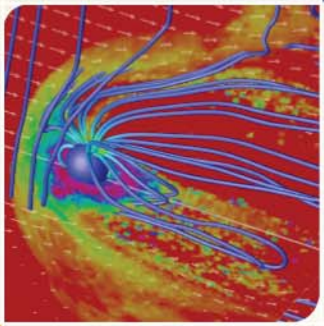

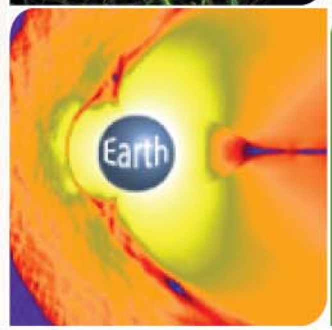

Massive-Scale Streaming Graph Analytics

David A. Bader

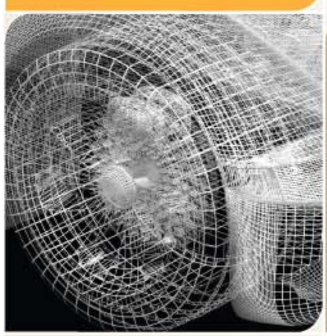

Georgia Tech College of Computing
Computational Science and Engineering

Dr. David A. Bader

- Full Professor, Computational Science and Engineering
- Executive Director for High Performance Computing.
- IEEE Fellow, AAAS Fellow
- interests are at the intersection of high-performance computing and real-world applications, including computational biology and genomics and massive-scale data analytics.
- Over $64M of externally sponsored projects since 2005
  - NSF, NIH, DARPA, DoD, DOE, IBM, NVIDIA, Intel, Cray, Oracle/Sun Microsystems, Microsoft Research, ExxonMobil.
- Steering Committees of the major HPC conferences, IPDPS and HiPC
- Multiple editorial boards in parallel and high performance computing
- Elected chair of IEEE and SIAM committees on HPC
- 215+ publications
- National Science Foundation CAREER Award recipient
- Directed the Sony-Toshiba-IBM Center for the Cell/B.E. Processor
- Founder of the Graph500 List for benchmarking "Big Data" computing platforms
- Recognized as a “RockStar”of High Performance Computing by InsideHPC and as HPCwire's People to Watch in 2012.

Exascale Streaming Data Analytics:
Real-world challenges

All involve analyzing massive streaming complex networks:

- Health care → disease spread, detection and prevention of epidemics/pandemics (e.g. SARS, Avian flu, H1N1 “swine” flu)
- Massive social networks → understanding communities, intentions, population dynamics, pandemic spread, transportation and evacuation
- Intelligence → business analytics, anomaly detection, security, knowledge discovery from massive data sets
- Systems Biology → understanding complex life systems, drug design, microbial research, unravel the mysteries of the HIV virus; understand life, disease,
- Electric Power Grid → communication, transportation, energy, water, food supply
- Modeling and Simulation → Perform full-scale economic-social-political simulations

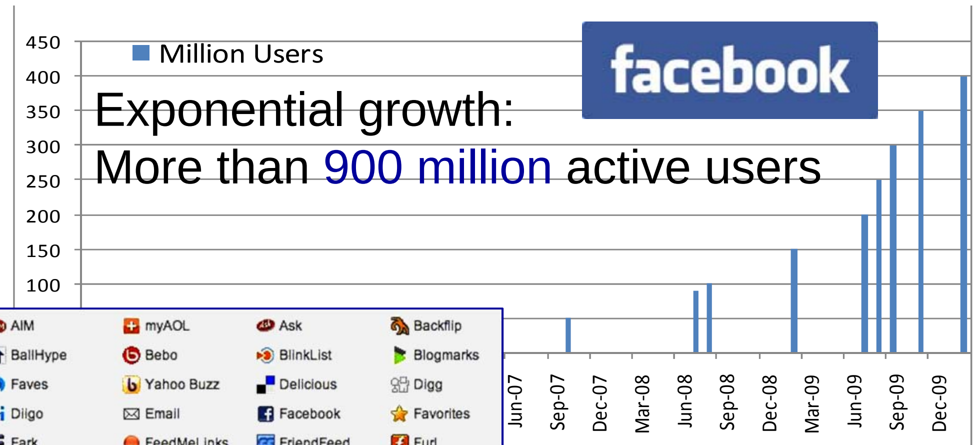

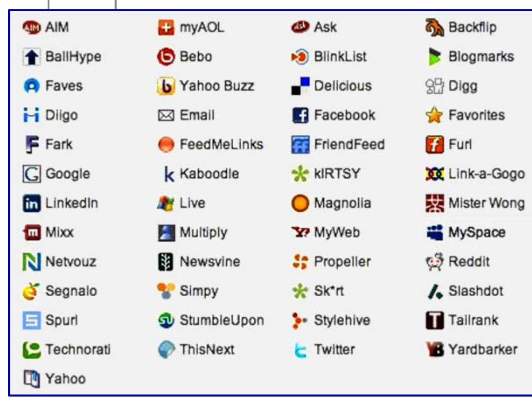

Ex: discovered minimal changes in O(billions)-size complex network that could hide or reveal top influencers in the community

Sample queries:

Allegiance switching: identify entities that switch communities.

