Herbrand Property, Finite Quasi-Herbrand Models, and a Chandra-Merlin Theorem for Quantified Conjunctive Queries

Simone Bova
Vienna University of Technology, Vienna, Austria

Fabio Mogavero
University of Oxford, Oxford, U.K.

***Abstract*—A structure enjoys the *Herbrand property* if, whenever it satisfies an equality between some terms, these terms are unifiable. On such structures the expressive power of equalities becomes trivial, as their semantic satisfiability is reduced to a purely syntactic check.**

In this work, we introduce the notion of Herbrand property and develop it in a finite model-theoretic perspective. We provide, indeed, a canonical realization of the new concept by what we call *quasi-Herbrand models* and observe that, in stark contrast with the naive implementation of the property via standard Herbrand models, their universe can be *finite* even in presence of functions in the vocabulary. We exploit this feature to decide and collapse the general and finite version of the satisfiability and entailment problems for previously unsettled fragments of first-order logic.

We take advantage of the Herbrand property also to establish novel and tight complexity results for the aforementioned decision questions. In particular, we show that the finite containment problem for quantified conjunctive queries is NPTIME-complete, tightening along two dimensions the known 3EXPTIME upper bound for the general version of the problem (Chen, Madelaine, and Martin, LICS'08). We finally present an alternative view on this result by generalizing to such queries the classic characterization of conjunctive query containment via polynomial-time verifiable homomorphisms (Chandra and Merlin, STOC'77).

I. Introduction

A fundamental theorem by Skolem [31] establishes that every first-order sentence without equality is *satisfiable* if and only if its *functional (Skolem) normal form* has a *canonical (Herbrand) model*. In this context, the universe of discourse is the set of ground terms over the vocabulary of the sentence and the interpretation of the functions is defined in an algebraically transparent way: each term denotes precisely itself.

A breakthrough technique by Büchi [10], further refined by Aanderaa [1] and Börger [6], exploits the structure of Herbrand models of relational first-order $\exists\forall\exists\forall$-sentences to prove the undecidability of the corresponding prefix class. Here, the Herbrand universe encodes the set of natural numbers with zero and successor. In that way, the data structure operated by a two-register machine are implemented transparently. This allows, therefore, an elementary reduction from the associated halting problem, which bypasses entirely the cumbersome axiomatization of the underlying register operations [7].

The transparency of classic Herbrand interpretations, which underlies their success as a tool for undecidability proofs, as well as numerous other applications in mathematical logic and theoretical computer science (e.g., in completeness theorems [16], semantic tableaux [21], alternative first-order semantics [17], automated reasoning [11], logic programming [23], and database theory [2]), comes at a price: their lack of succinctness. Indeed, as soon as the vocabulary contains a function symbol, the corresponding Herbrand universe becomes infinite. This phenomenon severely limits the effectiveness of Herbrand models in establishing the decidability of fragments of first-order logic with functions, not to mention in obtaining tight computational-complexity bounds or model-theoretic results like the finite-model property.

Aiming at decidability, however, more useful appears a property of Herbrand models implied by their transparency, rather than the transparency itself: *an equality between terms is satisfiable on a Herbrand model if and only if its terms are unifiable*. Intuitively, the particular interpretation of terms neutralizes the expressive power of equalities, by reducing their satisfiability, at first glance a hard, even infinitary, question, to a polynomial-time unifiability test. This observation has been exploited by Kozen to show that the validity problem of positive first-order logic is in NPTIME [22].

The present article is devoted to the study and application of the *Herbrand property*, a novel model-theoretic notion expressing the fact that the satisfiability of an equation boils down to the unifiability of its terms. In this terminology, the aforementioned observation by Kozen can be rephrased as follows: every Herbrand model enjoys the Herbrand property. Our work, though, tackles the concept per se, abstracting it from the specific implementation via Herbrand models, and investigates its consequences from both a finite model-theoretic and a structural-complexity perspective. We obtain non-trivial results on both levels.

A first part of the work (Section III and part of Section IV) is devoted to the finite model-theoretic development of the Herbrand property. The main result (Theorem 1) is a universal and finitary version of this concept, summarized as follows.

> *A set of terms is equalizable over all finite structures if and only if it is unifiable.*

Here, a set of terms is said to be *equalizable* over a structure if such a structure satisfies all pairwise equalities between terms of the set. The universal and finitary aspect of our characterization contrasts with the reduction observed by Kozen from *equalizability over Herbrand models* to unifiability, as we reduce *equalizability on all finite structures* to unifiability.

An easy corollary of the above result is the existence of finite models enjoying the Herbrand property, which we call *finite quasi-Herbrand models* (Definition 4 and Theorem 2). This can be seen as an evidence of the fact that the intrinsic infinitary nature of Herbrand models over vocabularies with functions is inessential. In other words, the latter can be seen as a naively verbose implementation of this fundamental concept.

The main consequence of our finitary version of the Herbrand property is that satisfiable universal single-binding sentences have finite quasi-Herbrand models (Theorem 5), i.e., more abstractly, the fragment of universal single-binding logic enjoys the *finite* (technically, *small* [32]) *model property*.

Universal single-binding logic is the language of positive Boolean combinations of universally quantified binding forms, where a *binding form* is, in turn, a Boolean combination of relational atoms over the same tuple of terms. This logic is syntactically contained in *conjunctive-binding logic* introduced in [27], a fragment of first-order logic that allows positive Boolean combinations of quantified conjunctions of binding forms. Since the satisfiability problems for the two logics are succinctly interreducible via skolemization, the following result holds (Corollary 2).

> *Conjunctive-binding logic enjoys the finite model property.*

In particular, its (finite) satisfiability problem is decidable, answering an open question in the literature and completing the decidability classification of binding fragments of first-order logic [27]. The result can also be read as a non-trivial generalization of the decidability proof for *Herbrand logic* [15], the language of quantified conjunctions of literals, as it is syntactically contained in the logic under analysis. On the other hand, conjunctive-binding logic is orthogonal to all known decidable fragments (prefix classes [7], two variable [19], [28], guarded fragments [3], [18], guarded negation [5], et cetera, see [27] for details) and its solution requires different ideas and techniques.

The rest of the work (part of Section IV and Section V) focuses on the consequences of the Herbrand property from the structural complexity viewpoint with respect to various satisfiability and entailment problems in conjunctive-binding logic and fragments thereof (Theorem 9 and Theorem 10).

Our first result (Theorem 5) is a characterization of the (finite) satisfiability problem for universal single-binding logic in terms of (finite) quasi-Herbrand models, placing this problem at the third level of the polynomial hierarchy (Corollary 3). The aforementioned interreducibility allows then to prove the following statement (Corollary 4).

> *The (finite) satisfiability problem for conjunctive-binding logic is* Σ_3^{P}*-complete.*

As opposed to satisfiability, the entailment problem for conjunctive-binding logic is, unfortunately, undecidable (Theorem 6). Interestingly enough, the prominent syntactic fragment of *quantified conjunctive queries (*QCQ*)* has been shown to have a decidable (*general*) entailment problem by Chen, Madelaine, and Martin [13]. This problem is closely related to QCQ *containment* in database theory. In this context, however, the notion of interest is *finite* entailment, i.e., entailment on all finite structures, as in most applications the database is finite. The question whether entailment and finite entailment in QCQ coincide, though, was left open in [13].

Our second result (Theorem 9) is a tight structural complexity classification of general and finite entailment in *positive Herbrand logic (*PH*)*, the logic of quantified conjunctions of *atoms*, which syntactically contains QCQ[1].

[1]Due to a terminological clash with [7, Definition 2.1.14], we avoid calling the fragment positive Horn first-order logic as in [13].

> *The (finite) entailment problem in positive Herbrand logic is* NPTIME*-complete.*

Our result has both a complexity-theoretic and a logical value: on the one hand it closes the previously standing gap between NPTIME-hardness and 3EXPTIME-membership for QCQ containment [13]; on the other hand, by exploiting our *finitary Herbrand property*, it actually pushes the finite version of the problem in NPTIME, even for PH. In retrospect, and not coincidentally, Chen, Madelaine, and Martin obtain their 3EXPTIME upper bound by reasoning on a finite substructure of an infinite Herbrand model associated with the Skolem normal form of the implicant sentence in the instance.

Our proof of Theorem 9, placing the problem in NPTIME, relies on the observation that positive instances of PH entailment have short resolution refutations. A careful inspection reveals that such small witnesses encode certain mappings from the consequent to the antecedent in the instance. In particular, in the special case of conjunctive queries (CQ), it is readily seen that these mappings are precisely homomorphisms. We have thus recovered the classic theorem by Chandra and Merlin [12], which places the (finite) containment question for CQ in NPTIME.

Our third and final result (Theorem 10), stemmed from this insight, consists in abstracting a lifted notion of homomorphism from short refutations of positive QCQ entailment instances. This notion characterizes the QCQ containment problem.

> *Given two* QCQ*s* ϕ *and* ψ*, it holds that* $\phi \models \psi$ *if and only if* ψ *admits a Skolem homomorphism to* ϕ*.*

A *Skolem homomorphism* (Definition 8) is a substitution of the variables in ψ by terms of the skolemization of ϕ, which is both sensitive to the dependencies induced by the quantifier prefix of ψ and faithful to the relational structure associated with ϕ. Besides, such an alleged Skolem homomorphism is efficiently checkable relative to ϕ and ψ, thus yielding an alternative view on the NPTIME-membership of the QCQ containment problem. Our result can be read, therefore, as an accurate lifting of Chandra-Merlin theorem to the QCQ realm.

The hard direction of Theorem 10 consists in proving that $\phi \models \psi$ implies the existence of a Skolem homomorphism from ψ to ϕ (Lemma 9). A first step (Lemma 7) reduces an instance of the QCQ containment problem to a satisfiability check of a universal single-binding sentence. Next, this check is converted into a unification problem between certain terms s and t derived from the original instance (Claim 4), appealing to the characterization of satisfiability in universal single-binding logic. Here, the application of the Herbrand property is reminiscent of analogous reductions in Kozen [22] and Denenberg and Lewis [15]. A last step (Lemma 10) extracts the desired Skolem homomorphism from the unifier of s and t.

The article is organized as follows. In Section III, we prove the equivalence of finite equalizability and unifiability and introduce the notion of quasi-Herbrand models. In Section IV, we characterize satisfiability in universal single-binding logic, obtaining the finite-model property and the complexity bounds for conjunctive-binding logic. Finally, in Section V, we first show that the (finite) entailment problem for PH is NPTIME-complete and then lift the Chandra-Merlin theorem to QCQ.

II. PRELIMINARIES

Let $\mathbb{N}$ and $\mathbb{N}_+ = \mathbb{N} \setminus \{0\}$ be the sets of natural and positive natural numbers and $[0, n] = \{0, \ldots, n\}$ and $[n] = \{1, \ldots, n\}$ with $[0] = \emptyset$ their initial segments, where $n \in \mathbb{N}$.

Given a set A, we let $\mathbf{a} = (a_1, \ldots, a_n) \in A^n$ denote a tuple of dimension $n \in \mathbb{N}$, where $\mathbf{a}$ is the empty tuple in case $n = 0$. Moreover, $\mathbf{a}_i = a_i$ indicates its i-th projection with $i \in [n]$. We also view a 1-dimensional tuple $\mathbf{a} = (a)$ over A as a. The length of an n-dimensional tuple $\mathbf{a}$ is denoted by $|\mathbf{a}| = n$. Finally, we freely identify a tuple $\mathbf{a} = (a_1, \ldots, a_n)$ with the word $a_1 \cdots a_n$ over the alphabet A, so that the empty word ε is identified with the empty tuple.

In the sequel, X is a countable set of variables and Σ is a first-order vocabulary split into its functional and relational components Σ_{f} and Σ_{r}, respectively. A Σ-structure $\mathcal{A}$ is defined by a universe A together with an interpretation of Σ over A, i.e., every function symbol $f \in \Sigma_{\mathrm{f}}$ is interpreted by a function $f^{\mathcal{A}} \colon A^{\mathsf{ar}(f)} \to A$ and every relation symbol $r \in \Sigma_{\mathrm{r}}$ is interpreted by a relation $r^{\mathcal{A}} \subseteq A^{\mathsf{ar}(r)}$, where $\mathsf{ar} \colon \Sigma \to \mathbb{N}$ associates each symbol in the vocabulary with its arity. The order of $\mathcal{A}$ is the cardinality $|A|$ of its universe.

III. QUASI-HERBRAND STRUCTURES

A *Herbrand structure* $\mathcal{H}$ over a vocabulary Σ is a first-order structure whose universe H is the set of all ground terms built over the functional component Σ_{f} of Σ and where each function symbol $f \in \Sigma_{\mathrm{f}}$ is interpreted as the function $f^{\mathcal{H}} \colon H^{\mathsf{ar}(f)} \to H$ applying that symbol to its argument, i.e., $f^{\mathcal{H}}(\mathbf{h}) = f\mathbf{h}$, with $\mathbf{h} \in H^{\mathsf{ar}(f)}$. Because of the algebraic transparency of its interpretation, it is readily verified that all equalities trivialize over $\mathcal{H}$, in the sense that two terms t_1 and t_2 *equalize* over $\mathcal{H}$, i.e., $\mathcal{H}$ satisfies the equation $t_1 = t_2$, iff these terms are *unifiable*. When the semantic notion of equalizability and the syntactic notion of unifiability coincide on a structure, we say that such a structure enjoys the *Herbrand property*. This novel model-theoretic property is the focus of the current section. We prove that, beside Herbrand structures, which are infinite over vocabularies with functions, there are finite structures, which we call *quasi-Herbrand structures*, enjoying the Herbrand property (Theorem 2). The existence of this type of structures follows from the main result of the section, a characterization of finite equalizability via unifiability (Theorem 1).

