

# Curriculum Vitae – Ioan Raicu, Ph.D.



## Assistant Professor

Illinois Institute of Technology (IIT)  
Department of Computer Science (CS)  
Data-Intensive Distributed Systems Laboratory (DataSys)

10 W. 31st Street  
Stuart Building 237D  
Chicago, IL 60616

Cellular: 1-847-722-0876  
Office: 1-312-567-5704  
Email: [iraicu@cs.iit.edu](mailto:iraicu@cs.iit.edu)



## Guest Research Faculty

Argonne National Laboratory (ANL)  
Math and Computer Science Division (MCS)  
Distributed Systems Laboratory (DSL)

Web: <http://www.cs.iit.edu/~iraicu/>

Lab Web: <http://datasys.cs.iit.edu/>

LinkedIn: <http://www.linkedin.com/in/ioanraicu>

Google: <http://scholar.google.com/citations?user=jE73HYAAAAAJ>

## Education

<b>Ph.D. in Computer Science</b>	University of Chicago	09/2005 – 03/2009
Dissertation: <i>“Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing”</i> Research Advisor: Ian Foster		
<b>Master of Science in Computer Science</b>	University of Chicago	09/2003 – 06/2005
Thesis: <i>“A Performance Study of the Globus Toolkit and Grid Services via DiPerF, an automated Distributed PERFORMANCE testing Framework”</i> Research Advisor: Ian Foster		
<b>Master of Science in Computer Science</b>	Wayne State University	09/2000 – 05/2002
Thesis: <i>“An Empirical Analysis of Internet Protocol version 6 (IPv6)”</i> Research Advisor: Sherali Zeadally		
<b>Bachelor of Science in Computer Science</b>	Wayne State University	09/1997 – 05/2000

## Awards

<b>NSF CAREER Award</b>	National Science Foundation, Office of Cyber Infrastructure	2011
<b>Computation Innovation Fellow</b>	National Science Foundation & Computing Research Association	2009
<b>GSRP Fellowship</b>	NASA, Ames Research Center	2006
<b>Graduate Student Fellowship</b>	Purdue University	2003
<b>Presidential Scholarship</b>	Wayne State University	1997
<b>Goldy Gemu Scholarship</b>	American Romanian Orthodox Youth	1997
<b>Stanitz Scholarship</b>	American Romanian Orthodox Youth	1997

## Work Experience

<b>Guest Research Faculty</b>	Argonne National Laboratory, Math and Computer Science Division	01/2011 – Present
<b>Assistant Professor</b>	Illinois Institute of Technology, Dept. of Computer Science	08/2010 – Present
<b>Computation Innovation Fellow</b>	Northwestern Univ., Dept. of Electrical Eng. and Computer Science	08/2009 – 07/2010
<b>Research Visitor</b>	NASA, Ames Research Center, NASA Advanced Supercomputing (NAS)	03/2009 – 05/2009
<b>Teaching/Research Assistant</b>	University of Chicago, Department of Computer Science	09/2003 – 03/2009
<b>Researcher (Internship)</b>	Argonne National Laboratory, Math and Computer Science Division	Summer 2005 & 2006
<b>Researcher (Internship)</b>	Sun Microsystems, Sun Labs & Solaris Networking Technologies	Summer 2003
<b>Teaching Assistant</b>	Purdue University, Department of Computer Science	08/2002 – 05/2003
<b>Adjunct Assistant Professor</b>	Univ. of Michigan, Department of Computer and Information Science	06/2002 – 08/2002
<b>Teaching/Research Assistant</b>	Wayne State University, Department of Computer Science	08/2000 – 08/2002
<b>Researcher (Internship)</b>	Accenture Technology Labs	Summer 2001
<b>System Analyst (Internship)</b>	Ford Motor Company	Summer 1999
<b>Owner</b>	High Teck Computers	01/1997 – 03/2001

# Research

## Research Interests


*Distributed Systems*  
*Grid Computing*  
*Many-Core Computing*

*Many-Task Computing*  
*Cloud Computing*  
*High-Performance Computing*

*Data-Intensive Computing*  
*Supercomputing*  
*Parallel Programming Languages*

## Publications

Over the past decade, I have co-authored many dozens of peer reviewed articles, book chapters, books, theses, and dissertations, which received 2219 citations. My [H-index](#) is 20, [G-Index](#) is 45, and [E-Index](#) is 38; these metrics are often used together as a useful index to characterize the scientific output of researchers. My publications can be found at <http://www.cs.iit.edu/~iraicu/research/publications/index.html>, and my Google Scholar Profile which includes all citation information can be found at <http://scholar.google.com/citations?hl=en&user=jE73HYAAAAAJ>.



### High-Impact Publications (ordered by citation counts):

<b>616 cites</b>	1. Ian Foster, Yong Zhao, <b>Ioan Raicu</b> , Shiyong Lu. "Cloud Computing and Grid Computing 360-Degree Compared", IEEE Grid Computing Environments (GCE08) 2008
<b>357 cites</b>	2. William Allcock, John Bresnahan, Rajkumar Kettimuthu, Michael Link, Catalin Dumitrescu, <b>Ioan Raicu</b> , Ian Foster, "The Globus Striped GridFTP Framework and Server," IEEE/ACM Supercomputing/SC, 2005
<b>190 cites</b>	3. Yong Zhao, Mihael Hategan, Ben Clifford, Ian Foster, Gregor von Laszewski, <b>Ioan Raicu</b> , Tibi Stef-Praun, Mike Wilde. "Swift: Fast, Reliable, Loosely Coupled Parallel Computation", IEEE Workshop on Scientific Workflows (SWF07) 2007
<b>155 cites</b>	4. <b>Ioan Raicu</b> , Yong Zhao, Catalin Dumitrescu, Ian Foster, Mike Wilde. "Falkon: a Fast and Light-weight task executiON framework", IEEE/ACM SuperComputing/SC, 2007
<b>154 cites</b>	5. Jennifer M. Schopf, <b>Ioan Raicu</b> , Laura Pearlman, Neill Miller, Carl Kesselman, Ian Foster, Mike D'Arcy. "Monitoring and Discovery in a Web Services Framework: Functionality and Performance of Globus Toolkit MDS4", Technical Report, Argonne National Laboratory, MCS Preprint #ANL/MCS-P1315-0106, January 2006
<b>78 cites</b>	6. <b>Ioan Raicu</b> , Ian Foster, Yong Zhao. "Many-Task Computing for Grids and Supercomputers", Invited Paper, IEEE Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS08), 2008
<b>52 cites</b>	7. Catalin Dumitrescu, <b>Ioan Raicu</b> , Matei Ripeanu, Ian Foster. "DiPerF: an automated Distributed PERFORMANCE testing Framework", IEEE/ACM GRID04, 2004
<b>48 cites</b>	8. Catalin Dumitrescu, <b>Ioan Raicu</b> , Ian Foster. "DI-GRUBER: A Distributed Approach for Grid Resource Brokering", IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SuperComputing/SC05), 2005
<b>48 cites</b>	9. <b>Ioan Raicu</b> , Zhao Zhang, Mike Wilde, Ian Foster, Pete Beckman, Kamil Iskra, Ben Clifford. "Towards Loosely-Coupled Programming on Petascale Systems", IEEE/ACM SC 2008
<b>45 cites</b>	10. Sherali Zeadally, <b>Ioan Raicu</b> . "Evaluating IPV6 on Windows and Solaris", IEEE Internet Computing, Volume 7, Issue 3, May/June 2003, pp 51 – 57
<b>43 cites</b>	11. Michael Wilde, Ian Foster, Kamil Iskra, Pete Beckman, Zhao Zhang, Allan Espinosa, Mihael Hategan, Ben Clifford, <b>Ioan Raicu</b> . "Parallel Scripting for Applications at the Petascale and Beyond", IEEE Computer Nov. 2009 Special Issue on Extreme Scale Computing, 2009
<b>36 cites</b>	12. <b>Ioan Raicu</b> , Yong Zhao, Ian Foster, Alex Szalay. "Accelerating Large-scale Data Exploration through Data Diffusion", IEEE International Workshop on Data-Aware Distributed Computing (DADC08) 2008
<b>31 cites</b>	13. Yong Zhao, <b>Ioan Raicu</b> , Ian Foster. "Scientific Workflow Systems for 21st Century e-Science, New Bottle or New Wine?", Invited Paper, IEEE Workshop on Scientific Workflows (SWF08) 2008
<b>27 cites</b>	14. <b>Ioan Raicu</b> , Sherali Zeadally. "Evaluating IPv4 to IPv6 Transition Mechanisms", IEEE International Conference on Telecommunications 2003, ICT'2003, Volume 2, pp 1091 – 1098, Feb 2003
<b>26 cites</b>	15. Sherali Zeadally, R. Wasseem, <b>Ioan Raicu</b> . "Comparison of End-System IPv6 Protocol Stacks", IEE Proceedings Communications, Special issue on Internet Protocols, Technology and Applications (VoIP), Vol. 151, No. 3, June 2004