Community structure: identify the genesis and dissipation of communities

Phase change: identify significant change in the network structure

REQUIRES PREDICTING / INFLUENCE CHANGE IN REAL-TIME AT SCALE

[**GRATEFUL**: Graph Analysis Tackling power Efficiency, Uncertainty, and Locality]

David A. Bader and Jason Riedy

[**OBJECTIVE**]

Research and develop new algorithms and software for crucial graph analysis problems in cybersecurity, intelligence integration, and network analysis.

Photo © CTL Corp.

Photo © Intel

[**DESCRIPTION**]

GRATEFUL will contribute to achieving resilience against errors and also extreme power efficiency through DARPA's Power Efficiency Revolution for Embedded Computing Technologies (PERFECT) program, continuing the trend of algorithms and software contributing along side of hardware advances to reduce power and increasing performance.
GRATEFUL enables deployment of advanced, mission-critical graph analysis applications within DARPA's power, performance, and resilience constraints.

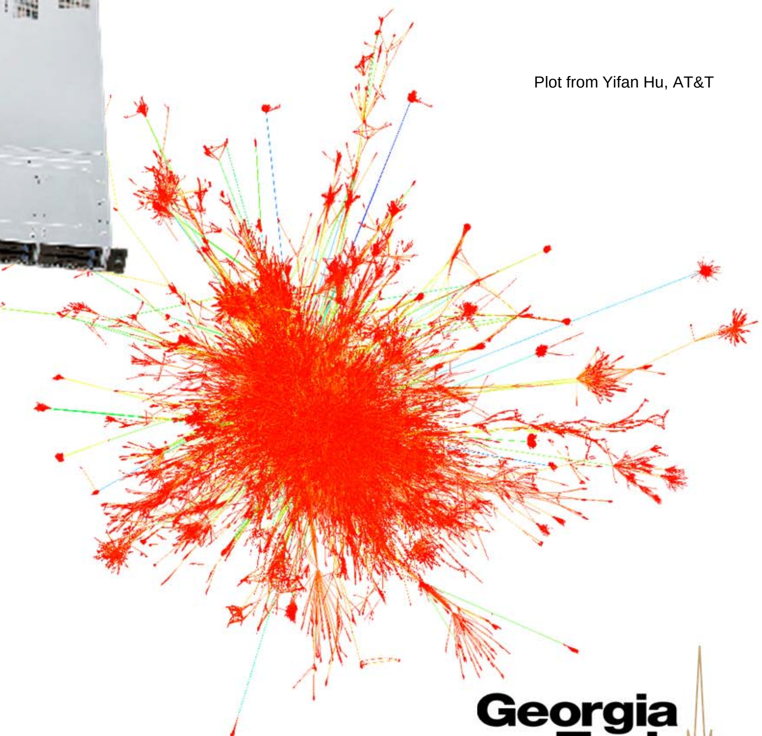

Plot from Yifan Hu, AT&T

Sponsored under the DARPA PERFECT program, Contract HR0011-13-2-0001

Example: Mining Twitter for Social Good

ICPP 2010

Massive Social Network Analysis: Mining Twitter for Social Good

David Ediger Karl Jiang Jason Riedy David A. Bader
Georgia Institute of Technology
Atlanta, GA, USA

Courtney Corley Rob Farber
Pacific Northwest National Lab.
Richland, WA, USA

William N. Reynolds
Least Squares Software, In
Albuquerque, NM, USA

Abstract—Social networks produce an enormous quantity of data. Facebook consists of over 400 million active users sharing over 5 *billion* pieces of information each month. Analyzing this vast quantity of unstructured data presents challenges for software and hardware. We present GraphCT, a *Graph* Characterization *T*oolkit for massive graphs representing social network data. On a 128-processor Cray XMT, GraphCT estimates the betweenness centrality of an artificially generated (R-MAT) 537 million vertex, 8.6 billion edge graph in 55 minutes and a real-world graph (Kwak, *et al.*) with 61.6 million vertices and 1.47 billion edges in 105 minutes. We use GraphCT to analyze public data from Twitter, a microblogging network. Twitter's message connections appear primarily tree-structured as a news dissemination system. Within the

involves over 400 million active users with an ave 120 'friendship' connections each and sharing 5 references to items each month [11].

