SAND2012-3441P

In-Situ Visualization with Catalyst

Abstract

In-situ visualization is a term for running a solver in tandem with visualization. Catalyst is the new name for ParaView's coprocessing library. ParaView is a powerful open-source turnkey application for analyzing and visualizing large data sets in parallel. By coupling these together, we can utilize HPC platforms for analysis while circumventing bottlenecks associated with storing and retrieving data in disk storage. We demonstrate two methods for *in-situ* visualization using Catalyst. The first is linking Catalyst directly with simulation codes. It simplifies integration with the codes by providing a programmatic interface to algorithms in ParaView. Attendees will learn how to build pipelines for Catalyst, how the API is structured, how to bind it to C, C++, Fortran, and Python and how to build Catalyst for HPC architectures. The second method uses a variety of techniques, known as data staging or in-transit visualization, that involve passing the data through the network to a second running job. Data analysis applications, written using Catalyst, can operate on this networked data from within this second job minimizing interference with the simulation but also avoiding disk I/O. Attendees will learn three methods of handling this procedure as well as the APIs for ADIOS and NESSIE.

Detailed Description

General description and tutorial goals

We are proposing a half-day tutorial on adding analysis and visualization to existing codes. The tutorial presents several methods to do this within the Catalyst framework. The first is through the interface presented by the Catalyst API and the visualization and analysis capabilities it provides. Attendees learn how to integrate the coprocessing library into their existing codes and to configure and apply the analysis. The remaining methods will use data staging and in-transit visualization, using ParaView's Live Covisualization, ADIOS and NESSIE. Live Covisualization makes use of ParaView's client server architecture to transfer server data from a Catalyst pipeline to a running ParaView client directly. This allows someone to view and modify the coprocessing pipeline during a simulation run. ADIOS is an I/O framework where the API is detached from the available methods of writing. Staging methods transfer data to another application's memory instead of to a file. ParaView has a generic plugin to read data from the ADIOS-BP file format or from staging areas written by ADIOS. NESSIE provides a network layer to transfer function calls from the simulation to a set of services running in a separate job. This is made transparent to the coprocessing pipeline which functions as though it is running directly linked with the simulation's original function calls. Attendees will learn how to use these methods to indirectly couple their visualizations with their simulation.

Targeted audience, content level, and relevance to SC attendees

We expect the content break down to be: 10% beginner, 40% intermediate, 50% advanced.

This tutorial has an appeal for many levels and types of attendees, but is primarily focused on those working with large-scale simulation codes. Attendees currently programming simulations or other large-scale codes will directly benefit from the expansive capabilities of ParaView in general and specifically within Catalyst. Attendees familiar with visualization, graphics, or VTK will learn how to programmatically interface to these components in large-scale applications.

Beginners benefit by getting an overview of the capabilities of the ParaView framework and an introduction to applying it. For attendees that want a more comprehensive introduction to using ParaView, we are also proposing a companion tutorial titled "Large Scale Visualization with ParaView". The understanding they will gain from attending only this tutorial will help them choose and apply their visualization tools. This tutorial also provides the necessary information for using the coprocessing tools already integrated into other codes.

Intermediate and advanced users are given the instruction necessary to leverage Catalyst within their codes. We give an overview of the API and detailed instruction on how to use it. We also apply instruction on how to configure the analysis within ParaView and build it into Catalyst, and how to create a workflow that provides analysis both running *in situ* with the simulation and interactively at a user's convenience. We demonstrate how simulation codes can directly send their results to a visualization server using data staging and in-transit and then use Catalyst or ParaView for the analysis and visualization as well.

The majority of the Supercomputing attendees are from universities or government labs. In this environment it is important to be able share tools and applications. Since this suite of tools is open source, there are no barriers to collaboration between diverse organizations. The fact that there are

no license fees for these applications is also important. Some university research groups may not be able to easily purchase expensive licenses for proprietary visualization applications.

ParaView is a world leader in high performance visualization on distributed clusters. Expensive shared memory computers have succumbed to the economically superior distributed cluster computer. Many Supercomputing attendees are researching the application of distributed clusters toward high performance computation and visualization. ParaView has been designed from inception to run well on distributed computing platforms.

