

EWA DEELMAN, Ph.D.

RESEARCH INTERESTS

Managing large-scale workflows and data in distributed environments. Design and exploration of collaborative scientific environments based on Grid and Cloud technologies. Performance optimization in parallel and distributed environments.

APPOINTMENTS

2009-present

Research Associate Professor, USC Computer Science Department

2003-2009

Research Assistant Professor, USC Computer Science Department

University of Southern California, Los Angeles, CA

2016-present

Research Director, Science Automation Technologies, USC Information Sciences Institute

2012-2015

Assistant Director, Science Automation Technologies, USC Information Sciences Institute

2007-2012

Project Leader, USC Information Sciences Institute

2002-2007

Research Team Leader, USC Information Sciences Institute

2000-2002

Computer Scientist, USC Information Sciences Institute

Marina del Rey, CA

Scientific applications need to make use of a number of heterogeneous and often distributed resources that include both data and computations. My research focuses on developing tools and techniques that enable easy, efficient, and reliable execution of complex applications in distributed environments. Particular emphasis is given to managing applications specified as computational workflows. My research encompasses a number of areas including task scheduling, performance optimization, reliability techniques, resource provisioning, and workflow monitoring and troubleshooting. One of the tools developed in my group is the Pegasus Workflow Management System that is being used day-to-day in a number of scientific applications including astronomy, biology, earthquake science, gravitational-wave physics and others. The applications make use of their own computing clusters, national cyberinfrastructure, and Cloud computing resources. My research interests also include data and metadata management in scientific applications.

1997-2000

Senior Software Developer, University of California

Los Angeles, CA

Conducted research in parallel simulation of message-passing programs and performance modeling.

Participated in the POEMS project, which aims to design a system for Performance Oriented End-to-End Modeling of Large Heterogeneous Adaptive Parallel/Distributed Computer/Communication Systems. Major development was geared towards efficient and accurate simulation of message-passing programs (such as those using MPI) on high performance systems (such as the IBM SP). New SMP-cluster architectures were modeled. Compiler optimizations which improved the efficiency and scalability of the simulator were investigated. The POEMS project involves a collaboration between several universities: Purdue, Rice, UCLA, U.T. Austin, U.T. El Paso, U. of Wisconsin and LANL.

Guest lectured in parallel computation courses.
Guided the research of graduate students.

EDUCATION

1997

Doctor of Philosophy in Computer Science

Rensselaer Polytechnic Institute Troy, NY

"Performance Optimization of Parallel Discrete Event Simulation of Spatially Explicit Problems"

Developed an object oriented Parallel Discrete Event Simulation System that uses the optimistic approach to event processing. The system was designed to run on an MIMD machine and tested on an IBM SP. In order to achieve good performance, new synchronization techniques were formulated. Developed a new algorithm to continuously monitor the progress of all simulation processes. Designed a new approach for rolling back computation in spatially explicit problems. This approach resulted in a speedup close to linear over a single processor run. Load balancing techniques were also investigated.

Thesis Advisor: Prof. Boleslaw Szymanski

1991

Master of Science in Computer Science

State University of New York, New Paltz, NY

1987

Bachelor of Arts in Mathematics

Wells College, Aurora, NY

PROFESSIONAL ACTIVITIES

- Workshops Chair, SC'16 Technical Program, Salt Lake City, November 2016
- Posters Chair, Cluster 2016, Taipei Taiwan, September 2016
- Editorial Board Member, International Journal of High Performance Computing Application, 2015 -
- Editorial Board Member, ACM International Conference Proceeding Series, 2015-
- Co-Technical Papers Chairs for the SC'15 Technical Program, Austin, TX, November 2015
- 11th International Conference on Parallel Processing and Applied Mathematics (PPAM), Vice-Chair of Program Committee, Krakow, Poland, September 2015
- 8th IEEE/ACM International Conference on Utility and Cloud Computing, Tutorials co-chair, Cyprus, December 2015
- Co-area chair for the SC'14 Technical Program (Clouds and Distributed Computing area), New Orleans, 2014
- Best Paper Committee Member, SC'14 New Orleans
- Member of the Editorial board, Journal of Grid Computing, 2013-
- Steering Committee Member, IEEE eScience Conference, 2013-
- Steering Committee Member, IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, 2013-
- Steering Committee Member, e-Science All Hands Meeting Foundation, 2008-
- Associate Editor responsible for Grid and Cloud Computing for the Scientific Programming Journal, 2002-2014
- Vice Program Chair, International Conference on Parallel Processing and Applied Mathematics. PPAM, Poland, September 2013

- Chair, Global Grid Forum Workflow Management Research Group, 2003-2010
 - Co-Chair NSF Workshop SDCI/STCI as the Software Supply Chain of the National Cyberinfrastructure Workshop, January 2010
 - Mentor for the CRA-W Distributed Research Experiences for Undergraduates (DREU) program, 2008, 2010
 - Co-Chair, Workshop on Workflows in Support of Large-Scale Science, in conjunction with SC 2008, SC 2009, and 2010
 - Founder and Chair, Workshop on Workflows in Support of Large-Scale Science, in conjunction with HPDC 2006, and 2007 June 2006 & 2007
 - Co-Guest Editor, Scientific Programming Journal, Special issue on Dynamic Computational Workflows: Discovery, Optimization, and Scheduling, 2007
 - Co-Guest Editor Journal of Grid Computing, Special Issue dedicated to Workflow Management in Grids, 2006
 - Co-Editor, Workflows in e-Science book, Springer 2007
 - Workshop Chair Applications and Middleware Grid Workshop, September 2005
 - Lecturer, Global Grid Forum Summer School, July 2004
 - Organizing Committee Member, GGF Workshop on Workflows in Grids, 2004.
 - Co-Editor of the Special Issue of the Scientific Programming Journal devoted to Grid Computing, Volume 10, Number 2, 2002.
 - Workshop Chair, Applications Grid Workshop, September 2003, 2005, 2007
 - Panel member at the International Conference on Dependable Systems and Network panel on "Dependability and the Grid", June 2002.
 - Technical Board member of the Gridlab project, an EU funded project.
 - PhD Committee member, Rensselaer Polytechnic Institute, 2002.
 - External Examiner for Masters' Candidate, McGill University, 2002.
 - General Co-Chair, of the 15th Workshop on Parallel and Distributed Simulation, PADS2001, Lake Arrowhead, CA, 2001.
 - Organizer So-Cal Seminar on Parallel Computing Systems, January 22, 1999.
 - Organizing Committee Chair, Third Workshop on Languages, Compilers, and Run-Time Systems for Scalable Computers held at Rensselaer Polytechnic Institute, May 1995.
 - Pipelink Project Member, Lectured high-school students about research performed in Computer Science. The goal of this project is to interest and support young women in Computer Science, 1995.
-
- The International Supercomputing Conference (ISC), 2016, Germany
 - 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2016), Columbia
 - 30th IEEE International Parallel & Distributed Processing Symposium (IPDPS), 2016, Chicago
 - 11th IEEE International Conference on eScience, 2015, Germany
 - 10th IEEE International Conference on eScience, 2014, Brazil
 - 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, 2014
 - 2nd Workshop on Sustainable Software for Science: Practice and Experiences in conjunction with SC'14, 2014
 - 9th IEEE International Conference on eScience, 2013
 - 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing
 - The 8th Workshop on Workflows in Support of Large-Scale Science, Denver, November 2013
 - First Workshop on Sustainable Software for Science: Practice and Experiences (WSSSPE), Denver, CO 2013
 - IEEE Big Data 2013, Santa Clara, CA, October 2013
 - The 8th International Conference on P2P, Parallel, Grid, Cloud and Internet Computing, Compiègne, France, October 2013
 - International Conference on Cloud and Green Computing, Karlsruhe, Germany, September,

Program
Committee
Member
(selected)

2013

- AMGCC'13 (The 1st International Workshop on Autonomic Management of Grid and Cloud Computing), Miami, Florida, August, 2013
- Grid and Cloud Computing in Biomedicine and Life Sciences, Porto, Portugal, June 2013
- BIGPROV'13: International Workshop on Managing and Querying Provenance Data at Scale, Genoa, March 2013, Italy
- Workshop on Workflows in Support of Large-Scale Science, in conjunction with SC 2012, Salt Lake City, November 2012
- The 5th IEEE/ACM International Conference on Utility and Cloud Computing, Chicago, IL November 2012
- The 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid) 2012, Ottawa, Canada, May 2012
- Scientific and Statistical Database Management Conference 2011, Portland Oregon
- HPDC 2011, San Jose, CA, 2011
- Workshop on Workflows in Support of Large-Scale Science, in conjunction with SC 2011
- Data Intensive Computing in the Clouds (DataCloud2011), San Jose, CA, 2011
- IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2011), Newport Beach, CA, 2011
- IPDPS 2011 PhD forum, Anchorage, Alaska, 2011
- IPDPS 2011, Anchorage Alaska, 2011
- e-Science 2010, Brisbane, Australia, 2010
- SC 2010, New Orleans, 2010
- CCGrid 2010, Australia, May 2010
- Grid 2010, Brussels, Belgium, October 2010
- 4th Int'l Workshop on Workflow systems in e-Science (WSES09), Shanghai, China, 2009
- High Performance Distributed Computing (HPDC 2009), Garching, Germany, June 2009
- SC 2008, Austin Texas, November 2008
- Grid 2008, Tsukuba, Japan, September 2008
- High Performance Distributed Computing (HPDC 2008), Boston, MA, June 2008
- e-Science 2007, Bangalore, India, December 2007
- NSF/Mellon Workshop on Scientific and Scholarly Workflow. Oct 3-5, 2007, Baltimore, MD
- CoreGRID Workshop on Grid Middleware, Dresden, June 18-19, 2007
- High Performance Distributed Computing (HPDC 2007), Monterey Bay California, June 27-29
- IEEE International Conference on Grid Computing (Grid 2007) (Vice Program chair for Scheduling, Resource Management and Runtime Environments), Grid 2006, Grid 2005
- Second International Conference on Grid computing, high-performance and Distributed Applications" (GADA'07)
- Workflow Systems in e-Science 2007, International Conference on Computational Science (ICCS2007)
- The 13th International Conference on Parallel and Distributed Systems, Hsintu, Taiwan, December 2007
- 3rd International Workshop on Grid and Peer-to-Peer based Workflows (GPWW), Brisbane, Australia, 2007
- IEEE International Symposium on Cluster Computing and the Grid, CCGrid 2007, 2006, 2005
- CoreGRID Workshop on Grid Middleware 2006
- First International Multiconference on Computer Science and Information Systems
- The 2006 International Conference on High Performance Computing and Communications, Grid and cluster computing topic. 2006.
- International Workshop on Workflow Systems in Grid Environments (WSGE06), October 21-23, 2006, Changsha, China.
- Challenges of Large Applications in Distributed Environments Workshop, CLADE 2009, 2008, 2007, 2006, 2005, 2004, 2003

- Fourth International Symposium on Parallel and Processing and Applications (ISPA'06), (Software and Applications Track), Sorrento, Italy, December 2006.
- 2nd International Workshop on Grid and Peer-to-peer based Workflows (GPWW 2006) ,Vienna, Austria on September, 2006, and 2005
- Scientific workflow management in e-Science, in conjunction with the International Conference on Computational Science (ICCS2006), May 2006.
- Statistical and Scientific Database Management (SSDBM) , 2006, 2005
- ICIW 2006 Web Service-based Systems and Applications track, February 2006
- International Workshop on Scientific Instruments and Sensors on the Grid, December 2005.
- Semantic Infrastructure for Grid Applications at CCGrid 2005
- European Grid Conference 2005
- Workshop on Web and Grid Services for Scientific Data Analysis (WAGSSDA), to be held in conjunction with the International Conference on Parallel Processing (ICPP-2005)
- IEEE International Conference on Services Computing (SCC 2004)
- International Conference on Semantics for a Networked World 2004
- 2nd European Across Grids Conference, 2004
- SC 2003, November 2003
- Heterogeneous Computing Workshop, 2003, 2004

Reviewer

Reviewer for journals such as the Journal of Grid Computing, Future Generations of Computer Systems, Journal of Parallel and Distributed Computing, IEEE Intelligent Systems, Informatica, IEEE Transactions on Parallel and Distributed Systems, Concurrency and Computation: Practice and Experience, ACM Computing Surveys

Reviewer for DOE, NSF, University of California Office of the President, and European Commission, and other EU funding agencies.