One analysis approach treats the interactions as and applies tools from graph theory, social r analysis, and scale-free networks [29]. Howe volume of data that must be processed to appl techniques overwhelms current computational cap: Even well-understood analytic methodologies advances in both hardware and software to proc growing corpus of social media.

Social media provides staggering amounts

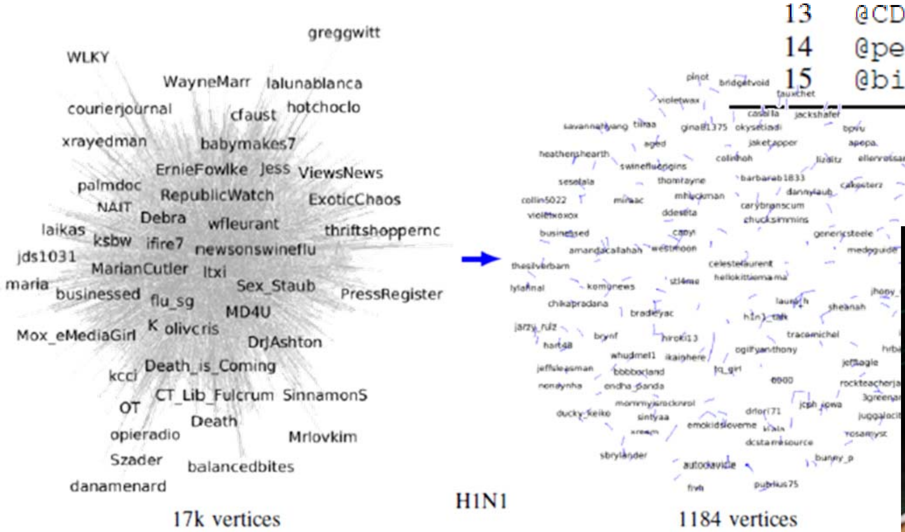

Fig. 3. Subcommunity filtering on Twitter data sets

TOP 15 USERS BY BETWEENNESS CENTRALITY

Rank	Data Set	
	H1N1	atlflood
1	@CDCFlu	@ajc
2	@addthis	@driveafaste
3	@Official_PAX	@ATLCheap
4	@FluGov	@TWCi
5	@nytimes	@HelloNorthGA
6	@tweetmeme	@11AliveNews
7	@mercola	@WSB_TV
8	@CNN	@shaunking
9	@backstreetboys	@Carl
10	@EllieSmith_x	@SpaceyG
11	@TIME	@ATLINtownPa
12	@CDCemergency	@TJsDJs
13	@CDC_eHealth	@ATLien
14	@perezhilton	@MarshallRamsey
15	@billmaher	@Kanye