A. Terms and Unification

Before proceeding in our study of the Herbrand property, we introduce a generalization of the standard concept of term and, accordingly, revise the notions of replacement, substitution, and unification [4]. This generalization is based on a concatenations of terms as a technical expedient yielding a substantial simplification of the notation. Intuitively, a generalized term is interpreted on a tuple of elements in a structure, with the dimension of the tuple equal to the dimension of the term, so that terms in the standard sense coincide with generalized terms of dimension one. The formal definition follows.

Definition 1 (Term). *The set of* (generalized) terms Tr_X^{Σ} *over vocabulary Σ and variable set X is the smallest set of finite words on $X \cup \Sigma_{\mathrm{f}}$ for which the following conditions hold.*

- *The empty word ε is a* term $\varepsilon \in \mathrm{Tr}_X^{\Sigma}$ *of* dimension 0, *i.e.,* $\mathsf{dim}(\varepsilon) = 0$, *having neither* occurring variables *nor* occurring functions, *i.e.,* $\mathsf{var}(\varepsilon) = \mathsf{fun}(\varepsilon) = \emptyset$.
- *Every variable $x \in X$ is a* term $x \in \mathrm{Tr}_X^{\Sigma}$ *of* dimension 1, *i.e.,* $\mathsf{dim}(x) = 1$, *having x as* occurring variable, *i.e.,* $\mathsf{var}(x) = \{x\}$, *but no* occurring functions, *i.e.,* $\mathsf{fun}(x) = \emptyset$.
- *For all functions $f \in \Sigma_{\mathrm{f}}$ and terms $t \in \mathrm{Tr}_X^{\Sigma}$ with* $\mathsf{ar}(f) = \mathsf{dim}(t)$, *the juxtaposition $ft \in \mathrm{Tr}_X^{\Sigma}$ is a* term *of* dimension $\mathsf{dim}(ft) = 1$, *having* occurring variables $\mathsf{var}(ft) = \mathsf{var}(t)$ *and* occurring functions $\mathsf{fun}(ft) = \{f\} \cup \mathsf{fun}(t)$.
- *For all terms $t_1, t_2 \in \mathrm{Tr}_X^{\Sigma}$, the concatenation $t_1t_2 \in \mathrm{Tr}_X^{\Sigma}$ is a* term *of* dimension $\mathsf{dim}(t_1t_2) = \mathsf{dim}(t_1) + \mathsf{dim}(t_2)$, *having* occurring variables $\mathsf{var}(t_1t_2) = \mathsf{var}(t_1) \cup \mathsf{var}(t_2)$ *and* occurring functions $\mathsf{fun}(t_1t_2) = \mathsf{fun}(t_1) \cup \mathsf{fun}(t_2)$.

For a term $t \in \mathrm{Tr}_X^{\Sigma}$, its *length* $|t|$ and *arity* $\mathsf{ar}(t)$ are defined, respectively, as the number of symbols and different variables $|\mathsf{var}(t)|$ occurring in it. A set of terms $T \subseteq \mathrm{Tr}_X^{\Sigma}$ is *rectangular* if $\mathsf{var}(t_1) \cap \mathsf{var}(t_2) = \emptyset$, for all $t_1, t_2 \in T$ with $t_1 \neq t_2$.

All terms in this article are generalized terms so that we freely avoid to qualify them as such. For instance, the word $xfxgyhxyzz$, where x, y, and z are variables and f, g, and h are unary, binary, and ternary functions, respectively, is a term of length 10, dimension 4, and arity 3; we also write $x, f(x), g(y, h(x, y, z)), z$ for the sake of readability.

Let $\mathrm{Ps}^{\Sigma} = (\{\varepsilon\} \cup \mathbb{N}) \cdot (\{(f, i) \in \Sigma_{\mathrm{f}} \times \mathbb{N} \,:\, i \in [\mathsf{ar}(f)]\}^*)$. The set of *positions* $\mathsf{pos}(t) \subseteq \mathrm{Ps}^{\Sigma}$ for a term t is the set of finite words defined as follows: *(i)* $\mathsf{pos}(\varepsilon) = \mathsf{pos}(x) = \{\varepsilon\}$; *(ii)* $\mathsf{pos}(ft) = \{f\} \times [\mathsf{ar}(f)] \cup \{(f, i) \cdot (p)_i \,:\, i \in [\mathsf{ar}(f)], p \in \mathsf{pos}(t)\}$; *(iii)* $\mathsf{pos}(t) = [\mathsf{dim}(t)] \cup \{i \cdot p \,:\, i \in [\mathsf{dim}(t)], p \in \mathsf{pos}((t)_i)\}$ if t has dimension greater than 1. By $t{\restriction_p}$, with $p \in \mathsf{pos}(t)$, we denote the subterm of t in position p. Notice that $t{\restriction_\varepsilon} = t$. For instance, $t{\restriction_2} = f(x)$ and $t{\restriction_{3(g,2)(h,1)}} = x$, if $t = x, f(x), g(y, h(x, y, z)), z$.

A *replacement* of variables $x_1, \ldots, x_n \in X$ with 1-dimensional terms $t_1, \ldots, t_n \in \mathrm{Tr}_X^{\Sigma}$ in a term $t \in \mathrm{Tr}_X^{\Sigma}$ is the term $t[x_1/t_1, \ldots, x_n/t_n]$ obtained by uniformly replacing every occurrence of x_i in t with t_i. A *substitution* is a function $\sigma \colon X \to \mathrm{Tr}_X^{\Sigma}$ such that *(i)* it has a support $\mathsf{sup}(\sigma) = \{x \in X \,:\, \sigma(x) \neq x\}$ of finite cardinality $|\mathsf{sup}(\sigma)| < \omega$, i.e., it is the identity almost everywhere, and *(ii)* $\mathsf{dim}(\sigma(x)) = 1$, for all variables $x \in X$. By $t^{\sigma} = t[x/\sigma(x) : x \in \mathsf{sup}(\sigma)]$, we denote the application of the substitution σ to the term t, i.e., the replacement of all variables x in the support of σ with the corresponding terms $\sigma(x)$. We also set $T^{\sigma} = \{t^{\sigma} : t \in T\}$, for any set of terms $T \subseteq \mathrm{Tr}_X^{\Sigma}$. In the following, we only consider substitutions that are *idempotent*, i.e., $(t^{\sigma})^{\sigma} = t^{\sigma}$, for every term $t \in \mathrm{Tr}_X^{\Sigma}$. This means that $\sigma(y) = y$, for all variables $y \in \mathsf{var}(\sigma(x))$ occurring in a term $\sigma(x)$ associated via σ with an arbitrary variable $x \in X$. Under the above condition, it holds that $t^{\sigma} = t[x_1/\sigma(x_1)] \cdots [x_n/\sigma(x_n)]$, where $\mathsf{sup}(\sigma) = \{x_1, \ldots, x_n\}$.

We now introduce the notion of *unification* for sets of terms, naturally lifting the unification of two terms, which represents the syntactic side of the Herbrand property.

Definition 2 (Unification). *A set of terms $T \subseteq \mathrm{Tr}_X^{\Sigma}$ is* unifiable *if it admits a* unifier, *i.e., a substitution $\mu \colon X \to \mathrm{Tr}_X^{\Sigma}$ such that ${t_1}^{\mu} = {t_2}^{\mu}$, for all $t_1, t_2 \in T$, i.e., $|T^{\mu}| = 1$.*

As an example, let $T = \{t_1, t_2, t_3\}$ be the rectangular set of 4-dimensional terms $t_1 = x_1, y_1, a(y_1, x_1), z_1$, $t_2 = x_2, s(x_2), y_2, a(z_2, s(x_2))$, and $t_3 = 0, x_3, y_3, a(y_3, z_3)$. By direct inspection, the following substitution μ unifies T:

$$\mu\colon \begin{cases} x_1, x_2 & \mapsto 0; \\ x_3, y_1, z_3 & \mapsto s(0); \end{cases} \begin{cases} y_2, y_3, z_2 & \mapsto a(s(0), 0); \\ z_1 & \mapsto a(a(s(0), 0), s(0)). \end{cases}$$

Indeed, we have $T^\mu = \{0, s(0), a(s(0), 0), a(a(s(0), 0), s(0))\}$.

The notion of substitution induces a *canonical preorder* on the set of terms defined as follows: for all $t_1, t_2 \in \mathrm{Tr}_X^\Sigma$, we say that t_1 is *at least as specific as* t_2, in symbols $t_1 \preccurlyeq t_2$, if there is a substitution $\sigma\colon X \to \mathrm{Tr}_X^\Sigma$ such that $t_1 = t_2^\sigma$. A corresponding preorder can be defined between substitutions: for all $\sigma_1, \sigma_2\colon X \to \mathrm{Tr}_X^\Sigma$, we say that σ_1 is *at least as specific as* σ_2, in symbols $\sigma_1 \preccurlyeq \sigma_2$, if $t^{\sigma_1} \preccurlyeq t^{\sigma_2}$, for all terms $t \in \mathrm{Tr}_X^\Sigma$.

The *most-general unifier* (mgu, for short) for a unifiable set of terms $T \subseteq \mathrm{Tr}_X^\Sigma$ is the substitution $\mu\colon X \to \mathrm{Tr}_X^\Sigma$, unique up to variable renaming, that is maximum w.r.t. $\preccurlyeq$ among the unifiers for T. A classic result by Robinson [30] shows that, if T only contains 1-dimensional terms, T has an mgu. This fact lifts to the case where T contains arbitrary k-dimensional terms using Martelli-Montanari rule-based unification algorithm [24] on the set of term equations $E_T = \{t_1\restriction_i = t_2\restriction_i : t_1, t_2 \in T, i \in [k]\}$, for $k \in \mathbb{N}$. By convention, a set of terms of different dimensions is not unifiable. Notice that, due to the definition, the functions employed in a term $\mu(x)$ of the mgu μ for a set of terms T are only those occurring in T, i.e., $\mathsf{fun}(\mu(x)) \subseteq \mathsf{fun}(T) = \bigcup\{\mathsf{fun}(t) : t \in T\}$, for all $x \in X$. Moreover, w.l.o.g., we can also ensure that the above property holds for variables too, i.e., $\mathsf{var}(\mu(x)) \subseteq \mathsf{var}(T) = \bigcup\{\mathsf{var}(t) : t \in T\}$, for all $x \in X$.

A *partial unification* for T is a set of terms T^σ, where σ is the substitution obtained from the solution of some subset E of all term equations derived from E_T by applying the rewriting rules prescribed by the Martelli-Montanari algorithm. For instance, two partial unifications for the set of terms $T = \{t_1, t_2, t_3\}$ of the previous example are induced by the substitutions $\sigma_1 = \{x_2 \mapsto 0;\ x_3, z_3 \mapsto s(0);\ y_2, y_3 \mapsto z_2\}$ and $\sigma_2 = \{x_2 \mapsto x_1;\ y_1, z_3 \mapsto s(x_1);\ y_3, z_2 \mapsto y_2\}$:

$$T^{\sigma_1} = \left\{\begin{matrix} x_1, y_1, a(y_1, x_1), z_1 \\ 0, s(0), z_2, a(z_2, s(0)) \end{matrix}\right\} \qquad T^{\sigma_2} = \left\{\begin{matrix} x_1, s(x_1), a(s(x_1), x_1), z_1 \\ x_1, s(x_1), y_2, a(y_2, s(x_1)) \\ 0, x_3, y_2, a(y_2, s(x_1)) \end{matrix}\right\}$$

Note that non-unifiable sets of terms might have partial unifications. For instance, consider the set $T = \{t_1, t_2\}$ of 3-dimensional terms $t_1 = x_1, y_1, a(s(x_1), y_1)$ and $t_2 = s(x_2), 0, a(x_2, s(y_2))$. A straightforward analysis of the unification rules, shows that T does not unify. Indeed, a hypothetical unifier should be a solution of both the syntactically unsolvable equations $0 = s(y_2)$ and $x = s(s(x))$, with either x_1 or x_2 in place of x, which are known in the literature as *symbol clash* and *occurs check failure* [4], respectively. However, the substitutions $\sigma_1 = \{y_1 \mapsto 0\}$ and $\sigma_2 = \{x_2 \mapsto s(x_1)\}$, solving $\{y_1 = 0\}$ and $\{x_2 = s(x_1)\}$ respectively, induce the partial unifications $T^{\sigma_1} = \{t_1{}^{\sigma_1}, t_2\}$ and $T^{\sigma_2} = \{t_1, t_2{}^{\sigma_2}\}$, where $t_1{}^{\sigma_1} = x_1, 0, a(s(x_1), 0)$ and $t_2{}^{\sigma_2} = s(s(x_1)), 0, a(s(x_1), s(y_2))$.

