

Dynamic Direct Access of MSO Query Evaluation over Strings

Pierre Bourhis  

Univ. Lille, CNRS, Inria, Centrale Lille, UMR 9189 CRISTAL, F-59000 Lille, France

Florent Capelli  

Univ. Artois, CNRS, UMR 8188, Centre de Recherche en Informatique de Lens (CRIL),
F-62300 Lens, France

Stefan Mengel  

Univ. Artois, CNRS, UMR 8188, Centre de Recherche en Informatique de Lens (CRIL),
F-62300 Lens, France

Cristian Riveros  

Pontificia Universidad Católica de Chile, Santiago, Chile
Millennium Institute for Foundational Research on Data, Santiago, Chile

Abstract

We study the problem of evaluating a Monadic Second Order (MSO) query over strings under updates in the setting of direct access. We present an algorithm that, given an MSO query with first-order free variables represented by an unambiguous variable-set automaton $\mathcal{A}$ with state set Q and variables X and a string s , computes a data structure in time $\mathcal{O}(|Q|^\omega \cdot |X|^2 \cdot |s|)$ and, then, given an index i retrieves, using the data structure, the i -th output of the evaluation of $\mathcal{A}$ over s in time $\mathcal{O}(|Q|^\omega \cdot |X|^3 \cdot \log(|s|)^2)$ where ω is the exponent for matrix multiplication. Ours is the first efficient direct access algorithm for MSO query evaluation over strings; such algorithms so far had only been studied for first-order queries and conjunctive queries over relational data.

Our algorithm gives the answers in lexicographic order where, in contrast to the setting of conjunctive queries, the order between variables can be freely chosen by the user without degrading the runtime. Moreover, our data structure can be updated efficiently after changes to the input string, allowing more powerful updates than in the enumeration literature, e.g. efficient deletion of substrings, concatenation and splitting of strings, and cut-and-paste operations. Our approach combines a matrix representation of MSO queries and a novel data structure for dynamic word problems over semi-groups which yields an overall algorithm that is elegant and easy to formulate.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Database theory

Keywords and phrases Query evaluation, direct access, MSO queries

Digital Object Identifier 10.4230/LIPIcs.ICDT.2025.26

Related Version *Full version with missing proofs*: <https://arxiv.org/abs/2409.17329>

Funding This work was partially funded by ANR projects CQFD 18-CE23-0003, EQUUS ANR-19-CE48-0019 and KCODA ANR-20-CE48-0004, by ANID - Millennium Science Initiative Program - Code ICN17_002, by ANID Fondecyt Regular project 1230935, and by the Alexander von Humboldt Foundation.

Acknowledgements This work benefited from Dagstuhl Seminar 24032, *Representation, Provenance, and Explanations in Database Theory and Logic*.

1 Introduction

The aim of direct access algorithms in query answering is to represent query answers in a compact and efficiently computable way without materializing them explicitly, while still allowing efficient access as if the answers were stored in an array. This is modeled in a



© Pierre Bourhis, Florent Capelli, Stefan Mengel, and Cristian Riveros;
licensed under Creative Commons License CC-BY 4.0

28th International Conference on Database Theory (ICDT 2025).

Editors: Sudeepa Roy and Ahmet Kara; Article No. 26; pp. 26:1–26:18

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

two-stage framework where, first, in a preprocessing phase, given an input, one computes a data structure which then, in the so-called access phase, allows efficiently accessing individual answers to the query by the position they would have in the materialized query result.

For this approach, there is a trade-off between the runtime of the preprocessing and the access phase: on the one hand, one could simply materialize the answers and then answer all queries in constant time; on the other hand, one could just not preprocess at all and then for every access request evaluate from scratch. When designing direct access algorithms, one thus tries to find interesting intermediate points between these extremes, mostly giving more importance to the access time than the preprocessing, since the latter has to be performed only once for every input while the former can be performed an arbitrary number of times.

Direct access algorithms were first introduced by Bagan, Durand, Grandjean, and Olive in [8] in the context of sampling and enumeration for first-order queries (both can easily be reduced to direct access). However, it was arguably the very influential thesis of Brault-Baron [14] that established direct access as an independent regime for query answering. Since then, there has been a large amount of work extending and generalizing Brault-Baron's work in several directions [19, 18, 15, 23, 17].

Curiously absent from this line of work is the query evaluation of Monadic Second Order logic (MSO), which in other contexts, like enumeration algorithms, has often been studied in parallel with conjunctive queries. In this paper, we make a first step in that direction, giving a direct access algorithm for MSO queries with free first-order variables over strings.

Our approach to direct access decomposes into several steps: first, we reduce direct access to a counting problem by using binary search to find the answers to be accessed as it is done, e.g., in [17, 18, 15]. We then express this counting problem in terms of matrix multiplication where we have to compute the result of a product of a linear number of small matrices. To enable the binary search we then require that this product can be efficiently maintained under substitutions of some of the matrices. This type of problem is known in the literature as *dynamic word problems* where the input is a word whose letters are elements from some semi-group whose product has to be maintained under element substitutions. Since we want to allow more powerful updates to the input later on, we cannot use the known data structures for dynamic word problems directly. We instead opt for a simpler approach that uses an extension of binary search trees to store our matrix product. The price of this is modest as it leads to a data structure that on only a doubly logarithmic factor slower than the provably optimal data structures for dynamic word problems. Plugging these ingredients together leads to a direct access algorithm with preprocessing time linear in the size of the input word and polylogarithmic access time. Moreover, if the query to evaluate is given in the form of an unambiguous automaton – which is well known to be equivalent in expressivity to MSO-queries [16] – then all runtime dependencies on the query are polynomial.

One advantage of our search tree based data structure is that, by relying heavily on known tree balancing techniques, it allows updating the input string efficiently without running the preprocessing from scratch. Such update operations have been studied before in the enumeration literature for first-order logic [12], conjunctive queries [28, 11, 12] and, most intensively for MSO on words and trees [10, 32, 36, 3, 35, 5, 29]. Compared to this latter line of work, we support more powerful updates: instead of single letter, resp. node, additions, and deletions as in most of these works, our data structure efficiently allows typical text operations like deletion of whole parts of the text, cut and paste, and concatenation of documents. The most similar update operations in the literature are those from [39] which beside our operations allow persistence of the data structure and thus, in contrast to us, also

copy and paste operations. Still, our updates are vastly more powerful than in all the rest of the literature. Moreover, we are the first to propose updates to direct access data structures; all works with updates mentioned above only deal with enumeration.

Another advantage of our algorithm is that it easily allows direct access to the query result for any lexicographic order. It has been observed that ranking the orders by a preference stated by the users is desirable. Thus, there are several works that tackle this question for direct access, enumeration, and related problems, see e.g. [15, 18, 9, 21, 20, 40, 41]. While the lexicographic orders we consider are more restricted, we still consider them very natural. Also, it is interesting to see that, while there is an unavoidable runtime price to pay for ranked access with respect to lexicographic orders to conjunctive queries [18, 15], we here get arbitrary lexicographic access for free by simply changing the order in which variables are treated during the binary search. As a consequence, we can even choose the order at access time as the preprocessing is completely agnostic to the order, similarly to what was shown for the ranked enumeration of conjunctive queries answers [20] which however does not consider a dynamic setting.

2 Preliminaries

Sets and strings. For a set A , we denote by 2^A the power set of A . For $n \in \mathbb{N}$ with $n \geq 1$, we denote by $[n]$ the set $\{1, \dots, n\}$. We use Σ to denote a finite alphabet and Σ^* all strings (also called words) with symbols in Σ . We will usually denote strings by s, s', s_i and similar. For every $s_1, s_2 \in \Sigma^*$, we write $s_1 \cdot s_2$ (or $s_1 s_2$ for short) for the concatenation of s_1 and s_2 . If $s = a_1 \dots a_n \in \Sigma^*$, we write $|s| = n$, that is, the length of s . We denote by $\epsilon \in \Sigma^*$ the string of 0 length (i.e., $|\epsilon| = 0$), also called the empty string. For $s = a_1 \dots a_n \in \Sigma^*$, we denote by $s[i, j]$ the substring $a_i \dots a_j$, by $s[..i]$ the prefix $a_1 \dots a_i$ and by $s[i..]$ the suffix $a_i \dots a_n$.

Mappings. Given a finite set X of variables and $n \in \mathbb{N}$, in this work we will heavily work with *mappings* of the form $\mu : X \rightarrow [n]$ as our outputs. We denote by $\text{dom}(\mu) = X$ the *domain* of μ and by $\text{range}(\mu) = \{i \in [n] \mid \exists x \in X : \mu(x) = i\}$ the *range* of μ . Further, we denote by $\mu_\emptyset$ the *empty mapping*, which is the unique mapping such that $\text{dom}(\mu) = \emptyset$. In our examples, we will usually write $(x_1 \mapsto i_1, \dots, x_\ell \mapsto i_\ell)$ to denote the mapping $\mu : \{x_1, \dots, x_\ell\} \rightarrow [n]$ such that $\mu(x_1) = i_1, \dots, \mu(x_\ell) = i_\ell$. We also write $\mu \cup (x \mapsto i)$ with $x \notin \text{dom}(\mu)$ for the mapping μ' such that $\mu'(x) = i$ and $\mu'(y) = \mu(y)$ for every $y \in \text{dom}(\mu)$.

