

Sublinear Algorithms for In-situ and In-transit Data Analysis at the Extreme-Scale

Janine Bennett, C. Seshadhri, Ali Pinar; Sandia National Laboratories* Livermore, CA
 David Thompson; Kitware, Inc., Carrboro, North Carolina

Abstract

Post-Moore's law scaling is creating a disruptive shift in simulation workflows, as saving the entirety of raw data to persistent storage becomes expensive. We are moving away from a post-process centric data analysis paradigm towards a concurrent analysis framework, in which raw simulation data is processed as it is computed. Algorithms must adapt to machines with extreme concurrency, low communication bandwidth, and high memory latency, while operating within the time constraints prescribed by the simulation. Furthermore, input parameters are often data dependent and cannot always be prescribed. The study of sublinear algorithms is a recent development in theoretical computer science and discrete mathematics that has significant potential to provide solutions for these challenges. The approaches of sublinear algorithms address the fundamental mathematical problem of understanding global features of a data set using limited resources. These theoretical ideas align with practical challenges of in-situ and in-transit computation where vast amounts of data must be processed under severe communication and memory constraints.

Introduction: Steady improvements in computing resources enable ever more enhanced simulations, but I/O constraints are impeding their impact. Despite increases in temporal resolution, the gap between time steps saved to disk keeps increasing. This compromise in fidelity makes it impossible to track features with timescales smaller than that of I/O frequency. This discrepancy will become more pressing on future architectures as increases in computational power significantly outpace I/O capabilities. This motivates a fundamental shift away from a post-process centric data analyses. Concurrent analysis frameworks are a promising direction wherein raw simulation output is processed as it is computed, decoupling the analysis from I/O. Both in-situ [YWG⁺10, BWM11, FMT⁺11] and in-transit [VHP11, AEW⁺11, BAB⁺12] processing are based on performing analyses as the simulation is running, storing only the results, which are several orders of magnitude smaller than the raw data. This mitigates the effects of limited disk bandwidth and capacity. Operations sharing primary resources of the simulation are considered in-situ, while in-transit processing involves asynchronous data transfers to secondary resources.

These workflows pose unique challenges as algorithms must be re-designed to operate within tight memory, communication, and I/O constraints. *Sublinear algorithms* may provide a solution. This field is a recent development in theoretical computer science and discrete mathematics, which addresses the mathematical problem of understanding global features of a data set using limited resources. These theoretical ideas are directly aligned with practical challenges of in-situ and in-transit computation and show great promise in addressing these issues.

*Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sublinear algorithms: Sublinear methods [Rub06, Ron09] are designed around the following question: given access to a function f over a massive discrete domain, how does one choose a tiny fraction of the domain space that provides information about detailed properties of f ? Often enough, to determine important features of an input, one does not need to actually look at the entire input. The field of sublinear algorithms makes precise the circumstances when this is possible and provides a host of algorithmic techniques to perform this task. The samples are chosen through randomization but often have no closed-form description and are much more intricate than statistical random sampling based on standard distributions. Sublinear algorithms allow for quantification of the error or uncertainty introduced by using only a sample of the data. This confidence measure is necessary for adoption of such techniques by the large-scale scientific computing community, whose scientific results are often used to make high-impact decisions that could have large financial or public policy implications.

From theory to practice: At a high level, any question that can be framed in terms of determining global properties of a large domain is subject to a sublinear analysis. However, because this theory is relatively new and highly mathematical, implementations of sublinear techniques in real applications are practically non-existent. Recent work by our group on graph analysis showcased the immense potential of transferring the advances in theoretical science into a practical realm. We designed sampling-based algorithms to estimate various triadic measures on graphs [SPK13]. The proposed techniques are effective in practice and come with theoretical error/confidence bounds for a given number of samples (e.g., 38K samples guarantee an error $< 10^{-2}$ with 99.9% confidence). This algorithm is faster than standard enumeration methods by many orders of magnitude. It possesses excellent scalability, since the number of samples required for a given accuracy is independent of input size. We have extended our ideas to a Hadoop implementation [KPP⁺13] that enabled analyzing triangle analysis on a graph with 4.6 billion edges (the largest ever reported).

Our aim is to apply these sampling-based techniques to large-scale, physics-based simulations, although the data analysis challenges are significantly different. Our initial investigations in application-independent generation of colormaps show promise. The visualization of large-scale data can be difficult because determining a good colormap requires information about the distribution of values in the data. The size of the data often makes this process computationally expensive. Using sampling methods, we have devised an algorithm to efficiently generate color maps that support the visualization of features of interest on large-scale data.

