



# 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](mailto:iraicu@cs.iit.edu)

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

Laboratory: <http://datasys.cs.iit.edu/>

Google Scholar: <http://tinyurl.com/ioan-raicu>



## Education

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

## Work Experience

Associate Professor	<a href="#">Illinois Institute of Technology (CS)</a> , Chicago IL	08/2016 – Present
Guest Research Faculty	<a href="#">Argonne National Laboratory (MCS)</a> , Lemont IL	01/2011 – Present
Council Member	<a href="#">Upsilon Pi Epsilon (UPE)</a> , Chico, CA	03/2020 – Present
Technical Advisory Board Member	<a href="#">FusionBlock</a> , Singapore	08/2018 – 12/2022
Technical Advisory Board Member	<a href="#">Ocient LLC</a> , Chicago IL	04/2016 – 04/2020
Visiting Scholar (Sabbatical)	<a href="#">Northwestern University (EECS)</a> , Evanston IL	09/2016 – 06/2017
Assistant Professor	<a href="#">Illinois Institute of Technology (CS)</a> , Chicago IL	08/2010 – 07/2016
Computation Innovation Postdoc Fellow	<a href="#">Northwestern University (EECS)</a> , Evanston IL	08/2009 – 07/2010
Research Visitor	<a href="#">NASA (ARC/NAS)</a> , Moffet Field CA	03/2009 – 05/2009
Teaching/Research Assistant	<a href="#">University of Chicago (CS)</a> , Chicago IL	09/2003 – 03/2009
Researcher (Internship)	<a href="#">Argonne National Laboratory (MCS)</a> , Lemont IL	Summer 2005 & 2006
Researcher (Internship)	<a href="#">Sun Microsystems</a> (Sun Labs & SNT), Menlo Park CA	Summer 2003
Teaching Assistant	<a href="#">Purdue University (CS)</a> , West Lafayette IN	08/2002 – 05/2003
Adjunct Assistant Professor	<a href="#">University of Michigan (CIS)</a> , Dearborn MI	06/2002 – 08/2002
Teaching/Research Assistant	<a href="#">Wayne State University (CS)</a> , Detroit MI	08/2000 – 08/2002
Researcher (Internship)	<a href="#">Accenture Technology Labs</a> , Palo Alto CA	Summer 2001
System Analyst (Internship)	<a href="#">Ford Motor Company</a> , Dearborn MI	Summer 1999
Owner	<a href="#">High Teck Computers</a> , Westland MI	01/1997 – 03/2001

## Awards

<b>Teacher of the Year</b>	Illinois Institute of Technology, Computer Science Department	2016
<b>Young Achievers in Scalable Computing</b>	IEEE Technical Committee on Scalable Computing (TCSC)	2014
<b>Outstanding Service Award</b>	IEEE/ACM CCGrid	2014
<b>Junior Faculty Research Award</b>	Illinois Institute of Technology	2013
<b>NSF CAREER Award</b>	National Science Foundation, Office of Cyber Infrastructure	2011-2018
<b>Computation Innovation Fellow</b>	National Science Foundation & Computing Research Assoc.	2009-2010
<b>GSRP Fellowship</b>	NASA, Ames Research Center	2006-2008
<b>Presidential Scholarship</b>	Wayne State University	1997

---

## Research

---

### Research Interests

*Distributed Systems*

*Computer Systems*

*Blockchain Technologies*

### Publications

I have published 140 peer-reviewed bodies of work, receiving 12657 citations, yielding a [H-index](https://scholar.google.com/citations?hl=en&user=jE73HYAAAAAJ) of 46. More information: <http://scholar.google.com/citations?hl=en&user=jE73HYAAAAAJ>.

#### Top 20 High-Impact Publications with 9487 citations:

5012 cites	1. Ian Foster, Yong Zhao, Ioan Raicu, Shiyong Lu. "Cloud Computing and Grid Computing 360-Degree Compared", IEEE GCE 2008
884 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
560 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 SWF 2007
453 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
433 cites	5. Ioan Raicu, Ian Foster, Yong Zhao. "Many-Task Computing for Grids and Supercomputers", IEEE MTAGS 2008
261 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, January 2006
212 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 IPDPS 2013
172 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 SuperComputing/SC 2008
163 cites	9. Ke Wang, Xiaobing Zhou, Tonglin Li, Michael Lang, Ioan Raicu. "Optimizing Load Balancing and Data-Locality with Data-aware Scheduling", IEEE Big Data 2014
153 cites	10. 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 Big Data 2014
147 cites	11. 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
140 cites	12. Yong Zhao, Ioan Raicu, Shiyong Lu, Xubo Fei. "Opportunities and Challenges in Running Scientific Workflows on the Cloud", IEEE CyberC 2011
132 cites	13. 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, 2009
126 cites	14. Iman Sadooghi, Jesús Hernández Martín, 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), 2015
115 cites	15. Ioan Raicu, Yong Zhao, Ian Foster, Alex Szalay. "Accelerating Large-scale Data Exploration through Data Diffusion", IEEE DADC 2008
113 cites	16. Yong Zhao, Ioan Raicu, Ian Foster. "Scientific Workflow Systems for 21st Century e-Science, New Bottle or New Wine?", Invited Paper, IEEE SWF 2008
110 cites	17. Hui Jin, Xi Yang, Xian-He Sun, Ioan Raicu. "ADAPT: Availability-aware MapReduce Data Placement in Non-Dedicated Distributed Computing", IEEE ICDCS 2012
106 cites	18. Ioan Raicu, Pete Beckman, Ian Foster. "Making a Case for Distributed File Systems at Exascale", ACM LSAP 2011
103 cites	19. Ioan Raicu, Sherali Zeadally. "Evaluating IPv4 to IPv6 Transition Mechanisms", IEEE ICT 2003
92 cites	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

