A Guide to $\mathcal{F}$LORA-*2* Packages

Version 1.2

(Monstera deliciosa)

January 31, 2017

Contents

Chapter 1

JAVA-to-$\mathcal{F}$LORA-*2* Interfaces

by Aditi Pandit and Michael Kifer

This chapter documents the API for accessing $\mathcal{F}$LORA-*2* from Java programs. The API has two versions: a *low-level API*, which enables Java programs to send arbitrary queries to $\mathcal{F}$LORA-*2* and get results, and a *high-level API*, which is more limited, but is easier to use. The high-level API establishes a correspondence between Java classes and $\mathcal{F}$LORA-*2* classes, which enables manipulation of $\mathcal{F}$LORA-*2* classes by executing appropriate methods on the corresponding Java classes. Both interfaces rely on the Java-XSB interface, called *Interprolog*, developed by `Interprolog.com`.

The API assumes that a Java program is started first and then it invokes XSB/$\mathcal{F}$LORA-*2* as a subprocess. The XSB/$\mathcal{F}$LORA-*2* side is passive: it only responds to the queries sent by the Java side. Queries can be anything that is accepted at the $\mathcal{F}$LORA-*2* shell prompt: queries, insert/delete commands, control switches, etc., are all fine. One thing to remember is that the backslash is used in Java as an escape symbol and in $\mathcal{F}$LORA-*2* as a prefix of the builtin operators and commands. Therefore, each backslash must be escaped with another backslash. That is, instead of a query like "`p(?X) \and q(?X).`" the API requires "`p(?X) \\and q(?X).`".

1.1 The Low-level Interface

The low-level API enables Java programs to send arbitrary queries to $\mathcal{F}$LORA-*2* and get results. It is assumed that the following environment variables are set:

`JAVA_BIN`: This variable points to the folder containing the javac and java executable programs. This variable is set in the `windowsVariables.bat` and `unixVariables.sh` scripts in the `java` subfolder of the $\mathcal{F}$LORA-*2* distribution.

`PROLOGDIR`: This variable points to the folder containing the XSB executable. This variable is set in the `flora_settings.bat` and `flora_settings.sh` scripts in the `java` folder.

`FLORADIR`: This variable must point to the folder containing the $\mathcal{F}$LORA-*2* installation. It is set by the `flora_settings.bat` and `flora_settings.sh` files in the `java` subfolder and this is where users should look in order to get the correct values for their systems. Both of the above files are generated automatically by the system installation scripts.

In order to be able to access $\mathcal{F}$LORA-*2*, the Java program must first establish a session for a running instance of $\mathcal{F}$LORA-*2*. Multiple sessions can be active at the same time. The knowledge bases in the different running instances are completely independent. Sessions are instances of the class `javaAPI.src.FloraSession`. This class provides methods for opening/closing sessions and loading $\mathcal{F}$LORA-*2* knowledge bases (which are also used in the high-level interface). In addition, a session provides methods for executing arbitrary $\mathcal{F}$LORA-*2* queries. The following is the complete list of the methods that are available in that class.

- `public FloraSession()`

  This method creates a connection to an instance of $\mathcal{F}$LORA-*2*.

- `close()`

  This method must be called to terminate a $\mathcal{F}$LORA-*2* session. Note that this does not terminate the Java program that initiated the session: to exit the Java program that talks to $\mathcal{F}$LORA-*2*, one needs to execute the statement

  ```
  System.exit();
  ```

  Note that just returning from the `main` method is not enough.

- `public Iterator<FloraObject> ExecuteQuery(String command)`

  This method executes the $\mathcal{F}$LORA-*2* command given by the parameter `command`. It is used to execute $\mathcal{F}$LORA-*2* queries that do not require variable bindings to be returned back to Java or queries that have only a single variable to be returned. Each binding is represented as an instance of the class `javaAPI.src.FloraObject`. The examples below illustrate how to process the results returned by this method.

- `public Iterator<HashMap<String,FloraObject>> ExecuteQuery(String query, Vector vars)`

  This method executes the $\mathcal{F}$LORA-*2* query given by the first argument. The Vector `vars` (of strings) specifies the names of all the variables in the query for which bindings need to be returned. These variables are added to the vector using the method `add` before calling `ExecuteQuery`. For instance, `vars.add("?X")`.

  This version of `ExecuteQuery` returns an iterator over all bindings returned by the $\mathcal{F}$LORA-*2* query. Each binding is represented by a `HashMap<String,FloraObject>` object which can be used to obtain the value of each variable in the query (using the `get()` method). The value of each variable returned is an instance of `javaAPI.src.FloraObject`.

  See the examples below for how to handle the results returned by this method.

- `void loadFile(String fileName,String moduleName)`

  This method loads the $\mathcal{F}$LORA-*2* program, specified by the parameter `fileName` into the $\mathcal{F}$LORA-*2* module specified in `moduleName`.

- `void compileFile(String fileName,String moduleName)`

  This method compiles (but does not load) the $\mathcal{F}$LORA-*2* program, specified by the parameter `fileName` for the $\mathcal{F}$LORA-*2* module specified in `moduleName`.

- `void addFile(String fileName,String moduleName)`

  This method adds the $\mathcal{F}$LORA-2 program, specified by the parameter `fileName` to an existing $\mathcal{F}$LORA-2 module specified in `moduleName`.

- `void compileaddFile(String fileName,String moduleName)`

  This method compiles the $\mathcal{F}$LORA-2 program, specified by the parameter `fileName` for addition to the $\mathcal{F}$LORA-2 module specified in `moduleName`.

The code snippet below illustrates the low-level API.

