FORTH
INSTITUTE OF COMPUTER SCIENCE

TeraHeap: Reducing Memory Pressure in Managed Big Data Frameworks

Iacovos G. Kolokasis, Giannos Evdorou, Shoaib Akram, Christos Kozanitis, Anastasios Papagiannis, Foivos S. Zakkak, Polyvios Pratikakis, and Angelos Bilas

Analytics Frameworks Need Large Heaps

- Analytics frameworks use managed runtimes
- To process large amount of data they need large heaps
- Large heaps are expensive and increase GC cost!
- DRAM is expensive in dollar cost, energy, and power
- GC requires expensive scans over large heaps

Today: Move Objects Off-heap

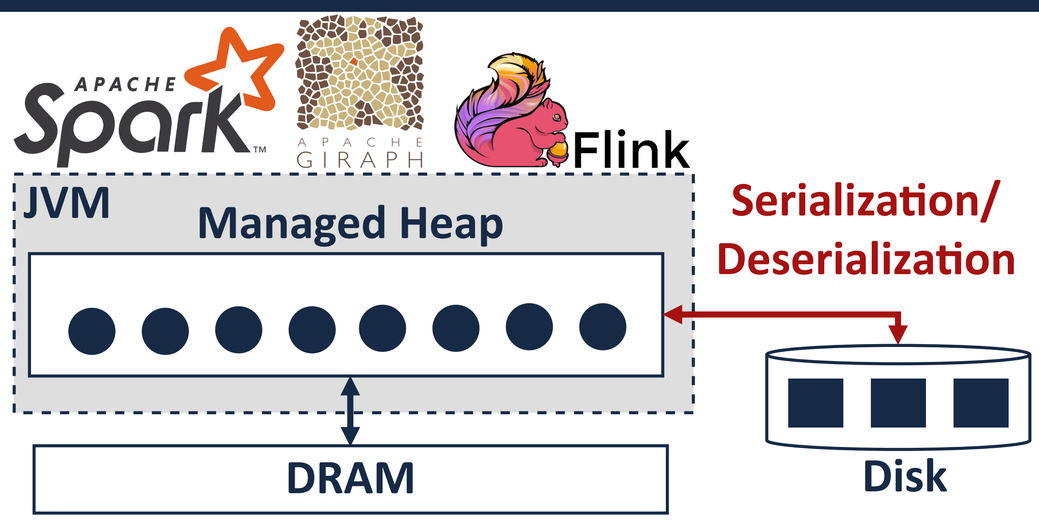

Serialization/Deserialization (S/D) is Terrible!

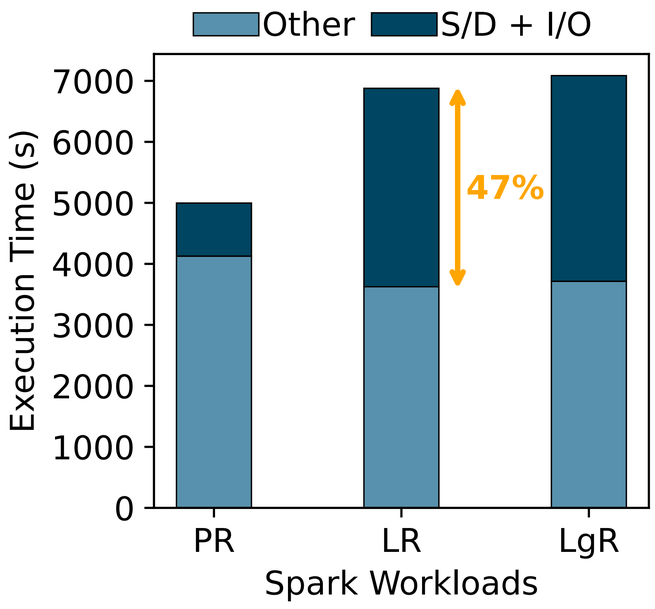

Additional complications:

☹ Serialization is not supported for all objects

☹ Moving objects off-heap can be unsafe

Extend the Heap Over Storage

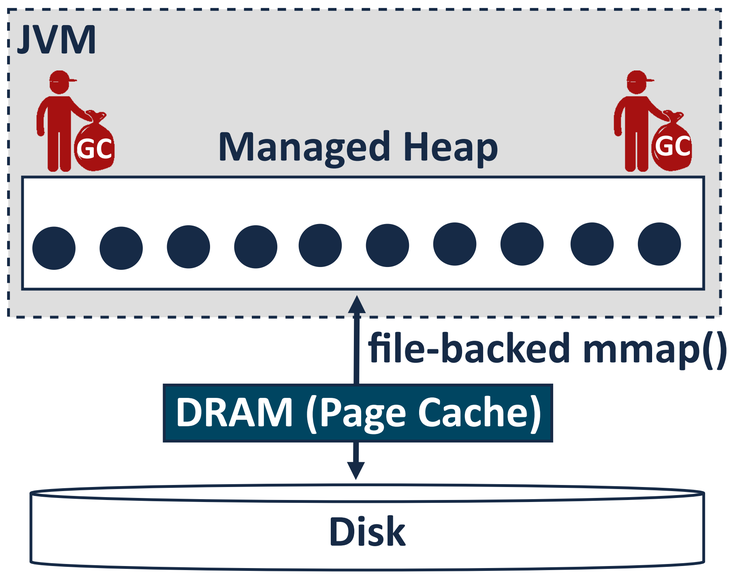

☺ Eliminate S/D cost

☹ Increase GC cost

TeraHeap: Eliminates S/D Without Increasing GC Cost!

