

Getting Started With Memcached Soliman Ahmed

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

The past decade has witnessed the rapid growth of data in large-scale distributed storage systems. Triplication, a reliability mechanism with 3x storage overhead and adopted by large-scale distributed storage systems, introduces heavy storage cost as data amount in storage systems keep growing. Consequently, erasure codes have been introduced in many storage systems because they can provide a higher storage efficiency and fault tolerance than data replication. However, erasure coding has many performance degradation factors in both I/O and computation operations, resulting in great performance degradation in large-scale erasure-coded storage systems. In this thesis, we investigate how to eliminate some key performance

issues in I/O and computation operations for applying erasure coding in large-scale storage systems. We also propose a prototype named ESetStore to improve the recovery performance of erasure-coded storage systems. We introduce our studies as follows. First, we study the encoding and decoding performance of the erasure coding, which can be a key bottleneck with the state-of-the-art disk I/O throughput and network bandwidth. We propose a graphics processing unit (GPU)-based implementation of erasure coding named G-CRS, which employs the Cauchy Reed-Solomon (CRS) code, to improve the encoding and decoding performance. To maximize the coding performance of G-CRS by fully utilizing the GPU computational power, we designed and implemented a set of optimization strategies. Our evaluation results demonstrated that G-CRS is 10 times faster than most of the other coding libraries. Second, we investigate the performance degradation introduced by intensive I/O operations in recovery for large-scale erasure-coded storage systems. To improve the recovery performance, we propose a data placement algorithm named ESet. We define a configurable parameter named overlapping factor for system administrators to easily achieve desirable recovery I/O parallelism. Our simulation results show that ESet can significantly improve the data recovery performance without violating the reliability requirement by distributing data and code blocks across different failure domains. Third, we take a look at the performance of applying coding techniques to in-memory storage. A reliable in-memory cache for key-value stores named R-Memcached is designed and proposed. This work can be served as a prelude of applying erasure coding to in-memory metadata storage. R-Memcached exploits coding techniques to achieve reliability, and can tolerate up to two node failures. Our experimental results show that R-Memcached can maintain very good latency and throughput performance even during the

period of node failures. At last, we design and implement a prototype named ESetStore for erasure-coded storage systems. The ESetStore integrates our data placement algorithm ESet to bring fast data recovery for storage systems.

This book constitutes the refereed proceedings of the 30th International Conference, ISC High Performance 2015, [formerly known as the International Supercomputing Conference] held in Frankfurt, Germany, in July 2015. The 27 revised full papers presented together with 10 short papers were carefully reviewed and selected from 67 submissions. The papers cover the following topics: cost-efficient data centers, scalable applications, advances in algorithms, scientific libraries, programming models, architectures, performance models and analysis, automatic performance optimization, parallel I/O and energy efficiency.

This book starts with an introduction to process modeling and process paradigms, then explains how to query and analyze process models, and how to analyze the process execution data. In this way, readers receive a comprehensive overview of what is needed to identify, understand and improve business processes. The book chiefly focuses on concepts, techniques and methods. It covers a large body of knowledge on process analytics – including process data querying, analysis, matching and correlating process data and models – to help practitioners and researchers understand the underlying concepts, problems, methods, tools and techniques involved in modern process analytics. Following an introduction to basic business process and process analytics concepts, it describes the state of the art in this area before examining different analytics techniques in detail. In this regard, the book covers analytics over different levels of process abstractions, from process execution data and methods for linking and correlating process execution data, to inferring process models,

querying process execution data and process models, and scalable process data analytics methods. In addition, it provides a review of commercial process analytics tools and their practical applications. The book is intended for a broad readership interested in business process management and process analytics. It provides researchers with an introduction to these fields by comprehensively classifying the current state of research, by describing in-depth techniques and methods, and by highlighting future research directions. Lecturers will find a wealth of material to choose from for a variety of courses, ranging from undergraduate courses in business process management to graduate courses in business process analytics. Lastly, it offers professionals a reference guide to the state of the art in commercial tools and techniques, complemented by many real-world use case scenarios.

The computing world today is in the middle of a revolution: mobile clients and cloud computing have emerged as the dominant paradigms driving programming and hardware innovation today. The Fifth Edition of Computer Architecture focuses on this dramatic shift, exploring the ways in which software and technology in the cloud are accessed by cell phones, tablets, laptops, and other mobile computing devices. Each chapter includes two real-world examples, one mobile and one datacenter, to illustrate this revolutionary change. Updated to cover the mobile computing revolution Emphasizes the two most important topics in architecture today: memory hierarchy and parallelism in all its forms. Develops common themes throughout each chapter: power, performance, cost, dependability, protection, programming models, and emerging trends ("What's Next") Includes three review appendices in the printed text. Additional reference appendices are available online. Includes updated Case Studies and completely new exercises.