Step 1: Writing a $\mathcal{F}$LORA-2 program. Let us assume that we have a file, called `flogic_basics.flr`, which contains the following information:

```
person :: object.
dangerous_hobby :: object.
john:employee.
employee::person.

bob:person.
tim:person.
betty:employee.

person[|age=>integer,
       kids=>person,
       salary(year)=>value,
       hobbies=>hobby,
       believes_in=>something,
       instances => person
|].

mary:employee[
    age->29,
    kids -> {tim,leo,betty},
    salary(1998) -> a_lot
].

tim[hobbies -> {stamps, snowboard}].
betty[hobbies->{fishing,diving}].

snowboard:dangerous_hobby.
diving:dangerous_hobby.

?_X[self-> ?_X].

person[|believes_in -> {something, something_else}|].
```

Step 2: Writing a JAVA application to interface with $\mathcal{F}$LORA-*2*. The following code loads a $\mathcal{F}$LORA-*2* program from a file and then passes queries to the knowledge base.

```
import java.util.*;
import net.sf.flora2.API.*;
import net.sf.flora2.API.util.*;

public class flogicbasicsExample {

    public static void main(String[] args) {
        // create a new session for a running instance of the engine
        FloraSession session = new FloraSession();
        System.out.println("Engine session started");

        // Assume that Java was called with -DINPUT_FILE=the-file-name
        String fileName = System.getProperty("INPUT_FILE");
        if(fileName == null || fileName.trim().length() == 0) {
            System.out.println("Invalid path to example file!");
            System.exit(0);

        }
        // load the program into module basic_mod
        session.loadFile(fileName,"basic_mod");

        /* Running queries from flogic_basics.flr */

        /* Query for persons */
        String command = "?X:person@basic_mod.";
        System.out.println("Query:"+command);
        Iterator<FloraObject> personObjs = session.ExecuteQuery(command);

        /* Printing out the person names and information about their kids */
        while (personObjs.hasNext()) {
            FloraObject personObj = personObjs.next();
            System.out.println("Person name:"+personObj);


        command = "person[instances -> ?X]@basic_mod.";
        System.out.println("Query:"+command);
        personObjs = session.ExecuteQuery(command);

        /* Prining out the person names  */
        while (personObjs.hasNext()) {
            Object personObj = personObjs.next();
            System.out.println("Person Id: "+personObj);
        }
```

```
        /* Example of ExecuteQuery with two arguments */
        Vector<String> vars = new Vector<String>();
        vars.add("?X");
        vars.add("?Y");

        Iterator<HashMap<String,FloraObject>> allmatches =
            session.ExecuteQuery("?X[believes_in -> ?Y]@basic_mod.",vars);
        System.out.println("Query:?X[believes_in -> ?Y]@basic_mod.");
        while(allmatches.hasNext()) {
            HashMap<String,FloraObject> firstmatch = allmatches.next();
            Object Xobj = firstmatch.get("?X");
            Object Yobj = firstmatch.get("?Y");
            System.out.println(Xobj+" believes in: "+?Yobj);
        }
        // quit the system
        session.close();
        System.exit(0);
    }
}
```

For the information on how to invoke the above Java class in the context of the Java-$\mathcal{F}$LORA-*2* API, please see Section 1.3.

1.2 The High-Level Interface

The high-level API operates by creating proxy Java classes for $\mathcal{F}$LORA-*2* classes selected by the user. This enables the Java program to operate on $\mathcal{F}$LORA-*2* classes by executing appropriate methods on the corresponding proxy Java classes. The use of the high-level API involves a number of steps, as described below.

Readers who intend to use only the low-level Java-$\mathcal{F}$LORA-*2* interface can skip this section.

Note: This interface will not work for $\mathcal{F}$LORA-*2* programs that use *non-alphanumeric* names for methods and predicates. For instance, if a program involves statements like `foo['bar$#123'->456]` then the interface might generate syntactically incorrect Java proxy classes and errors will be issued during the compilation.

Stage 1: Writing a $\mathcal{F}$LORA-*2* file. We assume the same `flogic_basics.flr` file as in the previous example.

Stage 2: Generating Java classes that serve as proxies for $\mathcal{F}$LORA-*2* classes. The $\mathcal{F}$LORA-*2* side of the Java-to-$\mathcal{F}$LORA-*2* high level API provides a predicate to generate Java proxy classes for each F-logic class which have a signature declaration in the $\mathcal{F}$LORA-*2* knowledge base. A proxy class gets defined so that it would have methods to manipulate the attributes and methods of the corresponding F-logic class for which signature declarations are available. If an F-logic class has a declared value-returning attribute `foobar` then the proxy class will have the following methods. Each method name has the form *action*$S_1S_2S_3$_foobar, where *action* is either get, set, or delete. The specifier S_1 indicates the type of the method — V for value-returning, B for Boolean, and P for procedural. The specifier S_2 tells whether the operation applies to the signature of the method (S), e.g., person[foobar=>string], or to the actual data (D), for example, john[foobar->3]. Finally, the specifier S_3 tells if the operation applies to the inheritable variant of the method (I) or its non-inheritable variant (N).

1. public Iterator<FloraObject> getVDI_foobar()
   public Iterator<FloraObject> getVDN_foobar()
   public Iterator<FloraObject> getVSI_foobar()
   public Iterator<FloraObject> getVSN_foobar()
   The above methods query the knowledge base and get all answers for the attribute foobar. They return iterators through which these answers can be processed one-by-one. Each object returned by the iterator is of type FloraObject. The getVDN form queries non-inheritable data methods and getVDI the inheritable ones. The getVSI and getVSN forms query the signatures of the attribute foobar.

2. public boolean setVDI_foobar(Vector value)
   public boolean setVDN_foobar(Vector value)
   public boolean setVSI_foobar(Vector value)
   public boolean setVSN_foobar(Vector value)
   These methods add values to the set of values returned by the attribute foobar. The values must be placed in the vector parameter passed these methods. Again, setVDN adds data for non-inheritable methods and setVDI is used for inheritable methods. setVSI and setVSN add types to signatures.

3. public boolean setVDI_foobar(Object value)
   public boolean setVDN_foobar(Object value)
   public boolean setVSI_foobar(Object value)
   public boolean setVSN_foobar(Object value)
   These methods provide a simplified interface when only one value needs to be added. It works like the earlier set_* methods, except that only one value given as an argument is added.

4. public boolean deleteVDI_foobar(Vector value)
   public boolean deleteVDN_foobar(Vector value)
   public boolean deleteVSI_foobar(Vector value)
   public boolean deleteVSN_foobar(Vector value)
   Delete a set of values of the attribute foobar. The set is specified in the vector argument.

5. public boolean deleteVDI_foobar(Object value)
   public boolean deleteVDN_foobar(Object value)
   public boolean deleteVSI_foobar(Object value)
   public boolean deleteVSN_foobar(Object value)
   A simplified interface for the case when only one value needs to be deleted.

6. public boolean deleteVDI_foobar()
   public boolean deleteVDN_foobar()

public boolean deleteVSI_foobar()
public boolean deleteVSN_foobar()
Delete all values for the attribute `foobar`.

