

Es/search/node

Right here, we have countless book **es/search/node** and collections to check out. We additionally present variant types and then type of the books to browse. The customary book, fiction, history, novel, scientific research, as well as various extra sorts of books are readily understandable here.

As this es/search/node, it ends happening swine one of the favored ebook es/search/node collections that we have. This is why you remain in the best website to look the amazing book to have.

Graph-Theoretic Concepts in Computer Science Andreas Brandstädt 2001-09-26 This book constitutes the thoroughly refereed post-workshop proceedings of the 27th International Workshop on Graph-Theoretic Concepts in Computer Science, WG 2001, held in Boltenhagen, Germany, in June 2001. The 27 revised full papers presented together with two invited contributions were carefully reviewed and selected from numerous submissions. The papers provide a wealth of new results for various classes of graphs, graph computations, graph algorithms and graph-theoretical applications in various fields.

Automated Verification of Concurrent Search Structures Krishna Siddharth 2022-05-31 Search structures support the fundamental data storage primitives on key-value pairs: insert a pair, delete by key, search by key, and update the value associated with a key. Concurrent search structures are parallel algorithms to speed access to search structures on multicore and distributed servers. These sophisticated algorithms perform fine-grained synchronization between threads, making them notoriously difficult to design correctly. Indeed, bugs have been found both in actual implementations and in the designs proposed by experts in peer-reviewed publications. The rapid development and

deployment of these concurrent algorithms has resulted in a rift between the algorithms that can be verified by the state-of-the-art techniques and those being developed and used today. The goal of this book is to show how to bridge this gap in order to bring the certified safety of formal verification to high-performance concurrent search structures. Similar techniques and frameworks can be applied to concurrent graph and network algorithms beyond search structures.

Foundations of Multidimensional and Metric Data

Structures Hanan Samet 2006-08-22 Publisher Description
Database Systems for Advanced Applications '97 Rodney Topor 1997 This volume contains the proceedings of the Fifth International Conference on Database Systems for Advanced Applications (DASFAA '97). DASFAA '97 focused on advanced database technologies and their applications. The 55 papers in this volume cover a wide range of areas in the field of database systems and applications ? including the rapidly emerging areas of the Internet, multimedia, and document database systems ? and should be of great interest to all database system researchers and developers, and practitioners.

Pearls of Functional Algorithm Design

Richard Bird 2010-09-16 Richard Bird takes a radical approach to algorithm

design, namely, design by calculation. These 30 short chapters each deal with a particular programming problem drawn from sources as diverse as games and puzzles, intriguing combinatorial tasks, and more familiar areas such as data compression and string matching. Each pearl starts with the statement of the problem expressed using the functional programming language Haskell, a powerful yet succinct language for capturing algorithmic ideas clearly and simply. The novel aspect of the book is that each solution is calculated from an initial formulation of the problem in Haskell by appealing to the laws of functional programming. Pearls of Functional Algorithm Design will appeal to the aspiring functional programmer, students and teachers interested in the principles of algorithm design, and anyone seeking to master the techniques of reasoning about programs in an equational style.

Parameterized and Exact Computation Frank Dehne 2004-10-14
 The central challenge of theoretical computer science is to deploy mathematical insights in ways that serve the creation of useful algorithms. In recent years there has been a growing interest in the two-dimensional framework of parameterized complexity, where, in addition to the overall input size, one also considers a parameter, with a focus on how these two dimensions interact in problem complexity. This book presents the proceedings of the 1st International Workshop on Parameterized and Exact Computation (IWPEC 2004, <http://www.iwpec.org>), which took place in Bergen, Norway, on September 14-16, 2004. The workshop was organized as part of ALGO 2004. There were seven previous workshops on the theory and applications of parameterized complexity. The first was organized at the Institute for the Mathematical Sciences in Chennai, India, in September, 2000. The second was held at Dagstuhl Castle, Germany, in July, 2001. In December, 2002, a workshop on parameterized complexity was held in conjunction with the FST-TCS meeting in Kanpur, India. A second Dagstuhl workshop on parameterized complexity was held in July, 2003. Another wo-

shop on the subject was held in Ottawa, Canada, in August, 2003, in conjunction with the WADS 2003 meeting. There have also been two Barbados workshops on applications of parameterized complexity. In response to the IWPEC 2004 call for papers, 47 papers were submitted, and from these the program committee selected 25 for presentation at the workshop.

In addition, invited lectures were accepted by the distinguished researchers Michael Langston and Gerhard Woeginger.

Computing and Combinatorics China) COCOON 97 (1997 : Shanghai 1997-07-30 The book is aimed at graduate students, researchers, engineers and physicists involved in fluid computations. An up-to-date account is given of the present state of the art of numerical methods employed in computational fluid dynamics. The underlying numerical principles are treated with a fair amount of detail, using elementary methods. Attention is given to the difficulties arising from geometric complexity of the flow domain. Uniform accuracy for singular perturbation problems is studied, pointing the way to accurate computation of flows at high Reynolds number. Unified methods for compressible and incompressible flows are discussed. A treatment of the shallow-water equations is included. A basic introduction is given to efficient iterative solution methods. Many pointers are given to the current literature, facilitating further study.

