Distributed Systems Concepts Design 4th Edition Solution Manual

Provides a broad and up-to-date account of the principles and practice of distributed system design.

Distributed SystemsConcepts and DesignAddison-Wesley Longman

The new edition of this bestselling title on Distributed Systems has been thoroughly revised throughout to reflect the state of the art in this rapidly developing field. It emphasizes the principles used in the design and construction of distributed computer systems based on networks of workstations and server computers.

The two-volume set LNCS 5072 and 5073 constitutes the refereed proceedings of the International Conference on Computational Science and Its Applications, ICCSA 2008, held in Perugia, Italy, in June/July, 2008. The two volumes contain papers presenting a wealth of original research results in the field of computational science, from foundational issues in computer science and mathematics to advanced applications in virtually all sciences making use of computational techniques. The topics of the fully refereed papers are structured according to the five major conference themes: computational methods, algorithms and scientific applications, high performance technical computing and networks, advanced and emerging applications, geometric modelling, graphics and visualization, as well as information systems and information technologies. Moreover, submissions from more than 20 workshops and technical sessions in the areas, such as embedded systems, geographical analysis, computational geometry, computational geomatics, computer graphics, virtual reality, computer modeling, computer algebra, mobile communications, wireless networks, computational forensics, data storage, information security, web learning, software engineering, computational intelligence, digital security, biometrics, molecular structures, material design, ubiquitous computing, symbolic computations, web systems and intelligence, and e-education contribute to this publication.

CD-ROM with a simulation system and numerous solved models is attached to the book. Distributed systems are a continuously expanding area of computer science and computer engineering. This book addresses the need for literature on modeling and simulation techniques for distributed systems. For simulation modeling of distributed systems in the book, a specific class of extended Petri nets is used that allows to easily represent the fundamental processes of any distributed system. The book is intended, first of all, as a text for related graduate-level university courses on distributed systems in computer science and computer engineering. Other computer science and computer engineering courses would also find the book useful as a source of practical information for a broad community of those graduate students who are busy with simulation in their study and research. The book can be useful also to academics who give related graduate courses or deliver research-oriented modules for graduate students. Further, the book can be helpful to system architects and developers who apply modeling and simulation techniques as a step in the design and implementation of their systems. Containing a large number of models, with commented source texts and simulation results on the attached CD-ROM, it can also serve as valuable reference book for researchers who want to develop their own models in terms of Petri nets.

This book offers a collection of high-quality, peer-reviewed research papers presented at the International Conference on Intelligent Computing, Communication and Devices (ICCD 2017), discussing all dimensions of intelligent sciences – intelligent computing, intelligent communication, and intelligent devices. Intelligent computing addresses areas such as intelligent and distributed computing, intelligent grid and cloud computing, internet of things, soft computing and engineering applications, data mining and knowledge discovery, semantic and web technology, hybrid systems, agent computing, bioinformatics, and recommendation systems. Intelligent communication is concerned with communication and network technologies, such as mobile broadband and all optical networks that are the key to groundbreaking inventions of intelligent communication technologies. It includes communication hardware, software and networked intelligence, mobile technologies, machine-to-machine communication networks, speech and natural language processing, routing techniques and network analytics, wireless ad hoc and sensor networks, communications and information security, signal, image and video processing, network management, and traffic engineering. Lastly, intelligent devices are any equipment, instruments, or machines that have their own computing capability. As computing technology becomes more advanced and less expensive, it can be incorporated an increasing number of devices of all kinds. This area covers such as embedded systems, radiofrequency identification (RFID), radiofrequency microelectromechanical system (RF MEMS), very-large-scale integration (VLSI) design and electronic devices, analog and mixed-signal integrated circuit (IC) design and testing, microelectromechanical system (MEMS) and microsystems, solar cells and photonics, nanodevices, single electron and spintronics devices, space electronics, and intelligent robotics.

A fascinating bird’s eye view on a hugely relevant topic. This book constitutes the refereed proceedings of the 4th International Conference on Ubiquitous Intelligence and Computing held in Hong Kong, China in 2007, co-located with ATC 2007, the 4th International Conference on Autonomic and Trusted Computing. The 119 revised full papers presented together with 1 keynote paper and 1 invited paper were carefully reviewed and selected from 463 submissions. The papers are organized in topical sections.

