




1. Problem Space




Example configuration of the Intel Optane DC PMEM Platform.

1. Hybrid Memories: New platforms accelerate analytics and simulations of ever exploding data sizes, by incorporating new memory technologies, creating TeraBytes of memory capacity at $\frac{1}{3}$ of the cost compared to DRAM-only systems.

2. Complexity Increases: The difference in access speeds and characteristics is more distinct, compared to traditional DRAM-only systems.

3. Traditional Approaches Break: Assumptions and heuristics that were proposed and fine-tuned for traditional memory systems, now break, resulting in significant loss in performance and efficiency.

2. Thesis Statement




Existing approaches build a system-level page scheduler that dynamically monitors application memory accesses, identifies frequently accessed pages given a history of data access behavior and moves them from PMEM to DRAM.

Challenges: Which pages to move and when to move them?

The contributions of this thesis combine *machine intelligent* selection of data movements with *fine-tuned* data movement time intervals, to bridge the performance gap left by related works, while allowing for practical system-level integration.

3. Why is it important?




Which pages to move?

Existing solutions cause up to **40%** application performance loss, caused by a wrong choice of which pages to move, since historical data access information can not always accurately predict future access behavior.

When to move the pages?

Existing solutions empirically tune the frequency of periodic data movements. This may lead to an average performance loss of **70%** for PMEM platforms and **5x** for disaggregated DRAM systems. No single frequency choice seems to work best across applications and across platforms.




Existing solutions leave a significant performance gap from the case of optimal selection of when and which pages to move. There is need for more sophisticated solutions compared to current history-based and empirically-tuned solutions.

4. Thesis Highlight 1

Kleio: a Hybrid Memory Page Scheduler with Machine Intelligence. [Best Paper Award Finalist at HPDC '19](#).


Summary: Kleio uses a combination of lightweight history-based and machine learning methods to predict and identify **which pages** are frequently accessed and periodically move them across DRAM and PMEM.



Insight: Not all pages need machine intelligent management, when history-based predictions work well. Kleio identifies a small page subset whose timely allocation in DRAM will boost application performance.

Machine Learning: Kleio deploys Recurrent Neural Networks for the pages of the selected subset to learn the pattern of their access count frequency.

Closing the Performance Gap: Kleio closes the gap between current history-based approaches and oracular ones by more than **80%**.





Practical Integration: Kleio's design and sophisticated page selection lays the grounds for the practical integration of machine intelligent memory management solutions in future systems.

5. Thesis Highlight 2

The Case for Optimizing the Frequency of Periodic Data Movements over Hybrid Memory Systems. [MEMSYS '20]

Cori: Dancing to the Right Beat of Periodic Data Movements over Hybrid Memory Systems. [Ongoing Work]

Summary: Cori is a system-level profiler that makes a sophisticated suggestion of the length of periodic time intervals **when moving** the selected pages will enable maximum performance improvements.



Observation: When the duration of periodic data monitoring and migration intervals is shorter than the median data reuse time, the effectiveness of history-based page schedulers degrades and application performance suffers.

Insight: Insights regarding data reuse lead to an informed choice of data movement frequencies that improve performance and resource efficiency.

Solution: Ongoing Work.

6. Remaining Work

Design of a comprehensive page scheduler, that leverages insights from both thesis highlights and adjusts the data movement selection and frequency to the upcoming data access patterns, creating an online adaptive solution.

Prior Publications optimize static hybrid memory allocations.

CoMerge: Toward Efficient Data Placement in Shared Heterogeneous Memory Systems. [MEMSYS '17]

Mnemo: Boosting Memory Cost Efficiency in Hybrid Memory Systems. [HPBDC workshop of IPDPS '19]

Funding: This thesis was partially supported by NSF award SPX-1822972, the DOE ECP project on Simple Interfaces for Complex Memories (SICM) and the DOE SSIO Unity project.