

Python XML and Database APIs

Christopher Simpkins

chris.simpkins@gatech.edu

Outline



XML Parsing

- SAX Parsers
- DOM Parsers
- ElementTree - the Pythonic way



Working with databases

- Connections
- Cursors
- Statements
- Transactions

XML Parsing

- ❏ CSV files are simply structured text-based data files
- ❏ XML files are more highly structured text-based data files, allowing nested, or “tree” structured data to be represented in text
- ❏ An XML Parser reads an XML file and extracts its structure. Three kinds of XML parsers:
 - SAX (Simple API for XML): a state machine that produces events for each element read from an XML file. Parser responds to events to process these elements
 - DOM (Document Object Model): the DOM standard specifies an object-tree representation of an XML document
 - » Tip: In Google Chrome you can see the DOM tree of any web page by clicking View->Developer->Developer Tools in the menu and clicking the Elements tab in the lower pane. Expand elements of the DOM tree and cursor over them to see the rendered parts of the page they represent. In Firefox you can install a plug-in to do the same thing
- ❏ We'll use ElementTree, which is essentially a Python-specific DOM parser

ElementTree

- ❏ An `ElementTree` is a representation of an XML document as a tree of `Element` objects with easy to use methods:
 - `parse("fileName")` reads an XML file and create an `ElementTree` representation of its document
 - `find("elementName")` gets a single child of an element
 - `findall("elementName")` gets an iterator over the like-named children of an element
- ❏ That's it. Really. It's that simple.

A Complete XML Parsing Example

- ❏ Say we have an XML file named `people.xml`

```
<?xml version="1.0"?>

<people>
  <person>
    <firstName>Alan</firstName>
    <lastName>Turing</lastName>
    <profession>Computer Scientist</profession>
  </person>
  <person>
    <firstName>Stephen</firstName>
    <lastName>Hawking</lastName>
    <profession>Physicist</profession>
  </person>
</people>
```

- ❏ Our XML document will have a root element named `people` and child elements named `person`
- ❏ Each `person` element will have three child elements named `firstName`, `lastName`, and `profession`

Parsing with ElementTree

❏ Parsing `people.xml` with `ElementTree` is this easy:

```
>>> import xml.etree.ElementTree
>>> people = xml.etree.ElementTree.parse("people.xml")
>>> persons = people.findall("person")
>>> for person in persons:
...     print person.find("firstName").text
...     print person.find("lastName").text
...     print person.find("profession").text
...
Alan
Turing
Computer Scientist
Stephen
Hawking
Physicist
>>>
```

Working with Databases

- ❏ Connection objects represent a connection to a database
 - provide transaction methods `rollback()` and `commit()`
 - provide a cursor object via `cursor()` to access the database
- ❏ Cursor object is a pointer to a part of the database
 - Connection's `cursor()` method returns a pointer to the database itself on which we can then execute statements
- ❏ SQL Statements are submitted to the `execute()` method of a database cursor
 - the `execute` method returns a cursor to its result
 - If the statement was a `select`, then the cursor can return the rows as a Python list of tuples via its `fetchall()` method

SQLite Databases

- ☒ Say we have the following database schema in `people-create.sql`:

```
create table if not exists person (  
    person_id integer primary key autoincrement,  
    first_name text,  
    last_name text,  
    profession text  
);
```

- ☒ And we create an empty SQLite database with it:

```
$ sqlite3 people.sqlite3  
SQLite version 3.7.9 2011-11-01 00:52:41  
Enter ".help" for instructions  
Enter SQL statements terminated with a ";"  
sqlite> .read people-create.sql  
sqlite> .exit  
$ ls  
people.sqlite3    people-create.sql
```


Inserting Data into an SQLite Database

- ❏ We can insert data into the database from within a Python program
- ❏ Use the cursor object to execute an SQL insert statement
- ❏ The insert statement uses ? markers for the values, and the values are supplied with a tuple

```
>>> import sqlite3
>>> conn = sqlite3.connect("people.sqlite3")
>>> curs = conn.cursor()
>>> curs.execute("insert into person values (?, ?, ?, ?)", (1, "Andy", "Register", "EE"))
<sqlite3.Cursor object at 0x1004dd1e8>
>>> curs.rowcount
1
>>> conn.commit()
```

- ❏ The `rowcount` attribute of the cursor object indicates how many rows of the database were affected
- ❏ The `commit()` method on the connection object causes the data to be written to the database. `rollback()` would undo all changes since the last `commit()`

Getting Data Out of an SQLite Database

- ❑ We can get data with an SQL select query
- ❑ After executing a select, the cursor's `fetchall()` method returns the results as a list of tuples

```
>>> curs.execute("insert into person values (?, ?, ?, ?)", (2, "Carlee", "Bishop", "Sys Eng"))
<sqlite3.Cursor object at 0x1004dd1e8>
>>> conn.commit()
>>> curs.execute("select * from person")
<sqlite3.Cursor object at 0x1004dd1e8>
>>> for row in curs.fetchall():
...     print(row)
...
(1, u'Andy', u'Register', u'EE')
(2, u'Carlee', u'Bishop', u'Sys Eng')
>>>
```

A Complete Example: Reading XML Data into a Database

- ❏ Say we have a system that receives data externally in the form of XML files and inserts the data from them into a database
- ❏ We can read the data from the XML file using ElementTree parsing shown earlier
- ❏ We can insert the data into our database using the database APIs we just saw
- ❏ Assume we have an empty database created with the people-create.sql file we saw earlier
- ❏ Let's write a program to read an XML file, extract the data from it, and insert the data into the database

