

The End of the x86 Dominance in Databases?

Huanchen Zhang
Carnegie Mellon University
huanche1@cs.cmu.edu

The explosion of “big-data” analytics imposes great challenges on traditional relational databases that are optimized for the CPU architecture. CPUs follow the von Neumann model, which is a control-flow architecture, and do not perform well on complex analytical queries backed by statistics and machine learning that involve massive data and parallelism. Therefore, emerging databases often offload the data-intensive part of the queries to hardware accelerators to achieve better performance and energy efficiency. Common hardware accelerators include GPUs and FPGAs. GPUs are highly-parallel and are efficient in performing sorting and sequential scans with predicates. They have been used in commercial products such as MapD (now Omnicore) [3] and Kinetica [2]. FPGAs are reconfigurable hardware and can be used to accelerate a variety of compute-intensive tasks such as data compression/decompression, predicate evaluation [6] and pattern matching [5].

The problem with today’s hardware accelerators, however, is that they are ad-hoc solutions for specific compute/data-intensive queries. Although FPGAs are programmable, reconfiguration including compilation, synthesis, and routing usually takes hours [6]. Therefore, applications must know the query patterns in advance so that they can prepare the hardware for the tasks.

This problem might get solved by Intel’s new hardware design called the Configurable Spatial Accelerator (CSA) [1]. CSA is a dataflow architecture designed for the high-performance computing to assist or even replace the traditional superscalar out-of-order CPUs in supercomputers. Figure 1 shows the microarchitecture of CSA. CSA contains dense arrays of heterogeneous processing elements (PEs) that can handle integer and floating-point arithmetic, as well as communication routing and in-fabric storage. They are connected using the on-chip configurable network to form a desirable dataflow. PE executions are asynchronous. A PE is executed once the inputs are ready, and the outputs are immediately forwarded to the downstream PE. Intel claims that CSA can provide an order-of-magnitude improvement in performance and energy efficiency.

But what is really amazing about this hardware are 1) it can directly map and execute the dataflow graphs generated by compilers (e.g., C++ compilers), and 2) it only takes several hundreds of nanoseconds to a few microseconds to reconfigure the PEs. Although this new hardware just got patented and is still under development (we have no idea when it will be available), it is not too early to think about how these new features will dramatically change the way reconfigurable hardware is used in a database. First, because of the short reconfiguration time, databases can determine

This article is published under a Creative Commons Attribution License (<http://creativecommons.org/licenses/by/3.0/>), which permits distribution and reproduction in any medium as well as allowing derivative works, provided that you attribute the original work to the author(s) and CIDR 2019, 9th Biennial Conference on Innovative Data Systems Research (CIDR '19), January 13-16, 2019, Asilomar, California, USA.

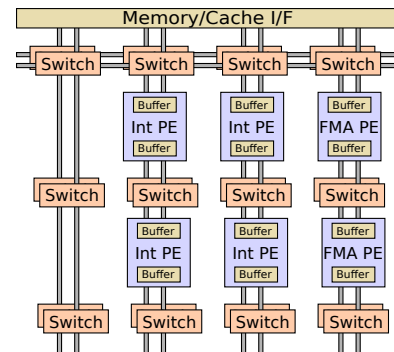


Figure 1: CSA Microarchitecture [1].

what tasks to accelerate on-the-fly. That means query optimizers will be involved (unlike today’s FPGA-accelerated databases that use UDFs to bypass query optimizers [4, 5]) with new cost models to make decisions in a heterogeneous environment. Second, because CSA directly executes the compiler-generated graphs, it can easily benefit from technological advances in query compilers (e.g., codegen, JIT). Mapping (partial) dataflow graphs to the silicon is claimed to be straightforward with one-to-one correspondence between the nodes in the graph and the PEs on the board.

One step further, rumor says that Intel is treating CSA as a processor or coprocessor rather than an accelerator because CSA is a full dataflow engine that can directly execute the graphs created by programs. It is probable that Intel is going to put multiple CSAs along with one or multiple CPUs on the same Xeon-like die. If this happens, CSAs will not be bottlenecked by the PCIe bandwidth (as is the case for FPGAs) and will have cache-coherence. This allows CSAs to efficiently handle a wider range of queries, including those in OLTP applications. For example, we can map important indexes to CSAs by hard-coding branching keys and the comparison logic in the PEs so that index lookups can be performed at bare-metal speed. Under this architecture, most database queries will be executed by CSAs with CPUs acting as coordinators. We do not know when the x86 era will end, but the CSA might challenge its dominance in databases soon.

REFERENCES

- [1] 2018. Configurable Spatial Accelerator (CSA) - Intel. https://en.wikichip.org/wiki/intel/configurable_spatial_accelerator. (2018).
- [2] 2018. Kinetica. <https://www.kinetica.com/>. (2018).
- [3] 2018. Omnicore. <https://www.omnicore.com/>. (2018).
- [4] Muhsen Owaid, David Sidler, Kaan Kara, and Gustavo Alonso. 2017. Centaur: A framework for hybrid CPU-FPGA databases. In *Field-Programmable Custom Computing Machines (FCCM), 2017 IEEE 25th Annual International Symposium on*. IEEE, 211–218.
- [5] David Sidler, Zsolt István, Muhsen Owaid, and Gustavo Alonso. 2017. Accelerating pattern matching queries in hybrid CPU-FPGA architectures. In *Proceedings of the 2017 ACM International Conference on Management of Data*. ACM, 403–415.
- [6] Bharat Sukhwani, Hong Min, Mathew Thoenes, Parijat Dube, Balakrishna Iyer, Bernard Brezzo, Donna Dillenberger, and Sameh Asaad. 2012. Database analytics acceleration using FPGAs. In *Proceedings of the 21st international conference on Parallel architectures and compilation techniques*. ACM, 411–420.