

Streamable Fragments of Forward XPath

Olivier Gauwin³ and Joachim Niehren^{1,2}

¹ Mostrare project, INRIA & LIFL (CNRS UMR8022)

² INRIA, Lille

³ University of Mons

Abstract. We present a query answering algorithm for a fragment of Forward XPath on XML streams that we obtain by compilation to deterministic nested word automata. Our algorithm is earliest and in polynomial time. This proves the finite streamability of the fragment of Forward XPath with child steps, outermost-descendant steps, label tests, negation, and conjunction (aka filters), under the reasonable assumption that the number of conjunctions is bounded. We also prove that finite streamability fails without this assumption except if $P=NP$.

Keywords : tree automata, pushdown automata, query answering, XML streams, XPath, temporal logics for unranked trees.

1 Introduction

Query answering algorithms for XPath on XML streams received much interest in the database and document processing communities [2,24,3,22,5,12,29,6,21] and are currently in the focus of the W3C working groups on XSLT and XPROC [14]. A little surprisingly, the topic is far from being settled given the large remaining gap between known streamable and non-streamable fragments. The objective of this paper is to narrow this gap by providing new positive and negative results for fragments of Forward XPath. Our approach relies on the relationship between temporal logics for unranked trees [16], which abstracts from the concrete syntax of XPath, and tree automata for XML streams [1,20,18].

Streaming is particularly relevant for data collections that are too large to be stored in main memory. Instead, incremental processing is needed in order to buffer only small parts of the data collection at every time point. In the easiest case, a stream is a word over some finite alphabet and a query selects some elements of this word, for instance all a -positions with two subsequent b 's. Usually, a query is considered streamable if there exists a one pass algorithm (see e.g. [26]) that computes the set of query answers with constant memory, independently of the input stream [28,27]. Note however, that streaming algorithms for element selection queries need to buffer all *alive* elements, i.e. those positions which might be selected in some continuations of the stream but not in others. In the above example, there exists at most one alive a -element at every time point, so this query can indeed be answered with bounded memory for all possible input streams.




Fig. 1: Sample XML document describing a bibliography.

The case of XML streams is similar except that they contain linearizations of unranked data trees and that queries select nodes in such trees. Consider for instance the XPath query $/bib/book[author="Ullman"]/author$ that selects all co-authors of Ullman (including himself) in all books of some bibliography (as illustrated in Fig. 1), or more precisely, all *author*-children of *book*-nodes that have at least one *author*-child with data value “Ullman”. An *author*-child of a *book*-node is alive, once the corresponding opening tag was seen on the stream, and as long as the closing *book* tag was not met and no *author*-node with data “Ullman” has been read. For bibliographies, in which all books have a bounded number of authors, the maximal number of alive nodes is bounded, so that the above query can be answered with bounded memory. For unusual bibliographies, however, the number of alive candidates may grow without any bound. As a consequence, the above query is not streamable in the usual sense even though it should be intuitively.

We propose the more liberal notion of *finite streamability* for languages of node selection queries on unranked trees. Finite streamability allows the memory to grow polynomially with the number of alive candidates, the size of the query, and the depth of the tree. In order to enable negative results, we assume in addition that the computation time per step is polynomial in the above parameters, and that the memory grows at least linearly with the number of alive candidates. The latter assumptions hold for all streaming algorithms without compression tricks for representing sets of alive candidates, an assumption that is satisfied by all streaming XPath algorithms in the literature so far.

An overview on finite streamability results for XPath fragments is given in Fig. 2. Despite of the intended weakness of this notion, only few positive results exist so far. Backward XPath (BXP) was proved finitely streamable based on transducers networks [5]. BXP queries never have any alive candidate since node selection is always determined at opening time. The second positive result [3] applies to FXP $(ch, o-ch_a^*, \wedge)^{thin}$, a thin fragment of positive Forward XPath on non-recursive documents, with star-restricted child steps, label-guarded (and thus outermost) descendants steps, and conjunctions (and thus filters in official XPath syntax). The only negative result so far got established for FXP (ch, ns^*, \wedge, \vee) , the fragment of positive Forward XPath with child and following-sibling axes,

	bounded number of \wedge	unbounded number of \wedge
BXP	yes [5]	yes [5]
FXP $(ch, o-ch_a^*, \wedge)^{thin}$	yes [3]	yes [3]
FXP (ch, \wedge, \neg)	yes	no
FXP $(ch, o-ch_a^*, \wedge, \neg)$	yes	no
FXP $(ch, o-ch_a^*, ns, \wedge, \neg)$?	no
FXP (ch, ns^*, \wedge, \vee)	?	no [5]
FXP (ch, ch^*, \wedge, \neg)	?	no

Colored results derive from the present paper. We assume here that $P \neq NP$.

