Active Database Systems Triggers And Rules For Advanced Database Processing The Morgan Kaufmann Series In Data Management Systems

This book constitutes the proceedings of the 5th International Conference on Web Information Systems Engineering, WISE 2004, held in Brisbane, Australia in November 2004. The 45 revised full papers and 29 revised short papers presented together with 3 invited contributions were carefully reviewed and selected from 198 submissions. The papers are organized in topical sections on Web information modeling; payment and security; information extraction; advanced applications; performance issues; linkage analysis and document clustering; Web caching and content analysis; XML query processing; Web search and personalization; workflow management and enterprise information systems; business processes; deep Web and dynamic content; Web information systems design; ontologies and applicatoins; multimedia, user interfaces, and languages; and peer-to-peer and grid systems. Content Description #Includes bibliographical references and index.

This book constitutes the strictly refereed post-workshop proceedings of the Third International Workshop on Rules in Database Systems, RIDS '97, held in Skövde, June 1997. The 13 revised papers presented in the book were carefully reviewed and selected from 33 submissions. The book documents the state-of-the art in the area. The papers are devoted to deductive databases, active database systems architectures, events in workflow management,

rule modelling and simulation, rule confluence, rule termination analysis, rule testing and validation, active database systems design.

This book constitutes the refereed proceedings of the 19th International Conference on Conceptual Modeling, ER 2000, held in Salt Lake City, Utah, USA in October 2000.The 37 revised full papers presented together with three invited papers and eight industrial abstracts were carefully reviewed and selected from a total of 140 submitted papers. The book offers topical sections on database integration, temporal and active database modeling, database and data warehouse design techniques, analysis patterns and ontologies, Web-based information systems, business process modeling, conceptual modeling and XML, engineering and multimedia application modeling, object-oriented modeling, applying object-oriented technology, quality in conceptual modeling, and application design using UML.

Active Database SystemsTriggers and Rules for Advanced Database ProcessingMorgan Kaufmann

Video segmentation is the most fundamental process for appropriate index ing and retrieval of video intervals. In general, video streams are composed 1 of shots delimited by physical shot boundaries. Substantial work has been done on how to detect such shot boundaries automatically (Arman et al. , 1993) (Zhang et al. , 1993) (Zhang et al. , 1995) (Kobla et al. , 1997). Through the inte gration of technologies such as image processing, speech/character recognition and natural language understanding, keywords can be extracted and associated with these shots for indexing (Wactlar et al. , 1996). A single shot, however, rarely carries enough amount of information to be meaningful by itself. Usu ally, it is a semantically meaningful interval that most users are interested in re trieving. Generally, such meaningful

intervals span several consecutive shots. There hardly exists any efficient and reliable technique, either automatic or manual, to identify all semantically meaningful intervals within a video stream. Works by (Smith and Davenport, 1992) (Oomoto and Tanaka, 1993) (Weiss et al. , 1995) (Hjelsvold et al. , 1996) suggest manually defining all such inter vals in the database in advance. However, even an hour long video may have an indefinite number of meaningful intervals. Moreover, video data is multi interpretative. Therefore, given a query, what is a meaningful interval to an annotator may not be meaningful to the user who issues the query. In practice, manual indexing of meaningful intervals is labour intensive and inadequate.

This book is the proceedings of a workshop held at Heriot-Watt University in Edinburgh in August 1993. The central theme of the workshop was rules in database systems, and the papers presented covered a range of different aspects of database rule systems. These aspects are reflected in the sessions of the workshop, which are the same as the sections in this proceedings: Active Databases Architectures Incorporating Temporal Rules Rules and Transactions Analysis and Debugging of Active Rules Integrating Graphs/Objects with Deduction Integrating Deductive and Active Rules Integrity Constraints Deductive Databases The incorporation of rules into database systems is an important area of research, as it is a major component in the integration of behavioural information with the structural data with which commercial databases have traditionally been associated. This integration of the behavioural aspects of an application with the data to which it applies in database systems leads to more straightforward application development and more efficient processing of data. Many novel applications seem to need database systems in which structural and behavioural information are fully integrated. Rules are only one means of expressing behavioural

information, but it is clear that different types of rule can be used to capture directly different properties of an application which are cumbersome to support using conventional database architectures. In recent years there has been a surge of research activity focusing upon active database systems, and this volume opens with a collection of papers devoted specifically to this topic.

