Give Me Another One!

Mike Behrisch[1], Miki Hermann[2(✉)], Stefan Mengel[2], and Gernot Salzer[1]

[1] Technische Universität Wien, Vienna, Austria
{behrisch,salzer}@logic.at
[2] LIX (UMR CNRS 7161), École Polytechnique, Palaiseau, France
{hermann,mengel}@lix.polytechnique.fr

Abstract. We investigate the complexity of an optimization problem in Boolean propositional logic related to information theory: Given a conjunctive formula over a set of relations, find a satisfying assignment with minimal Hamming distance to a given assignment that satisfies the formula (NearestOtherSolution, NOSol).

We present a complete classification with respect to the relations admitted in the formula. We give polynomial-time algorithms for several classes of constraint languages. For all other cases we prove hardness or completeness regarding poly-APX, NPO, or equivalence to a well-known hard optimization problem.

1 Introduction

We investigate the solution spaces of Boolean constraint satisfaction problems built from atomic constraints by means of conjunction and variable identification. We study a minimization problem in connection with Hamming distance: Given an instance of a constraint satisfaction problem in the form of a generalized conjunctive formula over a set of atomic constraints, the problem asks to find a satisfying assignment with minimal Hamming distance to a given assignment that satisfies the formula (NearestOtherSolution, NOSol).

As it is common, we analyze the complexity of our optimization problem through a parameter, representing the atomic constraints allowed to be used in the constraint satisfaction problem. We give a complete classification of the complexity of approximation with respect to this parameterization. It turns out that our problems can either be solved in polynomial time, or they are complete for a well-known optimization class, or else they are equivalent to well-known hard optimization problems.

Our study can be understood as a continuation of the minimization problems investigated by Khanna et al. in [11], especially that of MinOnes. The MinOnes

Mike Behrisch and Gernot Salzer—Supported by Austrian Science Fund (FWF) grant I836-N23.
Miki Hermann—Supported by ANR-11-ISO2-003-01 Blanc International grant ALCOCLAN.
Stefan Mengel—Supported by QUALCOMM grant administered by École Polytechnique.

K. Elbassioni and K. Makino (Eds.): ISAAC 2015, LNCS 9472, pp. 664–676, 2015.
DOI: 10.1007/978-3-662-48971-0_56

Table 1. Boolean co-clones with bases.

iS_0^k	$\{\mathrm{or}^k\}$	iL	$\{\mathrm{even}^4\}$	iN	$\{\mathrm{dup}^3\}$
iS_1^k	$\{\mathrm{nand}^k\}$	iL_2	$\{\mathrm{even}^4, \neg x, x\}$	iN_2	$\{\mathrm{nae}^3\}$
iS_{00}^k	$\{\mathrm{or}^k, x \to y, \neg x, x\}$	iV	$\{x \vee y \vee \neg z\}$	iI	$\{\mathrm{even}^4, x \to y\}$
iS_{10}^k	$\{\mathrm{nand}^k, \neg x, x, x \to y\}$	iV_2	$\{x \vee y \vee \neg z, \neg x, x\}$	iI_0	$\{\mathrm{even}^4, x \to y, \neg x\}$
iD_1	$\{x \oplus y, x\}$	iE	$\{\neg x \vee \neg y \vee z\}$	iI_1	$\{\mathrm{even}^4, x \to y, x\}$
iD_2	$\{x \oplus y, x \to y\}$	iE_2	$\{\neg x \vee \neg y \vee z, \neg x, x\}$	iM_2	$\{x \to y, \neg x, x\}$

optimization problem asks for a solution of a constraint satisfaction problem with the minimal Hamming weight, i.e., minimal Hamming distance to the 0-vector. Our work generalizes this by allowing the given vector to be any, potentially also non-0-vector. Moreover, our work can also be seen as a generalization of questions in coding theory.

It turns out that our problem NOSol lacks compatibility with existential quantification, which makes classical clone theory inapplicable. Therefore, we have to resort to weak co-clones requiring only closure under conjunction and equality. To dispose of the latter we apply the theory developed in [14], as well as minimal weak bases of Boolean co-clones from [12].

2 Preliminaries

An n-ary *Boolean relation* R is a subset of $\{0,1\}^n$; its elements $(b_1, \ldots, b_n)$ are also written as $b_1 \cdots b_n$. Let V be a set of variables. An *atomic constraint*, or an *atom*, is an expression $R(\boldsymbol{x})$, where R is an n-ary relation and $\boldsymbol{x}$ is an n-tuple of variables from V. Let $\mathcal{L}$ be the collection of all non-empty finite sets of Boolean relations, also called *constraint languages*. For $\Gamma \in \mathcal{L}$, a *Γ-formula* is a finite conjunction of atoms $R_1(\boldsymbol{x_1}) \wedge \cdots \wedge R_k(\boldsymbol{x_k})$, where the R_i are relations from Γ and the $\boldsymbol{x_i}$ are variable tuples of suitable arity.

An *assignment* is a mapping $m\colon V \to \{0,1\}$ assigning a Boolean value $m(x)$ to each variable $x \in V$. If we arrange the variables in some arbitrary but fixed order, say as a tuple $(x_1, \ldots, x_n)$, then the assignments can be identified with vectors from $\{0,1\}^n$. The i-th component of a vector m is denoted by $m[i]$ and corresponds to the value of the i-th variable, i.e., $m[i] = m(x_i)$. The *Hamming weight* $\mathrm{hw}(m) = |\{i \mid m[i] = 1\}|$ of m is the number of 1s in the vector m. The *Hamming distance* $\mathrm{hd}(m, m') = |\{i \mid m[i] \neq m'[i]\}|$ of m and m' is the number of coordinates on which the vectors disagree. The *complement* $\overline{m}$ of a vector m is its pointwise complement, $\overline{m}[i] = 1 - m[i]$.

An assignment m satisfies the constraint $R(x_1, \ldots, x_n)$ if $(m(x_1), \ldots, m(x_n)) \in R$ holds. It satisfies the formula φ if it satisfies all of its atoms; m is said to be a *model* or *solution* of φ in this case. We use $[\varphi]$ to denote the set of models of φ. Note that $[\varphi]$ represents a Boolean relation. In sets of relations represented this way we usually omit the brackets. A *literal* is a variable v, or its negation $\neg v$. Assignments m are extended to literals by defining $m(\neg v) = 1 - m(v)$.