ADIOS has grown out from the I/O requirements of the largest supercomputing applications using the Leadership Computing Facilities. It is a framework of write and read plugins so that an application can choose the best performing I/O strategy in a particular environment. ADIOS has been used by applications comprising a quarter of all INCITE (Innovative and Novel Computational Impact on Theory and Experiment) resource allocations of the Department of Energy in 2010.

The NEtwork Scalable Service Interface (Nessie) is a framework for developing in transit analysis capabilities. It provides a remote-procedure call (RPC) abstraction that allows the application-developer to create custom data services to match the specific needs of the application.

Audience prerequisites

Much of the discussion in this tutorial involves using the programming interface to the coprocessing library. Attendees will get the most out of the tutorial if they have enough programming experience to follow along with the discussions. The tutorial provides examples in C++, Fortran, and Python. For our target audience, familiarity with at least one of these languages is prevailing.

Ensuring cohesive content

The presenting institutions, Sandia National Laboratories, Kitware, Inc., and Oak Ridge National Laboratory, have worked closely together in the development of much of the software presented in this tutorial, and will continue to work closely together in building the tutorial content

Tutorial updates for SC

We have given many introductory tutorials to ParaView in different forums, including several past Supercomputing conferences. In addition, at SC11 we gave a similar tutorial on *in-situ* visualization. The intent of this tutorial is similar to last year's tutorial in that it demonstrates a way to provide a service to those developing other large-scale parallel applications, and Supercomputing is an excellent opportunity for that. The difference between the proposed tutorial and last year's tutorial is a more distinct split between the capabilities of direct in-situ and the evolving techniques for handling indirect coupling through data staging and in-transit visualization.

We are also submitting an introductory companion tutorial on using the ParaView framework for general large scale visualization. The reason that these proposals are being submitted separately is that current ParaView users may not be interested in attending the introductory tutorial and non-programmer users may not be interested in this content of this tutorial.

Acknowledgements

Sandia National Laboratories is a multi-program laboratory operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Company, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Description of Demonstration

The information in this tutorial is communicated via slide presentations and live demonstrations, both from a laptop provided by the presenter.

Required Hardware

Standard A/V equipment is sufficient for this tutorial.

Outline

1. Introduction (30 minutes).
    a. What is coprocessing
    b. Uses and examples
    c. Coprocessing ParaView plugin
2. Coprocessing API (1 hr).
    a. Establishing a pipeline.
        i. Python scripting.
        ii. Hard coding.
    b. Specifying/querying input data.
    c. VTK data structures.
    d. Linking to Fortran or C.
    e. Compiling.
        i. Cross compiling.
        ii. Compiling without shared libraries.
        iii. Compiling with and without Python.
3. Data Staging (1.5 hours).
    a. Through Catalyst/ParaView
    b. Through ADIOS
        i. Write API.
        ii. Linking to FORTRAN or C.
        iii. Choosing the write method.
        iv. Read API (to understand how ParaView gets the data).
        v. Establishing the pipeline in a job submission.
    c. Through NESSIE
        i. The data-service model
        ii. Developing a simple data-transfer service
        iii. Integrating data services into an application workflow (how to launch and use data services)

Tutorial Notes Release

The presenters of this tutorial agree to release the tutorial notes on the SC12 USB stick.

Travel Support Request

Some of the presenters will request support for travel.

Résumé: Nathan Fabian

Education

- Master of Science Computer Science from the University of New Mexico, 2008.
- Bachelor of Science Computer Science from the University of New Mexico, 2001.