The growing complexity of modern software systems increases the di?culty of ensuring the overall dependability of software-intensive systems. Complexity of environments, in which systems operate, high dependability requirements that systems have to meet, as well as the complexity of infrastructures on which they rely make system design a true engineering challenge. Mastering system complexity requires design techniques that support clear thinking and rigorous validation and veri?cation. Formal design methods help to achieve this. Coping with complexity also requires architectures that are t- erant of faults and of unpredictable changes in environment. This issue can be addressed by fault-tolerant design techniques. Therefore, there is a clear need of methods enabling rigorous modelling and development of complex fault-tolerant systems. This bookaddressessuchacuteissues indevelopingfault-tolerantsystemsas: – Veri?cation and re?nement of fault-tolerant systems – Integrated approaches to developing fault-tolerant systems – Formal foundations for error detection, error recovery, exception and fault handling – Abstractions, styles and patterns for rigorousdevelopment of fault tolerance – Fault-tolerant software architectures – Development and application of

tools supporting rigorous design of depe- able systems – Integrated platforms for developing dependable systems – Rigorous approaches to speci?cation and design of fault tolerance in novel computing systems
TheeditorsofthisbookwereinvolvedintheEU(FP-6)projectRODIN(R- orous Open Development Environment for Complex Systems), which brought together researchers from the fault tolerance and formal methods communi- 1 ties. In 2007 RODIN organized the MeMoT workshop held in conjunction with the Integrated Formal Methods 2007 Conference at Oxford University.
The LNCS series reports state-of-the-art results in computer science research, development, and education, at a high level and in both printed and electronic form. Enjoying tight cooperation with the R&D community, with numerous individuals, as well as with prestigious organizations and societies, LNCS has grown into the most comprehensive computer science research forum available. The scope of LNCS, including its subseries LNAI and LNBI, spans the whole range of computer science and information technology including interdisciplinary topics in a variety of application field.
Cloud computing is the latest market-oriented computing paradigm which brings software design and development into a new era characterized by “XaaS”, i.e. everything as a service. Cloud workflows, as typical software applications in the cloud, are composed of a set of partially ordered cloud software services to achieve specific goals. However, due to the low QoS (quality of service) nature of the cloud environment, the design of workflow systems in the cloud becomes a challenging issue for the delivery of high quality cloud workflow applications. To address such an issue, this book presents a systematic investigation to the three critical aspects for the design of a cloud workflow system, viz. system architecture, system functionality and quality of service. Specifically, the system architecture for a cloud workflow system is designed based on the general four-layer cloud architecture, viz. application layer, platform layer, unified resources layer and fabric layer. The system functionality for a cloud workflow system is designed based on the general workflow reference model but with significant extensions to accommodate software services in the cloud. The support of QoS is critical for the quality of cloud workflow applications. This book presents a generic framework to facilitate a unified design and development process for software components that deliver lifecycle support for different QoS requirements. While the general QoS requirements for cloud workflow applications can have many dimensions, this book mainly focuses on three of the most important ones, viz. performance, reliability and security. In this book, the architecture, functionality and QoS management of our SwinDeW-C prototype cloud workflow system are demonstrated in detail as a case study to evaluate our generic design for cloud workflow systems. To conclude, this book offers a general overview of cloud workflow systems and provides comprehensive introductions to the design of the system architecture, system functionality and QoS management.
This book constitutes the refereed proceedings of the Third International Conference on Grid and Pervasive Computing,

GPC 2008, held in Kunming, China, in May 2008. The 45 revised full papers presented together with 2 keynote lectures were carefully reviewed and selected from 184 submissions. The papers cover all current issues of grid and pervasive computing and focus on topics such as cluster computing, grid computing, high performance computing, network storage, peer-to-peer computing, pervasive computing, the Semantic Web and the Semantic Grid, and service-oriented computing.

WSC2008Chair's Welcome Message Dear Colleague, The World Soft Computing (WSC) conference is an annual international online conference on applied and theoretical soft computing technology. This WSC 2008 is the thirteenth conference in this series and it has been a great success. We received a lot of excellent paper submissions which were peer-reviewed by an international team of experts. Only60 papers out of111 submissions were selected for online publication. This assured a high quality standard for this online conference. The corresponding online statistics are a proof of the great world-wide interest in the WSC 2008 conference. The conference website had a total of33,367di?erent human user accessesfrom43 countries with around100 visitors every day,151 people signed up to WSC to discuss their scienti?c disciplines in our chat rooms and the forum. Also audio and slide presentations allowed a detailed discussion of the papers. The submissions and discussions showed that there is a wide range of soft computing applications to date. The topics covered by the conference range from applied to theoretical aspects of fuzzy, neuro-fuzzy and rough sets over to neural networks to single and multi-objective optimisation. Contributions aboutparticleswarmoptimisation,geneexpressionprogramming,clustering, classi?cation,supportvectormachines,quantumevolutionandagentsystems have also been received. One whole session was devoted to soft computing techniques in computer graphics, imaging, vision and signal processing.

