

Design Tradeoffs in CXL-Based Memory Pools for Public Cloud Platforms

Daniel S. Berger  and Daniel Ernst, *Microsoft Azure, Redmond, WA, 98052, USA*

Huaicheng Li, *Virginia Tech, Blacksburg, VA, 24061, USA*

Pantea Zardoshti, Monish Shah, Samir Rajadnya, Scott Lee, and Lisa Hsu, *Microsoft Azure, Redmond, WA, 98052, USA*

Ishwar Agarwal, *Intel Corporation, Santa Clara, CA, 95054, USA*

Mark D. Hill, *Microsoft Azure, Redmond, WA, 98052, USA and also with the University of Wisconsin-Madison, Madison, WI, 53715, USA*

Ricardo Bianchini, *Microsoft Azure, Redmond, WA, 98052, USA*

Dynamic random-access memory (DRAM) is a key driver of performance and cost in public cloud servers. At the same time, a significant amount of DRAM is underutilized due to fragmented use across servers. Emerging interconnects such as Compute Express Link (CXL) offers a path toward improving utilization through memory pooling. However, the design space of CXL-based memory systems is large, with key questions around the size, reach, and topology of the memory pool. At the same time, using pools require navigating complex design constraints around performance, virtualization, and management. This article discusses why cloud providers should deploy CXL memory pools, key design constraints, and observations in designing toward practical deployment. We identify configuration examples with significant positive return of investment.

Many public cloud customers deploy their workloads via virtual machines (VMs). VMs enable performance comparable to on-premises datacenters without the need to manage datacenters. Cloud providers face the challenge of achieving excellent performance at a competitive hardware cost.

A key driver of both performance and cost is main memory. The gold standard for memory performance is to preallocate a VM with cores and memory on the same socket. This leads to memory latency below 100 ns and facilitates virtualization acceleration. At the same time, Dynamic random-access memory (DRAM) has become a major portion of hardware cost due to its poor scaling properties with only nascent alternatives.¹ For example, DRAM can be over 50% of server cost.²

Through analysis of Azure VM traces, we identify *memory stranding* as a dominant source of memory waste and a potential source of cost savings. Stranding happens when all server cores are rented (i.e., allocated to customer VMs) but unallocated memory capacity remains and cannot be rented. We find that up to 30% of DRAM becomes stranded as more cores become allocated to VMs.

Limitations of the state-of-the-art: Reducing DRAM usage in the public cloud is challenging due to its stringent performance requirements. Pooling memory via memory disaggregation is a promising approach because stranded memory can be returned to the disaggregated pool and used by other servers. Unfortunately, existing pooling systems have microsecond access latencies and require page faults or changes to the VM guest.^{3,4}

The emerging CXL interconnect: The emerging Compute Express Link (CXL) interconnect⁵ enables cacheable load/store (ld/st) accesses to pooled memory on many current processors. Pool-memory accesses via loads/stores is a game changer for cloud computing as it allows

memory to remain statically preallocated while physically being located in a shared pool. However, CXL access latency depends on the overall system design, especially the pool size [the number of central processing unit (CPU) sockets able to use a given pool] and topology. Larger pools require traversing switching levels, which adds significant latency. In addition, each CXL component adds to the system cost, which must be balanced against stranding savings.

THROUGH ANALYSIS OF AZURE VM TRACES, WE IDENTIFY MEMORY STRANDING AS A DOMINANT SOURCE OF MEMORY WASTE AND A POTENTIAL SOURCE OF COST SAVINGS.

This work: This work is motivated by the memory stranding problem identified in Pond² and we paraphrase the stranding analysis in the “Cloud Workload Characterization” section. While Pond focuses on system software policies and mechanisms for allocating/managing pooled memory, this work focuses on design tradeoffs in the pool’s hardware configuration. First, we characterize pool components, possible topologies, and associate memory access latencies. We derive a set of design recommendations from this analysis. Second, we compare savings from memory pooling to the cost of its components for different pool sizes and CXL device types. We find that CXL-based memory pooling can yield significant positive returns on investment. Contrary to the focus of existing literature, smaller pools may be attractive. Third, we discuss future directions for the industry as well as academic research.

