



6-13-2011

Provenance for Aggregate Queries

Yael Amsterdamer
University of Pennsylvania

Daniel Deutch
Ben Gurion University

Val Tannen
University of Pennsylvania, val@cis.upenn.edu

Follow this and additional works at: http://repository.upenn.edu/cis_papers

 Part of the [Computer Sciences Commons](#)

Recommended Citation

Yael Amsterdamer, Daniel Deutch, and Val Tannen, "Provenance for Aggregate Queries", . June 2011.

Yael Amsterdamer, Daniel Deutch, and Val Tannen. 2011. Provenance for aggregate queries. In *Proceedings of the thirtieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems (PODS '11)*. ACM, New York, NY, USA, 153-164.

doi:10.1145/1989284.1989302

© ACM, 2011. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *Proceedings of the thirtieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, {(2011)} <http://doi.acm.org/10.1145/1989284.1989302> Email permissions@acm.org

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/cis_papers/645
For more information, please contact libraryrepository@pobox.upenn.edu.

Provenance for Aggregate Queries

Abstract

We study in this paper provenance information for queries with aggregation. Provenance information was studied in the context of various query languages that do not allow for aggregation, and recent work has suggested to capture provenance by annotating the different database tuples with elements of a commutative semiring and propagating the annotations through query evaluation. We show that aggregate queries pose novel challenges rendering this approach inapplicable. Consequently, we propose a new approach, where we annotate with provenance information not just tuples but also the individual values within tuples, using provenance to describe the values computation. We realize this approach in a concrete construction, first for “simple” queries where the aggregation operator is the last one applied, and then for arbitrary (positive) relational algebra queries with aggregation; the latter queries are shown to be more challenging in this context. Finally, we use aggregation to encode queries with difference, and study the semantics obtained for such queries on provenance annotated databases.

Disciplines

Computer Sciences

Comments

Yael Amerdamer, Daniel Deutch, and Val Tannen. 2011. Provenance for aggregate queries. In *Proceedings of the thirtieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems (PODS '11)*. ACM, New York, NY, USA, 153-164.

doi:[10.1145/1989284.1989302](https://doi.org/10.1145/1989284.1989302)

© ACM, 2011. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *Proceedings of the thirtieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, {(2011)}* <http://doi.acm.org/10.1145/1989284.1989302> Email permissions@acm.org

Provenance for Aggregate Queries

Yael Amsterdamer
University of Pennsylvania
and Tel Aviv University

Daniel Deutch
University of Pennsylvania
and Ben Gurion University

Val Tannen
University of Pennsylvania

ABSTRACT

We study in this paper provenance information for queries with *aggregation*. Provenance information was studied in the context of various query languages that do not allow for aggregation, and recent work has suggested to capture provenance by annotating the different database tuples with elements of a *commutative semiring* and propagating the annotations through query evaluation. We show that aggregate queries pose novel challenges rendering this approach inapplicable. Consequently, we propose a new approach, where we annotate with provenance information not just tuples but also the *individual values* within tuples, using provenance to describe the values computation. We realize this approach in a concrete construction, first for “simple” queries where the aggregation operator is the last one applied, and then for arbitrary (positive) relational algebra queries with aggregation; the latter queries are shown to be more challenging in this context. Finally, we use aggregation to encode queries with *difference*, and study the semantics obtained for such queries on provenance annotated databases.

1. INTRODUCTION

The annotation of the results of database transformations with provenance information has quite a few applications [19, 6, 35, 10, 11, 38, 34, 25, 27, 39, 37, 2, 4]. Recent work [24, 17, 21] has proposed a framework of *semiring annotations* that allows us to state formally what is expected of such provenance information. These papers have developed the framework for the positive fragment of the relational algebra (as well as for Datalog, the positive Nested Relational Calculus, and some query languages on trees/XML). The main goal of this paper is to extend the framework to *aggregate operations*.

In the perspective promoted by these papers, *provenance* is a general form of annotation information that can be specialized for different purposes, such as multiplicity, trust, cost, security, or identification of “possible worlds” which in turn applies to incomplete databases, deletion propagation, and probabilistic databases. In fact, the introduction of the framework in [24] was motivated by the need to track trust and deletion propagation in the Orchestra system [23]. What makes such a diversity of applications possible is that each is captured by a different semiring, while provenance is represented by elements of a semiring of *polynomials*. One then relies on the property that any semiring-annotation semantics *factors* through the provenance polynomials semantics. This means that storing provenance polynomials allows for many other practical applications. For example, to capture access control, where the access to different tuples require different security credentials, we can simply evalu-

R			
$EmpId$	$Dept$	Sal	
1	d_1	20	p_1
2	d_1	10	p_2
3	d_1	15	p_3
4	d_2	10	r_1
5	d_2	15	r_2

$Dept$	
d_1	$p_1 + p_2 + p_3$
d_2	$r_1 + r_2$

Figure 1: Projection on annotated relations

ate the polynomials in the *security semiring*, and propagate the security annotations through query evaluation (see Section 2.1), assigning security levels to query results.

Let us briefly illustrate *deletion propagation* as an application of provenance. Consider a simple example of an employee/department/salary relation R shown in Figure 1(a).

The variables p_1, p_2, p_3, r_1, r_2 can be thought of as tuple identifiers and in the framework of *provenance polynomials* [24] they are the “provenance tokens” or “indeterminates” out of which provenance is built. We denote by $\mathbb{N}[X]$ the set of provenance polynomials (here $X = \{p_1, p_2, p_3, r_1, r_2\}$). R can be seen as an $\mathbb{N}[X]$ -annotated relation; as defined in [24] the evaluation of query, for example $\Pi_{Dept}R$, produces another $\mathbb{N}[X]$ -annotated relation, in this example the one shown in Figure 1(b). Intuitively, in this simple example, the summation in the annotation of every result tuple is over the identifiers of its alternative origins¹.

Now, the result of propagating the deletions of tuples with $EmpId$ 3 and 5 in R is obtained by simply setting $p_3 = r_2 = 0$ in the answer. We get the same two tuples in the query answer but their provenances change to $p_1 + p_2$ and r_1 , respectively. If the tuple with $EmpId$ 4 is also deleted from R then we also set $r_1 = 0$, and the second tuple in the answer is deleted because its provenance has now become 0. This algebraic treatment of deletions is related to the counting algorithm for view maintenance [26], but is more general as it incrementally maintains not just the data but also the provenance.

An intuitive way of understanding what happens is that provenance-aware evaluation of queries conveniently “commutes” with deletions. In fact, in [24, 17] this intuition is captured formally by theorems that state that query evaluation commutes with semiring homomorphisms. The factorization through provenance relies on this and on the fact that the polynomial provenance semiring is “freely generated”. All applications of provenance polynomials we have listed, for trust, security, etc., are based on these theorems.

¹we explain how the annotations of query results are computed in Section 2.1

Dept	SalMass	
d ₁	45	$p_1 p_2 p_3$
d ₁	30	$p_1 p_2 \hat{p}_3$
d ₁	35	$p_1 \hat{p}_2 p_3$
d ₁	25	$\hat{p}_1 p_2 p_3$
d ₁	20	$p_1 \hat{p}_2 \hat{p}_3$
d ₁	10	$\hat{p}_1 p_2 \hat{p}_3$
d ₁	15	$\hat{p}_1 \hat{p}_2 p_3$
...

Dept	SalMass	
d ₁	30	$p_1 p_2$
d ₁	20	$p_1 \hat{p}_2$
d ₁	10	$\hat{p}_1 p_2$
...

(a)

(b)

Figure 2: A naive approach to aggregation

Thus, commutation with homomorphisms is an essential criterion for our proposed framework extension to aggregate operations. However, in Section 3.1 we prove that the framework of semiring-annotated relations introduced in [24] cannot be extended to handle aggregation while both satisfying commutation with homomorphisms and working as usual on set or bag relations.

If the semiring operations are not enough then perhaps we can add others? This is a natural idea so we illustrate it on the same R in Figure 1(a) and we use again the necessity to support deletion propagation to guide the approach. Consider the query that groups by *Dept* and sums *Sal*. The result of the summation depends on which tuples participate in it. To provide enough information to obtain all the possible summation results for all possible sets of deletions, we could use the representation in Figure 2(a) where we add to the semiring operations an unary operation $\hat{\cdot}$ with the property that $\hat{p} = 1$ whenever $p = 0$. This will indeed satisfy the deletion criterion. For example when the tuple with *Id* 3 is deleted we get the relation in Figure 2(b). In fact, there exist semirings with the additional structure needed to define $\hat{\cdot}$. For example in the semiring of polynomials with integer coefficients, $\mathbb{Z}[X]$, we can take $\hat{p} = 1 - p$ while in the semiring of boolean expressions with variables from X , $\text{BoolExp}(X)$, we can take $\hat{p} = \neg p$ ². The latter is essentially the approach taken in [31]. However, whether we use $\mathbb{Z}[X]$ or $\text{BoolExp}(X)$, we still have, in the worst case, *exponentially many different results* to account for, at least in the case of summation (a lower bound recognized also in [31]). It follows that summation in particular (and therefore any uniform approach to aggregation) cannot be represented with a feasible amount of annotation as long as the annotation stays at the tuple level.

Instead, we will present a provenance representation for aggregation results that leads only to a *poly-size increase*, one that we believe is tractably implementable using methods similar to the ones used in Orchestra [23]. We achieve this via a more radical approach: we annotate with provenance information not just the tuples of the answer but also *the manner in which the values in those tuples are computed*.

We can gain intuition towards our representation from the particular case of bags, which are in fact \mathbb{N} -relations, i.e., relations whose tuples are annotated with elements of the semiring $(\mathbb{N}, +, \cdot, 0, 1)$. Assume that R in Figure 1(a) is such a relation, i.e., $p_1, \dots, p_2 \in \mathbb{N}$ are tuple multiplicities. Then, after sum-aggregation the value of the attribute *SalMass* in the tuple with *Dept* d_1 is computed by $p_1 \times 20 + p_2 \times 10 + p_3 \times 15$. Now, if the multiplicities are, for example $p_1 = 2, p_2 = 3, p_3 = 1$ then the aggregate value is 85. But what if R is a relation annotated with provenance polynomials

²Both of these also give natural semantics to relational *difference*. $\mathbb{Z}[X]$ is used in [20] following the use of \mathbb{Z} in [22], while $\text{BoolExp}(X)$ is used in the seminal paper [28].

rather than multiplicities? Then, the aggregate value does not correspond to any number.

We will make $p_1 \times 20$ into an *abstract* construction, $p_1 \otimes 20$ and the aggregate value will be the formal expression $p_1 \otimes 20 + p_2 \otimes 10 + p_3 \otimes 15$.

Intuitively, we are embedding the domain of sum-aggregates, i.e., the reals \mathbb{R} , into a larger domain of formal expressions that capture how the sum-aggregates are computed from values annotated with provenance. We do the same for other kinds of aggregation, for instance min-aggregation gives $p_1 \otimes 20 \min p_2 \otimes 10 \min p_3 \otimes 15$. We call these *annotated aggregate* expressions.

In this paper we consider only aggregations defined by commutative-associative operations³. Specifically, our framework can accommodate aggregation based on any *commutative monoid*. For example the commutative monoid for summation is $(\mathbb{R}, +, 0)$ while the one for min is $(\mathbb{R}^\infty, \min, \infty)$ ⁴.

To combine an aggregation monoid M with an annotation commutative semiring K , in a way capturing aggregates over K -relations, we propose the use of the algebraic structure of *K -semimodules* (see Section 2.2). Semimodules are a way of generalizing (a lot!) the operations considered in linear algebra. Its “vectors” form only a commutative monoid (rather than an abelian group) and its “scalars” are the elements of K which is only a commutative semiring (rather than a field).

In general, a commutative monoid M does not have an obvious structure of K -semimodule. To make it such we may need to add new elements corresponding to the scalar multiplication of elements of M with elements from K , thus ending up with the formal expressions that represent aggregate computations, motivated above, as elements of a *tensor product* construction $K \otimes M$. We show that the use of tensor product expressions as a formal representation of aggregation result is effective in managing the provenance of “simple aggregate” queries, namely queries where the aggregation operators are the last ones applied.

We show that certain semirings are “compatible” with certain monoids, in the sense that the results of computation done in $K \otimes M$ may be mapped to M , faithfully representing the aggregation results. Interestingly, compatibility is aligned with common wisdom: it is known that some (idempotent) aggregation functions such as MIN and MAX work well for set relations, while SUM and PROD require the treatment of relations as bags. We show that non-idempotent monoids are compatible only with “bag” semirings, i.e. semirings from which there exists an homomorphism to \mathbb{N} .

In general, aggregation results may be used by the query as the input to further operators, such as value-based joins or selections. Here the formal representation of values leads to a difficulty: the truth values of comparison operators on such formal expressions is undetermined! Consequently, we extend our framework and construct semirings in which formal comparison expressions between elements of the corresponding semimodule are elements of the semiring. This means that an expression like $[p_1 \otimes 20 = p_2 \otimes 10 + p_3 \otimes 15]$ may be used as part of the provenance of a tuples in the join result. This expression is simply treated as a new provenance token (with constraints), until p_1, p_2, p_3 are assigned e.g. values from \mathbb{B} or \mathbb{N} , in which case we can interpret

³As shown in [32], for list collections it also makes sense to consider non-commutative aggregations.

⁴ K -annotated relations with *union* (see Section 2.1) also form such a structure.

both sides of the equality as elements of the monoid and determine the truth value of the equality (see Section 4). We show in Section 4 that this construction allows us to manage provenance information for arbitrary queries with aggregation, while *keeping the representation size polynomial in the size of the input database*. Our construction is robust: if further queries are applied, the token $[p_1 \otimes 20 = p_2 \otimes 10 + p_3 \otimes 15]$ can be used as part of a more complex expression, just as any other provenance token.

The main result of this paper is providing, for the first time, a semantics for aggregation (including group by) on semiring-annotated relations that:

- Coincides with the usual semantics on set/bag relations for min/max/sum/prod.
- Commutes with homomorphisms of semirings (hence all the ensuing applications)
- Is representable with only poly-size overhead.

A second result of this paper is a new semantics for *difference* on relations annotated with elements from *any* commutative semiring. This is done via an encoding of relational difference using nested aggregation. The fact that such an encoding can be done is known (see e.g. [29, 9]), but combined with our provenance framework, the encoding gives a semantics for “provenance-aware” difference. Our new semantics for $R - S$ is a hybrid of bag-style and set-style semantics, in the sense that tuples of R that appear in S do not appear in $R - S$ (i.e. a boolean negative condition is used), while those that do not appear in S appear in $R - S$ with the same annotation (multiplicity, if $K = \mathbb{N}$ is used) that they had in R . This makes the semantics different from the bag-monus semantics and its generalization to “monus-semirings” in [19] as well as from the “negative multiplicities” semantics in [22] (more discussion in Section 6). We examine the equational laws entailed by this new semantics, in contrast to those of previously proposed semantics for difference. In our opinion, this semantics is probably not the last word on difference of annotated relations, but we hope that it will help inform and calibrate future work on the topic.

