

A comparison of object-triple mapping libraries

Martin Ledvinka^{a,*} and Petr Křemen^a

^a *Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague, Technická 2, 166 27 Prague 6 - Dejvice, Czech Republic*

E-mails: martin.ledvinka@fel.cvut.cz, petr.kremen@fel.cvut.cz

Abstract. Domain-independent information systems like ontology editors provide only limited usability for non-experts when domain-specific linked data need to be created. On the contrary, domain-specific applications require adequate architecture for data authoring and validation, typically using the object-oriented paradigm. So far, several software libraries mapping the RDF model (representing linked data) to the object model have been introduced in the literature. In this paper, we develop a novel framework for comparison of object-triple mapping solutions in terms of features and performance. For feature comparison, we designed a set of qualitative criteria reflecting object-oriented application developer's needs. For the performance comparison, we introduce a benchmark based on a real-world information system that we implemented using one of the compared object-triple mapping (OTM) solutions – JOPA. We present a detailed evaluation of a selected set of OTM libraries, showing how they differ in terms of features. We further narrow down the set of selected OTM libraries to contain only RDF4J-compatible ones and use the benchmark to measure their time and memory efficiency.

Keywords: Object-triple Mapping, Object-ontological Mapping, Object-oriented Programming, RDF, Benchmark

1. Introduction

The idea of the *Semantic Web* [1] might not have seen such an explosive success as the original *Web*, but with technology leaders like Facebook¹, Google², IBM [2] or Microsoft³, as well as emerging national linked data [3] platforms⁴, adopting its principles, it seems it has finally taken root. Semantically meaningful data which allow to *infer* implicit knowledge, described, stored and queried using standardized languages and interconnected via global identifiers can provide a huge benefit for application users on many levels.

1.1. Motivating scenario

Let's consider a fictive national civil aviation authority (CAA), which is a public governmental body and is obliged to publish its data as open data. The CAA decides to develop a safety management system for the oversight of aviation organizations under its jurisdiction, involving safety occurrence reporting, safety issues definition, and aviation organization performance dashboard. The CAA can then concentrate on suspicious patterns identified by the system (e.g., from frequently occurring events) during its inspections at the aviation organizations or during safety guidelines preparation. To establish data and knowledge sharing between the CAA and the aviation organizations, CAA decides to share the safety data as 5-star Linked Data [4]. To achieve this, the system is designed on top of a proper ontological conceptualization of the domain. This allows capturing safety occurrence reports with deep insight into the actual safety situation, allowing to model causal dependencies between safety occurrences, describe event par-

* Corresponding author. E-mail: martin.ledvinka@fel.cvut.cz.

¹ <http://ogp.me>, accessed 2018-03-01

² <https://tinyurl.com/g-knowledge-graph>, accessed 2018-03-01

³ <https://tinyurl.com/ms-concept-graph>, accessed 2018-03-01.

⁴ e.g. the Czech linked open data cloud <https://linked.opendata.cz>, the Austrian one <https://www.data.gv.at/linked-data> (both accessed 2018-04-04).

ticipants, use rich typologies of occurrences, aircraft types, etc. Furthermore, integration of other available Linked Data sources, e.g., a register of aviation companies (airlines, airports etc.) and aircraft helps to reveal problematic aviation organizations. A subset of the described functionality has been implemented in the SISel system [5]. Another system with related functionality aimed at the spacecraft accident investigation domain was built at NASA by Carvalho et al. [6, 7].

To be usable by non-experts, the system is domain-specific, not a generic ontology editor/browser. No matter how the user interface is implemented, it is supposed to communicate with the backend through a well-designed domain-specific API. Business logic implementing such API requires efficient access to the underlying data. Using common Semantic Web libraries like Jena [8] or RDF4J [9] for such access ends up with a large amount of boilerplate code, as these libraries work on the level of triples/statements/axioms and do not provide frame-based knowledge structures. The solution is to allow developers to use the well-known object-oriented paradigm, which is well supported in IDEs and which provides compile-time checking of the object model structure and usage. However, to match objects to triples, one needs an *object-triple mapping* (OTM) that defines the contract between the application and the underlying knowledge structure. Many libraries trying to offer such a contract exist. However, these libraries differ a lot in their capabilities as well as efficiency, neither of which is systematically documented. In this survey paper, we provide a framework for comparing various OTM solutions in terms of features and performance. In addition, we use this framework and try to systematize the existing OTM approaches with the aim of helping developers to choose the most suitable one for their use case.

1.2. Contribution

The fundamental question of this paper is: “Which object-triple mapping solution is suitable for creating and managing object-oriented applications backed by an RDF triple store?”. The actual contribution of this paper can be split into three particular goals:

1. Select a set of *qualitative criteria* which can be used to compare object-triple mapping library features relevant for object-oriented applications.
2. Design and develop a benchmark for performance comparison of object-triple mapping libraries. This benchmark should be easy to use

and require as little effort as possible to accommodate a new library into the comparison.

3. Compare object-triple mapping solutions compatible with a selected storage in terms of their features and performance.

The remainder of this paper is structured as follows. Section 2 presents the necessary background on the RDF language, OTM systems, and Java performance benchmarks. Section 3 reviews existing approaches to comparing OTM libraries and benchmarking in the Semantic Web world in general. Section 4 presents the framework designed in this work. Section 5 introduces the OTM libraries chosen for the comparison and discusses the reasons for their selection. Sections 6 and 7 actually compare the selected libraries using the framework designed in Section 4. The paper is concluded in Section 8. Appendices A.1 and A.2 contain full reports of time performance, resp. scalability comparison.

2. Background

The fundamental standard for the Semantic Web is the Resource Description Framework (RDF) [10]. It is a data modeling language built upon the notion of *statements* about *resources*. These statements consist of three parts, the *subject* of the description, the *predicate* describing the subject, and the *object*, i.e., the predicate value. Such statements – *triples* – represent elements of a labeled directed graph, an *RDF graph*, where the nodes are resources (IRIs or blank nodes) or literal values (in the roles of subjects or objects) and the edges are properties (in the role of predicates) connecting them. RDF can be serialized in many formats, including RDF/XML, Turtle or N-triples. Taking an RDF graph and a set of *named graphs* (RDF graphs identified by IRIs), we get an *RDF dataset*.

The expressive power of RDF is relatively small. To be able to create an *ontology* (a formal, shared conceptualization) of a domain, more expressive languages like RDF Schema (RDFS) [11], OWL [12] and OWL 2 [13] were introduced. They extend RDF with constructs like classes, domains and ranges of properties, and class or property restrictions. A simplistic example of an RDFS ontological schema can be seen in Figure 1, its serialization in Turtle [14] is then displayed in Listing 1⁵. An important feature of the Se-

⁵Prefixes introduced in listings are reused in subsequent listings and figures as well.




Fig. 1. RDF Graph representing a simple RDFS ontology. Prefixes used in this Figure are defined in Listing 1

semantic Web is also that RDF is used to represent both the schema and the actual data.

Listing 1: Turtle serialization of an RDFS schema.

```

@prefix rdf:
<http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs:
<http://www.w3.org/2000/01/rdf-schema#> .
@prefix xsd:
<http://www.w3.org/2001/XMLSchema#> .
@prefix doc: <http://onto.fel.cvut.cz/ontologies/
documentation/>.
@prefix as:
<http://onto.fel.cvut.cz/ontologies/a-s/> .
@prefix ufo:
<http://onto.fel.cvut.cz/ontologies/ufo/> .

doc:documents
  a rdf:Property ;
  rdfs:range as:occurrence ;
  rdfs:domain doc:occurrence_report .

doc:has_file_number
  a rdf:Property ;
  rdfs:range xsd:long .

as:occurrence
  a rdfs:Class ;
  rdfs:subClassOf ufo:event .

doc:occurrence_report a rdfs:Class .
ufo:event a rdfs:Class .

```

The SPARQL Query Language (SPARQL) [15] and SPARQL Update [16] are standardized languages used to query and manipulate RDF data. A SPARQL query example is shown in Listing 2; it selects all classes and, if present, also their immediate super types.

The primary focus of this paper (as well as most of the discussed OTM solutions) is on RDFS. However, OWL (2) is referenced in several places to provide further context.

Listing 2: A SPARQL query example.

```

SELECT DISTINCT ?cls ?supcls WHERE {
  ?cls a rdfs:Class .
  OPTIONAL {
    ?cls rdfs:subClassOf ?supcls .
  }
}

```

2.1. Accessing semantic data

Ontological data are typically stored in dedicated semantic databases or *triple stores*, such as RDF4J [9], GraphDB [17] or Virtuoso [18]. To be able to make use of these data, an application needs means of accessing the underlying triple store. Unfortunately, there is no common standard for accessing a semantic database (like ODBC [19] for relational database access). This lack of standardization gave rise to multiple approaches and libraries that can be basically split into two categories (as discussed in [20]):

Domain-independent libraries like Jena [8], OWL API [21], or RDF4J (former Sesame API) [9]. Such libraries are suitable for generic tools like ontology editors or vocabulary explorers that do not make any assumptions about the ontology behind the data the application works with.

Domain-specific libraries like AliBaba [22], Empire [23], or JOPA [24] are examples of libraries which employ some form of *object-triple* (OTM) (or *object-ontological* (OOM)) *mapping* to map semantic data to domain-specific objects, binding the application with some assumption about the form of data.

Domain-specific libraries are more suitable for creating object-oriented applications. Such applications express the domain model in the form of object classes, and their properties (attributes or relationships). OTM, in this case, allows to automatically map the notions from the object-oriented realm to the semantic data model and vice versa. Domain-independent libraries perform no such thing. They treat the underlying data at the same level as the RDF language, i.e., at the level of statements, without any assumption about their meaning. Therefore, a domain-specific application attempting to use such libraries would either need to treat its data in the same manner, which makes it extremely difficult to understand, develop, and maintain; or the developer would end up writing their own version of an object-triple mapping library. Either way, it is more favorable to leverage dedicated object-triple mapping

Table 1

Simplified mapping between RDF(S) terms and object model artifacts

Ontology	Object
RDFS class	Entity class
RDF property	Attribute
RDFS class instance	Entity (object)

solutions. In addition, some OTM libraries allow accessing various data sources in a unified way, support transactional processing and compile-time checking of correct model usage. Table 1 shows a simplified mapping between RDF(S) terms and object model artifacts considered by the majority of the OTM solutions. We shall use these notions throughout this paper.

An analogous division where *indirect* and *direct* models correspond to the domain-independent and domain-specific approaches respectively is used in [25]. The authors argue that direct models are more suitable for domain-specific applications and propose a hybrid approach where parts of the application model are domain independent and parts are domain specific.

The difference between domain-independent and domain-specific libraries is similar to the difference between JDBC [26] and JPA [27] in the relational database world. JDBC is a row-based access technology, which requires a lot of boilerplate code. In addition, not all databases fully adhere to the SQL⁶ standard, so using JDBC, which is based on SQL, may bring compatibility issues. JPA, on the other hand, provides *object-relational mapping* (ORM) and makes access to various databases transparent.

There exists a number of OTM libraries. However, their user base is still relatively small, which is reflected in the small amount of information available about them. Some libraries are tied to a specific underlying storage, others' APIs do not reflect the specifics of semantic data access. Ultimately, one would like to be able to determine which of the libraries is the most suitable for his/her object-oriented application.

2.2. Performance benchmarking in Java

As was already stated, the developed comparison framework consists of a set of qualitative criteria and a performance benchmark. The benchmark is written in Java, which is hardly a surprise since Java is the most

⁶Structured Query Language, a data query, definition, manipulation and control language for relational databases.

prominent language in the Semantic Web community. Therefore, approaches to performance testing in Java were also reviewed.

Java-based performance testing has certain specifics which are given by the nature of the Java platform where each program is compiled into an intermediate language – the *bytecode* – and interpreted at runtime by the *Java virtual machine* (JVM). The JVM is able to optimize frequently executed parts of a program. These optimizations occur during execution of the application. The JVM also automatically manages application memory via a mechanism called *garbage collection* which automatically frees memory occupied by objects that are no longer referenced. Of course, if the application forgets to discard objects it does not use anymore – it has a so-called *memory leak* – it may eventually run out of free memory. Although garbage collection is nowadays extremely efficient, it still requires CPU time and may influence the application performance. Java benchmarks thus usually take into account application throughput and its dependence on the amount of available memory (*heap size*). The best known benchmarks in this area are SPEC [28] and DaCapo [29]. DaCapo measures application throughput w.r.t. heap size. In addition, it samples heap occupation and composition and new object allocation rate. The authors of [30] count garbage collection events and the time JVM spends executing them.

Generally, when benchmarking Java applications, it is necessary to execute a *warm-up* period which allows the JVM to perform code optimizations. Otherwise, the results could register a significant speedup after the initial unoptimized phase. Georges et al. [31] discuss further practices for executing Java application benchmarks and interpreting their results, introducing a statistically rigorous methodology for Java benchmarking. Their methodology was used in the experiments in this work.

3. Related work

This section overviews existing related approaches. First, we present existing object-triple mapping library comparisons (both feature and performance-wise) and then take a look at established storage benchmarks and discuss how our benchmark framework complements them.

There are very few comparisons of OTM libraries. Most benchmarks and comparisons concentrate on the underlying storage and its performance. While this is

certainly a very important aspect, one needs to be able to make a qualified decision when choosing from a set of object-triple mapping libraries. Holanda et al. [32] provide a performance comparison between their solution – JOINT-DE – and AliBaba. However, their benchmark uses a minimalistic model (create operation works with a single entity with one string attribute) and no other library is taken into account. Cristofaro [33] provides a short feature comparison of various OTM libraries supporting the Virtuoso storage server [18]. A more recent and elaborate feature-based comparison can be found in [34]. It compares a set of selected libraries from the point of development activity, documentation, ease of use and querying capabilities. In [35], authors of the well-known LUBM benchmark [36] do not present any actual comparison but they introduce a set of requirements a Semantic Web knowledge base system benchmark should fulfill. Although they again target mainly semantic database benchmarks, the criteria defined in their article apply to a large extent to this work as well and their satisfaction will be shown in Section 4.

The number of benchmarks comparing storage systems is, on the other hand, relatively large⁷. The best known benchmarks in this area are the Berlin SPARQL Benchmark [37], LUBM [36] and UOBM [38]. These benchmarks evaluate the performance of the storage engines and, in case of the more expressive LUBM and UOBM, also allow to compare the engines in terms of their inference capabilities. Our work in [39] represents a step towards the application-oriented usage because it leverages the UOBM model and data and uses a set of queries specifically designed to simulate the behavior of OTM libraries. All these benchmarks use the SPARQL endpoint provided by the evaluated system, which makes them easy to apply to various storage technologies. The benchmark developed here concentrates more on the overhead OTM libraries represent on top of the storage execution time and how effectively they use the storage. That is why, to eliminate its influence on the results, only one underlying repository was used in the evaluation in Section 7.

4. Design of the comparison framework

This section introduces the comparison framework which consists of two parts:

1. A set of criteria used to evaluate the features supported by the compared libraries.
2. A performance benchmark.

4.1. Feature criteria

We have designed a set of criteria, which can be used to evaluate an OTM library's suitability in a particular use case. The criteria can be split into three main categories:

General (GC) General criteria are based on the principles known from application development and relevant object-relational mapping features.

Ontology-specific (OC) These criteria take into account specific features of the Semantic Web language stack based on RDF and its design. They are not specific to OTM libraries and could be applied to libraries used to access semantic data in general.

Mapping (MC) Mapping criteria concern important techniques used for the object-triple mapping. They are motivated by the differences of ontologies on the one side and the object-model on the other [40].

The selection is based on existing knowledge from other domains (e.g., ORM), literature and our own experience with developing ontology-based applications. We intended to incorporate also other ontology-based information systems but we were not able to find publicly available source code of any ontology-based application which uses an OTM library. For each criterion, we describe its genesis, as well as the condition under which we consider the criterion to be satisfied.

GCI – Transactions Transactional processing, i.e., splitting work into individual, indivisible operations, is one of the fundamental paradigms of computer science [41]. It originated in databases but, due to its universally applicable principles, it is used throughout information systems with various levels of granularity. Nevertheless, support for transactions is one of the main features an OTM library has to provide, be it using an internal transaction engine or relying on the underlying storage.

Transactions are characterized by the *ACID* acronym, which stands for *Atomicity* of the transaction, eventual *Consistency* of data, *Isolation* of concurrent transactions and *Durability* of the transaction results. Given the complex nature of transaction isolation where sev-

⁷See for example the list at <https://www.w3.org/wiki/RdfStoreBenchmarking>, accessed 2018-11-21.

eral isolation levels exist⁸, and the lack of its treatment in all evaluated tools, our goal when examining transactional support in the OTM libraries is to evaluate their implementation of ACD.

Fully satisfied if: The library supports ACD user-controlled transactions.

GC2 – Storage access variability Storage access variability refers to the library’s ability to connect to semantic storages of various vendors. While application access to relational databases is standardized in two ways – the *Open Database Connectivity* (ODBC) [19] is a platform-independent database access API; the *Java Database Connectivity* (JDBC) [26] is a Java-based database access API (a JDBC driver can use an ODBC driver internally) – the semantic world lacks a corresponding standard. For instance, although RDF4J-based access is supported by most of the industry’s biggest repository developers, it is not possible to use it to connect to a Jena SDB store, which is backed by a relational database. To allow access to different storages, the OTM library has to account for this either by having modules for access to various storages or by defining an API whose implementations would provide access to the repositories. Another possibility, although probably not very efficient, could be accessing the storage using a SPARQL [42] or *Linked Data Platform* [43] endpoint.

Fully satisfied if: The library supports connecting to a triple store by means of at least two different APIs (e.g., Jena, RDF4J) or via a generic SPARQL endpoint API.

GC3 – Query result mapping Besides reading individual instances with known identifiers, it is often necessary to read a set of instances corresponding to some search criteria (for example, read all reports with the given severity assessment which document events that occurred during the last month). Naturally, mapping of SPARQL query results to objects is expected, as SPARQL is the standard query language for the Semantic Web, but other languages like SeRQL (a Sesame query language) can be supported. A similar feature is supported by JPA [27] where SQL results can be mapped to entities or dedicated Java objects.

Fully satisfied if: SPARQL query results can be declaratively mapped to the object model (entities).

