

CURRICULUM VITAE

SHRIDEEP PALLICKARA

ASSOCIATE PROFESSOR, MONFORT PROFESSOR

Department of Computer Science, Colorado State University

Office: (970) 492 4209

Home Page: <http://www.cs.colostate.edu/~shrideep/>

Email: shrideep@cs.colostate.edu

1 RESEARCH INTERESTS

Broadly in cloud computing and analytics with a focus on using machine learning, probabilistic, and statistical techniques to address scaling, autonomy, and tractability issues. These efforts target processing, routing, storing, and mining data streams generated by networked sensors, observational equipment, and programs.

2 EDUCATION

June 01	<i>Syracuse University</i> , Syracuse, New York. Ph.D., Computer Engineering.
December 98	<i>Syracuse University</i> , Syracuse, New York. M.S., Computer Engineering.
May 94	<i>Bombay University</i> , Bombay, India. Bachelor of Engineering, Electronics & Telecommunications.

3 AWARDS & HONORS

* I am the first recipient of both the Monfort Professorship and the Board of Governors' Award at CSU.

Honors

- Monfort Professor [2016-2018]
The Monfort Professorship is among the highest honors that Colorado State University bestows on faculty. The award is given primarily for research accomplishments. Two faculty are named Monfort Professors every year.

Teaching Awards

- Board of Governors' Excellence in Undergraduate Teaching Award. 2017.
The Board of Governors' Award is the University's most prestigious teaching award and is awarded to one faculty every year.
- *Online Innovative Educator Award (The OLIE)*. 2018.
Colorado State University Online Programs
- *Faculty Excellence in Undergraduate Teaching Award*. 2014.
College Of Natural Sciences, Colorado State University.
- *Effective and Innovative Teaching Award*, Dept. of Computer Science, Colorado State University. 2012.

Research Awards

- Shortlisted for the *ACM-DEBS Grand Challenge Award* with Amila Suriarachchi (graduate advisee). 2015.
- U.S. National Science Foundation CAREER Award. 2013.
CISE/CNS/Computer Systems Research
- Best Paper Award. *IEEE/ACM Conference on Utility and Cloud Computing* [27% acceptance rate] with Matthew Malensek (graduate advisee) and Sangmi Lee Pallickara. 2012.
- Best Student Paper Award at the *IEEE International Conference on Cloud Computing Technology and*

Science [24.9% acceptance rate] with Kate Ericson (graduate advisee) and Charles Anderson. 2010.

- One of Six Best Papers at ACM/IEEE GRID [18.8% acceptance rate] with Geoffrey Fox and Harshawardhan Gadgil. 2005.

4 PROFESSIONAL EXPERIENCE

July 2016 – 18	Monfort Professor	Colorado State University
July 13 –	Associate Professor	Department of Computer Science, Colorado State University
Aug 09 – June 13	Assistant Professor	Department of Computer Science, Colorado State University
Oct 02 – Aug 09	Research Associate	Community Grids Lab, Indiana University.
Oct 01 – Oct 02	Postdoctoral Fellow	Community Grids Lab, Indiana University.
Jan 99 – Sep 99	Intern	IBM T. J. Watson Research Center, New York.
May 97 – Sep 97	Intern	IBM T. J. Watson Research Center, New York.
May 96 – Jun 01	Research Assistant	Northeast Parallel Architectures Center, Department of Electrical Engineering and Computer Science, Syracuse University
Jan 96 – May 96	Teaching Assistant	Department of Electrical Engineering and Computer Science, Syracuse University
Jun 94 – July 95	Software Analyst	Mahindra- British Telecom Plc, Bombay India.

5 RESEARCH FUNDING

My research has been funded by agencies in the United States and the United Kingdom. U.S. federal agencies that I have been funded by include the National Science Foundation, the Department of Homeland Security, the Environmental Protection Agency, and the Department of Agriculture. My funding from the U.K. was from the Open Middleware Infrastructure Institute that was funded by the Engineering and Physical Sciences Research Council's (EPSRC) e-Science program.

5.1 Funding from the United States

- [F18] *Monfort Professorship*. Shrideep Pallickara (**PI**). Monfort Family Foundation and Colorado State University. \$150,000 2016–.
- [F17] *Assessing Epidemiological and Economic Impacts of Countermeasures and Vaccination Strategies in Disease Outbreaks at the National Scale*. Shrideep Pallickara (**PI**). Co-PIs: Sangmi Lee Pallickara, Jay Breidt, Neil Harvey, Mike Sanderson, and Dustin Pendell. Science and Technology Directorate. Department of Homeland Security. \$2,044,978 9/2015–9/2019.
- [F16] *Advancing Math and Computer Science Education for Middle School Students*. Shrideep Pallickara (**PI**). National Science Foundation. \$40,000 6/2014 – 6/2018.
- [F15] *Autonomous Scaling of Multitier Soil Erosion Models in IaaS Clouds*. Shrideep Pallickara (**PI**). Amazon Web Services Education Research Grant. \$5000 3/2014–3/2015.
- [F14] *Big Data Analytics of Epidemic Outbreaks*. Shrideep Pallickara (**PI**). Co-PIs: Sangmi Lee Pallickara, Neil Harvey, and Jay Breidt. LONG RANGE PROGRAM, Science and Technology Directorate. Department of Homeland Security. \$1,176,860 9/2013–6/2017.

- [F13] *Accessible Nutrient Data, Analysis and Modeling Dashboard for the Nation*. Shrideep Pallickara (**Co-PI**), with Mazdak Arabi (PI) and Co-PIs: Jeff Arnold, Richard Alexander, Nathaniel Booth. Part of an EPA Center for Effective Abatement of Nutrients. US Environmental Protection Agency. \$256,936 09/2013–08/2017.
- [F12] *CAREER: Robust Processing of Data Streams in Real Time*. Shrideep Pallickara (**PI**). National Science Foundation. CISE/CNS/Computer Systems Research. NSF CNS-1253908. \$400,000 3/2013–2/2019.
- [F11] REU Supplement: *Reliable Processing Semantics for Fault Tolerant Stream Processing Computations*. Shrideep Pallickara (**PI**). National Science Foundation. \$16,000 7/2016–6/2017.
- [F10] REU Supplement: *Fault Tolerant Processing for Stateless and Stateful Stream Computations*. Shrideep Pallickara (**PI**). National Science Foundation. \$16,000 7/2015–6/2016.
- [F9] REU Supplement: *Predicting Stream Data Arrivals Using Linear and Nonlinear Methods*. Shrideep Pallickara (**PI**). National Science Foundation. \$16,000 7/2014–6/2015.
- [F8] REU Supplement: *Simulating Clinical Settings and Fast Data Quality Assertions*. Shrideep Pallickara (**PI**). National Science Foundation. CISE/CNS/Computer Systems Research. \$16,000 6/2013–5/2014.
- [F7] *Boosting Tracing and Assimilation of Disease Profiles in the Distributed Orchestration of Epidemiological Simulations*. Shrideep Pallickara (**PI**). Co-PIs: Wim Böhm, Neil Harvey, Sanjay Rajopadhye, and Aaron Reeves. LONG RANGE PROGRAM, Science and Technology Directorate. Department of Homeland Security. \$352,713 8/2011–8/2013.
- [F6] *Enabling Scalable and Fault Tolerant Regional Epidemiological Simulations in the Cloud*. Shrideep Pallickara (**PI**). Co-PIs: Wim Böhm, Neil Harvey, Sanjay Rajopadhye, and Aaron Reeves. LONG RANGE PROGRAM, Science and Technology Directorate. Department of Homeland Security. \$738,139 8/2010–8/2013.
- [F5] *Evaluating the Computational Soundness and Enabling the Coarse Grained Parallelization of the NAADSM Epidemiological Simulation Model*. Shrideep Pallickara (**PI**). Co-PIs: Wim Böhm and Sanjay Rajopadhye. From the USDA/ Animal and Plant Health Inspection Service (APHIS) through the Department of Homeland Security. \$86,952 8/2009–8/2010.
- [F4] *Collaborative Research: Development of Middleware/Software to allow Visualization and Analysis of Large and Complex 4-D Geoscience Data Sets*. Shrideep Pallickara (**PI**). National Science Foundation. EAR-0446610. \$199,020 6/2005–5/2009.
- [F3] *Collaborative Research: High-Performance Techniques, Designs and Implementation of Software Infrastructure for Change Detection and Mining*. Shrideep Pallickara: **Co-PI**, with Geoffrey Fox (PI) and Marlon Pierce (Co-PI). National Science Foundation. IIS-0536947. \$371,850 9/2004–9/2009.

5.2 Funding from the United Kingdom

- [F2] *FINS – Support for Web Services based publish/subscribe notifications within Grid Applications*. Shrideep Pallickara: **Co-PI**, with Geoffrey Fox (PI). Open Middleware Infrastructure Institute of the United Kingdom. \$307,247 10/2004–1/2006.
- [F1] *FIRMS – Support for Web Services based reliable messaging within Grid Applications*. Shrideep Pallickara: **Co-PI**, with Geoffrey Fox (PI). Open Middleware Infrastructure Institute of the United Kingdom. \$158,570 10/2004–10/2005.

