Complexity of Reasoning with Cardinality Minimality Conditions

Nadia Creignou[1], **Frédéric Olive**[1], **Johannes Schmidt**[2]

[1] Aix Marseille Univ, CNRS, LIS, Marseille, France
[2] Jönköping University, Department of Computer Science and Informatics, School of Engineering, Sweden
{nadia.creignou, frederic.olive}@lis-lab.fr, johannes.schmidt@ju.se

Abstract

Many AI-related reasoning problems are based on the problem of satisfiability of propositional formulas with some cardinality-minimality condition. While the complexity of the satisfiability problem (SAT) is well understood when considering systematically all fragments of propositional logic within Schaefer's framework, this is not the case when such minimality condition is added. We consider the CARDMINSAT problem, which asks, given a formula ϕ and an atom x, whether x is true in some cardinality-minimal model of ϕ. We completely classify the computational complexity of the CARDMINSAT problem within Schaefer's framework, thus paving the way for a better understanding of the tractability frontier of many AI-related reasoning problems. To this end we use advanced algebraic tools.

Introduction

In many AI-related reasoning problems some notion of minimality is involved. Typically in belief change, e.g. revision or update, one of the basic principles is the principle of minimal change. We want to revise/update an agent's belief set by some new information. To this end we retain only those models of new information that have minimal distance to the models of the original agent's belief set. In the belief revision context distance between models is defined by the symmetric set difference of the atoms assigned to true in the compared models, and Dalal's operator (Dalal 1988), for instance, seeks to minimize the cardinality of this set. In abduction we search for an explanation (a set of literals) that is consistent with a given theory and which, together with this theory, logically entails all manifestations. It is natural to be interested not in all explanations but only in the minimal ones. Different notions of minimality might be considered, in particular minimality w.r.t set inclusion or w.r.t. cardinality (Eiter and Gottlob 1995).

In this paper, we focus on cardinality-minimality. With such a minimality condition the related reasoning tasks often give rise to Θ_2^{P}-complete problems (the class Θ_2^{P} is located at the second level of the polynomial hierarchy: polynomial time with only a logarithmic number of calls to the NP-oracle). For instance, model checking and implication are Θ_2^{P}-complete for Dalal's revision operator (Eiter and Gottlob 1992; Liberatore and Schaerf 2001). The relevance problem for abduction with a cardinality-minimality condition, deciding whether a literal belongs to a cardinality-minimal explanation, is Θ_2^{P}-complete when dealing with Horn formulas (Eiter and Gottlob 1995).

Propositional formulas play an important role in AI-reasoning problems. Since most relevant problems are intractable in full propositional logic, it is a natural question whether syntactic restrictions on the involved formulas can lead to tractable problems. Schaefer's framework offers an ideal framework to investigate this issue. It considers formulas in generalized conjunctive normal form and allows to systematically consider all fragments of propositional logic. Indeed, Schaefer's famous theorem (Schaefer 1978) shows that the SAT problem becomes tractable under some syntactic restrictions such as Horn, dual Horn, Krom or affine formulas, and remains intractable in all other, nontrivial, cases. Since then Schaefer's approach has been taken on numerous problems, among others on circumscription, abduction, and argumentation problems (Nordh 2004; Creignou and Zanuttini 2006; Nordh and Zanuttini 2008; Creignou, Egly, and Schmidt 2014). Tools from universal algebra prove to be a valuable tool for such endeavors, in particular when the problem questions are stable under introduction of existentially quantified variables and equality constraints (Creignou and Vollmer 2008). Unfortunately, cardinality is not preserved under such introduction. Therefore, in this paper, we resort to advanced algebraic tools built around the concept of a *weak base* (Schnoor and Schnoor 2008; Lagerkvist 2014).

There is a prototypical satisfiability problem for the class Θ_2^{P}, that could enlighten the complexity of many reasoning problems involving a cardinality-minimality condition: It is the CARDMINSAT problem, which asks, given a formula ϕ and an atom x, whether x is true in some cardinality-minimal model of ϕ. It provides a standard hard problem that can be useful to prove hardness results, especially in the context of knowledge representation and belief change. For instance, in (Creignou, Pichler, and Woltran 2018) the relevance problem for abduction mentioned above was proved to be Θ_2^{P}-complete for the combined Horn-Krom case. The Θ_2^{P}-hardness reduction used in (Creignou, Pichler, and Woltran 2018) is much easier than the one previously obtained in (Eiter and Gottlob 1995) for the Horn case, because it starts from the more closely related problem CARD-MINSAT, restricted to conjunctions of positive 2-clauses. Similarly, the model checking and implication problems associated with Dalal's operator were proved in (Creignou, Pichler, and Woltran 2018) to be Θ_2^{P}-complete for the combined Horn-Krom case by a reduction from the CARDMINSAT problem. Also the model checking problem associated with a syntactic revision operator for belief bases using a cardinality-maximality criterion was proved to be Θ_2^{P}-complete for the combined Horn-Krom case in this way (Creignou, Ktari, and Papini 2017).

Our main contribution is a complete complexity classification of the CARDMINSAT problem in Schaefer's framework, which opens the door for a better understanding of the complexity of many reasoning problems. As an illustration we prove that the above mentioned relevance problem for abduction remains Θ_2^{P}-complete when restricted to affine formulas (conjunctions of XOR-clauses). Due to space limitations, for results marked with a * the full proof can be found in the technical report of the paper (Creignou, Olive, and Schmidt 2023).

Preliminaries

Propositional Logic. We assume familiarity with propositional logic. A *literal* is a variable (or an atom) x (positive literal) or its negation $\neg x$ (negative literal). A *clause* is a disjunction of literals. For any integer $k \geq 1$, a k-clause is a clause containing at most k literals. An *XOR-clause* is a clause in which the usual connective "or" is replaced by the exclusive-or connective, denoted by $\oplus$. A CNF-formula (resp., an XOR-CNF-formula) is a conjunction of clauses (resp., XOR-clauses), a k-CNF-formula is a conjunction of k-clauses. For space economy we use occasionally the shorthands $\overline{x} := \neg x$ and $xy := x \wedge y$. Given a formula ϕ, we denote by $\mathrm{var}(\phi)$ the variables of ϕ. A mapping $\sigma\colon\ \mathrm{var}(\phi) \mapsto \{0,1\}$ is called an *assignment* to the variables of ϕ. An assignment σ satisfies a (XOR-)CNF-formula ϕ if σ satisfies all (XOR-)clauses simultaneously. In this case σ is called a *model* of ϕ. We call a variable $x \in \mathrm{var}(\phi)$ *frozen* if x is assigned the same value in all models of ϕ. The *weight* or *cardinality* of an assignment σ, denoted by $|\sigma|$, is the number of variables x such that $\sigma(x) = 1$. A *cardinality minimal* model of ϕ is a model of ϕ of minimum cardinality among all models of ϕ. For two formulas ψ, ϕ we write $\psi \models \phi$ if every model of ψ also satisfies ϕ. The two formulas are equivalent, $\psi \equiv \phi$, if they have the same set of variables and the same set of models. Observe that any XOR-clause is equivalent to a linear equation over the two-elements field, of the form $x_1 \oplus \ldots \oplus x_n = a$ where $a \in \{0,1\}$.

