \rightarrow \cdots \rightarrow \alpha_m$ and $\alpha_0 \geq \alpha$, then $\alpha_m \geq \alpha$.*

Proof. By induction over m, the length of the derivation. For $m = 0$ there is nothing to prove. Now, if $m > 0$ the step $\alpha_0 \rightarrow \alpha_1$ is by an atom $\max(x, y) + k \geq z$. Let us prove that $\alpha_1 \geq \alpha$. We only need to show that the inequality holds for the variable that changes, which is z; and indeed $\alpha_1(z) = \max(\alpha_0(x), \alpha_0(y)) + k \geq \max(\alpha(x), \alpha(y)) + k \geq \alpha(z)$. Now, by induction hypothesis $\alpha_m \geq \alpha$. □

The previous lemma, together with the Small Model Property (Lemma 2), provides us with a weakly polynomial algorithm, i.e., an algorithm whose runtime is polynomial in the input size if numbers are encoded in unary.

Theorem 1. *The Max-atom problem over the integers is weakly polynomial.*

Proof. Let S be a conjunction of max-atoms, with variables V, where $|V| = n$. For deciding the satisfiability of S one can construct an arbitrary max-derivation, starting, e.g., from the assignment α_0 with $\alpha_0(x) = 0$ for all x in V. At each step, one variable decreases by at least one. If S is satisfiable, by the Small Model Property and by Lemma 1, there is a model α such that $-K_S \leq \alpha(x) \leq 0$ for all x in V. Moreover, by the previous lemma, no variable x will ever get lower than $\alpha(x)$ in the derivation. Altogether this means that, if no model is found after $n \cdot K_S$ steps, then S is unsatisfiable. □

Note that the previous result does not directly extend to the case of the rationals since the transformation described in the introduction may produce an exponential blow-up in the *value* of the offsets.

As a corollary of the proof of the previous theorem, we obtain a PTIME decision procedure for sets of atoms of the forms $\max(x, y) \geq z$ or $\max(x, y) > z$. More generally, this also applies to K-*bounded sets*, where in S the absolute values of all offsets are bounded by a given constant K.

Example 1. Let S be the following set of max-atoms:

$$S = \{u - 10 \geq x, \quad z \geq y, \quad \max(x, y) - 1 \geq z, \quad \max(x, u) + 25 \geq z\},$$

and let α_0 be the assignment with $\alpha_0(x) = \alpha_0(y) = \alpha_0(z) = \alpha_0(u) = 0$. This initial assignment α_0 violates $u - 10 \geq x$, which allows us to decrease x and assign it the value -10: in terms of max-derivations $\alpha_0 \to \alpha_1$, where α_1 is the assignment with $\alpha_1(x) = -10$, $\alpha_1(y) = \alpha_1(z) = \alpha_1(u) = 0$.

Now the assignment α_1 only violates $\max(x, y)-1 \geq z$, which forces z to take the value -1: in terms of max-derivations, $\alpha_1 \to \alpha_2$, where α_2 is the assignment with $\alpha_2(x) = -10$, $\alpha_2(y) = 0$, $\alpha_2(z) = -1$, $\alpha_2(u) = 0$. Then α_2 only violates $z \geq y$, which forces y to take the value -1 too: $\alpha_2 \to \alpha_3$, where α_3 is the assignment with $\alpha_3(x) = -10$, $\alpha_3(y) = \alpha_3(z) = -1$, $\alpha_3(u) = 0$.

It is easy to see that 11 iterations of each of the last two steps will be needed to find a model: finally we will have a derivation $\alpha_0 \to^* \alpha$ with $\alpha(x) = -10$, $\alpha(y) = \alpha(z) = -11$, $\alpha(u) = 0$; since there is no α' such that $\alpha \to \alpha'$, α is a model of S, hence S is satisfiable.

Notice that, if we replace 10 in S by larger powers of 10, we get a family of inputs whose sizes increase linearly, but for which the number of steps of the max-derivations reaching to a model grows exponentially. Since the number of steps is polynomial in the *value* of the offsets, and not in the *sizes* of the offsets, the algorithm based on max-derivations can be *weakly* polynomial but not polynomial.

