ELECTRONIC WORKSHOPS IN COMPUTING

Series edited by Professor C.J. van Rijsbergen

Paolo Atzeni and Val Tannen (Eds)

Database Programming Languages (DBPL-5)

Proceedings of the Fifth International Workshop on Database
Programming Languages, Gubbio, Umbria, Italy, 6-8 September 1995

Paper:

Finite Query Languages for Sequence Databases

Giansalvatore Mecca and Anthony J. Bonner

Published in collaboration with the
British Computer Society

©Copyright in this paper belongs to the author(s)

Finite Query Languages for Sequence Databases

Giansalvatore Mecca*

D.I.F.A., Università della Basilicata

via della Tecnica, 3

85100 Potenza, Italy

mecca@dis.uniroma1.it

Anthony J. Bonner†

Department of Computer Science

University of Toronto

Toronto, Ontario, Canada

bonner@db.toronto.edu

Abstract

This paper develops a query language for sequence databases, such as genome databases and text databases. Unlike relational data, queries over sequential data can easily produce infinite answer sets, since the universe of sequences is infinite, even for a finite alphabet. The challenge is to develop query languages that are both highly expressive and finite. This paper develops such a language. It is a subset of a recently developed logic called *Sequence Datalog* [19]. Sequence Datalog distinguishes syntactically between *subsequence extraction* and *sequence construction*. Extraction creates sequences of bounded length, and leads to *safe* recursion; while construction can create sequences of arbitrary length, and leads to *unsafe* recursion. In this paper, we develop syntactic restrictions for Sequence Datalog that allow sequence construction but preserve finiteness. The main idea is to use safe recursion to control and limit unsafe recursion. The main results are the definition of a finite form of recursion, called *domain bounded recursion*, and a characterization of its complexity and expressive power. Although finite, the resulting class of programs is highly expressive, since its data complexity is complete for the elementary functions.

1 Introduction

It is widely accepted that relational databases do not provide enough support for many of today's advanced applications. In some cases, object-oriented databases [4] are the right solution. However, in other cases, such as genome databases [12] and text databases [14], there is still a need for more flexibility in data representation and manipulation. In these applications, much of the data has an inherently sequential structure. This has several implications for database management systems. First, a DBMS should provide a *sequence* type; that is, it should be able to manipulate sequences of unbounded length over a fixed alphabet. Second, the query languages provided to the user must have powerful primitives and operators for analyzing and restructuring sequences.

Sequences represent a particularly interesting domain for query languages. In contrast to sets, computations over sequences can easily become infinite, even when the underlying alphabet is finite. This is because repetitions of symbols are allowed, so that the number of possible sequences over any finite alphabet is infinite. The researcher thus faces an interesting challenge: on the one hand, the language should provide powerful primitives for restructuring sequences; on the other hand, the expressive power of the language should be carefully limited, to avoid infinite computations.

In [19], we developed a logic called *Sequence Datalog* for querying sequence databases. Two safe subsets of the logic were defined, based on a new computational model called *Generalized Sequence Transducers*. These machines are a simple yet powerful device for computing sequence mappings. In [19], we showed how networks of these machines could be expressed in Sequence Datalog. Moreover, any Sequence Datalog program constructed in this way is guaranteed to be safe and finite. In this paper, we take a different approach: instead of computational definitions, we develop *syntactic restrictions* that guarantee finiteness and safety. This provides an alternate view of finite computations in the logic. The main idea is to use structural recursion (which is guaranteed to terminate) to limit the construction

*Research partially supported by MURST and Consiglio Nazionale delle Ricerche (CNR).

†Research partially supported by an operating grant from the Natural Sciences and Engineering Research Council of Canada (NSERC).

of new sequences. The first result is a syntactically defined class of Sequence Datalog programs that guarantees finiteness and safety. We call these programs *domain bounded programs*. The second result is a characterization of their complexity and expressive power. We prove that domain bounded programs can express any sequence mapping with hyper-exponential time complexity. Thus, although finite, these programs are still highly expressive.

1.1 Overview of Sequence Datalog

Sequence Datalog is an extension of Datalog for manipulating sequences. It uses a simple data model that extends the relational model by allowing tuples of *sequences* in relations, instead of just tuples of constant symbols. This section provides an informal overview of the syntax and semantics of Sequence Datalog. A formal development can be found in [19].

To manipulate sequences, *Sequence Datalog* has two interpreted function symbols for constructing complex terms, one for concatenating sequences and one for extracting subsequences. Intuitively, if X and Y are sequences, and I and J are integers, then the term $X \bullet Y$ denotes the concatenation of X and Y , and the term $X[I : J]$ denotes the subsequence of X extending from position I to position J . To be more precise, the language of terms uses three countable, disjoint sets: a set of constant symbols, $a, b, c, \dots$, called the *alphabet* and denoted Σ ; a set of variables, $R, S, T, \dots$, called *sequence variables* and denoted V_Σ ; and another set of variables, $I, J, K, \dots$, called *index variables* and denoted V_I . A constant sequence (or *sequence*, for short) is an element of Σ^* . From these sets, we construct two kinds of term as follows:

- *index terms* are built from integers, index variables, and the special symbol *end*, by combining them recursively using the binary connectives $+$ and $-$. Thus, if N and M are index variables, then 3 , $N+3$, $N-M$, $end-5$ and $end-5+M$ are all index terms.
- *sequence terms* are built from constant sequences, sequence variables and index terms, by combining them recursively into *indexed terms* and *constructive terms*, as follows:
 - If s is a sequence variable and n_1, n_2 are *index terms*, then $s[n_1:n_2]$ is an *indexed sequence term*. n_1 and n_2 are called the *indexes* of s . As a shorthand, each sequence term of the form $s[n_i:n_i]$ is written $s[n_i]$.
 - If s_1, s_2 are *sequence terms*, then $s_1 \bullet s_2$ is a *constructive sequence term*.

Thus, if S_1 and S_2 are sequence variables, and N is an index variable, then $S_1[4]$, $S_1[1:N]$, and $ecgt \bullet S_1[1:end-1] \bullet S_2$ are all sequence terms.