6 PUBLICATIONS

These are listed in roughly reverse chronological order. In these listings, I denote my students with an * and my mentees with a †. Listings with the symbol § signify publications where I am the primary author of the research manuscript; in some cases, the authors are listed in alphabetical order. PDF versions of all papers are available for download at: <http://www.cs.colostate.edu/~shrideep/publications.html>

6.1 Journals

- [J41] Matthew Malensek*, Walid Budgaga*, Ryan Stern*, Shrideep Pallickara, and Sangmi Lee Pallickara. Trident: Distributed Storage, Analysis, and Exploration of Multidimensional Phenomena. (To appear) *IEEE Transactions on Big Data*. 2018.
- [J40] Wes Lloyd*, Shrideep Pallickara, Olaf David, Mazdak Arabi, Tyler Wible, Jeffrey Ditty, and Ken Rojas. Demystifying the Clouds: Harnessing Resource Utilization Models for Cost Effective Infrastructure Alternatives. *IEEE Transactions on Cloud Computing*. Vol. 5(4) pp 667-680. 2017. [Impact Factor: 3.77]
- [J39] Thilina Buddhika*, Ryan Stern*, Kira Lindburg*, Kathleen Ericson*, and Shrideep Pallickara. Online Scheduling and Interference Alleviation for Low-latency, High-throughput Processing of Data Streams. *IEEE Transactions on Parallel and Distributed Systems*. Vol. 28(12) pp 3553-3569. 2017. [Impact Factor: 4.181]
- [J38] Thilina Buddhika*, Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Synopsis: A Distributed Sketch over Voluminous Spatiotemporal Observational Streams. *IEEE Transactions on Knowledge and Data Engineering*. Vol. 29 (11) pp 2552-2566. 2017. [Impact Factor: 3.438]
- [J37] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Hermes: Federating Fog and Cloud Nodes to Support Query Evaluations in Continuous Sensing Environments. *IEEE Cloud Computing*. Vol. 4 (5) pp 54-62. 2017. [Impact Factor: 1.86]
- [J36] Walid Budgaga*, Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. A Framework for Scalable Real-Time Anomaly Detection over Voluminous, Geospatial Data Streams. *Concurrency and Computation: Practice & Experience*. Vol. 29(12) pp 1-16. John-Wiley. 2017. [Impact Factor: 1.133]
- [J35] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Fast, Ad Hoc Query Evaluations over Multidimensional Geospatial Datasets. *IEEE Transactions on Cloud Computing*. Vol. 5(1): 28-42. 2017. [Impact Factor: 3.77]
- [J34] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Autonomous Data Management and Federation to Support High-throughput Query Evaluations over Voluminous Datasets. *IEEE Cloud Computing*. Special Issue on Autonomic Clouds. Vol. 3 (3). pp 40-49. 2016. [Impact Factor: 1.86]
- [J33] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Analytic Queries over Geospatial Time-Series Data using Distributed Hash Tables. *IEEE Transactions on Knowledge and Data Engineering*. Vol 28(6) pp 1408-1422. 2016. [Impact Factor: 3.438]
- [J32] Walid Budgaga*, Matthew Malensek*, Sangmi Lee Pallickara, Neil Harvey, Jay Breidt, and Shrideep Pallickara. Predictive Analytics Using Statistical, Learning, and Ensemble Methods to Support Real-Time Exploration of Discrete Event Simulations. *Future Generation Computer Systems*. Elsevier. Volume 56, March, Pages 360–374. 2016. [Impact Factor: 3.997]
- [J31] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Minerva: Proactive Disk Scheduling for QoS in Multi-Tier, Multi-Tenant Cloud Environments. *IEEE Internet Computing*. Vol. 20 (3) pp 19-27. 2016. [Impact Factor: 1.521]

- [J30] Zhiquan Sui*, Matthew Malensek*, Neil Harvey, and Shrideep Pallickara. Autonomous Orchestration of Distributed Discrete Event Simulations in the Presence of Resource Uncertainty. *ACM Transactions on Autonomous and Adaptive Systems*. Vol. 10(3): 18. 2015. [Impact Factor: 2.48]
- [J29] Kathleen Ericson*, Shrideep Pallickara, and Charles Anderson. Failure-Resilient Real-Time Processing of Health Streams. Vol. 27(7) pp 1695–1717. *Concurrency and Computation: Practice & Experience*. John-Wiley. 2015. [Impact Factor: 1.133]
- [J28] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Evaluating Geospatial Geometry and Proximity Queries Using Distributed Hash Tables. *IEEE Computing in Science and Engineering (CISE)*. Special Issue on Extreme Data. Vol. 16(4) pp 53-60. 2014. [Impact Factor: 2.074]
- [J27] Zhiquan Sui*, Neil Harvey, and Shrideep Pallickara. On the Distributed Orchestration of Stochastic Discrete Event Simulations. *Concurrency and Computation: Practice & Experience*. John-Wiley. Vol. 26(11), pp 1889–1907. 2014. [Impact Factor: 1.133]
- [J26] Wes Lloyd*, Shrideep Pallickara, Olaf David, Jim Lyon, Mazdak Arabi and Ken Rojas. Performance Implications of Multi-Tier Application Deployments on Infrastructure-as-a-Service Clouds: Towards Performance Modeling. *Future Generation Computer Systems*. Vol. 29(5), pp 1254-1264. Elsevier. 2013. [Impact Factor: 3.997]
- [J25] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Exploiting Geospatial and Chronological Characteristics in Data Streams to Enable Efficient Storage and Retrievals. *Future Generation Computer Systems*. Vol. 29(4), pp 1049-1061. Elsevier. 2013. [Impact Factor: 3.997]
- [J24] Kathleen Ericson* and Shrideep Pallickara. On the Performance of High Dimensional Data Clustering and Classification Algorithms. *Future Generation Computer Systems*. Vol. 29(4), pp 1024-1034. Elsevier. 2013. [Impact Factor: 3.997]
- [J23] § Shrideep Pallickara and Geoffrey Fox. Enabling Hierarchical Dissemination of Streams in Content Distribution Networks. *Concurrency and Computation: Practice & Experience*. Vol. 24 (14), pp. 1594-1606. 2012. John-Wiley. [Impact Factor: 1.133]
- [J22] Kathleen Ericson* and Shrideep Pallickara. Adaptive Heterogeneous Language Support within a Cloud Runtime. *Future Generation Computer Systems*. Vol. 28(1), pp. 128-135. 2012. Elsevier. [Impact Factor: 3.997]
- [J21] Sangmi Lee Pallickara, Shrideep Pallickara, and Milija Zupanski. Enabling Efficient Data Subsetting over Large-scale Atmospheric Datasets. *Future Generation Computer Systems*. Vol. 28(1), pp. 112-118. 2012. Elsevier. [Impact Factor: 3.997]
- [J20] Geoffrey C. Fox, Mehmet S. Aktas, Galip Aydin, Harshawardhan Gadgil†, Shrideep Pallickara, Marlon E. Pierce, and Ahmet Sayar. Algorithms and the Grid. *Springer Journal of Computing and Visualization in Science (CVS)*. Vol 12 (3). pp 115-124. 2009. Springer Berlin-Heidelberg. [Impact Factor: 1.038]
- [J19] § Shrideep Pallickara, Geoffrey Fox and Harshawardhan Gadgil†. On the Secure Creation, Organization and Discovery of Topics in Distributed Publish/Subscribe Systems. *International Journal of High Performance Computing and Networking (IJHPCN)*. Vol. 5, No. 3, pp.156–167. 2008. [Impact Factor: 2.18]
- [J18] Geoffrey Fox, Galip Aydin, Harshawardhan Gadgil†, Shrideep Pallickara, Marlon Pierce, and Wenjun Wu. Management of Real-Time Streaming Data Grid Services. *Concurrency and Computation: Practice & Experience*. Volume 19(7) pp 983-998, 2007. [Impact Factor: 1.133]
- [J17] G. Erlebacher, D. Yuen, Z. Lu, E. Bollig, M. Pierce and Shrideep Pallickara. A Grid Framework for Visualization Services in the Earth Sciences. *Journal of Pure and Applied Geophysics*. Volume 163, Numbers 11-12. pp 2467-2483. Birkhäuser Verlag. 2006. [Impact Factor: 1.591]

