

# SPARQLES: Monitoring Public SPARQL Endpoints

**Editor(s):** Jens Lehmann, Universität Bonn, Fraunhofer IAIS & Universität Leipzig, Germany

**Solicited review(s):** Ivan Ermilov, Universität Leipzig, Germany; Two anonymous reviewers

Pierre-Yves Vandenbussche<sup>a</sup>, Jürgen Umbrich<sup>b</sup>, Luca Matteis<sup>c</sup>, Aidan Hogan<sup>d,\*</sup>, Carlos Buil-Aranda<sup>e</sup>

<sup>a</sup>*Fujitsu (Ireland) Limited, Swords, Co. Dublin, Ireland — E-mail: pierre-yves.vandenbussche@ie.fujitsu.com*

<sup>b</sup>*Vienna University of Economy and Business (WU), Austria — E-mail: juergen.umbrich@wu.ac.at*

<sup>c</sup>*Department of Computer Science, Sapienza University of Rome, Italy — E-mail: matteis@di.uniroma1.it*

<sup>d</sup>*Center for Semantic Web Research, Department of Computer Science, University of Chile, Chile — E-mail: ahogan@dcc.uchile.cl*

<sup>e</sup>*Center for Semantic Web Research, Departamento de Informática, Universidad Técnica Federico Santa María, Chile — E-mail: cbuil@inf.utfsm.cl*

**Abstract.** We describe SPARQLES: an online system that monitors the health of public SPARQL endpoints on the Web by probing them with custom-designed queries at regular intervals. We present the architecture of SPARQLES and the variety of analytics that it runs over public SPARQL endpoints, categorised by availability, discoverability, performance and interoperability. We also detail the interfaces that the system provides for human and software agents to learn more about the recent history and current state of an individual SPARQL endpoint or about overall trends concerning the maturity of all endpoints monitored by the system. We likewise present some details of the performance of the system and the impact it has had thus far.

**Keywords:** SPARQL endpoints, Linked Data, Semantic Web, Web of Data

## 1. Introduction

Thousands of Linked Datasets have been made publicly available in recent years.<sup>1</sup> These datasets span a plethora of topics, varying from general-interest datasets like DBpedia [20]<sup>2</sup> or GeoNames<sup>3</sup>, to more niche datasets on topics like proteins<sup>4</sup> or Pokémon<sup>5</sup>. Each dataset follows the Semantic Web standards [31,19] for describing its content, and the Linked

Data principles [8] for making that content accessible on the Web. The goal is to enable clients to access these diverse datasets in an automated and uniform way, and also to combine content from multiple locations in a similarly automated fashion.

To entice new consumers, many publishers began hosting public SPARQL endpoints over their datasets such that clients can pose complex queries to the server as a single request and retrieve direct answers. Hundreds of public SPARQL endpoints have thus emerged on the Web in recent years [11]. These endpoints index content with a variety of topics and sizes and (in theory at least) accept arbitrary SPARQL queries from remote clients over the Web. However, applications using these endpoints have been slow to emerge. The convenience of SPARQL queries for clients translates into significant server-side costs maintaining such heavyweight query services, which translate into a va-

---

\* Corresponding author.

<sup>1</sup>At the time of writing, LODstats [14] reports 9,960 datasets: <http://stats.lod2.eu/> (all URLs in this paper were last accessed on 2016/05/24). The most recent “State of the LOD Cloud” report found 1,014 datasets.

<sup>2</sup><http://datahub.io/dataset/dbpedia>

<sup>3</sup><http://datahub.io/dataset/geonames-semantic-web>

<sup>4</sup><http://datahub.io/dataset/uniprot-databases>

<sup>5</sup><http://datahub.io/dataset/pokepedia-fr>

riety of technical problems on the level of the SPARQL infrastructure itself [11].

With respect to what that SPARQL infrastructure consists of, the recent SPARQL 1.1 standard issued recommendations relating to the following:

**Query** The SPARQL 1.1 Query Language recommendation [18] extends the original SPARQL Query Language [26] with features such as property paths, sub-queries, aggregates, etc. The related SPARQL 1.1 Federated Query recommendation [25] specifies how a SPARQL engine can invoke a remote endpoint at runtime.

**Protocol** The SPARQL 1.1 Protocol recommendation [15] specifies how clients should interact with a SPARQL endpoint over HTTP, including how GET/POST requests should be structured, what sorts of responses should be returned, etc. Three output result formats have also been recommended, extending the XML format introduced in the original standard with options for returning data in JSON or CSV/TSV.

**Description** The SPARQL 1.1 Service Description recommendation [35] provides a vocabulary with which the capabilities and configuration of a SPARQL endpoint can be described in RDF such that, for example, clients can discover endpoints with the features they need.

**Update** The SPARQL 1.1 Update specification [16] describes a language for inserting, deleting and updating the data present in a SPARQL engine.

**Entailment** The SPARQL 1.1 Entailment Regimes recommendation [17] describes how ontological entailments can be included when computing the answers for a SPARQL query.

With respect to public SPARQL endpoints, the latter two aspects of the SPARQL infrastructure are *currently* of lesser interest: updates to data are unlikely to be enabled on a public query service and we do not yet know of any public SPARQL endpoint supporting entailment regimes. Thus, for clients of current public endpoints, the former three aspects – query, protocol, and description – appear to be of most relevance.

With respect to these three infrastructural aspects, in previous work [11] we performed an empirical investigation of the maturity of public SPARQL endpoints from the perspective of the client, who we argue needs endpoints that are: (i) highly-available through the SPARQL protocol, thus allowing queries to be an-

swered reliably at any time; (ii) described using standard vocabularies in well-known locations, thus allowing for the (automatic) discovery of relevant endpoints over the Web; (iii) capable of answering queries in acceptable time, thus enabling their use in real-time applications; (iv) compliant with respect to supporting the query features of SPARQL (1.1), thus enabling them to be interrogated alongside other endpoints in a uniform manner. If an endpoint has high availability, is well-described, supports all SPARQL 1.1 features, and returns query answers quickly, we consider it to meet all of the basic infrastructural requirements a client would have. However, these goals fall on a continuum rather than being binary or discrete: for example, endpoints may support a majority of features of SPARQL 1.1, or may only be able to answer certain queries within a given expected response time. Thus the question is to *what extent* endpoints are mature. We thus defined four general dimensions for assessing the maturity of public SPARQL endpoints, as follows.

First, we looked at *AVAILABILITY: the ratio of time for which a given endpoint is responsive through the SPARQL protocol, or alternatively, the probability of a SPARQL endpoint being able to successfully respond to a (simple valid) SPARQL query at a given point in time*. Endpoints are often provided on a not-for-profit basis, where the resources available to host and maintain them may be limited and thus services may go offline temporarily or even permanently without warning. Likewise, executing SPARQL queries can be expensive for a server, which may reach its capacity and be unable to respond to further requests. An application relying on a given endpoint would inherit these underlying availability issues; the situation can be even worse if an application relies on multiple endpoints.

Second, we looked at *DISCOVERABILITY: the degree to which an endpoint provides descriptions of its content, configuration and functionality in well-known locations using well-known vocabularies, such that clients can (automatically) discover that endpoint based on criteria such as the classes and properties its content pertains to, the amount of data it contains, the query features it supports, etc*. Without these metadata, clients may struggle to find endpoints with content and features relevant for their needs.

Third, we looked at *PERFORMANCE: the amount of time taken for an endpoint to answer a query over HTTP using standard SPARQL protocol methods*. Evaluating a SPARQL 1.0 query is PSPACE-complete [24]; the analogous complexity for SPARQL 1.1 evaluation is at least as hard. Of course, these types

of worst-case queries are likely to be quite rare [4], but even “PTIME queries” can require huge amounts of processing to satisfy over even moderately-sized datasets, which may make services slow and potentially unsuitable for use in real-time applications.

Finally, we looked at INTEROPERABILITY: *how compliant the endpoint is with respect to the features of the SPARQL 1.1 query language*. In particular, an endpoint that does not support some query features of SPARQL 1.1 may not be interoperable with other endpoints or applications that expect these features to be supported. This diversity in supported features means that a client may not have a uniform query interface common to all endpoints against which they can program the logic of their application(s).

Given mixed results in our initial experiments [11], we foresaw the need for an online system to track such aspects of public endpoints over time, and to help clients assess for themselves the maturity of individual endpoints based on empirical data. Along these lines, we initiated work on the SPARQL Endpoint Status (SPARQLES) system, which is currently available at <http://sparql.es.org/>.<sup>6</sup> The SPARQLES system has been online since October 2013 (two years at the time of writing), during which time we have made various refinements based on community feedback, and have made the system more reliable and less expensive for the public endpoints we monitor.

The SPARQLES system currently focuses on the same dimensions we investigated in previous work [11]: AVAILABILITY, DISCOVERABILITY, PERFORMANCE and INTEROPERABILITY. While these dimensions were selected by us in a somewhat ad hoc manner based on our experience of working with endpoints – rather than say the result of a systematic study – we argue that they provide a comprehensive insight into the maturity of a given public SPARQL endpoint, touching upon the query, protocol and descriptive aspects of the SPARQL infrastructure. While we cannot claim that SPARQLES captures a complete set of metrics, we are open to extending the system in future to consider other measures, possibly with the help of the community through our open source code-base, or perhaps as the result of further research on the topic of evaluating such public endpoint services. That is to say, we see much potential for further work on this topic.

