

Julian Graham

joolean@undecidable.net

<http://undecidable.net/joolean/>

Skills

- Java 8, RxJava, Guava, Spring, JPA
- Big data analytics / distributed computing: Hadoop Map/Reduce, HBase, Hive, Kafka, Cassandra
- Amazon Web Services: ElastiCache, EMR, Lambda, S3, SQS
- UNIX systems programming: C
- Functional programming: Scheme
- Build and deployment: GNU Autotools, Maven, Jenkins, GitLab, Docker

Professional experience

Principal Software Engineer; 2016-present Conductor, Inc.

(Key technologies: RxJava; Amazon Lambda, S3, SQS; Cassandra; Docker; Kubernetes)

I'm currently taking a step back from leading engineering teams at Conductor in order to focus on cross-cutting, difficult-to-staff product and architectural issues. In the past year, I've built some targeted features to help our sales team close several strategic accounts; and I've ported our customer-facing application off of a legacy database and onto a reporting framework based on Reactive Streams, saving us many hundreds of thousands of dollars in COGS and scaling our reporting capabilities up by an order of magnitude.

Right now, I'm developing architecture and application prototypes to help guide the team through a large-scale migration to microservices.

Engineering Manager; 2012-2016 Conductor, Inc.

(Key technologies: Kafka, Map/Reduce, HBase, Hive, Amazon S3)

During my tenure as an Engineering Manager at Conductor, I established three different back-end / data engineering teams, supervising up to 11 direct reports. I defined a software development lifecycle process for a part of the engineering organization previously operating without one, and acted as a technical product manager to define architecture and requirements for back-end software services and data products. I doubled the size of the team in a competitive hiring market, ran a successful summer internship program, and developed training materials to ramp new hires on distributed computing technologies.

As a technical manager, I guided my teams through the design and implementation of exciting projects, including:

- Transitioning our data collection pipeline to an actor model based on Apache Kafka, allowing us to scale the number of CPUs in our cluster in response to changes in data volume
- Creating a distributed rate-limiting system with locking based on Redis to control the behavior of our crawlers and maximize our utilization of scarce proxy resources
- Evolving our traffic analytics ETL services to increase their throughput by an order of magnitude - via Map/Reduce and Hive - and to make them highly available and auditable, while extracting an SPI that's since been used to integrate three additional analytics vendors
- Scaling the performance of our report publishing infrastructure by a factor of 10 by transitioning to a new report generation framework based on Amazon Elastic Map/Reduce and S3, and safely migrating many terabytes of historical customer data to our new cloud data warehouse - without interruption to service or degradation in data quality.

**Lead Software Engineer; 2009-2012
Conductor, Inc.**

(Key technologies: Java, Spring, JPA, Map/Reduce, HBase)

Conductor makes Searchlight, a web presence management application for marketers that helps them promote content via unpaid channels on the web. In my initial engineering role at Conductor, I worked across the stack, building out the data access layer for new features in Searchlight, and developing data visualizations for the application's front-end.

After moving to the platform engineering team, I was a key contributor on a project to transition the product's information extraction, recommendation generation, and ETL systems to Hadoop Map/Reduce and HBase, increasing the capacity of our report generation pipeline from the order of gigabytes to terabytes. I led the design and development efforts to integrate terabytes of traffic analytics data into the product via an automated collection and ETL pipeline and a sharded data warehouse.

**Senior Development Engineer, Technical Producer; 2007-present
Rebel Monkey, Inc.**

(Key technologies: Java, Spring, Jersey, Project Darkstar, CAS, ActionScript)

I joined Rebel Monkey to lead the server development for a cooperative, massively multiplayer online game - a dream project for me. We failed to build an audience for the game, but its technology platform was arguably a success. The game world ran inside a specialized middleware container created by Sun Microsystems (right before the takeover by Oracle) that executed microtransactional fragments of game code with strong guarantees for durability and consistency.

As a smaller project, I helped build a paper-doll animation system in ActionScript to allow players to customize the appearances of their avatars with different outfits and hairstyles.

Senior Software Engineer; 2003-2007
DataSynapse, Inc.

(Key technologies: Java, C#, C++)

In the days before Apache Hadoop, DataSynapse (now TIBCO) built a distributed computing platform to enable low-latency processing of compute-intensive tasks across clusters of commodity hardware, for clients in financial services and scientific research.

I contributed to the Java and C++ client libraries and task execution engines, and led the engineering effort to build a .NET version of the same.

Free Software projects

- gzochi [www.nongnu.org/gzochi] - Scalable middleware for MMOs
- r6rs-protobuf [gitlab.com/joolean/r6rs-protobuf] - Protocol Buffers implementation in Scheme
- r6rs-thrift [gitlab.com/joolean/r6rs-thrift] - Apache Thrift implementation in Scheme
- GNU Guile [www.gnu.org/software/guile] - Scheme language platform; contributed majority of R6RS standard library and pthreads threading implementation, misc. patches
- libRUIIN [www.nongnu.org/libruin] - Text console renderer for user interfaces
- SCSS [www.nongnu.org/scss] - CSS implementation in Scheme
- SDOM [www.nongnu.org/sdom] - DOM implementation in Scheme

Education

Yale University

Received Bachelor of Science in Computer Science, 2003.

Relevant coursework includes: Systems Programming & Computer Organization, Cryptography, Design & Analysis of Algorithms, Operating Systems, Compilers & Interpreters, Artificial Intelligence, Computer Networks