Low Latency Storage for Stream Data

Keywords: Big Data, Streaming, Storage

An increasing number of Big Data applications deal with small data (in the orders of Bytes or KiloBytes). This trend is easily observed in domains like finance, weather forecast, IoT, insurance or social networks. In such contexts, small items are continuously collected from the stream sources or are received from other stream processing computations. Even if the stream engines running the applications are processing such stream data on the fly, by passing it through the topology of stream operators, there is an increasing need to be able to store such items efficiently. Unlike traditional storage, the main challenge of storing stream data is the large number of small items (arriving at rates easily reaching tens of millions of parameters per second).

There is clear need for a dedicated solution for low latency stream storage. Such a solution should provide on the one hand traditional storage functionality and on the other hand stream-like performance (i.e., low latency I/O access to items/range of items). The goal of this internship is to explore the plausible paths towards such a dedicated storage solution: the main requirements and challenges, the design choices (e.g., a standalone component vs. an extension of an existing Big Data solution like HDFS) an architectural overview and experimental validation on the Grid5000 platform as well as public clouds (Microsoft Azure, Amazon WS).

To achieve this we plan to use several design principles: data partitioning schemes and the ability to deal with billions of small items, inspired by the Kafka approach for streaming; techniques for managing data in distributed caches that enable ns access time across large collections of items, as the ones introduced in RamCloud and DXRam; dynamic metadata partitioning for increased scalability and support for high concurrency, as used in Ceph, Giraffa and CalvinFS; enhanced I/O through microbatches for reading/writing from/to HDFS, inspired by the Spark approach for data processing.

References:

[1] Kreps Jay, Narkhede Neha, and Rao Jun. Kafka: A distributed messaging system for log processing. In Proceedings of 6th International Workshop on Networking Meets Databases, NetDB'11, 2011.
[2] RamCloud http://web.stanford.edu/~ouster/cgi-bin/projects.php
[3] Klein Florian and Schottner Michael. Dxram: A persistent in-memory storage for billions of small objects, PDCAT, IEEE, December 2013.
[4] Ceph http://ceph.com
[5] Giraffa https://github.com/GiraffaFS/giraffa/wiki
[6] Alexander Thomson, Thaddeus Diamond, Shu-Chun Weng, Kun Ren, Philip Shao, and Daniel J. Abadi. 2012. Calvin: fast distributed transactions for partitioned database systems. In Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data (SIGMOD '12).
[7] Apache Spark http://spark.apache.org
[8] Irina Botan, Gustavo Alonso, Peter M. Fischer, Donald Kossmann, and Nesime Tatbul. Flexible and scalable storage management for data-intensive stream processing. In

Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology, EDBT '09, pages 934–945
[9] Fangjin Yang, Eric Tschetter, Xavier L´eaut´e, Nelson Ray, Gian Merlino, and Deep Ganguli. Druid: A real-time analytical data store. In Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data, SIGMOD '14, pages 157–168
[10] Grid5000 www.grid.5000.fr

Main advisor:
Ovidiu Cristian Marcu (Inria) ovidiu-cristian.marcu@inria.fr

Other advisors:
Alexandru Costan (IRISA) Alexandru.costan@irisa.fr
Gabriel Antoniu (Inria) Gabriel.antoniu@inria.fr
Radu Tudoran (Huawei Research) radu.tudoran@huawei.com
Bogdan Nicolae (Argonne National Laboratory) bogdan.nicolae@acm.org

Required skills:
• Strong knowledge of computer networks and distributed systems
• Knowledge on storage and (distributed) file systems
• Strong programming skills (e.g. C/C++, Java)
• Working experience in the areas of Big Data management, Cloud computing, HPC, is an advantage
• Very good communication skills in oral and written English.

To apply, please email a CV to Dr. Gabriel Antoniu and Dr. Alexandru Costan.