Orchestrating asynchronous function calls in JavaScript often leads to callback hell, but there is a reliable way to avoid this painful state of affairs. With this concise and simple guide, you'll learn how to use jQuery deferreds and promises, an elegant approach for managing asynchronous calls in both client and server applications. This book contains 18 examples that use deferreds to solve progressively challenging real-world programming problems, along with 75 stimulating puzzles (and their solutions) that will help you understand how and when to use deferreds. Experienced JavaScript programmers will learn new tricks in a fun way, and become immersed in the practice of event-based programming. Understand the logic behind creating deferreds and returning promises Get a structured explanation of jQuery's deferred API Delve into the dynamics of using deferreds Explore a broad collection of useful deferred recipes developed by the authors Gain hands-on experience by solving challenges that accompany each recipe Go deeper into deferreds: encounter novel abstractions and mind-bending use cases

Handbook of Big Data provides a state-of-the-art overview of the analysis of large-scale datasets. Featuring contributions from well-known experts in statistics and computer science, this handbook presents a carefully curated collection of techniques from both industry and academia. Thus, the text instills a working understanding of key statistical

This book constitutes the refereed proceedings of the 18th International Conference on Runtime Verification, RV 2018, held in Limassol, Cyprus, in November 2018. The 21 full papers presented together with 3 short papers and 3 tool papers were carefully reviewed and selected from 49 submissions. The RV conference is concerned with all

aspects of monitoring and analysis of hardware, software and more general system executions. Runtime verification techniques are lightweight techniques to assess correctness, reliability, and robustness; these techniques are significantly more powerful and versatile than conventional testing, and more practical than exhaustive formal verification. Chapter “Hardware-based Runtime Verification with Embedded Tracing Units and Stream Processing” is available open access under a Creative Commons Attribution 4.0 International License via link.springer.com.

Large-scale, highly interconnected networks, which are often modeled as graphs, pervade both our society and the natural world around us. Uncertainty, on the other hand, is inherent in the underlying data due to a variety of reasons, such as noisy measurements, lack of precise information needs, inference and prediction models, or explicit manipulation, e.g., for privacy purposes. Therefore, uncertain, or probabilistic, graphs are increasingly used to represent noisy linked data in many emerging application scenarios, and they have recently become a hot topic in the database and data mining communities. Many classical algorithms such as reachability and shortest path queries become $\#P$ -complete and, thus, more expensive over uncertain graphs. Moreover, various complex queries and analytics are also emerging over uncertain networks, such as pattern matching, information diffusion, and influence maximization queries. In this book, we discuss the sources of uncertain graphs and their applications, uncertainty modeling, as well as the complexities and algorithmic advances on

uncertain graphs processing in the context of both classical and emerging graph queries and analytics. We emphasize the current challenges and highlight some future research directions.

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.

This book constitutes the proceedings of the Second International Conference on Cryptology in Africa, AFRICACRYPT 2009, held in Gammarth, Tunisia, on June 21-25, 2009. The 25 papers presented together with one invited talk were carefully reviewed and selected from 70 submissions. The topics covered are hash functions, block ciphers, asymmetric encryption, digital signatures, asymmetric encryption and anonymity, key agreement protocols, cryptographic protocols, efficient implementations, and implementation attacks.

ASPLOS '17: Architectural Support for Programming Languages and Operating Systems Apr 08, 2017-Apr 12, 2017 Xi'an, China. You can view more information about this proceeding and all of ACMs other published conference proceedings from the ACM Digital Library: <http://www.acm.org/dl>.

With Learning JavaScript Design Patterns, you'll learn how to write beautiful, structured, and maintainable JavaScript by applying classical and modern design patterns to the language. If you want to keep your code efficient, more manageable, and up-to-date with the latest best practices, this book is for you. Explore many popular design patterns, including Modules, Observers, Facades, and Mediators. Learn how modern architectural patterns—such as MVC, MVP, and MVVM—are useful from the

perspective of a modern web application developer. This book also walks experienced JavaScript developers through modern module formats, how to namespace code effectively, and other essential topics. Learn the structure of design patterns and how they are written Understand different pattern categories, including creational, structural, and behavioral Walk through more than 20 classical and modern design patterns in JavaScript Use several options for writing modular code—including the Module pattern, Asynchronous Module Definition (AMD), and CommonJS Discover design patterns implemented in the jQuery library Learn popular design patterns for writing maintainable jQuery plug-ins "This book should be in every JavaScript developer's hands. It's the go-to book on JavaScript patterns that will be read and referenced many times in the future."—Andrée Hansson, Lead Front-End Developer, *presis!*

