CXL Memory Performance for In-Memory Data Processing

VLDB 2025 | September 1 – 5 | London, United Kingdom

Marcel Weisgut, Daniel Ritter, Pınar Tözün, Lawrence Benson, Tilmann Rabl

Disaggregated Infrastructure

Allows individual resource scaling:
→ improves utilization & reduces TCO.

Storage disaggregation widely adopted.

Increasing efforts on **memory disaggregation.**

CXL — Technology to implement memory disaggregation.

Our Work: CXL Memory Performance Study

- Memory access performance evaluation on a server with four CXL memory devices.
- Column scans & B⁺-Tree workloads with data across up to four CXL devices.
- TPC-H evaluation with simple data placement heuristics on the in-memory DBMS Hyrise.
- Economic viability study of CXL memory devices.

Setup

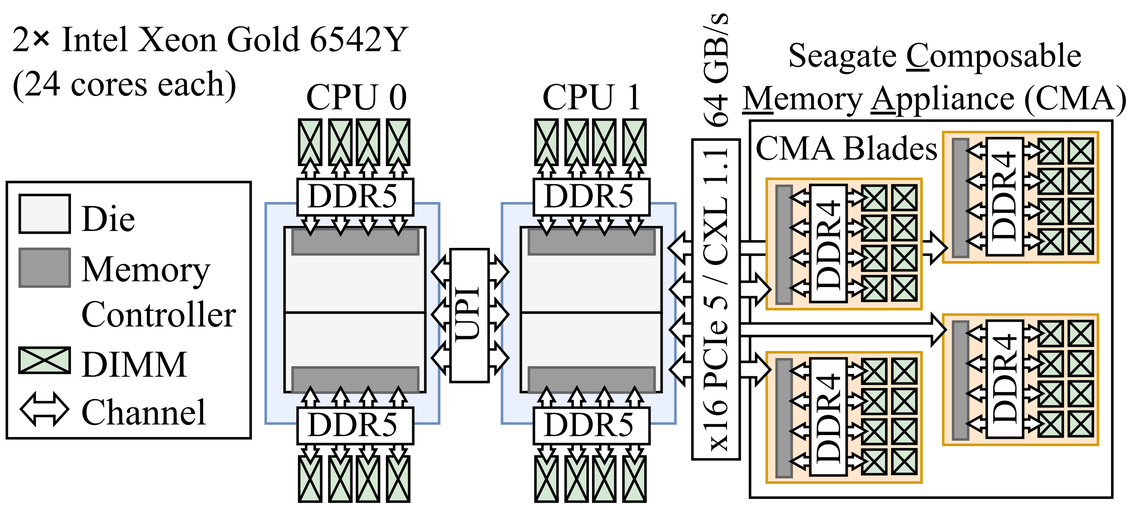

Memory per CPU:
8 x 32 GiB DDR5 (4800 MT/s)

Memory per CMA blade:
8 x 128 GiB DDR4 (1866 MT/s)

SIMD Column Scan

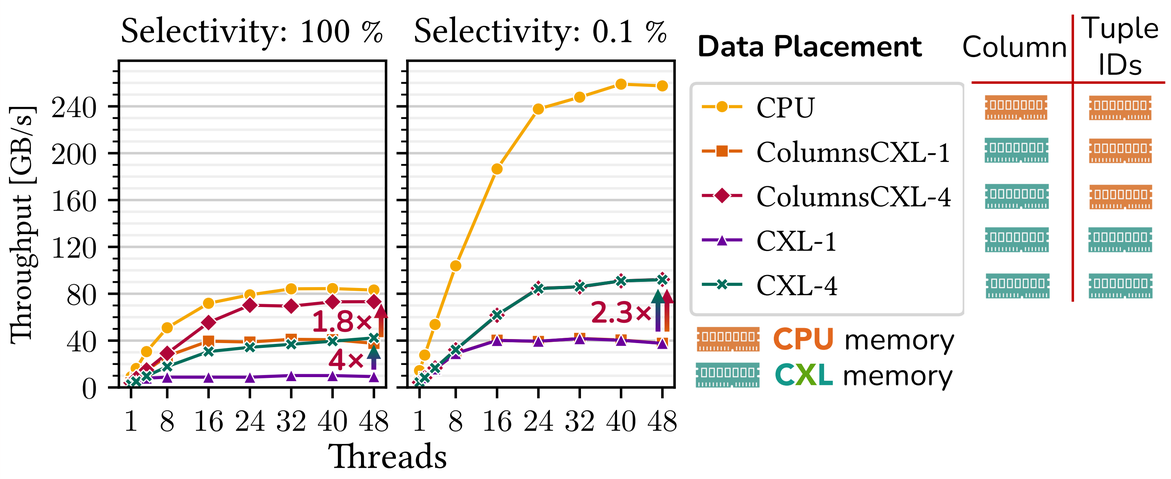

Lower selectivity → fewer writes → higher throughput.
Multiple devices yield higher throughput.

TPC-H Evaluation on the In-Memory DBMS Hyrise

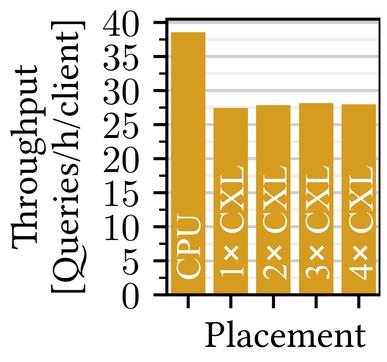

Indicates:
Hyrise is mainly latency-bound.

