Synchronization Algorithms And Concurrent Programming

Foundations of Multithreaded, Parallel, and Distributed Programming
Concurrent Programming
On Concurrent Programming
Java Concurrency in Practice
Introduction to Concurrency in Programming Languages
Concurrent and Distributed Computing in Java
Shared-Memory Synchronization
Distributed Computing Pearls
Concurrent Programming: Algorithms, Principles, and Foundations
Parallel and Distributed Computation: Numerical Methods
Shared-Memory Synchronization
Synchronization Algorithms and Concurrent Programming
Concurrent Programming in Java
Transactional Memory. Foundations, Algorithms, Tools, and Applications
Principles of Transactional Memory
Pro TBB
Nonsequential and Distributed Programming with Go
Topics in Parallel and Distributed Computing
Programming Concurrency on the JVM
Parallel and Distributed Processing
The Art of Multiprocessor Programming, Revised Reprint
Concurrent Patterns and Best Practices
Learn Concurrent Programming with Go
Principles of Concurrent Programming
Algorithms for Concurrent Systems
Mastering Concurrency Programming with Java 8
Principles of Concurrent and Distributed Programming
Algorithms for Concurrent Systems
Concurrent Programming on Windows
Transactional Memory, Second Edition
Encyclopedia of Parallel Computing
The Origin of Concurrent Programming
Concurrent, Real-Time and Distributed Programming in Java
Current Trends in Concurrency
Structured Concurrent Programming with Operating Systems Applications
The Art of Concurrency
Introduction to Parallel Programming
The SR Programming Language