ECAI 2012 L. De Raedt 2012-08-15 Artificial intelligence (AI) plays a vital part in the continued development of computer science and informatics. The AI applications employed in fields such as medicine, economics, linguistics, philosophy, psychology and logical analysis, not forgetting industry, are now indispensable for the effective functioning of a multitude of systems. This book presents the papers from the 20th biennial European Conference on Artificial Intelligence, ECAI 2012, held in Montpellier, France, in August 2012. The ECAI conference remains Europe's principal opportunity for researchers and practitioners of Artificial Intelligence to gather and to discuss the latest trends and

challenges in all subfields of AI, as well as to demonstrate innovative applications and uses of advanced AI technology. ECAI 2012 featured four keynote speakers, an extensive workshop program, seven invited tutorials and the new Frontiers of Artificial Intelligence track, in which six invited speakers delivered perspective talks on particularly interesting new research results, directions and trends in Artificial Intelligence or in one of its related fields. The proceedings of PAIS 2012 and the System Demonstrations Track are also included in this volume, which will be of interest to all those wishing to keep abreast of the latest developments in the field of AI.

Managing Distributed Cloud Applications and Infrastructure

Theo Lynn 2020-08-21 The emergence of the Internet of Things (IoT), combined with greater heterogeneity not only online in cloud computing architectures but across the cloud-to-edge continuum, is introducing new challenges for managing applications and infrastructure across this continuum. The scale and complexity is simply so complex that it is no longer realistic for IT teams to manually foresee the potential issues and manage the dynamism and dependencies across an increasing inter-dependent chain of service provision. This Open Access Pivot explores these challenges and offers a solution for the intelligent and reliable management of physical infrastructure and the optimal placement of applications for the provision of services on distributed clouds. This book provides a conceptual reference model for reliable capacity provisioning for distributed clouds and discusses how data analytics and machine learning, application and infrastructure optimization, and simulation can deliver quality of service requirements cost-efficiently in this complex feature space. These are illustrated through a series of case studies in cloud computing, telecommunications, big data analytics, and smart cities.

Elasticsearch for Hadoop Vishal Shukla 2015-10-27 Integrate Elasticsearch into Hadoop to effectively visualize and analyze your

data About This Book Build production-ready analytics applications by integrating the Hadoop ecosystem with Elasticsearch Learn complex Elasticsearch queries and develop real-time monitoring Kibana dashboards to visualize your data Use Elasticsearch and Kibana to search data in Hadoop easily with this comprehensive, step-by-step guide Who This Book Is For This book is targeted at Java developers with basic knowledge on Hadoop. No prior Elasticsearch experience is expected. What You Will Learn Set up the Elasticsearch-Hadoop environment Import HDFS data into Elasticsearch with MapReduce jobs Perform full-text search and aggregations efficiently using Elasticsearch Visualize data and create interactive dashboards using Kibana Check and detect anomalies in streaming data using Storm and Elasticsearch Inject and classify real-time streaming data into Elasticsearch Get production-ready for Elasticsearch-Hadoop based projects Integrate with Hadoop eco-system such as Pig, Storm, Hive, and Spark In Detail The Hadoop ecosystem is a de-facto standard for processing terra-bytes and peta-bytes of data. Lucene-enabled Elasticsearch is becoming an industry standard for its full-text search and aggregation capabilities. Elasticsearch-Hadoop serves as a perfect tool to bridge the worlds of Elasticsearch and Hadoop ecosystem to get best out of both the worlds. Powered with Kibana, this stack makes it a cakewalk to get surprising insights out of your massive amount of Hadoop ecosystem in a flash. In this book, you'll learn to use Elasticsearch, Kibana and Elasticsearch-Hadoop effectively to analyze and understand your HDFS and streaming data. You begin with an in-depth understanding of the Hadoop, Elasticsearch, Marvel, and Kibana setup. Right after this, you will learn to successfully import Hadoop data into Elasticsearch by writing MapReduce job in a real-world example. This is then followed by a comprehensive look at Elasticsearch essentials, such as full-text search analysis, queries, filters and aggregations; after which you gain an understanding of creating various visualizations and interactive dashboard using

Kibana. Classifying your real-world streaming data and identifying trends in it using Storm and Elasticsearch are some of the other topics that we'll cover. You will also gain an insight about key concepts of Elasticsearch and Elasticsearch-hadoop in distributed mode, advanced configurations along with some common configuration presets you may need for your production deployments. You will have "Go production checklist" and high-level view for cluster administration for post-production. Towards the end, you will learn to integrate Elasticsearch with other Hadoop eco-system tools, such as Pig, Hive and Spark. Style and approach A concise yet comprehensive approach has been adopted with real-time examples to help you grasp the concepts easily.

Real Time Computing Alexander D. Stoyenko 2013-12-14 NATO's Division of Scientific and Environmental Affairs sponsored this Advanced Study Institute because it was felt to be timely to cover this important and challenging subject for the first time in the framework of NATO's ASI programme. The significance of real-time systems in everyone's life is rapidly growing. The vast spectrum of these systems can be characterised by just a few examples of increasing complexity: controllers in washing machines, air traffic control systems, control and safety systems of nuclear power plants and, finally, future military systems like the Strategic Defense Initiative (SDI). The importance of such systems for the well-being of people requires considerable efforts in research and development of highly reliable real-time systems. Furthermore, the competitiveness and prosperity of entire nations now depend on the early application and efficient utilisation of computer integrated manufacturing systems (CIM), of which real-time systems are an essential and decisive part. Owing to its key significance in computerised defence systems, real-time computing has also a special importance for the Alliance. The early research and development activities in this field in the 1960s and 1970s aimed towards improving the then unsatisfactory

software situation. Thus, the first high-level real-time languages were defined and developed: RTL/2, Coral 66, Procol, LTR, and PEARL. In close connection with these language developments and with the utilisation of special purpose process control peripherals, the research on real-time operating systems advanced considerably.