twitter
public tweets

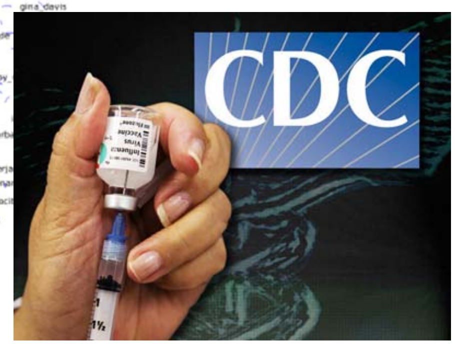

Image credit: bioethicsinstitute.org

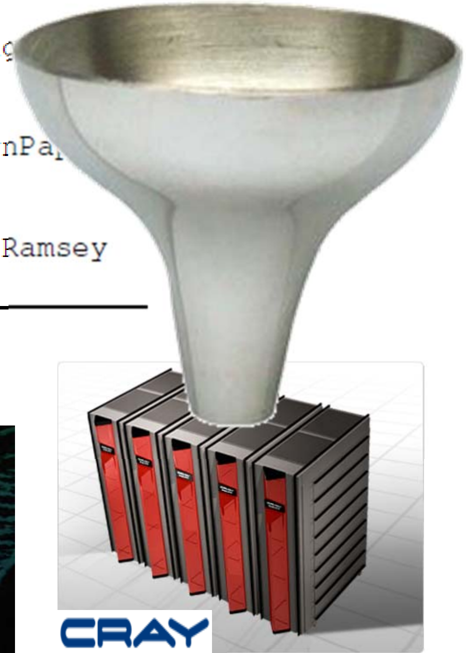

Pacific Northwest
NATIONAL LABORATORY

Georgia Tech College of Computing

Emcien: Automatically Connects the Dots To Deliver Insight

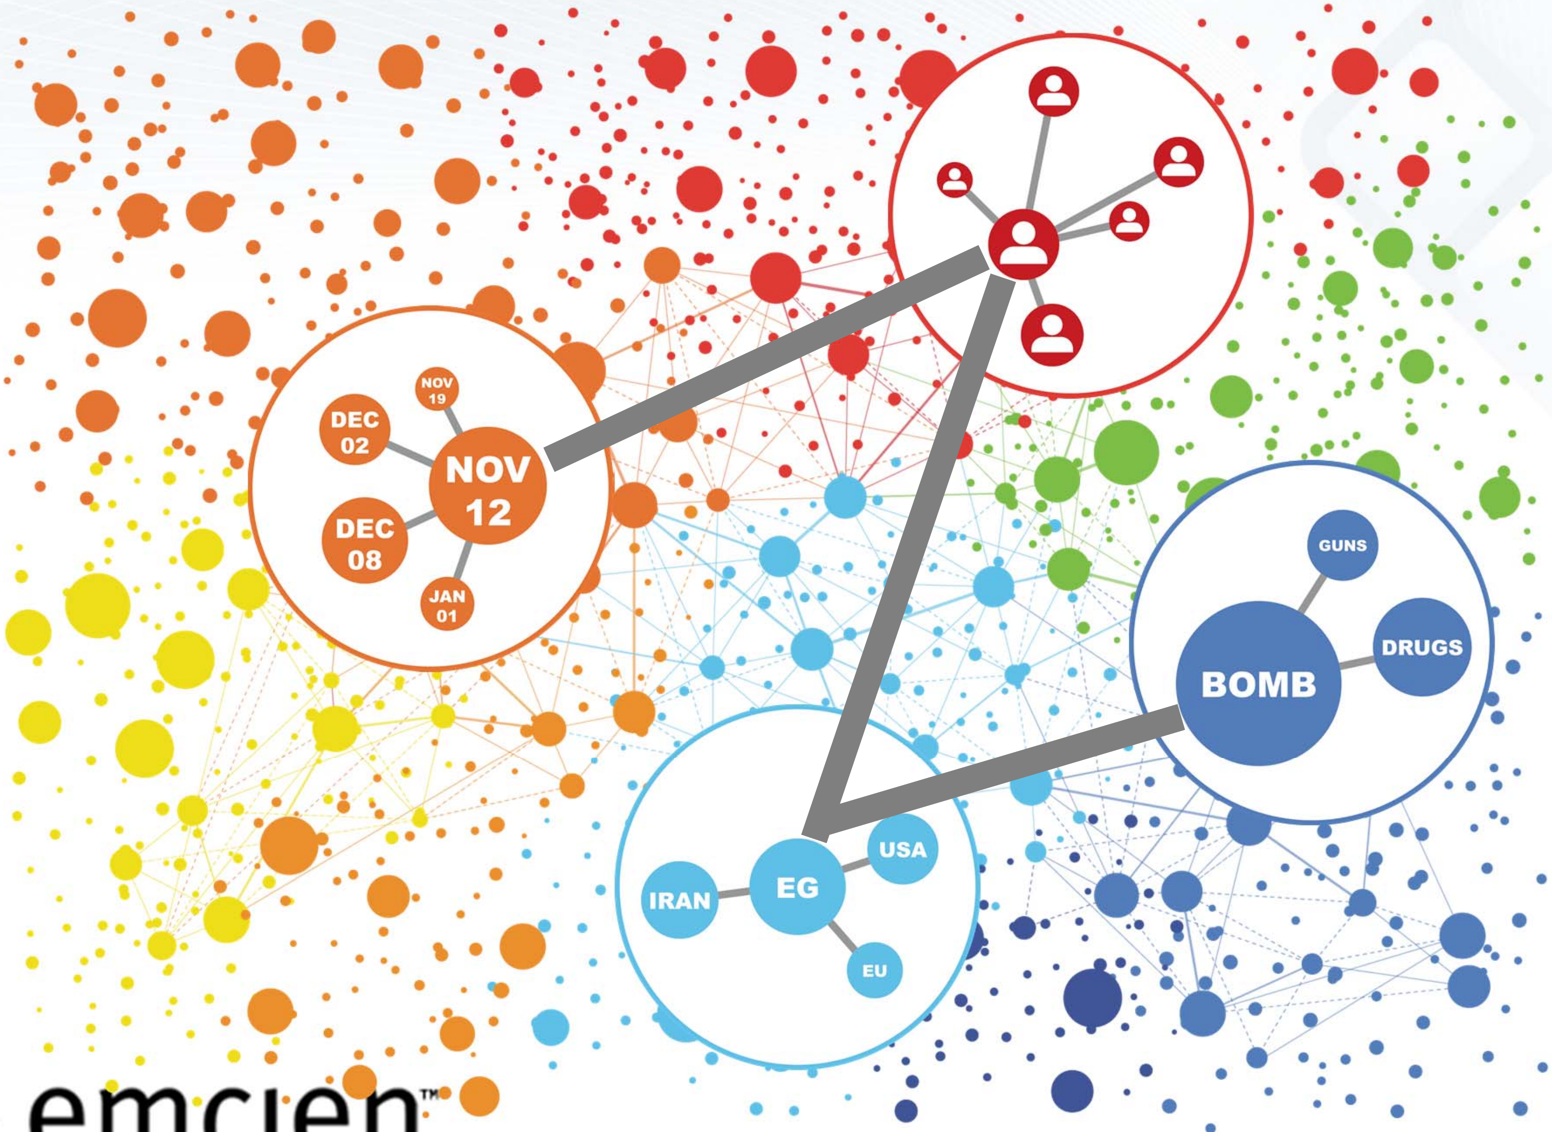

emcien™
Converting Data into Value

Graph500 Benchmark, www.graph500.org

Defining a new set of benchmarks to guide the design of hardware architectures and software systems intended to support such applications and to help procurements. Graph algorithms are a core part of many analytics workloads.

Executive Committee: D.A. Bader, R. Murphy, M. Snir, A. Lumsdaine

- Five Business Area Data Sets:
  - Cybersecurity
    - 15 Billion Log Entires/Day (for large enterprises)
    - Full Data Scan with End-to-End Join Required
  - Medical Informatics
    - 50M patient records, 20-200 records/patient, billions of individuals
    - Entity Resolution Important
  - Social Networks
    - Example, Facebook, Twitter
    - Nearly Unbounded Dataset Size
  - Data Enrichment
    - Easily PB of data
    - Example: Maritime Domain Awareness
      - Hundreds of Millions of Transponders
      - Tens of Thousands of Cargo Ships
      - Tens of Millions of Pieces of Bulk Cargo
      - May involve additional data (images, etc.)
  - Symbolic Networks
    - Example, the Human Brain
    - 25B Neurons
    - 7,000+ Connections/Neuron

The Cray XMT

- Tolerates latency by massive multithreading
  - Hardware support for 128 threads on each processor
  - Globally hashed address space
  - No data cache
  - Single cycle context switch
  - Multiple outstanding memory requests
