

DaeMon: Architectural Support for Efficient Data Movement in Disaggregated Systems

CHRISTINA GIANNOULA, University of Toronto, Canada and National Technical University of Athens, Greece

KAILONG HUANG*, University of Toronto, Canada

JONATHAN TANG*, University of Toronto, Canada

NECTARIOS KOZIRIS, National Technical University of Athens, Greece

GEORGIOS GOUMAS, National Technical University of Athens, Greece

ZESHAN CHISHTI, Intel Corporation, USA

NANDITA VIJAYKUMAR, University of Toronto, Canada

Resource disaggregation offers a cost effective solution to resource scaling, utilization, and failure-handling in data centers by physically separating hardware devices in a server. Servers are architected as pools of processor, memory, and storage devices, organized as independent failure-isolated components interconnected by a high-bandwidth network. A critical challenge, however, is the high performance penalty of accessing data from a remote memory module over the network. Addressing this challenge is difficult as disaggregated systems have high runtime variability in network latencies/bandwidth, and page migration can significantly delay critical path cache line accesses in other pages.

This paper conducts a characterization analysis on different data movement strategies in fully disaggregated systems, evaluates their performance overheads in a variety of workloads, and introduces *DaeMon*, the first software-transparent mechanism to significantly alleviate data movement overheads in fully disaggregated systems. First, to enable scalability to multiple hardware components in the system, we enhance each compute and memory unit with specialized engines that transparently handle data migrations. Second, to achieve high performance and provide robustness across various network, architecture and application characteristics, we implement a synergistic approach of bandwidth partitioning, link compression, decoupled data movement of multiple granularities, and adaptive granularity selection in data movements. We evaluate *DaeMon* in a wide variety of workloads at different network and architecture configurations using a state-of-the-art simulator. *DaeMon* improves system performance and data access costs by 2.39× and 3.06×, respectively, over the widely-adopted approach of moving data at page granularity.

Key Words: data movement, data access, memory access, hardware support, hardware mechanism, high performance, memory systems, memory disaggregation, resource disaggregation, disaggregated systems, workload characterization, benchmarking, performance characterization

1 Introduction

With recent advances in network technologies [33, 42, 64, 80, 81] that enable high bandwidth networks, *resource disaggregation* [33, 88] has emerged as a promising technology for data centers [16, 33, 34, 38, 55, 88, 95]. Resource disaggregation proposes the physical separation of hardware devices (CPU, accelerator, memory, and disk) in a server as independent and failure-isolated components connected over a high-bandwidth network such as RDMA [64] and Gen-Z [81]. Compared to monolithic servers that tightly integrate these components (Figure 1a), *disaggregated systems* can greatly improve *resource utilization*, as memory/storage components can be shared across applications; *resource scaling*, as hardware components can be flexibly added, removed, or upgraded; and *failure handling*, as the entire server does not need to be replaced in the event of a fault in a device. Thus, resource disaggregation can significantly decrease data center costs.

Disaggregated systems comprise multiple compute, memory and storage components, interconnected over a high-bandwidth network (Figure 1b), each independently managed by a specialized

*Equal contribution to this work.

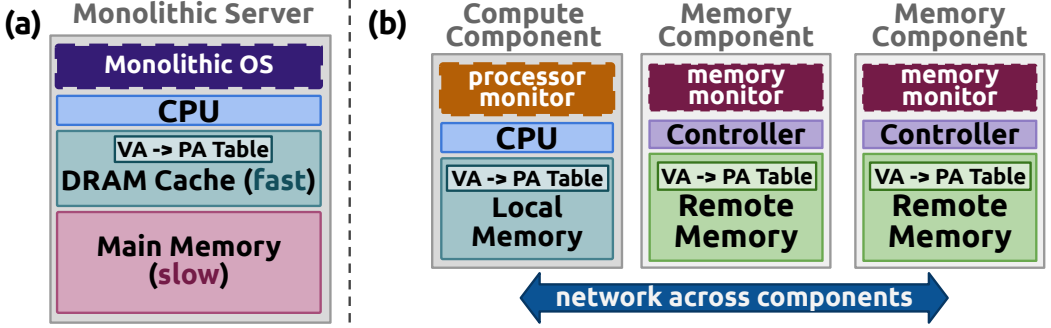


Fig. 1. (a) A monolithic system versus (b) a disaggregated system.

kernel module (monitor). Typically, each **compute component** includes a small amount (a few GBs) of main memory (henceforth referred to as **local memory**) to improve memory performance. However, almost all the memory in the data center is separated as network-attached **disaggregated memory** components to maximize resource sharing and independence in failure handling (different from typical hybrid memory architectures). Thus, the **majority** of the application working sets is accessed from the **disaggregated memory** components (henceforth referred to as **remote memory**). Each **memory component** includes its own controller and can be flexibly shared by many compute components. Thus, disaggregated systems can provide high memory capacity for applications with large working sets (e.g., bioinformatics, graph processing and neural networks) at lower cost. Fine-grain microsecond-latency networking technologies [37, 42, 64, 80, 81] that interconnect all hardware components have made fully disaggregated systems feasible, being only 2-8× slower than DRAM bus bandwidth. However, since a large fraction of the application’s data (typically ~80%) [33, 55, 88] is located and accessed from remote memory, the higher latencies of remotely accessing data over the network can cause large performance penalties.

Alleviating data access overheads is challenging in disaggregated systems for the following reasons. First, disaggregated systems are not monolithic and comprise independently managed entities: each component has its own hardware controller, its resource allocation is transparent from other components and a specialized kernel monitor uses its own functionality/implementation to manage the component it runs on (only communicates with other monitors via network messaging if there is a need to access remote resources). This characteristic necessitates a distributed and disaggregated solution that can scale to a large number of independent components in the system. Second, there is high variability in data access latencies as they depend on the location of the remote memory component and the contention with other compute components that share the same memory components and network. Data placements can also vary during runtime or between multiple executions, since data is dynamically allocated in one or more remote memory components and hardware updates can flexibly change the architecture of the memory component and the network topology. Third, data is typically migrated at page granularity [5, 6, 10, 36, 55, 58, 88, 104, 110] as it enables: (i) transparency to avoid modifications to existing OS/applications; (ii) low metadata overheads for address translation; and (iii) leveraging spatial locality within pages. However, we observe in §2.2 that moving memory pages in disaggregated systems, i.e., moving data at a large granularity over the network, can significantly increase bandwidth consumption and slow down accesses to cache lines in other concurrently accessed pages.

Recent works on hybrid memory systems [3, 4, 23, 26–29, 43, 46, 48, 53, 59–62, 65, 83, 84, 104], for example, those that integrate die-stacked DRAM [45] caches aim to address the high page movement costs between main memory and the DRAM cache [23, 43, 61, 62, 84] with mechanisms to move data

at smaller granularities [23, 44, 61, 62, 77, 97, 98], e.g., cache line, or by using page placement/hot page selection mechanisms [3, 4, 26–29, 46, 48, 53, 59, 60, 65, 83, 104]. However, these prior works are tailored for a monolithic tightly-integrated architecture (Figure 1a), and are not suitable for disaggregated systems (See § 7). These works assume *centralized* data management/allocation (unlike in disaggregated systems). For instance, software runtimes [3, 4, 26–29, 46, 48, 53, 59, 60, 65, 83, 104] running on CPUs in hybrid systems leverage TLBs/page tables to track page hotness and move pages across different memory devices (Figure 1a). Instead, in fully disaggregated systems all hardware memory functionalities (e.g., TLBs, page tables) of remote pages are moved to the memory components themselves [38, 88] (Figure 1b). Thus they cannot be used to track page hotness at the CPU side to implement intelligent page placement/movement in local memory. Similarly, hardware-based approaches [44, 56, 77, 97] add *centralized* hardware units in the CPU to track metadata for pages in second-tier memory. This however would incur high hardware costs in disaggregated systems that enable large amounts (e.g., TBs) of remote memory [88]. Requiring *each* compute component to control/track a large number of pages in remote memory components would impose significant hardware costs and scalability challenges, and thus might annihilate the benefits of resource disaggregation. Moreover, disaggregated systems incur significant variations in access latencies and bandwidth based on the current network architecture and concurrent jobs sharing the memory components/network, which are not addressed by prior work. This necessitates a solution primarily designed for robustness to this variability.

In this work, we analyze different data movement strategies in fully disaggregated systems, and introduce *DaeMon*, an efficient software-transparent mechanism to alleviate data movement overheads in disaggregated systems. *DaeMon* provides (i) high performance on dynamic workload demands, (ii) robustness to variations in architectures, network characteristics and application behavior, and (iii) independence and scalability to multiple compute components and memory components that are managed transparently to each other and are flexibly added/removed in the system.

DaeMon consists of two key ideas. First, we offload data migrations to dedicated hardware engines, named *DaeMon* compute and memory engine, that are added at each compute component and memory component, respectively. This key idea enables independence and scalability to a large number of compute components and memory components of disaggregated systems. Compared to a centralized design, *DaeMon*'s distributed management of data movement enables simultaneous processing of data movement across multiple components and decreases the processing costs and queuing delays to serve data requests. Second, we leverage the synergy of three key techniques to provide robustness to the high variability in network latencies/bandwidth. 1) We use a bandwidth partitioning approach to enable the decoupled movement of data at *two granularities*, i.e., page and cache line, and *prioritize* cache line granularity data moves over page moves. This design enables low access latencies to remote memory for the cache line requests on the critical path, while the associated pages can be still be moved independently at slower rates to retain the benefits of spatial locality. 2) We design an adaptive approach to decide on-the-fly if a request should be served by a cache line, page or *both*. Via selective granularity data movement, we provide robustness to variations in network, architectures and application characteristics. 3) We leverage hardware link compression when migrating pages to reduce network bandwidth consumption and alleviate queuing delays.

The synergy of the aforementioned key techniques provides a robust solution for disaggregated systems: decoupled multiple granularity data movement effectively prioritizes cache line requests on the critical path, and migrates pages at a slower rate leveraging compression to reduce bandwidth consumption. The adaptive granularity selection mechanism effectively adapts to the characteristics of the application data, e.g., by favoring moving more pages if application data is highly compressible.

The decoupled cache line granularity movement also enables the use of more sophisticated and effective compression algorithms (with relatively high compression latency) for page migrations.

We evaluate *DaeMon* using a range of capacity intensive workloads with different memory access patterns from machine learning, high-performance-computing, graph processing, and bioinformatics domains. Over the widely-adopted approach of moving data at page granularity, *DaeMon* decreases memory access latencies by $3.06\times$ on average, and improves system performance by $2.39\times$ on average. We demonstrate that *DaeMon* provides (i) robustness and significant performance benefits on various network/architecture configurations and application behavior (Figures 8 and 13), (ii) scalability to *multiple* hardware components and networks, (Figure 17), and (iii) adaptivity to dynamic workload demands, even when multiple heterogeneous jobs are concurrently executed in the disaggregated system (Figure 18).

This paper makes the following contributions:

- We heavily modify a state-of-the-art simulator to develop and evaluate the overheads of different data movement strategies in fully disaggregated systems, analyze the challenges of providing efficient data movement in such systems, and develop *DaeMon*, an adaptive distributed data movement mechanism for fully disaggregated systems.
- We enable decoupled data movement at two granularities, and migrate the requested critical data quickly at cache line granularity and the corresponding pages opportunistically without stalling critical cache line requests. We dynamically control the data movement granularity to effectively adapt to the current system load and application behavior. We employ a high-latency compression scheme to further reduce bandwidth consumption during page migrations.
- We evaluate *DaeMon* using a wide range of capacity intensive workloads, various architecture/network configurations, and in multi-workload executions of concurrent heterogeneous jobs. We demonstrate that *DaeMon* significantly outperforms the state-of-the-art data movement strategy, and constitutes a robust and scalable approach for data movement in fully disaggregated systems.

2 Background and Motivation

2.1 Baseline Disaggregated System

Figure 2 shows the baseline organization of the disaggregated system, which includes several compute components and memory components as network-attached components. To improve performance, each compute component tightly includes a few GBs of main memory, referred to as **local memory**, which can typically host 20-25% of the application's memory footprint [33, 85, 88]. Each memory component includes its own controller and connects multiple DIMM modules, referred to as **remote memory**.

