

Cheng Li

Present Address

MPI-SWS
Campus E1 4
D-66123 Saarbruecken, Germany

Contact Info

<http://www.mpi-sws.org/~chengli/>
(+49)176-62342761
chengli@mpi-sws.org

Interests

My research interest primarily focuses on building dependable and scalable distributed systems. In my dissertation work, I have developed principles and techniques to address the tension between maintaining strong consistency semantics and improving performance in replicated services.

In the past, I also developed techniques to detect concurrency bugs.

Education

Max Planck Institute for Software Systems (MPI-SWS) 2009 - 2015 (Expected)
Saarland University (UdS)
PhD candidate in Computer Science
Advisors: Rodrigo Rodrigues, Allen Clement

University of California, Los Angeles (UCLA) Summer 2008
Exchange bachelor student in Computer Science
Advisor: Mario Gerla

Nankai University (NKU), Tianjin, China 2005 - 2009
B.S. in Computer Science

Research Projects

- Consistency v. performance in replicated systems:** 2011 - Present
- *RedBlue Consistency* [OSDI'12] allows us to build scalable replicated systems by supporting the coexistence of strongly consistent (slow but crucial) operations and weakly consistent (fast) operations. Alongside this model, we provided programmers with a set of classification principles to associate various operations with their appropriate consistency level. In addition, we built a distributed storage infrastructure to replicate a few applications with RedBlue Consistency and evaluate what are the performance benefits when compared to always using strong consistency, and what is the overhead when compared to only using weak consistency.
 - *SIEVE* [ATC'14] is a lightweight tool that automatically assigns different consistency levels to various operations by requiring only a small amount of input from developers. The core of SIEVE is a combination of static analysis and runtime verification to detect the potential of breaking application-specific invariants.
 - *PoR Consistency* [PaPoC'15] generalizes the trade-offs between performance and consistency semantics by connecting consistency semantics to visibility restrictions between a pair of operations. It allows programmers to add or remove visibility restrictions among operations to strengthen or weaken consistency semantics.
 - *Eventually consistent relational database* leverages Commutative Replicated Data Types (CRDTs) to make every SQL statement commute for a possible eventually consistent replication, while preserving integrity constraints.

Fault tolerance at data center scale:

2012 - 2014

- *Visigoth fault tolerance* [EuroSys'15] is a new technique for designing resource-efficient distributed protocols for building reliable stateful services. It leverages the characteristics of data center systems, namely their secure environment and predictable performance, and hence makes less pessimistic assumptions than traditional solutions.

Concurrency bugs:

2009 - 2011

- We conducted a comprehensive *concurrency bug study* [DSN'10] and discovered two interesting types of bugs that receive little attention, namely, a) semantic bugs that produce plausible but incorrect results; and b) latent bugs that silently corrupt internal state when they first occur.
- *Pike* [EuroSys'11] is a concurrency bug detector that effectively finds semantic and latent bugs, by checking, given two requests, whether the output and the final internal state of a concurrent execution are equivalent to these of either of the two sequential executions.

Publications

- *WeaQL: Scaling Relational Databases through Weak Consistency*.
David Lopes, Cheng Li, João Leitão, Nuno Preguiça, and Rodrigo Rodrigues.
Submitted to ICDCS'16
- *Minimizing Coordination in Replicated Systems*.
Cheng Li, João Leitão, Allen Clement, Nuno Preguiça, and Rodrigo Rodrigues.
PaPoC'15, Apr. 2015, Bordeaux, France
- *Visigoth Fault tolerance*.
Daniel Porto, João Leitão, Cheng Li, Allen Clement, Aniket Kate, Flavio Junqueira, and Rodrigo Rodrigues.
EuroSys'15, Apr. 2015, Bordeaux, France
- *Lower Bound and Correctness Proofs for Consensus in the Visigoth Model*.
Daniel Porto, João Leitão, Cheng Li, Allen Clement, Aniket Kate, Flavio Junqueira, and Rodrigo Rodrigues.
Technical Report, Nova University of Lisbon, 2015
- *Automating the Choice of Consistency Levels in Replicated Systems*.
Cheng Li, João Leitão, Allen Clement, Nuno Preguiça, Rodrigo Rodrigues and Viktor Vafeiadis.
USENIX ATC'14, Jun. 2014, Philadelphia, PA, USA
- *Making Geo-Replicated Systems Fast as Possible, Consistent when Necessary*.
Cheng Li, Daniel Porto, Allen Clement, Rodrigo Rodrigues, Nuno Preguiça, and Johannes Gehrke.
OSDI'12, Oct. 2012, Hollywood, CA, USA
- *Finding Complex Concurrency Bugs in Large Multi-Threaded Applications*.
Pedro Fonseca, Cheng Li, and Rodrigo Rodrigues.
EuroSys'11, Apr. 2011, Salzburg, Austria
- *A study of the Internal and External Effects of Concurrency Bugs*.
Pedro Fonseca, Cheng Li, Vishal Singhal and Rodrigo Rodrigues.
DSN'10, Jun. 2010, Chicago, USA