The partial unifications T^{σ_1} and T^{σ_2} make explicit the aforementioned reason why T does not unify. Indeed, consider the subterms of T^{σ_1} at position $p_1 = 3(a, 2)$, i.e., $t_1{}^{\sigma_1}\restriction_{p_1} = 0$ and $t_2\restriction_{p_1} = s(y_2)$. A unifier must make them syntactically identical, which is impossible because of the symbol clash $0 = s(y_2)$. Similarly, a unifier must make syntactically identical the subterms $t_1\restriction_{p_2} = x_1$ and $t_2{}^{\sigma_2}\restriction_{p_2} = s(s(x_1))$ of T^{σ_2} at position $p_2 = 1$, which is impossible because of the occurs check failure $x_1 = s(s(x_1))$. As we now formalize, these two configurations are obstructions to the attempt of unifying T.

For two 1-dimensional terms $f_1 t_1, f_2 t_2 \in \mathrm{Tr}_X^\Sigma$ with $f_1, f_2 \in \Sigma_\mathrm{f}$, we say that the pair $(f_1 t_1, f_2 t_2)$ is a *function obstruction*, if $f_1 \neq f_2$. Given a variable $x \in X$ and a 1-dimensional term $t \in \mathrm{Tr}_X^\Sigma$, we say that the pair (x, t) is a *variable obstruction* of depth $d \in \mathbb{N}$, if there is a position $p \in \mathsf{pos}(t)$ with $|p| = d$ such that $t\restriction_p = x$. A pair of terms (t_1, t_2) is an *obstruction* if it is either a function obstruction or a variable obstruction. Thus the pair $(t_1{}^{\sigma_1}\restriction_{p_1}, t_2\restriction_{p_1})$ obtained from the terms in the previous example is a function obstruction, and $(t_1\restriction_{p_2}, t_2{}^{\sigma_2}\restriction_{p_2})$ is a variable obstruction of depth 2.

The following proposition, summarizing our findings on the unifiability of a set of terms, is proved by an inductive analysis of the Martelli-Montanari unification algorithm.

Proposition 1. *For any non-unifiable set of terms $T \subseteq \mathrm{Tr}_X^\Sigma$, one of the following two conditions hold:* (i) $\mathsf{dim}(t_1) \neq \mathsf{dim}(t_2)$, *for two terms $t_1, t_2 \in T$;* (ii) *there are a partial unification U of T, two terms $t_1, t_2 \in U$, and a position $p \in \mathsf{pos}(t_1) \cap \mathsf{pos}(t_2)$ for which $(t_1\restriction_p, t_2\restriction_p)$ is an obstruction.*

B. Equalization via Unification

We introduce the notion of *equalization* for a set of terms, i.e., the semantic side of the Herbrand property. In preparation, let $\mathcal{A}$ be a Σ-structure, $\chi\colon X \to A$ an assignment, and $t \in \mathrm{Tr}_X^\Sigma$ a term. By $t^{\mathcal{A},\chi} \in A^{\mathsf{ar}(t)}$ we denote the interpretation of t in $\mathcal{A}$ under χ defined inductively on the syntactic structure of the term as follows: *(i)* $\varepsilon^{\mathcal{A},\chi} = \varepsilon$; *(ii)* $x^{\mathcal{A},\chi} = \chi(x)$, where $x \in X$; *(iii)* $(ft)^{\mathcal{A},\chi} = f^{\mathcal{A}}(t^{\mathcal{A},\chi})$, where $f \in \Sigma_\mathrm{f}$; *(iv)* $(t_1 t_2)^{\mathcal{A},\chi} = t_1{}^{\mathcal{A},\chi} t_2{}^{\mathcal{A},\chi}$. We let $T^{\mathcal{A},\chi} = \{t^{\mathcal{A},\chi} : t \in T\}$ for any $T \subseteq \mathrm{Tr}_X^\Sigma$.

Definition 3 (Equalization)**.** *A set of terms $T \subseteq \mathrm{Tr}_X^\Sigma$ is* equalizable over *a Σ-structure $\mathcal{A}$ if it admits an* equalizer over *$\mathcal{A}$, i.e., an assignment $\xi\colon X \to A$ such that $t_1{}^{\mathcal{A},\xi} = t_2{}^{\mathcal{A},\xi}$, for all $t_1, t_2 \in T$, i.e., $\left|T^{\mathcal{A},\xi}\right| = 1$. Moreover, T is* equalizable *(resp.,* finitely equalizable*) if it is equalizable over all Σ-structures (resp., finite Σ-structures).*

As an example, consider again the set $T = \{t_1, t_2, t_3\}$ of 4-dimensional terms introduced after Definition 2. Let $\mathcal{A}$ be the standard model of Presburger arithmetic, i.e., the structure over $\mathbb{N}$ interpreting 0 as the number zero, s as the successor operation, and a as the addition operation. A direct computation reveals that the assignment $\xi\colon X \to \mathbb{N}$ sending z_1 to 2, both x_1 and x_2 to 0, and otherwise identically 1, is an equalizer for T over $\mathcal{A}$. Indeed, $T^{\mathcal{A},\xi} = \{(0, 1, 1, 2)\}$. This equalizability, however, is nothing specific to Presburger arithmetic. It follows, instead, from the unifiability of T, which implies its equalization over all structures interpreting its vocabulary. Now consider the set $T = \{t_1, t_2\}$ of 3-dimensional terms from the example introduced before Proposition 1. A bit of reflection shows that T does not equalize over $\mathcal{A}$, since an equalizer must solve, in particular, the equation $0 = s(y_2)$, which is impossible in Presburger arithmetic. Notice, however, that a finite structure would suffice to avoid the equalizability of T: just interpret, on the universe $\{0, 1\}$, the function s as the identically 1 function and the constant 0 as 0. Indeed, we exploit here the fact that T has a partial unification whose terms encapsulate the function obstruction $(0, s(y_2))$. We also observed that T has a partial unification whose terms encapsulate the variable obstruction $(x_1, s(s(x_1)))$. We can similarly exploit this fact to construct a different, but still finite, structure $\mathcal{A}$ over which T does not equalize. Let $A = \{0, 1, 2\}$, where a is the projection to the first argument, i.e., $a^{\mathcal{A}}(i, j) = i$, the function s is the successor modulo 3, i.e., $s^{\mathcal{A}}(i) = i + 1 \bmod 3$, and 0 is interpreted arbitrarily. Again an equalizer must solve the equation $x_1 = s(s(x_1))$, derived from the above variable obstruction, which is impossible because $x_1 = x_1 + 2 \bmod 3$ has no solutions.

The connection between unification and equalization observed in the previous examples is no coincidence; it is a fundamental fact at the very core of our work.

Theorem 1 (Universal Herbrand Property)**.** *Let $T \subseteq \mathrm{Tr}_X^\Sigma$ be a set of terms. The following statements are equivalent:*

1) *T is unifiable;*
2) *T is equalizable;*
3) *T is finitely equalizable.*

Implication $1 \Rightarrow 2$ is an easy step, which exploits a unifier μ for T in order to construct an equalizer ξ for T. The idea is to assign an arbitrary value to all variables occurring in the terms $\mu(x)$ and then use the corresponding interpretation as the value $\xi(x)$ that the equalizer assigns to the variable x.

Lemma 1. *Every unifiable set of terms is equalizable.*

Proof. Let $T \subseteq \mathrm{Tr}_X^\Sigma$ be an unifiable set of terms having unifier $\mu\colon X \to \mathrm{Tr}_X^\Sigma$ and $\mathcal{A}$ an arbitrary Σ-structure with $\chi\colon X \to A$ as one of its assignments. By definition, we known that ${t_1}^\mu = {t_2}^\mu$, for all terms $t_1, t_2 \in T$. Thus, we obviously have $({t_1}^\mu)^{\mathcal{A},\chi} = ({t_2}^\mu)^{\mathcal{A},\chi}$. Now, consider the assignment $\xi\colon X \to A$ defined by $\xi(x) = \mu(x)^{\mathcal{A},\chi}$, for all variables $x \in X$. Then, the following claim can be stated.

Claim 1. *For all terms $t \in T$, it holds that $t^{\mathcal{A},\xi} = (t^\mu)^{\mathcal{A},\chi}$.*

At this point, it is easy to see that ξ is an equalizer for T over $\mathcal{A}$, since ${t_1}^{\mathcal{A},\xi} = ({t_1}^\mu)^{\mathcal{A},\chi} = ({t_2}^\mu)^{\mathcal{A},\chi} = {t_2}^{\mathcal{A},\xi}$. ☐

Since Implication $2 \Rightarrow 3$ is an immediate consequence of the definition, we just need to focus on Implication $3 \Rightarrow 1$.

Lemma 2. *Every finitely equalizable set of terms is unifiable.*

The proof of the above lemma, split among several more specific lemmas, proceeds by contraposition. We use Proposition 1 as a guidance to construct a structure over which non-unifiable terms do not equalize either. An obstruction (t_1, t_2) to the unifiability of a set of uniform-dimension terms T can be nested deep inside two terms of some partial unification U of T itself. The idea to construct a structure over which T does not equalize is the following: *(i)* by means of Lemmas 3 and 4, build a structure over which t_1 and t_2 do not equalize; *(ii)* use Lemma 5 to transform the previously obtained structure into one over which U does not equalize; *(iii)* use Lemma 6 to extend the latter structure for U to the desired structure for T. Notice that, since terms of different dimension have interpretations of different dimensions, the case where T does not have uniform dimension is trivial.

Fact 1 (Dimension Clash)**.** *Let $t_1, t_2 \in \mathrm{Tr}_X^\Sigma$ be two terms of different dimension. There is no Σ-structure over which t_1 and t_2 equalize.*

As already observed, function obstructions (t_1, t_2) are easy to deal with, by interpreting the outermost functions in t_1 and t_2 to two different constant values.

Lemma 3 (Function Obstruction)**.** *Let (t_1, t_2) be a function obstruction over Σ. There exists a Σ-structure of order 2 over which t_1 and t_2 do not equalize.*

Variable obstructions (x, t) are harder to deal with. In particular, the order of the structure depends on the type of nesting of x into t. For instance, let $t = f(t_1, t_2, g(h(t_3, x), t_4))$. The occurrence of x into t has depth 3 in position $p = (f, 3)(g, 1)(h, 2)$. A structure $\mathcal{A}$ over which x and t do not equalize has universe $\{0, 1\}$, where f projects to its third argument, g projects to its first argument, and h swaps its second projection, i.e., $h^{\mathcal{A}}(i, j) = 1 - j$. By direct computation, $(f(t_1, t_2, g(h(t_3, x), t_4)))^{\mathcal{A},\chi} = (g(h(t_3, x), t_4))^{\mathcal{A},\chi} = (h(t_3, x))^{\mathcal{A},\chi} = 1 - x^{\mathcal{A},\chi} \neq x^{\mathcal{A},\chi}$, for all assignments χ. This construction easily lifts to all terms t where x occurs in a position p that does not contain repeated functions. However, a more careful treatment needs the case where a function appears more than once in p. For instance, let $t = f(t_1, f(x, t_2))$, so that x is in position $p = (f, 2)(f, 1)$. Here the interpretation of f cannot be as simple as a projection or a swapping-projection, since it must depend on both its arguments. The idea is then to define f so that it behaves as a projection or a swapping-projection depending on the nesting of the argument along the position p. We therefore construct the structure $\mathcal{A}$ with universe $A = \{0, 1\} \times \{0, 1\}$, where $f^{\mathcal{A}}((i_1, j_1), (i_1, j_2)) = (j_2, 1 - i_1)$. Intuitively, f is both a swapping-projection on the first coordinate of the first argument and a projection on the second coordinate of the second argument. Then x and t do not equalize over $\mathcal{A}$, since $(f(t_1, f(x, t_2))^{\mathcal{A},\chi})_1 = (f(x, t_2)^{\mathcal{A},\chi})_2 = 1 - (x^{\mathcal{A},\chi})_1 \neq (x^{\mathcal{A},\chi})_1$, for all assignments χ. This approach is described in full generality in the proof of the following lemma.

Lemma 4 (Variable Obstruction)**.** *Let (x, t) be a variable obstruction of depth $d \in \mathbb{N}_+$ over Σ. There exists a Σ-structure of order 2^d over which x and t do not equalize.*

Proof. Since (x, t) is a variable obstruction of depth $d \in \mathbb{N}_+$, there is a non-empty position $p \in \mathsf{pos}(t)$ with length d such that $t{\upharpoonright}_p = x$. Let $P = \{p_{<i} : i \in [d]\}$ be the set of non-empty prefixes of p, where, for technical convenience, we assume that $p_{<0} = p_{<d} = p$. Moreover, let $\mathcal{A}$ be the Σ-structure having universe $A = \{0, 1\}^P$, where the interpretation of a function $f \in \Sigma_{\mathrm{f}}$ over a tuple $\mathbf{a} \in A^{\mathsf{ar}(f)}$ is defined as follows, with $i \in [d]$:

$$f^{\mathcal{A}}(\mathbf{a})(p_{<i-1}) = \begin{cases} 1 - (\mathbf{a})_j(p_{<i}), & \text{if } (p)_i = (f, j) \text{ and } i = d; \\ (\mathbf{a})_j(p_{<i}), & \text{if } (p)_i = (f, j) \text{ and } i < d; \\ a \in \{0, 1\}, & \text{otherwise.} \end{cases}$$

Observe that $|A| = 2^d$. At this point, the following claim can be stated, whose proof may be obtained via a standard induction on the depth of the variable obstruction.

