Curriculum Vitae – Ioan Raicu, Ph.D.

Associate Professor
Illinois Institute of Technology (IIT)
Department of Computer Science (CS)

Guest Research Faculty
Argonne National Laboratory (ANL)
Math and Computer Science Div. (MCS)

10 W. 31st Street
Stuart Building 226B
Chicago, IL 60616

Office: 1-312-567-5704
Email: iraicu@cs.iit.edu
Website: http://www.cs.iit.edu/~iraicu/
Laboratory: http://datasys.cs.iit.edu/
Google Scholar: http://scholar.google.com/citations?user=jE73HYAAAAAJ

Education

NSF/CRA CIFellow Postdoc in EECS	**Northwestern University**	**08/2009 – 07/2010**
Project: "*Resource Management in Large-Scale Distributed Systems*"		
Mentor: Alok Choudhary		
Ph.D. in Computer Science	**University of Chicago**	**09/2005 – 03/2009**
Dissertation: "*Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing*"		
Research Advisor: Ian Foster		
Master of Science in Computer Science	**University of Chicago**	**09/2003 – 06/2005**
Thesis: "*A Performance Study of the Globus Toolkit and Grid Services via DiPerF, an automated DIstributed PERformance testing Framework*"		
Research Advisor: Ian Foster		
Master of Science in Computer Science	**Wayne State University**	**09/2000 – 05/2002**
Thesis: "*An Empirical Analysis of Internet Protocol version 6 (IPv6)*"		
Research Advisor: Sherali Zeadally		
Bachelor of Science in Computer Science	**Wayne State University**	**09/1997 – 05/2000**

Work Experience

Associate Professor	**Illinois Institute of Technology** (CS), Chicago IL	08/2016 – Present
Guest Research Faculty	**Argonne National Laboratory** (MCS), Lemont IL	01/2011 – Present
Technical Advisory Board Member	**Ocient LLC** **(formerly Xeograph LLC)**, Chicago IL	04/2016 – Present
Technical Advisory Board Member	**FusionBlock**, Singapore	08/2018 – Present
Visiting Scholar (Sabbatical)	**Northwestern University** (EECS), Evanston IL	09/2016 – 06/2017
Assistant Professor	**Illinois Institute of Technology** (CS), Chicago IL	08/2010 – 07/2016
Computation Innovation Postdoc Fellow	**Northwestern University** (EECS), Evanston IL	08/2009 – 07/2010
Research Visitor	**NASA** (ARC/NAS), Moffet Field CA	03/2009 – 05/2009
Teaching/Research Assistant	**University of Chicago** (CS), Chicago IL	09/2003 – 03/2009
Researcher (Internship)	**Argonne National Laboratory** (MCS), Lemont IL	Summer 2005 & 2006
Researcher (Internship)	**Sun Microsystems** (Sun Labs & SNT), Menlo Park CA	Summer 2003
Teaching Assistant	**Purdue University** (CS), West Lafayette IN	08/2002 – 05/2003
Adjunct Assistant Professor	**University of Michigan** (CIS), Dearborn MI	06/2002 – 08/2002
Teaching/Research Assistant	**Wayne State University** (CS), Detroit MI	08/2000 – 08/2002
Researcher (Internship)	**Accenture Technology Labs**, Palo Alto CA	Summer 2001
System Analyst (Internship)	**Ford Motor Company**, Dearborn MI	Summer 1999
Owner	**High Teck Computers**, Westland MI	01/1997 – 03/2001

Awards

Teacher of the Year	Illinois Institute of Technology, Computer Science Department	2016
Young Achievers in Scalable Computing	IEEE Technical Committee on Scalable Computing (TCSC)	2014
Outstanding Service Award	IEEE/ACM CCGrid	2014
Junior Faculty Research Award	Illinois Institute of Technology	2013
NSF CAREER Award	National Science Foundation, Office of Cyber Infrastructure	2011
Computation Innovation Fellow	National Science Foundation & Computing Research Association	2009
GSRP Fellowship	NASA, Ames Research Center	2006-2008
Presidential Scholarship	Wayne State University	1997

Research

Research Interests

Distributed Systems	*Many-Task Computing*	*Data-Intensive Computing*
Grid Computing	*Cloud Computing*	*Supercomputing*
Many-Core Computing	*Big Data Computing*	*Parallel Programming Systems*

Publications

My publications have appeared in over 100 peer-reviewed articles, which received 10289 citations (according to Google Scholar, see Figure to thex right), yielding a H-index of 40. My complete list of my publications can be found at http://datasys.cs.iit.edu/publications/, and my Google Scholar Profile which includes all citation information can be found at http://scholar.google.com/citations?hl=en&user=jE73HYAAAAAJ.

Top 10 High-Impact Publications (ordered by citation counts):

4086 cites	1.	Ian Foster, Yong Zhao, Ioan Raicu, Shiyong Lu. “Cloud Computing and Grid Computing 360-Degree Compared”, IEEE Grid Computing Environments (GCE08) 2008
724 cites	2.	William Allcock, John Bresnahan, Rajkumar Kettimuthu, Michael Link, Catalin Dumitrescu, Ioan Raicu, Ian Foster, "The Globus Striped GridFTP Framework and Server," IEEE/ACM Supercomputing/SC, 2005
465 cites	3.	Yong Zhao, Mihael Hategan, Ben Clifford, Ian Foster, Gregor von Laszewski, Ioan Raicu, Tibi Stef-Praun, Mike Wilde. “Swift: Fast, Reliable, Loosely Coupled Parallel Computation”, IEEE Workshop on Scientific Workflows (SWF07) 2007
402 cites	4.	Ioan Raicu, Yong Zhao, Catalin Dumitrescu, Ian Foster, Mike Wilde. “Falkon: a Fast and Light-weight tasK executiON framework”, IEEE/ACM SuperComputing/SC, 2007
372 cites	5.	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
258 cites	6.	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
182 cites	7.	Tonglin Li, Xiaobing Zhou, Kevin Brandstatter, Dongfang Zhao, Ke Wang, Anupam Rajendran, Zhao Zhang, Ioan Raicu. “ZHT: A Light-weight Reliable Persistent Dynamic Scalable Zero-hop Distributed Hash Table”, IEEE International Parallel & Distributed Processing Symposium (IPDPS) 2013
155 cites	8.	Ioan Raicu, Zhao Zhang, Mike Wilde, Ian Foster, Pete Beckman, Kamil Iskra, Ben Clifford. “Toward Loosely-Coupled Programming on Petascale Systems”, IEEE/ACM SC 2008
122 cites	9.	Dongfang Zhao, Zhao Zhang, Xiaobing Zhou, Tonglin Li, Ke Wang, Dries Kimpe, Philip Carns, Robert Ross, and Ioan Raicu. "FusionFS: Towards Supporting Data-Intensive Scientific Applications on Extreme-Scale High-Performance Computing Systems", IEEE International Conference on Big Data 2014
119 cites	10.	Ke Wang, Xiaobing Zhou, Tonglin Li, Michael Lang, Ioan Raicu. "Optimizing Load Balancing and Data-Locality with Data-aware Scheduling", IEEE BigData 2014

Books (1)

1. 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

Theses/Proposals/Qualifiers (22)

1. Alex Orhean, Ioan Raicu (advisor). "XSearch: Distributed Search in Scientific Large-Scale Data Sets", Illinois Institute of Technology, Department of Computer Science, PhD Oral Qualifier, 2018
2. Poornima Nookala, Ioan Raicu (advisor). "XTASK - eXTreme fine-grAined concurrent taSK invocation runtime", Illinois Institute of Technology, Department of Computer Science, PhD Oral Qualifier, 2017
3. Jian Peng, Ioan Raicu (advisor). "Burst Buffer Simulation in Dragonfly Network", Illinois Institute of Technology, Department of Computer Science, PhD Oral Qualifier, 2017
4. Iman Sadooghi, Ioan Raicu (advisor). "Scalable Resource Management in Cloud Computing", Illinois Institute of Technology, Computer Science Department, Doctorate Dissertation, September 2016
5. Itua Ijagbone, Ioan Raicu. "Scalable Indexing and Searching on Distributed File Systems", Department of Computer Science, Illinois Institute of Technology, MS Thesis, 2016
6. Jason Arnold, Boris Glavic, Ioan Raicu. "HRDBMS: Combining the Best of Modern and Traditional Relational Databases", Illinois Institute of Technology, Department of Computer Science, PhD Oral Qualifier, 2015
7. Iman Sadooghi, Ioan Raicu (advisor). "Scalable Resource Management in Cloud Computing", Illinois Institute of Technology, Computer Science Department, PhD Proposal, 2015
8. Tonglin Li, Ioan Raicu (advisor). "Distributed NoSQL Storage for Extreme-SCale System Services in Supercomputers and Clouds", Illinois Institute of Technology, Computer Science Department, PhD Dissertation, 2015
9. Dongfang Zhao, Ioan Raicu. "Big Data System Infrastructure at Extreme Scales", Illinois Institute of Technology, Computer Science Department, Doctorate Dissertation, July 2015
10. Ke Wang, Ioan Raicu. "Scalable Resource Management System Software for Extreme-Scale", Illinois Institute of Technology, Computer Science Department, Doctorate Dissertation, June 2015
11. Tonglin Li, Ioan Raicu. "A Convergence of NoSQL Storage Systems from Clouds to Supercomputers", Illinois Institute of Technology, Computer Science Department, PhD Proposal, 2014
12. Ke Wang, Ioan Raicu. "Towards Next Generation Resource Management at Extreme-Scales", Illinois Institute of Technology, Computer Science Department, PhD Proposal, 2014
13. Dongfang Zhao, Ioan Raicu. "Towards Supporting Data-Intensive Scientific Applications on Extreme-Scale High-Performance Computing Systems," Illinois Institute of Technology, Computer Science Department, PhD Proposal, 2014
14. Iman Sadooghi, Ioan Raicu. "CloudKon: a Cloud enabled Distributed task executiON framework", Illinois Institute of Technology, Department of Computer Science, PhD Oral Qualifier, 2013
15. Scott Krieder, Ioan Raicu. "GEMTC: GPU Enabled Many-Task Computing", Illinois Institute of Technology, Department of Computer Science, PhD Oral Qualifier, 2013
16. Anupam Rajendran, Ioan Raicu. "MATRIX: Many-Task Computing Execution Fabric for Extreme Scales", Department of Computer Science, Illinois Institute of Technology, MS Thesis, 2013
17. Dongfang Zhao, Ioan Raicu. "HyCache: A Hybrid User-Level File System with SSD Caching", Illinois Institute of Technology, Department of Computer Science, PhD Oral Qualifier, 2012
18. Ke Wang, Ioan Raicu. "SimMatrix: SIMulator for MAny-Task computing execution fabRIc at eXascales", Illinois Institute of Technology, Department of Computer Science, PhD Oral Qualifier, 2012
19. Tonglin Li, Ioan Raicu. "ZHT: a Zero-hop DHT for High-End Computing Environment", Illinois Institute of Technology, Department of Computer Science, PhD Oral Qualifier, 2012
20. Ioan Raicu, Ian Foster. "Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing", Computer Science Dept., University of Chicago, Doctorate Dissertation, March 2009
21. Ioan Raicu, Ian Foster. "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

