



BIROn - Birkbeck Institutional Research Online

Aamer, H. and Hidders, Jan and Paredaens, J. and Bussche, J.V.d. (2021) Expressiveness within Sequence Datalog. In: UNSPECIFIED (ed.) PODS'21: Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems. ACM, pp. 70-81. ISBN 9781450383813.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/53502/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>
contact lib-eprints@bbk.ac.uk.

or alternatively

Expressiveness within Sequence Datalog

Heba Aamer
Hasselt University
heba.mohamed@uhasselt.be

Jan Hidders
Birkbeck, University of London
jan@dcs.bbk.ac.uk

Jan Paredaens
Universiteit Antwerpen
jan.paredaens@uantwerpen.be

Jan Van den Bussche
Hasselt University
jan.vandenbussche@uhasselt.be

Abstract

Motivated by old and new applications, we investigate Datalog as a language for sequence databases. We reconsider classical features of Datalog programs, such as negation, recursion, intermediate predicates, and relations of higher arities. We also consider new features that are useful for sequences, notably, equations between path expressions, and “packing”. Our goal is to clarify the relative expressiveness of all these different features, in the context of sequences. Towards our goal, we establish a number of redundancy and primitivity results, showing that certain features can, or cannot, be expressed in terms of other features. These results paint a complete picture of the expressiveness relationships among all possible Sequence Datalog fragments that can be formed using the six features that we consider.

This paper is the extended version of a paper presented at PODS 2021 [1]

1 Introduction

Interest in sequence databases dates back for at least three decades [13]. For clarity, here, by sequence databases, we do not mean relations where the tuples are ordered by some sequence number or timestamp, possibly arriving in a streaming fashion (e.g., [12, 43, 39, 28]). Rather, we mean databases that allow the management of large *collections of sequences*.

In the early years, sequence databases were motivated by applications in object-oriented software engineering [5] and in genomics [9, 26]. While these applications remain relevant, more recent applications of sequence databases include the following.

- Process mining [25] operates on event logs, which are sets of sequences. Thus, sequence databases, and sequence database query languages, can serve as enabling technology for process mining. For example, a typical

query one may want to be able to support is look for all logs in which every occurrence of ‘complete order’ is followed by ‘receive payment’.

- Graph databases have as main advantage over relational databases that they offer convenient query primitives for retrieving paths. Paths are, of course, sequences. For example, the G-CORE graph query language proposal [29] supports the querying of sequences stored in the database, separately from the graph; these sequences do not even have to correspond to actual paths in the graph. An example query in such a context could be to return the nodes that belong to all paths in a given set of paths.
- JSON Schema [36] is based on the notion of JSON pointers, which are sequences of keys navigating into nested JSON objects. In J-Logic [22] we showed that modeling JSON databases as sequence databases is very convenient for defining JSON-to-JSON transformations in a logical, declarative manner.

For a simple example, consider a JSON object *Sales* that is a set of key–value pairs, where keys are items; the value for an item is a nested object holding the sales volumes for the item by year. Specifically, the nested object is again a set of key–value pairs, where keys are years and values are numbers. We can naturally view *Sales* as a set of length-3 sequences of the form item–year–value. Restructuring the object to group sales by year, rather than by item, then simply amounts to swapping the first two elements of every sequence. For another example, checking if two JSON objects are deep-equal amounts to checking equality of the corresponding sets of sequences.

- Logical approaches to information extraction [44, 17] model the result of an information extraction as a sequence database.

Given the importance of sequences in various advanced database applications, our research goal in this paper is to obtain a thorough understanding of the role that different language features play in querying sequence databases. For such an investigation, we need an encompassing query language in which these features are already present, or can be added. For this purpose we adopt Datalog, a logical framework that is well established in database theory research, and that has continued practical relevance [14, 7, 4].

Indeed, Datalog for sequence databases, or Sequence Datalog, was already introduced and studied by Bonner and Mecca in the late 1990s [9, 32]. They showed that, to make Datalog work with sequence databases, all we have to do is to add terms built from sequence variables using the concatenation operator. In our work we refer to such terms as *path expressions* and refer to sequence variables as *path variables*.¹ Bonner and Mecca studied computational

¹We actually work with a minor variant of Bonner and Mecca’s language; while they additionally introduce index terms, but only allow path expressions in the heads of rules, we allow path expressions also in rule bodies, and additionally introduce atomic variables. The two variants are equivalent in that one can be simulated by the other requiring no additional features such as negation or recursion.

completeness, complexity, and termination guarantees for Sequence Datalog, and showed how to combine Sequence Datalog with subcomputations expressed using transducers.

Sequence Datalog was recently also considered for information extraction (“document spanners”), with regular expression matching built-in as a primitive [33, 35]. Such regular expressions may be viewed as very useful syntactic sugar, as they are also expressible using recursion. Adding regular matching directly may be compared to Bonner and Mecca’s transducer extensions; the PTIME capturing result reported by Peterfreund et al. [35] may be compared to Corollary 3 of Bonner and Mecca [9].

In the present work, we study the relative expressiveness of query language features in the context of Sequence Datalog. Some of the features we consider are standard Datalog, namely, recursion, stratified negation, and intermediate predicates. The latter feature actually comprises two features, since we distinguish between monadic intermediate predicates and intermediate predicates of higher arities. While we omit regular expression matching as a feature, we consider two further features that are specific to sequences:

- Equalities between path expressions, which we call *equations*, allow for the elegant expression of pattern matching on sequences.
- *Packing*, a feature introduced in J-Logic, is a versatile tool that allows for subsequences to be “bracketed” and temporarily treated as atomic values; they can be unpacked later.

The standard Datalog features, whose expressiveness is well understood on classical relational structures [3, 16], need to be re-examined in the presence of sequences; moreover, their interaction with the new features needs to be understood as well. For example, consider recursion versus equations, and the query that checks whether an input sequence x consists exclusively of a ’s. (Path variables are prefixed by a dollar sign.) With an equation we can simply write $x \cdot a = a \cdot x$ (using the dot for concatenation). Without equations (or other means to simulate equations), however, this query can only be expressed using recursion. For another example, consider monadic versus higher-arity intermediate predicates. Classically, there are well-known arity hierarchies for Datalog [21]. In our setting, however, a unary relation can already hold arbitrary-length sequences, and indeed, using a simple coding trick, we will see that the arity feature is actually redundant.

In our work, we have chosen to define expressiveness in terms of the baseline class of “flat unary queries”, namely, functions from unary relations to unary relations, where both the input and the output are just sets of plain, unpacked sequences. In this way, we avoid trivial tautologies such as “packing is a primitive feature, because without it, we cannot create packed sequences”. As a matter of fact, we will show that packing, although it certainly is a convenient feature, is actually redundant for expressing these flat unary queries. A result in this direction was already stated for J-Logic [22], but the technique used there

to simulate packing requires recursion. In the present paper, we show that packing is redundant also in the absence of recursion. Our proof technique leverages associative unification [2], and more specifically, the termination of associative unification for particular cases of word equations [15].

Our further results can be summarized as follows.

1. At first sight, equations seem to be a redundant feature, at least in the presence of intermediate predicates. Indeed, instead of using an equation $e_1 = e_2$ as a subgoal, we can introduce an auxiliary relation $T(e_1, e_2)$, and replace the equation by the subgoal $T(e_1, e_1)$. (Our notation here is not precise but hopefully enough to convey the idea). With negated equations and recursion, however, this simple trick does not work as it violates stratification. We still show, however, that equations are redundant in the presence of both intermediate predicates and negation.
2. In the absence of intermediate predicates, however, equations are a primitive feature. Indeed, the “only a ’s” query mentioned above, easily expressed with an equation, is not expressible in the absence of intermediate predicates.
3. One can also, conversely, simulate intermediate predicates using equations: a simple folding transformation works in the absence of negation and recursion. In the presence of negation or recursion, however, intermediate predicates do add power. This is fairly easy to show for recursion: the squaring query “for every path p in the input, output a^{n^2} , where n is the length of p ” requires an intermediate predicate in which the output can be constructed recursively. In the presence of negation, the primitivity of intermediate predicates can be seen to follow from the corresponding result for classical Datalog (by quantifier alternation). Some work has still to be done, however, since the classical proof has to be extended to account for path expressions and equations.
4. It will not surprise the reader that recursion is primitive in Sequence Datalog. This can be seen in many ways; probably the easiest is to use the above squaring query, and to observe that without recursion, the length of output sequences is at most linear in the length of input sequences. Another proof, that also works for boolean queries, is by reduction to the classical inexpressibility of graph connectivity in first-order logic. As in the previous paragraph, the reduction must account for the use of path expressions and equations.
5. A classical fact is that nonrecursive Datalog with stratified negation is equivalent to the relational algebra. We extend the standard relational algebra by allowing path expressions in selection and projection, and adding operators for unpacking and for subsequences. We obtain a language equivalent to nonrecursive Sequence Datalog.




Figure 1: Relative expressiveness of the different sets of Sequence Datalog features (Negation, Equations, Intermediate predicates, and Recursion; features Arity and Packing will turn out to be entirely redundant). An ascending path denotes subsumption; absence of such a path denotes non-subsumption.

Our results allow us to completely classify the sixteen possible Sequence Datalog fragments in a Hasse diagram with respect to their expressive power, as shown in Figure 1. Some fragments are equivalent, as shown; also, the features for packing and higher-arity intermediate predicates are omitted, since they are redundant independently of the presence or absence of other features.

This paper is organized as follows. In Section 2 we define the sequence database model and the syntax and semantics of Sequence Datalog. In Section 3 we introduce the language features and rigorously define what we mean by one fragment (set of features) being subsumed in expressive power by another fragment. Section 4 presents our redundancy (expressibility) results, and Section 5 presents our primitivity (inexpressibility) results. The Hasse diagram of Figure 1 is assembled in Section 6. Section 7 presents the relational algebra for sequence databases. We conclude in Section 8, where we also discuss additional related work.

2 Sequence databases and Sequence Datalog

In this section we formally define the sequence database model and the syntax and semantics of Sequence Datalog. We do assume some familiarity with the basic notions of classical Datalog [3].

2.1 Data model for sequence databases

A *schema* Γ is a finite set of relation names, each name with an associated *arity* (a natural number). We fix a countably infinite universe **dom** of atomic data elements, called *atomic values*. The sets of *packed values*, *values*, and *paths* are defined as the smallest sets satisfying the following:

1. Every atomic value is a value.

2. Every finite sequence of values is a path. The empty path is denoted by ϵ .

When writing down paths, we will separate the elements by dots, where the \cdot symbol also serves as the usual symbol for concatenation. Recall that concatenation is associative.

3. If p is a path, then $\langle p \rangle$ is a packed value.
4. Every packed value is a value.

The set of all paths is denoted by Π .

For example, if a , b and c are atomic values, then $a \cdot b \cdot a$ is a path; $\langle a \cdot b \cdot a \rangle$ is a packed value; and $c \cdot \langle a \cdot b \cdot a \rangle$ is again a path.

An *instance* I of a schema Γ is a function that assigns to each relation name $R \in \Gamma$ a finite n -ary relation on Π , with n the arity of R .

