



C o m m u n i t y E x p e r i e n c e D i s t i l l e d

Clojure High Performance Programming

Understand performance aspects and write high performance code with Clojure

Shantanu Kumar

[PACKT] open source*
PUBLISHING community experience distilled

Clojure High Performance Programming

Understand performance aspects and write high performance code with Clojure

Shantanu Kumar

[PACKT] open source 
PUBLISHING
community experience distilled
BIRMINGHAM - MUMBAI

Clojure High Performance Programming

Copyright © 2013 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First published: November 2013

Production Reference: 1131113

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham B3 2PB, UK.

ISBN 978-1-78216-560-6

www.packtpub.com

Cover Image by Duraid Fatouhi (duraidfatouhi@yahoo.com)

Credits

Author

Shantanu Kumar

Project Coordinator

Amey Sawant

Reviewers

Jan Borgelin

Mimmo Cosenza aka Magomimmo

Paul Stadig

Miki Tebeka

Proofreader

Paul Hindle

Indexers

Hemangini Bari

Mehreen Deshmukh

Acquisition Editors

Sam Birch

Andrew Duckworth

Graphics

Ronak Dhruv

Yuvraj Mannari

Commissioning Editors

Priyanka Shah

Meeta Rajani

Llewellyn Rozario

Production Coordinator

Kyle Albuquerque

Technical Editors

Jalasha D'costa

Monica John

Cover Work

Kyle Albuquerque

Copy Editors

Alisha Aranha

Roshni Banerjee

Tanvi Gaitonde

Alfida Paiva

Lavina Pereira

About the Author

Shantanu Kumar is a software developer living in Bangalore, India, with his wife. He started learning programming in 1991, using BASIC on MS DOS when he was at school. There, he developed a keen interest in the x86 hardware and assembly language, and he dabbled in it for a good while. Later, he programmed professionally in various business domains and technologies while working with the Indian Air Force and several IT companies.

In recent years, Shantanu has worked on high performance and distributed systems. Having used Java for a long time, he discovered Clojure in early 2009 and has been a fan ever since. Clojure's pragmatism and fine-grained orthogonality continues to amaze him, and he believes he is a better developer because of this.

When not busy with programming or reading up on technical subjects, he enjoys reading non-fiction, riding his bike, and occasionally just lazing in his free time. Shantanu is an active participant in the Bangalore Clojure users group and develops several open source Clojure projects on GitHub.

Acknowledgments

I would like to thank Rich Hickey for creating Clojure and making it available as open source, and for his awesome talk videos. I would also like to thank Alex Miller for arranging so many Clojure talks and for making their videos accessible to all; and Alex Ott, Michael Klishin, and others from the Clojure community for their hard work in making Clojure documentation aggregated and available.

While I was working at the Bangalore office of Runa (now Staples Lab) earlier, several colleagues shared valuable input about Clojure performance. Most notably, Zach Tellman shared his insight about Clojure and JVM performance, Isaac Praveen and Abhijith Gopal shared a great deal of information about Clojure application behavior under load, and Philippe Hanrigou shared his ideas about high performance API design and queue systems. I want to thank all of them.

This book would not have become a reality without the fine people at Packt Publishing. I would like to thank Ashvini Sharma for contacting me and convincing me to take up writing this book, Anish Ramchandani and Amey Sawant for coordinating the writing process, and the Commissioning Editors Meeta Rajani, Llewellyn Rozario, and Priyanka Shah for shaping up this book as I engaged in my debut writing. Technical Editors Jalasha D'costa and Monica John helped me disambiguate and refine the language in this book. I also owe my gratitude to the technical reviewers Jan Borgelin, Mimmo Cosenza, Paul Stadig, and Miki Tebeka – their feedback made the content so much better. Any errors or omissions, however, are only due to me.

Writing this book has been an arduous task. I want to thank my wife Binita for putting up with me while I was immersed far too many days, nights, and weekends into the book. If not for her support, I would not have been able to do justice to this book.

About the Reviewers

Jan Borgelin is the co-founder and CTO of BA Group Ltd., a Finnish IT consultancy providing services for global enterprise clients. With over 10 years of professional software development experience, Jan has had the chance to work with different technologies and programming languages in international projects where performance requirements have always been critical to the success of the project.

Mimmo Cosenza aka Magomimmo is a programmer and entrepreneur living in Milan, Italy. In the eighties, after graduating in Philosophy of Language, he worked for Rank Xerox and IBM. Then, he joined the Artificial Intelligence lab of ENI S.p.A, the Italian national oil company.

He designed and developed very successful LISP-based applications for the exploration and production departments of ENI. In 1995, after having been in Los Angeles during the rise of the Internet, he founded Sinapsi – an Italian software boutique. In his own country, he is very well known for his involvement in open source communities. In 2012, he founded SmartRM Inc., a startup that applies the Digital Right Management technology for protecting privacy to share confidential information without losing the control of their circulation.