The semantics of terms is formalized in [19]. Constructive sequence terms have a semantics of concatenation; e.g., $abc \bullet def = abcdef$. Indexed sequence terms have a semantics of subsequence extraction; e.g., $abcdef[2:5] = bcde$ and $abcdef[4:end] = def$. However, there are some subtleties when the index terms take on “fringe” values, as illustrated by the following examples, where ϵ denotes the empty sequence:

$uvwxyz[3 : 6]$ is undefined	$pqrst[3 : end + 1]$ is undefined
$uvwxyz[3 : 5] = wxy$	$pqrst[3 : end] = rst$
$uvwxyz[3 : 4] = wx$	$pqrst[3 : end - 1] = rs$
$uvwxyz[3 : 3] = w$	$pqrst[3 : end - 2] = r$
$uvwxyz[3 : 2] = \epsilon$	$pqrst[3 : end - 3] = \epsilon$
$uvwxyz[3 : 1]$ is undefined	$pqrst[3 : end - 4]$ is undefined

Note that $s[n+1, n] = \epsilon$, while $s[n+i, n]$ is undefined for $i > 1$.

As in most logics, the language of formulas for *Sequence Datalog* includes a countable set of predicate symbols, $p, q, r, \dots$, where each predicate symbol has an associated arity. If p is a predicate symbol of arity n , and $s_1, \dots, s_n$ are sequence terms, then $p(s_1, \dots, s_n)$ is an atom. Moreover, if s_1 and s_2 are sequence terms, then $s_1 = s_2$ and $s_1 \neq s_2$ are also atoms. From atoms, we build *rules*, *facts* and *clauses* in the usual way [17]. The head and body of a clause, γ , are denoted $HEAD(\gamma)$ and $BODY(\gamma)$, respectively. A clause that contains a constructive term in its head is called a

constructive clause. A Sequence Datalog *program* is a set of Sequence Datalog rules in which constructive terms may appear in rule heads, but not in rule bodies.

We say that a variable, X , is *guarded* in a clause if X occurs in the body of the clause as an argument of some predicate. Otherwise, we say that X is *unguarded*. For example, X is guarded in $p(X[1]) \leftarrow q(X)$, whereas it is unguarded in $p(X) \leftarrow q(X[1])$. Because of the active domain semantics, variables in Sequence Datalog clauses need not be guarded.

The semantics of rules is formalized in [19] in terms of least fixpoints. As in classical logic programming [23, 3, 17], each Sequence Datalog program, P , has an associated operator, T_P , that maps databases to databases. Each application of T_P may create new atoms, which may contain new sequences. The operator T_P is monotonic and continuous and has a least fixpoint [19]. This paper develops conditions under which the least fixpoint is finite. In such cases, we say that P has a *finite semantics*.

Example 1.1 [Indexed Terms] The following rule extracts all prefixes of sequences in relation R :

$$prefix(X[1:N]) \leftarrow R(X).$$

For each sequence, X , in R , this rule says that a prefix of X is any subsequence starting with the first element and ending with the N -th element, so long as N is no longer than the length of X . $\square$

The universe of sequences over the alphabet, Σ , is infinite. Thus, to keep the semantics of programs finite, we do not evaluate rules over the entire universe, Σ^* . Instead, we introduce a new active domain for sequence databases, called the *extended active domain*. This domain contains all the sequences occurring in the database, plus all their *subsequences*.¹ Substitutions range over this domain when rules are evaluated.²

The extended active domain is not fixed during query evaluation. Instead, whenever a new sequence is created (by the concatenation operator, $\bullet$), the new sequence—and its subsequences—are added to the extended active domain. The fixpoint theory of Sequence Datalog provides a declarative semantics for this apparently procedural notion. In the fixpoint theory, the extended active domain of the least fixpoint is larger than the extended active domain of the database. For the database, the domain consists of the sequences in the database and all their subsequences. For the least fixpoint, the domain consists of the sequences in the database and any new sequences created during rule evaluation, and all their subsequences.

Example 1.2 [Constructive Terms] The following rule constructs all possible concatenations of sequences in relation R :

$$answer(X \bullet Y) \leftarrow R(X), R(Y).$$

This rule takes any pair of sequences, X and Y , in relation R , concatenates them, and stores the result in *answer*, thereby adding new sequences to the extended active domain. The concatenated sequences (and their subsequences) form the extended active domain of the least fixpoint. $\square$

Compared to Datalog with function symbols, or Prolog, two differences are apparent. The first is that Sequence Datalog has no uninterpreted function symbols, so it is not possible to build arbitrarily nested structures. On the other hand, Sequence Datalog has a richer syntax than the *[Head|Tail]* list constructor of Prolog. This richer syntax is motivated by a natural distinction between two types of recursion, one safe and the other unsafe. Recursion through construction of new sequences is *inherently unsafe* since it can create longer sequences, which can make the active domain grow indefinitely. On the other hand, structural recursion over existing sequences is *inherently safe*, since it only creates shorter sequences, so that growth in the active domain is bounded. In fact, it is bounded by the set of all subsequences of the active domain, which we call the extended active domain. Typically, languages for list manipulation do not discriminate between these two types of recursion. Sequence Datalog does: constructive recursion is performed using constructive terms, of the form $X \bullet Y$, while structural recursion is performed using indexed terms, of the form $X[n_1:n_2]$.

¹ In this paper, we always refer to *contiguous subsequences*, that is, subsequences specified by a start and end position in some other sequence. Thus, *bcd* is a contiguous subsequence of *abcde*, whereas *bd* is not.

² Note that the size of the *extended domain* is at most quadratic in the size of the database domain. In fact, the number of different contiguous subsequences of a given sequence of length k is at most $\sum_{i=0}^{k-1} (k-i)$, that is, $\frac{k(k+1)}{2} + 1$.

1.2 Controlling Constructive Recursion

This section illustrates how sequences are manipulated in Sequence Datalog. The examples develop the idea that constructive recursion (which is unsafe) can be limited and controlled by structural recursion (which is always safe). This is the main idea of this paper, and the basis for the syntactic restrictions developed in Section 4.

