Award No: DE-SC0016313
Title: UNITY: Unified Memory and Storage Space
Organization: Georgia Tech Research Corporation (Georgia Tech, GT)
PI: Ada Gavrilovska

Personnel: Ada Gavrilovska (GT PI), Greg Eisenhauer, Pradeep Fernando, Thaleia Doudali, Ranjan Sarpangala Venkatesh, James Tyson

Summary: UNITY is a 36-month project focused on providing design and evaluate a new distributed storage paradigm that unifies the traditionally distinct application views of memory- and file-based data storage into a single scalable and resilient environment. The project is a collaboration among Oak Ridge National Laboratory (ORNL, Lead institution), Los Alamos National Labs (LANL), and Georgia Tech (GT).

The main contributions of the GT team have been around development of low-level systems software for best leveraging the capabilities of new types of persistent memory technologies, and for development of methods for intelligent data management across memory/storage substrates with heterogeneous components. GT contributed the Phoenix library for optimized checkpoint/restart for HPC I/O for systems with non-volatile memory (NVM), the NVStream library for NVM-specialized streaming I/O for HPC workflows, the CoMerge, Mnemo and Kleio solutions for intelligent data management on NVM-based systems. These contributions result in significant improvements in both application performance and system efficiency.

Publications, Presentations, Products:

Year 3:

1. Ada Gavrilovska, "Real PMEM in Research: How-Tos and How To?-s", Programming Persistent Memory in Real Life (PIRL), San Diego, CA, Jul. 2019.
2. Ranjan Sarpangala Venkatesh, Till Smejkal, Ada Gavrilovska, Dejan Milojicic, "Fast Container Scaling with Multiple Virtual Address Spaces", ACM Symposium on Memory Systems (MEMSYS'19), Washington, DC, Sep. 2019.
3. Thaleia Dimitra Doudali, Sergey Blagodurov, Abhinav Vishnu, Sudhanva Gurumurthi, Ada Gavrilovska, "Kleio: A Hybrid Memory Page Scheduler with Machine Intelligence", Proceedings of the 28th ACM Symposium on High Performance Parallel and Distributed Computing (HPDC'19), Phoenix, AZ, Jun. 2019. (best paper finalist, best poster based on research paper award)
4. Thaleia Doudali, Ada Gavrilovska, "Mnemo: Boosting Memory Cost Efficiency in Hybrid Memory Systems", 5th Workshop on High Performance Big Data and Cloud Computing (HPBDC'19), in conjunction with the International Parallel and Distributed Processing Symposim (IPDPSW), Rio de Janeiro, Brazil, May 2019.
5. Thaleia Doudali, Ada Gavrilovska, Sergey Blagodurov, Abhinav Vishnu, Sudhanva Gurumurthi, "Page Scheduling on Hybrid Memory Systems", Poster at European Conference on Systems Research (EuroSys'19), Dresden, Germany, Mar. 2019.
6. Ada Gavrilovska, "Memory Fabric: Systems Software for Seamless Scaling Across Complex Memory Hierarchies", University of New Mexico and Sandia National Labs

seminar, Feb. 2019.

7. Sudarsun Kannan, Nitish Bhat, Ada Gavrilovska, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau, "Redesigning LSMs for Non-Volatile Memory with NoveLSM", Proceedings of the USENIX Annual Technical Conference (USENIX ATC'18), Boston, MA, Jul. 2018.
8. Pradeep Fernando, Ada Gavrilovska, Sudarsun Kannan, Greg Eisenhauer, "NVStream: Accelerating HPC Workflows with NVRAM-based Streaming Persistence", Proceedings of the 27th ACM Symposium on High Performance Parallel and Distributed Computing (HPDC'18), Tempe, AZ, Jun. 2018.