<b>23 cites</b>	16. Zhao Zhang, Allan Espinosa, Kamil Iskra, <b>Ioan Raicu</b> , Ian Foster, Michael Wilde. “ <i>Design and Evaluation of a Collective I/O Model for Loosely-coupled Petascale Programming</i> ”, IEEE Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS08), 2008
<b>22 cites</b>	17. <b>Ioan Raicu</b> , Loren Schwiebert, Scott Fowler, Sandeep K.S. Gupta. “ <i>Local Load Balancing for Globally Efficient Routing in Wireless Sensor Networks</i> ”, International Journal of Distributed Sensor Networks, 1: 163–185, 2005
<b>22 cites</b>	18. Alex Szalay, Julian Bunn, Jim Gray, Ian Foster, <b>Ioan Raicu</b> . “ <i>The Importance of Data Locality in Distributed Computing Applications</i> ”, NSF Workflow Workshop 2006
<b>21 cites</b>	19. Catalin Dumitrescu, <b>Ioan Raicu</b> , Ian Foster. “ <i>Experiences in Running Workloads over Grid3</i> ”, The 4th International Conference on Grid and Cooperative Computing (GCC 2005)
<b>20 cites</b>	20. <b>Ioan Raicu</b> , Ian Foster, Yong Zhao, Philip Little, Christopher Moretti, Amitabh Chaudhary, Douglas Thain. “ <i>The Quest for Scalable Support of Data Intensive Workloads in Distributed Systems</i> ”, ACM International Symposium on High Performance Distributed Computing (HPDC09), 2009

The list below includes all peer-reviewed publications.

### Books

- Ioan Raicu**. “*Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing*”, ISBN: 978-3-639-15614-0, VDM Verlag Dr. Muller Publisher, 2009

### Editorials

- Tevfik Kosar, **Ioan Raicu**. “Guest Editors' Introduction: Special Issue on Data-Intensive Computing in the Clouds”, Springer Journal of Grid Computing, 2012
- Ioan Raicu**, Ian T. Foster, Yong Zhao. “*Guest Editors' Introduction: Special Issue on Many-Task Computing*”, IEEE Transactions on Parallel and Distributed Systems, 2011
- Ivona Brandic, **Ioan Raicu**. “*Guest Editors' Introduction: Special Issue on Science-driven Cloud Computing*”, Scientific Programming Journal, 2011

### Journal Articles

- Michael Wilde, Ian Foster, Pete Beckman, **Ioan Raicu**. *Swift: Scalable Parallel Scripting for Scientific Computing*, SciDAC Review, pp 38-53, Spring 2010
- Ioan Raicu**, Ian Foster, Mike Wilde, Zhao Zhang, Alex Szalay, Kamil Iskra, Pete Beckman, Yong Zhao, Alok Choudhary, Philip Little, Christopher Moretti, Amitabh Chaudhary, Douglas Thain. “*Middleware Support for Many-Task Computing*”, Cluster Computing, The Journal of Networks, Software Tools and Applications, 2010
- Michael Wilde, Ian Foster, Kamil Iskra, Pete Beckman, Zhao Zhang, Allan Espinosa, Mihael Hategan, Ben Clifford, **Ioan Raicu**. “*Parallel Scripting for Applications at the Petascale and Beyond*”, IEEE Computer Nov. 2009 Special Issue on Extreme Scale Computing, 2009
- Catalin Dumitrescu, **Ioan Raicu**, Ian Foster. “*The Design, Usage, and Performance of GRUBER: A Grid uSLA-based Brokering Infrastructure*”, International Journal of Grid Computing, 2007
- Ioan Raicu**, Catalin Dumitrescu, Matei Ripeanu, Ian Foster. “*The Design, Performance, and Use of DiPerF: An automated Distributed PERFORMANCE testing Framework*”, International Journal of Grid Computing, Special Issue on Global and Peer-to-Peer Computing, 2006
- Catalin Dumitrescu, **Ioan Raicu**, Ian Foster. “*Usage SLA-based Scheduling in Grids*”, Journal on Concurrency and Computation: Practice and Experience, 2006
- Ioan Raicu**, Loren Schwiebert, Scott Fowler, Sandeep K.S. Gupta. “*Local Load Balancing for Globally Efficient Routing in Wireless Sensor Networks*”, International Journal of Distributed Sensor Networks, 1: 163–185, 2005
- Sheralli Zeadally, R. Wasseem, **Ioan Raicu**. “*Comparison of End-System IPv6 Protocol Stacks*”, IEE Proceedings Communications, Special issue on Internet Protocols, Technology and Applications (VoIP), Vol. 151, No. 3, June 2004
- Sherali Zeadally, **Ioan Raicu**. “*Evaluating IPV6 on Windows and Solaris*”, IEEE Internet Computing, Volume 7, Issue 3, May June 2003, pp 51 – 57

### Book Chapters

- Ioan Raicu**, Ian Foster, Yong Zhao, Alex Szalay, Philip Little, Christopher M. Moretti, Amitabh Chaudhary, Douglas Thain. “*Towards Data Intensive Many-Task Computing*”, book chapter in Data Intensive Distributed Computing: Challenges and Solutions for Large-Scale Information Management, IGI Global Publishers, 2011
- Yong Zhao, **Ioan Raicu**, Ian Foster, Mihael Hategan, Veronika Nefedova, Mike Wilde. “*Realizing Fast, Scalable and Reliable Scientific Computations in Grid Environments*”, book chapter in Grid Computing Research Progress, ISBN:

16. Catalin Dumitrescu, Jan Dünneweber, Philipp Lüdeking, Sergei Gorlatch, **Ioan Raicu**, Ian Foster. *"Simplifying Grid Application Programming Using Web-Enabled Code Transfer Tools"*. Toward Next Generation Grids, Chapter 6, Springer Verlag, 2007
- 

