Coughlin Associates

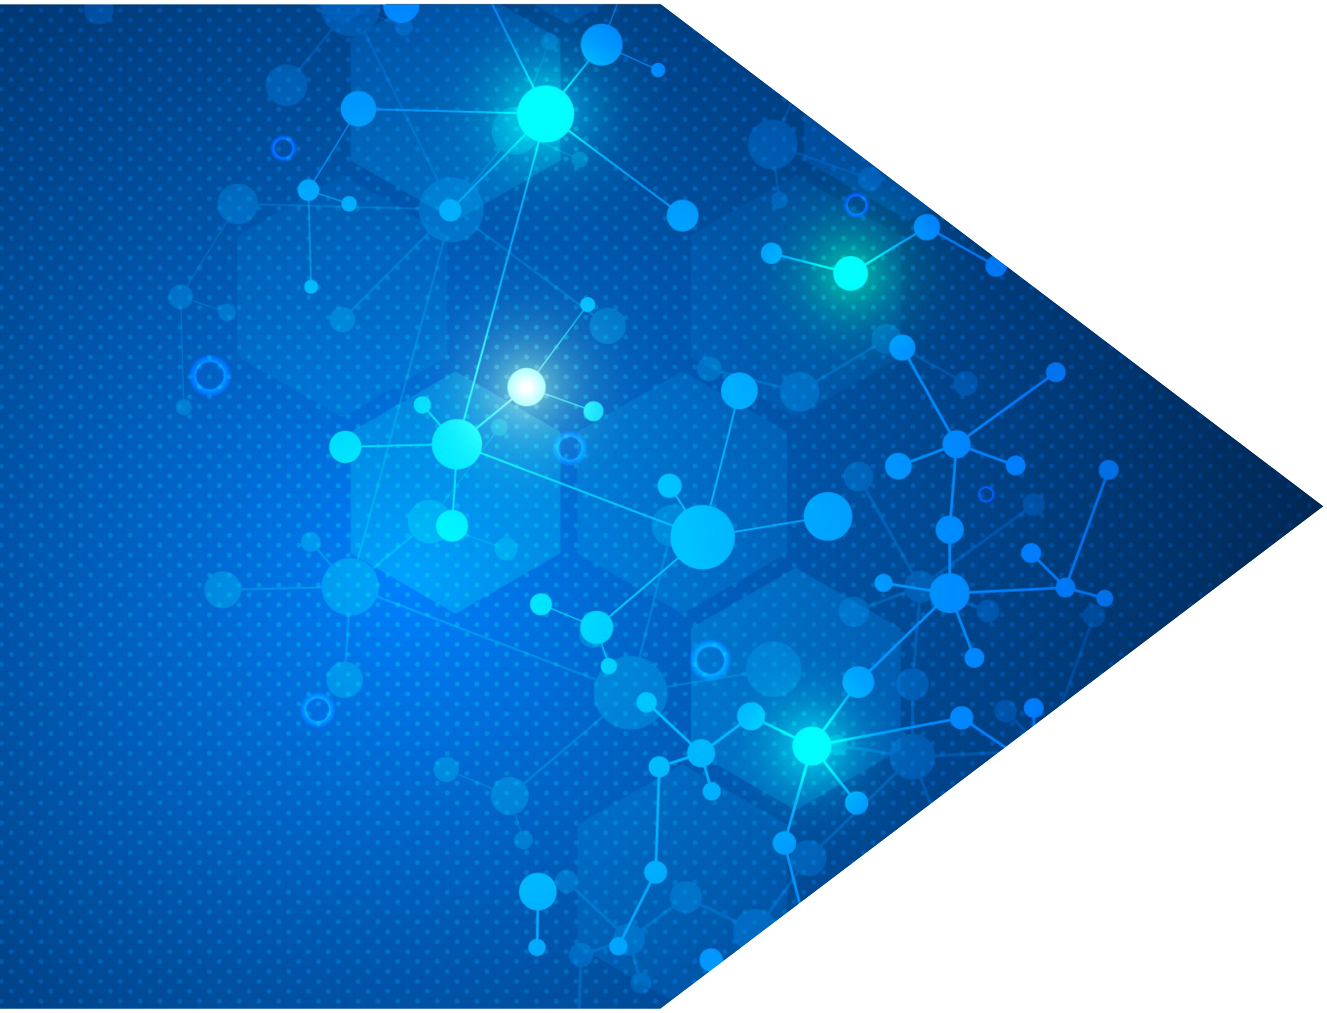

ANALYST REPORT

Memory Fabric Management For Big Data

Tom Coughlin, Coughlin Associates

Introduction

The amount of big data processed in real-time is growing exponentially as a result of the growth of applications such as fraud detection in financial services, animation and visual effects in the entertainment industry, recommendation engines in retail, and genomic research in health sciences, to name just a few. For example, the figure below shows cell samples used for genomic research have grown from 100 cells in 2009 to 1 million today. At the same time the modalities of data[1] have increased from a single modality to dozens.

Exponential Growth of Cell Data Used in Single-Cell Sequencing
(source: Analytical Biosciences)

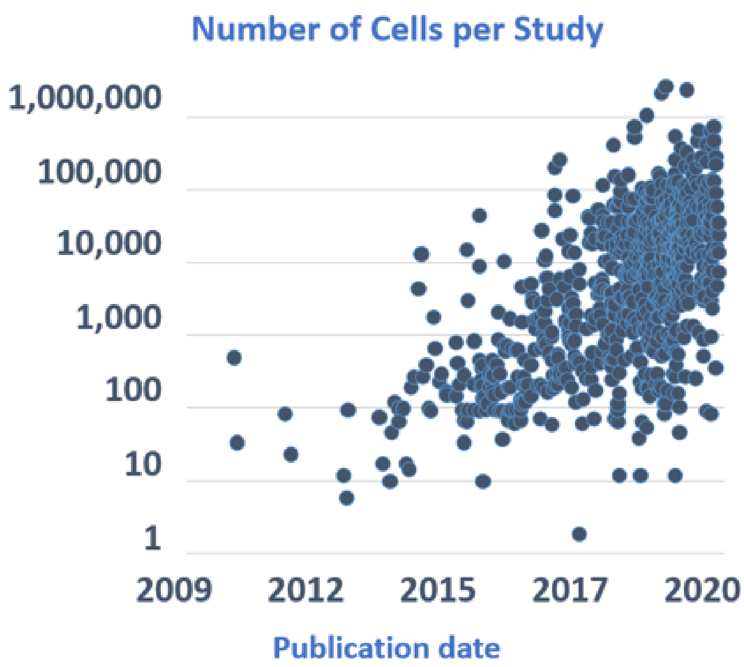

Modalities of Data

- **RNA**
- **DNA (CNV)**
- **DNA (SNV)**
- **DNA (immune repertoire)**
- **Chromatin Accessibility**
- **Surface Protein**
- **Spatial Localization**
- **Etc.**

Also contributing to the growth of data that needs to be processed is the use of various artificial intelligence (AI) tools such as machine learning (ML). This includes the use of AI models running on the network edge and in digital device endpoints. The more the data, the better the AI training model can classify the data.

Massive data sets yield the greatest value when making real-time decisions such as in fraud detection, or after long-running jobs like single-cell sequencing to reveal a COVID-19 variant. Both cases require rapid analysis and critical decision making is based upon this data. That is driving the growth of processing, memory and networking capabilities at the network edge to support low latency requirements as well as the growth of capabilities in data centers for more complex processing tasks.

[1] Data modalities in bioinformatics refers to the type of behavior, expression or way of life that belongs to a particular person or group of people.

New more efficient computer architectures are being developed to support the use of this tsunami of digital data. In particular, we need to move from a von Neumann computer architecture that involved lots of data movement from storage, takes too long and isn't responsive enough for these applications. A memory-centric model, with significant amounts of working content stored directly in memory, provides a more efficient and effective computing model to meet modern processing requirements.

