Lazy Evaluation of SPARQL Queries with Caching

Robin Textor-Falconi

QLever & SPARQL

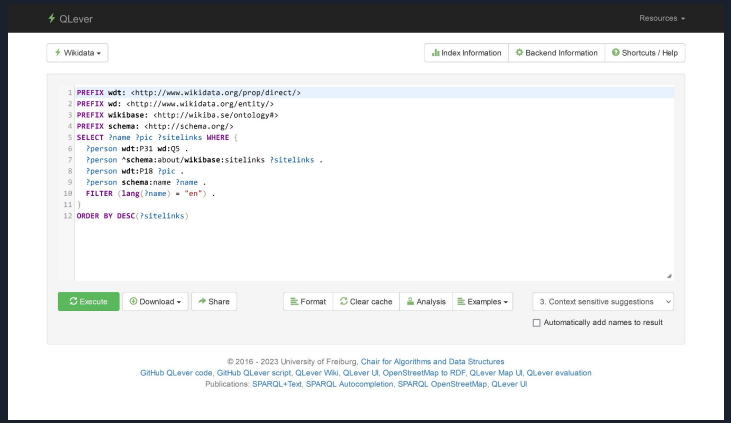

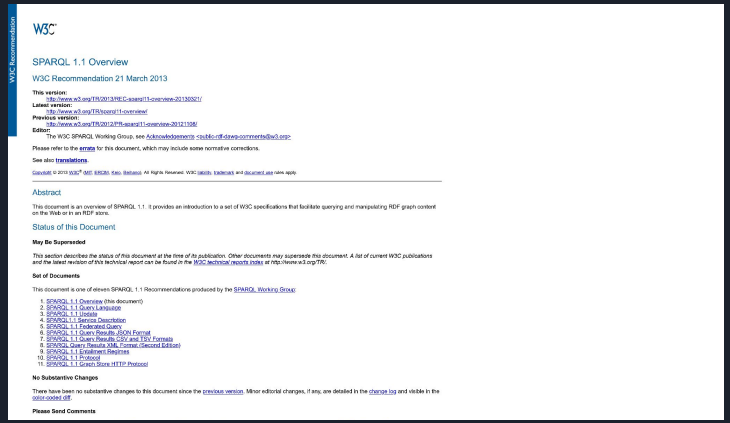

https://qlever.cs.uni-freiburg.de

https://www.w3.org/TR/2013/REC-sparql11-overview-20130321/

Problem Definition

Current Processing Model

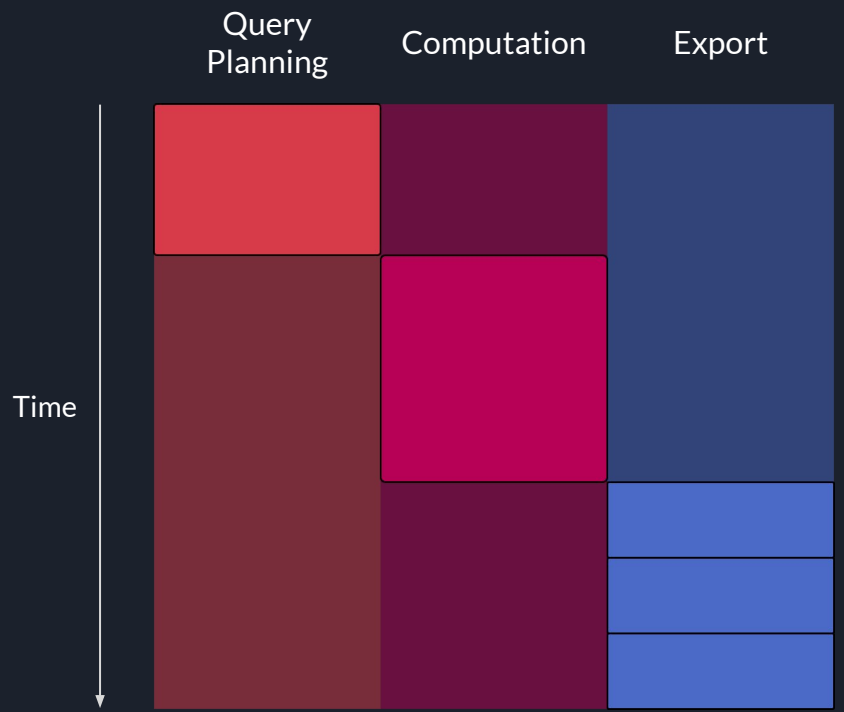

Practical Example

```
SELECT * WHERE { ?s ?p ?o }
```

INDEX SCAN ?s ?p ?o
Cols: ?o, ?p, ?s
Size: 20,052,950,074 x 3 [~ 20,052,968,255]
Time: 51,357ms [~ 20,052,968,255]

