A Low-Latency Multi-Version Key-Value Store Using B-tree on an FPGA-CPU Platform

Yuchen Ren, Jinyu Xie, Yunhui Qiu, Hankun Lv, Wenbo Yin, Lingli Wang
State Key Laboratory of ASIC and System, Fudan University

Bowei Yu, Hua Chen, Xianjun He, Zhijian Liao, Xiaozhong Shi
IT R&D Dept., Chengdu Research Institute, Huawei Technologies Co., Ltd.

FPL'19, Barcelona, September 11th, 2019

Introduction - Background

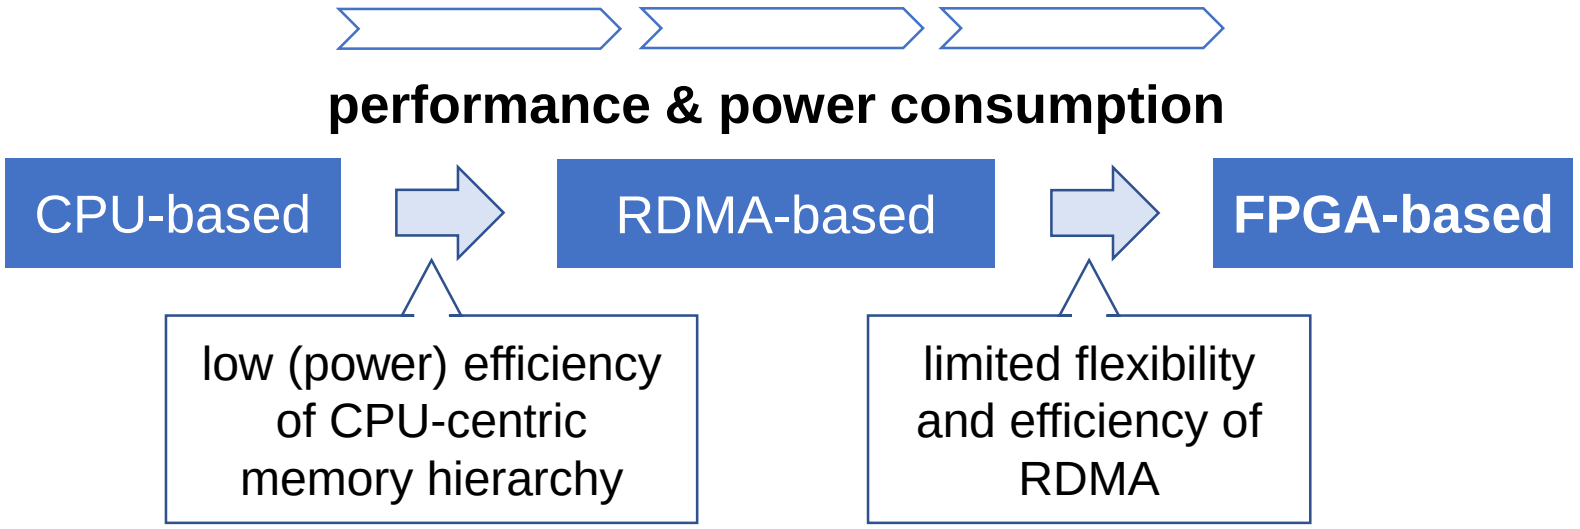

**RDMA: Remote Direct Memory Access*

Multi-Version KVS

(Key-Value Store)

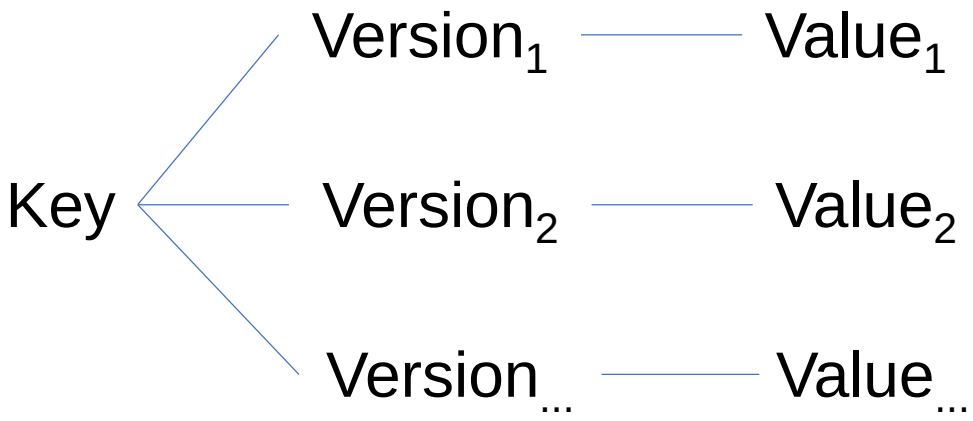

Introduction - Contribution

Design

- a low-latency multi-version in-memory KVS
- FPGA-CPU heterogeneous architecture

Storage

- keys – hash table – FPGA board (Cuckoo hashing)
- version-value pairs (VVPs) – B-trees – host memory

Operation

- *get*, *put*, *delete*, *CAS*, *getPredecessor* – bypassing the CPU
- range query – with the help of the CPU

