# Survey of Disaggregated Memory: Cross-layer Technique Insights for Next-Generation Datacenters

## JING WANG, CHAO LI\*, TAOLEI WANG, JINYANG GUO, HANZHANG YANG, YIMING ZHUANSUN, MINYI GUO, Shanghai Jiao Tong University, China

The growing scale of data requires efficient memory subsystems with large memory capacity and high memory performance. Disaggregated architecture has become a promising solution for today's cloud and edge computing for its scalability and elasticity. As a critical part of disaggregation, disaggregated memory faces many design challenges in many dimensions, including hardware scalability, architecture structure, software system design, application programmability, resource allocation, power management, etc. These challenges inspire a number of novel solutions at different system levels to improve system efficiency. In this paper, we provide a comprehensive review of disaggregated memory, including the methodology and technologies of disaggregated memory system foundation, optimization, and management. We study the technical essentials of disaggregated memory systems and analyze them from the hardware, architecture, system, and application levels. Then, we compare the design details of typical cross-layer designs on disaggregated memory. Finally, we discuss the challenges and opportunities of future disaggregated memory works that serve better for next-generation elastic and efficient datacenters.

Additional Key Words and Phrases: disaggregated memory, far memory, heterogeneous architecture, runtime optimization

## 1 INTRODUCTION

The scale of data-driven intelligent services today has grown exponentially, including video processing, natural language processing, graph processing, etc. The large amount of data used for applications provides better quality of analysis results, which requires careful memory resource management. Nevertheless, this aggravates the problem of extremely insufficient memory resources. Handling in-memory data poses great changes to today's hardware and software system designs. Both academic research and industrial deployment require efficient subsystems with larger memory capacity, higher computation performance, and higher resource utilization.

Most data center environments face memory bottleneck problems that block performance improvement of memory and storage devices. For example, hardware limitations necessitate the adoption of distributed computing and out-of-core processing in applications. One has to divide data into smaller parts or compress data size to handle memory shortage. In addition, today's cloud applications often rely on diverse memory-intensive services and runtimes, occupying nearly 70% of the workload [39]. The memory-related resource consumption/overhead of these services is significantly growing, which is known as the *memory tax* [152] phenomenon.

To address the above issues, one solution is to implement a large disaggregated memory pool and provide an unlimited memory resource, thereby alleviating the resource pressure. *Disaggregated architecture* is proposed as a flexible and composable infrastructure [84, 97], breaking the limits of scalability, performance, and efficiency on classic monolithic servers. Disaggregation allows applications to use any required ratio of resources of computing units, memory, storage, networks, accelerators, etc. One can re-aggregate disaggregated resources to compose entities and build upper-layer resource abstraction to utilize it efficiently and flexibly. Disaggregation can also solve resource fragmentation problems due to unbalanced resource scheduling on each machine.

Specifically, *disaggregated memory* (DM) aims to decouple CPU and memory resources so that one can flexibly manage these memory resources in a large and shared memory pool. Different from disaggregated storage and networks, the overhead of remote memory accesses can seriously slow down computation. To improve DM access performance, existing works try hard to adopt and design large-capacity memory devices, low-latency storage devices, and high-speed networks [3, 13, 101]. These DM devices are seen as a larger but slower "far memory" (local memory is its counterpart) and are added into the original memory hierarchy, as Figure 1 shows.

Author's address: Jing Wang, Chao Li\*, Taolei Wang, Jinyang Guo, Hanzhang Yang, Yiming Zhuansun, Minyi Guo, {jing618,sjtuwtl,lazarus, linqinluli,zsym2019}@sjtu.edu.cn,{lichao,guo-my}@cs.sjtu.edu.cn, Shanghai Jiao Tong University, Shanghai, China.



Fig. 1. This survey focuses on disaggregated memory(DM), which acts as an additional layer in the existing memory hierarchy. The survey mainly studies the design goal, design level, and key design points of the DM system.

Advanced memory and network devices may not guarantee high memory performance if the software systems are poorly designed. Existing operating systems (OS) and runtime systems face the problem of adapting to new DM hardware. Conventional cache-based hierarchical memory management methods can not exploit the high bandwidth and large capacity of emerging devices. Thus, learning about the key technologies, optimization opportunities, and future challenges is vital when designing high-performance and high-efficiency DM systems. There is an urgent need for a comprehensive and in-depth survey to provide a full-stack description of DM systems in the hardware, system, software, and applications layers.

This paper explores the hardware usage, architecture design, system adaption, runtime optimization, and future opportunities of DM systems. To our knowledge, it is the first survey to provide a comprehensive classification of technique optimization methods for DM systems. We discuss not only key techniques that lay the foundation for DM systems, but also the latest technologies related to memory hierarchies, heterogeneous devices, and direct-connect networks. We present the trend in heterogeneous memory design, the construction of disaggregated memory pools supporting heterogeneous environments with accelerators, as well as optimization methods for emerging AI applications. We show that, DM systems require well-designed architectural organization methods to better meet the needs of new types of applications and hardware.

This survey is organized as follows. As shown in Figure 1, we first give an overview of disaggregated memory systems in Section 2. We then introduce hardware usage (Section 3), architecture design (Section 4), OS and runtime layer development (Section 5), and application-oriented optimization methods (Section 6) of Disaggregated memory. We summarize some classic DM systems with cross-layer design and give a general analysis in Section 7. Finally, we discuss future work (Section 8) and conclude the paper (Section 9).

## 2 OVERVIEW

In this section, we first introduce the related survey works on disaggregated memory. We then discuss general design considerations of DM systems, including the design goals, features, and design basis, etc.

#### 2.1 Related Survey on Disaggregated Memory

Disaggregated memory was first discussed by Lim et al. [85] in 2009. They propose a general-purpose architectural building block, i.e., a memory blade, that allows memory to be "disaggregated" across a system ensemble. About a decade ago, in the 2010s, resource disaggregation was proposed. In 2016, Maciej et al. [23] conducted a survey on memory interconnection and sharing methods for HPC systems concerning scalability and virtualization support. Gao et al. [45] identified several research opportunities in this area, such as low-latency network,

disaggregation-aware scheduler, and new programming models. Li et al. [76] provided a survey on architecture design of edge-oriented computing paradigms, with a dedicated analysis of memory allocation.

The concept has been widely-used that disaggregation is for aggregation. Composable Disaggregated Infrastructure (CDI) [97] is introduced and becomes well known, which maintains disaggregated memory, disaggregated storage, and disaggregated network [59] and reorganizes the necessary resources. A cluster comprising multiple monolithic machines is established in this setup, underpinned by a Software-Defined Data Center (SDC).

In 2018, software-defined hardware infrastructures (SDHI) are proposed based on resource disaggregation [114], which aims to bring greater modularity, flexibility and extensibility to cloud infrastructures. Compared with conventional server-oriented software-defined infrastructure (SDI), SDHI can decompose the network control planes [25], storage control planes [37], software abstraction layers with computing function virtualization [59], etc. Wang et al. [144] argued that the next wave in database system innovation should be shared-memory designs enabled by RDMA-based memory disaggregation.

The most recent surveys discuss the opportunities and challenges of memory disaggregation on advanced memory devices. Li et al. [77] analyzed the taxonomy and key techniques regarding cloud-native disaggregated storage management, etc. Liu et al. [88] discussed some demands and challenges of memory disaggregation in cloud data centers and examined some promising research issues to show the opportunity of memory disaggregation in virtualized environments. Ewais et al. [42] provided a list of disaggregated memory works and discussed their basic architectures, implementations, and requirements. Tom et al. [95] gave an introduction to CXL memory [3], which is a representative protocol to enable the creation of pools of memory and accelerators.

We observe that related surveys have not summarized the optimization techniques for better application performance and system efficiency on disaggregated memory. This requires a deeper analysis of the design methodology of the system. In this survey, we take the first step to study the state-of-the-art techniques on disaggregated memory at each system level, providing a comprehensive overview of the current landscape.

#### 2.2 Design Considerations of Disaggregated Memory

In this subsection, we summarize the design considerations of creating an efficient DM system for different purposes. We first introduce the design goals of DM systems and then list important features of DM devices.

*2.2.1 Design goals.* A well-designed disaggregated memory system aims to combine the benefits of a large memory pool with the performance, flexibility, and reliability required to support data-hungry applications. These systems are envisioned to be particularly beneficial in environments with dynamic workloads and heterogeneous computing resources. We summarize the design goals of the disaggregated memory system as follows.

- Large available memory capacity: A primary objective is to provide a substantially larger memory pool than what is typically available in traditional monolithic servers, overcoming the physical limitations of in-server memory capacity[104]. Providing ample memory space to the central processing units, including CPUs and GPUs, can bring more chances for big data computation.
- High resource usage flexibility: Disaggregated memory should allow more granular and efficient allocations of memory resources, adapting to varying workload requirements. This flexibility enables better resource utilization, reduces wastage, and can save costs, especially in multi-tenant or cloud-based environments where resource demands fluctuate significantly[120].
- Reliable network and memory subsystems: When memory resources are accessed over a network, its network architecture must be robust and fault-tolerant. Similarly, the memory subsystem must be reliable, with mechanisms to protect against data loss or corruption[166]. This includes implementing redundancy, error detection, and correction techniques to maintain the integrity and availability of data.
- High application performance: Disaggregated memory is supposed to provide low-latency, high-throughput access to memory resources, even when these resources are not physically co-located with the

CPU. The system should minimize performance overheads associated with remote memory access, ensuring that applications can leverage the large memory space without compromising speed and efficiency[134].

2.2.2 Features of DM. Unlike the integrated monolithic servers in traditional architecture, research on DM architecture focuses on the novel design of *computing nodes* and *memory nodes*. In other words, servers will serve the applications as different roles. In the past, computing resources obtained data from memory and storage devices hierarchically, managed by the memory control unit inside each CPU. In disaggregated architecture, memory and storage resources are generally managed in pools. From the view of the computing units, disaggregated memory constitutes a memory entity far from local. Therefore, in most academic work, "*far memory*" is often used to refer to disaggregated memory.

Firstly, the computing nodes should keep the following additional capabilities to support high-performance data offloading to far memory.

- **DM latency awareness:** Accessing a "far" memory device on the computing nodes involves longer access latency and lower bandwidth than using local DRAM[85]. One should leverage and amortize this overhead by carefully designing the far memory access patterns of local data processing.
- Smart data offloading: Compute nodes are responsible for the overall computational flow control of the program[18]. Programs often have complex iterations that control the lifecycle of copies of the required data, and these characteristics make it very difficult to offload data to far memory.
- Fine-grained resource orchestration: Application utilizes multiple and heterogeneous computing, memory, network, and storage resources to complete user requests[120]. Improving system utilization requires efficient scheduling methods to reclaim and allocate these resources well.

Secondly, the memory nodes provide key functions to support large-capacity data storage and efficient data management strategy.

- Data-related workflow control: Memory nodes are responsible for storing the arriving data and fetching the corresponding data[145]. However, far memory access has higher latency compared with a storage system. Thus, users must be careful with each step of the data flow process.
- High-performance far memory access: DM serves as near-computing memory devices that show much higher performance than traditional data providers. For example, solid state disks (SSD) with much higher bandwidth and IOPS than disks can be recognized as a far memory space[152].
- Heterogeneous memory management: Memory nodes often consist of heterogeneous memory devices on different far-memory access paths, especially when applications need to access large amounts of memory. Heterogeneous memory management is important for overall bandwidth and efficiency[138].

2.2.3 DM system levels. Generally, existing works often design DM systems at three levels.

- Architecture-level design: One can design the DM system in a bottom-up way. The key objective is to fully utilize hardware resources in a cost-effective manner. Computing and memory nodes are expected to be connected through high-bandwidth networks. A good organization of them can provide high task performance and high data throughput.
- **OS-level design:** With appropriate hardware support, the OS-level DM design aims to utilize hardware efficiently. Given drivers and protocols for each new-added memory and network device, one can adjust the existing operation systems to provide application-transparent DM management. One can also design virtualized DM runtimes for resource isolation and task throughput.
- **Application-level design:** DM systems can also be designed in a top-down way. For example, one can design application-specific optimization methods that improve the efficiency of computing workflow and far memory access. Accelerators with bare metal performance also need to consider memory extension.

#### Survey of Disaggregated Memory • 5



Fig. 2. Overview of this survey.

These designs often provide non-transparent far memory access with easy-to-use and high-performance programming interfaces.

#### 2.3 Paper Overview

We summarize the key words related to disaggregated memory in this paper, referring to Figure 2.

**DM Hardware:** Disaggregated memory requires new architecture/hardware designs for better elasticity and efficiency. Research on DM devices can be classified into three types. Memory devices usually communicate with CPUs through on-chip buses and interfaces. Generally, two types of memory devices, RAM and NVRAM, can act as disaggregated memory, with different data persistence capabilities. The network connections used for disaggregated memory (or far memory) mainly involve network interface cards (NIC) and programmable-chip-based smart NICs. In addition, hardware feature analysis is also a major part of previous work, including bandwidth, capacity and power.

**DM Architectrue:** Disaggregated memory architecture refers to the organization ways of Computing Nodes (CM) and Memory nodes (MN). There are many ways for computing nodes to access external memory spaces in the existing server architecture. We summarize the DM architecture into 4 categories according to the end number of each far memory path. The number of compute nodes and memory nodes involved in the far-memory access path determines the design and system implementation of different disaggregated memory architectures.

**FM System/Runtime:** On top of disaggregated memory architectures, runtime design and optimization strategies for far memory access is crucial. They provide the necessary support for high-performance far memory access and efficient memory resource scheduling. We summarize the software environment of disaggregated memory, including system-level OS adaption, runtime design and task scheduling. These systems are designed to support highly flexible resource usage, low overhead data processing, and high task/data throughput.

**Application Optimization:** There is a trend towards adapting and redesigning far memory systems for specific scenarios, i.e., service-aware far-memory design, including serverless, elastic computing, fault-tolerant computing, and power-optimized far memory systems. In addition, related works also design corresponding specialized far-memory systems for dedicated big data tasks, such as graph processing and neural network training and inference. Application acceleration mainly adopts data management in the computing node, which considers

how to identify, partition, cache, and prefetch the cold data at fine-grained granularity. Domain-specific far memory system design can optimize the performance of elastic computing, graph processing, and DNN-supported systems.

#### 3 DM HARDWARE

Disaggregated memory requires new architecture/hardware designs for better elasticity and efficiency. This section first introduces cluster-level and server-level DM architectures. We then discuss DM scheduling methods.

To figure out how to deploy disaggregated memory on physical servers, this section introduces hardware that can be recognized as disaggregated memory, i.e., far memory. General-purpose computing nodes, such as CPUs and GPUs, can access far memory devices directly across hardware bus and slots, or through a network card. Thus, the key components of disaggregated memory are memory devices and network devices, i.e., **disaggregated memory** = (**network connections**) + **memory devices**. The "()" refers that network connections can be not included. We also provide a performance comparison of advanced memory, storage, and network devices.

## 3.1 Memory devices

Memory devices usually communicate with CPUs through on-chip buses and interfaces. The first column of Table 1 shows the commercially used memory devices, the slots, channel, and protocol. We compare different memory and memory connection devices according to the plug-in slots, data channels, and data transfer protocols, as shown in the rows of Table 1. Slots refer to the way memory devices and compute motherboards are interconnected and are usually hardware interfaces. The channel refers to the available data channels decided by the devices or the slots. Based on typical slots and channels, supported data transfer protocols are also required.

Generally, there are two types of memory device that can act as disaggregated memory, with different data persistence capabilities. *Volatile Random Access Memory* (RAM) includes both CPU DRAM, and GPU DRAM. The high-bandwidth memory (HBM) [70] and hybrid memory cube (HMC) [109] are proposed for higher density of DRAM memory. In addition, *Non-Volatile Random Access Memory* (NVRAM, short as NVM) can store data even when the power is off. Flash memory is a type of NVRAM commonly used in the design of SSDs [13]. Persistent Memory (PM) is a type of NVRAM; taking the Intel Optane NVM [151] as an example, PM is located in the system DIMM slots with an extra power control module. Furthermore, academia and industry are researching and experimenting with promising candidates for next-generation memory ("x"RAM), including magnetoresistive memories (MRAM), resistive memories (ReRAM), phase-change memories (PCRAM), and ferroelectric memories (FeRAM), etc.

