

# TeraCache: Efficient Caching over Fast Storage Devices

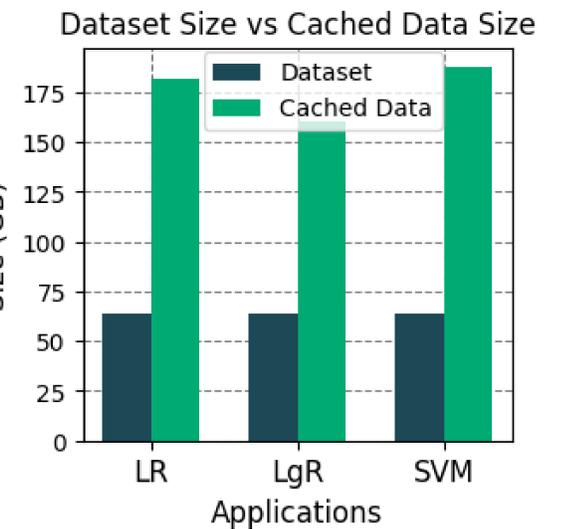
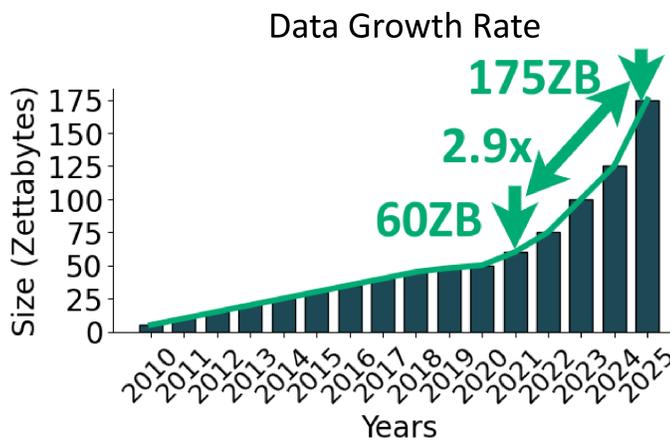
Iacovos G. Kolokasis<sup>1,2</sup>, Anastasios Papagiannis<sup>1,2</sup>, Foivos Zakkak<sup>3</sup>,  
Shoaib Akram<sup>4</sup>, Christos Kozanitis<sup>1</sup>, Polyvios Pratikakis<sup>1,2</sup>, and Angelos Bilas<sup>1,2</sup>

<sup>1</sup>Foundation for Research and Technology – Hellas GR, <sup>2</sup>University of Crete GR, <sup>3</sup>Red Hat Inc., <sup>4</sup>Australian National University AU

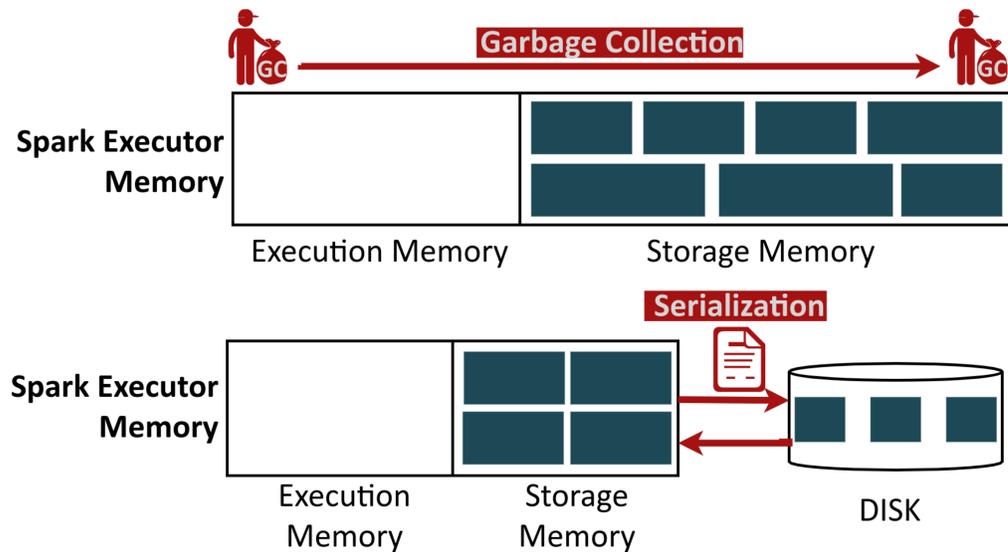
## Fast Storage Devices Available but Analytic Stacks Not Ready