INVITED TALKS

- CRA-W/CDC Distinguished Lecture event at the Southeast Women in Computing Conference (SEWIC), November, 2015
- Panelist at the White House National Strategic Computing Initiative Workshop, October 2015
- 11th IEEE International Conference on eScience, August 2015, Germany
- Supercomputing Frontiers 2015, March 2015
- Oak Ridge National Laboratory, Seminar, September 2014
- University of Queensland /Monash University (MURPA/QURPA) Seminar, Summer 2014
- Sonoma State University, CA, Computer Science Colloquium, Spring 2014
- International Conference on Parallel Processing and Applied Mathematics. PPAM, Torun, Poland, September 2013 (keynote)
- Focused Technical Workshop on Network Issue for Life Sciences Research – LLNL, July 2013 (invited talk)
- Laurence Berkeley National Laboratory, “Workflow Technologies for Science Automation”, Berkeley, CA, May 2013 (invited talk)
- AGH University of Science and Technology, “Science Platforms: Managing Workloads and Resources”, Krakow, Poland, March 2013 (invited talk)
- Panel Member, CENIC 100G and Beyond Workshop, San Diego, CA, March 2013
- Polish-Japanese Institute of Information Technology, “Science Platforms: Research and Technologies to Computational Science in Virtual Environments”, Warsaw, Poland, March 2013, (seminar)
- Institute of Fundamental Technological Research, Polish Academy of Sciences, “Science Platforms: Managing Computational Workloads and Resources”, Warsaw, Poland, March 2013, (Distinguished Lecture)
- HUBbub'12, “Managing Workflows Within HUBzero: How to Use Pegasus to Execute

- Computational Pipelines”, Indianapolis, Indiana, September 2012 (keynote)
- Clusters, Clouds, and Data for Scientific Computing, 2012, “Hosted Science: How to support complex scientific applications on the cloud”, Lyon Franc, September 2012 (invited talk)
 - EuroPar’12, Rhodes, Greece, August 2012, (keynote)
 - Invited Lecture Series, AGH University of Science and Technology, Krakow, Poland, March 2012
 - The 6th Workshop on Workflows in Support of Large-Scale Science, Seattle, November 2011, (keynote)
 - USC Global Conference, Hong Kong, October 2011 (panelist, The Grand Challenges of Engineering)
 - International Supercomputing Conference (ISC) Cloud Conference, Mannheim, Germany, September 2011
 - International Conference on Parallel Processing and Applied Mathematics. PPAM, Torun, Poland, September 2011 (keynote)
 - Cracow Grid Workshop, Cracow, Poland, October 2010. (keynote)
 - International Conference on Parallel Processing and Applied Mathematics. PPAM, Wroclaw, Poland, September 2009
 - Cracow Grid Workshop, Cracow, Poland, October 2008. (keynote)
 - CoreGrid Workshop, Crete, Greece, April 2008.
 - AAI Spring Symposium on Semantic Scientific Knowledge Integration, Stanford University, March 2008.
 - The International Workshop on Performance Analysis and Optimization of High-End Computing Systems, in conjunction with SC’07, November 2007
 - International Conference on Parallel Processing and Applied Mathematics. PPAM, Gdansk, Poland, September 2007
 - Computer Science Department Colloquium, Southampton University, UK April 2007
 - Computer Science Department Colloquium, King’s College, London, UK, April 2007
 - Colloquium, University of Vienna, Austria, June 2007
 - Workshop on Scientific Workflows and Business workflow standards in e-Science, in Conjunction with e-Science 2006, Amsterdam, December 2006
 - Workflow Optimization in Distributed Environments, Edinburgh, October 2006
 - Manchester University, School of Computer Science Seminar, October 2006
 - Louisiana State University, Computer Science Department Seminar, September 2006
 - Cyberinfrastructure for Ocean Observations Workshop, September 2006
 - Partnerships in Innovation: Serving a Networked Nation conference, National Archives and Records Administration, November 2004
 - Astronomical Data Analysis Software & Systems (ADASS), October 2004
 - SC4DEVO: Service Composition for Data Exploration in the Virtual Observatory, July 2004
 - SSDBM 04 panel on Data Management on the Grid, June 2004
 - Ground System Architectures Workshop (GSAW 2004), April 2004
 - SDSC Computational Science Seminar Series (CSSS)., January 2004
 - Rensselaer Polytechnic Institute, *GriPhyN: Data-Intensive Science in Grid Environments*, Troy, NY, January 2002.
 - VIRGO Laboratory, GriPhyN (Grid Physics Network) and LIGO: Data-Intensive Science in Grid Environments, Cascina, Italy, September 2001.
 - International Conference on Parallel Processing and Applied Mathematics. GriPhyN (Grid Physics Network): Enabling Data Intensive Science in Grid Environments. Naleczow, Poland, September 2001.
 - Poznan Supercomputing and Networking Center, GriPhyN: Data-Intensive Science in Grid Environments, Poznan, Poland, September 2001.
 - Mardi Gras Conference, GriPhyN (Grid Physics Network): Building a Data Grid Infrastructure for Experimental Physics, Baton Rouge, LA, February 2001.
 - University of Texas at El Paso, MPI-Sim, an MPI Simulator in POEMS (Performance Oriented End-to-End Modeling System), El Paso, TX, August 2000.

- USC/ISI, POEMS: Performance Oriented End-to-End Modeling System, Marina Del Rey, CA, July 2000.
- Aerospace Corporation, Optimization in Parallel Discrete Event Simulation, Los Angeles, CA, July 1997.
- Institute of Fundamental Technological Research, Polish Academy of Sciences, Performance Optimization in Parallel Discrete Event Simulation, Warsaw, Poland, June 1997.

TEACHING ACTIVITIES

- Lecturer at the 2nd Karlsruhe Summer School on Service Research, September 2013
- Guest Lectures at the AGH University of Science and Technology, March 2013
- USC, Computer Science Department, CS 599 Introduction to Grid Computing (Graduate level), Fall 2007
- UCLA, Department of Computer Science, Guest Lecturer
 - CS 133, Parallel and Distributed Programming, an undergraduate level course in parallel programming. It covers architectural and language models of parallelism. The emphasis is learning to write parallel programs. Winter Term 1998, and 1999.
 - CS 233A, Parallel Programming, graduate course in parallel programming that covers in more depth the topics of CS 133. Emphasis is on examining state of the art research and individual research projects. Winter Term 1998, and 1999.
 - CS 239, Parallel Simulation, a follow-up course to CS 233A. Spring Term 1998.

FELLOWSHIPS HONORS

- High-Performance Parallel and Distributed Computing (HPDC) 2015 achievement award, which is presented to an individual who has made long-lasting, influential contributions to the foundations or practice of the field of high-performance parallel and distributed computing, June 2015
- CENIC Star Performer for the ADAMANT project, January 2014
- Corporation for Education Network Initiatives in California's (CENIC) 2013 Innovations in Networking Award for Experimental/Developmental Applications for the ADAMANT (Adaptive Data-Aware Multi-domain Application Network Topologies) Project, March 2013
- Best Paper award, e-Science 2006, Amsterdam, Netherlands: E. Deelman, S. Callaghan, E. Field, H. Francoeur, R. Graves, N. Gupta, V. Gupta, T. H. Jordan, C. Kesselman, P. Maechling, J. Mehringer, G. Mehta, D. Okaya, K. Vahi, and L. Zhao, Managing Large-Scale Workflow Execution from Resource Provisioning to Provenance tracking: The CyberShake Example, Proceedings of e-Science, Amsterdam, The Netherlands, 2006.
- Best Paper award, 15th Workshop on Parallel and Distributed Simulation, Lake Arrowhead, CA, USA; 15-18 May 2001: E. Deelman, R. Bagrodia, R. Sakellariou, V. Adve. Improving Lookahead in Parallel Discrete Event Simulations of Large-Scale Applications Using Compiler Analysis. Proceedings 15th Workshop on Parallel and Distributed Simulation p. 5-13, Lake Arrowhead, CA, USA; 15-18 May 2001.
- General Electric Foundation Fellowship, Rensselaer Polytechnic Institute, 1993-1994.
- Special Distinction in the field of Mathematics, Wells College, 1987.

MEDIA CITATIONS

- "How ISI's Pegasus helped scientists make the discovery of the century", USC Viterbi School of Engineering, <http://viterbi.usc.edu/news/news/2016/isi-gravitational-waves-software-pegasus.htm>
- "Caltech wasn't the only SoCal school helping discover gravitational waves", mentions the use of Pegasus for managing workflows used to discover gravitational waves, Southern California NPR Station: KPCC, February 11, 2016 <http://www.scpr.org/news/2016/02/11/57446/caltech-wasn-t-the-only-socal-school-helping-disco/>
- Ewa Deelman Awarded Prestigious HPDC Honor, September 2015, <http://www.isi.edu/news/story/251>
- CENIC Star Performer: Ewa Deelman, USC, January 2014, <http://www.cenic.org/p=1554/>

- Project ADAMANT Wins Innovations in Networking Award for XD Apps, March 13, 2013, <http://www.cenic.org/p=356/>
- How can we use HPC platforms to help dig out new exoplanets?, iSGTW, April 3, 2013, <http://www.isgtw.org/feature/how-can-we-use-hpc-platforms-help-dig-out-new-exoplanets>
- Pegasus project leader, White House advanced manufacturing adviser to speak at HUBzero conference, iTap Newsroom (Purdue IT News), July 2012 <https://hubzero.org/news/hubbub2012-speakers>
- \$27 million award bolsters research computing grid, R&D Magazine, June 2012 <http://www.rdmag.com/news/2012/06/27-million-award-bolsters-research-computing-grid>
- NHGRI broadens sequencing program focus on inherited diseases, medical applications, NIH News, July 2012 <https://www.genome.gov/27546261>
- Grants to USC Faculty Top \$100 Million, USC News, November 11, 2009, http://uscnews.usc.edu/university/grants_to_usc_faculty_top_100_million.html
- ISI Researchers Will Support New Brain Gene Expression Project, ISI News, November 2, 2009 http://www3.isi.edu/about-news_story.htm?s=220
- USC Neuroscientists to Map Gene Expression, USC News, October 5, 2009 http://uscnews.usc.edu/science_technology/usc_neuroscientists_to_map_gene_expression.html
- Viterbi School's ISI Part of FutureGrid Test Bed, USC Viterbi News, September 2009 <http://viterbi.usc.edu/news/news/2009/viterbi-school-s213206.htm>
- ISI Part of FutureGrid Test Bed, ISI News, September 11, 2009 http://www3.isi.edu/about-news_story.htm?s=216
- Image of the week - Earth-quaking science in Hollywood, international Science Grid this Week, iSGTW online, January 2008, <http://www.isgtw.org/?pid=1000848>
- Feature - Montage a rising star in grid-enabled sky mosaics, international Science Grid this Week, iSGTW online, December 2007, <http://www.isgtw.org/?pid=1000731>
- Feature - Pegasus invites new communities to saddle up, international Science Grid this Week, iSGTW online, September 2007, <http://www.isgtw.org/?pid=1000664>
- ISI Leads \$13.8 Million E-Science Effort to Tame Terabyte Torrents, ISI News, <http://www.isi.edu/news/news.php?story=165>
- Rensselaer Supercomputers Battle Lyme Disease, Rensselaer Polytechnic Institute Review Vol. 17 No. 18, June 14, 1996

