AMOS NG

Contact

Address:
3621A Interlake Ave N
Seattle, WA 98103
USA

Phone:
+1 (650) 279 2894

Email:
jobs@amos.ng

Social

GitHub – amosjyng
LinkedIn – Amos Ng

Languages

English – Native
Mandarin – Conversational
Cantonese – Conversational

Summary

I'm a senior full-stack developer with an emphasis on the backend. I've worked across a number of consumer-facing products, completing projects end-to-end from initial requirements gathering to final implementation, QA, monitoring, and reporting.

Technologies used

Programming Languages

C, C++, Java, Javascript (JS), Python, Go (golang), Scala, Ruby, Rust, Clojure, Common Lisp, Haskell, HTML, CSS

Web Frameworks

React, Ruby on Rails, Django, Flask, Play Framework

Other Web Technologies

jQuery, REST, gRPC, Selenium

Databases

PostgreSQL, MySQL, SQLite, MongoDB, Elasticsearch

Automated Testing

Selenium, Junit, pytest, Jest

Cloud Services

AWS, Heroku, Digital Ocean, Google Cloud

Platforms Developed For

Mobile (iOS and Android), server (Linux), web (all major browsers)

Containerization/Provisioning

Docker, Vagrant, NixOps

Education

Bachelor of Science in **Physics** at the **University of California, Los Angeles**	2010 - 2013
Master of Science in **Language Technologies** at **Carnegie Mellon University**	2013 - 2015

Project: 3rd PARTY BACKEND INTEGRATIONS

Context

Company
Reserve

Products
Reserve iOS/Android apps
Internal compliance webapp

Industry
Software Development
(Full-Stack)

Time Frame
September 2021 – November 2021

Objective

Various tasks that were not part of the core business competency (namely, OFAC compliance and marketing campaign tracking) needed to be outsourced to external companies, Sumsub and AppsFlyer, respectively.

As such, the objective was to develop behind-the-scenes integration with third party systems without disrupting regular user activity or internal operations team workflows.

Tasks

- Design and implementation of all required changes throughout the stack:
  - o Database migrations to store new data required to interact with third party services
  - o Addition of server-side API endpoints to retrieve and record the new integration data on the database
  - o Rerouted existing mobile app user flow to use the new server-side APIs
  - o Built entirely new internal dashboards and user management pages for the internal operations team to continue solving customer support issues
- Design and implementation of automated unit and integration tests at every part of the stack
- Handled communications with external developers on integration issues
- Continuous reporting and monitoring of feature usage signals

Tech Stack

- PostgreSQL database
- Python backend server
- NodeJS React webapp
- NodeJS React Native mobile app
- Sumo Logic log monitoring

Project: NOTIFICATIONS PIPELINE REWRITE

Context

Company
Google

Products
Internal Duo Notifications System

Industry
Software Development
(Backend)

Time Frame
January 2020 – April 2020

Objective

Users of Duo, a consumer video calling app, are sent notifications as part of various internal marketing campaigns. The scheduling, sending, and tracking of these notifications and associated campaigns was handled by an outdated legacy system.

As such, the objective was to do a full rewrite of the system, without any downtime and without making any changes to downstream consumers of the data produced.

Tasks

- Design and implementation of fully rewritten new system that ran 500% faster than the old one
- Replicating undocumented code without access to original maintainers
- Design and implementation of automated unit and end-to-end pipeline tests
- Completion under a tight deadline, before support for deprecated dependencies was pulled
- Daily reporting and monitoring of new notifications processing cronjob
  - Ensured new system produced output that was byte-for-byte consistent with old system, for a few days before finalizing switchover

Tech Stack

- Java "Flume" (internal equivalent of Spark) big data processing
- Bigtable NoSQL data storage

Project: RING LATENCY MINIMIZATION

Context

Company
Google

Product
Duo iOS/Android apps

Industry
Software Development
(Full-Stack)

Time Frame
July 2017 – July 2018

Objective

Duo is a consumer video calling app. It was observed that the more users miss calls, the more they lose engagement with the app. This problem proved especially severe in countries with poor internet and older, slower phone models.

As such, the objective was to try to find a variety of methods to maximize call connection rates.

Tasks

- Drove the entire project from beginning to end, including:
  - Initial identification of ring latency in particular (the delay between the pressing of the call button and the ringing of the recipient's phone) as a major area of improvement
  - Proposing, designing, and implementing improvements to said ring latency metrics
  - Post-implementation QA and monitoring
- Design and implementation of all engineering changes:
  - Cross-stack metrics collection from both client and server
  - Asynchronous big data processing of newly collected metrics for subsequent analysis
  - Added new gRPC endpoints for server-side actions, and modified client behavior to use new endpoints
- Design and implementation of automated unit and integration tests at every part of the stack
- Extensive A/B testing to verify improvements and monitor for tradeoffs. Biggest wins include:
  - -10% in ring latency
  - +3.1% in call connection rates
  - +3 additional years of human interaction taking place on Duo every single week
  - No discernible drop in battery life or other negative metrics

Tech Stack

- Go backend
- Android mobile app
- gRPC endpoints
- "Sawmill" logging
- "Flume" (internal equivalent of Spark) big data processing

Project: PRIVACY CONSULTING

Context

Company
Google

Product
Various internal products

Industry
Privacy Consulting

Time Frame
April 2020 – June 2021

Objective

Review internal software projects for privacy concerns, ensuring compliance with US and EU laws, as well as Google's FTC consent decree.

Tasks

- Unblocked over 100 privacy or anonymization reviews to ensure the timely launch of internal services while maintaining employee privacy
- Assisted in the documentation of numerous systems for future audit compliance purposes
- Emphasized controls on the collection, access, retention, and automated deletion of employee data
- Where possible, explored more privacy-preserving approaches to accomplishing stated goals
- Educated employees and leaders from various internal organizations on privacy risks, including the differences between pseudonymization and anonymization
- Helped guide internal software tools towards implementing change management compliance
- Implementation of Gaussian and post-aggregation partition selection algorithms, as well as "distinct per key" counts, for the internal differential privacy large-scale data processing library

Tech Stack

- Differentially private "Flume" (internal version of Spark)

Project: NEUROSCIENCE INDEX

Context

Company
Allen Institute for Artificial Intelligence

Product
Semantic Scholar

Industry
Software Development (Full-Stack)

Time Frame
January 2016 – July 2016

Objective

Allow Semantic Scholar, a search engine for scientific papers, to support searching for neuroscience as well as computer science papers.

Tasks

- Optimization and simplification of the large-scale data-processing Spark pipeline to finish in 1 hour instead of 9, allowing daily ingestion of new neuroscience data sources to happen in a timely manner
- Improvements to GROBID for citation extraction from PDFs, including debugging upstream issues that stemmed from a bug in the JVM itself
- Design and implementation of new server API endpoints that allowed for new neuroscience-specific filters on searches
- Design and implementation of frontend UI filtering changes that hit the newly implemented backend API endpoints
- Design and implementation of automated unit and integration tests at every part of the stack

Tech Stack

- Scala backend
- React webapp
- Elasticsearch index
- Spark big data processing
- GROBID PDF processing