

BIG CPU, BIG DATA

**Solving the World's Toughest
Computational Problems
with Parallel Computing**

Second Edition

Alan Kaminsky

Department of Computer Science
B. Thomas Golisano College of Computing and Information Sciences
Rochester Institute of Technology

Copyright © 2019 by Alan Kaminsky. All rights reserved.
Second edition, January 2019
Published by Barnes & Noble Press
ISBN 9781987019742

The program source files listed in this book are part of the Parallel Java 2 Library (“The Library”). The Library is copyright © 2013–2019 by Alan Kaminsky. All rights reserved.

The Library is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

The Library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with The Library. If not, see <https://www.gnu.org/licenses/>.

You can get the Parallel Java 2 Library at <https://www.cs.rit.edu/~ark/pj2.shtml>.

Front cover image: The U.S. Department of Energy’s Sequoia supercomputer located at the Lawrence Livermore National Laboratory, in Livermore, California, USA. Photograph courtesy of Lawrence Livermore National Laboratory.

Alan Kaminsky
Professor
Department of Computer Science
B. Thomas Golisano College of Computing and Information Sciences
Rochester Institute of Technology
ark@cs.rit.edu
<https://www.cs.rit.edu/~ark/>

Preface

With the book *BIG CPU, BIG DATA*, my goal is to teach you how to write parallel programs that take full advantage of the vast processing power of modern multicore computers, compute clusters, and graphics processing unit (GPU) accelerators. Everyone's computer nowadays is a parallel computer, and so every programmer needs to know how to write parallel programs.

I'm not going to teach you parallel programming using popular parallel libraries like MPI, OpenMP, and OpenCL. (If you're interested in learning those, plenty of other books are available.) Why? Two reasons:

- I prefer to program in Java. The aforementioned libraries do not, and in my belief never will, support Java.
- In my experience, teaching and learning parallel programming with the aforementioned libraries is more difficult than with Java.

Instead, I'm going to use my *Parallel Java 2 Library* (PJ2) in this book. PJ2 is free, GNU GPL licensed software available from my web site (<https://www.cs.rit.edu/~ark/pj2.shtml>). You can download the complete source files, compiled class files, and Javadoc documentation. PJ2 requires Java Development Kit (JDK) 1.7 or higher. Installation instructions are included in the Javadoc.

PJ2 is suitable both for teaching and learning parallel programming and for real-world parallel program development. I use PJ2 and its predecessor, the Parallel Java Library (PJ), in my cryptography research and in my research on combinatorial optimization problems. Others have used PJ to do page rank calculations, ocean ecosystem modeling, salmon population modeling and analysis, medication scheduling for patients in long term care facilities, three-dimensional complex-valued fast Fourier transforms for electronic structure analysis and X-ray crystallography, and Monte Carlo simulation of electricity and gas markets. PJ was also incorporated into the IQM open source Java image processing application.

I am happy to answer general questions about PJ2, receive bug reports, and entertain requests for additional features. Please contact me by email at ark@cs.rit.edu. I regret that I am unable to provide technical support, specific

installation instructions for your system, or advice about configuring your parallel computer hardware.

More fundamental than the language or library, however, are parallel programming concepts and patterns, such as work sharing parallel loops, parallel reduction, and communication and coordination. Whether you use OpenMP's compiler directives, MPI's message passing subroutines, or PJ2's Java classes, the concepts and patterns are the same. Only the syntax differs. Once you've learned parallel programming in Java with PJ2, you'll be able to apply the same concepts and patterns in C, Fortran, or other languages with OpenMP, MPI, or other libraries.

To study parallel programming with this book, you'll need the following prerequisite knowledge: Java programming; C programming (for GPU programs); computer organization concepts (CPU, memory, cache, and so on); operating system concepts (threads, thread synchronization).

My pedagogical style is to teach by example. Accordingly, this book consists of a series of complete parallel program examples that illustrate various aspects of parallel programming. The example programs' source code is listed on the right-hand pages, and explanatory narrative is on the left-hand pages. The example source code is also included in the PJ2 download. To write programs well, you must first learn to read programs; so please avoid the temptation to gloss over the source code listings, and carefully study both the source code and the explanations.

Also study the PJ2 Javadoc documentation for the various classes used in the example programs. The Javadoc includes comprehensive descriptions of each class and method. Space does not permit describing all the classes in detail in this book.