Work Experience

Member of Technical Staff — Sandia National Laboratories

Scalable Visualization and Analysis

Selected Publications

- Barrett, Brian W, Richard F Barrett, James M Brandt, Ron B Brightwell, Matthew L Curry, Nathan D Fabian, Kurt B Ferreira, Ann C Gentile, K Scott Hemmert, Suzanne M Kelly, Ruth A Klundt, James H Laros, III, Vitus J Leung, Michael J Levenhagen, Gerald F Lofstead, Kenneth D Moreland, Ron A Oldfield, Kevin T Pedretti, Arun F Rodrigues, David Thompson, Tom Tucker, Lee H Ward, John P Van Dyke, Courtenay T Vaughan, Kyle B Wheeler, "Report of Experiments and Evidence for ASC L2 Milestone 4467 - Demonstration of a Legacy Application's Path to Exascale," Sandia National Laboratories. 2012**.**
- Kenneth Moreland, Ron Oldfield, Pat Marion, Sebastien Jourdain, Norbert Podhorszki, Venkatram Vishwanath, Nathan Fabian, Ciprian Docan, Manish Parashar, Mark Hereld, Michael E. Papka and Scott Klasky. "Examples of *In Transit* Visualization." Petascale Data Analytics: Challenges and Opportunities (PDAC-11). 2011.
- Nathan Fabian and Kenneth Moreland and David Thompson and Andrew C. Bauer and Pat Marion and Berk Geveci and Michel Rasquin and Kenneth E. Jansen. "The ParaView Coprocessing Library: A Scalable, General Purpose In Situ Visualization Library." Proceedings of the IEEE Symposium on Large-Scale Data Analysis and Visualization. 2011
- Kenneth Moreland, Nathan Fabian, Pat Marion, and Berk Geveci. "Visualization on Supercomputing Platform Level II ASC Milestone Results." Sandia National Laboratories. 2010.
- Elaine Raybourn, Nathan Fabian, Eilish Tucker, Matthew Willis. "Beyond Game Effectiveness Part II: A Qualitative Study of Multi-Role Experiental Learning". In Interservice/Industry Training, Simulation and Education Conference, 2010.
- D. Thompson, N. D. Fabian, K. D. Moreland, and L. G. Ice, "Design issues for performing in situ analysis of simulation data," Tech. Rep. SAND2009-2014, Sandia National Laboratories, 2009.
- D. Thompson, R. W. Grout, N. D. Fabian, and J. C. Bennett, "Detecting combustion and flow features in situ using principal component analysis," Tech. Rep. SAND2009-2017, Sandia National Laboratories, 2009.

Selected Presentations

- Moreland, Kenneth, Nathan Fabian, Berk Geveci, Utkarsh Ayachit, James Ahrens, "Next-Generation Capabilities for Large-Scale Scientific Visualization," SIAM Conference on Parallel Processing for Scientific Computing, February 2012.
- Andrew Bauer, Nathan Fabian, Norbert Podhorszki, Pat Marion. *"In-Situ* Visualization with the ParaView Coprocessing Library." *Supercomputing 2011*, November 2011.
- Kenneth Moreland, Andrew Bauer, Pat Marion, and Nathan Fabian. *"In-Situ* Visualization with the ParaView Coprocessing Library." *Supercomputing 2010*, November 2010.

- Elaine Raybourn, Nathan Fabian, Eilish Tucker, Matthew Willis. "Real-time Individualized Training Vectors for Experiential Learning". Poster Presentation at LDRD Day, Sept 9, 2010.

- K. Moreland, N. Fabian, D. Thompson, P. Pebay, D. DeMarle, J. Woodring, "ParaView Tutorial," IEEE Visualization, 2009.

Résumé: Andrew Bauer

Education

- Doctor of Philosophy Mechanical Engineering from University at Buffalo, 2003.
- Bachelor of Science Mechanical Engineering from Binghamton University, 1996.

Work Experience

Research Staff — Kitware Inc.

Lead developer of the ParaView CoProcessing library.

Research Staff — Rensselaer Polytechnic Institute

Implemented adaptive mesh technologies for a variety of FEM simulations.

Research Experience (Academic)

Doctoral Research — University at Buffalo

Conducted research on efficient parallel solvers and preconditioners for *hp* adaptive simulations of 2D and 3D linear elasticity. Worked on Hilbert Space Filling Curves method for load-balancing partitioned grids.

