

Quantified Constraints and Containment Problems*

Hubie Chen

Departament de Tecnologia, Universitat Pompeu Fabra,
Barcelona, Spain.
hubie.chen@upf.edu

Florent Madelaine

Univ Clermont1, EA2146, Laboratoire d'algorithmique et d'image de Clermont-Ferrand,
Aubière, F-63170, France
fmadelaine@laic.u-clermont1.fr

Barnaby Martin

Department of Computer Science, University of Durham,
Durham DH1 3LE, UK.
b.d.martin@durham.ac.uk

Abstract

We study two containment problems related to the quantified constraint satisfaction problem (QCSP).

Firstly, we give a combinatorial condition on finite structures \mathcal{A} and \mathcal{B} that is necessary and sufficient to render $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$. The required condition is the existence of a positive integer r such that there is a surjective homomorphism from the power structure \mathcal{A}^r to \mathcal{B} . We note that this condition is already necessary to guarantee containment of the Π_2 restriction of QCSP, that is $\Pi_2\text{-CSP}(\mathcal{A}) \subseteq \Pi_2\text{-CSP}(\mathcal{B})$. Since we are able to give an effective bound on such an r , we provide a decision procedure for the model containment problem with non-deterministic double-exponential time complexity.

Secondly, we prove that the entailment problem for quantified conjunctive-positive first-order logic is decidable. That is, given two sentences φ and ψ of first-order logic with no instances of negation or disjunction, we give an algorithm that determines whether $\varphi \rightarrow \psi$ is true in all structures (models). Our result is in some sense tight, since we show that the entailment problem for positive first-order logic (i.e. quantified conjunctive-positive logic plus disjunction) is undecidable.

Introduction

The *constraint satisfaction problem* (CSP), much studied in artificial intelligence, is known to admit several equivalent formulations, two of the most popular of which are the model-checking problem for existential conjunctive-positive first-order (FO) sentences and the homomorphism problem (see, e.g., [13]). The CSP is NP-complete in general, and a great deal of effort has been expended in classifying its complexity for certain restricted cases, in particular where it is parameterised by the *constraint language* (which corresponds to the model in the model-checking problem and the right-hand structure of the homomorphism problem). The problems $\text{CSP}(\mathcal{A})$ thereby obtained, sometimes termed non-uniform [9], are conjectured [9, 3] to be always either polynomial-time tractable or NP-complete. While this has not been settled in general, a number of partial results are known (e.g. over structures of size ≤ 3 [18, 4] and over undirected graphs [10]).

The model containment problem for CSP is the question, for finite structures \mathcal{A} and \mathcal{B} , whether $\text{CSP}(\mathcal{A}) \subseteq \text{CSP}(\mathcal{B})$? It is easy to see that this is equivalent to the question of existence of a homomorphism from \mathcal{A} to \mathcal{B} . Thus the model containment problem for CSP is, essentially, a CSP itself. The condition for $\text{CSP}(\mathcal{A}) = \text{CSP}(\mathcal{B})$ is, therefore, that \mathcal{A} and \mathcal{B} are homomorphically equivalent. This in turn is equivalent to the condition that \mathcal{A} and \mathcal{B} share the same, or rather isomorphic, cores (where the core of a structure \mathcal{A} is a minimal substructure that is homomorphically equivalent to \mathcal{A}). The complexity classification problem for $\text{CSP}(\mathcal{A})$ is

*The latter two authors were supported by EPSRC grant EP/C54384X/1, most of the research being undertaken while the second author was at the University of Durham.

greatly facilitated by the fact that we may, therefore, assume that \mathcal{A} is a core – i.e. that \mathcal{A} is a minimal representative of its equivalence class under the equivalence relation induced by homomorphic equivalence.

A useful generalisation of the CSP involves considering the model-checking problem for conjunctive-positive FO sentences with both quantifiers permitted in the prefix. This allows for a broader class of problems, used in artificial intelligence to capture non-monotonic reasoning, whose complexities rise through the polynomial hierarchy up to Pspace. When the quantifier prefix is restricted to Π_2 , with all universal quantifiers preceding existential quantifiers, we obtain the Π_2 -CSP; when the prefix is unrestricted, we obtain the *quantified constraint satisfaction problem* (QCSP). In general, the Π_2 -CSP and QCSP are Π_2^P -complete and Pspace-complete, respectively (for more on these complexity classes, we direct the reader to [17]).

As with the CSP, it has become popular to consider the QCSP parameterised by the constraint language, i.e. the model in the model-checking problem, and there is a growing body of results delineating the tractable instances from those that are (probably) intractable [2, 5]. The model containment problem for QCSP takes as input two finite structures \mathcal{A} and \mathcal{B} and asks whether $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$. Unlike the situation with the CSP, it is not apparent that this containment problem is in any way similar to the QCSP itself. As far as we know, neither a characterisation nor an algorithm for this problem had been known. In this paper we provide both, i.e. we settle the question as to when exactly $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$ by giving a characterising morphism from \mathcal{A} to \mathcal{B} . It turns out that $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$ exactly when there exists a positive integer r s.t. there is a surjective homomorphism from the power structure \mathcal{A}^r to \mathcal{B} . We note that this condition is already necessary to guarantee containment of $\Pi_2\text{-CSP}(\mathcal{A}) \subseteq \Pi_2\text{-CSP}(\mathcal{B})$. If the sizes of the structures \mathcal{A} and \mathcal{B} are $|\mathcal{A}|$ and $|\mathcal{B}|$, respectively, then we may take $r := |\mathcal{A}|^{|\mathcal{B}|}$. Thus to decide whether $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$, it suffices to verify whether or not there is a surjective homomorphism from $\mathcal{A}^{|\mathcal{A}|^{|\mathcal{B}|}}$ to \mathcal{B} . This provides a decision procedure for the model containment problem with non-deterministic double-exponential time complexity.

The Classical Decision Problem, known also as Hilbert’s *Entscheidungsproblem*, is the question, given a FO sentence φ , whether φ is true in all models (is logically valid) or, dually, is true in some model (is satisfiable). It is well-known that this problem is undecidable in general. The entailment problem for FO asks, given sentences φ and ψ , whether we have the logical validity of $\varphi \rightarrow \psi$ (denoted $\models \varphi \rightarrow \psi$). The equivalence problem is defined similarly, with \rightarrow substituted by \leftrightarrow . Both problems are easily seen to be equivalent to the Classical Decision Problem, and are therefore undecidable. A great literature exists on decidable and un-

decidable cases of the Classical Decision Problem, particularly under restrictions of quantifier prefixes and (arity and number of) relation and function symbols – see the monograph [1]. However, for certain natural fragments of FO, it seems the entailment and equivalence problems are not well-studied. The query containment problem is closely related to the entailment problem, but with truth in all finite models substituted for truth in all models. Query containment problems are fundamental to many aspects of database systems, including query optimisation, determining independence of queries and rewriting queries using views. The query containment problem for FO is also undecidable.

The sentence containment problem for the CSP – a.k.a. the query containment problem for existential conjunctive-positive FO – is the question, given existential conjunctive-positive sentences φ and ψ , whether, for all finite structures \mathcal{A} , $\mathcal{A} \models \varphi$ implies $\mathcal{A} \models \psi$ (i.e. $\models_{\text{FIN}} \varphi \rightarrow \psi$). It is easily seen that this problem is decidable and NP-complete, in fact it is an instance of the homomorphism problem (equivalently, the CSP itself). It is also easy to demonstrate, in this case, that the condition of finiteness is irrelevant. That is, $\models_{\text{FIN}} \varphi \rightarrow \psi$ if, and only if, $\models \varphi \rightarrow \psi$. Thus we have here the decidability and NP-completeness of the entailment problem for existential conjunctive-positive FO logic.

The second part of this paper is motivated by the sentence containment problem for the QCSP – a.k.a. the query containment problem for quantified conjunctive-positive FO – that is, given quantified conjunctive-positive sentences φ and ψ , to determine whether $\models_{\text{FIN}} \varphi \rightarrow \psi$. In this case it is not clear as to whether this coincides with the condition of entailment, $\models \varphi \rightarrow \psi$. Our principle contribution is to give a decision procedure, with triple-exponential time complexity, for the entailment problem, i.e. the problem to determine, for two quantified conjunctive-positive FO sentences φ and ψ , whether $\models \varphi \rightarrow \psi$. Since existential conjunctive-positive sentences are quantified conjunctive-positive, it follows from the comments of the previous paragraph that this entailment problem is NP-hard.

We will make particular use of a certain canonical model for the sentence φ , built on the Herbrand universe of terms derived from Skolem functions over a countably infinite set of (new) constants. Herbrand models are commonplace in algorithmic results on logical validity and equivalence in both first-order logic (e.g. [14]) and logic programming (e.g. [15, 7, 8]). However, our method differs significantly from those in the citations.

We also prove that the related entailment problem for positive FO – even without equality – is undecidable. Since the difference between quantified conjunctive-positive FO and positive FO is simply the addition of disjunction, we suggest that our decidability result is somehow tight.

Related work. Students of the algebraic method will appreciate the aesthetic of our model containment result, which appears to mirror that of the relationship between quantified conjunctive-positive FO and surjective polymorphisms. Surjective polymorphisms are nothing but surjective homomorphisms from a power of a structure to itself. Let $\text{sur-pol}(\mathcal{A})$ and $\text{sur-pol}(\mathcal{B})$ be the set of surjective polymorphisms of finite structures \mathcal{A} and \mathcal{B} , respectively. It was proved in [2] that, if $\text{sur-pol}(\mathcal{B}) \subseteq \text{sur-pol}(\mathcal{A})$, then $\text{QCSP}(\mathcal{A}) \leq_P \text{QCSP}(\mathcal{B})$ (where \leq_P denotes polynomial-time reduction).