It is natural to identify a value v with the one-length sequence v . Thus values, in particular atomic values, are also paths. Hence, classical relational database instances are a special case of instances as defined here. We refer to such instances as *classical*. So, in a classical instance, each relation name R is assigned a finite relation on **dom**.

2.2 Syntax of Sequence Datalog

We assume disjoint supplies of *atomic variables* (ranging over atomic values) and *path variables* (ranging over paths). The set of all variables is also disjoint from **dom**. We indicate atomic variables as $@x$ and path variables as $\$x$. *Path expressions* are defined just like paths, but with variables added in. Formally, we define the set of path expressions to be the smallest set such that:

1. Every atomic value is a path expression;
2. Every variable is a path expression;
3. If e is a path expression, then $\langle e \rangle$ is a path expression;
4. Every finite sequence of path expressions is a path expression.

A *predicate* is an expression of the form $P(e_1, \dots, e_n)$, with P a relation name of arity n , and each e_i a path expression. We call e_i the i th component of the predicate. An *equation* is an expression of the form $e_1 = e_2$, with e_1 and e_2 path expressions.

Many of the following definitions adapt well-known Datalog notions to our data model.

An *atom* is a predicate or an equation. A *negated atom* is an expression of the form $\neg A$ with A an atom. We write a negated equation $\neg e_1 = e_2$ also as a nonequality $e_1 \neq e_2$. A *literal* is an atom (also called a positive literal) or a negated atom (a negative literal). A *body* is a finite set of literals (possibly empty). A *rule* is an expression of the form $H \leftarrow B$, where H is a predicate, called the *head* of the rule, and B is a body.

We define the *limited variables* of a rule as the smallest set such that:

1. every variable occurring in a positive predicate in the body is limited; and
2. if all variables occurring in one of the sides of a positive equation in the body are limited, then all variables occurring in the other side are also limited.

A rule is called *safe* if all variables occurring in the rule are limited.

Finally, a *program* is a finite sequence of *strata*, which are finite sets of safe rules, so that use of negation in the program is stratified. Recall that stratified negation means that when a negated predicate $\neg P(e_1, \dots, e_n)$ occurs in some stratum, then no rule in that stratum or later strata can use P in the head predicate.

Note that classical Datalog programs with stratified negation are a special case of our notion of programs, where the only path expressions used are atomic values or atomic variables.

Example 2.1. An NFA can be represented by a unary relation N (initial states), a ternary relation D (transitions), and a unary relation F (final states). These would be classical relations. Now consider a unary relation R containing paths without packing, i.e., strings of atomic values. Then the following program, consisting of a single stratum, computes in relation A the strings from R that are accepted by the NFA. Recall that atomic variables are prefixed with @, and path variables with \$.

```
S(@q.$x, ε) ← R($x), N(@q).
S(@q2.$y, $z.@a) ← S(@q1.@a.$y, $z), D(@q1, @a, @q2).
A($x) ← S(@q, $x), F(@q).
```

Example 2.2. Consider unary relations R and S . The following program, again in a single stratum, uses packing and nonequalities to check whether there are at least three different occurrences of a string from S as a substring in strings from R . The boolean result is computed in the nullary relation A .

```
T($u.<$s>.$v) ← R($u.$s.$v), S($s).
A ← T($x), T($y), T($z), $x≠$y, $x≠$z, $y≠$z.
```

2.3 Semantics

We have defined the notion of instance as an assignment of relations over Π to relation names. A convenient equivalent view of instances is as sets of facts. A *fact* is an expression of the form $R(p_1, \dots, p_n)$ with R a relation name of arity n , and each p_i a path. An instance I of a schema Γ is viewed as the set of facts $\{R(p_1, \dots, p_n) \mid R \in \Gamma \text{ and } (p_1, \dots, p_n) \in I(R)\}$.

A *valuation* ν is a function that maps atomic variables to atomic values and path variables to paths. We say that ν is *appropriate* for a syntactical construct (such as a path expression, a literal, or a rule) if ν is defined on all variables in that syntactical construct. We can apply an appropriate valuation ν to a path

expression e by substituting each variable in e by its image under ν and obtain the path $\nu(e)$. Likewise, we can apply an appropriate valuation to a predicate and obtain a fact.

Let L be a literal, ν a valuation appropriate for L , and I an instance. The definition of when I, ν satisfies L is as expected: if L is a predicate, then the fact $\nu(L)$ must be in I ; if L is an equation $e_1 = e_2$, then $\nu(e_1)$ and $\nu(e_2)$ must be the same value. If L is a negated atom $\neg A$, then I, ν must not satisfy A .

A body B is satisfied by I, ν if all its literals are. Now a rule $r = H \leftarrow B$ is satisfied in I if for every valuation ν appropriate for r such that I, ν satisfies B , also I, ν satisfies H .

The relation names occurring in a program are traditionally divided into EDB and IDB relation names. The IDB relation names are the relation names used in the head of some rules; the other relation names are the EDB relation names. Given a schema Γ , a program is said to be over Γ if all its EDB relation names belong to Γ , and its IDB relation names do not. Now the semantics of programs is defined as usual. A program is called *semipositive* if negated predicates only use EDB relation names. We first apply the first stratum, which is semipositive, and then apply each subsequent stratum as a semipositive program to the result of the preceding strata. So we only need to give semantics for semipositive programs. Let \mathbf{P} be a semipositive program over Γ , and let I be an instance over Γ . Let Γ' be the set of IDB relation names of \mathbf{P} . Then $\mathbf{P}(I)$ is the smallest instance over $\Gamma \cup \Gamma'$ that satisfies all the rules of \mathbf{P} , and that agrees with I on Γ .

Due to recursion, for some programs or instances, $\mathbf{P}(I)$ may be undefined, since instances are required to be finite. We also say in this case that \mathbf{P} does not *terminate* on I . If, in the course of evaluating a program \mathbf{P} with several strata on an instance I , one of the strata does not terminate, we agree that the entire program \mathbf{P} is undefined on I . As mentioned in the Introduction, Bonner and Mecca have done substantial work on the question of guaranteeing termination for Sequence Datalog programs. In this paper, we only consider programs that always terminate.

Example 2.3. The program from Example 2.1, while recursive, is guaranteed to terminate on every instance. In contrast, the following two-rule program will not terminate on any instance:

```
T(a).
T(a.$x) ← T($x).
```

3 Features, fragments, and queries

In this paper, we consider six possible features that a program may use, each identified by a letter, spelled out as follows.

Arity A program *uses arity* (has the A-feature) if it contains at least one predicate of arity greater than one.

Recursion A program *uses recursion* (has the R-feature) if there is a cycle in its dependency graph.²

Equations A program *uses equations* (has the E-feature) if it contains at least one equation in some rule.

Negation A program *uses negation* (has the N-feature) if it contains at least one negated atom in some rule.

Packing A program *uses packing* (has the P-feature) if a path expression of the form $\langle e \rangle$ occurs in some rule.

Intermediate predicates A program *uses intermediate predicates* (has the I-feature) if it involves at least two different IDB relation names.

Let $\Phi = \{A, E, I, N, P, R\}$ be the set of all features. A subset of Φ is called a *fragment*. A program \mathbf{P} is said to belong to a fragment F if it uses only features from F .

Example 3.1. The following program belongs to fragment $\{E\}$. It computes, in relation S , all paths from R that consist exclusively of a 's.

$$S(\$x) \leftarrow R(\$x), a \cdot \$x = \$x \cdot a.$$

The following program does the same, but belongs to fragment $\{A, I, R\}$:

$$\begin{aligned} T(\$x, \$x) &\leftarrow R(\$x). \\ T(\$x, \$y) &\leftarrow T(\$x, \$y \cdot a). \\ S(\$x) &\leftarrow T(\$x, \epsilon). \end{aligned}$$

3.1 Queries and subsumption among fragments

Our goal is to compare the different fragments with respect to their power in expressing queries. Our methodology is to do this relative to a baseline class of queries that do not presuppose any feature to begin with. We next define these queries formally.

We call a schema *monadic* if each of its relation names has arity zero or one. Also, we call an instance *flat* if it contains no occurrences of packed values.

Given a monadic schema Γ and relation name $S \notin \Gamma$ of arity at most one, a *query* from Γ to S is a total mapping from flat instances over Γ to flat instances over $\{S\}$. A program \mathbf{P} is said to *compute* such a query if

1. \mathbf{P} is over Γ ;
2. \mathbf{P} terminates on every flat instance of Γ ;
3. S is an IDB relation of \mathbf{P} ; and

²The nodes of this graph are the IDB relation names, and there is an edge from R_1 to R_2 if R_2 occurs in the body of a rule with R_1 in its head predicate.

4. $\mathbf{P}(I)(S)$ equals $Q(I)$ for every flat instance I of Γ .

We now say that fragment F_1 is *subsumed* by fragment F_2 , denoted by $F_1 \leq F_2$, if every query computable by a program in F_1 is also computable by a program in F_2 . Note that it is possible, for different F_1 and F_2 , that $F_1 \leq F_2$ and $F_2 \leq F_1$. Such two fragments are equivalent in expressive power. There will turn out to be 11 equivalence classes; in Section 6 we will give a theorem that will characterize the subsumption relation as shown in Figure 1.

3.2 Redundancy and primitivity

We will explore the subsumption relation by investigating the redundancy or primitivity of the different features with respect to other features. A feature might be redundant in an absolute sense, in that it can be dropped from any fragment without decrease in expressive power. This is a very strong notion of redundancy, and we cannot expect it to hold for most features. Yet a more relative notion of redundancy may hold, meaning that some feature does not contribute to expressive power, on condition that some other features are already present, or are absent. This leads to the following notions.

Definition 3.2 (Redundancy). Let X be a feature and let Y and Z be sets of features.

- X is redundant if $F \leq F - \{X\}$ for every fragment F .
- X is redundant in the presence of Y if $F \leq F - \{X\}$ for every fragment F such that $Y \subseteq F$.
- X is redundant in the absence of Z if $F \leq F - \{X\}$ for every fragment F such that Z is disjoint from F .
- X is redundant in the presence of Y and absence of Z if $F \leq F - \{X\}$ for every fragment F such that $Y \subseteq F$ and Z is disjoint from F .

Similarly, but conversely, a feature might be primitive in an absolute sense, in that dropping it from a fragment always strictly decreases the expressive power. Then again, for other features only more relative notions of primitivity may hold.

Definition 3.3 (Primitivity). Let X be a feature and let Y and Z be sets of features. Recall that Φ is the set of all features.

- X is primitive if $\{X\} \not\leq \Phi - \{X\}$.
- X is primitive in the presence of Y if $\{X\} \cup Y \not\leq \Phi - \{X\}$.
- X is primitive in the absence of Z if $\{X\} \not\leq \Phi - (\{X\} \cup Z)$.

4 Expressibility results

In this section we show various expressibility results that lead to absolute or relative redundancy results for various features.