Fig. 2: Finite streamability of fragments of XPath.

conjunction, and disjunction [5]. There, a counter example from online verification [19] was adapted in order to show for a family of queries in this fragment, that every streaming algorithm answering them must produce a doubly exponential number of states, and thus be of exponential size at least. This result applies even to Boolean queries (without node selection).

In this paper we study $\text{FXP}(ch, o-ch_a^*, \wedge, \neg)$, the fragment of Forward XPath with child axis, outermost descendant axis, conjunction, and negation. An outermost descendant axis $o-ch_a^*$ selects all a -descendants reachable via non- a -descendants. Outermost constraints on descendant steps are a natural restriction for streaming algorithms as noticed for instance in the XSLT 2.1 definition [15]. Our first main result is a streaming algorithm for $\text{FXP}(ch, o-ch_a^*, \wedge, \neg)$ that shows that this query language becomes finitely streamable if its queries are restricted to a bounded number of conjunctions. This result is relevant for the W3C pipeline language XPROC, for instance, where Forward XPath queries with at most 3 filters (and thus conjunctions) appear to be enough. Our second main result is the failure of finite streamability for $\text{FXP}(ch, \wedge, \neg)$ except if $P=NP$. It shows the necessity to bound the number of conjunctions theoretically.

We obtain our streaming algorithm by compiling $\text{FXP}(ch, o-ch_a^*, \wedge, \neg)$ to deterministic nested word automata (dNWAs) [1]. These are tree automata processing linearizations of unranked trees in preorder in a single pass, while mixing top-down and bottom-up determinism. For queries with a fixed number of conjunctions, our compiler is in polynomial time. Otherwise it is in exponential time, while still avoiding the usual doubly-exponential blow-up for translating XPath to deterministic automata [7]. Since the query language defined by dNWA is finitely streamable [11], the finite streamability follows for all fragments of $\text{FXP}(ch, o-ch_a^*, \wedge, \neg)$ with a bounded number of conjunctions.

Outline. Section 2 introduces FXP and Section 3 recalls dNWAs. In Section 4 we present our compiler from FXP to dNWAs. Section 5 introduces the notion of finite streamability and states our main results, positive and negative. Further related work is discussed in Section 6. The short CIAA version contains only sketches or ideas of proofs. Complete proofs are available in the long version [9].

$$\begin{array}{ll}
\llbracket F_1 \wedge F_2 \rrbracket_{t,\mu} = \llbracket F_1 \rrbracket_{t,\mu} \cap \llbracket F_2 \rrbracket_{t,\mu} & \llbracket d(F) \rrbracket_{t,\mu} = \{\pi \mid \exists \pi' \in \llbracket F \rrbracket_{t,\mu}. (\pi, \pi') \in d^t\} \\
\llbracket \neg F \rrbracket_{t,\mu} = \text{nod}(t) - \llbracket F \rrbracket_{t,\mu} & \llbracket a(F) \rrbracket_{t,\mu} = \{\pi \mid a = \text{lab}^t(\pi)\} \cap \llbracket F \rrbracket_{t,\mu} \\
\llbracket \text{true} \rrbracket_{t,\mu} = \text{nod}(t) & \llbracket x \rrbracket_{t,\mu} = \{\mu(x)\}
\end{array}$$

Fig. 3: Semantics of FXP(ch, ch^*, \wedge, \neg) formulas.

2 FXP

We present FXP temporal logics for unranked trees, which abstract from various aspects of the Forward XPath concrete syntax. More general temporal logics are reviewed by Libkin in [16] for instance (except for variables that we use for node selection here such as in hybrid logic).

For a finite label set Σ , we define the set of unranked trees \mathcal{T}_Σ to be the least set such that $a(t_1, \dots, t_k) \in \mathcal{T}_\Sigma$ if $a \in \Sigma$, $k \geq 0$ and $t_i \in \mathcal{T}_\Sigma$ for all $1 \leq i \leq k$. We write $\text{nod}(t)$ for the set of nodes of the tree t , ϵ for its root node, and $\text{lab}^t(\pi)$ for the label of node π of t . By ch^t and ch^{*t} we denote the child and descendant relations of t respectively. We will also use the outermost descendant relation $(o\text{-}ch_a^*)^t$ which navigates to all a -descendants reachable over non- a -descendants. A monadic node selection query Φ over Σ is a total function that maps trees $t \in \mathcal{T}_\Sigma$ to set of tuples of nodes $\Phi(t) \subseteq \text{nod}(t)$.