Now, if we consider the set of max-atoms $S' = S \cup \{\max(x, y) + 9 \geq u\}$, we note that α above does not satisfy the new constraint. So we can decrease u and assign it the value -1, which makes $u - 10 \geq x$ false and forces x to take the value -11. Then $\max(x, y) - 1 \geq z$ is violated, and z is decreased to -12. Finally $z \geq y$ becomes false, so y is assigned -12. The loop of these four steps can be repeated over and over, making all variables decrease indefinitely. Thus, S' is unsatisfiable as no model is found within the bound of $n \cdot K_S$ steps given in the previous theorem.

4 Chaining Inference System and Membership in Co-NP

In this section we deal with the (equivalent in expressive power) language of max-atoms of the form $\max(x_1{+}k_1, ..., x_n{+}k_n) \geq z$. Here, T always stands for a max-expression of the form $\max(y_1{+}k'_1, \ldots, y_m{+}k'_m)$ with $m \geq 0$; when written inside a max-expression the whole expression is considered flattened, so then $\max(T, z{+}k)$ represents $\max(\ y_1{+}k'_1, \ldots, y_m{+}k'_m,\ z{+}k\ )$.

Definition 6. *The* Max-chaining *inference rule is the following:*

$$\frac{\max(x_1{+}k_1, ..., x_n{+}k_n) \geq y \qquad \max(T, y{+}k) \geq z}{\max(T, x_1{+}k_1{+}k, \ldots, x_n{+}k_n{+}k) \geq z} \qquad \textit{(Max-chaining)}$$

Definition 7. *The* Max-atom simplification rules *are as follows:*

$$\frac{\max(T, x{+}k) \geq x}{\max(T) \geq x} \qquad \textit{if } k < 0 \qquad \textit{(Max-atom simplification-1)}$$

$$\frac{\max(T, x{+}k, x{+}k') \geq y}{\max(T, x{+}k') \geq y} \qquad \textit{if } k \leq k' \qquad \textit{(Max-atom simplification-2)}$$

Theorem 2. *The Max-chaining rule and the Max-atom simplification rules are sound, i.e., the conclusions of the inference rules are logical consequences of their respective premises. Moreover, for each one of the Max-atom simplification rules, the conclusion and the premise are logically equivalent.*

Theorem 3. *Max-chaining, together with the Max-atom simplifications rules, is refutation complete. That is, if S is an unsatisfiable set of max-atoms that is closed under the Max-chaining and Max-atom simplification rules, then there is a* contradiction *in S, i.e., a max-atom of the form* $\max() \geq x$.

Proof. We prove a slightly stronger result, namely the refutation completeness with a concrete *ordered* application strategy, assuming an ordering on the variables $x_1 > \ldots > x_n$ occurring in S. We prove that if there is no contradiction in S then S is satisfiable. This is done by induction on n.
Base case: if $n = 1$ all atoms in S are of the form $\max(x{+}k_1, \ldots, x{+}k_m) \geq x$, with $m \geq 1$, and where at least one of the k_i is positive (otherwise *Max-atom simplification-1* generates the contradiction $\max() \geq x$). Therefore these max-atoms are tautologies and hence satisfiable.
Induction step. Assume $n > 1$. Let S_1 be the subset of S of its max-atoms in which the variable x_1 occurs. Let SR_1 and SL_1 be the subsets of S_1 of max-atoms in which x_1 occurs exactly once, only at the right-hand sides and only at the left-hand sides, respectively. By an easy induction applying the previous theorem, all max-atoms in S_1 are logical consequences of the ones in SR_1 and SL_1, since S is closed under the Max-Simplification rules. Let S_1' be the set of the $|SR_1| \cdot |SL_1|$ max-atoms that can be obtained by applying the max-chaining steps on x_1 between max-atoms of these two sets. Now let S_2 be the set $S \setminus S_1$. Note that it is closed under the Max-chaining and Max-atom simplification rules and that $S_2 \supseteq S_1'$. Since S_2 has one variable (x_1) less than S, by induction hypothesis there exists a model α for S_2.