- [J16] § Geoffrey Fox, Shrideep Pallickara, Marlon Pierce, Harshawardhan Gadgil[†]. Building Messaging Substrates for Web and Grid Applications. *Philosophical Transactions of the Royal Society: Mathematical, Physical and Engineering Sciences*. Volume 363, Number 1833, pp 1757-1773. Aug, 2005. [Impact Factor: 2.97]
- [J15] § Geoffrey Fox and Shrideep Pallickara. Deploying the NaradaBrokering Substrate in Aiding Efficient Web & Grid Service Interactions. *Proceedings of the IEEE*. Vol 93, No 3. pp 564-577. March 2005. [Impact Factor: 9.237]
- [J14] § Shrideep Pallickara, Geoffrey Fox, Ahmet Uyar[†], Hongbin Liu, Xi Rao[†], David Walker and Beytullah Yildiz[†]. Performance of a Possible Grid Message Infrastructure. *Concurrency and Computation: Practice & Experience*. Volume 17, Number 2-4. pp 193-214. [Impact Factor: 1.133]
- [J13] Geoffrey Fox, Sang Lim, Shrideep Pallickara and Marlon Pierce. Message-Based Cellular Peer-to-Peer Grids: Foundations for Secure Federation and Autonomic Services. *Future Generation Computer Systems*. Volume 21, Issue 3, pp 401-415. 2005. Elsevier. [Impact Factor: 3.997]
- [J12] § Geoffrey Fox, Shrideep Pallickara and Xi Rao[†]. Towards Enabling Peer-to-Peer Grids. *Concurrency and Computation: Practice & Experience*. Volume 17, Issue 7-8, Pages 1109-1131. 2005. [Impact Factor: 1.133]
- [J11] Mehmet Aktas, Galip Aydin, Andrea Donnellan, Geoffrey Fox, Robert Granat, Lisa Grant, Greg Lyzenga, Dennis McLeod, Shrideep Pallickara, Jay Parker, Marlon Pierce, John Rundle, Ahmet Sayar, and Terry Tullis. iSERVO: Implementing the International Solid Earth Research Virtual Observatory by Integrating Computational Grid and Geographical Information Web Services. Special Issue of the *Journal of Pure and Applied Geophysics*. Vol. 163(11-12), pp. 2281-2296. Birkhäuser Verlag AG. [Impact Factor: 1.591].
- [J10] § Shrideep Pallickara and Geoffrey Fox. Efficient Matching Of Events in Distributed Middleware Systems. *Journal Of Digital Information Management*. Volume 2(2). pp 79-87. 2004. [Impact Factor: 0.22]
- [J9] Minjun Wang, Geoffrey Fox and Shrideep Pallickara. Demonstrations of Collaborative Web Services and Peer-to-Peer Grids. *Journal Of Digital Information Management*. Volume 2(2). pp 93-97. 2004. [Impact Factor: 0.22]
- [J8] § Geoffrey Fox, Shrideep Pallickara, Marlon Pierce and David Walker. Towards Dependable Grid and Web Services. *ACM Ubiquity*. Volume 4 Issue 25. August 2003.
- [J7] Geoffrey Fox, Wenjun Wu, Ahmet Uyar[†], Hasan Bulut and Shrideep Pallickara. Global Multimedia Collaboration System. *Concurrency and Computation: Practice & Experience*. Volume 16(5). pp 441-447. 2004. [Impact Factor: 1.133]
- [J6] § Shrideep Pallickara and Geoffrey Fox. Efficient Support for Sophisticated Interactions between Entities in Distributed Brokering Systems. *ACM Ubiquity* Volume 4 Issue 16. June 2003.
- [J5] § Geoffrey Fox and Shrideep Pallickara. Support for Peer-to-Peer Interactions in Web Brokering Systems. *ACM Ubiquity*. Volume 3 Issue 15. May 2002.
- [J4] § Geoffrey Fox and Shrideep Pallickara. An Event Service to Support Grid Computational Environments. *Concurrency and Computation: Practice & Experience*. Special Issue on Grid Computing Environments. Volume 14(13-15) pp 1097-1129. 2002. [Impact Factor: 1.133]
- [J3] Geoffrey Fox, Sung Hoon Ko, Marlon E. Pierce, Ozgur Balsoy, Jake Kim, Sangmi Lee, Kang-Seok Kim, Sangyoon Oh, Xi Rao, Mustafa Varank, Hasan Bulut, Gurhan Gunduz, Xiaohong Qiu, Shrideep Pallickara, Ahmet Uyar, Choon-Han Youn. Grid Services for Earthquake Science. *Concurrency and Computation: Practice & Experience*. ACES Computational Environments for Earthquake Science Special Issue. Volume 14(6-7) pp 371-393. 2002. [Impact Factor: 1.133]

- [J2] Geoffrey Fox and Shrideep Pallickara. An Approach to High Performance Distributed Web Brokering *ACM Ubiquity* Volume 2 Issue 38. November 2001.
- [J1] G. Fox, W. Furmanski, B. Natarajan, H. T. Ozdemir, Z. Odcikin Ozdemir, S. Pallickara and T. Pulikal. Integrating Web, Desktop, Enterprise and Military Simulation Technologies To Enable World-Wide Scalable Televirtual (TVR) Environments. *Information & Security-An International Journal*, Volume 3-1999 : Modeling and Simulation. ISSN 1311-1493.

6.2 Book Chapters

- [B9] Zhiquan Sui* and Shrideep Pallickara. A Survey Of Load Balancing Techniques for Data Intensive Computing. *Handbook of Data Intensive Computing*. Chapter 6, pp 157-168. Springer. 2011. ISBN 978-1-4614-1415-5.
- [B8] Sangmi Lee Pallickara, Matthew Malensek*, and Shrideep Pallickara. On the Processing of Extreme Scale Datasets in the Geosciences. *Handbook of Data Intensive Computing*. Chapter 20, pp 521-538. Springer. 2011. ISBN 978-1-4614-1415-5.
- [B7] Kathleen Ericson* and Shrideep Pallickara. A Survey Of Storage and Fault Tolerance Strategies Used in Cloud Computing. Chapter 6: pp 137-158. *Handbook of Cloud Computing*. Springer. 2010. ISBN: 978-1-4419-6523-3.
- [B6] Sangmi Lee Pallickara, Shrideep Pallickara and Marlon Pierce. Scientific Data Management in the Cloud: A Survey of Technologies, Approaches and Challenges. Chapter 22: 517-534. *Handbook of Cloud Computing*. Springer. 2010. ISBN: 978-1-4419-6523-3.
- [B5] Shrideep Pallickara, Jaliya Ekanayake†, and Geoffrey Fox. Granules: A Lightweight Runtime for Scalable Computing With Support for Map-Reduce. Chapter 9: pp 201-226. *Handbook on Cloud Computing and Software Services: Theory and Techniques*. CRC Press. July 2010. ISBN: 1439803153/978-1439803158
- [B4] Shrideep Pallickara, Geoffrey Fox, Mehmet Aktas, Harshawardhan Gadgil†, Beytullah Yildiz†, Sangyoon Oh, Sima Patel, Marlon Pierce and Damodar Yemme A Retrospective on the Development of Web Service Specifications. Chapter II (pp 22-49) in *Securing Web Services: Practical Usage of Standards and Specifications*. Editor Periorellis Panos, University of Newcastle Upon Tyne. Published by IGI Global. ISBN-13: 978-1599046396, ISBN-10: 1599046393. 2007.
- [B3] Geoffrey Fox, Shrideep Pallickara, Galip Aydin and Marlon Pierce. Messaging in Web Service Grid with Applications to Geographical Information Systems. Chapter in *Grid Computing: The New Frontier Of High Performance Computing*. (ed. L. Grandinetti, published by Elsevier, Amsterdam, 2005) pp 305-331. ISBN: 0-444-51999-8
- [B2] Geoffrey Fox and Shrideep Pallickara. NaradaBrokering: An Event Based Infrastructure for Building Scaleable Durable Peer-to-Peer Grids. Chapter 22 of *Grid Computing: Making the Global Infrastructure a Reality Grid*. Published by John Wiley, West Sussex, England. ISBN 0-470-85319-0. 2003.
- [B1] Geoffrey Fox, Dennis Gannon, Sung-Hoon Ko, Sangmi Lee, Shrideep Pallickara, Marlon Pierce, Xiaohong Qiu, Xi Rao, Ahmet Uyar†, Minjun Wang, Wenjun Wu . Peer-to-Peer Grids. Chapter 18 of *Grid Computing: Making the Global Infrastructure a Reality Grid*. Published by John Wiley, West Sussex, England. ISBN 0-470-85319-0. 2003.