GC4 – Object-level query language Related to *GC3* is the question of query languages supported by the OTM library. While SPARQL is adequate in most situations, it can be cumbersome to work with, especially given the fact that one has to deal with IRIs and prefixes. When querying the data from an object-oriented application, it is convenient to be able to leverage the object model in the queries. This is again supported by JPA [27], where one can use the *Java Persistence Query Language* (JPQL) – a query language with SQL-like syntax but exploiting classes, attributes, and relationships instead of tables and columns. Furthermore, it also defines the *Criteria API* which allows one to construct queries dynamically at runtime using builder objects. Criteria API is especially suitable when a large number of optional filtering criteria can be used in the query and juggling with string concatenation would be difficult and error-prone. In addition, Criteria API provides the benefit of compile-time syntax checking, because it uses regular Java classes and methods.

Prototypical solutions already exist in the Semantic Web community. For instance, Hillairet et al. [44] use the Hibernate Query Language, an implementation of the standard Object Query Language [45] (OQL)⁹ used by one of the most popular ORM vendors, to translate queries to/from SPARQL. Stadler and Lehmann [46], on the other hand, present an engine rewriting Criteria API queries to SPARQL.

Fully satisfied if: The library provides an object model-based query language which can be used to access the data (e.g., like JPQL).

GC5 – Detached objects As pointed out in [32], applications need to be able to *detach* persistent objects from their connection to the storage in order to work with them, for instance, when the entity is passed up through web application layers and transferred over the network in response to a client request or stored in an application-level cache. The opposite (an object which is always *managed*) can lead to an excessive amount of retained connections to the storage, blocking the resources of the machine. The difference between detached and managed entities is indeed in that managed entities are tracked by the persistence library which watches for changes in these objects.

Fully satisfied if: The persistence library allows to attach/detach objects to/from the repository connection.

⁸E.g. <https://tinyurl.com/transaction-isolation>, accessed 2018-11-21.

⁹JPQL is another implementation of a subset of OQL.

GC6 – Code/ontology generator Setting up a database schema or an object model may seem like a one-time work but as the application evolves, so does often the schema. Keeping both the ontology schema and the object model in sync automatically can spare the developer a decent amount of time and bugs (e.g., a typo in an IRI). For this reason, JPA allows one to generate database schema from an object model [27]. In addition, many JPA implementations support also the converse – generating an object model from database tables. There also exist libraries like Owl2Java [47] which allow generating an object model from an ontology. To this end, three approaches to object model vs ontology schema transformation are discussed in [48] – manual mapping of an object model, its automatic generation from an ontological schema and a hybrid approach, which the authors deem most useful based on its correctness/efficiency ratio. In the hybrid scenario, a basic model is generated automatically and the developer then fine tunes it manually. Conversely, one could start with an object model, generate an ontology schema from it and continue extending the schema further.

Fully satisfied if: The library provides a generator to synchronize the object model with the ontology.

OC1 – Explicit inference treatment An RDFS repository may respond to a query with asserted and inferred statements. While they are indistinguishable from the formal point of view [49], this holds only for the retrieval scenario. Statement removal (possibly as part of an update) works differently: asserted triples can be directly deleted but the existence of inferred statements is only implied by the presence of other statements. Consider triples

```
@prefix skos:
<http://www.w3.org/2004/02/skos/core#> .
@prefix ex:
<http://onto.fel.cvut.cz/ontologies/example/> .

ex:falcon-1 skos:prefLabel "Falcon" .
skos:prefLabel rdfs:subPropertyOf rdfs:label .
```

and an application which maps values of `rdfs:label`. If an OTM library does not differentiate asserted and inferred statements, it will let the application update `ex:falcon-1`'s label. However, since the application works with `rdfs:label` and its value is inferred, updating it will actually not remove the asserted `skos:prefLabel` value. Instead, it will just assert a value of `rdfs:label`, leading to two labels asserted for the same object:

```
ex:falcon-1 skos:prefLabel "Falcon" .
ex:falcon-1 rdfs:label "Updated Falcon" .
```

This is hardly the expected behavior. An OTM library has to deal with such situations, e.g., by marking inferred attributes as read-only or by keeping track of inferred values and indicating an error whenever the application attempts to remove or modify such a value. This issue would be even more important if the repository supported SPIN [50] rules or OWL, which allow richer inference.

Fully satisfied if: The library gives its users explicit control over whether they are working with asserted or inferred statements.

OC2 – Named graphs Named graphs [51] are an important feature of RDF, as they allow to split RDF data into meaningful and identifiable parts inside an RDF dataset. Although there exist different semantics for treatment of named graphs, triple stores mostly adopt the one where the dataset's default graph represents a union or merge of its named graphs [52]. This strategy makes named graphs suitable for logical structuring of data. Consider a company repository containing information about projects, business contracts and employees. It is sensible to let the employees occupy a different named graph than projects, possibly together with a different degree of access control.

Fully satisfied if: The OTM library provides a way to access different graphs (default graph as well as named graphs) of an RDF dataset.

OC3 – Automatic provenance management RDF allows to record *provenance* information about the data using the RDF *reification* vocabulary [11] (in addition, an OWL 2-compliant PROV-O ontology [53] has been introduced by W3C). This approach, although not without flaws [54], provides an interesting alternative to *auditing* in JPA, which is not standardized and is done in an ad hoc manner. Quasthoff and Meinel [55] introduce a prototype able to generate provenance data in RDF processing. More specifically, it connects newly created statements to triples upon which they are based. An example of benefits of automatic provenance tracking can be formulated in terms of the benchmark model introduced later in Section 4.2.1 – instead of manually assigning the last editor or author of a report, the OTM could set it based on user session information available in the application. It should be possible to configure for which entities or which attributes the provenance data should be generated and how to retrieve relevant information (e.g., user).

Fully satisfied if: The library provides a way to automatically generate provenance statements about the data it created using a standard vocabulary. The granularity of the provenance data and its selection may be configurable, e.g., at entity or statement level.

MC1 – Inheritance mapping Class hierarchies represent a major component of every domain conceptualization. Although RDFS does not support advanced class declaration expressions like disjointness, intersection, or union, which are part of OWL [12], it allows building class hierarchies using the RDFS `subClassOf` property. When mapping class hierarchies to an object model, the developer may encounter one technical issue: some programming languages, including Java, do not support multiple inheritance on the class level, whereas RDFS does. This is often resolved using *interfaces*, which do not have this restriction.

Mapping ontological inheritance to an object model may bring subtle conceptual issues. For example, take an ontology containing the following statements:

```
ex:Person rdfs:type rdfs:Class .
ex:Employee rdfs:type rdfs:Class ;
    rdfs:subClassOf ex:Person .
ex:Student rdfs:type rdfs:Class ;
    rdfs:subClassOf ex:Person .
ex:a rdfs:type ex:Employee, ex:Student .
```

and a corresponding object model, i.e., entity classes `Person`¹⁰, `Employee`, `Student`; where `Employee` and `Student` are subclasses of `Person`. An application that would attempt to load an entity `a` (mapped to an RDFS class instance `ex:a`) as an instance of `Person`, leveraging the mechanism of *polymorphism*¹¹ supported by object-oriented languages, might experience the following three problematic scenarios:

1. If inference is not enabled, `ex:a` is not an instance of `ex:Person` so `a` cannot be loaded as an instance of `Person`.
2. If inference is enabled and the application attempts to remove the loaded entity `a` of type `Person`, it will in effect attempt to remove an inferred class assertion `ex:a rdfs:type ex:Person`.
3. The OTM library might attempt to look for an explicit class assertion in the same RDFS class

¹⁰We shall use the sans serif font to express object language constructs throughout this work.

¹¹Simply put, being able to treat instances of different classes using the same interface.

inheritance hierarchy. However, that would lead to an ambiguous result because both `ex:a rdfs:type ex:Employee` and `ex:a rdfs:type ex:Student` are explicit class assertions and they are at the same level in the hierarchy.

Thus, OTM libraries need to support inheritance mapping due to its ubiquity, yet attention has to be paid to its treatment.

Fully satisfied if: The library supports both single-class and multiple-class inheritance mapping.

MC2 – Unmapped data access An application's object model may represent only a part of the domain ontology. Similarly, one may start building an application using a simplified object model and gradually extend its coverage of the underlying ontology. Yet, there exist situations in which access to the portion of the ontology not mapped by the object model may be required. Consider an aviation safety occurrence reporting application (will be discussed in the next section) based on an aviation safety occurrence reporting ontology which extends a generic safety reporting ontology [56]. The generic safety ontology can be used for other high-risk industries such as power engineering, railroad transportation, etc. There may exist an application used, for example, by a national safety board, which allows browsing safety occurrence reports from different industries using an object model based on the generic ontology. Yet, the application could provide access to the industry-specific property values without them being mapped by the object model.

Another example of the benefits of accessing unmapped ontological data is ontology-based classification. In the aviation safety reporting tool, events can be classified using RDFS classes from a predefined vocabulary [57]. However, the application does not need to map these classes in order to use them if the OTM library allows access to these unmapped types.

Fully satisfied if: The library allows to access data which are not mapped by the OTM.

MC3 – RDF collections and containers By default, all RDF properties are plural. In addition, RDF does not impose any kind of ordering restriction on their values, giving them effectively the semantics of a mathematical *set*. RDF defines the notions of *containers* and *collections*, which allow to represent groups of values with different kinds of semantics. For instance, an RDF container `rdfs:Seq` represents an ordered unbounded sequence of elements [11]. On the other hand,

an RDF collection `rdf:List` is a sequence of elements with a known length [11]. Lists (or sequences of data in general) are ubiquitous in programming, so it is important to support mapping of such RDF structures.

Fully satisfied if: The library allows to manage both RDF collections and containers.

4.2. Performance benchmark

To provide a full picture of the object-triple mapping landscape, we complement the (qualitative) feature comparison with a (quantitative) performance comparison. To make the comparison more usable and repeatable, it was decided to design a benchmark framework which could be used to compare as large variety of OTM libraries as possible.

As [35] and [58] point out, a benchmark should aim at simulating real-world conditions as much as possible, especially avoiding microbenchmarks, which are easy to get wrong. Luckily, we had an experience of developing ontology-based safety information solutions for the aviation domain in the INBAS¹² and B-INBAS¹³ [5] projects. A part of these solutions is a safety occurrence reporting tool (RT)¹⁴. The RT is a web application written in Java, backed by ontologies and ontological data and its functionality and model were used as a base for the benchmark framework. One might ask why not reuse the object model and data generator from one of the established benchmarks. There are several reasons for this decision. As far as the generator is concerned, the main goal was to restrict unwanted randomization. We wanted the benchmark to generate precise numbers of instances so that all the libraries would have the same conditions. Also, the generator had to create object instances, whereas, for example, the LUBM generator outputs RDF. The model is based on a real-world application, it is small enough to enable adaptation to new benchmarked libraries, yet it exercises most of the common mapping features (described in Section 4.2.1).

4.2.1. Benchmark model

We took an excerpt of the RT model and modified it by removing some of the attributes which use features already present in other parts of the model. It is built upon a conceptualization of the aviation domain

called the *Aviation Safety Ontology* [59] which is based on the *Unified Foundational Ontology* (UFO) [60]. RDFS serialization of the model ontology is visualized in Figure 2. The model primarily concerns occurrence reports which document safety occurrences. Each occurrence may consist of several events structured in a part-of hierarchy. Events may be factors of other events, for example, one event may cause or mitigate another event. Reports can reference resources, which are typically file and document attachments. A question-answer tree represents a dynamically generated form structure. Basic provenance data (author, last editor, date of creation and last edit) are stored for each report as well.

The goal of the benchmark is to exercise both common mapping features known from JPA and features specific to OTM revealed in our experience with building ontology-based applications. Therefore, the model contains both singular and plural relationships, properties with literal values, inheritance, and references to resources not mapped by the object model (the `ufo:has_event_type` property references an external vocabulary, a `doc:logical_record` references an unmapped `doc:question-doc:answer` tree). The last point illustrates how the object model can be connected to a larger domain model which exceeds the application's area of operation. The corresponding object model is illustrated in Figure 3.

We conclude this part with a remark regarding instance identifiers. The identifier attribute was not included in declarations of the entity classes in the model because various libraries use various types for the identifier, e.g., String, URI, or a vendor-specific class. Moreover, since some of the libraries do not make a mapped object's identifier (the RDF resource) accessible at all, a *key* property was added which explicitly specifies a unique identifier of the instance. The notion of a unique key identifying an instance was formalized in the OWL 2 specification [13] in the form of `owl:hasKey` axioms. In contrast to them, our key has no formal semantics and its uniqueness is purely conventional. It is used to find matching objects when verifying operation results.

4.2.2. Benchmark operations

The set of operations executed by the benchmark is supposed to represent a common domain-specific application functions. Therefore, the scenarios include a basic set of create, retrieve, update and delete (CRUD) operations, plus a batch persist and a find all query. The CRUD operations represent a typical form of op-

¹²<https://www.inbas.cz>, accessed 2018-11-21.

¹³<https://www.inbas.cz/web/binbas>, accessed 2018-11-21.

¹⁴Source code is available at <https://github.com/kbss-cvut/reporting-tool>, accessed 2018-11-21.

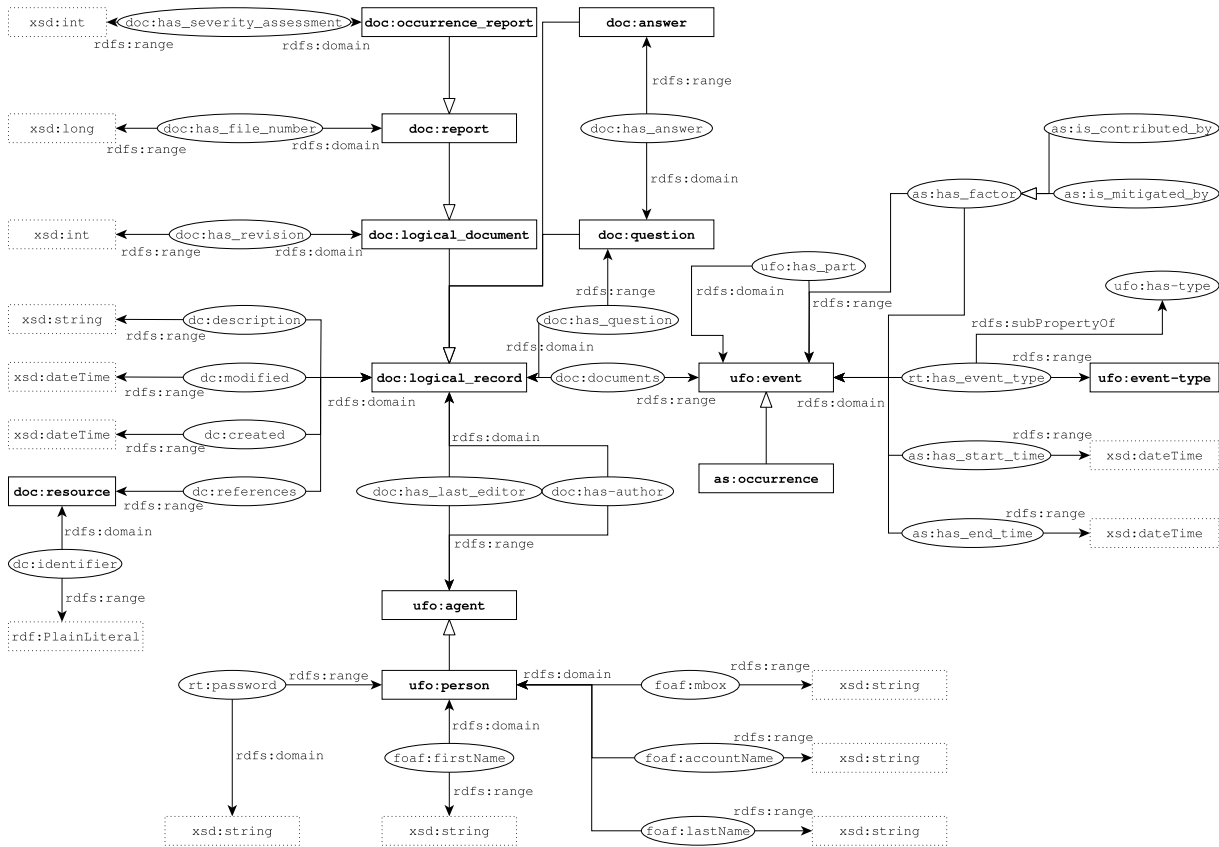


Fig. 2. Benchmark ontology visualization. Rectangles with solid line represent RDFS classes, rectangles with dotted line denote literal datatypes and ellipses are RDF properties used as RDF triple subjects/object. Each edge represents an RDF triple with a source/target node label representing the triple subject/object and the edge label representing the triple predicate. Unlabeled edges ended with a hollow arrow denote triples with the `rdfs:subClassOf` predicate. Prefixes `ufo`, `as`, `doc`, `xsd`, `rdf` and `rdfs` have been introduced in previous listings. Prefix `dc` denotes the Dublin Core namespace <http://purl.org/dc/terms/>, prefix `foaf` denotes the FOAF namespace <http://xmlns.com/foaf/0.1/> and `rt` denotes an application specific namespace <http://onto.fel.cvut.cz/ontologies/reporting-tool>.

eration of a business application, where data are persisted, updated, queried and occasionally deleted. The find all query is another example where the user requests a list of all instances of some type, e.g., for a table of all reports in the system. The batch persist, on the other hand, may represent a larger amount of data being imported by an application at once.

Each operation consists of three phases; it has optional *set up* and *tear down* phases, which are not measured but are used to prepare test data and verify results respectively. Between these optional phases is the actual *execution* phase, duration of which is measured. Test data are generated before each operation.

OPI – Create OPI represents a typical operation performed by a domain-specific business application. It creates an object graph centered around an entity

(an `OccurrenceReport` instance in this case) and then persists it. This simulates data being received, for instance, via REST services, connecting them to the logged-in user and persisting them into the repository. To make the results more representative, this transaction is repeated multiple times for separately created objects – a longer runtime reduces the influence of things like just-in-time compilation, garbage collection and possibly imprecise small-scale time measurement in the JVM. The operation specification is as follows:

1. *Set up* Persist `Person` instances to allow newly created reports to connect to them (simulating interconnection of existing and newly added data).
2. *Execution* Assign each generated report a random author and last editor and persist it in a separate transaction.




Fig. 3. UML class diagram of the benchmark model.

