Networking & Information Technology Research and Development Program

High End Computing Interagency Working Group (HEC IWG) & Big Data Senior Steering Group (BDSSG)

Peter Lyster, Deputy Director
National Coordination Office NITRD

Supercomputing and Big Data: From Collision to Convergence

Networking and Information Technology Research and Development Program

- Created by the High-Performance Computing (HPC) Act of 1991 (Public Law 102-194)
- Purpose: *To assure U.S. leadership in, and accelerate development and deployment of, advanced networking, computing systems, software, and associated information technologies*.
- Overseen by the *National Coordination Office* (NCO)
    - **Purpose:** *Provides support for the NITRD Program by providing technical expertise, planning, and coordination and by serving as the Program's central point of contact*.
    - **Vision:** *To be a catalyst for collaboration, information exchange, and outreach to foster knowledge, methods, R&D, technology transfer, and innovation to meet the NITRD Program goals.*

NITRD

Organization

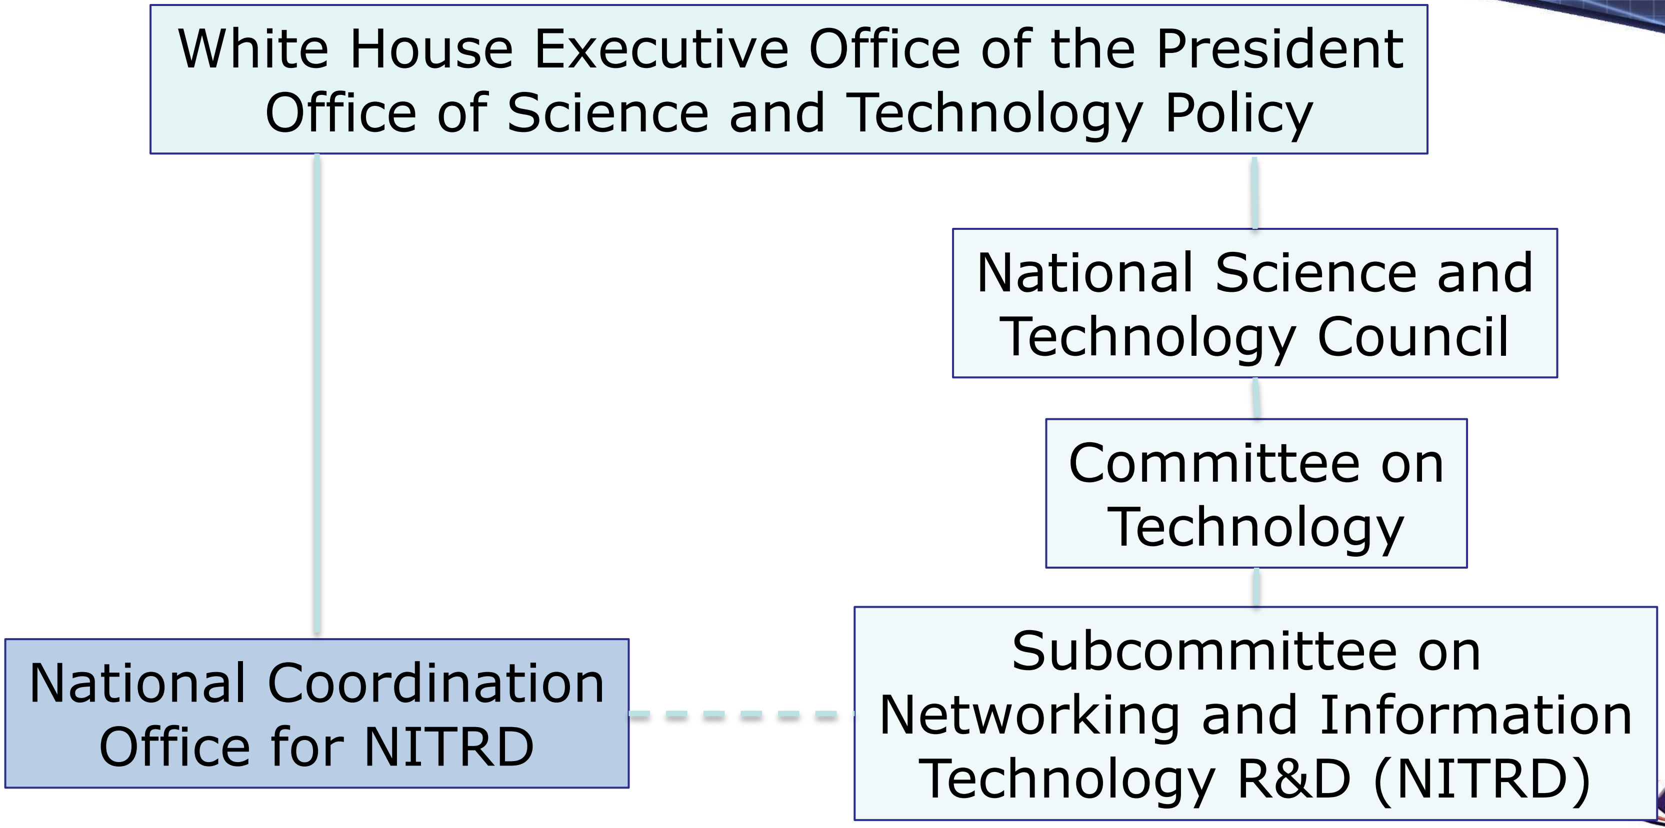

Goals

- What can supercomputing and big data communities learn from each other, and how can this be done?
- Can the technology for big data and high-fidelity HPC simulation really merge? If so, how may it happen, and when?
- What are the potential outcomes and impacts from such a merger?
- What research is needed to investigate the challenges and opportunities presented by the convergence of supercomputing and big data?

Panelists

- Randal Bryant, White House Office of Science and Technology Policy
- Andrew Moore, Carnegie Mellon University
- George Biros, University of Texas at Austin
- Ian Foster, Argonne National Laboratory & University of Chicago
- David Bader, Georgia Tech

High End Computing Interagency Working Group (HEC IWG)
&
Big Data Senior Steering Group (BDSSG)

Supercomputing and Big Data: From Collision to Convergence

NITRD

High Performance 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, understand life, unravel mysteries of disease

Electric Power Grid

- communication, transportation, energy, water, food supply

Modeling and Simulation

- Perform full-scale economic-social-political simulations

Unlike traditional applications in computational science and engineering, solving these problems at scale often raises new research challenges because of

- sparsity and the lack of locality in the massive data,
- design of parallel algorithms for massive, streaming data analytics, and
- the need for new HPDA supercomputers that are energy-efficient, resilient, and easy-to-program.

REQUIRES PREDICTING / INFLUENCE CHANGE IN REAL-TIME AT SCALE

NITRD

Recommendations

- **High Performance Data Analytics will require new**
  - High-performance computing platforms
  - Streaming algorithms
  - Energy-efficient implementations

  and are promising to solve real-world challenges!

- **Mapping applications to high performance architectures may yield 6 or more orders of magnitude performance improvement**

Backup Slides

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 $165M of research awards
- Steering Committees of the major HPC conferences, IPDPS and HiPC
- Multiple editorial boards in parallel and high performance computing
  - EIC of IEEE Transactions on Parallel and Distributed Systems
- Elected chair of IEEE and SIAM committees on HPC
- 230+ publications, ≥ 4,790 citations, *h*-index ≥ 38
- 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 in 2012 and as HPCwire's People to Watch in 2012 and 2014.

GEORGIA TECH

NITRD

Georgia Tech School of Computational Science and Engineering

THE CSE INNOVATION ECOSYSTEM:

CREATING SOLUTIONS AND LEADERS

Innovate. Collaborate. Problem Solved.

CSE is a diverse, interdisciplinary innovation ecosystem composed of award-winning faculty, researchers and students that

- Solves real-world problems and creates future leaders
- Enables breakthroughs in scientific discovery and engineering practice
- Uses the most advanced resources, techniques and ideas
- Is highly collaborative with an impressive roster of GT and industry partners

Ten Years of Success

- Founded: 2005
- Chair: David Bader
- Faculty:
  - 11 tenure track (FY 16)
  - 4 joint appointments
  - 6 adjunct faculty
  - 5 research scientists
- Administrative staff: 5
- Research expenditures: $5.6 million (FY 2015)
- High impact: $463K expenditure per faculty member

Award-Winning Faculty

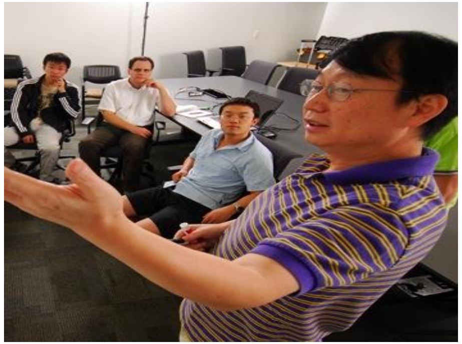

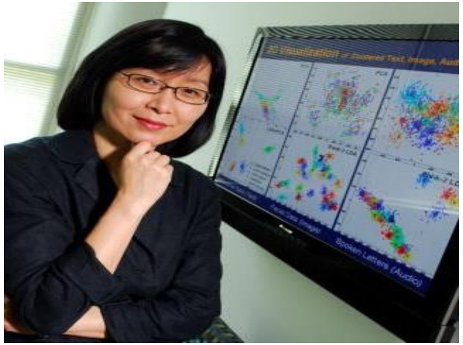

- 11 tenure-track faculty members (FY 16)
- 1 Regents' professor
- 5 NSF CAREER awards
- 2 IEEE fellows, 2 AAAS fellows, and 1 SIAM fellow
- 3 recent best paper awards and 2 finalists from SIAM, IEEE, etc.
- Several recent awards from industry:

    Accenture
    IBM
    Google
    NVIDIA
    Intel
    Lockheed Martin
    Yahoo! Labs
    Raytheon
    LexisNexis
    Microsoft Research
    Sony
    Cray
    Exxon Mobil

Faculty: Interdisciplinary Innovators

Georgia Tech | School of Computational Science and Engineering


Srinivas Aluru
Professor


David Bader
Professor and Chair

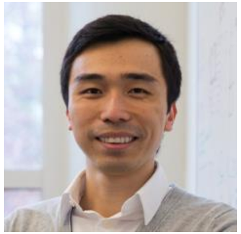
Polo Chau
Assistant Professor


Edmond Chow
Associate


Bistra Dilkina
Assistant Professor


Richard Fujimoto
Regents' Professor

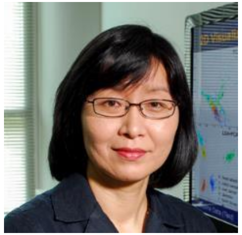
Haesun Park
Professor

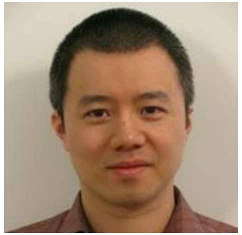
Le Song
Assistant


Jimeng Sun
Associate Professor

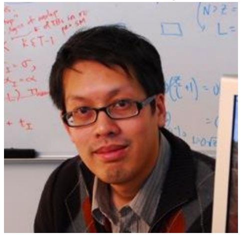
Richard Vuduc
Associate Professor

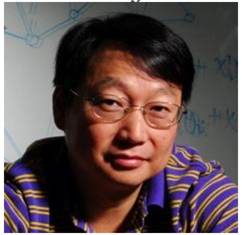
Hongyuan Zha
Professor


Kenneth Brown
Chemistry


Mark Borodovsky
BME


David Sherrill
Chemistry


Surya Kalidindi
Mech. Engr.