6.3 Refereed Conference Proceedings

- [C67] Wes Lloyd*, Shruti Ramesh, Swetha Chinthalapati, Lan Ly, and Shrideep Pallickara. Serverless Computing: An Investigation of Factors Influencing Microservice Performance. (To appear) Proceedings of the IEEE International Conference on Cloud Engineering (IC2E). Orlando, Florida, USA. 2018. [19% acceptance rate]
- [C66] Wes Lloyd*, Shrideep Pallickara, Olaf David, Mazdak Arabi, and Ken Rojas. Mitigating Resource Contention and Heterogeneity in Public Clouds for Scientific Modeling Services. Proceedings of the IEEE International Conference on Cloud Engineering (IC2E). Vancouver, Canada. pp 159-166. 2017. [33% acceptance rate]
- [C65] Naman Shah, Harshil Shah*, Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Network Analysis for Identifying and Characterizing Disease Outbreak Influence from Voluminous Epidemiology Data. Proceedings of the IEEE International Conference on Big Data (IEEE BigData). pp 1222-1231. Washington D.C., USA. 2016. [18.68% acceptance rate]
- [C64] Thilina Buddhika* and Shrideep Pallickara. Neptune: Real Time Stream Processing for Internet of Things and Sensing Environments. Proceedings of the 30th IEEE International Parallel & Distributed Processing Symposium. pp 1143-1152. Chicago, USA. 2016. [23% acceptance rate]
- [C63] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Alleviation of Disk I/O Contention in Virtualized Settings for Data-Intensive Computing. Proceedings of the IEEE/ACM International Symposium on Big Data Computing. Cyprus. 2015. [16% acceptance rate]
- [C62] Amila Suriarachchi* and Shrideep Pallickara. DEBS Grand Challenge: A High-Throughput, Scalable Solution for Calculating Frequent Routes and Profitability of New York Taxis. Proceedings of the ACM International Conference on Distributed Event-Based Systems. Oslo, Norway. 2015. pp 301-308.
** Shortlisted for the ACM DEBS Grand Challenge Award.
- [C61] Zhiquan Sui*, Neil Harvey, and Shrideep Pallickara. Learning Based Distributed Orchestration of Stochastic Discrete Event Simulations. Proceedings of the IEEE/ACM Conference on Utility and Cloud Computing. London, UK. 2014. pp 99-108. [19% acceptance rate]
- [C60] Matthew Malensek*, Walid Budgaga*, Sangmi Lee Pallickara, Neil Harvey, Jay Breidt, and Shrideep Pallickara. Using Distributed Analytics to Enable Real-Time Exploration of Discrete Event Simulations. Proceedings of the IEEE/ACM Conference on Utility and Cloud Computing. London, UK. 2014. pp 49-58. [19% acceptance rate]
- [C59] Wes Lloyd*, Shrideep Pallickara, Olaf David, Mazdak Arabi, and Ken Rojas. Dynamic Scaling for Service Oriented Applications: Implications of Virtual Machine Placement on IaaS Clouds. Proceedings of the IEEE Conference on Conference on Cloud Engineering. Boston, MA. 2014. pp 271-276. [20.9% acceptance rate]
- [C58] Matthew Malensek*, Sangmi Lee Pallickara and Shrideep Pallickara. Polygon-Based Query Evaluation over Geospatial Data Using Distributed Hash Tables. Proceedings of the IEEE/ACM Conference on Utility and Cloud Computing. Dresden, Germany. pp 219-226. 2013. [24% acceptance rate]
- [C57] Ryan Stern* and Shrideep Pallickara. On the Role of Topology in Autonomously Coping with Failures in Content Dissemination Systems. Proceedings of the ACM Cloud and Autonomic Computing Conference. Miami, USA. 2013. [35% acceptance rate]
- [C56] Matthew Malensek*, Sangmi Lee Pallickara and Shrideep Pallickara. Autonomously Improving Query Evaluations over Multidimensional Data in Distributed Hash Tables. Proceedings of the ACM Cloud and Autonomic Computing Conference. Miami, USA. 2013. [35% acceptance rate]
- [C55] Matthew Malensek*, Zhiquan Sui*, Neil Harvey and Shrideep Pallickara. Autonomous, Failure-resilient Orchestration of Distributed Discrete Event Simulations. Proceedings of the ACM Cloud and

- Autonomic Computing Conference*. Miami, USA. 2013. [35% acceptance rate]
- [C54] Wes Lloyd*, Shrideep Pallickara, Olaf David, Jim Lyon, Mazdak Arabi, and Ken Rojas. Service Isolation vs. Consolidation: Implications for IaaS Cloud Application Deployment. Proceedings of the *IEEE Conference on Cloud Engineering*. San Francisco, CA. 2013. [20.5% acceptance rate]
- [C53] Kathleen Ericson*, Shrideep Pallickara, and Charles Anderson. Cloud-Based Analysis of EEG Signals for BCI Applications. Proceedings of the 5th *International Brain Computer Interfaces (BCI) Meeting*. Pacific Grove, CA. USA. 2013.
- [C52] Kathleen Ericson* and Shrideep Pallickara. On the Performance of Virtualized Infrastructures for Processing Real time Streaming Data. Proceedings of the *IEEE/ACM Conference on Utility and Cloud Computing*. Chicago, USA. 2012. [27% acceptance rate]
- [C51] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Expressive Query Support for Multidimensional Data in Distributed Hash Tables. Proceedings of the *IEEE/ACM Conference on Utility and Cloud Computing*. Chicago, USA. 2012. [27% acceptance rate].
** Best Paper Award.
- [C50] Wes Lloyd*, Shrideep Pallickara, Olaf David, Jim Lyon, Mazdak Arabi and Ken Rojas. Performance Modeling to Support Multi-Tier Application Deployment to Infrastructure-as-a-Service Clouds. Proceedings of the *IEEE/ACM Conference on Utility and Cloud Computing*. Chicago, USA. 2012. [27% acceptance rate]
- [C49] Sangmi Lee Pallickara, Matthew Malensek* and Shrideep Pallickara. Enabling Access to Time-Series, Geospatial Data for On Demand Visualization. Proceedings of the *IEEE Symposium on Large-Scale Data Analysis and Visualization*, Providence, Rhode Island. 2011.
- [C48] Wes Lloyd*, Shrideep Pallickara, Olaf David, Jim Lyon, Mazdak Arabi and Ken Rojas. Migration of Multi-tier Applications to Infrastructure-as-a-Service Clouds: An Investigation Using Kernel-based Virtual Machines. Proceedings of the 12th *IEEE/ACM International Conference on Grid Computing*. pp 137-144. Lyon, France. 2011. [29% acceptance rate]
- [C47] Matthew Malensek*, Sangmi Lee Pallickara, and Shrideep Pallickara. Galileo: A Framework for Distributed Storage of High-Throughput Data Streams. Proceedings of the *IEEE/ACM Conference on Utility and Cloud Computing*. pp. 17-24. Melbourne, Australia. 2011. [26.7% acceptance rate]
- [C46] Kathleen Ericson* and Shrideep Pallickara. On the Performance of Distributed Clustering Algorithms in File and Streaming Processing Systems. Proceedings of the *IEEE/ACM Conference on Utility and Cloud Computing*. pp. 33-40. Melbourne, Australia. 2011. [26.7% acceptance rate]
- [C45] Kathleen Ericson*, Shrideep Pallickara, and Charles Anderson. Analyzing Electroencephalograms Using Cloud Computing Techniques. Proceedings of the *IEEE International Conference on Cloud Computing Technology and Science*. pp. 185-192. Indianapolis. November 2010. [24.9% acceptance rate]
** Best Student Paper Award
- [C44] Sangmi Lee Pallickara, Shrideep Pallickara, Milija Zupanski, and Stephen Sullivan. Efficient Metadata Generation to Enable Interactive Data Discovery over Large-scale Scientific Data Collections. Proceedings of the *IEEE International Conference on Cloud Computing Technology and Science*. pp. 573-580. Indianapolis. November 2010. [24.9% acceptance rate]
- [C43] Shrideep Pallickara, Jaliya Ekanayake[‡] and Geoffrey Fox. Granules: A Lightweight, Streaming Runtime for Cloud Computing With Support for Map-Reduce. Proceedings of the *IEEE International Conference on Cluster Computing (CLUSTER 2009)*. New Orleans, LA. [48% acceptance rate]
- [C42] Tao Huang, Shrideep Pallickara and Geoffrey Fox. A Framework for the Collaborative Annotation of Streams. Proceedings of the *IEEE International Symposium on Collaborative Technologies and*

- Systems. Baltimore, MD. 2009.
- [C41] Jaliya Ekanayake[†], Shrideep Pallickara, and Geoffrey Fox. Map-Reduce for Data Intensive Scientific Analyses. Proceedings of the *IEEE International Conference on e-Science*. Indianapolis. 2008. [29.6% acceptance rate]
- [C40] § Shrideep Pallickara, Jaliya Ekanayake[†], and Geoffrey Fox. An Overview of the Granules Runtime for Cloud Computing. Proceedings of the *IEEE International Conference on e-Science*. Indianapolis. 2008. [29.6% acceptance rate]
- [C39] Jaliya Ekanayake[†], Shrideep Pallickara, and Geoffrey Fox. A Collaborative Framework for Scientific Data Analysis. Proceedings of the *IEEE International Symposium on Collaborative Technologies and Systems*. Irvine, California. 2008.
- [C38] Beytullah Yildiz[†], Geoffrey Fox, Shrideep Pallickara. An Orchestration for Distributed Web Service Handlers. Proceedings of the *Third International Conference on Internet and Web Applications and Services ICIW 2008 June 8-13, 2008 - Athens, Greece*. [30% acceptance rate]
- [C37] § Shrideep Pallickara, Hasan Bulut, and Geoffrey Fox. Fault-Tolerant Reliable Delivery of Messages in Distributed Publish/Subscribe Systems. Proceedings of the *4th IEEE International Conference on Autonomic Computing*. Jacksonville, Florida. 2007. [14% acceptance rate]
- [C36] Harshawardhan Gadgil[†], Geoffrey Fox, Shrideep Pallickara, and Marlon Pierce. Scalable, Fault-tolerant Management of Grid Services. Proceedings of the *IEEE Cluster 2007 Conference*. Austin, Texas. [39.6% acceptance rate]
- [C35] § Shrideep Pallickara, Jaliya Ekanayake[†], and Geoffrey Fox. A Scalable Approach for the Secure and Authorized Tracking of the Availability of Entities in Distributed Systems. Proceedings of the *21st IEEE International Parallel & Distributed Processing Symposium (IPDPS)*. Longbeach, California. 2007. [26% acceptance rate]
- [C34] Harshawardhan Gadgil[†], Geoffrey Fox, Shrideep Pallickara, and Marlon Pierce. Scalable, Fault Tolerant Management in a Service-Oriented Architecture. Proceedings of the *2007 ACM International Symposium on High-Performance Distributed Computing (HPDC)*. Monterey Bay, CA. [20% acceptance rate]
- [C33] § Shrideep Pallickara, Marlon Pierce, Harshawardhan Gadgil[†], Geoffrey Fox, Yan Yan, Yi Huang. A Framework for Secure End-to-End Delivery of Messages in Publish/Subscribe Systems. Proceedings of the *7th IEEE/ACM International Conference on Grid Computing (GRID 2006)*. Barcelona, Spain. [20.2% acceptance rate]
- [C32] Beytullah Yildiz[†], Shrideep Pallickara and Geoffrey Fox. Experiences with deploying services within the Axis container. Proceedings of the *2006 IEEE International Conference on Internet and Web Applications and Services*. French Caribbean. [40.9% acceptance rate]
- [C31] § Shrideep Pallickara, Geoffrey Fox and Harshawardhan Gadgil[†]. On the Creation & Discovery of Topics in Distributed Publish/Subscribe systems. Proceedings of the *ACM/IEEE GRID 2005*, pp 25-32. Seattle, WA. [18.8% acceptance rate]
** Selected as one of the six Best Papers.
- [C30] Geoffrey Fox, Alex Ho, Shrideep Pallickara, Marlon Pierce, Wenjun Wu. Grids for the GiG and Real Time Simulations. Proceedings of the *IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications (DS-RT 2005)*, pp 129-138. Montreal, Canada. [33% acceptance rate]
- [C29] § Shrideep Pallickara, Geoffrey Fox, Beytullah Yildiz[†], Sangmi Lee Pallickara, Sima Patel and Damodar Yemme. On the Costs for Reliable Messaging in Web/Grid Service Environments. Proceedings of the *IEEE International Conference on e-Science & Grid Computing*. Melbourne, Australia. pp 344-351. [31.6% acceptance rate]

