

Database Query Processing Using Finite Cursor Machines

Martin Grohe¹, Yuri Gurevich², Dirk Leinders³, Nicole Schweikardt¹,
Jerzy Tyszkiewicz⁴, and Jan Van den Bussche³

¹ Humboldt-University Berlin

² Microsoft Research

³ Hasselt University and Transnational University of Limburg

⁴ University of Warsaw

Abstract. We introduce a new abstract model of database query processing, *finite cursor machines*, that incorporates certain data streaming aspects. The model describes quite faithfully what happens in so-called “one-pass” and “two-pass query processing”. Technically, the model is described in the framework of abstract state machines. Our main results are upper and lower bounds for processing relational algebra queries in this model, specifically, queries of the semijoin fragment of the relational algebra.

1 Introduction

We introduce and analyze *finite cursor machines*, an abstract model of database query processing. Data elements are viewed as “indivisible” abstract objects with a vocabulary of arbitrary, but fixed, functions. Relational databases consist of finitely many finite relations over the data elements. Relations are considered as tables whose rows are the tuples in the relation. Finite cursor machines can operate in a finite number of *modes* using an *internal memory* in which they can store bit strings. They access each relation through finitely many cursors, each of which can read one row of a table at any time. The answer to a query, which is also a relation, can be given through a suitable output mechanism. The model incorporates certain “streaming” or “sequential processing” aspects by imposing two restrictions: First, the cursors can only move on the tables sequentially in one direction. Thus once the last cursor has left a row of a table, this row can never be accessed again during the computation. Second, the internal memory is limited. For our lower bounds, it will be sufficient to put an $o(n)$ restriction on the internal memory size, where n is the size (that is, the number of entries) of the input database. For the upper bounds, no internal memory will be needed. The model is clearly inspired by the *abstract state machine (ASM)* methodology [16], and indeed we will formally define our model using this methodology. The model was first presented in a talk at the ASM 2004 workshop [29].

Algorithms and lower bounds in various data stream models have received considerable attention in recent years both in the theory community (e.g., [1, 2, 5, 6,

13, 14, 18, 25]) and the database systems community (e.g., [3, 4, 7, 12, 15, 20, 26]). Note that our model is fairly powerful; for example, the multiple cursors can easily be used to perform multiple sequential scans of the input data. But more than that; by moving several cursors asynchronously over the same table, entries in different, possibly far apart, regions of the table can be read and processed simultaneously. This way, different regions of the same or of different tables can “communicate” with each other without requiring any internal memory, which makes it difficult to use communication complexity to establish lower bounds. The model is also powerful in that it allows arbitrary functions to access and process data elements. This feature is very convenient to model “built in” standard operations on data types like integers, floating point numbers, or strings, which may all be part of the universe of data elements.

Despite these powerful features, the model is weak in many respects. We show that a finite cursor machine with internal memory size $o(n)$ cannot even test whether two sets A and B , given as lists, are disjoint, even if besides the lists A and B , also their reversals are given as input. However, if two sets A and B are given as *sorted* lists, a machine can easily compute the intersection. Aggarwal et al. [1] have already made a convincing case for combining streaming computations with sorting, and we will consider an extension of the model with a sorting primitive.

Our main results are concerned with evaluating *relational algebra queries* in the finite cursor machine model. Relational algebra forms the core of the standard query language SQL and is thus of fundamental importance for databases. We prove that, when all sorted versions of the database relations are provided as input, every operator of the relational algebra can be computed, except for the *join*. The latter exception, however, is only because the output size of a join can be quadratic, while finite cursor machines by their very definition can output only a linear number of different tuples. A *semijoin* is a projection of a join between two relations to the columns of one of the two relations (note that the projection prevents the result of a semijoin from getting larger than the relations to which the semijoin operation is applied). The *semijoin algebra* is then a natural fragment of the relational algebra that may be viewed as a generalization of acyclic conjunctive queries [9, 22, 21, 30]. When sorted versions of the database relations are provided as input, semijoins can be computed by finite cursor machines. Consequently, every query in the semijoin fragment of the relational algebra can be computed by a composition of finite cursor machines and sorting operations. This is interesting because it models quite faithfully what is called “one-pass” and “two-pass processing” in database systems [11]. The question then arises: are intermediate sorting operations really needed? Equivalently, can every semijoin-algebra query already be computed by a single machine on sorted inputs? We answer this question negatively in a very strong way, and this is our main technical result: Just a composition of two semijoins $R \times (S \times T)$ with R and T unary relations and S a binary relation is not computable by a finite cursor machine with internal memory size $o(n)$ working on sorted inputs. This result is quite sharp, as we will indicate.

The paper is structured as follows: After fixing some notation in Section 2, the notion of finite cursor machines is introduced in Section 3. The power of $O(1)$ -FCMs and of $o(n)$ -FCMs is investigated in Sections 4 and 5. Some concluding remarks and open questions can be found in Section 6.

Due to space limitations, some technical details of our proofs had to be deferred to the full version of this paper, available on the authors' websites.

2 Preliminaries

Throughout the paper we fix an arbitrary, typically infinite, universe \mathbb{E} of “data elements”, and we fix a database schema \mathcal{S} . I.e., \mathcal{S} is a finite set of relation names, where each relation name has an associated arity, which is a natural number. A database \mathbf{D} with schema \mathcal{S} assigns to each $R \in \mathcal{S}$ a finite, nonempty set $\mathbf{D}(R)$ of k -tuples of data elements, where k is the arity of R . In database terminology the tuples are often called *rows*. The *size* of database \mathbf{D} is defined as the total number of rows in \mathbf{D} .