**DIMM slots and protocols**: The commercially used DRAMs are Double Data Rate (DDR) DRAMs. With the expansion of the bit width of data readable by DDR and the upgrading of hardware structure, DDR2 can accomplish 4-bit data pre-reading, DDR3 is 8-bit, DDR4 is 16-bit, DDR 5 is 32-bit, and DDR6 is 64-bit. Every next version of DDR maintains double data bit width and thus has double data transmission bandwidth. During this evolution process, the working voltage of DDR has also been gradually lowered, and the working frequency has been gradually increased. Adding the block number can also add data bandwidth. With a single block achieving 34 GB/s in DDR4, one can have at least 200GB/s data bandwidth on eight DDR4 banks. The latest commercial DDR version is DDR5 (until 2024), while DDR6 is coming soon[1].

**PCIe slots and protocols**: PCIe bandwidth is keeping an increasing trend, which is roughly faster than double times every three years [11]. For example, PCIe 3.0 is twice as fast as PCIe 2.0, and PCIe 5.0 is twice as fast as PCIe 4.0 with backward compatibility. Currently, the commonly used duplex-lane PCIe 4.0 x16 version commonly has about 64GB/s of bandwidth. PCIe Gen 5.0 with total throughput of 128GB/s and PCIe 6.0x16 with 256GB/s bandwidth is considered. The rapid development of PCIe has allowed disaggregated memory devices to approach the bandwidth of traditional DRAM. In addition, PCIe also has better interface scalability with extendable roots.

<sup>,</sup> Vol. 1, No. 1, Article . Publication date: August 2024.

#### Survey of Disaggregated Memory • 7

Components	Device	Slot	Channel	Protocol				
	DRAM	DIMM/HBM/HMC	Bus	DDRx, (x is up to 6)				
	GPU DRAM	PCIe and NVLink	NVLink and NVLink switch	GDDR5, GDDR6				
Memory Device	xPU DRAM	OAM	OAM on x16 links	PCle Gen5/6 x16				
	NVM	DIMM with	Bus	Customized				
(Sec. 3.1)	144101	extra power support	Dus					
(500.5.1)	Persistent memory	DIMM or PCIe	Bus or PCIe switch	Customized				
	CXL memory	PCIe/DIMM	SATA, PCIe x32, x64, etc., bus	CXL 1.0, 2.0, 3.0, etc.				
	SSD	PCIe, M.2, U.2	SATA, PCIe x4, x8, etc.	AHCI, NVMe				
	"x"RAM							
	(MRAM, ReRAM,	DIMM or PCIe	Bus or PCIe switch	Customized				
	PCRAM, FeRAM, etc.)							
Network	NIC	PCIe	PCIe x4, x8, x16, etc.	TCP/IP				
Connection	RDMA NIC	PCIe	PCIe x4, x8, x16, x32, etc.	RDMA				
Device	DPU (for SmartNIC)	PCIe	PCIe x16, x32, etc.	Customized				
(Sec. 3.2)	FPGA (for SmartNIC)	PCIe	PCIe x16, x32, etc.	Customized				

Table 1. Memory and Connections that can be organized as Disaggregated Memory.

**Dedicated slots and protocols**: Recent works propose direct connection devices, such as NVLink, which can support direct memory communication between GPUs [103]. Tavallaei et al. propose an open accelerator module (OAM) [110] supporting various interconnect topologies for new types of hardware accelerators (such as GPU, FPGA, ASIC, NPU, TPU, IPU, et al). IBM Power 9 allows direct memory access between servers via specific cables and the OpenCAPI protocol [10]. Recently, Intel has developed CXL CPU motherboards, which are based on the enhanced protocol PCIe 5.0 and are capable of supporting approximate access latency for cross-core (socket) accesses, as well as interconnecting memory pools with accelerator pools [3]. Several works [46, 169] have proposed the use of electronic signals instead of photonic signals for memory interconnections.

#### 3.2 Network connections

The network connections used for disaggregated memory (or far memory) mainly involve network interface cards (NIC) and programmable-chip-based smart NICs. Details are shown in Table 1.

**RDMA NIC**: RDMA technology is one of the main methods of far memory access [18, 29, 49, 117, 134]. RDMA is widely used in data-intensive and compute-intensive scenarios as a remote memory direct access technique that bypasses the kernel. RDMA has three network protocols: Infiniband, RoCE, and iWARP. The main manufacturer of InfiniBand network is Mellanox, which is now acquired by NVIDIA. RoCE (RDMA over Converged Ethernet) is a network protocol that allows RDMA to be performed over Ethernet with lossless Ethernet for low-latency operation and complex configuration process. The AIPs provided by OFED [9] and IB face problems with usage, performance, and security. Thus, some works provide high-level programming interfaces [15, 41, 129].

**Smart NICs**: By designing high-speed on-chip memory networks with routing logic, one can deploy thousands of memory modules on a single chip and realize on-chip interconnections [104]. Meanwhile, there have been works to design customized smart NIC devices to replace generic NIC devices and reduce the overhead of far memory access [50]. Also, there is some work designing programmable switches [72] to work with server-level RDMA or smart NIC devices for cabinet-level memory interconnections. MIND [72] places memory management logic in the network fabric, and they find that centralizing memory management in the network permits bandwidth and latency-efficient realization of in-network cache coherence protocols. At the same time, programmable switch ASICs support other memory management logic at line rate.

#### 3.3 Characteristic of DM Hardware

We illustrate the performance of the mentioned devices with different versions in bandwidth (Figure 3), capacity (Figure 4), and power (Figure 5). In these figures, GDDR is the main memory for GPUs. LPDDR DRAM is used in edge devices such as mobile phones and DDR DRAM serves as the main memory for servers. CXL memory, persistent memory, and SSDs are memory devices that can store data. DPUs, RDMA, and NICs represent network cards used for data transfer with optional data caches or control logic.

**Bandwidth analysis:** Available data bandwidth is one of the key metrics of data access performance. As local memory, DRAMs have large memory bandwidth with hundreds of GB per second (GB/s) to support high-performance data access for computing units. GDDR memory, designed for GPUs, delivers the highest memory bandwidth to handle the large-scale parallel processing demands. LPDDR DRAMs for mobile devices with relatively lower power consumption and less data bandwidth. Persistent memory and SSDs are non-volatile storage devices with large memory capacities but offer lower bandwidth. Oftentimes, new-generation memory device exhibit wider bandwidth. For example, CXL 3.0's 70GB/s far exceeds CXL 1.0's 15.2GB/s in the same socket physical interface. An important trend is that networking devices and storage devices are evolving at a rapid pace, much faster than native DRAM-based DDR. while individual DRAMs still lead in performance today, their overall bandwidth may reach the upper bound due to limited on-chip DIMM slots. State-of-the-art interconnects is establishing scalable plug-in interconnect methods based on PCIe, OCP Accelerator Module (OAM), and others.

**Capacity analysis:** GDDR memories used in GPUs have higher bandwidth than DDR DRAMs used in servers, but with lower capacity due to specific application requirements. GDDR memory and LPDDR DRAM can work with lower capacities compared to the DDR DRAM in servers, since GPUs and mobile devices focus on performance and power efficiency, whereas servers need larger memory capacities to handle high data processing demands. CXL is an interconnect technology that not only enables data transfers but also incorporates a fixed memory capacity of 128GB. Non-volatile storage types, such as persistent memory and SSDs, are known for their large storage capacity compared to other memory devices.

**Power analysis:** Power consumption varies significantly across different types of memory and interconnect technologies, reflecting their different design goals. For instance, DPU has a power draw of around 25 W. However, a LPDDR DRAM consumes less than 1 W. When assessing the power consumption of a system, we need to consider the total power consumption rather than focusing solely on individual device power. For instance, while a single LPDDR DRAM may consume less power, there may be numerous such modules in a mobile device, leading to a significant cumulative power.

#### 4 DM ARCHITECTURE

Disaggregated memory architecture refers to the organization ways of Computing Nodes (CM) and Memory nodes (MN). There are many ways for computing nodes to access external memory spaces in the existing server architecture. Generally, it is commonly used that far memory path = processing units + local memory + connections + far memory.

In real physical machines, FMs can be on the same machine or on different machines interconnected by a network. Figure 6 shows an example of the hardware deployments of a computing node and a memory node. DM can be built on a single server or paired servers. Compute nodes usually include computation units, caches, and local memory, while memory nodes mainly include memory control units, memory interconnects, and far memory. Far memory space often refers to isolated free-accessed memory regions located on certain memory devices or servers.

In this section, we summarize the DM architecture into 4 categories according to the end number of each far memory path. The number of compute nodes and memory nodes involved in the far-memory access path determines the design and system implementation of different disaggregated memory architectures. Similarly to

<sup>,</sup> Vol. 1, No. 1, Article . Publication date: August 2024.

#### Survey of Disaggregated Memory • 9







Fig. 4. Memory capacity of different far memory devices, including memory, storage and network card.



Fig. 5. Power usage of different far memory devices, including memory, storage and network card.

the mapping of instruction and data stream in CPU architecture, we use single/multiple computing node (CN) Single/Multiple memory node (MN) to describe the disaggregated memory architecture.

10 • Jing Wang, et al.



Fig. 6. Physical disaggregated memory testbed with various memory and network devices.

#### 4.1 Single CN Single MN

The basic DM architecture is one single computing node with one single memory node. The computing node can be CPU processors, GPU processors, or any other accelerators. We use FM paths to represent the data paths from one computing node and a memory node.

*4.1.1 CPU far memory.* Currently, there are mainly three types of far memory access paths: the memory-based FM paths, network-attached FM paths, and controller-involved FM paths, as shown in Table 2.

**Memory-based FM path:** This type of FM often utilizes persistent memory like NVM and fast storage like SSD, as shown in the second column of Table 2. These works extend the memory inside each server, which are called *vertical* or *intra-node* far memory. Some works focus on scaling one's own memory to special memory nodes with a high-capacity NVM in the rack [20, 71, 128, 151, 160]. They try to design persistent, replicated, fault-tolerant coherence protocols on distributed file systems for Optent DC persistent memory modules. Many works adopts SSD [19, 29, 60, 68, 152, 162] as data offloading space. These works take the fast data read of SSD to accelerate I/O operation and design related resource management methods for high efficiency.

**Network-attached FM path:** It often connects memory devices through network cards, as shown in the third column of Table 2. These works extend the memory between servers, which are called *horizontal* or *inter-node* far memory. Many works utilize RDMA NIC [15, 18, 27, 41, 49, 88, 117, 129, 135, 160] as far memory. They explore the advantages of the RDMA protocol, including the non-CPU-involved computing workflows, one-sided operations, and variable data transfer chunk size. Some recent works adopt smart NICs implemented by FPGA or DPU as far as memory [50, 108]. These works offload some data processing workloads to the low-power devices to improve memory efficiency. However, the proposed smartNIC-based works often suffer performance bottlenecks compared with RDMA [148]. For cross-rack far memory access, a programmable switch is required. For example, one can place memory management logic in the network fabric and utilize a PCIe switch to enable rack-level high-bandwidth memory communication [72].

**Controller-involved FM path:** DM devices based on reconfigurable hardware involve specific controllers for highly efficient memory management, , as shown in the forth column of Table 2. For example, String Figure [104] builds a large memory pool with thousands of memory nodes and tens of CPUs, designing a hybrid routing protocol and a set of network reconfiguration mechanisms. Beacon [53] proposes a memory management framework to enable memory expansion with unmodified CXL-DIMMs. Mind [72] uses a control plane to track the total amount of memory allocated on each memory blade with leveraging its global view of allocations for load-balanced memory pool management.

4.1.2 *GPU far memory.* Heterogeneous far memory systems and runtime designs are new trends for today's memory-hungry workloads, such as training and inference of large foundation models. These workloads often rely on general-purpose accelerators, i.e. GPUs for high parallelism. Therefore, GPUs require orders-of-magnitude higher memory bandwidth than traditional CPU-based systems [75]. However, the memory space of GPU is still limited and the capacity of such high-bandwidth memory tends to be relatively small. There are several GPU far memory access paths, including CPU memory, NVLink-based remote GPU memory, direct-storage-based

#### Survey of Disaggregated Memory • 11



Table 2. Taxonomy of far memory (FM) access paths.

storage space, direct-rdma-based remote memory, etc. Thus, GPUs can access far memory through several paths, including unified memory (UM), GPU direct storage (GDS), and GPU Direct RDMA.

**Unified memory:** In GPU far memory architecture, Unified Virtual Memory (UVM) is an essential way to integrate GPU and CPU memory spaces. This arrangement facilitates data sharing between CPU and GPU, with no transfer exposed to apps and users. Recent work focuses on leveraging CPU memory for specific purposes and optimizing current management strategies. Swap Advisor [52] proposes a general swapping mechanism that can swap deeper and wider neural networks on limited GPU memory. The memory consumption is reduced by smartly swapping the model parameters that are updated at the end of each iteration and used by the next iteration. To optimize GPU's memory access overhead, Grus [141] co-utilizes different memory-sharing operations of CUDA UVM when running graph workloads. Shao et al. [123] investigated the implications of GPU UVM oversubscription. They evaluate the performance sensitivity of applications with over-subscribed memory size on GPU. Some works adopt memory efficient techniques with static and dynamic scheduling strategies to execute large-scale graph random walk on GPU and CPU memory[58, 140]. Soroush et al [22] explored the hidden performance implication of GPU memory management methods and developed a performance model that can predict system overhead for each memory management method.

**GPU direct memory access:** The integration of GPU memory with high-speed storage devices like SSDs and NVMe drives is another frontier in GPU far memory management. Besides storage devices with sufficient bandwidth for memory management, novel network technology and hardware like RDMA allow data access faster and lower-costly. Current work focuses on reducing data transfer volume by achieving higher offload and prefetch efficiency and achieving higher bandwidth and lower latency with hardware and software methods.

GPU Direct-storage (GDS) techniques [102] create a direct data path between local or remote storage, such as NVMe or NVMe over Fabrics (NVMe-oF), and GPU memory. By enabling direct-memory access (DMA) near the network adapter or storage, it moves data into or out of GPU memory—without burdening the CPU. GPUDirect RDMA [100] is a technology introduced in Kepler-class GPUs and CUDA 5.0 that enables a direct path for data exchange between the GPU and a third-party peer device using standard features of PCI Express. BAM [111] proposes a big accelerator memory system that enables GPUs to orchestrate high-throughput, fine-grain SSD storage access. It mitigates I/O traffic amplification by enabling the GPU threads to read or write small amounts of data on-demand through highly concurrent I/O queues.

## 4.2 Single CN Multiple MN

One can utilize multiple far memory nodes for a single computing node to improve data throughput. The compute nodes need to be carefully designed to utilize external heterogeneous memory resources. If there are multiple far memory devices, one may hierarchically utilize the hybrid memory devices or use the multiple far memory backends in parallel.

*4.2.1 Hierarchical hybrid memory.* Hybrid memory combining fast and slow heterogeneous memories provides a promising direction to increase memory capacity.

**General-purpose hybrid memory design:** Some works focus on hybrid memory management for general works. Hybrid<sup>2</sup> [130] considers a hybrid memory system composed of memory technologies with different characteristics that use only a small fraction of the near memory as a DRAM cache and further leverages DRAM cache for large data migration. Hybrid TLB Coalescing [105] proposes a novel HW-SW hybrid translation architecture, which can adapt to different memory mappings efficiently. UH-MEM [81] proposes utility-based hybrid memory management, a new page management mechanism for various hybrid memories that systematically estimates the utility of migrating a page between different memory types and uses it to guide data placement.