GRANTS

- PI, Predictive Modeling and Diagnostic Monitoring of Extreme Science Workflows, DOE, 9/2014-8/2017, \$3,824,927
- Co-PI, CIF21 DIBBs: Domain-aware management of heterogeneous workflows: Active data management for gravitational-wave science workflows, NSF, 10/2014-9/2017, \$444,850
- PI, XSEDE Scientific Workflow Application Support Specialist, (XSEDE), 9/2013 – 06/2015, NSF, \$220,000
- Co-PI, The Center for Genomic Studies of Mental Disorders, 06/2013 – 05/2018, NIMH, \$1,879,412
- Co-PI, PAGE II Coordinating Center, 06/13 – 05/17, NHGRI, \$435,999
- PI, Collaborative Research: CC-NIE Integration: Transforming Computational Science with ADAMANT (Adaptive Data-Aware Multi-domain Application Network Topologies), January 2013-December 2014, NSF, \$204,070
Integrating resource provisioning and workflow execution on Clouds.
- Sr. Personnel, The Open Science Grid The Next Five Years: Distributed High Throughput Computing for the Nation's Scientists, Researchers, Educators, and Students, August 2012-July 2017, NSF, \$913,748
Supporting computations on the Open Science Grid.
- Co-PI dV/dt – Accelerating the rate of progress towards extreme scale collaborative Science, September 2012-August 2015, DOE, \$744,606
Conducting research in resource planning on dynamic resources for complex applications.

- PI, SI2-SSI: Distributed Workflow Management Research and Software in Support of Science, April 2012-March 2016, NSF, \$2,153,597
Continued development and support of the Pegasus Workflow Management System.
- Co-PI rSeq: Robust and Portable Workflow-based tools for mRNA and Genome re-sequencing, February 2012-December 2014, NIH, \$338,330
Developing and supporting RNA_Seq workflows.
- Co-PI, EarthCube Community Workshop: Designing A Roadmap for Workflows in Geosciences, April 2012-March 2013, NSF, \$15,432
Organizing and participating in discussion on the future of the Geosciences cyberinfrastructure.
- PI, The Role of Software and Software Institutes in Computational Science Over Time, October 2012-September 2013, NSF \$74,430
Organizing a workshop regarding the future of Cyberinfrastructure.
- Sr. Personnel, Geoinformatics: A Petascale Cyberfacility for Physics-Based Seismic Hazard Analysis (SCEC PetaSHA3 Project), NSF, July 2010-June 2012, \$174,000
Supporting the earthquake science workflows.
- Sr. Personnel, Sustaining and Extending the Open Science Grid: Science Innovation on a PetaScale Nationwide Facility, NSF December 2009-August 2011, \$128,000
Providing user support for the project
- PI, SDCl/STCl as the Software Supply Chain of the National Cyberinfrastructure Workshop, NSF, 2/1/2010- 1/31/2011, \$89,968
Organizing a workshop for the PI meeting for the NSF Office of Cyberinfrastructure SDCl and STCl programs
- Co-PI, Transcriptional Atlas of Human Brain Development, NIH, 9/1/2009-8/31/2011, \$492,900
Developing workflows for gene expression studies
- Sr. Personnel, FutureGrid: An Experimental, High-Performance Grid Test-bed, NSF, 10/1/2009-9/30/2013, \$725,000
Developing workflow technologies for virtual execution environments
- PI, STCl: Middleware for Monitoring and Troubleshooting of Large-Scale Applications on National Cyberinfrastructure, NSF, 9/1/2009-8/31/2012, \$1,875,830
Developing efficient, online workflow monitoring and troubleshooting capabilities.
- PI, STCl: Integrated Resource Provisioning Across the National Cyberinfrastructure in Support of Scientific Workloads, NSF, 9/1/09-8/31/2012, \$1,611,038
Developing tools that can provision computational resources in static and dynamics manner.
- Co-PI DC: Medium: Intelligent Data Placement in Support of Scientific Workflows, NSF, 09/01/2009-08/31/2012, \$405,000
Developing algorithms for coupling workflow and data management systems.
- Co-PI, Center for Genomic Studies of Mental Disorders, National Institutes of Health, 10/01/08-9/30/2013, \$2,250,000
Development of the Cyberinfrastructure and workflow technologies for the coordinating center.
- Sr. Personnel, Genomic Psychiatry Cohort, National Institutes of Health, 10/2008-9/2013, \$300,000
Development of the Cyberinfrastructure for a psychiatric study.
- Co-PI, NHGRI EpiGenVar Coordinating Center, National Institutes of Health, 7/1/2008-6/30/2012, \$500,000
Development of the Cyberinfrastructure and workflow technologies for the coordinating center.
- PI, Supporting Ocean Modeling With Workflow Technologies, Jet Propulsion Laboratory, 6/1/2008-5/31/2009, \$70,000
Development of workflows for and ocean forecasting system.
- PI, SDCl NMI Improvement: Pegasus: From Concept to Execution- - Mapping Scientific Workflows onto the National Cyberinfrastructure, National Science Foundation, 09/01/2007 - 08/31/2012, \$1,700,000

Hardening of the Pegasus workflow software that manages execution of complex scientific workflows on distributed resources.

- PI, Designing Scientific Software one Workflow at a Time, National Science Foundation, 10/1/2007-9/30/2010, \$313,556.
Exploring the use of the workflow paradigm in the context of large-scale software design.
- Sr. Personnel, Cyberinfrastructure in Support of Research: A New Imperative, National Science Foundation, 10/01/2004 - 09/30/2008, \$250,000
This project, lead by NCSA developed technologies necessary to support a wide range of NSF applications.
- Co-PI, Scalable Cross-Organization Threat and Event Discovery through Grid Workflows, Air Force Research Lab, 10/1/2006 - 9/30/2010, \$2,104,000
This project, lead by ISI develops new techniques to optimize workflows representing large-scale data mining applications and to provide seamless access to distributed data sources.
- Co-PI, CSR-AES: Collaborative Research: Intelligent Optimization of Parallel and Distributed Applications, National Science Foundation, 08/01/2006 – 07/31/2009, \$215,452.
This project, lead by ISI develops new techniques to optimize workflows in distributed Systems.
- Sr. Personnel, Enabling Earthquake System Science Through Petascale Calculations (PetaShake), National Science Foundation, 10/01/2007-9/30/2009, \$200,000.
Providing a cyberinfrastructure and workflow technologies for earthquake science applications.
- PI, Pegasus: Supporting LIGO Workflows on the Open Science Grid, National Science Foundation, 11/1/2006 – 10/31/2007, \$102,000
This project, lead by ISI will enable efficient execution of gravitational-wave physics applications on large grid deployments in conjunction with the efficient data management.
- Co-PI, Neuroscience: Collaborative Research (CRCNS): Assembling Visible Neurons for Simulations: Merging of High-throughput 3D Microscopies with Advanced Computational Tools, National Institute of Health (NIH), 10/01/2002 – 09/30/2006, \$45,000
This project, lead by SDSC, developed workflow software necessary to support complex tomographic image reconstruction on distributed resources.
- PI, Towards Cognitive Grids: Knowledge-Rich Services for Autonomous Workflow refinement and Robust Execution, National Science Foundation, 12/15/2004 – 11/30/2006, \$200,000
This project examined the use of semantic technologies to describe resources and services on the grid.
- Co-PI, WSW-06: 2006 Workshop on Challenges of Scientific Workflows, National Science Foundation, 05/01/2006 to 10/31/2007, \$44,000
Funding for the Workshop on the Challenges of Scientific Workflows

PUBLICATIONS

BOOK EDITOR

1. Workflows for e-Science: Scientific Workflows for Grids, Ian J. Taylor, Ewa Deelman, Dennis B. Gannon, Matthew Shields (Editors), Springer, January 2007

JOURNALS

1. Ewa Deelman, Karan Vahi, Mats Rynge, Gideon Juve, Rajiv Mayani and Rafael Ferreira da Silva, Pegasus in the Cloud: Science Automation through Workflow Technologies, IEEE Internet Computing, 20:1, pp. 70-76, 2016.
2. Idafen Santana-Perez, Rafael Ferreira da Silva, Mats Rynge, Ewa Deelman, Maria S. Perez-Hernandez and Oscar Corcho, Reproducibility of Execution Environments in Computational Science Using Semantics and Clouds, Future Generation Computer Systems, January 2016.
3. Piotr Bryk, Maciej Malawski, Gideon Juve and Ewa Deelman, Storage-aware Algorithms for Scheduling of Workflow Ensembles in Clouds, Journal of Grid Computing, p 1-20, 2015.
4. Ewa Deelman, Christopher Carothers, Anirban Mandal, Brian Tierney, Jeffrey S. Vetter, Ilya Baldin, Claris Castillo, Gideon Juve, Dariusz Krol, Vickie Lynch, Ben Mayer, Jeremy Meredith, Thomas Proffen, Paul Ruth and Rafael Ferreira da Silva, PANORAMA: An Approach to Performance Modeling and Diagnosis of Extreme Scale Workflows, International Journal of High Performance Computing Applications, July 2015.