Claim 2. *For all indexes $i \in [d]$, assignments $\chi\colon X \to A$, and variable obstructions (x, s) over Σ having x at position $p_{\geq i}$ (suffix of p from i) in s, it holds that $s^{\mathcal{A},\chi}(p_{<i-1}) \neq x^{\mathcal{A},\chi}(p)$.*

The required thesis immediately follows by setting $s = t$ and $i = 1$ in the above claim. Indeed, $t^{\mathcal{A},\chi}(p) \neq x^{\mathcal{A},\chi}(p)$ implies $t^{\mathcal{A},\chi} \neq x^{\mathcal{A},\chi}$, for all assignments $\chi\colon X \to A$. Hence, x and t do not equalize over $\mathcal{A}$. □

We now focus on how to percolate the non-equalizability of an obstruction to the original terms in T via a partial unification U. The following construction serves the purpose.

Construction 1 (Structure Extension)**.** *The* extension *of a Σ-structure $\mathcal{A}$ w.r.t. a position $p \in \mathrm{Ps}^{\Sigma}$ is the Σ-structure $\mathcal{A}^*$ having universe $A^* = A^P$, where $P = \{p_{<i} : i \in [0, |p|]\}$ is the set of prefixes of p including the empty one $p_{<0} = \varepsilon$ and the interpretation of a function $f \in \Sigma_{\mathrm{f}}$ over a tuple $\mathbf{a}^* \in A^{*\mathsf{ar}(f)}$ is defined as follows:*

$$f^{\mathcal{A}^*}(\mathbf{a}^*)(p) = f^{\mathcal{A}}(\mathbf{a});$$

$$f^{\mathcal{A}^*}(\mathbf{a}^*)(p_{<i-1}) = \begin{cases} (\mathbf{a}^*)_j(p_{<i}), & \text{if } (p)_i = (f, j); \\ \text{arbitrary } a \in A, & \text{otherwise}; \end{cases}$$

where $i \in [|p|]$ and $\mathbf{a} \in A^{\mathsf{ar}(f)}$ is the tuple of elements in $\mathcal{A}$ such that $(\mathbf{a})_j = (\mathbf{a}^*)_j(p)$, for each coordinate $j \in [\mathsf{ar}(f)]$.

Intuitively, the extended structure $\mathcal{A}^*$ can be seen as the Cartesian product of the original structure $\mathcal{A}$, whose values are indexed in the universe A^* by the position p, with a structure in which the interpretation of the functions is a generalized version of the projection operation, reminiscent of the approach used in Lemma 4. In the following proposition, we collect the properties of the construction needed in the next two lemmas: Item 1 captures how the value of a subterm $t{\restriction}_p$ in position p percolates to the containing term t; Item 2 describes the embedding of $\mathcal{A}$ into $\mathcal{A}^*$.

Proposition 2. *Let $\mathcal{A}^*$ be the extension of a Σ-structure $\mathcal{A}$ w.r.t. a position $p \in \mathrm{Ps}^{\Sigma}$. For all terms $t \in \mathrm{Tr}_X^{\Sigma}$ and assignments $\chi^*\colon X \to A^*$, the following hold:*

1) *$(t)_j^{\mathcal{A}^*,\chi^*}(j) = (t{\restriction}_p)^{\mathcal{A}^*,\chi^*}(p)$, if $p = j \cdot q$ with $j \in \mathbb{N}$, and otherwise $t^{\mathcal{A}^*,\chi^*}(\varepsilon) = (t{\restriction}_p)^{\mathcal{A}^*,\chi^*}(p)$, where $p \in \mathsf{pos}(t)$;*
2) *$(t)_i^{\mathcal{A}^*,\chi^*}(p) = (t)_i^{\mathcal{A},\chi}$, where $i \in [\mathsf{dim}(t)]$ and $\chi\colon X \to A$ is such that $\chi(x) = \chi^*(x)(p)$, for all $x \in X$.*

If two terms do not equalize over a structure $\mathcal{B}$, we can apply Construction 1 to $\mathcal{B}$ w.r.t. an arbitrary position p to build a new structure $\mathcal{A}$ over which no set of terms T, containing the first two at position p, can equalize. Intuitively, by Items 1 and 2 of Proposition 2, the non-equal values of the two terms in $\mathcal{B}$ flow in $\mathcal{A}$ from position p to the terms in T.

Lemma 5 (Non-Equalizability Preservation I)**.** *Let $t_1, t_2 \in T \subseteq \mathrm{Tr}_X^{\Sigma}$ be two terms and $p \in \mathsf{pos}(t_1) \cap \mathsf{pos}(t_2)$ one of their positions. If $t_1{\restriction}_p$ and $t_2{\restriction}_p$ do not equalize over a Σ-structure of order n, there is a Σ-structure of order $n^{|p|+1}$ over which T does not equalize as well.*

Proof of Lemma 5. Let $\mathcal{B}$ be a Σ-structure over which $t_1{\restriction}_p$ and $t_2{\restriction}_p$ do not equalize and $\mathcal{A}$ the extension of $\mathcal{B}$ w.r.t. the position p as defined in Construction 1. By the hypothesis on $\mathcal{B}$, it holds that $(t_1{\restriction}_p)^{\mathcal{B},\chi_{\mathcal{B}}} \neq (t_2{\restriction}_p)^{\mathcal{B},\chi_{\mathcal{B}}}$, for all assignments $\chi_{\mathcal{B}}\colon X \to B$ in $\mathcal{B}$. If t_1 and t_2 are 1-dimensional (the multi-dimensional case is similar), by Items 1 and 2 of Proposition 2, we have that $t_1{}^{\mathcal{A},\chi_{\mathcal{A}}}(\varepsilon) = (t_1{\restriction}_p)^{\mathcal{A},\chi_{\mathcal{A}}}(p) = (t_1{\restriction}_p)^{\mathcal{B},\chi_{\mathcal{B}}} \neq (t_2{\restriction}_p)^{\mathcal{B},\chi_{\mathcal{B}}} = (t_2{\restriction}_p)^{\mathcal{A},\chi_{\mathcal{A}}}(p) = t_2{}^{\mathcal{A},\chi_{\mathcal{A}}}(\varepsilon)$, for all the assignments $\chi_{\mathcal{A}}\colon X \to A$ in $\mathcal{A}$ and $\chi_{\mathcal{B}}\colon X \to B$ in $\mathcal{B}$ such that $\chi_{\mathcal{B}}(x) = \chi_{\mathcal{A}}(x)(p)$, for all variables $x \in X$. This implies that $t_1{}^{\mathcal{A},\chi_{\mathcal{A}}} \neq t_2{}^{\mathcal{A},\chi_{\mathcal{A}}}$. Consequently, t_1 and t_2 do not equalize over $\mathcal{A}$, from which we derive that T does not equalize over $\mathcal{A}$ as well. We conclude noticing that, by construction, $\mathcal{A}$ has order $n^{|p|+1}$. □

Consider now the case where a set $T[x/t]$ is obtained from the set of terms T by replacing a variable x with a term t, where both x and t occur at a position p. If $T[x/t]$ does not equalize over a structure $\mathcal{B}$, we can apply again Construction 1 to $\mathcal{B}$ w.r.t. p in order to build a structure $\mathcal{A}$ over which T cannot equalize. The idea is that, if T admits an equalizer $\xi_{\mathcal{A}}$ over $\mathcal{A}$, then the values of x and t are equal under $\xi_{\mathcal{A}}$. Thus, by Items 1 and 2 of Proposition 2, it follows that $\xi_{\mathcal{A}}$ embeds an equalizer for $T[x/t]$ over $\mathcal{B}$, which is impossible.

Lemma 6 (Non-Equalizability Preservation II)**.** *Let $t_1, t_2 \in T \subseteq \mathrm{Tr}_X^{\Sigma}$ be two terms and $p \in \mathsf{pos}(t_1) \cap \mathsf{pos}(t_2)$ one of their positions such that $t_1{\restriction}_p$ is a variable. If $T[t_1{\restriction}_p/t_2{\restriction}_p]$ does not equalize over a Σ-structure of order n, there is a Σ-structure of order $n^{|p|+1}$ over which T does not equalize as well.*

We are now ready to prove the main technical lemma which, in turn, settles the proof of the universal Herbrand property.

Proof of Lemma 2. The thesis is proved by contraposition. Let $T \subseteq \mathrm{Tr}_X^{\Sigma}$ be a non-unifiable set of terms. By Proposition 1, one of the following two points hold: *(i)* T contains two terms of different dimension; *(ii)* there is a partial unification U of T together with two terms $t_1, t_2 \in U$ having an obstruction $(t_1{\restriction}_p, t_2{\restriction}_p)$ at some position $p \in \mathsf{pos}(t_1) \cap \mathsf{pos}(t_2)$. In the first case, as observed in Fact 1, there is no Σ-structure on which T can equalize. Hence, an arbitrary finite Σ-structure suffices to satisfy the thesis. In the second case, instead, by Lemmas 3 and 4, there is a finite Σ-structure over which $t_1{\restriction}_p$ and $t_2{\restriction}_p$ do not equalize. Consequently, due to Lemma 5, there exists another finite Σ-structure over which U does not equalize as well. Now, since U is a partial unification of T, there is a finite sequence of sets of terms $T = T_1, \ldots, T_n = U$, pairs of terms $t_i^1, t_i^2 \in T_i$, and positions $p_i \in \mathsf{pos}(t_i^1) \cap \mathsf{pos}(t_i^2)$, with $i \in [n-1]$, such that *(i)* $t_i^1{\restriction}_{p_i}$ is a variable and *(ii)* $T_{i+1} = T_i[t_i^1{\restriction}_{p_i}/t_i^2{\restriction}_{p_i}]$. Now, by inductively applying Lemma 6 on every set T_i with $i \in [n-1]$, from $i = n-1$ to $i = 1$, we obtain the existence of a finite Σ-structure over which T cannot equalize. Hence, T is not finitely equalizable. □

C. Finite Quasi-Herbrand Structures

We now introduce the concept of *quasi-Herbrand structures*, a somehow succinct implementations of the Herbrand property.

Definition 4 (Quasi-Herbrand Structures)**.** *A Σ-structure $\mathcal{H}$ is* quasi Herbrand *w.r.t. a set of terms $T \subseteq \mathrm{Tr}_X^{\Sigma}$ if, for every* *subset $U \subseteq T$, it holds that U is unifiable, whenever U is equalizable over $\mathcal{H}$.*

Thanks to Implication $3 \Rightarrow 1$ of Theorem 1, we can prove the existence of a finite quasi-Herbrand structure w.r.t. every finite set of terms T. The intuition is first *(i)* to exploit that implication to construct, for all non-unifiable subsets S of T, a structure $\mathcal{A}_S$ over which S does not equalize and then *(ii)* to compute the direct product $\mathcal{H} = \prod_S \mathcal{A}_S$ for the desired result. The following statement can be actually strengthened. Indeed, an inspection of the proofs shows a double-exponential bound on the order of the structure w.r.t. $\sum_{t \in T} |t|$.

Theorem 2 (Quasi-Herbrand Structures)**.** *For every set (resp., finite set) of terms $T \subseteq \mathrm{Tr}_X^{\Sigma}$, there exists a Σ-structure (resp., finite Σ-structure) which is quasi Herbrand w.r.t. T.*

Proof. Let $N = \{S_1, \ldots, S_n\} \subseteq 2^T$ be the family of subsets of T that are non-unifiable. Due to Fact 1, we can focus on subsets of uniform dimension only. By Lemma 2, for each $i \in [n]$, there is a finite Σ-structure $\mathcal{A}_i$ over which S_i does not equalize. Consider now the direct product $\mathcal{H} = \prod_{i=1}^{n} \mathcal{A}_i$ defined as usual, i.e., *(i)* $H = \bigtimes_{i=1}^{n} A_i$ and *(ii)* $(f^{\mathcal{H}}(\mathbf{a}))_i = f^{\mathcal{A}_i}(\mathbf{b}_i)$, for all indexes $i \in [n]$, functions $f \in \Sigma_{\mathrm{f}}$ of arity $k = \mathsf{ar}(f)$, and k-tuples $\mathbf{a} \in H^k$ and $\mathbf{b}_i \in {A_i}^k$ such that $(\mathbf{b}_i)_j = ((\mathbf{a})_j)_i$, for each coordinate $j \in [k]$. We show that $\mathcal{H}$ is the desired Σ-structure. First note that, if T is finite, H is finite as well. Now, let $S \subseteq T$ be a non-unifiable set of terms. We prove that S cannot equalize over $\mathcal{H}$. By definition of N, there is an index $i \in [n]$ such that $S = S_i$. Suppose by contradiction that S admits the equalizer $\xi \colon X \to H$ over $\mathcal{H}$, i.e., ${t_1}^{\mathcal{H},\xi} = {t_2}^{\mathcal{H},\xi}$, for all terms $t_1, t_2 \in S_i$. Finally, consider the assignment $\xi_i \colon X \to A_i$ defined by $\xi_i(x) = (\xi(x))_i$, for all variables $x \in X$. As $\mathcal{H}$ is a direct product, we have that ${t_1}^{\mathcal{A}_i,\xi_i} = {t_2}^{\mathcal{A}_i,\xi_i}$. Thus, ξ_i is an equalizer for S_i over $\mathcal{A}_i$, which is impossible. □

We conclude the section providing an exponential lower bound on the order of a quasi-Herbrand structure. The idea is the following. Consider the set T_1 containing the terms $t_1^1 = f(x_1, x_1, x_3)$, $t_2^1 = f(x_1, x_2, x_2)$, and $t_3^1 = f(x_3, x_2, x_3)$ together with the set T_2 containing the terms $t_1^2 = f(c_1, x_2, x_3)$, $t_2^2 = f(x_1, c_2, x_3)$, and $t_3^2 = f(x_1, x_2, c_3)$, where c_1, c_2, and c_3 are constants. A bit of reflection reveals that, for each of the 7 nonempty subset I of $\{1, 2, 3\}$, the set of terms $U_I = \{t_i^1 \in T_1 : i \notin I\} \cup \{t_i^2 \in T_2 : i \in I\}$ is a maximal unifiable subset of $T = T_1 \cup T_2$. Then every quasi-Herbrand structure for T has order at least 7 since, otherwise, two subsets would assume the same value although their union is a non-unifiable set. The idea generalizes, even in a parsimonious vocabulary, as follows.