- Support for fine-grained, word-level synchronization
  - Full/empty bit associated with every memory word
- Flexibly supports dynamic load balancing

Image Source: cray.com

- Our graph SW currently tested on a 512 processor XMT: 64K threads
  - 4 TB of globally shared memory
- XMT2 supports 512 processors and 64TB of globally shared memory

Comparing Performance of Large Graph Algorithms on MapReduce and the Cray XMT

- Bin Wu [Wu & Du, AICI 2010; Wu *et al.* ICNC 2011] present results on connected components and clustering coefficients using **Hadoop** on a commodity cluster.
  - Wu *et al.* use label propagation for the determination of the connected components of a static graph.
  - The test platform is a cluster containing 8 servers with 4 GB main memory each
  - The test data set is extracted from a telephone call graph containing 1.2 million vertices and 16 million edges.
  - The execution time presented in Figure 3 of their paper indicates that the computation of connected components required **approximately 40 minutes.**

- Using our code running on a 128-processor Cray XMT.
  - Using a synthetic RMAT graph generator with a power-law distribution in the number of neighbors, we obtained a graph with 134 million vertices and 2.1 billion edges.
  - The connected components computation took **15.1 seconds.**
  - A good estimate of I/O time is about 3 minutes.

→ Processing rates: 6K edges/sec vs. 139M edges/sec !!

→ Our implementation on the **Cray XMT** is **159 times faster** than the **Hadoop** results presented, on a graph **with 133 times as many edges.**

  → Over 4 orders of magnitude FASTER edge processing rate!

