Lecture #04: Database Storage (Part II)

15-445/645 Database Systems (Fall 2018)
https://15445.courses.cs.cmu.edu/fall2018/
Carnegie Mellon University
Prof. Andy Pavlo

1 Data Representation

This determines how a DBMS stores the actual bits for a value in-memory.

All integers are stored in their "native" C/C++ as specified by the IEEE-754 standard.

Variable Precision Numbers

- Inexact, variable-precision numeric type that uses the "native" C/C++ types specified by IEEE-754 standard.
- Faster than arbitrary precision numbers because the CPU can execute instructions on them directly.
- Example: `FLOAT`, `REAL`

Fixed Point Precision Numbers

- Numeric data types with arbitrary precision and scale. Typically stored in exact, variable-length binary representation with additional meta-data.
- Used when rounding errors are unacceptable.
- Example: `NUMERIC`, `DECIMAL`

2 Workloads

OLTP: On-line Transaction Processing

- Short lived txns
- Small footprint
- Repetitive operations
- Usually the kind of application that people build first

OLAP: On-line Analyitical Processing

- Long running queries
- Complex joins
- Exploratory queries

3 Storage Models

There are different ways to store tuples in pages. We have assumed the **n-ary storage model** so far.

N-Ary Storage Model (NSM)

The DBMS stores all of the attributes for a single tuple contiguously. Also known as a "row store". This approach is ideal for OLTP workloads where transactions tend to operate only an individual entity and insert heavy workloads.

Advantages:

- Fast inserts, updates, and deletes.
- Good for queries that need the entire tuple.

Disadvantages:

- Not good for scanning large portions of the table and/or a subset of the attributes. This is because it pollutes the buffer pool by fetching data that is not needed for processing the query.

There are two different ways to organize a NSM database:

- **Heap Organized Tables:** Tuples are stored in blocks called a heap, and the heap does not necessarily define an order.
- **Index-Organized Tables:** Tuples are stored in the primary key index itself, but different from a clustered index.

Decomposition Storage Model (DSM)

The DBMS stores a single attribute for all tuples contiguously in a block of data. Also known as a "column store". This model is ideal for OLAP workloads where read-only queries perform large scans over a subset of the table's attributes.

Advantages:

- Reduces the amount of wasted work during query execution because the DBMS only reads the data that it needs for that query.
- Enables better compression because all of the values for the same attribute are store contiguously in a single column.

Disadvantages:

- Slow for point queries, inserts, updates, and deletes because of tuple splitting/stitching.