Many applications follow the distributed computing paradigm, in which parts of the application are executed on different network-interconnected computers. The extension of these applications in terms of number of users or size has led to an unprecedented increase in the scale of the infrastructure that supports them. Large-Scale Distributed Computing and Applications: Models and Trends offers a coherent and realistic image of today's research results in large scale distributed systems, explains state-of-the-art technological solutions for the main issues regarding large scale distributed systems, and presents the benefits of using large scale distributed systems and the development process of scientific and commercial distributed applications.

In modern computing a program is usually distributed among several processes. The fundamental challenge when developing reliable and secure distributed programs is to support the cooperation of processes required to execute a common task, even when some of these processes fail. Failures may range from crashes to adversarial attacks by

malicious processes. Cachin, Guerraoui, and Rodrigues present an introductory description of fundamental distributed programming abstractions together with algorithms to implement them in distributed systems, where processes are subject to crashes and malicious attacks. The authors follow an incremental approach by first introducing basic abstractions in simple distributed environments, before moving to more sophisticated abstractions and more challenging environments. Each core chapter is devoted to one topic, covering reliable broadcast, shared memory, consensus, and extensions of consensus. For every topic, many exercises and their solutions enhance the understanding This book represents the second edition of "Introduction to Reliable Distributed Programming". Its scope has been extended to include security against malicious actions by non-cooperating processes. This important domain has become widely known under the name "Byzantine fault-tolerance".

This book covers all you need to know to model and design software applications from use cases to software architectures in UML and shows how to apply the COMET UML-based modeling and design method to real-world problems. The author describes architectural patterns for various architectures, such as broker, discovery, and transaction patterns for service-oriented architectures, and addresses software quality attributes including maintainability, modifiability, testability, traceability, scalability, reusability, performance, availability, and security. Complete case studies illustrate design issues for different software architectures: a banking system for client/server architecture, an online shopping system for service-oriented architecture, an emergency monitoring system for component-based software architecture, and an automated guided vehicle for real-time software architecture. Organized as an introduction followed by several short, self-contained chapters, the book is perfect for senior undergraduate or graduate courses in software engineering and design, and for experienced software engineers wanting a quick reference at each stage of the analysis, design, and development of large-scale software systems.

This book intends to inculcate the innovative ideas for the scheduling aspect in distributed computing systems. Although the models in this book have been designed for distributed systems, the same information is applicable for any type of system. The book will dramatically improve the design and management of the processes for industry professionals. It deals exclusively with the scheduling aspect, which finds little space in other distributed operating system books. Structured for a professional audience composed of researchers and practitioners in industry, this book is also suitable as a reference for graduate-level students.

Most applications in distributed computing center around a set of common subproblems. Distributed Systems: An Algorithmic Approach presents the algorithmic issues and necessary background theory that are needed to properly understand these challenges. Achieving a balance between theory and practice, this book bridges the gap between theoreticians and practitioners. With a set of exercises featured in

each chapter, the book begins with background information that contains various interprocess communication techniques and middleware services, followed by foundational topics that cover system models, correctness criteria, and proof techniques. The book also presents numerous important paradigms in distributed systems, including logical clocks, distributed snapshots, deadlock detection, termination detection, election, and several graph algorithms. The author then addresses failures and fault-tolerance techniques in diverse applications, such as consensus, transactions, group communication, replicated data management, and self-stabilization. He concludes with an exploration of real-world issues, including distributed discrete-event simulation and security, sensor networks, and peer-to-peer networks. By covering foundational matters of distributed systems and their relationships to real-world applications, Distributed Systems provides insight into common distributed computing subproblems,

The field of health informatics (or medical informatics as it is sometimes called) is still a relatively young one compared to other areas of biomedicine and the health sciences. Nevertheless, its impact on the quality and efficiency of healthcare is crucial. This second, extensively revised and updated edition of Health Informatics: An Overview includes new topics which address contemporary issues and challenges and shift the focus on the health problem space towards a computer perspective. An overview is provided of the health informatics discipline and the book is suitable for use as a basic text in both undergraduate and postgraduate curricula. Preparing students for practice as health professionals in any discipline, it deliberately avoids focusing on any one speciality. The publication is divided into six sections: an overview, basic concepts, applications supporting clinical practice, service delivery, management and clinical research and education. With contributions from many distinguished authors, this book is a valuable resource for healthcare professionals and students of health informatics alike.