4.1 Arity

Using a simple encoding trick we can see that arity is redundant. Indeed, let a and b be two different atomic values. For any paths s_1 , s_2 , s'_1 and s'_2 , we have the following:

Lemma 4.1. $(s_1, s_2) = (s'_1, s'_2)$ if and only if

$$s_1 \cdot a \cdot s_2 \cdot a \cdot s_1 \cdot b \cdot s_2 = s'_1 \cdot a \cdot s'_2 \cdot a \cdot s'_1 \cdot b \cdot s'_2.$$

Proof. The if-direction is trivial. For the only-if direction, we consider $s_1 \cdot a \cdot s_2 \cdot a \cdot s_1 \cdot b \cdot s_2 = s'_1 \cdot a \cdot s'_2 \cdot a \cdot s'_1 \cdot b \cdot s'_2$ and we observe that a appears in the middle of both sequences. Hence,

(a) $s_1 \cdot a \cdot s_2 = s'_1 \cdot a \cdot s'_2$ and

(b) $s_1 \cdot b \cdot s_2 = s'_1 \cdot b \cdot s'_2$.

For the sake of contradiction, let us assume $|s_1| < |s'_1|$. Then $s'_1 = s_1 \cdot x$ for a nonempty sequence x . Thus, equation (a) can be rewritten as $s_1 \cdot a \cdot s_2 = s_1 \cdot x \cdot a \cdot s'_2$, which simplifies to $a \cdot s_2 = x \cdot a \cdot s'_2$. Hence the sequence x must start with a . In the same way, however, we can deduce from (b) that x must start with b . Hence, the assumption we made is false.

Analogously, $|s_1| > |s'_1|$ can be seen to be false as well, so we know that $|s_1| = |s'_1|$. Then clearly $|s_2| = |s'_2|$ as well. Hence, from (a) and (b) we get that $s_1 = s'_1$ and $s_2 = s'_2$. \square

Using this encoding, arities higher than one can be reduced by one. Since we can do this repeatedly, we obtain:

Theorem 4.2. *Arity is redundant.*

Example 4.3. Consider the following program which computes in S the reversals of the paths in R :

$T(\$x, \epsilon) \leftarrow R(\$x)$.
 $T(\$x, \$y \cdot @u) \leftarrow T(\$x \cdot @u, \$y)$.
 $S(\$x) \leftarrow T(\epsilon, \$x)$.

The same query can be expressed without arity as follows:

$T(\$x \cdot a \cdot a \cdot \$x \cdot b) \leftarrow R(\$x)$.
 $T(\$x \cdot a \cdot \$y \cdot @u \cdot a \cdot \$x \cdot b \cdot \$y \cdot @u) \leftarrow T(\$x \cdot @u \cdot a \cdot \$y \cdot a \cdot \$x \cdot @u \cdot b \cdot \$y)$.
 $S(\$x) \leftarrow T(a \cdot \$x \cdot a \cdot b \cdot \$x)$.

4.2 Equations

In the presence of l and A , positive equations are readily seen to be redundant, by introducing an auxiliary intermediate predicate in the program. We only give an example:

Example 4.4. Recall the program from Example 3.1:

$$S(\$x) \leftarrow R(\$x), \quad a \cdot \$x = \$x \cdot a.$$

The same query can be computed without equations as follows:

$$\begin{aligned} T(a \cdot \$x, \$x) &\leftarrow R(\$x). \\ S(\$x) &\leftarrow T(\$x \cdot a, \$x). \end{aligned}$$

This simple method works only in the absence of negation, because, when applied to a negated equation in a rule that belongs to a recursive stratum, stratification is violated. However, negated equations can be handled by another method:

Lemma 4.5. *E is redundant in the presence of l , A and N .*

Proof. Positive equations can be handled as above. For each stratum Δ that contains negated equations, we insert a new stratum Δ' , right before Δ , consisting of the following rules. Let ρ be a renaming that maps each head relation name in Δ to a fresh relation name; relation names that occur only in bodies in Δ are mapped to themselves by ρ .

For each rule $H \leftarrow B$ in Δ without negated equations, we add the rule $\rho(H) \leftarrow \rho(B)$ to Δ' .

For each rule $r : H \leftarrow B \wedge e_1 \neq e'_1 \wedge \dots \wedge e_n \neq e'_n$ in Δ with n negated equations, we again add $\rho(H) \leftarrow \rho(B)$ to Δ' . Moreover, using a fresh relation name T , we add the following n rules for $i = 1, \dots, n$:

$$T(v_1, \dots, v_m) \leftarrow \rho(B) \wedge e_i = e'_i$$

Here, the v 's are all variables appearing in B .

Finally, in Δ , we replace r by the following rule:

$$H \leftarrow B \wedge \neg T(v_1, \dots, v_m).$$

□

Example 4.6. The following program retrieves in S those paths from R that can be written as $a_1 \cdots a_n b_n \cdots b_1$ with $a_i \neq b_i$ for $i = 1, \dots, n$:

$$\begin{aligned} U(\$x, \$x) &\leftarrow R(\$x). \\ U(\$x, \$y) &\leftarrow U(\$x, @a \cdot \$y \cdot @b), \quad @a \neq @b. \\ S(\$x) &\leftarrow U(\$x, \epsilon). \end{aligned}$$

Applying the method to eliminate negated equations, we obtain:

$$\begin{aligned}
\mathbf{U1}(\$x, \$x) &\leftarrow \mathbf{R}(\$x). \\
\mathbf{U1}(\$x, \$y) &\leftarrow \mathbf{U1}(\$x, @a.\$y.@b). \\
\mathbf{T}(\$x, \$y, @a, @b) &\leftarrow \mathbf{U1}(\$x, @a.\$y.@b), @a=@b. \\
\mathbf{S1}(\$x) &\leftarrow \mathbf{U1}(\$x, \epsilon). \\
\mathbf{U}(\$x, \$x) &\leftarrow \mathbf{R}(\$x). \\
\mathbf{U}(\$x, \$y) &\leftarrow \mathbf{U}(\$x, @a.\$y.@b), \neg\mathbf{T}(\$x, \$y, @a, @b). \\
\mathbf{S}(\$x) &\leftarrow \mathbf{U}(\$x, \epsilon).
\end{aligned}$$

From the above we conclude that E is redundant in the presence of I and A. Since we already know that arity is redundant, we obtain:

Theorem 4.7. *E is redundant in the presence of I.*

4.3 Packing

In this section we show that packing is redundant. The main task will be to eliminate packing from equations in nonrecursive programs. We will follow the following strategy to achieve this task:

1. In Section 4.3.3 we show how to eliminate all variables that can hold values with packing. We will call such variables *impure*. The elimination is achieved by “solving” equations involving impure variables.
2. Thereto, we will extend a known method for solving word equations. We begin by recalling this method in Section 4.3.1. In Section 4.3.2 we present the extension to path expressions.
3. When all variables are pure, equations involving packing can only be satisfiable if the two sides have a similar “shape”, called *packing structure*. We formalize this in Section 4.3.4.

The main result concerning packing is then proven in Section 4.3.5.

4.3.1 Solving equations

Consider an equation $e_1 = e_2$ and let X be the set of variables occurring in the equation. A valuation ν on X is called a *solution* if $\nu(e_1)$ and $\nu(e_2)$ are the same path. The set of solutions is typically infinite, so we would like a way to represent this set in a finite manner.

Thereto one can use variable substitutions: partial functions that map variables to path expressions over X . Such a variable substitution ρ is called a *symbolic solution* to the equation if $\rho(e_1)$ and $\rho(e_2)$ are the same path expression. Every symbolic solution ρ represents a set of solutions

$$[\rho] := \{\nu \circ \rho \mid \nu \text{ a valuation on } X\}.$$

A set R of symbolic solutions is called *complete* if $\bigcup\{[\rho] \mid \rho \in R\}$ yields the complete set of solutions to the equation.

The classical setting of *word equations* [2] can be seen as a special case of the situation just described. A word equation corresponds to the case where e_1 and e_2 contain no packing, and no atomic variables, i.e., all variables are path variables.

Plotkin’s “pig-pug” procedure for associative unification [38] generates a complete set of symbolic solutions to any word equation. However, not every word equation admits a *finite* complete set of symbolic solutions; a simple example is our familiar equation $x \cdot a = a \cdot x$. Hence, in general, the procedure may not terminate.³ Nevertheless, pig-pug is guaranteed to terminate on “one-sided nonlinear” equations [15]. These are word equations where all variables that occur more than once in the equation, only occur in one side of the equation.

We briefly review the pig-pug procedure. The procedure constructs a search tree whose nodes are labeled with word equations; the root is labeled with the original word equation. For each node we generate children according to a rewriting relation, \Rightarrow , on word equations. Specifically, we have the following rewrite rules:

1. Cancellation rule: $(x \cdot w_1 = x \cdot w_2) \Rightarrow (w_1 = w_2)$, for $x \in \mathbf{dom} \cup X$.
2. Main rules: each one of the rules is associated with a substitution, ρ . Let \mathbf{x} and \mathbf{y} be distinct variables and let a be an atomic value.
 - (a) $(\mathbf{x} \cdot w_1 = \mathbf{y} \cdot w_2) \Rightarrow (\mathbf{x} \cdot \rho(w_1) = \rho(w_2))$ with $\rho(\mathbf{x}) = \mathbf{y} \cdot \mathbf{x}$
 - (b) $(\mathbf{x} \cdot w_1 = \mathbf{y} \cdot w_2) \Rightarrow (\rho(w_1) = \rho(w_2))$ with $\rho(\mathbf{x}) = \mathbf{y}$
 - (c) $(\mathbf{x} \cdot w_1 = \mathbf{y} \cdot w_2) \Rightarrow (\rho(w_1) = \mathbf{y} \cdot \rho(w_2))$ with $\rho(\mathbf{y}) = \mathbf{x} \cdot \mathbf{y}$
 - (d) $(\mathbf{x} \cdot w_1 = a \cdot w_2) \Rightarrow (\mathbf{x} \cdot \rho(w_1) = \rho(w_2))$ with $\rho(\mathbf{x}) = a \cdot \mathbf{x}$
 - (e) $(\mathbf{x} \cdot w_1 = a \cdot w_2) \Rightarrow (\rho(w_1) = \rho(w_2))$ with $\rho(\mathbf{x}) = a$
 - (f) $(a \cdot w_1 = \mathbf{y} \cdot w_2) \Rightarrow (\rho(w_1) = \mathbf{y} \cdot \rho(w_2))$ with $\rho(\mathbf{y}) = a \cdot \mathbf{y}$
 - (g) $(a \cdot w_1 = \mathbf{y} \cdot w_2) \Rightarrow (\rho(w_1) = \rho(w_2))$ with $\rho(\mathbf{y}) = a$

When no rule is applicable to an equation, we have reached a leaf node in the search tree. There are three possible cases for such a leaf equation:

1. $(\epsilon = \epsilon)$.
2. $(a \cdot w_1) = (b \cdot w_2)$, for atomic values $a \neq b$.
3. $(\epsilon = w)$ or $(w = \epsilon)$, for nonempty w .

The first case is successful, while the other two are not. Each path from the root to a leaf node of the form $(\epsilon = \epsilon)$ yields a symbolic solution, formed by composing the substitutions given by the rewritings along the path. When starting from a one-side nonlinear equation, the tree is finite and we obtain a complete finite set of symbolic solutions.⁴

³The reader may be interested to know that other means of finite representation (different

4.3.2 Extension to path expressions

Our equations differ from word equations in that path expressions can involve packing as well as atomic variables. To this end, we extend the rewriting system as follows.