**Persistent-memory-based design:** Some works focus on hybrid DRAM and NVM memory hierarchy. Memos [89] proposes a memory management framework that can hierarchically schedule memory resources over the entire memory hierarchy, including cache, channels, and main memory comprising DRAM and NVM simultaneously. Panthera [131] proposes a semantics-aware automated memory management technique for Big Data processing over hybrid memories (DRAM and NVM) by analyzing user programs to infer their coarsegrained access patterns.MNEME [125] proposes an architectural change where each long bit line in DRAM and NVM is split into two segments by an isolation transistor.

*4.2.2 Multi-path far memory*. The coexistence of multiple far memory backends is a boost for offloading efficiency and system flexibility. Multiple far memory access paths on multiple memory backends (multi-end DM) may become the future solution for improving data throughput.

**Multiple far memory paths for CPU:** XMemPod [29] established access to the memory of other VMs through host- and RDMA-coordinated memory hierarchy. On the other hand, TMO [152] implements its memory pool through zswap and swap. Buddy Compression [34] compresses and allocates data between GPU device memory and buddy memory such as NVLink [103], to reduce reallocations within device memory and adapt to compressibility changes. XDM[138] design parallel multiple far memory backends and can adaptively switch the backends with dedicated far memory path configurations.

**Multiple far memory paths for GPU:** Buddy Compression [34] leverages compression to utilize a larger buddy memory from the host or DM, effectively increasing the memory capacity of the GPU. It compresses the offloaded memory locally as compressed memory and utilizes buddy storage with alternatives, including CPU memory, DM, unused peer GPU memory, and NVMe SSDs. ZeRO-Infinity [112] is a novel heterogeneous system technology that leverages GPU, CPU, and NVMe memory to allow for unprecedented model scale on limited resources without requiring model code refactoring. G10 [158] scale the GPU memory capacity with transparent data migrations by integrating the host memory, GPU memory, and flash memory into a unified memory space. GMT[31] design tiered GPU Memory systems with 3-tier hierarchy comprising GPU memory, host memory and SSDs and selectively chooses the provided far memory paths.

## 4.3 Multiple CN Single MN

Oftentimes, one can build a large virtualized memory pool that serves multiple computing nodes. In this way, the computing node may concentrate on the computing and memory access procedure design, which offloads the

<sup>,</sup> Vol. 1, No. 1, Article . Publication date: August 2024.

#### Survey of Disaggregated Memory • 13



Fig. 7. Typical far memory access paths on disaggregated memory.

coordination of the multiple memory nodes to the single memory node, i.e. memory pool. We list four categories of remote memory access (RMA) and far memory access (FMA) methods, as shown in Figure 7.

Far memory pool for CPU. (1) Cross-socket RMA/FMA: One can access the remote memory space of 4.3.1 another CPU sockets through the corresponding Integrated Memory Controller (IMC) on Non-Uniform Memory Access (NUMA) architecture, involving the latency and bandwidth difference between Local Memory Access (LMA) and Remote Memory Access (RMA)[69]. (2) Network-based RMA/FMA: Network-based RMA is usually used to transmit instructions and data through networks, such as TCP or UDP transmission [5]. The problem is that data transfer involves many CPU operations, interrupting the original computing on CPUs. Direct Memory Access (DMA) and Remote Direct Memory Access (RDMA) are often adopted for few-CPU-involved features and high-speed data transfers [82]. RDMA-based far memory has an order of magnitude difference in memory access latency and data access bandwidth compared to local DRAM memory [137]. Typically, the network interface cards (NIC) of the memory device need to be interconnected with the NIC of the memory device in another server through specific network cables. (3) Memory-based FMA: Persistent memory devices usually interact directly with the CPU and have their own data processing patterns and programming models[151]. The newly proposed CXL [3] technology is based on PCIe interfaces and it treats memory device as a special NUMA node (memory node) without CPU. NUMA awareness significantly impacts performance and requires more thread scheduling and memory allocation optimizations [159]. (4) Storage-based FMA: To ease the performance gap between CPUs and I/O devices, operating systems reserve buffers in memory managed with virtual file systems (VFS) and reduce actual hard disk I/O reads and writes [168]. Caching and evicting of I/O data can be controlled by LRU list configuration and other system parameters [7]. One usually accesses far memory by page swapping, swapping infrequently-used page data to a storage device such as a disk. As storage access performance improves, one can also access far memory directly through an I/O interface [86].

4.3.2 Far memory pool for GPU. Multi-GPU configurations amplify the computational power of many GPUs for serving large-scale applications like Large Language Model (LLM). However, they introduce challenges in managing memory throughout GPUs and accelerate computation and communication. Recent work aims to use a multi-GPU system for distributed deep learning and optimize memory management and communication mechanisms. COARSE [147] is a disaggregated memory extension for distributed deep learning (DL) training on multi-GPU. It proposes dynamic tensor routing and partitioning to fully utilize the non-uniform serial bus bandwidth varied across different cloud computing systems. MC-DLA [67] proposes a memory-centric deep

learning system architecture that aggregates a pool of capacity-optimized memory modules within the device-side interconnect for transparent memory capacity expansion. Griffin [21] is a holistic hardware-software solution to improve the performance of NUMA multi-GPU systems. Griffin introduces programmer transparent modifications to both the IOMMU and GPU architecture, supporting efficient runtime page migration.

## 4.4 Multiple CN Multiple MN

The Multiple CN Multiple MN architecture means that we organize multiple computing nodes with multiple memory nodes on multiple servers. We summarize them into two types, the homogeneous CDM architecture and heterogeneous CDM architecture, as shown in Table 3.

Table 3. The two architectures of cluster-level disaggregated memory. In this Table, each Computing components represent the computing units(CPU, GPU, etc.), caches (L1, L2, last-level cache, etc.), and local memory (DRAM, HBM, etc.). Disaggregated memory (DM) represents far/disaggregated devices.



4.4.1 Homogeneous CDM architecture. In a homogeneous CDM, the disaggregated memory resources are usually deployed on the same physical machine as the compute components, which can be interpreted as the compute nodes adding extra memory space. It can achieve high performance, but it requires careful management due to the complex interconnection between computing components and DM components. Since the network channels need to handle hybrid tasks, one should manage the network channel carefully to maintain data consistency and address the bandwidth competition problem.

**Computing and memory networks (CMN):** The computing component often accesses far memory through intra-server I/O fabrics on the same mainboard. The most commonly used fabric is PCIe, which is used on both CXL memory [3, 78] and SSD [68, 152, 154]. The computing component can also access memory resources of the pass-through server through computing and memory networks, as shown in Table 3-a. This network is responsible for far memory data transfer with remote servers. Commonly used devices include TCP/IP NIC Switch [5, 82], RDMA NIC switch [18, 29, 34, 49, 137], FPGA switch [108], CXL PCIe switch [53, 72, 94].

Thymesisflow [108] and Memtrade [94] are mainly based on CMN. Thymesisflow [108] features two switchable roles: computing role and memory-stealing role. The computing role introduces remote memory to a physical address space range; the memory-stealing role reserves a portion of the local memory on the host system, which is exposed as disaggregated memory to a neighbor host. Memtrade [94] adopts a producer-and-consumer model that allows memory producers to lease both their unallocated memory and allocated-but-idle application memory to remote consumers for a limited period of time.

Looking ahead, Advancements in CMN pave the way for homogeneous architectures to provide high-performing all-hardware disaggregated memory solutions. For example, Gen-Z [35] aims to provide a disaggregated I/O fabric. In addition to CXL [3], OpenCAPI [10] and NVLink2 [103] are also promising cache-coherent attachment technology for off-chip peripherals.

4.4.2 Heterogeneous CDM architecture. In heterogeneous CDM, DM resources are often decoupled with the computing components. It maintains higher flexibility of DM management due to the complete disaggregation of memory resources. The architecture also leaves data control and data consistency to memory pool. However, it may suffer a noticeable overhead of far memory management, especially when the far memory pool is large. Different from designs in homogeneous CDM architecture, one can adopt different NIC devices or isolated data transfer channels for distributed computing network (DCN) and memory access network (MAN).

**Distributed computing network (DCN):** The computing units rely on DCN to handle data communication to support parallel designs (Table 3-b). The DCN is mainly realized through a classical distributed computing network, responsible for computation distribution and task distribution [5, 24, 32].

**Memory access networks (MAN):** Similar to CMN, MAN handles data transfer with far memory devices (Table 3-c). There have been works utilize RDMA to access remote DRAM [18, 29, 49]. Some prior works design programmable FPGA devices or smart NICs to access the remote DRAM [50, 108, 133]. Recent works are more focused on utilizing PCIe-based CXL switches (CXL 2.0) to access remote CXL memory or SSDs directly [3, 53, 154]. Pond simulates CXL as a NUMA node without CPU cores and designs a CXL-based full-stack memory pool with an accurate memory allocation prediction model [78]. The CXL-ANNS enables highly scalable approximate nearest neighbor search (ANNS) services by disaggregating DRAM via CXL and placing all essential datasets into its memory pool [56]. Yang et al. [154] explores the design space of flash devices instead of memory devices on new CXL technologies to overcome the memory wall.

**Memory network (MN):** The disaggregated memory pool is organized by the MN, which is often based on on-chip design with high bandwidth and high memory density [104, 130] (Table 3-d). Many works design large memory nodes based on memory networks. Kim et al. propose a distributor-based network on HMC to reduce the network diameter while properly distributing the bandwidth across different routers [64]. Zhan et al. propose a DRAM-based unified memory network architecture on Hybrid Memory Cubes (HMC) that provide an intelligent I/O interface to reuse the intra-memory NoC as the network switches for inter-memory communication [92, 156]. Poremba et al. [109] observe that placing NVM cubes in a specific order in the MN improves performance by reducing the network size/diameter up to a certain NVM to DRAM ratio.

## 5 SYSTEM-LEVEL ADAPTION AND DESIGN

On top of disaggregated memory architectures, runtime design and optimization strategies for far memory access is crucial. They provide the necessary support for high-performance far memory access and efficient memory resource scheduling. Based on the above architecture-level design, this section summarizes the software environment of disaggregated memory, including system-level OS adaption, runtime design and task scheduling. These systems are designed to support highly flexible resource usage, low-overhead data processing, and high task/data throughput.

## 5.1 Operating System Adaption

The OS-level design leaves all the control to the OS without application-specific adaptions and thus is suitable for any general-purpose programs. One can add new control modules into the OS to manage the DM, detailed in Sec. 5.1.1. One can also modify the existing modules to support the now far memory access path, detailed in Sec. 5.1.2.

*5.1.1 Addition of control modules.* Adding new control modules into the OS can manage the disaggregated memory devices well. We categorize related methods into three groups, including the separate components control, separated power control, and separated memory control, as shown in Table 4.



Table 4. Different ways of control module addition.

**Separated components control:** A basic approach is to use isolated control subsystem for separated CPU, memory, network, and storage. Each device monitor can operate locally for its functionality and only communicates with other monitors for real-time resource requirements. LegoOS [120] introduces a modular system named *splitkernel* to decouple the traditional OS functionalities into loosely coupled monitors, each of which runs on and manages a hardware component. Each control module has only two global tasks: orchestrating resource allocation across components and handling component failure.

**Separated power control:** Energy costs can be saved if the power of CPU and memory are managed separately and provisioned on demand. Zombieland system [99] provides an effortless way for disaggregating the CPU-memory couple at the power supply domain level. CPU and memory still share the same board, but their power supply domains are separated. It also uses a new ACPI sleep state (called zombie), which allows the suspension of a server (thus saving energy) while making its memory remotely accessible. In addition, the consumed memory node is equipped with specialized hardware controllers with separated power control [50, 72].

**Separated memory control:** Independent control modules can manage memory devices, handling data read and write on memory blocks and the data coherency of the memory pool. At the hardwre level, a classic solution is to build a sizeable on-chip memory pool organized with routers and managed by a global memory controller. For example, String Figure [104] proposes a high-throughput, elastic, and scalable memory network architecture. It generates random topology for high network throughput and near-optimal path lengths in large-scale memory networks. At the software level, one can use appropriate software to manage memory resources in both memory blades and switches. For example, BEACON [53] is located near the memory pool and it can process data in the memory pool and accelerate data routing in the switch. This design can leverage the high communication bandwidth provided by CXL without modifying cost-sensitive DRAM. Further, there are several works adopting hardware-software co-design ways to manage the memory pool. For example, ThymesisFlow [108] introduces a software-defined hardware/software co-designed datapath on top of the IBM POWER9 [54]. It embraces the OpenCAPI [10] cache coherent attachment technology that is available today on IBM POWER9 processors [54]. It uses a Remote Memory Management Unit (RMMU) to manage the address translation and make each



Table 5. Differerent ways of existing module modifications.

memory section (viewed as far memory space) independently and "hotplugged" at runtime. Clio [50] proposes a hardware-based memory disaggregation solution that has the right amount of processing power at memory nodes. Implemented on FPGA hardware, Clio virtualizes and manages disaggregated memory at the memory node, including a new hardware-based virtual memory system, a customized network system, and a framework for computation offloading.

*5.1.2 Modification of existing modules.* Many prior works propose new technologies at the software level to fully use existing mechanisms in the OS. These designs can be quickly adapted to deploy on commercial hardware architectures while ensuring transparency to the task. We summarize prior works into three categories, including network card replacement, storage replacement, and memory replacement, as shown in Table 5. Traditionally, the swap space is a block partition on the disk, which is managed as a file. In recent years, researchers have replaced the original swap backend with higher-performance memory devices to optimize the data swap performance, such as RDMA network cards [18, 29, 49, 88] and SSD [29, 68, 152]. A basic implementation involves adding a kernel patch and a new module into the OS, utilizing data transfer procedures with far memory access.

**Network card replacement:** Many works have been devoted to adapting page swap mechanisms in existing systems to RDMA-based far-memory environments, as shown in Table 5. For example, Liang et al. [83] designs a high-performance networking block device (HPBD) over InfiniBand fabric, which serves as a swap device for efficient page transfer to/from remote memory servers. Infiniswap [49] is the first far memory system based on a traditional page-swap mechanism that utilizes RDMA as swap backends. Cao et al. [29, 88] designed the FastSwap tool to efficient page swapping via RDMA in VMs. FastSwap can leverage idle host memory and redirect the VM swapping traffic to the host-guest compressed shared memory swap area. Fastswap [18] explores the role of RDMA-based far memory in enhancing task throughput. Sherman [142] introduces a write-optimized B+ Tree index for disaggregated memory to boost the write performance of RDMA-based far memory access.

**Storage replacement:** Commercial cloud-native corporations often use storage to offload data due to the high cost of RDMA deployment, as shown in the third column of Table 5. Google in its work Zswap [68] showed that software-defined far-memory systems can actively compress cold memory pages in the DRAM. Meta's data centers use TMO system [152] to offload task data to far memory in heterogeneous datacenter environments. It designs a Pressure Stall Information (PSI) kernel to automatically adjust how much memory to offload to heterogeneous devices according to the device's performance characteristics and the application's sensitivity to memory-access slowdown. Koutsovasilis et al. [66] presented a memory balancing policy that autonomously

migrates memory pages across local and disaggregated memory. Xmempod [29] uses different software managers to support data swap on inter-VM DRAM-based far memory, RDMA-based far memory, and block-device-based far memory. Hyfarm [137] utilizes both SSD-based intra-server far memory and RDMA-based inter-server far memory that replaces the disk-based swap backend and thus balances the memory usage between servers well. Several works have utilized persistent memory as a remote storage space to expand memory resources further, taking advantage of its much better performance benefits and persistence over disk to store data [121, 167]. Clover [128] proposes to separate persistent memory and connect them to compute nodes in a remote storage fashion, allowing all compute nodes to directly access and manage storage nodes.