The design of knowledge systems is finding myriad applications from corporate databases to general decision support in areas as diverse as engineering, manufacturing and other industrial processes, medicine, business, and economics. In engineering, for example, knowledge bases can be utilized for reliable electric power system operation. In medicine they support complex diagnoses, while in business they inform the process of strategic planning. Programmed securities trading and the defeat of chess champion Kasparov by IBM's Big Blue are two familiar examples of dedicated knowledge bases in combination with an expert system for decision-making. With volumes covering "Implementation," "Optimization," "Computer Techniques," and "Systems and Applications," this comprehensive set constitutes a unique reference source for students, practitioners, and researchers in computer science, engineering, and the broad range of applications areas for knowledge-based systems.

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 book constitutes the refereed proceedings of the First International Workshop on Cooperative Information Agents - DAI Meets Databases, CIA-97, held in Kiel, Germany, in February 1997. The book opens with 6 invited full papers by internationally leading researchers surveying the state of the art in the area. The 16 revised full research papers presented were carefully selected during a highly competitive round of reviewing. The papers are organized in topical sections on databases and agent technology, agents for database search and knowledge discovery, communication and cooperation among information agents, and agent-based access to heterogeneous information sources.

This book constitutes the refereed proceedings of the 9th International Conference on Database and Expert Systems Applications, DEXA'98, held in Vienna, Austria, in August 1998. The 81 revised full papers presented were carefully selected from a total of more than 200 submissions. The papers are organized in sections on active

databases, object-oriented systems, data engineering, information retrieval, workflow and cooperative systems, spatial and temporal aspects, document management, spatial databases, adaptation and view updates, genetic algorithms, cooperative and distributed environments, interaction and communication, transcation, advanced applications, temporal aspects, oriented systems, partitioning and fragmentation, database queries, data, data warehouses, knowledge discovery and data mining, knowledge extraction, and knowledge base reduction for comprehension and reuse. These proceedings contain 25 contributed papers presented at the 13th East-EuropeanConferenceAdvances on Databases and InformationSystems (ADBIS 2009) held September 7-10, 2009, in Riga, Latvia. The Call for Papers attracted 93 submissions from 28 countries. In a rigorous reviewing process the inter- tional Program Committee of 64 members from 29 countries selected these 25 contributions for publication in this volume; in addition, there is the abstract of an invited talk by Matthias Brantner. Furthermore, 18 additional contributions were selected for short presentations and have been published in a separate v- ume of local proceedings by the organizing institution. Topically, the accepted paperscoverawidespectrumofdatabaseandinformationsystemtopicsranging from query processing and optimization via query languages, design methods, data integration, indexing and caching to business processes, data mining, and application oriented topics like XML and data on the Web. The ADBIS 2009conference continued the series

of ADBIS conferencesor- nized every year in di?erent countries of Easternand Central Europe, beginning in St. Petersburg (Russia, 1997), Poznan (Poland, 1998), Maribor (Slovenia, 1999), Prague (Czech Republic, as a joint ADBIS-DASFAA conference, 2000), Vilnius(Lithuania,2001), Bratislava(Slovakia,2002), Dresden(Germany,2003), Budapest(Hungary,2004), Tallinn(Estonia,2005), Thessaloniki(Greece,2006), Varna (Bulgaria, 2007), and Pori (Finland, 2008). The conferences are initiated and supervised by an international Steering Committee, which consists of r- resentatives from Armenia, Austria, Bulgaria, Czech Republic, Greece, Estonia, Germany, Hungary, Israel, Italy, Latvia, Lithuania, Poland, Russia, Serbia, S- vakia, Slovenia, and Ukraine, and is chaired by Professor Leonid Kalinichenko.