We assume distributed OS modules that coordinate and communicate with each other via network messaging when needed, similar to [38, 55, 88]: processor and memory kernel monitors run at compute components and memory components, respectively. The memory allocation/management of remote memory is performed at the memory component itself [38, 88], transparently to compute components, enabling the different components to be *independent*. The on-chip caches and the local memory of compute components are typically indexed by virtual addresses [88], and remote data is requested from memory components using virtual addresses [16, 38, 55, 88] (unlike in hybrid memory systems). The data management is typically performed at page granularity [16, 55, 88] (e.g., 4KB). The local memory of the compute component can be treated as a cache with tags [88] or a *local* virtual to physical translation mapping [55] can be used (either approach works with *DaeMon*, however we assume the second approach in our evaluation). The physical memory addresses of the local memory can be found by accessing and traversing metadata (tags or local page tables) kept in a dedicated (pre-reserved) DRAM memory space (kernel metadata is directly indexed via

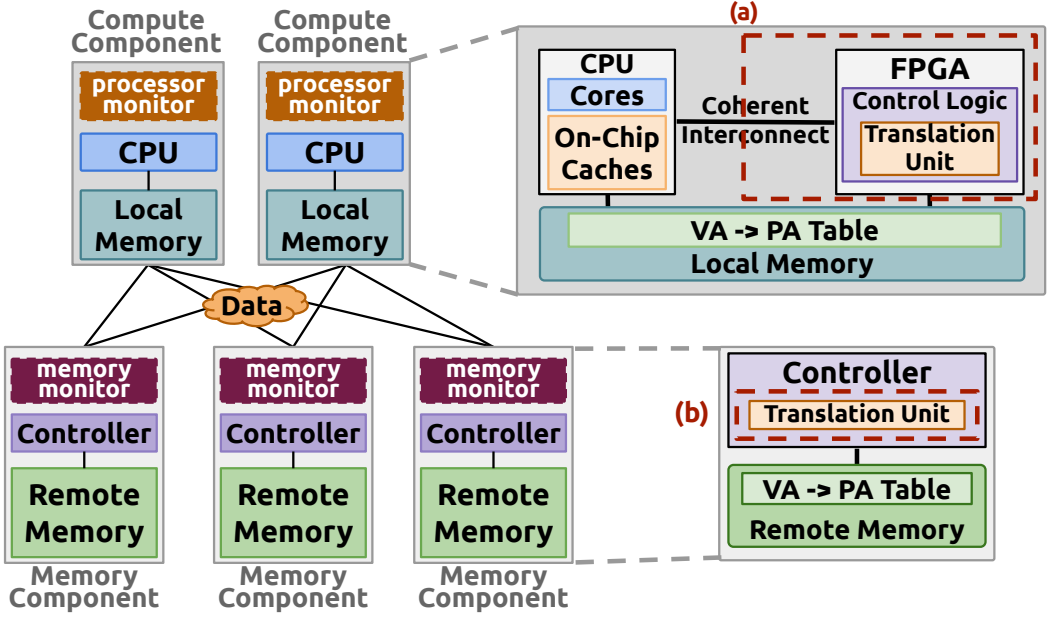


Fig. 2. High-level organization of a disaggregated system.

physical addresses). When an local memory *miss* happens, either (i) the processor kernel module of the compute component triggers a page fault and fetches the requested page from remote memory components [88], or (ii) dedicated software runtimes co-designed with hardware primitives [16] (e.g., supported in FPGA-based controllers as shown in Figure 2a) handle remote data requests on demand completely eliminating expensive page faults. Either approach works with *DaeMon*. We assume that the controllers of memory components implement *hardware-based* address translation (Figure 2b) to access pages in remote memory as proposed in [38].

Jobs running at different compute components can share *read-only* pages located at multiple memory components. Similarly to prior state-of-the-art works [16, 38, 88], we assume that the system does not support writable shared pages across compute components, since they are rare across datacenter jobs [38, 88].

2.2 Data Movement Overheads in Fully Disaggregated Systems

Prior state-of-the-art works [14, 36, 38, 47, 55, 75, 76, 100, 110, 113] typically enable data management at page granularity for three compelling reasons. First, the memory allocation and management is *transparent*, i.e., requires little to no modification to OS or application code. Second, the coarse granularity enables *low metadata overheads* for address translation in local memory and remote memory. Managing local memory as a cache at cache line granularity would incur prohibitively high metadata overheads [88]. Third, page movements enable exploiting *spatial locality* in common memory access patterns [44, 57, 103], and increase the number of accesses served from the lower cost local memory instead of remote memory.

Figure 3 compares the performance of different data movement strategies in disaggregated systems across various workloads (See Table 3). We evaluate one memory component and one compute component having local memory to fit ~20% of the application’s working set. We use 100 ns/400 ns latency [33, 55] to model the propagation and switching delays on the network (referred to as **switch latency**), and configure the network bandwidth between the compute

component and the memory component (referred to as **bandwidth factor**) to be $1/4\times$ the DRAM bus bandwidth [33, 88] of the local memory or remote memory. We compare six configurations: (i) Local: all accesses are served from the local memory; (ii) *cache-line*: accessing data from remote memory at cache line granularity, and directly writing data to the Last Level Cache (LLC) of the compute component (local memory is not used), (iii) Remote: accessing data from remote memory at page granularity (moving pages to local memory) accounting for all network-related overheads, (iv) *page-free*: remote accesses incur the latency of one cache line granularity remote access and the corresponding page is transferred to local memory at *zero cost* (spatial locality is leveraged), (v) *cache-line+page*: requesting data from remote memory at both cache line (moved to LLC) and page granularity (moved to local memory) and servicing data requests using the latency of the packet that arrives earlier to compute component (accounting for all network-related overheads), and (vi) DaeMon: accessing data from remote memory using *DaeMon* (accounting for all network-related overheads).




Fig. 3. Data movement overheads in disaggregated systems.

We make four observations. First, Remote, i.e., the typically-used approach of moving data at page granularity, incurs significant performance slowdowns, on average $3.86\times$, over the monolithic Local configuration due to transferring large amounts of data over the network. In addition to the large network bandwidth consumption, migrating pages can slow down critical path accesses to data in other concurrently accessed pages. Second, *page-free* achieves almost the same performance as the Local scheme. A small penalty is incurred as the first access to a page in remote memory incurs cache line granularity latency access cost. However, since the whole page is migrated to local memory *for free*, performance significantly improves thanks to spatial locality benefits of migrating pages in addition to the requested cache line. Thus, migrating pages to local memory is critical to achieving high performance. Third, *cache-line* outperforms Remote in some latency-bound workloads with poor spatial locality, however its performance benefits depend on network characteristics. For example, in *tr*, *cache-line* outperforms Remote by $1.42\times$ when the switch latency is 100 ns, while it incurs $1.82\times$ performance slowdown over Remote with 400 ns switch latency.

Fourth, the *cache-line+page* scheme, i.e., *naively* moving data at *both* granularities, is still inefficient (only $1.11\times$ better than Remote), since critical cache lines are still queued behind large pages.

Overall, we draw two conclusions. (i) Page migrations incur high performance penalties and can significantly slow down the critical path cache line requests to other concurrently accessed pages. However, if the overheads of migrating pages can be mitigated, moving data at page granularity offers a critical opportunity to alleviate remote access costs. (ii) There is no *one-size-fits-all* granularity in data movements to always perform best across all network configurations and applications. Depending on the spatial locality and the network load, some applications benefit from cache line-only accesses that avoid unnecessary congestion of pages in the network, while some applications significantly benefit from page movements that leverage spatial locality. To this end, we design *DaeMon* to significantly reduce data movement costs across various application, network and architecture characteristics. Figure 3 demonstrates that *DaeMon* significantly outperforms the Remote and *cache-line+page* schemes by on average $2.38\times$ and $2.14\times$, respectively.

3 DaeMon: Our Approach

DaeMon is an adaptive and scalable data movement mechanism for fully disaggregated systems that supports low-overhead page migration, enables software transparency, and provides robustness to variations in memory component placements, network architectures and application behavior. *DaeMon* comprises two key ideas:

(1) Disaggregated Hardware Support for Data Movement Acceleration. We enhance each compute component and memory component with specialized engines, i.e., *DaeMon* compute and memory engine (Figure 4), respectively, to manage data movements across the network of disaggregated systems. *DaeMon* engines enable independence and high scalability to a large number of compute components/memory components that are flexibly added/removed in disaggregated systems. Moreover, distributed management of data migrations at multiple *DaeMon* engines increases the execution parallelism and decreases the processing costs and queuing delays to serve data requests.

(2) Synergy of Three Key Techniques. *DaeMon* incorporates three synergistic key techniques shown in Figure 4:

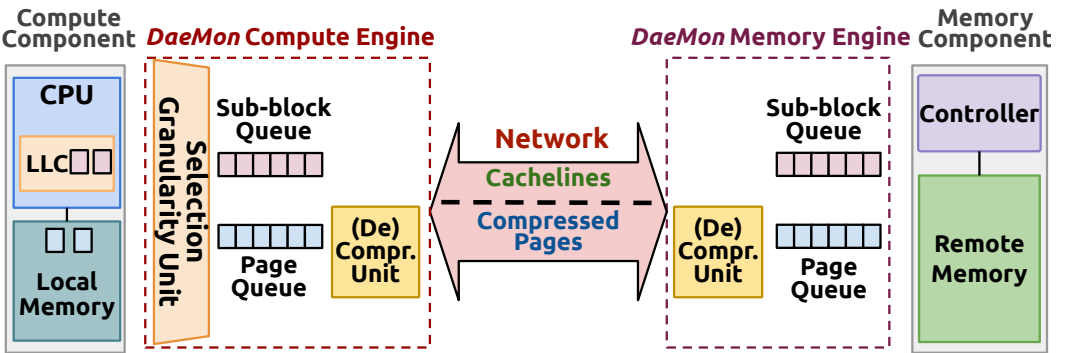


Fig. 4. High-level overview of *DaeMon*.

(I) Decoupled Multiple Granularity Data Movement. First, we integrate two separate hardware queues to manage and serve data requests from remote memory at two granularities, i.e., cache line (via the sub-block queue) and page (via the page queue) granularity. Cache line requests are directly moved to Last Level Cache (LLC) of the compute component to avoid additional metadata overheads

and eliminate memory latency. Page requests are moved to local memory of the compute component. Second, we *prioritize* moving cache lines over moving pages via a bandwidth partitioning approach: a queue controller serves cache line and page requests with a *predefined fixed ratio* to ensure that at any given time a certain fraction of the bandwidth resources is *always allocated* to serve cache line requests quickly. *DaeMon* implements both network and remote memory bus bandwidth partitioning.

This technique provides two benefits. First, retaining page migrations in *DaeMon* (i) enables software-transparency, (ii) allows maintaining metadata for DRAM at page granularity (thus incurring low metadata overheads), and (iii) exploits the performance benefits of data (spatial) locality within pages. Second, cache line data movements that are on the critical path are quickly served, and have fewer slowdowns from expensive page movements that may have been previously triggered, since *DaeMon* effectively prioritizes cache line movements.

(II) Selection Granularity Data Movement. To handle network, architecture and application variability in disaggregated systems, we design a dynamic approach to decide whether a request should be served by a cache line, page, or *both*, depending on application and network characteristics. At *DaeMon*'s engine of each compute component, we include two separate hardware buffers to track *pending* data migrations for both cache line and page granularity, and a selection granularity unit to control the granularity of upcoming data requests based on the utilization of the above buffers. The utilization of these buffers allows us to capture dynamic information regarding the current traffic in the system and the application behavior (i.e., locality). Our proposed selection granularity data movement enables robustness against fluctuations in network, architecture and application characteristics (we explain how this is implemented in §4.2).

(III) Link Compression on Page Movements. We leverage the decoupled page movement to use a high-latency link compression scheme (with high compression ratio), when moving pages across the network. We integrate hardware compression units at both the compute components and memory components to highly compress pages moved over the network: the page is compressed before it is being transferred over the network, and decompressed when it arrives at the destination (before it is written in memory modules). *Link compression* on page movements reduces the network bandwidth consumption and alleviates network bottlenecks.

Overall, *DaeMon* cooperatively integrates all three key techniques, the synergy of which provides a versatile solution:

- (1) Prioritizing requested cache lines helps *DaeMon* to tolerate high (de)compression latencies in page migrations over the network, while also leveraging benefits of page migrations (low metadata overheads, spatial locality).
- (2) Moving compressed pages consumes less network bandwidth, helping *DaeMon* to reserve part of the bandwidth to effectively prioritize critical path cache line accesses.
- (3) Selection granularity movement helps *DaeMon* to adapt to the application data compressibility: if the pages are highly compressible, the number of *pending* page migrations is relatively low, thus *DaeMon* favors moving data more at page granularity instead of cache line granularity (and vice-versa).

4 *DaeMon*: Detailed Design

We design *DaeMon* to be a disaggregated solution: a *DaeMon* compute engine is added at *each* compute component of the system to handle data requests to remote memory, and a *DaeMon* memory engine is integrated at the controller of *each* memory component of the system. Figure 5 shows our proposed architecture.

The baseline architecture of each compute component includes a chiplet-based CPU+FPGA architecture (this CPU+FPGA integrated design has also been proposed to prior state-of-the-art

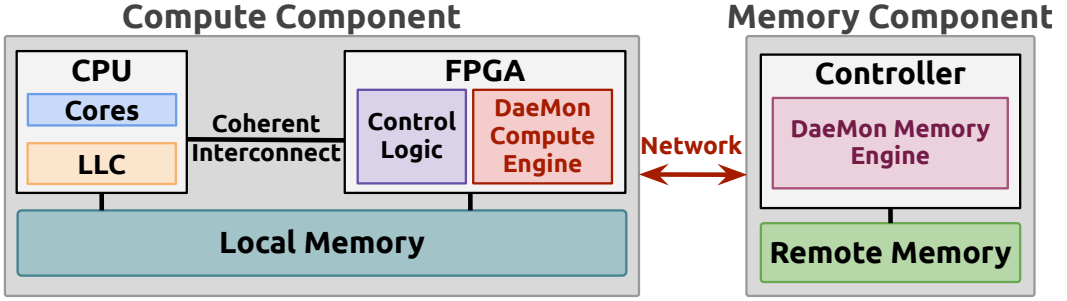


Fig. 5. Proposed architecture for compute component and memory component.

work [16, 41]), which is expected to have small cost [16] compared to the overall cost savings enabled by disaggregated systems, while it is also socket compatible to current systems [25, 41]. The FPGA has three communication paths: i) a coherent path, i.e., CPU-FPGA coherent links, to access the CPU on-chip cache hierarchy, ii) an interface (channel-based connection) to access the local memory, and iii) an external connection to the network controller to move data to/from remote memory. We propose extending the FPGA by adding a new lightweight hardware component to handle data requests, i.e., the *DaeMon* compute engine.

Each memory component includes its own controller [38, 55, 88], that has two communication paths: a channel-based connection to DIMM modules of remote memory, and an external connection to the network, which is used to move data from/to compute components. We propose extending the controller of each memory component by adding a new hardware component to handle data movements, i.e., the *DaeMon* memory engine.

In our study, we assume that the local memory is an inclusive cache for the remote memory, which contains all application data. The local memory implements an approximate LRU replacement policy, similar to prior state-of-the-art work [88].

4.1 Enabling Decoupled Multiple Granularity Data Movement

Figure 6 shows the detailed design of the *DaeMon* compute engine and *DaeMon* memory engine. *DaeMon* engine includes two queues to handle requests at each granularity: cache line granularity via the sub-block queue ①, and page granularity via the page queue ②. It also includes a queue controller ⑦ to serve requests from both queues, and a packet buffer ⑥ to temporarily keep arrived packets, while they are being processed.