12 Pinnacle Projects > US$1M

Georgia Tech | School of Computational Science and Engineering

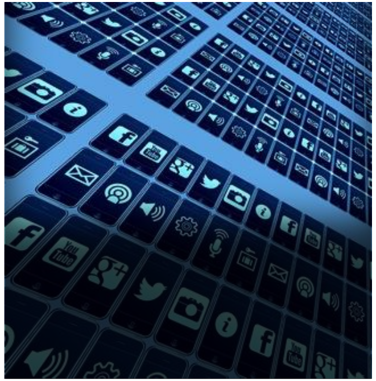

S. Aluru (PI), W. Feng, K. Olukotun, P. Schnable, C. Sing, and J. Zola, "BIGDATA: Mid-Scale: DA: Collaborative Research: Genomes Galore - Core Techniques, Libraries, and Domain Specific Languages for High-Throughput DNA Sequencing," NSF/NIH Bigdata Initiative, **$2M**

A. Somani, **S. Aluru** (Co-PI), R. Fox, E. Takle, and M. Gordon, "MRI: Acquisition of a HPC system for Data-Driven Discovery in Science and Engineering," National Science Foundation, **$1.8M**

S. Aluru (PI), K. Dorman, and P.S. Schnable, "AF:Medium: Parallel Algorithms and Software for High-throughput Sequence Assembly," National Science Foundation, **$1M**

Polo Chau (Co-PI), "Center of Excellence for Mobile Sensor Data-to-Knowledge (MD2K)," National Institute of Health, **$1.25M**

R. Fujimoto (Co-PI) and J. Crittenden (PI), "Participatory Modeling of Complex Urban Infrastructure Systems," National Science Foundation, **$2.5M**

R. Fujimoto (PI), T. Blum, **S. Kalidindi**, W. Newstetter, and **H. Zha**, "Computation-Enabled Design and Manufacturing of High Performance Materials," National Science Foundation, **$2.8M**

H. Park (PI), **H. Zha** (Co-PI), B. Drake (Co-PI), J. Choo (Co-PI), and J. Poulson (Co-PI), "Fast Algorithms on Imperfect, Heterogeneous, Distributed Data for Interactive Analysis," DARPA, **$2.7M**

H. Park (PI), J. Stasko (Co-PI), A. Gray (Co-PI), J. Monteiro (Co-PI), V. Koltchinskii (Co-PI), "FODAVA-lead: Dimension Reduction and Data Reduction: Foundations for Visualization," National Science Foundation and Department of Homeland Security, **$3.5M**

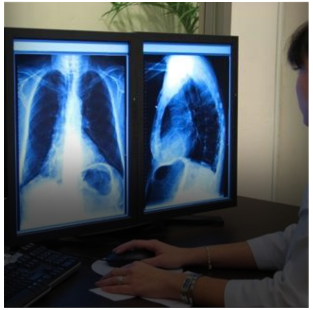

D. Bader (PI), **E.J. Riedy** (Co-PI), **R. Vuduc** (Co-PI), and V. Prasanna (PI), "SI2-SSI: Collaborative: The XScala Project: A Community Repository for Model-Driven Design and Tuning of Data-Intensive Applications for Extreme-Scale Accelerator-Based Systems," National Science Foundation, **$1.2M**

D. Bader (PI), **E.J. Riedy** (Co-PI), "GRATEFUL: GRaph Analysis Tackling power EFficiency, Uncertainty, and Locality, Power Efficiency Revolution for Embedded Computing Technologies (PERFECT) Program," DARPA, **$2.9M**

J. Sun, Smart Connect Health Project Award, National Science Foundation, **$2.1M**

H. Zha (Co-PI), "TWC SBE: TTP Option: Medium: Collaborative: EPICA: Empowering People to Overcome Information Controls and Attacks," National Science Foundation, **$1.1M** *...and more good news pending...*

Big Data Analytics

Answering the need for algorithms that scale to massive, complex data sets

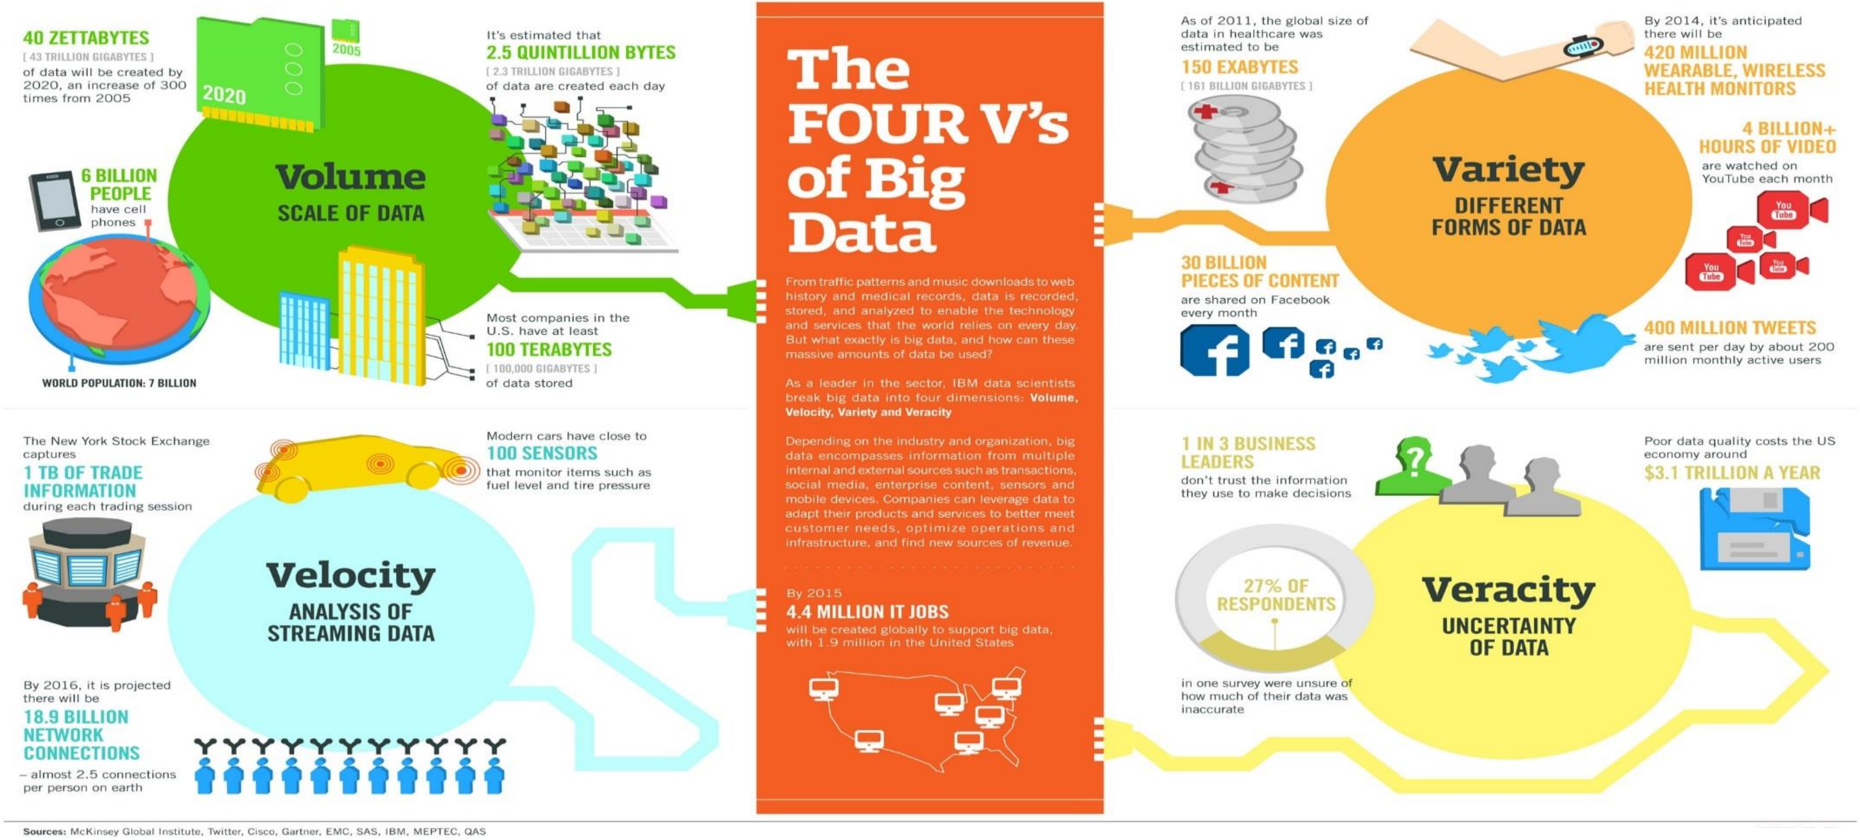

Core Research Areas

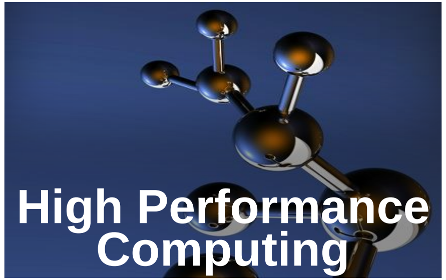

Devise computing solutions at the absolute limits of scale and speed using efficient, reliable and fast algorithms, software, tools and applications