**CAS: Compare and Swap*

Architecture

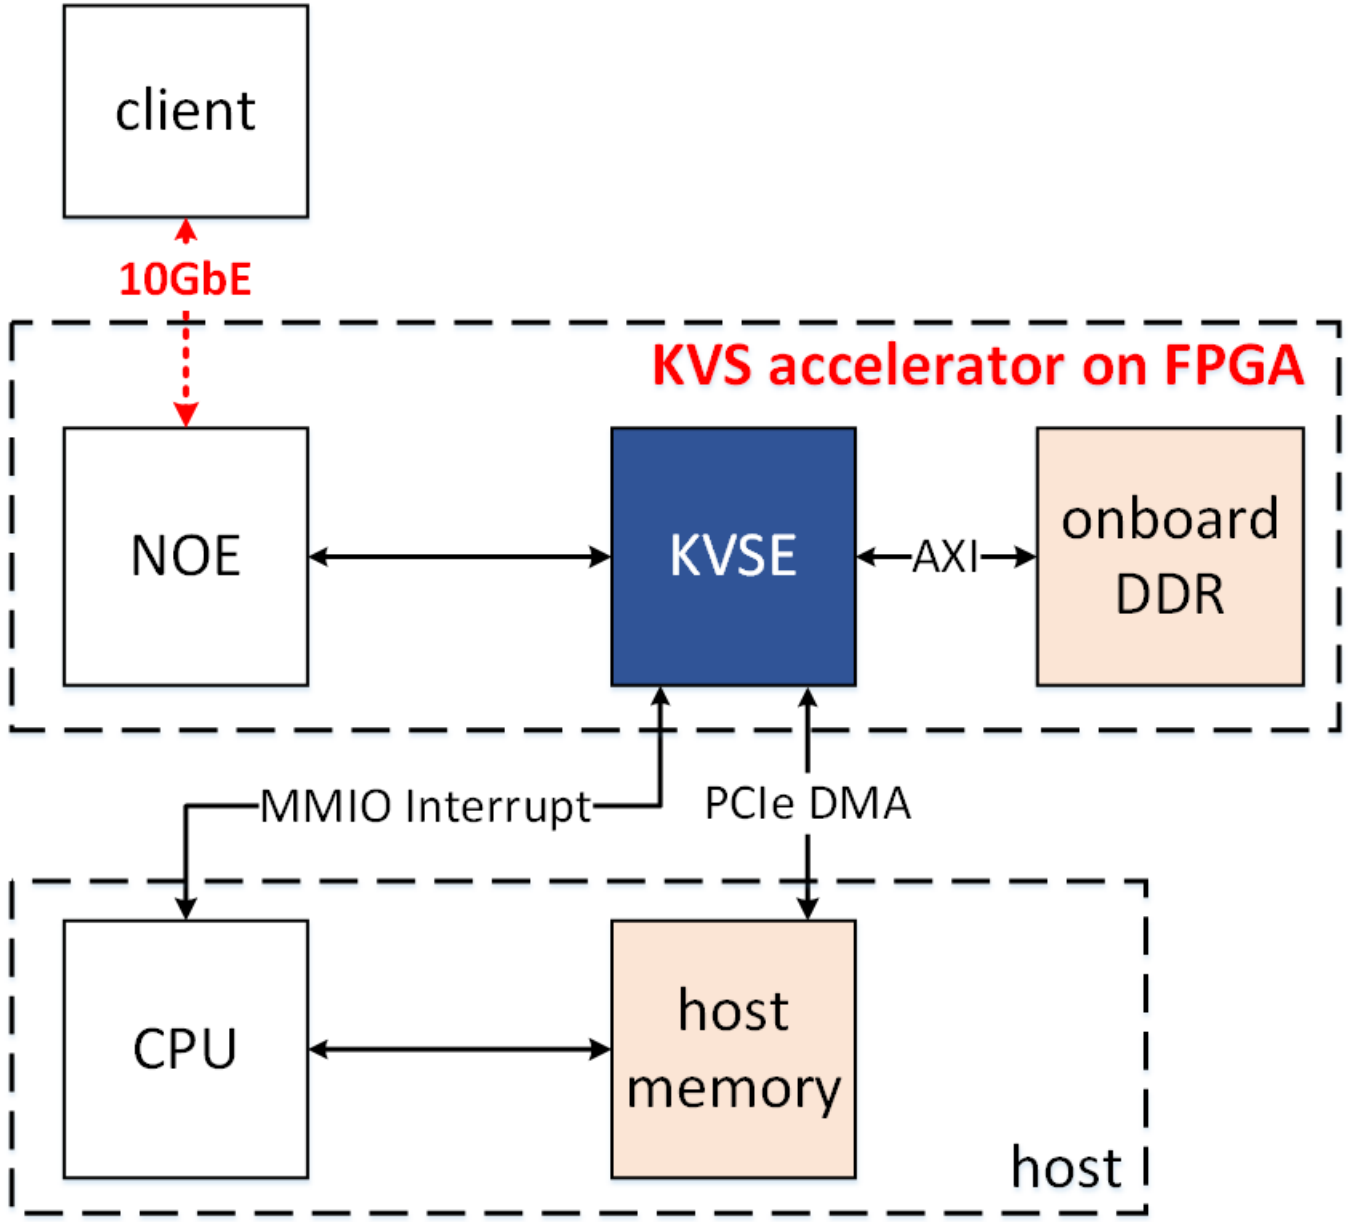

Architecture - *Network Offload Engine*

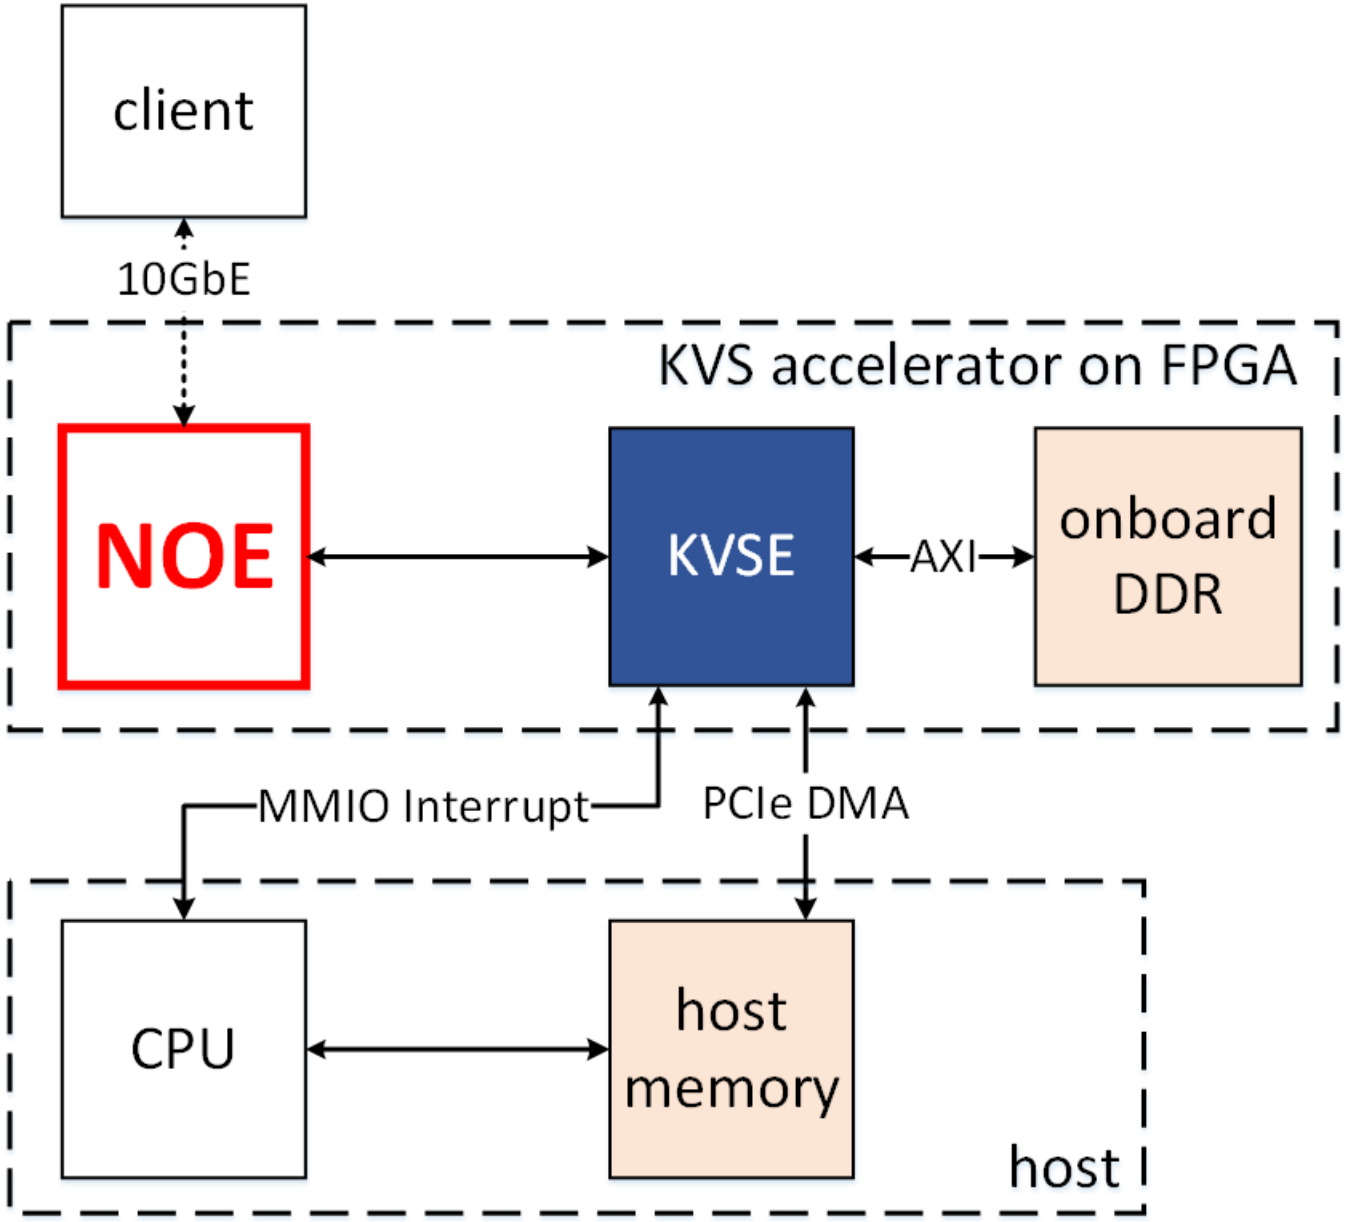

Architecture - *Key-Value Store Engine*