Further improvements in system performance are enabled by protocols that allow larger amounts of shared memory over a fast interconnect and networking protocol. Some big data applications today are running out of memory needed to provide fast data analysis and more and more applications will be processing data directly from memory as the number of these fast big data applications proliferate.

Shared memory allows running more processes in memory with fewer accesses to digital storage, enabling faster operations. Protocols such as Compute Express Link (CXL) are starting to be deployed in devices and these will be part of memory systems that will enable pools of memory and composable infrastructure.

CXL

The concept of shared memory is changing the way working memory is organized and how it is used. Memory directly connected to a processor may be called low latency near memory and involves DDR, LPDDR, GDDR, HBM and OMI interconnects. Larger memory pools with higher latency may be called far memory and are enabled by technologies such as CXL, GenZ, CCIX and OpenCAPI. CXL runs on the fast PCIe bus and is supported by a strong industry coalition. CXL enabled memory devices are starting to appear with full CXL-based systems available starting in 2022.

CXL is an open architecture based upon industry standards that allows dynamically multiplexing input/output (I/O), cache and regular memory. Using CXL fabrics, servers can efficiently share memory among CPUs, GPUs, DPUs and other accelerators. CXL supports tiers of memory allowing trade-off of cost and performance. The figure below shows some of the data flows that CXL enables.