Construct and study algorithms that build models, and make efficient data-driven predictions or decisions

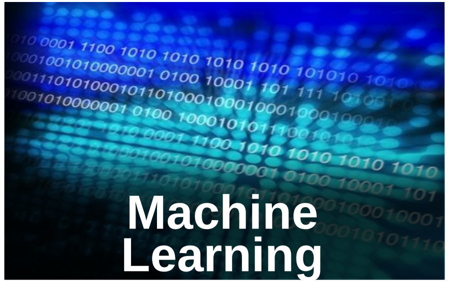

Develop new methods to analyze large and complex data sets, transforming data into value and solve grand challenges

Design fast theoretic algorithms on large-scale graphs, and detect malicious activity

Present data in ways that best yield insight and support decisions as problems scale and complexity increase

Graduate Education

- **Ph.D. and MS in Computational Science and Engineering**
- Ph.D. and MS in Bioengineering, Ph.D. in Bioinformatics, MS in Analytics

Strength in Diversity: CSE Home Units

School of Aerospace Engineering
School of Biology
Coulter Department of Biomedical Engineering
School of Chemistry and Biochemistry
School of Civil and Environmental Engineering
School of Computational Science and Engineering
School of Industrial and Systems Engineering
School of Mathematics

Students select a **Home** – unit & (*if applicable*) advisor
Coursework – **Core** + **Computation** + **Application**
Research – **Dissertation**
(*MS thesis option* + *PhD only*)

Georgia Tech School of Computational Science and Engineering

ANNIVERSARY

10 Years

ANNIVERSARY

2005 - 2015

Bader, Related Recent Publications (2005-2009)

Georgia Tech School of Computational Science and Engineering

- 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.
- 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.

Bader, Related Recent Publications (2010-2011)

Georgia Tech School of Computational Science and Engineering

- 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.
- 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.
- 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.
- 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.

Bader, Related Recent Publications (2012)

- 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.
- 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 Conferenc 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.
- Y. Chai, Z. Du, D.A. Bader, and X. Qin, "**Efficient Data Migration to Conserve Energy in Streaming Media Storage Systems,**" *IEEE Transactions on Parallel & Distributed Systems*, 2012
- M. S. Swenson, J. Anderson, A. Ash, P. Gaurav, Z. Sükösd, D.A. Bader, S.C. Harvey and C.E Heitsch, "**GTfold: Enabling parallel RNA secondary structure prediction on multi-core desktops**," *BMC Research Notes*, 5:341, 2012.
- D. Ediger, K. Jiang, E.J. Riedy, and D.A. Bader, "**GraphCT: Multithreaded Algorithms for Massive Graph Analysis**," *IEEE Transactions on Parallel & Distributed Systems*, 2012.
- D.A. Bader and K. Madduri, "**Computational Challenges in Emerging Combinatorial Scientific Computing Applications**," in O. Schenk, editor, *Combinatorial Scientific Computing*, Chapman & Hall / CRC Press, Chapter 17, pages 471-494, 2012.
- O. Green, R. McColl, and D.A. Bader, "**GPU Merge Path -- A GPU Merging Algorithm**," *26th ACM International Conference on Supercomputing* (ICS), San Servolo Island, Venice, Italy, June 25-29, 2012.
- O. Green, R. McColl, and D.A. Bader, "**A Fast Algorithm for Streaming Betweenness Centrality**," *4th ASE/IEEE International Conference on Social Computing* (SocialCom), Amsterdam, The Netherlands, September 3-5, 2012.
- D. Ediger, R. McColl, J. Riedy, and D.A. Bader, "**STINGER: High Performance Data Structure for Streaming Graphs**," *The IEEE High Performance Extreme Computing Conference* (HPEC), Waltham, MA, September 20-22, 2012. **Best Paper Award.**
- J. Marandola, S. Louise, L. Cudennec, J.-T. Acquaviva and D.A. Bader, **"Enhancing Cache Coherent Architecture with Access Patterns for Embedded Manycore Systems**," *14th IEEE International Symposium on System-on-Chip* (SoC), Tampere, Finland, October 11-12, 2012.
- L.M. Munguía, E. Ayguade, and D.A. Bader, "**Task-based Parallel Breadth-First Search in Heterogeneous Environments**," *The 19th Annual IEEE International Conference on High Performance Computing* (HiPC), Pune, India, December 18-21, 2012.

Bader, Related Recent Publications (2013)

Georgia Tech School of Computational Science and Engineering

- X. Liu, P. Pande, H. Meyerhenke, and D.A. Bader, "**PASQUAL: Parallel Techniques for Next Generation Genome Sequence Assembly,"** *IEEE Transactions on Parallel & Distributed Systems*, 24(5):977-986, 2013.
- David A. Bader, Henning Meyerhenke, Peter Sanders, and Dorothea Wagner (eds.), ***Graph Partitioning and Graph Clustering***, American Mathematical Society, 2013.
- E. Jason Riedy, Henning Meyerhenke, David Ediger and David A. Bader, "**Parallel Community Detection for Massive Graphs**," in David A. Bader, Henning Meyerhenke, Peter Sanders, and Dorothea Wagner (eds.), *Graph Partitioning and Graph Clustering*, American Mathematical Society, Chapter 14, pages 207-222, 2013.
- S. Kang, D.A. Bader, and R. Vuduc, "**Energy-Efficient Scheduling for Best-Effort Interactive Services to Achieve High Response Quality**," *27th IEEE International Parallel and Distributed Processing Symposium* (IPDPS), Boston, MA, May 20-24, 2013.
- J. Riedy and D.A. Bader, "**Multithreaded Community Monitoring for Massive Streaming Graph Data**," *7th Workshop on Multithreaded Architectures and Applications* (MTAAP), Boston, MA, May 24, 2013.
- D. Ediger and D.A. Bader, "**Investigating Graph Algorithms in the BSP Model on the Cray XMT**," *7th Workshop on Multithreaded Architectures and Applications* (MTAAP), Boston, MA, May 24, 2013.
- O. Green and D.A. Bader, "**Faster Betweenness Centrality Based on Data Structure Experimentation**," *International Conference on Computational Science* (ICCS), Barcelona, Spain, June 5-7, 2013.
- Z. Yin, J. Tang, S. Schaeffer, and D.A. Bader, "**Streaming Breakpoint Graph Analytics for Accelerating and Parallelizing the Computation of DCJ Median of Three Genomes**," *International Conference on Computational Science* (ICCS), Barcelona, Spain, June 5-7, 2013.
- T. Senator, D.A. Bader, et al., "**Detecting Insider Threats in a Real Corporate Database of Computer Usage Activities**," *19th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (KDD), Chicago, IL, August 11-14, 2013.
- J. Fairbanks, D. Ediger, R. McColl, D.A. Bader and E. Gilbert, "**A Statistical Framework for Streaming Graph Analysis**," *IEEE/ACM International Conference on Advances in Social Networks Analysis and Modeling* (ASONAM), Niagara Falls, Canada, August 25-28, 2013.
- A. Zakrzewska and D.A. Bader, "**Measuring the Sensitivity of Graph Metrics to Missing Data**," *10th International Conference on Parallel Processing and Applied Mathematics* (PPAM), Warsaw, Poland, September 8-11, 2013.
- O. Green and D.A. Bader, "**A Fast Algorithm for Streaming Betweenness Centrality**," *5th ASE/IEEE International Conference on Social Computing* (SocialCom), Washington, DC, September 8-14, 2013.
- R. McColl, O. Green, and D.A. Bader, "**A New Parallel Algorithm for Connected Components in Dynamic Graphs**," *The 20th Annual IEEE International Conference on High Performance Computing* (HiPC), Bangalore, India, December 18-21, 2013.

Bader, Related Recent Publications (2014-2015)

Georgia Tech School of Computational Science and Engineering

- R. McColl, D. Ediger, J. Poovey, D. Campbell, and D.A. Bader, "**A Performance Evaluation of Open Source Graph Databases**," *The 1st Workshop on Parallel Programming for Analytics Applications* (PPAA 2014) held in conjunction with the *19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (PPoPP 2014), Orlando, Florida, February 16, 2014.
- O. Green, L.M. Munguia, and D.A. Bader, "**Load Balanced Clustering Coefficients**," *The 1st Workshop on Parallel Programming for Analytics Applications* (PPAA 2014) held in conjunction with the *19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (PPoPP 2014), Orlando, Florida, February 16, 2014.
- A. McLaughlin and D.A. Bader, "**Revisiting Edge and Node Parallelism for Dynamic GPU Graph Analytics**," *8th Workshop on Multithreaded Architectures and Applications* (MTAAP), held in conjuntion with *The IEEE International Parallel and Distributed Processing Symposium (IPDPS 2014)*, Phoenix, AZ, May 23, 2014.
- Z. Yin, J. Tang, S. Schaeffer, D.A. Bader, "**A Lin-Kernighan Heuristic for the DCJ Median Problem of Genomes with Unequal Contents**," *20th International Computing and Combinatorics Conference* (COCOON), Atlanta, GA, August 4-6, 2014.
- Y. You, D.A. Bader and M.M. Dehnavi, "**Designing an Adaptive Cross-Architecture Combination for Graph Traversal,**" *The 43rd International Conference on Parallel Processing* (ICPP 2014), Minneapolis, MN, September 9-12, 2014.
- A. McLaughlin, J. Riedy, and D.A. Bader, "**Optimizing Energy Consumption and Parallel Performance for Betweenness Centrality using GPUs**," *The 18th Annual IEEE High Performance Extreme Computing Conference* (HPEC), Waltham, MA, September 9-11, 2014.
- A. McLaughlin and D.A. Bader, "**Scalable and High Performance Betweenness Centrality on the GPU**," *The 26th IEEE and ACM Supercomputing Conference* (SC14), New Orleans, LA, November 16-21, 2014. **Best Student Paper Finalist.**
- D. Dauwe, E. Jonardi, R. Friese, S. Pasricha, A.A. Maciejewski, D.A. Bader, and H.J. Siegel, "**A Methodology for Co-Location Aware Application Performance Modeling in Multicore Computing**," 17th Workshop on Advances on Parallel and Distributed Processing Symposium (APDCM), Hyderabad, India, May 25, 2015.
- A. Zakrzewska and D.A. Bader, "**Fast Incremental Community Detection on Dynamic Graphs**," 11th International Conference on Parallel Processing and Applied Mathematics (PPAM), Krakow, Poland, September 6-9, 2015.
- A. McLaughlin, J. Riedy, and D.A. Bader, "**An Energy-Efficient Abstraction for Simultaneous Breadth-First Searches**," The 19th Annual IEEE High Performance Extreme Computing Conference (HPEC), Waltham, MA, September 15-17, 2015.
- A. McLaughlin, D. Merrill, M. Garland and D.A. Bader, "**Parallel Methods for Verifying the Consistency of Weakly-Ordered Architectures**," The 24th International Conference on Parallel Architectures and Compilation Techniques (PACT), San Francisco, CA, October 18-21, 2015.

