

O'REILLY®



Optimized C++

C++性能优化手册 (影印版)

東南大學出版社

Kurt Guntheroth 著

C++性能优化手册 (影印版)

Optimized C++

Kurt Guntheroth 著

Beijing • Boston • Farnham • Sebastopol • Tokyo

O'REILLY®

O'Reilly Media, Inc. 授权东南大学出版社出版

南京 东南大学出版社

图书在版编目(CIP)数据

C++ 性能优化手册:英文/(美)库尔特·冈瑟罗思
(Kurt Guntheroth)著. —影印本. —南京:东南大学出版社,
2017.2 (2017.10重印)

书名原文:Optimized C++

ISBN 978-7-5641-6999-2

I. ①C… II. ①库… III. ①C 语言—程序设计—
手册—英文 IV. ①TP312.8-62

中国版本图书馆 CIP 数据核字(2017)第 008142 号

图字:10-2015-256 号

© 2016 by O'Reilly Media, Inc.

Reprint of the English Edition, jointly published by O'Reilly Media, Inc. and Southeast University Press,
2017. Authorized reprint of the original English edition, 2015 O'Reilly Media, Inc., the owner of all rights to
publish and sell the same.

All rights reserved including the rights of reproduction in whole or in part in any form.

英文原版由 O'Reilly Media, Inc. 出版 2016。

英文影印版由东南大学出版社出版 2017。此影印版的出版和销售得到出版权和销售权的所有者
—— O'Reilly Media, Inc. 的许可。

版权所有,未得书面许可,本书的任何部分和全部不得以任何形式重制。

C++ 性能优化手册(影印版)

出版发行:东南大学出版社

地 址:南京四牌楼 2 号 邮编:210096

出 版 人:江建中

网 址: <http://www.seupress.com>

电子邮件: press@seupress.com

印 刷:江苏凤凰数码印务有限公司

开 本:787 毫米×980 毫米 16 开本

印 张:24.25

字 数:475 千字

版 次:2017 年 2 月第 1 版

印 次:2017 年 10 月第 3 次印刷

书 号:ISBN 978-7-5641-6999-2

定 价:89.00 元

本社图书若有印装质量问题,请直接与营销部联系。电话(传真):025-83791830

Everyone thanks their spouse for helping make a book possible. It's trite, I know. My wife Renee Ostler made this book possible by giving me permission to take months off work, by giving me time and space to focus on writing, and by staying up late asking me questions about optimizing C++ code, even though she wasn't particularly engaged by the topic, just to show her support. She made this project important to her because it was important to me. No author could ask for more.

Preface

Hi. My name is Kurt, and I'm a code-aholic.

I have been writing software for over 35 years. I've never worked at Microsoft, or Google, Facebook, Apple, or anywhere else famous. But beyond a few short vacations, I have written code every day of that time. I have spent the last 20 years almost exclusively writing C++ and talking to other very bright developers about C++. This is my qualification to write a book about optimizing C++ code. I have also written *a lot* of English prose, including specifications, comments, manuals, notes, and blog posts (<http://oldhandsblog.blogspot.com>). It has amazed me from time to time that only half of the bright, competent developers I have worked with can string two grammatical English sentences together.

One of my favorite quotes comes by way of a letter from Sir Isaac Newton, in which he writes, "If I have seen farther, it is by standing on the shoulders of giants." I too have stood on the shoulders of giants, and particularly have read their book: elegant little books, like Brian Kernighan and Dennis Ritchie's *The C Programming Language*; smart, ahead-of-the-curve books, like Scott Meyers's *Effective C++* series; challenging, mind-expanding books, like Andrei Alexandrescu's *Modern C++ Design*; careful, precise books, like Bjarne Stroustrup and Margaret Ellis's *The Annotated C++ Reference Manual*. For most of my career, it never crossed my mind that I might someday write a book. Then one day, quite suddenly, I found I needed to write this one.

So why write a book about performance tuning in C++?

At the dawn of the 21st century, C++ was under assault. Fans of C pointed to C++ programs whose performance was inferior to supposedly equivalent code written in C. Famous corporations with big marketing budgets touted proprietary object-oriented languages, claiming C++ was too hard to use, and that their tools were the future. Universities settled on Java for teaching because it came with a free toolchain. As a result of all this buzz, big companies made big-money bets on coding websites and operating systems in Java or C# or PHP. C++ seemed to be on the wane. It was an uncomfortable time for anyone who believed C++ was a powerful, useful tool.