**Memory replacement:** Recent works also utilize volatile and non-volatile memory devices as far memory space, as shown in the fourth column of Table 5. Data units are often called objects with variable data sizes, which can be stored on NVM, CXL, and optical memory devices. A key idea is to design data indexing methods based on specific efficient data structures since the data offloading and fetching in the DM system is similar to traditional key-value store operations. For example, classic DM indexing works [142, 170] use B+ tree to build range indexes. dLSM [143, 145] proposes optimized LSM-tree compared with the traditional B-tree indexing for disaggregated memory, leveraging near-data computing with RDMA-specific customizations and tuning in byte address granularity with high system performance. Skadi [51] and SMART [90] argue that radix tree is more suitable for DM than the B+ tree due to smaller read and write amplifications. They design high-performance locks for the leaf nodes and read-delegation and write-combining techniques to reduce redundant I/Os.

## 5.2 Runtime Design and Optimization

Far memory systems for virtual environments are designed for better runtime management. Different from designs at the OS level, virtualization-level design is software-defined and non-transparent for applications. They often maintain more flexible memory management strategies, including memory reclamation, memory allocation, and memory calls.

Virtualized environments are adapting to more far-memory devices, which will be one of the main directions for research on disaggregated memory systems in the following decades. The virtualization method includes bare-hardware virtualization, virtual machines with inner OSs, container-based virtualization, and Java virtual machines to provide isolated resources.

*5.2.1 Memory and network virtualization.* In a virtualization environment, memory and network resources can be abstracted through memory address mapping, memory region isolation, and network channel allocation.

- *Memory address mapping*: For memory virtualization, the system establishes a complex mapping of virtual addresses to physical addresses to support flexible memory management from inside the VM to the Host. Works [72] can build a global virtual address space shared by all processes and range partitioned across memory blades to minimize the number of address translation entries.
- Memory region isolation: Typically, memory exists as multiple individual memory sticks that can be divided into different memory regions based on their distance from the CPU socket, i.e., NUMA architecture [69]. Memory accesses in the same socket have the same performance, but memory accesses across sockets will have lower bandwidth and longer latency.
- *Lightweight network channel allocation:* On general TCP/IP and RDMA network cards, the network channels are often controlled by the CPU cores since they are bonded with CPU threads. On RDMA, the Single-Root I/O Virtualization (SR-IOV) [144] is a standard mechanism of PCIe device virtualization technology. Each virtual PCIe device Z(such as RDMA) has its own PCIe configuration space and provides services for upper-layer software just like physical PCIe devices.

• *Easy-to-use interfaces abstraction*: people have designed diverse far memory access interfaces to easy use RDMA-based transfer[15, 41, 62, 129]. A proper implicit data structure is used behind the data transfer on memory spaces across network.

*5.2.2 Design in virtual machine (VM).* In VMs with host OS and without host OS, there are significant differences on far memory allocation ways. Each VM will be assigned CPU and memory in advance, while a hypervisor or monitor mainly manages storage I/O and network management. In scenarios with host OS, such as VMware [14], KVM [6], etc., the local memory can be allocated to each VM by the host, and the VMs can access different remote memories, both from other VMs and by exchanging data with the host memory. In scenarios without host OS, e.g., Xen, etc., the storage I/O, and network of each VM are managed by a particular VM (also called VM 0). In these scenarios, QEMU [12] is the core component for I/O and network management.

The disaggregated memory devices described in this paper include memory-like [3, 78], storage-like [151, 152, 154] and network-involved [18, 49, 72, 134] memory nodes. Thus, the design of a VM-level far memory system will consider the memory interactions between VMs, memory reclamation for each VM, and I/O and storage management of each VM. The host can allocate memory already reclaimed from VMs as well as additional memory that has not yet been allocated. In a network-based detached environment, the host can also allocate far memory for VMs on other machines.

Optimizing the memory access hierarchy in a virtual machine environment can improve the efficiency of memory resource allocation. Existing works concentrate on host swapping and ballooning for memory consolidation and over-commitment. VSwapper [19] addresses the challenge of poor performance caused by various types of superfluous swap operations, ruined swap files sequentially, and uninformed prefetch decisions upon page faults. It provides a Swap Mapper to monitor the disk I/O mapping with memory and a false read preventer to eliminate false reads. HybridSwap [162] addresses the problem that guest OS cannot utilize free pages in the host directly. They propose a distributed scalable framework to organize surplus memory in all hosts into virtual pools for swapping. XmemPod [29] can dynamically expand the memory capacity hierarchically memory by expanding its memory demand over virtualized host memory first and remote memory next before resorting to an external disk. Pond [78] can create machine learning models that can accurately predict local and pool memory allocation rates on VMs with same-NUMA-node memory performance.

*5.2.3 Design in containers.* In the cloud, far memory management in containers is attracting more attention, such as the main services Microservices and Serverless functions. There are optimization methods designed for containerized clouds. Some works optimize RDMA-0based far memory access in container environments. Freeflow [63] virtualizes RDMA environment for containerized clouds. It supports multi-tenancy isolation, container portability, and control and data plane controllability. Mitosis [150] proposes a fast RDMA remote fork that can quickly launch containers on multiple remote machines. Some works study the disadvantages and overheads when deploying *serverless* functions. Medes [118] optimizes the system performance by identifying memory redundancy and designing a simple sandbox management policy that supports smooth navigation in the trade-off space for operators. Unlike prior systems that allocate memory at job granularity, Jiffy [61] can achieve task-level resource scheduling. Jiffy efficiently multiplexes memory capacity across concurrently running jobs, reducing the overheads of reads and writes to slower persistent storage. Skadi [51] is a distributed runtime to mitigate the failure of deploying resource disaggregation in a harmony system. Beldi [157] proposes a stateful serverless functions to guarantee exactly once semantics.

*5.2.4 Design in Java virtual machine (JVM).* Some work focuses on invocating far-memory resources in Java virtualized environments, designing strategies for their synergy with Java garbage collection (GC) mechanisms. For example, MemLiner [133] proposes a runtime technique that "lines up" memory accesses from the application



Fig. 8. The centralized and decentralized DM management method.

and the GC, thereby reducing the local-memory working set and improving remote-memory prefetching. Semeru [132] designs a distributed JVM that provides a unified abstraction of virtual memory across CPU and memory servers and a distributed GC that offloads object tracing to memory servers. Some work has been devoted to building application non-transparent far-memory frameworks based on custom APIs to support far-memory calls in virtualized environments. AIFM [117] proposes application-integrated far memory based on remoteable and hybrid near/far memory data structures, which makes far memory available to applications through a simple API and without read-and-write amplification.

## 5.3 Task Scheduling Design

The above runtime systems can support various tasks in the virtualized environment, where tasks share the computing, memory, network, and storage resources of servers in the datacenter. Many works design and optimize task scheduling to better utilize disaggregated resources with high flexibility and high quality of service (QoS).

*5.3.1 DM scheduling architecture.* There are two types of management methods on disaggregated memory in the clusters, which we termed centralized management architecture and decentralized management architecture, as shown in Figure 8. The centralized DM management is easy for resource management and low-cost memory updating, while the decentralized DM management performs better for resource sharing and cost efficiency.

**Centralized DM scheduling architecture:** The centralized DM scheduling [46, 104, 153] maintains a central management entity, such as a large memory pool with multiple memory nodes, to provide far memory resources to each computing node. Actually, this concept was not proposed first; earlier in 1989, Kai Li et al. [80] stated that, in the shared memory system, one can use a shared virtual memory as the memory management monitor. The monitor consists of a data structure [16] and some procedures that provide mutually exclusive access to the data structure. Centralized DM has the advantage of data coherence management, which the memory pool monitor can strongly protect. However, it faces the problem of scaling up and out. One can find it hard to build a larger memory pool of more than 1000 memory nodes [104].

**Decentralized DM scheduling architecture:** In decentralized DM scheduling [29, 72, 108], the memory resource is distributed physically and shared virtually. Each server node is equipped with far memory resources and can provide additional memory resources for its own computing units and other computing nodes. As shown in Figure 8-(3), each server can have SRAM-based caches, DDR DRAMs, RAM-based non-volatile memories, flash-based SSDs, and disk to hierarchically store the data for heterogeneous computing units. This pattern can not only solve the scalability problem of centralized DM but also alleviate the performance bottleneck. Building a decentralized data communication scheme can decouple memory control and reduce data synchronization traffic. For example, COARSE [147] is a disaggregated memory extension that improves communication performance of GPUs. It builds on modern cache-coherent interconnect (CCI) protocols and MPI-like collective communication for synchronization, allowing low-latency and parallel access to training data and model parameters shared

among worker GPUs. Memtrade [94] employs a central coordinator that manages the disaggregated memory market and matches memory producers and memory consumers based on their supply and demand.

In cluster-level DM architectures, homogeneous and heterogeneous architectures have their advantages and disadvantages. Both can be used for resource scheduling and management through centralized and decentralized methods. Just like the shared-stade scheduling architecture [26, 119, 146], homogeneous and heterogeneous architectures may reasonably coexist and work in complementary ways in future disaggregated clusters.

*5.3.2 Resource management strategies.* Many works optimize the system from the view of hardware resources. We list several commonly used optimization methods for efficient resource management.

Adaptive data distribution: Optimizing resource allocation and access patterns in disaggregated memory systems requires an awareness of applications' or data types' latency sensitivity. Identifying various sensitivities helps the system optimize memory access and bandwidth allocation. HyFarM [137], CFM [18], and TMO [152] both focus on profiling and measuring the impact of far memory on application performance. CFM [18] develops a degradation profile for each application by measuring runtime across different local memory ratios. TMO [152], on the other hand, leverages the Pressure Stall Information (PSI) mechanism to dynamically adjust offloading decisions. HyFarM [137] senses sensitivity to collocate FM-sensitive and FM-tolerant tasks. Comparatively, while fastswap [18] and TMO [152] dive into performance degradation of individual tasks, HyFarM [137] applies the sensitivity to multi-task orchestration in a hybrid FM environment.

**I/O bandwidth expansion:** Existing works try to add new direct memory/storage access paths to improve the memory efficiency and I/O bandwidth utilization. GPU direct storage (GDS) supports direct storage that skips copies in CPU memory. However, the parallelization of GDS's I/O requests is limited by the CPU. BAM framework [111] adds NVMe-based I/O read and write logic to the cuda core, enabling the GPU to directly initiate more threads. The work GIDS [106] further adds asynchronous I/O processing modes based on BAM, thus further reducing I/O overheads. Direct memory access is also a solution to reduce data copy and duplication. For example, Zero-Infinity [112] can offload all of the partitioned model states to CPU or NVMe memory or keep them on the GPU based on memory requirements, thus handling huge parameters for training on current GPU clusters.

**Data throughput improvement:** The nature of disaggregated memory systems supports parallel data transfer and parallel execution of applications by allowing concurrent access to a shared memory pool. Fargraph+ [136] proposes a parallelism control strategy on graph processing on disaggregated memory architecture. It controls the multi-threading design of graph applications on the basis of saving RDMA-related resources (CPU core, memory space, RDMA queues, etc. Canvas [134] redesigns the swap system with fully isolated swap paths for remote-memory applications, thus improving the swap parallelism of data swap paths.

**Resource preemption reduction:** To reduce resource preemption in disaggregated memory systems, effective resource isolation among different tenants or applications is essential. Canvas [134] isolates swap paths for applications, each processing its own swap partition, cache, and RDMA bandwidth. Remote Region [15] supports exporting memory as files, supplementing RDMA functions lacking, such as namespace and access control. To balance the diverse needs of applications, ensuring high-priority tasks with more local memory resources is essential. HyFarM [137] works hard on balancing the resource pressure between servers by analyzing the far memory sensitiveness of each application and assigning more local memory to those far-memory-sensitive tasks.

**Reliability guarantee:** Existing work trends to improve the system's reliability include failure recovery, fault tolerance, resource isolation, etc. 1). Some work has been devoted to failure recovery designs to improve the reliability of far memory systems. Hydra [74] presents a failure resilience mechanism for remote memory, adopting an in-memory erasure coding scheme that achieves single-digit µs tail memory access latency by analyzing load balancing and availability trade-offs for distributed erasure codes. Carbink [?] proposes a far memory system that uses erasure-coding, remote memory compaction, one-sided RMAs, and offloadable parity calculations to achieve fast, storage-efficient fault tolerance. 2). Some works adopt coherence protocols for better

fault tolerance. Apta [107] designs a fault-tolerant object-granular CXL disaggregated memory for accelerating Function-as-a-Service (FaaS). Apta introduces a novel fault-tolerant coherence protocol for keeping the cached objects consistent without compromising availability when facing computing node failures. 3). Plenty of works implement isolated far memory access to avoid failure and interference. For example, the vVEEs [65] construct trusted and virtual Trusted Execution Environments (TEEs) by adding unified interfaces and standard security primitives that can run on heterogeneous hardware components. It also builds secure domain isolation across disaggregated components at the user level without losing flexibility and elasticity. Xmempod [29], Memory trades [94], and Remote reagions [15] use VMs and RDMA memory regions to isolate data of each program to avoid performance interference.

*5.3.3 Far-memory-aware scheduling.* There are some works designing far memory aware scheduling methods for higher memory utilization and efficiency.

**Memory resource scheduling:** The traditional scheduler needs to ensure strict Quality of Service (QoS) of tasks. Some tasks use elastic scaling to flexibly schedule resources [116], improve resource management efficiency through task understanding and resource occupancy prediction [36]. Some works try to dynamically capture accurate resource requirements from execution runtime for fine-grained scheduling [57], and optimizing jobs with complex dependency structures and heterogeneous resource requirements [48]. Others are geared towards goals such as scheduling fairness, energy efficiency, and cost control.

**Dynamic data placement:** Most of the recent scheduling approaches focusing on dynamic data placement. Memtrade [94] harvests producer memory using an application-aware control loop to form a distributed transient remote memory pool with minimal performance impact. For example, CFM [18] computes the minimum amount of data that can be retained in local memory for each task's SLO requirements, thus freeing up more memory space than is available to accommodate more tasks. Software-defined remote memory [68] compresses the swapped-out pages by the zswap kernel [86] and selectively evict low-priority jobs by killing them and rescheduling them on other machines when memory is reaching its limit. xMemPod [29] employs a hierarchical remote memory system that extends the VM's memory to the VM's mainframe memory, to the remote memory, and finally to the hard disk. Allot [33] proposes memory resource abstraction and data placement policies for distributed hybrid memory pools (DHMP) with RDMA capabilities to efficiently manage hybrid NVM and DRAM memory.

**Netural-network-based scheduling:** Some works design hybrid memory schedulers for natural networks. Sentinel [113] states that HM imposes challenges on tensor migration and allocation for high-performance DNN training. Sentinel uses dynamic profiling and coordinates OS and runtime-level profiling to bridge the semantic gap between OS and applications, which enables tensor-level profiling and co-allocating tensors with similar lifetime and memory access frequency into the same pages. Kleio [40] proposes a page scheduler with machine intelligence for applications that execute across hybrid memory components. The scheduler combines existing, lightweight, history-based data tiering methods for hybrid memory, with novel intelligent placement decisions based on deep neural networks.

**Application-aware scheduling:** Some works take the application sensitivity of far memory usage into account. By directly measuring an application's sensitivity to memory access slowdown, TMO [152] can effectively offload memory across diverse workloads with minimal impact on application performance. CFM [18] designs an application-aware cluster scheduler that leverages far memory to improve job throughput. HyFarM [137] proposes a task management strategy for hybrid far memory clusters that cooperatively co-locate tasks according to the far memory sensitivity analysis. It features an intra-/inter-node joint memory adaptation approach for balancing the memory usage at both the task and server levels.

**NUMA-aware scheduling:** Plenty of works concentrate on NUMA architecture to optimize application performance. Pond [78] relies on the CXL interconnect standard and takes the CXL memory as a NUMA node. It exposes the pool memory to a VM's guest OS as a zero-core virtual NUMA (zNUMA) node, i.e., a node

with memory but no cores, like Linux's CPU-less NUMA [69]. With the combination of hardware and systems techniques, it can achieve both same-NUMA-node memory performance and competitive cost for public cloud platforms. ThymesisFlow [108] designs the run-time attachment/detachment of byte-addressable disaggregated memory to a running Linux Kernel exploiting dynamically created NUMA nodes to host the remote memory. This work [66] utilizes Linux's NUMA memory page migration policy that automatically promotes and moves hot/cold memory pages between local and far memory.