- (h) Given an equation of the form $(@x \cdot w_1 = @y \cdot w_2)$, the only possibility is for $@x$ and $@y$ to be the same. Thus we add the rule $(@x \cdot w_1 = @y \cdot w_2) \Rightarrow (\rho(w_1) = \rho(w_2))$ with $\rho(@x) = @y$.
- (i) An equation of the form $(@x \cdot w_1 = \$y \cdot w_2)$ is not very different from the case where we have a constant instead of $@x$. Thus, we add two rules similar to rules (f) and (g):
 - $(@x \cdot w_1 = \$y \cdot w_2) \Rightarrow (\rho(w_1) = \$y \cdot \rho(w_2))$ with $\rho(\$y) = @x \cdot \y
 - $(@x \cdot w_1 = \$y \cdot w_2) \Rightarrow (\rho(w_1) = \rho(w_2))$ with $\rho(\$y) = @x$
- (j) Analogously, we add rules similar to rules (d) and (e):
 - $(\$x \cdot w_1 = @y \cdot w_2) \Rightarrow (\$x \cdot \rho(w_1) = \rho(w_2))$ with $\rho(\$x) = @y \cdot \x
 - $(\$x \cdot w_1 = @y \cdot w_2) \Rightarrow (\rho(w_1) = \rho(w_2))$ with $\rho(\$x) = @y$
- (k) Given an equation of the form $(\langle w_1 \rangle \cdot w_2 = \langle w_3 \rangle \cdot w_4)$, we work inductively and solve the equation $w_1 = w_3$ first. Assuming we can find a finite complete set R of symbolic solutions for this equation, we then add the rules $(\langle w_1 \rangle \cdot w_2 = \langle w_3 \rangle \cdot w_4) \Rightarrow (\rho(w_2) = \rho(w_4))$ for $\rho \in R$.
- (l) An equation of the form $(\langle w_1 \rangle \cdot w_2 = \$y \cdot w_3)$ is again not very different from the case where we have a constant instead of $\langle w_1 \rangle$. Thus, we add two rules similar to rules (f) and (g):
 - $(\langle w_1 \rangle \cdot w_2 = \$y \cdot w_3) \Rightarrow (\rho(w_2) = \$y \cdot \rho(w_3))$ with $\rho(\$y) = \langle w_1 \rangle \cdot \y
 - $(\langle w_1 \rangle \cdot w_2 = \$y \cdot w_3) \Rightarrow (\rho(w_2) = \rho(w_3))$ with $\rho(\$y) = \langle w_1 \rangle$
- (m) Analogously, we again add rules similar to rules (d) and (e):
 - $(\$x \cdot w_1 = \langle w_2 \rangle \cdot w_3) \Rightarrow (\$x \cdot \rho(w_1) = \rho(w_3))$ with $\rho(\$x) = \langle w_2 \rangle \cdot \x
 - $(\$x \cdot w_1 = \langle w_2 \rangle \cdot w_3) \Rightarrow (\rho(w_1) = \rho(w_3))$ with $\rho(\$x) = \langle w_2 \rangle$

from a finite set of substitutions) have been discovered, that work for arbitrary word equations [37].

⁴It is standard in the literature on word equations to consider only solutions that map variables to nonempty words. The above procedure is only complete under that assumption. However, allowing the empty word can be easily accommodated. For any equation eq on a set of variables X , and any subset Y of X , let eq_Y be the equation obtained from eq by replacing the variables in Y by the empty word. Let R_Y be a complete set of symbolic solutions for eq_Y where we extend each substitution to X by mapping every variable from Y to the empty word. Then the union of the R_Y is a complete set of symbolic solutions for eq , allowing the empty word. If eq is one-sided nonlinear, then eq_Y is too. This remark equally applies to the extension to path expressions presented in Section 4.3.2.

Furthermore, we now have extra non-successful cases for leaf equations, namely all equations of the form $(@x \cdot w_1 = \langle w_2 \rangle \cdot w_3)$ or $(\langle w_2 \rangle \cdot w_3 = @y \cdot w_1)$.

Extending known arguments, one can see that on any one-sided nonlinear equation, our extended rewriting system terminates and yields a finite complete set of symbolic solutions.

Example 4.8. Figure 2 shows a DAG representation of the search tree for the equation $\$x \cdot \langle @y \cdot \$z \rangle \cdot @w = \$u \cdot \$v \cdot \$u$. There are four successful branches, so the following substitutions comprise a complete set of symbolic solutions:

$$\begin{aligned} &\{\$x \mapsto @w, \$u \mapsto @w, \$v \mapsto \langle @y \cdot \$z \rangle\} \\ &\{\$x \mapsto @w \cdot \$x, \$v \mapsto \$x \cdot \langle @y \cdot \$z \rangle, \$u \mapsto @w\} \\ &\{\$x \mapsto \langle @y \cdot \$z \rangle \cdot @w \cdot \$v, \$u \mapsto \langle @y \cdot \$z \rangle \cdot @w\} \\ &\{\$x \mapsto \$x \cdot \langle @y \cdot \$z \rangle \cdot @w \cdot \$v \cdot \$x, \$u \mapsto \$x \cdot \langle @y \cdot \$z \rangle \cdot @w\} \end{aligned}$$

4.3.3 Pure variables and pure equations

We introduce a syntactic “purity check” on variables, that guarantees that they can only take values that do not contain packed values. Since later we will work stratum per stratum, it is sufficient in what follows to focus on semipositive, nonrecursive programs with only one IDB relation name.

Consider a rule in such a program. When a variable appears in some positive EDB predicate, we call the variable a *source variable* of the rule. Now we inductively define a variable in the rule to be *pure* if

1. it is a source variable (since we focus on flat input instances); or
2. it appears in one side of a positive equation, such that
 - all the variables in the other side of the equation are pure, and
 - the other side of the equation has no packing.

By leveraging associative unification, we are going to show that we can always eliminate impure variables. The method is based on a division of the positive equations of a rule into three categories:

Pure equations involve only pure variables.

Half-pure equations have all variables in one side pure, and at least one of the variables in the other side is impure.

Fully impure equations have impure variables in both sides.

Example 4.9. The three equations in the rule

$$S(\$x) \leftarrow R(\$x, \$y) \wedge \langle \$x \rangle = \langle \$y \rangle \wedge a \cdot \$x = \$z \wedge \$y = \langle \$u \rangle$$

are pure. The two equations in the rule

$$S(\$x) \leftarrow R(\$x, \$y) \wedge \langle \$y \rangle = \$z \wedge \langle \$x \rangle = \langle \$z \rangle$$

are half-pure. The equation $\langle \$t \rangle = \langle \$z \rangle$ in the rule

$$S(\$x) \leftarrow R(\$x, \$y) \wedge \langle \$t \rangle = \langle \$z \rangle \wedge \$z = \langle \$y \rangle \wedge \$t = \langle \$x \rangle$$

is fully impure.

It is instructive to compare the notion of pure variable with that of limited variable, used to define the notion of safe rule. Indeed, the set of limited variables can be equivalently defined as follows, where we only change the base case of the induction to immediately include all pure variables:

- Every pure variable is limited; and
- If all the variables occurring in one side of the sides of an equation in the rule are limited, then all the variables occurring in the other side are also limited.

Therefore, if there is at least one impure variable in a safe rule, then there must be at least one half-pure equation in the rule. In other words, it is not possible for a rule to have fully impure equations without having half-pure ones.

Lemma 4.10. *Let r be a rule in a semi-positive, nonrecursive program \mathbf{P} with only one IDB relation name. Then there exists a finite set of rules, equivalent to r on flat instances, in which all positive equations are pure.*

Proof. By induction on the number of half-pure equations. Let $r : H \leftarrow B \wedge e_1 = e_2$, where $e_1 = e_2$ is half-pure with e_1 the pure side and e_2 the impure side. Let u_1, \dots, u_n be the list of all *occurrences* of variables in e_1 . Let v_1, \dots, v_n be n fresh variables, and let e'_1 be e_1 with each u_i replaced by v_i . Now replace $e_1 = e_2$ by the following conjunction of $n + 1$ equations:

$$u_1 = v_1 \wedge \dots \wedge u_n = v_n \wedge e'_1 = e_2$$

Here, abusing notation, we use the same notation u_i for the variable that occurs at u_i .

Denote the result of this replacement by r' . The equation $e'_1 = e_2$ is one-sided nonlinear; by Section 4.3.2, there exists a finite complete set R of solutions. If we let r'' be r' without $e'_1 = e_2$, then clearly r is equivalent to the set of rules $\{\rho(r'') \mid \rho \in R\}$. However, some of these rules may not have strictly less half-pure equations than r , which is necessary for the induction to work.

We can solve this problem as follows. Call $\rho \in R$ *valid* if it maps variables that are pure in r'' to expressions without packing. Since all u_i and v_i are pure in r'' , the equations $\rho(u_i) = \rho(v_i)$ in $\rho(r'')$ are all pure, so $\rho(r'')$ does have strictly less half-pure equations than r .

Fortunately, we can restrict attention to the valid $\rho \in R$, so the induction goes through. Indeed, following the definition of pure variable, one can readily verify that for nonvalid ρ , the rule $\rho(r'')$ is unsatisfiable on flat instances. \square

4.3.4 Packing structures

By Lemma 4.10, all positive equations can be taken to be pure. We now reduce this further so that all positive equations are free of packing. Thereto we introduce the *packing structure* of a path expression e , denoted by $\delta(e)$, and defined as follows:

- $\delta(\epsilon) = *$.
- $\delta(a) = *$, with a a variable or an atomic value.
- $\delta(\langle e \rangle) = * \cdot \langle \delta(e) \rangle \cdot *$.
- $\delta(e_1 \cdot e_2)$ equals $\delta(e_1) \cdot \delta(e_2)$, in which we replace any consecutive sequence of stars by a single star.

Assume $\delta(e)$ has n stars. Then e can be constructed from $\delta(e)$ by replacing each star by a unique (possibly empty) subexpression of e . We call these subexpressions the components of e . Crucially, they do not use packing.

If e does not use packing, $\delta(e)$ is simply $*$. If e begins or ends with packing, or if some packing in e begins or ends with another packing, then some components will be empty.

Example 4.11. Let $e = @a \cdot \langle \$x \cdot \$y \cdot \$z \rangle \cdot \langle \epsilon \rangle$. Then $\delta(e) = * \cdot \langle * \cdot \langle * \rangle \cdot * \rangle \cdot * \cdot \langle * \rangle \cdot *$. The seven components of e are $@a$, ϵ , $\$x \cdot \y , $\$z$, ϵ , ϵ , and ϵ .

A pure equation $e_1 = e_2$ can only be satisfiable on flat instances if e_1 and e_2 have the same packing structure. Suppose there are n stars in this packing structure. Then, the equation can be replaced by the conjunction of n equations, where we equate the corresponding components of e_1 and e_2 . These equations are still pure, and free of packing.

Moreover, when all positive equations are pure, then all variables in the rule are pure, since the rule is safe. Now a negated equation $e_1 \neq e_2$ over pure variables is equivalent to the disjunction of the nonequalities between the corresponding components of e_1 and e_2 . Then the rule can be replaced by a set of rules, one for each disjunct, and the component nonequalities are free of packing. We can repeat this for all negated equations.