Selected Publications

- Fabian N., Moreland K., Thompson D., Bauer A**.**, Geveci B., Rasquin M., Jansen K., Marion P., The ParaView Coprocessing Library: A scalable, general purpose in situ visualization library , LDAV 2011 IEEE Symposium on Large-Scale Data Analysis and Visualization, Oct-2011
- S.C. Jardin, N. Ferraro J. Breslau, A. Bauer, K. Jansen, M. Shephard, and the M3D team, "Two-fluid Extended-MHD Calculations of Collisionless Reconnection in Magnetized Plasmas and Toroidal Equilibrium", *34th European Physical Society Conference on Plasma Physics*, 1-4, 2007.
- L. Diachin, A. Bauer, B. Fix, J. Kraftcheck, K. Jansen, X. Luo, M. Miller, C. Ollivier-Gooch, M.S. Shephard, T. Tautges and H. Trease, "Interoperable mesh and geometry tools for advanced petascale simulations", *Journal of Physics: Conference Series, SciDAC 2007*, Institute of Physics Publishing, Vol. 78, No. 012015, 1-6, 2007
- A.K. Patra, C.C. Nichita, A.C. Bauer, E.B. Pitman, M. Bursik and M.F. Sheridan, "Parallel adaptive discontinuous Galerkin approximation for thin layer avalanche modeling", *Computers & Geosciences*, Volume 32, Issue 7, 912-926, 2006.
- M.S. Shephard, E. Seol, J. Wan and A.C. Bauer, "Component-based Adaptive Mesh Control Procedures", *Conf. on Analysis, Modeling & Computation of PDE & Multiphase Flow Proceedings*, Kluwer Academic Press, 1-10, 2005.
- A.C. Bauer and A.K. Patra, "Robust and efficient Domain Decomposition preconditioners for adaptive hp finite element approximations of linear elasticity with and without

discontinuous coefficients", *International Journal for Numerical Methods in Engineering*, Volume 59, Issue 3, 337-364, 2004.

- E.B. Pitman, C.C. Nichita, A. Patra, A. Bauer, M. Sheridan and M. Bursik, "Computing granular avalanches and landslides", *Physics of Fluids*, Volume 15, Issue 12, 3638-3646, 2003.
- A.C. Bauer, S. Sanjanwala and A.K. Patra, "Portable Efficient Solvers for Adaptive Finite Element Simulations of Elastostatics in Two and Three Dimensions", *Lecture Notes in Computational Science and Engineering: 23, Recent Developments in Domain Decomposition Methods*, L.F. Pavarino, A. Toselli, eds., Springer, 223-243, 2002.
- A.C. Bauer and A.K. Patra, "Performance of parallel preconditioners for adaptive hp FEM discretization of incompressible flows", *Communications in Numerical Methods in Engineering*, 18:305-313, 2002.

Selected Presentations

- Kenneth Moreland, Andrew Bauer, Pat Marion, and Nathan Fabian. *"In-Situ* Visualization with the ParaView Coprocessing Library." *Supercomputing 2010*, November 2010.

- Andrew Bauer. "Introduction to Visualization with VTK and ParaView." *Princeton Institute for Computational Science and Engineering, Princeton University*, May 2010.

- Andrew Bauer. "VTK Advanced Track Filtering." *Kitware Developer's Training Week*, October 2009.
- "Developments to Improve Extended MHD Simulations in M3D-C1", *Computational Plasma Physics Group Seminar*, Princeton Plasma Physics Laboratory, 2006.
- "High Performance Scientific Computing: Applications to Granular Flow Simulations", *Pegrum Lecture Series*, Department of Geology, University at Buffalo, 2002.

Professional Affiliations

- American Society of Mechanical Engineers (ASME)
- Association for Computing Machinery (ACM)

Résumé: Utkarsh Ayachit

Education

- Master of Science, Computer Science from University of Maryland, Baltimore County, 2004.
- Bachelor of Engineering, Computer Engineering from University of Pune, India, 2002.