Paper Organization. The rest of the paper is organized as follows. Section 2 describes and exemplifies the main mathematical ingredients used throughout the paper. Section 3 describes our proposed framework for “simple” aggregation queries, and this framework is extended in Section 4 to nested aggregation queries. We consider difference queries in Section 5. Related Work is discussed in Section 6, and we conclude in Section 7.

2. PRELIMINARIES

We provide in this section the algebraic foundations that will be used throughout the paper. We start by recalling the notion of semiring and its use in [24] to capture provenance for the SPJU algebra queries. We then consider aggregates, and show the new algebraic construction that is required to accurately support it.

2.1 Semirings and SPJU

We briefly review the basic framework introduced in [24]. A *commutative monoid* is an algebraic structures $(M, +_M, 0_M)$ where $+_M$ is an associative and commutative binary operation and 0_M is an identity for $+_M$. A monoid homomorphism is a mapping $h : M \rightarrow M'$ where M, M' are monoids, and $h(0_M) = 0_{M'}, h(a + b) = h(a) + h(b)$. We will consider database operations on relations whose tuples

are annotated with elements from *commutative semirings*. These are structures $(K, +_K, \cdot_K, 0_K, 1_K)$ where $(K, +_K, 0_K)$ and $(K, \cdot_K, 1_K)$ are commutative monoids, \cdot_K is distributive over $+_K$, and $a \cdot_K 0_K = 0 \cdot_K a = 0_K$. A semiring homomorphism is a mapping $h : K \rightarrow K'$ where K, K' are semirings, and $h(0_K) = 0_{K'}, h(1_K) = 1_{K'}, h(a + b) = h(a) + h(b), h(a \cdot b) = h(a) \cdot h(b)$. Examples of commutative semirings are any commutative ring (of course) but also any distributive lattice, hence any boolean algebra. Examples of particular interest to us include the boolean semiring $(\mathbb{B}, \vee, \wedge, \perp, \top)$ (for usual set semantics), the natural numbers semiring $(\mathbb{N}, +, \cdot, 0, 1)$ (its elements are multiplicities, i.e., annotations that give bag semantics), and the *security* semiring $(\mathbb{S}, \min, \max, 0_{\mathbb{S}}, 1_{\mathbb{S}})$ where \mathbb{S} is the ordered set, $1_{\mathbb{S}} < C < S < T < 0_{\mathbb{S}}$ whose elements have the following meaning when used as annotations: $1_{\mathbb{S}}$: public (“always available”), C : confidential, S : secret, T : top secret, and $0_{\mathbb{S}}$ means “never available”.

Certain semirings play an essential role in capturing provenance information. Given a set X of *provenance tokens* which correspond to “atomic” provenance information, e.g., tuple identifiers, the semiring of *polynomials* $(\mathbb{N}[X], +, \cdot, 0, 1)$ was shown in [24] to adequately, and most generally, capture provenance for positive relational queries. The *provenance interpretation* of the semiring structure is the following. The $+$ operation on annotations corresponds to *alternative use* of data, the \cdot operation to *joint use* of data, 1 annotates data that is always and unrestrictedly available, and 0 annotates absent data. The definition of the K -relational algebra (see below for union, projection and join) fits indeed this interpretation. Algebraically, $\mathbb{N}[X]$ is the commutative semiring freely generated by X , i.e., for any other commutative semiring K , any valuation of the provenance tokens $X \rightarrow K$ extends uniquely to a semiring homomorphism $\mathbb{N}[X] \rightarrow K$ (an evaluation in K of the polynomials). We say that any semiring annotation semantics *factors* through the provenance polynomials semantics, which means that for practical purposes storing provenance information suffices for many other applications too. Other semirings can also be used to capture certain forms of provenance, albeit less generally than $\mathbb{N}[X]$ [24, 21]. For example, boolean expressions capture enough provenance to serve in the intensional semantics of queries on incomplete [28] and probabilistic data [18, 40].

To define annotated relations we use here the named perspective of the relational model [1]. Fix a countably infinite *domain* \mathbb{D} of values (constants). For any finite set U of attributes a tuple is a function $t : U \rightarrow \mathbb{D}$ and we denote the set of all such possible tuples by \mathbb{D}^U . Given a commutative semiring K , a K -*relation* (with schema U) is a function $R : \mathbb{D}^U \rightarrow K$ whose *support*, $\text{supp}(R) = \{t \mid R(t) \neq 0_K\}$ is *finite*. For a fixed set of attributes U we denote by $K\text{-Rel}$ (when U is clear from the context) the set of K -relations with schema U . We also define a K -*set* to be a function $S : \mathbb{D} \rightarrow K$ again of finite support. We then define:

Union If $R_i : \mathbb{D}^U \rightarrow K, i = 1, 2$ then $R_1 \cup_K R_2 : \mathbb{D}^U \rightarrow K$ is defined by $(R_1 \cup_K R_2)(t) = R_1(t) +_K R_2(t)$. The definition of union of K -sets follows similarly.

We also define the *empty* K -relation (K -set) by $\emptyset_K(t) = 0_K$. It is easy to see that $(K\text{-Rel}, \cup_K, \emptyset_K)$ is a commutative monoid ⁵. Similarly, we get the commutative monoid of K -sets $(K\text{-Set}, \cup_K, \emptyset_K)$.

Given a named relational schema, K -databases are defined from K -relations just as relational databases are defined

⁵In fact, it also has a semiring structure.

from usual relations, and in fact the usual (set semantics) databases correspond to the particular case $K = \mathbb{B}$. The (positive) K -relational algebra defined in [24] corresponds to a semantics on K -databases for the usual operations of the relational algebra. We have already defined the semantics of union above and we give here just two other cases leaving the rest for Appendix A (for a tuple t and an attributes set U' , $t|_{U'}$ is the restriction of t to U'):

Projection If $R : \mathbb{D}^U \rightarrow K$ and $U' \subseteq U$ then $\Pi_{U'} R : \mathbb{D}^{U'} \rightarrow K$ is defined by $(\Pi_{U'} R)(t) = \sum_K R(t')$ where the $+_K$ sum is over all $t' \in \text{supp}(R)$ such that $t'|_{U'} = t$.

Natural Join If $R_i : \mathbb{D}^{U_i} \rightarrow K$, $i = 1, 2$ then $R_1 \bowtie R_2 : \mathbb{D}^{U_1 \cup U_2} \rightarrow K$ is defined by $(R_1 \bowtie R_2)(t) = R_1(t_1) \cdot_K R_2(t_2)$ where $t_i = t|_{U_i}$, $i = 1, 2$.

2.2 Semimodules and aggregates

We will consider aggregates defined by commutative monoids. Some examples are SUM = $(\mathbb{R}, +, 0)$ for summation⁶, MIN = $(\mathbb{R}^{\pm\infty}, \min, +\infty)$ for min, MAX = $(\mathbb{R}^{\pm\infty}, \max, -\infty)$ for max, and PROD = $(\mathbb{R}, \times, 1)$ for product.

In dealing with aggregates we have to extend the operation of a commutative monoid to operations on relations annotated with elements of semirings. This interaction will be captured by *semimodules*.

DEFINITION 2.1. *Given a commutative semiring K , a structure $(W, +_W, 0_W, *_W)$ is a K -semimodule if $(W, +_W, 0_W)$ is a commutative monoid and $*_W$ is a binary operation $K \times W \rightarrow W$ such that (for all $k, k_1, k_2 \in K$ and $w, w_1, w_2 \in W$):*

$$k *_W (w_1 +_W w_2) = k *_W w_1 +_W k *_W w_2 \quad (1)$$

$$k *_W 0_W = 0_W \quad (2)$$

$$(k_1 +_K k_2) *_W w = k_1 *_W w +_W k_2 *_W w \quad (3)$$

$$0_K *_W w = 0_W \quad (4)$$

$$(k_1 \cdot_K k_2) *_W w = k_1 *_W (k_2 *_W w) \quad (5)$$

$$1_K *_W w = w \quad (6)$$

In any (commutative) monoid $(M, +_M, 0_M)$ define for any $n \in \mathbb{N}$ and $x \in M$

$$nx = x +_M \dots +_M x \quad (n \text{ times})$$

in particular $0x = 0_M$. Thus M has a canonical structure of \mathbb{N} -semimodule. Moreover, it is easy to check that a commutative monoid M is a \mathbb{B} -semimodule if and only if its operation is *idempotent*: $x +_M x = x$. The K -relations themselves form a K -semimodule $(K\text{-Rel}, \cup_K, \emptyset_K, *_K)$ where $(k *_K R)(t) = k \cdot_K R(t)$ ⁷.

We now show, for any K -semimodule W , how to define W -aggregation of a K -set of elements from W . We assume that $W \subseteq \mathbb{D}$ and that we have just one attribute, whose values are all from W . Consider the K -set S such that $\text{supp}(S) = \{w_1, \dots, w_n\}$ and $S(w_i) = k_i \in K$, $i = 1, \dots, n$ (i.e., each w_i is annotated with k_i). Then, the result of W -aggregating S is defined as

$$\text{SetAgg}_W(S) = k_1 *_W w_1 +_W \dots +_W k_n *_W w_n \in W$$

For the empty K -set we define $\text{SetAgg}_W(\emptyset_K) = 0_W$. Clearly, SetAgg_W is a *semimodule homomorphism*⁸. Since all commutative monoids are \mathbb{N} -semimodules this gives the usual

⁶COUNT is particular case of summation and AVG is obtained from summation and COUNT.

⁷In fact, it is the K semimodule freely generated by \mathbb{D}^U .

⁸In fact, it is the free homomorphism determined by the identity function on W .

sum, prod, min, and max aggregations on bags. Since MIN and MAX are \mathbb{B} -semimodules this gives the usual min and max aggregation on sets⁹.

Note that SetAgg is an operation on sets, not an operation on relations. In the sequel we show how to extend it to one.

2.3 A tensor product construction

More generally, we want to investigate M -aggregation on K -relations where M is a commutative monoid and K is some commutative semiring. Since M may not have enough elements to represent K -annotated aggregations we construct a K -semimodule in which M can be embedded, by transferring to semimodules the basic idea behind a standard algebraic construction, as follows.

Let K be any commutative semiring and M be any commutative monoid. We start with $K \times M$, denote its elements $k \otimes m$ instead of $\langle k, m \rangle$ and call them “simple tensors”. Next we consider (finite) bags of such simple tensors, which, with bag union and the empty bag, form a commutative monoid. It will be convenient to denote bag union by $+_{K \otimes M}$, the empty bag by $0_{K \otimes M}$ and to abuse notation denoting singleton bags by the unique element they contain. Then, every non-empty bag of simple tensors can be written (repeating summands by multiplicity) $k_1 \otimes m_1 +_{K \otimes M} \dots +_{K \otimes M} k_n \otimes m_n$. Now we define

$$k \ *_{K \otimes M} \sum k_i \otimes m_i = \sum (k \cdot_K k_i) \otimes m_i$$

Let \sim be the smallest congruence w.r.t. $+_{K \otimes M}$ and $*_{K \otimes M}$ that satisfies (for all k, k', m, m'):

$$(k +_K k') \otimes m \sim k \otimes m +_{K \otimes M} k' \otimes m$$

$$0_K \otimes m \sim 0_{K \otimes M}$$

$$k \otimes (m +_M m') \sim k \otimes m +_{K \otimes M} k \otimes m'$$

$$k \otimes 0_M \sim 0_{K \otimes M}$$

We denote by $K \otimes M$ the set of *tensors* i.e., equivalence classes of bags of simple tensors modulo \sim . We show in Appendix B that $K \otimes M$ forms a K -semimodule.

Lifting homomorphisms. Given a homomorphism of semirings $h : K \rightarrow K'$, and some commutative monoid M , we can “lift” h to a homomorphism of monoids in a natural way. The lifted homomorphism is denoted $h^M : K \otimes M \rightarrow K' \otimes M$ and defined by:

$$h^M(\sum k_i \otimes m_i) = \sum h(k_i) \otimes m_i$$

3. SIMPLE AGGREGATION QUERIES

In this section we begin our study of the “provenance-aware” evaluation of aggregate queries, focusing on “simple” such queries, that is, queries in which aggregations are done last; for example, un-nested SELECT FROM WHERE GROUP BY queries. This avoids the need to compare values which are the result of annotated aggregations and simplifies the treatment. The restriction is relaxed in the more general framework presented in Section 4.

The section is organized as follows. We list the desired features of a provenance-aware semantics for aggregation, and first try to design a semantics with these features, *without* using the tensor product construction, i.e. by simply

⁹The fact that the right algebraic structure to use for aggregates is that of *semimodules* can be justified in the same way in which using semirings was justified in [24]: by showing how the laws of semimodules follow from desired equivalences between aggregation queries.

working with K -relations as done in [24]. We show that this is impossible. Consequently, we turn to semantics that are based on combining aggregation with values via the tensor product construction. We propose such semantics that do satisfy the desired features, first for relational algebra with an additional AGG operator on relations (that allows aggregation of all values in a chosen attributes, but no grouping); and then for GROUP BY queries.

3.1 Semantic desiderata and first attempts

We next explain the desired features of a provenance-aware semantics for aggregation. To illustrate the difficulties and the need for a more complex construction, we will first attempt to define a semantics on K -relations, without using the tensor product construction of Section 2.3.

We consider a commutative semiring K (e.g., $\mathbb{B}, \mathbb{N}, \mathbb{N}[X], \mathbb{S}$, etc.) for tuple annotations and a commutative monoid M (e.g., $\text{SUM} = (\mathbb{R}, +, 0)$, $\text{PROD} = (\mathbb{R}, \times, 1)$, $\text{MAX} = (\mathbb{R}^{-\infty}, \max, -\infty)$, $\text{MIN} = (\mathbb{R}^{\infty}, \min, \infty)$ etc.) for aggregation. We will assume that the elements of M already belong to the database domain, $M \subseteq \mathbb{D}$.

We have recalled the semantics of SPJU queries in Section 2.1. Now we wish to add an M -aggregation operation AGG on relations. We then denote by SPJU-A the restricted class of queries consisting of any SPJU-expression followed possibly by just one application of AGG. This corresponds to SELECT AGG(*) FROM WHERE queries (no grouping).

For the moment, we do not give a concrete semantics to $\text{AGG}_M(R)$, allowing any possible semantics where the result of $\text{AGG}_M(R)$ is a K -relation. We note that $\text{AGG}_M(R)$ should be defined iff R is a K -relation with one attribute whose values are in M .

What properties do we expect of a reasonable semantics for SPJU-A (including, of course, a semantics for $\text{AGG}_M(R)$)? A basic sanity check is

Set/Bag Compatibility The semantics coincides with the usual one when $K = \mathbb{B}$ (sets) and $M = \text{MAX}$ or MIN , and when $K = \mathbb{N}$ (bags) and $M = \text{SUM}$ or PROD .