We shall need the following Boolean functions and relations later: By $x \oplus y$ we denote addition modulo 2 and $x \equiv y$ means $x \oplus y \oplus 1$. Further, we let $\mathrm{nae}^3 := \{0,1\}^3 \smallsetminus \{000, 111\}$, $\mathrm{dup}^3 := \{0,1\}^3 \smallsetminus \{010, 101\}$ and $\mathrm{even}^4 := \{(a_1, a_2, a_3, a_4) \in \{0,1\}^4 \mid \oplus_{i=1}^4 a_i = 0\}$, as well as $S_0 := [x_1 \wedge x_4 \equiv x_2 \wedge x_3]$, $S_1 := [S_0(x_1, x_2, x_3, x_1)]$ and $S_2 := [x_1 \vee x_2 \rightarrow x_3]$. Moreover, for $k \geq 1$ we define $\mathrm{or}^k := \{0,1\}^k \smallsetminus \{0 \cdots 0\}$ and $\mathrm{nand}^k := \{0,1\}^k \smallsetminus \{1 \cdots 1\}$.

Throughout the text we refer to different types of Boolean constraint relations following Schaefer's terminology [13] (see also [4,6]). A Boolean relation R is (1) 1-*valid* if $1 \cdots 1 \in R$ and it is 0-*valid* if $0 \cdots 0 \in R$, (2) *Horn* (*dual Horn*) if R can be represented by a formula in conjunctive normal form (CNF) having at most one unnegated (negated) variable in each clause, (3) *monotone* if it is both Horn and dual Horn, (4) *bijunctive* if it can be represented by a CNF having at most two variables in each clause, (5) *affine* if it can be represented by an affine system of equations $Ax = b$ over $\mathbb{Z}_2$, (6) *complementive* if for each $m \in R$ also $\overline{m} \in R$. A set Γ of Boolean relations is called 0-valid (1-valid, Horn, dual Horn, monotone, affine, bijunctive, complementive) if *every* $R \in \Gamma$ satisfies that property.

A formula constructed from atoms by conjunction, variable identification, and existential quantification is called a *primitive positive formula* (*pp-formula*). We denote by $\langle \Gamma \rangle$ the set of all relations that can be expressed using relations from $\Gamma \cup \{=\}$, conjunction, variable identification, and existential quantification. The set $\langle \Gamma \rangle$ is called the *co-clone* generated by Γ. A *base* of a co-clone $\mathcal{B}$ is a set of relations Γ, such that $\langle \Gamma \rangle = \mathcal{B}$. All co-clones, ordered by set inclusion, form a lattice. Together with their respective bases, which were studied in [5], some of them are listed in Table 1. In particular the sets of relations being 0-valid, 1-valid, complementive, Horn, dual Horn, affine, bijunctive, 2affine (both bijunctive and affine), and monotone each form a co-clone denoted by iI_0, iI_1, iN_2, iE_2, iV_2, iL_2, iD_2, iD_1, and iM_2, respectively.

We will also use a weaker closure than $\langle \Gamma \rangle$, called *conjunctive closure* and denoted by $\langle \Gamma \rangle_\wedge$, where the constraint language Γ is closed under conjunctive definitions, but not under existential quantification or addition of explicit equality constraints.

Minimal weak bases of co-clones are bases with certain additional properties. Since we rely on only some of them, we shall not define this term but refer the reader to [12,14].

Theorem 1. *If Γ is a minimal weak base of a co-clone, then $\Gamma \subseteq \langle \Gamma' \rangle_\wedge$ for any base Γ'.*

Lagerkvist computed weak bases for all Boolean co-clones in [12]. From there we infer that each co-clone $\mathcal{B} \in \{\mathrm{iE}, \mathrm{iE}_0, \mathrm{iE}_1, \mathrm{iE}_2, \mathrm{iN}, \mathrm{iN}_2, \mathrm{iI}\}$ has a singleton minimal weak base $\{R_{\mathcal{B}}\}$, in which $R_{\mathrm{iE}} := (S_1 \times \{0,1\}) \cap (\{0,1\} \times S_2)$, $R_{\mathrm{iE}_0} := R_{\mathrm{iE}} \times \{0\}$, $R_{\mathrm{iE}_1} := S_1 \times \{1\}$, $R_{\mathrm{iE}_2} := S_1 \times \{0\} \times \{1\}$, $R_{\mathrm{iN}} := \mathrm{even}^4 \cap S_0$, $R_{\mathrm{iN}_2} := [R_{\mathrm{iN}}(x_1, \ldots, x_4) \wedge \bigwedge_{i=1}^4 x_{i+4} = \neg x_i]$ and $R_{\mathrm{iI}} := [S_1(x_1, x_2, x_3) \wedge S_1(\neg x_4, \neg x_2, \neg x_3)]$.

We assume that the reader has a basic knowledge of approximation algorithms and complexity theory, see e.g. [1,6]. For reductions among decision problems we use polynomial-time many-one reduction denoted by $\leq_{\mathrm{m}}$. Many-one equivalence between decision problems is written as $\equiv_{\mathrm{m}}$. For reductions among optimization problems we employ approximation preserving reductions (AP-reductions), represented by $\leq_{\mathrm{AP}}$. AP-equivalence of optimization problems is stated as $\equiv_{\mathrm{AP}}$. Besides, the following approximation complexity classes in the hierarchy PO $\subseteq$ APX $\subseteq$ poly-APX $\subseteq$ NPO occur.

We also need a slightly non-standard variation of AP-reductions between optimization problems $\mathcal{P}_1$, $\mathcal{P}_2$: Viz., $\mathcal{P}_1$ AP-*Turing-reduces* to $\mathcal{P}_2$ if there is a polynomial-time oracle algorithm $\mathbb{A}$ and a constant $\alpha \geq 1$ such that for all $r > 1$ on any input x for $\mathcal{P}_1$ we have

- if all oracle calls within $\mathbb{A}$ upon inputs for $\mathcal{P}_2$ are answered with feasible solutions for $\mathcal{P}_2$, then $\mathbb{A}$ outputs a feasible solution for $\mathcal{P}_1$ on input x, and
- if for every call in $\mathbb{A}$ the oracle answers with an r-approximate solution, then $\mathbb{A}$ computes a $(1 + (r-1)\alpha + o(1))$-approximate solution for $\mathcal{P}_1$ on input x.