Then a funny thing happened. Processor cores stopped getting faster, but workloads kept growing. Those same companies began hiring C++ programmers to solve their scaling issues. The cost of rewriting code from scratch in C++ became less than the cost of the electricity going into their data centers. All of a sudden, C++ was popular again.

Uniquely among programming languages in wide use in early 2016, C++ offers developers a continuum of implementation choices, ranging from hands-off, automated support to fine manual control. C++ empowers developers to take control of performance trade-offs. This control makes optimization possible.

There are not many books on optimization of C++ code. One of the few is Bulka and Mayhew's meticulously researched but now somewhat dated *Optimizing C++*. The authors appear to have had similar career experiences to mine, and discovered many of the same principles. For readers who are interested in another take on the issues in this book, their book is a good place to start. Also, Scott Meyers, among many others, covers avoiding copy construction extensively and well.

There are enough different things to know about optimization to fill 10 books. I have tried to pick and choose things that seemed to occur frequently in my own work, or that offered the biggest performance wins. To the many readers with their own performance tuning war stories who may wonder why I've said nothing about strategies that worked miracles for them, all I can say is, *so little time, so much to tell*.

I welcome your errata, comments, and favorite optimization strategies at antelope_book@guntheroth.com.

I love the craft of software development. I enjoy endlessly practicing the *kata* of each new loop or interface. At the corner of Sonnet and Science, writing code is a skill so esoteric, an art form so internal, that almost nobody but another practitioner can appreciate it. There is beauty in an elegantly coded function, and wisdom in a powerful idiom well used. Sadly, though, for every epic software poem like Stepanov's Standard Template Library, there are 10,000 drab tomes of uninspired code.

The root purpose of this book is to give every reader permission to think a little harder about the beauty of well-tuned software. Take it and run with it. See farther!

Apology for the Code in This Book

Although I have been writing and optimizing C++ code for over 20 years, most of the code appearing in this book was developed specifically *for* this book. Like all new code, it surely contains defects. I offer my apologies.

I have developed for Windows, Linux, and various embedded systems over the years. The code presented in this book was developed on Windows. The code and the book no doubt show a Windows bias. The lessons of how to optimize C++ code that are

illustrated using Visual Studio on Windows apply equally to Linux, Mac OS X, or any other C++ environment. However, the precise timings of different optimizations depend on the compiler and standard library implementation, and the processor on which the code is tested. Optimization is an experimental science. Taking optimization advice on faith is fraught with negative surprises.

I am aware that compatibility with various other compilers, and with other Unix and embedded systems, can be challenging, and I apologize if the code does not compile on your favorite system. Since this book is not about cross-system compatibility, I have erred on the side of presenting simple code.

The curly-brace indentation style shown here is not my favorite:

```
if (bool_condition) {  
    controlled_statement();  
}
```

However, because it has the advantage of putting the most lines possible on the printed page, I have chosen to use it for the examples throughout this book.

Using Code Examples

Supplemental material (code examples, sample solutions, etc.) is available for download at www.guntheroth.com.

This book is here to help you get your job done. In general, if example code is offered with this book, you may use it in your programs and documentation. You do not need to contact us for permission unless you're reproducing a significant portion of the code. For example, writing a program that uses several chunks of code from this book does not require permission. Selling or distributing a CD-ROM of examples from O'Reilly books does require permission. Answering a question by citing this book and quoting example code does not require permission. Incorporating a significant amount of example code from this book into your product's documentation does require permission.

We appreciate, but do not require, attribution. An attribution usually includes the title, author, publisher, and ISBN. For example: "*Optimized C++* by Kurt Guntheroth (O'Reilly). Copyright 2016 Kurt Guntheroth, 978-1-491-92206-4."

If you feel your use of code examples falls outside fair use or the permission given above, feel free to contact us at permissions@oreilly.com.

Conventions Used in This Book

The following typographical conventions are used in this book:

Plain text

Used for menu titles, menu options, menu buttons, and keyboard accelerators (such as Alt and Control).

Italic

Indicates new terms, URLs, email addresses, pathnames, filenames, and file extensions.

Constant width

Used for program listings, as well as within paragraphs to refer to program elements such as variable or function names, databases, data types, environment variables, statements, and keywords.