Note that we associate min and max with sets and sum and product with bags. Min and max work fine with bags too, but we get the same result if we convert a bag to a set (eliminate duplicates) and then apply them. Sum and product (in the context of other operations such as projection) require us to use bags semantics in order to work properly. This is well-known, but our general approach sheds further light on the issue by discussing such “compatibility” for arbitrary semirings and monoids in Section 3.4.

As discussed in the introduction, a fundamental desideratum with many applications is *commutation with homomorphisms*. Note that a semiring homomorphism $h : K \rightarrow K'$ naturally extends to a mapping $h_{\text{Rel}} : K\text{-Rel} \rightarrow K'\text{-Rel}$ via $h_{\text{Rel}}(R) = h \circ R$ (i.e. the homomorphism is applied on the annotation of every tuple), which then further extends to K -databases. With this, the second desired property is

Commutation with Homomorphisms Given any two commutative semirings K, K' and any homomorphism $h : K \rightarrow K'$, for any query Q , its semantics on K -databases and on K' -databases satisfy $h_{\text{Rel}}(Q(D)) = Q(h_{\text{Rel}}(D))$ for any K -database D .

It turns out that this property determines quite precisely the way in which tuple annotations are defined. We say that the semantics of an operation Ω on K -databases is *algebraically uniform* with respect to the class of commutative semirings if

the annotations of the output $\Omega(D)$ are defined by the same (for all K) $\{+_K, \cdot_K, 0_K, 1_K\}$ -expressions, where the elements in the expressions are the annotations of the input D . One can see that the definition of the SPJU-algebra is indeed algebraically uniform and was shown in [24, 17] to commute with homomorphisms. The connection between the two properties is general (proof deferred to the Appendix):

PROPOSITION 3.1. *A semantics commutes with homomorphisms iff it is algebraically uniform.*

After stating two of the desired properties, namely set/bag compatibility and commutation with homomorphisms we can already show that it is not possible to give a satisfactory semantics to the SPJU-A algebra within the framework used in [24] for the SPJU-algebra.

PROPOSITION 3.2. *There is no K -relation semantics for MAX-(or MIN-)aggregation that is both set-compatible and commutes with homomorphisms. Similarly, there is no K -relation semantics for SUM-aggregation that is both bag-compatible and commutes with homomorphisms.*

PROOF. Assume by contradiction the existence of such semantics. Consider the $\mathbb{N}[X]$ -relation R with one attribute and two tuples with values 10 and 20, with the corresponding tuple annotations being $x, y \in X$. Let R' be $\text{AGG}_{\text{MAX}}(R)$ according to the assumed semantics; R' is also an $\mathbb{N}[X]$ -relation. Because a tuple t with a value 10 is a possible answer to the MAX-aggregation (when we set $y = 0$) it must occur in $\text{supp}(R')$. Let $p \in \mathbb{N}[X]$ be the annotation of the tuple t (having value 10) in R' . By algebraic uniformity the only variables that can occur in p are x and y , and we consequently denote it $p(x, y)$. Consider two homomorphisms $h', h'' : \mathbb{N}[X] \rightarrow \mathbb{B}$ defined by $h'(x) = h'(y) = \top$ and by $h''(x) = \top, h''(y) = \perp$. Applying AGG_{MAX} to $h'_{\text{Rel}}(R)$ and $h''_{\text{Rel}}(R)$ should, by set-compatibility, work as usual. Hence, by commutation with homomorphisms $h'(p) = \perp$ and $h''(p) = \top$. Functions on \mathbb{B} defined by polynomials in $\mathbb{N}[X]$ are monotone in each variable. But $\perp = h'(p(x, y)) = p(h'(x), h'(y)) = p(\top, \top)$ and $\top = h''(p(x, y)) = p(h''(x), h''(y)) = p(\top, \perp)$, in contradiction to the monotonicity.

□

Alternatively, one may consider going beyond semirings, to algebraic structures with additional operations. We have briefly explored the use of “negative” information in the introduction. As we show there, one could use the ring structure on $\mathbb{Z}[X]$ (the additional subtraction operation) or the boolean algebra structure on $\text{BoolExp}(X)$ (the additional complement operation) but the use of negative operation does not avoid the need to enumerate in separate tuples of the answer all the possible aggregation results given by subsets of the input. In the case of summation, at least, there are exponentially many such tuples. We reject such an approach and we state as an additional desideratum:

Poly-Size Overhead For any query Q and database D , the size of $Q(D)$, including annotations, should be only polynomial in the size of D .

We shall next show a semantics to the SPJU-A -algebra that satisfies all three properties we have listed.

3.2 Annotations \otimes values and SPJU-A

Let us fix a commutative monoid M (for aggregation) and a commutative semiring K (for annotation). The *inputs*

of our queries are as before: K -databases whose domain \mathbb{D} includes the values M over which we aggregate. However, the outputs are more complicated. The basic idea for the semantics of aggregation was already shown in Section 2.2 where it is assumed that the domain of aggregation has a K -semimodule structure. As we have shown in Section 2.3, we can give a tensor product construction that embeds M in the K -semimodule $K \otimes M$ (note that this embedding is not always *faithful*, as discussed in Section 3.4).

For the output relations of our algebra queries, we thus need results of aggregation (i.e., the elements of $K \otimes M$) to also be part of the domain out of which the tuples are constructed. Thus for the output domain we will assume that $K \otimes M \subseteq \mathbb{D}$, i.e. the result “combines annotations with values”. The elements of M (e.g., real numbers for sum or max aggregation) are still present, but only via the embedding $\iota : M \rightarrow K \otimes M$ defined by $\iota(m) = 1_K \otimes m$.

Having annotations from K appear in the values will change the way in which we apply homomorphisms to query results, so to emphasize the change we will call (M, K) -relations the K -annotated relations over such that the data domain \mathbb{D} that includes $K \otimes M$. To summarize, the semantics of the SPJU-A -algebra will map databases of K -relations (with $M \subseteq \mathbb{D}$) to (M, K) -relations (with $K \otimes M \subseteq \mathbb{D}$).

As we define the semantics of the SPJU-A -algebra, we first note that for selection, projection, join and union the definition is the same as for the SPJU-algebra on K -databases. The last step of the query is aggregation, denoted $\text{AGG}_M(R)$, and is well-defined iff R is a K -relation with one attribute whose values are in the M subset of \mathbb{D} . To apply the definition that uses the semimodule structure (shown in Section 2.2), we convert R to an (M, K) -relation $\iota(R)$ by replacing each $m \in M$ with $\iota(m) = 1_K \otimes m \in K \otimes M$. Then, if $\text{supp}(R) = \{m_1, \dots, m_n\}$ and $R(m_i) = k_i \in K, i = 1, \dots, n$ (i.e., each m_i is annotated with k_i) we define $\text{AGG}_M(R)$ as a one-attribute relation with one tuple annotation is 1_K and whose content is $\text{SetAgg}_{K \otimes M}(\iota(R))$, which is equal to

$$\begin{aligned} k_1 *_{K \otimes M} \iota(m_1) +_{K \otimes M} \dots +_{K \otimes M} k_n *_{K \otimes M} \iota(m_n) \\ = k_1 \otimes m_1 +_{K \otimes M} \dots +_{K \otimes M} k_n \otimes m_n \end{aligned}$$

We define the annotation of the only tuple in the output of AGG_M to be 1_K , since this tuple is always available. However, the *content* of this tuple does depend on R . For example, even when R is empty the output is not empty: by the semimodule laws, its content is $0_{K \otimes M} = \iota(0_M)$.

Commutation with Homomorphisms. We have explained in Section 2.3 how to lift a homomorphism $h : K \rightarrow K'$ to a homomorphism $h^M : K \otimes M \rightarrow K' \otimes M$. Via this we can lift h to a homomorphism h_{Rel} on (M, K) -relations: let R be such a relation and recall that some values in R are elements of $K \otimes M$, and the annotations of these tuples are elements of K . Then $h_{\text{Rel}}(R)$ denotes the relation obtained from R by replacing every $k \in K$ with $h(k)$, and additionally replacing every $k \otimes M \in K \otimes M$ with $h^M(k \otimes m)$. All other values in R stay intact. Applying h_{Rel} on a (M, K) -database D amounts to applying h_{Rel} on each (M, K) -relation in D .

We can now state the main result for our SPJU-A -algebra:

THEOREM 3.3. *Let K, K' be semirings, $h : K \rightarrow K', Q$ an SPJU-A query and let M be a commutative monoid. For every (M, K) -database D , $Q(h_{\text{Rel}}(D)) = h_{\text{Rel}}(Q(D))$ if and only if h is a semiring homomorphism.*

The proof is by induction on the query structure, and is straightforward given that for the constructs of SPJU

queries homomorphism commutation was shown in [24], while commutation for the new AGG_M construct follows directly from the definition.

EXAMPLE 3.4. *Consider the following $\mathbb{N}[X]$ -relation R :*

Sal	
20	r_1
10	r_2
30	r_3

Let M be some commutative monoid, then $\text{AGG}_M(R)$ consist of a single tuple with value $r_1 \otimes 20 +_{K \otimes M} r_2 \otimes 10 +_{K \otimes M} r_3 \otimes 30$. The intuition is that this value captures multiple possible aggregation values, each of which may be obtained by mapping the r_i annotations to \mathbb{N} , standing for the multiplicity of the corresponding tuple. The commutation with homomorphism allows us to first evaluate the query and only then map the r_i 's, changing directly the expression in the query result. For example, if $M = \text{SUM}$ and we map r_1 to 1, r_2 to 0, r_3 to 2, we obtain $1 \otimes 20 +_{K \otimes M} 2 \otimes 30 = 1 \otimes 20 +_{K \otimes M} 1 \otimes 30 +_{K \otimes M} 1 \otimes 30 = 1 \otimes 80$ (which corresponds to the M element 80). As another example, the commutation with homomorphisms allows us to propagate the deletion of the first tuple in R , by simply setting in the aggregation result $r_1 = 0$ (keeping the other annotations intact) and obtaining $2 \otimes 30 = (1+1) \otimes 30 = 1 \otimes 30 + 1 \otimes 30 = 1 \otimes (30+30) = 1 \otimes 60$.

We further demonstrate an application for *security*.

EXAMPLE 3.5. *Consider the following relation R , annotated by elements from the security semiring \mathbb{S} .*

Sal	
20	\mathbb{S}
10	$1_{\mathbb{S}}$
30	\mathbb{S}

Recall (from Section 2.1) the order relation $1_{\mathbb{S}} < \mathbb{C} < \mathbb{S} < \mathbb{T} < 0_{\mathbb{S}}$; a user with credentials cred can only view tuples annotated with security level equal or less than cred . Now let $M = \text{MAX}$ and we obtain: $\text{AGG}_{\text{MAX}}(R) = \mathbb{S} \otimes 20 +_{K \otimes M} 1_{\mathbb{S}} \otimes 10 +_{K \otimes M} \mathbb{S} \otimes 30 = \mathbb{S} \otimes (20 +_{\text{MAX}} 30) + 1_{\mathbb{S}} \otimes 10$ and we get $\mathbb{S} \otimes 30 + 1_{\mathbb{S}} \otimes 10$.

Assume now that we wish to compute the query results as viewed by a user with security credentials cred . A naive computation would delete from R all tuples that require higher credentials, and re-evaluate the query (which in general may be complex). But observe that the deletion of tuples is equivalent to applying to R a homomorphism that maps every annotation $t > \text{cred}$ to 0, and $t \leq \text{cred}$ to 1. Using homomorphism commutation we can do better by applying this homomorphism only on the result representation (namely $\mathbb{S} \otimes 30 + 1_{\mathbb{S}} \otimes 10$). For example, for a user with credentials \mathbb{C} , we map \mathbb{S} to 0 and $1_{\mathbb{S}}$ to 1, and obtain $0 \otimes 30 + 1 \otimes 10 = 1 \otimes 10$; similarly for a user with credentials \mathbb{S} we get $1 \otimes 30 + 1 \otimes 10 = 1 \otimes (30 +_{\text{MAX}} 10) = 1 \otimes 30$.

From the above definition of the semantics for aggregation, it is obvious that the poly-size overhead property is fulfilled. Indeed, consider the case of provenance for summation as in Example 3.4, and compare it to the naive representation provided in the Introduction. Instead of having to list all (exponentially many) options for the sum of salaries, we used an expression in $K \otimes \text{SUM}$ that is of linear size with respect to the input to the aggregation. As exemplified, the possible aggregate answers now correspond to different valuations for the provenance tokens, applied to this expression.

3.3 Group By

So far we have considered aggregation in a limited context, where the input relation contains a single attribute. In common cases, however, aggregation is used on arbitrary relations and in conjunction with grouping, so we next extend the algebra to handle such an operation. The idea behind the construction is quite simple: we separately group the tuples according to the values of their “group-by” attributes, and the aggregated values for each such group are computed similarly to the computation for the *AGG* operator. When considering the *annotation* of the aggregated tuple, we encounter a technical difficulty: we want this annotation to be equal 1_K if the input relation includes at least one tuple in the corresponding group, and 0_K otherwise (for intuition, consider the case of bag relations, in which the aggregated result can have at most multiplicity 1); we consequently enrich our structure to include an additional construct δ that will capture that, as follows:

DEFINITION 3.6. *A (commutative) δ -semiring is an algebraic structure $(K, +_K, \cdot_K, 0_K, 1_K, \delta_K)$ where $(K, +_K, \cdot_K, 0_K, 1_K)$ is a commutative semiring and $\delta_K : K \rightarrow K$ is a unary operation satisfying the “ δ -laws” $\delta_K(0_K) = 0_K$ and $\delta_K(n1_K) = 1_K$ for all $n \geq 1$. If K and K' are δ -semirings then a homomorphism between them is a semiring homomorphism $h : K \rightarrow K'$, for which we have in addition $h(\delta_K(k)) = \delta_{K'}(h(k))$.*

The δ -laws completely determine $\delta_{\mathbb{B}}$ and $\delta_{\mathbb{N}}$. But they leave a lot of freedom for the definition of δ_K in other semirings; in particular for the security semiring, a reasonable choice for δ_S is the identity function.

As with any equational axiomatization, we can construct the commutative δ -semiring *freely generated* by a set X , denoted $\mathbb{N}[X, \delta]$, by taking the quotient of the set of $\{+, \cdot, 0, 1, \delta\}$ -algebraic expressions by the congruence generated by the equational laws of commutative semirings and the δ -laws. For example, if e and e' are elements of $\mathbb{N}[X, \delta]$ (i.e., congruence classes of expressions given by some representatives) then $e +_{\mathbb{N}[X, \delta]} e'$ is the congruence class of the expression $e + e'$. The elements of $\mathbb{N}[X, \delta]$ are not standard polynomials but certain subexpressions can be put in polynomial form, for example $\delta(2 + 3xy^2)$ or $3 + 2\delta(x^2 + 2y)z^2$.