We do not wish to define a quorum of algebraic notions, but instead quote the following result, which, although unpublished in this form, is more or less known in the community. For a finite algebra \mathbb{A} , let $\text{inv}(\mathbb{A})$ be the set of relations (on the domain of \mathbb{A}) that are invariant under the operations of \mathbb{A} . For finite algebras \mathbb{A} and \mathbb{B} , if there exists an r s.t. there is a surjective homomorphism from \mathbb{A}^r to \mathbb{B} , then $\text{QCSP}(\text{inv}(\mathbb{B})) \leq_P \text{QCSP}(\text{inv}(\mathbb{A}))$. Compare this with our result which states that there exists an r s.t. there is a surjective homomorphism from \mathcal{A}^r to \mathcal{B} iff $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$. Aside from the right-hand sides being inverted (which one might expect from the duality of algebras and relations), these results are somewhat similar. However, our result is tight, i.e. holds in converse, and relies on set inclusion and not computational reduction.

For a structure \mathcal{A} , let $\text{rel}(\mathcal{A})$ be the set of relations definable on \mathcal{A} in quantified conjunctive-positive FO. Let $\Pi_2\text{-rel}(\mathcal{A})$ be that subset of relations that are already definable in the Π_2 fragment. It follows from [2] (although see [6] for details) that, for all \mathcal{A} , $\Pi_2\text{-rel}(\mathcal{A})$ and $\text{rel}(\mathcal{A})$ actually coincide. Although this is not the same as our observation that $\Pi_2\text{-CSP}(\mathcal{A}) \subseteq \Pi_2\text{-CSP}(\mathcal{B})$ iff $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$, the reader may once again appreciate a similar form.

Finally, we mention a result of classical model theory due to Keisler. In [12] a result of considerable generality appears whose projection onto our domain of discourse yields that $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$ iff there is a surjective homomorphism from \mathcal{A}^ω to \mathcal{B} (where ω is the set of natural numbers). Keisler's result goes beyond the situation in which \mathcal{A} and \mathcal{B} are finite (although then the power may be higher than ω), but our work may be seen as providing an effective finite bound on the power where this is the case.

Organisation of the paper. Having introduced some global preliminaries, the main body of the paper sits in two parts.

The first part, Section 1, covers the model containment problem for QCSP. After introducing the basic concepts involved in this result, we state and prove our characterisation in Section 1.1 and consider the properties of the ensuing algorithm.

The second part, Section 2, covers the entailment prob-

lem for quantified conjunctive-positive FO. Through the preliminaries, we introduce the canonical model for a quantified conjunctive-positive sentence. In Section 2.1, we give a methodology theorem that is the basis of our algorithm, and establish the remaining necessary machinery for our result. Section 2.2 details the complexity of our algorithm and Section 2.3 gives the undecidability of entailment for positive FO.

We conclude the paper with a section of final remarks.

Global Preliminaries

Throughout, let σ be a fixed, finite relational signature. If \mathcal{B} is a σ -structure, then its domain is denoted B and the cardinality of that domain $|B|$. The stipulation that σ contains no constants is purely for technical convenience, as we will occasionally wish to consider structures over the expanded signature $\sigma \cup C_\alpha$, where C_α is a set of α constant symbols.

A *homomorphism* from \mathcal{A} to \mathcal{B} is a function $h : A \rightarrow B$ that preserves positive relations. That is, if R is a p -ary relation symbol of σ , if $R(x_1, \dots, x_p) \in \mathcal{A}$ then $R(h(x_1), \dots, h(x_p)) \in \mathcal{B}$. Existence of a homomorphism (resp., surjective homomorphism) from \mathcal{A} to \mathcal{B} is denoted $\mathcal{A} \rightarrow \mathcal{B}$ (resp., $\mathcal{A} \twoheadrightarrow \mathcal{B}$). If both $\mathcal{A} \rightarrow \mathcal{B}$ and $\mathcal{B} \rightarrow \mathcal{A}$, then we describe \mathcal{A} and \mathcal{B} as *homomorphically equivalent*. If $f : A \rightarrow B$ is a function, and $A' \subseteq A$ then we denote by $\text{im}(A')$ the image of A' under f (i.e. $\{f(x) : x \in A'\}$).

A FO sentence φ is *quantified conjunctive-positive* if it contains no instances of negation or disjunction. It is clear that such a sentence may be put in the prenex normal form

$$\varphi := \forall \mathbf{x}_1 \exists \mathbf{y}_1 \dots \forall \mathbf{x}_k \exists \mathbf{y}_k P(\mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{x}_k, \mathbf{y}_k),$$

where P is a conjunction of positive atoms. If φ contains only variables \mathbf{x}_1 and \mathbf{y}_1 (i.e. one quantifier alternation) then it is said to be Π_2 ; if φ contains only (the existential) variables \mathbf{x}_1 then it is said to be Σ_1 . The *quantified constraint satisfaction problem* $\text{QCSP}(\mathcal{A})$ has

- Input: a quantified conjunctive-positive sentence φ .
- Question: does $\mathcal{A} \models \varphi$?

If φ is restricted to being Π_2 (resp., Σ_1) then the resulting problem is $\Pi_2\text{-CSP}(\mathcal{A})$ (resp., $\text{CSP}(\mathcal{A})$). The model containment problem for QCSP takes as input two finite structures \mathcal{A} and \mathcal{B} , and has as its yes-instances those pairs for which $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$. The model containment problem for CSP and $\Pi_2\text{-CSP}$ is defined analogously.

Let φ be a sentence of the form $\forall \mathbf{x}_1 \exists \mathbf{y}_1 \dots \forall \mathbf{x}_k \exists \mathbf{y}_k P(\mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{x}_k, \mathbf{y}_k)$, and let \mathcal{A} be a finite structure. Rather sloppily, we will identify a variable tuple \mathbf{x} with its underlying set of variables. The φ -game on \mathcal{A} is a two-player game that pitches *Universal* (male) against *Existential* (female). The game goes as follows. For $1 \leq i \leq k$ ascending:

- for every variable in \mathbf{x}_i , Universal chooses an element in A : i.e. he gives a function $f_{\forall_i} : \mathbf{x}_i \rightarrow A$; and,
- for every variable in \mathbf{y}_i , Existential chooses an element in A : i.e. she gives a function $f_{\exists_i} : \mathbf{y}_i \rightarrow A$.

Existential wins if, and only if,

$$\mathcal{A} \models P(f_{\forall_1}(\mathbf{x}_1), f_{\exists_1}(\mathbf{y}_1), \dots, f_{\forall_k}(\mathbf{x}_k), f_{\exists_k}(\mathbf{y}_k)),$$

where $f(\mathbf{x})$ is the natural pointwise action of f on the coordinates of \mathbf{x} .

A *strategy* $\varepsilon := (\varepsilon_1, \dots, \varepsilon_k)$ for Existential (resp., $v := (v_1, \dots, v_k)$ for Universal) tells her (resp., him) how to play a variable tuple given what has been played before. That is, ε_l is a function from $A^{(\mathbf{x}_1 \cup \mathbf{y}_1 \cup \dots \cup \mathbf{x}_{l-1})} \times \mathbf{y}_l$ to A and v_l is a function from $A^{(\mathbf{x}_1 \cup \mathbf{y}_1 \cup \dots \cup \mathbf{x}_{l-1} \cup \mathbf{y}_{l-1})} \times \mathbf{x}_l$ to A (note that $A^{(\mathbf{x}_1 \cup \mathbf{y}_1 \cup \dots \cup \mathbf{x}_{l-1})}$ and $A^{(\mathbf{x}_1 \cup \mathbf{y}_1 \cup \dots \cup \mathbf{x}_{l-1} \cup \mathbf{y}_{l-1})}$ are themselves functions specifying how the game was played on the previous variable tuples). A strategy for Existential is *winning* if it beats all possible strategies of Universal. The φ -game on \mathcal{A} is nothing other than a model-checking (Hinitikka) game, and it is straightforward to verify that Existential has a winning strategy if, and only if, $\mathcal{A} \models \varphi$.

1 The Model Containment Problem


For the set $C_m := \{c_1, \dots, c_m\}$ of constant symbols, we denote structures \mathfrak{A} over the signature $\sigma \cup C_m$ in Fraktur, whereupon \mathcal{A} denotes the restriction of \mathfrak{A} to σ , in the obvious way. Of course, a homomorphism $h : \mathfrak{A} \rightarrow \mathfrak{B}$ must also preserve the constants, i.e. if $c_i := x$ in \mathfrak{A} then $c_i := h(x)$ in \mathfrak{B} . Given two $\sigma \cup C_m$ -structures \mathfrak{A} and \mathfrak{B} , we define their (categorical) *product* $\mathfrak{A} \otimes \mathfrak{B}$ to have domain $A \times B$ and relations $R_i((x_1, y_1), \dots, (x_{a_i}, y_{a_i}))$ iff $R_i(x_1, \dots, x_{a_i}) \in \mathfrak{A}$ and $R_i(y_1, \dots, y_{a_i}) \in \mathfrak{B}$. The constant c_i in $\mathfrak{A} \otimes \mathfrak{B}$ is the element (x_i, y_i) s.t. $c_i := x_i \in \mathfrak{A}$ and $c_i := y_i \in \mathfrak{B}$.

The following lemma is a restricted version of the well-known fact that surjective homomorphisms preserve positive formulae (see, e.g., [11]).

