

Wei Xu 徐葳

wei.xu.0@gmail.com

+1 (510) 6121601

EDUCATION

University of California at Berkeley

Ph.D. in Computer Science

Advisors: Prof. David Patterson and Prof. Armando Fox

Dissertation: Detecting Large Scale System Problems by Mining Console Logs

Certificate in Management of Technology (MOT)

Berkeley, CA

Dec. 2005 – July, 2010

University of California at Berkeley

M.S. in Computer Science, GPA: 3.95 / 4.00

Research Advisor: Prof. David Patterson

Berkeley, CA

Aug. 2003- Dec. 2005

University of Pennsylvania

B.S. Eng. in Computer Science, Summa Cum Laude, GPA 4.00 / 4.00

with Electrical Engineering Minor

Philadelphia, PA

Sept. 2001- May. 2003

Tsinghua University

Major in Computer Science

Beijing, China

Sept. 1999 - Aug. 2001

PUBLICATIONS

Advances and Challenges in Log Analysis

Adam Oliner, Archana Ganapathi, Wei Xu

In the *Communications of ACM (CACM) and ACM Queue*, Invited article, Feb, 2012

Experience on Mining Google's Production Console Logs

Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael Jordan

In the *Workshop on Managing Systems via Log Analysis and Machine Learning Techniques (SLAML '10)*, Vancouver, BC Oct, 2010

A graphical representation for identifier structure in logs

Ariel Rabkin, Avani Wildani, Randy Katz, Wei Xu, Armando Fox

In the *Workshop on Managing Systems via Log Analysis and Machine Learning Techniques (SLAML '10)*, Vancouver, BC Oct, 2010

Detecting Large Scale System Problems by Mining Console Logs

Wei Xu

PhD dissertation, UC Berkeley, July, 2010

Using Machine Learning Techniques in Console Log Analysis

Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael Jordan

Invited paper in the *27th International Conference on Machine Learning (ICML '10)*, Haifa, Israel, June 2010

Online System Problem Detection by Mining Patterns of Console Logs

Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael Jordan

In Proc. of the *IEEE International Conference on Data Mining (ICDM '09)*, Miami, FL, December 2009

Large-scale system problem detection by mining console logs

Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael Jordan

In Proc. of the *22nd ACM Symposium on Operating Systems Principles (SOSP'09)*, Big Sky, MT, October 2009

Mining console logs for large-scale system problem detection

Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael Jordan

In Proc. of the *3rd workshop on Tackling Computer Systems Problems with Machine Learning Techniques (SysML'08)*, San Diego, CA, December 2008

Regulating workload in J2EE application servers

Wei Xu, Zhangxi Tan, Armando Fox and David Patterson

In Proc. of the *1st International Workshop on Feedback Control Implementation and Design in Computing Systems and Networks (FeBID'06)*, Vancouver, Canada, April 2006

Predictive control for dynamic resource allocation in enterprise data centers

Wei Xu, Xiaoyun Zhu, Sharad Singhal, and Zhikui Wang

In Proc. of the *10th IEEE/IFIP Network Operations & Management Symposium (NOMS'06)*, Vancouver, BC, Apr. 2006

Feedback control theory and processing system log streams

Wei Xu

Master thesis, EECS Department, UC Berkeley, December, 2005

Control considerations for scaling event processing

Wei Xu, Joseph L. Hellerstein, Bill Kramer and David Patterson

In Proc. of the *16th IFIP/IEEE Distributed Systems: Operations and Management (DSOM'05)*, Barcelona, Spain, October 2005

A flexible framework for statistical learning and data mining from system log streams

Wei Xu, Peter Bodik and David Patterson

In Proc. of *Workshop on Temporal Data Mining: Algorithms, Theory and Applications at The Fourth IEEE International Conference on Data Mining (ICDM'04)*, Brighton, UK, Nov, 2004

Peer-to-Peer support for massively multiplayer games

Bjorn Knutsson, Honghui Lu, Wei Xu and Bryan Hopkins

In Proc. of the *23rd Conference of the IEEE Communications Society (INFOCOM'04)*, Hong Kong, March 2004

WORKING EXPERIENCE

Google

Software Engineer

-Design and implementation of Google's next generation debug logging and cluster monitoring infrastructure.

