

Dr. Jeremy V. Kepner
Supercomputing Center Fellow, Founder, and Head

MIT Lincoln Laboratory, 244 Wood St., Lexington, MA, kepner@ll.mit.edu, <http://www.mit.edu/~kepner>
Princeton University, PhD. in Astrophysics (1998); Pomona College, B.A. in Astrophysics (1991)

Personal Statement

Dr. Kepner's career has been distinguished by major contributions to (I) parallel computing, (II) graph algorithms, (III) green supercomputing, and (IV) big data

- (I) His creation of the pMatlab software library and his authorship of the associated best-selling SIAM book demonstrated the viability of parallel Matlab, which is now used by thousands of scientists and engineers worldwide; he led the development of the PVTOL (Parallel Vector Tile Optimizing Library) parallel mathematics and embedded signal processing library that is widely used across government signal processing applications and was selected for an R&D 100 Award; he is a founding sponsor of the Julia language used by millions of programmers; his chairmanship of the IEEE High Performance Extreme Computing conference has helped make it the largest computing conference in New England and a critical venue for parallel computing research.
- (II) His best-selling SIAM book *Graph Algorithms in the Language of Linear Algebra* has become the standard text on matrix-based graph analysis. This work is the foundation of the GraphBLAS.org standard developed by Intel, IBM, and Nvidia to enable computer manufacturers to build new hardware for the purpose of accelerating graph computations. The GraphBLAS.org standard has been integrated into the SuiteSparse library used by millions of programmers worldwide.
- (III) His vision in creating the largest (\$100M, 9 acres, 100,000 square feet, 600 racks, 15,000 servers, 24 megawatts) and greenest (zero carbon emissions) supercomputing center in the world (the Massachusetts Green High Performance Computing Center) has provided supercomputing to one of the largest research communities (5,000+ researchers) in the world. He led the deployment of the largest zero-carbon emission supercomputer on the Top 500.
- (IV) His innovation in developing and patenting a novel database management mathematics, language, and schema enabled big data databases to exploit supercomputers to achieve world-record database performance as described in his MIT Press book: *Mathematics of Big Data*.

Positions and Employment

2016– MIT Lincoln Laboratory Supercomputing Center Founder and Head.
2015– Lincoln Fellow (one of six Fellows in the 60+ year history of MIT Lincoln Laboratory).
2010–2015 Senior Scientist, Computing and Analytics Group, MIT Lincoln Laboratory.
2006–2010 Senior Scientist, Embedded Digital Systems Group, MIT Lincoln Laboratory.
2006– Research Affiliate, MIT Mathematics Department.
2001– Research Affiliate, MIT Computer Science and AI Laboratory.
1998–2006 Scientist, Embedded Digital Systems Group, MIT Lincoln Laboratory.
1993–1996 Department of Energy Computational Science Graduate Fellow at Princeton University

Selected Accomplishments

Most published author in the 60+ year history of the United States' largest Department of Defense research laboratory: Peer-reviewed publications span the application of interactive supercomputing and big data technologies to astronomy, astrophysics, signal processing, radar, sonar, image processing, plasma physics, data mining, databases, graph algorithms, cyber security, visualization, cloud computing, random matrix theory, abstract algebra, and health sciences; many of these works have received best paper awards.

Supercomputing Center Founder: Co-inventor of the Massachusetts Green HPC Center (MGHPCC), the largest and greenest open-research supercomputing center in the world; founder and head of the MIT Lincoln Laboratory Supercomputing team that designs and delivers supercomputing data centers, hardware, software, and training to thousands of researchers across MIT.

Julia Language: founding sponsor of the Julia programming language used by millions.

Parallel Matlab: basis of MathWorks parallel computing solution used by thousands of scientists and engineers worldwide; author of the SIAM bestselling book on Parallel Matlab.

PVTOL (Parallel Vector Tile Optimizing Library): parallel mathematics and embedded signal processing library widely used across government signal processing applications and selected for an R&D 100 Award.

Graph Algorithms: Co-creator of GraphProcessor hardware and GraphBLAS.org software standard with Intel, IBM, and Nvidia based on the bestselling SIAM book *Graph Algorithms in the Language of Linear Algebra* by Dr. Kepner and Prof. Gilbert; deployed in the SuiteSparse library used by millions of programmers.

Big Data: World record holder in graph database ingest performance; achieved 100,000,000 database inserts per second using Apache Accumulo and patented software (d4m.mit.edu) written by Dr. Kepner; basis of forthcoming book *Mathematics of Big Data* from MIT Press.

HPC Challenge: Co-chair (with Prof. Jack Dongarra) of a multi-institutional DARPA team that developed several HPC procurement benchmarks (e.g., HPCChallenge.org) widely used for acceptance testing of many of the worlds' largest supercomputers.