Black Box Approach (Page Granularity)

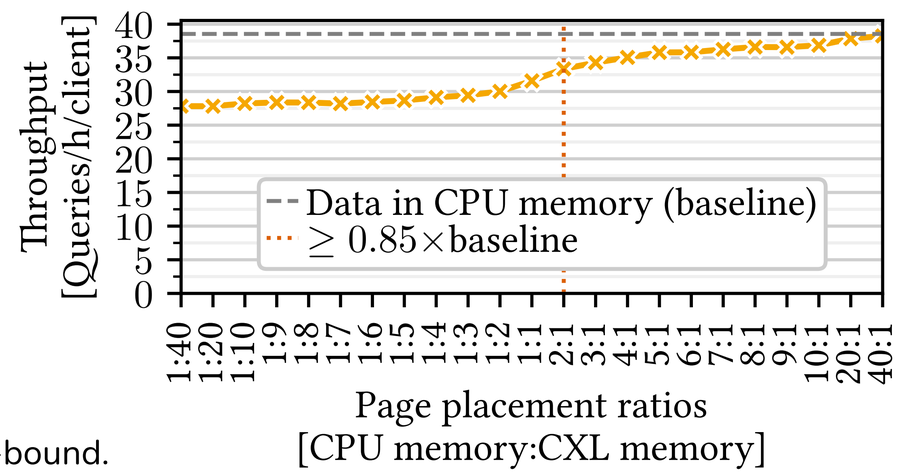

Access Frequency-Based (Column Granularity)

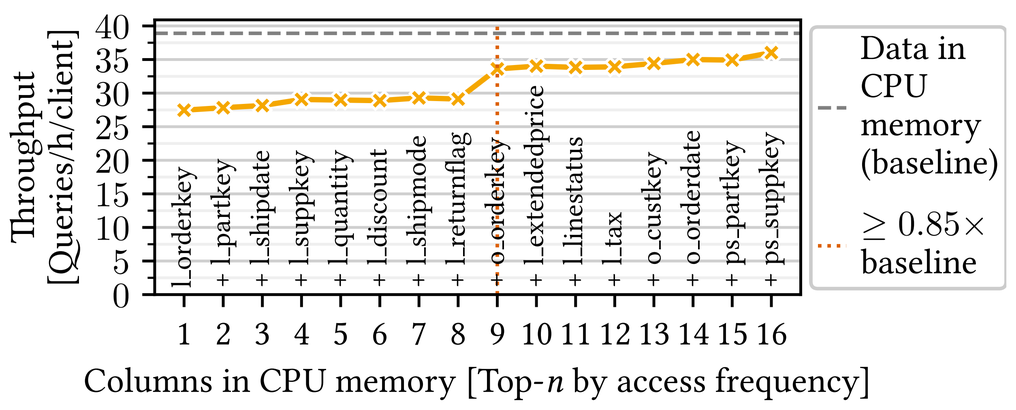

Both strategies use CXL memory of one device.

CXL memory for cold and warm data with hot data in CPU memory leads to high performance when a DBMS uses both CPU and CXL memory.

Economic Viability

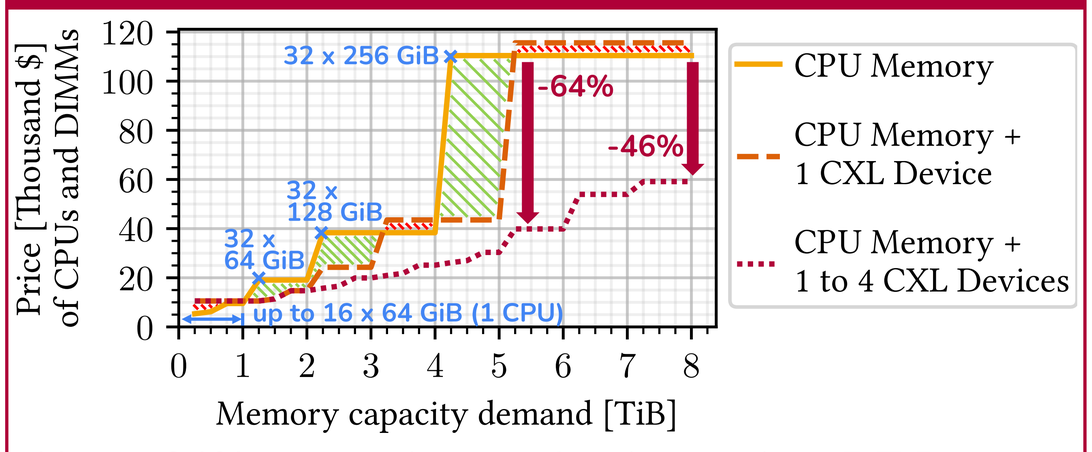

Up to 64% cost reduction for demands > 5 TiB.
Reusing DIMMs in CXL devices further reduces cost.

Conclusion

- CXL memory can reduce cost while providing high OLAP throughput with hot data in CPU memory.
- Bandwidth-bound workloads benefit from CXL memory bandwidth of multiple devices.
- Write-throughput limited workloads can benefit from writes to CPU memory and reads from CXL memory.
- Changing data structures according to their memory bottleneck improves performance.

HPI Hasso Plattner Institut
Database Group
Hasso Plattner Institute, University of Potsdam
Potsdam, Germany

We thank Seagate Technology LLC and Intel Corporation for their support, and Martin Boissier and the anonymous reviewers for their feedback. This work was partially funded by SAP, the German Research Foundation (ref. 414984028), the European Union's Horizon 2020 research and innovation programme (ref. 957407), and the Independent Research Fund Denmark's Inge Lehmann program (grant agreement number 0171-00062B).