# Bag Containment of Join-on-Free Queries

**George Konstantinidis**
University of Southampton

**Fabio Mogavero**
Università degli Studi di Napoli Federico II, Italy

## Abstract

*Bag-semantics* allows for atomic relations and query answers to contain multiple copies of the same data tuple, reflecting real-world database systems more accurately. Deciding *containment under bag-semantics* (or simply, *bag-containment*) for two *conjunctive queries* (*CQ*s) requires determining whether the answer of the first query, taking multiplicities into account, is contained within the answer of the second query, across all databases. Despite numerous attempts in the last thirty years, this problem of determining decidability and complexity of this task remains open as one of the prominent challenges in database theory, given its relevance in important applications.

Previous works have established the decidability of the problem for specific classes of queries, among which is the the bag-containment of *projection-free queries* (*PFQ*s), *i.e.*, queries without existentially quantified variables, into general *CQ*s. In this work, we push the boundaries further by addressing a broader, yet natural, fragment of *CQ*s, called *join-on-free queries* (*JoFQ*), which allows existential variables, while prohibiting joins involving them. We prove decidability of bag-containment of a *JoFQ* within a general *CQ*, placing the complexity of the problem in the first non-deterministic layer of the exponential hierarchy. The approach involves a homomorphism-counting reduction to the solution of a system of Diophantine inequalities with a specific structure (an undecidable problem in its general form) and an algorithm designed to address this category of inequalities.

**2012 ACM Subject Classification** Theory of computation → Database query languages (principles); Theory of computation → Database query processing and optimization (theory); Theory of computation → Logic and databases; Mathematics of computing → Combinatorics

**Keywords and phrases** Query Containment; Bag Semantics; Bag Containment; Diophantine Problems

**Digital Object Identifier** 10.4230/LIPIcs.ICDT.2025.2

**Funding** G. Konstantinidis was partially funded by the Horizon Europe projects DATAPACT (No. 101189771), RAISE (No. 101058479) and UPCAST (No. 101093216). F. Mogavero is member of the Gruppo Nazionale Calcolo Scientifico-Istituto Nazionale di Alta Matematica (GNCS-INdAM) and acknowledges a partial support by GNCS 2024 project "Certificazione, Monitoraggio, ed Interpretabilità in Sistemi di Intelligenza Artificiale".

## 1 Introduction

The *query containment problem* under *bag-semantics* has persisted as an unsolved conundrum in database theory for the past thirty years. It refers to deciding whether the answer of one query (*a.k.a. containee*) is always contained within the answer of another query (*a.k.a. containing*), across all possible database instances. The problem has significant applications in various areas of computer science, such as query optimisation [4, 22, 19], data integration [26, 21], knowledge representation and reasoning [3], and data privacy [27, 13, 23]. The simplest and most-studied form of the problem concerns *conjunctive queries* (*CQ*s), which constitute the core of every structured query language, under *set-semantics*, where duplicated tuples within database relations and query answers are not considered. Under bag-semantics, instead, relations and answers are *bags*, *i.e.*, multi-sets, allowing for multiple copies of the same tuple, a possibility incidentally considered already in [15]. This semantics constitutes the default behaviour on most RDBMSs, motivating the identification and study


28th International Conference on Database Theory (ICDT 2025).
Editors: Sudeepa Roy and Ahmet Kara; Article No. 2; pp. 2:1–2:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of the associated decision problems [5, 14], which, despite numerous attempts and significant strides [20], remain open to date.

The containment problem for *CQ*s under set-semantics was shown to be NPTIME-complete in [4], equivalent to the *homomorphism detection problem* between the two queries, which, in turn, corresponds to the evaluation of the containing query on the *canonical instance* of the containee one. In a series of articles [12, 9, 10, 2, 6, 7, 8], starting with [5, 14], partial positive decidability and complexity results were presented for the bag-semantics version of the problem. These works, most relying on homomorphism manipulations, either offered robust yet incomplete criteria for containment/non-containment, or focused on providing complete decision procedures for significant subclasses of *CQ*s. On the other hand, interesting extensions of the bag-containment problem were proven to be undecidable [14, 16, 28], commonly establishing connections to the validity problem of *Diophantine inequalities*, a question tightly related with *Hilbert's 10th Problem* [11, 30, 29].

In 2011, Kopparty and Rossman [25] gave start to a new line of research that tries to solve the equivalent *homomorphism domination problem* by means of more sophisticated techniques grounded in *information theory*. They proved the decidability for particular classes of *CQ*s enjoying specific graph-theoretic structures, known as *chordal* and *series-parallel*. More recently, Khamis *et al.* [17, 18] expanded upon this idea, proving that the bag-containment problem involving arbitrary containee queries into acyclic containing ones is equivalent to an unresolved question regarding the validity of certain *information inequalities*. They also showed decidability for the chordal containing queries having a *simple junction tree*.

In [24] we studied a natural fragment of *CQ*s, and proved that deciding the bag containment of a *projection-free query* (*PFQ*), *i.e.*, a query without existentially quantified variables, into an arbitrary *CQ* is in $\Pi_2^{\mathrm{P}}$ and NPTIME-HARD. This work first reduced the containment problem to solving a special case of Diophantine inequalities, known as *monomial-polynomial inequalities* (*MPI*), and then presented a solution procedure for the latter, based on the resolution of a linear system. Notably, this approach marks the first use of Diophantine inequalities to prove decidability of a bag-semantics question.

In this work, we continue in the steps of [24] by presenting a decidability result for a broader, yet equally natural, class of containee queries, called *join-on-free queries* (*JoFQ*), which allows existential variables, but forbids their participation in joins. Surprisingly, despite being often used for examples in bag-containment literature [5, 2], this class has never been formally identified or studied. To tackle the problem, first, we prove that we can focus on bag-set containment of Boolean *JoFQ*s (*JoFBQ*) into Boolean *CQ*s (*BCQ*) (Theorems 2, 3 and Problem 1). Then, we devise the notion of *multicanonical instance* (Def. 8) induced by the *minimal unification closure* (Def. 4) of the containee query. Such an instance represents the simplest structure that potentially disproves containment. In addition, we introduce the notion of *net-images* (Def. 5), which are sets of image tuples exclusive to different query atoms. This enables the counting (Thm. 9) and comparison (Thm. 10) of the number of homomorphisms of the two queries, modelling the containee as a *monomial function* and the containing as a *polynomial function* (Thm. 15). This results in an *MPI* of a more general form than those of [24], also coupled with additional linear constraints (Thm. 16). Then, we delve into solving the resulting system of Diophantine inequalities, expanding upon the technique in [24], where an algorithm is described for *MPI*s with only positive coefficients. Nevertheless, coming to the rescue, the addressed polynomials enjoy a property that we identify as *strong non-negativeness* (Def. 18). The linear constraints which supplement the *MPI* make a direct reduction to the solution of a linear system unfeasible and so we engage in a series of intermediate results to achieve the reduction. In particular, we first solve a more specific instance of the problem (Prob. 3 and Thm. 23) and then leverage this to attack the original one (Prob. 2 and Thm. 27), providing so an $\exists\forall$-alternating polynomial-time solution (Thm. 28). To the best of our knowledge, this particular Diophantine problem has not been explored in the existing literature. Summing up, we prove that both bag-bag and bag-set containments of *JoFQ*s into *CQ*s are decidable in CoNExpTime (Thm. 30).

## 2 Preliminaries

We consider the standard logic notions of *relations*, *variables*, *constants*, *terms* (variables or constants), $n$-tuples $\bar{t}$ of terms, with $|\bar{t}| = n$ their *arity*, *atoms* $R(\bar{t})$, and *ground atoms*, *a.k.a. facts*, $R(\bar{c})$ where $\bar{c}$ is a tuple of constants [1]. Relations are sets of facts. $\mathrm{Vr}(s)$, $\mathrm{Cn}(s)$, and $\mathrm{Tr}(s)$ denote the set of variables, constants, and terms in any syntactic expression $s$. A function $\mathsf{h}$ between terms is called *mapping*. By $\mathsf{h}(s)$ we denote the expression obtained by replacing all terms in $s$ with their $\mathsf{h}$-image; terms outside the domain of $\mathsf{h}$ remain unchanged.

Given $\mathrm{C} \subseteq \mathrm{Cn}$, $\mathrm{T}_1, \mathrm{T}_2 \subseteq \mathrm{Tr}$, a mapping $\mathsf{h}\colon \mathrm{T}_1 \to \mathrm{T}_2$ is *C-preserving* if $\mathsf{h}(c) = c$ for all $c \in \mathrm{C} \cap \mathrm{T}_1$. Given two sets of atoms $\mathrm{S}_1$ and $\mathrm{S}_2$, a *homomorphism* $\mathsf{h}$ *from* $\mathrm{S}_1$ *to* $\mathrm{S}_2$ is a $\mathrm{Cn}(\mathrm{S}_1)$-preserving mapping $\mathsf{h}\colon \mathrm{Tr}(\mathrm{S}_1) \to \mathrm{Tr}(\mathrm{S}_2)$ such that $\mathsf{h}(\mathrm{S}_1) \subseteq \mathrm{S}_2$. The set of homomorphisms from $\mathrm{S}_1$ to $\mathrm{S}_2$ is denoted by $\mathrm{Hom}(\mathrm{S}_1, \mathrm{S}_2)$. For two atoms $a_1$ and $a_2$, we write $a_1 \preceq_{\mathrm{C}} a_2$ if there is a C-preserving mapping $\mathsf{h}$ such that $\mathsf{h}(a_2) = a_1$; $\preceq$ is called the *homomorphic ordering* over atoms. Obviously, $a_1 \preceq_{\mathrm{Cn}(a_2)} a_2$ implies the existence of a homomorphism from $a_2$ to $a_1$. The expressions $a_1 \approx_{\mathrm{C}} a_2$ and $a_1 \prec_{\mathrm{C}} a_2$ have the obvious meaning. For a syntactic expression $s$, we assume $\preceq_s$ to mean $\preceq_{\mathrm{Cn}(s)}$. Given atoms $a_1, a_2$ and $\mathrm{C} = \mathrm{Cn}(a_1) \cup \mathrm{Cn}(a_2)$, a mapping $\mathsf{u}\colon \mathrm{Tr}(a_1) \cup \mathrm{Tr}(a_2) \to \mathrm{Tr}(a_1) \cup \mathrm{Tr}(a_2)$ is a *unifier* for $a_1$ and $a_2$ if it is C-preserving and $\mathsf{u}(a_1) = \mathsf{u}(a_2)$ (unifiers do not always exist). The *most general unifier* for $a_1$ and $a_2$ is a unifier denoted $\mathsf{mgu}(a_1, a_2)$ such that, for every other unifier $\mathsf{u}$ for $a_1$ and $a_2$, it holds that $\mathsf{u} \preceq_{\mathrm{C}} \mathsf{mgu}(a_1, a_2)$ (meaning, for every atom $a_3$, $\mathsf{u}(a_3) \preceq_{\mathrm{C}} \mathsf{mgu}(a_1, a_2)(a_3)$). We lift unifiers to sets and write $\mathsf{mgu}(\mathrm{S})$ for the atom that is the result of the unification on S.

A *conjunctive query* (*CQ*) $q(\bar{x})$ is a first-order formula of the form $\exists \bar{y} \bigwedge_{i=1}^{n} R_i(\bar{x_i}, \bar{y_i})$, where $\bigwedge_{i=1}^{n} R_i(\bar{x_i}, \bar{y_i})$ is a conjunction of $n$ (possibly repeated) atoms. We often use the datalog notation for a *CQ*: $q(\bar{x}) \leftarrow R_1(\bar{x_1}, \bar{y_1}), \ldots, R_n(\bar{x_n}, \bar{y_n})$. The set of atoms in a query $q$ is denoted by $\mathsf{body}(q)$. When clear from context we use $q$, $q(\bar{x})$ or *body*$(q)$ interchangeably. When $\bar{x} = \emptyset$ the query is Boolean (*BCQ*). A *set database instance* (or instance) $I$ is a set of facts (belonging to relation instances). The *answer under set semantics* of a CQ $q(\bar{x})$ over an instance $I$, denoted $q(I)$ is a set of $|q|$- tuples $\bar{c}$ of constants *s.t.*, there is a homomorphism $\mathsf{h}$ from $\mathsf{body}(q)$ to $I$ where $\mathsf{h}(\bar{x}) = \bar{c}$.

