Size 60 69mb Parallel Computer Organization And Design

Ambition will fuel him. Competition will drive him. But power has its price. It is the morning of the reaping that will kick off the tenth annual Hunger Games. In the Capitol, eighteen-year-old Coriolanus Snow is preparing for his one shot at glory as a mentor in the Games. The once-mighty house of Snow has fallen on hard times, its fate hanging on the slender chance that Coriolan outcharm, outwit, and outmaneuver his fellow students to mentor the winning tribute. The odds are against him. He's been given the humiliating assignment of mentoring the female tribute from District 12, the lowest of the low. Their fates are now completely intertwined - every choice Coriolanus makes could lead to favor or failure, triumph or ruin. Inside the arena, it will be a figh the arena, Coriolanus starts to feel for his doomed tribute . . . and must weigh his need to follow the rules against his desire to survive no matter what it takes.

This book, like the first and second editions, addresses the fundamental principles of interaction between radiation and matter and the principles of particle detection and detectors in a wide scope of fields, from low to high energy, including space physics and medical environment. It provides abundant information about the processes of electromagnetic and hadronic energy dep systems, performance of detectors and their optimization. The third edition includes additional material covering, for instance: mechanisms of energy loss like the inverse Compton scattering, corrections due to the Landau–Pomeranchuk–Migdal effect, an extended relativistic treatment of nucleus–nucleus screened Coulomb scattering, and transport of charged particles inside th the displacement damage (NIEL) in semiconductors has been revisited to account for recent experimental data and more comprehensive comparisons with results previously obtained. This book will be of great use to graduate students and final-year undergraduates as a reference and supplement for courses in particle, astroparticle, space physics and instrumentation. A part o courses in medical physics. The book can also be used by researchers in experimental particle physics at low, medium, and high energy who are dealing with instrumentation. Errata(s) Errata Contents:Electromagnetic Interaction of Radiation in MatterNuclear Interactions in MatterRadiation Environments and Damage in Silicon SemiconductorsScintillating Media and Scintillator DetectorsDisplacement Damage and Particle Interactions in Silicon DevicesGas Filled ChambersPrinciples of Particle Energy DeterminationSuperheated Droplet (Bubble) Detectors and CDM SearchMedical Physics Applications Readership: Researchers, academics, graduate students and professionals in accelerator, particle, astroparticle, space, applied and medical physics. Between Radiation/Particles and Matter;High;Intermediate and Low Energy Particle Physics;Medical Physics;Radiation/Particle Detection;Space Physics;Detectors;Semiconductors;Calorimeters;Chambers;Scintillators;Silicon Pixels;Radiation Damage;Single Event Effects;Solar CellsKey Features:Covers state-of-the-art detection techniques and underlying theoriesAddresses to for professionals in medical physics, nuclear engineering, and environmental studiesContains an updated reference table set of physical properties

The richly illustrated Interactive Web-Based Data Visualization with R, plotly, and shiny focuses on the process of programming interactive web graphics for multidimensional data analysis. It is written for the data analyst who wants to leverage the capabilities of interactive web graphics without having to learn web programming. Through many R code examples, you will learn ho functionality of these tools to enhance the presentation and exploration of data. By mastering these concepts and tools, you will impress your colleagues with your ability to quickly generate more informative, engaging, and reproducible interactive graphics using free and open source software that you can share over email, export to pdf, and more. Key Features: Convert static g interactive web-based form Link, animate, and arrange multiple plots in standalone HTML from R Embed, modify, and respond to plotly graphics in a shiny app Learn best practices for visualizing continuous, discrete, and multivariate data Learn numerous ways to visualize geo-spatial data This book makes heavy use of plotly for graphical rendering, but you will also learn about support different phases of a data science workflow, such as tidyr, dplyr, and tidyverse. Along the way, you will gain insight into best practices for visualization of high-dimensional data, statistical graphics, and graphical perception. The printed book is complemented by an interactive website where readers can view movies demonstrating the examples and interact with graphics.