Approximate Bandwidth Partitioning. To prioritize cache line data movements while also ensuring that page movements are not aggressively stalled, we design an approximate bandwidth partitioning approach between the cache line and page movements, and configure the queue controller to serve cache line and page requests with a *predefined fixed ratio*. Assuming that cache line and page requests transfer 64B and 4KB of data, respectively, and having a bandwidth partitioning ratio of 25% (Figure 11 presents a sensitivity study on this ratio), 25% of the bandwidth is reserved for cache lines as follows: for each page request issued through the network, which results in transferring 4KB data, the queue controller needs to serve $4096/64 * 0.25/(1 - 0.25) \approx 21$ cache line requests, each transferring 64B of data. To ensure this approximate partitioning is always maintained, we retain this alternate serving of page and cache line requests even if either queue is empty (i.e., requests may be not issued in all cycles). *DaeMon* implements an approximate bandwidth partitioning both in the network across components of the system and when accessing data from remote memory modules.

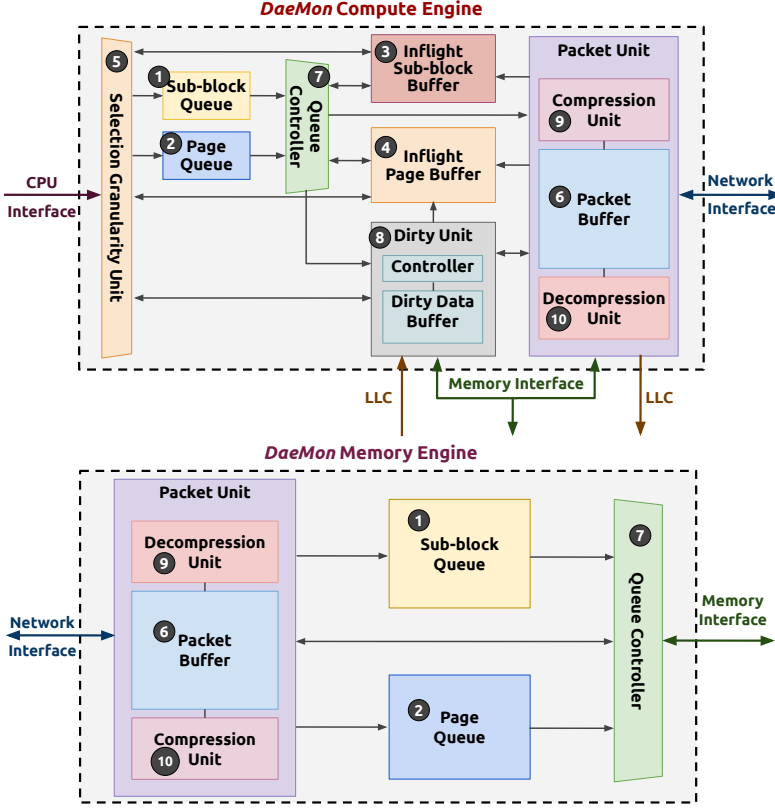


Fig. 6. Detailed design of *DaeMon* engines for the compute (left) and memory (right) components.

4.2 Selecting the Data Movement Granularity

DaeMon compute engine additionally includes two separate hardware buffers to track data requests which are scheduled to be moved or in the process of being migrated (henceforth referred to as *inflight*): (i) the inflight sub-block buffer for the cache line granularity requests **3**, and (ii) the inflight page buffer for the page granularity requests **4**. Both buffers are used to track *pending* data migrations and avoid requesting the same data multiple times. *DaeMon* compute engine includes a selection granularity unit **5** which throttles data requests to avoid requesting the same data multiple times, and decides at which granularity the request should be served (cache line, page, or both granularities).

Scheduling Page Granularity Data Movements. When the *DaeMon* compute engine receives a data request, the selection granularity unit checks (i) the utilization of the inflight page buffer, and (ii) if the corresponding page has already been scheduled to be moved. If the page has already been requested or the inflight page buffer is full, the selection granularity unit *does not request the page*. Thus at any given time, the number of pages scheduled to be moved is automatically limited by the selection granularity unit, also limiting storage/area overheads to track the pending page migrations.

If the inflight page buffer is not full, the selection granularity units schedules the page migration by adding a new entry in the page queue and the inflight page buffer, marking the page as *scheduled*.

When the queue controller issues the movement, the corresponding entry is released in the page queue, and the page entry in the inflight page buffer is marked as *moved*. When the requested page arrives, the corresponding entry is released (*invalid* state) in the inflight page buffer. The page is written to local memory and all pending requests are serviced via local memory. Any entries in the inflight sub-block buffer with requests to cache lines in the same page are removed and thus, any data packets that arrive in the future with cache lines from the same page are simply ignored. In *DaeMon*, we retain existing data management and address translation mechanisms at page granularity. Local page table updates at the compute component are only performed in page migrations.

Scheduling Cache Line Granularity Data Movements. To decide whether a cache line granularity movement should be made, the selection granularity unit checks (i) the utilization of the inflight sub-block buffer and (ii) if the corresponding page was already scheduled to be moved (by a previous request). There are two cases. First, if the corresponding page is *not* scheduled to be moved according to the inflight page buffer, the selection granularity unit always schedules a cache line granularity data movement. Second, if the corresponding page is already scheduled to be moved, the selection granularity unit sends the cache line only if: (i) the sub-block buffer has lower utilization than the page buffer and (ii) the page is not already in the process of migration (i.e., the page is in the page queue). Otherwise, it drops the request as the page has already been requested. This avoids unnecessarily sending cache lines when the corresponding page is likely to arrive faster and when the sub-block queue is likely to be slow due to oversaturation.

If a cache line is scheduled, a new entry is added both in the sub-block queue and the inflight sub-block buffer. When the queue controller issues the movement, the corresponding entry is released in the sub-block queue. When the requested cache line arrives at the compute component, the corresponding entry is released in the sub-block buffer, and the data is *directly* written to LLC through the FPGA-based coherent interconnect.

The above mechanism enables an adaptive approach for the data movement granularity based on the dynamic network/architecture and application characteristics:

- (1) If there is *high locality* within pages, there are fewer pages requested, and the sub-block buffer fills up faster than the page buffer. Thus, *DaeMon* favors issuing pages and throttles cache line requests. If there is *low locality* within pages, the page buffer fills up faster than the sub-block buffer, since cache line requests are served at a higher rate than page requests (e.g., 21:1 cache lines versus pages requests for 25% bandwidth ratio). Thus, *DaeMon* favors issuing cache line movements and throttles page migrations.
- (2) If both the page and sub-block buffers are fully utilized, *DaeMon* detects *bandwidth constrained* scenarios. In bandwidth constrained scenarios, *DaeMon* favors issuing more cache line movements to alleviate bandwidth bottlenecks. When the bottleneck is mitigated, (inflight buffers are not fully utilized), *DaeMon* schedules more page movements to obtain locality benefits.
- (3) Additionally, when using link compression to transfer pages, *DaeMon* is able to adapt to the compressibility of the application data: if the pages are highly compressible, the inflight page buffer empties at a faster rate and thus *DaeMon* favors sending more page migrations (and vice versa).

4.3 Handling Dirty Data

Dirty data (cache lines/pages) is always directly written to remote memory. Data (cache line or page granularity) can be in one of the three states: (i) *local*: when data is cached in on-chip caches (for cache lines) or local memory (for pages), (ii) *remote*: when data is only in remote memory, and (iii) *inflight*: when data is being migrated. With *DaeMon*, data can be present simultaneously in two states: for example, local as a cache line (in the cache hierarchy of the compute component) and

inflight as a page or vice versa. This poses coherence issues if the processor writes to data in the above state.

There are two scenarios: (i) if a page arrives to compute component *before* a prioritized cache line, any modifications to the page may be overwritten by the stale cache line that arrives later, and (ii) if a dirty cache line is evicted from the LLC while the corresponding page is in transit, the modifications would be lost when the page arrives to compute component. As explained, in the (i) scenario, when a page arrives, the corresponding entries in the inflight sub-block buffer with requests to cache lines in the same page are removed and thus, any data packets that arrive in the future with cache lines from the same page are simply ignored. In the (ii) scenario, for every dirty cache line that gets evicted by the LLC and also misses in the local memory, its corresponding page can be either inflight or in remote memory. To ensure correctness, *DaeMon* compute engine first checks if there is an inflight page request in the inflight page buffer. If there is *no* inflight page request (according to the inflight page buffer), the evicted dirty cache line is directly migrated to remote memory. In the other case, we need to retain the dirty cache line until the page arrives. We include a dirty unit **8** in the *DaeMon* compute engine with a *dirty data buffer* that temporarily stores these dirty cache lines. When the corresponding inflight page arrives, the *DaeMon* compute engine flushes the dirty cache line(s) from the dirty buffer to local memory.

Prior works [2, 16] observe that typically a few cache lines (1-8 cache lines) or all cache lines of a page are accessed. Thus, when the evicted dirty cache lines of the same page increase beyond a predefined threshold (e.g., 8 cache lines), the *DaeMon* compute engine flushes all dirty cache lines to remote memory, and marks the corresponding entry for that page in the inflight page buffer as *throttled*. When the inflight page arrives, the *DaeMon* compute engine ignores it, since its entry is in the *throttled* state, and sends a new request for that page to receive the up-to-date data. This enables lower area/storage overheads for the dirty data buffer.

4.4 Link Compression in Page Migrations

Approaches for data compression are typically of two types: (i) latency-optimized compression schemes [8, 21, 74, 105, 106], which optimize/minimize the (de)compression latencies, and (ii) ratio-optimized compression schemes [1, 50, 94, 112], which provide higher compression ratios while incurring relatively high (de)compression latencies. We select a ratio-optimized compression scheme in *DaeMon* based on two observations (§6): (i) in disaggregated systems, queueing delays and network latencies can be significant, thus compression benefits outweigh the high (de)compression latencies, and (ii) *DaeMon* prioritizes cache lines that are on the critical path, thus we can tolerate relatively high (de)compression latencies for page migrations.

DaeMon engines include (de)compression units **9** **10** that compress pages transferred through the network. We implement a hardware design similar to IBM MXT [1, 94], using the LZ77 compression algorithm [112], and operating at 1KB granularity at a time. (De)Compression units include 4 engines, each of which operates on 256B of data and uses a 256B shared dictionary, incurring in total a 64-cycle latency according to [1, 94].

4.5 *DaeMon*'s Hardware Structures

We estimate the overheads of *DaeMon*'s hardware structures for each compute component assuming a 64-core CPU, using CACTI [67]. The sizes of the *DaeMon* sub-block and page queues and the sub-block and page buffers have been selected based on the maximum possible number of *pending* data migrations at a time, which is determined by the number of the available LLC MSHRs (Miss Status Holding Registers) in a typical CPU system, and is independent of the workloads' patterns and the mix of workloads that are running at each time. For the hardware structures at each

memory component, we scale the sizes of the *DaeMon* sub-block and page queues, assuming that each memory component can concurrently serve up to 4 compute components. Table 1 presents the hardware overheads of *DaeMon* compute engine (C) and *DaeMon* memory engine (M). Figure 7 shows an inflight sub-block buffer entry, an inflight page buffer entry, and a dirty data buffer entry.

Hardware Structure	Entries	Size (KB)	Access Cost (ns)	Area Cost (mm ²)	Energy Cost (nJ)
Sub-block Queue (C)	128	0.5	0.34	0.084	0.038
Sub-block Queue (M)	512	2	0.38	0.093	0.039
Page Queue (C)	256	1	0.35	0.087	0.038
Page Queue (M)	1024	4	0.40	0.105	0.041
Inflight Sub-block Buffer (C)	128	1.625	0.56	0.041	0.046
Inflight Page Buffer (C)	256	3.25	0.77	0.089	0.096
Dirty Data Buffer (C)	256	17	0.62	0.168	0.046
Packet Buffer (C)	-	8	0.538	0.137	0.044
Packet Buffer (M)	-	32	1.032	0.263	0.047
2 × Dictionary Table (C,M)	1024	1	0.28	0.015	0.020

Table 1. *DaeMon*'s hardware overheads for C: compute engine and M: memory engine.




Fig. 7. An inflight sub-block buffer entry, an inflight page buffer entry, and a dirty data buffer entry.

Sub-block Queue (SRAM), 128 entries: The sub-block queue size is limited by the available LLC MSHRs of the compute component.

Page Queue (SRAM) - 256 entries: The page queue has 256 entries, since *DaeMon* serves requests from the page queue at a smaller rate than the sub-block queue.

Inflight Sub-block Buffer (CAM) - 128 entries: Similar to the sub-block queue, this buffer has 128 entries. We design this hardware structure to be indexed using the corresponding page address to achieve smaller area costs, since at a given time there may be multiple inflight cache line requests to the same page. Each entry (Figure 7a) includes the page address, the state (*scheduled* or *invalid*), and a 64-bit queue that is used to indicate the offsets within the page of the inflight cache requests by (re)setting the corresponding bits.

Inflight Page Buffer (CAM) - 256 entries: An inflight page buffer entry (Figure 7b) includes the page address, the state that can be *scheduled*, *moved*, *throttled* (when the page needs to be re-requested) or *invalid*, and a 64-bit queue to indicate the offsets of the dirty cache lines of the inflight page that are temporarily kept in the dirty data buffer.

Dirty Data Buffer (SRAM) - 256 entries: A dirty data buffer entry (Figure 7c) includes the evicted cache line and its address.

Packet Buffer (SRAM) - 8KB: We use an 8KB buffer to temporarily store arrived data packets

until they are processed.

Dictionary Tables for (De)Compression (CAM) - 2KB: *DaeMon* proposes 4 engines at each (de)compression unit, each of them has 256B CAM [1, 94]. In total, we estimate each dictionary table as 1KB CAM.