The temporal logic FXP($ch, o\text{-}ch_a^*, \wedge, \neg$) is a query language for node selection in unranked trees, in which one can talk about outermost a -descendants and children while using negation and conjunction. The expressions of this logic are terms with a single fixed free variable x (for the selecting position) over the ranked signature $\Delta = \{\wedge, \neg, \text{true}, x\} \cup \mathcal{D} \cup \Sigma$ where $\mathcal{D} = \{ch\} \cup \{o\text{-}ch_a^* \mid a \in \Sigma\}$. These terms have the following form where $d \in \mathcal{D}$ and $a \in \Sigma$.

$$F ::= F_1 \wedge F_2 \mid \neg F \mid \text{true} \mid d(F) \mid a(F) \mid x$$

FXP($ch, o\text{-}ch_a^*, \wedge, \neg$) corresponds to a natural class of Forward XPath expressions in the official XPath syntax modulo linear time transformations. The XPath expression $/ch^*::a[ch::b]/ch::*$ for instance becomes $ch^*(a(ch(x) \wedge ch(b(\text{true}))))$. Note that XPath filters are mapped to conjunctions in FXP.

Given a tree t and a variable assignment $\mu : \{x\} \rightarrow \text{nod}(t)$, we define a set valued semantics $\llbracket F \rrbracket_{t,\mu} \subseteq \text{nod}(t)$ for all formulas in Fig. 3. Path expression F defines the monadic query $\llbracket F \rrbracket$ that selects the following nodes for $t \in \mathcal{T}_\Sigma$:

$$\llbracket F \rrbracket(t) = \{\mu(x) \mid \epsilon \in \llbracket F \rrbracket_{t,\mu}, \mu : \{x\} \rightarrow \text{nod}(t)\}$$

The size $|F|$ is the usual size of term F and its (conjunction) *width* is the number of leaves in F .

Smaller fragments of FXP($ch, o\text{-}ch_a^*, \wedge, \neg$) can be obtained by removing some of the operators. For instance, we will write FXP(ch, \wedge, \neg) for the fragment using only the ch axis, conjunction and negation. The dialect of FXP(ch, ch^*, \wedge, \neg) is obtained by allowing for arbitrary descendant axis instead of only outermost a -descendants.

3 Deterministic Automata for XML Streams

We recall the notion of deterministic nested word automata (dNWAs) [1] following their presentation as streaming tree automata [8], and illustrate how to run them on XML streams. Similar kinds of tree automata were proposed for processing XML streams already in [20,18,17]. Note that these tree automata provide an explicit “visual” stack in contrast to standard tree automata.

XML streams are linearizations of unranked trees. The unranked tree $a(b, c)$ for instance becomes the XML stream $\langle a \rangle \langle b \rangle \langle /b \rangle \langle c \rangle \langle /c \rangle \langle /a \rangle$ where $\langle a \rangle$ is an opening tag and $\langle /a \rangle$ a closing tag. The events of the preorder traversal of a tree t are defined as follows (where **op** marks opening and **cl** closing events):

$$eve(t) = \{\mathbf{start}\} \cup (\{\mathbf{op}, \mathbf{cl}\} \times nod(t))$$

Hence, $eve(a(b, c)) = \{\mathbf{start}, (\mathbf{op}, \epsilon), (\mathbf{op}, \pi_1), (\mathbf{cl}, \pi_1), (\mathbf{op}, \pi_2), (\mathbf{cl}, \pi_2), (\mathbf{cl}, \epsilon)\}$, where π_i denotes here the i th child of the root. All events in $eve(t)$ except for **start** can be identified with a precise position in the XML stream for t . The event set is totally ordered with **start** as least element. We denote this order by \preceq and for an event $\eta \neq \mathbf{start}$ we write $pr_{\preceq}(\eta)$ for the immediately preceding event wrt. \preceq .

Definition 1. A dNWA is a tuple $(\Sigma, Q, \Gamma, i, F, \delta)$ where Σ is a finite alphabet, Q a finite set of states with a distinguished initial state $i \in Q$ and final states $F \subseteq Q$, Γ a finite set of stack symbols, and δ a set of rules. For each state $q_0 \in Q$ and letter $a \in \Sigma$, there is at most one rule $q_0 \xrightarrow{\mathbf{op} \ a: \gamma} q_1$ in δ , and for each $q_0 \in Q$, $a \in \Sigma$, and $\gamma \in \Gamma$, it contains at most one rule $q_0 \xrightarrow{\mathbf{cl} \ a: \gamma} q_1$.