<sup>6</sup>SPARQLES is also a predecessor of an older system that tracked only availability [32]: <http://labs.mondeca.com/sparqlEndpointsStatus.html>.

This paper then extends upon previous works [11, 32] and describes the current SPARQLES system itself in detail: how it is constructed, what sorts of tests it performs, what queries it issues, what data it collects, what kinds of conclusions can be drawn, what interfaces and visualisations are provided, etc. In Section 2, we first discuss works relating to studies of public endpoints and monitoring Web services. In Section 3, we introduce the high-level SPARQLES architecture. In Section 4 we describe in more detail the analytics that SPARQLES runs over public endpoints and in Section 5 we describe the interfaces that we provide for agents to interact with the data collected. In Section 6 we present evaluation of the system including runtimes of analytics, growth in storage overheads, and A.P.I. performance. We later discuss the impact, limitations, and sustainability of SPARQLES in Section 7 before concluding with Section 8.

The SPARQLES system – both code and data – is published under a Creative Commons 4.0 license (CC BY), with code available from <https://github.com/pyvandebussche/sparql.es> and data available through interfaces described in Section 5.

## 2. Related Work

A number of works and systems have dealt with issues relating to public SPARQL endpoints. The DataHub<sup>7</sup> catalogue lists hundreds of Linked Datasets, many of which link to a SPARQL endpoint; it is the resulting list of endpoints that we monitor. Lorey [21] proposed a number of metrics for determining the performance of SPARQL endpoints, such as latency, throughput, random access time and join execution; experiments were performed in controlled settings rather than over public endpoints on the Web. Paulheim & Hertling [23] tried to find relevant SPARQL endpoints for a random sample of ten thousand IRIs; using VoID descriptions and the DataHub catalogue, they were successful in about 15% of cases. Mehdi et al. [22] proposed best-effort methods to find public endpoints relevant to a list of domain-specific keywords by querying endpoints for RDF literals generated from the terms. With respect to how public SPARQL endpoints are being used, a number of works have also performed analyses of the query logs of prominent endpoints [4,28], which were made avail-

<sup>7</sup><http://datahub.io/>

able by the USEWOD initiative [7] and later by LSQ [30]. While all of these works have helped to build a more detailed picture of the current state-of-the-art with respect to public SPARQL endpoints, none provide an online system like SPARQLES, nor does any one work look at the range of analytics we provide.

More generally, since public SPARQL endpoints can be considered as Web Services, our work also relates to the topic of monitoring Web Services, and in particular, the notion of Quality of Service (QoS). One of the seminal works in this area was by Ran [27], who proposed an influential list of twenty-three QoS dimensions for Web services in four categories: runtime (R), transaction support (T), configuration (C), and security (S). In this context, our work touches upon the following dimensions:

- Performance (R)** The response time, latency and throughput of the service.
- Availability (R)** The probability of a system being operational at any given point in time.
- Robustness/Flexibility (R)** The degree to which a service can cope with diverse inputs.
- Exception Handling (R)** The gracefulness with which errors are handled and explained.
- Supported Standard (C)** The degree of compliance of the service with respect to some standards.
- Completeness (C)** The ratio of advertised features that are found to work in practice.

Some of the other QoS dimensions defined by Ran [27] are either not currently relevant for public SPARQL endpoints – such as authentication, authorization, etc. – or cannot be easily tested in our setting – for example, the accuracy of results. On the other hand, we also consider the discoverability of endpoints, which refers to how well the service describes itself, which was not explicitly mentioned by Ran [27].

### 3. SPARQLES Architecture

The SPARQLES system is designed to observe a set of public SPARQL endpoints over time. Currently SPARQLES is tracking all of the endpoints listed in the DataHub catalogue found using the DataHub APIs; we thus align the inclusion criteria of SPARQLES with that of the DataHub. SPARQLES performs a fixed set of analytics against each listed endpoint at fixed intervals, stores the historical results and allows these results to be accessed through online interfaces.

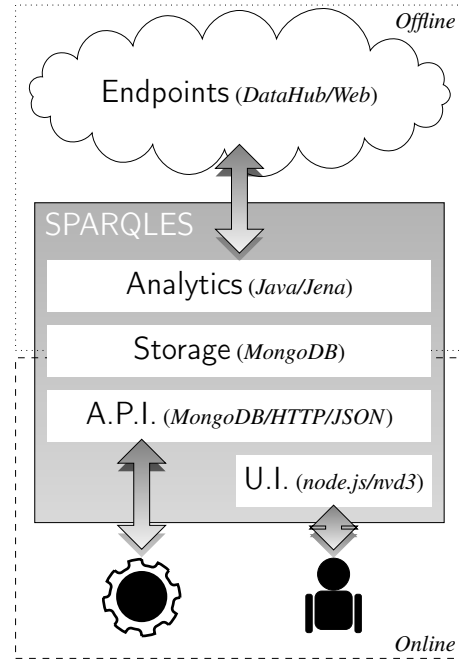


Fig. 1. High-level System Architecture

The high-level architecture for observing the selected endpoints is depicted in Figure 1, where we show the offline and online parts of the system. The offline parts are responsible for collecting information about the endpoints. The online parts are responsible for presenting the results to the clients of the system. The main components are as follows:

**Analytics (offline):** responsible for performing analysis over endpoints at regular intervals, thus producing the raw observational data. This component is implemented with custom Java code that uses Jena as a query client to interact with endpoints over the SPARQL protocol. Analytics are scheduled using cron jobs.

**Storage (both):** offers persistence over the results of the offline Analytics component and enables online querying and aggregation. For this, we use MongoDB, which stores both low-level data – such as the responses of endpoints to individual queries – as well as higher-level, aggregated data – such as the number of queries that succeeded in the past month.

**A.P.I. (online):** offers software agents a RESTful application programming interface through which to query key data about endpoints. Agents may access the A.P.I. through simple HTTP GET calls,

which return JSON-formatted data from the storage back-end.

U.I. (**online**): *offers human agents a user interface with a mix of aggregate visualisations and per-endpoint visualisations.* The search and user interactions provided by the U.I. are built on top of the data delivered by the A.P.I. The user interface is implemented using the Node.js<sup>8</sup> Javascript runtime engine and the nvd3<sup>9</sup> Javascript library for rendering interactive visualisations.

We describe these components in more detail in the following sections. In Section 4, we focus on the offline phase, and in particular, the types of analytics we run. Thereafter, in Section 5, we describe the storage and online parts of the system, including the types of interfaces that we provide for the public to interact with the collected data by the SPARQLES system.

## 4. Analytics

We now provide details of the offline phase, and in particular, the analytics performed by the system.

The following analytics are implemented with custom Java code that uses Apache Jena (2.12.2) to coordinate making requests to and collecting responses from SPARQL endpoints.<sup>10</sup> These analytics are scheduled to run at regular intervals using cron jobs. SPARQLES is hosted in the Vienna University of Economy and Business (WU), served by a 1 Gigabit network.

### 4.1. Availability

In our previous experiments we found that many endpoints have significant periods of downtime [11]. Some downtimes may be temporary, caused by network failures, sporadic high server loads, engine crashes, and so forth. Other downtimes appear permanent, indicating that an endpoint has probably been discontinued. Anticipating downtimes or distinguishing reliable endpoints from unreliable ones can be crucial for many clients. Hence SPARQLES closely monitors the historical availability of endpoints.

*Availability analytics* We define an endpoint as available if it can respond to a simple SPARQL query with some compliant response through the SPARQL protocol. For this reason, rather than perform a ping (which would only ensure that the host is accessible) or a generic HTTP lookup (which would only ensure that some web-page is available that could be, e.g., an under-maintenance notice), we wish to ensure that the endpoint can respond to basic queries. To avoid unnecessary load on remote servers, we send queries that should, in general, be simple to compute responses for.

Along these lines, to check availability, the system first issues a generic ASK query as follows:

```
ASK WHERE { ?s ?p ?o . }
```

Responding to this query should be trivial (is the index empty or not?). As soon as a valid response (positive or negative) is received, the system considers the request successful and concludes that the endpoint is available. However, some SPARQL endpoints cannot handle this ASK query. For such endpoints, we try a second query using the SELECT operation as follows:

```
SELECT ?s WHERE { ?s ?p ?o . } LIMIT 1
```

Again, this query should be cheap to compute: return any triple from the index (if any). We deem any endpoint responding to either query with any valid SPARQL response as available at that time.

*Frequency* We run availability tests once an hour, which allows us to monitor, e.g., the uptimes at different times of the day, including hours of peak Web-usage (performance will be discussed later). Availability results can then be aggregated per endpoint into a success rate for fixed time intervals, e.g., to compute availability over the past day, week, month, etc.

*Limitations* The local SPARQLES server may experience some downtimes or local network issues that may lead to remote endpoints being falsely reported as unavailable. In general however, when errors known to be local are omitted and when hourly results are aggregated into larger time intervals, such as weeks or months, such local effects should be smoothed out.

