

RAFAEL FERREIRA DA SILVA, Ph.D.

USC Information Sciences Institute
4676 Admiralty Way, Suite 1001
Marina del Rey, CA 90292
310-448-8352 | www.rafaelsilva.com
rafsilva@isi.edu

APPOINTMENTS

- 2016-present **Research Assistant Professor**
University of Southern California, Department of Computer Science
- 2013-present **Computer Scientist**
University of Southern California, Information Sciences Institute

My main research topic is on the optimization of the execution of scientific workflows on distributed computing systems. In particular, my work is focused on the execution of eScience applications with the Pegasus workflow management system. Other research topics include optimization methods and applications for high performance and high throughput computing, data science, and reproducibility.

- 2010-2013 **Software Engineer**
Centre National de la Recherche Scientifique (CNRS), France

As a software engineer of the Virtual Imaging Platform (VIP) I was responsible for the life cycle of the project software from the user- to middleware-level. My work included the design, management, development, and deployment of the software stack from the VIP platform. A non-exhaustive list of software includes a web portal, a workflow engine, a grid middleware, and their associated services. Additionally, I also implemented software development and maintenance processes and methods.

EDUCATION

- 2013 **Doctor of Philosophy in Computer Science**
Institut National des Sciences Appliquées de Lyon, France
“A science-gateway for workflow executions: online and non-clairvoyant self-healing of workflow executions on grids”
- My PhD thesis aimed to reach a general model of a scientific gateway that could autonomously detect and handle operational incidents during workflow executions. Our approach is based on the MAPE-K loop to determine incident degrees and to perform fault-tolerance actions. The strategy was implemented and deployed in a production environment. Results show that the proposed method speeds up execution up to a factor of 4, and properly detects unrecoverable errors.
- 2010 **Master of Science in Computer Science**
Federal University of Campina Grande (UFCG), Brazil
- 2007 **Bachelor of Science in Computer Science**
Federal University of Paraiba (UFPB), Brazil

**PROFESSIONAL
ACTIVITIES**

Funding
Agencies
Reviewer

- National Science Foundation (NSF), 2017-present.

Program
Committee
Member

- 18th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid'18), Washington DC, USA, 2018.
- IEEE Cluster, Honolulu, USA, 2017.
- 17th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid'17), Madrid, Spain, 2017.
- 12th Workflows in Support of Large-Scale Science (WORKS), Denver, USA, 2017.
- 29th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), Campinas, Brazil, 2017.
- International Workshop on Clouds and (eScience) Applications Management (CloudAM'17), Austin, USA, 2017.
- 11th Workflows in Support of Large-Scale Science (WORKS), Salt Lake City, USA, 2016.
- 28th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), Los Angeles, USA, 2016.
- 1st International Workshop on Reproducible Open Science (RepScience), Hannover, Germany, 2016.
- International Workshop on Clouds and (eScience) Applications Management (CloudAM'16), Shanghai, China, 2016.
- 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'16), Cartagena, Colombia, 2016.
- 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'15), Shenzhen, Guangdong, China, 2015.
- International Workshop on Clouds and (eScience) Applications Management (CloudAM'15), Cyprus, UK, 2015.
- International Workshop on Clouds and (eScience) Applications Management (CloudAM'14), London, UK, 2014.

Chair

- Organizer, 2018 National Science Foundation Software Infrastructure for Sustained Innovation Principal Investigators Meeting (NSF SI2 PI Meeting), Washington, USA, 2018.
- Publicity Chair, 12th Workflows in Support of Large-Scale Science (WORKS), Denver, USA, 2017.
- Publicity Chair, 11th Workflows in Support of Large-Scale Science (WORKS), Salt Lake City, USA, 2016.
- Local Arrangements Chair, 28th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), Los Angeles, USA, 2016.

Journal
Reviewer

- ACM Transactions on Parallel Computing, 2017-present.
- Journal of Grid Computing, 2016-present.
- Computers & Electrical Engineering, 2015-present.
- Future Generation Computer Systems, 2013-present.
- Concurrency and Computation: Practice and Experience, 2013-present.
- IEEE Transactions on Parallel and Distributed Systems, 2014-present.
- Journal of Parallel and Distributed Computing, 2015.
- Computing and Informatics, 2015.
- IEEE Communications Magazine, 2014.
- Scientific Programming, 2014.
- The Computer Journal, 2014.

