

Performance Comparison of select and construct queries of triplestores

-

on the example of the JVMG project

Tobias Malmsheimer Stuttgart Media University

Japanese Visual Media Graph

- Research project at Stuttgart Media University
- Goal: Build a knowledge graph on the domain of Japanese visual media (manga, anime, computer games)
- Data source: Enthusiast communities on the web


Our Workflow

- Ingestion workflow
 - Retrieve data dump / load data from API in proprietary data format
 - Assess data structure and organisation and create an OWL ontology
 - Convert into RDF triples
 - Load into Fuseki database
- Integration workflow
 - Match semantically equivalent entities into clusters
 - Merge information from all sources (create a merged entity)
 - Match semantically equivalent entity properties
 - Merge property values (reduce redundancy)

Why we need a Database

- Need to store and access our knowledge graph
- Our frontend expects a sparql-endpoint
 - User can explore our knowledge graph
- Tiny use cases (TUCs)
 - Media researcher use the knowledge graph for their research questions
 - In this way, we can improve the knowledge graph and the supporting tools

vndb abstract data model



Source: <https://zenodo.org/record/5506936>


How does vndb look like

Type	Count
Character	90077
Trait	2237
Tag	2053
Producer	10394
Staff	21164
Visual Novel	28190
Visual Novel Release	71349

