

## Getting Started With Memcached Soliman Ahmed

2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS) Things I Overheard While Talking to Myself Smart Computing Paradigms: New Progresses and Challenges Cardbox The Art of Agent-oriented Modeling Handbook of IoT and Big Data A Quarter Century of UNIX Industrial Applications Of Electron Microscopy Four Meetings Qt5 Python GUI Programming Cookbook Advanced Intelligent Systems for Sustainable Development (AI2SD'2019) Internetware Security, Privacy, and Anonymity in Computation, Communication, and Storage Developing Multi-Agent Systems with JADE Principles and Practices of Interconnection Networks Handbook of Big Data Technologies Digital Color Image Processing Getting Started with Memcached Big Data Data-intensive Workflow Management Algorithms and Architectures for Parallel Processing Data Management in the Cloud New Trends in Image Analysis and Processing -- ICIAP 2015 Workshops Developing Web Applications with Apache, MySQL, memcached, and Perl Applications in Ubiquitous Computing Large Scale and Big Data Big Data and Hadoop The Art of Application Performance Testing Linked Data Management Big Data Analytics and Knowledge Discovery The Myth of Homeland Security Data Mining Big Data 2.0 Processing Systems Ontology-Based Multi-Agent Systems Linked Data Algorithms and Architectures for Parallel Processing Wireless Rechargeable Sensor Networks Information and Decision Sciences Tcl/Tk Getting Started with Memcached

### 2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)

An introduction to color in three-dimensional image processing and the emerging area of multi-spectral image processing. The importance of color information in digital image processing is greater than ever. However, the transition from scalar to vector-valued image functions has not yet been generally covered in most textbooks. Now, Digital Color Image Processing fills this pressing need with a detailed introduction to this important topic. In four comprehensive sections, this book covers: The fundamentals and requirements for color image processing from a vector-valued viewpoint Techniques for preprocessing color images Three-dimensional scene analysis using color information, as well as the emerging area of multi-spectral imaging Applications of color image processing, presented via the examination of two case studies In addition to introducing readers to important new technologies in the field, Digital Color Image Processing also contains novel topics such as: techniques for improving three-dimensional reconstruction, three-dimensional computer vision, and emerging areas of safety and security applications in luggage inspection and video surveillance of high-security facilities. Complete with full-color illustrations and two applications chapters, Digital Color Image Processing is the only book that covers the breadth of the subject under one convenient cover. It is written at a level that is accessible for first- and second-year graduate students in electrical and computer engineering and computer science courses, and that is also appropriate for researchers who wish to extend their knowledge in the area of color image processing.

## **Things I Overheard While Talking to Myself**

This SpringerBrief provides a concise guide to applying wireless energy transfer techniques in traditional battery-powered sensor networks. It examines the benefits and challenges of wireless power including efficiency and reliability. The authors build a wireless rechargeable sensor networks from scratch and aim to provide perpetual network operation. Chapters cover a wide range of topics from the collection of energy information and recharge scheduling to joint design with typical sensing applications such as data gathering. Problems are approached using a natural combination of probability theory, optimization, algorithm and protocol designs. All proposed mechanisms are evaluated by extensive simulations. Wireless Rechargeable Sensor Networks targets professionals and researchers working in networks, wireless communications, energy technology and information technology. Advanced-level students studying electrical engineering and computer science will also find this material useful as a study guide.

## **Smart Computing Paradigms: New Progresses and Challenges**

Linked Data Management presents techniques for querying and managing Linked Data that is available on today's Web. The book shows how the abundance of Linked Data can serve as fertile ground for research and commercial applications. The text focuses on aspects of managing large-scale collections of Linked Data. It offers a detailed introduction to Linked Data and related standards, including the main principles distinguishing Linked Data from standard database technology. Chapters also describe how to generate links between datasets and explain the overall architecture of data integration systems based on Linked Data. A large part of the text is devoted to query processing in different setups. After presenting methods to publish relational data as Linked Data and efficient centralized processing, the book explores lookup-based, distributed, and parallel solutions. It then addresses advanced topics, such as reasoning, and discusses work related to read-write Linked Data for system interoperation. Despite the publication of many papers since Tim Berners-Lee developed the Linked Data principles in 2006, the field lacks a comprehensive, unified overview of the state of the art. Suitable for both researchers and practitioners, this book provides a thorough, consolidated account of the new data publishing and data integration paradigm. While the book covers query processing extensively, the Linked Data abstraction furnishes more than a mechanism for collecting, integrating, and querying data from the open Web—the Linked Data technology stack also allows for controlled, sophisticated applications deployed in an enterprise environment.