- Conference Reviewer
- 13th IEEE International Conference on e-Science (eScience'17), Auckland, New Zealand, 2017.
 - IEEE Cluster, Taipei, Taiwan, 2016.
 - 30th IEEE International Parallel & Distributed Processing Symposium (IPDPS), Chicago, USA, 2016.
 - 11th IEEE International Conference on e-Science (eScience'15), Munich, Germany, 2015.
 - 10th Workflows in Support of Large-Scale Science (WORKS), Austin, USA, 2015.
 - Euro-Par, Vienna, Austria, 2015.
 - IEEE International Conference on Cloud Engineering (IC2E), Tempe, USA, 2014.
 - 10th IEEE International Conference on e-Science (eScience'14), Guarujá, Brazil, 2014.
 - SuperComputing'14, New Orleans, USA, 2014.
 - 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2014), Chicago, USA, 2014.
 - SuperComputing'13, Denver, USA, 2013.
 - 10th International Conference on Parallel Processing and Applied Mathematics (PPAM), Warsaw, Poland, 2013.
 - 1st International Workshop on Autonomic Management of Grid and Cloud Computing (AMGCC13), Miami, USA, 2013.

TEACHING

- ACTIVITIES
- USC, Computer Science Department, CS 104 Data Structures and Object-Oriented Design (Undergraduate level), Spring 2017.

INVITED

- TALKS
- Automating Real-Time Seismic Analysis: Through Streaming and High Throughput Workflows, *Workshop on Environmental Computing Applications – State of the art*, ICCS, San Diego, CA, USA, 2016.
 - Automating Scientific Computations: From the User's Desktop to Clouds and World-Class Supercomputers, *IRIS Workshop*, Vancouver, WA, USA, 2016.

GRANTS

- Co-PI, *MINT: Model Integration through Knowledge-Rich Data and Process Composition*, DARPA, 12/2017–11/2021, \$12,979,881, grant #W911NF-18-1-0027.
- Co-PI, *BIGDATA: IA: Collaborative Research: In Situ Data Analytics for Next Generation Molecular Dynamics Workflows*, NSF, 10/2017–09/2021, \$1,993,043, grant #1741057, #1740990, and #1741040.
- Co-PI, *Collaborative Research: SI2-SSE: WRENCH: A Simulation Workbench for Scientific Workflow for Users, Developers, and Researchers*, NSF, 01/2017–12/2019, \$497,956, grant #1642369 and #1642335.

PUBLICATIONS

- Journals
1. A. Brinckman, E. Deelman, S. Gupta, J. Nabrzyski, S. Park, R. Ferreira da Silva, I. J. Taylor, and K. Vahi, Collaborative Circuit Designs using the CRAFT Repository, *Future Generation Computer Systems*, accepted, 2018.
 2. B. Tovar, R. Ferreira da Silva, G. Juve, E. Deelman, W. Allcock, D. Thain, and M. Livny, A Job Sizing Strategy for High-Throughput Scientific Workflows, *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, iss. 2, pp. 240-253, 2018.
 3. R. Ferreira da Silva, R. Filgueira, I. Pietri, M. Jiang, R. Sakellariou, and E. Deelman, A Characterization of Workflow Management Systems for Extreme-Scale Applications, *Future Generation Computer Systems*, vol. 75, pp. 228-238, 2017.
 4. T. Glatard, M. Rousseau, S. Camarasu-Pop, R. Adalat, N. Beck, S. Das, R. Ferreira da Silva,