## 6 APPLICATION-LEVEL OPTIMIZATION METHODS

Faced with different applications and different far memory access patterns, we summarize general-purpose execution optimization and domain-specific system design.





## 6.1 General-Purpose Execution Optimization

The first consideration for application performance improvement in far memory environment is the data management policy during execution. Particularly, the optimization strategies of application access to far memory mainly occur in three processes: 1) far-memory job distribution and local-memory data caching in compute nodes, 2) data transferring between local and far memory, i.e. data communication between compute nodes and memory nodes, and 3) data caching and storing procedure through hierarchical memory devices on memory nodes. To simplify this, we summarize that **FMA-based data processing = data distribution + data communication + data preservation**.

*6.1.1 Data distribution.* Far-memory-involved computation is a new type of computing pattern that is complementary to the existing computation mode. We discuss how far memory acts in program running procedures and the relationship between far memory and existing computing patterns.

**In-memory data offloading:** To save local memory space, large-scale applications, such as graph processing [124, 164], video processing [43], AI model training [123, 139] applications, often load many tax-like few-used data into local memory. Based on fine-grained program analysis, far memory can help to offload the in-memory data into far memory, including offloading cold and read-only data [117, 135], evicting one-off data [18, 152], compressing cold data [34, 86], et al.

**out-of-core data distribution:** Existing works often [91, 115, 161, 165, 168] use classic storages (i.e. HDD) to store large dataset. Far memory systems can accelerate the I/O performance by replacing the HDD-based access

paths and programming interfaces with RDMA-based[18, 135] or SSD-based[88, 152] far memory access paths, leveraging much larger data bandwidth and shorter access latency.

**Data partition in distributed computation:** At the cluster level, distributed refers to dividing the task process and task data into separate servers and running them in parallel[32]. With far memory involved, one can only partition task data with few job distributions, which is more friend to distributed computing[5, 15, 117, 129]. One can also provide flexible and additional memory spaces for each sub-task to enhance overall efficiency[136].

**Task working set partition:** In memory processing applications load all the working set [38, 44] into local main memory and then execute them in memory and caches. The working set mainly includes two parts: the required software environment and source data. In the cloud, most of the serverless functions and micro-service applications use this model. For latency-critical applications, the in-memory processing model can have high performance but occupy much more memory space, always several times the least local memory size since there is part of the memory tax [152]. Fastswap [18] employs Linux control groups (cgroups) [4] for working set partition. It sets limitations on physical memory allocation to a process to control its local memory ratio. FAM-Graph [155] proposes a partition scheme that separates vertex data and retrieves edge data chunks only when needed. Another method is to compress cold data to save more time and capacity [86]. Google's software-defined far memory [68] leverages zswap mechanism to compress cold pages proactively, then move them to slower storage.

6.1.2 Data communication. The key operations of data communication between local memory (computing nodes) and far memory (memory nodes) are data offloading (to write and store data on the far memory) and data fetching (to read data from far memory and get them into local memory for computing). Some works design specific memory access approaches based on read and write behavior and utilize data transfer channels to improve memory bandwidth utilization [30, 141]. (1). Data/Memory offloading is a common term used to indicate that part of the data in a program can be offloaded to far memory and will be retrieved from the same memory address. Due to the limited local memory, the unused data in local memory needs to be reclaimed to free up memory space (i.e., memory reclamation or garbage collection). One can swap data from the far memory into the local memory. Designing an efficient fetching mechanism is necessary due to the limited memory or I/O bandwidth and high demand for access speed. One can leverage a memory management method along with a network communication protocol to fetch the remote data.

**User-friendly program interfaces.** Program interfaces significantly simplify access to disaggregated memory systems, utilizing RDMA and other technologies. LITE [129] and FaRM [41] both abstract RDMA to facilitate application-level interactions. LITE [129] virtualizes native RDMA into a local indirection tier to achieve both safety and performance. FaRM [41] uses RDMA to build a memory-distributed platform in which address space is shared. Remote Region [15] and AIFM [117] both offer concise APIs for memory manipulation, making it easier to manage remote and local data. Besides, FreeFlow [62], AR-gRPC [24], and Mitosis [150] support high-performance applications by providing far-memory access frameworks for specific scenarios.

Adaptive data granularity: Data granularity refers to the level or size of data division. The way data is stored and processed across various hardware and software platforms significantly influences efficiency. Some works try to design smaller or more flexible data access granularity for applications to improve application performance. Fargraph [135] transparently segment data into data chunks in RDMA transfer to facilitate RDMA read/write. Increasing data granularity to enable accurate operations is also a trend. AIFM [117] utilizes far memory at object granularity and manages local offloadable memory at log granularity. Kona [27] uses cache coherence instead of virtual memory for tracking applications' memory access behavior transparently at cache-line granularity. Some work addresses the problem of high dirty data amplification that comes from using page granularity for tracking changes to the cached data (4KB or higher). Remote regions [15] consider the trade-off between granularity and flexibility in its region map, which tracks the location of regions.

**Parallel data swap:** The coexistence of multiple applications can lead to significant performance degradation due to interference among swap operations. This challenge underscores the importance of developing robust parallel data swap mechanisms that can accommodate the simultaneous demands of diverse applications while minimizing the adverse impacts of such interference. Recent works aim to optimize swap efficiency and minimize slowdowns due to cross-application interactions. Canvas [134] introduces a fully isolated swap environment for each application, encompassing resources such as swap partitions, caches, prefetchers, and RDMA bandwidth. They design dynamic swap entry allocation, semantics-aware prefetching, and strategic RDMA scheduling, collectively enhancing swap efficiency and application performance. VSwapper [19] complements this approach by focusing on the granularity of swap operations. Through its Swap Mapper, VSwapper addresses issues like silent and stale writes by establishing a clear correspondence between disk blocks and guest memory pages. This is achieved via *mmap* system calls, which streamline the swap process and reduce operational overhead, ensuring that the host kernel efficiently manages and names pages.

*6.1.3 Data preservation.* In multi-level memory systems that involve both fast and slow storage spaces, there are two data preservation methods: *inclusive* cache hierarchy and *exclusive* cache hierarchy, as shown in Table 6. Oftentimes, the fast memory entities have a small memory capacity based on volatile memory cells. On the contrary, slow memory entities often have a large capacity on non-volatile memory/storage blocks.

In an *inclusive* cache hierarchy, the data stored in the fast memory (like RAM or cache) is a subset of the data stored in the slow memory (such as HDDs or SSDs)[55]. Traditionally, the performance gap between devices in the memory hierarchy is relatively large. Thus, the off-the-shelf often uses inclusive cache hierarchy to reduce end-to-end memory access latency[2?]. However, as the analysis before (in Sec. 3.3), the performance gap has been narrowing with the development of different memory devices. The use of data swap is adequate for better utilization of memory and storage space. In an *exclusive* cache hierarchy, the data stored in the fast and slow memory spaces is distinct[55]. This arrangement entails that the data stored in the fast memory is not duplicated in the slow memory space. This method can ensure the maximum utilization of these two types of memory resources. This setup can significantly improve system performance, as it prevents unnecessary duplication and better uses faster memory for critical tasks.

**Cache optimization:** Recent works have been focusing on designing data caches between local and far memory. For example, Hybrid<sup>2</sup> [130] designs a memory system architecture that combines a small high-bandwidth 3D-stacked DRAM and a large lower-bandwidth off-chip DRAM. Hydra [73] introduces a configurable resilience mechanism that applies online erasure coding to individual remote memory pages while maintaining high availability. Hydra can be integrated into the existing far memory frameworks such as Infiniswap [49], Leap [17] and Remote Regions [15], to have better performance.

**Selective data prefeching:** Prefetching data into local memory before the actual access instruction is often adopted. Efficient prefetching optimizations can happen at the application, system, and hardware levels. AIFM [117] leverages the semantics of data structures to support customized prefetching and caching policies. Fargraph [135] prioritizes data segments based on far memory ratio to minimize transmission overhead. The application's host server uses local memory as a data cache. Fastswap design a prefetch method for higher data access performances [18]. Leap [17] implements its prefetching algorithm with a data path that segregates application data transfer to far memory. In contrast, Canvas [134] schedules packets of each application in an RDMA manner. Furthermore, some works save space and reduce data transfer overhead by compressing cold data [68]. HoPP [79] designs a hardware-based memory controller to prefetch sufficient hot pages to OS on a separate data path alongside the normal remote data path. Assise [20] maximizes data locality for all file IO by carrying out IO on process-local, socket-local, and client-local persistent memory. Some works utilize better indexing strategies to accelerate the data fetching. Sherman [142] leverages distributed B+ tree in a DM system to accelerate indexing.

### 6.2 Domain-Specific Design

*6.2.1 Elastic computing on far memory.* Elastic applications typically require on-demand dynamic networking. For instance, computing nodes in a disaggregated storage system access the data stored at the storage nodes across the network. A elastic computing system needs to scale automatically according to application demands. For example, KRCore [149] designs a microsecond-scale control plane on commodity RDMA hardware for elastic computing. The key ideas include virtualizing pre-initialized kernel-space RDMA connections instead of creating one from scratch and retrofitting advanced RDMA dynamic connected transport with static transport for low overhead connection and high networking speed.

*6.2.2 Graph processing on far memory.* Fargraph [135] analyzes the impact of graph processing workload on disaggregated architecture. Specifically, it reduces the overall data movement through a well-crafted, graph-aware data segment offloading mechanism. In addition, it uses optimal data segment splitting and asynchronous data buffering to achieve graph iteration-friendly far memory access. FAMgraph [155] leverages application-level properties, such as read-only edge data, to efficiently tier data between local and remote memory and prefetch remote data for local computation.

6.2.3 DNN-supported far memory. As the model and dataset for production-scale recommendation systems scale up, the size of the embeddings is approaching the memory capacity limit. Recent studies offload the embedding lookups into SSDs, which target the embedding-dominated recommendation models. RM-SSD [127] offloads the entire recommendation system into SSD with in-storage computing capability. The proposed SSD-side FPGA solution leverages a low-end FPGA to speed up both the embedding-dominated and MLP-dominated models with high resource efficiency. ReCXL[87] is a CXL memory disaggregation system that utilizes near-memory processing for scalable recommendation model training, utilizing data transferring bandwidth optimizations.

*6.2.4 Large-scale model training system.* Due to the huge size of parameters, both training and inference of large language model (LLM) need to be deployed in a distributed manner to improve parallelism. The cost of their hardware and power consumption is exceptionally high. Recent work has investigated high-performance scaling of memory pools and efficient management of hierarchical memory architectures. For example, HET [96] is a new system framework with a fine-grained cache consistency design that significantly improves the scalability of massive embedding model training. It embraces skewed popularity distributions of embeddings and leverages it to address the communication bottleneck with an embedding cache.Angel-PTM [98] proposes a productive deep learning system that can train extremely large-scale models pre-training and fine-tuning with hierarchical memory. The key designs of Angel-PTM are fine-grained memory management via page abstraction and a unified scheduling method that coordinates computations, data movements, and communications. Furthermore, Angel-PTM has also designed a lock-free updating mechanism to address SSD I/O bottlenecks.

## 7 SUMMARY OF CROSS-LAYER DM WORKS

In this section, we give an overall summary about representative disaggregated memory works with cross-layer designs, i.e. architecture level, system level, and application level. We summarize the far memory types (architecture level), runtime design level (system level), design considerations including data granularity, programming difficulty, and design metric (application level).

**Far memory type:** The second column in Table 7 shows far memory types used in recent works. From the perspective of hardware, most works choose DRAM as local memory. There are plenty of works supporting NIC-based far memory. A subset of studies explore alternative NIC designs except RDMA NIC, including PCIe Switch, OpenCAPI, NVLink, DIMM, Optical, FPGA, CCI, and CXL Switch. Particularly serving for GPU-like accelerators, GPU far memory is becoming a new trend. Related works have supported the heterogeneous memory types of NUMA, HMC, 3D-stacked DRAM, OCM, and HBM, which can also viewed as far memory. With the

Table 7. A summary of function on related works.  $\checkmark$  means this aspect is carefully designed in the paper,  $\circ$  means this aspect is considered in the paper, - is not considered. "NIC" means Network Interface Card. "CXL" is Computing Express Link. "TCO" means Total Cost of Ownership. "OCM" is Optically Connected Memory. "VM" represents Virtual Machine and "JVM" is Java Virtual Machine. "DL" is Deep Learning and "ML" is Machine Learning. "A/F" represents that the system considers anonymous pages and File-backed pages. "Fa" represents Failure and "Is" represents Isolation.