The success of the World Wide Web depends on the ability of users to store, p- cess and retrieve digital information regardless of distance boundaries, languages and domains of knowledge. The universality and ?exibility of the World Wide Web have also enabled the rapid growth of a variety of new services and applications based on human–machine interaction. The semantics of exchanged information and services should be useful not only for human to human communications, but also in that machines would be able to understand and automatically process web content. Semanticsgive well-de?nedmeaningto web content and enable computersand p- ple to work in cooperation. Today, the crucial challenge becomes the development of languages to express information in a machine processable format. Now more than ever, new advanced techniques and intelligent approaches are required to tra- form the Web into a universal reasoning and computing machine. Web intelligence attempts to deal with this challenge by exploiting information technologies and - ti?cial intelligence approaches to design the next generation of web-empowered systems and services.

This book gathers selected papers presented at the conference of the Forum for Interdisciplinary Mathematics (FIM), held at Palau Macaya, Barcelona, on 18 to 20 November, 2015. The event was co-organized by the University of Barcelona (Spain), the Spanish Royal Academy of Economic and Financial Sciences (Spain) and the Forum for Interdisciplinary Mathematics (India). This instalment of the conference was presented with the title “Applied Mathematics and Computational Intelligence” and particularly focused on the use of Mathematics and Computational Intelligence techniques in a diverse range of scientific disciplines, as well as their applications in real-world problems. The book presents thirty peer-reviewed research papers, organised into four topical sections: on Mathematical Foundations; Computational Intelligence and Optimization Techniques; Modelling and Simulation Techniques; and Applications in Business and Engineering. This book will be of great interest to anyone working in the area of applied mathematics and computational intelligence and will be especially useful for

scientists and graduate students pursuing research in these fields.

This book constitutes the thoroughly refereed post-conference proceedings of the Second International ICST Conference on Sensor Systems and Software, S-Cube 2010, held in Miami, Florida, USA, in December 2010. The 17 revised full papers presented were carefully reviewed and selected and cover a wide range of topics including sensor application programming paradigms, novel sensor applications, sensor network middleware, trust security and privacy, wireless sensor network management and monitoring, and sensor application development support systems.

In recent years, the field of Universal Access has made significant progress in consolidating theoretical approaches, scientific methods and technologies, as well as in exploring new application domains. Increasingly, professionals in this rapidly maturing area require a comprehensive and multidisciplinary resource that addresses current principles, methods, and tools. Written by leading international authorities from academic, research, and industrial organizations and nonmarket institutions, The Universal Access Handbook covers the unfolding scientific, methodological, technological, and policy issues involved in the process of achieving universal access in the information society. In a collection of 61 chapters, the book discusses how to systematically apply universal design principles to information technologies. It explains the various dimensions of diversity in the technological platforms and contexts of use, including trends in mobile interaction and ambient intelligence environments. The implications of Universal Access on the development life cycle of interactive applications and services are unfolded, addressing user interface architectures and related components. Novel interaction methods and techniques for Universal Access are analyzed, and a variety of applications in diverse domains are discussed. The book reflects recent developments, consolidates present knowledge, and points towards new perspectives for the future. A quick glance through the contents demonstrates not only the breadth and depth of coverage but also the caliber of the contributions. An indispensible source of information for interdisciplinary and cross-thematic study, the book provides a baseline for further in-depth studies, as well as an important educational tool in an increasingly globalized research and development environment.

Global information retrieval and anywhere, anytime information access has stimulated a need to design and model the personalized information search in a flexible and agile way that can use the specific personalization techniques, algorithms, and available technology infrastructure to satisfy high-level functional requirements for personalization. Personalized Information Retrieval and Access: Concepts, Methods and Practices surveys the main concepts, methods, and practices of personalized information retrieval and access in today's data intensive, dynamic, and distributed environment, and provides students, researchers, and practitioners with authoritative coverage of recent technological advances that are shaping the future of globally distributed information retrieval and anywhere, anytime information access.