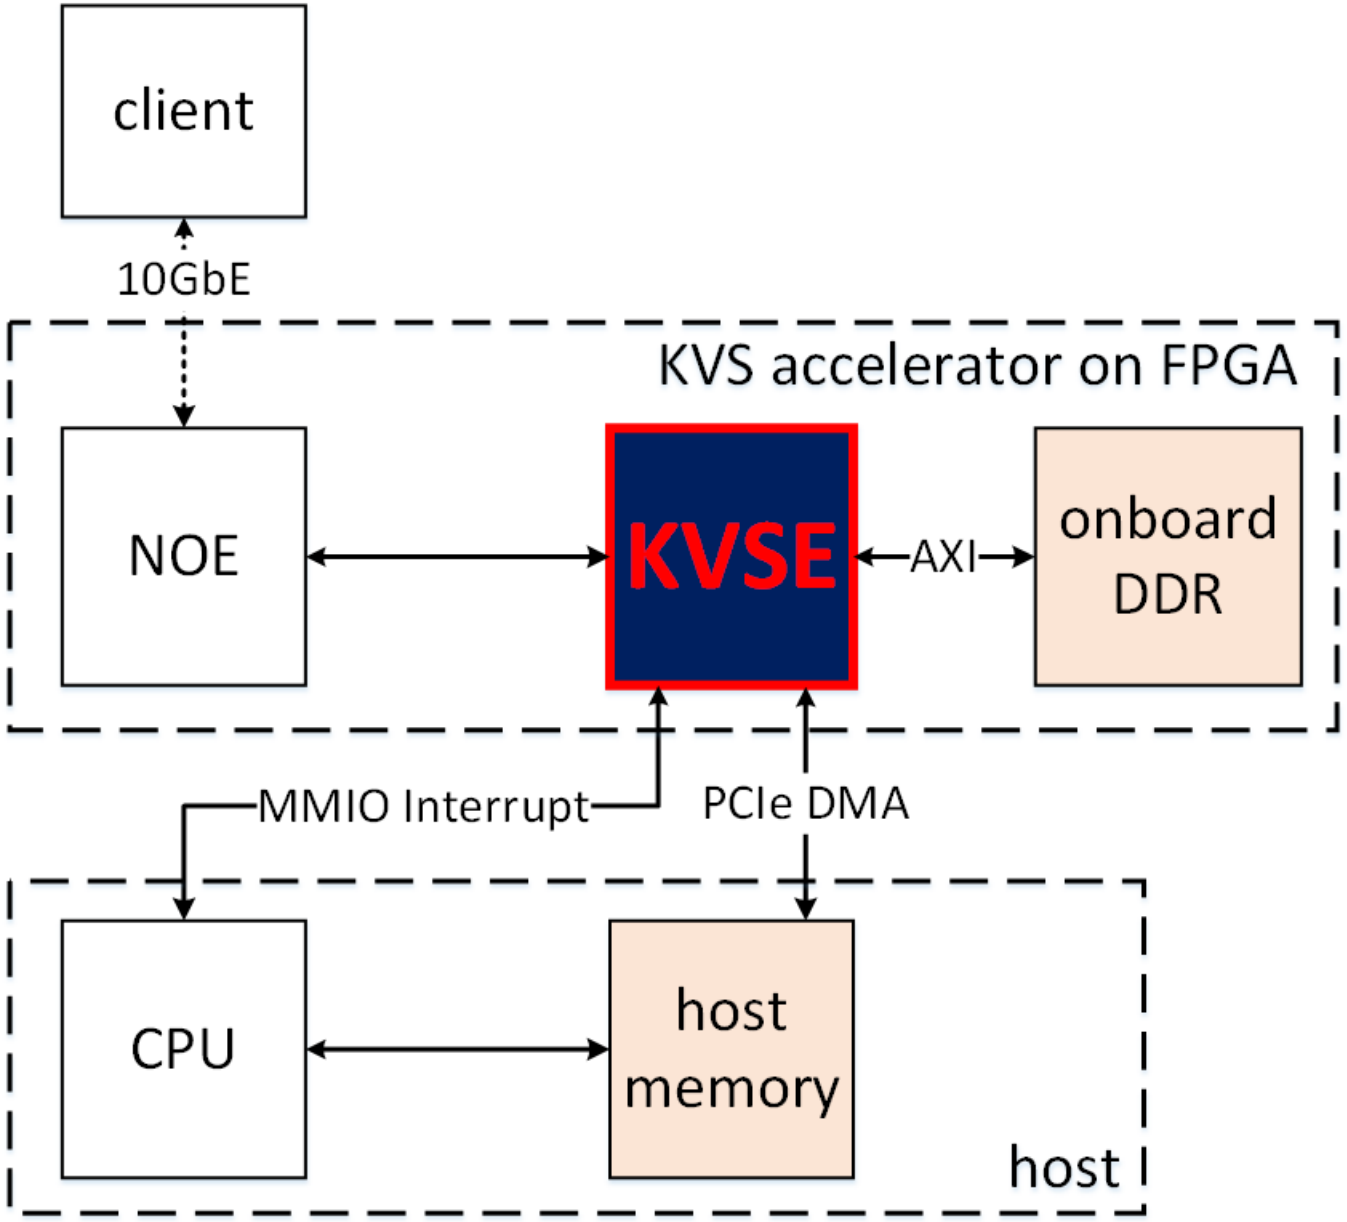

Architecture - First-level indexing by key

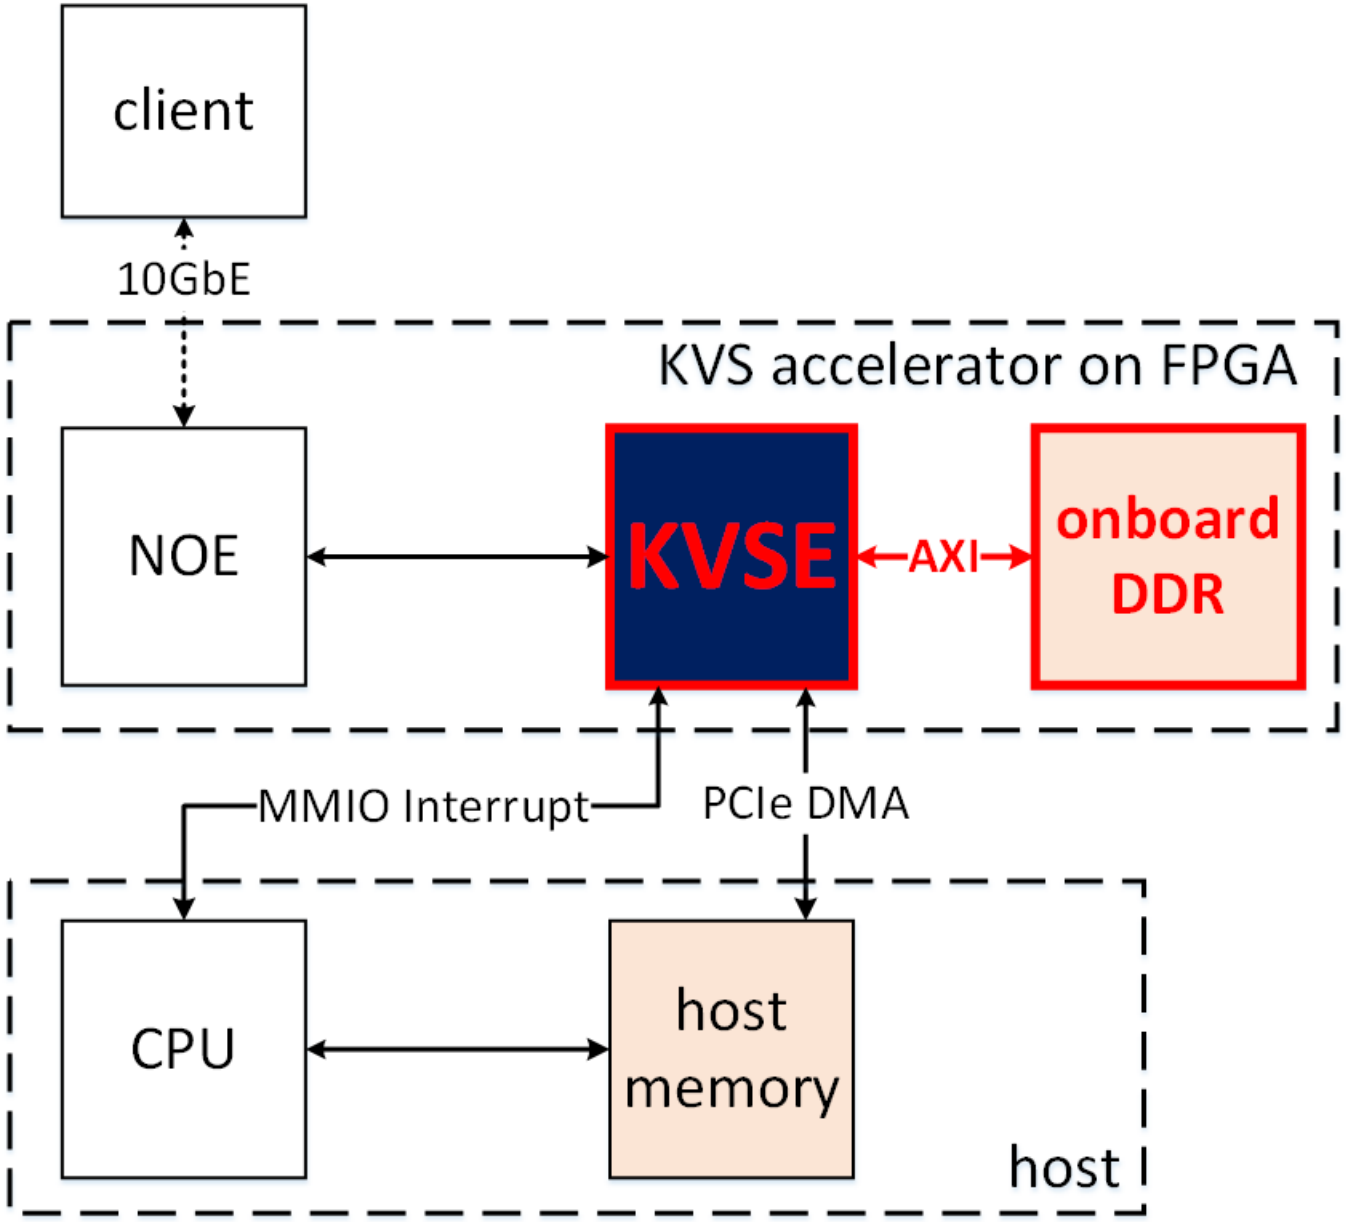

Architecture - Second-level indexing by version

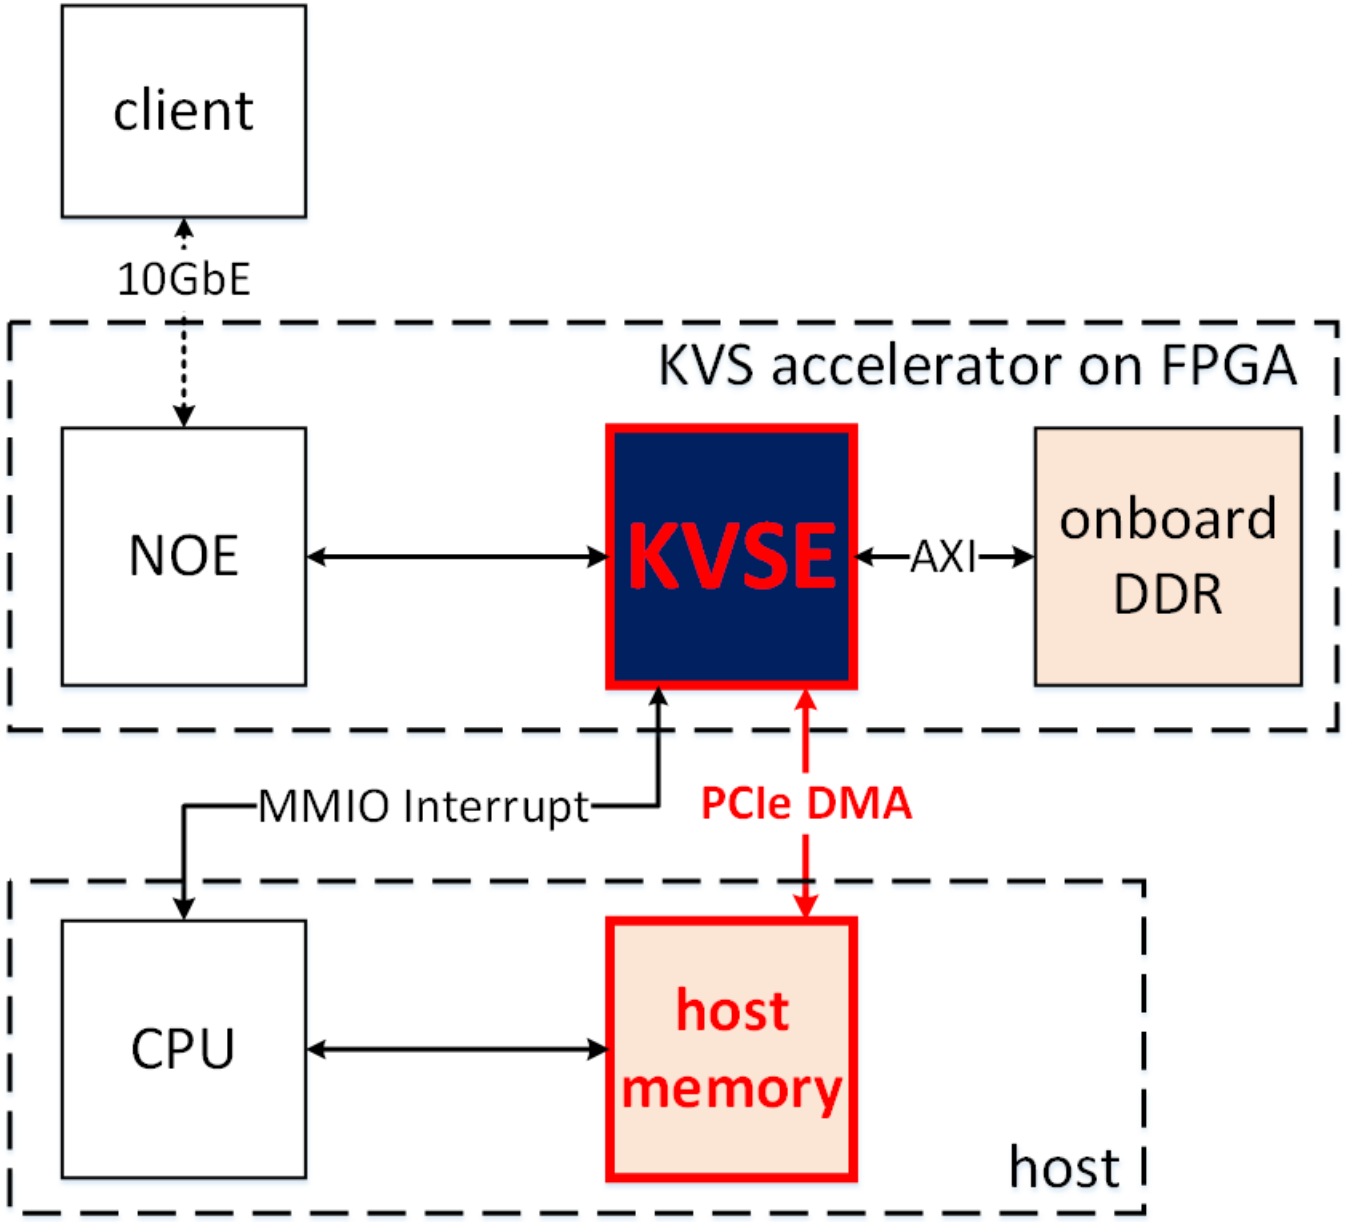

Implementation