A bag or *multiplicity* over a set $I$ is function $\mu\colon I \to \mathbb{N}$. The *answer under bag-bag semantics* of a CQ $q(\bar{x}) = \exists \bar{y} \bigwedge_{i=1}^{n} R_i(\bar{x_i}, \bar{y_i})$ over bag instance $\mu$, is the bag $q_{bb}^{I,\mu}\colon q(I) \to \mathbb{N}$ *s.t.*, for any tuple $\bar{c} \in q(I)$, the multiplicity of $\bar{c}$ is $q_{bb}^{I,\mu}(\bar{c}) = \sum_{\mathsf{h} \in \mathrm{Hom}(q(\bar{x}), I)}^{\mathsf{h}(\bar{x}) = \bar{c}} \prod_i \mu(\mathsf{h}(R_i(\bar{x_i}, \bar{y_i}))))$.

Notice that syntactic repetitions in $\exists \bar{y} \bigwedge_{i=1}^{n} R_i(\bar{x_i}, \bar{y_i})$ correspond to different factors in the product of the multiplicity. The *answer under bag-set semantics* of $q(\bar{x})$ over $I$, denoted $q_{bs}^{I}$, is the same as $q_{bb}^{I,\mu}\colon q(I) \to \mathbb{N}$ but with $\mu$ defined to have range $\{1\}$, simply counting the number of homomorphisms for an answer tuple of a $q$. In this semantics, repetitions of ground atoms do not matter – so we assume no atom repetitions.

Given *CQ*s $q$ and $p$ we say that $q$ is *set contained* in $p$, in symbols $q \sqsubseteq_s p$, if, for all set instances $I$, it holds that $q(I) \subseteq p(I)$. Similarly, we say that $q$ is *bag-bag (bag-set) contained* in $p$, in symbols $q \sqsubseteq_{bb} p$ ($q \sqsubseteq_{bs} p$) , if for all bags $\mu$ over instances $I$, it holds that $q_{bb}^{I,\mu} \subseteq p_{bb}^{I,\mu}$ ($q_{bs}^{I} \subseteq p_{bs}^{I}$), which means $q_{bb}^{I,\mu}(\bar{t}) \leq p_{bb}^{I,\mu}(\bar{t})$ ($q_{bs}^{I}(\bar{t}) \leq p_{bs}^{I}(\bar{t})$), for all $t \in I$. Recall that bag-bag containment implies bag-set containment, which implies set-containment [5].

## 3 Join-on-Free Query Containment

*Join-on-Free Queries* are a large and natural class of *CQ*s, whose join variables are all free, *i.e.*, joins do not involve existential variables.

▶ **Definition 1.** *A* CQ $\exists \overline{y} \bigwedge_{i=1}^{n} R_i(\overline{x}_i, \overline{y}_i)$ *is* join-on-free *(*JoFQ*), if* $\overline{y}_i \cap \overline{y}_j = \emptyset$*, for all* $i, j \in [1, n]$ *with* $i \neq j$.

This class contains queries for which bag-containment has been discussed, but remained unsolved, such as the the classic example from [5].

$$q(x, z) \leftarrow P(x), Q(u, x), Q(v, z), R(z); \quad p(x, z) \leftarrow P(x), Q(u, y), Q(v, y), R(z).$$

Note that $q$ is *JoFQ* since its existential variables do not join. The authors of [5] claimed that $q$ is bag-contained in $p$, but no algorithm has been provided. We give a positive answer to this open problem, by presenting a decidability result for such containment questions, for both bag-bag and bag-set semantics. In fact, bag-bag and bag-set containments are poly-time reducible to each other (claimed in [5], then proved in [16]). To do these reductions one needs to change the queries. Here, we provide a linear-time reduction from the first to the second problem that also maintains the property of the containee query being join-on-free.

▶ **Theorem 2.** *For every pair of* CQ*s* $(q, p)$*, with* $q$ *a* JoFQ*, we can obtain in linear-time a pair of* CQ*s* $(q', p')$*, with* $q'$ *a* JoFQ *as well, such that* $q \sqsubseteq_{\mathtt{bb}} p$ iff $q' \sqsubseteq_{\mathtt{bs}} p'$.

The next theorem states that bag-set containment of *CQ*s reduces to bag-set containment of BCQs. For this, we use *probe tuples* introduced in [24], which are tuples of fresh constants not appearing in the two queries. Note that, there are exponentially many probe tuples simulating all patterns of repetitions of constants in the positions of free variables.

▶ **Theorem 3.** *For every pair of* CQ*s* $(q(\overline{x}), p(\overline{x}))$*, it holds that* $q \sqsubseteq_{\mathtt{bs}} p$ iff $q(\overline{t}) \sqsubseteq_{\mathtt{bs}} p(\overline{t})$*, for all probe tuples* $\overline{t}$ *for* $q$.

Thus, we subsequently focus on the bag-set containment problem for boolean *JoFQ*s (i.e., *JoFBQ*s). We also assume that all constants in the containing query appear in the containee one; otherwise, the queries are not even set contained. Fig. 1 introduces a running example used along this work.

▶ **Problem 1.** *Given a pair of* BCQ*s* $(q, p)$*, where* $q$ *is a* JoFBQ *and* $\mathrm{Cn}(p) \subseteq \mathrm{Cn}(q)$*, the* join-on-free containment problem under bag-set semantics *(*JoF-BCP*) is the problem of deciding whether* $q \sqsubseteq_{\mathtt{bs}} p$ *holds true.*

**Images & Net Images of Atoms.** Essentially, a *JoFBQ* does not contain any joins, as the only common terms among atoms are constants. Due to this, we can show that its bag-set multiplicity, *i.e.*, the number of homomorphisms it has to an instance, can be computed as the product of the number of homomorphisms each different atom in the query has to that instance. A *JoFBQ* might still contain atoms of the same predicate, known as self-joins (note, containment on self-join-free queries is decidable [2, 14]). Thus, two atoms in a *JoFBQ* might have the same image (*i.e.*, atoms are unifiable). Still, we can compute their multiplicity through a monomial function whose parameters correspond to the number of tuples in the image of an atom. As an example, the 140 multiplicity of the query $q$ on the instance $I$ of Fig. 1 can be obtained as the products of the cardinalities $g_1 = 5$, $g_2 = 7$, and $g_3 = 4$ of

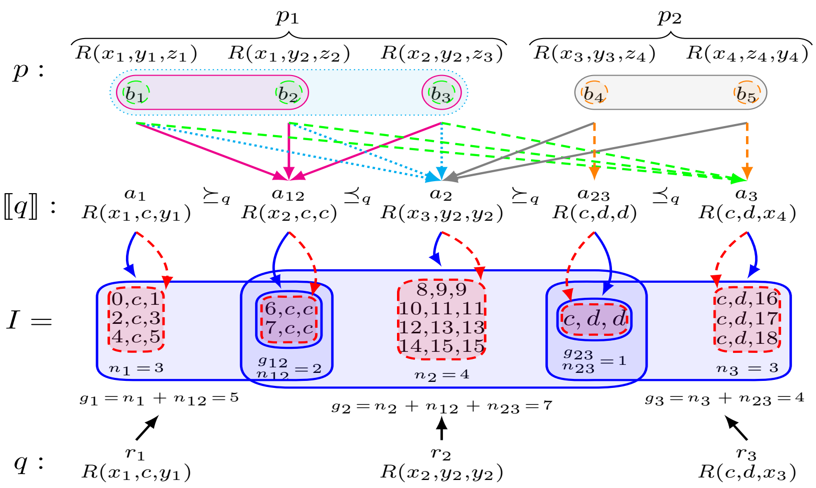

**Figure 1** A three-atom *JoFBQ* $q$, the associated minimal unification closure $[\![q]\!]$ with homomorphic ordering $\preceq_q$ among its five atoms $a_i$, a multicanonical instance $I$ for $q$, and a *BCQ* $p$ obtained as the cross-product of the two *BCQ*s $p_1$ and $p_2$.

the three sets of tuples images (enclosed in the solid blue regions) of its three atoms. Such simple closed expressions are not directly derivable, however, for the arbitrary *BCQ*s $p$ used as containing queries in Prob. 1. To address this, we need a notion of *canonical instance* (see [1]) that precisely reflects the structure of the containee query. In particular, in this context, it has to reflect the closure of the containee query under unification of its atoms, per the following definition.

▶ **Definition 4.** *The* minimal unification closure *of a* JoFBQ $q$*, denoted by* $[\![q]\!]$*, is a minimal set of atoms* w.r.t. *set inclusion such that: (a) for all unifiable sets of atoms* $\mathrm{S} \subseteq q$*, there exists an atom* $a \in [\![q]\!]$ *with* $a \approx_q \mathsf{mgu}(\mathrm{S})$*; (b) no two atoms in* $[\![q]\!]$ *share a variable.*

Observe that the minimal unification closure is unique up to isomorphism. As an example, a possible set $[\![q]\!]$ for the *JoFBQ* $q$ is reported in Fig. 1. It contains the three atoms $a_i$, with $i \in \{1, 2, 3\}$, isomorphic to the atoms $r_i$ in $q$, and the two atoms $a_{12}$ and $a_{23}$, obtained as the unification of the two subqueries $\{r_1, r_2\}$ and $\{r_2, r_3\}$. These are all the unifications one can perform in $q$; any additional atom would make $[\![q]\!]$ non minimal.

Due to the minimality of the closure, every unifiable subset S of $q$ has a unique element in $[\![q]\!]$ isomorphic to $\mathsf{mgu}(\mathrm{S})$. Thus, by $[\![a]\!]_q$, with $a \in q$, we can denote the unique atom in $[\![q]\!]$ such that $[\![a]\!]_q \approx_q \mathsf{mgu}(\{a\}) \approx_q a$. For example, in Fig. 1, we have $[\![r_i]\!]_q = a_i$. For the same reason, $[\![q]\!]$ cannot contain two different isomorphic atoms. Hence, if two atoms $a, \widehat{a} \in [\![q]\!]$ have a common image tuple in an instance $I$, it must be the case that either $\widehat{a} \prec_q a$ or $a \prec_q \widehat{a}$. This allows us to define a decomposition of $I$ into sets of tuples that are exclusively in the image of a single atom from $[\![q]\!]$. We call these sets *net-images*. In order to stress the difference with the sets of image tuples, we sometimes refer to the latter as *gross-images*.

▶ **Definition 5.** *For all* JoFBQ*s* $q$*, atoms* $a \in [\![q]\!]$*, and instances* $I$*:*

- *the* image *of* $a$ *on* $I$ w.r.t. $q$ *is the set* $\mathsf{img}_q(a, I) \triangleq \{b \in I \mid b \preceq_q a\}$*;*
- *the* net-image *of* $a$ *on* $I$ w.r.t. $q$ *is the set* $\mathsf{nimg}_q(a, I) \triangleq \mathsf{img}_q(a, I) \setminus \bigcup_{\widehat{a} \in [\![q]\!]}^{\widehat{a} \prec_q a} \mathsf{img}_q(\widehat{a}, I)$.

In Fig. 1, the solid blue (*resp.*, dashed red) arrows connect the atoms $a_i$ to the corresponding images $\mathsf{img}_q(a_i, I)$ (*resp.*, net-images $\mathsf{nimg}_q(a_i, I)$) in $I$ *w.r.t.* $q$. These sets of tuples are enclosed by the solid blue (*resp.*, dashed red) regions having sizes $g_i$ (*resp.*, $n_i$). For example, the tuple $\langle 0, c, 1\rangle$ is both in the image and net-image of $a_1$, while the tuple $\langle 6, c, c\rangle$ is in the images of $a_1$, $a_2$, and $a_{12}$, but in the net-image of $a_{12}$ only. Similarly, $\langle c, d, d\rangle$ is in the images of $a_2$, $a_3$, and $a_{23}$, but in the net-image of $a_{23}$ only. It should be evident that the image of an atom $a$ in $[\![q]\!]$ is completely partitionable into the net-images of all atoms $\widehat{a}$ that are related to $a$ by the homomorphic ordering $\prec_q$, as prescribed by the following theorem.