### 4.2. Discoverability

For a client, finding a SPARQL endpoint that contains content relevant for their needs [23,11] and the features that they require [11] can be challenging. The

<sup>8</sup><http://nodejs.org/>

<sup>9</sup><http://nvd3.org/>

<sup>10</sup><https://jena.apache.org/>

goal of the discoverability analytics is to determine the degree to which endpoints offer descriptions of themselves and their contents using (de facto) standards: to what extent an endpoint offers descriptions – in well-known locations using well-known vocabularies – of (i) its content and (ii) the features it supports. The SPARQLES system thus checks if a client can automatically find, for a given endpoint:

1. An SD description of its configuration [35].
2. A VoID description of its content [3].

The type of engine (Fuseki, Virtuoso, etc.) powering a SPARQL endpoint can also be an important information for a client; for example, some engines support non-standard keyword search functions that a client may be interested in. We thus also look for:

3. The type of engine powering the endpoint, sometimes mentioned in the HTTP header [11].

*SD Analytics* Endpoint capabilities – such as the version of SPARQL supported, query and update features, I/O formats, custom functions, and/or entailment regimes – can be described in RDF using the SPARQL 1.1 Service Description (SD) vocabulary, which became a W3C Recommendation in March 2013 [35]. Such descriptions, if made widely available, could help a client find public endpoints that support the features it needs (e.g., find SPARQL 1.1 endpoints).

The service description for an endpoint is retrieved by simply dereferencing the endpoint IRI itself [35]. As such, the SPARQLES system performs a HTTP GET request for an endpoint IRI, follows redirects and uses content negotiation to request RDF formats (viz. RDF/XML, N-Triples, Turtle or RDFa).

*VoID Analytics* The Vocabulary of Interlinked Datasets (VoID) [3] has become the de facto standard for describing RDF datasets (in RDF). The vocabulary allows for specifying, e.g., the number of triples a dataset contains, the number of unique subjects, a list of properties and classes used, the number of triples with a given property as predicate, the number of instances of a given class, the number of triples used to describe instances of a given class, and so forth. If VoID descriptions were widely available for SPARQL endpoints, a client could leverage them to discover endpoints with potentially relevant content.

There are a number of best-practices regarding how VoID should be published; SPARQLES looks in three locations. First, the system looks in the content gotten by dereferencing the endpoint URL (i.e., the same

document as the SD description). Second, the system checks the location denoted by the Well-Known IRI pattern `http://{domain}/.well-known/void` recommended for use with VoID, where {domain} is replaced with the fully-qualified domain name (FQDN) extracted from the endpoint URL [11]. Third, SPARQLES uses the following query to detect if the endpoint indexes its own VoID description, where %ep is replaced with the URL of the endpoint in question:

```
PREFIX void: <http://rdfs.org/ns/void#>
SELECT DISTINCT ?ds
WHERE { ?ds a void:Dataset ;
        void:sparqlEndpoint %ep . }
```

*Server Name Analytics* A variety of options are now available for SPARQL engines, including Virtuoso [13], Sesame [10], etc. However, performance and compliance across different vendors can vary quite dramatically. Knowing which engine – or even which version of an engine – powers a given SPARQL endpoint may be useful for (expert) clients to know which version of a query to send. For example, in previous works we found that certain analogous strategies for processing joins in a federated setting worked well for certain SPARQL engines but performed poorly or even outright failed for others [12].

Unfortunately, neither VoID nor the SD vocabulary provide terms for specifying an engine or version number to a client. Hints are available, such as scanning the frontpage or an error page for mention of a fixed list of engines. However, when dereferencing the endpoint URL, the type of engine and the version number is often (though not always) provided in the `Server` field of the HTTP header. Although not always provided – perhaps since it may require low-level server configuration – this is the cleanest method we have found to currently establish which implementation powers an endpoint without requiring hard-coded, engine-specific heuristics.

*Frequency* When compared with availability, we do not expect discoverability to be so dynamic: once descriptions are published, they are likely to stay published (and as discussed later, we do not check that the descriptions are up-to-date). For this reason, we run discoverability analytics once a week; we have received no complaints from endpoint maintainers about the remote expense of these analytics.

*Limitations* SPARQLES only checks for the existence of meta-data, but does not attempt to validate the meta-data itself, nor does it try to measure the completeness of descriptions. Additionally, VoID descriptions or engine information may be extracted from locations or with vocabularies not checked by SPARQLES: however, to help clients, we believe it is important for publishers to offer such information using well-known vocabularies in discoverable locations.

### 4.3. Performance

SPARQLES runs a set of performance-related analytics that aims to compare the runtimes of different public endpoints for comparable queries from a client’s perspective (i.e., including HTTP overhead). Since we cannot control or know in detail about the content of endpoints, for the purpose of comparability, we must rely on generic queries that would execute in a similar manner independent of the exact content indexed by the endpoint. We test three fundamental aspects of a query engine: lookups, streaming and joins.

*Lookup Analytics* The goal is to measure the time taken to perform an atomic lookup (according to different triple patterns). The query template is as follows:

```
ASK {<x> ?p ?o}
```

Here <x> is replaced with an arbitrary IRI that is not expected to exist in the remote data (a lookup still needs to be performed to ensure this). To mitigate caching effects, we generate a fresh IRI each time.

Since in the above example the subject is set, we call it an ASK<sub>s</sub> query. We also run ASK<sub>p</sub>, ASK<sub>o</sub>, ASK<sub>sp</sub>, ASK<sub>so</sub>, ASK<sub>po</sub>, ASK<sub>s-po</sub> versions of the query.

Given that an atomic lookup should be fast to execute, we consider this query as estimating the latency of querying the endpoint, which would most likely be dominated by the HTTP network overhead.

*Streaming Analytics* We measure the time taken for an endpoint to stream a large result-set that should be trivial to compute. The query is as follows:

```
SELECT * {?s ?p ?o} LIMIT 100001
```

Here we ask to stream 100,001 results. Since we have found that public endpoints may limit maximum result sizes to a “round number” – say 100,000 – we ask for one hundred thousand *and one* results to detect such a case. We also send queries for limits with

50,000, 25,000, 12,500, 6,250 and 3,125 results. Since the endpoint should be able to stream results contiguously from its index, we consider such queries as estimating the maximum throughput of the service.

*Join Analytics* We use the following three queries to measure a generic notion of join performance:

```
SELECT DISTINCT ?s ?q
WHERE {?s ?p ?o OPTIONAL {?s ?q <x>}} LIMIT 1000
```

```
SELECT DISTINCT ?s ?q
WHERE {?s ?p ?o OPTIONAL {<x> ?q ?s}} LIMIT 1000
```

```
SELECT DISTINCT ?o ?q
WHERE {?s ?p ?o OPTIONAL {<x> ?q ?o}} LIMIT 1000
```

These queries are designed – insofar as possible – to be comparable across endpoints no matter what content is indexed. In these queries, <x> is a fresh IRI not expected to appear in the data. For example, the first query requests that 1,000 unique subjects be joined with a pattern that generates no answers: this join must still be executed to check that ?q is indeed unbound. The result will return 1,000 distinct subject-unbound pairs. While the first query looks at *s-s* joins, the second performs an *s-o* join and the third an *o-o* join; these three join-types were the most common found in analyses of real-world logs [4,30].

*Frequency* Like availability, we expect performance to vary for different times of the day, different days of the week, etc. For this reason, like availability, we would like to have frequent experiments. However, unlike availability, the queries required to test performance are not so trivial for endpoints to compute (we present more details on this later). For this reason, we opted to run performance experiments once a day; at this level of frequency, we have received no complaints from endpoint maintainers.

*Limitations* The performance results do not indicate *why* specific queries are slow: is it due to the engine, the HTTP overhead, the content indexed? In general, we try to make the query load balanced irrespective of the content and our goal is to measure the costs from the perspective of a client who is concerned about the “bottom line” of response times.

Performance results may also be affected by local issues. For example, slow runtimes may be due to a busy network on the SPARQLES end (e.g., if other analytics happen to run simultaneously); to help mitigate this issue, in the system’s interfaces, we display the

median value of the last ten performance runs. Other factors may be more difficult to control for; e.g., endpoints on servers that are geographically closer to the SPARQLES host may be given an advantage. Still, the performance results should serve as a useful guide.

#### 4.4. Interoperability

If available, the Service Description of an endpoint should describe the query features and the version of SPARQL that an endpoint supports. However, we have seen that SD meta-data are often unavailable and, in any case, an endpoint may claim to support features that it does not, or may claim support for SPARQL 1.1 while only supporting a subset of new features.

SPARQLES thus offers analytics for interoperability, whose goal is to verify which SPARQL features – i.e. specific operators, solution modifiers, etc. – are supported, gathering data about what SPARQL features are available for the users of various endpoints.

Along these lines, SPARQLES takes a subset of queries from the W3C Data Access Working Group test-cases – designed to test all features from both versions of the standard – and issues them on a weekly basis to SPARQL endpoints. We consider the test as passed if a valid SPARQL response is returned. Since we cannot control the content of endpoints, we cannot verify that the returned response is actually correct; hence we may overestimate compliance with the standard. We expect that if an endpoint does not support a feature, an exception will be thrown (e.g., a parse exception). However, since an endpoint may time-out on a given query, we may also underestimate compliance where the feature may be supported but the endpoint cannot answer the query instance provided.