## **Cardbox**

Cloud computing has emerged as a successful paradigm of service-oriented computing and has revolutionized the way computing infrastructure is used. This success has seen a proliferation in the number of applications that are being

deployed in various cloud platforms. There has also been an increase in the scale of the data generated as well as consumed by such applications. Scalable database management systems form a critical part of the cloud infrastructure. The attempt to address the challenges posed by the management of big data has led to a plethora of systems. This book aims to clarify some of the important concepts in the design space of scalable data management in cloud computing infrastructures. Some of the questions that this book aims to answer are: the appropriate systems for a specific set of application requirements, the research challenges in data management for the cloud, and what is novel in the cloud for database researchers? We also aim to address one basic question: whether cloud computing poses new challenges in scalable data management or it is just a reincarnation of old problems? We provide a comprehensive background study of state-of-the-art systems for scalable data management and analysis. We also identify important aspects in the design of different systems and the applicability and scope of these systems. A thorough understanding of current solutions and a precise characterization of the design space are essential for clearing the "cloudy skies of data management" and ensuring the success of DBMSs in the cloud, thus emulating the success enjoyed by relational databases in traditional enterprise settings.

### **The Art of Agent-oriented Modeling**

This book presents a comprehensive introduction to Internetware, covering aspects ranging from the fundamental principles and engineering methodologies to operational platforms, quality measurements and assurance and future directions. It also includes guidelines and numerous representative real-world case studies that serve as an invaluable reference resource for software engineers involved in the development of Internetware applications. Providing a detailed analysis of current trends in modern software engineering in the Internet, it offers an essential blueprint and an important contribution to the research on software engineering and systems for future Internet computing.

### **Handbook of IoT and Big Data**

This book describes efficient and effective techniques for harnessing the power of Linked Data by tackling the various aspects of managing its growing volume: storing, querying, reasoning, provenance management and benchmarking. To this end, Chapter 1 introduces the main concepts of the Semantic Web and Linked Data and provides a roadmap for the book. Next, Chapter 2 briefly presents the basic concepts underpinning Linked Data technologies that are discussed in the book. Chapter 3 then offers an overview of various techniques and systems for centrally querying RDF datasets, and Chapter 4 outlines various techniques and systems for efficiently querying large RDF datasets in distributed environments. Subsequently, Chapter 5 explores how streaming requirements are addressed in current, state-of-the-art RDF stream data processing. Chapter 6 covers performance and scaling issues of distributed RDF reasoning systems, while Chapter 7 details

benchmarks for RDF query engines and instance matching systems. Chapter 8 addresses the provenance management for Linked Data and presents the different provenance models developed. Lastly, Chapter 9 offers a brief summary, highlighting and providing insights into some of the open challenges and research directions. Providing an updated overview of methods, technologies and systems related to Linked Data this book is mainly intended for students and researchers who are interested in the Linked Data domain. It enables students to gain an understanding of the foundations and underpinning technologies and standards for Linked Data, while researchers benefit from the in-depth coverage of the emerging and ongoing advances in Linked Data storing, querying, reasoning, and provenance management systems. Further, it serves as a starting point to tackle the next research challenges in the domain of Linked Data management.

## **A Quarter Century of UNIX**

Workflows may be defined as abstractions used to model the coherent flow of activities in the context of an in silico scientific experiment. They are employed in many domains of science such as bioinformatics, astronomy, and engineering. Such workflows usually present a considerable number of activities and activations (i.e., tasks associated with activities) and may need a long time for execution. Due to the continuous need to store and process data efficiently (making them data-intensive workflows), high-performance computing environments allied to parallelization techniques are used to run these workflows. At the beginning of the 2010s, cloud technologies emerged as a promising environment to run scientific workflows. By using clouds, scientists have expanded beyond single parallel computers to hundreds or even thousands of virtual machines. More recently, Data-Intensive Scalable Computing (DISC) frameworks (e.g., Apache Spark and Hadoop) and environments emerged and are being used to execute data-intensive workflows. DISC environments are composed of processors and disks in large-commodity computing clusters connected using high-speed communications switches and networks. The main advantage of DISC frameworks is that they support and grant efficient in-memory data management for large-scale applications, such as data-intensive workflows. However, the execution of workflows in cloud and DISC environments raise many challenges such as scheduling workflow activities and activations, managing produced data, collecting provenance data, etc. Several existing approaches deal with the challenges mentioned earlier. This way, there is a real need for understanding how to manage these workflows and various big data platforms that have been developed and introduced. As such, this book can help researchers understand how linking workflow management with Data-Intensive Scalable Computing can help in understanding and analyzing scientific big data. In this book, we aim to identify and distill the body of work on workflow management in clouds and DISC environments. We start by discussing the basic principles of data-intensive scientific workflows. Next, we present two workflows that are executed in a single site and multi-site clouds taking advantage of provenance. Afterward, we go towards workflow management in DISC environments, and we present, in detail, solutions that enable the optimized execution of the workflow using frameworks such as Apache Spark and its extensions.