Schaefer's Framework. A *Boolean relation* of arity $k \in \mathbb{N}$ is a relation $R \subseteq \{0,1\}^k$, and a *constraint* C is a formula $C = R(x_1,\ldots,x_k)$, where R is a k-ary Boolean relation, and $x_1,\ldots,x_k$ are (not necessarily distinct) variables. An assignment σ *satisfies* C, if $(\sigma(x_1),\ldots,\sigma(x_k)) \in R$. A *constraint language* Γ is a finite set of Boolean relations, and a Γ-*formula* is a conjunction of constraints using relations from Γ. Note that we do not consider infinite constraint languages in this paper. Finally, a Γ-formula ϕ is *satisfied* by an assignment σ, if σ simultaneously satisfies all constraints in it. In such a case σ is also called a *model of* ϕ. We say that a k-ary relation R is *defined by* a formula ϕ if ϕ is a formula over k distinct variables $x_1,\ldots,x_k$ and $\phi \equiv R(x_1,\ldots,x_k)$.

Moreover, we say that a Boolean relation R is:

- *Horn* (resp., *dual-Horn*) if it is definable by a CNF-formula ϕ that contains at most one positive (resp., negative) literal per each clause,
- *Krom* if it is definable by a 2-CNF-formula,
- *affine* it is definable by an XOR-CNF formula, or equivalently by a formula ϕ that is a conjunction of linear equations of the form $x_1 \oplus \ldots \oplus x_n = a$, where $a \in \{0,1\}$,
- *width-2-affine* it is definable by an XOR-2-CNF formula, or equivalently by a formula ϕ that is a conjunction of linear equations involving each at most two variables, that is either of the form $x_1 = a$ or of the form $x_1 \oplus x_2 = a$, where $a \in \{0,1\}$.
- 1-*valid* (resp., 0-*valid*) if $(1,\ldots,1) \in R$ (resp., $(0,\ldots,0) \in R$).
- *complementive* if for every tuple $(t_1,\ldots,t_k) \in R$ also $(1-t_1,\ldots,1-t_k) \in R$.

Furthermore, we say a relation is *Schaefer* if it is Horn, dual-Horn, Krom, or affine. Finally, for a property $\mathcal{P}$ of a relation, we say that a constraint language Γ is $\mathcal{P}$ if all relations in Γ are $\mathcal{P}$.

We define the unary relations $\mathrm{T} = \{1\}$, $\mathrm{F} = \{0\}$, and the 6-ary relation $R^{1/3}_{3\neq} = \{100011, 010101, 001110\}$. We denote by OR^k the k-ary OR, by NAND^k the k-ary NAND, and by XOR^k the k-ary XOR. The relation EVEN^k contains all k-ary tuples which contain an even number of 1's. The relation $\mathrm{EVEN}^k_{k\neq}$ denotes the $2k$-ary relation defined by $\mathrm{EVEN}^k(x_1,\ldots,x_k) \wedge (x_1 \neq x_{k+1}) \wedge \cdots \wedge (x_k \neq x_{2k})$.

In the following definition we introduce different notions of closure for a constraint language.

Definition 1. *1. The set $\langle\Gamma\rangle$ is the smallest set of relations that contains Γ, the equality constraint, $=$, and which is closed under primitive positive first order definitions, that is, if ϕ is a $\Gamma \cup \{=\}$-formula and $R(x_1,\ldots,x_n) \equiv \exists y_1 \ldots \exists y_l \phi(x_1,\ldots,x_n,y_1,\ldots,y_l)$, then $R \in \langle\Gamma\rangle$. In other words, $\langle\Gamma\rangle$ is the set of relations that can be expressed as a $\Gamma \cup \{=\}$-formula with existentially quantified variables.*

2. The set $\langle\Gamma\rangle_{\nexists}$ is the set of relations that can be expressed as a $\Gamma \cup \{=\}$-formula (no existentially quantified variables are allowed).

3. The set $\langle\Gamma\rangle_{\nexists,\neq}$ is the set of relations that can be expressed as a Γ-formula (neither the equality relation nor existentially quantified variables are allowed).

Example 1. *Let $\Gamma = \{R\}$, $R(x_1,x_2) = (x_1 \rightarrow x_2)$, and $S(x_1,x_2) = (x_1 = x_2)$. We can express S as Γ-formula via $S(x,y) \equiv R(x,y) \wedge R(y,x)$. Thus, $S \in \langle\Gamma\rangle_{\nexists,\neq}$.*

The set $\langle\Gamma\rangle$ is called a *relational clone* or a *co-clone* with *base* Γ (Böhler et al. 2005). Notice that for a co-clone C and a constraint language Γ the statements $\Gamma \subseteq \mathsf{C}$, $\langle\Gamma\rangle \subseteq \mathsf{C}$,

$\langle\Gamma\rangle_{\nexists} \subseteq \mathsf{C}$, and $\langle\Gamma\rangle_{\nexists,\neq} \subseteq \mathsf{C}$ are equivalent. Throughout the paper, we refer to different types of Boolean relations and corresponding co-clones following Schaefer's terminology (Schaefer 1978). Important co-clones and weak bases for this paper are given in table 1. For a complete list we refer to (Lagerkvist 2014) or to the technical report (Creignou, Olive, and Schmidt 2023). For clause type descriptions and simpler bases (not minimal weak bases) we refer to (Nordh and Zanuttini 2008) and (Böhler et al. 2005), respectively.

A graph representation of the co-clone structure can be found in figure 1. This graph is usually called Post's lattice (Post 1941). Some important properties/names are labeled besides the respective co-clone. Informally explained, every vertex corresponds to a co-clone while the edges model the containment relation in this lattice structure.

Complexity Classes. All complexity results in this paper refer to classes in the Polynomial Hierarchy (PH) (Papadimitriou 1994). The building blocks of PH are the classes P and NP of decision problems solvable in deterministic, resp. non-deterministic, polynomial time. The class Δ_2^{P} is the class of decision problems that can be decided by a deterministic Turing machine in polynomial time using an oracle for the class NP. One can put restrictions on the number of oracle calls. If on input x with $|x| = n$ at most $O(\log n)$ calls to the NP oracles are allowed, then we get the class $\mathrm{P}^{\mathrm{NP}[O(\log n)]}$, which is also referred to as Θ_2^{P}. A large collection of Θ_2^{P}-complete problems can be obtained from (Krentel 1988; Gasarch, Krentel, and Rappoport 1995). For the reductions we employ polynomial many-one reductions, denoted by $\leq_{\mathrm{m}}^{\mathrm{P}}$.

CardMinSat

We aim at studying the following natural variant of SAT and analyzing its complexity.

CardMinSat
Instance : A propositional formula ϕ and an atom x.
Question : Is x true in a cardinality-minimal model of ϕ?

This problem is one of the prototypical problems of the class Θ_2^{P}, see (Wagner 1988; Creignou, Pichler, and Woltran 2018). It makes sense to study whether syntactic restrictions on the formulas make the problem easier and to go through a more fine-grained complexity study of CardMinSat, in the following also denoted Cms. To this aim we propose to investigate this problem within Schaefer's framework. Hence we consider the following problem, in which Γ is a constraint language, i.e., a finite set of Boolean relations.