Understanding Control Flow
Concurrent Programming

CARTER MYLA

Foundations of Multithreaded, Parallel, and Distributed Programming Pearson Education
Here, one of the leading figures in the field provides a comprehensive survey of the subject, beginning with prepositional logic and concluding with concurrent programming. It is based on graduate courses taught at Cornell University and is designed for use as a graduate text. Professor Schneier emphasises the use of formal methods and assertional reasoning using notation and paradigms drawn from programming to drive the exposition, while exercises at the end of each chapter extend and illustrate the main themes covered. As a result, all those interested in studying concurrent computing will find this an invaluable approach to the subject.
Concurrent Programming Springer Science & Business Media
In modern computer science, there exists no truly sequential computing system; and most advanced programming is parallel programming. This is particularly evident in modern application domains like scientific computation, data science, machine intelligence, etc. This lucid introductory textbook will be invaluable to students of computer science and technology, acting as a self-contained primer to parallel programming. It takes the reader from introduction to expertise, addressing a broad gamut of issues. It covers different parallel programming styles, describes parallel architecture, includes parallel programming frameworks and techniques, presents algorithmic and analysis techniques and discusses parallel design and performance issues. With its broad coverage, the book can be useful in a wide range of courses; and can also prove useful as a ready reckoner for professionals in the field.
On Concurrent Programming Prentice Hall
From driving, flying, and swimming, to digging for unknown objects in space exploration, autonomous robots take on varied shapes and sizes. In part, autonomous robots are designed to perform tasks that are too dirty, dull, or dangerous for humans. With nontrivial autonomy and volition, they may soon claim their own place in human society. These robots will be our allies as we strive for understanding our natural and man-made environments and build positive synergies around us. Although we may never perfect replication of biological capabilities in robots, we must harness the inevitable emergence of robots that synchronizes with our own capacities to live, learn, and grow. This book is a snapshot of motivations and methodologies for our collective attempts to transform our lives and enable us to cohabit with robots that work with and for us. It reviews and guides the reader to seminal and continual developments that are the foundations for successful paradigms. It attempts to demystify the abilities and limitations of robots. It is a progress report on the continuing work that will fuel future endeavors. Table of Contents: Part I: Preliminaries/Agency, Motion, and Anatomy/Behaviors / Architectures / Affect/Sensors / Manipulators/Part II: Mobility/Potential Fields/Roadmaps / Reactive Navigation / Multi-Robot Mapping: Brick and Mortar Strategy / Part III: State of the Art / Multi-Robotics Phenomena / Human-Robot Interaction / Fuzzy Control / Decision Theory and Game Theory / Part IV: On the Horizon / Applications: Macro and Micro Robots / References / Author Biography / Discussion
Java Concurrency in Practice Prentice Hall
Principles of Concurrent and Distributed Programming provides an introduction to concurrent programming focusing on general principles and not on specific systems. Software today is inherently concurrent or distributed - from event-based GUI designs to operating and real-time systems to Internet applications. This edition is an introduction to concurrency and examines the growing importance of concurrency constructs embedded in programming languages and of formal methods such as model checking.
Introduction to Concurrency in Programming Languages Elsevier
Mathematics of Computing -- Parallelism.
Concurrent and Distributed Computing in Java John Wiley & Sons
This book provides an introduction to concurrent, real-time, distributed programming with Java object-oriented language support as an algorithm description tool. It describes in particular the mechanisms of synchronization (cooperative and competitive) and sharing of data (internal class, static variables) between threads in Java. He then discusses the use of Java for real-time applications. Consequently, a presentation of the RTSJ (Real Time Specification for Java) specification dedicated to the development of real-time applications in Java is also introduced in this book. Finally, a presentation of programming distributed in Java is presented in this book. We are particularly interested in communication using the TCP Sockets and high-level communication using Java Remote Method Invocation (RMI). The book also contains an annex which contains a practical set of application exercises in relation to the theme of the book. Knowledge of the Java language is a prerequisite for understanding the book.
Shared-Memory Synchronization Reading, Mass. ; Don Mills, Ont. : Addison-Wesley Publishing Company
Transactional memory (TM) is an appealing paradigm for concurrent programming on shared memory architectures. With a TM, threads of an application communicate, and synchronize their actions, via in-memory transactions. Each transaction can perform any number of operations on shared data, and then either commit or abort. When the transaction commits, the effects of all its operations become immediately visible to other transactions; when it aborts, however, those effects are entirely discarded. Transactions are atomic: programmers get the illusion that every transaction executes all its operations instantaneously, at some single and unique point in time. Yet, a TM runs transactions concurrently to leverage the parallelism offered by modern processors. The aim of this book is to provide theoretical foundations for transactional memory. This includes defining a model of a TM, as well as answering precisely when a TM implementation is correct, what kind of properties it can ensure, what are the power and limitations of a TM, and what inherent trade-offs are involved in designing a TM algorithm. While the focus of this book is on the fundamental principles, its goal is to capture the common intuition behind the semantics of TMs and the properties of existing TM

implementations. Table of Contents: Introduction / Shared Memory Systems / Transactional Memory: A Primer / TM Correctness Issues / Implementing a TM / Further Reading / Opacity / Proving Opacity: An Example / Opacity vs.\ Atomicity / Further Reading / The Liveness of a TM / Lock-Based TMs / Obstruction-Free TMs / General Liveness of TMs / Further Reading / Conclusions