Lemma 1. *If $\mathfrak{A} \rightarrow \mathfrak{B}$ then $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$.*

Sketch Proof. If $s : \mathfrak{A} \rightarrow \mathfrak{B}$ is a surjective homomorphism, then let $s^{-1} : \mathfrak{B} \rightarrow \mathfrak{A}$ be s.t. $s \circ s^{-1}$ is the identity on \mathfrak{B} . Let φ be of the form $\forall \mathbf{x}_1 \exists \mathbf{y}_1 \dots \forall \mathbf{x}_k \exists \mathbf{y}_k P(\mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{x}_k, \mathbf{y}_k)$. Given a winning strategy ε for Existential in the φ -game on \mathfrak{A} , we build a winning strategy ε' for her in the φ -game on \mathfrak{B} , whereupon the result follows. For $1 \leq i \leq k$, let g be a mapping from $(\mathbf{x}_1 \cup \mathbf{y}_1 \cup \dots \cup \mathbf{x}_{i-1})$ to \mathfrak{B} and let y be a variable of \mathbf{y}_i . We set $\varepsilon'_i(g, y) := s \circ \varepsilon_i(s^{-1} \circ g, y)$. \square

Example 1. Consider the graphs drawn below. Both \mathcal{H}_1 and \mathcal{H}_2 have a surjective homomorphism to \mathcal{K}_3 ; therefore we can derive both $\text{QCSP}(\mathcal{H}_1) \subseteq \text{QCSP}(\mathcal{K}_3)$ and $\text{QCSP}(\mathcal{H}_2) \subseteq \text{QCSP}(\mathcal{K}_3)$.



1.1 Characterisation

As the following shows, perhaps surprisingly, the model containment of QCSP is already determined by the model containment of Π_2 -CSP, the restriction of QCSP to Π_2 sentences.

Theorem 1. *Let \mathcal{A} and \mathcal{B} be σ -structures. The following are equivalent.*

- (i) $\mathcal{A}^{|\mathcal{A}|^{|\mathcal{B}|}} \rightarrow \mathcal{B}$.
- (ii) There exists r s.t. $\mathcal{A}^r \rightarrow \mathcal{B}$.
- (iii) $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$.
- (iv) $\Pi_2\text{-CSP}(\mathcal{A}) \subseteq \Pi_2\text{-CSP}(\mathcal{B})$.

The proof of this result relies on two novel constructions. Firstly, we show how to build a *product strategy* for the φ -game on \mathcal{A}^r from a strategy for the φ -game on \mathcal{A} in Section 1.1.1. Secondly, we build *indicator structures* for the containment of Π_2 -CSP in Section 1.1.2.

1.1.1 Product and winning strategy

In this section we will show that a structure and any of its powers share the same QCSP. In the following, let \mathcal{A} be a structure and let $r \geq 1$. Let φ be a sentence of the form $\forall \mathbf{x}_1 \exists \mathbf{y}_1 \dots \forall \mathbf{x}_k \exists \mathbf{y}_k P(\mathbf{x}_1, \mathbf{y}_1, \dots, \mathbf{x}_k, \mathbf{y}_k)$ and let ε be a strategy for Existential in the φ -game on \mathcal{A} . The *product strategy* ε^r for Existential in the φ -game on \mathcal{A}^r is defined as follows. For $1 \leq i \leq k$, let g be a mapping from $(\mathbf{x}_1 \cup \mathbf{y}_1 \cup \dots \cup \mathbf{x}_{i-1})$ to \mathcal{A}^r and let y be a variable of \mathbf{y}_i . We set $\varepsilon^r_i(g, y) := (\varepsilon_i(\text{pr}_1 \circ g, y), \dots, \varepsilon_i(\text{pr}_r \circ g, y))$, where $\text{pr}_1, \dots, \text{pr}_r$ denote the natural projections from \mathcal{A}^r to \mathcal{A} .

Lemma 2. $\text{QCSP}(\mathcal{A}) = \text{QCSP}(\mathcal{A}^r)$.

Sketch Proof. The backward containment follows from Lemma 1, since $\mathcal{A}^r \rightarrow \mathcal{A}$. For the forward containment, if ε is a winning strategy for Existential in the φ -game on \mathcal{A} then ε^r is a winning strategy for her in the φ -game on \mathcal{A}^r . The result follows. \square


1.1.2 Indicator structure for Π_2 -CSP

Recall the signature $\sigma \cup C_m$, where $C_m := \{c_1, \dots, c_m\}$. We will associate C_m with $[m] := \{1, \dots, m\}$, in the natural way. Given a mapping λ from $[m]$ to a structure \mathcal{A} , we write \mathfrak{A}_λ for the $\sigma \cup C_m$ -structure induced naturally by \mathcal{A} and the interpretation of the constant symbols given by λ . Let $A^{[m]}$ denote the set of all possible interpretations. We call *indicator structure* the $\sigma \cup C_m$ -structure $\mathfrak{A}^{[A]^m} := \bigotimes_{\lambda \in A^{[m]}} \mathfrak{A}_\lambda$ (note that this is well-defined since \bigotimes is associative and commutative, up to isomorphism).

There is a natural correspondence between Π_2 quantified conjunctive-positive sentences φ with m universally quantified variables and $\sigma \cup C_m$ -structures. Recall φ is of the form $\forall \mathbf{x}_1 \exists \mathbf{y}_1 P(\mathbf{x}_1, \mathbf{y}_1)$, where $\mathbf{x}_1 := (x_1^1, \dots, x_1^m)$. From φ , we build the $\sigma \cup C_m$ -structure \mathfrak{D}_φ in the following way. The elements of \mathfrak{D}_φ are the variables of φ , and the relation tuples of \mathfrak{D}_φ are exactly the facts of the conjunction $P(\mathbf{x}_1, \mathbf{y}_1)$ (indeed if all the quantifiers of φ were switched to being existential then one would obtain the so-called canonical query – see [13] – of the structure \mathfrak{D}_φ , the restriction of \mathfrak{D}_φ to σ). Finally, the elements x_1^1, \dots, x_1^m interpret the constants c_1, \dots, c_m . Conversely, given a $\sigma \cup C_m$ -structure \mathfrak{D} , we build the Π_2 quantified conjunctive-positive sentence $\varphi_{\mathfrak{D}}$ as follows. The variables of $\varphi_{\mathfrak{D}}$ are the elements of \mathfrak{D} , and the quantifier-free part of $\varphi_{\mathfrak{D}}$ is the conjunction of the facts of \mathfrak{D} . Finally, the variables (whose elements interpreted the constants) c_1, \dots, c_m are universally quantified, while all other variables are existentially quantified (to the inside of the universal quantification). This correspondence is essentially bijective, and is illustrated in the following example.

Example 2. $\varphi := \forall x_1^1, x_1^2, x_1^3 \exists y_1^1, y_1^2, y_1^3, y_1^4 E(y_1^1, x_1^1) \wedge E(x_1^1, y_1^2) \wedge E(x_1^1, y_1^3) \wedge E(y_1^2, x_1^2) \wedge E(y_1^3, x_1^3) \wedge E(x_1^2, y_1^4) \wedge E(x_1^3, y_1^4)$.

The sentence φ , depicted on the left, gives rise to the $\sigma \cup C_3$ -structure \mathfrak{D}_φ , depicted on the right.



Theorem 2 (Methodology I). *Let φ be of the form $\forall \mathbf{x}_1 \exists \mathbf{y}_1 P(\mathbf{x}_1, \mathbf{y}_1)$, where P is a conjunction of positive atoms and $\mathbf{x}_1 := (x_1^1, \dots, x_1^m)$. Let \mathfrak{D}_φ be φ 's corresponding $\sigma \cup C_m$ -structure. The following are equivalent:*

- (i) $\mathcal{A} \models \varphi$
- (ii) $\mathfrak{D}_\varphi \longrightarrow \mathfrak{A}^{[A]^m}$

Proof. $\mathcal{A} \models \varphi$ iff for every mapping f_{\forall_1} from \mathbf{x}_1 to \mathcal{A} , there exists a mapping f_{\exists_1} from \mathbf{y}_1 to \mathcal{A} such that $\mathcal{A} \models P(f_{\forall_1}(\mathbf{x}_1), f_{\exists_1}(\mathbf{y}_1))$. From the definition, this is equivalent to there existing a homomorphism from \mathfrak{D}_φ to \mathfrak{A}_λ , for

every $\lambda \in A^{[m]}$ (indeed, when λ coincides with f_{\forall_1} , under the natural substitution of the domain $[m]$ by (x_1^1, \dots, x_1^m) , then $f_{\forall_1} \cup f_{\exists_1}$ provides the homomorphism). By construction of $\mathfrak{A}^{[A]^m}$ as a product of such \mathfrak{A}_λ , we have equivalently that there exists a homomorphism from \mathfrak{D} to $\mathfrak{A}^{[A]^m}$. \square

We will shortly prove an extension of this result that relates model containment of Π_2 -CSP with homomorphism between indicators. First, we need the following technical lemma, which allows us to restrict ourselves to Π_2 quantified conjunctive-positive sentences with a bounded number of universal variables.

Lemma 3. *Let \mathcal{A} and \mathcal{B} be two σ -structures. Π_2 -CSP(\mathcal{A}) \subseteq Π_2 -CSP(\mathcal{B}) if, and only if, for every Π_2 quantified conjunctive-positive sentence φ with at most $|B|$ universal variables, $\mathcal{A} \models \varphi$ implies $\mathcal{B} \models \varphi$.*

Proof. The forward direction is trivial; we prove the backward direction by contraposition. Suppose that Π_2 -CSP(\mathcal{A}) $\not\subseteq$ Π_2 -CSP(\mathcal{B}), i.e. there is a φ s.t. $\mathcal{A} \models \varphi$ but $\mathcal{B} \not\models \varphi$. Let $\mathbf{x}_1 := (x_1^1, \dots, x_1^m)$ be the universal variables of φ and let $B^{[m]}$ be the set of mappings μ from $[m]$ to B . For each such mapping μ , let $\mu(\varphi)$ be the Π_2 quantified conjunctive-positive sentence obtained from φ by identifying variables of \mathbf{x}_1 that share the same image under μ (more precisely, μ under the natural substitution of the domain $[m]$ by (x_1^1, \dots, x_1^m)). Note that $\mu(\varphi)$ has at most $|B|$ universally quantified variables. For any structure \mathcal{C} , $\mathcal{C} \models \varphi$ implies, for all $\mu \in B^{[m]}$ that $\mathcal{C} \models \mu(\varphi)$. Furthermore, if $|C| \leq |B|$, the converse implication also holds, since every play by Universal in the φ -game on \mathcal{C} can be cast, for some mapping μ , as a play of the $\mu(\varphi)$ -game on \mathcal{C} . Since $\mathcal{B} \not\models \varphi$, we may deduce a μ_0 s.t. $\mathcal{B} \not\models \mu_0(\varphi)$. However, since $\mathcal{A} \models \mu_0(\varphi)$, the result follows. \square

For σ -structures \mathcal{A} and \mathcal{B} , the indicator $\sigma \cup C_{|B|}$ -structure $\mathfrak{A}^{[A]^{|B|}}$ will play a particular role in our proof. Note that its restriction to σ is exactly the structure $\mathcal{A}^{[A]^{|B|}}$.

