A

Classification of Annotation Semirings over Containment of Conjunctive Queries

EGOR V. KOSTYLEV, University of Edinburgh
JUAN L. REUTTER, University of Edinburgh and Pontificia Universidad Católica de Chile
ANDRÁS Z. SALAMON, University of Edinburgh

We study the problem of query containment of conjunctive queries over annotated databases. Annotations are typically attached to tuples and represent metadata such as probability, multiplicity, comments, or provenance. It is usually assumed that annotations are drawn from a commutative semiring. Such databases pose new challenges in query optimization, since many related fundamental tasks, such as query containment, have to be reconsidered in the presence of propagation of annotations.

We axiomatize several classes of semirings for each of which containment of conjunctive queries is equivalent to existence of a particular type of homomorphism. For each of these types we also specify all semirings for which existence of a corresponding homomorphism is a sufficient (or necessary) condition for the containment. We develop new decision procedures for containment for some semirings which are not in any of these classes. This generalizes and systematizes previous approaches.

Categories and Subject Descriptors: H.2.4 [**Database Management**]: Systems—*query processing*; H.2.1 [**Database Management**]: Data Models

General Terms: Design, Algorithms, Theory

Additional Key Words and Phrases: Annotation, Provenance, Query Optimization

ACM Reference Format:
Kostylev, E. V., Reutter, J. L., and Salamon, A. Z. 2013. Classification of Annotation Semirings over Containment of Conjunctive Queries *ACM Trans. Datab. Syst.* V, N, Article A (YYYY), 39 pages.
DOI: http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

Relational database annotation is rapidly coming to market. The expressive power of curated [Buneman et al. 2008] and probabilistic databases [Fuhr and Rölleke 1997; Zimányi 1997], various forms of provenance [Cui et al. 2000; Buneman et al. 2001; Green et al. 2007], and even bag (multiset) semantics as a way to model standard SQL [Chaudhuri and Vardi 1993], derives from an annotation attribute with special behaviour. Green et al. [2007] observed that in all of these cases annotations propagate through queries as we expect if the domain of annotations has the structure of a *commutative semiring*. Karvounarakis and Green [2012] recently surveyed work building on this model.

This work is supported under SOCIAM: The Theory and Practice of Social Machines, a project funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/J017728/1. This work was also supported by FET-Open Project FoX, grant agreement 233599; EPSRC grants EP/F028288/1, G049165 and J015377; and the Laboratory for Foundations of Computer Science. Authors' addresses: Laboratory for Foundations of Computer Science, School of Informatics, University of Edinburgh; and Departamento de Ciencia de la Computación, Escuela de Ingeniería, Pontificia Universidad Católica de Chile. Authors' email addresses: ekostyle@inf.ed.ac.uk, Juan.Reutter@ed.ac.uk and Andras.Salamon@ed.ac.uk.

DOI: http://dx.doi.org/10.1145/0000000.0000000

To perform standard tasks such as query rewriting and query optimization, it must be possible to compare queries in some appropriate manner. Every application that supports annotations should therefore also support comparisons between queries. However, as noted by Ioannidis and Ramakrishnan [1995] and Chaudhuri and Vardi [1993] for the particular case of bag semantics and quite generally by Green [2011], the introduction of annotations requires a complete rethinking of these kinds of tasks: a pair of queries may behave differently when posed over ordinary relations or over annotated relations; the behaviour can be different even for different semirings. Hence a general theory is needed to explain how queries behave over annotated relations, and to provide query optimization and query rewriting techniques, regardless of the semiring chosen for annotations.

In this paper, we study the problem of *containment* of queries, specifically for the classes of *conjunctive queries* (CQs). For this purpose we formally generalize the standard notion of containment for relational databases [Chandra and Merlin 1977] so that it subsumes previously studied containments for bag semantics [Ioannidis and Ramakrishnan 1995; Chaudhuri and Vardi 1993] and several other semirings [Green 2011]. We study in our view the most general reasonable notion of containment, based on a few intuitive axioms which any containment should satisfy.

The ideal would be to obtain a procedure to decide containment of CQs for an arbitrary annotation semiring. However, there is evidence that obtaining such a procedure for all semirings is a truly challenging, if not impossible, task. Indeed, this would require solving containment for bag semantics, which is a long-standing open problem for CQs [Chaudhuri and Vardi 1993; Ioannidis and Ramakrishnan 1995; Afrati et al. 2010; Chirkova 2012], and is even undecidable for unions of CQs [Ioannidis and Ramakrishnan 1995] or CQs with inequalities [Jayram et al. 2006]. With these observations in mind, we instead ask the following, narrower question: are there reasonable classes of semirings for which we can prove that containment of CQs is decidable? In this paper we answer this question positively, by finding several such classes. Our main results generalize and extend previous work [Green 2011; Grahne et al. 1997] unifying how semantic properties of query containment link to syntactic properties of different types of *homomorphisms* between queries. We also show that these classes are of importance in practice, as they contain the majority of the annotation semirings that have been proposed.

For standard relational databases (which can be modelled by a set semantics semiring consisting of just two elements **true** and **false**), query containment corresponds precisely to the NP-complete problem of deciding whether there exists a homomorphism between these queries [Chandra and Merlin 1977]. Thus, the natural starting point of our search for decidable classes is to ask for which semirings the CQ containment problem coincides with CQ containment for the usual set semantics. This question was partially answered by Ioannidis and Ramakrishnan [1995], where for semirings which are so called type A systems, containment was shown to be equivalent to the existence of a homomorphism. We show that it is possible to describe the class $\mathbf{C}_{\mathrm{hom}}$ of *all* such semirings by two simple axioms: idempotence of multiplication and annihilation of the multiplicative identity. (The latter property informally means that the multiplicative identity is the greatest element in the semiring.) Notably, this class corresponds precisely to the class of type A′ systems [Ioannidis and Ramakrishnan 1995], for which such a characterization was left open.

Continuing our search for decidable classes, in Sec. 4 we consider those classes obtained by relaxing the axioms for $\mathbf{C}_{\mathrm{hom}}$. In Sec. 4.1 to 4.4 we show that for each of these classes there exists a well-known natural type of homomorphism that is associated with the class. For these classes, existence of an appropriate type of homomorphism between two CQs is *sufficient* to conclude that the one CQ is contained in the other. As an example, consider the class of semirings that satisfy only the annihilation axiom. In Sec. 4.2 we demonstrate that this class contains precisely all the semirings for which the existence of an *injective homomorphism* is a sufficient condition for containment of two CQs. A sufficient condition does not guarantee the decidability of the containment problem; one needs a necessary condition as well. For this purpose, we describe the *largest* class for which an injective homomorphism is necessary for containment of CQs. Thereby, we have that for all semirings in the intersection of these two classes, the existence of an injective homomorphism is both a necessary and sufficient condition for the containment of two CQs, resulting in a class $\mathbf{C}_{\text{in}}$ of semirings for which containment is decidable.

We establish similar results for several other classes of semirings obtained by relaxing the axioms that define the class $\mathbf{C}_{\text{hom}}$, and show how these classes are characterized by other well-known types of homomorphisms. This yields NP decision procedures for containment of CQs for the corresponding classes of semirings. We provide matching complexity lower bounds: all of these decision problems are NP-complete. We also prove a more general result that the decision problem is NP-hard for all semirings considered in this paper.

To axiomatize some of these classes, in Sec. 4.1 we introduce the notion of CQ-*admissible polynomials*. Intuitively, a polynomial is CQ-admissible if it can be obtained by evaluating a CQ over a database annotated with variables. In Sec. 5 we give a syntactic characterization of these polynomials. This novel concept is of independent interest; for instance Olteanu and Závodný [2012] implicitly use the properties of such polynomials for effectively storing and manipulating the provenance of CQ results.

Moving beyond homomorphisms, in Sec. 6 we also find several semirings for which containment of CQs can be solved via a *small model property*, by looking for a small enough database witness for absence of containment. More precisely, we show that, if a semiring satisfies the idempotence of addition axiom, then two CQs are contained with respect to this semiring if and only if they are contained on all instances of size no greater than the size of the pair of queries. Using this property, we show that in this case the containment problem can be cast as the problem of deciding whether the evaluation of a CQ-admissible polynomial is greater than or equal to the evaluation of another such polynomial, for any assignment of values to the variables from the corresponding semiring. Thus, the decidability of such an order on polynomials implies the decidability for containment of CQs, under any semiring that satisfies our idempotence axiom. This results in new decision procedures to solve containment of CQs, for a wide range of semirings that had not been previously addressed. As an example of how to use this machinery, we study the problem of the order on polynomials for two well-known semirings – the *tropical semiring* and the *max-plus algebra* – and use these results to provide novel complexity bounds to decide containment of CQs under these semirings.

It follows from our definition of containment that two queries are equivalent if and only if they are contained in each other. Thus, all of our upper bounds for query containment naturally translate into upper bounds for deciding the equivalence of queries. However, lower bounds need not be the same. For instance, while the decidability problem of containment under bag semantics remains open, and is Π_2^p-hard according to Chaudhuri and Vardi [1993], the equivalence problem in this case can be solved simply by checking for an isomorphism between queries. Therefore deciding the equivalence of queries under annotated relations is a different problem from the containment problem that we study in this paper, and it is an interesting, non-trivial problem that deserves to be studied on its own.

Most of our results were previously announced in a conference paper ([Kostylev et al. 2012]). Here we include detailed proofs and several new results. We also present a number of alternative definitions and characterizations that are not only useful for the understanding of the complete picture behind the study, but are also interesting in their own right.

In particular, the new material includes the following. Sec. 3.3 contains a more detailed analysis of the properties of the class $\mathbf{C}_{\mathrm{hom}}$ of semirings. In Sec. 4.1 we show (Prop. 4.6) how CQ-admissible polynomials can be defined only in terms of queries without free variables, and use this property to define a novel, alternative characterization for the class $\mathbf{N}_{\mathrm{hcov}}$ of semirings, which is based on CQ-admissible polynomials (Lem. 4.7). In Sec. 4.4 we show how our machinery can be used to prove that containment of conjunctive queries is NP-hard under any semiring considered in this paper. Sec. 5 now gives a detailed proof of the syntactic characterization of CQ-admissible polynomials, and along with the proof we include the intuition behind this characterization. Finally, in Sec. 6 we describe (Prop. 6.5) a completely new technique for deciding the order on CQ-admissible polynomials under some particular semirings, such as the tropical semiring and the max-plus algebra. Our approach draws upon results in the area of linear integer programming, and in particular enables us to improve the upper bounds for containment of CQs under these semirings, from the PSPACE bound presented by Kostylev et al. [2012], to Π_2^p.

We would also like to note that some of the results in [Kostylev et al. 2012] are not included in this version. To be more precise, this paper contains only results regarding containment of conjunctive queries, while the conference version also investigates the problem of containment of unions of conjunctive queries. As much as we would have liked to include all these results, due to the space limitations it was not possible to include them with the same level of detail as the rest of the results of this paper. We intend to publish these results in an extended version dedicated solely to the problem of containment of unions of conjunctive queries. For now, we refer the reader to [Kostylev et al. 2012].

2. PRELIMINARIES

Commutative semirings An algebraic structure $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$ with binary operations *sum* $\oplus$ and *product* $\otimes$ and constants $\mathbb{0}$ and $\mathbb{1}$ is a (*commutative*) *semiring* iff $\langle K, \oplus, \mathbb{0}\rangle$ and $\langle K, \otimes, \mathbb{1}\rangle$ are commutative monoids[1] with identities $\mathbb{0}$ and $\mathbb{1}$ respectively, $\otimes$ is distributive over $\oplus$, and $a \otimes \mathbb{0} = \mathbb{0}$ holds for each $a \in K$. It will be convenient for us to consider only *nontrivial* semirings, i.e. semirings such that $\mathbb{0} \neq \mathbb{1}$. We use the symbols $\sum$ and $\prod$ to denote sum and product of sets of semiring elements, i.e. using operations $\oplus$ and $\otimes$.

In the paper we will discuss many examples of semirings, such as the semiring of natural numbers, where the abstract operations $\oplus$ and $\otimes$ instantiate to the usual $+$ and $\times$; or the tropical semiring, where these operations instantiate to $\min$ and $+$, respectively.

$\mathcal{K}$-relations A *schema* $\mathbb{S}$ is a finite set of *relational symbols*, each of which is assigned a non-negative *arity*. For a semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$ and a domain $\mathbb{D}$ of *constants*, a *$\mathcal{K}$-instance* I over a schema $\mathbb{S}$ assigns to each relational symbol R from $\mathbb{S}$ of arity m a *$\mathcal{K}$-relation* R^I, which is a (total) function from the set of *tuples* $\mathbb{D}^m$ to K such that its *support*, i.e. the set $\{\mathbf{t} \mid \mathbf{t} \in \mathbb{D}^m, R^I(\mathbf{t}) \neq \mathbb{0}\}$, is finite.[2] We call $R^I(\mathbf{t})$ the *annotation* of the tuple $\mathbf{t}$ in the $\mathcal{K}$-relation R^I.

Queries A *conjunctive query* (or *CQ*, for short) Q over a schema $\mathbb{S}$ is an expression of the form $\exists \mathbf{v}\ \phi(\mathbf{u}, \mathbf{v})$, where $\mathbf{u}$ is a list of *free* variables, $\mathbf{v}$ is a list of *existential* variables

[1]A commutative monoid is a set with an associative and commutative binary operation and an identity element.
[2]Hence, in this model every tuple of appropriate arity is annotated in every relation.

and $\phi(\mathbf{u}, \mathbf{v})$ is a multiset of relational atoms over $\mathbb{S}$ using variables $\mathbf{u} \cup \mathbf{v}$. As usual we write $\phi(\mathbf{u}, \mathbf{v}) = R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_n(\mathbf{u}_n, \mathbf{v}_n)$, where $\mathbf{u}_1 \cup \ldots \cup \mathbf{u}_n = \mathbf{u}$ and $\mathbf{v}_1 \cup \ldots \cup \mathbf{v}_n = \mathbf{v}$, keeping in mind that R_i and R_j in this expression can be the same symbol even if $i \neq j$. A *union* of conjunctive queries (*UCQ*) $\mathbf{Q}$ is a multiset of CQs over the same schema and the same set of free variables.

Evaluations For a CQ $Q = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and a tuple $\mathbf{t}$, denote by $\mathcal{V}(Q, \mathbf{t})$ the set of all mappings f from $\mathbf{u} \cup \mathbf{v}$ to the domain $\mathbb{D}$ such that $f(\mathbf{u}) = \mathbf{t}$. Given a $\mathcal{K}$-instance I, the *evaluation* of Q on I for $\mathbf{t}$ is the value

$$Q^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q,\mathbf{t})} \prod_{1 \leq i \leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)).$$

Similarly, the *evaluation* of a UCQ $\mathbf{Q}$ on I for $\mathbf{t}$ is the value

$$\mathbf{Q}^I(\mathbf{t}) = \sum_{Q \in \mathbf{Q}} Q^I(\mathbf{t}).$$

Note, that from this definition it follows that if $\mathbf{Q} = \emptyset$ then $\mathbf{Q}^I(\mathbf{t}) = \mathbb{0}$.

3. GENERAL FRAMEWORK

3.1. $\mathcal{K}$-containment and $\preceq$-positive semirings

As noted by Green et al. [2007], the introduction of annotations on relations requires a complete rethinking of the notions of query optimization and query rewriting. For the case of bag semantics, Chaudhuri and Vardi [1993] demonstrated that two queries that are equivalent when posed over ordinary relations may not be equivalent when evaluated on $\mathcal{K}$-relations. Furthermore, for two different semirings $\mathcal{K}_1$ and $\mathcal{K}_2$, two queries may be equivalent under $\mathcal{K}_1$-relations, but not equivalent under $\mathcal{K}_2$-relations.

Our main aim is to explore the problem of query containment over different $\mathcal{K}$-relations. First we need to formally specify what we mean by “equivalence” and “containment” of queries. The notion of equivalence is naturally formalised as follows: given a semiring $\mathcal{K}$, UCQs $\mathbf{Q}_1$ and $\mathbf{Q}_2$ over the same schema are $\mathcal{K}$-*equivalent* (denoted $\mathbf{Q}_1 \equiv_{\mathcal{K}} \mathbf{Q}_2$) iff for every $\mathcal{K}$-instance I and tuple $\mathbf{t}$ it holds that $\mathbf{Q}_1^I(\mathbf{t}) = \mathbf{Q}_2^I(\mathbf{t})$. However, to study containment of queries over some semiring $\mathcal{K}$, we should be able to compare elements of $\mathcal{K}$ not only for equality. Therefore, we assume that the semiring $\mathcal{K}$ is equipped with a *partial order*[3] $\preceq_{\mathcal{K}}$. This allows us to define when a UCQ $\mathbf{Q}_1$ is $\mathcal{K}$-*contained* in a UCQ $\mathbf{Q}_2$, which we denote by $\mathbf{Q}_1 \subseteq_{\mathcal{K}} \mathbf{Q}_2$:

$$\mathbf{Q}_1 \subseteq_{\mathcal{K}} \mathbf{Q}_2 \iff \forall I\ \forall \mathbf{t}\ \mathbf{Q}_1^I(\mathbf{t}) \preceq_{\mathcal{K}} \mathbf{Q}_2^I(\mathbf{t}).$$

Note that by this definition different partial orders may produce the same $\mathcal{K}$-containment. However, for every $\mathcal{K}$-containment there exists a unique *minimal* order among these, i.e. the partial order $\preceq_{\mathcal{K}}$ such that there is no subrelation of $\preceq_{\mathcal{K}}$ that produces the same $\mathcal{K}$-containment. It is a reasonable assumption that $\preceq_{\mathcal{K}}$ is minimal with respect to $\subseteq_{\mathcal{K}}$, and indeed we will make this assumption for the rest of the paper.[4]

However, for some partial orders the above definition results in a rather spartan notion of $\mathcal{K}$-containment. For example, by considering the usual order $\leq$ on the semiring $\mathbb{Z}$ of *integers*, one can easily verify that the empty UCQ is not $\mathbb{Z}$-contained in any non-empty UCQ.

Thus, we need to restrict the class of partially ordered semirings that we consider for our study. In order to do so, we list four intuitive requirements that, in our view, any definition of $\mathcal{K}$-containment should satisfy, and then identify all the semirings $\mathcal{K}$

[3] A partial order is a transitive, reflexive and antisymmetric binary relation.
[4] We use this fact in the proof of Prop. 3.1.

equipped with partial orders $\preceq_{\mathcal{K}}$ for which the definition of $\mathcal{K}$-containment is guaranteed to satisfy our requirements. These requirements are as follows:

(C1) $\subseteq_{\mathcal{K}}$ is a *preorder*, i.e. reflexive and transitive;
(C2) $\mathbf{Q}_1 \equiv_{\mathcal{K}} \mathbf{Q}_2$ iff $\mathbf{Q}_1 \subseteq_{\mathcal{K}} \mathbf{Q}_2$ and $\mathbf{Q}_2 \subseteq_{\mathcal{K}} \mathbf{Q}_1$;
(C3) $\emptyset \subseteq_{\mathcal{K}} \mathbf{Q}$ holds for all $\mathbf{Q}$;
(C4) if $\mathbf{Q}_1 \subseteq_{\mathcal{K}} \mathbf{Q}_2$ then $\mathbf{Q}_1 \cup \mathbf{Q}_3 \subseteq_{\mathcal{K}} \mathbf{Q}_2 \cup \mathbf{Q}_3$ for any $\mathbf{Q}_3$.

Requirements (C1) and (C2) essentially state that our notion of containment behaves as a partial order with respect to the equality we have defined previously. Requirements (C3) and (C4) impose further conditions to ensure that the notion of containment behaves in a natural way. For example, requirement (C3) rules out the example with $\mathbb{Z}$ and $\leq$; and requirement (C4) is typically needed when considering query processing tasks such as query rewriting.

It turns out that we can easily axiomatize the class of semirings with partial orders that have $\mathcal{K}$-containments satisfying (C1) – (C4). The following proposition says that this class consists of all $\preceq$*-positive*[5] semirings, i.e. semirings $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$ equipped with a partial order $\preceq_{\mathcal{K}}$, such that

(1) $\mathbb{0} \preceq_{\mathcal{K}} a$ for all $a \in K$, and
(2) $a \preceq_{\mathcal{K}} b \Rightarrow a \oplus c \preceq_{\mathcal{K}} b \oplus c$ for all $a, b, c \in K$.

PROPOSITION 3.1. *A semiring $\mathcal{K}$ equipped with a partial order $\preceq_{\mathcal{K}}$ is $\preceq$-positive iff the corresponding $\mathcal{K}$-containment $\subseteq_{\mathcal{K}}$ satisfies the requirements* (C1) – (C4).

PROOF. First we show that if $\subseteq_{\mathcal{K}}$ satisfies the requirements (C1) – (C4) then the partial order $\preceq_{\mathcal{K}}$ satisfies $\mathbb{0} \preceq_{\mathcal{K}} a$ and $a \preceq_{\mathcal{K}} b \Rightarrow a \oplus c \preceq_{\mathcal{K}} b \oplus c$ for all a, b, and c from $\mathcal{K}$.

To show that $\mathbb{0} \preceq_{\mathcal{K}} a$ for all $a \in K$, consider a UCQ $\mathbf{Q} = \{\exists v\ R(v)\}$, and a $\mathcal{K}$-instance I such that $R^I(c) = a$ for some constant $c \in \mathbb{D}$ and $R^I(c') = \mathbb{0}$ for all $c' \in \mathbb{D}$, $c' \neq c$. By requirement (C3) we have that $\emptyset \subseteq_{\mathcal{K}} \mathbf{Q}$. Hence, we have that

$$\mathbb{0} = \emptyset^I() \preceq_{\mathcal{K}} \mathbf{Q}^I() = a.$$

To show that $a \oplus c \preceq_{\mathcal{K}} b \oplus c$ for all $a, b, c \in K$ such that $a \preceq_{\mathcal{K}} b$, consider $\mathbf{Q}_1$ and $\mathbf{Q}_2$ over some schema $\mathbb{S}$, such that $\mathbf{Q}_1 \subseteq_{\mathcal{K}} \mathbf{Q}_2$ and for some $\mathcal{K}$-instance I and tuple $\mathbf{t}$ it holds that $\mathbf{Q}_1^I(\mathbf{t}) = a$ and $\mathbf{Q}_2^I(\mathbf{t}) = b$. Such UCQs exist by the minimality of $\preceq_{\mathcal{K}}$. Extend $\mathbb{S}$ to a new schema $\mathbb{S}'$ with a new relation R of arity equal to the size of $\mathbf{t}$. If we consider $\mathbf{Q}_1$ and $\mathbf{Q}_2$ as UCQs over $\mathbb{S}'$ then $\mathbf{Q}_1 \subseteq_{\mathcal{K}} \mathbf{Q}_2$ still holds, since these CQs do not use R. Let J be an extension of I on R such that $R^J(\mathbf{t}) = c$ and $R^J(\mathbf{t}') = \mathbb{0}$ for all $\mathbf{t}' \neq \mathbf{t}$. Let also $\mathbf{Q}_3 = \{Q\}$ where $Q = R(\mathbf{u})$. From requirement (C4) we have that $\mathbf{Q}_1 \cup \mathbf{Q}_3 \subseteq_{\mathcal{K}} \mathbf{Q}_2 \cup \mathbf{Q}_3$. Hence the following holds, which suffices for the proof:

$$a \oplus c = \mathbf{Q}_1^J(\mathbf{t}) \oplus \mathbf{Q}_3^J(\mathbf{t}) \preceq_{\mathcal{K}} \mathbf{Q}_2^J(\mathbf{t}) \oplus \mathbf{Q}_3^J(\mathbf{t}) = b \oplus c.$$

It is left to show that if $\preceq_{\mathcal{K}}$ is a partial order satisfying $\mathbb{0} \preceq_{\mathcal{K}} a$ and $a \preceq_{\mathcal{K}} b \Rightarrow a \oplus c \preceq_{\mathcal{K}} b \oplus c$ for all a, b, and c from $\mathcal{K}$, then $\subseteq_{\mathcal{K}}$ satisfies the requirements (C1) – (C4).