The three volume set LNAI 5177, LNAI 5178, and LNAI 5179, constitutes the refereed proceedings of the 12th International Conference on Knowledge-Based Intelligent Information and Engineering Systems, KES 2008, held in Zagreb, Croatia, in September 2008. The 316 revised papers presented were carefully reviewed and selected. The papers present a wealth of original research results from the field of intelligent information processing in the broadest sense; topics covered in the second volume are artificial intelligence driven engineering design optimization; biomedical informatics: intelligent information management from nanomedicine to public health; communicative intelligence; computational intelligence for image processing and pattern recognition; computational intelligence in human cancer research;

computational intelligence techniques for Web personalization; computational intelligent techniques for bioprocess modelling, monitoring and control; intelligent computing for Grid; intelligent security techniques; intelligent utilization of soft computing techniques; reasoning-based intelligent systems: relevant reasoning for discovery and prediction; spatio-temporal database concept support for organizing virtual earth; advanced knowledge-based systems; chance discovery; innovation-oriented knowledge management platform; knowledge-based creativity support systems; knowledge-based interface systems; knowledge-based multi-criteria decision support; and knowledge-based systems for e-business.

The World Wide Web has become a ubiquitous global tool, used for finding infor mation, communicating ideas, carrying out distributed computation and conducting business, learning and science. The Web is highly dynamic in both the content and quantity of the information that it encompasses. In order to fully exploit its enormous potential as a global repository of information, we need to understand how its size, topology and content are evolv ing. This then allows the development of new techniques for locating and retrieving information that are better able to adapt and scale to its change and growth. The Web's users are highly diverse and can access the Web from a variety of devices and interfaces, at different places and times, and for varying purposes. We thus also need techniques for personalising the presentation and content of Web based information depending on how it is being accessed and on the specific

user's requirements. As well as being accessed by human users, the Web is also accessed by appli cations. New applications in areas such as e-business, sensor networks, and mobile and ubiquitous computing need to be able to detect and react quickly to events and changes in Web-based information. Traditional approaches using query-based 'pull' of information to find out if events or changes of interest have occurred may not be able to scale to the quantity and frequency of events and changes being generated, and new 'push' -based techniques are needed.

This book constitutes the refereed proceedings of the 5th International Conference on Deductive and Object-Oriented Databases, DOOD'97, held in Montreux, Switzerland, in December 1997. The 22 revised full papers presented in this book were selected from a total of 59 submissions. Also included are abstracts or full versions of three invited talks and three tutorials and six short presentations. The papers are organized in topical sections on materialized view maintenance, extending DBMs features, database updates, managing change in object databases, semantics of active databases, formal semantics, and new directions.

This book contains the refereed proceedings of the 8th International Conference on Database and Expert Systems Applications, DEXA '97, held in Toulouse, France, September 1997. The 62 revised full papers presented in the book, together with three invited contributions, were selected from a total of 159 submissions. The papers are organized in sections on modeling, object-oriented databases, active and temporal

aspects, images, integrity constraints, multimedia databases, deductive databases and knowledge-based systems, allocation concepts, data interchange, digital libraries, transaction concepts, learning issues, optimization and performance, query languages, maintenance, federated databases, uncertainty handling and qualitative reasoning, and software engineering and reusable software.

This book constitutes the strictly refereed post-workshop proceedings of the International Workshop on Logic in Databases, LID'96, held in San Miniato, Italy, in July 1996, as the final meeting of an EC-US cooperative activity. The volume presents 21 revised full papers selected from 49 submissions as well as 3 invited contributions and a summary of a panel discussion on deductive databases: challenges, opportunities and future directions. The retrospective survey on logic and databases by Jack Minker deserves a special mention: it is a 56-page overview and lists 357 references. The papers are organized in sections on uncertainty, temporal and spatial reasoning, updates, active databases, semantics, advanced applications, query evaluation, language extensions, and logic constructs and expressive power.