Year 2:

9. Pradeep Fernando, Ada Gavrilovska, Sudarsun Kannan, Greg Eisenhauer, "NVStream: Accelerating HPC Workflows with NVRAM-based Streaming Persistence", Proceedings of the 27th ACM Symposium on High Performance Parallel and Distributed Computing (HPDC'18), Tempe, AZ, Jun. 2018.
10. Alexander Merritt, Ada Gavrilovska, Yuan Chen, Dejan Milojicic, "Concurrent Log-Structured Memory for Many-Core Key-Value Stores", Proceedings of the VLDB Endowment (PVLDB), vol. 11, no. 4, Dec. 2017. Part of the 44th International Conference on Very Large Data Bases (VLDB'18), Rio de Janeiro, Aug. 2018.
11. Thaleia Doudali, Ada Gavrilovska, "Data Management in Heterogeneous Memory Systems", Poster at Salishan High Performance Computing Conference, Salishan, OR, Apr. 2018.
12. Thaleia Doudali, Ada Gavrilovska, "Mnemo: Boosting Memory Cost Efficiency in Hybrid Memory Clouds", Poster at the ACM Symposium on Cloud Computing (SOCC'18), Carlsbad, CA, Oct. 2018.
13. Sudarsun Kannan, Ada Gavrilovska, Vishal Gupta, Karsten Schwan, "Heterogeneous Memory Management in Datacenters", 9th Annual Non-Volatile Memories Workshop (NVMW'18), San Diego, CA, Mar. 2018.
14. Thaleia Doudali, Ada Gavrilovska, "CoMerge: Toward Efficient Data Placement in Shared Heterogeneous Memory Servers", ACM International Symposium on Memory Systems (MEMSYS'17), Washington, DC, Oct. 2017.
15. Ada Gavrilovska, "Memory Fabric: Systems Software for Seamless Scaling Across Complex Memory Hierarchies", Facebook Infrastructure Team Seminar, Invited Talk, Jul. 2017.
16. Sudarsun Kannan, Ada Gavrilovska, Karsten Schwan, "HeteroOS – OS Design for Heterogeneous Memory Management in Datacenters", 44th International Symposium on Computer Architecture (ISCA'17), Toronto, CA, Jun. 2017. (16.7%)
17. Terry Jones, Michael Lang, Ada Gavrilovska, Michael J. Brim, Geoffroy Vallee, Benjamin Mayer, Aaron Welch, Tonglin Li, Latchesar Ionkov, Douglas Otstott, Greg Eisenhauer, Thaleia Doudali, Pradeep Fernando, "UNITY: Unified Memory and File Space", 7th International Workshop on Runtime and Operating Systems for Supercomputers (ROSS'17), in conjunction with HPDC, Washington, DC, June, 2017. *
18. Alexis Champsaur, Jai Dayal, Matthew Wolf, Ada Gavrilovska, Greg Eisenhauer, Jay Lofstead, Patrick Widener, "SmartBlock: An Approach to Standardizing In Situ Workflow Components", 2nd Workshop on Emerging Parallel and Distributed Runtime

Systems and Middleware (IPDRM'17), in conjunction with IPDPS, Orlando, FL, May, 2017.

Year 1:

19. Ada Gavrilovska, "Accelerating HPC Services with Heterogeneous Memory", Invited Seminar, Oak Ridge National Laboratory, May 2017.
20. Ada Gavrilovska, "Re-architecting Systems Software for Memory Heterogeneity and Scale", Salishan Conference on High Speed Computing, Invited Speaker, Gleneden Beach, OR, Apr. 2016.
21. Pradeep Fernando, Sudarsun Kannan, Ada Gavrilovska, Karsten Schwan, "Memory Speed HPC I/O with NVM", Poster and presentation at the 8th Annual Non-Volatile Memories Workshop (NVMW'17), San Diego, CA, Mar. 2017.
22. Pradeep Fernando, Sudarsun Kannan, Ada Gavrilovska, Karsten Schwan, "Phoenix: Memory-Speed HPC I/O with NVM", 23rd Annual International Conference on High Performance Computing, Data and Analytics (HiPC'16), Hyderabad, India, Dec. 2016.
23. Ada Gavrilovska, "Implication of Heterogeneous Memories in Next Generation Server Systems", Keynote Address, 25th International Symposium on High Performance Parallel and Distributed Computing (HPDC'16), Kyoto, Japan, Jun. 2016.

Software:

24. **NVStream:** https://github.com/pradeepfn/nvs
25. **Mnemo:** https://github.com/Thaleia-DimitraDoudali/mnemo
26. **CoMerge and datasets:** https://github.com/Thaleia-DimitraDoudali/CoMerge
27. **Phoenix:** https://github.com/pradeepfn/nvs/releases/tag/hipc-paper
28. **NoveLSM**: https://gitlab.com/sudarsunkannan/leveldb-nvm
29. **Nibble**: https:// github.com/GTkernel/nibble-lsm