Flying Ad Hoc Networks Jingjing Wang 2022-02-10 Relying on unmanned autonomous flight control programs, unmanned aerial vehicles (UAVs) equipped with radio communication devices have been actively developed around the world. Given their low cost, flexible maneuvering and unmanned operation, UAVs have been widely used in both civilian operations and military missions, including environmental monitoring, emergency communications, express distribution, even military surveillance and attacks, for example. Given that a range of standards and protocols used in terrestrial wireless networks are not applicable to UAV networks, and that some practical constraints such as battery power and no-fly zone hinder the maneuverability capability of a single UAV, we need to explore advanced communication and networking theories and methods for the sake of supporting future ultra-reliable and low-latency applications. Typically, the full potential of UAV network's functionalities can be tapped with the aid of the cooperation of multiple drones relying on their ad hoc networking, in-network communications and coordinated control. Furthermore, some swarm intelligence models and algorithms conceived for dynamic negotiation, path programming, formation flight and task assignment of multiple cooperative drones are also beneficial in terms of extending UAV's functionalities and coverage, as well as of increasing their efficiency. We call the networking and cooperation of multiple drones as the terminology 'flying ad hoc network (FANET)', and there indeed are numerous new challenges to be overcome before the widespread of so-called heterogeneous FANETs. In this book, we examine a range of technical issues in FANETs, from physical-layer channel modeling to MAC-layer

resource allocation, while also introducing readers to UAV aided mobile edge computing techniques.

Optimization Problems in Self-Organizing Networks Steffen Wolf 2010 Modern computer networks or wireless ad-hoc networks offer a wide range of interesting optimization problems. Usual optimization goals are the minimization of the message delay in a Peer-to-Peer system or the minimization of the energy consumption of a wireless network. This thesis presents different kinds of algorithms to solve such optimization problems. Starting from the mathematical formulations for these problems, various global view optimization algorithms are presented. These algorithms are based on evolutionary algorithms and local search or similar heuristics. They can be used to quickly find near-optimal solutions, if a global view of the network is possible. As the participants in a computer network or a wireless ad-hoc network are autonomous nodes, distributed algorithms can be designed that enable these nodes to collectively solve the optimization problem. Four distributed algorithms are formulated and evaluated in this thesis, thus laying grounds for distributed optimization of networks. Using these algorithms, the network can be modelled as a self-optimizing network and the optimization problem can be approached without global view.

Kubernetes - A Complete DevOps Cookbook Murat Karslioglu 2020-03-13 Leverage Kubernetes and container architecture to successfully run production-ready workloads Key FeaturesImplement Kubernetes to orchestrate and scale applications proficientlyLeverage the latest features of Kubernetes to resolve common as well as complex problems in a cloud-native environmentGain hands-on experience in securing, monitoring, and troubleshooting your applicationBook Description Kubernetes is a popular open source orchestration platform for managing containers in a cluster environment. With this Kubernetes cookbook, you'll learn how to implement Kubernetes using a recipe-based approach. The book will prepare you to create highly

available Kubernetes clusters on multiple clouds such as Amazon Web Services (AWS), Google Cloud Platform (GCP), Azure, Alibaba, and on-premises data centers. Starting with recipes for installing and configuring Kubernetes instances, you'll discover how to work with Kubernetes clients, services, and key metadata. You'll then learn how to build continuous integration/continuous delivery (CI/CD) pipelines for your applications, and understand various methods to manage containers. As you advance, you'll delve into Kubernetes' integration with Docker and Jenkins, and even perform a batch process and configure data volumes. You'll get to grips with methods for scaling, security, monitoring, logging, and troubleshooting. Additionally, this book will take you through the latest updates in Kubernetes, including volume snapshots, creating high availability clusters with kops, running workload operators, new inclusions around kubectl and more. By the end of this book, you'll have developed the skills required to implement Kubernetes in production and manage containers proficiently. What you will learnDeploy cloud-native applications on KubernetesAutomate testing in the DevOps workflowDiscover and troubleshoot common storage issuesDynamically scale containerized services to manage fluctuating traffic needsUnderstand how to monitor your containerized DevOps environmentBuild DevSecOps into CI/CD pipelinesWho this book is for This Kubernetes book is for developers, IT professionals, and DevOps engineers and teams who want to use Kubernetes to manage, scale, and orchestrate applications in their organization. Basic understanding of Kubernetes and containerization is necessary.

