


Massively Parallel Databases and MapReduce Systems

Timely and cost-effective analytics over "big data" has emerged as a key ingredient for success in many businesses, scientific and engineering disciplines, and government endeavors. Web clicks, social media, scientific experiments, and datacenter monitoring are among data sources that generate vast amounts of raw data every day. The need to convert this raw data into useful information has spawned considerable innovation in systems for large-scale data analytics, especially over the last decade. Massively Parallel Databases and MapReduce Systems addresses the design principles and core features of systems for analyzing very large datasets using massively-parallel computation and storage techniques on large clusters of nodes. It first discusses how the requirements of data analytics have evolved since the early work on parallel database systems. It then describes some of the major technological innovations that have each spawned a distinct category of systems for data analytics. Each unique system category is described along a number of dimensions including data model and query interface, storage layer, execution engine, query optimization, scheduling, resource management, and fault tolerance. It concludes with a summary of present trends in large-scale data analytics. Massively Parallel Databases and MapReduce Systems is an ideal reference for anyone with a research or professional interest in large-scale data analytics.



85,00 €

79,44 € (zzgl. MwSt.)

Nicht mehr lieferbar

Artikelnummer: 9781601987501

Medium: Buch

ISBN: 978-1-60198-750-1

Verlag: Now Publishers

Erscheinungstermin: 20.11.2013

Sprache(n): Englisch

Auflage: Erscheinungsjahr 2013

Serie: Foundations and Trends® in Databases

Produktform: Kartoniert

Seiten: 120

Format (B x H): 156 x 234 mm

