

# Longbin Lai

☎ (+86) 18067950576 | ✉ longbin.lai@gmail.com

## Summary

---

An enthusiastic and dedicated researcher and engineer with exceptional work ethics. Working and leading experiences in both academia and industry. Over ten years' academic and engineering experiences with Hadoop (MapReduce), Spark and their variants. Proficient in Algorithms. Outstanding academic publications in big graph/data processing. Leadership in developing distributed graph processing system.

## Education

---

### The University of New South Wales, Australia (UNSW)

*Sydney, Australia*

PH.D. IN COMPUTER SCIENCE

*Jul. 2013 - Jul. 2017*

- All courses Highly Distinguished

### Shanghai Jiao Tong University (SJTU)

*Shanghai, China*

M.S. IN COMPUTER TECHNOLOGY

*Sep. 2010 - Mar. 2013*

- GPA 3.8 / 4.0, China's National scholarship, Top 2%

### Shanghai Jiao Tong University (SJTU)

*Shanghai, China*

B.S. IN INFORMATION SECURITY

*Sep. 2006 - Jun. 2010*

- GPA 3.6 / 4.0, twice B-class SJTU academic scholarship, Top 15%

## Experience

---

### Damo Academy, Alibaba Corporation

*Hangzhou, China*

STAFF ENGINEER

*Dec. 2019 - Present*

- Research on large-scale data/graph processing system, graph query language, and query optimization.
- Lead a team to build Gaia, an interactive graph query system in the distributed runtime.
- Develop GraphScope (<https://github.com/alibaba/GraphScope>), a large-scale one-stop graph processing system, that incorporates processing the workloads of graph analytics, interactive graph queries, graph pattern matching and graph learning.

### School of computer science and software engineering, East Normal China University

*Shanghai, China*

RESEARCH ASSISTANT

*May. 2019 - Dec. 2019*

- Lead a team to develop graph pattern matching system.
- Research on large-scale graph database, especially distributed query processing and storage.

### School of Computer Science and Engineering, UNSW

*Sydney, Australia*

RESEARCH ASSISTANT

*May. 2017 - May. 2019*

- Design and implement big graph processing primitives and languages.
- Lead a team to develop graph pattern matching system.

### Google Inc.

*Mountain View, CA, USA*

TECH INTERN

*Jan. 2017 - Apr. 2017*

- Designed and implemented an emulator that simulates the Google backbone network and the routing strategies for testing, debugging and routing validation.

## School of Computer Science and Engineering, UNSW

Sydney, Australia

PHD CANDIDATE, INDEPENDENT RESEARCH PROJECT

Jul. 2013 - May. 2017

- (**TwinTwigJoin**) Increased the performance of subgraph enumeration by up to an order of magnitude compared to the state-of-the-art by applying a decomposition-and-join framework in MapReduce.
- (**SEED**) Further improved the TwinTwigJoin by more than one order of magnitude by using a more advanced graph data storage mechanism (extending the traditional adjacency list) and an optimal join structure.

## Department of Advertising and Searching, Alibaba Cloud

HangZhou, China

### Computing Corporation

RESEARCH INTERN, TEAM PROJECT

Jan. 2012 - Sep. 2012

- Designed and implemented a web recommendation system based on Alibaba cloud computing system (MapReduce-like system), which serves over 1000 top websites in China.
- Improved the throughput of the recommendation system to over 2 billion records per hour via a well-designed MapReduce data flow.
- Implemented a prototype of web classification algorithm that is twice faster than existing algorithm by solely using the url of the web page.

## IBM Share-With-University Project

Shanghai, China

RESEARCH ASSISTANT, PROJECT LEADER

Oct. 2009 - Oct. 2010

- Saved the storage overhead of Hadoop File System (HDFS) by up to 30% without compromising the storage reliability by replacing the full replication mechanism with erasure coding.
- Improved the performance of Hadoop streaming utility (allowing coding with languages other than Java) by over 60% by replacing the synchronized inter-process communication module in Linux with desynchronized single-read-single-write queue.