Given a mapping $\mu : X \rightarrow [n]$, we define $\mu^{-1} : [n] \rightarrow 2^X$ as the *set-inverse mapping* of μ , defined by $\mu^{-1}(i) := \{x \in X \mid \mu(x) = i\}$. Note that $i \in \text{range}(\mu)$ if, and only if, $\mu^{-1}(i) \neq \emptyset$. Moreover, given a subset $Y \subseteq X$, we define the projection $\pi_Y(\mu)$ as the mapping $\mu' : Y \rightarrow [n]$ such that $\mu'(y) = \mu(y)$ for every $y \in Y$. Remark that if $Y = \emptyset$, then $\pi_Y(\mu) = \mu_\emptyset$.

Variable-set automata. A *variable-set automaton* [24, 25, 4] (or *vset automaton* for short) is a tuple $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$ where Q is a finite set of states, X is a finite set of variables, $q_0 \in Q$ is an initial state, $F \subseteq Q$ is a finite set of states, and $\Delta \subseteq Q \times \Sigma \times 2^X \times Q$ is the transition relation. Given a string $s = a_1 \dots a_n$, a run ρ of $\mathcal{A}$ over s is a sequence:

$$\rho := q_0 \xrightarrow{a_1/X_1} q_1 \xrightarrow{a_2/X_2} \dots \xrightarrow{a_n/X_n} q_n \quad (1)$$

such that $(q_{i-1}, a_i, X_i, q_{i+1}) \in \Delta$ for every $i \leq n$. We say that a run ρ like (1) is *valid* if, and only if, $\bigcup_{i=1}^n X_i = X$ and $X_i \cap X_j = \emptyset$ for every $i < j \leq n$; in other words, each variable in X appears exactly once in ρ . If ρ is valid, one can define the mapping $\mu_\rho : X \rightarrow [n]$ such

that, for every $x \in X$, $\mu_\rho(x) = i$ where i is the unique position that satisfies $x \in X_i$. As usual, we also say that a run ρ like (1) is *accepting* if, and only if, $q_n \in F$. Then we define the *output* of $\mathcal{A}$ over a string s as the set of mappings:

$$\llbracket \mathcal{A} \rrbracket(s) = \{\mu_\rho \mid \rho \text{ is a valid and accepting run of } \mathcal{A} \text{ over } s\}.$$

We define a *partial run* ρ of $\mathcal{A}$ as a sequence of transitions $\rho := p_0 \xrightarrow{b_1/Y_1} p_1 \xrightarrow{b_2/Y_2} \dots \xrightarrow{b_n/Y_n} p_n$ such that $(p_{i-1}, b_i, Y_i, p_i) \in \Delta$. We say that ρ is a partial run *from* p_0 *to* p_n *over the string* $b_1 \dots b_n$. Note that a run is also a partial run where we additionally assumed that $p_0 = q_0$. We say that a partial run is valid if $Y_i \cap Y_j = \emptyset$ for all $i \leq j \leq n$. We define the length of ρ as $|\rho| = n$, and we make the convention that a single state p_0 is a partial run of length 0. We define the set of variables $\text{vars}(\rho)$ of a partial run ρ as $\text{vars}(\rho) = \bigcup_{i=1}^n Y_i$. Given two partial runs $\rho = p_0 \xrightarrow{b_1/Y_1} \dots \xrightarrow{b_n/Y_n} p_n$ and $\sigma = r_0 \xrightarrow{c_1/Z_1} \dots \xrightarrow{c_m/Z_m} r_m$ such that $p_n = r_0$ we define the run $\rho \cdot \sigma$ as the concatenation of ρ and σ , i.e., the partial run $\rho \cdot \sigma := p_0 \xrightarrow{b_1/Y_1} \dots \xrightarrow{b_n/Y_n} p_n \xrightarrow{c_1/Z_1} \dots \xrightarrow{c_m/Z_m} r_m$. Note that $|\rho \cdot \sigma| = |\rho| + |\sigma|$ and $\text{vars}(\rho \cdot \sigma) = \text{vars}(\rho) \cup \text{vars}(\sigma)$.

We define the size $|\mathcal{A}|$ of a vset automaton $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$ as the number of states and transitions, so $|\mathcal{A}| := |Q| + |\Delta|$. In the following, we assume that all vset automata are *trimmed*, i.e., for every $q \in Q$ there exists a run ρ that reaches q , and there exists a partial run σ from q to some state in F . We can make this assumption without loss of generality, since removing unreachable states can be done in time $\mathcal{O}(|\mathcal{A}|)$ without modifying $\llbracket \mathcal{A} \rrbracket$.

MSO and vset automata. In this paper, we usually refer to *Monadic Second Order Logic (MSO)* as our language for query evaluation; however, we will not define it formally here since we will use vset automata as an equivalent model for MSO. Precisely, when we refer to MSO, we mean MSO formulas of the form $\varphi(x_1, \dots, x_n)$ over strings where $x_1, \dots, x_n$ are first-order open variables (see, e.g., [31] for a formal definition). Then, given a string s as a logical structure, the *MSO query evaluation problem* refers to computing the set $\llbracket \varphi(x_1, \dots, x_n) \rrbracket(s) := \{\mu : \{x_1, \dots, x_n\} \rightarrow [|s|] \mid s, \mu \models \varphi(x_1, \dots, x_n)\}$, which is the set of all first-order assignments (i.e., mappings) that satisfy φ over s . One can show that MSO over strings with first-order open variables is equally expressible to vset automata, basically, by following the same construction as for the Büchi-Elgot-Trakhtenbrot theorem [16] (see also [34]). Furthermore, vset automata (as defined here) are equally expressive to regular document spanners [24] for information extraction (see, e.g., [6, 34]). For this reason, we can use vset automata to define MSO queries, which also applies to the setting of document spanners, or any other query language equally expressive to MSO.

Functional and unambiguous vset automata. It is useful to consider vset automata that have no accepting runs that are not valid, so we make the following definition: we say that a vset automaton is *functional* if, and only if, for every $s \in \Sigma^*$, every accepting run ρ of $\mathcal{A}$ over s is also valid. In [24], it was shown that there is an algorithm that, given a vset automaton $\mathcal{A}$, constructs in exponential time a functional vset automaton $\mathcal{A}'$ of exponential size with respect to $\mathcal{A}$ such that $\llbracket \mathcal{A} \rrbracket = \llbracket \mathcal{A}' \rrbracket$. We will thus restrict our analysis and algorithms to functional vset automata, since we can extend them to non-functional vset automata, incurring an exponential blow-up. One useful property of functional vset automata is that each state determines the variables that are assigned before reaching it in the following sense.




Figure 1 A running example of a vset automaton $\mathcal{A}_0$ that will be used throughout this work.

► **Lemma 1.** Let $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$ be a functional vset automaton. For every $q \in Q$ there exists a set $X_q \subseteq X$ such that for every partial run ρ from q_0 to q it holds that $\text{vars}(\rho) = X_q$.

Proof (sketch). The lemma is standard in the literature for similar models, see e.g. [33]. We provide a short proof. By way of contradiction, suppose that there exists a state $q \in Q$ and two runs ρ and ρ' that reach q such that $\text{vars}(\rho) \neq \text{vars}(\rho')$. Given that $\mathcal{A}$ is trimmed, there exists a partial run σ starting from q that reaches some final state in F . Then the runs $\rho \cdot \sigma$ and $\rho' \cdot \sigma$ are accepting. Since $\mathcal{A}$ is functional, both runs are also valid and should satisfy $\text{vars}(\rho \cdot \sigma) = X = \text{vars}(\rho' \cdot \sigma)$. However, $\text{vars}(\rho \cdot \sigma) \neq \text{vars}(\rho' \cdot \sigma)$ which is a contradiction. ◀

By Lemma 1, for every functional vset automaton $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$ and every $q \in Q$ we define its sets of variables as X_q to be the set as in the lemma. One consequence of Lemma 1 is that for every transition $(q, a, Y, q') \in \Delta$, we have $Y = X_{q'} \setminus X_q$. In particular, there can only be $|\Sigma|$ transitions for every pair q, q' of states such that overall $|\Delta| = O(|Q|^2 \cdot |\Sigma|)$.

We will also restrict to *unambiguous* vset automata. We say that a vset automaton $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$ is *unambiguous* if for every string $s \in \Sigma^*$ and every $\mu \in \llbracket \mathcal{A} \rrbracket(s)$ there exists exactly one accepting run ρ of $\mathcal{A}$ over s such that $\mu = \mu_\rho$. It is well-known that for every vset automaton $\mathcal{A}$ one can construct an equivalent unambiguous functional vset automaton $\mathcal{A}'$ of exponential size with respect to $\mathcal{A}$. Therefore, we can restrict our analysis to the class of unambiguous functional vset automaton without losing expressive power. Both are standard assumptions in the literature for MSO query evaluation and have been used to find efficient enumeration algorithms for computing $\llbracket \mathcal{A} \rrbracket(s)$, see e.g. [25, 38, 34].