Heuristic Search Stefan Edelkamp 2011-05-31 Search has been vital to artificial intelligence from the very beginning as a core technique in problem solving. The authors present a thorough overview of heuristic search with a balance of discussion between theoretical analysis and efficient implementation and application to real-world problems. Current developments in search such as

pattern databases and search with efficient use of external memory and parallel processing units on main boards and graphics cards are detailed. Heuristic search as a problem solving tool is demonstrated in applications for puzzle solving, game playing, constraint satisfaction and machine learning. While no previous familiarity with heuristic search is necessary the reader should have a basic knowledge of algorithms, data structures, and calculus. Real-world case studies and chapter ending exercises help to create a full and realized picture of how search fits into the world of artificial intelligence and the one around us. Provides real-world success stories and case studies for heuristic search algorithms Includes many AI developments not yet covered in textbooks such as pattern databases, symbolic search, and parallel processing units

Photonic Slot Routing in Optical Transport Networks Gosse Wedzinga 2012-12-06 All-optical networking is generally believed to be the only solution for coping with the ever-increasing demands in bandwidth, such as the World Wide Web application. Optical backbone networks efficiently achieve a high level of traffic aggregation by multiplexing numerous users on circuit-switched wavelength paths - the so-called wavelength routing approach. In contrast, the reduced level of traffic aggregation in access and metro networks makes wavelength routing solutions not adequate. In these network areas, packet-interleaved optical time-division multiplexing with its finer and more dynamic bandwidth allocation is advocated. The book presents such an approach, known as photonic slot routing. It illustrates how this approach may provide a cost-effective solution to deploying all-optical transport networks, using today's optical device technology. To that end, the author combines DWDM-technology with fixed slot optical switching, and gives a comprehensive description of this approach in which slots are aligned across the wavelengths to form groups of data-flows that propagate as a whole inside the network. Operating algorithms are developed, and network performance is

analyzed, both by means of theoretical analysis and many simulations of sample networks. This work will be of particular interest to researchers and professionals who are active in photonic networking.

Learning Couchbase Henry Potsangbam 2015-11-23 Design documents and implement real world e-commerce applications with Couchbase About This Book Get acquainted with Couchbase architecture and design your document-based data schema Implement full text search using industry standard elastic search plugins Develop critical and high performance applications using this hands-on tutorial guide Who This Book Is For If you are new to the NoSQL document system or have little or no experience in NoSQL development and administration and are planning to deploy Couchbase for your next project, then this book is for you. It would be helpful to have a bit of familiarity with Java. What You Will Learn Get acquainted with the concept of NoSQL databases and configure your Couchbase database cluster Maintain Couchbase effectively using the web-based administrative console with ease Enable partition capabilities by making use of Buckets Analyze important design considerations for maintaining relationship between various documents Use Couchbase SDK Java API to store and retrieve document Write views using map/reduce to retrieve documents efficiently Get familiar with N1QL and how to use it in Java applications Integrate Couchbase with Elasticsearch to implement full text search Configure XDCR for disaster recovery and develop ecommerce application using Couchbase In Detail NoSQL database systems have changed application development in terms of adaptability to dynamics schema and scalability. Compared with the currently available NoSQL database systems, Couchbase is the fastest. Its ease of configuration and powerful features for storing different schema structures, retrieval using map reduce and inbuilt disaster recovery by replicating document across the geographical region, make it one of the most powerful, scalable and comprehensive

NoSQL in the market. Couchbase also introduces smart client API for various programming language to integrate the database with the application easily, yet providing very complex features like cluster health awareness. This book achieves its goal by taking up an end-to-end development structure, right from understanding NOSQL document design to implementing full fledged eCommerce application design using Couchbase as a backend. Starting with the architecture of Couchbase to get you up and running, this book quickly takes you through designing a NoSQL document and implementing highly scalable applications using Java API. You will then be introduced to document design and get to know the various ways to administer Couchbase. Followed by this, learn to store documents using bucket. Moving on, you will then learn to store, retrieve and delete documents using smart client base on Java API. You will then retrieve documents using SQL like syntax call N1QL. Next, you will learn how to write map reduce base views. Finally, you will configure XDCR for disaster recovery and implement an eCommerce application using Couchbase. Style and approach The book starts from absolute basics and slowly moves to more advanced topics ensuring at every step that all concepts and terms are understood by the reader to have complete understanding at every stage. Technical and complex terms are explained in clear and simple language, thus making this book a perfect companion for those who have started their journey to NoSQL using Couchbase

Learning Elasticsearch Abhishek Andhavarapu 2017-06-30 Store, search, and analyze your data with ease using Elasticsearch 5.x About This Book Get to grips with the basics of Elasticsearch concepts and its APIs, and use them to create efficient applications Create large-scale Elasticsearch clusters and perform analytics using aggregation This comprehensive guide will get you up and running with Elasticsearch 5.x in no time Who This Book Is For If you want to build efficient search and analytics applications using Elasticsearch, this book is for you. It will also benefit developers

who have worked with Lucene or Solr before and now want to work with Elasticsearch. No previous knowledge of Elasticsearch is expected. What You Will Learn See how to set up and configure Elasticsearch and Kibana Know how to ingest structured and unstructured data using Elasticsearch Understand how a search engine works and the concepts of relevance and scoring Find out how to query Elasticsearch with a high degree of performance and scalability Improve the user experience by using autocomplete, geolocation queries, and much more See how to slice and dice your data using Elasticsearch aggregations. Grasp how to use Kibana to explore and visualize your data Know how to host on Elastic Cloud and how to use the latest X-Pack features such as Graph and Alerting In Detail Elasticsearch is a modern, fast, distributed, scalable, fault tolerant, and open source search and analytics engine. You can use Elasticsearch for small or large applications with billions of documents. It is built to scale horizontally and can handle both structured and unstructured data. Packed with easy-to- follow examples, this book will ensure you will have a firm understanding of the basics of Elasticsearch and know how to utilize its capabilities efficiently. You will install and set up Elasticsearch and Kibana, and handle documents using the Distributed Document Store. You will see how to query, search, and index your data, and perform aggregation-based analytics with ease. You will see how to use Kibana to explore and visualize your data. Further on, you will learn to handle document relationships, work with geospatial data, and much more, with this easy-to-follow guide. Finally, you will see how you can set up and scale your Elasticsearch clusters in production environments. Style and approach This comprehensive guide will get you started with Elasticsearch 5.x, so you build a solid understanding of the basics. Every topic is explained in depth and is supplemented with practical examples to enhance your understanding.