We are now ready to define the group by (denoted *GB*) operation; subsequently we exemplify its use, including in particular the role of δ :

DEFINITION 3.7. *Let R be a K -relation on set of attributes U , let $U' \subseteq U$ be a subset of attributes that will be grouped and $U'' \subseteq U$ be the subset of attributes with values in M (to be aggregated). We assume that $U' \cap U'' = \emptyset$. For a tuple t , we define $T = \{t' \in \text{supp}(R) \mid \forall u \in U'' \ t'(u) = t(u)\}$.*

We then define the aggregation result $R' = GB_{U', U''}(R)$ as follows:

$$R'(t) = \begin{cases} \delta_K(\sum_{t' \in T} R(t')) & T \neq \emptyset, \text{ and} \\ & \forall u \in U'' \ t(u) = \sum_{t' \in T} R(t') \otimes t'(u) \\ 0 & \text{Otherwise.} \end{cases}$$

EXAMPLE 3.8. *Consider the relation R :*

Dept	Sal	
d ₁	20	r ₁
d ₁	10	r ₂
d ₂	10	r ₃

and a query $GB_{\{\text{Dept}\}, \text{Sal}}R$, where the monoid used is *SUM*. The result (denoted R') is:

Dept	Sal	
d ₁	$r_1 \otimes 20 +_K \otimes_{SUM} r_2 \otimes 10$	$\delta_K(r_1 +_K r_2)$
d ₂	$r_3 \otimes 10$	$\delta_K(r_3)$

Each aggregated value (for each department) is computed very similarly to the computation in Example 3.4. Consider the provenance annotation of the first tuple: intuitively, we expect it to be 1_K if at least one of the first two tuples of R exists, i.e. if at least one out of r_1 or r_2 is non-zero. Indeed the expression is $\delta(r_1 +_K r_2)$ and if we map r_1, r_2 to e.g. 2 and 1 respectively, we obtain $\delta_{\mathbb{N}}(3) = 1$.

We use SPJU-AGB as the name for relational algebra with the two new operators *AGG* and *GB*. We note that the poly-size overhead property is still fulfilled for queries in SPJU-AGB; commutation with homomorphism also extends to SPJU-AGB (see proof in the Appendix).

Recall that an additional desideratum from the semantics was bag / set compatibility. Recall that sets and bags are modeled by $K = \mathbb{N}$ and $K = \mathbb{B}$ respectively. We next study compatibility in a more general way, for arbitrary K and M .

3.4 Annotation-aggregation compatibility

The first desideratum we listed was an obvious sanity check: whatever semantics we define, when specialized to the familiar aggregates of max, min and summation, it should produce familiar results. Since we had to take an excursion through the tensor product $K \otimes M$, this familiarity is not immediate. However, the following proposition holds (its correctness will follow from theorems 3.12 and 3.13).

PROPOSITION 3.9. *In the following constructions: $\mathbb{B} \otimes \text{MAX}$, $\mathbb{B} \otimes \text{MIN}$, and $\mathbb{N} \otimes \text{SUM}$, $\iota : M \rightarrow K \otimes M$ where $\iota(m) = 1_K \otimes m$ is a monoid isomorphism.*

and this means our semantics satisfies the set/bag compatibility property because in these cases computing in $K \otimes M$ exactly mirrors computing in M .

But of course, we are also interested in working with other semirings, in particular the provenance semiring, for which $\mathbb{N}[X] \otimes M$ and M are in general not isomorphic (in particular, ι is not surjective and thus not an isomorphism). In fact, the whole point of working in $\mathbb{N}[X] \otimes \text{MAX}$, for example, is to add annotated aggregate computations to the domain of values. Most of these do not correspond to actual real numbers as e.g. $\iota(\text{MAX})$ is a strict subset of $\mathbb{N}[X] \otimes \text{MAX}$ (and similarly $\iota(\text{SUM})$ is a strict subset of $\mathbb{N}[X] \otimes \text{SUM}$ etc.). However, when provenance tokens are valued to obtain set (or bag) instances, we can go back into $\iota(\text{MAX})$ (or $\iota(\text{SUM})$ etc.), and then we should obtain familiar results by “stripping off” the ι . It turns out that this works correctly with $\mathbb{N}[X]$ but not necessarily with arbitrary commutative semirings K . The reason is that not only that ι is not an isomorphism, but in general it may be *unfaithful* (not injective). Indeed $\iota : \text{SUM} \rightarrow \mathbb{B} \otimes \text{SUM}$ is not injective:

$$\begin{aligned} \iota(4) &= \iota(2 + 2) = \iota(2) +_{K \otimes M} \iota(2) = \top \otimes 2 +_{K \otimes M} \top \otimes 2 = \\ &= (\top \vee \top) \otimes 2 = \top \otimes 2 = \iota(2) \end{aligned}$$

This is not surprising, as it is related to the well-known difficulty of making summation work properly with set semantics. In general, we thus define compatibility as follows:

DEFINITION 3.10. *We say that a commutative semiring K and a commutative monoid M are compatible if ι is injective.*

The point of the definition is that when there is compatibility, we can work in $K \otimes M$ and whenever the results belong to $\iota(M)$, we can *safely* read them as familiar answers from M . We give three theorems that capture some general conditions for compatibility.

First, we note that if we work with a semiring in which $+_K$ is idempotent, such as \mathbb{B} or \mathbb{S} , a compatible monoid must also be idempotent (e.g. MIN or MAX but not SUM):

PROPOSITION 3.11. *Let K be some commutative semiring such that $+_K$ is idempotent, and let M be some commutative monoid. If M is compatible with K , then $+_M$ is idempotent.*

PROOF. $\iota(m) = 1_K \otimes m = (1_K +_K 1_K) \otimes m = 1_K \otimes m +_{K \otimes M} 1_K \otimes m = 1_K \otimes (m +_M m) = \iota(m +_M m)$ \square

Nicely enough, idempotent aggregations are compatible with every annotation semiring K that is *positive with respect to* $+_K$. K is said to be positive with respect to $+_K$ if $k +_K k' = 0_K \Rightarrow k = k' = 0_K$. For instance, \mathbb{B} , \mathbb{N} , \mathbb{S} and $\mathbb{N}[X]$ are such semirings (but not $(\mathbb{Z}, +, \cdot, 0, 1)$). The following theorem holds:

THEOREM 3.12. *If M is a commutative monoid such that $+_M$ is idempotent, then M is compatible with any commutative semiring K which is positive with respect to $+_K$.*

PROOF SKETCH. We define $h : K \otimes M \rightarrow M$ as $h(\sum_{i \in I} k_i \otimes m_i) = \sum_{j \in J} m_j$ where $J = \{j \in I | k_j \neq 0\}$. We can show that h is well-defined (details deferred to the Appendix); since $\forall m \in M \ h \circ \iota(m) = m$, ι is injective and thus K and M are compatible. \square

For general (and in particular non-idempotent) monoids (e.g. SUM) we identify a sufficient condition on K (which in particular holds for $\mathbb{N}[X]$), that allows for compatibility:

THEOREM 3.13. *Let K be a commutative semiring. If there exists a semiring homomorphism from K to \mathbb{N} then K is compatible with all commutative monoids.*

PROOF SKETCH. Let h' be a homomorphism from K to \mathbb{N} , and M be an arbitrary commutative monoid. We define a mapping $h : K \otimes M \rightarrow M$ by $h(\sum k_i \otimes m_i) = \sum h'(k_i) m_i$. We show in the Appendix that h is well-defined and that $h \circ \iota$ is the identity function hence ι is injective. \square

COROLLARY 3.14. *The semiring of provenance polynomials $\mathbb{N}[X]$ is compatible with all commutative monoids.*

Now consider the security semiring \mathbb{S} . It is idempotent, and therefore not compatible with non-idempotent monoids such as SUM. Still, we want to be able to use \mathbb{S} and other non-idempotent semirings, while allowing the evaluation of aggregation queries with non-idempotent aggregates. This would work if we could construct annotations that would allow us to use Theorem 3.13, in other words, if we could *combine* annotations from \mathbb{S} , with multiplicity annotations (i.e. annotations from \mathbb{N}). We explain next the construction of such a semiring \mathbb{SN} (for security-bag), and its compatibility with any commutative monoid M will follow from the existence of a homomorphism $h : \mathbb{SN} \rightarrow \mathbb{N}$.

Constructing a compatible semiring. We start with the semiring of polynomials $\mathbb{N}[\mathbb{S}]$, i.e. polynomials where instead of indeterminates (variables) we have members of \mathbb{S} , and the coefficients are natural numbers. Already $\mathbb{N}[\mathbb{S}]$ is compatible with any commutative monoid M , as there exists a homomorphism $h : \mathbb{N}[\mathbb{S}] \rightarrow \mathbb{N}$; but if we work with $\mathbb{N}[\mathbb{S}]$ we lose the

ability to use the identities that hold in \mathbb{S} and to thus reduce the size of annotations in query results. We can do better by taking the quotient of $\mathbb{N}[\mathbb{S}]$ by the smallest congruence containing the following identities:

- $\forall s_1, s_2 \in \mathbb{S} \ s_1 \geq s_2 \implies s_1 \cdot_{\mathbb{N}[\mathbb{S}]} s_2 = s_1$.
- $\forall c \in \mathbb{N}, s \in \mathbb{S} \ 0 \cdot_{\mathbb{N}[\mathbb{S}]} s = c \cdot_{\mathbb{N}[\mathbb{S}]} 0_{\mathbb{S}} = 0$.
- $\forall c \in \mathbb{N} \ c \cdot_{\mathbb{N}[\mathbb{S}]} 1_{\mathbb{S}} = c$.

We will denote the resulting quotient semiring by \mathbb{SN} . It is easy to check that the faithfulness of the embeddings of \mathbb{N} and \mathbb{S} in $\mathbb{N}[\mathbb{S}]$ is preserved by taking the quotient. Most importantly, \mathbb{SN} is still homomorphic to \mathbb{N} . Thus,

COROLLARY 3.15. *\mathbb{SN} is compatible with any commutative monoid M .*

EXAMPLE 3.16. *Consider the SUM monoid. Let R, S be the following \mathbb{S} -relations which by the embedding of \mathbb{S} we take as \mathbb{SN} -relations:*

A		A	
30	S	30	T
		10	1 _S

$R \qquad S$

Consider the query: $AGG(R \cup \Pi_{S,A}(S \bowtie R))$. Ignoring the annotations, the expected result (under bag semantics) is 70. Working within the (compatible) semantics defined by $\mathbb{SN} \otimes SUM$, the query result contains an aggregated value of $(T \cdot_{\mathbb{SN}} S +_{\mathbb{SN}} S) \otimes 30 + S \otimes 10$. We can further simplify this to $T \otimes 30 + S \otimes 30 + S \otimes 10 = T \otimes 30 + S \otimes 40$. This means that e.g. for a user with credentials T the query result is $1_{\mathbb{SN}} \otimes 70$, and we can use the inverse of ι to map it to \mathbb{N} and obtain 70. Similarly, for a user with credentials S , the query result is mapped to 40. These are indeed the expected results.

Note that if we would have used in the above example \mathbb{S} instead of \mathbb{SN} we would have $(T +_S S) = S$ so $(T +_S S) \otimes 30$ would be the same as $S \otimes 30$. For a user with credentials T we could either use this, leading to the result of $1_S \otimes 40$, or use the same computation done in the example, to obtain $1_S \otimes 70$. Indeed, in $\mathbb{S} \otimes SUM$, we have $1_S \otimes 40 = 1_S \otimes 70$. This is the same phenomenon demonstrated in the beginning of this subsection for \mathbb{B} , where ι is not injective, preventing us from stripping it away.

Note also that if we would have used $\mathbb{N}[\mathbb{S}]$ instead of \mathbb{SN} then we could not have done the illustrated simplifications.

4. NESTED AGGREGATION QUERIES

So far we have studied only queries where the aggregation operator is the last one performed. In this section we extend the discussion to queries that involve comparisons on aggregate values. We first demonstrate the difficulties that arise in designing an algebra for such queries, then explain how to extend the construction to overcome these difficulties.

Note. For simplicity, all results and examples are presented for queries in which the comparison operator is equality ($=$). However the results can easily be extended to arbitrary comparison predicates, that can be decided for elements of M .

4.1 Difficulties

We start by exemplifying where the algebra proposed for restricted aggregation queries, fails here:

EXAMPLE 4.1. Reconsider the relation (denoted R') which is the result of aggregation query, depicted in Example 3.8. Further consider a query Q_{select} that selects from R' all tuples for which the aggregated salary equals 20. The crux is that deciding the truth value of the selection condition involves interpreting the comparison operator on symbolic representation of values in R' ; so far, we have no way of interpreting the obtained comparison expression, for instance $r_1 \otimes 20 + r_2 \otimes 10$ “equals” 20, and thus we cannot decide the existence of tuples in the selection result.

Note that in the above example, the truth value of the comparison (and consequently the set of tuples in the query result) depends in a *non-monotonic* way on the existence of tuples in the (original) input relation R : note that if we map r_1 to 1 and r_2 to 0 then the tuple with dept. d_1 appears in the query result, but if we map both to 1, it does not. The challenge that this non-monotonicity poses is fundamental, and is encountered by any algebra on (M, K) -relations. The following proposition, which is the counterpart of proposition 3.2, holds (proof deferred to the Appendix):

PROPOSITION 4.2. *There is no (M, K) -relation semantics for nested aggregation queries with MAX-(or MIN-) aggregation that is both set-compatible and commutes with homomorphisms. Similarly for SUM-aggregation and bag-compatibility.*

Consequently, a more intricate construction is required for nested aggregation queries.

4.2 An Extended Structure

We start with an example of our treatment of nested aggregation queries, then give the formal construction.

EXAMPLE 4.3. Reconsider example 4.1, and recall that the challenge in query evaluation lies in comparing elements of $K \otimes M$ with elements of M (or $K \otimes M$, e.g. in case of joins). Our solution is to introduce to the semiring K new elements, of the form $[x = y]$ where $x, y \in K \otimes M$ (if we need to compare with $m \in M$, we use $\iota(m)$ instead). The result of evaluating the query in example 4.1 (using $M = \text{SUM}$) will then be captured by:

Dept	Sal	
d_1	$r_1 \otimes 20$ $+_{K \otimes M} r_2 \otimes 10$	$\delta(r_1 +_K r_2)_K$ $[r_1 \otimes 20 +_{K \otimes M} r_2 \otimes 10 = 1_K \otimes 20]$
d_2	$r_3 \otimes 10$	$\delta(r_3)_K [r_3 \otimes 10 = 1_K \otimes 20]$

Intuitively, since we do not know which tuples will satisfy the selection criterion, we keep both tuples and multiply the provenance annotation of each of them by a symbolic equality expression. These equality expressions are kept as symbols until we can embed the values in $M = \text{SUM}$ and decide the equality (e.g. if $K = \mathbb{N}$), in which case we “replace” it by 1_K if it holds or 0_K otherwise. For example, given a homomorphism $h : \mathbb{N}[X] \rightarrow \mathbb{N}$, $h(r_1) = h(r_2) = 1$, then $h^M(r_1 \otimes 20 +_{K \otimes M} r_2 \otimes 10) = h(r_1) \otimes 20 +_{K \otimes M} h(r_2) \otimes 10 = 1 \otimes 30 \neq 1 \otimes 20$, thus the equality expression is replaced with (i.e. mapped by the homomorphism to) 0_K .

