

# On the Complexity of Nonrecursive XQuery and Functional Query Languages on Complex Values

Christoph Koch  
DBAI, Institut für Informationssysteme  
Technische Universität Wien  
A-1040 Vienna, Austria  
koch@dbai.tuwien.ac.at

## ABSTRACT

*This paper studies the complexity of evaluating functional query languages for complex values such as monad algebra and the recursion-free fragment of XQuery.*

*We show that monad algebra with equality restricted to atomic values is complete for the class  $TA[2^{O(n)}, O(n)]$  of problems solvable in linear exponential time with a linear number of alternations. The monotone fragment of monad algebra with atomic value equality but without negation is complete for nondeterministic exponential time. For monad algebra with deep equality, we establish  $TA[2^{O(n)}, O(n)]$  lower and exponential-space upper bounds.*

*Then we study a fragment of XQuery, Core XQuery, that seems to incorporate all the features of a query language on complex values that are traditionally deemed essential. A close connection between monad algebra on lists and Core XQuery (with “child” as the only axis) is exhibited, and it is shown that these languages are expressively equivalent up to representation issues. We show that Core XQuery is just as hard as monad algebra w.r.t. combined complexity, and that it is in  $TC_0$  if the query is assumed fixed.*

## 1. INTRODUCTION

Complex values form part of various data models for advanced database applications, such as object-oriented, object-relational, and semistructured data models. A large amount of theoretical work on query languages for complex values has been done (e.g. [25, 28, 1, 23, 20, 37, 19, 24, 6, 2, 34, 41, 9, 30, 39, 10]), and this has laid the foundations for object-oriented query languages as well as SQL 1999 or XQuery.

Earlier complexity studies on query languages for complex values have almost entirely focused on logic- [29] and particularly logic programming-based query languages [39, 10, 8], and fixpoint languages (e.g. [20]). However, the query languages considered by many researchers to be most natural for complex values (such as *complex value algebra without powerset* [1, 3], its syntactic variant *monad algebra* [37, 6], and XQuery) are functional.

**Monad algebra.** Monad algebra was shown expressively

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PODS 2005 June 13-15, 2005, Baltimore, Maryland.

Copyright 2005 ACM 1-59593-062-0/05/06 ... \$5.00.

equivalent to a number of other important complex-value query languages such as *nested relational algebra* [25] and complex value algebra without powerset in previous research [37]. (Complex value algebra *with powerset* [28, 1, 19] can take hyperexponential runtime. Queries that really need the powerset operator are usually too costly to evaluate.)

Since some of these languages were developed driven by practical requirements rather than from first principles as is the case for monad algebra, it appears that the expressiveness of these languages on complex values is “the right one” to many researchers and plays a role analogous to that of the power of first-order logic (or relational algebra) on the relational model.

One known result [36] is that monad algebra is in  $TC_0$  w.r.t. *data complexity* (i.e., if the query is assumed fixed [38]). However, the complexity of monad algebra if the query is assumed variable (*query/combined complexity* [38]) is open. In this paper, we study the complexity of monad algebra under the latter assumption.

**XQuery.** XQuery is destined to become the dominant data-transformation query language for XML data and to take a role analogous to the one occupied by SQL for relational databases.

It is folklore that full XQuery is Turing-complete, but it is also obvious that queries without recursion are guaranteed to terminate in straightforward functional implementations of the XQuery language. Recursion in XQuery is rarely used in practice (see also [43]); recursive XML transformations are usually implemented in XSLT.

In essence, XQuery is a quite natural typed functional programming language for XML; still it is sometimes criticized by the research community as huge and clumsy. In this paper we study a substantial recursion-free fragment of XQuery, which we call Core XQuery<sup>1</sup>. It seems that Core XQuery contains all and only the features one would expect from a functional query language for unranked trees in the spirit of complex-value algebra without powerset.

Little foundational research on XQuery has been done to date. There are only some cautious first attempts at finding clean formalizations of and algebras for the language [22, 11, 42]. Most other recent work has focused on engineering good query processors for XQuery [31, 32, 14, 12, 27].

In this paper, we attempt a first closer look at the complexity of XQuery, or more precisely, of the Core XQuery

<sup>1</sup>This fragment is not to be confused with the XQuery Core [42], which is a much larger fragment of XQuery that is also harder to study.

	with negation	without negation
deep equality	in EXPSPACE; TA[ $2^{O(n)}$ , $O(n)$ ]-hard	
equality on atomic values	TA[ $2^{O(n)}$ , $O(n)$ ]-complete	NEXPTIME-complete

**Table 1: Summary of results on query/combined complexity for monad algebra and Core XQuery.**

fragment. We attempt to do this in a principled manner, establishing connections to earlier, well-studied formalisms for functional queries on complex-value databases [34, 37, 6, 41, 18, 30]. Indeed our results on the complexity of monad algebra quite directly yield a characterization of the complexity of Core XQuery.

The technical contributions of this paper are as follows.

- We show that monad algebra on sets, lists and bags is complete for TA[ $2^{O(n)}$ ,  $O(n)$ ] w.r.t. combined complexity in the presence of negation and equality on atomic values.
- We show that monad algebra on sets, lists and bags with equality on atomic values but without negation is NEXPTIME-complete.
- For the case of monad algebra with deep equality, we obtain an EXPSPACE upper bound. A TA[ $2^{O(n)}$ ,  $O(n)$ ] lower bound follows from the fact that negation is easily definable using deep equality.
- We introduce the Core XQuery language, a simple yet powerful nonrecursive fragment of XQuery.
- We exhibit a close connection between XQuery and monad algebra on lists and show that the Core XQuery queries that use only the child axis for navigation in data trees capture monad algebra on lists up to representation issues.<sup>2</sup>

The established mappings are efficiently computable. This allows us to make use of our complexity characterization of monad algebra, but it also gives a very concise formal semantics to Core XQuery.

- We show that if equality is restricted to atomic values, Core XQuery is complete for TA[ $2^{O(n)}$ ,  $O(n)$ ].
- Core XQuery with deep equality is in EXPSPACE and hard for the class TA[ $2^{O(n)}$ ,  $O(n)$ ]. Since we can again very directly express negation using deep equality, this result holds even if negation and universal quantification (“every”) are ruled out from the language.
- We show that the monotone fragment of Core XQuery – without negation and with equality restricted to atomic values – is NEXPTIME-complete.
- Finally, we show that Core XQuery is in TC<sub>0</sub> w.r.t. data complexity.

Table 1 summarizes our complexity results for query and combined complexity.

<sup>2</sup>Monad algebra on nested lists (which are in fact the same as unranked, ordered, *unlabeled* trees) is incomparable with Core XQuery in the strict sense because we have excluded position arithmetics and attributes from Core XQuery and thus cannot simulate the tuples of monad algebra that are essential to its expressiveness.

To the best of the author’s knowledge, this is the first work characterizing the complexity of XQuery. The mappings to and from monad algebra also give an argument that Core XQuery is a well-designed language that offers the “right” degree of expressive power.

**Related work.** It seems that the most relevant work regarding this problem – apart from the characterization of the data complexity of monad algebra in [36] – is on the complexity of nonrecursive logic programming.

For nonrecursive logic programming, a full complexity characterization [8] has been obtained for the most common forms of complex values (that is, values built from sets, lists, bags, tuples, and atomic values.) and various classes of logic programs (with and without negation, range-restriction, and types). It turns out that the complexity of nonrecursive logic programming is robustly (for various kinds of complex values, and with or without range-restriction) NEXPTIME-complete. In the presence of negation (and necessarily range-restriction), nonrecursive logic programming is known to be in the class TA[ $2^{O(n)}$ ,  $O(n)$ ] [39] and hard for the class TA[ $2^{O(n/\log n)}$ ,  $O(n/\log n)$ ] [40, 8].

A main difference between functional languages such as monad algebra and XQuery and logic programming as studied in [9, 39] is the form of nonmonotonicity employed. In functional languages that have the power to check the equality of complex values, negation is usually a redundant operation. Equality introduces nonmonotonicity into the functional languages, while the seemingly same deep equality is innocuous in logic programming languages. Nonmonotonicity in the functional languages is different and seemingly more powerful than that obtained through negation in nonrecursive normal logic programming. For example, in monad algebra, we can compute two sets of doubly exponential size in two different ways (using a “map” operation that applies a transformation to each element of a set) and then check their equality, while the upper bounds on the complexity of nonrecursive logic programming rely on the fact that unifiers in SLD resolution proofs of nonrecursive logic programs cannot grow beyond singly exponential size.

The work [20, 29] is on more expressive query languages. [29] proves *LDM logic* without powerset complete for the class TA[ $2^{O(n)}$ ,  $O(n)$ ]. Differently from monad algebra, LDM logic is a logical language with quantification, operates on cyclic data, and cannot express deep equality.

**Structure.** The structure of this paper is as follows. Section 2 discusses some necessary complexity-theoretic background and introduces complex values and monad algebra (on sets). Section 3 studies the complexity of monad algebra on sets, focusing on upper bounds, while Section 4 provides the corresponding lower bounds. Section 5 briefly discusses the complexity of monad algebra on lists and bags. Section 6 defines the Core XQuery fragment and provides efficiently computable mappings between monad algebra on lists and XQuery. Sections 7 and 8 present our results on combined and data complexity of XQuery, respectively.

## 2. PRELIMINARIES

### 2.1 Complexity-Theoretic Background

By  $AC_0$  we refer to the class of languages recognizable by LOGSPACE-uniform families of circuits using and- and or-gates of unbounded fan-in of polynomial size and constant depth. By  $TC_0$  we refer to the same class except that in addition so-called majority-gates are permitted, which compute “true” iff more than half of their inputs are true. For details on circuit complexity and the notion of uniformity we refer to [17, 26].

We assume deterministic, nondeterministic, and alternating Turing machines known and refer to e.g. [26] for definitions. By  $DTIME[t(n)]$  and  $NTIME[t(n)]$ , we denote the classes of all problems solvable in time  $t(n)$  (where  $n$  is the size of the input) on deterministic and nondeterministic Turing machines, respectively. By  $DSPACE[s(n)]$ , we denote the classes of all problems solvable in space  $s(n)$  on deterministic Turing machines. By  $TA[t(n), a(n)]$ , we denote the class of problems solvable in time  $t(n)$  using  $a(n)$  alternations on alternating Turing machines.

We will use the following abbreviations for complexity classes in this paper:

$$\begin{aligned} \text{NETIME} &= \text{NTIME}[2^{O(n)}] \\ \text{NEXPTIME} &= \text{NTIME}[2^{n^{O(1)}}] \\ \text{2ETIME} &= \text{DTIME}[2^{2^{O(n)}}] \\ \text{2EXPTIME} &= \text{DTIME}[2^{2^{n^{O(1)}}}] \\ \text{LOGSPACE} &= \text{DSPACE}[O(\log n)] \\ \text{EXSPACE} &= \text{DSPACE}[2^{n^{O(1)}}] \end{aligned}$$

It is known that  $AC_0 \subseteq TC_0 \subseteq \text{LOGSPACE} \subseteq \text{NEXPTIME} \subseteq \text{TA}[2^{n^{O(1)}}, 1] \subseteq \text{TA}[2^{n^{O(1)}}, n^{O(1)}] \subseteq \text{TA}[2^{n^{O(1)}}, 2^{n^{O(1)}}] = \text{EXSPACE} \subseteq \text{2EXPTIME}$ . Moreover,

$$\text{NETIME} \subseteq \text{TA}[2^{O(n)}, O(n)] \subseteq \text{2ETIME} \subseteq \text{2EXPTIME}$$

(cf. e.g. [26, 7]).

$\text{NETIME}$  and  $\text{2ETIME}$  are not robust complexity classes – they are not closed under LOGSPACE-reductions, as can be verified using a simple padding argument and the Time Hierarchy theorem [21]. We will consider completeness for those classes as well as of  $\text{TA}[2^{O(n)}, O(n)]$  under *LOGLIN-reductions*, under which they are known to be closed (cf. e.g. [8]). By a LOGLIN reduction, we denote a LOGSPACE reduction that produces output of linear size.  $\text{TA}[2^{O(n)}, O(n)]$  is known to be closed under LOGLIN reductions and has important complete problems from logic, such as deciding the Theory of Real Addition [5, 13].

### 2.2 Complex Values and Monad Algebra

We now introduce monad algebra on sets; monad algebra on lists and bags will be briefly sketched in Section 5.