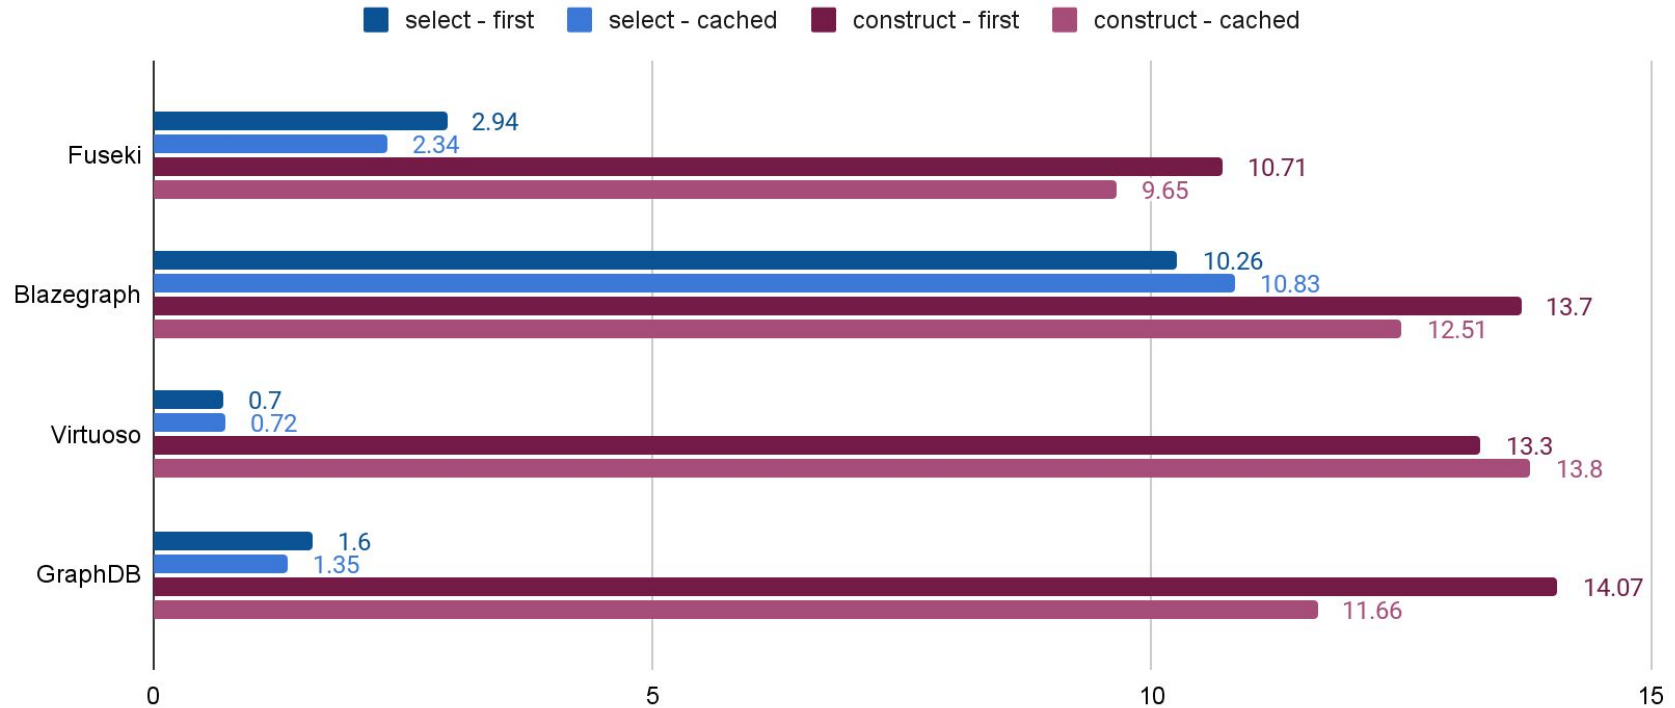
Disclaimer

- Just my opinion
- I did not optimize everything for each database
- Interested in the “out of the box” experience
- Therefore critique and bad numbers do not mean the database is bad

Small Query - First Evaluation



Medium Query - First Evaluation



Performance - Queries

- 2 different queries
 - Usual Queries we often use
- Small Query
 - Character page from Phorni
 - 96 triple
 - Small character page
- Medium Query
 - Overview of all vndb Characters
 - ~180k triple
 - Name of every entity which has type Character

Performance - CPU

- Python
 - Sparqlwrapper and rdflib for sparql endpoints
- Complete round trip
 - Executing sparql query
 - Serializing result
 - Average of 5 Queries, after warm-up queries
- Two kinds of queries
 - Select, assumption: better optimized
 - Construct
- Serialization methods
 - XML, JSON, CSV, JSONLD, Turtle, n-triples

Query - Select

```

1 PREFIX label: <http://www.w3.org/2000/01/rdf-schema#label>
2 PREFIX graph_label: <http://mediagraph.link/jvmg/ont/shortLabel>
3 SELECT
4   ?subject ?predicate ?object ?subjectLabel ?predicateLabel ?objectLabel ?graph ?graphLabel
5 WHERE {
6   {
7     GRAPH ?graph {
8       ?subject ?predicate ?object . FILTER ( ?subject = <http://mediagraph.link/vndb/character/1> )
9     }
10    OPTIONAL { ?graph      graph_label: ?graphLabel }
11    OPTIONAL { ?object     label:        ?objectLabel }
12    OPTIONAL { ?predicate  label:        ?predicateLabel }
13  } UNION {
14    GRAPH ?graph {
15      ?subject ?predicate ?object . FILTER ( ?object = <http://mediagraph.link/vndb/character/1> )
16    }
17    OPTIONAL { ?graph      graph_label: ?graphLabel }
18    OPTIONAL { ?subject    label:        ?subjectLabel }
19    OPTIONAL { ?predicate  label:        ?predicateLabel }
20  }
21 }

```


Query - Construct

```


1 PREFIX label: <http://www.w3.org/2000/01/rdf-schema#label>
2 PREFIX graph_label: <http://mediagraph.link/jvmg/ont/shortLabel>
3 CONSTRUCT {
4     ?subject ?predicate ?object .
5     ?graph graph_label: ?graphLabel .
6     ?object label: ?objectLabel .
7     ?subject label: ?subjectLabel .
8     ?predicate label: ?predicateLabel .
9 } WHERE {
10 {
11     GRAPH ?graph {
12         ?subject ?predicate ?object . FILTER ( ?subject = <http://mediagraph.link/vndb/character/1> )
13     }
14     OPTIONAL { ?graph graph_label: ?graphLabel }
15     OPTIONAL { ?object label: ?objectLabel }
16     OPTIONAL { ?predicate label: ?predicateLabel }
17 } UNION {
18     GRAPH ?graph {
19         ?subject ?predicate ?object . FILTER ( ?object = <http://mediagraph.link/vndb/character/1> )
20     }
21     OPTIONAL { ?graph graph_label: ?graphLabel }
22     OPTIONAL { ?subject label: ?subjectLabel }
23     OPTIONAL { ?predicate label: ?predicateLabel }
24 }
25 }

```


Fuseki - Small Query




Fuseki - Medium Query




Blazegraph - Small Query




Blazegraph - Medium Query




Virtuoso - Small Query




Virtuoso - Medium Query




GraphDB - Small Query




GraphDB - Medium Query



Small Query



Medium Query



Conclusion

- Select queries tends to be a bit faster
- Huge performance difference when using different serialization formats
 - But: not all serialization formats allow all features (quads in fuseki)
- Your mileage may vary: test combinations of query types and serialization formats

Questions?