3. Tear down Verify the persisted data.

OP2 – Batch create As was already stated, the batch create operation represents, for example, data being imported or processed by an application in one transaction. Thus, the mode of operation is almost the same as *OP1*, only now all the entities are persisted in one large transaction.

1. *Set up* Persist Person instances to allow newly created reports to connect to them (simulating interconnection of existing and newly added data).
2. *Execution* Assign each generated report a random author and last editor and persist all of them in one large transaction.
3. *Tear down* Verify the persisted data.

OP3 – Retrieve *OP3* stands for the application requesting a specific entity together with its object graph (an instance and its property values in RDF terms). Again, to increase the clarity of the measurement, multiple objects are read one by one. This operation also verifies that all the required data were loaded by the persistence library.

1. *Set up* Persist test data using the same process described in *set up* and *execution* of *OP2*.
2. *Execution* Iterate through all existing reports, read each report using the existing report's iden-

tifier. Verify that the loaded report corresponds to the existing one.

OP4 – Retrieve all A “find all instances of a type” query is a typical operation for many applications. Its implementations can vary. Some libraries support mapping SPARQL query results to objects, so a SPARQL SELECT query is used, others contain such a method directly in their API (although a query may be more favorable for implementation of paging support). The goal is ultimately the same – retrieve a relatively large amount of entities together with their references at once.

1. *Set up* Persist test data using the same process described in *set up* of *OP3*.
2. *Execution* Read all reports and verify them. If the library API provides a dedicated method for this task, use it. Otherwise, use a SPARQL SELECT query.

OP5 – Update Update merges a modified entity into the repository. Several of its attributes and attributes of objects it references are updated, with both literal values and references to other objects being changed. Also, a reference to a new object that needs to be persisted is added. This operation is done in a single transaction. Several objects are updated in this way. Note that this scenario simulates an entity being updated externally and then being merged into the repository. For libraries which do not support detached objects, an instance is first loaded from the repository and the changes are made on this managed instance.

1. *Set up* Persist test data using the same process described in *set up* of *OP3*.
2. *Execution* For one half of the existing reports, perform the following updates (each object in a separate transaction):
 - set a different last editor (singular reference attribute) and last modified date (singular date-time attribute),
 - change occurrence name (singular string attribute),
 - change severity assessment (singular integer attribute),
 - increment revision number by one (singular integer attribute),
 - add a newly generated attachment (plural reference attribute),
 and merge the updated entity into the repository.
3. *Tear down* Verify the updates.

OP6 – Delete While modern triple stores do not particularly concentrate on removal operations since storage capacity is relatively cheap and many systems tend to create new versions instead of modifying existing data, business applications sometimes need to delete data, e.g., for privacy or security reasons. *OP6* represents deletion of an object and most of its references. Once again, multiple objects are deleted in separate transactions to improve result robustness.

1. *Set up* Persist test data using the same process described in *set up* of *OP3*.
2. *Execution* Delete one half of the existing reports, including their dependencies, i.e., everything from their object graph except the author and last editor.
3. *Tear down* Verify that all relevant data were removed.

A benchmark runner is used to execute the operations according to a configuration. The runner collects a number of statistics:

- The fastest and slowest round execution time,
- Mean round execution time,
- Q1, Q2 (median) and Q3 of round execution times,
- Total execution time,
- Standard deviation of round execution times.

In addition, it allows a file to be configured into which raw execution times for rounds are output so that they can be further analyzed.

The whole benchmark framework is available online¹⁵. Table 2 shows conformance of the comparison framework developed here to the Semantic Web knowledge base system benchmark requirements defined in [35] and discussed in Section 3.

5. Libraries selected for comparison

In this section, we introduce the libraries selected for the comparison. Several OTM libraries were selected based on their popularity, level of maturity, and development activity – libraries with latest changes older than five years were omitted. The feature comparison is platform agnostic, so a diverse set of libraries was chosen:

- ActiveRDF
- AliBaba
- AutoRDF
- Empire
- JAOb
- JOPA
- KOMMA
- RDFBeans
- RDFReactor
- The Semantic Framework
- Spira
- SuRF

However, since the benchmark part of the comparison framework is based on Java, only Java OTM libraries were evaluated so that platform specifics (e.g., code being compiled to native code like in C++, to bytecode like in Java/C# or directly interpreted like in Ruby) do not influence the results. The selection is further narrowed by requiring the libraries to support triple store access – using plain RDF files for real-world applications is not suitable in terms of performance and data management. Finally, to improve the comparability, we decided to concentrate only on libraries which support access to a specific triple store. This eliminates the influence of the storage technology and its implementation on the performance benchmark where, for example, a Jena relational database-based storage could have different performance characteristics than an RDF4J storage. We picked RDF4J¹⁶ [9] because the amount of mature tools for it is the greatest, including industry-grade storage implementations like GraphDB [17]. As a result, the following libraries were evaluated in the benchmark:

- AliBaba
- Empire
- JOPA
- KOMMA
- RDFBeans

RDFReactor was a candidate for the performance benchmark as well, but its RDF4J/Sesame storage connector works only with an embedded in-memory store and cannot connect to a remote repository. Thus, it was excluded from the benchmark selection.

¹⁵<https://github.com/kbss-cvut/otm-benchmark>, accessed 2018-11-17.

¹⁶Since RDF4J is an evolution of the Sesame API, we will use the terms interchangeably throughout the paper. We will stress the difference when necessary.

Table 2

Evaluation of benchmark criteria defined in [35] on our benchmark. SUT stands for system under test

Req#	Description	Evaluation
G1	Scalability	The data generator allows generation of datasets of arbitrary size.
G2	Broad architectural scope	The comparison framework is platform and storage independent. The benchmark requires Java.
G3	Easy to add new SUT	The benchmark API requires implementation of approximately ten short classes to add a new SUT.
G4	Meaningful metrics	The benchmark collects commonly used performance metrics and statistics (mean, standard deviation, etc.).
G5	Controlled measurements	Data generation is repeatable, warm-up period is included and configurable.
S1	Real-world workload	The benchmark measures performance of typical domain-specific object-oriented application operations.
S2	Metrics supporting various (possibly conflicting) requirements	The benchmark measures execution time and memory consumption. Plus it provides feature comparison criteria.
S3	Metrics capture result conformance to expectations	Not fully applicable (original intention was for inference completeness/soundness). The benchmark verifies correct results of the operations against predefined data.

The mapping in the selected Java libraries is realized via annotations which specify the mapping of Java entity classes to RDFS classes and attributes to RDF properties. Listing 3 shows an example entity class declaration in Empire.

Listing 3: Example of an Empire Java entity class with annotations specifying mapping to the RDF(S) data model. *Getters* and *setters* are methods used to get and set the values of attributes respectively.

```

@Entity
@Namespaces({"as",
"http://onto.fel.cvut.cz/ontologies/a-s/",
"ufo", "http://onto.fel.cvut.cz/ontologies/ufo/",
"rt", "http://onto.fel.cvut.cz/ontologies/reporting-tool/"})
@RdfsClass("ufo:event")
public class Event implements SupportsRdfId {

    @RdfProperty("as:has_start_time")
    private Date startTime;

    @RdfProperty("as:has_end_time")
    private Date endTime;

    @OneToMany(fetch = FetchType.EAGER, cascade = CascadeType.ALL)
    @RdfProperty("ufo:has_part")
    private Set<Event> subEvents;

    @RdfProperty("rt:has_event_type")
    private URI eventType;

    // Getters and setters omitted for brevity
}

```

The following paragraphs introduce all the evaluated OTM solutions.

5.1. ActiveRDF

ActiveRDF [61] is an object-oriented API for Semantic Web data written in Ruby. It is built around the premise that dynamically-typed languages like Ruby are more suitable for semantic data representation because, like RDF, they do not require objects to conform to a specific class prescription. ActiveRDF represents resources by object proxies, attaching to them methods representing properties. Invocations of these methods translate to read/write queries to the storage. In addition, it supports an object-oriented query syntax based on PathLog [62] and is able to actually generate convenience filter methods automatically.

In fact, ActiveRDF is not a typical OTM library like AliBaba or Empire where RDFS classes are mapped to entity classes. Instead, it treats RDF resources as objects to which it dynamically attaches methods corresponding to their properties. One caveat of the highly dynamic nature of ActiveRDF is pointed out in [63] – it does not offer type correctness and typographical error verification present in libraries based on statically-typed languages such as Java.

5.2. AliBaba

AliBaba [22] is a Java library for developing complex RDF-based applications. It is an evolution of the Elmo [64] library. AliBaba is an extension to the original Sesame API, so it supports only storages accessible via a Sesame/RDF4J connector. Its API is cen-

tered around the `ObjectConnection` class, which allows to persist and retrieve objects mapped by AliBaba's OTM. It uses dynamically generated proxy objects to track updates. AliBaba also allows to use SPARQL queries to support more complex strategies of mapping values to Java attributes. On the other hand, the entities managed by AliBaba are not able to store IRIs of the resources they represent, which is a severe limitation requiring the developer to maintain a map of objects and identifiers externally.

In addition to OTM, AliBaba contains an HTTP server which can automatically make resources accessible as REST services, providing querying, inserting and deleting capabilities.

5.3. *AutoRDF*

AutoRDF [65] is a library facilitating handling of RDF data written in C++. The fact that it is written in C++ allows it to run, for instance, on embedded devices – actually, one of its intended usages is in ID document verification. It is built on top of the Redland RDF library¹⁷. Similar to ActiveRDF, AutoRDF does not really store any data in the runtime objects. Instead, they act as proxies and all operations are directed to the underlying RDF. Besides RDFS, AutoRDF supports several OWL features like `owl:oneOf` (mapped to enumerations), cardinality restrictions and classes being restrictions of other classes. It also supports anonymous resources (identified by RDF blank nodes), and automatic documentation generation based on selected RDFS vocabulary terms (e.g., `rdfs:label`, `rdfs:comment` or `rdfs:isDefinedBy`).

5.4. *Empire*

Empire [23] is an RDF-based implementation of the JPA standard [27]. Therefore, its API should be familiar to developers used to working with relational databases in Java. However, since parts of JPA deal with the physical model of the relational database under consideration (e.g., the `@Table` and `@Column` annotations), Empire is not fully JPA-compliant. On the other hand, it does support the basic `EntityManager` API, query API, and entity lifecycle. Unfortunately, Empire's documentation is limited and it is often unclear which parts of JPA are implemented.

It adds a set of RDF-specific annotations which are used to express the mapping of Java classes and attributes to RDFS classes and properties (see Listing 3). It does not support anonymous resources, so all instances it works with must have an IRI.

5.5. *JAOB*

The Java Architecture for OWL Binding (JAOB) [66], as its name suggests, is primarily intended for OWL ontologies. However, with minor adjustments, RDF(S) ontologies can be manipulated by it. First, the annotation used for declaring class mapping is called `OWLClass`, but an RDFS class IRI can be used as its value. Then, in contrast to RDF(S), OWL defines three kinds of properties - object properties, datatype properties and annotation properties [12]. These properties have equivalent annotations in JAOB. To use them with RDF(S), it is necessary to use `OWLObjectProperty` for attributes whose value is expected to be a reference to another object (RDF resource), `OWLDataProperty` can be used for other attributes.

JAOB is built on top of OWL API, so it supports only ontologies stored in files. It basically works as *(un)marshaller*, so it is able to load or save collections of mapped objects. However, there is, for example, no direct way to update an entity in JAOB.

JAOB also contains a generator which allows generating Java classes based on an ontology schema and vice versa.

5.6. *JOPA*

The Java OWL Persistence API (JOPA) [20, 24] is a persistence library primarily designed for accessing OWL ontologies. However, it can be used to work with RDF data as well. Its API is designed so that it resembles the JPA as much as possible, but it does take semantic data specifics like named graphs or inference into account. JOPA tries to bridge the gap between the object-oriented application world, which uses the *closed-world assumption* (missing knowledge is assumed false), and the ontological world, which is based on the *open-world assumption* (missing knowledge can be true in some worlds and false in others), by splitting the model into a static and dynamic part. The static part is mapped to the object model and its validity is guarded by *integrity constraints* with closed-world semantics [67]. The dynamic part is not directly mapped (although it is accessible to a limited

¹⁷<http://librdf.org/>, accessed 2018-11-26.

extent) and may evolve freely without affecting the object model.

Using JOPA with RDF(S) ontologies requires minor adjustments because JOPA is OWL-based. These adjustments are similar to those described for JAOb.

5.7. KOMMA

The Knowledge Management and Modeling Architecture (KOMMA) [68] is a Java library for building applications based on semantic data. A part of this system is an object-triple mapping module, but the library itself has much richer functionality, including support for building graphical ontology editors based on the Eclipse Modeling Framework¹⁸. OTM in KOMMA is based on Java interfaces representing RDFS classes, for which it generates dynamic implementations at runtime. This allows, for instance, to accommodate support for multiple inheritance into the model. In addition, KOMMA also supports transaction management, caching and RDF named graphs. Parts of the KOMMA code base reuse code from AliBaba since both frameworks are based on similar principles.

5.8. RDFBeans

RDFBeans [69] is another OTM library built on top of RDF4J. It allows two forms of an object model: 1) the object model may consist of Java classes representing the RDFS classes; 2) the object model may consist of interface declarations forming the mapping, RDFBeans will generate appropriate implementations at runtime using the *dynamic proxy* mechanism. RDFBeans supports, besides the basic features like transactions and inheritance, also mapping of Java collections to RDF containers, or operation cascading, e.g., when an updated entity is merged into the storage, objects referenced by this entity are merged as well. Instances of entity classes without an explicitly declared identifier attribute are mapped to RDF blank nodes.

5.9. RDFReactor

RDFReactor [70] is yet another Java OTM library. It supports code generation and uses proxy objects to access RDF data directly, instead of storing attribute values in entities. In addition to the mapped properties, RDFReactor entity classes also contain an API for type information manipulation – methods `getAll`, `add`

and `remove` allow to retrieve and update types of an instance at runtime. Besides RDFS, RDFReactor also supports selected OWL features like cardinality constraints, inverse properties.

All entity classes mapped by RDFReactor extend the `ReactorBase` class which acts as a generic representation of an RDF resource, allowing untyped access to all its properties. Subclasses can then explicitly map these properties to provide type-safe access.

5.10. The Semantic Framework

The Semantic Framework (SF) [71] is a Java library for object-oriented access to ontologies. Similar to JAOb or JOPA, the SF is primarily built for OWL, but it can, again, be used with RDFS as well. It internally uses OWL API to process RDF files and its core is an extension of the Jenabean library [72]. The combination of Jena (internally used by Jenabean) and OWL API is, in our opinion, rather odd given the fact that Jena is also able to process RDF files. The mapping is realized using annotations in entity classes.

5.11. Spira

Spira [73] is a library for viewing RDF data as model objects written in Ruby. It allows to access RDF data both as domain objects (entities) and as RDF statements. In particular, one may view and edit statements representing an entity directly from its object representation. On the other hand, it is also possible to create RDF statements directly, without any object representation of the resource. Spira allows to ‘view’ resources as instances of various classes without requiring them to explicitly have the corresponding RDF type. It also allows to map one Ruby class to multiple ontological types, to work both with named and anonymous resources, or to use localized property values, i.e., a single-valued property can have a value in multiple languages.

Spira uses `RDF.rb`¹⁹ to access the underlying RDF data. It can work with multiple repositories at once, all of them accessed in a thread-safe manner.

5.12. SuRF

SuRF [74] is a Python object RDF mapper. SuRF is built on top of the `RDFLib`²⁰ library and allows ac-

¹⁸<http://www.eclipse.org/modeling/emf/>, accessed 2018-01-02.

¹⁹<https://github.com/ruby-rdf/rdf>, accessed 2018-11-21.

²⁰<https://github.com/RDFLib/rdfib>, accessed 2018-11-21.

cess to various triple stores, including Sesame [9], Virtuoso [18], or a generic SPARQL endpoint. Attributes of SuRF entities are by default loaded lazily and references to other entities are automatically accompanied by convenience inverse versions, i.e., a `knows` attribute of an entity is complemented by an `is_knows_of` attribute on the target instances (yes, the wording is often awkward).

SuRF contains a search API which can be used to retrieve instances of the mapped classes and filter them by their attribute values. One subtle feature of SuRF is that it allows to retrieve resources regardless of whether they are actually present in the triple store. To check for their existence, a dedicated `is_present` method has to be called.

5.13. Libraries omitted from selection

There are several OTM libraries which were omitted from the selection. Elmo [64], Jastor [75], Jenabean [72], Owl2Java [47] or Sommer [76] are obsolete, with their last versions published more than five years ago.

We also wanted to include a purely SPARQL endpoint-based solution which would allow us to compare the performance of a storage agnostic library to libraries exploiting vendor-specific optimizations. However, RAN-Map [77] is not published online (neither sources nor binaries) and TRIOO [78] is buggy and we were not able to adapt it to the benchmark application without significant changes to its code base (for example, it does not insert resource type statements on persist).

Libraries like JOINT-DE [32] (and the original JOINT), Sapphire [63] and the Semantic Object Framework [79] were excluded because their source code is not publicly available and the articles describing them do not provide enough details to evaluate the comparison criteria.

6. Feature comparison

Comparison of features of the selected libraries according to the criteria defined in Section 4 is summarized in Table 3. In the table, ✓ signifies that the criterion is fully satisfied according to the condition defined in Section 4. ◦ signifies partial satisfaction and the text should be consulted for further details, and × means the criterion is not satisfied. We now expand on the results.

6.1. GC1 – Transactions

AutoRDF, JAOb and the Semantic Framework do not support transactions likely due to their lack of support for triple stores in general (more details in Section 6.2). While ActiveRDF, RDFReactor and Spira do support access to regular triple stores and, for example, RDFReactor internally makes use of the store’s transactional mechanism, they do not allow the programmer to control the transactions externally. Empire’s API hints towards support for transactional processing but its implementation is rather strange. As will be discussed in Section 7, operations which insert/update/remove data are automatically committed without any way of preventing this behavior. So, for example, it is not possible to bind updates to multiple entities into one transaction.