Table of Contents

Preface.....	xv
1. An Overview of Optimization.....	1
Optimization Is Part of Software Development	2
Optimization Is Effective	3
It's OK to Optimize	3
A Nanosecond Here, a Nanosecond There	6
Summary of Strategies for Optimizing C++ Code	6
Use a Better Compiler, Use Your Compiler Better	7
Use Better Algorithms	8
Use Better Libraries	9
Reduce Memory Allocation and Copying	10
Remove Computation	11
Use Better Data Structures	12
Increase Concurrency	12
Optimize Memory Management	12
Summary	12
2. Computer Behavior Affecting Optimization.....	15
Lies C++ Believes About Computers	16
The Truth About Computers	17
Memory Is Slow	17
Memory Is Not Accessed in Bytes	18
Some Memory Accesses Are Slower than Others	19
Memory Words Have a Big End and a Little End	20
Memory Has Finite Capacity	20
Instruction Execution Is Slow	21
Making Decisions Is Hard for Computers	22

There Are Multiple Streams of Program Execution	22
Calling into the Operating System Is Expensive	24
C++ Tells Lies Too	24
All Statements Are Not Equally Expensive	24
Statements Are Not Executed in Order	25
Summary	26
3. Measure Performance.....	27
The Optimizing Mindset	28
Performance Must Be Measured	28
Optimizers Are Big Game Hunters	29
The 90/10 Rule	29
Amdahl's Law	31
Perform Experiments	32
Keep a Lab Notebook	34
Measure Baseline Performance and Set Goals	35
You Can Improve Only What You Measure	37
Profile Program Execution	37
Time Long-Running Code	40
"A Little Learning" About Measuring Time	40
Measuring Time with Computers	46
Overcoming Measurement Obstacles	54
Create a Stopwatch Class	58
Time Hot Functions in a Test Harness	62
Estimate Code Cost to Find Hot Code	63
Estimate the Cost of Individual C++ Statements	63
Estimate the Cost of Loops	64
Other Ways to Find Hot Spots	66
Summary	67
4. Optimize String Use: A Case Study.....	69
Why Strings Are a Problem	69
Strings Are Dynamically Allocated	70
Strings Are Values	70
Strings Do a Lot of Copying	71
First Attempt at Optimizing Strings	72
Use Mutating String Operations to Eliminate Temporaries	74
Reduce Reallocation by Reserving Storage	74
Eliminate Copying of String Arguments	75
Eliminate Pointer Dereference Using Iterators	76
Eliminate Copying of Returned String Values	77
Use Character Arrays Instead of Strings	78

Summary of First Optimization Attempt	80
Second Attempt at Optimizing Strings	80
Use a Better Algorithm	80
Use a Better Compiler	82
Use a Better String Library	83
Use a Better Allocator	87
Eliminate String Conversion	88
Conversion from C String to std::string	89
Converting Between Character Encodings	89
Summary	90
5. Optimize Algorithms.....	91
Time Cost of Algorithms	92
Best-Case, Average, and Worst-Case Time Cost	95
Amortized Time Cost	95
Other Costs	96
Toolkit to Optimize Searching and Sorting	96
Efficient Search Algorithms	96
Time Cost of Searching Algorithms	97
All Searches Are Equal When n Is Small	98
Efficient Sort Algorithms	98
Time Cost of Sorting Algorithms	99
Replace Sorts Having Poor Worst-Case Performance	99
Exploit Known Properties of the Input Data	100
Optimization Patterns	100
Precomputation	101
Lazy Computation	102
Batching	102
Caching	103
Specialization	104
Taking Bigger Bites	104
Hinting	105
Optimizing the Expected Path	105
Hashing	105
Double-Checking	105
Summary	106
6. Optimize Dynamically Allocated Variables.....	107
C++ Variables Refresher	108
Storage Duration of Variables	108
Ownership of Variables	111
Value Objects and Entity Objects	112

