

Joint thesis proposal at
MODALIS (I3S) and WIMMICS (INRIA / I3S) teams
Distributed Query Processing over a wide-area network

1 Advisors

- Johan Montagnat, DR CNRS, MODALIS research team leader
<http://www.i3s.unice.fr/~johan>
- Olivier Corby, CR INRIA, Wimmics research team
<http://www-sop.inria.fr/members/Olivier.Corby>

The candidate will be hired by the I3S laboratory.

2 Context

Scientific data volumes are growing at an unprecedented scale, data is often distributed over many acquisition sites, and there are a strong incentives to publish scientific data online for cross-analysis and secondary reuse. In the medical domain for instance, medical imaging devices acquire ever higher (spatial and temporal) resolution data, data is acquired and archived in plenty clinical and research centres, and huge international-scale data cohorts are assembled to address requirements of modern medical studies on diseases and treatment effects.

Exploiting the wealth of data available over the Internet becomes challenging considering the amount of data to process and the heterogeneity of data models and data sources used to represent and archive data. At the heart of this process, data search and federation engines are needed to align heterogeneous data models and federate distributed data sources, thus overcoming the partitioning of data over several locations and leveraging inlined data distributed over the wide-area network.

The Wimmics research team at INRIA has been studying knowledge graph representation models and semantic data search techniques over the past decade. It produces the CORESE/KGRAM search engine¹ [3, 4], which is compliant with the W3C standards (RDF, RDFS, SPARQL), making it well suited for querying data structured through rich and explicit data models. Furthermore, CORESE/KGRAM flexible software architecture makes it highly adaptable for multiple data sources querying and data models alignment. The MODALIS research team is specialised in distributed computing. It has been working in extending CORESE/KGRAM core with Distributed Query Processing (DQP) capabilities, in order to federate heterogeneous data sources distributed over multiple locations [7, 5]. Furthermore, MODALIS is implied in several medical data federation projects, including the CrEDIBLE initiative² [8] for biomedical data modelling and federation, and the GINSENG ANR project³ [2] for epidemiological studies.

This PhD work focuses on the development of performing and reliable query distribution techniques in a large-scale distribution context where network is not so reliable and exhibits varying performance. The primary target is the federation of medical data stores distributed at an international scale to help setting up multi-centric medical studies.

¹CORESE/KGRAM, <http://www-sop.inria.fr/edelweiss/software/corese/kgram/>

²CrEDIBLE: federating distributed and heterogeneous biomedical data and knowledge.
<http://credible.i3s.unice.fr/>

³GINSENG: Health Watch and Epidemiology Platform.
<https://ui.grid.creatis.insa-lyon.fr/fr/web/guest/home>

3 Goal and challenges

The goal of this PhD work is twofolds:

1. In the distributed computing domain: to optimise the parallel performance and reliability of the query engine.
2. In the e-health domain: to demonstrate the relevance of DQP techniques with realistic use cases and define a medical benchmark suitable to evaluate the work done.

The main challenge addressed is the optimisation of the query plan parallel processing. Theoretically distributed querying could lead to performance improvement as compared to a single integrated data repository, as query workload is distributed over several backends. In practice though, the need for joining results from several sources and the number of messages exchanged over the network to compute complete queries often has a severe impact on the overall query engine performance. Maximal parallelisation of the query plan processing is needed to compensate for this distribution overhead. In addition, targeting large-scale distribution yields additional performance limiting factors such as network variable performance and backends failure. Strategies to make the query processing resilient against such degraded conditions become mandatory.

In addition, evaluating the performance of a distributed query engine is challenging given the variable nature of a distributed infrastructure and the significant impact of the specific use cases addressed. A benchmark in the medical domain will be defined to measure the performance of the strategies and techniques developed. Alternative tools will be tested against this benchmark to make performance comparisons.

4 Workplan and expected results

The work will start with a bibliography study of distributed query processing techniques for semantic data store (*e.g.* FedX [11], ANAPSID [1], DARQ [9], SPLENDID [6]...). The survey will more particularly cover parallel query plan execution techniques and software architectures developed to improve parallelism. This will be compared to the actual query plan executor of KGRAM-DQP. An analysis of medical data stores federation requirements will be conducted to identify the most common query pattern required and the most stringent optimisation requirements implied. Large-scale data federation initiatives such as ADNI in the Alzheimer's domain ⁴, OFSEP in Multiple Sclerosis ⁵, BIRN ⁶ and CATI ⁷ in neurosciences or GINSENG in epidemiology [2] will be considered.

Techniques to optimise the query plan parallel execution will be developed. The most expected query plan patterns will be considered (in particular, query plans dealing with either horizontally or vertically partitioned data sets) and algorithms for parallel execution of these patterns will be developed. These algorithms will be integrated in KGRAM-DQP for experimentation.

Network reliability and variable performance will be studied. If needed, query plan reevaluation techniques at runtime will be considered to reconsider query processing order or optimisation strategies when some network connections underperform. Network connection reliability will also be dealt with, making it possible

⁴Alzheimer's Disease Neuroimaging Initiative: <http://www.adni-info.org>

⁵Observatoire Français de la Sclérose en Plaques <http://www.edmus.org/en/proj/observatoire.html>

⁶Biomedical Informatics Research Network <http://www.birncommunity.org/>

⁷Centre d'Acquisition et de Traitement d'Images <http://www.fondation-alzheimer.org/content/centre-dacquisition-et-de-traitement-automatis-de-limage>

to reconsider the data sources set if a connection failure prevent from getting results from one of the sources. The query result should then include incompleteness warning.