It is straightforward to check that AP-Turing-reductions are transitive. Moreover, if $\mathcal{P}_1$ AP-Turing-reduces to $\mathcal{P}_2$ with constant α and $\mathcal{P}_2$ has an $f(n)$-approximation algorithm, then there is an $\alpha f(n)$-approximation algorithm for $\mathcal{P}_1$.

To relate our problem to well-known optimization problems we make the following convention: For optimization problems $\mathcal{P}$ and $\mathcal{Q}$ we say that $\mathcal{Q}$ is $\mathcal{P}$-*hard* if $\mathcal{P} \leq_{\mathrm{AP}} \mathcal{Q}$, i.e. if $\mathcal{P}$ reduces to it. Moreover, $\mathcal{Q}$ is called $\mathcal{P}$-*complete* if $\mathcal{P} \equiv_{\mathrm{AP}} \mathcal{Q}$. We use these notions in particular with respect to the following problems from [11], taking parameters $\Gamma \in \mathcal{L}$.

Problem $\mathsf{MinOnes}(\Gamma)$**.** Given a conjunctive formula φ over relations from Γ, a solution is any assignment m satisfying φ. The goal is to minimize the Hamming weight $\mathrm{hw}(m)$.

Problem $\mathsf{WeightedMinOnes}(\Gamma)$**.** Given a conjunctive formula φ over relations from Γ and a weight function $w\colon V \to \mathbb{N}$ on the variables V of φ, a solution is again any assignment m satisfying φ. The objective is to minimize the value $\sum_{x:m(x)=1} w(x)$.

We now define some well-studied problems to which we will relate our problems. Note that these problems do not depend on any parameter.

Problem $\mathsf{MinDistance}$**.** Given a matrix $A \in \mathbb{Z}_2^{k \times l}$ any non-zero vector $x \in \mathbb{Z}_2^l$ with $Ax = 0$ is considered a solution. The aim is to minimize the Hamming weight $\mathrm{hw}(x)$.

Problem $\mathsf{MinHornDeletion}$**.** For a conjunctive formula φ over relations from the constraint language $\{[x \vee y \vee \neg z], [x], [\neg x]\}$, an assignment m satisfying φ is feasible. The objective is given by the minimum number of unsatisfied conjuncts of φ.

$\mathsf{MinDistance}$ and $\mathsf{MinHornDeletion}$ are NP-hard to approximate within $2^{\Omega(\log^{1-\varepsilon}(n))}$ for all $\varepsilon > 0$ [8,11]. Thus, unless P = NP, both are inequivalent to any problem $\mathcal{P} \in$ APX.

We also use the classic satisfiability problem $\mathsf{SAT}(\Gamma)$, asking for a conjunctive formula φ over a $\Gamma \in \mathcal{L}$, if φ is satisfiable. Schaefer presented in [13] a complete classification of complexity for $\mathsf{SAT}(\Gamma)$. His dichotomy theorem proves that

$\mathsf{SAT}(\Gamma)$ is in P if Γ is 0-valid ($\Gamma \subseteq \mathrm{iI}_0$), 1-valid ($\Gamma \subseteq \mathrm{iI}_1$), Horn ($\Gamma \subseteq \mathrm{iE}_2$), dual Horn ($\Gamma \subseteq \mathrm{iV}_2$), bijunctive ($\Gamma \subseteq \mathrm{iD}_2$), or affine ($\Gamma \subseteq \mathrm{iL}_2$); otherwise it is NP-complete. Moreover, we need the decision problem $\mathsf{AnotherSAT}(\Gamma)$, asking for a conjunctive formula φ over Γ and a model m, if there is another model $m' \neq m$ for φ. In [10] Juban completely classified the complexity of AnotherSAT. His dichotomy result shows $\mathsf{AnotherSAT}(\Gamma)$ to be polynomial-time decidable if Γ is both 0- and 1-valid ($\Gamma \subseteq \mathrm{iI}$), complementive ($\Gamma \subseteq \mathrm{iN}_2$), Horn ($\Gamma \subseteq \mathrm{iE}_2$), dual Horn ($\Gamma \subseteq \mathrm{iV}_2$), bijunctive ($\Gamma \subseteq \mathrm{iD}_2$), or affine ($\Gamma \subseteq \mathrm{iL}_2$); or else to be NP-complete.

3 Results

Here we present the formal definition of our considered problem, with parameter $\Gamma \in \mathcal{L}$, and our results; the proofs follow in subsequent sections.

Problem $\mathsf{NearestOtherSolution}(\Gamma)$, $\mathsf{NOSol}(\Gamma)$
Input: A conjunctive formula φ over relations from Γ and an assignment m satisfying φ.
Solution: Another assignment m' satisfying φ.
Objective: Minimum Hamming distance $\mathrm{hd}(m, m')$.

Theorem 2. *For every $\Gamma \in \mathcal{L}$ the optimization problem $\mathsf{NOSol}(\Gamma)$ is*

(i) in PO *if*
  (a) Γ is bijunctive ($\Gamma \subseteq \mathrm{iD}_2$) *or*
  (b) $\Gamma \subseteq \langle x_1 \vee \cdots \vee x_k, x \to y, \neg x, x\rangle$ for some $k \in \mathbb{N}$, $k \geq 2$ ($\Gamma \subseteq \mathrm{iS}_{00}^k$) *or*
  (c) $\Gamma \subseteq \langle \neg x_1 \vee \cdots \vee \neg x_k, x \to y, \neg x, x\rangle$ for some $k \in \mathbb{N}$, $k \geq 2$ ($\Gamma \subseteq \mathrm{iS}_{10}^k$)*;*
(ii) MinDistance*-complete if Γ is exactly affine* ($\mathrm{iL} \subseteq \langle\Gamma\rangle \subseteq \mathrm{iL}_2$)*;*
(iii) MinHornDeletion*-complete under* AP*-Turing-reductions if Γ is*
  (a) exactly Horn ($\mathrm{iE} \subseteq \langle\Gamma\rangle \subseteq \mathrm{iE}_2$) *or*
  (b) exactly dual Horn ($\mathrm{iV} \subseteq \langle\Gamma\rangle \subseteq \mathrm{iV}_2$)*;*
(iv) in poly-APX *if Γ is*
  (a) exactly both 0*-valid and* 1*-valid* ($\langle\Gamma\rangle = \mathrm{iI}$) *or*
  (b) exactly complementive ($\mathrm{iN} \subseteq \langle\Gamma\rangle \subseteq \mathrm{iN}_2$)*,*
  where $\mathsf{NOSol}(\Gamma)$ is n-approximable but not $(n^{1-\varepsilon})$-approximable unless P = NP*;*
(v) and NPO*-complete otherwise* ($\mathrm{iI}_0 \subseteq \langle\Gamma\rangle$ *or* $\mathrm{iI}_1 \subseteq \langle\Gamma\rangle$)*.*

The optimization problem can be transformed into a decision problem as usual. We add a bound $k \in \mathbb{N}$ to the input and ask if $\mathrm{hd}(m, m') \leq k$. This way we obtain the corresponding decision problem $\mathsf{NOSol}^{\mathrm{d}}$. Its complexity follows immediately from the theorems above. All cases in PO become polynomial-time decidable, whereas the other cases, which are APX-hard, become NP-complete. This way we obtain a dichotomy theorem classifying the decision problems as polynomial or NP-complete for all finite sets of relations Γ.