## **Industrial Applications Of Electron Microscopy**

This textbook explores the different aspects of data mining from the fundamentals to the complex data types and their applications, capturing the wide diversity of problem domains for data mining issues. It goes beyond the traditional focus on data mining problems to introduce advanced data types such as text, time series, discrete sequences, spatial data, graph data, and social networks. Until now, no single book has addressed all these topics in a comprehensive and integrated way. The chapters of this book fall into one of three categories: Fundamental chapters: Data mining has four main problems, which correspond to clustering, classification, association pattern mining, and outlier analysis. These chapters comprehensively discuss a wide variety of methods for these problems. Domain chapters: These chapters discuss the specific methods used for different domains of data such as text data, time-series data, sequence data, graph data, and spatial data. Application chapters: These chapters study important applications such as stream mining, Web mining, ranking, recommendations, social networks, and privacy preservation. The domain chapters also have an applied flavor. Appropriate for both introductory and advanced data mining courses, Data Mining: The Textbook balances mathematical details and intuition. It contains the necessary mathematical details for professors and researchers, but it is presented in a simple and intuitive style to improve accessibility for students and industrial practitioners (including those with a limited mathematical background). Numerous illustrations, examples, and exercises are included, with an emphasis on semantically interpretable examples. Praise for Data Mining: The Textbook - "As I read through this book, I have already decided to use it in my classes. This is a book written by an outstanding researcher who has made fundamental contributions to data mining, in a way that is both accessible and up to date. The book is complete with theory and practical use cases. It's a must-have for students and professors alike!" -- Qiang Yang, Chair of Computer Science and Engineering at Hong Kong University of Science and Technology "This is the most amazing and comprehensive text book on data mining. It covers not only the fundamental problems, such as clustering, classification, outliers and frequent patterns, and different data types, including text, time series, sequences, spatial data and graphs, but also various applications, such as recommenders, Web, social network and privacy. It is a great book for graduate students and researchers as well as practitioners." -- Philip S. Yu, UIC Distinguished Professor and Wexler Chair in Information Technology at University of Illinois at Chicago

## **Four Meetings**

## **Qt5 Python GUI Programming Cookbook**

This practical book provides a step-by-step approach to testing mission-critical applications for scalability and performance

before they're deployed -- a vital topic to which other books devote one chapter, if that. Businesses today live and die by network applications and web services. Because of the increasing complexity of these programs, and the pressure to deploy them quickly, many professionals don't take the time to ensure that they'll perform well and scale effectively. The Art of Application Performance Testing explains the complete life cycle of the testing process, and demonstrates best practices to help you plan, gain approval for, coordinate, and conduct performance tests on your applications. With this book, you'll learn to: Set realistic performance testing goals Implement an effective application performance testing strategy Interpret performance test results Cope with different application technologies and architectures Use automated performance testing tools Test traditional local applications, web-based applications, and web services (SOAs) Recognize and resolves issues that are often overlooked in performance tests Written by a consultant with 30 years of experience in the IT industry and over 12 years experience with performance testing, this easy-to-read book is illustrated with real-world examples and packed with practical advice. The Art of Application Performance Testing thoroughly explains the pitfalls of an inadequate testing strategy and offers you a robust, structured approach for ensuring that your applications perform well and scale effectively when the need arises. "Ian has maintained a vendor-agnostic methodology beautifully in this material. The metrics and graphs, along with background information provided in his case studies, eloquently convey to the reader, 'Methodology above all, tools at your discretion' Ian's expertise shines through throughout the entire reading experience."-- Matt St. Onge, Enterprise Solution Architect, HCL Technologies America / Teradyne

## **Advanced Intelligent Systems for Sustainable Development (AI2SD'2019)**

This multi-contributed handbook focuses on the latest workings of IoT (internet of Things) and Big Data. As the resources are limited, it's the endeavor of the authors to support and bring the information into one resource. The book is divided into 4 sections that covers IoT and technologies, the future of Big Data, algorithms, and case studies showing IoT and Big Data in various fields such as health care, manufacturing and automation. Features Focuses on the latest workings of IoT and Big Data Discusses the emerging role of technologies and the fast-growing market of Big Data Covers the movement toward automation with hardware, software, and sensors, and trying to save on energy resources Offers the latest technology on IoT Presents the future horizons on Big Data

## **Internetware**

## **Security, Privacy, and Anonymity in Computation, Communication, and Storage**

Based on interviews with the key software engineers who invented and built the powerful UNIX operating system, this book

provides unique insight into the operating system that dominates the modern computing environment. Originating from a small project in a backroom at AT &T Bell Labs, UNIX has grown to be a dominant operating system in the commercial computing world -the operating system responsible for the development of the C programming language and the modern networked environment. Peter Salus is a longtime and well-recognized promoter and spokesman for UNIX and the UNIX community.