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

Scaling FLOP/SYNC-intensive HP Data Analytics (HPDA)

1. Need for end-to-end HPDA for CS&E
2. HPDA can be FLOPS/SYNC-intensive
3. Challenges:
   algorithms
   productivity

GEORGE BIROS
padas.ices.utexas.edu

INSTITUTE FOR COMPUTATIONAL ENGINEERING & SCIENCES
THE UNIVERSITY OF TEXAS AT AUSTIN

1. End-to-end HPDA in CS&E

TRADITIONAL	END-TO-END
Check-pointing	Uncertainty quantification
Sampling	Coupled sampling
Visualization	Adjoint-based assimilation
Data-assimilation	Pattern recognition
Trajectory analysis	Model reduction
3D/time correlations	Streaming

2. FLOPS/SYNC-intensive HPDA

METHODS	ALGORITHMS
Nearest-neighbors	Linear algebra
Kernel methods	Optimization
Logistic regression	Geometry
Support vectors	Sampling
Graphical models	Indexing/searching
Deep learning	Graphs/Trees

now: hadoop, async stochastic gradient, but: jacobi vs N-body

N-body methods $O(N^2) \rightarrow O(N)$

simulation

- Gravity & Coulomb
- Waves & Scattering
- Fluids & Transport

data analysis

- Graphics
- Machine learning
- Kriging
- Image analysis

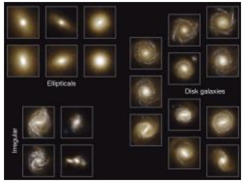

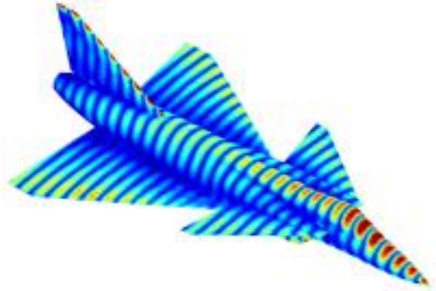

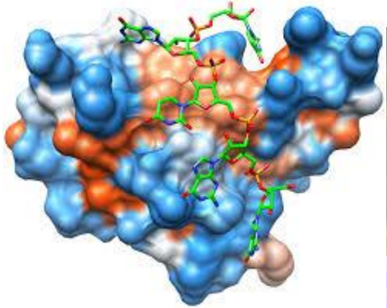

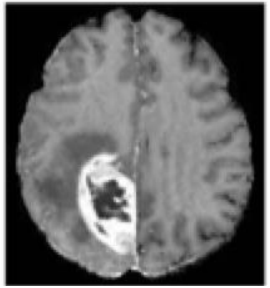

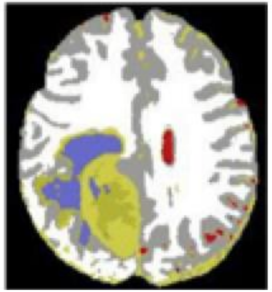

Scaling TACC's STAMPEDE 16 sandy bridge/ node

FLUIDS 12B / 3D ~300GB

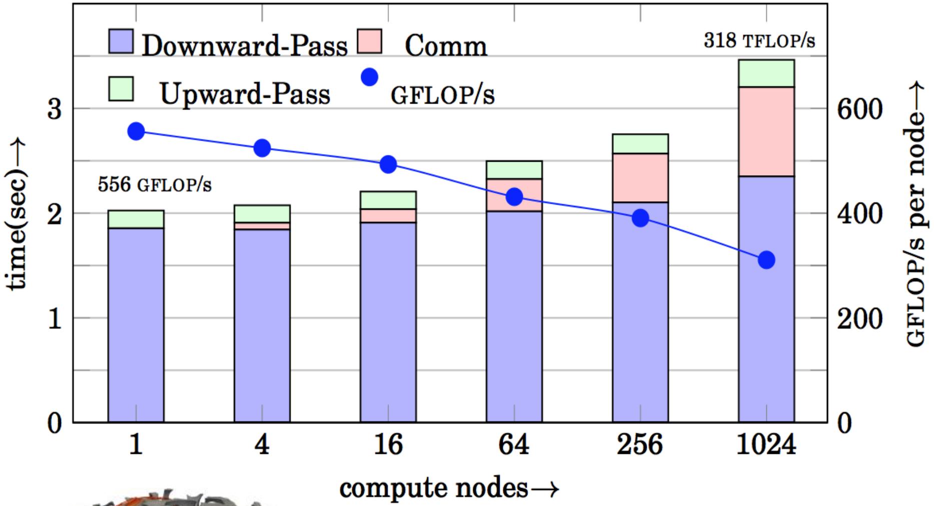

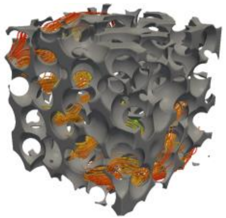

CLASSIFICATION 1B / 128D ~500GB

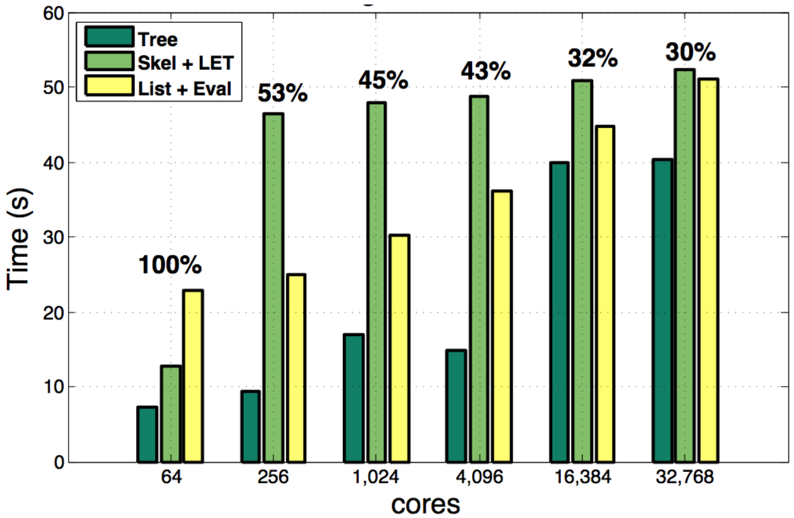

Malhotra, Gholami, B., SC'14 *Best Student Paper Finalist*

March W, Yu C , B. et al, IPDPS 15, KDD'15, **SC'15**

3. Challenges productivity/reproducibility/performance gaps

Next generation HPDA algorithms
Large evolving design space
Expanding complexity of
    Algorithms / APIs / Hardware
No parallel machine model
Scheduling / Streaming
End-to-end scalability

MPI+X
Open{MP,CL,ACC}
Pthreads, TBB
CUDA/SSE/AVX
PREFETCHING
NVRAMS
FPGAS

C++/MPI+X/BLAS/PETSc VS Java/Hadoop/SQL/SPARK

Supercomputing & Big Data
A Convergence?

Randal E. Bryant
Office of Science and Technology Policy

Two Classes of Large-Scale Systems

- **Modern supercomputer**

  - Run programs in hours or days that would require decades or centuries on normal machine
  - Designed for numerically-intensive applications

- **Internet Data Center**

  - Support millions of customers
    - Mostly small transactions
    - + large-scale analytics
  - Designed for data collection, storage, and analysis