A *query* is a mapping Q from databases to relations, such that the relation $Q(\mathbf{D})$ is the answer of the query Q to database \mathbf{D} . The *relational algebra* is a basic language used in database theory to express exactly those queries that can be composed from the actual database relations by applying a sequence of the following operations: union, intersection, difference, projection, selection, and joins. The meaning of the first three operations should be clear, the *projection* operator $\pi_{i_1, \dots, i_k}(R)$ returns the projection of a relation R to its components i_1, \dots, i_k , the *selection* operator $\sigma_{p(i_1, \dots, i_k)}(R)$ returns those tuples from R whose i_1 th, \dots , i_k th components satisfy the predicate p , and the *join* operator $R \bowtie_{\theta} S$ (where θ is a conjunction of equalities of the form $\bigwedge_{s=1}^k x_{i_s} = y_{j_s}$) is defined as $\{(\bar{a}, \bar{b}) : \bar{a} \in R, \bar{b} \in S, a_{i_s} = b_{j_s} \text{ for all } s \in \{1, \dots, k\}\}$. A natural sub-language of the relational algebra is the so-called *semijoin algebra* where, instead of ordinary joins, only *semijoin* operations of the form $R \ltimes_{\theta} S$ are allowed, defined as $\{\bar{a} \in R : \exists \bar{b} \in S : a_{i_s} = b_{j_s} \text{ for all } s \in \{1, \dots, k\}\}$.

To formally introduce our computation model, we need some basic notions from mathematical logic such as (many-sorted) vocabularies, structures, terms, and atomic formulas.

3 Finite Cursor Machines

In this section we formally define *finite cursor machines* using the methodology of Abstract State Machines (ASMs). Intuitively, an ASM can be thought of as a transition system whose states are described by many-sorted first-order structures (or algebras)¹. Transitions change the interpretation of some of the symbols—those in the *dynamic* part of the vocabulary—and leave the remaining

¹ Beware that “state” refers here to what for Turing machines is typically called “configuration”; the term “mode” is used for what for Turing machines is typically called “state”.

symbols—those in the *static* part of the vocabulary—unchanged. Transitions are described by a finite collection of simple update rules, which are “fired” simultaneously (if they are inconsistent, no update is carried out). A crucial property of the sequential ASM model, which we consider here, is that in each transition only a limited part of the state is changed. The detailed definition of sequential ASMs is given in the Lipari guide [16], but our presentation will be largely self-contained.

We now describe the formal model of finite cursor machines.

The Vocabulary: The *static vocabulary* of a finite cursor machine (FCM) consists of two parts, \mathcal{Y}_0 (providing the background structure) and \mathcal{Y}_S (providing the particular input).

\mathcal{Y}_0 consists of three sorts: *Element*, *Bitstring*, and *Mode*. Furthermore, \mathcal{Y}_0 may contain an arbitrary number of functions and predicates, as long as the output sort of each function is *Bitstring*. Finally, \mathcal{Y}_0 contains an arbitrary but finite number of constant symbols of sort *Mode*, called *modes*. The modes *init*, *accept*, and *reject* are always in \mathcal{Y}_0 .

\mathcal{Y}_S provides the input. For each relation name $R \in \mathcal{S}$, there is a sort Row_R in \mathcal{Y}_S . Moreover, if the arity of R is k , we have function symbols $\text{attribute}_R^i: \text{Row}_R \rightarrow \text{Element}$ for $i = 1, \dots, k$. Furthermore, we have a constant symbol \perp_R of sort Row_R . Finally, we have a function symbol $\text{next}_R: \text{Row}_R \rightarrow \text{Row}_R$ in \mathcal{Y}_S .

The *dynamic vocabulary* \mathcal{Y}_M of an FCM M contains only constant symbols. This vocabulary always contains the symbol *mode* of sort *Mode*. Furthermore, there can be a finite number of symbols of sort *Bitstring*, called *registers*. Moreover, for each relation name R in the database schema, there are a finite number of symbols of sort Row_R , called *cursors on R*.

The Initial State: Our intention is that FCMs will work on databases. Database relations, however, are sets, while FCMs expect lists of tuples as inputs. Therefore, formally, the input to a machine is an *enumeration* of a database, which consists of enumerations of the database relations, where an enumeration of a relation is simply a listing of all tuples in some order. An FCM M that is set to run on an enumeration of a database \mathbf{D} then starts with the following structure \mathcal{M} over the vocabulary $\mathcal{Y}_0 \cup \mathcal{Y}_S \cup \mathcal{Y}_M$: The interpretation of *Element* is \mathbb{E} ; the interpretation of *Bitstring* is the set of all finite bitstrings; and the interpretation of *Mode* is simply given by the set of modes themselves. For technical reasons, we must assume that \mathbb{E} contains an element \perp . For each $R \in \mathcal{S}$, the sort Row_R is interpreted by the set $\mathbf{D}(R) \cup \{\perp_R\}$; the function attribute_R^i is defined by $(x_1, \dots, x_k) \mapsto x_i$, and $\perp_R \mapsto \perp$; finally, the function next_R maps each row to its successor in the list, and maps the last row to \perp_R . The dynamic symbol *mode* initially is interpreted by the constant *init*; every register contains the empty bitstring; and every cursor on a relation R contains the first row of R .

The Program of an FCM: A *program* for the machine M is now a program as defined as a basic sequential program in the sense of ASM theory, with the important restriction that all basic updates concerning a cursor c on R must be of the form $c := \text{next}_R(c)$.

Thus, basic update rules of the following three forms are rules: $mode := t$, $r := t$, and $c := next_R(c)$, where t is a term over $\mathcal{Y}_0 \cup \mathcal{Y}_S \cup \mathcal{Y}_M$, and r is a register and c is a cursor on R . The semantics of these rules is the obvious one: Update the dynamic constant by the value of the term. Update rules r_1, \dots, r_m can be combined to a new rule $\text{par } r_1 \dots r_m \text{ endpar}$, the semantics of which is: Fire rules r_1, \dots, r_m in parallel; if they are inconsistent do nothing. Furthermore, if r_1 and r_2 are rules and φ is an atomic formula over $\mathcal{Y}_0 \cup \mathcal{Y}_S \cup \mathcal{Y}_M$, then also $\text{if } \varphi \text{ then } r_1 \text{ else } r_2 \text{ endif}$ is a rule. The semantics is obvious.

