

PETER BAKKUM

+1 757-812-2365 • pbbakkum@gmail.com • linkedin.com/in/pbbakkum

Stripe

San Francisco, California

Engineering Manager, Staff Engineer

May 2018 - June 2022

Stripe is a financial technology firm specializing in online B2B payments. My primary focus was managing the Ledger engineering team; we migrated Stripe's business accounting from a manual process to an automatic, audited pipeline. We integrated 25+ money-moving teams (100s of millions of daily events) into a custom service used heavily by Ops and Accounting, which now powers Stripe's financials. My 3 years as a manager was bookended by stints as an individual contributor on the Treasury Engineering (automating internal treasury functions) and Payment Flows (Stripe API infrastructure) teams.

Quizlet

San Francisco, California

Platform Lead

May 2014 - September 2017

Quizlet is an ed-tech app popular with high-school and college students and is a Top-50 U.S. website by traffic. I joined Quizlet as the 11th employee and helped scale the company to 65 employees and 25M monthly active users by leading the Platform Team, responsible for operating the website. I split my time between technical work (including on-call), recruiting, managing 7 team-members, technical decision-making, and managing vendors. Quizlet grew ~50% a year during my tenure, prompting a focus on infrastructure scaling and performance. Major projects included migrating clouds to Google Cloud Platform, building out a data pipeline and analytics warehouse, and optimizing and scaling relational databases.

Groupon

Palo Alto, California

Member of Technical Staff

February 2012 - May 2014

After entering Groupon through the acquisition of Hyperpublic, I continued our work sourcing and cleaning a sophisticated database of local businesses as part of the Merchant Data team. I focused on building the data platform for the team, including deployments of Hadoop, HBase, Storm, Elasticsearch, and MongoDB. I spent time helping other teams integrate with our Hadoop package, and ultimately made it the standard method for deploying Hadoop within Groupon. Following that project, I ran Groupon's internal Taxonomy Service, involving both product management and engineering. Finally, I spent time prototyping technology for the future of Groupon's platform, working with Linux containers and Mesos.

Hyperpublic

New York, New York

Software Engineer

November 2011 - February 2012

I joined Hyperpublic, an early-stage startup funded by Lerer Ventures, as the 5th employee. I focused on sourcing and integrating social data to Hyperpublic's data set. Hyperpublic's core product was a sourced, aggregated, and cleaned database of local businesses accessible to developers through an API. Ultimately Hyperpublic was sold to Groupon which leveraged the team and its technology to improve its internal partner intelligence capability.

NEC Laboratories America

Princeton, New Jersey

Consultant

August 2010 - August 2012

Research Assistant

Summer 2010

I was brought into a heterogeneous computing research group at NEC led by Srimat Chakradhar to continue work I published as my undergraduate thesis. My work involved leveraging Nvidia GPUs to accelerate database query processing. The outcome of my work there was an experimental SQL database capable of executing queries on a GPU, which produces a speedup of 2-4x for many workloads when compared to multi-core CPU execution.

Investment Technology Group

New York, New York

Quantitative Analyst

August 2010 - November 2011

Intern

Summer 2008, Summer 2009

At ITG I worked in the Algorithmic Trading group to analyze equity trading performance for clients of the firm, including complex situations such as dark pool trades, hidden market orders, and algorithm gaming. This involved understanding the landscape of electronic equity trading in the U.S. and using quantitative analysis of large data sets and statistical modeling to suggest optimizations for our clients' trading strategies. I also spend time developing trading strategies and working with software developers to implement and experiment with them. Prior to my full-time role I interned in the Market Data group, then in the Algorithmic Trading group.

Darden School of Business

Charlottesville, Virginia

Research Assistant

Spring 2010

At Darden I helped manage data for a professor's research project on Community Development Financial Institutions. The project was funded by the John D. & Catherine T. MacArthur Foundation.

The Forgetful Gentleman

Charlottesville, Virginia

Web Programmer

Spring 2009

During my third year at U.Va. I developed a website for a startup company incubated at the Darden School of Business, which is now a vendor of stationary both online and at retail locations like Neiman Marcus.

Miller Center of Public Affairs

Charlottesville, Virginia

Web Programmer

October 2006 - May 2008

During my first and second years at U.Va. I collaborated with other students and Miller Center employees to plan and implement a complete website redesign involving hundreds of dynamic pages and a huge media library.

Northrop Grumman

Newport News, Virginia

Intern

Summer 2007

Unpaid Intern

October 2004 - June 2005

I worked in the Defense Mission Systems branch of Northrop Grumman as an intern in high school, and later during the summer while at U.Va. While there I developed user interface components for the Global Command & Control System, a software product contracted to Northrop by the Department of Defense. The DoD uses this software to track and visualize units across a global battlefield map.

University of Virginia

School of Engineering & Applied Science, 2010

Bachelor of Science, Computer Science, Economics.

Graduated with distinction in Computer Science.

Graduated with a Finance concentration in Economics.

Recipient, The Timothy M. Grembowski Scholarship.

Finalist, 23rd Annual Undergraduate Research & Design Symposium.

Champion, 8th McIntire School of Commerce Hedge Fund Tournament.

Technical Proficiency

Most proficient: Go, SQL, Linux, Google Cloud Platform

Experience with: Java, Python, Ruby, HTML/CSS, Javascript, C/C++, CUDA, Hadoop, Spark, Kafka, Flink, MongoDB, MySQL, PostgreSQL, HBase, Pinot, BigQuery, Spanner

Publications

U.S. Patent Application US20220147977A1: Database with data integrity maintenance
Carl Jackson, Peter Bakkum, Usman Masood. November 11, 2020.

MongoDB in Action, Second Edition
Kyle Banker, Peter Bakkum, Shaun Verch, Douglas Garrett, and Tim Hawkins. March 2016.
Manning Publications Inc. ISBN 9781617291609.

Accelerating SQL Database Operations on a GPU with CUDA
Peter Bakkum and Kevin Skadron. In *GPGPU '10: Proceedings of the Third Workshop on General-Purpose Computation on Graphics Processing Units*, pages 94-103, New York, NY, USA, March 2010. ACM.

Talks

Real-time Analytics at Stripe
Pushing the Boundaries of Real-time Analytics (StarTree Inc. Launch Event), July 2021.

Running Open-Source Databases on Google Cloud Platform
Google Cloud Next 2017, San Francisco, CA, March 2017.

Cloud Spanner 201: Getting the most out of Cloud Spanner
Google Cloud Next 2017, San Francisco, CA, March 2017.

Scaling Quizlet Fall 2016
SF DevOps for Startups Meetup, San Francisco, CA, January 2017.

Graduating with Honors: How Quizlet is Scaling Elasticsearch to Help Children Learn
Elasticon 2015, San Francisco, CA, March 2015.

Using MongoDB for Groupon's Place Data
MongoDB SF, San Francisco, CA, May 2013.

A GPU Database Architecture
GPU Technology Conference, Santa Clara, CA, March 2013.

Accelerating SQL Database Operations on a GPU with CUDA
Third Workshop on General-Purpose Computation on Graphics Processing Units, Pittsburgh, PA, March 2010.

Blogging

Quizlet Tests Cloud Spanner — The Most Sophisticated Cloud Database. February 2017.

What's the Best Cloud? Probably GCP. March 2016.

MongoDB Data Locality. June 2013.