Mountain View, CA

Aug, 2010 to now

Netease Beijing R&D Center / Yodao.com

Technical Consultant

- Lead the team for development and maintenance of the distributed storage system

Berkeley, CA

Feb - Jun, 2006

Jan - Dec, 2007

Netease Beijing R&D Center

Software Research Engineer

- Designer and lead engineer for a large scale, distributed storage system

- Designer and lead engineer for image search (<http://image.yodao.com>)

Beijing, China

Jun - Dec, 2006

HP Labs

Research Intern

- Research on predictive control in virtual machine resource allocations

Palo Alto, CA

Jun - Aug, 2005

Institute of Software, Chinese Academy of Science (ISCAS)
Research Assistant

Beijing, China
Jul-Aug, 2002

PROFESSIONAL ACTIVITIES

Panelist, Improving Logging Practice <i>Workshop on Managing Systems via Log Analysis and Machine Learning Techniques (SLAML'11)</i>	2011
Invited speaker <i>the 27th International Conference on Machine Learning (ICML'10)</i>	2010
Program Committee Member <i>Workshop on Managing Systems via Log Analysis and Machine Learning Techniques (SLAML)</i>	2010-2011
Program Committee Member <i>The 2nd Workshop on the Analysis of System Logs (WASL) 2009</i>	2009

OTHER EXPERIENCE

Internet-in-a-Box - Porting standard Linux applications to a soft-core CPUs on FPGA - Aiming at emulating large scale inter-network on a single FPGA board	2006-2007
Applying control theory to software systems - Controlling data flow in distributed data stream processing systems - Controlling request handling in three-tier J2EE application servers	2004-2006
Peer-to-Peer systems - Querying DHT under Extreme Network Conditions - Supporting Massively Multi-player Games with Peer-to-Peer Systems	2002-2004

Programming and IT

<i>Assistant Network Administrator</i> , Law School, Upenn	Feb.- June, 2002
Microsoft Certified System Engineer on Windows 2000 (Early Achiever)	March, 2001
Instructional equipment design and implementation Digital Logics Circuit Lab, Dept. of CS&T, Tsinghua Univ.	Summer, 2001
<i>Programmer</i> , ANJIA Corp. Ltd. (www.anjia.com), Shanghai, China	Summer, 2000
<i>Programmer and Network Support</i> , Dept. of Biology, Tsinghua Univ.	1999 -- 2001

Teaching

<i>Teaching Assistant</i> for CS186 (undergrad database) EECS department, UC Berkeley	Spring, 2003
<i>Lab Assistant</i> of Digital System Design Course Dept. of CS&T, Tsinghua Univ.	Summer, 2001

HONORS AND AWARDS

Department Fellowship EECS Department, UC Berkeley,	2004
---	------

Faculty Appreciation Award	2003
School of Engineering and Applied Science, Univ. of Penn	
Dean's List	2002-2003
School of Engineering and Applied Science, Univ. of Penn	
Tau Beta Pi Honor Society and Golden Key Honor Society	2002

SKILLS

Human Languages:	Proficient in English, native in Mandarin Chinese, beginning level of Japanese
Programming:	Proficient in Java, C++, C, Pascal, and various scripting/web developing languages; Proficient in Windows and Linux servers and virtual machines; Proficient in server development frameworks and tools; Extensive experience in cloud computing (both Amazon's EC2 and Google's AppEngine), and in large-scale cluster computing systems;
Hardware:	micro-controller, x86/SPARC Assembly languages

REFERENCES

David Patterson pattsrn@eecs.berkeley.edu UC Berkeley
Armando Fox fox@eecs.berkeley.edu UC Berkeley
Michael Jordan jordan@cs.berkeley.edu UC Berkeley
Ling Huang ling.huang@intel.com Intel Labs, Berkeley