This lecture presents various aspects of uncertainty in schema matching within a single unified framework. We introduce basic formulations of uncertainty and provide several alternative representations of schema matching uncertainty. Then, we cover two common methods that have been proposed to deal with uncertainty in schema matching, namely ensembles and top-K matchings, and analyze them in this context. We conclude with a set of real-world applications. Schema matching is the task of providing correspondences between concepts describing the meaning of data in various heterogeneous, distributed data sources. Schema matching is one of the basic operations required by the process of data and schema integration, and thus has a

great effect on its outcomes, whether these involve targeted content delivery, view integration, database integration, query rewriting over heterogeneous sources, duplicate data elimination, or automatic streamlining of workflow activities that involve heterogeneous data sources. Schema matching research has been going on for more than 25 years. Over the years, a significant body of work has been devoted to the identification of schema matchers, heuristics for schema matching. The main objective of schema matchers is to provide correspondences that will be effective from the user's point of view. Over the years, a realization has emerged that schema matchers are inherently uncertain. Since 2003, work on the uncertainty in schema matching has picked up, along with research on uncertainty in other areas of data management. The Indus region, comprising the northwestern part of the Indian subcontinent (now Pakistan), has always had its distinct identity - racially, ethnically, linguistically and culturally. In the last five thousand years, this region has been a part of India, politically, for only five hundred years. Pakistan, then, is no 'artificial' state conjured up by the disaffected Muslim elite of British India. Aitzaz Ahsan surveys the history of Indus - as he refers to this region - right from the time of the Harappan civilization to the era of the British Raj, concluding with independence and the creation of Pakistan. Ahsan's message is aimed both at Indians still nostalgic about 'undivided' India and their Pakistani compatriots who

narrowly tend to define their identity by their 'un-Indianness'.

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.

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.

The Internet and World Wide Web have revolutionized access to information. Users now store information across multiple platforms from personal computers to smartphones and websites. As a consequence, data management concepts, methods and techniques are increasingly focused on distribution concerns. Now that information largely resides in the network, so do the tools that process this information. This book explains the foundations of XML with a focus on data distribution. It covers the many facets of distributed data management on the Web, such as description logics, that are already emerging in today's data integration applications and herald tomorrow's semantic Web. It also introduces the machinery used to manipulate the unprecedented amount of data collected on the Web. Several 'Putting into Practice' chapters describe detailed practical applications of the technologies and techniques. The book will serve as an introduction to the new, global, information systems for Web professionals and master's level courses.

Getting Started with MemcachedPackt Pub Limited

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 design systems that carry out 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 technologies. We also present a methodology that will guide us in the design of the integrated 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 issues.
- 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.

What happens when something is sucked into a black hole? Does it disappear? Three decades ago, a young physicist named Stephen Hawking claimed it did—and in doing so put at risk everything we know about physics and the fundamental laws of the universe. Most scientists didn't recognize the import of Hawking's claims, but Leonard Susskind and Gerard 'tHooft realized the threat, and responded with a counterattack that changed the course of physics. **THE BLACK HOLE WAR** is the thrilling story of their united effort to reconcile Hawking's revolutionary theories of black holes with their own sense of reality—effort that would eventually result in Hawking admitting he was wrong, paying up, and Susskind and 'tHooft realizing that our world is a hologram projected from the outer boundaries of space. A brilliant book about modern physics, quantum mechanics, the fate of stars and the deep mysteries of black holes, Leonard Susskind's account of the Black Hole War is mind-bending and exhilarating reading.

The bug bounty hunting community is full of technical resources. However, any successful hunter will tell you that succeeding in this industry takes more than technical knowledge. Without the proper mindset, the effective tactics and the key soft skills, here is the hard truth: You won't last in the bug bounty hunting game. You might find few

bugs at first, but you won't stand the lack of motivation and self-esteem when you can't find bugs for few weeks. After months, the situation may even develop to burnout. If you understand and exploit known security vulnerabilities in CTF challenges but still struggle to find bugs in real-world targets, this book is for you. I wrote this book with a single purpose in mind: Help you understand and master essential skills to become a successful bug bounty hunter, in an entertaining way. To achieve this goal, I designed the book around the story of Anna, a fictitious Junior Security Engineer who has just heard of bug bounty hunting. Throughout her fascinating journey, you will witness all the steps she took to get started the right way. You will observe all the limits she discovers about herself, and you will grasp all the proven solutions she came up with to overcome them, collect 1000 reputation points and earn her first \$5000 along the way. Whether you have just started or have spent years in this industry, you will undoubtedly identify with the different hurdles of the story. I am sure you will add some missing tricks to your toolset to succeed in bug bounty hunting. At the end of the story, you will find technical appendices that support Anna's journey. There, you will find how to approach a bug bounty program for the first time, and how to perform in-depth web application hacking to increase your chances of finding bugs. You can read this book from cover to cover while bookmarking the pivot points along the story. Then, you can go back to each crucial moment whenever you face the same situation. Sit tight and enjoy the ride!