Requirements (C1) and (C2) follow immediately from the fact that $\preceq_{\mathcal{K}}$ is a reflexive, transitive and antisymmetric relation.

To prove requirement (C3), we need to show that for each $\mathbf{Q}$ it holds that $\emptyset \subseteq_{\mathcal{K}} \mathbf{Q}$. Consider an arbitrary UCQ $\mathbf{Q}$. From the fact that $\mathbb{0} \preceq_{\mathcal{K}} a$ for each $a \in \mathcal{K}$, we have that $\mathbb{0} \preceq_{\mathcal{K}} \mathbf{Q}^I(\mathbf{t})$ for any instance I and tuple $\mathbf{t}$. As noted above, for any I and $\mathbf{t}$ we have

[5]In the conference version of this paper the term "positive semiring" was used for this notion. To avoid confusion with other usage in the literature (for instance, Green [2011] uses the definition of Eilenberg [1974]), a refined term is used in this paper.

$\emptyset^I(\mathbf{t}) = \mathbb{0}$. Thus, it holds that $\emptyset^I(\mathbf{t}) \preceq_{\mathcal{K}} \mathbf{Q}^I(\mathbf{t})$ for any I and $\mathbf{t}$, and by definition this means that $\emptyset \subseteq_{\mathcal{K}} \mathbf{Q}$.

To prove requirement (C4), we show that for each $\mathbf{Q}_1, \mathbf{Q}_2$ such that $\mathbf{Q}_1 \subseteq_{\mathcal{K}} \mathbf{Q}_2$, and each $\mathbf{Q}_3$, it holds that $\mathbf{Q}_1 \cup \mathbf{Q}_3 \subseteq_{\mathcal{K}} \mathbf{Q}_2 \cup \mathbf{Q}_3$. Assume then that for queries $\mathbf{Q}_1$ and $\mathbf{Q}_2$, and any instance I and tuple $\mathbf{t}$, we have that $\mathbf{Q}_1^I(\mathbf{t}) \preceq_{\mathcal{K}} \mathbf{Q}_2^I(\mathbf{t})$. Then from the properties of $\preceq_{\mathcal{K}}$ we have that $\mathbf{Q}_1^I(\mathbf{t}) \oplus c \preceq_{\mathcal{K}} \mathbf{Q}_2^I(\mathbf{t}) \oplus c$, for any instance I and tuple $\mathbf{t}$, and for any $c \in \mathcal{K}$. Hence, $\mathbf{Q}_1^I(\mathbf{t}) \oplus \mathbf{Q}_3^I(\mathbf{t}) \preceq_{\mathcal{K}} \mathbf{Q}_2^I(\mathbf{t}) \oplus \mathbf{Q}_3^I(\mathbf{t})$, for any such I and $\mathbf{t}$. Hereby, by definition we obtain $\mathbf{Q}_1 \cup \mathbf{Q}_3 \subseteq_{\mathcal{K}} \mathbf{Q}_2 \cup \mathbf{Q}_3$. $\square$

We assume for the rest of the paper that all semirings are $\preceq$-positive and denote the class of such semirings by $\mathbf{S}_{\preceq}$.

We focus in this work on the following decision problem:

CQ $\mathcal{K}$-CONTAINMENT:	
Input:	CQs Q_1, Q_2.
Question:	Is $Q_1 \subseteq_{\mathcal{K}} Q_2$?

In particular, we are interested in classifying the semirings in $\mathbf{S}_{\preceq}$ for which different conditions on CQs are sufficient for $\mathcal{K}$-containment, and also for which semirings they are necessary. If for a semiring $\mathcal{K}$ such a condition is both sufficient and necessary, and it is possible to check the condition algorithmically, then we have a decision procedure for $\mathcal{K}$-containment.

3.2. Naturally ordered semirings and provenance polynomials

A semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$ is *naturally ordered* iff the preorder $\preceq_{\mathcal{K}}^{\text{nat}}$, defined as $a \preceq_{\mathcal{K}}^{\text{nat}} b \iff \exists c\ a \oplus c = b$, is a partial order. Green [2011] noted that in most semantics considered so far, including set and bag semantics, the notion of containment is based on natural orders of the semirings. In principle, this condition appears to be too restrictive, and for this reason we have opted for the more general approach based on $\preceq$-positive semirings. It is straightforward to show that any naturally ordered semiring is a $\preceq$-positive semiring. However, it is also possible to show that every $\preceq$-positive semiring with a partial order $\preceq_{\mathcal{K}}$ is a naturally ordered semiring, but $\preceq_{\mathcal{K}}$ is an extension of $\preceq_{\mathcal{K}}^{\text{nat}}$ (i.e. $\preceq_{\mathcal{K}}^{\text{nat}}$ is a subrelation of $\preceq_{\mathcal{K}}$). Thus, our approach is general enough to include all previous work, as far as we are aware.

In [Green 2011] the problem of $\mathcal{K}$-containment of CQs and UCQs was considered for several naturally ordered semirings, including the one known as the semiring of *provenance polynomials*, $\mathcal{N}[X] = \langle \mathbb{N}[X], +, \times, 0, 1\rangle$. This is the set $\mathbb{N}[X]$ of polynomials over a set of variables X, with natural number coefficients, equipped with the usual operations $+$ and $\times$. Green et al. [2007] pointed out that this semiring (without any order) is special among all semirings since it is "most general", i.e. possesses the *universal property*: for any (unordered) semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$ any function $\nu \colon X \to K$ can be uniquely extended to a *morphism* $\mathsf{Eval}_\nu \colon \mathbb{N}[X] \to K$, i.e. a mapping between semirings which preserves all the operations and relations (including constants 0 and 1). Conceptually, this property means that any semantical behaviour of the universal semiring is also the behaviour of any other semiring (see Green et al. [2007] for details). Green [2011] showed that $\mathcal{N}[X]$, now with its natural order, is universal for all naturally ordered semirings. It turns out that this is also true for all ($\preceq$-positive) semirings.

PROPOSITION 3.2. *Given a set of variables X, $\mathcal{N}[X]$ is universal for the class $\mathbf{S}_{\preceq}$ of all ($\preceq$-positive) semirings.*

PROOF. From [Green 2011] we know that $\mathcal{N}[X]$ is universal for all naturally ordered semirings. Also, as mentioned just before this proposition, every ($\preceq$-positive) semiring

$\mathcal{K}$ is naturally ordered with the order $\preceq_{\mathcal{K}}$ extending its natural order $\preceq_{\mathcal{K}}^{\mathrm{nat}}$. From these facts we conclude that any morphism $\mathsf{Eval}_\nu : \mathbb{N}[X] \to K$ preserves the natural order on $\mathcal{N}[X]$, i.e. $\mathsf{P} \preceq_{\mathcal{N}[X]} \mathsf{P}'$ implies that $\mathsf{Eval}_\nu(\mathsf{P}) \preceq_{\mathcal{K}}^{\mathrm{nat}} \mathsf{Eval}_\nu(\mathsf{P}')$ and, hence, $\mathsf{Eval}_\nu(\mathsf{P}) \preceq_{\mathcal{K}} \mathsf{Eval}_\nu(\mathsf{P}')$. □

Based on this property, we can formulate different universal axioms on semirings, involving the order $\preceq_{\mathcal{K}}$, in terms of $\mathcal{N}[X]$. Given a semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$ from $\mathbf{S}_{\preceq}$, a set X of n variables, and polynomials P_1 and P_2 from $\mathbb{N}[X]$, we write $\mathsf{P}_1 \preceq_{\mathcal{K}} \mathsf{P}_2$ iff for *each* function $\nu \colon X \to K$ the inequality $\mathsf{Eval}_\nu(\mathsf{P}_1) \preceq_{\mathcal{K}} \mathsf{Eval}_\nu(\mathsf{P}_2)$ holds for the morphism Eval_ν, i.e. the order holds for every valuation of these polynomials. Since $\preceq_{\mathcal{K}}$ is a partial order, we can also write $\mathsf{P}_1 =_{\mathcal{K}} \mathsf{P}_2$ for $\mathsf{P}_1 \preceq_{\mathcal{K}} \mathsf{P}_2 \wedge \mathsf{P}_2 \preceq_{\mathcal{K}} \mathsf{P}_1$. Polynomials of this kind will play an important role in this paper, and we will extensively use such polynomial notation. Sometimes we will also refer to *monomials*, by which we mean products of variables (without coefficients).

3.3. Containment by homomorphisms

The study of query containment in the context of query optimization had begun for relational databases by the 1970s [Chandra and Merlin 1977]. These databases can be naturally modelled by $\mathcal{B}$-relations, where $\mathcal{B} = \langle\{\mathbf{false}, \mathbf{true}\}, \vee, \wedge, \mathbf{false}, \mathbf{true}\rangle$ is the *set semantics* semiring. Here a tuple is annotated with **true** iff it is in the relation and **false** otherwise. For $\mathcal{B}$-containment the natural order $\preceq_{\mathcal{B}}$ is assumed, which is defined as $\mathbf{false} \preceq_{\mathcal{B}} \mathbf{true}$. A CQ Q_1 is $\mathcal{B}$-contained in a CQ Q_2 iff one can find a homomorphism from Q_2 to Q_1, by the classical result of Chandra and Merlin [1977]. Given CQs $Q_1 = \exists \mathbf{v}_1\ \phi_1(\mathbf{u}_1, \mathbf{v}_1)$ and $Q_2 = \exists \mathbf{v}_2\ \phi_2(\mathbf{u}_2, \mathbf{v}_2)$, a *homomorphism* (also known as *containment mapping*) from Q_2 to Q_1 is a function $h \colon \mathbf{u}_2 \cup \mathbf{v}_2 \to \mathbf{u}_1 \cup \mathbf{v}_1$ such that $h(\mathbf{u}_2) = \mathbf{u}_1$ and for each atom $R(\mathbf{u}, \mathbf{v})$ from $\phi_2(\mathbf{u}_2, \mathbf{v}_2)$, the atom $R(h(\mathbf{u}, \mathbf{v}))$ is in $\phi_1(\mathbf{u}_1, \mathbf{v}_1)$. A homomorphism extends to atoms and sets of atoms in the usual way. We write $Q_2 \to Q_1$ iff there exists a homomorphism from Q_2 to Q_1.

Based on the results of Green [2011] or Ioannidis and Ramakrishnan [1995] it is not difficult to show that the existence of a homomorphism between CQs is necessary for their $\mathcal{K}$-containment over any $\preceq$-positive nontrivial semiring $\mathcal{K}$. For the proof of this fact one would use the following notion, which we exploit extensively in the rest of the paper.

Fix a set of variables X. A *canonical instance* ([Green et al. 2007]) $[\![Q]\!]$ of a CQ Q is an $\mathcal{N}[X]$-instance with the same schema as Q and with the set of variables of Q as its domain, such that for every $\mathcal{N}[X]$-relation $R^{[\![Q]\!]}$ and for every tuple $\mathbf{u}, \mathbf{v}$ it holds that $R^{[\![Q]\!]}(\mathbf{u}, \mathbf{v}) = x_1 + \ldots + x_n$, where $n \geq 0$ is the number of atoms in Q of the form $R(\mathbf{u}, \mathbf{v})$, and $x_1, \ldots, x_n$ are unique (over all $[\![Q]\!]$) variables from X.

While there may be infinitely many canonical instances for any given query, they are all isomorphic up to renaming of the variables in the domain of the annotations $\mathcal{N}[X]$. This allows us to speak of the canonical instance of a query, as if it were a unique instance. Next we give a simple example of a canonical instance.

Example 3.3. For the CQ $Q_1 = \exists u, v, w\ R(u, v), R(u, w), S(v, w), S(v, w)$ we have

$$R^{[\![Q_1]\!]}(u, v) = x_1, \quad R^{[\![Q_1]\!]}(u, w) = x_2,$$
$$S^{[\![Q_1]\!]}(v, w) = x_3 + x_4,$$

i.e. in the relation R of the canonical instance $[\![Q_1]\!]$ the tuple (u, v) is annotated by x_1, the tuple (u, w) by x_2, and all other tuples by 0; also, in the relation S of this instance the tuple (v, w) is annotated by $x_3 + x_4$, and all other tuples again receive the 0 annotation.

Having this notion, we can prove the fact that $\mathcal{N}[X]$-containment of CQs implies the existence of a homomorphism between them.

FACT 3.4 ([GREEN 2011; IOANNIDIS AND RAMAKRISHNAN 1995]).
For any semiring $\mathcal{K}$ from $\mathbf{S}_{\preceq}$, if $Q_1 \subseteq_{\mathcal{K}} Q_2$ then $Q_2 \to Q_1$.

PROOF. By universality of the semiring $\mathcal{N}[X]$ for $\mathbf{S}_{\preceq}$, each of the semirings in $\mathbf{S}_{\preceq}$ inherits the natural order of $\mathcal{N}[X]$. It is then enough to show that $Q_1 \subseteq_{\mathcal{N}[X]} Q_2$ implies $Q_2 \to Q_1$ for any CQs Q_1 and Q_2.

Let Q_1 and Q_2 be CQs with free variables $\mathbf{u}$ such that $Q_1 \subseteq_{\mathcal{N}[X]} Q_2$. We have, in particular, that $Q_1^{[\![Q_1]\!]}(\mathbf{u}) \preceq_{\mathcal{N}[X]} Q_2^{[\![Q_1]\!]}(\mathbf{u})$ for the canonical instance $[\![Q_1]\!]$. Clearly, $Q_1^{[\![Q_1]\!]}(\mathbf{u}) \neq 0$, which implies that $Q_2^{[\![Q_1]\!]}(\mathbf{u}) \neq 0$. But this means that there exists a mapping h in $\mathcal{V}(Q_2, \mathbf{u})$ such that for every atom $R(\mathbf{u}', \mathbf{v}')$ in Q_2, the atom $R(h(\mathbf{u}', \mathbf{v}'))$ is in Q_1. Then h is the desired homomorphism from Q_2 to Q_1. □

The previous result shows that all semirings $\mathcal{K}$ in $\mathbf{S}_{\preceq}$ share with the set semantics $\mathcal{B}$ the property that existence of a homomorphism is a necessary condition for $\mathcal{K}$-containment. Yet, as mentioned before, for the specific case of set semantics we have that the existence of a homomorphism is also a sufficient condition for containment. Thus, a first natural question to ask is: which semirings behave like $\mathcal{B}$ with respect to containment of CQs, i.e. for which semirings $\mathcal{K}$ is it the case that $Q_2 \to Q_1$ is sufficient (and necessary) for $Q_1 \subseteq_{\mathcal{K}} Q_2$? This question was answered partially in [Green et al. 2007; Green 2011; Ioannidis and Ramakrishnan 1995], and Grahne et al. [1997] showed that this correspondence holds if $\mathcal{K}$ is a distributive bilattice. As the main result of this section we show that it is possible to axiomatize the class of all semirings for which $\mathcal{K}$-containment of CQs coincides with the usual set semantics containment.

Definition 3.5 (*Class* $\mathbf{C}_{\mathrm{hom}}$ *of semirings*). Denote by $\mathbf{C}_{\mathrm{hom}}$ the class of semirings $\mathcal{K}$ that satisfy the following axioms (using the convenient polynomial notation introduced at the end of Sec. 3.2, i.e. assuming that all variables are universally quantified):

(1) ($\otimes$-*idempotence*) $x \times x =_{\mathcal{K}} x$;
(2) ($\mathbb{1}$-*annihilation*) $1 + x =_{\mathcal{K}} 1$.

Next we show that $\mathbf{C}_{\mathrm{hom}}$ contains exactly all semirings that behave like set semantics, w.r.t. $\mathcal{K}$-containment of CQs. In order to do that, we need the following characterizations of the $\mathbb{1}$-*annihilation* axiom. We use these characterizations throughout the paper.

LEMMA 3.6. *Given a semiring $\mathcal{K}$,*

(*1*) *if $\mathcal{K}$ satisfies the $\mathbb{1}$-annihilation axiom then $x_1 \times y_1 + \ldots + x_n \times y_n \preceq_{\mathcal{K}} x_1 + \ldots + x_n$ for every non-negative integer n;*
(*2*) *if $\mathcal{K}$ does not satisfy $\mathbb{1}$-annihilation, then $x \times y \not\preceq_{\mathcal{K}} x$.*

Note that the statement of this lemma uses the polynomial notation introduced in the last part of Sec. 3.2, with implicit universal quantification.

PROOF. Let $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$.

For **Part 1**, we know that $\mathbb{1} \oplus a = \mathbb{1}$ for all $a \in K$. Multiplying by b we conclude that for all $a, b \in K$ it holds that

$$b \oplus (b \otimes a) = b. \tag{1}$$

Since the first requirement of positivity implies $\mathbb{0} \preceq_{\mathcal{K}} b$, by the second requirement we can conclude that $b \otimes a \preceq_{\mathcal{K}} b \oplus (b \otimes a)$. Applying this inequality to (1), we obtain that

$b \otimes a \preceq_{\mathcal{K}} b$ for all $a, b \in K$. Again, by the second requirement of positivity we have that for every $n \geq 0$ and every $a_1, \ldots, a_n, b_1, \ldots, b_n \in K$ it holds that

$$(b_1 \otimes a_1) \oplus \ldots \oplus (b_n \otimes a_n) \preceq_{\mathcal{K}} b_1 \oplus \ldots \oplus b_n.$$

It means that the desired inequality $x_1 \times y_1 + \ldots + x_n \times y_n \preceq_{\mathcal{K}} x_1 + \ldots + x_n$ holds for every integer $n \geq 0$.

For **Part 2**, assume for the sake of contradiction that K satisfies $x \times y \preceq_{\mathcal{K}} x$, i.e. $a \otimes b \preceq_{\mathcal{K}} a$ for all $a, b \in K$. Particularly, for all $a, c \in K$ it holds that $a \oplus (a \otimes c) = a \otimes (\mathbb{1} \oplus c) \preceq_{\mathcal{K}} a$. Take $a = \mathbb{1}$. Then $\mathbb{1} \oplus c \preceq_{\mathcal{K}} \mathbb{1}$ for all $c \in K$. However, by positivity also $\mathbb{1} \preceq_{\mathcal{K}} \mathbb{1} \oplus c$. Hence $\mathbb{1} \oplus c = \mathbb{1}$ for all $c \in K$, which contradicts the assumption that $\mathcal{K}$ does not satisfy $\mathbb{1}$-annihilation. □

We are now ready to present the main result of this section.

THEOREM 3.7. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{C}_{\mathrm{hom}}$;*
- *$Q_1 \subseteq_{\mathcal{K}} Q_2$ iff $Q_2 \to Q_1$, for all CQs Q_1 and Q_2.*

PROOF. By Prop. 3.4, for any ($\preceq$-positive, nontrivial) semiring $\mathcal{K}$ and CQs Q_1, Q_2 it holds that if $Q_1 \subseteq_{\mathcal{K}} Q_2$ then $Q_2 \to Q_1$. Hence, we only need to show that given a semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1} \rangle$,

(1) if $\mathcal{K} \in \mathbf{C}_{\mathrm{hom}}$ and Q_1, Q_2 are CQs such that $Q_2 \to Q_1$, then $Q_1 \subseteq_{\mathcal{K}} Q_2$;
(2) if $\mathcal{K} \notin \mathbf{C}_{\mathrm{hom}}$, then there exist CQs Q_1 and Q_2 such that $Q_2 \to Q_1$, but $Q_1 \not\subseteq_{\mathcal{K}} Q_2$.

For **Part 1**, assume that $Q_1 = \exists \mathbf{v}\, R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and $Q_2 = \exists \mathbf{w}\, S_1(\mathbf{q}_1, \mathbf{w}_1), \ldots, S_m(\mathbf{q}_m, \mathbf{w}_m)$, where $\mathbf{u}$ is the tuple of free variables of Q_1 and $\mathbf{q}$ is the tuple of free variables of Q_2, each $\mathbf{u}_i$ and $\mathbf{q}_j$ consist of variables from $\mathbf{u}$ and $\mathbf{q}$, respectively, and each $\mathbf{v}_i$ and $\mathbf{w}_j$ consist of variables from $\mathbf{v}$ and $\mathbf{w}$, respectively. We need to show that for an arbitrary $\mathcal{K}$-instance I and a tuple $\mathbf{t}$ the following holds:

$$Q_1^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \prod_{1 \leq i \leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \preceq_{\mathcal{K}} \sum_{g \in \mathcal{V}(Q_2, \mathbf{t})} \prod_{1 \leq j \leq m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) = Q_2^I(\mathbf{t}). \quad (2)$$

It is given that $Q_2 \to Q_1$, i.e. there exists a homomorphism h from Q_2 to Q_1. Without loss of generality, let us assume that when applying h to (the atoms of) Q_2 one obtains $R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$, or, in other words, that the first ℓ atoms of our enumeration of Q_1 are the image of h in Q_1. Let us write

$$Q_1^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \prod_{1 \leq i \leq \ell} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \prod_{\ell < i \leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)).$$

Let $\mathcal{V}(Q_1, \mathbf{t}) = \mathcal{V}_1 \cup \ldots \cup \mathcal{V}_k \cup \ldots$ be a (disjoint) partitioning of the set of mappings $\mathcal{V}(Q_1, \mathbf{t})$, such that f and f' are in the same $\mathcal{V}_k$ iff $f(\mathbf{u}_i, \mathbf{v}_i) = f'(\mathbf{u}_i, \mathbf{v}_i)$ for each $1 \leq i \leq \ell$. Rearranging the equation, we obtain

$$Q_1^I(\mathbf{t}) = \sum_{k \geq 1} \sum_{f \in \mathcal{V}_k} \prod_{1 \leq i \leq \ell} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \prod_{\ell < i \leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)).$$

Since $\otimes$ distributes over $\oplus$ and each mapping from every $\mathcal{V}_k$ maps the variables of the first ℓ atoms of our enumeration of Q_2 to the same constants, we have that

$$Q_1^I(\mathbf{t}) = \sum_{k \geq 1} \prod_{1 \leq i \leq \ell} R_i^I(f_k(\mathbf{u}_i, \mathbf{v}_i)) \left(\sum_{f \in \mathcal{V}_k} \prod_{\ell < i \leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \right),$$

where f_k is just an arbitrary representative from $\mathcal{V}_k$. Since $R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ is the image of the atoms of Q_2 by h, just as done in [Ioannidis and Ramakrishnan 1995] from $\otimes$-idempotence we conclude that

$$Q_1^I(\mathbf{t}) = \sum_{k\geq 1} \prod_{1\leq j\leq m} S_j^I(f_k \circ h(\mathbf{q}_j, \mathbf{w}_j)) \left(\sum_{f\in\mathcal{V}_k} \prod_{\ell<i\leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \right).$$

We consider only instances with finite support, so only a finite number of the outer summands are not equal to $\mathbb{0}$. Hence, we can apply **Part 1** of Lem. 3.6, so

$$Q_1^I(\mathbf{t}) \preceq_{\mathcal{K}} \sum_{k\geq 1} \prod_{1\leq j\leq m} S_j^I(f_k \circ h(\mathbf{q}_j, \mathbf{w}_j)).$$

Since for all $k \geq 1$ we have $f_k \circ h \in \mathcal{V}(Q_2, t)$, the desired inequality (2) holds.

For **Part 2**, we need to show that given a semiring $\mathcal{K} \notin \mathbf{C}_{\text{hom}}$ there exist CQs Q_1 and Q_2 such that there is a homomorphism from Q_2 to Q_1, but $Q_1 \not\subseteq_{\mathcal{K}} Q_2$. There are three possibilities.

(a). Semiring $\mathcal{K}$ may fail to satisfy $x \times x =_{\mathcal{K}} x$, witnessed by $a \in K$ such that $a \not\preceq_{\mathcal{K}} a \otimes a$. Consider a schema with a unary relation R and CQs $Q_1 = \exists u\ R(u)$ and $Q_2 = \exists u, v\ R(u), R(v)$. Clearly $Q_2 \to Q_1$. However, for the $\mathcal{K}$-instance I such that $R^I(c) = a$ for some constant c and $R^I(c') = \mathbb{0}$ for all $c' \neq c$ it follows that $Q_1^I() = a$ and $Q_2^I() = a \otimes a$, so $Q_1 \not\subseteq_{\mathcal{K}} Q_2$.

