

## Michael Allen Heroux

Center for Computing Research, Sandia National Laboratories,

+1 505 379 5518, maheou@sandia.gov

Department of Computer Science, Saint John's University,

+1 320 363 3394, mheroux@csbsju.edu

March 2017

### Experience

**Senior Scientist.** Center for Computing Research Sandia National Labs, May 1998-present. Principle member 1998-2004, Distinguished member 2004-2016, Senior Scientist 2016 – present. Conduct research and development of numerical methods for scientific and engineering applications on large-scale parallel computers. Participate on program and standards committees in areas of expertise. Lead the Trilinos libraries project, the Mantevo applications performance-modeling project and HPCG Benchmark.

**Scientist in Residence and Adjunct Faculty Member.** Department of Computer Science, Saint John's University, September 1998-present, Scientist in Residence 2004-present. Teach courses in Numerical Analysis, Parallel Computing, Computer Science Research Methodologies and Software Engineering. Direct undergraduate research theses in parallel computing and related areas. Participate in curriculum development.

**Group Leader.** Scalable Computing, Algorithms and Capability Prototyping Groups, SGI/Cray Research, March 1995-May 1998. Led a team of specialists in scientific computing. Directed activities and participated in development, porting and optimization of large-scale parallel applications for SGI/Cray systems. Participated in and led standardization efforts for scientific computing. Led efforts in development of new application capabilities. Provided applications analysis and requirements to future computer systems development including the Cray T3E, T90, J90, SV1 and SV2.

**Numerical Analyst.** CFD Group, Engineering Applications, Cray Research, September 1993-February 1995. Responsible for research and development of numerical methods for engineering applications in CFD, structural analysis, electronics and reservoir simulation. Worked with application developers on Cray vector multiprocessors and distributed memory machines. Particular areas of interest were the solution of sparse and dense linear systems, iterative methods, parallel algorithms and large-scale scientific computation. Served as consultant on numerical methods for Cray Research customers and application specialists.

**Numerical Analyst.** Mathematical Software Research Group, Cray Research, October 1988-September 1993. Conducted research and development of numerical linear algebra libraries. Served as consultant on numerical methods for Cray Research customers and application specialists. Developed libraries of high-performance software for Cray Research computer systems.

## **Education**

Ph.D. Mathematics. May 1989, Colorado State University, Fort Collins, Colorado.

M.S. Mathematics. August 1986, Colorado State University, Fort Collins, Colorado.

B.A. Mathematics. December 1983, Saint John's University, Collegeville, Minnesota.

## **Professional Awards**

- 2015 HPCWire “People to Watch 2015” selection.
- 2014 FLC Regional Technology Transfer Award for Mantevo.
- R&D 100 Award for Mantevo 1.0, 2013, project initiator and leader.
- Best Poster Award, SC11 Conference, November 2011.
- Distinguished Member of the Association for Computing Machinery, October 2009.
- ASC Salutes Profile, NNSA/ASC profile, September 2007.
- R&D 100 Award for Trilinos 3.1, 2004, project initiator and leader.
- SC2004 HPC Software Challenge Award, 2004.
- Member of Cray Research Gordon Bell Prize Finalist Team, 1996.
- Sandia Employee Awards:
  - Mantevo Team R&D 100 Award and External Impact on HPC Co-Design, 2013.
  - Winning X-caliber proposal for the DARPA/UHPC Program, 2011.
  - Educating the next generation of computational scientists, 2010.
  - IAA Algorithms Team, 2009.
  - Organizing Next-generation Applications Workshop, 2008.
  - Xyce/Charon/Algorithms Team, 2008.
  - Supercomputing Architecture & Programming Environment Team, 2008.
  - Leading Trilinos 7.0 Release, 2006.
  - Leadership of Trilinos Project, 2004.
  - Xyce Development, 2004.
  - Efforts in Nanosciences Initiative, 2003.
  - Algorithms for Circuit Simulation, 2001.
  - Parallel Circuit Simulation Code, 2000.

