Jose Manuel Faleiro

Department of Computer Science
Yale University
51 Prospect St
New Haven, CT 06511

www.jmfaleiro.com
jose.faleiro@yale.edu
+1 (203) 676 9255

Research Interests

I am interested in data management systems, multi-core systems, and distributed systems.

My dissertation research focussed on using deterministic transaction scheduling to address fundamental bottlenecks in state-of-the-art database systems on modern hardware. Additionally, I have worked on a variety of projects on multi-core and distributed systems. More recently, I have become interested in combining ideas from the database systems and program languages communities; particularly, in using database systems techniques to parallelize program analysis.

Education

Yale University *Aug 2012 - July 2018 (expected)*
PhD Computer Science
Advised by Daniel J. Abadi

Birla Institute of Technology and Science, Pilani, India *Aug 2007-Jun 2011*
B.E.(Honors) Computer Science

Awards and Honors

- *Alan J. Perlis Fellowship 2013-2014* at Yale University.
- *Alan J. Perlis Fellowship 2012-2013* at Yale University.
- *Microsoft Research Tech Transfer Award*
- *BITS Merit Scholarship, Spring 2010*

Projects

Multi-core database systems

- *Multi-core database replication*, a mechanism for exploiting multi-core parallelism while replaying updates in log-shipping replication protocols.
- *Early write visibility*, a recoverability mechanism for deterministic databases that permits transactions to safely expose writes prior to the end of their execution.
- *Multi-core synchronization tradeoffs*, a study which finds that lock-free and lock-based algorithms suffer from the same bottlenecks on multi-core hardware.
- *Orthrus*, a multi-core database system which uses explicit message-passing as a communication primitive to bound contention for shared data-structures.
- *Bohm*, a serializable deterministic multi-version concurrency control protocol which ensures that reads never block writes on update transactions.
- *Lazy transaction evaluation*, a deterministic scheduling mechanism that defers the execution of transactions in order to improve data locality and reduce peak provisioning requirements.

Distributed systems

- *FlexLog*, a partially ordered shared log for building stateful distributed control plane services.

- *Indy*, a key-value store which maintains state in lattice-based data-structures to scale across both multi-core hardware and distributed systems.
- *Geo-distributed actor directory*, an optimistic linearizable naming mechanism for geo-distributing virtual actors in Microsoft Orleans.
- *Fairness-Isolation-Throughput tradeoff*, guidelines for reasoning about when distributed transactions can limit a system's achievable throughput.

Pre-PhD

- *Generalized Lattice Agreement*, a wait-free distributed coordination mechanism to serve consistent reads in eventually consistent database systems.

PUBLICATIONS

High Performance Transactions via Early Write Visibility
Jose M. Faleiro, Daniel J. Abadi, Joseph M. Hellerstein
VLDB 2017

Geo-Distribution of Actor-Based Services
Philip A. Bernstein, Sebastian Burckhardt, Sergey Bykov, Natacha Crooks, Jose M. Faleiro, Gabriel Kliot, Alok Kumbhare, Muntasir Raihan Rahman, Vivek Shah, Adriana Szekeres, Jorgen Thelin
OOPSLA 2017

Latch-free Synchronization in Database Systems: Silver Bullet or Fool's Gold?
Jose M. Faleiro, Daniel J. Abadi
CIDR 2017

Design Principles for Scaling Multi-core OLTP Under High Contention
Kun Ren, Jose M. Faleiro, Daniel J. Abadi
SIGMOD 2016

Rethinking serializable multiversion concurrency control
Jose M. Faleiro, Daniel J. Abadi
VLDB 2015

FIT: A Distributed Database Performance Tradeoff
Jose M. Faleiro, Daniel J. Abadi
IEEE Data Engineering Bulletin, March 2015

Lazy Evaluation of Transactions in Database Systems
Jose M. Faleiro, Alexander Thomson, Daniel J. Abadi
SIGMOD 2014

Generalized Lattice Agreement
Jose M. Faleiro, Sriram Rajamani, Kaushik Rajan, Ganesan Ramalingam, Kapil Vaswani
PODC 2012

CScale - A Programming Model for Scalable and Reliable Distributed Applications
Jose Faleiro, Sriram Rajamani, Kaushik Rajan, Ganesan Ramalingam, Kapil Vaswani
Monterey Workshop 2012

UNDER REVIEW

The FlexLog: A Partially Ordered Shared Log
Joshua Lockerman, Jose M. Faleiro, Juno Kim, Soham Sankaran, Daniel J. Abadi, James Aspnes, Mahesh Balakrishnan

Indy: A KVS for any Scale
Chenggang Wu, Jose M. Faleiro, Yihan Lin, Joseph M. Hellerstein