	Far Memory Type (Local is DRAM)					Runtime Design Level					Granularity				Programming				Design Metric								
Related Works	PM (NVM)	CXL	Network Fabric (NIC)	Memory Controller	SSD	disk	GPU	Hybrid Memory	Hardware-oriented	Architecture-oriented	Operating System (OS)	VM-oriented	Service-oriented	Cache Line (64B)	Page Size (4KB)	Large Page (2MB)	Object (varied)	Transparency	Cache Coherence	Customized API usage	Parameter Conf.	Ownership Cost (TCO)	Execution Latency	Data/task Throughput	Power Consumption	Isolation & Failure	Resource Utilization
Memory Blade'09 [85]	-	-	-	FGRA	-	-	-	-	0	$\checkmark$	-	-	-	-	$\checkmark$	-	-	0	$\checkmark$	-	-	$\checkmark$	$\checkmark$	- 1		-	-
FaRM'14 [41]	-	-	RDMA	-	0	-	-	-	-	-	-	$\checkmark$	Read/Write	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	- 1	-	-
Holistic'15 [66]	-	-	-	-	-	-	-	NUMA	-	-	$\checkmark$	VM	-	-	$\checkmark$	-	-	$\checkmark$	-	-	-	-	-	-	-	-	$\checkmark$
HybridSwap'15 [162]	-	-	-	-	$\checkmark$	-	-	-	-	-	-	VM	-	-	1	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-	-	-
LITE'17 [129]	-	-	RDMA	-	-	-	-	-	-	-	$\checkmark$	$\checkmark$	Read/Write	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	-	-	$\checkmark$
Infiniswap'17 [49]	-	-	RDMA	$\checkmark$	-	-	-	-	-	-	$\checkmark$	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-	Fa	-
LegoOS'18 [120]	$\checkmark$	-	-	-	0	0	0	-	-	0	$\checkmark$	0	-	-	$\checkmark$	-	-	$\checkmark$	-	-	-	$\checkmark$	$\checkmark$	-		Fa	-
Zombieland'18 [99]	-	-	RDMA	ACPI	-	-	-	-	-	-	0	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	$\checkmark$	-	-
Remote Region'18 [15]	-	-	RDMA	-	0	-	-	-	-	-	-	$\checkmark$	-	-	<ul> <li>✓</li> </ul>	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	-	$\checkmark$	$\checkmark$	-	Is	-
AR-gRPC'18 [24]	-	-	RDMA	-	-	-	-	-	-	-	-	~	Tensorflow	-	-	-	<b>√</b>	-	-	$\checkmark$	<b>√</b>	-	$\checkmark$	$\checkmark$	-	Is	-
MC-DLA'18 [67]	-	-	PCle Switch	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	DL Training	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	$\checkmark$	-	-
String Figure'19 [104]	-	-	-	RTL	-	-	-	HMC	$\checkmark$	$\checkmark$	-	-		$\checkmark$	-	-	-	-	$\checkmark$	-	-	-	-	$\checkmark$	$\checkmark$	-	-
Zswap'19 [68, 86]	-	-	-	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	-	$\checkmark$	$\checkmark$	-	-	-	-
FreeFlow'19 [62]	-	-	vRDMA	-	-	-	-	-	-	-	-	$\checkmark$	Container	-	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	-	$\checkmark$	$\checkmark$	-	Is	-
Leap'19 [17]	-	-	RDMA	$\checkmark$	$\checkmark$	-	-	-	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	-	-	-
Clover'20 [128]	$\checkmark$	-	RDMA	-	-	-	-	-	-	-	-	$\checkmark$	Kv Store	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	$\checkmark$	-	-
ThymesisFlow'20 [108]	-	0	OpenCAPI	FPGA	-	-	0	-	$\checkmark$	0	$\checkmark$	-	-	0	$\checkmark$	-	-	$\checkmark$	0	-	-	-	-	$\checkmark$	-	-	$\checkmark$
Fastswap, CFM'20 [18]	-	-	RDMA	$\checkmark$	$\checkmark$	-	-	-	-	-	$\checkmark$	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	-	-	-
Semeru'20 [132]	-	-	RDMA	✓	0	-	-	-	-	-	-	JVM	-	-	$\checkmark$	-	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	-	-	$\checkmark$
AIFM'20 [117]	-	-	RDMA	-	-	-	-	-	-	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	-	-	-
Swap Advisor'20[52]	-	-	NVLink	-	$\checkmark$	-	$\checkmark$	-	-	-	-	$\checkmark$	DL	-	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	-	$\checkmark$	$\checkmark$	-	-	-
Buddy Compression'20 [34]	-	-	NVLink	-	$\checkmark$	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	-	$\checkmark$	-	-	-	$\checkmark$	$\checkmark$	-	-	-
Griffin'20 [21]	-	-	PCIe	IOMMU	-	-	V	NUMA	$\checkmark$	$\checkmark$	-	-	-	-	✓	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-		-
Hybrid <sup>2</sup> '20 [130]	-	-	DIMM	-	-	-	-	3D-stacked DRAM	$\checkmark$	$\checkmark$	-	-	-	-	$\checkmark$	-	$\checkmark$	-	-	$\checkmark$	~	$\checkmark$	$\checkmark$	$\checkmark$	-	-	-
MIND'21 [72]	-	-	Switch	ASIC	-	-	-	-	0	$\checkmark$	-	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-	Fa	-
XMemPod'21 [28, 29, 88]	-	-	RDMA	$\checkmark$	$\checkmark$	-	-	-	-	-	$\checkmark$	VM	-	-	$\checkmark$	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	-	Is	-
Kona'21 [27]	-	<b>√</b>	RDMA	FPGA	-	-	-	-	-	$\checkmark$	-	$\checkmark$	-	$\checkmark$	-	-	-	$\checkmark$	$\checkmark$	$\checkmark$	-	-	$\checkmark$	-	-	Fa	-
ZeRO-Infinity'21 [112]	-	-	NVLink	-	V	-	<b>↓</b>	-	-	-	-	$\checkmark$	DL	-	<ul> <li>✓</li> </ul>	-	-	-	$\checkmark$	-	-	-	$\checkmark$	$\checkmark$	-		-
	-	<b>↓</b>	-	-	$\checkmark$	-	-	-	-	$\checkmark$	$\checkmark$	-	-	-										_			
HET'21 [96]	-	-	-	-	-	$\checkmark$	$\checkmark$	-	-	-	-	$\checkmark$	ML Training	-	$\checkmark$	-	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	-	-	-	-
Memtrade'21 [94]	-	-	-	-	$\checkmark$	$\checkmark$	-	-	-	-	-	VM	-	-	$\checkmark$	-	-	-	-	$\checkmark$	-	$\checkmark$	$\checkmark$	-	-	-	$\checkmark$
TMO'22 [152]	-	-	-	-	$\checkmark$	$\checkmark$	-	-	-	-	$\checkmark$	-	-	-	A/F	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	-	-	-
Sherman'22 [142]	0	-	RDMA	-	-	-	-	-	-	-	-	$\checkmark$	Read/Write	-	-	-	$\checkmark$	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	-	-	-
Teleport'22[163]	-	-	RDMA	-	0	0	0	-	-	0	$\checkmark$	0	-	-	$\checkmark$	-	-	-	$\checkmark$	$\checkmark$	-	$\checkmark$	$\checkmark$	-	-	-	-
Jiffy'22 [61]	-	-	RDMA	~	$\checkmark$	-	-	-	-	-	-	$\checkmark$	Serverless	-	-	-	$\checkmark$	-	-	~	$\checkmark$	$\checkmark$	$\checkmark$	-	-	-	$\checkmark$
Clio'22 [50]	-	-	RDMA	FPGA	-	-	-	-	0	~	-	-	-	-	0	-	<ul> <li>✓</li> </ul>	-	-	~	-	-	$\checkmark$	$\checkmark$	$\checkmark$	-	-
OCM'22 [46]	-	-	Optical	-	-	-	-	OCM	$\checkmark$	$\checkmark$	-	-	-	$\checkmark$	-	-	-	~		-	-	-	$\checkmark$		$\checkmark$	<u> </u>	-
BEACON'22 [53]	-	$\checkmark$	-	-	-	-	-	0	0	$\checkmark$	-	-	Analysis	$\checkmark$	-	-	-	-	-	$\checkmark$	~	-	$\checkmark$	-	•	-	-
MemLiner'22 [133]	-	-	RDMA	√	-	-	-	-	-	-	-	JVM	-	-	-	-	$\checkmark$	-	0	$\checkmark$	-	$\checkmark$	$\checkmark$	-		-	-
Mitosis'22[150]	-	-	RDMA	-	-	-	-	-	-	-	$\checkmark$	-	Serverless	-	-	-	$\checkmark$	-	-	$\checkmark$	$\checkmark$	-	$\checkmark$	$\checkmark$	<u> </u>	Is	-
Medes'22 [118]	-	-	RDMA	-	-	-	-	-	-	-	-	$\checkmark$	Serverless	-	<ul> <li>✓</li> </ul>	-	0	-	-	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	-	Is	-
Fargraph'22 [135, 136]	-	-	RDMA		<ul> <li>✓</li> </ul>	-	-		-	-	-	<ul> <li>✓</li> </ul>	Graph	-	-	-	V.	-		✓	-	√	$\checkmark$	$\checkmark$	-	<u> </u>	-
FAMgraph 22[155]	-	-	RDMA	~	<b>√</b>	-	-		-	-	-	~	Graph	-	-	-	V.	-	-	<ul> <li>✓</li> </ul>	-	~	$\checkmark$			-	-
KRCore 22 [149]	-	-	RDMA FDCA	-	-	-	-	-	-	-	~	-	Elastic	-	-	-	<b>↓</b>	-	-	√ 	-	-	V			-	-
RM-SSD 22[127]	-	-	FPGA	~	-	-	-	-	~	0	-	-	-	-	<b>_</b> √	-	-	-	-	~	-	-	$\sim$	$\vdash$		<u> </u>	-
COARSE'22 [147]	-	-	CCI	FPGA	-	-	$\checkmark$	-	-	-	-	-	Training	-	~	-	-	-	$\checkmark$	-	-	-	$\checkmark$	~	-	-	-
HyFarM'22 [137]	-	-	RDMA	-	1	1	-	-	-	-	-	↓ ✓	-	-	A/F	-	-	√ ,	-	-	Į√.	$\checkmark$	L	<u> </u>		- -	$\vdash$
Canvas 23 [134]	-	-	RDMA	-	-	-	-	-	-	-	<b>√</b>	<b>√</b>	-	-	A/F	-	-	~	<u> -</u>	-	<b>↓</b>	-	Ľ	<u>-</u>	⊢	ls	
BAM 23 [111]	-	-		-	√	-	<b>↓</b> ✓	-	-	-	-	<ul> <li>✓</li> <li>✓</li> </ul>	-		<b>↓</b> ✓	-	-	-	<b>↓</b>	-	-	-	Ľ,	۲÷	<b>⊢</b> –∣	<u> </u>	<u> -</u>
GIDS 23 [106]	-	-	-	-	V /	-	V /	-	-	-	-	V	-	-	V	-	-	-	V	-	-	-	$ \downarrow $			<u> </u>	-
Angel-PTM [98]	-	_	_	_	✓ ✓	-	× ./	- HBM	_	_	_		LLM	_	✓ √	4MB	_	_	× ./	_	_	-	× 		-	-	-
Dan d'on [70]	-	,	CXL		Ľ.	-	<u> </u>	NUM	/	-	-	VD4	Training	-	-		-	/			-	,	Ľ,		$\vdash$	-	<u> </u>
xDM'24 [138]	-	✓   ✓	Switch RDMA	-	-	-	-	NUMA	✓ -	-	-	VM	-	-		- 2MB	-	$\checkmark$	0	-	-	$\checkmark$	$\checkmark$	-	-	Is	
GMT'24 [31]	-	<u> </u>	-	-	ŀ.	ŀ		-	$\checkmark$	, V	-	-	-	$\checkmark$	, V	-	-	, V	0	-	Í	-	+ T	-		-	<u>† -</u>
D. 0072-1	1				†	† ·							AI	· .													<u> </u>
ReCXL'24 [87]	-	✓	-	PNM	-	-	✓	✓	$\checkmark$	$\checkmark$	-	-	training	√	✓	-	-	~	0	-	✓ _	- 1	~	-	$  \vee  $	-	-

advancement of storage devices such as NVM, PM, SSD, and disk, storage-based far memory has started attracting interest. Furthermore, there has been a gradual increase in research supporting CXL. The design of memory controllers for far memory pool management is increasing, such as FGRA, ACPI, RTL, FPGA, IOMMU, and ASIC.

**Runtime design level:** We list 5 design considerations in the third column of Table 7, including specific hardware-oriented system designs (Sec. 3.1 and 3.3), architecture designs (Sec. 4.4), OS-based adaptions (Sec. 5.1), VM-centric designs on virtualized environments (Sec. 5.2), and services-aware designs (Sec. 6.2). Most works implement optimization methods in virtualized environments, including virtual machines and Java virtual machines. Nearly a third of works design for a specific service, including optimized systems on data read/write access, tensorflow, deep learning training, container-based optimization, key-value store, serverless, genome analysis, graph processing, and elastic computing services.

**Data granularity:** We list the supported data granularity in the fourth columns of Table 7. There are three general types of fixed-size data granularity, the cacheline size in 64B, the page size in 4KB, and the large page size in 2MB. Most of the works support 4KB page size, with more papers tents to use a large page size as 4MB in Angel-PTM [98]. A few works enable 64-byte cache line architecture in data transfer: String Figure [104], Kona [27], OCM [46], and BEACON [53]. Many works use a flexible and varied data block size.

**Program programming difficulty:** As shown in the fifth columns of Table 7, when analyzing the programming difficulties of programs in the related works, we consider four parts: programming transparency, cache coherence, customized API usage, and parameter configurations. We observe that quite a few works are designed to be transparent and enable configuration parameters. Recent related works have started to design cache coherence protocols, especially on CXL-based memory. Most application-oriented works design customized API usage to utilize RDMA, or far memory calls for user-friendly programming.

**System design Metric:** We list typical design metrics in the sixth column of Table 7. In general, the design goals of the listed works aim at better functionality and higher performance. The detailed categories include the total cost of ownership (TCO), execution latency, data/task throughput, power consumption, isolation and failure, and resource utilization. The majority of systems are designed to reduce execution latency and improve data or task throughput. Some works mainly focus on reducing total cost of ownership (TCO) and power consumption by using efficient hardware or design power management software. Systems can also add secure-aware functions by realizing resource isolation (marked as Is) and handling failures (marked as Fa).

## 8 FUTURE WORK

## 8.1 Emerging technology of disaggregated memory dystem

The design of the DM system requires lower-power and high-performance hardware devices, as well as support for higher concurrency and larger-scale applications.

Hardware design for DM: Existing works try to design and develop memory hardware with the ability of neardata processing, direct connection, hot swap/plugging, data consistency, and security guarantee. CXL-supported memory still needs further design on cross-node and cross-rack memory infrastructure. This will increasingly rely on the design of higher-speed network transfers and more efficient network bandwidth allocation.

**DM for cloud services:** Adapting far memory to platform as a service (PaaS), software as a service SaaS, function as a service (FaaS), and model as a service (MaaS) is essential. This will also lead to many promising research problems and novel solutions on virtualization environment design for services. DM system design needs to reduce the overhead of far memory access in virtualized environments.

**Memory pool for xPUs:** Today, the computation center is gradually migrating from CPUs to GPUs, FPGA, and processing-in-memory (PIM) accelerators, also named as x processing units (xPUs). Existing work uses distributed shared memory to design systems, and there is still a blue ocean of collaboration for heterogeneous

<sup>,</sup> Vol. 1, No. 1, Article . Publication date: August 2024.

compute resources with memory pools. Providing GPUs with larger memory storage space and higher-speed data access will lead to explosive performance gains.

**DM for micro datacenter:** In the era of micro datacenter, i.e. edge datacenter, edge devices tend to have more restricted memory and storage space and thus require fine-grained data offload design. Micro datacenters can provide higher computing power and a more flexible allocation of computing resources than end devices. Deploying disaggregated architecture in micro data centers will improve overall resource efficiency.

**High memory bandwidth and utilization:** Multiple modality natural networks have become popular in today's applications. It is essential to build highly parallel and high-bandwidth GPU remote access pathways. Designing high-performance GPU remote memory access patterns can be a main trend in future work. Further, improving multiple data transfer channels is also a good solution with good coordination of complex memory hierarchies with high parallelism.

#### 8.2 Disaggregated Memory for Emerging Applications

The explosive growth of AI model size inevitably increases the computation and memory costs, requiring large amount of GPUs. Managing memory resources is especially important for the pre-training and fine-tuning retraining process of large models, which needs to ensure the training latency and effectiveness.

**Design for large foundation model:** In sparse LFM scenarios, the number of longest tokens determines the matrix's length and width, so many zeros are in the matrix. This will occupy a lot of valuable GPU memory resources. One solution is to utilize multi-level memory to offload cold data to slow storage devices to utilize GPU memory space more efficiently. Using a polarized de-redundancy approach, the flexible memory extension method for LFM can accommodate more important computational data.

**Design for multi-modal neural network:** Disaggregated memory provides significant performance improvements, higher memory utilization efficiency, and robust model training capabilities for multi-modal neural networks by separating memory resources from processing units. It allows one to handle large-scale datasets from different data sources, supporting more complex network structures while reducing bottlenecks caused by memory limitations. It can greatly enhance the scalability and flexibility of the model.

**Design for metaverse and cloud gaming:** Disaggregated memory technology significantly enhances performance and optimizes metaverse and cloud gaming resources. By enabling seamless, cross-server sharing of memory resources within data centers, this technology allows metaverse and cloud gaming platforms to handle large-scale user data and complex graphic rendering tasks with greater efficiency.

**Design for embodied artificial intelligence:** DM holds transformative potential for embodied artificial intelligence, enhancing systems like autonomous vehicles and interactive robots by enabling dynamic scalability, improved memory utilization, and reduced latency. The technology allows AI systems to dynamically adjust and allocate memory resources across multiple nodes, which is essential for efficiently handling large, unpredictable data streams and complex computations.

**Design for scientific computing:** Disaggregated memory significantly enhances the efficiency and flexibility of high-performance computing systems in scientific computing, such as genomics, by allowing dynamic allocation of memory resources independent of compute nodes. This adaptability improves data load management effectiveness, reduces memory bottlenecks, and optimizes the overall system performance and cost efficiency.

#### 9 CONCLUSION

In this paper, we summarize the concept and implementation of disaggregated memory. We have developed a detailed taxonomy of disaggregated memory systems with systematically analysis. We introduce the technique insights of disaggregated memory design and optimization at different system layers, including hardware, architecture, system, and application layers. Lastly, we delved into the potential future opportunities of disaggregated

memory, aiming to offer insightful guidance for forthcoming research and developments in this field. This comprehensive review serves as a foundation for understanding on how to design disaggregated memory for the next-generation elastic and efficient datacenters.