BACKGROUND

Cloud resource allocation: Public cloud workloads run inside VMs. To offer performance close to dedicated (nonvirtualized) resources, VM resources are statically allocated by reserving each resource (CPU, DRAM, network bandwidth, etc.) for a VM’s lifetime. Additionally, providers optimize input/output performance with virtualization accelerators that bypass the hypervisor.⁶ For example, accelerated networking is enabled by default on AWS and Azure. Virtualization acceleration requires statically preallocating (or “pinning”) a VM’s entire address space.⁷

Cloud resource scheduling: Scheduling VMs with heterogeneous multidimensional resource demands onto servers leads to a challenging bin-packing problem.^{8,9}

Scheduling is further complicated by constraints such as spreading VMs across multiple failure domains.

A simplified view of Azure’s VM scheduling is that VMs are first assigned to a compute cluster and then placed on a specific server within this cluster. A cluster roughly corresponds to a row of racks with homogeneous server configurations. We use the unit of a cluster to characterize our workloads.

Memory stranding: It is often difficult to provision servers that closely match the resource demands of the incoming VM mix. A common reason is that the DRAM-to-core ratio of a server that will last years must be determined at platform design time and is statically fixed over its lifetime. Additionally, fixed-size DIMMs over limited freedom in determining the DRAM-to-core ratio.

When the DRAM-to-core ratio of VM arrivals and a cluster’s server resources do not match, tight packing becomes especially difficult. We define a resource as *stranded* when it is technically available to be rented to a customer, but is practically unavailable as some other resource has been exhausted. The typical scenario for *memory stranding* is that all cores have been rented, but there is still memory available in the server.

Reducing stranding via pooling: This work proposes to break the fixed hardware configuration of servers by disaggregating memory into a pool that is accessible by multiple hosts.¹⁰ By dynamically reassigning memory to different hosts at different times, we can shift memory resources to where they are needed. Thus, we can provision servers close to the average DRAM-to-core ratios and tackle deviations via the memory pool.

Pooling via CXL: The CXL.mem protocol for *ld/st* memory semantics maps device memory to the system address space. Last-level cache (LLC) misses to CXL memory addresses translate into requests on a CXL port whose responses bring the missing cache-lines. Similarly, LLC write-backs translate into CXL data writes. CXL memory is virtualized using hypervisor page tables and the memory-management unit and is thus compatible with virtualization acceleration. CXL.mem uses PCIe’s electrical interface with custom link and transaction layers for low latency. Intel measures CXL port latencies at 25-ns round-trip.¹¹ With PCIe 5.0, the bandwidth of a bidirectional $\times 8$ -CXL port at a typical 2:1 read:write ratio roughly matches an 80-bit DDR5-4800 channel.

CLOUD WORKLOAD CHARACTERIZATION

Stranding at Azure

We summarize previous analysis on stranding.²

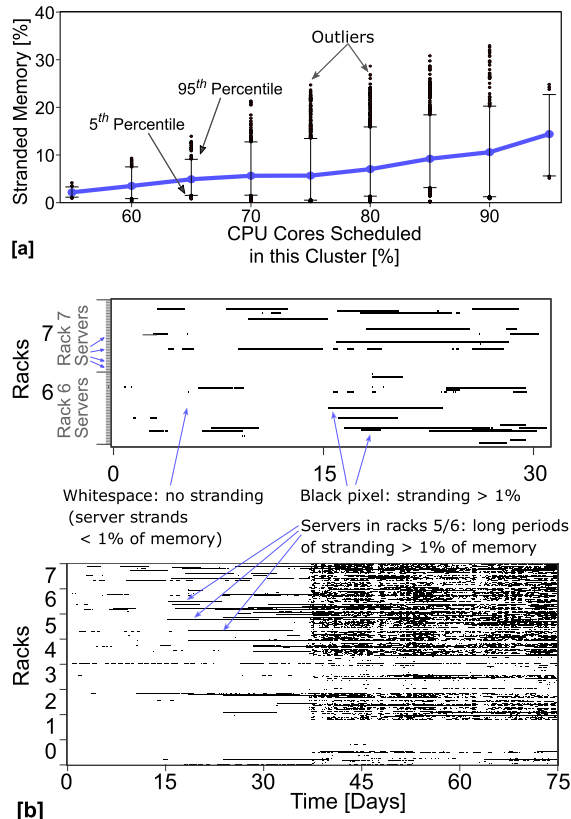


FIGURE 1. Memory stranding. (a) Stranding increases significantly as more CPU cores are scheduled; (b) Stranding changes dynamically over time.

