Speeding up Big Data Operations with Data Center Memory Disaggregation

Robin Abrahamse, Zaid Al-Ars & Akos Hadnagy

Accelerated Big Data Systems, University of Technology Delft

Introduction

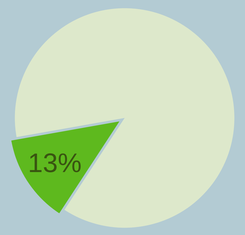

Up to 13% of global electricity consumed by data centers in 2030 [1]

⬇

Improving data center efficiency matters!

Data center workloads are often hindered by limited available memory and therefore require performance-leaking scalability techniques

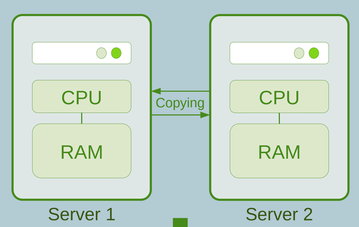

Traditional scale-out approach involves copying large amounts of data over the network

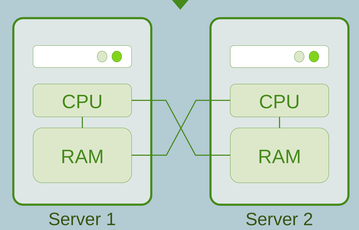

Memory disaggregation allows servers to directly access memory of adjacent servers through dedicated hardware [2]

Potential to reduce network load and improve overall data center efficiency and cost!

Memory disaggregation will be incorporated in the next generation data center processors (such as POWER10)

The goal is to investigate the potential for memory disaggregation to assist in speeding up big data operations (in particular ‘wide-dependency’ operations like shuffles)

Memory disaggregation is integrated in the commonly used Plasma in-memory object store to allow efficient memory access across multiple compute nodes [3][4]

System Proposal

Plasma framework can commit and seal objects to make them immutable

Sealed objects are made available for access by Plasma clients via buffers

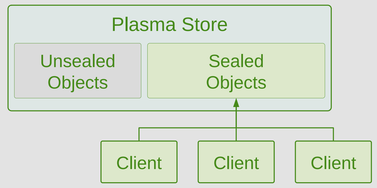

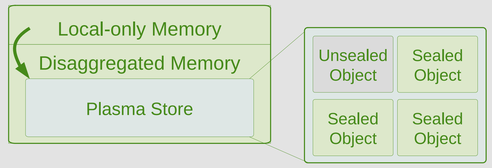

Step 1
Plasma object allocation in disaggregated memory region for access by remote servers

Step 2
Sharing Plasma object buffers with remote Plasma stores via RPC upon request by Plasma clients

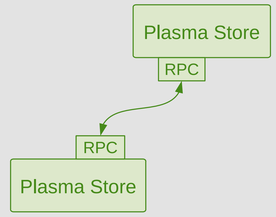

Performance microbenchmarks of memory-fetching

Discussion

Memory disaggregated enables a new paradigm for dynamically pooling processing and memory resources

Memory disaggregated Plasma provides an open framework for distributed data management in big data analysis

Current memory disaggregation technology is limited due to e.g. cache coherency and latency characteristics

The proposed framework lays a stepping stone for future work; lots of opportunities in research and industry

Conclusions

Memory disaggregation is rapidly becoming more relevant

A new paradigm in distributed computing to leverage high volume, high performance data management

Large potential for memory disaggregation to transform big data analysis

Memory disaggregated Plasma allows faster big data analysis solutions by leveraging remote compute node memory

Proposed framework enables new research, more work on the way!

References

[1] Andrae, A. S., & Edler, T. (2015). On global electricity usage of communication technology: trends to 2030.

[2] Pinto, C. et al. (2020). ThymesisFlow: A Software-Defined, HW/SW co-Designed Interconnect Stack for Rack-Scale Memory Disaggregation.

[3] Apache Software Foundation. (2016). Arrow: A cross-language development platform for in-memory analytics. https://arrow.apache.org

[4] Ahmad, T. et al. (2020). ArrowSAM: In-Memory Genomics Data Processing Using Apache Arrow.

More Information