

# Exam

## Web Data Management

### Master Parisien de Recherche en Informatique

March 1st, 2019

This is the final exam for the Web Data Management class, which will determine 50% of your grade for this class (the other 50% being given by the project). The exam consists of 3 independent exercises + 1 bonus question. **You must write your answer to both exercise 3 and the bonus question on a *separate sheet of paper*.** You can choose to answer the questions in English or in French, as you prefer.

Write your name clearly on the top right of every sheet used for your exam answers, and number every page.

You are provided with an XPath “cheat sheet” which gives you a summary about the syntax of XPath (but which you probably shouldn’t take at face value; see the Bonus question at the end). You are additionally allowed **one A4 sheet** (i.e., two pages, one on each side) with the content of your choice. You may not use any other written material.

The exam is **strictly personal**: any communication or influence between students, or use of outside help, is prohibited. No electronic devices such as calculators, computers, or mobile phones, are permitted. Any violation of the rules may result in a grade of 0 and/or disciplinary action.

## Exercise 1: Compiling DTDs to tree automata (11 points)

In the entire exercise, we will consider XML documents that contain no textual content or entities, and have no document type declaration, comments, or processing instructions. The documents only use tag names in a finite fixed alphabet  $\Sigma$ . Further, each tag is either empty or contains exactly two child tags. In other words, the DOM representation of an XML document is always a full binary tree with node labels in  $\Sigma$ , which we call a  $\Sigma$ -tree. For now, we also assume that there are no attribute nodes (we will revisit this assumption further on).

Remember that a *DTD* can be used to describe the schema of an XML document. We will study a simplified version of the DTD language called *SimpleDTD* that only allows two kinds of declarations:

- First, declarations of the form:

`<!ELEMENT eltname (E)>`

where `eltname` is a name in  $\Sigma$ , and where  $E$  is a non-empty *deterministic* regular expression (regex) built with concatenation (`,`) disjunction (`|`), element names (from  $\Sigma$ ), and parentheses. This declaration specifies that the sequence of children of the element must satisfy the regex. Remember that a *deterministic* regex is one where, for every position in the regex, for every element name in  $\Sigma$ , the next position in the regex is uniquely defined.

- Second, declarations of the form:

`<!ELEMENT eltname EMPTY>`

where `eltname` is a name in  $\Sigma$ . This declaration specifies that the element must be empty.

A *SimpleDTD* is a sequence of such declarations with exactly one declaration overall per element name.

**Question 1 (1 point).** Consider the alphabet  $\Sigma_1 = \{a, b, c\}$  and the following constraints:

- elements labeled  $b$  and elements labeled  $c$  have no children
- elements labeled  $a$  have either two children labeled  $b$  or two children labeled  $c$ .

(a) Write a SimpleDTD  $\Delta_1$  that expresses these constraints.

(b) Now, we also allow elements labeled  $a$  to have one first child labeled  $b$  and a second child labeled  $c$ , in addition to what  $\Delta_1$  allows. Write a SimpleDTD  $\Delta'_1$  that expresses these constraints.

Our goal will be to translate DTD variants to tree automata. A *bottom-up deterministic finite tree automaton on  $\Sigma$ -trees* (or  $\Sigma$ -TA) is a tuple  $A = (Q, F, \iota, \delta)$  where  $Q$  is a finite set of *states*,  $F \subseteq Q$  is the set of *final states*,  $\iota : \Sigma \rightarrow Q$  is the *initialization function*, and  $\delta : Q \times Q \times \Sigma \rightarrow Q$  is the *transition function*. The *run* of  $A$  on a  $\Sigma$ -tree  $T$  is defined as the unique function  $\rho$  that maps each node of  $T$  to  $Q$  such that:

- for each leaf  $n$  of  $T$  with label  $l \in \Sigma$ , we have  $\rho(n) := \iota(l)$ ;
- for each internal node  $n$  of  $T$  with label  $l \in \Sigma$  and children  $n_1$  and  $n_2$ , we have  $\rho(n) := \delta(\rho(n_1), \rho(n_2), l)$ .

We say that the run  $\rho$  is *accepting* if  $\rho(r) \in F$  where  $r$  is the root of  $T$ , and we say that  $A$  *accepts*  $T$  if the run of  $A$  on  $T$  is accepting.