Distributed Computing Pearls Pearson

Containing over 300 entries in an A-Z format, the Encyclopedia of Parallel Computing provides easy, intuitive access to relevant information for professionals and researchers seeking access to any aspect within the broad field of parallel computing. Topics for this comprehensive reference were selected, written, and peer-reviewed by an international pool of distinguished researchers in the field. The Encyclopedia is broad in scope, covering machine organization, programming languages, algorithms, and applications. Within each area, concepts, designs, and specific implementations are presented. The highly-structured essays in this work comprise synonyms, a definition and discussion of the topic, bibliographies, and links to related literature. Extensive cross-references to other entries within the Encyclopedia support efficient, user-friendly searchers for immediate access to useful information. Key concepts presented in the Encyclopedia of Parallel Computing include; laws and metrics; specific numerical and non-numerical algorithms; asynchronous algorithms; libraries of subroutines; benchmark suites; applications; sequential consistency and cache coherency; machine classes such as clusters, shared-memory multiprocessors, special-purpose machines and dataflow machines; specific machines such as Cray supercomputers, IBM’s cell processor and Intel’s multicore machines; race detection and auto parallelization; parallel programming languages, synchronization primitives, collective operations, message passing libraries, checkpointing, and operating systems. Topics covered: Speedup, Efficiency, Isoefficiency, Redundancy, Amdahls law, Computer Architecture Concepts, Parallel Machine Designs, Benmarks, Parallel Programming concepts & design, Algorithms, Parallel applications. This authoritative reference will be published in two formats: print and online. The online edition features hyperlinks to cross-references and to additional significant research. Related Subjects: supercomputing, high-performance computing, distributed computing

Concurrent Programming: Algorithms, Principles, and Foundations Pearson Education

Zusammenfassung: This book offers a comprehensive survey of shared-memory synchronization, with an emphasis on "systems-level" issues. It includes sufficient coverage of architectural details to understand correctness and performance on modern multicore machines, and sufficient coverage of higher-level issues to understand how synchronization is embedded in modern programming languages. The primary intended audience for this book is "systems programmers"--the authors of operating systems, library packages, language run-time systems, concurrent data structures, and server and utility programs. Much of the discussion should also be of interest to application programmers who want to make good use of the synchronization mechanisms available to them, and to computer architects who want to understand the ramifications of their design decisions on systems-level code

Parallel and Distributed Computation: Numerical Methods Springer Nature

Software -- Programming Languages.

Shared-Memory Synchronization Springer Nature

Topics in Parallel and Distributed Computing provides resources and guidance for those learning PDC as well as those teaching students new to the discipline. The pervasiveness of computing devices containing multicore CPUs and GPUs, including home and office PCs, laptops, and mobile devices, is making even common users dependent on parallel processing. Certainly, it is no longer sufficient for even basic programmers to acquire only the traditional sequential programming skills. The preceding trends point to the need for imparting a broad-based skill set in PDC technology. However, the rapid changes in computing hardware platforms and devices, languages, supporting programming environments, and research advances, poses a challenge both for newcomers and seasoned computer scientists. This edited collection has been developed over the past several years in conjunction with the IEEE technical committee on parallel processing (TCPP), which held several workshops and discussions on learning parallel computing and integrating parallel concepts into courses throughout computer science curricula. Contributed and developed by the leading minds in parallel computing research and instruction Provides resources and guidance for those learning PDC as well as those teaching students new to the discipline Succinctly addresses a range of parallel and distributed computing topics Pedagogically designed to ensure understanding by experienced engineers and newcomers Developed over the past several years in conjunction with the IEEE technical committee on parallel processing (TCPP), which held several workshops and discussions on learning parallel computing and integrating parallel concepts

Synchronization Algorithms and Concurrent Programming Simon and Schuster

This highly acclaimed work, first published by Prentice Hall in 1989, is a comprehensive and theoretically sound treatment of parallel and distributed numerical methods. It focuses on algorithms that are naturally suited for massive parallelization, and it explores the fundamental convergence, rate of convergence, communication, and synchronization issues associated with such algorithms. This is an extensive book, which aside from its focus on parallel and distributed algorithms, contains a wealth of material on a broad variety of computation and optimization topics. It is an excellent supplement to several of our other books, including Convex Optimization Algorithms (Athena Scientific, 2015), Nonlinear Programming (Athena Scientific, 1999), Dynamic Programming and Optimal Control (Athena Scientific, 2012), Neuro-Dynamic Programming (Athena Scientific, 1996), and Network Optimization (Athena Scientific, 1998). The on-line edition of the book contains a 95-page solutions manual.