- [C28] Geoffrey Fox et al. Building Sensor Filter Grids: Information Architecture for the Data Deluge. Proceedings of the *IEEE International Conference on Semantics, Knowledge and Grid (SKG2005)*, Beijing, China, November, 27-29, 2005. [15% acceptance rate]
- [C27] § Shrideep Pallickara, Harshawardhan Gadgil[†] and Geoffrey Fox. On the Discovery of Brokers in Distributed Messaging Infrastructures. Proceedings of the *IEEE Cluster 2005 Conference*. Boston, MA. [32.6% acceptance rate]
- [C26] Sang Boem Lim, Geoffrey Fox, Ali Kaplan, Shrideep Pallickara and Marlon Pierce. GridFTP and Parallel TCP Support in NaradaBrokering. Proceedings of the *Sixth International Conference on Algorithms and Architectures for Parallel Processing*. Lecture Notes in Computer Science (LNCS) published by Springer-Verlag. ISBN 3-540-29235-7 pp 93-102. [27.5% acceptance rate]
- [C25] Geoffrey Fox, Galip Aydin, Harshawardhan Gadgil[†], Shrideep Pallickara, Marlon Pierce, and Wenjun Wu. Management of Real-Time Streaming Data Grid Services. Fourth *International Conference on Grid and Cooperative Computing (GCC2005)*, held in Beijing, China. Lecture Notes in Computer Science (LNCS) published by Springer-Verlag. ISSN: 0302-9743. Volume 3795. pp 3-12. [10% acceptance rate]
- [C24] Pete Burnap, Hasan Bulut, Shrideep Pallickara, Geoffrey Fox, David Walker, Beytullah Yildiz[†], A Kaplan and M Nacar. Worldwide Messaging Support for High Performance Real-time Collaboration. Proceedings of the UK e-Science Programme All Hands Meeting 2005 (AHM2005). Nottingham, UK.
- [C23] § Shrideep Pallickara, Geoffrey Fox and Sangmi Lee Pallickara. An Analysis of Reliable Delivery Specifications for Web Services. Proceedings of the *IEEE ITCC Conference on Information Technology 2005*, pp 360-365.
- [C22] § Shrideep Pallickara and Geoffrey Fox. An Analysis of Notification Related Specifications for Web/Grid applications. Proceedings of the *IEEE ITCC Conference on Information Technology 2005*. pp 762-763.
- [C21] Harshawardhan Gadgil[†], Geoffrey Fox, Shrideep Pallickara, Marlon Pierce, Robert Granat . A Scripting based Architecture for Management of Streams and Services in Real-time Grid Applications. Proceedings of the *IEEE/ACM Cluster Computing and Grid 2005 Conference (CCGrid 2005)*. Cardiff, UK. [32.3% acceptance rate]
- [C20] Harshawardhan Gadgil[†], Geoffrey Fox, Shrideep Pallickara. HPSearch for Managing Distributed Services. Work In Progress paper presented at the *IEEE/ACM Cluster Computing and Grid 2005 Conference (CCGrid 2005)*. Cardiff, UK. [32.3% acceptance rate]
- [C19] § Geoffrey Fox, Shrideep Pallickara and Savas Parastatidis. Towards Flexible Messaging for SOAP Based Services. Proceedings of the *IEEE/ACM Supercomputing Conference 2004*. Pittsburgh, PA. [30.7% acceptance rate]
- [C18] Xiaohong Qiu, Shrideep Pallickara, and Ahmet Uyar[†]. Making SVG a Web Service in a Message-based MVC Architecture. Proceedings of the *Scalable Vector Graphics Conference* Tokyo, Japan. 2004
- [C17] Hasan Bulut, Geoffrey Fox, Wenjun Wu, Ahmet Uyar[†], Shrideep Pallickara and Harun Altay. A Web Services Based Streaming Gateway for Heterogeneous A/V Collaboration. Proceedings of the 2004 *International Conference on Internet Computing*. Volume II, pp 493-499.
- [C16] § Shrideep Pallickara and Geoffrey Fox. A Scheme for Reliable Delivery of Events in Distributed Middleware Systems. Proceedings of the *IEEE International Conference on Autonomic Computing*. pp 328-329. 2004. [25% acceptance rate]
- [C15] Hasan Bulut, Shrideep Pallickara and Geoffrey Fox. Implementing a NTP-Based Time Service within a Distributed Brokering System. Proceedings of the *ACM International Conference on the Principles and Practice of Programming in Java*. pp 126-134. 2004.
- [C14] Shrideep Pallickara and Geoffrey Fox. On the Matching Of Events in Distributed Brokering Systems.

- § Proceedings of the *IEEE ITCC Conference on Information Technology*. April 2004. pp 68-76.
- [C13] Amey Dharurkar[†], Shrideep Pallickara and Geoffrey Fox. A Topology Viewer for Distributed Brokering Systems. Proceedings of the *IEEE ITCC Conference on Information Technology*. 2004. pp 228-234.
- [C12] Minjun Wang, Geoffrey Fox and Shrideep Pallickara. A Demonstration of Collaborative Web Services and Peer-to-Peer Grids. Proceedings of the *IEEE ITCC2004 International Conference on Information Technology*. Volume II pp 62-67.
- [C11] Shrideep Pallickara and Geoffrey Fox. NaradaBrokering: A Middleware Framework and Architecture for Enabling Durable Peer-to-Peer Grids. Proceedings of *ACM/IFIP/USENIX International Middleware Conference Middleware-2003*. pp 41-61. Lecture Notes in Computer Science (LNCS) published by Springer-Verlag 2003, ISBN 3-540-40317-5. [15.8% acceptance rate]
- [C10] Gurhan Gunduz[†], Shrideep Pallickara and Geoffrey Fox. A Portal based Approach to Aggregating Network Performance. Proceedings of the 2003 *International Conference on Internet Computing*. Volume II pp 495-501.
- [C9] Shrideep Pallickara, Geoffrey Fox, and Marlon Pierce. Incorporating an XML Matching Engine into Distributed Brokering Systems. Proceedings of the *International Conference on Parallel and Distributed Processing Techniques and Applications*. (PDPTA'03). Volume IV pp 1511-1517.
- [C8] Ahmet Uyar[†], Shrideep Pallickara and Geoffrey Fox. Towards an Architecture for Audio/Video Conferencing in Distributed Brokering Systems. Proceedings of the 2003 *International Conference on Communications in Computing*. pp 17-23.
- [C7] Shrideep Pallickara, Geoffrey Fox, John Yin, Gurhan Gunduz[†], Hongbin Liu, A Uyar, M Varank. A Transport Framework for Distributed Brokering Systems. Proceedings of the *International Conference on Parallel and Distributed Processing Techniques and Applications*. pp 772-778. 2003.
- [C6] Geoffrey Fox, Hasan Bulut, Kangseok Kim, Sung-Hoon Ko, Sangmi Lee, Sangyoon Oh, Shrideep Pallickara, Xiaohong Qiu, Ahmet Uyar, Minjun Wang, Wenjun Wu. Collaborative Web Services and Peer-to-Peer Grids. Proceedings of the 2003 *IEEE Collaborative Technologies Symposium*.
- [C5] Geoffrey Fox, Shrideep Pallickara, Xi Rao[†]. A Scaleable Event Infrastructure for Peer to Peer Grids. Proceedings of the *ACM Java Grande ISCOPE Conference 2002*. pp 66-75. Seattle, WA. [32% acceptance rate]
- [C4] Hasan Bulut, Geoffrey Fox, Dennis Gannon, Kangseok Kim, Sung-Hoon Ko, Sangmi Lee, Sangyoon Oh, Xi Rao, Shrideep Pallickara, Qinglin Pei, Marlon Pierce, Aleksander Slominski, Ahmet Uyar, Wenjun Wu, Choonhan Youn. An Architecture for e-Science and its Implications. Proceedings of the 2002 *International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS)*. International Society for Modeling & Simulation, pp 14-24.
- [C3] Geoffrey Fox and Shrideep Pallickara. The Narada Event Brokering System: Overview and Extensions Proceedings of the 2002 *International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'02)*. Volume I pages 353-359. [38% acceptance rate]
- [C2] Geoffrey Fox, Ozgur Balsoy, Shrideep Pallickara, Ahmet Uyar, Dennis Gannon, Aleksander Slominski. Community Grids. Proceedings of the *International Conference on Computational Science (ICCS 2002)*. Amsterdam, Netherlands April 2002. pp 22-38. Lecture Notes in Computer Science (LNCS) published by Springer-Verlag 2002, ISBN 3-540-43591-3. [46% acceptance rate]
- [C1] Geoffrey Fox and Shrideep Pallickara. JMS Compliance in the Narada Event Brokering System. Proceedings of the 2002 *International Conference on Internet Computing (IC-02)*. Vol 2 pp 391-397. [30% acceptance rate]

