

Education	Massachusetts Institute of Technology Sept. 2011 - Present Ph.D. student in Electrical Engineering and Computer Science (EECS) Advised by Professor Samuel Madden Master's thesis: <i>Fast Transactions for Multicore In-Memory Databases</i> GPA: 5.0/5.0
	University of California, Berkeley Aug. 2006 - Dec. 2010 B.A. in Computer Science and B.S. in Mechanical Engineering GPA: 3.97/4.0
Publications	Anti-Caching: A New Approach to Swapping in Main Memory OLTP Database Systems. Justin DeBrabant, Andrew Pavlo, <i>Stephen Tu</i> , Michael Stonebraker, Stan Zdonik <i>To appear in VLDB 2014.</i>
	Speedy Transactions in Multicore In-Memory Databases. <i>Stephen Tu</i> , Wenting Zheng, Eddie Kohler, Barbara Liskov, and Samuel Madden <i>SOSP 2013.</i>
	Processing Analytical Queries over Encrypted Data. <i>Stephen Tu</i> , M. Frans Kaashoek, Samuel Madden, and Nickolai Zeldovich <i>VLDB 2013.</i>
	The HipHop Compiler for PHP. Haiping Zhao, Iain Proctor, Minghui Yang, Xin Qi, Mark Williams, Guilherme Ottoni, Charlie Gao, Andrew Paroski, Scott MacVicar, Jason Evans, and <i>Stephen Tu</i> <i>OOPSLA 2012.</i>
	The case for PIQL: a performance insightful query language. Michael Armbrust, Nick Lanham, <i>Stephen Tu</i> , Armando Fox, Michael Franklin and David Patterson <i>SoCC 2010.</i>
	PIQL: A Performance Insightful Query Language For Interactive Applications. Michael Armbrust, <i>Stephen Tu</i> , Armando Fox, Michael J. Franklin, David A. Patterson, Nick Lanham, Beth Trushkowsky, and Jesse Trutna <i>SIGMOD Demo, 2010.</i>
Research Experience	Silo - MIT CSAIL Spring 2013 - Summer 2013 <ul style="list-style-type: none">• Developed a new epoch based optimistic concurrency control (OCC) algorithm designed for extremely high throughput serializable transactions on modern multicores.• Protocol achieves near linear scalability on popular database benchmarks, and beats out existing commercial and research systems by several factors.
	Monomi - MIT CSAIL Fall 2011 - Summer 2012 <ul style="list-style-type: none">• Developed a system for analytical query processing on encrypted databases. Monomi incurs a modest 2x overhead on the popular TPC-H benchmark in the worst case compared to plaintext (insecure) database, and much less in the average case.• Techniques include an algorithm for optimal split client/server query execution, and an integer linear program (ILP) formulation for workload-driven physical design of encrypted databases.
	PIQL/SCADS - RADLab, UC Berkeley Fall 2009 - Fall 2010 <ul style="list-style-type: none">• Developed a Scala compiler plugin for ORM-like language integration with SCADS (a distributed key/value store).• Implemented various operations for SCADS, including data repartitioning and deletion, and also parallel/async query executors.

Work Experience

Software Engineer Intern - HipHop Compiler Team, Facebook April 2011 - Aug 2011

- Worked on various projects, including parallelizing the type-inference phase of the compiler, improving the performance of closures in the HipHop runtime, more aggressive type propagation, and various other code generation optimizations based on data-flow analysis.

Software Engineer Intern - Datacenters Team, Facebook Jan 2011 - April 2011

- Turned prototype of MySQL database consistency checker into a working system capable of checking all of Facebook's user databases. Found real cases of inconsistencies in replicated data.
- This work was presented at the O'Reilly MySQL Conference in 2011.

Scala Remote Actors - Google Summer of Code Summer 2010

- Worked on an experimental reimplement of the standard remote actors library. Addressed issues of scalability by using a non-blocking asynchronous network architecture and allowing for pluggable serialization frameworks.

Software Engineer Intern - Intuit Summer 2009

- Designed and implemented a rule editor for tax analyst programmers to encode IRS tax specifications in XPath 2.0, used to validate electronic filings.
- Built RESTful web-service which took electronically filed tax documents and validated the input against the rules generated via the editor.

Lead Programmer - RSSP-IT, UC Berkeley Summer 2008 - Summer 2010

- Prototyped transition from legacy `mod_perl` environment to a J2EE stack, using Spring/Hibernate for the web framework and Maven/Nexus for build management and deployment.
- Integrated Calnet Central Authentication Service into existing legacy applications, including registration for in-room connection and helpdesk (for residence hall students).

Relevant Skills

Programming Languages

- C, C++, Java, Scala, Python, PHP, Matlab/Octave, SQL

Systems and Tools

- Basic Linux systems administration, basic MySQL performance tuning knowledge, Git, Subversion, Bison, Flex, ANTLR, \LaTeX

Awards

Honorable Mention - CRA's Outstanding Undergraduate Researcher Award 2011

Github

<https://github.com/stephentu>