We consider complex values constructed from sets, tuples, and atomic values from a single-sorted domain<sup>3</sup>. Types are terms of the grammar

$$\tau ::= \text{Dom} \mid \{\tau\} \mid \langle A_1 : \tau_1, \dots, A_k : \tau_k \rangle$$

where  $k \geq 0$ .

<sup>3</sup>All results in this paper immediately generalize to many-sorted domains.

Consider the query language on complex values consisting of expressions built from the following operations (the types of the operations are provided as well):

1. identity

$$\text{id} : x \mapsto x \quad \tau \rightarrow \tau$$

2. composition<sup>4</sup>

$$f \circ g : x \mapsto g(f(x)) \quad \frac{f : \tau \rightarrow \tau', g : \tau' \rightarrow \tau''}{f \circ g : \tau \rightarrow \tau''}$$

3. constants from  $\text{Dom} \cup \{\emptyset, \langle \rangle\}$  ( $\langle \rangle$  is the nullary tuple)

4. singleton set construction

$$\text{sng} : x \mapsto \{x\} \quad \tau \rightarrow \{\tau\}$$

5. application of a function to every member of a set

$$\text{map}(f) : X \mapsto \{f(x) \mid x \in X\}$$

$$\frac{f : \tau \rightarrow \tau'}{\text{map}(f) : \{\tau\} \rightarrow \{\tau'\}}$$

6. flatten:  $X \mapsto \bigcup X \quad \{\{\tau\}\} \rightarrow \{\tau\}$

7. pairing

$$\text{pairwith}_{A_1} : \langle A_1 : X_1, A_2 : x_2, \dots, A_n : x_n \rangle \mapsto \langle \langle A_1 : x_1, A_2 : x_2, \dots, A_n : x_n \rangle \mid x_1 \in X_1 \rangle$$

$$\langle A_1 : \{\tau_1\}, A_2 : \tau_2, \dots, A_n : \tau_n \rangle \rightarrow \langle \langle A_1 : \tau_1, \dots, A_n : \tau_n \rangle \rangle$$

( $\text{pairwith}_{A_i}$  for  $i > 1$  is defined analogously.)

8. tuple formation

$$\langle A_1 : f_1, \dots, A_n : f_n \rangle : x \mapsto \langle A_1 : f_1(x), \dots, A_n : f_n(x) \rangle$$

$$\frac{f_1 : \tau \rightarrow \tau_1, \dots, f_n : \tau \rightarrow \tau_n}{\langle A_1 : f_1, \dots, A_n : f_n \rangle : \tau \rightarrow \langle A_1 : \tau_1, \dots, A_n : \tau_n \rangle}$$

9. projection

$$\pi_{A_i} : \langle A_1 : x_1, \dots, A_i : x_i, \dots, A_n : x_n \rangle \mapsto x_i$$

$$\pi_{A_i} : \langle A_1 : \tau_1, \dots, A_n : \tau_n \rangle \rightarrow \tau_i$$

The language has a nice theoretical foundation from programming language theory, that of structural recursion on sets extended by a small amount of machinery for creating and destroying tuples [37]. Formally, the language above is a Cartesian category with a *strong monad* on it (where “strong” refers to so-called *tensorial strength* introduced by the “pairwith” operation). We call this language *monad algebra* [37], or  $\mathcal{M}$  for short.

We will use  $\text{flatmap}(f)$  as a shortcut for  $\text{map}(f) \circ \text{flatten}$ . Observe that projection  $\pi$  is applied to tuples rather than to sets of tuples as in relational algebra. For example, the relational algebra expression  $\pi_{AB}$  corresponds to the expression  $\text{map}(\langle A : \pi_A, B : \pi_B \rangle)$  in  $\mathcal{M}$ .

<sup>4</sup>Again, our convention throughout the paper is that  $(f \circ g)(x) = g(f(x))$ , not  $f(g(x))$ .

EXAMPLE 2.1. The Cartesian product  $f \times g$  can be defined as  $\langle 1 : f, 2 : g \rangle \circ \text{pairwith}_1 \circ \text{flatmap}(\text{pairwith}_2)$ .

Observe the difference from the product of relational algebra. For instance, the query  $\text{id} \times \text{id}$  on a set of pairs  $S$  computes the set  $\{\langle \langle x_1, x_2 \rangle, \langle x_3, x_4 \rangle \rangle \mid \langle x_1, x_2 \rangle, \langle x_3, x_4 \rangle \in S\}$  rather than  $\{\langle x_1, x_2, x_3, x_4 \rangle \mid \langle x_1, x_2 \rangle, \langle x_3, x_4 \rangle \in S\}$ .  $\square$

It is customary to define Boolean queries (“predicates”) as queries that produce values of type  $\{\langle \rangle\}$ , i.e., that either return  $\{\langle \rangle\}$  (“true”) or  $\emptyset$  (“false”) [37]. Note that the logical conjunction  $\gamma \wedge \delta$  of two predicates  $\gamma$  and  $\delta$  can be computed as  $\gamma \times \delta$ .

By *positive monad algebra*  $\mathcal{M}_\cup$ , we denote  $\mathcal{M}$  extended by the set union operation  $\cup$ . This language has a number of nice properties [37, 6], but it is known to be incomplete as a practical query language because it cannot yet express an equality predicate

$$(A_i = A_j) : \langle A_1 : \tau_1, \dots, A_k : \tau_k \rangle \rightarrow \{\langle \rangle\}.$$

However, if we extend  $\mathcal{M}_\cup$  by any nonempty subset of the operations equality ( $A = B$ ), testing set membership ( $A \in B$ ) or containment ( $A \subseteq B$ ), selection  $\sigma_{A=B}$ , set difference “ $-$ ”, set intersection  $\cap$ , or nesting<sup>5</sup>, we always get the same expressive power.<sup>6</sup> We will call any one of these extended languages *full monad algebra*.

THEOREM 2.2 ([37]).  $\mathcal{M}_\cup[=] \equiv \mathcal{M}_\cup[\sigma] \equiv \mathcal{M}_\cup[-] \equiv \mathcal{M}_\cup[\cap] \equiv \mathcal{M}_\cup[\subseteq] \equiv \mathcal{M}_\cup[\in] \equiv \mathcal{M}_\cup[\text{nest}]$ .

Moreover, generalizing selections to test against constants or to support “ $\in$ ”, “ $\subseteq$ ”, or Boolean combinations of conditions does not increase the expressiveness of full monad algebra [37].

EXAMPLE 2.3. Given a Boolean predicate  $\gamma$ , selection  $\sigma_\gamma$  can be expressed as  $\text{flatmap}(\langle 1 : \text{id}, 2 : \text{id} \circ \gamma \rangle \circ \text{pairwith}_2 \circ \text{map}(\pi_1))$ .

Predicate ( $A \subseteq B$ ) can be expressed in  $\mathcal{M}_\cup[=]$  as

$$\langle A : \pi_A, A' : \pi_A \cap \pi_B \rangle \circ (A = A')$$

where  $f \cap g := (f \times g) \circ \sigma_{1=2} \circ \text{map}(\pi_1)$ . A predicate ( $f \subseteq g$ ) can of course be expressed as  $\langle 1 : f, 2 : g \rangle \circ (1 \subseteq 2)$ .  $\square$

EXAMPLE 2.4. Given a complex value of type

$$\langle R : \{\tau\}, S : \{\tau\} \rangle,$$

difference  $R - S$  can be implemented in  $\mathcal{M}_\cup[\sigma]$  as

$$\text{pairwith}_R \circ \text{map}(\langle R : \pi_R, S_R : \langle R : \pi_R, S : \pi_S \rangle \circ \text{pairwith}_S \circ \sigma_{R=S} \rangle) \circ \sigma_{S=\emptyset} \circ \text{map}(\pi_R).$$

The idea is to compute, for each element  $r$  of  $R$ , the set  $S_R$  of elements in  $S$  that are equal to  $r$  and then to select those elements  $r$  of  $R$  for which  $S_R$  is empty.  $\square$

<sup>5</sup>The “nest” operation of complex value algebra without powerset [1] groups tuples by some of their attributes. For example,  $\text{nest}_{C=(B)}(\bar{R})$  on relation  $R(AB)$  computes the value  $\{\langle A : x, C : \{\langle B : y \rangle \mid \langle A : x, B : y \rangle \in R \} \rangle \mid (\exists y) \langle A : x, B : y \rangle \in R\}$ .

<sup>6</sup>No analogous statement can be made about flat relational algebra.

Theorem 2.2 demonstrates that full monad algebra (w.l.o.g.,  $\mathcal{M}_\cup[=]$ ) is a very robust notion. It can serve as an “expressiveness benchmark” for query languages on complex-value databases. Indeed, it has been shown that full monad algebra is a *conservative extension* of relational algebra.<sup>7</sup>

THEOREM 2.5 ([34]). *A mapping from a (flat) relational database to a (flat) relation is expressible in  $\mathcal{M}_\cup[=]$  if and only if it is expressible in relational algebra.*

There are a number of alternative ways of stating the query evaluation problem. In this paper, we study the complexity of Boolean queries. For XQuery, we will study the problem of deciding whether the root node (which must be always present in a valid XQuery result) of the resulting XML tree has children.

In the following, we will discuss three kinds of complexity of query evaluation, *data complexity* (where queries are assumed to be fixed and data variable), *query complexity* (where the query is variable and the data is assumed to be fixed), and *combined complexity* (where both data and query are considered variable) [38].

### 3. COMPLEXITY OF MONAD ALGEBRA

We will study the complexity of full monad algebra  $\mathcal{M}_\cup[=]$  as well as monotone fragments. It is folklore that by extending  $\mathcal{M}_\cup$  by equality on atomic values  $=_{\text{atomic}}$ , we still cannot express nonmonotone operations such as equality of sets or negation. We can safely generalize  $=_{\text{atomic}}$  to equality of arbitrary complex values that do not include sets,  $=_{\text{mon}}$ , defined inductively as  $=_{\text{atomic}}$  on atomic values and  $v_1 =_{\text{mon}} w_1 \wedge \dots \wedge v_k =_{\text{mon}} w_k$  on tuples  $\langle v_1, \dots, v_k \rangle$  and  $\langle w_1, \dots, w_k \rangle$ . Of course this generalization does not improve upon the expressiveness of  $\mathcal{M}_\cup[=_{\text{atomic}}]$ .

PROPOSITION 3.1.  $\mathcal{M}_\cup[=_{\text{atomic}}]$  captures  $\mathcal{M}_\cup[=_{\text{mon}}]$ .

We start our complexity study with data complexity. It is quite easy to conclude from Theorem 2.5 (conservativity over relational algebra) that the data complexity of  $\mathcal{M}_\cup[\sigma]$  must be rather low.

PROPOSITION 3.2 (FOLKLORE, [36]).  $\mathcal{M}_\cup[=]$  is in TC0 w.r.t. data complexity.

We will now show that the query complexity of monad algebra is substantially higher. Actually, it is possible to write queries that compute values of doubly exponential size.

PROPOSITION 3.3. *There is an  $\mathcal{M}_\cup$  query  $Q$  that computes a value of size  $2^{2^{\Omega(|Q|)}}$ .*

**Proof.** Consider the query  $Q$

$$\phi_{\{0,1\}} \circ \underbrace{(\text{id} \times \text{id}) \circ \dots \circ (\text{id} \times \text{id})}_{m \text{ times}}$$

where  $\phi_{\{0,1\}} = (0 \circ \text{sng}) \cup (1 \circ \text{sng})$  computes the set  $\{0, 1\}$  and  $m$  is linear in  $|Q|$ . Query  $Q$  computes the set of all nested pairs ( $=$ binary trees) of depth  $m$  with labels from  $\{0, 1\}$  at the leaves. There are  $2^{2^m}$  such nested pairs.  $\square$

For the converse,

<sup>7</sup>A generalized version of Theorem 2.5 can be found in [41].

PROPOSITION 3.4. *The values computable by  $\mathcal{M}_{\cup[=]}$  queries are of size  $2^{2^{O(n)}}$ , where  $n$  is the size of the input (i.e., database and query).*

COROLLARY 3.5.  *$\mathcal{M}_{\cup[=]}$  is in 2ETIME w.r.t. combined complexity.*

This is easy to see because given an input value of size  $2^{2^{O(n)}}$ , each operation of  $\mathcal{M}_{\cup[=]}$  can be evaluated on the input in time  $2^{2^{O(n)}}$  on a random access machine. There are  $|Q| \leq n$  operations, and  $|Q| \cdot 2^{2^{O(n)}} = 2^{2^{O(n)}}$ .