Example 1.3 [Pattern Matching] Suppose we are interested in sequences of the form $a^n b^n c^n$ in relation R . The query $answer(X)$ retrieves all such sequences, where the predicate $answer$ is defined by the following rules:

$$\begin{aligned}
 answer(X) &\leftarrow R(X), \\
 &\quad abc_n(X[1:N_1], X[N_1 + 1:N_2], X[N_2 + 1:end]). \\
 abc_n(\epsilon, \epsilon, \epsilon) &\leftarrow true. \\
 abc_n(X, Y, Z) &\leftarrow X[1] = a, Y[1] = b, Z[1] = c, \\
 &\quad abc_n(X[2:end], Y[2:end], Z[2:end]).
 \end{aligned}$$

The formula $answer(X)$ is true iff X is a sequence in R and it is possible to split X in three parts such that abc_n is true. Predicate abc_n is true for every triple of sequences of the form (a^n, b^n, c^n) in the extended active domain of the database. $\square$

In Example 1.3, the semantics of the rulebase is trivially finite for every database, since the rules contain no constructive terms. Thus, any sequence in the least fixpoint is a subsequence of a sequence in the database. In contrast, the next two examples restructure the sequences in the database, producing new sequences longer than any in the database. Example 1.4 does this in a straightforward way, but has an infinite semantics. Example 1.5 solves the same problem, but with a finite semantics.

Example 1.4 [Infinite Semantics] Suppose R is a unary relation containing a set of sequences. For each sequence X in R , we want the sequence obtained by repeating each symbol in X twice. For example, given the sequence $abcd$, we want the sequence $aabbccdd$. We call these sequences *echo* sequences. The easiest way to define echo sequences is with the following program:

$$\begin{aligned}
 answer(X, Y) &\leftarrow R(X), echo(X, Y). \\
 echo(\epsilon, \epsilon) &\leftarrow true. \\
 echo(X, X[1] \bullet X[1] \bullet Z) &\leftarrow echo(X[2:end], Z).
 \end{aligned}$$

The first rule retrieves every sequence in relation R and its echo, by invoking the predicate $echo(X, Y)$. The last two rules specify what an echo sequence is. For every sequence, X , in the extended active domain, these rules generate its echo sequence, Y . Starting with $X = \epsilon$ and $Y = \epsilon$, they recursively concatenate single characters onto X while concatenating two copies of the same character onto Y . As new sequences are generated, they are added to the active domain, which expands indefinitely. $\square$

The program in Example 1.4 has an infinite semantics over every database that contains a non-empty sequence. This is because the rules defining $echo(X, Y)$ recursively generate longer and longer sequences without bound. For example, suppose the input database contains only one tuple, $\{R(aa)\}$. Its extended active domain consists of the sequences ϵ, a, aa . The table below shows how the inferred facts and the extended domain both grow during a bottom up computation of the least fixpoint. Each row in the table is the result of one additional application of the T operator. In each row, the inferred facts contain one more echo entry, and the extended active domain contains one more sequence, consisting entirely of a 's. The least fixpoint of the T operator is therefore infinite, and its extended active domain is the set of all sequences made of a 's. Note that the query answer consists of a single atom, $answer(aa, aaaa)$, which is computed during the fourth step. Thus, although the least fixpoint is infinite, the query answer is not. The next example expresses the query in such a way that both the answer and the least fixpoint are finite.

step	inferred facts	extended domain
0	$R(aa)$	ϵ, a, aa
1	$R(aa), \text{echo}(\epsilon, \epsilon)$	ϵ, a, aa
2	$R(aa), \text{echo}(\epsilon, \epsilon), \text{echo}(a, aa)$	ϵ, a, aa
3	$R(aa), \text{echo}(\epsilon, \epsilon), \text{echo}(a, aa), \text{echo}(aa, aaaa)$	$\epsilon, a, aa, aaa, aaaa$
4	$R(aa), \text{echo}(\epsilon, \epsilon), \text{echo}(a, aa), \text{echo}(aa, aaaa)$ $\text{echo}(aaa, aaaaaa), \text{answer}(aa, aaaa)$	$\epsilon, a, aa, aaa, aaaa$ $aaaaa, aaaaaa$
5	$R(aa), \text{echo}(\epsilon, \epsilon), \text{echo}(a, aa), \text{echo}(aa, aaaa)$ $\text{echo}(aaa, aaaaaa), \text{echo}(aaaa, aaaaaaaa),$ $\text{answer}(aa, aaaa)$	$\epsilon, a, aa, aaa, aaaa$ $aaaaa, aaaaaa$ $aaaaaaa, aaaaaaaa$
...	...	...

Example 1.5 [Finite Semantics] The following program is another way of expressing the query in Example 1.4.

$$\begin{aligned}
 \text{answer}(X, Y) &\leftarrow R(X), \text{echo}(X, Y). \\
 \text{echo}(\epsilon, \epsilon) &\leftarrow \text{true}. \\
 \text{echo}(X[1:N+1], Z \bullet X[N+1] \bullet X[N+1]) &\leftarrow R(X), \text{echo}(X[1:N], Z).
 \end{aligned}$$

In this program, the sequences in relation R act as input for the third rule, which defines the predicate $\text{echo}(X, Y)$. This rule recursively scans each input sequence, X , while constructing an output sequence, Y . For each character in the input sequence, two copies of the character are appended to the output sequence. The rule computes the echo of every prefix of every sequence in R . The first rule then retrieves the echoes of the sequences in R . $\square$

Like Example 1.4, the program in Example 1.5 involves constructive recursion. However, in Example 1.5, the least fixpoint is finite. This is because constructive recursion does not go on indefinitely, but terminates as soon as the input sequences have been scanned. In essence, growing terms of upwardly bounded length are used to guarantee termination: these terms “grow” at each recursive evaluation of the rule, and recursion stops when the upper bound has been reached. In this way, structural recursion over the first argument controls and limits constructive recursion over the second argument. The bottom-up computation is similar to the one in the table above, except that no more echo facts are inferred after the fourth step, whence the least fixpoint is reached and the computation stops.