Dataset: We measure stranding in 100 general-purpose clusters over a 75-day period. A general-purpose cluster hosts a mix of first-party and third-party VM workloads that do not require special hardware (such as graphics processing units). We select clusters with similar deployment years, spanning major regions on the planet. Each cluster trace contains millions of per-VM arrival/departure events.

Memory stranding: Figure 1(a) shows the hourly average amount of stranded DRAM across our cluster sample, bucketed by the percentage of scheduled CPU cores. In clusters where 75% of CPU cores are scheduled for VMs, 6% of memory is stranded. This grows to over 10% when ~85% of CPU cores are allocated to VMs. This makes sense since stranding is an artifact of highly utilized nodes, which correlates with highly utilized clusters. Outliers are shown by the error bars, representing 5th and 95th percentiles. At 95th, stranding reaches 25% during high utilization periods. Individual outliers reach more than 30% stranding. Figure 1(b) shows stranding over time across eight adjacent racks. Every row shows a server within each

rack. A workload change (around day 36) suddenly increased stranding significantly. Furthermore, stranding can affect many racks concurrently (e.g., racks 2, 4–7) and it is generally hard to predict which clusters/racks will have stranded memory.

VM Memory Utilization in Azure

Dataset: We perform measurements on the same 100 general-purpose production clusters. For untouched memory, we rely on guest-reported memory usage counters cross-referenced with hypervisor page table access bit scans. We sample memory bandwidth counters using Intel RDT¹² for a subset of clusters with compatible hardware. Finally, we use hypervisor counters to measure nonuniform memory access (NUMA) spanning in dual-socket servers, where a VM has cores on one socket and some memory from another socket.

Memory bandwidth: Memory bandwidth usage of general-purpose workloads is generally low with average bandwidth utilization below 10 GB/s. VMs on a small number of hosts do, however, use 100% of memory bandwidth.

NUMA spanning: Most VMs are small and can fit on a single socket. Azure’s hypervisor aims to schedule VMs on dual-socket servers such that they fit entirely (cores and memory) on a single NUMA node. We find that spanning occurs for only 2%–3% of VMs.

Overall, untouched memory and low memory bandwidth requirements make VM workloads a good fit for memory pooling. However, with 97%–98% of VMs using NUMA-local memory, performance parity for pooled memory will be challenging.

Workload Sensitivity to Memory Latency

We summarize previous experiments on latency sensitivity.²

Dataset: We evaluate 158 workloads across proprietary workloads, in-memory stores, data processing, and benchmark suites. They run on dual-socket Intel Skylake 8157 M, with a 182% latency increase for socket-remote memory, or AMD EPYC 7452, with 222% latency increase. We normalize performance as slowdown relative to NUMA-local performance.

Latency sensitivity: Figure 2 surveys workload slowdowns. Under a 182% increase in memory latency, we find that 26% of the 158 workloads experience less than 1% slowdown under CXL. At the same time, some workloads are severely affected with 21% of the workloads facing > 25% slowdowns. Overall, every workload class has at least one workload with less than 5% slowdown and one workload with more than 25% slowdown (except SPLASH2x). Our proprietary




FIGURE 2. Performance slowdowns when memory latency increases by 182%–222% (see the “Workload Sensitivity to Memory Latency” section). Workloads have different sensitivity to increased memory latency as they would see with CXL. X-axis shows 158 representative workloads; Y represents the normalized performance slowdown, i.e., performance under higher (remote) latency relative to all local memory. “Proprietary” denotes production workloads at Azure.

workloads are less impacted than the overall workload set with almost half seeing $< 1\%$ slowdown. These production workloads are NUMA-aware and often include data placement optimizations.

Under a 222% increase in memory latency, we find that 23% of the 158 workloads experience less than 1% slowdown under CXL. More than 37% of workloads face $> 25\%$ slowdowns—a significantly higher fraction than on the 182% emulated latency increase. We find that the processing pipeline for some workloads, like VoltDB, seems to have just enough slack to accommodate the smaller 182% latency increase with significant pipeline stalls for 222% latency increase. Other workload classes like graph processing (GAPBS) are sensitive to both latency and bandwidth, and both effects are worsened on the 222% system.