Since monad algebra has the power to construct arbitrary values from scratch, we will use the following proposition and will subsequently focus on query complexity. As all complexity classes we will consider for query and combined complexity throughout this paper will be closed under LOGLIN-reductions, combined complexity is no harder than query complexity.

PROPOSITION 3.6. *There is a LOGLIN reduction that, given a complex value  $v$ , computes an  $\mathcal{M}_{\cup}$  expression that evaluates to  $v$  on an arbitrary (e.g. empty) database.*

The main upper bound results of this paper follow next.

THEOREM 3.7.  *$\mathcal{M}_{\cup[=atomic]}$  is in NEXPTIME w.r.t. query complexity.*

**Proof Sketch.** Without loss of generality, we may assume that all operations of the given monad algebra query are unary. This requires only a slight change of notation when we use the union operation  $\cup$ : Rather than writing  $f \cup g$ , we write  $\langle A : f, B : g \rangle \circ \cup$ .

The proof is by a LOGSPACE-reduction to the *success problem* of nonrecursive logic programming with function symbols (but without sets), i.e. the problem of deciding whether a distinguished boolean predicate evaluates to true. This problem is known to be NEXPTIME-complete [9].

We now come to address the observation made in the introduction that while monad algebra queries may compute complex values of doubly exponential size (Proposition 3.3), resolution proofs for nonrecursive logic programs are always of only singly exponential size [9]. We show that every  $\mathcal{M}_{\cup[=atomic]}$  query can be reduced to a non-recursive logic program with a single binary function symbol  $f$ . For convenience, we write terms built using  $f$  in a contrived list representation (paths); for example, the term  $f(f(x, y), f(z, f(u, v)))$  will be written as  $(x.y).z.u.v$ . Left  $f$ -term children are considered Skolem functions generating new path labels. For example,  $(x.y).z.u.v$  is understood as a path  $w.z.u.v$  where  $w$  is a label generated from and identified by  $x.y$ .

We view every complex value as a deterministic tree, i.e., a tree in which each node  $v$  is uniquely identified by the path of labels from the root to  $v$ . We are able to uniquely assign such labels – even the elements of an index set to the elements of a set value, as we are considering query complexity and construct every set value from scratch (see Proposition 3.6). Such a deterministic tree is of course fully described by the set of root-to-leaf paths occurring in it.

We can now give an alternative semantics of  $\mathcal{M}_{\cup[=atomic]}$  in terms of deterministic trees, that is, each query maps a

deterministic tree given as a set of paths to a deterministic tree given as a set of paths, with

$$\begin{aligned} \llbracket \text{id} \rrbracket(V) &:= V \\ \llbracket c \rrbracket(V) &:= \{c\} \\ \llbracket \pi_A \rrbracket(V) &:= \{v \mid A.v \in V\} \\ \llbracket \text{sng} \rrbracket(V) &:= \{s.v \mid v \in V\} \\ \llbracket f \circ g \rrbracket(V) &:= \llbracket g \rrbracket(\llbracket f \rrbracket(V)) \\ \llbracket \text{flatten} \rrbracket(V) &:= \{(i.j).v \mid i.j.v \in V\} \\ \llbracket A =_{atomic} B \rrbracket(V) &:= \{\langle \rangle \mid A.v, B.v \in V\} \end{aligned}$$

$$\begin{aligned} \llbracket \pi_A \cup \pi_B \rrbracket(V) &:= \\ &\{(1.i).v \mid A.i.v \in V\} \cup \{(2.i).v \mid B.i.v \in V\} \end{aligned}$$

$$\begin{aligned} \llbracket \langle A_1 : f_1, \dots, A_k : f_k \rangle \rrbracket(V) &:= \\ &\{A_1.v_1, \dots, A_k.v_k \mid v_1 \in \llbracket f_1 \rrbracket(V) \wedge \dots \wedge v_k \in \llbracket f_k \rrbracket(V)\} \end{aligned}$$

$$\begin{aligned} \llbracket \text{map}(f) \rrbracket(V) &:= \\ &\{i.w \mid \exists u : i.u \in V \wedge w \in \llbracket f \rrbracket(\{v \mid i.v \in V\})\} \end{aligned}$$

$$\begin{aligned} \llbracket \text{pairwith}_{A_j} \rrbracket(V) &:= \{i.A_j.v \mid A_j.i.v \in V\} \cup \\ &\{i.A_k.w \mid A_j.i.v, A_k.w \in V \wedge j \neq k\} \end{aligned}$$

Here,  $V$  always denotes a set of paths,  $u, v, w, v_1, \dots, v_k$  denote paths, and  $i, j$  denotes indexes of set members. The symbol  $\langle \rangle$  in the definition for the equality predicate is to be understood as a constant and a path of length one. Observe how the flatten operation merges two set member indexes  $i$  and  $j$  into one exploiting our *binary* function symbol for encoding paths.

For an example, consider the query

$$\{0, 1\} = \langle A : 0 \circ \text{sng}, B : 1 \circ \text{sng} \rangle \circ \cup$$

which evaluates to a deterministic tree that can be uniquely specified by its set of root-to-leaf paths  $\{(1.s).0, (2.s).1\}$ .

The reduction of monad algebra queries to nonrecursive logic programming is now technical but not difficult.

Our predicates are binary and of the form  $p(X, v)$ , where  $X$  is a path prefix identifying a node  $w$  of the deterministic tree representation of our complex value, and  $v$  denotes one of the paths to leaves emanating from  $w$ , which taken together fully specify the complex value below node  $w$ .

- We translate an expression  $\text{map}(f)$  on path  $X$  represented by predicate  $\llbracket Q \rrbracket$  into

$$\begin{aligned} \llbracket Q; \text{start\_map} \rrbracket(X.i, v) &\leftarrow \llbracket Q \rrbracket(X, i.v). \\ \llbracket Q; \text{map}(f) \rrbracket(X, i.v) &\leftarrow \llbracket Q; \text{start\_map}; f \rrbracket(X.i, v). \end{aligned}$$

plus the translation of  $f$  mapping from predicate  $\llbracket Q; \text{start\_map} \rrbracket$  to  $\llbracket Q; \text{start\_map}; f \rrbracket$ .

That is, on a value identified by path prefix  $X$ , we move down to the set member children of  $X$ , the  $X.i$ . Then we apply  $f$  on the values  $X.i$ , and finally, we return to  $X$ .

- We translate an expression  $\langle A_1 : f_1, \dots, A_k : f_k \rangle$  on path  $X$  represented by predicate  $\llbracket Q \rrbracket$  into

$$\begin{aligned} \llbracket Q; \langle A_1 : f_1, \dots, A_k : f_k \rangle \rrbracket(X, A_1.v) &\leftarrow \llbracket Q; f_1 \rrbracket(X, v) \\ &\vdots \\ \llbracket Q; \langle A_1 : f_1, \dots, A_k : f_k \rangle \rrbracket(X, A_k.v) &\leftarrow \llbracket Q; f_k \rrbracket(X, v) \end{aligned}$$

plus, for each  $1 \leq i \leq k$ , the translation of  $f_i$  mapping from predicate  $\llbracket Q \rrbracket$  to predicate  $\llbracket Q; f_i \rrbracket$ .

- Compositions  $f \circ g$ , are read as  $f;g$  and  $f$  and  $g$  are translated separately. The result predicate of  $f$  is used as the input predicate of  $g$ .
- The remaining operations are translated as follows.

$$\begin{aligned} \llbracket Q; c \rrbracket(X, c) &\leftarrow \llbracket Q \rrbracket(X, v). \\ \llbracket Q; \text{pairwith}_B \rrbracket(X, i.B.v) &\leftarrow \llbracket Q \rrbracket(X, B.i.v). \\ \llbracket Q; \text{pairwith}_B \rrbracket(X, i.A.v) &\leftarrow \llbracket Q \rrbracket(X, A.v), \\ &\quad \llbracket Q \rrbracket(X, B.i.v). \\ \llbracket Q; \text{flatten} \rrbracket(X, (i.j).v) &\leftarrow \llbracket Q \rrbracket(X, i.j.v). \\ \llbracket Q; (A =_{\text{atomic}} B) \rrbracket(X, s.\langle \rangle) &\leftarrow \llbracket Q \rrbracket(X, A.v), \\ &\quad \llbracket Q \rrbracket(X, B.v). \\ \llbracket Q; \pi_A \cup \pi_B \rrbracket(X, (1.i).v) &\leftarrow \llbracket Q \rrbracket(X, A.i.v) \\ \llbracket Q; \pi_A \cup \pi_B \rrbracket(X, (2.i).v) &\leftarrow \llbracket Q \rrbracket(X, B.i.v) \\ \llbracket Q; \pi_{A_i} \rrbracket(X, v) &\leftarrow \llbracket Q \rrbracket(X, A_i.v) \\ \llbracket Q; \text{sng} \rrbracket(X, s.v) &\leftarrow \llbracket Q \rrbracket(X, v). \end{aligned}$$

By Proposition 3.6, we may assume that our query ignores the input data; so we assume a predicate  $\llbracket \epsilon \rrbracket$  and a fact  $\llbracket \epsilon \rrbracket(\epsilon, \text{dummy}) \leftarrow \cdot$  as part of our logic program.

It is not hard to verify that this translation of a query  $Q$  in  $\mathcal{M}_{\cup[=atomic]}$  into a nonrecursive logic program can be effected in LOGSPACE and that indeed the goal  $\llbracket Q \rrbracket(\epsilon, i.\langle \rangle)$  is true iff  $Q$  evaluates to true.  $\square$

The reduction to nonrecursive logic programming of the proof of Theorem 3.7 can be rather easily extended to a reduction from  $\mathcal{M}_{\cup[=atomic, not]}$  to nonrecursive normal logic programming (that is, with negation). All we need to do is encode the operation “not” as

$$\begin{aligned} \llbracket Q; \text{not} \rrbracket(X, s.\langle \rangle) &\leftarrow \text{set}[\llbracket Q \rrbracket(X), \\ &\quad \text{not nonempty}[\llbracket Q \rrbracket(X)]. \\ \text{nonempty}[\llbracket Q \rrbracket(X)] &\leftarrow \llbracket Q \rrbracket(X, v). \end{aligned}$$

where the “set $\llbracket Q \rrbracket$ ” predicates are defined alongside the  $\llbracket Q \rrbracket$  predicates such that  $\text{set}[\llbracket Q \rrbracket(X)]$  is true iff  $X$  is the path prefix of a set, empty or not. This reduction is not in LOGLIN because of the size of the predicates generated. Even if we replace the predicate names by shorter ones of the form  $p_i$  (where  $i$  is an integer), they are of size  $\log n$  each (where  $n$  is the size of the input query in monad algebra) and the overall size of the logic program is  $O(n \cdot \log n)$ . (There are linearly many rules.) But since we can compose this preparation with an ATM run and nonrecursive range-restricted normal logic programming is known to be in  $\text{TA}[2^{O(n)}, O(n)]$  [39], this shows that

**COROLLARY 3.8.**  $\mathcal{M}_{\cup[=atomic, not]}$  is in  $\text{TA}[2^{O(n \cdot \log n)}, O(n \cdot \log n)]$  w.r.t. query complexity.

We can improve this to

**THEOREM 3.9.**  $\mathcal{M}_{\cup[=atomic, not]}$  is in  $\text{TA}[2^{O(n)}, O(n)]$  w.r.t. query complexity.

**Proof Sketch.** The proof is direct, using alternating Turing machines, but again incorporates the deterministic tree technique and the idea of evaluating a logic program, now with negation. We will sketch a fixed alternating Turing machine  $M$  that recognizes the  $\mathcal{M}_{\cup[=]}$  queries that evaluate to true. Consider the proof of NEXPTIME-membership of nonrecursive logic programming without negation of [9]. It is by an argument that SLD resolution for such a program is in NEXPTIME because we can start from the goal and then always guess and verify unifiers until we have a proof that the goal is true. Unifiers are of singly exponential size; this is particularly easy to see for the special programs produced in the proof of Theorem 3.7, because there all predicates are over paths, and each path is of size  $O(n \cdot 2^{O(n)}) = 2^{O(n)}$ , where  $n$  is the size of the input query. (There are  $O(n)$  steps in the paths and each step is a value of size  $2^{O(n)}$  - no greater tuples can be computed by a  $\mathcal{M}_{\cup[=]}$  query of size  $n$ , and thus by our logic programs.)