Now, an FCM program is just a single rule. (Since finitely many rules can be combined to one using the $\text{par} \dots \text{end}$ construction, one rule is enough.)

The Computation of an FCM: Starting with the initial state, successively apply the (single rule of the FCM's) program until $mode$ is equal to *accept* or to *reject*. Accordingly, we say that M terminates and *accepts*, respectively, *rejects* its input.

Given that inputs are *enumerations* of databases, we must be careful to define the result of a computation on a database. We agree that an FCM *accepts* a database \mathbf{D} if it accepts *every* enumeration of \mathbf{D} . This already allows us to use FCMs to compute decision queries. In the next paragraph we will see how FCMs can output lists of tuples. We then say that an FCM M computes a query Q if on each database \mathbf{D} , the output of M on *any* enumeration of \mathbf{D} is an enumeration of the relation $Q(\mathbf{D})$. Note that later we will also consider FCMs working only on sorted versions of database relations: in that case there is no ambiguity.

Producing Output: We can extend the basic model so that the machine can output a list of tuples. To this end, we expand the dynamic vocabulary \mathcal{Y}_M with a finite number of constant symbols of sort **Element**, called *output registers*, and with a constant of sort **Mode**, called the *output mode*. The output registers can be updated following the normal rules of ASMs. In each state of the finite cursor machine, when the output mode is equal to the special value *out*, the tuple consisting of the values in the output registers (in some predefined order) is output; when the output mode is different from *out*, no tuple is output. The initial settings of the output registers and the output mode are as follows: each output register contains the value \perp ; the output mode is equal to *init*. We denote the output of a machine M working on a database \mathbf{D} by $M(\mathbf{D})$.

Space Restrictions: For considering FCMs whose bitstring registers are restricted in size, we use the following notation: Let M be a finite cursor machine and \mathcal{F} a class of functions from \mathbb{N} to \mathbb{N} . Then we say that M is an \mathcal{F} -*machine* (or, an \mathcal{F} -*FCM*) if there is a function $f \in \mathcal{F}$ such that, on each database enumeration \mathbf{D} of size n , the machine only stores bitstrings of length $f(n)$ in its registers. We are mostly interested in $O(1)$ -FCMs and $o(n)$ -FCMs. Note that the latter are quite powerful. For example, such machines can easily store the positions of the cursors. On the other hand, $O(1)$ -machines are equivalent to FCMs that do not use registers at all (because bitstrings of constant length could also be simulated by finitely many *modes*).

Example 1. The following FCM program works on a ternary relation $R(A, B, C)$ and produces the sum of attributes A and B for each row with C at least 100.

```

if outputmode = out then
par outputmode := init, c := nextR(c) endpar
else if outputmode <> out and attributeR3(c) > 100 then
par outputmode := out, out1 := attributeR1(c) + attributeR2(c) endpar
else c := nextR(c) endif endif

```

3.1 Discussion of the Model

Storing Bitstrings instead of Data Elements: An important question about our model is the strict separation between data elements and bitstrings. Indeed, data elements are abstract entities, and our background structure may contain arbitrary functions and predicates, mixing data elements and bitstrings, with the important restriction that the output of a function is always a bitstring. At first sight, a simpler way to arrive at our model would be without bitstrings, simply considering an arbitrary structure on the universe of data elements. Let us call this variation of our model the “universal model”.

Note that the universal model can easily become computationally complete. It suffices that finite strings of data elements can somehow be represented by other data elements, and that the background structure supplies the necessary manipulation functions for that purpose. Simple examples are the natural numbers with standard arithmetic, or the strings over some finite alphabet with concatenation. Thus, if we would want to prove complexity lower bounds in the universal model, while retaining the abstract nature of data elements and operations on them, it would be necessary to formulate certain logical restrictions on the available functions and predicates on the data elements. Finding interesting such restrictions is not clear to us. In the model with bitstrings, however, one can simply impose restrictions on the length of the bitstrings stored in registers, and that is precisely what we will do. Of course, the unlimited model with bitstrings can also be computationally complete. It suffices that the background structure provides a coding of data elements by bitstrings.

Element Registers: The above discussion notwithstanding, it might still be interesting to allow for registers that can remember certain data elements that have been seen by the cursors, but without arbitrary operations on them. Formally, we would expand the dynamic vocabulary \mathcal{Y}_M with a finite number of constant symbols of sort *Element*, called *element registers*. It is easy to see, however, that such element registers can already be simulated by using additional cursors, and thus do not add anything to the basic model.

Running Time and Output Size: A crucial property of FCMs is that all cursors are one-way. In particular, an FCM can perform only a linear number of steps where a cursor is advanced. As a consequence, an FCM with output can output only a linear number of different tuples. On the other hand, if the background structure is not restricted in any way, arbitrary computations on the register contents can occur in between cursor advancements. As a matter of fact, in this paper we will present a number of positive results and a number of negative results. For the positive results, registers will never be needed,

and in particular, FCMs run in linear time. For the negative results, arbitrary computations on the registers will be allowed.

Look-ahead: Note that the terms in the program of an FCM can contain nested applications of the function $next_R$, such as $next_R(next_R(c))$. In some sense, such nestings of depth up to d correspond to a *look-ahead* where the machine can access the current cursor position as well as the next d positions. It is, however, straightforward to see that every k -cursor FCM with look-ahead $\leq d$ can be simulated by a $(k \times d)$ -cursor FCM with look-ahead 0. Thus, throughout the remainder of this paper we will w.l.o.g. restrict attention to FCMs that have look-ahead 0, i.e., to FCMs where the function $next_R$ never occurs in if-conditions or in update rules of the form $mode := t$ or $r := t$.

The Number of Cursors: In principle we could allow more than constantly many cursors, which would enable us to store that many data elements. We stick with the constant version for the sake of technical simplicity, and also because our *upper* bounds only need a constant number of cursors. Note, however, that our main *lower* bound result can be extended to a fairly big number of cursors (cf. Remark 11).