22. Ioan Raicu, Sherali Zeadally. "An Empirical Analysis of Internet Protocol version 6 (IPv6)", Wayne State University, Computer Science Department, MS Thesis, May 2002

Editorials (4)

1. Kate Keahey, Ioan Raicu, Kyle Chard, and Bogdan Nicolae. "Guest Editors Introduction: Special Issue on Scientific Cloud Computing", IEEE Transaction on Cloud Computing, 2015
2. Tevfik Kosar, Ioan Raicu. "Guest Editors' Introduction: Special Issue on Data-Intensive Computing in the Clouds", Springer Journal of Grid Computing, 2012
3. Ioan Raicu. Ian T. Foster, Yong Zhao. "Guest Editors' Introduction: Special Issue on Many-Task Computing", IEEE Transactions on Parallel and Distributed Systems, 2011
4. Ivona Brandic, Ioan Raicu. "Guest Editors' Introduction: Special Issue on Science-driven Cloud Computing", Scientific Programming Journal, 2011

Journal Articles (25)

1. Alexandru Iulian Orhean, Florin Pop, Ioan Raicu. "New Scheduling Approach using Reinforcement Learning for Heterogeneous Distributed Systems", Elsevier Journal of Parallel and Distributed Computing (JPDC), 2018
2. S Timm, G Cooper, S Fuess, G Garzoglio, B Holzman, R Kennedy, D Grassano, A Tiradani, R Krishnamurthy, S Vinayagam, I Raicu, H Wu, S Ren and S-Y Noh. "Virtual machine provisioning, code management, and data movement design for the Fermilab HEPCloud Facility", Journal of Physics: Conference Series 2017
3. Dongfang Zhao, Ke Wang, Kan Qiao, Tonglin Li, Iman Sadooghi, Ioan Raicu. Toward High-performance Key-value Stores through GPU Encoding and Locality-aware Encoding, Elsevier Journal of Parallel and Distributed Computing (JPDC), Special Issue on Scalable Computing Systems for Big Data Applications, 2016
4. Pedro Valero-Lara, Poornima Nookala, Fernando L. Pelayo, Johan Jansson, Serapheim Dimitropoulos, Ioan Raicu. "Many-Task Computing on Many-Core Architectures," Special Issue on High Performance Computing Solutions for Complex Problems, Scientific International Journal for Parallel and Distributed Computing, Scalable Computing: Practice and Experience (SCPE), 2016
5. Dongfang Zhao, Kan Qiao, Zhou Zhou, Tonglin Li, Xiaobing Zhou, Ioan Raicu. "Exploiting Multi-cores for Efficient Interchange of Large Messages in Distributed Systems", Concurrency and Computation: Practice and Experience (CCPE), 2015 (Impact Factor 1.0)
6. Dongfang Zhao, Ning Liu, Dries Kimpe, Robert Ross, Xian-He Sun, and Ioan Raicu. "Towards Exploring Data-Intensive Scientific Applications at Extreme Scales through Systems and Simulations", IEEE Transaction on Parallel and Distributed Systems (TPDS), no. 1, pp. 1 – 14, PrePrints, 2015; DOI:10.1109/TPDS.2015.2456896 (Impact Factor 2.173)
7. Ke Wang, Abhishek Kulkarni, Michael Lang, Dorian Arnold, and Ioan Raicu. "Exploring the Design Tradeoffs for Extreme-Scale High-Performance Computing System Software", IEEE Transaction on Parallel and Distributed Systems (TPDS), Issue 99, pp 1 – 14, PrePrints, 2015; DOI:10.1109/TPDS.2015.2430852 (Impact Factor 2.173)
8. Dongfang Zhao , Kan Qiao , Jian Yin , and Ioan Raicu. "Dynamic Virtual Chunks: On Supporting Efficient Accesses to Compressed Scientific Data", IEEE Transaction on Service Computing (TSC), SI on Big Data, pp 1 – 14, PrePrints, 2015; DOI:10.1109/TSC.2015.2456889 (Impact Factor 3.05)
9. Iman Sadooghi, Jesús Hernández Martin, Tonglin Li, Kevin Brandstatter, Ketan Maheshwari, Tiago Pais Pitta de Lacerda Ruivo, Gabriele Garzoglio, Steven Timm,Yong Zhao, Ioan Raicu. "Understanding the Performance and Potential of Cloud Computing for Scientific Applications", IEEE Transaction on Cloud Computing (TCC), Issue 99, pp 1 – 14, PrePrints, 2015; DOI:10.1109/TCC.2015.2404821
10. Ke Wang, Kan Qiao, Iman Sadooghi, Xiaobing Zhou, Tonglin Li, Michael Lang, Ioan Raicu. "Load-balanced and locality-aware scheduling for data-intensive workloads at extreme scales", Concurrency and Computation: Practice and Experience (CCPE), pp 1 – 29, 2015; DOI:10.1002/cpe.3617 (Impact Factor 1.0)
11. Tonglin Li, Xiaobing Zhou, Ke Wang, Dongfang Zhao, Iman Sadooghi, Zhao Zhang, Ioan Raicu. "A Convergence of Key-Value Storage Systems from Clouds to Supercomputers", Concurency and Computation: Practice and Experience (CCPE), 2015; DOI:10.1002/cpe.3614 (Impact Factor 1.0)

12. Dongfang Zhao, Kan Qiao, Ioan Raicu. "Towards Cost-Effective and High-Performance Caching Middleware for Distributed Systems", International Journal of Big Data Intelligence (IJBDI), SI on High-Performance Data Intensive Computing, 2015
13. Steven Timm, Gabriele Garzoglio, Parag Mhashilkar, Joseph Boyd, G. Bernabeu, Neha Sharma, N. Peregonow, H.W. Kim, S. Noh, Sandeep Palur, Ioan Raicu. "Cloud Services for the Fermilab scientific stakeholders", Journal of Physics: Conference Series 2015
14. Yong Zhao, Youfu Li, Ioan Raicu, Shiyong Lu, Wenhong Tian, Heng Liu. "Enabling Scalable Scientific Workflow Management in the Cloud", Future Generation Computer Systems, Volume 46, May 2015, Pages 3–16, DOI:10.1016/j.future.2014.10.023 (Impact Factor 2.786)
15. Yong Zhao, Youfu Li, Ioan Raicu, Shiyong Lu, Cui Lin, Yanzhe Zhang, Wenhong Tian, Ruini Xue. "A Service Framework for Scientific Workflow Management in the Cloud", IEEE Trans. on Services Computing (TSC), Issue 99, pp 1 – 14, PrePrints, 2014; DOI:10.1109/TSC.2014.2341235 (Impact Factor 3.05)
16. Michael Wilde, Ian Foster, Pete Beckman, Ioan Raicu. “Scalable Parallel Scripting for Scientific Computing”, SciDAC Review Spring 2010, pp 38 – 53
17. 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, Volume 13 Issue 3, September 2010, pp 291 - 314, DOI: 10.1007/s10586-010-0132-9 (Impact Factor 0.679)
18. 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, SI on Extreme Scale Computing, pp 50 – 60, 2009; DOI 10.1007/s10723-006-9060-6 (Impact Factor 2.205)
19. 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”, Journal of Physics, Volume 180, Number 1, pp 1 – 5, 2009; DOI:10.1088/1742-6596/180/1/012046
20. Catalin Dumitrescu, Ioan Raicu, Ian Foster. “The Design, Usage, and Performance of GRUBER: A Grid uSLA-based Brokering Infrastructure”, International Journal of Grid Computing (JGC), March 2007, Volume 5, Issue 1, pp 99-126; DOI 10.1007/s10723-006-9060-6 (Impact Factor 1.507)
21. Catalin Dumitrescu, Ioan Raicu, Ian Foster. “Usage SLA-based Scheduling in Grids”, Journal on Concurrency and Computation: Practice and Experience (CCPE), 19:945–963, 2007; DOI: 10.1002/cpe.1091 (Impact Factor 1.154)
22. 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 (JGC), September 2006, Volume 4, Issue 3, pp 287-309; DOI 10.1007/s10723-006-9037-5 (Impact Factor 1.507)
23. 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: pp 163 – 185, 2005; DOI: 10.1080/15501320590966431 (Impact Factor 0.665)
24. 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, pp 238 – 242; DOI:10.1049/ip-com:20040283
25. Sherali Zeadally, Ioan Raicu. “Evaluating IPV6 on Windows and Solaris”, IEEE Internet Computing (IC), Volume 7, Issue 3, May June 2003, pp 51 – 57 (Impact Factor 2.579)

Book Chapters (5)

1. Dongfang Zhao, Akash Mahakode, Sandip Lakshminarasaiah, Ioan Raicu. “High-performance Storage Support for Scientific Big Data Applications on the Cloud”, Springer’s Resource Management for Big-Data Platforms: Algorithms, Modelling, and High-Performance Computing Techniques, 2016
2. Yong Zhao, Youfu Li, Ioan Raicu, Cui Lin, Wenhong Tian, Ruini Xue. “Migrating Scientific Workflow Management Systems from the Grid to the Cloud”, Springer’s Cloud Computing for Data Intensive Applications, 2014

3. Ioan Raicu, Ian Foster, Yong Zhao, Alex Szalay, Philip Little, Christopher M. Moretti, Amitabh Chaudhary, Douglas Thain. “Towards Data Intensive Many-Task Computing”, Data Intensive Distributed Computing: Challenges and Solutions for Large-Scale Information Management, 2012
4. Yong Zhao, Ioan Raicu, Ian Foster, Mihael Hategan, Veronika Nefedova, Mike Wilde. “Realizing Fast, Scalable and Reliable Scientific Computations in Grid Environments”, Grid Computing Research Progress, 2008
5. Catalin Dumitrescu, Jan Dünnweber, Philipp Lüdeking, Sergei Gorlatch, Ioan Raicu, Ian Foster. “Simplifying Grid Application Programming Using Web-Enabled Code Transfer Tools”. Toward Next Generation Grids, 2007

Conference Papers (39)