Knowledge Base Systems are an integration of conventional database systems with Artificial Intelligence techniques. They provide inference capabilities to the database system by encapsulating the knowledge of the application domain

within the database. Knowledge is the most valuable of all corporate resources that must be captured, stored, re-used and continuously improved, in much the same way as database systems were important in the previous decade. Flexible, extensible, and yet efficient Knowledge Base Systems are needed to capture the increasing demand for knowledge-based applications which will become a significant market in the next decade. Knowledge can be expressed in many static and dynamic forms; the most prominent being domain objects, their relationships, and their rules of evolution and transformation. It is important to express and seamlessly use all types of knowledge in a single Knowledge Base System. Parallel, Object-Oriented, and Active Knowledge Base Systems presents in detail features that a Knowledge Base System should have in order to fulfill the above requirements. Parallel, Object-Oriented, and Active Knowledge Base Systems covers in detail the following topics: Integration of deductive, production, and active rules in sequential database systems. Integration and inter-operation of multiple rule types into the same Knowledge Base System. Parallel rule matching and execution, for deductive, production, and active rules, in parallel Export, Knowledge Base, and Database Systems. In-depth description of a Parallel, Object-Oriented, and Active Knowledge Base System that integrates all rule paradigms into a single database system without hindering performance.

Parallel, Object-Oriented, and Active Knowledge Base Systems is intended as a graduate-level text for a course on Knowledge Base Systems and as a reference for researchers and practitioners in the areas of database systems, knowledge base systems and Artificial Intelligence.

A timely survey of the field from the point of view of some of the subject's most active researchers. Divided into several parts organized by theme, the book first covers the underlying methodology regarding active rules, followed by formal specification, rule analysis, performance analysis, and support tools. It then moves on to the implementation of active rules in a number of commercial systems, before concluding with applications and future directions for research. All researchers in databases will find this a valuable overview of the topic.

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.

Doctoral Thesis / Dissertation from the year 2006 in the subject Computer

Science - Applied, grade: 1.0, University of Dusseldorf "Heinrich Heine," 127 entries in the bibliography, language: English, abstract: Federated information systems provide access to interrelated data that is distributed over multiple autonomous and heterogeneous data sources. The integration of these sources demands for flexible and extensible architectures that balance both, the highest possible autonomy and a reasonable degree of information sharing. In current federated information systems, the integrated data sources do only have passive functionality with regard to the federation. However, continuous improvements take the functionality of modern databases beyond former limits. The significant improvement, on which this work is based on, is the ability of modern active database systems to execute programs written in a standalone programming language as user-defined functions or stored procedures from within their database management systems. We introduce Enhanced Active Database Systems as a new subclass of active databases that are able to interact with other components of a federation using external program calls from within triggers. We present several concepts and architectures that are specifically developed for Enhanced Active Databases to improve interoperability and consistency in federated information systems. As the basic concept we describe Active Event Notifications to provide an information system with synchronous and

asynchronous update notifications in real-time. Based on this functionality, Enhanced Active Databases are able to actively participate in global integrity maintenance executing partial constraint checks on interrelated remote data. Furthermore, we present an architecture for a universal wrapper component that especially supports Active Event Notifications, which makes it perfectly suitable for eve

This book constitutes the thoroughly refereed post-workshop proceedings of the Second International Workshop on Active, Real-Time, and Temporal Database Systems, ARTDB'97, held in Como, Italy in September 1997. The nine revised full papers presented were carefully reviewed and selected from a total of 17 submissions. Also included are an introductory survey chapter and three invited papers written by prominent researchers in the field, as well as two summaries of the panel discussions held at the workshop. The papers are devoted to the issue of building database systems supporting reactive behaviour, while enforcing timeliness and predictability.

"Focused on the latest research on text and document management, this guide addresses the information management needs of organizations by providing the most recent findings. How the need for effective databases to house information is impacting organizations worldwide and how some organizations that possess a

vast amount of data are not able to use the data in an economic and efficient manner is demonstrated. A taxonomy for object-oriented databases, metrics for controlling database complexity, and a guide to accommodating hierarchies in relational databases are provided. Also covered is how to apply Java-triggers for X-Link management and how to build signatures."