(b). Alternatively, $\mathcal{K}$ may fail to satisfy $x \times x =_{\mathcal{K}} x$, witnessed by $a \in K$ such that $a \otimes a \neq a$ but $a \preceq_{\mathcal{K}} a \otimes a$. For the CQs Q_1 and Q_2 from case *(a)* we have that $Q_1 \to Q_2$, but since $Q_2^I() = a \otimes a \not\preceq_{\mathcal{K}} a = Q_1^I()$ for the $\mathcal{K}$-instance I, we conclude that $Q_2 \not\subseteq_{\mathcal{K}} Q_1$.

(c). Finally, $\mathcal{K}$ may fail to satisfy $1 + x =_{\mathcal{K}} 1$. Then by **Part 2** of Lem. 3.6 there exist $a, b \in K$ such that $a \otimes b \not\preceq_{\mathcal{K}} a$. In this case, consider CQs $Q_1 = \exists v\ R(v), S(v)$ and $Q_2 = \exists v\ R(v)$ over a schema with two unary relations R and S. Clearly $Q_2 \to Q_1$. However, for the $\mathcal{K}$-instance I such that $R^I(c) = a$, $S^I(c) = b$ for some constant c and $R^I(c') = \mathbb{0}$, $S^I(c') = \mathbb{0}$ for all $c' \neq c$, we have that $Q_1^I() = a \otimes b \not\preceq_{\mathcal{K}} a = Q_2^I()$. Hence, $Q_1 \not\subseteq_{\mathcal{K}} Q_2$.

This ends the proof of the theorem. $\square$

Deciding the existence of a homomorphism between CQs is well-known to be NP-complete [Aho et al. 1979, Thm. 7(1)]. We therefore obtain the following corollary.

COROLLARY 3.8. *If $\mathcal{K} \in \mathbf{C}_{\text{hom}}$ then* CQ $\mathcal{K}$-CONTAINMENT *is* NP-*complete.*

Many semirings used for annotations are distributive lattices, and hence belong to $\mathbf{C}_{\text{hom}}$. Besides the set semantics $\mathcal{B}$, they include the semiring of positive boolean expressions $\texttt{PosBool}[X]$ described by Green et al. [2007], which is used in incomplete databases [Imieliński and Lipski 1984], and the probabilistic semiring $\mathcal{P}[\Omega]$ used in event tables [Fuhr and Rölleke 1997; Zimányi 1997]. To the best of our knowledge none of the semirings that belong to $\mathbf{C}_{\text{hom}}$ but are not distributive lattices have been proposed for use in practice, although one can easily construct an infinite number of them. This can be done by taking any of the distributive lattices mentioned above, equipped with a partial order that is not natural, i.e. any order that does not satisfy the axiom $a \preceq_{\mathcal{K}}^{\text{nat}} b \iff \exists c\ a \oplus c = b$ (but of course that still satisfies our positivity requirements).

Also, the class $\mathbf{C}_{\text{hom}}$ corresponds precisely to the class of *type* A$'$ *systems* introduced by Ioannidis and Ramakrishnan [1995]. They raised the question of what the decision procedure is for CQ containment over such systems. Our Thm. 3.7 answers this question. However, many annotation semirings do not belong to $\mathbf{C}_{\mathrm{hom}}$, including provenance polynomials $\mathcal{N}[X]$, the why-provenance semiring $\texttt{Why}[X]$ discussed by Buneman et al. [2001], or bag semantics $\mathcal{N}$ [Chaudhuri and Vardi 1993]. In the next section, we study what happens when we relax the conditions for $\mathbf{C}_{\mathrm{hom}}$.

4. $\mathcal{K}$-CONTAINMENT OF CQS

From a practical point of view, it would be useful to have a decision procedure for $\mathcal{K}$-containment of CQs for an arbitrary semiring $\mathcal{K}$. However, as we have mentioned in the introduction, there is evidence that obtaining such a procedure for all semirings not in $\mathbf{C}_{\mathrm{hom}}$ is a truly challenging, if not impossible, task. The semiring $\mathcal{N} = \langle \mathbb{N}_0, +, \times, 0, 1\rangle$ of natural numbers with zero, with the usual arithmetic operations and the natural order, is used to model bag semantics [Green et al. 2007]. A universal decision procedure for CQ $\mathcal{K}$-CONTAINMENT would thus require being able to solve this problem for the special case of bag semantics $\mathcal{N}$, which is a long-standing open problem [Chaudhuri and Vardi 1993; Ioannidis and Ramakrishnan 1995]. It is also not difficult to show that there are infinitely many semirings $\mathcal{K}$ for which deciding $\mathcal{K}$-containment of CQs is at least as hard as for bag semantics, in terms of computational complexity.

With these observations in mind, we instead ask the following, narrower question: are there any reasonable classes of semirings for which we can prove that $\mathcal{K}$-containment of CQs is decidable? We have already pointed out that this is the case for the class $\mathbf{C}_{\mathrm{hom}}$, since for all semirings $\mathcal{K}$ in $\mathbf{C}_{\mathrm{hom}}$ the problem of $\mathcal{K}$-containment can be solved by deciding the existence of a homomorphism. A natural starting point for our search is therefore to relax the axioms of the class $\mathbf{C}_{\mathrm{hom}}$. We thus obtain the class of semirings that satisfy the $\otimes$-idempotence axiom, that we denote by $\mathbf{S}_{\mathrm{hcov}}$; the class of semirings that satisfy the $\mathbb{1}$-annihilation axiom, denoted by $\mathbf{S}_{\mathrm{in}}$; and, if we relax both axioms, the class $\mathbf{S}_{\preceq}$ of all ($\preceq$-positive) semirings.

We show that for each of these classes there exists a natural type of homomorphism that characterizes the class, but only as a *sufficient condition* for $\mathcal{K}$-containment of CQs. In the search for classes similar to $\mathbf{C}_{\mathrm{hom}}$, we then provide the *largest* class of semirings for which each of these conditions is *necessary* for $\mathcal{K}$-containment, resulting in analogues of Thm. 3.7 for different classes of semirings and different types of homomorphisms.

Besides $\mathbf{S}_{\mathrm{hcov}}$, $\mathbf{S}_{\mathrm{in}}$, and $\mathbf{S}_{\preceq}$ we look at one more class, that we denote by $\mathbf{S}_{\mathrm{sur}}$. This class lies "between" $\mathbf{S}_{\mathrm{hcov}}$ and $\mathbf{S}_{\preceq}$, in the sense that it can be obtained from $\mathbf{S}_{\mathrm{hcov}}$ by a partial, instead of complete, relaxation of the $\otimes$-idempotence axiom. The class $\mathbf{S}_{\mathrm{sur}}$ is interesting in its own right, since it can be characterized by the well studied notion of *surjective* homomorphism ([Chaudhuri and Vardi 1993; Ioannidis and Ramakrishnan 1995]) as yielding a sufficient condition for CQ $\mathcal{N}$-containment. In the same fashion, we identify the largest class of semirings for which this condition is also necessary.

All the axioms for necessary classes are based on the notion of CQ-admissible polynomials. We first opt for a conceptual non-constructive definition, but give a syntactic characterization in a separate subsection.

The results about the classes above are summarized in Tab. I on page 37, which can be used as a roadmap of Sec. 4.1–4.4.

Notice that, up to this point, we have only considered solving the $\mathcal{K}$-containment problem by means of finding different types of homomorphisms between CQs. Thus, it is natural to ask whether there exists a different approach for solving this problem. We address this question at the end of this section, and show that there exists a large class of semirings which possesses a *small model property*: if a CQ Q_1 is not $\mathcal{K}$-contained in a CQ Q_2, then this is witnessed by a small enough $\mathcal{K}$-instance.

4.1. Containment by homomorphic covering

We begin with the class of $\otimes$-idempotent semirings.

Definition 4.1 (*Class* $\mathbf{S}_{\mathrm{hcov}}$ *of semirings*). Let $\mathbf{S}_{\mathrm{hcov}}$ be the class of all semirings in $\mathbf{S}_{\preceq}$ that satisfy the $\otimes$-idempotence axiom:

$$x \times x =_{\mathcal{K}} x.$$

For these semirings, we exploit the notion of homomorphic covering: given CQs Q_1 and Q_2, we say that Q_2 *homomorphically covers* Q_1, and write $Q_2 \rightrightarrows Q_1$, if for every atom $R(\mathbf{u})$ in Q_1 there exists a homomorphism h from Q_2 to Q_1 with $R(\mathbf{u})$ in the image of h; or, more formally, if for every such atom in Q_1 there is a homomorphism h from Q_2 to Q_1 and an atom $R(\mathbf{v})$ in Q_2 such that $h(\mathbf{v}) = \mathbf{u}$.

This type of homomorphism arose in the context of query optimization as a necessary condition for $\mathcal{N}$-containment of CQs over bag semantics $\mathcal{N}$ [Chaudhuri and Vardi 1993]. It was also noted that existence of a homomorphic covering is not sufficient to guarantee $\mathcal{N}$-containment. Homomorphic coverings were also used by Green [2011] to show that $Q_2 \rightrightarrows Q_1$ is both necessary and sufficient for $Q_1 \subseteq_{\mathtt{Lin}[X]} Q_2$, where $\mathtt{Lin}[X]$ is the lineage semiring [Cui et al. 2000; Buneman et al. 2001]. This semiring is used to model propagation of comments of arbitrary nature.

In this section we establish axiomatic bounds for semirings to have homomorphic covering as a sufficient and as a necessary condition for $\mathcal{K}$-containment of CQs. We start with the first part and show that the class $\mathbf{S}_{\mathrm{hcov}}$ captures precisely all semirings for which $Q_2 \rightrightarrows Q_1$ is a sufficient condition.

PROPOSITION 4.2. *The following are equivalent:*

- *semiring* $\mathcal{K}$ *belongs to* $\mathbf{S}_{\mathrm{hcov}}$*;*
- $Q_2 \rightrightarrows Q_1$ *implies* $Q_1 \subseteq_{\mathcal{K}} Q_2$*, for all CQs* Q_1, Q_2.

PROOF. We need to show that given a semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1} \rangle$,

(1) if $\mathcal{K} \in \mathbf{S}_{\mathrm{hcov}}$ and Q_1, Q_2 are CQs such that $Q_2 \rightrightarrows Q_1$, then $Q_1 \subseteq_{\mathcal{K}} Q_2$;
(2) if $\mathcal{K} \notin \mathbf{S}_{\mathrm{hcov}}$, then there exist CQs Q_1 and Q_2 such that $Q_2 \rightrightarrows Q_1$, but $Q_1 \not\subseteq_{\mathcal{K}} Q_2$.

For **Part 1**, we assume that $Q_1(\mathbf{u}) = \exists \mathbf{v}\, R_1(\mathbf{u}_1, \mathbf{v}_1), \dots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and $Q_2(\mathbf{q}) = \exists \mathbf{w}\, S_1(\mathbf{q}_1, \mathbf{w}_1), \dots, S_m(\mathbf{q}_m, \mathbf{w}_m)$, where $\mathbf{u}$ is the tuple of free variables of Q_1 and $\mathbf{q}$ is the tuple of free variables of Q_2, each $\mathbf{u}_i$ and $\mathbf{q}_j$ consist of variables from $\mathbf{u}$ and $\mathbf{q}$, respectively, and each $\mathbf{v}_i$ and $\mathbf{w}_j$ consist of variables from $\mathbf{v}$ and $\mathbf{w}$, respectively. We need to show that for an arbitrary $\mathcal{K}$-instance I and a tuple $\mathbf{t}$, we have

$$Q_1^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \prod_{1 \le i \le n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \preceq_{\mathcal{K}} \sum_{g \in \mathcal{V}(Q_2, \mathbf{t})} \prod_{1 \le j \le m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) = Q_2^I(\mathbf{t}). \tag{3}$$

We need some extra notation. Since Q_2 homomorphically covers Q_1, let $h_1, \dots, h_n$ be the (not necessarily distinct) homomorphisms from Q_2 to Q_1 such that, for each $1 \le i \le n$, the atom $R_i(\mathbf{u}_i, \mathbf{v}_i)$ is in the image of h_i. Consider the following set $\mathcal{V}^1(Q_2, \mathbf{t})$ of mappings:

$$\mathcal{V}^1(Q_2, \mathbf{t}) = \{g \mid g = f \circ h_i, \text{ where } f \in \mathcal{V}(Q_1, \mathbf{t}), 1 \le i \le n\}.$$

It is easy to show that $\mathcal{V}^1(Q_2, \mathbf{t}) \subseteq \mathcal{V}(Q_2, \mathbf{t})$ (this was also proved by Ioannidis and Ramakrishnan [1995]). For the left part of the inequality (3) we have

$$Q_1^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \prod_{1 \le i \le n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \prod_{1 \le i \le n} \prod_{1 \le j \le m} S_j^I(f \circ h_i(\mathbf{q}_j, \mathbf{w}_j)). \tag{4}$$

The second of these equalities holds by $\otimes$-idempotence of $\mathcal{K}$: indeed, for each $f \in \mathcal{V}(Q_1, \mathbf{t})$ the product of $S_j^I(f \circ h_i(\mathbf{q}_j, \mathbf{w}_j))$ multiplies the same values as the corresponding product of $R_i^I(f(\mathbf{u}_i, \mathbf{v}_i))$, but with greater or equal exponents. If we define, for mappings $g_1, \ldots, g_n$ from $\mathcal{V}^1(Q_2, \mathbf{t})$,

$$\hat{\Pi}(g_1, \ldots, g_n) = \prod_{1 \leq i \leq n} \prod_{1 \leq j \leq m} S_j^I(g_i(\mathbf{q}_j, \mathbf{w}_j)),$$

then we can arrange the right part of (4) to obtain:

$$Q_1^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \hat{\Pi}(f \circ h_1, \ldots, f \circ h_n).$$

From positivity of $\mathcal{K}$ and the fact that the number of tuples in I with annotation greater than 0 is finite, one can show that:

$$\sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \hat{\Pi}(f \circ h_1, \ldots, f \circ h_n) \preceq_{\mathcal{K}} \sum_{g_1, \ldots, g_n \in \mathcal{V}^1(Q_2, \mathbf{t})} \hat{\Pi}(g_1, \ldots, g_n). \tag{5}$$

Indeed, since for each $f \in \mathcal{V}(Q_1, \mathbf{t})$ all of $f \circ h_1, \ldots, f \circ h_n$ are in $\mathcal{V}^1(Q_2, \mathbf{t})$, to show that equation (5) holds we only need to prove that on the left side no summand occurs more times than on the right side, i.e. that for every different $f, f' \in \mathcal{V}(Q_1, \mathbf{t})$ it holds that $f \circ h_1, \ldots, f \circ h_n$ and $f' \circ h_1, \ldots, f' \circ h_n$ are not completely the same. Assume the contrary. Since $f \neq f'$, there exists a variable v in Q_1 such that $f(v) \neq f'(v)$. Let v be among the variables of an atom $R_i(\mathbf{u}_i, \mathbf{v}_i)$, $1 \leq i \leq n$, of Q_1. But this atom is in the image of h_i, and thus Q_2 contains an atom $S_j(\mathbf{q}_j, \mathbf{w}_j)$, $1 \leq j \leq m$, such that $h_i(\mathbf{q}_j, \mathbf{w}_j) = (\mathbf{u}_i, \mathbf{v}_i)$. Hence $f \circ h_i$ differs from $f' \circ h_i$, which is the desired contradiction.

Continuing with the proof, notice that the right part of equation (5) can be rearranged as follows:

$$\sum_{g_1, \ldots, g_n \in \mathcal{V}^1(Q_2, \mathbf{t})} \hat{\Pi}(g_1, \ldots, g_n) = \left(\sum_{g \in \mathcal{V}^1(Q_2, \mathbf{t})} \prod_{1 \leq j \leq m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) \right)^n.$$

Finally, we obtain equation (3) from $\otimes$-idempotence and positivity of $\mathcal{K}$:

$$\left(\sum_{g \in \mathcal{V}^1(Q_2, \mathbf{t})} \prod_{1 \leq j \leq m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) \right)^n = \sum_{g \in \mathcal{V}^1(Q_2, \mathbf{t})} \prod_{1 \leq j \leq m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) \preceq_{\mathcal{K}}$$

$$\sum_{g \in \mathcal{V}(Q_2, \mathbf{t})} \prod_{1 \leq j \leq m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) = Q_2^I(\mathbf{t}).$$

This shows **Part 1** of the proposition.

For **Part 2** we reuse the CQs Q_1 and Q_2 from the cases *(a)* and *(b)* of the proof for **Part 2** of Thm. 3.7. Clearly, we have $Q_2 \rightrightarrows Q_1$ and $Q_1 \rightrightarrows Q_2$ but by the same reasons $Q_1 \not\subseteq_{\mathcal{K}} Q_2$ and $Q_2 \not\subseteq_{\mathcal{K}} Q_1$. □

Of course, a sufficient condition itself does not guarantee the decidability of the $\mathcal{K}$-containment problem; one needs such a condition to be necessary as well. Since one can easily find semirings in $\mathbf{S}_{\text{hcov}}$ for which the existence of a homomorphic covering is not a necessary condition (for example, any semiring in $\mathbf{C}_{\text{hom}}$), our only hope is to describe the *largest* class for which a homomorphic covering is necessary for $\mathcal{K}$-containment of CQs. Next we present two different characterizations of this class. The first one is just an axiomatization similar to the definition of previous classes of semirings, and the second one is based on an interesting characterization of polynomials. We show that both definitions end up being equivalent, and then proceed to show that this class satisfies our desiderata.

Definition 4.3 (*Class* $\mathbf{N}_{\text{hcov}}$ *of semirings*). Denote by $\mathbf{N}_{\text{hcov}}$ the class of semirings $\mathcal{K}$ such that for every $n, k \geq 1$ it holds that

$$x_1 \times \ldots \times x_n \times y \not\preceq_{\mathcal{K}} (x_1 + \ldots + x_n)^k$$

(again, assuming all variables to be universally quantified).

The above definition implies, in particular, that $x \times y \not\preceq_{\mathcal{K}} x$. By Lem. 3.6 none of the semirings in $\mathbf{N}_{\text{hcov}}$ can then satisfy the $\mathbb{1}$-*annihilation* axiom. Hence $\mathbf{N}_{\text{hcov}} \cap \mathbf{C}_{\text{hom}} = \emptyset$. This is to be expected, since we are describing a class of semirings for which homomorphic covering is a necessary condition for $\mathcal{K}$-containment of CQs. Nevertheless, one can verify that many interesting semirings belong to $\mathbf{N}_{\text{hcov}}$, such as bag semantics $\mathcal{N}$. In fact, it was already proved by Chaudhuri and Vardi [1993] that homomorphic covering is a necessary condition for $\mathcal{N}$-containment of CQs.

Next we give a different characterization for the class $\mathbf{N}_{\text{hcov}}$. It is based on the following definition.

Definition 4.4. A polynomial P from $\mathbb{N}[X]$ is ***CQ-admissible*** iff there exists a CQ Q, an $\mathcal{N}[X]$-instance I each tuple of which is annotated with either a unique variable from X or 0, and a tuple $\mathbf{t}$, such that $Q^I(\mathbf{t}) = \mathsf{P}$.

Essentially, a polynomial is CQ-admissible if it is possible to obtain it by a CQ on an *abstractly tagged* instance ([Green et al. 2007]).

Example 4.5. The following shows that the polynomial $\mathsf{P} = x^2 + xy$ is CQ-admissible. Consider the $\mathcal{N}[X]$-instance I over a schema with one binary relation R, in which $R^I(a,a) = x$, $R^I(a,b) = y$ for some elements $a, b \in \mathbb{D}$, and all other annotations are set to 0. Then for the query $Q = \exists u, v\ R(u,u), R(u,v)$ we have that $Q^I() = \mathsf{P}$.

In Sec. 5 we will see that the polynomial $x^2 + xy + y^2$, for instance, is not CQ-admissible.

We write $\mathbb{N}^{cq}[X]$ for the set of all CQ-admissible polynomials with variables X. We will use this notion intensively in the rest of this paper; for now, we have opted to give a non-constructive definition, but we will give an algebraic characterization of $\mathbb{N}^{cq}[X]$ in Sec. 5. In this section, however, we present an interesting observation regarding CQ-admissible polynomials. The following proposition essentially shows that the above definition can be stated only in terms of conjunctive queries without free variables, such as the query Q in the example above.

PROPOSITION 4.6. *For every CQ-admissible polynomial* P *there exists a CQ* Q *without free variables and* $\mathcal{N}[X]$*-instance* I *with only unique variables or* 0 *as annotations such that* $Q^I() = \mathsf{P}$.

PROOF. This proposition is an immediate corollary of Prop. 5.1 from Sec. 5, which is somewhat technical. That is why, for the sake of clarity we left Prop. 5.1 for a separate section. □

Using the notion of CQ-admissible polynomials, we are ready to give the alternative characterization of class $\mathbf{N}_{\text{hcov}}$.

LEMMA 4.7. *A semiring $\mathcal{K}$ belongs to $\mathbf{N}_{\text{hcov}}$ iff for every CQ-admissible polynomial* P *over a set of variables $X = \{x_1, \dots, x_n\}$ the inequality*

$$x_1 \times \ldots \times x_n \preceq_{\mathcal{K}} \mathsf{P}$$

implies that every variable from X occurs in P.

PROOF. Let Q be a CQ consisting of k atoms, and I be an $\mathcal{N}[X]$-instance with tuples annotated with unique variables $x_1, \dots, x_n$ (or zero), such that the mappings in the set $\mathcal{V}(Q, \mathbf{t})$ allow us to obtain any possible combination of images of the atoms of Q to non-zero annotated tuples of I. Then $Q^I(\mathbf{t}) = (x_1 + \ldots + x_n)^k$. Therefore, this polynomial is CQ-admissible. Moreover, every P from $\mathbb{N}^{cq}[X]$ of degree k satisfies $\mathsf{P} \preceq_{\mathcal{N}[X]} (x_1 + \ldots + x_n)^k$. By universality of $\mathcal{N}[X]$ for $\mathbf{S}_{\preceq}$, each of the semirings in $\mathbf{S}_{\preceq}$ inherits the natural order of $\mathcal{N}[X]$, so for every polynomial $\mathsf{P} \in \mathbb{N}^{cq}[X]$ of degree $k \geq 1$ we have that

$$\mathsf{P} \preceq_{\mathcal{K}} (x_1 + \ldots + x_n)^k. \tag{6}$$

By the definition, $\mathcal{K} \in \mathbf{N}_{\text{hcov}}$ iff for all $n, k \geq 1$ it holds that

$$x_1 \times \ldots \times x_n \times y \not\preceq_{\mathcal{K}} (x_1 + \ldots + x_n)^k.$$

Since the polynomial $(x_1 + \cdots + x_n)^k$ is CQ-admissible, from (6) and transitivity of $\preceq_{\mathcal{K}}$ we conclude that $\mathcal{K} \in \mathbf{N}_{\text{hcov}}$ iff for all $n, k \geq 1$, set of variables $X = \{x_1, \dots, x_n\}$, another variable y, and $\mathsf{P} \in \mathbb{N}^{cq}[X]$ of degree k it holds that

$$x_1 \times \ldots \times x_n \times y \not\preceq_{\mathcal{K}} \mathsf{P}. \tag{7}$$

This means that $\mathcal{K} \in \mathbf{N}_{\text{hcov}}$ iff (7) holds for all $n \geq 1$, set of variables $X = \{x_1, \dots, x_n\}$ and variable y, and $\mathsf{P} \in \mathbb{N}^{cq}[X]$.

Next we will show that the second part of this statement is equivalent to the statement of the lemma: for every polynomial $\mathsf{P} \in \mathbb{N}^{cq}[X]$ over a set of variables $X = \{x_1, \dots, x_n\}$, the inequality

$$x_1 \times \ldots \times x_n \preceq_{\mathcal{K}} \mathsf{P} \tag{8}$$

implies that P uses all the variables in X.