Theorem 3 (Minimal Quasi-Herbrand Structures)**.** *Let Σ be a vocabulary with at least two constants and one binary function. For every $n \in \mathbb{N}$, there exists a set of $2n$ terms $T \subseteq \mathrm{Tr}_X^{\Sigma}$ of dimension 1 and length $\mathrm{O}(n \log n)$ such that every Σ-structure that is quasi Herbrand w.r.t. T has order $\Omega(2^n)$.*

IV. Satisfiability in Conjunctive-Binding Logic

Motivated by a deeper understanding of the decidability of some powerful extensions of modal logic [25], [26], a family of syntactic fragments of relational first-order logic, called *binding fragments*, has been recently proposed in the literature [27]. The novel idea is to classify first-order sentences relative to the *binding forms* they admit, i.e., the Boolean combinations of variable patterns they allow, which naturally led to a diamond-shaped hierarchy of four syntactic classes (Definition 5). Relative to this classification, the finite-model property and the satisfiability problem for *conjunctive-binding logic* remained open. In this section, we successfully address both questions in the more general context where functions can occur in the vocabulary (Corollary 2 and Corollary 4). The key tool here is a syntactic characterization of satisfiability in conjunctive-binding logic (Item 4 of Theorem 5), which paves the way for the entailment results reported later in Section V.

A. Binding Forms in First-Order Logic

Let $\Sigma = \Sigma_{\mathrm{f}} \uplus \Sigma_{\mathrm{r}}$ be a vocabulary. For a relation $r \in \Sigma_{\mathrm{r}}$ and a term $t \in \mathrm{Tr}_X^{\Sigma_{\mathrm{f}}}$ with $\mathsf{ar}(r) = \mathsf{dim}(t)$, the juxtaposition of symbols rt is called a *t-atom*. Let $\mathcal{A}$ be a Σ-structure and $\chi \colon X \to A$ an assignment. We say that $\mathcal{A}$ *satisfies* rt under χ, in symbols $\mathcal{A}, \chi \models rt$, if $t^{\mathcal{A},\chi} \in r^{\mathcal{A}}$. Given this singularity, the remaining first-order syntax and semantics are defined as usual. A *t-binding Σ-form* γ is a Boolean combination of t-atoms. We let Bn_X^{Σ} denote all t-binding Σ-forms with $t \in \mathrm{Tr}_X^{\Sigma_{\mathrm{f}}}$. A *quantifier prefix* $\wp \in \mathrm{Qn}_X$ over the set of variables X is a finite sequence of the existential and universal quantifiers $\exists$ and $\forall$, each binding a different variable in X.

Definition 5 (Binding Fragments)**.** Boolean binding logic *(*BB*) is the set of first-order Σ-formulas built accordingly to the following context-free grammar, where $\wp \in \mathrm{Qn}_X$ and $\gamma \in \mathrm{Bn}_X^{\Sigma}$:*

1) $\varphi := \bot \mid \top \mid \wp\psi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi)$*;*
2) $\psi := \gamma \mid (\psi \wedge \psi) \mid (\psi \vee \psi)$.

Further binding fragments can be obtained by restricting the the above grammar as follows.

- ($\wedge$B) Conjunctive-binding logic *is the set of all formulas obtained by weakening Item 2 to $\psi := \gamma \mid (\psi \wedge \psi)$.*
- ($\vee$B) Disjunctive-binding logic *is the set of all formulas obtained by weakening Item 2 to $\psi := \gamma \mid (\psi \vee \psi)$.*
- (1B) Single-binding logic *is the set of all formulas obtained by weakening Item 2 to $\psi := \gamma$.*
- ($\forall$1B) Universal single-binding logic *is the set of all single-binding formulas in which the prefix $\wp$ is universal.*

There are interesting (graph) properties, expressible via $\forall$1B sentences, that are impossible to express in other logics, e.g., in the guarded negation fragment [5]. Among them, we have the *completeness* and the *k-edge-colorability* of a graph: $\forall x_1 \forall x_2 rt$ and $\forall x_1 \forall x_2 (\neg rt \vee \bigvee_{i=1}^{k} c_i t \wedge \bigwedge_{i=1}^{k-1} \bigwedge_{j=i+1}^{k} \neg(c_i t \wedge c_j t))$, where r denote the edge relation, c_i the i-th coloring relation, and $t = x_1, x_2$. The fact that every vertex in the graph is neither a *sink* nor a *source* or the existence of a vertex that is a *predecessor* of all vertices can be expressed in 1B: $\forall x_1 \exists x_2 rt \wedge \forall x_2 \exists x_1 rt$ and $\exists x_1 \forall x_2 rt$. The *transitivity* of the graph can be easily formalized in $\vee$B: $\forall x_1 \forall x_2 \forall x_3 (\neg r(x_1, x_2) \vee \neg r(x_2, x_3) \vee r(x_1, x_3))$. Finally, the existence of an *isolated vertex* or the fact that every vertex has a *successor* that is the predecessor of all the vertices can be expressed in $\wedge$B: $\exists x_1 \forall x_2 (\neg r(x_1, x_2) \wedge \neg r(x_2, x_1))$ and $\forall x_1 \exists x_2 \forall x_3 (r(x_1, x_2) \wedge r(x_2, x_3))$.

Every first-order sentence has an equisatisfiable canonical form, namely, its functional (Skolem) normal form [31]. The two tame binding fragments 1B and $\wedge$B enjoy an even nicer canonization, which is represented by $\forall$1B. Indeed, the standard skolemization procedure [20] applied to $\wedge$B sentences yields, in essence, $\forall$1B sentences. This fact explains the better model- and complexity-theoretic behaviors of $\wedge$B versus $\vee$B [27].

Theorem 4 ($\wedge$B Skolemization)**.** *Let φ be a $\wedge$B Σ-sentence. There is a $\forall$1B sentence φ^* over an extended vocabulary $\Sigma^* \supseteq \Sigma$, with length $|\varphi^*| = \mathrm{O}(|\varphi|)$, for which the following hold:*

1) *if $\mathcal{A}^* \models \varphi^*$ then $\mathcal{A}^* \models \varphi$, for all Σ^*-structures $\mathcal{A}^*$;*
2) *for any Σ-structure $\mathcal{A}$ with $\mathcal{A} \models \varphi$, there is a Σ^*-structure $\mathcal{A}^*$ with $\mathcal{A}^* \models \varphi^*$ that is a Σ^*-extension of $\mathcal{A}$.*

B. The Complexity of the Satisfiability Problem

The skolemization result reported in Theorem 4 allows to focus on $\forall$1B in order to study the model-theoretic properties and solve the satisfiability problem for $\wedge$B and its fragments. The key idea here is a characterization yielding a *finite quasi-Herbrand model* for every satisfiable $\forall$1B sentence.

We first need to introduce some notation. An *implicant* for a positive Boolean formula β is a subset $I \in \mathrm{Im}(\beta)$ of the propositions occurring in β such that $I \models \beta$. This notion can be easily lifted to any BB sentence φ, by considering it as positive Boolean formula over the set of sentences in prenex form $\wp\psi$ occurring in φ. For example, consider the BB sentence $\varphi = \wp_1\psi_1 \vee (\wp_2\psi_2 \wedge (\wp_3\psi_3 \vee \wp_4\psi_4))$. We have that $\mathrm{Im}(\varphi) = \{\{\wp_1\psi_1\}, \{\wp_2\psi_2, \wp_3\psi_3\}, \{\wp_2\psi_2, \wp_4\psi_4\}\}$. Given a set of prenex sentences $U = \{\wp_1\psi_1, \ldots, \wp_n\psi_n\}$, by $\mathrm{Tr}(U)$ and $\mathrm{Bn}(U)$ we denote, respectively, the set of terms and maximal quantifier-free subformulas ψ_i occurring in these formulas. For instance, $\mathrm{Tr}(U) = \{t_1, t_2, t_3\}$ and $\mathrm{Bn}(U) = \{\psi_1, \psi_2\}$, where $\psi_1 = r_1t_1 \wedge r_2t_1$, $\psi_2 = r_1t_2 \vee (r_2t_2 \wedge r_3t_3)$, and $U = \{\wp_1\psi_1, \wp_2\psi_2\}$. Finally, for an arbitrary Boolean formula over t-atoms, with $t \in \mathrm{Tr}_X^{\Sigma_\mathrm{f}}$, we indicate by $\mathsf{bool}(\gamma)$ the Boolean formula over the relations in Σ_r obtained from γ by erasing all the occurrences of the terms, e.g., $\mathsf{bool}(\psi_2) = r_1 \vee (r_2 \wedge r_3)$.

Theorem 5 ($\forall$1B Satisfiability Characterization)**.** *Let φ be a $\forall$1B Σ-sentence. The following statements are equivalent:*

1) *φ is satisfiable;*
2) *φ is finitely satisfiable;*
3) *φ admits a finite quasi-Herbrand Σ-model;*
4) *there is an implicant $I \in \mathrm{Im}(\varphi)$ such that, for all subsets $U \subseteq I$ for which the set of terms $\mathrm{Tr}(U)$ is unifiable, the Boolean formula $\bigwedge_{\gamma \in \mathrm{Bn}(U)} \mathsf{bool}(\gamma)$ is satisfiable.*

This result relies on the universal Herbrand property (Section III). Indeed, Implication 1 $\Rightarrow$ 4 directly depends on Implication 1 $\Rightarrow$ 2 of Theorem 1. Implication 4 $\Rightarrow$ 3 follows from the existence of a finite Herbrand structure, i.e., Theorem 2, which is in turn derived from Implication 3 $\Rightarrow$ 1 of Theorem 1 again. The chain 3 $\Rightarrow$ 2 $\Rightarrow$ 1 holds by definition.

Proof. (1 $\Rightarrow$ 4). Let $\mathcal{A}$ be a Σ-model of φ and $I \in \mathrm{Im}(\varphi)$ one of its implicants such that $\mathcal{A} \models \phi$, for all $\phi \in I$. Moreover, consider a subset $U \subseteq I$ for which the set of n-dimensional terms $T = \mathrm{Tr}(U)$ is unifiable, for some $n \in \mathbb{N}$. By Lemma 1, T equalizes over $\mathcal{A}$, say to the n-tuple $\mathbf{a} \in A^n$ via the equalizer $\xi \colon X \to A$, i.e., $t^{\mathcal{A},\xi} = \mathbf{a}$, for all terms $t \in T$. Observe that, because of the particular choice of the implicant I, it holds that $\mathcal{A}, \xi \models \bigwedge_{\gamma \in \mathrm{Bn}(U)} \gamma$, where φ is a $\forall$1B sentence. Now, let $\beta \colon \{r \in \Sigma_\mathrm{r} : \mathsf{ar}(r) = n\} \to \{\bot, \top\}$ be the Boolean valuation such that $\beta(r) = \top$, if $\mathbf{a} \in r^{\mathcal{A}}$, and $\beta(r) = \bot$, otherwise. Due to the above observation, it is easy to see that $\beta \models \bigwedge_{\gamma \in \mathrm{Bn}(U)} \mathsf{bool}(\gamma)$. Hence, $\bigwedge_{\gamma \in \mathrm{Bn}(U)} \mathsf{bool}(\gamma)$ is satisfiable, as required by the statement.