This book closes the gap for beginners who want to study the Amharic language and had difficulties in finding the right grammar for this purpose: The first grammar of Amharic, the national language of Ethiopia, was published by Hiob Ludolf in 1698. The Amharic grammar published by Praetorius in 1879 is based on Amharic religious texts and on scattered material, usually com milestone in the study of Amharic is Marcel Cohen's Traite de langue amharique (1936), but this grammar, too is not completely suited for beginners since the author's generalizations are at times aimed at linguists. The grammar that comes closest to the concept of a beginner's grammar is that of C.H. Dawkin (1960), yet this grammar is extremely short, does not give examples student to the intricacies of the language.The new book gives all the grammatical forms and the sentences of the present grammar in Amharic script and in phonetic transcription. The illustrative examples have a free and a literal translation. This procedure should likewise prove to be useful for the Semitist as well as for the general linguist.

The FEniCS Tutorial I
Interactive Web-Based Data Visualization with R, plotly, and shiny
Social Perception and Social Reality
Modeling Failure in Materials
Tensor Network Contractions
Neural Fields
Inferring Patterns and Dynamics of Species Occurrence

This book includes an extended version of selected papers presented at the 11th Industry Symposium 2021 held during January 7–10, 2021. The book covers contributions ranging from theoretical and foundation research, platforms, methods, applications, and tools in all areas. It provides theory and practices in the area of data science, which add a social, geographical, and temporal dimension to data science research. It also includes application-oriented papers that prepare and use data in discovery research. This book contains chapters from academia as well as practitioners on big data technologies, artificial intelligence, machine learning, deep learning, data representation and visualization, business analytics, healthcare analytics, bioinformatics, etc. This book is helpful for the students, practitioners, researchers as well as industry professional.

This is an essential read for anyone interested in the way language is used in the world of politics. Based on Aristotle's premise that we are all political animals, able to use language to pursue our own ends, the book uses the theoretical framework of linguistics to explore the ways in which we think and behave politically. Contemporary and high profile case studies of politicians and other speakers are used, including an examination of the dangerous influence of a politician's words on the defendants in the Stephen Lawrence murder trial. International in its perspective, Analysing Political Discourse also considers the changing landscape of political language post-September 11, including the increasing use of religious imagery in the political discourse of, amongst others, George Bush. Written in a lively and engaging style, this book provides an essential introduction to political discourse analysis.

Coding and testing are often considered separate areas of expertise. In this comprehensive guide, author and Java expert Scott Oaks takes the approach that anyone who works with Java should be equally adept at understanding how code behaves in the JVM, as well as the tunings likely to help its performance. You'll gain in-depth knowledge of Java application performance, using the Java Virtual Machine (JVM) and the Java platform, including the language and API. Developers and performance engineers alike will learn a variety of features, tools, and processes for improving the way Java 7 and 8 applications perform. Apply four principles for obtaining the best results from performance testing Use JDK tools to collect data on how a Java application is performing Understand the advantages and disadvantages of using a JIT compiler Tune JVM garbage collectors to affect programs as little as possible Use techniques to manage heap memory and JVM native memory Maximize Java threading and synchronization performance features Tackle performance issues in Java EE and Java SE APIs Improve Java-driven database application performance

This dissertation considers two combinatorial problems arising in large-scale, sparse optimizaton. The first is the problem of approximating the Hessian matrix of a smooth, non-linear function by finite differencing, where the object is to minimize the required number of gradient evaluations. The second is to find as sparse a representation as possible of a given set of linear constraints.

The Google Story
PostgreSQL for Data Architects
Architecture of Computing Systems - ARCS 2009
High Performance Computing Systems
The Carnation Genome
Consumption, Harm and Policy Approaches
Trends of Data Science and Applications

Essential Mathematics for Games and Interactive Applications, 2nd edition presents the core mathematics necessary for sophisticated 3D graphics and interactive physical simulations. The book begins with linear algebra and matrix multiplication and expands on this foundation to cover such topics as color and lighting, interpolation, animation and basic game physics. Essential Mathematics focuses on the issues of 3D game development important to programmers and includes optimization guidance throughout. The new edition Windows code will now use Visual Studio.NET. There will also be DirectX support provided, along with OpenGL - due to its cross-platform nature. Programmers will find more concrete examples included in this edition, as well as additional information on tuning, optimization and robustness. The book has a companion CD-ROM with exercises and a test bank for the academic secondary market, and for main market: code examples built around a shared code base, including a math library covering all the topics presented in the book, a core vector/matrix math engine, and libraries to support basic 3D rendering and interaction.

Real-Time Collision DetectionCRC Press

As part of Packt's cookbook series, each recipe offers a practical, step-by-step solution to common problems found in HBase administration. This book is for HBase administrators, developers, and will even help Hadoop administrators. You are not required to have HBase experience, but are expected to have a basic understanding of Hadoop and MapReduce.

