

## Advanced Database Systems Lecture Notes

This volume constitutes the proceedings of the 13th International Conference on the Entity-Relationship Approach, ER '94, held in Manchester, UK in December 1994. The ER '94 book is devoted to business modelling and re-engineering and provides a balanced view between research and practical experience. The 34 full revised papers presented are organized in sections on business process modelling, enterprise modelling, systems evolution, modelling integrity constraints, object-oriented databases, active databases, CASE, reverse engineering, information system modelling, schema coordination, and re-engineering.

This book constitutes the workshop proceedings of the 24th International Conference on Database Systems for Advanced Applications, DASFAA 2019, held in Chiang Mai, Thailand, in April 2019. The 14 full papers presented were carefully selected and reviewed from 26 submissions to the three following workshops: the 6th International Workshop on Big Data Management and Service, BDMS 2019; the 4th International Workshop on Big Data Quality Management, BDQM 2019; and the Third International Workshop on Graph Data Management and Analysis, GDMA 2019. This volume also includes the short papers, demo papers, and tutorial papers of the main conference DASFAA 2019.

Coverage in this proceedings includes XML schemas, data mining, spatial data, indexes and cubes, data streams, P2P and transactions, complex pattern processing, IR techniques, queries and transactions, XML databases, data warehouses, and distributed data.

This two volume set LNCS 7238 and LNCS 7239 constitutes the refereed proceedings of the 17th International Conference on Database Systems for Advanced Applications, DASFAA 2012, held in Busan, South Korea, in April 2012. The 44 revised full papers and 8 short papers presented together with 2 invited keynote papers, 8 industrial papers, 8 demo presentations, 4 tutorials and 1 panel paper were carefully reviewed and selected from a total of 159 submissions. The topics covered are query processing and optimization, data semantics, XML and semi-structured data, data mining and knowledge discovery, privacy and anonymity, data management in the Web, graphs and data mining applications, temporal and spatial data, top-k and skyline query processing, information retrieval and recommendation, indexing and search systems, cloud computing and scalability, memory-based query processing, semantic and decision support systems, social data, data mining.

This volume constitutes the proceedings of the First International Symposium organized by the Japan Society for Software Science and Technology. The symposium was held in Kanazawa, Japan, November 4-6, 1993 and attracted many researchers from academia and industry as well as ambitious practitioners. Object technologies, in particular object-oriented programming, object-oriented databases, and software object bases, currently attract much attention and hold a great promise of future research and development in diverse areas of advanced software. The volume contains besides 6 invited presentations by renown researchers and 25 contributed papers carefully selected by an international program committee from a total of 92 submissions.

This book presents recent research in intelligent information and database systems. The carefully selected contributions were initially accepted for presentation as posters at the 9th Asian Conference on Intelligent Information and Database Systems (ACIIDS 2017) held from to 5 April 2017 in Kanazawa, Japan. While the contributions are of an advanced scientific level, several are accessible for non-expert readers. The book brings together 47 chapters divided into six main parts: • Part I. From Machine Learning to Data Mining, • Part II. Big Data and Collaborative Decision Support Systems, • Part III. Computer Vision Analysis, Detection, Tracking and Recognition, • Part IV. Data-Intensive Text Processing, • Part V. Innovations in Web and Internet Technologies, and • Part VI. New Methods and Applications in Information and Software Engineering. The book is an excellent resource for researchers and those working in algorithmics, artificial and computational intelligence, collaborative systems, decision management and support systems, natural language processing, image and text processing, Internet technologies, and information and software engineering, as well as for students interested in such research areas.

Database management is attracting wide interest in both academic and industrial contexts. New application areas such as CAD/CAM, geographic information systems, and multimedia are emerging. The needs of these application areas are far more complex than those of conventional business applications. The purpose of this book is to bring together a set of current research issues that addresses a broad spectrum of topics related to database systems and applications. The book is divided into four parts: - object-oriented databases, - temporal/historical database systems, - query processing in database systems, - heterogeneity, interoperability, open system architectures, multimedia database systems.