CXL Data Flows

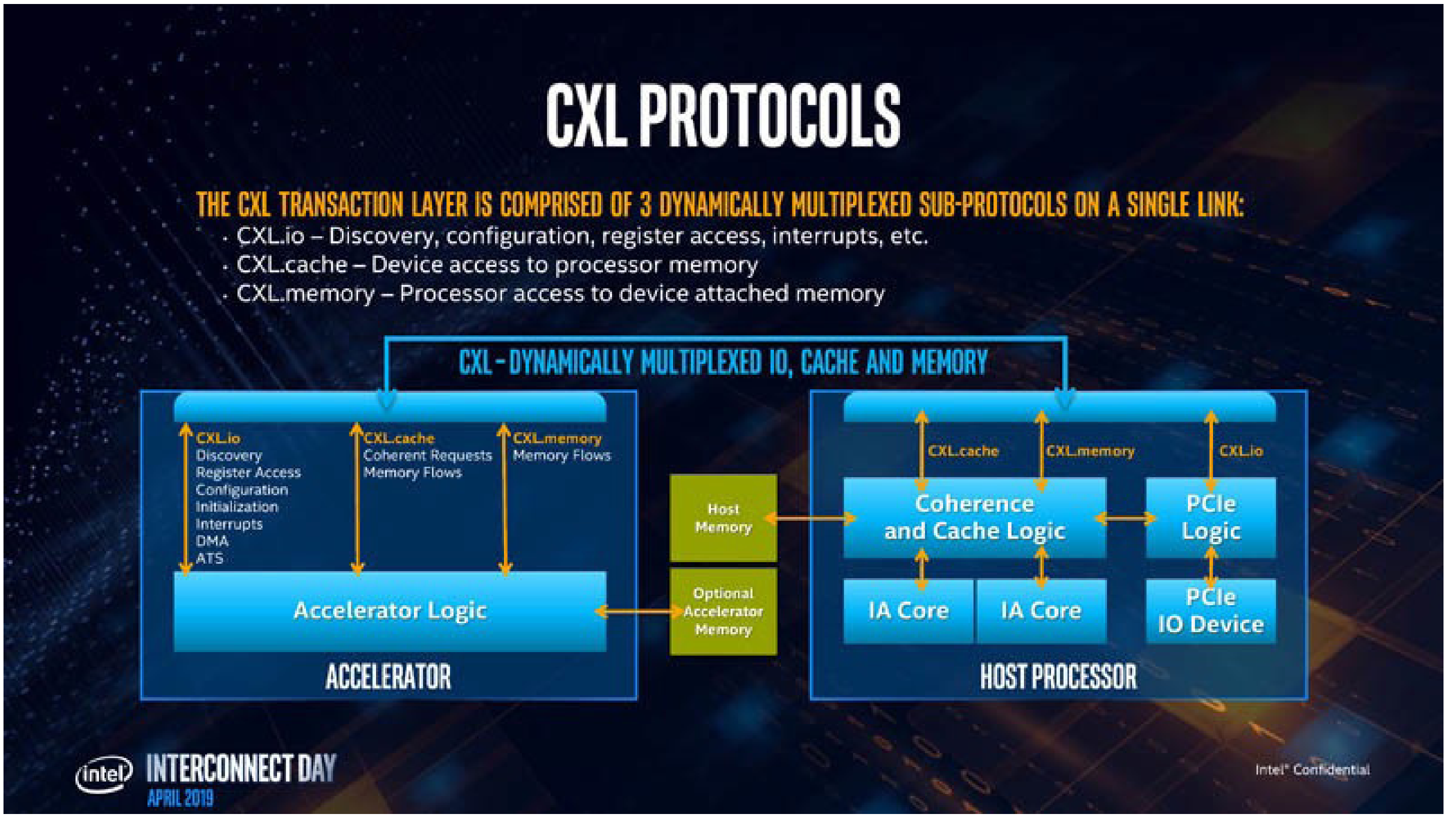

CXL enables working with memory having various performance characteristics (unlike the near memory DDR interconnect), this allows using traditional DRAM as well as other memory technologies, such as Intel's Optane memory (which sells for less than DRAM and allows larger memory pools at a lower cost). The increased use of persistent memory technologies (PMEM) such as Optane, is enabling memory centric computing with fewer data swaps with primary storage. Coughlin Associates projects that by about 2028 persistent memory technologies (such as Optane and MRAM) will equal DRAM in petabytes shipped[2].

CXL will transform the use of memory, similar to the way NVMe is transforming storage. It will enable sharing heterogenous processors and memory in a CXL fabric and running increasing numbers of big data applications entirely within memory, supporting near real-time processing. CXL enables pooling petabytes of memory in a CXL switched fabric by 2023. Creating such dense memory pools will require more compact memory form factors and today's common M.2 and U.2 form factors will migrate to E1.S and E3 form factors for future dense memory platforms.

The figure below shows how a 2.3PB memory system can be built with new memory form factors and CXL. This system uses external memory arrays with 48 X 4TB E3 modules, servers with 20TB of memory using E1 and E3 modules and a 3 rack big data analytics cluster tied together with a CXL fabric.

---

[2] Emerging Memories Find Their Direction, Tom Coughlin and Jim Handy, Coughlin Associates, 2020.