Elasticsearch 7.0 Cookbook Alberto Paro 2019-04-30 Search, analyze, and manage data effectively with Elasticsearch 7 Key

FeaturesExtend Elasticsearch functionalities and learn how to deploy on Elastic CloudDeploy and manage simple Elasticsearch nodes as well as complex cluster topologiesExplore the capabilities of Elasticsearch 7 with easy-to-follow recipesBook Description Elasticsearch is a Lucene-based distributed search server that allows users to index and search unstructured content with petabytes of data. With this book, you'll be guided through comprehensive recipes on what's new in Elasticsearch 7, and see how to create and run complex queries and analytics. Packed with recipes on performing index mapping, aggregation, and scripting using Elasticsearch, this fourth edition of Elasticsearch Cookbook will get you acquainted with numerous solutions and quick techniques for performing both every day and uncommon tasks such as deploying Elasticsearch nodes, integrating other tools to Elasticsearch, and creating different visualizations. You will install Kibana to monitor a cluster and also extend it using a variety of plugins. Finally, you will integrate your Java, Scala, Python, and big data applications such as Apache Spark and Pig with Elasticsearch, and create efficient data applications powered by enhanced functionalities and custom plugins. By the end of this book, you will have gained in-depth knowledge of implementing Elasticsearch architecture, and you'll be able to manage, search, and store data efficiently and effectively using Elasticsearch. What you will learnCreate an efficient architecture with ElasticsearchOptimize search results by executing analytics aggregationsBuild complex queries by managing indices and documentsMonitor the performance of your cluster and nodesDesign advanced mapping to take full control of index stepsIntegrate Elasticsearch in Java, Scala, Python, and big data applicationsInstall Kibana to monitor clusters and extend it for pluginsWho this book is for If you're a software engineer, big data infrastructure engineer, or Elasticsearch developer, you'll find this book useful. This Elasticsearch book will also help data professionals working in the e-commerce and FMCG industry who

use Elastic for metrics evaluation and search analytics to get deeper insights for better business decisions. Prior experience with Elasticsearch will help you get the most out of this book.

Security and Privacy in New Computing Environments

Wenbo Shi

Functional and Logic Programming Aart Middeldorp

1999-10-27 This volume contains the papers presented at the 4th Fuji International Symposium on Functional and Logic Programming (FLOPS'99) held in Tsukuba, Japan, November 11-13, 1999, and hosted by the Electrotechnical Laboratory (ETL). FLOPS is a forum for presenting and discussing all issues concerning functional programming, logic programming, and their integration. The symposium takes place about every 1.5 years in Japan. Previous FLOPS meetings were held in Fuji Susuno (1995), Shonan Village (1996), and Kyoto (1998). 1 There were 51 submissions from Austria (), Belgium (2), Brazil(3), China 3 3 1 7 (1), Denmark (2), France (3), Germany (8), Ireland (1), Israel (), Italy (1), 4 3 12 1 Japan (9), Korea (1), Morocco (1), The Netherlands (1), New Zealand (1), 3 1 1 3 5 Portugal (), Singapore (), Slovakia (1), Spain (4), Sweden (1), UK (4), 2 3 4 6 1 and USA (2), of which the program committee selected 21 for presentation. In 4 addition, this volume contains full papers by the two invited speakers, Atsushi Ohori and Mario Rodr'iguez-Artalejo.

Case-Based Reasoning Research and Development Rosina O.

Weber 2007-07-27 The International Conference on Case-Based Reasoning (ICCBR) is the pre-eminent international meeting on case-based reasoning (CBR). ICCBR 2007 (<http://www.iccbr.org/iccbr07/>) was the seventh in this series, presenting the most significant contributions in the field of CBR. The conference took place in Belfast, Northern Ireland, UK, during August 13-16, 2007. ICCBR and its sister conferences ECCBR (European Conference on Case-Based Reasoning) alternate every year. ICCBR 2007 followed a series of six successful international conferences previously held in Sesimbra, Portugal (1995);

Providence, Rhode Island, USA (1997); Seeon, Germany (1999); Vancouver, Canada (2001); Trondheim, Norway (2003); and Chicago, Illinois, USA (2005). The European Conferences on Case-Based Reasoning (ECCBR) were held as European workshops in Kaiserslautern, Germany (1993); Chantilly, France (1994); Lausanne, Switzerland (1996); Dublin, Ireland (1998); and Trento, Italy (2000); and as European conferences in Aberdeen, UK (2002); Madrid, Spain (2004); and Lykia World, Turkey (2006). Days one, two, and four comprised presentations and posters on theoretical and applied CBR research. In order to emphasize the importance of applications, the traditional industry day was converted into an Industry Program held on the second day, in the middle of the conference. Day three was devoted to five workshops: Case-Based Reasoning and Context-Awareness; Case-Based Reasoning in the Health Sciences; Textual Case-Based Reasoning: Beyond Retrieval; Uncertainty and Fuzziness in Case-Based Reasoning; and Knowledge Discovery and Similarity.

