A Prototype Service for Benchmarking Power Consumption of Mobile Semantic Applications

Evan W. Patton and Deborah L. McGuinness

Rensselaer Polytechnic Institute
110 8th Street, Troy NY 12180 USA
{pattoe, dlm}@cs.rpi.edu
http://tw.rpi.edu/

Abstract. We present a prototype web service that enables researchers to evaluate the performance per watt of semantic web tools. The web service provides access to a hardware platform for collecting power consumption data for a mobile device. Experiments are specified using RDF to define the conditions of the experiment, the operations that compose those conditions, and how they are combined into individual execution plans. Further, experimental descriptions and their provenance are published as linked data, allowing others to easily repeat experiments. We will demonstrate how we have used the system to date, how others can use it, and discuss its potential to revolutionize design and development of semantically enabled mobile applications.

Keywords: reasoning, mobile, power, performance, resource-constrained

1 Introduction

One challenge that semantic technologies face when deployed on mobile platforms like smartphones is the amount of energy available for the device to compute and communicate with other agents. For example, the Google Nexus One, one of the first Android smartphones, had a single core processor operating at 1 GHz and 512 MB of RAM. Samsung's latest offering, the Galaxy S5, has a quad core, 2.5 GHz processor and 2 GB of RAM, more than a 8-fold increase in processing power and 4-fold increase in memory in 5 years. However, the battery capacity of the two phones are 1400 mAh and 2800 mAh, resp., which indicates that battery technology is progressing more slowly than processing technology. Further, application complexity has also increased. Tools are needed to help developers understand how semantic tools consume power so as to identify when they can use local reasoning on mobile platforms or when off-device computation is more practical.

We introduced a broadly reusable methodology [3] motivated by these concerns to evaluate the performance of reasoners relative to the amount of energy consumed during operation. Ultimately, these metrics will provide developers deeper insight into power consumption and enable next-generation applications of semantic technologies for power constrained devices. We present a prototype

Table 1. A data sample for query 14 (Listing 1.1) from LUBM executed on the Samsung Galaxy S4. Times are in milliseconds, memory is in kilobytes, and power is in milliwatts.

Reasoner	Init.	Ont. Load	Data Load	Query Plan	Answer	Memory	Power
Jena	0.122	372.6	7076	2.594	233.2	35023	944
Pellet	0.152	355.7	8872	1.984	12350	59418	1024
HermiT	0.427	407.8	17442	0.092	21205	58720	995

ontology-driven web service for researchers to use our reference hardware setup to perform analysis of semantic web tools' power consumption.

2 Web Service for Power-Performance Evaluation

Our power benchmarking methodology [3] bypasses the removable battery in a Samsung Galaxy S4 to collect power data during reasoning and query answering tasks using three reasoning engines. Because our methodology requires a hardware setup, we are developing and will demonstrate a web service to execute experiments using our existing infrastructure. The web service is based on the Semantic Automated Discovery and Integration (SADI) Framework[1] and accepts jobs described using RDF and the ontology we will discuss in Section 3. On completion, it provides a ZIP file containing runtime information, raw and processed power measurements, power and energy consumption statistics, and provenance capturing information about the process. Table 1 shows a sample data point for each of three different reasoners on the Lehigh University Benchmark [2], query 14 (shown in Listing 1.1).

Listing 1.1. Lehigh University Benchmark query 14

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX ub: <http://swat.cse.lehigh.edu/onto/univ-bench.owl#>
SELECT ?X WHERE {?X rdf:type ub:UndergraduateStudent}
```

3 Toward an Ontology for Experiment Descriptions

We will demonstrate our experimental ontology for declaratively describing the operational constraints of an experiment, which is then executed on the target device. The experiment description is published as linked data, along with metadata about the experiment output, andprovenance modeled using the PROV-O ontology [1].These metadata are published in a triple store to enable meta-analysis, recombination, and extension of power experiments.

[1] http://sadiframework.org/content/

Experiment. Experiment provides the root of an experiment description. Listing 1.2 shows an example experiment. Core to an experiment description are conditions, which are grouped together based on some common dimension. If an experiment defines more than one condition group, the engine performing the experiment can generate specific conditions through the use of a *conditionFillStrategy*. We are currently investigating two different strategies, *CrossJoin* and *Paired*, that evaluate the cross product and paired conditions across groups, respectively. To provide control over what data are returned, the author of the experiment can declare which variables are of interest.

Listing 1.2. An example description of an experiment