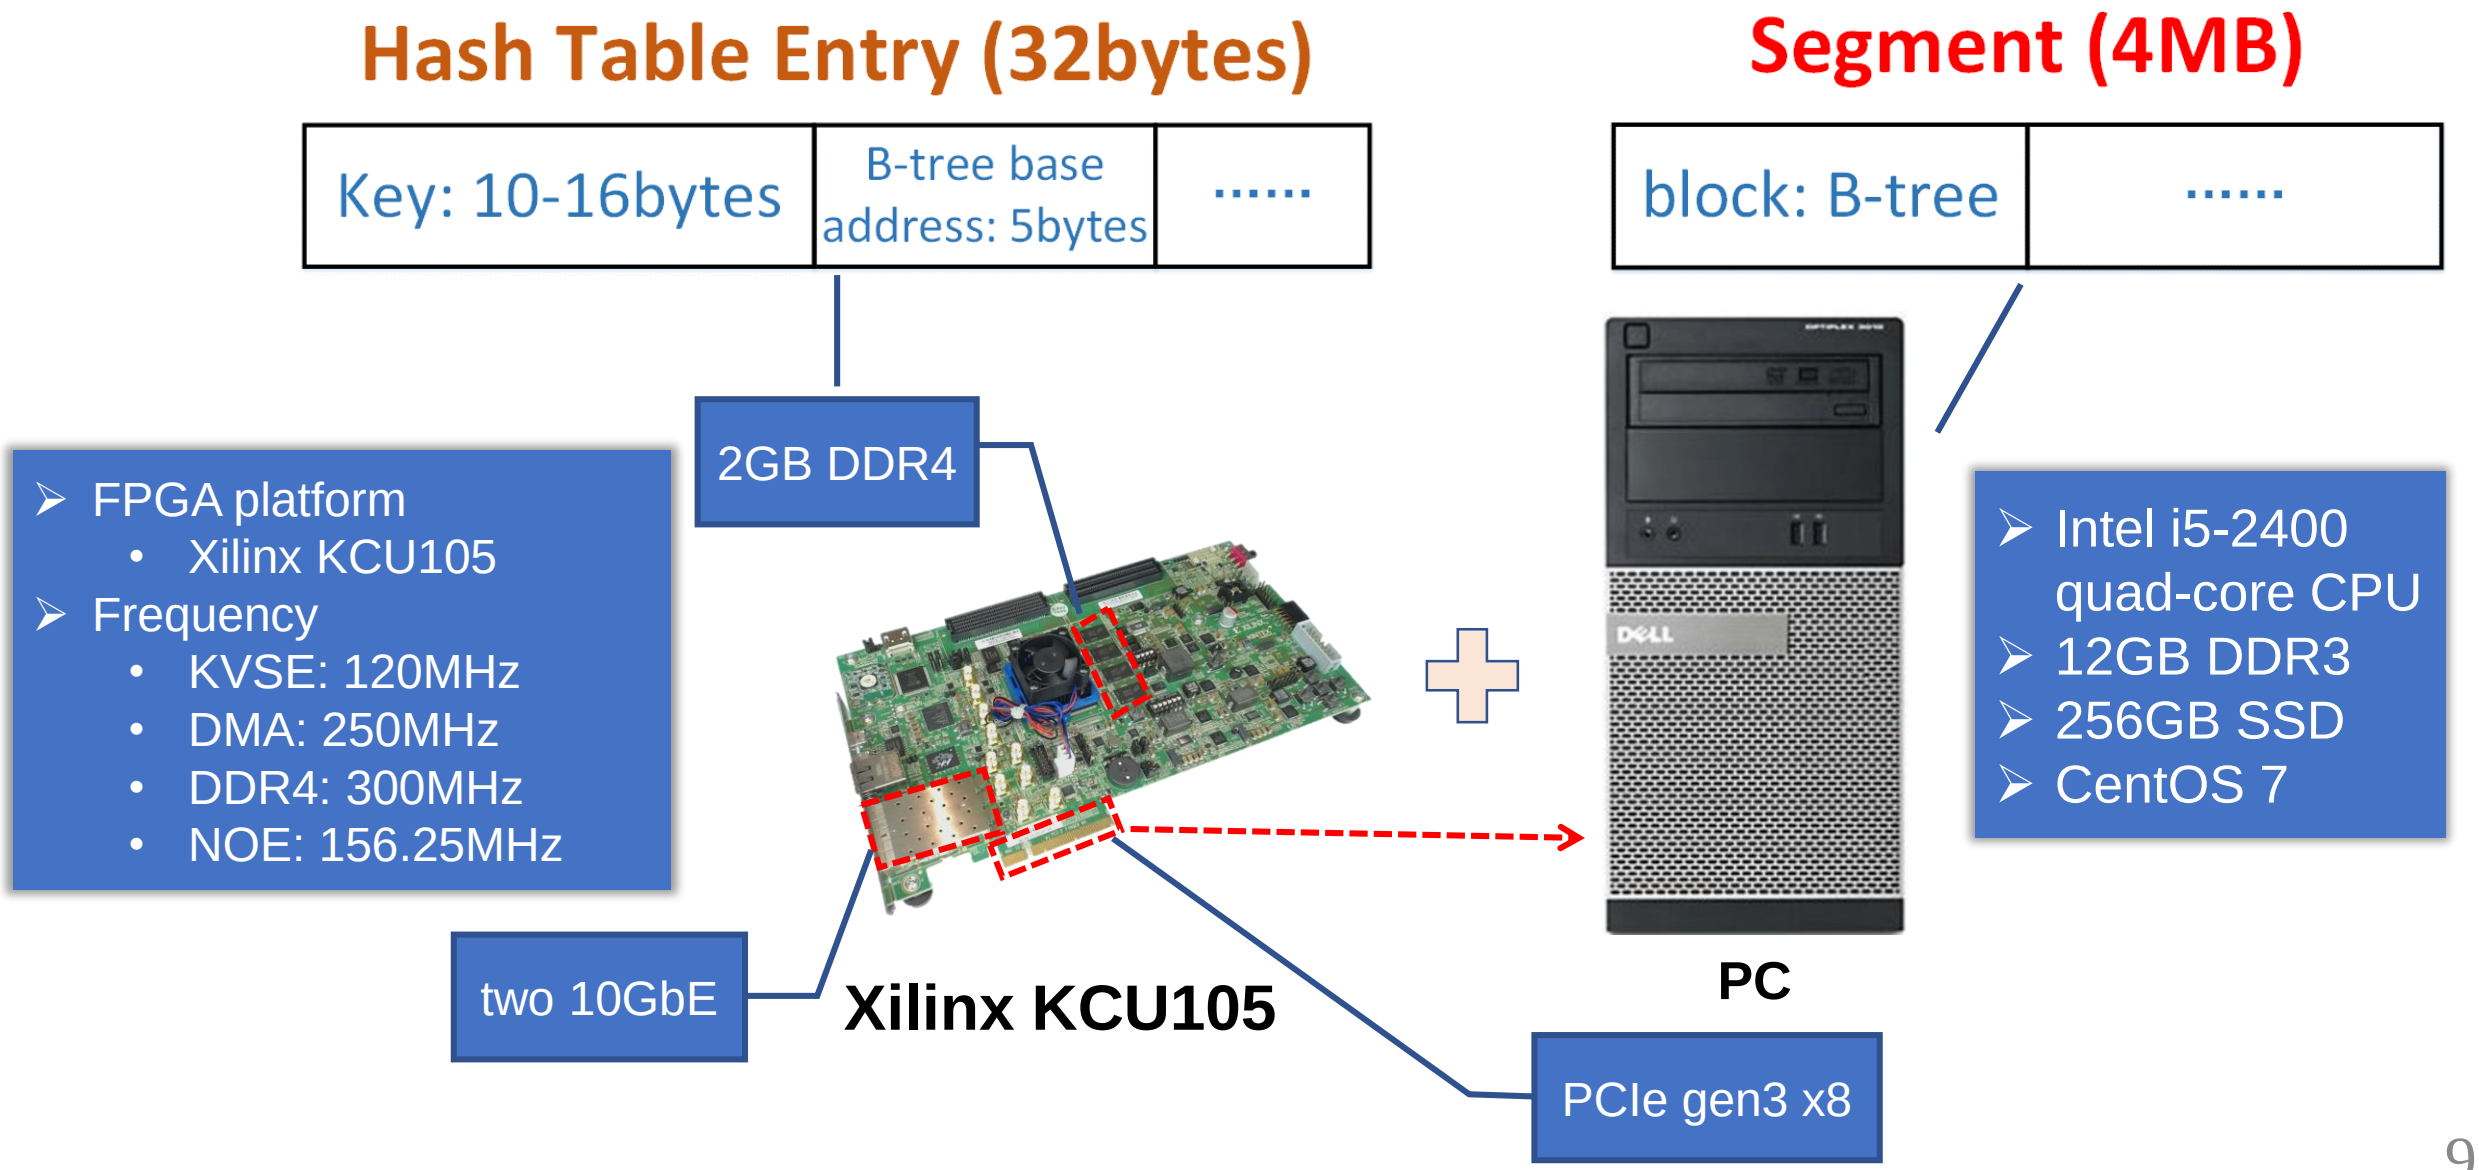

Evaluation - *Key-Value Store Message Generator* in FPGA hardware

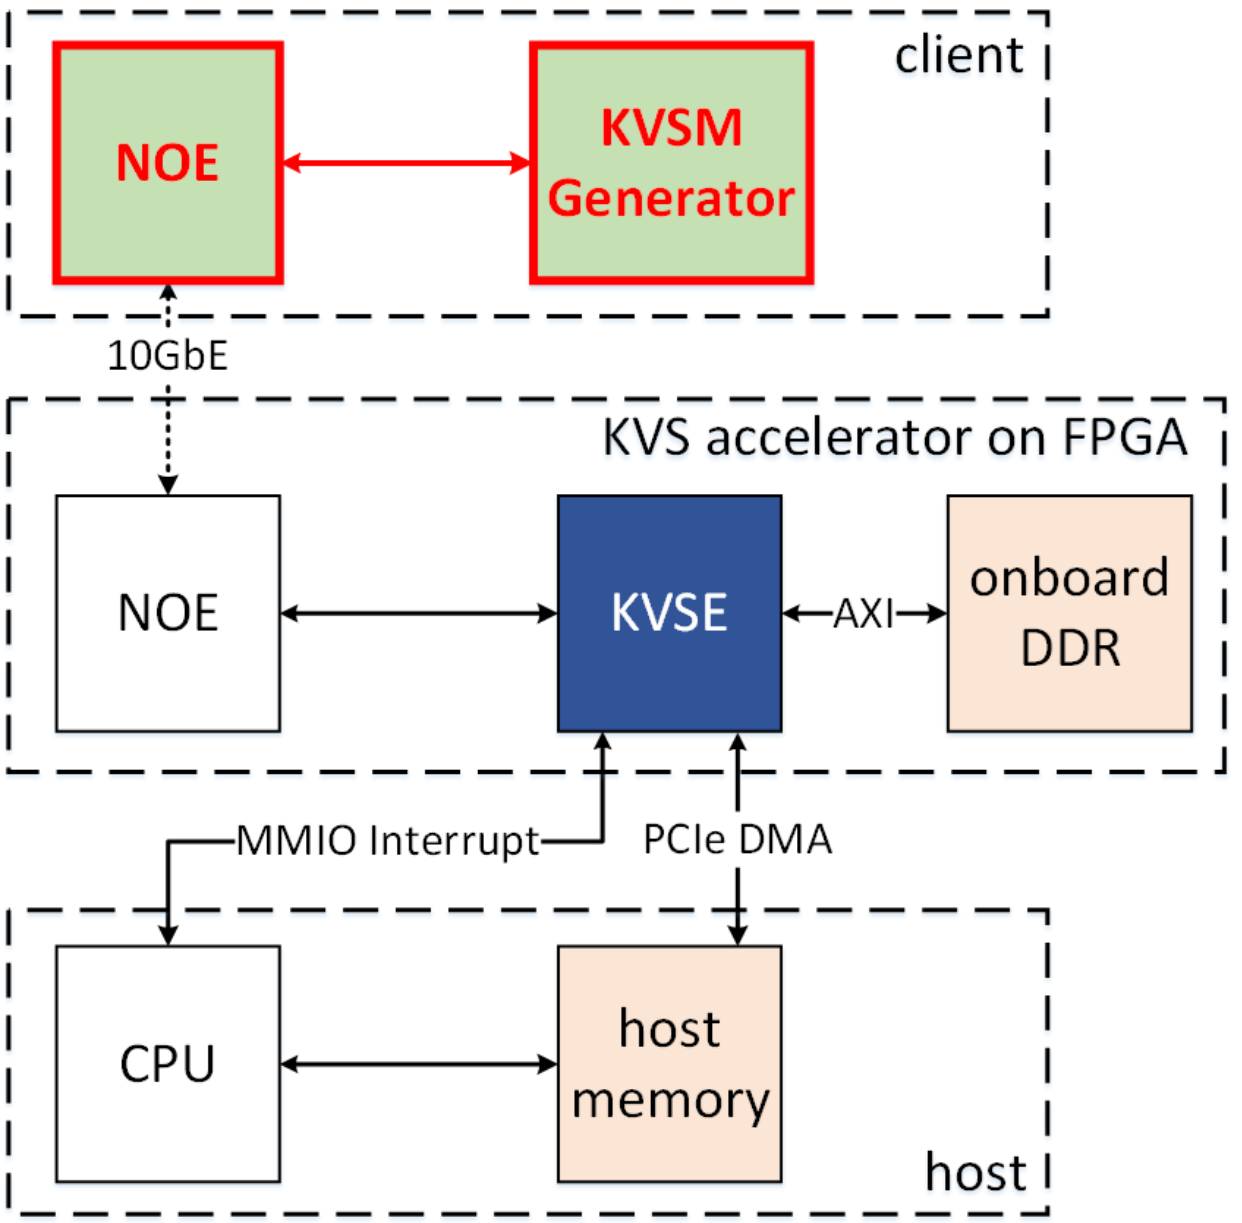

Evaluation - Results

- **Latency increases almost linearly**