► **Example 2.** Figure 1a depicts an unambiguous functional vset automaton $\mathcal{A}_0$ on variables $\{x_1, x_2\}$ and alphabet $\{a, b\}$. It can be verified that $X_{q_0} = \emptyset$, $X_{q_1} = \{x_1\}$, $X_{q_2} = \{x_2\}$ and $X_{q_3} = \{x_1, x_2\}$. Let $s = abbab$. The following are all valid and accepting runs of $\mathcal{A}_0$ over s :

- $\rho_0 := q_0 \xrightarrow{a/\{x_1\}} q_1 \xrightarrow{b/\{x_2\}} q_2 \xrightarrow{b/\emptyset} q_2 \xrightarrow{a/\{x_1\}} q_3 \xrightarrow{b/\emptyset} q_3$ and $\mu_{\rho_0} = \{x_1 \mapsto 1, x_2 \mapsto 2\}$
- $\rho_1 := q_0 \xrightarrow{a/\emptyset} q_0 \xrightarrow{b/\{x_2\}} q_2 \xrightarrow{b/\emptyset} q_2 \xrightarrow{a/\{x_1\}} q_3 \xrightarrow{b/\emptyset} q_3$ and $\mu_{\rho_1} = \{x_1 \mapsto 4, x_2 \mapsto 2\}$
- $\rho_2 := q_0 \xrightarrow{a/\emptyset} q_0 \xrightarrow{b/\emptyset} q_0 \xrightarrow{b/\{x_2\}} q_2 \xrightarrow{a/\{x_1\}} q_3 \xrightarrow{b/\emptyset} q_3$ and $\mu_{\rho_2} = \{x_1 \mapsto 4, x_2 \mapsto 5\}$
- $\rho_3 := q_0 \xrightarrow{a/\emptyset} q_0 \xrightarrow{b/\emptyset} q_0 \xrightarrow{b/\{x_2\}} q_2 \xrightarrow{a/\{x_1\}} q_3 \xrightarrow{b/\emptyset} q_3$ and $\mu_{\rho_3} = \{x_1 \mapsto 4, x_2 \mapsto 3\}$

In Figure 1b, we show a summary of $\llbracket \mathcal{A}_0 \rrbracket(s)$. It can be verified that $\llbracket \mathcal{A}_0 \rrbracket(s)$ is the set of tuples μ such that $\mu(x_1)$ is a position where there is an a in s , $\mu(x_2)$ is a position where there

is a b in s . Moreover if $\mu(x_1) < \mu(x_2)$ then $\mu(x_2)$ is the first position after $\mu(x_1)$ containing a b . Otherwise, if $\mu(x_2) < \mu(x_1)$, then $\mu(x_1)$ is the first position after $\mu(x_2)$ containing an a .

Computational model. We work with the usual RAM model with unit costs and logarithmic size registers, see e.g. [27]. All data structures that we will encounter in our algorithms will be of polynomial size, so all addresses and pointers will fit into a constant number of registers. Also, all atomic memory operations like random memory access and following pointers can be performed in constant time. The numbers we consider will be of value at most $|s|^{|X|}$ where X is the variable set of the automaton we consider. Thus, all numbers can be stored in at most $|X|$ memory cells and arithmetic operations on them can be performed in time $O(|X|^2)$ naively¹. In the sequel, we use ω to denote the exponent for matrix multiplication.

3 Direct access for MSO queries

Let X be a set of variables. Given a total order $\prec$ on X , we extend $\prec$ to a *lexicographic order* on mappings as usual: for mappings $\mu, \mu' : X \rightarrow [n]$ we define $\mu \prec \mu'$ if, and only if, there exists $x \in X$ such that $\mu(x) < \mu'(x)$ and, for every $y \in X$, if $y \prec x$, then $\mu(y) = \mu'(y)$.

Fix a total order $\prec$ on X . Given a vset automaton $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$ and an input string s , consider the set of outputs $\llbracket \mathcal{A} \rrbracket(s) = \{\mu_1, \mu_2, \dots\}$ such that $\mu_1 \prec \mu_2 \prec \dots$. We define the *i -th output of $\mathcal{A}$ over s* , denoted by $\llbracket \mathcal{A} \rrbracket(s)[i]$, as the mapping μ_i . Intuitively, we see $\llbracket \mathcal{A} \rrbracket(s)$ as an array where the outputs are ordered by $\prec$ and we retrieve the i -th element of this array. The *direct access problem* for vset automata is the following: given a vset automaton $\mathcal{A}$, a string s , and an index i , compute the i -th output $\llbracket \mathcal{A} \rrbracket(s)[i]$; if the index i is larger than the number of solutions then return an out-of-bound error. Without loss of generality, in the following we always assume that $i \leq |\llbracket \mathcal{A} \rrbracket(s)|$, since $|\llbracket \mathcal{A} \rrbracket(s)|$ is easy to compute for unambiguous functional vset automata, see Section 4. As usual, we split the computation into two phases, called the *preprocessing* phase and the *access* phase:

Problem: MSODIRECTACCESS[$\prec$]	
Preprocessing:	$\left\{ \begin{array}{ll} \text{input:} & \text{a vset automaton } \mathcal{A} \text{ and } s \in \Sigma^* \\ \text{result:} & \text{a data structure } D_{\mathcal{A},s} \end{array} \right.$
Access:	$\left\{ \begin{array}{ll} \text{input:} & \text{an index } i \text{ and } D_{\mathcal{A},s} \\ \text{output:} & \text{the } i\text{-th output } \llbracket \mathcal{A} \rrbracket(s)[i] \end{array} \right.$

During the preprocessing phase, the algorithm receives as input a vset automaton $\mathcal{A}$ and a string $s \in \Sigma^*$ and computes a data structure $D_{\mathcal{A},s}$. After preprocessing, there is the access phase where the algorithm receives any index i as input and computes the i -th output $\llbracket \mathcal{A} \rrbracket(s)[i]$ by using the precomputed data structure $D_{\mathcal{A},s}$.

To measure the efficiency of a direct access algorithm, we say that an algorithm for the problem MSODIRECTACCESS[$\prec$] has *f -preprocessing* and *g -access time* for some functions f and g if, and only if, the running time of the preprocessing phase and the access phase is in $O(f(\mathcal{A}, s))$ and $O(g(\mathcal{A}, s))$, respectively. Note that the running time of the access phase does not depend on i , given that its bitsize is bounded by $|X| \cdot \log(|s|)$ so that it can be stored in $|X|$ registers in the RAM model and arithmetic on it can be done in time $O(|X|^2)$ naively.

Given a class $\mathcal{C}$ of vset automata, we define the problem MSODIRECTACCESS[$\prec$] for $\mathcal{C}$ as the problem above when we restrict $\mathcal{A}$ to $\mathcal{C}$. The main result of this work is the following.

¹ We remark that there are better algorithms for arithmetic operations, but we will not follow this direction here. If the reader is willing to use faster algorithms for multiplication and addition, then in all runtime bounds below the factor $|X|^2$ can be reduced accordingly.

► **Theorem 3.** *There is an algorithm that solves $\text{MSODIRECTACCESS}[\prec]$ for the class of unambiguous functional vset automata for any variable order $\prec$ with preprocessing $O(|Q|^\omega \cdot |X|^2 \cdot |s|)$ and access time $O(|Q|^\omega \cdot |X|^3 \cdot \log^2(|s|))$. These bounds remain even true when we assume that the order $\prec$ is only given as an additional input in the access phase.*

In Theorem 3, we restrict to the class of unambiguous vset automata. A natural question is what happens if we consider the class of all functional vset automata. We first remark that the *data complexity* of the problem will not change as it is possible to translate any functional vset automaton to an unambiguous one with an exponential blow-up. Unfortunately, in combined complexity, there is no polynomial time algorithm under standard assumptions.

► **Proposition 4.** *If $\text{MSODIRECTACCESS}[\prec]$ for the class of functional vset automata has an algorithm with polynomial time preprocessing and access, then $\text{SPANL} \subseteq \text{FP}$.*

SPANL is the class of functions computable as $|R|$, where R is the set of output values returned by the accepting paths of an NL machine, see [2] for a formal definition. In [2], it is shown that, if the functions in SPANL is in $\text{SPANL} \subseteq \text{FP}$, i.e., computable in polynomial time, then $\text{P} = \text{NP}$. Therefore, by Theorem 4 and [2] it is unlikely that there is an efficient algorithm for $\text{MSODIRECTACCESS}[\prec]$ for all functional vset automata in combined complexity.

In the remainder of this paper, we will prove Theorem 3. To make the proof more digestible, in Section 4 we show how to reduce the direct access problem to a counting problem and introducing a matrix approach for it. In Section 5, we will present the data structure constructed during the preprocessing and used in the access phase, showing Theorem 3. In Section 6, we then show how to integrate updates to the input string into our approach.

4 From direct access to counting: a matrix approach

In this section, we present the main algorithmic ideas for Theorem 3. In the next section, we use these ideas to develop a data structure for the preprocessing and access algorithm.

From direct access to counting. Let $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$ be a vset automaton and let $s = a_1 \dots a_n$ be a string of length n . Assume that $X = \{x_1, \dots, x_\ell\}$ has ℓ -variables ordered by a total order $\prec$, w.l.o.g., $x_1 \prec x_2 \prec \dots \prec x_\ell$. For an index $k \in [\ell]$ and a mapping $\tau : \{x_1, \dots, x_k\} \rightarrow [n]$, we define the set:

$$\llbracket \mathcal{A} \rrbracket(s, \tau) = \{\mu \in \llbracket \mathcal{A} \rrbracket(s) \mid \pi_{x_1, \dots, x_{k-1}}(\mu) = \pi_{x_1, \dots, x_{k-1}}(\tau) \wedge \mu(x_k) \leq \tau(x_k)\}.$$

Intuitively, the set $\llbracket \mathcal{A} \rrbracket(s, \tau)$ restricts the output set $\llbracket \mathcal{A} \rrbracket(s)$ to all outputs in $\llbracket \mathcal{A} \rrbracket(s)$ which coincide with τ for variables x_i before the variable x_k , and that have a position before $\tau(x_k)$ for the variable x_k . If $\tau = \mu_\emptyset$ is the empty mapping, we define $\llbracket \mathcal{A} \rrbracket(s, \tau) = \llbracket \mathcal{A} \rrbracket(s)$.