*SPARQL 1.0 Analytics* First, the SPARQLES system tests the endpoints for the core SPARQL 1.0 query features that it supports. We issue endpoints a subset of the Data Access Working Group test-cases for SPARQL 1.0,<sup>11</sup> omitting syntax tests and focusing on core functionalities.<sup>12</sup> This test-set checks a range of aspects of the SPARQL 1.0 standard including query types SELECT, CONSTRUCT and ASK (we omit DESCRIBE since it is an optional feature); filter features, such as REGEX, IRI and blank node checks, etc.; support for datatypes, such as numerics, strings and

booleans; support for graph selection features, including FROM (NAMED) and GRAPH; and the solution modifiers, ORDER BY, LIMIT and OFFSET (DESC|ASC), as well as DISTINCT and REDUCED modifiers.

*SPARQL 1.1 Analytics* SPARQLES also performs tests on SPARQL 1.1 features using a test suite taken from the W3C SPARQL Working Group.<sup>13</sup>

We first test support for aggregates, where expressions such as average, maximum, minimum, sum and count can be applied over groups of solutions (possibly using an explicit GROUP BY clause). We then test support for sub-queries in combination with other features. Next we test support for property-paths, binding of individual variables, and support for binding tuples of variables (VALUES). We also check support for filter features that check for the existence of some data (MINUS, EXISTS), and some new operator expressions (STRSTARTS and STRCONTAINS for strings; ABS for numerics). Finally, the last three queries test a miscellany of features including NOT IN used to check a variable binding against a list of filtered values, an abbreviated version of CONSTRUCT queries whereby the WHERE clause can be omitted, and support for the SPARQL SERVICE keyword, which invokes a federated request from the remote endpoint being tested to a local endpoint we have set up on the SPARQLES server.

*Frequency* Much like discoverability, we do not expect interoperability to be a highly dynamic property of an endpoint; for example, we suppose that once an endpoint adds support for SPARQL 1.1 features, it will continue to support these features until it is discontinued. For this reason, we schedule interoperability analytics to run once a week. During the first year of operation, we began to receive complaints that the queries we were using – based on the W3C test-cases – were causing a high load on remote servers.<sup>14</sup> We considered lowering the frequency of these analytics but instead decided to make the queries less demanding by refactoring them to include fresh IRIs in such a way that it should be efficient for the server to compute that the result is empty; as usual, we then simply monitor for exceptions. With these new queries, we have received no complaints thus far.

<sup>11</sup><http://www.w3.org/2001/sw/DataAccess/tests/r2>

<sup>12</sup>Queries available at <https://github.com/pyvandenbussche/sparqles>.

<sup>13</sup><http://www.w3.org/2009/sparql/docs/tests/data-sparql11/>

<sup>14</sup>See, e.g., <https://github.com/pyvandenbussche/sparqles/issues/23>.



*Limitations* As aforementioned, the main limitation of these experiments is that we classify an endpoint as implementing a specific SPARQL 1.1 feature if that endpoint returns any valid response without throwing an exception. If an endpoint times out, we will classify it as not implementing that feature, and conversely, if it returns an incorrect solution, we will count it as supporting that feature.<sup>15</sup> Our recent refactoring of the test-case queries has helped to reduce false negatives due to time-outs. Another limitation is that the SERVICE call relies on our local endpoint being accessible; to mitigate problems, we designed the federated call to be cheap and simple and to restart our local endpoint just before invoking these analytics.

Another limitation is that we do not test some features of the SPARQL 1.1 standard, such as SPARQL 1.1 Update – since we presume we should not have write privileges for public endpoints – or SPARQL 1.1 Entailment Regimes – since we do not know of any public SPARQL endpoints with this feature.

## 5. Storage & Interfaces

We now describe how SPARQLES manages the experimental data gathered during these analyses and the public interfaces through which software agents and users can interact with these data.

### 5.1. Storage

As tests are performed, the results and metrics collected are serialised using the Apache AVRO (1.7.5) library and sent to a MongoDB instance for storage. The MongoDB instance maintains 11 different collections that, loosely speaking, represent different materialised views over the data collected:

- 4 collections store the “raw” version of the data collected for the four analytical dimensions;
- 1 collection maintains the current list of endpoints registered in the DataHub;
- 6 collections correspond to aggregated views of the raw data as required by the User Interface.

The aggregate views are recomputed at regular intervals using cron jobs: these views return the data re-

<sup>15</sup>Strictly speaking, a query timing out is not compliant with SPARQL 1.1; however, in spirit, we are more interested about whether a feature is supported in general and not about if a specific query instance runs or not.

quired by the U.I. in a single lookup and thus avoid running aggregations while the user waits.

### 5.2. Application Programming Interfaces

SPARQLES provides seven RESTful APIs that facilitate remote access to the data collection. These APIs are designed to provide clients with both (i) information relating to specific endpoints, as well as (ii) information about all endpoints relating to a specific type of analytical experiment. Likewise we split the APIs into two groups, as follows.

The first group contains *endpoint-specific APIs*, which helps to find a specific endpoint, or returns detailed results for a specific endpoint:

**LIST** takes no input and returns a list of all endpoints in SPARQLES: the URL of the endpoint, as well as the name and URL of the dataset they are associated with in the DataHub catalogue.

**AUTOCOMPLETE** takes as input a string, such as “dbpedia”, and returns all endpoints in SPARQLES whose URL, dataset label, or dataset URL contains the input as a sub-string.

**INFO** takes as input the URL of a specific endpoint, and returns all data for that endpoint, including dataset URL and label; availability over the past day, week, month, and overall; performance results for cold/warm runs of ask and join queries and suspected threshold size; interoperability information regarding support for SPARQL and SPARQL 1.1 query features; and discoverability information regarding locations (if any) where VoID and SD descriptions could be found, and the server-name extracted from the HTTP header.

The second group contains *analytical APIs*, which return aggregate results for all endpoints on a given analytical dimension:

**AVAILABILITY** returns uptimes for the last test, day, week, month – and overall – for all endpoints.

**DISCOVERABILITY** provides the server name, VoID locations and SD availability found (if any) for each endpoint in the most recent experiment.

**INTEROPERABILITY** counts the SPARQL 1.0 and SPARQL 1.1 queries passed by each endpoint in the most recent experiment.

**PERFORMANCE** provides the mean performance for join and ask queries in cold/warm runs from the most recent experiment for all endpoints.

The APIs are provided and described online at: <http://sparqles.org/api>.

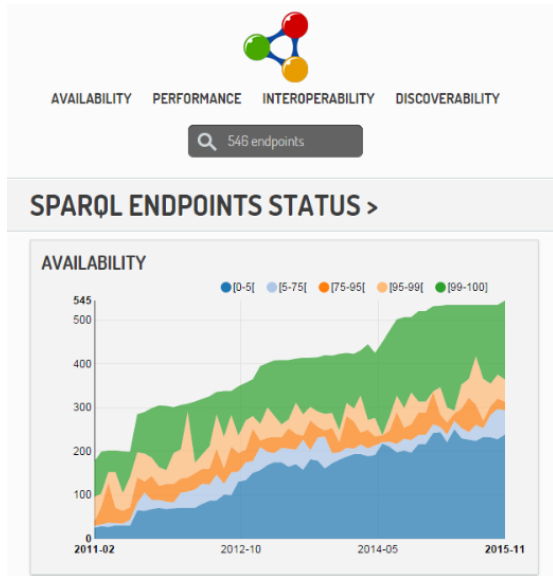


Fig. 2. Screen capture of SPARQLES’ homepage focusing on the availability overview. The homepage offers “at-a-glance” aggregated views for the four dimensions of analytics described.

### 5.3. User Interface

The SPARQLES user interface – available on the Web at <http://sparql.es.org/> – offers an entry point for human users interested in the experimental results. The interface is implemented using the Node.js Javascript runtime engine and the nvd3 Javascript library for rendering interactive visualisations.

The homepage offers “at-a-glance” aggregated views of the four dimensions computed across all endpoints. In Figure 2, for example, we see the aggregate view for availability, which shows the evolution of the number of endpoints falling into five different availability intervals ( $[0-5]$ ,  $[5-75]$ ,  $[75-95]$ ,  $[95-99]$ ,  $[99-100]$ ). Other aggregate views likewise provide an overview of performance, interoperability and discoverability.

From the homepage, the user has a number of possible navigation steps, as illustrated in Figure 3.

The user can navigate to a page dedicated to each dimension to get an overview of key results for all endpoints in a list view, as follows:

**AVAILABILITY** lists the availability for each endpoint over the past 24 hours and the past 7 days.

**PERFORMANCE** lists the suspected result-size threshold and the median cold/warm run-times for ASK and JOIN queries over the last 10 runs.