### Conference/Workshop Papers

17. Hui Jin, Xi Yang, Xian-He Sun, **Ioan Raicu**. *"ADAPT: Availability-aware MapReduce Data Placement in Non-Dedicated Distributed Computing Environment"*, IEEE International Conference on Distributed Computing Systems (ICDCS) 2012
18. Yong Zhao, **Ioan Raicu**, Shiyong Lu, Xubo Fei. *"Opportunities and Challenges in Running Scientific Workflows on the Cloud"*, IEEE International Conference on Network-based Distributed Computing and Knowledge Discovery (CyberC) 2011
19. **Ioan Raicu**, Pete Beckman, Ian Foster. *"Making a Case for Distributed File Systems at Exascale"*, Invited Paper, ACM Workshop on Large-scale System and Application Performance (LSAP), 2011
20. Michael Wilde, **Ioan Raicu**, Allan Espinosa, Zhao Zhang, Ben Clifford, Mihael Hategan, Kamil Iskra, Pete Beckman, Ian Foster. *"Extreme-scale scripting: Opportunities for large task-parallel applications on petascale computers"*, Scientific Discovery through Advanced Computing Conference (SciDAC09), 2009
21. **Ioan Raicu**, Ian Foster, Yong Zhao, Philip Little, Christopher Moretti, Amitabh Chaudhary, Douglas Thain. *"The Quest for Scalable Support of Data Intensive Workloads in Distributed Systems"*, ACM International Symposium on High Performance Distributed Computing (HPDC09), 2009
22. **Ioan Raicu**, Ian Foster, Yong Zhao. *"Many-Task Computing for Grids and Supercomputers"*, Invited Paper, IEEE Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS08), 2008
23. Ian Foster, Yong Zhao, **Ioan Raicu**, Shiyong Lu. *"Cloud Computing and Grid Computing 360-Degree Compared"*, IEEE Grid Computing Environments (GCE08) 2008
24. **Ioan Raicu**, Zhao Zhang, Mike Wilde, Ian Foster, Pete Beckman, Kamil Iskra, Ben Clifford. *"Towards Loosely-Coupled Programming on Petascale Systems"*, IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SuperComputing/SC08), 2008
25. Zhao Zhang, Allan Espinosa, Kamil Iskra, **Ioan Raicu**, Ian Foster, Michael Wilde. *"Design and Evaluation of a Collective I/O Model for Loosely-coupled Petascale Programming"*, IEEE Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS08), 2008
26. Yong Zhao, **Ioan Raicu**, Ian Foster. *"Scientific Workflow Systems for 21st Century e-Science, New Bottle or New Wine?"*, Invited Paper, IEEE Workshop on Scientific Workflows (SWF08) 2008
27. **Ioan Raicu**, Yong Zhao, Ian Foster, Alex Szalay. *"Accelerating Large-scale Data Exploration through Data Diffusion"*, IEEE International Workshop on Data-Aware Distributed Computing (DADC08) 2008
28. **Ioan Raicu**, Yong Zhao, Catalin Dumitrescu, Ian Foster, Mike Wilde. *"Falcon: a Fast and Light-weight task executiON framework"*, IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SuperComputing/SC07), 2007
29. Yong Zhao, Mihael Hategan, Ben Clifford, Ian Foster, Gregor von Laszewski, **Ioan Raicu**, Tiberiu Stef-Praun, Mike Wilde. *"Swift: Fast, Reliable, Loosely Coupled Parallel Computation"*, IEEE Workshop on Scientific Workflows (SWF07) 2007
30. **Ioan Raicu**, Ian Foster, Alex Szalay, Gabriela Turcu. *"AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis"*, TeraGrid Conference 2006, June 2006
31. Alex Szalay, Julian Bunn, Jim Gray, Ian Foster, **Ioan Raicu**. *"The Importance of Data Locality in Distributed Computing Applications"*, NSF Workflow Workshop 2006
32. William Allcock, John Bresnahan, Rajkumar Kettimuthu, Michael Link, Catalin Dumitrescu, **Ioan Raicu**, Ian Foster, *"The Globus Striped GridFTP Framework and Server,"* sc, p. 54, IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SuperComputing/SC05), 2005
33. Catalin Dumitrescu, **Ioan Raicu**, Ian Foster. *"DI-GRUBER: A Distributed Approach for Grid Resource Brokering"*, IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SuperComputing/SC05), 2005
34. Catalin Dumitrescu, **Ioan Raicu**, Ian Foster. *"Experiences in Running Workloads over Grid3"*, The 4th International Conference on Grid and Cooperative Computing (GCC 2005)
35. Catalin Dumitrescu, **Ioan Raicu**, Matei Ripeanu, Ian Foster. *"DiPerF: an automated Distributed PERformance testing Framework"*, IEEE/ACM GRID04, pp 289 – 296, 2004
36. **Ioan Raicu**, Loren Schwiebert, Scott Fowler, Sandeep K.S. Gupta. *"e3D: An Energy-Efficient Routing Algorithm for Wireless Sensor Networks"*, IEEE ISSNIP 2004 (The International Conference on Intelligent Sensors, Sensor
-

---

Networks and Information Processing), 2004

37. **Ioan Raicu**, Sherali Zeadally. "Impact of IPv6 on End-User Applications", IEEE International Conference on Telecommunications 2003, ICT'2003, Volume 2, pp 973 – 980, Feb 2003
  38. **Ioan Raicu**, Sherali Zeadally. "Evaluating IPv4 to IPv6 Transition Mechanisms", IEEE International Conference on Telecommunications 2003, ICT'2003, Volume 2, pp 1091 – 1098, Feb 2003
  39. **Ioan Raicu**. "Efficient Even Distribution of Power Consumption in Wireless Sensor Networks", ISCA 18th International Conference on Computers and Their Applications, CATA 2003
  40. **Ioan Raicu**, Owen Richter, Loren Schwiebert, Sherali Zeadally. "Using Wireless Sensor Networks to Narrow the Gap between Low-Level Information and Context-Awareness", Proceedings of the ISCA 17th International Conference, Computers and their Applications, 2002
- 

### Posters / Extended Abstracts

41. Tonglin Li, Raman Verma, Xi Duan, Hui Jin, **Ioan Raicu**. "Exploring Distributed Hash Tables in High-End Computing", SIGMETRICS Performance Evaluation Review-Measurement and Evaluation 39(3), 128, 2011
  42. Hui Jin, Xi Yang, Xian-He Sun, **Ioan Raicu**. "An Empirical Evaluation of MapReduce under Interruptions", Cloud Computing and its Applications (CCA) 2011
  43. **Ioan Raicu**, Yong Zhao, Ian Foster, Mike Wilde, Zhao Zhang, Ben Clifford, Mihael Hategan, Sarah Kenny. "Managing and Executing Loosely Coupled Large Scale Applications on Clusters, Grids, and Supercomputers", Extended Abstract, GlobusWorld08, part of Open Source Grid and Cluster Conference 2008
  44. Quan T. Pham, Atilla S. Balkir, Jing Tie, Ian Foster, Mike Wilde, **Ioan Raicu**. "Data Intensive Scalable Computing on TeraGrid: A Comparison of MapReduce and Swift", TeraGrid Conference (TG08) 2008
  45. **Ioan Raicu**, Yong Zhao, Ian Foster, Alex Szalay. "A Data Diffusion Approach to Large Scale Scientific Exploration", Extended Abstract, Microsoft Research eScience Workshop (MSES07) 2007
  46. Catalin Dumitrescu, Alexandru Iosup, H. Mohamed, Dick H.J. Epema, Matei Ripeanu, Nicolae Tapus, **Ioan Raicu**, Ian Foster. "ServMark: A Framework for Testing Grids Services", IEEE Grid 2007
  47. **Ioan Raicu**, Catalin Dumitrescu, Ian Foster. "Dynamic Resource Provisioning in Grid Environments", TeraGrid Conference (TG07) 2007
  48. **Ioan Raicu**, Ian Foster, Alex Szalay. "Harnessing Grid Resources to Enable the Dynamic Analysis of Large Astronomy Datasets", IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC06), 2006
  49. **Ioan Raicu**. "Routing Algorithms for Wireless Sensor Networks", Poster Presentation, Grace Hopper Celebration of Women in Computing (GHC02), 2002
- 

