

Paper Review - Ray: A Distributed Framework for Emerging AI Applications

Summary

This paper introduces Ray, a distributed framework built to meet the needs of reinforcement learning (RL) and other emerging AI applications. Existing systems like Map Reduce and Apache Spark fall short because RL requires fine-grained, stateful, and dynamic computations that mix simulation, distributed training, and real-time serving. Ray provides a unified interface for both task parallel and actors-based computation on top of a dynamic execution engine. Its main contributions include the Global Control Store (GCS) for scalable, fault-tolerant state management and a bottom-up scheduler for low-latency, high-throughput task placement. Experiments show Ray can scale to over 1.8 million tasks per second and perform better than specialized systems in RL workloads but also demonstrates efficiency on large RL workloads such as Evolution Strategies and Proximal Policy Optimization, achieving better performance with fewer resources and reducing costs compared to specialized systems.

Evaluation

The paper addresses an important gap: existing frameworks cannot meet the fine-grained, stateful, and dynamic needs of RL. Ray's solution is valid and impactful, unifying tasks and actors under a dynamic task graph model. Its novelty lies in the fully distributed control plane (GCS + bottom-up scheduler) and the in-memory object store for fast, fault-tolerant data sharing. The evaluation is strong, combining microbenchmarks with end-to-end comparisons against systems like TensorFlow, Horovod, Clipper, and MPI, where Ray shows competitive or superior results. The manuscript is clear, well-structured, and effectively uses figures and pseudocode.

Main takeaways

- Emerging AI workloads like RL need distributed systems that unify simulation, training, and serving with low latency.
- Ray addresses this through a dynamic task graph that combines tasks and actors, plus architectural innovations like the GCS and bottom-up scheduler.
- The result is a scalable, fault-tolerant system that outperforms, or rivals specialized frameworks while remaining flexible and cost (and compute) efficient.

Strengths

- Ray framework provides the ability to express both task-parallel and actor-based computations providing a strong flexibility for developing complex RL applications, accommodating both stateless and stateful computations.
- Ray matches or outperforms specialized systems in RL by combining the fault-tolerant GCS, bottom-up scheduling, and an in-memory object store, enabling fast, scalable, and efficient execution.

Weaknesses

- Ray's generality limits specialized optimizations, as scheduling decisions often lack full computation graph knowledge and may need complex runtime profiling.
- Ray is slower than OpenMPI for small-object allreduce, since it lacks the low-overhead algorithms that make MPI faster (future work).
- GCS lineage storage can grow unbounded, requiring garbage collection policies that are still under active development.

Discussions:

- When might stitching together specialized systems (example: TensorFlow + MPI + Clipper) still be more practical than a unified framework like Ray, and could fault tolerance be offloaded to another service?
- What garbage collection policies could manage lineage growth effectively beyond simple flushing?