

MISQ Archivist

A Time-Based Dynamic Synchronization Policy for Consolidated Database Systems

Xinxue (Shawn) Qu and Zhengrui Jiang

Abstract

Data is becoming an increasingly important asset in today's organizations. Various challenges in the age of big data, such as high volume and high velocity, call for efficient database maintenance policies. This study focuses on deriving an optimal synchronization policy for consolidated database systems. By trading off synchronization cost with data staleness cost, we propose a time-based dynamic synchronization (TDS) policy, which evaluates the system state at predetermined checkpoints and synchronizes a database only if given staleness thresholds are crossed. Although several database synchronization and knowledge refreshment policies have been proposed in the prior literature, the TDS policy retains their advantages and overcomes their inflexibility in that, under the TDS policy, system check and synchronization are easy to schedule, disruptions to business operations can be avoided, and synchronization is run only if necessary. Experimental results show that the TDS policy consistently outperforms benchmark policies, leading to substantial cost savings. In particular, the performance gap between the TDS policy and a static periodic policy is greater when data changes arrive less frequently but carry higher unit staleness costs, or when queries arrive more frequently and suffer higher unit staleness costs.

Keywords: Consolidated database (CDB), database maintenance, database synchronization, Markov decision process, dynamic programming