### Other (Thesis, Technical Reports)

50. **Ioan Raicu**. "Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing", Computer Science Dept., University of Chicago, Doctorate Dissertation, March 2009
  51. Jennifer M. Schopf, **Ioan Raicu**, Laura Pearlman, Neill Miller, Carl Kesselman, Ian Foster, Mike D'Arcy. "Monitoring and Discovery in a Web Services Framework: Functionality and Performance of Globus Toolkit MDS4", Technical Report, Argonne National Laboratory, MCS Preprint #ANL/MCS-P1315-0106, January 2006
  52. **Ioan Raicu**. "A Performance Study of the Globus Toolkit® and Grid Services via DiPerF, an automated Distributed PERFORMANCE testing Framework", University of Chicago, Computer Science Department, MS Thesis, May 2005
  53. **Ioan Raicu**. "An Empirical Analysis of Internet Protocol version 6 (IPv6)", Wayne State University, Computer Science Department, MS Thesis, May 2002
- 

### Under Review / Preparation

54. Tonglin Li, Hui Jin, Antonio Perez De Tejada, Kevin Brandstatter, Zhao Zhang, **Ioan Raicu**. "ZHT: A Zero-Hop Distributed Hash Table for High-End Computing", under preparation for IEEE/ACM SC 2012
  55. Ke Yue, **Ioan Raicu**. "Scheduling Direct Acyclic Graphs on Massively Parallel 1K-core Processors", under preparation for IEEE/ACM SC 2012
  56. Ke Wang, **Ioan Raicu**. "SimMatrix: Exploring Many-Task Computing through Simulations at Exascales", under review at ACM HPDC 2012
  57. Yong Zhao, **Ioan Raicu**, Shiyong Lu, Xubo Fei. "Cloud Workflows: A Marriage between Scientific Workflows and Cloud Computing", under review at the Journal of Concurrency and Computation: Practice and Experience (CCPE)
  58. **Ioan Raicu**, Ian Foster, Yong Zhao. "Many-Task Computing: A Survey", under preparation for IEEE TPDS
  59. Zhao Zhang, Ian Foster, Daniel Katz, Mike Wilde, **Ioan Raicu**. "AME: An Anyscale Many-Task Computing Engine", under preparation
-

## Presentations

1. Introduction to Computer Science, the What, How, and Why of CS, CAMRAS Interview Weekend, Illinois Institute of Technology, February 10<sup>th</sup>, 2012
2. Undergraduate Research in the DataSys Laboratory, CAMRAS Students Luncheon, Illinois Institute of Technology, January 18<sup>th</sup>, 2012
3. Distributed and Cloud Computing Specialization at IIT, Introduction to the MCS Specializations Seminar, Illinois Institute of Technology, September 20<sup>th</sup>, 2011
4. An Overview of Distributed Systems, Illinois Institute of Technology, September 13<sup>th</sup>, 2011
5. Challenges and Opportunities in Large-Scale Storage Systems, ECE Seminar, Illinois Institute of Technology, September 9<sup>th</sup>, 2011
6. Making a Case for Distributed File Systems at Exascales, SCS Seminar, Illinois Institute of Technology, June 17<sup>th</sup>, 2011
7. Making a Case for Distributed File Systems at Exascales, ACM LSAP 2011, June 8th, 2011
8. CCGrid 2014 Proposal, CCGrid 2011 Steering Committee Meeting, May 25<sup>th</sup>, 2011
9. Common Challenges between Exascales and Cloud Computing, IEEE DataCloud 2011, May 16th, 2011
10. Making a Case for Distributed File Systems at Exascales, University of Chicago, May 13th, 2011
11. Making a Case for Distributed File Systems at Exascales, Indiana University, April 27th, 2011
12. DataSys: Data-Intensive Distributed Systems Laboratory Overview, Research Forum, Illinois Institute of Technology, March 22nd, 2011
13. Cloud Computing and Grid Computing 360-Degree Compared, Computer Science Department, Loyola University, March 17th, 2011
14. Reading/Writing, Mathematics, and Computing, CAMRAS Interview Weekend, Illinois Institute of Technology, February 17<sup>th</sup>, 2011
15. Cloud Computing and Grid Computing 360-Degree Compared, IEEE Fox Valley South Section, Illinois Institute of Technology, January 26th, 2011
16. Avoiding Achilles' Heel in Exascale Computing with Distributed File Systems, DLS Seminar, January 14th, 2011
17. Research Overview, Research Forum, Illinois Institute of Technology, October 18th, 2010
18. Common Challenges in Manycore, Exascale, and Cloud Computing, Computer Science Department, Illinois Institute of Technology, October 4th, 2010
19. Rethinking Storage Systems for Exascale Computing, HPDC 2010, Wild and Crazy Ideas Session, June 25th, 2010
20. Grid, cloud, and science: Accelerating discovery A View and Practice from University of Chicago, The Forum on Cloud in Academia at Illinois Institute of Technology, April 28th, 2010
21. Exascale Many-Task Computing with a Billion Processors, HPDC Program Committee Workshop 2010, March 22nd, 2010
22. Many-Task Computing on Grids, Clouds, and Supercomputers, Department of Computer Science, Illinois Institute of Technology, January 22nd, 2010
23. Scalable Resource Management in Cloud Computing, Grid Computing and Supercomputing, College of Computing and Digital Media, DePaul University, January 20th, 2010
24. Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing, CUCIS Seminar, Department of Electrical Engineering and Computer Science, Northwestern University, September 21st, 2009
25. The Quest for Scalable Support of Data Intensive Workloads in Distributed Systems, ACM HPDC09, June 13th, 2009
26. Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing, Department of Computer Science, University of Nebraska at Omaha, March 10th, 2009
27. The Quest for Scalable Support of Data Intensive Workloads in Distributed Systems, Motorola Labs, March 4th, 2009
28. Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing, Department of Mathematics, Computer Science, and Statistics, Purdue University Calumet, March 3rd, 2009
29. Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing, Department of Computer Science, Colorado State University, March 2nd, 2009
30. Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing, The Center of Advanced Computer Studies, University of Louisiana at Lafayette, February 27th, 2009

31. Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing, Dissertation Defense, Computer Science Department, University of Chicago, February 12th, 2009
32. Scalable Resource Management in Clouds and Grids, Motorola Labs, December 5th, 2008.
33. The Quest for Scalable Support of Data Intensive Applications through Data Diffusion, IEEE/ACM Supercomputing 2008, Argonne National Laboratory Booth, November 20th, 2008.
34. Falcon, a Fast and Light-weight task executiON framework for Clusters, Grids, and Supercomputers, IEEE/ACM Supercomputing 2008, Argonne National Laboratory Booth, November 19th, 2008.
35. Running 1 Million Jobs in 10 Minutes via the Falcon Fast and Light-weight task executiON framework, IEEE/ACM Supercomputing 2008, Megajob BOF, November 18th, 2008.
36. Toward Loosely Coupled Programming on Petascale Systems, IEEE/ACM Supercomputing 2008, November 18th, 2008.
37. Cloud Computing and Grid Computing 360-Degree Compared, IEEE Grid Computing Environments (GCE08) 2008, November 16th, 2008.
38. The Quest for Scalable Support of Data Intensive Applications through Data Diffusion, University of Chicago, Petascale Active Data Storage (PADS) Seminar, November 12th, 2008.
39. Cloud Computing and Grid Computing 360-Degree Compared, University of Chicago, Graduate CS Seminar, November 10th, 2008.
40. Scalable Resource Management in Clouds and Grids, Accenture Technology Labs, October 24th, 2008.
41. Systems at University of Chicago, University of Chicago, Department of Computer Science, September 25th, 2008.
42. Scientific Workflow Systems for 21st Century, MWGS08: MidWest Grid School 2008, September 17th, 2008.
43. Harnessing Grid Resources with Data Data-Centric Task Farms, Notre Dame University, CSE Department, August 20th, 2008.
44. Scientific Workflow Systems for 21st Century, New Bottle or New Wine?, IEEE Workshop on Scientific Workflows 2008, July 2008.
45. Accelerating Large-scale Data Exploration through Data Diffusion, ACM/IEEE International Workshop on Data-Aware Distributed Computing 2008, June 2008.
46. Accelerating Large-Scale Data Exploration through Data Diffusion, DSLW 2008, May 22nd, 2008.
47. Managing and Executing Loosely-Coupled Large-Scale Applications on Clusters, Grids, and Supercomputers, GlobusWorld 2008, May 15th, 2008.
48. Harnessing Grid Resources with Data Data-Centric Task Farms, NASA Ames Research Center, May 14th, 2008.
49. Harnessing Grid Resources with Data Data-Centric Task Farms, Hyde Park Global Investments LLC, April 18th, 2008.
50. Harnessing Grid Resources with Data Data-Centric Task Farms, University of Chicago, Department of Computer Science, Dissertation Proposal, December 12th, 2007.
51. Harnessing Grid Resources with Data Data-Centric Task Farms, University of Chicago, Department of Computer Science, Distributed Systems Lab Seminar, December 5th, 2007.
52. Falcon: a Fast and Light-weight task executiON framework for Grid Environments, IEEE/ACM SuperComputing 2007, November 15th, 2007.
53. Accelerating Large Scale Scientific Exploration with Falcon, IEEE/ACM SuperComputing 2007, Argonne National Laboratory Booth, November 14th, 2007.
54. A Data Diffusion Approach to Large Scale Scientific Exploration, University of Chicago, CS Department, Distributed Systems Lab Seminar, October 24th, 2007.
55. A Data Diffusion Approach to Large Scale Scientific Exploration, 2007 Microsoft eScience Workshop at RENCi, October 21st, 2007.
56. Falcon: a Fast and Light-weight task executiON framework, DSLW 2007, May 29th, 2007.
57. Falcon: a Fast and Light-weight task executiON framework for Grid Environments, CS Seminar, University of Chicago, CS Department, April 30th, 2007.
58. AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis, California Institute of Technology, AstroGrid 2007, February 12th, 2007.
59. AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis, IEEE/ACM SuperComputing 2006, November 2006.
60. Storage and Compute Resource Management via DYRE, 3DcacheGrid, and CompuStore, University of Chicago, Department of Computer Science, Distributed Systems Lab Seminar, November 2006.
61. Harnessing Grid Resources to Enable the Dynamic Analysis of Large Astronomy Datasets, DSLW 2006, June 2006
62. AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis, TeraGrid 2006, June 2006

63. Harnessing Grid Resources to Enable the Dynamic Analysis of Large Astronomy Datasets, University of Chicago, Department of Computer Science, Graduate Seminar, February 2006.
64. AstroPortal: A Science Portal to Grid Resources, University of Chicago, Department of Computer Science, Distributed Systems Lab Seminar, January 2006.
65. The Design, Performance, and Utility of DiPerF, an automated Distributed Performance testing Framework, University of Chicago, Department of Computer Science, Distributed Systems Lab Seminar, October 2005.
66. "A Performance Study of the Globus Toolkit® and Grid Services via DiPerF, an automated Distributed Performance testing Framework", University of Chicago: Master Thesis Defense, May 2005.
67. "Decreasing End-to-End Job Execution Times by Increasing Resource Utilization using Predictive Scheduling in the Grid", University of Chicago: Grid Computing Seminar, March 2005.
68. "DiPerF: an automated Distributed Performance testing Framework", University of Chicago: Graduate Seminar, November 2004.
69. "DiPerF: an automated Distributed Performance testing Framework", IEEE/ACM GRID2004, November 2004.
70. "A Study between Networks and General Purpose Systems for High Bandwidth Applications", University of Chicago: Computer Architecture, June 2004.
71. "DiPerF: an automated Distributed Performance testing Framework", University of Chicago: Grid Computing Seminar, March 2004.
72. "Searching Large Image Databases", University of Chicago: Artificial Intelligence (AI), March 2004.
73. "Mobile IPv6", Sun Microsystems Laboratories: September 2003.
74. "Mobile IPv6", Sun Microsystems Laboratories: August 2003.
75. "NEon", Sun Microsystems Laboratories: August 2003.
76. "Routing Algorithms in Wireless Sensor Networks", Purdue University: Seminar in Peer-to-Peer Networks, October 2002.
77. "Evaluating IPv6 using the Agere 2.5 PayloadPlus Network Processor", Purdue University: Seminar in Network Processors, October 2002.
78. "Routing Algorithms for Wireless Sensor Networks", Grace Hopper Celebration of Women in Computing, October 2002.
79. "An Empirical Analysis of Internet Protocol version 6 (IPv6)", Wayne State University: Master Thesis Defense, April 2002.
80. "Using Wireless Sensor Networks to Narrow the Gap between Low-Level Information and Context-Awareness", ISCA 17th International Conference on Computers and Their Applications, April 2002.
81. "MEMS Technology Overview and Limitations", Wayne State University: BioComputing Seminar, April 2002.
82. "Energy-Efficient Routing Algorithms in Wireless Sensor Networks", Wayne State University: Networking Wireless Sensor Lab, February 2002.
83. "IP Encapsulation Methods: A Transition Mechanism to Deploy IPv6 Networks", Wayne State University: Seminar in High Performance Networks, December 2001.
84. "Mobility Support in IPv6", Wayne State University: Seminar in High Performance Networks, October 2001.
85. "Wireless Sensor Networks and their Applications", Wayne State Univ.: Seminar in Wireless Sensor Networks, September 2001.
86. "Proximity Detection using Wireless RF Sensor", Accenture Technology Labs: August 2001.
87. "BSD vs. Streams Protocols", IEEE International Conference on Telecommunication (ICT2001): June 2001.

---

## Proposals/Funding

---

### Active

---

- NSF 2011 - 2015:** *"Avoiding Achilles' Heel in Exascale Computing with Distributed File Systems"*  
Institution: NSF, OCI  
Award: \$450,000  
Period: 01/2011 - 12/2015  
Collaborators: Arthur Barney Maccabe, Marc Snir, Rob Ross, Mike Wilde, Kamil Iskra, Jacob Furst, Mary Cummane
- DOE 2011 - 2012:** *"FusionFS: Distributed File Systems for Exascale Computing"*  
Institution: DOE ANL ALCF  
Award: 50,000 hours on the IBM BlueGene/P  
Period: 03/2011 - 02/2012  
Collaborators: Mike Wilde, Kamil Iskra, Pete Beckman, Rob Ross

### Pending

---

- NSF 2012 - 2013:** *"REU Supplement: Avoiding Achilles' Heel in Exascale Computing with Distributed File Systems"*  
Institution: NSF, OCI  
Award: \$16,625  
Period: 03/2012 - 02/2013
- DOE 2012 - 2017:** *"Paving the Road to Exascales with Many-Task Computing"*  
Institution: DOE, ASCR  
Award (pending): \$750,000  
Date Submitted: 09/01/2011  
Period: 06/2012 - 06/2017  
Collaborators: Arthur Barney Maccabe, Pete Beckman, Mike Wilde, Marc Snir, Ian Foster
- MSR 2012 - 2014:** *"Cost Effective Cloud Storage – The Tipping Point for Mainstream Adoption"*  
Institution: Microsoft Research Faculty Fellowship  
Award (pending): \$200,000  
Date Submitted: 10/31/2011  
Period: 07/2012 - 06/2014
- NSF 2012 - 2017:** *"MRI: Development and Deployment of Sustainable, Reliable, Secure, and Portable Intelligent Cyber Physical System to Promote Effective and Efficient Wastewater Treatment"*  
Institution: NSF, CNS  
Award (pending): \$1,527,003  
Date Submitted: 01/26/2012  
Period: 10/2012 - 09/2017  
Collaborators: Xiangyang Li, Jafar Saniie, Paul Anderson, Victor Perez-Luna, Erdal Oruklu, Ioan Raicu, Hamid Arastoopour, Jamshid Mohammadi, and Thomas M. Jacobius

