Robust Query Processing in Co-Processor-accelerated Databases

Sebastian Breß[1,2⋆], Henning Funke[2], Jens Teubner[2], Volker Markl[1,3]

DFKI GmbH, Intelligente Analytik für Massendaten, Alt-Moabit 91c, 10559 Berlin[1]
Technische Universität Dortmund, FG DBIS, Otto-Hahn-Straße 14, 44227 Dortmund[2]
Technische Universität Berlin, FG DIMA, Einsteinufer 17, 10587 Berlin[3]
sebastian.bress@dfki.de, henning.funke@tu-dormund.de,
jens.teubner@tu-dormund.de, volker.markl@tu-berlin.de

Abstract. Technology limitations are making the use of *heterogeneous computing devices* much more than an academic curiosity. In fact, the use of such devices is widely acknowledged to be the only promising way to achieve application-speedups that users urgently need and expect. However, building a robust and efficient query engine for heterogeneous co-processor environments is still a significant challenge.

In our latest work [1], we identify two effects that limit performance in case co-processor resources become scarce. *Cache thrashing* occurs when the working set of queries does not fit into the co-processor's data cache, resulting in performance degradations up to a factor of 24. *Heap contention* occurs when multiple operators run in parallel on a co-processor and when their accumulated memory footprint exceeds the main memory capacity of the co-processor, slowing down query execution by up to a factor of six.

We propose solutions for both effects. *Data-driven operator placement* avoids data movements when they might be harmful; *query chopping* limits co-processor memory usage and thus avoids contention. The combined approach—*data-driven query chopping*—achieves robust and scalable performance on co-processors. We validate our proposal with our open-source GPU-accelerated database engine CoGaDB and the popular star schema and TPC-H benchmarks.

Acknowledgments. The work has received funding from the Deutsche Forschungsgemeinschaft (DFG), Collaborative Research Center SFB 876, project C5, from the European Union's Horizon2020 Research & Innovation Program under grant agreement 671500 (project "SAGE"), and by the German Ministry for Education and Research as Berlin Big Data Center BBDC (funding mark 01IS14013A).

References

1. S. Breß, H. Funke, and J. Teubner. Robust query processing in co-processor-accelerated databases. In *Proceedings of the International Conference on Management of Data (SIGMOD)*, pages 1891–1906. ACM, 2016.

⋆ Work done when author was working at TU Dortmund.