We have arrived at the following:

Lemma 4.12. *Let r be a rule in a semi-positive, nonrecursive program \mathbf{P} with only one IDB relation name. Then there exists a finite set of rules, equivalent to r on flat instances, in which all variables are pure, and all equations (positive or negated) are free of packing.*

4.3.5 Redundancy of packing

We are now ready for the following result. The proof further leverages packing structures.

Lemma 4.13. *Packing is redundant in the absence of recursion.*

Proof. Consider a query computed by a nonrecursive program \mathbf{P} . We must show that \mathbf{P} can be equivalently rewritten without packing. If \mathbf{P} has only one IDB predicate, Lemma 4.12 gives us what we want. Indeed, by the Lemma, we may assume that equations are already free of packing. Now since the input is a flat instance, any positive (negated) EDB predicate that contains packing may be taken to be always false (true). Also, the result of the query is a flat instance, so IDB predicates containing packing are false as well. We thus obtain a program free of packing as desired.

When \mathbf{P} uses intermediate predicates, the elimination of packing from IDB predicates requires more work. Since \mathbf{P} is nonrecursive, we may assume that every stratum involves only one IDB relation name. Since arity is redundant, we may assume that \mathbf{P} does not use arity, but feel free to use arity in the rewriting of \mathbf{P} .

Let us consider the first stratum. For every rule, we proceed as follows. Let $R(e)$ be the head of the rule. Let m be the number of stars in $\delta(e)$ and let e_1, \dots, e_m be the components of e . Replace the head with $R_{\delta(e)}(e_1, \dots, e_m)$ where $R_{\delta(e)}$ is a fresh relation name.

After this step, the rules in the first stratum no longer contain packing in the head. Of course, R -predicates in rules in later strata must now be updated to call the new relation names. So, assume $R(e)$ appears in the body of some later rule r . For each of the packing structures ps introduced for R , we make a copy of r in which we replace $R(e)$ by the conjunction $R_{ps}(\$e_1, \dots, \$e_m) \wedge e = e'$, where

- m is the number of stars in ps ;
- $\$e_1, \dots, \e_m are fresh path variables; and
- e' is obtained from the packing structure ps by replacing the i th star by $\$e_i$, for $i = 1, \dots, m$.

This rewriting introduces equations in later strata, which is necessary because these later strata have not yet been purified per Lemma 4.12.

We do the above for every stratum. So, stratum by stratum, we first remove packing from equations, leaving only pure variables in rules; we replace head predicates; and rewrite calls to these head predicates in later rules.

After this transformation, packing still appears in negated IDB predicates, which have been untouched so far. Fortunately, all rules have pure variables at this point. Thus, a literal $\neg R(e)$, where $\delta(e)$ matches one of the packing structures of R , say ps , with m stars, can now be replaced by $\neg R_{ps}(e_1, \dots, e_m)$, where e_i is the i th component of e . If $\delta(e)$ does not match any of the packing structures introduced for R , the negative literal is true on flat instances and can be omitted.

Observing that packing in EDB predicates can be handled as in the semi-positive case, we are done. \square

Example 4.14. Rewriting the program from Example 2.2 without packing yields a program with 28 rules:

```

T($u, $s, $v) ← R($u·$s·$v), S($s).
A ← T($x1, $x2, $x3), T($y1, $y2, $y3), T($z1, $z2, $z3),
    $xi ≠ $yi, $xj ≠ $zj, $yk ≠ $zk.
% for i=1,2,3, j=1,2,3, k=1,2,3

```

To get from Lemma 4.13 to the following theorem, it remains to show that packing is redundant in the presence of recursion. Building on the flat–flat theorem for J-Logic [22, 23] we can close that gap and we obtain:

Theorem 4.15. *Packing is redundant.*

Proof. It remains to show that P is redundant in the presence of R. Earlier work on J-Logic (flat–flat theorem [22, 23]) is readily adapted to Sequence Datalog and shows that P is redundant in the presence of R and N. The general idea of the rewriting used in that proof is as follows:

1. We add a new stratum at the beginning of the program, where we preprocess the input relations as follows: every path $k_1 \cdot k_2 \cdot \dots \cdot k_n$ is replaced by its doubled version $k_1 \cdot k_1 \cdot k_2 \cdot k_2 \cdot \dots \cdot k_n \cdot k_n$.
2. We modify the program so that it works with doubled EDB and IDB relations. Packing is simulated using a technique of simulated delimiters, which relies on the doubled encoding.
3. In the last step, we undouble the doubled output.

Steps 1 and 3 as published introduce negation even if the original program does not use negation. We next show that this can be avoided. Instead, we introduce arity, which is harmless as arity is redundant.

We double an EDB relation R into R' as follows:

```

T(ε, $x) ← R($x).
T($x·@y·@y, $z) ← T($x, @y·$z).
R'($x) ← T($x, ε).

```

We undouble a doubled output relation S' into S as follows:

```

T($x, ε) ← S'($x).
T($x, @y·$z) ← T($x·@y·@y, $z).
S($x) ← T(ε, $x).

```

□

4.4 Intermediate predicates

The following result is straightforward: intermediate predicates can be eliminated by folding in the bodies of the intermediate rules, using equations to unify calling predicates with intermediate head predicates.

Theorem 4.16. *I is redundant in the presence of E and the absence of N and R.*

5 Inexpressibility results

In this section we show various inexpressibility results that lead to absolute or relative primitive results for various features.

5.1 Recursion

To see that recursion is primitive also in the context of Sequence Datalog, we can make the following observation.

Lemma 5.1. *Let Q be a query that can be computed by a nonrecursive program. Then for any input instance I , the lengths of paths in $Q(I)$ are bounded by a linear function of the maximal length of a path in I .*

Proof. Let \mathbf{P} be a nonrecursive program computing a query Q . Let \mathbf{P}' be \mathbf{P} with all negated literals removed. The Q' query computed by \mathbf{P}' contains Q , so if we can prove the claim for Q' , it also holds for Q .

By Theorem 4.16, we know that Q' is computable by a program \mathbf{P}'' that does not use intermediate predicates. Let n be the number of rules, and for $i = 1, \dots, n$, let $S(e_i)$ be the head of the i th rule; a_i the number of path variables in e_i ; and b_i the number of atomic values and variables in e_i . Then the length of sequences returned by the i th rule is at most $a_i x + b_i$, with x the maximal length of a sequence in the input. The desired linear function can now be taken to be $ax + b$, where $a = \max\{a_i \mid 1 \leq i \leq n\}$ and $b = \max\{b_i \mid 1 \leq i \leq n\}$. \square

We immediately get:

Proposition 5.2. *Let a be a fixed atomic value and let Q be any query from $\{R\}$ to S satisfying the property that for every instance I and every natural number n such that $R(a^n) \in I$, the string a^{n^2} is a substring of a path in $Q(I)$. Then Q is not expressible without recursion.*

We readily obtain:

Theorem 5.3. *Recursion is primitive.*

Proof. First, we show that R is primitive in the presence of l . Consider the following recursive program \mathbf{P} , computing the query Q returning all paths a^{n^2} where n is a natural number such that $R(a^n)$ is in the input:

```

T( $\epsilon$ ,  $\$x$ ,  $\$x$ )  $\leftarrow$  R( $\$x$ ).
T( $\$y\$x$ ,  $\$x$ ,  $\$z$ )  $\leftarrow$  T( $\$y$ ,  $\$x$ ,  $a \cdot \$z$ ).
S( $\$y$ )  $\leftarrow$  T( $\$y$ ,  $\$x$ ,  $\epsilon$ ).

```

By Proposition 5.2, query Q is not expressible without recursion.

The above program uses intermediate predicates. In the absence of this feature, consider just the program \mathbf{P}' consisting of the first two rules. Strictly, this program does not compute a query, as T is ternary. However, we can turn \mathbf{P}' into a program \mathbf{P}'' using the arity simulation technique of Lemma 4.1.

Program \mathbf{P}'' computes a well-defined query Q'' from $\{R\}$ to T . Although Q'' is not a natural query, Proposition 5.2 applies to it, so it is not expressible without recursion. \square

5.1.1 Boolean queries

The above queries showing primitivity of recursion are unary. What about boolean queries? It turns out that for boolean queries, in the presence of intermediate predicates, recursion is still primitive. In the absence of intermediate predicates, however, recursion is redundant for boolean queries, for trivial reasons.

Let us go in a bit more detail. Let R be a binary relation viewed a directed graph. Let $Q_{a \rightarrow b?}$ be the boolean query from $\{R\}$ to S that checks whether b is reachable from a . It is well-known that $Q_{a \rightarrow b?}$, as a classical relational query, is not computable in classical Datalog without recursion. We can view $Q_{a \rightarrow b?}$ as a query on sequence databases by encoding edges (a, b) by paths $a \cdot b$ of length two. Under this encoding, the query is clearly computable by a Sequence Datalog program in the fragment $\{l, R\}$:

```

T( $@x \cdot @y$ )  $\leftarrow$  R( $@x \cdot @y$ ) .
T( $@x \cdot @z$ )  $\leftarrow$  T( $@x \cdot @y$ ), R( $@y \cdot @z$ ) .
S  $\leftarrow$  T( $a \cdot b$ ) .

```

We can now show that $Q_{a \rightarrow b?}$ is not computable without recursion in Sequence Datalog by showing that, on input instances containing only sequences of length two, any nonrecursive Sequence Datalog program can be simulated by a classical nonrecursive Datalog program. This simulation is similar to the one shown in Lemma 5.4 appearing later. The only added complication is that, due to intermediate predicates, sequences of lengths longer than two can appear. However, since there is no recursion, these lengths are bounded by a constant depending only on the program.

In the absence of the l -feature, we note that any boolean query, computed by a recursive program without intermediate predicates, is already computed by the nonrecursive rules only. Indeed, if the result of the query is false, then none of the rules is fired. If, on the other hand, the result of the query is true, then at least one rule is fired; however, no recursive rule can be fired before at least one nonrecursive rule is fired.

5.2 Intermediate predicates

It is well known that in classical Datalog, without intermediate predicates, we can not express queries that require universal quantifiers [11]. We can transfer this result to Sequence Datalog by a simulation technique.

Let Γ be a monadic schema and let I be an instance of Γ . We say that I is “two-bounded” if only paths of lengths one or two occur in I . We can encode two-bounded instances by classical instances as follows. Let Γ^c (c for classical)

be the schema that has two relation names R^1 and R^2 for each $R \in \Gamma$. For I two-bounded as above, we define the classical instance I^c of Γ^c as follows:

- $I^c(R^1) = \{a \in \mathbf{dom} \mid a \in I(R)\};$
- $I^c(R^2) = \{(a, b) \mid a \cdot b \in I(R)\}.$

Lemma 5.4. *Let \mathbf{P} be a program in the fragment $\{E, N, R\}$, with IDB relation name S , such that the result of \mathbf{P} on a two-bounded instance is still two-bounded. Then there exists a semipositive classical Datalog program \mathbf{P}^c using only the IDB relation names S^1 and S^2 , such that for every two-bounded instance I of Γ , we have $\mathbf{P}^c(I^c) = (\mathbf{P}(I))^c$.*

Proof. Our goal is to eliminate path variables as well as concatenations in path expressions. We start with path variables. In any rule containing a head predicate or positive predicate of the form $S(e_1 \cdot \$x \cdot e_2)$ or $R(e_1 \cdot \$x \cdot e_2)$, we can replace $\$x$ either by ϵ , $@x$, or $@x_1 \cdot @x_2$ (splitting the rule in three versions).