### Completed

---

- Fermi 2011:** *"Networking and Distributed Systems in High-Energy Physics"*  
Institution: Fermi National Accelerator Laboratory  
Award: \$28,372  
Period: 05/2011 – 09/2011  
Collaborators: Gabriele Garzoglio
- DOE 2011 - 2012:** *"HPCcloud: Exploring HPC Fault Tolerance in the Cloud"*  
Institution: DOE ANL ALCF  
Award: 100,000 hours on the Magellan Cloud  
Period: 03/2011 - 02/2012

Collaborators: Xian-He Sun, Hui Jin, Tao Ke

**NSF 2011 - 2012:** *"Student Travel Support for ACM HPDC 2011"*

Institution: NSF, CCF

Award: \$10,000

Period: 05/2011 - 04/2012

Collaborators: Arthur Barney Maccabe

**NSF/CRA 2009 - 2010:** *"Resource Management in Large-Scale Distributed Systems"*

Institution: NSF and CRA, CIFellows

Award: \$140,000

Period: 08/2009 - 07/2010

Collaborators: Alok Choudhary

**NASA 2006 - 2009:** *"Harnessing Grid Resources to Enable the Dynamic Analysis of Large Astronomy Datasets"*

Institution: NASA, Ames Research Center, Graduate Student Researchers Program

Award: \$84,000

Period: 10/2006 - 5/2009

Collaborators: Ian Foster, Jerry C. Yan

---

# Projects

---

## MTC: Many-Task Computing

---

<b>MATRIX</b>	<b><i>MAny-Task computing execution fabRlc at eXascales</i></b> Period: 01/2011 – Present People: Ke Wang, Zhao Zhang Web Site: <a href="http://datasys.cs.iit.edu/projects/MATRIX/index.html">http://datasys.cs.iit.edu/projects/MATRIX/index.html</a>
<b>SimMatrix</b>	<b><i>SIMulator for MAny-Task computing execution fabRlc at eXascales</i></b> Period: 06/2011 – Present People: Ke Wang Web Site: <a href="http://datasys.cs.iit.edu/projects/SimMatrix/index.html">http://datasys.cs.iit.edu/projects/SimMatrix/index.html</a>
<b>Falkon</b>	<b><i>A Fast and Light-weight task executiON framework</i></b> Period: 12/2006 – Present People: Zhao Zhang, Mike Wilde Web Site: <a href="http://datasys.cs.iit.edu/projects/Falkon/index.html">http://datasys.cs.iit.edu/projects/Falkon/index.html</a>
<b>Swift</b>	<b><i>Parallel Programming System</i></b> Period: 12/2006 – Present People: Zhao Zhang, Mike Wilde Web Site: <a href="http://www.ci.uchicago.edu/swift/">http://www.ci.uchicago.edu/swift/</a>

## Storage

---

<b>FusionFS</b>	<b><i>Fusion distributed File System</i></b> Period: 08/2010 - Present People: Tonglin Li, Iman Sadooghi Web Site: <a href="http://datasys.cs.iit.edu/projects/FusionFS/index.html">http://datasys.cs.iit.edu/projects/FusionFS/index.html</a>
<b>ZHT</b>	<b><i>Zero-Hop Distributed Hash Table</i></b> Period: 03/2010 – Present People: Tonglin Li, Antonio Perez, Kevin Brandstatter, Zhao Zhang Web Site: <a href="http://datasys.cs.iit.edu/projects/ZHT/index.html">http://datasys.cs.iit.edu/projects/ZHT/index.html</a>
<b>PAFS</b>	<b><i>Provenance-Aware Distributed File System</i></b> Period: 08/2011 – Present People: Zhangjie Ma, Tanu Malik Web Site: <a href="http://datasys.cs.iit.edu/projects/PAFS_summary12.pdf">http://datasys.cs.iit.edu/projects/PAFS_summary12.pdf</a>
<b>SCH-FS</b>	<b><i>SSD-Cached Hybrid File System</i></b> Period: 01/2012 – Present People: Dongfang Zhao Web Site: <a href="http://datasys.cs.iit.edu/projects/SCH-FS_summary12.pdf">http://datasys.cs.iit.edu/projects/SCH-FS_summary12.pdf</a>
<b>NoVoHT</b>	<b><i>Non-Volatile Hash Table</i></b> Period: 10/2011 – Present People: Kevin Brandstatter Web Site: <a href="http://datasys.cs.iit.edu/">http://datasys.cs.iit.edu/</a>

## Cloud Computing

---

<b>CloudStorage</b>	<b><i>Understanding the Cost of Cloud Storage</i></b> Period: 08/2011 – Present People: Iman Sadooghi Web Site: <a href="http://datasys.cs.iit.edu/projects/CloudStorage_summary12.pdf">http://datasys.cs.iit.edu/projects/CloudStorage_summary12.pdf</a>
---------------------	--

---

## Many-Core Computing

---

<b>ManyCoreSim</b>	<b><i>Scheduling Direct Acyclic Graphs on Massively Parallel Processors</i></b> Period: 08/2011 – Present People: Ke Yue Web Site: <a href="http://datasys.cs.iit.edu/projects/ManyCoreSim_summary12.pdf">http://datasys.cs.iit.edu/projects/ManyCoreSim_summary12.pdf</a>
<b>vGPGPU</b>	<b><i>Virtualizing GPUs to Support MTC Applications</i></b> Period: 08/2011 – Present People: Scott Krieder Web Site: <a href="http://datasys.cs.iit.edu/">http://datasys.cs.iit.edu/</a>

---

## HPC: High-Performance Computing

---

<b>SimHEC</b>	<b><i>Simulator for High-End Computing Systems</i></b> Period: 08/2011 – Present People: Da Zhang Web Site: <a href="http://datasys.cs.iit.edu/projects/SimHEC_summary12.pdf">http://datasys.cs.iit.edu/projects/SimHEC_summary12.pdf</a>
---------------	--

---

## Mobile Computing

---

<b>CiteSearcher</b>	<b><i>a Google Scholar frontend for iOS and Android mobile devices</i></b> Period: 06/2011 – Present People: Kevin Brandstatter Web Site: <a href="http://datasys.cs.iit.edu/projects/CiteSearcher/index.html">http://datasys.cs.iit.edu/projects/CiteSearcher/index.html</a>
---------------------	--

---

## Completed Projects

---

<b>AstroPortal</b>	<b><i>A Science Gateway for Large-scale Astronomy Data Analysis</i></b> Period: 06/2005 – 03/2009 Web Site: <a href="http://datasys.cs.iit.edu/projects/Falkon/astro_portal.htm">http://datasys.cs.iit.edu/projects/Falkon/astro_portal.htm</a>
<b>DiPerF</b>	<b><i>An automated Distributed PERFORMANCE testing Framework</i></b> Period: 09/2003 – 06/2005 Web Site: <a href="http://datasys.cs.iit.edu/projects/diperf/">http://datasys.cs.iit.edu/projects/diperf/</a>
<b>ServMark</b>	<b><i>An Architecture for Testing Grid Services</i></b> Period: 01/2006 – 06/2008 Web Site: <a href="http://dev.globus.org/wiki/Incubator/ServMark">http://dev.globus.org/wiki/Incubator/ServMark</a>
<b>HOC-SA</b>	<b><i>Higher-Order Components-Service Architecture</i></b> Period: 06/2006 – 12/2007 Web Site: <a href="http://dev.globus.org/wiki/Incubator/HOC-SA">http://dev.globus.org/wiki/Incubator/HOC-SA</a>
<b>DI-GRUBER:</b>	<b><i>A Distributed Grid Resource Broker</i></b> Period: 09/2003 – 12/2006 Web Site: <a href="http://people.cs.uchicago.edu/~cldumitr/GRUBER/">http://people.cs.uchicago.edu/~cldumitr/GRUBER/</a>