AliBaba, KOMMA, and RDFBeans support transactions by relying on the underlying storage’s transaction management, i.e., starting an OTM-level transaction begins a database-level transaction and its end causes the database-level transaction to end as well. SuRF, on the other hand, tracks the changes to objects by marking their updated attributes dirty and allows the programmer to commit the changes manually. JOPA handles transactions by creating clones of the manipulated objects and tracking their changes locally, pushing them into the storage on commit. In addition, the changes are stored in transaction-local models and are used to enhance results of subsequent operations in the same transaction, e.g., when an entity is removed, it will not be in the results of a find all query executed later in the same transaction.

6.2. GC2 – Storage access variability

AutoRDF, JAOb and the Semantic Framework support only access to RDF data stored in text files, albeit with various serializations. AliBaba and RDFBeans are tightly bound to the Sesame API, so they can access exclusively Sesame/RDF4J storages.

While KOMMA does not have any other implementation than RDF4J storage access, its internal APIs are designed so that new storage connector implementations can be added.

ActiveRDF supports access to various storages, including Jena, Sesame or a generic SPARQL endpoint. Empire contains storage access modules which can implement the required access API. For instance,

Table 3

Selected OTM libraries compared using criteria defined in Section 4. × means no support, ○ represents partial support (consult the main text for further details), ✓ is full support of the feature and N/A signifies that the feature cannot be evaluated in the particular case

Criterion	ActiveRDF	AliBaba	AutoRDF	Empire	JAOB	JOPA
GC1 (Transactions)	×	✓	×	○	×	✓
GC2 (Storage-agnostic)	✓	×	×	✓	×	✓
GC3 (Query result mapping)	○	✓	×	✓	×	✓
GC4 (Object query language)	✓	×	×	×	×	×
GC5 (Detached objects)	×	×	×	✓	✓	✓
GC6 (Model Generator)	N/A	✓	✓	×	○	○
OC1 (Explicit inference)	×	✓	×	×	×	✓
OC2 (Named graphs)	✓	×	×	✓	×	✓
OC3 (Provenance management)	×	○	×	×	×	×
MC1 (Inheritance)	N/A	✓	✓	○	×	○
MC2 (Unmapped data access)	✓	✓	×	×	×	✓
MC3 (RDF containers/collections)	×	✓	×	×	×	○

Criterion	KOMMA	RDFBeans	RDFReactor	SF	Spira	SuRF
GC1 (Transactions)	✓	✓	×	×	×	✓
GC2 (Storage-agnostic)	○	×	✓	×	✓	✓
GC3 (Query result mapping)	✓	×	×	×	×	×
GC4 (Object query language)	×	×	×	×	×	✓
GC5 (Detached objects)	×	✓	×	✓	×	×
GC6 (Model Generator)	✓	×	✓	×	×	N/A
OC1 (Explicit inference)	✓	×	✓	×	×	×
OC2 (Named graphs)	✓	×	✓	×	×	×
OC3 (Provenance management)	○	×	×	×	○	×
MC1 (Inheritance)	○	○	○	○	○	N/A
MC2 (Unmapped data access)	✓	×	✓	×	✓	✓
MC3 (RDF containers/collections)	✓	✓	×	○	×	×

it provides a connector to the Stardog database²¹. JOPA defines a storage access layer called the *OntoDriver*. This allows switching the underlying storage easily. Currently, JOPA supports Jena, OWL API and Sesame/RDF4J storages. RDFReactor uses a similar mechanism – an underlying storage access layer called *RDF2Go* which, at the time of writing, supports Jena and RDF4J access. Spira, thanks to its reliance on RDF.rb, can access various storages, including Sesame, AllegroGraph²² or MongoDB²³. Simi-

larly, SuRF can use the RDFLib to connect to AllegroGraph, Sesame or a SPARQL endpoint.

6.3. GC3 – Query result mapping

In contrast to JAOB, which does not provide any query API at all, AutoRDF, RDFBeans, the Semantic Framework and Spira do not provide a query API either but contain at least a find all method, which allows to retrieve all instances of the specified type. In case of RDFBeans the situation is peculiar because its RDFBeanManager wraps an instance of RDF4J RepositoryConnection which supports SPARQL query execution, yet, RDFBeanManager does not expose this func-

²¹<https://www.stardog.com/>, accessed 2018-11-21.

²²<https://franz.com/agraph/allegrograph/>, accessed 2018-11-21.

²³<https://www.mongodb.com/>, accessed 2018-11-21.

tionality. SuRF allows to execute arbitrary SPARQL queries, but cannot map their results to entities.

While ActiveRDF does not support mapping SPARQL query results to entities, its advanced query API obviates this issue by supporting almost complete SPARQL grammar [61].

Finally, AliBaba, Empire, JOPA and KOMMA support mapping SPARQL query results to entities by allowing to specify target entity class when retrieving query results.

6.4. GC4 – Object-level query language

ActiveRDF contains a variation of a criteria query API containing methods representing query operations like selection, filtering and projection. Additionally, it automatically generates methods for filtering objects by their properties, e.g., it will provide a method `Person.find_by_name` for a class `Person` with an attribute name. A similar API exists in SuRF.

None of the other libraries support any object-level query language, so their users have to resort to SPARQL queries, if available.

6.5. GC5 – Detached objects

ActiveRDF, AliBaba, AutoRDF, RDFReactor, Spira, and SuRF do not support detached objects because their entities act as proxies which load attribute values from the repository when they are accessed for the first time, or on each access, depending on the internal implementation. Conversely, setting attribute values causes the changes to be written into the storage immediately (except for SuRF which tracks changes locally). Therefore, they have to hold onto a connection to the storage in order to provide basic data access functions. Similarly, although it would appear that KOMMA does support detached objects because it is possible to close a `KOMMA IEntityManager` and still access attributes of an object read by the closed manager, the contrary is true. As pointed out in [32], the generated proxy objects retain a connection to the underlying storage. Closing an `IEntityManager` only closes a delegate object.

Empire, JAOb, JOPA, RDFBeans and the Semantic Framework, on the other hand, allow working with the entities completely independently of the persistence context from which they were retrieved because they store the data in the actual objects.

6.6. GC6 – Code/ontology generator

Empire, RDFBeans, the Semantic Framework and Spira do not contain any generator capable of creating entity classes from ontology schema or vice versa. ActiveRDF and SuRF do not contain code generators either, but they do not actually use any static object model and all entity classes are generated on demand at runtime.

AliBaba, AutoRDF, JAOb, JOPA and KOMMA contain generators able to build an object model from the ontology. The AliBaba generator supports both RDFS and OWL, but it is less configurable. AutoRDF supports several OWL constructs, but its core is RDFS-based. On the other hand, the generators in JAOb and JOPA are OWL-based and expect the ontology to contain declarations of OWL classes and object, data and annotation properties. KOMMA reuses the AliBaba code generator but the feature is undocumented.

JAOb is the only library with a generator capable of creating an ontology schema based on an object model. This generator is able to create classes and their hierarchies, plus OWL data and object property declarations with domains and ranges and information about whether the property is functional or not [12].

6.7. OC1 – Explicit inference treatment

JOPA allows marking entity attributes as possibly containing inferred data. Currently, this makes them effectively read-only. However, in the future, the implementation may permit additive changes. AliBaba, KOMMA and RDFReactor allow marking attributes as read-only, a mechanism which could be used to prevent attempts at modification of inferred statements.

6.8. OC2 – Named graphs

AliBaba, AutoRDF, JAOb, RDFBeans, the Semantic Framework, Spira, and SuRF are not aware of RDF named graphs and work only with the default graph.

Empire supports named graphs to a limited extent. Named graph access is specified via the `@NamedGraph` annotation, which has two modes of operation. The first mode leads to each instance of the annotated class being stored in its own named graph. The second mode then allows to specify an IRI of a named graph into which all instances of the annotated class are saved. JOPA and ActiveRDF, on the other hand, allow one to specify a named graph per instance and for individual attributes. RDFReactor uses model in-

stances which represent individual named graphs of the repository. Each entity is then associated with a single model.

6.9. OC3 – Automatic provenance management

None of the evaluated libraries fully supports automatic provenance management. However, AliBaba and KOMMA support the notion of *behaviors* which could be used to generate provenance data. A behavior adds custom functionality to entities and may be shared by multiple entity classes. One such implementation could be invoked by entity setters, extract user information and insert corresponding provenance statements into the repository. In Spira, Active Model²⁴ callbacks may be used to achieve the same goal – the callbacks can be attached to arbitrary methods, so, for example, a callback invoked after the `save` method could create provenance data. JOPA supports the notion of entity lifecycle listeners introduced in JPA [27], so, for instance, a `prePersist/preUpdate` listener may set provenance data for an entity. Nonetheless, none of the libraries has such a provenance generator.

6.10. MC1 – Inheritance mapping

Inheritance mapping is probably the most complex feature in the comparison framework, with many subtle issues. Despite this, several evaluated libraries take a straightforward approach which can often lead to unexpected results.

Consider Empire which does support multiple inheritance in that it is able to work with classes which inherit mapped getters/setters from multiple interfaces. However, it is unable to persist more than one type for an entity. So if each interface declares a mapped type, only one of them gets persisted. KOMMA and RDFBeans suffer from the same issue, i.e., they correctly interpret inherited attributes, but always persist only a single type. To illustrate the issue, let us have interfaces `Student` and `Employee`, which are mapped to RDFS classes `ex:Student` and `ex:Employee` respectively. Then, we declare an entity class `WorkingStudent`, which implements both `Student` and `Employee`. Persisting an object `c` of type `WorkingStudent` into the repository would result in either `ex:c rdf:type ex:Student` or `ex:c rdf:type ex:Employee` being inserted, de-

pending on the order of declarations in the `implements` clause of `WorkingStudent`, where `c` is mapped to `ex:c`. Conversely, a straightforward implementation of loading an instance of `WorkingStudent` would require the corresponding resource to have types `ex:Student` and `ex:Employee`, mapping `WorkingStudent` to an unnamed intersection of the two. However, this is not formalized in any way. A reasonable solution may be to require mapping of concrete entity classes to a single RDFS class only.

The Semantic Framework appears to also support multiple inheritance thanks to its use of Java classes and interfaces. However, the implementation extracts only attributes declared in a particular class, without regard for any fields inherited from its superclass. This in effect means that any values of fields declared in a superclass are ignored when persisting an instance.

JAOB does not support inheritance mapping. SuRF allows to specify multiple superclasses when loading a resource but they do not represent mapped ontological classes, they are regular Python classes adding custom behavior to the instance. On the other hand, similarly to ActiveRDF, since loaded instances contain attributes corresponding to all properties found on the resource, the concept of inheritance mapping does not really apply in this case. Appropriate type assertion presence on a resource can be of importance w.r.t. the type hierarchy, but such an issue is more closely related to the way libraries handle explicit and inferred statements.

JOPA currently supports only class-based inheritance, so mapping classes with multiple supertypes is not possible. However, the semantics of its treatment of inheritance is precisely defined – it requires an explicit type assertion and supports polymorphic instance access, i.e., an inferred supertype may be specified for loading as long as an unambiguous asserted subtype can be found for a resource. Similarly, RDFReactor and Spira support only single type inheritance. On the other hand, Spira allows an entity to specify multiple types, so multiple ontological classes can be combined into a single mapped Ruby class, which can be used as another class's parent.

AliBaba and AutoRDF are thus the only libraries fully supporting class hierarchy mapping. Since AliBaba is a Java library, it relies on interfaces to support multiple inheritance. AutoRDF is able to exploit the built-in support for multiple inheritance in C++.

²⁴See https://guides.rubyonrails.org/active_model_basics.html, accessed 2018-11-17.

6.11. MC2 – Unmapped data access

Providing access to unmapped properties is more challenging in statically typed languages like Java or C++. AutoRDF, Empire, JAOb, RDFBeans and the Semantic Framework do not support such a feature.

AliBaba, JOPA, KOMMA and RDFReactor, although being Java libraries, do attempt to supply this access. In JOPA, there are two relevant attribute annotations – a `@Types` attribute allows to read and modify an instance’s ontological types and `@Properties` denotes a map of property-set of values pairs, representing the properties not mapped by the rest of the entity attributes. AliBaba allows to retrieve the RDFS resource represented by an entity and, since it is an extension of the Sesame API, to access any data related to the resource via regular Sesame API calls. KOMMA contains the `IResource` interface which represents RDFS resources and has utility methods for manipulating statements related to particular resources. RDFReactor solves the issue by making all entity classes implement the `ReactorBase` interface (actually, the classes usually extend the `ReactorBaseImpl` abstract class) whose API contains generic methods for setting and retrieving arbitrary properties attached to the instance. In fact, the mapped attributes make use of this API internally.

ActiveRDF and SuRF, due to their lack of a static model, naturally allow to access all the properties of a resource. In Spira, all model objects can be manipulated also at the RDF statement level thanks to the integration with RDF.rb.

6.12. MC3 – RDF collections and containers

Support for RDF collections and containers is scarce among the evaluated libraries. The only libraries which fully implement this feature are AliBaba, KOMMA and RDFBeans. AliBaba supports RDF lists and RDF containers. Both are mapped to instances of `java.util.List`. If the list root has the type `rdf:List`, it is treated correspondingly, otherwise, it is assumed to be a RDF container. This root type management, however, has to be done on RDF statement level, no dedicated API is present for it in AliBaba. KOMMA’s support for RDF lists and containers is virtually the same as in AliBaba. RDFBeans also allows to work with both RDF containers and lists. The type of the target RDF construct is specified via Java annotations.

The approaches of the Semantic Framework and JOPA may be regarded as partial support. The Se-

semantic Framework automatically stores array and `java.util.List` attributes as RDF sequences, without any way to configure this behavior. JOPA supports the OWL design pattern of linked lists and linked lists with content nodes [80], but it does not support any of the RDF containers. The type of the list is specified via an annotation.

7. Performance comparison

All the benchmark operations revolve around an object graph whose template is shown in Figure 4. As can be seen, there is one central `OccurrenceReport` object connected to an `Occurrence`, which is the root of a three level-deep balanced binary tree of `Events`. Each report also has a set of three `Resources` attached to it and is connected to a randomly selected author and last editor (instances of `Person`).

7.1. Experiment setup

The experiments were conducted on a machine with the following setup:

- OS: Linux Mint 17.3 64-bit
- CPU: Intel Core i5-750 2.67 GHz (4 cores)
- RAM: 16 GB
- Disk: Samsung SSD 850 EVO 250 GB
- Java: Java HotSpot JDK²⁵ 8u161, 64-bit

We decided to use *GraphDB 8.4.1 Free* as the storage. It was running on the same machine as the benchmark to eliminate possible network lag. GraphDB [17] is an industry-grade semantic graph database supporting the RDF4J API. The free version is limited to two queries in parallel, but this is of no concern, since the benchmark is not multithreaded. To ensure that none of the evaluated libraries had an unfair advantage by being specifically optimized for GraphDB, an additional experiment on a native RDF4J storage was conducted. Its results showed similar relative differences between the OTM libraries as the measurements on GraphDB. The latest available versions of the libraries were used. Since some of them were not present in the Maven Central Repository²⁶ (an online repository of Java libraries), they had to be built from source code locally. The versions were (as of 16th November 2018) as follows:

²⁵Java Development Kit

²⁶<http://search.maven.org/>, accessed 2018-11-16.

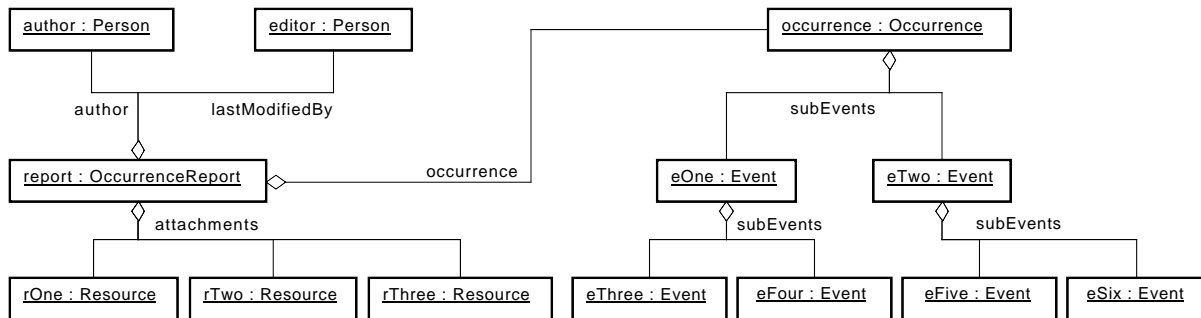


Fig. 4. UML object diagram of the object model used by the benchmark.

- *AliBaba 2.1*, built locally, sources downloaded from <https://bitbucket.org/openrdf/alibaba/>,
- *Empire 1.0*, built locally, sources downloaded from <https://github.com/mhgrove/empire>,
- *JOPA 0.10.7*, retrieved from Maven Central,
- *KOMMA 1.3.3-SNAPSHOT*, built locally, sources retrieved from <https://github.com/komma/komma>,
- *RDFBeans 2.2*, retrieved from Maven Central.

After publishing the initial performance results, we received optimization suggestions from the authors of KOMMA. They provided a significant performance boost for KOMMA and were incorporated into the benchmark setup. However, these optimizations are undocumented.

7.1.1. Performance

The scalability of the performance benchmark allows custom specification of the size of the data. For our experiments, we used the following values:

- OP1 – Create** 300 Persons to choose the authors and last editors from, 300 OccurrenceReport instances together with their associated object graphs.
- OP2 – Batch Create** The same setup as *OP1*.
- OP3 – Retrieve** 300 OccurrenceReports together with their object graphs and 300 Persons pre-persisted (each report is assigned a random author and last editor from the Person instances).
- OP4 – Retrieve All** The same setup as *OP3*.
- OP5 – Update** The same setup as *OP3*. One half of the reports (every odd one, to be precise) are updated.
- OP6 – Delete** The same setup as *OP3*. One half of the reports (again, every odd one) are deleted, together with their dependencies.

The actual size of the dataset (e.g., 300 Persons above) was chosen arbitrarily, just to provide a reasonable ex-

ecution time on the machine running the evaluation. Each operation was executed in 120 rounds, where the first twenty rounds represented a warm-up period which was not measured and was intended for the Java virtual machine to stabilize and perform optimizations. The following 100 rounds were measured. This was executed in three separate runs, i.e., three executions of the JVM, to account for different ways the code could be optimized, as discussed in [31]. The whole benchmark, i.e., all runs of all operations for all libraries, was executed automatically by a shell script. The repository was cleared after each round and GraphDB was restarted between the operations.

