Linhai Song

Last update on November 18, 2017

College of Information Sciences and Technology
The Pennsylvania State University
http://songlh.github.io/

Email: songlinhai0543@gmail.com
Alt: songlh@ist.psu.edu

Research Interests

- Performance Optimization & Tuning for Large Software
- Applied Machine Learning
- Program Analysis, Security, and Software Testing

Employment

College of Information Sciences and Technology, PSU — Pennsylvania, USA
Assistant Professor — *2017.08 – Present*

FireEye Research Labs, FireEye, Inc. — California, USA
Staff Research Scientist — *2015.11 – 2017.07*

Education

University of Wisconsin-Madison — Wisconsin, USA
Ph.D. in Computer Science (M.S. along the way) — *2010.08 – 2015.10*
Advisor: Shan Lu
Thesis: Understanding, Detecting, and Diagnosing Real-World Performance Bugs

Institute of Computing Technology, Chinese Academy of Sciences — Beijing, China
M.S. in Computer Science — *2007.08 – 2010.06*
Advisor: Xueqi Cheng

Huazhong University of Science and Technology — Hubei, China
B.E. in Software Engineering — *2003.08 – 2007.06*

Academic Awards

MICRO Best Paper Runner Up, 2014

- "COMP: Compiler Optimizations for Manycore Processors" published in MICRO'2014
- One of five papers selected from 273 MICRO'2014 submissions

ACM SIGPLAN Research Highlights Award, 2011

- "Automated Atomicity-Violation Fixing" published in PLDI'2011
- One of eight papers selected from all papers published in 13 ACM SIGPLAN conferences in 2011

Research Experience

Assistant Professor, PSU — 2017.08 – present

- Design algorithms for automated protocol subsetting and dialect generation;
- Design algorithms for production-run algorithmic profiling;
- Apply deep learning to program analysis;
- Explore performance bottlenecks for GO programming languages.

Staff Research Scientist, FireEye Research Labs — 2015.11 – 2017.07

- Conduct data mining for the security repository on VirusTotal;
- Design and implement an end-point anti-virus system;

- Design and implement algorithms to calculate similarity between JavaScript programs.

Research Assistant, University of Wisconsin-Madison 2011.01 – 2015.10

- Design and implement a series of static-dynamic hybrid analysis for inefficient loops;
- Study the correlation between features of critical sections and their change histories;
- Explore the design space of applying statistical debugging to performance failure diagnosis;
- Implement a dynamic technique to detect inefficient nested loops for C/C++ programs;
- Design and implement a series of static rule-based detectors for performance bugs;
- Conduct a comprehensive study on 110 real-world performance bugs;
- Implement the deadlock detection module in an atomicity violation concurrency bug fixing project.

Research Intern, FutureWei Technologies Inc. 2014.05 – 2014.09

- Demonstrate a static bug detection technique for inefficient loops with Cond-Break fixes;
- Demonstrate a failure diagnosis technique built on hardware performance counters.

Research Intern, NEC Labs America 2013.05 – 2013.08

- Explore performance bottlenecks for Intel Xeon Phi manycore coprocessors (MIC);
- Design and implement three source-to-source compiler optimizations for parallel loops which offload computation to MIC.

Research Intern, Microsoft Research Asia 2010.05 – 2010.07

- Design and implement a toolkit for graphical model inference based on secondary development on Visio;
- Get **excellent** assessment for this project.

Research Assistant, Institute of Computing Technology, Chinese Academy of Sciences 2007.09 – 2010.05

- Design two separate algorithms to extract news articles and blog posts respectively;
- Implement the web content extraction module in a web retrieve system.

Publications

Refereed Conference and Workshop Publications

1. **Linhai Song**, Shan Lu
   Program Analysis for Inefficient Loops, **ICSE'2017**.
2. **Linhai Song**, Heqing Huang, Wu Zhou, Wenfei Wu, Yiying Zhang
   Learning from Big Malware, **APSys'2016**.
3. Rui Gu, Guoliang Jin, **Linhai Song**, Linjie Zhu, Shan Lu
   What Change History Tells Us About Thread Synchronization, **FSE'2015**.
4. **Linhai Song**, Min Feng, Nishkam Ravi, Yi Yang, Srimat Chakradhar
   COMP: Compiler Optimizations for Manycore Processors, **MICRO'2014**.
   Won MICRO'2014 Best Paper Runner Up
5. **Linhai Song**, Shan Lu
   Statistical Debugging for Real-World Performance Problems, **OOPSLA'2014**.
6. Adrian Nistor, **Linhai Song**, Darko Marinov, Shan Lu
   Toddler: Detecting Performance Problems via Similar Memory-Access Patterns, **ICSE'2013**.
7. Guoliang Jin*, **Linhai Song***, Xiaoming Shi, Joel Scherpelz, Shan Lu
   Understanding and Detecting Real-World Performance Bugs, **PLDI'2012**.
   (*: alphabetical order of surnames)
8. Guoliang Jin, **Linhai Song**, Wei Zhang, Shan Lu, Ben Liblit
   Automated Atomicity-Violation Fixing, **PLDI'2011**.
   Won ACM SIGPLAN Research Highlights Award

Other Publications

1. **Linhai Song**, Shan Lu

Program Analysis for Inefficient Loops, UChicago CS **Technical Report TR-2016-06**.