Work Experience

Technical Lead — Kitware, Inc., 2009 – Present

R&D Engineer — Kitware, Inc. 2004 –2009

Selected Publications

- Ayachit U., Moreland K., Geveci B., Ma K-L., Dax Toolkit: A Proposed Framework for Data Analysis and Visualization at Extreme Scale , LDAV 2011 IEEE Symposium on Large-Scale Data Analysis and Visualization, Oct-2011
- Jomier J., Jourdain S., Ayachit U., Marion C., Remote Visualization of Large Datasets with MIDAS and ParaViewWeb, Web3D Conference 2011, Jun-2011
- Jourdain S., Ayachit U., Geveci B., ParaViewWeb, A web framework for 3D Visualization and Data Processing, IADIS International Conference on Web Virtual Reality and Three-Dimensional Worlds, Jul-2010
- Ayachit U., DeMarle D., Customizing ParaView, IEEE Revise Workshop, Oct-2009

Selected Presentations

- “DIY Vis Applications.” Berk Geveci, Utkarsh Ayachit, Jeffrey Baumes, Michael Bostock, Vadim Ogievetsky, Brian Wylie, Timothy M. Shead, Emanuele Santos, Timo Ropinski, Jörg-Stefan Prßini. Tutorial IEEE Visualization 2010, October 2010.
- "Large Scale Visualization with ParaView." Kenneth Moreland, John Greenfield, W. Alan Scott, Utkarsh Ayachit, and Berk Geveci. Tutorial Supercomputing 2009, November 2009.
- "Large Scale Visualization with ParaView." Kenneth Moreland, John Greenfield, W. Alan Scott, Utkarsh Ayachit, Berk Geveci, and David DeMarle. Tutorial Supercomputing 2008, November 2008.
- "Advanced ParaView Visualization." Kenneth Moreland, Utkarsh Ayachit, Timothy Shead, John Biddiscombe, and David Thompson. Tutorial IEEE Visualization 2008, October 2008.

Résumé: Norbert Podhorszki

Education

Doctor of Philosophy Information Science and Technology, Eötvös Loránd University of Budapest, Hungary, 2005.

Master of Science Computer Science, Eötvös Loránd University of Budapest, Hungary, 1995.

Work Experience

Research Scientist ORNL/NCCS

Scientific workflows for fusion code coupling. ADIOS framework read API and read API for staging applications. Fusion code coupling using the ADIOS framework.

Post-doctorate Researcher University of California, Davis

Kepler scientific workflow system extension for HPC. Workflow development for HPC monitoring and post-processing applications. Given Kepler tutorials at SC06, SC07, SC08, SciDAC-2008.

Research Scientist Computer and Automation Research Institute, Budapest, Hungary

R&D in performance monitoring and visualization for cluster and Grid applications. Lead developer of SZTAKI Desktop Grid. Presenter at Grid summer schools for the European Grid for EscienE project.

Selected Publications

- F. Zhang, C. Docan, M. Parashar, S. Klasky, N. Podhorszki, H. Abbasi: "Enabling In-situ Execution of Coupled Scientific Workflow on Multi-core Platform", Proceedings of the 26th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2012), Shanghai, China, May 2012.
- K. Moreland, R. Oldfield, P. Marion, S. Jourdain, N. Podhorszki, C. Docan, M. Parashar, M. Hereld, M. E. Papka and S. Klasky. "Examples of In Transit Visualization", Proceedings of the Workshop on Petascale Data Analytics: Challenges and Opportunities (PDAC-11), in conjunction with ACM/IEEE SC11, Seattle, WA, USA, November 2011.
- F. Zheng, H. Abbasi, C. Docan, J. Lofstead, Q. Liu, S. Klasky, M. Parashar, N. Podhorszki, K. Schwan, M. Wolf, "PreDatA - Preparatory Data Analytics on Peta-Scale Machines", IPDPS 2010, IEEE Computer Society Press 2010.
- C. Docan, J. Cummings , S. Klasky , M. Parashar, N. Podhorszki, F. Zhang, "Experiments with Memory-to-Memory Coupling for End-to-End fusion Simulation Workflows", ccGrid2010, IEEE Computer Society Press 2010.
- Cummings, Klasky, Podhorszki, Barreto, Lofstead, Schwan, Docan, Parashar, Sim, Shoshani, "EFFIS: and End-to-end Framework for Fusion Integrated Simulation", The 18th

Euromicro International Conference on Parallel, Distributed and Network-Based Computing (PDP) 2010.