Our ATM  $M$  basically follows such a resolution strategy to prove that the query evaluates to true. We first compute the logic program of Theorem 3.7 (and its extension to support negation described above) and write it to our worktape. Then we start proving the goal  $\llbracket Q \rrbracket(\epsilon, i.\langle \rangle)$ . Inductively, to prove a goal, for a given unifier, we guess a rule, adapt our unifier to the body atoms of the rule (both using existential configurations of the ATM), and then branch out using universal ATM computation to check the body atoms of the rule in parallel. Whenever we encounter a negated atom in a rule body, we employ universal computation to verify that this atom cannot become true. Constructed appropriately,  $M$  of course accepts if and only if the goal  $\llbracket Q \rrbracket(\epsilon, i.\langle \rangle)$  is true, and thus iff our  $\mathcal{M}_{\cup[=]}$  query evaluates to true.

Let us study the depth of the computation trees of  $M$ . The depth of the proof tree of the logic program is only linear in the size of the query. The paths in the computation tree of  $M$  are of length  $2^{O(n)}$  because all we need to do is choose rules and unify very special terms (our deterministic tree paths) of size  $2^{O(n)}$ . (One can verify by inspection of the construction of the logic program that this is feasible in linear time in the size of the paths.) The number of alternations used is bounded by the number of predicates in the program (There are  $O(n)$  many because there are linearly many rules.) plus the number of negation symbols in the program, which is again  $O(n)$ . Thus,  $M$  is an ATM that runs in time  $2^{O(n)}$  with  $O(n)$  alternations, and our result is shown.  $\square$

**REMARK 3.10.** The previous proof amounts in no way to a claim that we can close any gap we want - here, the gap between  $\text{TA}[2^{O(n)}, O(n)]$  and  $\text{TA}[2^{O(n \cdot \log n)}, O(n \cdot \log n)]$  - by just claiming that there is an appropriate Turing machine that performs our reduction *and* then solves the problem in the desired complexity class. But our proof shows that the predicate names introduced by our reduction from monad algebra to logic programming occupy space while not contributing to the power of the logic programs. This suggests that nonrecursive logic programming “wastes” some succinctness. This is also supported by the fact that, because of the space blow-up caused by the predicates, there is

a gap between the currently best known upper bound on the complexity of normal logic programming of TA $[2^{O(n)}, O(n)]$  [39] and the best lower bound of TA $[2^{O(n/\log n)}, O(n/\log n)]$  [40, 8].  $\square$

**THEOREM 3.11.**  $\mathcal{M}_{\cup[=]}$  is in EXPSPACE w.r.t. combined complexity.

**Proof Sketch.** We lack the space to prove this, but a brief argument can be given. Since all values computable in  $\mathcal{M}_{\cup[=]}$  are of at most doubly exponential size (see Proposition 3.4), we can represent an index of a set member (or even a path in the deterministic tree representation of complex values of the proof of Theorem 3.7) in a “register” of singly exponential size. We only use polynomially many (in the size of the query) such registers to evaluate the query using a strategy of recomputation of values on demand.

For example, an operation  $\text{map}(f)$  on a subvalue identified by path prefix  $\pi$  can be executed by computing each of the path prefixes  $\pi.i$ , where  $i$  identifies an element of the set  $\pi$ , (we can do this anytime we want because we know the query) and applying  $f$  to each of the  $\pi.i$ .

Deep equality of values identified by path prefixes  $\pi_1$  and  $\pi_2$  can be checked by verifying for each value identified by path prefix  $\pi_1.i$  whether there is an *equal* value identified by some path prefix  $\pi_2.j$ , and vice-versa. Equality here in general is again deep, so we must employ this procedure recursively, but only up to the at most linear depth of the values; thus we only need linearly many exp-sized registers for checking deep equality.  $\square$

## 4. LOWER BOUNDS

In this section we establish lower bounds matching the upper bounds of Theorems 3.7 and 3.9.

**THEOREM 4.1.**  $\mathcal{M}_{\cup[=atomic]}$  is NEXPTIME-hard w.r.t. query complexity.

**Proof Sketch.** The proof is by a LOGSPACE-reduction from NEXPTIME Turing machine acceptance. Let  $M = (Q^M, q_0^M, \delta^M, F^M)$  be a nondeterministic Turing machine (NTM) that runs in time  $2^{n^{O(1)}}$  on inputs of size  $n$ . We simulate the computation of  $M$  in  $\mathcal{M}_{\cup[=atomic]}$ . Each run of  $M$  is a sequence of configurations of length  $2^{K(n)}$ , for a suitable  $k$  and  $K(n) = n^k$ . (We may assume w.l.o.g. that terminating computation paths of  $M$  remain in a final state until time  $2^{K(n)}$  by appropriate design of  $M$ .) Each configuration of  $M$  consists of a read/write tape, a current state, and a position marker on the tape. Of course every  $2^{K(n)}$  time NTM computation uses tape space bounded by  $2^{K(n)}$ .

There are two main difficulties that we face in this reduction: We have to (i) deal with Turing machine tapes and configurations of exponential size and have to (ii) model the accepting computations of  $M$  of exponential length succinctly – the  $\mathcal{M}_{\cup[=atomic]}$  query that must achieve this has to be computable in LOGSPACE and thus must be of polynomial size.

### Modeling configurations.

- Each tape of a configuration is modeled as a tuple of arity  $2^{K(n)}$  (or more precisely, nested pairs of nesting depth  $K(n)$ ) of tape symbols.

Let  $\Sigma = \{s_1, \dots, s_c\}$  be the (fixed) tape alphabet of  $M$ . Rather than representing the current position of the read/write head on the tape separately from the tape, we will assume a *valid tape* over extended tape alphabet  $\Sigma' = \Sigma \cup \{\triangleright s \triangleleft \mid s \in \Sigma\}$  to contain a single symbol  $\triangleright s \triangleleft$  (with  $s \in \Sigma$ ) on the tape that indicates that this tape position stores symbol  $s$  and is the current position of the read/write head.

We can compute the set of all  $(2 \cdot c)^{2^{K(n)}}$  such tapes in  $\mathcal{M}_{\cup}$  as

$$\text{Tapes} := \phi_{\Sigma'} \circ \underbrace{(\text{id} \times \text{id}) \circ \dots \circ (\text{id} \times \text{id})}_{K(n) \text{ times}}$$

where  $\phi_{\Sigma'}$  is an appropriate  $\mathcal{M}_{\cup}$  expression that computes  $\Sigma'$ .<sup>8</sup>

As a result of this construction, some elements of set *Tapes* do not correspond to valid Turing tapes because they contain either zero or more than two markers indicating the current position of the read/write head on the tape. We will deal with this later.

- A superset of all possible configurations is

$$\text{Configs} := (\text{Tapes} \times Q^M) \circ \text{map}(\langle t : \pi_1, q : \pi_2 \rangle).$$

- The start configuration, consisting of the input tape, the start state, and the position marker at position 0 of the tape is obtained as follows.

We compute the start tape as the input  $x$ , with  $|x| = n$ , padded with  $(2^{K(n)} - n)$   $\#$ -symbols (denoting unused tape space) and with the first position marked, but in our nested pairs representation.

Let query  $\phi_x$  define the nested pair of depth  $\lceil \log_2 n \rceil$  representing  $x$  padded by  $(2^{\lceil \log_2 n \rceil} - n)$   $\#$ -symbols, and with the first position marked. (This is easy to compute in LOGSPACE.) For example, for input  $x = 01101$ , the value computed<sup>9</sup> is

$$\langle \langle \triangleright 0 \triangleleft, 1 \rangle, \langle 1, 0 \rangle \rangle, \langle \langle 1, \# \rangle, \langle \#, \# \rangle \rangle.$$

The start tape is

$$\phi_{\text{start}} := \langle 1 : \phi_x, 2 : \phi_{\text{empty}} \rangle \circ \underbrace{\phi_{\text{pad}} \circ \dots \circ \phi_{\text{pad}}}_{K(n) - \lceil \log_2 n \rceil - 1 \text{ times}}$$

with


$$\phi_{\text{pad}} = \langle 1 : \text{id}, 2 : \langle 1 : \pi_2, 2 : \pi_2 \rangle \rangle,$$

$$\phi_{\text{empty}} := \# \circ \underbrace{\langle \text{id}, \text{id} \rangle \circ \langle \text{id}, \text{id} \rangle \circ \dots \circ \langle \text{id}, \text{id} \rangle}_{\lceil \log_2 n \rceil \text{ times}}.$$

This takes the value computed by  $\phi_x$  – which contains the input and some padding up to  $2^{\lceil \log_2 n \rceil}$  symbols, pairs it with a sequence of  $\#$ -symbols of the same length (computed by  $\phi_{\text{empty}}$ ), and then iteratively doubles the length of the tape by appending two copies of the second half of the already computed tape (because

<sup>8</sup>I.e.,  $\phi_{\Sigma'} := s_1 \circ \text{sng} \cup \dots \cup s_c \circ \text{sng} \cup \triangleright s_1 \triangleleft \circ \text{sng} \cup \dots \cup \triangleright s_c \triangleleft \circ \text{sng}$ .

<sup>9</sup>It may be advisable to have a special symbol indicating the left end of the tape on its leftmost position to help the machine avoid running out of bounds. We assume such a symbol part of the input, rather than of our construction.



**Figure 1: Zooming into the tapes to find a valid tape change resulting from a computation step of  $M$ .**

the second half consists exclusively of #-symbols). By this trick, there is a *fixed* expression  $\phi_{pad}$  that we can compose our query with to double the length of the value produced.

The start configuration is

$$C_{start} := \langle t : \phi_{start}, q : q_0^M \rangle.$$

Observe that  $C_{start}$  is a valid configuration with precisely one tape head position marker.

- The accepting configurations are those configurations in which the state is an element of the set  $F^M = \{f_1, \dots, f_{|F^M|}\}$  of accepting states of  $M$ :

$$\begin{aligned} \text{AcceptingConfigs} &:= \\ \text{Configs} \circ (\sigma_{q=atomic} f_1 \cup \dots \cup \sigma_{q=atomic} f_{|F^M|}). \end{aligned}$$

- In the following, we will test equality of nested pairs (tape segments) and configurations of exponential size. We can define an equality test  $=_{mon}$  of linear size on tapes and tape segments using only  $=_{atomic}$  inductively as follows. On values of type  $\text{Dom}$ ,  $=_{mon}$  is  $=_{atomic}$ . Otherwise, on pairs  $\langle 1 : \tau_1, 2 : \tau_2 \rangle$ ,

$$\begin{aligned} (A =_{mon} B) &:= ((\pi_A \circ \phi) \times (\pi_B \circ \phi)) \circ \\ &\sigma_{1.T=atomic} 2.T \circ \sigma_{1.V=mon} 2.V \circ (\text{id} \times \text{id}) \circ \\ &\sigma_{1.1.T=atomic} \text{"1"} \circ \sigma_{2.1.T=atomic} \text{"2"} \circ \text{map}(\langle \rangle) \end{aligned}$$

where  $\phi := (\langle T : 1, V : \pi_1 \rangle \circ \text{sng} \cup \langle T : 2, V : \pi_2 \rangle \circ \text{sng})$ . For configurations  $C, C'$ ,

$$(C =_{mon} C') \Leftrightarrow (C.t =_{mon} C'.t \wedge C.q =_{atomic} C'.q).$$

- Next we define an  $\mathcal{M}_U[=_{atomic}]$  expression  $\phi_{succ}$  that computes the pairs of configurations  $\langle C, C' \rangle$  such that  $C'$  is a possible successor of  $C$ , i.e., computable using the transition relation  $\delta^M$  of  $M$  in one step.

Here the exponential size of the configurations is a problem;  $\phi_{succ}$  has to be chosen carefully in order not to be of exponential size. We achieve this as follows. We start with the Cartesian product of *Configs* – all pairs of configurations, even many that are invalid because they have zero or more than two head markers. For each pair, we make working copies  $w, w'$  of the

tapes. We achieve this by  $\mathcal{M}_U$  expression

$$\begin{aligned} \phi_{prepare-succ} &:= \text{Configs} \circ (\text{id} \times \text{id}) \circ \\ &\text{map}(\langle s : \text{id}, w : \pi_{C.t}, w' : \pi_{C'.t} \rangle). \end{aligned}$$

For  $w'$  to be a possible successor of  $w$ , the two tapes may differ at at most two consecutive tape positions (if the tape head moved, otherwise they may only differ at at most one position), and these positions must contain the read/write head position marker. We synchronously “zoom into” the working copies to find these two positions using the following three rules:

1. If  $w.2 = w'.2$  (i.e., the second halves of the tapes are equal), replace  $w$  by  $w.1$  and  $w'$  by  $w'.1$ .
2. If  $w.1 = w'.1$  (i.e., the first halves of the tapes are equal), replace  $w$  by  $w.2$  and  $w'$  by  $w'.2$ .
3. If  $w.1.1 = w'.1.1$  and  $w.2.2 = w'.2.2$  (i.e., the first and last quarters of the tapes are equal) replace  $w$  by newly constructed pair  $\langle 1 : w.1.2, 2 : w.2.1 \rangle$  and  $w'$  by  $\langle 1 : w'.1.2, 2 : w'.2.1 \rangle$  (that is, by the second and third quarters).

