

## TING CHEN

Software Engineer, LinkedIn Corporation  
chenboat@gmail.com  
www.cs.wisc.edu/~tchen

### Education

- 1998.8-2002.7 *National University of Singapore* *B.Comp in Computer Sciences*  
First Class Honors GPA: 3.82
- 2002.7-2004.12 *National University of Singapore* *M.Sc in Computer Sciences*  
Thesis Title: View Processing for a Native XML DBMS  
Winner of Singapore Computer Society (SCS) best Master Thesis Award 2005
- 2005.8-2011 *University of Wisconsin, Madison* *M.Sc and PhD Student in Computer Sciences*  
Thesis topic: indexing for regex evaluation; current GPA: 3.81; Minor in Statistics

### Experiences

- 2010.9-Now *LinkedIn Corp.* *Software Engineer*  
Build Java-based recruiting tools on LinkedIn's professional network which enable recruiters to find and reach ideal candidates in a scalable and economical manner. The tool has its own website visible only to paid recruiters. The website is built using Groovy-on-Grails framework in the front end, Oracle database as the backend storage and wired with other components of LinkedIn codebase using Spring dependency injection. My projects in our team include: (1) Provide web-based CRM features to LinkedIn's big corporate customers including contract data migration. We used Apache ActiveMQ to process users' web requests and perform retry and error recovery if necessary. I wrote PL/SQL statements to move data around in the backend. Groovy, a Java dialect and Javascript/Ajax are used to construct front end webpages. (2) Show Who View My Profile(WVMP) data to LinkedIn corporate customers. WVMP stores viewer-viewee relationship data in LinkedIn's Voldemort distributed key-value store. We query and anonymize the data before presenting it to the users. (3) Debug production issues on LinkedIn's website by examining production logs using tools Splunk/GLU, checking RPC call performance and suggesting database query tuning opportunities.
- 2008.5-2008.8 *IBM Almaden Research Center* *Intern*  
Continued to design and implement novel indexing schemes for regular expression evaluation. We build several versions of the index engine: on a Lucene index engine, a plain Linux file system and a IBM DB2 database and compared their relative performances. We design an algorithm to build index on the Linux FS with bounded memory guarantee, 80% less build time compared with state-of-the-art system and small index sizes. To use the index engine, we invented a new index query algorithm which provides more pruning of text documents. We also studied how to incrementally update the built index.
- 2007.5-2007.8 *IBM Almaden Research Center* *Intern*  
Designed and implemented novel indexing schemes for regular expression evaluation
- 2006.10-2007.12 *DBLife Research Group* *Research Assistant (Advisor: AnHai Doan)*  
Designed and Implemented the researcher citation subsystem of DBLife community information system prototype (dblfe.cs.wisc.edu) using perl; Helped to fix system bugs of the production DBLife server
- 2006.5-2006.8 *Google Inc.* *Intern )*  
Worked in Google's Ads Backend group to develop C++ middle-ware to improve the efficiency and usability of Google's database server clusters

## **Skills**

Java, C/C++, Perl, XSLT, PL/SQL, MySQL, PostgreSQL, DB2, CVS/SVN, Spring, Groovy/Grails

## **Publications**

1. Jiansheng Huang, Ting Chen, AnHai Doan, Jeffrey F. Naughton: On the provenance of non-answers to queries over extracted data. VLDB 2008
2. Eric Chu, A. Baid, Ting Chen, Anhai Doan, Jeff Naughton: A Relational Approach to Incrementally Extracting and Querying Structure in Unstructured Data. VLDB 2007
3. Ting Chen, Jiaheng Lu, Tok Wang Ling: On Boosting Holism in XML Twig Pattern Matching using Structural Indexing Techniques. SIGMOD 2005