5. Rafael Ferreira da Silva, Gideon Juve, Mats Rynge, Ewa Deelman and Miron Livny, Online Task Resource Consumption Prediction for Scientific Workflows, *Parallel Processing Letters*, 25:3, 2015.
6. Weiwei Chen, Rafael Ferreira da Silva, Ewa Deelman and Thomas Fahringer, Dynamic and Fault-Tolerant Clustering for Scientific Workflows, *IEEE Transactions on Cloud Computing*, vol.PP, no.99, pp.1-1, 2015.
7. D, Shungin, ... Ewa Deelman (author #69 out of ~300), New genetic loci link adipose and insulin biology to body fat distribution, **Nature**. 2015 Feb 12;518(7538):197-206. doi: 10.1038/nature14177.
8. Maciej Malawski, Kamil Figiela, Marian Bubak, Ewa Deelman and Jarek Nabrzyski, Scheduling Multilevel Deadline-Constrained Scientific Workflows on Clouds Based on Cost Optimization, *Scientific Programming*, 10.1155/2015/680271, 2015.
9. Maciej Malawski, Gideon Juve, Ewa Deelman and Jarek Nabrzyski, Algorithms for Cost- and Deadline-Constrained Provisioning for Scientific Workflow Ensembles in IaaS Clouds, *Future Generation Computer Systems*, 02/2015; 48:1–18. DOI: 10.1016/j.future.2015.01.004.
10. Ewa Deelman, Karan Vahi, Gideon Juve, Mats Rynge, Scott Callaghan, Philip J Maechling, Rajiv Mayani, Weiwei Chen, Rafael Ferreira da Silva, Miron Livny and Kent Wenger, Pegasus: a Workflow Management System for Science Automation, *Future Generation Computer Systems*, 2014.
11. Weiwei Chen, Rafael Ferreira da Silva, Ewa Deelman and Rizos Sakellariou, Using Imbalance Metrics to Optimize Task Clustering in Scientific Workflow Executions, *Future Generation Computer Systems*, 2014.
12. McLennan, Michael, Steven Clark, Ewa Deelman, Mats Rynge, Karan Vahi, Frank McKenna, Derrick Kearney, and Carol Song. "HUBzero and Pegasus: integrating scientific workflows into science gateways." *Concurrency and Computation: Practice and Experience* (2014).
13. Seyerle AA, Young AM, Jeff JM, Melton PE, Jorgensen NW, Lin Y, Carty CL, Deelman E, Heckbert SR, Hindorff LA, Jackson RD, Martin LW, Okin PM, Perez MV, Psaty BM, Soliman EZ, Whitsel EA, North KE, Laston S, Kooperberg C, Avery CL. Evidence of heterogeneity by race/ethnicity in genetic determinants of QT interval. *Epidemiology*. 2014 Nov;25(6):790-8.
14. Park SL, Caberto CP, Lin Y, Goodloe RJ, Dumitrescu L, Love SA, Matisse TC, Hindorff LA, Fowke JH, Schumacher FR, Beebe-Dimmer J, Chen C, Hou L, Thomas F, Deelman E, Han Y, Peters U, North KE, Heiss G, Crawford DC, Haiman CA, Wilkens LR, Bush WS, Kooperberg C, Cheng I, Le Marchand L. Association of Cancer Susceptibility Variants with Risk of Multiple Primary Cancers: The Population Architecture using Genomics and Epidemiology Study. *Cancer Epidemiol Biomarkers Prev*. 2014 Nov;23(11):2568-78.
15. Park SL, Fesinmeyer MD, Timofeeva M, Caberto CP, Kocarnik JM, Han Y, Love SA, Young A, Dumitrescu L, Lin Y, Goodloe R, Wilkens LR, Hindorff L, Fowke JH, Carty C, Buyske S, Schumacher FR, Butler A, Dilks H, Deelman E, Cote ML, Chen W, Pande M, Christiani DC, Field JK, Bickebller H, Risch A, Heinrich J, Brennan P, Wang Y, Eisen T, Houlston RS, Thun M, Albanes D, Caporaso N, Peters U, North KE, Heiss G, Crawford DC, Bush WS, Haiman CA, Landi MT, Hung RJ, Kooperberg C, Amos CI, Le Marchand L, Cheng I."Pleiotropic Associations of risk Variants identified for Other cancers With lung cancer risk: the PAGE and tricl consortia." *Journal of the National Cancer Institute* 106, no. 4 (2014): dju061.
16. Kocarnik JM, Pendergrass SA, Carty CL, Pankow JS, Schumacher FR, Cheng I, Durda P, Ambite JL, Deelman E, Cook NR, Liu S, Wactawski-Wende J, Hutter C, Brown-Gentry K, Wilson S, Best LG, Pankratz N, Hong CP, Cole SA, Voruganti VS, Bůžkova P, Jorgensen NW, Jenny NS, Wilkens LR, Haiman CA, Kolonel LN, Lacroix A, North K, Jackson R, Le Marchand L, Hindorff LA, Crawford DC, Gross M, Peters U. "Multi-Ancestral Analysis of Inflammation-Related Genetic Variants and C-Reactive Protein in the Population Architecture using Genomics and Epidemiology (PAGE) Study." *Circulation: Cardiovascular Genetics* (2014): CIRCGENETICS-113.
17. Lim U, Kocarnik JM, Bush WS, Matisse TC, Caberto C, Park SL, Carlson CS, Deelman E, Duggan D, Fesinmeyer M, Haiman CA, Henderson BE, Hindorff LA, Kolonel LN, Peters U, Stram DO, Tiirikainen M, Wilkens LR, Wu C, Kooperberg C, Le Marchand L., Pleiotropy of cancer susceptibility variants on the risk of non-Hodgkin lymphoma: the PAGE consortium., *PLoS*

- One. 2014 Mar 5;9(3):e89791. doi: 10.1371/journal.pone.0089791. eCollection 2014.
18. Malawski, Maciej, Kamil Figiela, Marian Bubak, Ewa Deelman, and Jarek Nabrzyski. "Cost Optimization of Execution of Multi-level Deadline-Constrained Scientific Workflows on Clouds." In *Parallel Processing and Applied Mathematics*, pp. 251-260. Springer Berlin Heidelberg, 2014.
 19. Gideon Juve, Mats Rynge, Ewa Deelman, Jens-S. Vockler, G. Bruce Berriman. Comparing FutureGrid, Amazon EC2, and Open Science Grid for Scientific Workflows. *Computing in Science and Engineering*, 15:4, pp. 20-29, 2013.
 20. Ewa Deelman, Gideon Juve, Maciej Malawski, Jarek Nabrzyski. Hosted Science: Managing Computational Workloads in the Cloud. *Parallel Processing Letters*, 23:2, June 2013.
 21. Karan Vahi, Ian Harvey, Taghrid Samak, Daniel Gunter, Kieran Evans, David Rogers, Ian Taylor, Monte Goode, Fabio Silva, Eddie Al-Shakarchi, Gaurang Mehta, Ewa Deelman, Andrew Jones. A Case Study into Using Common Real-Time Workflow Monitoring Infrastructure for Scientific Workflows. *Journal of Grid Computing: Volume 11, Issue 3 (2013)*, Page 381-406.
 22. Zhang L, Buzkova P, Wassel CL, Roman MJ, North KE, Crawford DC, Boston J, Brown-Gentry KD, Cole SA, Deelman E, Goodloe R, Wilson S, Heiss G, Jenny NS, Jorgensen NW, Matise TC, McClellan BE Jr, Nato AQ Jr, Ritchie MD, Franceschini N, Kao WH., Lack of associations of ten candidate coronary heart disease risk genetic variants and subclinical atherosclerosis in four U.S. populations: The Population Architecture using Genomics and Epidemiology (PAGE) study. *Atherosclerosis*. 2013 Jun;228(2):390-9. doi: 10.1016/j.atherosclerosis.2013.02.038. Epub 2013 Mar 13.
 23. SA Pendergrass, Brown-Gentry K, Dudek S, Frase A, Torstenson ES, Goodloe R, Ambite JL, Avery CL, Buyske S, Bůžková P, Deelman E, Fesinmeyer MD, Haiman CA, Heiss G, Hindorff LA, Hsu CN, Jackson RD, Kooperberg C, Le Marchand L, Lin Y, Matise TC, Monroe KR, Moreland L, Park SL, Reiner A, Wallace R, Wilkens LR, Crawford DC, Ritchie MD. Phenome-Wide Association Study (PheWAS) for Detection of Pleiotropy within the Population Architecture using Genomics and Epidemiology (PAGE) Network. *PLoS Genet*. 2013 Jan;9(1):e1003087. doi:10.1371/journal.pgen.1003087. Epub 2013 Jan 31. PubMed PMID: 23382687.
 24. Gideon Juve, Ann L. Chervenak, Ewa Deelman, Shishir Bharathi, Gaurang Mehta, Karan Vahi: Characterizing and profiling scientific workflows. *Future Generation Comp. Syst.* 29(3): 682-692, 2013
 25. G Bruce Berriman, Ewa Deelman, Gideon Juve, Mats Rynge, Jens Vöckler, The application of cloud computing to scientific workflows: a study of cost and performance. *Philos Transact A Math Phys Eng Sci*. 2012 Dec 10;371(1983):20120066. doi: 10.1098/rsta.2012.0066. Print 2013 Jan 28.
 26. Gideon Juve, Ewa Deelman, Bruce Berriman, Benjamin P. Berman, Phil Maechling, An Evaluation of the Cost and Performance of Scientific Workflows on Amazon EC2, *Journal of Grid Computing*, vol. 10, no. 1, pp. 5-21, 2012.
 27. Setiawan VW, Haessler J, Schumacher F, Cote ML, Deelman E, Fesinmeyer MD, Henderson BE, Jackson RD, Vöckler JS, Wilkens LR, Yasmeen S, Haiman CA, Peters U, Le Marchand L, Kooperberg C., . HNF1B and endometrial cancer risk: results from the PAGE study. *PLoS One*. 2012;7(1):e30390. doi: 10.1371/journal.pone.0030390. Epub 2012 Jan 27
 28. Gurmeet Singh, Ewa Deelman, The interplay of resource provisioning and workflow optimization in scientific applications, *Computation: Practice and Experience, Special Issue: Special Issue Fourth International Workshop On Workflow Management (ICWM2009)*, 23 (16), pages 1969–1989, November 2011
 29. Franceschini N, Carty C, Buzkova P, Reiner AP, Garrett T, Lin Y, Voeckler JS, Hindorff LA, Cole SA, Boerwinkle E, Lin DY, Bookman E, Best LG, Bella JN, Eaton C, Greenland P, Jenny N, North KE, Taverna D, Young AM, Deelman E, Kooperberg C, Psaty B, Heiss G. Association of genetic variants and incident coronary heart disease in multiethnic cohorts: the PAGE study. *Circ Cardiovasc Genet*. 2011 Dec;4(6):661-72. Epub 2011 Oct 31. PubMed PMID: 22042884.
 30. Setiawan VW, Haessler J, Schumacher F, Cote ML, Deelman E, Fesinmeyer MD, Henderson BE, Jackson RD, Voeckler JS, Wilkens LR, Yasmeen S, Haiman CA, Peters U, Le Marchand L, Kooperberg C. HNF1B and Endometrial Cancer Risk: Results from the PAGE study. *PLoS One*.

- 2012;7(1):e30390. Epub 2012 Jan 27. PubMed PMID: 22299039;
31. Faris J, Kolker E, Szalay A, Bradlow L, Deelman E, Feng W, Qiu J, Russell D, Stewart E, Kolker E. Communication and data-intensive science in the beginning of the 21st century. *OMICS*. 2011 Apr;15(4):213-5. PubMed PMID: 21476843.
 32. Ying Wang, Gaurang Mehta, Rajiv Mayani, Jingxi Lu, Tade Souaiaia, Yangho Chen, Andrew Clark, Hee Jae Yoon, Lin Wan, Oleg V. Evgrafov, James A. Knowles, Ewa Deelman, and Ting Chen, RseqFlow: Workflows for RNA-Seq data analysis, *Bioinformatics* (2011) first published online July 27, 2011 doi:10.1093/bioinformatics/btr441
 33. Yolanda Gil, Varun Ratnakar, Jihie Kim, Pedro Gonzalez-Calero, Paul Groth, Joshua Moody, and Ewa Deelman, Wings: Intelligent Workflow-Based Design of Computational Experiments. *IEEE Intelligent Systems*, 26(1), 2011.
 34. Eun-Kyu Byun, Yang-Suk Kee, Jin-Soo Kim, Ewa Deelman, Seungryoul Maeng: BTS: Resource capacity estimate for time-targeted science workflows. *J. Parallel Distrib. Comput.* (JPDC) 71(6):848-862 (2011)
 35. Saccone SF, Quan J, Mehta G, Bolze R, Thomas P, Deelman E, Tischfield JA, Rice JP, New tools and methods for direct programmatic access to the dbSNP relational database. *Nucleic Acids Res.* 2010 Oct 30.
 36. Saccone SF, Bolze R, Thomas P, Quan J, Mehta G, Deelman E, Tischfield JA, Rice JP. SPOT: a web-based tool for using biological databases to prioritize SNPs after a genome-wide association study. *Nucleic Acids Res.* June 2010
 37. Y. Gil; Ratnakar, V; Kim, J; Gonzalez-Calero, P; Groth, P; Moody, J; Deelman, E; WINGS: Intelligent Workflow-Based Design of Computational Experiments, *Intelligent Systems, IEEE*, 2010
 38. Gideon Juve, Ewa Deelman, Karan Vahi, Gaurang Mehta, Experiences with Resource Provisioning for Scientific Workflows Using Corral, *Scientific Programming*, 18:2, pp. 77-92, April 2010.
 39. Robert Graves, Thomas Jordan, Scott Callaghan, Ewa Deelman, Edward Field, Gideon Juve, Carl Kesselman, Philip Maechling, Gaurang Mehta, Kevin Milner, David Okaya, Patrick Small, Karan Vahi, CyberShake: A Physics-Based Seismic Hazard Model for Southern California, *Pure and Applied Geophysics*, May 2010.
 40. Vijay S. Kumar, Tahsin Kurc, Varun Ratnakar, Jihie Kim, Gaurang Mehta, Karan Vahi, Yoonju Lee Nelson, P. Sadayappan, Ewa Deelman, Yolanda Gil, Mary Hall and Joel Saltz Parameterized specification, configuration and execution of data-intensive scientific workflows, *Cluster Computing*, Volume 13, Number 3 / September, 2010
 41. Gideon Juve, Ewa Deelman, Scientific Workflows and Clouds, *ACM Crossroads*, 16:3, pp. 14-18, Spring 2010
 42. Ewa Deelman, Grids and Clouds: Making Workflow Applications Work in Heterogeneous Distributed Environments *International Journal of High Performance Computing Applications* Online First, published on December 4, 2009 as doi:10.1177/1094342009356432
 43. Scott Callaghan, Ewa Deelman, Dan Gunter, Gideon Juve, Philip Maechling, Christopher Brooks, Karan Vahi, Kevin Milner, Robert Graves, Edward Field, David Okaya, Thomas Jordan, Scaling up workflow-based applications, *Journal of Computer and System Sciences*, In Press, Corrected Proof, Available online 22 November 2009, ISSN 0022-0000, DOI: 10.1016/j.jcss.2009.11.005.
 44. Kevin Lee, Norman W. Paton, Rizos Sakellariou, Ewa Deelman, Alvaro A. A. Fernandes, Gaurang Mehta, Adaptive Workflow Processing and Execution in Pegasus, *Concurrency Practice and Experience*, volume 21, issue 16, 2009, pages 1965-1981.
 45. Gurmeet Singh, Carl Kesselman, and Ewa Deelman, An End-to-End Framework for Provisioning-Based Resource and Application Management, *IEEE Systems Journal*, (3) 1, March 2009
 46. Ewa Deelman, Dennis Gannon, Mathew Shields, Ian Taylor, Workflows and e-Science: An overview of workflow system features and capabilities, *Future Generations of Computer Systems*, July 2008.
 47. Simon Miles, Paul Groth, Ewa Deelman, Karan Vahi, Gaurang Mehta, and Luc Moreau.