### **Developing Multi-Agent Systems with JADE**

One of the greatest challenges faced by designers of digital systems is optimizing the communication and interconnection between system components. Interconnection networks offer an attractive and economical solution to this communication crisis and are fast becoming pervasive in digital systems. Current trends suggest that this communication bottleneck will be even more problematic when designing future generations of machines. Consequently, the anatomy of an interconnection network router and science of interconnection network design will only grow in importance in the coming years. This book offers a detailed and comprehensive presentation of the basic principles of interconnection network design, clearly illustrating them with numerous examples, chapter exercises, and case studies. It incorporates hardware-level descriptions of concepts, allowing a designer to see all the steps of the process from abstract design to concrete implementation. Case studies throughout the book draw on extensive author experience in designing interconnection networks over a period of more than twenty years, providing real world examples of what works, and what doesn't. Tightly couples concepts with implementation costs to facilitate a deeper understanding of the tradeoffs in the design of a practical network. A set of examples and exercises in every chapter help the reader to fully understand all the implications of every design decision.

### **Principles and Practices of Interconnection Networks**

During the last two decades, the idea of Semantic Web has received a great deal of attention. An extensive body of knowledge has emerged to describe technologies that seek to help us create and use aspects of the Semantic Web. Ontology and agent-based technologies are understood to be the two important technologies here. A large number of articles and a number of books exist to describe the use individually of the two technologies and the design of systems that use each of these technologies individually, but little focus has been given on how one can - sign systems that carryout integrated use of the two different technologies. In this book we describe ontology and agent-based systems individually, and highlight advantages of integration of the two different and complementary te- nologies. We also present a methodology that will guide us in the design of the - tegrated ontology-based multi-agent systems and illustrate this methodology on two use cases from the health and software engineering domain. This book is organized as follows: • Chapter I, Current issues and the need for ontologies and agents, describes existing problems associated with

uncontrollable information overload and explains how ontologies and agent-based systems can help address these - sues. • Chapter II, Introduction to multi-agent systems, defines agents and their main characteristics and features including mobility, communications and collaboration between different agents. It also presents different types of agents on the basis of classifications done by different authors.

### **Handbook of Big Data Technologies**

Large Scale and Big Data: Processing and Management provides readers with a central source of reference on the data management techniques currently available for large-scale data processing. Presenting chapters written by leading researchers, academics, and practitioners, it addresses the fundamental challenges associated with Big Data processing tools and techniques across a range of computing environments. The book begins by discussing the basic concepts and tools of large-scale Big Data processing and cloud computing. It also provides an overview of different programming models and cloud-based deployment models. The book's second section examines the usage of advanced Big Data processing techniques in different domains, including semantic web, graph processing, and stream processing. The third section discusses advanced topics of Big Data processing such as consistency management, privacy, and security. Supplying a comprehensive summary from both the research and applied perspectives, the book covers recent research discoveries and applications, making it an ideal reference for a wide range of audiences, including researchers and academics working on databases, data mining, and web scale data processing. After reading this book, you will gain a fundamental understanding of how to use Big Data-processing tools and techniques effectively across application domains. Coverage includes cloud data management architectures, big data analytics visualization, data management, analytics for vast amounts of unstructured data, clustering, classification, link analysis of big data, scalable data mining, and machine learning techniques.

### **Digital Color Image Processing**

This book is an easy-to-follow guide full of hands-on examples of real-world administration tasks, which will help you design and build highly scalable and very fast web applications. This book is ideal for web application developers looking for hands-on, practical, and to-the-point recipes to integrate the Memcached service into their day-to-day programming tasks. No prior knowledge of Memcached is assumed for using this book.