2 Preliminary Definitions

This section introduces technical definitions used in the rest of the paper, including *sequence database*, *sequence query* and *sequence function*.

Let Σ be a countable set of symbols, called the *alphabet*. Σ^* denotes the set of all possible sequences over Σ , including the *empty sequence*, ϵ . $\sigma_1\sigma_2$ denotes the concatenation of two sequences, $\sigma_1, \sigma_2 \in \Sigma^*$. $\text{LEN}(\sigma)$ denotes the length of sequence σ , and $\sigma(i)$ denotes its i -th element. With an abuse of notation, we blur the distinction between elements of the alphabet and 1-ary sequences.

We say that a sequence, σ' , of length k is a *contiguous subsequence* of sequence σ if for some integer, $i \geq 0$, $\sigma'(j) = \sigma(i+j)$ for $j = 1, \dots, k$. Note that for each sequence of length k over Σ , there are at most $\frac{k(k+1)}{2} + 1$ different contiguous subsequences (including the empty sequence). For example, the contiguous subsequences of the sequence abc are: $\epsilon, a, b, c, ab, bc, abc$.

We now describe an extension of the relational model, in the spirit of [13, 15]. The model allows for tuples containing sequences of elements, instead of just constant symbols. A *relation* of arity k over Σ is a finite subset of the k -fold cartesian product of Σ^* with itself. A *database* over Σ is a finite set of relations over Σ . We assign a distinct predicate symbol, r , of appropriate arity to each relation in a database.

A *sequence query* is a partial mapping from the set of databases over Σ to itself. Given a sequence query, Q , and a database, DB , the result of evaluating Q over DB is denoted $Q(\text{DB})$. Similarly, a *sequence function* [11] is a partial mapping from Σ^* to itself. A sequence function is *computable* if it is partial recursive. Usually, a notion of *genericity* [10] is introduced for queries. The notion can be extended to sequence queries in a natural way. We say that a sequence query Q is *computable* [10] if it is generic and partial recursive.

Sequence functions can be thought of as queries from a database, $\{input(\sigma_{in})\}$, containing a single sequence tuple, to a database, $\{output(\sigma_{out})\}$, containing a single sequence tuple. Expressibility results formulated in terms of sequence functions are especially meaningful for sequence query languages, since they provide a clear characterization of the power of the language to manipulate sequences. A sequence query language cannot express complex queries over sequence databases if it cannot express complex sequence functions. In short, function expressibility is necessary for query expressibility.

This paper addresses the complexity of sequence functions, and the data complexity [25] of sequence queries. Given a sequence function, f , the *complexity* of f is defined in the usual way, as the complexity of computing $f(\sigma)$, measured with respect to the length of the sequence σ . Given a sequence query, Q , a database, DB , and a suitable encoding of DB as a Turing machine tape, the *data complexity* of Q is the complexity of computing an encoding of $Q(DB)$, measured with respect to the size of DB . A query language, L , is *complete* in the complexity class C if: (i) each query expressible in L has complexity in C ; (ii) there is a query, Q , expressible in L such that computing $Q(DB)$ is a complete problem for the complexity class C .

This paper also addresses the ability of query languages to express sequence functions. A query language, L , is said to *express* a complexity class, C , of sequence functions if: (i) each sequence function expressible in L has complexity in C and conversely, (ii) each sequence function with complexity in C can be expressed in L .

3 The Finiteness Problem

As discussed in the previous sections, computations over sequences may become infinite even when the underlying alphabet is finite. We are interested in studying *finite* programs, that is, programs that have a finite semantics over every input database.

Definition 1 [Finite Programs] *A program is finite if it has a finite semantics (i.e. a finite least fixpoint) over every database.*

As it is typical of powerful logics, the *finiteness property* for Sequence Datalog programs is in general undecidable [19]. Thus, our aim is to develop *subsets* of the logic that are finite. In [19], we took what might be called a “semantic” approach, defining finite subsets in terms of abstract computing devices, called *generalized sequence transducers*. In this paper, we take a syntactic approach, defining finite subsets in terms of syntactic restrictions.

We first note that a necessary condition for infiniteness is the generation of sequences of unbounded length, as in Example 1.4. To do this, programs must use recursion through construction. That is, newly computed sequences must be used recursively to construct more new sequences. This kind of computation is closely related to a particular form of constructive rule, which we call *recursive constructive rules*. In such rules, the predicate in the head depends on itself. To formalize this concept, we introduce the notion of a *predicate dependency graph* of a Sequence Datalog program. This notion, and several others, are closely related:

- Predicate p is a *constructive predicate* in program P if P contains a constructive rule for p , that is, a rule with a constructive term (a term containing $\bullet$) in its head. Note that constructive predicates cause new sequences to be added to the domain during query evaluation.
- Predicate p *depends on* predicate q in program P if P contains a rule in which p is the predicate symbol in the head and q is a predicate symbol in the body. If the rule is constructive, then p *depends constructively* on q .³
- The *predicate dependency graph*, PDG_P , of program P is a directed graph representing the binary relation “*depends on*” over the predicates of P . An edge (p, q) in this graph is a *constructive edge* if p depends constructively on q .
- Predicate p is *recursive with respect to construction* in program P if the predicate dependency graph for P contains a cycle passing through p with a constructive edge.

³For clarity’s sake, this paper supposes that only guarded variables occur in programs. With some added complexity, the definitions can be adjusted to the more general case.

Example 3.1 Consider the following two Sequence Datalog programs, P_1 and P_2 :

$$P_1 : \begin{cases} p(X) \leftarrow r(X, Y), q(Y[2:end]). \\ q(X) \leftarrow r(X, Y), p(Y[1:N]). \\ r(X \bullet Y, Y \bullet X) \leftarrow a(X, Y). \end{cases}$$

$$P_2 : \begin{cases} p(X) \leftarrow q(X). \\ q(X) \leftarrow r(X). \\ r(X \bullet Y) \leftarrow p(X), t(X, Y). \end{cases}$$

Both programs are recursive, so their predicate dependency graphs both have cycles. The graph of P_2 has a constructive cycle, while the graph of P_1 does not. Thus, P_2 is recursive wrt constructions, while P_1 is not. $\square$

The simplest way to enforce finiteness in the presence of constructive rules is to disallow recursion through construction. This means forbidding programs whose predicate dependency graph contains cycles with constructive edges. Intuitively, this means we can find a *stratification* of the program with respect to constructive rules. In this case, the least fixpoint of the program is finite for every database, since there is no way to construct new sequences of unbounded length. We have the following result about the resulting language, which we call *Stratified Sequence Datalog*.