► **Example 5.** Going back to the example from Figure 1a with $s = \text{abbab}$ and $\tau = (x_1 \mapsto 2)$ then $\llbracket \mathcal{A}_0 \rrbracket(s, \tau)$ contains all tuples from $\llbracket \mathcal{A}_0 \rrbracket(s)$ which map x_1 to a position before 2. As seen in Figure 1b we have $\llbracket \mathcal{A}_0 \rrbracket(s, \tau) = \{(x_1 \mapsto 1, x_2 \mapsto 2)\}$. Now if $\tau = (x_1 \mapsto 2, x_2 \mapsto 3)$, we have $\llbracket \mathcal{A}_0 \rrbracket(s, \tau) = \emptyset$ since we only keep tuples where x_1 is set to 2 and no such tuple can be found. If $\tau = (x_1 \mapsto 4, x_2 \mapsto 3)$, we have $\llbracket \mathcal{A}_0 \rrbracket(s, \tau) = \{(x_1 \mapsto 4, x_2 \mapsto 2), (x_1 \mapsto 4, x_2 \mapsto 3)\}$.

For the sake of presentation, we denote by $\#\llbracket \mathcal{A} \rrbracket(s)$ (resp. $\#\llbracket \mathcal{A} \rrbracket(s, \tau)$) the number $|\llbracket \mathcal{A} \rrbracket(s)|$ (resp. $|\llbracket \mathcal{A} \rrbracket(s, \tau)|$) of outputs in $\llbracket \mathcal{A} \rrbracket(s)$ (resp. $\llbracket \mathcal{A} \rrbracket(s, \tau)$). For direct access, we are interested in finding efficient algorithms for computing $\#\llbracket \mathcal{A} \rrbracket(s, \tau)$ because of the following connection.

► **Lemma 6** (Lemma 7 in [17]). *If there is an algorithm that computes $\#[\mathcal{A}](s, \tau)$ in time T for every $k \in [\ell]$ and every $\tau : \{x_1, \dots, x_k\} \rightarrow [n]$, then there is an algorithm that retrieves $[\mathcal{A}](s)[i]$ in time $O(T \cdot \ell \cdot \log(n))$ for every index i .*

Proof (sketch). A similar proof can be found in [17]. For the convenience of the reader and because we will use a slightly more complicated variant later, we give a quick sketch here. Let $\mu_i = [\mathcal{A}](s)[i]$. The idea is to first compute $\mu_i(x_1)$ by observing the following: $\mu_i(x_1)$ is the smallest value $j_1 \in [n]$ such that $\#[\mathcal{A}](s, (x_1 \mapsto j_1)) \geq i$. This value can be found in time $O(T \cdot \log n)$ by performing a binary search on $\{\#[\mathcal{A}](s, (x_1 \mapsto j_1)) \mid j_1 \leq n\}$. Once we have found $\mu_i(x_1)$, we compute $\mu_i(x_2)$ similarly by observing that it is the smallest value j_2 such that $\#[\mathcal{A}](s, (x_1 \mapsto j_1, x_2 \mapsto j_2)) \geq i - \#[\mathcal{A}](s, (x_1 \mapsto j_1 - 1))$. The claim then follows by a simple induction. ◀

Given Lemma 6, in the following we concentrate our efforts on developing an index structure for efficiently computing $\#[\mathcal{A}](s, \tau)$ for every $k \in [\ell]$ and every $\tau : \{x_1, \dots, x_k\} \rightarrow [n]$.

A matrix representation for the counting problem. Let $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$ be an unambiguous functional vset automaton. For any letter $a \in \Sigma$, we define the matrix $M_a \in \mathbb{N}^{Q \times Q}$ such that for every $p, q \in Q$:

$$M_a[p, q] := \begin{cases} 1 & (p, a, S, q) \in \Delta \text{ for some } S \subseteq X \\ 0 & \text{otherwise.} \end{cases}$$

Strictly speaking, the matrix M_a depends on $\mathcal{A}$ and we should write $M_a^{\mathcal{A}}$; however, $\mathcal{A}$ will always be clear from the context and, thus, we omit $\mathcal{A}$ as a superscript from M_a . Since $\mathcal{A}$ is functional, for every pair of states $p, q \in Q$ and $a \in \Sigma$ there exists at most one transition $(p, a, S, q) \in \Delta$. Then M_a contains a 1 for exactly the pairs of states that have a transition with the letter a . For this reason, one can construct M_a in time $O(|Q|^2)$ for every $a \in \Sigma$. Figure 1c gives an example of matrices M_a and M_b for the automaton $\mathcal{A}_0$ from Example 2.

We can map strings to matrices in $\mathbb{N}^{Q \times Q}$ by homomorphically extending the mapping $a \mapsto M_a$ from letters to strings by mapping every string $s = a_1 \dots a_n$ to a matrix M_s defined by the product $M_s := M_{a_1} \cdot M_{a_2} \cdot \dots \cdot M_{a_n}$. For ϵ , we define $M_\epsilon = I$ where I is the identity matrix in $\mathbb{N}^{Q \times Q}$. Note that this forms an homomorphism from strings to matrices where $M_{s_1 s_2} = M_{s_1} \cdot M_{s_2}$ for every pair of strings $s_1, s_2 \in \Sigma^*$. It is easy to verify that for all states $p, q \in Q$ we have that $M_s[p, q]$ is the number of partial runs from p to q of $\mathcal{A}$ over s . Furthermore, if we define $\vec{q}_0$ as the (row) vector such that $\vec{q}_0[p] = 1$ if $p = q_0$ and 0, otherwise, and $\vec{F}$ as the (column) vector such that $\vec{F}[p] = 1$ if $p \in F$ and 0 otherwise, then we have the following equality between number of outputs and matrix products: $\#[\mathcal{A}](s) = \vec{q}_0 \cdot M_s \cdot \vec{F}$.

Our goal is to have a similar result for calculating $\#[\mathcal{A}](s, \tau)$ given a mapping $\tau : \{x_1, \dots, x_k\} \rightarrow [n]$. For this purpose, for every string $s = a_1 \dots a_n$, we define the matrix M_s^τ as follows. Recall that τ^{-1} is the set-inverse of τ . Then, for every $a \in \Sigma$ and $i \in [n]$, define the matrix $M_{a,i}^\tau \in \mathbb{N}^{Q \times Q}$ such that for every $p, q \in Q$:

$$M_{a,i}^\tau[p, q] = \begin{cases} 1 & \tau^{-1}(i) \setminus \{x_k\} \subseteq S \text{ for some } (p, a, S, q) \in \Delta \text{ and} & (1) \\ & \text{if } \tau(x_k) = i \text{ then } x_k \in X_q; & (2) \\ 0 & \text{otherwise.} \end{cases}$$

Finally, we define $M_s^\tau = M_{a_1,1}^\tau \cdot M_{a_2,2}^\tau \cdot \dots \cdot M_{a_n,n}^\tau$. Intuitively, Condition (1) for $M_{a,i}^\tau[p, q] = 1$ makes sure that all runs of $\mathcal{A}$ over s counted by M_s^τ must take transitions at position i that contain all variables x_j with $\tau(x_j) = i$ and $j < k$. Note that we remove x_k from $\tau^{-1}(i)$ since

x_k has a special meaning in τ . Indeed, Condition (2) for $M_{a,i}^\tau[p, q] = 1$ restricts x_k such that its assignment must be before or equal to position i . For this second condition, we exploit the set X_q of a functional vset automata, that gives us information of all variables that have been used until state q .

It is important to notice that if $i \notin \text{range}(\tau)$ then $M_{a,i}^\tau = M_{a,i}$. Since $\text{range}(\tau)$ has at most k elements, the sequences $M_{a_1}, \dots, M_{a_n}$ and $M_{a_1,1}^\tau, \dots, M_{a_n,n}^\tau$ differ in at most k matrices. In particular, if τ is the empty mapping, then $M_s^\tau = M_s$ as expected. Finally, similarly to M_a one can compute $M_{a,i}^\tau$ in time $O(|Q|^2)$ for any $a \in \Sigma$ assuming that we have τ and τ^{-1} .

► **Example 7.** Consider again the example from Figure 1a with $s = abbab$. Observe that for any τ , whenever $M_a[p, q] = 0$ then necessarily $M_{a,i}^\tau[p, q] = 0$. Hence in this example, we focus on example where the entry $M_a[p, q]$ goes from 1 to 0 after applying a partial mapping. We let $\tau = (x_1 \mapsto 4, x_2 \mapsto 4)$. We have that $M_{a,4}^\tau[q_0, q_1] = 0$ because even if there is a transition (q_0, a, S, q_1) with $x_1 \in S$, we do not have $x_2 \in X_{q_1}$. In other words, a run compatible with τ cannot take transition $(q_0, a, \{x_1\}, q_1)$ at position 4 since it does not allow to map x_2 to a position before 4. On the other hand, $M_{a,4}^\tau[q_2, q_3] = 1$ because it is possible to set x_1 to position 4 by taking transition $(q_2, a, \{x_1\}, q_3)$. Moreover, since $x_2 \in X_{q_3}$, it means that x_2 has been necessarily set by an earlier transition, hence at a position preceding 4. Finally $M_{a,4}^\tau[q_3, q_3] = 0$ because there is not transition from q_3 to q_3 setting variable x_1 .

Similarly to the relation between $\#[\mathcal{A}](s)$ and M_s , we can compute $\#[\mathcal{A}](s, \tau)$ by using M_s^τ as the following result shows.

► **Lemma 8.** *For every unambiguous functional vset automaton $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$, every string $s \in \Sigma^*$, and every mapping $\tau : \{x_1, \dots, x_k\} \rightarrow [n]$ it holds that: $\#[\mathcal{A}](s, \tau) = \vec{q}_0 \cdot M_s^\tau \cdot \vec{F}$.*

The algorithm. Now that we have defined the main technical tools, we can present the algorithm for the preprocessing and direct access of an unambiguous functional vset automaton $\mathcal{A}$ over a string $s = a_1 \dots a_n$. For this purpose, we will assume here the existence of a data structure for maintaining the product of matrices $\{M_{a_i,i}^\tau\}_{i \in [n]}$ to then present and analyze the algorithms. In Section 5, we will explain this data structure in full detail.