For the **“if”** direction, consider a polynomial $\mathsf{P} \in \mathbb{N}^{cq}[X]$ such that (8) holds, and assume for the sake of contradiction that P uses only $x_1, \dots, x_\ell$ for some $\ell < n$. Since (8) is an universal axiom, we may assume that $\ell = n - 1$. But then this contradicts (7). Hence, every variable in X must occur in P.

The **“only if”** direction is immediate. □

Lastly, we are now able to present a formal proof that the class $\mathbf{N}_{\text{hcov}}$ is the largest class of semirings for which the notion of homomorphic covering is a necessary condition for $\mathcal{K}$-containment of CQs.

PROPOSITION 4.8. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{N}_{\text{hcov}}$;*
- *$Q_1 \subseteq_{\mathcal{K}} Q_2$ implies $Q_2 \rightrightarrows Q_1$, for all CQs Q_1, Q_2.*

PROOF. We need to show, that given a semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1} \rangle$,

(1) if $\mathcal{K} \in \mathbf{N}_{\text{hcov}}$ and Q_1, Q_2 are CQs such that $Q_1 \subseteq_{\mathcal{K}} Q_2$, then $Q_2 \rightrightarrows Q_1$;
(2) if $\mathcal{K} \notin \mathbf{N}_{\text{hcov}}$, then there are CQs Q_1 and Q_2 such that $Q_1 \subseteq_{\mathcal{K}} Q_2$, but there is no homomorphic covering from Q_2 to Q_1.

For **Part 1**, we assume that $Q_1(\mathbf{u}) = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \dots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and $Q_2(\mathbf{q}) = \exists \mathbf{w}\ S_1(\mathbf{q}_1, \mathbf{w}_1), \dots, S_m(\mathbf{q}_m, \mathbf{w}_m)$, where $\mathbf{u}$ is the tuple of free variables of Q_1 and $\mathbf{q}$ is the tuple of free variables of Q_2, each $\mathbf{u}_i$ and $\mathbf{q}_j$ consist of variables from $\mathbf{u}$ and $\mathbf{q}$, respectively, and each $\mathbf{v}_i$ and $\mathbf{w}_j$ consist of variables from $\mathbf{v}$ and $\mathbf{w}$, respectively. Let $X = \{x_1, \ldots, x_n\}$ be a set of variables. Consider a canonical $\mathcal{N}[X]$-instance $[\![Q_1]\!]$ of Q_1 (see the definition in Sec. 3.3). Denote $\mathsf{P}_1 = Q_1^{[\![Q_1]\!]}(\mathbf{u})$ and $\mathsf{P}_2 = Q_2^{[\![Q_1]\!]}(\mathbf{u})$. By construction of $[\![Q_1]\!]$, using the identity mapping in $\mathcal{V}(Q_1, \mathbf{u})$ it is clear that $x_1 \times \ldots \times x_n$ is a monomial in P_1 and by positivity of $\mathcal{K}$ we have that $x_1 \times \ldots \times x_n \preceq_{\mathcal{K}} \mathsf{P}_1$.

Moreover, the polynomial P_2 is CQ-admissible. This does not follow directly from the fact that $\mathsf{P}_2 = Q_2^{[\![Q_1]\!]}(\mathbf{u})$, because the tuples in $[\![Q_1]\!]$ may not be annotated by single variables. However, extensions Q_2' of Q_2 and Q_1' of Q_1 can be constructed to show that P_2 is CQ-admissible, i.e. so that all the tuples in $[\![Q_1']\!]$ are annotated with distinct variables and $(Q_2')^{[\![Q_1']\!]}(\mathbf{u}) = \mathsf{P}_2$. In order to do this, construct the query Q_1' by replacing each atom of form $R_i(\mathbf{u}_i, \mathbf{v}_i)$, $1 \le i \le n$, by an atom $R_i'(\mathbf{u}_i, \mathbf{v}_i, s)$, where s is a fresh existentially quantified variable, and query Q_2' by replacing $S_j(\mathbf{q}_i, \mathbf{w}_i)$, for $1 \le j \le m$, by an atom $S_j'(\mathbf{q}_i, \mathbf{w}_i, t)$, where t is again a fresh existentially quantified variable. One can then verify that the evaluation of Q_2' over $[\![Q_1']\!]$ yields exactly the polynomial P_2.

Since $Q_1 \subseteq_{\mathcal{K}} Q_2$, for every $\mathcal{K}$-instance I and tuple $\mathbf{t}$ we have that $Q_1^I(\mathbf{t}) \preceq_{\mathcal{K}} Q_2^I(\mathbf{t})$. In particular, this holds for every $\mathcal{K}$-instance obtained from $[\![Q_1]\!]$ by an evaluation $X \to K$. Hence $\mathsf{P}_1 \preceq_{\mathcal{K}} \mathsf{P}_2$, and in particular, $x_1 \times \ldots \times x_n \preceq_{\mathcal{K}} \mathsf{P}_2$. Since $\mathcal{K} \in \mathbf{N}_{\text{hcov}}$, by Lem. 4.7 we have that every variable in X occurs in P_2. By definition,

$$\mathsf{P}_2 = Q_2^{[\![Q_1]\!]}(\mathbf{u}) = \sum_{f \in \mathcal{V}(Q_2, \mathbf{u})} \prod_{1 \le j \le m} S_j^{[\![Q_1]\!]}(f(\mathbf{q}_j, \mathbf{w}_j)).$$

Since $[\![Q_1]\!]$ is an $\mathcal{N}[X]$-instance, for every atom $R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ from Q_1 (for which $R_\ell^{[\![Q_1]\!]}(\mathbf{u}_\ell, \mathbf{v}_\ell) = x_{\ell_1} + \ldots + x_{\ell_k}$), there exists a mapping $f \in \mathcal{V}(Q_2, \mathbf{u})$ such that $f(\mathbf{q}) = \mathbf{u}$ and $R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ is in the image of f. Hence, from every such mapping one can construct a homomorphism from Q_2 to Q_1 which has $R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ in its image. This holds for every atom $R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ from Q_1, and thus all these homomorphisms form a homomorphic covering of Q_1 by Q_2.

For **Part 2**, we need to construct CQs Q_1 and Q_2 such that $Q_1 \subseteq_{\mathcal{K}} Q_2$ but $Q_2 \not\Rightarrow Q_1$. Let $\mathcal{K}$ be a semiring not in $\mathbf{N}_{\text{hcov}}$. By Lem. 4.7 there is a CQ-admissible polynomial $\mathsf{P} \in \mathbb{N}^{cq}[X]$ such that

$$x_1 \times \ldots \times x_n \preceq_{\mathcal{K}} \mathsf{P}, \tag{9}$$

but P does not use all the variables $x_1, \ldots, x_n$. Without loss of generality, we may assume that $X = \{x_1, \ldots, x_n\}$.

By Prop. 4.6, there exists a CQ Q without free variables and an $\mathcal{N}[X]$-instance I with only unique variables or 0 as annotations, such that $Q^I() = \mathsf{P}$. Assume that the schema of I has k relations $R_1, \ldots, R_k$, and let $\mathbf{u}$ be the tuple of all elements of the domain of I which occur in tuples having annotations different from 0 in the $\mathcal{N}[X]$-relations of I. Construct a schema $\mathbb{S}$ as follows: for every $\mathcal{N}[X]$-relation R_i of arity n_i, $1 \le i \le k$, in the schema of I, add to $\mathbb{S}$ a relational symbol S_i of arity $n_i + |\mathbf{u}|$.

Suppose $Q = \exists \mathbf{v}\, R_1'(\mathbf{v}_1), \ldots, R_m'(\mathbf{v}_m)$. Then let $Q_2 = \exists \mathbf{v}\, S_1'(\mathbf{u}, \mathbf{v}_1), \ldots, S_m'(\mathbf{u}, \mathbf{v}_m)$, be a query with free variables $\mathbf{u}$, where each of S_j', $1 \le j \le m$, has the relational symbol corresponding to the relational symbol constructed from R_j' (here we abuse notation and look at the constants $\mathbf{u}$ as variables). In turn, we construct a CQ Q_1 with free variables $\mathbf{u}$ without existential variables in the following way: for each relation R_i and each tuple $\mathbf{q}$ such that $R_i^I(\mathbf{q}) \in X$ the query Q_1 contains the atom $S_i(\mathbf{q}, \mathbf{u})$.

Next we show that $Q_2 \not\Rightarrow Q_1$. First, note that the instance I and the canonical $\mathcal{N}[X]$-instance $[\![Q_1]\!]$ of Q_1 are related in the following way: for every $1 \le i \le k$ and tuple $\mathbf{q}$ it holds that $S_i^{[\![Q_1]\!]}(\mathbf{q}, \mathbf{u}) = R_i^I(\mathbf{q})$, and all other annotations in I or $[\![Q_1]\!]$ are 0. Moreover, from the construction of Q_1 and Q_2 we have

$$Q_1^{[\![Q_1]\!]}(\mathbf{u}) = x_1 \times \ldots \times x_n \text{ and } Q_2^{[\![Q_1]\!]}(\mathbf{u}) = \mathsf{P}. \tag{10}$$

There exists a bijective correspondence between mappings f from $\mathcal{V}(Q_2, \mathbf{u})$ which produce monomials in $Q_2^{[\![Q_1]\!]}(\mathbf{u}) = \mathsf{P}$ (not equal to 0) and homomorphisms from Q_2 to Q_1. Since P does not use all the variables $x_1, \ldots, x_n$, there exists an atom $S(\mathbf{q}, \mathbf{u})$ in Q_1, which is annotated in $[\![Q_1]\!]$ by a variable $x \in X$ missed in P, such that there are no homomorphisms from Q_2 to Q_1 with $S(\mathbf{q}, \mathbf{u})$ in the images. This means that there is no homomorphic covering of Q_1 by Q_2.

Finally, we need to show that $Q_1 \subseteq_{\mathcal{K}} Q_2$, i.e. for each $\mathcal{K}$-instance J and tuple $\mathbf{t}$ of size $|\mathbf{u}|$,

$$Q_1^J(\mathbf{t}) \preceq_{\mathcal{K}} Q_2^J(\mathbf{t}). \tag{11}$$

Let J and $\mathbf{t}$ be arbitrary $\mathcal{K}$-instance and tuple. If $Q_1^J(\mathbf{t}) = \mathbb{0}$ then (11) automatically holds. So, let $Q_1^J(\mathbf{t}) \neq \mathbb{0}$. Since Q_1 does not have existential variables, the set $\mathcal{V}(Q_1, \mathbf{t})$ contains only one mapping f. Hence, $Q_1^J(\mathbf{t}) = a_1 \otimes \ldots \otimes a_n$ where $a_1, \ldots, a_n$ are the annotations of the images by f of atoms from Q_1 in such an order that for every ℓ, $1 \leq \ell \leq n$, if f maps an atom of Q_1 to a_ℓ then the only mapping from Q_1 to $[\![Q_1]\!]$ maps this atom to x_ℓ. This correspondence defines a function $\nu\colon X \to K$ by $\nu(x_\ell) = a_\ell$ for every ℓ. Moreover, by the fact that $\mathcal{K}$ is not trivial and (10) we have that $Q_2^J(\mathbf{t}) = \mathsf{Eval}_\nu(\mathsf{P}) \oplus a$, where a is some value from $\mathcal{K}$ obtained by a sum of valuations of all mappings from $\mathcal{V}(Q_2, \mathbf{t})$, which give images on J not completely containing in the image of f. Hereby, from (9) and positivity of $\mathcal{K}$ we conclude (11). This finishes the proof of the proposition. $\square$

Therefore, bag semantics $\mathcal{N}$ is in $\mathbf{N}_{\mathrm{hcov}}$, but not in $\mathbf{S}_{\mathrm{hcov}}$. However, $\mathtt{Lin}[X]$ is in both, and we have the following result for the class $\mathbf{C}_{\mathrm{hcov}} = \mathbf{S}_{\mathrm{hcov}} \cap \mathbf{N}_{\mathrm{hcov}}$[6] of all semirings which behave like $\mathtt{Lin}[X]$ w.r.t. $\mathcal{K}$-containment of CQs.

THEOREM 4.9. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{C}_{\mathrm{hcov}}$;*
- *$Q_1 \subseteq_{\mathcal{K}} Q_2$ iff $Q_2 \rightrightarrows Q_1$, for all CQs Q_1 and Q_2.*

We also know that checking for homomorphic covering between CQs is an NP-complete problem [Green 2011]. This gives us the following result.

COROLLARY 4.10. *If $\mathcal{K} \in \mathbf{C}_{\mathrm{hcov}}$ then* CQ $\mathcal{K}$-CONTAINMENT *is* NP*-complete.*

4.2. Containment by injective homomorphism

In this section we consider the class of semirings which satisfy the $\mathbb{1}$-annihilation axiom.

Definition 4.11 (*Class $\mathbf{S}_{\mathrm{in}}$ of semirings*). Denote by $\mathbf{S}_{\mathrm{in}}$ the class of all semirings in $\mathbf{S}_{\preceq}$ that satisfy the $\mathbb{1}$-annihilation axiom:

$$1 + x =_{\mathcal{K}} 1.$$

This class was considered implicitly in previous studies of containment on $\mathcal{K}$-relations [Green et al. 2007; Green 2011; Ioannidis and Ramakrishnan 1995], and has notable applications. In the context of the Semantic Web it was shown by Buneman

[6]Since $\otimes$-idempotence defines $\mathbf{S}_{\mathrm{hcov}}$, the exponent k may be omitted from the necessary condition of $\mathbf{C}_{\mathrm{hcov}}$.

and Kostylev [2010] that $\mathbf{S}_{\text{in}}$ is the class of all semirings which can be safely used as annotation domains for RDF data while respecting the inference system of RDFS. An extension of the SPARQL query language for querying annotated RDF data then followed in [Zimmermann et al. 2011], entailing a need to solve optimization problems for this class of semirings. As an example of a semiring which is in $\mathbf{S}_{\text{in}}$, but not in $\mathbf{C}_{\text{hom}}$, we give the *tropical* semiring $\mathcal{T}^+ = \langle \mathbb{N}_0 \cup \{\infty\}, \min, +, \infty, 0\rangle$ (with its natural order).

To study the class $\mathbf{S}_{\text{in}}$, we introduce the notion of injective homomorphism: given CQs $Q_1 = \exists \mathbf{v}_1\ \phi_1(\mathbf{u}_1, \mathbf{v}_1)$ and $Q_2 = \exists \mathbf{v}_2\ \phi_2(\mathbf{u}_2, \mathbf{v}_2)$, a homomorphism h from Q_2 to Q_1 is *injective* (or *one-to-one*) if h is injective on atoms, i.e. the multiset of atoms $h(\phi_2(\mathbf{u}_2, \mathbf{v}_2))$ is contained in the multiset of atoms $\phi_1(\mathbf{u}_1, \mathbf{v}_1)$. We write $Q_2 \hookrightarrow Q_1$ iff there exists an injective homomorphism from Q_2 to Q_1. Similar to the case of $\mathbf{S}_{\text{hcov}}$, the following proposition shows that the class $\mathbf{S}_{\text{in}}$ is precisely the class of semirings for which the existence of an injective homomorphism is a sufficient condition for $\mathcal{K}$-containment of CQs.

PROPOSITION 4.12. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{S}_{\text{in}}$;*
- *$Q_2 \hookrightarrow Q_1$ implies $Q_1 \subseteq_{\mathcal{K}} Q_2$, for all CQs Q_1, Q_2.*

PROOF. We need to show that given a semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$,

(1) if $\mathcal{K} \in \mathbf{S}_{\text{in}}$ and Q_1, Q_2 are CQs such that $Q_2 \hookrightarrow Q_1$, then $Q_1 \subseteq_{\mathcal{K}} Q_2$;
(2) if $\mathcal{K} \notin \mathbf{S}_{\text{in}}$, then there exist CQs Q_1 and Q_2 such that $Q_2 \hookrightarrow Q_1$, but $Q_1 \not\subseteq_{\mathcal{K}} Q_2$.

For **Part 1**, we assume that $Q_1(\mathbf{u}) = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and $Q_2(\mathbf{q}) = \exists \mathbf{w}\ S_1(\mathbf{q}_1, \mathbf{w}_1), \ldots, S_m(\mathbf{q}_m, \mathbf{w}_m)$, where $\mathbf{u}$ is the tuple of free variables of Q_1 and $\mathbf{q}$ is the tuple of free variables of Q_2, each $\mathbf{u}_i$ and $\mathbf{q}_j$ consist of variables from $\mathbf{u}$ and $\mathbf{q}$, respectively, and each $\mathbf{v}_i$ and $\mathbf{w}_j$ consist of variables from $\mathbf{v}$ and $\mathbf{w}$, respectively. We need to show that for an arbitrary $\mathcal{K}$-instance I and a tuple $\mathbf{t}$, we have

$$Q_1^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \prod_{1 \leq i \leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \preceq_{\mathcal{K}} \sum_{g \in \mathcal{V}(Q_2, \mathbf{t})} \prod_{1 \leq j \leq m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) = Q_2^I(\mathbf{t}). \quad (12)$$

Since $Q_2 \hookrightarrow Q_1$, let h be this injective homomorphism. Then we have that the multiset $h(S_1(\mathbf{q}_1, \mathbf{w}_1)), \ldots, h(S_m(\mathbf{q}_m, \mathbf{w}_m))$ is contained in the multiset of atoms of Q_1. Without loss of generality, we assume that this multiset corresponds to the first m atoms of Q_1, i.e. for each $1 \leq i \leq m$, we have that $h(S_i(\mathbf{q}_i, \mathbf{w}_i)) = R_i(\mathbf{u}_i, \mathbf{v}_i)$.

Let $\mathcal{V}(Q_1, \mathbf{t}) = \mathcal{V}_1 \cup \ldots \cup \mathcal{V}_k \cup \ldots$ be a (disjoint) partitioning of the set of mappings $\mathcal{V}(Q_1, \mathbf{t})$, such that f and f' are in the same $\mathcal{V}_k$ iff $f(\mathbf{u}_i, \mathbf{v}_i) = f'(\mathbf{u}_i, \mathbf{v}_i)$ for each $1 \leq i \leq m$. Similarly to **Part 1** of the proof of Thm. 3.7 we have

$$\begin{aligned} Q_1^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \prod_{1 \leq i \leq m} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \prod_{m < i \leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) = \\ \sum_{k \geq 1} \sum_{f \in \mathcal{V}_k} \prod_{1 \leq i \leq m} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \prod_{m < i \leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) = \\ \sum_{k \geq 1} \prod_{1 \leq i \leq m} R_i^I(f_k(\mathbf{u}_i, \mathbf{v}_i)) \left(\sum_{f \in \mathcal{V}_k} \prod_{m < i \leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \right), \end{aligned}$$

where f_k is just an arbitrary representative from $\mathcal{V}_k$. Since for each $1 \leq i \leq m$, we have that $h(S_i(\mathbf{q}_i, \mathbf{w}_i)) = R_i(\mathbf{u}_i, \mathbf{v}_i)$, we can instead write

$$Q_1^I(\mathbf{t}) = \sum_{k\geq 1} \prod_{1\leq j\leq m} S_j^I(f_k \circ h(\mathbf{q}_j, \mathbf{w}_j)) \left(\sum_{f\in\mathcal{V}_k} \prod_{m<i\leq n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \right).$$

Th proof now goes along the same lines as the first part of the proof of Thm. 3.7. We only consider instances with finite support, and thus only a finite number of the outer summands are non-zero. Hence, we can apply **Part 1** of Lem. 3.6 and get

$$Q_1^I(\mathbf{t}) \preceq_{\mathcal{K}} \sum_{k\geq 1} \prod_{1\leq j\leq m} S_j^I(f_k \circ h(\mathbf{q}_j, \mathbf{w}_j)).$$

Since for all $k \geq 1$ we have $f_k \circ h \in \mathcal{V}(Q_2, t)$, the desired inequality (12) follows from positivity of $\mathcal{K}$.

For **Part 2** we just reuse the case *(c)* of **Part 2** of the proof of Thm. 3.7. $\square$

Unfortunately, as shown in the following example, $Q_2 \hookrightarrow Q_1$ is just a sufficient, but not always necessary condition for CQ $\mathcal{K}$-containment for a semiring $\mathcal{K}$ from $\mathbf{S}_{\text{in}} \backslash \mathbf{C}_{\text{hom}}$.

Example 4.13. Consider the conjunctive queries

$$Q_1 = \exists u, v, w\ R(u, v), R(u, w), \quad Q_2 = \exists u, v\ R(u, v), R(u, v).$$

We will see in Sec. 6 that Q_1 is $\mathcal{T}^+$-contained in Q_2. However, there is no injective homomorphism from Q_2 to Q_1.

Next we exploit the connection between CQ $\mathcal{K}$-containment and comparison of polynomials from $\mathbb{N}^{cq}[X]$ to define precisely the class of semirings for which an injective homomorphism is a corresponding necessary condition.

Definition 4.14 (*Class* $\mathbf{N}_{\text{in}}$ *of semirings*). Denote by $\mathbf{N}_{\text{in}}$ the class of semirings $\mathcal{K}$ for which for every polynomial P from $\mathbb{N}^{cq}[X]$ and any set of variables $x_1, \ldots, x_n$, the inequality

$$x_1 \times \ldots \times x_n \preceq_{\mathcal{K}} \mathsf{P}$$

implies that there exists a subset $x_{i_1}, \ldots, x_{i_m}$ of the variables $x_1, \ldots, x_n$ such that P contains the monomial $x_{i_1} \times \ldots \times x_{i_m}$.

PROPOSITION 4.15. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{N}_{\text{in}}$;*
- *$Q_1 \subseteq_{\mathcal{K}} Q_2$ implies $Q_2 \hookrightarrow Q_1$, for all CQs Q_1, Q_2.*

PROOF. We need to show, that given a semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1} \rangle$,

(1) if $\mathcal{K} \in \mathbf{N}_{\text{in}}$ and Q_1, Q_2 are CQs such that $Q_1 \subseteq_{\mathcal{K}} Q_2$, then $Q_2 \hookrightarrow Q_1$;
(2) if $\mathcal{K} \notin \mathbf{N}_{\text{in}}$, then there are CQs Q_1 and Q_2 such that $Q_1 \subseteq_{\mathcal{K}} Q_2$, but there is no injective homomorphism from Q_2 to Q_1.

The proof for this proposition is very similar to the proof of Prop. 4.8.

For **Part 1**, we assume again that $Q_1(\mathbf{u}) = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and $Q_2(\mathbf{q}) = \exists \mathbf{w}\ S_1(\mathbf{q}_1, \mathbf{w}_1), \ldots, S_m(\mathbf{q}_m, \mathbf{w}_m)$, where $\mathbf{u}$ is the tuple of free variables of Q_1 and $\mathbf{q}$ is the tuple of free variables of Q_2, each $\mathbf{u}_i$ and $\mathbf{q}_j$ consist of variables from $\mathbf{u}$ and $\mathbf{q}$, respectively, and each $\mathbf{v}_i$ and $\mathbf{w}_j$ consist of variables from $\mathbf{v}$ and $\mathbf{w}$, respectively. Let $X = \{x_1, \ldots, x_n\}$ be a set of variables. Consider a canonical $\mathcal{N}[X]$-instance $[\![Q_1]\!]$ of Q_1 (see the definition in Sec. 3.3). Denote $\mathsf{P}_1 = Q_1^{[\![Q_1]\!]}(\mathbf{u})$ and $\mathsf{P}_2 = Q_2^{[\![Q_1]\!]}(\mathbf{u})$.