PRIMA 2013: Principles and Practice of Multi-Agent

Systems Guido Boella 2013-11-19 This book constitutes the refereed proceedings of the 16th International Conference on Principles and Practice of Multi-Agent Systems, PRIMA 2013, held in Dunedin, New Zealand, in December 2013. The conference was co-located with the 26th Australasian Artificial Intelligence Conference, AI 2013. The 24 revised full papers presented together with 18 short papers and 2 invited papers were carefully reviewed and selected from 81 submissions. The papers are organized in topical sections on foundations of agents and multi-agent systems; agent and multi-agent system architectures; agent-oriented software engineering; agent-based modelling and simulation; cooperation/collaboration, coordination/communication; hybrid technologies, application domains; and applications.

KI 2009: Advances in Artificial Intelligence Bärbel Mertsching 2009-09-29 The 32nd Annual German Conference on Artificial

Intelligence, KI 2009 (KI being the German acronym for AI), was held at the University of Paderborn, Germany on September 15-18, 2009, continuing a series of successful events. Starting back in 1975 as a national meeting, the conference now gathers researchers and developers from academic fields and industries worldwide to share their research results covering all aspects of artificial intelligence. This year we received submissions from 23 countries and 4 continents. Besides the international orientation, we made a major effort to include as many branches of AI as possible under the roof of the KI conference. A total of 21 area chairs representing different communities within the field of AI selected further members of the program committee and helped the local organizers to acquire papers. The new approach appealed to the AI community: we had 126 submissions, which constituted an increase of more than 50%, and which resulted in 14 parallel sessions on the following topics agents and intelligent virtual environments AI and engineering automated reasoning cognition evolutionary computation Robotics experience and knowledge management history and philosophical foundations knowledge representation and reasoning machine learning and mining natural language processing planning and scheduling spatial and temporal reasoning vision and perception offering cutting edge presentations and discussions with leading experts. Thirty-one percent of the contributions came from outside German-speaking countries.

Frontiers in Algorithmics Franco P. Preparata 2008-06-07 The Annual International Frontiers in Algorithmics Workshop is a focused forum on current trends in research on algorithms, discrete structures, and their applications. It intends to bring together international experts at the research frontiers in those areas to exchange ideas and to present significant new results. The mission of the workshop is to stimulate the various fields for which algorithmics can become a crucial enabler, and to strengthen the ties between the Eastern and Western algorithmics research

communities. The Second International Frontiers in Algorithmics Workshop (FAW 2008) took place in Changsha, China, June 19–21, 2008. In response to the Call for Papers, 80 papers were submitted from 15 countries and regions: Canada, China, France, Germany, Greece, Hong Kong, India, Iran, Japan, Mexico, Norway, Singapore, South Korea, Taiwan, and the USA. After a six-week period of careful reviewing and discussion, the Program Committee accepted 32 submissions for presentation at the conference. These papers were selected for nine special focus tracks in the areas of biomedical informatics, discrete structures, geometric information processing and communication, games and incentive analysis, graph algorithms, internet algorithms and protocols, parameterized algorithms, design and analysis of heuristics, approximate and online algorithms, and machine learning. The program of FAW 2008 also included three keynote talks by Xiaotie Deng, John E. Hopcroft, and Milan Sonka.

Selected Writings on Computing: A personal Perspective Edsger W. Dijkstra 2012-12-06 Since the summer of 1973, when I became a Burroughs Research Fellow, my life has been very different from what it had been before. The daily routine changed: instead of going to the University each day, where I used to spend most of my time in the company of others, I now went there only one day a week and was most of the time -that is, when not travelling!- alone in my study. In my solitude, mail and the written word in general became more and more important. The circumstance that my employer and I had the Atlantic Ocean between us was a further incentive to keep a fairly complete record of what I was doing. The public part of that output found its place in what became known as "the EWD series", which can be viewed as a form of scientific correspondence, possible since the advent of the copier. (That same copier makes it hard to estimate its actual distribution: I myself made about two dozen copies of my texts, but their recipients were welcome to act as further nodes of the distribution tree.) The decision to publish a selection from the

EWD series in book form was at first highly embarrassing, but as the months went by I got used to the idea. As soon as some guiding principles had been adopted -preferably not published elsewhere, as varied and as representative as possible, etc. Fog/Edge Computing For Security, Privacy, and Applications Wei Chang 2021-01-04 This book provides the state-of-the-art development on security and privacy for fog/edge computing, together with their system architectural support and applications. This book is organized into five parts with a total of 15 chapters. Each area corresponds to an important snapshot. The first part of this book presents an overview of fog/edge computing, focusing on its relationship with cloud technology and the future with the use of 5G communication. Several applications of edge computing are discussed. The second part of this book considers several security issues in fog/edge computing, including the secure storage and search services, collaborative intrusion detection method on IoT-fog computing, and the feasibility of deploying Byzantine agreement protocols in untrusted environments. The third part of this book studies the privacy issues in fog/edge computing. It first investigates the unique privacy challenges in fog/edge computing, and then discusses a privacy-preserving framework for the edge-based video analysis, a popular machine learning application on fog/edge. This book also covers the security architectural design of fog/edge computing, including a comprehensive overview of vulnerabilities in fog/edge computing within multiple architectural levels, the security and intelligent management, the implementation of network-function-virtualization-enabled multicasting in part four. It explains how to use the blockchain to realize security services. The last part of this book surveys applications of fog/edge computing, including the fog/edge computing in Industrial IoT, edge-based augmented reality, data streaming in fog/edge computing, and the blockchain-based application for edge-IoT. This book is designed for academics, researchers and government officials, working in the

field of fog/edge computing and cloud computing. Practitioners, and business organizations (e.g., executives, system designers, and marketing professionals), who conduct teaching, research, decision making, and designing fog/edge technology will also benefit from this book. The content of this book will be particularly useful for advanced-level students studying computer science, computer technology, and information systems, but also applies to students in business, education, and economics, who would benefit from the information, models, and case studies therein.