Path variables may still occur in equations. By safety, they must appear in positive equations, and inductively we may assume that any remaining path variable $\$x$ occurs in a positive equation $e_1 = e_2$ where e_1 contains no path variables. This equation is then of the form $a_1 \cdots a_n = b_1 \cdots b_m \cdot \$x \cdot e$, where the a s and b s are atomic variables or values.

- If $m = n$, replace $\$x$ by the empty path.
- If $m > n$, the equation is unsatisfiable and the rule can be removed.
- If $m < n$, replace $\$x$ by $a_{m+1} \cdots a_i$, for $m < i \leq n$ (splitting the rule in $n - m + 1$ versions).

After these steps, all equations (positive or negated) are of the form $a_1 \cdots a_n = b_1 \cdots b_m$, where the a s and b s are atomic variables or values. Such equations can be easily eliminated. Moreover, any predicates, possibly negated, that are of the form $R(e)$ with e empty or strictly longer than two, can be eliminated as well.

We finally replace every remaining predicate (head or body) of the form $R(a)$ by $R^1(a)$ and every predicate of the form $R(a_1 \cdot a_2)$ by $R^2(a_1, a_2)$, and we are done. \square

As a consequence, the query computed by the following program, belonging to the fragment $\{I, N\}$, cannot be expressed without intermediate predicates:

$$\begin{aligned} W(@x) &\leftarrow R(@x \cdot @y), \neg B(@y). \\ S(@x) &\leftarrow R(@x \cdot @y), \neg W(@x). \end{aligned}$$

Indeed, the classical counterpart of this query is the query asking, on any directed graph where some nodes are “black”, for all nodes with only edges to black nodes. That query is well-known not to be expressible in semipositive Datalog [11].

We thus obtain:

Theorem 5.5. *I is primitive in the presence of N.*

We also have the following primitivity result in the presence of recursion. The proof merely combines some observations we have already made.

Theorem 5.6. *I is primitive in the presence of R.*

Proof. Recall the squaring query Q from the proof of Theorem 5.3, which is expressible in the fragment $\{I, R\}$. Suppose, for the sake of contradiction, that Q can be computed by a program without intermediate predicates. Consider the behavior of this program on the family of singleton instances $I_n = \{R(a^n)\}$, for all natural numbers n . Since $Q(I_n)$ is nonempty, at least one of the rules must fire, which is only possible if at least one of the nonrecursive rules fires. Since there are no intermediate predicates, however, this nonrecursive rule must already produce the correct output $S(a^{n^2})$. This contradicts Lemma 5.1. \square

5.3 Equations

The two theorems in the previous subsection provide counterparts to Theorem 4.16. The following theorem confirms that the presence of equations is necessary for Theorem 4.16, and implies that the fragments $\{I\}$ and $\{E\}$ are actually equivalent.

Theorem 5.7. *E is primitive in the absence of I.*

This result follows immediately from the following lemma.

Lemma 5.8. *Let a be an atomic value. The boolean query that checks if the input relation R contains a path consisting exclusively of a 's, cannot be computed by a program that lacks features I and E .*

Proof. By the redundancy of packing and arity, we may ignore these features. Also, in Section 5.1.1, we already noted that in the absence of intermediate predicates, recursion does not help in expressing boolean queries. Hence, it suffices to show that the query cannot be computed by a program in the fragment $\{N\}$. For the sake of contradiction, suppose such a program exists.

Take any rule from the program, and consider the instance J obtained from the positive predicates in the body by “freezing” all variables, i.e., viewing them as atomic values distinct from the atomic values already occurring in the rule. Unless the rule is unsatisfiable (in which case we may ignore it), it will fire on J . So the query is true on J and the body must contain a positive predicate of the form $R(a^\ell)$.

Now consider the instance $I = \{R(a^n)\}$ where n is strictly larger than all values ℓ as above found in the rules. Then no rule can fire on I , but the query is true on I , so we have the desired contradiction. \square

Indeed, that query is readily expressed using an equation, as we well know.




Figure 3: If-direction of Theorem 6.1.

6 Putting it all together

The results from the previous two sections allow us to characterize the subsumption relation among fragments (defined in Section 3) as follows.

Theorem 6.1. *For any fragments F_1 and F_2 , we have $F_1 \leq F_2$ if and only if the following five conditions are satisfied:*

1. $N \in F_1 \Rightarrow N \in F_2$;
2. $R \in F_1 \Rightarrow R \in F_2$;
3. $E \in F_1 \Rightarrow (E \in F_2 \vee I \in F_2)$;
4. $(I \in F_1 \wedge R \notin F_1 \wedge N \notin F_1) \Rightarrow (I \in F_2 \vee E \in F_2)$;
5. $(I \in F_1 \wedge (R \in F_1 \vee N \in F_1)) \Rightarrow I \in F_2$.

Proof. For the only-if direction, we verify the five conditions, assuming $F_1 \leq F_2$.

1. Immediate from the primitivity of negation. We have not stated this primitivity as a theorem because it is so clear (any fragment without negation can express only monotone queries; with negation we can express set difference which is not monotone).
2. Immediate from primitivity of recursion.

3. Immediate from Theorem 5.7.
4. Assume $\mathsf{I} \in F_1 \wedge \mathsf{R} \notin F_1 \wedge \mathsf{N} \notin F_1 \wedge \mathsf{E} \notin F_2 \wedge \mathsf{I} \notin F_2$. By Theorem 4.7, we have $\{\mathsf{E}\} \leq F_1$. Now Theorem 5.7 leads to a contradiction with $F_1 \leq F_2$.
5. Immediate from Theorems 5.5 and 5.6.

For the if-direction, since arity and packing are redundant, $F_1 \leq F_2$ if and only if $\hat{F}_1 \leq \hat{F}_2$, where $\hat{F} = F - \{\mathsf{A}, \mathsf{P}\}$. Now Figure 3 (on page 26) infers $\hat{F}_1 \leq \hat{F}_2$ from the five conditions and the redundancy results. \square

7 Sequence relational algebra

Given the importance of algebraic query plans for database query execution, we show here how to extend the classical relational algebra to obtain a language equivalent to recursive-free Sequence Datalog programs. We note that a similar language, while calculus-based rather than algebra-based, is the language StriQuel proposed by Grahne and Waller [20].

The relational algebra, with operators projection; equality selection; union; difference; and cartesian product, is well known [3, 45]. To extend this algebra to our data model (Section 2.1), we generalize the selection and projection operators and add two extraction operators. Let R be an n -ary relation.

Selection: The classical equality selection $\sigma_{\mathsf{s}_i=\mathsf{s}_j}(R)$, with $i, j \in \{1, \dots, n\}$, returns $\{t \in R \mid t_i = t_j\}$. We now allow path expressions α and β over the variables $\mathsf{s}_1, \dots, \mathsf{s}_n$ and have the selection operator

$$\sigma_{\alpha=\beta}(R) := \{t \in R \mid t(\alpha) = t(\beta)\}.$$

Here, t is viewed as the valuation that maps s_i to t_i for $i = 1, \dots, n$.

Projection: For path expressions $\alpha_1, \dots, \alpha_p$ over variables $\mathsf{s}_1, \dots, \mathsf{s}_n$ as above, we define

$$\pi_{\alpha_1, \dots, \alpha_p}(R) := \{(t(\alpha_1), \dots, t(\alpha_p)) \mid t \in R\}.$$

Unpacking: For $i \in \{1, \dots, n\}$, the operator $UNPACK_i(R)$ returns

$$\{(t_1, \dots, t_{i-1}, s, t_{i+1}, \dots, t_n) \mid (t_1, \dots, t_{i-1}, \langle s \rangle, t_{i+1}, \dots, t_n) \in R\}.$$

Substrings: $SUB_i(R)$ equals

$$\{(t_1, \dots, t_n, s) \mid (t_1, \dots, t_n) \in R \text{ and } s \text{ is a substring of } t_i\}.$$

We could also have defined a more powerful unpacking operator, which extracts components from paths using path expressions, similar to the use of path expressions in Sequence Datalog. Such an operator is useful in practice but can for theoretical purposes be simulated using the given operators, as we will show.

“Sequence relational algebra” expressions over a schema Γ , built up using the above operators from the relation names of Γ and constant relations, are defined as usual. We have, as expected, the following theorem. Note that this result applies for arbitrary instances, not only for flat inputs and flat outputs.

Theorem 7.1. *For every program \mathbf{P} without recursion and every IDB relation name T , there exists a sequence relational algebra expression E such that for every instance I , we have $\mathbf{P}(I)(T) = E(I)$. The converse statement holds as well.*

That sequence relational algebra can be translated to Sequence Datalog is clear. Our approach to translate in the other direction is for the most part standard. We can make use of the following normal form.

Lemma 7.2. *Let \mathbf{P} be a nonrecursive Sequence Datalog program that does not use equations. Then there is a nonrecursive program \mathbf{P}' computing the same query as \mathbf{P} where each rule in \mathbf{P}' has one of the following six forms:*

1. $R_1(v_1, \dots, v_n) \leftarrow R_2(e_1, \dots, e_m);$
2. $R_1(v_1, \dots, v_n, e) \leftarrow R_2(v_1, \dots, v_n);$
3. $R_1(v_1, \dots, v_n) \leftarrow R_2(x_1, \dots, x_k), R_3(y_1, \dots, y_\ell);$
4. $R_1(v_1, \dots, v_n) \leftarrow R_2(v_1, \dots, v_n), \neg R_3(v'_1, \dots, v'_m);$
5. $R_1(v'_1, \dots, v'_m) \leftarrow R_2(v_1, \dots, v_n);$
6. $R(p) \leftarrow .$

The following restrictions apply:

- In all forms, v_1, \dots, v_n are distinct variables. Moreover, in forms 2 to 6, each v_i must be a path variable.
- In form 3, the x_i and y_j are path variables and $\{v_1, \dots, v_n\}$ is contained in $\{x_1, \dots, x_k\} \cup \{y_1, \dots, y_\ell\}$.
- In forms 4 and 5, v'_1, \dots, v'_m are distinct variables taken from $\{v_1, \dots, v_n\}$.
- In form 6, p is a path (constant relation).

This lemma is stated for programs without equations, since we know that equations are redundant in the presence of intermediate predicates. Given the normal form, extraction rules of the first form can be expressed in the algebra as follows. First, by compositions of unpacking and substring operations, we can generate all subpaths until the maximum packing depth of the expressions appearing in the rule. Using cartesian product and selection, we then select the desired paths. Rules of the second form are generalized projections. The other rules are handled as in the classical relational algebra. It remains to prove:

Proof of Lemma 7.2. The conversion to normal form is best described on a general example. Consider the following one-rule Sequence Datalog program:

```
T(a.b.c, @x.c.$y, $z.$z) ← P1($y.$y, $z.a, @u.d), P2($z.@x.c, d),
                             -N1(@x.$y.$z, a.@x), -N2(a.b, $y).
```

In what follows, we call the rule that we process the main rule and its stratum the main stratum.