### **Getting Started with Memcached**

On the heels of his acclaimed memoir, *Never Have Your Dog Stuffed*, beloved actor and bestselling author Alan Alda has

written *Things I Overheard While Talking to Myself*, an insightful and funny look at some of the impossible questions he's asked himself over the years: What do I value? What, exactly, is the good life? (And what does that even mean?) Picking up where his bestselling memoir left off—having been saved by emergency surgery after nearly dying on a mountaintop in Chile—Alda finds himself not only glad to be alive but searching for a way to squeeze the most juice out of his new life. Looking for a sense of meaning that would make this extra time count, he listens in on things he's heard himself saying in private and in public at critical points in his life—from the turbulence of the sixties, to his first Broadway show, to the birth of his children, to the ache of September 11, and beyond. Reflecting on the transitions in his life and in all our lives, he notices that “doorways are where the truth is told,” and wonders if there's one thing—art, activism, family, money, fame—that could lead to a “life of meaning.” In a book that is candid, wise, and as questioning as it is incisive, Alda amuses and moves us with his unique and hilarious meditations on questions great and small. *Things I Overheard While Talking to Myself* is another superb Alan Alda performance, as inspiring and entertaining as the man himself.

### **Big Data**

This book takes a deep dive into ubiquitous computing for applications in health, business, education, tourism, and transportation. The rich interdisciplinary contents of the book appeal to readers from diverse disciplines who aspire to create new and innovative research initiatives and applications in ubiquitous computing. Topics include condition monitoring and diagnostics; multi-objective optimization in design, multi-objective optimization of machining parameters, and more. The book benefits researchers, advanced students, as well as practitioners interested in applications of ubiquitous computing. Features practical, tested applications in ubiquitous computing Includes applications such as health, business, education, electronics, tourism, and transportation Applicable to researchers, academics, students, and professionals

### **Data-intensive Workflow Management**

This book constitutes the refereed proceedings of seven workshops held at the 18th International Conference on Image Analysis and Processing, ICIAP 2015, in Genoa, Italy, in September 2015: International Workshop on Recent Advances in Digital Security: Biometrics and Forensics, BioFor 2015; International Workshop on Color in Texture and Material Recognition, CTMR 2015; International Workshop on Medical Imaging in Rheumatology: Advanced applications for the analysis of inflammation and damage in the rheumatoid joint, RHEUMA 2015; International Workshop on Image-Based Smart City Application, ISCA 2015; International Workshop on Multimedia Assisted Dietary Management, MADiMa 2015; International Workshop on Scene Background Modeling and Initialization, SBMI 2015; and International Workshop on Image and Video Processing for Quality of Multimedia Experience, QoEM 2015.

## **Algorithms and Architectures for Parallel Processing**

This book gathers papers presented at the second installment of the International Conference on Advanced Intelligent Systems for Sustainable Development (AI2SD-2019), which was held on July 08-11, 2019 in Marrakech, Morocco. It offers comprehensive coverage of recent advances in big data, data analytics and related paradigms. The book consists of fifty-two chapters, each of which shares the latest research in the fields of big data and data science, and describes use cases and applications of big data technologies in various domains, such as social networks and health care. All parts of the book discuss open research problems and potential opportunities that have arisen from the rapid advances in big data technologies. In addition, the book surveys the state of the art in data science, and provides practical guidance on big data analytics and data science. Expert perspectives are provided by authoritative researchers and practitioners from around the world, who discuss research developments and emerging trends, present case studies on helpful frameworks and innovative methodologies, and suggest best practices for efficient and effective data analytics. Chiefly intended for researchers, IT professionals and graduate students, the book represents a timely contribution to the growing field of big data, which has been recognized as one of the leading emerging technologies that will have a major impact on various fields of science and various aspects of human society over the next several decades. Therefore, the content in this book is an essential tool to help readers understand current developments, and provides them with an extensive overview of the field of big data analytics as it is practiced today. The chapters cover technical aspects of key areas that generate and use big data, such as management and finance, medicine and health care, networks, the Internet of Things, big data standards, benchmarking of systems, and others. In addition to a diverse range of applications, key algorithmic approaches such as graph partitioning, clustering and finite mixture modeling of high-dimensional data are also covered. The varied collection of topics addressed introduces readers to the richness of the emerging field of big data analytics.