Finally, a distributed medical data query benchmark will be developed to assess the performance of the strategies and the parallel executor and compare it to existing query tools. This benchmark will be inspired from existing benches, such as FedBench [10], although strengthening the test with more medical cases. Realistic use cases and corresponding data stores will be taken from various projects in which MODALIS is involved (e.g. CrEDIBLE [8], GINSENG [2] or CATI/Pitié Salpêtrière hospital databases). Experiments will be designed and implemented on the Grid5000 national platform⁸ to ensure realistic execution conditions.

The final result of this work will be an improved KGRAM-DQP architecture, supporting parallel query plan execution through several optimisation strategies, and a methodology for assessing the performance of the query engine in the medical domain through a realistic benchmark. Secondary results are expected in the medical projects providing the test cases.

5 Relation to Labex UCN@Sophia

This thesis enters in two of the Labex themes:

- *Distributed and ubiquitous computing.* The Distributed Query Processing engine exploits parallel optimisation techniques and interacts with multiple data source components distributed over the network. Problems of network reliability and performance over a wide-area network for improving user experience when exploiting the data federation are explicitly tackled.
- *e-health.* Assumptions on the kind of data models and data distribution strategies may have a very significant impact on the system performance. This work is grounded into the medical domain for realistic use cases definition, validation, and reuse in medical-oriented project.

References

- [1] Maribel Acosta, Maria-Esther Vidal, Tomas Lampo, Julio Castillo, and Edna Ruckhaus. ANAPSID: An Adaptive Query Processing Engine for SPARQL Endpoints. In *International Semantic Web Conference (ISWC'11)*, volume LNCS 7031 of , pages 18–34, Bonn, Germany, October 2011. Springer.
- [2] Sébastien Cypièrre, Paul De Vlieger, David Sarramia, D. Hill, and Lydia Maigne. Development of a metamodel for medical database management on a grid network. Application to Health Watch and epidemiology for cancer and perinatal health. In *International Symposium on Cluster, Cloud and Grid Computing (CCGrid'12)*, , pages 892–897, Ottawa, Canada, May 2012.
- [3] Olivier Corby, Rose Dieng, Catherine Faron-Zucker, and Fabien Gandon. Searching the Semantic Web: Approximate Query Processing Based on Ontologies. *IEEE Intelligent System*, 21(1):20–27, January 2006.
- [4] Olivier Corby, Alban Gaignard, Catherine Faron-Zucker, and Johan Montagnat. KGRAM Versatile Inference and Query Engine for the Web of Linked Data. In *IEEE/WIC/ACM International Conference on Web Intelligence (WI'12)*, , Macao, China, December 2012.

⁸Grid5000, <https://www.grid5000.fr>

- [5] Alban Gaignard, Johan Montagnat, Catherine Faron-Zucker, and Olivier Corby. Semantic Federation of Distributed Neurodata. In *MICCAI Workshop on Data- and Compute-Intensive Clinical and Translational Imaging Applications(DCICTIA-MICCAI 2012)*, , pages 41–50, Nice, France, October 2012.
- [6] Olaf Görlitz and Steffen Staab. SPLENDID: SPARQL Endpoint Federation Exploiting VOID Descriptions. In *2nd International Workshop on Consuming Linked Data(COLD'11)*, , Bonn, Germany, October 2011.
- [7] Franck Michel, Alban Gaignard, Farooq Ahmad, Christian Barillot, Bénédicte Batrancourt, Michel Dojat, Bernard Gibaud, Pascal Girard, David Godard, Gilles Kassel, Diane Lingrand, Grégoire Malandain, Johan Montagnat, M. Péligrini-Issac, Xavier Pennec, Javier Rojas Balderrama, and Bacem Wali. Grid-wide neuroimaging data federation in the context of the NeuroLOG project. In *HealthGrid'10*, , pages 112–123, Paris, France, June 2010. IOS Press.
- [8] Corby Olivier, Faron Zucker Catherine, Gaignard Alban, Gibaud Bernard, Kassel Gilles, Michel Franck, and Johan Montagnat. CrEDIBLE Multi-disciplinary workshop. Technical Report CrEDIBLE-13-3-v1, I3S, Sophia Antipolis, France, December 2013.
- [9] Bastian Quilitz and Ulf Leser. *Querying Distributed RDF Data Sources with SPARQL*, pages 524–538. . Springer, 2008.
- [10] Michael Schmidt, Olaf Görlitz, Peter Haase, Günter Ladwig, Andreas Schwarte, and Than Tran. Fedbench: A benchmark suite for federated semantic data query processing. In *International Semantic Web Conference(ISWC'11)*, volume LNCS 7031 of , pages 585–600, Bonn, Germany, October 2011. Springer.
- [11] Andreas Schwarte, Peter Haase, Katja Hose, Ralf Schenkel, and Michael Schmidt. Fedx: optimization techniques for federated query processing on linked data. In *international Semantic Web conference(ISWC'11)*, , pages 601–616, Bonn, Germany, October 2011.