---

# Teaching

---

**Assistant Professor** Illinois Institute of Technology, Department of Computer Science

Fall 2012

[CS 495: Introduction to Distributed Systems](#)

Spring 2012

[CS 553: Cloud Computing](#)

[CS 695: Doctoral Seminar](#)

Fall 2011

[CS 595: Data-Intensive Computing](#)

Spring 2011

[CS 550: Advanced Operating Systems](#)

[CS 695: Doctoral Seminar](#)

Fall 2010

[CS 595: Data-Intensive Computing](#)

**Lecturer** Northwestern Univ., Dept. of Electrical Eng. and Computer Science

01/2010 – 06/2010

[EECS 211: Fundamentals of Computer Programming](#)

[EECS 495: Hot Topics in Distributed Systems: Data-Intensive Computing](#)

**Teaching Assistant** University of Chicago, Department of Computer Science

09/2003 – 06/2005

Networking and Distributed Systems

Advanced Network Design

Introduction to Programming for the World Wide Web I

Honors Introduction to Computer Science 2

Introduction to Computer Systems

Fundamentals of Computer Programming I in Scheme

**Teaching Assistant** Purdue University, Department of Computer Science

08/2002 – 05/2003

Introduction to Networking

**Adjunct Assistant Professor** Univ. of Michigan, Department of Computer and Information Science

Summer 2002

Data Structures and Algorithm Analysis in C++

**Teaching Assistant** Wayne State University, Department of Computer Science

08/2000 – 12/2001

Problem Solving & Programming in C++

Data Structures & Abstraction in C++

---

## Students

---

### PhD Students (Advisor) at Illinois Institute of Technology

---

**Tonglin Li**, Expected graduation 2014

**Ke Wang**, Expected graduation 2015

**Iman Sadooghi**, Expected graduation 2016

**Scott Krieder**, Expected graduation 2016

**Dongfang Zhao**, Expected graduation 2016

### PhD Students (Dissertation Committee)

---

#### Hui Jin

Institution: Illinois Institute of Technology (CS)

Title: System Support for Resilience in Large-Scale Parallel Systems

Primary Advisor: Dr. Xian-He Sun

Graduation: Expected May 2012

#### Wei Tang

Institution: Illinois Institute of Technology (CS)

Title: Improving Job Scheduling and Resource Management on Production Supercomputers

Primary Advisor: Dr. Zhiling Lan

Graduation: Expected May 2012

#### Bin Zhang

Institution: DePaul University (SoC)

Title: Configuration Synthesis and Automated Alarm Correlation Structure of Security System

Primary Advisor: Dr. James Yu

Graduation: Expected 2012

### Master Students (Research Advisor) at Illinois Institute of Technology

---

**Zhangjie Ma**, Expected graduation 05/2012

**Da Zhang**, Expected graduation 05/2012

**Pedro Alvarez-Tabio**, Expected graduation 08/2012

**Corentin Debains**, Expected graduation 08/2012

### Undergraduate Students (Research Advisor) at Illinois Institute of Technology

---

**Antonio Perez De Tejada**, Expected graduation 05/2012

**Kevin Brandstatter**, Expected graduation 2015

### Alumni

---

**Ke Yue**, Illinois Institute of Technology, MS 12/2011

**Juan Carlos Hernández Munuera**, Illinois Institute of Technology, MS 08/2011

**Raman Verma**, Illinois Institute of Technology, MS 12/2010

**Kyle Chard**, Victoria University of Wellington, PhD 03/2011; "*DRIVE: A Distributed Economic Meta-Scheduler for the Federation of Grid and Cloud Systems*", Primary Advisor: Kris Bubendorfer

**Xi Duan**, Illinois Institute of Technology, MS 05/2011

**Yuchi Tsao**, Illinois Institute of Technology (ECE), PhD 12/2011; "*Architecture-Level Hardware Optimization for Digital Signal Processing*", Primary Advisor: Dr. Kyu-won "Ken" Choi