Most modern-day organizations have a need to record data relevant to their everyday activities and many choose to organise and store some of this information in an electronic database. Database Systems provides an essential introduction to modern database technology and the development of database systems. This new edition has been fully updated to include new developments in the field, and features new chapters on: e-business, database development process, requirements for databases, and distributed processing. In addition, a wealth of new examples and exercises have been added to each chapter to make the book more practically useful to students, and full lecturer support will be available online.

As the information contained in databases has become a critical resource in organizations, efficient access to that information and the ability to share it among different users and across different systems has become an urgent need. The interoperability of heterogeneous database systems-literally, the ability to access

information between or among differing types of databases, is the topic of this timely book. In the last two decades, tremendous improvements in tools and technologies have resulted in new products that provide distributed data processing capabilities. This book describes these tools and emerging technologies, explaining the essential concepts behind the topics but focusing on practical applications. Selected products are discussed to illustrate the characteristics of the different technologies. This is an ideal source for anyone who needs a broad perspective on heterogeneous database integration and related technologies.

"Addresses the evolution of database management, technologies and applications along with the progress and endeavors of new research areas."--P. xiii.

Component Database Systems is a collection of invited chapters by the researchers making the most influential contributions in the database industry's trend toward componentization This book represents the sometimes-divergent, sometimes-convergent approaches taken by leading database vendors as they seek to establish commercially viable componentization strategies. Together, these contributions form the first book devoted entirely to the technical and architectural design of component-based database systems. In addition to detailing the current state of their research, the authors also take up many of the issues affecting the likely future directions of component databases. If you have a stake in the evolution of any of today's leading database systems, this book will make fascinating reading. It will also help prepare you

for the technology that is likely to become widely available over the next several years. * Is comprised of contributions from the field's most highly respected researchers, including key figures at IBM, Oracle, Informix, Microsoft, and POET. * Represents the entire spectrum of approaches taken by leading software companies working on DBMS componentization strategies. * Covers component-focused architectures, methods for hooking components into an overall system, and support for component development. * Examines the component technologies that are most valuable to Web-based and multimedia databases. * Presents a thorough classification and overview of component database systems.

Comprehensive coverage of critical issues related to information science and technology.

This book constitutes the thoroughly refereed and revised post-workshop proceedings of two international meetings devoted to deductive databases and logic programming. The technical papers presenting original research results were carefully reviewed and selected for inclusion in the present book, together with several survey articles summarizing the state of the art in the area. Thus, the book is the ultimate reference for anybody interested in the theory and application of deductive databases and logic programming.

IDT (Intelligent Decision Technologies) seeks an interchange of research on intelligent systems and intelligent technologies which enhance or improve decision making in

industry, government and academia. The focus is interdisciplinary in nature, and includes research on all aspects of intelligent decision technologies, from fundamental development to the applied system. It constitutes a great honor and pleasure for us to publish the works and new research results of scholars from the First KES International Symposium on Intelligent Decision Technologies (KES IDT'09), hosted and organized by University of Hyogo in conjunction with KES International (Himeji, Japan, April, 2009). The symposium was concerned with theory, design, development, implementation, testing and evaluation of intelligent decision systems. Its topics included intelligent agents, fuzzy logic, multi-agent systems, artificial neural networks, genetic algorithms, expert systems, intelligent decision making support systems, information retrieval systems, geographic information systems, and knowledge management systems. These technologies have the potential to support decision making in many areas of management, international business, finance, accounting, marketing, healthcare, military applications, production, networks, traffic management, crisis response, and human interfaces.

modelling large-scale pr- lems in computing and biochemistry.