Cms(Γ)
Instance : A Γ-formula ϕ and an atom x.
Question : Is x true in a cardinality-minimal model of ϕ?

Analogously we denote by Sat(Γ) the Boolean satisfiability problem for Γ-formulas. Our goal is to obtain a complete complexity classification of Cms(Γ), depending on Γ. This issue has already been settled in the literature within the Krom fragment.

Theorem 1. (Creignou, Pichler, and Woltran 2018) *Let Γ be a Krom constraint language. If Γ is width-2 affine or Horn, then* Cms(Γ) *is decidable in polynomial time. Otherwise it is Θ_2^{P}-complete.*

We extend this result and obtain a complete complexity classification in all fragments of propositional logic.

Theorem 2. *Let Γ be a constraint language. If Γ is width-2 affine or Horn or 0-valid, then* Cms(Γ) *is decidable in polynomial time. Otherwise it is Θ_2^{P}-complete.*

Note that Cms(Γ) is trivial for 0-valid formulas (the answer is always "no"). The complexity classification of Cms in the Krom fragment had been obtained by means of partial frozen co-clones. While these partial frozen co-clones are well described within the Krom fragment (Nordh and Zanuttini 2009), they are only partially known in the full range of propositional logic. For this reason in order to get the complete classification we use another set of tools. In particular we will use a restricted notion of closure, and build on the notion of weak bases introduced in (Schnoor and Schnoor 2008). This is described in the next section.

Technical Tools

Proof's Method

The above introduced closure operator $\langle\cdot\rangle$ on sets of Boolean relations is relevant in order to obtain complexity results for the satisfiability problem. Indeed, assume that $\Gamma_1 \subseteq \langle\Gamma_2\rangle$. Then a Γ_1-formula can be transformed into a satisfiability-equivalent Γ_2-formula, thus showing that Sat(Γ_1) can be reduced in polynomial time to Sat(Γ_2) (Jeavons 1998). Hence, the complexity of Sat(Γ) depends only on the co-clone $\langle\Gamma\rangle$. Accordingly, in order to obtain a full complexity classification for the satisfiability problem one only has to study the co-clones.

Unfortunately, since we are here interested in cardinality-minimal models, we cannot a priori only study the co-clones. Indeed, existential variables and equality constraints that may occur when transforming a Γ_1-formula into a satisfiability-equivalent Γ_2-formula are problematic, as they can change the set of models and the cardinality of each model. Therefore, we will use a more restricted notion of closure, namely the above introduced closure operator $\langle\cdot\rangle_{\nexists,\neq}$. This operator avoids existential quantifiers and equality constraints. The only operation to express relations in $\langle\Gamma\rangle_{\nexists,\neq}$ is conjunction of Γ-constraints (see e.g. example 1). Consequently, when replacing in a reduction a relation $R \in \langle\Gamma\rangle_{\nexists,\neq}$ by its representing Γ-formula, R is represented *exactly*: no new variables are introduced and no constraints other than those built on Γ are allowed (in particular no equality constraints). Therefore, any reduction based on the closure operator $\langle\cdot\rangle_{\nexists,\neq}$ preserves exactly the set of models, and, a fortiori, all cardinality-minimal models. Hence, we obtain the following property.

Proposition 3. *Let Γ be a constraint language and R be a relation.*
If $R \in \langle\Gamma\rangle_{\nexists,\neq}$ then Cms(R) $\leq_m^P$ Cms(Γ).

The proof of our complete classification will consist in a systematic exploration of the co-clones lattice, yet reduc-

Co-clone	Minimal weak base	Name/Indication
II_2	$R^{1/3}_{3\neq}(x_1,\ldots,x_6)\wedge \mathrm{F}(x_7)\wedge \mathrm{T}(x_8)$	all Boolean relations
II_1	$(x_1\vee x_2)\wedge(x_1x_2\leftrightarrow x_3)\wedge \mathrm{T}(x_4)$	1-valid
IN_2	$\mathrm{EVEN}^4_{4\neq}(x_1,\ldots,x_8)\wedge x_1x_4\leftrightarrow x_2x_3$	complementive
IL_1	$\mathrm{XOR}^3(x_1,x_2,x_3)\wedge \mathrm{T}(x_4)$	affine and 1-valid
IL_2	$\mathrm{EVEN}^3_{3\neq}(x_1,\ldots,x_6)\wedge \mathrm{F}(x_7)\wedge \mathrm{T}(x_8)$	affine
IL_3	$\mathrm{EVEN}^4_{4\neq}(x_1,\ldots,x_8)$	-
IV_1	$(\overline{x_1}\leftrightarrow \overline{x_2}\overline{x_3})\wedge(\overline{x_2}\vee\overline{x_3}\rightarrow\overline{x_4})\wedge \mathrm{T}(x_5)$	dualHorn and 1-valid
IV_2	$(\overline{x_1}\leftrightarrow \overline{x_2}\overline{x_3})\wedge \mathrm{F}(x_4)\wedge \mathrm{T}(x_5)$	dualHorn
$\mathsf{IS}^k_0, k\geq 2$	$\mathrm{OR}^k(x_1,\ldots,x_k)\wedge \mathrm{T}(x_{k+1})$	positive of width k
$\mathsf{IS}^k_{02}, k\geq 2$	$\mathrm{OR}^k(x_1,\ldots,x_k)\wedge \mathrm{F}(x_{k+1})\wedge \mathrm{T}(x_{k+2})$	essentially positive of width k
$\mathsf{IS}^k_{01}, k\geq 2$	$\mathrm{OR}^k(x_1,\ldots,x_k)\wedge(x_{k+1}\rightarrow x_1\cdots x_k)\wedge \mathrm{T}(x_{k+2})$	-
$\mathsf{IS}^k_{00}, k\geq 2$	$\mathrm{OR}^k(x_1,\ldots,x_k)\wedge(x_{k+1}\rightarrow x_1\cdots x_k)\wedge \mathrm{F}(x_{k+2})\wedge \mathrm{T}(x_{k+3})$	IHS-B+ of width k

Table 1: Important co-clones with minimal weak bases from Lagerkvist (2014).

tions can only be obtained via the restrictive operator $\langle\cdot\rangle_{\nexists,\neq}$, not via the more expressive, co-clone generating, operator $\langle\cdot\rangle$. In this context, the concept of a *weak base* is important (Schnoor and Schnoor 2008). A weak base B for a co-clone C has the property that (1) $\langle B\rangle = \mathsf{C}$, and (2) $B \in \langle\Gamma\rangle_{\nexists}$ for any Γ such that $\langle\Gamma\rangle = \mathsf{C}$. The existence of a weak base for each co-clone has been shown by Schnoor and Schnoor (2008). For a finitely generated co-clone C there even exists a single relation weak base. If such a weak base B is in addition *irredundant* (that is, the matrix representation does not contain redundant columns), it holds even that $B \in \langle\Gamma\rangle_{\nexists,\neq}$ for any Γ such that $\langle\Gamma\rangle = \mathsf{C}$. Lagerkvist (2014) has identified *minimal* weak bases for all finitely generated co-clones. A relation R is minimal, if (1) R is irredundant, (2) R contains no fictitious coordinates, (3) there is no $R' \subsetneq R$, such that $\langle R\rangle = \langle R'\rangle$. A coordinate i is called *fictitious* if its value has no influence on the membership of a tuple, that is, $(x_1,\ldots,x_{i-1},0,x_{i+1},\ldots,x_k) \in R$ if and only if $(x_1,\ldots,x_{i-1},1,x_{i+1},\ldots,x_k) \in R$. Table 1 contains important co-clones with minimal weak bases, a complete list is found in (Lagerkvist 2014), or in the technical report.