Overall, *DaeMon*'s hardware overheads are due to the cache memories corresponding to the sub-block and page queues, the sub-block and page buffers, and the dictionary tables used for data compression. The total sizes of the *DaeMon* cache memories are ~34KB and 40KB for the *DaeMon* compute and memory engine, respectively. Therefore, *DaeMon*'s hardware overheads are similar to that of the small L1 cache memory of a modern state-of-the-art processor (e.g., Intel Xeon). We conclude that our proposed hardware structures incur very modest hardware and financial costs to be integrated into the compute components and memory components of disaggregated systems.

4.6 Handling Failures

DaeMon handles compute component, memory component and network failures using fault-tolerance approaches of prior works [16, 55, 88]. If the compute component fails (CPU or *DaeMon* compute engine), the application needs to be restarted potentially on a different compute component of the system. Network failures are handled using timeouts: *DaeMon* engines can trigger timeouts when pending page or cache line requests have not arrived after a long time, or when ACK messages have not been received for migrations of dirty data. The exploration of the timeout period value is left for future work. Finally, memory component failures are handled via data replication, similarly to prior work [88]: *DaeMon* can send the evicted dirty data to more than one memory component, and wait to receive ACK messages from all of them.

4.7 *DaeMon* Extensions

Prefetching. *DaeMon* can flexibly support hardware/software-based prefetchers. Existing CPU prefetchers might generate data requests, which *DaeMon* can normally serve by migrating the prefetched data at a cache line granularity, page granularity or both granularities, via on our proposed selection granularity scheme. Page prefetchers [63] might generate page-granularity data requests, which *DaeMon* can serve by migrating the prefetched data at page granularity or throttling the page request based on our selection granularity scheme.

Large Pages. *DaeMon* can be easily extended to support large granularity pages (e.g., 2MB). To effectively prioritize cache line requests over page requests, *DaeMon*'s predefined ratio for the approximate bandwidth partitioning needs to be properly configured based on the size of the large page. To enable multiple page sizes (e.g., both 4KB and 2MB), we could enhance *DaeMon* to split large pages (e.g., 2MB) to consecutive page requests of smaller sizes (e.g., 4KB) issued in the page queue.

5 Methodology

Simulation Methodology. We use Sniper [17, 18], a state-of-the-art accurate simulator, and we heavily modified it to model a disaggregated system with one compute component and multiple memory components interconnected across the network. We present detailed evaluation results using one memory component and provide a characterization study of multiple memory components with various network configurations in Figure 17. For the network across components, we use both (i) a fixed latency of 100 ns/400 ns [33, 55] to model propagation and switching delays inside network (referred to as *switch latency*), and (ii) a variable latency of modeling the current bandwidth utilization at each simulation interval (100K ns) when configuring the network bandwidth to be 2-8× less than DRAM bandwidth [33, 88] (referred to as *bandwidth factor*). For the compute component, we configure a state-of-the-art CPU server with on-chip cache memories of typical sizes and

x86 OoO cores of 3.6GHz frequency. The local memory size is configured to fit $\sim 20\%$ of each application’s working set, and we evaluate LRU replacement policy [88] in local memory, unless otherwise stated. The aforementioned configuration is consistent with prior state-of-the-art works in disaggregated systems [33, 55, 88]. For both the local memory and remote memory, we evaluate a DDR4 memory model with 17GB/s bus bandwidth, and we simulate hardware-based address translation for memory pages, having overhead as one DRAM access cost per lookup, as explained in prior state-of-the-art work [38]. We evaluate access overheads in *DaeMon* queues/buffers using CACTI [67] (See Table 1). Table 2 lists the parameters of our simulated system.

CPU	3.6 GHz, 4-way OoO x86 cores, 224-entry ROB;
L1 Instr. Cache	32 KB, 4-way associativity, LRU;
L1 Data Cache	32 KB, 8-way associativity, 4-cycle access latency, LRU;
L2 Cache	256 KB, 8-way associativity, 8-cycle access latency, LRU;
LLC	4MB, 16-way associativity, 30-cycle access latency, LRU;
Local Memory	2400MHz, 15ns process. latency, 17GB/s bus bandwidth [13];
Network	2-8 \times less than bus bandwidth, 100-400 ns switching latency [33, 88];
Remote Memory	2400MHz, 15ns process. latency, 17GB/s bus bandwidth [13];

Table 2. Configuration of simulated system.

Workloads. We evaluate various workloads with different memory access patterns from various application domains including graph processing, machine learning, bioinformatics, linear algebra, data analytics, and HPC domains, shown in Table 3. The dynamic working sets at any given point at runtime range from 43.2MB to 1.32GB. In a fully disaggregated system, the application working set (irrespective of the size) is primarily housed in remote memory to provide the benefits of improved elasticity, heterogeneity, and failure isolation. Therefore, we configure the local memory size to fit $\sim 20\%$ of each application’s working set (similar to prior state-of-the-art work [33, 55, 88]). All data is initially located in remote memory. We simulate most workloads to full execution and for slower long running workloads, we simulate 1B instructions.

Workload	Domain	Input Data
K-Core Decomposition (kc) [89]	Graph Processing	1M vertices x 10M edges
Triangle Counting (tr) [89]	Graph Processing	1M vertices x 10M edges
Page Rank (pr) [89]	Graph Processing	1M vertices x 10M edges
Needle Wunsch (nw) [20]	Bioinformatics	4096 base pairs per sequence
Breath First Search (bf) [89]	Graph Processing	1M vertices x 10M edges
Betweenness Centrality (bc) [89]	Graph Processing	1M vertices x 10M edges
Timeseries (ts) [107]	Data Analytics	262144 elements in sequence
Sparse Matrix Vector Multipl. (sp) [52]	Linear Algebra	pkustk14 matrix
Sparse Lengths Sum (sl) [68]	Machine Learning	Kaggle Criteo 10GB Dataset
High Perf. Conjugate Gradient (hp) [40]	HPC	104 x 104 x 104
Particle Filter (pf) [20]	HPC	4096 x 4096, 30000 particles
Darknet19 (dr) [82]	Machine Learning	dog.jpg (768 x 576 pixels)
Resnet50 (rs) [82]	Machine Learning	dog.jpg (768 x 576 pixels)

Table 3. Summary of workloads.

6 Evaluation

We evaluate six schemes: (i) **Remote**: the typically-used approach [16, 55, 88] of moving data to/from remote memory at page granularity; (ii) **LC**: *DaeMon*'s link compression for page movement without enabling cache line granularity data movement, i.e., moving data at page granularity with LZ-based link compression enabled; (iii) **BP**: enabling only *DaeMon*'s decoupled multiple granularity data movement with 25% bandwidth partitioning ratio for cache line movements, i.e., moving data *always* at both granularities; (iv) **PQ**: enabling *DaeMon*'s *both* decoupled multiple granularity and selection granularity data movement with 25% bandwidth partitioning ratio for cache line movements (without enabling data compression in page migrations); (v) **DaeMon**: *DaeMon*'s complete design enabling all its three techniques (25% bandwidth partitioning ratio); and (vi) **Local**: the monolithic approach where all the data fits in local memory of the compute component.

6.1 Performance

Figure 8 compares all schemes with different network configurations. Our evaluated workloads exhibit three patterns and we make the following observations.

First, *kc*, *tr*, *pr*, and *nw* exhibit relatively poor spatial locality within pages. In such workloads, BP effectively prioritizes critical cache line requests. However, PQ provides significant benefits thanks to dynamically selecting the data movement granularity: the page buffer saturates faster than the sub-block buffer given the poor locality and the higher servicing rate of the cache line requests in the queue controller, thus the selection granularity unit enables the movement of more cache lines and fewer pages. This results in reduced access latencies as critical path cache line requests are no longer stalled behind many page migrations.

Second, *bf*, *bc*, and *ts* exhibit medium spatial locality within pages. In such workloads, both LC and PQ decrease data access costs using different approaches: LC enables exploiting more spatial locality by moving more pages, while PQ accelerates accesses to the critical path cache line requests, both of which benefit these workloads.

Third, the remaining workloads exhibit high spatial locality within pages, thus page migration is critical to leverage data locality. In these workloads, BP incurs high performance slowdowns, since it is *oblivious* to application behavior. Instead, PQ effectively enables more page movements and throttles cache line movements by tracking pending data requests, thus achieving similar system performance to Remote. LC performs better for *sp*, *sl*, *hp*, and *pf*, since these workloads have higher data compressibility than *dr* and *rs*.

Fourth, when network bandwidth is more constrained, LC provides *even higher* performance over Remote, while PQ is unaffected by bandwidth as the bandwidth partitioning approach prioritizes cache line movements even with low available bandwidth.

Fifth, PQ is slightly affected by the switch latency (Please also see Figure 20 in Appendix § A): PQ outperforms Remote by 1.60× and 1.51× for 100 ns and 400 ns switch latency, respectively. The slightly lower benefits are due to PQ's inability to hide network switch latencies in critical cache line movements. Instead, LC is unaffected by switch latency, as page movement incurs much higher overheads (due to very high network processing and queueing delays) over the smaller switch latency, which link compression is able to alleviate.

Finally, *DaeMon* provides high performance benefits for all three classes of workloads with different locality characteristics thanks to synergistically integrating both LC and PQ: (i) PQ helps hide the (de)compression latencies in LC and migrate fewer pages in order to prioritize critical path cache line movements, and (ii) LC releases network bandwidth resources and helps recover the lost spatial locality in pages by moving more pages with the available network bandwidth. *dr* and *rs* show only 1.05× speedup over Remote as neither LC nor PQ is able to provide speedups due to the




Fig. 8. Speedup in all workloads normalized to Remote using various network configurations.

poor application data compressibility and high spatial locality within pages (which favors moving pages rather than cache lines). *DaeMon*'s adaptive approach also provides high performance benefits across all network configurations: (i) when the switch latencies are high, cache lines movements are slowed down and the sub-block queue fills up faster, thus *DaeMon* favors moving more pages, which

is more effective at high network switch latencies; and (ii) the approximate bandwidth partitioning approach effectively prioritizes cache line over page movements even when network bandwidth is constrained. Therefore, *DaeMon* significantly outperforms the state-of-the-art Remote scheme by $1.85\times$, $2.36\times$, $2.97\times$ for $1/2$, $1/4$, and $1/8$ bandwidth factor, respectively.

Overall, we conclude that *DaeMon*'s cooperative techniques provide a robust approach to alleviate data movement overheads across various network characteristics and application behavior.

6.2 Memory Access Costs

Figure 9 compares the average data access costs (latencies) achieved by various schemes normalized to Remote. Due to space limitations, in the remaining plots, we present a *representative* subset of our evaluated workloads, but we report geometric mean values across *all* evaluated workloads. Please also see Figure 19 in Appendix § A, which compares the network bandwidth utilization achieved by the various data movement schemes.




Fig. 9. Data access costs achieved by various schemes normalized to Remote.

We make three observations. First, LC improves data access costs over Remote by $2.12\times$ across all network configurations (not graphed), because it reduces the network processing costs and queueing delays by sending fewer bytes through the network. PQ improves access costs ($2.06\times$ over Remote across all configurations) by prioritizing critical path cache line movements. Second, PQ *significantly* reduces data access costs in workloads with poor page locality (e.g., *pr*, *nw*), since critical path cache line movements are not stalled by migrating pages. However, in applications with high data locality (e.g., *dr*, *rs*), although PQ reduces data access costs by $1.43\times$ over Remote, it improves performance by only $1.05\times$, because the selection granularity unit favors sending pages for workloads with high locality and a few requests are served at cache line granularity. Third, *DaeMon* significantly reduces data access costs by $3.06\times$ over Remote. *DaeMon* employs link compression to migrate more pages with lower network overhead over PQ, thus exploiting more data locality, while also leveraging the ability to prioritize critical cache line requests. In *pr*, *DaeMon* can achieve lower access latency than Local, since serving requests from both local memory and remote memory increases the effective aggregate memory bandwidth.

6.3 Hit Ratio in Local Memory

Figure 10 presents the hit ratio in local memory, and is thus a measure of the page movement benefits. To prioritize cache lines, PQ throttles some page migrations, thus reducing the local memory hit rate as a tradeoff for reduced access latencies to critical path cache line requests. However, *DaeMon* enables moving more pages over PQ thanks to link compression, while still retaining the cache line prioritization benefits of PQ. The numbers shown over each bar for *DaeMon* present the additional pages that were moved in *DaeMon* as a percentage over PQ, thanks to the reduced bandwidth consumption provided by link compression. A zero value indicates that neither PQ nor *DaeMon* has throttled any page movement.




Fig. 10. Hit ratio in local memory achieved by various schemes.

We draw three findings. First, Remote has on average 97.7% hit ratio in local memory. Thanks to high spatial locality, all workloads benefit from page migration, leading to high hit rates: even workloads with relatively poor spatial locality (e.g., *nw*) have 90% hit ratio in local memory. Second, PQ decreases the hit ratio in local memory by up to 18.4% over Remote, because PQ throttles page movements in some workloads to prioritize cache line requests, thus increasing the number of accesses to remote memory. Third, *DaeMon* recovers most of the lost local memory hits, achieving on average *only* 0.4% worse hit ratio over Remote. Leveraging link compression in *DaeMon* reduces network bandwidth consumption and significantly increases the number of pages that can be migrated over PQ. Across all configurations (not graphed) *DaeMon* migrates 68.9% of the pages throttled by PQ via leveraging link compression. We conclude that *DaeMon* enables both leveraging the benefits of data locality within pages and the prioritization of critical path cache line requests.

6.4 Sensitivity Study to Bandwidth Partitioning Ratio

Figure 11 presents a sensitivity study on the bandwidth partitioning ratio between the cache line and page movements.