- N. Khalili-Mahani, V. Korkhov, P. Quirion, P. Rioux, S. D. Olabarriaga, P. Bellec, and A. C. Evans, Software architectures to integrate workflow engines in science gateways, *Future Generation Computer Systems*, vol. 75, pp. 239-255, 2017.
5. I. Santana-Perez, R. Ferreira da Silva, M. Rynge, E. Deelman, M. S. Pérez-Hernández, and O. Corcho, Reproducibility of execution environments in computational science using semantics and clouds, *Future Generation Computer Systems*, vol. 67, pp.354-367, 2017.
 6. E. Deelman, C. Carothers, A. Mandal, B. Tierney, J. S. Vetter, I. Baldin, C. Castillo, G. Juve, D. Krol, V. Lynch, B. Mayer, J. Meredith, T. Proffen, P. Ruth, and R. Ferreira da Silva, PANORAMA: An Approach to Performance Modeling and Diagnosis of Extreme Scale Workflows, *International Journal of High Performance Computing Applications*, vol. 31, iss. 1, pp. 4-18, 2017.
 7. E. Deelman, K. Vahi, M. Rynge, G. Juve, R. Mayani, and R. Ferreira da Silva, Pegasus in the Cloud: Science Automation through Workflow Technologies, *IEEE Internet Computing*, 20(1), pp. 70-76, 2016.
 8. R. Ferreira da Silva, G. Juve, M. Rynge, E. Deelman, and M. Livny, Online Task Resource Consumption Prediction for Scientific Workflows, *Parallel Processing Letters*, 25(3), 2015.
 9. W. Chen, R. Ferreira da Silva, E. Deelman, and R. Sakellariou, Using imbalance metrics to optimize task clustering in scientific workflow executions, *Future Generation Computer Systems*, vol. 46, pp. 69-84, 2015.
 10. E. Deelman, K. Vahi, G. Juve, M. Rynge, S. Callaghan, P. J. Maechling, R. Mayani, W. Chen, R. Ferreira da Silva, M. Livny, and K. Wenger, Pegasus, a workflow management system for science automation, *Future Generation Computer Systems*, vol. 46, pp. 17-35, 2015.
 11. T. Glatard, L. B. Lewis, R. Ferreira da Silva, R. Adalat, N. Beck, C. Lepage, P. Rioux, M. Rousseau, T. Sherif, E. Deelman, N. Khalili-Mahani, and A. C. Evans, Reproducibility of neuroimaging analyses across operating systems, *Frontiers in Neuroinformatics*, 9(12), 2015.
 12. J. Howison, E. Deelman, M. J. McLennan, R. Ferreira da Silva, and J. D. Herbsleb, Understanding the scientific software ecosystem and its impact: Current and future measures, Research Evaluation, 2015.
 13. W. Chen, R. Ferreira da Silva, E. Deelman, and T. Fahringer, Dynamic and fault-tolerant clustering for scientific workflows, *IEEE Transactions on Cloud Computing*, 2015.
 14. T. Glatard, L. B. Lewis, R. Ferreira da Silva, R. Adalat, N. Beck, C. Lepage, P. Rioux, M. Rousseau, T. Sherif, E. Deelman, N. Khalili-Mahani, and A. C. Evans, Reproducibility of neuroimaging analyses across operating systems, *Frontiers in Neuroinformatics*, 9(12), 2015.
 15. R. Ferreira da Silva, T. Glatard, and F. Desprez, Controlling fairness and task granularity in distributed, online, non-clairvoyant workflow executions, *Concurrency and Computation: Practice and Experience*, 26(14), pp. 2347-2366, 2014.
 16. R. Ferreira da Silva, T. Glatard, and F. Desprez, Self-healing of workflow activity incidents on distributed computing infrastructures, *Future Generation Computer Systems*, 29(8), pp. 2284-2294, 2013.
 17. T. Glatard, C. Lartizien, B. Gibaud, R. Ferreira da Silva, G. Forestier, F. Cervenansky, M. Alessandrini, H. Benoit-Cattin, O. Bernard, S. Camarasu-Pop, N. Cerezo, P. Clarysse, A. Gaignard, P. Hugonnard, H. Liebgott, S. Marache, A. Marion, J. Montagnat, J. Tabary, and D. Friboulet, A virtual imaging platform for multi-modality medical image simulation, *IEEE Transactions on Medical Imaging*, 32(1), pp. 110-118, 2013.
 18. S. Camarasu-Pop, T. Glatard, R. Ferreira da Silva, P. Gueth, D. Sarrut, and H. Benoit-Cattin, Monte carlo simulation on heterogeneous distributed systems: a computing framework with parallel merging and checkpointing strategies, *Future Generation Computer Systems*, 29(3), pp. 728-738, 2013.
 19. F. Brasileiro, M. Gaudencio, R. Ferreira da Silva, A. Duarte, D. Carvalho, D. Scardaci, L. Ciuffo, R. Mayo, H. Hoeger, M. Stanton, R. Ramos, R. Barbera, B. Marechal, and P. Gavillet, Using a simple prioritisation mechanism to effectively interoperate service and opportunistic grids in the eela-2 e-infrastructure, *Journal of Grid Computing*, vol. 9, pp.