All three cases may apply at the same time because the tapes of two valid configurations  $C, C'$ , where  $C'$  is a successor of  $C$ , can be equal. An example of iterative zooming is shown in Figure 1. There, we look at a nested pair term of depth four (covering a tape of length 16) and the tape change occurs at positions 6 and 7. We first zoom into the left (using Rule 1) and from there into the right half (using Rule 2). Now both halves differ, but the first and fourth quarter do not, so we can use Rule 3 to zoom down to the differing tape positions 6 and 7. In general, we obtain a tape sequence of length two by zooming into a tape (which is of length  $2^{K(n)}$ )  $K(n) - 1$  times.

In our encoding in  $\mathcal{M}_U[=_{atomic}]$ , we will compute the union of all triples  $(s, w, w')$  such that  $s$  is a pair of configurations with tapes  $t = uvw$  and  $t = uw'v$  and  $w$  and  $w'$  are of length 2 (i.e.,  $w$  and  $w'$  are – if any – the only corresponding sequences in  $t, t'$  that differ). Now we have to make sure that  $w$  and  $w'$  contain the position marker.

This can be expressed in  $\mathcal{M}_U[=_{atomic}]$  as follows:

$$\begin{aligned} \phi_{witness-succ} &:= \phi_{prepare-succ} \circ \\ &\underbrace{\phi_{zoom-in} \circ \dots \circ \phi_{zoom-in}}_{K(n)-1 \text{ times}} \circ \phi_{marker} \end{aligned}$$

where<sup>10</sup>

$$\begin{aligned} \phi_{zoom-in} &:= (\sigma_{12\triangleright 34\triangleleft} \circ \pi_{12\triangleright 34\triangleleft} \cup \\ &\sigma_{\triangleright 12\triangleleft 34} \circ \pi_{\triangleright 12\triangleleft 34} \cup \sigma_{1\triangleright 23\triangleleft 4} \circ \pi_{1\triangleright 23\triangleleft 4}) \\ \sigma_{12\triangleright 34\triangleleft} &:= \sigma_{w.1=mon} w'.1 \\ \pi_{12\triangleright 34\triangleleft} &:= \text{map}(\langle s : \pi_s, w : \pi_{w.2}, w' : \pi_{w'.2} \rangle) \\ \sigma_{\triangleright 12\triangleleft 34} &:= \sigma_{w.2=mon} w'.2 \\ \pi_{\triangleright 12\triangleleft 34} &:= \text{map}(\langle s : \pi_s, w : \pi_{w.1}, w' : \pi_{w'.1} \rangle) \\ \sigma_{1\triangleright 23\triangleleft 4} &:= \sigma_{w.1.1=mon} w'.1.1 \circ \sigma_{w.2.2=mon} w'.2.2 \end{aligned}$$

<sup>10</sup>Here and later,  $\pi_{A_1, \dots, A_m} := \pi_{A_1} \circ \dots \circ \pi_{A_m}$ .

$$\begin{aligned} \pi_{1 \triangleright 23 \triangleleft 4} &:= \text{map}(\langle s : \pi_s, \\ &\quad w : \pi_w \circ \langle 1 : \pi_{1.2}, 2 : \pi_{2.1} \rangle, \\ &\quad w' : \pi_{w'} \circ \langle 1 : \pi_{1.2}, 2 : \pi_{2.1} \rangle \rangle) \end{aligned}$$

and  $\phi_{\text{marker}}$  selects those tuples for which  $w, w'$  are two tapes of length two that contain the read/write head marker:

$$\begin{aligned} \phi_{\text{marker}} &:= (\sigma_{w.1=\text{atomic} \triangleright s_1 \triangleleft} \cup \dots \cup \sigma_{w.1=\text{atomic} \triangleright s_c \triangleleft} \cup \\ &\quad \sigma_{w.2=\text{atomic} \triangleright s_1 \triangleleft} \cup \dots \cup \sigma_{w.2=\text{atomic} \triangleright s_c \triangleleft}). \end{aligned}$$

Now, for each  $\langle s : X, w : Y, w' : Z \rangle \in \llbracket \phi_{\text{witness-succ}} \rrbracket$ , either  $C = C'$  or  $Y, Z$  are precisely the at most two adjacent positions of the tapes of  $C$  and  $C'$  that can differ if  $C'$  is to be a successor of  $C$ . We can easily encode the valid successors with respect to transition relation  $\delta_M$  by a union of expressions that amount to selecting every pair of tapes that matches one of the transition rules encoded in  $\delta^M$ :

$$\phi_{\text{succ}} := \phi_{\text{witness-succ}} \circ (\sigma_{\gamma_1} \cup \dots \cup \sigma_{\gamma_m}) \circ \text{map}(\pi_s)$$

For instance, if  $\langle q', b, +1 \rangle \in \delta(q, a)$ , one  $\sigma_{\gamma_i}$  is to select the triples  $\langle s : \langle S : \langle t : u \triangleright a \triangleleft sv, q : q \rangle, S' : \langle t : ub \triangleright s \triangleleft v, q : q' \rangle \rangle, w : \triangleright a \triangleleft s, w' : b \triangleright s \triangleleft \rangle \in \llbracket \phi_{\text{witness-succ}} \rrbracket$ . (Details are omitted for lack of space, but it is important to note that the values that we are dealing with are atomic, so we only need equality on atomic values.)

As for *Configs*,  $\phi_{\text{succ}}$  contains pairs  $\langle C, C' \rangle$  of invalid configurations. However, whenever  $C$  is a valid configuration,  $C'$  is indeed a possible successor configuration on  $M$ . It follows by induction that, starting from valid configuration  $C_{\text{start}}$ , we will only reach valid configurations via the successor relation  $\phi_{\text{succ}}$ .

**Modeling computations.** Now we are ready to model accepting computations of  $M$ . Here the problem is the possibly exponential running time. We use a simple recursive divide-and-conquer approach in the spirit of the usual proof of Savitch's theorem (cf. e.g. [33]). Let  $\psi_i$  the pairs of configurations  $\langle C, C' \rangle$  such that  $C$  is reachable from  $C'$  in  $2^i$  steps. We define  $\psi_i$  as

$$\begin{aligned} \psi_0 &:= \phi_{\text{succ}} \\ \psi_{i+1} &:= \psi_i \circ (\text{id} \times \text{id}) \circ \sigma_{1.C'=2.C} \circ \\ &\quad \text{map}(\langle C : \pi_{1.C}, C' : \pi_{2.C'} \rangle) \end{aligned}$$

Note that the definition of  $\psi_{i+1}$  uses  $\psi_i$  only once, thus the formula remains computable in LOGSPACE.

There is an accepting computation path of length  $2^{K(n)}$  iff there is a pair  $\langle C, C' \rangle$  in  $\psi_{K(n)}$  such that  $C = C_{\text{start}}$  and the state of  $C'$  is in  $F^M$ . We can phrase this as

$$\begin{aligned} \phi_{\text{accept}} &:= \left( \langle 1 : C_{\text{start}}, 2 : \psi_{K(n)} \rangle \circ \text{pairwith2} \circ \right. \\ &\quad \left. \sigma_{1=\text{mon} 2.C} \circ \text{map}(\pi_{2.C'}) \right) \times \text{AcceptingConfigs} \Big) \circ \\ &\quad \text{map}([1 =_{\text{mon}} 2]) \circ \text{flatten}. \end{aligned}$$

(Again we only employ equality on configurations.)

It is not difficult to see that the  $\mathcal{M}_{\cup} [=_{\text{atomic}}]$  query  $\phi_{\text{accept}}$  constructed is of polynomial size and can be computed in LOGSPACE. The entire problem is formulated as the query (e.g., the input  $x$  is constructed from constants and pairs)

and  $\phi_{\text{accept}}$  will not make use of an input value. Thus we have shown that  $\mathcal{M}_{\cup} [=_{\text{atomic}}]$  is NEXPTIME-hard with respect to query complexity (i.e., for a fixed database).  $\square$

It is not hard to verify that

**LEMMA 4.2.** *For the construction of the proof of Theorem 4.1, (a)  $|\phi_{\text{accept}}| = O(K(n))^2$ . (b) If  $=_{\text{mon}}$  is available as a built-in,  $\phi_{\text{accept}}$  can be defined such that  $|\phi_{\text{accept}}| = O(K(n))$ .*

**COROLLARY 4.3.**  $\mathcal{M}_{\cup} [=_{\text{mon}}]$  is NETIME-hard under LOGLIN-reductions (query complexity).

**THEOREM 4.4.**  $\mathcal{M}_{\cup} [=_{\text{mon}}, \text{not}]$  is  $\text{TA}[2^{O(n)}, O(n)]$ -hard under LOGLIN-reductions (query complexity).

**Proof Sketch.** The proof is by a LOGLIN-reduction from  $\text{TA}[2^{O(n)}, O(n)]$  Turing machine acceptance. Let

$$M = (Q_{\exists}^M, Q_{\forall}^M, q_0^M, \delta^M, F^M)$$

be an alternating Turing machine (ATM) that runs in time  $2^{O(n)}$  with  $O(n)$  alternations on inputs of size  $n$ .

We simulate the computation of  $M$  in  $\mathcal{M}_{\cup} [=_{\text{mon}}, \text{not}]$ . Each run of  $M$  is a tree of configurations of depth  $2^{k \cdot n}$ , for a suitable constant  $k$ . We may assume w.l.o.g. that terminating computation paths of  $M$  are no longer than  $2^{k \cdot n}$ , i.e., the depth of the computation tree of  $M$  is bounded by  $2^{k \cdot n}$ .

By Lemma 4.2, if we use a built-in operation  $=_{\text{mon}}$ , the sizes of all formulas of the proof of Theorem 4.1 are linear in the size of  $K(n)$ . Now, we fix  $K(n) = k \cdot n$ , for some constant  $k$ .

We use the formulae  $C_{\text{start}}$ , *Configs*, *AcceptingConfigs*, and  $\phi_{\text{succ}}$  constructed as described. We define a modified version of  $\psi_{k \cdot n}$  which computes the set of computation paths of length up to  $2^{k \cdot n}$  (this can be realized by adding ‘‘stay transitions’’  $\langle C, C \rangle$ , for  $C \in \text{Configs}$  to  $\phi_{\text{succ}}$ ) and where the states of the intermediate configurations are all from  $Q_{\exists}$  if the state of the first configuration is from  $Q_{\exists}$  and are all from  $Q_{\forall}$  otherwise. We can define this as

$$\begin{aligned} \psi_{i+1} &:= \psi_i \circ (\text{id} \times \text{id}) \circ \sigma_{1.C'=2.C} \circ \\ &\quad \sigma_{1.C.q \in Q_{\exists}^M \Leftrightarrow 2.C.q \in Q_{\exists}^M} \circ \text{map}(\langle C : \pi_{1.C}, C' : \pi_{2.C'} \rangle). \end{aligned}$$

Now that we only need to consider tapes of size  $2^{k \cdot n}$ , the monad algebra expressions only occupy space  $O(n)$ , thus so far we have a LOGLIN reduction.

Let the sets of configurations  $A_i$  be inductively defined as

$$\begin{aligned} A_1 &:= \{C \mid \exists C' (C, C') \in \psi_{k \cdot n} \wedge \\ &\quad C' \in \text{AcceptingConfigs} \wedge C.q \in Q_{\exists}^M\} \\ A_{i+1} &:= \{C \mid \exists C' (C, C') \in \psi_{k \cdot n} \wedge \\ &\quad C' \in (\text{Configs} - A_i) \wedge C.q \in Q_{\exists}^M \Leftrightarrow C'.q \notin Q_{\exists}^M\} \end{aligned}$$

Clearly,  $C \in A_i$  for odd  $i$  means that  $C.q \in Q_{\exists}^M$  and that  $C$  is eventually accepting;  $C \in A_i$  for even  $i$  means that  $C.q \in Q_{\forall}^M$  and that  $C$  is not eventually accepting (both via  $i$  alternations and in  $2^i$  steps).