He loves to teach the art of programming. He is the author of *Modern-cljs*, an open source book on the Clojure and ClojureScript programming languages. The book is hosted on <https://github.com/magomimmo/modern-cljs>.

Currently, he is applying machine learning techniques to Big Data by using Clojure on the server-side and ClojureScript on the client-side.

Paul Stadig is a professional software developer living in Crozet, VA, with his wife and three children. He has a B.S. and an M.S. in Computer Science from George Mason University, and he has 16 years of software development experience. He has an insatiable curiosity about the world in general and about programming languages in particular.

He has been involved in the Clojure community since 2008, he was a reviewer for the first edition of *Programming Clojure*, and he is also a contributor to the language. Since 2010, he has been employed at Sonian, where he builds cloud-based distributed systems in Clojure.

Miki Tebeka has been shipping software for more than 10 years. He has developed a wide variety of products from assemblers and linkers to news trading systems and cloud infrastructures. Miki currently works on the data pipeline at Demand Media. In his free time, Miki is active in several open source communities.

www.PacktPub.com

Support files, eBooks, discount offers and more

You might want to visit www.PacktPub.com for support files and downloads related to your book.

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.PacktPub.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at service@packtpub.com for more details.

At www.PacktPub.com, you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.



<http://PacktLib.PacktPub.com>

Do you need instant solutions to your IT questions? PacktLib is Packt's online digital book library. Here, you can access, read and search across Packt's entire library of books.

Why Subscribe?

- Fully searchable across every book published by Packt
- Copy and paste, print and bookmark content
- On demand and accessible via web browser

Free Access for Packt account holders

If you have an account with Packt at www.PacktPub.com, you can use this to access PacktLib today and view nine entirely free books. Simply use your login credentials for immediate access.

Table of Contents

| | |
|-----------------------------------------------------|-----------|
| Preface | 1 |
| Chapter 1: Performance by Design | 5 |
| Usecase classification | 5 |
| User-facing software | 6 |
| Computational and data-processing tasks | 6 |
| CPU bound | 6 |
| Memory bound | 7 |
| Cache bound | 7 |
| Input/Output (I/O) bound | 7 |
| Online transaction processing (OLTP) | 8 |
| Online analytical processing (OLAP) | 8 |
| Batch processing | 9 |
| Structured approach for performance | 9 |
| Performance vocabulary | 10 |
| Latency | 10 |
| Throughput | 11 |
| Bandwidth | 11 |
| Baseline and benchmark | 12 |
| Profiling | 12 |
| Performance optimization | 13 |
| Concurrency and parallelism | 13 |
| Resource utilization | 14 |
| Workload | 14 |
| Latency numbers every programmer should know | 14 |
| Summary | 15 |

| | |
|------------------------------------------------------------|-----------|
| Chapter 2: Clojure Abstractions | 17 |
| Non-numeric scalars and interning | 18 |
| Identity, value, and epochal time model | 19 |
| Variables and mutation | 20 |
| Collection types | 21 |
| Persistent data structures | 21 |
| Constructing less-used data structures | 22 |
| Complexity guarantee | 23 |
| Concatenation of persistent data structures | 24 |
| Sequences and laziness | 25 |
| Laziness | 25 |
| Laziness in data structure operations | 26 |
| Constructing lazy sequences | 27 |
| Transients | 29 |
| Fast repetition | 30 |
| Performance miscellanea | 31 |
| Disabling assertions in production | 31 |
| Destructuring | 31 |
| Recursion and tail-call optimization (TCO) | 32 |
| Premature end in reduce | 33 |
| Multimethods versus protocols | 33 |
| Inlining | 33 |
| Summary | 34 |
| Chapter 3: Leaning on Java | 35 |
| Inspect the equivalent Java source for Clojure code | 35 |
| Create a new project | 36 |
| Compile Clojure sources into Java bytecode | 36 |
| Decompile the .class files into Java source | 36 |
| Numerics, boxing, and primitives | 38 |
| Arrays | 39 |
| Reflection and type hints | 42 |
| Array of primitives | 43 |
| Primitives | 43 |
| Macros and metadata | 44 |
| Miscellaneous | 44 |
| Using array/numeric libraries for efficiency | 45 |
| HipHip | 45 |
| primitive-math | 48 |
| Resorting to Java and native code | 48 |
| Proteus – mutable locals in Clojure | 49 |

