Persistence of ontologies, triple stores, programmatic access

Martin Ledvinka

May 2, 2014

1 Sesame

1.1 Introduction

Sesame is a framework for processing RDF data. Consists of two main components:

- Repository API – application access
- SAIL API – storage and inference

Sesame provides a triple store, an API for application developers to access the data, a SPARQL endpoint, REST API for remote access to the repository and a web-based management application.

Sesame (server-side) requires a Java servlet and JSP container, e. g. Apache Tomcat.

1.2 Triple Store

Data persistence is realized by implementations of Sesame's SAIL[1] API. There exist multiple choices of the storage implementation:

Memory Store

- Stores all data in main memory
- The data can by synced to disk
- Very fast for small and medium volume datasets

[1] Storage and Inference Layer

Native Store

- Data saved in special files on disk
- Slower than memory store, but scales much better for larger datasets
- Various indexes, all based on B-Trees
- By default SPOC and POSC indexes are used

RDBMS Store Bindings for RDBMS, e. g. MySQL Store, PostgreSQL Store.

Third-party Stores E. g. OWLIM – persistence and inference engine, supports OWL 2 QL (RL) reasoning.

1.3 Repository API

Java API for:

- Accessing remote Sesame repositories
- Creating and accessing local Sesame repositories (memory and native)
- Processing RDF data

1.4 Sesame Workbench

A web application for managing Sesame repositories. UI with functionality covering most of the HTTP API of Sesame.

2 Tasks

We will be using Firefox add-on RESTClient, an easy-to-use REST testing and debugging tool.

Task One Explore the Sesame architecture, see system documentation at http://openrdf.callimachus.net/sesame/2.7/docs/system.docbook?view.

Task Two Load data into your Sesame repository and try using the Sesame REST API to query and manipulate the data.

- Get the size of the repository
- Get all statements where FullProfessor0 is subject/object
- Remove all statements where Publication13 is subject/object
- Use a SPARQL query to get all predicates associated with FullProfessor0 (whether it is there as subject or object)

Task Three Explore the Sesame Workbench and try using it to accomplish the same results as in Task Two.

3 Reference

- http://openrdf.callimachus.net/sesame/2.7/docs/system.docbook?view
- http://openrdf.callimachus.net/sesame/2.7/docs/users.docbook?view
- http://openrdf.callimachus.net/sesame/2.7/apidocs/index.html