Theorem 1 *Stratified Sequence Datalog is data complete for PTIME.*

Although Stratified Sequence Datalog is complete for PTIME, it has a very limited ability to restructure sequences. Many natural, simple and low-complexity restructurings—such as reversing a sequence or computing its complement—require constructive recursion, and cannot be expressed in Stratified Sequence Datalog. Intuitively, reversing a sequence of length n requires n concatenations (to concatenate the n characters together in reverse order). Likewise for complementation. For these operations, the number of concatenations depends on the database. In contrast, in Stratified Sequence Datalog, the number of concatenations is fixed for each program and is independent of the database.

To increase expressiveness while preserving finiteness, the next section develops syntactically restricted forms of recursion.

4 Domain Bounded Recursion

In some patterns of recursion, the length of newly constructed sequences is bounded above; that is, the recursive construction of new sequences proceeds up to a certain length and then stops, so that the semantics is finite. This section develops one of these patterns, which we call *domain bounded recursion*. In this form of recursion, the length of constructed sequences is bounded above by the size of the active domain of the database, that is, by the sum of the lengths of all sequences in the database. Recursion therefore stops after a finite amount of time, depending on the size of the domain.

In previous examples of finite programs, the length of a computed sequence depended only on the length of a single sequence in the database (as when computing the echo of a sequence, in Example 1.5). In the next example, the length of a computed sequence depends on *two* database sequences.

Example 4.1 [Shuffle] We define a *shuffle* of two sequences, σ_1 and σ_2 , to be any sequence obtained by interleaving the symbols of σ_1 with those of σ_2 . For example, the shuffles of ab and 12 are $ab12, a1b2, a12b, 12ab, 1a2b, 1ab2$. The program below computes all the shuffles of pairs of sequences from the unary predicate R . It defines a predicate $shuffle(X, Y, Z)$ to be true iff Z is a shuffle of X and Y , where X and Y are sequences in R . The length of Z is thus the sum of the lengths of X and Y . For this reason, recursion stops after a finite amount of time.

$$\begin{aligned} \gamma_1 : \quad & shuffle(\epsilon, \epsilon, \epsilon) \leftarrow true. \\ \gamma_2 : \quad & shuffle(X[1:N+1], Y[1:M], X[N+1] \bullet Z) \leftarrow R(X), R(Y), \\ & \quad \quad \quad shuffle(X[1:N], Y[1:M], Z). \\ \gamma_3 : \quad & shuffle(X[1:N], Y[1:M+1], Y[M+1] \bullet Z) \leftarrow R(X), R(Y), \\ & \quad \quad \quad shuffle(X[1:N], Y[1:M], Z). \end{aligned}$$

In this program, pairs of sequences from relation R act as inputs to the last two rules, which define $shuffle(X, Y, Z)$. These rules recursively scan the input sequences X and Y while constructing an output sequence, Z . Starting with the empty sequence, these rules nondeterministically scan a symbol in X or a symbol in Y , and then append this symbol to the growing shuffle sequence, Z . $\square$

Before defining *Domain Bounded Recursion*, we need to develop some preliminary notions.

4.1 Reasoning about Length

To determine if a program is finite, we need to reason about what it does. In particular, we need to reason about the lengths of any new sequences created by the program. If these lengths can be bounded, then the program is finite. This section develops a simple formalism for comparing the “lengths” of two sequence terms. The idea is that terms can be compared on the basis of their possible instantiations. For instance, if X is a sequence variable, then we would say that the term $X \bullet abc$ is “longer” than the term X . This means that any instantiation of the one term is longer than the corresponding instantiation of the other term. This idea will allow us to reason about programs, and to develop conditions under which they are guaranteed to be finite.

As a first step, we develop the notion of the *symbolic length* of a sequence term. This is an arithmetic expression in which symbols and numbers can appear. For example, if X is a sequence variable, then its symbolic length is the symbol L_X . Likewise, if $X \bullet Y$ is a sequence term, then its symbolic length is $L_X + L_Y$. The symbolic length of a constant sequence (e.g., $actg$) is its actual length (e.g., 4). Such expressions allow us to reason about the lengths of partially specified sequences. The reasoning is tractable because we are dealing with just a tiny subset of arithmetic.

To reason about sequence terms such as $X[N : M]$, we need to reason about the index terms N and M . We therefore introduce the notion of the *symbolic value* of an index term. Like symbolic lengths, symbolic values are arithmetic expressions in which numbers and symbols can appear. For example, if N is an index variable, then its symbolic value is the symbol V_N . In general, the symbolic value of an index term depends on the sequence term in which it is embedded. For example, in the sequence term $X[N : end]$, the index term end represents the last character in the sequence X . Thus, in the term $actg[2 : end]$, the symbolic value of end is 4, while in the term $actgactg[2 : end]$, its symbolic value is 8. The following definition makes these ideas precise.

Definition 2 [Symbolic Length and Value] *The symbolic length of a sequence term, s , is an arithmetic expression, denoted $L(s)$. The symbolic value of an index term, n , in the context of s is also an arithmetic expression, denoted $V(n, s)$. These expressions are built from integers, two binary connectives (+ and -), and a collection of symbols. They are constructed in a mutually recursive fashion as follows:*

– Symbolic Lengths:

- If s is a constant sequence in Σ^* , then $L(s)$ is the length of s .
e.g., $L(actg) = 4$.
- If X is a sequence variable in V_Σ , then $L(X)$ is the symbol L_X .
- $L(s_1 \bullet s_2)$ is the expression $L(s_1) + L(s_2)$.
e.g., $L(X \bullet atcg) = L_X + 4$
- $L(s[n_1:n_2])$ is the expression $V(n_2, s) - V(n_1, s) + 1$.