A configuration of a dNWA A on a tree t consists of an event of t , a state of Q , and a stack of elements in Γ . An opening rule $q_0 \xrightarrow{\mathbf{op} \ a: \gamma} q_1$ can be applied to a configuration that opens some a -node in state q_0 . In this case, the subsequent configuration is reached by pushing γ to the current stack, changing the state to q_1 , and advancing to the next event. A closing rule $q_0 \xrightarrow{\mathbf{cl} \ a: \gamma} q_1$ can be applied to a configuration that closes some a -node in state q_0 . The symbol γ is then popped from the stack, the current state is changed to q_1 , and the current event is advanced by one. It should be noticed that transitions on configurations are always deterministic.

There is exactly one initial configuration: its event is **start**, its state i , and its stack is empty. Furthermore, note that the current stack is always the sequence of symbols that were pushed to the stack by the ancestors of the current node and itself. A configuration is accepting if the current event is the closing event of the root, the current state is final, and the current stack is empty.

More formally, a run r of an dNWA A on a tree t is a pair of functions $r_e : eve(t) \rightarrow Q$ and $r_n : nod(t) \rightarrow \Gamma$, such that $r_e(\mathbf{start}) = i$ and that δ contains the following rules for all $\pi \in nod(t)$ with $a = lab^t(\pi)$, $\alpha \in \{\mathbf{op}, \mathbf{cl}\}$ and $\eta = (\alpha, \pi)$:

$$r_e(pr_{\preceq}(\eta)) \xrightarrow{\alpha \ a: r_n(\pi)} r_e(\eta)$$




Fig. 4: A dNWA over $\Sigma = \{a, b\}$ with $Q = \{0, B, 1\}$ and $F = \{0, B\}$.




Fig. 5: Run of the dNWA of Fig. 4 on an input XML document.

A run r is successful if $r_e((c1, \epsilon)) \in F$. The recognized language $L(A)$ is the set of trees on which A has a successful run. We call an dNWA *pseudo-complete* if there is a run on every tree $t \in \mathcal{T}_\Sigma$.

For illustration, consider the dNWA in Fig. 4, which recognizes all trees containing some a -node with some b -child. This Boolean query is $ch^*(a(ch(b(true))))$ in FXP or $[//a/b]$ in XPath syntax. We will freely use the symbol $*$ to stand either for an arbitrary letter or an arbitrary stack symbol. The idea of this automaton is to move to state B when ever closing some b -node and to propagate this state by passing B to all closing events of following-siblings (except if some of them contains some a -descendant with some b -child, so that the automaton can safely go into the successful state 1). The automaton can move to the successful state 1 when closing some a -node from state B , since state B can only be assigned to closing events of children with a previous b -sibling. The run of this dNWA on tree $a(a, b, a)$ is illustrated in Fig. 5. Stack symbols can be either annotated to nodes of trees or to edges from opening to corresponding closing events on XML streams. The horizontal propagation of B works as follows: at opening time B is pushed onto the stack and at closing time it is popped from there.

In order to compute the run of a dNWA A on an XML stream with tree t , the current configuration of A needs to be stored at each event of t . This configuration contains the state of the current event and the sequence of states annotated to the ancestors of the current node, i.e., the current stack. Note that the size of the stack is at most $depth(t)$, so that membership to $L(A)$ can be decided by a streaming algorithm with a memory of size $O(|A| + depth(t))$.

Evaluation of dNWAs encoding DTDs or other XML schemas performs streaming schema validation. A weakness of naive evaluation for testing membership $t \in L(A)$ is the laziness of A in streaming mode: it only detects a -nodes with b -children when closing the a -node, but could already do so when opening the b -child. For tree $a(a, b, a)$ for instance, the earliest event is (op, π_2) when reading the first tag $\langle b \rangle$. The streaming algorithm from [11] improves on this situation: it decides membership $t \in L(A)$ for dNWAs A at the earliest possible event of tree t while remaining in PTIME. In order to find this earliest event, this algorithm needs to inspect the whole configuration at every event, not only the state.

Automata can also be used to define monadic queries. As before, we fix a variable x . For every tree $t \in \mathcal{T}_\Sigma$ and node $\pi \in \text{nod}(t)$, we define the *canonical tree* $t * \pi \in \mathcal{T}_{\Sigma \times 2^{\{x\}}}$ obtained from t by relabeling π with $(\text{lab}^t(\pi), \{x\})$ and all other nodes π' with $(\text{lab}^t(\pi'), \emptyset)$. More generally, a tree $t \in \mathcal{T}_{\Sigma \times 2^{\{x\}}}$ is *canonical* if exactly one of its nodes has a label in $\Sigma \times \{x\}$. A dNWA A with signature $\Sigma \times 2^{\{x\}}$ defines the query $\llbracket A \rrbracket$ on trees over Σ with $\llbracket A \rrbracket(t) = \{\pi \in \text{nod}(t) \mid t * \pi \in L(A)\}$.

