複数データ構造を統合するトランザクション処理向け インデックス構築手法

Designing Database Indexes by Integrating Multiple Data Structures

Sho Nakazono^{1,a)} Tatsuhiro Nakamori^{2,b)} Hideyuki Kawashima^{2,c)}

Abstract: Database indexes process workloads that are mixtures of point and range queries. The optimal data structure is different for each workload, and a variety of indexes have been proposed. Concurrent B+ trees have been widely used in database systems, due to versatility in handling various workloads. However, concurrent B+ trees suffer from two problems: 1) low performance for point queries and 2) a performance penalty for range queries to avoid phantom anomalies. To solve these two problems, this paper proposes PLI, an optimization method for existing indexes such as concurrent B+ trees. PLI adds two data structures to an existing index: 1) a hash table for point queries and 2) a precision locking for avoiding phantom anomalies. PLI synchronizes all data structures so as to transparently handle them as a single database index. In this paper, we applied PLI to OpenBwTree, one of the state-of-the-art implementations of concurrent B+ tree. This optimization resulted in a performance improvement of about 2.5x for YCSB-A workloads and about 6.6x for YCSB-E workloads compared to the stand-alone OpenBwTree.

Keywords: Database Index, Transaction Processing, Phantom Anomaly

¹ NTT コンピュータ&データサイエンス研究所

NTT Computer and Data Science Laboratories

² 慶應義塾大学環境情報学部

Keio University

a) syou.nakazono.nu@hco.ntt.co.jp

b) t19601tn@sfc.keio.ac.jp

c) river@sfc.keio.ac.jp