## **Data Management in the Cloud**

Learn how to employ JADE to build multi-agent systems! JADE (Java Agent DEvelopment framework) is a middleware for the development of applications, both in the mobile and fixed environment, based on the Peer-to-Peer intelligent autonomous agent approach. JADE enables developers to implement and deploy multi-agent systems, including agents running on wireless networks and limited-resource devices. Developing Multi-Agent Systems with JADE is a practical guide to using JADE. The text will give an introduction to agent technologies and the JADE Platform, before proceeding to give a comprehensive guide to programming with JADE. Basic features such as creating agents, agent tasks, agent communication, agent discovery and GUIs are covered, as well as more advanced features including ontologies and content languages, complex behaviours, interaction protocols, agent mobility, and the in-process interface. Issues such as JADE internals, running JADE agents on mobile devices, deploying a fault tolerant JADE platform, and main add-ons are also

covered in depth. Developing Multi-Agent Systems with JADE: Comprehensive guide to using JADE to build multi-agent systems and agent orientated programming. Describes and explains ontologies and content language, interaction protocols and complex behaviour. Includes material on persistence, security and a semantics framework. Contains numerous examples, problems, and illustrations to enhance learning. Presents a case study demonstrating the use of JADE in practice. Offers an accompanying website with additional learning resources such as sample code, exercises and PPT-slides. This invaluable resource will provide multi-agent systems practitioners, programmers working in the software industry with an interest on multi-agent systems as well as final year undergraduate and postgraduate students in CS and advanced networking and telecoms courses with a comprehensive guide to using JADE to employ multi agent systems. With contributions from experts in JADE and multi agent technology.

## **New Trends in Image Analysis and Processing -- ICIAP 2015 Workshops**

This book presents the proceedings of the 6th International Conference on Frontiers of Intelligent Computing: Theory and Applications (FICTA-2017), held in Bhubaneswar, Odisha. The event brought together researchers, scientists, engineers, and practitioners to exchange their new ideas and experiences in the domain of intelligent computing theories with prospective applications to various engineering disciplines. The book is divided into two volumes: Information and Decision Sciences, and Intelligent Engineering Informatics. This volume covers broad areas of Information and Decision Sciences, with papers exploring both the theoretical and practical aspects of data-intensive computing, data mining, evolutionary computation, knowledge management & networks, sensor networks, signal processing, wireless networks, protocols & architectures etc. The book also offers a valuable resource for students at the post-graduate level in various engineering disciplines.

## **Developing Web Applications with Apache, MySQL, memcached, and Perl**

The four-volume set LNCS 11334-11337 constitutes the proceedings of the 18th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2018, held in Guangzhou, China, in November 2018. The 141 full and 50 short papers presented were carefully reviewed and selected from numerous submissions. The papers are organized in topical sections on Distributed and Parallel Computing; High Performance Computing; Big Data and Information Processing; Internet of Things and Cloud Computing; and Security and Privacy in Computing.

## **Applications in Ubiquitous Computing**

This two-volume book focuses on both theory and applications in the broad areas of communication technology, computer science and information security. It brings together contributions from scientists, professors, scholars and students, and

presents essential information on computing, networking, and informatics. It also discusses the practical challenges encountered and the solutions used to overcome them, the goal being to promote the “translation” of basic research into applied research, and of applied research into practice. The works presented here will also demonstrate the importance of basic scientific research in a range of fields.

### **Large Scale and Big Data**

"The Art of Agent-Oriented Modeling is an introduction to agent-oriented software development for students and for software developers who are interested in learning about new software engineering techniques." --FOREWORD.

### **Big Data and Hadoop**

In just a few chapters you will learn about Tcl features that allow you to isolate and protect your code from being damaged in large applications. You will even learn how to extend the language itself. Tcl/Tk: A Developer's Guide clearly discusses development tools, proven techniques, and existing extensions. It shows how to use Tcl/Tk effectively and provides many code examples. This fully revised new edition is the complete resource for computer professionals, from systems administrators to programmers. It covers versions 7.4 to 8.4 and includes a CD-ROM containing the interpreters, libraries, and tutorials to get you started quickly. Additional materials in the book include case studies and discussions of techniques for the advanced user. On the CD-ROM \*Distributions for Tcl 8.3 and 8.4 for Linux, Solaris, Macintosh, and Windows. \*A copy of ActiveTcl from ActiveState. \*The latest release of TclTutor. \*How-to's and tutorials as well as copies of all the tools discussed in the book.

### **The Art of Application Performance Testing**

Big Data: Principles and Paradigms captures the state-of-the-art research on the architectural aspects, technologies, and applications of Big Data. The book identifies potential future directions and technologies that facilitate insight into numerous scientific, business, and consumer applications. To help realize Big Data's full potential, the book addresses numerous challenges, offering the conceptual and technological solutions for tackling them. These challenges include life-cycle data management, large-scale storage, flexible processing infrastructure, data modeling, scalable machine learning, data analysis algorithms, sampling techniques, and privacy and ethical issues. Covers computational platforms supporting Big Data applications Addresses key principles underlying Big Data computing Examines key developments supporting next generation Big Data platforms Explores the challenges in Big Data computing and ways to overcome them Contains expert contributors from both academia and industry

