

Preface

Key-value stores widely served as data platforms for various web applications, ranging from Facebook social networking based on Cassandra to Amazon online shopping hosted by Dynamo, providing a distributed solution to cloud computing and big data management. In modern web applications, user experience satisfaction determines the applications' success rate. In order to answer web queries quickly, key-value stores generally employ weak consistency model. This model relaxes the data consistency to improve system performance on query response. However, a drawback is that data accessed by users might be stale. Hence, there is an intrinsic trade-off between query latency and data consistency, which has become a key factor in the design of large-scale data management systems.

In particular, data consistency expresses as the freshness of data accessed by queries at a local node. Clearly, the latency/consistency trade-off at node level boils down to finding the trade-off between query latency (i.e., *Quality of Service (QoS)*) and data freshness (i.e., *Quality of Data (QoD)*). In real application, different web queries or users would have different expectations in terms of QoS and QoD. Hence, how to optimize QoS and QoD by scheduling queries and updates in key-value stores becomes an important research issue.

Our book comprehensively illustrates quality-aware scheduling in key-value stores. In addition, this work provides scheduling strategies and a prototype framework of quality-aware scheduler as well as a demonstration of online applications. The book offers a rich blend of theory and practice. It is suitable for students, researchers, and practitioners interested in distributed systems, NoSQL key-value stores, and scheduling.

Shanghai, China
April 2015

Chen Xu
Aoying Zhou

Quality-aware Scheduling for Key-value Data Stores

Xu, C.; Zhou, A.

2015, XI, 97 p. 32 illus., Softcover

ISBN: 978-3-662-47305-4