▶ **Theorem 6.** *Given an instance $I$, a* JoFBQ *$q$, and an atom $a \in [\![q]\!]$, the set of all subsumed net-images $\left\{\mathsf{nimg}_q(\widehat{a}, I) \,\middle|\, \widehat{a} \in [\![q]\!], \widehat{a} \preceq_q a\right\}$ of $a$ on $I$* w.r.t. *$q$ forms a partition of the corresponding image $\mathsf{img}_q(a, I)$.*

In the next section we will show how to compute the multiplicity of a containing query $p$ by just looking at the sizes of gross- and net-images of each atom in the minimal unification closure $[\![q]\!]$ of the containee query $q$. It is therefore important to characterise the values that these sizes can assume. Obviously, ground atoms have at most one net-image tuple in any instance (*e.g.*, $a_{23}$ in Fig. 1 has only $\langle c, d, d\rangle$ in both its gross- and net-image). Other than this, net-images can have arbitrary sizes. On the contrary, due to Thm. 6, the number of tuples in any gross-image is uniquely dictated by those of the included net-images (*e.g.*, the size $g_2$ of the gross-image of $a_2$ is the sum of the sizes $n_2$, $n_{12}$, and $n_{23}$ of the net-images of $a_2$, $a_{12}$, and $a_{23}$). The following definition formalises these requirements.

▶ **Definition 7** (Net-Image/Image Function)**.** *Given a* JoFBQ *$q$:*

- *a* net-image function *for $q$ is a function $\mathsf{n}\colon [\![q]\!] \to \mathbb{N}$ such that $\mathsf{n}(a) \leq 1$, for every ground atom $a \in [\![q]\!]$;*
- *an* image function *for $q$ is a function $\mathsf{g}\colon [\![q]\!] \to \mathbb{N}$ for which there exists a net-image function $\mathsf{n}$ such that $\mathsf{g}(a) = \sum_{\widehat{a} \in [\![q]\!]}^{\widehat{a} \preceq_q a} \mathsf{n}(\widehat{a})$.*

**Multicanonical Instances.** We can now introduce the notion of *multicanonical (database) instance* of a *JoFQ* as an instance that precisely reflects the structure of its minimal unification closure up to isomorphism and number of tuples. Besides forcing all tuples in the instance to be part of the gross/net-image of an atom of the query, we also minimise joins by avoiding the sharing of non-query-constant values between different tuples.

▶ **Definition 8.** *A* multicanonical instance *for a* JoFBQ *$q$ is an instance $I$ satisfying the following constraints:*

*(1) for all data tuples $d \in I$, there exists an atom $a \in [\![q]\!]$ with $a \approx_q d$;*
*(2) $\mathrm{Tr}(d_1) \cap \mathrm{Tr}(d_2) \subseteq \mathrm{Cn}(q)$, for all data tuples $d_1, d_2 \in I$.*

Clearly, the only essential difference between two different multicanonical instances resides in the number of tuples in the gross/net-images of the atoms. This means that, for each image (*resp.*, net-image) function $\mathsf{g}$ (*resp.*, $\mathsf{n}$) for a *JoFBQ* $q$, there exists a unique up to isomorphism multicanonical instance for $q$ with the sizes of images (*resp.*, net-images) prescribed by $\mathsf{g}$ (*resp.*, $\mathsf{n}$), and *vice versa*.

Instance $I$ of Fig. 1 is an example of a multicanonical instance for the query $q$. In fact, we can transform any instance $I$ into a multicanonical one $I^*_q$ for a *JoFBQ* $q$, still preserving the multiplicity of $q$. The intuition behind this transformation is to generalise the net-image tuples of an atom $a$ in $[\![q]\!]$ until they are isomorphic to $a$ itself, while breaking, at the same time, all joins that do not correspond to query constants. For example, the extension $I \cup \{R(8, c, 8), R(8, c, c)\}$ of $I$ in Fig. 1 can be made multicanonical by changing the two additional tuples into $R(19, c, 20)$ and $R(21, c, c)\}$. The next theorem formalises this idea, by also showing that any other *BCQ* $p$ can only decrease its multiplicity passing from $I$ to $I^*_q$.

▶ **Theorem 9.** *For all* JoFBQ*s $q$ and instances $I$, there exists a multicanonical instance $I_q^*$ for $q$ such that:*

*(a)* $|\text{Hom}(q, I)| = |\text{Hom}(q, I_q^*)| = \prod_{a \in q} |\mathsf{img}_q([\![a]\!]_q, I)|$*;*
*(b)* $|\text{Hom}(p, I)| \geq |\text{Hom}(p, I_q^*)|$*, for all* BCQ*s $p$.*

This result implies that Prob. 1, *i.e.*, the bag-set containment of *JoFBQ*s into *BCQ*s, can be decided by just searching for violations within the realm of multicanonical instances.

▶ **Theorem 10.** *For every pair $(q, p) \in$ JoF-BCP, the following statements are equivalent: (a) $q \not\sqsubseteq_{\mathsf{bs}} p$; (b) there exists a multicanonical instance $I$ for $q$ such that $q_{bs}^I \not\subseteq p_{bs}^I$.*

## 4 Counting Homomorphisms

Thm. 9a states that the multiplicity of a *JoFBQ* $q$ on any instance is given by the product of the sizes of its atom images. That is, given a multicanonical instance $I$ and its image function $\mathsf{g}$, the following monomial function computes the multiplicity of $q$ on $I$:

$$\mathsf{M}_q(\mathsf{g}) \triangleq \prod_{a \in q} \mathsf{g}([\![a]\!]_q).$$

**Counting Through Lenses.** Similar to [24], the idea is to solve a polynomial inequality characterising the containment problem. Specifically, given a multicanonical instance $I$ for the containee *JoFBQ* $q$, we need to compute the multiplicity of any given containing *BCQ* $p$ by constructing a polynomial function for it. Differently to [24], however, the presence of existential variables in $q$ considerably complicates the analysis.

Consider, for example, the multiplicity of the containing subquery $p_1$ of Fig. 1, on just the $g_2 = 7$ tuples in the image of $a_2$ (the middle solid blue region). On this subinstance, $p_1$ has nine homomorphisms: four of them map the three atoms $b_1$, $b_2$, and $b_3$ to the $n_{12} = 2$ tuples in the net-image of $a_{12}$; four additional homomorphisms map these atoms to the $n_2 = 4$ tuples in the net-image of $a_2$; finally, a single homomorphism goes on the tuple in the net-image of $a_{23}$. It should be now clear that, to count these homomorphisms, we cannot just look at the gross-image of $a_2$ (as the approach in [24] would do), without considering its partitioning into the contained net-images. In fact, as elaborated later, the correct count is given by the polynomial function $n_{12}^2 + n_2 + n_{23}^3 = 2^2 + 4 + 1^3 = 9$. Observe however that, in order to compare the polynomial function for the containing query against the monomial function described above for the containee one, we need to express the multiplicity in terms of the sizes of the gross-images. This would mean, for the previous example, to evaluate the following polynomial with negative coefficients: $g_{12}^2 + (g_2 - g_{12} - g_{23}) + g_{23}^3 = 2^2 + (7 - 2 - 1) + 1^3 = 9$.

The example also shows how homomorphisms $\widehat{\mathsf{h}} \in \text{Hom}(p, I)$ from $p$ to the data in $I$ can be clustered based on their shape, *i.e.*, on the homomorphism $\mathsf{h} \in \text{Hom}(p, [\![q]\!])$ from $p$ to $[\![q]\!]$ that, by abstracting from the actual values of the data, identifies the symbolic tuple $\mathsf{h}(b)$ in $[\![q]\!]$ whose net-image contains the actual tuple $\widehat{\mathsf{h}}(b)$, for each atom $b$ in $p$. In fact, $\mathsf{h}$ is a lens on $\text{Hom}(p, I)$ that allows to identify and count all homomorphisms with the same shape.

▶ **Definition 11.** *Let $(q, p) \in$ JoF-BCP. The* lens *of a homomorphism* $\mathsf{h} \in \text{Hom}(p, [\![q]\!])$ *over an instance $I$ is the set of homomorphisms* $\mathrm{L}_{q,p}^{I,\mathsf{h}} \triangleq \left\{ \widehat{\mathsf{h}} \in \text{Hom}(p, I) \;\middle|\; \forall a \in p.\, \widehat{\mathsf{h}}(a) \in \mathsf{nimg}_q(\mathsf{h}(a), I) \right\}$.

Consider in Fig. 1 the homomorphism, say $\mathsf{h}_1$, from $p$ to $[\![q]\!]$ that maps the subquery $p_1$ onto the atom $a_{12}$ (solid magenta arrows), and the subquery $p_2$ onto the atoms $a_{23}$ and $a_3$ (dashed orange arrows). The lens for this homomorphism contains all six homomorphisms that map $b_1, b_2$, $b_3$ onto the net-image of $a_{12}$ and $b_4$ and $b_5$ onto the net-images of $a_{23}$ and

$a_3$. Dually, every homomorphism from $p$ to $I$ belongs to some lens. *E.g.*, the homomorphism $\{x_1, x_2 \mapsto c; y_1, y_2 \mapsto d; z_1, z_2, z_3 \mapsto 16; x_3, x_4 \mapsto 8; y_3, y_4, z_4 \mapsto 9\}$ is in the lens of the homomorphism, say $\mathsf{h}_2$, from $p$ to $[\![q]\!]$ mapping the subquery $p_1$ onto $a_3$ (dashed green arrows) and the subquery $p_2$ onto $a_2$ (solid grey arrows). Thus, the lenses cover the set $\mathrm{Hom}(p, I)$; even more, they fully partition it, as stated in the following.

▶ **Theorem 12.** *For each pair* $(q, p) \in \text{JoF-BCP}$ *and multicanonical instance* $I$ *for* $q$*, the set of lenses* $\{\mathrm{L}^{I,\mathsf{h}}_{q,p} \mid \mathsf{h} \in \mathrm{Hom}(p, [\![q]\!])\}$ *over* $I$ *is a partition of the set of homomorphisms* $\mathrm{Hom}(p, I)$.

Clearly, Thm. 12 lets us compute the multiplicity of the containing query $p$ by summing up the sizes of all lenses $\mathrm{L}^{I,\mathsf{h}}_{q,p}$, with $\mathsf{h} \in \mathrm{Hom}(p, [\![q]\!])$. Towards estimating these sizes, first note that each atom $a = \mathsf{h}(b)$, with $b \in p$, gives rise to at least $\mathsf{n}(a)$ homomorphisms in the lens. For example, in Fig. 1, atom $a_2$ has $n_2 = 4$ net-image tuples in $I$. This forces the lens of the homomorphisms from $p_1$ onto $a_2$ (dotted blue arrows), say $\mathsf{h}_3$, to contain at least four different homomorphisms from $p_1$ to $I$, each mapping one of the three atoms $b_i$ to one of the four tuples in $\mathsf{nimg}_q(a_2, I)$. Obviously, we can have at most $4^3$ such homomorphisms. However, not all atoms may map independently onto these tuples. For this particular lens, in fact, all atoms of $p_1$ have to map to the exact same tuple, giving rise just to four homomorphisms. In more detail, if $b_1$ maps to $\langle 8, 9, 9\rangle$, then, due to join on $x_1 \mapsto 8$, atom $b_2$ maps onto the same tuple, which in turns forces $b_3$ to do the same, due to join on $y_2 \mapsto 9$. Therefore, $|\mathrm{L}^{I,\mathsf{h}_3}_{q,p_1}| = n_2$. The fact that values 8 and 9 are unique to the first tuple in the net-image is no coincidence, since they correspond to existential variables in $a_2$, which always map to unique values in a multicanonical instance. For two joined atoms of $p$ to differentiate in their image through $\mathsf{h}$, they have to join only on constants of $q$. Thus, two atoms of $p$ containing a join variable mapped by $\mathsf{h}$ to a variable of $[\![q]\!]$ do not contribute to the total size of its lens in a multiplicative way, but rather they are equivalent *w.r.t.* $\mathsf{h}$, in a sense formalised as follows.