- Provenance: The bridge between experiments and data. *Computing in Science and Engineering*, 2008.
48. Yolanda Gil, Ewa Deelman, Mark Ellisman, Thomas Fahringer, Geoffrey Fox, Dennis Gannon, Carole Goble, Miron Livny, Luc Moreau, Jim Myers, Examining the Challenges of Scientific Workflows, *IEEE Computer*, *IEEE Computer*, vol. 40, pp. 24-32, 2007.
 49. Gurmeet Singh, Karan Vahi, Arun Ramakrishnan, Gaurang Mehta, Ewa Deelman, Henan Zhao, Rizos Sakellariou, Kent Blackburn, Duncan Brown, Stephen Fairhurst, David Meyers, G. Bruce Berriman, John Good, Daniel S. Katz, Optimizing Workflow Data Footprint, *Special issue of the Scientific Programming Journal dedicated to Dynamic Computational Workflows: Discovery, Optimisation and Scheduling*, 2007
 50. Jihie Kim, Ewa Deelman, Yolanda Gil, Gaurang Mehta, Varun Ratnakar. Provenance Trails in the Wings/Pegasus Workflow System, *Concurrency and Computation: Practice and Experience, Special Issue on the First Provenance Challenge*, 2007.
 51. Joseph C. Jacob, Daniel S. Katz, G. Bruce Berriman, John Good, Anastasia C. Laity, Ewa Deelman, Carl Kesselman, Gurmeet Singh, Mei-Hui Su, Thomas A. Prince, Roy Williams, Montage: a grid portal and software toolkit for science-grade astronomical image mosaicking, *IJCE*, 2006
 52. Ewa Deelman, Tevfik Kosar, Carl Kesselman and Miron Livny, What Makes Workflows Work in an Opportunistic Environment? *Concurrency and Computation: Practice and Experience*, Volume 18, Issue 10, Pages 1187 – 1199. 2005
 53. Gurmeet Singh, Carl Kesselman, Ewa Deelman, Optimizing Grid-Based Workflow Execution, *Journal of Grid Computing*, Vol. 3, No. 3-4. (September 2005), pp. 201-219.
 54. Ewa Deelman, Gurmeet Singh, Mei-Hui Su, James Blythe, Yolanda Gil, Carl Kesselman, Gaurang Mehta, Karan Vahi, G. Bruce Berriman, John Good, Anastasia Laity, Joseph C. Jacob, Daniel S. Katz, Pegasus: a Framework for Mapping Complex Scientific Workflows onto Distributed Systems, *Scientific Programming Journal*, Volume 13, Number 3, 2005
 55. P. Maechling, H. Chalupsky, M. Dougherty, E. Deelman, Y. Gil, S. Gullapalli, V. Gupta, C. Kesselman, J. Kim, G. Mehta, B. Mendenhall, T. Russ, G. Singh, M. Spraragen, G. Staples, K. Vahi, Simplifying Construction of Complex Workflows for Non-Expert Users of the Southern California Earthquake Center Community Modeling Environment, *SIGMOD Record*, Volume 34 Number 3, September 2005
 56. Jim Blythe, Ewa Deelman, and Yolanda Gil, Automatically Composed Workflows for Grid Environments Described with Varying Levels of Detail, *IEEE Intelligent Systems*, 19(4): 16-23 (2004)
 57. Y. Gil, E. Deelman, J. Blythe, C. Kesselman, H. Tangmunarunkit. Artificial intelligence and grids: workflow planning and beyond, *IEEE Intelligent Systems*, Volume: 19, Issue: 1, Jan.-Feb. 2004, Pages:26 – 33
 58. A. Chervenak, E. Deelman, C. Kesselman, B. Allcock, I. Foster, V. Nefedova, J. Lee, A. Sim, A. Shoshani, B. Drach, D. Williams, D. Middleton, High-performance remote access to climate simulation data: a challenge problem for data grid technologies, *Parallel Computing* Volume 29, Issue 10 (October 2003), Pages: 1335-1356, 2003.
 59. E. Deelman, J. Blythe, Y. Gil, C. Kesselman, G. Mehta, K. Vahi, K. Blackburn, A. Lazzarini, A. Arbee, R. Cavanaugh, S. Koranda, Mapping Complex Workflows Onto Grid Environments, *Journal of Grid Computing*, Vol.1, No. 1, 2003., pp25-39.
 60. R. Williams, B. Berriman, E. Deelman, J. Good, J. Jacob, C. Kesselman, C. Lonsdale, S. Oliver, T. Prince. Multi-Wavelength Image Space: Another Grid-Enabled Science, *Journal of Concurrency and Computation: Practice and Experience*, Wiley, March 2003
 61. V.S. Adve, R. Bagrodia, E. Deelman, Rizos Sakellariou. Compiler-Optimized Simulation of Large-Scale Applications on High Performance Architectures. *Journal of Parallel and Distributed Computing*, Vol. 62, No. 3, Mar 2002, pp. 393-426.
 62. E. Deelman, B. K. Szymanski. Simulating Spatially Explicit Problems on High Performance Architectures, *Journal of Parallel and Distributed Computing*, Vol. 62, No. 3, Mar 2002, pp. 446-467.
 63. R. Bagrodia, E. Deelman, and T. Phan. Parallel Simulation of Large Scale Parallel Applications,

International Journal of High-Performance Computing Applications. Volume 15, Number, 1, Spring 2001.

64. V.S. Adve, R.Bagrodia, J.C. Browne, E. Deelman, A. Dube, E. Houstis, J.Rice, R. Sakellariou, D. Sundaram-Stukel, P. J. Teller, and M. K. Vernon. POEMS: End-to-end Performance Design of Large Parallel Adaptive Computational Systems, *IEEE Transactions on Software Engineering*, vol.26, no.11 p. 1027-48, November 2000.
65. S. Prakash, E. Deelman, and R. Bagrodia. Asynchronous Parallel Simulation of Parallel Programs. *IEEE Transactions on Software Engineering*, 26(5), May 2000.
66. T. Caraco, G. Gardner, W. Maniatty, E. Deelman, and B. K. Szymanski, Lyme disease: self-regulation and pathogen invasion. *Journal of Theoretical Biology*, 1998 Aug 21, 193(4):561-75.

CONFERENCE PAPERS

1. Anirban Mandal, Paul Ruth, Ilya Baldin, Yufeng Xin, Claris Castillo, Gideon Juve, Mats Rynge, Ewa Deelman and Jeff Chase, Adapting Scientific Workflows on Networked Clouds Using Proactive Introspection, IEEE/ACM Utility and Cloud Computing (UCC), 2015
2. Gideon Juve, Benjamin Tovar, Rafael Ferreira da Silva, Dariusz Krol, Douglas Thain, Ewa Deelman, William Allcock and Miron Livny, Practical Resource Monitoring for Robust High Throughput Computing, Workshop on Monitoring and Analysis for High Performance Computing Systems Plus Applications, 2015.
3. Rafael Ferreira da Silva, Mats Rynge, Gideon Juve, Igor Sfiligoi, Ewa Deelman, James Letts, Frank Wurthwein and Miron Livny, Characterizing a High Throughput Computing Workload: The Compact Muon Solenoid (CMS) Experiment at LHC, 2015 International Conference on Computational Science, 2015.
4. Sudarshan Srinivasan, Gideon Juve, Rafael Ferreira da Silva, Karan Vahi and Ewa Deelman, A Cleanup Algorithm for Implementing Storage Constraints in Scientific Workflow Executions, 9th Workshop on Workflows in Support of Large-Scale Science (WORKS), 2014.
5. Ilia Pietri, Gideon Juve, Ewa Deelman and Rizos Sakellariou, A Performance Model to Estimate Execution Time of Scientific Workflows on the Cloud, 9th Workshop on Workflows in Support of Large-Scale Science (WORKS), 2014.
6. Rafael Ferreira da Silva, Weiwei Chen, Gideon Juve, Karan Vahi, Ewa Deelman. Community Resources for Enabling Research in Distributed Scientific Workflows. 10th IEEE International Conference on e-Science (eScience 2014), Guarujá, Brazil, 2014
7. Idafen Santana-Perez, Rafael Ferreira da Silva, Mats Rynge, Ewa Deelman, Maria S. Perez-Hernandez, and Oscar Corcho. A Semantic-Based Approach to Attain Reproducibility of Computational Environments in Scientific Workflows: A Case Study. 1st International Workshop on Reproducibility in Parallel Computing (REPPAR), in conjunction with Euro-Par 2014, Porto, Portugal, 2014
8. Tristan Glatard, Lindsay B Lewis, Rafael Ferreira da Silva, Marc-Etienne Rousseau, Claude Lepage, Pierre Rioux, Najmeh Mahani, Ewa Deelman, Alan C Evans. Extending provenance information in CBRAIN to address reproducibility issues across computing platforms. NeuroInformatics 2014, Leiden, The Netherlands, 2014.
9. Idafen Santana-Perez, Rafael Ferreira da Silva, Mats Rynge, Ewa Deelman, Maria S. Perez-Hernandez, and Oscar Corcho. Leveraging Semantics to Improve Reproducibility in Scientific Workflows. The reproducibility at XSEDE workshop, Atlanta, USA, 2014
10. Weiwei Chen, Rafael Ferreira da Silva, Ewa Deelman, and Rizos Sakellariou. Balanced Task Clustering in Scientific Workflows. 9th IEEE International Conference on e-Science (eScience 2013), Beijing, China, Oct 24, 2013
11. Mats Rynge, Gideon Juve, Jamie Kinney, John Good, Bruce Berriman, Ann Merrihew, and Ewa Deelman. Producing an Infrared Multiwavelength Galactic Plane Atlas using Montage, Pegasus and Amazon Web Services. 23rd Annual Astronomical Data Analysis Software and Systems (ADASS) Conference, September 2013
12. Ilia Pietri, Maciej Malawski, Gideon Juve, Ewa Deelman, Jarek Nabrzyski, Rizos Sakellariou. Energy-Constrained Provisioning for Scientific Workflow Ensembles. IEEE International Conference on Cloud and Green Computing (CGC'13), September 2013.