---

## Theses & Dissertations (11)

---

1. Alexandru Iulian Orhean, Ioan Raicu (advisor). "Scalable Indexing and Search in High-End Computing Systems", Illinois Institute of Technology, Department of Computer Science, PhD Dissertation, 2023
  2. Poornima Nookala, Ioan Raicu (advisor). "Extreme Fine-grained Parallelism on Modern Many-Core Architectures", Illinois Institute of Technology, Department of Computer Science, PhD Dissertation, 2022
  3. Iman Sadooghi, Ioan Raicu (advisor). "Scalable Resource Management in Cloud Computing", Illinois Institute of Technology, Computer Science Department, Doctorate Dissertation, September, 2016
  4. Itua Ijagbone, Ioan Raicu. "Scalable Indexing and Searching on Distributed File Systems", Department of Computer Science, Illinois Institute of Technology, MS Thesis, 2016
  5. 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
  6. Dongfang Zhao, Ioan Raicu. "Big Data System Infrastructure at Extreme Scales", Illinois Institute of Technology, Computer Science Department, Doctorate Dissertation, July 2015
  7. Ke Wang, Ioan Raicu. "Scalable Resource Management System Software for Extreme-Scale", Illinois Institute of Technology, Computer Science Department, Doctorate Dissertation, June 2015
  8. Anupam Rajendran, Ioan Raicu. "MATRIX: Many-Task Computing Execution Fabric for Extreme Scales", Department of Computer Science, Illinois Institute of Technology, MS Thesis, 2013
  9. 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
  10. 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
  11. Ioan Raicu, Sherali Zeadally. "An Empirical Analysis of Internet Protocol version 6 (IPv6)", Wayne State University, Computer Science Department, MS Thesis, May 2002
- 

## Journal Articles (26)

---

12. Alexandru Iulian Orhean, Anna Giannakou, Katie Antypas, Ioan Raicu, Lavanya Ramakrishnan. "Evaluation of a Scientific Data Search Infrastructure", Concurrency and Computation: Practice and Experience (CCPE), 2022
  13. 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
  14. 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
  15. 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
  16. 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
  17. 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)
  18. 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)
  19. 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)
-

- 
20. 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)
  21. 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
  22. 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)
  23. 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", Concurrency and Computation: Practice and Experience (CCPE), 2015; DOI:10.1002/cpe.3614 (Impact Factor 1.0)
  24. 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
  25. 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
  26. 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)
  27. 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)
  28. Michael Wilde, Ian Foster, Pete Beckman, Ioan Raicu. "Scalable Parallel Scripting for Scientific Computing", SciDAC Review Spring 2010, pp 38 – 53
  29. 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)
  30. 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)
  31. 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
  32. 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)
  33. 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)
  34. 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)
  35. 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)
  36. 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

- 
37. 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)
- 

### Conference Papers (41)

---

38. Alexandru Iulian Orhean, Anna Giannakou, Lavanya Ramakrishnan, Kyle Chard, Ioan Raicu. "SCANNS: Towards Scalable and Concurrent Data Indexing and Searching in High-End Computing System", The 22nd IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid) 2022 (25% acceptance rate)
39. Poornima Nookala, Peter Dinda, Kyle C Hale, Kyle Chard, Ioan Raicu. "Enabling Extremely Fine-grained Parallelism via Scalable Concurrent Queues on Modern Many-core Architectures", 29th International Symposium on Modelling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS) 2021 (24% acceptance rate)
40. 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)
41. 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)
42. 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
43. 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)
44. 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)
45. 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)
46. 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)
47. 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)
48. 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)
49. 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)
50. Yong Zhao, Youfu Li, Ioan Raicu, Shiyong Lu, Xuan Zhang. "Architecting Cloud Workflow: Theory and Practice", IEEE CIT 2014 (26% acceptance rate)
51. 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)
52. Ke Wang, Xiaobing Zhou, Hao Chen, Michael Lang, Ioan Raicu. "Next Generation Job Management Systems for Extreme Scales", ACM HPDC 2014 (24% acceptance rate)
53. Tonglin Li, Ioan Raicu, Lavanya Ramakrishnan. "Scalable State Management for Scientific Applications in the Cloud", IEEE BigData 2014 (19% acceptance rate)
54. 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)
55. Dongfang Zhao, Kan Qiao, Ioan Raicu. "HyCache+: Towards Scalable High-Performance Caching Middleware for Parallel File Systems", IEEE/ACM CCGrid 2014 (19% acceptance rate)
56. 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
-



- 
57. Dongfang Zhao, Chen Shou, Tanu Malik, Ioan Raicu. "Distributed Data Provenance for Large-Scale Data-Intensive Computing", IEEE Cluster 2013 (31% acceptance rate)
  58. 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)
  59. 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)
  60. 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)
  61. Anupam Rajendran, Parag Mhashikar, Hyunwoo Kim, Dave Dykstra, Gabriele Garzoglio, Ioan Raicu. "Optimizing Large Data Transfers over 100Gbps Wide Area Networks", IEEE/ACM CCGrid 2013 (22% acceptance rate)
  62. Ke Wang, Kevin Brandstatter, Ioan Raicu. "SimMatrix: Simulator for MAny-Task computing execution fabRIc at eXascales", ACM HPC 2013
  63. Dongfang Zhao, Da Zhang, Ke Wang, Ioan Raicu. "Exploring Reliability of Exascale Systems through Simulations", ACM HPC 2013
  64. 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)
  65. Yong Zhao, Ioan Raicu, Shiyong Lu, Xubo Fei. "Opportunities and Challenges in Running Scientific Workflows on the Cloud", IEEE CyberC 2011 (21% acceptance rate)
  66. 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)
  67. 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)
  68. Ioan Raicu, Yong Zhao, Catalin Dumitrescu, Ian Foster, Mike Wilde. "Falcon: a Fast and Light-weight task execution framework", IEEE/ACM SuperComputing/SC 2007 (20% acceptance rate)
  69. Ioan Raicu, Ian Foster, Alex Szalay, Gabriela Turcu. "AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis", TeraGrid 2006
  70. 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)
  71. Catalin Dumitrescu, Ioan Raicu, Ian Foster. "DI-GRUBER: A Distributed Approach for Grid Resource Brokering", IEEE/ACM SuperComputing/SC 2005 (22% acceptance rate)
  72. Catalin Dumitrescu, Ioan Raicu, Ian Foster. "Experiences in Running Workloads over Grid3", GCC 2005
  73. Catalin Dumitrescu, Ioan Raicu, Matei Ripeanu, Ian Foster. "DiPerF: an automated Distributed PERFORMANCE testing Framework", IEEE/ACM GRID 2004 (22% acceptance rate)
  74. Ioan Raicu, Loren Schwiebert, Scott Fowler, Sandeep K.S. Gupta. "e3D: An Energy-Efficient Routing Algorithm for Wireless Sensor Networks", IEEE ISSNIP 2004
  75. Ioan Raicu, Sherali Zeadally. "Impact of IPv6 on End-User Applications", IEEE ICT 2003
  76. Ioan Raicu, Sherali Zeadally. "Evaluating IPv4 to IPv6 Transition Mechanisms", IEEE ICT 2003
  77. Ioan Raicu. "Efficient Even Distribution of Power Consumption in Wireless Sensor Networks", ISCA CATA 2003
  78. 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
- 