## **Professional Leadership**

- Technical Papers Chair, Supercomputing 2017 Conference.
- Gordon Bell Prize Committee member, 2016 – 2021.
- Chair, SC16 Test of Time Award Committee.
- Scientific Libraries Lead, DOE Exascale Computing Project, 2016 – present.
- Chair of the NITRD, multi-agency workshop on Computational Science and Engineering Sustainability and Software Productivity (CSESSP) Challenges., October 15 – 16, 2015, Washington, DC.
- Principle architect and developer of the HPCG benchmark code.
- Editor-in-Chief, ACM Transactions on Mathematical Software, 2010 – 2017.
- Created the Replicated Computational Results review for ACM Transactions on Mathematical Software, 2013 – 2015.
- Member of SC Conference Test of Time Award Committee, 2014 – 2015.
- Applications Program Chair, SC13 Technical Program, 2013.
- Editor, SIAM Book Series on Software, Environments and Tools, 2012-present.
- Associate Editor, SIAM Journal on Scientific Computing, 2010-present.
- Subject Area Editor, Journal on Parallel and Distributed Computing, 2011-present.
- Lead writer of Software Section in the International Exascale Software Project (IESP), 2011.
- Chair of DOE Application readiness review for Titan 20PF computer system, 2010.
- Created Career and Junior Scientist Awards for SIAM SIAG/SC, 2009-2010.
- Led SIAG/SC committee to select Career/Junior Scientist winners, 2009-2010.
- Led SIAG/SC committee to select 2010-2011 officers, 2009.
- Wrote whitepaper for NSF on sustainable software engineering, 2009.
- Member, International Exascale Software Project (IESP), 2008-present.
- Sandia rep, DOE/ASCR Breakthroughs Report, 2009.
- Sandia PI, The Exascale Software Center (ESC), 2010-present.
- Sandia PI, The SciDAC-2 TOPS-2 project, 2005-present.
- Sandia PI, The Extreme-scale Algorithms & Software Institute (EASI), 2009-present.
- Sandia PI, Institute for Advanced Architectures & Algorithms, 2008-present.
- Associate Editor for SIAM Journal on Scientific Computing, Jan 2010-present.
- Chair of the SIAM Supercomputing Special Interest Group, 2008-2009.

- Program Director for SIAM Supercomputing Special Interest Group, 2000-2003.
- Program Chair for 2004 SIAM Parallel Processing Conference.

## **Professional Service**

- Program committee member (past 3 years): SIAM PP18, SC16, SC15, CCGrid 2015, EX14, ICCS 2014, ICS 2014, ScalA 2016, 2015, SEHPCCSE17, 16, WACCPD14.
- Visiting committee member, CEA high performance computing review.
- Advisory Board Member, NSF-funded FLAME project, U of Texas at Austin.
- Reviewer for NSF in computational science and scalable computing, 2003-present.
- PhD committee, Daniel Sunderland, University of Utah, 2014-present.
- PhD committee, Fan Ye, Masion de Simulation, Saclay-Paris, 2014-present.
- PhD committee, France Boillod-Cerneax, University of Lille, Paris, 2014.
- PhD committee, Radu Popescu, Ecole Polytechnique Federale de Lausanne, 2012-2013.
- PhD committee, Sarah Knepper, Emory University, 2010-2011.
- PhD committee, Bryan Marker, University of Texas at Austin, 2010-2011.
- Referee for SIAM Journal of Scientific Computing, SIAM Review, ACM Transactions on Mathematical Software, IEEE Transactions on Parallel and Distributed Systems, 1999-present.

## **Professional Memberships**

- The Society for Industrial and Applied Mathematics.
- Distinguished Member, The Association for Computing Machinery.

## **Community Contributions and Impact**

- **Community scientific software research and development:**
  - I initiated the Trilinos project 15 years ago as an effort to produce a compatible collection of independently developed mathematical software tools.
  - I have led the Trilinos project through several distinct transition phases: First, the expansion of project functionality beyond solvers; then the transition from a Sandia-centric project to one that includes non-Sandians as first class developers; then the (still ongoing) transition to scalable manycore, accelerator and heterogeneous systems and now the expansion of the Trilinos ecosystem to include non-native packages.
  - Trilinos has grown from the original 3 packages to 60 and it represents the single largest active mathematical software project in the world.
  - Trilinos has user communities across the world. We host several tutorial events and host annual user group meetings in the US and Europe.
  - Trilinos provides the collaboration and delivery framework for many Sandia activities including most ASC algorithms efforts, numerous LDRDs, Office of Science projects and CRADAs.
- **Numerical Linear Algebra:**
  - Underlapping for domain decomposition: I developed the original concept of using “underlapped” subgraphs and subdomains that permit the use of standard preconditioners for communication avoiding (s-step) Krylov solvers. This idea led to the first practical preconditioning strategy for CA iterative methods.
  - Complex linear systems solution methods: I developed (with David Day) new algorithms and spectral theory for solving complex-valued linear systems via equivalent real formulations. These formulations permit the use of commonly available real-valued math software for which there is no similar complex-valued version (which is very common). This work has been wide used, especially in the electromagnetics community.