13. Anirban Mandal, Paul Ruth, Ilya Baldin, Yufeng Xin, Claris Castillo, Mats Rynge, Ewa Deelman. Evaluating I/O Aware Network Management for Scientific Workflows on Networked Clouds. The 3rd International Workshop on Network-aware Data Management, in conjunction with SC'13, Denver, CO., November 2013
14. Rafael Ferreira Da Silva, Gideon Juve, Ewa Deelman, Tristan Glatard, Frederic Desprez, Douglas Thain, Benjamín Tovar and Miron Livny. Toward Fine-Grained Online Task Characteristics Estimation in Scientific Workflows. 8th Workshop On Workflows in Support of Large-Scale Science, November 2013.
15. Sepideh Azarnoosh, Mats Rynge, Gideon Juve, Ewa Deelman Michal Nieć, Maciej Malawski, Rafael Ferreira da Silva. Introducing PRECIP: An API for Managing Repeatable Experiments in the Cloud. Workshop on Cloud Computing for Research Collaborations (CRC), 2013.
16. Karan Vahi, Mats Rynge, Gideon Juve, Rajiv Mayani, and Ewa Deelman. Rethinking Data Management for Big Data Scientific Workflows. Workshop on Big Data and Science: Infrastructure and Services, September 2013
17. Michael McLennan, Steve Clark, Frank McKenna, Ewa Deelman, Mats Rynge, Karan Vahi, Derrick Kearney, Carol Song: Bringing Scientific Workflow to the Masses via Pegasus and HUBzero. Proceedings of the 5th International Workshop on Science Gateways, Zurich, Switzerland, 3-5 June, 2013.
18. G. Bruce Berriman, Carolyn Brinkworth, Dawn Gelino, Dennis K. Wittman, Ewa Deelman, Gideon Juve, Mats Rynge, Jamie Kinney, A Tale Of 160 Scientists, Three Applications, A Workshop and A Cloud, Astronomical Data Analysis Software and Systems XXII, 2012.
19. Ann L. Chervenak, David E. Smith, Weiwei Chen, Ewa Deelman, Integrating Policy with Scientific Workflow Management for Data-Intensive Applications, The 7th Workshop on Workflows in Support of Large-Scale Science (WORKS'12), Salt Lake City, November 10-16, 2012
20. Rohit Agarwal, Gideon Juve, Ewa Deelman, Peer-to-Peer Data Sharing for Scientific Workflows on Amazon EC2, 7th Workshop on Workflows in Support of Large-Scale Science (WORKS'12), 2012.
21. Weiwei Chen, Ewa Deelman, WorkflowSim: A Toolkit for Simulating Scientific Workflows in Distributed Environments, The 8th IEEE International Conference on eScience 2012, Chicago, Oct 8-12, 2012.
22. Taghrid Samak, Dan Gunter, Monte Goode, Ewa Deelman, Gideon Juve, Fabio Silva, Failure Analysis of Distributed Scientific Workflows Executing in the Cloud, 8th International Conference on Network and Service Management (CNSM 2012), 2012.
23. Maciej Malawski, Gideon Juve, Ewa Deelman, Jarek Nabrzyski, Cost- and Deadline-Constrained Provisioning for Scientific Workflow Ensembles in IaaS Clouds, 24th IEEE/ACM International Conference on Supercomputing (SC12), 2012.
24. Ewa Deelman, Gideon Juve, G. Bruce Berriman, Using Clouds for Science, Is it Just Kicking the Can Down The Road?, 2nd International Conference on Cloud Computing and Services Science (CLOSER 2012), 2012.
25. Mats Rynge, Gideon Juve, Karan Vahi, Scott Callaghan, Gaurang Mehta, Philip J. Maechling, Ewa Deelman, Enabling Large-scale Scientific Workflows on Petascale Resources Using MPI Master/Worker, XSEDE'12, July 2012.
26. Weiwei Chen, Ewa Deelman, Fault Tolerant Clustering in Scientific Workflows, IEEE 6th International Workshop on Scientific Workflows (SWF 2012) in conjunction with IEEE 8th World Congress on Services (SERVICES 2012), Honolulu, Hawaii, June 24, 2012
27. A Cloud-based Dynamic Workflow for Mass Spectrometry Data Analysis, Ashish Nagavaram, Gagan Agrawal, Michael Freitas, Gaurang Mehta, Rajiv Mayani, Ewa Deelman, Proceedings of the 7th IEEE International Conference on e-Science (e-Science 2011), December 2011.
28. Weiwei Chen, Ewa Deelman, Workflow Overhead Analysis and Optimizations, 6th Workshop on Workflows in Support of Large-Scale Science (WORKS 11), Seattle, Washington, November 14th, 2011.

29. Taghrid Samak, Dan Gunter, Monte Goode, Ewa Deelman, Gaurang Mehta, Fabio Silva, Karan Vahi, Failure Prediction and Localization in Large Scientific Workflows, 6th Workshop on Workflows in Support of Large-Scale Science (WORKS 11), Seattle, Washington, November 14th, 2011.
30. Mats Rynge, Gideon Juve, Gaurang Mehta, Ewa Deelman, Krista Larson, Burt Holzman, Igor Sfiligoi, Frank Würthwein, G. Bruce Berriman, Scott Callaghan, Experiences Using GlideinWMS and the Corral Frontend Across Cyberinfrastructures, Proceedings of the 7th IEEE International Conference on e-Science (e-Science 2011), December 2011.
31. Gideon Juve, Ewa Deelman Automating Application Deployment in Infrastructure Clouds, 3rd IEEE International Conference on Cloud Computing Technology and Science (CloudCom) 2011.
32. Dan Gunter, Christopher H. Brooks, Ewa Deelman, Monte Good, Gideon Juve, Gaurang Mehta, Priscilla Moaes, Taghrid Samak, Fabio Silva, Martin Swamy, Karan Vahi, Online Workflow Management and Performance Analysis with STAMPEDE, Proceedings of the 7th International Conference on Network and Service Management (CNSM 2011), 2011.
33. Taghrid Samak, Dan Gunter, Ewa Deelman, Gideon Juve, Gaurang Mehta, Fabio Silva, Karan Vahi, Online Fault and Anomaly Detection for Large-Scale Scientific Workflows, 13th IEEE International Conference on High Performance Computing and Communications (HPCC 2011), 2011.
34. BTS: Resource capacity estimate for time-targeted science workflows, Journal of Parallel and Distributed Computing, Eun-Kyu Byun, Yang-Suk Kee, Jin-Soo Kim, Ewa Deelman, Seungryoul Maeng, Volume 71, Issue 6, Special Issue on Cloud Computing, June 2011, Pages 848-862, ISSN 0743-7315, DOI: 10.1016/j.jpdc.2011.01.008.
35. Ten years of software sustainability at the Infrared Processing and Analysis Center, G. Bruce Berriman, John Good, Ewa Deelman and Anastasia Alexov, Phil. Trans. R. Soc. A 2011 369, 3384-3397, doi: 10.1098/rsta.2011.0136
36. Wrangler: Virtual Cluster Provisioning for the Cloud, Gideon Juve and Ewa Deelman, short paper, Proceedings of the 20th International Symposium on High Performance Distributed Computing (HPDC'11), 2011.
37. Experiences Using Cloud Computing for A Scientific Workflow Application, Jens-S. Vöckler, Gideon Juve, Ewa Deelman, Mats Rynge, G. Bruce Berriman, Proceedings of 2nd Workshop on Scientific Cloud Computing (ScienceCloud 2011), 2011.
38. G. Bruce Berriman, Gideon Juve, Ewa Deelman, Moira Regelson, Peter Plavchan, The Application of Cloud Computing to Astronomy: A Study of Cost and Performance, Workshop on e-Science challenges in Astronomy and Astrophysics in conjunction with the 6th IEEE International Conference on e-Science (e-Science 2010), December 2010.
39. Mirko Sonntag Dimka Karastoyanova and Ewa Deelman, Bridging The Gap Between Business And Scientific Workflows, e-Science 2010, Brisbane, Australia
40. Mirko Sonntag Dimka Karastoyanova and Ewa Deelman, BPEL4Pegasus: Combining Business and Scientific Workflows, International Conference on Service-Oriented Computing (ICSOC), San Francisco, California, December 2010.
41. Rizos Sakellariou, Henan Zhao, Ewa Deelman. Mapping Workflows on Grid Resources: Experiments with the Montage Workflow. In Grids, P2P and Services Computing, Springer, 2010, pp. 119-132.
42. G. Bruce Berriman, Ewa Deelman, Paul Groth, and Gideon Juve, The Application of Cloud Computing to the Creation of Image Mosaics and Management of Their Provenance SPIE Conference 7740: Software and Cyberinfrastructure for Astronomy, 2010.
43. Gideon Juve, Ewa Deelman, Karan Vahi, Gaurang Mehta, Bruce Berriman, Benjamin P. Berman, Phil Maechling, Data Sharing Options for Scientific Workflows on Amazon EC2, 22nd IEEE/ACM Conference on Supercomputing (SC10), New Orleans, Louisiana, November 2010.
44. Paul Groth, Ewa Deelman, Gideon Juve, Gaurang Mehta and Bruce Berriman, "A Pipeline-Centric Provenance Model", WORKS 2009