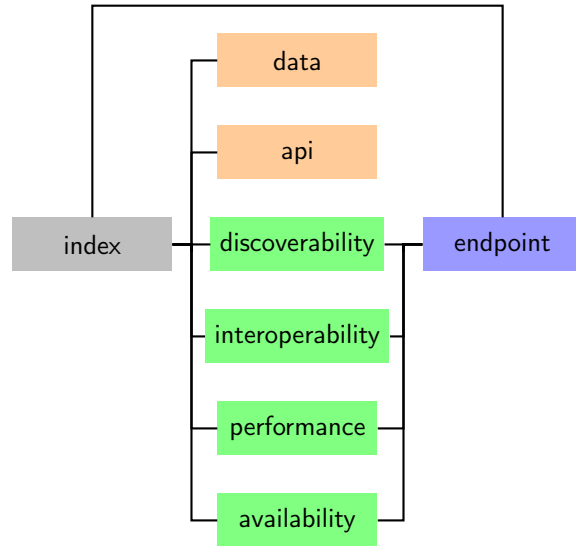


Fig. 3. SPARQL user-interface sitemap

**INTEROPERABILITY** lists the ratio of SPARQL 1.0 and SPARQL 1.1 query test-cases passed by each endpoint in the most recent run.

**DISCOVERABILITY** indicates whether or not a VoID and/or SD description is available for each endpoint in some location, and what server name could be found (if any), in the most recent run.

Otherwise, either by using the auto-complete search function (available on all pages), or by clicking on a specific endpoint mentioned in one of the previous four list-view pages, the user can arrive to an endpoint-specific view with detailed information about all four dimensions for a given endpoint. An example for the main DBpedia endpoint is provided in Figures 4 and 5 (referring to the same page but split here for formatting purposes). The views provide information on the weekly availability for the past year, median performance for the past ten runs, the interoperability test queries failed or passed<sup>16</sup> as well as the locations in which VoID/SD descriptions can be found and the server name detected (if any).

From the homepage, there are also links to download data dumps or view the A.P.I. documentation.

<sup>16</sup>On hovering the mouse pointer over the query name, the user can see the full query, and on hovering over a red icon, the user can see details of the exception encountered.

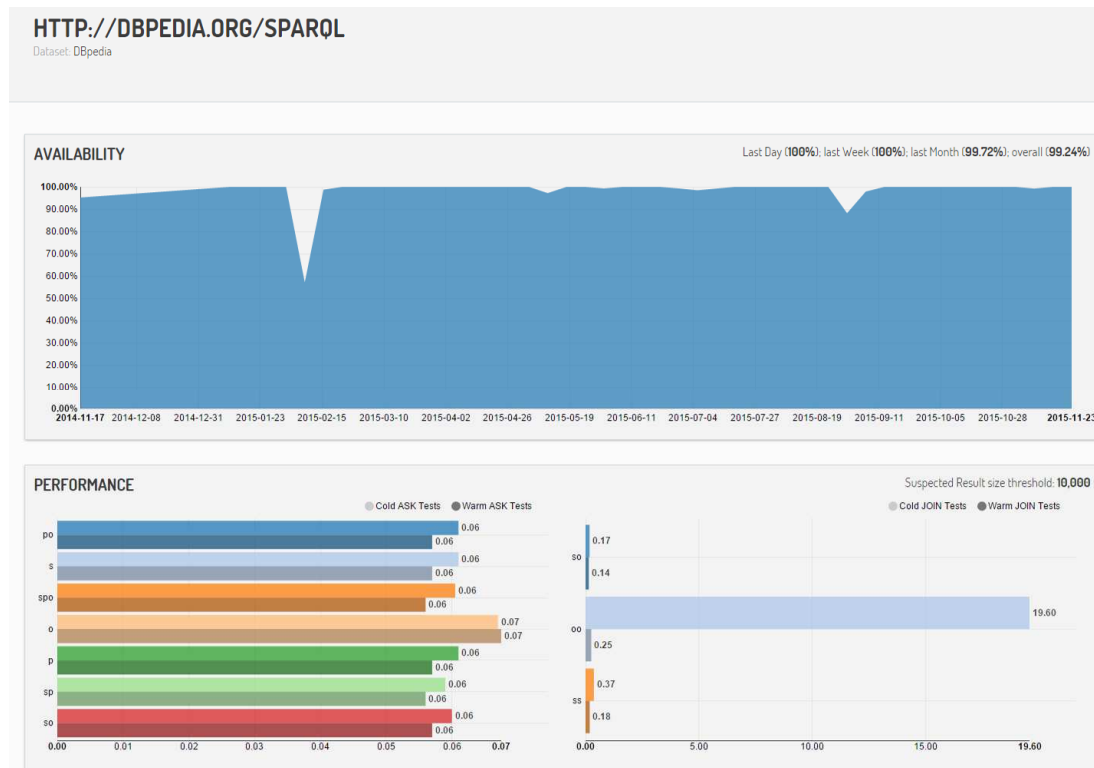


Fig. 4. Screen capture of the detailed view for the DBpedia endpoint (<http://dbpedia.org/sparql>) showing the results of availability tests for the past year and performance tests based on the median result for the most recent ten runs.

## 6. Evaluation

SPARQLES has been running since October 2013 where, as of November 2015, it monitors 545 public endpoints. We now present some high-level results with respect to operating, maintaining and improving various aspects of the SPARQLES system.

*Analytics* With respect to running the individual analytical tasks, we measured the time of the four most recent runs within the SPARQLES system.

The fastest were the AVAILABILITY experiments, which took between 41–44 minutes ( $\mu = 41.8$  minutes,  $\sigma = 1.1$  minutes) in the four most recent runs. These experiments involve sending either one or two simple queries to each server. This runtime fits within the current one hour interval; however, if we were to increase the number of endpoints observed by about 33%, we would need to increase the interval to (say) two hours, to ensure that the analytics terminate in time.

Second were the DISCOVERABILITY experiments, which took between 65–67 minutes ( $\mu = 65.5$  minutes,  $\sigma = 1$  minute) in the most recent four runs. These ex-

periments involve sending a single query and checking two Web documents for metadata about each endpoint. These experiments could easily be maintained at the weekly interval, even if the number of endpoints observed were to increase, e.g., a hundredfold.

Third were the INTEROPERABILITY experiments, which took between 65–87 minutes ( $\mu = 72.6$  minutes,  $\sigma = 9$  minutes) in the four most recent runs. These experiments involve running 42 queries against each remote endpoints. However, since we have refactored these queries to return empty results, they should be trivial to run. These experiments are likewise easily maintained within their current weekly interval.

Fourth were the PERFORMANCE experiments, which took between 537–614 minutes ( $\mu = 583.5$  minutes,  $\sigma = 28.8$  minutes) in the four most recent runs. Against each endpoint, we run 17 queries twice (once cold and once warm). These experiments take significantly longer than the others: although INTEROPERABILITY has more queries, the queries in PERFORMANCE are designed to be non-trivial to answer, returning thousands of results and requiring a significant amount of processing for the endpoint to service. Within the cur-

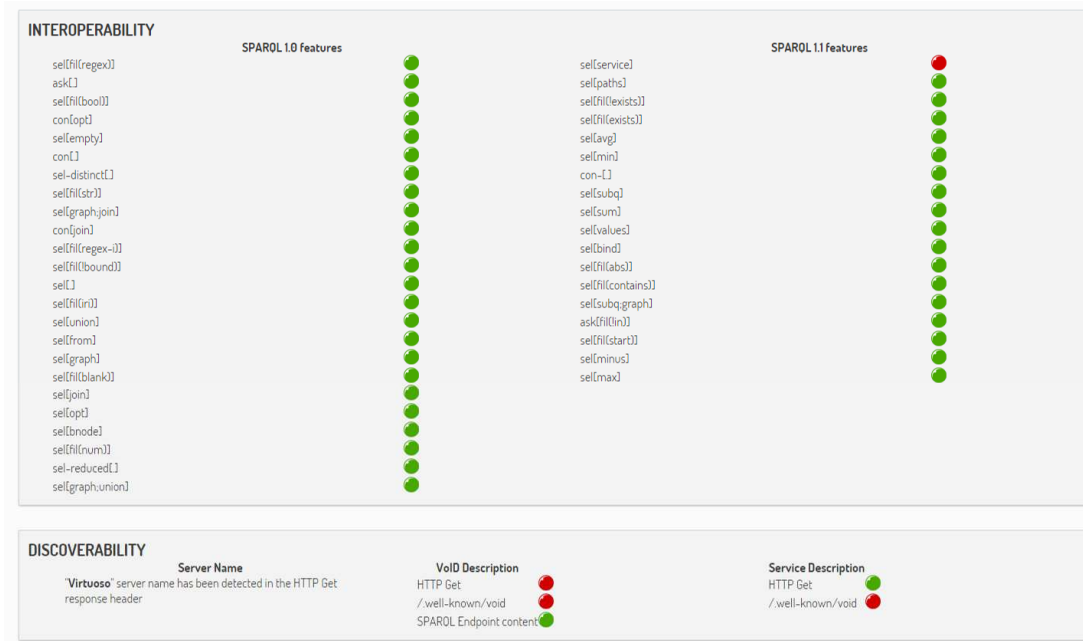


Fig. 5. Screen capture of the detailed view for the DBpedia endpoint (<http://dbpedia.org/sparql>) showing the results of interoperability (individual queries passed or failed) and discoverability tests (locations of VoID and SD descriptions) in the most recent run.

rent daily interval and settings, we could support these analytics for about twice the number of endpoints.