The benchmark was run for several values of heap size. The results should then reveal how sensitive the libraries are to restricted heap size and how they scale as the heap size changes. The heap size was fixed (using the `-Xms` and `-Xmx` JVM configuration parameters) to:

- 32 MB
- 64 MB
- 128 MB
- 256 MB
- 512 MB
- 1 GB

The maximum number of statements generated by the benchmark for *OP3* (the other operations are either the same or very similar) was 23 700 for Empire, JOPA, KOMMA and RDFBeans and 23 700 or 41 026 for AliBaba. The pair of numbers for AliBaba is due to the rather strange cascading strategy employed by AliBaba. When persisting an object, AliBaba automatically cascades the persist recursively to all references of the persisted entity. Such a strategy generates *blank nodes* for the referenced instances, leading to the 23 700 statements in total. If, on the other hand, one wanted to be able to denote the referenced instances

(e.g., the occurrence or person) using a regular IRI, they would have to be persisted separately. This, however, leads to a large duplication of data, because Alibaba will perform the cascading anyway, resulting in the 41 026 statements.

7.1.2. Memory

Measuring memory utilization turned out to be a bit tricky. Most of the literature concentrates on obtaining and analyzing heap dumps in the search for memory leaks (e.g., [58]) which was not the goal of this comparison.

Sampling heap occupation or object allocation rate (as DaCapo [29] does) is difficult to set up and its information value is arguable. Peak or average heap size or occupation also provide unclear benefits. In the end, a different strategy was used – a separate runner was developed which executed a series of CRUD operations similar to the ones used in the performance benchmark for a specific period of time. More precisely, the runner executed *OP1*, *OP4*, *OP5* and *OP6* in a sequence over and over until a specified timeout. This way, all the libraries were used for the same amount of time. To collect the results, the Java virtual machine was configured to output detailed information about garbage collection events (using the `-XX:+PrintGCDetails` flag) and a fixed heap size was set, so that the application had a limited amount of memory available.

The particular setup for the memory benchmark was as follows:

- Runtime: 4 h,
- Heap size: 40 MB,
- Default garbage collector settings for the JVM used in the benchmark.

After the execution, *GCViewer*²⁷ was used to analyze the results and collect the following values:

- The number of young generation garbage collection events (YGC),
- The number of full garbage collection events (FGC),
- Accumulated garbage collection pause time (GCT),
- Application throughput.

To put these terms in context: since most objects in applications are short-lived, YGC is able to dispose of them quickly [58]. FGC events occur when the Java

virtual machine needs to reclaim memory occupied by older objects and it is more demanding in terms of CPU time. Generally, an application benefits from as few FGC's as possible. Parts of both young generation and full garbage collection can happen in parallel to the application execution (if the CPU allows it), but sometimes the application has to be stopped so that the *garbage collector* can mark objects for disposal and remove them. It is these so-called *stop-the-world* events which are represented by GCT and which influence the application throughput. In particular, the throughput indicates how much of the total runtime the application spent actually performing its tasks, the rest being spent in garbage collection. More detailed explanation of garbage collection principles and strategies employed by the JVM can be found, for example, in [58, 81]. All in all, these values, together with the limited amount of available memory and the fixed execution time, offer a reasonable picture of how an application using these libraries is utilizing the available memory. A small heap size was used in order to put the garbage collector under more stress so that its overhead would be more noticeable. However, a 32 MB heap was not used since most evaluated libraries had issues using it (see Section 7.2.2).

7.2. Results

Let us now delve into the results of the experiments and attempt to interpret them. First, performance results are detailed, followed by a discussion of scalability of the evaluated libraries. Finally, memory utilization experiments are examined. Complete results can be found in Appendix A and also – together with the underlying data – online²⁸.

7.2.1. Performance

In this part, relative performance differences between the evaluated libraries are discussed. As a representative example, results for benchmark with heap size set to 1 GB are visualized in a plot in Figure 5. Table 4 complements the plot visualization with information on execution time mean, standard deviation, the 95 % confidence interval and the application average throughput. Throughput illustrates how many *reports* was the OTM library able to process per second. A report in this context denotes the object graph shown in Figure 4. More specifically, this means:

²⁷<https://github.com/chewiebug/GCViewer>, accessed 2018-11-21.

²⁸At <https://kbss.felk.cvut.cz/web/kbss/otm-benchmark>, accessed 2018-11-22.

- OP1, OP2** Each persisted report consists of 11 entities mapped to 68 triples (Person instances are pre-persisted in the set up phase and thus not counted).
- OP3, OP4** Each loaded report consists of 13 entities mapped to 90 triples (Person instances are loaded together with the report).
- OP5** Each updated report consists of 90 triples loaded for the merge and 9 triples updated.
- OP6** Each removed report represents 90 triples loaded and 68 of them removed (Person instances were not removed).

Performance on other heap sizes is discussed in Section 7.2.2, Appendix A.1 contains complete results.

All the libraries performed relatively consistently, with low deviation and the confidence interval size within 2% of the execution time mean, as suggested by [31]. Only Empire exhibited larger standard deviation for the retrieval operations (*OP3* and *OP4*). KOMMA displayed large standard deviation and many outliers when running with 32 MB heap but this is most likely due to the high stress put on the garbage collector which has to pause the application relatively often. The following paragraphs elaborate on individual operation results.

Create (OP1) and batch create (OP2) The character of *OP1* and *OP2* suggests that the libraries should perform better in *OP2* because they have to execute only a single transaction compared to a linear number of transactions in *OP1*. The only library that defied this expectation and performed virtually the same in both *OP1* and *OP2* was Empire. This is due to the fact that Empire actually commits the transaction after each data modifying operation of its *EntityManager*, more precisely, after *merge*, *persist* and *remove*. Not only is this highly inefficient in the batch scenario, but it can also lead to incorrect behavior because it effectively prohibits a rollback of transactional changes or grouping multiple operations into one transaction. The rather poor performance of Empire is partially due to the fact that it iterates over all attributes of the object being persisted using Java *reflection* instead of relying on a *metamodel* (model of entity classes and their attributes) [27] built during application persistence initialization. Reflection is a powerful mechanism which allows a program to examine and modify itself and the environment in which it operates [82] but it represents a performance overhead. While middleware libraries in Java typically use reflection, it is important to be aware of its performance impacts and to avoid

them when possible. The other factor in Empire's create performance is its strategy for checking that an instance does not already exist in the repository. It uses a SPARQL SELECT query with unbound subject, predicate and object, filtering results by the subject being equal to the persisted individual's identifier. First, our experience with the RDF4J API indicates that executing SPARQL queries is in most cases slower than using the statement filtering API. Second, a simpler ASK query could have been used.

The performance of KOMMA depends on whether a data change tracker is enabled. If it is enabled, it attempts to remove old attribute values on change. This leads to significantly worse performance. When it is disabled, KOMMA's performance is comparable to that of Empire and RDFBeans. The data change tracker relevance for *persist* is given by the fact that KOMMA does not implement a dedicated *persist* operation, instead, it allows to register an identifier to a particular type. Persisting attribute values then requires calling setters on the managed instance returned by the registration method. Whenever an attribute is set, the value is inserted into the repository and the object state and KOMMA's internal attribute cache are refreshed.

RDFBeans performed slightly better than Empire and KOMMA. Two factors of the implementation of its *add* method are of interest. One is that the method is *synchronized*. This incurs a slight performance overhead because Java has to manage the synchronization mechanism even for the single-threaded benchmark. And it is somewhat unexpected, given that the documentation of RDF4J specifically declares the *RepositoryConnection*, around which RDFBeans *BeanManager* is built, not to be thread-safe and recommends multi-threaded applications to obtain separate connections for each thread²⁹. Second, RDFBeans uses the RDF4J API *hasStatement* method with a bound subject to check for a resource's existence in the repository. In contrast, JOPA uses the same method, but specifies the whole triple, checking whether the subject is an instance of the specified RDFS class. This is not only arguably faster, but it also allows to persist the same entity as an instance of different RDFS classes.

It can be seen that AliBaba and JOPA have very similar performance in *OP1*. In the batch scenario (*OP2*), on the other hand, AliBaba was able to perform a round in under one second, an order of magnitude faster than

²⁹See <https://tinyurl.com/repository-connection>, accessed 2018-11-21.

Table 4

Results of the performance benchmark with all libraries running on a 1 GB heap. (\bar{T}) denotes mean execution time, (σ) standard deviation, CI_{95} 95% confidence interval, and TPS represents average throughput per second (details in Section 7.2.1). O stands for the operation executed and L for the evaluated library. For each operation, the fastest library is underlined

O	OP1 – Create				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	9 155.46	27 628.78	8 975.33	29 121.85	25 599.89
σ (ms)	174.77	124.39	82.86	628074	137.00
CI_{95} (ms)	(9 135.68; 9 175.23)	(27 614.70; 27 642.86)	(8 965.95; 8 984.703)	(29 050.70; 29 193.00)	(25 584.38; 25 615.39)
\overline{TPS} (# of reports)	32.77	10.86	33.42	10.30	11.72
O	OP2 – Batch create				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	738.16	27 624.64	4 725.72	22 875.85	17 325.64
σ (ms)	65.98	176.91	103.95	593.35	121.09
CI_{95} (ms)	(730.70; 745.63)	(27 604.62; 27 644.66)	(4 713.96; 4 737.48)	(22 808.70; 22 942.99)	(17 311.97; 17 339.37)
\overline{TPS} (# of reports)	406.42	10.86	63.48	13.11	171.32
O	OP3 – Retrieve				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	19 838.94	36 632.24	5 894.10	11 819.63	59 486.47
σ (ms)	307.56	2 808.46	88.75	141.01	280.45
CI_{95} (ms)	(19 804.13; 19 873.74)	(36 314.44; 36 950.04)	(5 884.06; 5 904.15)	(11 803.67; 11 835.58)	(59 454.74; 59 518.21)
\overline{TPS} (# of reports)	15.12	8.19	50.90	25.38	5.04
O	OP4 – Retrieve all				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	21 811.98	35 852.54	8 806.90	36 402.01	57 521.18
σ (ms)	729.50	2 782.01	155.64	346.75	391.03
CI_{95} (ms)	(21 729.43; 21 894.53)	(35 537.73; 36 167.35)	(8 789.29; 8 824.51)	(36 362.77; 36 441.25)	(57 476.93; 57 565.43)
\overline{TPS} (# of reports)	13.75	8.19	34.06	8.24	5.22
O	OP5 – Update				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	11 127.57	19 154.26	7 833.98	12 181.53	12 802.06
σ (ms)	176.45	86.91	137.68	253.78	205.95
CI_{95} (ms)	(11 107.60; 11 147.53)	(19 144.43; 19 164.10)	(7 818.40; 7 849.56)	(12 152.83; 12 210.22)	(12 778.76; 12 825.37)
\overline{TPS} (# of reports)	13.48	7.83	19.15	12.31	11.72
O	OP6 – Delete				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	13 479.67	33 567.58	11 800.06	10 271.81	15 912.55
σ (ms)	216.33	739.66	96.92	177.13	106.94
CI_{95} (ms)	(13 455.19; 13 504.15)	(33 483.88; 33 651.28)	(11 789.10; 11 811.03)	(10 251.79; 10 291.88)	(15 900.45; 15 924.65)
\overline{TPS} (# of reports)	11.13	4.47	12.71	14.60	9.43




Fig. 5. Performance of the individual libraries on a 1 GB heap. The plots are grouped by the respective operations. Lower is better.

in *OP1*. The large difference is given by the transaction commit mechanism, which synchronizes a binary object store with its persistent version, potentially stored on the system disk.

Retrieve (OP3) and retrieve all (OP4) JOPA performed the best in both retrieval scenarios. Looking through the source code, all the libraries use different strategies for loading entity attributes.

Alibaba uses a single SPARQL SELECT query to get the subject and all the properties mapped by the entity class. To avoid failed joins, all the property patterns are wrapped in OPTIONAL. However, OPTIONAL increases the query complexity to PSPACE [83].

Empire adopts a different strategy. A SPARQL CONSTRUCT query is issued, retrieving a graph containing all statements concerning the specified subject. Then, it processes the attributes mapped by the target entity class and filters values for the corresponding properties from the constructed graph.

JOPA does not use SPARQL to retrieve entities for two reasons. One is our experience showing that filtering using the RDF4J API is more efficient. The other, and more important, is that a SPARQL query does not allow to specify whether inferred statements should be included or not and the result does not indicate the origin of the resulting statements (see *OCI* in Section 4). Instead, JOPA uses a single call to get all statements concerning the specified subject and then processes these statements, populating the resulting object with appropriate values. While this strategy may turn out inefficient in cases where the object model represents only a tiny portion of the graph around the loaded subject, it does sufficiently well in most cases. Indeed, the results show that it outperforms all the other strategies by a considerable margin.

KOMMA supports two strategies for loading an entity. One relies on a dedicated find method which returns initially an empty object. When a getter is called

on this object, KOMMA loads the corresponding attribute value using a SPARQL SELECT query. Since KOMMA has an internal cache of values, it may happen that the value is already cached and no query has to be performed. The other strategy, employed in the benchmark, uses an explicit SPARQL CONSTRUCT query, which fetches all relevant property values in one call and significantly speeds up instance retrieval. However, the query could not be used in *OP4*, so KOMMA's performance in this operation is considerably worse than in *OP3*.

RDFBeans returns a fully initialized object, but it loads its attributes one by one using the RDF4J filtering API. Moreover, similarly to persist, it first checks for the subject's existence in the repository, which involves a rather costly statement filter with a bound subject. Thus, it was the slowest for both *OP3* and *OP4*.

Update (OP5) *OP5* showed comparable performance by AliBaba, JOPA, KOMMA and RDFBeans, with JOPA outperforming the others. Empire performed the worst. The update is also interesting in terms of the internal implementation in the particular libraries.

AliBaba, because it does not support detached objects, requires the user to first load the required entity and then update it manually, essentially corresponding to the behavior Empire and JOPA implement internally. However, this can become extremely cumbersome – for instance, consider an object coming to the REST API of an application in a PUT request, signifying a data update. Then the application would either have to determine the changes by scanning the incoming object and comparing it to an existing object in the database, or use the incoming object as the new state and replace the existing data completely by iterating over all its attributes and merging their values into the repository. The approach of JOPA (see below) is much more flexible because it allows, for example, to specify whether the merge operation should be cascaded to objects referenced by the entity being merged.

Empire internally loads all the data concerning the subject which is being merged, removes them and then inserts the new, updated, data. While efficient, it is hardly the correct behavior. Consider a situation when the application is using a legacy RDF triple store and the object model covers only a part of the data schema. Invoking merge in Empire will not only delete the mapped data but also all the other statements which have the same subject as the updated instance. The automatic commit of changes in Empire, which, given the setup of *OP5*, did not influence the performance but

which can be harmful from the transactional behavior point of view nonetheless, should also be stressed.

JOPA internally loads the instance into which updates will be merged (the same behavior is expected in JPA [27]). It then compares the merged instance with the loaded one and propagates the changes into the repository. The semantics of update in JOPA is supposed to correspond to JPA as much as possible. For instance, JPA specifies that when merging an entity with references to other objects, these references should not be merged and, actually, if they contain changes, these changes should be overwritten with the data loaded from the repository [27]. Of the benchmarked libraries, only JOPA behaves in this way.

Since KOMMA does not support detached objects, updating an object requires loading a corresponding instance from the repository and updating the required attributes in the same manner as with AliBaba. The optimized instance loading query from *OP3* implementation was used to retrieve the instance to update.

RDFBeans works similarly to Empire in that it simply removes all statements concerning the updated subject and inserts new statements corresponding to the updated object. This, on the one hand, gives RDFBeans a significant performance boost, on the other hand, it can lead to incorrect behavior. As mentioned when discussing Empire, it can lead to the loss of data not mapped by the object model. RDFBeans performed better because it does not load the data to be removed before the actual removal, it directly issues the delete command.

Delete (OP6) Except for Empire, all the libraries exhibited comparable performance for *OP6*, with KOMMA winning by a narrow margin.

AliBaba does not have a dedicated remove method, so entity removal in its case is a little complicated. It consists of loading the object to be removed, setting all its attributes to null and then removing the class designation, i.e., the assertion of the instance's type.

The performance of Empire suffers from the fact that it uses the same strategy for loading statements as in the case of *OP3*, i.e., a SPARQL CONSTRUCT query. This query retrieves all statements with the specified subject. Empire then removes these statements from the storage. Similar to update, this has the potentially harmful effect of removing statements which are not a part of the object model.

To avoid this issue, JOPA employs the *epistemic remove* strategy. Epistemic remove deletes only statements asserting properties mapped by the object model.

Although less efficient than removing all statements concerning an individual, it is safer in terms of data consistency.

Unfortunately, AliBaba, Empire and JOPA do not deal with the situation where the removed resource is an object of other statements, i.e., it could be referenced by other instances (whether mapped by the object model or not). Relational databases resolve this problem with *referential integrity* which does not allow removal of a certain record as long as there are references to it in the form of *foreign keys*. In the open world of ontologies, a reference to an individual completely suffices to prove its existence. However, for an object model, this may not be (and often is not) the case. Therefore, OTM libraries should face this issue. In this regard, KOMMA and RDFBeans come with a solution.

Both KOMMA and RDFBeans perform entity removal by removing all statements whose subject *or* object is the resource being removed. This sufficiently deals with the issue of referential integrity. On the other hand, this solution can lead to unintended data removal. For example, one may remove an instance without knowing by how many other resources (mapped or not) it is referenced. Another user can then load an object which formerly referenced the removed one and find out that the connection has been severed for reasons unknown. We believe that this issue deserves a more thorough and systematic approach.

7.2.2. Scalability

There are multiple ways of testing the scalability of an application or a library. In the case of this research, it was tested how well the libraries scale with regards to the memory available to the application.

Most of the libraries had issues running on the smallest, 32 MB heap. The only library able to perform the whole benchmark on a 32 MB heap was RDFBeans. AliBaba and JOPA were able to execute *OP1*, but failed on the other operations. KOMMA managed to execute *OP1* and *OP2*, but with large standard deviation and many outliers. Empire finished the *OP6* benchmark, but its standard deviation was so large that the results are all but inconclusive. In addition, Empire ran out of memory also for *OP4* on a 64 MB heap. The fact that it did not happen for *OP3*, which is similar to *OP4*, is likely due to the individual instances loaded by *OP3* being discarded right after the verification phase, whereas for *OP4* all the instances are loaded at once and then verified.