C++ Dynamic Variable API Refresher	113
Smart Pointers Automate Ownership of Dynamic Variables	116
Dynamic Variables Have Runtime Cost	118
Reduce Use of Dynamic Variables	119
Create Class Instances Staticallly	119
Use Static Data Structures	121
Use <code>std::make_shared</code> Instead of <code>new</code>	124
Don't Share Ownership Unnecessarily	125
Use a "Master Pointer" to Own Dynamic Variables	126
Reduce Reallocation of Dynamic Variables	127
Preallocate Dynamic Variables to Prevent Reallocation	127
Create Dynamic Variables Outside of Loops	128
Eliminate Unneeded Copying	129
Disable Unwanted Copying in the Class Definition	130
Eliminate Copying on Function Call	131
Eliminate Copying on Function Return	132
Copy Free Libraries	134
Implement the "Copy on Write" Idiom	136
Slice Data Structures	137
Implement Move Semantics	137
Nonstandard Copy Semantics: A Painful Hack	138
<code>std::swap()</code> : The Poor Man's Move Semantics	138
Shared Ownership of Entities	139
The Moving Parts of Move Semantics	140
Update Code to Use Move Semantics	141
Subtleties of Move Semantics	142
Flatten Data Structures	145
Summary	146
7. Optimize Hot Statements.....	147
Remove Code from Loops	148
Cache the Loop End Value	149
Use More Efficient Loop Statements	149
Count Down Instead of Up	150
Remove Invariant Code from Loops	151
Remove Unneeded Function Calls from Loops	152
Remove Hidden Function Calls from Loops	155
Remove Expensive, Slow-Changing Calls from Loops	156
Push Loops Down into Functions to Reduce Call Overhead	157
Do Some Actions Less Frequently	158
What About Everything Else?	160
Remove Code from Functions	160

Cost of Function Calls	161
Declare Brief Functions Inline	165
Define Functions Before First Use	165
Eliminate Unused Polymorphism	165
Discard Unused Interfaces	166
Select Implementation at Compile Time with Templates	170
Eliminate Uses of the PIMPL Idiom	171
Eliminate Calls into DLLs	173
Use Static Member Functions Instead of Member Functions	173
Move Virtual Destructor to Base Class	174
Optimize Expressions	174
Simplify Expressions	175
Group Constants Together	176
Use Less-Expensive Operators	177
Use Integer Arithmetic Instead of Floating Arithmetic	177
Double May Be Faster than Float	179
Replace Iterative Computations with Closed Forms	180
Optimize Control Flow Idioms	182
Use switch Instead of if-elseif-else	182
Use Virtual Functions Instead of switch or if	182
Use No-Cost Exception Handling	183
Summary	185
8. Use Better Libraries	187
Optimize Standard Library Use	187
Philosophy of the C++ Standard Library	188
Issues in Use of the C++ Standard Library	188
Optimize Existing Libraries	191
Change as Little as Possible	191
Add Functions Rather than Change Functionality	192
Design Optimized Libraries	192
Code in Haste, Repent at Leisure	193
Parsimony Is a Virtue in Library Design	194
Make Memory Allocation Decisions Outside the Library	194
When in Doubt, Code Libraries for Speed	195
Functions Are Easier to Optimize than Frameworks	195
Flatten Inheritance Hierarchies	196
Flatten Calling Chains	196
Flatten Layered Designs	196
Avoid Dynamic Lookup	198
Beware of 'God Functions'	199
Summary	200

9. Optimize Searching and Sorting.....	201
Key/Value Tables Using <code>std::map</code> and <code>std::string</code>	202
Toolkit to Improve Search Performance	203
Make a Baseline Measurement	204
Identify the Activity to Be Optimized	204
Decompose the Activity to Be Optimized	205
Change or Replace Algorithms and Data Structures	206
Using the Optimization Process on Custom Abstractions	208
Optimize Search Using <code>std::map</code>	208
Use Fixed-Size Character Array Keys with <code>std::map</code>	208
Use C-Style String Keys with <code>std::map</code>	210
Using Map's Cousin <code>std::set</code> When the Key Is in the Value	212
Optimize Search Using the <code><algorithm></code> Header	213
Key/Value Table for Search in Sequence Containers	214
<code>std::find()</code> : Obvious Name, $O(n)$ Time Cost	215
<code>std::binary_search()</code> : Does Not Return Values	216
Binary Search Using <code>std::equal_range()</code>	216
Binary Search Using <code>std::lower_bound()</code>	217
Handcoded Binary Search	218
Handcoded Binary Search using <code>strcmp()</code>	219
Optimize Search in Hashed Key/Value Tables	220
Hashing with a <code>std::unordered_map</code>	221
Hashing with Fixed Character Array Keys	221
Hashing with Null-Terminated String Keys	222
Hashing with a Custom Hash Table	224
Stepanov's Abstraction Penalty	225
Optimize Sorting with the C++ Standard Library	226
Summary	228
10. Optimize Data Structures.....	229
Get to Know the Standard Library Containers	229
Sequence Containers	230
Associative Containers	230
Experimenting with the Standard Library Containers	231
<code>std::vector</code> and <code>std::string</code>	236
Performance Consequences of Reallocation	237
Inserting and Deleting in <code>std::vector</code>	238
Iterating in <code>std::vector</code>	240
Sorting <code>std::vector</code>	241
Lookup with <code>std::vector</code>	241
<code>std::deque</code>	242
Inserting and Deleting in <code>std::deque</code>	243