Advances in Information Retrieval Matthias Hagen

Algorithms - ESA 2000 Mike Paterson 2003-07-31 This book constitutes the refereed proceedings of the 8th Annual European Symposium on Algorithms, ESA 2000, held in Saarbrücken, Germany in September 2000. The 39 revised full papers presented together with two invited papers were carefully reviewed and selected for inclusion in the book. Among the topics addressed are parallelism, distributed systems, approximation, combinatorial optimization, computational biology, computational geometry, external-memory algorithms, graph algorithms, network algorithms, online algorithms, data compression, symbolic computation, pattern matching, and randomized algorithms.

Frontiers in Algorithmics Xiaotie Deng 2009-06-08 The Third International Frontiers of Algorithmics Workshop (FAW 2009), held during June 20–23, 2009 at Hefei University of Technology, Hefei, Anhui, China, continued to provide a focused forum on current trends in research on algorithmics, including discrete structures, and their applications. We aim at stimulating the various fields for which algorithmics can become a crucial enabler, and to strengthen the ties between the Eastern and Western algorithmics research communities as well as theory and practice of algorithmics. We had three distinguished invited speakers: Guoliang Chen, Andrew Chi-Chih Yao and Frances Foong Yao, speaking on parallel computing, communication complexity and applications, and computer and

network power management. The final program also included 33 peer-reviewed papers selected out of 87 contributed submissions, covering topics including approximation and online algorithms; computational geometry; graph theory and graph algorithms; games and applications; heuristics; large-scale data mining; machine learning; pattern recognition algorithms; and parameterized algorithms. April 2009 Xiaotie Deng John Hopcroft Jinyun Xue Organization FAW 2009 was organized by Hefei University of Technology, China.

Practical Enterprise Data Lake Insights Saurabh Gupta 2018-07-29 Use this practical guide to successfully handle the challenges encountered when designing an enterprise data lake and learn industry best practices to resolve issues. When designing an enterprise data lake you often hit a roadblock when you must leave the comfort of the relational world and learn the nuances of handling non-relational data. Starting from sourcing data into the Hadoop ecosystem, you will go through stages that can bring up tough questions such as data processing, data querying, and security. Concepts such as change data capture and data streaming are covered. The book takes an end-to-end solution approach in a data lake environment that includes data security, high availability, data processing, data streaming, and more. Each chapter includes application of a concept, code snippets, and use case demonstrations to provide you with a practical approach. You will learn the concept, scope, application, and starting point. What You'll Learn Get to know data lake architecture and design principles Implement data capture and streaming strategies Implement data processing strategies in Hadoop Understand the data lake security framework and availability model Who This Book Is For Big data architects and solution architects

Database and Expert Systems Applications ROLAND P AUTOR WAGNER 1996-08-28 Content Description #Includes bibliographical references and index.

Readings in Computer Vision Martin A. Fischler 2014-06-28 The field of computer vision combines techniques from physics, mathematics, psychology, artificial intelligence, and computer science to examine how machines might construct meaningful descriptions of their surrounding environment. The editors of this volume, prominent researchers and leaders of the SRI International AI Center Perception Group, have selected sixty papers, most published since 1980, with the viewpoint that computer vision is concerned with solving seven basic problems: Reconstructing 3D scenes from 2D images Decomposing images into their component parts Recognizing and assigning labels to scene objects Deducing and describing relations among scene objects Determining the nature of computer architectures that can support the visual function Representing abstractions in the world of computer memory Matching stored descriptions to image representation Each chapter of this volume addresses one of these problems through an introductory discussion, which identifies major ideas and summarizes approaches, and through reprints of key research papers. Two appendices on crucial assumptions in image interpretation and on parallel architectures for vision applications, a glossary of technical terms, and a comprehensive bibliography and index complete the volume.

EVOLVE - A Bridge between Probability, Set Oriented Numerics, and Evolutionary Computation VI Alexandru-Adrian Tantar 2017-11-09 This book comprises selected research papers from the 2015 edition of the EVOLVE conference, which was held on June 18–June 24, 2015 in Iași, Romania. It presents the latest research on Probability, Set Oriented Numerics, and Evolutionary Computation. The aim of the EVOLVE conference was to provide a bridge between probability, set oriented numerics and evolutionary computation and to bring together experts from these disciplines. The broad focus of the EVOLVE conference made it possible to discuss the connection between these related fields of study computational science. The selected papers published in

the proceedings book were peer reviewed by an international committee of reviewers (at least three reviews per paper) and were revised and enhanced by the authors after the conference. The contributions are categorized into five major parts, which are: Multicriteria and Set-Oriented Optimization; Evolution in ICT Security; Computational Game Theory; Theory on Evolutionary Computation; Applications of Evolutionary Algorithms. The 2015 edition shows a major progress in the aim to bring disciplines together and the research on a number of topics that have been discussed in previous editions of the conference matured over time and methods have found their ways in applications. In this sense the book can be considered an important milestone in bridging and thereby advancing state-of-the-art computational methods.