We will now *extend* α to a model α' for S. That is, we will have $\alpha'(x_i) = \alpha(x_i)$ for all $i > 1$, and in addition α' will also be defined for x_1, in such a way that $\alpha' \models SR_1 \cup SL_1$, which implies $\alpha' \models S_1$, and hence, since $\alpha \models S_2$, we will obtain $\alpha' \models S$.

Let SR_1 be of the form $\{\ T_1 \geq x_1,\ \ldots,\ T_m \geq x_1\ \}$ $(m > 0)$, and let $\alpha(T)$ denote the evaluation of T under the assignment α. [2] Now we define $\alpha'(x_1)$ to be $\min(\ \alpha(T_1), \ldots, \alpha(T_m)\ )$. W.l.o.g., say, $\alpha'(x_1) = \alpha(T_1)$. Let T_1 be of the form $\max(y_1 + k_1, \ldots, y_m + k_m)$, so that $\alpha'(x_1) = \max(\alpha(y_1) + k_1, \ldots, \alpha(y_m) + k_m)$. Clearly α' satisfies by construction all atoms in SR_1. It only remains to show that α' is also a model of SL_1, i.e., of the atoms of the form $\max(x_1+k, T) \geq z$. For each such atom in SL_1, the corresponding conclusion by max-chaining with the atom $\max(y_1+k_1, \ldots, y_m+k_m) \geq x_1$ is the atom $\max(y_1+k_1+k, \ldots, y_m+k_m+k,\ T\ ) \geq z$, which is in S_2 and is hence satisfied by α. So, as $\alpha'(x_1) = \max(\alpha(y_1)+k_1, \ldots, \alpha(y_m)+k_m)$, also $\max(x_1+k, T) \geq z$ is satisfied by α'. $\square$

Notice that the algorithm described in the proof of the previous theorem is a generalization of the Fourier-Motzkin elimination procedure.

Example 2. Let us consider again the system introduced in Example 1 extended with $\max(x, y) + 9 \geq u$, which makes it unsatisfiable. Atoms are written now in the format used in this section.

$$\begin{array}{llll} \{\ \max(u-10) & \geq x, & \max(z) & \geq y, \\ \ \ \max(x-1, y-1) & \geq z, & \max(x+25, u+25) & \geq z \\ \ \ \max(x+9, y+9) \geq u\} & & & \end{array}$$

By applying a closure strategy as described in the proof of the previous theorem, we get a contradiction:

Rule	**Set of Max-Atoms**
	$\max(u-10) \geq x,\quad \max(z) \geq y,$ $\max(x-1, y-1) \geq z,\ \max(x+25, u+25) \geq z$ $\max(x+9, y+9) \geq u$
max-chaining x	$\max(z) \geq y,\ \max(u-11, y-1) \geq z,$ $\max(u+15, u+25) \geq z, \max(u-1, y+9) \geq u$
atom-simplification-2	$\max(z) \geq y,\ \max(u-11, y-1) \geq z,$ $\max(u+25) \geq z,\ \max(u-1, y+9) \geq u$
atom-simplification-1	$\max(z) \geq y,\ \max(u-11, y-1) \geq z,$ $\max(u+25) \geq z,\ \max(y+9) \geq u$
max-chaining y	$\max(u-11, z-1) \geq z,\ \max(u+25) \geq z,$ $\max(z+9) \geq u$
atom-simplification-1	$\max(u-11) \geq z,\ \max(u+25) \geq z,\ \max(z+9) \geq u$
max-chaining z	$\max(u-2) \geq u,\ \max(u+34) \geq u$
atom-simplification-1	$\max() \geq u$

Theorem 4. *The Max-atom problem for right-distinct sets is decidable in polynomial time.*

[2] Note that when $SR_1 = \emptyset$, if SL_1 has the form $\{\ \max(x_1+k_1, T_1) \geq z_1,\ \ldots,\ \max(x_1+k_n, T_n) \geq z_n\ \}$ $(n > 0)$ one just needs to define $\alpha'(x_1) = \max(\ \alpha(z_1)-k_1, \ldots, \alpha(z_n)-k_n\ )$. If $SL_1 = \emptyset$ too, then $\alpha'(x_1)$ can be defined arbitrarily.