Similar to the proof of Prop. 4.8, one can conclude that P_1 contains the monomial $x_1 \times \ldots \times x_n$, P_2 is CQ-admissible, and from the fact that $Q_1 \subseteq_{\mathcal{K}} Q_2$ it must be the case that $\mathsf{P}_1 \preceq_{\mathcal{K}} \mathsf{P}_2$. From positivity of the semiring, we therefore have that $x_1 \times \ldots \times x_n \preceq_{\mathcal{K}} \mathsf{P}_2$. Since $\mathcal{K}$ belongs to $\mathbf{N}_{\text{in}}$, from the definition of $\mathbf{N}_{\text{in}}$ we have that P_2 contains a monomial $x_{i_1} \times \ldots \times x_{i_h}$ for some $1 \leq i_1 < \ldots < i_h \leq n$. Since $[\![Q_1]\!]$ is an $\mathcal{N}[X]$-instance, there exist a mapping $f \in \mathcal{V}(Q_2, \mathbf{u})$, such that $f(\mathbf{q}) = \mathbf{u}$ and for every syntactically distinct atom $R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ from Q_1 (for which $R_\ell^{[\![Q_1]\!]}(\mathbf{u}_\ell, \mathbf{v}_\ell) = x_{\ell_1} + \ldots + x_{\ell_k}$) its preimage by f has N elements, where N is the size of the set $\{\ell_1, \ldots, \ell_k\} \cap \{i_1, \ldots, i_h\}$. Since $R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ has k duplicates in Q_1 and $N \leq k$ we can construct an injective function from Q_2 to Q_1 which maps every atom $S_j(\mathbf{q}_j, \mathbf{w}_j)$ of Q_2 to an atom of the form $S_j(f(\mathbf{q}_j, \mathbf{w}_j))$. Since $f(\mathbf{q}) = \mathbf{u}$, this function is our desired injective homomorphism.

For **Part 2** we also follow a path similar to the proof of Prop. 4.8. We construct Q_1 and Q_2 exactly as in that proof: the assumption that $\mathcal{K} \notin \mathbf{N}_{\text{in}}$ yields a polynomial P over variables $X = \{x_1, \ldots, x_n\}$ such that $x_1 \times \cdots \times x_n \preceq_{\mathcal{K}} \mathsf{P}$, yet P does not have a monomial $x_{i_1} \times \ldots \times x_{i_h}$ for any distinct variables $x_{i_1}, \ldots, x_{i_h}$ from X. From the fact that P is CQ-admissible we construct queries Q_1 and Q_2 with free variables $\mathbf{u}$, an instance I and a tuple $\mathbf{u}$ of elements, such that $Q_1 \subseteq_{\mathcal{K}} Q_2$, and

$$Q_1^{[\![Q_1]\!]}(\mathbf{u}) = x_1 \times \ldots \times x_n \text{ and } Q_2^{[\![Q_1]\!]}(\mathbf{u}) = \mathsf{P}. \tag{13}$$

It is then straightforward to conclude from (13) that there is no injective homomorphism from Q_2 to Q_1, since the existence of such would imply that there is a monomial $x_{i_1} \times \ldots \times x_{i_h}$ in P, for some distinct variables $x_{i_1}, \ldots, x_{i_h}$ from X. $\square$

Prop. 4.12 and 4.15 give us decidability of CQ $\mathcal{K}$-CONTAINMENT for all semirings $\mathcal{K}$ from $\mathbf{C}_{\text{in}} = \mathbf{S}_{\text{in}} \cap \mathbf{N}_{\text{in}}$.

THEOREM 4.16. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{C}_{\text{in}}$;*
- *$Q_1 \subseteq_{\mathcal{K}} Q_2$ iff $Q_2 \hookrightarrow Q_1$, for all CQs Q_1 and Q_2.*

By showing that deciding the existence of an injective homomorphism between queries is NP-complete, we can state the same about $\mathcal{K}$-containment of CQs for any $\mathcal{K} \in \mathbf{C}_{\text{in}}$.

PROPOSITION 4.17. *If $\mathcal{K} \in \mathbf{C}_{\text{in}}$ then* CQ $\mathcal{K}$-CONTAINMENT *is* NP*-complete.*

PROOF. By Thm. 4.16, it is enough to consider the following decision problem.

CQ-INJ:	
Input:	CQs Q_1, Q_2.
Question:	Does $Q_2 \hookrightarrow Q_1$ hold?

It is clear that CQ-INJ is in NP, since a homomorphism forms a certificate of membership, and can be checked in polynomial time. To show NP-hardness, we many-one reduce CLIQUE to CQ-INJ by encoding the input graph and a clique of appropriate size as a pair of queries. We also add dummy edges (on fresh variables), one for each edge in the input graph that is not part of the subgraph induced by the clique vertices.[7]

Formally, suppose we are given an instance of CLIQUE. This consists of an input graph G with vertices $V(G)$ and edges $E(G)$ (containing no self-loops or repeated

[7]Note that the addition of dummy edges is a technical device that will allow us to deduce a stronger Thm. 4.28.

edges), and a desired clique size k. Our aim is to produce an instance $(Q_1^{G,k}, Q_2^{G,k})$ of CQ-INJ, such that $Q_2^{G,k} \hookrightarrow Q_1^{G,k}$ iff G contains a k-clique.

We may assume that the vertices of G are enumerated as $\{w_1, w_2, \ldots, w_n\}$, and that G contains m edges $\{e_1, e'_1\}, \{e_2, e'_2\}, \ldots, \{e_m, e'_m\}$, each a subset of size 2 of $V(G)$.

If $m < k(k-1)/2$, then the graph does not contain enough edges to contain the desired clique, so we output a hardcoded pair of CQs $Q_1^{G,k}$ and $Q_2^{G,k}$ such that $Q_2^{G,k} \not\hookrightarrow Q_1^{G,k}$; it suffices to let $Q_1^{G,k} = \exists v\ R(v)$ and $Q_2^{G,k} = \exists v\ S(v)$ over a schema with two unary symbols R and S. If $m = k(k-1)/2$ then G must be a clique, so we again output a hardcoded pair of CQs $Q_1^{G,k}$ and $Q_2^{G,k}$ such that $Q_2^{G,k} \hookrightarrow Q_1^{G,k}$, and it suffices to let $Q_1^{G,k} = Q_2^{G,k} = \exists v\ R(v)$. So now assume that $m > k(k-1)/2$, and let $M = m - k(k-1)/2$.

Consider a schema with a single binary relation R.

Let $Q_1^{G,k}$ be the CQ with no free variables, with n distinct existential variables $w_1, w_2, \ldots, w_n$, and with an atom $R(w_i, w_j)$ whenever $\{w_i, w_j\}$ is an edge of G. Note that $Q_1^{G,k}$ has $2m$ atoms.

Also, let $Q_2^{G,k}$ be a CQ with no free variables, and with $k + 2M$ distinct existential variables $\mathbf{v} = v_1, v_2, \ldots, v_k, w'_1, w'_2, \ldots, w'_M, w''_1, w''_2, \ldots, w''_M$, defined as

$$\begin{aligned} Q_2^{G,k} = \exists \mathbf{v}\ & R(v_1, v_2), R(v_1, v_3), \ldots, R(v_1, v_k), R(v_2, v_1), R(v_2, v_3), \ldots, R(v_k, v_{k-1}), \\ & R(w'_1, w''_1), R(w''_1, w'_1), R(w'_2, w''_2), R(w''_2, w'_2), \ldots, R(w'_M, w''_M), R(w''_M, w'_M), \end{aligned}$$

which has $2M + k(k-1) = 2m$ atoms.

The queries $Q_1^{G,k}$ and $Q_2^{G,k}$ can be computed using logarithmic space by considering each edge of G separately. To complete the logspace many-one reduction, we now show that there exists a clique with at least k vertices in G iff $Q_2^{G,k} \hookrightarrow Q_1^{G,k}$.

For the forward implication, suppose there is a clique with k vertices in G, consisting of the distinct vertices $w_{i_1}, \ldots, w_{i_k}$ from G. Enumerate the edges of G that are not in the complete subgraph induced by the clique, $E(G) \setminus \{\{w_{i_j}, w_{i_l}\} \mid 1 \le j < l \le k\}$, as $\{\{e_{c_1}, e'_{c_1}\}, \{e_{c_2}, e'_{c_2}\}, \ldots, \{e_{c_M}, e'_{c_M}\}\}$. Here $\{c_1, c_2, \ldots, c_M\}$ forms a subset of $\{1, 2, \ldots, m\}$ of size M.

Define the map h as

$$\begin{aligned} h(v_j) &= w_{i_j}, && \text{for each } j = 1, 2, \ldots, k, \text{ and} \\ \left.\begin{aligned} h(w'_j) &= e_{c_j}, \\ h(w''_j) &= e'_{c_j} \end{aligned}\right\} & && \text{for each } j = 1, 2, \ldots, M. \end{aligned}$$

This defines an injective homomorphism from $Q_2^{G,k}$ to $Q_1^{G,k}$.

For the reverse implication, suppose $Q_2^{G,k} \hookrightarrow Q_1^{G,k}$. Then there is an injective homomorphism $h\colon \{v_1, v_2, \ldots, v_k, w'_1, w'_2, \ldots, w'_m, w''_1, w''_2, \ldots, w''_m\} \to \{w_1, \ldots, w_n\}$. We have to show that $\{h(v_i), h(v_j)\}$ is an edge of G for every distinct pair v_i and v_j. Without loss of generality, suppose $1 \le i < j \le k$. By construction, $R(v_i, v_j)$ and $R(v_j, v_i)$ are both atoms of $Q_2^{G,k}$. Since h is a homomorphism, $R(h(v_i), h(v_j))$ and $R(h(v_j), h(v_i))$ are then both atoms of $Q_1^{G,k}$. Hence $\{h(v_i), h(v_j)\}$ is an edge of G. □

Notwithstanding this result, there are interesting semirings (including the tropical semiring $\mathcal{T}^+$), which lie in $\mathbf{S}_{\text{in}}$, but neither in $\mathbf{C}_{\text{hom}}$ nor in $\mathbf{C}_{\text{in}}$. In Sec. 6 we will see how to obtain decidability for some semirings in $\mathbf{S}_{\text{in}}$, but at the cost of higher complexity.

4.3. Containment by surjective homomorphism

Looking back to the bag semantics semiring $\mathcal{N}$, we know that it lies in the class $\mathbf{N}_{\text{hcov}}$ for which homomorphic covering is necessary, but it does not lie in $\mathbf{C}_{\text{hcov}}$. However, there does exist a well-known sufficient condition for $\mathcal{N}$-containment. This condition is the existence of a surjective homomorphism ([Chaudhuri and Vardi 1993; Ioannidis and Ramakrishnan 1995]): given CQs $Q_1 = \exists \mathbf{v}_1\ \phi_1(\mathbf{u}_1, \mathbf{v}_1)$ and $Q_2 = \exists \mathbf{v}_2\ \phi_2(\mathbf{u}_2, \mathbf{v}_2)$ a homomorphism h from Q_2 to Q_1 is *surjective* (or *onto*) if h is a surjection on atoms, i.e. the multiset of atoms $\phi_1(\mathbf{u}_1, \mathbf{v}_1)$ is contained in the multiset of atoms $h(\phi_2(\mathbf{u}_2, \mathbf{v}_2))$. We write $Q_2 \twoheadrightarrow Q_1$ iff there exists a surjective homomorphism from Q_2 to Q_1.

It is therefore natural to ask for which semirings $Q_2 \twoheadrightarrow Q_1$ is sufficient for $\mathcal{K}$-containment of CQs, and for which this is necessary. Besides $\mathcal{N}$, this condition is sufficient for a larger class of semirings denoted *type* B *systems* [Ioannidis and Ramakrishnan 1995]. From [Green 2011] it is known that $Q_2 \twoheadrightarrow Q_1$ is equivalent to (i.e. both necessary and sufficient for) $\mathtt{Why}[X]$- and $\mathtt{Trio}[X]$-containment of CQs, where $\mathtt{Why}[X]$ is a semiring capturing *why provenance* of [Buneman et al. 2001], and $\mathtt{Trio}[X]$ is a semiring for the provenance model used in the Trio project [Das Sarma et al. 2008]. However, the exact axiomatic bounds for these classes of semirings were not previously known.

As usual, we start by axiomatizing semirings which have $Q_2 \twoheadrightarrow Q_1$ as a sufficient condition.

Definition 4.18 (*Class* $\mathbf{S}_{\text{sur}}$ *of semirings*). Denote by $\mathbf{S}_{\text{sur}}$ the class of semirings that satisfy the axiom:

1. (*$\otimes$-semi-idempotence*) $x \times y \preceq_{\mathcal{K}} x \times x \times y$.

This class can be obtained by relaxing the $\otimes$-idempotence axiom of $\mathbf{S}_{\text{hcov}}$, but only partially, i.e. $\mathbf{S}_{\text{hcov}} \subset \mathbf{S}_{\text{sur}}$. Other than the semirings already mentioned as belonging to $\mathbf{S}_{\text{hcov}}$, it contains the semiring $\mathcal{T}^- = \langle \mathbb{N}_0 \cup \{-\infty\}, \max, +, -\infty, 0\rangle$ known as the *max-plus* (or *schedule*) *algebra* (with its natural order). As desired, the class $\mathbf{S}_{\text{sur}}$ corresponds to all the semirings for which the existence of a surjective homomorphism is a sufficient condition for $\mathcal{K}$-containment of CQs.

PROPOSITION 4.19. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{S}_{\text{sur}}$;*
- *$Q_2 \twoheadrightarrow Q_1$ implies $Q_1 \subseteq_{\mathcal{K}} Q_2$, for all CQs Q_1, Q_2.*

PROOF. We need to show, that given a semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$,

(1) if $\mathcal{K} \in \mathbf{S}_{\text{sur}}$ and Q_1, Q_2 are CQs such that $Q_2 \twoheadrightarrow Q_1$, then $Q_1 \subseteq_{\mathcal{K}} Q_2$;
(2) if $\mathcal{K} \notin \mathbf{S}_{\text{sur}}$, then there exist CQs Q_1 and Q_2 such that $Q_2 \twoheadrightarrow Q_1$, but $Q_1 \not\subseteq_{\mathcal{K}} Q_2$.

For **Part 1**, we assume that $Q_1(\mathbf{u}) = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \dots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and $Q_2(\mathbf{q}) = \exists \mathbf{w}\ S_1(\mathbf{q}_1, \mathbf{w}_1), \dots, S_m(\mathbf{q}_m, \mathbf{w}_m)$, where $\mathbf{u}$ is the tuple of free variables of Q_1 and $\mathbf{q}$ is the tuple of free variables of Q_2, each $\mathbf{u}_i$ and $\mathbf{q}_j$ consist of variables from $\mathbf{u}$ and $\mathbf{q}$, respectively, and each $\mathbf{v}_i$ and $\mathbf{w}_j$ consist of variables from $\mathbf{v}$ and $\mathbf{w}$, respectively. We need to show that for an arbitrary $\mathcal{K}$-instance I and a tuple $\mathbf{t}$, it holds that

$$Q_1^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \prod_{1 \le i \le n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)) \preceq_{\mathcal{K}} \sum_{g \in \mathcal{V}(Q_2, \mathbf{t})} \prod_{1 \le j \le m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) = Q_2^I(\mathbf{t}). \quad (14)$$

There exists a surjective homomorphism h from Q_2 to Q_1, and a set of atoms in Q_2 which h maps to $R_1(\mathbf{u}_1, \mathbf{v}_1), \dots, R_n(\mathbf{u}_n, \mathbf{v}_n)$. Without loss of generality, we can assume that this set of atoms is $S_1(\mathbf{q}_1, \mathbf{w}_1), \dots, S_n(\mathbf{q}_n, \mathbf{w}_n)$. Hence we have that

$$Q_1^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q_1, \mathbf{t})} \prod_{1 \le i \le n} S_i^I(f \circ h(\mathbf{q}_i, \mathbf{w}_i)).$$

Since h is a homomorphism, from semiring positivity and $\otimes$-semi-idempotence we conclude that

$$Q_1^I(\mathbf{t}) \preceq_{\mathcal{K}} \sum_{f \in \mathcal{V}(Q_1,\mathbf{t})} \prod_{1 \le j \le m} S_j^I(f \circ h(\mathbf{q}_j, \mathbf{w}_j)).$$

Let $\mathcal{V}^h = \{f \circ h \mid f \in \mathcal{V}(Q_1, \mathbf{t})\}$. Note that the cardinality of $\mathcal{V}^h$ and $\mathcal{V}(Q_1, \mathbf{t})$ is the same, i.e. for every pair of functions f, f' in $\mathcal{V}(Q_1, \mathbf{t})$, $f \circ h$ and $f' \circ h$ are actually different mappings (this fact was first shown by Ioannidis and Ramakrishnan [1995]). We thus obtain

$$Q_1^I(\mathbf{t}) \preceq_{\mathcal{K}} \sum_{f \in \mathcal{V}(Q_1,\mathbf{t})} \prod_{1 \le j \le m} S_j^I(f \circ h(\mathbf{q}_j, \mathbf{w}_j)) = \sum_{g \in \mathcal{V}^h} \prod_{1 \le j \le m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)).$$

However, $\mathcal{V}^h \subseteq \mathcal{V}(Q_2, \mathbf{t})$, and hence from positivity of the semiring we conclude

$$Q_1^I(\mathbf{t}) \preceq_{\mathcal{K}} \sum_{g \in \mathcal{V}^h} \prod_{1 \le j \le m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) \preceq_{\mathcal{K}} \sum_{g \in \mathcal{V}(Q_2,\mathbf{t})} \prod_{1 \le j \le m} S_j^I(g(\mathbf{q}_j, \mathbf{w}_j)) = Q_2^I(\mathbf{t}),$$

i.e. the desired inequality (14) holds.

For **Part 2** we assume that $\mathcal{K} \notin \mathbf{S}_{\text{sur}}$, and need to show that there are CQs Q_1 and Q_2 such that there is a surjective homomorphism from Q_2 to Q_1, but $Q_1 \not\subseteq_{\mathcal{K}} Q_2$.

Then semiring $\mathcal{K}$ fails to satisfy $x \times y \preceq_{\mathcal{K}} x \times x \times y$, witnessed by some $a, b \in K$ such that $a \otimes b \not\preceq_{\mathcal{K}} a \otimes a \otimes b$. Consider a schema with unary relations R and S and queries $Q_1 = \exists v\ R(v), S(v)$ and $Q_2 = \exists u, v\ R(u), R(v), S(v)$. Clearly $Q_2 \twoheadrightarrow Q_1$. However, for the $\mathcal{K}$-instance I such that $R^I(c) = a$, $S^I(c) = b$ for some constant c and $R^I(c') = S^I(c') = \mathbb{0}$ for all $c' \neq c$, we have that $Q_1^I() \not\preceq_{\mathcal{K}} Q_2^I()$, which means that $Q_1 \not\subseteq_{\mathcal{K}} Q_2$. □

As we saw for the bag semantics semiring $\mathcal{N}$, the existence of a surjective homomorphism is not necessary for $\mathcal{N}$-containment, but homomorphic covering is. The same can be said about the max-plus algebra $\mathcal{T}^-$, but not for other semirings, such as $\mathtt{Why}[X]$ or $\mathtt{Trio}[X]$. Hence, again we need to axiomatize the class of semirings for which $Q_2 \twoheadrightarrow Q_1$ is necessary for $\mathcal{K}$-containment of CQs. For this we exploit once more the notion of CQ-admissible polynomials.

Definition 4.20 (*Class* $\mathbf{N}_{\text{sur}}$ *of semirings*). Denote by $\mathbf{N}_{\text{sur}}$ the class of semirings $\mathcal{K}$ for which for every polynomial P from $\mathbb{N}^{cq}[X]$ and any set of variables $x_1, \ldots, x_n$, the inequality

$$x_1 \times \ldots \times x_n \preceq_{\mathcal{K}} \mathsf{P}$$

implies that there exist exponents $m_1, \ldots, m_n \geq 1$ such that P contains the monomial $x_1^{m_1} \times \ldots \times x_n^{m_n}$.

PROPOSITION 4.21. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{N}_{\text{sur}}$;*
- *$Q_1 \subseteq_{\mathcal{K}} Q_2$ implies $Q_2 \twoheadrightarrow Q_1$, for all CQs Q_1, Q_2.*

PROOF. We need to show, that given a semiring $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1} \rangle$,

(1) if $\mathcal{K} \in \mathbf{N}_{\text{sur}}$ and Q_1, Q_2 are CQs such that $Q_1 \subseteq_{\mathcal{K}} Q_2$, then $Q_2 \twoheadrightarrow Q_1$;
(2) if $\mathcal{K} \notin \mathbf{N}_{\text{sur}}$, then there are CQs Q_1 and Q_2 such that $Q_1 \subseteq_{\mathcal{K}} Q_2$, but there is no surjective homomorphism from Q_2 to Q_1.

The proof for this proposition is again very similar to the proof of Prop. 4.8 (and the proof of Prop. 4.15).

For **Part 1**, we assume again that $Q_1(\mathbf{u}) = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \dots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and $Q_2(\mathbf{q}) = \exists \mathbf{w}\ S_1(\mathbf{q}_1, \mathbf{w}_1), \dots, S_m(\mathbf{q}_m, \mathbf{w}_m)$, where $\mathbf{u}$ is the tuple of free variables of Q_1 and $\mathbf{q}$ is the tuple of free variables of Q_2, each $\mathbf{u}_i$ and $\mathbf{q}_j$ consist of variables from $\mathbf{u}$ and $\mathbf{q}$, respectively, and each $\mathbf{v}_i$ and $\mathbf{w}_j$ consist of variables from $\mathbf{v}$ and $\mathbf{w}$, respectively. Let $X = \{x_1, \dots, x_n\}$ be a set of variables. Consider a canonical $\mathcal{N}[X]$-instance $[\![Q_1]\!]$ of Q_1 Denote $\mathsf{P}_1 = Q_1^{[\![Q_1]\!]}(\mathbf{u})$ and $\mathsf{P}_2 = Q_2^{[\![Q_1]\!]}(\mathbf{u})$.

Similar to the proof of Prop. 4.8, one can conclude that P_1 contains the monomial $x_1 \times \dots \times x_n$, and from the fact that $Q_1 \subseteq_{\mathcal{K}} Q_2$ it must be the case that $\mathsf{P}_1 \preceq_{\mathcal{K}} \mathsf{P}_2$. From positivity of the semiring, we therefore have that $x_1 \times \dots \times x_n \preceq_{\mathcal{K}} \mathsf{P}_2$. Since $\mathcal{K}$ belongs to $\mathbf{N}_{\text{sur}}$, from the definition of $\mathbf{N}_{\text{sur}}$ we have that P_2 contains a monomial $x_1^{m_1} \times \dots \times x_n^{m_n}$ for some $m_1, \dots, m_n \geq 1$.

Since $[\![Q_1]\!]$ is an $\mathcal{N}[X]$-instance, there exists a mapping $f \in \mathcal{V}(Q_2, \mathbf{u})$, such that $f(\mathbf{q}) = \mathbf{u}$ and for every syntactically distinct atom $R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ from Q_1 (for which $R_\ell^{[\![Q_1]\!]}(\mathbf{u}_\ell, \mathbf{v}_\ell) = x_{\ell_1} + \dots + x_{\ell_k}$, where k is the number of occurrences of the atom $R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ in Q_1) its preimage by f has $m_{\ell_1} + \dots + m_{\ell_k}$ elements. Since $R_\ell(\mathbf{u}_\ell, \mathbf{v}_\ell)$ has k duplicates in Q_1 and $k \leq m_{\ell_1} + \dots + m_{\ell_k}$ we can construct a surjective function from Q_2 to Q_1 which maps every atom $S_j(\mathbf{q}_j, \mathbf{w}_j)$ of Q_2 to an atom of the form $S_j(f(\mathbf{q}_j, \mathbf{w}_j))$. Since $f(\mathbf{q}) = \mathbf{u}$, this function is our desired surjective homomorphism.