▶ **Definition 13.** *Given a pair* $(q, p) \in \text{JoF-BCP}$*, two atoms* $b_1, b_2 \in p$ *are* existentially joined through *a homomorphism* $\mathsf{h} \in \mathrm{Hom}(p, [\![q]\!])$*, in symbols* $b_1 \bowtie_{\mathsf{h}} b_2$*, if there exists a variable* $x \in \mathrm{Vr}(b_1) \cap \mathrm{Vr}(b_2)$ *such that* $\mathsf{h}(x) \notin \mathrm{Cn}(q)$*. The equivalence relation obtained as transitive closure of the reflexive and symmetric relation* $\bowtie_{\mathsf{h}}$ *is denoted by* $\equiv_{\mathsf{h}}$.

Intuitively, equivalent atoms *w.r.t.* a homomorphism $\mathsf{h}$ essentially act as single atom, since their image coincide in every homomorphism in the lens of $\mathsf{h}$. On the contrary, all different equivalent classes *w.r.t.* $\equiv_{\mathsf{h}}$ act as independent atoms. Consider again the homomorphism $\mathsf{h}_3$ discussed above. Clearly, both $b_1 \bowtie_{\mathsf{h}_3} b_2$ and $b_2 \bowtie_{\mathsf{h}_3} b_3$, thus $b_1 \equiv_{\mathsf{h}_3} b_2 \equiv_{\mathsf{h}_3} b_3$. This is graphically represented in Fig. 1 by the dotted cyan area grouping $b_1$, $b_2$, and $b_3$. Consider now the homomorphisms, say $\mathsf{h}_4$, from $p_1$ onto $a_{12}$ (solid magenta arrows). In this case, we have $b_1 \bowtie_{\mathsf{h}_4} b_2$, but neither $b_1 \bowtie_{\mathsf{h}_4} b_3$ nor $b_2 \bowtie_{\mathsf{h}_4} b_3$. Thus, $b_1 \equiv_{\mathsf{h}_4} b_2$ is the only possible equivalence, as represented by the two solid magenta areas containing $b_1$ and $b_2$, on one side, and $b_3$ alone, on the other one. As a consequence of this equivalence, all homomorphisms in the lens of $\mathsf{h}_4$ map $b_1$ and $b_2$ to the same tuple in $\mathsf{nimg}_q(a_{12}, I)$, still allowing $b_3$ to be mapped to a different tuple. Therefore, $|\mathrm{L}^{I,\mathsf{h}_4}_{q,p_1}| = n_{12}^2$.

Given a homomorphism $\mathsf{h} \in \mathrm{Hom}(p, [\![q]\!])$ and an atom $a \in [\![q]\!]$, let $\mathsf{h}^{-1}(a) = \{b \in p \mid \mathsf{h}(b) = a\}$ be the inverse image of $a$. The number $|\mathsf{h}^{-1}(a)/\equiv_{\mathsf{h}}|$ of equivalence classes *w.r.t.* $\mathsf{h}$ and $a$ over that set is denoted by $\eta_p(\mathsf{h}, a)$.

As an example, for the homomorphisms $\mathsf{h}_1$ and $\mathsf{h}_2$ from $p$ to $[\![q]\!]$ previously discussed, we have $\eta_p(\mathsf{h}_1, a_{12}) = 2$ and $\eta_p(\mathsf{h}_1, a_{23}) = \eta_p(\mathsf{h}_1, a_3) = 1$, while $\eta_p(\mathsf{h}_2, a_2) = 1$ and $\eta_p(\mathsf{h}_2, a_3) = 3$. For this reason, the sizes of the two corresponding lenses $\mathrm{L}^{I,\mathsf{h}_1}_{q,p}$ and $\mathrm{L}^{I,\mathsf{h}_2}_{q,p}$ can be computed via the monomial functions $n_{12}^2 \cdot n_{23}^1 \cdot n_3^1 = 2^2 \cdot 1^1 \cdot 3^1 = 12$ and $n_2^1 \cdot n_3^3 = 4^1 \cdot 3^3 = 36$, as prescribed by the next theorem.

▶ **Theorem 14.** *For all pairs* $(q, p) \in$ JoF-BCP*, multicanonical instance* $I$ *for* $q$*, homomorphism* $\mathsf{h} \in \mathrm{Hom}(p, [\![q]\!])$*, and lens* $\mathrm{L}^{I,\mathsf{h}}_{q,p}$ *of* $\mathsf{h}$ *over* $I$*, it holds* $|\mathrm{L}^{I,\mathsf{h}}_{q,p}| = \prod_{a \in \mathsf{h}(p)} |\mathsf{nimg}_q(a, I)|^{\eta_p(\mathsf{h},a)}$.

**Polynomial-Function Characterisation.** As an immediate consequence of Thms. 12 and 14, the multiplicity of the containing query $p$ can be computed as a polynomial function taking as parameter the net-image function $\mathsf{n}$ of the multicanonical instance $I$ for the containee query $q$ as follow:

$$\mathsf{N}^p_q(\mathsf{n}) \triangleq \sum_{\mathsf{h} \in \mathrm{Hom}(p,[\![q]\!])} \prod_{a \in \mathsf{h}(p)} \mathsf{n}(a)^{\eta_p(\mathsf{h},a)}.$$

*E.g.*, for the two subqueries $p_1$ and $p_2$ of Fig. 1, after some grouping, we have $\mathsf{N}^{p_1}_q(\mathsf{n}) = (\mathsf{n}(a_1) + \mathsf{n}(a_{12}))^2 + \mathsf{n}(a_2) + (\mathsf{n}(a_{23}) + \mathsf{n}(a_3))^3$ and $\mathsf{N}^{p_2}_q(\mathsf{n}) = (\mathsf{n}(a_1) + \mathsf{n}(a_{12})) \cdot \mathsf{n}(a_{12}) + \mathsf{n}(a_2) + \mathsf{n}(a_{23}) \cdot (\mathsf{n}(a_{23}) + \mathsf{n}(a_3))$.

However, the monomial function $\mathsf{M}_q(\mathsf{g})$ previously introduced for $q$ takes as parameter the gross-image function $\mathsf{g}$ of the multicanonical instance $I$. Thus, to properly compare the two multiplicities, we need to obtain a polynomial function for $p$ parameterised on $\mathsf{g}$ as well. Fortunately, Def. 7 ensures that every gross-image function $\mathsf{g}$ has an associated net-image function $\mathsf{n}_\mathsf{g}$ recursively defined as follows: $\mathsf{n}_\mathsf{g}(a) = \mathsf{g}(a) - \sum^{\widehat{a} \preceq_q a}_{\widehat{a} \in [\![q]\!]} \mathsf{n}_\mathsf{g}(\widehat{a})$. The expert reader might here note that this transformation is essentially an application of the Möbius inversion formula on posets [31, 33]. This allows to obtain the following polynomial function for $p$:

$$\mathsf{P}^p_q(\mathsf{g}) \triangleq \mathsf{N}^p_q(\mathsf{n}_\mathsf{g}) = \sum_{\mathsf{h} \in \mathrm{Hom}(p,[\![q]\!])} \prod_{a \in \mathsf{h}(p)} \mathsf{n}_\mathsf{g}(a)^{\eta_p(\mathsf{h},a)}.$$

*E.g.*, for the two subqueries $p_1$ and $p_2$ of Fig. 1, after exploiting the equalities $\mathsf{n}_\mathsf{g}(a_1) = \mathsf{g}(a_1) - \mathsf{g}(a_{12})$, $\mathsf{n}_\mathsf{g}(a_{12}) = \mathsf{g}(a_{12})$, $\mathsf{n}_\mathsf{g}(a_2) = \mathsf{g}(a_2) - \mathsf{g}(a_{12}) - \mathsf{g}(a_{23})$, $\mathsf{n}_\mathsf{g}(a_{23}) = \mathsf{g}(a_{23})$, and $\mathsf{n}_\mathsf{g}(a_3) = \mathsf{g}(a_3) - \mathsf{g}(a_{23})$, we obtain $\mathsf{P}^{p_1}_q(\mathsf{g}) = \mathsf{g}(a_1)^2 + (\mathsf{g}(a_2) - \mathsf{g}(a_{12}) - \mathsf{g}(a_{23})) + \mathsf{g}(a_3)^3$ and $\mathsf{P}^{p_2}_q(\mathsf{g}) = \mathsf{g}(a_1) \cdot \mathsf{g}(a_{12}) + (\mathsf{g}(a_2) - \mathsf{g}(a_{12}) - \mathsf{g}(a_{23})) + \mathsf{g}(a_{23}) \cdot \mathsf{g}(a_3)$.

Thus, thanks to Thm. 10 and the characterisations of the multiplicities of the two queries, a solution for Prob. 1 reduces to the following non-existence of an image function.

▶ **Theorem 15.** *For every pair* $(q, p) \in$ JoF-BCP*, the following statements are equivalent: (a)* $q \not\sqsubseteq_{\mathsf{bs}} p$*; (b) there exists an image function* $\mathsf{g}$ *for* $q$ *such that* $\mathsf{M}_q(\mathsf{g}) > \mathsf{P}^p_q(\mathsf{g})$.

**Monomial-Polynomial Characterisation.** Modelling the values in the range of an image function $\mathsf{g}$ as unknowns $\overline{\boldsymbol{u}}$, allow us to introduce the symbolic monomial $\mathsf{M}_q(\overline{\boldsymbol{u}})$ and polynomial $\mathsf{P}^p_q(\overline{\boldsymbol{u}})$ for the functions $\mathsf{M}_q(\mathsf{g})$ and $\mathsf{P}^p_q(\mathsf{g})$. Note that $\mathsf{P}^p_q(\mathsf{g})$ does not have a polynomial form, since products and summations are nested inside each other due to the recursive definition of the net-image function $\mathsf{n}_\mathsf{g}$. Nevertheless, there is always an equivalent polynomial form, which can be obtained via standard syntactic manipulations.

Thm. 16 below links a negative answer to Prob. 1 with a solution of an inequality between $\mathsf{M}_q(\overline{\boldsymbol{u}})$ and $\mathsf{P}^p_q(\overline{\boldsymbol{u}})$. Intuitively, the values of this solution correspond to the values of an image function that induces an instance disproving containment. To ensure such a correspondence, we need to verify the constraints specified in Def. 7. We consider some linear constraints, one per unknown, based on the following symbolic version of the formula used to compute the net-image function $\mathsf{n}_\mathsf{g}$ from the image function $\mathsf{g}$: $\mathsf{C}^a_q(\overline{\boldsymbol{u}}) \triangleq u_a - \sum^{\widehat{a} \prec_q a}_{\widehat{a} \in [\![q]\!]} \mathsf{C}^{\widehat{a}}_q(\overline{\boldsymbol{u}})$. Note that, $\mathsf{C}^a_q(\overline{\boldsymbol{u}}) = u_a$, for $a \in \mathsf{min}_{\prec_q}[\![q]\!]$.

▶ **Theorem 16.** *For every pair* $(q, p) \in$ JoF-BCP*, the following statements are equivalent: (a)* $q \not\sqsubseteq_{\mathsf{bs}} p$*; (b) there is a Diophantine solution of* $\mathsf{M}_q(\overline{\boldsymbol{u}}) > \mathsf{P}^p_q(\overline{\boldsymbol{u}})$ *that satisfies the inequalities* $\{\mathsf{C}^a_q(\overline{\boldsymbol{u}}) \geq 0\}_{a \in [\![q]\!]}$.

