Carnegie Mellon University ADVANCED DATABASE SYSTEMS

 \bigcirc

tu

 \bigcirc

#FREB PAVLO

Storage Models & Data Layout

@Andy_Pavlo // 15-721 // Spring 2019

DATA ORGANIZATION



DATA ORGANIZATION

One can think of an in-memory database as just a large array of bytes.

- \rightarrow The schema tells the DBMS how to convert the bytes into the appropriate type.
- \rightarrow Each tuple is prefixed with a header that contains its meta-data.

Storing tuples with as fixed-length data makes it easy to compute the starting point of any tuple.



TODAY'S AGENDA

Type Representation Data Layout / Alignment Storage Models System Catalogs





DATA REPRESENTATION

INTEGER/BIGINT/SMALLINT/TINYINT

 \rightarrow C/C++ Representation

FLOAT/REAL vs. NUMERIC/DECIMAL

 \rightarrow IEEE-754 Standard / Fixed-point Decimals

TIME/DATE/TIMESTAMP

 \rightarrow 32/64-bit int of (micro/milli)seconds since Unix epoch

VARCHAR/VARBINARY/TEXT/BLOB

- \rightarrow Pointer to other location if type is \geq 64-bits
- → Header with length and address to next location (if segmented), followed by data bytes.

VARIABLE PRECISION NUMBERS

Inexact, variable-precision numeric type that uses the "native" C/C++ types.

Store directly as specified by <u>IEEE-754</u>.

Typically faster than arbitrary precision numbers. \rightarrow Example: FLOAT, REAL/DOUBLE



VARIABLE PRECISION NUMBERS

Rounding Example

#include <stdio.h>

Output

x+y	=	0.30000001192092895508
0.3	=	0.299999999999999998890

int main(int argc, char* argv[]) {
 float x = 0.1;
 float y = 0.2;
 printf("x+y = %.20f\n", x+y);
 printf("0.3 = %.20f\n", 0.3);



FIXED PRECISION NUMBERS

Numeric data types with arbitrary precision and scale. Used when round errors are unacceptable. \rightarrow Example: NUMERIC, DECIMAL

Typically stored in a exact, variable-length binary representation with additional meta-data. \rightarrow Like a VARCHAR but not stored as a string

Postgres Demo



POSTGRES: NUMERIC









DATA LAYOUT





VARIABLE-LENGTH FIELDS



VARIABLE-LENGTH FIELDS



NULL DATA TYPES

Choice #1: Special Values

→ Designate a value to represent NULL for a particular data type (e.g., INT32_MIN).

Choice #2: Null Column Bitmap Header

→ Store a bitmap in the tuple header that specifies what attributes are null.

Choice #3: Per Attribute Null Flag

- \rightarrow Store a flag that marks that a value is null.
- \rightarrow Have to use more space than just a single bit because this messes up with word alignment.



NULL DATA TYPES

Integer l	Numbers
-----------	---------

Data Type	Size	Size (Not Null)	Synonyms	Min Value	Max Value
BOOL	2 bytes	1 byte	BOOLEAN	0	1
BIT	9 bytes	8 bytes			
TINYINT	2 bytes	1 byte		-128	127
SMALLINT	4 bytes	2 bytes		-32768	32767
MEDIUMINT	4 bytes	3 bytes		-8388608	8388607
INT	8 bytes	4 bytes	INTEGER	-2147483648	2147483647
BIGINT	12 bytes	8 bytes		-2 ** 63	(2 ** 63) - 1
				0	

messes up with word alignment.



DISCLAIMER

The truth is that you only need to worry about word-alignment for cache lines (e.g., 64 bytes).

I'm going to show you the basic idea using 64-bit words since it's easier to see...





ABASE GROUP

WORD-ALIGNED TUPLES

14

```
CREATE TABLE AndySux (

32-bits id INT PRIMARY KEY,

cdate TIMESTAMP,

color CHAR(2),

zipcode INT

);

CREATE TABLE AndySux (

cdate TIMESTAMP,

color CHAR(2),

zipcode INT

);
```



TABASE GROUP

WORD-ALIGNED TUPLES

14





All attributes in a tuple must be word aligned to enable the CPU to access it without any unexpected behavior or additional work.

```
CREATE TABLE AndySux (
32-bits id INT PRIMARY KEY,
64-bits cdate TIMESTAMP,
16-bits color CHAR(2),
32-bits zipcode INT
);
```

ATABASE GROUP



Approach #1: Perform Extra Reads

 \rightarrow Execute two reads to load the appropriate parts of the data word and reassemble them.

Approach #2: Random Reads

 \rightarrow Read some unexpected combination of bytes assembled into a 64-bit word.

Approach #3: Reject

 \rightarrow Throw an exception

Source: Levente Kurusa

DATABASE GROUP



WORD-ALIGNMENT: PADDING

Add empty bits after attributes to ensure that tuple is word aligned.

```
CREATE TABLE AndySux (
32-bits id INT PRIMARY KEY,
64-bits cdate TIMESTAMP,
16-bits color CHAR(2),
32-bits zipcode INT
);
```



WORD-ALIGNMENT: REORDERING

Switch the order of attributes in the tuples' physical layout to make sure they are aligned. \rightarrow May still have to use padding.