For **Part 2** we also follow a path similar to the proof of Prop. 4.8. We construct Q_1 and Q_2 exactly as in that proof: the assumption that $\mathcal{K} \notin \mathbf{N}_{\text{sur}}$ yields a polynomial P over variables $X = \{x_1, \dots, x_n\}$ such that $x_1 \times \dots \times x_n \preceq_{\mathcal{K}} \mathsf{P}$, yet P does not have a monomial $x_1^{m_1} \times \dots \times x_n^{m_n}$ for any exponents $m_1, \dots, m_n \geq 1$. From the fact that P is CQ-admissible we construct queries Q_1 and Q_2 with free variables $\mathbf{u}$, an instance I and a tuple $\mathbf{u}$ of elements, such that $Q_1 \subseteq_{\mathcal{K}} Q_2$, and

$$Q_1^{[\![Q_1]\!]}(\mathbf{u}) = x_1 \times \dots \times x_n \text{ and } Q_2^{[\![Q_1]\!]}(\mathbf{u}) = \mathsf{P}. \tag{15}$$

It is then straightforward to conclude from (15) that there is no surjective homomorphism from Q_2 to Q_1, since the existence of such would imply that there is a monomial $x_1^{m_1} \times \dots \times x_n^{m_n}$ in P, for some $m_1, \dots, m_n \geq 1$. □

For those semirings $\mathcal{K}$ that do belong to $\mathbf{C}_{\text{sur}} = \mathbf{S}_{\text{sur}} \cap \mathbf{N}_{\text{sur}}$ (like $\texttt{Why}[X]$ and $\texttt{Trio}[X]$), we have once again a decision procedure for $\mathcal{K}$-containment of CQs. This is summarized by the following theorem.

THEOREM 4.22. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{C}_{\text{sur}}$;*
- *$Q_1 \subseteq_{\mathcal{K}} Q_2$ iff $Q_2 \twoheadrightarrow Q_1$, for all CQs Q_1 and Q_2.*

Checking a surjective homomorphism between CQs is NP-complete [Chaudhuri and Vardi 1993]. The complexity of CQ $\mathcal{K}$-CONTAINMENT for $\mathbf{C}_{\text{sur}}$ then follows.

COROLLARY 4.23. *If $\mathcal{K} \in \mathbf{C}_{\text{sur}}$ then* CQ $\mathcal{K}$-CONTAINMENT *is* NP-*complete.*

Note that NP-hardness will also follow from the general result Thm. 4.28.

As mentioned in the introduction, we leave open the problem of finding decision procedures for all semirings that belong to $\mathbf{S}_{\text{sur}}$, but not to $\mathbf{N}_{\text{sur}}$, such as $\mathcal{N}$ or $\mathcal{T}^-$. In Sec. 6 we show that for some of these semirings, such as $\mathcal{T}^-$, the problem of $\mathcal{K}$-containment of CQs can be solved using a different approach, albeit with higher computational complexity.

4.4. Containment by bijective homomorphism

Finally, we deal with the class obtained from $\mathbf{C}_{\text{hom}}$ by relaxing both of its axioms. This is just the class of all ($\preceq$-positive) semirings $\mathbf{S}_{\preceq}$. For this class we use again the notion of bijective homomorphism: Given CQs $Q = \exists \mathbf{v}_1\ \phi_1(\mathbf{u}_1, \mathbf{v}_1)$ and $Q_2 = \exists \mathbf{v}_2\ \phi_2(\mathbf{u}_2, \mathbf{v}_2)$, we say that a homomorphism h from Q_2 to Q_1 is *bijective* (or *exact*) if it is a bijection on atoms, i.e. the multiset of atoms $h(\phi_2(\mathbf{u}_2, \mathbf{v}_2))$ is the same as the multiset of atoms $\phi_1(\mathbf{u}_1, \mathbf{v}_1)$. We write $Q_2 \hookrightarrow\!\!\!\!\rightarrow Q_1$ if there exists a bijective homomorphism from Q_2 to Q_1.

Note that a bijective homomorphism is not necessarily an isomorphism, since it can identify variables. However, a bijective homomorphism can exist between two CQs only if they contain the same number of atoms. Also, a homomorphism is bijective iff it is both injective and surjective. We use this fact further in this section, to characterize a class of semirings for which a bijective homomorphism is a necessary condition for $\mathcal{K}$-containment of CQs.

From the results of Green [2011] and Prop. 3.2 we can immediately obtain that the existence of a bijective homomorphism is sufficient for CQ $\mathcal{K}$-containment for an arbitrary $\preceq$-positive nontrivial semiring $\mathcal{K}$.

PROPOSITION 4.24. *For any semiring $\mathcal{K}$ from $\mathbf{S}_{\preceq}$, if $Q_2 \hookrightarrow\!\!\!\!\rightarrow Q_1$ then $Q_1 \subseteq_{\mathcal{K}} Q_2$.*

PROOF. Indeed, as shown by Green [2011], the condition $Q_2 \hookrightarrow\!\!\!\!\rightarrow Q_1$ is both sufficient and necessary for $\mathcal{N}[X]$-containment of CQs over the provenance polynomials semiring $\mathcal{N}[X]$. Since $\mathcal{N}[X]$ is universal for $\mathbf{S}_{\preceq}$ by Prop. 3.2, we can conclude that this condition is sufficient for $\mathcal{K}$-containment of CQs for any $\mathcal{K}$ from $\mathbf{S}_{\preceq}$. □

From [Green 2011] we also know that existence of a bijective homomorphism is necessary for $\mathcal{B}[X]$-containment of CQs, where $\mathcal{B}[X] = \langle \mathbb{B}[X], +, \times, 0, 1\rangle$ is the semiring of *boolean provenance polynomials*, i.e. polynomials over X with boolean coefficients from $\mathbb{B} = \{\mathbf{false}, \mathbf{true}\}$. This means that $\mathcal{B}[X]$ behaves like $\mathcal{N}[X]$ w.r.t. $\mathcal{K}$-containment of CQs. As we have seen in previous sections, this is not the case for all semirings. Also, one can easily show that the existence of a bijective homomorphism is not necessary for bag semantics $\mathcal{N}$, or even for the semiring $\mathcal{R}^+$ of non-negative reals with the usual operations and order.

Our next aim is to identify all semirings which behave as $\mathcal{N}[X]$. To do so we again exploit the notion of CQ-admissible polynomials.

Definition 4.25 (*Class $\mathbf{C}_{\rm bi}$ of semirings*). Denote by $\mathbf{C}_{\rm bi}$ the class of all semirings $\mathcal{K}$ for which for every polynomial P from $\mathbb{N}^{cq}[X]$ and any set of variables $x_1, \ldots, x_n$, the inequality

$$x_1 \times \ldots \times x_n \preceq_{\mathcal{K}} \mathsf{P}$$

implies that P contains the monomial $x_1 \times \ldots \times x_n$.

THEOREM 4.26. *The following are equivalent:*

- *semiring $\mathcal{K}$ belongs to $\mathbf{C}_{\rm bi}$;*
- *$Q_1 \subseteq_{\mathcal{K}} Q_2$ iff $Q_2 \hookrightarrow\!\!\!\!\rightarrow Q_1$, for all CQs Q_1 and Q_2.*

PROOF. By Prop. 4.24, if $Q_2 \hookrightarrow\!\!\!\!\rightarrow Q_1$ then $Q_1 \subseteq_{\mathcal{K}} Q_2$ for any semiring $\mathcal{K}$. Hence we only prove the remaining direction. Looking at the condition for the classes $\mathbf{N}_{\rm sur}$ and $\mathbf{N}_{\rm in}$, we conclude that $\mathbf{C}_{\rm bi} = \mathbf{N}_{\rm in} \cap \mathbf{N}_{\rm sur}$. From Prop. 4.15 and 4.21 we have that $\mathbf{N}_{\rm in}$ corresponds to all the semirings for which the existence of an injective homomorphism is necessary for CQ $\mathcal{K}$-containment, and $\mathbf{N}_{\rm in}$ – to all the semirings for which the existence of a surjective homomorphism is necessary. Since as noted above, a homomorphism is bijective iff it is both injective and surjective, we have that $\mathbf{C}_{\rm bi}$ consists of all the semirings $\mathcal{K}$ for which $Q_2 \hookrightarrow\!\!\!\!\rightarrow Q_1$ is necessary for $Q_1 \subseteq_{\mathcal{K}} Q_2$. □

In particular, notice that both $\mathcal{B}[X]$ and $\mathcal{N}[X]$ belong to $\mathbf{C}_{\rm bi}$. Thus, this theorem can be seen as a generalization of the results of Green [2011]. There it was also shown that

$\mathcal{N}[X]$-containment of CQs is an NP-complete problem. We can now extend this result to the entire class $\mathbf{C}_{\rm bi}$.

PROPOSITION 4.27. *If* $\mathcal{K} \in \mathbf{C}_{\rm bi}$ *then* CQ $\mathcal{K}$-CONTAINMENT *is* NP-*complete.*

PROOF. The proof is essentially the same as the proof for Prop. 4.17. Note that our construction of $Q_1^{G,k}$ and $Q_2^{G,k}$ forces the homomorphism h from $Q_2^{G,k}$ to $Q_1^{G,k}$ to be not only injective, but bijective as well. The result then follows by Thm. 4.26. □

Combining Prop. 3.4 and 4.24 with the proof of Prop. 4.17 actually leads to the following general result for all $\preceq$-positive semirings. Previously known special cases include [Aho et al. 1979, Thm. 7(1)] and [Green 2011, Cor. 7.4], as well as results implicit in [Chandra and Merlin 1977] and [Grahne et al. 1997].

THEOREM 4.28. CQ $\mathcal{K}$-CONTAINMENT *is* NP-*hard for any semiring* $\mathcal{K}$ *from* $\mathbf{S}_{\preceq}$.

PROOF. In the proof of Prop. 4.17 we provided a logspace reduction by constructing CQs $Q_1^{G,k}$ and $Q_2^{G,k}$ for any input graph G and positive integer k.

The identical construction also defines a logspace many-one reduction from CLIQUE to CQ $\mathcal{K}$-CONTAINMENT. To prove this claim, we show that when G is a graph and k a positive integer, then G contains a k-clique iff $Q_1^{G,k} \subseteq_{\mathcal{K}} Q_2^{G,k}$. (In the case of $m < k(k-1)/2$, note that $Q_1^{G,k} \not\subseteq_{\mathcal{K}} Q_2^{G,k}$ for the hardcoded queries $Q_1^{G,k}$ and $Q_2^{G,k}$, since $\mathcal{K}$ is nontrivial. Moreover, if $m = k(k-1)/2$ then $Q_1^{G,k} \subseteq_{\mathcal{K}} Q_2^{G,k}$ for the hardcoded queries in this case.)

Suppose first that the input graph G contains a k-clique. Hence $Q_2^{G,k} \hookrightarrow\!\!\!\!\rightarrow Q_1^{G,k}$, as in the proof of Prop. 4.17. By Prop. 4.24 it follows that $Q_1^{G,k} \subseteq_{\mathcal{K}} Q_2^{G,k}$.

Now recall that, by Prop. 3.4, the existence of a homomorphism between CQs is necessary for their $\mathcal{K}$-containment for any semiring $\mathcal{K}$ from $\mathbf{S}_{\preceq}$. In other words, if $Q_1^{G,k} \subseteq_{\mathcal{K}} Q_2^{G,k}$ then $Q_2^{G,k} \to Q_1^{G,k}$. Hence G contains a k-clique. By the same method as in the proof of Prop. 4.17 we then have $Q_2^{G,k} \hookrightarrow\!\!\!\!\rightarrow Q_1^{G,k}$. □

Thm. 4.28 applies to all the semirings that we study here. For the special case of the semiring $\mathcal{N}$, Chaudhuri and Vardi [1993, Thm. 4.9] state the stronger result of Π_2^p-hardness for $\mathcal{N}$-containment of CQs.

This completes our study of $\mathcal{K}$-containment of CQs for the classes of semirings obtained from $\mathbf{C}_{\rm hom}$ by relaxing its axioms.

5. CQ-ADMISSIBLE POLYNOMIALS

In Sec. 2 we defined the *evaluation* of a CQ $Q = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ on a $\mathcal{K}$-instance I for a tuple $\mathbf{t}$ as

$$Q^I(\mathbf{t}) = \sum_{f \in \mathcal{V}(Q,\mathbf{t})} \prod_{1 \le i \le n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)).$$

Thus, the evaluation of a CQ on an $\mathcal{N}[X]$-instance with unique variables from the set X as annotations is a polynomial over X. In Def. 4.4 we called such polynomials *CQ-admissible*. We heavily used this notion in the definitions of the classes $\mathbf{N}_{\rm in}$, $\mathbf{N}_{\rm bi}$, and $\mathbf{N}_{\rm sur}$. The goal of this section is to give a constructive algebraic characterization of the set $\mathbb{N}^{cq}[X]$ of all CQ-admissible polynomials. As we mentioned in the introduction, this notion is of independent interest: for instance, it was implicitly used in [Olteanu and Závodný 2012].

From the definition of evaluation we immediately obtain that every CQ-admissible polynomial must be *homogeneous* (i.e. all non-zero monomials have the same degree). For this last statement, let Q be a CQ consisting of k atoms, and I be an $\mathcal{N}[X]$-instance with tuples annotated with variables $X = \{x_1, \ldots, x_n\}$, such that the mappings in the set $\mathcal{V}(Q, \mathbf{t})$ allow us to obtain any possible combination of images of the atoms of Q to non-zero annotated tuples of I. It follows that $Q^I(\mathbf{t})$ corresponds precisely to the expansion of the multinomial expression $(x_1 + \ldots + x_n)^k$, and therefore $Q^I(\mathbf{t}) = (x_1 + \ldots + x_n)^k$. As a justifying example, consider a CQ $Q = \exists u, v\ R(u), R(v)$, and an $\mathcal{N}[X]$-instance I and two elements a and b such that $R^I(a) = x_1$ and $R^I(b) = x_2$, with all other tuples annotated by 0. Then $Q^I() = x_1^2 + x_1x_2 + x_2x_1 + x_2^2 = (x_1 + x_2)^2$. In fact, this property allowed us to formulate the axiom for the class $\mathbf{N}_{\mathrm{hcov}}$ in Def. 4.3 without reference to CQ-admissible polynomials. As some negative examples, note that polynomials such as $2x$ and $x^2 + y$ do not satisfy this property and are not in $\mathbb{N}^{cq}[X]$.

The polynomials x^2, $2xy$ and $x + y$ satisfy the requirements above, and it is not difficult to construct CQs which admit them. Unfortunately, these are not the only requirements: the polynomial $x^2 + xy + y^2$ satisfies them, but can be proved not to be in $\mathbb{N}^{cq}[X]$. In order to present the precise characterization, we need an auxiliary notion: for a set of variables X an *ordered monomial of degree* n (or *o-monomial*) is a string from X^n. For an o-monomial $\overrightarrow{\mathsf{M}}$ we denote by $\overrightarrow{\mathsf{M}}[i]$ the variable appearing in its i-th position.

PROPOSITION 5.1. *A non-zero polynomial* P *is in* $\mathbb{N}^{cq}[X]$ *iff it can be represented in a form*

$$\overrightarrow{\mathsf{P}} = \sum_{1 \leq \ell \leq m} \overrightarrow{\mathsf{M}}_\ell, \text{ such that}$$

1. $\overrightarrow{\mathsf{M}}_\ell$, $1 \leq \ell \leq m$, *are pairwise distinct o-monomials over* X *of the same degree* n *(here concatenation in* $\overrightarrow{\mathsf{M}}_\ell$ *as a string is interpreted as product in* P*), and*

2. if for an o-monomial $\overrightarrow{\mathsf{M}}$ *of degree* n*, and for each* i, j *such that* $1 \leq i \leq n$ *and* $1 \leq j \leq n$*, the representation* $\overrightarrow{\mathsf{P}}$ *contains o-monomials (each of degree* n*)* $\overrightarrow{\mathsf{M}}_1, \ldots, \overrightarrow{\mathsf{M}}_{2k+1}$, $k \geq 0$*, such that*

- $\overrightarrow{\mathsf{M}}_1[i] = \overrightarrow{\mathsf{M}}[i]$, $\overrightarrow{\mathsf{M}}_{2k+1}[j] = \overrightarrow{\mathsf{M}}[j]$, *and*
- $\overrightarrow{\mathsf{M}}_{2\ell-1}[j] = \overrightarrow{\mathsf{M}}_{2\ell}[j]$, $\overrightarrow{\mathsf{M}}_{2\ell}[i] = \overrightarrow{\mathsf{M}}_{2\ell+1}[i]$ *for all* $1 \leq \ell \leq k$,

then $\overrightarrow{\mathsf{M}}$ *is contained in* $\overrightarrow{\mathsf{P}}$.

Before the formal proof of this proposition we give a short explanation why the requirements above do not hold for the polynomial $x^2 + xy + y^2$. Consider the only representation $\overrightarrow{\mathsf{P}} = xx + xy + yy$ of this polynomial. The first requirement clearly holds. However, the second requirement does not. Indeed, there exists the o-monomial $\overrightarrow{\mathsf{M}} = yx$, which is not in $\overrightarrow{\mathsf{P}}$, but for every i and j the representation $\overrightarrow{\mathsf{P}}$ contains required o-monomials: for $i = j = 1$ we have $k = 0$ and $\overrightarrow{\mathsf{M}}_1 = yy$; for $i = 1$ and $j = 2$ we have $k = 1$, $\overrightarrow{\mathsf{M}}_1 = yy$, $\overrightarrow{\mathsf{M}}_2 = xy$, and $\overrightarrow{\mathsf{M}}_3 = xx$; the cases $i = j = 2$ and $i = 2, j = 1$ are just symmetrical to the cases above. Hence, the polynomial $x^2 + xy + y^2$ is indeed not CQ-admissible.

PROOF. To prove the **"if"** direction, it is enough to construct a CQ $Q = \phi(\mathbf{v})$ without free variables over some schema $\mathbb{S}$, and an $\mathcal{N}[X]$-instance I with tuples annotated with unique variables from X such that $Q^I() = \mathsf{P}$. Since P is CQ-admissible, fix a representation $\overrightarrow{\mathsf{P}}$ which satisfies requirements 1 and 2, and denote by n the degree of P.

In our construction the schema and the CQ depend only on the degree n, but not on the exact form of P (the last one is used in the construction of the instance). Particularly, the schema $\mathbb{S}$ consists of a single relation R of arity $n(n-1)+1$. The first attribute of this relation plays no other role in the proof than to allow the existence of two tuples with exactly the same constants in the other $n(n-1)$ attributes. So for every tuple this attribute contains a unique id value and it will not be mentioned further in the proof. The other $n(n-1)$ attributes correspond intuitively to all pairs (i,j) such that $1 \leq i \leq n$, $1 \leq j \leq n$, and $i \neq j$, ordered first by the i-th coordinate and then by the j-th coordinate. (Notice that there are exactly $n(n-1)$ of these pairs). Let $\mathbf{t}$ be a tuple of constants in $\mathbb{D}$ of arity $n(n-1)+1$. We will abuse notation, and speak of the *projection* of $\mathbf{t}$ with respect to a pair (i,j), denoted by $\pi_{(i,j)}\mathbf{t}$, to refer to the element of $\mathbf{t}$ in the position corresponding to the pair (i,j), according to the order given in the intuitive interpretation of R. For example, $\pi_{(1,2)}\mathbf{t}$ corresponds to the first element of $\mathbf{t}$ (since we require $i \neq j$ the first element in our order of pairs is $(1,2)$) and $\pi_{(n,n-2)}\mathbf{t}$ corresponds to the next to last element of $\mathbf{t}$.

The CQ Q contains n atoms $R(\mathbf{v}_1), \dots, R(\mathbf{v}_n)$, where for each $1 \leq m \leq n$ the tuple $\mathbf{v}_m$ contains variables $v_m^{(i,j)}$, for $1 \leq i \leq n$, $1 \leq j \leq n$, $i \neq j$, arranged in the same order as in the intuitive interpretation of R. All variables in the atoms of Q are distinct, except for the following rule: for each $1 \leq i \leq n$ and $1 \leq j \leq n$ such that $i \neq j$ we have

$$v_i^{(i,j)} = v_j^{(j,i)}. \tag{16}$$

Finally, we create the instance I on the base of the representation $\overrightarrow{\mathsf{P}}$ as follows. Let $\{\mathbf{t}_x \mid x \in X\}$ be a set of tuples of constants in $\mathbb{D}$ all of size n such that the constants in these tuples obey the following rule: if $\overrightarrow{\mathsf{M}}$ is an o-monomial in $\overrightarrow{\mathsf{P}}$ then for each $1 \leq i \leq n$ and $1 \leq j \leq n$ such that $i \neq j$ it holds that

$$\pi_{(i,j)}\mathbf{t}_{\overrightarrow{\mathsf{M}}[i]} = \pi_{(j,i)}\mathbf{t}_{\overrightarrow{\mathsf{M}}[j]}; \tag{17}$$

and all the values which are not forced to be equal by this rule, are different. Then, create I by setting $R^I(\mathbf{t}_x) = x$ for each $x \in X$, and annotating all other tuples with 0.

Next we prove that $Q^I() = \mathsf{P}$. It is enough to show that an o-monomial $\overrightarrow{\mathsf{M}}$ of degree n is in $\overrightarrow{\mathsf{P}}$ iff there exists a mapping f in $\mathcal{V}(Q)$ such that

$$f(\mathbf{v}_i) = \mathbf{t}_{\overrightarrow{\mathsf{M}}[i]} \text{ for every } 1 \leq i \leq n. \tag{18}$$

For the "only if" direction of this statement, we will show that the mapping f defined by (18) belongs to $\mathcal{V}(Q)$. Since all the repeated variables in Q are defined by (16), we only need to check that for each $1 \leq i \leq n$ and $1 \leq j \leq n$ such that $i \neq j$ it holds that $f(v_i^{(i,j)}) = f(v_j^{(j,i)})$. This is true by the construction (17), since we have:

$$f(v_i^{(i,j)}) = \pi_{(i,j)}\mathbf{t}_{\overrightarrow{\mathsf{M}}[i]} = \pi_{(j,i)}\mathbf{t}_{\overrightarrow{\mathsf{M}}[j]} = f(v_j^{(j,i)}).$$

For the "if" direction of the statement consider a mapping $f \in \mathcal{V}(Q)$ conforming to (18). We need to show that for the monomial $\overrightarrow{\mathsf{M}}$ and for every i, j the conditions of requirement 2 hold. Consider such a pair (i,j). If $i = j$ then we can take $k = 0$ and $\overrightarrow{\mathsf{M}}_1 = \overrightarrow{\mathsf{M}}$ for which the conditions trivially hold. Let now $i \neq j$. From (16) we have that $f(v_i^{(i,j)}) = f(v_j^{(j,i)})$ which implies that

$$\pi_{(i,j)}\mathbf{t}_{\mathsf{M}[i]} = \pi_{(j,i)}\mathbf{t}_{\mathsf{M}[j]}. \tag{19}$$

By the construction, constants in I may coincide if they belong to the same position in some tuple $\mathbf{t}_x$ (i.e. they are the *same* constant), or if they were enforced by (17). Thus, if (19) holds then there must exist o-monomials $\overrightarrow{\mathsf{M}}_1, \dots, \overrightarrow{\mathsf{M}}_{2k+1}$, $k \geq 0$, in $\overrightarrow{\mathsf{P}}$, such that $\overrightarrow{\mathsf{M}}_1[i] = \overrightarrow{\mathsf{M}}[i]$, $\overrightarrow{\mathsf{M}}_{2k+1}[j] = \overrightarrow{\mathsf{M}}[j]$ and for each $1 \leq \ell \leq k$ the equalities $\overrightarrow{\mathsf{M}}_{2\ell-1}[j] = \overrightarrow{\mathsf{M}}_{2\ell}[j]$

and $\overrightarrow{\mathsf{M}}_{2\ell}[i] = \overrightarrow{\mathsf{M}}_{2\ell+1}[i]$ hold. Since $\overrightarrow{\mathsf{P}}$ satisfies requirement 2 and our choice of i, j was arbitrary, we can conclude that $\overrightarrow{\mathsf{M}}$ is in $\overrightarrow{\mathsf{P}}$.