The book consists of these parts:

- Part I covers introductory concepts.
- Part II covers parallel programming for multicore computers.
- Part III covers parallel programming for compute clusters.
- Part IV covers parallel programming on GPUs.
- Part V covers big data parallel programming using map-reduce.
- Appendix A compares the performance of Java/PJ2 with C/OpenMP.

Instructors: There are no PowerPoint slides to go with this book. Slide shows have their place, but the classroom is not it. Nothing is guaranteed to put students to sleep faster than a PowerPoint lecture. An archive containing all the book's illustrations in PNG format is available from the book's web site (<https://www.cs.rit.edu/~ark/bcbd/>); please feel free to use these to develop your own instructional materials.

Preface to the Second Edition

In the three years since the first edition of *BIG CPU, BIG DATA* was published, I have used the book as a text for several parallel computing courses. I have continued to use the Parallel Java 2 Library (PJ2) to develop parallel programs for my research on cryptography and combinatorial optimization problems, and I have added new classes to PJ2 as I have encountered the need for additional capabilities. The second edition of *BIG CPU, BIG DATA* incorporates feedback from classroom use of the book and adds new material reflecting the new capabilities of PJ2. Most of the text, however, is the same as the first edition.

Major revisions to the book include the following.

I broke Part I, Preliminaries, into several chapters and added a new chapter on supercomputers and their benchmarks.

I changed the design of the minimum vertex cover programs in Chapters 15, 16, and 26 and the Hamiltonian cycle programs in Chapters 17 and 27. The programs now obtain the graph to be analyzed from a *graph specification* object.

I further changed the design of the cluster parallel Hamiltonian cycle program in Chapter 27 to use PJ2's new *cluster work queue* capability.

I added a new Chapter 28 with a cluster parallel big prime number finding program to illustrate PJ2's *on-demand tasks* capability.

The Rochester Institute of Technology Computer Science Department's *tardis* cluster parallel computer, which I used for all the running time measurements in the first edition, has been upgraded. The old *tardis* machine had ten quad-core nodes, each with 8 GB of main memory. The new *tardis* machine has ten dodeca-core (12-core) nodes, each with 64 GB of main memory. I have executed all the programs in the book on the new *tardis* machine and have updated the running time measurements. For the scaling studies, I have scaled the parallel programs up to 12 cores (on a single node) or 120 cores (on the entire cluster).

The RIT Computer Science Department's *kraken* GPU accelerated parallel computer has also been upgraded. The old *kraken* machine had four Nvidia Tesla C2075 448-core GPU cards. The new *kraken* machine has four Nvidia Tesla K40c 2880-core GPU cards. I have executed all the GPU pro-

grams in the book on the new kraken machine and have updated the running time measurements.

As with the first edition, there are no PowerPoint slides to go with this book. An archive containing all the book's illustrations in PNG format is available from the book's web site (https://www.cs.rit.edu/~ark/bcbd_2/).

I dedicate the second edition of *BIG CPU, BIG DATA* to the folks who run the Equal Grounds Coffee House in Rochester, New York (<http://equal-grounds.com/>). Much of the work of writing this edition took place at their tables, using their WiFi connection, fortified with their café mochas, iced lattes, and fruit smoothies.

Alan Kaminsky
January 2019

Table of Contents

Preface	iii
Preface to the Second Edition	v
Table of Contents	vii
Source Code Listings	ix

Part I. Preliminaries

Chapter 1. What Is Parallel Computing?	3
Chapter 2. Parallel Hardware	11
Chapter 3. Parallel Software	21
Chapter 4. Parallel Applications	31
Chapter 5. Supercomputers	35

Part II. Tightly Coupled Multicore

Chapter 6. Parallel Loops	45
Chapter 7. Parallel Loop Schedules	57
Chapter 8. Parallel Reduction	69
Chapter 9. Reduction Variables	79
Chapter 10. Load Balancing	97
Chapter 11. Overlapping	109
Chapter 12. Sequential Dependencies	123
Chapter 13. Strong Scaling	135
Chapter 14. Weak Scaling	149
Chapter 15. Exhaustive Search	161
Chapter 16. Heuristic Search	185
Chapter 17. Parallel Work Queues	197

Part III. Loosely Coupled Cluster

Chapter 18. Massively Parallel	219
Chapter 19. Hybrid Parallel	231
Chapter 20. Tuple Space	239
Chapter 21. Cluster Parallel Loops	251
Chapter 22. Cluster Parallel Reduction	267
Chapter 23. Cluster Load Balancing	277

Part III. Loosely Coupled Cluster (cont.)