Computing Trends

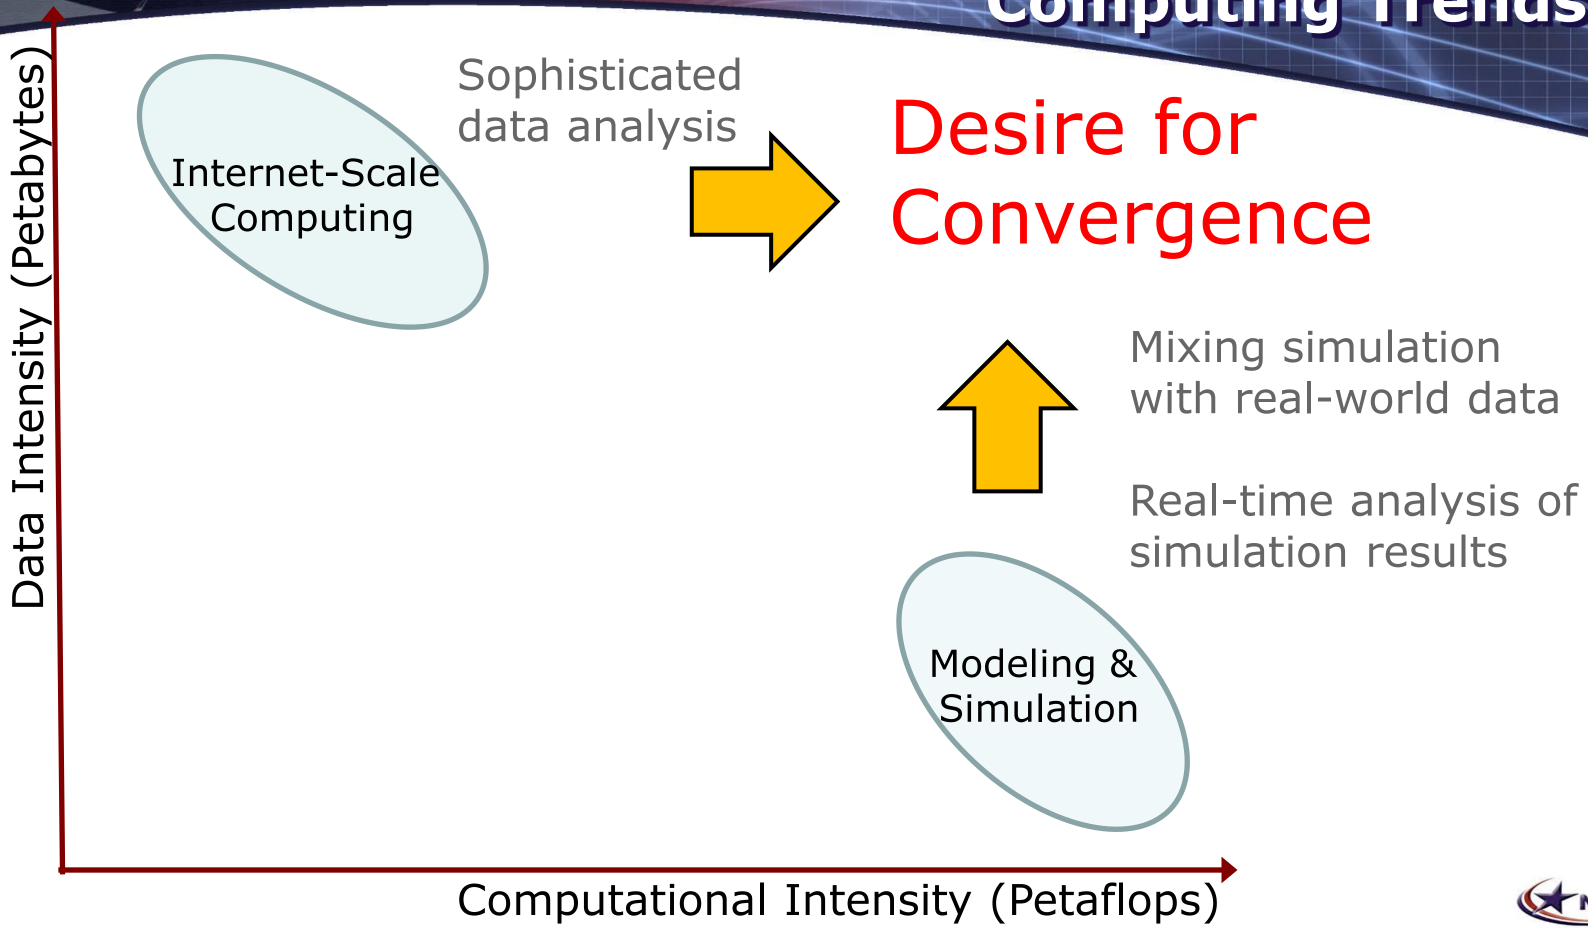

Titan Hardware

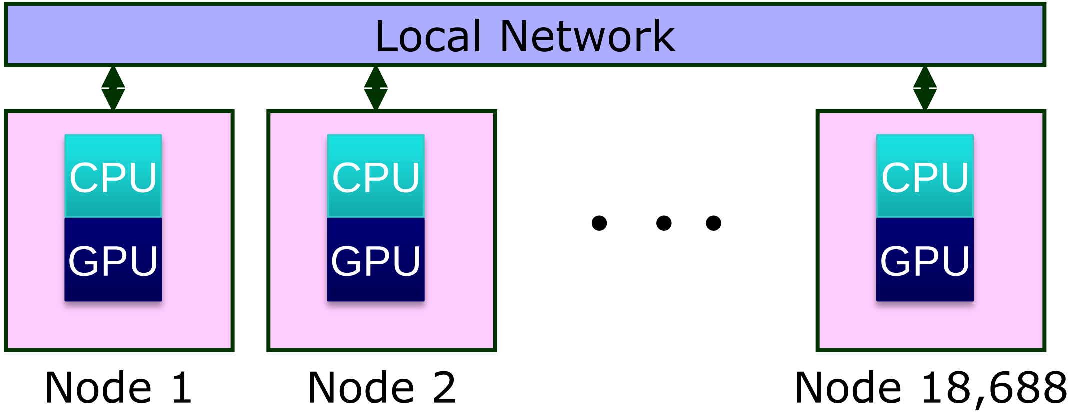

- **Each Node**
  - AMD 16-core processor
  - nVidia Graphics Processing Unit
  - 38 GB DRAM
  - *No disk drive*

Network Cluster

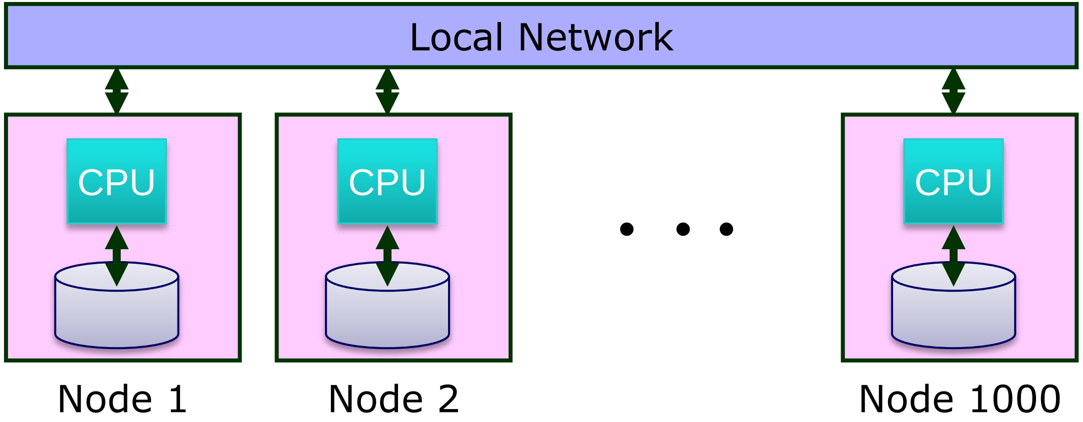

- **Typical Node**
  - 2 multicore CPUs
  - 2 disk drives
  - DRAM
- **Enhancements**
  - Flash memory
  - GPUs
  - Fast / high-bandwidth networking

Challenges for Convergence

Supercomputers	Network Clusters
Hardware	
▪ Customized	▪ Consumer grade
▪ Optimized for reliability	▪ Optimized for low cost
Run-Time System	
▪ Source of "noise"	▪ Provides reliability
▪ Static scheduling	▪ Dynamic allocation
▪ Strive for homogeneity	▪ Tolerate variability
Application Programming	
▪ Low-level, processor-centric model (e.g., MPI)	▪ High level, data-centric model (e.g., Hadoop, Spark)

NITRD

Issues to Be Discussed

- What can the supercomputing and big-data communities learn from each other?
- Can and should the technologies for big data and high-fidelity HPC simulation really merge?
- What new classes of applications would arise through a convergence?
- What research is needed to enable a convergence?

Ian Foster (foster@anl.gov)
Department of Computer Science, University of Chicago
Math and Computer Science, Argonne National Laboratory

Discovery engines for 21st Century Science

NITRD

Data-driven discovery requires discovery engines

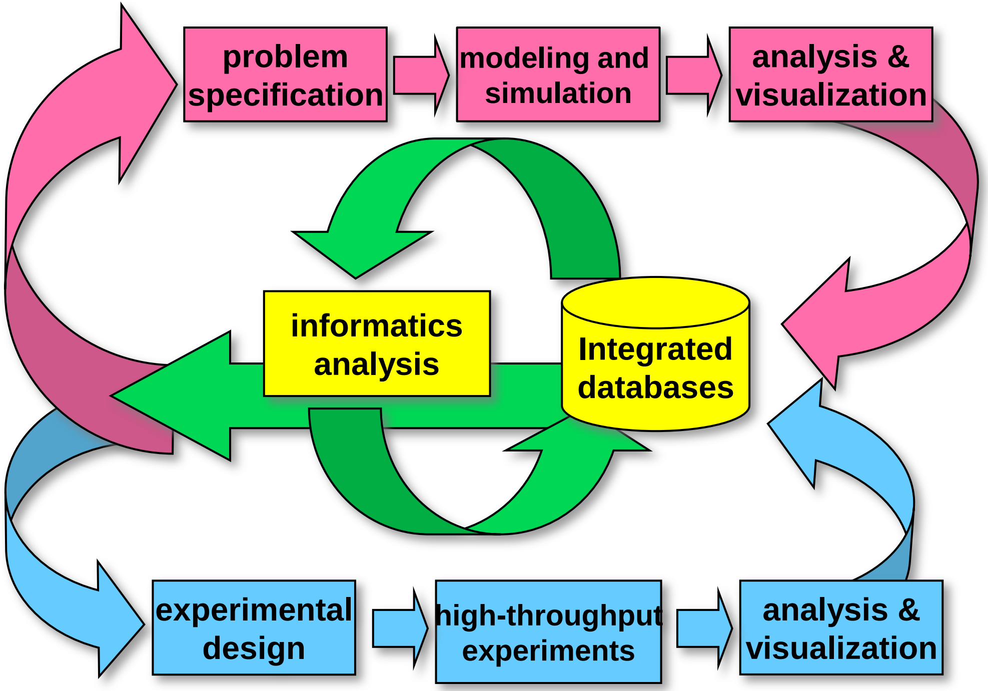

Discovery engine

- Knowledge base & computing engine for a disciplinary research program
- Tight coupling and automation for high-throughput discovery
- Centralized for economies of scale, knowledge sharing, and collaboration

Such systems exist, but are specialized and expensive

Ad pricing

- Ingest stream of web pages, searches, click stream; compute model offline; run experiments (?)