4 Duality and Inapplicability of Clone Closure

The problem NOSol is not compatible with existential quantification as the following shows:

Example 3. Consider the relation $R = \{00000, 01111, 10101\}$ and let (φ_R, m) be an instance of NOSol with $\varphi_R = R(x_1, \ldots, x_5)$ and $m = 10101$. Both $m_1 = 00000$ and $m_2 = 01111$ are feasible solutions of φ_R and $\mathrm{hd}(m, m_1) = \mathrm{hd}(m, m_2) = 3$. Hence m_2 is an optimal solution of (φ_R, m). Let $m' = 1010$, $m'_1 = 0000$, and $m'_2 = 0111$ be new tuples, constructed from m, m_1, and m_2 respectively, by truncating the last coordinate. Hence, they are the solutions of $(\exists x_5\, \varphi_R, m')$. However, note that $\mathrm{hd}(m', m'_1) = 2$ and $\mathrm{hd}(m', m'_2) = 3$. The tuple m'_2 is not an optimal solution of $(\exists x_5\, \varphi_R, m')$.

Because of this incompatibility, we cannot prove an AP-equivalence result between any two NOSol problems parametrized by constraint languages generating the same co-clone. Yet, similar results hold for the conjunctive closure.

Proposition 4. *Let Γ and Γ' be constraint languages. If $\Gamma' \subseteq \langle\Gamma\rangle_\wedge$ holds then we have the reductions $\mathsf{NOSol}^{\mathrm{d}}(\Gamma') \leq_{\mathrm{m}} \mathsf{NOSol}^{\mathrm{d}}(\Gamma)$ and $\mathsf{NOSol}(\Gamma') \leq_{\mathrm{AP}} \mathsf{NOSol}(\Gamma)$.*

Proof. For similarity it suffices to show that $\mathsf{NOSol}(\Gamma') \leq_{\mathrm{AP}} \mathsf{NOSol}(\Gamma)$ if $\Gamma' \subseteq \langle\Gamma\rangle_\wedge$.

Let a formula φ with a model m be an instance of $\mathsf{NOSol}(\Gamma')$. As $\Gamma' \subseteq \langle\Gamma\rangle_\wedge$, every constraint $R(x_1, \ldots, x_k)$ of φ can be written as a conjunction of constraints upon relations from Γ. Substitute the latter into φ, obtaining φ'. Now (φ', m) is an instance of $\mathsf{NOSol}(\Gamma)$, where φ' is only polynomially larger than φ. For φ and φ' have the same variables and hence the same models, also the nearest other models of φ and φ' are the same. □

For a relation $R \subseteq \{0,1\}^n$, its *dual* relation is $\mathrm{dual}(R) = \{\overline{m} \mid m \in R\}$, i.e., the relation containing the complements of tuples from R. We naturally extend this to sets of relations Γ by putting $\mathrm{dual}(\Gamma) = \{\mathrm{dual}(R) \mid R \in \Gamma\}$. Since taking complements is involutive, duality is a symmetric relation. By inspecting the bases of co-clones in Table 1, we deduce that many co-clones are duals of each other, e.g. iE_2 and iV_2.

We now show that it suffices to consider one half of Post's lattice of co-clones.

Lemma 5. *For every Boolean constraint language Γ we have the mutual reductions $\mathsf{NOSol}^{\mathrm{d}}(\Gamma) \equiv_{\mathrm{m}} \mathsf{NOSol}^{\mathrm{d}}(\mathrm{dual}(\Gamma))$ and $\mathsf{NOSol}(\Gamma) \equiv_{\mathrm{AP}} \mathsf{NOSol}(\mathrm{dual}(\Gamma))$.*

Proof. For a Γ-formula φ and an assignment m to φ we construct a $\mathrm{dual}(\Gamma)$-formula φ' by substitution of every atom $R(\boldsymbol{x})$ by $\mathrm{dual}(R)(\boldsymbol{x})$. Then m satisfies φ if and only if $\overline{m}$ satisfies φ', $\overline{m}$ being the complement of m. Moreover, $\mathrm{hd}(m, m') = \mathrm{hd}(\overline{m}, \overline{m'})$. □

5 Finding Another Solution Closest to the Given One

5.1 Polynomial-Time Cases

Since we cannot take advantage of the clone closure, we must proceed differently. We use the following result based on a previous theorem of Baker and Pixley [2].

Proposition 6 (Jeavons et al. [9]). *Every bijunctive constraint $R(x_1, \ldots, x_n)$ is equivalent to $\bigwedge_{1 \leq i \leq j} R_{ij}(x_i, x_j)$, where R_{ij} is the projection of R to the coordinates i and j.*

Proposition 7. *If Γ is bijunctive ($\Gamma \subseteq \mathrm{iD}_2$) then* $\mathsf{NOSol}(\Gamma)$ *is in* PO.

Proof. According to Proposition 6 we may assume that the formula φ is a conjunction of atoms $R(x, y)$ or a unary constraint $R(x, x)$ in the form $[x]$ or $[\neg x]$. Unary constraints can be eliminated and their value propagated into the other clauses, since they fix the value for a given variable.

For each variable x we construct a model m_x of φ with $m_x(x) \neq m(x)$ such that $\mathrm{hd}(m_x, m)$ is minimal among all models with this property. Initially we set $m_x(x)$ to $1 - m(x)$ and $m_x(y) := m(y)$ for all variables $y \neq x$ and mark x as flipped. If m_x satisfies all atoms we are done. Otherwise let $R(u, v)$ be an atom falsified by m_x. If u and v are marked as flipped, the construction fails, a model m_x with the property $m_x(x) \neq m(x)$ does not exist. Otherwise the uniquely determined variable v in $R(u, v)$ is not marked as flipped. Set $m_x(v) := 1 - m(v)$, mark v as flipped, and repeat the process.