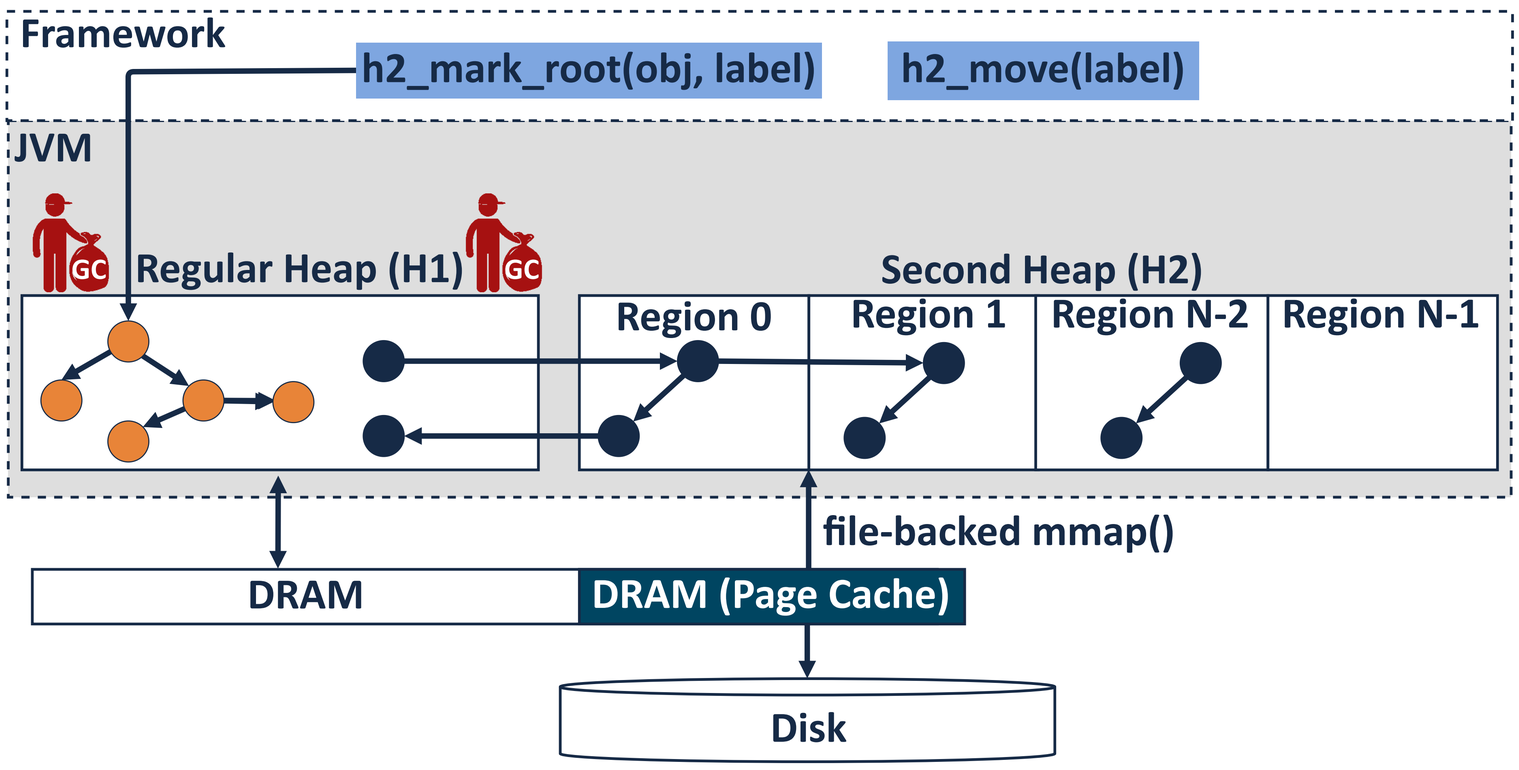

☺ Provide the illusion of a single managed heap

☺ Eliminate S/D cost

☺ Avoid GC scans in the device heap

Challenges

- Identify which objects to move to H2
- Reclaim dead objects in H2 without GC scans
- Track cross-heap references with low I/O cost

Comparison With Same and Less Amount of DRAM

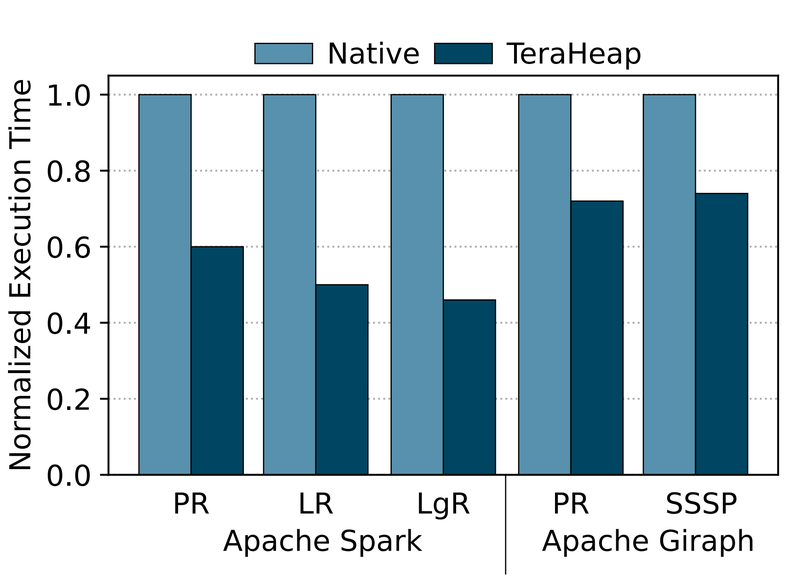

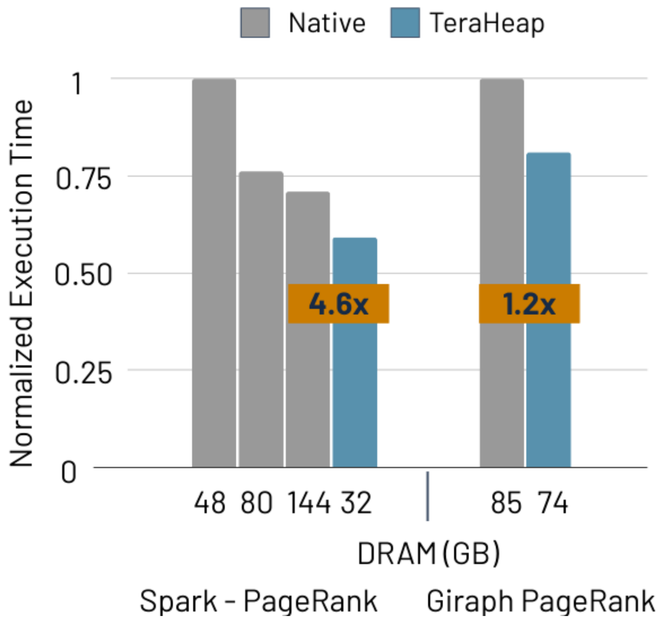

Key Takeaways

- Analytics frameworks deal with large datasets using S/D
- TeraHeap improves Spark performance by up to 54%
- TeraHeap improves Giraph performance by up to 28%
- TeraHeap requires up to 4.6x less DRAM

evolve
Leading the Big Data Revolution

EUPEX
European Pilot for Exascale