

Jun Woo Park

PERSONAL DATA

WORK ADDRESS: Google Sunnyvale
PERSONAL EMAIL: johnpa@gmail.com

EDUCATION - CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PA, USA

MAY 2010 Bachelor of Science (Hons.) in ELECTRICAL AND COMPUTER ENGINEERING
DECEMBER 2016 Master of Science in COMPUTER SCIENCE
MAY 2019 Doctor of Philosophy in COMPUTER SCIENCE
Advisor and committee chair: Gregory R. GANGER
Committee members: Phillip B. GIBBONS, George AMVROSIADIS, Michael A. KOZUCH (Intel labs)
Topics: Cluster scheduling, cloud computing, cluster workload analysis

WORK EXPERIENCE

JUL 2019 - CURRENT	<p>Software Engineer at GOOGLE LLC., San Francisco Bay Area <i>Platforms Cloud Efficiency team</i></p> <p>Efficiency lead for Cloud container runtime environment for internal services. Driving benefit cost analysis and design for planning and development entry.</p> <p>VM family recommendation for current GCP customers, to scheduler, and for prospective GCP customers</p> <p>Cloud workload analysis, bare-metal performance testing with regards to AMD SoftSKU, etc.</p>
JUN-AUG 2016	<p>Intern at TWO SIGMA INVESTMENT LP, New York <i>Job scheduler and Job scheduling log analysis</i></p> <p>Job scheduling data analysis using Zeppelin/Spark, thinking about jobs running on VMs to containers, and collecting job scheduling data for subsequent publications for Eurosys'18, ATC'18, SOCC'18.</p>
JUN-AUG 2014	<p>Intern at GOOGLE Inc., San Francisco Bay Area <i>Machine learning based automation of infrastructure monitoring rule tuning</i></p> <p>Analysis of infrastructure monitoring rule changes and bug reports of cluster application pipeline. Devised an anomaly detector based rule tuning algorithm for SREs.</p>
MAR-JUN 2013	<p>Software Engineer at KULCLOUD NETWORKS, South Korea <i>Software Defined Networking</i></p> <p>Openflow controller and SDN-based network management system for OpenStack.</p>
MAY 2010-MAR 2013 MAR 2009-DEC 2009	<p>Financial Software Engineer at KOREA ASSET PRICING, South Korea <i>Risk management system for credit derivatives</i></p> <p>Credit derivative risk management software system providing services including daily Marked-to-Market valuation, stress testing, and other analytics data to financial firms including major investment banks and Insurance companies in Korea. Researched and drafted internal technical reports on credit linked derivative valuation methodology with a special focus on hybrid derivative products (payoff linked to combination of credit risk, stock prices, interest rates, and foreign exchange rates)</p>

RESEARCH PROJECTS AND PUBLICATION

SOCC'18 Stratus: cost-aware container scheduling in the public cloud

Authors: Chung A., **Park J.W.**, Ganger G.R.

A new cluster scheduler specialized for orchestrating job execution on virtual clusters, which focuses primarily on dollar cost considerations by aggressively packing tasks onto machines, trying to make allocated resources be either mostly full (highly utilized) or empty (so they can be released to save money).

[Best Student Paper Award]

USENIX ATC'18 On the diversity of cluster workloads and its impact on research results

Authors: Amvrosiadis G., **Park J.W.**, Ganger G.R., Gibson G.A., Baseman E., DeBardeleben N.

An analysis of the differences and similarities between new traces consisting of job scheduler logs from one private and two HPC clusters, drawing comparisons to well-known Google cluster trace.

EUROSYS'18 3Sigma: distribution-based cluster scheduling with runtime uncertainty

Authors: **Park J.W.**, Tumanov, A., Jiang A., Kozuch, M.A., Ganger, G. R.

A scheduler that uses full distributions of relevant runtime history instead of single-point estimates (e.g., mean or median of relevant subset of historical runtimes) to cope with inherent variability of the job runtimes.

EUROSYS'16 TetriSched: global rescheduling with adaptive plan-ahead in dynamic heterogeneous clusters

Authors: Tumanov, A., Zhu, T., **Park, J.W.**, Kozuch, M. A., Harchol-Balter M., Ganger, G. R.

A scheduler that leverages information supplied by the reservation system about jobs' deadlines and estimated runtimes to plan ahead in deciding whether to wait for a busy preferred resource type (e.g., machine with a GPU) or fall back to less preferred placement options.

[Best Student Paper Award]

OSDI'14 Scaling distributed machine learning with the parameter server

Authors: Li, M., Andersen, D.G., **Park, J.W.**, Smola, A.J., ...

A parameter server framework for distributed machine learning problems. The framework manages asynchronous data communication between nodes, and supports flexible consistency models, elastic scalability, and continuous fault tolerance.

SKILLS

Languages: English, Korean, C/C++, Java, Python, Go, Scala, SQL, Visual Basic
Frameworks: Hadoop YARN, Spark, Kubernetes, Node.js (express), Flask
Other skills: Optimization (ILP), Credit linked derivatives

SCHOLARSHIPS AND CERTIFICATES

Samsung Scholarship (\$50,000 per year for academic years 2013-2018)

Passed level 2 of Chartered Financial Analyst exam

REFERENCE

Available upon request