– Symbolic Values:

- If n is an integer, then $V(n, s) = n$.
- If N is an index variable in V_I , then $V(N, s)$ is the symbol V_N .
- $V(end, s) = L(s)$.
e.g., $V(end, X) = L_X$ $V(end, atcgatcg) = 8$
- $V(n_1 \pm n_2, s)$ is the expression $V(n_1, s) \pm V(n_2, s)$.
e.g., $V(N + 3, X) = V_N + 3$ $V(end - N, X) = L_X - V_N$

Here are some sequence terms and their symbolic lengths:

s	$L(s)$
$X[3:N]$	$V_N - 3 + 1$
$X[3:end]$	$L_X - 3 + 1$
$X[3:N] \bullet Y[N:end]$	$(V_N - 3 + 1) + (L_Y - V_N + 1)$
$atcgatcgatcg[3 + N:end - M]$	$(12 - V_M) - (3 + V_N) + 1$

Symbolic lengths can be manipulated in a variety of ways. For instance, we can add and subtract two symbolic lengths to obtain another symbolic length. In some situations, we can also evaluate a symbolic length to obtain an integer. For example, if a symbolic length contains only integers and no symbols, then it can be evaluated in the normal way. Even if a symbolic length contains symbols, these symbols may cancel out, so the expression can be evaluated; *e.g.*, the value of $L_X + 4 - L_X - 2$ is 2. This gives two well-defined situations in which symbolic lengths can be evaluated. In fact, these are the only such situations. After all, if a symbol does not cancel itself out, then a symbolic length cannot be evaluated; *e.g.*, the expression $L_X - L_Y + 3$ cannot be evaluated. This idea gives us a mechanism with which to compare two symbolic lengths.

Definition 3 [Comparability] Two sequence terms s_1, s_2 are comparable if the symbolic expression $L(s_1) - L(s_2)$ can be evaluated,⁴ to yield an integer, k . If $k > 0$ then s_1 is longer than s_2 . If $k = 0$ then s_1 is the same length as s_2 .

Example 4.2 [Comparing Sequence Terms] The sequence term $s_1 = actg \bullet Y$ is longer than the sequence term $s_2 = Y$. To see this, note that $L(s_1) = 4 + L_Y$ and $L(s_2) = L_Y$, so $L(s_1) - L(s_2)$ evaluates to 4, a positive integer. Similarly, the term $s_3 = X[5:N]$ is longer than the term $s_4 = X[5:N - 3]$. In this case, $L(s_3) = V_N - 5 + 1$ and $L(s_4) = V_N - 3 - 5 + 1$, so $L(s_3) - L(s_4)$ evaluates to 3, a positive integer.

On the other hand, the terms $s_5 = S[1:N]$ and $s_6 = S[1:M]$ are incomparable. To see this, note that $L(s_5) - L(s_6)$ reduces to $V_N - V_M$, which cannot be evaluated. Similarly, the terms $s_7 = X \bullet Y$ and $s_8 = Y$ are incomparable. In this case, $L(s_7) - L(s_8)$ reduces to L_X , which cannot be evaluated. $\square$

4.2 Constrained Variables

Another notion that we need is *constrained variables*. Intuitively, we need to infer when a variable ranges over a fixed domain that does not grow during query evaluation [24]. For example, in the rule $p(X[1 : 3]) \leftarrow q(X)$, the variable X is constrained, since it is forced to range over sequences in relation q . However, in the rule $p(X) \leftarrow q(X[1 : 3])$, variable X is *not* constrained. To see this, suppose the database contains the fact $q(abc)$. Then the index term $X[1 : 3]$ can take on the value abc , which means that X can be any sequence that has abc as a *prefix*. Thus, X can range over an infinite domain, including sequences of unbounded length. As another example, consider the following rule:

$$p(Y) \leftarrow q(X), Y = X[N : M]$$

In this case, X is a constrained variable, and so is Y . First, X is constrained to range over the sequences in relation q ; and then Y is constrained to range over subsequences of X .

These ideas motivate the following definition. In this definition, and throughout this paper, we use the notation (p, i) to refer to the i^{th} attribute (or argument) of predicate p .

Definition 4 [Constrained Variables] Let γ be a rule, let S be a sequence variable, and let p be a predicate occurring in the body of γ . We say that S is constrained by attribute (p, i) in γ if at least one of the following holds:

- variable S is the i^{th} argument of some occurrence of predicate p in the body of γ ;

⁴ Sometimes, two symbolic lengths can be intuitively compared even if their difference cannot be evaluated. For example, we can easily conclude that every instantiation of the term $s_1 = X \bullet Y \bullet a$ is longer than the corresponding instantiation of the term $s_2 = X$, even though their difference $L(s_1) - L(s_2) = L_Y + 1$ cannot be evaluated. However, we shall ignore this possibility, since taking it into account would complicate the definitions and the theoretical development without increasing the expressibility of the formalism.

- the body of γ contains an equality atom of the form $S = S_1[N_1:N_2]$ where S_1 is constrained by (p, i) in γ .

The ideas of constrained variables and symbolic length lead to a simple formulation of structural recursion in Sequence Datalog. For example, consider the following rules, where q is a base predicate:

$$\begin{aligned} p(X, \epsilon) &\leftarrow q(X). \\ p(X, X[1 : N + 1]) &\leftarrow q(X), p(X, X[1 : N]). \end{aligned}$$

Here, $p(X, Y)$ is true iff X is a sequence in q and Y is a prefix of X . To see this, note that if X is a sequence in q , then $p(X, \epsilon)$ is true, by the first rule. Then, using the second rule, $p(X, X[1 : 1])$ is true, then $p(X, X[1 : 2])$ is true, then $p(X, X[1 : 3])$, and so on up to $p(X, X)$. After this, $X[1 : N]$ is undefined, so recursion stops. The rules thus scan each sequence in q from beginning to end, which is a canonical example of structural recursion. There are two points to observe here. (i) In both rules, variable X is *constrained* by the predicate q . (ii) The second argument of p *grows* with each bottom-up application of the rules. The notion of “growth” can be made precise by comparing the symbolic lengths of terms in the head and body of a rule using Definition 3. In this case, $X[1 : N + 1]$ is longer than $X[1 : N]$. The following definition generalizes this idea.

