Roman Leventov

Github StackOverflow Twitter key-value-stories.blogspot.com koloboke.com LinkedIn

Areas of expertise

Java SE, JVM (JIT compiler), concurrency, JMM, data structures (esp. hash tables), off-heap, performance optimization, benchmarking (JMH), API design

Interests: design and implementation of data storing and data processing systems, databases, data structures, language design and runtimes (their tradeoffs), Linux, performance optimization, mechanical sympathy.

Experience

Software Engineer at Metamarkets, 2016 - present

- Responsible for management of a Druid cluster of 350+ instances in AWS
- Finding and resolving bottlenecks, reliability and GC issues with Druid

Software Engineer at Higher Frequency Trading, 2014 - 2016

- Engineering of the family of OpenHFT/Chronicle libraries and frameworks

Projects

I'm the author of

- Chronicle Map 3: the fastest, highly concurrent off-heap key-value store for Java

- Koloboke Collections: the fastest primitive hash maps for Java.
Koloboke Compile is a compile-time code generator which combines features of collections in the most efficient way. Koloboke Compile squeezes the last 5% of collections performance.

- Zero Allocation Hashing: the fastest hashing library for Java

- SmoothieMap: lower-footprint, latency spike-free java.util.Map implementation

- Yarr: data flow framework in Haskell, a faster rewrite of the Repa framework with better API

I'm a committer to Druid: performance-orientied, highly scalable time-series DB in Java. Improving query processing performance, reducing allocations and making changes aiming to improve performance in regard to Java's GC, reducing locking, fixing concurrency bugs.

Public Speaking

Highload++, November 2015, Chronicle Map - key-value store for trading in Java