Petascale Memory Configuration with a CXL Fabric

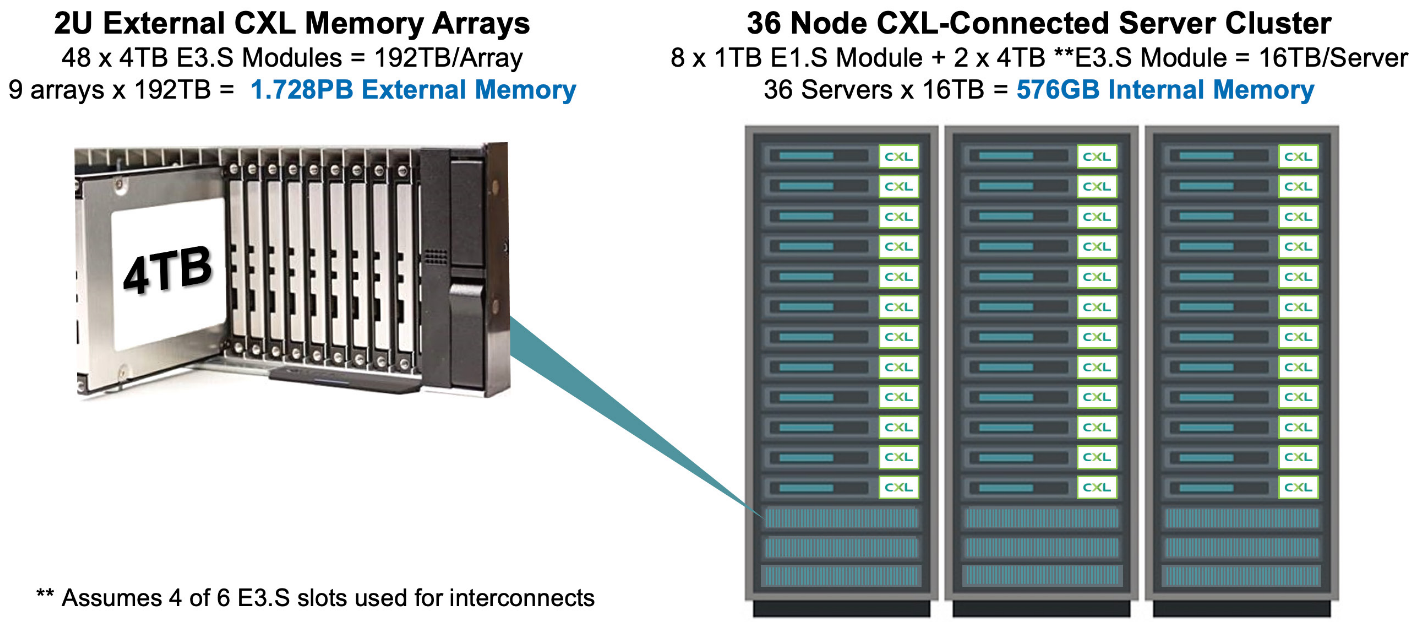

Memory Fabrics

With the rise of vast pools of memory, many concepts common in storage systems will be applied to memory. Sharing of these memory pools through memory fabrics such as CXL will become common and these pools will include DRAM, Optane and other non-volatile memory technologies. In addition to many different types of memories, specialized computational accelerators will be located within the memory pools, providing accelerated and off-loaded data processing.

Availability of these hardware advances, combined with memory fabric management software, will enable fast efficient processing of data. This will enable new applications and over time more and more applications will likely move from traditional Von Neumann computer architectures to memory-centric computing, using technologies such as CXL.

MemVerge’s Memory Machine

The figure below shows MemVerge’s view of the memory infrastructure landscape when CXL is available. Heterogeneous processors and memory alongside fabric software that virtualizes the memory and provides a suite of in-memory data services to provision the capacity, performance, availability, security, and mobility of the data in memory.