## Book Chapters (5)

- 
79. Dongfang Zhao, Akash Mahakode, Sandip Lakshminarasiah, 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
  80. 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
-

- 
81. 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
  82. 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
  83. Catalin Dumitrescu, Jan D nnweber, Philipp L deking, Sergei Gorlatch, Ioan Raicu, Ian Foster. "Simplifying Grid Application Programming Using Web-Enabled Code Transfer Tools". Toward Next Generation Grids, 2007
- 

## Workshop Papers (19)

---

84. Jamison Kerney, John Raicu, Kyle Chard, Ioan Raicu. "Towards Fine-grained Parallelism in Parallel and Distributed Python Libraries", IEEE International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS) 2024, co-located with IEEE IPDPS 2024
  85. Ioan Raicu, Alexandru Iulian Orhean, Kyle Chard, Lavanya Ramakrishnan, Anna Giannakou. "Next Generation Indexing and Search in Large-Scale Scientific Storage Systems", ASCR Workshop on the Management and Storage of Scientific Data 2022
  86. 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
  87. Tonglin Li, Kate Keahey, Ke Wang, Dongfang Zhao, Ioan Raicu. "A Dynamically Scalable Cloud Data Infrastructure for Sensor Networks", Invited Paper, ACM ScienceCloud 2015
  88. 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
  89. 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
  90. Chen Shou, Dongfang Zhao, Tanu Malik, Ioan Raicu. "Towards a Provenance-Aware a Distributed File System", USENIX TaPP 2013
  91. Ke Wang, Zhangjie Ma, Ioan Raicu. "Modelling Many-Task Computing Workloads on a Petaflop IBM BlueGene/P Supercomputer", IEEE CloudFlow 2013
  92. Dongfang Zhao, Ioan Raicu. "HyCache: A User-Level Caching Middleware for Distributed File Systems", IEEE HPDIC 2013
  93. 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
  94. 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
  95. Ian Foster, Yong Zhao, Ioan Raicu, Shiyong Lu. "Cloud Computing and Grid Computing 360-Degree Compared", IEEE Grid Computing Environments (GCE08) 2008
  96. 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
  97. 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
  98. 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
  99. 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
  100. Ioan Raicu, Ian Foster, Alex Szalay, Gabriela Turcu. "AstroPortal: A Science Gateway for Large-scale Astronomy Data Analysis", TeraGrid Conference 2006, June 2006
  101. Alex Szalay, Julian Bunn, Jim Gray, Ian Foster, Ioan Raicu. "The Importance of Data Locality in Distributed Computing Applications", NSF Workflow Workshop 2006
-

- 
102. 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] (40)**

---

103. Lan Nguyen, Ioan Raicu. "Accelerating CRUD with Chrono Dilation for Time-Series Storage Systems", IEEE/ACM SuperComputing/SC 2023
104. Aahad Abubaker, Tanmay Anand, Sonal Gaikwad, Mahad Haider, Jackie Mcaninch, Lan Nguyen, Alexandru Orhean, Ioan Raicu. "Exploring Green Cryptographic Hashing Algorithms", IEEE/ACM SuperComputing/SC 2023
105. Caleb Lehman, Poornima Nookala, Ioan Raicu. "Scalable Load-Balancing Concurrent Queues on Many-Core Architectures", IEEE/ACM SuperComputing/SC 2019
106. 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
107. 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
108. Vineeth Remanan Pillai, Daniela Stan Raicu, Ioan Raicu. "Accelerating Worm Segmentation through Inter-node Parallelism", Engineering in Medicine and Biology Conference (EMBC) 2018
109. 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
110. 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
111. 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
112. Prajakt Shastri, Daniel Parker, Sanjiv Kapoor, Ioan Raicu. "Exploring Randomized Multipath Routing on Multi-Dimensional Torus Networks", IEEE/ACM SuperComputing/SC 2016
113. Jian Peng, Sugghosh Divanji, Ioan Raicu, Mike Lang. "Simulating the Burst Buffer Storage Architecture on an IBM BlueGene/Q Supercomputer", IEEE/ACM SuperComputing/SC 2016
114. Jonathan Wu, Suraj Chafle, Ioan Raicu, Kyle Chard. "Optimizing Search in Un-Sharded Large-Scale Distributed Systems", IEEE/ACM SuperComputing/SC 2016
115. Alex Ballmer, Ioan Raicu. "FemtoGraph: A Pregel Based Shared-memory Graph Processing Library", IEEE/ACM SuperComputing/SC 2016
116. 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
117. 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
118. Tonglin Li, Ioan Raicu. "Distributed NoSQL Storage for Extreme-Scale System Services", Doctoral Showcase, IEEE/ACM Supercomputing/SC 2015
119. 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
120. Jason Arnold, Boris Glavic, Ioan Raicu. "HRDBMS: A NewSQL Database for Analytics", IEEE Cluster 2015
121. Chaoqi Ma, Jiabao Li, Tonglin Li, Ioan Raicu. "GRAPH/Z: A Key-Value Store Based Scalable Graph Processing System", IEEE Cluster 2015
-