The theme of this book is the potential of new advanced database systems. The volume presents the proceedings of the 10th British National Conference on Databases, held in Aberdeen, Scotland, in July 1992. The volume contains two invited papers, one on the promise of distributed computing and the challenges of legacy systems by M.L. Brodie, and the other on object-oriented requirements capture and analysis and the Orca project by D.J.L. Gradwell. The following four parts each contain three submitted papers selected from a total of 36 submissions. The parts are entitled: - Object-oriented databases - Parallel implementations and industrial systems - Non-relational data models - Logic programming and databases

Modern software systems are becoming more complex in many ways and have to cope with a growing number of abnormal situations which, in turn, are increasingly complex to handle. The most general way of dealing with these problems is by incorporating exception handling techniques in software design. In the past, various exception handling models and techniques have been proposed and many of them are part of practical languages and software composition technologies. This book is composed of five parts, which deal with topics related to exception handling in the context of programming language models, design methodologies, concurrent and distributed systems, applications and experiences, and large-scale systems such as database and workflow process management systems. The 17 coherently written chapters by leading researchers competently address a wide range of issues in exception handling.

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that may come packaged with the bound book. Database Systems: The Complete Book is ideal for Database Systems and Database Design and Application courses offered at the junior, senior and graduate levels in Computer Science departments. A basic understanding of algebraic expressions and laws, logic, basic data structure, OOP concepts, and programming environments is implied. Written by well-known computer scientists, this introduction to database systems offers a comprehensive approach, focusing on database design, database use, and implementation of database applications and database management systems. The first half of the book provides in-depth coverage of databases from the point of view of the database designer, user, and application programmer. It covers the latest database standards SQL:1999, SQL/PSM, SQL/CLI, JDBC, ODL, and XML, with broader coverage of SQL than most other texts. The second half of the book provides in-depth coverage of databases from the point of view of the DBMS implementor. It focuses on storage structures, query processing, and transaction management. The book covers the main techniques in these areas with broader coverage of query optimization than most other texts, along with advanced topics including multidimensional and bitmap indexes, distributed transactions, and information integration techniques.

Tuning your database for optimal performance means more than following a few short steps in a vendor-specific guide. For maximum improvement, you need a broad and deep knowledge of basic tuning principles, the ability to gather data in a systematic way, and the skill to make your system run faster. This is an art as well as a science, and Database Tuning: Principles, Experiments, and Troubleshooting Techniques will help you develop portable skills that will allow you to tune a wide variety of database systems on a multitude of hardware and operating systems. Further, these skills, combined with the scripts provided for validating results, are exactly what you need to evaluate competing database products and to choose the right one. Forward by Jim Gray, with invited chapters by Joe Celko and Alberto Lerner

Includes industrial contributions by Bill McKenna (RedBrick/Informix), Hany Saleeb (Oracle), Tim Shetler (TimesTen), Judy Smith (Deutsche Bank), and Ron Yorita (IBM) Covers the entire system environment: hardware, operating system, transactions, indexes, queries, table design, and application analysis Contains experiments (scripts available on the author's site) to help you verify a system's effectiveness in your own environment Presents special topics, including data warehousing, Web support, main memory databases, specialized databases, and financial time series Describes performance-monitoring techniques that will help you recognize and troubleshoot problems

This book constitutes the refereed proceedings of the 13th International Conference on Database Systems for Advanced Applications, DASFAA 2008, held in New Delhi, India, in March 2008. The 30 revised full papers and 27 revised short papers presented together with the abstracts of 3 invited talks as well as 8 demonstration papers and a panel discussion motivation were carefully reviewed and selected from 173 submissions. The papers are organized in topical sections on XML schemas, data mining, spatial data, indexes and cubes, data streams, P2P and transactions, XML processing, complex pattern processing, IR techniques, queries and transactions, data mining, XML databases, data warehouses and industrial applications, as well as mobile and distributed data.

Calendar and time units and specialized units, such as business days and academic years, play a major role in a wide range of information system applications. System support for reasoning about these units, called granularities, is important for the efficient design, use, and implementation of such applications. This book deals with several aspects of temporal information and provides a unifying model for granularities. Practitioners can learn about critical aspects that must be taken into account when designing and implementing databases supporting temporal information.

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.

This two volume set LNCS 10177 and 10178 constitutes the refereed proceedings of the 22nd International Conference on Database Systems for Advanced Applications, DASFAA 2017, held in Suzhou, China, in March 2017. The 73 full papers, 9 industry papers, 4 demo papers and 3 tutorials were carefully selected from a total of 300 submissions. The papers are organized around the following topics: semantic web and knowledge management; indexing and distributed systems; network embedding; trajectory and time series data processing; data mining; query processing and optimization; text mining; recommendation; security, privacy, sensor and cloud; social network analytics; map matching and spatial keywords; query processing and optimization; search and information retrieval; string and sequence processing; stream data processing; graph and network data processing; spatial databases; real time data processing; big data; social networks and graphs.