### Increasing Memory Demands for Analytics

- Analytic servers use caches for avoiding recomputation (compute caches)
- Cache size is often several times the input dataset size
- DRAM scaling is limited (more \$ per GB)
- Analytics resort to high capacity, fast storage devices, such as NVMe

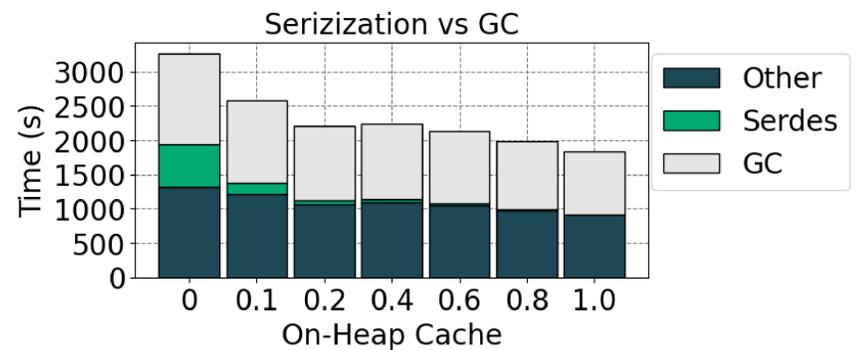


## Spark Offers On-Heap and Off-Heap Caching

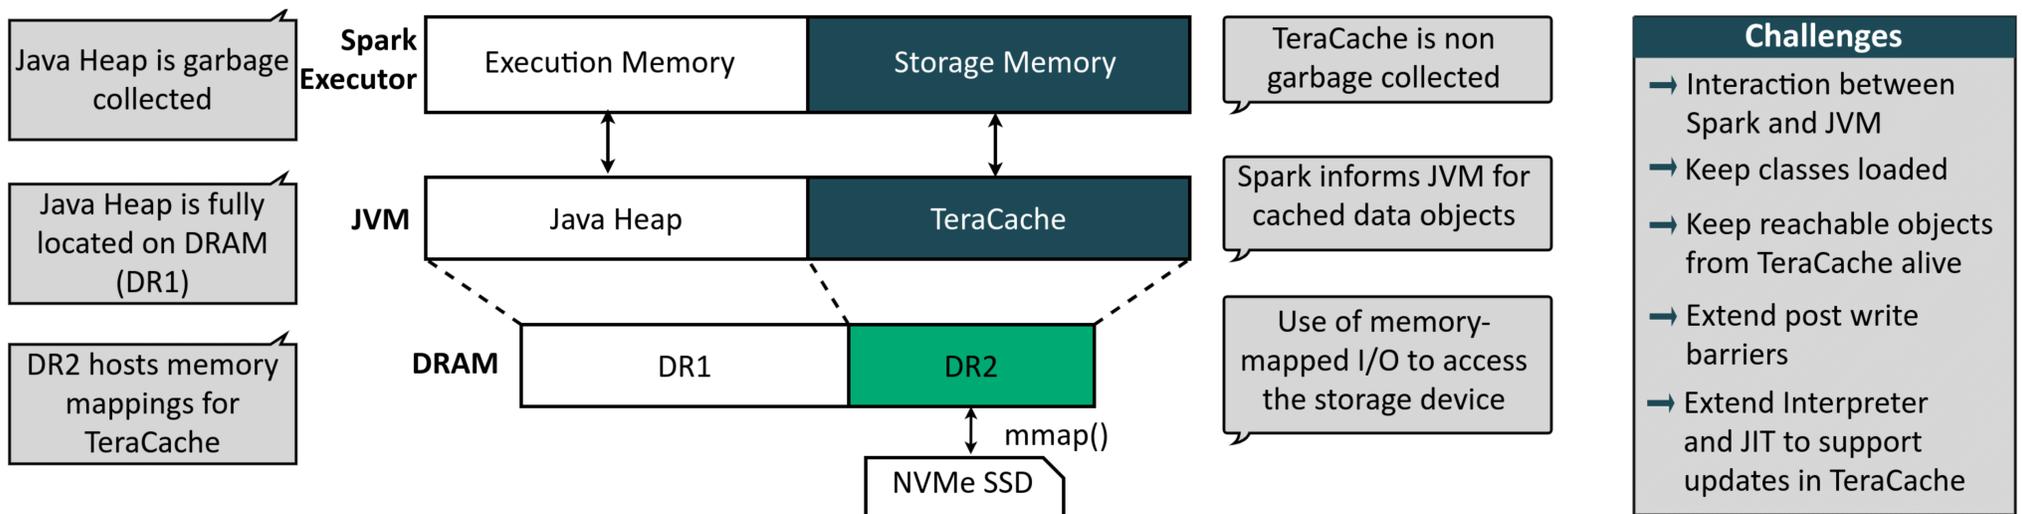


## Merging On-Heap and Off-Heap Benefits

	Pros	Cons
[HTML]D6D6D6		
[HTML]D6D6D6On-heap	No Serialization	High GC Time
[HTML]D6D6D6Off-heap	Low GC Time	High Serialization