Proof. For right-distinct sets, the closure process eliminating variables one by one, as explained in the refutation completeness proof, can be done in polynomial time if the Max-atom simplification rules are applied eagerly. The proof shows that after each Max-atom simplification step, its premise can be ignored (i.e., removed) once the conclusion has been added, and that tautologies of the form $\max(\ldots, x+k, \ldots) \geq x$ with $k \geq 0$ can also be ignored. Eliminating one variable x can then be done in polynomial time, since there is only one leftmost premise of chaining with x. After eliminating x, a new right-distinct set of max-atoms with one variable less and at least one atom less is obtained, in which each atom has arity bounded by the number of variables and the size of the offsets is bounded by the sum of the sizes of the offsets in the input. □

Example 3. In the previous example, an unsatisfiable right-distinct subset is:

$$\{\ \max(u-10) \geq x,\ \ \max(z) \geq y,\ \ \max(x-1, y-1) \geq z,\ \ \max(x+9, y+9) \geq u\ \}.$$

Applying the polynomial-time closure we get a contradiction:

Rule	**Set of Max-Atoms**
	$\max(u-10) \geq x$ $\max(z) \geq y$ $\max(x-1, y-1) \geq z$ $\max(x+9, y+9) \geq u$
max-chaining x	$\max(z) \geq y$ $\max(u-11, y-1) \geq z$ $\max(u-1, y+9) \geq u$
atom-simplification-1	$\max(z) \geq y$ $\max(u-11, y-1) \geq z$ $\max(y+9) \geq u$
max-chaining y	$\max(u-11, z-1) \geq z$ $\max(z+9) \geq u$
atom-simplification-1	$\max(u-11) \geq z$ $\max(z+9) \geq u$
max-chaining z	$\max(u-2) \geq u$
atom-simplification-1	$\max() \geq u$

Theorem 5. *The Max-atom problem is in co-NP.*

Proof. By Lemma 5, if a set of max-atoms is unsatisfiable, it has a right-distinct unsatisfiable subset. This subset is a small unsatisfiability certificate, which, by the previous theorem, can be verified in polynomial time. □

Since there are few interesting problems in NP ∩ co-NP that are not known to be polynomial, this problem (in its several equivalent forms) appears to be relevant. Moreover, given the history of problems in this class, such as deciding primality [AKS04], there is hope for a polynomial-time algorithm.

5 PTIME Equivalences

In this section we show the polynomial reducibility between the Max-atom problem, the satisfiability problem for two-sided linear max-plus systems, and the existence problem of shortest hyperpaths in hypergraphs.

Theorem 6. *The Max-atom problem and the problem of satisfiability of a two-sided linear max-plus system are polynomially reducible to each other.*

Proof. Reducing this kind of max-equations to max-atoms can be done as explained in the introduction. For the reverse reduction, by the Small Model Property, if S is satisfiable then it has a model α such that size(α) $\leq K_S$ (notice that K_S can be computed in polynomial time). Let $V = \{x_1, \ldots, x_n\}$ be the set of variables over which S is defined. Now, for each variable x_i, we consider the equation

$$\begin{array}{l} \max(\ x_1 - 1, ..., x_{i-1} - 1, x_i + K_S, x_{i+1} - 1, ..., x_n - 1) = \\ \max(\ x_1, \quad ..., x_{i-1}, \quad x_i + K_S, x_{i+1}, \quad ..., x_n), \end{array}$$

which is equivalent to $x_i + K_S \geq x_j$, i.e., $K_S \geq x_j - x_i$ for all j in $1 \ldots n$, $j \neq i$. Let S'_0 be the two-sided linear max-plus system consisting of these n equations. Now we add new equations to S'_0 to obtain a system S' which is equisatisfiable to S. This is achieved by replacing every max-atom $\max(x_{i_1}, x_{i_2}) + k \geq x_{i_3}$ in S by the equation

