

5. XQuery

The purpose of this lab is to practice

- writing XQuery queries.

Hint: Run your XQuery queries using Kernow. You can download Kernow from

<http://kernowforsaxon.sourceforge.net/>

Kernow is a Java package and can be run under Linux, Windows and Mac OS. Kernow is a graphical frontend to the Saxon XQuery/XSLT processor.

You can develop your queries in the XQuery Sandbox. To tell Kernow that the output should be properly indented, type as the first line:

```
declare option saxon:output "indent=yes";
```

The version of Kernow that you download allows you to run 100 queries. Then it will ask you to enter a code in order to proceed. I have the code and will send it to you if you ask me by email.

Country Queries

The following are queries over the `countries.xml` document. Formulate your queries in such a way that they return the correct result for all possible documents that satisfy `countries.dtd`.

1. Return a list of `<country>` elements, containing elements `<name>` and `<population>`, alphabetically ordered by name.
2. Return a similar list, ordered by `<population>` in descending order.
3. Return a similar list, again ordered by population in descending order, with an additional attribute `rank`, which indicates the position of the country on the list.

Try two approaches, one using the construct

```
for $var at $pos in <expression>
```

and another one without. Note that you may have to use the casting functions `number` and `string`.

4. Return a list of city elements, containing the name of the city, such that each city has an attribute `population` and another attribute `country`. The cities are returned according to their population, in descending order.

5. Restructure the document by listing countries according to population, cities within each country according to population, and languages within each country according to percentage, all in descending order. Try to reuse as much as possible existing elements and attributes instead of creating new ones. (Creation of nodes is usually a costly operation in XQuery engines.)
6. Add to each country element new attributes that record the number of cities in the country and the number of languages spoken in the country.
7. Return an element `<languages>` with a list of `<language>` elements, alphabetically sorted, where each language from the document occurs exactly once.
8. Modify the previous query so that you return only those languages that are spoken by at least 10 percent of the population in those countries where they occur.
9. Extend query number 7 so that each language element contains a list of country elements, such that the language is spoken in the country.
10. Extend the previous query so that you return with every country element below a language the number of speakers of the language in that country. You may again need a casting function (`xs:int` casts an argument as an integer, whenever that is possible; `xs` is the default prefix for the XML Schema namespace).