The proof method to obtain our complete classification will use the minimal weak bases as follows. In order to show a hardness result for all constraint languages generating a certain co-clone C, we pick a minimal weak base B of C and show that $\textsc{Cms}(B)$ is hard. This implies then hardness of $\textsc{Cms}(\Gamma)$ for any Γ such that $\langle\Gamma\rangle = \mathsf{C}$ by applying Proposition 3 (because B is a minimal weak base, it is irredundant, and we hence have that $B \in \langle\Gamma\rangle_{\nexists,\neq}$). We state this in the following proposition.

Proposition 4. *Let C be a co-clone and B be a minimal weak base of C. Then it holds that $\textsc{Cms}(B) \leq^P_m \textsc{Cms}(\Gamma)$ for any Γ such that $\langle\Gamma\rangle = \mathsf{C}$.*

To start with, we need hardness results for some specific relations, and this is the aim of the next section.

Specific Hardness Results

We give here some hardness results for some specific relations, they will be used in order to get hardness results for co-clones in the next section. The classification obtained in Theorem 1 for the Krom fragments implies the following result.

Lemma 5. $\textsc{Cms}(\mathrm{OR}^2)$ *is Θ^{P}_2-hard.*

The next result will also be a cornerstone in our classification proof.

Lemma 6 (*). $\textsc{Cms}(\mathrm{XOR}^3)$ *is Θ^{P}_2-hard.*

Proof sketch. Recall that $\mathrm{XOR}^3(x,y,z) \equiv (x\oplus y\oplus z)$ and $\mathrm{XOR}^4(x,y,z,u) \equiv (x\oplus y\oplus z\oplus u)$. Here we will also use the ternary relation $\mathrm{NAE}^3 = \{0,1\}^3 \setminus \{000,111\}$ and the problem $\textsc{Cms}^*(\Gamma)$, defined as follows:

$\textsc{Cms}^*(\Gamma)$
Instance : A Γ-formula ϕ, atom x, integer k.
Question : Is x true in a cardinality-minimal model of ϕ and is this cardinality $\leq k$?

The proof consists of the following sequence of reductions.

1. $\textsc{Cms}(\mathrm{OR}^2) \leq^{\mathrm{P}}_{\mathrm{m}} \textsc{Cms}(\mathrm{NAE}^3)$
2. $\textsc{Cms}(\mathrm{NAE}^3) \leq^{\mathrm{P}}_{\mathrm{m}} \textsc{Cms}^*(\mathrm{XOR}^3)$
3. $\textsc{Cms}^*(\mathrm{XOR}^3) \leq^{\mathrm{P}}_{\mathrm{m}} \textsc{Cms}(\mathrm{XOR}^4)$
4. $\textsc{Cms}(\mathrm{XOR}^4) \leq^{\mathrm{P}}_{\mathrm{m}} \textsc{Cms}(\mathrm{XOR}^3, \mathrm{XOR}^2)$
5. $\textsc{Cms}(\mathrm{XOR}^3, \mathrm{XOR}^2) \leq^{\mathrm{P}}_{\mathrm{m}} \textsc{Cms}(\mathrm{XOR}^3)$

Then the result follows from Lemma 5.

We now give the reductions 1., 4., and 5. Reductions 2. and 3. can be found in the technical report.

1. $\textsc{Cms}(\mathrm{OR}^2) \leq^{\mathrm{P}}_{\mathrm{m}} \textsc{Cms}(\mathrm{NAE}^3)$.

To each constraint $\mathrm{OR}^2(x,y)$ we associate the constraint $\mathrm{NAE}^3(x,y,f)$ where f is a fresh variable. Observe that $\mathrm{OR}^2(x,y) \equiv \mathrm{NAE}^3(x,y,0)$. Therefore the idea is to use f in the place of 0 as a global variable (that is the same for all constraints) and to force it to take value 0 in all cardinality-minimal models. This can be done by giving a weight N to f big enough. For this we add the constraint $\mathrm{NAE}^3(f,f,t)$, which expresses $f\neq t$, and N constraints $\mathrm{NAE}^3(f_j,f_j,t)$, where the f_j, for $j=1,\ldots,N$, are fresh variables. This ensures that if $f=1$ then $f_1,\ldots,f_N=1$. Observe moreover that since NAE^3 is a complementive relation, the built NAE^3-formula is satisfiable if and only if it has a model with $f=0$. Taking $N>n$ where n is the number of