$$\begin{array}{l} \max(\ x_{i_1} + k, x_{i_2} + k, x_{i_3}, \quad x_j - K_S - |k| - 1, ...) = \\ \max(\ x_{i_1} + k, x_{i_2} + k, x_{i_3} - 1, x_j - K_S - |k| - 1, ...), \end{array}$$

where j ranges over all variable indices different from i_1, i_2, i_3 (if any of the indices i_1, i_2 or i_3 coincide, an obvious simplification must be applied). The offset $-K_S - |k| - 1$ has been chosen so that variables with this offset do not play a role in the maxima. If we leave them out, it is clear that the resulting constraint $\max(x_{i_1} + k, x_{i_2} + k, x_{i_3}) = \max(x_{i_1} + k, x_{i_2} + k, x_{i_3} - 1)$ is equivalent to the max-atom $\max(x_{i_1}, x_{i_2}) + k \geq x_{i_3}$. □

For the relationship with shortest hyperpaths, first some preliminary notions on hypergraphs are presented. We do this by contrasting them with the analogous concepts for graphs.

A (directed, weighted) *graph* is a tuple $G = (V, E, W)$ where V is the set of *vertices*, E is the set of *edges* and $W : E \to \mathbb{Z}$ is the *weight function*. Each edge is a pair (s, t) from a vertex $s \in V$ called the *source vertex* to a vertex $t \in V$ called the *target vertex*.

A (directed, weighted) *hypergraph* is a tuple $H = (V, E, W)$ where V is the set of *vertices*, E is the set of *hyperedges* and $W : E \to \mathbb{Z}$ is the *weight function*. Each hyperedge is a pair (S, t) from a non-empty finite subset of vertices $S \subseteq V$ called the *source set* to a vertex $t \in V$ called the *target vertex*. Thus, a graph is a hypergraph in which for all hyperedges the source set consists of a single element.

Given a graph $G = (V, E, W)$ and vertices $x, y \in V$, a *path from x to y* is a sequence of edges defined recursively as follows: (i) if $y = x$, then the empty sequence $\emptyset$ is a path from x to y; (ii) if there is an edge $(z, y) \in E$ and a path $s_{x,z}$ from x to z, then the sequence $s_{x,y}$ obtained by appending (z, y) to the sequence $s_{x,z}$ is a path from x to y.

Given a hypergraph $H = (V, E, W)$, a subset of vertices $X \subseteq V$, $X \neq \emptyset$ and $y \in V$, a *hyperpath from X to y* is a tree defined recursively as follows: (i) if $y \in X$, then the empty tree $\emptyset$ is a hyperpath from X to y; (ii) if there is a hyperedge $(Z, y) \in E$ and hyperpaths t_{X,z_i} from X to z_i for each $z_i \in Z$, then the tree $t_{X,y}$ with root (Z, y) and children the trees t_{X,z_i} for each vertex $z_i \in Z$, is a hyperpath from X to y. Therefore, when viewing graphs as hypergraphs, a path is just a hyperpath where the tree has degenerated into a sequence of edges. This notion of hyperpath corresponds to the *unfolded hyperpaths* or *hyperpath trees* of [AIN92].

Using the weight function W on the edges E of a graph, one can extend the notion of weight to paths. Namely, the *weight of a path p*, denoted by $\omega(p)$, can be defined naturally as follows: (i) if p is $\emptyset$, then $\omega(p) = 0$; (ii) if p is the result of appending the edge e to the path q, then $\omega(p) = W(e) + \omega(q)$.

On the other hand, in the case of hypergraphs several notions of hyperpath weight have been studied [AIN92]. In this paper we consider the one of *rank* (also called the *distance* [GLPN93]) of a hyperpath p, which is defined as: (i) if p is $\emptyset$, then $\omega(p) = 0$; (ii) if p is a tree with root the hyperedge e and children $p_1, ..., p_m$, then $\omega(p) = W(e) + \max\big(\omega(p_1), \ldots, \omega(p_m)\big)$. This natural notion intuitively corresponds to the heaviest path in the tree.