- **Proxy applications for HPC co-design:**
  - I implemented the first miniapp (HPCCG) and led the Mantevo project from the beginning. With colleagues at Sandia I demonstrated the value of miniapps in co-design activities.
  - Every co-design effort across DOE uses miniapps (or more generally proxy apps) in the way that Mantevo does.
  - Mantevo miniapps are cited in more than 200 publications over the past four years.
  - Mantevo has expanded to be an international community project with contributions from 7 institutions outside of Sandia, and a community web portal at <http://www.mantevo.org>.
  - Mantevo 3.0 was released during SC'14, including 16 packages, five of which are new since release 2.0 from a year ago.
  
- **Resilience:**
  - I have developed a taxonomy for application-driven resilient computing models, including two new approaches that I pioneered:
    - **Local-failure-local-recovery (LFLR):** This approach promotes a recovery model whose cost and scope is proportional to scope of failure. LFLR has emerged as a promising next step for practical application resilience for cases where global checkpoint-restart is costly or infeasible.
    - **Selective Reliability:** This model permits application and library developers to declare data and compute regions to be more (or less) reliable than the default execution environment. This kind of selectivity enables the development of new algorithms where the majority of data and computation are in low reliability mode, but some portion of data and computation are in high reliability mode, ensuring the resilience of application execution.
  - I have developed and promoted an additional model, relaxed bulk synchronous parallel (rBSP), that some application teams have implemented as a way to mitigate performance variability on emerging systems.
  - LFLR is integrated into ASC product R&D plans and its scalable recovery has been demonstrated on more than 10,000 processes.

- **Community benchmark for HPC Systems:**
  - Two years ago I started a new benchmarking effort called HPCG at the request of NNSA, in collaboration with Jack Dongarra. I initiated the HPCG strategy to complement the LINPACK benchmark for the TOP 500 list.
  - I am the architect and implementer of the reference version of the benchmark code and have worked directly with community members and vendors on design and implementation features.
  - I organized a series of community meeting to build understanding and acceptance of the benchmark in the international community.
  - We have announced two lists of results, first at ISC'14 with 15 of the top machines participating, then at SC'14 with 25 machines on the list.
  - All major computer vendors have an optimized version of HPCG, displaying the critical features we want probed for future systems.
  - HPCG has received considerable coverage in the HPC press and is the subject of more than 20 publications since its release one year ago.
  
- **Scientific Productivity:**
  - I have participated in the definition, scoping and strategic discussions for a productivity-focused approach to advancing computational science and engineering.
  - I have participated in six workshops on productivity over the past two years and the writing of 3 DOE reports.
  - I was the keynote speaker at an inter-agency workshop on productivity in August 2014 and an invited member for an SC'14 panel on scientific productivity.
  - I lead (with Lois McInnes, Argonne and David Moulton, LANL) the IDEAS Project, the first DOE project funded on scientific productivity.
  - Trilinos project efforts on software engineering processes, lifecycles and community collaboration models serve as a reference for the IDEAS project.



## Publicly-available Software

- **The Trilinos Project (trilinos.org):** Open Source (LGPL/BSD), Initiated and lead the project, 2001-present. *Trilinos is a 2004 R&D 100 winner and the world's largest open source computational science and engineering libraries project.* It is a collection of nearly sixty open source software packages supported by a common software engineering infrastructure and community development model.
  - **Trilinos package development:** Each Trilinos package is a self-contained software product with its own scope of development. These are the packages I have designed and developed:
    - **Epetra:** Principal designer and implementer. Epetra is the predecessor to Tpetra and is one of the two most popular scalable data class libraries on the planet (PETSc is the other). Epetra is used by thousands of application and library developers for constructing and using scalable sparse and dense linear algebra objects.
    - **AztecOO:** Principal designer and implementer. An object-oriented version of the popular Aztec linear solver library. AztecOO is the most widely used iterative solver package in Trilinos, used by thousands of people, providing the core linear solver capabilities for many Sandia and DOE applications.
    - **Tpetra and Kokkos:** Initial designer and developer; remain an algorithm designer and funding source.
    - **Ifpack:** Principle designer and implementer. A collection of algebraic sparse preconditioners and smoothers. Widely used in Sandia and DOE applications.
    - **Ifpack2:** Initial designer. Next-generation of Ifpack targeting scalable manycore architectures.
    - **Amesos:** Principal designer. A package of interfaces to common direct sparse solvers. Widely used at Sandia and other DOE labs.
    - **Amesos2:** Initial designer. Next-generation of Amesos targeting scalable manycore architectures.
    - **Belos:** Designer. Follow-on to AztecOO as a collection of scalable, state-of-the-art iterative methods.
    - **Komplex:** Principal designer and implementer. A package of solvers for complex-valued systems using equivalent real formulations.
    - **Teuchos:** Designer and developer. The core services package in Trilinos. Widely used.
- **The Mantevo Project (mantevo.org):** Open Source (LGPL), Initiated and continue to lead the project, 2006-present. *Mantevo is the first project to concretely define the*

**concept of a miniapplication as a co-design vehicle for next generation applications and computer systems.** Mantevo is a collection of 16 open-source, stand-alone miniapplications that serve as performance proxies for Sandia's large-scale applications.