Chapter 24. File Output on a Cluster	285
Chapter 25. Interacting Tasks	299
Chapter 26. Cluster Heuristic Search	321
Chapter 27. Cluster Work Queues	329
Chapter 28. On-Demand Tasks	339

Part IV. GPU Acceleration

Chapter 29. GPU Massively Parallel	355
Chapter 30. GPU Parallel Reduction	375
Chapter 31. Multi-GPU Programming	389
Chapter 32. GPU Sequential Dependencies	397
Chapter 33. Objects on the GPU	413
Chapter 34. GPU Heuristic Search	429

Part V. Big Data

Chapter 35. Basic Map-Reduce	455
Chapter 36. Cluster Map-Reduce	471
Chapter 37. Big Data Analysis	487
Chapter 38. Map-Reduce Meets GIS	497

Appendices

Appendix A. Clash of the Titans: C vs. Java	513
Index	523

Source Code Listings

Part II. Tightly Coupled Multicore

Listing 6.1. PrimeSeq.java	47
Listing 6.2. PrimeSmp.java	51
Listing 7.1. MineCoinSeq.java	59
Listing 7.2. MineCoinSmp.java	61
Listing 8.1. PiSeq.java	71
Listing 8.2. PiSmp.java	73
Listing 9.1. Histogram.java	83
Listing 9.2. StatTestSeq.java	87
Listing 9.3. HistogramVbl.java	91
Listing 9.4. StatTestSmp.java	91
Listing 10.1. TotientSeq.java	99
Listing 10.2. TotientSmp.java	101
Listing 11.1. MandelbrotSmp.java	113
Listing 11.2. MandelbrotSmp2.java (relevant portion)	119
Listing 12.1. ZombieSmp.java	129
Listing 15.1. GraphSpec.java	165
Listing 15.2. RandomGraph.java	165
Listing 15.3. AMGraph64.java	171
Listing 15.4. MinVerCovSmp.java	173
Listing 16.1. MinVerCovSmp3.java	189
Listing 17.1. HamCycState.java	205
Listing 17.2. HamCycStateSmp.java	209
Listing 17.3. HamCycSmp.java	211

Part III. Loosely Coupled Cluster

Listing 18.1. MineCoinClu.java	223
Listing 19.1. MineCoinClu2.java	235
Listing 20.1. Tuple subclass MineCoinResult	247
Listing 21.1. MineCoinClu3.java	255
Listing 22.1. PiClu.java	271
Listing 23.1. TotientClu.java	279
Listing 24.1. MandelbrotClu.java	289

Part III. Loosely Coupled Cluster (cont.)

Listing 25.1. ZombieClu.java	305
Listing 26.1. MinVerCovClu3.java	325
Listing 27.1. HamCycClu.java	333
Listing 28.1. OddPrimeList.java	341
Listing 28.2. BigPrimeClu.java	343

Part IV. GPU Acceleration

Listing 29.1. OuterProductGpu.cu	363
Listing 29.2. OuterProductGpu.java	367
Listing 30.1. PiGpu.java	377
Listing 30.2. PiGpu.cu	381
Listing 31.1. PiGpu2.java	391
Listing 32.1. ZombieGpu.cu	401
Listing 32.2. ZombieGpu.java	403
Listing 33.1. ZombieGpu2.cu	415
Listing 33.2. ZombieGpu2.java	419
Listing 34.1. WV.cu	433
Listing 34.2. WalkSackGpu.cu	435
Listing 34.3. WV.java	439
Listing 34.4. KnapsackProblem.java	443
Listing 34.5. KnapsackSC.java	443
Listing 34.6. WalkSackGpu.java	445

Part V. Big Data

Listing 35.1. Concordance01.java	463
Listing 35.2. Combiner.java	465
Listing 36.1. WebLog01.java	477
Listing 36.2. WebLog05.java, class MyMapper	481
Listing 37.1. MaxTemp01.java	489
Listing 38.1. Nola911Block.java	499
Listing 38.2. Nola911.java	503

Appendices

Listing A.1. PiSeq.c	515
Listing A.2. PiSmp.c	517

Everything about big data. All posts. Top. ClickHouse stores data in compressed form. When running queries, ClickHouse tries to do as little as possible, in order to conserve CPU resources. In many cases, all the potentially time-consuming computations are already well optimized, plus the user wrote a well thought-out query. Then all that's left to do is to perform decompression. So why does LZ4 decompression becomes a bottleneck? LZ4 seems like an extremely light algorithm: the data decompression rate is usually from 1 to 3 GB/s per processor core, depending on the data. This is much faster than the typical disk subsystem. Moreov