This book constitutes the refereed proceedings of the 14th International Conference on Database Systems for Advanced Applications, DASFAA 2009, held in Brisbane, Australia, in April 2009. The 39 revised full papers and 22 revised short papers presented together with 3 invited keynote papers, 9 demonstration papers, 3 tutorial abstracts, and one panel abstract were carefully reviewed and selected from 186 submissions. The papers are organized in topical sections on uncertain data and ranking, sensor networks, graphs, RFID and data streams, skyline and rising stars, parallel and distributed processing, mining and analysis, XML query, privacy, XML keyword search and ranking, Web and Web services, XML data processing, and multimedia.

The database field has experienced a rapid and incessant growth since the development of relational databases. The progress in database systems and applications has produced a diverse landscape of specialized technology areas that have often become the exclusive domain of research specialists. Examples include active databases, temporal databases, object-oriented databases, deductive databases, imprecise reasoning and queries, and multimedia information systems. This book provides a systematic introduction to and an in-depth treatment of these advanced database areas. It supplies practitioners and researchers with authoritative coverage of recent technological advances that are shaping the future of commercial database systems and intelligent information systems. Advanced Database Systems was written by a team of six leading specialists who have made significant contributions to the development of the technology areas covered in the book. Benefiting from the authors' long experience teaching graduate and professional courses, this book is designed to provide a gradual introduction to advanced research topics and includes many examples and exercises to support its use for individual study, desk reference, and graduate classroom teaching.

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.

The Handbook provides practitioners, scientists and graduate students with a good overview of basic notions, methods and techniques, as well as important issues and trends across the broad spectrum of data management. In particular, the book covers fundamental topics in the field such as distributed databases, parallel databases, advanced databases, object-oriented databases, advanced transaction management, workflow management, data warehousing, data mining, mobile computing, data integration and the Web. Summing up, the Handbook is a valuable source of information for academics and practitioners who are interested in learning the key ideas in the considered area.

Semantic Models for Multimedia Database Searching and Browsing begins with the introduction of multimedia information applications, the need for the development of the multimedia database management systems (MDBMSs), and the important issues and challenges of multimedia systems. The temporal relations, the spatial relations, the spatio-temporal relations, and several semantic models for multimedia information systems are also introduced. In addition, this book discusses recent advances in multimedia database searching and multimedia database browsing. More specifically, issues such as image/video segmentation, motion detection, object tracking, object recognition, knowledge-based event modeling, content-based retrieval, and key frame selections are presented for the first time in a single book. Two case studies consisting of two semantic models are included in the book to illustrate how to use semantic models to design multimedia information systems. Semantic Models for Multimedia Database Searching and Browsing is an excellent reference and can be used in advanced level courses for researchers, scientists, industry professionals, software engineers, students, and general readers who are interested in the issues, challenges, and ideas underlying the current practice of multimedia presentation, multimedia database searching, and multimedia browsing in multimedia information systems.

Nontraditional Database Systems is the fifth volume in the Advanced Information Processing Technology series. It brings together the results of research carried out by the Japanese database research community in the field of nontraditional database systems. The book examines nontraditional types of applications, data types, systems and environments together with high-performance architecture to support nontraditional applications, such as web mining, data engineering and object processing.

Increasingly, formal specification is being used by database researchers to describe and understand the systems they are designing and implementing. Similarly, those working on formal specification techniques have recognised that the database field provides a rich context for developing their ideas. However, as experts in one field often have a relatively limited knowledge of the other, there is a growing need for discussion about the relationship between these two fields and how they can be usefully combined. This volume contains the 16 papers which were presented at the International Workshop on Specification on Database Systems, held in Glasgow, 3-5 July 1991. The purpose of

the workshop was to bring together these fields and to examine, through a series of invited talks, presentations and working groups, the role that formal specification can play in developing database systems. The papers describe current research into topics such as the formal specification of data models, query languages and transaction handling and the use of formal specification techniques to understand problems which arise in database systems. The working groups, which are summarised at the end of the volume, covered a variety of issues including the role of graphical notations in database specification, the use of specification techniques in enabling "open" or extensible database systems and the education of the database community in specification techniques. This volume will be invaluable to the increasing number of researchers who are using both database systems and formal specification techniques in their work, and who wish to gain a more detailed knowledge of these two fields and the issues which affect them.