**Storage** With respect to the growth of data in the SPARQLES system, the maximum amount of data we collect per endpoint is given in Table 1, where, each week, we store an upper limit of 177 JSON objects and 4,415 JSON key–value pairs per endpoint (we store fewer key–value pairs if an endpoint is unavailable). In terms of the growth of endpoints, one can see from Figure 2 that the number of endpoints increases gradually. Hence we can conclude that the storage overhead of SPARQLES is sustainable, particularly given the ability of MongoDB to scale horizontally across machines. For example, as of November 2015, SPARQLES was monitoring 545 endpoints, and thus collecting 96,465 JSON objects and a max. of 2,406,1175 JSON key–value pairs per week, where the compressed dump of all historical data was 397 megabytes, corresponding to a live MongoDB index of 14 gigabytes.

**A.P.I.** In order to ensure that our A.P.I.’s would perform well under high loads, we sent 1,000 requests to each of our A.P.I.’s from 10, 50 and 100 parallel clients respectively. Figure 6 provides box-plots of the individual runtimes encountered. Under this type of load,

Table 1

Data items stored by SPARQLES per endpoint: **Freq.** denotes the frequency of experiments, **KV** denotes the maximum number of key–value pairs generated per experiment, **O/W** denotes the number of JSON objects generated per week (each experiment generates one JSON object), and **KV/W** denotes the maximum number of key–value pairs per week (i.e.,  $\text{KV/W} = \text{KV} \times \text{O/W}$ )

Analysis	Freq.	KV	O/W	KV/W
AVAILABILITY	hourly	11	168	1,848
PERFORMANCE	daily	306	7	2,142
DISCOVERABILITY	weekly	35	1	35
INTEROPERABILITY	weekly	390	1	390
<b>Total</b>	—	—	177	4,415

although we see that the slowest request can range in the tens of seconds (especially for a higher number of clients), typical performance in the lower three quartiles remains reliably below half a second. In summary, we encounter a few slow requests that require up to 20 seconds to complete, but the majority of requests are answered within 0.5 seconds, even with 100 clients simultaneously issuing 1,000 requests.

**Usability** We collect and react to user feedback relating to bugs, feature-requests and usability on the

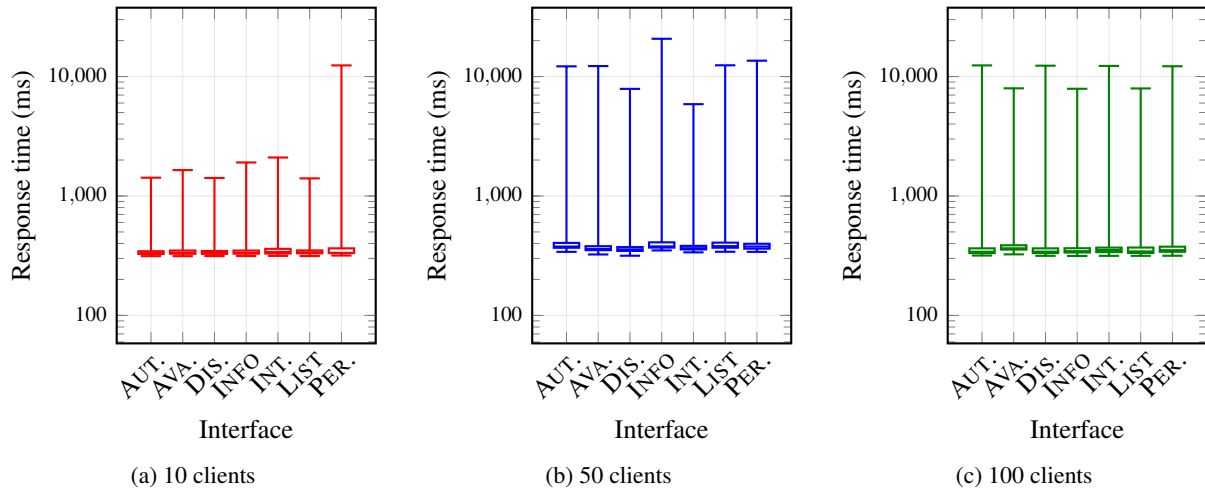


Fig. 6. A.P.I. response times for 1,000 requests with 10, 50 and 100 parallel clients. On the  $x$ -axis, AUT. denotes Autocomplete, AVA. denotes Availability, DIS. denotes Discoverability, INT. denotes Interoperability, PERF. denotes performance. The  $y$ -axis shows the response times in milliseconds; the axis is given in log scale and aligned horizontally for comparison across the three plots. The box plots are drawn for the maximum, third quartile (75<sup>th</sup> percentile), median, first quartile (25<sup>th</sup> percentile) and minimum times from the individual responses.

issue tracker for the project, which as of 2015/01/30 contained 21 open issues and 24 closed issues; we refer the interested reader to <https://github.com/pyvandenbussche/sparqles/issues> for details on individual comments. Feedback has related to varied aspects, be it misreported statistics for the endpoints maintained by users, or the expense of certain queries for remote services, or problems with characters/escaping in the interface, or requests for various enhancements. This feedback has been invaluable for improving the usability, correctness and sustainability of the service. For example, one of the most important changes we have made based on this feedback was to modify the INTEROPERABILITY queries to reduce the computational strain they were placing on the public endpoints that SPARQLES monitors. We are continuing to address the open issues and consider the other feedback and requests for enhancements that we receive from the community.

## 7. Discussion

We now discuss the use-cases, impact thus far, limitations and sustainability of SPARQLES.

### 7.1. Impact

One of the main goals of the system is to disseminate timely information about the health of individual endpoints. To help characterise the impact of SPAR-

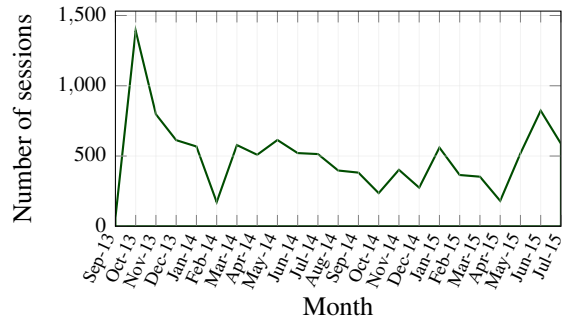


Fig. 7. Number of unique sessions per month for SPARQLES

QLES, in the following we present some statistics collected from the Google Analytics for the site.

Over a 23 month period, SPARQLES has seen a total of 11,420 user sessions, averaging about 497 sessions per month. Figure 7 presents the data for the past 23 months, where we see a peak in October 2013, after which the number of user sessions was between 172 (February 2014) and 823 (June 2015). Table 2 presents the number of sessions along with the most visited first pages, second pages, and third pages.<sup>17</sup>

Another indirect goal of the system is to encourage endpoints to follow best practices: we would hope

<sup>17</sup>Please note that since we cannot access raw data, some of the figures may be rounded (for example, the Google Analytics system reports “3.4K” rather than an exact figure). Likewise we only have details of visits to the top twenty pages, hence we may only have an upper-bound for other pages.

Table 2  
Number of SPARQLES sessions over a 22 month period spanning  
from 2013/09/27–2015/07/27.

START PAGE		1 <sup>ST</sup> INTERACTION		2 <sup>ND</sup> INTERACTION	
Total	11,420	Total	3,400	Total	2,300
/	8,300	/availability	1,500	specific endpoint	1,202
specific endpoint	1,247	specific endpoint	949	/availability	340
/availability	1,200	/discoverability	286	/	326
/discoverability	221	/interoperability	214	/interoperability	179
/interoperability	172	/performance	201	/performance	151
/performance	45	/	133	/discoverability	126
/api	16	/api	27	/api	<12
/data	<11	/data	<7	/data	<12

that by tracking such metrics about endpoints, maintainers might be made aware of shortcomings with the SPARQL services they offer and rectify these accordingly. Though from personal communications with some endpoint maintainers we know that there have been anecdotal instances of this,<sup>18</sup> it is difficult to ascertain to what degree SPARQLES has had an impact on the maturity of SPARQL endpoints in this respect.

Perhaps the most important impact of this work has been to formally acknowledge the kinks in the current public SPARQL infrastructure, which has helped motivate new lines of research. We can, for example, point to work on Linked Data Fragments, which points to our AVAILABILITY and PERFORMANCE statistics to justify why alternatives to SPARQL are needed [33, 34]. We can also point to our own work which takes into account the result-size limitations of endpoints from our PERFORMANCE results to motivate custom federation strategies [12]. Aside from using the statistics from SPARQLES to motivate new lines of research, a number of research works have also used the SPARQLES system itself, for example, Benedetti et al. [6] use our endpoint list, and in particular our AVAILABILITY and INTEROPERABILITY statistics, to find active endpoints supporting SPARQL 1.1 features (as their profiling tool requires); Rietveld & Hoekstra [29] use the SPARQLES search API to provide endpoint search functionality in their SPARQL client; Atemezing & Troncy [5] use our API to collect a list of SPARQL endpoints for their experiments; Acosta et al. [1,2] use our PERFORMANCE statistics to design hybrid query plans that take the limitations of endpoints into account; and so forth.

## 7.2. Limitations

For each of the analytics presented in Section 4, we discussed a variety of specific limitations, referring, e.g., to the difficulty in distinguishing local problems from remote problems. There are also a couple of global limitations of the system worth mentioning.

First, SPARQLES is subject to Goodhart’s law, which states: “*When a measure becomes a target, it ceases to be a good measure.*” An over-eager endpoint maintainer could, e.g., detect and artificially respond to SPARQLES’ queries so as to improve how the endpoint is “rated” by the system. We know of no such example of this happening but it is possible.