Suppose the existence of a data structure D that given values $n, m \in \mathbb{N}$ maintains n square matrices $M_1, \dots, M_n$ of size $m \times m$. This data structure supports three methods called **init**, **set**, and **out** which are defined as follows:

- $D \leftarrow \text{init}(M_1, \dots, M_n)$: receive as input the initial n matrices $M_1, \dots, M_n$ of size $m \times m$ and initialize a data structure D storing the matrices $M_1, \dots, M_n$;
 - $D' \leftarrow \text{set}(D, j, M)$: receive as input a data structure D , a position $j \leq n$, an $m \times m$ -matrix M , and output a data structure D' that is equivalent to D but its j -th matrix M_j is replaced by M ; and
 - $M \leftarrow \text{out}(D)$: receive as input a data structure D storing matrices $M_1, \dots, M_n$ and outputs a pointer to the matrix $M := M_1 \cdot \dots \cdot M_n$, i.e., the product of the n matrices.
- Further, we assume that method **set** is *persistent* [22], meaning that each call to **set** produces a new data structure D' without modifying the previous data structure D .

Assuming the existence of the data structure D , in Algorithm 1 we present all steps for the preprocessing and direct access of an unambiguous functional vset automaton $\mathcal{A}$ over a string $s = a_1 \dots a_n$. For the sake of presentation, we assume that $\mathcal{A}$, s , and the data structure D are available globally by all procedures.

Algorithm 1 uses all the tools developed in this section for solving the MSO direct access problem. The preprocessing (Algorithm 1, left) receives as input the vset automaton and the string and constructs the data structure D by calling the method **init** with matrices

■ **Algorithm 1** The preprocessing and direct access algorithms of an unambiguous functional vset automaton $\mathcal{A} = (Q, \Sigma, X, \Delta, q_0, F)$ with variables $x_1, \dots, x_\ell$ and string $s = a_1 \dots a_n$.

1: procedure PREPROCESSING($\mathcal{A}, s$)	15: procedure DIRECTACCESS(i)
2: $D \leftarrow \text{init}(M_{a_1}, \dots, M_{a_n})$	16: $\tau \leftarrow \mu_\emptyset$
3: procedure BINARYSEARCH(x, i, τ)	17: for $k = 1, \dots, \ell$ do
4: $L \leftarrow 0$	18: $j \leftarrow \text{BINARYSEARCH}(x_k, i, \tau)$
5: $R \leftarrow n$	19: $i \leftarrow i - \text{CALCULATEDIFF}(x_k, j, \tau)$
6: while $L \neq R$ do	20: $\tau \leftarrow \tau \cup (x_k \mapsto j)$
7: $j \leftarrow \lceil (L + R)/2 \rceil$	21: UPDATESTRUCT(x_{k+1}, j, τ)
8: $\tau' \leftarrow \tau \cup (x \mapsto j)$	22: return τ
9: $D' \leftarrow \text{set}(D, j, M_{a_j, j}^{\tau'})$	23: procedure CALCULATEDIFF(x, j, τ)
10: if $\vec{q}_0 \cdot \text{out}(D') \cdot \vec{F} \geq i$ then	24: $\tau_{\text{prev}} \leftarrow \tau \cup (x \mapsto j - 1)$
11: $R \leftarrow j$	25: $D_{\text{prev}} \leftarrow \text{set}(D, j, M_{a_{j-1}, j-1}^{\tau_{\text{prev}}})$
12: else	26: return $\vec{q}_0 \cdot \text{out}(D_{\text{prev}}) \cdot \vec{F}$
13: $L \leftarrow j$	27: procedure UPDATESTRUCT(x, j, τ)
14: return R	28: $\tau_{\text{next}} \leftarrow \tau \cup (x \mapsto j + 1)$
	29: $D \leftarrow \text{set}(D, j, M_{a_j, j}^{\tau_{\text{next}}})$

$M_{a_1}, \dots, M_{a_n}$. The direct access (Algorithm 1, right) receives any index i and outputs the i -th mapping τ in $\llbracket \mathcal{A} \rrbracket(s)$ by following the strategy of Lemma 6. Specifically, starting from an empty mapping τ (line 16), it finds the positions for variables $x_1, \dots, x_\ell$ by binary search for each variable x_k (line 16-17). After finding the value j for x_k , it decreases the index i by calculating the difference of outputs just before j (line 19), updates τ with $(x_k \mapsto j)$ (line 20), and updates D with the new value of x_k (line 21). We use the auxiliary procedures CALCULATEDIFF and UPDATESTRUCT to simplify the presentation of these steps. The workhorse of the direct access is the procedure BINARYSEARCH. It performs a standard binary search strategy for finding the value for variable x , by using the reduction of Lemma 6 to the counting problem and the matrix characterization $\# \llbracket \mathcal{A} \rrbracket(s, \tau)$ of Lemma 8.

The correctness of Algorithm 1 follows from Lemma 6 and 8. Regarding the running time, suppose that methods `init`, `set`, and `out` of the data structure D take time t_{init} , t_{set} , and t_{out} , respectively. For the preprocessing, one can check that the running time is $O(|Q|^2 \cdot |s| + t_{\text{init}})$ where it takes $O(|Q|^2 \cdot |s|)$ for creating the matrices $M_{a_1}, \dots, M_{a_n}$. For the direct access, one can check that it takes time $O(|X| \cdot \log(|s|) \cdot (t_{\text{set}} + t_{\text{out}}))$ for each index i .

The next section shows how to implement the data structure D and its methods `init`, `set`, and `out`. In particular, we show that the running time of these methods will be $t_{\text{init}} = O(|Q|^\omega \cdot |X|^2 \cdot |s|)$, $t_{\text{set}} = O(|Q|^\omega \cdot |X|^2 \cdot \log(|s|))$, and $t_{\text{out}} = O(1)$. Overall, the total running time of Algorithm 1 will be $O(|Q|^\omega \cdot |X|^2 \cdot |s|)$ for the preprocessing and $O(|Q|^\omega \cdot |X|^3 \cdot \log(|s|)^2)$ for each direct access as stated in Theorem 3.

5 Maintaining semi-groups products

In this section, we will show how to implement the data structure required in Section 4, which in particular has a `set` operation that allows to change some of the elements in a sequence of matrices while maintaining their product efficiently. Similar problems have been studied under the name of *dynamic word problems* in the literature, see e.g. [26, 7]. In that setting, one is given a sequence of elements from a semi-group and wants to maintain the

product of this sequence under substitution of the elements. Depending on the algebraic structure of the semi-group, there are algorithms of different efficiencies, and there is by now quite a good understanding, see again [26, 7].

We could likely use the results of [26] directly to define our data structure (though [26] assumes a finite semiring which does not correspond to our setting). However, later we will want to support more powerful changes to the strings than just substitutions, and it is not clear how the approach from [26] could be adapted for this. So we choose to use a less technically involved approach that allows for more powerful update operations while only losing a factor of $\log \log(n)$ compared to [26]. So fix a semi-group $\mathbb{G}$, i.e., a set $\mathbb{G}$ with an associative operation $\circ$. The semi-group that will be of most interest to us are the $(m \times m)$ -matrices with elements in $\mathbb{N}$ with the usual matrix multiplication, but our approach here works for any semi-group. As it is common, we often leave out the multiplication symbol $\circ$ and write $g_1 g_2$ for $g_1 \circ g_2$. We next introduce our data structure.


We will store sequences $\vec{g} = g_1, \dots, g_n$ over $\mathbb{G}$ in binary trees. To this end, let $T = (V, E)$ be a rooted binary tree in which the vertices are labeled with elements from $\mathbb{G}$. We assume that the children of each internal node of T are either a left or a right child. The label of a vertex v is denoted by $\text{label}(v)$. Remember that the in-order traversal of a rooted binary tree, when encountering a vertex v with left child v_1 and right child v_2 , first recursively visits the subtree of v_1 , then v , and finally recursively the subtree of v_2 . For every vertex v of T , we define the sequence of v , in symbols $\text{seq}(v)$, to be the sequence of elements in the subtree rooted in v read in in-order traversal. We also define the *product* of v , in symbols $\text{prod}(v)$ as the product of the elements in $\text{seq}(v)$ in the order of the sequence. Note that $\text{prod}(v)$ is a single element from $\mathbb{G}$ while $\text{seq}(v)$ is a sequence of elements from $\mathbb{G}$; we never store $\text{seq}(v)$ explicitly in the vertex v but store $\text{prod}(v)$. We also write $\text{seq}(T)$ for the sequence of the root of T and analogously define $\text{prod}(T)$.

In the remainder, we want to maintain $\text{prod}(T)$ under changes to $\text{seq}(T)$. To this end, it will be useful to access positions in $\text{seq}(T)$ efficiently. If the only changes were substitutions of elements as needed for Algorithm 1, we could easily do this by letting T be a binary search tree in which the keys are the position in $\text{seq}(T)$. However, we later also want to be able to delete, add, and move strings, and for these operations we would have to change a linear number of keys which is too expensive for our setting. Instead, we use the linear list representation with search-trees, see e.g. [30, Chapter 6.2.3]: we only store the *size* of the subtree of T rooted in v , which we denote by $\text{size}(v)$. Remark that $\text{size}(v)$ is also the length of $\text{seq}(v)$. The *size* information lets us locate the element at desired positions in $\text{seq}(v)$.