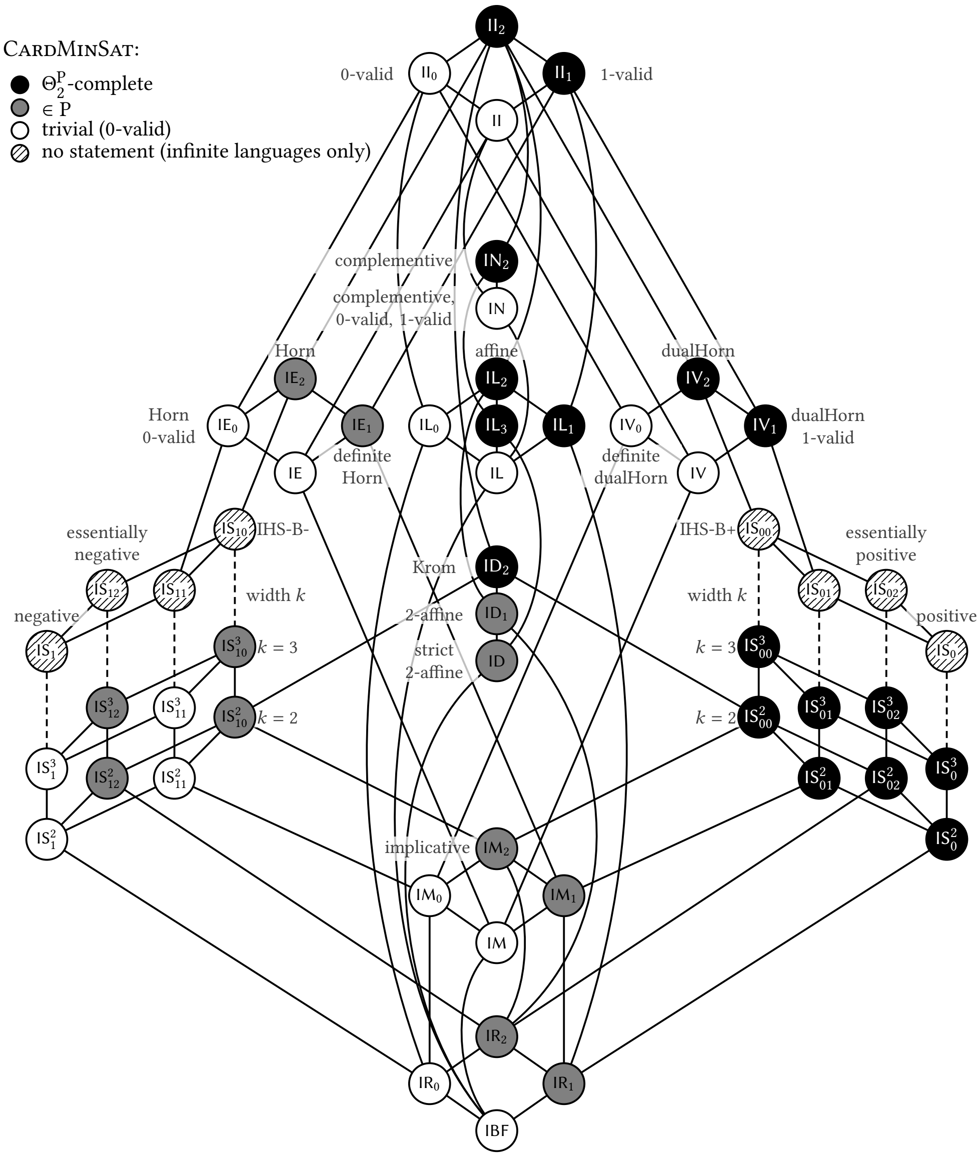

Figure 1: Complexity overview for CardMinSat illustrated on Post's Lattice.

variables of the original formula ensures that $f = 0$ in any cardinality-minimal model.

4. $\textsc{Cms}(\mathrm{XOR}^4) \leq_m^P \textsc{Cms}(\mathrm{XOR}^3,\, \mathrm{XOR}^2)$.

Observe that $\mathrm{XOR}^4(x_1, x_2, x_3, x_4) \equiv \exists y, z : \mathrm{XOR}^3(x_1, x_2, y) \wedge \mathrm{XOR}^3(x_3, x_4, z) \wedge \mathrm{XOR}^2(y, z)$. The two fresh variables y and z take complementary values, so they will together contribute a weight 1 in any case.

5. $\textsc{Cms}(\mathrm{XOR}^3,\, \mathrm{XOR}^2) \leq_m^P \textsc{Cms}(\mathrm{XOR}^3)$.

Let (ϕ, x) be an instance of $\textsc{Cms}(\mathrm{XOR}^3,\, \mathrm{XOR}^2)$. If ϕ is unsatisfiable, we map (ϕ, x) to a trivial negative instance of $\textsc{Cms}(\mathrm{XOR}^3)$, e.g. $(\mathrm{XOR}^3(x_1, x_2, x_3), x)$.

Otherwise, we replace any constraint $\mathrm{XOR}^2(x, y)$ by $\mathrm{XOR}^3(x, y, w)$ where w is a fresh variable of weight impact N big enough, say bigger than the number of variables of the original formula. This assures that the cardinality-minimal models of the formula are the models of ϕ extended with $w = 0$. The variable w can be given the needed weight impact by adding the constraints $\mathrm{XOR}^3(t, t, t) \wedge \bigwedge_{i=1}^{N} \mathrm{XOR}^3(t, w, w_i)$ where t and the w_i's are fresh variables. □

Proof of Main Theorem

We prove here Theorem 2. The classification can be visualized on Post's Lattice, see figure 1. The classification obeys the borders among co-clones and, as discussed in the previous section, will be obtained by a systematic exploration of the co-clones.

Observe that Theorem 1, the previously obtained classification in the Krom fragment, concerns co-clones in the lower part of the lattice, namely every co-clone C such that $\mathsf{C} \subseteq \mathsf{ID}_2$. In the depiction of Post's Lattice in figure 1 the color coding is as follows. The "white" co-clones, for which the problem CMS is trivial, are the co-clones that contain only 0-valid relations. For those ones the cardinality-minimum solution is the all-0 solution, and the answer is always "no". The "grey" co-clones, for which the problem CMS is decidable in polynomial time, correspond to co-clones C such that either $\mathsf{C} \subseteq \mathsf{IE}_2$ or $\mathsf{C} \subseteq \mathsf{ID}_1$. In the first case, all relations are Horn, and therefore there exists a unique cardinality-minimal model that can be found by unit propagation in polynomial time. In the second case, all relations are width-2-affine and the tractability result follows from Theorem 1.

Finally, to obtain the complexity classification it remains to prove hardness for the "black" co-clones, namely II_2, II_1, IN_2, IL_2, IL_3, IL_1, IV_2, IV_1, and, for any $k \geq 2$, for the co-clones IS_{00}^{k}, IS_{01}^{k}, IS_{02}^{k}, IS_{0}^{k}. The "black" co-clone ID_2 is dealt with by Theorem 1.

As we have discussed in the previous section, for each remaining co-clone C, given one of its weak bases B we will show that $\mathrm{CMS}(B)$ is hard. This will be done by a reduction from a known hard problem, either $\mathrm{CMS}(\mathrm{OR}^2)$ or $\mathrm{CMS}(\mathrm{XOR}^3)$. For example, given an instance (ϕ, x) of $\mathrm{CMS}(\mathrm{OR}^2)$, where ϕ is a conjunction of OR^2-clauses, we will build a B-formula ϕ' such that x belongs to a cardinality-minimal model of ϕ if and only if x belongs to a cardinality-minimal model of ϕ'. The construction of ϕ' is obtained by a local replacement of each clause of ϕ by an equivalent B-formula. Usually this requires introduction of fresh (existentially quantified) variables. Some of these additional variables will be *frozen*, which means that their truth value is the same in all models, and thus their contribution to the weight of any model is fixed. In order to be sure that the weight of the non-frozen additional variables will not compromise the cardinality-minimal models, the trick is to neutralize them by adding for each such variable y, another one y' and to force them to take complementary values, i.e. $y \neq y'$. Thus the weight contribution of y and y' together will always be 1 in all models. Sometimes, to do so we will have to express the truth value 0. When this is not possible directly, the idea is to replace 0 by a variable f, and then introduce a big number of copies of f such that any cardinality-minimal model of the formula has to set f to 0.

In the following when we speak about *the* minimal weak base of a co-clone we mean the weak base given in (Lagerkvist 2014) (cf. also table 1 and the technical report). In the proofs, we will always restate the exact definition of the corresponding weak base, and, where convenient, also its matrix representation.

The following proposition provides the missing hardness results.

Proposition 7 (*). *Let Γ be a constraint language. Then $\mathrm{CMS}(\Gamma)$ is Θ_2^{P}-hard if $\langle\Gamma\rangle \in \{\mathsf{II}_2, \mathsf{II}_1, \mathsf{IN}_2, \mathsf{IL}_2, \mathsf{IL}_3, \mathsf{IL}_1, \mathsf{IV}_2, \mathsf{IV}_1, \mathsf{IS}_{00}^{k}, \mathsf{IS}_{01}^{k}, \mathsf{IS}_{02}^{k}, \mathsf{IS}_{0}^{k}\}$, for any $k \geq 2$.*

The first five cases are proven in the following lemmas. Proofs for the remaining cases can be found in the technical report.

Lemma 8. *Let $\langle\Gamma\rangle = \mathsf{II}_2$. Then $\mathrm{CMS}(\Gamma)$ is Θ_2^{P}-hard.*

Proof. Let R_{II_2} be the minimal weak base of II_2, that is, $R_{3\neq}^{1/3}(x_1, \ldots, x_6) \wedge \mathrm{F}(x_7) \wedge \mathrm{T}(x_8)$, where $R_{3\neq}^{1/3} = \{100011, 010101, 001110\}$.

