Optimizing Distributed Reinforcement Learning with Reactor Model and Lingua Franca

Pranav Atreya, Jacky Kwok, Amil Dravid

Overview

- Reinforcement learning (RL) workloads come with unique considerations that common distributed computing design patterns (e.g., MapReduce) don't satisfy well
- Typical SoTA frameworks for distributed RL (e.g., Ray[1]) employ the actor model of concurrency
- **Hypothesis: Can the reactor model of computation and thread-based parallelism accelerate distributed RL workloads?**

Actor Model (Ray)

- "Actor" as primary unit of computation, communication through asynchronous message passing
- Message processing by actors need not be strictly ordered
- We use Ray's implementation of the actor model

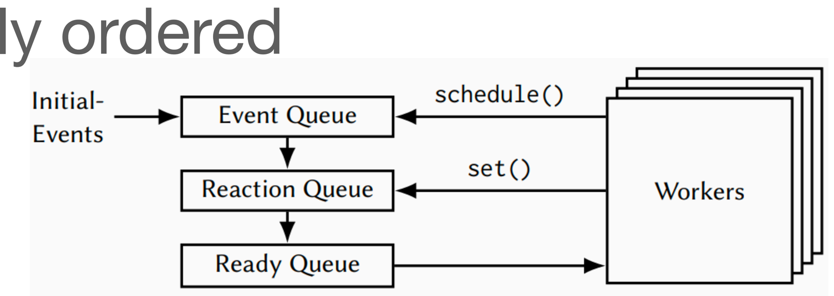

Reactor Model (Lingua Franca)

- Enables deterministic concurrency through "reactors" and "reactions". Instead of messages, reactors react to discrete events each linked to a logical time by triggering reactions. Can be thought of as "sparse synchronous model".
- Reactions can modify state shared by other reactions in the same reactor, but communication across reactors is solely through events.
- Reactors don't directly reference peers, enabling hierarchy
- "Actions" bridge internal determinism with external nondeterminism
- Lingua Franca [2] (partially contributed to by us) implements the reactor model