3D Data Visualization: Pioneer in using 3D gaming and printing technology to visualize supercomputing operations, cyber computer networks, and social media data.

Center for Engaging Supercomputing: Founder and co-director of center focused on delivering interactive supercomputing capabilities.

SuperCloud: Principal Investigator of DoD-funded project to integrate HPC, big data, database, and virtualized environments; SuperCloud is operationally deployed on all Lincoln Laboratory Supercomputing Center systems.

DoD HPC Modernization Program: Principal Investigator for DoD HPC investments in Persistent Surveillance Supercomputing and Interactive Algorithm Development for Weapons of Mass Destruction Defense.

Dr. Kepner has an excellent record of mentoring. During the past ten years, he has mentored 25 undergraduate, master's, and PhD students. Many of his students are now faculty or scientists at leading institutions (CalTech, Univ. of Chicago, Amherst College, Georgetown, Lawrence Berkeley Lab, and Lawrence Livermore Lab). Seven of his undergraduate students have been awarded graduate Fellowships (NSF, NDSEG) to pursue their PhDs.

Books Authored

Mathematics of Big Data: Spreadsheets, Databases, Matrices, and Graph, MIT Press, 2018.

Graph Algorithms in the Language of Linear Algebra, 2011 best seller, SIAM series on software, environments and tools (series editor: Prof. Jack Dongarra); basis of the GraphBLAS.org standard.

Parallel Matlab for Multicore and Multinode Computers, 2009 best seller, SIAM series on software, environments and tools (series editor, Prof. Jack Dongarra).

Special Issues Edited

HPCS and the Path Towards Usable Petascale Computing: User Productivity Challenges, CT Watch, Vol 2, Number 4, November 2006.

International Journal of HPC Applications: Special Issue on High Level Programming Languages and Models, Winter 2006 (November).

International Journal of HPC Applications: Special Issue on HPC Productivity, Volume 18, Number 4, Winter 2004 (November).

Patents

Extreme Virtual Memory, U.S. Patent #9,852,079 (2017).

System and Method of Masking and Computing on Masked Data in a Data Store, Provisional Patent (2016).

Multidimensional Associative Array Database, U.S. Patent #8,631,031 (2011).

Online Courses

Signal Processing on Databases, MIT Open Courseware, 20K+ views since online debut in summer of 2016.

Honors and Awards

2017 Lincoln Team Award for deploying the largest supercomputer in New England, the 3rd largest at a U.S. university, and the largest zero-carbon emissions supercomputer on the Top 500.

2017 *R&D Magazine* R&D 100 Award Finalist for BigDAWG federated database system.

2014 MIT Lincoln Laboratory Technical Excellence Award (Lincoln Laboratory's highest technical award) "For his leadership and vision in bringing supercomputing to Lincoln Laboratory through the establishment of [LLSC]; his pivotal role in open systems for embedded computing; his creativity in developing a novel database management language and schema; and his contributions to the field of graph analytics."

2013 Elected to St. Anthony Village High School Hall of Fame for outstanding post-graduate accomplishment.

2012 Lincoln Team Award for Green Supercomputer Deployment.

2011 *R&D Magazine* R&D 100 Award for PVTOL (Parallel Vector Tile Optimizing Library) parallel mathematics and embedded signal processing library widely used across government signal processing applications.

- 2008 Lincoln Team Award for Office of Naval Intelligence Experiment.
- 2006 Lincoln Team Award for DoD Supercomputer Deployment.

Affiliations and Service

- 2015–2017 Chair, SIAM Data Mining Activity Group.
- 2017 Co-chair, IEEE Big Data Conference.
- 2014–2015 Vice-Chair, SIAM Data Mining Activity Group.
- 2013 Co-chair, SIAM Symposium on Big Databases for Big Computations.
- 2012– Chair, IEEE High Performance Extreme Computing (HPEC) Conference.
- 2012 Co-chair, SIAM Symposium on Massive Graphs: Big Compute meets Big Data.
- 2011– MIT SIAM section co-founding faculty advisor.
- 2011– IEEE Senior Member.
- 2009 Co-chair, SIAM Symposium on High Performance Computing on Massive Real-World Graphs.
- 2008 Co-chair, SIAM Symposium on High Performance Computing on Large Graphs.
- 2008 Co-chair, SIAM Symposium on Parallel Computing with Matlab.
- 2006–2011 IEEE member.
- 2006– SIAM member.
- 2005– Co-chair and co-founder, HPC Challenge benchmark suite and contest; HPC Challenge is used to validate most supercomputing systems in the world.
- 2001–2011 Technical Chair, High Performance Embedded Computing Workshop.

