

EDUARDO BERROCAL, PH. D.

Web: <https://eduberrocal.net>
LinkedIn: <https://www.linkedin.com/in/eduberrocal>

EXPERIENCE

- Intel, Hillsboro, OR** Jul/2017 –
Cloud and Enterprise Server Software Engineer
- Bell Labs (Nokia), Murray Hill, NJ** 2016
Summer Internship
Worked on the development of a new Big Data analysis platform as part of a new data-driven approach for helping virtualization of components in future 5G End2End Wireless Networks. In this project, some of the packages of the Berkeley Data Analysis Stack, such as Mesos, Kafka, Cassandra, Marathon and Spark were used, as well as programming in Java.
- Argonne National Laboratory, Argonne, IL** 2013 – 2015
Research Aide (Summers of 2014 and 2015)
Worked in a novel approach to protect HPC applications against Silent Data Corruptions (SDC) in a lightweight manner based on the applications' data behavior. This work produced multiple publications, and work experience in C and Python.
Research Aide (Summer 2013)
Collaborated in the development of the AWE project (a workflow engine coded in Golang), adding new features. Also worked in a bioinformatics project whose goal was to prove that using the wrong software parameters (not tuned) in metagenomics analyses can disrupt critical biological signals in the analyses' results. For the latter, some exposure to the language R was required.
- University of Chicago, Chicago, IL** May/2009 – Aug/2012
Applications' Programmer
Web application developer for Lynx (<http://lynx.ci.uchicago.edu>), a web portal to help geneticists (specially human geneticists) with their informatics needs during their research. This system was built using a backend in Java, frontend in JavaScript, HTML and CSS, databases in MySQL, and remote access to data using web services.
- Universidad Politécnica de Madrid, Madrid, Spain** Jan/2006 – Feb/2007
Internship at Madrid Super Computation and Visualization Center (CESVIMA)
Worked in the Blue Brain Project, a large scale collaboration between industry and multiple universities in Europe, developing software for scientific visualization (neurons' state) in C++.

EDUCATION

- Massive Open On-line Courses (MOOCs)** 2012 – Present
(for a list of courses, see appendix A)
- Ph.D., Computer Science**, Illinois Institute of Technology, Chicago, IL 2012 – 2017
(for a list of courses, see appendix B)

GPA: 3.62/4.0

Awards:

Best Paper Award in *Proceedings of IEEE Cluster'14*, 2014.

Scholarships:

NSF travel grant for ACM HPDC conference, 2015.

NSF travel grant for IEEE Cluster conference, 2014.

Victor Tsao travel Scholarship, 2014.

Research Assistant:

Research has focused on applying data analysis techniques to solve problems in High Performance Computing (HPC) systems. The main project involved the design of a new density-based machine learning algorithm that can predict hardware faults in supercomputers. As part of this work, experience in Python, C, MPI, OpenMP, CUDA, Java, and Hadoop Map-reduce was gained.

Teaching Assistant:

Introduction to Parallel and Distributed Processing (Spring 2015).

Data Integration, Warehousing, and Provenance (Spring 2015).

Advance Operating Systems (Fall 2013, Fall 2014).

Parallel and Distributed Processing (Spring 2013, Spring 2014).

CUDA Teaching Center at IIT:

Part of the CUDA TC at IIT. Responsibilities included building, and maintaining, a 19-node GPU computing cluster used for both teaching and research, as well as teaching introductory lectures on GPGPU programming with CUDA for courses in systems that include parallel programming as part of their curriculum.

M.S., Computer Science (join-degree with UPM), Illinois Institute of Technology, Chicago, IL 2008 – 2009
(for a list of courses, see appendix C)

GPA: 3.59/4.0

M.S., Computer Engineering, Universidad Politécnica de Madrid, Madrid, Spain 2004 – 2009
(for a list of courses, see appendix D)

Awards:

Salutatorian (second of 380 students).

Scholarships:

Scholarship from *Vodafone* to earn a master's degree in the United States or Canada, 2008.

High achievement from the region of Madrid (Spain), 2005.