- Using a carefully architected and tuned pipeline on large and specialized computer system

Weather forecasting

- Ingest small number of data streams; run ensembles; propose expts; prepare standard products

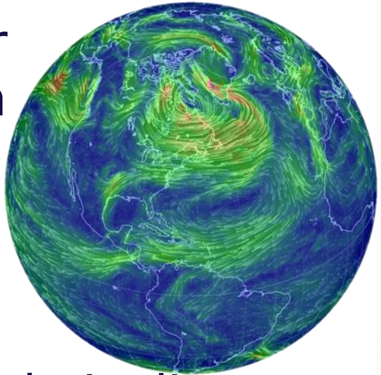

- … using a carefully architected and tuned pipeline on a large, specialized computer system

Science needs many discovery engines

Understanding dark matter and dark energy

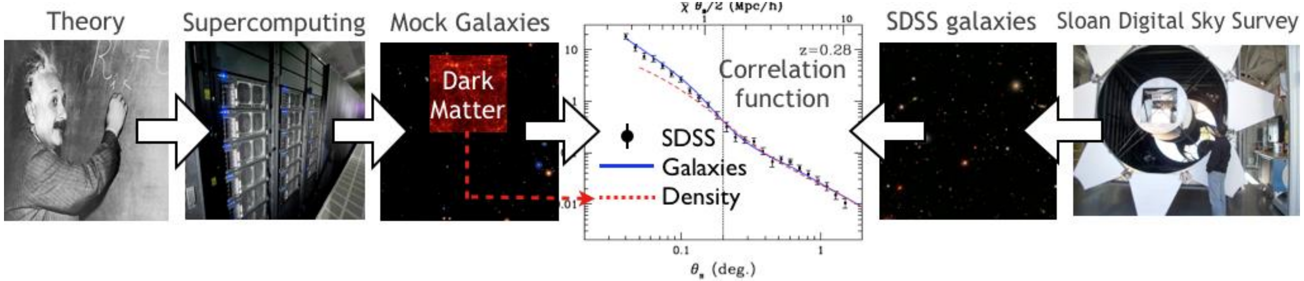

Understanding & designing disordered structures

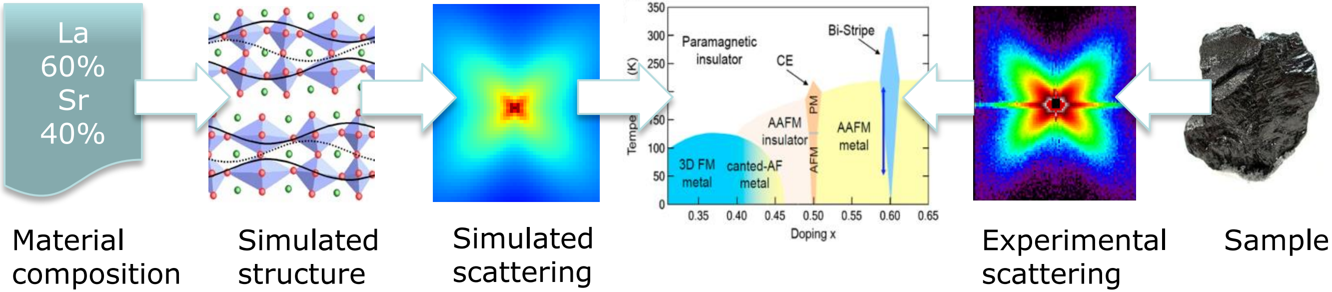

And more more:
- Health care
- Climate change
- Environment
- Ecosystems
- …
- …

Example: A discovery engine for disordered structures

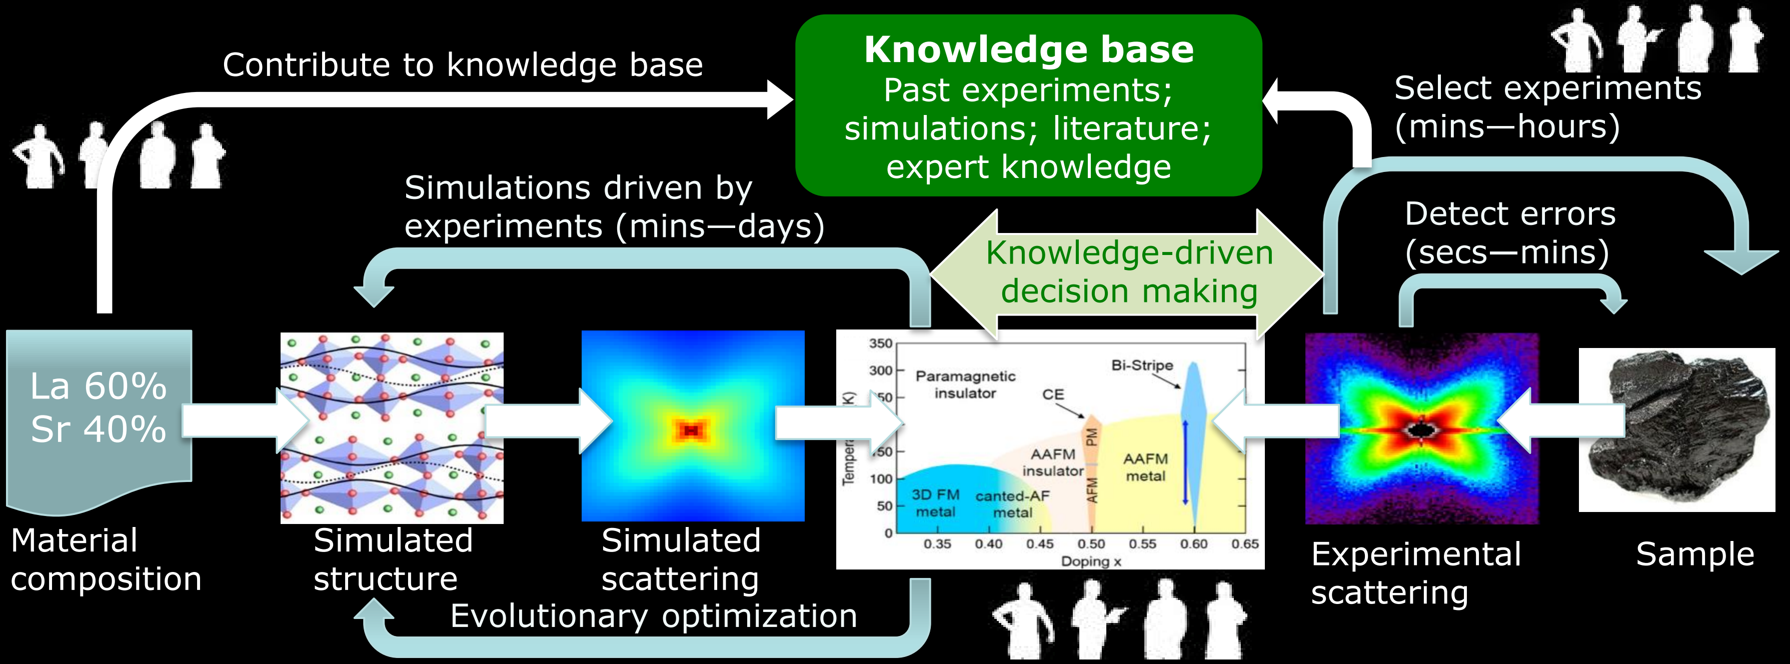

Diffuse scattering images from Ray Osborn et al., Argonne

Discovery engines and extreme-scale computing

KBase
PREDICTIVE BIOLOGY

Opportunity:
Reach many more researchers than extreme-scale simulation

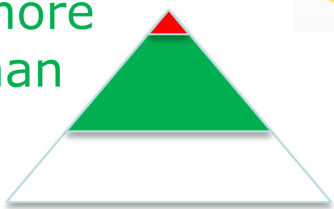

Exciting research agenda

- End-to-end automation to slash costs

- Massive knowledge management and fusion
- Rapid inference and knowledge-based response

Challenges for exascale technologies

- Reliable, secure, high-speed system integration beyond the machine room

- On-demand scheduling to align with human decision taking timelines

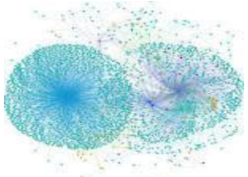

- New computational problems that stress computer architectures in new ways

NITRD

Networking & Information Technology Research and Development Program

Andrew W. Moore (awm@cs.cmu.edu)
Carnegie Mellon, School of Computer Science
Formerly VP Engineering, Google

Good natured provocative remarks: It is, basically, a collision

NITRD

A Big Data Machine Learning Problem

high performance computing

Web News Books Images Videos More Search tools

About 117,000,000 results (0.52 seconds)

Change the Way You Engage - ibmmarketingcloud.com
Ad www.ibmmarketingcloud.com/
The IBM Marketing Cloud Built on Silverpop Engage. Learn More.

High Performance Computing - nsghpc.com
Ad www.nsghpc.com/
Experts in 2-64 node HPC clusters for finance, oil/gas & engineering.
What Is HPC - Operating Systems - Software Tools - Cluster Platforms

Scholarly articles for **high performance computing**
... transformations for **high-performance computing** - Bacon - Cited by 1074
... -scale study of failures in **high-performance computing** ... - Schroeder - Cited by 701
GPU cluster for **high performance computing** - Fan - Cited by 492

High Performance Computing most generally refers to the practice of aggregating **computing** power in a way that delivers much higher **performance** than one could get out of a typical desktop **computer** or workstation in order to solve large problems in science, engineering, or business.