| | |
|-------------------------------------------------------|-----------|
| Summary | 50 |
| Chapter 4: Host Performance | 51 |
| The hardware | 51 |
| Processors | 52 |
| Branch prediction | 52 |
| Instruction scheduling | 52 |
| Threads and cores | 53 |
| Memory systems | 54 |
| Cache | 55 |
| Interconnect | 55 |
| Storage and networking | 56 |
| The Java Virtual Machine | 56 |
| The just-in-time (JIT) compiler | 56 |
| Memory organization | 58 |
| HotSpot heap and garbage collection | 60 |
| Measuring memory (heap/stack) usage | 60 |
| Measuring latency with Criterion | 62 |
| Criterion and Leiningen | 63 |
| Summary | 64 |
| Chapter 5: Concurrency | 65 |
| Low-level concurrency | 65 |
| Hardware memory barrier instructions | 66 |
| Java support and its Clojure equivalent | 66 |
| Atomic updates and state | 68 |
| Atomic updates in Java | 68 |
| Clojure's support for atomic updates | 69 |
| Asynchronous agents and state | 70 |
| Asynchrony, queuing, and error handling | 72 |
| Advantages of agents | 73 |
| Nesting | 74 |
| Coordinated transactional ref and state | 74 |
| Ref characteristics | 75 |
| Ref history and intransaction deref operations | 76 |
| Transaction retries and barging | 77 |
| Upping transaction consistency with ensure | 77 |
| Fewer transaction retries with commutative operations | 78 |
| Agents can participate in transactions | 78 |
| Nested transactions | 79 |
| Performance considerations | 80 |
| Dynamic var binding and state | 80 |

| | |
|---------------------------------------------------------|------------|
| Validating and watching the reference types | 81 |
| Java concurrent data structures | 82 |
| Concurrent maps | 83 |
| Concurrent queues | 84 |
| Clojure support for concurrent queues | 86 |
| Concurrency with threads | 86 |
| JVM support for threads | 87 |
| Thread pools in the JVM | 87 |
| Clojure concurrency support | 88 |
| Asynchronous execution with Futures | 88 |
| Anticipated asynchronous execution result with promises | 90 |
| Clojure parallelization and the JVM | 90 |
| Moore's law | 90 |
| Amdahl's law | 91 |
| Clojure support for parallelization | 91 |
| pmap | 91 |
| pcalls | 92 |
| pvalues | 92 |
| Java 7's fork/join framework | 92 |
| Parallelism with reducers | 93 |
| Reducible, reducer function, reduction transformation | 93 |
| Realizing reducible collections | 94 |
| Foldable collections and parallelism | 94 |
| Summary | 95 |
| Chapter 6: Optimizing Performance | 97 |
| A tiny statistics terminology primer | 98 |
| Median, first quartile, and third quartile | 98 |
| Percentile | 99 |
| Variance and standard deviation | 100 |
| Understanding criterium output | 101 |
| Guided performance objectives | 102 |
| Performance testing | 102 |
| Test environment | 102 |
| What to test | 103 |
| Measuring latency | 103 |
| Measuring throughput | 104 |
| Load, stress, and endurance tests | 104 |
| Performance monitoring | 105 |
| Introspection | 105 |
| JVM instrumentation via JMX | 106 |

| | |
|-------------------------------------------|------------|
| Profiling | 106 |
| OS and CPU-cache-level profiling | 108 |
| I/O profiling | 108 |
| Performance tuning | 108 |
| JVM tuning | 109 |
| I/O tuning and backpressure | 110 |
| Summary | 110 |
| Chapter 7: Application Performance | 111 |
| Data sizing | 111 |
| Reduced serialization | 112 |
| Chunking to reduce memory pressure | 113 |
| Sizing for file/network operations | 113 |
| Sizing for JDBC query results | 114 |
| Resource pooling | 115 |
| JDBC resource pooling | 116 |
| I/O batching and throttling | 116 |
| JDBC batch operations | 117 |
| Batch support at API level | 118 |
| Throttling requests to services | 119 |
| Precomputing and caching | 119 |
| Concurrent pipelines | 120 |
| Distributed pipelines | 121 |
| Applying back pressure | 121 |
| Thread pool queues | 122 |
| Servlet containers like Tomcat and Jetty | 122 |
| HTTP Kit | 123 |
| Performance and queuing theory | 123 |
| Little's Law | 124 |
| Summary | 124 |
| Index | 125 |

Preface

Clojure is a remarkably high-performance language despite its dynamic nature. What really strikes you though is the fact that it combines performance with fundamental simplicity and pragmatism, which makes it such a joy to program in. Over the last six years since its first public release, Clojure has been heavily tested and deployed in production by many people and organizations across various domains. Its user base has grown rapidly during this period.

Clojure High Performance Programming is all about Clojure running on the Java Virtual Machine. The JVM has a reputation of being a robust platform to develop and deploy applications on. In this book, we take a deeper look at the performance characteristics of various features of Clojure and the underlying environment. We also explore what it takes to build well-performing software. We begin with the performance fundamentals and gradually proceed over to Clojure and other matters you may have to deal with while writing high-performance applications.

Understanding and achieving performance is both an art and a science, just like writing good software. Remember the big picture in the back of your mind but also be prepared to get into the details with measurement tools. More importantly, know how the software works and keenly study the environment in which it runs. I hope this book will help you on that path.

What this book covers

Chapter 1, Performance by Design, classifies the various use cases with respect to performance and analyzes how to interpret their performance aspects and needs.

Chapter 2, Clojure Abstractions, is a guided tour of various Clojure data structures, abstractions (persistent data structures, vars, macros, and so on), and their performance characteristics.