This book constitutes the refereed proceedings of the ACM/IFIP/USENIX 8th International Middleware Conference 2007, held in Newport Beach, CA, USA, in November 2007. The 22 revised full papers presented were carefully reviewed and selected from 108 submissions. The papers are organized in topical sections on component-based middleware, mobile and ubiquitous computing, grid and cluster computing, enhancing communication, resource management, reliability and fault tolerance.

This book constitutes the thoroughly refereed proceedings of the 4th International Conference on Mobile Wireless Middleware, Operating Systems, and Applications, Mobilware 2011, held in London, UK, in June 2011. The 21 revised full papers presented were carefully reviewed and selected from numerous contributions. The papers are organized in

topical sections on mobile systems in education, SOC for mobile Apps (SOC), networking platforms (NW), mobile execution frameworks (MFW), mobile cloud (MC) and distributed execution, and mobile sensor networks.

This book constitutes the symposia and workshops of the 10th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP. Each of the sympois and workshops focuses on a particular theme and complements the spectrum of the main conference.

This book constitutes the refereed proceedings of the 8th International Conference, ICSOC 2010, held in San Francisco, CA, USA, in December 2010. The 33 revised full papers and three full industry papers, presented together with 18 short papers, three PhD symposium posters and four regular posters, were carefully reviewed and selected from 238 submissions. The papers are organized in topical sections on Service and Business Process Modeling; Service Management; Quality of Service; Service Science and Design; Service Development and Run-time Management; High-level Description Languages; Service Level Agreements; Service Engineering Methodologies; Service Security, Privacy, and Trust; Business Service Modeling; Formal Methods; and Service Applications.

This book constitutes the proceedings of the second International Workshop on Advanced Computational Intelligence (IWACI 2009), with a sequel of IWACI 2008 successfully held in Macao, China. IWACI 2009 provided a high-level international forum for scientists, engineers, and educators to present state-of-the-art research in computational intelligence and related fields. Over the past decades, computational intelligence community has witnessed t- mendous efforts and developments in all aspects of theoretical foundations, archit- tures and network organizations, modelling and simulation, empirical study, as well as a wide range of applications across different domains. IWACI 2009 provided a great platform for the community to share their latest research results, discuss critical future research directions, stimulate innovative research ideas, as well as facilitate inter- tional multidisciplinary collaborations. IWACI 2009 received 146 submissions from about 373 authors in 26 countries and regions (Australia, Brazil, Canada, China, Chile, Hong Kong, India, Islamic Republic of Iran, Japan, Jordan, Macao, Malaysia, Mexico, Pakistan, Philippines, Qatar, Republic of Korea, Singapore, South Africa, Sri Lanka, Spain, Taiwan, Thailand, UK, USA, Ve- zuela, Vietnam, and Yemen) across six continents (Asia, Europe, North America, South America, Africa, and Oceania). Based on the rigorous peer reviews by the Program Committee members, 52 high-quality papers were selected for publication in this book, with an acceptance rate of 36.3%. These papers cover major topics of the theoretical research, empirical study, and applications of computational intelligence.

Research in multi-agent systems offers a promising technology for problems with networks, online trading and negotiations but also social structures and communication. This is a book on agent and multi-agent technology for

internet and enterprise systems. The book is a pioneer in the combination of the fields and is based on the concept of developing a platform to share ideas and presents research in technology in the field and application to real problems. The chapters range over both applications, illustrating the possible uses of agents in an enterprise domain, and design and analytic methods, needed to provide the solid foundation required for practical systems.

"This book explores the latest empirical research and best real-world practices for preventing, weathering, and recovering from disasters such as earthquakes or tsunamis to nuclear disasters and cyber terrorism"--Provided by publisher.

This book constitutes the refereed proceedings of the first International Conference on Internet of Vehicles, IOV 2014, held in Beijing, China, in September 2014. The 41 full papers presented were carefully reviewed and selected from 160 submissions. They focus on the following topics: IOV systems and applications; wireless communications, ad-hoc and sensor networks; security, privacy, IoT and big data intelligence; cloud and services computing.

This tutorial reference takes the reader from use cases to complete architectures for real-time embedded systems using SysML, UML, and MARTE and shows how to apply the COMET/RTE design method to real-world problems. The author covers key topics such as architectural patterns for distributed and hierarchical real-time control and other real-time software architectures, performance analysis of real-time designs using real-time scheduling, and timing analysis on single and multiple processor systems. Complete case studies illustrating design issues include a light rail control system, a microwave oven control system, and an automated highway toll system. Organized as an introduction followed by several self-contained chapters, the book is perfect for experienced software engineers wanting a quick reference at each stage of the analysis, design, and development of large-scale real-time embedded systems, as well as for advanced undergraduate or graduate courses in software engineering, computer engineering, and software design.