## **Linked Data Management**

This book introduces you to the Big Data processing techniques addressing but not limited to various BI (business intelligence) requirements, such as reporting, batch analytics, online analytical processing (OLAP), data mining and Warehousing, and predictive analytics. The book has been written on IBMs Platform of Hadoop framework. IBM Infosphere BigInsight has the highest amount of tutorial matter available free of cost on Internet which makes it easy to acquire proficiency in this technique. This therefore becomes highly vulnerable coaching materials in easy to learn steps. The book optimally provides the courseware as per MCA and M. Tech Level Syllabi of most of the Universities. All components of big Data Platform like Jaql, Hive Pig, Sqoop, Flume , Hadoop Streaming, Oozie: HBase, HDFS, FlumeNG, Whirr, Cloudera, Fuse , Zookeeper and Mahout: Machine learning for Hadoop has been discussed in sufficient Detail with hands on Exercises on each.

## **Big Data Analytics and Knowledge Discovery**

This book is an easy-to-follow guide full of hands-on examples of real-world administration tasks, which will help you design and build highly scalable and very fast web applications. This book is ideal for web application developers looking for hands-on, practical, and to-the-point recipes to integrate the Memcached service into their day-to-day programming tasks. No prior knowledge of Memcached is assumed for using this book.

## **The Myth of Homeland Security**

This handbook offers comprehensive coverage of recent advancements in Big Data technologies and related paradigms. Chapters are authored by international leading experts in the field, and have been reviewed and revised for maximum reader value. The volume consists of twenty-five chapters organized into four main parts. Part one covers the fundamental concepts of Big Data technologies including data curation mechanisms, data models, storage models, programming models and programming platforms. It also dives into the details of implementing Big SQL query engines and big stream processing systems. Part Two focuses on the semantic aspects of Big Data management including data integration and exploratory ad hoc analysis in addition to structured querying and pattern matching techniques. Part Three presents a comprehensive overview of large scale graph processing. It covers the most recent research in large scale graph processing platforms, introducing several scalable graph querying and mining mechanisms in domains such as social networks. Part Four details novel applications that have been made possible by the rapid emergence of Big Data technologies such as Internet-of-Things (IOT), Cognitive Computing and SCADA Systems. All parts of the book discuss open research problems, including potential opportunities, that have arisen from the rapid progress of Big Data technologies and the associated increasing

requirements of application domains. Designed for researchers, IT professionals and graduate students, this book is a timely contribution to the growing Big Data field. Big Data has been recognized as one of leading emerging technologies that will have a major contribution and impact on the various fields of science and various aspects of the human society over the coming decades. Therefore, the content in this book will be an essential tool to help readers understand the development and future of the field.

### **Data Mining**

Over 60 recipes to help you design interactive, smart, and cross-platform GUI applications  
Key Features  
Get succinct QT solutions to pressing GUI programming problems in Python  
Learn how to effectively implement reactive programming  
Build customized applications that are robust and reliable  
Book Description  
PyQt is one of the best cross-platform interface toolkits currently available; it's stable, mature, and completely native. If you want control over all aspects of UI elements, PyQt is what you need. This book will guide you through every concept necessary to create fully functional GUI applications using PyQt, with only a few lines of code. As you expand your GUI using more widgets, you will cover networks, databases, and graphical libraries that greatly enhance its functionality. Next, the book guides you in using Qt Designer to design user interfaces and implementing and testing dialogs, events, the clipboard, and drag and drop functionality to customize your GUI. You will learn a variety of topics, such as look and feel customization, GUI animation, graphics rendering, implementing Google Maps, and more. Lastly, the book takes you through how Qt5 can help you to create cross-platform apps that are compatible with Android and iOS. You will be able to develop functional and appealing software using PyQt through interesting and fun recipes that will expand your knowledge of GUIs  
What you will learn  
Use basic Qt components, such as a radio button, combo box, and sliders  
Use QSpinBox and sliders to handle different signals generated on mouse clicks  
Work with different Qt layouts to meet user interface requirements  
Create custom widgets and set up customizations in your GUI  
Perform asynchronous I/O operations and thread handling in the Python GUI  
Employ network concepts, internet browsing, and Google Maps in UI  
Use graphics rendering and implement animation in your GUI  
Make your GUI application compatible with Android and iOS devices  
Who this book is for  
If you're an intermediate Python programmer wishing to enhance your coding skills by writing powerful GUIs in Python using PyQt, this is the book for you.

