XML-
XPath and XQuery

Introduction to Databases

CompSci 316 Spring 2017

DUKE
COMPUTER SCIENCE

2

Announcements (Wed., Apr. 12)

- Homework #4 due Monday, April 24, 11:55 pm
  - 4.1, 4.2, 4.3, X1 is posted
  - Please start early
  - There may be another extra credit problem
- Projects
  - keep working on them and write your final report
  - Demo in the week of April 24
- Google Cloud use?
  - If anyone is planning to use or using google cloud for HW/project, please send me an email

3

Relational Mapping

4

Approaches to XML processing

- Text files/messages
- Specialized XML DBMS
  - Tamino (Software AG), BaseX, eXist, Sedna, …
  - Not as mature as relational DBMS
- Relational (and object-relational) DBMS
  - Middleware and/or extensions
  - IBM DB2's pureXML, PostgreSQL's XML type/functions…

5

Mapping XML to relational

- Store XML in a column
  - Simple, compact
  - CLOB (Character Large OBject) type + full-text indexing, or better, special XML type + functions
  - Poor integration with relational query processing
  - Updates are expensive
- Alternatives?
  - Schema-oblivious mapping: ← *Focus of this lecture*
    well-formed XML → generic relational schema
    1. Node/edge-based mapping for graphs
    2. Interval-based mapping for trees
    3. Path-based mapping for trees
  - Schema-aware mapping:
    valid XML → special relational schema based on DTD

6

1. Node/edge-based: schema

- *Element*(*eid*, *tag*)
- *Attribute*(*eid*, *attrName*, *attrValue*) Key: (*eid*, *attrName*)
  - Attribute order does not matter
- *ElementChild*(*eid*, *pos*, *child*) Keys: (*eid*, *pos*), (*child*)
  - *pos* specifies the ordering of children
  - *child* references either *Element*(*eid*) or *Text*(*tid*)
- *Text*(*tid*, *value*)
  - *tid* cannot be the same as any *eid*

☞Need to "invent" lots of *id*'s

☞Need indexes for efficiency, e.g., *Element*(*tag*), *Text*(*value*)

7

Node/edge-based: example

```
<bibliography>
 <book ISBN="ISBN-10" price="80.00">
  <title>Foundations of Databases</title>
  <author>Abiteboul</author>
  <author>Hull</author>
  <author>Vianu</author>
  <publisher>Addison Wesley</publisher>
  <year>1995</year>
 </book>...
</bibliography>
```

Element

eid	*tag*
e0	bibliography
e1	book
e2	title
e3	author
e4	author
e5	author
e6	publisher
e7	year

ElementChild

eid	*pos*	*child*
e0	1	e1
e1	1	e2
e1	2	e3
e1	3	e4
e1	4	e5
e1	5	e6
e1	6	e7
e2	1	t0
e3	1	t1
e4	1	t2
e5	1	t3
e6	1	t4
e7	1	t5

Attribute

eid	*attrName*	*attrValue*
e1	ISBN	ISBN-10
e1	price	80

Text

tid	*value*
t0	Foundations of Databases
t1	Abiteboul
t2	Hull
t3	Vianu
t4	Addison Wesley
t5	1995

8

Node/edge-based: simple paths

- //title
  - SELECT eid FROM Element WHERE tag = 'title';
- //section/title
  - SELECT e2.eid
    FROM Element e1, ElementChild c, Element e2
    WHERE e1.tag = 'section'
    AND e2.tag = 'title'
    AND e1.eid = c.eid
    AND c.child = e2.eid;

☞Path expression becomes joins!

- Number of joins is proportional to the length of the path expression

9

Node/edge-based: complex paths

- //bibliography/book[author="Abiteboul"]/@price
  - SELECT a.attrValue
    FROM Element e1, ElementChild c1,
    Element e2, Attribute a
    WHERE e1.tag = 'bibliography'
    AND e1.eid = c1.eid AND c1.child = e2.eid
    AND e2.tag = 'book'
    AND EXISTS (SELECT * FROM ElementChild c2,
    Element e3, ElementChild c3, Text t
    WHERE e2.eid = c2.eid AND c2.child = e3.eid
    AND e3.tag = 'author'
    AND e3.eid = c3.eid AND c3.child = t.tid
    AND t.value = 'Abiteboul')
    AND e2.eid = a.eid
    AND a.attrName = 'price';

some author of e2 is 'Abiteboul'

10

Node/edge-based: descendent-or-self

- //book//title
  - Requires SQL3 recursion
  - WITH RECURSIVE ReachableFromBook(id) AS
    ((SELECT eid FROM Element WHERE tag = 'book')
    UNION
    (SELECT c.child
    FROM ReachableFromBook r, ElementChild c
    WHERE r.eid = c.eid))
    SELECT eid
    FROM Element
    WHERE eid IN (SELECT * FROM ReachableFromBook)
    AND tag = 'title';