4 FXP to Deterministic Automata

In this section, we propose a translation of $\text{FXP}(ch, o\text{-}ch_a^*, \wedge, \neg)$ to dNWAs. It runs in polynomial time if we assume a bound on the number of conjunctions. Our translation works by induction on the structure of formulas.

In order to avoid exponential blowups, our dNWAs will evaluate at most one subformula at every time point. Consider for instance the formula $ch(F')$. As all axes in F' are downwards (this would fail with the next-sibling axis), the algorithm can always know when closing a child, whether F' holds there or not. Thus, when opening the next child, the test for the previous child is finished. Therefore F' is tested for at most one child at a time. Note that an unbounded number of overlapping tests would end up in an exponential blowup. The same invariant also holds for $o\text{-}ch_a^*(F')$ formulas: no nested a -descendants need to be tested simultaneously for F' ; considering outermost a -descendants is enough.

Proposition 1. *For every formula F of $\text{FXP}(ch, o\text{-}ch_a^*, \wedge, \neg)$, we can build a dNWA A such that $\llbracket A \rrbracket = \llbracket F \rrbracket$ in time $O(|F|^{2 \cdot \text{width}(F)} \cdot |\Sigma|^{\text{width}(F)+1} \cdot 45^{\text{width}(F)})$.*

The automaton construction is by induction on the structure of formulas. Here we only highlight the main trick necessary that makes the construction polynomial when fixing $\text{width}(F)$. Conjunctions are mapped to automata intersection and negations to automata complementation, by swapping final states while assuming pseudo-complete dNWAs. Determinism is essential here. Note




Fig. 6: Successful run of the dNWA recognizing $F = ch(F')$.

that the compilation of conjunctions might produce dNWA of size exponential in $width(F)$. The translations of label tests and variables is straightforward.

The main point, where we avoid an important blow up, appears already in the construction of the automaton for $ch(F')$ and similarly for $o-ch_a^*(F')$. The idea of the dNWA for $ch(F')$ is to run the dNWA A' testing F' on every child of the root until finding one that satisfies F' . When running on the subtree rooted by some child, the algorithm must know when the child will be closed. In order to do so, it must push a special symbol to the stack when opening the child. It could do so by pushing a tagged version of the stack symbol γ pushed by A' . However, this would double the number of node states at each ch operator (as we also have to use γ below), leading to a global size increase of 2^n for formula $ch^n(true)$. The trick here, is to push a single new symbol 0, and to recompute node state γ corresponding to the current run due to determinism: knowing the initial state of A' and the label of the child, we can infer the rule of A' applied to open this child, and thus γ .

Let $A' = (\Sigma \times 2^{\{x\}}, Q', \Gamma', i', F', \delta')$ be the automaton built for F' . Automaton $A = (\Sigma \times 2^{\{x\}}, Q, \Gamma, i, F, \delta)$ for F will produce runs of the form in Fig. 6. It has three new states $Q = Q' \uplus \{\mathbf{start}, 0, 1\}$ and one additional stack symbol $\Gamma = \Gamma' \uplus \{0\}$.

1. State **start** is only used as initial state, to open the root node: $i = \{\mathbf{start}\}$ and a rule $\mathbf{start} \xrightarrow{\text{op } (a,V):0} 0$ is added to δ for all possible $(a, V) \in \Sigma \times 2^{\{x\}}$.
2. State 0 is used when closing a child of the root, if no matching for F' has been found so far. When a child is opened from 0, we start testing F' and assign node state 0 to this child. We have to add new rules, from rules starting from the initial state of A' (note that stack symbol γ are lost):

$$\frac{q_1 \in i' \quad q_1 \xrightarrow{\text{op } (a,V):\gamma} q_2 \in \delta'}{0 \xrightarrow{\text{op } (a,V):0} q_2 \in \delta}$$




Fig. 7: dNWA constructed for $ch(a(ch(b(x))))$ with $\Sigma = \{a, b\}$.

3. State 1 is universally accepting, so we always stay there once a matching has been found: $1 \xrightarrow{\alpha(a,V):0} 1 \in \delta$ for all $(\alpha, a, V) \in \{\text{op}, \text{c1}\} \times \Sigma \times 2^{\{x\}}$, and $F = \{1\}$.
4. Then a test of F' is launched: the set of new rules δ subsumes δ' .
5. When closing a child of the root, we have to check whether the test of F' succeeded or not. As argued before, A pushes state 0 when opening a child, so that stack symbol γ pushed by A' is lost temporarily. But A can recompute this symbol when closing the child. In case of success, A closes in state 1, otherwise in state 0.