Before turning to the algorithmic solution of the Diophantine problem identified in the above theorem, let us consider for a final time the two queries $q$ and $p$ reported in Fig. 1. The corresponding Diophantine inequality $\mathsf{M}_q(\overline{\boldsymbol{u}}) > \mathsf{P}_q^p(\overline{\boldsymbol{u}})$, with $\mathsf{P}_q^p(\overline{\boldsymbol{u}}) = \mathsf{P}_q^{p_1}(\overline{\boldsymbol{u}}) \cdot \mathsf{P}_q^{p_2}(\overline{\boldsymbol{u}})$, is

$$u_{a_1} u_{a_2} u_{a_3} > (u_{a_1}^2 + u_{a_2} - u_{a_{12}} - u_{a_{23}} + u_{a_3}^3)(u_{a_1} u_{a_{12}} + u_{a_2} - u_{a_{12}} - u_{a_{23}} + u_{a_{23}} u_{a_3}).$$

Note that the monomial on the left hand side of our inequality only uses a subset of the unknowns, namely, the gross image cardinalities that correspond to atoms of the minimal unification closure. The associated system $\{\mathsf{C}_q^a(\overline{\boldsymbol{u}}) \geq 0\}_{a \in [\![q]\!]}$ of linear inequalities induced by the homomorphic ordering $\preceq_q$ over the set of atoms $[\![q]\!] = \{a_1, a_{12}, a_2, a_{23}, a_3\}$ is, instead,

$$\mathsf{C}_q^{a_1}(\overline{\boldsymbol{u}}) = u_{a_1} - u_{a_{12}} \geq 0, \qquad \mathsf{C}_q^{a_{12}}(\overline{\boldsymbol{u}}) = u_{a_{12}} \geq 0,$$
$$\mathsf{C}_q^{a_2}(\overline{\boldsymbol{u}}) = u_{a_2} - u_{a_{12}} - u_{a_{23}} \geq 0, \qquad \mathsf{C}_q^{a_{23}}(\overline{\boldsymbol{u}}) = u_{a_{23}} \geq 0, \qquad \mathsf{C}_q^{a_3}(\overline{\boldsymbol{u}}) = u_{a_3} - u_{a_{23}} \geq 0.$$

## 5 A Diophantine Problem

In the preceding sections, we established a connection between the bag-set containment problem of *JoFBQ*s into *BCQ*s and the unsolvability of a specific Diophantine inequality, when conjoined with a set of linear constraints. Here we perform a broader analysis of the structural underpinnings of this mathematical encoding, which will provide the sought algorithmic solution to the entire question. It is important to recall that deciding the existence of a Diophantine solution for a polynomial inequality is, in general, impossible [14], due to its tight relationship with Hilbert's 10th problem [29]. Therefore, we need to identify sufficient conditions under which decidability can be achieved. The technical development of this work generalises the approach proposed in [24]. Thus, we employ a similar notation.

Given two $n$-vectors of *unknowns* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ and natural numbers $\overline{\boldsymbol{e}} = \{e_u\}_{u \in \overline{\boldsymbol{u}}} \in \mathbb{N}^n$, for some $n \in \mathbb{N}_+$, we denote by $\overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}}$ the *unitary monomial* $\prod_{u \in \overline{\boldsymbol{u}}} u^{e_u}$ in the polynomial ring $\mathbb{Z}[\overline{\boldsymbol{u}}]$. A *Monomial-Polynomial Inequality* (*MPI*) is an expression of the form $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$, where $\mathsf{M}(\overline{\boldsymbol{u}}) = \overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}} \in \mathbb{Z}[\overline{\boldsymbol{u}}]$ is a unitary monomial and $\mathsf{P}(\overline{\boldsymbol{u}}) = \sum_{i=1}^{m} \alpha_i \overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}_i} \in \mathbb{Z}[\overline{\boldsymbol{u}}]$ is a *polynomial* with coefficients $\{\alpha_i\}_{i=1}^{m} \subseteq \mathbb{Z} \setminus \{0\}$, for some $m \in \mathbb{N}_+$. A *Generalised* MPI (*GMPI*) allows for non-negative real exponent vectors $\{\overline{\boldsymbol{e}}\} \cup \{\overline{\boldsymbol{e}}_i\}_{i=1}^{m} \subseteq \mathbb{R}_{\geq 0}^n$.

As already observed in the last example of the previous section, the linear constraints on the sizes of images given in Thm. 16 are unambiguously induced by the homomorphic ordering $\preceq_q$ of the atoms contained in the minimal unification closure $[\![q]\!]$ of the containee *JoFBQ* $q$. Here, we generalise this concept. Given a *partially ordered set* (*poset*) $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ on the $n$-unknowns $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$, we introduce the following recursive definition of a *distinguished set* $\{\mathsf{C}_{\mathfrak{P}}^u(\overline{\boldsymbol{u}}) \in \mathbb{Z}[\overline{\boldsymbol{u}}]\}_{u \in \overline{\boldsymbol{u}}}$ *of linear polynomials*, one per unknown in $\overline{\boldsymbol{u}}$: $\mathsf{C}_{\mathfrak{P}}^u(\overline{\boldsymbol{u}}) \triangleq u - \sum_{u' \in \overline{\boldsymbol{u}}}^{u' \prec u} \mathsf{C}_{\mathfrak{P}}^{u'}(\overline{\boldsymbol{u}})$. Note that $\mathsf{C}_{\mathfrak{P}}^u(\overline{\boldsymbol{u}}) = u$, when $u \in \mathsf{min}_{\preceq} \overline{\boldsymbol{u}}$. Again, this is just another application of a Möbius inversion formula. The linear polynomials just introduced represent the number of net-image tuples discussed earlier, once the unknowns are substituted with the count of images attributed to the respective atom.

To accurately represent the sizes of images, the linear polynomials must yield non-negative values. To solve an *MPI*, we look for vectors of values, called $\mathfrak{P}$-*coherent* values, adhering to these constraints.

▶ **Definition 17.** *An $n$-vector of values $\overline{\boldsymbol{\xi}} = \{\xi_u\}_{u \in \overline{\boldsymbol{u}}} \in \mathbb{N}^n$ is* $\mathfrak{P}$-coherent*, given a* poset $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *on $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$, if it satisfies the system of linear polynomial inequalities $\{\mathsf{C}_{\mathfrak{P}}^u(\overline{\boldsymbol{u}}) \geq 0\}_{u \in \overline{\boldsymbol{u}}}$.*

To tame the undecidability of the general case, we focus on polynomials with a particular trait: they maintain non-negativity when evaluated on coherent vectors, even upon the removal of the unitary form of one of their positive monomials. We will later show that the polynomial $\mathsf{P}_q^p(\overline{\boldsymbol{u}})$, for any *BCQ* $p$ and *JoFBQ* $q$, enjoys this desirable property.

▶ **Definition 18.** *A polynomial* $\mathsf{P}(\overline{\boldsymbol{u}}) = \sum_{i=1}^{m} \alpha_i \overline{\boldsymbol{u}}^{\overline{e}_i} \in \mathbb{Z}[\overline{\boldsymbol{u}}]$ *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ *is* strongly non-negative *w.r.t. a* poset $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *whenever (a)* $\alpha_i \geq 1$*, for some* $i \in [m]$*, and (b) if* $\alpha_i \geq 1$ *then the associated polynomial* $\mathsf{P}(\overline{\boldsymbol{u}}) - \overline{\boldsymbol{u}}^{\overline{e}_i}$ *is non-negative when evaluated on* $\mathfrak{P}$*-coherent vectors in* $\mathbb{N}^n$*, for all* $i \in [m]$*.*

Clearly, the homomorphic ordering $\preceq_q$ applied to the atoms within $[\![q]\!]$ induces a *poset*. On a closer examination, it becomes apparent that this structure closely resembles a meet semilattice, due to its defining property: when considering any two atoms that unify, their most general unification is also an element of $[\![q]\!]$. Nevertheless, it is important to acknowledge the potential presence of non-unifiable atoms. Thus, in order to complete this structure, making it a meet semilattice, we have to include a distinct minimal element. A 0*-completable meet semilattice (*0-cmsem*)* $\mathfrak{P} = \langle \mathrm{U}, \preceq \rangle$ is a *poset* such that, for every subset $\mathrm{V} \subseteq \mathrm{U}$, there exists an element in $\{0\} \cup \mathrm{U}$, called *greatest lower bound* of V, *a.k.a. infimum* or *meet*, usually denoted by $\bigwedge \mathrm{V}$, such that $\bigwedge \mathrm{V} \preceq_0 u$, for all $u \in \mathrm{V}$, where $\preceq_0$ is the extension of the order relation $\preceq$ having 0 as smallest element. We can now state the Diophantine decision problem.

▶ **Problem 2.** *Let* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$ *be an* MPI *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ *and* $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *a* 0-cmsem*, where* $\mathsf{P}(\overline{\boldsymbol{u}})$ *is strongly non-negative* w.r.t. $\mathfrak{P}$*. Is there a* $\mathfrak{P}$*-coherent Diophantine solution for the* MPI*?*

Consider the *MPI* $\mathsf{M}_q(\overline{\boldsymbol{u}}) > \mathsf{P}_q^p(\overline{\boldsymbol{u}})$, the linear inequality system $\{\mathsf{C}_q^a(\overline{\boldsymbol{u}}) \geq 0\}_{a \in [\![q]\!]}$, and the 0-*cmsem* $\mathfrak{P} = \langle [\![q]\!], \preceq_q \rangle$, with $[\![q]\!] = \{a_1, a_{12}, a_2, a_{23}, a_3\}$, discussed after Thm. 16 in the example related to Fig. 1. The existence of a simultaneous solution of the *MPI* and the linear system would result in a positive answer to the associated instance of Prob. 2 and *vice versa*.

In line with the methodology presented in [24], the solution of the introduced decision problem involves a reduction to and a resolution of a homogeneous (*i.e.*, without constant terms) linear system of inequalities. Yet, unlike the direct reduction employed there, here we do this in two different phases, in order to deal with the additional linear constraints. Initially, our focus lies in solving a more specific problem (see Prob. 3 below), seeking solutions that adhere to (i) being positive, and (ii) having distinct values for any neighbouring pair of unknowns within the given *poset*. Subsequently, we reduce Prob. 2 to potentially several instances of Prob. 3, facilitating a systematic handling of these constraints.

**Strict Solutions.** To formalise the restrictions on the solutions discussed above, first let us introduce some notation. For a *poset* $\mathfrak{P} = \langle \mathrm{U}, \preceq \rangle$ on an arbitrary set U, we denote by $\lhd$ the *(immediate) predecessor relation* compatible with $\preceq$, *i.e.*, $u_1 \lhd u_2$ *iff* $u_1 \prec u_2$ and there is no $u \in \mathrm{U}$ such that $u_1 \prec u \prec u_2$. We can now strengthen the coherence property.

▶ **Definition 19.** *An* $n$*-vector of values* $\overline{\boldsymbol{\xi}} = \{\xi_u\}_{u \in \overline{\boldsymbol{u}}} \in \mathbb{N}^n$ *is* $\mathfrak{P}$-strict*, given a* poset $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$*, if it is* $\mathfrak{P}$*-coherent and satisfies the inequalities* $\{u > 0\}_{u \in \min_{\preceq} \overline{\boldsymbol{u}}} \cup \{u_1 < u_2\}_{u_1, u_2 \in \overline{\boldsymbol{u}}}^{u_1 \lhd u_2}$*.*

Thanks to the above restriction on the space of solutions, we can formally state the second Diophantine decision problem.

▶ **Problem 3.** *Let* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$ *be an* MPI *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ *and* $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *a* poset*, where* $\mathsf{P}(\overline{\boldsymbol{u}})$ *is strongly non-negative* w.r.t. $\mathfrak{P}$*. Is there a* $\mathfrak{P}$*-strict Diophantine solution for the* MPI*?*

Consider again the example after Thm. 16. The existence of a simultaneous solution of the *MPI*, the linear system, and the set of inequalities $\{u_{a_1} > u_{a_{12}}, u_{a_2} > u_{a_{12}}, u_{a_2} >$

$u_{a_{23}}, u_{a_3} > u_{a_{23}}, u_{a_{12}} > 0, u_{a_{23}} > 0\}$ would result in a positive answer to the associated instance of Prob. 3 and *vice versa*.

Clearly, the existence of a solution for an *MPI* cannot generally be inferred by evaluating the relative order of the degrees of its monomial and polynomial. However, by examining *MPI*s of a single unknown, one may observe a strong correlation between this order and the presence of a solution. This observation, initially employed in [24] for polynomials with only positive coefficients, is extended here to accommodate the broader spectrum of strictly non-negative polynomials.