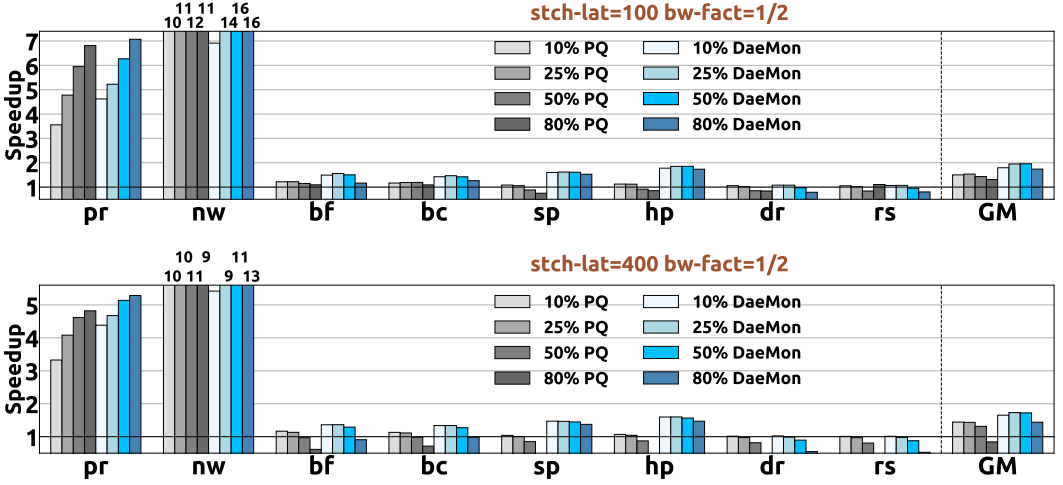


Fig. 11. Performance of PQ and *DaeMon* normalized to Remote varying the bandwidth partitioning ratio.

We draw three findings. First, a higher bandwidth partitioning ratio (e.g., 50%) than *DaeMon*'s default 25% ratio, incurs slowdowns in workloads of medium and high spatial locality, and only improves performance in workloads with very low locality within pages (e.g., *pr*, *nw*). This is because high bandwidth partitioning ratios favor cache line movements and throttle a higher number of page movements. Second, since cache line data movements are affected more by the switch latency compared to page movements, the performance benefits of higher bandwidth partitioning ratios reduce at higher switch latencies. For example, in *pr*, the 50% bandwidth partitioning ratio outperforms 25% ratio by $1.19\times$ and $1.08\times$ using *DaeMon* at 100 ns and 400 ns switch latency, respectively. Finally, across all different bandwidth factors (not graphed), *DaeMon*'s default 25% ratio outperforms the 50% ratio by $1.02\times$ and $1.04\times$ for 100 ns and 400 ns switch latency, respectively, and the 80% ratio by $1.07\times$ and $1.33\times$ for 100 ns and 400 ns switch latency, respectively. We conclude that *DaeMon*'s default 25% ratio on average performs best across all various network and application characteristics.

6.5 Sensitivity Study to Various Compression Algorithms

Figure 12 compares the performance of LC normalized to Remote with three compression schemes: (i) *fpcbdi*: a latency-optimized hybrid scheme of BDI [74] and FPC [8] with 4-cycle (de)compression latency per cache line [50]; (ii) *fve*: the latency-optimized FVE [92] scheme using a 256B dictionary table and having 6-cycle (de)compression latency per cache line [92]; and (iii) *LZ*: *DaeMon*'s compression ratio-optimized LZ-based scheme [50, 94] (See details on § 4.4).

We observe that *LZ* always outperforms Remote, despite the high (de)compression latencies, because the network overheads are significantly higher, indicating that link compression is a highly effective solution for disaggregated systems. *dr* and *rs* show little performance improvement with *LZ*, because the application data is less compressible (their compression ratio is $1.42\times$ versus $4.47\times$ on average across all evaluated workloads). Moreover, *LZ* outperforms *fpcbdi* and *fve* across all network configurations (not graphed) by $1.54\times$ and $1.44\times$ on average, respectively, since it achieves higher compression ratios (on average $2.92\times$ and $2.73\times$ higher compression ratio than *fpcbdi* and *fve* respectively). The benefits of *LZ* over *fpcbdi* and *fve* are even higher in the more bandwidth limited configurations (e.g., with $1/8$ bandwidth factor). Therefore, we conclude that the high network overheads in disaggregated systems favor compression algorithms that provide higher




Fig. 12. Performance of LC varying the compression scheme.

compression ratios, since the benefits of the reduced bandwidth consumption outweigh the higher (de)compression latencies.

6.6 Network Disturbance Study

Figures 13 and 14 compare the IPC and the hit ratio in local memory respectively, of LC, PQ and *DaeMon*, when the network traffic varies during runtime: we simulate contention from other compute components that share the same network, by artificially injecting packets inside the network. We evaluate *pr* and *nw*, as they incur the highest data movement costs.

DaeMon outperforms both LC and PQ by 2.85× and 1.19×, respectively, even when network traffic varies during runtime. *DaeMon* effectively adapts to varying application behavior and network conditions at runtime. For example, in *nw*, in the first 50M instructions, *DaeMon* benefits more from LC as the application has high bandwidth consumption and higher locality within pages. In the next 100M instructions, the workload exhibits less data locality within pages, and *DaeMon* benefits more from PQ, which provides significant performance benefits over LC by effectively prioritizing critical path cache line requests. In the last part of execution, *DaeMon* again leverages the benefits of LC. Therefore, we conclude that *DaeMon* provides a versatile approach to dynamic and variable runtime application and network characteristics.

6.7 Multithreaded Performance

Figure 15 shows *DaeMon*'s performance benefits for multithreaded workloads on 8 OoO cores, thus evaluating more bandwidth-limited executions compared to that of Figure 8. Please also see Figure 21 in Appendix § A which evaluates even more bandwidth-limited executions. Across *all* workloads and network configurations (not graphed), *DaeMon* outperforms the typically-used Remote scheme by 2.73× on average. When network bandwidth is very limited, e.g., 1/16 bandwidth factor (Figure 21), *DaeMon*'s benefits are even higher, by 3.95× over Remote.

6.8 Sensitivity Study to Replacement Policy in Local Memory

Figure 16 compares *DaeMon* and Local normalized to Remote, when using First-In-First-Out (FIFO) replacement policy in local memory.




Fig. 13. Performance of LC, PQ, *DaeMon*, when creating artificial disturbance in the network during runtime.




Fig. 14. Hit ratio in local memory of LC, PQ and *DaeMon*, when creating artificial disturbance in the network during runtime.

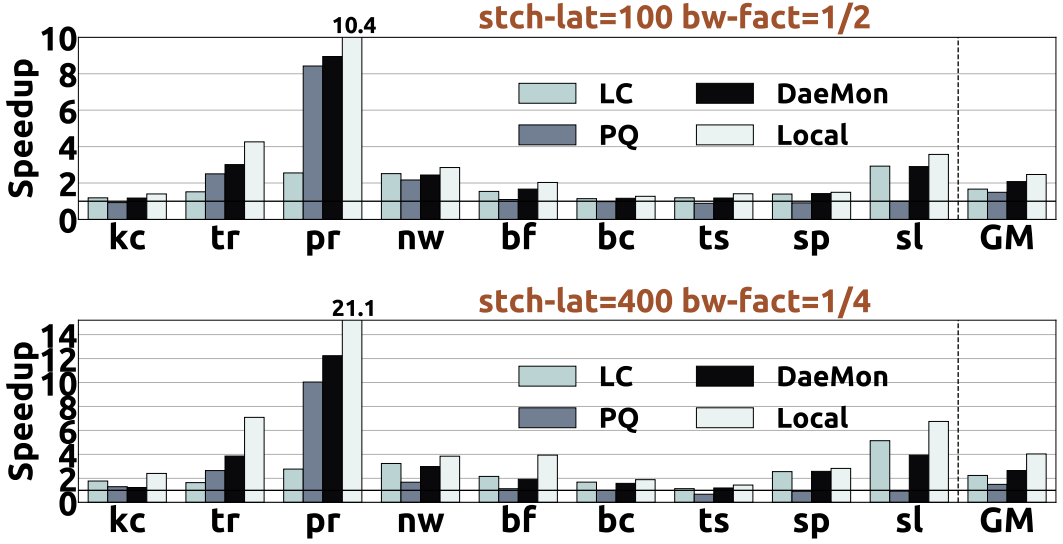
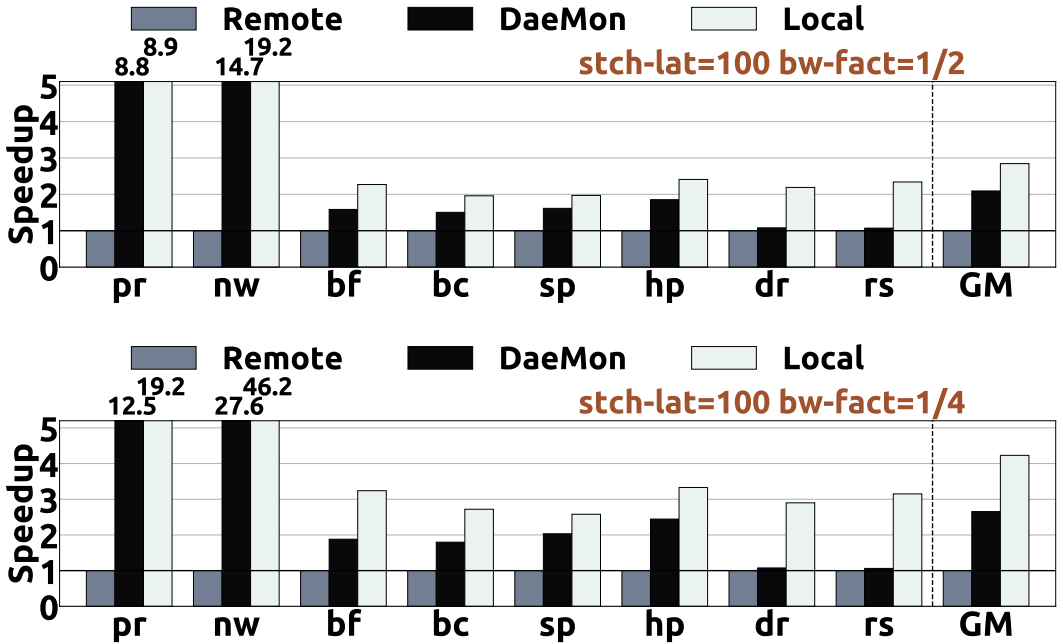


Fig. 15. Speedup achieved by various schemes in multithreaded workloads normalized to Remote.

Fig. 16. Performance of Local and *DaeMon* over Remote, when using FIFO replacement policy in local memory.

Across *all* workloads and network configurations (not graphed), *DaeMon* outperforms the widely-adopted Remote scheme by 2.63 \times , when using a FIFO replacement policy in local memory. *DaeMon* is orthogonal to the replacement policy used in local memory, and can be used synergistically with any arbitrary replacement policy in local memory to even further reduce data access costs.

Overall, *DaeMon* can significantly mitigate the data movement overheads in fully disaggregated systems independently on the number of data migrations happens during runtime: even when a small number of data migrations happens during runtime (e.g., thanks to sophisticated approaches such as intelligent replacement policies in local memory, hot page placement/selection techniques, page prefetchers), *DaeMon* can even further alleviate the data movement costs by dynamically selecting the granularity of data movements, prioritizing the critical cache line requests, and opportunistically moving compressed pages at slower rates. Therefore, we conclude that *DaeMon* can work synergistically with sophisticated replacement policies in local memory, page prefetchers and intelligent page placement/movement techniques to even further improve system performance.

6.9 Sensitivity Study to Multiple Memory Components

Figure 17 compares Remote and *DaeMon* normalized to Local, when varying the number of memory components and having a different network configuration for each memory component. Please also see Figure 22 in Appendix § A. We evaluated distributing memory pages with either a round-robin way or randomly across remote memory components, and draw the same key observation for both distributions. When adding more memory components using the same network configuration with that of when having one memory component (e.g., having 100 ns switch latency and 1/4 bandwidth factor for each memory component), performance of both Remote and *DaeMon* improves over Local: memory pages are distributed across *multiple* memory components and the system provides larger aggregate network and memory bandwidth, thus data migrations incur smaller overheads. Finally, *DaeMon* significantly outperforms Remote by 3.25 \times across all workload-architecture combinations, and constitutes a scalable solution for large-scale disaggregated systems with multiple hardware components and various architectures.

6.10 Sensitivity Study to Multiple Concurrent Workloads

Figure 18 shows *DaeMon*'s performance benefits when concurrently running multiple workloads on a compute component with 4 OoO cores. The performance of *each* core is normalized to that of the same core using Remote. The local memory hosts $\sim 15\%$ and $\sim 9\%$ of *each* application's working set, when running 2 and 4 workloads, respectively. *DaeMon* outperforms Remote by 1.96 \times across all multiple-workload experiments, thus being highly efficient and performant when multiple *heterogeneous* jobs concurrently run in the disaggregated system.

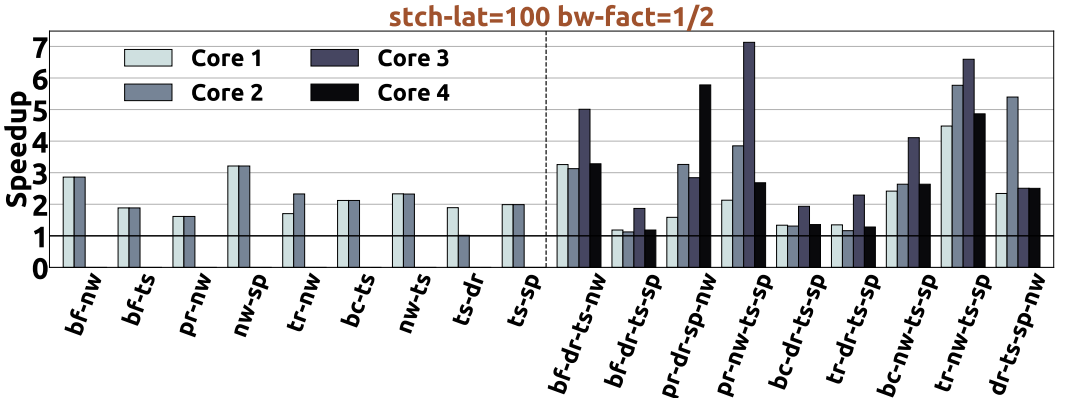


Fig. 18. Performance of *DaeMon* over Remote when running multiple concurrent workloads in a 4-CPU compute component and a memory component.