Concurrent Programming in Java Pragmatic Bookshelf

The advent of multicore processors has renewed interest in the idea of incorporating transactions into the programming model used to write parallel programs. This approach, known as transactional memory, offers an alternative, and hopefully better, way to coordinate concurrent threads. The ACI (atomicity, consistency, isolation) properties of transactions provide a foundation to ensure that concurrent reads and writes of shared data do not produce inconsistent or incorrect results. At a higher level, a computation wrapped in a transaction executes atomically - either it completes successfully and commits its result in its entirety or it aborts. In addition, isolation ensures the transaction produces the same result as if no other transactions were executing concurrently. Although transactions are not a parallel programming panacea, they shift much of the burden of synchronizing and coordinating parallel computations from a programmer to a compiler, to a

language runtime system, or to hardware. The challenge for the system implementers is to build an efficient transactional memory infrastructure. This book presents an overview of the state of the art in the design and implementation of transactional memory systems, as of early spring 2010. Table of Contents: Introduction / Basic Transactions / Building on Basic Transactions / Software Transactional Memory / Hardware-Supported Transactional Memory / Conclusions

Transactional Memory. Foundations, Algorithms, Tools, and Applications Springer Science & Business Media

Der Band bietet eine kompakte Einführung in die Nichtsequentielle Programmierung als gemeinsamen Kern von Vorlesungen über Betriebssysteme, Verteilte Systeme, Parallele Algorithmen, Echtzeitprogrammierung und Datenbanktransaktionen. Basiskonzepte zur Synchronisation und Kommunikation nebenläufiger Prozesse werden systematisch dargestellt: Schlösser, Semaphore, Monitore, lokaler und netzweiter Botschaftenaustausch. Die Algorithmen sind in der Programmiersprache Google Go formuliert, mit der viele Synchronisationskonzepte ausgedrückt werden können.

Principles of Transactional Memory Cambridge University Press

Master the principles and techniques of multithreaded programming with the Java 8 Concurrency API About This Book Implement concurrent applications using the Java 8 Concurrency API and its new components Improve the performance of your applications or process more data at the same time, taking advantage of all of your resources. Construct real-world examples related to machine learning, data mining, image processing, and client/server environments Who This Book Is For If you are a competent Java developer with a good understanding of concurrency but have no knowledge of how to effectively implement concurrent programs or use streams to make processes more efficient, then this book is for you. What You Will Learn Design concurrent applications by converting a sequential algorithm into a concurrent one Discover how to avoid all the possible problems you can get in concurrent algorithms Use the Executor framework to manage concurrent tasks without creating threads Extend and modify Executors to adapt their behavior to your needs Solve problems using the divide and conquer technique and the Fork/Join framework Process massive data sets with parallel streams and Map/Reduce implementation Control data-race conditions using concurrent data structures and synchronization mechanisms Test and monitor concurrent applications In Detail Concurrency programming allows several large tasks to be divided into smaller sub-tasks, which are further processed as individual tasks that run in parallel. All the sub-tasks are combined together once the required results are achieved; they are then merged to get the final output. The whole process is very complex. This process goes from the design of concurrent algorithms to the testing phase where concurrent applications need extra attention. Java includes a comprehensive API with a lot of ready-to-use components to implement powerful concurrency applications in an easy way, but with a high flexibility to adapt these components to your needs. The book starts with a full description of design principles of concurrent applications and how to parallelize a sequential algorithm. We'll show you how to use all the components of the Java Concurrency API from basics to the most advanced techniques to implement them in powerful concurrency applications in Java. You will be using real-world examples of complex algorithms related to machine learning, data mining, natural language processing, image processing in client / server environments. Next, you will learn how to use the most important components of the Java 8 Concurrency API: the Executor framework to execute multiple tasks in your applications, the phaser class to implement concurrent tasks divided into phases, and the Fork/Join framework to implement concurrent tasks that can be split into smaller problems (using the divide and conquer technique). Toward the end, we will cover the new inclusions in Java 8 API, the Map and Reduce model, and the Map and Collect model. The book will also teach you about the data structures and synchronization utilities to avoid data-race conditions and other critical problems. Finally, the book ends with a detailed description of the tools and techniques that you can use to test a Java concurrent application. Style and approach A complete guide implementing real-world examples with algorithms related to machine learning, data mining, and natural language processing in client/server environments. All the examples are explained in a step-by-step approach.