Distributed Reinforcement Learning

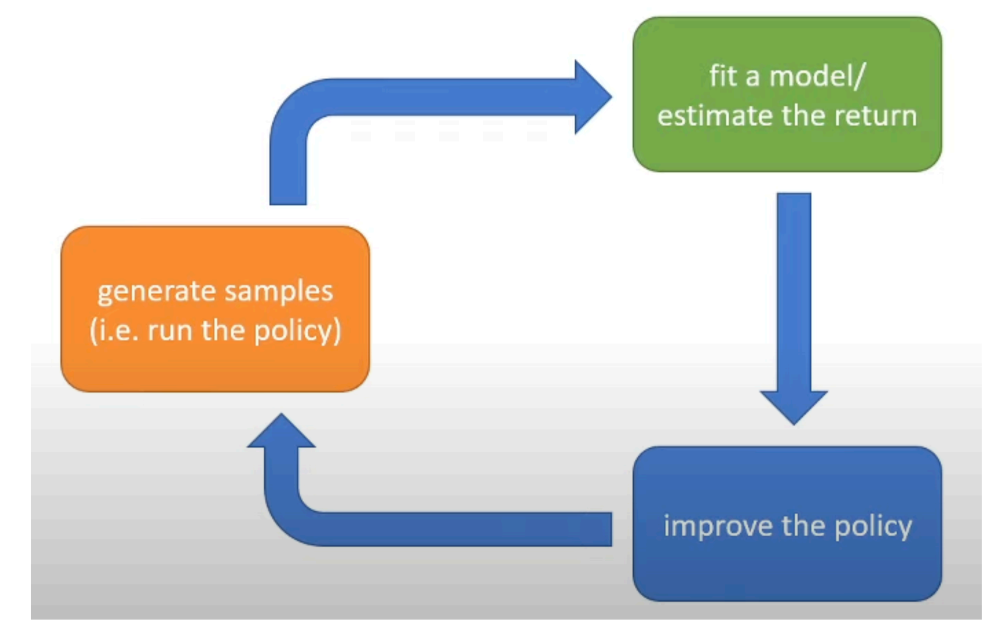

RL Pipeline

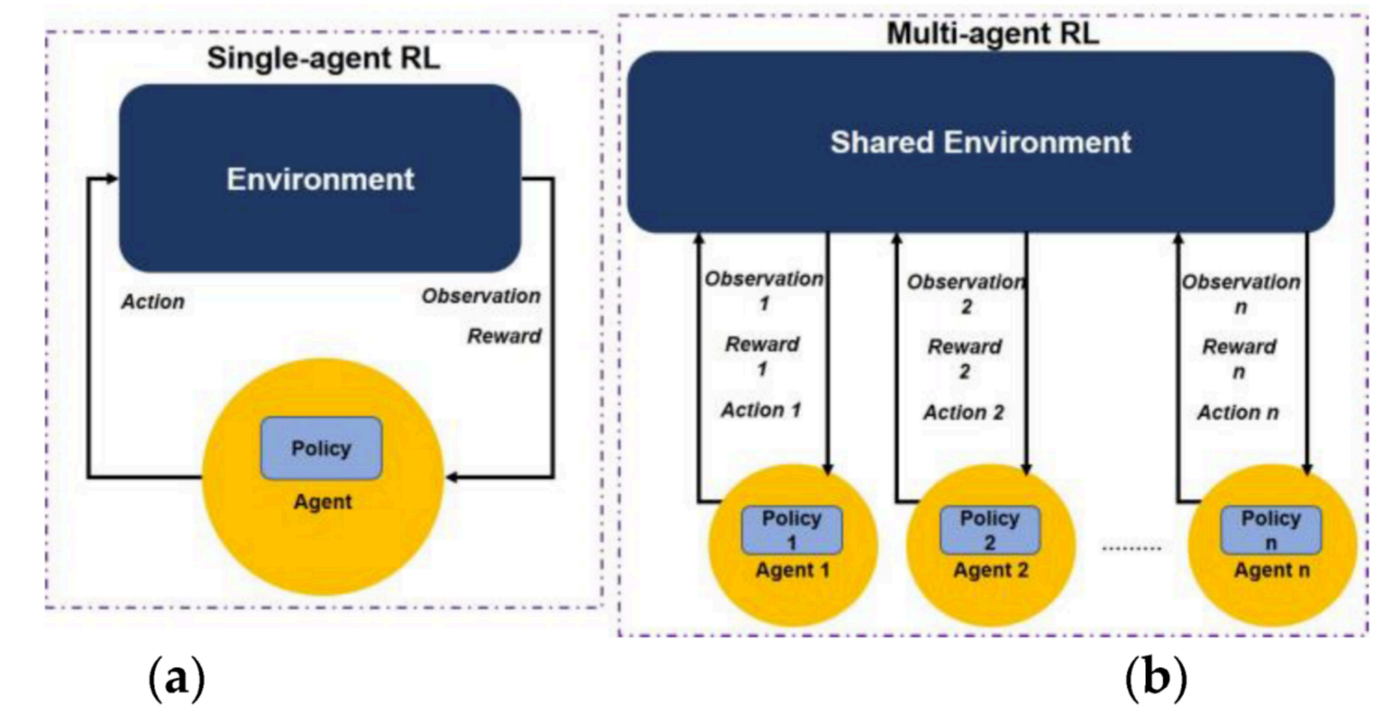

(a) Single-agent model of RL, (b) multi-agent model

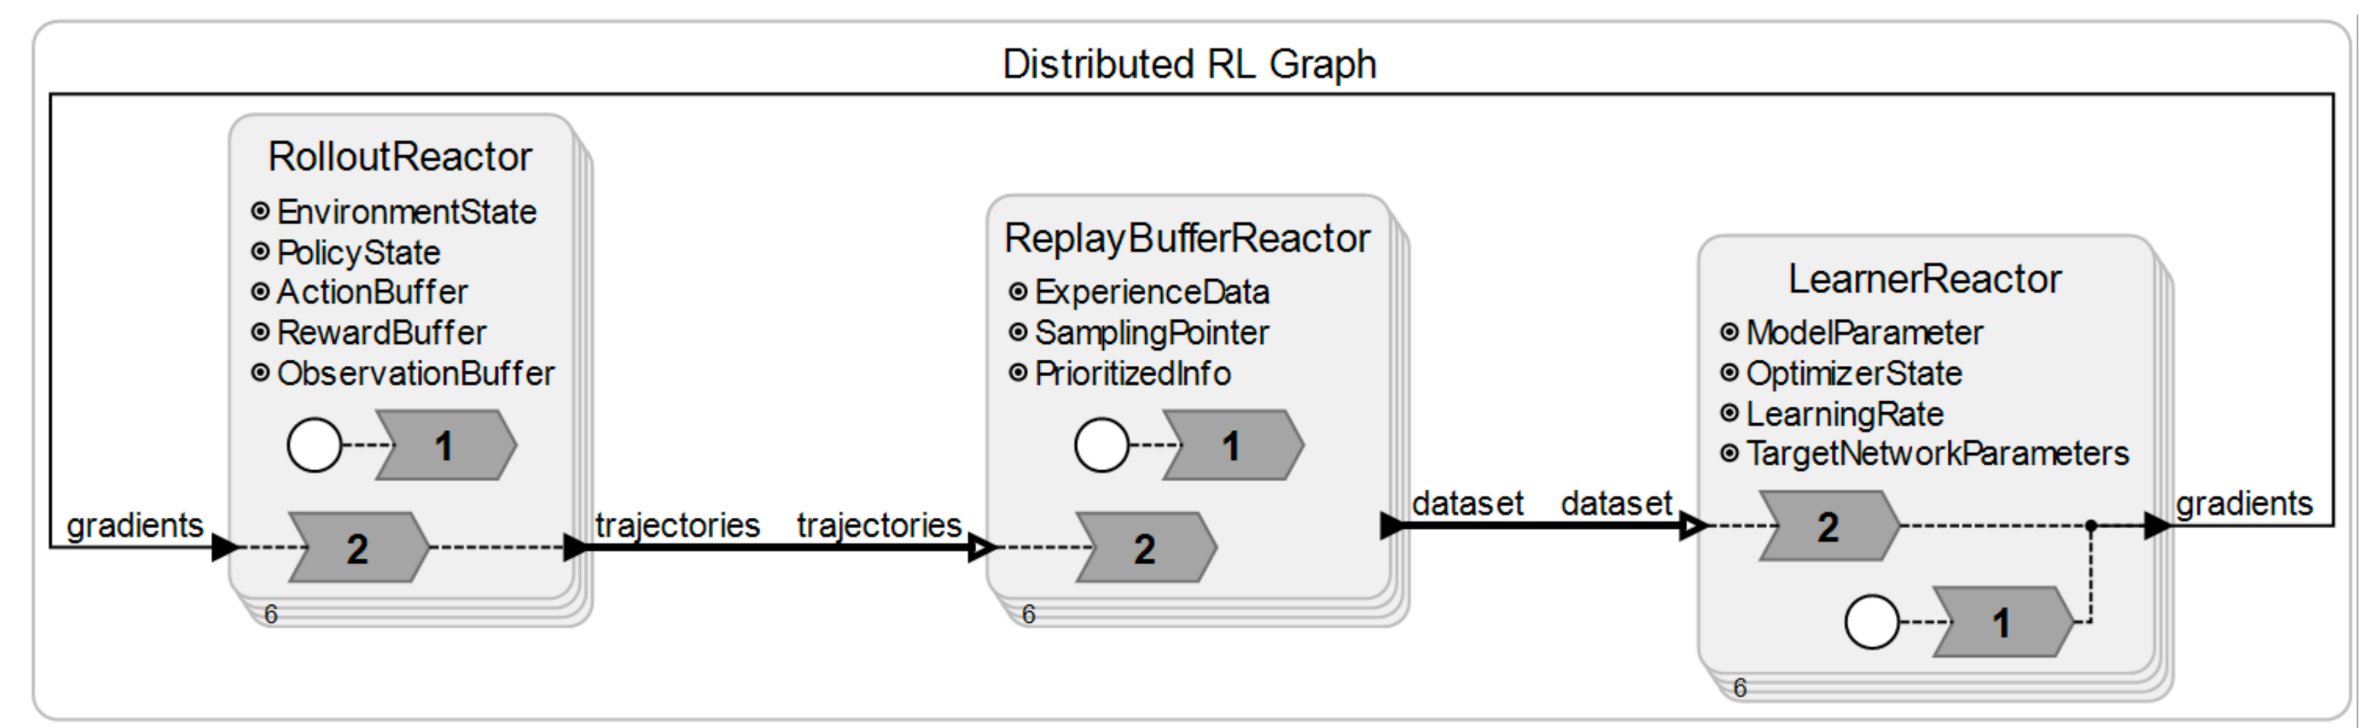

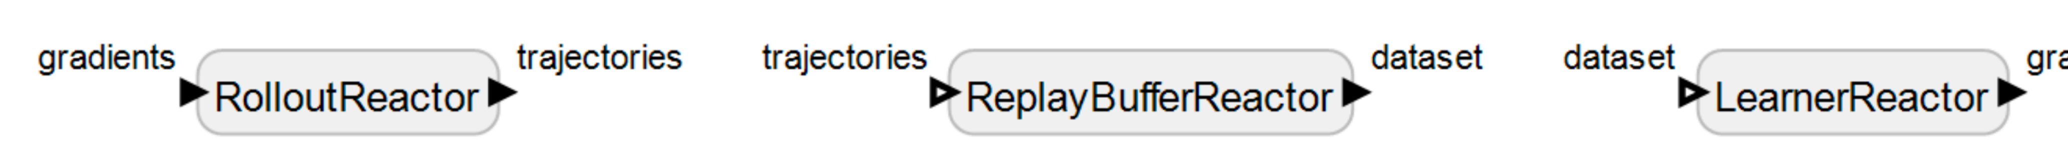