For F-logic methods with arguments, the high-level API provides Java methods as above, but they take more arguments to accommodate the parameters that F-logic methods take. Let us assume that the F-logic method is called `foobar2` and it takes parameters `arg1` and `arg2`. As before the `getVDI_*`, `setVDI_*`, etc., forms of the Java methods are for dealing with inheritable $\mathcal{F}$LORA-*2* methods and the `getVDN_*`, `setVDN_*`, etc., forms are for dealing with non-inheritable $\mathcal{F}$LORA-*2* methods.

1. `public Iterator<FloraObject> getVDI_foobar2(Object arg1, Object arg2)`
   `public Iterator<FloraObject> getVDN_foobar2(Object arg1, Object arg2)`
   Obtain all values for the F-logic method invocation `foobar2(arg1,arg2)`.
2. `public boolean setVDI_foobar2(Object arg1, Object arg2, Vector value)`
   `public boolean setVDN_foobar2(Object arg1, Object arg2, Vector value)`
   Add a set of methods specified in the parameter `value` for the method invocation `foobar2(arg1,arg2)`.
3. `public boolean setVDI_foobar2(Object arg1, Object arg2, Object value)`
   `public boolean setVDN_foobar2(Object arg1, Object arg2, Object value)`
   A simplified interface when only one value is to be added.
4. `public boolean deleteVDI_foobar2(Object arg1, Object arg2, Vector value)`
   `public boolean deleteVDN_foobar2(Object arg1, Object arg2, Vector value)`
   Delete a set of values from `foobar2(arg1,arg2)`. The set is given by the vector parameter `value`.
5. `public boolean deleteVDI_foobar2(Object arg1, Object arg2, Object value)`
   `public boolean deleteVDN_foobar2(Object arg1, Object arg2, Object value)`
   A simplified interface for deleting a single value.
6. `public boolean deleteVDI_foobar2(Object arg1, Object arg2)`
   `public boolean deleteVDN_foobar2(Object arg1, Object arg2)`
   Delete all values for the method invocation `foobar2(arg1,arg2)`.

For Boolean and procedural methods, the generated methods are similar except that there is only one version for the set and delete methods. In addition, Boolean inheritable methods use the `getBDI_*`, `setBDI_*`, etc., form, while non-inheritable methods use the `getBDN_*`, etc., form. Procedural methods use the `getPDI_*`, `getPDN_*`, etc., forms. For instance,

1. `public boolean getBDI_foobar3()`
   `public boolean getBDN_foobar3()`
   `public boolean getPDI_foobar3()`
   `public boolean getPDN_foobar3()`
2. `public boolean setBDI_foobar3()`
   `public boolean setBDN_foobar3()`

```
public boolean setPDI_foobar3()
public boolean setPDN_foobar3()
```

3. ```
   public boolean deleteBDI_foobar3()
   public boolean deleteBDN_foobar3()
   public boolean deletePDI_foobar3()
   public boolean deletePDN_foobar3()
   ```

In addition, the methods to query the ISA hierarchy are available:

- `public Iterator<FloraObject> getDirectInstances()`
- `public Iterator<FloraObject> getInstances()`
- `public Iterator<FloraObject> getDirectSubClasses()`
- `public Iterator<FloraObject> getSubClasses()`
- `public Iterator<FloraObject> getSuperClasses()`
- `public Iterator<FloraObject> getDirectSuperClasses()`

These methods apply to the java proxy object that corresponds to the F-logic class person.

All these methods are generated automatically by executing the following $\mathcal{F}$LORA-*2* query (defined in the `javaAPI` package). All arguments in the query must be bound:

```
// write(?Class,?Module,?ProxyClassFileName).
?- write(foo,example,'myproject/foo.java').
```

The first argument specifies the class for which to generate the methods, the file name tells where to put the Java file for the proxy object, and the model argument tells which $\mathcal{F}$LORA-*2* model to load this program to. The result of this execution will be the file `foo.java` which should be included with your java program (the program that is going to interface with $\mathcal{F}$LORA-*2*). Note that because of the Java conventions, the file name must have the same name as the class name. It is important to remember, however, that proxy methods will be generated only for those F-logic methods that have been declared using signatures.

Let us now come back to our program `flogic_basics.flr` for which we want to use the high-level API. Suppose we want to query the person class. To generate the proxy declarations for that class, we create the file `person.java` for the module `basic_mod` as follows.

```
?- load{'examples/flogic_basics'>>basic_mod}.
?- load{javaAPI}.
?- write(person,basic_mod,'examples/person.java')@\prolog
```

The `write` method will create the file `person.java` shown below. The methods defined in `person.java` are the class constructors for `person`, the methods to query the ISA hierarchy, and the "get", "set" and "delete" methods for each method and attribute declared in the $\mathcal{F}$LORA-*2* class `person`. The parameters for the "get", "set" and "delete" Java methods
```

are the same as for the corresponding $\mathcal{F}$LORA-*2* methods. The first constructor for class `person` takes a low-level object of class `javaAPI.src.FloraObject` as a parameter. The second parameter is the $\mathcal{F}$LORA-*2* module for which the proxy object is to be created. The second `person`-constructor takes F-logic object Id instead of a low-level `FloraObject`. It also takes the module name, as before, but, in addition, it takes a session for a running $\mathcal{F}$LORA-*2* instance. The session parameter was not needed for the first `person`-constructor because `FloraObject` is already attached to a concrete session.