W.l.o.g., we may assume that  $F^M \subseteq Q_{\exists}^M$ . By this assumption  $F^M \subseteq A_1$ , and thus the final transitions leading to accepting states may be universal, rather than just existential. We will now be somewhat sloppy and assume that the number of alternations  $K(n) = O(n)$  we ask for is always odd. This is to keep the argument short, but a slight

modification of the construction allows to eliminate the assumption.

Then,  $M$  accepts its input precisely if  $C_{start}$  is eventually accepting with  $K(n)$  alternations, that is, iff  $C_{start} \in A_{K(n)}$ .

It is not difficult to construct  $A_{K(n)}$  in monad algebra. We only remark that difference  $A - B$  on sets of nested tuples can be defined using  $=_{mon}$  and “not” as

$$\{a \in A \mid \exists b \in B \wedge a =_{mon} b\}$$

or, in monad algebra on pair  $\langle 1 : A, 2 : B \rangle$ ,

$$\text{pairwith}_1 \circ \text{flatMap}(\langle a : \pi_1, c : \langle a : \pi_1, B : \pi_2 \rangle \circ \text{pairwith}_B \circ \text{flatMap}(a =_{mon} B) \circ \text{not} \rangle \circ \text{pairwith}_c \circ \text{map}(\pi_a))$$

Formula  $\phi_{accept}$  is obviously of linear size, and thus the construction in LOGLIN. This concludes our proof.  $\square$

Considering again  $=_{atomic}$  as a built-in, our LOGSPACE-reduction of the proof of Theorem 4.1 for configurations and  $\phi_{succ}$  (with  $K(n) = n^k$ ) in combination with the construction for computations ( $A_i$  and  $\phi_{accept}$ ) of the proof of Theorem 4.4 yields

**COROLLARY 4.5.**  $\mathcal{M}_\cup[=_{atomic}, \text{not}]$  is  $TA[2^{n^{O(1)}}, n^{O(1)}]$ -hard under LOGSPACE-reductions (query complexity).

We can give a more precise lower bound.

**THEOREM 4.6.**  $\mathcal{M}_\cup[=_{atomic}, \text{not}]$  is  $TA[2^{O(n)}, O(n)]$ -hard under LOGLIN-reductions (query complexity).

**Proof Sketch.** To allow for a query  $\phi_{accept}$  of linear size overall, we have to rephrase both  $\phi_{witness-succ}$  and  $\psi_{K(n)}$  to use our formula defining  $=_{mon}$  via  $=_{atomic}$  only a constant number of times. We can do this now that we have negation and thus equality of a set of nested tuples available. We only sketch the idea here briefly, but it is the same for the two cases. Rather than testing equality linearly many times, we postpone the testing of equality on pairs of tuples until we have collected all the pairs in a set and we can test equality of them all at once. We demonstrate the idea for  $\psi_{K(n)}$ . Let

$$\begin{aligned} \psi'_0 &:= \phi_{succ} \circ \text{map}(\langle 1 : \text{id}, 2 : \emptyset \rangle) \\ \psi'_{i+1} &:= (\text{id} \times \text{id}) \circ \sigma_{1.C.q \in Q_{\exists}^M \Leftrightarrow 2.C.q \in Q_{\exists}^M} \circ \\ &\quad \text{map}(\langle 1 : \langle C : \pi_{1.1.C}, C' : \pi_{2.1.C'} \rangle, \\ &\quad 2 : \pi_{1.2} \cup \pi_{2.2} \cup \langle 1 : \pi_{1.1.C'}, 2 : \pi_{2.1.C'} \rangle \circ \text{sng} \rangle) \end{aligned}$$

Now we are interested in those pairs of configurations  $(c, c')$  s.t.  $\langle 1 : \langle C : c, C' : c' \rangle, 2 : S \rangle \in \psi'_{K(n)}$  and for all  $\langle 1 : t, 2 : t' \rangle \in S$ ,  $t =_{mon} t'$ . We can define this as

$$\psi_{K(n)} := \psi'_{K(n)} \circ \text{map}(\langle 1 : \pi_1, 2 : \pi_2 \circ \text{all-equal} \rangle \circ \text{pairwith}_2 \circ \text{map}(\pi_1)) \circ \text{flatten}$$

where  $\text{all-equal} := \text{map}((1 =_{mon} 2) \circ [\text{not}]) \circ \text{flatten} \circ \text{not}$ .

For  $\phi_{witness-succ}$ , we proceed analogously. We define  $\phi_{zoom-in}$  to be a mapping from sets of tuples  $\langle s : \langle C, C' \rangle, w : t, w' : t', mbe : S \rangle$  (where  $(s : \langle C, C' \rangle, w : t, w' : t')$  is as in the proof of Theorem 4.1 and  $S$  is a set of pairs yet to be checked to be equal –  $mbe$  is short for “must be equal”) to sets of tuples of the same type. We replace e.g.

$\sigma_{12 \triangleright 34 \triangleleft} \circ \pi_{12 \triangleright 34 \triangleleft}$  in  $\phi_{zoom-in}$  by

$$\begin{aligned} &\text{map}(\langle s : \pi_s, w : \pi_{w.2}, w' : \pi_{w'.2}, \\ &\quad mbe : \langle w : \pi_{w.1}, w' : \pi_{w'.1} \rangle \circ \text{sng} \cup \\ &\quad \pi_{mbe} \circ \text{map}(\langle w : \pi_{w.1}, w' : \pi_{w'.1} \rangle \circ \text{sng} \cup \\ &\quad \langle w : \pi_{w.2}, w' : \pi_{w'.2} \rangle \circ \text{sng} \rangle) \circ \text{flatten} \rangle) \end{aligned}$$

That is, in each such step we add the values that were checked to be equal using  $=_{mon}$  in  $\phi_{zoom-in}$  – here, for  $\sigma_{12 \triangleright 34 \triangleleft} \circ \pi_{12 \triangleright 34 \triangleleft}$ ,  $w.1$  and  $w'.1$  – and add them to  $mbe$ . Before we do that, we split the pairs  $(t, t')$  of  $mbe$  into their immediate constituents (as shown in the bottom two lines of the monad algebra expression above). This is necessary to assure that all members of  $mbe$  are of the same type. However, it has a nice side-effect. By this restructuring, after the last zoom-in step, the members of  $mbe$  are pairs of *atomic values*, and we actually do not need  $=_{mon}$  here at all and can use  $=_{atomic}$  instead.

Note that we could not have used this construction in the proof of Theorem 4.1 because now we need negation to check that for each pair  $(t, t')$  in  $mbe$ ,  $t =_{atomic} t'$ . (This can be done using the “all-equal” predicate defined above, with  $=_{atomic}$  replacing  $=_{mon}$ .)  $\square$

Since “not” is equivalent to  $(\text{id} = \emptyset)$ ,

**COROLLARY 4.7.**  $\mathcal{M}_\cup[=]$  is  $TA[2^{O(n)}, n]$ -hard under LOGLIN-reductions (query complexity).

The queries constructed in our lower bound proofs are from flat relations to flat relations (we may assume this since we actually use no input data value). Since relational algebra is in PSPACE w.r.t. combined complexity (cf. e.g. [3]) and presumably PSPACE  $\neq$  NEXPTIME, it seems unlikely that there is even a PSPACE reduction from  $\mathcal{M}_\cup[=]$  on flat relations to relational algebra in the spirit of the Conservativity Theorem of Paredaens and Van Gucht ([34], Theorem 2.5).

## 5. LISTS AND BAGS

In this section, we study the complexity of monad algebra on lists  $\mathcal{M}_\cup^{\square}$  and bags  $\mathcal{M}_\cup^{\square \square}$ . A formal definition of these languages is beyond the scope of this paper, but see [37, 6, 30] for full formal definitions. We will use the same syntax and operation names as for monad algebra on sets, but now, for instance,  $\cup$  on lists means to append two lists and “flatten” appends the list-typed members of a list in order of appearance. For bags, these operations ignore order but preserve duplicates. Of course, two lists are equal iff they are of the same length and for each  $i$ , the  $i$ -th members of the two lists are equal. Two bags are equal iff each member of either bag occurs the same number of times in both bags.

For bags, we will also consider the additional operations “monus” (a powerful version of difference which allows to express arithmetics in monad algebra on bags) and “unique”, an operation that eliminates duplicates from bags. In [30] it was shown that adding either of these two operations strictly increases the expressive power of the language (and adding both makes the language yet stronger).

First we again look at data complexity.

**PROPOSITION 5.1 (FOLKLORE).**  $\mathcal{M}_\cup^{\square \square}[=, \text{monus}]$  and  $\mathcal{M}_\cup^{\square}[=]$  are in  $TC_0$  w.r.t. data complexity.

There is no space to provide a proof for this, but “parsing” and accessing nested data is described in the proof of Proposition 3.2 and implementing the various operations of monad algebra is not difficult. (See also the similar proof that XQuery is in TC<sub>0</sub> – Theorem 8.3). It is folklore that the majority gates of TC<sub>0</sub> circuits are powerful enough to support the arithmetics required to implement bag operations such as bag difference.

For a result that suggests that this is a good bound,

PROPOSITION 5.2 ([18]). *There are  $\mathcal{M}_{\cup}^{\{\}}[=, \text{minus}]$  queries that are not in AC<sub>0</sub>.*

Regarding query/combined complexity, we can show that

PROPOSITION 5.3. *The languages  $\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}, \text{not}]$  and  $\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}, \text{not}]$  are TA[2<sup>O(n)</sup>, O(n)]-hard under LOGLIN-reductions (query complexity).*

PROPOSITION 5.4.  *$\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}]$  and  $\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}]$  are NEXPTIME-hard w.r.t. query complexity.*

The lower bound proofs for sets work without modifications on lists and bags – we actually do not compare collections except in the definition of  $A_i$  of the proof of Theorem 4.4, where we compute differences. But here, we define difference  $R - S$  as a *filter* (using “map”) that computes those elements of  $R$  for which no element of  $S$  with the same value exists. For lists, this will preserve order of the elements in  $R$  and for bags it will preserve their multiplicities. For the correctness of our reduction, this does not matter, as long as we interpret nonempty collections of type  $\{\langle \rangle\}$  resp.  $\{\langle \rangle\}$  (possibly with duplicates) as truth and empty collections as falsity.

For the upper bounds,

THEOREM 5.5.  *$\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}]$  and  $\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}]$  are in NEXPTIME w.r.t. query complexity.*

THEOREM 5.6.  *$\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}, \text{not}]$  and  $\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}, \text{not}]$  are in TA[2<sup>O(n)</sup>, O(n)] under LOGLIN-reductions (query complexity).*

Here, the proofs of Theorems 3.7 and 3.9 work without modifications for lists and bags. Actually, our encoding using deterministic trees treats collections as lists, and thus preserves both order and multiplicities of members. If only equality on atomic values is available, however, we cannot distinguish between sets, lists, and bags in queries. Finally,

THEOREM 5.7.  *$\mathcal{M}_{\cup}^{\{\}}[=, \text{minus}, \text{unique}]$  and  $\mathcal{M}_{\cup}^{\{\}}[=]$  are in EXPSPACE w.r.t. query complexity.*

**Proof Sketch.** The proof is the same as for Theorem 3.11, but now we also have to check deep list and bag equality in EXPSPACE. However, these checks can be effected by easy modifications of the equality test described in the proof of Theorem 3.11. (For instance, for bag equality, we must check for each value in either bag that the number of occurrences of the value in the two bags is the same. Since we can enumerate and re-enumerate all paths of the deterministic tree representation of a bag in some unchanging – say lexicographic – order, this is straightforward.)  $\square$

In the remainder of this paper, we will assume that  $\mathcal{M}_{\cup}^{\{\}}$  has one further operation, “true”, which evaluates to  $\{\langle \rangle\}$  (true) on a list if it is nonempty and to  $\{\}$  (false) otherwise. We will use “true” to eliminate duplicate entries from truth values (i.e., from  $\{\langle \rangle, \dots, \langle \rangle\}$ ). It is easy to verify that this does not increase the complexity of  $\mathcal{M}_{\cup}^{\{\}}[=]$ ,  $\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}]$ , or  $\mathcal{M}_{\cup}^{\{\}}[=_{\text{atomic}}, \text{not}]$ . For  $\mathcal{M}_{\cup}^{\{\}}[=]$ , in the proof of Theorem 5.7, we only need a small number of exp-size registers to recompute the collection (we can stop early if we find at least one member). In the proofs of the upper bounds of the other two fragments, “true” is a non-operation because for proving a goal, duplicates do not matter.