If m_x does not exist for any variable x, then m is the sole model of φ and the problem is not solvable. Otherwise choose one of the variables x for which $\mathrm{hd}(m_x, m)$ is minimal and return m_x as second solution m'. □

Proposition 8. *If $\Gamma \subseteq \mathrm{iS}_{00}^k$ or $\Gamma \subseteq \mathrm{iS}_{10}^k$ for some $k \geq 2$ then* $\mathsf{NOSol}(\Gamma)$ *is in* PO.

Proof. We perform the proof only for iS_{00}^k. Lemma 5 implies the same result for iS_{10}^k.

The co-clone iS_{00}^k is generated by $\Gamma' := \{\mathrm{or}^k, [x \rightarrow y], [x], [\neg x]\}$. According to [7], this set Γ' is also a so-called *plain basis* of iS_{00}^k, i.e. we may assume that our inputs (φ, m) contain conjunctive formulas φ over these relations and equality, without existential quantification.

Note that $x \vee y$ is a polymorphism of Γ, i.e., for any two solutions m_1, m_2 of φ we have that the assignment $m_1 \vee m_2$ which is defined by $(m_1 \vee m_2)(x) = m_1(x) \vee m_2(x)$ for every x is also a solution of φ. It follows that we get the optimal solution m' for the instance φ and m by either flipping some values 1 of m to 0 or flipping some values 0 of m to 1 but not both. To see this, assume the optimal solution m' flips both ones and zeros, then $m' \vee m$ is a solution of φ that is closer to m than m' which is a contradiction.

The main idea is to compute for each variable x of φ the distance of the solution m_x, which is minimal among the solutions of φ which differ from m

on the variable x, and flip only ones or only zeros. Then the algorithm chooses one m_x closest to m as m' and returns it. Since m and m' differ in at least one variable, this yields the correct result.

We describe the computation of m_x. If $m(x) = 0$, we flip x to 1 and propagate iteratively along equalities $x = z$ and $x \rightarrow y$-constraints, i.e., if $x \rightarrow y$ is a constraint of φ and $m(y) = 0$, we flip y to 1 and propagate. This process terminates after at most n flips, as we only flip from 0 to 1 and no variable is flipped more than once. If the resulting assignment satisfies φ, this is our m_x. Otherwise, there is no satisfying assignment which we get by flipping x and only flipping 0 to 1 and thus no candidate m_x with the desired properties. If $m(x) = 1$, we flip x to 0 and propagate backward along equalities $x = z$ and binary implications, i.e., if $y \rightarrow x$ is a constraint of φ and $m(y) = 1$, we flip y to 0 and iterate. Again, if the result satisfies φ, this is our m_x; else, there is no candidate m_x for this variable. Finally, return the candidate m_x being closest to m if it exists, otherwise there is no feasible solution. □

5.2 Hard Cases

Lemma 9. *Let Γ be a constraint language. If* $\mathrm{iI}_1 \subseteq \langle\Gamma\rangle$ *or* $\mathrm{iI}_0 \subseteq \langle\Gamma\rangle$ *holds then finding a feasible solution for* $\mathsf{NOSol}(\Gamma)$ *is* NPO*-hard. Otherwise,* $\mathsf{NOSol}(\Gamma) \in$ poly-APX.

Proof. Finding a feasible solution to $\mathsf{NOSol}(\Gamma)$ is exactly the problem $\mathsf{AnotherSAT}(\Gamma)$ which is NP-hard if and only if $\mathrm{iI}_1 \subseteq \langle\Gamma\rangle$ or $\mathrm{iI}_0 \subseteq \langle\Gamma\rangle$ according to Juban [10]. If $\mathsf{AnotherSAT}(\Gamma)$ is polynomial-time decidable, we can always find a feasible solution for $\mathsf{NOSol}(\Gamma)$ if it exists. Obviously, every feasible solution is an n-approximation of the optimal solution, where n is the number of variables of the input. □

Tightness Results. It will be convenient to consider the following decision problem.

Problem: $\mathsf{AnotherSAT}_{<n}(\Gamma)$
Input: A conjunctive formula φ over relations from Γ and an assignment m satisfying φ.
Question: Is there another satisfying assignment m' of φ, different from m, such that $\mathrm{hd}(m, m') < n$, where n is the number of variables of φ?

Note that $\mathsf{AnotherSAT}_{<n}(\Gamma)$ is not compatible with existential quantification. Let $\varphi(y, x_1, \ldots, x_n)$ with the model m be an instance of $\mathsf{AnotherSAT}_{<n}(\Gamma)$ and m' its solution satisfying $\mathrm{hd}(m, m') < n + 1$. Let m_1 and m_1' be the corresponding vectors to m and m', respectively, with the first coordinate truncated. When we existentially quantify the variable y in φ, producing $\varphi_1(x_1, \ldots, x_n) = \exists y\, \varphi(y, x_1, \ldots, x_n)$, then both m_1 and m_1' are solutions of φ', but we cannot guarantee $\mathrm{hd}(m_1, m_1') < n$. Hence we need the equivalent of Proposition 4 for this problem, whose proof is analogous.

Proposition 10. $\mathsf{AnotherSAT}_{<n}(\Gamma') \leq_{\mathrm{m}} \mathsf{AnotherSAT}_{<n}(\Gamma)$ *for* $\Gamma, \Gamma' \in \mathcal{L}$, $\Gamma' \subseteq \langle \Gamma \rangle_{\wedge}$.

Proposition 11. *If* $\Gamma \in \mathcal{L}$ *with* $\langle \Gamma \rangle = \mathrm{iI}$ *or* $\mathrm{iN} \subseteq \langle \Gamma \rangle \subseteq \mathrm{iN}_2$, *then* $\mathsf{AnotherSAT}_{<n}(\Gamma)$ *is* NP*-complete.*

Proof. Containment in NP is clear, so it only remains to show hardness. Since the considered problem is not compatible with existential quantification, we cannot use clone theory and therefore we will consider the three co-clones iN_2, iN and iI individually, making use of minimal weak bases.