It can be seen from the form of the proxy object constructors that proxy objects are attached to specific $\mathcal{F}$LORA-*2* modules, which may seem to go against the general philosophy that F-logic objects do not belong to any module — only their methods do. On closer examination, however, attaching high-level proxy Java objects to modules makes perfect sense. Indeed, a proxy object encapsulates operations for manipulating F-logic attributes and methods, which belong to concrete $\mathcal{F}$LORA-*2* modules, so the proxy object needs to know which module it operates upon.

person.java file

```
import java.util.*;
import net.sf.flora2.API.*;
import net.sf.flora2.API.util.*;

public class person {

    public FloraObject sourceFloraObject;

    // proxy objects' constructors
    public person(FloraObject sourceFloraObject, String moduleName) { ... }
    public person(String floraOID,String moduleName, FloraSession session) { ... }

    // ISA hierarchy queries
    public Iterator<FloraObject> getDirectInstances() { ... }
    public Iterator<FloraObject> getInstances() { ... }
    public Iterator<FloraObject> getDirectSubClasses() { ... }
    public Iterator<FloraObject> getSubClasses() { ... }
    public Iterator<FloraObject> getDirectSuperClasses() { ... }
    public Iterator<FloraObject> getSuperClasses() { ... }

    // Java methods for manipulating methods
    public boolean setVDI_age(Object value) { ... }
    public boolean setVDN_age(Object value) { ... }
    public Iterator<FloraObject> getVDI_age(){ ... }
    public Iterator<FloraObject> getVDN_age(){ ... }
    public boolean deleteVDI_age(Object value) { ... }
    public boolean deleteVDN_age(Object value) { ... }
    public boolean deleteVDI_age() { ... }
    public boolean deleteVDN_age() { ... }
    public boolean setVDI_salary(Object year,Object value) { ... }
```

```
    public boolean setVDN_salary(Object year,Object value) { ... }
    public Iterator<FloraObject> getVDI_salary(Object year) { ... }
    public Iterator<FloraObject> getVDN_salary(Object year) { ... }
    public boolean deleteVDI_salary(Object year,Object value) { ... }
    public boolean deleteVDN_salary(Object year,Object value) { ... }
    public boolean deleteVDI_salary(Object year) { ... }
    public boolean deleteVDN_salary(Object year) { ... }
    public boolean setVDI_hobbies(Vector value) { ... }
    public boolean setVDN_hobbies(Vector value) { ... }
    public Iterator<FloraObject> getVDI_hobbies(){ ... }
    public Iterator<FloraObject> getVDN_hobbies(){ ... }
    public boolean deleteVDI_hobbies(Vector value) { ... }
    public boolean deleteVDN_hobbies(Vector value) { ... }
    public boolean deleteVDI_hobbies(){ ... }
    public boolean deleteVDN_hobbies(){ ... }
    public boolean setVDI_instances(Vector value) { ... }
    public boolean setVDN_instances(Vector value) { ... }
    public Iterator<FloraObject> getVDI_instances(){ ... }
    public Iterator<FloraObject> getVDN_instances(){ ... }
    public boolean deleteVDI_instances(Vector value) { ... }
    public boolean deleteVDN_instances(Vector value) { ... }
    public boolean deleteVDI_instances(){ ... }
    public boolean deleteVDN_instances(){ ... }
    public boolean setVDI_kids(Vector value) { ... }
    public boolean setVDN_kids(Vector value) { ... }
    public Iterator<FloraObject> getVDI_kids(){ ... }
    public Iterator<FloraObject> getVDN_kids(){ ... }
    public boolean deleteVDI_kids(Vector value) { ... }
    public boolean deleteVDN_kids(Vector value) { ... }
    public boolean deleteVDI_kids(){ ... }
    public boolean deleteVDN_kids(){ ... }
    public boolean setVDI_believes_in(Vector value) { ... }
    public boolean setVDN_believes_in(Vector value) { ... }
    public Iterator<FloraObject> getVDI_believes_in(){ ... }
    public Iterator<FloraObject> getVDN_believes_in(){ ... }
    public boolean deleteVDI_believes_in(Vector value) { ... }
    public boolean deleteVDN_believes_in(Vector value) { ... }
    public boolean deleteVDI_believes_in(){ ... }
    public boolean deleteVDN_believes_in(){ ... }
}
```

Stage 3: Writing Java applications that use the high-level API. The following program (`flogicbasicsExample.java`) shows several queries that use the high-level interface. The class `person.java` is generated at the previous stage. The methods of the high-level interface operate on Java objects that are proxies for $\mathcal{F}$LORA-*2* objects. These Java objects are members of the class `javaAPI.src.FloraObject`. Therefore, before one can use the high-level methods one need to first retrieve the appropriate proxy objects on which to operate. This is done by sending an appropriate query through the method `ExecuteQuery`—the same method that was used in the low-level interface. Alternatively, `person`-objects could be constructed using the 3-argument proxy constructor, which takes F-logic oids.

```
import java.util.*;
import net.sf.flora2.API.*;
import net.sf.flora2.API.util.*;

public class flogicbasicsExample {

    public static void main(String[] args) {
        /* Initializing the session */
        FloraSession session = new FloraSession();
        System.out.println("Flora session started");

        String fileName = "examples/flogic_basics"; // must be a valid path
        /* Loading the flora file */
        session.loadFile(fileName,"basic_mod");

        // Retrieving instances of the class person through low-level API
        String command = "?X:person@basic_mod.";
        System.out.println("Query:"+command);
        Iterator<FloraObject> personObjs = session.ExecuteQuery(command);

        /* Print out person names and information about their kids */
        person currPerson = null;
        while (personObjs.hasNext()) {
            FloraObject personObj = personObjs.next();
            // Elevate personObj to the higher-level person-object
            currPerson =new person(personObj,"basic_mod");

            /* Set that person's age to 50 */
            currPerson.setVDN_age("50");

            /* Get this person's kids */
            Iterator<FloraObject> kidsItr = currPerson.getVDN_kids();
            while (kidsItr.hasNext()) {
                FloraObject kidObj = kidsItr.next();
                System.out.println("Person: " + personObj + " has kid: " +kidObj);

                person kidPerson = null;
                // Elevate kidObj to kidPerson
                kidPerson = new person(kidObj,"basic_mod");

                /* Get kidPerson's hobbies */
                Iterator<FloraObject> hobbiesItr = kidPerson.getVDN_hobbies();
```

```
                while(hobbiesItr.hasNext()) {
                    FloraObject hobbyObj = hobbiesItr.next();
                    System.out.println("Kid:"+kidObj + " has hobby:" +hobbyObj);
                }
            }
        }

        FloraObject age;
        // create a person-object directly by supplying its F-logic OID
        // father(mary)
        currPerson = new person("father(mary)", "example", session);
        Iterator<FloraObject> maryfatherItr = currPerson.getVDN_age();
        age = maryfatherItr.next();
        System.out.println("Mary's father is " + age + " years old");

        // create a proxy object for the F-logic class person itself
        person personClass = new person("person", "example", session);
        // query its instances through the high-level interface
        Iterator<FloraObject> instanceIter = personClass.getInstances();
        System.out.println("Person instances using high-level API:");
        while (instanceIter.hasNext())
            System.out.println("    " + instanceIter.next());

        session.close();
        System.exit();
    }
}
```

1.3 Executing Java Application Programs with $\mathcal{F}$LORA-*2*

To run Java programs that interface with $\mathcal{F}$LORA-*2*, follow the following guidelines.

