

Packing a Punch: Building Intelligent Distributed Storage with Hardware

Zsolt István

ETH Zurich

**Monday, 03.07.2017, 13:00 – 14:30, TU Dresden, Fakultät Informatik
Nöthnitzer Str. 46, Raum APB-3105**

Abstract:

Storage in data centers is evolving very quickly: boundaries between main-memory and persistent storage are becoming blurred, persistent storage no longer has a block-based interface, it can provide random access at high rates, and it is often exposed over the network. Yet, the ever increasing amount of data being handled causes an intrinsic

inefficiency: moving data around is expensive in terms of bandwidth, latency, and power consumption. One solution to the data movement bottleneck is bringing computation closer to the source, i.e., offloading part of the processing directly to the storage nodes. This talk presents Caribou, a distributed storage layer for near-data processing in database engines built using FPGAs. Caribou exposes DRAM through a simple key-value store interface and each storage node provides high-bandwidth near-data processing (e.g., predicate evaluation and pattern matching) and fault tolerance through replication. The result is an efficient intelligent distributed data store that can be used not only to boost performance but also to reduce power consumption and real estate usage in the data center thanks to the micro-server architecture adopted. Caribou can also act as an exploration platform for developing novel near-data processing techniques for the larger area of data sciences. Thanks to its modular design, its functionality can be easily extended and the existing management and processing modules can be tailored to the applications' needs independently.

Bio:

I am currently studying as a PhD student at the Systems Group at ETH Zurich under the supervision of Prof. Dr. Gustavo Alonso. I have a Master's degree in Distributed Systems from ETH Zurich (2013), and an Engineering degree in Computer Science from the Technical University of Cluj-Napoca (2011).