```
[] a exp:Experiment ;
   exp:name 'LUBM on Android' ;
   exp:version '1.0' ;
   dc:creator <http://www.evanpatton.com/evan.rdf#me> ;
   exp:trials 30 ;
   exp:conditions :ReasonerConditionGroup,
      OntologyConditionGroup ;
   exp:conditionFillStrategy exp:CrossJoin ;
   exp:dependentVariable exp:ExecutionInterval,
      exp:AveragePowerConsumption, exp:MaxPowerConsumption .
```

*Conditions and Condition Groups. Condition*s are the highest unit of test in our experiment ontology. They are composed of collections of operations that specify a sequence of actions to take on the device. Listing 1.3 shows an example of an ontology condition group that specifies two different ontology operations. Currently, we only support nominal values, but future versions of the ontology will also support ordinal, scalar, and ratio level inputs.

Listing 1.3. An example of a condition group with conditions

```
:OntologyConditionGroup a exp:ConditionGroup ;
  exp:name 'Ontology Condition' ;
  exp:varies exp:OntologyDatasetQueryOperation ;
  exp:nominalValues ( :SchemaOrgOperations :LUBMOperations )
```

*Operations. Operation*s encapsulate actions to be performed on the experimental device. In Listing 1.4, we show an example of how an operation would define tests for the LUBM query set. The *measurePowerDuringDownload* property can be used to evaluate the performance of communication channels while retrieving the content required for performing the experiment. In addition to loading ontologies, datasets, and executing queries, our ontology supports modeling reasoners, parallel and sequential operations, and randomization of operations.[2]

[2] Due to space constraints, we cannot elaborate on the details of modeling each operation type. For more information and examples, please see http://power.tw.rpi.edu

Listing 1.4. An example of an combined operation on an ontology, dataset, and queries

```
:LUMBOperations a exp:OntologyDatasetQueryOperation ;
   exp:name 'LUBM' ;
   exp:measurePowerDuringDownload false ;
   exp:ontology lubm:univ-bench.owl ;
   exp:dataset lubm:lubm-100k.ttl ;
   exp:query lubm:query1.rq , lubm:query2.rq , ... .
```

4 Discussion and Conclusions

As semantic technologies become more prevalent, we need to ensure that tools are available to assist in their deployment on a variety of devices including mobile platforms, which are often power constrained. While our initial work provides direction for this effort, we recognize that more widespread adoption requires lower barriers to entry. We described a web service under active development to provide access to a reference implementation of the hardware described in [3] where we found that, while compute time accounts for most energy consumption, significant memory consumption may affect power consumption during reasoning. With this investigation, we are working to enable semantic web researchers and implementers to obtain insight into the power requirements for semantic technology stacks. We will demonstrate a variety of example experiments and discuss broader usage with attendees. In future work we intend to use expand this web service to provide a means of easily repeating experiments as well as further support for modeling the execution of experiments using We also intend to provide example code that present an analysis of more reasoners, e.g. by utilizing the work in [4].

Acknowledgements

Mr. Patton was funded by an NSF Graduate Research Fellowship. RPIs Tetherless World Constellation is supported in part by Fujitsu, Lockheed Martin, LGS, Microsoft Research, Qualcomm, in addition to sponsored research from DARPA, IARPA, NASA, NIST, NSF, and USGS.

References

1. Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., Soiland-Reyes, S., Zednik, S., Zhao, J.: PROV-O: The PROV ontology. Tech. rep., W3C (2013), http://www.w3.org/TR/prov-o/
2. Guo, Y., Pan, Z., Heflin, J.: LUBM: A benchmarking for OWL knowledge base systems. Web Semantics 3(2), 158–182 (2005)
3. Patton, E.W., McGuinness, D.L.: A power consumption benchmark for reasoners on mobile devices. In: Proceedings of the 13th International Semantic Web Conference: Replication, Benchmark, Data & Software Track (2014)
4. Yus, R., Bobed, C., Esteban, G., Bobillo, F., Mena, E.: Android goes semantic: DL reasoners on smartphones. In: OWL Reasoner Evaluation Workshop 2013 (2013)