From now on, we will assume that hypergraphs are *finite*, i.e., the set of vertices V is finite.

Example 4. Fig. 1 (a) shows an example of a hypergraph. E.g., the hyperedge $(\{u\}, x)$ has weight -10, while the weight of the hyperedge $(\{u, x\}, z)$ is 25 . The empty tree is a hyperpath from $\{u, y\}$ to y with rank 0; Fig. 1 (b) shows another hyperpath from $\{u, y\}$ to y, with rank 24.

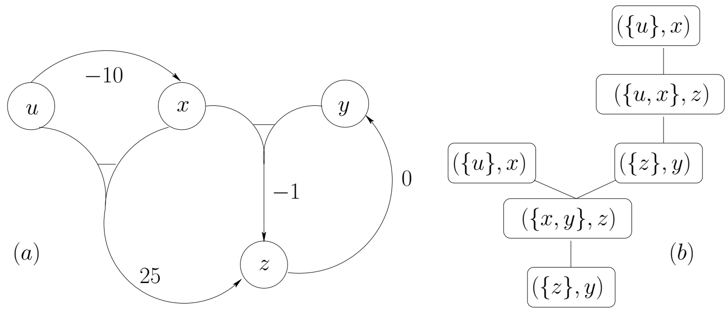

Fig. 1. Example of hypergraph.

We now show that the Max-atom problem and the problem of existence of shortest hyperpaths (i.e., with the least rank) in hypergraphs are equivalent, in the sense that they can polynomially be reduced to one another.

Definition 8. *Let $H = (V, E, W)$ be a hypergraph. Given a subset of vertices $X \subseteq V$, $X \neq \emptyset$, the* distance *function $\delta_X : V \to \mathbb{Z} \cup \{\pm\infty\}$ is defined as*

$$\delta_X(y) = \inf\{\omega(t_{X,y}) \mid t_{X,y} \text{ is a hyperpath from } X \text{ to } y\},$$

where for $S \subseteq \mathbb{R}$, we denote by $\inf(S) \in \mathbb{R} \cup \{\pm\infty\}$ the infimum of S.

The distance function δ_X is said to be well-defined *if $\delta_X(y) > -\infty$ for all vertices $y \in V$.*

With this definition, intuitively $+\infty$ means "no hyperpath" and $-\infty$ is related to negative cycles, for instance in the presence of an hyperedge such as $W(\{x\}, x) = -1$.

Our goal is to show that the satisfiability of sets of max-atoms is equivalent to the problem of, given a hypergraph $H = (V, E, W)$, decide if for all non-empty $X \subseteq V$ the distance function δ_X is well-defined, i.e., for all $y \in V$ there exists a shortest hyperpath from X to y. To that end, we need the following lemmas:

Lemma 8. *Let $H = (V, E, W)$ be a hypergraph and $X \subseteq V$, $X \neq \emptyset$ be a set of vertices such that $-\infty < \delta_X(y) < +\infty$ for all $y \in V$. If $(Z, y) \in E$, then $\delta_X(y) \leq W(Z, y) + \max_{z \in Z}(\delta_X(z))$.*

Proof. By hypothesis for all $y \in V$ we have $-\infty < \delta_X(y) < +\infty$. Thus, in particular, for all $z \in Z$ there exists a hyperpath t_z from X to z such that $\omega(t_z) = \delta_X(z)$. Now the tree t with root (Z, y) and children the trees t_z for each $z \in Z$ is a hyperpath from X to y. So $\delta_X(y) \leq \omega(t) = W(Z, y) + \max_{z \in Z}(\omega(t_z)) = W(Z, y) + \max_{z \in Z}(\delta_X(z))$. □