Iterating in <code>std::deque</code>	245
Sorting <code>std::deque</code>	245
Lookup with <code>std::deque</code>	245
<code>std::list</code>	245
Inserting and Deleting in <code>std::list</code>	247
Iterating in <code>std::list</code>	248
Sorting <code>std::list</code>	248
Lookup with <code>std::list</code>	249
<code>std::forward_list</code>	249
Inserting and Deleting in <code>std::forward_list</code>	250
Iterating in <code>std::forward_list</code>	251
Sorting <code>std::forward_list</code>	251
Lookup in <code>std::forward_list</code>	251
<code>std::map</code> and <code>std::multimap</code>	251
Inserting and Deleting in <code>std::map</code>	252
Iterating in <code>std::map</code>	255
Sorting <code>std::map</code>	255
Lookup with <code>std::map</code>	255
<code>std::set</code> and <code>std::multiset</code>	255
<code>std::unordered_map</code> and <code>std::unordered_multimap</code>	256
Inserting and Deleting in <code>std::unordered_map</code>	260
Iterating in <code>std::unordered_map</code>	260
Lookup with <code>std::unordered_map</code>	261
Other Data Structures	261
Summary	263
11. Optimize I/O.....	265
A Recipe for Reading Files	265
Create a Parsimonious Function Signature	267
Shorten Calling Chains	269
Reduce Reallocation	269
Take Bigger Bites—Use a Bigger Input Buffer	272
Take Bigger Bites—Read a Line at a Time	272
Shorten Calling Chains Again	274
Things That Didn't Help	275
Writing Files	276
Reading from <code>std::cin</code> and Writing to <code>std::cout</code>	277
Summary	278
12. Optimize Concurrency.....	279
Concurrency Refresher	280
A Walk Through the Concurrency Zoo	281

Interleaved Execution	285
Sequential Consistency	286
Races	287
Synchronization	288
Atomicity	289
C++ Concurrency Facilities Refresher	291
Threads	291
Promises and Futures	292
Asynchronous Tasks	295
Mutexes	296
Locks	297
Condition Variables	298
Atomic Operations on Shared Variables	301
On Deck: Future C++ Concurrency Features	304
Optimize Threaded C++ Programs	305
Prefer <code>std::async</code> to <code>std::thread</code>	306
Create as Many Runnable Threads as Cores	308
Implement a Task Queue and Thread Pool	309
Perform I/O in a Separate Thread	310
Program Without Synchronization	310
Remove Code from Startup and Shutdown	313
Make Synchronization More Efficient	314
Reduce the Scope of Critical Sections	314
Limit the Number of Concurrent Threads	315
Avoid the Thundering Herd	316
Avoid Lock Convoys	317
Reduce Contention	317
Don't Busy-Wait on a Single-Core System	319
Don't Wait Forever	319
Rolling Your Own Mutex May Be Ineffective	319
Limit Producer Output Queue Length	320
Concurrency Libraries	320
Summary	322
13. Optimize Memory Management.....	323
C++ Memory Management API Refresher	324
The Life Cycle of Dynamic Variables	324
Memory Management Functions Allocate and Free Memory	325
New-Expressions Construct Dynamic Variables	328
Delete-Expressions Dispose of Dynamic Variables	331
Explicit Destructor Calls Destroy Dynamic Variables	332
High-Performance Memory Managers	333

Provide Class-Specific Memory Managers	335
Fixed-Size-Block Memory Manager	336
Block Arena	338
Adding a Class-Specific operator new()	340
Performance of the Fixed-Block Memory Manager	342
Variations on the Fixed-Block Memory Manager	342
Non-Thread Safe Memory Managers Are Efficient	343
Provide Custom Standard Library Allocators	343
Minimal C++11 Allocator	346
Additional Definitions for C++98 Allocator	347
A Fixed-Block Allocator	352
A Fixed-Block Allocator for Strings	354
Summary	355
Index.....	357