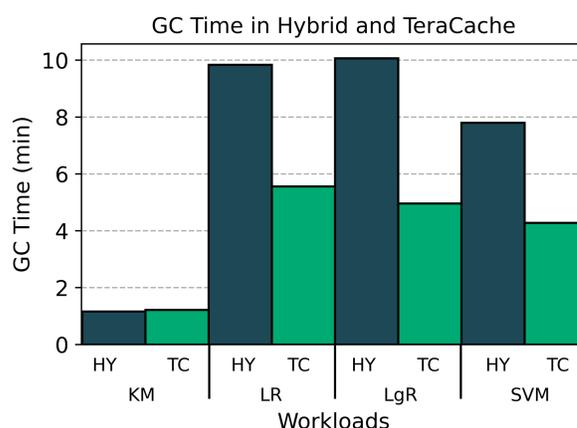
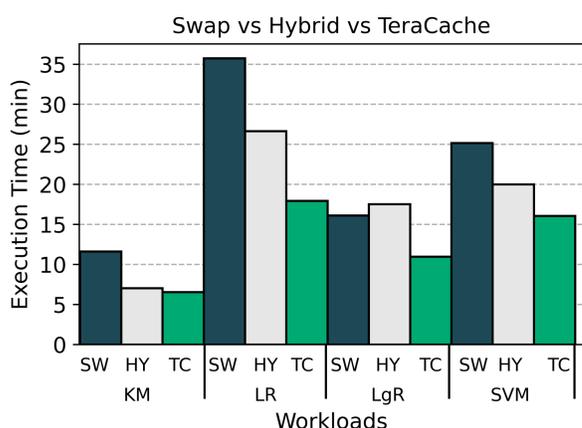


## TeraCache: Best of Both Worlds!



- ### Challenges
- Interaction between Spark and JVM
  - Keep classes loaded
  - Keep reachable objects from TeraCache alive
  - Extend post write barriers
  - Extend Interpreter and JIT to support updates in TeraCache

## Machine Learning Workloads



## Key Takeaways

- RDD caching is critical in Spark
- GC and serialization introduce significant overhead
- TeraCache improves ML workloads performance by 25% over the state-of-the-art