## 6. CORE XQUERY

We consider the fragment of XQuery with abstract syntax

$$\begin{aligned} \text{query} & ::= \langle a \rangle \mid \langle a \rangle \text{query} \langle /a \rangle \mid \text{query query} \\ & \mid \text{var} \mid \text{var}/\text{axis} :: \phi \\ & \mid \text{for } \text{var} \text{ in } \text{query} \text{ return } \text{query} \\ & \mid \text{if } \text{cond} \text{ then } \text{query} \\ & \mid (\text{let } \text{var} := \text{query}) \text{ query} \\ \text{cond} & ::= \text{var} = \text{var} \mid \text{query} \end{aligned}$$

where  $a$  denotes the XML tags,  $\text{axis}$  the XPath axes<sup>11</sup>,  $\text{var}$  a set of XQuery variables  $\$x, \$x_1, \$x_2, \dots, \$y, \$z, \dots$ , and  $\phi$  a *node test* (either a tag name or “\*”). We refer to this fragment as *Core XQuery*, or *XQ* for short.

For simplicity, we will work with pure node-labeled unranked ordered trees, and by atomic values, we will refer to leaves (or equivalently, their labels).

XQuery supports several forms of equality. We will not try to use the same syntax (=, eq, or deep\_equal) as in the current standards proposal – it is not clear whether the syntax has stabilized. Throughout this paper, equality is by value (that is, by value as a tree rather than by the *yield* of strings at leaf nodes of the tree). We will write  $=_{\text{deep}}$  and  $=_{\text{atomic}}$  for deep and atomic equality, respectively. We will use = for statements that apply to both forms of equality.

The semantics of XQ is given in Figure 2. As for  $\mathcal{M}_{\cup}^{\{\}}$ ,  $\cup$  here denotes list concatenation;  $<_{\text{doc}}$  is the depth-first left-to-right traversal order through the tree,  $\chi^t$  is the axis relation  $\chi$  on tree  $t$ ,  $\text{lab}_*^t$  is true on all nodes of  $t$ , and  $\text{lab}_a^t$ , for  $a$  a tag name, is true on those nodes of  $t$  labeled  $a$ . All XQ queries evaluate to lists of nodes; however, we assume that XQ variables always bind to single nodes rather than lists; to assure this, we require that for expressions “(let  $\$x_{k+1} := \alpha$ )  $\beta$ ”,  $\alpha$  is either of the form  $\langle a \rangle$  or  $\langle a \rangle \alpha_0 \langle /a \rangle$  (i.e., this gives a easy syntactic condition that  $\alpha$  always evaluates to a singleton list). This semantics is (observationally) consistent with [42] restricted to Core XQuery.

In our definition of the syntax of Core XQuery, we have been economical with operators introduced. Since conditions are true iff they evaluate to a nonempty collection,

$$\begin{aligned} \phi \text{ or } \psi & ::= \phi \psi \\ \phi \text{ and } \psi & ::= \text{if } \phi \text{ then } \psi \\ \text{some } \$x \text{ in } \alpha \text{ satisfies } \phi & ::= \text{for } \$x \text{ in } \alpha \text{ return } \phi \end{aligned}$$

<sup>11</sup>For simplicity, particularly of the following semantics, we will only consider the child and the descendant axis; but all complexity upper bounds will hold for all XPath axes. Our theorems will refer to “all axes”, but proofs will assume that this notion refers just to the two axes child and descendant.

$$\begin{aligned}
\llbracket \langle a / \rangle \rrbracket_k(\vec{e}) &:= \llbracket \langle a / \rangle \rrbracket \\
\llbracket \langle a \rangle \alpha \langle / a \rangle \rrbracket_k(\vec{e}) &:= \llbracket \langle a \rangle \llbracket \alpha \rrbracket(\vec{e}) \langle / a \rangle \rrbracket \\
\llbracket \alpha \beta \rrbracket_k(\vec{e}) &:= \llbracket \alpha \rrbracket(\vec{e}) \cup \llbracket \beta \rrbracket(\vec{e}) \\
\llbracket \text{for } \$x_{k+1} \text{ in } \alpha \\
&\text{return } \beta \rrbracket_k(\vec{e}) &:= \text{let } l = \llbracket \alpha \rrbracket_k(\vec{e}), n = |l|; \\
&\text{return } \bigcup_{1 \leq i \leq n} \llbracket \beta \rrbracket_{k+1}(\vec{e}, l_i) \\
\llbracket (\text{let } \$x_{k+1} := \alpha) \beta \rrbracket_k(\vec{e}) &:= \llbracket \beta \rrbracket_{k+1}(\vec{e}, \llbracket \alpha \rrbracket_k(\vec{e})) \\
\llbracket \$x_i \rrbracket_k(t_1, \dots, t_n) &:= t_i \\
\llbracket \$x_i / \chi :: \phi \rrbracket_k(\vec{e}) &:= \text{list of nodes } v \text{ of tree } t \text{ s.t.} \\
&\llbracket \$x_i \rrbracket_k(\vec{e}) = [t] \wedge \\
&\chi^t(\text{root}^t, v) \wedge \text{lab}_\phi^t(v) \\
&\text{in order } <_{doc} \\
\llbracket \text{if } \gamma \text{ then } \alpha \rrbracket_k(\vec{e}) &:= \text{if } \llbracket \gamma \rrbracket_k(\vec{e}) \text{ then } \llbracket \alpha \rrbracket_k(\vec{e}) \text{ else } \llbracket \rrbracket_k(\vec{e}) \\
\llbracket \$x_i = \$x_j \rrbracket_k(e_1, \dots, e_k) &:= \text{if } e_i = e_j \text{ then } \llbracket \langle \text{yes} / \rangle \rrbracket \text{ else } \llbracket \rrbracket_k(\vec{e})
\end{aligned}$$

**Figure 2: Semantics of Core XQuery.**

Using deep equality, we can define negation,

$$\text{not } \phi := ((\langle a \rangle \{ \text{if } \phi \text{ then } \langle b / \rangle \} \langle / a \rangle) =_{deep} \langle a / \rangle).$$

Conditions “every  $\$x$  in  $\alpha$  satisfies  $\phi$ ” can be defined using “not” and “some”. We can even assume that expressions of the form  $\text{var}/\text{path}$  are supported, where  $\text{path}$  is any expression in navigational XPath (aka. Core XPath [15, 16]).

It is clear that

**PROPOSITION 6.1.** *Let  $\mathbf{X}$  be a set of operations and axes. Then, each  $XQ[=_{deep}, \text{not}, \text{every}, \mathbf{X}]$  query can be translated in LOGLIN into an equivalent  $XQ[=_{deep}, \mathbf{X}]$  query and each  $XQ[\text{and}, \text{or}, \text{some}, \mathbf{X}]$  query can be translated in LOGLIN into an equivalent  $XQ[\mathbf{X}]$  query.*

Next, we provide mappings between Core XQuery (using only the child axis) and monad algebra on lists. These show the equivalence of these languages up to representation issues, but our main aim is to provide reductions for the study of the complexity of XQuery.

### Translation from Core XQuery to $\mathcal{M}_{\cup}^{\square}$

We recursively map the data tree  $T$  to a complex value  $C(T)$  as follows: Each tree node with label  $t$  and children  $v_1, \dots, v_n$  is mapped to a tuple

$$\langle \text{label} : t, \text{children} : \{C(v_1), \dots, C(v_n)\} \rangle.$$

Modulo representation issues captured in the tree translation function  $C$ , there is an equivalent monad algebra query for each  $XQ[=, \text{child}]$  query, for  $=$  either  $=_{deep}$  or  $=_{atomic}$ :

**THEOREM 6.2.** *There is a mapping*

$$MA : XQ[=, \text{child}, \text{not}] \rightarrow \mathcal{M}_{\cup}^{\square}[=, \text{not}]$$

such that for each  $XQ[=, \text{child}, \text{not}]$  query  $Q$ ,

1. for any XML document tree  $T$ ,

$$[C(Q(T))] = MA(Q)(\{\{N : \$ROOT, V : [C(T)]\}\}),$$

2.  $MA(Q)$  can be computed in space  $O(\log |Q|)$ , and

3.  $|MA(Q)| = O(|Q|)$ .

**Proof Sketch.** It is easy to verify that the function  $MA$  of Figure 3 satisfies conditions (1) to (3) of our theorem.  $\square$

For atomic equality,  $\sigma_{1.V=2.V}$  in the definition of  $MA$  is to be implemented as  $\sigma_{1.V.\text{label}=\text{atomic}.2.V.\text{label}}$ . Note that on a  $XQ[=, \text{child}]$  query  $Q$ ,  $MA(Q)$  is a  $\mathcal{M}_{\cup}^{\square}[=]$  query.

### Translation from $\mathcal{M}_{\cup}^{\square}$ to Core XQuery

Let  $T$  be the following canonical translation from complex values to trees:

$$T(\langle A_1 : v_1, A_2 : v_2 \rangle) = \langle \text{tup} \langle a_1 \rangle T(v_1) \langle / a_1 \rangle \langle a_2 \rangle T(v_2) \langle / a_2 \rangle \langle / \text{tup} \rangle$$

$$T(\{v_1, \dots, v_n\}) = \langle \text{list} \rangle T(v_1) \dots T(v_n) \langle / \text{list} \rangle$$

Note that  $T$  is not the inverse of the mapping  $C$  that we introduced above to map from XML trees to complex values constructed from tuples and lists.

Let  $\mathcal{M}_{\cup}^{\square, (\cdot, \cdot)}$  denote the monad algebra queries on lists and pairs (rather than on tuples of arbitrary arity). For both  $=_{deep}$  and  $=_{atomic}$ , we have

**THEOREM 6.3.** *There is a mapping*

$$XQ : \mathcal{M}_{\cup}^{\square}[=] \rightarrow XQ[=, \text{child}]$$

such that for each  $\mathcal{M}_{\cup}^{\square}[=]$  query  $Q$ ,

1. for any complex value  $v$ ,

$$T(Q(v)) = XQ(Q)(\$ROOT)(T(v)),$$

2.  $XQ(Q)$  can be computed in space  $O(\log |Q|)$ , and

3. If  $Q$  is a  $\mathcal{M}_{\cup}^{\square, (\cdot, \cdot)}$  query,  $|XQ(Q)| = O(|Q|)$ .

Due to space limitations, the mapping is not presented here. However, it is rather straightforward to define.

## 7. COMBINED COMPLEXITY OF XQ

Theorems 6.2 and 6.3 provide LOGLIN-reductions back and forth between queries in monad algebra on lists and XQ. Note that each XML document is an XQuery. We can compose query and data into a single query (certainly in LOGLIN). All the complexity results of this section involve complexity classes that are closed under LOGLIN-reductions. Thus combined complexity for Core XQuery with composition will never be harder than query complexity, and we only need to study the latter.