- 
122. Poornima Nookala, Serapheim Dimitropoulos, Karl Stough, Ioan Raicu. "Evaluating the Support of MTC Applications on Intel Xeon Phi Many-Core Accelerators", IEEE Cluster 2015
  123. 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
  124. Dongfang Zhao and Ioan Raicu. "Storage Support for Data-Intensive Applications on Extreme-Scale HPC Systems", Doctoral Showcase, IEEE/ACM Supercomputing/SC 2014
  125. Tonglin Li, Kate Keahey, Ioan Raicu. "A Cloud-based Interactive Data Infrastructure for Sensor Networks", IEEE/ACM Supercomputing/SC 2014
  126. 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
  127. Dongfang Zhao, Jian Yin, Ioan Raicu. "Improving the I/O Throughput for Data-Intensive Scientific Applications with Efficient Compression Mechanisms", IEEE/ACM Supercomputing 2013
  128. Benjamin Grimmer, Scott Krieder, Ioan Raicu. "Enabling Dynamic Memory Management Support for MTC on NVIDIA GPUs", EuroSys 2013
  129. Ke Wang, Ioan Raicu. "Paving the Road to Exascale with Many-Task Computing", Doctoral Showcase, IEEE/ACM Supercomputing/SC 2012
  130. Dongfang Zhao, Ioan Raicu. "Distributed File Systems for Exascale Computing", Doctoral Showcase, IEEE/ACM Supercomputing/SC 2012
  131. Iman Sadooghi, Ioan Raicu. "Towards Scalable and Efficient Scientific Cloud Computing", Doctoral Showcase, IEEE/ACM Supercomputing/SC 2012
  132. Scott Krieder, Ioan Raicu. "Towards the Support for Many-Task Computing on Many-Core Computing Platforms", Doctoral Showcase, IEEE/ACM Supercomputing/SC 2012
  133. Scott Krieder, Ben Grimmer, Ioan Raicu. "Early Experiences in running Many-Task Computing workloads on GPGPUs", XSEDE 2012
  134. 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
  135. Hui Jin, Xi Yang, Xian-He Sun, Ioan Raicu. "An Empirical Evaluation of MapReduce under Interruptions", Cloud Computing and its Applications (CCA) 2011
  136. 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
  137. 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
  138. 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
  139. 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
  140. Ioan Raicu, Catalin Dumitrescu, Ian Foster. "Dynamic Resource Provisioning in Grid Environments", TeraGrid Conference (TG07) 2007
  141. 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
  142. Ioan Raicu. "Routing Algorithms for Wireless Sensor Networks", Poster Presentation, Grace Hopper Celebration of Women in Computing (GHC02), 2002
-

---

## Papers under review (2)

143. Poornima Nookala, Kyle Chard, and Ioan Raicu. “X-OpenMP – eXtreme fine-grained tasking using lock-less work stealing”, under review at Future Generation Computer Systems (FGCS) 2023; latest review was Major Revision, plan to submit it 03/02/24
144. Alexandru Iulian Orhean, Anna Giannakou, Lavanya Ramakrishnan, Kyle Chard, Boris Glavic, Ioan Raicu. “SCIPIS: Scalable and Concurrent Persistent Indexing and Search in High-End Computing Systems”, under review at Journal of Parallel and Distributed Computing (JPDC) 2023; latest review was Minor Revision, resubmitted on 12/13/23
- 

---

## 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, Amazon, and Intel). I have received 16 awards totaling over \$4.1M in funds, \$776K in hardware donation, \$100K in credits on public clouds, and 8M hours on large-scale distributed resources.

### Pending Grants (2024 – 2027, \$1.25M in funds)

---

<b>NSF 2024 - 2027:</b>	<b>“XSearch: High-Performance Distributed Indexing and Search in Large-Scale Systems”</b>
Institution:	NSF CNS CSR Core Medium
Award:	\$1,248,123
Period:	07/2024 – 06/2027
Web:	<a href="http://datasys.cs.iit.edu/grants/xsearch/">http://datasys.cs.iit.edu/grants/xsearch/</a>
People:	Ioan Raicu (IIT/PI Lead Institution), Alexandru Orhean (DePaul/PI), Anna Giannakou (LBL/PI), Kyle Chard (UChicago/PI)
Summary:	This work will design, implement, and deploy a scalable distributed indexing system—XSearch, which is designed to support powerful search mechanisms across file-based data. XSearch will support file metadata as well as data content indexing and search. Scientific data comes in many flavors, from free-text data, to numerical data, to image data, to video data. XSearch will be integrated with existing parallel and distributed file systems to provide efficient search capabilities implicitly for a wide range of data types with minimal user effort. This work is split into four main thrusts: 1. Data Structures & Algorithms: Explore scalable data structures and algorithms for indexing/search; 2. Metadata & Interfaces: I. Metadata extraction from scientific datasets; II. Inferring metadata from data; III. Explore interfaces for search; 3. Distributed Search Infrastructure: Distributed XSearch and Integration into HPC Systems; 4. Science: Solve data search for real science in material science, environmental science & cosmology.

### Active Grants (2021 – 2025, \$904K in funds)

---

<b>NSF 2022 - 2025:</b>	<b>“REU Site: BigDataX: From Theory to Practice in Big Data Computing at eXtreme Scales”</b>
Institution:	NSF OAC REU
Award:	\$362,878 (total \$404,437)
Period:	07/2022 – 06/2025
Web:	<a href="http://datasys.cs.iit.edu/grants/BigDataX/">http://datasys.cs.iit.edu/grants/BigDataX/</a>
People:	Ioan Raicu (IIT/PI), Kyle Hale (IIT/Co-PI), Zhiling Lan (IIT), Stefan Muller (IIT), Kyle Chard (UChicago/PI), Katarzyna Keahey (UChicago)
Summary:	This award aims to continue the 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. The site funds 10 students for 10 weeks each summer to be immersed in research in big data computing at extreme scales.

**NSF 2021 - 2025:** **“Enabling Extremely Fine-grained Parallelism on Modern Many-core Architectures”**  
 Institution: NSF OAC Small  
 Award: \$333,684 (total \$500,000)  
 Period: 07/2021 – 06/2025  
 Web: <http://datasys.cs.iit.edu/grants/xtask/>  
 People: Ioan Raicu (IIT/PI), Kyle Chard (UChicago/PI)  
 Summary: This proposal seeks to explore novel data-structures and algorithms that allow for scalable runtime and execution models for fine-grained parallelism at sub-microsecond timescales. Preliminary work by the PIs at the language and runtime levels suggests a path to achieving this. The project objectives are: 1. Unifying runtime enabling task granularities measured in cycles: design, analysis, and implementation of building blocks for efficient fine-grained computing on diverse node hardware; 2. Evaluating performance of these building blocks in the context of real parallel systems and application kernels on a range of computer architectures; 3. Measuring performance and scalability impact of runtime on benchmark kernels and real applications; 4. Integrating this research with education programs from undergraduate to graduate levels through new course material on parallel computing.