Conference
Papers

1. R. Ferreira da Silva, S. Callaghan, and E. Deelman, *On the Use of Burst Buffers for Accelerating Data-Intensive Scientific Workflows*, in 12th Workshop on Workflows in Support of Large-Scale Science (WORKS'17), 2017.
2. A. Mandal, P. Ruth, I. Baldin, R. Ferreira da Silva, and E. Deelman, *Toward Prioritization of Data Flows for Scientific Workflows Using Virtual Software Defined Exchanges*, in First International Workshop on Workflow Science (WoWS 2017), 2017.
3. V. Lynch, J. B. Calvo, E. Deelman, R. Ferreira da Silva, M. Goswami, Y. Hui, E. Lingerfelt, and J. Vetter, *Distributed Workflows for Modeling Experimental Data*, in 2017 IEEE High Performance Extreme Computing Conference, 2017.
4. I. J. Taylor, A. Brinckman, E. Deelman, R. Ferreira da Silva, S. Gupta, J. Nabrzyski, S. Park, and K. Vahi, *Accelerating Circuit Realization via a Collaborative Gateway of Innovations*, in 9th International Workshop on Science Gateways (IWSG 2017), 2017.
5. R. Ferreira da Silva, R. Filgueira, E. Deelman, E. Pairo-Castineira, I. M. Overton, and M. Atkinson, *Using Simple PID Controllers to Prevent and Mitigate Faults in Scientific Workflows*, in 11th Workflows in Support of Large-Scale Science (WORKS), 2016.
6. R. Filgueira, R. Ferreira da Silva, A. Krause, E. Deelman, and M. Atkinson, *Asterism: Pegasus and dispel4py hybrid workflows for data-intensive science*, in 7th International Workshop on Data-Intensive Computing in the Clouds (DataCloud), 2016, pp. 1-8.
7. D. Krol, R. Ferreira da Silva, E. Deelman, and V. E. Lynch, *Workflow Performance Profiles: Development and Analysis*, in Euro-Par 2014: Parallel Processing Workshops, 2016.
8. R. Ferreira da Silva, E. Deelman, R. Filgueira, K. Vahi, M. Rynge, R. Mayani, and B. Mayer, *Automating Environmental Computing Applications with Scientific Workflows*, in Environmental Computing Workshop (ECW), 2016.
9. D. Krol, R. Ferreira da Silva, E. Deelman, and V. E. Lynch, *Science Automation in Practice: Performance Data Farming in Workflows*, in Euro-Par 2014: Parallel Processing Workshops, 2016.
10. S. Schlagkamp, R. Ferreira da Silva, W. Allcock, E. Deelman, and U. Schwiegelshohn, *Consecutive Job Submission Behavior at Mira Supercomputer*, in 25th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC), 2016.
11. S. Schlagkamp, R. Ferreira da Silva, E. Deelman and U. Schwiegelshohn, *Understanding User Behavior: from HPC to HTC*, in International Conference on Computational Science (ICCS), 2016.
12. S. Schlagkamp, R. Ferreira da Silva, J. Renker, and G. Rinkenauer, *Analyzing Users in Parallel Computing: A User-Oriented Study*, in 14th International Conference on High Performance Computing & Simulation (HPCS), 2016.
13. S. Schlagkamp, M. Hofmann, L. Eufinger, and R. Ferreira da Silva, *Increasing Waiting Time Satisfaction in Parallel Job Scheduling via a Flexible MILP Approach*, in 14th International Conference on High Performance Computing & Simulation (HPCS), 2016.
14. H. Nawaz, G. Juve, R. Ferreira da Silva, and E. Deelman, *Performance Analysis of an I/O-Intensive Workflow executing on Google Cloud and Amazon Web Services*, in 18th Workshop on Advances in Parallel and Distributed Computational Models (APDCM), 2016.
15. A. Mandal, P. Ruth, I. Baldin, D. Krol, G. Juve, R. Mayani, R. Ferreira da Silva, E. Deelman, J. Meredith, J. Vetter, V. Lynch, B. Mayer, J. Wynne III, M. Blanco, C. Carothers, J. LaPre, and B. Tierney, *Toward and End-to-end Framework for Modeling, Monitoring, and Anomaly Detection for Scientific Workflows*, in Workshop on Large-Scale Parallel Processing (LSPP), 2016.
16. R. Ferreira da Silva, M. Rynge, G. Juve, I. Sfiligoi, E. Deelman, J. Letts, F. Würthwein, and M. Livny, *Characterizing a high throughput computing workload: the compact muon solenoid (CMS) experiment at LHC*, in 2015 International Conference on Computational Science (ICCS), 2015.
17. G. Juve, B. Tovar, R. Ferreira da Silva, D. Krol, D. Thain, E. Deelman, W. Allcock, and M.

