

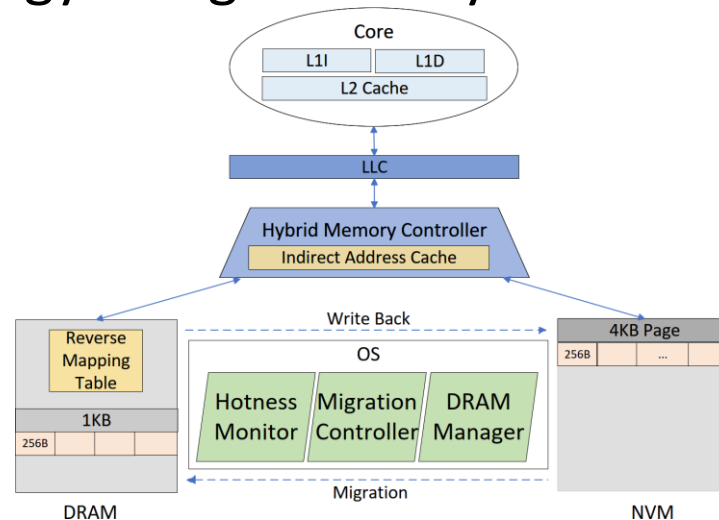
# A Hybrid Memory Architecture Supporting Fine-grained Data Migration

**Ye CHI, Jianhui YUE, Xiaofei LIAO, Haikun LIU, Hai JIN**

Frontiers of Computer Science, DOI: [10.1007/s11704-023-2675-y](https://doi.org/10.1007/s11704-023-2675-y)

# Problems & Ideas


- Problems of page migration for hybrid memory systems:
  - Coarse-grained page migration at a granularity of 4 KB pages wastes memory bandwidth and DRAM resources;
  - Fine-grained page migration at a granularity of 64B cache-line introduces significant storage overhead and hardware modifications.
- Ideas: A non-hierarchical caching architecture that organizes DRAM and NVM in a flat address space physically, but manages them in a cache/memory hierarchy, with a page migration strategy at a granularity of 256B.



The architecture of Mocha

# Main Contributions

- Contributions:
  - A non-hierarchical hybrid memory architecture that organizes a portion of DRAM as the cache of NVM;
  - A reverse mapping table in the DRAM to record the physical address mappings between NVM and DRAM, and an IAC in the hybrid memory controller to speed up the address translation;
  - A fine-grained data migration strategy at the granularity of 256 bytes to adapt to the physical feature of Optane media;
  - A dynamic threshold adjustment scheme for hot data identification to improve the efficiency of the utility-based hot data migration.



Comparison of performance metrics between different systems. Left: the normalized IPC based on baseline; Right: the normalized data migration traffic based on HSCC.