4 The Power of $O(1)$ -Machines

We start with a few simple observations on the database query processing capabilities of FCMs, with or without sorting, and show that sorting is really needed.

Let us first consider *compositions* of FCMs in the sense that one machine works on the outputs of several machines working on a common database.

Proposition 2. *Let M_1, \dots, M_r be FCMs working on a schema \mathcal{S} , let \mathcal{S}' be the output schema consisting of the names and arities of the output lists of M_1, \dots, M_r , and let M_0 be an FCM working on schema \mathcal{S}' . Then there exists an FCM M working on schema \mathcal{S} , such that $M(\mathbf{D}) = M_0(\mathbf{D}')$, for each database \mathbf{D} with schema \mathcal{S} and the database \mathbf{D}' that consists of the output relations $M_1(\mathbf{D}), \dots, M_r(\mathbf{D})$.*

The proof is obvious: Each row in a relation R_i of database \mathbf{D}' is an output row of a machine M_i working on \mathbf{D} . Therefore, each time M_0 moves a cursor on R_i , the desired finite cursor machine M will simulate that part of the computation of M_i on \mathbf{D} until M_i outputs a next row.

Let us now consider the operators from relational algebra: Clearly, *selection* can be implemented by an $O(1)$ -FCM. Also, *projection* and *union* can easily be accomplished if either duplicate elimination is abandoned or the input is given in a suitable order. *Joins*, however, are *not* computable by an FCM, simply because the output size of a join can be quadratic, while finite cursor machines can output only a linear number of different tuples.

In stream data management research [4], one often restricts attention to *sliding window joins* for a fixed window size w . This means that the join operator is successively applied to portions of the data, each portion consisting of a number w of consecutive rows of the input relations. It is then straightforward to obtain the following:

Proposition 3. *For every fixed window size $w \in \mathbb{N}$ there is an $O(1)$ -FCM that implements the sliding window join operator of width w . However, no FCM (with registers of arbitrary size) can compute the full join of two relations of arity ≥ 2 .*

Using more elaborate methods, we can moreover show that even checking whether the join is nonempty (so that output size is not an issue) is hard for FCMs. Specifically, we will consider the problem whether two sets intersect, which is the simplest kind of join. We will give two proofs: an elegant one for $O(1)$ -machines, using a proof technique that is simple to apply, and an intricate one for more general $o(n)$ -machines (Theorem 12). Note that the following result is valid for *arbitrary* (but fixed) background structures.

Theorem 4. *There is no $O(1)$ -FCM that checks for two sets R and S whether $R \cap S \neq \emptyset$. (This holds even if also the reversals of R and S are supplied as input.)*

Proof. We give here the proof without the reversals; the proof with reversals can be obtained using the proof technique of our main result (Theorem 10). Let M be an $O(1)$ -FCM that is supposed to check whether $R \cap S \neq \emptyset$. Without loss of generality, we assume that \mathbb{E} is totally ordered by a predicate $<$ in \mathcal{Y}_0 . Using Ramsey’s theorem, we can find an infinite set $V \subseteq \mathbb{E}$ over which the truth of the atomic formulas in M ’s program on tuples of data elements only depends on the way these data elements compare w.r.t. $<$ (details on this can be found, e.g., in Libkin’s textbook [24, Section 13.3]). Now choose $2n$ elements in V , for n large enough, satisfying $a_1 < a'_1 < \dots < a_n < a'_n$, and consider the run of M on $R = \{a_1, \dots, a_n\}$ (listed in that order) and $S = \{a'_n, \dots, a'_1\}$. We say that a pair of cursors “checks” i if in some state during the run, one of the cursors is on a_i and the other one is on a'_i . By the way the lists are ordered, every pair of cursors can check only one i . Hence, some j is not checked. Now replace a'_j in S by a_j . The machine will not notice this, because a_j and a'_j have the same relative order with respect to the other elements in the lists. The intersection of R and S , however, is now nonempty, so M is wrong. \square

Of course, when the sets R and S are given as *sorted* lists, an FCM can easily compute $R \cap S$ by performing one simultaneous scan over the two lists. Moreover, while the full join is still not computable simply because its output is too large, the semijoin $R \times S$ is also easily computed by an FCM on sorted inputs. Furthermore, the same holds for the difference $R - S$. These easy observations motivate us to extend FCMs with sorting, in the spirit of “two-pass query processing” based on sorting [11].

Formally, assume that \mathbb{E} is totally ordered by a predicate $<$ in \mathcal{Y}_0 . Then a relation of arity p can be sorted “lexicographically” in $p!$ different ways: for any permutation ρ of $\{1, \dots, p\}$, let sort_ρ denote the operation that sorts a p -ary relation $\rho(1)$ -th column first, $\rho(2)$ -th column second, and $\rho(p)$ -th column last. By an FCM *working on sorted inputs* of a database \mathbf{D} , we mean an FCM that gets all possible sorted orders of all relations of \mathbf{D} as input lists. We then summarize the above discussion as follows:

Proposition 5. *Each operator of the semijoin algebra (i.e., union, intersection, difference, projection, selection, and semijoin) can be computed by an $O(1)$ -FCM on sorted inputs.*

Corollary 6. *Every semijoin algebra query can be computed by a composition of $O(1)$ -FCMs and sorting operations.*

Proof. Starting from the given semijoin algebra expression we replace each operator by a composition of one FCM with the required sorting operations. \square

The simple proof of the above corollary introduces a lot of intermediate sorting operations. In some cases, intermediate sorting can be avoided by choosing in the beginning a particularly suitable ordering that can be used by *all* the operations in the expression [28].