1. Jason Arnold, Boris Glavic, Ioan Raicu. "A High-Performance Distributed Relational Database System for Scalable OLAP Processing", IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2019 (24% acceptance rate)
2. Alexandru Iulian Orhean, Itua Ijagbone, Dongfang Zhao, Kyle Chard, Ioan Raicu. “Toward Scalable Indexing and Search on Distributed and Unstructured Data”, IEEE Big Data Congress 2017 (23% acceptance rate)
3. Iman Sadooghi, Geet Kumar, Ke Wang, Dongfang Zhao, Tonglin Li, Ioan Raicu. “Albatross: an Efficient Cloud-enabled Task Scheduling and Execution Framework using Distributed Message Queues”, eScience 2016
4. Tonglin Li, Ke Wang, Shiva Srivastava, Dongfang Zhao, Kan Qiao, Iman Sadooghi, Xiaobing Zhou, Ioan Raicu. "A Flexible QoS Fortified Distributed Key-Value Storage System for the Cloud", IEEE Big Data 2015 (20% acceptance rate)
5. Iman Sadooghi, Ke Wang, Dharmit Patel, Dongfang Zhao, Tonglin Li, Shiva Srivastava, Ioan Raicu. “FaBRiQ: Leveraging Distributed Hash Tables towards Distributed Publish-Subscribe Message Queues”, IEEE/ACM BDC 2015 (16% acceptance rate)
6. Ke Wang, Ning Liu, Iman Sadooghi, Xi Yang, Xiaobing Zhou, Michael Lang, Xian-He Sun, Ioan Raicu. "Overcoming Hadoop Scaling Limitations through Distributed Task Execution", IEEE Cluster 2015 (24% acceptance rate)
7. Ke Wang, Xiaobing Zhou, Kan Qiao, Michael Lang, Benjamin McClelland, Ioan Raicu. "Towards Scalable Distributed Workload Manager with Monitoring-Based Weakly Consistent Resource Stealing", ACM HPDC 2015 (27% acceptance rate)
8. Dongfang Zhao, Zhao Zhang, Chen Shou, Ioan Raicu. “Towards Supporting Data-Intensive Scientific Applications on Extreme-Scale High Performance Computing Systems”, IEEE BigData 2014 (18% acceptance rate)
9. Ke Wang, Xiaobing Zhou, Tonglin Li, Michael Lang, Ioan Raicu. "Optimizing Load Balancing and Data-Locality with Data-aware Scheduling", IEEE BigData 2014 (18% acceptance rate)
10. Dongfang Zhao, Jian Yin, Kan Qiao, Ioan Raicu. "Virtual Chunks: On Supporting Random Accesses to Scientific Data in Compressible Storage Systems", IEEE BigData 2014 (18% acceptance rate)
11. Yong Zhao, Youfu Li, Ioan Raicu, Shiyong Lu, Xuan Zhang. "Architecting Cloud Workflow: Theory and Practice", IEEE CIT 2014 (26% acceptance rate)
12. Scott J. Krieder, Justin M. Wozniak, Timothy Armstrong, Michael Wilde, Daniel S. Katz, Benjamin Grimmer, Ian T. Foster, Ioan Raicu. “Design and Evaluation of the GeMTC Framework for GPU-enabled Many-Task Computing”, ACM HPDC 2014 (16% acceptance rate)
13. Ke Wang, Xiaobing Zhou, Hao Chen, Michael Lang, Ioan Raicu. “Next Generation Job Management Systems for Extreme Scales”, ACM HPDC 2014 (24% acceptance rate)
14. Tonglin Li, Ioan Raicu, Lavanya Ramakrishnan. "Scalable State Management for Scientific Applications in the Cloud", IEEE BigData 2014 (19% acceptance rate)
15. Iman Sadooghi, Sandeep Palur, Ajay Anthony, Isha Kapur, Karthik Belagodu, Pankaj Purandare, Kiran Ramamurty, Ke Wang, Ioan Raicu. "Achieving Efficient Distributed Scheduling with Message Queues in the Cloud for Many-Task Computing and High-Performance Computing", IEEE/ACM CCGrid 2014 (19% acceptance rate)
16. Dongfang Zhao, Kan Qiao, Ioan Raicu. “HyCache+: Towards Scalable High-Performance Caching Middleware for Parallel File Systems”, IEEE/ACM CCGrid 2014 (19% acceptance rate)
17. Tiago Pais Pitta de Lacerda Ruivo, Gerard Bernabeu Altayo, Gabriele Garzoglio, Steven Timm, Hyun Woo Kim, Seo-Young Noh, Ioan Raicu. "Exploring Infiniband Hardware Virtualization in OpenNebula towards Efficient High-Performance Computing", SCALE Challenge, IEEE/ACM CCGrid 2014

18. Dongfang Zhao, Chen Shou, Tanu Malik, Ioan Raicu. “Distributed Data Provenance for Large-Scale Data-Intensive Computing”, IEEE Cluster 2013 (31% acceptance rate)
19. Dongfang Zhao, Corentin Debains, Pedro Alvarez-Tabio, Kent Burlingame, Ioan Raicu. “Towards High-Performance and Cost-Effective Distributed Storage Systems with Information Dispersal Algorithms”, IEEE Cluster 2013 (31% acceptance rate)
20. Ke Wang, Abhishek Kulkarni, Dorian Arnold, Michael Lang, Ioan Raicu. “Using Simulation to Explore Distributed Key-Value Stores for Extreme-Scale System Services”, IEEE/ACM Supercomputing/SC 2013 (20% acceptance rate)
21. Tonglin Li, Xiaobing Zhou, Kevin Brandstatter, Dongfang Zhao, Ke Wang, Anupam Rajendran, Zhao Zhang, Ioan Raicu. “ZHT: A Light-weight Reliable Persistent Dynamic Scalable Zero-hop Distributed Hash Table”, IEEE IPDPS 2013 (21% acceptance rate)
22. Anupam Rajendran, Parag Mhashilkar, Hyunwoo Kim, Dave Dykstra, Gabriele Garzoglio, Ioan Raicu. “Optimizing Large Data Transfers over 100Gbps Wide Area Networks”, IEEE/ACM CCGrid 2013 (22% acceptance rate)
23. Ke Wang, Kevin Brandstatter, Ioan Raicu. “SimMatrix: Simulator for MAny-Task computing execution fabRIc at eXascales”, ACM HPC 2013
24. Dongfang Zhao, Da Zhang, Ke Wang, Ioan Raicu. “Exploring Reliability of Exascale Systems through Simulations”, ACM HPC 2013
25. Hui Jin, Xi Yang, Xian-He Sun, Ioan Raicu. “ADAPT: Availability-aware MapReduce Data Placement in Non-Dedicated Distributed Computing”, IEEE ICDCS 2012 (13% acceptance rate)
26. Yong Zhao, Ioan Raicu, Shiyong Lu, Xubo Fei. "Opportunities and Challenges in Running Scientific Workflows on the Cloud", IEEE CyberC 2011 (21% acceptance rate)
27. 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 HPDC 2009 (29% acceptance rate)
28. Ioan Raicu, Zhao Zhang, Mike Wilde, Ian Foster, Pete Beckman, Kamil Iskra, Ben Clifford. “Toward Loosely-Coupled Programming on Petascale Systems”, IEEE/ACM SuperComputing/SC 2008 (21% acceptance rate)
29. Ioan Raicu, Yong Zhao, Catalin Dumitrescu, Ian Foster, Mike Wilde. “Falkon: a Fast and Light-weight tasK executiON framework”, IEEE/ACM SuperComputing/SC 2007 (20% acceptance rate)
30. Ioan Raicu, Ian Foster, Alex Szalay, Gabriela Turcu. “AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis”, TeraGrid 2006
31. 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 SuperComputing/SC 2005 (22% acceptance rate)
32. Catalin Dumitrescu, Ioan Raicu, Ian Foster. "DI-GRUBER: A Distributed Approach for Grid Resource Brokering", IEEE/ACM SuperComputing/SC 2005 (22% acceptance rate)
33. Catalin Dumitrescu, Ioan Raicu, Ian Foster. “Experiences in Running Workloads over Grid3”, GCC 2005
34. Catalin Dumitrescu, Ioan Raicu, Matei Ripeanu, Ian Foster. "DiPerF: an automated DIstributed PERformance testing Framework", IEEE/ACM GRID 2004 (22% acceptance rate)
35. Ioan Raicu, Loren Schwiebert, Scott Fowler, Sandeep K.S. Gupta. “e3D: An Energy-Efficient Routing Algorithm for Wireless Sensor Networks”, IEEE ISSNIP 2004
36. Ioan Raicu, Sherali Zeadally. “Impact of IPv6 on End-User Applications”, IEEE ICT 2003
37. Ioan Raicu, Sherali Zeadally. “Evaluating IPv4 to IPv6 Transition Mechanisms” , IEEE ICT 2003
38. Ioan Raicu. “Efficient Even Distribution of Power Consumption in Wireless Sensor Networks”, ISCA CATA 2003
39. Ioan Raicu, Owen Richter, Loren Schwiebert, Sherali Zeadally. “Using Wireless Sensor Networks to Narrow the Gap between Low-Level Information and Context-Awareness”, ISCA CATA 2002

Workshop Papers (17)

1. Dongfang Zhao, Xu Yang, Iman Sadooghi, Gabriele Garzoglio, Steven Timm, Ioan Raicu. "High-Performance Storage Support for Scientific Applications on the Cloud", Invited Paper, ACM ScienceCloud 2015
2. Tonglin Li, Kate Keahey, Ke Wang, Dongfang Zhao, Ioan Raicu. "A Dynamically Scalable Cloud Data Infrastructure for Sensor Networks", Invited Paper, ACM ScienceCloud 2015
3. Yong Zhao, Youfu Li, Shiyong Lu, Ioan Raicu, Cui Lin. "Devising a Cloud Scientific Workflow Platform for Big Data", IEEE International Symposium on Scientific Workflows and Big Data Science (SWF) 2014

4. Dharmit Patel, Faraj Khasib, Iman Sadooghi, Ioan Raicu. "Towards In-Order and Exactly-Once Delivery using Hierarchical Distributed Message Queues", 1st International Workshop on Scalable Computing For Real-Time Big Data Applications (SCRAMBL'14) at IEEE/ACM CCGrid 2014
5. Chen Shou, Dongfang Zhao, Tanu Malik, Ioan Raicu. “Towards a Provenance-Aware a Distributed File System”, USENIX TaPP 2013
6. Ke Wang, Zhangjie Ma, Ioan Raicu. “Modelling Many-Task Computing Workloads on a Petaflop IBM BlueGene/P Supercomputer”, IEEE CloudFlow 2013
7. Dongfang Zhao, Ioan Raicu. “HyCache: A User-Level Caching Middleware for Distributed File Systems”, IEEE HPDIC 2013
8. 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
9. 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
10. Ian Foster, Yong Zhao, Ioan Raicu, Shiyong Lu. “Cloud Computing and Grid Computing 360-Degree Compared”, IEEE Grid Computing Environments (GCE08) 2008
11. 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
12. 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
13. 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
14. 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
15. Ioan Raicu, Ian Foster, Alex Szalay, Gabriela Turcu. “AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis”, TeraGrid Conference 2006, June 2006
16. Alex Szalay, Julian Bunn, Jim Gray, Ian Foster, Ioan Raicu. “The Importance of Data Locality in Distributed Computing Applications”, NSF Workflow Workshop 2006
17. 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", Argonne National Laboratory, MCS Preprint #ANL/MCS-P1315-0106, January 2006

Posters / Extended Abstracts [Peer Reviewed] (38)