Definition 5 [Growing Attributes] Suppose predicate p occurs in the head of a rule and occurs once in the rule body. For each k , attribute (p, k) grows in the rule if the sequence term in attribute (p, k) in the head is longer than the sequence term in attribute (p, k) in the body. In addition, attribute (p, k) does not shrink in the rule if the sequence term in attribute (p, k) in the head is longer than or the same length as the sequence term in attribute (p, k) in the body.

4.3 Domain Bounded Recursion: Definition

We have now developed the concepts needed to define Domain Bounded Recursion. The idea is to allow recursion through construction, but in a controlled and limited way. The result is a class of Sequence Datalog programs that we call *domain bounded programs*. This subset is defined in terms of four restrictions on Sequence Datalog programs. The first two restrictions are not strictly necessary, since they can be generalized without much difficulty; however, they simplify the theoretical development. The last two restrictions are the heart of domain bounded recursion. They are based on the ideas of constrained variables and symbolic length, developed above. The restrictions all apply to recursive constructive rules, that is, to recursive rules that have a constructive term in the head.

The first restriction we impose is that in recursive constructive rules, the head must have exactly one constructive argument. A constructive argument is an argument that contains a constructive term. Thus, the left-hand rule below is allowed, but the right-hand rule is not:

$$p(X, X \bullet Y) \leftarrow q(X), p(X, Y). \quad p(X \bullet Y, X \bullet Y) \leftarrow q(X), p(X, Y).$$

The second restriction we impose on Sequence Datalog is that recursive constructive rules be linear. Recall that a rule is *linear* iff the predicate in the head is mutually recursive with the predicate of at most one atom in the body [6]. Actually, we require more than mere linearity, since we disallow *mutual* recursion through construction.⁵ Thus, the predicate symbol in the head of a recursive constructive rule must also occur in the body of the rule. We call this *restricted linear recursion* through construction. For example, if q is a base predicate, then the following rule is restricted linear:

$$p(X, X \bullet Y) \leftarrow q(X), p(X, Y).$$

The rulebases in Examples 4.1 and 1.5, defining the predicates *shuffle* and *echo*, are both restricted linear. This property of a program can easily be checked in polynomial time (polynomial in the number of rules).

Note that mutual recursion and non-linear recursion are still allowed. However, they are not allowed in constructive rules. We thus have all the power of classical Datalog at our disposal (since Datalog is a subset of Sequence Datalog). Moreover, abolishing mutual recursion through construction does not limit our expressive power, since mutual recursion can always be reduced to non-mutual recursion.

⁵Technically, all constructive cycles in the attribute dependency graph of a program can involve only a single node. Thus, mutually recursive rules are allowed (as in Example 4.1), but mutually recursive predicates are not.

The third restriction we impose is based on an idea we call *constructive variables*. Consider a recursive constructive rule. As above, suppose the rule is restricted linear and has exactly one constructive term in the head. Suppose this term is in argument (p, k) . Because the rule is restricted linear, predicate p also occurs in the body, exactly once. Argument (p, k) in the body contains at most one sequence variable, Z (since constructive terms are not allowed in rule bodies). If Z also occurs in argument (p, k) in the head, then we call Z the *constructive variable* of the rule, e.g. Variable Z is the constructive variable of predicate *shuffle* in Example 4.1. In other words, Z is passed from argument (p, k) in the body to argument (p, k) in the head, and in the process, some other sequence is appended to it. This property allows the rule to recursively construct new sequences, which can lead to the unbounded generation of new sequences. It is this kind of behavior that we want to limit and control using structural recursion.

Definition 6 [Domain Bounded Program] A Sequence Datalog program is domain bounded if every recursive constructive rule in the program satisfies all the following conditions:

1. the head has exactly one constructive argument;
2. the rule is restricted linear;
3. except for the constructive variable (if any), every sequence variable is constrained;
4. there is some non-constructive argument that grows in the rule, and every other non-constructive argument does not shrink in the rule.

As mentioned earlier, only items 3 and 4 in this definition are strictly necessary. The others serve to simplify the development. (A more general treatment will soon be available in [18].)

Example 4.3 [Continued Shuffling] Consider the predicate *shuffle* in Example 4.1. There are two recursive constructive rules this program, γ_2 and γ_3 . Both are restricted linear and both have one constructive argument in the head. For both rules, Z is the constructive variable, and all other sequence variables (X and Y) are constrained by predicate R . Finally, in rule γ_2 , the first argument grows and the second does not shrink; and in rule γ_3 , the second argument grows and the first does not shrink. The program is therefore domain bounded. $\square$

5 Complexity and Expressibility

In this section we prove that domain bounded programs are finite. We shall actually prove a stronger result, showing that, although finite, domain bounded recursion is highly expressive. In fact, domain bounded programs can generate sequences of exponential length, as the following example shows.

Example 5.1 [Long Sequences] The following program is domain bounded:

$$\begin{aligned} \text{doubling}(\epsilon, 1) &\leftarrow \text{true}. \\ \text{doubling}(X[1:N+1], Y \bullet Y) &\leftarrow \text{input}(X), \text{doubling}(X[1:N], Y). \end{aligned}$$

Given a sequence, σ , of length n in predicate *input*, the predicate *doubling* computes a sequence of 1's of length 2^n by doubling the length of a unit sequence n times. $\square$

We now prove complexity and expressibility results for domain bounded programs, and show that they capture exactly the class of *elementary sequence functions* [20], that is, the class of sequence functions with hyper-exponential time complexity.

