SAMS: SmartNIC-Assisted Disaggregated Memory System

Seunghyeok Choi, Youngbin Im

Ulsan National Institute of Science and Technology

choiish98@unist.ac.kr, ybim@unist.ac.kr

SmartNIC 기반 메모리 분리 시스템

최승혁, 임영빈

울산과학기술원

Abstract

We propose the SmartNIC-Assisted Disaggregated Memory System (SAMS), a new system that improves application performance and lowers CPU utilization in Disaggregated memory environments by leveraging SmartNICs. Our evaluation demonstrates that SAMS reduces execution time by up to 74% and host CPU utilization by up to 34% compared with Linux, and by up to 23% and 48%, respectively, compared with Fastswap.

Ⅰ. Introduction

The rapid growth of datacenter workloads has intensified pressure on memory capacity. Memory disaggregation has emerged as a remedy by pooling the memory resources and letting the compute nodes access the memory via low-latency interconnects such as RDMA. However, the host-side page-fault path during remote memory access often becomes a bottleneck that degrades overall system performance.

Ⅱ. Method

To mitigate these bottlenecks, we present SAMS, a SmartNIC-based Disaggregated memory system. Figure 1 illustrates the SAMS architecture.

Path Separation. In Linux kernel, when a page fault occurs, the kernel handles both the faulted page and any prefetch pages together. SAMS separates pages by access urgency: pages needed immediately are served via a synchronous path, while pages that can tolerate latency are handled via an asynchronous path. The synchronous path keeps tail latency low by serving requests on the host. The asynchronous path uses the sNIC (SmartNIC): given host-provided metadata, the sNIC proactively fetches prefetch pages from remote memory and delivers them to the host.

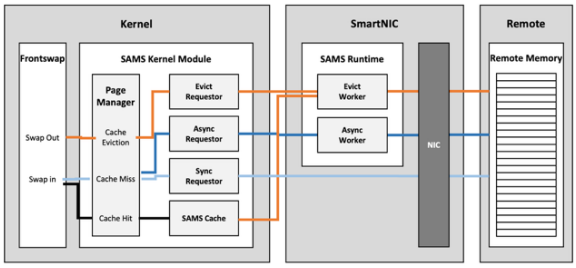

Fig 1. SAMS Architecture.

Efficient Prefetching. The Linux prefetching retrieves pages sequentially after a faulted page. Recent work improves on this by analyzing process-level memory access patterns for prefetching [1]. We extend this idea to the thread level, enabling fine-grained and more accurate prefetching.

Page Manager (PM). Even if a prefetched page has arrived from the sNIC, it may not be visible to the kernel until its virtual address mapping is registered. Upon a page fault, the Page Manager (PM) checks whether the requested page has already been prefetched by the sNIC and staged in host memory. By registering prefetched pages in Page Cache before application access, SAMS reduces CPU load.

Effective Page Eviction. When the kernel issues a swap-out request, the PM first unmaps the page's virtual address and caches the page. When the cache reaches capacity, the host batches eviction requests and sends them to the sNIC. The sNIC reads the page data from the cache in the host and writes it to the designated offset in remote memory, then signals completion to the PM. After receiving the signal, the PM reclaims the cached pages. This pipeline reduces I/O latency from page replacement and alleviates host CPU pressure for page management.

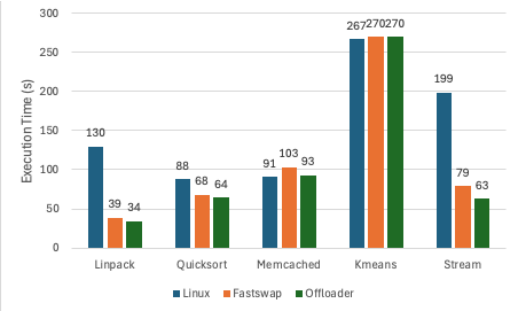

Fig 2. Execution time comparison.

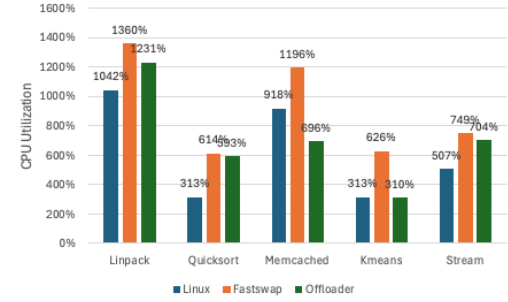

Fig 3. CPU utilization comparison.

Evaluation. We evaluate Linux, Fastswap [2], and SAMS. Our testbed consists of two servers equipped with an Intel Xeon Silver 4215R CPU and 128 GB of DDR4 memory and a 200-Gbps BlueField-2 sNIC. We set one server to use 32 GB of remote memory on the other server. Workloads include Linpack, Quicksort, Memcached, K-means, and STREAM.

With workloads exhibiting largely sequential access patterns, SAMS improves application performance while reducing CPU utilization. For more random-access patterns, SAMS achieves CPU-utilization reductions comparable to or better than Fastswap [2], while maintaining similar performance. For sequential patterns, the sNIC's asynchronous path prefetches contiguous pages with minimal host CPU involvement, thereby reducing host page-fault frequency and freeing CPU cycles for application computation. Batching and cached offloading of page eviction operations onto the sNIC further improve the performance.

Ⅲ. Conclusion

We presented SAMS, a SmartNIC-assisted Disaggregated memory system that lowers host CPU utilization and improves application performance across diverse workloads.

ACKNOWLEDGMENT

This work was partly supported by the Institute of Information & Communications Technology Planning & Evaluation(IITP)-ITRC(Information Technology Research Center) grant funded by the Korea government(MSIT)(IITP-2025-II211817, 25%), the Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (RS-2025-00349594, 25%), the Institute of Information & Communications Technology Planning & Evaluation(IITP) under Next-generation Cloud-native Cellular Network Leadership Program grant funded by the Korea government(MSIT) (RS-2025-00418784, 25%), and the Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (RS-2025-00405128).

REFERENCE

[1] Hasan Al Maruf, "Effectively Prefetching Remote Memory with Leap," Proceedings of the 2020 USENIX Annual Technical Conference, pp. 1-14, July 15-17, 2020.

[2] Emmanuel Amaro, "Can Far Memory Improve Job Throughput?", Proceedings of the Fifteenth European Conference on Computer Systems (EuroSys '20), pp. 1-16, April 27-30, 2020.