Case $\langle \Gamma \rangle = \mathrm{iN}$*:* We show a reduction from $\mathsf{AnotherSAT}(R)$ where $R = \{000, 101, 110\}$ which is NP-hard by [10]. Since R is 0-valid, $\mathsf{AnotherSAT}(R)$ is still NP-complete if we restrict it to instances $(\varphi, \mathbf{0})$, where φ is a conjunctive formula over R and $\mathbf{0}$ is the constant 0-assignment. Thus we can perform a reduction from this restricted problem.

By Theorem 1 and Proposition 10 we may assume that Γ contains the minimal weak base relation R_{iN}. Given a formula φ over R, we construct another formula φ' over R_{iN} by replacing every constraint $R(x_i, x_j, x_k)$ with $R_{\mathrm{iN}}(x_i, x_j, x_k, w)$, where w is a new global variable. Moreover, set m to the constant 0-assignment. This construction is a many-one reduction from the restricted version of $\mathsf{AnotherSAT}(R)$ to $\mathsf{AnotherSAT}_{<n}(\Gamma)$.

To see this, observe that the tuples in R_{iN} that have a 0 in the last coordinate are exactly those in $R \times \{0\}$. Thus any solution of φ can be extended to a solution of φ' by assigning 0 to w. Assume that φ' has a solution m which is not constant $\mathbf{0}$ or constant $\mathbf{1}$. Because R_{iN} is complementive, we may assume that $m(w) = 0$. But then m restricted to the variables of φ is not the constant 0-assignment and satisfies all constraints of φ. This completes the proof of the first case.

Case $\langle \Gamma \rangle = \mathrm{iN}_2$*:* We show a reduction from $\mathsf{AnotherSAT}_{<n}(R_{\mathrm{iN}})$ which is NP-hard by the previous case. Reasoning as before, we may assume that Γ contains $R_{\mathrm{iN}_2} = \{m\overline{m} \mid m \in R_{\mathrm{iN}}\}$. Given an R_{iN}-formula φ over the variables $x_1, \ldots, x_n$, we construct an R_{iN_2}-formula over the variables $x_1, \ldots, x_n, x'_1, \ldots, x'_n$ by replacing $R_{\mathrm{iN}}(x_i, x_j, x_k, x_\ell)$ with $R_{\mathrm{iN}_2}(x_i, x_j, x_k, x_\ell, x'_i, x'_j, x'_k, x'_\ell)$. Moreover, we define an assignment m' to φ' by setting $m'(x_i) := m(x_i)$ and $m'(x'_i) := \overline{m}(x_i)$. It is easy to see that this construction is a reduction from $\mathsf{AnotherSAT}_{<n}(R_{\mathrm{iN}})$ to $\mathsf{AnotherSAT}_{<n}(\Gamma)$.

Case $\langle \Gamma \rangle = \mathrm{iI}$*:* Note that by restricting the first argument of the minimal weak base relation R_{iI} to 0, we get the relation $\{0\} \times R$ with $R := \{000, 011, 101\}$. By [10] we have that $\mathsf{AnotherSAT}(R)$ is NP-complete. Now we proceed similarly to the first case, observing that the only solution m such that $m(w) = 1$ is the constant 1-assignment. □

Proposition 12. *For* $\Gamma \in \mathcal{L}$ *such that* $\langle \Gamma \rangle = \mathrm{iI}$ *or* $\mathrm{iN} \subseteq \langle \Gamma \rangle \subseteq \mathrm{iN}_2$ *and any* $\varepsilon > 0$ *there is no polynomial-time* $n^{1-\varepsilon}$*-approximation algorithm for* $\mathsf{NOSol}(\Gamma)$, *unless* P = NP.

Proof. Assume that there is a constant $\varepsilon > 0$ with a polynomial-time $n^{1-\varepsilon}$-approximation algorithm for $\mathsf{NOSol}(\Gamma)$. We will show how to use this algorithm to solve $\mathsf{AnotherSAT}_{<n}(\Gamma)$ in polynomial time. Proposition 11 completes the proof.

Let (φ, m) be an instance of $\mathsf{AnotherSAT}_{<n}(\Gamma)$ with n variables. If $n = 1$, then we reject the instance. Otherwise, we construct a new formula φ' and a new assignment m' as follows. Let k be the smallest integer greater than $1/\varepsilon$. Choose a variable x of φ and introduce $n^k - n$ new variables x^i for $i = 1, \ldots, n^k - n$. For every $i \in \{1, \ldots, n^k - n\}$ and every constraint $R(y_1, \ldots, y_\ell)$ in φ, such that $x \in \{y_1, \ldots, y_\ell\}$, construct a new constraint $R(z_1^i, \ldots, z_\ell^i)$ by $z_j^i = x^i$ if $y_j = x$ and $z_j^i = y_j$ otherwise; add all the newly constructed constraints to φ in order to get φ'. Moreover, we extend m to an assignment of φ' by setting $m'(x^i) = m(x)$. Now run the $n^{1-\varepsilon}$-approximation algorithm for $\mathsf{NOSol}(\Gamma)$ on (φ', m'). If the answer is $\overline{m'}$ then reject, otherwise accept.

We claim that the algorithm described above is a correct polynomial-time algorithm for the decision problem $\mathsf{AnotherSAT}_{<n}(\Gamma)$ when Γ is complementive. Polynomial runtime is clear. It remains to show its correctness. If the only solutions to φ are m and $\overline{m}$, then, as $n > 1$, the approximation algorithm must answer $\overline{m'}$ and the output is correct. Assume that there is a satisfying assignment m_s different from m and $\overline{m}$. The relation Γ is complementive, hence we may assume that $m_s(x) = m(x)$. It follows that φ' has a satisfying assignment m_s' for which $\mathrm{hd}(m_s', m') < n$ holds. But then the approximation algorithm must find a satisfying assignment m'' for φ' with $\mathrm{hd}(m', m'') < n \cdot (n^k)^{1-\varepsilon} = n^{k(1-\varepsilon)+1}$. Since the inequality $k > 1/\varepsilon$ holds, it follows that $\mathrm{hd}(m', m'') < n^k$. Consequently, m'' is not the complement of m' and the output of our algorithm is again correct.

When Γ is not complementive but both 0-valid and 1-valid ($\langle \Gamma \rangle = \mathrm{iI}$), we perform the expansion algorithm described above for each variable of the formula φ and reject if the result is the complement for each run. The runtime remains polynomial. □

MinDistance-Equivalent Cases. In this section we show that affine co-clones give rise to problems equivalent to MinDistance. The upper bound is easy.