We next define the construction formally; the idea underlying the construction is to define a semiring whose elements are polynomials, in which equation elements are additional indeterminates. To achieve that, we introduce for any semiring K and any commutative monoid M , the “domain” equation $\widehat{K} = \mathbb{N}[K \cup \{[c_1 = c_2] \mid c_1, c_2 \in \widehat{K} \otimes M\}]$.

The right-hand-side is a monotone, in fact continuous w.r.t. the usual set inclusion operator, hence this equation has a set-theoretic least solution (no need for order-theoretic domain theory). The solution also has an obvious commutative semiring structure induced by that of polynomials. The solution semiring is $\widehat{K} = (X, +_{\widehat{K}}, \cdot_{\widehat{K}}, 0_{\widehat{K}}, 1_{\widehat{K}})$, and we continue by taking the quotient on \widehat{K} defined by the following axioms.

For all $k_1, k_2 \in K, c_1, c_2, c_3, c_4 \in \widehat{K} \otimes M$:

$$\begin{aligned} 0_{\widehat{K}} &\sim 0_K \\ 1_{\widehat{K}} &\sim 1_K \\ k_1 +_{\widehat{K}} k_2 &\sim k_1 +_K k_2 \\ k_1 \cdot_{\widehat{K}} k_2 &\sim k_1 +_K k_2 \\ [c_1 = c_3] &\sim [c_2 = c_4] \text{ (if } c_1 = \widehat{K} \otimes M c_2, c_3 = \widehat{K} \otimes M c_4) \end{aligned}$$

and if K and M are such that ι defined by $\iota(m) = 1_K \otimes m$ is an isomorphism (and let h be its inverse), we further take the quotient defined by: for all $a, b \in K \otimes M$,

$$\begin{aligned} (*) [a = b] &\sim 1_K \text{ (if } h(a) =_M h(b)) \\ [a = b] &\sim 0_K \text{ (if } h(a) \neq_M h(b)) \end{aligned}$$

We use K^M to denote the semiring obtained by applying the above construction on a semiring K and a commutative monoid M . A key property is that, when we are able to interpret the equalities in M , K^M collapses to K . Formally,

PROPOSITION 4.4. *If K and M are such that $K \otimes M$ and M are isomorphic via ι then $K^M = K$.*

The proof (deferred to the Appendix) is by induction on the structure of elements in K^M , showing that at each step we can “solve” an equality sub-expression, and replace it with 0_K or 1_K .

Lifting homomorphisms. To conclude the description of the construction we explain how to lift a semiring homomorphism from $h : K \rightarrow K'$ to $h^M : K^M \rightarrow K'^M$, for any commutative monoid M and semirings K, K' . h^M is defined recursively on the structure of $a \in K^M$: if $a \in K$ we define $h^M(a) = h(a)$, otherwise $a = [b \otimes m_1 = c \otimes m_2]$ for some $b, c \in K^M$ and $m_1, m_2 \in M$ and we define $h^M(a) = [h^M(b) \otimes m_1 = h^M(c) \otimes m_2]$. Note that the application of a homomorphism h^M maps equality expressions to equality expressions (in which elements in K' appear instead of elements of K appeared before). If K' and M are such that their corresponding $\iota : M \rightarrow K \otimes M$ defined by $\iota(m) = 1_K \otimes m$ is injective, then we may “resolve the equalities”, otherwise the (new) equality expression remains.

4.3 The Extended Semantics

The extended semiring construction allows us to design a semantics for general aggregation queries. Intuitively, when the existence of a tuple in the result relies on the result of a comparison involving aggregate values (as in the result of applying selection or joins), we multiply the tuple annotation by the corresponding equation annotation.

In the sequel we assume, to simplify the definition, that the query aggregates and compares only values of $K^M \otimes M$ (a value $m \in M$ is first replaced by $\iota(m) = 1_K \otimes m$). In what follows, let $R(R_1, R_2)$ be (M, K^M) -relations on an attributes set U . Recall that for a tuple t , $t(u)$ (where $u \in U$) is the value of the attribute u in t ; also for $U' \subseteq U$, recall that we

use $t|_{U'}$ to denote the restriction of t to the attributes in U' . Last, we use $(K^M \otimes M)^U$ to denote the set of all tuples on attributes set U , with values from $K^M \otimes M$. The semantics follows:

1. *empty relation*: $\forall t \ \phi(t) = 0$.

2. *union*: $(R_1 \cup R_2)(t) =$

$$\begin{cases} \sum_{t' \in \text{supp}(R_1)} R_1(t') \cdot \prod_{u \in U'} [t'(u) = t(u)] & \text{if } t \in \text{supp}(R_1) \\ + \sum_{t' \in \text{supp}(R_2)} R_2(t') \cdot \prod_{u \in U'} [t'(u) = t(u)] & \cup \text{supp}(R_2) \\ 0 & \text{Otherwise.} \end{cases}$$

3. *projection*: Let $U' \subseteq U$, and let $T = \{t|_{U'} \mid t \in \text{supp}(R)\}$. Then $\Pi_{U'}(t) =$

$$\begin{cases} \sum_{t' \in \text{Supp}(R)} R(t') \cdot \prod_{u \in U'} [t(u) = t'(u)] & \text{if } t \in T \\ 0 & \text{Otherwise.} \end{cases}$$

4. *selection*: If P is an equality predicate involving the equation of some attribute $u \in U$ and a value $m \in M$ then $(\sigma_P(R))(t) = R(t) \cdot [t(u) = \iota(m)]$.

5. *value based join*: We assume for simplicity that R_1 and R_2 have disjoint sets of attributes, U_1 and U_2 resp., and that the join is based on comparing a single attribute of each relation. Let $u'_1 \in U_1$ and $u'_2 \in U_2$ be the attributes to join on. For every $t \in (K^M \otimes M)^{U_1 \cup U_2}$:

$$\begin{aligned} & (R_1 \bowtie_{R_1.u_1=R_2.u_2} R_2)(t) = \\ & R_1(t|_{U_1}) \cdot R_2(t|_{U_2}) \cdot_K [t(u_1) = t(u_2)]. \end{aligned}$$

Simple Variants. Natural join (when U_1 and U_2 are not necessarily disjoint) is captured by a similar expression, with the equality sub-expression on the attributes common to U_1 and U_2 ; join on multiple values is captured by multiplication by the corresponding multiple equality expressions; in the representation of cartesian product (denoted by \times) no equality expressions appear (only $R_1(t|_{U_1}) \cdot R_2(t|_{U_2})$).

6. *Aggregation*: $AGG_M(R)(t) =$

$$\begin{cases} 1 & t(u) = \sum_{t' \in \text{supp}(R)} R(t') *_{K \otimes M} t'(u) \\ 0 & \text{otherwise} \end{cases}$$

7. *Group By*: Let $U' \subseteq U$ be a subset of attributes that will be grouped and $u \in U \setminus U'$ be the aggregated attribute. Then for every $t \in (K^M \otimes M)^{U' \cup \{u\}}$:

$$GB_{U',u}R(t) =$$

$$\begin{cases} \delta((\Pi_{U'}R)(t|_{U'})) & t(u) = \sum_{t' \in \text{supp}(R)} (R(t') \cdot_K \\ & \prod_{u \in U'} [t'(u) = t(u)]) *_{K \otimes M} t'(u) \\ 0 & \text{otherwise} \end{cases}$$

It is straightforward to show that the algebra satisfies set/bag compatibility and poly-size overhead; commutation with homomorphism is proved in the Appendix.

EXAMPLE 4.5. Reconsider the relation in Example 4.3, and let us perform another sum aggregation on *Sal*. The value in the result now contains equation expressions:

$$\begin{aligned} & \delta(r_1 +_K r_2) \cdot_K [r_1 \otimes 20 +_{K \otimes M} r_2 \otimes 10 = 1_K \otimes 20] \\ & *_{K \otimes M} (r_1 \otimes 20 +_{K \otimes M} r_2 \otimes 10) \\ & +_{K \otimes M} \delta(r_3) \cdot_K [r_3 \otimes 10 = 1_K \otimes 20] *_{K \otimes M} r_3 \otimes 10 \end{aligned}$$

Given a homomorphism $h : \mathbb{N}[X] \rightarrow \mathbb{N}$ we can “solve” the equations, e.g. if $h(r_1) = 1$, $h(r_2) = 0$ and $h(r_3) = 2$, we obtain an aggregated value of $1 \otimes 40$. Note that the aggregation value is not monotone in r_1, r_2, r_3 : map r_2 to 1 (and keep r_1, r_3 as before), to obtain $1 \otimes 20$.

5. DIFFERENCE

We next show that via our semantics for aggregation, we can obtain for the first time a semantics for arbitrary queries with difference on K -relations. We describe the obtained semantics and study some of its properties.

5.1 Semantics for Difference

We first note that difference queries may be encoded as queries with aggregation, using the monoid $\mathbb{B} = (\{\perp, \top\}, \vee, \perp)$ (the following encoding was inspired by [29, 9]):

$$R - S = \Pi_{a_1 \dots a_n} \{ (GB_{\{a_1, \dots, a_n\}, b}(R \times \perp_b \cup S \times \top_b)) \wedge_{a_1, \dots, a_n} (R \times \perp_b) \}.$$

\perp_b and \top_b are relations on a single attribute b , containing a single tuple (\perp) and (\top) respectively, with provenance 1_K . Using the semantics of Section 4, we obtain a semantics for the difference operation.

Interestingly, we next show that the obtained semantics can be captured by a simple and intuitive expression. First, we note that since \mathbb{B} is idempotent, every semiring K positive with respect to $+_K$ is compatible with \mathbb{B} (see Theorem 3.12). The following proposition then holds for every K, K' and every two (\mathbb{B}, K) -relations R, S (proof deferred to the Appendix):

PROPOSITION 5.1. For every tuple t , semirings K, K' such that $K' \otimes \mathbb{B}$ is isomorphic to \mathbb{B} via $\iota(m) = 1_K \otimes m$, if $h : K \rightarrow K'$ is a semiring homomorphism then:

$$h^{\mathbb{B}}([(R - S)(t)]) = h^{\mathbb{B}}([S(t) \otimes \top = 0] \cdot_K R(t)).$$

The obtained provenance expression is thus “equivalent” (in the precise sense of Proposition 5.1) to $[S(t) \otimes \top = 0] \cdot R(t)$. The following lemma helps us to understand the meaning of the obtained equality expression:

LEMMA 5.2. For every semiring K which is positive w.r.t. $+_K$ and $h : K \rightarrow \mathbb{B}$, $h^M([S(t) \otimes \top = 0]) = \top$ iff $h(S(t)) = \perp$.

PROOF. It is clear that if $h(S(t)) = \perp$, $h^M([S(t) \otimes \top = 0]) = [h(S(t)) \otimes \top = 0] = [\perp \otimes \top = 0] = \top$. For the other direction, assume that $h(S(t)) = \top$. Thus $[h(S(t)) \otimes \top = 0] = [\top \otimes \top = 0]$. Since \mathbb{B} and $\widehat{\mathbb{B}}$ are compatible, $\iota : \mathbb{B} \rightarrow \mathbb{B} \otimes \widehat{\mathbb{B}}$ is injective; thus $\iota(\perp) \neq \iota(\top)$; consequently $h^M([S(t) \otimes \top = 0]) = [\top \otimes \top = 0] = \perp$. \square

Consequently, the semantics can be interpreted as follows: a tuple t appears in the result of $R - S$ if it appears in R , but does not appear in S . When the tuple appears in the result of $R - S$, it carries its original annotation from R . I.e. the existence of t in S is used as a boolean condition.

EXAMPLE 5.3. Let R, S be the following relations, where R contains employees and their departments and S containing departments that are designated to be closed:

ID	Dep	
1	d ₁	t ₁
2	d ₁	t ₂
2	d ₂	t ₃

R

Dep	
d ₁	t ₄

S

To obtain a relation with all departments that remains active, we can use the query $(\Pi_{Dep}R) - S$, resulting in:

Dep	
d ₁	$[t_4 \otimes \top = 0] \cdot (t_1 + t_2)$
d ₂	$[0 = 0] \cdot t_3 \quad (= t_3)$

Now consider some homomorphism $h : \mathbb{N}[X] \rightarrow \mathbb{N}$ (multiplicity e.g. stands for number of employees in the department). Note that if $h(t_4) > 0$ then the department d_1 is closed and indeed d_1 is omitted from the support of the difference query result, otherwise it retains each original annotation that it had in R . Assume now that we decide to revoke the decision of closing the department d_1 . This corresponds to mapping t_4 to 0; we can easily propagate this deletion to the query results; the equality appearing in the annotation of the first tuple is now $[0 = 0] = 1_K$ and we obtain as expected:

Dep	
d ₁	$t_1 + t_2$
d ₂	t_3

In particular, we obtain a semantics for the entire Relational Algebra, including difference. It is interesting to study the specialization of the obtained semantics for particular semirings: $\mathbb{B}, \mathbb{N}, \mathbb{Z}$, and to compare it to previously studied semantics for difference.

5.2 Comparison with other semantics

For a semiring K and a commutative monoid M we say that two queries Q, Q' are equivalent if for every input (M, K) -database D , the results (including annotations) $Q(D)$ and $Q'(D)$ are congruent (namely the corresponding values and annotations are congruent) according to the axioms of $K^M \otimes M$ and K^M . In the sequel we fix $M = \widehat{B}$ and consider different instances of K , exemplifying different equivalence axioms for queries with difference while comparing them with previously suggested semantics. We use $Q \equiv_K Q'$ to denote the equivalence of Q, Q' with respect to K and \widehat{B} .