- N. Podhorszki, S. Klasky, Q. Liu, H. Abbasi, J. Lofstead, K. Schwan, M. Wolf, F. Zheng, C. Docan, M. Parashar and J. Cummings, "Plasma fusion code coupling using scalable I/O services and scientific workflows", The 4th Workshop on Workflows in Support of Large-Scale Science (WORKS'09), Portland, OR, Nov. 2009
- C S Chang, S Ku, P Diamond, M Adams, R Barreto, Y Chen, J Cummings, E D'Azevedo, G Dif-Pradalier, S Ethier, L Greengard, T S Hahm, F Hinton, D Keyes, S Klasky, Z Lin, J Lofstead, G Park, S Parker, N Podhorszki, K Schwan, A Shoshani, D Silver, M Wolf, P Worley, H Weitzner, E Yoon and D Zorin: "Whole-volume integrated gyrokinetic simulation of plasma turbulence in realistic diverted-tokamak geometry." *IOP conference series SciDAC, 2009.*
- J H Chen, A Choudhary, B de Supinski, M DeVries, E R Hawkes, S Klasky, W K Liao, K L Ma, J Mellor-Crummey, N Podhorszki, R Sankaran, S Shende and C S Yoo , "Terascale direct numerical simulations of turbulent combustion using S3D", 2009 *Comput. Sci. Disc.* **2** 015001
- J. Lofstead, S. Klasky, K. Schwan, N. Podhorszki, C. Jin, "Flexible IO and Integration for Scientific Codes Through the adaptable IO System", Proc. of the 6th intl. workshop on Challenges of Large Applications in Distributed Environments (CLADE). 2008, Boston, MA, June 23 - 23, 2008.

Selected Presentations

- "Introduction to Scientific Workflow Management and the Kepler System", Scientific Discovery through Advanced Computing Program (SciDAC) tutorials on Friday, July 18th, 2008, Microsoft Research, Richmond, WA, http://blogs.msdn.com/b/escience/archive/2008/07/11/scientific-discovery-through-advanced-computing-program-scidac-tutorials-on-friday-july-18th.aspx#Kepler

- "Introduction to Scientific Workflow Management", at Supercomputing'08, Nov. 16, Austin, TX. http://sc08.supercomputing.org/scyourway/conference/view/tut130.html

- "Introduction to Scientific Workflow Management and the Kepler System", at Supercomputing'07, Nov. 11, 2007, Reno, NV http://sc07.supercomputing.org/schedule/event_detail.php?evid=11036

- "Introduction to Scientific Workflow Management and the Kepler System", at Supercomputing'06, Nov. 12, 2006, Tampa, FL http://sc06.supercomputing.org/schedule/event_detail.php?evid=9229

- Grid'05 EGEE Summer School, July 11-16, 2005, Budapest, Hungary. http://egee.hu/grid05/?m=2

Résumé: Ron A. Oldfield

Education

- Dartmouth College, Ph.D. Computer Science, 2003
- University of New Mexico, B.S. Computer Science, 1993

Employment

- Sandia National Laboratories, Principal Member of Technical Staff, 1/2011–Present
- Sandia National Laboratories, Senior Member of Technical Staff, 9/2003–1/2011
- Cold Regions Research and Engineering Laboratory, Technical Consultant, 7/1998–07/2000
- Sandia National Laboratories, Limited Term Member of Technical Staff, 1/1994–5/1997
- Sandia National Laboratories, Associated Western Universities Student Fellow, 6/1992–12/1993

Current Research Activities

- PI of Scalable I/O services project: ASC project to exploit available system resources for in-situ analysis and I/O-optimization services for applications on HPC systems.
- Project lead for all ASC-funded Scalable I/O activities at Sandia.