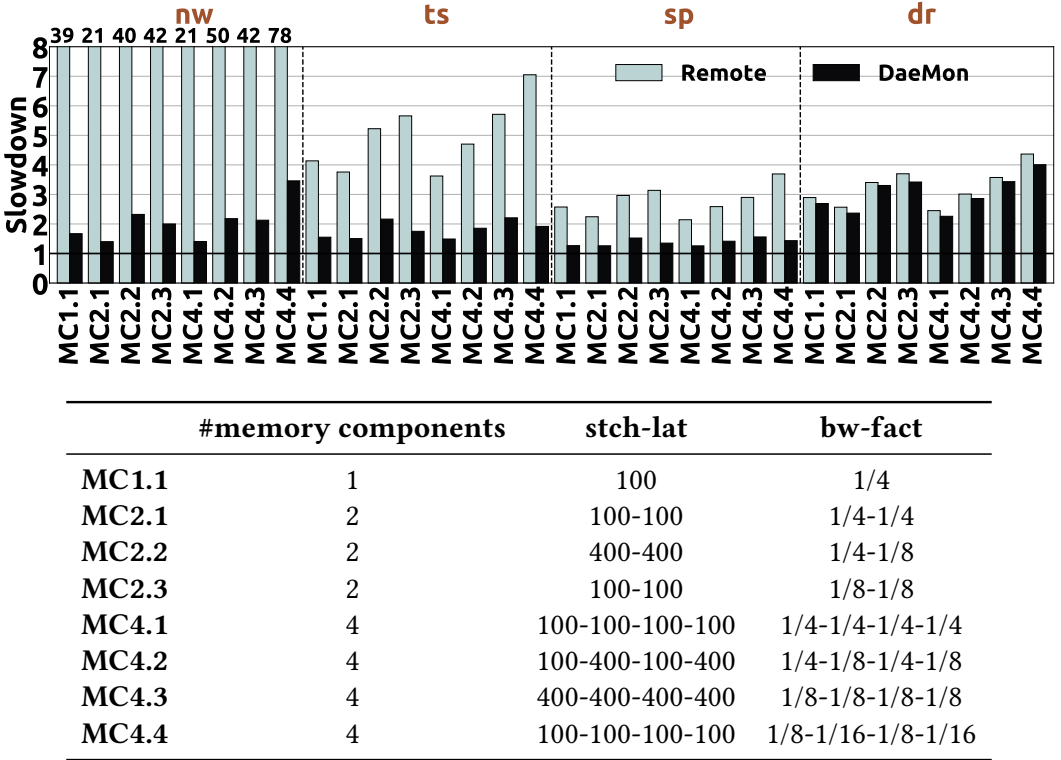


Fig. 17. Performance of Remote and *DaeMon* over Local when using multiple memory components.

6.11 Key Takeaways and Recommendations

This section summarizes our key takeaways and recommendations extracted from our evaluations.

Key Takeaway #1. *There is no one-size-fits-all granularity in data movements: the best-performing granularity at each time depends on the network/system load and the application data access patterns, which can significantly vary across applications and within application during runtime.* Figure 8 demonstrates that some applications significantly benefit from the prioritization of critical cache line data movements (e.g., *pr*, *nw*), and some applications only benefit from page migrations that leverage data locality (e.g., *dr*, *rs*). Figure 12 shows that some applications have highly compressible data, and thus greatly benefit from *compressed* page granularity data movements. Finally, Figure 13 proves that the application behavior and network traffic can highly vary during runtime, and thus the best-performing data movement granularity needs to adapt to the application characteristics and network/system conditions. Therefore, we recommend that system and hardware designers of disaggregated systems implement system-level solutions and hardware mechanisms that dynamically change and adapt their configurations and selection methods to the availability of the system resources and the runtime behavior of the heterogeneous applications.

Key Takeaway #2. *Typical datacenter applications exhibit high data locality within memory pages (e.g., 4KB).* Figure 10 shows that Remote achieves high data locality, i.e., always has at least 90% hit ratio in local memory, across a wide variety of datacenter workloads with diverse access patterns. Therefore, migrating data at a large granularity, e.g., page granularity, is very effective and critical to achieving high system performance in fully disaggregated systems. To this end, we suggest that hardware and system designers of disaggregated systems retain coarse-grained data migration

(i.e., page granularity data migration), since it both enables high performance and maintains low metadata overheads for address translation in local memory and remote memory.

Key Takeaway #3. *Aggressively prioritizing the cache line granularity data movements that are on the critical path might hurt performance.* Figure 11 shows that a high bandwidth partitioning ratio, e.g., 50% or 80% bandwidth partitioning ratio, which significantly prioritizes the cache line granularity data movements over the page granularity data movements, incurs significant performance slowdowns in workloads with medium and high spatial locality. As a result, we suggest that hardware and system designers of data movement solutions tailored for disaggregated systems always ensure that page migrations are not aggressively stalled.

Key Takeaway #4. *Distributed and disaggregated data movements solutions are highly effective and efficient in fully disaggregated systems.* Disaggregated systems are distributed architectures and comprise multiple hardware devices, each of them is independently and transparently managed from other hardware components in the system. Our evaluations in Figures 17 and 22 show that distributed and disaggregated solutions for data movement (i.e., *DaeMon*) better leverage the available aggregate network and memory bandwidth in the system, and enable high scalability to large-scale disaggregated systems with multiple hardware components. To this end, we recommend that hardware architects design distributed hardware mechanisms for fully disaggregated systems.

7 Related Work

To our knowledge, this is first work to (i) analyze and alleviate the data movement problem in fully disaggregated systems; (ii) enable prioritized and decoupled movement of data at multiple granularities simultaneously to reduce access latencies; (iii) propose a dynamic selection granularity mechanism with approximate bandwidth partitioning to effectively leverage both cache line and page movement depending on application and network characteristics; and (iv) implement a synergistic solution of link compression, bandwidth partitioning, and adaptive granularity selection in data movements. We discuss prior work.

Disaggregated Systems. Several prior works [5, 6, 10, 14–16, 33, 35, 36, 38, 39, 47, 55, 58, 71, 75, 76, 79, 88, 100, 109–111, 113] propose OS modules, system-level solutions, programming frameworks, software management systems, architectures and emulators for disaggregated systems. These works do not tackle the data movement problem in disaggregated systems, and thus *DaeMon* is orthogonal to these proposals. MIND [55] proposes memory sharing among compute components by implementing coherence and address translation in network switches. Kona [16] is a software runtime to track cache line granularity accesses to remote memory, and eliminate page faults by decoupling the application memory access tracking from the virtual memory page size. However, Kona and MIND do not mitigate data movement overheads in disaggregated systems, as data is always moved at page granularity. Thus, *DaeMon* is largely orthogonal to these works and could be used to further improve performance. Clio [38] proposes a disaggregated system that virtualizes and manages remote memory at the hardware level (independently to compute components), and eliminates expensive page faults in memory components. Clio accesses *remote* data at a byte granularity via dedicated API, however *not* being transparent to programmers. As explained in § 2.2, moving data *always* at a small granularity can cause significant performance penalties in many applications, and does not provide robustness against fluctuations in network characteristics. Instead, *DaeMon* is software-transparent, robust and significantly alleviates data movement costs via decoupled and selective data movement at *multiple* granularities. Lim et al. [57] propose a disaggregated architecture and characterize moving data *only* at cache line or page granularity. The authors show that the page-based configuration outperforms the cache line configuration at most common patterns (as observed in § 2.2), however it does not address the high performance penalties of page migrations. Maruf and Chowdhury [63] propose a page prefetching scheme for

disaggregated systems, which however can only help applications with high locality within pages, and does not capture the significant variability in data access costs of fully disaggregated systems. *DaeMon* is orthogonal to page prefetchers and can work synergistically with them to even further improve performance, as described in Section 4.7. We leave the experimentation of their synergy for future work.

Hybrid Memory Systems. Numerous works for hybrid memory systems propose data placement schemes [3, 19, 22, 28, 29, 32, 46, 48, 60, 83, 101], or selection methods [4, 26, 27, 44, 53, 54, 59, 65, 77, 90, 97, 104] to identify hot memory pages that are migrated to die-stacked DRAM, that is organized as a cache of a larger main memory. Compared to these approaches, first, intelligent page placement/movement is orthogonal to *DaeMon*, and cannot by itself address the high overheads caused by remote page migrations across the network, that can be significantly slower than that within the server and more latency/bandwidth-constrained in the context of fully disaggregated systems. Second, these prior works assume a monolithic centralized system where TLBs/page tables can be leveraged to track page hotness of remote pages (e.g., [4, 26, 27, 53, 59, 65, 104]) or that memory allocation/placement is handled by the server itself (e.g., [3, 28, 29, 32, 46, 48, 56, 60, 83]). However, in disaggregated systems, address translation and memory management are distributed across memory components and cannot be used to track pages at the CPU server side, while compute components and memory components are managed by independent kernel monitors that have no visibility/control of other components or data management/placement across components. Similarly, hardware-based approaches [44, 56, 77, 97] for hybrid systems add *centralized* hardware units at the server side to store page tracking metadata for the second-tier main memory. For example, Chop [44] adds 4MB of metadata to track 16GB of second-tier memory. These schemes would incur significant area overheads (in the order of GBs) to track large amounts of remote memory (in the order of TBs) enabled by disaggregated systems [88]. Requiring *each* compute component to track a large number of pages enabled by multiple remote memory components would cause scalability issues and significantly limit the benefits of resource disaggregation. Thus, designing an effective scalable hot page selection scheme for fully disaggregated systems is an open challenge, and *DaeMon* could work in conjunction with such schemes to further improve performance. Third, all these prior works do not handle variability in data access costs of disaggregated systems. Disaggregated systems necessitate an adaptive mechanism given the significant variations in access latencies and bandwidth. Fourth, applying/adapting the design of prior schemes tailored for tightly-integrated hybrid systems in disaggregated systems might incur significantly higher overheads and require important modifications than that described in the original papers.

A few recent works design hardware schemes for commodity servers to enable moving data *only* at cache line granularity [23, 61, 62, 98] or a larger sub-block granularity (a few cache lines) [43, 84]. Ekman et al. [30] evaluate a critical-block first approach, where each 8KB page is split in blocks of 2KB data, and the requested (critical) 2KB block of data is transferred first, and written in DRAM cache. As we show in § 2.2, moving data at a single granularity (page or cache line) can incur high performance costs and does not provide robustness towards significant variations in network bandwidth and latencies.

Hardware Compression. Prior works propose compression schemes [1, 8, 12, 21, 24, 31, 49, 51, 69, 70, 72–74, 78, 87, 92, 93, 102, 105, 106, 108] for cache memory, main memory and memory bus links in CPUs/GPUs [66, 86, 91, 99], and selection methods to dynamically enable/disable compression [7, 9, 96], or find the best-performing compression scheme [11, 50]. These works integrate ratio-optimized or latency-optimized compression schemes depending on the particular context and system’s characteristics they target. Our work enables link compression in page movements synergistically with decoupled multiple granularity data movement, which allows us to tolerate the high compression latencies of ratio-optimized compression schemes such as LZ [112].

8 Conclusion

DaeMon is the first adaptive data movement solution for fully disaggregated systems. *DaeMon* supports low-cost page migration, scales elastically to multiple hardware components, enables software transparency, and provides robustness across various architecture/network characteristics and the application behavior by effectively monitoring *pending* cache line and page movements. Our evaluations using a state-of-the-art simulator show that *DaeMon* significantly improves system performance and data access costs for a wide range of applications under various architecture and network configurations, and when multiple jobs are simultaneously running in the system. We conclude that *DaeMon* is an efficient, scalable and robust solution to alleviate data movement overheads in disaggregated systems, and hope that this work encourages further studies of the data movement problem in disaggregated systems.

Acknowledgments

We thank the anonymous reviewers from SIGMETRICS 2023, and our shepherd, Abhishek Chandra, for their comments and suggestions. We also thank Konstantinos Kanellopoulos and Ivan Fernandez for their help on technical aspects of this work.