Video games represent a unique blend of programming, art, music, and unbridled creativity. To the general public, they are perhaps the most exciting computer applications ever undertaken. In the field of computer science, they have been the impetus for a continuous stream of innovations designed to provide gaming enthusiasts with the most realistic and enjoyable gaming experience possible. Algorithmic and Architectural Gaming Design: Implementation and Development discusses the most recent advances in the field of video game design, with particular emphasis on practical examples of game development, including design and implementation. The target audience of this book includes educators, students, practitioners, professionals, and researchers working in the area of video game design and development. Anyone actively developing video games will benefit from the practical application of fundamental computer science concepts demonstrated in this book.

??????????????????????,????????????????,??????????????

This book constitutes the thoroughly refereed post-conference proceedings of the IFIP WG 11.4 International Workshop on Open Problems in Network Security, iNetSec 2011, held in Lucerne, Switzerland, in June 2011, co-located and under the auspices of IFIP SEC 2011, the 26th IFIP TC-11 International Information Security Conference. The 12 revised full papers were carefully reviewed and selected from 28 initial submissions; they are fully revised to incorporate reviewers' comments and discussions at the workshop. The volume is organized in topical sections on assisting users, malware detection, saving energy, policies, and problems in the cloud.

"This book provides a general overview about research on ubiquitous and pervasive computing and its applications, discussing the recent progress in this area and pointing out to scholars what they should do (best practices) and should not do (bad practices)"--Provided by publisher.

For a large, complex system, the amount of test cases in a regression test suite can range from a few hundred to several thousands, which can take hours or even days to execute. Regression testing also requires considerable resources that are often not readily available. This precludes their use in an interactive setting, further contributing to an inefficient testing process. Cloud computing offers the use of virtualized hardware, effectively unlimited storage, and software services that can help reduce the execution time of large test suites in a cost-effective manner. The research presented by Tilley and Parveen leverages the resources provided by cloud computing infrastructure to facilitate the concurrent execution of test cases. They introduce a decision framework called SMART-T to support migration of software testing to the cloud, a distributed environment called HadoopUnit for the concurrent execution of test cases in the cloud, and a series of case studies illustrating the use of the framework and the environment. Experimental results indicate a significant reduction in test execution time is possible when compared with a typical sequential environment. Software testing in the cloud is a subject of high interest for advanced practitioners and academic researchers alike. For advanced practitioners, the issue of cloud computing and its impact on the field of software testing is becoming increasingly relevant. For academic researchers, this is a subject that is replete with interesting challenges; there are so many open problems that graduate students will be busy for years to come. To further disseminate results in this field, the authors created a community of interest called “Software Testing in the Cloud” (www.STITC.org), and they encourage all readers to get involved in this exciting new area.

Following from the very successful First KES Symposium on Agent and Multi-Agent Systems – Technologies and Applications (KES-AMSTA 2007), held in Wroclaw, Poland, 31 May–1 June 2007, the second event in the KES-AMSTA symposium series (KES-AMSTA 2008) was held in Incheon, Korea, March 26–28, 2008. The symposium was organized by the School of Computer and Information Engineering, Inha University, KES International and the KES Focus Group on

Agent and Mul- agent Systems. The KES-AMSTA Symposium Series is a sub-series of the KES Conference Series. The aim of the symposium was to provide an international forum for scientific research into the technologies and applications of agent and multi-agent systems. Agent and multi-agent systems are related to the modern software which has long been recognized as a promising technology for constructing autonomous, complex and intelligent systems. A key development in the field of agent and multi-agent systems has been the specification of agent communication languages and formalization of ontologies. Agent communication languages are intended to provide standard declarative mechanisms for agents to communicate knowledge and make requests of each other, whereas ontologies are intended for conceptualization of the knowledge domain. The symposium attracted a very large number of scientists and practitioners who submitted their papers for nine main tracks concerning the methodology and applications of agent and multi-agent systems, a doctoral track and two special sessions.

Timo Warns has developed tractable fault models that, while being non-probabilistic, are accurate for dependent and propagating faults. Using seminal problems such as consensus and constructing coteries, he demonstrates how the new models can be used to design and evaluate effective and efficient means of fault tolerance.