Project Title: In Situ Indexing and Query Processing of AMR Data
DOE Award Number: SN10016 DE-SC0012389
Project PI: Nagiza F. Samatova

Tasks that have been accomplished (detailed/technical):
We have developed the first, holistic, Adaptive Mesh Refinement (AMR) capable, *in situ* query framework to ensure the tractability of the current- and next-generation AMR simulation analysis on the path to extreme scale. We have addressed key issues pertaining to *in situ analysis*, memory hierarchy, generalized parallel indexing, index-compression, and real-time access pattern detection. We have published the following research papers in high profile peer reviewed computer science conferences:

1. ***Parallel In Situ Detection of Connected Components in Adaptive Mesh Refinement Data*** [1]:
   We proposed the first connected component detection methodology for structured AMR data that is applicable in a distributed and *in situ* context. Our key strategy is to use a parallel, *in situ*, AMR-aware communication method to synchronize component connectivity across the AMR structure distributed over many processes. By using a hierarchical process grouping conforming to the AMR refinement level structure, we are able to compute global connectivity over arbitrary distributions of AMR data without expensive all-to-all communication.

2. ***AMR-aware In Situ Indexing and Scalable Querying*** [2]:
   We proposed the first work to support scalable AMR-aware analysis. We proposed an AMR-aware hybrid index for supporting two common forms (i.e., spatial and value-based query selections) in query-driven analytics. To sustainably support future-scale analyses, we designed an in-situ (run-time) index building strategy with minimized performance impact to the co-located simulation. Moreover, we presented a parallel post-processing query method with an adaptive workload-balanced strategy.

3. ***Pattern-driven dynamic data layout optimization*** [3]:
   We proposed a dynamic data reorganization framework that performs dynamic data access pattern tracing and identification functions, flexible layout reorganization strategies, efficient storage of partial replicas to support multiple read patterns, and redirection of read accesses to a favorable layout at runtime. This framework can perform the optimization over AMR data.

4. ***In situ Storage Layout Optimization for AMR Spatio-temporal Read Accesses*** [4]:
   To enable efficient AMR read accesses, we developed an *in situ* AMR data layout optimization framework. Our framework automatically selects from a set of candidate layouts based on a performance model, and reorganizes the data before writing to storage.

5. ***AMRZone: A Runtime AMR Data Sharing Framework for Scientific Applications*** [5]:
   We proposed a framework, AMRZone, which is specifically designed to facilitate runtime AMR data sharing across multiple scientific applications. The framework consists of three major contributions. First, it presents an architecture with centralized metadata management to facilitate runtime AMR data management. Second, it achieves runtime balanced workload distribution at the staging space, by

adopting an AMR boxes-based runtime workload assignment policy. Third, it builds a polytree-based online spatial index to facilitate spatial constrained AMR data retrieval.

6. ***Exploring Memory Hierarchy and Network Topology for Runtime AMR Data Sharing Across Scientific Applications*** [6]:
   In this paper, we proposed methods of how to enable an in-memory AMR data staging framework [5] to utilize SSDs on supercomputer as an overflow space. We also study how to improve data access performance over the staging space via data prefetching and network-topology-aware data distribution.

Reference:

1. Zou, Xiaocheng, et al. "Parallel in situ detection of connected components in adaptive mesh refinement data." *Cluster, Cloud and Grid Computing (CCGrid), 2015 15th IEEE/ACM International Symposium on*. IEEE, 2015.
2. Zou, Xiaocheng, et al. "AMR-aware in situ indexing and scalable querying." *Proceedings of the 24th High Performance Computing Symposium*. Society for Computer Simulation International, 2016.
3. Tang, Houjun, et al. "Usage pattern-driven dynamic data layout reorganization." *Cluster, Cloud and Grid Computing (CCGrid), 2016 16th IEEE/ACM International Symposium on*. IEEE, 2016.
4. Tang, Houjun, et al. "In Situ Storage Layout Optimization for AMR Spatio-temporal Read Accesses." *Parallel Processing (ICPP), 2016 45th International Conference on*. IEEE, 2016.
5. Zhang, Wenzhao, et al. "AMRZone: A Runtime AMR Data Sharing Framework for Scientific Applications." *Cluster, Cloud and Grid Computing (CCGrid), 2016 16th IEEE/ACM International Symposium on*. IEEE, 2016.
6. Zhang, Wenzhao, et al. "Exploring memory hierarchy and network topology for runtime AMR data sharing across scientific applications." *Big Data (Big Data), 2016 IEEE International Conference on*. IEEE, 2016.
7. Stephen L Cornford, Daniel F Martin, Daniel T Graves, Douglas F Ranken, Anne M Le Brocq, Rupert M Gladstone, Antony J Payne, Esmond G Ng, and William H Lipscomb. Adaptive mesh, finite volume modeling of marine ice sheets. Journal of Computational Physics, 232(1):529–549, 2013.
8. M. Adams, P. Colella, D. T. Graves, J.N. Johnson, N.D. Keen, T. J. Ligocki, D. F. Martin, P.W. McCorquodale, D. Modiano, P.O. Schwartz, T.D. Sternberg, and B. Van Straalen. Chombo software package for amr applications-design document, 2000.
9. Atchley, Scott, et al. "The common communication interface (CCI)." *High Performance Interconnects (HOTI), 2011 IEEE 19th Annual Symposium on*. IEEE, 2011.