45. Gideon Juve, Ewa Deelman, Karan Vahi, Gaurang Mehta, Bruce Berriman, Benjamin P. Berman, Philip Maechling Scientific Workflow Applications on Amazon EC2, Cloud-based Services and Applications Workshop in Conjunction with e-Science 2009
46. Philip Maechling, Ewa Deelman, Yifeng Cui: Implementing Software Acceptance Tests as Scientific Workflows. PDPTA 2009: 317-323
47. Vijay S. Kumar, P. Sadayappan, Gaurang Mehta, Karan Vahi, Ewa Deelman, Varun Ratnakar, Jihie Kim, Yolanda Gil, Mary W. Hall, Tahsin M. Kurç, Joel H. Saltz: An integrated framework for performance-based optimization of scientific workflows. HPDC 2009: 177-186
48. Rizos Sakellariou and Henan Zhao and Ewa Deelman, Mapping Workflows on Grid Resources: Experiments with the Montage Workflow, CoreGrid 2009
49. R. Grave, S. Callaghan, E. Deelman, E. Field, N. Gupta, T.H. Jordan, G. Juve, C. Kesselman, P. Maechling, G. Mehta, D. Meyers, D. Okaya, K. Vahi, Physics Based Probabilistic Seismic Hazard Calculations for Southern California., 14th World Conference on Earthquake Engineering: Innovation Practice Safety. 2008
50. Scott Callaghan, Philip Maechling, Ewa Deelman, Karan Vahi, Gaurang Mehta, Gideon Juve, Kevin Milner, Robert Graves, Edward Field, David Okaya, Dan Gunter, Keith Beattie, Thomas Jordan, "Reducing Time-to-Solution Using Distributed High-Throughput Mega-Workflows - Experiences from SCEC CyberShake", Fourth IEEE International Conference on e-Science (e-Science 2008), 10-12 December 2008 in Indianapolis, Indiana, USA,
51. Gideon Juve, Ewa Deelman, "Resource Provisioning Options for Large-Scale Scientific Workflows", *Third International Workshop on Scientific Workflows and Business Workflow Standards in e-Science (SWBES)* in conjunction with Fourth IEEE International Conference on e-Science (e-Science 2008), 10 December 2008 in Indianapolis, Indiana, USA
52. Christina Hoffa, Gaurang Mehta, Timothy Freeman, Ewa Deelman, Kate Keahey, Bruce Berriman, John Good, "On the Use of Cloud Computing for Scientific Workflows", 3rd International Workshop on Scientific Workflows and Business Workflow Standards in e-Science (SWBES) in conjunction with Fourth IEEE International Conference on e-Science (e-Science 2008), 10 December 2008 in Indianapolis, Indiana, USA
53. Ewa Deelman, Gurmeet Singh, Miron Livny, Bruce Berriman, John Good. "The Cost of Doing Science on the Cloud: The Montage Example" Proceeding of Super Computing 2008, Austin , Texas, November 2008
54. Shishir Bharathi, Ann Chervenak, Ewa Deelman, Gaurang Mehta, Mei-Hui Su, Karan Vahi, "Characterization of Scientific Workflows", 3rd Workshop on Workflows in Support of Large-Scale Science (WORKS08), Austin, TX, November, 2008
55. Yang-Suk Kee, Ewa Deelman, Karan Vahi, Eun-kyu Byun, Jin-Soo Kim, "Pegasus on Virtual Grid: A Case Study of Workflow Planning over Captive Resources", 3rd Workshop on Workflows in Support of Large-Scale Science (WORKS08), Austin, TX, November 2008
56. Ewa Deelman, Ann Chervenak "Data Management Challenges of Large-Scale, Data-Intensive Scientific workflows" in WSES 08: 3rd International Workshop on Workflow Systems in e-Science, in conjunction with CCGrid 2008, May 2008, Lyon, France
57. K.Lee, N. W. Paton, R. Sakellariou, E. Deelman, A. A. A. Fernandes, G. Mehta "Adaptive Workflow Processing and Execution in Pegasus", 3rd International Workshop on Workflow Management and Applications in Grid Environments (WaGe08), in Proceedings of the Third International Conference on Grid and Pervasive Computing Symposia/Workshops, Pages 99-106, May 25-28 2008, Kunming, China
58. Vijay Kumar, Mary Hall, Jihie Kim, Yolanda Gil, Tahsin Kurc, Ewa Deelman and Joel Saltz, "Designing and Parameterizing a Workflow for Optimization: A Case Study in Biomedical Imaging", In Proc. of the Workshop on Next Generation Software, held in conjunction with IPDPS '08, April, 2008
59. Gurmeet Singh, Mei-Hui Su, Karan Vahi, Ewa Deelman, Bruce Berriman, John Good, Daniel S. Katz, and Gaurang Mehta, "Workflow Task Clustering for Best Effort Systems with Pegasus", Mardi Gras Conference, Baton Rouge, LA, January 2008

60. Simon Miles, Ewa Deelman, Paul Groth, Karan Vahi, Gaurang Mehta, Luc Moreau. Connecting Scientific Data to Scientific Experiments with Provenance, Third IEEE International Conference on e-Science and Grid Computing (e-Science 2007) Bangalore, India. , 2007
61. Ann Chervenak, Ewa Deelman, Miron Livny, Mei-Hui Su, Rob Schuler, Shishir Bharathi, Gaurang Mehta, Karan Vahi, Data Placement for Scientific Applications in Distributed Environments, Proceedings of Grid Conference 2007, Austin, Texas, September 2007.
62. Gurmeet Singh, Carl Kesselman, Ewa Deelman, Adaptive Pricing for Resource Reservations, 8th IEEE/ACM International Conference on Grid Computing (Grid 2007), Austin, Texas, September 2007.
63. Yolanda Gil, Varun Ratnakar, Ewa Deelman, Gaurang Mehta, and Jihie Kim. Wings for Pegasus: Creating Large-Scale Scientific Applications Using Semantic Representations of Computational Workflows, Proceedings of the 19th Annual Conference on Innovative Applications of Artificial Intelligence (IAAI), Vancouver, British Columbia, Canada, July 22-26, 2007.
64. Nandita Mandal, Ewa Deelman, Gaurang Mehta, Mei-Hui Su, and Karan Vahi, Integrating Existing Scientific Workflow Systems: The Kepler/Pegasus Example, Proceedings of the Second Workshop on Workflows in Support of Large-Scale Science (WORKS'07), in conjunction with the IEEE International Symposium on High Performance Distributed Computing Monterey, CA, June 2007.
65. Yolanda Gil, Pedro A. Gonzalez-Calero, Ewa Deelman. On the Black Art of Designing Computational Workflows, Proceedings of the Second Workshop on Workflows in Support of Large-Scale Science (WORKS'07), in conjunction with the IEEE International Symposium on High Performance Distributed Computing Monterey, CA, June 2007.
66. Gurmeet Singh, Carl Kesselman, Ewa Deelman, Adaptive Pricing for Resource Reservations The 8th IEEE/ACM International Conference on Grid Computing (Grid 2007)
67. Gurmeet Singh, Carl Kesselman, Ewa Deelman, A Provisioning Model and its Comparison with Best-Effort for Performance-Cost Optimization in Grids, in IEEE International Symposium on High Performance Distributed Computing (HPDC) 2007
68. Arun Ramakrishnan, Gurmeet Singh, Henan Zhao, Ewa Deelman, Rizos Sakellariou, K. Vahi, K. Blackburn, D. Meyers, and M. Samidi, Scheduling Data -Intensive Workflows onto Storage-Constrained Distributed Resources, in Seventh IEEE International Symposium on Cluster Computing and the Grid — CCGrid 2007
69. Ewa Deelman and Yolanda Gil, Managing Large-Scale Scientific Workflows in Distributed Environments: Experiences and Challenges, Proceedings of Workflows in e-Science, Amsterdam, The Netherlands, 2006.
70. E. Deelman, S. Callaghan, E. Field, H. Francoeur, R. Graves, N. Gupta, V. Gupta, T. H. Jordan, C. Kesselman, P. Maechling, J. Mehringer, G. Mehta, D. Okaya, K. Vahi, and L. Zhao, Managing Large-Scale Workflow Execution from Resource Provisioning to Provenance tracking: The CyberShake Example, Proceedings of e-Science, Amsterdam, The Netherlands, 2006. **(best paper)**
71. G. Singh, C. Kesselman, and E. Deelman, Application-level Resource Provisioning on the Grid, Proceedings of e-Science, Amsterdam, The Netherlands, 2006.
72. V. Nefedova, R. Jacob, I. Foster, Z. Liu, Y. Liu, E. Deelman, G. Mehta, M.-H. Su, and K. Vahi, Automating Climate Science: Large Ensemble Simulations on the TeraGrid with the GriPhyN Virtual Data System, Proceedings of e-Science, Amsterdam, The Netherlands, 2006.
73. A. Lathers, M.-H. Su, A. Kulungowski, A. W. Lin, G. Mehta, S. T. Peltier, Ewa Deelman, and M. H. Ellisman, Enabling Parallel Scientific Applications with Workflow Tools, Proceedings of Challenges of Large Applications in Distributed Environments (CLADE), Paris, 2006.
74. Y. Gil, V. Ratnakar, and E. Deelman, Metadata Catalogs with Semantic Representation, Proceedings of International Provenance Annotation Workshop (IPAW-06), Chicago, IL, 2006.
75. J. Muench, H. Francoeur, D. Okaya, Y. Cui, P. Maechling, E. Deelman, G. Mehta, R. Moore, and T. Jordan, SCEC Earthworks Science Gateway: Widening SCEC Community Access to the TeraGrid. TeraGrid 2006 Conference, 2006.

76. Yolanda Gil, Varun Ratnakar, and Ewa Deelman. Virtual Metadata Catalogs: Augmenting Metadata Catalogs with Semantic Representations, Short paper at the Fourth International Semantic Web Conference (ISWC-05), Galway, Ireland, November 7-10, 2005.
77. E. Deelman, A. Galstyan, Y. Gil, M. Hall, K. Lerman, A. Nakano, P. Vashista, and J. Saltz, A Systematic Approach to Composing and Optimizing Application Workflows, Proceedings of Workshop on Patterns in High Performance Computing, Urbana-Champaign, 2005.
78. Daniel S. Katz, Joseph C. Jacob, G. Bruce Berriman, John Good, Anastasia C. Laity, Ewa Deelman, Carl Kesselman, Gurmeet Singh, Mei-Hui Su, Thomas A. Prince, A Comparison of Two Methods for Building Astronomical Image Mosaics on a Grid. ICPP Workshops 2005: 85-94
79. Seung-Hye Jang, Valerie Taylor, Xingfu Wu, Mieke Prajugo, Ewa Deelman, Gaurang Mehta, and Karan Vahi, Performance Prediction-based versus Load-based Site Selection: Quantifying the Difference, the 18th International Conference on Parallel and Distributed Computing Systems (PDCS-2005), Las Vegas, Nevada, 12-14 September 2005.
80. Jim Blythe, Sonal Jain, Ewa Deelman, Karan Vahi, Yolanda Gil, Anirban Mandal, Ken Kennedy, Task Scheduling Strategies for Workflow-based Applications in Grids, IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2005).
81. Gurmeet Singh, Ewa Deelman, Gaurang Mehta, Karan Vahi, Mei-Hui Su, G. Bruce Berriman, John Good, Joseph C. Jacob, Daniel S. Katz, Albert Lazzarini, Kent Blackburn, and Scott Koranda, The Pegasus Portal: Web Based Grid Computing, The 20th Annual ACM Symposium on Applied Computing, SAC 2005.
82. Jim Blythe, Ewa Deelman, Yolanda Gil. Planning and Metadata on the Computational Grid, In AAAI Spring Symposium on Semantic Web Services, Palo Alto, California, USA, 2004.
83. Rattapoom Tuchinda, Snehal Thakkar, Yolanda Gil, and Ewa Deelman, Artemis: Integrating Scientific Data on the Grid, In the proceedings of the Sixteenth Innovative Applications of Artificial Intelligence, San Jose, CA, July 2004
84. Ewa Deelman, Gurmeet Singh, Malcolm P. Atkinson, Ann Chervenak, Neil P Chue Hong, Carl Kesselman, Sonal Patil, Laura Pearlman, Mei-Hui Su, Grid-Based Metadata Services, SSDBM04, Santorini, Greece, June 2004
85. I. Foster and others in the iVDGL project, The Grid2003 Production Grid: Principles and Practice, High Performance Distributed Computing, 2004. HPDC-13 2004.
86. J. C. Jacob, D. S. Katz, T. Prince, G. B. Berriman, J. C. Good, A. C. Laity, E. Deelman, G. Singh, and M.-H. Su, The Montage Architecture for Grid-Enabled Science Processing of Large, Distributed Datasets, Proceedings of the Earth Science Technology Conference (ESTC) 2004, June 2004.
87. G. B. Berriman, E. Deelman, J. Good, J. Jacob, D. S. Katz, C. Kesselman, A. Laity, T. A. Prince, G. Singh, and M. Su, Montage: a Grid Enabled Engine for Delivering Custom Science-Grade Image Mosaics on Demand, Proceedings of the SPIE Conference on Astronomical Telescopes and Instrumentation, June 2004.
88. Ewa Deelman, James Blythe, Yolanda Gil, Carl Kesselman, Gaurang Mehta, Sonal Patil, Mei-Hui Su, Karan Vahi, Miron Livny, Pegasus: Mapping Scientific Workflows onto the Grid, Across Grids Conference 2004, Nicosia, Cyprus.
89. Jim Blythe, Ewa Deelman, Yolanda Gil. Planning and Metadata on the Computational Grid, In AAAI Spring Symposium on Semantic Web Services, 2004.
90. Ewa Deelman, Raymond Plante, Carl Kesselman, Gurmeet Singh, Mei-Hui Su, Gretchen Greene, Robert Hanisch, Niall Gaffney, Antonio Volpicelli, James Annis, Vijay Sekhri, Tamas Budavari, Maria Nieto-Santisteban, William O'Mullane, David Bohlender, Tom McGlynn, Arnold Rots, Olga Pevunova, Grid-Based Galaxy Morphology Analysis for the National Virtual Observatory, Proceedings of SC 2003.
91. G. B. Berriman, J. C. Good, A. C. Laity, A. Bergou, J. Jacob, D. S. Katz, E. Deelman, C. Kesselman, G. Singh, M.-H. Su, R. Williams, Montage A Grid Enabled Image Mosaic Service for the National Virtual Observatory, ADASS XIII, ASP Conference Series Vol XXX, F Ochsenein M