- Place the files `flogicsbasicsExample.java` (the program you have written) and `person.java` (the automatically generated file) in the same directory and compile them using the `javac` command. Add the jar-files containing the API code and `interprolog.jar` to the Java classpath:
  - *FLORADIR*`/java/flora2java.jar`
  - *FLORADIR*`/java/interprolog.jar`

  FLORADIR here should be replaced with the value of the variable `FLORADIR`, which is set by the scripts `flora_settings.sh` (Linux/Mac) or `flora_settings.bat` (Windows), as mentioned in Section 1.1 on page 1.
- Generally, Java programs that call $\mathcal{F}$LORA-*2* should be invoked using the following command. For Unix-like systems (Linux, Mac, etc.), change `%`*VAR*`%` to `$`*VAR*:

```
%JAVA_BIN%\java -DPROLOGDIR=%PROLOGDIR%
                -DFLORADIR=%FLORADIR%
                -Djava.library.path=%PROLOGDIR%
                -classpath %MYCLASSPATH% flogicbasicsExample
```

The above command uses several shell variables, which are explained below. Instead of using the variables, one can substitute their values directly.

`JAVA_BIN`: This variable should point to the directory containing the `java` and `javac` executables of the JDK.

`PROLOGDIR`: This variable should be set to the directory containing the XSB executable.

`FLORADIR`: This variable should be set to the directory containing the $\mathcal{F}$LORA-*2* system.

`MYCLASSPATH`: This variable should include the jar files containing the API code, i.e., `.../java/flora2java.jar` and file `.../java/interprolog.jar`, plus the above `flogicbasicsExample` class. For instance, it can be set to `%CLASSPATH%;`*FLORADIR*`/java/flora2java.jar;`*FLORADIR*`/java/interprolog.jar;flogicbasicsExample`. For Linux and Mac, use ':' instead of ';' as a separator. As before, *FLORADIR* should be replaced with a proper value, as explained above.

- Some Java applications may employ additional shell variables. For instance, the program that uses the low-level API in Section 1.1 (in Step 2) has the line

  ```
  String fileName = System.getProperty("INPUT_FILE");
  ```

  which means that it expects the shell variable `INPUT_FILE` to be set. In this particular case, it expects that variable to have the address of the `flogic_basics.flr` $\mathcal{F}$LORA-*2* file, which it then loads. Therefore, the java command shown above would also need this parameter:

  ```
  -DINPUT_FILE=%INPUT_FILE%
  ```

  In general, one such additional parameter is needed for each property that the Java application queries using the `getProperty()` method.

1.4 Summary of the Variables Used by the Interface

The Java-$\mathcal{F}$LORA-*2* interface needs the following shell variables to be set:

- `JAVA_HOME` - this is normally set when you install Java. If not, set this variables manually.
- The following variables can be set by executing the scripts `flora_settings.bat` (Windows) or `flora_settings.sh` (Linux/Mac) located in `flora2/java/`:
  - `FLORADIR` — the path to the $\mathcal{F}$LORA-*2* installation directory.
  - `PROLOGDIR` — the path to the folder containing XSB executable.

If you need to set the above variables in some other way, look inside the above scripts to get the exact values these variables should be set to.

- The following variable is set by the scripts `unixVariables.sh` or `windowsVariables.bat`:
  - `JAVA_BIN` — the directory where Java executables are (java, javac).

  If you need to set this variable without running the aforesaid script, you need to know the correct value for that variable. The simplest way is to execute the script and then check the value of environment variable `JAVA_BIN`.

1.5 Building the Prepackaged Examples

Sample applications of the Java-$\mathcal{F}$LORA-*2* interface are found in the `java/API/examples` folder. To build the code for the interface, use the scripts `build.bat` or `build.sh` (or `build.bat` on Windows) in the `java/API` folder. To build the the examples, use the scripts `buildExample.sh` or `buildExample.bat` in the `java/API/examples` folder, whichever applies. For instance, to build the `flogicbasicsExample` example, use these commands on Linux, Mac, and other Unix-like systems:

```
cd examples
buildExample.sh flogicbasicsExample
```

On Windows, use this:

```
cd examples
buildExample.bat flogicbasicsExample
```

To run the demos, use the scripts `runExample.sh` or `runExample.bat` in the `java/API/examples` folder. For instance, to run the `flogicbasicsExample`, use this command on Linux, Mac, and the like:

```
runExample.sh flogicbasicsExample
```

On Windows, use this:

```
runExample.bat flogicbasicsExample
```

Chapter 2

Persistent Modules

by Vishal Chowdhary

This chapter describes a $\mathcal{F}$LORA-*2* package that enables persistent modules. A *persistent module* (abbr., PM) is like any other $\mathcal{F}$LORA-*2* module except that it is associated with a database. Any insertion or deletion of base facts in such a module results in a corresponding operation on the associated database. This data persists across $\mathcal{F}$LORA-*2* sessions, so the data that was present in such a module is restored when the system restarts and the module is reloaded.

2.1 PM Interface

A module becomes persistent by executing a statement that associates the module with an ODBC data source described by a DSN. To start using the module persistence feature, first load the following package into some module. For instance:

```
?- [persistentmodules>>pm].
```

The following API is available. Note that if you load `persistentmodules` into some other module, say `foo`, then `foo` should be used instead of `pm` in the examples below.

- `?- ?Module[attach(?DSN,?DB,?User,?Password)]@pm.`
  This action associates the data source described by an ODBC `DSN` with the module. If `?DB` is a variable then the database is taken from the DSN. If `?DB` is bound to an atomic string, then that particular database is used. Not all DBMSs support the operation of replacing the DSN's database at run time. For instance, MS Access or PostgresSQL do not. In this case, `?DB` must stay unbound or else an error will be issued. For other DBMS, such as MySQL, SQL Server, and Oracle, `?DB` can be bound.

  The `?User` and `?Password` must be bound to the user name and the password to be used to connect to the database.

  The database specified by the DSN must already exist and must be created by a previous call to the method `attachNew` described below. Otherwise, the operation is aborted.