Another scientific application that shows promise is the aggregation of feature-based statistics [BKL⁺11]. For example, given functions defined over the 3-dimensional grid, combustion scientists are interested in the average size of the connected regions with a value greater than some threshold [MGB⁺11]. A sublinear algorithm would ascertain this value without examining the whole domain by choosing a small subset (using randomization) through complex algorithmic schemes. The output would be approximate, but with a quantified time-error tradeoff. In applications with an underlying turbulent transport, it is important to track features of interest over time as they are advected by the flow. The full analysis is too computationally expensive at the frequency required by the simulation timescales, and the sketches provided by sublinear algorithms might be utilized as a surrogate for a subset of the time steps. Finally, many algorithms running on massive inputs work much faster if prior knowledge of the input is incorporated. This is often encoded in terms of algorithm parameters, which are set by domain experts based on their estimates of global properties of the input. Sublinear algorithms can explore a tiny portion of the domain to provide accurate estimates of these desired global properties.

In summary, sublinear algorithms demonstrate enormous potential in providing solutions to data analysis challenges in extreme-scale scientific computing environments. This revolutionary approach tells us how to perform computations by looking at small portions of the data, and this could lead to huge reductions in required communication. Initial results in transitioning these technologies from theory to practice show promise for current architectures. However, further study is required to understand both their broad applicability and performance on potential future architectures.

References

- [AEW⁺11] Hasan Abbasi, Greg Eisenhauer, Matthew Wolf, Karsten Schwan, and Scott Klasky. Just In Time: Adding Value to The IO Pipelines of High Performance Applications with JITStaging. In *Proc. of 20th International Symposium on High Performance Distributed Computing (HPDC'11)*, June 2011.
- [BAB⁺12] Janine C. Bennett, Hasan Abbasi, Peer-Timo Bremer, Ray Grout, Attila Gyulassy, Tong Jin, Scott Klasky, Hemanth Kolla, Manish Parashar, Valerio Pascucci, Philippe Pebay, David Thompson, Hongfeng Yu, Fan Zhang, and Jacqueline Chen. Combining in-situ and in-transit processing to enable extreme-scale scientific analysis. In Jeffrey Hollingsworth, editor, *SC '12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, Salt Lake Convention Center, Salt Lake City, UT, USA, November 10–16, 2012*, pages 49:1–49:9, pub-IEEE:adr, 2012. IEEE Computer Society Press.
- [BKL⁺11] J. Bennett, V. Krishnamoorthy, S. Liu, R. Grout, E. R. Hawkes, J. H. Chen, J. Shepherd, V. Pascucci, and P.-T. Bremer. Feature-based statistical analysis of combustion simulation data. *IEEE Trans. Vis. Comp. Graph.*, 17(12):1822–1831, 2011.
- [BWM11] Jean-Marie Favre Brad Whitlock and Jeremy S. Meredith. Parallel In Situ Coupling of Simulation with a Fully Featured Visualization System. In *Proc. of 11th Eurographics Symposium on Parallel Graphics and Visualization (EGPGV'11)*, April 2011.
- [FMT⁺11] N. Fabian, K. Moreland, D. Thompson, A.C. Bauer, P. Marion, B. Gevecik, M. Rasquin, and K.E. Jansen. The paraview coprocessing library: A scalable, general purpose in situ visualization library. In *Proc. of IEEE Symposium on Large Data Analysis and Visualization (LDAV)*, pages 89–96, October 2011.
- [KPP⁺13] Tamara G. Kolda, Ali Pinar, Todd Plantenga, C. Seshadhri, and Christine Task. Counting triangles in massive graphs with mapreduce. Technical Report 1301.5887, Arxiv, 2013.
- [MGB⁺11] Ajith Mascarenhas, Ray W. Grout, Peer-Timo Bremer, Evatt R. Hawkes, Valerio Pascucci, and Jacqueline H. Chen. Topological feature extraction for comparison of terascale combustion simulation data. In Valerio Pascucci, Xavier Tricoche, Hans Hagen, and Julien Tierny, editors, *Topological Methods in Data Analysis and Visualization*, Mathematics and Visualization, pages 229–240. Springer Berlin Heidelberg, 2011.
- [Ron09] D. Ron. Algorithmic and analysis techniques in property testing. *Foundations and Trends in Theoretical Computer Science*, 5(2):73–205, 2009.
- [Rub06] Ronitt Rubinfeld. Sublinear time algorithms. *International Conference of Mathematicians*, 2006.
- [SPK13] C. Seshadhri, Ali Pinar, and Tamara G. Kolda. Triadic measures on graphs: The power of wedge sampling. In *Proceedings of the SIAM Conference on Data Mining (SDM)*, 2013.
- [VHP11] V. Vishwanath, M. Hereld, and M.E. Papka. Toward simulation-time data analysis and i/o acceleration on leadership-class systems. In *Proc. of IEEE Symposium on Large Data Analysis and Visualization (LDAV)*, October 2011.
- [YWG⁺10] Hongfeng Yu, Chaoli Wang, Ray W. Grout, Jacqueline H. Chen, and Kwan-Liu Ma. In-situ visualization for large-scale combustion simulations. *IEEE Computer Graphics and Applications*, 30:45–57, 2010.