The term “stringology” is a popular nickname for text algorithms, or algorithms on strings. This book deals with the most basic algorithms in the area. Most of them can be viewed as “algorithmic jewels” and deserve reader-friendly presentation. One of the main aims of the book is to present several of the most celebrated algorithms in a simple way by omitting obscuring details and separating algorithmic structure from combinatorial theoretical background. The book reflects the relationships between applications of text-algorithmic techniques and the classification of algorithms according to the measures of complexity considered. The text can be viewed as a parade of algorithms in which the main purpose is to discuss the foundations of the algorithms and their interconnections. One can partition the algorithmic problems discussed into practical and theoretical problems. Certainly, string matching and data compression are in the former class, while most problems related to symmetries and repetitions in texts are in the latter. However, all the problems are interesting from an algorithmic point of view and enable the reader to appreciate the importance of combinatorics on words as a tool in the design of efficient text algorithms.In most textbooks on algorithms and data structures, the presentation of efficient algorithms on words is quite short as compared to issues in graph theory, sorting, searching, and some other areas. At the same time, there are many presentations of interesting algorithms on words accessible only in journals and in a form directed mainly at specialists. This book fills the gap in the book literature on algorithms on words, and brings together the many results presently dispersed in the masses of journal articles. The presentation is reader-friendly; many examples and about two hundred figures illustrate nicely the behaviour of otherwise very complex algorithms.

Alcohol in the European Union
Samphire Marshes of the Peel-Harvey Estuarine System
Jewels Of Stringology: Text Algorithms
Methods and Applications to Quantum Many-Body Systems
Principles of Asynchronous Circuit Design
Western Australia
Capture-recapture and Removal Methods for Sampling Closed Populations

Social Perception and Social Reality contests the received wisdom in the field of social psychology that suggests that social perception and judgment are generally flawed, biased, and powerfully self-fulfilling. Jussim reviews a wealth of real world, survey, and experimental data collected over the last century to show that in fact, social psychological research consistently demonstrates that biases and self-fulfilling prophecies are generally weak, fragile, and fleeting. Furthermore, research in the social sciences has shown stereotypes to be accurate. Jussim overturns the received wisdom concerning social perception in several ways. He critically reviews studies that are highly cited darlings of the bias conclusion and shows how these studies demonstrate far more accuracy than bias, or are not replicable in subsequent research. Studies of equal or higher quality, which have been replicated consistently, are shown to demonstrate high accuracy, low bias, or both. The book is peppered with discussions suggesting that theoretical and political blinders have led to an odd state of affairs in which the flawed or misinterpreted bias studies receive a great deal of attention, while stronger and more replicable accuracy studies receive relatively little attention. In addition, the author presents both personal and real world examples (such as stock market prices, sporting events, and political elections) that routinely undermine heavy-handed emphases on error and bias, but are generally indicative of high levels of rationality and accuracy. He fully embraces scientific data, even when that data yields unpopular conclusions or contests prevailing conventions or the received wisdom in psychology, in other social sciences, and in broader society.

This book constitutes the thoroughly refereed post-conference proceedings of the 6th International Conference on Mobile, Secure and Programmable Networking, held in Paris, France, in October 2020. The 16 full papers presented in this volume were carefully reviewed and selected from 31 submissions. They discuss new trends in networking infrastructures, security, services and applications while focusing on virtualization and cloud computing for networks, network programming, software defined networks (SDN) and their security.

This book constitutes the refereed proceedings of the 19th Symposium on High Performance Computing System, WSCAD 2018, held in São Paulo, Brazil, in October 2018. The 12 revised full papers presented were carefully reviewed and selected out of 61 submissions. The papers included in this book are organized according to the following topics: cloud computing; performance; processors and memory architectures; power and energy.

Written by an expert in the game industry, Christer Ericson's new book is a comprehensive guide to the components of efficient real-time collision detection systems. The book provides the tools and know-how needed to implement industrial-strength collision detection for the highly detailed dynamic environments of applications such as 3D games, virt

Quarkus 1.x
Analysing Political Discourse
FreeBSD Handbook
Getting the Most Out of Your Code
Java Performance: The Definitive Guide
Real-Time Collision Detection
The Ballad of Songbirds and Snakes (A Hunger Games Novel)

The FreeBSD Handbook is a comprehensive FreeBSD tutorial and reference. It covers installation, day-to-day use of FreeBSD, and mach more, such as the Ports collection, creating a custom kernel, security topics, the X Window System, how to use FreeBSD's Linux binary compatibility, and how to upgrade your system from source using the 'make world' command, to name a few.