- **Mantevo package development:** Each Mantevo miniapplication is a self-contained software product. These are the packages that I have designed and developed:
  - **HPCCG:** Principal designer and developer. Performance proxy for a scalable finite-volume/finite-difference single physics PDE application. ***HPCCG has been used in dozens of performance studies for new system design. Rewritten 6 times using new programming languages and programming models.***
  - **MiniFE:** Designer and developer. Follow-on to HPCCG as a proxy for unstructured finite element single physics applications. Used to prototype manycore algorithms and parallel pattern implementations that are now in production use in Trilinos. Used in numerous systems performance studies on mixed precision and hybrid MPI+threading programming environments.
- **Tramonto ([software.sandia.gov/tramonto](http://software.sandia.gov/tramonto)):** Open Source (LGPL), Lead scalable algorithms designer and developer, 2004-present. Tramonto is an open source application for modeling and simulation of inhomogeneous fluids using classical density functional theories. ***Tramonto has unique modeling capabilities for a wide variety of applications, including biophysics applications for new pharmaceuticals based on anti-microbial peptides.***
- **Aztec ([www.cs.sandia.gov/CRF/aztec1.html](http://www.cs.sandia.gov/CRF/aztec1.html)):** Open Source (Special license), Lead developer, 1998-2000. ***Popular open source preconditioned iterative solver package that is still download frequently (250 downloads this year).***
- **Sparse BLAS ([math.nist.gov/spblas](http://math.nist.gov/spblas)):** Open Source (no license), Lead designer, 1999-2002. The sparse BLAS are a *de facto* standard for sparse kernel computations.
- **BPKIT ([sourceforge.net/projects/bpkit](http://sourceforge.net/projects/bpkit)):** Open Source (LGPL), Lead designer, 1995-1996. BPKIT was one of the first object-oriented math software packages, and it remains a popular prototyping environment for preconditioned iterative methods.
- **GEMMW ([www.mgnet.org/~douglas/ccd-free-software.html](http://www.mgnet.org/~douglas/ccd-free-software.html)):** Open Source (no license), Developer, 1994. GEMMW is a portable parallel implementation of Strassen-Winograd dense matrix-matrix multiplication.
- **Cray Sparse Solvers:** Distributed with Cray Scientific Libraries (LIBSCI), Principal designer and developers of the preconditioned sparse solvers, 1989-1993. ***Provided optimized libraries for sparse linear systems on Cray vector multiprocessor and MPP machines.***

- **Cray Optimized BLAS/LAPACK:** Distributed with LIBSCI, developer of YMP/C90 kernels for vector multiprocessor systems, 1989-1993. *Developed unique hybrid implementation for single vector processor and multiple vector processors.*
- **Cray vectorized tridiagonal solvers:** Distributed with LIBSCI, principal developer, 1989-1993. *Developed 3:1 cyclic reduction and burn-at-both-ends algorithms for vector processors.*