1. Caleb Lehman, Poornima Nookala, Ioan Raicu. “Scalable Load-Balancing Concurrent Queues on Many-Core Architectures”, IEEE/ACM SuperComputing/SC 2019
2. Alex Ballmer, Brendan Batliner, Anna Benson, Blake Ehrenbeck, Zhen Huang, Parker Joncus, Travis Koehring, Alexandru Orhean, William Scullin, Ben Allen, Ioan Raicu. “Reaching for 100TFlops at 3KW Power with Intel Scalable Processors and NVIDIA V100 NVLINK GPUs”, Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2018
3. Ioan Raicu, William Scullin, Ben Allen, Kyle Hale, Kyle Chard, Alexandru Iulian Orhean. Breaking 100TFlops at 3KW Power with IBM Power9 and NVIDIA V100 GPUs over NVLink and 200GbE Mesh Interconnect”, Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2018
4. Vineeth Remanan Pillai, Daniela Stan Raicu, Ioan Raicu. “Accelerating Worm Segmentation through Inter-node Parallelism”, Engineering in Medicine and Biology Conference (EMBC) 2018
5. Alexander Ballmer, David Ghiurco, Iva Veseli, Hasan Rizvi, Ryan Mitchell, Ryan Prendergast, William Scullin, Ben Allen, Alexandru Iulian Orhean, Ioan Raicu. “Scalable High-performance Computing in a 3000-Watt Power Envelope”, Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2017
6. Ioan Raicu, William Scullin, Ben Allen, Kyle Hale, Kyle Chard, Simone Campanoni. “Maximizing Computation per Power Ratios in High-Performance Computing: from Aggressive Power Management to Approximate Computing”, Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2017

7. Anna Blue Keleher, Kyle Chard, Ian Foster, Alex Orhean, Ioan Raicu. “Finding a Needle in a Field of Haystacks: Metadata Search for Distributed Research Repositories”, IEEE/ACM SuperComputing/SC 2017
8. Prajakt Shastry, Daniel Parker, Sanjiv Kapoor, Ioan Raicu. “Exploring Randomized Multipath Routing on Multi-Dimensional Torus Networks”, IEEE/ACM SuperComputing/SC 2016
9. Jian Peng, Sughosh Divanji, Ioan Raicu, Mike Lang. “Simulating the Burst Buffer Storage Architecture on an IBM BlueGene/Q Supercomputer”, IEEE/ACM SuperComputing/SC 2016
10. Jonathan Wu, Suraj Chafle, Ioan Raicu, Kyle Chard. “Optimizing Search in Un-Sharded Large-Scale Distributed Systems”, IEEE/ACM SuperComputing/SC 2016
11. Alex Ballmer, Ioan Raicu. “FemtoGraph: A Pregel Based Shared-memory Graph Processing Library”, IEEE/ACM SuperComputing/SC 2016
12. Ben Walters, Alexander Ballmer, Adnan Haider, Andrei Dumitru, Keshav Kapoor, Calin Segarceau, William Scullin, Ben Allen, Ioan Raicu. “Illinois Institute of Technology – SC15 Student Cluster Competition”, Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2015
13. Ben Walters, Alex Ballmer, Andrei Dumitru, Adnan Haider, Serapheim Dimitropoulos, Ariel Young, William Scullin, Ben Allen, Ioan Raicu. "15 TFlops Haswell vs. 60 TFlops Knight Landing for HPC Scientific Computing Applications", Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2015
14. Tonglin Li, Ioan Raicu. "Distributed NoSQL Storage for Extreme-Scale System Services", Doctoral Showcase, IEEE/ACM Supercomputing/SC 2015
15. Xiaobing Zhou, Tonglin Li, Ke Wang, Dongfang Zhao, Iman Sadooghi, Ioan Raicu. "MHT: A Light-weight Scalable Zero-hop MPI Enabling Distributed Hash Table", IEEE Big Data 2015
16. Jason Arnold, Boris Glavic, Ioan Raicu. “HRDBMS: A NewSQL Database for Analytics”, IEEE Cluster 2015
17. Chaoqi Ma, Jiabao Li, TongLin Li, Ioan Raicu. “GRAPH/Z: A Key-Value Store Based Scalable Graph Processing System”, IEEE Cluster 2015
18. Poornima Nookala, Serapheim Dimitropoulos, Karl Stough, Ioan Raicu. “Evaluating the Support of MTC Applications on Intel Xeon Phi Many-Core Accelerators”, IEEE Cluster 2015
19. Thomas Dubucq, Tony Forlini, Virgile Landeiro Dos Reis, Isabelle Santos, Ke Wang, Ioan Raicu. “Benchmarking State-of-the-art Many-Task Computing Runtime Systems”, ACM HPDC 2015
20. Dongfang Zhao and Ioan Raicu. "Storage Support for Data-Intensive Applications on Extreme-Scale HPC Systems", Doctoral Showcase, IEEE/ACM Supercomputing/SC 2014
21. Tonglin Li, Kate Keahey, Ioan Raicu. “A Cloud-based Interactive Data Infrastructure for Sensor Networks”, IEEE/ACM Supercomputing/SC 2014
22. Kevin Brandstatter, Jason DiBabbo, Daniel Gordon, Ben Walters, Alex Ballmer, Lauren Ribordy, Ioan Raicu. "Delivering 3.5 Double Precision GFlops/Watt and 200Gb/sec Bi-Section Bandwidth with Intel Xeon Phi-based Cisco Servers", Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2014
23. Dongfang Zhao, Jian Yin, Ioan Raicu. “Improving the I/O Throughput for Data-Intensive Scientific Applications with Efficient Compression Mechanisms”, IEEE/ACM Supercomputing 2013
24. Benjamin Grimmer, Scott Krieder, Ioan Raicu. "Enabling Dynamic Memory Management Support for MTC on NVIDIA GPUs", EuroSys 2013
25. Ke Wang, Ioan Raicu. “Paving the Road to Exascale with Many-Task Computing”, Doctoral Showcase, IEEE/ACM Supercomputing/SC 2012
26. Dongfang Zhao, Ioan Raicu. “Distributed File Systems for Exascale Computing”, Doctoral Showcase, IEEE/ACM Supercomputing/SC 2012
27. Iman Sadooghi, Ioan Raicu. “Towards Scalable and Efficient Scientific Cloud Computing”, Doctoral Showcase, IEEE/ACM Supercomputing/SC 2012
28. Scott Krieder, Ioan Raicu. “Towards the Support for Many-Task Computing on Many-Core Computing Platforms”, Doctoral Showcase, IEEE/ACM Supercomputing/SC 2012
29. Scott Krieder, Ben Grimmer, Ioan Raicu. “Early Experiences in running Many-Task Computing workloads on GPGPUs”, XSEDE 2012

30. 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
31. Hui Jin, Xi Yang, Xian-He Sun, Ioan Raicu. "An Empirical Evaluation of MapReduce under Interruptions", Cloud Computing and its Applications (CCA) 2011
32. 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
33. 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
34. 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
35. 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
36. Ioan Raicu, Catalin Dumitrescu, Ian Foster. "Dynamic Resource Provisioning in Grid Environments", TeraGrid Conference (TG07) 2007
37. 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
38. Ioan Raicu. "Routing Algorithms for Wireless Sensor Networks", Poster Presentation, Grace Hopper Celebration of Women in Computing (GHC02), 2002

Presentations (116)

1. Mystic: Programmable Systems Research Testbed to Explore a Stack-WIde Adaptive System fabriC, Trends in HPDC Workshop 2019, March 14th, 2019
2. Mystic: Programmable Systems Research Testbed to Explore a Stack-WIde Adaptive System fabriC, Illinois Institute of Technology, November 29th, 2018
3. Fusion: Teaching Cluster at IIT in the CS Department, Illinois Institute of Technology, September 27th, 2018
4. Accelerating Worm Segmentation through Inter-node Parallelism, Engineering in Medicine and Biology Conference (EMBC) 2018, July 18th, 2018
5. Panel: Communicating Student Research, NSF REU PI Meeting, April 23rd, 2018
6. Big Data System Infrastructure at Extreme Scales, Department of Electrical Engineering and Computer Science, Northwestern University, May 1st, 2017
7. The Big Bang of Computing, TEDxIIT, April 10th, 2016
8. The Difference Between Failure and Success is +1, Keynote Presentation at the 1st Workshop on E-science ReseaRch leading tO negative Results (ERROR) 2015, September 3rd, 2015
9. Avoiding Achilles' Heel in Exascale Computing with Distributed File Systems, NSF CyberBridges 2015, September 1st, 2015
10. Exploring Extreme-Scale Computing through Simulations, Argonne National Laboratory, July 14th, 2015
11. Big Data Computing - The 4th Paradigm, Illinois Tech Preview Weekend, Illinois Institute of Technology, April 10th, 2015
12. Introduction to Computer Science, the What, How, and Why of CS, CAMRAS Interview Weekend, Illinois Institute of Technology, February 6th, 2015
13. Distributed Storage Systems for Extreme-Scale Data-Intensive Computing, Fermi National Accelerator Laboratory, January 8th, 2015
14. Distributed Storage Systems for Extreme-Scale Data-Intensive Computing, GCASR 2014, May 19th, 2014
15. Distributed Storage Systems for Extreme-Scale Data-Intensive Computing, HPDC PC Workshop, March 20th, 2014
16. Resource Management for Extreme-Scale Data-Intensive Computing, Illinois Institute of Technology, September 13th, 2013
17. Distributed Storage Systems for Extreme-Scale Data-Intensive Computing, CHANGES 2013, September 10th, 2013
18. Resource Management in Extreme Scales Distributed Systems, CS / Illinois Institute of Technology, August 21st, 2013
19. Resource Management in Extreme Scales Distributed Systems, University of New Mexico, August 12th, 2013

20. Supporting Data-Intensive Distributed Systems, DePaul University, July 18th, 2013
21. Distributed Storage Systems for Extreme-Scale Data-Intensive Computing, NSF CyberBridges 2013, July 15th, 2013
22. Distributed Storage Systems for Extreme-Scale Data-Intensive Computing, Google, July 12th, 2013
23. Optimizing Large Data Transfers over 100Gbps Wide Area Networks, IEEE/ACM CCGrid 2013, May 14th,2013
24. Supporting Data-Intensive Computing at Extreme Scales, Big Data Research Forum, Illinois Institute of Technology, March 12th, 2013
25. Introduction to Computer Science, the What, How, and Why of CS, CAMRAS Interview Weekend, Illinois Institute of Technology, February 14th, 2013
26. GeMTC: GPU enabled Many-Task Computing, Argonne National Laboratory, October 4th, 2012
27. Data-Intensive Computing at the Intersection of Cloud Computing and Supercomputing, Illinois Institute of Technology, October 4th, 2012
28. Building Blocks for Scalable Distributed Storage Systems, 1st Greater Chicago Area System Research Workshop, May 22nd, 2012
29. Challenges and Opportunities in Large-Scale Storage Systems, Fermi National Accelerator Laboratory, March 29th, 2012
30. Introduction to Computer Science, the What, How, and Why of CS, CAMRAS Interview Weekend, Illinois Institute of Technology, February 10th, 2012
31. Undergraduate Research in the DataSys Laboratory, CAMRAS Students Luncheon, Illinois Institute of Technology, January 18th, 2012
32. Distributed and Cloud Computing Specialization at IIT, Introduction to the MCS Specializations Seminar, Illinois Institute of Technology, September 20th, 2011
33. An Overview of Distributed Systems, Illinois Institute of Technology, September 13th, 2011
34. Challenges and Opportunities in Large-Scale Storage Systems, Illinois Institute of Technology, September 9th, 2011
35. Making a Case for Distributed File Systems at Exascales, SCS Seminar, Illinois Institute of Technology, June 17th, 2011
36. Making a Case for Distributed File Systems at Exascales, ACM LSAP 2011, June 8th, 2011
37. CCGrid 2014 Proposal, CCGrid 2011 Steering Committee Meeting, May 25th, 2011
38. Common Challenges between Exascales and Cloud Computing, IEEE DataCloud 2011, May 16th, 2011
39. Making a Case for Distributed File Systems at Exascales, University of Chicago, May 13th, 2011
40. Making a Case for Distributed File Systems at Exascales, Indiana University, April 27th, 2011
41. DataSys: Data-Intensive Distributed Systems Laboratory Overview, Research Forum, Illinois Institute of Technology, March 22nd, 2011
42. Cloud Computing and Grid Computing 360-Degree Compared, Computer Science Department, Loyola University, March 17th, 2011
43. Reading/Writing, Mathematics, and Computing, CAMRAS Interview Weekend, Illinois Institute of Technology, February 17th, 2011
44. Cloud Computing and Grid Computing 360-Degree Compared, IEEE Fox Valley South Section, Illinois Institute of Technology, January 26th, 2011
45. Avoiding Achilles' Heel in Exascale Computing with Distributed File Systems, DLS Seminar, January 14th, 2011
46. Research Overview, Research Forum, Illinois Institute of Technology, October 18th, 2010
47. Common Challenges in Manycore, Exascale, and Cloud Computing, Computer Science Department, Illinois Institute of Technology, October 4th, 2010
48. Rethinking Storage Systems for Exascale Computing, HPDC 2010, Wild and Crazy Ideas Session, June 25th, 2010
49. 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
50. Exascale Many-Task Computing with a Billion Processors, HPDC Program Committee Workshop 2010, March 2010
51. Many-Task Computing on Grids, Clouds, and Supercomputers, Department of Computer Science, Illinois Institute of Technology, January 22nd, 2010