Stabilization, Safety, and Security of Distributed Systems

Borzoo Bonakdarpour 2016-11-01 This book constitutes the refereed proceedings of the 18th International Symposium on Stabilization, Safety, and Security of Distributed Systems, SSS 2016, held in Lyon, France, in November 2016. This year the Program Committee was organized into three groups reflecting the major trends related to self-* systems: (a) Self-* and Autonomic Computing, (b) Foundations, and (c) Networks, Multi-Agent Systems, and Mobility.

Reliability of Computer and Communication Networks Fred S. Roberts 1991-01-01 Reliability problems arise with increasing frequency as our modern systems of telecommunications, information transmission, transportation, and distribution become more and more complex. In December 1989 at DIMACS at Rutgers University, a Workshop on Reliability of Computer and Communications Networks was held to examine the discrete mathematical methods relevant to these problems. There were nearly ninety participants, including theoretical mathematicians, computer scientists, and electrical engineers from academia and industry, as well as network practitioners, engineers, and reliability

planners from leading companies involved in the use of computer and communications networks. This volume, published jointly with the Association for Computing Machinery, contains the proceedings from this Workshop. The aim of the Workshop was to identify the latest trends and important open problems, as well as to survey potential practical applications. The Workshop explored questions of computation of reliability of existing systems and of creating new designs to insure high reliability, in addition to the closely related notion of survivability. Redundancy, single stage and multistage networks, interconnected networks, and fault tolerance were also covered. The Workshop emphasized practical applications, with many invited speakers from a variety of companies which are dealing with practical network reliability problems. The success of the Workshop in fostering many new interactions among researchers and practitioners is reflected in the proceedings, which provide an exciting look at some of the major advances at the forefront of this important field of research.

Service-Oriented Computing E. Michael Maximilien 2011-03-04 This book constitutes the joint post-proceedings of four topical workshops held as satellite meetings of the 8th International Conference on service-oriented computing, ICSOC 2010, held in San Francisco, CA, USA in December 2010. The 23 revised papers presented together with four introductory descriptions are organized in topical sections corresponding to the individual workshops: performance assessment and auditing in service computing (PAASC 2010), engineering service-oriented applications (WESOA 2010), services, energy and ecosystems (SEE 2010), and service-oriented computing in logistics (SOC-LOG 2010)

Dissemination of Information in Optical Networks: Subir Bandyopadhyay 2007-10-24 This book offers a broad overview of techniques used in the design of Wavelength Division Multiplexing (WDM) networks for efficient dissemination of information in computer networks. It starts with an overview of the hardware components then provides a thorough review of WDM. Each topic

is covered rigorously with emphasis on detailed explanations of the approaches used. Numerous exercises are included.

Advanced Elasticsearch 7.0 Wai Tak Wong 2019-08-23 Master the intricacies of Elasticsearch 7.0 and use it to create flexible and scalable search solutions Key Features Master the latest distributed search and analytics capabilities of Elasticsearch 7.0 Perform searching, indexing, and aggregation of your data at scale Discover tips and techniques for speeding up your search query performance Book Description Building enterprise-grade distributed applications and executing systematic search operations call for a strong understanding of Elasticsearch and expertise in using its core APIs and latest features. This book will help you master the advanced functionalities of Elasticsearch and understand how you can develop a sophisticated, real-time search engine confidently. In addition to this, you'll also learn to run machine learning jobs in Elasticsearch to speed up routine tasks. You'll get started by learning to use Elasticsearch features on Hadoop and Spark and make search results faster, thereby improving the speed of query results and enhancing the customer experience. You'll then get up to speed with performing analytics by building a metrics pipeline, defining queries, and using Kibana for intuitive visualizations that help provide decision-makers with better insights. The book will later guide you through using Logstash with examples to collect, parse, and enrich logs before indexing them in Elasticsearch. By the end of this book, you will have comprehensive knowledge of advanced topics such as Apache Spark support, machine learning using Elasticsearch and scikit-learn, and real-time analytics, along with the expertise you need to increase business productivity, perform analytics, and get the very best out of Elasticsearch. What you will learn Pre-process documents before indexing in ingest pipelines Learn how to model your data in the real world Get to grips with using Elasticsearch for exploratory data analysis Understand how to build analytics and RESTful services Use Kibana, Logstash, and Beats for dashboard

applicationsGet up to speed with Spark and Elasticsearch for real-time analyticsExplore the basics of Spring Data Elasticsearch, and understand how to index, search, and query in a Spring applicationWho this book is for This book is for Elasticsearch developers and data engineers who want to take their basic knowledge of Elasticsearch to the next level and use it to build enterprise-grade distributed search applications. Prior experience of working with Elasticsearch will be useful to get the most out of this book.

Database Systems For Advanced Applications '97 - Proceedings Of The 5th International Conference On Database Systems For

Advanced Applications Rodney Topor 1997-03-15 This volume contains the proceedings of the Fifth International Conference on Database Systems for Advanced Applications (DASFAA '97). DASFAA '97 focused on advanced database technologies and their applications. The 55 papers in this volume cover a wide range of areas in the field of database systems and applications - including the rapidly emerging areas of the Internet, multimedia, and document database systems - and should be of great interest to all database system researchers and developers, and practitioners.