THE MEMORY POOL DESIGN SPACE

Designing a memory pool involves multiple hardware components and design choices that expand with every new CXL release. To limit complexity, we focus on two design aspects: 1) whether to provide connectivity via CXL switches or through CXL multiheaded devices (MHDs) (see Sec. 2.5)⁵ and 2) how large the constructed pool should be to maximize return-on-investment (ROI). We discuss a particular set of choices suitable for general-purpose cloud computing. Other use cases may see different sets of choices and tradeoffs.

Components

CXL memory controller (MC) devices act as a bridge between the CXL protocol and memory devices such as DDR5 DRAMs. Today’s MCs typically bridge between 1-2 CXL $\times 8$ ports and 1-2 80b channels of DDR5.¹³

CXL switches behave similar to other network switches in that they forward requests and data, without serving as an endpoint. Physically, CXL switches will likely share many characteristics (e.g., port count) with PCIe switches, due to using the same physical interface. For the purposes of this analysis, we assume that switches with 128 lanes (16 ports) of CXL are used to build a fabric layer.

A CXL MHD essentially combines a switch and a memory controller in a single device. Specifically, the MHD offers multiple CXL ports and appears to each connected host as a single logical memory device.⁵ The most significant tradeoffs for MHD designs are the number of incoming CXL ports and DDR channels. A useful design comparison is a modern server CPU IO-die (IOD), such as the one in AMD Genoa.¹⁴ The Genoa IOD offers 128 PCIe5 lanes as well as 12 DDR5 channels. With the $\times 8$ -CXL requirement, this would be analogous to a 16-headed device with at least 8 channels of DDR5. In our analysis, we consider both this 16-headed device as well as a smaller 8-headed device.

Pool Size Versus Latency

At a high level, the first design decision is whether cloud compute servers can pool all of their memory. With 21%–37% of workloads facing significant slowdowns on pool-only configurations (see the “Cloud Workload Characterization” section), we do not recommend fully disaggregating compute and memory. Servers need to retain significant amounts of local DRAM to maintain performance expectations, which will likely go beyond the scope of on-die memory. Further, achieving maximum memory bandwidth requires CPUs to populate all available local DDR channels, creating a practical minimum for local memory capacity.

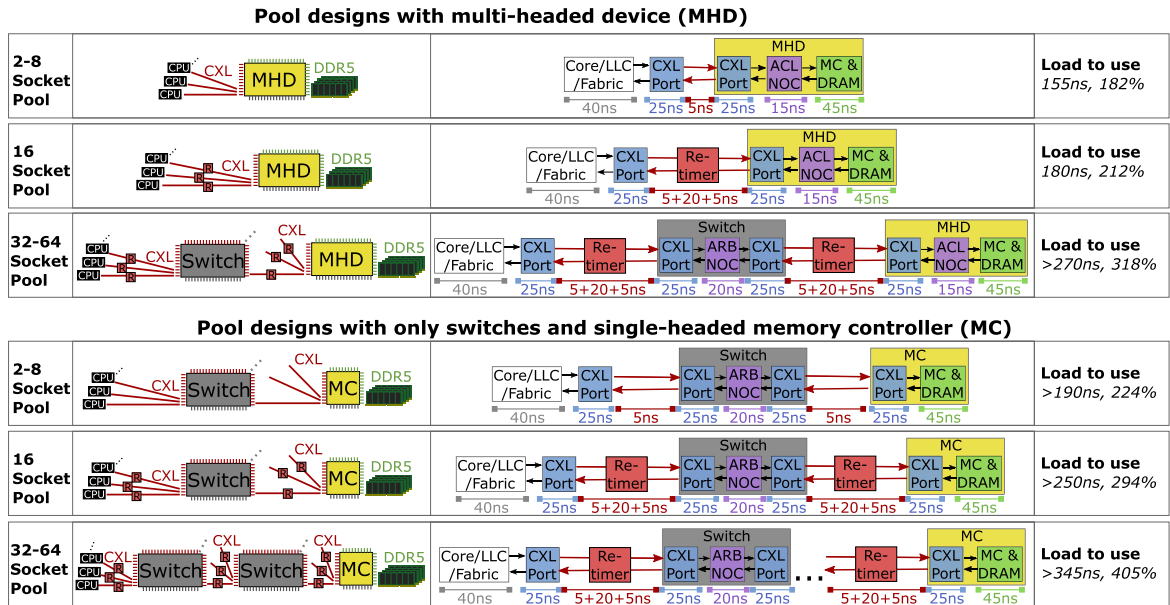


FIGURE 3. Pool size and latency tradeoffs. Small pools of 8–16 sockets add only 75–90 ns relative to NUMA-local DRAM. Latency increases for larger pools that require retimers and a switch.