Tensor network is a fundamental mathematical tool with a huge range of applications in physics, such as condensed matter physics, statistic physics, high energy physics, and quantum information sciences. This open access book aims to explain the tensor network contraction approaches in a systematic way, from the basic definitions to the important applications. This book is also useful to those who apply tensor networks in areas beyond physics, such as machine learning and the big-data analysis. Tensor network originates from the numerical renormalization group approach proposed by K.G. Wilson in 1975. Through a rapid development in the last two decades, tensor network has become a powerful numerical tool that can efficiently simulate a wide range of scientific problems, with particular success in quantum many-body physics. Varieties of tensor network algorithms have been proposed for different problems. However, the connections among different algorithms are not well discussed or reviewed. To fill this gap, this book explains the fundamental concepts and basic ideas that connect and/or unify different strategies of the tensor network contraction algorithms. In addition, some of the recent progresses in dealing with tensor decomposition techniques and quantum simulations are also represented in this book to help the readers to better understand tensor network. This open access book is intended for graduated students, but can also be used as a professional book for researchers in the related fields. To understand most of the contents in the book, only basic knowledge of quantum mechanics and linear algebra is required. In order to fully understand some advanced parts, the reader will need to be familiar with notion of condensed matter physics and quantum information, that however are not necessary to understand the main parts of the book. This book is a good source for non-specialists on quantum physics to understand tensor network algorithms and the related mathematics.

Corporate workgroups, distributed enterprises, and small to medium-sized companies are increasingly seeking to network and consolidate storage to improve availability, share information, reduce costs, and protect and secure information. These organizations require enterprise-class solutions capable of addressing immediate storage needs cost-effectively, while providing an upgrade path for future requirements. IBM® System Storage® N series storage systems and their software capabilities are designed to meet these requirements. IBM System Storage N series storage systems offer an excellent solution for a broad range of deployment scenarios. IBM System Storage N series storage systems function as a multiprotocol storage device that is designed to allow you to simultaneously serve both file and block-level data across a single network. These activities are demanding procedures that, for some solutions, require multiple, separately managed systems. The flexibility of IBM System Storage N series storage systems, however, allows them to address the storage needs of a wide range of organizations, including distributed enterprises and data centers for midrange enterprises. IBM System Storage N series storage systems also support sites with computer and data-intensive enterprise applications, such as database, data warehousing, workgroup collaboration, and messaging. This IBM Redbooks® publication explains the software features of the IBM System Storage N series storage systems. This book also covers topics such as installation, setup, and administration of those software features from the IBM System Storage N series storage systems and clients and provides example scenarios.

Anomalous transport is a ubiquitous phenomenon in astrophysical, geophysical and laboratory plasmas; and is a key topic in controlled nuclear fusion research. Despite its fundamental importance and ongoing research interest, a full understanding of anomalous transport in plasmas is still incomplete, due to the complexity of the nonlinear phenomena involved. Aspects in Anomalous Transport in Plasmas is the first book to systematically consider anomalous plasma transport theory and provides a unification of the many theoretical models by emphasizing interrelations between seemingly different methodologies. It is not intended as a catalogue of the vast number of plasma instabilities leading to anomalous transport; instead it chooses a number of these and emphasizes the aspects specifically due to turbulence. After a brief introduction, the microscopic theory of turbulence is discussed, including quasilinear theory and various aspects of renormalization methods, which leads to an understanding of resonance broadening, mode coupling, trajectory correlation and clumps. The second half of the book is devoted to stochiastic tramsport, using methods based on the Langevin equations and on Random Walk theory. This treatment aims at going beyond the traditional limits of weak turbulence, by introducing the recently developed method of decorrelation trajectories, and its application to electrostatic turbulence, magnetic turbulence and zonal flow generation. The final chapter includes very recent work on the nonlocal transport phenomenon.