#### REFERENCES

- [1] [n.d.]. Intel Performance Counter Monitor A Better Way to Measure CPU Utilization. https://www.intel.com/content/www/us/en/ developer/articles/technical/performance-counter-monitor.html.
- [2] [n.d.]. Model-specific registers (MSRs). https://man7.org/linux/man-pages/man4/msr.4.html.
- [3] 2024. Compute Express Link. https://www.computeexpresslink.org/.
- [4] 2024. Control Group v2. https://www.kernel.org/doc/Documentation/cgroup-v2.txt.
- [5] 2024. High-performance remote procedure call framework. https://github.com/grpc.
- [6] 2024. KVM (for Kernel-based Virtual Machine). https://www.linux-kvm.org/page/Main\_Page.
- [7] 2024. MGLRU. https://lore.kernel.org/all/20230714194610.828210-1-hannes@cmpxchg.org/.
- [8] 2024. Nvidia DPU. https://www.nvidia.cn/networking/products/data-processing-unit.
- [9] 2024. OFED from OpenFabrics. https://network.nvidia.com/products/infiniband-drivers/linux/mlnx\_ofed/.
- [10] 2024. OpenCAPI Specification. https://opencapi.org/.
- [12] 2024. QEMU. https://github.com/qemu/qemu.
- [13] 2024. Sumsung Solid State Drives (SSD). https://www.samsung.com/us/computing/memory-storage/solid-state-drives.
- [14] 2024. VMware. https://www.vmware.com.
- [15] Marcos K Aguilera, Nadav Amit, et al. 2018. Remote regions: a simple abstraction for remote memory. In USENIX Annual Technical Conference (USENIX ATC).
- [16] Marcos K. Aguilera, Kimberly Keeton, et al. 2019. Designing Far Memory Data Structures: Think Outside the Box. In Workshop on Hot Topics in Operating Systems (HotOS).
- [17] Hasan Al Maruf and Mosharaf Chowdhury. 2020. Effectively prefetching remote memory with leap. In USENIX Conference on Usenix Annual Technical Conference (USENIX ATC).
- [18] Emmanuel Amaro, Branner Augmon, et al. 2020. Can far memory improve job throughput?. In European Conference on Computer Systems (Eurosys).
- [19] Nadav Amit, Dan Tsafrir, and Assaf Schuster. 2014. Vswapper: A memory swapper for virtualized environments. Sigplan Notices (2014).
- [20] Thomas E Anderson, Marco Canini, et al. 2020. Assise: Performance and Availability via Client-local NVM in a Distributed File System. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [21] Trinayan Baruah, Yifan Sun, Ali Tolga Dinçer, et al. 2020. Griffin: Hardware-software support for efficient page migration in multi-gpu systems. In *IEEE International Symposium on High Performance Computer Architecture (HPCA)*.
- [22] Soroush Bateni, Zhendong Wang, et al. 2020. Co-optimizing performance and memory footprint via integrated cpu/gpu memory management, an implementation on autonomous driving platform. In *Real-Time and Embedded Technology and Applications Symposium* (*RTAS*).
- [23] Maciej Bielski, Christian Pinto, et al. 2016. Survey on memory and devices disaggregation solutions for HPC systems. In IEEE Intl Conference on Computational Science and Engineering (CSE) and IEEE Intl Conference on Embedded and Ubiquitous Computing (EUC) and Intl Symposium on Distributed Computing and Applications for Business Engineering (DCABES).
- [24] Rajarshi Biswas, Xiaoyi Lu, and Dhabaleswar K Panda. 2018. Accelerating tensorflow with adaptive rdma-based grpc. In IEEE International Conference on High Performance Computing (HiPC).
- [25] Andreas Blenk, Arsany Basta, et al. 2015. Survey on network virtualization hypervisors for software defined networking. IEEE Communications Surveys & Tutorials (2015).
- [26] Eric Boutin, Jaliya Ekanayake, et al. 2014. Apollo: Scalable and coordinated scheduling for cloud-scale computing. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [27] Irina Calciu, M Talha Imran, et al. 2021. Rethinking software runtimes for disaggregated memory. In Architectural Support for Programming Languages and Operating Systems (ASPLOS).
- [28] Wenqi Cao et al. 2017. FastSwap. https://github.com/git-disl/FastSwap.
- [29] Wenqi Cao and Ling Liu. 2020. Hierarchical Orchestration of Disaggregated Memory. IEEE Transactions on Computers (TC) (2020).
- [30] Yuhai Cao, Chao Li, et al. 2018. DR DRAM: Accelerating memory-read-intensive applications. In International Conference on Computer Design (ICCD).
- [31] Chia-Hao Chang, Jihoon Han, Anand Sivasubramaniam, Vikram Sharma Mailthody, Zaid Qureshi, and Wen-Mei Hwu. 2024. GMT: GPU Orchestrated Memory Tiering for the Big Data Era. In Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3. 464–478.
- , Vol. 1, No. 1, Article . Publication date: August 2024.

- [32] Rong Chen, Jiaxin Shi, et al. 2019. PowerLyra: Differentiated Graph Computation and Partitioning on Skewed Graphs. ACM Transactions on Parallel Computing (TPC) (2019).
- [33] Tingting Chen, Haikun Liu, et al. 2021. Resource abstraction and data placement for distributed hybrid memory pool. Frontiers of Computer Science (FCS) (2021).
- [34] Esha Choukse, Michael B Sullivan, et al. 2020. Buddy compression: Enabling larger memory for deep learning and HPC workloads on gpus. In *International Symposium on Computer Architecture (ISCA)*.
- [35] CCIX Consortium. 2023. Cache Coherent Interconnect for Accelerators. https://genzconsortium.org.
- [36] Eli Cortez, Anand Bonde, et al. 2017. Resource central: Understanding and predicting workloads for improved resource management in large cloud platforms. In Symposium on Operating Systems Principles (SOSP).
- [37] Ala Darabseh, Mahmoud Al-Ayyoub, et al. 2015. Sdstorage: a software defined storage experimental framework. In 2015 IEEE International Conference on Cloud Engineering.
- [38] P. J. Denning. 1980. Working Sets Past and Present. IEEE Transaction of Software Engineering (TSE) (1980).
- [39] Data Dog. 2023. The state of Serverless. https://www.datadoghq.com/state-of-serverless/.
- [40] Thaleia Dimitra Doudali, Sergey Blagodurov, et al. 2019. Kleio: A hybrid memory page scheduler with machine intelligence. In International Symposium on High-Performance Parallel and Distributed Computing.
- [41] Aleksandar Dragojević, Dushyanth Narayanan, et al. 2014. FaRM: Fast remote memory. In USENIX Symposium on Networked Systems Design and Implementation (NSDI).
- [42] Mohammad Ewais and Paul Chow. 2023. Disaggregated Memory in the Datacenter: A Survey. IEEE Access (2023).
- [43] FFmpeg. 2013. A complete, cross-platform solution to record, convert and stream audio and video. http://ffmpeg.org.
- [44] MPI Forum. 2023. Working set. https://en.wikipedia.org/wiki/Working\_set.
- [45] Peter X Gao, Akshay Narayan, et al. 2016. Network requirements for resource disaggregation. In USENIX symposium on operating systems design and implementation (OSDI). 249–264.
- [46] Jorge Gonzalez, Mauricio G Palma, et al. 2022. Optically connected memory for disaggregated data centers. Journal of Parallel and Distributed Computing (JPDC) (2022).
- [47] Donghyun Gouk, Sangwon Lee, et al. 2022. Direct Access, High-Performance Memory Disaggregation with DirectCXL. In USENIX Annual Technical Conference (USENIX ATC).
- [48] Robert Grandl, Srikanth Kandula, et al. 2016. GRAPHENE: Packing and Dependency-Aware scheduling for Data-Parallel clusters. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [49] Juncheng Gu, Youngmoon Lee, et al. 2017. Efficient memory disaggregation with infiniswap. In USENIX Symposium on Networked Systems Design and Implementation (NSDI).
- [50] Zhiyuan Guo, Yizhou Shan, et al. 2022. Clio: A hardware-software co-designed disaggregated memory system. In ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS).
- [51] Cunchen Hu, Chenxi Wang, et al. 2023. Skadi: Building a Distributed Runtime for Data Systems in Disaggregated Data Centers. In Workshop on Hot Topics in Operating Systems (HotOS).
- [52] Chien-Chin Huang, Gu Jin, and Jinyang Li. 2020. Swapadvisor: Pushing deep learning beyond the gpu memory limit via smart swapping. In International Conference on Architectural Support for Programming Languages and Operating Systems.
- [53] Wenqin Huangfu and Krishna T and the others Malladi. 2022. BEACON: Scalable Near-Data-Processing Accelerators for Genome Analysis near Memory Pool with the CXL Support. In IEEE/ACM International Symposium on Microarchitecture (MICRO).
- [54] IBM. 2024. IBM Power 9 CPU. https://www.ibm.com/it-infrastructure/power/power9.
- [55] Bruce Jacob, David Wang, and Spencer Ng. 2010. Memory systems: cache, DRAM, disk. Morgan Kaufmann.
- [56] Junhyeok Jang, Hanjin Choi, et al. 2023. CXL-ANNS: Software-Hardware Collaborative Memory Disaggregation and Computation for Billion-Scale Approximate Nearest Neighbor Search. In USENIX Annual Technical Conference (USENIX ATC).
- [57] Tatiana Jin, Zhenkun Cai, et al. 2020. Improving resource utilization by timely fine-grained scheduling. In European Conference on Computer Systems (EuroSys).
- [58] and others Junyi Mei, Shixuan Sun. 2024. FlowWalker: A Memory-efficient and High-performance GPU-based Dynamic Graph Random Walk Framework. Proc. VLDB Endow. (2024).
- [59] G Kandiraju, Hubertus Franke, et al. 2014. Software defined infrastructures. IBM Journal of Research and Development.
- [60] Hiwot Tadese Kassa, Jason Akers, et al. 2021. Improving Performance of Flash Based Key-Value Stores Using Storage Class Memory as a Volatile Memory Extension. In USENIX Annual Technical Conference (USENIX ATC).
- [61] Anurag Khandelwal, Yupeng Tang, et al. 2022. Jiffy: Elastic far-memory for stateful serverless analytics. In European Conference on Computer Systems (EuroSys).
- [62] Daehyeok Kim, Tianlong Yu, et al. 2019. FreeFlow: Software-based Virtual RDMA Networking for Containerized Clouds. In USENIX Symposium on Networked Systems Design and Implementation (NSDI).
- [63] Daehyeok Kim, Tianlong Yu, et al. 2019. {FreeFlow}: Software-based Virtual {RDMA} Networking for Containerized Clouds. In USENIX Symposium on Networked Systems Design and Implementation (NSDI).

- [64] Gwangsun Kim, John Kim, et al. 2013. Memory-centric system interconnect design with hybrid memory cubes. In International conference on Parallel architectures and compilation techniques (PACT).
- [65] Atsushi Koshiba, Felix Gust, et al. 2023. Trusted Heterogeneous Disaggregated Architectures. In ACM SIGOPS Asia-Pacific Workshop on Systems (APSys).
- [66] Panos Koutsovasilis, Michele Gazzetti, and Christian Pinto. 2021. A Holistic System Software Integration of Disaggregated Memory for Next-Generation Cloud Infrastructures. In IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid).
- [67] Youngeun Kwon and Minsoo Rhu. 2018. Beyond the memory wall: A case for memory-centric hpc system for deep learning. In IEEE/ACM International Symposium on Microarchitecture (MICRO).
- [68] Andres Lagar-Cavilla, Junwhan Ahn, et al. 2019. Software-defined far memory in warehouse-scale computers. In Architectural Support for Programming Languages and Operating Systems (ASPLOS).
- [69] Christopher Lameter. 2024. Flavors of Memory supported by Linux, Their Use and Benefit. https://events.linuxfoundation.org/.
- [70] Jong-Chern Lee, Jihwan Kim, et al. 2016. High bandwidth memory(HBM) with TSV technique. In International SoC Design Conference (ISOCC).
- [71] Sekwon Lee, Soujanya Ponnapalli, et al. 2022. DINOMO: an elastic, scalable, high-performance key-value store for disaggregated persistent memory. *Proceedings of the VLDB Endowment* (2022).
- [72] Seung-seob Lee, Yanpeng Yu, et al. 2021. Mind: In-network memory management for disaggregated data centers. In Symposium on Operating Systems Principles (SOSP).
- [73] Youngmoon Lee, Hasan Al Maruf, et al. 2022. Hydra: Resilient and highly available remote memory. In Proceedings of the 20th USENIX Conference on File and Storage Technologies (FAST).
- [74] Youngmoon Lee, Hassan Al Maruf, and the others. 2019. Mitigating the Performance-Efficiency Tradeoff in Resilient Memory Disaggregation. Arxiv (2019).
- [75] Baolin Li, Rohin Arora, et al. 2022. AI-Enabling Workloads on Large-Scale GPU-Accelerated System: Characterization, Opportunities, and Implications. In IEEE International Symposium on High-Performance Computer Architecture (HPCA).
- [76] Chao Li, Yushu Xue, et al. 2018. Edge-Oriented Computing Paradigms: A Survey on Architecture Design and System Management. ACM Computer Survey (2018).
- [77] Guoliang Li, Haowen Dong, and Chao Zhang. 2022. Cloud databases: new techniques, challenges, and opportunities. *Proceedings of the* VLDB Endowment (2022).
- [78] Huaicheng Li, Daniel S Berger, et al. 2023. Pond: CXL-Based Memory Pooling Systems for Cloud Platforms. In Architectural Support for Programming Languages and Operating Systems (ASPLOS).
- [79] Haifeng Li, Ke Liu, et al. 2023. HoPP: Hardware-Software Co-Designed Page Prefetching for Disaggregated Memory. In IEEE International Symposium on High-Performance Computer Architecture (HPCA).
- [80] Kai Li and Paul Hudak. 1989. Memory coherence in shared virtual memory systems. ACM Transactions on Computer Systems (TOCS) (1989).
- [81] Yang Li, Saugata Ghose, et al. 2017. Utility-based hybrid memory management. In *IEEE International Conference on Cluster Computing* (*CLUSTER*).
- [82] Zhaogeng Li, Ning Liu, and Jiaoren Wu. 2019. Toward a production-ready general-purpose RDMA-enabled RPC. In ACM SIGCOMM 2019 Conference Posters and Demos (SIGCOMM).
- [83] Shuang Liang, Ranjit Noronha, and Dhabaleswar K Panda. 2005. Swapping to remote memory over infiniband: An approach using a high performance network block device. In IEEE International Conference on Cluster Computing (CLUSTER).
- [84] Kevin Lim, Jichuan Chang, et al. 2009. Disaggregated memory for expansion and sharing in blade servers. In *International Symposium* on *Computer Architecture (ISCA)*.
- [85] Kevin Lim, Jichuan Chang, et al. 2009. Disaggregated memory for expansion and sharing in blade servers. In International Symposium on Computer Architecture (ISCA).
- [86] Linux. 2024. Zswap kernel. https://www.kernel.org/doc/html/latest/admin-guide/mm/zswap.html.
- [87] Haifeng Liu, Long Zheng, Yu Huang, Jingyi Zhou, Chaoqiang Liu, Runze Wang, Xiaofei Liaot, Hai Jinf, and Jingling Xue. 2024. Enabling Efficient Large Recommendation Model Training with Near CXL Memory Processing. In 2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA). 382–395. https://doi.org/10.1109/ISCA59077.2024.00036
- [88] Ling Liu, Wenqi Cao, et al. 2019. Memory Disaggregation: Research Problems and Opportunities. In IEEE International Conference on Distributed Computing Systems (ICDCS).
- [89] Lei Liu, Shengjie Yang, et al. 2019. Hierarchical hybrid memory management in OS for tiered memory systems. IEEE Transactions on Parallel and Distributed Systems (TPDS) (2019).
- [90] Xuchuan Luo, Pengfei Zuo, et al. 2023. SMART: A High-Performance Adaptive Radix Tree for Disaggregated Memory. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [91] Steffen Maass, Changwoo Min, et al. 2017. Mosaic: Processing a trillion-edge graph on a single machine. In *European Conference on Computer Systems (EuroSys)*.
- , Vol. 1, No. 1, Article . Publication date: August 2024.