Second, as a more pragmatic issue, since we first put the system online in October 2013, we have had various local reliability issues, where data were not collected for certain weeks, where data were lost due to server migration, and where the site itself was offline. During this period, we have been resolving various issues as they occur such that, although there are still some known issues, we now believe that the system is reaching maturity. Likewise, we have received a lot of feedback from the community, which has been invaluable for improving the service in the past years.

Third, some of the analytics may be biased towards servers that are closer geographically to the SPARQLES host in Austria. One option to mitigate this bias – as well as local reliability issues – would be to replicate SPARQLES analytics in multiple remote locations and create a mechanism for aggregating a global consensus across all remote instances. Currently we do not have the resources available to host another instance of SPARQLES. However, the SPARQLES code is available for download, where the community can

<sup>18</sup>See for example <https://github.com/pyvandenbussche/sparqles/issues/42>.

download and install their own instances, perhaps targeted at those endpoints of interest to them.

### 7.3. Sustainability

One indirect but important aspect of sustainability is the load that SPARQLES puts on the public SPARQL infrastructure. For example, we discussed before about how the original versions of the interoperability queries were causing a heavy load for a number of SPARQL services. To mitigate this, we run more expensive tasks less frequently: while simple availability tests are done hourly, performance analytics are run daily and interoperability tests are run weekly. Likewise we have revised the interoperability queries to make them less costly and have been attentive in addressing all complaints raised in our issue tracker relating to the cost SPARQL puts on remote servers.

With respect to the long term sustainability of SPARQLES, the system is currently hosted on a university server in Austria and we have the commitment of the university – including help from system administrators – to continue hosting the service for the foreseeable future. We also provide a weekly dump of historical data, which is backed-up internally, as well as the open-source code and instructions needed to generate a new SPARQLES instance. Assuming a worst case scenario where the university would no longer be able to host the service or the data and we were no longer contactable, someone from the community could create a fresh instance of SPARQLES, albeit (in that unlikely scenario) without historical data. We of course welcome the community to mirror the service and also the historical data. Again, all code and data are available under CC-BY (© ⓘ), Version 4.0.

### 7.4. Future Directions

As the system has been maturing, we have started to consider adding some new features as requested by the community. One of the most popularly requested features is to have data collected by the SPARQLES tool made available as Linked Data. Though we are (perhaps ironically) reluctant to make a SPARQL endpoint available, as a starting point, we are looking into creating Linked Data IRIs for individual endpoints that dereference to SPARQLES statistics about them. Other requested features included offering an email notification system to contact endpoint administrators when their system was not available, or offering badges for endpoints with high availability, and so forth. There are

numerous directions in which SPARQLES could still be improved, which we will tackle as time progresses.

More in terms of research, we foresee a number of potential future directions relating to SPARQLES and the monitoring of public SPARQL endpoints. One possible direction would be to look, more systematically, at further dimensions of quality specific to SPARQL endpoints – perhaps in the broader context of the survey of Linked Data quality by Zaveri et al. [36] – and to design practical tests and more detailed metrics by which to monitor endpoints. Another potentially fruitful area for future investigation would be to perform formal usability studies for the system – testing times users need to perform set tasks, conducting AB tests, etc. – and adapting the SPARQLES infrastructure accordingly. Although we have been reacting to community feedback through the issue tracker and other sources, this has been in an ad hoc manner: controlled user studies would offer a more structured approach to improving SPARQLES.

We remain very much open to future collaboration with the community along these or other lines.

## 8. Conclusion

In this paper, we have presented the SPARQL Endpoint Status (SPARQLES) system for keeping track of the health and maturity of public SPARQL endpoints. We presented the high-level architecture, which consists of an offline component for running tests over endpoints, and an online component for providing visualisations and A.P.I.'s for the collected results. We presented four dimensions of analytics that the system runs over public endpoints. Thereafter, we presented some of the details of the interfaces SPARQLES provides for human and automated agents to interact with the underlying data collection. We also presented some measures with respect to the overall runtime of analytics, the growth in storage requirements, and the performance of our A.P.I.'s under load. Finally, we discussed some aspects relating to the high-level impact, limitations and sustainability of the tool.

In general, we believe that the SPARQLES system provides the community with a unique perspective on public SPARQL endpoints. The online system helps to shed light not only on some of the cobwebs and cracks in the SPARQL infrastructure, but also on the cream of the crop: those SPARQL endpoints that are highly-available, readily-discoverable, highly-performant and highly-interoperable.

**Acknowledgements** This work was supported by Fujitsu Laboratories Limited, by CONICYT/FONDECYT Project no. 3130617, by FONDECYT Project no. 11140900, by DGIP Project no. 116.24.1, and by the Millennium Nucleus Center for Semantic Web Research under Grant NC120004. We would like to thank the reviewers of this and our previous paper, and various other members of the Linked Data community, for their feedback on SPARQLES. We also wish to thank the Open Knowledge Foundation for hosting the project for over a year.

## References

- [1] M. Acosta, M. Vidal, F. Flöck, S. Castillo, C. Buil-Aranda, and A. Harth. SHEPHERD: A shipping-based query processor to enhance SPARQL endpoint performance. In M. Horridge, M. Rospocher, and J. van Ossenbruggen, editors, *Proceedings of the ISWC 2014 Posters & Demonstrations Track – a track within the 13th International Semantic Web Conference, ISWC 2014, Riva del Garda, Italy, October 21, 2014.*, volume 1272 of *CEUR Workshop Proceedings*, pages 453–456. CEUR-WS.org, 2014. URL [http://ceur-ws.org/Vol-1272/paper\\_147.pdf](http://ceur-ws.org/Vol-1272/paper_147.pdf).
- [2] M. Acosta, M. Vidal, F. Flöck, S. Castillo, and A. Harth. PLANET: query plan visualizer for shipping policies against single SPARQL endpoints. In M. Horridge, M. Rospocher, and J. van Ossenbruggen, editors, *Proceedings of the ISWC 2014 Posters & Demonstrations Track – a track within the 13th International Semantic Web Conference, ISWC 2014, Riva del Garda, Italy, October 21, 2014.*, volume 1272 of *CEUR Workshop Proceedings*, pages 189–192. CEUR-WS.org, 2014. URL [http://ceur-ws.org/Vol-1272/paper\\_148.pdf](http://ceur-ws.org/Vol-1272/paper_148.pdf).
- [3] K. Alexander, R. Cyganiak, M. Hausenblas, and J. Zhao. Describing linked datasets. In C. Bizer, T. Heath, T. Berners-Lee, and K. Idehen, editors, *Proceedings of the WWW2009 Workshop on Linked Data on the Web, (LDOW) 2009, Madrid, Spain, April 20, 2009.*, volume 538 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2009. URL [http://ceur-ws.org/Vol-538/ldow2009\\_paper20.pdf](http://ceur-ws.org/Vol-538/ldow2009_paper20.pdf).
- [4] M. Arias Gallego, J. D. Fernández, M. A. Martínez-Prieto, and P. de la Fuente. An empirical study of real-world SPARQL queries. In B. Berendt, L. Hollink, V. Hollink, M. Luczak-Rösch, K. Möller, and D. Vallet, editors, *Proceedings of the 1st International Workshop on Usage Analysis and the Web of Data (USEWOD 2011) co-located with the 20th International World Wide Web Conference (WWW 2011), Hyderabad, India, March 28th, 2011.*, 2011.
- [5] G. A. Atemez and R. Troncy. Towards a linked-data based visualization wizard. In O. Hartig, A. Hogan, and J. Sequeda, editors, *Proceedings of the 5th International Workshop on Consuming Linked Data (COLD 2014) co-located with the 13th International Semantic Web Conference (ISWC 2014), Riva del Garda, Italy, October 20, 2014.*, volume 1264 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2014. URL [http://ceur-ws.org/Vol-1264/cold2014\\_AtemezT.pdf](http://ceur-ws.org/Vol-1264/cold2014_AtemezT.pdf).
- [6] F. Benedetti, S. Bergamaschi, and L. Po. Visual querying LOD sources with LODeX. In K. Barker and J. M. Gómez-Pérez, editors, *Proceedings of the 8th International Conference on Knowledge Capture, K-CAP 2015, Palisades, NY, USA, October 7-10, 2015*, pages 12:1–12:8. ACM, 2015. DOI: 10.1145/2815833.2815849.
- [7] B. Berendt, L. Hollink, V. Hollink, M. Luczak-Rösch, K. Möller, and D. Vallet. Usage analysis and the Web of Data. *SIGIR Forum*, 45(1):63–69, 2011. DOI: 10.1145/1988852.1988864.
- [8] T. Berners-Lee. Linked Data. Design issues for the World Wide Web, World Wide Web Consortium, 2006. URL <https://www.w3.org/DesignIssues/LinkedData.html>.
- [9] E. Blomqvist and T. Groza, editors. *Proceedings of the ISWC 2013 Posters & Demonstrations Track, Sydney, Australia, October 23, 2013*, volume 1035 of *CEUR Workshop Proceedings*, 2013. CEUR-WS.org. URL <http://ceur-ws.org/Vol-1035>.
- [10] J. Broekstra, A. Kampman, and F. van Harmelen. Sesame: A generic architecture for storing and querying RDF and RDF schema. In I. Horrocks and J. A. Hendler, editors, *The Semantic Web – ISWC 2002, First International Semantic Web Conference, Sardinia, Italy, June 9–12, 2002, Proceedings*, volume 2342 of *Lecture Notes in Computer Science*, pages 54–68. Springer, 2002. ISBN 3-540-43760-6. DOI: 10.1007/3-540-48005-6\_7.
- [11] C. Buil-Aranda, A. Hogan, J. Umbrich, and P. Vandenbussche. SPARQL web-querying infrastructure: Ready for action? In H. Alani, L. Kagal, A. Fokoue, P. T. Groth, C. Biemann, J. X. Parreira, L. Aroyo, N. F. Noy, C. Welty, and K. Janowicz, editors, *The Semantic Web – ISWC 2013 – 12th International Semantic Web Conference, Sydney, NSW, Australia, October 21–25, 2013, Proceedings, Part II*, volume 8219 of *Lecture Notes in Computer Science*, pages 277–293. Springer, 2013. DOI: 10.1007/978-3-642-41338-4\_18.
- [12] C. Buil-Aranda, A. Polleres, and J. Umbrich. Strategies for executing federated queries in SPARQL1.1. In P. Mika, T. Tudorache, A. Bernstein, C. Welty, C. A. Knoblock, D. Vrandečić, P. T. Groth, N. F. Noy, K. Janowicz, and C. A. Goble, editors, *The Semantic Web – ISWC 2014 – 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part II*, volume 8797 of *Lecture Notes in Computer Science*, pages 390–405. Springer, 2014. DOI: 10.1007/978-3-319-11915-1\_25.
- [13] O. Erling. Virtuoso, a Hybrid RDBMS/Graph Column Store. *IEEE Data Eng. Bull.*, 35(1):3–8, 2012. URL <http://sites.computer.org/debull/A12mar/vicol.pdf>.
- [14] I. Ermilov, M. Martin, J. Lehmann, and S. Auer. Linked Open Data Statistics: Collection and exploitation. In P. Klinov and D. Mouromtsev, editors, *Knowledge Engineering and the Semantic Web – 4th International Conference, KESW 2013, St. Petersburg, Russia, October 7–9, 2013. Proceedings*, volume 394 of *Communications in Computer and Information Science*, pages 242–249. Springer, 2013. DOI: 10.1007/978-3-642-41360-5\_19.
- [15] L. Feigenbaum, G. T. Williams, K. G. Clark, and E. Torres. SPARQL 1.1 Protocol. W3C Recommendation 21 March 2013, 2013. URL <https://www.w3.org/TR/sparql11-protocol/>.
- [16] P. Gearon, A. Passant, and A. Polleres. SPARQL 1.1 Update. W3C Recommendation 21 March 2013, 2013. URL <https://www.w3.org/TR/sparql11-update/>.