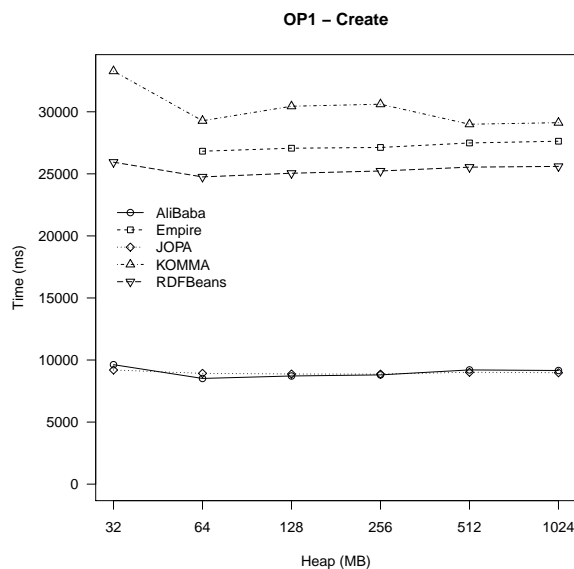


Fig. 6. Benchmark performance of *OP1* w.r.t. heap size. Lower is better.

Contrary to the expectations, with increasing heap size, the benchmark runtime did not decrease. In fact, in some cases, it tended to increase slightly. This increase could be explained by the garbage collector having to manage a larger amount of memory. Overall, once the heap size passed the threshold after which the benchmark did not run out of memory, there was no definitive trend in terms of application performance with respect to the heap size. This can be seen in an example plot for *OP1* in Figure 6 (all plots regarding scalability w.r.t. heap size can be found in Appendix A.2).

To summarize, the benchmark developed as part of the OTM comparison framework appears not suitable for measuring scalability of the OTM libraries. In hindsight, this seems logical. Concerning scalability w.r.t. heap size, once the benchmark data fit into the available memory, increasing the heap size brings no performance benefit. Especially since garbage collection is triggered after the setup of each round. A more useful scalability comparison could involve a variable number of concurrent clients. However, this is not supported by the presented benchmark. On the other hand, a concurrent access-based benchmark would have to be carefully designed in terms of data correctness and objectivity.

7.2.3. Memory

Memory utilization is summarized in Table 5. It shows that both KOMMA and RDFBeans are relatively memory efficient, creating large numbers of short-lived objects which can be disposed of by the less-impactful young generation garbage collection. AliBaba required significantly more full garbage collections but the application throughput remained over 96 %. JOPA's memory utilization fits between AliBaba and KOMMA with throughput nearing 99 %. This has to be considered in the context of performance, where both AliBaba and JOPA clearly outperform RDFBeans and, in most operations, also KOMMA, indicating that especially JOPA uses longer-living objects as a means to cache relevant data or to fetch the data more efficiently in terms of execution time.

Finally, the throughput of the benchmark application running Empire was approximately 85 %, significantly less than the other libraries. However, when we tried the same experiment, but only with a 1 h runtime, its throughput was comparable to the other libraries. We suspect that there may be a memory leak in Empire. For the shorter runtime, the leak had not enough time to fully manifest. It would also explain Empire failing to execute *OP4* on a 64 MB heap.

8. Conclusions

We have introduced a novel framework for comparison of object-triple mapping libraries. It consists of a set of qualitative criteria describing features important for developers of domain-specific object-oriented applications and a performance benchmark written in Java. The framework was used to compare a diverse set of OTM libraries. The results indicate that significant differences exist between the evaluated OTM providers in terms of features, performance and their treatment of semantic data – for instance, some expect the object model to completely cover the schema of the data, others also support access to a subset of the schema.

8.1. Feature Comparison Conclusions

The evaluated libraries may be split into several groups based on different points of view. One contains solutions which appear to be proof-of-concept prototypes with limited functionality. This group contains AutoRDF, JAOb, and the Semantic Framework – their usability is restricted to a rather small set of specific

cases. Another group consists of libraries in which entities are merely proxies of the underlying RDF data. Such libraries could be problematic especially in larger (web) applications because they generally do not support detached objects. On the other hand, thanks to the fact that they use dynamically generated proxies, they are able to support a wide variety of shared behavior in entities. In case of ActiveRDF and SuRF, this is enhanced by the fact that they do not require a predefined object model. The last group may be denoted as general purpose persistence libraries and it contains Empire, JOPA and RDFBeans. These libraries attempt to provide functionality similar to JPA but in the world of ontologies.

In general, most of the libraries suffer from the fact that there is no formalization of the object-triple mapping. Such a gap leads to potentially inconsistent behavior and straightforward solutions which fail to handle less frequent cases. To give a final recommendation, for multiuser applications with non-trivial business logic, we believe that support for transactions (*GC1*), detached objects (*GC2*), query result mapping (*GC3*) and inheritance mapping (*MC1*) are a must have for an OTM library. Explicit treatment of inferred knowledge (*OCI*) would be important in scenarios where the mapped data are involved in inference.

8.2. Performance Comparison Conclusions

Several conclusions may be derived from the performance comparison results. AliBaba demonstrated very good performance and supports a rich set of features. However, while restriction to Sesame-based storage is not a big problem thanks to its wide support among triple store implementors, the unintuitive cascading strategy and lack of support for detached objects can be of concern. Empire exhibited worse than average time performance and a potential memory leak in query mapping. On the other hand, its API directly implements portions of the JPA standard [27] which can make it suitable for existing projects migrating from relational to semantic databases. JOPA provides rich features as well as excellent performance, memory utilization, and sound data access patterns for the general case. The performance of KOMMA critically depends on the use of optimization hints provided by its authors. Without them, KOMMA performed in most cases significantly worse than the other libraries. With them, its performance rivals the performance of AliBaba and JOPA. However, its API made it more difficult to use in the benchmark than the other libraries.

Table 5

Memory utilization summary. *YGC* (*FGC*) is young generation (full) garbage collection event count, *GCT* is the total time spent in garbage collection and *Throughput* is the application throughput

Measure	AliBaba	Empire	JOPA	KOMMA	RDFBeans
YGC	14 311	14 169	26 049	44 706	29 658
FGC	11 685	13 320	2 586	98	51
GCT (s)	496.98	1 326.58	147.78	65.36	40.4
Throughput (%)	96.37	84.56	98.97	99.55	99.72

RDFBeans is the most suitable library for environments with extremely limited memory. In comparison to the other four libraries evaluated in the performance benchmark, its time performance is, apart from retrieval operations, average.

Overall, while the slowest player in an ontology-based application stack will always be the triple store (given the size and complexity of data in processes), the benchmark showed that OTM libraries significantly differ in terms of performance. These differences are given by the effectiveness with which OTM solutions access the underlying storage. Application developers thus may choose a library suitable for their case, e.g., RDFBeans for applications with limited resources, AliBaba for batch processing scenarios, JOPA for general cases. However, we argue that more important than performance differences are the differences in the semantics of the implementation, where, for example, Empire and RDFBeans can have unexpected side-effects in scenarios where the mapping does not cover the whole RDF data schema.

Regarding the libraries' memory scalability and efficiency, the benchmark indicates that heap size is of little concern. Most real-world applications nowadays use vastly more than 32 MB of heap. None of the evaluated solutions showed any memory-related performance issues after crossing this threshold (with the exception of a possible memory leak in Empire). However, as was stated earlier, a more realistic concurrent access-based scenario running for a longer period of time could be more revealing.

The performance benchmark was designed so that adding new libraries into the comparison would be relatively straightforward. However, attention must be paid to providing equal conditions to all compared libraries. For this reason, the evaluation provided in the paper concentrated on a single triple store implementation. The performance benchmark is published online and so are the results of the comparison presented in this work. We believe that these results might help po-

tential users of OTM libraries in deciding which one to choose. In addition, the benchmark can be considered an example application for each of the libraries, complete with an object model declaration and the persistence layer setup and usage. Given the scarce documentation of most of the evaluated libraries, this can be a welcome benefit.

In the future, the benchmark should be extended to allow evaluating scalability w.r.t. multiple concurrent users, providing an even more realistic test case for the OTM libraries.

Acknowledgements

We would like to thank the reviewers for their detailed and insightful comments. This work was supported by grants No. GA 16-09713S Efficient Exploration of Linked Data Cloud of the Grant Agency of the Czech Republic and No.SGS16/229/OHK3/3T/13 Supporting ontological data quality in information systems of the Czech Technical University in Prague.

Appendix A. Complete benchmark results

A.1. Performance

This section contains plots depicting performance of the evaluated libraries in the benchmark. Figure 7 shows results for a 32 MB heap, Figure 8 for a 64 MB heap, Figure 9 for a 128 MB heap, Figure 10 for a 256 MB heap, Figure 11 for a 512 MB heap and Figure 12 for a 1 GB heap. The tables contain the mean execution time, standard deviation, confidence intervals, and throughput in reports per second for each operation and each library. Each table represents benchmark results for a particular heap size. Table 6 contains results for a 32 MB heap, Table 7 for a 64 MB heap, Table 8 for a 128 MB heap, Table 9 for a 256 MB heap, Table 10 for a 512 MB heap, and Table 11 for a 1 GB heap.




Fig. 7. Performance of the individual libraries on a 32 MB heap. The plots are grouped by the respective operations. Lower is better.




Fig. 8. Performance of the individual libraries on a 64 MB heap. The plots are grouped by the respective operations. Lower is better.




Fig. 9. Performance of the individual libraries on a 128 MB heap. The plots are grouped by the respective operations. Lower is better.




Fig. 10. Performance of the individual libraries on a 256 MB heap. The plots are grouped by the respective operations. Lower is better.




Fig. 11. Performance of the individual libraries on a 512 MB heap. The plots are grouped by the respective operations. Lower is better.




Fig. 12. Performance of the individual libraries on a 1 GB heap. The plots are grouped by the respective operations. Lower is better.

Table 6

Results of the performance benchmark with all libraries running on a 32 MB heap. (\bar{T}) denotes mean execution time, (σ) standard deviation, CI_{95} 95% confidence interval, and TPS represents average throughput per second (details in Section 7.2.1). O stands for the operation executed and L for the evaluated library. For each operation, the fastest library is underlined

O <i>OP1 – Create</i>					
L	AliBaba	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	9 624.75	×	9 195.53	33 273.05	25 940.71
σ (ms)	231.66	×	122.52	7 356.89	123.18
CI_{95} (ms)	(9 598.54; 9 650.97)	×	(9 181.66; 9 209.39)	(32 440.56; 34 105.55)	(25 926.77; 25 954.65)
\overline{TPS} (# of reports)	31.17	×	32.64	9.01	11.56
O <i>OP2 – Batch create</i>					
L	AliBaba	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	×	×	×	26 316.05	17 398.23
σ (ms)	×	×	×	6 372.84	181.18
CI_{95} (ms)	×	×	×	(25 594.91; 27 037.19)	(17 377.73; 17 418.73)
\overline{TPS} (# of reports)	×	×	×	11.40	17.24
O <i>OP3 – Retrieve</i>					
L	AliBaba	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	×	×	×	×	59 810.57
σ (ms)	×	×	×	×	304.38
CI_{95} (ms)	×	×	×	×	(59 776.13; 59 845.01)
\overline{TPS} (# of reports)	×	×	×	×	5.02
O <i>OP4 – Retrieve all</i>					
L	AliBaba	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	×	×	×	×	57 778.39
σ (ms)	×	×	×	×	362.33
CI_{95} (ms)	×	×	×	×	(57 737.39; 57 819.39)
\overline{TPS} (# of reports)	×	×	×	×	5.19
O <i>OP5 – Update</i>					
L	AliBaba	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	×	×	×	×	12 982.18
σ (ms)	×	×	×	×	142.54
CI_{95} (ms)	×	×	×	×	(12 966.05; 12 998.31)
\overline{TPS} (# of reports)	×	×	×	×	11.55
O <i>OP6 – Delete</i>					
L	AliBaba	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	×	40 753.17	×	×	16 169.85
σ (ms)	×	8 865.69	×	×	168.68
CI_{95} (ms)	×	(39 749.95; 41 756.40)	×	×	(16 150.77; 16 188.94)
\overline{TPS} (# of reports)	×	3.68	×	×	9.28

Table 7

Results of the performance benchmark with all libraries running on a 64MB heap. (\bar{T}) denotes mean execution time, (σ) standard deviation, CI_{95} 95% confidence interval, and TPS represents average throughput per second (details in Section 7.2.1). O stands for the operation executed and L for the evaluated library. For each operation, the fastest library is underlined

O	<i>OP1 – Create</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	8 469.15	26 816.89	8 925.29	29 276.15	24 752.90
σ (ms)	180.32	271.77	117.75	549.30	168.57
CI_{95} (ms)	(8 448.77; 8 489.56)	(26 786.14; 26 847.64)	(8 911.96; 8 938.61)	(29 213.99; 29 338.31)	(24 733.82; 24 771.97)
\overline{TPS} (# of reports)	35.42	11.19	33.61	10.25	12.12
O	<i>OP2 – Batch create</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	816.98	26 817.00	4 846.23	22 949.73	17 288.31
σ (ms)	54.10	183.64	101.83	435.75	146.06
CI_{95} (ms)	(810.86; 823.10)	(26 796.22; 26 837.78)	(4 834.71; 4 857.76)	(22 900.42; 22 999.04)	(17 271.78; 17 304.84)
\overline{TPS} (# of reports)	367.20	11.12	61.90	13.07	17.35
O	<i>OP3 – Retrieve</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	21 531.18	39 095.80	6 057.05	12 247.85	59 460.36
σ (ms)	286.97	3 620.09	94.56	1 019.82	353.73
CI_{95} (ms)	(21 498.71; 21 563.66)	(38 686.15; 39 505.44)	(6 046.35; 6 067.75)	(12 132.44; 12 363.25)	(59 420.33; 59 500.39)
\overline{TPS} (# of reports)	13.93	7.67	49.53	24.49	5.05
O	<i>OP4 – Retrieve all</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	21 904.37	×	8 905.65	27 237.13	57 410.77
σ (ms)	725.06	×	137.40	3 876.60	260.54
CI_{95} (ms)	(21 822.33; 21 986.42)	×	(8 890.01; 8 921.19)	(37 798.46; 38 675.80)	(57 381.28; 57 440.25)
\overline{TPS} (# of reports)	13.70	×	33.69	11.01	5.23
O	<i>OP5 – Update</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	10 694.26	18 356.25	7 933.59	12 374.00	12 339.22
σ (ms)	133.35	118.52	127.87	624.33	109.96
CI_{95} (ms)	(10 679.17; 10 709.35)	(18 342.84; 18 369.66)	(7 919.12; 7 948.06)	(12 303.35; 12 444.65)	(12 326.77; 12 351.66)
\overline{TPS} (# of reports)	14.03	8.17	18.91	12.12	12.16
O	<i>OP6 – Delete</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	13 052.14	33 955.28	11 902.21	10 337.81	15 491.79
σ (ms)	169.87	1 298.76	105.45	379.44	117.26
CI_{95} (ms)	(13 032.92; 13 071.36)	(33 808.31; 34 102.24)	(11 890.28; 11 914.14)	(10 294.87; 10 380.74)	(15 478.52; 15 505.06)
\overline{TPS} (# of reports)	11.49	4.42	12.60	14.51	9.68

Table 8

Results of the performance benchmark with all libraries running on a 128 MB heap. (\bar{T}) denotes mean execution time, (σ) standard deviation, CI_{95} 95% confidence interval, and TPS represents average throughput per second (details in Section 7.2.1). O stands for the operation executed and L for the evaluated library. For each operation, the fastest library is underlined

O	<i>OP1 – Create</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	8 713.70	27 064.14	8 873.12	30 446.10	25 048.47
σ (ms)	127.39	189.82	106.47	496.49	225.96
CI_{95} (ms)	(8 699.28; 8 728.11)	(27 042.66; 27 085.62)	(8 861.07; 8 885.16)	(30 389.92; 30 502.28)	(25 022.90; 25 074.04)
\overline{TPS} (# of reports)	34.43	11.08	33.81	9.85	11.98
O	<i>OP2 – Batch create</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	739.42	27 153.00	4 746.37	23 175.81	17 304.30
σ (ms)	72.39	159.27	106.78	626.51	85.99
CI_{95} (ms)	(731.23; 747.61)	(27 134.97; 27 171.02)	(4 734.28; 4 758.45)	(23 104.91; 23 246.70)	(17 294.57; 17 314.03)
\overline{TPS} (# of reports)	405.72	11.05	63.12	12.94	17.34
O	<i>OP3 – Retrieve</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	21 527.06	37 651.89	5 959.37	11 990.78	59 569.55
σ (ms)	226.91	3 184.72	84.30	183.12	268.05
CI_{95} (ms)	(21 501.38; 21 552.73)	(37 291.51; 38 012.26)	(5 949.83; 5 968.91)	(11 970.06; 12 011.50)	(59 539.22; 59 599.89)
\overline{TPS} (# of reports)	13.94	7.97	50.34	25.02	5.04
O	<i>OP4 – Retrieve all</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	21 875.11	36 230.50	8 858.86	36 507.41	57 370.60
σ (ms)	784.56	1 752.28	139.63	550.50	263.63
CI_{95} (ms)	(21 786.33; 21 963.89)	(36 032.21; 36 428.78)	(8 843.06; 8 874.67)	(36 445.12; 36 569.70)	(57 340.76; 57 400.43)
\overline{TPS} (# of reports)	13.71	8.28	33.86	8.22	5.23
O	<i>OP5 – Update</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	10 906.91	18 693.81	7 811.44	12 169.35	12 524.04
σ (ms)	145.07	133.32	121.96	293.84	120.72
CI_{95} (ms)	(10 890.50; 10 923.33)	(18 678.73; 18 708.90)	(7 797.64; 7 825.25)	(12 136.10; 12 202.60)	(12 510.38; 12 537.70)
\overline{TPS} (# of reports)	13.75	8.02	19.20	12.33	11.98
O	<i>OP6 – Delete</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	13 231.23	33 625.23	11 821.90	10 209.39	15 717.66
σ (ms)	151.42	994.90	105.45	159.48	157.61
CI_{95} (ms)	(13 214.10; 13 248.36)	(33 512.65; 33 737.81)	(11 809.97; 11 833.84)	(10 191.35; 10 227.44)	(15 699.82; 15 735.49)
\overline{TPS} (# of reports)	11.34	4.46	12.69	14.69	9.54