The database used in the `attach` statement must not be accessed directly—only through the persistent modules interface. The above statement will create the necessary tables in the database, if they are not already present.

Note that the same database can be associated with several different modules. The package will not mix up the facts that belong to different modules.

- `?- ?Module[attachNew(?DSN,?DB,?User,?Password)]@pm.`
  Like `attach`, but a new database is created as specified by `?DSN`. If the same database already exists, an exception of the form `FLORA_DB_EXCEPTION(?ErrorMsg)` is thrown. (In a program, include `flora_exceptions.flh` to define `FLORA_DB_EXCEPTION`; in the shell, use the symbol `'_$flora_db_error'`.) This method creates all the necessary tables, if they are not already present.

  Note that this command works only with database systems that understand the SQL command `CREATE DATABASE`. For instance, MS Access does not support this command and will cause an error.

- `?- ?Module[detach]@pm.`
  Detaches the module from its database. The module is no longer persistent in the sense that subsequent changes are not reflected in any database. However, the earlier data is not lost. It stays in the database and the module can be reattached to that database.

- `?- ?Module[loadDB]@pm.`
  On re-associating a module with a database (i.e., when `?Module[attach(?DSN, ?DB,?User,?Password)]@pm` is called in a new $\mathcal{F}$LORA-*2* session), database facts previously associated with the module are loaded back into it. However, since the database may be large, $\mathcal{F}$LORA-*2* does not preload it into the main memory. Instead, facts are loaded on-demand. If it is desired to have all these facts in main memory at once, the user can execute the above command. If no previous association between the module and a database is found, an exception is thrown.

- `?- set_field_type(?Type)@pm.`
  By default, $\mathcal{F}$LORA-*2* creates tables with the VARCHAR field type because this is the only type that is accepted by all major database systems. However, ideally, the CLOB (character large object) type should be used because VARCHAR fields are limited to 4000-7000 characters, which is usually inadequate for most needs. Unfortunately, the different database systems differ in how they support CLOBs, so the above call is provided to let the user specify the field types that would be acceptable to the system(s) at hand. The call should be made right before `attachNew` is used. Examples:

```
?- set_field_type('TEXT DEFAULT NULL')@pm.   // MySQL, PostgreSQL
?- set_field_type('CLOB DEFAULT NULL')@pm.   // Oracle, DB2
```

Once a database is associated with the module, querying and insertion of the data into the module is done as in the case of regular (transient) modules. Therefore PM's provide a transparent and natural access to the database and every query or update may, in principle, involve a database operation. For example, a query like `?- ?D[dept -> ped]@StonyBrook.` may invoke the SQL `SELECT` operation if module `StonyBrook` is associated with a database. Similarly `insert{a[b -> c]@stonyBrook}` and `delete{a[e -> f]@stonyBrook}` will invoke

SQL INSERT and DELETE commands, respectively. Thus, PM's provide a high-level abstraction over the external database.

Note that if ?Module[loadDB]@pm has been previously executed, queries to a persistent module will *not* access the database since $\mathcal{F}$LORA-*2* will use its in-memory cache instead. However, insertion and deletion of facts in such a module will still cause database operations.

2.2 Examples

Consider the following scenario sequence of operations.

```
// Create new modules mod, db_mod1, db_mod2.
flora2 ?-  newmodulemod, newmoduledb_mod1, newmoduledb_mod2.
flora2 ?-  [persistentmodules>>pm].

// insert data into all three modules.
flora2 ?-  insertq(a)@mod,q(b)@mod,p(a,a)@mod.
flora2 ?-  insertp(a,a)@db_mod1, p(a,b)@db_mod1.
flora2 ?-  insertq(a)@db_mod2,q(b)@db_mod2,q(c)@db_mod2.

//  Associate modules db_mod1, db_mod2 with an existing database db
//  The data source is described by the DSN mydb.
flora2 ?-  db_mod1[attach(mydb,db,user,pwd)]@pm.
flora2 ?-  db_mod2[attach(mydb,db,user,pwd)]@pm.

// insert more data into db_mod2 and mod.
flora2 ?-  inserta(p(a,b,c),d)@db_mod2.
flora2 ?-  insertq(a)@mod,q(b)@mod,p(a,a)@mod.

// shut down the engine
flora2 ?-  \halt.
```

Restart the $\mathcal{F}$LORA-*2* system.

```
// Create the same modules again
flora2 ?-  newmodulemod, newmoduledb_mod1, newmoduledb_mod2.

// try to query the data in any of these modules.
flora2 ?-  q(?X)@mod.
No.

flora2 ?-  p(?X,?Y)@db_mod1.
No.

//  Attach the earlier database to db_mod1.
flora2 ?-  [persistentmodules>>pm].
```

```
flora2 ?-  db_mod1[attach(mydb,db,user,pwd)]@pm.

// try querying again...

// Module mod is still not associated with any database and nothing was
// inserted there even transiently, we have:
flora2 ?-  q(?X)@mod.
No.

// But the following query retrieves data from the database associated
// with db_mod1.
flora2 ?-  p(?X,?Y)@db_mod1.
?X = a,
?Y = a.

?X = a,
?Y = b.

Yes.

// Since db_mod2 was not re-attached to its database,
// it still has no data, and the query fails.
flora2 ?-  q(?X)@db_mod2.

No.
```

Chapter 3

SGML and XML Parsers for $\mathcal{F}$LORA-*2*

by Rohan Shirwaikar

This chapter documents the $\mathcal{F}$LORA-*2* package that provides SGML and XPath parsing capabilities. One set of predicates supports parsing SGML, XML, and HTML documents, which creates $\mathcal{F}$LORA-*2* objects in the user specified module. Other predicates evaluate XPath queries on XML documents and create $\mathcal{F}$LORA-*2* objects in user specified modules. The predicates make use of the **sgml** and **xpath** packages of XSB.