For the **"only if"** direction of the proposition, consider a CQ $Q = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_n(\mathbf{u}_n, \mathbf{v}_n)$, where $\mathbf{u}$ is the tuple of free variables of Q, an $\mathcal{N}[X]$-instance I, each tuple of which is annotated either with a unique variable from X, or 0, and a tuple $\mathbf{t}$ of size $|\mathbf{u}|$, all over some schema $\mathbb{S}$. We need to show that the polynomial $\mathsf{P} = Q^I(\mathbf{t})$ satisfies requirements 1 and 2.

Without loss of generality we assume that all variables from X are used as annotations in I, and for each $x \in X$ denote by $\mathbf{t}_x$ and R_x the tuple and the relation such that $R_x^I(\mathbf{t}_x) = x$. From the form of Q we have that

$$\mathsf{P} = \sum_{f \in \mathcal{V}(Q,\mathbf{t})} \prod_{1 \le i \le n} R_i^I(f(\mathbf{u}_i, \mathbf{v}_i)).$$

Obviously we are only interested in those mappings from $\mathcal{V}(Q, \mathbf{t})$ such that the corresponding product in the sum above is non-zero. We can view all these non-zero products as o-monomials, ordered by the conventional order on atoms in the CQ Q. Thus, P can be represented in a form

$$\overrightarrow{\mathsf{P}} = \sum_{1 \le \ell \le m} \overrightarrow{\mathsf{M}}_\ell,$$

where each $\overrightarrow{\mathsf{M}}_\ell$, $1 \le \ell \le m$, is an o-monomial for which there exists a mapping f_ℓ in $\mathcal{V}(Q, \mathbf{t})$ such that for each $1 \le i \le n$ it holds that

$$R_i(f_\ell(\mathbf{u}_i, \mathbf{v}_i)) = R_{\overrightarrow{\mathsf{M}}_\ell[i]}(\mathbf{t}_{\overrightarrow{\mathsf{M}}_\ell[i]}). \tag{20}$$

Next we check the requirements 1 and 2 of the definition of CQ-admissible polynomial for this representation.

For requirement 1, note that the degree of each $\overrightarrow{\mathsf{M}}_\ell$ is equal to the number of atoms in Q which means that all of them have the degree n. Also, all of them are different, since distinct mappings from $\mathcal{V}(Q, \mathbf{t})$ cannot give equal o-monomials. Hence requirement 1 holds for $\overrightarrow{\mathsf{P}}$.

For requirement 2, consider an o-monomial $\overrightarrow{\mathsf{M}}$ of degree n such that the precondition of this requirement holds, i.e. for each $1 \le i \le n$ and $1 \le j \le n$ there exist o-monomials $\overrightarrow{\mathsf{M}}_1, \ldots, \overrightarrow{\mathsf{M}}_{2k+1}$, $k \ge 0$, in $\overrightarrow{\mathsf{P}}$, such that $\overrightarrow{\mathsf{M}}_1[i] = \overrightarrow{\mathsf{M}}[i]$, $\overrightarrow{\mathsf{M}}_{2k+1}[j] = \overrightarrow{\mathsf{M}}[j]$, and for each $1 \le \ell \le k$ it holds that $\overrightarrow{\mathsf{M}}_{2\ell-1}[j] = \overrightarrow{\mathsf{M}}_{2\ell}[j]$ and $\overrightarrow{\mathsf{M}}_{2\ell}[i] = \overrightarrow{\mathsf{M}}_{2\ell+1}[i]$.

We need to prove that $\overrightarrow{\mathsf{M}}$ is also contained in $\overrightarrow{\mathsf{P}}$. According to (20), we need to show that there exists a mapping f in $\mathcal{V}(Q, \mathbf{t})$ such that for each $1 \le i \le n$ it holds that

$$R_i(f(\mathbf{u}_i, \mathbf{v}_i)) = R_{\overrightarrow{\mathsf{M}}[i]}(\mathbf{t}_{\overrightarrow{\mathsf{M}}[i]}).$$

Note that by the precondition of requirement 2, for all $1 \le i \le n$ we have that $R_i = R_{\overrightarrow{\mathsf{M}}[i]}$. Hence, it is enough to prove that the multimapping f, defined as $f(\mathbf{u}_i, \mathbf{v}_i) = \mathbf{t}_{\overrightarrow{\mathsf{M}}[i]}$, $1 \le i \le n$, is a mapping such that $f(\mathbf{u}) = \mathbf{t}$.

Let us first show that f is indeed a mapping. Denote $\pi_m \hat{\mathbf{t}}$ the constant in the m-th position of the tuple $\hat{\mathbf{t}}$. Since for every i we have that $R_i = R_{\overrightarrow{\mathsf{M}}[i]}$, the size of the tuple of variables $\mathbf{u}_i, \mathbf{v}_i$ is the same as the size of the tuple of constants $\mathbf{t}_{\overrightarrow{\mathsf{M}}[i]}$. So, f maps every variable from $\mathbf{u}, \mathbf{v}$ to a nonempty set of constants from $\mathbb{D}$. Hence, it is enough to show that for each $1 \le i \le n$, $1 \le j \le n$, and for each pair of positions m, m' such that the variable in the m-th position of the tuple $\mathbf{u}_i, \mathbf{v}_i$ and the variable in the m'-th position of the tuple $\mathbf{u}_j, \mathbf{v}_j$ coincide, we have that

$$\pi_m \mathbf{t}_{\vec{\mathsf{M}}[i]} = \pi_{m'} \mathbf{t}_{\vec{\mathsf{M}}[j]}. \tag{21}$$

On the one hand, from the precondition of the requirement 2, the fact that the $\vec{\mathsf{M}}_1$, corresponding to the pair (i, j), is in $\vec{\mathsf{P}}$, and equation (20) we know that there exists a mapping f_1 in $\mathcal{V}(Q, \mathbf{t})$ such that $f_1(\mathbf{u}_i, \mathbf{v}_i) = \mathbf{t}_{\vec{\mathsf{M}}_1[i]} = \mathbf{t}_{\vec{\mathsf{M}}[i]}$. In particular, we have that $\pi_m \mathbf{t}_{\vec{\mathsf{M}}[i]} = \pi_m \mathbf{t}_{\vec{\mathsf{M}}_1[i]}$. On the other hand, since by our assumption the m-th variable of $\mathbf{u}_i, \mathbf{v}_i$ and the m'-th variable of $\mathbf{u}_j, \mathbf{v}_j$ coincide and f_1 is a function, we have that $\pi_m \mathbf{t}_{\vec{\mathsf{M}}_1[i]} = \pi_{m'} \mathbf{t}_{\vec{\mathsf{M}}_1[j]}$. Continuing such reasoning, we have the chain

$$\begin{aligned}\pi_m \mathbf{t}_{\vec{\mathsf{M}}[i]} = \pi_m \mathbf{t}_{\vec{\mathsf{M}}_1[i]} = \pi_{m'} \mathbf{t}_{\vec{\mathsf{M}}_1[j]} = \pi_{m'} \mathbf{t}_{\vec{\mathsf{M}}_2[j]} = \pi_m \mathbf{t}_{\vec{\mathsf{M}}_2[i]} = \dots \\ = \pi_m \mathbf{t}_{\vec{\mathsf{M}}_{2k+1}[i]} = \pi_{m'} \mathbf{t}_{\vec{\mathsf{M}}_{2k+1}[j]} = \pi_{m'} \mathbf{t}_{\vec{\mathsf{M}}[j]},\end{aligned}$$

which justifies equation (21).

To complete the proof of the proposition we need to show that $f(\mathbf{u}) = \mathbf{t}$. Let u be a free variable from $\mathbf{u}$, and i be the number of an atom in Q such that u appears in the tuple $\mathbf{u}_i$. Consider the o-monomial $\vec{\mathsf{M}}_1$ corresponding to the pair (i, i). We know that there exists a mapping f_1 in $\mathcal{V}(Q, \mathbf{t})$ such that $f_1(\mathbf{u}_i, \mathbf{v}_i) = \mathbf{t}_{\vec{\mathsf{M}}_1[i]} = \mathbf{t}_{\vec{\mathsf{M}}[i]} = f(\mathbf{u}_i, \mathbf{v}_i)$. In particular, this means that $f_1(u) = f(u)$. By the definition of $\mathcal{V}(Q, \mathbf{t})$ we have that $f_1(\mathbf{u}) = \mathbf{t}$. The free choice of u in $\mathbf{u}$ implies the required $f(\mathbf{u}) = \mathbf{t}$ (i.e. $f \in \mathcal{V}(Q, \mathbf{t})$). $\square$

As promised in Sec. 4.1, we obtain Prop. 4.6, which states that for every CQ-admissible polynomial P there exists a CQ Q without free variables and $\mathcal{N}[X]$-instance I with only unique variables or 0 as annotations such that $Q^I() = \mathsf{P}$, as an immediate corollary of the proposition above.

With this characterization of CQ-admissible polynomials in hand, it is straightforward to design an NP procedure which checks whether a polynomial is in $\mathbb{N}^{cq}[X]$. We leave the question of the exact complexity of this problem open.

6. CONTAINMENT VIA SMALL MODELS

Up to now we have studied how to decide $\mathcal{K}$-containment of CQs by analyzing their structure, resulting in several classes of semirings $\mathcal{K}$ for which the existence of a homomorphism of a corresponding type between the CQs is equivalent to their $\mathcal{K}$-containment. It is natural to ask whether the problem of decidability of CQ $\mathcal{K}$-CONTAINMENT can be solved by different techniques for some semirings which are not in any of these classes. Indeed, several other approaches have appeared in the literature. Green [2011] suggested a PSPACE algorithm for checking $\mathcal{N}[X]$-containment of UCQs over provenance polynomials $\mathcal{N}[X]$, based on the fact that if a UCQ $\mathbf{Q}_1$ is not $\mathcal{N}[X]$-contained in a UCQ $\mathbf{Q}_2$, then there exists a witnessing $\mathcal{N}[X]$-instance, with its size bounded by the size of $\mathbf{Q}_1$ and $\mathbf{Q}_2$. Another approach is to cast the problem of decidability of $\mathcal{K}$-containment as the problem of checking the corresponding order $\preceq_{\mathcal{K}}$ on polynomials, as done by Ioannidis and Ramakrishnan [1995] to show undecidability of UCQ $\mathcal{N}$-CONTAINMENT for bag semantics, or by Green [2011] to design algorithms for containment of CQs over different types of provenance.

The main result of this section is that by combining these ideas one can obtain new decidability results for $\mathcal{K}$-containment of CQs over different semirings $\mathcal{K}$. Particularly, we concentrate on the $\oplus$-*idempotent* semirings, i.e. the semirings where $x =_{\mathcal{K}} x + x$ holds. We denote by $\mathbf{S}^1_{\preceq}$ the class of all such semirings. This is quite a large class: for example, since by the simple reduction

$$\mathbb{1} + x =_{\mathcal{K}} \mathbb{1} \implies (\mathbb{1} + x)y =_{\mathcal{K}} y \implies y + xy =_{\mathcal{K}} y \implies y + y =_{\mathcal{K}} y$$

for any $\mathcal{K} \in \mathbf{S}_{\text{in}}$, we can obtain that $\mathbb{1}$-annihilation implies $\oplus$-idempotence, we have that $\mathbf{S}^{\mathbb{1}}_{\preceq}$ contains the whole class $\mathbf{S}_{\text{in}}$, and, hence, the classes $\mathbf{C}_{\text{hom}}$ and $\mathbf{C}_{\text{in}}$. Also, $\mathbf{S}^{\mathbb{1}}_{\preceq}$ has non-empty intersections with other classes we have considered so far – $\mathbf{C}_{\text{hcov}}$, $\mathbf{C}_{\text{sur}}$, and $\mathbf{C}_{\text{bi}}$.

The goal of this section is to introduce a small model property for semirings in $\mathbf{S}^{\mathbb{1}}_{\preceq}$. In order to do so, we will reduce the problem of containment of CQs over such semirings to the problem of checking the order on polynomials. We will then show how to check such an order for several individual semirings. In contrast to the rest of this paper, we leave a comprehensive description of semirings for which this approach works for future research. Let us start by describing the intuition behind our approach.

It is straightforward to show that for any semiring $\mathcal{K}$ from $\mathbf{S}^{\mathbb{1}}_{\preceq}$ and two CQs Q_1, Q_2 with the same set of free variables $\mathbf{u}$ it holds that

$$Q_1 \subseteq_{\mathcal{K}} Q_2 \Longrightarrow Q_1^{[\![Q_1]\!]}(\mathbf{u}) \preceq_{\mathcal{K}} Q_2^{[\![Q_1]\!]}(\mathbf{u}), \tag{22}$$

i.e. the inequality $Q_1^{[\![Q_1]\!]}(\mathbf{u}) \preceq_{\mathcal{K}} Q_2^{[\![Q_1]\!]}(\mathbf{u})$ is a necessary condition for $\mathcal{K}$-containment of CQs. Here $Q_1^{[\![Q_1]\!]}(\mathbf{u})$ and $Q_2^{[\![Q_1]\!]}(\mathbf{u})$ are again polynomials from $\mathcal{N}[X]$, and this inequality is the polynomial notation (defined in Sec. 3.2) for the statement that for any values $a_1, \ldots, a_n$ from $\mathcal{K}$, it holds that the valuation of the polynomial $Q_1^{[\![Q_1]\!]}(\mathbf{u})$ over $a_1, \ldots, a_n$ is less or equal (according to the partial order $\preceq_{\mathcal{K}}$) than the valuation of $Q_2^{[\![Q_1]\!]}(\mathbf{u})$ over the same values $a_1, \ldots, a_n$.

Unfortunately, it is not that difficult to construct an example where the other direction of (22) does not hold for a semiring $\mathcal{K}$ from $\mathbf{S}^{\mathbb{1}}_{\preceq}$, i.e. the inequality $Q_1^{[\![Q_1]\!]}(\mathbf{u}) \preceq_{\mathcal{K}} Q_2^{[\![Q_1]\!]}(\mathbf{u})$ is not a sufficient condition for $\mathcal{K}$-containment of CQs Q_1 and Q_2. As the main result of this section we show that it is possible to extend the idea above and obtain a condition based on comparison of polynomials, which is both necessary and sufficient for $\mathcal{K}$-containment for any semiring $\mathcal{K}$ from $\mathbf{S}^{\mathbb{1}}_{\preceq}$. Particularly, this condition consists of comparing not only the evaluations of Q_1 and Q_2 over the canonical instance $[\![Q_1]\!]$, but over all instances obtained from $[\![Q_1]\!]$ by identifying some variables from its domain $\mathbf{u} \cup \mathbf{v}_1$. In the following theorem we use the set of such instances as a way of describing all possibilities of assigning the variables of Q_1 to the elements of the domain of a $\mathcal{K}$-instance I by a mapping from Q_1 to I. However, we start with a bit of new notation and a lemma.

Recall that for a CQ $Q = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and a tuple $\mathbf{t}$ we have denoted by $\mathcal{V}(Q, \mathbf{t})$ the set of all mappings f from $\mathbf{u} \cup \mathbf{v}$ to the domain of I such that $f(\mathbf{u}) = \mathbf{t}$. For every such a mapping $f \in \mathcal{V}(Q, \mathbf{t})$ we also denote by I_f a $\mathcal{K}$-instance over the same domain defined for every relation R and tuple s by

$$R^{I_f}(\mathbf{s}) = \begin{cases} R^I(\mathbf{s}), & \text{if there exists } i, 1 \leq i \leq n, \text{ such that } R_i(f(\mathbf{u}_i, \mathbf{v}_i)) = R(\mathbf{s}), \\ \mathbb{0}, & \text{otherwise.} \end{cases}$$

Essentially, the instance I_f is the restriction of I on the image of the mapping f.

LEMMA 6.1. *Let $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1} \rangle$ be a semiring from $\mathbf{S}^{\mathbb{1}}_{\preceq}$, Q be a CQ, I be a $\mathcal{K}$-instance, and $\mathbf{t}$ be a tuple. Then*

(*1*) *for every CQ Q_1 with the same number of free variables as Q it holds that*

$$\sum_{f \in \mathcal{V}(Q,\mathbf{t})} Q_1^{I_f}(\mathbf{t}) \preceq_{\mathcal{K}} Q_1^I(\mathbf{t});$$

(2) *it holds that*

$$\sum_{f \in \mathcal{V}(Q,\mathbf{t})} Q^{I_f}(\mathbf{t}) = Q^I(\mathbf{t}).$$

PROOF. For **Part 1**, note that since each $\mathcal{K}$-instance I_f is a restriction of I, every product in the sum of the left hand side is also in the sum of the right hand side. However, $\mathcal{K}$ is $\oplus$-idempotent, so the multiplicities of these products are not important. Hence we have the required inequality from positivity of $\mathcal{K}$.

For **Part 2** it is left to show that $\sum_{f \in \mathcal{V}(Q,\mathbf{t})} Q^{I_f}(\mathbf{t}) \succeq_{\mathcal{K}} Q^I(\mathbf{t})$. But this holds since every product in the right hand side corresponds to some mapping f from $\mathcal{V}(Q,\mathbf{t})$, and hence this product is equal to a product in $Q^{I_f}(\mathbf{t})$. Again, by positivity we obtain the desired inequality. $\square$

Now we are ready to state the theorem.

THEOREM 6.2. *Let $\mathcal{K} = \langle K, \oplus, \otimes, \mathbb{0}, \mathbb{1}\rangle$ be a semiring from $\mathbf{S}^1_{\preceq}$ and Q_1, Q_2 be two CQs. Then $Q_1 \subseteq_{\mathcal{K}} Q_2$ iff $Q_1^{[\![\pi(Q_1)]\!]}(\mathrm{s}) \preceq_{\mathcal{K}} Q_2^{[\![\pi(Q_1)]\!]}(\mathrm{s})$ holds for every function $\pi : \mathbf{u} \cup \mathbf{v} \to \mathbf{u} \cup \mathbf{v}$ and every tuple* s *of variables from* $\mathbf{u} \cup \mathbf{v}$.

PROOF. Let $Q_1 = \exists \mathbf{v}\ \phi(\mathbf{u}, \mathbf{v})$.

We begin with the **"only if"** direction. Let $Q_1 \subseteq_{\mathcal{K}} Q_2$ and, for the sake of contradiction, $\pi : \mathbf{u} \cup \mathbf{v} \to \mathbf{u} \cup \mathbf{v}$ be a function and s be a tuple of variables from $\mathbf{u} \cup \mathbf{v}$ such that

$$Q_1^{[\![\pi(Q_1)]\!]}(\mathrm{s}) \not\preceq_{\mathcal{K}} Q_2^{[\![\pi(Q_1)]\!]}(\mathrm{s}).$$

This inequality makes sense, since the variables of the tuple s are among the elements of the domain of $[\![\pi(Q_1)]\!]$, and the expressions $Q_1^{[\![\pi(Q_1)]\!]}(\mathrm{s})$, $Q_2^{[\![\pi(Q_1)]\!]}(\mathrm{s})$ are polynomials from $\mathcal{N}[X]$. By the polynomial notation this inequality means that there exists an assignment $\tau : X \to K$ such that the value of the first polynomial for τ is not less than the value of the second one. We can write it as

$$\tau(Q_1^{[\![\pi(Q_1)]\!]}(\mathrm{s})) \not\preceq_{\mathcal{K}} \tau(Q_2^{[\![\pi(Q_1)]\!]}(\mathrm{s})).$$

Consider now the $\mathcal{K}$-instance $\tau([\![\pi(Q_1)]\!])$ obtained from $[\![\pi(Q_1)]\!]$ by substituting the abstract annotations of tuples for the corresponding values from τ. For the CQs Q_1 and Q_2 we have that $Q_1^{\tau([\![\pi(Q_1)]\!])}(\mathrm{s}) = \tau(Q_1^{[\![\pi(Q_1)]\!]}(\mathrm{s}))$ and $Q_2^{\tau([\![\pi(Q_1)]\!])}(\mathrm{s}) = \tau(Q_2^{[\![\pi(Q_1)]\!]}(\mathrm{s}))$. Hence, it holds that

$$Q_1^{\tau([\![\pi(Q_1)]\!])}(\mathrm{s}) \not\preceq_{\mathcal{K}} Q_2^{\tau([\![\pi(Q_1)]\!])}(\mathrm{s}).$$

However, this contradicts the fact that $Q_1 \subseteq_{\mathcal{K}} Q_2$, since we have a witnessing $\mathcal{K}$-instance $\tau([\![\pi(Q_1)]\!])$ for which the containment does not hold.

Next we show the **"if"** direction. Assume that for every function $\pi : \mathbf{u} \cup \mathbf{v} \to \mathbf{u} \cup \mathbf{v}$ and every tuple s of variables from $\mathbf{u} \cup \mathbf{v}$ it holds that

$$Q_1^{[\![\pi(Q_1)]\!]}(\mathrm{s}) \preceq_{\mathcal{K}} Q_2^{[\![\pi(Q_1)]\!]}(\mathrm{s}).$$

We need to show that for an arbitrary $\mathcal{K}$-instance I and tuple $\mathbf{t}$ it holds that

$$Q_1^I(\mathbf{t}) \preceq_{\mathcal{K}} Q_2^I(\mathbf{t}).$$

From Lem. 6.1 we have that

$$\sum_{f \in \mathcal{V}(Q_1,\mathbf{t})} Q_1^{I_f}(\mathbf{t}) = Q_1^I(\mathbf{t}) \quad \text{and} \quad \sum_{f \in \mathcal{V}(Q_1,\mathbf{t})} Q_2^{I_f}(\mathbf{t}) \preceq_{\mathcal{K}} Q_2^I(\mathbf{t}).$$

Thus, it is left to show that

$$\sum_{f \in \mathcal{V}(Q_1,\mathbf{t})} Q_1^{I_f}(\mathbf{t}) \preceq_{\mathcal{K}} \sum_{f \in \mathcal{V}(Q_1,\mathbf{t})} Q_2^{I_f}(\mathbf{t}).$$

The sums on both sides here are over the same set of mappings. From positivity of the semiring $\mathcal{K}$, it is therefore sufficient to prove that for every $f \in \mathcal{V}(Q_1, \mathbf{t})$ it holds that

$$Q_1^{I_f}(\mathbf{t}) \preceq_{\mathcal{K}} Q_2^{I_f}(\mathbf{t}). \tag{23}$$

We will show this by contradiction. Let f be a mapping from $\mathcal{V}(Q_1, \mathbf{t})$ such that $Q_1^{I_f}(\mathbf{t}) \not\preceq_{\mathcal{K}} Q_2^{I_f}(\mathbf{t})$. Let also $\tau : X \to K$ be an assignment such that for its extension to 0 by $\tau(0) = \mathbb{0}$ it holds that $\tau([\![f(Q_1)]\!]) = I_f$ (such an assignment exists since I_f is the image of f).

Without loss of generality we can assume that the domain $\mathbb{D}$ of the $\mathcal{K}$-instance I_f contains exactly $|\mathbf{u} \cup \mathbf{v}|$ elements: indeed, the supports $\{\mathbf{t} \mid \mathbf{t} \in \mathbb{D}^m, R^{I_f}(\mathbf{t})) \neq \mathbb{0}\}$ of all relations R cannot contain more than $|\mathbf{u} \cup \mathbf{v}|$ elements by the definition of I_f, so we can safely remove some of the others, if there are more of them; if there are fewer elements, then we can introduce new artificial ones with a $\mathbb{0}$ annotation for each tuple containing them.

Consider now an arbitrary bijection $h : \mathbb{D} \to \mathbf{u} \cup \mathbf{v}$ (which exists by the assumption above) and the function $\pi = h \circ f$, which maps $\mathbf{u} \cup \mathbf{v}$ to itself. For this function we have that

$$\tau(Q_1^{[\![\pi(Q_1)]\!]}(\pi(\mathbf{u}))) = Q_1^{I_f}(\mathbf{t}) \quad \text{and} \quad \tau(Q_2^{[\![\pi(Q_1)]\!]}(\pi(\mathbf{u}))) = Q_2^{I_f}(\mathbf{t}).$$

By the initial assumption we have that

$$\tau(Q_1^{[\![\pi(Q_1)]\!]}(\pi(\mathbf{u}))) = Q_1^{I_f}(\mathbf{t}) \not\preceq_{\mathcal{K}} Q_2^{I_f}(\mathbf{t}) = \tau(Q_2^{[\![\pi(Q_1)]\!]}(\pi(\mathbf{u}))).$$

However, this contradicts the fact that for the function π and the tuple of variables $\pi(\mathbf{u})$ it holds that

$$Q_1^{[\![\pi(Q_1)]\!]}(\pi(\mathbf{u})) \preceq_{\mathcal{K}} Q_2^{[\![\pi(Q_1)]\!]}(\pi(\mathbf{u})).$$

Hence, the inequality (23) holds and the proof of the theorem is completed. □

This theorem shows that for $\mathcal{K} \in \mathbf{S}^1_{\preceq}$ the CQ $\mathcal{K}$-CONTAINMENT problem can be reduced to a small number of problems of checking the order $\preceq_{\mathcal{K}}$ between CQ-admissible polynomials.