Contributions to Parallel Programming, Graph Algorithms, Supercomputing, and Big Data (Selected from 100+ publications authored or about Dr. Kepner contributions)

I. Parallel Programming

Dr. Kepner has been a world leader in creating software tools such as pMatlab, Parallel VSIPL++ and PVTOL that allow users to interact with parallel computing systems in a way to meet their needs. pMatlab was the first parallel Matlab implementation that demonstrated the viability of parallel Matlab and directly lead to the Mathworks Parallel Computing Toolbox which is used by thousands of scientists and engineers today. PVTOL received an R&D 100 award in 2011.

1. *Parallel MATLAB: The next generation*, Kepner & Travinin, HPEC 2003.
2. *MatlabMPI*, Kepner & Ahalt, Journal of Parallel and Distributed Computing, 2004.
3. *pMatlab Parallel Matlab Library*, Bliss & Kepner, International Journal of High Performance Computing Applications, 2004.
4. *Parallel VSIPL++: An open standard software library for high-performance parallel signal processing*, Lebak, Kepner, Hoffmann & Rutledge, Proceedings of the IEEE, 2005.
5. *Parallel Matlab for Multicore and Multinode Computers*, Kepner, SIAM Press, 2009.
6. R&D 100 Award Winners, R&D Magazine, 2011.

II. Graph Algorithms

Dr. Kepner has been a pioneer in applying matrix mathematics to graph problems. This work is the focus of his second SIAM book *Graph Algorithms in the Language of Linear Algebra*. The sparse linear algebra approach to graphs is gaining wide acceptance and is the basis of the GraphBLAS.org standard initiated by Intel, IBM, and Nvidia.

7. *Standards for Graph Algorithm Primitives*, Mattson, Bader, Berry, Buluc, Dongarra, Faloutsos, Feo, Gilbert, Gonzalez, Hendrickson, Kepner, Leiserson, Lumsdaine, Padua, Poole, Reinhardt, Stonebraker, Wallach, & Yo, IEEE/SIAM HPEC, 2013.
8. *Graphulo Implementation of Server-Side Sparse Matrix Multiply in the Accumulo Database*, Hutchison, Kepner, Gadepally & Fuchs, IEEE/SIAM HPEC 2015, Best Student Paper Award Finalist (top 10% of ~100 submissions).
9. *From NoSQL Accumulo to NewSQL Graphulo: Design and Utility of Graph Algorithms inside a BigTable Database*, Hutchison, Kepner, Gadepally & Howe, IEEE/SIAM HPEC 2016, Outstanding Student Paper Award, (top 5% of ~100 submissions).

III. Supercomputing and the Massachusetts Green High Performance Computing Center (MGHPCC)

Dr. Kepner has been a leader in the field of supercomputing benchmarking, supercomputing integration with big data, and large-scale computer monitoring and visualization. In 2005, Dr. Kepner first developed and championed the idea of placing a large supercomputing center next to a hydroelectric dam on the Connecticut River to take

advantage of its green energy. In 2012, the \$100M MGHPCC was completed and provides capacity for over 20,000 computer servers to Harvard, MIT, University of Massachusetts, Boston University, and Northeastern University.

10. *HPC productivity: An overarching view*, Kepner, International Journal of High Performance Computing Applications, 2004.
11. *The HPC Challenge (HPCC) Benchmark Suite*, Luszczek, Bailey, Dongarra, Kepner, Lucas, Rabenseifner & Takahashi, Proceedings of the ACM/IEEE Conference on Supercomputing, 2006.
12. *Designing scalable synthetic compact applications for benchmarking high productivity computing systems*, Bader, Madduri, Gilbert, Shah, Kepner, Meuse & Krishnamurthy, Cyberinfrastructure Technology Watch, 2006.
13. *Large Scale Network Situational Awareness Via 3D Gaming*, Hubbell & Kepner, IEEE/SIAM HPEC, 2012 (see also popular press coverage in: *The real Tron: IT security as a shoot 'em up*, Hodson, New Scientist, Oct 3, 2012; *MIT's 'game' that lets you patrol a computer network*, Critical Gamer, 2012).
14. *Benchmarking SciDB Data Import on HPC Systems*, Samsi, Brattain, Arcand, Bestor, Bergeron, Byun, Gadepally, Hubbell, Klein, Michaleas, Milechin, Mullen, Prout, Rosa, Yee, Kepner & Reuther, IEEE/SIAM HPEC 2016, Best Paper Award Finalist (top 5% of ~100 submissions).
15. *Getting up to Supercomputing Speed in Holyoke*, Hedayat-Zadeh, University of Massachusetts, May, 2012
16. *Massachusetts Miracle*, Fleschner, Pomona College Magazine, Spring 2013.
17. *Green Computing is super: MIT Lincoln Laboratory's new supercomputing facility reduces energy impacts*, Ryan, MIT News, July 18, 2013.