\mathbb{B} -relations. For $K = \mathbb{B}$, our semantics is the same as set-semantics, thus the following proposition holds:

PROPOSITION 5.4. *For $Q, Q' \in RA$ it holds that $Q \equiv_{\mathbb{B}} Q'$ if and only if $Q \equiv Q'$ under set semantics.*

\mathbb{N} -relations. For $K = \mathbb{N}$, we compare our semantics to bag equivalence and observe that they are different (for queries with difference, even without aggregation). Intuitively this is because in our semantics, the righthand side of the difference is treated as a boolean condition, rather than having the effect of decreasing the multiplicity. Formally,

PROPOSITION 5.5. *$Q \equiv_{\mathbb{N}} Q'$ does not imply that $Q \equiv Q'$ under bag semantics, and vice versa.*

PROOF. Observe that $A - (B \cup B) \equiv_{\mathbb{N}} A - B$; but this does not hold for bag semantics. In contrast, under bag semantics $(A \cup B) - B \equiv A, B$, but not for our semantics. \square

EXAMPLE 5.6. *Reconsider Example 5.3, and let $t_1 = t_2 = t_3 = t_4 = 1$. Under bag semantics, after projecting R on the department attribute, the multiplicity of the department d_1 becomes 2; after applying the difference the department d_1 is still in the result, but now with multiplicity 1; in contrast under our semantics the department d_1 does not appear in the support of the result.*

\mathbb{Z} -relations. Finally, in [22] the authors have presented \mathbb{Z} semantics for difference, and have shown that it leads to equivalence axioms that are different from those that hold for queries with bag difference. It is also different from the equivalence axioms that we have here for \mathbb{Z} -relations:

PROPOSITION 5.7. *$Q \equiv_{\mathbb{Z}} Q'$ does not imply $Q \equiv Q'$ under \mathbb{Z} semantics¹⁰, and vice versa.*

PROOF. Under \mathbb{Z} semantics it was shown in [22] that $(A - (B - C)) \equiv (A \cup C) - B$. This does not hold for our semantics. In contrast $A - (B \cup B) \equiv_{\mathbb{Z}} A - B$, but this equivalence does not hold under \mathbb{Z} semantics. \square

Deciding Query Equivalence. We conclude with a note on the decidability of equivalence of queries using our semantics. It turns out that for semirings such as \mathbb{B}, \mathbb{N} for which we can interpret the results in $\widehat{\mathbb{B}}$ (in the sense of proposition 5.1 above), query equivalence is undecidable.

PROPOSITION 5.8. *Let K be such that $K^{\widehat{\mathbb{B}}} \otimes \widehat{\mathbb{B}}$ is isomorphic to $\widehat{\mathbb{B}}$. Equivalence of Relational Algebra queries on K -relations is undecidable.*

PROOF. The proof is by reduction from equivalence under set semantics: let ϕ be the empty query, i.e. a query whose answer always the empty relation. Given two RA queries Q, Q' (note that Q and Q' can include difference), their equivalence under set semantics holds if and only if $Q - Q' \equiv_{\mathbb{K}} \phi$ and $Q' - Q \equiv_{\mathbb{K}} \phi$. \square

6. RELATED WORK

Provenance information has been extensively studied in the database literature. Different provenance management techniques are introduced in [14, 7, 8, 6], etc., and it was shown in [24, 21] that these approaches can be compared in the semiring framework. To our knowledge, this work is the first to study aggregate queries in the context of provenance semirings. Provenance information has a variety of applications (see introduction) and we believe that our novel framework for aggregate queries will benefit all of these. Specifically, queries with aggregation play a key role in modeling the operational logic of scientific workflows (see e.g. [5, 16]) and our framework is likely to facilitate a more fine-grained approach to workflow provenance.

Aggregate queries have been extensively studied in e.g. [12, 13] for bag and set semantics. As explained in [12], such queries are fundamental in many applications: OLAP queries, mobile computing, the analysis of streaming data, etc. We note that Monoids are used to capture general aggregation operators in [13], but our paper seems to be the first to study their interaction with provenance.

Several semantics of *difference* on relations with annotations have been proposed, starting with the c -tables of [28]. The semirings with monus of [19] generalize this as well as bag-semantics. Difference on relations with annotations from \mathbb{Z} are considered in [22] and from $\mathbb{Z}[X]$ in [20]. As explained in Section 5, the semantics for difference defined in this paper is different from all of these.

There are interesting connections between provenance management and query evaluation on uncertain (and probabilistic) databases (e.g. [30, 15, 6, 3]), as observed in [24]. Evaluation of aggregate queries on probabilistic databases has been studied in e.g. [36, 33]. Trying to optimize the performance of aggregate query evaluation on probabilistic databases via provenance management is an intriguing future research challenge.

¹⁰As defined in [22].

7. CONCLUSION

We have studied in this paper provenance information for queries with *aggregation* in the semiring framework. We have identified three desiderata for the assessment of candidate approaches: compatibility with the usual set/bag semantics, commutation with semiring homomorphisms and poly-size overhead. After showing that approaches using provenance only to annotate the database tuples do not satisfy all desiderata simultaneously, we considered a different framework in which the computation of aggregate *values* is itself annotated with provenance. This has led us to the algebraic structure of *semimodules* over commutative semirings of annotations and to a tensor product construction for the semantics of annotated aggregation. The first product of this approach is a “good” (i.e. satisfying the desiderata) semantics for SPJU queries followed by an aggregation or a group-by with aggregation. We have further studied the challenges that arise in evaluation of queries that apply *comparisons* on aggregation results, e.g., joins over aggregate values, and shown that by careful adaptation of the semi-module framework these challenges can be overcome with a semantics that satisfies the desiderata. Finally, we noted that *difference queries* may be encoded as queries with aggregation, and studied the algebra induced for such queries.

We have exemplified in the paper the application of our approach for deletion propagation and security annotations. As mentioned in the Introduction and Related Work sections, there are various other areas in which provenance is useful. Future research will focus on applying our framework to the research tasks tackled in these areas.

8. REFERENCES

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [2] Foto N. Afrati and Angelos Vasilakopoulos. Managing lineage and uncertainty under a data exchange setting. In *SUM*, 2010.
- [3] L. Antova, C. Koch, and D. Olteanu. 10^{10^6} worlds and beyond: efficient representation and processing of incomplete information. *VLDB J.*, 18(5), 2009.
- [4] David W. Archer, Lois M. L. Delcambre, and D. Maier. A framework for fine-grained data integration and curation, with provenance, in a dataspace. In *First workshop on Theory and practice of provenance*, 2009.
- [5] Z. Bao, S. B. Davidson, S. Khanna, and S. Roy. An optimal labeling scheme for workflow provenance using skeleton labels. In *SIGMOD Conference*, 2010.
- [6] O. Benjelloun, A.D. Sarma, A.Y. Halevy, M. Theobald, and J. Widom. Databases with uncertainty and lineage. *VLDB J.*, 17, 2008.
- [7] P. Buneman, J. Cheney, and S. Vansummeren. On the expressiveness of implicit provenance in query and update languages. *ACM Trans. Database Syst.*, 33(4), 2008.
- [8] P. Buneman, S. Khanna, and W.C. Tan. Why and where: A characterization of data provenance. In *ICDT*, 2001.
- [9] P. Buneman, S. Naqvi, V. Tannen, and L. Wong. Principles of programming with complex objects and collection types. *TCS*, 149(1), 1995.
- [10] J. Cheney, L. Chiticariu, and W. C. Tan. Provenance in databases: Why, how, and where. *Foundations and Trends in Databases*, 1(4), 2009.
- [11] J. Cheney, S. Chong, N. Foster, M. I. Seltzer, and S. Vansummeren. Provenance: a future history. In *OOPSLA Companion*, 2009.
- [12] S. Cohen. Containment of aggregate queries. *SIGMOD Record*, 34(1), 2005.
- [13] S. Cohen, W. Nutt, and Y. Sagiv. Rewriting queries with arbitrary aggregation functions using views. *ACM Trans. Database Syst.*, 31(2), 2006.
- [14] Y. Cui, J. Widom, and J.L. Wiener. Tracing the lineage of view data in a warehousing environment. *ACM Transactions on Database Systems*, 25(2), 2000.
- [15] Nilesh N. Dalvi, Christopher Ré, and Dan Suciu. Probabilistic databases: diamonds in the dirt. *Commun. ACM*, 52(7), 2009.
- [16] E. Deelman, G. Singh, M. Su, and et. Al. Pegasus: A framework for mapping complex scientific workflows onto distributed systems. *Scientific Programming*, 13(3), 2005.
- [17] J.N. Foster, T.J. Green, and V. Tannen. Annotated XML: queries and provenance. In *PODS*, 2008.
- [18] N. Fuhr and T. Rölleke. A probabilistic relational algebra for the integration of information retrieval and database systems. *ACM Trans. Inf. Syst.*, 15(1), 1997.
- [19] F. Geerts and A. Poggi. On database query languages for k-relations. *J. Applied Logic*, 8(2), 2010.
- [20] T.J. Green. *Collaborative data sharing with mappings and provenance*. PhD thesis, University of Pennsylvania, 2009.
- [21] T.J. Green. Containment of conjunctive queries on annotated relations. In *ICDT*, 2009.
- [22] T.J. Green, Z. Ives, and V. Tannen. Reconcilable differences. In *ICDT*, 2009.
- [23] T.J. Green, G. Karvounarakis, Z. Ives, and V. Tannen. Update exchange with mappings and provenance. In *VLDB*, 2007.
- [24] T.J. Green, G. Karvounarakis, and V. Tannen. Provenance semirings. In *PODS*, 2007.
- [25] T.J. Green, G. Karvounarakis, N. E. Taylor, O. Biton, Z. Ives, and V. Tannen. Orchestra: facilitating collaborative data sharing. In *Proc. of SIGMOD*, 2007.
- [26] A. Gupta, I.S. Mumick, and V.S. Subrahmanian. Maintaining views incrementally. In *SIGMOD*, 1993.
- [27] J. Huang, T. Chen, A. Doan, and J. F. Naughton. On the provenance of non-answers to queries over extracted data. *Proc. VLDB*, 1, 2008.
- [28] T. Imielinski and W. Lipski. Incomplete information in relational databases. *J. ACM*, 31(4), 1984.
- [29] C. Koch. Incremental query evaluation in a ring of databases. In *PODS*, 2010.
- [30] C. Koch and D. Olteanu. Conditioning probabilistic databases. *PVLDB*, 1(1), 2008.
- [31] J. Lechtenböcker, H. Shu, and G. Vossen. Aggregate queries over conditional tables. *J. Intell. Inf. Syst.*, 19(3), 2002.
- [32] S.K. Lellahi and V. Tannen. A calculus for collections and aggregates. In *Cat. Theory and Comp. Science*, 1997.
- [33] J. Li, B. Saha, and A. Deshpande. A unified approach to ranking in probabilistic databases. *Proc. VLDB*, 2, 2009.
- [34] B. Liu, L. Chiticariu, V. Chu, H.V. Jagadish, and F. Reiss. Refining information extraction rules using data provenance. *IEEE Data Eng. Bull.*, 33(3), 2010.
- [35] A. Meliou, W. Gatterbauer, K. F. Moore, and D. Suciu. The complexity of causality and responsibility for query answers and non-answers. *PVLDB*, 4(1), 2010.
- [36] C. Ré and D. Suciu. Efficient evaluation of having queries on a probabilistic database. In *DBPL*, 2007.
- [37] Issam S., Adrian F., and Vladimiro S. A formal model of provenance in distributed systems. In *First workshop on Theory and practice of provenance*, 2009.
- [38] S. Vansummeren and J. Cheney. Recording provenance for sql queries and updates. *IEEE Data Eng. Bull.*, 30(4), 2007.
- [39] W. Zhou, M. Sherr, T. Tao, X. Li, B. T. Loo, and Y. Mao. Efficient querying and maintenance of network provenance at internet-scale. In *SIGMOD*, 2010.
- [40] E. Zimányi. Query evaluation in probabilistic relational databases. *Theor. Comput. Sci.*, 171(1-2), 1997.

APPENDIX

A. SPJU ALGEBRA FOR K-RELATIONS

We recall the full definition of the SPJU algebra for K-relations from [24]:

Empty Relation $\forall t \phi(t) = 0$.

Selection If $R : \mathbb{D}^U \rightarrow K$ and the selection predicate P maps each U-tuple to either 0 or 1 then $\sigma_P R : \mathbb{D}^U \rightarrow K$ is defined by $(\sigma_P R)(t) = R(t) \cdot P(t)$.

Projection If $R : \mathbb{D}^U \rightarrow K$ and $U' \subseteq U$ then $\Pi_{U'} R : \mathbb{D}^{U'} \rightarrow K$ is defined by $(\Pi_{U'} R)(t) = \sum_K R(t')$ where the $+$ _K sum is over all $t' \in \text{supp}(R)$ such that $t'|_{U'} = t$.

Natural Join If $R_i : \mathbb{D}^{U_i} \rightarrow K$, $i = 1, 2$ then $R_1 \bowtie R_2 : \mathbb{D}^{U_1 \cup U_2} \rightarrow K$ is defined by $(R_1 \bowtie R_2)(t) = R_1(t_1) \cdot_K R_2(t_2)$ where $t_i = t|_{U_i}$, $i = 1, 2$.

Union If $R_i : \mathbb{D}^U \rightarrow K$, $i = 1, 2$ then $R_1 \cup_K R_2 : \mathbb{D}^U \rightarrow K$ is defined by $(R_1 \cup_K R_2)(t) = R_1(t) +_K R_2(t)$.

B. PROPERTIES OF $K \otimes M$

We show here that the $K \otimes M$ constructed in Section 2.3 forms a K -semimodule, and highlight some of its basic properties.

PROPOSITION B.1. $K \otimes M$ is a K -semimodule.

PROOF. We show that the six semimodule axioms (definition 2.1) hold. Four of them hold already for bags of simple tensors. For example

$$\begin{aligned} (k \cdot_K k') *_{K \otimes M} \sum k_i \otimes m_i &= \sum (k \cdot_K k' \cdot_K k_i) \otimes m_i = \\ &= k *_{K \otimes M} \sum (k' \cdot_K k_i) \otimes m_i = k *_{K \otimes M} (k' *_{K \otimes M} \sum k_i \otimes m_i) \end{aligned}$$

By taking the quotient by congruence defined in Section 2.3 we also get the remaining two axioms:

$$\begin{aligned} (k +_K k') *_{K \otimes M} \sum k_i \otimes m_i &= \sum (k \cdot_K k_i +_K k' \cdot_K k_i) \otimes m_i \sim \\ &\sim \sum (k \cdot_K k_i) \otimes m_i +_{K \otimes M} \sum (k' \cdot_K k_i) \otimes m_i = \\ &= \sum (k \cdot_K k_i) \otimes m_i +_{K \otimes M} \sum (k' \cdot_K k_i) \otimes m_i = \\ &= k *_{K \otimes M} \sum k_i \otimes m_i +_{K \otimes M} k' *_{K \otimes M} \sum k_i \otimes m_i \end{aligned}$$

and

$$0_K *_{K \otimes M} \sum k_i \otimes m_i = \sum 0_K \otimes m_i \sim \sum 0_{K \otimes M} = 0_{K \otimes M}$$

This concludes the proof. \square

Furthermore, $K \otimes M$ is the “most economical” K -semimodule in the following sense. Define $\iota : M \rightarrow K \otimes M$ such that $\iota(m) = 1_K \otimes m$. Every tensor is a linear (with respect to K) combination of simple tensors from $\iota(M)$. More precisely,

PROPOSITION B.2. $K \otimes M$ satisfies the “universality” property, i.e. for any K -semimodule W and any homomorphism of monoids $f : M \rightarrow W$ there exists a unique homomorphism of K -semimodules $f^* : K \otimes M \rightarrow W$ such that $f^* \circ \iota = f$.