Lemma 9. *Let $H = (V, E, W)$ be a hypergraph and $\alpha : V \to \mathbb{Z}$ be such that $\alpha(y) \leq \max_{z \in Z}(\alpha(z)) + W(Z, y)$ for all hyperedges $(Z, y) \in E$. If t is a hyperpath from a non-empty $X \subseteq V$ to $y \in V$, then $\alpha(y) \leq \max_{x \in X}(\alpha(x)) + \omega(t)$.*

Proof. Let us prove it by induction over the depth of t. In the base case $t = \emptyset$, and therefore $y \in X$. Since $\omega(\emptyset) = 0$, trivially $\alpha(y) \leq \max_{x \in X}(\alpha(x)) = \max_{x \in X}(\alpha(x)) + \omega(\emptyset)$. Now, if t has positive depth, its root is a hyperedge $(Z, y) \in E$, and its children are trees t_1, ..., t_m connecting X to z_1, ..., z_m respectively, where $Z = \{z_1, ..., z_m\}$. By induction hypothesis, for each i in $1 \ldots m$ we have $\alpha(z_i) \leq \max_{x \in X}(\alpha(x)) + \omega(t_i)$. Now:

$$\alpha(y) \leq \max_{1 \leq i \leq n}(\alpha(z_i)) + W(Z, y) \leq \max_{1 \leq i \leq n}(\max_{x \in X}(\alpha(x)) + \omega(t_i)) + W(Z, y) =$$

$$= \max_{x \in X}(\alpha(x)) + \max_{1 \leq i \leq n}(\omega(t_i)) + W(Z, y) = \max_{x \in X}(\alpha(x)) + \omega(t). \qquad \square$$

Finally we are in condition to prove the equivalence of the two problems. For convenience, in what remains of this section we assume max-atoms to be of the form $\max_{1 \leq i \leq n}(x_i) + k \geq z$.

Theorem 7. *The Max-atom problem and the problem of well-definedness of the distance functions of all subsets of vertices of a hypergraph are polynomially reducible to each other.*

Proof. First we prove that, given a set S of max-atoms, one can compute in polynomial time a hypergraph $H(S)$ whose distance functions are well-defined if and only if S is satisfiable.

Let S be a set of max-atoms over the variables V. We can assume w.l.o.g. that there exists a variable $x \in V$ such that there are max-atoms $x \geq y \in S$ for every $y \in V$ (adding a fresh variable with these properties preserves satisfiability). The hypergraph $H(S)$ is defined as follows: its set of vertices is V; and for each max-atom $\max_{z \in Z}(z) + k \geq y$, we define a hyperedge $e = (Z, y)$ with weight $W(e) = k$.

Let us see that the distance function δ_x in $H(S)$ is well-defined if and only if S is satisfiable (we write δ_x instead of $\delta_{\{x\}}$ for the sake of clarity). Let us prove that if δ_x is well-defined then S is satisfiable. By construction, for each max-atom $\max_{z \in Z}(z) + k \geq y \in S$ there exists a hyperedge $e = (Z, y)$ in $H(S)$ with weight $W(e) = k$. Now, since δ_x is well-defined and all vertices are hyperconnected to $\{x\}$, by Lemma 8 we have $\max_{z \in Z}(\delta_x(z)) + W(Z, y) \geq \delta_x(y)$, and so $\delta_x \models S$. Let us prove the converse, i.e., that if S is satisfiable then δ_x is well-defined, by contradiction. Let us assume that δ_x is not well-defined and let α be a model of S. Then there is $y \in V$ such that $\delta_x(y) = -\infty$. This implies that for all $w \in \mathbb{R}$ there exists a hyperpath t_w from $\{x\}$ to y such that $\omega(t_w) < w$; in particular, this holds for $w = \alpha(y) - \alpha(x)$. As $\alpha \models S$, by Lemma 9 we have $\alpha(x) + \omega(t_w) \geq \alpha(y)$, i.e., $\omega(t_w) \geq \alpha(y) - \alpha(x)$, which is a contradiction.

Finally, as in $H(S)$ all vertices are hyperconnected to $\{x\}$ by a hyperedge, it is clear that δ_x is well-defined if and only if so is δ_X for all $X \subseteq V$, $X \neq \emptyset$.