Theorem 3 (Indicator). *Let \mathcal{A} and \mathcal{B} be two σ -structures. The following are equivalent.*

- (i) Π_2 -CSP(\mathcal{A}) \subseteq Π_2 -CSP(\mathcal{B})
- (ii) $\mathfrak{A}^{[A]^{|B|}} \longrightarrow \mathfrak{B}^{[B]^{|B|}}$

Proof. (Downwards). Since trivially $\mathfrak{A}^{[A]^{|B|}} \longrightarrow \mathfrak{A}^{[A]^{|B|}}$, by Theorem 2, $\mathcal{A} \models \varphi_{\mathfrak{A}^{[A]^{|B|}}}$. Thus, from our hypothesis, it follows that $\mathcal{B} \models \varphi_{\mathfrak{A}^{[A]^{|B|}}}$. Applying Theorem 2 in the other direction, we get that $\mathfrak{A}^{[A]^{|B|}} \longrightarrow \mathfrak{B}^{[B]^{|B|}}$, as desired.

(Upwards). By the previous lemma, it suffices to consider Π_2 quantified conjunctive-positive sentences with at most $|B|$ universally quantified variables. Let φ be such a sentence. Applying Theorem 2 we have that $\mathcal{A} \models \varphi$ implies $\mathfrak{D}_\varphi \longrightarrow \mathfrak{A}^{[A]^{|B|}}$. By composition, since $\mathfrak{A}^{[A]^{|B|}} \longrightarrow$




Figure 1. Depiction of the proof of Lemma 3 using the sentence of Example 2, in the case $b = 2$. Note that the sentence has 3 universal variables, i.e. $c = 3$. The three new sentences result from identifying x_1^1 and x_1^2 , x_1^2 and x_1^3 and x_1^1 and x_1^3 , respectively.

$\mathfrak{B}^{|B|^{B|}}$, we have that $\mathfrak{D}_\varphi \rightarrow \mathfrak{B}^{|B|^{B|}}$. Hence, applying Theorem 2 in the converse direction, we get that $\mathcal{B} \models \varphi$. \square

We are now in a position to prove **Theorem 1**, whose statement we reiterate for the benefit of the reader. Let \mathcal{A} and \mathcal{B} be σ -structures; the following are equivalent.

- (i) $\mathcal{A}^{|A|^{B|}} \twoheadrightarrow \mathcal{B}$.
- (ii) There exists r s.t. $\mathcal{A}^r \twoheadrightarrow \mathcal{B}$.
- (iii) $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$.
- (iv) $\Pi_2\text{-CSP}(\mathcal{A}) \subseteq \Pi_2\text{-CSP}(\mathcal{B})$.

Proof of Theorem 1. (i) \Rightarrow (ii) is trivial. (ii) \Rightarrow (iii) follows from Lemmas 1 and 2. (iii) \Rightarrow (iv) is trivial. Finally, (iv) \Rightarrow (i) follows from the Indicator Theorem, in the following way. Let us assume $\Pi_2\text{-CSP}(\mathcal{A}) \subseteq \Pi_2\text{-CSP}(\mathcal{B})$ and, consequently, $\mathfrak{A}^{|A|^{B|}} \twoheadrightarrow \mathfrak{B}^{|B|^{B|}}$. By construction of the indicator, this implies that for each μ in $B^{|B|}$, we have $\mathfrak{A}^{|A|^{B|}} \twoheadrightarrow \mathfrak{B}_\mu$. By choosing a μ_0 that is surjective, we derive $\mathfrak{A}^{|A|^{B|}} \twoheadrightarrow \mathfrak{B}_{\mu_0}$ and, by forgetting the constant symbols, the result follows. \square

The Indicator Theorem is interesting because it allows us to relate $\Pi_2\text{-CSP}$ model containment, and through Theorem 1, QCSP model containment with ordinary homomorphisms. In fact we could have bypassed that observation and taken a more direct route to the proof of the outstanding case of Theorem 1, (iv) \Rightarrow (i). By contraposition, suppose that $\mathcal{A}^{|A|^{B|}} \not\rightarrow \mathcal{B}$. It can be shown that the Π_2 sentence $\varphi_{\mathfrak{A}^{|A|^{B|}}}$ separates $\Pi_2\text{-CSP}(\mathcal{A})$ and $\Pi_2\text{-CSP}(\mathcal{B})$, i.e. $\mathcal{A} \models \varphi_{\mathfrak{A}^{|A|^{B|}}}$ but $\mathcal{B} \not\models \varphi_{\mathfrak{A}^{|A|^{B|}}}$.

Example 3. Consider, again, the graphs of Example 1. It can be shown that, for each r , $\mathcal{K}_3^r \not\rightarrow \mathcal{H}_1$, while $\mathcal{K}_3^2 \rightarrow \mathcal{H}_2$. It follows that $\text{QCSP}(\mathcal{K}_3) = \text{QCSP}(\mathcal{H}_2)$. In fact, \mathcal{K}_3 and \mathcal{H}_2 not only agree on all sentences of quantified conjunctive-positive FO, but actually on all sentences of FO in which equality does not appear [16].

1.1.3 Complexity

Having established a combinatorial characterisation for the QCSP model containment problem, we make the following

observation as to its complexity – as can be seen the twin bounds are far from tight.

Theorem 4. *The model containment problem for QCSP, that is the problem which, given finite structures \mathcal{A} and \mathcal{B} , decides whether $\text{QCSP}(\mathcal{A}) \subseteq \text{QCSP}(\mathcal{B})$ is 1.) in nondeterministic double-exponential time, and 2.) is NP-hard (under polynomial-time reductions).*

Proof. Membership of nondeterministic double-exponential time follows from Theorem 1 by building $\mathcal{A}^{|A|^{B|}}$ and guessing a surjective homomorphism to \mathcal{B} (which can easily be verified as such in double-exponential time). NP-hardness follows by a reduction from the problem *graph 3-colourability*, as we will demonstrate.

Let \mathcal{K}_1 and \mathcal{K}_3 be the (antireflexive) 1- and 3-clique, respectively. That is, \mathcal{K}_1 is a single loopless vertex and \mathcal{K}_3 is the triangle. For graphs \mathcal{G} and \mathcal{H} , let $\mathcal{G} \uplus \mathcal{H}$ be the disjoint union of \mathcal{G} and \mathcal{H} . Let $3.\mathcal{K}_1$ be $\mathcal{K}_1 \uplus \mathcal{K}_1 \uplus \mathcal{K}_1$. It is well-known that \mathcal{G} is 3-colourable iff $\mathcal{G} \rightarrow \mathcal{K}_3$. It is easy to see that this is equivalent to $(\mathcal{G} \uplus 3.\mathcal{K}_1) \rightarrow \mathcal{K}_3$. It is relatively straightforward now to see that this is equivalent to the existence of an r s.t. $(\mathcal{G} \uplus 3.\mathcal{K}_1)^r \rightarrow \mathcal{K}_3$. The result now follows from Theorem 1. \square

2 The Entailment Problem

For a simpler exposition, we will assume throughout this section that all quantified conjunctive-positive sentences have strict quantifier alternation, i.e. are of the form

$$\varphi := \forall x_1 \exists y_1 \dots \forall x_k \exists y_k P(x_1, y_1, \dots, x_k, y_k),$$

where P is a conjunction of positive atoms. Of course, any quantified conjunctive-positive sentence may be readily converted to an equivalent sentence in this form by the introduction of dummy variables. If P contains any atomic instance $x_i = x_j$ ($i \neq j$) or $y_i = x_j$ ($i < j$) then we describe φ as *degenerate*. It is clear that all models of a degenerate φ are of cardinality 1, and that there is a finite set of normalised σ -structures over the domain $\{1\}$. It follows that, if φ is degenerate, we may establish directly whether $\models \varphi \rightarrow \psi$ by evaluating ψ over all normalised models of φ .

Note that instances of equality in a non-degenerate φ may be propagated out by substitution. In order to answer the question $\models \varphi \rightarrow \psi$ in general, we will wish to build a canonical model of φ . Henceforth, we will assume that φ (but not necessarily ψ) contains no instances of equality.

The Canonical Model

Let φ be a quantified conjunctive-positive sentence of the form $\forall x_1 \exists y_1 \dots \forall x_k \exists y_k P(x_1, y_1, \dots, x_k, y_k)$. We consider k to be the *depth* of φ , denoted $\text{depth}(\varphi)$. We wish to build a *canonical model* of φ , and we shall do this via its Skolem normal form. Let $F := \{f_1, \dots, f_k\}$ be a set of function symbols, in which the arity of f_i is i . Let $\text{Skolem}(\varphi) :=$

$$\forall x_1 \dots \forall x_k P(x_1, f_1(x_1), \dots, x_k, f_k(x_1, \dots, x_k)),$$

be the derivative sentence over the signature $\sigma \cup F$. Each atom of P induces what we designate a *quantified atom* in $\text{Skolem}(\varphi)$. It is well-known that the models of φ and $\text{Skolem}(\varphi)$ are intimately related, indeed they are identical up to the additional interpretation of the new function symbols of F .