► **Observation 9.** *For every vertex v in T we can, given an index $i \in [|\text{seq}(v)|]$ and correct size-labels, decide in constant time if the entry of position i in $\text{seq}(v)$ is in the left subtree, the right subtree or in v itself. In the two former cases, we can also compute at which position it is in the contiguous subsequence of that subtree.*

Observation 9 directly gives an easy algorithm that descends from the root of T to any desired position in $\text{seq}(T)$ in time linear in the height of T .

As the final information, we also add, for each vertex v of T , the height of v , denoted by $\text{height}(v)$. This information is not needed in this section, but it will be useful in Section 6 for maintaining balanced trees, so we use it already here. Since the height of T will determine the runtime of our algorithm, we aim to bound it as for binary search trees. We say that a tree T satisfies the *AVL condition* if for every vertex v with children v_1, v_2 we have that $|\text{height}(v_1) - \text{height}(v_2)| \leq 1$ and if v has only one child v_1 , then v_1 is a leaf. The AVL condition allows us to bound the height of trees by the following classical result.



■ **Figure 2** AVL-product representation for the string *abbbab* and automaton $\mathcal{A}_0$ from Figure 1a.

► **Theorem 10** ([1]). *Every tree of size n that satisfies the AVL condition has height $O(\log(n))$.*

We now have everything to define our data structure: we call a tree T in which all vertices are labeled with elements from $\mathbb{G}$, which have correct attributes **size**, **height**, and **prod** as above and which satisfies the AVL-condition an *AVL-product representation of $\text{seq}(T)$* .

► **Example 11.** We consider again the automaton $\mathcal{A}_0$ from Figure 1a and string *abbbab* (this is a different string than the one considered in previous examples as a longer string makes illustration more relevant for the data structure). In this case, we show an AVL-product representation of $M_a M_b M_b M_b M_a M_b = M_{a,1}^{\mu_0} M_{b,2}^{\mu_0} M_{b,3}^{\mu_0} M_{b,4}^{\mu_0} M_{a,5}^{\mu_0} M_{b,6}^{\mu_0}$ in Figure 2. For convenience, despite this not being mandatory, we assume each vertex v of the tree also is identified by a unique number $\text{id}(v)$. In the example, we can verify that **prod** values are computed from the labels. For example, if we denote by v_i the vertex having $\text{id}(v_i) = i$, we have $\text{prod}(v_2) = \text{prod}(v_3) \cdot M_b \cdot \text{prod}(v_4)$ which recursively is equal to $M_a M_b M_b$. Similarly, $\text{prod}(v_1) = \text{prod}(v_2) \cdot \text{prod}(v_6)$, that is, $M_a M_b M_b M_b M_a M_b$, which is the product we want to maintain.

We now show that AVL-product representations are useful in our setting. We start by observing that out-queries are trivial since we store the necessary information **prod**(T) explicitly in the root.

► **Observation 12.** *Given an AVL-product representation T , one can return $\text{prod}(T)$ in constant time. So in particular, if $\text{seq}(T)$ is a sequence of matrices, AVL-product representations implement the method **out** in constant time.*

We next show that there is an efficient initialization algorithm.

► **Lemma 13.** *There is an algorithm that, given a sequence $\vec{g}$ over $\mathbb{G}$ of length n , computes in time $O(n)$ and with $O(n)$ semi-group products an AVL-product representation of $\vec{g}$.*

Specializing to matrices, we get the following for the **init**-operation of Algorithm 1.

► **Corollary 14.** *We can perform **init** for AVL-product representations in time $O(m^\omega n)$.*

Finally, we show that AVL-product representations can be used for the substitution updates that we need for direct access when using the **set** operation in a persistent manner.

► **Lemma 15.** *There is an algorithm that, given an AVL-product representation T , a position i and a semi-group element g , outputs a new AVL-product representation of the sequence we get from $\text{seq}(T)$ by substitution the element in position i by g . The runtime is $O(\log(|\text{seq}(T)|))$ plus the time to compute $O(\log(|\text{seq}(T)|))$ products in the semi-group.*

► **Corollary 16.** *The set-method from Algorithm 1 can be implemented in time $O(m^\omega \log(n))$ with AVL-product representations.*

► **Example 17.** Going back to the example from Figure 2. Assume now that one wants to update the data structure to only keep runs where x_1 is set to a position lower or equal than 2. To this end, one needs to compute $M_{a,1}^\tau M_{b,2}^\tau M_{b,3}^\tau M_{b,4}^\tau M_{a,5}^\tau M_{b,6}^\tau = M_{a,1}^{\mu_0} M_{b,2}^\tau M_{b,3}^{\mu_0} M_{b,4}^{\mu_0} M_{a,5}^{\mu_0} M_{b,6}^{\mu_0}$ where $\tau = (x_1 \mapsto 2)$. Hence, we only have to change one matrix in the original product.

To do so, going down in the tree using $\text{size}(\cdot)$, we can identify that the vertex where the second matrix is introduced in the product is vertex v_2 . Hence, we only have to update $\text{prod}(v_2)$ and $\text{prod}(v_1)$. To do so, observe that the new value M'_2 of $\text{prod}(v_2)$ is $\text{prod}(v_3) \cdot M_{b,2}^\tau \cdot \text{prod}(v_4)$, which can be obtained with two matrix products, $\text{prod}(v_3)$ and $\text{prod}(v_4)$ being already computed. Similarly, the new value of $\text{prod}(v_1)$ is $M'_2 \cdot M_b \cdot \text{prod}(v_6)$ which can also be computed with two extra matrix products, provided that we compute M'_2 first. Hence, in general, this update scheme can be performed with at most $2 \cdot \text{height}(T)$ matrix products.

6 Dynamic direct access under complex editing operations

In the previous sections, we have established that, given a string $s \in \Sigma^*$ of length n and an unambiguous functional automaton $\mathcal{A}$, we can answer direct access queries on the relation $\llbracket \mathcal{A} \rrbracket(s)$ in time $O(\log(n)^2)$ after $O(n)$ preprocessing. Now, assume that the string s is modified. If one wants to perform direct access on the relation induced by $\mathcal{A}$ on the updated string so far one has to perform the initialization of Lemma 13 again which will have $O(n)$ complexity for each editing step. In this section, we will show that the AVL-product representations from Section 5 can be updated efficiently so that we can perform direct access on the edited string without redoing the preprocessing from scratch.

In the following, we formalize the scenario of complex edits over strings. Then we define the problem of dynamic MSO direct access. We finish by showing how to easily extend our data structure and approach to give a solution for this scenario.

Editing programs over strings. Fix an alphabet Σ . Let $\text{SDB} = \{s_1, \dots, s_N\}$ be a set of strings over Σ called a *strings database*. For the sake of simplicity, we use $s_i \in \text{SDB}$ both as a string and as a *string name* (i.e., as a label referring to s_i that is considered of constant size). Let $\mathcal{S}$ be a set of *string variables* $S_1, S_2, \dots$ disjoint from SDB and Σ . A *literal* l is either a string name $s_i \in \text{SDB}$, a string variable $S \in \mathcal{S}$, or a symbol $a \in \Sigma$. An *editing rule* is a command of the following four types:

$$S := \text{concat}(l, l') \quad (S, S') := \text{split}(l, i) \quad (S, S') := \text{cut}(l, i, j) \quad S := \text{paste}(l, l', i)$$

where S and S' are string variables, l and l' are literals, and $i, j \in \mathbb{N}$. Intuitively, string variables will be assigned to strings and then concat will allow to concatenate two strings, split to split a string at position i , cut to extract the substring between positions i and j , and paste to insert a string inside another at position i . An *editing program* Π is a pair (S, P) where $S \in \mathcal{S}$ is the output string variable and P is a sequence of editing rules $R_1; R_2; \dots; R_n$ such that each string name appears at most once in the right-hand side of some rule, and each string variable appears at most once in the right-hand side of a rule R_i after it appears

in the left-hand side of a rule R_j with $j < i$. In other words, each string in the database can be used only once and each variable can be used at most once after it is defined. We define the size of an editing program Π as the number of rules $|\Pi| = n$.

► **Example 18.** Suppose that we have a strings database $\text{SDB}_0 = \{s_1, s_2\}$. We define the editing program $\Pi_0 = (S_6, P_0)$ with P_0 as the sequence:

$$(S_1, S_2) := \text{split}(s_1, 4); (S_3, S_4) := \text{split}(S_2, 1); S_5 := \text{concat}(S_1, a); S_6 := \text{concat}(S_5, S_4).$$