$$\frac{q'_1 \in i' \quad q'_1 \xrightarrow{\text{op}(a,V):\gamma} q'_2 \in \delta' \quad q_1 \xrightarrow{\text{c1}(a,V):\gamma} q_2 \in \delta' \quad q_2 \in F'}{q_1 \xrightarrow{\text{c1}(a,V):0} 1 \in \delta}$$

$$\frac{q'_1 \in i' \quad q'_1 \xrightarrow{\text{op}(a,V):\gamma} q'_2 \in \delta' \quad q_1 \xrightarrow{\text{c1}(a,V):\gamma} q_2 \in \delta' \quad q_2 \notin F'}{q_1 \xrightarrow{\text{c1}(a,V):0} 0 \in \delta}$$

6. Finally, to remain pseudo-complete, we have to propagate state 0 when closing the root node: $0 \xrightarrow{\text{c1}(a,V):0} 0 \in \delta$ for all $(a, V) \in \Sigma \times 2^{\{x\}}$.

Even though the ideas of the constructions are rather simple, it should be noticed that dNWAs obtained by this construction are often hard to understand. This is mainly due to the recomputation trick. See Fig. 7 for an example.

5 Streamability of Query Languages

We present the notion of finite streamability of query languages, and apply it to the query languages defined by dNWA's and fragments of Forward XPath.

Definition 2. A monadic query language for unranked trees in \mathcal{T}_Σ is a triple $(E, \llbracket \cdot \rrbracket, |\cdot|)$ that consists of a set E whose elements are called query definitions, a function from definitions $e \in E$ to monadic query $\llbracket e \rrbracket$, that we call the query defined by e , and a mapping of query definitions $e \in E$ to natural numbers $|e| \in \mathbb{N}$, that we call the size of e .

How many candidates must be buffered when answering a query Φ on a tree t ? Intuitively, at least all alive candidates need to be stored, where a candidate $\pi \in \text{nod}(t)$ is called *alive* at an event $\eta \in \text{eve}(t)$ if it can be selected in some continuation of the stream and rejected in other ones. The concurrency $\text{concur}_\Phi(t)$ of Φ on t is the maximal number of alive candidates at all events.

The main idea of finite streamability is to require that the number of buffered candidates must be polynomially bounded in the concurrency. In order to do so, aliveness of some candidates must be decided at some point. Doing this in PTIME in the size of query definitions imposes a serious restriction, that all finitely streamable query languages must satisfy. In order to obtain lower bounds we assume that candidate sets are always stored without compression. This property is satisfied by all streaming algorithms in the literature.

Definition 3. We call a query language $(E, \llbracket \cdot \rrbracket, |\cdot|)$ finitely streamable if there exists polynomials p_0, p_1, p_2 such that for all query definitions $e \in E$ one can compute in time $p_0(|e|)$ a RAM machine \mathcal{M}_e computing $\llbracket e \rrbracket$, such that

- the space used by \mathcal{M}_e per step on $t \in \mathcal{T}_\Sigma$ is at most $p_1(|e|, \text{concur}_{\llbracket e \rrbracket}(t), \text{depth}(t))$ and at least $\text{concur}_{\llbracket e \rrbracket}(t)$, and
- the time used by \mathcal{M}_e per step on $t \in \mathcal{T}_\Sigma$ is at most $p_2(|e|, \text{concur}_{\llbracket e \rrbracket}(t), \text{depth}(t))$.

Prior work on earliest query answering provides our first positive result on streamability for dNWA's.

Theorem 1 ([11]). The language of monadic queries defined by dNWA's over $\Sigma \times 2^{\{x\}}$ is finitely streamable.

Proof. For monadic queries, the streaming algorithm in [11] has the following costs per step: $O(c \cdot |A|^2)$ in time and $O(c \cdot d \cdot |A|)$ in space, where $c = \text{concur}_{\llbracket A \rrbracket}(t)$ and $d = \text{depth}(t)$. This algorithm requires the dNWA A to accept only canonical trees, which can be obtained by intersecting it with a dNWA checking canonicity (this can be done in polynomial time). A RAM machine implementing this algorithm can be built in PTIME.

We define the query language $\text{FXP}(ch, o-ch_a^*, \wedge^{(k)}, \neg)$ which expressions are formulas F of $\text{FXP}(ch, o-ch_a^*, \wedge, \neg)$ with less than k conjunctions, i.e. such that $\text{width}(F) \leq k$. For this fragment, the translation provided in Section 4 is in polynomial time, and thus avoids more general doubly exponential compilation schemas of XPath expressions into deterministic tree automata [7].

