

Nicolas Schiper
phone : +41.21.702.5567
nicolas.schiper@gmail.com

Education

- Ph.D.** Informatics, University of Lugano (USI), Switzerland, October 2009.
Thesis title: *On Multicast Primitives in Large Networks and Partial Replication Protocols*
Supervisor: Fernando Pedone
- M.Sc.** Computer Science, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, April 2005.
Supervisor: Sam Toueg, University of Toronto, September–December 2004.

Employment History

- 02/2012 – 05/2014 **Postdoctoral associate**
Cornell University, USA
- 12/2010 – 12/2011 **Software Engineer**
Google, Zürich, Switzerland
- 11/2009 – 10/2010 **Postdoc**
USI, Switzerland
- 07/2007 – 08/2007 **Research Intern**
Microsoft Research, Beijing, China
- 10/2005 – 10/2009 **Research and Teaching Assistant**
USI, Switzerland
- 03/2005 – 05/2005 **Software Engineer Intern**
ELCA, Switzerland

Expertise and Technical Knowledge

Expertise: Distributed systems, transactional systems, fault-tolerance, group communication, middleware
benchmarking

Programming languages: Java / C / Scala / SQL

Databases: HSQLDB / H2 / MySQL / Derby

Development environment: Git / Eclipse

OS: Linux system programming (inter-process synchronization & communication, networking, asynchronous I/O)

Working Experience

- Cornell University Development of energy-efficient database replication protocols and measurement of their efficiency on a Linux cluster under the TPC-C and Twitter benchmarks. The protocols include state machine replication, primary-backup, and deferred-update replication (~13 KLOCs when excluding the benchmarks and JPaxos).
Design and implementation of ShadowDB (~5 KLOCs), the first replicated database whose code comes with formal correctness guarantees. ShadowDB is available as open-source at <https://github.com/nschiper/ShadowDB>.

Google	Improvement of the spam console of Google Trader, a classified for emerging markets. Integration of the console with automatic spam detection tools.
Microsoft Research	Continuation of the research on database replication protocols that support strong data consistency and data sharding.
University of Lugano	Research on scalable database replication protocols and group communication protocols. Implementation of atomic multicast protocols optimized for WANs and P-Store (a few KLOCs each), one of the first transactional key-value store supporting data sharding. Some of these protocols are available as open-source and are now maintained by Pierre Sutra: https://github.com/otrack/Batelier Evaluation carried out on a Linux cluster using Modelnet to emulate WAN links, and under the TPC-B and TPC-C benchmarks. Development of the first prototype of Ring-Paxos, a super high-throughput Paxos protocol.
ELCA	Development of a plugin to cache the result of RMI calls in the Spring framework.

Honors and Awards

1. Best paper award at the 5th Latin-American Symposium on Dependable Computing (LADC 2011) for “Byzantine Fault-Tolerant Deferred Update Replication” with F. Pedone and J. E. Armendariz-Inigo.

Publications

Journals

1. R. Van Renesse, N. Schiper, and Fred B. Schneider. Vive la Différence: Paxos vs. Viewstamped Replication vs. Zab. *To appear in IEEE Transactions on Dependable and Secure Computing*.
2. F. Pedone and N. Schiper. Byzantine Fault-Tolerant Deferred Update Replication. *Journal of the Brazilian Computer Society*, 2012, pp. 7-16.

Conferences, Symposia, and Workshops

1. N. Schiper, F. Pedone, R. van Renesse. The Energy Efficiency of Database Replication Protocols. *To appear in 44th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'2014)*.
2. N. Schiper, V. Rahli, R. van Renesse, M. Bickford, R. L. Constable. Developing Correctly Replicated Databases using Formal Tools. *To appear in 44th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'2014)*.
3. S. Nikolau, R. Van Renesse, and N. Schiper. Cooperative client caching strategies for social and web applications. *In 7th Workshop on Large Scale Distributed Systems and Middleware*, 6 pages (online).
4. R. Van Renesse, C. Ho, and N. Schiper. Byzantine Chain Replication. *In 16th International Conference on Principles of Distributed Systems (OPODIS 2012)*, pp. 345-359. **acceptance rate**: 26.9% (24/89).
5. V. Rahli, N. Schiper, R. Van Renesse, M. Bickford, and R. L. Constable. A Diversified and Correct-by-Construction Broadcast Service. *In 2nd Workshop on Rigorous Protocol Engineering (WRiPE 2012)*, pp. 1-6.
6. N. Schiper, V. Rahli, R. Van Renesse, M. Bickford, and R. L. Constable. ShadowDB: A Replicated Database on a Synthesized Consensus Core. *In 8th Workshop on Hot Topics in System Dependability (HotDep 2012)*, 6 pages (online).
7. P. Raykov, N. Schiper, and F. Pedone. Byzantine Fault-Tolerance with Commutative Commands. *In 15th International Conference On Principles Of Distributed Systems (OPODIS 2011)*, pp. 320-342. **acceptance rate**: 27.1% (26/96).