Pro TBB Springer Nature

Concurrent and Distributed Computing in Java addresses fundamental concepts in concurrent computing with Java examples. The book consists of two parts. The first part deals with techniques for programming in shared-memory based systems. The book covers concepts in Java such as threads, synchronized methods, waits, and notify to expose students to basic concepts for multi-threaded programming. It also includes algorithms for mutual exclusion, consensus, atomic objects, and wait-free data structures. The second part of the book deals with programming in a message-passing system. This part covers resource allocation problems, logical clocks, global property detection, leader election, message ordering, agreement algorithms, checkpointing, and message logging. Primarily a textbook for upper-level undergraduates and graduate students, this thorough treatment will also be of interest to professional programmers.

Nonsequential and Distributed Programming with Go Packt Publishing Ltd

Concurrent programming and operating systems. Concurrency problems and language features. A sequential programming language: SP/k; A concurrent programming language CSP/k. Examples of concurrent programs. Design of an operating system: input and output spoolers. Design of an operating system: the executive; Implementing a Kernel. Appendices.

Topics in Parallel and Distributed Computing Morgan Kaufmann

This book constitutes the refereed proceedings of 10 international workshops held in conjunction with the merged 1998 IPPS/SPDP symposia, held in Orlando, Florida, US in March/April 1998. The volume comprises 118 revised full papers presenting cutting-edge research or work in progress. In accordance with the workshops covered, the papers are organized in topical sections on reconfigurable architectures, run-time systems for parallel programming, biologically inspired solutions to parallel processing problems, randomized parallel computing, solving combinatorial optimization problems in parallel, PC based networks of workstations, fault-tolerant parallel and distributed systems, formal methods for parallel programming, embedded HPC systems and applications, and parallel and distributed real-time systems.

Programming Concurrency on the JVM Addison Wesley

An essential reader containing 19 important papers on the invention and early development of concurrent programming and its relevance to computer science and computer engineering. All of them are written by the pioneers in concurrent programming, including Brinch Hansen himself, and

have introductions added that summarize the papers and put them in perspective. The editor provides an overview chapter and neatly places all developments in perspective with chapter introductions and expository apparatus. Essential resource for graduates, professionals, and researchers in CS with an interest in concurrent programming principles. A familiarity with operating system principles is assumed.

Parallel and Distributed Processing Springer Science & Business Media

Threads are a fundamental part of the Java platform. As multicore processors become the norm, using concurrency effectively becomes essential for building high-performance applications. Java SE 5 and 6 are a huge step forward for the development of concurrent applications, with improvements to the Java Virtual Machine to support high-performance, highly scalable concurrent classes and a rich set of new concurrency building blocks. In Java Concurrency in Practice, the creators of these new facilities explain not only how they work and how to use them, but also the motivation and design patterns behind them. However, developing, testing, and debugging multithreaded programs can still be very difficult; it is all too easy to create concurrent programs that appear to work, but fail when it matters most: in production, under heavy load. Java Concurrency in Practice arms readers with both the theoretical underpinnings and concrete techniques for building reliable, scalable, maintainable concurrent applications. Rather than simply offering an inventory of concurrency APIs and mechanisms, it provides design rules, patterns, and mental models that make it easier to build concurrent programs that are both correct and performant. This book covers: Basic concepts of concurrency and thread safety Techniques for building and composing thread-safe classes Using the concurrency building blocks in java.util.concurrent Performance optimization dos and don'ts Testing concurrent programs Advanced topics such as atomic variables, nonblocking algorithms, and the Java Memory Model