PROOF. Define f^* first on bags of simple tensors as follows

$$f^*(\sum k_i \otimes m) = \sum k_i *_{K \otimes M} f(m)$$

where the second sum is taken with $+$ _W (in particular, the empty such sum is by convention 0_W). Thus $f^*(\iota(m)) =$

$f^*(1_K \otimes m) = 1_K *_{K \otimes M} f(m) = f(m)$. Then, one can check that f^* is a homomorphism with respect to $+$ _{$K \otimes M$} , $0_{K \otimes M}$ and $*_{K \otimes M}$. This implies that for f^* to preserve \sim it suffices to preserve the four laws of the congruence given in Section 2.2, which is readily checked, for example

$$\begin{aligned} f^*(k \otimes m +_{K \otimes M} k \otimes m') &= k *_{K \otimes M} f(m) +_W k *_{K \otimes M} f(m') = \\ &= k *_{K \otimes M} (f(m) +_W f(m')) = k *_{K \otimes M} f(m +_M m') = f^*(k \otimes (m +_M m')) \end{aligned}$$

Since f^* preserves \sim it can be defined as above by picking a representative from each equivalence class. Now let $g : K \otimes M \rightarrow W$ be another linear function such that $g \circ \iota = f$. Then

$$\begin{aligned} g(\sum k_i \otimes m) &= g(\sum k_i *_{K \otimes M} (1_K \otimes m)) = \sum k_i *_{K \otimes M} g(1_K \otimes m) = \\ &= \sum k_i *_{K \otimes M} f(m) = f^*(\sum k_i \otimes m) \end{aligned}$$

hence $g = f^*$, thus verifying the uniqueness of f^* . In particular, any linear function on $K \otimes M$ is completely determined by its behavior on the tensors in $\iota(M)$. \square

We also say that $K \otimes M$ is the K -semimodule *freely generated* by M (thus the “most economical” appellation).

Recall that we have defined in Section 2.3 the “lifting” of a homomorphism of semirings $h : K \rightarrow K'$ to a homomorphism of monoids $h^M : K \otimes M \rightarrow K' \otimes M$. Its definition is an immediate consequence of Proposition B.2: indeed $K' \otimes M$ becomes a K -semimodule via h so we can define h^M as the unique homomorphism of K -semimodules that by the proposition extends $\iota' : M \rightarrow K' \otimes M$. Note that this yields the definition in Section 2.3:

$$h^M(\sum k_i \otimes m_i) = \sum h(k_i) \otimes m_i$$

C. ADDITIONAL PROOFS

PROOF. (Proposition 3.1)

The “if” direction follows from the fact that homomorphisms, by definition, preserve $\{+, \cdot, 0, 1\}$ -expressions. For the “only if” direction we use *abstractly tagged* [24] databases. These are $\mathbb{N}[X]$ -databases in which each tuple is annotated by just an indeterminate in X , and a different one at that. It is as if tuples are annotated by their distinct id. Clearly, there is a canonical way of choosing a large enough X and producing a canonical abstractly tagged database that is completely determined by its support. Now, fix an operation Ω , and consider its semantics on K -databases, for various K . For any K and any K -database D , let D^a be the abstractly tagged database determined by $\text{supp}(D)$. Let $h : \mathbb{N}[X] \rightarrow K$ be the homomorphism uniquely determined by mapping the abstract tags in X to the actual annotations of D . Fix a tuple t in $\text{supp}(\Omega(D^a))$ and let $p_t = \Omega(D^a)(t) \in \mathbb{N}[X]$ be the polynomial that annotates it in $\Omega(D^a)$. By commutation with homomorphisms and the definition of h_{Rel} we have $\Omega(D)(t) = \Omega(h_{Rel}(D^a))(t) = h_{Rel}(\Omega(D^a))(t) = h(p_t)$. Using the homomorphism properties we move h inside p_t until it applies to just indeterminates. For an indeterminate, say x , $h(x)$ is the K -annotation of the unique tuple in $\text{supp}(D)$ that is annotated with x in D^a . It follows that $\Omega(D)(t)$ is given by the $\{+, \cdot, 0, 1\}$ -expression p_t in terms of the annotations of D . But p_t only depends on $\text{supp}(D)$ and Ω while it is the same for all K . This shows the algebraic uniformity of the semantics. \square

Commutation with homomorphism for SPJU-AGB queries.

We next prove that the semantics proposed for restricted aggregation queries (in Section 3) satisfies commutation with homomorphisms:

PROOF. The proof is by induction on the query structure, but since commutation with homomorphisms was already shown for SPJU queries, we need only to prove that GB commutes with homomorphisms as well. Let R be a K -Rel on the set of attributes U , where $U', U'' \subseteq U$ and $U' \cap U'' = \emptyset$. R may be the result of applying any sequence of SPJU operations that appear in the query Q , followed by $GB_{U', U''}(R)$.

Consider the result when first applying the GB operation followed by h_{Rel} . According to definition 3.7, the result of applying GB will be a relation R' , whose support contains every tuple t such that:

1. t is defined on the attributes $U' \cup U''$;
2. For some non-empty subset $T = \{t'_1, \dots, t'_n\}$ of $\text{supp}(R)$, the restriction of t to attributes of U' is equal to the restriction of every tuple $t'_i \in T$ to U' , and not equal to the restriction to U' of any other tuple in $\text{supp}(R) - T$;
3. For each $u \in U''$, $t(u) = \sum_{t'_i \in T} R(t'_i) \otimes t'_i(u)$; and
4. $R'(t) = \delta \left(\sum_{t'_i \in T} R(t'_i) \right)$.

The effect of applying h_{Rel} on such t would then be:

1. $(h_{Rel}(R'))(t)$ is equal by definition to $h(R'(t)) = h \left(\delta \left(\sum_{t'_i \in T} R(t'_i) \right) \right) = \delta \left(\sum_{t'_i \in T} h(R(t'_i)) \right)$.
2. For every $u \in U''$, $t(u)$ will be replaced by $h^M(t(u)) = h^M \left(\sum_{t'_i \in T} R(t'_i) \otimes t'_i(u) \right) = \sum_{t'_i \in T} h(R(t'_i)) \otimes t'_i(u)$.
3. The rest of the values in t remain intact.

Then, let us check the result of applying h_{Rel} before the GB operation, and compare it to the above result. Applying h_{Rel} on R will only affect the tuple provenance annotations; for every tuple t' , $(h_{Rel}(R))(t') = h(R(t'))$. Now let us apply $GB_{U', U''}(h_{Rel}(R))$. Again, according to our semantics, the result will be a relation R'' , whose support contains every tuple t such that:

1. t is defined on the attributes $U' \cup U''$ (as before);
2. For some non-empty subset $T = \{t'_1, \dots, t'_n\}$ of $\text{supp}(h_{Rel}(R))$, the restriction of t to attributes of U' is equal to the restriction of every tuple $t'_i \in T$ to U' , and not equal to the restriction to U' of any other tuple in $\text{supp}(h_{Rel}(R)) - T$;
3. for each $u \in U''$, $t(u) = \sum_{t'_i \in T} (h_{Rel}(R)(t'_i)) \otimes t'_i(u) = \sum_{t'_i \in T} h(R(t'_i)) \otimes t'_i(u)$; and
4. $R''(t) = \delta \left(\sum_{t'_i \in T} h(R(t'_i)) \right)$.

We now need to verify that these results are indeed equivalent. Note first that applying h_{Rel} on R before applying the GB operation, only affects the tuple annotations, and not their values. We then employ a “by-case” analysis to verify equivalence. For any tuple t that is both in the support of R and in the support of $h_{Rel}(R)$, it is easy to observe from the above equations that the “contribution” of t to both the aggregation value and its annotation is the same in both cases. Consequently we can focus on tuples in $\text{supp}(R)$ for which h_{Rel} sets their annotations to 0, thus deleting them from their groups or even deleting a whole group by deleting all its members. Those tuples and groups contribute to the result in R' , when applying GB first (before applying h_{Rel}).

However, this means that the summands corresponding to the annotations of those tuples in the δ annotation of the groups in R' will be later set to 0 by h_{Rel} ; as for the aggregation results, for every tuple t' that is deleted by h , its summand $h(R(t')) \otimes t'(u)$ in each aggregation result will be set to $0 \otimes t'_i(u) = 0_{K \otimes M}$ and thus it will have no effect on the aggregation results. We thus conclude that if no group has been deleted altogether, the results are equivalent. The last case to consider is that where all the annotations of tuples in group T are set to zero. In this case, its δ expression will be equal to zero as well, so the group is effectively deleted also by h_{Rel} after the GB.

This concludes the proof. \square

PROOF. (Thm. 3.12)

Let K be a commutative semiring which is positive with respect to $+_K$ and define $h : K \otimes M \rightarrow M$ as $h(\sum_{i \in I} k_i \otimes m_i) = \sum_{j \in J} m_j$ where $J = \{j \in I \mid k_j \neq 0\}$. We can show that h is well-defined (see below); since $\forall m \in M \quad h \circ \iota(m) = m$, ι is injective and thus K and M are compatible.

We need to verify that h is a well-defined mapping, and for that we check that it is well-defined on $K \otimes M$ after taking the quotient (as defined in Section 2.2):

- (For $k, k' \neq 0_K$) $h((k +_K k') \otimes m) = m = m +_M m = h(k \otimes m +_{K \otimes M} k' \otimes m)$.
- $h(0_K \otimes m) = 0_M$, and also the empty bag is mapped to the “empty sum” i.e. $h(0_{K \otimes M}) = 0_M$.
- (For $k \neq 0_K$) $h(k \otimes (m +_M m')) = m +_M m' = h(k \otimes m +_{K \otimes M} k \otimes m')$.
- $h(k \otimes 0_M) = 0_M$, and again $h(0_{K \otimes M}) = 0_M$.

Note that we assumed that k and k' are non-zero in the first and third axioms. Since K is positive with respect to $+_K$, no such k, k' can satisfy $k +_K k' = 0_K$, thus the case of $0_K \otimes m$ is *uniquely* defined to be mapped to 0_M , by the second axiom.

This concludes the proof. \square

PROOF. (Thm. 3.13)

Let h' be a homomorphism from K to \mathbb{N} , and M be an arbitrary commutative monoid. We define a mapping $h : K \otimes M \rightarrow M$, as follows.

$h(\sum k_i \otimes m_i) = \sum h'(k_i)m_i$. We show that h is well-defined and that $h \circ \iota$ is the identity function.

We first show that this mapping is well-defined, i.e. that every pair of elements from $K \otimes M$ which are equated by the axioms of the tensor construction (as defined in Section 2.2), are mapped by h to the same values.

1st axiom. Left side: $h((k +_K k') \otimes m)$ is equal by the definition of h to $h'(k +_K k')m$. Since h' is a homomorphism, this is equal to $(h'(k) + h'(k'))m$. Right side: $h(k \otimes m +_{K \otimes M} k' \otimes m) = h'(k)m +_M h'(k')m$ by h definition. Since $h'(k), h'(k')$ are natural numbers, this is equal to the result of the left hand side.

2nd axiom. Left side: $h(0_K \otimes m) = h'(0_K)m$. Since h' is a homomorphism, $h'(0_K) = 0$ and thus the expression is equal to $0m = 0_M$. Right side: by definition of h , the “empty” sum in $K \otimes M$ must be mapped to the “empty” sum in M , which is 0_M .

3rd axiom. Left side: $h(k \otimes (m +_M m')) = h'(k)(m +_M m') = h'(k)m +_M h'(k)m'$. Right side: $h(k \otimes m +_{K \otimes M} k \otimes m') = h'(k)m +_M h'(k)m'$.

4th axiom. Left side: $h(k \otimes 0_M) = h'(k)0_M = 0_M$. Right side: same as the 2nd axiom.

Since h is well-defined such that $h(a+b) = h(a) + h(b)$, it is a homomorphism from $K \otimes M$ to M . Now we need to show that $h \circ \iota$ is the identity function, implying that ι is injective and thus that M and K are compatible. This is easy: since h' is a homomorphism it must map 1_K to 1; then for every $m \in M$, $h(\iota(m)) = h(1_K \otimes m) = h'(1_K)m = 1m = m$. \square

PROOF. (Proposition 4.2)

We show the proof for SUM and the proof for MAX (MIN) is similar. Reconsider the relation R' and the query Q_{select} in Example 4.1, and assume that $R'' = Q_{select}(R')$ is a $(M, \mathbb{N}[X])$ -relation capturing the query result according to some algebra. Assume by contradiction that the algebra commutes with homomorphism. Let h, h' be homomorphisms from $\mathbb{N}[X]$ to \mathbb{N} corresponding to those in Example 3.2, I.e. $h(r_1) = h(r_3) = h'(r_1) = h'(r_2) = h'(r_3) = 1$ and $h(r_2) = 0$. We saw in Example 4.1 that the aggregation result when h is applied is 20; thus, in order to be bag-compatible, $h^M(R'')$ must include a tuple t'' representing this tuple which matched the selection condition. Let $p_{t''} \in \mathbb{N}[X]$ be its provenance annotation, then $h(p_{t''}) = 1$. However, $h^M(R'')$ is empty, since no aggregation result is equal to 20 in that case. Thus $h'(p_{t''}) = 0$. Similarly to the proof of prop. 3.2, observe that there exists no such polynomial $p_{t''} \in \mathbb{N}[X]$. \square

PROOF. (Theorem 4.4)

The proof is by induction on the structure of elements in K^M . We say that an expression $exp \in K^M$ has a nesting level of 0 if for every expression $[c_1 \otimes m_1 = c_2 \otimes m_2]$ appearing in exp , $c_1, c_2 \in K$ (and $m_1, m_2 \in M$); exp has a nesting level n if each such c_1, c_2 are of nesting level $n-1$ or less. For nesting level of 0, axiom (*) above allows us to replace $[c_1 = c_2]$ with 1_K or 0_K . Now, assume that the theorem holds for expressions with nesting level $n-1$ or less, and let exp be of nesting level n . Then for each sub-expression $[c_1 \otimes m_1 = c_2 \otimes m_2]$, we can replace, by the induction hypothesis (and using the axioms above), c_1, c_2 with elements of K and then apply axiom (*) to replace the equality expression with an element of K . We can repeat for every equality sub-expression of exp , obtaining an element of K . \square

Commutation with homomorphism for the extended semantics. We next prove that the semantics proposed for nested aggregation queries (in Section 4) satisfies commutation with homomorphisms:

PROOF. The proof is by induction on the query structure. For each operation we consider two cases: (I) applying the homomorphism h_{Rel} after applying the operation; (II) applying it before the operation. Both cases must yield equal results.

In what follows we use the same notations, R, R_1, R_2 and so on, as used in the definition of the extended semantics.