Example 7. Consider the query $(R - S) \ltimes_{x_2=y_2} T$, where R, S and T are binary relations. Since the semijoin compares the second columns, it needs its inputs sorted on second columns first. Hence, if $R - S$ is computed on $\text{sort}_{(2,1)}(R)$ and $\text{sort}_{(2,1)}(S)$ by some machine M , then the output of M can be piped directly to a machine M' that computes the semijoin on that output and on $\text{sort}_{(2,1)}(T)$. By compositionality (Proposition 2), we can then even compose M and M' into a single FCM. A stupid way to compute the same query would be to compute $R - S$ on $\text{sort}_{(1,2)}(R)$ and $\text{sort}_{(1,2)}(S)$, thus requiring a re-sorting of the output.

The question then arises: can intermediate sorting operations always be avoided? Equivalently, can every semijoin algebra query already be computed by a single machine on sorted inputs? We can answer this negatively. Our proof applies a known result from the classical topic of multihead automata, which is indeed to be expected given the similarity between multihead automata and FCMs.

Specifically, the *monochromatic 2-cycle* query about a binary relation E and a unary relation C asks whether the directed graph formed by the edges in E consists of a disjoint union of 2-cycles where the two nodes on each cycle either both belong to C or both do not belong to C . Note that this query is indeed expressible in the semijoin algebra as “*Is $e_1 \cup e_2 \cup e_3$ empty?*”, where $e_1 := E - (E \ltimes_{\substack{x_2=y_1 \\ x_1=y_2}} E)$, where $e_2 := E \ltimes_{\substack{x_2=y_1 \\ x_1 \neq y_2}} E$, and where $e_3 := (E \ltimes_{x_1=y_1} C) \ltimes_{x_2=y_1}$
 $((\pi_1(E) \cup \pi_2(E)) - C)$

(We use a nonequality in the semijoin condition, but that is easily incorporated in our formalism as well as computed by an FCM on sorted inputs.)

Theorem 8. *The monochromatic 2-cycle query is not computable by an $O(1)$ -FCM on sorted inputs.*

Proof sketch. The proof is via a reduction from the Palindrome problem. As was proved by Hromkovič [19], the set of Palindromes cannot be decided by a one-way multi-head deterministic sensing finite state automaton (1DSeFA(k)). It can be shown that Hromkovič’s proof can be generalized to the presence of an arbitrary but finite number of oblivious right-to-left heads that can only move from right

to left on the input tape sensing other heads, but not read the symbols on the tape. Now let M be an $O(1)$ -FCM that is supposed to solve the monochromatic 2-cycle query. Again using Ramsey’s theorem, we can find an infinite set $V \subseteq \mathbb{E}$ over which the truth of the atomic formulas in M ’s program on tuples of data elements only depends on the way these data elements compare w.r.t. $<$. Hence, there is an $O(1)$ -FCM M' with only $<$ in its rules, and equivalent to M over V . We now come to the reduction. For $a_1 < \dots < a_n \in V$, with n large enough, fix relation E as $\{(a_i, a_{n-i+1}) \mid 1 \leq i \leq n\}$. Given a string $w = w_1 \dots w_n$ over $\{0, 1\}$, we define relation $C = \{a_i \mid w_i = 1\}$. It is then clear that w is a palindrome if and only if E and C form a positive instance to the monochromatic 2-cycle query. From FCM M' we can then construct a $1DSeFA(k)$ that would decide Palindrome, and thus arrive at a contradiction. \square

An important remark is that the above proof only works if the set C is only given in ascending order. In practice, however, one might as well consider sorting operations in descending order, or, for relations of higher arity, arbitrary mixes of ascending and descending orders on different columns. Indeed, that is the general format of sorting operations in the database language SQL. We thus extend our scope to sorting in descending order, and to much more powerful $o(n)$ -machines, in the next section.

5 Descending Orders and the Power of $o(n)$ -Machines

We already know that the computation of semijoin algebra queries by FCMs and sortings in ascending order only requires intermediate sortings. So, the next question is whether the use of descending orders can avoid intermediate sorting. We will answer this question negatively, and will do this even for $o(n)$ -machines (whereas Theorem 8 is proven only for $O(1)$ -machines).

Formally, on a p -ary relation, we now have sorting operations $\text{sort}_{\rho, f}$, where ρ is as before, and $f: \{1, \dots, p\} \rightarrow \{\uparrow, \downarrow\}$ indicates ascending or descending. To distinguish from the terminology of the previous section, we talk about an FCM working on *AD-sorted inputs* to make clear that both ascending and descending orders are available.

Before we show our main technical result, we remark that the availability of sorted inputs using descending order allows $O(1)$ -machines to compute more relational algebra queries. Indeed, we can extract such a query from the proof of Theorem 8. Specifically, the “Palindrome” query about a binary relation R and a unary relation C asks whether R is of the form $\{(a_i, a_{n-i+1}) \mid i = 1, \dots, n\}$ with $a_1 < \dots < a_n$, and $C \subseteq \{a_1, \dots, a_n\}$ such that $a_i \in C \Leftrightarrow a_{n-i+1} \in C$. We can express this query in the relational algebra (using the order predicate in selections). In the following proposition, the lower bound was already shown in Theorem 8, and the upper bound is easy.

Proposition 9. *The “Palindrome” query cannot be solved by an $O(1)$ -FCM on sorted inputs, but can be solved by an $O(1)$ -FCM on AD-sorted inputs.*

We now establish:

Theorem 10. *The query $RST := \text{“Is } R \times_{x_1=y_1} (S \times_{x_2=y_1} T) \text{ nonempty?”}$, where R and T are unary and S is binary, is not computable by any $o(n)$ -FCM working on AD-sorted inputs.*

Proof. For the sake of contradiction, suppose M is a $o(n)$ -FCM computing RST on sorted inputs. Without loss of generality, we can assume that M accepts or rejects the input only when all cursors are positioned at the end of their lists.

Let k be the total number of cursors of M , let r be the number of registers and let m be the number of modes occurring in M 's program. Let $v := \binom{k}{2} + 1$.