What is high performance computing? - insideHPC
insidehpc.com/hpc-basic-training/what-is-hpc/

More about Supercomputer

Feedback

What is high performance computing? - insideHPC
insidehpc.com/hpc-basic-training/what-is-hpc/
High Performance Computing most generally refers to the practice of aggregating **computing** power in a way that delivers much higher **performance** than one could get

NITRD

A Big Data Machine Learning Problem

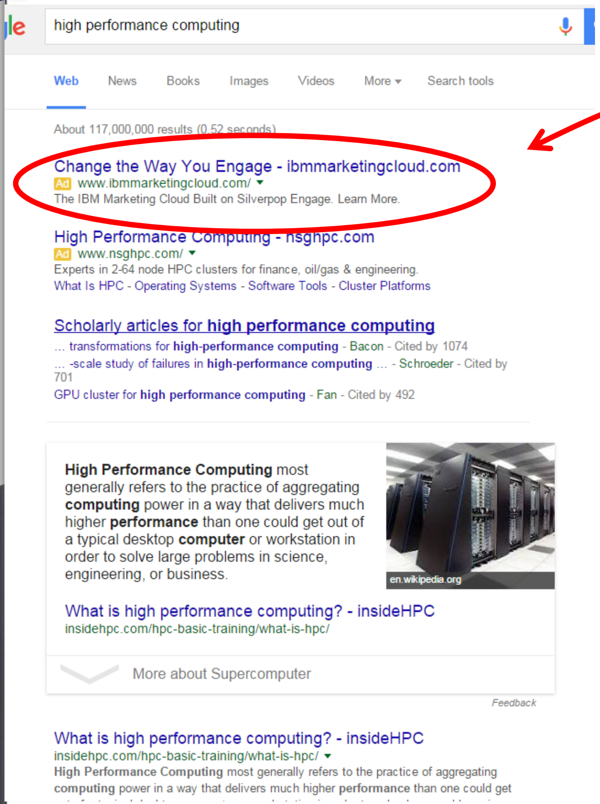

Serving-time task:

$$P(\text{user clicks ad} \mid \text{query} = \text{"hpc" and context})$$

A Big Data Machine Learning Problem

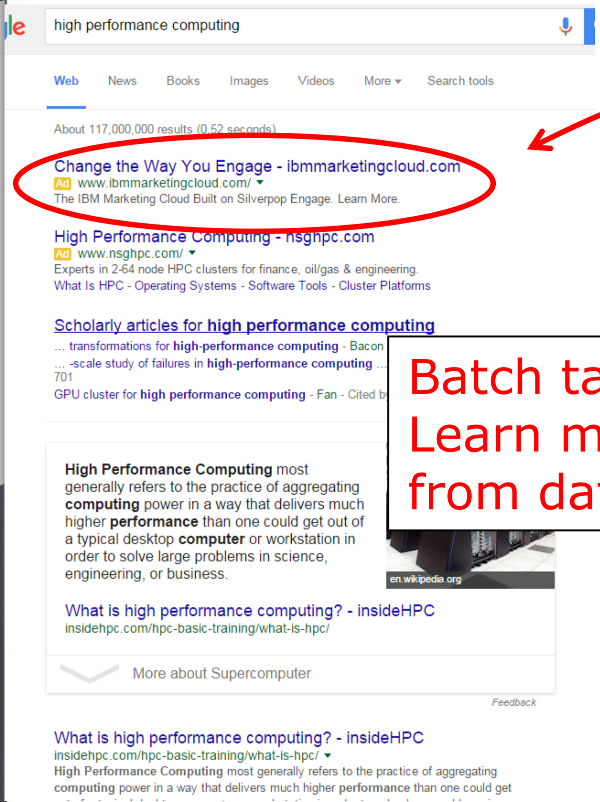

Serving-time task:

$$P(\text{user clicks ad} \mid \text{query} = \text{"hpc" and context})$$

Batch task:
Learn model from data

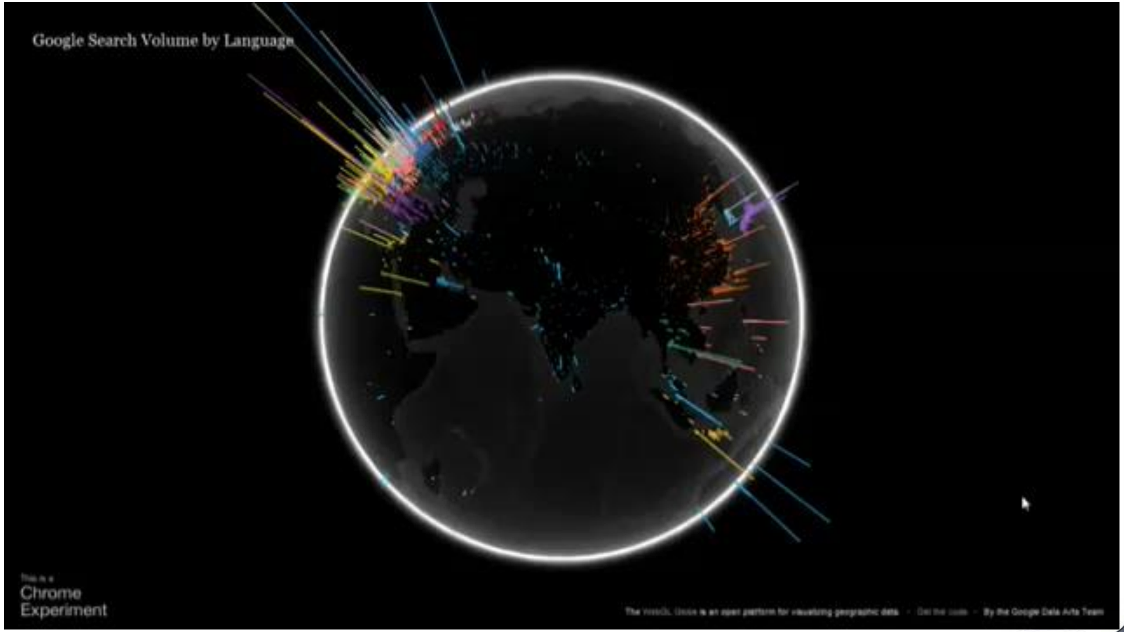

This is Machine Learning

Training Data consists of *R* records:

$$\overbrace{\begin{pmatrix}\boldsymbol{x}_1\\ \boldsymbol{x}_2\\ \vdots\\ \boldsymbol{x}_R\end{pmatrix}}^{\text{Input vectors}} = \overbrace{\begin{pmatrix} x_{11} & x_{12} & \cdots & x_{1M}\\ x_{21} & x_{22} & & x_{2M}\\ \vdots & & \ddots & \\ x_{R1} & x_{R2} & & x_{RM}\end{pmatrix}}^{\text{Each input has M components}} \rightarrow \overbrace{\begin{pmatrix}y_1\\ y_2\\ \vdots\\ y_R\end{pmatrix}}^{\text{Output Observations}}$$

$$\vdots$$

This is Machine Learning

Training Data consists of *R* records:

$$\overbrace{\begin{pmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \\ \vdots \\ \boldsymbol{x}_R \end{pmatrix}}^{\text{Input vectors}} = \overbrace{\begin{pmatrix} x_{11} & x_{12} & \cdots & x_{1M} \\ x_{21} & x_{22} & & x_{2M} \\ \vdots & & \ddots & \\ x_{R1} & x_{R2} & & x_{RM} \end{pmatrix}}^{\text{Each input has M components}} \rightarrow \overbrace{\begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_R \end{pmatrix}}^{\text{Output Observations}}$$

$$\vdots$$

Learning Task:

Find a model represented by a set of *N* weights

$$\boldsymbol{w} = w_1 , w_2 .. w_N$$

which accurately predicts

$$P(y_{new} \mid \boldsymbol{x}_{new} , \boldsymbol{w})$$

This is Machine Learning

Training Data consists of *R* records:

Ou

Ob

$$\underbrace{\begin{pmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \\ \vdots \\ \boldsymbol{x}_R \end{pmatrix}}_{\text{Input vectors}} = \underbrace{\begin{pmatrix} x_{11} & x_{12} & \cdots & x_{1M} \\ x_{21} & x_{22} & & x_{2M} \\ \vdots & & \ddots & \\ x_{R1} & x_{R2} & & x_{RM} \end{pmatrix}}_{\text{Each input has M components}} \rightarrow$$

$$\vdots$$

$$\operatorname*{argmax}_{w} P(\boldsymbol{w} | \boldsymbol{x}_1, \boldsymbol{x}_2, \ldots \boldsymbol{x}_R, y_1, y_2, \ldots y_R)$$

Learning Task:

Find a model represented by a set of *N* weights

$\boldsymbol{w} = w_1 , w_2 .. w_N$

which accurately predicts

$P(y_{new} \mid \boldsymbol{x}_{new} , \boldsymbol{w})$

This is Machine Learning

Training Data consists of *R* records:

Input vectors

Each input has M components

Ou

Ob

$$\overbrace{\begin{pmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \\ \vdots \\ \boldsymbol{x}_R \end{pmatrix}}^{\text{Input vectors}} = \begin{pmatrix} \overbrace{x_{11} & x_{12} & \cdots & x_{1M}}^{\text{Each input has M components}} \\ x_{21} & x_{22} & & x_{2M} \\ \vdots & & \ddots & \\ x_{R1} & x_{R2} & & x_{RM} \end{pmatrix} \rightarrow$$

$\vdots$

$$\underset{w}{\operatorname{argmax}}\, P(\boldsymbol{w}|\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots \boldsymbol{x}_R, y_1, y_2, \ldots y_R)$$

$$=$$

$$\underset{w}{\operatorname{argmax}} \log P(\boldsymbol{w}) + \sum_{i=1}^{R} \log P(y_i|\boldsymbol{x}_i, \boldsymbol{w})$$

Learning Task:

Find a model represented by a set of *N* weights

$\boldsymbol{w} = w_1 , w_2 .. w_N$

which accurately predicts

$P(y_{new} \mid \boldsymbol{x}_{new} , \boldsymbol{w})$

This is Machine Learning

Training Data consists of R records:

Ou

Ob

$$\underbrace{\begin{pmatrix} \boldsymbol{x}_1 \\ \boldsymbol{x}_2 \\ \vdots \\ \boldsymbol{x}_R \end{pmatrix}}_{\text{Input vectors}} = \overbrace{\begin{pmatrix} x_{11} & x_{12} & \cdots & x_{1M} \\ x_{21} & x_{22} & & x_{2M} \\ \vdots & & \ddots & \\ x_{R1} & x_{R2} & & x_{RM} \end{pmatrix}}^{\text{Each input has M components}} \rightarrow$$

$\vdots$

$$\underset{w}{\operatorname{argmax}} P(\boldsymbol{w}|\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots \boldsymbol{x}_R, y_1, y_2, \ldots y_R)$$

$$=$$

$$\underset{w}{\operatorname{argmax}} \log P(\boldsymbol{w}) + \sum_{i=1}^{R} \log P(y_i|\boldsymbol{x}_i, \boldsymbol{w})$$

If w_k is one of the current weights $\boldsymbol{w}$ then a better guess for w_k is

$$w_k \leftarrow w_k + h \frac{\frac{\partial}{\partial w_k} P(\boldsymbol{w})}{P(\boldsymbol{w})} + h \sum_{i=1}^{R} \frac{\frac{\partial}{\partial w_k} P(y_i|\boldsymbol{x}_i, \boldsymbol{w})}{P(y_i|\boldsymbol{x}_i, \boldsymbol{w})}$$