Observation 1: A significant percentage (more than 25%) of datacenter memory needs to remain local to compute servers.

To understand pool latencies, we first characterize the impact on latency of achievable topologies given viable components.

Observation 2: When using at least a $\times 8$ -CXL port for each host, pool sizes beyond 16–32 hosts will require at least one level of switches if MHDs are used or two levels of switches if using only individual MCs.

Access latencies derive from multiple parameters. Port latency plays a dominant role with initial measurements indicating 25 ns.¹¹ Retimers are devices used to maintain CXL/PCIe signal integrity over distances above roughly half a meter, depending on the implementation of the signal path. They add about 10 ns of latency in each direction.¹⁵ Each switch will add at least 70 ns of latency due to ports, arbitration, and network-on-chip (NOC).

Figure 3 shows a range of CXL path types based on pool sizes and the use of MHDs versus switches with single-headed devices. We find that small 8 and 16-socket pools using MHDs increase latencies to 182%–212% (155–180 ns) relative to NUMA-local DRAM. Latency when using only switches and single-headed memory controllers would further increase by 23%–38%.

Rack-scale pooling with 64 sockets would increase latencies by 318%–405% (270–345 ns) and pooling across multiple racks would require yet another level of switching and potentially multiple retimers, increasing latencies by more than 465% (395 ns). Comparing these latencies to the slowdowns observable at 182%–222% (see the “Cloud Workload Characterization” section), we observe that large-scale pooling will likely be prohibitive from a performance perspective.

Observation 3: The size of CXL-based memory pools will likely be a subset of a rack to minimize the performance impact of access latencies.

Modern CPUs can connect to multiple MHDs or switches, which allows scaling to meet bandwidth and capacity goals for different clusters.

Pool Size Versus DRAM Savings

We analyze VM-to-server traces from Azure (see the “Cloud Workload Characterization” section) to estimate the amount of DRAM that could be saved via pools of different sizes. The reduction in DRAM comes from averaging host’s peak memory needs across the pool. Our simulation plays back VM traces while assigning a fixed percentage of pool memory. We repeatedly run cluster simulations while decreasing overall memory in 0.5% steps until the first VM is rejected. The minimum amount of cluster memory

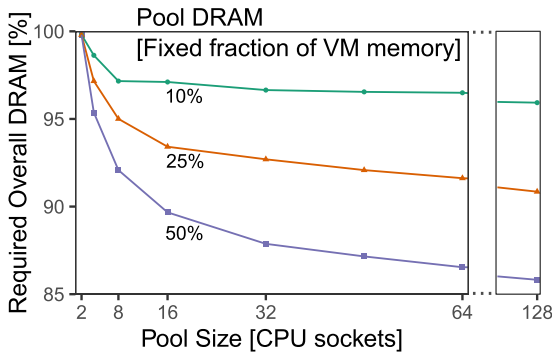


FIGURE 4. Impact of pool size. Small pools of 32 sockets are sufficient to significantly reduce overall memory needs.

corresponds to the “required overall DRAM” reported below.

Figure 4 presents cluster DRAM requirements when VMs are assigned either 10%, 30%, or 50% of pool DRAM. As the pool size increases, the figure shows that required overall DRAM decreases. However, this effect diminishes for larger pools. For example, with a fixed 50% pool DRAM, a pool with 32 sockets saves 12% of DRAM while a pool with 64 sockets saves 13% of DRAM. Note that allocating 50% of VM memory to pool DRAM requires latency mitigation techniques (see the “Discussion and Conclusion” section). Besides low latency, feasible configurations also must be ROI positive, as discussed next.

Pool Size Versus System Cost

System cost depends on many factors. We consider a simplified model that focuses on key hardware components: DRAM, memory controllers, cables, and the memory blade enclosure/printed circuit board (PCB). Our model ignores factors of time, scale, and market competition. Specifically, our model calculates cost relative to a nonpooled server’s bill of materials (BOM) based on the following set of parameters.