(4 $\Rightarrow$ 3). Assume the vocabulary $\Sigma = \Sigma_\mathrm{f} \uplus \Sigma_\mathrm{r}$ as decomposed into its functional Σ_f and relational Σ_r part and consider a finite Σ_f-structure $\mathcal{M}$ which is quasi Herbrand w.r.t. the finite rectangular set of terms $T = \mathrm{Tr}(I)$, where $I \in \mathrm{Im}(\varphi)$ is the implicant of φ derived from the hypothesis. Remind that T can be assumed to be rectangular, since φ is a $\forall$1B sentence. Moreover, the existence of such a structure is ensured by Theorem 2. We now show how to construct a Σ-extension $\mathcal{H}$ of $\mathcal{M}$ such that $\mathcal{H} \models \varphi$. For every n-tuple $\mathbf{a} \in A^n$, with $n \in \{\mathsf{ar}(r) : r \in \Sigma_\mathrm{r}\}$, let $T_\mathbf{a} = \{t \in T : \exists \chi \in A^X . t^{\mathcal{H},\chi} = \mathbf{a}\}$ be the set of n-dimension terms assuming $\mathbf{a}$ as possible value. Since T is rectangular, so is $T_\mathbf{a}$. Thus, there exists an equalizer $\xi \colon X \to A$ such that $t^{\mathcal{H},\xi} = \mathbf{a}$, for all terms $t \in T_\mathbf{a}$. Consequently, $T_\mathbf{a}$ equalizes over $\mathcal{M}$ and, so, it unifies too, since $\mathcal{M}$ is quasi Herbrand w.r.t. T. Due to the hypothesis, the Boolean formula $\bigwedge_{\gamma \in \mathrm{Bn}(U_\mathbf{a})} \mathsf{bool}(\gamma)$ is satisfiable, for all sets $U_\mathbf{a} \subseteq I$ with $\mathrm{Tr}(U_\mathbf{a}) = T_\mathbf{a}$. Therefore, there exists a Boolean valuation $\beta_\mathbf{a} \colon \{r \in \Sigma_\mathrm{r} : \mathsf{ar}(r) = n\} \to \{\bot, \top\}$ for which it holds that $\beta_\mathbf{a} \models \bigwedge_{\gamma \in \mathrm{Bn}(U_\mathbf{a})} \mathsf{bool}(\gamma)$. For an arbitrarily chosen such $\beta_\mathbf{a}$, the interpretation of each relation $r \in \Sigma_\mathrm{r}$ in $\mathcal{H}$ is now defined as $r^{\mathcal{H}} = \{\mathbf{a} \in A^{\mathsf{ar}(r)} : \beta_\mathbf{a}(r) = \top\}$. At this point, the following statement can be claimed.

Claim 3. *For all sentences $\phi \in I$, it holds that $\mathcal{H} \models \phi$.*

Since *(i)* $\mathcal{H}$ is finite, being a Σ-extension of $\mathcal{M}$, and *(ii)* $\mathcal{H} \models \varphi$, as I is an implicant of φ, the thesis clearly follows. ☐

The *finite-model property* for $\forall$1B is a consequence of Equivalence 1 $\Leftrightarrow$ 2 of the above result. It is, actually, a *small-model property*, due to the double-exponential bound on the order of a quasi-Herbrand structure observed before Theorem 2.

Corollary 1 ($\forall$1B FMP)**.** $\forall$1B *enjoys the finite-model property.*

Thanks to Theorem 4, the same property holds for $\wedge$B.

Corollary 2 ($\wedge$B FMP)**.** $\wedge$B *enjoys the finite-model property.*

The decidability of the *(finite) satisfiability problem* for $\forall$1B follows by observing that the property described in Item 4 of the previous theorem can be checked by an $\exists\forall\exists$-alternating Turing machine in polynomial time. For the lower bound we appeal to a hardness result from [27].

Corollary 3 ($\forall$1B SAT)**.** *The (finite) satisfiability problem for* $\forall$1B *is* Σ_3^P*-complete.*

Again by Theorem 4, the satisfiability problem for $\forall$1B and $\wedge$B are linear-time interreducible. Hence, the following holds.

Corollary 4 ($\wedge$B SAT)**.** *The (finite) satisfiability problem for* $\wedge$B *is* Σ_3^P*-complete.*

V. Entailment in Conjunctive-Binding Logic

As opposed to satisfiability, the entailment problem for $\wedge$B is undecidable, since it subsumes the validity problem for $\wedge$B and, so, the complement of the satisfiability problem for $\vee$B, which has already been proved to be undecidable [27]. This is because the negation of a $\vee$B sentence is a $\wedge$B one. The same observation holds for the finite version of the question.

Theorem 6 ($\wedge$B Entailment)**.** *The (finite) entailment problem for* $\wedge$B *is undecidable.*

It is readily noticed, however, that certain entailment problems in the vicinity of $\wedge$B are decidable as a direct consequence of the decidability of its satisfiability. A first example of this fact is obtained by focusing on 1B. Indeed, differently from $\wedge$B, 1B is closed under negation, thus, the problem of interest linearly reduces to the unsatisfiability problem for 1B.

Theorem 7 (1B Entailment)**.** *The (finite) entailment problem for* 1B *is* Π_3^{P}*-complete.*

A second example is the entailment from a $\wedge$B sentence φ to a $\vee$B one θ, which is orthogonal to the same problem for $\wedge$B. By the above observation on the negation of $\vee$B sentences, the verification of $\phi \models \theta$ immediately reduces to the unsatisfiability check for $\phi \wedge \neg\theta$.

Theorem 8 ($\wedge$B $\models$ $\vee$B Entailment)**.** *The (finite) entailment problem from* $\wedge$B *to* $\vee$B *is* Π_3^{P}*-complete.*

A fragment of $\wedge$B where its decidable satisfiability still helps, but not directly, is *positive Herbrand logic* (PH), i.e., $\wedge$B without negations and disjunctions [7]. It is especially intriguing because its relational fragment coincides with the language of *quantified conjunctive queries* (QCQ) in database theory [13]. The main results of this section are a complexity classification of the (finite) entailment problem in PH (Theorem 9), whose relevance in database theory is outlined in the introduction, and a Chandra-Merlin-style theorem for QCQ (Theorem 10), whose impact is mentioned in the discussions.

A. The Complexity of Finite Entailment in PH

We now prove that *(i)* general and finite entailment in PH coincide and that *(ii)* the complexity of these two problems is NPTIME-complete.

As a guidance throughout the technical development of the section, we give a running example of an entailment in PH, namely, the positive instance $\varphi \models \vartheta$, where

$$\begin{aligned} \varphi &= \forall x_1 \exists y_1 \forall z_1\, \psi_1 \wedge \forall w_1 \forall v_1 \exists u_1\, \psi_2, \\ \psi_1 &= r(x_1, y_1) \wedge p(y_1, a(z_1, y_1), z_1), \\ \psi_2 &= r(w_1, u_1) \wedge q(u_1, v_1), \\ \vartheta &= \forall x_2 \exists y_2 \forall z_2 \exists w_2 \exists v_2\, \eta, \\ \eta &= r(x_2, y_2) \wedge r(y_2, v_2) \wedge q(v_2, z_2) \wedge p(y_2, w_2, z_2). \end{aligned}$$

As a first manipulation, reminiscent of a similar reduction applied while studying the satisfiability of $\wedge$B, we cast the PH antecedent of an entailment instance to a universal PH ($\forall$PH) sentence by skolemization (Theorem 4), which allows to focus on the quantifier prefix of the consequent only.

Corollary 5 (Skolemized Entailment)**.** *Let φ and ϑ be two Σ-sentences, the first being* PH*. There exists a* $\forall$PH *sentence φ^* over an extended vocabulary $\Sigma^* \supseteq \Sigma$, with length $|\varphi^*| = \mathrm{O}(|\varphi|)$, such that $\varphi \models_{(\mathrm{fin})} \vartheta$ iff $\varphi^* \models_{(\mathrm{fin})} \vartheta$.*

Continuing our running example, the action of Corollary 5 on φ gives $\varphi^* = \varphi_1 \wedge \varphi_2 \wedge \varphi_3 \wedge \varphi_4$, where $\varphi_1 = \forall x_1\, r(x_1, f_1(x_1))$, $\varphi_2 = \forall x_1 \forall z_1\, p(f_1(x_1), a(z_1, f_1(x_1)), z_1)$, $\varphi_3 = \forall w_1 \forall v_1 r(w_1, f_2(w_1, v_1))$, and $\varphi_4 = \forall w_1 \forall v_1 q(f_2(w_1, v_1), v_1)$. Obviously, we obtain that $\varphi^* \models \vartheta$.

A second crucial manipulation, which we call *binding canonization*, allows to reduce $\forall$PH-to-PH entailment instances (φ, ϑ) to $\forall$1B ones. Roughly, a "canonized" term t is mined from the consequent ϑ, by concatenating all terms in an arbitrary ordering of the atoms, and correspondingly, a family of "canonized" terms S respecting the same sequence of atoms is mined from the antecedent φ (Definition 6). The canonized terms are then posed in a fresh relation in such a way that the entailment between the original conjunctions of atoms is encoded (Lemma 7). The details follow.

We use a countable family of renaming substitutions $\sigma_{i,j} : X \to X$ to render any set of terms rectangular. Formally, $t_1{}^{\sigma_{i,j}}$ and $t_2{}^{\sigma_{i',j'}}$ are rectangular, for all terms $t_1, t_2 \in \mathrm{Tr}_X^{\Sigma}$ and pairs of indexes $(i,j), (i',j') \in \mathbb{N} \times \mathbb{N}$ such that $(i,j) \neq (i',j')$.

Definition 6 (Binding Canonization)**.** *Let $Z_1 \subseteq \mathrm{At}_X^{\Sigma}$ and $Z_2 = \{r_1 t_{1,1}, \ldots, r_1 t_{1,m_1}, \ldots, r_n t_{n,1}, \ldots, r_n t_{n,m_n}\} \subseteq \mathrm{At}_X^{\Sigma}$ be two finite sets of atoms. A* binding canonization *for Z_1 w.r.t. Z_2 is a pair (t, S) of a term $t \in \mathrm{Tr}_X^{\Sigma}$ and a set of terms $S \subseteq \mathrm{Tr}_X^{\Sigma}$ defined as follows:*

- $t = \prod_{i=1}^{n} \prod_{j=1}^{m_i} t_{i,j}$;
- $S = \left\{ \prod_{i=1}^{n} \prod_{j=1}^{m_i} s_{i,j}{}^{\sigma_{i,j}} : r_i s_{i,j} \in Z_1 \right\}$.

The canonized term t in the above definition encodes the information about all relations and their bindings in the consequent, while the terms in S represent the canonized terms in the antecedent which comply with the sequence of atoms dictated by t.

A binding canonization of the atoms in Z_1 of the skolemized antecedent φ^* w.r.t. the atoms in Z_2 of the consequent ϑ for the running example is given in Table I, where $S = \{s_1, s_2, s_3, s_4\}$.

Lemma 7 (Entailment Canonization)**.** *Let $\varphi_1 = \wp_1 \psi_1$ and $\varphi_2 = \wp_2 \psi_2$ be two* PH *Σ-sentences, the first being universal. Moreover, let (t, S) be the binding canonization of $\mathrm{At}(\varphi_1)$ w.r.t. $\mathrm{At}(\varphi_2)$ and $r \notin \Sigma_{\mathrm{r}}$ a fresh relation with $\mathsf{ar}(r) = \mathsf{dim}(t)$. Then, $\varphi_1 \models_{(\mathrm{fin})} \varphi_2$ iff $\bigwedge_{s \in S} \wp_1 r s \models_{(\mathrm{fin})} \wp_2 r t$.*

We are now ready to settle the complexity of the general and finite entailment problems for PH and, thus, for QCQ.

Theorem 9 (PH Entailment)**.** *The (finite) entailment problem for* PH *is* NPTIME*-complete.*

Proof. We first prove the complexity result and than show that the finite and general version of the problem coincide. The lower bound directly follows from the NPTIME-hardness of the containment problem for classic (Boolean) conjunctive queries [12]. The NPTIME-membership is derived by observing that a PH-entailment instance of the form $\wp_1(\alpha_1 \wedge \cdots \wedge \alpha_m) \models \wp_2(\beta_1 \wedge \cdots \wedge \beta_n)$ holds iff the sentence $\wp_1(\alpha_1 \wedge \cdots \wedge \alpha_m) \wedge \overline{\wp_2}(\neg\beta_1 \vee \cdots \vee \neg\beta_n)$ is unsatisfiable, which in its turn is equivalent to the fact that the set of clauses $\{\alpha_1, \ldots, \alpha_m, \neg\beta'_1 \vee \cdots \vee \neg\beta'_n\}$ is refutable via first-order unit resolution [11], where $\wp_1$ is assumed to be universal, thanks to Corollary 5, and $\beta'_j = \beta_j^{\sigma_{\overline{\wp_2}}}$ are obtained via skolemization. In case of

	r		r		q		p		
$t =$	$x_2,$	$y_2,$	$y_2,$	$v_2,$	$v_2,$	$z_2,$	$y_2,$	$w_2,$	z_2
$t^* =$	$c_{x_2},$	$y_2,$	$y_2,$	$v_2,$	$v_2,$	$f_{z_2}(y_2),$	$y_2,$	$w_2,$	$f_{z_2}(y_2)$
$s_1 =$	$x_{11},$	$f_1(x_{11}),$	$x_{12},$	$f_1(x_{12}),$	$f_2(w_{21}, v_{21}),$	$v_{21},$	$f_1(x_{31}),$	$a(z_{31}, f_1(x_{31})),$	z_{31}
$s_2 =$	$x_{11},$	$f_1(x_{11}),$	$w_{12},$	$f_2(w_{12}, v_{12}),$	$f_2(w_{21}, v_{21}),$	$v_{21},$	$f_1(x_{31}),$	$a(z_{31}, f_1(x_{31})),$	z_{31}
$s_3 =$	$w_{11},$	$f_2(w_{11}, v_{11}),$	$x_{12},$	$f_1(x_{12}),$	$f_2(w_{21}, v_{21}),$	$v_{21},$	$f_1(x_{31}),$	$a(z_{31}, f_1(x_{31})),$	z_{31}
$s_4 =$	$w_{11},$	$f_2(w_{11}, v_{11}),$	$w_{12},$	$f_2(w_{12}, v_{12}),$	$f_2(w_{21}, v_{21}),$	$v_{21},$	$f_1(x_{31}),$	$a(z_{31}, f_1(x_{31})),$	z_{31}
$u =$	$c_{x_2},$	$f_1(c_{x_2}),$	$\cdots,$	$f_2(f_1(c_{x_2}), f_{z_2}(f_1(c_{x_2}))),$	$\cdots,$	$f_{z_2}(f_1(c_{x_2})),$	$\cdots,$	$a(f_{z_2}(f_1(c_{x_2})), f_1(c_{x_2})),$	$\cdots$

Table I: Binding canonization for $Z_1 = \{r(x_1, f_1(x_1)), r(w_1, f_2(w_1, v_1)), q(f_2(w_1, v_1), v_1), p(f_1(x_1), a(z_1, f_1(x_1))\}$ w.r.t. $Z_2 = \{r(x_2, y_2), r(y_2, v_2), q(v_2, z_2), p(y_2, w_2, z_2)\}$.

a positive instance, a refutation of length polynomial in the size of the instance necessarily exists. Indeed, it consists of a sequence of n applications of the resolution rule (actually, modus ponens) deriving, say, $(\neg\beta_2' \vee \cdots \vee \neg\beta_n')^{\mu_1}$, where μ_1 is an mgu of β_1' and some α_{i_1}, then $(\neg\beta_3' \vee \cdots \vee \neg\beta_n')^{\mu_1\mu_2}$, where μ_2 is an mgu of ${\beta_2'}^{\rho_1}$ and some (not necessarily different) α_{i_2}, and so on until ${\beta_n'}^{\mu_1 \cdots \mu_n}$ is resolved.