If α is a positive integer, let $C_\alpha := \{c_1, \dots, c_\alpha\}$; if $\alpha := \omega$, let $C_\alpha := \{c_1, \dots\}$. Define $T_\varphi(C_\alpha)$ to be the set of (closed) *terms* obtained from all compositions of the functions of F on themselves and on the constants of C_α . The *rank* of a term $t \in T_\varphi(C_\alpha)$, denoted $\text{rank}(t)$, is the maximum nesting depth of its function symbols; C_α is precisely that subset of $T_\varphi(C_\alpha)$ of terms of rank 0. Define $T_\varphi^m(C_\alpha)$ to be the subset of $T_\varphi(C_\alpha)$ induced by terms whose rank is $\leq m$. Note that $T_\varphi(C_\alpha)$ is exactly the domain of the *term algebra* of $\sigma \cup F \cup C_\alpha$ (see, e.g., [11]).

Considering all instantiations of x_1, \dots, x_k by the terms of $T_\varphi(C_\alpha)$, we see that $\text{Skolem}(\varphi)$ becomes an infinite set of positive atoms Φ , exactly the instantiations of the quantified atoms of $\text{Skolem}(\varphi)$. These immediately give rise to a canonical (sometimes known as Herbrand) model of $\text{Skolem}(\varphi)$ over the domain $T_\varphi(C_\alpha)$ in the standard way (see, e.g., [11]); we denote this model $\mathcal{T}_\varphi(C_\alpha)$. Note that Φ is the positive (Robinson) diagram of $\mathcal{T}_\varphi(C_\alpha)$. Rather sloppily, we will consider $\mathcal{T}_\varphi(C_\alpha)$ to be at once a σ -structure (a bona fide model of φ) and a $\sigma \cup F \cup C_\alpha$ -structure – this should cause no confusion. By further abuse of nomenclature, we will also continue referring to the elements of $T_\varphi(C_\alpha)$ as ‘terms’ and elements of $C_\alpha \subseteq T_\varphi(C_\alpha)$ as ‘constants’. Let $\mathcal{T}_\varphi^m(C_\alpha)$ be the *truncation* (submodel) of $\mathcal{T}_\varphi(C_\alpha)$ induced by the domain $T_\varphi^m(C_\alpha)$. Note that $\mathcal{T}_\varphi^m(C_\alpha)$ is generally not a model of φ ; however, the following is immediate from the construction.

Fact 1. For all α , $\mathcal{T}_\varphi(C_\alpha) \models \varphi$.



Example 4. Let $\sigma := \langle E \rangle$ contain a single binary relation (i.e. σ -structures are digraphs). Let $\varphi := \forall x \forall z \exists y E(x, y) \wedge E(y, z)$. In this case,¹

$$\text{Skolem}(\varphi) := \forall x \forall z E(x, f(x, z)) \wedge E(f(x, z), z).$$

The quantified atoms of $\text{Skolem}(\varphi)$ are

$$\forall x \forall z E(x, f(x, z)) \text{ and } \forall x \forall z E(f(x, z), z).$$

The following are depictions of the truncations $\mathcal{T}_\varphi^2(C_1)$ and $\mathcal{T}_\varphi^1(C_2)$, respectively.



A Surjective Diagram Lemma

Let φ be a quantified conjunctive-positive sentence, F its associated set of Skolem functions and $\text{Skolem}(\varphi)$ its Skolem normal form. The canonical model $\mathcal{T}_\varphi(C_\omega)$, with a countably infinite set of constants, plays a key role in our discourse. The following is a variant of the Diagram Lemma (see, e.g., [11]).

Lemma 4. Let φ be a quantified conjunctive-positive sentence. Then, for all countable (not necessarily infinite) structures \mathcal{B} , if $\mathcal{B} \models \varphi$ then there is a surjective homomorphism $h : \mathcal{T}_\varphi(C_\omega) \rightarrow \mathcal{B}$ s.t. $h(C_\omega) = B$.

Proof. Let b_1, \dots be an enumeration of the elements of \mathcal{B} . Let \mathfrak{B} be the expansion of \mathcal{B} , over the signature $\sigma \cup C_\omega$ s.t. the elements b_1, \dots interpret the constants c_1, \dots (if \mathcal{B} is finite interpret all remaining constants as, e.g., b_1). Since φ contains no constants, $\mathfrak{B} \models \varphi$. It follows that there is a further expansion $\overline{\mathfrak{B}}$ over the signature $\sigma \cup F \cup C_\omega$, s.t. $\overline{\mathfrak{B}} \models \text{Skolem}(\varphi)$

¹The reader may notice that φ is not in the correct form as it fails to have strict alternation of quantifiers. While the introduction of a dummy existential quantifier (and consequent dummy unary Skolem function in $\text{Skolem}(\varphi)$) would rectify this, it would also make the example rather hard to follow.

Considering $\mathcal{T}_\varphi(C_\omega)$ as a $\sigma \cup F \cup C_\omega$ -structure, we now uncover the canonical function $h : \mathcal{T}_\varphi(C_\omega) \rightarrow \overline{\mathfrak{B}}$. Each $t \in \mathcal{T}_\varphi(C_\omega)$ is a syntactic term over $F \cup C_\omega$. Set $h(t)$ to be the element (which interprets) t in $\overline{\mathfrak{B}}$.

The function h is manifestly a homomorphism, since $\overline{\mathfrak{B}} \models \text{Skolem}(\varphi)$ (actually, it is also unique).

By once again considering $\mathcal{T}_\varphi(C_\omega)$ to be a σ -structure, we see that h is a surjective homomorphism from $\mathcal{T}_\varphi(C_\omega)$ to \mathcal{B} , s.t. $h(C_\omega) = B$. \square

2.1. Characterisation

We are now in a position to derive a model-theoretic characterisation for $\models \varphi \rightarrow \psi$.

Theorem 5 (Methodology II). *Let φ and ψ be quantified conjunctive-positive sentences. The following are equivalent:*

- $\models \varphi \rightarrow \psi$, i.e. $\varphi \rightarrow \psi$ is logically valid, and
- $\mathcal{T}_\varphi(C_\omega) \models \psi$.

Proof. (Downwards.) Since $\models \varphi \rightarrow \psi$, we derive $\mathcal{T}_\varphi(C_\omega) \models \varphi \rightarrow \psi$, whence, since $\mathcal{T}_\varphi(C_\omega) \models \varphi$, we derive $\mathcal{T}_\varphi(C_\omega) \models \psi$.

(Upwards.) This direction requires a little subtlety; we proceed by contraposition. Suppose $\not\models \varphi \rightarrow \psi$; it follows that there is a model for $\varphi \wedge \neg\psi$. By the Downward Löwenheim-Skolem Theorem, it follows that there is a countable model for $\varphi \wedge \neg\psi$, say \mathcal{A} , whereupon $\mathcal{A} \models \varphi$ but $\mathcal{A} \not\models \psi$.

Since $\mathcal{A} \models \varphi$, it follows from Lemma 4 that there is a surjective homomorphism $h : \mathcal{T}_\varphi(C_\omega) \rightarrow \mathcal{A}$. Now, if it were the case that $\mathcal{T}_\varphi(C_\omega) \models \psi$, then we may deduce the contradiction $\mathcal{A} \models \psi$ by Lemma 1. The result $\mathcal{T}_\varphi(C_\omega) \not\models \psi$ follows. \square

2.1.1 Restricting Universal's Play

Now let φ be a quantified conjunctive-positive sentence of which $\mathcal{T}_\varphi(C_\alpha)$ is a canonical model. Let ψ be a quantified conjunctive-positive sentence of the form $\forall x_1 \exists y_1 \dots \forall x_l \exists y_l Q(x_1, y_1, \dots, x_l, y_l)$. The ψ -rel-game on $\mathcal{T}_\varphi(C_\alpha)$ is defined similarly to the ψ -game on $\mathcal{T}_\varphi(C_\alpha)$, except Universal is now restricted to playing elements of $C_\alpha \subseteq \mathcal{T}_\varphi(C_\alpha)$. In this case, Existential has a winning strategy in the ψ -rel-game on $\mathcal{T}_\varphi(C_\alpha)$ iff $\mathcal{T}_\varphi(C_\alpha) \models$

$$\forall x_1 \in C_\alpha \exists y_1 \dots \forall x_l \in C_\alpha \exists y_l Q(x_1, y_1, \dots, x_l, y_l),$$

that is, if $\mathcal{T}_\varphi(C_\alpha)$ models ψ with the universal variables relativised to C_α .

Proposition 1. *Let φ and ψ be quantified conjunctive-positive sentences, with $\mathcal{T}_\varphi(C_\alpha)$ a canonical model of φ .*

Then, Existential has a winning strategy in the ψ -game on $\mathcal{T}_\varphi(C_\alpha)$, i.e. $\mathcal{T}_\varphi(C_\alpha) \models \psi$, iff Existential has a winning strategy in the ψ -rel-game on $\mathcal{T}_\varphi(C_\alpha)$.

Proof. The forward direction is trivial. The backward direction may be proved in a similar manner to Lemma 1, given that Lemma 4 provides us with a surjective endomorphism $s : \mathcal{T}_\varphi(C_\omega) \rightarrow \mathcal{T}_\varphi(C_\omega)$ s.t. $s(C_\omega) = \mathcal{T}_\varphi(C_\omega)$. \square

2.1.2 Substitution Lemmas

Given a term $t \in \mathcal{T}_\varphi(C_\omega)$ one may consider the various *subterms* of which it is composed. For example, the term $f(f(c_1, c_2), f(f(c_1, c_1), c_2))$ of rank 3 contains both c_2 and $f(c_1, c_1)$ as subterms. We will talk of a term t as *containing* the constants that are its subterms. We adopt the notation $t[t'/t'']$ to denote the term obtained by replacing all instances of t' in t by t'' .