In a one-variable *GMPI*, when the *degree* $\mathsf{deg}(\mathsf{M}(u))$ of the monomial $\mathsf{M}(u)$ is strictly greater than the *positive degree* $\mathsf{deg}_+(\mathsf{P}(u))$ of the polynomial $\mathsf{P}(u)$ (*i.e.*, the maximum degree of its monomials with positive coefficients), a solution must exist. In fact, there is an infinite continuous range of potential solutions.

▶ **Lemma 20.** *For each* GMPI $\mathsf{M}(v) > \mathsf{P}(v)$ *on* $v \in \mathrm{Un}$*, with* $\mathsf{deg}(\mathsf{M}(v)) > \mathsf{deg}_+(\mathsf{P}(v))$*, there is a value* $\ell \in \mathbb{R}_{\geq 1}$ *such that every other value* $\xi \in \mathbb{R}$ *with* $\xi > \ell$ *is a solution of the* GMPI.

The converse is not universally valid, since a solution might exist even when the monomial has a strictly lower degree than the polynomial. However, by demanding the latter to remain non-negative after the removal of the unitary form of the leading monomial, such a converse correlation can be also established. The latter property is intimately connected with the strongly non-negativeness of the multi-dimensional polynomial.

▶ **Lemma 21.** *If a* GMPI $\mathsf{M}(v) > \mathsf{P}(v)$ *on* $v \in \mathrm{Un}$ *admits a solution in* $\mathbb{R}_{\geq 1}$ *on which the polynomial* $\mathsf{P}(v) - v^{\mathsf{deg}_+(\mathsf{P}(v))}$ *is non-negative, then* $\mathsf{deg}(\mathsf{M}(v)) > \mathsf{deg}_+(\mathsf{P}(v))$*.*

The core idea behind the technique of [24] is to initially correlate the solutions $\overline{\boldsymbol{\xi}}$ of a multi-dimensional *MPI* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$ with those $\xi$ of a related one-dimensional *GMPI* $\mathsf{M}^*(v) > \mathsf{P}^*(v)$. This correlation is then used together with the established relationship between solutions and degrees, to reduce the problem into solving a homogeneous linear system whose coefficients are the exponents of the original monomial and polynomial. Expanding upon this idea, we incorporate the additional linear constraints imposed on the unknowns. The additional strictness requirements play here a pivotal role, enabling the application of the approach in the current more complex setting.

Let us consider an *MPI* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$ on the $n$-unknowns $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$, where $\mathsf{M}(\overline{\boldsymbol{u}}) = \overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}} \in \mathbb{Z}[\overline{\boldsymbol{u}}]$ and $\mathsf{P}(\overline{\boldsymbol{u}}) = \sum_{i=1}^{m} \alpha_i \overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}_i} \in \mathbb{Z}[\overline{\boldsymbol{u}}]$. Moreover, let $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ be a *poset* on $\overline{\boldsymbol{u}}$. We can then introduce the associated $\langle \mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}}), \mathfrak{P} \rangle$*-system of (homogeneous linear) inequalities* defined as follows:

$$\left\{ (\overline{\boldsymbol{e}} - \overline{\boldsymbol{e}}_i)^{\mathsf{T}} \cdot \overline{\boldsymbol{u}} > 0 \right\}_{i \in [m]}^{\alpha_i \geq 1} \cup \left\{ u > 0 \right\}_{u \in \mathsf{min}_{\preceq} \overline{\boldsymbol{u}}} \cup \left\{ u_1 < u_2 \right\}_{u_1, u_2 \in \overline{\boldsymbol{u}}}^{u_1 \lhd u_2}.$$

It is worth noting that the coherence requirements on the solutions of the *MPI* do not directly manifest in the structure of the linear system. Indeed, the first part deals with the relative order between the degrees, while the other two components only enforce the strictness restrictions. Nevertheless, we show that any solution satisfying the inequalities $\{u > 0\}_{u \in \mathsf{min}_{\preceq} \overline{\boldsymbol{u}}} \cup \{u_1 < u_2\}_{u_1, u_2 \in \overline{\boldsymbol{u}}}^{u_1 \lhd u_2}$ always leads to one satisfying the coherence requirements.

Before presenting the theorem stating the correctness and completeness of the proposed solution approach, we need to establish a polynomial bound on the representation size of the potential solution of the linear system, which can be derived from well-known results from the literature [32, *e.g.*, Corollary 17.1c]. Such a bound is then exploited in the decision procedure provided at the end of this section.

An $n$-vector of values $\overline{\boldsymbol{d}} = \{d_u\}_{u \in \overline{\boldsymbol{u}}} \in \mathbb{N}^n$ is *$s$-bit-bounded*, for some number $s \in \mathbb{N}$, if $\mathsf{size}(\overline{\boldsymbol{d}}) \triangleq \sum_{u \in \overline{\boldsymbol{u}}} (1 + \lceil \log_2(d_u + 1) \rceil) \leq s$. Also, recall that the *facet complexity* of a linear system is the maximum length of the binary encoding of its inequalities, which corresponds to the length of the binary representations of all rational numbers appearing as coefficients.

▶ **Lemma 22.** *Let* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$ *be an* MPI *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ *and* $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *a* poset*. The* $\langle \mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}}), \mathfrak{P} \rangle$*-system of inequalities admits a solution* iff *it admits a* $(6\phi n^3)$*-bit-bounded Diophantine solution, with* $\phi$ *its facet complexity.*

Summing up, we can formalise the aforementioned reduction of Prob. 3 as follows.

▶ **Theorem 23.** *Let* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$ *be an* MPI *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ *and* $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *a* poset*. If the* $\langle \mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}}), \mathfrak{P} \rangle$*-system of inequalities admits a solution then the* MPI *admits a* $\mathfrak{P}$*-strict Diophantine solution. Conversely, under the assumption that* $\mathsf{P}(\overline{\boldsymbol{u}})$ *is strongly non-negative* w.r.t. $\mathfrak{P}$*, if the* MPI *admits a* $\mathfrak{P}$*-strict Diophantine solution then the* $\langle \mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}}), \mathfrak{P} \rangle$*-system of inequalities admits a* $(6\phi n^3)$*-sized Diophantine solution, with* $\phi$ *its facet complexity.*

The high-level idea behind the proof of the above result is as follows. If the $\langle \mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}}), \mathfrak{P} \rangle$-system of inequalities admits a solution, then, by Lemma 22, there exists an integer solution $\overline{\boldsymbol{d}} \in \mathbb{N}^n$. From this, we build a *GMPI* $\mathsf{M}^*(v) > \mathsf{P}^*(v)$ on a single unknown $v$, using the same coefficients as the original *MPI* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$ and exponents derived from $\overline{\boldsymbol{d}}$, such that $\mathsf{deg}(\mathsf{M}^*(v)) > \mathsf{deg}_+(\mathsf{P}^*(v))$. Now, by Lemma 20, there exists a value $\ell \in \mathbb{R}_{\geq 1}$ such that every $\xi \in \mathbb{R}$ with $\xi > \ell$ is a solution of this *GMPI*. From each such $\xi$, we then construct an $n$-dimensional vector $\overline{\boldsymbol{\xi}} \in \mathbb{N}_+^n$ satisfying $\mathsf{M}(\overline{\boldsymbol{\xi}}) = \mathsf{M}^*(\xi) > \mathsf{P}^*(\xi) = \mathsf{P}(\overline{\boldsymbol{\xi}})$. In other words, $\overline{\boldsymbol{\xi}}$ is a solution. Since $\xi$ can be chosen arbitrarily large, we can guarantee the existence of a vector $\overline{\boldsymbol{\xi}}$ (actually, infinitely many) that is $\mathfrak{P}$-strict. For the converse direction, the reasoning is analogous. Starting from a $\mathfrak{P}$-strict solution $\overline{\boldsymbol{\xi}} \in \mathbb{N}_+^n$ of the *MPI* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$, we build a one-dimensional *GMPI* $\mathsf{M}^*(v) > \mathsf{P}^*(v)$, again using the same coefficients as the original *MPI* and exponents derived from $\overline{\boldsymbol{\xi}}$ (which may not necessarily be integers). For this *GMPI*, there exists a value $\xi \in \mathbb{R}$ such that $\mathsf{M}^*(\xi) = \mathsf{M}(\overline{\boldsymbol{\xi}}) > \mathsf{P}(\overline{\boldsymbol{\xi}}) = \mathsf{P}^*(\xi)$. By construction, the polynomial $\mathsf{P}^*(v) - v^{\mathsf{deg}_+(\mathsf{P}^*(v))}$ is non-negative on $\xi$, since $\mathsf{P}(\overline{\boldsymbol{u}})$ is strongly non-negative *w.r.t.* $\mathfrak{P}$. Now, by Lemma 21, $\mathsf{deg}(\mathsf{M}^*(v)) > \mathsf{deg}_+(\mathsf{P}^*(v))$. From this, we can derive a solution to the $\langle \mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}}), \mathfrak{P} \rangle$-system of inequalities of suitable size, again invoking Lemma 22.

**General Solutions.** We can now refocus on Prob. 2. The high-level idea behind the technique we employ here can be summarised as follows. Each coherent solution $\overline{\boldsymbol{\xi}}$ to the original *MPI* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$ can be condensed into a strict solution $\overline{\boldsymbol{\zeta}}$ of a reduced *MPI* $\mathsf{M}_\gamma(\overline{\boldsymbol{v}}) > \mathsf{P}_\gamma(\overline{\boldsymbol{v}})$, obtained from the first by mapping unknowns in $\overline{\boldsymbol{u}}$ to unknowns in $\overline{\boldsymbol{v}} \subseteq \overline{\boldsymbol{u}}$ via a function $\gamma$. Specifically, we merge all unknowns in $\overline{\boldsymbol{u}}$ that are related by the order and share the same value in $\overline{\boldsymbol{\xi}}$ into a single unknown in $\overline{\boldsymbol{v}}$. Conversely, every coherent solution $\overline{\boldsymbol{\zeta}}$ for the reduced *MPI* can be expanded into a coherent solution $\overline{\boldsymbol{\xi}}$ of the original *MPI* by duplicating the value of an unknown in $\overline{\boldsymbol{v}}$ across all corresponding unknowns in $\overline{\boldsymbol{u}}$. To implement this idea, we leverage the following concept.

For a *poset* $\mathfrak{P} = \langle \mathrm{U}, \preceq \rangle$ on an arbitrary set U, a *$\mathfrak{P}$-collapse* $\gamma \colon \mathrm{U} \to (\{0\} \cup \mathrm{U})$ is a function such that (a) it is deflationary, *i.e.*, $\gamma(u) \preceq_0 u$, for all $u \in \mathrm{U}$, (b) it is monotone, *i.e.*, $\gamma(u_1) \preceq_0 \gamma(u_2)$, for all $u_1, u_2 \in \mathrm{U}$ with $u_1 \prec u_2$, and (c) all elements of U in the range of $\gamma$ are fixpoints of $\gamma$ itself, *i.e.*, $\gamma(u) = u$, for all $u \in \mathsf{rng}(\gamma) \setminus \{0\}$.