Inserting Data into the Database

- First, we'll write a function to insert persons into the database in `people.py`

```
import sqlite3

def insert_person(db_conn, person_id, first_name, last_name, profession):
    curs = db_conn.cursor()
    curs.execute("insert into person values (?, ?, ?, ?)",
                  (person_id, first_name, last_name, profession))
    conn.commit()
```

- We can test our program in progress by importing it into the Python interactive shell

```
>>> import people
>>> import sqlite3
>>> conn = sqlite3.connect("people.sqlite3")
>>> people.insert_person(conn, 1, "George", "Burdell", "Student")
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "people.py", line 7, in insert_person
    conn.commit()
NameError: global name 'conn' is not defined
```

- Oops. It's `db_conn`, not `conn`

Debugging Cycle

☒ So we fix our function:

```
import sqlite3

def insert_person(db_conn, person_id, first_name, last_name, profession):
    curs = db_conn.cursor()
    curs.execute("insert into person values (?, ?, ?, ?)",
                  (person_id, first_name, last_name, profession))
    db_conn.commit()
```

☒ Recreate our empty database:

```
$ rm people.sqlite3; sqlite3 people.sqlite3 < people-create.sql
```

☒ And try our (hopefully) fixed function:

```
$ python
>>> import people
>>> import sqlite3
>>> conn = sqlite3.connect("people.sqlite3")
>>> people.insert_person(conn, 1, "George", "Burdell", "Student")
>>> conn.cursor().execute("select * from person").fetchall()
[(1, u'George', u'Burdell', u'Student')]
```

☒ Now insert_person() works. On to XML extraction

person_data_from_element

- Now let's add a function to extract person data from an XML element:

```
import xml.etree.ElementTree

def person_data_from_element(element):
    first = element.find("firstName").text
    last = element.find("lastName").text
    profession = element.find("profession").text
    return first, last, profession
```

- And test in the Python shell:

```
$ python
>>> import people
>>> import xml.etree.ElementTree
>>> peeps = xml.etree.ElementTree.parse("people.xml")
>>> first_person = peeps.findall("person")[0]
>>> people.person_data_from_element(first_person)
('Alan', 'Turing', 'Computer Scientist')
```

- Note that we called the XML element `peeps` so it wouldn't clobber the `people` module name

The Complete Program

Now we add the main loop to get all the persons from the XML file and our program is complete

```
import sqlite3
import xml.etree.ElementTree

def insert_person(db_conn, person_id, first_name, last_name, profession):
    curs = db_conn.cursor()
    curs.execute("insert into person values (?, ?, ?, ?)",
                  (person_id, first_name, last_name, profession))
    db_conn.commit()

def person_data_from_element(element):
    first = element.find("firstName").text
    last = element.find("lastName").text
    profession = element.find("profession").text
    return first, last, profession

if __name__ == "__main__":
    conn = sqlite3.connect("people.sqlite3")
    people = xml.etree.ElementTree.parse("people.xml")
    persons = people.findall("person")
    for index, element in enumerate(persons):
        first, last, profession = person_data_from_element(element)
        insert_person(conn, index, first, last, profession)
```

Our Program in Action

```
$ rm people.sqlite3; sqlite3 people.sqlite3 < people-create.sql
$ python people.py
$ sqlite3 people.sqlite3
SQLite version 3.7.9 2011-11-01 00:52:41
Enter ".help" for instructions
Enter SQL statements terminated with a ";"
sqlite> select * from person;
0|Alan|Turing|Computer Scientist
1|Stephen|Hawking|Physicist
sqlite>
```

- ☒ It works!
- ☒ But it's lame to have to manually insert primary keys

Determining a Safe Primary Key (Hard Way)

- ❑ If we insert a record with a duplicate key, it will fail
- ❑ If we know keys are integers, we can get the integers and add one to the highest one
 - Note: can't just use size of table - key values may be greater due to inserts and deletes
- ❑ Here's how we can do it in Python:

```
>>> import sqlite3
>>> conn = sqlite3.connect("people.sqlite3")
>>> curs = conn.cursor().execute("select person_id from person")
>>> results = curs.fetchall()
>>> results
[(0,), (1,)]
>>> keys = [tuple[0] for tuple in results]
>>> keys
[0, 1]
>>> keys.sort()
>>> largest_key = keys[len(keys) - 1]
>>> largest_key
1
```

Determining a Safe Primary Key (Easy Way)

- ❏ The easy way to determine a safe primary key is not to
- ❏ `person_id` is an auto increment field, so if we insert a new record into the table with no `person_id`, SQLite figures out a safe `person_id` using the logic we just described
- ❏ So instead of
 - `insert into person (person_id, first_name, last_name, profession) values (...`
- ❏ We use
 - `insert into person (first_name, last_name, profession) values (...`
- ❏ We simply leave out the PK and let SQLite figure it out.

A More Robust Program

☒ This version doesn't require an empty database

```
import sqlite3
import xml.etree.ElementTree

def insert_person(db_conn, first_name, last_name, profession):
    curs = db_conn.cursor()
    curs.execute("insert into person (first_name, last_name, profession) values (?, ?, ?)",
                  (first_name, last_name, profession))
    db_conn.commit()

def person_data_from_element(element):
    first = element.find("firstName").text
    last = element.find("lastName").text
    profession = element.find("profession").text
    return first, last, profession

if __name__ == "__main__":
    conn = sqlite3.connect("people.sqlite3")
    people = xml.etree.ElementTree.parse("people.xml")
    persons = people.findall("person")
    for element in persons:
        first, last, profession = person_data_from_element(element)
        insert_person(conn, first, last, profession)
```

Conclusion

- ❏ We learned
 - How to parse XML
 - How to work with databases
 - How to connect XML data and databases using Python programming
- ❏ Lots more to know, but we saw the basics
- ❏ Consult Python docs or books for more details