References

- [1] B. Abali, H. Franke, D. E. Poff, R. A. Saccone, C. O. Schulz, L. M. Herger, and T. B. Smith, *Memory Expansion Technology (MXT): Software Support and Performance*, IBM Journal of Research and Development, 2001.
- [2] Atul Adya, Robert Grandl, Daniel Myers, and Henry Qin, *Fast Key-Value Stores: An Idea Whose Time Has Come and Gone*, HotOS, 2019.
- [3] Neha Agarwal, David Nellans, Mark Stephenson, Mike O'Connor, and Stephen W. Keckler, *Page Placement Strategies for GPUs within Heterogeneous Memory Systems*, ASPLOS, 2015.
- [4] Neha Agarwal and Thomas F. Wenisch, *Thermostat: Application-Transparent Page Management for Two-Tiered Main Memory*, ASPLOS, 2017.
- [5] Marcos K. Aguilera, Nadav Amit, Irina Calciu, Xavier Deguillard, Jayneel Gandhi, Stanko Novaković, Arun Ramanathan, Pratap Subrahmanyam, Lalith Suresh, Kiran Tati, Rajesh Venkatasubramanian, and Michael Wei, *Remote Regions: A Simple Abstraction for Remote Memory*, ATC, 2018.
- [6] Marcos K. Aguilera, Nadav Amit, Irina Calciu, Xavier Deguillard, Jayneel Gandhi, Pratap Subrahmanyam, Lalith Suresh, Kiran Tati, Rajesh Venkatasubramanian, and Michael Wei, *Remote Memory in the Age of Fast Networks*, SoCC, 2017.
- [7] A.R. Alameldeen and D.A. Wood, *Adaptive Cache Compression for High-Performance Processors*, ISCA, 2004.
- [8] Alaa Alameldeen and David Wood, *Frequent Pattern Compression: A Significance-Based Compression Scheme for L2 Caches*, Technical Report, 2004.
- [9] Alaa R. Alameldeen and David A. Wood, *Interactions Between Compression and Prefetching in Chip Multiprocessors*, HPCA, 2007.
- [10] Sebastian Angel, Mihir Nanavati, and Siddhartha Sen, *Disaggregation and the Application*, HotCloud, 2020.
- [11] Angelos Arelakis, Fredrik Dahlgren, and Per Stenstrom, *HyComp: A Hybrid Cache Compression Method for Selection of Data-Type-Specific Compression Methods*, MICRO, 2015.
- [12] Angelos Arelakis and Per Stenstrom, *SC2: A Statistical Compression Cache Scheme*, ISCA, 2014.
- [13] JEDEC Solid State Technology Assn., *JESD79-4B: DDR4 SDRAM Standard*, <http://www.softnology.biz/pdf/JESD79-4B.pdf>, 2017.
- [14] Laurent Bindschaedler, Ashvin Goel, and Willy Zwaenepoel, *Hailstorm: Disaggregated Compute and Storage for Distributed LSM-Based Databases*, ASPLOS, 2020.
- [15] Dhantu Buragohain, Abhishek Ghogare, Trishal Patel, Mythili Vutukuru, and Purushottam Kulkarni, *DiME: A Performance Emulator for Disaggregated Memory Architectures*, APSys, 2017.
- [16] Irina Calciu, M. Talha Imran, Ivan Puddu, Sanidhya Kashyap, Hasan Al Maruf, Onur Mutlu, and Aasheesh Kolli, *Rethinking Software Runtimes for Disaggregated Memory*, ASPLOS, 2021.
- [17] Trevor E. Carlson, Wim Heirman, and Lieven Eeckhout, *Sniper: Exploring the Level of Abstraction for Scalable and Accurate Parallel Multi-Core Simulations*, SC, 2011.
- [18] Trevor E. Carlson, Wim Heirman, Stijn Eyerman, Ibrahim Hur, and Lieven Eeckhout, *An Evaluation of High-Level Mechanistic Core Models*, TACO, 2014.
- [19] Chia-Hao Chang, Adithya Kumar, and Anand Sivasubramaniam, *To Move or Not to Move? Page Migration for Irregular Applications in over-Subscribed GPU Memory Systems with DynaMap*, SYSTOR, 2021.
- [20] Shuai Che, Michael Boyer, Jiayuan Meng, David Tarjan, Jeremy W. Sheaffer, Sang-Ha Lee, and Kevin Skadron, *Rodinia: A Benchmark Suite for Heterogeneous Computing*, IISWC, 2009.
- [21] Xi Chen, Lei Yang, Robert P. Dick, Li Shang, and Haris Lekatsas, *C-Pack: A High-Performance Microprocessor Cache Compression Algorithm*, VLSI, 2010.
- [22] Chiachen Chou, Aamer Jaleel, and Moinuddin Qureshi, *BATMAN: Techniques for Maximizing System Bandwidth of Memory Systems with Stacked-DRAM*, MEMSYS, 2017.
- [23] Chia Chen Chou, Aamer Jaleel, and Moinuddin K. Qureshi, *CAMEO: A Two-Level Memory Organization with Capacity of Main Memory and Flexibility of Hardware-Managed Cache*, MICRO, 2014.
- [24] Esha Choukse, Mattan Erez, and Alaa R. Alameldeen, *Compresso: Pragmatic Main Memory Compression*, MICRO, 2018.
- [25] David Cock, Abishek Ramdas, Daniel Schwyn, Michael Giardino, Adam Turowski, Zhenhao He, Nora Hossle, Dario Korolija, Melissa Licciardello, Kristina Martsenko, Reto Acherermann, Gustavo Alonso, and Timothy Roscoe, *Enzian: An Open, General, CPU/FPGA Platform for Systems Software Research*, ASPLOS, 2022.
- [26] Xiangyu Dong, Yuan Xie, Naveen Muralimanohar, and Norman P. Jouppi, *Simple but Effective Heterogeneous Main Memory with On-Chip Memory Controller Support*, SC, 2010.
- [27] Thaleia Dimitra Doudali, Sergey Blagodurov, Abhinav Vishnu, Sudhanva Gurumurthi, and Ada Gavrilovska, *Kleio: A Hybrid Memory Page Scheduler with Machine Intelligence*, HPDC, 2019.
- [28] Thaleia Dimitra Doudali, Daniel Zahka, and Ada Gavrilovska, *Cori: Dancing to the Right Beat of Periodic Data Movements over Hybrid Memory Systems*, IPDPS, 2021.

- [29] Subramanya R. Dulloor, Amitabha Roy, Zheguang Zhao, Narayanan Sundaram, Nadathur Satish, Rajesh Sankaran, Jeff Jackson, and Karsten Schwan, *Data Tiering in Heterogeneous Memory Systems*, EuroSys, 2016.
- [30] Magnus Ekman and Per Stenstrom, *A Cost-Effective Main Memory Organization for Future Servers*, IPDPS, 2005.
- [31] M. Ekman and P. Stenstrom, *A Robust Main-Memory Compression Scheme*, ISCA, 2005.
- [32] M. J. Feeley, W. E. Morgan, E. P. Pighin, A. R. Karlin, H. M. Levy, and C. A. Thekkath, *Implementing Global Memory Management in a Workstation Cluster*, SOS, 1995.
- [33] Peter X. Gao, Akshay Narayan, Sagar Karandikar, Joao Carreira, Sangjin Han, Rachit Agarwal, Sylvia Ratnasamy, and Scott Shenker, *Network Requirements for Resource Disaggregation*, OSDI, 2016.
- [34] Christina Giannoula, *Accelerating Irregular Applications via Efficient Synchronization and Data Access Techniques*, <https://arxiv.org/abs/2211.05908>, 2022.
- [35] Donghyun Gouk, Sangwon Lee, Miryeong Kwon, and Myoungsoo Jung, *Direct Access, High-Performance Memory Disaggregation with DirectCXL*, ATC, 2022.
- [36] Juncheng Gu, Youngmoon Lee, Yiwen Zhang, Mosharaf Chowdhury, and Kang G. Shin, *Efficient Memory Disaggregation with Infiniswap*, NSDI, 2017.
- [37] Chuanxiong Guo, Haitao Wu, Zhong Deng, Gaurav Soni, Jianxi Ye, Jitu Padhye, and Marina Lipshteyn, *RDMA over Commodity Ethernet at Scale*, SIGCOMM, 2016.
- [38] Zhiyuan Guo, Yizhou Shan, Xuhao Luo, Yutong Huang, and Yiyang Zhang, *Clio: A Hardware-Software Co-Designed Disaggregated Memory System*, ASPLOS, 2022.
- [39] Sangjin Han, Norbert Egi, Aurojit Panda, Sylvia Ratnasamy, Guangyu Shi, and Scott Shenker, *Network Support for Resource Disaggregation in Next-Generation Datacenters*, HotNets, 2013.
- [40] HPCG, *High Performance Conjugate Gradient Benchmark*, <https://github.com/hpcg-benchmark/hpcg>, 2019.
- [41] Ranggi Hwang, Taehun Kim, Youngeun Kwon, and Minsoo Rhu, *Centaur: A Chiplet-Based, Hybrid Sparse-Dense Accelerator for Personalized Recommendations*, ISCA, 2020.
- [42] Intel, *Intel Omni-Path Architecture*, <https://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-driving-exascale-computing.html>, 2021.
- [43] Djordje Jevdjic, Stavros Volos, and Babak Falsafi, *Die-Stacked DRAM Caches for Servers: Hit Ratio, Latency, or Bandwidth? Have It All with Footprint Cache*, ISCA, 2013.
- [44] Xiaowei Jiang, Niti Madan, Li Zhao, Mike Upton, Ravishankar Iyer, Srihari Makineni, Donald Newell, Yan Solihin, and Rajeev Balasubramanian, *CHOP: Adaptive Filter-Based DRAM Caching for CMP Server Platforms*, HPCA, 2010.
- [45] Hongshin Jun, Jinhee Cho, Kangseol Lee, Ho-Young Son, Kwiwook Kim, Hanho Jin, and Keith Kim, *HBM DRAM Technology and Architecture*, IMW, 2017.
- [46] Sudarsun Kannan, Ada Gavrilovska, Vishal Gupta, and Karsten Schwan, *HeteroOS: OS Design for Heterogeneous Memory Management in Datacenter*, ISCA, 2017.
- [47] K. Katrinis, D. Syrivelis, D. Pnevmatikatos, G. Zervas, D. Theodoropoulos, I. Koutsopoulos, K. Hasharoni, D. Raho, C. Pinto, F. Espina, S. Lopez-Buedo, Q. Chen, M. Nemirovsky, D. Roca, H. Klos, and T. Berends, *Rack-Scale Disaggregated Cloud Data Centers: The dReDBox Project Vision*, DATE, 2016.
- [48] Jonghyeon Kim, Wonkyo Choe, and Jeongseob Ahn, *Exploring the Design Space of Page Management for Multi-Tiered Memory Systems*, ATC, 2021.
- [49] Jung-rae Kim, Michael Sullivan, Esha Choukse, and Mattan Erez, *Bit-Plane Compression: Transforming Data for Better Compression in Many-Core Architectures*, ISCA, 2016.
- [50] Seikwon Kim, Seonyoung Lee, Taehoon Kim, and Jaehyuk Huh, *Transparent Dual Memory Compression Architecture*, PACT, 2017.
- [51] M. Kjelsø, M. Gooch, and S. Jones, *Design and Performance of a Main Memory Hardware Data Compressor*, EUROMICRO, 1996.
- [52] Fredrik Kjolstad, Stephen Chou, David Lugato, Shoaib Kamil, and Saman Amarasinghe, *Taco: A Tool to Generate Tensor Algebra Kernels*, ASE, 2017.
- [53] Jagadish B. Kotra, Haibo Zhang, Alaa R. Alameldeen, Chris Wilkerson, and Mahmut T. Kandemir, *CHAMELEON: A Dynamically Reconfigurable Heterogeneous Memory System*, MICRO, 2018.
- [54] Andres Lagar-Cavilla, Junwhan Ahn, Suleiman Souhlal, Neha Agarwal, Radoslaw Burny, Shakeel Butt, Jichuan Chang, Ashwin Chaugule, Nan Deng, Junaid Shahid, Greg Thelen, Kamil Adam Yurtsever, Yu Zhao, and Parthasarathy Ranganathan, *Software-Defined Far Memory in Warehouse-Scale Computers*, ASPLOS, 2019.
- [55] Seung-seob Lee, Yanpeng Yu, Yupeng Tang, Anurag Khandelwal, Lin Zhong, and Abhishek Bhattacharjee, *MIND: In-Network Memory Management for Disaggregated Data Centers*, SOS, 2021.
- [56] Yang Li, Saugata Ghose, Jongmoo Choi, Jin Sun, Hui Wang, and Onur Mutlu, *Utility-Based Hybrid Memory Management*, CLUSTER, 2017.
- [57] Kevin Lim, Jichuan Chang, Trevor Mudge, Parthasarathy Ranganathan, Steven K. Reinhardt, and Thomas F. Wenisch, *Disaggregated Memory for Expansion and Sharing in Blade Servers*, ISCA, 2009.