Consider terms $t_1, t_2, \dots, t_r, t', t'' \in \mathcal{T}_\varphi(C_\omega)$. Suppose that $R(t_1, t_2, \dots, t_r)$ holds in the canonical model $\mathcal{T}_\varphi(C_\omega)$; might it always be the case that $R(t_1[t'/t''], t_2[t'/t''], \dots, t_r[t'/t''])$ holds in $\mathcal{T}_\varphi(C_\omega)$? The answer is no; for example, in the case of digraphs, if $E(c, f(c)) \in \mathcal{T}_\varphi(C_\omega)$, then we have no reason to conclude that $E(c, c) \in \mathcal{T}_\varphi(C_\omega)$, even though the latter corresponds to $E(c[f(c)/c], f(c)[f(c)/c])$. However, we can make substitutions subject to certain rules, as the following lemmas attest.

Lemma 5 (Substitution of terms of distinct rank). *Let R be a p -ary relation symbol of σ , and consider $t_1, \dots, t_p, t' \in \mathcal{T}_\varphi(C_\omega)$ s.t. $\text{rank}(t')$ is distinct from each of $\text{rank}(t_1), \dots, \text{rank}(t_p)$. For all terms t'' , if $R(t_1, \dots, t_p) \in \mathcal{T}_\varphi(C_\omega)$ then $R(t_1[t'/t''], \dots, t_p[t'/t'']) \in \mathcal{T}_\varphi(C_\omega)$.*

Proof. Consider the quantified atom of $\text{Skolem}(\varphi)$ that caused $R(t_1, \dots, t_p)$ to be in $\mathcal{T}_\varphi(C_\omega)$ (via its instantiation in the positive diagram Φ). It must have been of the form

$$\forall \bar{z}_1 \dots \forall \bar{z}_p R(g_1(\bar{z}_1), \dots, g_p(\bar{z}_p)),$$

where $\bar{z}_1, \dots, \bar{z}_p$ are not required to be disjoint, and each g_i is either

- the identity ι (in which case \bar{z}_i is a singleton) or
- some $f_j \in F$ (in which case \bar{z}_i is a j -tuple).

Since t' is distinct in rank from each of t_1, \dots, t_p , it can be easily seen that all occurrences of t' in the t_1, \dots, t_p of $R(t_1, \dots, t_p)$ must have come from occurrences of t' in the instantiations of the variables $\bar{z}_1, \dots, \bar{z}_p$. It follows that the related instantiation $\bar{z}_1[t'/t''], \dots, \bar{z}_p[t'/t'']$ yields $R(t_1[t'/t''], \dots, t_p[t'/t''])$, and the result follows. \square

Lemma 6 (Substitution of constants). *Let R be a p -ary relation symbol of σ , consider $t_1, \dots, t_p \in \mathcal{T}_\varphi(C_\omega)$ and $c, c' \in C_\omega$. If $R(t_1, \dots, t_p) \in \mathcal{T}_\varphi(C_\omega)$ then $R(t_1[c/c'], \dots, t_p[c/c']) \in \mathcal{T}_\varphi(C_\omega)$.*

Proof. Similar to the previous lemma. \square

Let $\pi : C_\omega \rightarrow C_\omega$ be some (partial) bijection. For a term $t \in \mathcal{T}_\varphi(C_\omega)$, let $\pi(t)$ be the term obtained by simultaneously switching each constant c_i for $\pi(c_i)$, in the obvious manner.

Lemma 7 (Permutation of constants). *Let R be a p -ary relation symbol of σ , and consider $t_1, \dots, t_p \in \mathcal{T}_\varphi(C_\omega)$. Then, $R(t_1, \dots, t_p) \in \mathcal{T}_\varphi(C_\omega)$ iff $R(\pi(t_1), \dots, \pi(t_p)) \in \mathcal{T}_\varphi(C_\omega)$.*

Proof. It is evident from the construction that, for each permutation π , $\mathcal{T}_\varphi(C_\omega)$ has an automorphism that maps each term t to $\pi(t)$. \square

The structure $\mathcal{T}_\varphi(C_\alpha)$ has the useful property that any finite substructure $\mathcal{A} \subseteq \mathcal{T}_\varphi(C_\alpha)$ has a homomorphism to the truncation $\mathcal{T}_\varphi^{|\mathcal{A}|}(C_\alpha)$. In fact, we are able to derive a stronger property. Call a partial function $f : \mathcal{T}_\varphi(C_\alpha) \rightarrow \mathcal{T}_\varphi(C_\alpha)$ *constant-conservative* if, for all $t \in \mathcal{T}_\varphi(C_\alpha)$, $f(t)$ contains no constants that are not contained in t .

Lemma 8. *For $\mathcal{A} \subseteq \mathcal{T}_\varphi(C_\alpha)$, there is a constant-conservative homomorphism $\mathcal{A} \rightarrow \mathcal{T}_\varphi^{|\mathcal{A}|}(C_\alpha)$.*

The general idea of the proof is, in the (worst) case that the terms of \mathcal{A} have distinct ranks, that they can still all be mapped to the first $|\mathcal{A}|$ ranks in a way that preserves the rank-order. The proof uses Lemma 5 in order to explain what we do when a rank has been ‘missed out’ in \mathcal{A} . Indeed, when a rank has been missed out, then we may reduce the rank of all higher terms in the rank-order, in an almost arbitrary way, while preserving homomorphism. However, to ensure that the homomorphism is constant-conservative, we reduce rank in a more particular manner.

Proof. Let $t_1, \dots, t_{|\mathcal{A}|}$ be the elements of \mathcal{A} ordered by increasing rank. If the maximal rank is $> |\mathcal{A}|$ then there exists some $t_i \in \mathcal{A}$ of rank r s.t. no $t \in \mathcal{A}$ is of rank $r - 1$, and t_i is of the form $f_j(s_1, \dots, s_j)$ for some terms s_1, \dots, s_j of which (at least) one is of rank $r - 1$. Suppose one that is of rank $r - 1$ is s_m . Pick any subterm s'_m of s_m of rank $r - 2$. Let \mathcal{A}' be that substructure of $\mathcal{T}_\varphi(C_\alpha)$ derived by substituting s'_m for s_m in all the terms of \mathcal{A} . Clearly this substitution is constant-conservative. We claim that the function from \mathcal{A} to \mathcal{A}' induced by this substitution is a homomorphism, whereupon we may iterate the above reasoning until the obtained structure has maximal rank $\leq |\mathcal{A}|$.

(Proof that $\mathcal{A} \rightarrow \mathcal{A}'$.) Consider the elements $t_1, \dots, t_{|\mathcal{A}|}$ of \mathcal{A} and the natural map that takes them

to $t_1[s_m/s'_m], \dots, t_{|\mathcal{A}|}[s_m/s'_m]$ in \mathcal{A}' . We will demonstrate that this is a homomorphism. Let R be a p -ary relation symbol of σ . Suppose $R(t_{\lambda_1}, \dots, t_{\lambda_p}) \in \mathcal{A} \subseteq \mathcal{T}_\varphi(C_\alpha) \subseteq \mathcal{T}_\varphi(C_\omega)$, by Lemma 5 we have $R(t_{\lambda_1}[s_m/s'_m], \dots, t_{\lambda_p}[s_m/s'_m]) \in \mathcal{T}_\varphi(C_\omega)$, whereupon the result follows (since \mathcal{A}' is an induced substructure of $\mathcal{T}_\varphi(C_\alpha) \subseteq \mathcal{T}_\varphi(C_\omega)$). \square

2.1.3 Restricting Existential’s Play

Proposition 1 tells us that we may consider Universal’s play restricted to the set C_α in the ψ -game on $\mathcal{T}_\varphi(C_\alpha)$. Now we detail how we may make a certain assumption about Existential’s play, without affecting her ability to win.

Let φ, ψ be quantified conjunctive-positive sentences, with ψ of the form $\forall x_1 \exists y_1 \dots \forall x_l \exists y_l Q(x_1, y_1, \dots, x_l, y_l)$, and let $\mathcal{T}_\varphi(C_\alpha)$ be a canonical model of φ . Define the ψ -rel-cc-game on $\mathcal{T}_\varphi(C_\alpha)$ as the ψ -rel-game on $\mathcal{T}_\varphi(C_\alpha)$ but now restrict Existential to only playing terms t containing constants that Universal has already played (the cc abbreviates constant-conservative). In other words, if Universal has played c_{j_1}, \dots, c_{j_i} for variables x_1, \dots, x_i , then Existential must play some $t \in \mathcal{T}_\varphi(\{c_{j_1}, \dots, c_{j_i}\})$ for y_i . Legitimate strategies for Existential in this game will be termed *constant-conservative*. Winning strategies for Existential in the ψ -rel-cc-game on $\mathcal{T}_\varphi(C_\alpha)$ are central to our discourse.

Consider the ψ -rel-game (resp., ψ -rel-cc-game) on the truncation $\mathcal{T}_\varphi^m(C_\alpha) \subseteq \mathcal{T}_\varphi(C_\alpha)$ defined in the obvious way.