Theorem 2. *For every fixed $k \geq 0$ and alphabet Σ , $\text{FXP}(ch, o\text{-}ch_a^*, \wedge^{(k)}, \neg)$ is finitely streamable.*

Proof. Let k be fixed. For every formula F in $\text{FXP}(ch, o\text{-}ch_a^*, \wedge^{(k)}, \neg)$, $\text{width}(F) \leq k$, so, according to the translation proposed in Section 4 (Proposition 1), there exists a polynomial p such that for all formulas F of $\text{FXP}(ch, o\text{-}ch_a^*, \wedge^{(k)}, \neg)$ we can build in time $O(p(|F|))$ a dNWA A such that $\llbracket A \rrbracket = \llbracket F \rrbracket$. Hence, finite streamability of queries by dNWA (Theorem 1) can be lifted to $\text{FXP}(ch, o\text{-}ch_a^*, \wedge^{(k)}, \neg)$.

The restriction on the width of formulas is necessary to remain in PTIME.

Theorem 3. *$\text{FXP}(ch, \wedge, \neg)$ is not finitely streamable, and remains non finitely streamable when restricted to non-recursive trees, unless $P = NP$.*

Here, we only give a brief sketch of the proof. We first show for all languages of descending queries that finite streamability implies that query satisfiability is in polynomial time. This can be shown by proving that aliveness of candidates must be decided for obtaining finite streamability, so that previous hardness results for earliest query answering carry over [6,11]. This works under the realistic assumption that the number of alive candidates is a space lower bound for streaming algorithms. We then show that satisfiability of $\text{FXP}(ch, \wedge, \neg)$ is NP-hard by strengthening results from [4]. Hence, without assuming $P=NP$ or a bound on the number of conjunctions, $\text{FXP}(ch, \wedge, \neg)$ cannot be finitely streamable, nor any larger query language.

6 Related Work

Our compiler from $\text{FXP}(ch, o\text{-}ch_a^*, \wedge^{(k)}, \neg)$ must avoid the usual doubly exponential blow-up when translating XPath expressions into deterministic tree automata [7]. One exponential goes away by bounding the number of conjunctions and all kinds of overlapping tests, for instance when adding ns or ns^* steps. The other exponential is circumvented by the restriction to outermost descendants steps since these can be checked deterministically.

As proved in the current paper, finite streamability of $\text{FXP}(ch, \wedge, \neg)$ continues to fail even if restricted to non-recursive documents. This shows that the memory consumption of the two algorithms of [2] and [12] cannot be polynomial in the number of alive candidates, in contrast to what is stated there⁴ except if $P=NP$. We also note that streaming algorithms for Forward XPath in [22] and [23,24] do not claim finite streamability. The complexity results stated there count the maximal number of candidates stored simultaneously by their algorithms, rather than the maximal number of alive candidates with respect to the query.

Space lower bounds for multi-pass streaming algorithms were shown in [13]. Previous space lower bounds for one-pass streaming algorithms for XPath were obtained by communication complexity arguments without any assumptions on

⁴ Authors of [2] and [12] have been notified. The journal version of [2] will take this remark into account.

compression tricks. Therefore, they remained limited to very specific fragments. In [2], wildcard-free queries in $\text{FXP}(ch, ch^*, \wedge, \neg)$ are considered under the assumption of an infinite signature. It is shown that the maximal number of closed simultaneously alive answer candidates is a lower bound for “mostly all” non-recursive trees in the sense of instance complexity. In [25], it is shown that for some queries in $\text{FXP}(ch, ch^*, \wedge)$ with independent ch predicates, the lower bound becomes $n \cdot c$ where n is the length of the selecting branch of the XPath expression, and c is maximal number of concurrently alive candidates. This shows that even compression tricks do not help for these query languages.

In [10] it was shown that it is decidable in polynomial time for queries defined by deterministic nested word automata, whether the maximal number of concurrently alive candidates is bounded. This result can be lifted to $\text{FXP}(ch, o\text{-}ch_a^*, \wedge^{(k)}, \neg)$ by using our P-time compiler to dNWAs.

7 Conclusion

We have shown that $\text{FXP}(ch, o\text{-}ch_a^*, \wedge, \neg)$ becomes finite streamability when fixing the number of conjunctions. Without such a bound, even $\text{FXP}(ch, \wedge, \neg)$ is not finitely streamable. Our results reveal some errors in previous work. This illustrates that they are nontrivial even though proofs are straightforward (once the translation is set up properly). It should also be noticed that our algorithm can be extended to support schemas (defined by DTDs or dNWAs) as well as for queries selecting tuples of nodes instead of nodes.

In QuiXProc (see www.quixproc.com), a transfer project of INRIA and INNOVIMAX, we are currently working on highly efficient streaming algorithms for $\text{FXP}(ch, o\text{-}ch_a^*, \wedge, \neg)$ based on similar dNWA constructions, which enable early node selection (not necessarily always earliest). First tests with our implementation, whose source code is freely available at fxp.lille.inria.fr, confirm this expectation. We are working on improving the integration of these algorithms into XPROC to industrial quality. We are thus confident to prove the practical relevance of the methods presented here in the near future.

