

Book Review

Transaction Processing on Modern Hardware

Mohammad Sadoghi
Spyros Blanas

Synthesis Lectures on Data Management
Morgan & Claypool Publishers
ISBN: 9781681734996 and 9781681735016
Copyright 2019

When transactions data are processed, it leads to the derivation of meaningful and useful inferences. Realizing this value, the authors have brought a good work on it. Transactions are stored in databases, which provide a trigger for analysing them. In the last ten years, the databases data enable the researchers to move to unearth the processing of the transactions. This book has eight chapters with the first chapter introduction that outlines the treatment of data transactions.

The second chapter on 'Transaction Concepts' have given a focus on basic concepts associated with the data transactions. The conceptual clarifications and the features and requirements are presented in this unit. The transactions characteristics include atomicity, consistency, isolation and durability. Concurrent transactions give rise to conflicts that can be controlled by executing the transactions serially. To do it so, concurrent protocols are established which are explained in this unit.

The next chapter on 'Multi-version Concurrency Revisited', brief the developments in this subdomain over the years by highlighting the changes, architecture and data models. Different concurrency models with data architecture and models form the core of this unit. In the fourth chapter on Coordination-avoidance concurrency, the new kind of deterministic transaction processing tries to remove any non-deterministic code path from the transaction execution logic which helps to eliminate all execution induced aborts. Here the possibility of logic-induced aborts is detailed.

In the next chapter, the concurrency protocols which have the multi-core architecture are discussed in addition to the description of various architectures for the transactional systems. The designs for the data partitioning and indexing are explained in the sixth chapter on Hardware-assisted Transactional Utilities. The database partitioning concepts and indexing fundamentals with features form the body of this unit.

In the seventh chapter on 'Transactions on Heterogeneous Hardware', the authors addressed the core issue of leveraging the heterogeneity and efficient data access in the network of heterogeneous systems are given. In the last chapter on 'Outlook', the authors have forecasted the future of hardware specialization and its possible impact. This book is supported with a bibliography of both the literature used and required for further understanding.

Hathairat Ketmaneechairat
King Mongkut's University of Technology North Bangkok
Thailand

Query Processing Over Incomplete Databases

Yunjun Gao and Xiaoye Miao
Synthesis Lectures on Data Management
Morgan & Claypool Publishers
www.morganclaypool.com
ISBN: 9781681734200 and 9781681734224

Query processing is a fundamental activity in database processing which has been extensively addressed. However, the incomplete data in some datasets lead to the birth of incomplete databases where query processing is not straight forward. Thus the need for addressing this issue is fulfilled by this book which has five chapters.

In the first chapter on Introduction, the authors basically explained the issue of incomplete data management. Further, the querying and indexing in the incomplete databases with illustrations are presented. Three challenges in this direction such as index structure design, form efficient algorithms and query answering are posed by the authors.

The missing value is the main character of the incomplete datasets. They outlined three methods to deal with the missing values such as data discarding model, observed data dependant model and data imputation model in the second chapter on how to handle incomplete data methods. Further, they focus on statistical methods, machine learning approaches and a few new advance imputation methods.

In the third chapter on Query Semantics on Incomplete data, the authors have present three representative query semantics such as k-neighbor search, skyline query, and top-k dominating query in the incomplete databases.

The fourth chapter on Advanced Techniques is very crucial in this book. The authors treated a few advanced techniques for processing queries over incomplete data with a discussion on novel index structures, effective pruning heuristics and the crowdsourcing techniques. The index structures are well explained with the help of many illustrations followed by an extensive discussion on crowdsourcing techniques. The ten algorithms presented in this chapter with numerous architectures enhance the value of the content of this book.

The last chapter on Conclusion summarize the contributions. This book is supported by a well-knitted bibliography of useful references.

Hathairat Ketmaneechairat
King Mongkut's University of Technology North Bangkok
Thailand