Practical Example

INDEX SCAN ?s ?p ?o
Cols: ?o, ?p, ?s
Size: 20,052,950,074 x 3 [~ 20,052,968,255]
Time: 51,357ms [~ 20,052,968,255]

$$
\begin{aligned}
& 3 \cdot 8 \cdot 20{,}052{,}950{,}074 \text{ bytes} \\
= \quad & 481{,}270{,}801{,}776 \text{ bytes} \\
\approx \quad & \textbf{481 GB !!}
\end{aligned}
$$

High memory requirement!

Existing Workaround

```
SELECT * WHERE { ?s ?p ?o } OFFSET           0 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET    10000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET    20000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET    30000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET    40000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET    50000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET    60000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET    70000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET    80000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET    90000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET   100000000 LIMIT 10000000
SELECT * WHERE { ?s ?p ?o } OFFSET   110000000 LIMIT 10000000
                                        …
SELECT * WHERE { ?s ?p ?o } OFFSET 20050000000 LIMIT 10000000
```

Solution

Proposed Processing Model

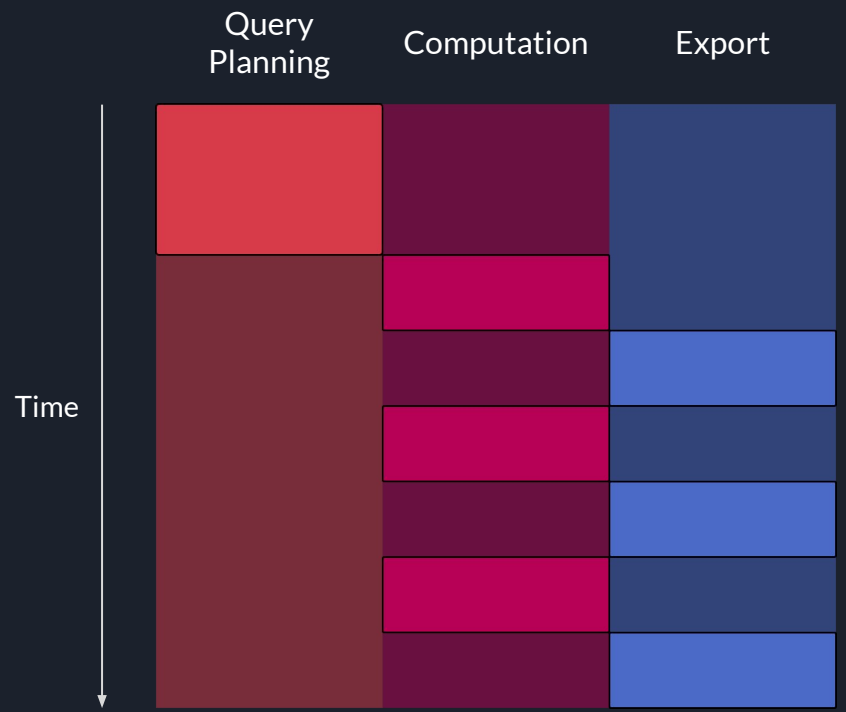

Code Examples in Different Languages

```cpp
#include <iostream>
#include <string>
#include <ranges>

int main() {
  std::ranges::istream_view<std::string> input{std::cin};
  for (const auto& line : input) {
    std::cout << line << std::endl;
  }
  return 0;
}
```

```java
import java.io.BufferedReader;
import java.io.InputStreamReader;
import java.io.IOException;
import java.nio.charset.StandardCharsets;
import java.util.stream.Stream;

public class Cat {
  public static void main(String[] args) throws IOException {
    try (var reader = new InputStreamReader(
          System.in, StandardCharsets.UTF_8);
         var bufferedReader = new BufferedReader(reader)) {
      bufferedReader.lines().forEach(System.out::println);
    }
  }
}
```

```cs
using System;

class Cat
{
  static void Main()
  {
    string? line;
    while ((line = Console.ReadLine()) is not null)
    {
      Console.WriteLine(line);
    }
  }
}
```

```hs
import System.IO (isEOF)

main :: IO ()
main = do
    eof <- isEOF
    if eof
        then return ()
        else do
            line <- getLine
            putStrLn line
            main
```