---

## Service (Post-Chronological Order)

---

### Panels

---

**NSF:** NSF Office of Cyberinfrastructure (OCI) Panel, 2011, 2012 – Ad-Hoc Reviewer

**NSF:** NSF Office of Cyberinfrastructure (OCI) Panel, 2010 – Panellist

### Journal Editor

---

**JoCCASA:** Journal of Cloud Computing Advances, Systems and Applications (Springer), 2011, 2012 – *Board Member*

**JGC:** Journal of Grid Computing, Special Issue on Data Intensive Computing in the Clouds, 2011 – *Editor*

**SPJ:** Scientific Programming Journal, Special Issue on Science-driven Cloud Computing, 2011 – *Editor*

**TPDS:** IEEE Transactions on Parallel and Distributed Systems, Special Issue on Many-Task Computing, 2011 -- *Editor*

### Journal Reviewer

---

**TPDS:** IEEE Transactions on Parallel and Distributed Systems, 2009, 2010, 2011

**FGCS:** The International Journal of Grid Computing and eScience, Future Generation Computer Systems, 2010, 2011

**JPDC:** Journal of Parallel and Distributed Computing, 2010, 2011

**TC:** IEEE Transactions on Computers, 2006, 2009, 2010

**TSC:** Transactions on Services Computing, 2010

**CPE:** Journal of Concurrency and Computation: Practice and Experience, 2006, 2009, 2010

**Cluster:** Cluster Computing, 2010

**IC:** IEEE Internet Computing, 2007, 2010

**IJCA:** ISCA International Journal of Computers and their Applications, 2009

**IJPDIM:** International Journal of Business Process Integration and Management, 2009

**CL:** IEEE Communication Letters, 2005

### Chair / Organizer

---

**CCGrid:** Cluster, Cloud, and Grid Computing Conference, 2014 – *Local Chair*

**TCSC:** IEEE Technical Committee on Scalable Computing (TCSC), 2012 -- *Cyber Infrastructure Chair*

**DataCloud:** IEEE International Workshop on Data-Intensive Computing in the Clouds (at IPDPS), 2011 -- *Chair*

**DataCloud-SC:** IEEE International Workshop on Data-Intensive Computing in the Clouds (at SC), 2011 -- *Chair*

**MTAGS:** ACM Work. on Many-Task Computing on Grids and Supercomputers (at SC), 2008, 2009, 2010, 2011 -- *Chair*

**ScienceCloud:** ACM Workshop on Scientific Cloud Computing (at HPDC), 2010, 2011 – *Chair*

**HPDC/SigMetrics:** HPDC/SigMetrics Student Research Posters Session 2011 – *Chair*

**CCA:** Cloud Computing and Its Applications Conference, 2011 – *Organizing Committee*

**HPDC:** ACM Symposium on High Performance Distributed Computing, 2010 -- *Student Activities Chair*

**DSLW:** Distributed Systems Laboratory Workshop, 2006, 2007, 2008, 2009 -- *Organizer*

**MegaJobs:** Bird of Feather Session – “How to Run One Million Jobs” (at SC), 2008 -- *Chair*

### Other Leading Roles (Steering Committee Member, Publicity Chair)

---

**CCGrid:** Cluster, Cloud, and Grid Computing Conference, 2012 – *Publicity Chair*

**eScience:** IEEE e-Science Conference, 2012 -- *Publicity Chair*

**Grid:** ACM/IEEE International Conference on Grid Computing, 2011, 2012 – *Publicity Chair*

**HPDC:** ACM Symp on High Performance Distributed Computing, 2011, 2012 -- *Publicity Chair*

**ICAC:** International Conference on Autonomic Computing, 2012 – *Publicity Chair*

**CloudFlow:** Work. on Workflow Models, Systems, Services and Applications in the Cloud, 2012 – *Steering Committee*

**ScienceCloud:** ACM Workshop on Scientific Cloud Computing (at HPDC), 2012 – *Steering Committee*

**SWF:** IEEE International Workshop on Scientific Workflows (at ICWS, SCC, CLOUD, SERVICES), 2010 -- *Publicity Chair*

---

### Conference Program Committee (selected)

---

**SC:** IEEE/ACM International Conference on High Performance Computing, Networking, Storage, and Analysis, 2012  
**HPDC:** ACM International Symposium on High Performance Distributed Computing, 2010, 2011, 2012  
**eScience:** IEEE e-Science Conference, 2010, 2011, 2012  
**CCGrid:** Cluster, Cloud, and Grid Computing Conference, 2012  
**Grid:** ACM/IEEE International Conference on Grid Computing, 2011, 2012  
**ICAC:** International Conference on Autonomic Computing, 2012  
**CLOUD:** IEEE International Conference on Cloud Computing, 2011, 2012  
**CLOSER:** International Conference on Cloud Computing and Services Science, 2012  
**FutureTech:** FTRA International Conference on Future Information Technology, 2012  
**ISPA:** IEEE International Symposium on Parallel and Distributed Processing with Applications, 2011  
**DHCS:** Digital Humanities and Computer Science Conference, 2011  
**ICWS:** IEEE International Conference on Web Services, 2011  
**TG:** TeraGrid Conference, 2009, 2010, 2011  
**CICIS:** Int. Conf. on Complex, Intelligent and Software Intensive Systems, 2010  
**CloudCom:** IEEE International Conference on Cloud Computing Technology and Science, 2010  
**CSE:** IEEE International Conference on Computational Science and Engineering, 2009, 2010  
**ICN:** IEEE International Conference on Networks, 2009, 2010  
**ICNS:** IEEE International Conference on Networking and Services, 2009, 2010  
**NBiS:** Int. Conference on Network-Based Information Systems, Grid, P2P Networks and Applications Track, 2010  
**RTSOAA:** IEEE International Workshop on Real-Time Service-Oriented Architecture and Applications, 2009  
**CTS:** IEEE/ACM International Symposium on Collaborative Technologies and Systems, 2009  
**ICIW:** IEEE International Conference on Internet and Web Applications and Services, 2008, 2009  
**ADVCOMP:** IEEE Conference on Advanced Engineering Computing and Applications in Sciences, 2008, 2009

### Workshop Program Committee (selected)

---

**HCW:** IEEE International Heterogeneity in Computing Workshop, 2010, 2011, 2012  
**DIDC:** ACM Workshop on Data-Intensive Distributed Computing, 2009, 2010, 2011  
**GCE:** IEEE Workshop on Grid Computing Portals and Science Gateways, 2007, 2008, 2009, 2010, 2011  
**SC-PhD:** Early Adopters Ph.D. Workshop: Building the Next Generation of Application Scientists, 2009, 2010, 2011  
**SWF:** IEEE International Workshop on Scientific Workflows, 2009, 2010, 2011  
**KDCloud:** IEEE/ICDM Int. Work. on Knowledge Discovery Using Cloud and Distributed Computing Platforms, 2010  
**CC:** IEEE International Workshop on Cloud Computing, 2009  
**HPDDM:** IEEE International Workshop on High Performance Distributed Data Management, 2009

---

## Collaborators (since 2008, not including students in DataSys lab)

---

**Atila S. Balkir** (UChicago), **Roger Barga** (MSR), **Pete Beckman** (UChicago/ANL), **Ivona Brandic** (VUT), **Amitabh Chaudhary** (ND), **Ben Clifford** (Consultant), **Mary Cummane** (Perspectives), **Ewa Deelman** (USC), **Catalin Dumitrescu** (FNAL), **Allan Espinosa** (UChicago), **Xubo Fei** (WSU), **Jacob Furst** (DePaul), **Gabriele Garzoglio** (FNAL), **Chris Gladwin** (Clevsafe), **Mihael Hategan** (UCI), **Kamil Iskra** (ANL), **Sarah Kenny** (UChicago), **Carl Kesselman** (USC), **Tevfik Kosar** (SUNY), **Philip Little** (ND), **Shiyong Lu** (WSU), **Arthur Barney Maccabe** (ORNL), **Tanu Malik** (UChicago), **John McGee** (RENCI), **Christopher Moretti** (ND), **Veronika Nefedova** (ANL), **Quan T. Pham** (UChicago), **Marlon Pierce** (IU), **Ruth Pordes** (FNAL), **Dick Repasky** (IU), **Matei Ripeanu** (UBC), **Rob Ross** (ANL), **Yogesh Simmhan** (USC), **Marc Snir** (ANL), **Rick Stevens** (UChicago/ANL), **Xian-He Sun** (IIT), **Alex Szalay** (JHU), **Douglas Thain** (ND), **Jing Tie** (UChicago), **Gabriela Turcu** (UChicago), **Michael Wilde** (UChicago/ANL), **Justin M. Wozniak** (ANL), **Xi Yang** (IIT), **Zhao Zhang** (UChicago), **Yong Zhao** (UESTC)

---

## Graduate Advisors and Postdoctoral Sponsors

---

**Alok Choudhary** (Northwestern University), **Ian T. Foster** (University of Chicago and Argonne National Laboratory), **Douglas Comer** (Purdue University), **Sherali Zeadally** (University of the District of Columbia)

---

## References

---

- Xian-He Sun** Professor and Chair, Dept. of Computer Science, **Illinois Institute of Technology**  
Contact: 1-312-567-5260, [sun@cs.iit.edu](mailto:sun@cs.iit.edu)
- Ian Foster** Arthur Holly Compton Distinguished Service Professor, Dept. of Computer Science, **University of Chicago**  
Director, **Computation Institute**, University of Chicago  
Senior Scientist, Mathematics and Computer Science Division, **Argonne National Laboratory**  
Co-Founder and Chief Scientist, **Univa UD**  
Contact: 1-630-252-4619, [foster@anl.gov](mailto:foster@anl.gov)
- Pete Beckman** Computer Scientist, Mathematics and Computer Science Division, **Argonne National Laboratory**  
Senior Fellow, **Computation Institute**, University of Chicago  
Division Director, Exascale Technology and Computing Institute, **Argonne National Laboratory**  
Contact: 1-630-252-5751, [beckman@mcs.anl.gov](mailto:beckman@mcs.anl.gov)
- Alex Szalay** Alumni Centennial Professor, Department of Physics and Astronomy, **The Johns Hopkins University**  
Contact: 1-410-516-7217, [szalay@jhu.edu](mailto:szalay@jhu.edu)
- Alok Choudhary** John G. Searle Professor, Dept. of Electrical Engineering and Computer Science, **Northwestern University**  
Contact: 1-847-467-4129, [choudhar@eecs.northwestern.edu](mailto:choudhar@eecs.northwestern.edu)
- Rick Stevens** Professor, Department of Computer Science, **University of Chicago**  
Senior Fellow, **Computation Institute**, University of Chicago  
Associate Laboratory Director, Computing and Life Sciences, **Argonne National Laboratory**  
Contact: 1-773-834-6816, [stevens@anl.gov](mailto:stevens@anl.gov)
- Carl Kesselman** Research Professor, Computer Science, **University of Southern California**  
Director, Fellow, Center for Grid Technologies, **Information Sciences Institute**  
Co-Founder and Chief Scientist, **Univa UD**  
Contact: 1-310-488-9338, [carl@isi.edu](mailto:carl@isi.edu)
- Mike Wilde** Fellow, **Computation Institute**, University of Chicago  
Software Architect, Mathematics and Computer Science Division, **Argonne National Laboratory**  
Contact: 1-708-203-9548, [wilde@mcs.anl.gov](mailto:wilde@mcs.anl.gov)
-