# Addressing the Big-Data Bottleneck through Emerging In-Storage Computing Technologies

James Annis[1], Thomas Junk[1], Eric Neilsen[1], Kevin Skadron[2], Steven Timm[1], Michael Wang*[1], and Yujun Wu[1]

[1]Fermi National Accelerator Laboratory, Batavia, IL 60510, USA
[2]Center for Research in Intelligent Storage and Processing in Memory, University of Virginia, Charlottesville, Virginia 22904, USA

August 31, 2020

## 1 Introduction

For over a decade, the semiconductor industry has dealt with the CPU clock speed limit brought about by the demise of Dennard scaling [1] by exploiting parallelism. This is the main reason much of the effort in HEP has been devoted to taking advantage of heterogeneous architectures and refactoring code bases for multithreaded operation. Nevertheless, recent years have seen continued improvement in processor performance. Things have not been quiet either on the storage front, where drive capacities have steadily increased in size, recently hitting the 100TB mark for a single device. While both processor and storage technologies continued to evolve, the technology associated with the data paths linking them has fallen woefully behind. The antiquated data paths will pose a serious bottleneck limiting the ability of the processing elements to efficiently access the data on the storage systems. This is especially urgent since HEP experiments and cosmological surveys in the coming decade will be acquiring unprecedented amounts of data. It is urgent for the community to begin paying closer attention now to find ways of addressing this issue, so that the physics potential of the upcoming experiments and surveys will not be negatively affected.

## 2 Big-Data Problems in HEP and Astro in the Coming Decade

In this section, we use the Large Synoptic Survey Telescope (LSST) [2] and the Deep Underground Neutrino Experiment (DUNE) [3] as examples of big-data efforts in the coming decade. We stress, however, that the ideas and opinions raised in this LOI, in no way, represent those of these collaborations.

The LSST survey will make use of the world's largest digital camera, with 3024 megapixels, to produce 12GB uncompressed images every $15^s$. One of its major pipelines, which involves computationally intensive steps to detect transient objects like supernovae or other cataclysmic events like optical counterparts to gravitational wave sources, needs to be performed in near real-time so time-critical information can be sent to an alert network in the $2^m$ allocated by LSST's requirement. Relying on traditional computing architectures, however, will result in a large fraction of the total time available for processing wasted on file transfers between storage or memory and the CPUs or GPUs.

A large LArTPC-based detector similar to the Deep Underground Neutrino Experiment (DUNE) faces a similar problem where, in the event of a supernova burst (SNB) candidate, $100^s$ worth of full-stream LArTPC data totaling 120TB/module, spanning the time envelope of the SNB, needs to be buffered in the upstream DAQ. The several hours required to transfer this amount of data from the underground caverns to the back-end DAQ on the surface for further refinement can negatively impact the ability to publish reliable results quickly on the SuperNova Early Warning System.

## 3 In-Storage Computing as a Possible Solution

Fortunately, HEP and Astro are not the only ones facing the big-data crunch. New technologies are appearing and continue to be defined and developed to address precisely the big-data problems confronting us. Essentially, the idea is

*corresponding author: mwang@fnal.gov

that if the data cannot be brought to the processors fast enough, why not bring the processors to the data instead? This is certainly not a new idea, but until recently, there have been no practical solutions. One promising solution is offered by the emerging technology of *computational storage* [4], which includes *computational storage devices* (CSDs) like Samsung's SmartSSD [5] which embeds processing elements (PEs) into a solid state drive (SSD), or *computational storage processors* (CSPs) that put ARM cores or FPGA accelerators directly on the NVMe storage bus, in close proximity to SSDs [6, 7]. By allowing in-situ data reduction, this approach minimizes the cycles wasted in transferring massive amounts of data from storage to the processors across the inefficient data channels linking them. This will make real-time constraints easier to meet and allow new algorithms to be implemented that enhance capabilities.

# 4 Conclusion

The emerging intelligent storage technologies described above and related efforts [8, 9] could hold the key to addressing the data bottleneck facing upcoming HEP experiments and cosmological surveys. Aside from the two challenging examples cited above, AI and offline applications will also benefit from such technologies. Emulators, prototypes, and even commercial products are becoming available now for evaluation. We believe it is important for our community to be aware of this and to take a more concerted effort in evaluating their feasibility for our needs, since this will give us a better idea of what is technically possible and help us chart the future.

# References


[1] R. H. Dennard *et al. Design of ion-implanted MOSFET's with very small physical dimensions*. IEEE Journal of Solid-State Circuits, vol. 9, no. 5, pp. 256-268, Oct. 1974. doi: 10.1109/JSSC.1974.1050511.

[2] Željko Ivezić *et al. LSST: From science drivers to reference design and anticipated data products*. The Astrophysical Journal, vol. 873, no. 2, p. 111, mar 2019. [Online]. Available: https://doi.org/10.3847%2F1538-4357%2Fab042c

[3] B. Abi *et al. Volume I. Introduction to DUNE*. Journal of Instrumentation, vol. 15, no. 08, pp. T08 008T08 008, Aug 2020. [Online]. Available: https://doi.org/10.1088%2F1748-0221%2F15%2F08%2Ft08008

[4] Storage Networking Industry Association: Computational Storage. https://www.snia.org/computational

[5] Samsung SmartSSD, Computational Storage Drive. https://samsungsemiconductor-us.com/smartssd/

[6] Eideticom. https://www.eideticom.com/

[7] BittWare: Computational Storage. https://www.bittware.com/fpga/storage/

[8] CRISP: Center for Research in Intelligent Storage and Processing in Memory, University of Virginia, Charlottesville, Virginia 22904. https://crisp.engineering.virginia.edu/

[9] Z. Ruan, T. He, and J. Cong. *INSIDER: Designing in-storage computing system for emerging high-performance drive*. USENIX ATC19: Proceedings of the 2019 USENIX Conference on Usenix Annual Technical Conference, July 2019, Pages 379-394.