Conference Talks

- **Minimizing Coordination in Replicated Systems**
EuroSys/PaPoC'15, Apr. 2015, Bordeaux, France
- **Automating the Choice of Consistency Levels in Replicated Systems**
USENIX ATC'14, Jun. 2014, Philadelphia, PA, USA

- **Making Geo-Replicated Systems Fast as Possible, Consistent when Necessary**
OSDI'12, Oct. 2012, Hollywood, CA, USA

Invited Talks

- **Minimizing Coordination in Geo-Replicated Systems**
INESC-ID/IST, Jul. 2015, Lisbon, Portugal
- **Consistency Trade-offs in Geo-Replicated Systems**
Oracle Labs, Jun. 2015, Zurich, Switzerland
- **Automating the Choice of Consistency Levels in Replicated Systems**
IMDEA Software Institute, Jun. 2015, Madrid, Spain
- **Building Fast and Consistent (Geo-)Replicated Systems**
Facebook, Apr. 2015, Menlo Park, CA, USA
Huawei European Research Center, Feb. 2015, Munich, Germany
IBM Research China, Feb. 2014, Beijing, China
Microsoft Research Asia, Feb. 2014, Beijing, China
- **Sub-second N-Hop Graph Traversal Queries**
Microsoft Research Cambridge, Oct. 2013, UK
- **SIEVE: A Tool for Making Geo-Replicated Services Fast and Consistent**
CITI Seminar at NOVALINCS, May 2013, Lisbon, Portugal

Industry Experience

Intern, Microsoft Research, Cambridge, UK Summer 2013
I worked with Flavio Paiva Junqueira and Miguel Castro to re-design breadth-first search (BFS) to be a fully decentralized algorithm, which then can be used as a primitive to build fast traversal queries over large-scale social graphs.

Intern, Visto Corporation (Good Technology), Tianjin, China 11/2008 - 3/2009
I worked with another intern to implement a web information system, which manages many devices across a few departments.

Honors

USENIX Student Travel Grant	6/2014
Excellent bachelor award	6/2009
Best bachelor thesis, Nankai University 2009	6/2009
China Government Scholarship	9/2008
Cross-disciplinary scholarship in Science and Technology initiative UCLA	7/2008
Honorable Mention, 2008 American Mathematical Contest in Modeling	2/2008

Activities

- Master thesis mentor, Nova Uni. 2014 - Present
I informally co-advise a master student, David Lopes, on his master thesis project, in which he extends the SIEVE work to create CRDTs covering the core features of SQL and to ensure integrity constraints atop of an eventually consistent relational database.
- Teaching assistant, Operating systems course, UdS 4/2011 - 10/2011
Responsible for course research project advising, tutorials, and exam creation
- Computer systems seminar organizer, Nova Uni. 3/2014 - 7/2014
Responsible for proposing the new rule and format, organizing weekly meetings, and assisting in paper selection
- Volunteer at USENIX FCW'14 6/2014
Summarized talks in two sessions of *HotCloud'14* for the magazine *USENIX ;login:*
- Student office assistant of CS Department, NKU 9/2007 - 2/2009
- Undergraduate Student Representatives of CS Department, NKU 9/2006 - 6/2008

References **Rodrigo Rodrigues (advisor)**
Professor of Computer Science, IST, Portugal
Former faculty, MPI-SWS, Germany
rodrigo.miragaia.rodrigues@tecnico.ulisboa.pt

Allen Clement (co-advisor)
Software Engineer at Google, Zurich
Former faculty, MPI-SWS, Germany
aclement@gmail.com

Johannes Gehrke (collaborator)
Tisch University Professor of computer science, Cornell University, USA
Distinguished Engineer at Microsoft
johannes@cs.cornell.edu

Nuno Preguiça (collaborator)
Associate professor of Computer Science, Universidade Nova de Lisboa, Portugal
nuno.preguica@fct.unl.pt