

Compact NUMA-aware Locks

Thursday, 23 September 2021 10:30 (30 minutes)

Lock throughput can be increased by handing a lock to a waiter on the same NUMA node as the lock holder, provided care is taken to avoid starvation of waiters on other NUMA nodes. This talk will discuss CNA (compact NUMA-aware lock) as the slow path alternative for the current implementation of qspinlocks in the kernel.

CNA is a NUMA-aware version of the MCS spin-lock. Spinning threads are organized in two queues, a main queue for threads running on the same node as the current lock holder, and a secondary queue for threads running on other nodes. Experimental results with micro and macrobenchmarks confirm that the throughput of a system with contended qspinlocks can increase up to ~3x with CNA, depending on the actual workload.

I agree to abide by the anti-harassment policy

I agree

Primary authors: KOGAN, Alex (Oracle Labs); DICE, Dave (Oracle Labs)

Presenters: KOGAN, Alex (Oracle Labs); DICE, Dave (Oracle Labs)

Session Classification: Performance and Scalability MC

Track Classification: Performance and Scalability MC