- Livny, *Practical Resource Monitoring for Robust High Throughput Computing*, Workshop on Monitoring and Analysis for High Performance Computing Systems Plus Applications (HPCMASPA), 2015.
18. T. Glatard, R. Ferreira da Silva, N. Boujelben, R. Adalat, N. Beck, P. Rioux, M. Rousseau, E. Deelman, and A. C. Evans, *Boutiques: an application-sharing system based on Linux containers*, NeuroInformatics 2015, 2015.
 19. R. Oda, D. Cordeiro, R. Ferreira da Silva, E. Deelman, and K. Braghetto, *The case for resource sharing in scientific workflow executions*, in XVI Simposio em Sistemas Computacionais de Alto Desempenho (WSCAD), 2015.
 20. R. Ferreira da Silva, W. Chen, G. Juve, K. Vahi, and E. Deelman, *Community resources for enabling and evaluating research on scientific workflows*, in 10th IEEE International Conference on e-Science, 2014, pp. 177-184.
 21. R. Ferreira da Silva, T. Fahringer, J. J. Durillo, and E. Deelman, *A unified approach for modeling and optimization of energy, makespan and reliability for scientific workflows on large-scale computing infrastructures*, in Workshop on modeling & simulation of systems and applications (MODSIM), 2014.
 22. T. Glatard, L. B. Lewis, R. Ferreira da Silva, M. Rousseau, C. Lepage, P. Rioux, N. Mahani, E. Deelman, and A. C. Evans, *Extending provenance information in cbrain to address reproducibility issues across computing platforms*, Frontiers in Neuroinformatics, iss. 76, 2014.
 23. I. Santana-Perez, R. Ferreira da Silva, M. Rynge, E. Deelman, M. S. Pérez-Hernández, and O. Corcho, *A semantic-based approach to attain reproducibility of computational environments in scientific workflows: a case study*, in Euro-par 2014: parallel processing workshops, 2014, vol. 8805, pp. 452-463.
 24. S. Srinivasan, G. Juve, R. Ferreira da Silva, K. Vahi, and E. Deelman, *A cleanup algorithm for implementing storage constraints in scientific workflow executions*, in 9th workshop on workflows in support of large-scale science (WORKS), 2014, pp. 41-49.
 25. S. Azarnoosh, M. Rynge, G. Juve, E. Deelman, M. Niec, M. Malawski, and R. Ferreira da Silva, *Introducing Precip: an API for managing repeatable experiments in the cloud*, in 2013 IEEE 5th international conference on cloud computing technology and science, 2013, pp. 19-26.
 26. W. Chen, R. Ferreira da Silva, E. Deelman, and R. Sakellariou, *Balanced task clustering in scientific workflows*, in 2013 IEEE 9th International Conference on e-Science, 2013, pp. 188-195.
 27. R. Ferreira da Silva and T. Glatard, *A science-gateway workload archive to study pilot jobs, user activity, bag of tasks, task sub-steps, and workflow executions*, in Euro-par 2012: parallel processing workshops, 2013, vol. 7640, pp. 79-88.
 28. R. Ferreira da Silva, T. Glatard, and F. Desprez, *Workflow fairness control on online and non-clairvoyant distributed computing platforms*, in Euro-Par 2013 parallel processing, 2013, vol. 8097, pp. 102-113.
 29. R. Ferreira da Silva, T. Glatard, and F. Desprez, *On-line, non-clairvoyant optimization of workflow activity granularity on grids*, in Euro-Par 2013 parallel processing, 2013, vol. 8097, pp. 255-266.
 30. R. Ferreira da Silva, G. Juve, E. Deelman, T. Glatard, F. Desprez, D. Thain, B. Tovar, and M. Livny, *Toward fine-grained online task characteristics estimation in scientific workflows*, in 8th workshop on workflows in support of large-scale science (WORKS), 2013, pp. 58-67.
 31. R. Ferreira da Silva, T. Glatard, and F. Desprez, *Self-healing of operational workflow incidents on distributed computing infrastructures*, in 12th IEEE/ACM international symposium on cluster, cloud and grid computing (CCGrid), 2012, pp. 318-325.
 32. T. Glatard, A. Marion, H. Benoit-Cattin, S. Camarasu-Pop, P. Clarysse, R. Ferreira da Silva, G. Forestier, B. Gibaud, C. Lartizien, H. Liebgott, K. Moulin, and D. Friboulet, *Multi-modality image simulation with the virtual imaging platform: illustration on cardiac echography and MRI*, in 9th IEEE international symposium on biomedical imaging, 2012, pp. 98-101.