- [92] The Machine. 2024. The Machine: A new kind of computer. https://www.labs.hpe.com/the-machine.
- [93] Pak Markthub, Mehmet E Belviranli, Seyong Lee, Jeffrey S Vetter, and Satoshi Matsuoka. 2018. DRAGON: breaking GPU memory capacity limits with direct NVM access. In SC18: International Conference for High Performance Computing, Networking, Storage and Analysis. IEEE, 414–426.
- [94] Hasan Al Maruf, Yuhong Zhong, et al. 2021. Memtrade: A Disaggregated-Memory Marketplace for Public Clouds. (2021).
- [95] Pankaj Mehra and Tom Coughlin. 2022. Taming Memory With Disaggregation. Computer (2022).
- [96] Xupeng Miao, Hailin Zhang, et al. 2021. HET: Scaling out Huge Embedding Model Training via Cache-Enabled Distributed Framework. (2021), 312–320.
- [97] Dave Montgomery. 2019. The Future of Data Infrastructure: CDI. https://www.datacenterknowledge.com/industry-perspectives/futuredata-infrastructure.
- [98] Xiaonan Nie, Yi Liu, et al. 2023. Angel-PTM: A Scalable and Economical Large-Scale Pre-Training System in Tencent. (2023).
- [99] Vlad Nitu, Boris Teabe, et al. 2018. Welcome to zombieland: Practical and energy-efficient memory disaggregation in a datacenter. In European Conference on Computer Systems (Eurosys).
- [100] NVIDIA. 2024. GPUDirect RDMA. https://docs.nvidia.com/cuda/gpudirect-rdma/index.html.
- [102] NVIDIA. 2024. Magnum IO GPUDirect Storage. https://developer.nvidia.com/gpudirect-storage.
- [103] NVIDIA. 2024. NvLink Interconnect. http://www.nvidia.com/object/nvlink.html.
- [104] Matheus Ogleari, Ye Yu, and the others. 2019. String figure: A scalable and elastic memory network architecture. In *IEEE International Symposium on High Performance Computer Architecture (HPCA)*.
- [105] Chang Hyun Park, Taekyung Heo, et al. 2017. Hybrid tlb coalescing: Improving tlb translation coverage under diverse fragmented memory allocations. In *International Symposium on Computer Architecture (ISCA)*.
- [106] Jeongmin Brian Park, Vikram Sharma Mailthody, Zaid Qureshi, and Wen-mei Hwu. 2024. Accelerating Sampling and Aggregation Operations in GNN Frameworks with GPU Initiated Direct Storage Accesses. Proc. VLDB Endow. (2024).
- [107] Adarsh Patil, Vijay Nagarajan, et al. 2023. Apta: Fault-tolerant object-granular CXL disaggregated memory for accelerating FaaS. In IEEE/IFIP International Conference on Dependable Systems and Networks (DSN).
- [108] Christian Pinto, Dimitris Syrivelis, et al. 2020. Thymesisflow: A software-defined, hw/sw co-designed interconnect stack for rack-scale memory disaggregation. In Proceedings of the 2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO).
- [109] Matthew Poremba, Itir Akgun, et al. 2017. There and back again: Optimizing the interconnect in networks of memory cubes. ACM SIGARCH Computer Architecture News (CAN) (2017).
- [110] Open Compute Project. 2023. Open Accelerator Infrastructure (OAI) OCP Accelerator Module (OAM) Base Specification. https: //www.opencompute.org/documents/oai-oam-base-specification-r2-0-v1-0-20230919-pdf.
- [111] Zaid Qureshi, Vikram Sharma Mailthody, et al. 2023. BaM: A Case for Enabling Fine-grain High Throughput GPU-Orchestrated Access to Storage. Architectural Support for Programming Languages and Operating Systems (ASPLOS).
- [112] Samyam Rajbhandari, Olatunji Ruwase, and Zero-infinity. 2021. Zero-infinity: Breaking the gpu memory wall for extreme scale deep learning. In International Conference for High Performance Computing, Networking, Storage and Analysis.
- [113] Jie Ren, Jiaolin Luo, et al. 2021. Sentinel: Efficient tensor migration and allocation on heterogeneous memory systems for deep learning. In IEEE International Symposium on High-Performance Computer Architecture (HPCA).
- [114] Amir Roozbeh, Joao Soares, et al. 2018. Software-defined "hardware" infrastructures: A survey on enabling technologies and open research directions. *IEEE Communications Surveys & Tutorials* (2018).
- [115] Amitabha Roy, Laurent Bindschaedler, et al. 2015. Chaos: Scale-out graph processing from secondary storage. In Symposium on Operating Systems Principles (SOSP).
- [116] Nilabja Roy, Abhishek Dubey, and Aniruddha Gokhale. 2011. Efficient autoscaling in the cloud using predictive models for workload forecasting. In IEEE International Conference on Cloud Computing (CLOUD).
- [117] Zhenyuan Ruan, Malte Schwarzkopf, et al. 2020. AIFM: High-Performance, Application-Integrated Far Memory. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [118] Divyanshu Saxena, Tao Ji, et al. 2022. Memory deduplication for serverless computing with Medes. In European Conference on Computer Systems (EuroSys).
- [119] Malte Schwarzkopf, Andy Konwinski, et al. 2013. Omega: flexible, scalable schedulers for large compute clusters. In ACM European Conference on Computer Systems (EuroSys).
- [120] Yizhou Shan, Yutong Huang, et al. 2018. Legoos: A disseminated, distributed OS for hardware resource disaggregation. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [121] Yizhou Shan, Shin-Yeh Tsai, and Yiying Zhang. 2017. Distributed shared persistent memory. In Symposium on Cloud Computing (SoCC).
- [122] Chuanming Shao, Jinyang Guo, et al. 2022. Oversubscribing GPU Unified Virtual Memory: Implications and Suggestions. In ACM/SPEC International Conference on Performance Engineering (ICPE).

- 34 Jing Wang, et al.
- [123] Chuanming Shao, Jinyang Guo, et al. 2022. Oversubscribing GPU Unified Virtual Memory: Implications and Suggestions. In International Conference on Performance Engineering (ICPE).
- [124] Julian Shun and Guy E Blelloch. 2013. Ligra: a lightweight graph processing framework for shared memory. In ACM SIGPLAN symposium on Principles and practice of parallel programming (PPOPP).
- [125] Shihao Song, Anup Das, and Nagarajan Kandasamy. 2020. Exploiting inter-and intra-memory asymmetries for data mapping in hybrid tiered-memories. In ACM SIGPLAN International Symposium on Memory Management (ISMM).
- [126] Jie Sun, Mo Sun, et al. 2023. Helios: An Efficient Out-of-core GNN Training System on Terabyte-scale Graphs with In-memory Performance. arXiv:2310.00837
- [127] Xuan Sun, Hu Wan, et al. 2022. RM-SSD: In-Storage Computing for Large-Scale Recommendation Inference. In IEEE International Symposium on High-Performance Computer Architecture (HPCA).
- [128] Shin-Yeh Tsai, Yizhou Shan, and Yiying Zhang. 2020. Disaggregating Persistent Memory and Controlling Them Remotely: An Exploration of Passive Disaggregated Key-ValueStores. In USENIX Annual Technical Conference (USENIX ATC).
- [129] Shin-Yeh Tsai and Yiying Zhang. 2017. LITE Kernel RDMA Support for Datacenter Applications. In Symposium on Operating Systems Principles (SOSP).
- [130] Evangelos Vasilakis, Vassilis Papaefstathiou, et al. 2020. Hybrid<sup>2</sup>: Combining Caching and Migration in Hybrid Memory Systems. In IEEE International Symposium on High Performance Computer Architecture (HPCA).
- [131] Chenxi Wang, Huimin Cui, et al. 2019. Panthera: Holistic memory management for big data processing over hybrid memories. In ACM SIGPLAN Conference on Programming Language Design and Implementation.
- [132] Chenxi Wang, Haoran Ma, et al. 2020. Semeru: A Memory-Disaggregated Managed Runtime. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [133] Chenxi Wang, Haoran Ma, et al. 2022. MemLiner: Lining up Tracing and Application for a Far-Memory-Friendly Runtime. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [134] Chenxi Wang, Yifan Qiao, et al. 2023. Canvas: Isolated and Adaptive Swapping for Multi-Applications on Remote Memory. In USENIX Symposium on Networked Systems Design and Implementation (NSDI 23).
- [135] Jing Wang, Li Chao, et al. 2022. Excavating the Potential of Graph Workload on RDMA-based Far Memory Architecture. In International Parallel and Distributed Processing Symposium (IPDPS).
- [136] Jing Wang, Li Chao, et al. 2023. Fargraph+: Excavating the Parallelism of Graph Processing Workload on RDMA-based Far Memory System. In *Journal of Parallel and Distributed Computing (JPDC)*.
- [137] Jing Wang, Chao Li, et al. 2022. HyFarM: Task Orchestration on Hybrid Far Memory for High Performance Per Bit. In International Conference on Computer Design (ICCD).
- [138] Jing Wang, Hanzhang Yang, Chao Li, Yiming Zhansun, Yuan Wang, Cheng Xu, Xiaofeng Hou, Minyi Guo, Yang Hu, and Yaqian Zhao. 2024. Boosting Data Center Performance via Intelligently Managed Multi-backend Disaggregated Memory. In International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'24, CCF A, Accepted).
- [139] Linnan Wang, Jinmian Ye, et al. 2018. Superneurons: Dynamic GPU memory management for training deep neural networks. In Principles and practice of parallel programming (PPoPP).
- [140] Pengyu Wang, Chao Li, et al. 2021. Skywalker: Efficient Alias-Method-Based Graph Sampling and Random Walk on GPUs. In Parallel Architectures and Compilation Techniques (PACT).
- [141] Pengyu Wang, Jing Wang, et al. 2021. Grus: Toward unified-memory-efficient high-performance graph processing on gpu. ACM Transactions on Architecture and Code Optimization (TACO) (2021).
- [142] Qing Wang, Youyou Lu, and Jiwu Shu. 2022. Sherman: A Write-Optimized Distributed B+Tree Index on Disaggregated Memory. In International Conference on Management of Data (SIGMOD).
- [143] Ruihong Wang, Chuqing Gao, Jianguo Wang, Prishita Kadam, M TamerÖzsu, and Walid G Aref. 2024. Optimizing LSM-based indexes for disaggregated memory. *The VLDB Journal* (2024), 1–24.
- [144] Ruihong Wang, Jianguo Wang, et al. 2022. The Case for Distributed Shared-Memory Databases with RDMA-Enabled Memory Disaggregation. VLDB (2022).
- [145] Ruihong Wang, Jianguo Wang, et al. 2023. dLSM: An LSM-Based Index for Memory Disaggregation. In IEEE International Conference on Data Engineering (ICDE).
- [146] Xinkai Wang, Hao He, et al. 2023. Not All Resources are Visible: Exploiting Fragmented Shadow Resources in Shared-State Scheduler Architecture. In ACM Symposium on Cloud Computing (SoCC).
- [147] Zixuan Wang, Joonseop Sim, et al. 2022. Enabling Efficient Large-Scale Deep Learning Training with Cache Coherent Disaggregated Memory Systems. In International Symposium on High-Performance Computer Architecture (HPCA).
- [148] Xingda Wei, Rongxin Cheng, et al. 2023. Characterizing Off-path SmartNIC for Accelerating Distributed Systems. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [149] Xingda Wei, Fangming Lu, et al. 2022. KRCORE: A Microsecond-scale RDMA Control Plane for Elastic Computing. In USENIX Annual Technical Conference (USENIX ATC).
- , Vol. 1, No. 1, Article . Publication date: August 2024.

- [150] Xingda Wei, Fangming Lu, et al. 2023. No Provisioned Concurrency: Fast RDMA-codesigned Remote Fork for Serverless Computing. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [151] Michele Weiland, Holger Brunst, et al. 2019. An early evaluation of Intel's optane DC persistent memory module and its impact on high-performance scientific applications. In International Conference for High Performance Computing, Networking, Storage and Analysis (SC).
- [152] Johannes Weiner, Niket Agarwal, et al. 2022. TMO: transparent memory offloading in datacenters. In ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS).
- [153] Du Xiang, Tao Liu, et al. 2018. Two-dimensional multibit optoelectronic memory with broadband spectrum distinction. Nature communications (2018).
- [154] Shao-Peng Yang, Minjae Kim, et al. 2023. Overcoming the Memory Wall with CXL-Enabled SSDs. In USENIX Annual Technical Conference (USENIX ATC).
- [155] Daniel Zahka and Ada Gavrilovska. 2022. FAM-Graph: Graph Analytics on Disaggregated Memory. In IEEE International Parallel and Distributed Processing Symposium (IPDPS).
- [156] Jia Zhan, Itir Akgun, et al. 2016. A unified memory network architecture for in-memory computing in commodity servers. In IEEE/ACM International Symposium on Microarchitecture (MICRO).
- [157] Haoran Zhang, Adney Cardoza, et al. 2020. Fault-tolerant and transactional stateful serverless workflows. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [158] Haoyang Zhang, Yirui Zhou, Yuqi Xue, Yiqi Liu, and Jian Huang. 2023. G10: Enabling an efficient unified gpu memory and storage architecture with smart tensor migrations. In Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture. 395–410.
- [159] Kaiyuan Zhang, Rong Chen, and Haibo Chen. 2015. NUMA-aware graph-structured analytics. In ACM SIGPLAN symposium on principles and practice of parallel programming (PPoPP).
- [160] Ming Zhang, Yu Hua, et al. 2022. FORD: Fast One-sided RDMA-based Distributed Transactions for Disaggregated Persistent Memory. In USENIX Conference on File and Storage Technologies (FAST).
- [161] Mingxing Zhang, Yongwei Wu, et al. 2018. Wonderland: A novel abstraction-based out-of-core graph processing system. ACM SIGPLAN Notices (2018).
- [162] Pengfei Zhang, Xi Li, et al. 2015. HybridSwap: A scalable and synthetic framework for guest swapping on virtualization platform. In IEEE Conference on Computer Communications (INFOCOMM).
- [163] Qizhen Zhang, Xinyi Chen, et al. 2022. Optimizing Data-Intensive Systems in Disaggregated Data Centers with TELEPORT. In International Conference on Management of Data (SIGMOD).
- [164] Yunming Zhang, Ajay Brahmakshatriya, et al. 2020. Optimizing Ordered Graph Algorithms with GraphIt. In International Symposium on Code Generation and Optimization (CGO).
- [165] Yunming Zhang, Vladimir Kiriansky, et al. 2017. Making caches work for graph analytics. In *IEEE International Conference on Big Data* (*Big Data*).
- [166] Yang Zhou, Hassan MG Wassel, et al. 2022. Carbink: {Fault-Tolerant} Far Memory. In USENIX Symposium on Operating Systems Design and Implementation (OSDI).
- [167] Bohong Zhu, Youmin Chen, et al. 2021. Octopus+: An rdma-enabled distributed persistent memory file system. ACM Transactions on Storage (TOS) (2021).
- [168] Xiaowei Zhu, Wentao Han, and Wenguang Chen. 2015. GridGraph: Large-scale graph processing on a single machine using 2-level hierarchical partitioning. In USENIX Annual Technical Conference (USENIX ATC).
- [169] Ziyi Zhu, Yiwen Shen, et al. 2019. Flexible resource allocation using photonic switched interconnects for disaggregated system architectures. In *Optical Fiber Communication Conference (OFC)*.
- [170] Tobias Ziegler, Sumukha Tumkur Vani, et al. 2019. Designing Distributed Tree-based Index Structures for Fast RDMA-capable Networks. In International Conference on Management of Data (SIGMOD).