Step 1: Get variables from positive literals

Step 1.1 Replace every occurrence of a positive atom $P(e_1, \dots, e_m)$ by a new predicate $H(v_1, \dots, v_n)$ where $\{v_1, \dots, v_n\}$ is the set of variables used in the atom. For each H add a new rule of the form $H(v_1, \dots, v_n) \leftarrow P(e_1, \dots, e_m)$. Note that these set of rules are guaranteed to be form 1. Moreover, every *atomic* variable in the main rule should be replaced by a new *path* variable.

In case the positive atom does not use variables, then we replace every occurrence by a new predicate $H(\$v)$ with a fresh variable $\$v$. To get this H , we add the two rules $H' \leftarrow P(e_1, \dots, e_m)$ and $H(a) \leftarrow H'$ for a new predicate H' and $a \in \mathbf{dom}$. Note that the first rule is of form 1, while the second added rule is of form 2.

```
H1($y, $z, @u) ← P1($y.$y, $z.a, @u.d).
H2($z, @x) ← P2($z.@x.c, d).
T(a.b.c, $x.c.$y, $z.$z) ← H1($y, $z, $u), H2($z, $x),
                             -N1($x.$y.$z, a.$x), -N2(a.b, $y).
```

Step 1.2

- When no positive atoms exist in the main rule, then the rule has no variables. Only in this case, we add to the main stratum a new rule of the form $R(a) \leftarrow$, where R is a new relation name and a is some value from the domain. This added rule is of form 6. Moreover, we add $R(\$v)$ to the body of the main rule, where $\$v$ is a fresh path variable.
- Otherwise, this step should be repeated until only one positive atom remains in the main rule. We remove two positive atoms $H_i(x_1, \dots, x_n)$ and $H_j(y_1, \dots, y_m)$, and replace them with $H(v_1, \dots, v_k)$, where H is a fresh predicate name, and the set of variables vs is the union of the set of xs and ys . In addition, we introduce a new rule of the form

$$H(v_1, \dots, v_k) \leftarrow H_i(x_1, \dots, x_n), H_j(y_1, \dots, y_m)$$

in the main stratum. This rule is of form 3.

```

H1($y, $z, @u) ← P1($y·$y, $z·a, @u·d).
H2($z, @x) ← P2($z·@x·c, d).
H($y, $z, $u, $x) ← H1($y, $z, $u), H2($z, $x).
T(a·b·c, $x·c·$y, $z·$z) ← H($y, $z, $u, $x),
                               ¬N1($x·$y·$z, a·$x), ¬N2(a·b, $y).

```

Step 2: Separate each negative literal in an intermediate rule.

Step 2.1 Let $H(v_1, \dots, v_n)$ be the only positive atom in the body of the rule. Every literal $\neg N(e_1, \dots, e_m)$ is replaced by a predicate $HN(v_1, \dots, v_n)$, where HN is a new relation name. Moreover, we add a rule of the form

$$HN(v_1, \dots, v_n) \leftarrow H(v_1, \dots, v_n), \neg N(e_1, \dots, e_m)$$

to the main stratum, and we remove $H(v_1, \dots, v_n)$ from the main rule.

```

H1($y, $z, @u) ← P1($y·$y, $z·a, @u·d).
H2($z, @x) ← P2($z·@x·c, d).
H($y, $z, $u, $x) ← H1($y, $z, $u), H2($z, $x).
HN1($y, $z, $u, $x) ← H($y, $z, $u, $x), ¬N1($x·$y·$z, a·$x).
HN2($y, $z, $u, $x) ← H($y, $z, $u, $x), ¬N2(a·b, $y).
T(a·b·c, $x·c·$y, $z·$z) ← HN1($y, $z, $u, $x),
                               HN2($y, $z, $u, $x).

```

Step 2.2 We do the same as in step 1.2, leaving us in the end with a single positive atom holding the variables from the original rule. All the rules introduced by this step are of form 3.

```

H1($y, $z, @u) ← P1($y·$y, $z·a, @u·d).
H2($z, @x) ← P2($z·@x·c, d).
H($y, $z, $u, $x) ← H1($y, $z, $u), H2($z, $x).
HN1($y, $z, $u, $x) ← H($y, $z, $u, $x), ¬N1($x·$y·$z, a·$x).
HN2($y, $z, $u, $x) ← H($y, $z, $u, $x), ¬N2(a·b, $y).
HN($y, $z, $u, $x) ← HN1($y, $z, $u, $x), HN2($y, $z, $u, $x).
T(a·b·c, $x·c·$y, $z·$z) ← HN($y, $z, $u, $x).

```

Step 3: Generate negated expressions. We next work on the rules that were introduced to deal with the negated atoms.

Step 3.1 In step 2.1 we added rules with negative literals:

$$HN(v_1, \dots, v_n) \leftarrow H(v_1, \dots, v_n), \neg N(e_1, \dots, e_m)$$

For each such added rule, we define a sequence of rules in order to generate the values for the expressions e_i . Since our rule is safe from the beginning, we are guaranteed that all the variables used in these expressions are among the v s.

We inductively generate m rules as follows (where the v 's are fresh variables) and add them to the main stratum:

1. $N_1(v_1, \dots, v_n, e_1) \leftarrow H(v_1, \dots, v_n)$

2. for $1 < i \leq m$, the rule

$$N_i(v_1, \dots, v_n, v'_1, \dots, v'_{i-1}, e_i) \leftarrow N_{i-1}(v_1, \dots, v_n, v'_1, \dots, v'_{i-1}).$$

Each one of the above rules is of form 2. In addition, we replace $H(v_1, \dots, v_n)$ in the rule under consideration by

$$N_m(v_1, \dots, v_n, v'_1, \dots, v'_m).$$

Moreover, we replace $\neg N(e_1, \dots, e_m)$ by $\neg N(v'_1, \dots, v'_m)$.

```

H1($y, $z, @u) ← P1($y·$y, $z·a, @u·d).
H2($z, @x) ← P2($z·@x·c, d).
H($y, $z, $u, $x) ← H1($y, $z, $u), H2($z, $x).
N11($y, $z, $u, $x, $x·$y·$z) ← H($y, $z, $u, $x).
N21($y, $z, $u, $x, a·b) ← H($y, $z, $u, $x).
N12($y, $z, $u, $x, $n11, a·$x) ← N11($y, $z, $u, $x, $n11).
N22($y, $z, $u, $x, $n21, $y) ← N21($y, $z, $u, $x, $n21).
HN1($y, $z, $u, $x) ← N12($y, $z, $u, $x, $n11, $n12),
    ¬N1($n11, $n12).
HN2($y, $z, $u, $x) ← N22($y, $z, $u, $x, $n21, $n22),
    ¬N2($n21, $n22).
HN($y, $z, $u, $x) ← HN1($y, $z, $u, $x), HN2($y, $z, $u, $x).
T(a·b·c, $x·c·$y, $z·$z) ← HN($y, $z, $u, $x).

```

Step 3.2 We have now obtained rules of the form

$$HN(v_1, \dots, v_n) \leftarrow N_m(v_1, \dots, v_n, v'_1, \dots, v'_m), \neg N(v'_1, \dots, v'_m).$$

We now further replace them with

$$HN(v_1, \dots, v_n) \leftarrow FN(v_1, \dots, v_n, v'_1, \dots, v'_m);$$

where FN is a new relation name. Now this rule is of form 5. Moreover, we add the rule

$$FN(v_1, \dots, v_n, v'_1, \dots, v'_m) \leftarrow N_m(v_1, \dots, v_n, v'_1, \dots, v'_m), \neg N(v'_1, \dots, v'_m),$$

which is of form 4, to the main stratum.


```

H1($y, $z, @u) ← P1($y·$y, $z·a, @u·d).
H2($z, @x) ← P2($z·@x·c, d).
H($y, $z, $u, $x) ← H1($y, $z, $u), H2($z, $x).
N11($y, $z, $u, $x, $x·$y·$z) ← H($y, $z, $u, $x).
N21($y, $z, $u, $x, a·b) ← H($y, $z, $u, $x).
N12($y, $z, $u, $x, $n11, a·$x) ← N11($y, $z, $u, $x, $n11).
N22($y, $z, $u, $x, $n21, $y) ← N21($y, $z, $u, $x, $n21).
FN1($y, $z, $u, $x, $n11, $n12) ←
    N12($y, $z, $u, $x, $n11, $n12), -N1($n11, $n12).
FN2($y, $z, $u, $x, $n21, $n22) ←
    N22($y, $z, $u, $x, $n21, $n22), -N2($n21, $n22).
HN1($y, $z, $u, $x) ← FN1($y, $z, $u, $x, $n11, $n12).
HN2($y, $z, $u, $x) ← FN2($y, $z, $u, $x, $n21, $n22).
HN($y, $z, $u, $x) ← HN1($y, $z, $u, $x), HN2($y, $z, $u, $x).
T(a·b·c, $x·c·$y, $z·$z) ← HN($y, $z, $u, $x).

```

Step 4: Generate final head expressions. We are now left to work on the final rule which is normalized in a similar way as step 3.1. The final rule is of the form $T(e_1, \dots, e_m) \leftarrow H(v_1, \dots, v_n)$, where by safety it is guaranteed that any variable appearing in any of the es is among the vs .

We inductively generate m rules as follows (where the v 's are fresh variables):

1. $T_1(v_1, \dots, v_n, e_1) \leftarrow H(v_1, \dots, v_n)$
2. for $1 < i \leq m$, the rule

$$T_i(v_1, \dots, v_n, v'_1, \dots, v'_{i-1}, e_i) \leftarrow T_{i-1}(v_1, \dots, v_n, v'_1, \dots, v'_{i-1}).$$

Each one of the above rules is of form 2. The last thing to be done is to update the main rule to

$$T(v'_1, \dots, v'_m) \leftarrow T_m(v_1, \dots, v_n, v'_1, \dots, v'_m).$$

Now, this rule is of form 5.

```

H1($y, $z, @u) ← P1($y·$y, $z·a, @u·d).
H2($z, @x) ← P2($z·@x·c, d).
H($y, $z, $u, $x) ← H1($y, $z, $u), H2($z, $x).
N11($y, $z, $u, $x, $x·$y·$z) ← H($y, $z, $u, $x).
N21($y, $z, $u, $x, a·b) ← H($y, $z, $u, $x).
N12($y, $z, $u, $x, $n11, a·$x) ← N11($y, $z, $u, $x, $n11).
N22($y, $z, $u, $x, $n21, $y) ← N21($y, $z, $u, $x, $n21).
FN1($y, $z, $u, $x, $n11, $n12) ←
    N12($y, $z, $u, $x, $n11, $n12), -N1($n11, $n12).
FN2($y, $z, $u, $x, $n21, $n22) ←
    N22($y, $z, $u, $x, $n21, $n22), -N2($n21, $n22).
HN1($y, $z, $u, $x) ← FN1($y, $z, $u, $x, $n11, $n12).

```

```

HN2($y, $z, $u, $x) ← FN2($y, $z, $u, $x, $n21, $n22).
HN($y, $z, $u, $x) ← HN1($y, $z, $u, $x), HN2($y, $z, $u, $x).
T1($y, $z, $u, $x, a·b·c) ← HN($y, $z, $u, $x).
T2($y, $z, $u, $x, $t1, $x·c·$y) ← T1($y, $z, $u, $x, $t1).
T3($y, $z, $u, $x, $t1, $t2, $z·$z) ←
    T2($y, $z, $u, $x, $t1, $t2).
T($t1, $t2, $t3) ← T3($y, $z, $u, $x, $t1, $t2, $t3).