Dataflow graph for distributed RL

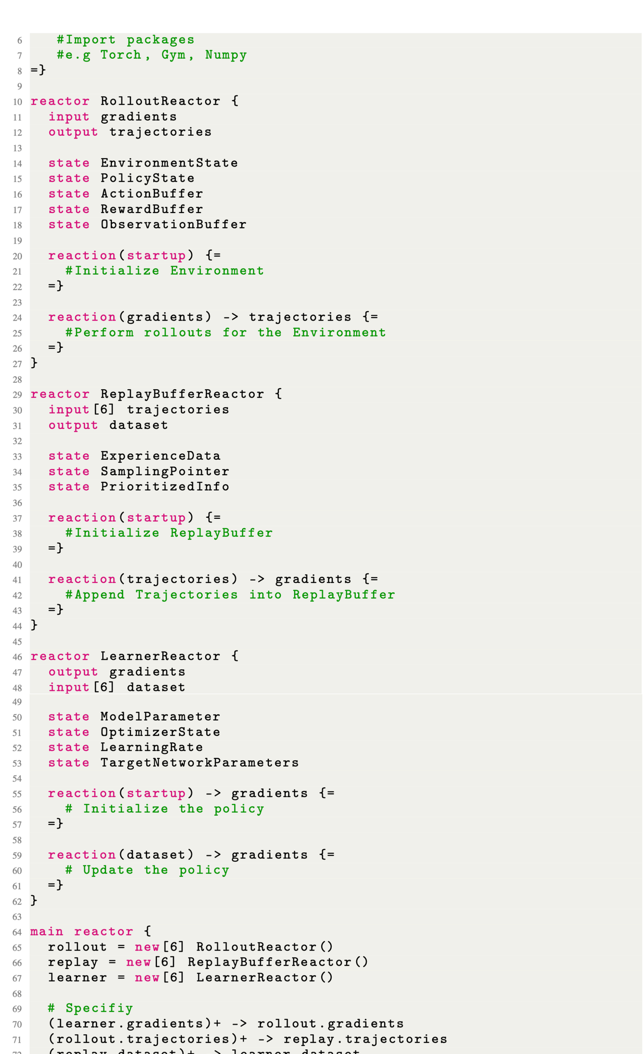

RL in LF

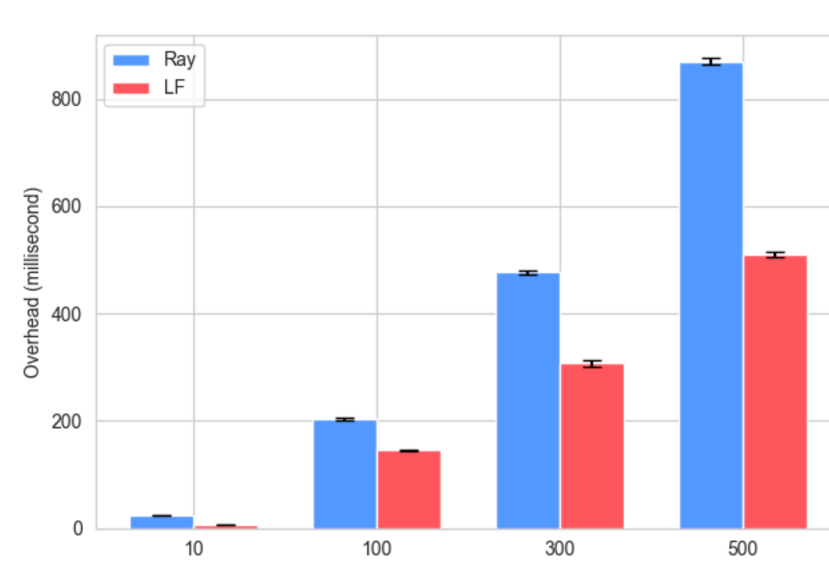

Broadcast & gather Times vs Object Size

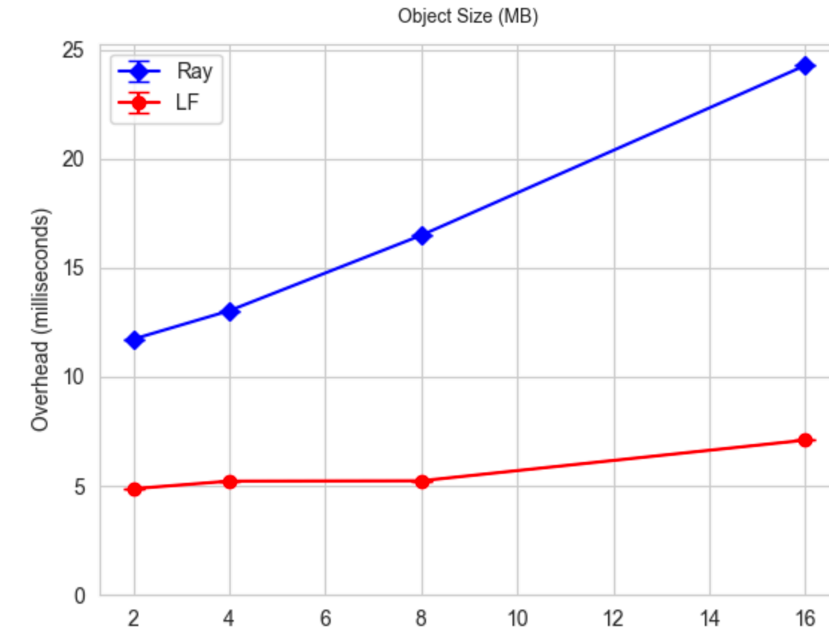

Broadcast & gather Times vs # Actors