IBM System Storage N series Software Guide
Why Accuracy Dominates Bias and Self-Fulfilling Prophecy
Introductory Grammar of Amharic
A Systems Perspective
A Handbook
Theory and Practice
Understanding Quarkus

Neural field theory has a long-standing tradition in the mathematical and computational neurosciences. Beginning almost 50 years ago with seminal work by Griffiths and culminating in the 1970ties with the models of Wilson and Cowan, Nunez and Amari, this important research area experienced a renaissance during the 1990ties by the groups of Ermentrout, Robinson, Bressloff, Wright and Haken. Since then, much progress has been made in both, the development of mathematical and numerical techniques and in physiological refinement und understanding. In contrast to large-scale neural network models described by huge connectivity matrices that are computationally expensive in numerical simulations, neural field models described by connectivity kernels allow for analytical treatment by means of methods from functional analysis. Thus, a number of rigorous results on the existence of bump and wave solutions or on inverse kernel construction problems are nowadays available. Moreover, neural fields provide an important interface for the coupling of neural activity to experimentally observable data, such as the electroencephalogram (EEG) or functional magnetic resonance imaging (fMRI). And finally, neural fields over rather abstract feature spaces, also called dynamic fields, found successful applications in the cognitive sciences and in robotics. Up to now, research results in neural field theory have been disseminated across a number of distinct journals from mathematics, computational neuroscience, biophysics, cognitive science and others. There is no comprehensive collection of results or reviews available yet. With our proposed book Neural Field Theory, we aim at filling this gap in the market. We received consent from some of the leading scientists in the field, who are willing to write contributions for the book, among them are two of the founding-fathers of neural field theory: Shun-ichi Amari and Jack Cowan.

At head of title: World Health Organization, regional office for Europe.

This book offers a concise and gentle introduction to finite element programming in Python based on the popular FEniCS software library. Using a series of examples, including the Poisson equation, the equations of linear elasticity, the incompressible Navier–Stokes equations, and systems of nonlinear advection–diffusion–reaction equations, it guides readers through the essential steps to quickly solving a PDE in FEniCS, such as how to define a finite variational problem, how to set boundary conditions, how to solve linear and nonlinear systems, and how to visualize solutions and structure finite element Python programs. This book is open access under a CC BY license.

An inside look at the billion-dollar enterprise reveals how the Internet icon grew from a concept to a social phenomenon with a bold mission: to organize all of the world's information and make it easily accessible to people in more than one hundred languages. Reprint. 50,000 first printing.

Solving PDEs in Python
Theory and Applications
Essential Mathematics for Games and Interactive Applications
DB2 9 for z/OS Performance Topics
19th Symposium, WSCAD 2018, São Paulo, Brazil, October 1–3, 2018, Revised Selected Papers
22nd International Conference, Delft, The Netherlands, March 10-13, 2009, Proceedings
Theory and Practices

The heart of any system that simulates the physical interaction between objects is collision detection-the ability to detect when two objects have come into contact. This system is also one of the most difficult aspects of a physical simulation to implement correctly, and invariably it is the main consumer of CPU cycles. Practitioners, new to the field or otherwise, quickly discover that the attempt to build a fast, accurate, and robust collision detection system takes them down a long path fraught with perils and pitfalls unlike most they have ever encountered. Without in-depth knowledge and understanding of the issues associated with engineering a collision detection system, the end of that path is an abyss that has swallowed many a good programmer! Gino van den Bergen's new book is the story of his successful journey down that path. The outcome is his well-known collision detection system, the SOftware Library for Interference Detection (SOLID). Along the way, he covers the topics of vector algebra and geometry, the various geometric primitives of interest in a collision system, the powerful method of separating axes for the purposes of intersection testing, and the equally powerful Gilbert-Johnson-Keerthi (GJK) algorithm for computing the distance between convex objects. But this book provides much more than a good compendium of the ideas that go into building a collision system. The curse of practical computational geometry is floating-point arithmetic. Algorithms with straightforward implementations when using exact arithmetic can have catastrophic failures in a floating-point system. Specifically, intersection and distance algorithms implemented in a floating-point system tend to fail exactly in the most important case in a collision system-when two objects are just touching. Great care must be taken to properly handle floating-point round off errors. Gino's ultimate accomplishment in this book is his presentation on how to correctly implement the GJK distance algorithm in the presence of single-precision floating-point arithmetic. And what better way to illustrate this than with a case study, the final chapter on the design and implementation of SOLID. About the CD-ROM The companion CD-ROM includes the full C++ source code of SOLID 3.5 as well as API documentation in HTML and PDF formats. Both single (32bit) and double (64bit) precision versions of the SOLID SDK plus example programs can be compiled for Linux platforms using GNU g++ version 2.95 to 3.3 and for Win32 platforms using Microsoft Visual C++ version 6.0 to 7.1. Use of the SOLID source code is governed by the terms of either the GNU GPL or the Trolltech QPL (see CD-ROM documentation for details). About the Author Gino van den Bergen is a game developer living and working in The Netherlands. He is the creator of SOLID and holds a Ph.D. in computing science from Eindhoven University of Technology. Gino implemented collision detection and physics in NaN Technologies' Blender, a creation suite for interactive 3D content.