Choose n to be a multiple of v^2 , and choose $4n$ values in \mathbb{E} satisfying $a_1 < a'_1 < a_2 < a'_2 < \dots < a_n < a'_n < b_1 < b'_1 < \dots < b_n < b'_n$.

Divide the ordered set $\{1, \dots, n\}$ evenly in v consecutive blocks, denoted by B_1, \dots, B_v . So, B_i equals the set $\{(i-1)\frac{n}{v} + 1, \dots, i\frac{n}{v}\}$. Consider the following permutation of $\{1, \dots, n\}$:

$$\pi : (i-1)\cdot\frac{n}{v} + s \mapsto (v-i)\cdot\frac{n}{v} + s$$

for $1 \leq i \leq v$ and $1 \leq s \leq \frac{n}{v}$. So, π maps subset B_i to subset B_{v-i+1} , and vice versa.

We fix the binary relation S of size $2n$ for the rest of this proof as follows:

$$S := \{(a_\ell, b_{\pi\ell}) : \ell \in \{1, \dots, n\}\} \cup \{(a'_\ell, b'_{\pi\ell}) : \ell \in \{1, \dots, n\}\}.$$

Furthermore, for all sets $I, J \subseteq \{1, \dots, n\}$, we define unary relations $R(I)$ and $T(J)$ of size n as follows:

$$\begin{aligned} R(I) &:= \{a_\ell : \ell \in I\} \cup \{a'_\ell : \ell \in I^c\} \\ T(J) &:= \{b_\ell : \ell \in J\} \cup \{b'_\ell : \ell \in J^c\}, \end{aligned}$$

where I^c denotes $\{1, \dots, n\} - I$. By $\mathbf{D}(I, J)$, we denote the database consisting of the lists $\text{sort}_\uparrow(R(I))$, $\text{sort}_\downarrow(R(I))$, $\text{sort}_\uparrow(T(J))$, $\text{sort}_\downarrow(T(J))$, and all sorted versions of S . It is easy to see that the nested semijoin of $R(I)$, S , and $T(J)$ is empty if, and only if, $(\pi(I) \cap J) \cup (\pi(I)^c \cap J^c) = \emptyset$. Therefore, for each I , the query RST returns *false* on instance $\mathbf{D}(I, \pi(I)^c)$, which we will denote by $\mathbf{D}(I)$ for short. Furthermore, we observe for later use:

the query RST on $\mathbf{D}(I, \pi(J)^c)$ returns *true* if, and only if, $I \neq J$. (*)

To simplify notation a bit, we will in the following use R_\uparrow and T_\uparrow to denote lists $\text{sort}_\uparrow(R(I))$ and $\text{sort}_\uparrow(T(I))$ sorted in ascending order, and we use R_\downarrow and T_\downarrow to denote the lists $\text{sort}_\downarrow(R(I))$ and $\text{sort}_\downarrow(T(I))$ sorted in descending order.

Consider a cursor c on list R_\uparrow of the machine M . In a certain state (i.e., configuration), we say that c is on position ℓ on R_\uparrow if M has executed $\ell-1$ update rules $c := \text{next}_{R_\uparrow}(c)$. I.e., if cursor c is on position ℓ on R_\uparrow , then c sees value a_ℓ or a'_ℓ . We use analogous notation for the sorted lists R_\downarrow , T_\uparrow , and T_\downarrow . I.e., if a cursor c is on position ℓ on R_\downarrow (resp. T_\uparrow , resp. T_\downarrow), then c sees value $a_{n-\ell+1}$ or $a'_{n-\ell+1}$ (resp. b_ℓ or b'_ℓ , resp. $b_{n-\ell+1}$ or $b'_{n-\ell+1}$).

Consider the run of M on $\mathbf{D}(I)$. We say that a pair of cursors of M *checks block* B_i if at some state during the run

- one cursor in the pair is on a position in B_i on R_\uparrow (i.e., the cursor reads an element a_ℓ or a'_ℓ , for some $\ell \in B_i$) and the other cursor in the pair is on a position in B_{v-i+1} on T_\uparrow (i.e., the cursor reads an element $b_{\pi\ell}$ or $b'_{\pi\ell}$, for some $\ell \in B_i$), or
- one cursor in the pair is on a position in B_{v-i+1} on R_\downarrow (i.e., the cursor reads an element a_ℓ or a'_ℓ , for some $\ell \in B_i$) and the other cursor in the pair is on a position in B_i on T_\downarrow (i.e., the cursor reads an element $b_{\pi\ell}$ or $b'_{\pi\ell}$, for some $\ell \in B_i$).

Note that each pair of cursors working on the ascendingly sorted lists R_\uparrow and T_\uparrow or on the descendingly sorted lists R_\downarrow and T_\downarrow , can check at most one block. There are v blocks and at most $\binom{k}{2} < v$ cursor pairs. Hence, there is one block B_{i_0} that is not checked by any pair of cursors working on R_\uparrow and T_\uparrow or on R_\downarrow and T_\downarrow . In order to also deal with pairs of cursors on R_\uparrow and T_\downarrow or on R_\downarrow and T_\uparrow , we further divide each block B_i evenly into v consecutive subblocks, denoted by B_i^1, \dots, B_i^v . So, B_i^j equals the set $\{(i-1)\frac{n}{v} + (j-1)\frac{n}{v^2} + 1, \dots, (i-1)\frac{n}{v} + j\frac{n}{v^2}\}$. We say that a pair of cursors of M checks subblock B_i^j if at some state during the run

- one cursor in the pair is on a position in B_i^j on R_\uparrow (thus reading an element a_ℓ or a'_ℓ , for some $\ell \in B_i^j$) and the other cursor in the pair is on a position in B_i^{v-j+1} on T_\downarrow (thus reading an element $b_{\pi\ell}$ or $b'_{\pi\ell}$, for some $\ell \in B_i^j$), or
- one cursor in the pair is on a position in B_{v-i+1}^{v-j+1} on R_\downarrow (thus reading an element a_ℓ or a'_ℓ , for some $\ell \in B_i^j$) and the other cursor in the pair is on a position in B_{v-i+1}^j on T_\uparrow (thus reading an element $b_{\pi\ell}$ or $b'_{\pi\ell}$, for some $\ell \in B_i^j$).

