

Press Release January 2018

HOBBIT

Holistic Benchmarking of
Big Linked Data

H2020 Research and Innovation
Grant Agreement No. 688227

Website

<http://project-hobbit.eu>

Contact

Dr. Axel-Cyrille Ngonga Ngomo
Institute for Applied Informatics
Hainstraße 11, 04109 Leipzig
Germany

Phone: +49 341 97 32362
ngonga@informatik.uni-leipzig.de



HOBBIT
Holistic Benchmarking
of Big Linked Data

HOBBIT celebrates its second year!

We are very happy to announce that HOBBIT is now two years old! The project consortium members were very active in this year and worked extensively towards the project's major results: *benchmark development, deployment of the HOBBIT platform and challenge organisation!*

More importantly, HOBBIT delivered the first version of its *Generation & Acquisition, Analysis & Processing, Storage & Curation and Visualization & Services* benchmarks. More specifically:

- **Generation & Acquisition benchmarks** measure the performance of *SPARQL query processing systems* when faced with streams of data from industrial machinery in terms of efficiency and completeness. In order to reflect the real loads on triple stores that are used in real applications, *Public Transport, Twitter, Traffic and Sensor data* from the plastic injection moulding industrial plants of Weidmüller were collected. In this context, HOBBIT provides benchmarks to measure the performance of extraction systems for unstructured streams of natural-language data. We took under account for the benchmarks, the real, curated unstructured datasets by experts and unstructured data streams by Bengal, a generic data generator.
- **Analysis & Processing benchmarks** focus on testing the performance of *link discovery systems and machine learning methods (supervised and unsupervised) for data analytics*. More specifically, the link discovery benchmarks developed in HOBBIT can be used to test the performance of (a) *instance matching tools* that implement string-based approaches for identifying matching entities and (b) systems that deal with *topological relations* proposed in the state of the art *DE-9IM* model (Dimensionally Extended nine-Intersection Model). The analysis benchmark is used to test the efficiency and effectiveness of machine learning (supervised and unsupervised) approaches on structured data.
- **Storage & Curation benchmarks** aim at testing the performance of *data storage and versioning systems* for Linked Data. The data storage benchmark focuses on the typical challenges faced by Linked Data storage systems and is based on the Social Network Benchmark (SNB) developed in the context of the EU FP7 LDBC Project. HOBBIT's versioning benchmark aims to test the ability of versioning systems to efficiently manage evolving Linked Data datasets and queries evaluated across multiple versions of such datasets; it extends the LDBC Semantic Publishing Benchmark (SPB), inspired from the publishing domain.



- **Visualization & Services benchmarks** aim at testing the performance of *query answering* and *faceted browsing systems* for Linked Data without involving users. The developed benchmarks are not intended to test user interfaces but focus on providing performance and accuracy measurements for approaches used in such interfaces. For the benchmarks, browsing scenarios that reflect an authentic use-case and challenge participating systems on different points of difficulty were developed.

All aforementioned benchmarks are available on HOBBIT's CKAN: <https://ckan.project-hobbit.eu/dataset>, along with their source code and related publications.

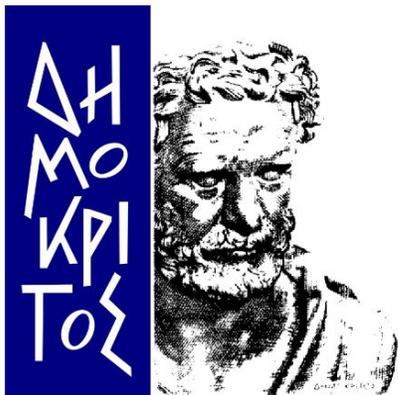
The HOBBIT *evaluation platform* was released in the beginning of the second year of the project. It is a distributed FAIR benchmarking platform for the Linked Data lifecycle that is open source and it can be downloaded and executed locally. The HOBBIT Platform can be accessed through its online instance that can be used for a) running public challenges and b) making sure that even people without the required infrastructure are able to run the benchmarks they are interested in.