Definition 7 [Elementary Functions] The class of elementary sequence functions, $\mathcal{E}$, is defined in terms of the hyper-exponential functions, $\text{hyp}_i(n)$. These latter functions are defined recursively as follows:

- $\text{hyp}_1(n) = n$
- $\text{hyp}_{i+1}(n) = 2^{\text{hyp}_i(n)}$, for $i > 1$

hyp_i is called the hyper-exponential function of level i . The set of elementary sequence functions is the set of sequence functions that have hyper-exponential time complexity, that is, the set of sequence functions:

$$\mathcal{E} = \bigcup_{i \geq 1} \text{DTIME}[hyp_i(\mathcal{O}(n))]$$

The following theorem characterizes the expressibility of domain bounded programs in terms of these functions. Recall from Section 2 that a language L is said to express a class of sequence functions, C , if (i) each sequence function expressible in L has complexity in C , and conversely, (ii) each sequence function with complexity in C can be expressed in L .

Theorem 2 (Expressibility of Domain Bounded Programs) *Domain bounded programs express the class $\mathcal{E}$ of elementary sequence functions.*

Proof: (Sketch) To prove the lower expressibility bound, we show that given any Turing machine that runs in hyper-exponential time, its computations can be encoded as a domain bounded program. The program generates a counter of hyper-exponential length using a technique similar to that in Example 5.1. This establishes a lower complexity bound, which leads directly to a lower expressibility bound, since we are dealing with sequence *functions*, and not more-general sequence *queries*.

To prove the upper expressibility bound, we show that domain bounded programs can compute at most a hyper-exponential number of sequences, and that the maximum length of these sequences is at most hyper-exponential in the size of the initial database. Thus, the domain size of the least fixpoint is hyper-exponential in domain size of the database. The least fixpoint can therefore be computed in hyper-exponential time. $\square$

Acknowledgments

The authors would like to thank Paolo Atzeni, Victor Vianu and the anonymous reviewers for useful suggestions, which helped improve both the technical development and the presentation.

References

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of databases*. Addison Wesley Publ. Co., Reading, Massachusetts, 1994.
- [2] S. Abiteboul and V. Vianu. Datalog extensions for database queries and updates. *Journal of Computing and System Sciences*, 43(1):62–124, August 1991.
- [3] K. Apt and M. H. Van Emden. Contributions to the theory of logic programming. *Journal of the ACM*, 29(3):841–862, 1982.
- [4] M. Atkinson, F. Bancilhon, D. DeWitt, K. Dittrich, D. Maier, and Z. Zdonik. The object-oriented database manifesto. In *First Intern. Conference on Deductive and Object Oriented Databases (DOOD'89), Kyoto, Japan*, pages 40–57, 1989.
- [5] P. Atzeni, editor. *LOGIDATA+: Deductive Databases with Complex Objects, Lecture Notes in Computer Science 701*. Springer-Verlag, 1993.
- [6] F. Bancilhon and R. Ramakrishnan. An amateur's introduction to recursive query processing strategies. In *ACM SIGMOD International Conf. on Management of Data*, pages 16–52, 1986.
- [7] A. J. Bonner. Hypothetical Datalog: complexity and expressibility. *Theoretical Computer Science*, 76:3–51, 1990.

- [8] V. Breazu-Tannen, P. Buneman, and S. Naqvi. Structural recursion as a query language. In *Third Intern. Workshop on Database Programming Languages (DBPL'91)*, pages 9–19, 1991.
- [9] V. Breazu-Tannen, P. Buneman, and L. Wong. Naturally embedded query languages. In *Fourth International Conference on Data Base Theory, (ICDT'92), Lecture Notes in Computer Science*, pages 140–154, 1992.
- [10] A. K. Chandra and D. Harel. Computable queries for relational databases. *Journal of Computing and System Sciences*, 21:333–347, 1980.
- [11] L. S. Colby, E. L. Robertson, L. V. Saxton, and D. Van Gucht. A query language for list-based complex objects. In *Thirteenth ACM SIGMOD Intern. Symposium on Principles of Database Systems (PODS'94)*, pages 179–189, 1994.
- [12] Communications of the ACM. Special issue on the Human Genome project. vol. 34(11), November 1991.
- [13] S. Ginsburg and X. Wang. Pattern matching by RS-operations: towards a unified approach to querying sequence data. In *Eleventh ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS'92)*, pages 293–300, 1992.
- [14] G. H. Gonnet. Text dominated databases: Theory, practice and experience. Tutorial presented at PODS, 1994.
- [15] G. Grahne, M. Nykanen, and E. Ukkonen. Reasoning about strings in databases. In *Thirteenth ACM SIGMOD Intern. Symposium on Principles of Database Systems (PODS'94)*, pages 303–312, 1994.
- [16] S. Grumbach and T. Milo. An algebra for POMSETS. In *Fifth International Conference on Data Base Theory, (ICDT'95), Prague, Lecture Notes in Computer Science*, pages 191–207, 1995.
- [17] J. W. Lloyd. *Foundations of Logic Programming*. Springer-Verlag, second edition, 1987.
- [18] G. Mecca and A. J. Bonner. Domain bounded recursion. In preparation.
- [19] G. Mecca and A. J. Bonner. Sequences, Datalog and Transducers. In *Fourteenth ACM SIGMOD Intern. Symposium on Principles of Database Systems (PODS'95), San Jose, California*, pages 23–35, 1995.
- [20] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [21] J. Richardson. Supporting lists in a data model (a timely approach). In *Eighteenth International Conference on Very Large Data Bases (VLDB'92), Vancouver, Canada*, pages 127–138, 1992.
- [22] D. Stott Parker, E. Simon, and P. Valduriez. SVP – a model capturing sets, streams and parallelism. In *Eighteenth International Conference on Very Large Data Bases (VLDB'92), Vancouver, Canada*, pages 115–126, 1992.
- [23] M. H. Van Emden and R. A. Kowalski. The semantics of predicate logic as a programming language. *Journal of the ACM*, 23(4):733–742, 1976.
- [24] A. Van Gelder and R. Topor. Safety and translation of relational calculus queries. *ACM Transactions on Database Systems*, 16:235–278, 1991.
- [25] M. Vardi. The complexity of relational query languages. In *Fourteenth ACM SIGACT Symp. on Theory of Computing*, pages 137–146, 1988.