```

□

8 Conclusion

Sequence databases and sequence query processing (e.g., [42]) were an active research topic twenty years ago or more. We hope our paper can revive interest in the topic, given its continued relevance for advanced database applications. Systems in use today do support sequences one way or another, but often only nominally, without high expressive power or performance. This situation may cause application builders to bypass the database system and solve their problem in an ad-hoc manner.

Of course, to support data science, there is much current research on database systems and query languages for arrays and tensors, e.g., [41, 34, 24, 27, 6]. However, in this domain, applications are typically focused on supporting linear algebra operations [31, 24, 6]. Such applications are qualitatively different from the more generic type of sequence database queries considered in this paper.

We note that other sequence query language approaches, not based on Datalog, deserve attention as well. There have been proposals based on functional programming [30], on structural recursion [40], and on transducers [10, 8, 19, 18]. On the other hand, a proposal very close in spirit to Sequence Datalog can be found in the work by Grahne and Waller [20] already mentioned in Section 7.

Sequence Datalog is also a very useful language for dealing with non-flat instances. In this paper, for reasons we have explained, we focused on queries from flat instances to flat instances. However, using packing, interesting data structures can be represented in a direct manner. For example, a tree with root label a and childtrees T_1, \dots, T_n can be represented by the path $a \cdot \langle T_1 \rangle \cdots \langle T_n \rangle$ (where each T_i is represented by a path in turn). Thus, Sequence Datalog can be used as an XML-to-XML query language and more.

We conclude by recalling an intriguing theoretical open problem already mentioned before [22]. It can be stated independently of Sequence Datalog, although we did stumble upon the problem while thinking about Sequence Datalog. Consider monadic Datalog with stratified negation over sets of *natural numbers*, with natural number constants and variables, and addition as the only operation. Which functions on finite sets of natural numbers are expressible in this language?

References

- [1] H. Aamer, J. Hidders, J. Paredaens, and J. Van den Bussche. Expressiveness within sequence datalog. In *Proceedings 40th/36th ACM Symposium on Principles of Databases*, pages 70–81. ACM, 2021.
- [2] H. Abdulrab and J.-P. Pécuchet. Solving word equations. *Journal of Symbolic Computation*, 8(5):499–521, 1989.
- [3] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [4] M. Alviano and A. Pieris, editors. *Datalog 2.0 2019: Third International Workshop on the Resurgence of Datalog in Academia and Industry*, volume 2368 of *CEUR Workshop Proceedings*, 2019.
- [5] M. Atkinson, F. Bancilhon, D. DeWitt, K. Dittrich, D. Maier, and S. Zdonik. The object-oriented database system manifesto. In W. Kim, J.-M. Nicolas, and S. Nishio, editors, *Proceedings 1st International Conference on Deductive and Object-Oriented Databases*, pages 40–57. Elsevier Science Publishers, 1989.
- [6] P. Barceló, N. Higeura, J. Pérez, and B. Suercaseaux. On the expressiveness of LARA: A unified language for linear and relational algebra. In C. Lutz and J.C. Jung, editors, *Proceedings 23rd International Conference on Database Theory*, volume 155 of *Leibniz International Proceedings in Informatics*, pages 6:1–6:20. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020.
- [7] P. Barceló and R. Pichler, editors. *Datalog in Academia and Industry: Second International Workshop, Datalog 2.0*, volume 7494 of *Lecture Notes in Computer Science*. Springer, 2012.
- [8] M. Benedikt, L. Libkin, T. Schwentick, and L. Segoufin. Definable relations and first-order query languages over strings. *Journal of the ACM*, 50(5):694–751, 2003.
- [9] A. Bonner and G. Mecca. Sequences, Datalog, and transducers. *Journal of Computer and System Sciences*, 57:234–259, 1998.
- [10] A.J. Bonner and G. Mecca. Querying sequence databases with transducers. *Acta Informatica*, 36:511–544, 2000.
- [11] A.K. Chandra and D. Harel. Structure and complexity of relational queries. *Journal of Computer and System Sciences*, 25(1):99–128, 1982.
- [12] J. Chomicki. Temporal query languages: a survey. In D.M. Gabbay and H.J. Ohlbach, editors, *Temporal Logic: ICTL’94*, volume 827 of *Lecture Notes in Computer Science*, pages 506–534. Springer-Verlag, 1994.

- [13] The Committee for Advanced DBMS Function. Third-generation database system manifesto. *SIGMOD Record*, 19(3):31–44, 1990.
- [14] O. de Moor, G. Gottlob, T. Furche, and A. Sellers, editors. *Datalog Reloaded: First International Workshop, Datalog 2010*, volume 6702 of *Lecture Notes in Computer Science*. Springer, 2011.
- [15] F. Durán, S. Eker, S. Escobar, N. Martí-Oliet, J. Meseguer, and C. Talcott. Associative unification and symbolic reasoning modulo associativity in Maude. In V. Rusu, editor, *Proceedings 12th International Workshop on Rewriting Logic and Its Applications*, volume 11152 of *Lecture Notes in Computer Science*, pages 98–114. Springer, 2018.
- [16] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, second edition, 1999.
- [17] R. Fagin, B. Kimelfeld, F. Reiss, and S. Vansummeren. Document spanners: A formal approach to information extraction. *Journal of the ACM*, 62(2):12:1–12:51, 2015.
- [18] S. Ginsburg and X.S. Wang. Regular sequence operations and their use in database queries. *Journal of Computer and System Sciences*, 56(1):1–26, 1998.
- [19] G. Grahne, M. Nykänen, and E. Ukkonen. Reasoning about strings in databases. *Journal of Computer and System Sciences*, 59:116–162, 1999.
- [20] G. Grahne and E. Waller. How to make SQL stand for String Query Language. In R.C.H. Connor and A.O. Mendelzon, editors, *Research Issues in Structured and Semistructured Database Programming*, volume 1949 of *Lecture Notes in Computer Science*, pages 61–79. Springer, 2000.
- [21] M. Grohe. Arity hierarchies. *Annals of Pure and Applied Logic*, 82(2):103–163, 1996.
- [22] J. Hidders, J. Paredaens, and J. Van den Bussche. J-logic: Logical foundations for json querying. In *Proceedings 36th ACM Symposium on Principles of Databases*, pages 137–149. ACM, 2017.
- [23] J. Hidders, J. Paredaens, and J. Van den Bussche. J-logic: a logic for querying json. arXiv:2006.04277, 2020.
- [24] D. Hutchison, B. Howe, and D. Suciu. LaraDB: A minimalist kernel for linear and relational algebra computation. In F.N. Afrati and J. Sroka, editors, *Proceedings 4th ACM SIGMOD Workshop on Algorithms and Systems for MapReduce and Beyond*, pages 2:1–2:10, 2017.
- [25] IEEE Task Force on Process Mining. Process mining manifesto, 2011. URL: <https://www.tf-pm.org/resources/manifesto>.

- [26] H.V. Jagadish and F. Olken. Database management for life science research. *SIGMOD Record*, 33(2):15–20, 2004.
- [27] H. Jananathan, Z. Zhou, et al. Polystore mathematics of relational algebra. In J.-Y. Nie, Z. Obradovic, T. Suzumura, et al., editors, *Proceedings IEEE International Conference on Big Data*, pages 3180–3189. IEEE, 2017.
- [28] Y. Law, H. Wang, and C. Zaniolo. Relational languages and data models for continuous queries on sequences and data streams. *ACM Transactions on Database Systems*, 36(2):8:1–8:32, 2011.
- [29] LDBC Graph Query Language Task Force. G-CORE: A core for future graph query languages. In *Proceedings 2018 International Conference on Management of Data*, pages 1421–1432. ACM, 2018.
- [30] L. Libkin, R. Machlin, and L. Wong. A query language for multidimensional arrays: design, implementations, and optimization techniques. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, volume 25:2 of *SIGMOD Record*, pages 228–239. ACM Press, 1996.
- [31] S. Luo, Z.J. Gao, M.N. Gubanov, L.L. Perez, D. Jankov, and C.M. Jermaine. Scalable linear algebra on a relational database system. *Communications of the ACM*, 63(8):93–101, 2020.
- [32] G. Mecca and A.J. Bonner. Query languages for sequence databases: Termination and complexity. *IEEE Transactions on Knowledge and Data Engineering*, 13(3):519–525, 2001.
- [33] Y. Nahshon, L. Peterfreund, and S. Vansummeren. Incorporating information extraction in the relational database model. In *Proceedings 19th International Conference on Web and Databases*, pages 6:1–6:7. ACM, 2019.
- [34] S. Papadopoulos et al. The TileDB array data storage manager. *Proceedings of the VLDB Endowment*, 10(4):349–360, 2016.
- [35] L. Peterfreund et al. Recursive programs for document spanners. In P. Barcelo and M. Calautti, editors, *Proceedings 22nd International Conference on Database Theory*, volume 127 of *LIPICs*, pages 13:1–13:18. Schloss Dagstuhl–Leibniz Center for Informatics, 2019.
- [36] F. Pezoa, J.L. Reutter, F. Suarez, M. Ugarte, and D. Vrgoč. Foundations of JSON Schema. In *Proceedings 25th International Conference on World Wide Web*, pages 263–273, 2016.
- [37] W. Plandowski. On PSPACE generation of a solution set of a word equation and its applications. *Theoretical Computer Science*, 792:20–61, 2019.
- [38] G. Plotkin. Building-in equational theories. In B. Meltzer and D. Michie, editors, *Machine Intelligence 7*, pages 73–90. Edinburgh University Press, 1972.

- [39] R. Ramakrishnan et al. SRQL: sorted relational query language. In M. Rafanelli and M. Jarke, editors, *Proceedings 10th International Conference on Scientific and Statistical Database Management*, pages 84–95. IEEE Computer Society, 1998.
- [40] E.L. Robertson, L.V. Saxton, D. Van Gucht, and S. Vansummeren. Structural recursion as a query language on lists and ordered trees. *Theory of Computing Systems*, 44:590–619, 2009.
- [41] F. Rusu and Y. Cheng. A survey on array storage, query languages, and systems. arXiv:1302.0103, 2013.
- [42] R. Sadri, C. Zaniolo, A. Zarkesh, et al. Expressing and optimizing sequence queries in database systems. *ACM Transactions on Database Systems*, 29(2):282–318, 2004.
- [43] P. Seshadri, M. Livny, and R. Ramakrishnan. SEQ: A model for sequence databases. In P.S. Yu and A.L.P. Chen, editors, *Proceedings 11th International Conference on Data Engineering*, pages 232–239. IEEE Computer Society, 1995.
- [44] W. Shen et al. Declarative information extraction using Datalog with embedded extraction. In Ch. Koch et al., editors, *Proceedings 33th International Conference on Very Large Data Bases*, pages 1033–1044. ACM, 2007.
- [45] J.D. Ullman. *Principles of Database and Knowledge-Base Systems*, volume I. Computer Science Press, 1988.