Proposition 2. *Let φ, ψ be quantified conjunctive-positive sentences, with ψ of the form $\forall x_1 \exists y_1 \dots \forall x_l \exists y_l Q(x_1, y_1, \dots, x_l, y_l)$. The following are equivalent.*

- (i) *Existential has a winning strategy in the ψ -rel-game on $\mathcal{T}_\varphi(C_\omega)$.*
- (ii) *Existential has a winning strategy in the ψ -rel-cc-game on $\mathcal{T}_\varphi(C_\omega)$.*
- (iii) *Existential has a winning strategy in the ψ -rel-cc-game on $\mathcal{T}_\varphi(C_l)$.*
- (iv) *Existential has a winning strategy in the ψ -rel-cc-game on $\mathcal{T}_\varphi^{l+2}(C_l)$.*

The proof of this proposition will be broken into natural constituent parts.

Proof of Proposition 2 (i \Rightarrow ii). Consider a *game tree* \mathcal{G}_ε for the ψ -rel-game on $\mathcal{T}_\varphi(C_\omega)$ under Existential strategy ε . \mathcal{G}_ε is an out-tree, branching on all possible Universal moves over C_ω when Existential plays according to ε . The branching factor of \mathcal{G}_ε from the root to the leaves is alternately ω and 1, and the distance from the root to the leaves is $2l$. The nodes at distance $2i - 1$ (resp., $2i$) from the root are labelled

with Universal's (resp., Existential's) i th move. The root is unlabelled. If ε is a winning strategy, then when we read off valuations for $x_1, y_1, \dots, x_l, y_l$ on a path, we will always have $\mathcal{T}_\varphi(C_\omega) \models Q(x_1, y_1, \dots, x_l, y_l)$.

We will modify \mathcal{G}_ε inductively from the root to the leaves, in such a way as to ultimately enforce that Existential's moves are constant-conservative while keeping her strategy winning. The property (*) that we will maintain is that, at distance $\leq 2i$ from the root, there is no node λ labelled by an Existential play t containing a constant c that Universal has not played on the path from the root to λ . When $i = 0$ this is clearly true; and when $i = 2l$ we have that Existential's play was always constant-conservative.

Suppose the inductive hypothesis (*) holds at distance $\leq 2i$ from the root. While there is a node λ , at distance $2(i+1)$ from the root, labelled by an Existential play t containing a constant c that Universal has not played on the path from the root to λ , we undertake the following procedure.

- Remove all subtrees beyond λ whose roots are labelled with Universal plays c .
- Pick a constant c' that has been already played by Universal on the path from the root to λ , and substitute all terms t labelling nodes in the subtree rooted at λ with $t[c/c']$.

It follows from Lemma 6 that this modified game tree still represents a winning strategy for Existential, *so long as Universal never plays c beyond node λ* .

Now consider all missing subtrees corresponding to Universal plays of c after λ . These follow Existential plays at nodes $\lambda_1 := \lambda, \lambda_2, \dots, \lambda_{(l-i-1)}$ at distances $0, 2, \dots, 2(l-i-1)$ beyond λ . For each $r \in \{0, 1, \dots, 2(l-i-1)\}$, consider what Universal plays for x_{i+1+r} :

- Pick some next Universal play that is a constant c'' s.t. c'' has not appeared on any path from the root to λ_r (such a constant must exist since only a finite number of constants can be mentioned on any path).
- Take the bijection $\pi : C_\omega \rightarrow C_\omega$ that swaps c and c'' . Duplicate the subtree corresponding to the choice c'' (i.e. rooted at the node labelled c'' immediately after λ) but reset all the node labels t to $\pi(t)$. Now reintroduce this subtree as the choice c (immediately after λ).

Since neither c'' nor c is mentioned before λ_r , it follows from Lemma 7 that this modified game tree still represents a winning strategy for Existential. \square

Proof of Proposition 2 (ii \Rightarrow iii). Existential may use the same winning strategy in the ψ -rel-cc-game on $\mathcal{T}_\varphi(C_l)$ as she used in the ψ -rel-cc-game on $\mathcal{T}_\varphi(C_\omega)$. This is because her play is constant-conservative. \square

Proof of Proposition 2 (iii \Rightarrow iv). Consider a winning strategy ε in the ψ -rel-cc-game on $\mathcal{T}_\varphi(C_l)$. We will construct a winning strategy ε' for her in the ψ -rel-cc-game on $\mathcal{T}_\varphi^{l+2}(C_l)$. Recall x_1, \dots, x_l are the ordered universal variables of ψ ; there are at most l^l ways in which they may be, in order, played on to the set C_l . This means that Existential needs at most $l \cdot l^l$ elements of $\mathcal{T}_\varphi(C_l)$ to beat any strategy of Universal. This means that there is a substructure $\mathcal{A} \subseteq \mathcal{T}_\varphi(C_l)$ that contains at most $l \cdot l^l$ elements other than those of C_l s.t. Existential has the winning strategy ε in the ψ -rel-cc-game on \mathcal{A} . Note that $|\mathcal{A}| \leq l + l \cdot l^l \leq l^{l+2}$.

Let $h : \mathcal{A} \rightarrow \mathcal{T}_\varphi^{l+2}(C_l)$ be a (constant-conservative) homomorphism, as guaranteed by Lemma 8. It follows that $\varepsilon' := h \circ \varepsilon$ suffices. \square

Proof of Proposition 2 (iv \Rightarrow i). Suppose Existential has a winning strategy ε in the ψ -rel-cc-game on $\mathcal{T}_\varphi^{l+2}(C_l)$, we will construct a (constant-conservative) winning strategy ε' for her in the ψ -rel-game on $\mathcal{T}_\varphi(C_\omega)$. At the j th round, Existential has in mind a partial bijection $\pi_j : C_\omega \rightarrow C_\omega$.

Universal plays first, with some constant c_{i_1} for x_1 . Existential sets $\pi_1 := (c_{i_1}, c_1)$ (i.e. the partial bijection that swaps c_{i_1} and c_1), and responds with $\pi_1^{-1} \circ \varepsilon_1(\pi_1(x_1)) = \pi_1^{-1} \circ \varepsilon_1(c_1)$ for y_1 . At the $j+1$ th round, Universal plays some $c_{i_{j+1}}$ for x_{j+1} . If Universal has already played this, then Existential sets $\pi_{j+1} := \pi_j$; otherwise Existential sets $\pi_{j+1} := (c_{i_{j+1}}, c_{j+1}) \circ \pi_j$. In both cases she responds with

$$\pi_{j+1}^{-1} \circ \varepsilon_{j+1}(\pi_{j+1}(x_1), \pi_{j+1}(y_1), \dots, \pi_{j+1}(x_{j+1}))$$

for y_{j+1} . Since the strategy ε is constant-conservative, no new constants are introduced through ε , and it follows from Lemma 7 that the strategy ε' is winning. \square

Remark 1. Although the constant-conservative nature of Existential's play is used in the proof of (ii \Rightarrow iii) above, it is only a truly vital component in the proof of (iv \Rightarrow i). Imagine the play were not constant-conservative in that proof. Universal begins in the ψ -rel-game on $\mathcal{T}_\varphi(C_\omega)$ by playing c_{i_1} for x_1 , and Existential sets $\pi_1 := (c_{i_1}, c_1)$. In the auxiliary ψ -rel-cc-game on $\mathcal{T}_\varphi^{l+2}(C_l)$, Existential now looks up what she would have played in her winning strategy if Universal had played c_1 for x_1 . But, she might have played a response for y_1 that contains $> l$ constants! Clearly there is no hope to extend the partial bijection π_1 s.t. the range involves elements only from C_l .

2.2 An Algorithm for Containment

Our decision procedure for the entailment problem makes use of the following fact, which may be proved by induction on m .

Fact 2. If φ is a quantified conjunctive-positive sentence of depth k , then $|T_\varphi^m(C_l)| \leq (l+1)^{(k+1)^m}$.

Theorem 6. The entailment problem for quantified conjunctive-positive sentences is decidable in triple exponential time.

Proof. Consider the input sentences φ and ψ of depth k and l , respectively. By Theorem 5 and Proposition 2, it suffices to verify whether Existential has a winning strategy in the ψ -rel-cc-game on $\mathcal{T}_\varphi^{l+2}(C_l)$. The structure $\mathcal{T}_\varphi^{l+2}(C_l)$ is of size bounded by

$$\zeta := (l+1) \uparrow (k+1) \uparrow (l) \uparrow (l+2),$$

where the \uparrow denotes exponentiation (with precedence to the right). We may search through all $2l$ -tuples that could be played in the ψ -rel-cc-game on $\mathcal{T}_\varphi^{l+2}(C_l)$, in time $O(\zeta^{2l})$ to determine whether Existential has a winning strategy. Noting that

$$\zeta^{2l} = O((l+1) \uparrow (k+1) \uparrow (l) \uparrow 2l(l+2)),$$

the result follows. \square

2.3 Undecidability of Entailment for Positive FO

The *entailment problem for positive FO* (EPPFO) is defined as follows.

- Input: two sentences φ and ψ of positive (equality-free) FO.
- Question: does $\models \varphi \rightarrow \psi$?

We consider also its dual problem, DUAL-EPPFO.

- Input: two sentences φ and ψ of positive (equality-free) FO.
- Question: is $\varphi \wedge \neg\psi$ satisfiable?

These problems are clearly Turing equivalent ($\varphi \wedge \neg\psi$ is satisfiable iff it is not the case that $\neg\varphi \vee \psi$ is valid), and undecidability of the latter implies undecidability of the former.

We introduce one further problem, which may be seen as the satisfiability version of the (pure predicate) Classical Decision Problem, SAT-CDP.

- Input: a sentence φ of (equality-free) FO.
- Question: is φ satisfiable?

It is well-known that this problem is undecidable (see, e.g., [1]). We are now in a position to prove the main result of this section.

Theorem 7. The entailment problem for positive (equality-free) FO, EPPFO, is undecidable.