3.1 Summary of the Predicates

load_xml_structure/3	Parse XML data into $\mathcal{F}$LORA-*2* objects
load_sgml_structure/3	Parse SGML data into $\mathcal{F}$LORA-*2* objects
load_html_structure/3	Parse HTML data into $\mathcal{F}$LORA-*2* objects
load_xhtml_structure/3	Parse XHTML data into $\mathcal{F}$LORA-*2* objects
parse_xpath_xml/3	Apply an XPath expression to an XML document and parse the result
parse_xpath_sgml/4	Apply an XPath expression to an SGML document and parse the result
parse_xpath_html/4	Apply an XPath expression to an HTML document and parse the result
parse_xpath_xhtml/4	Apply an XPath expression to an XHTML document and parse the result

3.2 Description

This package supports parsing SGML, XML, and HTML documents, converting them to sets of $\mathcal{F}$LORA-*2* objects stored in user-specified $\mathcal{F}$LORA-*2* modules. The SGML interface provides facilities to parse input in the form of files, URLs and strings (Prolog atoms).

For example, the following XML snippet

```
<greeting id='1'>
```

```
<first ssn=111'>
John
</first>
</greeting>
```

will be converted into the following $\mathcal{F}$LORA-*2* objects:

```
o1[ element -> {o2}]
o2[ id -> '1']
o2[ first -> {o3}]
o2[ ssn -> '111']
```

To load the `flrxml` package, the user should run the following command at the $\mathcal{F}$LORA-*2* prompt.

```
  flora2 ?-  [flrxml].
```

The following predicates are provided by the `flrxml` package. They take SGML, XML, HTML, or XHTML documents and create the corresponding $\mathcal{F}$LORA-*2* objects as specified in Section 3.4.

load_sgml_structure(`+?Source,-?Warn,+?Module`)**@flrxml**

load_xml_structure(`+?Source,-?Warn,+?Module`)**@flrxml**

load_html_structure(`+?Source,-?Warn,+?Module`)**@flrxml**

load_xhtml_structure(`+?Source,-?Warn,+?Module`)**@flrxml**

The arguments to these predicates have the following meaning:

`?Source` is an input SGML, XML, HTML, or XHTML document. It is of the form `url(url)`, `file('file name')` or `string('document as a string')`. `?Module` is the name of the $\mathcal{F}$LORA-*2* module where the objects created by `flrxml` should be placed. `?Module` must be bound. `?Warn` gets bound to a list of warnings, if any are generated, or to an empty list.

3.3 XPath Support

XPath support is based on the XSB `xpath` package, which must be configured as explained in the XSB manual. This package, in turn, relies on the XML parser called `libxml2`. It comes with most Linux distributions and is also available for Windows, MacOS, and other Unix-based systems from `http://xmlsoft.org`. Note that both the library itself and the `.h` files of that library must be installed.

The following predicates are provided. They select parts of the input document using the provided XPath expression and create $\mathcal{F}$LORA-*2* objects as specified in Section 3.4. These predicates handle XML, SGML, HTML, and XHTML, respectively.

parse_xpath_xml(+?Source,+?XPath,+?NamespacePrefixList,-?Warn,+?Module)**@flrxml**

parse_xpath_sgml(+?Source,+?XPath,+?NamespacePrefixList,-?Warn,+?Module)**@flrxml**

parse_xpath_html(+?Source,+?XPath,+?NamespacePrefixList,-?Warn,+?Module)**@flrxml**

parse_xpath_xhtml(+?Source,+?XPath,+?NamespacePrefixList,-?Warn,+?Module)**@flrxml**

The arguments have the following meaning:

`Source` specifies the input document; this parameter has the same format as in `load_structure`. `?XPath` is an XPath expression specified as a Prolog atom. *?Module* is the module where the resulting $\mathcal{F}$LORA-*2* objects should be placed. `?Module` must be bound. *?Warn* gets bound to a list of warnings, if any are generated during the processing, or to an empty list.

`?NamespacePrefixList` is a string that looks like a space separated list of items of the form *`prefix = namespaceURL`*. This allows one to use namespace prefixes in the XPath query given in the *?XPath* parameter. For example if the XPath expression is `'/x:html/x:head/x:meta'` where `x` stands for `'http://www.w3.org/1999/xhtml'`, then this prefix would have to be defined in `?NamespacePrefixList`:

```
parse_xpath_xhtml(url('http://w3.org'),
                  '/x:html/x:head/x:meta',
                  'x=http://www.w3.org/1999/xhtml',
                  ?Warnings,
                  foomodule)@flrxml.
```

3.4 Mapping XML to $\mathcal{F}$LORA-*2*

This mapping is based on a proposal by Guizhen Yang. It specifies how an XML parser can construct the corresponding F-logic objects as a result of parsing an input XML document.

The basic idea is as follows:

- Elements in XML are modeled as objects in F-logic.
- Subelements in XML are modeled as multivalued attributes in F-logic.
- Element attributes in XML are modeled as single-valued attributes in F-logic. This complies to the XML 1.0 specification which states that an attribute be defined only once for each element in an XML document.

This proposal deals with data-intensive XML documents, i.e., those that don't rely on the interpretation of comments or processing instructions to carry data. However, this proposal does consider the modeling of mixed element content in which text and subelements are interspersed.

We do not consider modeling XML entities either, assuming no entity references or that all entity references in the original XML document are already resolved by an XML parser.

3.4.1 Object Ids

According to the XML specification 1.0, an XML element can be defined with an oid that is unique across the document. Such an oid can be provided as the value of an element attribute of type ID, although this attribute can be arbitrarily named. Since an XML element is modeled as an F-logic object, we would like the oid of this object to take the value of any ID attribute if such value is defined. Otherwise, the oid must be automatically generated by the system.

Sitting on top of the XML root element, there is an additional root object which just functions as the access point to the entire object hierarchy.

Note:

- Only a validating XML parser can decide whether an attribute is an ID attribute since such definition is provided by a DTD.
- The oids of leaf nodes, which have no outgoing edges and carry plain text, are just the string values.

For example, the following XML document:

```
<?xml version="1.0"?>
<person ssn="111-22-3333">
  <name first="John"                      (Example 1)
        last="Smith"/>
</person>
```

is represented by the following F-logic objects, provided we already know that `ssn` is an ID attribute:

```
o1[person -> {'111-22-3333'}].
'111-22-3333'[ssn -> '111-22-3333',
              name -> {o2}
             ].
o2[first -> 'John', last -> 'Smith'].
```

3.4.2 Text and Mixed Element Content

The content of an XML element may consist of plain text, or subelements interspersed with plain text, for instance:

```
<greeting>Hi! My name is <first>John</first> <last>Smith</last>.</greeting>
```

Each text segment is modeled in F-logic as if it were referred to by a special tag, named `'$text'`. Corresponding to each text segment, a node is created and is referred to from the parent node by an edge labeled `'$text'`. The text becomes the value of a special single-valued attribute, named `'$string'`, of the newly created node. Moreover, if the content of an XML element consists solely of plain text, then the text also becomes the value of a special single-valued attribute of this element, named `'$content'`, which is introduced for convenience purpose.

