Outback: Fast and Communication-efficient Index for Key-Value Store on Disaggregated Memory

Disaggregated memory systems achieve resource utilization efficiency and system scalability by distributing computation and memory resources into distinct pools of nodes. RDMA is an attractive solution to support high-throughput communication between different disaggregated resource pools. However, existing RDMA solutions face a dilemma: one-sided RDMA completely bypasses computation at memory nodes, but its communication takes multiple round trips; two-sided RDMA achieves one-round-trip communication but requires non-trivial computation for index lookups at memory nodes, which violates the principle of disaggregated memory. This work presents Outback, a novel indexing solution for key-value stores with a one-round-trip RDMA-based network that does not incur computation-heavy tasks at memory nodes. Outback is the first to utilize dynamic minimal perfect hashing and separates its index into two components: one memory-efficient and compute-heavy component at compute nodes and the other memory-heavy and compute-efficient component at memory nodes.
We implement a prototype of Outback and evaluate its performance in a public cloud. The experimental results show that Outback achieves higher throughput than both the state-of-the-art one-sided RDMA and two-sided RDMA-based in-memory KVS by 1.06-5.03x, due to the unique strength of applying a separated perfect hashing index.