SKILLS

Computer Languages: C/C++, Java, Python, R.

Web: JavaScript, Java, PHP, HTML, CSS, MySQL.

Parallel Programming: MPI, Pthreads, OpenMP, CUDA.

Berkeley Data Analytics Stack: Hadoop MapReduce, Spark, Hbase, Cassandra, Kafka, Mesos, Marathon.

Operating Systems: Windows, Unix (MacOS, GNU/Linux).

Office: MS Office, LibreOffice, LaTeX.

Environment: Eclipse, Git, Maven, Ant, SVN, Trac.

Languages: Spanish (native), English (full professional proficiency).

ACTIVITIES AND COMMUNITY

Student Volunteer (*SCinet program*) for SC conference 2016 in Salt Lake City, Utah.

Technical Reviewer for Transactions on Parallel and Distributed Systems (TPDS), 2016.
Student Volunteer (*SCinet program*) for SC conference 2015 in Austin, Texas.
Technical Reviewer for Transactions on Parallel and Distributed Systems (TPDS), 2015.
Student Volunteer for SC conference 2014 in New Orleans, Louisiana.
Student Volunteer for SC conference 2013 in Denver, Colorado.
Finalist of the Illinois Technology Association (ITA) Programming Challenge, 2012.
Member of the Institute of Electrical and Electronic Engineers (IEEE).
Member of the Association of Computing Machinery (ACM).

PUBLICATIONS

- E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, and F. Cappello, "Toward General Software Level Silent Data Corruption Detection for Parallel Applications," *IEEE Transactions on Parallel and Distributed Systems* (issue 99), 2017.
- E. Berrocal, and Z. Lan, "Improving Fault Tolerance for Extreme Scale Systems," *Doctoral Showcase at SC'16 (poster)*, 2016.
- J. Wu, X. Xiong, E. Berrocal, J. Wang, and Z. Lan, "Topology mapping of irregular parallel applications on torus-connected supercomputers," *The Journal of Supercomputing* (2016), doi:10.1007/s11227-016-1876-7, 2016.
- E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, and F. Cappello, "Exploring Partial Replication to Improve Lightweight Silent Data Corruption Detection for HPC Applications," *Proc. of EuroPar'16*, 2016.
- E. Berrocal, L. Yu, S. Wallace, X. Yang, Z. Zhou, X. Wang, and Z. Lan, "SPeaR: Toward Smart HPC Through Active Learning and Intelligent Scheduling," *Proc. of HPDC'2015 (poster)*, 2015.
- Y. Wang, W. Shi, and E. Berrocal, "On Performance Resilient Scheduling for Scientific Workflows in HPC Systems with Constrained Storage Resources," *6th Workshop on Scientific Cloud Computing (part of Proc. Of HPDC'15)*, 2015.
- E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, and F. Cappello, "Lightweight Silent Data Corruption Detection Based on Runtime Data Analysis for HPC Applications," *Proc. of HPDC'15 (short paper)*, 2015.
- S. Di, E. Berrocal, and F. Cappello, "An Efficient Silent Data Corruption Detection Method with Error-Feedback Control and Even Sampling for HPC Applications," *Proc. of CCGRID'15*, 2015.
- S. Di, E. Berrocal, L. Bautista-Gomez, K. Heisey, R. Gupta, and F. Cappello, "Toward Fault Detection of Silent Data Corruptions for HPC Applications," *Proc. of SC'14 (poster)*, 2014.
- E. Berrocal, L. Bautista-Gomez, S. Di, Z. Lan, and F. Cappello, "Lightweight Silent Data Corruption Detection Based on Runtime Data Analysis for HPC Applications," *Report, Argonne National Laboratory*, 2014.
- E. Berrocal, L. Yu, S. Wallace, M. E. Papka, and Z. Lan, "Exploring Void Search for Fault Detection on Extreme Scale Systems," (*Best Paper Award*), *Proc. of IEEE Cluster'14*, 2014.
- D. Sulakhe, S. Balasubramanian, B. Xie, B. Feng, A. Taylor, S. Wang, E. Berrocal, U. Dave, J. Xu, D. Bornigen, T. Conrad Gilliam, and N. Maltsev, "Lynx: a database and knowledge extraction engine for integrative medicine", *Nucleic Acids Research*, December 2013.
- J. Wu, Y. Yu, E. Berrocal, and Z. Lan, "Toward Petascale Cosmology Simulations", *2nd Annual Greater Chicago Area Systems Research Workshop (GCASR)*, (*poster*) May 2013.
- D. Sulakhe, S. Balasubramanian, B. Xie, E. Berrocal, B. Feng, A. Taylor, B. Chitturi, U. Dave, G. Agam, J. Xu, D. Bornigen, I. Dubchak, T. Conrad Gilliam, and N. Maltsev, "High-throughput translational medicine: challenges and solutions", *In book: Systems Analysis of Human Multigene Disorders*, Pages 39-67 (Chapter 3).