52. Scalable Resource Management in Cloud Computing, Grid Computing and Supercomputing, College of Computing and Digital Media, DePaul University, January 20th, 2010
53. 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 2009
54. The Quest for Scalable Support of Data Intensive Workloads in Distributed Systems, ACM HPDC09, June 13th, 2009
55. 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
56. The Quest for Scalable Support of Data Intensive Workloads in Distributed Systems, Motorola Labs, March 4th, 2009
57. 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
58. Many-Task Computing: Bridging the Gap between High Throughput Computing and High Performance Computing, Department of Computer Science, Colorado State University, March 2nd, 2009
59. 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
60. 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
61. Scalable Resource Management in Clouds and Grids, Motorola Labs, December 5th, 2008.
62. The Quest for Scalable Support of Data Intensive Applications through Data Diffusion, IEEE/ACM Supercomputing 2008, Argonne National Laboratory Booth, November 20th, 2008.
63. Falkon, a Fast and Light-weight tasK executiON framework for Clusters, Grids, and Supercomputers, IEEE/ACM Supercomputing 2008, Argonne National Laboratory Booth, November 19th, 2008.
64. Running 1 Million Jobs in 10 Minutes via the Falkon Fast and Light-weight tasK executiON framework, IEEE/ACM Supercomputing 2008, Megajob BOF, November 18th, 2008.
65. Toward Loosely Coupled Programming on Petascale Systems, IEEE/ACM Supercomputing 2008, November 18th, 2008.
66. Cloud Computing and Grid Computing 360-Degree Compared, IEEE Grid Computing Environments (GCE08) 2008, November 16th, 2008.
67. The Quest for Scalable Support of Data Intensive Applications through Data Diffusion, University of Chicago, Petascale Active Data Storage (PADS) Seminar, November 12th, 2008.
68. Cloud Computing and Grid Computing 360-Degree Compared, University of Chicago, Grad CS Seminar, November, 2008
69. Scalable Resource Management in Clouds and Grids, Accenture Technology Labs, October 24th, 2008.
70. Systems at University of Chicago, University of Chicago, Department of Computer Science, September 25th, 2008.
71. Scientific Workflow Systems for 21st Century, MWGS08: MidWest Grid School 2008, September 17th, 2008.
72. Harnessing Grid Resources with Data Data-Centric Task Farms, Notre Dame University, CSE Dept., August 20th, 2008.
73. Scientific Workflow Systems for 21st Century, New Bottle or New Wine?, IEEE Workshop on Scientific Workflows, July 2008.
74. Accelerating Large-scale Data Exploration through Data Diffusion, ACM/IEEE International Workshop on Data-Aware Distributed Computing 2008, June 2008.
75. Accelerating Large-Scale Data Exploration through Data Diffusion, DSLW 2008, May 22nd, 2008.
76. Managing and Executing Loosely-Coupled Large-Scale Applications on Clusters, Grids, and Supercomputers, GlobusWorld 2008, May 15th, 2008.
77. Harnessing Grid Resources with Data Data-Centric Task Farms, NASA Ames Research Center, May 14th, 2008.
78. Harnessing Grid Resources with Data Data-Centric Task Farms, Hyde Park Global Investments LLC, April 18th, 2008.
79. Harnessing Grid Resources with Data Data-Centric Task Farms, University of Chicago, Department of Computer Science, Dissertation Proposal, December 12th, 2007.
80. Harnessing Grid Resources with Data Data-Centric Task Farms, University of Chicago, Department of Computer Science, Distributed Systems Lab Seminar, December 5th, 2007.
81. Falkon: a Fast and Light-weight tasK executiON framework for Grid Environments, IEEE/ACM SuperComputing 2007, November 15th, 2007.

82. Accelerating Large Scale Scientific Exploration with Falkon, IEEE/ACM SuperComputing 2007, Argonne National Laboratory Booth, November 14th, 2007.
83. A Data Diffusion Approach to Large Scale Scientific Exploration, University of Chicago, CS Department, Distributed Systems Lab Seminar, October 24th, 2007.
84. A Data Diffusion Approach to Large Scale Scientific Exploration, 2007 Microsoft eScience Workshop, October 2007.
85. Falkon: a Fast and Light-weight tasK executiON framework, DSLW 2007, May 29th, 2007.
86. Falkon: a Fast and Light-weight tasK executiON framework for Grid Environments, CS Seminar, University of Chicago, CS Department, April 30th, 2007.
87. AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis, California Institute of Technology, AstroGrid 2007, February 12th, 2007.
88. AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis, IEEE/ACM SC 2006, November 2006.
89. Storage and Compute Resource Management via DYRE, 3DcacheGrid, and CompuStore, University of Chicago, Department of Computer Science, Distributed Systems Lab Seminar, November 2006.
90. Harnessing Grid Resources to Enable the Dynamic Analysis of Large Astronomy Datasets, DSLW 2006, June 2006
91. AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis, TeraGrid 2006, June 2006
92. Harnessing Grid Resources to Enable the Dynamic Analysis of Large Astronomy Datasets, University of Chicago, Department of Computer Science, Graduate Seminar, February 2006.
93. AstroPortal: A Science Portal to Grid Resources, University of Chicago, Department of Computer Science, Distributed Systems Lab Seminar, January 2006.
94. 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.
95. "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.
96. "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.
97. "DiPerF: an automated DIstributed PERformance testing Framework", University of Chicago: Graduate Seminar, November 2004.
98. "DiPerF: an automated DIstributed PERformance testing Framework", IEEE/ACM GRID2004, November 2004.
99. "A Study between Networks and General Purpose Systems for High Bandwidth Applications", University of Chicago: Computer Architecture, June 2004.
100. "DiPerF: an automated DIstributed PERformance testing Framework", University of Chicago: Grid Computing Seminar, March 2004.
101. "Searching Large Image Databases", University of Chicago: Artificial Intelligence (AI), March 2004.
102. "Mobile IPv6", Sun Microsystems Laboratories: September 2003.
103. "Mobile IPv6", Sun Microsystems Laboratories: August 2003.
104. "NEon", Sun Microsystems Laboratories: August 2003.
105. "Routing Algorithms in Wireless Sensor Networks", Purdue University: Seminar in Peer-to-Peer Networks, October 2002.
106. "Evaluating IPv6 using the Agere 2.5 PayloadPlus Network Processor", Purdue University, October 2002.
107. "Routing Algorithms for Wireless Sensor Networks", Grace Hopper Celebration of Women in Computing, October 2002.
108. "An Empirical Analysis of Internet Protocol version 6 (IPv6)", Wayne State University: MS Defense, April 2002.
109. "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.
110. "MEMS Technology Overview and Limitations", Wayne State University: BioComputing Seminar, April 2002.
111. "Energy-Efficient Routing Algorithms in Wireless Sensor Networks", Wayne State University: Networking Wireless Sensor Lab, February 2002.
112. "IP Encapsulation Methods: A Transition Mechanism to Deploy IPv6 Networks", Wayne State University: Seminar in High Performance Networks, December 2001.

113."Mobility Support in IPv6", Wayne State University: Seminar in High Performance Networks, October 2001

114."Wireless Sensor Networks and their Applications", Wayne State University, September 2001.

115."Proximity Detection using Wireless RF Sensor", Accenture Technology Labs: August 2001.

116."BSD vs. Streams Protocols", IEEE International Conference on Telecommunication (ICT2001): June 2001.

Proposals/Funding

My work has been funded by various National Science Foundation (NSF) programs (CAREER, CRI, REU, CIFellows), various Department of Energy (DOE) national laboratories (ANL, LANL, FNAL), NASA ARC, and Industry (NVIDIA, Microsoft, and Amazon). I have received 26 awards totalling $3.1M in funds, $196K in hardware donation, $100K in credits on public clouds, and 8M hours on large-scale distributed resources.

Active Grants (2017 – 2021, $1.37M in funds)

NSF 2018 - 2021: **"REU Site: BigDataX: From Theory to Practice in Big Data Computing at eXtreme Scales"**

Institution: NSF, CNS
Award: $333,106 (total $370,000)
Period: 03/2018 – 02/2021
Web: http://datasys.cs.iit.edu/grants/BigDataX/
People: Ioan Raicu (IIT/PI), Kyle Hale (IIT/Co-PI), Gruia Calinescu (IIT), Kyle Chard (UChicago/PI), Justin Wozniak (UChicago/ANL), Aaron J. Elmore (UChicago)
Summary: This award aims to establish a Research Experiences for Undergraduates (REU) site named BigDataX, which will focus on undergraduate research in both theory and practice of big data computing at extreme scales. The primary objective of this award is to promote a data-centric view of scientific and technical computing, at the intersection of distributed systems theory and practice.

NSF 2017 - 2020: **"CRI: II-NEW: MYSTIC: prograMmable sYstems reSearch Testbed to explore a stack-wIde adaptive system fabriC"**

Institution: NSF CRI
Award: $1,000,000
Period: 07/2017 - 06/2020
Web: http://mystic.cs.iit.edu
People: Ioan Raicu (IIT/lead-PI), Kyle Hale (IIT/co-PI), Xian-He Sun (IIT/co-PI)
Summary: This proposal seeks to significantly improve programmable infrastructure in the CS department at IIT. The goal of this proposal is to fill the existing void in delivering an open testbed for experimenting with reconfigurable communication and I/O subsystems in order to perform low-level systems research.