IV. Big Data Supercomputing

Dr. Kepner's team set several world records in database performance by deploying big data technologies on supercomputers. Of particular importance is his role in developing the Accumulo Database and its widely used D4M schema (which was developed by the NSA to handle their massive data storage and retrieval challenges) and is the basis of a new class of polystore databases and his forthcoming MIT Press book *Mathematics of Big Data*. In addition, Dr. Kepner has recently patented novel mathematics and techniques for operating on big data while it is in an encrypted form.

18. *Multidimensional Associative Array Database*, Kepner, U.S. Patent #8,631,031 (2011).
19. *Dynamic Distributed Dimensional Data Model (D4M) Database and Computation System*, Kepner, Arcand, Bergeron, Bliss, Bond, Byun, Condon, Gregson, Hubbell, Kurz, McCabe, Michaleas, Prout, Reuther, Rosa & Yee, International Conference on Acoustics, Speech, and Signal Processing, 2012.
20. *D4M 2.0 Schema: A General Purpose High Performance Schema for the Accumulo Database*, Kepner, Anderson, Arcan, Bestor, Bergeron, Byun, Hubbell, Michaleas, Mullen, O'Gwynn, Prout, Reuther, Rosa, & Yee, IEEE/SIAM HPEC 2013.
21. *Improving Big Data Visual Analytics with Interactive Virtual Reality*, A. Moran, V. Gadepally, M. Hubbell & J. Kepner, IEEE/SIAM HPEC 2015, Best Student Paper Award Finalist (top 10% of ~100 submissions).
22. *Achieving 100,000,000 database inserts per second using Accumulo and D4M*, Kepner, Arcand, Bestor, Bergeron, Byun, Gadepally, Hubbell, Michaleas, Mullen, Prout, Reuther, Rosa & Yee, IEEE/SIAM HPEC, 2014.
23. *The BigDAWG polystore system*, Duggan, Elmore, Stonebraker, Balazinska, Howe, Kepner, Madden, Maier, Mattson & Zdonik, ACM Sigmod Record, 2015.
24. *Genetic Sequence Matching Using D4M Big Data Approaches*, Dodson, Ricke & Kepner, IEEE/SIAM HPEC 2014, Best Paper Award (#1 of ~100 submissions).
25. *System and Method of Masking and Computing on Masked Data in a Data Store*, Patent (provisional), 2016.
26. *Julia Implementation of the Dynamic Distributed Dimensional Data Model*, Chen, Edelman, Kepner, Gadepally & Hutchison, IEEE/SIAM HPEC 2016, Best Paper Award (#1 of ~100 submissions).

Selected Additional Publications

Radially global gyrokinetic simulation studies of transport barriers, Parker, Mynick, Artun, Cummings, Decyk, Kepner, Lee & Tang, Physics of Plasmas, 1996.

An automated cluster finder: the adaptive matched filter, Kepner, Fan, Bahcall, Gunn, Lupton & Xu, The Astrophysical Journal, 1998.

Detecting clusters of galaxies in the Sloan Digital Sky Survey, Kim, Kepner, Postman, Strauss, Bahcall, Gunn, Lupton, Annis, Nichol, Castander, Brinkmann, Brunner, Connolly, Csabai, Hindsley, Ivezić, Vogeley & York, The Astronomical Journal, 2002.

pMapper: Automatic Mapping of Parallel MATLAB Programs, Travinin, Hoffmann, Bond, Chan, & Wong, HPEC Workshop, 2005, Best Paper Finalist (top 10% of ~50 submissions).

Exploring the Cell with HPEC Challenge Benchmarks, Sacco, Schrader & Kepner, HPEC Workshop, 2006, Best Paper Finalist (top 10% of ~50 submissions).

Language, Dialect, and Speaker Recognition Using Gaussian Mixture Models on the Cell Processor, Malyska, Mohindra, Reynolds & Kepner, HPEC Workshop, 2008, Best Paper Finalist (top 10% of ~50 submissions).

Automatic registration of LIDAR and optical images of urban scenes, Mastin, Kepner & Fisher, IEEE Computer Vision and Pattern Recognition, 2009.

3-D Graph Processor, Song, Kepner, Nguyen, Kramer, Gleyzer, Mann, Horst, Retherford, Bond, Bliss, Robinson, Mohindra & Mullen, HPEC Workshop, 2010, Best Paper (#1 of ~50 submissions).