This book constitutes the refereed proceedings of the 22nd International Conference on Architecture of Computing Systems, ARCS 2009, held in Delft, The Netherlands, in March 2009. The 21 revised full papers presented together with 3 keynote papers were carefully reviewed and selected from 57 submissions. This year's special focus is set on energy awareness. The papers are organized in topical sections on compilation technologies, reconfigurable hardware and applications, massive parallel architectures, organic computing, memory architectures, enery awareness, Java processing, and chip-level multiprocessing.

Quantum measurement (Le., a measurement which is sufficiently precise for quantum effects to be essential) was always one of the most impor tant points in quantum mechanics because it most evidently revealed the difference between quantum and classical physics. Now quantum measure ment is again under active investigation, first of all because of the practical necessity of dealing with highly precise and complicated measurements. The nature of quantum measurement has become understood much bet ter during this new period of activity, the understanding being expressed by the concept of decoherence. This term means a physical process lead ing from a pure quantum state (wave function) of the system prior to the measurement to its state after the measurement which includes classical elements. More concretely, decoherence occurs as a result of the entangle ment of the measured system with its environment and results in the loss of phase relations between components of the wave function of the measured system. Decoherence is essentially nothing else than quantum measurement, but considered from the point of view of its physical mechanism and resolved in time. The present book is devoted to the two concepts of quantum measure ment and decoherence and to their interrelation, especially in the context of continuous quantum measurement.

Principles of Asynchronous Circuit Design - A Systems Perspective addresses the need for an introductory text on asynchronous circuit design. Part I is an 8-chapter tutorial which addresses the most important issues for the beginner, including how to think about asynchronous systems. Part II is a 4-chapter introduction to Balsa, a freely-available synthesis system for asynchronous circuits which will enable the reader to get hands-on experience of designing high-level asynchronous systems. Part III offers a number of examples of state-of-the-art asynchronous systems to illustrate what can be built using asynchronous techniques. The examples range from a complete commercial smart card chip to complex microprocessors. The objective in writing this book has been to enable industrial designers with a background in conventional (clocked) design to be able to understand asynchronous design sufficiently to assess what it has to offer and whether it might be advantageous in their next design task.

Build, Release and Distribute your Python App with Docker
Statistical Inference from Band Recovery Data
Mobile, Secure, and Programmable Networking
Hbase Administration Cookbook
Aspects of Anomalous Transport in Plasmas
Collision Detection in Interactive 3D Environments
Principles of Radiation Interaction in Matter and Detection

Microservices is an architectural style that structures an application as a collection of distributed services. Microservices are certainly appealing but there are many questions that should be asked prior to diving into this architectural style: How do I deal with an unreliable network in a distributed architecture? How do I test my services? How do I monitor them? How do I package and execute them?That's when Quarkus comes into play.In this fascicle, you will learn Quarkus but also its ecosystem. You will discover Quarkus internals and how you can use it to build REST and reactive microservices, bind and process JSON or access datastores in a transactional way. With Cloud Native and GraalVM in mind, Quarkus makes packaging and orchestrating your microservices with Docker and Kubernetes easy.This fascicle has a good mix of theory and practical examples. It is the companion book of Practising Quarkus 1.x where you learn how to develop an entire microservice architecture.