The online instance of the HOBBIT benchmarking platform is accessible at master.project-hobbit.eu and its code is accessible at <https://github.com/hobbit-project>.

The developed benchmarks and platform were extensively used in the Challenges organized by HOBBIT project: the Mighty Storage (MOCHA), Query Answering over Linked Data (QALD), Open Knowledge Extraction (OKE) and DEBS Grand Challenge. The MOCHA, QALD and OKE challenges were organized in the context of ESWC 2017, whereas the DEBS Grand Challenge was held in conjunction with the DEBS 2017 conference.

In addition to the aforementioned ones, HOBBIT organized the QALD-8 Challenge in conjunction with ISWC 2017 that was incorporated in the Natural Language Interfaces for Web of Data (NLIWoD) workshop. HOBBIT also proposed and co-organized a new track at the Ontology Matching (OM) 2017 workshop, that has been running under the auspices of OAEI; the workshop was held in conjunction with ISWC 2017. After this successful endeavour, the OM organizers and the HOBBIT consortium members decided to work towards replacing the SEALS platform, used for a number of years for running the OM benchmarks, with the HOBBIT platform. The OAEI 2017.5 campaign that will take place in conjunction with ESWC 2018, aims at ontology matchers to benchmark their systems using exclusively the HOBBIT platform. More information on the campaign can be found here: <http://oaei.ontologymatching.org/2017.5>

The aim of the Mighty Storage Challenge (MOCHA) was to test the performance of solutions for SPARQL processing in aspects that are relevant for modern applications. These include ingesting data, answering queries on large datasets and serving as backend for applications driven by Linked Data. Three system participated in the MOCHA tasks: (a) Virtuoso Open-Source Edition 7.2, developed by OpenLink Software, that served as the baseline system for all MOCHA 2017 tasks (MOCHA Baseline), (b) QUAD, developed



Learn about our Challenges



by Ontos, and (c) Virtuoso Commercial Edition 8.0 (beta), developed by OpenLink Software.

The Question Answering over Linked Data (QALD) challenge aimed at providing an up-to-date benchmark for assessing and comparing state-of-the-art systems that mediate between a user, expressing his or her information need in natural language, and RDF data. The Challenge was organized in four different tasks: multilingual question answering over DBpedia, hybrid question answering, large-scale question answering over RDF and finally, question answering over Wikidata. Three systems participated in the task: WDAqua, AMAL and ganswer.

and our Open Challenges



The Open Knowledge Extraction (OKE) Challenge was organized in four different tasks (a) Focused Named Entity Identification and Linking, (b) Broader Named Entity Identification and Linking (c) Focused Musical Named Entity Recognition and Linking and finally (d) Knowledge Extraction. The goal of the OKE Challenge is to test the performance of Knowledge Extraction Systems with respect to the Semantic Web. Adel and Fox systems participated in the evaluation of the OKE Challenge.

Finally, the focus of the DEBS 2017 Grand Challenge was on the analysis of the RDF streaming data generated by digital and analogue sensors embedded within manufacturing equipment. The goal of the challenge was to implement detection of anomalies in the behaviour of such manufacturing equipment. The challenge was co-organized with AGT International on behalf of the HOBBIT Project.

Join HOBBIT Community



Last but not least, HOBBIT launched a set of *Open Challenges*, namely the OKE, MOCHA, SQA (Scalable Question Answering) and STREAML (Stream Machine Learning) Challenges. OKE and MOCHA have been based on the challenges that were launched in the context of ESWC 2017 and are currently running. The main task of SQA Open Challenge is to challenge the systems in translating a user's information request into such a form that it can be efficiently evaluated using standard Semantic Web query processing and inferencing techniques. The StreamML Open Challenge focuses on the task related to the problem of automatic detection of anomalies for manufacturing equipment.

In addition to these challenges that mostly focus on pushing systems to their limits, HOBBIT organizes an open call for benchmarks that will be integrated in the platform so that a large number of people can benchmark their systems on a standardized hardware producing comparable results.

Information on the HOBBIT Project, the platform, benchmarks and challenges can be found here: <https://project-hobbit.eu>