This book constitutes the thoroughly refereed post-proceedings of the First International Workshop on Formal Modeling and Analysis of Timed Systems, FORMATS 2003, held in Marseille, France in September 2003. The 19 revised full papers presented together with an invited paper and the abstracts of two invited talks were carefully selected from 36 submissions during two rounds of reviewing and improvement. All current aspects of formal method for modeling and analyzing timed systems are addressed; among the timed systems dealt with are timed automata, timed Petri nets, max-plus algebras, real-time systems, discrete time systems, timed languages, and real-time operating systems. This book constitutes the refereed proceedings of the 12th International Conference on Database Systems for Advanced Applications, DASFAA 2007, held in Bangkok, Thailand, April 2007. Coverage includes query language and query optimization, data mining and knowledge discovery, P2P and grid-based data management, XML databases, database modeling and information retrieval, Web and information retrieval, database applications and security.

This book constitutes the workshop proceedings of the 15th International Conference on Database Systems for Advanced Applications, DASFAA 2010, held in Tsukuba, Japan, in April 2010. The volume contains six workshops, each focusing on specific research issues that contribute to the main themes of the DASFAA conference: The First International Workshop on Graph Data Management: Techniques and Applications (GDM 2010), The Second International Workshop on Benchmarking of Database Management Systems and Data-Oriented Web Technologies (BenchmarkX'10); The Third International Workshop on Managing Data Quality in Collaborative Information Systems (MCIS2010), The Workshop on Social Networks and Social Media Mining on the Web (SNSMW2010), The Data Intensive eScience Workshop (DIEW 2010), and The Second International Workshop on Ubiquitous Data Management (UDM2010).

This book constitutes the refereed proceedings of the 11th East European Conference on Advances in Databases and Information Systems, ADBIS 2007, held in Varna, Bulgaria, in September/October 2007. The 23 revised papers presented together with three invited lectures were carefully reviewed and selected from 77 submissions. The papers address current research on database theory, development of advanced DBMS technologies, and their advanced applications.

This two volume set LNCS 5981 and LNCS 5982 constitutes the refereed proceedings of the 15th International Conference on Database Systems for Advanced Applications, DASFAA 2010, held in Tsukuba, Japan, in April 2010. The 39 revised full papers and 16 revised short papers presented together with 3 invited keynote papers, 22 demonstration papers, 6 industrial papers, and 2 keynote talks were carefully reviewed and selected from 285 submissions. The papers of the first volume are organized in topical sections on P2P-based technologies, data mining technologies, XML search and matching, graphs, spatial databases, XML technologies, time series and streams, advanced data mining, query processing, Web, sensor networks and communications, information management, as well as communities and Web graphs. The second volume contains contributions related to trajectories and moving objects, skyline queries, privacy and security, data streams, similarity search and event processing, storage and advanced topics, industrial, demo papers, and tutorials and panels.

Geographic Information Systems (GIS) have been experiencing a steady and unprecedented growth in terms of general interest, theory development, and new applications in the last decade or so. GIS is an inter-disciplinary field that brings together many diverse areas such as computer science, geography, cartography, engineering, and urban planning. Database Issues in Geographic Information Systems approaches several important topics in GIS from a database perspective. Database management has a central role to play in most computer-based information systems, and is expected to have an equally important role to play in managing information in GIS as well. Existing database technology, however, focuses on the alphanumeric data that are required in business applications. GIS, like many other application areas, requires the ability to handle spatial as well as alphanumeric data. This requires new innovations in data management, which is the central theme of this monograph. The monograph begins with an overview of different application areas and their data and functional requirements. Next it addresses the following topics in the context of GIS: representation and manipulation of spatial data, data modeling, indexing, and query processing. Future research directions are outlined in each of the above topics. The last chapter discusses issues that are emerging as important areas of technological innovations in GIS. Database Issues in Geographic Information Systems is suitable as a secondary text for a graduate level course on Geographic Information Systems, Database Systems or Cartography, and as a reference for researchers and practitioners in industry.

This book constitutes the workshop proceedings of the 16th International Conference on Database Systems for Advanced Applications, DASFAA 2011, held in Hong Kong, China, in April 2011. The volume contains six workshops, each focusing on specific research issues that contribute to the main themes of the DASFAA conference: The First International Workshop on Graph-structured Data Bases (GDB 2011); the First International Workshop on Spatial Information Modeling, Management and Mining (SIM3 2011); the International Workshop on Flash-based Database Systems (FlashDB 2011); the Second International Workshop on Social Networks and Social Media Mining on the Web (SNSMW 2011); the First International Workshop on Data Management for Emerging Network Infrastructures (DaMEN 2011); and the Fourth International Workshop on Data Quality in Integration Systems (DQIS 2011).