## Skills

---

**Programming** Rust, Python, Java, C/C++, SQL, Scala

**Big Data** Hadoop, Spark, Timely dataflow system, Flink, AWS Infrastructure

**Big Graph** Giraph/Pregel, GraphX, Gelly, Neo4J, GraphLab

**Machine Learning** Tensorflow, Scikit-learn, Libsvm, Spark ML

**Others** Recommendation Systems, Yii2 on Php, Django on Python

## Selected Publications

---

### HUGE: An Efficient and Scalable Subgraph Enumeration System

*Sigmod 2021, ERA A\**

ZHENGYI YANG, LONGBIN LAI, XUEMIN LIN, KONGZHANG HAO, WENJIE ZHANG

*ACM SIGMOD/PODS International  
Conference on Management of Data*

### GraphScope: a unified engine for big graph processing

*VLDB 2021, ERA A\**

WENFEI FAN, TAO HE, LONGBIN LAI, XUE LI, YONG LI, ZHAO LI, ZHENGPING QIAN, CHAO TIAN,  
LEI WANG, ETC.

*PVLDB Volume 14, Issue 12*

### GAIA: A System for Interactive Analysis on Distributed Graphs Using a High-Level Language

*NSDI 2021, ERA A\**

ZHENGPING QIAN, CHENQIANG MIN, LONGBIN LAI, YONG FANG, GAOFENG LI, YOUYANG YAO,  
BINGQING LYU, ETC.

*18th USENIX Symposium on Networked  
Systems Design and Implementation*

### **A framework to quantify approximate simulation on graph data**

XIAOSHUANG CHEN, LONGBIN LAI, LU QIN, XUEMIN LIN, BOGE LIU

*ICDE 2021, ERA A\**  
2021 IEEE 37th International Conference  
on Data Engineering

### **Efficient structural node similarity computation on billion-scale graphs**

XIAOSHUANG CHEN, LONGBIN LAI, LU QIN, XUEMIN LIN

*The VLDB Journal, ERA A\**  
Volume 30, Issue 3

### **Distributed subgraph matching on timely dataflow**

LONGBIN LAI, ZHU QING, ZHENGYI YANG, XIN JIN, ZHENGMIN LAI, RAN WANG, ETC.

*VLDB 2019, ERA A\**  
PVLDB Volume 12, Issue 10

### **Scalable Distributed Subgraph Enumeration**

LONGBIN LAI, LU QIN, XUEMIN LIN, YING ZHANG, LIJUN CHANG

*VLDB 2017, ERA A\**  
PVLDB Volume 10, Issue 3

### **Scalable Subgraph Enumeration in MapReduce, a cost-oriented approach**

LONGBIN LAI, LU QIN, XUEMIN LIN, LIJUN CHANG

*VLDB Journal, ERA A\**  
The VLDB Journal, Volume 26, Issue 3

### **Scalable Subgraph Enumeration in MapReduce**

LONGBIN LAI, LU QIN, XUEMIN LIN, LIJUN CHANG

*VLDB 2015, ERA A\**  
PVLDB Volume 8, Issue 10

### **ShmStreaming: A shared memory approach for improving Hadoop streaming performance**

LONGBIN LAI, JINGYU ZHOU, LONG ZHENG, HUAKANG LI, YANCHAO LU

*AINA 2013, ERA B*

## **Honors & Awards**

---

2021	<b>N/A</b> , HangZhou 521 Project of Talent Introduction	<i>Hangzhou, China</i>
2012	<b>Top 1%</b> , China's National Scholarship	<i>SJTU, China</i>
2011	<b>Top 4%</b> , Tencent Academic Scholarship	<i>SJTU, China</i>
2010	<b>Top 10%</b> , Outstanding Graduate of Shanghai Jiao Tong University	<i>SJTU, China</i>
2009	<b>Top 6%</b> , Sony Academic Scholarship	<i>SJTU, China</i>
07, 08	<b>Top 15%</b> , B-Class SJTU Academic Scholarship	<i>SJTU, China</i>