- //www.w3.org/TR/sparql11-update/.
- [17] B. Glimm and C. Ogbuji. SPARQL 1.1 Entailment Regimes. W3C Recommendation 21 March 2013, 2013. URL <https://www.w3.org/TR/sparql11-entailment/>.
- [18] S. Harris and A. Seaborne. SPARQL 1.1 Query Language. W3C Recommendation 21 March 2013, 2013. URL <https://www.w3.org/TR/sparql11-query/>.
- [19] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, and S. Rudolph. OWL 2 Web Ontology Language Primer (Second Edition). W3C Recommendation 11 December 2012, 2012. URL <https://www.w3.org/TR/owl2-primer/>.
- [20] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, and C. Bizer. DBpedia – A large-scale, multilingual knowledge base extracted from Wikipedia. *Semantic Web*, 6(2):167–195, 2015. DOI: 10.3233/SW-140134.
- [21] J. Lorey. Identifying and determining SPARQL endpoint characteristics. *International Journal of Web Information Systems*, 10(3):226–244, 2014. DOI: 10.1108/IJWIS-03-2014-0007.
- [22] M. Mehdi, A. Iqbal, A. Hogan, A. Hasnain, Y. Khan, S. Decker, and R. Sahay. Discovering domain-specific public SPARQL endpoints: a life-sciences use-case. In B. C. Desai, A. M. de Almeida, J. Bernardino, and E. F. Gomes, editors, *18th International Database Engineering & Applications Symposium, IDEAS 2014, Porto, Portugal, July 7–9, 2014*, pages 39–45. ACM, 2014. DOI: 10.1145/2628194.2628220.
- [23] H. Paulheim and S. Hertling. Discoverability of SPARQL endpoints in linked open data. In Blomqvist and Groza [9], pages 245–248. URL [http://ceur-ws.org/Vol-1035/iswc2013\\_poster\\_17.pdf](http://ceur-ws.org/Vol-1035/iswc2013_poster_17.pdf).
- [24] J. Pérez, M. Arenas, and C. Gutierrez. Semantics and complexity of SPARQL. *ACM Transactions on Database Systems*, 34(3), 2009. DOI: 10.1145/1567274.1567278.
- [25] E. Prud’hommeaux and C. Buil-Aranda. SPARQL 1.1 Federated Query. W3C Recommendation 21 March 2013, 2013. URL <https://www.w3.org/TR/sparql11-federated-query/>.
- [26] E. Prud’hommeaux and A. Seaborne. SPARQL Query Language for RDF. W3C Recommendation, January 2008.
- [27] S. Ran. A model for web services discovery with QoS. *SIGecom Exchanges*, 4(1):1–10, 2003. DOI: 10.1145/844357.844360.
- [28] L. Rietveld and R. Hoekstra. Man vs. machine: Differences in SPARQL queries. In B. Berendt, L. Hollink, and M. Luczak-Rösch, editors, *Proceedings of the 4th International Workshop on Usage Analysis and the Web of Data (USEWOD 2014) co-located with the 11th Extended Semantic Web Conference (ESWC 2014), Crete, Greece, May 25th, 2014*, 2014.
- [29] L. Rietveld and R. Hoekstra. The YASGUI family of SPARQL clients. *Semantic Web*, 8(3):373–383, 2017. DOI: 10.3233/SW-150197.
- [30] M. Saleem, M. I. Ali, A. Hogan, Q. Mehmood, and A. Ngonga Ngomo. LSQ: the linked SPARQL queries dataset. In M. Arenas, Ó. Corcho, E. Simperl, M. Strohmaier, M. d’Aquin, K. Srinivas, P. T. Groth, M. Dumontier, J. Heflin, K. Thirunarayan, and S. Staab, editors, *The Semantic Web – ISWC 2015 – 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11–15, 2015, Proceedings, Part II*, volume 9367 of *Lecture Notes in Computer Science*, pages 261–269. Springer, 2015. DOI: 10.1007/978-3-319-25010-6\_15.
- [31] G. Schreiber and Y. Raimond. RDF 1.1 Primer. W3C Working Group Note 24 June 2014, 2014. URL <https://www.w3.org/TR/rdf11-primer/>.
- [32] P. Vandenbussche, C. Buil-Aranda, A. Hogan, and J. Umbrich. Monitoring SPARQL endpoint status. In Blomqvist and Groza [9], pages 81–84. URL [http://ceur-ws.org/Vol-1035/iswc2013\\_demo\\_21.pdf](http://ceur-ws.org/Vol-1035/iswc2013_demo_21.pdf).
- [33] R. Verborgh, O. Hartig, B. D. Meester, G. Haesendonck, L. D. Vocht, M. V. Sande, R. Cyganiak, P. Colpaert, E. Mannens, and R. V. de Walle. Low-cost queryable linked data through triple pattern fragments. In M. Horridge, M. Rospocher, and J. van Ossenbruggen, editors, *Proceedings of the ISWC 2014 Posters & Demonstrations Track – a track within the 13th International Semantic Web Conference, ISWC 2014, Riva del Garda, Italy, October 21, 2014.*, volume 1272 of *CEUR Workshop Proceedings*, pages 13–16. CEUR-WS.org, 2014. URL [http://ceur-ws.org/Vol-1272/paper\\_10.pdf](http://ceur-ws.org/Vol-1272/paper_10.pdf).
- [34] R. Verborgh, M. V. Sande, P. Colpaert, S. Coppens, E. Mannens, and R. V. de Walle. Web-scale querying through Linked Data Fragments. In C. Bizer, T. Heath, S. Auer, and T. Berners-Lee, editors, *Proceedings of the Workshop on Linked Data on the Web (LDOW) co-located with the 23rd International World Wide Web Conference (WWW 2014), Seoul, Korea, April 8, 2014.*, volume 1184 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2014. URL [http://ceur-ws.org/Vol-1184/ldow2014\\_paper\\_04.pdf](http://ceur-ws.org/Vol-1184/ldow2014_paper_04.pdf).
- [35] G. T. Williams. SPARQL 1.1 Service Description. W3C Recommendation 21 March 2013, 2013. URL <https://www.w3.org/TR/sparql11-service-description/>.
- [36] A. Zaveri, A. Rula, A. Maurino, R. Pietrobbon, J. Lehmann, and S. Auer. Quality assessment for Linked Data: A survey. *Semantic Web*, 7(1):63–93, 2016. DOI: 10.3233/SW-150175.