Challenges of the Thesis

- Implement general mechanism to existing codebase
  - Keep backwards compatibility
  - Keep code architecture largely untouched
- Keep Caching Mechanism
- Identify operations benefiting from new mechanism
- Add actual implementations for suited operations
  - Performance should ideally not regress by much

Demo

FILTER (?a = ?c)
Cols: ?c, ?b, ?a
Size: 17,716,426 x 3 [~ 20,052,968]
Time: 90,680ms [~ 20,073,021,223]

INDEX SCAN ?a ?b ?c
Cols: ?c, ?b, ?a
Size: 20,052,950,074 x 3 [~ 20,052,968,255]
Time: 1,415ms [~ 20,052,968,255]

qlever.cs.uni-freiburg.de/wikidata/CkGEt4

GROUP BY on ?a
Cols: ?a, ?count (U)
Size: 2,187,234,654 x 2 [~ 2,187,234,560]
Time: 574,101ms [~ 0]

INDEX SCAN ?a ?b ?c
Cols: ?a, ?b, ?c
Size: 20,052,950,073 x 3 [~ 20,052,968,255]
Time: 771ms [~ 20,052,968,255]

qlever.cs.uni-freiburg.de/wikidata/msUuJN

Demo

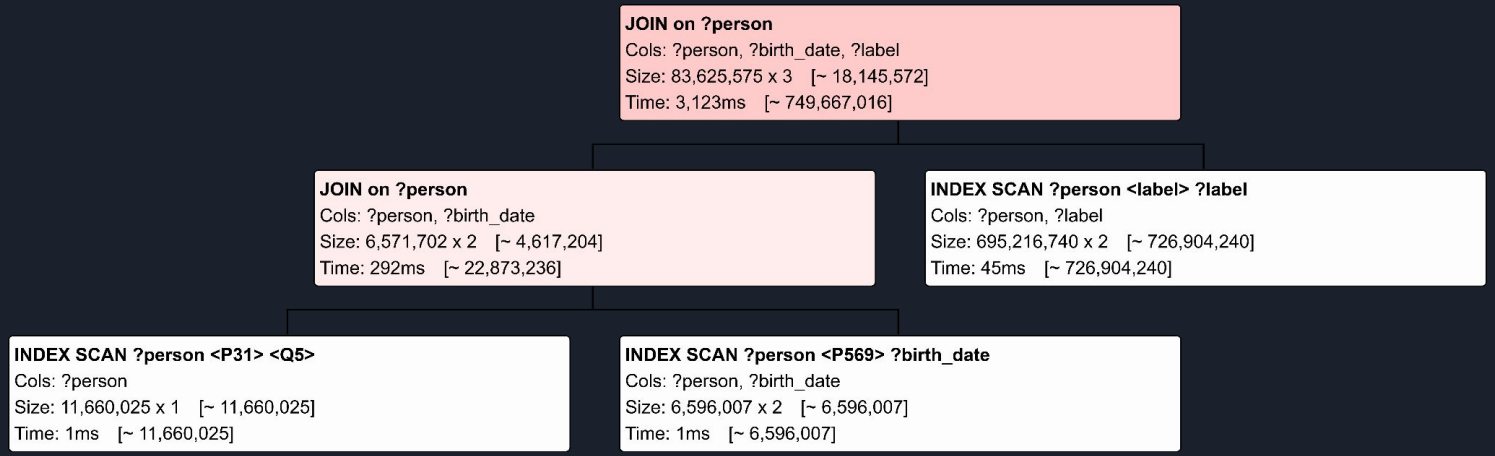

qlever.cs.uni-freiburg.de/wikidata/Zg778r

Performance Analysis

Laziness Performance Overhead

Query	Processing Time		Memory Delta	
	Lazy	Non-Lazy	Lazy	Non-Lazy
IndexScan example	20 s	-	702 MB	> 48 GB
Filter example	127 s	-	702 MB	> 48 GB
GroupBy example	608 s	-	702 MB	> 48 GB
Join example	3715 ms	4833 ms	702 MB	2151 MB
Bind example	447 min	-	1130 MB	> 48 GB
Union example	39 s	-	702 MB	> 48 GB
Distinct example	82 s	-	702 MB	> 48 GB
TransitivePath example	13 s	15 s	702 MB	2244 MB
CartesianProductJoin example	111 s	-	702 MB	> 48 GB
GroupBy example variant	385 ms	-	705 MB	> 48 GB
CartesianProductJoin example variant	144 ms	122 ms	702 MB	702 MB
People with pictures	2342 ms	2222 ms	851 MB	996 MB