Learning Task:

Find a model represented by a set of N weights

$\boldsymbol{w} = w_1, w_2 .. w_N$

which accurately predicts

$P(y_{new} \mid \boldsymbol{x}_{new}, \boldsymbol{w})$

This is Machine Learning

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_{100,000} \end{pmatrix} \rightarrow \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_{100,000} \end{pmatrix}$$

Core 1

$$\begin{pmatrix} x_{100,001} \\ x_{100,002} \\ \vdots \\ x_{200,000} \end{pmatrix} \rightarrow \begin{pmatrix} y_{100,001} \\ y_{100,002} \\ \vdots \\ y_{200,000} \end{pmatrix}$$

Core 2

⋮

$$\begin{pmatrix} \vdots \\ x_{10^9} \end{pmatrix} \rightarrow \begin{pmatrix} \vdots \\ y_{10^9} \end{pmatrix}$$

Core 10,000

f *R* records:

components

Ou

Ob

$$\begin{matrix} x_{1M} \\ x_{2M} \\ \\ x_{RM} \end{matrix}\Bigg) \rightarrow \Bigg($$

$$\underset{w}{\operatorname{argmax}} P(\mathbf{w} | \mathbf{x}_1, \mathbf{x}_2, \ldots \mathbf{x}_R, y_1, y_2, \ldots y_R)$$

$$=$$

$$\underset{w}{\operatorname{argmax}} \log P(\mathbf{w}) + \sum_{i=1}^{R} \log P(y_i | \mathbf{x}_i, \mathbf{w})$$

If w_k is one of the current weights $\boldsymbol{w}$ then a better guess for w_k is

$$w_k \leftarrow w_k + h \frac{\frac{\partial}{\partial w_k} P(\mathbf{w})}{P(\mathbf{w})} + h \sum_{i=1}^{R} \frac{\frac{\partial}{\partial w_k} P(y_i | \mathbf{x}_i, \mathbf{w})}{P(y_i | \mathbf{x}_i, \mathbf{w})}$$

$\boldsymbol{w} = w_1, w_2 .. w_N$

which accurately predicts $P(y_{new} \mid \boldsymbol{x}_{new}, \boldsymbol{w})$

This is Machine Learning

f R records:

components

Ou

Ob

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_{100{,}000} \end{pmatrix} \rightarrow \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_{100{,}000} \end{pmatrix}$$

Core 1

$$\begin{pmatrix} x_{100{,}001} \\ x_{100{,}002} \\ \vdots \\ x_{200{,}000} \end{pmatrix} \rightarrow \begin{pmatrix} y_{100{,}001} \\ y_{100{,}002} \\ \vdots \\ y_{200{,}000} \end{pmatrix}$$

Core 2

$\vdots$

$$\begin{pmatrix} \vdots \\ x_{10^9} \end{pmatrix} \rightarrow \begin{pmatrix} \vdots \\ y_{10^9} \end{pmatrix}$$

Core 10,000

$$\begin{matrix} x_{1M} \\ x_{2M} \\ \\ x_{RM} \end{matrix} \Bigg) \rightarrow \Bigg($$

$\boldsymbol{w} = w_1, w_2$ ..

which accurately pred

$P(y_{new} \mid \boldsymbol{x}_{new},$

$$\underset{w}{\operatorname{argmax}} P(\boldsymbol{w} \mid \boldsymbol{x}_1, \boldsymbol{x}_2, \dots \boldsymbol{x}_R, y_1, y_2, \dots y_R)$$

$$=$$

$$\underset{w}{\operatorname{argmax}} \log P(\boldsymbol{w}) + \sum_{i=1}^{R} \log P(y_i \mid \boldsymbol{x}_i, \boldsymbol{w})$$

If w_k is one of the current weights $\boldsymbol{w}$ then a better guess for w_k is

$$h \frac{\frac{\partial}{\partial w_k} P(\boldsymbol{w})}{P(\boldsymbol{w})} + h \sum_{i=1}^{R} \frac{\frac{\partial}{\partial w_k} P(y_i \mid \boldsymbol{x}_i, \boldsymbol{w})}{P(y_i \mid \boldsymbol{x}_i, \boldsymbol{w})}$$

Very well suited to racks and racks of boring commodity servers with decent network, GPUs and flash

Use Cases for R=10^{12} records, M=10^{10} input dimensions

Google AdWords

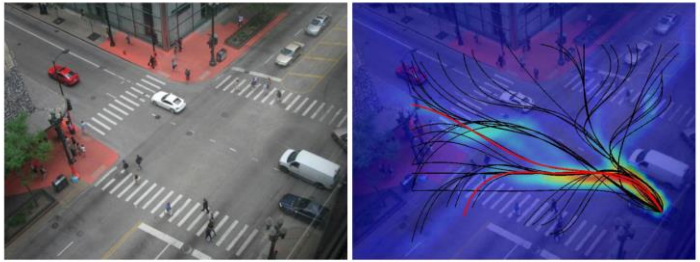

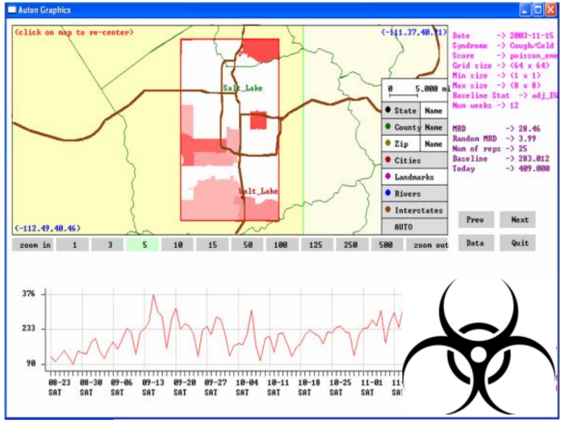

NETFLIX

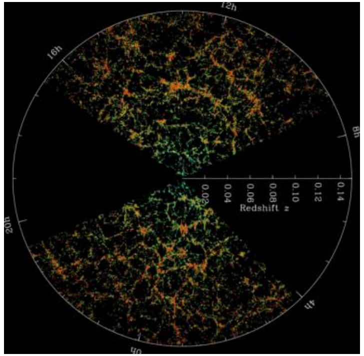

Use Cases for $R=10^{12}$ records, $M=10^{10}$ input dimensions

Google AdWords

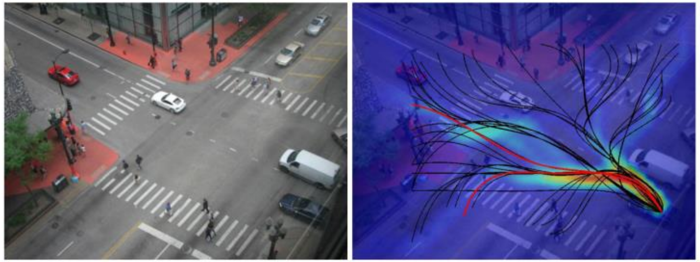

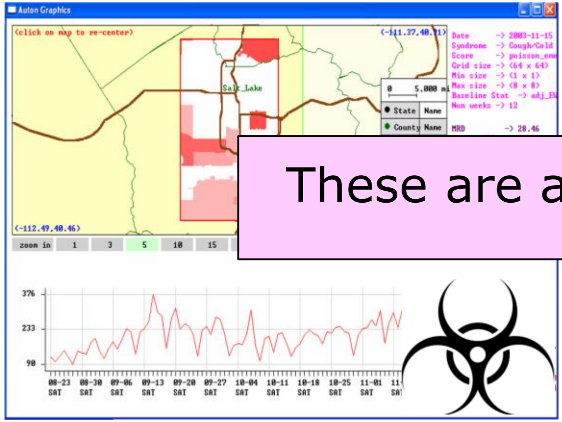

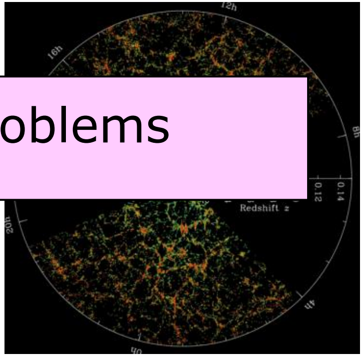

These are all $\underset{\mathbf{w}}{\operatorname{argmax}}\, P(\mathbf{w}|\text{data})$ problems

But there are some with more of a multipole flavor

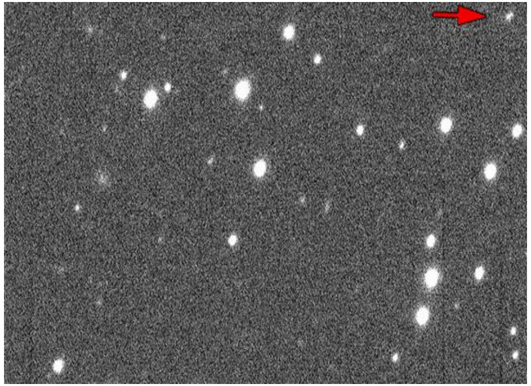

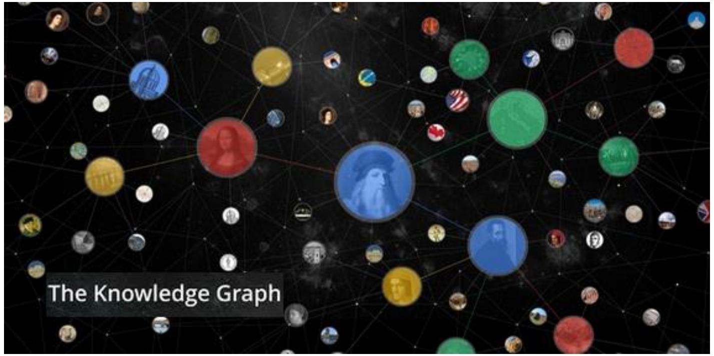

My Opinion

My gut feel on HPC and big data:

- Classic HPC is far removed from what is normally needed for big data

- But big data can use a lot of help with
  - Vector compute at nodes
  - Fast RAM cache over Flash
  - Does NOT need accurate RAM
  - Does NOT need reliable compute nodes

- Classic HPC will be much more important for the big-AI that will be built on the big-data

Questions, Comments, Ideas?
awm@cs.cmu.edu