**Question 2 (2 points).** Give an algorithm which, given a SimpleDTD  $\Delta$  on an alphabet  $\Sigma$ , constructs a  $\Sigma$ -TA  $A_\Delta$  such that for every  $\Sigma$ -tree  $T$ , the automaton  $A_\Delta$  accepts  $T$  if and only if  $T$  satisfies  $\Delta$ . Argue for the correctness of your algorithm, and discuss its complexity as a function of  $\Sigma$  and  $\Delta$ .

**Question 3 (1 point).** For a fixed SimpleDTD  $\Delta$ , what is the complexity, given an input  $\Sigma$ -tree  $T$ , to determine whether  $T$  satisfies  $\Delta$ ? Justify your answer.

**Question 4 (1 point).** Consider the alphabet  $\Sigma_4 = \{a, b, c, d\}$  and the following constraint: “for every  $a$ -labeled node  $n$ , there must be both a  $b$ -labeled node and a  $c$ -labeled node that are descendants of  $n$ ”.

Write an XPath query that returns precisely the  $a$ -labeled nodes that do *not* satisfy the constraint. (You may refer to the provided XPath cheat sheet.)

**Question 5 (1 point).** Design a  $\Sigma_4$ -TA that accepts exactly the trees that satisfy the constraint of question 4, or prove that no such tree automaton exists.

**Question 6 (1 point).** Design a DTD that accepts exactly the trees that satisfy the constraint of question 4, or prove that no such DTD exists.

We now add *attributes* to our XML language. The finite tree alphabet  $\Sigma$  will now be partitioned into three kind of labels: *element names* (as previously), *attribute names*, and the two special labels *attr* and *child* (to be explained).

We keep our assumptions on XML documents from the beginning of the exercise, except elements may now have attributes, whose name must be an attribute name in  $\Sigma$ , and whose value is *always the empty string*. We recall that well-formed XML documents cannot have two attributes with the same name.

We study another simplified version of the DTD language called *AttrDTD* that only allows two kinds of declarations (the **ELEMENT** declarations are no longer allowed, but we implicitly assume that there are no constraints on the children nodes allowed on a given element):

- First, declarations of the form

```
<!ATTLIST eltname attname "" #IMPLIED>
```

which specifies that element `eltname` (an element name in  $\Sigma$ ) *may* have an attribute named `attname` (an attribute name in  $\Sigma$ ),

- Second, declarations of the form

```
<!ATTLIST eltname attrname "" #REQUIRED>
```

which specifies that element *eltname* (an element name in  $\Sigma$ ) *must* have an attribute named *attrname* (an attribute name in  $\Sigma$ ).


An *AttrDTD* is a sequence of such declarations where there is at most one declaration per pair of an element name and an attribute name. (However, each element name may have zero, one, or many declarations of either type.)

We now explain how XML documents with attributes are represented as  $\Sigma$ -trees. Each element node  $n$  of the XML document with label  $l$  (an element name in  $\Sigma$ ) is coded by a node  $n'$  with label  $l$  in the tree, which has two children whose respective labels are the special labels *attr* and *child*, and whose children are defined as follows:

- The *attr* node is a leaf if  $n$  has no attributes. Otherwise if  $n$  has  $k > 0$  attributes then the *attr* node is the root of a chain of  $k + 1$  *attr* nodes: for  $1 \leq i \leq k$ , the  $i$ -th *attr* node has as right child a node whose label is that of the  $i$ -th attribute (in the order in which they appear in the textual representation of the XML file) and has as left child the  $(i + 1)$ -th *attr* node; note that the  $(k + 1)$ -th *attr* node is a leaf.
- The children of the *child* node code the children of  $n$  (and if  $n$  has no children then the *child* node is a leaf).

Here is an example on the alphabet  $\Sigma = \{a, b, c, x, y, attr, child\}$ :

```
<a x="">
  <b></b>
  <c x="" y=""></c>
</a>
```



We now denote by  $\Sigma$ -tree the trees whose nodes carry a label in  $\Sigma$  and which have been constructed in this way.

**Question 7 (2 points).** Give an algorithm which, given an *AttrDTD*  $\Delta$  on an alphabet  $\Sigma$ , constructs a  $\Sigma$ -TA  $A_\Delta$  such that for every  $\Sigma$ -tree  $T$ , the automaton  $A_\Delta$  accepts  $T$  if and only if  $T$  satisfies the DTD. Argue for the correctness of your algorithm, and discuss its complexity as a function of  $\Sigma$  and  $\Delta$ .