STINGER Software Dissemination

- http://www.cc.gatech.edu/stinger

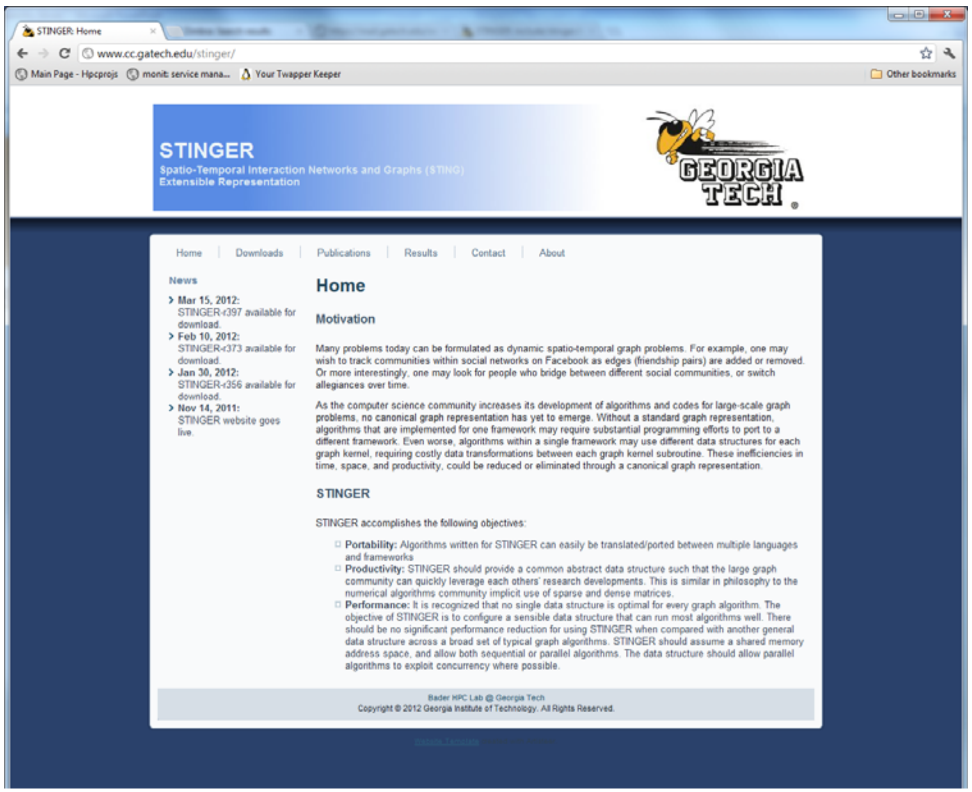

- Handles high-rates of streaming data, concurrently ingested with complex analytics
- Semantic relations (vertices and edges have types)
- Streaming analytics for maintaining graph properties, anomaly detection, temporal analytics

➔ Predictive analytics in real-time

Conclusions

- Need solutions that can handle unstructured 'Big Data' *in motion*, complex analytics, and fast transactions
- Solving massive scale analytic problems in real-time requires strategic investments in a portfolio of
  - high performance computing architectures
  - data-intensive programming models, and
  - parallel algorithms.
- Mapping applications to high performance architectures may yield 6 or more orders of magnitude performance improvements

Collaborators and Acknowledgments

- Jason Riedy, Research Scientist, (Georgia Tech)
- Graduate Students (Georgia Tech):
  - David Ediger
  - James Fairbanks
  - Oded Green
  - Rob McColl
  - Pushkar Pande
  - Anita Zakrzewska
  - Karl Jiang
- Bader PhD Graduates:
  - Seunghwa Kang (Pacific Northwest National Lab)
  - Kamesh Madduri (Penn State)
  - Guojing Cong (IBM TJ Watson Research Center)
- John Feo and Daniel Chavarría-Miranda (Pacific Northwest Nat'l Lab)

Bader, Related Recent Publications (2005-2008)