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

When, in 1984, Richard P. Feynman gave his famous course on computation at the California Institute of Technology, he asked Tony Hey to adapt his lecture notes into a book. Although led by Feynman, the course also featured, as occasional guest speakers, some of the most brilliant men in science at that time, including Marvin Minsky, Charles Bennett, and John Hopfield. Although the lectures are now thirteen years old, most of the material is timeless and presents a "Feynmanesque" overview of many standard and some not-so-standard topics in computer science such as reversible logic gates and quantum computers.

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.

If you understand basic mathematics and know how to program with Python, you're

ready to dive into signal processing. While most resources start with theory to teach this complex subject, this practical book introduces techniques by showing you how they're applied in the real world. In the first chapter alone, you'll be able to decompose a sound into its harmonics, modify the harmonics, and generate new sounds. Author Allen Downey explains techniques such as spectral decomposition, filtering, convolution, and the Fast Fourier Transform. This book also provides exercises and code examples to help you understand the material. You'll explore: Periodic signals and their spectrums Harmonic structure of simple waveforms Chirps and other sounds whose spectrum changes over time Noise signals and natural sources of noise The autocorrelation function for estimating pitch The discrete cosine transform (DCT) for compression The Fast Fourier Transform for spectral analysis Relating operations in time to filters in the frequency domain Linear time-invariant (LTI) system theory Amplitude modulation (AM) used in radio Other books in this series include Think Stats and Think Bayes, also by Allen Downey.

An illustrated, large-format edition of the best-seller has been expanded to encompass the remarkable advances that have occurred in science and technology over the past eight years, with a new chapter on Wormholes and Time Travel and more than 240 full-color, captioned illustrations. 100,000 first printing.

A work on turbulent premixed combustion is timely because of increased concern about the environmental impact of combustion and the search for new combustion concepts

and technologies. An improved understanding of lean fuel turbulent premixed flames must play a central role in the fundamental science of these new concepts. Lean premixed flames have the potential to offer ultra-low emission levels, but they are notoriously susceptible to combustion oscillations. Thus, sophisticated control measures are inevitably required. The editors' intent is to set out the modeling aspects in the field of turbulent premixed combustion. Good progress has been made recently on this topic. Thus, it is timely to edit a cohesive volume containing contributions from international experts on various subtopics of the lean premixed flame problem. 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.

For centuries, a few high caste Hindus practiced sati (or suttee), a ritual that required a widow to surrender herself to be burned on her husband's funeral pyre. Although sati is illegal and extremely rare today, remnants of the custom in a social sense survive throughout most of Nepali society. Red symbolizes life, vibrancy, and passion—a color

and existence that is often denied to women who have lost their husbands. Now, as then, by the time her husband's funeral fire is extinguished, the woman as wife no longer exists. Her separation from both family and society has begun.

This book constitutes the refereed proceedings of the 13th International Conference on Information Security, ISC 2010, held in Boca Raton, FL, USA, in October 2010. The 25 revised full papers and the 11 short papers presented were carefully reviewed and selected from 117 submissions. The papers are organized in topical sections on attacks and analysis; analysis; authentication, PIR and content identification; privacy; malware, crimeware and code injection; intrusion detection; side channels; cryptography; smartphones; biometrics; cryptography, application; buffer overflow; and cryptography, theory.

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. In today's competitive and highly dynamic environment, analyzing data to understand how the business is performing, to predict outcomes and trends, and to improve the effectiveness of business processes underlying business operations has become critical. The traditional approach to reporting is no longer adequate, users now demand easy-to-use intelligent platforms and applications capable of analyzing real-time business data to provide insight and actionable information at the right time. The end goal is to improve the enterprise performance by better and timelier decision making, - abled by the availability of up-to-date, high-quality information. As a response, the notion of "real-time enterprise" has emerged and is beginning to be recognized in the industry.

Gartner defines it as “using up-to-date information, getting rid of delays, and using speed for competitive advantage is what the real-time enterprise is all about. . . Indeed, the goal of the real-time enterprise is to act on events as they happen.” Although there has been progress in this direction and many companies are introducing products toward making this vision a reality, there is still a long way to go. In particular, the whole lifecycle of business intelligence requires new techniques and methodologies capable of dealing with the new requirements imposed by the real-time enterprise.

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 varies aspect 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.

"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.

[Copyright: e4ac3640a38bd12cb322797c9cd6b72e](https://www.pdfdrive.com/getting-started-with-memcached-soliman-ahmed-pdf/e4ac3640a38bd12cb322797c9cd6b72e)