Next, we define the semantics of editing programs. A *string assignment* is a partial function $\sigma: (\text{SDB} \cup \Sigma \cup \mathcal{S}) \rightarrow \Sigma^*$ such that $\sigma(a) = a$ for every $a \in \Sigma$ and $\sigma(s) = s$ for every $s \in \text{SDB}$. We define the trivial assignment σ_{SDB} such that $\text{dom}(\sigma_{\text{SDB}}) = \text{SDB} \cup \Sigma$ (i.e., no variable is mapped). Given an assignment σ , a string variable S , and a string $s \in \Sigma^*$, we write $\sigma[S \mapsto s]$ to denote the assignment σ' that replaces S with s in σ (i.e., $\sigma'(S) = s$ and $\sigma'(S') = \sigma(S')$ for every $S' \in \text{dom}(\sigma) \setminus \{S\}$). Then, we define the semantics of rules as a function $\llbracket \cdot \rrbracket$ that maps assignments to assignments such that for every assignment σ :

$$\begin{aligned} \llbracket S := \text{concat}(l, l') \rrbracket(\sigma) &= \sigma[S \mapsto \sigma(l) \cdot \sigma(l')] \\ \llbracket (S, S') := \text{split}(l, i) \rrbracket(\sigma) &= \sigma[S \mapsto \sigma(l)[..i], S' \mapsto \sigma(l)[i+1..]] \\ \llbracket (S, S') := \text{cut}(l, i, j) \rrbracket(\sigma) &= \sigma[S \mapsto \sigma(l)[i, j], S' \mapsto \sigma(l)[..i-1] \cdot \sigma(l)[j+1..]] \\ \llbracket S := \text{paste}(l, l', i) \rrbracket(\sigma) &= \sigma[S \mapsto \sigma(l)[..i] \cdot \sigma(l') \cdot \sigma(l)[i+1..]] \end{aligned}$$

where we use the syntax $\sigma[S \mapsto s, S' \mapsto s']$ as a shorthand for $(\sigma[S \mapsto s])[S' \mapsto s']$. We extend the semantics to sequences of editing rules $R_1; \dots; R_n$ recursively as follows:

$$\llbracket R_1; \dots; R_{n-1}; R_n \rrbracket(\sigma) = \llbracket R_n \rrbracket(\llbracket R_1; \dots; R_{n-1} \rrbracket(\sigma)).$$

Finally, we define the *semantics of an editing program* $\Pi = (S, P)$ as the string $\llbracket \Pi \rrbracket(\text{SDB}) = \llbracket P \rrbracket(\sigma_{\text{SDB}})(S)$, namely, the string at S after evaluating P starting from the trivial assignment. We assume that the above semantics is undefined if any of the previous definitions are undefined like, for example, indices i, j are out of range, the variable S is not defined, etc. Note that one can easily check in time $O(|\Pi|)$ whether the semantics will be undefined or not. For this reason, we assume that the semantics of an editing program is always defined.

► **Example 19 (continue).** If we take $s_1 = \text{bbbcb}$, we get $\llbracket \Pi_0 \rrbracket(\text{SDB}) = \text{bbbab}$. In particular, one can easily check that the program Π_0 is updating the fifth letter of s_1 with a letter a .

Our editing programs can easily simulate the insertion, deletion, or update of a symbol in a string. Further, they can simulate the complex document editing operations presented in [39] like concat, extract, delete, and insert (see [39] for a definition of these operations). So, our set of operations can be seen as nearly equivalent to the one presented in [39]; the only operation that we disallow is copying. Indeed, we do not allow editing programs to make copies of the input strings or the string variables by definition. An editing program that makes copies could produce strings of exponential size, and the input index i for direct access will be of polynomial size with respect to the input documents, breaking our assumptions regarding the RAM model. For this reason, we disallow copying in our programs and leave this extension for future work.

Dynamic direct access of MSO. Now that the notion of editing programs is clear, we can define the problem of dynamic direct access as an extension of the $\text{MSODIRECTACCESS}[\prec]$ introduced in Section 3, adding a phase of edits over the strings database.

Problem: DYNAMICMSODIRECTACCESS[$\prec$]	
Preprocessing:	input: a vset automaton $\mathcal{A}$ and a strings database $\text{SDB} = \{s_1, \dots, s_N\}$
	result: a data structure $D_{\mathcal{A}, \text{SDB}}$
Editing:	input: an editing program Π and $D_{\mathcal{A}, \text{SDB}}$
	result: a data structure $D_{\Pi, \mathcal{A}, \text{SDB}}$
Access:	input: an index i and $D_{\Pi, \mathcal{A}, \text{SDB}}$
	output: the i -th output $\llbracket \mathcal{A} \rrbracket(\llbracket \Pi \rrbracket(\text{SDB}))[i]$

Contrary to the standard direct access, we add an intermediate phase, called editing, where one can receive an editing program Π over SDB as input. Thus, the direct access is now performed over the resulting string $\llbracket \Pi \rrbracket(\text{SDB})$. As before, we measure each phase separately and say that an algorithm for DYNAMICMSODIRECTACCESS[$\prec$] has f -preprocessing, h -editing, and g -access time for some functions f , h , and g if, and only if, the running time of the preprocessing, editing, and access phases is in $O(f(\mathcal{A}, \text{SDB}))$, $O(h(\mathcal{A}, \text{SDB}, \Pi))$, and $O(g(\mathcal{A}, \text{SDB}, \Pi))$, respectively.

By taking advantage of our techniques for MSO direct access, we can show that DYNAMICMSODIRECTACCESS[$\prec$] can be solved with logarithmic running time during the editing phase and without increasing the running time of the preprocessing and access phases.

► **Theorem 20.** *There is an algorithm that solves DYNAMICMSODIRECTACCESS[$\prec$], for the class of unambiguous functional vset automata for any variable order $\prec$ with preprocessing time $O(|Q|^\omega \cdot |X|^2 \cdot \sum_{i=1}^N |s_i|)$, editing time $O(|Q|^\omega \cdot |X|^2 \cdot |\Pi| \cdot \max_{i=1}^N \log(|s_i|))$, and access time $O(|Q|^\omega \cdot |X|^3 \cdot \log^2(|\Pi| \cdot \max_{i=1}^N |s_i|))$. These bounds remain even true when we assume that the order $\prec$ is only given as an additional input in the access phase.*

The proof goes by showing how to implement each editing rule receiving as inputs AVL-products representations. Precisely, the preprocessing constructs an AVL-product representation for each string s_i in SDB (e.g., like in Section 5). Then, during the editing phase we produce new AVL-product representations from the initial ones for each editing operation. Finally, the resulting representation for the string $\llbracket \Pi \rrbracket(\text{SDB})$ is used for the direct access phase. These give the desired running time for the preprocessing and the direct access phases. Therefore, it is only left to show how to implement each editing rule in time $O(|Q|^\omega \cdot |X|^2 \cdot \max_{i=1}^N \log(|s_i|))$ in order to prove the running time for the editing phase.

Implementing the editing rules. It is easy to see that $(S, S') := \text{cut}(l, i, j)$ can be implemented by splitting twice (at i and j) and one concatenation. More formally, as the sequence of editing rules $(S_1, S_2) := \text{split}(l, j)$; $(S_3, S) := \text{split}(S_1, i)$; $S' \leftarrow \text{concat}(S_3, S_2)$. Similarly, $S := \text{paste}(l, l', i)$ can be obtained by one split and two concatenations. More formally, $(S_1, S_2) := \text{split}(l, i)$; $S_3 := \text{concat}(S_1, l')$; $S := \text{concat}(S_3, S_2)$. Finally, we can easily replace a single symbol by a newly initialized data structure. For this reason, we dedicate this subsection to show how to implement **concat** and **split** efficiently.

Assume that the inputs of **concat** and **split** are given by AVL-product representations. We will heavily rely on the corresponding operations on AVL trees as presented in [13]: there is it shown that many operations on ordered sets can be implemented with AVL trees by using only an operation called **join** that does the following: given two AVL trees T_1, T_2 and a key k , such that the maximal key in T_1 is smaller than k and the minimal key in T_2 is bigger than k , **join** returns an AVL tree T containing all keys in T_1 , T_2 and k . Besides the usual

tree navigation, the only basic operation that `join` uses is `node`, which takes a key k and two trees T_1, T_2 and returns a tree that has k in its root and T_1 and T_2 as left and right subtrees, respectively². We directly get the following result.

► **Lemma 21.** *There is an algorithm `join` that, given two AVL-product representations T_1, T_2 and a semi-group element g , computes in time $O(\log(\text{size}(T_1) + \text{size}(T_2)))$ and the same number of semi-group operations an AVL-product representation T of the sequence we get by concatenating $\text{seq}(T_1)$, g , and $\text{seq}(T_2)$.*

Using `join`, one can implement an operation called `split` that takes as input an AVL tree and a key k and returns two AVL trees T_1 and T_2 such that all keys in T smaller than k are in T_1 while all keys bigger than k are in T_2 . Besides `join`, the only operation that `split` performs is a descent in the tree to find the node with key k . This directly gives the following algorithms for operations on $\text{seq}(T)$.

► **Lemma 22.** *There is an algorithm `split` that, given an AVL-product representation T and a position i , computes in time $O(\log(\text{size}(T)))$ and the same number of semi-group operations two trees AVL-product representations T_1, T_2 such that $\text{seq}(T_1)$ is the prefix of $\text{seq}(T)$ up to but excluding the i -th entry and $\text{seq}(T_2)$ is the suffix of $\text{seq}(i)$ that starts after the i -th position.*

Applying both lemmas for matrices, we directly get bounds for `concat` and `split`.

► **Corollary 23.** *There are algorithms for both $S := \text{concat}(l, l')$ and $(S, S') := \text{split}(l, i)$ on AVL-product representations that, given inputs T_1, T_2 , resp., T and i , run in time $O(m^\omega \cdot \log(|T_1| + |T_2|))$, resp., $O(m^\omega \cdot \log(|T|))$.*

7 Future work

We have given a direct access algorithm for MSO queries on strings that allows for powerful updates and supports all lexicographic orders without especially having to prepare them in the preprocessing. Our result is certainly only a first step in understanding direct access for MSO queries that motivate new research questions. First, it would be interesting to see to which extent one can adapt our results to MSO on trees. Second, one could study how to adjust our approach to support the copying of strings like in [39]. Third, another question is if one could reduce the access time of our algorithm. Finally, it would be interesting to understand if direct access on strings (or trees) is also possible for more expressive queries, say the grammar-based extraction formalisms of [37, 6].