Proof. By reduction from the SAT-CDP to the problem DUAL-EPPFO defined above. Let φ be some input to the SAT-CDP, containing relation symbols R_1, \dots, R_r , of respective arities a_1, \dots, a_r . We introduce r new relation symbols S_1, \dots, S_r , also of respective arities a_1, \dots, a_r . We will now use these S -relations to axiomatise negation. Consider

$$\theta_0 := \bigwedge_{i=1}^r \forall \mathbf{x}_i S_i(\mathbf{x}_i) \leftrightarrow \neg R_i(\mathbf{x}_i)$$

$$\theta_1 := \bigwedge_{i=1}^r \forall \mathbf{x}_i S_i(\mathbf{x}_i) \vee R_i(\mathbf{x}_i)$$

$$\theta_2 := \bigwedge_{i=1}^r \forall \mathbf{x}_i \neg S_i(\mathbf{x}_i) \vee \neg R_i(\mathbf{x}_i),$$

where each \mathbf{x}_i is an a_i -tuple. Note that θ_0 is logically equivalent to $\theta_1 \wedge \theta_2$. Now note that θ_2 is logically equivalent to

$$\neg \bigvee_{i=1}^r \exists \mathbf{x}_i S_i(\mathbf{x}_i) \wedge R_i(\mathbf{x}_i),$$

which we designate $\neg\psi$ (where ψ is positive). Finally, derive φ' from φ by first propagating all negations to atomic level and then substituting any instances of negated relations $\neg R_i$ with S_i . It is easy to see that φ is satisfiable iff $(\varphi' \wedge \theta_1) \wedge \neg\psi$ is satisfiable. Furthermore, $\varphi' \wedge \theta_1$ and ψ are (equality-free) positive, and the result follows. \square

3 Final Remarks

The model containment problem. Two questions in particular arise from our discussion, and provide the most immediate challenge for further investigations.

We know that both the model containment problem for CSP and the CSP itself are NP-complete; indeed they are essentially the same problem. Given that the QCSP is Pspace-complete, it may be wondered what is the exact complexity of its associated model containment problem. It is far from clear that our algorithm is optimal; might the containment problem also be in Pspace, and, if so, might it be complete?

In the world of CSP, the core is a well-understood notion. For a structure \mathcal{A} , the minimal (w.r.t. size of domain) substructures of \mathcal{A} that are homomorphically equivalent with \mathcal{A} are necessarily isomorphic to one another; thus it is that we speak of *the* core. Suppose that we define a q -core of a structure \mathcal{A} to be a minimal substructure $\mathcal{A}' \subseteq \mathcal{A}$ s.t. $\text{QCSP}(\mathcal{A}') = \text{QCSP}(\mathcal{A})$. Thus far we have failed to prove, if \mathcal{A}' and \mathcal{A}'' are both q -cores of \mathcal{A} , then \mathcal{A}' and \mathcal{A}'' are isomorphic. It would suffice to settle the following. Let $\mathcal{A} \sim \mathcal{B}$ be the equivalence relation given to hold exactly

when $\text{QCSP}(\mathcal{A}) = \text{QCSP}(\mathcal{B})$. For any \mathcal{A} , if \mathcal{A}' and \mathcal{A}'' are minimal representatives of \mathcal{A} 's equivalence class under \sim , then \mathcal{A}' and \mathcal{A}'' are isomorphic.

The entailment problem. While $\models \varphi \rightarrow \psi$ is undecidable when both φ and ψ are positive FO, an analysis of our method yields that it is actually decidable for φ quantified conjunctive-positive and ψ positive. This is because we may still build the canonical model of φ , and our game semantics hold in the presence of disjunction.

It is unclear how our method might be brought to bear on the question, for quantified conjunctive-positive φ and ψ , as to whether $\models_{\text{FIN}} \varphi \rightarrow \psi$. If one could construct a finite canonical model \mathcal{F}_φ for each φ , i.e. a finite model that still respects Theorem 5 (Methodology II), one would have solved this. To see this, assume that we have a finite \mathcal{F}_φ s.t. $\models \varphi \rightarrow \psi$ iff $\mathcal{F}_\varphi \models \psi$. We now prove that these are equivalent to $\models_{\text{FIN}} \varphi \rightarrow \psi$, by showing that $\models_{\text{FIN}} \varphi \rightarrow \psi$ implies $\models \varphi \rightarrow \psi$. This latter implication is immediate, since $\mathcal{F}_\varphi \models \varphi$ and \mathcal{F}_φ is finite we can derive $\mathcal{F}_\varphi \models \psi$, and the result follows from our assumption.

However, even for some simple sentences, we can demonstrate that there can be no finite canonical model. Consider $\varphi_1 := \forall x \exists y E(x, y)$, whose canonical models $\mathcal{T}_{\varphi_1}(C_1)$ and $\mathcal{T}_{\varphi_1}(C_\omega)$ are the infinite directed path (\mathcal{DP}_ω) and ω disjoint copies of said path ($\mathcal{DP}_\omega \uplus \mathcal{DP}_\omega \uplus \dots$), respectively.

Suppose we had a finite model \mathcal{F}_{φ_1} of size d s.t., for all quantified conjunctive-positive ψ , $\mathcal{F}_{\varphi_1} \models \psi$ iff $\models \varphi_1 \rightarrow \psi$. Since $\mathcal{F}_{\varphi_1} \models \varphi_1$, \mathcal{F}_{φ_1} must contain (as a not-necessarily induced submodel) a directed cycle of length $e \leq d$ (\mathcal{DC}_e). It follows that the sentence $\psi' :=$

$$\exists x_1, x_2, \dots, x_{e-1}, x_e \quad E(x_1, x_2) \wedge \dots \wedge E(x_{e-1}, x_e) \wedge E(x_e, x_1)$$

is true on \mathcal{F}_{φ_1} . But $\varphi_1 \rightarrow \psi'$ is not logically valid, since \mathcal{DC}_{e+1} is a model of the former but not the latter.

On the other hand, for some sentences we can produce finite canonical models. For $\varphi_2 := \forall x \exists y E(x, y) \wedge E(y, x)$, the finite canonical model \mathcal{K}_2 (or $\mathcal{K}_2 \uplus \mathcal{K}_2$) exists. That \mathcal{K}_2 is sufficient for this task follows from the fact that, for all models \mathcal{A} of φ_2 , there exists a constant $k_{\mathcal{A}}$ s.t. $(\mathcal{K}_2)^{k_{\mathcal{A}}} \rightarrow \mathcal{A}$, and therefore $\text{QCSP}(\mathcal{K}_2) \subseteq \text{QCSP}(\mathcal{A})$. Similarly, for $\varphi_3 :=$

$$\forall x \exists y \exists z \quad E(x, y) \wedge E(y, x) \wedge E(y, z) \wedge E(z, y) \wedge E(z, x) \wedge E(x, z),$$

the canonical model $\mathcal{K}_3 \uplus \mathcal{K}_3$ exists. In the latter case \mathcal{K}_3 will not do: consider $\psi'' :=$

$$\forall x \forall y \exists w \exists z \quad E(x, y) \wedge E(y, w) \wedge E(w, z) \wedge E(z, y);$$

$\varphi_3 \rightarrow \psi''$ is not logically valid, as $\mathcal{K}_3 \uplus \mathcal{K}_3$ models the former but not the latter, but $\mathcal{K}_3 \models \psi''$.

These examples perhaps suggest a study of quantified conjunctive-positive sentences whose underlying digraphs are symmetric.

Acknowledgements

We are very grateful to Arnaud Durand for supplying the undecidability proof of the entailment problem for positive FO. We are also grateful to an anonymous referee for directing us to the paper of Keisler.

References

- [1] E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Springer-Verlag, Berlin, 1997.
- [2] F. Börner, A. Krokhin, A. Bulatov, and P. Jeavons. Quantified constraints and surjective polymorphisms. Technical Report PRG-RR-02-11, Oxford University, 2002.
- [3] A. Bulatov, A. Krokhin, and P. Jeavons. Constraint satisfaction problems and finite algebras. In *ICALP*, volume 1853 of *LNCS*, pages 272–282. Springer-Verlag, 2000.
- [4] A. A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM*, 53(1):66–120, 2006.
- [5] H. Chen. The complexity of quantified constraint satisfaction: Collapsibility, sink algebras, and the three-element case, 2006. To appear *SIAM J. Comput.*
- [6] H. Chen. A rendez-vous of logic, complexity, and algebra. *ACM SIGACT News*, 2006.
- [7] T. Eiter, M. Fink, H. Tompits, and S. Woltran. Complexity results for checking equivalence of stratified logic programs. In M. M. Veloso, editor, *IJCAI*, pages 330–335, 2007.
- [8] T. Eiter, M. Fink, and S. Woltran. Semantical characterizations and complexity of equivalences in answer set programming. *ACM Trans. Comput. Logic*, 8(3):17, 2007.
- [9] T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through datalog and group theory. *SIAM J. Comput.*, 28, 1999.
- [10] P. Hell and J. Nešetřil. On the complexity of H-coloring. *J. Combin. Theory Ser. B*, 48, 1990.
- [11] W. Hodges. *Model Theory*. Cambridge University Press, 1993.
- [12] H. J. Keisler. Reduced products and horn classes. *Trans. AMS*, 117:307–328, 1965.
- [13] P. G. Kolaitis and M. Y. Vardi. *Finite Model Theory and Its Applications*, chapter A logical Approach to Constraint Satisfaction. Springer-Verlag, 2005.
- [14] D. Kozen. Communication: Positive first-order logic is NP-complete. *IBM J. Res. Dev.*, 25(4):327–332, 1981.
- [15] V. Lifschitz, D. Pearce, and A. Valverde. Strongly equivalent logic programs. *Computational Logic*, 2(4):526–541, 2001.
- [16] B. D. Martin. *Logic, Computation and Constraint Satisfaction*. PhD thesis, University of Leicester, 2005.
- [17] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [18] T. Schaefer. The complexity of satisfiability problems. In *STOC*, 1978.