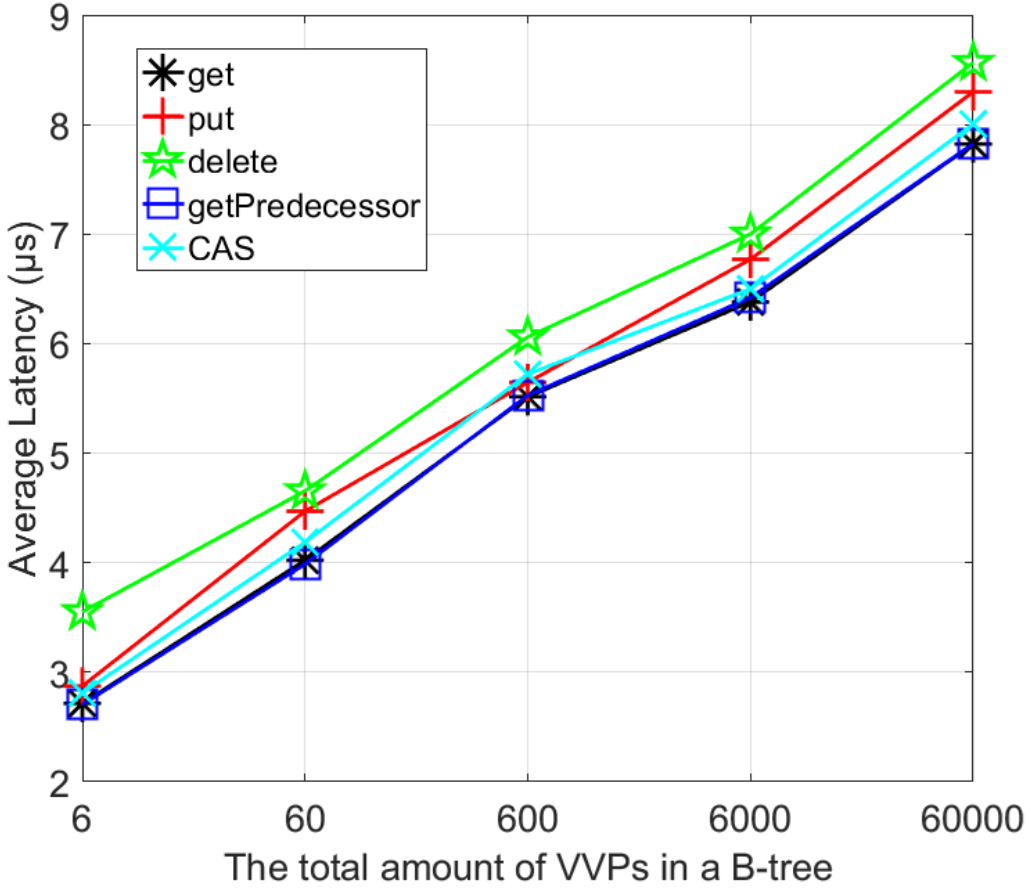

- **KVSE is the bottleneck**

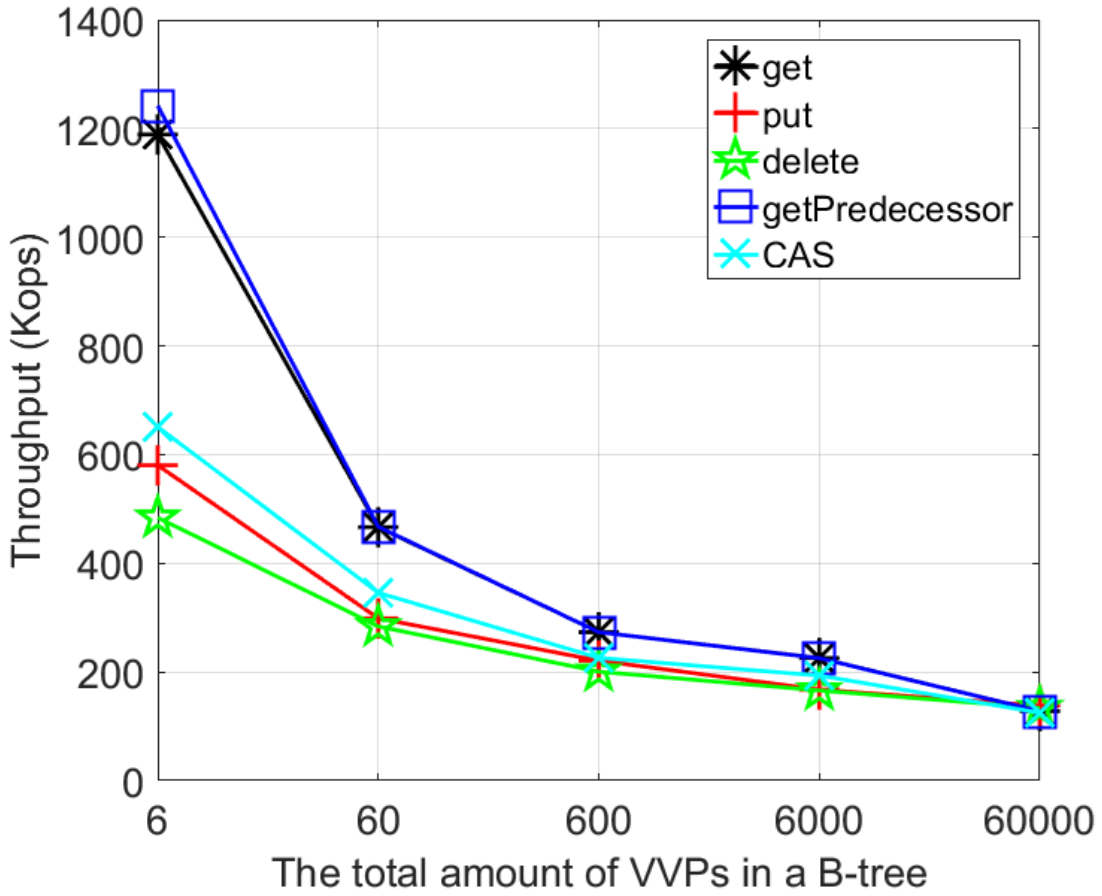

* *Kops:* Thousand operations per second

Conclusion

Comparison (latency, *get* operation)

- Our KVS: < 8μs (within a B-tree of 5 levels)
- Hybrid FPGA approach: ≈ 75μs (within a B^+-tree of 5 levels)
- Many software-based KVS systems: > 1ms (on the support of versioning)

Future work

- Optimize the system architecture of our multi-version KVS.
- Expand to a distributed KVS by setting up multiple storage hosts.

* *Hybrid FPGA approach:* D. Heinrich, S. Werner, M. Stelzner, C. Blochwitz, T. Pionteck and S. Groppe, "Hybrid FPGA approach for a B+ tree in a semantic Web database system," 2015 10th International Symposium on Reconfigurable Communication-centric Systems-on-Chip (ReCoSoC), Bremen, 2015, pp. 1-8.

Thanks!

Contact

ycren18@fudan.edu.cn wbyin@fudan.edu.cn llwang@fudan.edu.cn