- › **MC:** cost of a typical 2×8 CXL memory controller (e.g., 0.4%).
- › **MHD8:** cost of an 8-headed memory controller (e.g., 0.8%).
- › **MHD16:** cost of a 16-headed memory controller (e.g., 2.0%).
- › **Switch:** cost of a 16-port CXL switch (e.g., 1.6%).
- › **Ret:** cost of a CXL retimer (e.g., 0.02%).
- › **Infra:** cost of the supporting memory enclosure, PCBs, and cables expressed as a multiplier applied to MHD or switch cost (e.g., 0.5–2×).




FIGURE 5. Pool system cost tradeoffs. Both cost and savings increase with pool size. Infrastructure overheads also play a key factor in cost. Cost savings (black line) from Figure 4 are workload dependent and may look significantly different for other use cases. We advise practitioners to evaluate savings for their workloads.

The exemplary values for the parameters are roughly based on estimates of silicon area as well as connectivity and infrastructure necessary to support the memory pools. Note that there is significant room for these parameters to change between companies, server configurations, use cases, and over time.

Figure 5 presents cost overheads for pool sizes from 2–64 sockets and for pools encompassing two different capacity points relative to total system memory. The baseline for comparison is the full cost of a non-pooled server, including CPU, DRAM, and other standard infrastructure [e.g. network interface cards (NICs), power delivery, management controllers, boards, etc.]. Within this baseline, DRAM memory is assumed to account for approximately half of the total cost, with the CPU and other infrastructure splitting the other half. All other modeled configurations hold the total cost of the base system constant, but add the costs of the extra components required for pooling part of the memory. Our results are reported as a percentage of cost uplift versus the baseline configuration. We vary the infrastructure overhead cost to show that the overall costs are very sensitive to the ability for a design to cost-effectively provide connectivity to the pool. The analysis also shows that overhead for switch-based designs versus MHD designs is significant. As an example, an 8-socket pool

implemented with switches adds over two-and-a-half times the cost of an 8-socket pool based on MHDs.

This overhead is important, as the system-level goal is reaching a beneficial pooling configuration, which is one where the cost uplift of moving memory into the pool is less than the efficiency benefit of having flexible memory as outlined in the savings analysis above. In Figure 5, the black line plots the savings estimate from the earlier analysis (Figure 4). Configurations below this line are ROI positive, while those above the line are likely ROI negative unless further optimizations can be made to improve savings. Note in particular that most switch-based configurations are ROI negative, while many MHD-based configurations are ROI positive, especially for smaller pool sizes.

Observation 4: Positive ROI requires pool designers to navigate a complex tradeoff between pool size, topology, and savings, which is workload dependent. Infrastructure overheads may become a major hurdle to adopting CXL-based pooling as expensively designed configurations will not achieve beneficial ROI.

RELATED WORK

Low memory resource utilization and stranding has been observed at Google¹⁶ and Microsoft.¹⁷ This motivated at least three lines of prior research on memory pooling prior to CXL.

Hypervisor/OS level approaches such as in Gu et al.³ rely on page faults and access monitoring to maintain the working set in local DRAM. In the context of general-purpose cloud computing, these OS-based approaches bring too much overhead and jitter. They are also incompatible with virtualization acceleration (e.g., DDA).

Runtime-based disaggregation designs^{4,18} proposed customized application programming interfaces for remote memory access. While effective, this approach requires developers to explicitly use these mechanisms at the application level.

Hardware-based memory disaggregation have served as an inspiration for CXL but prior approaches were not available on commodity hardware.^{10,19} Prior analysis of requirements for disaggregation are related to our goals. However, network-based disaggregation²⁰ lead to a different design space, e.g., with latency considered in the range of 1–40 μ s, whereas we consider latencies lower by an order of magnitude.

DISCUSSION AND CONCLUSION

CXL-based memory pooling promises to reduce DRAM needs for general-purpose cloud platforms. This paper outlines the design space for memory pooling and offers a framework to evaluate different proposals.

WE HIGHLIGHT THAT SMALL POOLS, SPANNING UP TO 16 SOCKETS, CAN LEAD TO SIGNIFICANT DRAM SAVINGS. THIS REQUIRES KEEPING INFRASTRUCTURE COST OVERHEADS LOW, WHICH REINFORCES THE NEED FOR STANDARDIZATION OF POOLING INFRASTRUCTURE.