---

### Completed Grants (2006 – 2022, \$3.1M in funds)

---

**NSF 2018 - 2022:** **“REU Site: BigDataX: From Theory to Practice in Big Data Computing at eXtreme Scales”**  
 Institution: NSF CNS REU  
 Award: \$333,106 (total \$370,000)  
 Period: 03/2018 – 02/2022  
 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 - 2021:** **“CRI: II-NEW: MYSTIC: prograMmable sYstems reSearch Testbed to explore a stack-wlde adaptive system fabriC”**  
 Institution: NSF CRI  
 Award: \$1,042,238  
 Period: 07/2017 - 06/2021  
 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.

**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 REU  
 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 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 – 2023, \$82K in funds, \$776K hardware donation, \$100K in credits on public clouds)**

---

- Intel 2021 - 2023:** **“Exploring Concurrent Data Structures for Parallel Computing with Intel Hardware and Software”**  
 Institution: Intel  
 Award: \$70,000  
 Period: 12/2021 – 02/2023  
 Web: <http://datasys.cs.iit.edu/grants/xtask-intel/>  
 People: Ioan Raicu (IIT/PI)  
 Summary: This project seeks to address how to extend a programming language, runtime, and execution model to support parallelism at the node level for modern and future nodes. This work aims to supporting undergraduate research to explore fine-grained parallelism using the Intel software stack – OneAPI and SYCL.
- UChicago 2021:** **“Chi@IIT”**  
 Institution: University of Chicago  
 Award: Donated 3 racks from Chameleon testbed containing 84 compute nodes (24-core Intel Xeon CPUs, 128GB RAM, 10GbE), 8 storage nodes (24-core Intel Xeon CPUs, 128GB RAM, 24TB HDD, 10GbE), 2 control nodes, 3 40GbE switches; these nodes are valued at \$500,000 at the time of donation  
 Period: 07/2021 - 12/2023  
 People: Ioan Raicu (IIT/PI)
- 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 at the time of donation  
 Period: 10/2014 - 12/2018  
 People: Ioan Raicu (IIT/PI)



<b>Amazon 2011-2016:</b>	<b>“Distributed Systems on the Amazon Cloud”</b>
Institution:	Amazon AWS
Award:	\$60,700 in AWS Credits
Period:	01/2011 – 05/2016
<b>Microsoft 2014-2015:</b>	<b>“Implicitly-Parallel Functional Dataflow for Productive Cloud Programming on Windows Azure”</b>
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)
<b>NVIDIA 2013 - 2014:</b>	<b>“NVIDIA Teaching Center Program Proposal”</b>
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 – 2024, 8.2M CPU-hours on large-scale distributed resources)

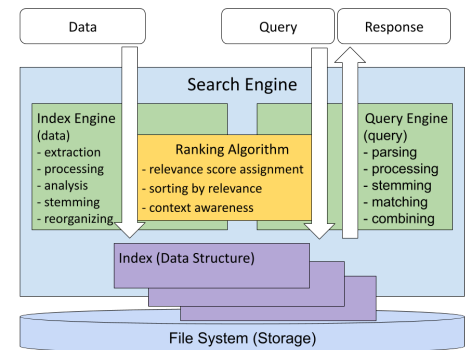
---

<b>DOE 2017 - 2024:</b>	<b>“Resource Management in Distributed Systems”</b>
Institution:	DOE ANL
Award:	1.5M hours on the Chameleon Resources
Period:	03/2017 - 06/2024
<b>XSEDE 2017 - 2018:</b>	<b>“Resource Management in Distributed Systems”</b>
Institution:	NSF XSEDE
Award:	200K hours on XSEDE
Period:	06/2017 - 05/2018
<b>DOE 2014 – 2018:</b>	<b>“Student Cluster Competition”</b>
Institution:	DOE ANL ALCF
Award:	300K hours on the ALCF Resources
Period:	03/2014 - 12/2018
<b>XSEDE 2013 - 2014:</b>	<b>“Many-Task Computing with Many-Core Accelerators on XSEDE”</b>
Institution:	NSF XSEDE
Award:	200K hours on XSEDE
Period:	01/2013 - 08/2014
<b>GLCPC 2013 - 2014:</b>	<b>“Implicitly-parallel functional dataflow for productive hybrid programming on Blue Waters”</b>
Institution:	Great Lakes Consortium for Petascale Computation (GLCPC)
Award:	6M hours on the Blue Waters Supercomputer
Period:	03/2013 - 08/2014
<b>NICS 2013 - 2014:</b>	<b>“Many-Task Computing with Many-Core Accelerators on Beacon”</b>
Institution:	National Institute for Computational Sciences (NICS)
Award:	320K hours on the Beacon system
Period:	06/2013 - 08/2014
<b>DOE 2011 - 2013:</b>	<b>“FusionFS: Distributed File Systems for Exascale Computing”</b>
Institution:	DOE ANL ALCF
Award:	450K hours on the IBM BlueGene/P
Period:	03/2011 - 10/2013
<b>DOE 2011 - 2012:</b>	<b>“HPCCloud: Exploring HPC Fault Tolerance in the Cloud”</b>
Institution:	DOE ANL ALCF
Award:	100K hours on the Magellan Cloud
Period:	03/2011 - 02/2012

## Active Projects & Synergistic Activities


### XSearch: Distributed Indexing and Search in Large-Scale Storage Systems

Rapid advances in digital sensors, networks, storage, and computation coupled with decreasing costs is leading to the creation of huge collections of data. Increasing data volumes, particularly in science and engineering, has resulted in the widespread adoption of parallel and distributed file systems for storing and accessing data efficiently. However, as file system sizes and the amount of data “owned” by users grows, it is increasingly difficult to discover and locate data amongst the petabytes of data. While much research effort has focused on the methods to efficiently store and process data, there has been relatively little focus on methods to efficiently explore, index, and search data using the same high-performance storage and compute systems. We proposed to design a new indexing framework from scratch to address the problem of efficiently indexing data in high-end systems, characterized by many-core architectures, with multiple NUMA nodes and multiple PCIe NVMe storage devices. We designed a single-node framework that can be used as a building block for implementing high-performance distributed indexed search engines, where the software architecture of the framework is scalable by design. The indexing pipeline is exposed and allows easy modification and tuning, allowing the proposed system to saturate storage, memory and compute resources on a variety of hardware. We have designed a novel tokenizer and inverted index to achieve high performance improvement both in terms of indexing and in terms of search latency. We evaluated the proposed indexing framework with configurations up to 192-cores, 768GiB of RAM, and up to 16 NVMe drives, and achieved performance improvements up to 19x better indexing while delivering up to 280X lower search latency when compared to Apache Lucene. We are exploring methods for distributing indexing and search to scale to some of the largest HPC storage systems available. Specifically, we have been investigating integration of the indexing/search framework into parallel and distributed storage systems (called XSearch) to enable automatic metadata and data indexing and search. We are investigating the support of different types of data, from unstructured text, to structured numerical data (e.g. 2D matrices), to time-series data.