33. R. Ferreira da Silva, S. Camarasu-Pop, B. Grenier, V. Hamar, D. Manset, J. Montagnat, J. Revillard, J. Rojas Balderrama, A. Tsaregorodtsev, and T. Glatard, *Multi-infrastructure workflow execution for medical simulation in the virtual imaging platform*, in 2011 healthgrid conference, 2011.
34. A. Marion, G. Forestier, H. Benoit-Cattin, S. Camarasu-Pop, P. Clarysse, R. Ferreira da Silva, B. Gibaud, T. Glatard, P. Hugonnard, C. Lartizien, H. Liebgott, S. Specovius, J. Tabary, S. Valette, and D. Friboulet, *Multi-modality medical image simulation of biological models with the virtual imaging platform (VIP)*, in 24th international symposium on computer-based medical systems, 2011, pp. 1-6.
- Book Chapters
1. R. Ferreira da Silva, T. Glatard, and F. Desprez, *Self-managing of operational issues for grid computing: the case of the virtual imaging platform*, in Emerging research in cloud distributed computing systems, S. Bagchi, Ed., IGI Global, 2015.

MEDIA

- *Computer Scientist wins NSF Grant*, <http://www.udel.edu/udaily/2017/december/michela-taufer-NSF-analytics-molecular-dynamics>, UDaily, December 2017.
- *ISI's Pegasus Program Contributed to New Gravitational-Wave Detector Discovery*, <https://www.isi.edu/news/story/323>, USC/ISI, October 2017.
- Nobel Prize-winning discovery on gravitational waves came about with contributions from USC scientists, <http://news.usc.edu/129550/nobel-prize-winning-discovery-on-gravitational-waves-came-about-with-contributions-from-usc-scientists>, USC News, October, 2017.
- *UH Mānoa Professor Casanova receives \$499,000 NSF grant*, <https://www.hawaii.edu/news/2016/09/28/uh-manoa-professor-casanova-receives-499000-nsf-grant/>, University of Hawai'i News, September 2016.