Related Publications (reverse chronological order)

[1] Jay Lofstead, Ron Oldfield, Todd Kordenbrock, and Charles Reiss. Extending scalability of collective I/O through Nessie and staging. In *Proceedings of the 6th Parallel Data Storage Workshop*, Seattle, WA, November 2011.

[2] Kenneth Moreland, Ron Oldfield, Pat Marion, Sebastien Joudain, Norbert Podhorszki, Venkatram Vishwanath, Nathan Fabian, Ciprian Docan, Manish Parashar, Mark Hereld, Michael E. Papka, and Scott Klasky. Examples of in transit visualization. In *Proceedings of the PDAC 2011 : 2nd International Workshop on Petascale Data Analytics: Challenges and Opportunities*, Seattle, WA, November 2011.

[3] Jay Lofstead, Fang Zheng, Qing Liu, Scott Klasky, Ron Oldfield, Todd Kordenbrock, Karsten Schwan, and Matthew Wolf. Managing variability in the IO performance of petascale storage systems. In *Proceedings of SC2010: High Performance Networking and Computing*, November 2010.

[4] Ron A. Oldfield, Todd Kordenbrock, and Patrick Widener. A survey of data-movement approaches for HPC storage systems. In Ada Gavrilovska, editor, *Attaining High Performance Communication: A Vertical Approach*, chapter 17, pages 329–351. CRC Press, 2009.

[5] Ron A. Oldfield, Andrew Wilson, George Davidson, and Craig Ulmer. Access to external resources using service-node proxies. In *Proceedings of the Cray User Group Meeting*, Atlanta, GA, May 2009.

[6] Charles Reiss, Gerald Lofstead, and Ron Oldfield. Implementation and evaluation of a staging

proxy for checkpoint I/O. Technical report, Sandia National Laboratories, Albuquerque, NM, August 2008.

[7] Ron A. Oldfield, Sarala Arunagiri, Patricia J. Teller, Seetharami Seelam, Rolf Riesen, Maria Ruiz Varela, and Philip C. Roth. Modeling the impact of checkpoints on next-generation systems. In *Proceedings of the 24th IEEE Conference on Mass Storage Systems and Technologies*, San Diego, CA, September 2007.

[8] Ron A. Oldfield. Investigating lightweight storage and overlay networks for fault tolerence. In *Proceedings of the High Availability and Performance Computing Workshop*, Santa Fe, NM, October 2006.

[9] Ron A. Oldfield, Patrick Widener, Arthur B. Maccabe, Lee Ward, and Todd Kordenbrock. Efficient data-movement for lightweight I/O. In *Proceedings of the 2006 International Workshop on High Performance I/O Techniques and Deployment of Very Large Scale I/O Systems*, Barcelona, Spain, September 2006.

[10] Ron A. Oldfield, Arthur B. Maccabe, Sarala Arunagiri, Todd Kordenbrock, Rolf Riesen, Lee Ward, and Patrick Widener. Lightweight I/O for scientific applications. In *Proceedings of the IEEE International Conference on Cluster Computing*, Barcelona, Spain, September 2006.

[11] Ron Oldfield and David Kotz. Scientific applications using parallel I/O. In Hai Jin, Toni Cortes, and Rajkumar Buyya, editors, *High Performance Mass Storage and Parallel I/O: Technologies and Applications*, chapter 45, pages 655–666. IEEE Computer Society Press and John Wiley & Sons, 2001.

[12] Ron A. Oldfield, David E. Womble, and Curtis C. Ober. Efficient parallel I/O in seismic imaging. *The International Journal of High Performance Computing Applications*, 12(3):333–344, Fall 1998.

Recent Collaborators

- Patrick Bridges, University of New Mexico
- Ada Gavrilovska, Georgia Institute of Technology
- Karsten Schwan, Georgia Institute of Technology
- Scott Klasky, Oak Ridge National Laboratory
- Hasan Abassi, Oak Ridge National Laboratory
- Gerald Lofstead, Georgia Institute of Technology
- Arthur B. Maccabe, Oak Ridge National Laboratory
- Philip C. Roth, Oak Ridge National Laboratory
- Patricia Teller, University of Texas at El Paso
- Patrick Widener, Emory University