### **Big Data 2.0 Processing Systems**

The conference provides a forum for engineers and scientists in academia, industry, and government to present their latest research findings in current and emerging areas of distributed computing  
Topics of interest include, but are not limited to  
Cloud Computing and Data Center Systems  
Big Data Systems and Distributed Data Management & Analytics  
Distributed Operating Systems and Middleware  
Distributed Algorithms and Theory  
Distributed Fault Tolerance and Dependability

Security, Privacy, and Trust in Distributed Systems Internet of Things, Smart Cities, and Cyber Physical Systems Distributed Green Computing and Energy Management Mobile and Wireless Computing Edge and Fog Computing Social Networks and Crowdsourcing The technical program consists of research tracks in 11 focused areas In addition, there are three special topic tracks the Industry and Experimentation track, the Applications and Experiences track, and the Vision and Blue Sky Thinking track

## **Ontology-Based Multi-Agent Systems**

Sometimes, it's not the lingering relationships that stick in our memories, but rather the brief, fleeting encounters. In the brilliant short story "Four Meetings," literary master Henry James explores an intense infatuation that ran its entire course in only four meetings.

## **Linked Data**

This book provides readers the “big picture” and a comprehensive survey of the domain of big data processing systems. For the past decade, the Hadoop framework has dominated the world of big data processing, yet recently academia and industry have started to recognize its limitations in several application domains and big data processing scenarios such as the large-scale processing of structured data, graph data and streaming data. Thus, it is now gradually being replaced by a collection of engines that are dedicated to specific verticals (e.g. structured data, graph data, and streaming data). The book explores this new wave of systems, which it refers to as Big Data 2.0 processing systems. After Chapter 1 presents the general background of the big data phenomena, Chapter 2 provides an overview of various general-purpose big data processing systems that allow their users to develop various big data processing jobs for different application domains. In turn, Chapter 3 examines various systems that have been introduced to support the SQL flavor on top of the Hadoop infrastructure and provide competing and scalable performance in the processing of large-scale structured data. Chapter 4 discusses several systems that have been designed to tackle the problem of large-scale graph processing, while the main focus of Chapter 5 is on several systems that have been designed to provide scalable solutions for processing big data streams, and on other sets of systems that have been introduced to support the development of data pipelines between various types of big data processing jobs and systems. Lastly, Chapter 6 shares conclusions and an outlook on future research challenges. Overall, the book offers a valuable reference guide for students, researchers and professionals in the domain of big data processing systems. Further, its comprehensive content will hopefully encourage readers to pursue further research on the subject.

## **Algorithms and Architectures for Parallel Processing**

## **Wireless Rechargeable Sensor Networks**

Providing proven strategies for solutions to research, development, and production dilemmas, this reference details the instrumentation and underlying principles for utilization of electron microscopy in the manufacturing, automotive, semiconductor, photographic film, pharmaceutical, chemical, mineral, forensic, glass, and pulp and paper industries. The book covers safety, calibration, and troubleshooting techniques, as well as methods in sample preparation and image collection, interpretation, and analysis. It includes contributions from microscopy experts based at major corporations and scientists from universities and major research centers.

## **Information and Decision Sciences**

This book constitutes the refereed proceedings of the 11th International Conference on Security, Privacy, and Anonymity in Computation, Communication, and Storage. The 45 revised full papers were carefully reviewed and selected from 120 submissions. The papers cover many dimensions including security algorithms and architectures, privacy-aware policies, regulations and techniques, anonymous computation and communication, encompassing fundamental theoretical approaches, practical experimental projects, and commercial application systems for computation, communication and storage.

## **Tcl/Tk**

This book constitutes the refereed proceedings of the 21st International Conference on Big Data Analytics and Knowledge Discovery, DaWaK 2019, held in Linz, Austria, in September 2019. The 12 full papers and 10 short papers presented were carefully reviewed and selected from 61 submissions. The papers are organized in the following topical sections: Applications; patterns; RDF and streams; big data systems; graphs and machine learning; databases.

## **Getting Started with Memcached**

The four-volume set LNCS 11334-11337 constitutes the proceedings of the 18th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2018, held in Guangzhou, China, in November 2018. The 141 full and 50 short papers presented were carefully reviewed and selected from numerous submissions. The papers are organized in topical sections on Distributed and Parallel Computing; High Performance Computing; Big Data and Information Processing; Internet of Things and Cloud Computing; and Security and Privacy in Computing.



[ROMANCE](#) [ACTION & ADVENTURE](#) [MYSTERY & THRILLER](#) [BIOGRAPHIES & HISTORY](#) [CHILDREN'S](#) [YOUNG ADULT](#) [FANTASY](#)  
[HISTORICAL FICTION](#) [HORROR](#) [LITERARY FICTION](#) [NON-FICTION](#) [SCIENCE FICTION](#)