Experiments

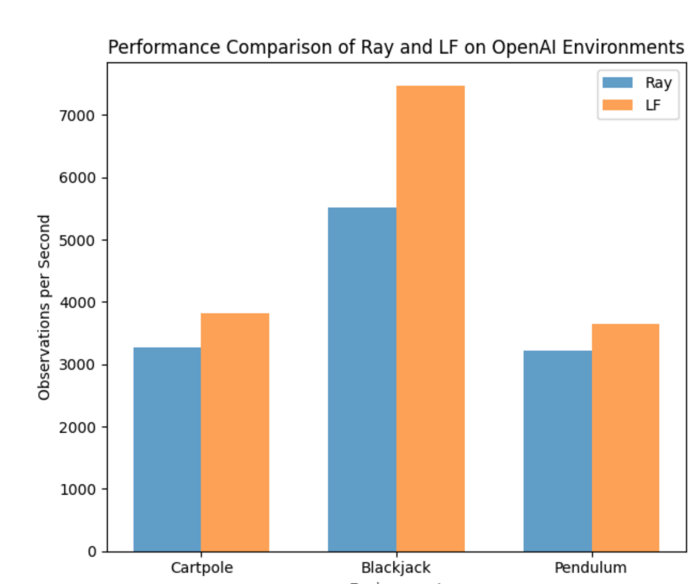

OpenAI Gym Environment Sampling Rates

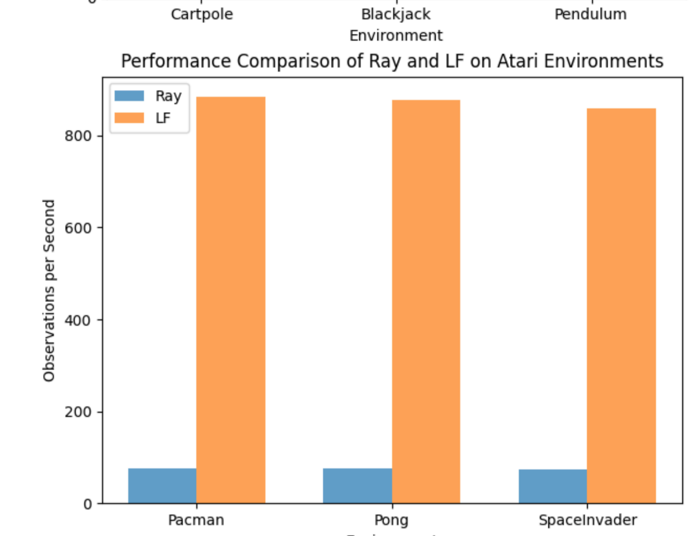

Atari Environment Sampling Rates

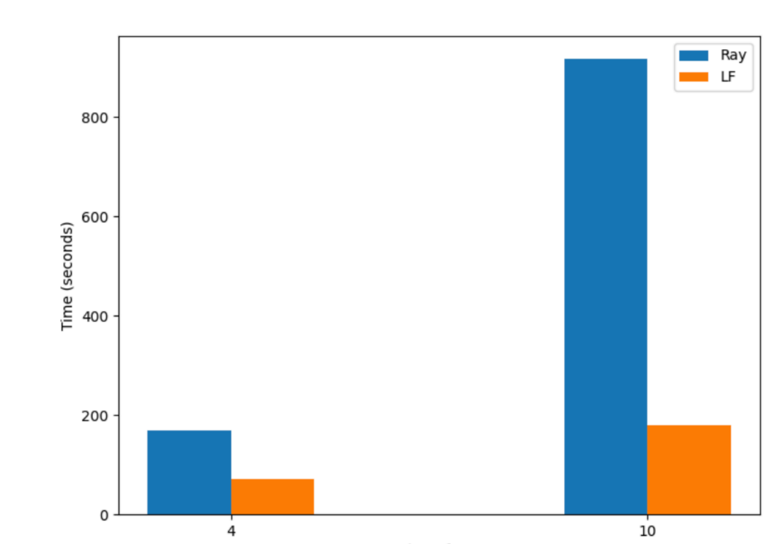

Scaling # agents in multi-agent RL