Caching Overhead

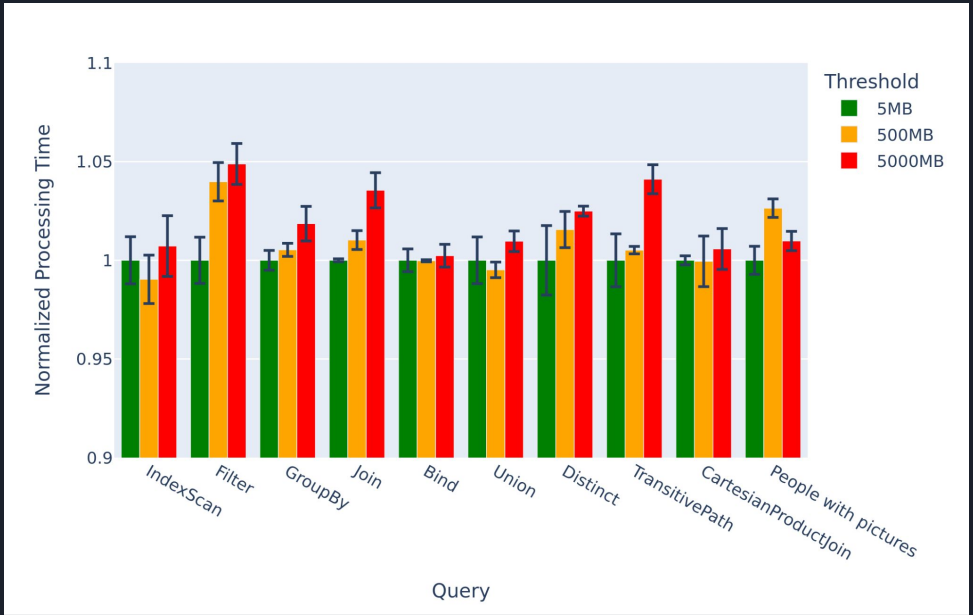

Chunk Size Impact

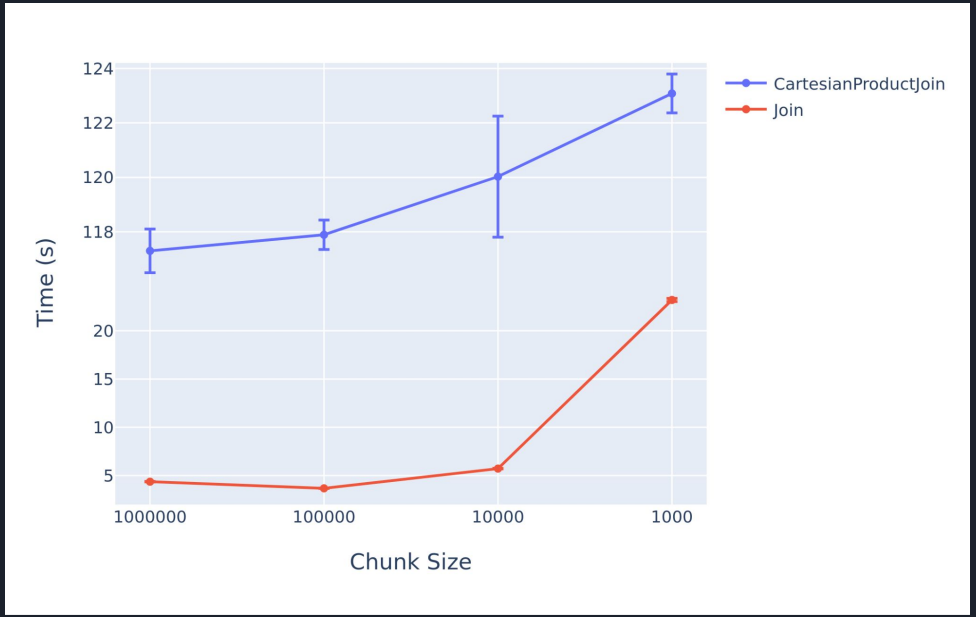

Lazy Evaluation of SPARQL Queries with Caching

Robin Textor-Falconi