- D.A. Bader, G. Cong, and J. Feo, "On the Architectural Requirements for Efficient Execution of Graph Algorithms," *The 34th International Conference on Parallel Processing* (ICPP 2005), pp. 547-556, Georg Sverdrups House, University of Oslo, Norway, June 14-17, 2005.
- D.A. Bader and K. Madduri, "Design and Implementation of the HPCS Graph Analysis Benchmark on Symmetric Multiprocessors," *The 12th International Conference on High Performance Computing* (HiPC 2005), D.A. Bader *et al.*, (eds.), Springer-Verlag LNCS 3769, 465-476, Goa, India, December 2005.
- D.A. Bader and K. Madduri, "Designing Multithreaded Algorithms for Breadth-First Search and st-connectivity on the Cray MTA-2," *The 35th International Conference on Parallel Processing* (ICPP 2006), Columbus, OH, August 14-18, 2006.
- D.A. Bader and K. Madduri, "Parallel Algorithms for Evaluating Centrality Indices in Real-world Networks," *The 35th International Conference on Parallel Processing* (ICPP 2006), Columbus, OH, August 14-18, 2006.
- K. Madduri, D.A. Bader, J.W. Berry, and J.R. Crobak, "Parallel Shortest Path Algorithms for Solving Large-Scale Instances," *9th DIMACS Implementation Challenge -- The Shortest Path Problem*, DIMACS Center, Rutgers University, Piscataway, NJ, November 13-14, 2006.
- K. Madduri, D.A. Bader, J.W. Berry, and J.R. Crobak, "An Experimental Study of A Parallel Shortest Path Algorithm for Solving Large-Scale Graph Instances," *Workshop on Algorithm Engineering and Experiments* (ALENEX), New Orleans, LA, January 6, 2007.
- J.R. Crobak, J.W. Berry, K. Madduri, and D.A. Bader, "Advanced Shortest Path Algorithms on a Massively-Multithreaded Architecture," *First Workshop on Multithreaded Architectures and Applications* (MTAAP), Long Beach, CA, March 30, 2007.
- D.A. Bader and K. Madduri, "High-Performance Combinatorial Techniques for Analyzing Massive Dynamic Interaction Networks," *DIMACS Workshop on Computational Methods for Dynamic Interaction Networks*, DIMACS Center, Rutgers University, Piscataway, NJ, September 24-25, 2007.
- D.A. Bader, S. Kintali, K. Madduri, and M. Mihail, "Approximating Betewenness Centrality," The *5th Workshop on Algorithms and Models for the Web-Graph* (WAW2007), San Diego, CA, December 11-12, 2007.
- David A. Bader, Kamesh Madduri, Guojing Cong, and John Feo, "Design of Multithreaded Algorithms for Combinatorial Problems," in S. Rajasekaran and J. Reif, editors, *Handbook of Parallel Computing: Models, Algorithms, and Applications*, CRC Press, Chapter 31, 2007.
- Kamesh Madduri, David A. Bader, Jonathan W. Berry, Joseph R. Crobak, and Bruce A. Hendrickson, "Multithreaded Algorithms for Processing Massive Graphs," in D.A. Bader, editor, *Petascale Computing: Algorithms and Applications*, Chapman & Hall / CRC Press, Chapter 12, 2007.
- D.A. Bader and K. Madduri, "SNAP, Small-world Network Analysis and Partitioning: an open-source parallel graph framework for the exploration of large-scale networks," *22nd IEEE International Parallel and Distributed Processing Symposium* (IPDPS), Miami, FL, April 14-18, 2008.

Bader, Related Recent Publications (2009-2010)

- S. Kang, D.A. Bader, "An Efficient Transactional Memory Algorithm for Computing Minimum Spanning Forest of Sparse Graphs," 14th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), Raleigh, NC, February 2009.
- Karl Jiang, David Ediger, and David A. Bader. "Generalizing k-Betweenness Centrality Using Short Paths and a Parallel Multithreaded Implementation." The 38th International Conference on Parallel Processing (ICPP), Vienna, Austria, September 2009.
- Kamesh Madduri, David Ediger, Karl Jiang, David A. Bader, Daniel Chavarría-Miranda. "A Faster Parallel Algorithm and Efficient Multithreaded Implementations for Evaluating Betweenness Centrality on Massive Datasets." 3rd Workshop on Multithreaded Architectures and Applications (MTAAP), Rome, Italy, May 2009.
- David A. Bader, et al. "STINGER: Spatio-Temporal Interaction Networks and Graphs (STING) Extensible Representation." 2009.
- David Ediger, Karl Jiang, E. Jason Riedy, and David A. Bader. "Massive Streaming Data Analytics: A Case Study with Clustering Coefficients," Fourth Workshop in Multithreaded Architectures and Applications (MTAAP), Atlanta, GA, April 2010.
- Seunghwa Kang, David A. Bader. "Large Scale Complex Network Analysis using the Hybrid Combination of a MapReduce cluster and a Highly Multithreaded System:," Fourth Workshop in Multithreaded Architectures and Applications (MTAAP), Atlanta, GA, April 2010.
- David Ediger, Karl Jiang, Jason Riedy, David A. Bader, Courtney Corley, Rob Farber and William N. Reynolds. "Massive Social Network Analysis: Mining Twitter for Social Good," The 39th International Conference on Parallel Processing (ICPP 2010), San Diego, CA, September 2010.
- Virat Agarwal, Fabrizio Petrini, Davide Pasetto and David A. Bader. "Scalable Graph Exploration on Multicore Processors," *The 22nd IEEE and ACM Supercomputing Conference* (SC10), New Orleans, LA, November 2010.
- Z. Du, Z. Yin, W. Liu, and D.A. Bader, "On Accelerating Iterative Algorithms with CUDA: A Case Study on Conditional Random Fields Training Algorithm for Biological Sequence Alignment," IEEE International Conference on Bioinformatics & Biomedicine, Workshop on Data-Mining of Next Generation Sequencing Data (NGS2010), Hong Kong, December 20, 2010.