Completed Grants (2006 – 2018, $1.7M in funds)

NSF 2011 - 2018: **"Avoiding Achilles' Heel in Exascale Computing with Distributed File Systems"**

Institution: NSF ACI CAREER
Award: $734,170
Period: 01/2011 - 06/2018
Web: http://datasys.cs.iit.edu/grants/NSF-CAREER/
People: Ioan Raicu (IIT/PI), Arthur Barney Maccabe (ORNL), Marc Snir (ANL/UIUC), Rob Ross (ANL), Mike Wilde (ANL/UC), Kamil Iskra (ANL), Jacob Furst (DePaul)
Summary: This award aims to study a radically different storage architecture for high-end computing systems, one which is designed with non-volatile memory on every compute node, and every node to actively participate in the metadata and data management.

NSF 2015 - 2018: **"REU Site: BigDataX: From Theory to Practice in Big Data Computing at eXtreme Scales"**

Institution: NSF, CNS
Award: $288,000
Period: 03/2015 – 02/2018
Web: http://datasys.cs.iit.edu/grants/BigDataX/
People: Ioan Raicu (IIT/PI), Gruia Calinescu (IIT/Co-PI), Justin Wozniak (UChicago/ANL), Mike Wilde (UChicago/ANL), Kyle Chard (UChicago)
Summary: This award aims to establish a Research Experiences for Undergraduates (REU) site named BigDataX, which will focus on undergraduate research in both theory and practice of big data computing at extreme scales. The primary objective of this award is to promote a data-centric view of scientific and technical computing, at the intersection of distributed systems theory and practice.

DOE ANL 2017-2018: **"Data Transfer Optimizations using Storage Layout-Awareness and Software- Defined Networking"**

Institution: DOE ANL
Award: $52,624
Period: 08/2017 – 05/2018
People: Ioan Raicu (IIT/PI), Rajkumar Kettimuthu (ANL)
Summary: This award focuses on research and development in optimizing wide-area data transfer for large-scale science datasets based on the organization of the data in the parallel storage system as well as through utilizing software-defined networking methods. The major goals are: 1) develop algorithms and tools to determine the optimal or near-optimal order for large-scale multi-file or directory transfers, and 2) setup a virtual science environment to evaluate software-defined networking (SDN) techniques to optimize science flows.

Fermi 2011 - 2017: **"Networking and Distributed Systems in High-Energy Physics"**

Institution: Fermi National Accelerator Laboratory, The Virtual Facility Project
Award: $192,360
Period: 05/2011 – 10/2017
People: Ioan Raicu (IIT/PI), Gabriele Garzoglio (FNAL)
Summary: This project has funded two to three graduate students per summer from 2011 to 2017 to conduct research on-site at Fermi National Accelerator Laboratory (FNAL). The work has involved storage systems, networking, and distributed systems research and development for more efficiently operating and scaling high-energy physics data-intensive scientific workflows on the Fermilab computing facility as well as commercial Clouds.

DOE LANL 2013-2015: **"Investigation of Distributed Systems for HPC System Services"**

Institution: DOE LANL
Award: $165,000
Period: 01/2013 – 12/2015
People: Ioan Raicu (IIT/PI), Mike Lang (LANL)
Summary: This award explored the design tradeoffs for scalable system software at extreme scales through (1) a general system software taxonomy, (2) simulations, and (3) real system implementations.

DOE ANL 2014: **"Exploring Distributed File Systems at Extreme Scales through the ROSS Simulator"**

Institution: DOE ANL
Award: $17,144
Period: 01/2014 – 05/2014
People: Ioan Raicu (IIT/PI), Rob Ross (ANL)
Summary: This award aimed to explore through simulations the new storage architecture which collocated persistent storage with compute nodes throughout high-end computing systems at extreme scales of millions of nodes.

IIT 2013 - 2014: **"Towards the Support for Many-Task Computing on Many-Core Computing Platforms"**
Institution: Illinois Institute of Technology, STARR Fellowship
Award: $15,000
Period: 08/2013 - 7/2014
People: Ioan Raicu (IIT/PI), Michael Wilde (ANL), Scott Krieder (IIT/Student)
Summary: This project explores a programming model and runtime environment, which addresses the urgent yet vexing problem of how to simplify the programming of complex hybrid systems architectures. This work specifically tackles the programmability challenges posed by NVIDIA GPUs.

NSF 2011 - 2012: **"Student Travel Support for ACM HPDC 2011"**
Institution: NSF, CCF
Award: $10,000
Period: 05/2011 - 04/2012
Summary: This award supported 20 students to attend the ACM HPDC conference in Chicago in June 2011.

NSF/CRA 2009 - 2010: **"Resource Management in Large-Scale Distributed Systems"**
Institution: NSF, CRA
Award: $140,000
Period: 08/2009 - 07/2010
Summary: This award funded my 1-year postdoc under Alok Choudhary at Northwestern University. I was the lead on this proposal.

NASA 2006 - 2009: **"Harnessing Grid Resources to Enable the Dynamic Analysis of Large Astronomy Datasets"**
Institution: NASA
Award: $84,000
Period: 09/2006 - 03/2009
Summary: This award funded my last 3 years of my PhD degree at University of Chicago. I was the lead on this proposal.

Industry Grants (2011 – 2018, $12K in funds, $276K hardware donation, $100K in credits on public clouds)

NVIDIA 2014 - 2018: **"Student Cluster Challenge"**
Institution: NVIDIA
Award: Donated 8 GPUs (Tesla K40) in 2014, 8 GPUs (Tesla V100 16GB) in 2017, and 8 GPUs (Tesla V100 32GB) in 2018; these GPUs are valued at $266,000 in hardware donations
Period: 10/2014 - 12/2018
People: Ioan Raicu (IIT/PI)

Amazon 2011-2016: **"Distributed Systems on the Amazon Cloud"**
Institution: Amazon AWS
Award: $60,700 in AWS Credits
Period: 01/2011 – 05/2016

Microsoft 2014-2015: **"Implicitly-Parallel Functional Dataflow for Productive Cloud Programming on Windows Azure"**
Institution: Microsoft Research
Award: $40,000
Period: 06/2014 – 05/2015
People: Ioan Raicu (IIT/PI), J. Wozniak (UC/ANL), M. Wilde (UC/ANL), K. Maheshwari (ANL)

NVIDIA 2013 - 2014: **"NVIDIA Teaching Center Program Proposal"**
Institution: NVIDIA
Award: $12,000 and $10,000 in hardware donations
Period: 06/2013 - 15/2014
People: Ioan Raicu (IIT/PI), Zhiling Lan (IIT/Co-PI), Eduardo Berrocal (IIT), Scott Krieder (IIT)

Resource Time Allocations (2011 – 2020, 8M CPU-hours on large-scale distributed resources)

DOE 2017 - 2020: **"Resource Management in Distributed Systems"**
Institution: DOE ANL
Award: 700K hours on the Chameleon Resources
Period: 03/2017 - 12/2020

XSEDE 2017 - 2018: **"Resource Management in Distributed Systems"**
Institution: NSF XSEDE
Award: 200K hours on XSEDE
Period: 06/2017 - 05/2018

DOE 2014 - 2017: **"Student Cluster Competition"**
Institution: DOE ANL ALCF
Award: 200K hours on the ALCF Resources
Period: 03/2014 - 12/2017

XSEDE 2013 - 2014: **"Many-Task Computing with Many-Core Accelerators on XSEDE"**
Institution: NSF XSEDE
Award: 200K hours on XSEDE
Period: 01/2013 - 08/2014

GLCPC 2013 - 2014: **"Implicitly-parallel functional dataflow for productive hybrid programming on Blue Waters"**
Institution: Great Lakes Consortium for Petascale Computation (GLCPC)
Award: 6M hours on the Blue Waters Supercomputer
Period: 03/2013 - 08/2014

NICS 2013 - 2014: **"Many-Task Computing with Many-Core Accelerators on Beacon"**
Institution: National Institute for Computational Sciences (NICS)
Award: 320K hours on the Beacon system
Period: 06/2013 - 08/2014

DOE 2011 - 2013: **"FusionFS: Distributed File Systems for Exascale Computing"**
Institution: DOE ANL ALCF
Award: 450K hours on the IBM BlueGene/P
Period: 03/2011 - 10/2013

DOE 2011 - 2012: **"HPCCloud: Exploring HPC Fault Tolerance in the Cloud"**
Institution: DOE ANL ALCF
Award: 100K hours on the Magellan Cloud
Period: 03/2011 - 02/2012

Projects

I have many research projects which have been under development over the last several years. Many projects are being led by my students in the DataSys lab, but some projects are much bigger and involve multiple institutions, with significant contributions from Argonne National Laboratory, University of Chicago, and Los Alamos National Laboratory.

Active Projects

Mystic — **Programmable Systems Research Testbed to Explore a Stack-WIde Adaptive System fabriC**
Period: 07/2017 – *Present*
Leading People: Ioan Raicu, Xian-He Sun, Kyle Hale
Web Site: http://mystic.cs.iit.edu

XSearch — **Distributed Indexing and Search in Large-Scale Storage Systems**
Period: 01/2015 – *Present*
Leading People: Ioan Raicu, Alexandru Orhean, Kyle Chard
Web Site: http://datasys.cs.iit.edu/projects/xsearch/

XTask — **eXTreme fine-grAined concurrent taSK invocation runtime**
Period: 01/2015 – *Present*
Leading People: Ioan Raicu, Poornima Nookala, Peter Dinda, Kyle Hale, Kyle Chard
Web Site: http://datasys.cs.iit.edu/projects/xtask/

FastWorm — **Studying C.Elegans Behaviour through Big Data Computing**
Period: 01/2015 – *Present*
Leading People: Ioan Raicu, Vineeth Pillai, Daniela Raicu, Jacob Furst, Hoky Kim
Web Site: http://datasys.cs.iit.edu/projects/fastworm/

Parsl — ***High performance workflows***
Period: *01/2019 – Present*
Leading People: Ioan Raicu, Kyle Chard
Web Site: http://parsl-project.org

Swift — ***A simple tool for fast, easy scripting on big machines***
Period: *12/2006 – Present*
Leading People: Ioan Raicu, Kyle Chard
Web Site: http://swift-lang.org/main/

InfinityTrader — **Accelerating Deep Neural Networks for Large-Scale Time Series Datasets**
Period: *01/2020 – Present*
Leading People: Ioan Raicu
Web Site: http://datasys.cs.iit.edu/projects/trader/

Completed Projects

uDNN — **Accelerating Deep Neural Networks through Variable Precision Arithmetic**
Web Site: http://datasys.cs.iit.edu/projects/udnn/

Vertico — **VolunteER cloud compuTIng researCh framewOrk**
Web Site: http://datasys.cs.iit.edu/projects/vertico/

HRDBMS — ***A Scalable Distributed Relational Database for Commodity Hardware*** [08/2013 – 12/2018]
Web Site: http://datasys.cs.iit.edu/projects/HRDBMS/

FusionFS — ***Fusion distributed File System*** [08/2010 – 06/2018]
Web Site: http://datasys.cs.iit.edu/projects/FusionFS/

ZHT — ***Zero-Hop Distributed Hash Table*** [03/2010 – 06/2018]
Web Site: http://datasys.cs.iit.edu/projects/ZHT/