The matrix representation is as follows.

$$R_{\mathsf{II}_2} = \begin{pmatrix} 10001101 \\ 01010101 \\ 00111001 \end{pmatrix}$$

We show that $\mathrm{CMS}(\mathrm{OR}^2) \leq_{\mathrm{m}}^{\mathrm{P}} \mathrm{CMS}(R_{\mathsf{II}_2})$. Then the result follows from Lemma 5 and Proposition 4.

Let (ϕ, x) be an instance of $\mathrm{CMS}(\mathrm{OR}^2)$, where $\phi = \bigwedge_{i=1}^{p}(x_i^1 \vee x_i^2)$. Let $\{a_i, b_i, c_i, d_i, a_i', b_i', c_i', d_i' \mid i = 1 \ldots p\} \cup \{t, f\}$ be fresh variables. For each constraint $(x_i^1 \vee x_i^2)$ we build the constraint $R_{\mathsf{II}_2}(a_i, b_i, c_i, d_i, x_i^1, x_i^2, f, t)$. Observe that $\mathrm{OR}^2(x_i^1, x_i^2) \equiv$

$$\exists a_i, b_i, c_i, d_i, f, t\ R_{\mathsf{II}_2}(a_i, b_i, c_i, d_i, x_i^1, x_i^2, f, t).$$

The existential variables are uniquely determined. The variables f and t are frozen, while the values of a_i, b_i, c_i, d_i are not. Nevertheless their values can be neutralized by the introduction of additional fresh variables a_i', b_i', c_i', d_i' who are forced to take complementary values. In the case of a_i and a_i' this can be achieved by the constraint $R_{\mathsf{II}_2}(a_i, a_i', f, a_i', a_i, t, f, t)$. Analogous constraints are added for b_i, b_i', c_i, c_i' and d_i, d_i'.

Consider ϕ' the conjunction of all these constraints. Observe that the formulas ϕ and ϕ' are equivalent when quantifying on the fresh variables. Moreover, the models of ϕ and ϕ' are in one-to-one correspondence. Each model σ of ϕ can be extended to a model σ' of ϕ' whose weight is $|\sigma'| = |\sigma| + 4p + 1$. Consequently, x belongs to a cardinality-minimal model of ϕ if and only if x belongs to a cardinality-minimal model of ϕ', thus concluding the proof.

□

Lemma 9. *Let $\langle\Gamma\rangle = \mathsf{II}_1$. Then $\mathrm{CMS}(\Gamma)$ is Θ_2^{P}-hard.*

Proof. Let R_{II_1} be the minimal weak base of II_1, that is, $R_{\mathsf{II}_1}(x_1, x_2, x_3, x_4) = (x_1 \vee x_2) \wedge (x_1 x_2 \leftrightarrow x_3) \wedge \mathrm{T}(x_4)$. The matrix representation is as follows.

$$R_{\mathsf{II}_1} = \begin{pmatrix} 0101 \\ 1001 \\ 1111 \end{pmatrix}$$

We show that $\mathrm{CMS}(\mathrm{OR}^2) \leq_{\mathrm{m}}^{\mathrm{P}} \mathrm{CMS}(R_{\mathsf{II}_1})$. Then the result follows from Lemma 5 and Proposition 4.

Let (ϕ, x) be an instance of $\mathrm{CMS}(\mathrm{OR}^2)$, where $\phi = \bigwedge_{i=1}^{p}(x_i^1 \vee x_i^2)$. For each constraint $(x_i^1 \vee x_i^2)$ we build the constraint $R_{\mathsf{II}_1}(x_i^1, x_i^2, y_i, t)$. Observe that $(x_i^1 \vee x_i^2) \equiv \exists y_i \exists t R_{\mathsf{II}_1}(x_i^1, x_i^2, y_i, t)$. The variable t is frozen to 1. The variable y_i is not, but we can neutralize its weight by adding the constraint $R_{\mathsf{II}_1}(y_i, z_i, f, t)$, which will force $z_i \equiv \neg y_i$ as soon as f is evaluated to 0. We force f to be evaluated to 0 in any cardinality-minimal model by adding the constraints $R_{\mathsf{II}_1}(f_j^1, f_j^2, f, t)$, for $j = 1, \ldots, N$. If $f = 1$, these constraints force all the f_j^1, f_j^2 to 1, that is, $1 + 2N$ variables. If $f = 0$, one of the f_j^1, f_j^2 is forced to 1 and the other to 0, that is, the weight contribution is only N.

Consider ϕ' the conjunction of all these constraints. Observe that the formulas ϕ and ϕ' are equivalent when quantifying on the fresh variables. Moreover, the models of ϕ and the models of ϕ' in which $f = 0$ are in one-to-one correspondence. Each model σ of ϕ can be extended to a model σ' of ϕ' with $\sigma'(f) = 0$, whose weight is $|\sigma'| = |\sigma| + p + N + 1$.

Observe that ϕ is always satisfiable and therefore, by the above observation, ϕ' always admits a model with $f = 0$. Moreover, the models of ϕ' in which $f = 1$ are of cardinality at least $2N + 2$, while the models of ϕ' in which $f = 0$ are of cardinality at most $n + p + N + 1$, where n is the number of variables of ϕ. Now, if we choose N big enough, e.g. $N \geq p + n$, we ensure that an assignment with $f = 1$ can never be a cardinality-minimal model. Consequently, putting all together shows that x belongs to a cardinality-minimal model of ϕ if and only if x belongs to a cardinality-minimal model of ϕ', thus concluding the proof. □

Lemma 10. *Let* $\langle\Gamma\rangle = \mathsf{IN}_2$. *Then* $\textsc{Cms}(\Gamma)$ *is* Θ_2^{P}*-hard.*

Proof. Let R_{IN_2} be the minimal weak base of IN_2, that is, $R_{\mathsf{IN}_2} = \mathrm{EVEN}^4_{4\neq}(x_1, \ldots, x_8) \wedge x_1x_4 \leftrightarrow x_2x_3$. The matrix representation is as follows.

$$R_{\mathsf{IN}_2} = \begin{pmatrix} 00001111 \\ 00110011 \\ 01010101 \\ 10101010 \\ 11001100 \\ 11110000 \end{pmatrix}$$

We show that $\textsc{Cms}(\mathrm{OR}^2) \leq^{\mathrm{P}}_{\mathrm{m}} \textsc{Cms}(R_{\mathsf{IN}_2})$. Then the result follows from Lemma 5 and Proposition 4. Observe that $\mathrm{OR}^2(x_i^1, x_i^2) \equiv$

$$\exists a_i, b_i, c_i, d_i \; R_{\mathsf{IN}_2}(0, a_i, b_i, c_i, d_i, x_i^1, x_i^2, 1).$$

In this co-clone we can express $f \neq t$, but not $f = 0$ and $t = 1$. The idea is to use f and t in place of 0 and 1 as global variables (that is, the same for all constraints) and to force them to take the appropriate values in all cardinality-minimal models. This can be done by adding the constraint $R_{\mathsf{IN}_2}(f, f, f, f, t, t, t, t)$, which expresses $f \neq t$, and by adding a number N big enough of constraints $R_{\mathsf{IN}_2}(f_j, f_j, f_j, f_j, t, t, t, t)$, where the f_j, for $j = 1, \ldots, N$, are fresh variables.

