Achieving High Throughput in Disaggregated KVS with the Lowest Computation in Remote Server

Yi Liu

Abstract:

The RDMA-based Key-Value Stores (KVS) plays an important role to optimize distributed storage systems through direct memory access between nodes. By minimizing latency and maximizing throughput, these systems enhance the efficiency of data retrieval and storage operations in interconnected environments, proving crucial for high-performance applications.
While one-sided RDMA excels in achieving high throughput without involving remote memory nodes' kernels, we observe that the throughput can also be achieved by two-sided RDMA even with limited CPU resources. Furthermore, we disclose the most CPU-consuming part in two-sided RDMA based KVS are lies the local KV lookup.

Based on the above two observations, we developed our solution OUTBACK, which significantly reduces the two-sided RDMA CPU cost on remote memory nodes, ensuring that the utilization of the limited cores remains on par with the efficiency of one-sided RDMA. These optimizations aim to augment the scalability and performance of disaggregated KV stores, thereby contributing to the advancement of more resource-efficient and scalable distributed storage architectures.