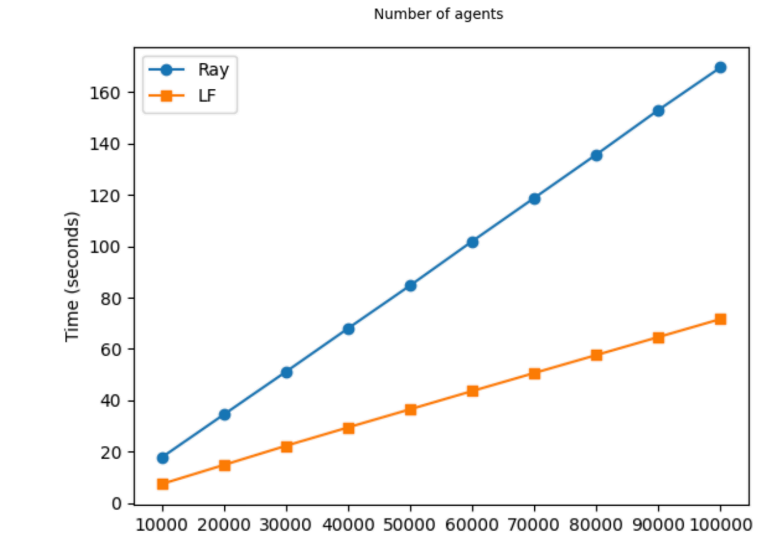

Communication Overhead vs Episodes

[1] MORITZ, P., NISHIHARA, R., WANG, S., TUMANOV, A., LIAW, R., LIANG, E., ELIBOL, M., YANG, Z., PAUL, W., JORDAN, M. I., ET AL. Ray: A distributed framework for emerging {AI} applications. In 13th USENIX symposium on operating systems design and implementation (OSDI 18) (2018), pp. 561–577.

[2] Menard, Christian, et al. "High-performance deterministic concurrency using lingua franca." ACM Transactions on Architecture and Code Optimization 20.4 (2023): 1-29.