Lemma 13. *For affine $\Gamma \in \mathcal{L}$ ($\Gamma \subseteq \mathrm{iL}_2$) the problem* $\mathsf{NOSol}(\Gamma)$ *reduces to* MinDistance.

Proof. Let the formula φ and the model m be an instance of $\mathsf{NOSol}(\Gamma)$ over the variables $x_1, \ldots, x_n$. Clearly, φ can be written as $A\boldsymbol{x} = \boldsymbol{b}$ and m is a solution of this affine system. As any solution of $A\boldsymbol{x} = \boldsymbol{b}$ can be written as $m' = m + m_0$ where m_0 is a solution of $A\boldsymbol{x} = \boldsymbol{0}$, the problem becomes equivalent to computing the solutions of this homogeneous system of small weight. But this is exactly the MinDistance problem. □

The following lemma can be easily proved, since the equivalence relation $[x \equiv y]$ is the solution set of the linear equation $x + y = 0$. The relation $[x]$ is represented by the equation $x = 1$ whereas the relation $[\neg x]$ is represented by $x = 0$.

Lemma 14. $\mathsf{NOSol}(\{\mathrm{even}^4\}) \equiv_{\mathrm{AP}} \mathsf{NOSol}(\{\mathrm{even}^4, [x], [\neg x]\})$.

Corollary 15. *For $\Gamma \in \mathcal{L}$ with* $\mathrm{iL} \subseteq \langle\Gamma\rangle \subseteq \mathrm{iL}_2$ *we have* MinDistance $\leq_{\mathrm{AP}}$ NOSol(Γ).

Proof. We show an AP-reduction to NOSol$(\{\mathrm{even}^4, [x], [\neg x]\})$. Since every system of linear equations can be written as a conjunction over relations in iL_2, the claim follows. □

MinHornDeletion-Equivalent Cases. As in Proposition 11 the need to use conjunctive closure instead of $\langle\ \rangle$ causes a case distinction in the proof of the following result.

Lemma 16. *If Γ is proper Horn* ($\mathrm{iE} \subseteq \langle\Gamma\rangle \subseteq \mathrm{iE}_2$) *then one of the following relations is in $\langle\Gamma\rangle_\wedge$: $[x \to y]$, $[x \to y] \times \{0\}$, $[x \to y] \times \{1\}$, or $[x \to y] \times \{01\}$.*

Proof. Supposing that $\langle\Gamma\rangle = \mathrm{iE}$, we get from Theorem 1 that R_{iE} belongs to $\langle\Gamma\rangle_\wedge$. Observe that $R_{\mathrm{iE}}(x_1, x_1, x_1, x_4) = [x_1 \to x_4]$ and thus $[x \to y] \in \langle R_{\mathrm{iE}}\rangle_\wedge \subseteq \langle\Gamma\rangle_\wedge$ which concludes this case. The case $\langle\Gamma\rangle = \mathrm{iE}_0$ leads to $[x \to y] \times \{0\} \in \langle\Gamma\rangle_\wedge$ in a completely analogous manner. The cases $\langle\Gamma\rangle = \mathrm{iE}_1$ and $\langle\Gamma\rangle = \mathrm{iE}_2$ lead to $[x \to y] \times \{1\} \in \langle\Gamma\rangle_\wedge$ and $[x \to y] \times \{01\} \in \langle\Gamma\rangle_\wedge$, respectively, by observing that $(x_1 \equiv x_1 \wedge x_3) = x_1 \to x_3$. □

Lemma 17. *If a constraint language $\Gamma \in \mathcal{L}$ is proper Horn* ($\mathrm{iE} \subseteq \langle\Gamma\rangle \subseteq \mathrm{iE}_2$), *then* NOSol$(\Gamma)$ *is* MinHornDeletion*-hard.*

Proof. Reduction from MinOnes$(\Gamma \cup \{[x]\})$ which is MinHornDeletion-hard by [11]. Consider first the case in which $[x \to y] \in \langle\Gamma\rangle_\wedge$. By Proposition 4 we may assume that $[x \to y] \in \Gamma$. Let φ be a $\Gamma \cup \{[x]\}$-formula. We construct φ' as follows. Replace each atomic formula $R(y_1, \ldots, y_k)$ in φ, where $R \in \Gamma$, by its conjunctive normal form decomposition, which yields a formula φ''. Since $R \in \Gamma \subseteq \mathrm{iE}_2$ holds, each clause occurring in this decomposition contains at most one unnegated variable. Those that contain negated variables are 0-valid, and so is their conjunction. The remaining ones, which are not 0-valid, are just single variables (literals). Next, replace all literals y from φ'' by $x \to y$, where x is a global new variable. Finally, add $v \to x$ for all variables v of φ to get φ'.

Observe that φ' is 0-valid. Moreover, the other solutions of φ' are exactly the solutions of φ extended by the assignment $x := 1$, because whenever one of the variables v takes the value 1, the clause $v \to x$ forces x to 1 which in turn enforces the unary clauses y of φ by the implications $x \to y$. It follows that $\mathrm{OPT}(\varphi) + 1 = \mathrm{OPT}(\varphi', \mathbf{0})$.

Moreover, for every r-approximate solution m' of φ' we first check whether $m = \mathbf{0}$ is a solution of φ. In case it is, $\mathrm{OPT}(\varphi) = 0$ and we trivially have $\mathrm{hw}(m) \leq 2r\mathrm{OPT}(\varphi)$. Otherwise, $\mathrm{OPT}(\varphi) \geq 1$ and we get a solution m of φ by restriction to the variables of φ with the weight $\mathrm{hw}(m) = \mathrm{hd}(\mathbf{0}, m') - 1 \leq r(\mathrm{OPT}(\varphi', \mathbf{0})) - 1 \leq r(\mathrm{OPT}(\varphi) + 1) - 1 \leq 2r\mathrm{OPT}(\varphi)$. In any case, we have thus $\mathrm{hw}(m) \leq 2r\mathrm{OPT}(\varphi)$ which shows that the construction is an AP-reduction with $\alpha = 2$.

For the other cases of Lemma 16 we argue similarly. The only difference is the introduction of some new variables, forced to constant values by the respective relation from Lemma 16. It is easy to see that these constants do not change the rest of the analysis. □

The proof of the following corollary requires a reduction to a similar problem, namely NearestSolution (NSol), which differs from NOSol in the point that the input assignment m does not need to satisfy the input formula φ; if it does, then m is the optimal solution (see [3] for details).