In choosing N bigger than the number of variables of the original formula, we can assure that in any cardinality-minimal model f is assigned 0 and t is assigned 1. In using this trick we can mimic the reduction proposed in the proof of Proposition 8 and hence transform an OR^2-formula into an R_{IN_2}-formula in preserving the cardinality-minimal models, thus providing the reduction from $\textsc{Cms}(\mathrm{OR}^2)$ to $\textsc{Cms}(R_{\mathsf{IN}_2})$. □

Lemma 11. *Let* $\langle\Gamma\rangle = \mathsf{IL}_2$. *Then* $\textsc{Cms}(\Gamma)$ *is* Θ_2^{P}*-hard.*

Proof. Let R_{IL_2} be the minimal weak base of IL_2, that is, $R_{\mathsf{IL}_2} = \mathrm{EVEN}^3_{3\neq}(x_1, \ldots, x_6) \wedge \mathrm{F}(x_7) \wedge \mathrm{T}(x_8)$. The matrix representation is as follows.

$$R_{\mathsf{IL}_2} = \begin{pmatrix} 00011101 \\ 01110001 \\ 10101001 \\ 11000101 \end{pmatrix}$$

We show that $\textsc{Cms}(\mathrm{XOR}^3) \leq^{\mathrm{P}}_{\mathrm{m}} \textsc{Cms}(R_{\mathsf{IL}_2})$. Then the result follows from Lemma 6 and Proposition 4.

Let (ϕ, x) be an instance of $\textsc{Cms}(\mathrm{XOR}^3)$, where $\phi = \bigwedge_{i=1}^{p}(x_i^1 \oplus x_i^2 \oplus x_i^3)$. Let $\{u_i, v_i, w_i, u_i', v_i', w_i' \mid i = 1 \ldots p\} \cup \{t, f\}$ be fresh variables. For each constraint $(x_i^1 \oplus x_i^2 \oplus x_i^3)$ we build the constraint $R_{\mathsf{IL}_2}(u_i, v_i, w_i, x_i^1, x_i^2, x_i^3, f, t)$. Observe that $\mathrm{XOR}^3(x_i^1, x_i^2, x_i^3) \equiv$

$$\exists f, t, u_i, v_i, w_i \; R_{\mathsf{IL}_2}(u_i, v_i, w_i, x_i^1, x_i^2, x_i^3, f, t).$$

The variables f, t are frozen, and the other existential variables are uniquely determined and can be neutralized by adding three additional variables u_i', v_i', w_i' and the constraint $R_{\mathsf{IL}_2}(u_i, v_i, w_i, u_i', v_i', w_i', f, t)$.

Consider ϕ' the conjunction of all these constraints. Observe that the formulas ϕ and ϕ' are equivalent when quantifying on the fresh variables. Moreover, the models of ϕ and ϕ' are in one-to-one correspondence. Each model σ of ϕ can be extended to a model σ' of ϕ' whose weight is $|\sigma'| = |\sigma| + 3p + 1$. Consequently, x belongs to a cardinality-minimal model of ϕ if and only if x belongs to a cardinality-minimal model of ϕ', thus concluding the proof. □

Lemma 12. *Let* $\langle\Gamma\rangle = \mathsf{IL}_3$. *Then* $\textsc{Cms}(\Gamma)$ *is* Θ_2^{P}*-hard.*

Proof. Let R_{IL_3} be the minimal weak base of IL_3, that is, $R_{\mathsf{IL}_3} = \mathrm{EVEN}^4_{4\neq}(x_1, \ldots, x_8)$.

We show that $\textsc{Cms}(\mathrm{XOR}^3) \leq^{\mathrm{P}}_{\mathrm{m}} \textsc{Cms}(R_{\mathsf{IL}_3})$. Then the result follows from Lemma 6 and Proposition 4.

All relations in the co-clone IL_3 are complementive. In particular, up to some permutation of the variables (columns) R_{IL_3} is the complementive closure of R_{IL_2}. Thus, we use the same reduction idea as for IL_2, and then proceed analogously as we have done for the case of IN_2: replace 0 and 1 by f and t, express $f \neq t \equiv R_{\mathsf{IL}_3}(f, f, f, f, t, t, t, t)$ and put a big weight on f, hence forcing $f = 0$ and $t = 1$ in any cardinality-minimal model.

□

This concludes the proof of Theorem 2. We restate here Theorem 2 in terms of co-clones.

Theorem 13. *Let* Γ *be a finite constraint language. The problem* $\textsc{Cms}(\Gamma)$ *is*

- Θ_2^{P}*-complete if* $\mathsf{C} \subseteq \langle\Gamma\rangle \subseteq \mathsf{II}_2$ *for* $\mathsf{C} \in \{\mathsf{IS}_0^2, \mathsf{IL}_3, \mathsf{IL}_1\}$,
- *polynomial time solvable if either* $\mathsf{ID} \subseteq \langle\Gamma\rangle \subseteq \mathsf{ID}_1$ *or* $\mathsf{IR}_1 \subseteq \langle\Gamma\rangle \subseteq \mathsf{IE}_2$,
- *trivial otherwise (*Γ *is 0-valid)*

Example of Application

Let us now consider the following relevance problem for abduction. A *propositional abduction problem* (PAP) $\mathcal{P}$ consists of a tuple $\langle V, H, M, T\rangle$, where V is a finite set of *variables*, $H \subseteq V$ is the set of *hypotheses*, $M \subseteq V$ is the set of *manifestations*, and T is a consistent *theory* in the form of a propositional formula. A set $\mathcal{S} \subseteq H$ is a *solution* (also called *explanation*) to $\mathcal{P}$ if $T \cup \mathcal{S}$ is consistent and $T \cup \mathcal{S} \models M$ holds. Often, one is not interested in any solution of a given PAP $\mathcal{P}$ but only in minimal solutions, where minimality is defined w.r.t. set inclusion or smaller cardinality.

For subset-minimality the relevance problem has been completely classified in Schaefer's framework by Creignou and Zanuttini (2006). Here we consider the following decision problem.

CARD-MIN-RELEVANCE
Instance : PAP $\mathcal{P} = \langle V, H, M, T\rangle$ and hypothesis $h \in H$.
Question : Is h relevant, i.e., does $\mathcal{P}$ admit a cardinality-minimal solution $\mathcal{S}$ such that $h \in \mathcal{S}$?

It is known that the CARD-MIN-RELEVANCE problem is Θ_3^{P}-complete in its full generality and Θ_2^{P}-complete in the Horn case (Eiter and Gottlob 1995). The Krom case has been considered afterwards (Creignou, Pichler, and Woltran 2018). The complexity results obtained so far for the CARD-MIN-RELEVANCE problem were restricted, due to an incomplete picture of the complexity of CARDMINSAT. With the help of Theorem 2 we extend these results in showing that the complexity of CARD-MIN-RELEVANCE in the affine case matches the Horn and Krom cases.