We now show that entailment and finite entailment in PH coincide. For the nontrivial direction, let ϕ_1 and ϕ_2 be two PH sentences such that $\phi_1 \not\models \phi_2$. By Corollary 5, let ϕ_1^* be the $\forall$PH sentence such that $\phi_1 \models \phi_2$ iff $\phi_1^* \models \phi_2$. Thus, $\phi_1^* \not\models \phi_2$. By applying Lemma 7 with ϕ_1^* in place of φ_1 and ϕ_2 in place of φ_2, we obtain that $\bigwedge_{s \in S} \wp_1 rs \not\models \wp_2 rt$. Therefore, $\bigwedge_{s \in S} \wp_1 rs \wedge \overline{\wp_2} \neg rt$ is satisfiable and indeed, by Corollary 2, finitely satisfiable. Hence, $\bigwedge_{s \in S} \wp_1 rs \not\models_{\mathsf{fin}} \wp_2 rt$. Again by Lemma 7, we have $\phi_1^* \not\models_{\mathsf{fin}} \phi_2$, which implies $\phi_1 \not\models_{\mathsf{fin}} \phi_2$, due to Corollary 5. □

B. A Chandra-Merlin Theorem for QCQ

We finally provide a homomorphism-based characterization of entailment problem for QCQ in the style of the Chandra-Merlin theorem for CQ [12].

We illustrate the idea via the running example. We first compute an interpolant φ' between φ and ϑ, which is obtained from φ by choosing the variables y_1 and z_1, respectively, as replacements for the universal variables w_1 and v_1 in ψ_2. In formulas, $\varphi' = \forall x_1 \exists y_1 \forall z_1 \exists u_1 (r(x_1, y_1) \wedge r(y_1, u_1) \wedge q(u_1, z_1) \wedge p(y_1, a(z_1, y_1), z_1))$. Next, we analyze the mappings from the variables of ϑ to those of φ', roughly aiming at respecting both the quantifier prefixes and the relations. These are generalizations of the standard notion of homomorphism used to characterize containment of CQ's [12]. In the case under scrutiny, ϑ can be derived from φ' via a variable renaming (where indexes 1 become 2) and the introduction of an existential variable w_2 as a replacement for the term $a(z_2, y_2)$. Since $a(z_2, y_2)$ depends on the universal variable z_2 and the existential variable y_2, the variable w_2 is quantified after both z_2 and y_2 in the quantifier prefix of ϑ.

We now start the technical development by introducing the key notion of *definable skolemization* which, given a quantifier prefix, returns a set of Skolem functions in a syntactically explicit fashion. This is done by assigning each existential variable to terms implementing their Skolem functions. Given a quantifier prefix $\wp \in \mathrm{Qn}_X$, we denote by $\mathsf{var}(\wp)$, $\mathsf{var}_\forall(\wp)$, and $\mathsf{var}_\exists(\wp)$, respectively, the set of variables, universal variables, and existential variables quantified in $\wp$. For a variable $x \in \mathsf{var}(\wp)$, the notation $\wp_{>x}$ (resp., $\wp_{<x}$) is the suffix (resp., prefix) of $\wp$ containing the variables that occur at the right (resp., left) of x in $\wp$.

Definition 7 (Definable Skolemization)**.** *A* definable skolemization *for a quantifier prefix* $\wp \in \mathrm{Qn}_X$ *is a substitution* $\tau\colon X \to \mathrm{Tr}_X^\Sigma$ *satisfying the following conditions:*

1) τ *injectively maps universal variables of* $\wp$ *to variables in* X*, i.e.,* $\tau(\mathsf{var}_\forall(\wp)) \subseteq X$ *and* $\tau\restriction_{\mathsf{var}_\forall(\wp)}$ *is injective;*
2) τ *maps existential variables* $x \in \mathsf{var}_\exists(\wp)$ *of* $\wp$ *to terms whose occurring variables avoid* $\tau(\mathsf{var}_\forall(\wp_{>x}))$*, i.e.,* $\mathsf{var}(\tau(x)) \cap \tau(\mathsf{var}_\forall(\wp_{>x})) = \emptyset$.

A definable skolemization for the quantifier prefix $\forall x_2 \exists y_2 \forall z_2 \exists w_2 \exists v_2$ of ϑ in our running example is represented by the substitution τ acting as $x_2 \mapsto x$, $y_2 \mapsto f_1(x)$, $z_2 \mapsto w$, $w_2 \mapsto a(w, f_1(x))$, $v_2 \mapsto f_2(f_1(x), w)$. Note that τ satisfies Definition 7 even swapping the quantification on x_2 to existential. However, τ violates the second condition w.r.t. $\forall x_2 \exists y_2 \exists w_2 \forall z_2 \exists v_2$, as the variable w, which is the image of the universal variable z_2 on the right of w_2, occurs in the term $\tau(w_2) = a(w, f_1(x))$.

We are finally ready to introduce the notion of *Skolem homomorphism*, a syntactic characterization of the $\forall$PH-to-QCQ entailment relation (Theorem 10). Restricted to conjunctive queries, it boils down to the standard notion of homomorphism. By $\mathrm{At}(\varphi)$ we denote the set of atoms occurring in the PH sentence φ.

Definition 8 (Skolem Homomorphism)**.** *Let* φ_1 *and* $\varphi_2 = \wp_2 \psi_2$ *be two* PH Σ*-sentences, the first being universal. A* Skolem homomorphism *from* φ_2 *to* φ_1 *is a definable skolemization* $\tau\colon X \to \mathrm{Tr}_X^\Sigma$ *for* $\wp_2$ *such that, for each atom* $rt_2 \in \mathrm{At}(\varphi_2)$ *in* φ_2*, there is an atom* $rt_1 \in \mathrm{At}(\varphi_1)$ *in* φ_1 *such that* ${t_2}^\tau \preccurlyeq t_1$.

Intuitively, Skolem homomorphisms map variables in the consequent to terms in the antecedent in such a way that the variable dependencies in the prefix are respected while, simultaneously, the interpretation of the atoms in the consequent get covered by that of the atoms in the antecedent. Continuing our example, the definable skolemization τ above is a Skolem homomorphism from ϑ to the skolemization φ^* of the antecedent φ, as shown by the following table.

	$\mathrm{At}(\vartheta)^\tau$		$\mathrm{At}(\varphi^*)$
r	$x, f_1(x)$	$\preccurlyeq$	$x_1, f_1(x_1)$
	$f_1(x), f_2(f_1(x), w)$	$\preccurlyeq$	$w_1, f_2(w_1, v_1)$
q	$f_2(f_1(x), w), w$	$\preccurlyeq$	$f_2(w_1, v_1), v_1$
p	$f_1(x), a(w, f_1(x)), w$	$\preccurlyeq$	$f_1(x_1), a(z_1, f_1(x_1)), z_1$

Theorem 10 ($\forall$PH$\models$QCQ Entailment Characterization)**.** *Let* φ_1 *be a* $\forall$PH *sentence and* φ_2 *a* QCQ *one. Then,* $\varphi_1 \models \varphi_2$ *iff* φ_2 *admits a Skolem homomorphism to* φ_1.

The easy (⇐) direction of Theorem 10 uses the Skolem homomorphism to extend every model of the antecedent with the interpretations of the Skolem functions for the existential variables in the quantifier prefix of the consequent. The extended structure, so, is a model of the skolemization of the consequent. Therefore, by Theorem 4, the original structure satisfies the consequent itself.

Lemma 8. *Let φ_1 and φ_2 be two* PH *sentences, the first being universal, such that φ_2 admits a Skolem homomorphism to φ_1. Then, $\varphi_1 \models \varphi_2$.*

The hard (⇒) direction of Theorem 10 starts by reducing a positive instance of the $\forall$PH$\models$QCQ entailment problem to a positive instance of the entailment problem for $\forall$1B (Definition 6 and Lemma 7). The characterization of satisfiability for $\forall$1B (Theorem 5) yields two unifiable critical terms s^* and t^* from the antecedent and the consequent, respectively, whose mgu encodes, in essence, the required Skolem homomorphism (Lemma 10).

Lemma 9. *Let φ_1 be a $\forall$PH sentence and φ_2 a* QCQ *one, such that $\varphi_1 \models \varphi_2$. Then, φ_2 admits a Skolem homomorphism to φ_1.*

The *Skolem substitution* for a quantifier prefix $\wp \in \mathrm{Qn}_X$ is the substitution $\sigma_\wp \colon X \to \mathrm{Tr}_X^{\Sigma^* \setminus \Sigma}$, with $\Sigma^* \supseteq \Sigma$, assigning each existential variable $x \in \mathsf{var}_\exists(\wp)$ in $\wp$ with a term $\sigma_\wp(x) = f_x\, z_1 \cdots z_{n_x}$, where $\{z_1, \ldots, z_{n_x}\} = \mathsf{var}_\forall(\wp_{<x})$ is the set of variables from which x depends in $\wp$.

If $\varphi_1 \models \varphi_2$, then by Lemma 7 the conjunction $\bigwedge_{s\in S} \wp_1 rs \wedge \overline{\wp_2}\neg rt$ is unsatisfiable, where $\overline{\wp_2}$ is the quantification prefix dual of $\wp_2$. An appeal to the satisfiability characterization for $\forall$1B (Theorem 5) yields a term $s^* \in S$ that unifies with the term $t^* = t^{\sigma_{\overline{\wp_2}}}$, where $\sigma_{\overline{\wp_2}}$ is the Skolem substitution obtained by the skolemization of $\overline{\wp_2}$. We report in Table I the term t^* for our running example. In this case s^* is the term s_2, and the term u is the result of the unification of t^* and s^*. The following key lemma identifies the properties of the mgu between t^* and s^*, that allow to derive the required Skolem homomorphism.

Lemma 10 (Quantified Unification). *Let $s \in \mathrm{Tr}_X^{\Sigma}$ and $t \in \mathrm{Tr}_X^{\emptyset}$ be two rectangular terms and $\sigma_\wp \colon X \to \mathrm{Tr}_X^{\Sigma^*\setminus\Sigma}$ a Skolem substitution for a quantifier prefix $\wp \in \mathrm{Qn}_X$ with $\mathsf{var}(\wp) = \mathsf{var}(t)$ and $\Sigma^* \supseteq \Sigma$. If s and $t^{\sigma_\wp}$ unify, then they have a* mgu *$\mu \colon X \to \mathrm{Tr}_X^{\Sigma^*}$ satisfying the following properties:*

1) *for every existential variable $x \in \mathsf{var}_\exists(\wp)$ of $\wp$, there is a variable $y \in \mathsf{var}(s)$ in s such that $\mu(y) = \sigma_\wp(x)^\mu$;*
2) *for every universal variable $x_1 \in \mathsf{var}_\forall(\wp)$ and existential variable $x_2 \in \mathsf{var}_\exists(\wp)$ of $\wp$ with $x_1 <_\wp x_2$, it holds that $\mathsf{fun}(\mu(x_1)) \cap \mathsf{fun}(\sigma_\wp(x_2)) = \emptyset$.*

We can now prove the main technical lemma which, in turn, settles the proof of the entailment characterization.

Proof of Lemma 9. Assume $\varphi_1 = \wp_1\psi_1$ and $\varphi_2 = \wp_2\psi_2$. Since $\varphi_1 \models \varphi_2$, by Lemma 7, we have that $\bigwedge_{s\in S} \wp_1 rs \models \wp_2 rt$, where (t, S) is the binding canonization of $\mathrm{At}(\psi_1)$ w.r.t. $\mathrm{At}(\psi_2)$ and $r \notin \Sigma_{\mathrm{r}}$ is a fresh relation of suitable arity. Due to the canonic entailment, it is immediate to see that $\bigwedge_{s\in S} \wp_1 rs \wedge \neg\wp_2 rt$ is unsatisfiable. So, by elementary logic, the single-binding sentence $\varphi = \bigwedge_{s\in S} \wp_1 rs \wedge \overline{\wp_2}\neg rt$ is unsatisfiable as well, where $\overline{\wp_2}$ is the quantifier prefix dual of $\wp_2$. By applying Theorem 4 to φ via the Skolem substitution $\sigma_{\overline{\wp_2}} \colon X \to \mathrm{Tr}_X^{\Sigma^*\setminus\Sigma}$ for $\overline{\wp_2}$, we obtain the unsatisfiable universal single-binding sentence $\varphi^* = \bigwedge_{s\in S} \wp_1 rs \wedge \overline{\wp_2^*}\neg rt^*$, where $t^* = t^{\sigma_{\overline{\wp_2}}}$ and $\wp_2^*$ is the quantifier prefix derived from $\wp_2$ by removing its universal variables. Notice that the skolemization procedure does not change the first part of the sentence, since $\wp_1$ is already universal. At this point, thanks to the satisfiability characterization of universal single-binding logic given in Theorem 5, we can prove the following claim.