Note that each pair of cursors working either on R_\uparrow and T_\downarrow or on R_\downarrow and T_\uparrow , can check at most one subblock in B_{i_0} . There are v subblocks in B_{i_0} and at most $\binom{k}{2} < v$ cursor pairs. Hence, there is at least one subblock $B_{i_0}^{j_0}$ that is not checked by any pair of cursors working either on R_\uparrow and T_\downarrow or on R_\downarrow and T_\uparrow . Note that, since the entire block B_{i_0} is not checked by any pair of cursors working either on R_\uparrow and T_\uparrow or on R_\downarrow and T_\downarrow , the subblock $B_{i_0}^{j_0}$ is thus not checked by *any* pair of cursors (on R_\uparrow , R_\downarrow , T_\uparrow , T_\downarrow).

We say that M checks subblock B_i^j if at least one pair of cursors of M checks subblock B_i^j .

At this point it is useful to introduce the following terminology. By “block $B_{i_0}^{j_0}$ on R ”, we refer to the positions in $B_{i_0}^{j_0}$ of list R_\uparrow and to the positions in $B_{v-i_0+1}^{v-j_0+1}$ of list R_\downarrow , i.e., “block $B_{i_0}^{j_0}$ on R ” contains values a_ℓ or a'_ℓ where $\ell \in B_{i_0}^{j_0}$. By “block $B_{i_0}^{j_0}$ on T ”, however, we refer to the positions in $B_{v-i_0+1}^{j_0}$ of list T_\uparrow and to the positions in $B_{i_0}^{v-j_0+1}$ of list T_\downarrow , i.e., “block $B_{i_0}^{j_0}$ on T ” contains values $b_{\pi\ell}$ where $\ell \in B_{i_0}^{j_0}$. Note that this terminology is consistent with the way we have defined the notion of “checking a block”.

It can be shown that there exist at least two different instances $\mathbf{D}(I)$ and $\mathbf{D}(J)$ with the following crucial properties:

1. The query RST returns *false* on $\mathbf{D}(I)$ and on $\mathbf{D}(J)$ (cf. (*));
2. M does not check block $B_{i_0}^{j_0}$ on $\mathbf{D}(I)$, nor on $\mathbf{D}(J)$;
3. $\mathbf{D}(I)$ and $\mathbf{D}(J)$ differ on R and T only in block $B_{i_0}^{j_0}$; and
4. For each cursor c , when c has just left block $B_{i_0}^{j_0}$ (on R or T) in the run on $\mathbf{D}(I)$, the machine M is in the same state as when c has just left block $B_{i_0}^{j_0}$ in the run on $\mathbf{D}(J)$.

Let $\mathcal{V}_0, \mathcal{V}_1, \dots$ be the sequence of states in the run of M on $\mathbf{D}(I)$ and let $\mathcal{W}_0, \mathcal{W}_1, \dots$ be the sequence of states in the run of M on $\mathbf{D}(J)$. Let t_c^I and t_c^J be the points in time when the cursor c of M has just left block $B_{i_0}^{j_0}$ in the run on $\mathbf{D}(I)$ and $\mathbf{D}(J)$, respectively. Because of Property 4 above, $\mathcal{V}_{t_c^I}$ equals $\mathcal{W}_{t_c^J}$ for each cursor c . Note that the start states \mathcal{V}_0 and \mathcal{W}_0 are equal.

Now consider instance $\mathbf{D}_{\text{err}} := \mathbf{D}(I, \pi(J)^c)$. So, \mathbf{D}_{err} has the same lists R_γ, R_λ as $\mathbf{D}(I)$ and the same lists T_γ, T_λ as $\mathbf{D}(J)$. It can now be shown that the (rejecting) runs of M on $\mathbf{D}(I)$ and $\mathbf{D}(J)$ can be combined to obtain a run of M on \mathbf{D}_{err} which rejects \mathbf{D}_{err} . This is wrong, however, because due to (*) the query RST returns *true* on \mathbf{D}_{err} . Finally, this completes the proof of Theorem 10. \square

Remark 11. (a) An analysis of the proof of Theorem 10 shows that we can make the following, more precise statement: *Let $k, m, r, s : \mathbb{N} \rightarrow \mathbb{N}$ such that*

$$k(n)^6 \cdot (\log m(n)) \cdot r(n) \cdot \max(s(n), \log n) = o(n).$$

Then for sufficiently large n , there is no FCM with at most $k(n)$ cursors, $m(n)$ modes, and $r(n)$ registers each holding bitstrings of length at most $s(n)$ that, for all unary relations R, T and binary relations S of size n decides if $R \times_{x_1=y_1} (S \times_{x_2=y_1} T)$ is nonempty. (In the statement of Theorem 10, k, m, r are constant.) This is interesting in particular because we can use a substantial number of cursors, polynomially related to the input size, to store data elements and still obtain the lower bound result.

(b) Note that Theorem 10 is sharp in terms of arity: if S would have been unary (and R and T of arbitrary arities), then the according RST query would have been computable on sorted inputs.

(c) Furthermore, Theorem 10 is also sharp in terms of register bitlength: Assume data elements are natural numbers, and focus on databases with elements from 1 to $O(n)$. If the background provides functions for setting and checking the i -th bit of a bitstring, the query RST is easily computed by an $O(n)$ -FCM.

By a variation of the proof of Theorem 10 we can also show the following strengthening of Theorem 4:

Theorem 12. *There is no $o(n)$ -FCM working on enumerations of unary relations R and S and their reversals, that checks whether $R \cap S \neq \emptyset$.*

Note that Theorems 10 and 12 are valid for arbitrary background structures.