Summary Solr in Action is a comprehensive guide to implementing scalable search using Apache Solr. This clearly written book walks you through well-documented examples ranging from basic keyword searching to scaling a system for billions of documents and queries. It will give you a deep understanding of how to implement core Solr capabilities. About the Book Whether you're handling big (or small) data, managing documents, or building a website, it is important to be able to quickly search through your content and discover meaning in it. Apache Solr is your tool: a ready-to-deploy, Lucene-based, open source, full-text search engine. Solr can scale across many servers to enable real-time queries and data analytics across billions of documents. Solr in Action teaches you to implement scalable search using Apache Solr. This easy-to-read guide balances conceptual discussions with practical examples to show you how to implement all of Solr's core capabilities. You'll master topics like text analysis, faceted search, hit highlighting, result grouping, query suggestions, multilingual search, advanced geospatial and data operations, and relevancy tuning. This book assumes basic knowledge of Java and standard database technology. No prior knowledge of Solr or Lucene is required. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. What's Inside How to scale Solr for big data Rich real-world examples Solr as a NoSQL data store Advanced multilingual, data, and relevancy tricks Coverage of versions through Solr 4.7 About the Authors Trey Grainger is a director of engineering at CareerBuilder. Timothy Potter is a senior member of the engineering team at LucidWorks. The authors work on the scalability and reliability of Solr, as well as on recommendation engine and big data analytics technologies. Table of Contents PART 1 MEET SOLR Introduction to Solr Getting to know Solr Key Solr concepts Configuring Solr Indexing Text analysis PART 2 CORE SOLR CAPABILITIES Performing queries and handling results Faceted search Hit highlighting Query suggestions Result grouping/field collapsing Taking Solr to production PART 3 TAKING SOLR TO THE NEXT LEVEL SolrCloud Multilingual search Complex query operations Mastering relevancy

The World Wide Web constitutes the largest existing source of texts written in a great variety of languages. A feasible and sound way of exploiting this data for linguistic research is to compile a static corpus for a given language. There are several adavantages of this approach: (i) Working with such corpora obviates the problems encountered when using Internet search engines in quantitative linguistic research (such as non-transparent ranking algorithms). (ii) Creating a corpus from web data is virtually free. (iii) The size of corpora compiled from the WWW may exceed by several orders of magnitudes the size of language resources offered elsewhere. (iv) The data is locally available to the user, and it can be linguistically post-processed and queried with the tools preferred by her/him. This book addresses the main practical tasks in the creation of web corpora up to giga-token size. Among these tasks are the sampling process (i.e., web crawling) and the usual cleanups including boilerplate removal and removal of duplicated content. Linguistic processing and problems with linguistic processing coming from the different kinds of noise in web corpora are also covered. Finally, the authors show how web corpora can be evaluated and compared to other corpora (such as traditionally compiled corpora). For additional material please visit the companion website: sites.morganclaypool.com/wcc Table of Contents: Preface / Acknowledgments / Web Corpora / Data Collection / Post-Processing / Linguistic Processing / Corpus Evaluation and Comparison / Bibliography / Authors' Biographies

DB2 9 for z/OS is an exciting new version, with many improvements in performance and little regression. DB2 V9 improves availability and security, as well as adds greatly to SQL and XML functions. Optimization improvements include more SQL functions to optimize, improved statistics for the optimizer, better optimization techniques, and a new approach to providing information for tuning. V8 SQL procedures were not eligible to run on the IBM System z9 Integrated Information Processor (zIIP), but changing to use the native SQL procedures on DB2 V9 makes the work eligible for zIIP processing. The performance of varying length data can improve substantially if there are large numbers of varying length columns. Several improvements in disk access can reduce the time for sequential disk access and improve data rates. The key DB2 9 for z/OS performance improvements include reduced CPU time in many utilities, deep synergy with IBM System z hardware and z/OS software, improved performance and scalability for inserts and LOBs, improved SQL optimization, zIIP processing for remote native SQL procedures, index compression, reduced CPU time for data with varying lengths, and better sequential access. Virtual storage use below the 2 GB bar is also improved. This IBM Redbooks publication provides an overview of the performance impact of DB2 9 for z/OS, especially performance scalability for transactions, CPU, and elapsed time for queries and utilities. We discuss the overall performance and possible impacts when moving from version to version. We include performance measurements that were made in the laboratory and provide some estimates. Keep in mind that your results are likely to vary, as the conditions and work will differ. In this book, we assume that you are familiar with DB2 V9. See DB2 9 for z/OS Technical Overview, SG24-7330, for an introduction to the new functions.