Claim 4. *There is a term $s^* \in S$ that unifies with t^*.*

Let $\mu \colon X \to \mathrm{Tr}_X^{\Sigma^*}$ be the mgu for s^* and t^* given by Lemma 10. Moreover, consider the derived function $\rho \colon \mathsf{var}_\forall(\wp_2) \to X$ such that $\mu(\rho(x)) = \sigma_{\overline{\wp_2}}(x)^\mu$, for all $x \in \mathsf{var}_\forall(\wp_2)$. The existence of such a function is ensured by Item 1 of the same lemma, where s^*, t, $\sigma_{\overline{\wp_2}}$, and $\overline{\wp_2}$ are in place of s, t, $\sigma_\wp$, and $\wp$, respectively. Intuitively, ρ implements the behavior of the candidate Skolem homomorphism τ we are looking for relative to the universal variables of the prefix $\wp_2$. The following claim can be proved by a direct inspection of the definition of ρ.

Claim 5. *The function ρ is injective.*

We can now focus on the definition of τ relative to the variables $x \in \mathsf{var}_\exists(\wp_2)$ that are existential in the prefix $\wp_2$. To do this, we make use of the corresponding term $\mu(x)$ in the unifier μ. However, because of the substitution $\sigma_{\overline{\wp_2}}$, it may be the case that t^* contains Skolem symbols in $\Sigma^*\setminus\Sigma$. Therefore, the unifier might contain such spurious symbols as well, which we need to replace by the corresponding variable. This replacing is put in practice by means of the lifting function $\mathsf{lift} \colon \mathrm{Tr}_X^{\Sigma^*} \to \mathrm{Tr}_X^{\Sigma}$ defined as follows: *(i)* $\mathsf{lift}(\varepsilon) = \varepsilon$; *(ii)* $\mathsf{lift}(x) = x \in X$; *(iii)* $\mathsf{lift}(fu) = f\mathsf{lift}(u)$, if $f \in \Sigma$; *(iv)* $\mathsf{lift}(u_1u_2) = \mathsf{lift}(u_1)\mathsf{lift}(u_2)$; *(v)* $\mathsf{lift}(\sigma_{\overline{\wp_2}}(x)^\mu) = \mathsf{lift}(f_x t_x{}^\mu) = \rho(x)$, for $x \in \mathsf{var}_\forall(\wp_2)$, where $f_x \in \Sigma^* \setminus \Sigma$. Intuitively, the function lift parses a term without modifying it until an instance $\sigma_{\overline{\wp_2}}(x)^\mu$ of a Skolem term $\sigma_{\overline{\wp_2}}(x)$ for an arbitrary universal variable $x \in \mathsf{var}_\forall(\wp_2)$ is found and replaced by the corresponding variable $\rho(x)$ occurring in the term s^*. The function $\tau \colon X \to \mathrm{Tr}_X^{\Sigma}$ is then defined as follows:

$$\tau(x) = \begin{cases} \rho(x), & \text{if } x \in \mathsf{var}_\forall(\wp_2); \\ \mathsf{lift}(\mu(x)), & \text{if } x \in \mathsf{var}_\exists(\wp_2); \\ x, & \text{otherwise.} \end{cases}$$

Thanks to the definition of the function ρ, Claim 5, and Item 2 of Lemma 10, we can prove that τ satisfies the properties required by Definition 7, as stated in the following claim.

Claim 6. *The function τ is a definable skolemization for $\wp_2$.*

To conclude the proof, it only remains to show that, for all atoms $r_i t_{i,j} \in \{r_1 t_{1,1}, \ldots, r_n t_{n,m_n}\} = \mathrm{At}(\psi_2)$ in φ_2, there is an atom $r_i s_{i,j} \in \mathrm{At}(\psi_1)$ in φ_1 for which $t_{i,j}{}^\tau \preccurlyeq s_{i,j}$ holds. The latter inequality requires the existence of a substitution $\sigma \colon X \to \mathrm{Tr}_X^{\Sigma}$ such that $t_{i,j}{}^\tau = s_{i,j}{}^\sigma$. First notice that, due to the binding canonization, we can choose $s_{i,j}$ to be the subterm of s^* corresponding to the position of $t_{i,j}$ in t. By Claim 4, we can also observe that $s_{i,j}$ and $t_{i,j}{}^{\sigma_{\overline{\wp_2}}}$ unify via μ. Now, consider the substitution σ defined as follows: $\sigma(x) = \mathsf{lift}(\mu(x))$, if $x \in \mathsf{var}(s^*)$, and $\sigma(x) = x$, otherwise. Thanks to the above observation, we can prove that σ is precisely the witness for the required inequality ${t_{i,j}}^{\tau} \preccurlyeq s_{i,j}$, as stated in the following.

Claim 7. ${t_{i,j}}^{\tau} = {s_{i,j}}^{\sigma}$*, for all indexes* $i \in [n]$ *and* $j \in [m_i]$*.*

By summing up, it is immediate to see that the function τ satisfies the property stated in Definition 8. Hence, φ_2 admits a Skolem homomorphism to φ_1. □

Note that the computation of a Skolem homomorphism, if any, can be done in nondeterministic polynomial time. As observed in the above proof, a Skolem homomorphism can be extracted from the mgu of the two critical terms s^* and t^* via a polynomial-time transformation. The term t^* is linear-time computable from the set of atoms of the consequent. Finally, s^* is a term to guess in the set S whose length, being equal to the one of t^*, is linear in the length of the consequent.

VI. Discussion

Placing the (finite) entailment problem for PH within NPTime not just closes the wide complexity-theoretic gap between the NPTime-hardness of CQ containment [12] and the 3ExpTime-membership of QCQ containment [13], it actually pushes the problem in the range of practically feasible computation, e.g., via SAT solvers. Interestingly, resolution-based first-order provers, once executed on QCQ-containment instances, implement in essence the behavior dictated by the proposed extension of the Chandra-Merlin theorem.

On the theoretical side, we conjecture that Skolem homomorphisms actually characterize entailment in PH. In fact, we already know that such homomorphisms are sufficient for entailment in PH (Lemma 8) and necessary for entailment from PH to QCQ (Lemma 9). A careful inspection of our argument reveals that a generalization of Lemma 10 would suffice to turn the conjecture into a theorem. To the best of our knowledge, the decidability of the more general entailment in full Herbrand logic is unsettled. Its complexity ranges between the Π_2^{P}-hardness of the containment problem for CQ with negation [29] and the undecidability of the entailment in $\wedge$B.

We believe that the ideas in this work have the potential for nontrivial developments. An intriguing problem is the rewriting of QCQs in order to minimize the number of distinct variables used in the query, which is the known algorithmic bottleneck for query evaluation. The issue, fully understood on CQs [8], [9], [14], is wide open on QCQs. Indeed, the problem is not even known to be decidable. Perhaps, the notion of Skolem homomorphism might eventually offer a viable approach.

Acknowledgments

The authors thank Michael Benedikt, Massimo Benerecetti, Michael Vanden Boom, and the anonymous reviewers for thorough suggestions which helped to improve earlier versions of the manuscript. The first author is supported by the FWF Austrian Science Fund (Parameterized Compilation, P26200).

References

[1] S. Aanderaa, "On the Decision Problem for Formulas in which all Disjunctions are Binary." in *SLS'71*. Elsevier, 1971, vol. 63, pp. 1–18.
[2] S. Abiteboul, R. Hull, and V. Vianu, *Foundations of Databases.* Addison-Wesley, 1995.
[3] H. Andréka, J. van Benthem, and I. Németi, "Modal Languages And Bounded Fragments Of Predicate Logic." *JPL*, vol. 27, no. 3, pp. 217–274, 1998.
[4] F. Baader and W. Snyder, "Unification Theory." in *Handbook of Automated Reasoning (vol. 1).* Elsevier & MIT Press, 2001, pp. 445–532.
[5] V. Bárány, B. ten Cate, and L. Segoufin, "Guarded Negation." *JACM*, vol. 62, no. 3, pp. 22:1–26, 2015.
[6] E. Börger, "Reduktionstypen in Krom- und Hornformeln." Ph.D. dissertation, University of Münster, Münster, Germany, 1971.
[7] E. Börger, E. Grädel, and Y. Gurevich, *The Classical Decision Problem.*, ser. Perspectives in Mathematical Logic. Springer, 1997.
[8] S. Bova and H. Chen, "The Complexity of Width Minimization for Existential Positive Queries." in *ICDT'14*. OpenProceedings.org, 2014, pp. 235–244.
[9] ——, "How Many Variables are Needed to Express an Existential Positive Query?" in *ICDT'17*. OpenProceedings.org, 2017, pp. 9:1–16.
[10] J. Büchi, "Turing-Machines and the Entscheidungsproblem." *MA*, vol. 148, no. 3, pp. 201–213, 1962.
[11] C.-L. Chang and R.C.T. Lee, *Symbolic Logic and Mechanical Theorem Proving.*, ser. Computer Science Classics. Academic Press, 1973.
[12] A. Chandra and P. Merlin, "Optimal Implementation of Conjunctive Queries in Relational Data Bases." in *STOC'77*. Association for Computing Machinery, 1977, pp. 77–90.
[13] H. Chen, F. Madelaine, and B. Martin, "Quantified Constraints and Containment Problems." *LMCS*, vol. 11, no. 3, pp. 1–28, 2015.
[14] V. Dalmau, P. Kolaitis, and M. Vardi, "Constraint Satisfaction, Bounded Treewidth, and Infinite-Variable Logics." in *CP'02*, ser. LNCS 2470. Springer, 2002, pp. 310–326.
[15] L. Denenberg and H. Lewis, "The Complexity of the Satisfiability Problem for Krom Formulas." *TCS*, vol. 30, no. 3, pp. 319–341, 1984.
[16] H.-D. Ebbinghaus, J. Flum, and W. Thomas, *Mathematical Logic.*, ser. Undergraduate Texts in Mathematics. Springer, 1984.
[17] M. Genesereth and E. Kao, *Introduction to Logic.*, ser. Synthesis Lectures on Computer Science. Morgan & Claypool Publishers, 2013.
[18] E. Grädel, "Decision Procedures for Guarded Logics." in *CADE'99*, ser. LNCS 1632. Springer, 1999, pp. 31–51.
[19] E. Grädel, P. Kolaitis, and M. Vardi, "On the Decision Problem for Two-Variable First-Order Logic." *BSL*, vol. 3, no. 1, pp. 53–69, 1997.
[20] W. Hodges, *A Shorter Model Theory.* Cambridge University Press, 1997.
[21] R. Jeffrey, *Formal Logic: Its Scope and Limits.* McGraw-Hill, 1981.
[22] D. Kozen, "Positive First-order Logic is NP-complete." *IBMJRD*, vol. 25, no. 4, pp. 327–332, 1981.
[23] J. Lassez, M. Maher, and K. Marriott, "Unification Revisited." in *FLFP'86*, ser. LNCS 306. Springer, 1988, pp. 67–113.
[24] A. Martelli and U. Montanari, "An Efficient Unification Algorithm." *TOPLAS*, vol. 4, no. 2, pp. 258–282, 1982.
[25] F. Mogavero, A. Murano, G. Perelli, and M. Vardi, "What Makes ATL* Decidable? A Decidable Fragment of Strategy Logic." in *CONCUR'12*, ser. LNCS 7454. Springer, 2012, pp. 193–208.
[26] ——, "Reasoning About Strategies: On the Satisfiability Problem." *LMCS*, vol. 13, no. 1:9, pp. 1–37, 2017.
[27] F. Mogavero and G. Perelli, "Binding Forms in First-Order Logic." in *CSL'15*, ser. LIPIcs 41. Leibniz-Zentrum fuer Informatik, 2015, pp. 648–665.
[28] M. Mortimer, "On Languages with Two Variables." *MLQ*, vol. 21, no. 1, pp. 135–140, 1975.
[29] M.-L. Mugnier, G. Simonet, and M. Thomazo, "On the Complexity of Entailment in Existential Conjunctive First-Order Logic with Atomic Negation." *IC*, vol. 215, pp. 8–31, 2012.
[30] J. Robinson, "A Machine-Oriented Logic Based on the Resolution Principle." *JACM*, vol. 12, no. 1, pp. 23–41, 1965.
[31] T. Skolem, "Logisch-kombinatorische Untersuchungen über die Erfüllbarkeit oder Beweisbarkeit mathematischer Sätze nebst einem Theorem über dichte Mengen." vol. 4, 1920.
[32] A. Urquhart, "Decidability and the Finite Model Property." *JPL*, vol. 10, no. 3, pp. 367–370, 1981.