### XTask: eXTreme fine-grAined concurrent taSK invocation runtime

Processors with 100s of threads of execution and GPUs with 1000s of cores are among the state-of-the-art in high-end computing systems. This transition to many-core computing has required the community to develop new algorithms to overcome significant latency bottlenecks through massive concurrency. Implementing efficient parallel runtimes that can scale up to hundreds of threads with extremely fine-grained tasks (less than 100µs) remains a challenge. We proposed a novel lockless concurrent queueing system that can scale up to hundreds of threads in a system called XQueue. We integrate XQueue into LLVM OpenMP (called X-OpenMP), a library for lightweight tasking on modern many-core systems with hundreds of cores. We show that it is possible to implement a parallel execution model using lock-less techniques for



enabling applications to strongly scale on many-core architectures. We extend our work to include dynamic load balancing using work stealing to efficiently distributes the load across worker threads. We implement a lock-less algorithm for work stealing and evaluate the performance using micro and macro benchmarks. We compare the performance of X-OpenMP with native LLVM OpenMP and GNU OpenMP implementations using task-based linear algebra routines from PLASMA numerical library, Strassen's matrix multiplication from the BOTS Benchmark Suite, and the Unbalanced Tree Search benchmark. Applications parallelized using OpenMP can run without modification by simply linking against the X-OpenMP library. X-OpenMP achieves up to 40X speedup compared to GNU OpenMP and up to 6X speedup compared to the native LLVM OpenMP implementations on fine-grained parallel workloads.

### XStore: Efficient Time Series Storage System

In recent years, we have seen an un-precedented growth of data in our daily lives ranging from health data from an Apple Watch, financial stock price data, volatile crypto-currency data, to diagnostic data of nuclear/rocket simulations. The increase in high-precision, high-sample-rate timeseries data is a challenge to existing database technologies. We have developed a novel technique that utilizes sparse-file support to achieve  $O(1)$  time complexity in create, read, update, and delete (CRUD) operations while supporting time granularity down to 1 millisecond. We designed and implemented XStore to be lightweight and offer high performance without the need to maintain an index of the timeseries data. We are conducting a detailed evaluation between XStore and existing best-of-breed systems such as MongoDB and InfluxDB using real-world cryptocurrency data across dozens of coins, years of data, with millisecond granularity, totaling over a trillion datapoints. We found significant speedups in both read and write performance compared to existing systems by as much as an order of magnitude.

## [MEMO: High-Throughput Blockchain using Memoization](#)

Blockchain technologies are behind the success of digital currencies as they are difficult to change, hack, or cheat the system while operating in a trustless and decentralized environment. However, blockchains are orders of magnitude slower and less energy efficient than existing centralized payment processing approaches such as credit cards run by VISA. Proof-of-work based blockchains (e.g. Bitcoin) have a disproportionate carbon footprint and slow transaction speeds (e.g. 7 tps). Our new digital currency MEMO is a high-throughput blockchain using memorization (formerly known as GIGI: Green Lightning Coin). It achieves over one thousand tps with confirmations in seconds through larger blocks with smaller block times, and has the potential to achieve over a million transactions per second through sharding. MEMO achieves energy efficiency on par with centralized solutions as well as proof of stake blockchains using an improved Proof of Space consensus using the novel CryptoMemoiz algorithm to implement cryptographic puzzle memoization. CryptoMemoiz leverages XSearch indexing to reduce I/O, large memory, and significant compute resource requirements. MEMO was designed to be lightweight enough to operate on small nodes such as Raspberry Pis, as well as making use of high-end servers with 100+ cores, hundreds of gigabytes of memory, and multiple NVMe storage devices. This work has resulted in winning a 1st place in the inaugural Grainger Computing Innovation Prize that came with a \$15K winning prize.



## [Mystic: Programmable Systems Research Testbed to Explore a Stack-Wide Adaptive System fabric](#)

I served as the lead PI on a \$1M+ 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. The testbed includes diverse architectures from different manufactures and various types of computing devices. The system houses many of the largest many-core platforms available, including NVIDIA accelerators (e.g. Tesla V100 GPUs w/ 5120-cores), Intel processors (Xeon Phi w/ 288HT and Xeon Scalable w/ 384HT), AMD Epyc (w/ 128HT), IBM Power9 [w/ 160HT], ARM ThunderX2 [w/ 224HT], and FPGAs (Intel 510T Altera Arria 10). The network interconnect is composed of 100GbE in a Fat-Tree topology. The storage system houses a 388TB Ceph-based distributed storage system. The software stack is composed in multiple layers: (1) Overcloud: OpenStack (PerconaDB, Keystone, Horizon, Cinder, Glance, Nova, Neutron, RabbitMQ, Swift, Ironic, Blazar, LXD); (2) Undercloud: MaaS, Juju, LDAP, Ceph Cluster, User Login, Monitoring Service (Grafana+InfluxDB); (3) Protocloud: CentOS 7.6, Gateway, Firewall, IP Failover, DHCP, DNS, HW/SW RAID, KVM Hypervisor, SaltStack.



## [Chameleon: A configurable experimental environment for large-scale edge to cloud research](#)

Chameleon is a large-scale, deeply reconfigurable experimental platform built to support Computer Sciences systems research. Community projects range from systems research developing new operating systems, virtualization methods, performance variability studies, and power management research to projects in software defined networking, artificial intelligence, and resource management. Chameleon is an open testbed which has supported 3,000+ users working on 500+ projects. The project has been online since 2015 with about \$30M in funding support. In 2021, I received a donation from University of Chicago for 3 racks of systems that were decommissioned from the Chameleon testbed (originally acquired in 2015). In 2022, I helped bring the 84 compute nodes and 8 storage nodes online as a new satellite site Chi@IIT, joining three additional satellite sites from Northwestern University, National Center for Atmospheric Research, and University of Illinois at Chicago. The main two sites with Chameleon hardware are at University of Chicago and the University of Austin Texas.