Therefore, the above XML segment would generate the following F-logic objects, where `o1, ..., o9` are new oids:

```
o1[greeting -> {o2}].
o2['$text' -> {o3},
   first -> {o4},
   '$text' -> {o6},
   last  -> {o7},
   '$text' -> {o9}
  ].
o3['$string' -> 'Hi! My name is '].
o4['$content' -> 'John', '$text' -> {o5}].
o5['$string' -> 'John'].
o6['$string' -> ' '].
o7['$content' -> 'Smith', '$text' -> {o8}].
o8['$string' -> 'Smith'}.
o9['$string' -> '.'].
```

Note:

- Handling of the whitespaces depends on the application. In the examples shown here, we assume that "insignificant" whitespaces have been omitted by the XML parser.
- `'$content'` can also be defined for every element. Its value is just the ASCII string as it appeared in the original document.

3.4.3 Multivalued XML Attributes

XML element attributes of type IDREFS are multivalued, in the sense that their value is a string consisting of one or more oids separated by whitespaces. Therefore, the value of such an attribute is a set. To stick to the convention that element attributes are modeled as single-valued attributes in F-logic, the value of an XML IDREFS attribute is represented as a list.

For example, the following XML segment:

```
<paper id="yk00" references="klw95 ckw91">
  <title>paper title</title>
</paper>
```

will generate the following F-logic atoms, assuming that the `reference` attribute is of type IDREFS:

```
yk00[references -> [klw95,ckw91],
       title -> {o1}
      ].
o1['$content' -> 'paper title', '$text' -> {o2}].
o2['$string' -> 'paper title'].
```

Note that since attribute definitions are provided by DTDs, only a validating XML parser can decide whether an element attribute has the type IDREFS. A non-validating parser should just output the attribute values as strings. The semantics of these strings is subject to further interpretation by the applications.

3.4.4 Ordering

XML is order-sensitive, since XML DTDs impose order among elements. For example, the following DTD element definition

```
  <!ELEMENT book (author+, title, ISBN?)>
```

states that the content of a `book` element consists of one or more `author` element, followed by one `title` element, followed by an optional `ISBN` element. Note that XML 1.0 does not prescribe any order among element attributes.

Preserving the order of elements can also be useful for translating F-logic objects back into an XML document.

A total order among the elements in an XML document should suffice to establish the order in which the elements appear sequentially in the XML document. In other words, such a total order should correspond to the order in which the element open tags appear. In addition, the ordering should also take into account the mixed element contents in which each text segment is referred to by a `'$text'` attribute.

```
<?xml version="1.0"?>
<person ssn="111-22-3333">
  <name>
    <first>John</first>                         (Example 2)
    <last>Smith</last>
  </name>
  <email>jsmith@abc.com</email>
</person>
```

For example, the XML document in Example 2 can be represented by the following tree, in which the integers enclosed by parentheses beside the nodes represent the order assigned to the elements:

```
                o1 (0)
                |
                | person
                |
           '111-22-3333' (1)
               / \
         name /   \ email
             /     \
        (2) o2     o7 (7)
           / \       \
    first /   \last   \ '$text'
         /     \       \
    (3) o3     o5 (5)  o8 (8)
        |      |       'jsmith@abc.com'
'$text' |      |'$text'
        |      |
    (4) o4     o6 (6)
       'John'  'Smith'
```

Pre-order traversal of the XML tree will generate the total order for the XML elements and text segments.

The ordering information that exists in XML documents is encoded as follows. For each node, we introduce a special single-valued attribute, named `'$order'`, which stores the order number.

```
<bibliography>
<paper id="sb97">
  <author>
    <first>John</first>
    <last>Smith</last>
  </author>
  <author>
    <first>David</first>
    <last>Brown</last>
  </author>                                  (Example 3)
</paper>

<paper id="s91">
  <author>
    <first>John</first>
    <last>Smith</last>
  </author>
</paper>
</bibliography>
```

3.4.5 More on Special Attributes

The following special attributes are all single-valued. Note that this proposal does not intend to eliminate redundant information that may exist across various attributes. It is up to the implementation to decide whether to generate these attributes extensionally or intentionally.

1. `'$tag'`
   For each node, `'$tag'` returns the unordered label of the edge pointing to this node. `'$tag'` can be defined as follows:

   ```
   ?O['$tag' -> ?Tag] :- ?[Tag -> ?O].
   ```

   Note that for a node representing a text segment, the value of its `'$tag'` attribute is `'$text'`.

2. `'$parent'`
   For each node, `'$parent'` returns the oid of the parent node.

3. `'$leftSibling'`
   For each node, `'$leftSibling'` returns the oid of the node appearing immediately before the current node. This attribute is not defined for the nodes without a left sibling.

4. `'$rightSibling'`
   For each node, `'$rightSibling'` returns the oid of the node appearing immediately after the current node. This attribute is not defined for those nodes without a right sibling.

5. `'$childrenNum'`
   For each node, `'$childrenNum'` returns the number of children including nodes representing text segments.

6. `'$childrenList'`
   For each node, `'$childrenList'` returns a list, which is ordered, of the oids of its children. Note that each text segment is also counted as a child node.

7. `'$child'(N)`
   For each node, `'$child'(N)` returns the `N`-th child, where `1` $\leq$ `N` and `N` $\leq$ `'$chidlrenNum'`.

8. `'$tagList'`
   For each node, `'$tagList'` returns an ordered list of the tags of its children. Note that each text segment is also counted as if it were enclosed by a `'$text'` tag.

9. `'$tag'(N)`
   For each node, `'$tag'(N)` returns the tag of the `N`-th child, where `1` $\leq$ `N` and `N` $\leq$ `'$chidlrenNum'`. This attribute can be defined as follows:

```
O['$tag'(N) -> Tag] :- O['$child'(N) -> V['$tag' -> Tag]].
```

Note: The aforesaid attribute `'$content'` can be defined for the nodes whose content is pure text as follows:

```
O['$content' -> String] :-
 O['$childrenNum' -> 1].text(1)['$string' -> String].
```