References

- 1 Georgii M Adel'son-Vel'skii and Evgenii Landis. An algorithm for the organization of information. *Soviet Math.*, 3:1259–1263, 1962.
- 2 Carme Àlvarez and Birgit Jenner. A very hard log-space counting class. *Theor. Comput. Sci.*, 107(1):3–30, 1993. doi:10.1016/0304-3975(93)90252-0.
- 3 Antoine Amarilli, Pierre Bourhis, and Stefan Mengel. Enumeration on trees under relabelings. In *ICDT*, volume 98, pages 5:1–5:18, 2018. doi:10.4230/LIPIcs.ICDT.2018.5.
- 4 Antoine Amarilli, Pierre Bourhis, Stefan Mengel, and Matthias Niewerth. Constant-delay enumeration for nondeterministic document spanners. In *ICDT*, pages 22:1–22:19, 2019. doi:10.4230/LIPIcs.ICDT.2019.22.

² We remark that the usual rotations in search trees can be simulated by a constant number of `node`-operations.

- 5 Antoine Amarilli, Pierre Bourhis, Stefan Mengel, and Matthias Niewerth. Enumeration on trees with tractable combined complexity and efficient updates. In *PODS*, pages 89–103, 2019. doi:10.1145/3294052.3319702.
- 6 Antoine Amarilli, Louis Jachiet, Martin Muñoz, and Cristian Riveros. Efficient enumeration for annotated grammars. In *PODS*, pages 291–300, 2022. doi:10.1145/3517804.3526232.
- 7 Antoine Amarilli, Louis Jachiet, and Charles Paperman. Dynamic membership for regular languages. In *ICALP*, volume 198 of *LIPICs*, pages 116:1–116:17, 2021. doi:10.4230/LIPICs.ICALP.2021.116.
- 8 Guillaume Bagan, Arnaud Durand, Etienne Grandjean, and Frédéric Olive. Computing the j th solution of a first-order query. *RAIRO Theor. Informatics Appl.*, 42(1):147–164, 2008. doi:10.1051/ita:2007046.
- 9 Nurzhan Bakibayev, Tomás Kociský, Dan Olteanu, and Jakub Zavodny. Aggregation and ordering in factorised databases. *Proc. VLDB Endow.*, 6(14):1990–2001, 2013. doi:10.14778/2556549.2556579.
- 10 Andrey Balmin, Yannis Papakonstantinou, and Victor Vianu. Incremental validation of XML documents. *ACM Trans. Database Syst.*, 29(4):710–751, 2004. doi:10.1145/1042046.1042050.
- 11 Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering conjunctive queries under updates. In *PODS*, pages 303–318, 2017. doi:10.1145/3034786.3034789.
- 12 Christoph Berkholz, Jens Keppeler, and Nicole Schweikardt. Answering FO+MOD queries under updates on bounded degree databases. *ACM Trans. Database Syst.*, 43(2):7:1–7:32, 2018. doi:10.1145/3232056.
- 13 Guy E. Blelloch, Daniel Ferizovic, and Yihan Sun. Just join for parallel ordered sets. In *SPAA*, pages 253–264, 2016. doi:10.1145/2935764.2935768.
- 14 Johann Brault-Baron. *De la pertinence de l'énumération : complexité en logiques propositionnelle et du premier ordre. (The relevance of the list: propositional logic and complexity of the first order)*. PhD thesis, University of Caen Normandy, France, 2013. URL: <https://tel.archives-ouvertes.fr/tel-01081392>.
- 15 Karl Bringmann, Nofar Carmeli, and Stefan Mengel. Tight fine-grained bounds for direct access on join queries. In *PODS*, pages 427–436, 2022. doi:10.1145/3517804.3526234.
- 16 J Richard Büchi. Weak second-order arithmetic and finite automata. *Mathematical Logic Quarterly*, 6(1-6), 1960. doi:10.1002/malq.19600060105.
- 17 Florent Capelli and Oliver Irwin. Direct access for conjunctive queries with negations. In *ICDT*, volume 290, pages 13:1–13:20, 2024. doi:10.4230/LIPICs.ICDT.2024.13.
- 18 Nofar Carmeli, Nikolaos Tziavelis, Wolfgang Gatterbauer, Benny Kimelfeld, and Mirek Riedewald. Tractable orders for direct access to ranked answers of conjunctive queries. *ACM Trans. Database Syst.*, 48(1):1:1–1:45, 2023. doi:10.1145/3578517.
- 19 Nofar Carmeli, Shai Zeevi, Christoph Berkholz, Alessio Conte, Benny Kimelfeld, and Nicole Schweikardt. Answering (unions of) conjunctive queries using random access and random-order enumeration. *ACM Trans. Database Syst.*, 47(3):9:1–9:49, 2022. doi:10.1145/3531055.
- 20 Shaleen Deep, Xiao Hu, and Paraschos Koutris. Ranked enumeration of join queries with projections. *Proc. VLDB Endow.*, 15(5):1024–1037, 2022. doi:10.14778/3510397.3510401.
- 21 Johannes Doleschal, Benny Kimelfeld, Wim Martens, and Liat Peterfreund. Weight annotation in information extraction. *Log. Methods Comput. Sci.*, 18(1), 2022. doi:10.46298/lmcs-18(1:21)2022.
- 22 James R Driscoll, Neil Sarnak, Daniel Dominic Sleator, and Robert Endre Tarjan. Making data structures persistent. In *STOC*, pages 109–121, 1986. doi:10.1145/12130.12142.
- 23 Idan Eldar, Nofar Carmeli, and Benny Kimelfeld. Direct access for answers to conjunctive queries with aggregation. In *ICDT*, volume 290, pages 4:1–4:20, 2024. doi:10.4230/LIPICs.ICDT.2024.4.
- 24 Ronald Fagin, Benny Kimelfeld, Frederick Reiss, and Stijn Vansummen. Document spanners: A formal approach to information extraction. *Journal of the ACM (JACM)*, 62(2):1–51, 2015. doi:10.1145/2699442.

- 25 Fernando Florenzano, Cristian Riveros, Martín Ugarte, Stijn Vansummeren, and Domagoj Vrgoč. Efficient enumeration algorithms for regular document spanners. *ACM Transactions on Database Systems (TODS)*, 45(1):1–42, 2020. doi:10.1145/3351451.
- 26 Gudmund Skovbjerg Frandsen, Peter Bro Miltersen, and Sven Skyum. Dynamic word problems. *Journal of the ACM (JACM)*, 44(2):257–271, 1997. doi:10.1145/256303.256309.
- 27 Etienne Grandjean and Louis Jachiet. Which arithmetic operations can be performed in constant time in the RAM model with addition? *CoRR*, abs/2206.13851, 2022. doi:10.48550/arXiv.2206.13851.
- 28 Muhammad Idris, Martín Ugarte, and Stijn Vansummeren. The dynamic yannakakis algorithm: Compact and efficient query processing under updates. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1259–1274, 2017. doi:10.1145/3035918.3064027.
- 29 Sarah Kleest-Meißner, Jonas Marasus, and Matthias Niewerth. MSO queries on trees: Enumerating answers under updates using forest algebras. *CoRR*, abs/2208.04180, 2022. doi:10.48550/arXiv.2208.04180.
- 30 Donald Ervin Knuth. *The art of computer programming, , Volume III, 2nd Edition*. Addison-Wesley, 1998. URL: <https://www.worldcat.org/oclc/312994415>.
- 31 Leonid Libkin. *Elements of finite model theory*, volume 41. Springer, 2004. doi:10.1007/978-3-662-07003-1.
- 32 Katja Losemann and Wim Martens. MSO queries on trees: enumerating answers under updates. In *CSL-LICS*, pages 67:1–67:10, 2014. doi:10.1145/2603088.2603137.
- 33 Francisco Maturana, Cristian Riveros, and Domagoj Vrgoc. Document spanners for extracting incomplete information: Expressiveness and complexity. In *PODS*, pages 125–136, 2018. doi:10.1145/3196959.3196968.
- 34 Martin Muñoz and Cristian Riveros. Streaming enumeration on nested documents. In *ICDT*, volume 220 of *LIPICs*, pages 19:1–19:18, 2022. doi:10.4230/LIPICs.ICDT.2022.19.
- 35 Matthias Niewerth. MSO queries on trees: Enumerating answers under updates using forest algebras. In *LICS*, pages 769–778, 2018. doi:10.1145/3209108.3209144.
- 36 Matthias Niewerth and Luc Segoufin. Enumeration of MSO queries on strings with constant delay and logarithmic updates. In Jan Van den Bussche and Marcelo Arenas, editors, *PODS*, pages 179–191, 2018. doi:10.1145/3196959.3196961.
- 37 Liat Peterfreund. Grammars for document spanners. In Ke Yi and Zhewei Wei, editors, *ICDT*, volume 186, pages 7:1–7:18, 2021. doi:10.4230/LIPICs.ICDT.2021.7.
- 38 Markus L. Schmid and Nicole Schweikardt. Spanner evaluation over slp-compressed documents. In *PODS*, pages 153–165, 2021. doi:10.1145/3452021.3458325.
- 39 Markus L. Schmid and Nicole Schweikardt. Query evaluation over slp-represented document databases with complex document editing. In *PODS*, pages 79–89, 2022. doi:10.1145/3517804.3524158.
- 40 Nikolaos Tziavelis, Wolfgang Gatterbauer, and Mirek Riedewald. Optimal join algorithms meet top-k. In *SIGMOD*, pages 2659–2665, 2020. doi:10.1145/3318464.3383132.
- 41 Nikolaos Tziavelis, Wolfgang Gatterbauer, and Mirek Riedewald. Any-k algorithms for enumerating ranked answers to conjunctive queries. *CoRR*, abs/2205.05649, 2022. doi:10.48550/arXiv.2205.05649.