As cloud datacenters are quickly evolving, some key parameters will differ significantly even among cloud providers and over time. The fraction of VM memory that can be allocated on CXL pools depends largely on the type of latency mitigation employed. For example, the recent Pond² system can allocate an average of 35%–44% of DRAM on CXL pools while satisfying stringent cloud performance goals. Future techniques for performance management may lead to significantly higher CXL pool usage. Another difference comes from server and infrastructure cost breakdowns, which lead to entirely different cost curves (Figure 5).

Regardless of the variability in system and cost parameters, we believe that Observations 1-4 broadly apply to general-purpose clouds. We highlight that small pools, spanning up to 16 sockets, can lead to significant DRAM savings. This requires keeping infrastructure cost overheads low, which reinforces the need for standardization of pooling infrastructure. Latency and cost increase quickly for larger pool sizes, while the efficiency benefits fall off, which may make large pools counterproductive in many scenarios.

Our savings model focuses on pooling itself, e.g., averaging peak DRAM demand across the pool, and for Azure specific workloads. CXL also enables other savings including using cheaper media behind a CXL controller, such as reusing DDR4 from decommissioned servers. We advise practitioners to create a savings model for their specific use cases, which might differ from ours.

CXL reopens memory controller architecture as a research frontier. With memory controllers decoupled from CPU sockets, new controller features can be more quickly explored and deployed. Cloud providers need improved reliability, availability, and serviceability (RAS) capabilities including memory error correction, management, and isolation. Tighter integration between memory chips, modules, and controllers can enable improvements along the Pareto frontier of RAS, memory bandwidth, and latency.

REFERENCES

1. S. Shiratake, "Scaling and performance challenges of future DRAM," in *Proc. IEEE Int. Memory Workshop*, 2020, pp. 1–3.
2. H. Li et al., "Pond: CXL-Based memory pooling systems for cloud platforms," in *Proc. Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, 2023, pp. 574–587.
3. J. Gu, Y. Lee, Y. Zhang, M. Chowdhury, and K. G. Shin, "Efficient memory disaggregation with INFINISWAP," in *Proc. 14th USENIX Symp. Netw. Syst. Des. Implementation*, 2017, pp. 649–667.
4. Z. Ruan, M. Schwarzkopf, M. K. Aguilera, and A. Belay, "AIFM: High-performance, application-integrated far memory," in *Proc. 14th USENIX Conf. Oper. Syst. Des. Implementation*, 2020, pp. 315–332.
5. "CXL Specification," 2020. Accessed: Dec. 2020. [Online]. Available: <https://www.computeexpresslink.org/download-the-specification>
6. H. Li et al., "LeapIO: Efficient and portable virtual NVMe storage on ARM SoCs," in *Proc. 25th Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, 2020, pp. 591–605.
7. I. Lesokhin et al., "Page fault support for network controllers," in *Proc. 22nd Int. Conf. Archit. Support Program. Lang. Operating Syst.*, 2017, pp. 449–466.
8. O. Hadary et al., "Protean: VM allocation service at scale," in *Proc. 14th USENIX Conf. Operating Syst. Des. Implementation*, 2020, pp. 845–861.
9. E. Cortez, A. Bonde, A. Muzio, M. Russinovich, M. Fontoura, and R. Bianchini, "Resource central: Understanding and predicting workloads for improved resource management in large cloud platforms," in *Proc. 26th Symp. Oper. Syst. Princ.*, 2017, pp. 153–167.
10. C. Pinto et al., "ThymesisFlow: A Software-defined, HW/SW co-designed interconnect stack for rack-scale memory disaggregation," in *Proc. 53rd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2020, pp. 868–880.
11. D. D. Sharma, "Compute express link: An open industry-standard interconnect enabling heterogenous data-centric computing," in *Proc. IEEE Symp. High-Perform. Interconnects*, 2020, pp. 5–12.
12. "Intel resource director technology (Intel RDT)," 2015. Accessed: Sep. 2022. [Online]. Available: <https://www.intel.com/content/www/us/en/architecture-and-technology/resource-director-technology.html>
13. Asteralabs Leo, "Memory connectivity platform for CXL 1.1 and 2.0," 2022. Accessed: Aug. 2022. [Online]. Available: https://www.asteralabs.com/wp-content/uploads/2022/08/Asteralabs_Leo_Aurora_Product_FINAL.pdf
14. L. Su, "AMD unveils workload-tailored innovations and products at the accelerated data center premiere," Nov. 2021. [Online]. Available: <https://www.amd.com/en/press-releases/2021-11-08-amd-unveils-workload-tailored-innovations-and-products-the-accelerated>
15. "CXL use-cases driving the need for low latency performance retimers," 2021. [Online]. Available: <https://www.microchip.com/en-us/about/blog/learning-center/cxl-use-cases-driving-the-need-for-low-latency-performance-reti>
16. M. Tirmazi et al., "Borg: The next generation," in *Proc. 15th Eur. Conf. Comput. Syst.*, 2020, pp. 1–14.
17. Q. Zhang, P. A. Bernstein, D. S. Berger, and B. Chandramouli, "Redy: Remote dynamic memory cache," *Proc. VLDB Endowment*, vol. 15, pp. 766–779, 2021.
18. I. Calciu et al., "Rethinking software runtimes for disaggregated memory," in *Proc. 26th ACM Int. Conf. Architectural Support Program. Lang. Oper. Syst.*, 2021, pp. 79–92.
19. Z. Guo, Y. Shan, X. Luo, Y. Huang, and Y. Zhang, "Clio: A hardware-software co-designed disaggregated memory system," in *Proc. 26th ACM Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, 2022, pp. 417–433.
20. X. Peter et al., "Network requirements for resource disaggregation," in *Proc. 2th USENIX Conf. Operating Syst. Des. Implementation*, 2016, pp. 249–264.