Given a polynomial $\mathsf{Q}(\overline{\boldsymbol{u}}) \in \mathbb{Z}[\overline{\boldsymbol{u}}]$, we denote by $\mathsf{Q}(\overline{\boldsymbol{u}})\lceil_\gamma$ the reduced version induced by $\gamma$. To derive this, we initially replace each unknown $u$ in $\overline{\boldsymbol{u}}$ with the corresponding element $\gamma(u)$, yielding $\mathsf{Q}(\overline{\boldsymbol{u}})[u \in \overline{\boldsymbol{u}}/\gamma(u)]$. Subsequently, we simplify the obtained expression by eliminating all monomials $\overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}}[u \in \overline{\boldsymbol{u}}/\gamma(u)]$ that reduce to 0 due to an unknown $u$ being mapped to 0 by $\gamma$ and raised to a non-zero exponent $e_u$. As usual when dealing with the polynomial ring $\mathbb{Z}[\overline{\boldsymbol{u}}]$, we conventionally assume $0^0 \triangleq 1$, thereby preventing cancellation for unknowns $u$ raised to a zero exponent. An alternative way to define $\mathsf{Q}(\overline{\boldsymbol{u}})\restriction_\gamma$ is as follows. For each unitary monomial $\overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}}$, if $\gamma(u) = 0$ for some $u \in \overline{\boldsymbol{u}}$ with $e_u \neq 0$, then $\overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}}\restriction_\gamma \triangleq 0$, else $\overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}}\restriction_\gamma \triangleq \overline{\boldsymbol{v}}^{\overline{\boldsymbol{e}}\restriction_\gamma}$, where $\overline{\boldsymbol{v}} \triangleq \mathsf{rng}(\gamma) \setminus \{0\}$ and $(\overline{\boldsymbol{e}}\restriction_\gamma)_v \triangleq \sum_{u \in \overline{\boldsymbol{u}}}^{\gamma(u)=v} \overline{\boldsymbol{e}}_u$, for all $v \in \overline{\boldsymbol{v}}$. Then, $\mathsf{Q}(\overline{\boldsymbol{u}})\restriction_\gamma \triangleq \sum_{i=1}^{m} \alpha_i \overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}_i}\restriction_\gamma$, where $\mathsf{Q}(\overline{\boldsymbol{u}}) = \sum_{i=1}^{m} \alpha_i \overline{\boldsymbol{u}}^{\overline{\boldsymbol{e}}_i}$.

Consider again the *MPI* $\mathsf{M}_q(\overline{\boldsymbol{u}}) > \mathsf{P}_q^p(\overline{\boldsymbol{u}})$, the linear inequality system $\{\mathsf{C}_q^a(\overline{\boldsymbol{u}}) \geq 0\}_{a \in [\![q]\!]}$, and the 0-*cmsem* $\mathfrak{P} = \langle [\![q]\!], \preceq_q \rangle$, with $[\![q]\!] = \{a_1, a_{12}, a_2, a_{23}, a_3\}$, discussed after Thm. 16. Moreover, let $\gamma$ be the function mapping (i) $u_{a_2}$ to $u_{a_{12}}$, (ii) $u_{a_{23}}$ to 0, and (iii) all the other unknowns $u_{a_i}$ to themselves. One can easily check that $\gamma$ is a $\mathfrak{P}$-collapse. The reduced *MPI* $\mathsf{M}_q(\overline{\boldsymbol{u}})\restriction_\gamma > \mathsf{P}_q^p(\overline{\boldsymbol{u}})\restriction_\gamma$ is, thus, $u_{a_1} u_{a_{12}} u_{a_3} > (u_{a_1}^2 + u_{a_3}^3) u_{a_1} u_{a_{12}}$, while the reduced linear system $\{\mathsf{C}_q^a(\overline{\boldsymbol{u}})\restriction_\gamma \geq 0\}_{a \in [\![q]\!]}$ is

$$\mathsf{C}_q^{a_1}(\overline{\boldsymbol{u}})\restriction_\gamma = u_{a_1} - u_{a_{12}} \geq 0, \qquad \mathsf{C}_q^{a_{12}}(\overline{\boldsymbol{u}})\restriction_\gamma = u_{a_{12}} \geq 0,$$
$$\mathsf{C}_q^{a_2}(\overline{\boldsymbol{u}})\restriction_\gamma = 0 \geq 0, \qquad \mathsf{C}_q^{a_{23}}(\overline{\boldsymbol{u}})\restriction_\gamma = 0 \geq 0, \qquad \mathsf{C}_q^{a_3}(\overline{\boldsymbol{u}})\restriction_\gamma = u_{a_3} \geq 0.$$

It should be clear that the value $\mathsf{Q}(\overline{\boldsymbol{\xi}})$ of a polynomial $\mathsf{Q}(\overline{\boldsymbol{u}})$ on an assignment $\overline{\boldsymbol{\xi}}$ of its unknowns always coincides with value $\mathsf{Q}_\gamma(\overline{\boldsymbol{\zeta}})$ of its $\gamma$-induced version $\mathsf{Q}_\gamma(\overline{\boldsymbol{v}})$ on the condensed assignment $\overline{\boldsymbol{\zeta}}$, as stated below.

▶ **Lemma 24.** *Let* $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *be a* poset *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ *and* $\gamma \colon \overline{\boldsymbol{u}} \to (\{0\} \cup \overline{\boldsymbol{u}})$ *a* $\mathfrak{P}$*-collapse, with* $\overline{\boldsymbol{v}} \triangleq \mathsf{rng}(\gamma) \setminus \{0\} \in \mathrm{Un}^k$*. Moreover, let* $\mathsf{Q}(\overline{\boldsymbol{u}}) \in \mathbb{Z}[\overline{\boldsymbol{u}}]$ *be a polynomial on* $\overline{\boldsymbol{u}}$ *and* $\mathsf{Q}_\gamma(\overline{\boldsymbol{v}}) \triangleq \mathsf{Q}(\overline{\boldsymbol{u}})\restriction_\gamma$ *its* $\gamma$*-induced version on* $\overline{\boldsymbol{v}}$*. Then,* $\mathsf{Q}(\overline{\boldsymbol{\xi}}) = \mathsf{Q}_\gamma(\overline{\boldsymbol{\zeta}})$*, for all* $\overline{\boldsymbol{\xi}} = \{\xi_u\}_{u \in \overline{\boldsymbol{u}}} \in \mathbb{N}^n$ *and* $\overline{\boldsymbol{\zeta}} = \{\zeta_v\}_{v \in \overline{\boldsymbol{v}}} \in \mathbb{N}^k$*, with* $\xi_u = \zeta_{\gamma(u)}$*, if* $\gamma(u) \in \overline{\boldsymbol{v}}$*, and* $\xi_u = 0$*, otherwise.*

We can also prove that every coherence constraint $\mathsf{C}_{\mathfrak{P}}^u(\overline{\boldsymbol{u}}) \geq 0$ imposed on an unknown $u \in \overline{\boldsymbol{u}}$ within a *poset* $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ either is trivially satisfied or reduces to a corresponding constraint $\mathsf{C}_{\mathfrak{P}\restriction_\gamma}^u(\overline{\boldsymbol{v}}) \geq 0$ of the $\gamma$-induced *poset* $\langle \mathrm{U}, \preceq \rangle\restriction_\gamma \triangleq \langle \mathsf{rng}(\gamma) \setminus \{0\}, \preceq \cap (\mathsf{rng}(\gamma) \times \mathsf{rng}(\gamma)) \rangle$.

Going back to the above example, by applying the $\gamma$-restriction to the 0-*cmsem* $\mathfrak{P}$, we obtain the $\gamma$-induced 0-*cmsem* $\mathfrak{P}\restriction_\gamma = \langle \{u_{a_1}, u_{a_{12}}, u_{a_3}\}, \{(u_{a_{12}}, u_{a_1})\} \rangle$ and the three constraints $\mathsf{C}_{\mathfrak{P}\restriction_\gamma}^{u_{a_1}}(\overline{\boldsymbol{v}}) = u_{a_1} - u_{a_{12}} \geq 0$, $\mathsf{C}_{\mathfrak{P}\restriction_\gamma}^{u_{a_{12}}}(\overline{\boldsymbol{v}}) = u_{a_{12}} \geq 0$, and $\mathsf{C}_{\mathfrak{P}\restriction_\gamma}^{u_{a_3}}(\overline{\boldsymbol{v}}) = u_{a_3} \geq 0$.

▶ **Lemma 25.** *Let* $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *be a* poset *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ *and* $\gamma \colon \overline{\boldsymbol{u}} \to (\{0\} \cup \overline{\boldsymbol{u}})$ *a* $\mathfrak{P}$*-collapse, with* $\overline{\boldsymbol{v}} \triangleq \mathsf{rng}(\gamma) \setminus \{0\} \in \mathrm{Un}^k$*. Then,* $\mathsf{C}_{\mathfrak{P}}^u(\overline{\boldsymbol{u}})\restriction_\gamma = \mathsf{C}_{\mathfrak{P}\restriction_\gamma}^u(\overline{\boldsymbol{v}})$*, if* $u \in \overline{\boldsymbol{v}}$*, and* $\mathsf{C}_{\mathfrak{P}}^u(\overline{\boldsymbol{u}})\restriction_\gamma = 0$*, otherwise.*

Building upon the above results, it becomes quite evident that any $\mathfrak{P}\restriction_\gamma$-coherent solution $\overline{\boldsymbol{\zeta}} = \{\zeta_v\}_{v \in \overline{\boldsymbol{v}}}$ of a $\gamma$-induced *MPI* $\mathsf{M}_\gamma(\overline{\boldsymbol{v}}) > \mathsf{P}_\gamma(\overline{\boldsymbol{v}})$, for some $\mathfrak{P}$-collapse $\gamma$, always translates to a $\mathfrak{P}$-coherent solution $\overline{\boldsymbol{\xi}} = \{\xi_u\}_{u \in \overline{\boldsymbol{u}}}$ of the original *MPI* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$, by just replicating the value $\zeta_v$ on all unknowns $u \in \overline{\boldsymbol{u}}$ with $\gamma(u) = v$.

For the converse direction, however, the situation is more intricate. In principle, given a $\mathfrak{P}$-coherent solution $\overline{\boldsymbol{\xi}} = \{\xi_u\}_{u \in \overline{\boldsymbol{u}}}$ of the original *MPI*, we would like to collapse two unknowns $u_1$ and $u_2$ in $\overline{\boldsymbol{u}}$, with $u_1 \prec u_2$ and $\xi_{u_1} = \xi_{u_2}$, into the same unknown in $\overline{\boldsymbol{v}}$, by defining a suitable $\mathfrak{P}$-collapse $\gamma$. Unfortunately, the existence of such a $\mathfrak{P}$-collapse is not guaranteed on arbitrary *poset* $\mathfrak{P}$, due to the potential occurrence of multiple maximal lower bounds for two elements. This is where the property of a 0-*cmsem* comes to the rescue. Indeed, by exploiting the following truncated version of the *inclusion/exclusion principle* on a meet semilattice, we can show that there always exists a $\mathfrak{P}$-collapse $\gamma$ aligned with the above requirements.

▶ **Lemma 26.** *Let* $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *be a* 0-cmsem *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ *and* $\overline{\boldsymbol{\xi}} = \{\xi_u\}_{u \in \overline{\boldsymbol{u}}} \in \mathbb{N}^n$ *a* $\mathfrak{P}$*-coherent* $n$*-vector of values. Then, for all* $u, u_1, u_2 \in \overline{\boldsymbol{u}}$ *with* $u_1 \prec u$ *and* $u_2 \prec u$*, if* $\widehat{u} \triangleq u_1 \wedge u_2 \neq 0$ *then* $\mathsf{C}^u_{\mathfrak{P}}(\overline{\boldsymbol{\xi}}) \leq \xi_u - \xi_{u_1} - \xi_{u_2} + \xi_{\widehat{u}}$ *else* $\mathsf{C}^u_{\mathfrak{P}}(\overline{\boldsymbol{\xi}}) \leq \xi_u - \xi_{u_1} - \xi_{u_2}$*.*

A reduction from Prob. 2 to Prob. 3 can now be proved.