Theorem 14. CARD-MIN-RELEVANCE *is Θ_2^{P}-complete even if the theory is restricted to XOR-CNF-formulas.*

Proof. Membership follows from the fact that one can decide the satisfiability of an XOR-CNF formula in polynomial time. The hardness proof is obtained via a reduction from $\mathrm{CMS}(\mathrm{XOR}^3)$.

Consider an arbitrary instance (ϕ, x_i) of $\mathrm{CMS}(\mathrm{XOR}^3)$. Let $\phi = \bigwedge_{i=1}^{p}(x_i^1 \oplus x_i^2 \oplus x_i^3)$ over variables $X = \{x_1, \ldots, x_n\}$ and let $G = \{g_1, \ldots, g_p\}$ be a set of fresh, pairwise distinct variables. We define the PAP $\mathcal{P} = \langle V, H, M, T\rangle$ as follows:

$$\begin{aligned} V &= X \cup G \\ H &= X \\ M &= G \\ T &= \{(x_i^1 \oplus x_i^2 \oplus x_i^3 \oplus \bar{g}_i) \mid 1 \leq i \leq p\} \end{aligned}$$

It is easy to verify that the models of ϕ coincide with the solutions of $\mathcal{P}$. Hence, x_i is in a cardinality-minimal model of ϕ if and only if x_i is in a cardinality-minimal solution of $\mathcal{P}$.

Note that, more precisely, the proof shows the hardness of CARD-MIN-RELEVANCE for XOR^4-formulas.

□

Conclusion

In this paper we obtained a complete complexity classification of the problem $\mathrm{CARDMINSAT}(\Gamma)$ for all finite constraint languages Γ: if Γ is width-2-affine, Horn or 0-valid, $\mathrm{CARDMINSAT}(\Gamma)$ is solvable in polynomial time, otherwise it is Θ_2^{P}-complete. The weak base method developed by Schnoor and Schnoor (2008), completed with the description of minimal weak bases for co-clones by Lagerkvist (2014) proved to be a valuable tool for this endeavor. As described in the introduction understanding the complexity of CARDMINSAT is crucial for the study of several reasoning tasks in artificial intelligence that are based on minimizing cardinality. As we have motivated and outlined above we believe the establishment of the complete complexity picture of $\mathrm{CARDMINSAT}(\Gamma)$ is a cornerstone for future research in this direction: it will allow the precise analysis of the computational complexity of problems such as relevance questions and belief revision operators. To obtain a richer picture we further plan to investigate the parametrized complexity of such problems. For instance, in (Mahmood, Meier, and Schmidt 2021) a rich picture of the parametrized complexity of abduction problems is obtained. Yet, the named abduction relevance problem in this picture is missing. With the now established complete complexity classification of CARDMINSAT it seems in reach to complete this picture.

Acknowledgments

This research has been supported by the Centre International de Recherche Mathématique (CIRM), Project reference *Research in pairs* 1886.

References

Böhler, E.; Reith, S.; Schnoor, H.; and Vollmer, H. 2005. Bases for Boolean co-clones. *Inf. Process. Lett.*, 96(2): 59–66.

Creignou, N.; Egly, U.; and Schmidt, J. 2014. Complexity Classifications for Logic-Based Argumentation. *ACM Trans. Comput. Log.*, 15(3): 19:1–19:20.

Creignou, N.; Ktari, R.; and Papini, O. 2017. Complexity of Model Checking for Cardinality-Based Belief Revision Operators. In *In Proc. ECSQARU 2017*, volume 10369 of *LNCS*, 387–397. Springer.

Creignou, N.; Olive, F.; and Schmidt, J. 2023. Complexity of Reasoning with Cardinality Minimality Conditions. *CoRR*, abs/2303.01571.

Creignou, N.; Pichler, R.; and Woltran, S. 2018. Do Hard SAT-Related Reasoning Tasks Become Easier in the Krom Fragment? *Log. Methods Comput. Sci.*, 14(4).

Creignou, N.; and Vollmer, H. 2008. Boolean Constraint Satisfaction Problems: When Does Post's Lattice Help? In Creignou, N.; Kolaitis, P. G.; and Vollmer, H., eds., *Complexity of Constraints - An Overview of Current Research Themes [Result of a Dagstuhl Seminar]*, volume 5250 of *Lecture Notes in Computer Science*, 3–37. Springer.

Creignou, N.; and Zanuttini, B. 2006. A Complete Classification of the Complexity of Propositional Abduction. *SIAM J. Comput.*, 36(1): 207–229.

Dalal, M. 1988. Investigations into Theory of Knowledge Base Revision. In *Proc. AAAI*, 475–479. AAAI Press / The MIT Press.

Eiter, T.; and Gottlob, G. 1992. On the Complexity of Propositional Knowledge Base Revision, Updates, and Counterfactuals. *Artif. Intell.*, 57(2-3): 227–270.

Eiter, T.; and Gottlob, G. 1995. The Complexity of Logic-Based Abduction. *J. ACM*, 42(1): 3–42.

Gasarch, W. I.; Krentel, M. W.; and Rappoport, K. J. 1995. OptP as the Normal Behavior of NP-Complete Problems. *Mathematical Systems Theory*, 28(6): 487–514.

Jeavons, P. G. 1998. On the algebraic structure of combinatorial problems. *Theor. Comput. Sci.*, 200: 185–204.

Krentel, M. W. 1988. The Complexity of Optimization Problems. *J. Comput. Syst. Sci.*, 36(3): 490–509.

Lagerkvist, V. 2014. Weak bases of Boolean co-clones. *Inf. Process. Lett.*, 114(9): 462–468.

Liberatore, P.; and Schaerf, M. 2001. Belief Revision and Update: Complexity of Model Checking. *J. Comput. Syst. Sci.*, 62(1): 43–72.

Mahmood, Y.; Meier, A.; and Schmidt, J. 2021. Parameterized complexity of abduction in Schaefer's framework. *J. Log. Comput.*, 31(1): 266–296.

Nordh, G. 2004. A Trichotomy in the Complexity of Propositional Circumscription. In Baader, F.; and Voronkov, A., eds., *Logic for Programming, Artificial Intelligence, and Reasoning, 11th International Conference, LPAR 2004, Montevideo, Uruguay, March 14-18, 2005, Proceedings*, volume 3452 of *Lecture Notes in Computer Science*, 257–269. Springer.

Nordh, G.; and Zanuttini, B. 2008. What makes propositional abduction tractable. *Artif. Intell.*, 172(10): 1245–1284.

Nordh, G.; and Zanuttini, B. 2009. Frozen boolean partial co-clones. In *Proc. ISMVL*, 120–125.

Papadimitriou, C. H. 1994. *Computational Complexity*. Addison-Wesley.

Post, E. L. 1941. The two-valued iterative systems of mathematical logic. *Annals of Mathematical Studies*, 5: 1–122.

Schaefer, T. J. 1978. The Complexity of Satisfiability Problems. In *Proc. STOC*, 216–226. ACM.

Schnoor, H.; and Schnoor, I. 2008. Partial polymorphisms and constraint satisfaction problems. In Creignou, N.; Kolaitis, P. G.; and Vollmer, H., eds., *Complexity of Constraints*, volume 5250, 229–254. Berlin Heidelberg: Springer Verlag.

Wagner, K. 1988. On Restricting the Access to an NP-Oracle. In *Proc. ICALP*, volume 317 of *LNCS*, 682–696. Springer.