**Question 8 (1 point).** Assume now that we want to extend our tree representation to XML documents where attributes can carry values. We want to allow arbitrary values to appear in the XML document, however we still want the tree representation to be on a fixed alphabet (i.e.,  $\Sigma$  should be finite, it should not contain infinitely many possible attribute values). We also want trees to remain binary and full. Present a possible way to extend our tree model to represent attribute values subject to these constraints.

**Question 9 (1 point).** We extend the *AttrDTD* language to add attribute declarations of the form:

```
<!ATTLIST eltname attrname ID #REQUIRED>
```

which defines a mandatory attribute *attrname* for element *eltname* which must be an XML ID (i.e., it must consist of only lowercase and uppercase ASCII letters, digits, '-', '.', ':', or '\_', and must start by a letter, by ':', or by '\_') and must be unique (i.e., the value used for an occurrence of an attribute declared with ID cannot be used as the value of another occurrence of an attribute declared with ID).

Explain how to extend the algorithm of question 7 to produce automata (with your extended encoding in question 8) that can handle *AttrDTD* featuring such declarations, or explain why it is not possible to do so.

## Exercise 2: Translating XPath to SPARQL (4 points)

The goal of this open-ended exercise is to convert XPath queries on trees to SPARQL queries on RDF graphs. For brevity, all URIs will be in an *m* namespace; you do not need to worry about how it is declared. You may refer to the provided cheat sheet for information about the syntax of XPath. For SPARQL, recall the following important features of the language:

- Basic queries: `SELECT ?x WHERE { ... }`
- The body contains statements of the form `subject predicate object` . where the subject, predicate and object may be URIs (e.g., `<m:42>`) or variables (e.g., `?foo`). The object may also be a literal (e.g., `"42"`).
- The predicate can also be a *property path*, i.e., a regular expression on paths, with the most common operators being disjunction (`|`), concatenation (`/`), Kleene star (`*`), one or more occurrences (`+`), and the inverse relation (e.g., `^<m:foo>` is the inverse relation of `<m:foo>`).
- The `{ ... } UNION { ... }` operator can be used to union the results of two queries.
- The `FILTER EXISTS { ... }` operator can be used to filter the preceding results based on the existence of matches to a subquery. There is also `FILTER NOT EXISTS`.
- We can add `OFFSET x LIMIT y` at the end of a `SELECT` query to return only *y* results starting at the *x*-th result.
- There is support for aggregates, e.g., `SELECT ?x (COUNT(?y) AS ?count) WHERE { ... }`.

We work with XML documents having no document type declaration, no comments, no processing instructions, no entities, no attributes, and no textual content. To convert an XML tree to an RDF graph, we generate one numerical identifier per element in the XML document order (starting at 0), and put them in the *m* namespace. We use three special relation names:

- *m:1*, where the subjects are the nodes of the tree, and the object for each subject is a literal containing the node label;
- *m:c*, where the subjects are the internal nodes of the XML tree, and the objects for a given subject are the children of that node;
- *m:n*, where the subjects are the nodes of the XML tree that are not the last child of their parent, and the object for a given subject is the next sibling of that node.

Here is an example, where for brevity we write the list of RDF facts in three columns:

XML document	XML tree	RDF representation (NTriples format)								
<code>&lt;a&gt;</code>	<pre>       a      /  \     b  c  a        / \       b  d           </pre>	<code>&lt;m:0&gt;</code>	<code>&lt;m:1&gt;</code>	<code>"a"</code> .	<code>&lt;m:0&gt;</code>	<code>&lt;m:c&gt;</code>	<code>&lt;m:1&gt;</code> .	<code>&lt;m:1&gt;</code>	<code>&lt;m:n&gt;</code>	<code>&lt;m:2&gt;</code> .
<code>&lt;b&gt;&lt;/b&gt;</code>		<code>&lt;m:1&gt;</code>	<code>&lt;m:1&gt;</code>	<code>"b"</code> .	<code>&lt;m:0&gt;</code>	<code>&lt;m:c&gt;</code>	<code>&lt;m:2&gt;</code> .	<code>&lt;m:2&gt;</code>	<code>&lt;m:n&gt;</code>	<code>&lt;m:5&gt;</code> .
<code>&lt;c&gt;</code>		<code>&lt;m:2&gt;</code>	<code>&lt;m:1&gt;</code>	<code>"c"</code> .	<code>&lt;m:0&gt;</code>	<code>&lt;m:c&gt;</code>	<code>&lt;m:5&gt;</code> .	<code>&lt;m:3&gt;</code>	<code>&lt;m:n&gt;</code>	<code>&lt;m:4&gt;</code> .
<code>&lt;b&gt;&lt;/b&gt;</code>		<code>&lt;m:3&gt;</code>	<code>&lt;m:1&gt;</code>	<code>"b"</code> .	<code>&lt;m:2&gt;</code>	<code>&lt;m:c&gt;</code>	<code>&lt;m:3&gt;</code> .			
<code>&lt;d&gt;&lt;/d&gt;</code>		<code>&lt;m:4&gt;</code>	<code>&lt;m:1&gt;</code>	<code>"d"</code> .	<code>&lt;m:2&gt;</code>	<code>&lt;m:c&gt;</code>	<code>&lt;m:4&gt;</code> .			
<code>&lt;/c&gt;</code>		<code>&lt;m:5&gt;</code>	<code>&lt;m:1&gt;</code>	<code>"a"</code> .						
<code>&lt;a&gt;&lt;/a&gt;</code>										
<code>&lt;/a&gt;</code>										

**Question 1 (1 point).** Translate the XPath query

`a//b[../c]`

to an equivalent SPARQL query of the form `SELECT ?x WHERE { ... }`, i.e., the RDF nodes *?x* returned by the query should be exactly the RDF nodes corresponding to the tree nodes returned by the XPath query.

**Question 2 (3 points).** Explain more generally how to translate XPath queries to SPARQL. In particular, you should mention the following features (the list is not exhaustive):

- Finding the root of the document
- Axes such as `descendant`, `following-sibling`, or `following`
- Predicates, nested predicates
- `position()`, `last()`, `count()`

*Please write your answer to Exercise 3 and the bonus question on a separate sheet of paper.*

### Exercise 3: The Deep Web (5 points)

*This exercise is fairly open ended. No fixed answer is expected, but you will be judged on the originality, feasibility, and soundness of your proposals.*

Sources from the deep Web (Web forms, Web services) typically expose some relational data, as is traditional in the data integration setting, with a twist: it is not possible to access the data without providing some input to the form or service. For example, a directory service such as *PagesBlanches* may expose relational data of the form `Contact(Last,First,City,Phone)` but only if a form is submitted with input consisting of both a `Last` name and a `City`. This is called an *access method* and it is denoted:

`Contactioio(Last,First,City,Phone)`

where ‘i’ stands for *input* and ‘o’ for *output*: if both a `Last` name and a `City` are provided as *input*, a list of matching records will be provided, including the value of the `First` name and `Phone` number for each record.

There are two common approaches to exploiting deep Web data:

- the *surfacing* approach (also called *extensional*, *siphoning*) where one aims to retrieve and store *all* content hidden behind the deep Web interface, in the spirit of data warehousing;
- the *intensional* approach where one aims to rewrite user queries at runtime to make use of the deep Web query interfaces, in the spirit of the mediator approach to data integration.

**Question 1 (2 points).** Assume we want to follow the surfacing approach. Propose a pragmatic strategy to retrieve as much content as possible from a directory deep Web source featuring the access method above. You need to describe in detail the architecture of a system that siphons the deep Web source: what content is submitted to the Web form? where to find input data? when to stop? how to assess how much data is left?

We now focus on the intensional approach. Assume we have a bunch of relational sources, each with one or several access methods that constrain access to the data. We also assume we have a Local-As-View description of each source as a conjunctive query over a global, mediator schema. Now consider a user query  $Q$ .

**Question 2 (1 point).** Give an example with at least two different sources where  $Q$  could be rewritten using at least two local sources with standard data integration techniques if all access methods had no input attribute, but where such a rewriting is not possible when access methods have input attributes.

**Question 3 (2 points).** Propose an approach to decide whether there exists a rewriting of  $Q$  in terms of local sources, that respects the access methods. Can standard data integration techniques, such as the Bucket or Inverse Rules algorithm, be adapted to this setting?

### Bonus Question (+2 points)

The XPath cheat sheet provided was found on the Internet, and as most things found on the Internet, it contains (at least) one major error. This is an error with respect to the semantics of XPath. You get 2 bonus points if you identify this error and provide an example XML document and query that illustrate why this is an error.