TALKS AND SEMINARS

- E. Berrocal, and Z. Lan, “Improving Fault Tolerance for Extreme Scale Systems,” *Doctoral Showcase at SC’16*, 2016.
- E. Berrocal, “Research Testbed for a 5G Distributed Native Analytics Framework”, *Bell Labs (Nokia)*, Aug 2016.
- E. Berrocal, “Studying the Impact of Partial Duplication on Silent Data Corruption Detection for HPC Applications”, *Argonne National Laboratory*, Aug 2015.
- E. Berrocal, “Detecting Silent Data Corruption through Time Series Prediction in HPC Applications”, *Argonne National Laboratory*, Aug 2014.

APPENDIX A

MOOCs courses.

- Statistical Learning*, Stanford University, 2016
Grade: Ongoing
- Machine Learning With Big Data*, University of California San Diego, 2016
Grade: 93.3%
- An Entire MBA in 1 Course*, Prof. Chris Haroun, 2016
Grade: Completed
- R Programming*, Johns Hopkins University, 2016
Grade: 94%
- Econometrics: Methods and Applications*, Erasmus University Rotterdam, 2015
Grade: 93%
- Heterogeneous Parallel Programming*, University of Illinois at Urbana-Champaign, 2012-2013
Grade: 100% with Distinction

APPENDIX B

Courses taken during my Ph.D. at Illinois Institute of Technology.

Design and Analysis of Algorithms (for Ph.D.)	Doctoral Seminar
Machine Learning	Data-Intensive Computing
Cloud Computing	Advanced Database Organization
Probabilistic Graphical Models	Research and Thesis Ph.D.

APPENDIX C

Courses taken during my M.S. at Illinois Institute of Technology.

Formal Languages	Advanced Operating Systems
Computer Networks I	Design and Analysis of Algorithms
Information Security	Software Systems Architecture
Science of Programming	Software Project Management
Parallel and Distributed Processing	Research and Thesis M.S.

APPENDIX D

Courses taken during my M.S. at Universidad Politécnica de Madrid.

Infinitesimal Calculus	Computer Structure Laboratory	Software Engineering I
Programming Methodology	Statistical Inference	Compilers
Discrete Math	Computational Logic	Network Architectures
Physics	Personal Computers Laboratory	Operating System Design
Linear Algebra	Domotics and Intelligent Buildings	Databases
Foundations of Computer Hardware	Numerical Calculus	Translating Computing Texts
Formal Logic	Operations Research	Distributed Operating Systems
General Computing	Concurrent Programming	Logic Programming Extensions
Theoretical Computer Science	English for CS I	Computer Systems

Probabilities and Statistics
Computer Technology
Computer Structure
Data Structures I
Mathematical Analysis
Data Structures II
Systematic Program Development
Operating Systems

Program Development Models
Digital System Design
Structuring Logic Design
Logic Programming
Windows Administration
Computer Networks
Computer Architecture
Artificial Intelligence

Knowledge Engineering
Software Engineering II
Multiprocessor Architectures
Information Protection
Distributed Systems:
 Communications Architecture
Discrete Control Systems Design
Fault Tolerance in Computers