▶ **Theorem 27.** *Let* $\mathsf{M}(\overline{\boldsymbol{u}}) > \mathsf{P}(\overline{\boldsymbol{u}})$ *be an* MPI *on* $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ *and* $\mathfrak{P} = \langle \overline{\boldsymbol{u}}, \preceq \rangle$ *a* poset*. If there exists a* $\mathfrak{P}$*-collapse* $\gamma$ *such that the* $\gamma$*-induced* MPI $\mathsf{M}(\overline{\boldsymbol{u}})\restriction_{\gamma} > \mathsf{P}(\overline{\boldsymbol{u}})\restriction_{\gamma}$ *admits a* $\mathfrak{P}\restriction_{\gamma}$*-coherent Diophantine solution then the original* MPI *admits a* $\mathfrak{P}$*-coherent Diophantine solution. Conversely, under the assumption that* $\mathfrak{P}$ *is a* 0-cmsem*, if the original* MPI *admits a* $\mathfrak{P}$*-coherent Diophantine solution then there exists a* $\mathfrak{P}$*-collapse* $\gamma$ *such that the* $\gamma$*-induced* MPI $\mathsf{M}(\overline{\boldsymbol{u}})\restriction_{\gamma} > \mathsf{P}(\overline{\boldsymbol{u}})\restriction_{\gamma}$ *admits a* $\mathfrak{P}\restriction_{\gamma}$*-strict Diophantine solution.*

**Deciding Solutions.** To effectively address Prob. 2, we leverage the combination of Thms. 23 and 27 via a $\exists\forall$-alternating polynomial-time algorithm that outputs true *iff* the problem has an affirmative answer. We adopt a classic guess-and-check strategy: first guess a potential $\mathfrak{P}$-collapse $\gamma$ alongside a polynomially-bounded solution $\overline{\boldsymbol{d}}$ for the induced $\langle \mathsf{M}_{\gamma}(\overline{\boldsymbol{u}}) > \mathsf{P}_{\gamma}(\overline{\boldsymbol{u}}), \mathfrak{P}\restriction_{\gamma} \rangle$-system and then check the correctness of these choices. Consequently, we derive the following upper bound on the complexity of our Diophantine problem.

▶ **Theorem 28.** *Prob. 2 can be decided in* $\exists\forall$*-alternating polynomial-time with existential-guess space polynomial in the number* $n$ *of unknowns and universal-guess space logarithmic in maximum between* $n$ *and the number* $m$ *of monomials.*

## 6 The Problem Solution

We can finally combine Thms. 16 and 28 to obtain a solution for the bag-set containment problem of *JoFBQ*s into *BCQ*s. Before doing so, we first need to show that the minimal unification closure $[\![q]\!]$ of a containee *JoFBQ* $q$ equipped with the homomorphic ordering $\preceq_q$ is a 0-*cmsem*. Then, we prove that the polynomial $\mathsf{P}^p_q(\overline{\boldsymbol{u}})$ tallying the homomorphisms of the containing *BCQ* $p$ is strongly non-negative *w.r.t.* $\mathfrak{P}$. The intuition behind this fact is that $\mathsf{P}^p_q(\overline{\boldsymbol{u}})$ represents a sum of positive monomials, where each monomial accounts for the homomorphisms of an atom in $[\![q]\!]$. Consequently, removing a single monomial, even more so if it is unitary, cannot result in a negative value.

▶ **Lemma 29.** *For every pair* $(q, p) \in$ JoF-BCP*, the following properties hold true: (a)* $\mathfrak{P} = \langle [\![q]\!], \preceq_q \rangle$ *is a* 0-cmsem*; (b) the polynomial* $\mathsf{P}^p_q(\overline{\boldsymbol{u}})$ *is strongly non-negative* w.r.t. $\mathfrak{P}$*.*

Since every atom in $[\![q]\!]$ corresponds to the most general unification of a unifiable subset of atoms from $q$, there are $n \leq 2^{|q|}$ many elements in this set and, thus, so many unknowns in the vector $\overline{\boldsymbol{u}} \in \mathrm{Un}^n$ of the *MPI* and linear constraints of Thm. 16. Moreover, there are $m \leq 2^{|q| \cdot |p|}$ many homomorphisms in the set $\mathrm{Hom}(p, [\![q]\!])$ and, thus, so many monomials in the polynomial $\mathsf{P}^p_q(\overline{\boldsymbol{u}})$. Therefore, by applying Thm. 28, we obtain a $\forall\exists$-alternating exponential-time algorithm with a universal-guess space exponential in $|q|$ and an existential-guess space polynomial in $|q| \cdot |p|$. Since the latter guess can be deterministically simulated in exponential time in $|q| \cdot |p|$, we obtain a universal exponential-time algorithm, witnessing the membership of the problem to CoNExpTime, *i.e.*, within the first non-trivial level of the exponential hierarchy. By relying on Thm. 2, we can extend the result to encompass the bag-bag containment problem. Furthermore, via Thm. 3, both problems can be addressed, even when dealing with non-Boolean queries.

▶ **Theorem 30.** *The bag-bag and bag-set containment problems for a join-on-free containee* CQ $p$ *against an arbitrary containing* CQ $q$ *can be solved in* CoNExpTime*.*

## 7 Discussion

We study a new class of conjunctive queries, namely *join-on-free queries* (*JoFQ*), which, to the best of our knowledge, has not been previously considered in the literature. Specifically, we establish the decidability of the containment problem of *JoFQ*s into arbitrary *CQ*s, under both bag-bag/bag-set semantics, proving membership in CoNExpTime. This class seems to represent the most comprehensive set of containee queries examined against arbitrary containing ones for which the problem has been studied and solved so far.

The solution strategy adopted in this work can be broken down into two independent, but synergistic, contributions. On the one hand, we provide a characterisation of the containment problem through the concept of *multicanonical instance*, which is, in turn, based on the notion of *minimal unification closure* of a *JoFQ*. This enabled us to reduce the problem between two queries to a direct comparison of the respective numbers of homomorphisms towards this instance. In particular, these numbers can be computed via a *monomial function*, for the containee query, and a *polynomial function*, for the containing one, the latter enjoying the crucial property of *strong non-negativeness*. On the other hand, we solve a special case of the *Diophantine inequality problem*. While in its general form it is known to be undecidable, we prove the decidability of the problem in the context of *monomial-polynomial inequalities* under the proviso of strong non-negativeness of the polynomials.

This work leaves several intriguing directions for future investigation. First, while the containment problems studied here are known to be NPTime-hard [24], identifying tighter lower bounds remains an open challenge. Notably, since the withdrawal of the $\Pi_2^{\mathrm{P}}$-hardness claim by Chaudhuri and Vardi for the general bag containment problem, progress in this area has stalled, making it an appealing avenue for further exploration. Second, the potential to improve the CoNExpTime upper-bound of our algorithm remains uncertain. Achieving decidability within this complexity class is already a significant tour de force, but it is not yet clear whether the current approach includes redundancies that could be removed to yield a more efficient algorithm. Finally, our results suggest that the techniques developed here could generalise to a broader class of queries. Relaxing the join-on-free condition to permit a restricted form of joins, while ensuring that the minimal unification closure remains a 0-completable meet semilattice *w.r.t.* bag-set inclusion, could significantly expand the applicability of our approach.

At the current state of the art, alongside early studies employing classic notions of database theory (*e.g.*, [5, 14, 12, 9, 10, 2, 6, 7, 8]), two main strategies have been considered in the literature to tackle the vanilla version of this difficult problem: one involves restricting the containing query, while keeping the containee query unrestricted; the other adopts the dual approach. Both strategies, however, exploit technical tools from different fields. In the first category, researches such as [25, 17, 18] examine the strong correlation between the bag-containment problem and the validity of specific information-theory inequalities. In the second category, instead, we can cast the present article and [24], upon which our work builds, that investigate the problem in correlation with the solution of specific Diophantine inequalities. It seems conceivable that a suitable integration of these diverse techniques may lead to the final solution of this challenging open problem.

## Acknowledgements

We would like to thank the anonymous reviewers for their insightful comments (*e.g.*, regarding the connection with the Möbius inversion formula) that considerably improved the paper.

## References


1 S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
2 F.N. Afrati, M. Damigos, and M. Gergatsoulis. Query Containment under Bag and Bag-Set Semantics. *IPL*, 110(10):360–369, 2010.
3 R.J. Brachman and H.J. Levesque. *Knowledge Representation and Reasoning*. Morgan Kaufmann, 2004.
4 A.K. Chandra and P.M. Merlin. Optimal Implementation of Conjunctive Queries in Relational Data Bases. In *STOC'77*, pages 77–90. ACM, 1977.
5 S. Chaudhuri and M.Y. Vardi. Optimization of Real Conjunctive Queries. In *PODS'93*, pages 59–70. ACM, 1993.
6 R. Chirkova. Equivalence and Minimization of Conjunctive Queries under Combined Semantics. In *ICDT'12*, pages 262–273. OpenProceedings.org, 2012.
7 R. Chirkova. Combined-Semantics Equivalence and Minimization of Conjunctive Queries. *TCJ*, 57(5):775–795, 2014.
8 R. Chirkova. Combined-Semantics Equivalence of Conjunctive Queries: Decidability and Tractability Results. *JCSS*, 82(3):395–465, 2016.
9 S. Cohen. Equivalence of Queries Combining Set and Bag-Set Semantics. In *PODS'06*, pages 70–79. ACM, 2006.
10 S. Cohen. Equivalence of Queries that are Sensitive to Multiplicities. *PVLDB*, 18(3):765–785, 2009.
11 M. Davis. Hilbert's Tenth Problem is Unsolvable. *AMM*, 80(3):233–269, 1973.
12 S. Grumbach and T. Milo. Towards Tractable Algebras for Bags. *JCSS*, 52(3):570–588, 1996.
13 H. Grunert and A. Heuer. Query Rewriting by Contract under Privacy Constraints. *OJIOT*, 4(1):54–69, 2018.
14 Y.E. Ioannidis and R. Ramakrishnan. Containment of Conjunctive Queries: Beyond Relations as Sets. *TODS*, 20(3):288–324, 1995.
15 Y.E. Ioannidis and E. Wong. Towards an Algebraic Theory of Recursion. *JACM*, 38(2):329–381, 1991.
16 T.S. Jayram, P.G. Kolaitis, and E. Vee. The Containment Problem for Real Conjunctive Queries with Inequalities. In *PODS'06*, pages 80–89. ACM, 2006.
17 M.A. Khamis, P.G. Kolaitis, H.Q. Ngo, and D. Suciu. Bag Query Containment and Information Theory. In *PODS'20*, pages 95–112. ACM, 2020.
18 M.A. Khamis, P.G. Kolaitis, H.Q. Ngo, and D. Suciu. Bag Query Containment and Information Theory. *TODS*, 46(3):1–39, 2021.
19 M.A. Khamis, H.Q. Ngo, and A. Rudra. FAQ: Questions Asked Frequently. In *PODS'16*, pages 13–28. ACM, 2016.
20 P.G. Kolaitis. The Query Containment Problem: Set Semantics vs. Bag Semantics. In *AMW'13*, CEUR-WS 1949, 2013.
21 G. Konstantinidis and J.L. Ambite. Scalable Query Rewriting: A Graph-Based Approach. In *SIGMOD'11*, pages 97–108. ACM, 2011.
22 G. Konstantinidis and J.L. Ambite. Scalable Containment for Unions of Conjunctive Queries under Constraints. In *SWIM'13*, pages 4:1–8. ACM, 2013.
23 G. Konstantinidis, J. Holt, and A. Chapman. Enabling Personal Consent in Databases. *PVLDB*, 15(2):375–387, 2021.
24 G. Konstantinidis and F. Mogavero. Attacking Diophantus: Solving a Special Case of Bag Containment. In *PODS'19*, pages 399–413. ACM, 2019.
25 S. Kopparty and B. Rossman. The Homomorphism Domination Exponent. *EJC*, 32(7):1097–1114, 2011.
26 A.Y. Levy, A.O. Mendelzon, Y. Sagiv, and D. Srivastava. Answering Queries Using Views. In *PODS'95*, pages 95–104. ACM, 1995.
27 A. Machanavajjhala and J. Gehrke. On the Efficiency of Checking Perfect Privacy. In *PODS'06*, pages 163–172. ACM, 2006.

**28** J. Marcinkowski and M. Orda. Bag Semantics Conjunctive Query Containment. Four Small Steps Towards Undecidability. *PACMMOD*, 2(2):103, 2024.

**29** Y. Matiyasevich. *Hilbert's 10th Problem*. MIT Press, 1993.

**30** J. Robinson. Solving Diophantine Equations. In *Studies in Logic and the Foundations of Mathematics*, volume 74, pages 63–67. Elsevier, 1973.

**31** G. Rota. On the Foundations of Combinatorial Theory I. Theory of Möbius Functions. In *Classic Papers in Combinatorics*, volume 2, pages 340–368. Springer, 1964.

**32** A. Schrijver. *Theory of Linear and Integer Programming*. Wiley, 1986.

**33** R.P. Stanley. *Enumerative Combinatorics: Volume 1*. CUP, 2011.