Table 9

Results of the performance benchmark with all libraries running on a 256 MB heap. (\bar{T}) denotes mean execution time, (σ) standard deviation, CI_{95} 95% confidence interval, and TPS represents average throughput per second (details in Section 7.2.1). O stands for the operation executed and L for the evaluated library. For each operation, the fastest library is underlined

O	<i>OPI – Create</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	8 801.82	27 126.05	8 866.54	30 610.41	25 229.68
σ (ms)	125.52	122.86	109.68	702.99	256.97
CI_{95} (ms)	(8 787.61; 8 816.02)	(27 112.15; 27 139.95)	(8 854.12; 8 878.95)	(30 530.86; 30 689.96)	(25 200.60; 25 258.75)
\overline{TPS} (# of reports)	34.08	11.06	33.84	9.80	11.89
O	<i>OP2 – Batch create</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	782.93	27 229.57	4 714.30	23 290.91	17 311.32
σ (ms)	67.60	103.90	142.41	430.83	94.71
CI_{95} (ms)	(775.28; 790.58)	(27 211.89; 27 247.24)	(4 702.54; 4 726.05)	(23 242.16; 23 339.66)	(17 300.61; 17 322.04)
\overline{TPS} (# of reports)	383.18	11.01	63.64	12.88	17.33
O	<i>OP3 – Retrieve</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	21 115.87	36 909.70	5 906.61	11 972.34	59 441.71
σ (ms)	494.35	2 868.89	83.56	143.42	294.53
CI_{95} (ms)	(21 059.93; 21 171.81)	(36 585.06; 37 234.34)	(5 897.16; 5 916.07)	(11 956.11; 11 988.56)	(59 408.38; 59 475.03)
\overline{TPS} (# of reports)	14.21	8.13	50.79	25.06	5.05
O	<i>OP4 – Retrieve all</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	21 802.91	36 421.56	8 833.37	36 659.15	57 491.72
σ (ms)	770.91	3 013.25	95.78	401.20	271.83
CI_{95} (ms)	(21 715.68; 21 890.15)	(36 080.59; 36 762.54)	(8 822.53; 8 844.21)	(36 613.75; 36 704.55)	(57 460.96; 57 522.48)
\overline{TPS} (# of reports)	13.76	8.24	33.96	8.18	5.22
O	<i>OP5 – Update</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	10 998.84	18 786.56	7 806.32	12 209.20	12 622.81
σ (ms)	195.20	110.43	126.62	196.00	126.15
CI_{95} (ms)	(10 976.75; 11 020.92)	(18 774.06; 18 799.05)	(7 791.99; 7 820.65)	(12 187.02; 12 231.38)	(12 608.53; 12 637.08)
\overline{TPS} (# of reports)	13.64	7.98	19.22	12.29	11.88
O	<i>OP6 – Delete</i>				
L	<u>AliBaba</u>	Empire	JOPA	KOMMA	RDFBeans
\bar{T} (ms)	13 360.98	33 359.14	11 774.40	10 270.25	15 806.75
σ (ms)	203.04	861.21	105.93	169.73	109.86
CI_{95} (ms)	(13 338.01; 13 383.96)	(33 261.69; 33 456.60)	(11 762.41; 11 786.39)	(10 251.04; 10 289.45)	(15 794.32; 15 819.19)
\overline{TPS} (# of reports)	11.23	4.50	12.74	14.61	9.49

Table 10

Results of the performance benchmark with all libraries running on a 512 MB heap. (\bar{T}) denotes mean execution time, (σ) standard deviation, CI_{95} 95% confidence interval, and TPS represents average throughput per second (details in Section 7.2.1). O stands for the operation executed and L for the evaluated library. For each operation, the fastest library is underlined

O	<i>OP1 – Create</i>				
L	<u>AliBaba</u>	<u>Empire</u>	<u>JOPA</u>	<u>KOMMA</u>	<u>RDFBeans</u>
\bar{T} (ms)	9 201.87	27 488.23	9 012.40	28 996.45	25 540.07
σ (ms)	149.49	202.24	155.23	537.78	174.80
CI_{95} (ms)	(9 184.95; 9 218.79)	(27 465.34; 27 511.11)	(8 994.83; 9 029.96)	(28 935.59; 29 057.30)	(25 520.29; 25 559.85)
\overline{TPS} (# of reports)	32.60	10.91	33.29	10.35	11.75
O	<i>OP2 – Batch create</i>				
L	<u>AliBaba</u>	<u>Empire</u>	<u>JOPA</u>	<u>KOMMA</u>	<u>RDFBeans</u>
\bar{T} (ms)	711.93	27 568.86	4 737.19	22 864.82	17 323.95
σ (ms)	64.86	303.20	105.23	561.91	120.04
CI_{95} (ms)	(704.59; 719.27)	(27 534.55; 27 603.17)	(4 725.28; 4 749.10)	(22 801.24; 22 928.41)	(17 310.36; 17 337.53)
\overline{TPS} (# of reports)	421.39	10.88	63.33	13.12	17.32
O	<i>OP3 – Retrieve</i>				
L	<u>AliBaba</u>	<u>Empire</u>	<u>JOPA</u>	<u>KOMMA</u>	<u>RDFBeans</u>
\bar{T} (ms)	20 507.68	36 729.19	5 882.51	11 828.92	59 516.95
σ (ms)	208.02	2 815.72	83.42	117.86	272.68
CI_{95} (ms)	(20 484.14; 20 531.22)	(36 410.57; 37 047.81)	(5 873.07; 5 891.95)	(11 815.59; 11 842.26)	(59 486.09; 59 547.80)
\overline{TPS} (# of reports)	14.63	8.17	51.00	25.36	5.04
O	<i>OP4 – Retrieve all</i>				
L	<u>AliBaba</u>	<u>Empire</u>	<u>JOPA</u>	<u>KOMMA</u>	<u>RDFBeans</u>
\bar{T} (ms)	21 905.64	36 096.34	8 866.55	36 594.84	57 470.96
σ (ms)	778.16	2 838.06	158.25	544.96	364.24
CI_{95} (ms)	(21 817.59; 21 993.70)	(35 775.19; 36 417.49)	(8 848.64; 8 884.46)	(36 533.19; 36 656.48)	(57 429.75; 57 512.18)
\overline{TPS} (# of reports)	13.70	8.31	33.84	8.20	5.22
O	<i>OP5 – Update</i>				
L	<u>AliBaba</u>	<u>Empire</u>	<u>JOPA</u>	<u>KOMMA</u>	<u>RDFBeans</u>
\bar{T} (ms)	11 121.83	19 154.90	7 847.03	12 204.42	12 765.91
σ (ms)	170.39	133.68	121.25	176.62	142.05
CI_{95} (ms)	(11 102.55; 11 141.11)	(19 139.77; 19 170.02)	(7 833.31; 7 860.75)	(12 184.43; 12 224.40)	(12 749.83; 12 781.98)
\overline{TPS} (# of reports)	13.49	7.83	19.12	12.29	11.75
O	<i>OP6 – Delete</i>				
L	<u>AliBaba</u>	<u>Empire</u>	<u>JOPA</u>	<u>KOMMA</u>	<u>RDFBeans</u>
\bar{T} (ms)	13 520.13	33 562.09	11 792.30	10 246.78	15 965.50
σ (ms)	243.72	768.82	110.50	200.21	122.67
CI_{95} (ms)	(13 492.55; 13 547.71)	(33 475.09; 33 649.09)	(11 779.79; 11 804.80)	(10 224.12; 10 269.43)	(15 951.61; 15 979.38)
\overline{TPS} (# of reports)	11.09	4.47	12.72	14.64	9.40

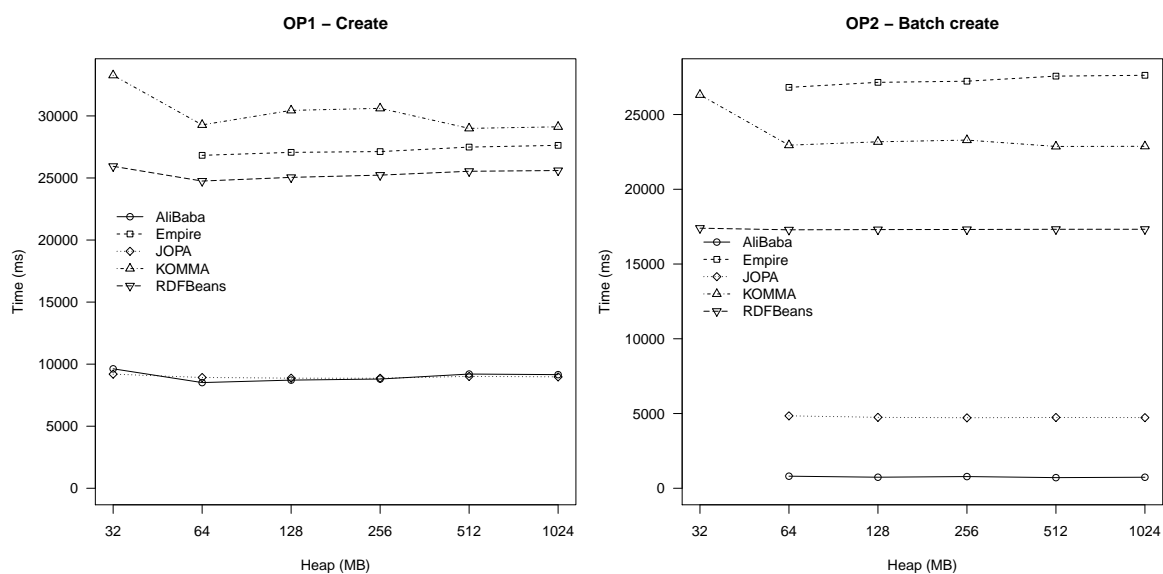
Table 11

Results of the performance benchmark with all libraries running on a 1 GB heap. (\bar{T}) denotes mean execution time, (σ) standard deviation, CI_{95} 95% confidence interval, and TPS represents average throughput per second (details in Section 7.2.1). O stands for the operation executed and L for the evaluated library. For each operation, the fastest library is underlined

O	<i>OP1 – Create</i>				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	9 155.46	27 628.78	8 975.33	29 121.85	25 599.89
σ (ms)	174.77	124.39	82.86	628074	137.00
CI_{95} (ms)	(9 135.68; 9 175.23)	(27 614.70; 27 642.86)	(8 965.95; 8 984.703)	(29 050.70; 29 193.00)	(25 584.38; 25 615.39)
\overline{TPS} (# of reports)	32.77	10.86	33.42	10.30	11.72
O	<i>OP2 – Batch create</i>				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	738.16	27 624.64	4 725.72	22 875.85	17 325.64
σ (ms)	65.98	176.91	103.95	593.35	121.09
CI_{95} (ms)	(730.70; 745.63)	(27 604.62; 27 644.66)	(4 713.96; 4 737.48)	(22 808.70; 22 942.99)	(17 311.97; 17 339.37)
\overline{TPS} (# of reports)	406.42	10.86	63.48	13.11	171.32
O	<i>OP3 – Retrieve</i>				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	19 838.94	36 632.24	5 894.10	11 819.63	59 486.47
σ (ms)	307.56	2 808.46	88.75	141.01	280.45
CI_{95} (ms)	(19 804.13; 19 873.74)	(36 314.44; 36 950.04)	(5 884.06; 5 904.15)	(11 803.67; 11 835.58)	(59 454.74; 59 518.21)
\overline{TPS} (# of reports)	15.12	8.19	50.90	25.38	5.04
O	<i>OP4 – Retrieve all</i>				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	21 811.98	35 852.54	8 806.90	36 402.01	57 521.18
σ (ms)	729.50	2 782.01	155.64	346.75	391.03
CI_{95} (ms)	(21 729.43; 21 894.53)	(35 537.73; 36 167.35)	(8 789.29; 8 824.51)	(36 362.77; 36 441.25)	(57 476.93; 57 565.43)
\overline{TPS} (# of reports)	13.75	8.19	34.06	8.24	5.22
textbfO	<i>OP5 – Update</i>				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	11 127.57	19 154.26	7 833.98	12 181.53	12 802.06
σ (ms)	176.45	86.91	137.68	253.78	205.95
CI_{95} (ms)	(11 107.60; 11 147.53)	(19 144.43; 19 164.10)	(7 818.40; 7 849.56)	(12 152.83; 12 210.22)	(12 778.76; 12 825.37)
\overline{TPS} (# of reports)	13.48	7.83	19.15	12.31	11.72
O	<i>OP6 – Delete</i>				
L	<u>AliBaba</u>	Empire	<u>JOPA</u>	KOMMA	RDFBeans
\bar{T} (ms)	13 479.67	33 567.58	11 800.06	10 271.81	15 912.55
σ (ms)	216.33	739.66	96.92	177.13	106.94
CI_{95} (ms)	(13 455.19; 13 504.15)	(33 483.88; 33 651.28)	(11 789.10; 11 811.03)	(10 251.79; 10 291.88)	(15 900.45; 15 924.65)
\overline{TPS} (# of reports)	11.13	4.47	12.71	14.60	9.43

A.2. Scalability

This section contains figures illustrating the scalability of the selected OTM libraries w.r.t. heap size for all operations. Figure 13 shows performance results depending on heap size for *OP1*, *OP2*, *OP3* and *OP4*. Figure 14 then for *OP5* and *OP6*.



(a) Benchmark performance of *OP1* w.r.t. heap size.

(b) Benchmark performance of *OP2* w.r.t. heap size.

Fig. 13. Scalability w.r.t. heap size. Lower is better.

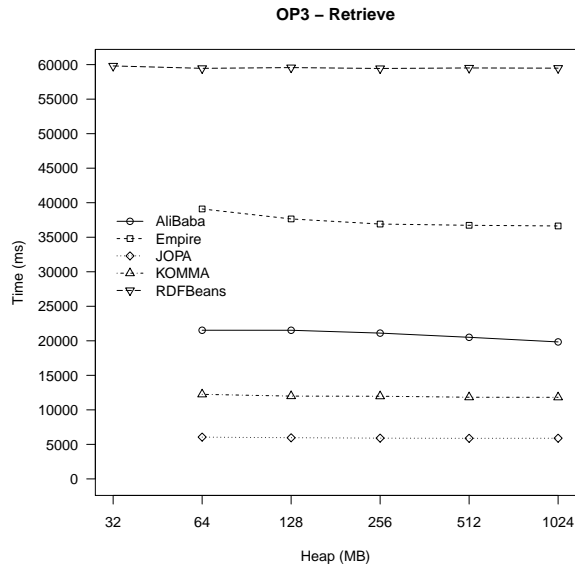
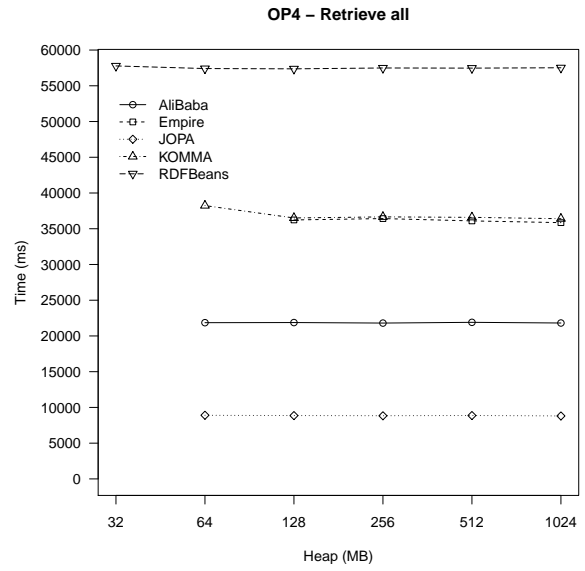


(a) Benchmark performance of *OP3* w.r.t. heap size.(b) Benchmark performance of *OP4* w.r.t. heap size.(c) Benchmark performance of *OP5* w.r.t. heap size.(d) Benchmark performance of *OP6* w.r.t. heap size.

Fig. 14. Scalability w.r.t. heap size. Lower is better.