References

1. R Alur and P Madhusudan. Adding nesting structure to words. *Journal of the ACM*, 56(3):1–43, 2009.
2. Z Bar-Yossef, M Fontoura, and V Josifovski. Buffering in query evaluation over XML streams. In *ACM PODS*, pages 216–227, 2005.
3. Z Bar-Yossef, M Fontoura, and V Josifovski. On the memory requirements of XPath evaluation over XML streams. *J. Comp. Syst. Sci.*, 73(3):391–441, 2007.
4. M Benedikt, W Fan, and F Geerts. XPath satisfiability in the presence of DTDs. *Journal of the ACM*, 55(2):1–79, 2008.
5. M Benedikt and A Jeffrey. Efficient and expressive tree filters. In *FST-TCS*, pages 461–472, 2007.
6. M Benedikt, A Jeffrey, and R Ley-Wild. Stream Firewalling of XML Constraints. In *ACM SIGMOD*, pages 487–498, 2008.

7. N Francis, C David, and L Libkin. A direct translation from XPath to nondeterministic automata. In *5th Alberto Mendelzon International Workshop on Foundations of Data Management*, 2011.
8. O Gauwin. *Streaming Tree Automata and XPath*. PhD thesis, Université Lille 1, 2009.
9. O Gauwin and J Niehren. Streamable fragments of Forward XPath. Long version. <http://hal.inria.fr/inria-00442250/en>. 2011.
10. O Gauwin, J Niehren, and S Tison. Bounded delay and concurrency for earliest query answering. In *LATA*, volume 5457 of *LNCS*, pages 350–361, 2009.
11. O Gauwin, J Niehren, and S Tison. Earliest query answering for deterministic nested word automata. In *FCT*, volume 5699 of *LNCS*, pages 121–132, 2009.
12. G Gou and R Chirkova. Efficient algorithms for evaluating XPath over streams. In *ACM SIGMOD*, pages 269–280, 2007.
13. M Grohe, C Koch, and N Schweikardt. Tight lower bounds for query processing on streaming and external memory data. In *ICALP*, *LNCS* 3580, 1076–1088, 2005.
14. M Kay. Saxon diaries, 2009. <http://saxonica.blogharbor.com/blog>.
15. M Kay. XSLT 2.1 – W3C working draft, May 2010.
16. L Libkin. Logics over unranked trees: an overview. *Logical Methods in Computer Science*, 3(2):1–31, 2006.
17. V Kumar, P Madhusudan, and M Viswanathan. Visibly pushdown automata for streaming XML. In *WWW*, pages 1053–1062, 2007.
18. C Koch, S Scherzinger, N Schweikardt, and B Stegmaier. Schema-based Scheduling of Event Processors and Buffer Minimization for Queries on Structured Data Streams. In *30th VLDB*, pages 228–239. Morgan Kaufmann, 2004.
19. O Kupferman and M Y. Vardi. Model checking of safety properties. *Form. Meth. in Syst. Design*, 19(3):291–314, 2001.
20. A Neumann and H Seidl. Locating matches of tree patterns in forests. In *FSTTCS*, volume 1530 of *LNCS*, pages 134–145, 1998.
21. A Nizar and S Kumar. Efficient Evaluation of Forward XPath Axes over XML Streams. In *COMAD*, pages 222–233, 2008.
22. D Olteanu. SPEX: Streamed and progressive evaluation of XPath. *IEEE Trans. on Know. Data Eng.*, 19(7):934–949, 2007.
23. P Ramanan. Evaluating an XPath Query on a Streaming XML Document. In *COMAD*, pages 41–52, 2005.
24. P Ramanan. Worst-case optimal algorithm for XPath evaluation over XML streams. *J. of Comp. Syst. Sci.*, 75:465–485, 2009.
25. P Ramanan. Memory lower bounds for XPath evaluation over XML streams. *J. of Comp. Syst. Sci.*, In Press, Corrected Proof:–, 2010.
26. N Schweikardt. Machine models and lower bounds for query processing. In *ACM PODS*, pages 41–52, 2007.
27. L Segoufin and C Sirangelo. Constant-memory validation of streaming XML documents against DTDs. In *ICDT*, pages 299–313, 2007.
28. L Segoufin and V Vianu. Validating streaming XML documents. In *ACM PODS*, pages 53–64, 2002.
29. X Wu and D Theodoratos. Evaluating Partial Tree-Pattern Queries on XML Streams. In *CIKM*, pages 1409–1410, 2008.