- [58] Kevin Lim, Yoshio Turner, Jose Renato Santos, Alvin AuYoung, Jichuan Chang, Parthasarathy Ranganathan, and Thomas F. Wenisch, *System-Level Implications of Disaggregated Memory*, HPCA, 2012.
- [59] Haikun Liu, Yujie Chen, Xiaofei Liao, Hai Jin, Bingsheng He, Long Zheng, and Rentong Guo, *Hardware/Software Cooperative Caching for Hybrid DRAM/NVM Memory Architectures*, ICS, 2017.
- [60] Lei Liu, Shengjie Yang, Lu Peng, and Xinyu Li, *Hierarchical Hybrid Memory Management in OS for Tiered Memory Systems*, TPDS, 2019.
- [61] Gabriel Loh and Mark D. Hill, *Supporting Very Large DRAM Caches with Compound-Access Scheduling and MissMap*, IEEE Micro, 2012.
- [62] Gabriel H. Loh and Mark D. Hill, *Efficiently Enabling Conventional Block Sizes for Very Large Die-Stacked DRAM Caches*, MICRO, 2011.
- [63] Hasan Al Maruf and Mosharaf Chowdhury, *Effectively Prefetching Remote Memory with Leap*, ATC, 2020.
- [64] Mellanox, *Mellanox Innova Adapters*, https://www.nvidia.com/en-us/networking/products/data-processing-unit/?mtag=programmable_adapter_cards, 2020.
- [65] Mitesh R. Meswani, Sergey Blagodurov, David Roberts, John Slice, Mike Ignatowski, and Gabriel H. Loh, *Heterogeneous Memory Architectures: A HW/SW Approach for Mixing Die-Stacked and Off-Package Memories*, HPCA, 2015.
- [66] Sparsh Mittal and Jeffrey S. Vetter, *A Survey Of Architectural Approaches for Data Compression in Cache and Main Memory Systems*, TPDS, 2016.
- [67] Naveen Muralimanohar, Rajeev Balasubramonian, and Norm Jouppi, *Optimizing NUCA Organizations and Wiring Alternatives for Large Caches with CACTI 6.0*, MICRO, 2007.
- [68] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G. Azzolini, Dmytro Dzhulgakov, Andrey Mallevich, Ilia Cherniavskii, Yinghai Lu, Raghuraman Krishnamoorthi, Ansha Yu, Volodymyr Kondratenko, Stephanie Pereira, Xianjie Chen, Wenlin Chen, Vijay Rao, Bill Jia, Liang Xiong, and Misha Smelyanskiy, *Deep Learning Recommendation Model for Personalization and Recommendation Systems*, arXiv, 2019.
- [69] Tri M. Nguyen, Adi Fuchs, and David Wentzlaff, *CABLE: A CAche-Based Link Encoder for Bandwidth-Starved Manycores*, MICRO, 2018.
- [70] Tri M. Nguyen and David Wentzlaff, *MORC: A Manycore-Oriented Compressed Cache*, MICRO, 2015.
- [71] Vlad Nitu, Boris Teabe, Alain Tchana, Canturk Isci, and Daniel Hagimont, *Welcome to ZombieLand: Practical and Energy-Efficient Memory Disaggregation in a Datacenter*, EuroSys, 2018.
- [72] Sungbo Park, Ingab Kang, Yaebin Moon, Jung Ho Ahn, and G. Edward Suh, *BCD Deduplication: Effective Memory Compression Using Partial Cache-Line Deduplication*, ASPLOS, 2021.
- [73] Gennady Pekhimenko, Vivek Seshadri, Yoongu Kim, Hongyi Xin, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry, *Linearly Compressed Pages: A Low-Complexity, Low-Latency Main Memory Compression Framework*, MICRO, 2013.
- [74] Gennady Pekhimenko, Vivek Seshadri, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry, *Base-Delta-Immediate Compression: Practical Data Compression for on-Chip Caches*, PACT, 2012.
- [75] Ivy Peng, Roger Pearce, and Maya Gokhale, *On the Memory Underutilization: Exploring Disaggregated Memory on HPC Systems*, SBAC-PAD, 2020.
- [76] Christian Pinto, Dimitris Syrivelis, Michele Gazzetti, Panos Koutsouvasilis, Andrea Reale, Kostas Katrinis, and H. Peter Hofstee, *ThymesisFlow: A Software-Defined, HW/SW co-Designed Interconnect Stack for Rack-Scale Memory Disaggregation*, MICRO, 2020.
- [77] Andreas Prodromou, Mitesh Meswani, Nuwan Jayasena, Gabriel Loh, and Dean M. Tullsen, *MemPod: A Clustered Architecture for Efficient and Scalable Migration in Flat Address Space Multi-level Memories*, HPCA, 2017.
- [78] Cheng Qian, Libo Huang, Qi Yu, Zhiying Wang, and Bruce Childers, *CMH: Compression Management for Improving Capacity in the Hybrid Memory Cube*, CF, 2018.
- [79] Pramod Subba Rao and George Porter, *Is Memory Disaggregation Feasible? A Case Study with Spark SQL*, ANCS, 2016.
- [80] RDMA, *RDMA Consortium*, <http://www.rdmaconsortium.org/>, 2019.
- [81] RDMA, *Gen-Z Core Specification*, <https://genzconsortium.org/>, 2022.
- [82] Joseph Redmon, *Darknet: Open Source Neural Networks in C*, <http://pjreddie.com/darknet/>, 2013–2016.
- [83] Zhenyuan Ruan, Malte Schwarzkopf, Marcos K. Aguilera, and Adam Belay, *AIFM: High-Performance, Application-Integrated Far Memory*, OSDI, 2020.
- [84] Jee Ho Ryoo, Mitesh R. Meswani, Andreas Prodromou, and Lizy K. John, *SILC-FM: Subblocked InterLeaved Cache-Like Flat Memory Organization*, HPCA, 2017.
- [85] Amedeo Sapio, Ibrahim Abdelaziz, Abdulla Aldilajan, Marco Canini, and Panos Kalnis, *In-Network Computation is a Dumb Idea Whose Time Has Come*, HotNets, 2017.
- [86] Vijay Sathish, Michael J. Schulte, and Nam Sung Kim, *Lossless and Lossy Memory I/O Link Compression for Improving Performance of GPGPU Workloads*, PACT, 2012.

- [87] Ali Shafiee, Meysam Taassori, Rajeev Balasubramonian, and Al Davis, *MemZip: Exploring Unconventional Benefits from Memory Compression*, HPCA, 2014.
- [88] Yizhou Shan, Yutong Huang, Yilun Chen, and Yiyang Zhang, *LegoOS: A Disseminated, Distributed OS for Hardware Resource Disaggregation*, OSDI, 2018.
- [89] Julian Shun and Guy E. Blelloch, *Ligra: A Lightweight Graph Processing Framework for Shared Memory*, PpopP, 2013.
- [90] Gagandeep Singh, Rakesh Nadig, Jisung Park, Rahul Bera, Nastaran Hajinazar, David Novo, Juan Gómez-Luna, Sander Stuijk, Henk Corporaal, and Onur Mutlu, *Sibyl: Adaptive and Extensible Data Placement in Hybrid Storage Systems Using Online Reinforcement Learning*, ISCA, 2022.
- [91] Martin Thuresson, Lawrence Spracklen, and Per Stenstrom, *Memory-Link Compression Schemes: A Value Locality Perspective*, IEEE Trans. Comput., 2008.
- [92] Martin Thuresson and Per Stenström, *Accommodation of the Bandwidth of Large Cache Blocks Using Cache/Memory Link Compression*, ICPP, 2008.
- [93] Yingying Tian, Samira M. Khan, Daniel A. Jiménez, and Gabriel H. Loh, *Last-Level Cache Deduplication*, ICS, 2014.
- [94] R.B. Tremaine, T.B. Smith, M. Wazlowski, D. Har, Kwok-Ken Mak, and S. Arramreddy, *Pinnacle: IBM MXT in a Memory Controller Chip*, IEEE Micro, 2001.
- [95] Shin-Yeh Tsai and Yiyang Zhang, *LITE Kernel RDMA Support for Datacenter Applications*, SOSp, 2017.
- [96] Irina Chihaiia Tudu and Thomas Gross, *Adaptive Main Memory Compression*, ATC, 2005.
- [97] Evangelos Vasilakis, Vassilis Papaefstathiou, Pedro Trancoso, and Ioannis Sourdis, *LLC-Guided Data Migration in Hybrid Memory Systems*, IPDPS, 2019.
- [98] Vasilakis, Evangelos and Papaefstathiou, Vassilis and Trancoso, Pedro and Sourdis, Ioannis, *Hybrid2: Combining Caching and Migration in Hybrid Memory Systems*, HPCA, 2020.
- [99] Nandita Vijaykumar, Gennady Pekhimenko, Adwait Jog, Abhishek Bhowmick, Rachata Ausavarungnirun, Chita Das, Mahmut Kandemir, Todd C. Mowry, and Onur Mutlu, *A Case for Core-Assisted Bottleneck Acceleration in GPUs: Enabling Flexible Data Compression with Assist Warps*, ISCA, 2015.
- [100] Chenxi Wang, Haoran Ma, Shi Liu, Yuanqi Li, Zhenyuan Ruan, Khanh Nguyen, Michael D. Bond, Ravi Netravali, Miryung Kim, and Guoqing Harry Xu, *Semeru: A Memory-Disaggregated Managed Runtime*, OSDI, 2020.
- [101] Johannes Weiner, Niket Agarwal, Dan Schatzberg, Leon Yang, Hao Wang, Blaise Sanouillet, Bikash Sharma, Tejun Heo, Mayank Jain, Chunqiang Tang, and Dimitrios Skarlatos, *TMO: Transparent Memory Offloading in Datacenters*, ASPLOS, 2022.
- [102] P. Wilson, Scott F. Kaplan, and Y. Smaragdakis, *The Case for Compressed Caching in Virtual Memory Systems*, ATC, 1999.
- [103] Dong Hyuk Woo, Nak Hee Seong, Dean L. Lewis, and Hsien-Hsin Sean Lee, *An Optimized 3D-stacked Memory Architecture by Exploiting Excessive, High-Density TSV Bandwidth*, HPCA, 2010.
- [104] Zi Yan, Daniel Lustig, David Nellans, and Abhishek Bhattacharjee, *Nimble Page Management for Tiered Memory Systems*, ASPLOS, 2019.
- [105] Jun Yang, Rajiv Gupta, and Chuanjun Zhang, *Frequent Value Encoding for Low Power Data Buses*, TODAES, 2004.
- [106] Jun Yang, Youtao Zhang, and R. Gupta, *Frequent Value Compression in Data Caches*, MICRO, 2000.
- [107] Chin-Chia Michael Yeh, Yan Zhu, Liudmila Ulanova, Nurjahan Begum, Yifei Ding, Hoang Anh Dau, Diego Furtado Silva, Abdullah Mueen, and Eamonn Keogh, *Matrix Profile I: All Pairs Similarity Joins for Time Series: A Unifying View that Includes Motifs, Discords and Shapelets*, ICDM, 2016.
- [108] Vinson Young, Sanjay Kariyappa, and Moinuddin K. Qureshi, *Enabling Transparent Memory-Compression for Commodity Memory Systems*, HPCA, 2019.
- [109] Georgios Zervas, Hui Yuan, Arsalan Saljoghei, Qianqiao Chen, and Vaibhawa Mishra, *Optically Disaggregated Data Centers with Minimal Remote Memory Latency: Technologies, Architectures, and Resource Allocation*, JOCN, 2018.
- [110] Qizhen Zhang, Yifan Cai, Sebastian G. Angel, Vincent Liu, Ang Chen, and B. T. Loo, *Rethinking Data Management Systems for Disaggregated Data Centers*, CIDR, [n. d.].
- [111] Yang Zhou, Hassan M. G. Wassel, Sihang Liu, Jiaqi Gao, James Mickens, Minlan Yu, Chris Kennelly, Paul Turner, David E. Culler, Henry M. Levy, and Amin Vahdat, *Carbink: Fault-Tolerant Far Memory*, OSDI, 2022.
- [112] J. Ziv and A. Lempel, *A Universal Algorithm for Sequential Data Compression*, IEEE Transactions on Information Theory, 1977.
- [113] Pengfei Zuo, Jiazhao Sun, Liu Yang, Shuangwu Zhang, and Yu Hua, *One-sided RDMA-Conscious Extendible Hashing for Disaggregated Memory*, ATC, 2021.

APPENDIX

A Extended Results

A.1 Network Bandwidth Utilization

Figure 19 compares the bandwidth utilization across the network of a compute component and a memory component achieved by various data movement schemes.




Fig. 19. Bandwidth utilization (%) across the network of a compute component and a memory component achieved by various data movement schemes.

We make three key observations. First, LC typically reduces the network bandwidth utilization over Remote (by 2.49× on average across all workloads and network configurations), because fewer bytes are transferred through the network, since remote pages are migrated in a compressed format. Note that LC improves the total execution time over Remote, and thus in a few workloads, e.g., *pr*, the network bandwidth utilization might be higher within a smaller execution time. Second, PQ decreases the network bandwidth utilization over Remote in workloads with poor spatial locality within pages (e.g., *nw*), since the selection granularity unit effectively schedules more cache line movements and fewer page migrations. Instead, PQ might slightly increase the network bandwidth utilization over Remote in workloads with medium spatial locality within pages (e.g., *bf*, *bc*), since the selection granularity unit enables both cache line and page migrations to leverage both the ability to prioritize critical cache line requests and the benefits of data locality within pages. In workloads with high spatial locality within pages (e.g., *dr*, *rs*), PQ favors more page migrations and fewer cache line movements, thus achieving similar network bandwidth utilization to Remote. Third, *DaeMon* greatly decreases the network bandwidth utilization over Remote by 2.32× on average across all workloads and network configurations (not graphed). *DaeMon* effectively transfers remote pages in a compressed format and on-the-fly selects the granularity of data migrations to *significantly* reduce the bandwidth consumption across the network of fully disaggregated systems.

A.2 Sensitivity Study to Switch Latency

Figure 20 compares *DaeMon*'s performance over Remote's performance averaged across all workloads, when varying the switch latency of the network. When the fixed switch latency becomes very high dominating the total data movement costs, *DaeMon* has lower benefits over Remote, since *DaeMon* does not hide the propagation and switching delays in network components (e.g., fixed processing costs of the packet inside network switches). However, even with a very high switch latency in the order of microsecond, i.e., $1\mu s$ ($=1000\text{ ns}$), *DaeMon* outperforms Remote by 1.49 \times on average across all workloads.




Fig. 20. Performance benefits of *DaeMon* over Remote, when varying the switch latency of the network.

A.3 Sensitivity Study to Network Bandwidth

To evaluate bandwidth-limited scenarios, Figure 21 compares *DaeMon*'s performance normalized to Remote's performance in multithreaded workloads running on 8 OoO cores of a compute component, when varying the bandwidth factor of the network, e.g., up to having a very low bandwidth factor of 1/16 (i.e., network bandwidth is 16 \times slower than the DRAM bus bandwidth) between a compute component and memory component. We find that on average *DaeMon*'s benefits increase over the widely-adopted approach of moving data at page granularity, i.e., Remote, since *DaeMon* even more significantly alleviates bandwidth bottlenecks and data movement overheads under bandwidth-constrained scenarios.




Fig. 21. Performance benefits of *DaeMon* normalized to Remote using multithreaded workloads, when varying the bandwidth factor of the network between a compute component and memory component.

A.4 Performance Benefits With Multiple Memory Components

Figure 22 evaluates the performance of *DaeMon* normalized to Remote's performance, when increasing the number of memory components in the system having the same network configuration for each memory component, i.e., 100 ns switch latency and a bandwidth factor of 1/4. We evaluated distributing memory pages with either a round-robin way or randomly across multiple remote memory components, and drew the same key observations for both distributions. Similarly to

Figure 17, we observe that when pages are distributed across *multiple* memory components and the system provides larger aggregate network and memory bandwidth, data access costs decrease. For example, when increasing the number of memory components from 2 to 4, the remote data access latency decreases by 1.39 \times on average across all workloads. However, even when data access costs affect less the total execution time of applications, *DaeMon* still further mitigates data access overheads: *DaeMon* outperforms the widely-adopted Remote approach by 2.09 \times and 1.88 \times on average across all workloads, when using 2 and 4 memory components, respectively.




Fig. 22. Performance benefits of *DaeMon* normalized to Remote, when increasing the number of memory components having 100 ns switch latency and a bandwidth factor of 1/4 for each memory component.