

Massive-scale Streaming Analytics: Models, Parallelism, & Real-world Applications

Keynote Talk

David A. Bader

School of Computational Science and Engineering
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332 USA

ABSTRACT

Emerging real-world graph problems include: detecting and preventing disease in human populations; revealing community structure in large social networks; and improving the resilience of the electric power grid. Unlike traditional applications in computational science and engineering, solving these social problems at scale often raises new challenges because of the sparsity and lack of locality in the data, the need for research on scalable algorithms and development of frameworks for solving these real-world problems on high performance computers, and for improved models that capture the noise and bias inherent in the torrential data streams. Highlighting this keynote talk, Bader will discuss the opportunities and challenges in massive data-intensive computing for applications in social sciences, physical sciences, and engineering.

Focusing on parallel algorithm design and implementation, Bader formalizes a practical model for graph analysis on streaming data. In this model, a massive graph undergoes changes from an input stream of edge insertions and removals. The model supports concurrent updating of the graph while algorithms execute concurrently on the dynamic data structure. The talk introduces a concept of validity: an algorithm is valid if the output is correct for a graph consisting of the initial graph

with some subset of concurrent changes. Practical examples of this model are given for valid implementations of breadth first search, connected components, PageRank, and triangle counting, all useful graph kernels in real-world applications. This is joint work with E. Jason Riedy and Chunxing Yin.

KEYWORDS

Graph Algorithms; Parallel Computing.

ACM Reference format:

D.A. Bader. 2018. In *Proceedings of 30th ACM Symposium on Parallelism in Algorithms and Architectures, Vienna, Austria, July 2018 (SPAA'18)*, 1 page. <https://doi.org/10.1145/3210377.3210379>

ACKNOWLEDGMENTS

This work is supported in part by the Defense Advanced Research Projects Agency (DARPA) under Contract Number FA8750-17-C-0086. The content of the information in this document does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

SPAA'18, July 16-18, 2018, Vienna, Austria.

© 2018 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-5799-9/18/07.

DOI: <https://doi.org/10.1145/3210377.3210379>