BBSim — **Exploring Burst Buffer Storage Architectures through CODES/ROSS Simulations** [01/2015 – 12/2018]
Web Site: http://datasys.cs.iit.edu/projects/bbsim/

MPathIC	**Multi-Path Network Support on Multi-Dimensional Network Architectures** [01/2014 – 12/2017] Web Site: http://datasys.cs.iit.edu/projects/
GraphZ++	**Lightweight and Scalable Graph Processing System** [01/2016 – 12/2017] Web Site: http://datasys.cs.iit.edu/projects/
Albatros	***Distributed Message Queuing System*** [01/2015 – 12/2016] Web Site: http://datasys.cs.iit.edu/projects/
Fabriq	***Distributed Message Queuing System*** [01/2015 – 12/2016] Web Site: http://datasys.cs.iit.edu/projects/
MATRIX	***MAny-Task computing execution fabRIc at eXascales*** [01/2011 – 06/2015] Web Site: http://datasys.cs.iit.edu/projects/MATRIX/
Slurm++	***Next Generation Distributed Batch Scheduling*** [01/2013 – 06/2015] Web Site: http://datasys.cs.iit.edu/projects/
CloudKon	**A Cloud-enabled Distributed tasK executiON framework** [05/2013 – 12/2016] Web Site: http://datasys.cs.iit.edu/projects/
GeMTC	**GPU enabled Many-Task Computing** [08/2011 – 01/2015] Web Site: http://datasys.cs.iit.edu/projects/GeMTC/
SimMatrix	**SIMulator for MAny-Task computing execution fabRIc at eXascales** [06/2011 – 06/2015] Web Site: http://datasys.cs.iit.edu/projects/SimMatrix/
NoVoHT	***Non-Volatile Hash Table*** [10/2011 – 05/2015] Web Site: http://datasys.cs.iit.edu/
Falkon	***A Fast and Light-weight tasK executiON framework*** [12/2006 – 12/2010] Web Site: http://datasys.cs.iit.edu/projects/Falkon/index.html
ManyCoreSim	***Scheduling Direct Acyclic Graphs on Massively Parallel Processors*** [08/2011 – 10/2013] Web Site: http://datasys.cs.iit.edu/projects/ManyCoreSim_summary12.pdf
IStore	**Information Dispersal Algorithms in Distributed Filesystems** [01/2012 – 08/2013] Web Site: http://datasys.cs.iit.edu/
HyCache	***SSD-Cached Hybrid File System*** [01/2012 – 05/2013] Web Site: http://datasys.cs.iit.edu/projects/SCH-FS_summary12.pdf
RXSim	***Simulator for High-End Computing Systems*** [08/2011 – 05/2013] Web Site: http://datasys.cs.iit.edu/projects/SimHEC_summary12.pdf
CiteSearcher	***A Google Scholar frontend for iOS and Android mobile devices*** [06/2011 – 12/2012] Web Site: http://datasys.cs.iit.edu/projects/CiteSearcher/index.html
AstroPortal	***A Science Gateway for Large-scale Astronomy Data Analysis*** [06/2005 – 03/2009] Web Site: http://datasys.cs.iit.edu/projects/Falkon/astro_portal.htm
DiPerF	***An automated DIstributed PERformance testing Framework*** [09/2003 – 06/2005] Web Site: http://datasys.cs.iit.edu/projects/diperf/
ServMark	***An Architecture for Testing Grid Services*** [01/2006 – 06/2008] Web Site: http://dev.globus.org/wiki/Incubator/ServMark
HOC-SA	***Higher-Order Components-Service Architecture*** [06/2006 – 12/2007] Web Site: http://dev.globus.org/wiki/Incubator/HOC-SA
DI-GRUBER:	***A Distributed Grid Resource Broker*** [09/2003 – 12/2006] Web Site: http://people.cs.uchicago.edu/~cldumitr/GRUBER/

Teaching

Associate Professor Illinois Institute of Technology, Department of Computer Science

08/2016 – Present

Fall 2020	CS 351: Systems Programming
Spring 2020	CS 553: Cloud Computing (64 students)
Fall 2019	CS 550: Advanced Operating Systems (37 students)
Spring 2019	CS 553: Cloud Computing (24 students)
Fall 2018	CS 554: Data-Intensive Computing (38 students)
Fall 2018	CS 495: Cluster Computing (7 students)
Spring 2018	CS 553: Cloud Computing (84 students)
Fall 2017	CS 553: Cloud Computing (123 students)
Fall 2017	CS 554: Data-Intensive Computing (48 students)
Summer 2017	CS 595: Special Topics in Distributed Systems (9 students)

Assistant Professor Illinois Institute of Technology, Department of Computer Science

08/2010 – 08/2016

Summer 2016	CS 595: Special Topics in Distributed Systems (8 students)
Spring 2016	CS 553: Cloud Computing (186 students)
Fall 2015	CS 550: Advanced Operating Systems (115 students)
Spring 2015	CS 554: Data-Intensive Computing (62 students)
Fall 2014	CS 553: Cloud Computing (165 students)
Spring 2014	CS 451: Introduction to Parallel and Distributed Computing (29 students)
Fall 2013	CS 554: Data-Intensive Computing (52 students)
Spring 2013	CS 553: Cloud Computing (75 students)
Fall 2012	CS 495: Introduction to Distributed Computing (18 students)
Spring 2012	CS 553: Cloud Computing (54 students)
Spring 2012	CS 695: Doctoral Seminar (16 students)
Fall 2011	CS 595: Data-Intensive Computing (32 students)
Spring 2011	CS 550: Advanced Operating Systems (53 students)
Spring 2011	CS 695: Doctoral Seminar (11 students)
Fall 2010	CS 595: Data-Intensive Computing (9 students)

Adjunct Assistant Professor Northwestern Univ., Dept. of Electrical Eng. and Computer Science

01/2010 – 06/2010

EECS 211: Fundamentals of Computer Programming (40 students)
EECS 495: Hot Topics in Distributed Systems: Data-Intensive Computing (10 students)

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

06/2002 – 08/2002

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 Mentoring and Advising

Current PhD Students (Advisor) at Illinois Institute of Technology

Alexandru Iulian Orhean (4th year), Expected graduation 2022, "Distributed Indexing and Search in Large-Scale Storage Systems"
Nookala Poornima (4th year), Expected graduation 2022, "Efficient Support for Many-Task Computing on Accelerators"
Vineeth Pillai (3rd year, part time), Expected graduation 2023, "Exploring Efficient Support for Big Data Scientific Computing"
Yi Qu (1st year), Expected graduation 2025
Hao Ding (1st year), Expected graduation 2025

Graduated PhD Students (Advisor) at Illinois Institute of Technology (IIT)

Iman Sadooghi, IIT/CS, PhD 12/2016, "Scalable Resource Management in Cloud Computing"
Cloud Architect @ Bank of America
Tonglin Li, IIT/CS, PhD 12/2015, "Distributed NoSQL Storage for Extreme-Scale System Services in Supercomputers and Clouds"
Research Scientist @ LBL
Ke Wang, IIT/CS, PhD 08/2015, "Towards Next Generation Resource Management at Extreme-Scales"
Engineer @ Microsoft
Dongfang Zhao, IIT/CS, PhD 08/2015, "Big Data System Infrastructure at Extreme Scales"
Assistant Professor @ University of Nevada, Reno

Graduated MS Students with Thesis (Advisor) at Illinois Institute of Technology

Itua E. Ijagbone, IIT/CS, MS 05/2016, "*Scalable Indexing and Searching on Distributed File Systems"*
Anupam Rajendran, IIT/CS, MS 05/2013, "*MATRIX: Many-Task Computing Execution Fabric for Extreme Scales" – Dell*

Alumni (with a thesis) – Committee Member

Xu Yang, Institution: IIT/CS, PhD 04/2017, "*Cooperative Batch Scheduling for HPC Systems*", Primary Advisor: Dr. Zhiling Lan
Eduardo Berrocal, Institution: IIT/CS, PhD 04/2017, "*Improving Fault Tolerance for Extreme Scale Systems*", Primary Advisor: Dr. Zhiling Lan
Sean Wallace, Institution: IIT/CS, PhD 04/2017, *"Power Profiling, Analysis, Learning, and Management for High-Performance Computing*", Primary Advisor: Dr. Zhiling Lan
Zhou Zhou, Institution: IIT/CS, PhD 12/2015, "*Multi-Dimensional Batch Scheduling Framework for High-End Supercomputers*", Primary Advisor: Dr. Zhiling Lan
Jingjin Wu, Institution: IIT/CS, PhD 08/2013, "*Performance Analysis and Optimization of Large-Scale Scientific Applications*", Primary Advisor: Dr. Zhiling Lan – UESTC, China
Wei Tang, Institution: IIT/CS, PhD 08/2012, "*An Integrated Resource Management and Scheduling Framework for Production Supercomputers*", Primary Advisor: Dr. Zhiling Lan – Argonne National Laboratory, Lemont, IL
Hui Jin, Institution: IIT/CS, PhD 05/2012, "*System Support for Resilience in Large-Scale Parallel Systems*", Primary Advisor: Dr. Xian-He Sun – Oracle, San Francisco, CA
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 – University of Chicago, Chicago, IL
Yuchi Tsao, IIT/ECE, PhD 12/2011; "*Architecture-Level Hardware Optimization for Digital Signal Processing*", Primary Advisor: Dr. Kyu-won Choi

Student Cluster Competition at IEEE/ACM Supercomputing/SC

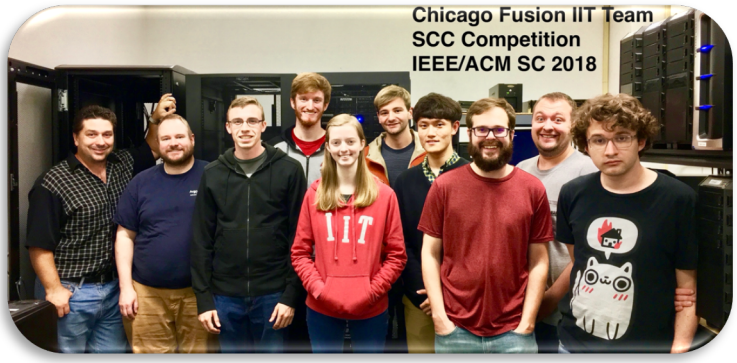

An out of the classroom activity that I found to have worked wonders engage undergraduate students at IIT, was to participate in the Student Cluster Competition (SCC) at the IEEE/ACM Supercomputing/SC 2014, 2015, 2017, and 2018 conferences. This competition brings together teams of undergraduate/highschool students from all around the world to compete in running 5 high-performance applications over a 48-hour period on sponsored hardware (Argonne, Intel, and Mellanox) that they have built and configured. The team ranked 8th out of 13 teams in 2014, 4th out of 9 teams in 2015, 6th out of 16 teams in 2017, and 9th out of 16 teams in 2018. The co-coaches were Ioan Raicu, William Scullin, and Ben Allen, with generous support from Argonne National Laboratory specifically from Michael Papka at ALCF, Intel Corporation, Mellanox, and NVIDIA. More online information about our participation in this competition can be found at: http://datasys.cs.iit.edu/grants/BigDataX/2018/scc.html

NSF REU BigDataX Program

In 2015, I was awarded a 3-year NSF REU site called BigDataX that funded 24 undergraduate students to do research in big data computing. In 2018, we were renewed for an additional 3-year period to support 30 more students in the same program. The BigDataX program hosts 10 students every summer for 10 weeks doing research in one of two interdisciplinary laboratories: 1) the DataSys Lab in CS at IIT or 2) the Systems Group in the CI/CS at UChicago. The two labs physically close (4 miles apart), and students will have weekly activities that place all students in a single lab to strengthen the cohort experience. The PIs aim to recruit most of the students from outside of the host institutions, focusing on recruiting students from institutions without research opportunities as well as women and minorities. The primary objective of this award is to promote a data-centric view of scientific and technical computing, at the intersection of distributed systems theory and practice. This award has five mentors with a variety of complementing expertise from theory to programming languages to distributed systems. This work includes a comprehensive educational plan integrating eight undergraduate students with senior PhD students with incremental manageable goals, aimed at allowing undergraduate students to achieve publishable results within the ten-week summer program. More information can be found at http://datasys.cs.iit.edu/grants/BigDataX/.