Quantum Measurements and Decoherence
Ghent, Belgium, 26-30 August 2019
Practical Docker with Python
6th International Conference, MSPN 2020, Paris, France, October 28–29, 2020, Revised Selected Papers
Occupancy Estimation and Modeling
A Programmer's Guide, Second Edition
Web Corpus Construction

Occupancy Estimation and Modeling: Inferring Patterns and Dynamics of Species Occurrence, Second Edition, provides a synthesis of model-based approaches for analyzing presence-absence data, allowing for imperfect detection. Beginning from the relatively simple case of estimating the proportion of area or sampling units occupied at the time of surveying, the authors describe a wide variety of extensions that have been developed since the early 2000s. This provides an improved insight about species and community ecology, including, detection heterogeneity; correlated detections; spatial autocorrelation; multiple states or classes of occupancy; changes in occupancy over time; species co-occurrence; community-level modeling, and more. Occupancy Estimation and Modeling: Inferring Patterns and Dynamics of Species Occurrence, Second Edition has been greatly expanded and detail is provided regarding the estimation methods and examples of their application are given. Important study design recommendations are also covered to give a well rounded view of modeling. Provides authoritative insights into the latest in occupancy modeling Examines the latest methods in analyzing detection/no detection data surveys Addresses critical issues of imperfect detectability and its effects on species occurrence estimation Discusses important study design considerations such as defining sample units, sample size determination and optimal effort allocation

This book summarizes recent advances in carnation genome research for large-scale transcriptome analysis, the draft genome sequence, DNA markers and genome mapping, flower color, mutations, flower opening, vase life, interspecific hybridization, fragrance. The carnation is one of the most important ornamental flowers in the world, along with the chrysanthemum and the rose. The genus Dianthus is a member of the Caryophyllaceae and includes more than 300 species of annuals and evergreen perennials. Modern carnation cultivars are the product of highly complex hybridization, owing to their long history of breeding. The carnation genome was first sequenced in ornamentals by a Japanese research team in 2013. The carnation has been genetically improved over the years, and there are various types of flower colors, shapes, patterns, and sizes. In this book, the molecular mechanism of flower color development and the transposable elements responsible for this diversity are studied in detail. In addition, it presents breeding and physiological research for improving flower vase life, one of the most important traits in ornamentals, based on a model of ethylene susceptible flowers. To improve selection efficiency, genomic analysis tools including DNA markers and genetic linkage maps are also highlighted. In closing, the book discusses mutation breeding technologies such as ion-beam irradiation and genetically modified carnations.

This book is for developers and data architects who have some exposure to databases. It is assumed that you understand the basic concepts of tables and common database objects, including privileges and security.

This Book of Abstracts is the main publication of the 70th Annual Meeting of the European Federation of Animal Science (EAAP). It contains abstracts of the invited papers and contributed presentations of the sessions of EAAP's eleven Commissions: Animal Genetics, Animal Nutrition, Animal Management and Health, Animal Physiology, Cattle Production, Sheep and Goat Production, Pig Production, Horse Production and Livestock Farming Systems, Insects and Precision Livestock Farming.

Models and Phenomenology
Book of Abstracts of the 70th Annual Meeting of the European Federation of Animal Science
Solr in Action
A Combinatorial Approach to Some Sparse Matrix Problems

The Fiber Bundle Model

Learn the key differences between containers and virtual machines. Adopting a project based approach, this book introduces you to a simple Python application to be developed and containerized with Docker. After an introduction to Containers and Docker you'll be guided through Docker installation and configuration. You'll also learn basic functions and commands used in Docker by running a simple container using Docker commands. The book then moves on to developing a Python based Messaging Bot using required libraries and virtual environment where you'll add Docker Volumes to your project, ensuring your container data is safe. You'll create a database container and link your project to it and finally, bring up the Bot-associated database all at once with Docker Compose. What You'll Learn Build, run, and distribute Docker containers Develop a Python App and containerize it Use Dockerfile to run the Python App Define and run multi-container applications with Docker Compose Work with persisting data generated by and used by Docker containers Who This Book Is For Intermediate developers/DevOps practitioners who are looking to improve their build and release workflow by containerizing applications

Gathering research from physics, mechanical engineering, and statistics in a single resource for the first time, this text presents the background to the model, its theoretical basis, and applications ranging from materials science to earth science. The authors start by explaining why disorder is important for fracture and then go on to introduce the fiber bundle model, backed by various different applications. Appendices present the necessary mathematical, computational and statistical background required. The structure of the book allows the reader to skip some material that is too specialized, making this topic accessible to the engineering, mechanics and materials science communities, in addition to providing further reading for graduate students in statistical physics.