References

- [1] T. Berners-Lee, J. Hendler and O. Lassila, The Semantic Web, *Scientific American* **284**(5) (2001), 28–37.
- [2] A. Kalyanpur, B.K. Boguraev, S. Patwardhan, J.W. Murdock, A. Lally, C. Welty, J.M. Prager, B. Coppola, A. Fokoue-Nkoutche, L. Zhang, Y. Pan and Z.M. Qiu, Structured Data and Inference in DeepQA, *IBM J. Res. Dev.* **56**(3) (2012), 351–364, ISSN 0018-8646. doi:10.1147/JRD.2012.2188737. <http://dx.doi.org/10.1147/JRD.2012.2188737>.
- [3] T. Heath and C. Bizer, *Linked Data: Evolving the Web into a Global Data Space*, 1st edn, Morgan & Claypool, 2011. ISBN 9781608454303. <http://linkeddatabook.com/>.
- [4] D. Wood, M. Zaidman, L. Ruth and M. Hausenblas, *Linked Data: Structured Data on the Web*, Manning Publications Co., 2014.
- [5] M. Ledvinka, A. Lališ and P. Křemen, Towards Data-Driven Safety: An Ontology-Based Information System, *Journal of Aerospace Information Systems* (2018), 1–15, ISSN 2327-3097. doi:10.2514/1.1010622.
- [6] R.E. Carvalho, J. Williams, I. Sturken, R. Keller and T. Panontin, Investigation Organizer: the development and testing of a Web-based tool to support mishap investigations, in: *2005 IEEE Aerospace Conference*, 2005, pp. 89–98, ISSN 1095-323X. doi:10.1109/AERO.2005.1559302.
- [7] R. Carvalho, S. Wolfe, D. Berrios and J. Williams, Ontology Development and Evolution in the Accident Investigation Domain, in: *2005 IEEE Aerospace Conference*, 2005, pp. 1–8, ISSN 1095-323X. doi:10.1109/AERO.2005.1559634.
- [8] J.J. Carroll, I. Dickinson, C. Dollin, D. Reynolds, A. Seaborne and K. Wilkinson, Jena: implementing the semantic web recommendations, in: *Proceedings of the 13th international World Wide Web conference (Alternate Track Papers & Posters)*, 2004, pp. 74–83.
- [9] J. Broekstra, A. Kampman and F. van Harmelen, Sesame: A Generic Architecture for Storing and Querying RDF and RDF Schema, in: *Proceedings of the First International Semantic Web Conference on The Semantic Web*, 2002, pp. 54–68.
- [10] R. Cyganiak, D. Wood and M. Lanthaler, RDF 1.1 Concepts and Abstract Syntax, Technical Report, W3C, 2014.
- [11] D. Brickley and R.V. Guha, RDF Schema 1.1, W3C Recommendation, W3C, 2014.
- [12] S. Bechhofer, F. van Harmelen, J. Hendler, I. Horrocks, D.L. McGuinness, P.F. Patel-Schneider and L.A. Stein, OWL Web Ontology Language, W3C Recommendation, W3C, 2004.
- [13] B. Motik, B. Parsia and P.F. Patel-Schneider, OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax, W3C Recommendation, W3C, 2009.
- [14] D. Beckett, T. Berners-Lee, E. Prud'hommeaux and G. Carothers, RDF 1.1 Turtle, W3C Recommendation, W3C, 2014, <https://www.w3.org/TR/turtle/>, accessed 08-08-2017..
- [15] S. Harris and A. Seaborne, SPARQL 1.1 Query Language, Technical Report, W3C, 2013.
- [16] P. Gearon, A. Passant and A. Polleres, SPARQL 1.1 Update, Technical Report, W3C, 2013, Accessed 2018-01-08. <https://www.w3.org/TR/sparql11-update>.
- [17] B. Bishop, A. Kiryakov, D. Ognyanoff, I. Peikov, Z. Tashev and R. Velkov, OWLIM: A family of scalable semantic repositories, *Semantic Web – Interoperability, Usability, Applicability* (2010).
- [18] O. Erling and I. Mikhailov, RDF Support in the Virtuoso DBMS., in: *Conference on Social Semantic Web*, S. Auer, C. Bizer, C. Müller and A.V. Zhdanova, eds, LNI, Vol. 113, GI, 2007, pp. 59–68. ISBN 978-3-88579-207-9. <http://dblp.uni-trier.de/db/conf/cssw/cssw2007.html#ErlingM07>.
- [19] Microsoft, Microsoft Open Database Connectivity (ODBC), Microsoft Corporation, 2016. https://cdn.simba.com/wp-content/uploads/2016/03/ODBC_specification.pdf.
- [20] P. Křemen and Z. Kouba, Ontology-Driven Information System Design, *IEEE Transactions on Systems, Man, and Cybernetics: Part C* **42**(3) (2012), 334–344, ISSN 1094-6977.
- [21] M. Horridge and S. Bechhofer, The OWL API: A Java API for OWL ontologies, *Semantic Web – Interoperability, Usability, Applicability* (2011).
- [22] J. Leigh, AliBaba, 2007, Accessed 2018-01-02. <https://bitbucket.org/openrdf/alibaba/>.
- [23] M. Grove, Empire: RDF & SPARQL Meet JPA, *semanticweb.com* (2010). http://semanticweb.com/empire-rdf-sparql-meet-jpa_b15617.
- [24] M. Ledvinka and P. Křemen, JOPA: Accessing Ontologies in an Object-oriented Way, in: *Proceedings of the 17th International Conference on Enterprise Information Systems*, 2015.
- [25] C. Puleston, B. Parsia, J. Cunningham and A. Rector, Integrating Object-Oriented and Ontological Representations: A Case Study in Java and OWL, in: *The Semantic Web - ISWC 2008*, A. Sheth, S. Staab, M. Dean, M. Paolucci, D. Maynard, T. Finin and K. Thirunarayan, eds, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 130–145. ISBN 978-3-540-88564-1.
- [26] JCP, JDBCTM 4.2 Specification, Oracle America, Inc., 2014.
- [27] JCP, JSR 317: JavaTM Persistence API, Version 2.0, Sun Microsystems, 2009.
- [28] Standard Performance Evaluation Corporation, SPECjbb®2015, 2015, Accessed 2018-01-04. <https://www.spec.org/jbb2015/>.
- [29] S.M. Blackburn, R. Garner, C. Hoffmann, A.M. Khang, K.S. McKinley, R. Bentzur, A. Diwan, D. Feinberg, D. Frampton, S.Z. Guyer, M. Hirzel, A. Hosking, M. Jump, H. Lee, J.E.B. Moss, A. Phansalkar, D. Stefanović, T. VanDrunen, D. von Dincklage and B. Wiedermann, The DaCapo Benchmarks: Java Benchmarking Development and Analysis, in: *Proceedings of the 21st Annual ACM SIGPLAN Conference on Object-oriented Programming Systems, Languages, and Applications*, OOPSLA '06, ACM, New York, NY, USA, 2006, pp. 169–190. ISBN 1-59593-348-4. doi:10.1145/1167473.1167488. <http://doi.acm.org/10.1145/1167473.1167488>.
- [30] G. Gousios, V. Karakoidas and D. Spinellis, Tuning Java's memory manager for high performance server applications, in: *Proceedings of the 5th International System Administration and Network Engineering Conference SANE 06*, A. Zavras, ed., Stichting SANE, 2006, pp. 69–83, NLUUG.
- [31] A. Georges, D. Buytaert and L. Eeckhout, Statistically Rigorous Java Performance Evaluation, in: *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object-oriented Programming Systems and Applications*, OOPSLA '07, ACM, New York, NY, USA, 2007, pp. 57–76. ISBN 978-1-59593-786-5. doi:10.1145/1297027.1297033. <http://doi.acm.org/10.1145/1297027.1297033>.
- [32] O. Holanda, S. Isotani, I.I. Bittencourt, D. Dermeval and W. Alcantara, An Object Triple Mapping System Supporting Detached Objects, *Engineering Applications of Arti-*

- cial Intelligence* **62**(C) (2017), 234–251, ISSN 0952-1976. doi:10.1016/j.engappai.2017.04.010. <https://doi.org/10.1016/j.engappai.2017.04.010>.
- [33] P. Cristofaro, Virtuoso RDF Triple Store Analysis Benchmark & mapping tools RDF / OO, 2013, Accessed 2018-01-02. <https://tinyurl.com/virt-rdf-map-tools>.
- [34] K. Schlegel, Selecting an RDF mapping library for cross-media enhancements, 2015, Accessed 2018-01-04. <https://tinyurl.com/sel-rdf-map-lib>.
- [35] Y. Guo, A. Qasem, Z. Pan and J. Heflin, A Requirements Driven Framework for Benchmarking Semantic Web Knowledge Base Systems, *IEEE Transactions on Knowledge and Data Engineering* **19**(2) (2007), 297–309.
- [36] Y. Guo, Z. Pan and J. Heflin, LUBM: A benchmark for OWL knowledge base systems, *Journal of Web Semantics* **3**(2–3) (2005), 158–182. <http://dx.doi.org/10.1016/j.websem.2005.06.005>; <http://www.bibsonomy.org/bibtex/2924e60509d7e1b45c6f38eaf9a5c6bb/gromgull>.
- [37] C. Bizer and A. Schultz, The Berlin SPARQL benchmark, *International Journal On Semantic Web and Information Systems* **5**(2) (2009), 1–24.
- [38] L. Ma, Y. Yang, Z. Qiu, G. Xie, Y. Pan and S. Liu, Towards a Complete OWL Ontology Benchmark, in: *ESWC'06 Proceedings of the 3rd European conference on The Semantic Web: research and applications*, 2006, pp. 125–139.
- [39] M. Ledvinka and P. Křemen, Object-UOBM: An Ontological Benchmark for Object-oriented Access, in: *Knowledge Engineering and the Semantic Web*, Springer, CCIS series, 2015.
- [40] W.V. Siricharoen, Ontologies and Object models in Object Oriented Software Engineering, in: *Proceedings of the International MultiConference of Engineers and Computer Scientists 2006, IMECS '06, June 20-22, 2006, Hong Kong, China*, 2006, pp. 856–861.
- [41] G. Weikum and G. Vossen, *Transactional Information Systems: Theory, Algorithms, and the Practice of Concurrency Control and Recovery*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002. ISBN 9780080519562.
- [42] L. Feigenbaum, G.T. Williams, K.G. Clark and E. Torres, SPARQL 1.1 Protocol, Technical Report, W3C, 2013, Accessed 2018-01-08. <https://www.w3.org/TR/sparql11-protocol>.
- [43] S. Speicher, J. Arwe and A. Malhotra, Linked Data Platform 1.0, Technical Report, W3C, 2013, Accessed 2018-01-08. <http://www.w3.org/TR/ldp/>.
- [44] G. Hillairet, F. Bertrand and J.Y. Lafaye, Rewriting Queries by Means of Model Transformations from SPARQL to OQL and Vice-Versa, in: *Theory and Practice of Model Transformations*, R.F. Paige, ed., Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 116–131. ISBN 978-3-642-02408-5.
- [45] R.G.G. Cattell and D.K. Barry, *The Object Data Standard: ODMG 3.0*, Morgan Kaufmann, 2000. ISBN 1-55860-647-5.
- [46] C. Stadler and J. Lehmann, JPA Criteria Queries over RDF Data, in: *Workshop on Querying the Web of Data co-located with the Extended Semantic Web Conference*, 2017. http://jens-lehmann.org/files/2017/quweda_jpa.pdf.
- [47] M. Zimmermann, Owl2Java – A Java Code Generator for OWL, 2009, Accessed 2018-01-02. <http://www.incunabulum.de/projects/it/owl2java/>.
- [48] M. Quasthoff, H. Sack and C. Meinel, Can Software Developers Use Linked Data Vocabulary?, in: *Proceedings of International Conference on Semantic Systems 2009 (i-semantics 2009)*, 2009.
- [49] D. Allemang and J. Hendler, *Semantic Web for the Working Ontologist: Effective Modeling in RDFS and OWL*, 2nd edn, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2011. ISBN 9780123859655, 9780123859662.
- [50] Holger Knublauch, SPIN - SPARQL Inferencing Notation, 2017, <http://spinrdf.org/>, accessed 2018-11-17. <http://spinrdf.org/>.
- [51] J.J. Carroll, C. Bizer, P. Hayes and P. Stickler, Named Graphs, Provenance and Trust, in: *Proceedings of the 14th International Conference on World Wide Web, WWW '05*, ACM, New York, NY, USA, 2005, pp. 613–622. ISBN 1-59593-046-9. doi:10.1145/1060745.1060835. <http://doi.acm.org/10.1145/1060745.1060835>.
- [52] A. Zimmermann, RDF 1.1: On Semantics of RDF Datasets, 2014, Accessed 2018-01-04. <https://www.w3.org/TR/rdf11-datasets/>.
- [53] T. Lebo, S. Sahoo and D. McGuinness, PROV-O: The PROV Ontology, W3C Recommendation, W3C, 2013, <http://www.w3.org/TR/prov-o/>, accessed 2018-11-17..
- [54] S.S. Sahoo, O. Bodenreider, P. Hitzler, A. Sheth and K. Thirunarayan, Provenance Context Entity (PaCE): Scalable Provenance Tracking for Scientific RDF Data, in: *Proceedings of the 22nd International Conference on Scientific and Statistical Database Management, SSDBM'10*, Springer-Verlag, Berlin, Heidelberg, 2010, pp. 461–470. ISBN 3-642-13817-9, 978-3-642-13817-1. <http://dl.acm.org/citation.cfm?id=1876037.1876075>.
- [55] M. Quasthoff and C. Meinel, Tracing the Provenance of Object-Oriented Computations on RDF Data, in: *Proceedings of the 2nd Workshop (in conjunction with 7th ESWC) on Trust and Privacy on the Social and Semantic Web (SPOT 2010)*, 2010.
- [56] P. Křemen, B. Kostov, M. Blaško, J. Ahmad, V. Plos, A. Lališ, S. Stojić and P. Vittek, Ontological Foundations of European Coordination Centre for Accident and Incident Reporting Systems, *Journal of Aerospace Information Systems* **14**(5) (2017), 279–292, <https://doi.org/10.2514/1.1010441>.
- [57] L. Saeeda, Iterative Approach for Information Extraction and Ontology Learning from Textual Aviation Safety Reports, in: *The Semantic Web*, E. Blomqvist, D. Maynard, A. Gangemi, R. Hoekstra, P. Hitzler and O. Hartig, eds, Springer International Publishing, Cham, 2017, pp. 236–245. ISBN 978-3-319-58451-5.
- [58] S. Oaks, *Java Performance: The Definitive Guide*, 1st edn, O'Reilly Media, Inc., 2014. ISBN 1449358454, 9781449358457.
- [59] B. Kostov, J. Ahmad and P. Křemen, Towards Ontology-Based Safety Information Management in the Aviation Industry, in: *On the Move to Meaningful Internet Systems: OTM 2016 Workshops: Confederated International Workshops: EI2N, FBM, ICSP, Meta4eS, and OTMA 2016, Rhodes, Greece, October 24–28, 2016, Revised Selected Papers*, I. Ciuciu, C. Debruyne, H. Panetto, G. Weichhart, P. Bollen, A. Fensel and M.-E. Vidal, eds, Springer International Publishing, Cham, 2017, pp. 242–251. ISBN 978-3-319-55961-2. doi:10.1007/978-3-319-55961-2_25. https://doi.org/10.1007/978-3-319-55961-2_25.
- [60] G. Guizzardi, Ontological Foundations for Structural Conceptual Models, PhD thesis, University of Twente, 2005.

- [61] E. Oren, B. Heitmann and S. Decker, ActiveRDF: Embedding SemanticWeb data into object-oriented languages, *Web Semantics: Science, Services and Agents on the World Wide Web* 6(3) (2008), ISSN 1570-8268. <http://www.websemanticsjournal.org/index.php/ps/article/view/143>.
- [62] J. Frohn, G. Lausen and H. Uphoff, Access to objects by path expressions and rules, in: *Proceedings of the International Conference on Very Large Data Bases (VLDB)*, 1994, pp. 273–284.
- [63] G. Stevenson and S. Dobson, Sapphire: Generating Java Runtime Artefacts from OWL Ontologies, in: *Advanced Information Systems Engineering Workshops*, C. Salinesi and O. Pastor, eds, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 425–436. ISBN 978-3-642-22056-2.
- [64] P. Mika, Social Networks and the Semantic Web: The Next Challenge, *IEEE Intelligent Systems* 20(1) (2005). <http://www.cs.vu.nl/~pmika/research/papers/IEEE-TrendsAndControversies.pdf>.
- [65] F. Chevalier, AutoRDF - Using OWL as an Object Graph Mapping (OGM) Specification Language, in: *The Semantic Web*, H. Sack, G. Rizzo, N. Steinmetz, D. Mladenić, S. Auer and C. Lange, eds, Springer International Publishing, Cham, 2016, pp. 151–155. ISBN 978-3-319-47602-5.
- [66] J. von Malottki, JAOb (Java Architecture for OWL Binding), 2008, Accessed 2018-01-02. <http://wiki.yoshtec.com/jaob>.
- [67] J. Tao, E. Sirin, J. Bao and D.L. McGuinness, Integrity Constraints in OWL, in: *AAAI*, M. Fox and D. Poole, eds, AAAI Press, 2010.
- [68] K. Wenzel, KOMMA: An Application Framework for Ontology-based Software Systems, *Semantic Web – Interoperability, Usability, Applicability* (2010).
- [69] A. Alishevskikh, RDFBeans, 2017, Accessed 2018-01-02. <https://rdfbeans.github.io>.
- [70] M. Völkel and Y. Sure, RDFReactor - From Ontologies to Programmatic Data Access, in: *Poster and Demo at International Semantic Web Conference (ISWC) 2005, Galway, Ireland*, 2005.
- [71] P. Ježek and R. Mouček, Semantic framework for mapping object-oriented model to semantic web languages, *Frontiers in Neuroinformatics* 9(3) (2015).
- [72] T. Cowan and D. Donohue, Jenabeau, 2010, Accessed 2018-01-04. <https://code.google.com/archive/p/jenabeau/>.
- [73] Ben Lavender, Spira: A Linked Data ORM for Ruby, 2010, Accessed 2018-03-20. <https://github.com/ruby-rdf/spira>.
- [74] Cosmin Basca, SuRF, 2009, Accessed 2018-03-20. <https://github.com/cosminbasca/surfrdf>.
- [75] B.H. Szekely and J. Betz, Jastor Homepage, 2006, Accessed 2018-01-02. <http://jastor.sourceforge.net>.
- [76] H. Story, Semantic Object (Metadata) Mapper, 2009, Accessed 2018-01-02. <https://github.com/bblfish/sommer>.
- [77] D. Toti and M. Rinelli, RAN-Map: a system for automatically producing API layers from RDF schemas, *Journal of Ambient Intelligence and Humanized Computing* 8(2) (2017), 291–299, ISSN 1868-5145. doi:10.1007/s12652-016-0394-z. <https://doi.org/10.1007/s12652-016-0394-z>.
- [78] S. Fernández, D. Berrueta, M.G. Rodríguez and J.E. Labra, TRIOO: Keeping the semantics of data safe and sound into object-oriented software, in: *ICSOF 2010 - Proceedings of the 5th International Conference on Software and Data Technologies*, Vol. 1, 2010, pp. 311–320.
- [79] P.-H. Chiu, C.-C. Lo and K.-M. Chao, Integrating Semantic Web and Object-Oriented Programming for Cooperative Design, *Journal of Universal Computer Science* 15(9) (2009), 1970–1990.
- [80] N. Drummond, A. Rector, R. Stevens, G. Moulton, M. Horridge, H.H. Wang and J. Seidenberg, Putting OWL in Order: Patterns for Sequences in OWL, in: *OWL Experiences and Directions (OWLEd 2006)*, Athens Georgia, 2006.
- [81] Oracle, Java Garbage Collection Basics, 2016, Accessed 2018-01-04. <http://www.oracle.com/webfolder/technetwork/tutorials/obe/java/gc01/index.html>.
- [82] I.R. Forman and N. Forman, *Java Reflection in Action (In Action Series)*, Manning Publications Co., Greenwich, CT, USA, 2004. ISBN 1932394184.
- [83] M. Schmidt, M. Meier and G. Lausen, Foundations of SPARQL Query Optimization, in: *Proceedings of the 13th International Conference on Database Theory, ICDT '10*, ACM, New York, NY, USA, 2010, pp. 4–33. ISBN 978-1-60558-947-3. doi:10.1145/1804669.1804675.