COROLLARY 6.3. *If $\mathcal{K} \in \mathbf{S}^1_{\preceq}$ and it is decidable to check whether $\mathsf{P}_1 \preceq_{\mathcal{K}} \mathsf{P}_2$ for any pair of polynomials $\mathsf{P}_1, \mathsf{P}_2$ from $\mathbb{N}^{cq}[X]$, then* CQ $\mathcal{K}$-CONTAINMENT *is decidable.*

We do not investigate the decidability of $\mathsf{P}_1 \preceq_{\mathcal{K}} \mathsf{P}_2$ for the entire class $\mathbf{S}^1_{\preceq}$, but do so for some of its most important members that do not have any corresponding type of homomorphism, considered above – the tropical semiring $\mathcal{T}^+$ and the max-plus algebra $\mathcal{T}^-$. For these particular semirings, we will show that containment of CQs can be decided in the second level of the polynomial hierarchy. To do so we need the following technical lemma.

LEMMA 6.4. *Let $\mathcal{K}$ be a semiring from $\mathbf{S}^1_{\preceq}$, $\mathsf{M}_1, \dots \mathsf{M}_n$ be monomials over a set of variables X, and P be a polynomial from $\mathcal{N}[\bar{X}]$. Then $\mathsf{M}_1 + \ldots + \mathsf{M}_n \preceq_{\mathcal{K}} \mathsf{P}$ iff $\mathsf{M}_i \preceq_{\mathcal{K}} \mathsf{P}$ holds for each $1 \leq i \leq n$.*

PROOF. The **"only if"** direction is straightforward, since $\mathsf{M}_i \preceq_{\mathcal{K}} \mathsf{M}_1 + \ldots + \mathsf{M}_n$ for each $1 \leq i \leq n$. The **"if"** direction follows from the chain

$$\mathsf{M}_1 + \ldots + \mathsf{M}_n \preceq_{\mathcal{K}} \underbrace{\mathsf{P} + \ldots + \mathsf{P}}_{n \text{ times}} =_{\mathcal{K}} \mathsf{P},$$

which holds by positivity and $\oplus$-idempotence of $\mathcal{K}$. □

PROPOSITION 6.5. CQ $\mathcal{T}^+$- *and* $\mathcal{T}^-$-CONTAINMENT *are in* Π_2^p.

PROOF. We concentrate on the case of max-plus algebra $\mathcal{T}^-$. For the case of the tropical semiring the proof differs only in instantiations of semiring operations and order.

Consider arbitrary conjunctive queries $Q_1(\mathbf{u}) = \exists \mathbf{v}\ R_1(\mathbf{u}_1, \mathbf{v}_1), \ldots, R_n(\mathbf{u}_n, \mathbf{v}_n)$ and $Q_2(\mathbf{q}) = \exists \mathbf{w}\ S_1(\mathbf{q}_1, \mathbf{w}_1), \ldots, S_m(\mathbf{q}_m, \mathbf{w}_m)$, where $\mathbf{u}$ is the tuple of free variables of Q_1 and $\mathbf{q}$ is the tuple of free variables of Q_2, each $\mathbf{u}_i$ and $\mathbf{q}_j$ consist of variables from $\mathbf{u}$ and $\mathbf{q}$, respectively, and each $\mathbf{v}_i$ and $\mathbf{w}_j$ consist of variables from $\mathbf{v}$ and $\mathbf{w}$, respectively. By Thm. 6.2 it suffices to develop an algorithm which decides whether for every function $\pi : \mathbf{u} \cup \mathbf{v} \to \mathbf{u} \cup \mathbf{v}$ and every tuple $\mathbf{s}$ of variables from $\mathbf{u} \cup \mathbf{v}$ it holds that

$$Q_1^{\pi(\llbracket Q_1 \rrbracket)}(\mathbf{s}) \preceq_{\mathcal{T}^-} Q_2^{\pi(\llbracket Q_1 \rrbracket)}(\mathbf{s}). \tag{24}$$

By the definition of evaluations, this inequality can be explicitly written as

$$\sum_{f \in \mathcal{V}(Q_1, \mathbf{s})} \prod_{1 \leq i \leq n} R_i^{\pi(\llbracket Q_1 \rrbracket)}(f(\mathbf{u}_i, \mathbf{v}_i)) \preceq_{\mathcal{T}^-} \sum_{h \in \mathcal{V}(Q_2, \mathbf{s})} \prod_{1 \leq j \leq m} S_j^{\pi(\llbracket Q_1 \rrbracket)}(h(\mathbf{q}_j, \mathbf{w}_j)).$$

The expressions on both sides are polynomials from $\mathcal{N}[X]$, where X is some set of variables. Without loss of generality we may assume that $X = x_1, \ldots, x_n$. Hence, by Lem. 6.4, this inequality is equivalent to the statement

$$\forall f \in \mathcal{V}(Q_1, \mathbf{s}) \left(\prod_{1 \leq i \leq n} R_i^{\pi(\llbracket Q_1 \rrbracket)}(f(\mathbf{u}_i, \mathbf{v}_i)) \preceq_{\mathcal{T}^-} \sum_{h \in \mathcal{V}(Q_2, \mathbf{s})} \prod_{1 \leq j \leq m} S_j^{\pi(\llbracket Q_1 \rrbracket)}(h(\mathbf{q}_j, \mathbf{w}_j)) \right),$$

which is polynomial notation for the following formula over the domain of the semiring $\mathcal{T}^-$ (recall that its natural order coincides with the usual order $\leq$ on integers):

$$\forall f \in \mathcal{V}(Q_1, \mathbf{s}) \left(\forall x_1, \ldots, x_n \sum_{1 \leq i \leq n} R_i^{\pi(\llbracket Q_1 \rrbracket)}(f(\mathbf{u}_i, \mathbf{v}_i)) \leq \max_{h \in \mathcal{V}(Q_2, \mathbf{s})} \sum_{1 \leq j \leq m} S_j^{\pi(\llbracket Q_1 \rrbracket)}(h(\mathbf{q}_j, \mathbf{w}_j)) \right).$$

Consider now the inequality within the parentheses. If a sum on the right hand side contains a variable which does not appear in the left hand side, then this sum does not influence the truth of the formula, since we can instantiate this variable to $-\infty$. So we can ignore this sum, and can assume that the left hand side contains all the variables from $x_1, \ldots, x_n$ which occur in the right hand side. If, in turn, one of these variables is equal to $-\infty$, then the formula clearly holds. Otherwise, the expression within the parentheses holds if and only if the following linear integer programming system Γ does not have a solution over $\mathbb{N}_0$:

$$\boxed{\begin{array}{rcl} x_1, \ldots, x_n & \geq & 0, \\ \sum_{1 \leq i \leq n} R_i^{\pi(\llbracket Q_1 \rrbracket)}(f(\mathbf{u}_i, \mathbf{v}_i)) & > & \sum_{1 \leq j \leq m} S_j^{\pi(\llbracket Q_1 \rrbracket)}(h(\mathbf{q}_j, \mathbf{w}_j)),\ h \in \mathcal{V}(Q_2, \mathbf{s}). \end{array}}$$

Since $\mathcal{V}(Q_2, \mathbf{s})$ can have an exponential number of mappings h, the system Γ can contain an exponential number of inequalities. However, the dimension (i.e. the number n

of variables) of this system is just polynomial in the size of the CQ Q_1. From [Kannan and Monma 1978; Lenstra 1983] it is known that there is a fixed polynomial p such that, if a system such as Γ has a solution, then it must have a solution where each x_i is assigned a value at most $O(2^{p(n)})$, where $p(n)$ is a valuation of p for n.[8] This means that such a solution has a polynomial binary representation.

These observations justify the following algorithm which checks that (24) does not hold:

(1) guess
- a function $\pi : \mathbf{u} \cup \mathbf{v} \to \mathbf{u} \cup \mathbf{v}$,
- a tuple s of variables from $\mathbf{u} \cup \mathbf{v}$,
- a mapping f in $\mathcal{V}(Q_1, \mathrm{s})$,
- a set of annotation values $a_1, \ldots, a_\ell$ from $\{-\infty, 0, \ldots, 2^{p(n)}\}$;

(2) call an oracle which guesses a mapping h in $\mathcal{V}(Q_2, \mathrm{s})$;

(3) check that

$$\sum_{1 \le i \le n} R_i^{\pi([\![Q_1]\!])}(f(\mathbf{u}_i, \mathbf{v}_i)) \le \sum_{1 \le j \le m} S_j^{\pi([\![Q_1]\!])}(h(\mathbf{q}_j, \mathbf{w}_j))$$

does not hold for the values $a_1, \ldots, a_\ell$ of variables $x_1, \ldots, x_\ell$.

All the guesses of this algorithm are of polynomial size in the size of the input CQs Q_1 and Q_2. Also, the check in step (3) can be done in polynomial time. Hence the algorithm demonstrates that the problem is in Σ_2^p. Therefore checking whether (24) holds for every function $\pi : \mathbf{u} \cup \mathbf{v} \to \mathbf{u} \cup \mathbf{v}$ and every tuple s of variables from $\mathbf{u} \cup \mathbf{v}$, is in Π_2^p. $\square$

This proposition along with Thm. 4.28 leaves a gap between upper and lower bounds for complexity of CQ $\mathcal{T}^+$- and $\mathcal{T}^-$-CONTAINMENT. We leave the exact complexity of these problems open.

To conclude this section we illustrate how to use Prop. 6.5 by a further extension of Ex. 4.13.

Example 6.13 (*Continued*). For the identity mapping Id on the variables $\mathbf{u} \cup \mathbf{v}$ of the CQ Q_1 we have that

$$Q_1^{Id([\![Q_1]\!])}() = x_1^2 + 2x_1x_2 + x_2^2, \text{ and } Q_2^{Id([\![Q_{11}]\!])}() = x_1^2 + x_2^2.$$

It is straightforward to see that

$$x_1^2 + 2x_1x_2 + x_2^2 =_{\mathcal{T}^+} x_1^2 + x_2^2.$$

The same can be shown for all other mappings $\pi : \mathbf{u} \cup \mathbf{v} \to \mathbf{u} \cup \mathbf{v}$. By Thm. 6.2 we have that $Q_1 \subseteq_{\mathcal{T}^+} Q_2$.

7. CONCLUSION

We have studied containment of CQs over relations with annotations of different types. We have established several interesting classes of semirings for which this problem is decidable by means of different syntactic criteria, developed by modifying and extending the well-known notion of homomorphism between CQs. Our work extends and systematizes previous results on the subject and should have practical implications,

[8] Actually, the bound is $O(|a|^{p(n)})$, where a is the coefficient in the system with the largest absolute value. We can use the simpler $O(2^{p(n)})$ bound because all the coefficients in our program are bounded by n.

Table I. Summary of semiring classes. The column "class" refers to the class of semirings for which the corresponding homomorphism type from the column "homomorphism type" is a sufficient and necessary condition for containment of CQs. "Sufficient class" column refers to the class of semirings for which the homomorphism type is sufficient for containment, and the "axioms of sufficient class" column defines this class. Similarly, "necessary class" refers to the class of semirings for which the homomorphism type is a necessary condition for containment, but for clarity references to the corresponding "definition of necessary class" are given instead of the axioms.

$\mathcal{K}$-**containment of CQs**						
class	homomorphism type	sufficient class	axioms of sufficient class	necessary class	definition of necessary class	Sec.
$\mathbf{C}_{\mathrm{hom}}$	$Q_2 \to Q_1$ (usual)	$\mathbf{C}_{\mathrm{hom}}$	$\otimes$-idempotence $\mathbb{1}$-annihilation	$\mathbf{S}_{\preceq}$	—	3.3
$\mathbf{C}_{\mathrm{hcov}}$	$Q_2 \rightrightarrows Q_1$ (hom. cov.)	$\mathbf{S}_{\mathrm{hcov}}$	$\otimes$-idempotence	$\mathbf{N}_{\mathrm{hcov}}$	Def. 4.3	4.1
$\mathbf{C}_{\mathrm{in}}$	$Q_2 \hookrightarrow Q_1$ (injective)	$\mathbf{S}_{\mathrm{in}}$	$\mathbb{1}$-annihilation	$\mathbf{N}_{\mathrm{in}}$	Def. 4.14	4.2
$\mathbf{C}_{\mathrm{sur}}$	$Q_2 \twoheadrightarrow Q_1$ (surjective)	$\mathbf{S}_{\mathrm{sur}}$	$\otimes$-semi-idempotence	$\mathbf{N}_{\mathrm{sur}}$	Def. 4.20	4.3
$\mathbf{C}_{\mathrm{bi}}$	$Q_2 \hookrightarrow\!\!\!\!\rightarrow Q_1$ (bijective)	$\mathbf{S}_{\preceq}$	—	$\mathbf{C}_{\mathrm{bi}}$	Def. 4.25	4.4

since most semirings used for annotations in the literature fall into one of these well-behaved classes. Tab. I provides a summary of the results on these classes. For semirings that do not fall into these classes, we have extended the range of available machinery for query optimization problems, by providing generalized or improved necessary and sufficient conditions. For some of these semirings we also suggest new decision procedures based on a small model property.

Many problems remain open. First, it would be natural to extend this research to study containment of *more powerful* queries than CQs. Set and bag containment for several classes of such more powerful queries have been considered in the literature, for instance unions of conjunctive queries [Ioannidis and Ramakrishnan 1995], queries with inequalities [Klug 1988; Jayram et al. 2006], aggregate queries [Cohen et al. 2007; 2003], and recursive Datalog programs [Shmueli 1987]. Some results for the case of unions are already presented in the conference paper [Kostylev et al. 2012]. The semantics of aggregate queries with provenance is given by Amsterdamer et al. [2011].

The second immediate direction is to study another fundamental optimization problem – *equivalence* of queries of different types over annotated relations. It is interesting to note that, for the case of equivalence, it is possible to extend the class of considered semirings by relaxing some axioms of $\preceq$-positive semirings without loss of a meaningful semantics. Indeed, Green et al. [2011] considers the equivalence of CQs and UCQs over integers $\mathbb{Z}$. In contrast, as discussed in Sec. 3.1, $\mathbb{Z}$-containment is vacuous.

Third, we would like to continue studying the *small model property* approach, either proving or disproving that such methods can work for semirings with non-idempotent addition.

Finally, it would also be interesting to study *CQ-admissible polynomials* on their own, and in particular how to decide containment over them. We believe that this study may have consequences for solving some of the fundamental open problems in the area of query optimization.

ACKNOWLEDGMENTS

We thank Peter Buneman, Jeff Egger, Diego Figueira, and Tony Tan for useful discussions, Todd J. Green for comments on previous results, and the anonymous reviewers for suggestions on presentation.

REFERENCES

Foto N. Afrati, Matthew Damigos, and Manolis Gergatsoulis. 2010. Query containment under bag and bag-set semantics. *Inform. Process. Lett.* 110, 10 (2010), 360–369. DOI: http://dx.doi.org/10.1016/j.ipl.2010.02.017

A. Aho, Y. Sagiv, and J. Ullman. 1979. Equivalences among Relational Expressions. *SIAM J. Comput.* 8, 2 (1979), 218–246. DOI:http://dx.doi.org/10.1137/0208017

Yael Amsterdamer, Daniel Deutch, and Val Tannen. 2011. Provenance for aggregate queries. In *PODS 2011: Proceedings of the thirtieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, Maurizio Lenzerini and Thomas Schwentick (Eds.). ACM, New York, NY, USA, 153–164. DOI:http://dx.doi.org/10.1145/1989284.1989302

Peter Buneman, James Cheney, Wang-Chiew Tan, and Stijn Vansummeren. 2008. Curated databases. In *PODS 2008: Proceedings of the twenty-seventh ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, New York, NY, USA, 1–12. DOI:http://dx.doi.org/10.1145/1376916.1376918

Peter Buneman, Sanjeev Khanna, and Wang-Chiew Tan. 2001. Why and Where: A Characterization of Data Provenance. In *ICDT 2001: Proceedings of the 8th International Conference on Database Theory (Lecture Notes in Computer Science)*, Jan Van den Bussche and Victor Vianu (Eds.), Vol. 1973. Springer-Verlag, Berlin, Germany, 316–330. DOI:http://dx.doi.org/10.1007/3-540-44503-X_20

Peter Buneman and Egor V. Kostylev. 2010. Annotation Algebras for RDFS. In *SWPM 2010: Proceedings of the Second International Workshop on the role of Semantic Web in Provenance Management*. CEUR Workshop Proceedings, Aachen, Germany. http://ceur-ws.org/Vol-670/paper_4.pdf

Ashok K. Chandra and Philip M. Merlin. 1977. Optimal Implementation of Conjunctive Queries in Relational Data Bases. In *STOC 1977: Proceedings of the ninth annual ACM symposium on Theory of computing*. ACM, New York, NY, USA, 77–90. DOI:http://dx.doi.org/10.1145/800105.803397

Surajit Chaudhuri and Moshe Y. Vardi. 1993. Optimization of *real* conjunctive queries. In *PODS 1993: Proceedings of the twelfth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*. ACM, New York, NY, USA, 59–70. DOI:http://dx.doi.org/10.1145/153850.153856

Rada Chirkova. 2012. Equivalence and Minimization of Conjunctive Queries under Combined Semantics. In *ICDT 2012: Proceedings of the 15th International Conference on Database Theory*. ACM, New York, NY, USA, 262–273. DOI:http://dx.doi.org/10.1145/2274576.2274604

Sara Cohen, Werner Nutt, and Yehoshua Sagiv. 2003. Containment of aggregate queries. In *ICDT 2003: Proceedings of the 9th International Conference on Database Theory*. Springer-Verlag, Berlin, Germany, 111–125.

Sara Cohen, Werner Nutt, and Yehoshua Sagiv. 2007. Deciding equivalences among conjunctive aggregate queries. *Journal of the ACM* 54, 2 (2007), 5:1–5:50. DOI:http://dx.doi.org/10.1145/1219092.1219093

Yingwei Cui, Jennifer Widom, and Janet L. Wiener. 2000. Tracing the lineage of view data in a warehousing environment. *ACM Transactions on Database Systems* 25, 2 (June 2000), 179–227. DOI:http://dx.doi.org/10.1145/357775.357777

A. Das Sarma, M. Theobald, and J. Widom. 2008. Exploiting Lineage for Confidence Computation in Uncertain and Probabilistic Databases. In *ICDE 2008: Proceedings of the 24th IEEE International Conference on Data Engineering*. IEEE Computer Society, Los Alamitos, CA, USA, 1023–1032. DOI:http://dx.doi.org/10.1109/ICDE.2008.4497511

Samuel Eilenberg. 1974. *Automata, Languages, and Machines*. Academic Press, New York, NY, USA.

Norbert Fuhr and Thomas Rölleke. 1997. A probabilistic relational algebra for the integration of information retrieval and database systems. *ACM Transactions on Information Systems* 15, 1 (Jan. 1997), 32–66. DOI:http://dx.doi.org/10.1145/239041.239045

Gösta Grahne, Nicolas Spyratos, and Daniel Stamate. 1997. Semantics and containment of queries with internal and external conjunctions. In *ICDT 1997: Proceedings of the 6th International Conference on Database Theory (Lecture Notes in Computer Science)*, Foto Afrati and Phokion Kolaitis (Eds.), Vol. 1186. Springer-Verlag, Berlin, Germany, 71–82. DOI:http://dx.doi.org/10.1007/3-540-62222-5_37

Todd Green, Zachary Ives, and Val Tannen. 2011. Reconcilable Differences. *Theory of Computing Systems* 49, 2 (2011), 460–488. DOI:http://dx.doi.org/10.1007/s00224-011-9323-x

Todd J. Green. 2011. Containment of Conjunctive Queries on Annotated Relations. *Theory of Computing Systems* 49, 2 (2011), 429–459. DOI:http://dx.doi.org/10.1007/s00224-011-9327-6

Todd J. Green, Grigoris Karvounarakis, and Val Tannen. 2007. Provenance semirings. In *PODS 2007: Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, New York, NY, USA, 31–40. DOI:http://dx.doi.org/10.1145/1265530.1265535

Tomasz Imieliński and Witold Lipski, Jr. 1984. Incomplete Information in Relational Databases. *Journal of the ACM* 31, 4 (1984), 761–791. DOI:http://dx.doi.org/10.1145/1634.1886

Yannis E. Ioannidis and Raghu Ramakrishnan. 1995. Containment of conjunctive queries: beyond relations as sets. *ACM Transactions on Database Systems* 20, 3 (Sept. 1995), 288–324. DOI:http://dx.doi.org/10.1145/211414.211419

T. S. Jayram, Phokion G. Kolaitis, and Erik Vee. 2006. The Containment Problem for *Real* Conjunctive Queries with Inequalities. In *PODS 2006: Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, New York, NY, USA, 80–89. DOI:http://dx.doi.org/10.1145/1142351.1142363

R. Kannan and C.L. Monma. 1978. On the computational complexity of integer programming problems. In *Proceedings of the workshop on Optimization and Operations Research (Lecture Notes in Economics and Mathematical Systems)*, R. Henn, B. Korte, and W. Oettli (Eds.), Vol. 157. Springer-Verlag, Berlin, Germany, 161–172.

Grigoris Karvounarakis and Todd J. Green. 2012. Semiring-annotated Data: Queries and Provenance Algebras. *SIGMOD Record* 41, 3 (Sept. 2012), 5–14. DOI:http://dx.doi.org/10.1145/1462571.1462577

Anthony Klug. 1988. On conjunctive queries containing inequalities. *Journal of the ACM* 35, 1 (Jan. 1988), 146–160. DOI:http://dx.doi.org/10.1145/42267.42273

Egor V. Kostylev, Juan L. Reutter, and András Z. Salamon. 2012. Classification of annotation semirings over query containment. In *PODS 2012: Proceedings of the 31st symposium on Principles of database systems*. ACM, New York, NY, USA, 237–248. DOI:http://dx.doi.org/10.1145/2213556.2213590

Jr. Lenstra, H. W. 1983. Integer Programming with a Fixed Number of Variables. *Mathematics of Operations Research* 8, 4 (1983), 538–548. DOI:http://dx.doi.org/10.1287/moor.8.4.538

Dan Olteanu and Jakub Závodný. 2012. Factorised Representations of Query Results: Size Bounds and Readability. In *ICDT 2012: Proceedings of the 15th International Conference on Database Theory*. ACM, New York, NY, USA. DOI:http://dx.doi.org/10.1145/2274576.2274607

O. Shmueli. 1987. Decidability and expressiveness aspects of logic queries. In *Proceedings of the sixth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems (PODS 1987: Proceedings of the sixth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems)*. ACM, New York, NY, USA, 237–249. DOI:http://dx.doi.org/10.1145/28659.28685

Esteban Zimányi. 1997. Query evaluation in probabilistic relational databases. *Theoretical Computer Science* 171, 1–2 (1997), 179–219. DOI:http://dx.doi.org/10.1016/S0304-3975(96)00129-6

Antoine Zimmermann, Nuno Lopes, Axel Polleres, and Umberto Straccia. 2011. A general framework for representing, reasoning and querying with annotated Semantic Web data. *Web Semantics: Science, Services and Agents on the World Wide Web* 11 (March 2011), 72–95. DOI:http://dx.doi.org/10.1016/j.websem.2011.08.006

Received Month Year; revised Month Year; accepted Month Year