8. F. Pedone, N. Schiper, and J. E. Armendáriz-Iñigo. Byzantine Fault-Tolerant Deferred Update Replication. In *5th Latin-American Symposium on Dependable Computing (LADC'2011)*, pp. 7-16. **Best paper award.**
acceptance rate: 34.7% (16/46).
9. N. Schiper, P. Sutra, and F. Pedone. P-Store : Genuine Partial Replication in Wide Area Networks. In *29th IEEE International Symposium on Reliable Distributed Systems (SRDS'2010)*, pp. 214-224. **acceptance rate:** 27.2% (21/77).
10. P. J. Marandi, M. Primi, N. Schiper, and F. Pedone. Ring-Paxos : A Highly Efficient Atomic Broadcast Protocol. In *40th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'2010)*, pp. 527-536.
acceptance rate: 23.9% (65/272).
11. N. Schiper and F. Pedone. Fast, Flexible, and Highly Resilient Genuine Fifo and Causal Multicast Algorithms. In *25th Symposium on Applied Computing (SAC'2010)*, pp. 418-422. **acceptance rate:** 26.9%.
12. N. Schiper, P. Sutra, and F. Pedone. Genuine versus Non-Genuine Atomic Multicast Protocols for Wide Area Networks : an Empirical Study. In *28th IEEE International Symposium on Reliable Distributed Systems (SRDS'2009)*, pp. 166-175. **acceptance rate:** 22.1% (23/104).
13. N. Schiper and F. Pedone. Solving Atomic Multicast when Groups Crash. In *12th International Conference on Principles of Distributed Systems (OPODIS'2008)*, pp. 481-495. **acceptance rate:** 29.4% (30/102).
14. N. Schiper and S. Toueg. A Robust and Lightweight Stable Leader Election Service for Dynamic Systems. In *38th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'2008)*, pp. 207-216. **acceptance rate:** 24.5% (58/237).
15. N. Schiper and F. Pedone. On the Inherent Cost of Atomic Broadcast and Multicast Algorithms in Wide Area Networks. In *9th International Conference on Distributed Computing and Networking (ICDCN'2008)*, pp. 145-157. **acceptance rate:** 16.2% (30/185).
16. N. Schiper and F. Pedone. Optimal Atomic Broadcast and Multicast Algorithms for Wide Area Networks, Brief Announcement. In *26th Symposium on Principles of Distributed Computing (PODC'2007)*, pp. 384-385.
17. N. Schiper and F. Pedone. Optimistic Algorithms for Partial Database Replication. In *10th International Conference on Principles of Distributed Systems (OPODIS'2006)*, pp 81-93. **acceptance rate:** 12.2% (28/230).
Also appeared as a Brief Announcement in *20th International Symposium on Distributed Computing (DISC'2006)*, pp. 557-559.

Theses

1. N. Schiper. *On Multicast Primitives in Large Networks and Partial Replication Protocols*. Ph.D. thesis, University of Lugano, Switzerland, 2009.

Research Grants

1. Swift-WAN2: Performance and Dependability for data center applications, Swiss National Science Foundation (2011 - 2012, 102K CHF), co-author with Fernando Pedone.

Teaching Experience

Lecturer

Distributed Algorithms, USI Fall 2009

Teaching assistant

Distributed Algorithms, USI Spring 2006, Fall 2008

Computer Systems Lab, USI Spring 2008

Computer Systems, USI Spring 2007, Spring 2008

Algorithms and Data Structures, USI Spring 2006

Advising Experience

Ph.D. students

At Cornell University

Stavros Nikolaou (under co-supervision with R. Van Renesse).

Qin Jia (co-supervised her Advanced Systems project with R. Van Renesse).

M.Sc. students

At Cornell University

T. Mao (co-supervised with V. Rahli), *Programming an Interpreter from EventML Specifications to Scala*, 2013.

At USI

D. Sciascia (co-supervised with F. Pedone), *dsmDB: Clustering in-memory Database Management Systems*, 2009.

M. Primi (co-supervised with F. Pedone), *Paxos made Code: Implementing a high throughput Atomic Broadcast*, 2009.

References

Dr. Robbert Van Renesse, ACM fellow.

Principal Research Scientist, Department of Computer Science, Cornell University, Ithaca, NY 14853, USA.

E-mail: rvr@cs.cornell.edu

Prof. Fernando Pedone.

Professor of Informatics, Faculty of Informatics, University of Lugano, Switzerland.

E-mail: fernando.pedone@usi.ch

Prof. Sam Toueg, recipient of the 2010 Edsger W. Dijkstra prize in Distributed Computing.

Department of Computer Science, University of Toronto, Ontario, M5S 3G4, Canada.

E-mail: sam@cs.toronto.edu.