**COROLLARY 7.1.** *With respect to query complexity,*

- $XQ[=_{deep}, \text{child}]$  is  $TA[2^{O(n)}, O(n)]$ -hard under LOGLIN-reductions and in EXPSPACE;
- $XQ[=_{atomic}, \text{child}, \text{not}]$  is  $TA[2^{O(n)}, O(n)]$ -complete under LOGLIN-reductions;
- $XQ[=_{atomic}, \text{child}]$  is NEXPTIME-complete.

$$\begin{aligned}
MA & : XQ[=, \text{child}, \text{not}] \rightarrow \{\langle N : \text{varname}, V : \tau \rangle\} \rightarrow \tau' \\
MA(\alpha \beta) & := MA(\alpha) \cup MA(\beta) \\
MA(\langle a/\rangle) & := \langle \text{label} : a, \text{children} : [] \rangle \circ \text{sng} \\
MA(\langle a \rangle \alpha \langle /a \rangle) & := \langle \text{label} : a, \text{children} : MA(\alpha) \rangle \circ \text{sng} \\
MA(\$x_i) & := \pi_i \\
MA(\$x_i/t) & := \pi_i \circ \text{flatmap}(\pi_{\text{children}} \circ \sigma_{\text{label}=t}) \\
MA(\text{for } \$x \text{ in } \alpha \text{ return } \beta) & := \langle 1 : \text{id}, 2 : MA(\alpha) \rangle \circ \text{pairwith}_2 \circ \text{flatmap}((\pi_1 \cup (\langle N : \$x, V : \pi_2 \rangle \circ \text{sng})) \circ MA(\beta)) \\
MA((\text{let } \$x := \alpha) \beta) & := \langle 1 : \text{id}, 2 : MA(\alpha) \rangle \circ \text{pairwith}_2 \circ \text{flatmap}((\pi_1 \cup (\langle N : \$x, V : \pi_2 \rangle \circ \text{sng})) \circ MA(\beta)) \\
MA(\text{if } \alpha \text{ then } \beta) & := \langle 1 : \text{id}, 2 : MA(\alpha) \circ \text{true} \rangle \circ \text{pairwith}_2 \circ \text{flatmap}(\pi_1 \circ MA(\beta)) \\
MA(\text{not } \alpha) & := MA(\alpha) \circ \text{map}(\langle \rangle) \circ \text{not} \\
MA(\$x = \$y) & := \langle 1 : \sigma_{N=\$x}, 2 : \sigma_{N=\$y} \rangle \circ \text{pairwith}_1 \circ \text{flatmap}(\text{pairwith}_2) \circ \sigma_{1.V=2.V}
\end{aligned}$$

**Figure 3: Mapping from  $XQ[=, \text{child}, \text{not}]$  to  $\mathcal{M}_{\cup}^{\square}[=, \text{not}]$ .**

**Proof.** This follows immediately from the results of Section 5 on  $\mathcal{M}_{\cup}^{\square}$  and, for the upper bounds, the LOGSPACE-reduction from  $XQ[=, \text{child}]$  to  $\mathcal{M}_{\cup}^{\square}[=]$  of Theorem 6.2 resp., for the lower bounds, the LOGLIN-reduction from  $\mathcal{M}_{\cup}^{\square}[=]$  to  $XQ[=, \text{child}]$  of Theorem 6.3.  $\square$

The complexity classes in which our XQuery evaluation problems reside are large enough that minor extensions, such as supporting all XPath axes, do not matter.

**THEOREM 7.2.** *W.r.t. combined complexity,*

- $XQ[=_{\text{deep}}, \text{all axes}]$  is in EXPSPACE,
- $XQ[=_{\text{atomic}}, \text{all axes}, \text{not}]$  is in  $TA[2^{O(n)}, O(n)]$ , and
- $XQ[=_{\text{atomic}}, \text{all axes}]$  is in NEXPTIME.

A proof of this is not difficult but is beyond the scope of this paper and will be given in its long version. For the EXPSPACE bound, for instance, the proof idea is that no value computable by an XQ query is greater than of doubly exponential size in the size of the input (see also Proposition 3.4 for the analogous fact for monad algebra on lists). Thus we can use “addresses” of singly exponential size to refer to subtrees of an intermediate result. While we cannot store the subtree as a whole, we can always recompute it from these addresses and the query. There are only linearly many variables in an XQuery, so we need only linearly many exp-size registers to evaluate the query. This yields an EXPSPACE algorithm.

Note that this EXPSPACE algorithm is quite robust and allows to add a number of XQuery features that we excluded from XQ, such as counting, (document position) arithmetics, and duplicate elimination.

## 8. DATA COMPLEXITY OF XQ

By Proposition 5.1, monad algebra is in LOGSPACE with respect to data complexity. Since LOGSPACE is closed under composition (cf. [33]) and the mapping  $C$  is clearly in LOGSPACE,

**COROLLARY 8.1.**  $XQ[=_{\text{deep}}, \text{child}]$  is in LOGSPACE w.r.t. data complexity.

We can improve on this result.

The data complexity of XQuery is so low that we need to be careful about how the XML data is represented. We distinguish the cases of representation by a DOM tree (i.e., a pointer structure) and representation by a string (an XML document). As we show, the complexity is (presumably) slightly lower on strings than on trees, even though the former require parsing the input. It turns out that the complexity bounds are precisely the same as for XPath [16, 35].

**THEOREM 8.2.**  $XQ[=_{\text{deep}}, \text{all axes}]$  is LOGSPACE-complete under  $NC_1$ -reductions (data complexity) if the input is given as a DOM tree.

**THEOREM 8.3.**  $XQ[=_{\text{deep}}, \text{all axes}]$  is in  $TC_0$  w.r.t. data complexity if the XML input is given as a character string.

The latter result is by reduction to first-order logic with majority quantifiers (FOM). It is known that  $TC_0$  is equivalent to the class of languages recognizable using FOM sentences [4].

Unfortunately, for lack of space, we have to refer to the long version of this paper for proofs of these two results.

## Acknowledgments

I thank Michael Benedikt, Peter Buneman, Leonid Libkin, Nicole Schweikardt, Val Tannen, and Andrei Voronkov for their remarks that helped to improve this paper.

## 9. REFERENCES

- [1] S. Abiteboul and C. Beeri. “The Power of Languages for the Manipulation of Complex Values”. *VLDB J.*, **4**(4):727–794, 1995.
- [2] S. Abiteboul and G. G. Hillebrand. “Space Usage in Functional Query Languages”. In *Proc. of the 5th International Conference on Database Theory (ICDT)*, pages 439–454, 1995.
- [3] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [4] D. A. M. Barrington, N. Immerman, and H. Straubing. “On Uniformity within  $NC_1$ ”. *Journal of Computer and System Sciences*, **41**(3):274–306, 1990.
- [5] L. Berman. “The Complexity of Logical Theories”. *Theoretical Comput. Sci.*, **11**:216–224, 1980.

- [6] P. Buneman, S. A. Naqvi, V. Tannen, and L. Wong. "Principles of Programming with Complex Objects and Collection Types". *Theor. Comput. Sci.*, **149**(1):3–48, 1995.
- [7] A. K. Chandra, D. C. Kozen, and L. J. Stockmeyer. "Alternation". *Journal of the ACM*, **28**:114–133, 1981.
- [8] E. Dantsin, T. Eiter, G. Gottlob, and A. Voronkov. "Complexity and Expressive Power of Logic Programming". *ACM Computing Surveys*, **33**(3):374–425, Sept. 2001.
- [9] E. Dantsin and A. Voronkov. "Complexity of Query Answering in Logic Databases with Complex Values". In *Proc. LFCS'97, LNCS 1234*, pages 56–66, 1997.
- [10] E. Dantsin and A. Voronkov. Expressive Power and Data Complexity of Query Languages for Trees and Lists. In *Proceedings of the 19th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'00)*, pages 157–165, Dallas, Texas, USA, 2000. ACM Press.
- [11] M. Fernandez, J. Siméon, and P. Wadler. An Algebra for XML Query. In *Proc. FSTTCS 2000, LNCS 1974*, pages 11–45. Springer-Verlag, 2000.
- [12] M. F. Fernandez and J. Siméon. "Building an Extensible XQuery Engine: Experiences with Galax (Extended Abstract)". In *Proc. XSYM*, pages 1–4, 2004.
- [13] J. Ferrante and C. Rackoff. "A Decision Procedure for the First Order Theory of Real Addition with Order". *SIAM J. Comput.*, **4**(1):69–76, 1975.
- [14] D. Florescu, C. Hillery, D. Kossmann, P. Lucas, F. Riccardi, T. Westmann, M. J. Carey, A. Sundararajan, and G. Agrawal. "The BEA/XQRL Streaming XQuery Processor". In *Proceedings of the 29th International Conference on Very Large Data Bases (VLDB)*, pages 997–1008, 2003.
- [15] G. Gottlob, C. Koch, and R. Pichler. "Efficient Algorithms for Processing XPath Queries". In *Proceedings of the 28th International Conference on Very Large Data Bases (VLDB)*, pages 95–106, Hong Kong, China, 2002.
- [16] G. Gottlob, C. Koch, and R. Pichler. "The Complexity of XPath Query Evaluation". In *Proceedings of the 22nd ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS)*, pages 179–190, San Diego, California, 2003.
- [17] R. Greenlaw, H. J. Hoover, and W. L. Ruzzo. *Limits to Parallel Computation: P-Completeness Theory*. Oxford University Press, 1995.
- [18] S. Grumbach, L. Libkin, T. Milo, and L. Wong. "Query Languages for Bags – Expressive Power and Complexity". *SIGACT News*, **27**(2):30–44, 1996.
- [19] S. Grumbach and T. Milo. "Towards Tractable Algebras for Bags". In *Proceedings of the 12th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'93)*, pages 49–58, 1993.
- [20] S. Grumbach and V. Vianu. "Tractable Query Languages for Complex Object Databases". In *Proceedings of the 10th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'91)*, pages 315–327, 1991.
- [21] J. Hartmanis, P. M. Lewis II, and R. E. Stearns. "Hierarchies of Memory Limited Computations". In *Proc. Sixth Annual IEEE Symposium on Switching Circuit Theory and Logical Design*, pages 179–190, 1965.
- [22] J. Hidders, J. Paredaens, R. Verkammen, and S. Demeyer. "A Light but Formal Introduction to XQuery". In *Proc. XSYM*, pages 5–20, 2004.
- [23] R. Hull and J. Su. "On Accessing Object-oriented Databases: Expressive Power, Complexity, and Restrictions". In *Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data (SIGMOD'89)*, pages 147–158, 1989.
- [24] R. Hull and J. Su. "Algebraic and Calculus Query Languages for Recursively Typed Complex Objects". *J. Comput. Syst. Sci.*, **47**(1):121–156, 1993.
- [25] G. Jaeschke and H.-J. Schek. "Remarks on the Algebra of Non First Normal Form Relations". In *Proc. PODS'82*, pages 124–138, 1982.
- [26] D. S. Johnson. "A Catalog of Complexity Classes". In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume 1, chapter 2, pages 67–161. Elsevier Science Publishers B.V., 1990.
- [27] C. Koch, S. Scherzinger, N. Schweikardt, and B. Stegmaier. "Schema-based Scheduling of Event Processors and Buffer Minimization for Queries on Structured Data Streams". In *Proceedings of the 30th International Conference on Very Large Data Bases (VLDB)*, Toronto, Canada, 2004.
- [28] G. Kuper and M. Y. Vardi. "A New Approach to Database Logic". In *Proceedings of the 3rd ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'84)*, pages 86–96, 1984.
- [29] G. Kuper and M. Y. Vardi. "On the Complexity of Queries in the Logical Data Model". *Theor. Comput. Sci.*, **116**(1&2):33–57, 1993.
- [30] L. Libkin and L. Wong. "Query Languages for Bags and Aggregate Functions". *Journal of Computer and System Sciences*, **55**(2):241–272, 1997.
- [31] B. Ludäscher, P. Mukhopadhyay, and Y. Papakonstantinou. "A Transducer-Based XML Query Processor". In *Proceedings of the 28th International Conference on Very Large Data Bases (VLDB)*, pages 227–238, 2002.
- [32] A. Marian and J. Siméon. "Projecting XML Documents". In *Proceedings of the 29th International Conference on Very Large Data Bases (VLDB)*, pages 213–224, 2003.
- [33] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [34] J. Paredaens and D. Van Gucht. "Possibilities and Limitations of Using Flat Operators in Nested Algebra Expressions". In *Proc. PODS*, pages 29–38, 1988.
- [35] L. Segoufin. "Typing and Querying XML Documents: Some Complexity Bounds". In *Proceedings of the 22nd ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS)*, pages 167–178, 2003.
- [36] D. Suciu and V. Tannen. "A Query Language for NC". In *Proceedings of the ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'94)*, pages 167–178, 1994.
- [37] V. Tannen, P. Buneman, and L. Wong. "Naturally Embedded Query Languages". In *Proc. of the 4th International Conference on Database Theory (ICDT)*, pages 140–154, 1992.
- [38] M. Y. Vardi. "The Complexity of Relational Query Languages". In *Proc. 14th Annual ACM Symposium on Theory of Computing (STOC'82)*, pages 137–146, San Francisco, CA USA, May 1982.
- [39] S. Vorobyov and A. Voronkov. "Complexity of Nonrecursive Logic Programs with Complex Values". In *Proceedings of the 17th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'98)*, 1998.
- [40] A. Voronkov. Personal communication, 2004.
- [41] L. Wong. "Normal Forms and Conservative Extension Properties for Query Languages over Collection Types". *Journal of Computer and System Sciences*, **52**(3):495–505, 1996.
- [42] World Wide Web Consortium. "XQuery 1.0 and XPath 2.0 Formal Semantics. W3C Working Draft (Aug. 16th 2002), 2002. <http://www.w3.org/TR/query-algebra/>.
- [43] "XML Query Use Cases. W3C Working Draft 02 May 2003", 2003. <http://www.w3.org/TR/xmlquery-use-cases/>.