6.4 Refereed Workshop Proceedings

- [W6] Kathleen Ericson* and Shrideep Pallickara. Robust Processing of Health Stream Data. *ACM/IEEE Super Computing Conference (SC)*. Doctoral Showcase. Denver, USA. 2013.
- [W5] Kathleen Ericson*, Shrideep Pallickara, and Charles Anderson. Handwriting Recognition Using a Cloud Runtime. *Proceedings of the Colorado Celebration of Women in Computing (CCWIC)*. 2010.
** Selected as one of the 8 showcased presentations.
- [W4] Harshawardhan Gadgil†, Geoffrey Fox, Shrideep Pallickara, Marlon Pierce. Managing Grid Messaging Middleware. *Proceedings of the IEEE Workshop on Challenges of Large Applications in Distributed Environments (CLADE)*, June 19, 2006, Paris, France. Held in conjunction with the IEEE HPDC 2006 conference.
- [W3] Geoffrey Fox, Wenjun Wu, Ahmet Uyar‡, Hasan Bulut, Shrideep Pallickara. A Web Services Framework for Collaboration and Videoconferencing. *Proceedings of the 3rd Annual Workshop on Advanced Collaborative Environments co-located with IEEE HPDC-12 & GGF8*. 2003.
- [W2] Geoffrey Fox, Wenjun Wu, Ahmet Uyar‡, Hasan Bulut and Shrideep Pallickara. Global Multimedia Collaboration System. *Proceedings of the 1st International Workshop on Middleware for Grid Computing. Co-located with ACM/IFIP Middleware Conference*. pp 245-250. 2003.
- [W1] Daniel Dias, Geoffrey Fox, Wojtek Furmanski, Vishal Mehra, Balaji Natarajan, H.Timucin Ozdemir, Shrideep Pallickara and Zeynep Ozdemir. Exploring JSDA, CORBA and HLA based MuTechs for Scalable Televirtual (TVR) Environments. *Workshop on Object Orientation and VRML in conjunction with the Virtual Reality Modeling Language Symposium*. Monterey, California 1998

6.5 Ph.D. Dissertation and Masters Thesis

- [G2] A Grid Event Service – *Ph.D. Dissertation, June 2001. Syracuse University*. Advisor: Geoffrey Fox
§
- [G1] Java Distributed Collaborative Environment (JDCE) – As a Test-bed for Distributed Object Technology *Masters Thesis, Aug'98 Syracuse University*. Advisor: Geoffrey Fox.
§

6.6 Non-overlapping Technical Reports [NOT PEER REVIEWED]

- [R4] Geoffrey Fox, Harshawardhan Gadgil†, Shrideep Pallickara, Marlon Pierce, John Rundle, Andrea Donnellan, Jay Parker, Robert Granat, Greg Lyzenga, Dennis McLeod and Anne Chen. Complexity Computational Environments (CCE) Architecture NASA Technical Report September 2004.
- [R3] Geoffrey Fox Harshawardhan Gadgil†, Shrideep Pallickara, Marlon Pierce, Robert L. Grossman, Yunhong Gu, David Hanley, Xinwei Hong . High Performance Data Streaming in Service Architectures. Technical Report Indiana University and University of Illinois at Chicago. July 2004
- [R2] G.C. Fox, W. Furmanski, H. T. Ozdemir, S. Pallickara. New Systems Technologies and Software Products for HPCC: Volume III – High Performance Commodity Computing on the Pragmatic Object Web Management White Paper for Research Consortium Inc (RCI). - Summer '98.
- [R1] Wojtek Furmanski, Daniel Dias, Balaji Natrajan, Vishal Mehra, Shrideep Pallickara. Prototype of a Scalable Tele-Virtual Environment on the Web Using VRML2.0 / JSDA IBM T. J. Watson Research Center, April 1997. Joint Project involving IBM-Watson Research Center & Syracuse University. Principal Investigator: Dr. Wojtek Furmanski.

7 PROJECTS

My research has led to projects in the area of large-scale distributed systems. All systems-software, produced as part of these projects, is in the open-source domain. The code-bases for these systems are large involving hundreds of classes and packages. The system software is robust and has been deployed in academic, commercial, and defense domains such as earthquake science, high-energy physics, epidemic modeling, environmental and ecological monitoring, atmospheric sciences, internet-scale conferencing systems, and brain computer interfaces.

7.1 Granules

Granules supports the processing of data streams over a distributed collection of processing elements. Such streams can be generated in settings involving observational and monitoring equipment, simulations, and computational workflows. In Granules these computations can be long running, with multiple rounds of execution, with the ability to retain state across successive rounds. Granules allows a collection of related computations to be expressed as directed graphs that have cycles in them, and orchestrates the completion of such distributed processing. Granules manages the lifecycle and finite state machine associated with computations. The system can orchestrate such stream processing computations within traditional clusters, collection of desktops, or IaaS VM-based settings. The processing encapsulated within these computations can be arbitrary, and encoded in C, C++, C#, Java, R and Python. Granules also incorporates support for variants of the MapReduce paradigm that make it amenable for scientific applications. By abstracting the complexities of doing I/O and the vagaries of execution in distributed settings, Granules allows a domain scientist to focus on the problem on at hand and not on the artifacts related to deployments in large-scale distributed systems. A broad class of compute and data intensive applications can benefit from the capabilities available in Granules. Some of the application domains that Granules is currently deployed in include brain computer interfaces, epidemiological modeling, handwriting recognition, data clustering algorithms, and bio-informatics (mRNA sequencing). Granules is funded by the Department of Homeland Security and the US National Science Foundation.

Granules includes subprojects, Neptune and Forager. Neptune is an effort that targets high-throughput stream processing in continuous sensing environments and Internet-of-Things settings. Forager is a framework for harvesting CPU cycles and memory for computationally-intensive, iterative task-based systems such as stochastic discrete event simulations, and learning and analytic algorithms.

URL: <http://granules.cs.colostate.edu>

7.2 Galileo

Time-series data occurs in settings such as observations initiated by radars and satellites, checkpointing data representing state of the system at regular intervals, and analytics representing the evolution of extracted knowledge over time. Galileo is a demonstrably scalable storage framework for managing such time-series data. The distributed storage system is incrementally scalable with the ability to assimilate new storage nodes as they become available. How data is stored and dispersed impacts the efficiency of subsequent retrievals. The data dispersion algorithm in Galileo stores similar data items in network proximity without introducing storage imbalances at individual storage nodes. This allows for a significant reduction in the search space for queries that may be performed on the stored data. Galileo is funded by the Department of Homeland Security and the Environmental Protection Agency.

URL: <http://galileo.cs.colostate.edu> [Joint effort with the Big Data group at CSU]

7.3 Glean

The Glean project focuses on performing analytics at scale over Big Data. The datasets we consider are in the order of Petabytes and encompass billions of files representing trillions of observations, measurements, or simulation datapoints. Glean achieves this by combing innovations in large-scale storage systems, cloud computing, machine learning, and statistics. A particular focus of this effort is to perform analytics in real-time over streaming data representing time-series observations. The effort makes use of machine learning, statistical techniques, Bayesian networks, and ensemble methods to achieve this.

URL: <http://glean.cs.colostate.edu> [Joint effort with the Big Data group at CSU]

7.4 Spindle

The Spindle effort focuses on issues in migration of applications to Infrastructure-as-a-Service clouds. We construct performance models for applications based on profiling their constituent components during execution. We extract several features relating to CPU processing, memory consumption, and I/O. We then use these performance models to inform our VM placement and composition decisions while also reducing economic costs. These performance models are continually updated to reflect the application's performance under conditions of varying load.

URL: <http://spindle.cs.colostate.edu>

7.5 Funnel

The Funnel project focuses on the scalable dissemination of streams generated by sensors and other observational devices. Funnel is a content dissemination system that combines aspects of peer-to-peer systems with that of publish/subscribe to provide infrastructureless, high-throughput routing of data streams. Unlike traditional publish/subscribe systems that require a dedicated middleware to perform the content dissemination, in Funnel the nodes (either producers or consumers) organize themselves into a topology that is then used for disseminating content. Unlike traditional P2P systems, the focus is on real-time, high-throughput routing of data streams between multiple producers and consumers in the system.

