Distributed storage and indexing for large-scale graph analytics

Angelos-Christos Anadiotis, Ioana Manolescu

May 19, 2021

1 Context

Graphs have been widely used for modelling connections between entities. Accordingly, graph-based analytics are used in several contexts, including detection of fraud and fake news as well as social network analysis. Regardless of the application, one of the major challenges in graph analytics is to discover paths of *reasonable* length and accuracy, which connect different entities.

However, the efficiency of graph analysis algorithms is affected not only by the complexity of the search space, but also by the shape of the underlying data. Specifically, graphs are expected to work on top of heterogeneous data sources, which are stored (a) in different formats, ranging from plain text and multimedia to fully structured data, and, (b) in different locations, where access might be constrained by a network connection. Compared to standard analytical processing technologies, graph analytics require to maintain both the underlying data sources and the graph built on top of them, whereas the organization of the data, naturally affects the performance of graph processing [5, 6].

This PhD project is part of the SourcesSay AI chair, a four-year project sponsored by the ANR (Agence nationale de la recherche) and DGA (Direction Générale de l'Armement), starting in 2020 within the CEDAR team joint between Inria and Ecole Polytechnique[1]. The project aims at storing, analyzing, enriching and integrated datasets of various data models under a

[1] `https://sourcessay.inria.fr`

graph model; it builds upon ideas and software previously developed within the ConnectionLens [4, 7] research project[2].

2 Objectives

The main goal of this research project is to provide a set of scalable storage and indexing technologies, which should enable powerful analytics operations on top of graphs built by ingesting heterogeneous data sources. The following objectives should be attained:

1. Propose a database design which should efficiently support graph-based queries.

2. Enable both in-memory and disk-based data storage and retrieval, in particular developing a buffer pool mechanism which should fit the requirements of graph-based queries.

3. Propose data compression and indexing technologies which improve the performance of graph queries, given the restrictions on the different data layouts and storage locations.

4. Enhance existing query planning and optimization technologies by extending them to support operators for distributed graph processing over heterogeneous data; these techniques should in particular cope with the very heterogeneous nature of the graphs built in the project.

The project will rely on recent developments of the CEDAR team in the context of scalable graph querying [2, 3].

3 Approach

The fundamental idea of this PhD thesis project is to create data processing pipelines across different settings (scale-up, scale-out, in-memory, hybrid, etc.) which maximize the overlap between processing and data manipulations, like transformations and transfers. Given that the graph co-exists with the underlying heterogeneous data, techniques that have been applied in modern online analytical engines, like column-oriented storage layout [8]

[2] https://team.inria.fr/cedar/connectionlens/

and late materialization [1, 12] as well as their combination [9] will be investigated. Further, state-of-the-art indexing techniques like [10] and [11] will also be considered.

This project will innovate in both the query optimization and the query execution level, while taking advantage of the hardware available in modern servers.

Accordingly, the successful candidate is expected to develop hardware-aware solutions, leveraging the intrinsic functionality provided by the underlying hardware. The software developed in the context of this project is expected to span from the operating system to the applications facing the end user.

References

[1] D. J. Abadi, D. S. Myers, D. J. DeWitt, and S. R. Madden. Materialization strategies in a column-oriented DBMS. In *2007 IEEE 23rd International Conference on Data Engineering*, pages 466–475, 2007.

[2] A. Anadiotis, M. Y. Haddad, and I. Manolescu. Graph-based keyword search in heterogeneous data sources. `https://arxiv.org/abs/2009.04283`, 2021.

[3] A.-C. Anadiotis, O. Balalau, T. Bouganim, F. Chimienti, H. Galhardas, M. Y. Haddad, S. Horel, I. Manolescu, and Y. Youssef. Empowering investigative journalism with graph-based heterogeneous data management. `https://arxiv.org/abs/2102.04141`, 2021.

[4] A.-C. Anadiotis, O. Balalau, C. Conceicao, H. Galhardas, M. Y. Haddad, I. Manolescu, T. Merabti, and J. You. Graph integration of structured, semistructured and unstructured data for data journalism. `https://arxiv.org/abs/2012.08830`, 2020.

[5] R. Angles, P. A. Boncz, J. Larriba-Pey, I. Fundulaki, T. Neumann, O. Erling, P. Neubauer, N. Martínez-Bazan, V. Kotsev, and I. Toma. The linked data benchmark council: a graph and RDF industry benchmarking effort. *SIGMOD Rec.*, 43(1):27–31, 2014.

[6] A. Bonifati, G. H. L. Fletcher, H. Voigt, and N. Yakovets. *Querying Graphs*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2018.

[7] C. Chanial, R. Dziri, H. Galhardas, J. Leblay, M. L. Nguyen, and I. Manolescu. Connectionlens: Finding connections across heterogeneous data sources. *Proc. VLDB Endow.*, 11(12):2030–2033, 2018.

[8] M. DiScala and D. J. Abadi. Automatic generation of normalized relational schemas from nested key-value data. In *Proceedings of the 2016 International Conference on Management of Data*, SIGMOD '16, page 295–310, New York, NY, USA, 2016. Association for Computing Machinery.

[9] S. Idreos, S. Manegold, H. Kuno, and G. Graefe. Merging what's cracked, cracking what's merged: Adaptive indexing in main-memory column-stores. *Proc. VLDB Endow.*, 4(9):586–597, June 2011.

[10] V. Leis, A. Kemper, and T. Neumann. The adaptive radix tree: Artful indexing for main-memory databases. In *Proceedings of the 2013 IEEE International Conference on Data Engineering (ICDE 2013)*, ICDE '13, page 38–49, USA, 2013. IEEE Computer Society.

[11] S. Luo, S. Chatterjee, R. Ketsetsidis, N. Dayan, W. Qin, and S. Idreos. Rosetta: A robust space-time optimized range filter for key-value stores. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, SIGMOD '20, page 2071–2086, New York, NY, USA, 2020. Association for Computing Machinery.

[12] L. Shrinivas, S. Bodagala, R. Varadarajan, A. Cary, V. Bharathan, and C. Bear. Materialization strategies in the vertica analytic database: Lessons learned. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*, pages 1196–1207, 2013.