## Selected Invited Presentations

- **Keynote: *Productive and Sustainable: More Effective CSE***, SIAM Conference on Computational Science and Engineering 2017, Atlanta, GA, February 2017.
- **Keynote: *Strategies for Next Generation HPC Applications and Systems***, ACSI Conference 2016, Fukuoka, Japan, January 2016.
- **Keynote: *A Task-centric/Dataflow Application Architecture for Scalable Systems***, SCALA Workshop 2015, SC'13, Austin, TX, November 2015.
- **Keynote: *Effective use of Miniapps for co-design***. WRAP Workshop, IEEE Cluster, September 2015.
- **Invited: *Efficiency or Productivity: Pick One***. Panelist. SC'14, New Orleans, LA, November 2014.
- **Tutorial: *Scalable Manycore Computing for Sparse Computation***, SC'11, SC'12, SC'13, SC'14.
- **Invited: *Improving Scientific Productivity: Practical Approaches Toward an Elusive Goal***, CEA Invited Talk, Paris, France, October 2014.
- **Keynote: *Productivity for Productivity***, Inter-agency Productivity Workshop, University of Indiana, Bloomington, IN, August 2014.
- **Keynote: *Challenges and Opportunities for Scalable Finite Element Setup & Assembly***, FE Assembly Workshop, Albuquerque, NM, May 2014.
- **Invited: *System Software: A Necessary but Ill-prepared Hero***, Salishan Conference, Salishan, OR, April 2014.
- **Invited: *Toward the Next Generation of Parallel and Resilient Algorithms & Libraries***, Advances in Numerical Algorithms and High Performance Computing, University College London, England, April 2014.
- **Keynote: *Toward the Next Generation of Scalable and Resilient Algorithms***, FP3C Workshop, Maison de Simulation, Paris, France, March 2014.
- **Keynote: *Toward the Next Generation of Parallel and Resilient Algorithms & Applications***, SPPEXA Workshop 2014, Cologne, Germany, December 2013.
- **Keynote: *Toward the Next Generation of Parallel and Resilient Algorithms***, SCALA Workshop 2013, SC'13, Denver, CO, November 2013.
- **Keynote: *Toward Resilient Algorithms and Applications***, FTXS 2013, New York, NY, June 2013.
- **Keynote: *Toward Effective Parallel Programming: What we Need and Don't Need***, HIPS Workshop, IPDPS 2013, Boston, MA, May 2013.
- **Invited: *The Virtues of Data Transparency***, SOS 17, Jekyll Island, SC, March 2013.

- **Invited: *What Every SIAM Member Should Know about Computing on Emerging Architectures***, 2012 SIAM Annual Meeting, Minneapolis, MN, July 2012.
- **Invited: *Scalability of Trilinos: People, Processes, Parallelism***, 2012 ESCO Conference, Pilsen, Czech Republic, June 2011.
- **Invited: *Numerical Libraries on Emerging Architectures***, 2011 Supercomputing Conference Tutorial, Seattle, WA, November 2011.
- **Invited: *Emerging Architectures and UQ: Implications and Opportunities***, IFIP Workshop on uncertainty quantification, Boulder, CO, August 2011.
- **Invited: *Building the Next Generation of Parallel Applications and Libraries***, INT Workshop on Exascale Computing, Seattle, WA, June 2011.
- **Invited: *Toward Portable Programming of numerical linear algebra on manycore nodes***, CEA-EDF-INRIA 2011 Summer School, Nice, France, June 2011.
- **Keynote: *Scalability of Trilinos: People, Processes, Parallelism***, 3<sup>rd</sup> International Conference on Computational Methods in Engineering and Science (FEMTEC 2011), South Lake Tahoe, NV, May 2011.
- **Invited: *Building the Next Generation of Parallel Applications***, Salishan Conference on High Speed Computing, April 2011.
- **Invited: *Miniapplications: Vehicles for Co-Design***, Engelberg, Switzerland, March 2011.
- **Invited: *Requirements on Next-Generation Programming Models***, U of Houston, January 2011.
- **Invited: *Trilinos for Extreme-scale for Computing***, U Texas, Austin, January 2011.
- **Invited: *Software Engineering for Computational Science and Engineering***, Cray, Inc., January 2011.
- **Invited: *Building the Next Generation of Parallel Applications and Libraries***, IAM, April 2011.
- **Invited: *Bi-modal MPI-only & MPI+threading***, Cray, Inc., December 2010.
- **Invited: *The Extreme-scale Algorithms & Software Institute***, Fall Creek Falls Conference, October 2010, Memphis, TN.
- **Invited: *Building the Next Generation of Scalable Applications***, Future of the Field Workshop, Snowbird, UT, July 2010.
- **Keynote: *Building the Next Generation of Parallel Applications***, Int'l Workshop on OpenMP, Tsukuba, Japan, June 2010.
- **Keynote: *Trilinos for Extreme-scale Computing***, SPEEDUP Workshop, ETH-Zurich, September 2010.
- **Invited: *Trilinos Overview and Tutorial***, Purdue University, September 2009.

- **Invited: *Software Needs for Next-generation systems***, SOS13, Hilton Head, SC, March 2009.
- **Invited: *Algorithms for 1M cores: What Might and Might not Work***, Simulating the Future Workshop, Paris, France, September 2008.
- **Organizer: *When MPI-only is not Enough: Building the Next Generation of Scalable Applications Workshop***, Santa Fe, NM, May 2008.
- **Invited: *Design Issues for Numerical Libraries on Multicore Systems***, SciDAC Conference, July 2008, Seattle, WA.
- **Invited: *An Overview of Trilinos***, Oak Ridge National Laboratory, TN, October 2007.
- **Keynote: *Optimal Kernels to Optimal Solutions: Algorithm and Software Issues in Solver Development***, PDP07, February 2007, Naples, Italy.