The latest edition of a popular text and reference on database research, with substantial new material and revision; covers classical literature and recent hot topics. Lessons from database research have been applied in academic fields ranging from bioinformatics to next-generation Internet architecture and in industrial uses including

Web-based e-commerce and search engines. The core ideas in the field have become increasingly influential. This text provides both students and professionals with a grounding in database research and a technical context for understanding recent innovations in the field. The readings included treat the most important issues in the database area--the basic material for any DBMS professional. This fourth edition has been substantially updated and revised, with 21 of the 48 papers new to the edition, four of them published for the first time. Many of the sections have been newly organized, and each section includes a new or substantially revised introduction that discusses the context, motivation, and controversies in a particular area, placing it in the broader perspective of database research. Two introductory articles, never before published, provide an organized, current introduction to basic knowledge of the field; one discusses the history of data models and query languages and the other offers an architectural overview of a database system. The remaining articles range from the classical literature on database research to treatments of current hot topics, including a paper on search engine architecture and a paper on application servers, both written expressly for this edition. The result is a collection of papers that are seminal and also accessible to a reader who has a basic familiarity with database systems.

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 book constitutes the refereed proceedings of the 23rd British National Conference on Databases, BNCOD 23, held in Belfast, Northern Ireland, UK in July 2006. The 12 revised full papers and 6 revised short papers presented together with 2 invited lectures and 13 poster papers were carefully reviewed and selected from 58 submissions. The papers are organized in topical sections on

data modelling and architectures and transaction management, data integration and interoperability and information retrieval, query processing and optimisation, data mining, data warehousing and decision-support systems, as well as data streaming.

Current Trends in Data Management Technology reports on the most recent, important advances in data management as it applies to diverse issues, such as Web information management, workflow systems, electronic commerce, reengineering business processes, object-oriented databases, and more. Geared toward designers and professionals interested in the conceptual aspects of integrity problems in different paradigms, Database Integrity: Challenges and Solutions successfully addresses these and a variety of other issues.

ABSTRACT (cont.): Active database semantics can be supported on an existing SQL Server (we use Oracle SQL Server as the test SQL Server) by the ECA Agent between the SQL Server and multiple clients. ECA rules are completely supported through the ECA Agent without changing applications in the SQL Server. Both primitive and composite events can be detected in the ECA Agent and actions are invoked in SQL Server. All events are persistent in RDBMS. The Java Local Event Detector (Java LED) is used to notify and detect both primitive events and composite events. The ECA Agent uses Java Database Connectivity

(JDBC) to connect to the SQL server. The architecture of the ECA Agent and implementation details are shown in this thesis. Alternative approaches are discussed in details, and the features and limitations are identified.

Time is ubiquitous in information systems. Almost every enterprise faces the problem of its data becoming out of date. However, such data is often valu able, so it should be archived and some means to access it should be provided. Also, some data may be inherently historical, e.g., medical, cadastral, or ju dicial records. Temporal databases provide a uniform and systematic way of dealing with historical data. Many languages have been proposed for tem poral databases, among others temporal logic. Temporal logic combines ab stract, formal semantics with the amenability to efficient implementation. This chapter shows how temporal logic can be used in temporal database applica tions. Rather than presenting new results, we report on recent developments and survey the field in a systematic way using a unified formal framework [GHR94; Ch094]. The handbook [GHR94] is a comprehensive reference on mathematical foundations of temporal logic. In this chapter we study how temporal logic is used as a query and integrity constraint language. Consequently, model-theoretic notions, particularly for mula satisfaction, are of primary interest. Axiomatic systems and proof meth ods for temporal logic [GHR94] have found so far

relatively few applications in the context of information systems. Moreover, one needs to bear in mind that for the standard linearly-ordered time domains temporal logic is not re cursively axiomatizable [GHR94]' so recursive axiomatizations are by necessity incomplete.

This book constitutes the refereed proceedings of the 5th East European Conference on Advances in Databases and Information Systems, ADBIS 2001, held in Vilnius, Lithuania, in September 2001. The 25 revised full papers presented together with one invited paper and two abstracts of invited talks were carefully reviewed and selected from 82 submissions. The papers are organized in topical sections on query optimization, multimedia and multilingual information systems, spatiotemporal aspects of databases, data mining, transaction processing, conceptual modeling and information systems specification, active databases, query methods, XML, and information systems design.

Copyright: 91de0dc42e573027610d53f1f339bc71