Corollary 18. *If* $\Gamma \in \mathcal{L}$ *is proper Horn* ($\mathrm{iE} \subseteq \langle \Gamma \rangle \subseteq \mathrm{iE}_2$) *or proper dual-Horn* ($\mathrm{iV} \subseteq \langle \Gamma \rangle \subseteq \mathrm{iV}_2$) *then* NOSol($\Gamma$) *is* MinHornDeletion*-complete under* AP*-Turing-reductions.*

Proof. Hardness follows from Lemma 17 and duality. Moreover, NOSol(Γ) can be AP-Turing-reduced to NSol($\Gamma \cup \{[x], [\neg x]\}$) as follows: Given a Γ-formula φ and a model m, we construct for every variable x of φ a formula $\varphi_x = \varphi \wedge (x = \overline{m}(x))$. Then for every x we run an oracle algorithm for NSol($\Gamma \cup \{[x], [\neg x]\}$) on (φ_x, m) and output one result of these oracle calls that is closest to m.

We claim that this algorithm is indeed an AP-Turing reduction. To see this observe first that the algorithm always computes a feasible solution, unless only m satisfies φ. Moreover, we have $\mathrm{OPT}(\varphi, m) = \min_x(\mathrm{OPT}(\varphi_x, m))$. Let $A(\varphi, m)$ be the answer of the algorithm on (φ, m) and let $B(\varphi_x, m)$ be the answers to the oracle calls. Consider a variable x^* such that $\mathrm{OPT}(\varphi, m) = \min_x(\mathrm{OPT}(\varphi_x, m)) = \mathrm{OPT}(\varphi_{x^*}, m)$, and assume that $B(\varphi_{x^*}, m)$ is an r-approximate solution of (φ_{x^*}, m). Then we get

$$\frac{\mathrm{hd}(m, A(\varphi, m))}{\mathrm{OPT}(\varphi, m)} = \frac{\min_y(\mathrm{hd}(m, B(\varphi_y, m))}{\mathrm{OPT}(\varphi_{x^*}, m)} \leq \frac{\mathrm{hd}(m, B(\varphi_{x^*}, m))}{\mathrm{OPT}(\varphi_{x^*}, m)} \leq r.$$

Thus the algorithm is indeed an AP-Turing-reduction from NOSol(Γ) to NSol($\Gamma \cup \{[x], [\neg x]\}$). Note that NSol($\Gamma \cup \{[x], [\neg x]\}$) reduces to MinHornDeletion (see [3]). Duality completes the proof. □

6 Concluding Remarks

The studied problem is in PO for bijunctive constraints. If the constraints are implication hitting set bounded by k for some $k \geq 2$, the problem NOSol still remains in PO. The situation is more complicated for Horn constraints and dual Horn constraints, where the task becomes equivalent to MinHornDeletion. The next complexity stage of the solution structure is characterized by affine constraints, where we can apply standard linear algebra techniques to prove equivalence with the MinDistance-problem. The penultimate stage of solution structure complexity is represented by constraints, for which the existence of a solution is guaranteed by their definition, but we do not have any other exploitable information. We need a guarantee of at least two solutions. The existence of a second solution is guaranteed by iN_2 being complementive. Our problem belongs to the class poly-APX for these constraints. We can even exactly pinpoint the polynomial (n, i.e. arity of the formula) for which we can get a polynomial-time approximation. This complexity result indicates that we cannot get a suitable approximation for these types of the considered optimization problem. All other cases cannot be approximated in polynomial time at all.

References

1. Ausiello, G., Crescenzi, P., Gambosi, G., Kann, V., Marchetti-Spaccamela, A., Protasi, M.: Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties. Springer-Verlag, Heidelberg (1999)
2. Baker, K.A., Pixley, A.F.: Polynomial interpolation and the Chinese Remainder Theorem for algebraic systems. Math. Z. **143**(2), 165–174 (1975)
3. Behrisch, M., Hermann, M., Mengel, S., Salzer, G.: Minimal distance of propositional models (2015). CoRR abs/1502.06761
4. Böhler, E., Creignou, N., Reith, S., Vollmer, H.: Playing with Boolean blocks, Part II: Constraint satisfaction problems. SIGACT News **35**(1), 22–35 (2004)
5. Böhler, E., Reith, S., Schnoor, H., Vollmer, H.: Bases for Boolean co-clones. Inf. Process. Lett. **96**(2), 59–66 (2005)
6. Creignou, N., Khanna, S., Sudan, M.: Complexity classifications of Boolean constraint satisfaction problems. In: SIAM Monographs on Discrete Mathematics and Applications, vol. 7. SIAM, Philadelphia (PA) (2001)
7. Creignou, N., Kolaitis, P.G., Zanuttini, B.: Structure identification of Boolean relations and plain bases for co-clones. J. Comput. Syst. Sci. **74**(7), 1103–1115 (2008)
8. Dumer, I., Micciancio, D., Sudan, M.: Hardness of approximating the minimum distance of a linear code. IEEE Trans. Inf. Theory **49**(1), 22–37 (2003)
9. Jeavons, P., Cohen, D., Gyssens, M.: Closure properties of constraints. J. Assoc. Comput. Mach. **44**(4), 527–548 (1997)
10. Juban, L.: Dichotomy theorem for the generalized unique satisfiability problem. In: Ciobanu, G., Păun, G. (eds.) FCT 1999. LNCS, vol. 1684, pp. 327–337. Springer, Heidelberg (1999)
11. Khanna, S., Sudan, M., Trevisan, L., Williamson, D.P.: The approximability of constraint satisfaction problems. SIAM J. Comput. **30**(6), 1863–1920 (2000)
12. Lagerkvist, V.: Weak bases of boolean co-clones. Inf. Process. Lette. **114**(9), 462–468 (2014)
13. Schaefer, T.J.: The complexity of satisfiability problems. In: Proceedings 10th Symposium on Theory of Computing (STOC 1978), San Diego (California, USA), pp. 216–226 (1978)
14. Schnoor, H., Schnoor, I.: Partial polymorphisms and constraint satisfaction problems. In: Creignou, N., Kolaitis, P.G., Vollmer, H. (eds.) Complexity of Constraints. LNCS, vol. 5250, pp. 229–254. Springer, Heidelberg (2008)