URL: <http://funnel.cs.colostate.edu>

7.6 NaradaBrokering

NaradaBrokering is an infrastructure for managing voluminous data streams being produced concurrently, at high rates, by a very large number of entities. The infrastructure allows entities to precisely specify constraints on portions of the streams that they are interested in consuming. By preferentially deploying links during disseminations, the routing solution ensures that underlying network is optimally utilized. NaradaBrokering provisions a rich set of easy-to-use capabilities. Since the provisioned capabilities-stack is mutable, scientists can instrument, permute and combine different capabilities to compose streaming systems. NaradaBrokering has been deployed in a diverse set of domains such as earthquake science, particle physics, ecological/environmental monitoring, geosciences, GIS systems, defense applications, and commercial internet conferencing systems. Research in NaradaBrokering has been funded through two grants from the National Science Foundation, two grants from the United Kingdom's Open Middleware Infrastructure Institute and an STTR grant from the Department of Energy.

URL: <http://www.naradabrokering.org>

8 TEACHING

Teaching is often the easy part – knowing what the students don't know is the hard part. I rely on micro-surveys at the end of every lecture to identify concepts that students grasped and those that they had problems with. Comprehension problems often stem from gaps in the student's knowledge; I have found that addressing such gaps helps students grasp and retain concepts better. These micro-surveys ask students to list concepts that they understood and concepts that they don't. I find that micro-surveys in *Computer Science* are particularly amenable to gap analysis, and I have used them to identify the deeper reasons why students don't understand a particular concept – this is akin to treating the cause rather than the symptom. In courses that I have taught multiple times, I have found that retaining micro-surveys from the past help me perform a time-series analysis to:

- (1) Subtly adjust explanation of difficult concepts. This has been very helpful in continually making targeted and incremental adjustments to my teaching materials.
- (2) Pose questions in class based on the misunderstandings that students have had in the past. These questions help anchor discussions that students are often reluctant to initiate.

Finally, I feel that to truly assimilate knowledge one must be able to build on it; my courses generally involve several programming assignments that are geared towards reinforcing concepts covered in class.

8.1 Curriculum Development

I created the track on *Distributed Systems* in the Computer Science curriculum at CSU. This involved the development of 3 new computer science courses: CS455, CS555, and CS655. CS455 is designated as a *capstone course* and is targeted at senior undergraduate students. First year graduate students take CS555, while CS655 is an advanced research seminar course typically taken by PhD students. The primary objective of these courses is for students to gain a deeper understanding of the algorithms, frameworks, and strategies that underpin complex, large-scale distributed systems. The x55 series of courses entail extensive programming and substantive systems building assignments in addition to a writing component that is generally split into deliverables to allow critical feedback on the content and organization of ideas.

8.2 Courses taught at Colorado State University

TERM	COURSE	ENROLLMENT	AUDIENCE
Spring 2018	CS455: <i>Introduction to Distributed Systems</i> [On-campus (55) and Online (18)]	73	Undergraduate CAPSTONE COURSE
Fall 2017	CS555: <i>Distributed Systems</i> [On-campus (7) and Online (8)]	15	Graduate Course
Spring 2017	CS455: <i>Introduction to Distributed Systems</i> [On-campus (42) and Online (17)]	59	Undergraduate CAPSTONE COURSE
Spring 2017	CS370: <i>Operating Systems</i>	123	Core undergraduate Course
Spring 2016	CS370: <i>Operating Systems</i>	102	Core undergraduate Course
Fall 2015	CS555: <i>Distributed Systems</i>	15	Graduate Course
Spring 2015	CS455: <i>Introduction to Distributed Systems</i>	52	Undergraduate CAPSTONE COURSE
Fall 2014	CS 370: <i>System Architecture and Software</i>	68	Core undergraduate Course
Spring 2014	CS455: <i>Introduction to Distributed Systems</i>	46	Undergraduate CAPSTONE COURSE
Fall 2013	CS655: <i>Advanced Topics in Distributed Systems</i>	4	Graduate Seminar Course
Spring 2013	CS455: <i>Introduction to Distributed Systems</i>	30	Undergraduate CAPSTONE COURSE
Fall 2012	CS555: <i>Distributed Systems</i>	13	Graduate Course

Fall 2012	CS370: System Architecture and Software	44	Core undergraduate Course
Spring 2012	CS455: Introduction to Distributed Systems.	41	Undergraduate CAPSTONE COURSE
Spring 2012	CS 370: System Architecture and Software	45	Core undergraduate Course
Fall 2011	CS655: Advanced Topics in Distributed Systems	7	Graduate Seminar Course
Spring 2011	CS 370: System Architecture and Software	55	Core undergraduate Course
Fall 2010	CS555: Distributed Systems	18	Graduate Course
Spring 2010	CS451: Operating Systems	28	Senior undergraduate Course
Fall 2009	CS670: Special Topics in Architecture/Systems {Distributed Systems}	10	Graduate Course

8.3 Cortez Middle School, Montezuma-Cortez School District Cortez, Colorado

Since 2013, I direct an NSF-funded summer camp, *Math in Action in Computer Science*, targeted at 6th and 7th graders from the Cortez Middle School. The camp involves Native American students who travel to Fort Collins; each student receives a stipend along with room-and-board. The weeklong camp focuses on reinforcing fundamental mathematical concepts alongside their applications in computer science. The CAMP includes sessions on applying for college and scholarship opportunities. The camp focuses on experiential learning and involves hands-on activities involving computer programming, individual, and group activities. Several of these students typically have limited access to computers at home and school. This is a five-year effort. <http://macs.cs.colostate.edu>

9 ACADEMIC ADVISING

Current Graduate Advisees

Thilina Buddhika, Ph.D. candidate

Walid Budgaga, Ph.D. candidate [Co-advised with Sangmi Lee Pallickara]

Sitakanta Mishra, M.S. student

Daniel Rammer, Ph.D. candidate

Ryan Stern, Ph.D. candidate

Current Undergraduate Advisees

Paul Bivrell

Kegan Strawn

Supervised Graduate Degrees

Doctoral Degree Matthew Malensek (Ph.D., 2017) [Co-advised with Sangmi Lee Pallickara]
Dissertation: *Low-latency, Query-Driven Analytics over Voluminous Multidimensional, Spatiotemporal Datasets*
Now: Assistant Professor, Computer Science – University of San Francisco

Wesley Lloyd (Ph.D., 2014)

Dissertation: *Autonomous Management of Cost, Performance, and Resource Uncertainty for Migration Of Applications to Infrastructure-as-a-Service (IaaS) Clouds*
Now: Assistant Professor, Computer Science & Systems – University of Washington, Tacoma.

Zhiquan Sui (Ph.D., 2014)

Dissertation: *Distributed Algorithms for the Orchestration of Stochastic Discrete Event Simulations.*
Now: Software Engineer, ThoughtSpot Inc, Sunnyvale CA.

Kathleen Ericson (Ph.D., 2014)

Dissertation: *Robust Health Stream Processing*
Now: Assistant Professor, Computer Science – University of Tennessee, Martin.

M.S. (Thesis)

Sudalaikkan Leo Vigneshwaran (M.S., 2016)

Preservation of Low Latency Service Request Processing in Dockerized Microservice Architectures

Harshil Shah (M.S., 2016) [Co-advised with Sangmi Lee Pallickara]

Thesis: Identification and Characterization of Super-spreaders from Voluminous Epidemiology Data

Mante Luo (M.S., 2016)

Thesis: *Horizontal Scaling of Video Conferencing Applications in Virtualized Environments*

Thilina Buddhika (M.S., 2015)

Thesis: *Real Time Stream Processing for Internet of Things and Sensing Environments*

Hanisha Koneru (M.S., 2015)

Thesis: *Implications of Storage Subsystem Interactions on Processing Efficiency in Data Intensive Computing*

Amila Suirarchchi (M.S., 2015)

Thesis: *Achieving High-Throughput Distributed, Graph-based Multi-stage Stream Processing*

Ryan Stern (M.S., 2014)

Thesis: *On the Role of Topology in Autonomously Coping with Failures in Content Dissemination Systems*

Walid Budgaga (M.S., 2014) [Co-advised with Sangmi Pallickara]

Thesis: *A Framework for Real-time, Autonomous Anomaly Detection over Voluminous Time-Series Geospatial Data Streams*

Matthew Malensek (M.S., 2012)

Thesis: *On the Evaluation of Exact-Match and Range Queries over Multidimensional Data in Distributed Hash Tables*

Kathleen Ericson (M.S., 2010)
 Thesis: *On the Support for Heterogeneous Languages in Cloud Runtimes*

Supervised Undergraduate Honors Theses

Amanda Carbonari (Dec 2015)
 Kevin Kowalski (Dec 2012)
 Hamilton Reed (Dec 2012)