Memory Infrastructure with CXL

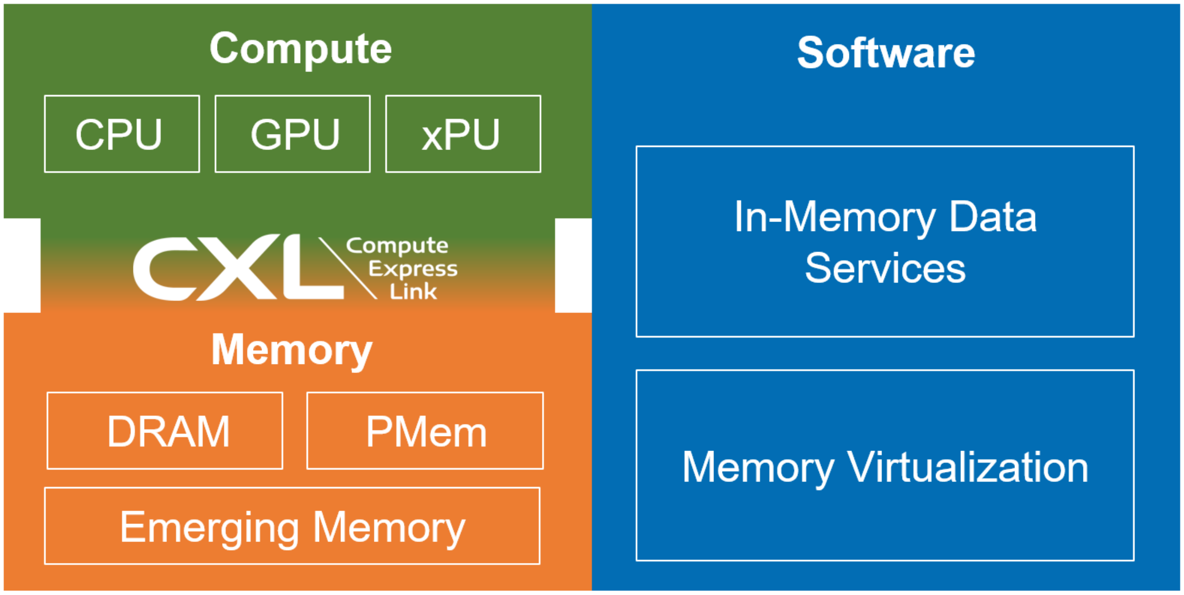

The figure below shows MemVerge's view of the evolution and future of big memory accelerating from today's memory systems to CXL-based memory systems in 2022.