11

2. Interval-based: schema

- *Element(left, right, level, tag)*
  - *left* is the start position of the element
  - *right* is the end position of the element
  - *level* is the nesting depth of the element
  - Key is *left*
- *Text(left, right, level, value)*
  - Key is *left*
- *Attribute(left, attrName, attrValue)*
  - Key is (*left*, *attrName*)

12

Interval-based: example

```
1<bibliography>
 2<book ISBN="ISBN-10" price="80.00">
  3<title>4Foundations of Databases</title>5
  6<author>7Abiteboul</author>8
  9<author>10Hull</author>11
  12<author>13Vianu</author>14
  15<publisher>16Addison Wesley</publisher>17
  18<year>191995</year>20
 <book>21...
</bibliography>999
```

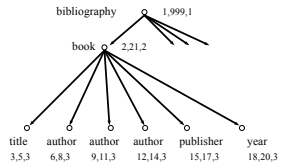

☞Where did *ElementChild* go?

- e_1 is the parent of e_2 iff:

  $[e_1.left, e_1.right] \supset [e_2.left, e_2.right]$, and
  $e_1.level = e_2.level - 1$

13

Interval-based: queries

- //section/title
  - SELECT e2.left
    FROM Element e1, Element e2
    WHERE e1.tag = 'section' AND e2.tag = 'title'
    AND e1.left < e2.left AND e2.right < e1.right
    AND e1.level = e2.level-1;

  ☞Path expression becomes "containment" joins!
    - Number of joins is proportional to path expression length
- //book//title
  - SELECT e2.left
    FROM Element e1, Element e2
    WHERE e1.tag = 'book' AND e2.tag = 'title'
    AND e1.left < e2.left AND e2.right < e1.right;

  ☞No recursion!

14

Summary so far

Node/edge-based vs. interval-based mapping

- Path expression steps
  - Equality vs. containment join
- Descendent-or-self
  - Recursion required vs. not required

15

3. Path-based mapping

Label-path encoding: paths as strings of labels

- *Element*(*pathid*, *left*, *right*, …), *Path*(*pathid*, *path*), …
  - *path* is a string containing the sequence of labels on a path starting from the root
  - Why are *left* and *right* still needed?

Element

pathid	*left*	*right*	…
1	1	999	…
2	2	21	…
3	3	5	…
4	6	8	…
4	9	11	…
4	12	14	…
…	…	…	…

Path

pathid	*path*
1	/bibliography
2	/bibliography/book
3	/bibliography/book/title
4	/bibliography/book/author
…	…

16

Label-path encoding: queries

- Simple path expressions with no conditions
  //book//title
  - Perform string matching on *Path*
  - Join qualified *pathid*'s with *Element*
- //book[publisher='Prentice Hall']/title
  - Evaluate //book/title
  - Evaluate //book/publisher[text()='Prentice Hall']
  - Must then ensure title and publisher belong to the same book (how?)

  ☞Path expression with attached conditions needs to be broken down, processed separately, and joined back

17

Another Path-based mapping

Dewey-order encoding

- Each component of the id represents the order of the child within its parent

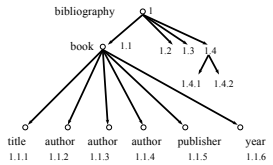

Element(*dewey_pid*, *tag*)
Text(*dewey_pid*, *value*)
Attribute(*dewey_pid*, *attrName*, *attrValue*)

18

Dewey-order encoding: queries

- Examples:
  //title
  //section/title
  //book//title
  //book[publisher='Prentice Hall']/title
  - Works similarly as interval-based mapping
    - Except parent/child and ancestor/descendant relationship are checked by prefix matching

19

An overview of XSLT, SAX, and DOM

20

XSLT

- XML-to-XML rule-based transformation language
  - Used most frequently as a stylesheet language
  - An XSLT program is an XML document itself

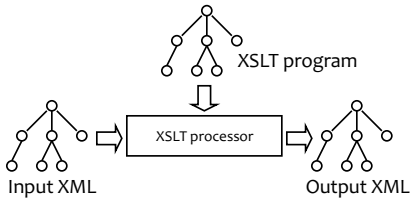

Actually, output does not need to be in XML in general

21

XSLT program

- An XSLT program is an XML document containing
  - Elements in the <xsl:> namespace
  - Elements in user namespace
- Roughly, result of evaluating an XSLT program on an input XML document = the XSLT document where each <xsl:> element is replaced with the result of its evaluation
- Basic ideas
  - Templates specify how to transform matching input nodes
  - Structural recursion applies templates to input trees recursively
- Uses XPath as a sub-language

22

XSLT elements

- Element describing transformation rules
  - <xsl:template>
- Elements describing rule execution control
  - <xsl:apply-templates>
  - <xsl:call-template>
- Elements describing instructions
  - <xsl:if>, <xsl:for-each>, <xsl:sort>, etc.
- Elements generating output
  - <xsl:value-of>, <xsl:copy-of>, <xsl:element>, <xsl:attribute>, <xsl:text>, etc.