## [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/high school 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 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. We (Michael Papka UIC, Alexandru Orhean DePaul, and I) are gearing up for another run in the 2024 SCC competition with the help of Argonne, Intel, and Dell as co-sponsors and a joint IIT/UIC/DePaul team.






### **[Upsilon Pi Epsilon – Executive Council](#)**

It is the express purpose of Upsilon Pi Epsilon (UPE) to promote the computing and information disciplines and to encourage their contribution to the enhancement of knowledge. The UPE was first established in 1967, and now consists of over 300 chapters in various colleges and universities in North America and overseas. UPE is the first and only, existing international honor society in the Computing and Information disciplines. The mission of UPE is to recognize academic excellence at both the undergraduate and graduate levels in the Computing and Information Disciplines. I am a member of the Executive Council overseeing new UPE chapter establishment, scholarship review and awards, as well as direct future directions of the UPE international organization.



### **[Upsilon Pi Epsilon – IIT Chapter](#)**

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 2023, we have about 50 active members, and over 80 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. After a COVID-induced pause, the IIT chapter restarted with an initiation ceremony in May 2023. We are planning a full slate of activities in the 2023-2024 academic year.



### **[NSF Research Experiences for Undergraduates \(REU\) BigDataX Program](#)**

I have been leading undergraduate research in big data computing at IIT with over \$1-million funding over a 10-year period [2015 – 2025], covering 84 domestic undergraduate students. 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 department at IIT or 2) the Globus Labs group in the CS department 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. The latest award has six 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.



---

## Teaching

---

### **Assistant/Associate Professor** Illinois Institute of Technology, Department of Computer Science

---

08/2010 – Present

I taught a total of 35 courses to 1649 students since 2010. I taught core courses at both undergraduate and graduate level. I also designed new courses and taught them multiple times each, denoted by an \* below.

#### **Undergraduate Courses**

CS351: **Systems Programming** [179 students: F20, S22]

\*CS495/CS451: **Introduction to Parallel and Distributed Computing** [47 students: F12, S14]

\*CS495: **Cluster Computing** [7 students: F18]

\*COMP495/SCI497: **Big Data Computing** [11 students: Sum19, Sum21]

#### **Graduate Courses**

CS550: **Advanced Operating Systems** [285 students: S11, F15, F19, F23]

\*CS553: **Cloud Computing** [939 students: S12, S13, F14, S16, F17, S18, S19, S20, S21, F22, S24]

\*CS595/CS554: **Data-Intensive Computing** [304 students: F10, F11, F13, S15, F17, F18, Sum20, Sum21, F21, S23]

\*CS595: **Special Topics in Distributed Systems** [17 students: Sum16, Sum17]

CS695: **Doctoral Seminar** [27 students: S11, S12]

### **Adjunct Assistant Professor** Northwestern University, Dept. of Electrical Engineering and Computer Science

---

01/2010 – 06/2010

EECS 211: Fundamentals of Computer Programming [40 students, S10]

\*EECS 495: Hot Topics in Distributed Systems: Data-Intensive Computing [10 students, W10]

### **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

Introduction to Computer Science 1

Introduction to Computer Science 2

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++ [50 students, Sum02]

### **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 Research Advising

---

### Current PhD Students (1) – Research Advisor at Illinois Institute of Technology

---

Lan Nguyen (3<sup>rd</sup> year), Expected graduation 05/26, “XStore: Efficient Time Series Storage System”

### Current MS Students (6) – Research Advisor at Illinois Institute of Technology

---

Adarsh Agrawal (2<sup>nd</sup> year), Expected graduation 2024, “MEMO: High-Throughput Blockchain using Memoization”

Chaudhari Shlok Mohan (2<sup>nd</sup> year), Expected graduation 2024, “MEMO: High-Throughput Blockchain using Memoization”

Abhishek Sharma (2<sup>nd</sup> year), Expected graduation 2024, “MEMO: High-Throughput Blockchain using Memoization” & “XStore: Efficient Time Series Storage System”

Thakur Bhagvan Sudarshan (2<sup>nd</sup> year), Expected graduation 2024, “System Administration on Mystic & Chameleon”

Pratik Bejagamwar (2<sup>nd</sup> year), Expected graduation 2024, “System Administration on Mystic & Chameleon”

Batkishig Dulamsurankhor (2<sup>nd</sup> year), Expected graduation 2025, “XStore: Efficient Time Series Storage System”

### Current Undergraduate Students (2) – Research Advisor at Illinois Institute of Technology

---

Nhan Hoang (2<sup>nd</sup> year), Expected graduation 2025, “MEMO: High-Throughput Blockchain using Memoization”

Jamison Kerney (4<sup>th</sup> year), Expected graduation 2023, “Exploring Fine-grained Parallelism in PARSL”

### Graduated PhD Students (6) – Research Advisor at Illinois Institute of Technology

---

Alexandru I. Orhean, IIT/CS, PhD 08/2023, “Scalable Indexing and Search in High-End Computing Systems”

Assistant Professor @ DePaul University

Poornima Nookala, IIT/CS, PhD 12/2022, “Extreme Fine-grained Parallelism on Modern Many-Core Architectures”

Research Engineer @ Intel

Iman Sadooghi, IIT/CS, PhD 12/2016, “Scalable Resource Management in Cloud Computing”

Software Engineer @ Box

Tonglin Li, IIT/CS, PhD 12/2015, “Distributed NoSQL Storage for Extreme-Scale System Services in Supercomputers and Clouds”

Software Engineer @ Amazon Web Services (AWS)

Ke Wang, IIT/CS, PhD 08/2015, “Towards Next Generation Resource Management at Extreme-Scales”

Senior Software Engineer @ Google

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 (2) – Research Advisor at Illinois Institute of Technology

---

Itua E. Ijagbone, IIT/CS, MS 05/2016, “Scalable Indexing and Searching on Distributed File Systems” – Braintree

Anupam Rajendran, IIT/CS, MS 05/2013, “MATRIX: Many-Task Computing Execution Fabric for Extreme Scales” – Meta

