Network-aware Optimization for Query Executions in Large Systems

Long Cheng
Department of Mathematics and Computer Science
Eindhoven University of Technology
L.Cheng@tue.nl

An increasing number of companies rely on the results of big data analytics to improve their operations, customer service and risk management. To cope with the added weight of Big Data, high-performance data management/analytics over large systems, such as datacenters equipped with thousands of servers, is becoming the mainstream in today's business world. As one of the key tasks in such scenarios, efficient execution of query operations is still challenging current techniques. One of the main performance challenges is network communications: expensive queries could spend more than half of their completion time on data transferring [4], and data center systems could consume more than one-third of their total energy on communications [3].

Over the past years, significant performance improvements on query executions have been achieved by using state-of-the-art methods, such as reducing network traffic designed in the data management domain [4], and data flow scheduling in the data communications domain [2]. However, almost all the techniques in the two fields just view each other as a black box, and the additional performance gains from a co-optimization perspective have not yet been explored. Moreover, though energy-efficient solutions, e.g., dynamic routing, are being studied in the network management domain, few of them have ever considered to improve energy issues from a data engineering angle.

In this presentation, I will talk about the general design of a novel query execution system NEO-QE, i.e., **NE**twork-aware **O**ptimization for **Q**uery **E**xecutions in Large Systems. Firstly, I will use a simple example to show the motivation of the design. Then, I will talk about the logical architecture of the system to demonstrate that how can we bridge the gap of optimization between high-level query executions and underlying data communications. Finally, I will show some preliminary results of the design [1]. In general, it is expected that the proposed system will be always highly efficient and robust in the presence of different workloads and networks in large-scale distributed scenarios.

References

[1] L. Cheng, Y. Wang, Y. Pei, and D. Epema. A coflow-based co-optimization framework for high-performance data analytics. In *ICPP*, 2017.

[2] M. Chowdhury, Y. Zhong, and I. Stoica. Efficient coflow scheduling with Varys. In *SIGCOMM*, 2014.

[3] A. Greenberg, J. R. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta. VL2: a scalable and flexible data center network. *Comm. ACM*, 2011.

[4] O. Polychroniou, R. Sen, and K. A. Ross. Track join: distributed joins with minimal network traffic. In *SIGMOD*, 2014.