DANIEL S. BERGER is a senior researcher in the Azure Systems Research Group, Microsoft Azure, Redmond, WA, 98052, USA. Berger received a Ph.D. degree in computer science from TU Kaiserslautern. Contact him at daberg@microsoft.com.

DANIEL ERNST is a principal architect in the Leading Edge Architecture Pathfinding (LEAP), Microsoft Azure, Redmond, WA, 98052, USA. Ernst received a Ph.D. degree in computer science and engineering from the University of Michigan. Contact him at danernst@microsoft.com.

HUAICHENG LI is an assistant professor at Virginia Tech, Blacksburg, VA, 24061, USA. Li received a Ph.D. degree in computer science from the University of Chicago. Contact him at huaicheng@cs.vt.edu.

PANTEA ZARDOSHTI is a research software development engineer in the AzSR Group, Microsoft Azure, Redmond, WA, 98052, USA. Zardoshti received a Ph.D. degree in computer science from the Lehigh University. Contact her at pzardoshti@microsoft.com.

MONISH SHAH is a senior principal hardware engineer in the LEAP Group at Microsoft Azure, Redmond, WA, 98052, USA. Shah received an M.Sc. degree in electrical engineering from Stanford University. Contact him at monish.shah@microsoft.com.

SAMIR RAJADNYA is a principal memory system engineer in the LEAP Group, Microsoft Azure, Redmond, WA, 98052, USA. Rajadnya received an M.Tech. degree in electrical engineering from IIT Bombay. Contact him at srajadnya@microsoft.com.

SCOTT LEE is a principal software engineer lead at Microsoft, Redmond, WA, 98052, USA. Lee received a B.Sc. degree in computer engineering from the University of Washington. Contact him at scolee@microsoft.com.

LISA HSU is a principal architect at Microsoft Azure, Redmond, WA, 98052, USA. Hsu received a Ph.D. degree in computer science from the University of Michigan. Contact her at lisa.hsu@microsoft.com.

ISHWAR AGARWAL is a senior principal engineer at Intel Corporation, Santa Clara, CA, 95054, USA. Agarwal received

an M.Sc. degree in electrical and computer engineering from Georgia Tech. Contact him at ishwar.agarwal@intel.com.

MARK D. HILL is a partner architect and leads the LEAP Group at Microsoft Azure, Redmond, WA, 98052, USA and also with the University of Wisconsin-Madison, Madison, WI, 53715, USA. Hill received a Ph.D. degree in computer science from UC Berkeley and served 32 years at University of Wisconsin Computer Science. Contact him at markhill@microsoft.com.

RICARDO BIANCHINI is a distinguished engineer at Microsoft Azure, Redmond, WA, 98052, USA. Bianchini received a Ph.D. degree in computer science from University of Rochester. Contact him at ricardob@microsoft.com.

Over the Rainbow: 21st Century Security & Privacy Podcast

Tune in with security leaders of academia, industry, and government.

Bob Blakley and **Lorrie Cranor**

OVER THE RAINBOW
by IEEE Security & Privacy

Subscribe Today
www.computer.org/over-the-rainbow-podcast