2. Dongdong Deng, Guoliang Jin, Marc de Kruijf, Ang Li, Ben Liblit, Shan Lu, Shanxiang Qi, Jinglei Ren, Karthikeyan Sankaralingam, **Linhai Song**, Yongwei Wu, Mingxing Zhang, Wei Zhang, Weimin Zheng
   Fixing, Preventing, and Recovering from Concurrency Bugs, **Science China Information Sciences, April 2015**.
3. **Linhai Song**, Shan Lu
   Statistical Debugging for Real-World Performance Problems, **GCASR'2015 Poster**.
4. **Linhai Song**, Shan Lu
   Statistical Debugging for Real-World Performance Problems, UW-Madison CS **Technical Report 1803**.

Publications before Ph.D

1. Yan Guo, Huifeng Tang, **Linhai Song**, Yu Wang, Guodong Ding
   ECON: An Approach to Extract Content from Web News Page, **APWeb'2010**.
2. **Linhai Song**, Xueqi Cheng, Yan Guo, Bo Wu, Yu Wang
   Blog Post Extraction Using Title Finding, **CCIR'2009**.
3. Yu Wang, Bingxing Fang, Bo Wu, **Linhai Song**, Yan Guo
   Schema Matching Incorporating with Attribute Distribution Features, **CCIR'2009**.
4. Feng Guan, Xiaoming Yu, Zeying Peng, Hongbo Xu, Yue Liu, **Linhai Song**, Xueqi Cheng
   ICTNET at Web Track 2009 Ad-hoc Task, **TREC'2009**.
5. Xueke Xu, Yue Liu, Hongbo Xu, Xiaoming Yu, **Linhai Song**, Feng Guan, Zeying Peng, Xueqi Cheng
   ICTNET at Blog Track TREC 2009, **TREC'2009**.
6. Bo Wu, Xueqi Cheng, Yu Wang, Yan Guo, **Linhai Song**
   Simultaneous Product Attribute Name and Value Extraction from Web Pages, **WI'2009 workshop**.
7. **Linhai Song**, Xueqi Cheng, Yan Guo, Yue Liu, Guodong Ding
   ContentEx: A framework for automatic content extraction programs, **ISI'2009 short**.

Patents

1. Min Feng, Srimat Chakradhar, **Linhai Song**
   Compiler Optimization for Many Integrated Core Processors, U.S. Patent No. 20150277877, Oct 1st, 2015.

Professional Services

- Reviewer for CCS'2017
- Reviewer for Transactions on Software Engineering
- Reviewer for Usenix ATC'2017
- Reviewer for Journal of Computer Science and Technology
- PC member of Artifact Evaluation session in PLDI'2015
- PC member of Artifact Evaluation session in ISSTA'2014

Talks

1. Performance Diagnosis for Inefficient Loops
   Conference Presentation in ICSE'2017, May 2016
2. Improve Software Security and Performance through Data Analytics
   The Pennsylvania State University, March 2016
3. Learning from Big Malware
   Conference Presentation in APSys'2016, August 2016
4. Understanding, Detecting, and Diagnosing Real-World Performance Bugs
   National University of Singapore, March 2016
5. Understanding, Detecting, and Diagnosing Real-World Performance Bugs
   Microsoft Research Asia, December 2015
6. Understanding, Detecting, and Diagnosing Real-World Performance Bugs

Peking University, June 2015

7. Understanding, Detecting, and Diagnosing Real-World Performance Bugs
   Pivotal Labs, May 2015
8. Statistical Debugging for Real-World Performance Problems
   Conference Presentation in OOPSLA'2014, October 2014
9. Statistical Debugging for Real-World Performance Problems
   WISDOM Workshop II, May 2014
10. Optimizing Memory Performance on Many Integrated Core Coprocessors
    NEC Labs America, August 2013
11. Understanding and Detecting Real-World Performance Bugs
    Conference Presentation in PLDI'2012, June 2012
12. Understanding and Detecting Real-World Performance Bugs
    Programming Languages Seminar, University of Wisconsin-Madison, May 2012

Skills

- **Languages:** C/C++, C#, Java, Python, PHP, SQL, HTML, JavaScript, Bash
- **Instrumentation & Analysis:** LLVM, PIN, GCC, GDB
- **Tools & Libraries:** Pthread, OMP, STL, SVN, GIT, MySQL, SQLite
- **Platforms:** Linux, Windows, Intel MIC