23

XSLT example

- Find titles of books authored by "Abiteboul"

```
<?xml version="1.0"?>
<xsl:stylesheet
 xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
 version="2.0">
<xsl:template match="book[author='Abiteboul']">
  <booktitle>
    <xsl:value-of select="title"/>
  </booktitle>
</xsl:template>
</xsl:stylesheet>
```

Standard header of an XSLT document

- Not quite; we will see why later

24

<xsl:template>

```
<xsl:template match="book[author='Abiteboul']">
 <booktitle>
  <xsl:value-of select="title"/>
 </booktitle>
</xsl:template>
```

- <xsl:template match="*match_expr*"> is the basic XSLT construct describing a transformation rule
  - *match_expr* is an XPath-like expression specifying which nodes this rule applies to
- <xsl:value-of select="*xpath_expr*"/> evaluates *xpath_expr* within the context of the node matching the template, and converts the result sequence to a string
- <booktitle> and </booktitle> simply get copied to the output for each node matched

25

Template in action

```
<xsl:template match="book[author='Abiteboul']">
  <booktitle>
    <xsl:value-of select="title"/>
  </booktitle>
</xsl:template>
```

- Example XML fragment

```
<book ISBN="ISBN-10" price="80.00">
  <title>Foundations of Databases</title>
  <author>Abiteboul</author>
  <author>Hull</author>
  <author>Vianu</author>
  <publisher>Addison Wesley</publisher>
  <year>1995</year>
  <section>…</section>…
</book>
<book ISBN="ISBN-20" price="40.00">
  <title>A First Course in Databases</title>
  <author>Ullman</author>
  <author>Widom</author>
  <publisher>Prentice-Hall</publisher>
  <year>2002</year>
  <section>…</section>…
</book>
```

Template applies

```
<booktitle>
 Foundations of Databases
</booktitle>
```

Template does not apply; default behavior is to process the node recursively and print all text nodes

```
A First Course in Databases
Ullman
Widom
Prentice-Hall
2002
…… …
```

26

Removing the extra output

- Add the following template:

```
<xsl:template match="text()|@*"/>
```

- This template matches all text and attributes
- XPath features
  - text() is a node test that matches any text node
  - @* matches any attribute
  - | means "or" in XPath
- Body of the rule is empty, so all text and attributes become empty string
  - This rule effectively filters out things not matched by the other rule

27

Other features of XSLT

- Loop and condition
- White space control, insertion of newline
- Calling templates with parameters
- Debugging and exiting the program
  - <xsl:message>, <xsl:message terminate="yes">
- Defining variables, keys, functions

28

SAX & DOM

Both are API's for XML processing

- SAX (Simple API for XML)
  - Started out as a Java API, but now exists for other languages too
- DOM (Document Object Model)
  - Language-neutral API with implementations in Java, C++, python, etc.

29

SAX processing model

- Serial access
  - XML document is processed as a stream
  - Only one look at the data
  - Cannot go back to an early portion of the document
- Event-driven
  - A parser generates events as it goes through the document (e.g., start of the document, end of an element, etc.)
  - Application defines event handlers that get invoked when events are generated

30

A simple SAX example

- Print out text contents of title elements

```
import sys
import xml.sax
from StringIO import StringIO

class PathHandler(xml.sax.ContentHandler):
   def startDocument(self):
      ……
   def startElement(self, name, attrs):
      ……
   ……

xml.sax.parse(sys.stdin, PathHandler())
```

31

SAX events

Most frequently used events:

- startDocument
- endDocument
- startElement
- endElement
- characters
  - Whenever the parser has processed a chunk of character data (without generating other kinds of events)
  - Warning: The parser may generate multiple characters events for one piece of text

```
<?xml version="1.0"?>                    → startDocument
<bibliography>                           → startElement
  <book ISBN="ISBN-10" price="80.00">    → startElement
    <title>Foundations of Databases</title>
    …                                    → startElement
  </book>                                → endElement
  …
</bibliography>                          → endElement
                                         → endDocument
```

Whitespace may come up as characters or ignorableWhitespace, depending on whether a DTD is present

32

A simple SAX example (cont'd)

```
def startDocument(self):
    self.outBuffer = None
def startElement(self, name, attrs):
    if name == 'title':
        self.outBuffer = StringIO()
def endElement(self, name):
    if name == 'title':
        print self.outBuffer.getvalue()
        self.outBuffer = None
def characters(self, content):
    if self.outBuffer is not None:
        self.outBuffer.write(content)
```

Tag name

A map from attribute names to values

Characters read

33

DOM processing model

- XML is parsed by a parser and converted into an in-memory DOM tree
- DOM API allows an application to
  - Construct a DOM tree from an XML document
  - Traverse and read a DOM tree
  - Construct a new, empty DOM tree from scratch
  - Modify an existing DOM tree
  - Copy subtrees from one DOM tree to another

  etc.

34

Summary

- XML data can be "shredded" into rows in a relational database
- XQueries can be translated into SQL queries
  - Queries can then benefit from smart relational indexing, optimization, and execution
- With schema-oblivious approaches, comprehensive XQuery-SQL translation can be easily automated
  - Different data mapping techniques lead to different styles of queries
- Schema-aware translation is also possible and potentially more efficient, but automation is more complex