This volume contains expanded and updated versions of papers presented at CIKM 92, the First International

Conference on Information and Knowledge Management, held in Baltimore, Maryland, in November 1992. The conference participants came from the fields of database management, information retrieval, and artificial intelligence. A major theme in the volume is how these fields are merging and enriching each other. The eight papers discuss: discovering context in a conceptual schema; a system for face information retrieval; queries in OODB systems; consistency checking in OODBs; integrity constraints representation in OODBs; a framework for temporal object databases; inductive dependencies and approximate databases; OODB design methodologies.

The 4 volume set LNCS 12112-12114 constitutes the papers of the 25th International Conference on Database Systems for Advanced Applications which will be held online in September 2020. The 119 full papers presented together with 19 short papers plus 15 demo papers and 4 industrial papers in this volume were carefully reviewed and selected from a total of 487 submissions. The conference program presents the state-of-the-art R&D activities in database systems and their applications. It provides a forum for technical presentations and discussions among database researchers, developers and users from academia, business and industry.

This book constitutes the refereed proceedings of the 11th International Conference on Database Systems for Advanced Applications, DASFAA 2006, held in Singapore in April 2006. 46 revised full papers and 16 revised short papers presented were carefully reviewed and selected from 188 submissions. Topics include sensor networks, subsequence matching and repeating patterns, spatial-temporal databases, data mining, XML compression and indexing, xpath query evaluation, uncertainty and streams, peer-to-peer and distributed networks and more.

This two volume set LNCS 9642 and LNCS 9643 constitutes the refereed proceedings of the 21st International Conference on Database Systems for Advanced Applications, DASFAA 2016, held in Dallas, TX, USA, in April 2016. The 61 full papers presented were carefully reviewed and selected from a total of 183 submissions. The papers cover the following topics: crowdsourcing, data quality, entity identification, data mining and machine learning, recommendation, semantics computing and knowledge base, textual data, social networks, complex queries, similarity computing, graph databases, and miscellaneous, advanced applications.

Advanced Database Systems Springer Science & Business Media

This two volume set LNCS 6587 and LNCS 6588 constitutes the refereed proceedings of the 16th International Conference on Database Systems for Advanced Applications, DASFAA 2011, held in Saarbrücken, Germany, in April 2010. The 53 revised full papers and 12 revised short papers presented together with 2 invited keynote papers, 22 demonstration papers, 4 industrial papers, 8 demo papers, and the abstract of 1 panel discussion, were carefully reviewed and selected from a total of 225 submissions. The topics covered are social network, social network and privacy, data mining, probability and uncertainty, stream processing, graph, XML, XML and graph, similarity, searching and digital preservation, spatial queries, query processing, as well as indexing and high performance.

On behalf of the Organizing Committee, we would like to welcome you to the proceedings of the 10th International Conference on Database Systems for Advanced Applications (DASFAA 2005).

Active database systems enhance traditional database functionality with powerful rule-processing capabilities, providing a uniform and efficient mechanism for many database system applications. Among these applications are integrity constraints, views, authorization, statistics gathering, monitoring and alerting, knowledge-based systems, expert systems, and workflow management. This significant collection focuses on the most prominent research projects in active database systems. The project leaders for each prototype system provide detailed discussions of their projects and the relevance of their results to the future of active database systems. Features: A broad overview of current active database systems and how they can be extended and improved A comprehensive introduction to the core topics of the field, including its motivation and history Coverage of active database (trigger) capabilities in commercial products Discussion of forthcoming standards

This volume contains three keynote papers and 51 technical papers from contributors around the world on topics in the research and development of database systems, such as Data Modelling, Object-Oriented Databases, Active Databases, Data Mining, Heterogeneous Databases, Distributed Databases, Parallel Query Processing, Multi-Media Databases, Transaction Management Systems, Document Databases, Temporal Databases, Deductive Databases, User Interface, and Advanced Database Applications.

This book constitutes the refereed proceedings of the 4th International Conference on Deductive and Object-Oriented Databases, DOOD '95, held in Singapore in December 1995. Besides two keynote papers by Stefano Ceri and Michael Kifer, the book contains revised full versions of 28 papers selected from a total of 88 submissions. The volume gives a highly competent state-of-the-art report on DOOD research and advanced applications. The papers are organized in sections on active databases, query processing, semantic query optimization, transaction management, authorization, implementation, and applications.

[Copyright: 9204435e1072febb78aacc6783a0f452](https://doi.org/10.1007/978-1-4020-0452-2)