Evolution of Memory Infrastructure

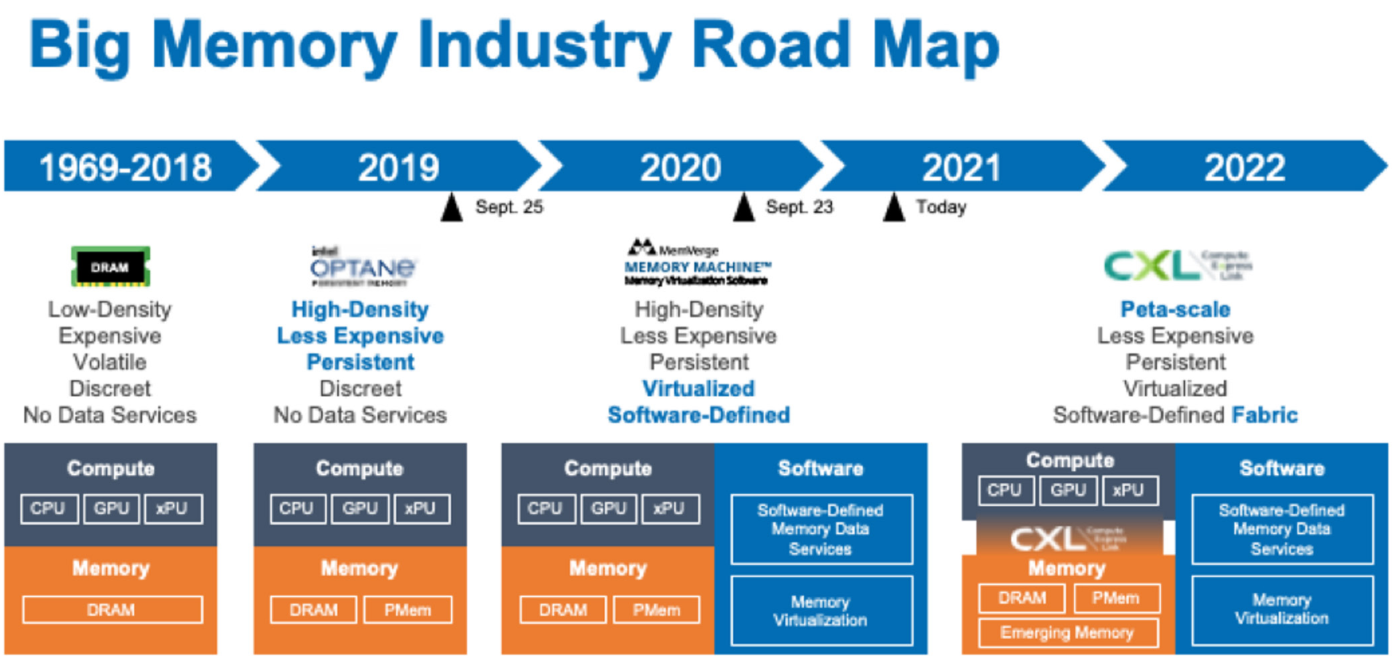

To handle all this communication with big memory, memory fabric management software will be needed, such as MemVerge's Persistent Virtualized Software-Defined Memory (which they call the Memory Machine). Memory fabric management can abstract applications from the complexity of provisioning memory capacity, performance, availability, mobility and security.

This memory fabric management allows provisioning to applications depending upon demand and tuning the available memory using dynamic provisioning and optimization to match the required memory for the application in the most cost-effective manner. MemVerge's management software also allows storage-like operations such as snapshots, tiering and fast recovery from in-memory application crashes.

However, there is a major difference between snapshots to storage and snapshots to memory. Snapshot and recovery of 100s of gigabytes to terabytes of data to and from storage takes minutes to hours. In-memory snapshots are needed in for very large data sets because terabytes of data can be loaded, saved, replicated, and recovered in a few seconds.

Existing applications can run safely and without changes on MemVerge's management software. MemVerge's memory fabric management software abstracts applications from changes needed to support new generations of processors, memory and interconnects. Today, many big data applications can take advantage of MemVerge's data management software. Tomorrow there will be even more.

Case Studies

Let's look at a couple of case studies on the use of MemVerge's big memory solutions for applications in animation and VFX as well as in genomic sequencing.

Much TV and movie content involved using state-of-the-art computer-generated imagery. In order to create this content IT organizations supporting modern studios often deal with fragile applications and plug-ins that crash frequently. The IT pros at Chapeau Studio in Los Angeles, CA say they have at most 30 seconds to recover from a crash before an artist is "out of the zone." However, recovering from a crash usually involves restoring data from storage and takes minutes to hours depending upon when the job was last saved.

A MemVerge managed in-memory solution consisting of Intel's Optane Persistent Memory (PMem) and MemVerge's Memory Machine enabled an in-memory animation application rapid crash recovery system. The MemVerge software enables ZeroIO in-memory snapshots from DRAM to PMEM. Recovering from these snapshots happens at memory speeds and recovery of hundreds of gigabytes of data is accomplished in a few seconds, enabling artists to stay "in the zone."

Genomic processing involves matrix-heavy computations and with conventional approaches involves data written on and read from digital primary storage. Writing and reading from storage is the biggest delay and just one genomic analysis can consume days or weeks.

With MemVerge's Memory Machine software, data is loaded from and written to a pool of DRAM and lower cost PMEM, eliminating the slower steps of writing and reading from storage. Analytical Biosciences (based in the US and China) used the Big Memory solution to increase load times by 800X and 60% faster job throughput across all their genomic process steps as shown in the figure below.

Task Completion Time (Seconds)