Union. First, consider case (I), where the union is applied first. According to the defined semantics, the result of $R_1 \cup R_2$ is a (M, K^M) -relation such that for every tuple t in its support it holds that:

1. t is defined (only) on the attributes in U .
2. $t \in \text{supp}(R_1) \cup \text{supp}(R_2)$
3. $(R_1 \cup R_2)(t) = \sum_{t' \in \text{supp}(R_1)} R_1(t') \cdot_K \prod_{u \in U} [t'(u) = t(u)] +_K \sum_{t' \in \text{supp}(R_2)} R_2(t') \cdot_K \prod_{u \in U} [t'(u) = t(u)]$.

Then, applying h_{Rel} on $R_1 \cup R_2$ has the following effect on that t .

1. For every value in t from $K^M \otimes M$, its value changes from $\Sigma k_i \otimes m_i$ to $h^M(\Sigma k_i \otimes m_i) = \Sigma h(k_i) \otimes m_i$.
2. The provenance annotation of t is changed according to the axioms of the homomorphism lifting, to $h_{Rel}((R_1 \cup R_2))(t) = \sum_{t' \in \text{supp}(R_1)} h^M(R_1(t')) \cdot \prod_{u \in U} [h^M(t'(u)) = h^M(t(u))] + \sum_{t' \in \text{supp}(R_2)} h^M(R_2(t')) \cdot \prod_{u \in U} [h^M(t'(u)) = h^M(t(u))]$

Now, for case (II), let us apply h_{Rel} first, on both R_1 and R_2 . This would affect only the provenance annotations of tuples within this relations, causing perhaps to the deletion of some tuples, and the values from $K^M \otimes M$, which change in the same manner as described in item (1) above. Let us compute the result of $h_{Rel}(R_1) \cup h_{Rel}(R_2)$. Every tuple t in the support of the obtained relation is such that:

1. t is defined (only) on the attributes in U .
2. $t \in \text{supp}(R_1) \cup \text{supp}(R_2)$ and it holds that either $h(R_1(t)) \neq 0_{K'}$ or $h(R_2(t)) \neq 0_{K'}$
3. $(h_{Rel}(R_1) \cup h_{Rel}(R_2))(t) = \sum_{t' \in \text{supp}(R_1)} h^M(R_1(t')) \cdot \prod_{u \in U} [h^M(t'(u)) = h^M(t(u))] +_K \sum_{t' \in \text{supp}(R_2)} h^M(R_2(t')) \cdot \prod_{u \in U} [h^M(t'(u)) = h^M(t(u))]$.

We next verify that the results are indeed equal. for every tuple t there are several options: it can be in the support of R_1, R_2 , neither or both; and h can set its provenance to 0 in none of them, one of them or both. Any tuple t which is not in $\text{supp}(R_1)$ or $\text{supp}(R_2)$ clearly does not affect the result. Any tuple who is at least in one of them, will be annotated in case (I) with a sum of each annotation of each tuple t' in $\text{supp}(R_1) \cup \text{supp}(R_2)$, multiplied by tokens that equate each value from $K^M \otimes M$ in t to the value of the same attribute in t' . In the worst case, where all the values are from $K^M \otimes M$, we do not know which tuples are equal and thus we compare each pair on each attribute. Then applying h might cause some of the original tuple annotation, hence some of the summands in the provenance of t to become 0. In case (II) those tuples for which h sets the annotations to 0 are removed from R_1, R_2 or both, and thus they do not appear in the sum to begin with. For the tuples that remain it is clear that the obtained annotations and $K^M \otimes M$ are the same in both case (I) and case (II).

One thing to note here is that different tuples in a relation, R_1 for instance, may be equated after applying h_{Rel} . This is true, for instance, when we have two tuples which differ only by some aggregation result, but after applying h those results turn out to be the same. This works well with the homomorphism commutation as well, because it is easy to see that in both cases the tuples will be equal, in case (I) after applying h on the union result and in case (II) before the union is applied.

The proof for projection is very similar to the one for union, thus it is not repeated here.

Selection. According to the algebra definition, to get $(\sigma_P(R))(t)$ we simply multiply the annotation of each tuple $t \in \text{supp}(R)$ by an expression equating the value of the relevant attribute u in t to some value m (embedded into $K^M \otimes M$ using ι).

In case (I) the provenance of some tuple t in the support of $\sigma_P(R)$ might be $R(t) \cdot_K [(\Sigma k_i \otimes m_i) = 1_K \otimes m]$, which would become, after applying h , $h(R(t)) \cdot_K [(\Sigma h(k_i) \otimes m_i) = 1_{K'} \otimes m]$.

In case (II) the annotation of tuple t will become $h(R(t))$ after applying h_{Rel} , and $t(u)$ would become $\Sigma h(k_i) \otimes m_i$. Thus after applying the selection, we would get the same result, $h(R(t)) \cdot_K [(\Sigma h(k_i) \otimes m_i) = 1_{K'} \otimes m]$.

Aggregation. In case (I), we first apply the GB operation $GB_{U',u}R$ and obtain a relation such that for every tuple t in its support it holds that:

1. t is defined (only) on the attributes in $U \cup \{u\}$.
2. There exists some tuple $t' \in \text{supp}(R)$ such that for every attribute $u' \in U'$, $t(u') = t'(u')$.
3. $t(u) = \Sigma_{t' \in \text{supp}(R)} (R(t')) \cdot_{\mathcal{K}} \prod_{u' \in U'} [t(u') = t'(u')] *_{\mathcal{K} \otimes M} t'(u)$
4. $GB_{U',u}R(t) = \delta(\Sigma_{t' \in \text{supp}(R)} R(t')) \cdot_{\mathcal{K}} \prod_{u' \in U'} [t(u') = t'(u')]$.

Now, after applying h_{Rel} on the result, the effect on such tuple t would be:

1. $t(u) = \Sigma_{t' \in \text{supp}(R)} (h^M(R(t')) \cdot_{\mathcal{K}} \prod_{u' \in U'} [h^M(t(u')) = h^M(t'(u'))]) *_{\mathcal{K} \otimes M} h^M(t'(u))$
2. $h_{Rel}(GB_{U',u}R)(t) = \delta(\Sigma_{t' \in \text{supp}(R)} h^M(R(t')) \cdot_{\mathcal{K}} \prod_{u' \in U'} [h^M(t(u')) = h^M(t'(u'))])$.

In case (II), we first apply h_{Rel} , which affects the tuple provenances and the values from $K^M \otimes M$. Then aggregation is applied on the result. Each tuple t in $\text{supp}(GB_{U',u}h_{Rel}(R))$ is such that:

1. t is defined (only) on the attributes in $U \cup \{u\}$.
2. There exists some tuple $t' \in \text{supp}(h_{Rel}(R))$ such that for every attribute $u' \in U'$, $t(u') = t'(u')$.
3. $t(u) = \Sigma_{t' \in \text{supp}(R)} (h^M(R(t')) \cdot_{\mathcal{K}} \prod_{u' \in U'} [t(u') = h^M(t'(u'))]) *_{\mathcal{K} \otimes M} h^M(t'(u))$
4. $GB_{U',u}h_{Rel}(R)(t) = \delta(\Sigma_{t' \in \text{supp}(R)} h^M(R(t')) \cdot_{\mathcal{K}} \prod_{u' \in U'} [t(u') = h^M(t'(u'))])$.

We now verify that the results in both cases are indeed equal. In the first case, according to the definition, every tuple t in $\text{supp}(R)$ is forming the basis of a group, which conditionally may contain every tuple in R (using equation expressions to verify that each tuple is indeed in that group only if its restriction to U' is equal to the restriction of t to U'). When we apply h , some tuple annotations may be set to 0, and thus their corresponding summands (in the group provenances and aggregation results) are set to 0, and do not affect the result. In case (II) some tuples may be removed by h from the relation even before the aggregation is performed. This has a similar effect to setting their corresponding summands to 0 as in case (I). There is a slight difference here: If $\text{supp}(R)$ was of size n , so will be the size of the support of $GB_{U',u}R$, maybe even after applying h_{Rel} on it; However, $\text{supp}(h_{Rel}(R))$ may be of size $m < n$, and thus so will be $GB_{U',u}h_{Rel}(R)$. The reason is that as long as we cannot evaluate equation values, we have to allow for n different groups (as many as there are in the support), thus if there are “actually” less, when we move to a $K^M \otimes M$ where equations may be evaluated, we may get the same tuple representing the group duplicated the number of times as the number of its group members. This is acceptable, since duplicates are ignored. However if we apply h first, we may know that there cannot be n groups to begin with, since some of the tuples are deleted. The effect of this will only be less group duplicates. The commutation of the AGG operation with homomorphism follows from the above proof as well.

This concludes the proof. \square

PROOF. (Proposition 5.1)

First, note that by the homomorphism commutation, we can check the equivalence of using both expressions on $h^{\hat{B}}(R)$

and $h^{\hat{B}}(S)$, i.e. verify that

$$((\prod_{a_1, \dots, a_n} \{GB_{\{a_1, \dots, a_n\}, b}(h^{\hat{B}}(R) \times \perp_b \cup h^{\hat{B}}(S) \times \top_b)\} \times (h^{\hat{B}}(R) \times \perp_b)))(t) = [h^{\hat{B}}(S)(t) \otimes \top = 0] \cdot_{\mathcal{K}'} h^{\hat{B}}(R)(t).$$

Note that since $K' \otimes \hat{B}$ is isomorphic via an isomorphism I to \hat{B} , instead of checking equality in $K' \otimes \hat{B}$ we can check for equality after application of the isomorphism (in \hat{B}); in particular this allows us to interpret the equalities and replace them with $1_{\mathcal{K}'}$ or $0_{\mathcal{K}'}$.

Now let us follow the operation of difference encoded by aggregation step-by-step. Let $R, S \in \mathbb{D}^U \rightarrow K'^{\hat{B}}$; let $\text{supp}(h^{\hat{B}}(R)) = \{r_1, \dots, r_n\}$ and $\text{supp}(h^{\hat{B}}(S)) = \{s_1, \dots, s_m\}$. Then $\text{supp}(h^{\hat{B}}(R) \times \perp_b) = \{r'_i \mid r'_i(b) = \perp \wedge \exists r_i \in \text{supp}(h^{\hat{B}}(R)) \forall u \in U \ r'_i(u) = r_i(u)\}$; since \times is equivalent to join with no attribute equations, the provenance $(h^{\hat{B}}(R) \times \perp_b)(t)$ is $h^{\hat{B}}(R)(t) \cdot_{\mathcal{K}'} \perp_b(t) = h^{\hat{B}}(R)(t)$ for any $t \in \text{supp}(h^{\hat{B}}(R) \times \perp_b)$, and $0_{\mathcal{K}'}$, otherwise; similarly for $h^{\hat{B}}(S) \times \top_b$. Now, note that the support sets of both relations are mutually exclusive, and that all the values in the tuples in the support of both $h^{\hat{B}}(R)$ and $h^{\hat{B}}(S) \times \top_b$ are from \mathbb{D} . Thus, it is easy to see that $(h^{\hat{B}}(R) \times \perp_b) \cup (h^{\hat{B}}(S) \times \top_b)(t)$ is $h^{\hat{B}}(R)(t|_U)$ for $t \in \text{supp}(h^{\hat{B}}(R) \times \perp_b)$, $h^{\hat{B}}(S)(t|_U)$ for $t \in \text{supp}(h^{\hat{B}}(S) \times \top_b)$, and $0_{\mathcal{K}'}$, otherwise. Now we apply group-by on the b attribute. There are four possible classes of tuples:

1. For every r'_i such that $r_i = s_j \in h^{\hat{B}}(R) \cap h^{\hat{B}}(S)$, we would get a tuple r''_i in the support of the result, such that $r''_i|_U = r'_i|_U = r_i = s_j = s'_j|_U$, and $r''_i(b) = (h^{\hat{B}}(R) \times \perp_b)(r'_i) \otimes \perp +_{\mathcal{K}' \otimes \hat{B}} h^{\hat{B}}(S) \times \top_b(s'_i) \otimes \top = h^{\hat{B}}(S)(r_i) \otimes \top$. The provenance of this tuple would be $\delta(h^{\hat{B}}(R)(r_i) +_{\mathcal{K}' \otimes \hat{B}} h^{\hat{B}}(S)(r_i))$.
2. For every r'_i such that $r_i \in h^{\hat{B}}(R) \setminus h^{\hat{B}}(S)$, there is a tuple r''_i in the group-by result such that $r''_i(b) = (h^{\hat{B}}(R) \times \perp_b(r'_i)) \otimes \perp = 0_{\mathcal{K}' \otimes \hat{B}}$, and where the provenance of r''_i is $\delta(h^{\hat{B}}(R)(r_i))$.
3. For every s'_i such that $s_i \in h^{\hat{B}}(S) \setminus h^{\hat{B}}(R)$, there is a tuple s''_i in the group-by result such that $s''_i(b) = (h^{\hat{B}}(S) \times \top_b(s'_i)) \otimes \top = h^{\hat{B}}(S)(s_i) \otimes \top$, and where the provenance of s''_i is $\delta(h^{\hat{B}}(S)(s_i))$.
4. Every other tuple has provenance $0_{\mathcal{K}'}$, i.e. it is not in the support of the aggregation result.

Now we need to perform a join of the group-by result and $h^{\hat{B}}(R) \times \perp_b$. For that we will first rename the attributes of the aggregation result to $U_2 \cup \{b'\}$. For each tuple t in the aggregation result such that $t|_U \in h^{\hat{B}}(R)$, i.e. tuples of cases (1) and (2) above, there will be a unique corresponding tuple in the join result (since all these tuples have unique U values, they will join with a unique tuple in $h^{\hat{B}}(R) \times \perp_b$). The obtained provenance of the join of tuples from case (1), would be $\delta(h^{\hat{B}}(R)(r_i) + h^{\hat{B}}(S)(r_i)) \cdot h^{\hat{B}}(R)(r_i) \cdot [h^{\hat{B}}(S)(r_i) \otimes \top = 0]$. However, $\delta(h^{\hat{B}}(R)(r_i) + h^{\hat{B}}(S)(r_i))$ is redundant here: if $h^{\hat{B}}(S)(r_i) \neq 0_{\mathcal{K}'}$ or $h^{\hat{B}}(R)(r_i) = 0_{\mathcal{K}'}$, the provenance is $0_{\mathcal{K}'}$; otherwise, $\delta(h^{\hat{B}}(R)(r_i) + h^{\hat{B}}(S)(r_i)) = \delta(h^{\hat{B}}(R)(r_i)) = 1_{\mathcal{K}'}$. As for the join of tuples from case (2), the case is simpler: the provenance is $\delta(h^{\hat{B}}(R)(r_i)) \cdot h^{\hat{B}}(R)(r_i) \cdot [0 = 0] = \delta(h^{\hat{B}}(R)(r_i)) \cdot h^{\hat{B}}(R)(r_i)$, and again $\delta(h^{\hat{B}}(R)(r_i))$ has no ef-

fect. For cases (3) and (4) the provenance is $0_{K'}$. This matches the result of $[(h^{\hat{B}}(S))(t) \otimes \top = 0]_{K'}(h^{\hat{B}}(R))(t)$. \square