Bader, Related Recent Publications (2011-2012)

- D. Ediger, J. Riedy, H. Meyerhenke, and D.A. Bader, "Tracking Structure of Streaming Social Networks," 5th Workshop on Multithreaded Architectures and Applications (MTAAP), Anchorage, AK, May 20, 2011.
- D. Mizell, D.A. Bader, E.L. Goodman, and D.J. Haglin, "Semantic Databases and Supercomputers," 2011 Semantic Technology Conference (SemTech), San Francisco, CA, June 5-9, 2011.
- E.J. Riedy, H. Meyerhenke, D. Ediger, and D.A. Bader, "Parallel Community Detection for Massive Graphs," The 9th International Conference on Parallel Processing and Applied Mathematics (PPAM 2011), Torun, Poland, September 11-14, 2011. Lecture Notes in Computer Science, 7203:286-296, 2012.
- E.J. Riedy, D. Ediger, D.A. Bader, and H. Meyerhenke, "Parallel Community Detection for Massive Graphs," 10th DIMACS Implementation Challenge -- Graph Partitioning and Graph Clustering, Atlanta, GA, February 13-14, 2012.
- E.J. Riedy, H. Meyerhenke, D.A. Bader, D. Ediger, and T. Mattson, "Analysis of Streaming Social Networks and Graphs on Multicore Architectures," The 37th IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Kyoto, Japan, March 25-30, 2012.
- J. Riedy, H. Meyerhenke, and D.A. Bader, "Scalable Multi-threaded Community Detection in Social Networks," 6th Workshop on Multithreaded Architectures and Applications (MTAAP), Shanghai, China, May 25, 2012.
- P. Pande and D.A. Bader, "Computing Betweenness Centrality for Small World Networks on a GPU," *The 15th Annual High Performance Embedded Computing Workshop* (HPEC), Lexington, MA, September 21-22, 2011.
- H. Meyerhenke, E.J. Riedy, and D.A. Bader, "Parallel Community Detection in Streaming Graphs," Minisymposium on Parallel Analysis of Massive Social Networks, *15th SIAM Conference on Parallel Processing for Scientific Computing* (PP12), Savannah, GA, February 15-17, 2012.
- D. Ediger, E.J. Riedy, H. Meyerhenke, and D.A. Bader, "Analyzing Massive Networks with GraphCT," Poster Session, *15th SIAM Conference on Parallel Processing for Scientific Computing* (PP12), Savannah, GA, February 15-17, 2012.
- R.C. McColl, D. Ediger, and D.A. Bader, "Many-Core Memory Hierarchies and Parallel Graph Analysis," Poster Session, *15th SIAM Conference on Parallel Processing for Scientific Computing* (PP12), Savannah, GA, February 15-17, 2012.
- E.J. Riedy, D. Ediger, H. Meyerhenke, and D.A. Bader, "STING: Software for Analysis of Spatio-Temporal Interaction Networks and Graphs," Poster Session, *15th SIAM Conference on Parallel Processing for Scientific Computing* (PP12), Savannah, GA, February 15-17, 2012.
- David A. Bader, Christine Heitsch, and Kamesh Madduri, "Large-Scale Network Analysis," in J. Kepner and J. Gilbert, editor, *Graph Algorithms in the Language of Linear Algebra*, SIAM Press, Chapter 12, pages 253-285, 2011.
- Jeremy Kepner, David A. Bader, Robert Bond, Nadya Bliss, Christos Faloutsos, Bruce Hendrickson, John Gilbert, and Eric Robinson, "Fundamental Questions in the Analysis of Large Graphs," in J. Kepner and J. Gilbert, editor, *Graph Algorithms in the Language of Linear Algebra*, SIAM Press, Chapter 16, pages 353-357, 2011.

Acknowledgment of Support

NATIONAL INSTITUTES OF HEALTH

DEPARTMENT OF ENERGY · UNITED STATES OF AMERICA

NASA

DEPARTMENT OF DEFENSE · UNITED STATES OF AMERICA

DARPA

OAK RIDGE National Laboratory

Pacific Northwest NATIONAL LABORATORY

Sandia National Laboratories

IBM®

Microsoft® Research

NVIDIA.

NORTHROP GRUMMAN

® LexisNexis®

SONY

intel®

CRAY

Sun microsystems

XILINX®

TOSHIBA