### Alumni with a thesis (12) – Committee Member

---

Xin Wang, IIT/CS, PhD 12/2023, “Heterogenous Workloads Study Towards Large-Scale Interconnect Network Simulation”, Advisor Zhiling Lan

Yao Kang, IIT/CS, PhD 11/2022, “Workload Interference Analysis and Mitigation on Dragonfly Class Networks”, Advisor Zhiling Lan

Ahmad Alnafoosi, DePaul/IS, PhD 06/2022, “Empirical Assessment of Big Data Technology Adoption Factors for Organizations with Data Storage Systems”, Advisor Theresa Steinbach

Xu Yang, IIT/CS, PhD 04/2017, “Cooperative Batch Scheduling for HPC Systems”, Advisor Zhiling Lan

Eduardo Berrocal, IIT/CS, PhD 04/2017, “Improving Fault Tolerance for Extreme Scale Systems”, Primary Advisor: Dr. Zhiling Lan

Sean Wallace, IIT/CS, PhD 04/2017, “Power Profiling, Analysis, Learning, and Management for High-Performance Computing”, Advisor Zhiling Lan

Zhou Zhou, IIT/CS, PhD 12/2015, “Multi-Dimensional Batch Scheduling Framework for High-End Supercomputers”, Advisor Z. Lan

Jingjin Wu, IIT/CS, PhD 08/2013, “Performance Analysis and Optimization of Large-Scale Scientific Applications”, Advisor Zhiling Lan

Wei Tang, IIT/CS, PhD 08/2012, “An Integrated Resource Management and Scheduling Framework for Production Supercomputers”, Advisor Zhiling Lan

Hui Jin, IIT/CS, PhD 05/2012, “System Support for Resilience in Large-Scale Parallel Systems”, Advisor Xian-He Sun

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

Yuchi Tsao, IIT/ECE, PhD 12/2011; “Architecture-Level Hardware Optimization for Digital Signal Processing”, Advisor Kyu-won Choi

---

## Journal Editor

**TCC:** IEEE Transactions on Cloud Computing, 2013 – 2021 – *Associate Editor*  
**Cluster:** Springer Cluster Computing Journal, 2013 – 2020 – *Editorial Board*  
**JoCCASA:** Springer Journal of Cloud Computing: Advances, Systems and Applications, 2011 – 2020 – *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, 2015 – *Guest Editor*

---

## Journal Reviewer

**TPDS:** IEEE Transactions on Parallel and Distributed Systems, 2009 - 2012, 2014 – 2017, 2021  
**TCC:** IEEE Transactions on Cloud Computing, 2014 – 2017, 2020  
**TC:** IEEE Transactions on Computers, 2006, 2009, 2010, 2015, 2019  
**CCPE:** Journal of Concurrency and Computation: Practice and Experience, 2006, 2009, 2010, 2014, 2015, 2018

---

## Chair / Organizer

**GCASR:** Greater Chicago Area System Research Workshop, 2012 – 2023 – *General Chair, Steering Committee*  
**REUPI:** NSF CISE REU Site PI Virtual Meeting, 2021 – *General Co-Chair*  
**SC:** IEEE/ACM Int. Conf. on High Performance Computing, Networking, Storage, and Analysis, 2018 – *Cloud/Grids Track Chair*  
**HPDC:** ACM Symp. on High Performance Distributed Computing, 2010, 2011 – 2015, 2016, 2020 – *Stud. Activities Chair, Workshop Chair, Publicity Chair*  
**ScienceCloud:** ACM Workshop on Scientific Cloud Computing (@HPDC), 2012, 2015 – 2017 – *Steering Committee*  
**CCGrid:** Cluster, Cloud, and Grid Computing Conference, 2012, 2013, 2014, 2017 – *Local Chair, Publicity Chair*  
**IPDPS:** IEEE International Parallel & Distributed Processing Symposium, 2016 – *Local Chair*  
**BDC:** IEEE/ACM International Symposium on Big Data Computing (BDC), 2014, 2015 – *Program Committee Chair*  
**DataCloud:** ACM/IEEE Int. Workshop on Data-Intensive Computing in the Clouds (@SC), 2013 – 2015 – *Steering Committee*  
**Cluster:** IEEE Cluster, 2015 – *Publicity Chair*  
**MTAGS:** ACM/IEEE Workshop on Many-Task Computing on Clouds, Grids, & Supercomputers (@SC), 2008 – 2015, *General Chair*

---

## Conference Program Committee

**SC:** IEEE/ACM Int. Conf. on High Performance Computing, Networking, Storage, & Analysis, 2012 - 2023  
**IPDPS:** IEEE International Parallel & Distributed Processing Symposium, 2014, 2019, 2023  
**CLOUD:** IEEE International Conference on Cloud Computing, 2011, 2012, 2014, 2015, 2017, 2018, 2023  
**HPDC:** ACM Int. Symposium on High Performance Distributed Computing, 2010 - 2014, 2016, 2019 - 2021  
**eScience:** IEEE e-Science Conference, 2010 - 2017  
**ICA3PP:** International Conference on Algorithms and Architectures for Parallel Processing, 2017  
**CCGrid:** Cluster, Cloud, and Grid Computing Conference, 2012 - 2014, 2016  
**Cluster:** IEEE Cluster, 2014 - 2016  
**CloudCom:** IEEE International Conference on Cloud Computing Technology and Science, 2010, 2014 - 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

---

## University Service

- **Master of Data Science, Associate Director** (Fall 2021 – Present)
- **Grainger Computing Innovation Prize, Chair** (Spring 2022 – Present), Member (Spring 2021 – Fall 2021)
- **Upsilon Pi Epsilon (UPE) Honor Society, Faculty Mentor** (Fall 2018 – Present)
- **UFC Academic Computing Committee, Representative** (Fall 2020 – Present)
- **AUCOPT Committee, Member** (Fall 2017 – Present)
- **Graduate Studies Committee, Member** (Fall 2011 – Spring 2013, Fall 2017 – Present)
- **Graduate Admissions, Member** (Spring 2011, Fall 2014 – Spring 2016, Fall 2021 – Present)
- **Lab Resources Group, Member** (Fall 2021 – Present)
- **High Performance Steering Committee, Member** (Summer 2022 – Present)
- **Faculty Search Committee, Member** (Fall 2010 – Spring 2016, 2023 - Present)