Past Undergraduate Advisees

NSF REU	Kira Lindberg , Sahan Jayasumana, Marion Watts, Amanda Carbonari, Lauren Bay, Owen Davis, Sarah Orrico, Gerhard Van Andel
MACS: Math in Action in Computer Science	Ian Bertolacci, Paul Bivrell, Conrad Christensen, Tanner Evans, Gareth Halladay, Samuel Haynes, Zaneb Mazeh, Michael McCaffery, Shannon McPherson, Kegan Strawn, Marion Watts
Sites-Regelson Fellow	Jordan Messac
DHS	Conrad Christensen

Ph.D. Committee Memberships

Department of Computer Science, Colorado State University [8]
 Stephanie DiBenedetto, Chengyu Fan, Kaustubh Gadhkari, Susmit Shannigrahi, Anant Shah, Dr. Andrew Stone, Dr. He Yan, and Dr. Han Zhang

Department of Electrical Engineering & Computer Engineering, Colorado State University [2]
 Ali Fatih Boud and Saideep Tiku

Department of Electrical Engineering & Computer Science, Syracuse University [2]
 Dr. Gurhan Gunduz and Dr. Ahmet Uyar

Department of Civil & Environmental Engineering, Colorado State University [1]
 Dr. André Dozier

Ph.D. External Examiner

Department of Computing & Information Systems, University of Melbourne, Australia. [1]
 Dr. William Voorsluys

Faculty of Information Technology, Monash University, Australia [2]
 Dr. Hasanul Ferdous, Dr. Joarder Kamal

Masters Thesis Committee Memberships

Department of Computer Science, Colorado State University
 Internal Committee Member: Kiril Beyalev, Chuanqi Huang, Jared Koontz, Saptashwa Mitra, Naman Shah

Department of Civil & Environmental Engineering, Colorado State University
 Outside Committee Member: Elizabeth Mahon

10 ACADEMIC SERVICE

Editorial Boards

Journal of Big Data. Publisher: Springer New York/Heidelberg, Germany.
Member of Editorial Board (2013 – present)

International Journal of Autonomic Computing. Publisher: Inderscience, Geneva, Switzerland.
Member of Editorial Board (2008 – 2015)

Panels

NSF Panel 2010, 2011, 2012

co-Chair for the Panel on *Network Measurements* at the 35th Annual *IEEE Conference on Local Computer Networks (LCN)* in Denver. 2010.

Panel on the *FUTURE OF CLOUD COMPUTING* at the *IEEE/ACM International Conference on Utility & Cloud Computing* in Chicago. 2012.

Conferences

Program Committee: *IEEE/ACM International Conference on Utility and Cloud Computing*, Austin, Texas. 2017.

Program Committee Co-CHAIR: *IEEE/ACM International Conference on Big Data Computing, Applications and Technologies*. Austin, Texas. 2017.

Steering Committee: *NSF/DOE STREAM* (Indianapolis, 2015) and (Washington, D.C. 2016).

Program Committee: *IEEE/ACM International Conference on Utility and Cloud Computing*, Shanghai, China. 2016.

Program Committee: *IEEE International Conference on Big Data and Cloud Computing*, Atlanta GA. 2016.

Program Committee: *IEEE International Conference on Parallel Processing (ICPP)*. Beijing, China. 2016.

Program Committee: *IEEE International Conference on Data Science and Data Intensive Systems*, Sydney, Australia, 2015.

Program Committee: *IEEE/ACM International Conference on Utility and Cloud Computing*, Cyprus, 2015.

Program Committee: *Autonomic Computing, Datacenters, and Advanced Cyberinfrastructure Track of the IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid)*. Shenzhen, China. 2015.

Program Committee: *IEEE International Conference on Cloud Engineering (IC2E)*. Tempe, AZ. USA. 2015.

Program Committee: *IEEE International Conference on Computer Communications and Networks (ICCCN)*. USA. 2015.

Program Committee: *IEEE International Conference on Big Data and Cloud Computing*, China. 2015.

Program Committee Co-CHAIR: *IEEE/ACM International Conference on Utility and Cloud Computing* London, UK. 2014.

Program Committee: *IEEE International Workshop on Big Data in Computational Epidemiology*. Washington DC. USA. 2014.

Program Committee: *IEEE International Conference on Big Data and Cloud Computing*, Sydney, Australia. 2014.

Program Committee, *International Workshop on Model-Driven Engineering for High Performance and Cloud computing*, Miami, USA. Co-Located with the 16th *ACM/IEEE Conference on Model Driven Engineering Languages & Systems (MODELS)* 2013

Program Committee; Workshops Co-Chair: *IEEE/ACM International Conference on Utility and Cloud Computing* Dresden, Germany. 2013.

Program Committee, *International Workshop on Resource Management of Cloud Computing*. Collocated with the IEEE International Conference on Distributed Computing Systems (ICDCS). Philadelphia, USA. 2013.

Program Committee. The 10th *IEEE International Symposium on Parallel and Distributed Processing with Applications*. Madrid, Spain. 2012.

Program Committee. The 14th *IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing*. Ottawa, Canada. 2012.

Program Committee, Workshops-Chair. : *IEEE/ACM International Conference on Utility and Cloud Computing* Chicago, USA. 2012.

Program Committee, *International Workshop on Model-Driven Engineering for High Performance and Cloud computing*, Innsbruck Austria. Co-Located with the 15th ACM/IEEE Conference on Model Driven Engineering Languages & Systems (MODELS) 2012.

Program Committee Co-CHAIR: *IEEE/ACM International Conference on Utility and Cloud Computing*. Melbourne, Australia. 2011.

Program Committee, The 36th *IEEE Conference on Local Computer Networks (LCN)* 2011. Bonn, Germany.

Publicity Co-Chair. The 18th *IEEE International Conference on High Performance Computing (HiPC)* 2011). Bangalore, India.

Program Committee, The *IEEE International Conference on Cloud and Green Computing*, 2011. Melbourne Australia.

Publicity co-Chair; Program Committee. *IEEE International Conference on. Utility and Cloud Computing (UCC 2010)*. Chennai, India.

Program Committee. The 12th *IEEE International Conference on High Performance Computing and Communications (HPCC-10)*. 2010. Melbourne, Australia For the research track on Web Services and Internet Computing.

Program Committee, The 11th *IEEE International Conference on High Performance Computing and Communications (HPCC-09)*. Seoul, Korea,. 2009. For the research track on Web Services and Internet Computing.

Program Committee, The 14th *IEEE International Conference on Parallel and Distributed Systems (ICPADS'08)*. Melbourne, Australia. For the track on Web and Peer-to-Peer Systems.

Program Committee, The 10th *IEEE International Conference on High Performance Computing and Communications (HPCC-08)*. DaLian, China. For the research track on Web Services and Internet Computing.

University Service

ISTeC Research Advisory Committee (2013 –)

The University's Information Science & Technology Center (ISTeC) is a university-wide organization for enhancing CSU's research, education, and outreach activities pertaining to the design and innovative application of computer, communication, and information systems.

College Service

Data Science Major Committee, College of Natural Sciences (2016 –)

The College of Natural Sciences is in the process of offering a Bachelors in Data Science. Unlike majors that are at the department level, the Data Science major is at the college-level and spans the departments

of Computer Science, Mathematics, and Statistics. This major will start admitting its first students in Fall 2018. Work on this committee has involved identification of learning objectives, design of new courses, streamlining course offerings with AUCC requirements so that students graduate in a timely fashion.

Teaching Awards Committee, College of Natural Sciences (2016)

The College of Natural Sciences solicits nominations annually for Teaching and Mentoring awards for faculty, post docs, and graduate students. Service on this committee involves perusing nomination packets for different categories across departments within the college and meeting to rank and make award recommendations to the Dean.

Departmental Service

Executive Committee, Computer Science Department (2013 –2016, 2017-2020)

Represents the interests of the Faculty and provide advice and assistance to the Department Chair in the administration of the Department

Code Committee, Computer Science Department (2015–2016, 2017–)

Undergraduate Committee, Computer Science Department (2013–2015, 2016–)

Administrative procedures relating to the undergraduate: curriculum and redesign, degree requirements, recruitment, adoption of textbooks, and consistency of core courses.

Graduate Recruitment Committee, Computer Science Department (2011– 2013)

Every year the department receives about 400 applications for admission into the Ph.D. and Masters program.

Strategic Planning Committee, Computer Science Department Spring 2013

Operations Committee, Computer Science Department (2010 – 2011)

Reviewing and finalization of scholarships and awards for undergraduate and graduate students. These included the *Patricia Mohilner Memorial Scholarship* and *Thomas J. Heidenfelder Scholarship* for undergraduate students and the *Anita Read Graduate Award*.

Approvals for use of the student tech fees for computer and networking equipment purchases within the various labs in the department.

Other Reviewing Activities

Reviewer, Austria Science Fund (FWF Der Wissenschaftsfonds): 2011, 2012, 2013

I also regularly review manuscripts for Journals including – *ACM Transactions on Autonomous and Adaptive Systems*, *Concurrency and Computation: Practice & Experience*, *Future Generation Computer Systems*, *IEEE Computing in Science and Engineering*, *IEEE Transactions on Cloud Computing*, *IEEE Transactions on Computers*, *IEEE Transactions on Knowledge and Data Engineering*, *IEEE Transactions on Parallel and Distributed Systems*, *IEEE Transactions on Services Computing*, *Journal of Parallel and Distributed Computing*, and the *Journal of Supercomputing*.