Secondly, let us prove that given a hypergraph H, one can compute in polynomial time a set $S(H)$ of max-atoms such that H has a well-defined distance function δ_X for all $X \subseteq V$, $X \neq \emptyset$ if and only if $S(H)$ is satisfiable. Given $H = (V, E, W)$, the variables of $S(H)$ are V, the vertices of H; and for each hyperedge $(Z, y) \in E$, we consider the max-atom $\max_{z \in Z}(z) + W(Z, y) \geq y$. The proof concludes by observing that H has a well-defined distance function δ_X for all $X \subseteq V$, $X \neq \emptyset$ if and only if the same property holds for $H(S(H))$, if and only if $S(H)$ is satisfiable. □

Example 5. The hypergraph corresponding to the set of max-atoms considered in Example 1 is the one shown in Example 4.

6 Conclusions and Future Directions

The contributions of this paper can be summarized as follows:

- First, we have shown that the Max-atom problem is in NP ∩ co-NP. As no PTIME algorithm for solving this problem has been found yet, this is relevant since there are few interesting problems in NP ∩ co-NP that are not known to be polynomial.

- We have given a *weakly* polynomial decision procedure for the problem (when the offsets are integers). This algorithm becomes polynomial under more restrictive conditions on the input, e.g. by imposing a bound on offsets.
- Finally, we have shown the equivalence of deciding the Max-atom problem with two other at first sight unrelated problems: namely, (i) the satisfiability of two-sided linear max-plus systems of equations, used in Control Theory for modeling Discrete Event Systems; and (ii) the existence for a given hypergraph of shortest paths from any non-empty subset of vertices to any vertex. Finding a PTIME algorithm for these problems has been open in the respective areas for more than 30 years, and is still unsolved.

As regards future work, in the short term we would like to find a weakly polynomial algorithm when the offsets may be arbitrary rational numbers. This would perhaps give new insights about the long-term goal of finding a polynomial algorithm for deciding the satisfiability of sets of max-atoms.

As noticed by an anonymous referee, the Max-atom problem is a special case of the problem of finding a super-fixed point of a min-max function. A super-fixed point of a function f on a (partially) ordered set A is an $a \in A$ such that $f(a) \geq a$. Now, for instance, the satisfiability of S in Example 1 from Section 3 is equivalent to finding a super-fixed point of

$$f(u, x, y, z) = (u, u - 10, z, \min(\max(x, y) - 1, \max(x, u) + 25))$$

with respect to the coordinate-wise partial order. More information on min-max functions can be found in [G94]. The referee further mentioned a connection with game theory in [C92]. We gratefully acknowledge these suggestions for future research.

References

[AIN92] G. Ausiello, G. Italiano, and U. Nanni. Optimal traversal of directed hypergraphs. Tech. Rep. TR-92-073, ICSI, Berkeley, CA, 1992.

[AKS04] M. Agrawal, N. Kayal, and N. Saxena. PRIMES is in P. *Annals of Mathematics*, 160(2):781–793, 2004.

[BNRC08] M. Bezem, R. Nieuwenhuis, and E. Rodriguez-Carbonell. Exponential behaviour of the Butkovič-Zimmermann algorithm for solving two-sided linear systems in max-algebra. To appear in *Discrete Applied Mathematics*.

[BZ06] P. Butkovič and K. Zimmermann. A strongly polynomial algorithm for solving two-sided linear systems in max-algebra. *Discrete Applied Mathematics*, 154(3):437–446, 2006.

[C92] A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.

[GLPN93] G. Gallo, G. Longo, S. Pallottino, and S. Nguyen. Directed hypergraphs and applications. *Discrete Applied Mathematics*, 42:177–201, 1993.

[G94] J. Gunawardena. Min-max functions. *Discrete Events Dynamic Systems*, 4:377–406, 1994.

[NOT06] R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT Modulo Theories: from an Abstract Davis-Putnam-Logemann-Loveland Procedure to DPLL(T). *Journal of the ACM*, 53(6):937–977, 2006.