

Hyunjung Park

hyunjung@cs.stanford.edu

Education

- 2008–2014 Ph.D. in Electrical Engineering, **Stanford University** (research area: database systems)
- 2006–2009 M.S. in Electrical Engineering, **Stanford University** (concentration: computer hardware)
- 1999–2002 B.S. in Electrical Engineering, **Seoul National University**

Employment

- 2013/11–present Researcher, **Microsoft Research** Redmond, WA
Conducting research on topics related to Microsoft Azure SQL Database, a multi-tenant relational database-as-a-service.
- 2008/9–2013/10 Research Assistant and Teaching Fellow, **Stanford University** Stanford, CA
Designed and implemented systems for crowdsourcing structured data.
Designed and implemented a Hadoop extension for fine-grained data provenance support.
Taught and co-taught CS346 Database System Implementation in Springs 2010 and 2011.
- Summer 2011 Research Intern, **Microsoft Research** Redmond, WA
Developed an I/O performance isolation mechanism for co-located tenants in Microsoft Azure SQL Database.
- Summer 2010 Research Intern, **Microsoft Research** Redmond, WA
Developed an automatic, workload-driven method for defragmenting B+-tree indexes in Microsoft SQL Server.
- Summer 2009 Research Intern, **IBM Research – Almaden** San Jose, CA
- Summer 2008 Software Engineering Intern, **Google, Inc.** Seoul, Korea
- 2004/11–2006/2 Senior Software Engineer, **Naver Corporation** (formerly known as NHN) Seongnam, Korea
- 2002/9–2004/10 Software Engineer, **Piolink, Inc.** Seoul, Korea
Developed a high-performance layer-7 switch that transparently separates HTTP requests from TCP connections and routes those requests, in the linux kernel.

Awards

- 2010 Winner, ACM SIGMOD Programming Contest (building a distributed query engine)
- 2006–2011 Samsung Graduate Fellowship
- 1998 Silver medal, 39th International Mathematical Olympiad (IMO)

Selected Publications

- 2014 **H. Park** and J. Widom. “CrowdFill: Collecting Structured Data from the Crowd”, *ACM SIGMOD International Conference on Management of Data*, Snowbird, Utah.
- 2013 **H. Park** and J. Widom. “Query Optimization over Crowdsourced Data”, *39th International Conference on Very Large Data Bases (VLDB)*, Trento, Italy.
- 2011 **H. Park**, R. Ikeda, and J. Widom. “RAMP: A System for Capturing and Tracing Provenance in MapReduce Workflows”, *37th International Conference on Very Large Data Bases (VLDB)*, Seattle, Washington. (Also presented in *Hadoop Summit 2011*, Santa Clara, California.)