

Professional Interests

Systems and applications for *data intensive computing at large scale*, including computational databases, parallel and distributed systems, scalable I/O and indexing techniques for large multi-dimensional spatial datasets, data compression and visualization.

Education

Ph.D. Electrical and Computer Engineering. Carnegie Mellon University, Pittsburgh PA, USA. Aug 2007.
Thesis: “*Methods for querying compressed wavefields*”.

M.Sc. Electrical and Computer Engineering. Carnegie Mellon University, Pittsburgh PA, USA. May 2000.
Thesis: “*Dv — A framework for remote visualization*”.

B.S. Computer Science (Ingeniero de Sistemas). Universidad EAFIT, Medellín, Colombia. June 1996.
Research Project: “*Mercury — A protocol for communications in distributed virtual environments*”.

Professional Experience

Carnegie Mellon University. Pittsburgh, PA, USA. Oct 2007 – present.

Systems Faculty, Computer Science Department. Research in Data Intensive Scalable Computing (DISC) for eScience. Architect and technical leader of the PDL data-intensive clusters (Open Cloud and Open Cirrus).

Hewlett-Packard Laboratories. Palo Alto CA, USA. Jun 2000 – Sep 2000.

Research intern: Implemented a mechanism for aggregating bandwidth of the long-range Internet connections of multiple portable devices such as laptops and PDAs.

Massachusetts Institute of Technology (MIT) / Center for Educational Computing Initiatives (CECI). Boston MA, USA. Aug 1997 – Jul 1998.

Visiting researcher: Designed and implemented a caching architecture for low bandwidth networks.

Universidad EAFIT, Medellín, Colombia. Jun 1995 – Jul 1997.

Adjunct instructor, Software engineer in “Proyecto Conexiones”, system and network administrator.

Corporación Nacional de Ahorro y Vivienda (CONAVI). Medellín, Colombia. Jan 1995 – Jun 1995.

EDP auditor: Corporate network performance and security evaluation.

Selected Publications

Kai Ren, Julio López and Garth Gibson. “*Otus: Resource Attribution and Metrics Correlation in Data-Intensive Clusters*”. ACM/IEEE International Workshop on MapReduce and its Applications. June 2011.

Bin Fu, Kai Ren, Julio López, Eugene Fink, and Garth Gibson. “*DiscFinder: A data-intensive scalable cluster finder for astrophysics*”. ACM Int. Symp. on High Performance Distributed Computing (HPDC), June, 2010.

Julio López, L. Ramirez-Guzman, J. Bielak, and D. O’Hallaron. “*BEMC: A Searchable, Compressed Representation for Large Seismic Wavefields*”. Scientific and Statistical Database Management (SSDBM), 2010.

S. Schlosser, M. Ryan, R. Taborda, J. López, D. O’Hallaron, and J. Bielak, “*Materialized community ground models for large-scale earthquake simulation*”, Supercomputing (SC’08), Nov 2008.

M. Mesnier, M. Wachs, J. López, R. Sambasivan, J. Hendricks and G. Ganger. “*||TRACE — Parallel trace replay with approximate causal events*”. File And Storage Technologies Conf. (FAST). 2007

S. Papadomanolakis, A. Ailamaki, J. López, T. Tu, D. O’Hallaron, G. Heber. “*Efficient query processing on unstructured tetrahedral meshes*”. SIGMOD Conf., 2006.

Additional Information

Programming: Hadoop, HBase, MPI, C/C++, Java, Python, Perl, Bash, Linux, SQL.

Implemented the parallel I/O in the CMU Quake ground-motion simulation numerical solver, which runs on various super computing platforms with tens of thousands of processors, including Cray XT5.

Supercomputing Conference 2006 Analytics Challenge award and Supercomputing 2003 Gordon Bell Prize for special achievement awarded to the CMU Quake project.

Professional affiliations: IEEE, ACM.

Languages: proficient English, native Spanish.