- Allen and D Egret eds, 2003.
92. Gurmeet Singh, Shishir Bharathi, Ann Chervenak, Ewa Deelman, Carl Kesselman, Mary Manohar, Sonal Patil, and Laura Pearlman. A Metadata Catalog Service for Data Intensive Applications, Proceedings of SC 2003.
 93. Ewa Deelman, James Blythe, Yolanda Gil, Carl Kesselman, Scott Koranda, Albert Lazzarini, Gaurang Mehta, Maria Alessandra Papa, Karan Vahi, Pegasus and the Pulsar Search: From Metadata to Execution on the Grid., Applications Grid Workshop, PPAM 2003, Czestochowa, Poland 2003.
 94. J. Blythe, E. Deelman, Y. Gil, C. Kesselman, A. Agarwal, G. Mehta, K. Vahi, The Role of Planning in Grid Computing, ICAPS, 2003
 95. J. Blythe, E. Deelman, Y. Gil, C. Kesselman, Transparent Grid Computing: a Knowledge-Based Approach. IAAI, 2003.
 96. A. Chervenak, E. Deelman, I. Foster, L. Guy, W. Hoschek, A. Iamnitich, C. Kesselman, P. Kunst, M. Ripeanu, B. Schwartzkopf, H. Stockinger, K. Stockinger, B. Tierney. Giggle: A Framework for Constructing Scalable Replica Location Services. Proceedings of Supercomputing 2002 (SC2002), November 2002.
 97. E. Deelman, C. Kesselman, G. Mehta, L. Meshkat, L. Pearlman, K. Blackburn, P. Ehrens, A. Lazzarini, R. Williams, S. Koranda. GriPhyN and LIGO, building a virtual data grid for gravitational wave scientists. High Performance Distributed Computing, 2002. HPDC-11 2002. Page(s): 225 -234
 98. E. Deelman, C. Kesselman, S. Koranda, K. Blackburn A. Lazzarini, and R. Williams. Applications of Virtual Data in the LIGO Experiment, Proceedings of the International Conference on Parallel Processing and Applied Mathematics, Naleczow, Poland, September 2001, Lecture Notes in Computer Science. Vol.2328; Berlin, Germany : Springer-Verlag, 2002, xix+915 p. (23-34)
 99. W. Allcock, I. Foster, V. Nefedova, A. Chervenak, E. Deelman, C. Kesselman, J. Lee, A. Sim, A. Shoshani, B. Drach, D. Williams. High-Performance Remote Access to Climate Simulation Data: A Challenge Problem for Data Grid Technologies. Proceedings of SC 2001, Denver, CO, November 2001.
 100. E. Deelman, R. Bagrodia, R. Sakellariou, V. Adve. Improving Lookahead in Parallel Discrete Event Simulations of Large-Scale Applications Using Compiler Analysis. Proceedings 15th Workshop on Parallel and Distributed Simulation p. 5-13, Lake Arrowhead, CA, USA; 15-18 May 2001. **(best paper)**
 101. E. Deelman, C. Kesselman, R.D. Williams, A. Lazzarini, T.A. Prince, J. Romano, B. Allen, A Virtual Data Grid for LIGO, Lecture Notes in Computer Science 2110 (2001) 3, also Proc. HPCN 2001, Amsterdam, June 2001.
 102. V. S. Adve, R. Bagrodia, E. Deelman, T. Phan, and R. Sakellariou, Compiler-Supported Simulation of Highly Scalable Parallel Applications, Proceedings of SC'99, Portland, OR, Nov 13 - 19, 1999.
 103. R. Bagrodia, E. Deelman, S. Doco, T. Phan; Performance Prediction of Large Parallel Applications Using Parallel Simulations. ACM SIGPLAN 1999 Symposium on Principles and Practice of Parallel Programming, Atlanta, Georgia, May 4-6, 1999.
 104. E. Deelman, A. Dube, A. Hoisie, Y. Luo, R. Oliver, D. Sundaram-Stukel, H. Wasserman, V. S. Adve, R. Bagrodia, J. C. Browne, E. Houstis, O. Lubeck, J. Rice, P. Teller, M. K. Vernon; POEMS: End-to-end Performance Design of Large Parallel Adaptive Computational Systems. Proceedings of the First International Workshop on Software and Performance '98 - WOSP '98, Santa Fe, New Mexico, October 12-16, 1998.
 105. E. Deelman, and B.K. Szymanski, Dynamic load balancing in parallel discrete event simulation for spatially explicit problems. Proceedings of the Twelfth Workshop on Parallel and Distributed Simulation, PADS '98, Banff, CA, May 26-29, 1998.
 106. E. Deelman and B.K. Szymanski. System Knowledge Acquisition in Parallel Discrete Event Simulation. Proceedings of the 1997 IEEE International Conference on Systems Man and

- Cybernetics, Orlando, FL, October 12-15, 1997
107. E. Deelman and B.K. Szymanski. Continuously Monitored Global Virtual Time. Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications, Las Vegas, NV, June 30- July 3, 1997.
 108. E. Deelman and B.K. Szymanski. Breadth-First Rollback in Spatially Explicit Simulations. Proceedings of the 11th Workshop on Parallel and Distributed Simulation, PADS'97, Lockenhaus, Austria, June 10-13, 1997.
 109. E. Deelman, B.K. Szymanski and T. Caraco. Simulating Lyme Disease Using Parallel Discrete Event Simulation. Proceedings of the 1996 Winter Simulation Conference, San Diego, CA, December 8-11, 1996.
 110. B.K. Szymanski, E. Deelman, J. Flaherty, C. Norton, J. Teresco and L. Ziantz. Parallel Scientific Computing on the IBM SP2 at SCOREC-Rensselaer Polytechnic Institute. Proceedings of SUP'EUR 96, Krakow, Poland, September 8-11, 1996.
 111. E. Deelman, T. Caraco and B.K. Szymanski. Parallel Discrete Event Simulation of Lyme Disease. Biocomputing: Proc. of the 1996 Pacific Symposium, Hawaii, January 3-6, 1996.
 112. E. Deelman, W. K. Kaplow, B. K. Szymanski, P. Tannenbaum and L. Ziantz. Integrating Data and Task Parallelism in Scientific Programs. Proceeding of the Workshop on Languages, Compilers and Run-Time Systems for Parallel Programming, Kluwer Academic Publishers, Boston, 1995.

BOOK CHAPTERS

1. Ewa Deelman and Ann Chervenak, Data Management in Scientific Workflows. *Data Intensive Distributed Computing: Challenges and Solutions for Large-scale Information Management*. Ed. Tefik Kosar. Hershey: IGI Global, 2012.
2. Raphael Bolze and Ewa Deelman, Exploiting the Cloud of Computing Environments: An Application's perspective, *Cloud Computing and Software Services: Theory and Techniques*, Editors Syed A. Ahson and Mohammad Ilyas, CRC Press, July 2010
3. Ewa Deelman, Bruce Berriman, Ann Chervenak, Oscar Corcho, Paul Groth, Luc Moreau, Metadata and provenance management, in *Scientific Data Management: Challenges, Technology, and Deployment*, Editors - Arie Shoshani, Doron Rotem, Series: Chapman & Hall/CRC Computational Science, 2009
4. Bertram Ludascher, Ilkay Altintas, Shawn Bowers, Julian Cummings, Terence Critchlow, Ewa Deelman, David De Roure, Juliana Freire, Carole Goble, Matthew Jones, Scott Klasky, Timothy McPhillips, Norbert Podhorszki, Claudio Silva, Ian Taylor, Mladen Vouk, Scientific Process Automation and Workflow Management, in *Scientific Data Management: Challenges, Technology, and Deployment*, Editors - Arie Shoshani, Doron Rotem, Series: Chapman & Hall/CRC Computational Science, 2009
5. Ewa Deelman, Miron Livny, Gaurang Mehta, Andy Pavlo, Gurmeet Singh, Mei-Hui Su, Karan Vahi, R. Kent Wenger, Pegasus and DAGMan From Concept to Execution: Mapping Scientific Workflows onto Today's Cyberinfrastructure, in "High Performance Computing and Grids in Action" (L. Grandinetti Editor), IOS Press, Amsterdam, volume 16 in the series "Advances in Parallel Computing". March 2008
6. G. Bruce Berriman, Ewa Deelman, John Good, Joseph C. Jacob, Daniel S. Katz, Anastasia C. Laity, Thomas A. Prince, Gurmeet Singh, and Mei-Hui Su, Generating Complex Astronomy Workflows, in *Workflows for e-Science*, 2007
7. Philip Maechling, Ewa Deelman, Li Zhao, Robert Graves, Gaurang Mehta, Nitin Gupta, John Mehringer, Carl Kesselman, Scott Callaghan, David Okaya, Hunter Francoeur, Vipin Gupta, Yifeng Cui, Karan Vahi, Thomas Jordan, and Edward Field, SCEC CyberShake Workflows—Automating Probabilistic Seismic Hazard Analysis Calculations, in *Workflows for e-Science*, 2007
8. Ewa Deelman, Gaurang Mehta, Gurmeet Singh, Mei-Hui Su, and Karan Vahi, Pegasus: Mapping Large-Scale Workflows to Distributed Resources, in *Workflows for e-Science*, 2007
9. Ewa Deelman, Looking into the Future of Workflows: The Challenges Ahead, in *Workflows for e-Science*, 2007

10. Craig A. Lee, B. Scott Michel, Ewa Deelman and Jim Blythe, From Event-Driven Workflows Towards a Posteriori Computing, In Future Generation Grids as part of the Springer LNCS series, Getov, Vladimir; Laforenza, Domenico; Reinefeld, Alexander (Eds.) , 2006.
11. D.S. Katz, N. Anagnostou, G.B. Berriman, E. Deelman, J. Good, J.C. Jacob, C. Kesselman, A. Laity, T.A. Prince, G.Singh, M.-H. Su, and R. Williams Astronomical Image Mosaicking on a Grid: Initial Experiences in Engineering the Grid-Status and Perspectives, B. Di. Martino, J. Dongarra, A. Hoisie, L. Yand, and H. Zima eds. Nova, 2005
12. Ewa Deelman, James Blythe, Yolanda Gil, and Carl Kesselman, Workflow Management in GriPhyN, in Grid Resource Management, J. Nabrzyski, J. Schopf, and J. Weglarz editors, Kluwer, 2003.