Upsilon Pi Epsilon (International Honor Society for the Computing and Information Disciplines)

In 2017, IIT Computer Science founded the local UPE chapter and inducted its first members. In 2018, I became the faculty advisor of the UPE chapter. As of 2020, we have over 50 active members, and over 30 alumni who have graduated. The mission of our UPE chapter is to recognize academic excellence at both the undergraduate and graduate levels in the computing field. Our students participate in several events throughout the year, from UPE general body meetings, to site visits at local companies, seminar speakers, advising sessions with incoming freshman students, to initiation ceremonies when new members are inducted into the UPE. More information can be found at http://upe.cs.iit.edu.

Mystic: Programmable Systems Research Testbed to Explore a Stack-WIde Adaptive System fabriC

I served as the lead PI from 2017 to 2020 on a NSF CRI testbed that established and maintained a dynamically configurable testbed called MYSTIC. The testbed aimed to study system re-configurability across the entire computing stack, from the processor to memory, storage, and the network. It allows low-level experimentation and reconfiguration at the level of networks-on-chip (NoC), universal memory, and the network interconnect with multidimensional network topologies. Dynamically reconfiguring interconnects, memory and storage will speed-up applications running on a heterogeneous computing environment. This is work joint with Xian-He Sun and Kyle Hale, among several other faculty and students. More information can be found at http://mystic.cs.iit.edu.

University Service

- Upsilon Pi Epsilon (UPE) Honor Society, Faculty Mentor (Fall 2018 - Present)
- Strategic Vision Committee, Member (Spring 2019 – Fall 2019)
- Graduate Studies Committee, Participant (Fall 2011 – Spring 2013, Fall 2017 – Present)
- Teaching Cluster Committee, Chair (Fall 2017 – Fall 2018)
- Undergraduate Studies Committee, Member (Fall 2010 – Spring 2013, Fall 2017 – Spring 2018)
- Seminar Committee, Chair (Fall 2010 – Spring 2012), Member (Fall 2012 – Spring 2016)
- Faculty Search Committee, Participant (Fall 2010 – Spring 2011), Member (Fall 2011 – Spring 2016)
- Qualifier Preparation (Spring 2011 - Spring 2016)
- Undergraduate Admissions (CAMRAS Interviews), Member (Spring 2011 - Spring 2016)
- Graduate Admissions, Member (Spring 2011, Fall 2014 – Spring 2016)

External Service (Post-Chronological Order, Since 2014)

I have founded and chaired several workshops, such as IEEE/ACM Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS) and the ACM Workshop on Scientific Cloud Computing (ScienceCloud). I am also the local organizer chair of the 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid) 2014. I am on the editorial board of the IEEE Transaction on Cloud Computing, the Springer Journal of Cloud Computing Advances, Systems and Applications (JoCCASA), and the Springer Cluster Computing Journal (Cluster). I have been in leadership roles in several high profile conferences, such as HPDC, CCGrid, Grid, eScience, ICAC, Cluster, and CloudCom.

Panels

NSF: NSF CISE CAREER Panel, 2017, 2019 – Panelist
NSF: NSF CISE CRI Panel, 2018 – Panelist
NSF: NSF CISE REU Panel, 2015, 2016 – Panelist
NSF: NSF CISE CSR Panel, 2013, 2015 – Panelist

Journal Editor

TCC: IEEE Transactions on Cloud Computing, 2013-Present – *Associate Editor*
Cluster: Springer Cluster Computing Journal, 2013-Present – *Editorial Board*
JoCCASA: Springer Journal of Cloud Computing: Advances, Systems and Applications, 2011-Present – *Assoc. Editor*
TCC-MTC: IEEE Transaction on Cloud Computing, SI on Many-Task Computing in the Cloud, 2015 – *Guest Editor*
TCC-SC: IEEE Transactions on Cloud Computing, SI on Scientific Cloud Computing, 2014-2015 – *Guest Editor*

Journal Reviewer

TPDS: IEEE Transactions on Parallel and Distributed Systems, 2009, 2010, 2011, 2012, 2014, 2015, 2016, 2017
TCC: IEEE Transactions on Cloud Computing, 2014, 2015, 2016, 2017
TC: IEEE Transactions on Computers, 2006, 2009, 2010, 2015
CCPE: Journal of Concurrency and Computation: Practice and Experience, 2006, 2009, 2010, 2014, 2015

Chair / Organizer

SC: IEEE/ACM Int. Conf. on High Performance Computing, Networking, Storage, and Analysis, '18 – *Cloud/Grids Track Chair*
GCASR: Greater Chicago Area System Research Workshop, 2013, 2017 – *General Chair*
HPDC: ACM Symp. on High Performance Distributed Computing, 2010, 2015, 2016 – *Stud. Activities Chair, Workshop Chair*
IPDPS: IEEE International Parallel & Distributed Processing Symposium, 2016 – *Local Chair*
BDC: IEEE/ACM International Symposium on Big Data Computing (BDC), 2014, 2015 – PC Chair
MTAGS: ACM/IEEE Workshop on Many-Task Computing on Clouds, Grids, and Supercomputers (at SC), 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015 -- *General Chair*
CCGrid: Cluster, Cloud, and Grid Computing Conference, 2013, 2014 – *Workshops Chair, Local Chair*
CloudCom: IEEE Int. Conference on Cloud Computing Technology and Science, 2014 – *Big Data Track Chair*
ScienceCloud: ACM Workshop on Scientific Cloud Computing (at HPDC), 2010, 2011, 2013, 2014 – *General Chair*
TCSC: IEEE Technical Committee on Scalable Computing (TCSC), 2012, 2013, 2014 **--** *Cyber Infrastructure Chair*
CASK: Workshop on Collaborative methodologies to Accelerate Scientific Knowledge discovery in big data, 2014 – General Chair

Other Leading Roles (Steering Committee Member, Publicity Chair)

HPDC: ACM Symp. on High Performance Distributed Computing, 2011, 2012, 2013, 2014, 2020 -- *Publicity* Chair
CCGrid: Cluster, Cloud, and Grid Computing Conference, 2012, 2013, 2017 – *Publicity Chair*
ScienceCloud: ACM Workshop on Scientific Cloud Computing (at HPDC), 2012, 2015, 2016, 2017 –*Steering Committee*
GCASR: Greater Chicago Area System Research Workshop, 2012, 2013, 2014, 2015, 2016, 2017 – *Steering Committee*
DataCloud: ACM/IEEE Int. Workshop on Data-Intensive Computing in the Clouds (at SC), 2013, 2014, 2015 – *Steering Committee*
Cluster: IEEE Cluster, 2015 – *Publicity Chair*

Conference Program Committee

SC: IEEE/ACM Int. Conf. on High Performance Computing, Networking, Storage, and Analysis, '12,'13,'14,'15,'16,'17,'18,'19,'20
HPDC: ACM Int. Symposium on High Performance Distributed Computing, 2010, 2011, 2012, 2013, 2014, 2016, 2019, 2020
IPDPS: IEEE International Parallel & Distributed Processing Symposium, 2014, 2019
CLOUD: IEEE International Conference on Cloud Computing, 2011, 2012, 2014, 2015, 2017, 2018
eScience: IEEE e-Science Conference, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017
ICA3PP: International Conference on Algorithms and Architectures for Parallel Processing, 2017
CCGrid: Cluster, Cloud, and Grid Computing Conference, 2012, 2013, 2014, 2016
Cluster: IEEE Cluster, 2014, 2015, 2016
CloudCom: IEEE International Conference on Cloud Computing Technology and Science, 2010, 2014, 2015, 2016
BDSEA: IEEE/ACM International Conference on Big Data Science, Engineering and Applications, 2016
MASCOTS: Modelling, Analysis, and Simulation on Computer and Telecommunication Systems, 2015
InterCloud-HPC: Int. Symposium on Cloud Computing and Services for High Performance Computing Systems, 2015
3PGCIC: International Conference on P2P, Parallel, Grid, Cloud, and Internet Computing, 2015
ICDCS: International Conference on Distributed Computing Systems, 2014
CLOSER: International Conference on Cloud Computing and Services Science, 2012, 2014
HPCC: IEEE International Conference on High Performance Computing and Communications, 2013, 2014

Workshop Program Committee

ScienceCloud: ACM Workshop on Scientific Cloud Computing (at HPDC), 2017
ARMS-CC: Workshop on Adaptive Resource Management and Scheduling for Cloud Computing, 2015, 2016, 2017
DIDC: ACM Workshop on Data-Intensive Distributed Computing, 2009, 2010, 2011, 2012, 2014, 2016
IPDRM: IEEE International Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware, 2016
SCRAMBL: Int. Workshop on Scalable Computing For Real-Time Big Data Applications, 2014, 2015, 2016
BigSystem: Int. Workshop on Software-Defined Ecosystems, 2014

Graduate Advisors and Postdoctoral Sponsors

- **Alok Choudhary** (Northwestern University, Postdoctoral Sponsor)
- **Ian T. Foster** (University of Chicago and Argonne National Laboratory, MS & PhD Thesis Advisor)
- **Sherali Zeadally** (University of the District of Columbia, MS Thesis Advisor)

References (close collaborators)

Ian Foster — Arthur Holly Compton Distinguished Service Professor, Computer Science Department, University of Chicago
Senior Fellow, Computation Institute, University of Chicago
Distinguished Fellow, Mathematics and Computer Science Division, Argonne National Laboratory
Contact: 1-630-252-4619, foster@anl.gov

Kyle Chard — Senior Research Project Professional, Computation Institute, University of Chicago
Contact: 1-773-681-2181, kyle@ci.uchicago.edu

Chris Gladwin — CEO, Ocient
Contact: cgladwin@ocient.com

Peter Dinda — Professor, Department of Electrical Engineering and Computer Science, Northwestern University
Contact: 1-847-467-7859, pdinda@northwestern.edu

Kyle Hale — Assistant Professor, Department of Computer Science, Illinois Institute of Technology
Contact: 1-312-567-5147, khale@cs.iit.edu

Zhiling Lan — Professor, Department of Computer Science, Illinois Institute of Technology
Contact: 312.567.5710, lan@iit.edu