TALKS

Addressing the Parallelism Gap in Replication Server Systems
StrangeLoop, September 2017

High Performance Transactions via Early Write Visibility
VLDB, August 2017

Modern Transaction Processing: Addressing the Elephants in the Room
Apple, April 2017
Harvard, March 2017

Latch-free Synchronization in Database Systems: Silver Bullet or Fool's Gold?
CIDR, January 2017

The Implications of Multi-core Synchronization on DBMS Implementations
Design Principles for Scaling Multi-core OLTP Under High Contention
Amazon Web Services Palo Alto, August 2016
SIGMOD, June 2016

Demystifying Distributed Transactions with the Fairness-Isolation-Throughput Trade-off
RICON, November 2015

Transaction Processing is Easy if You're God (short talk)
HPTS, September 2015

Rethinking Serializable Multiversion Concurrency Control
Multi-versioning in Main-memory Databases: Limitations and Opportunities
North East Database Day, MIT, January 2016
VLDB, September 2015
Stanford, August 2015
UC Berkeley, June 2015

Lazy Evaluation of Transactions in Database Systems
Yale University, October 2014
Microsoft Research Redmond, June 2014
SIGMOD, June 2014

Towards a Characterization of Mutual Exclusion Locking Primitives
Microsoft Research India, December 2013
Yale University, September 2013
Microsoft Research Silicon Valley, August 2013

SERVICE

- PC Member PaPoC 2018 (co-located with EuroSys '18)
- PC Member PaPoC 2016 (co-located with EuroSys '16)
- External Reviewer SIGMOD 2016, SIGCOMM 2016, VLDB 2015, SIGMOD 2014, SOSP 2013

Teaching Experience

Database System Implementation and Architectures (CS 438) *Spring 2016*
Responsible for creating course syllabus, occasional lecturing, creating assignments, proposing and mentoring student projects.

Operating Systems (CS 422) *Spring 2014*
Responsible for creating and grading exams, occasional lecturing, grading, and office hours.

Programming & Entrepreneurship (CS 113) *Fall 2014*
Responsible for grading exams and evaluating student projects.

Software Engineering (CS 439) *Spring 2014*
Responsible for evaluating student projects.

Introduction to Computer Science (CS 201) *Fall 2013*
Responsible for weekly office hours and grading exams.

Research Mentoring

PhD students

- Juno Kim, Yale *12/2016 - 05/2017*
  I mentored Juno on building applications on top of the FlexLog system.

- Joshua Lockerman, Yale *10/2016 - present*
  I mentored Joshua on the design and implementation of the FlexLog system, writing, and experimental evaluation.

- Chenggang Wu, UC Berkeley *10/2016 - 11/2017*
  I worked with Chenggang on the Indy key-value store, and advised him on design and implementation, writing, and experimental evaluation.

Undergraduates

- Matthew Brady, Yale *09-12/2016*
  Transaction scheduling mechanisms for deterministic database systems.

- David Hatch, Yale *09-12/2016*
  Asynchronous logging for deterministic database systems.

- Kshitijh Meelu, Yale *09-12/2016*
  Dynamic thread allocation mechanism to switch threads between various database system components based on utilization.

- Stanislaw Swidwinski, Yale *09/2016 - 06/2017*
  Transaction re-ordering mechanism to reduce the impact of conflicts on database throughput.

- Sylvan Zheng, Yale *09-12/2016*
  Parallel read- and write-set analysis for deterministic database systems.

Industry Experience

Contractor *Sep 2017 - present*
Facebook
I am building parallel log replay and read-only snapshot mechanisms in MySQL to address replication lag in primary-backup deployments.

Research Intern *Jun 2014 - Aug 2014*
Microsoft Research Redmond
I worked on geo-distributing Orleans, Microsoft's distributed actor-based cloud programming framework.

Research Intern *Jun 2013 - Aug 2013*
Microsoft Research Silicon Valley
Devised and evaluated a lightweight instrumentation technique to correlate poor performance in parallel programs with contention induced back-offs in lock implementations.

Research Developer *Jul 2011 - Jul 2012*
Microsoft Research India
Designed and implemented a distributed programming model built on commutative replicated data-types (CRDTs), a class of eventually consistent distributed data-structures.

Research Intern *Jan 2011 - Jun 2011*
Microsoft Research India
Devised and implemented a semantically rich graph-based representation of a software repository to leverage history to help developers fix bugs.

REFERENCES

Daniel J. Abadi
Darnell-Kanal Professor of Computer Science
University of Maryland

Mahesh Balakrishnan
Associate Professor of Computer Science
Yale University

Alan Fekete
Professor, School of Information Technologies
University of Sydney

Joseph M. Hellerstein
Jim Gray Professor of Computer Science
UC Berkeley