6 Concluding Remarks

A natural question arising from Corollary 6 is whether finite cursor machines with sorting are capable of computing relational algebra queries *beyond* the semijoin algebra. The answer is affirmative:

Proposition 13. *The boolean query over a binary relation R that asks if $R = \pi_1(R) \times \pi_2(R)$ can be computed by an $O(1)$ -FCM working on $\text{sort}_{(1,2),(\uparrow,\uparrow)}(R)$ and $\text{sort}_{(2,1),(\uparrow,\uparrow)}(R)$.*

The proof is straightforward. Note that, using an Ehrenfeucht-game argument, one can indeed prove that the query from Proposition 13 is not expressible in the semijoin algebra [23].

We have not been able to solve the following:

Problem 14. Is there a boolean relational algebra query that cannot be computed by any composition of $O(1)$ -FCMs (or even $o(n)$ -FCMs) and sorting operations?

Under a plausible assumption from parameterized complexity theory [10, 8] we can answer the $O(1)$ -version of this problem affirmatively for FCMs with a decidable background structure.

There are, however, many queries that are not definable in relational algebra, but computable by FCMs with sorting. By their sequential nature, FCMs can easily compare cardinalities of relations, check whether a directed graph is regular, or do modular counting—and all these tasks are not definable in relational algebra. One might be tempted to conjecture, however, that FCMs with sorting cannot go beyond relational algebra with counting and aggregation, but this is false:

Proposition 15. *On a ternary relation G and two unary relations S and T , the boolean query “Check that $G = \pi_{1,2}(G) \times (\pi_1(G) \cup \pi_2(G))$, that $\pi_{1,2}(G)$ is deterministic, and that T is reachable from S by a path in $\pi_{1,2}(G)$ viewed as a directed graph” is not expressible in relational algebra with counting and aggregation, but computable by an $O(1)$ -FCM working on sorted inputs.*

References

1. G. Aggarwal, M. Datar, S. Rajagopalan, and M. Ruhl. On the streaming model augmented with a sorting primitive. *FOCS 2004*, p 540–549.
2. N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *JCSS*, 58:137–147, 1999.
3. M. Altinel and M. Franklin. Efficient filtering of XML documents for selective dissemination of information. *VLDB 2000*, p 53–64.
4. B. Babcock, S. Babu, M. Datar, R. Motwani, and J. Widom. Models and issues in data stream systems. *PODS 2002*, p 1–16.
5. Z. Bar-Yossef, M. Fontoura, and V. Josifovski. On the memory requirements of XPath evaluation over XML streams. *PODS 2004*, p 177–188.
6. Z. Bar-Yossef, M. Fontoura, and V. Josifovski. Buffering in query evaluation over XML streams. *PODS 2005*, p 216–227.

7. C.Y. Chan, P. Felber, M.N. Garofalakis, and R. Rastogi. Efficient filtering of XML documents with XPath expressions. *The VLDB Journal*, 11:354–379, 2002.
8. R.G. Downey and M.R. Fellows. *Parameterized Complexity*. Springer, 1999.
9. R. Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *JACM*, 30:514–550, 1983.
10. J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
11. H. Garcia-Molina, J.D. Ullman, and J. Widom. *Database System Implementation*. Prentice Hall, 1999.
12. T.J. Green, G. Miklau, M. Onizuka, and D. Suciu. Processing XML streams with deterministic automata. *ICDT 2003*, p 173–189.
13. M. Grohe, C. Koch, and N. Schweikardt. Tight lower bounds for query processing on streaming and external memory data. *ICALP 2005*, p 1076–1088.
14. M. Grohe and N. Schweikardt. Lower bounds for sorting with few random accesses to external memory. *PODS 2005*, p 238–249.
15. A.K. Gupta and D. Suciu. Stream processing of XPath queries with predicates. *SIGMOD 2003*, p 419–430.
16. Y. Gurevich. Evolving algebras 1993: Lipari guide. In E. Börger, editor, *Specification and Validation Methods*, p 9–36. Oxford University Press, 1995.
17. L. Hella, L. Libkin, J. Nurmonen, and L. Wong. Logics with aggregate operators. *JACM*, 48(4):880–907, July 2001.
18. M. Henzinger, P. Raghavan, and S. Rajagopalan. Computing on data streams. *External Memory Algorithms. DIMACS Series In Discrete Mathematics And Theoretical Computer Science*, 50:107–118, 1999.
19. J. Hromkovič. One-way multihead deterministic finite automata. *Acta Informatica*, 19:377–384, 1983.
20. Y-N. Law, H. Wang, and C. Zaniolo. Query languages and data models for database sequences and data streams. *VLDB 2004*, p 492–503.
21. D. Leinders and J. Van den Bussche. On the complexity of division and set joins in the relational algebra. *PODS 2005*, p 76–83.
22. D. Leinders, M. Marx, J. Tyszkiewicz, and J. Van den Bussche. The semijoin algebra and the guarded fragment. *JoLLI*, 14(3):331–343, 2005.
23. D. Leinders, J. Tyszkiewicz, and J. Van den Bussche. On the expressive power of semijoin queries. *IPL*, 91(2):93–98, 2004.
24. L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
25. S. Muthukrishnan. *Data Streams: Algorithms and Applications*. Now Publishers Inc, 2005.
26. F. Peng and S.S. Chawathe. XPath queries on streaming data. *SIGMOD 2003*, p 431–442.
27. A.L. Rosenberg. On multi-head finite automata. In *Proceedings of the 6th IEEE Symposium on Switching Circuit Theory and Logical Design*, p 221–228, 1965.
28. D. Simmen, E. Shekita, and T. Malkemus. Fundamental techniques for order optimization. *SIGMOD 1996*, p 57–67.
29. J. Van den Bussche. Finite cursor machines in database query processing. In *Proceedings of the 11th International Workshop on ASMs*, p 61–61, 2004.
30. M. Yannakakis. Algorithms for acyclic database schemes. *VLDB 1981*, p 82–94.