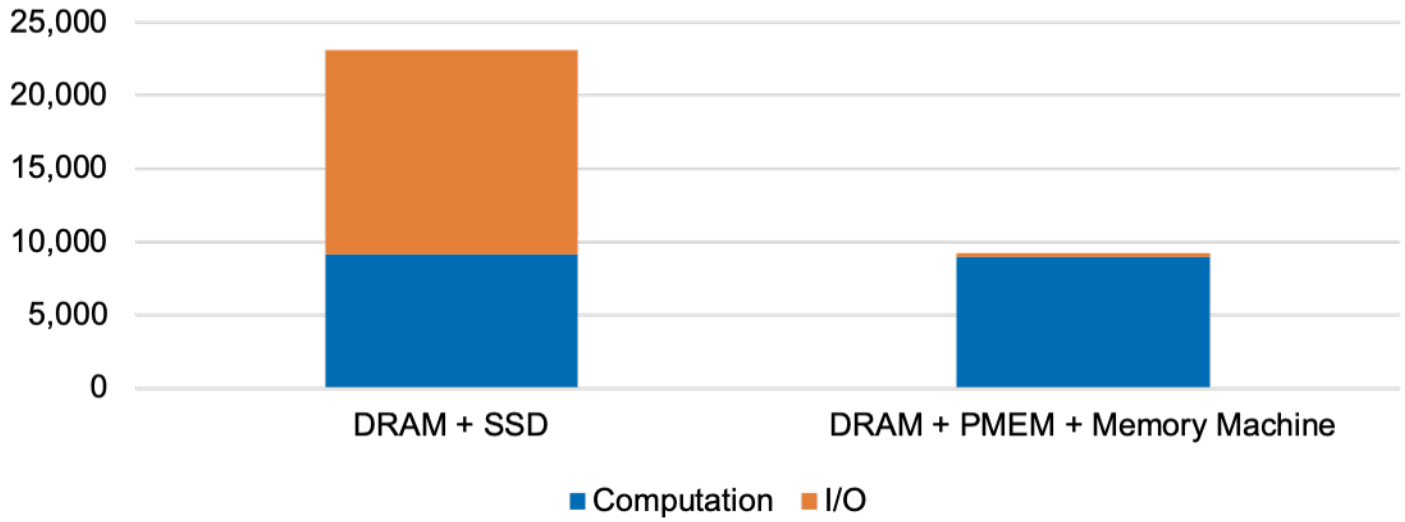

Summary

Digital data is increasing due to the growing use of sensors, monitors, cameras and other connected and smart devices. This growing big data is being processed with AI tools running at endpoints, the network edge and in big data centers and much of this data requires rapid analysis and interpretation to be useful.

CXL will enable a transformation for memory systems. CXL allows a richer heterogenous memory and processor environment. Applications run in faster near or far memory available through CXL are much quicker than those that require frequent data access in digital storage.

MemVerge has developed a memory fabric management software that it called the Memory Machine. This software enables virtualization, snapshots and many other functions with memory that are common in digital storage today. Using MemVerge's Memory Machine, studios have experienced faster crash recoveries in their animation and VFX workflows, allowing artists to "stay in the zone." Genomic analysis companies can improve their overall performance by 60% by avoiding reading and writing data on storage.

With future memory systems using CXL and new memory device form factors, petabyte-scale memory systems will be possible, powering an increasing number of applications that run in-memory with infrequent access of digital storage. These vast memory pools can be efficiently allocated and managed by memory management software, such as MemVerge's Memory Machine, to provide fast and cost-effective big data analysis solutions.

About the Author

Tom Coughlin, President, Coughlin Associates is a digital storage analyst as well as a business and technology consultant. He has over 40 years in the data storage industry with engineering and management positions at several companies.

Dr. Coughlin has many publications and six patents to his credit. Tom is also the author of Digital Storage in Consumer Electronics: The Essential Guide, which is now in its second edition with Springer. Coughlin Associates provides market and technology analysis as well as Data Storage Technical and Business Consulting services. Tom publishes the Digital Storage Technology Newsletter, the Media and Entertainment Storage Report, the Emerging Non-Volatile Memory Report and other industry reports. Tom is also a regular contributor on digital storage for Forbes.com and other blogs.

Tom is active with SMPTE, SNIA and the IEEE, (he is Past Director for IEEE Region 6, Past President of IEEE USA, Past Chair of the IEEE New Initiatives and Public Visibility Committees and active in the Consumer Electronics Society) and other professional organizations. Tom is the founder and organizer of the Storage Visions Conference (www.storagevisions.com) as well as the Creative Storage Conference (www.creativestorage.org). He was the general chairman of the annual Flash Memory Summit for 10 years. He is a Fellow of the IEEE and a board member of the Consultants Network of Silicon Valley (CNSV). For more information on Tom Coughlin and his publications and activities go to www.tomcoughlin.com.