



Título: Scan and join operators for asymmetric media

Data: 31/05/2019

Horário: 16:00h

Local: Sala de Seminários - Bloco 952

Resumo:

Solid state drives (SSDs) have emerged as an attractive alternative for storing large databases. A key reason for that is the fact that SSDs may provide IOPS rates up to two orders of magnitude greater than the rates delivered by HDDs. However, write operations consume more time (and have a higher cost) than read operations. On the other hand, database systems (DBSs) have been designed assuming read and write operations would be executed in the same amount of time (characteristic of HDDs). Thus, to fully exploit benefit provided by SSDs, components of DBSs should be aware of read/write asymmetry in SSDs. The join operation is the query operator which requires the highest amount of accesses (read/write operations) to secondary memory. In this paper, we present a new join algorithm, denoted DaC-Join whose key goal is to reduce the amount of write operations during the execution of algebraic join operations. In order to validate the novel approach, simulations with a prototype have been conducted over TPC-H benchmark database. The achieved results evidence the suitability of DaC-Join. For instance, DaC-Join can reduce up to 97% the amount of write operations w.r.t.

FlashJoin, a well-known join operator proposed to be deployed in SSDs, and, consequently, it can be up to 81% faster than FlashJoin.

Banca:

- Prof. Dr. José Maria da Silva Monteiro Filho(MDCC/UFC - Orientador)
- Prof. Dr. Angelo Roncalli Alencar Brayner (MDCC/UFC - Co-orientador)
- Prof. Dr. Jose de Aguiar Moraes Filho (Unifor)