CMU-DB ALIGNMENT EXPERIMENT

Processor: 1 socket, 4 cores w/ 2×HT Workload: Insert Microbenchmark

Avg. Throughput

No Alignment0.523 MB/secOptimization #111.7 MB/secOptimization #2814.8 MB/sec

Source: Tianyu Li

CARNEGIE MELLON DATABASE GROUP

STORAGE MODELS

N-ary Storage Model (NSM) Decomposition Storage Model (DSM) Hybrid Storage Model



DATABASE GROUP



CARNEGIE MELLON DATABASE GROUP

STORAGE N

N-ary Storage Model (NSM) Decomposition Storage Mod Hybrid Storage Model





ABSTRACT

There has been a significant amount of excitement and recent work on column-oriented database systems ("column-stores"). These database systems have been shown to perform more than an order of magnitude better than traditional row-oriented database systerns ("row-stores") on analytical workloads such as those found in data warehouses, decision support, and business intelligence applications. The elevator pitch behind this performance difference is straightforward: column-stores are more I/O efficient for read-only queries since they only have to read from disk (or from memory) those attributes accessed by a query.

This simplistic view leads to the assumption that one can obtain the performance benefits of a column-store using a row-store: either by vertically partitioning the schema, or by indexing every column so that columns can be accessed independently. In this paper, we demonstrate that this assumption is false. We compare the performance of a commercial row-store under a variety of different configurations with a column-store and show that the row-store performance is significantly slower on a recently proposed data warehouse benchmark. We then analyze the performance difference and show that there are some important differences between the two systems at the query executor level (in addition to the obvious differences at the storage layer level). Using the column-store, we then tease apart these differences, demonstrating the impact on performance of a variety of column-oriented query execution techniques, including vectorized query processing, compression, and a new join algorithm we introduce in this paper. We conclude that while it is not impossible for a row-store to achieve some of the performance advantages of a column-store, changes must be made to both the storage layer and the query executor to fully obtain the benefits of a column-oriented approach.

Categories and Subject Descriptors

H.2.4 [Database Management]: Systems-Query processing, Re-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific SIGMOD 08, June 9-12, 2008, Vancouver, BC, Canada,

Copyright 2008 ACM 978-1-60558-102-6/08/06 ...\$5.00

General Terms

Experimentation, Performance, Measurement

Keywords

C-Store, column-store, column-oriented DBMS, invisible join, compression, tuple reconstruction, tuple materialization.

1. INTRODUCTION

Recent years have seen the introduction of a number of columnoriented database systems, including MonetDB [9, 10] and C-Store [22]. The authors of these systems claim that their approach offers orderof-magnitude gains on certain workloads, particularly on read-intensive analytical processing workloads, such as those encountered in data

Indeed, papers describing column-oriented database systems usually include performance results showing such gains against traditional, row-oriented databases (either commercial or open source). These evaluations, however, typically benchmark against row-oriented systems that use a "conventional" physical design consisting of a collection of row-oriented tables with a more-or-less one-to-one mapping to the tables in the logical schema. Though such results clearly demonstrate the potential of a column-oriented approach, they leave open a key question: Are these performance gains due to something fundamental about the way column-oriented DBMSs are internally architected, or would such gains also be possible in a conventional system that used a more column-oriented physical

Often, designers of column-based systems claim there is a fundamental difference between a from-scratch column-store and a rowstore using column-oriented physical design without actually exploring alternate physical designs for the row-store system. Hence, one goal of this paper is to answer this question in a systematic way. One of the authors of this paper is a professional DBA specializing in a popular commercial row-oriented database. He has carefully implemented a number of different physical database designs for a recently proposed data warehousing benchmark, the Star Schema Benchmark (SSBM) [18, 19], exploring designs that are as "column-oriented" as possible (in addition to more traditional de-

Vertically partitioning the tables in the system into a collection of two-column tables consisting of (table key, attribute) pairs, so that only the necessary columns need to be read to answer a query.

Using index-only plans; by creating a collection of indices that cover all of the columns used in a query, it is possible

N-ARY STORAGE MODEL (NSM)

The DBMS stores all of the attributes for a single tuple contiguously.

Ideal for OLTP workloads where txns tend to operate only on an individual entity and insert-heavy workloads.

Use the tuple-at-a-time iterator model.



NSM: PHYSICAL STORAGE

Choice #1: Heap-Organized Tables

- \rightarrow Tuples are stored in blocks called a heap.
- \rightarrow The heap does not necessarily define an order.

Choice #2: Index-Organized Tables

- \rightarrow Tuples are stored in the primary key index itself.
- \rightarrow Not quite the same as a clustered index.



N-ARY STORAGE MODEL (NSM)

Advantages

- \rightarrow Fast inserts, updates, and deletes.
- \rightarrow Good for queries that need the entire tuple.
- \rightarrow Can use index-oriented physical storage.

Disadvantages

 \rightarrow Not good for scanning large portions of the table and/or a subset of the attributes.



DECOMPOSITION STORAGE MODEL (DSM)

The DBMS stores a single attribute for all tuples contiguously in a block of data.

Ideal for OLAP workloads where read-only queries perform large scans over a subset of the table's attributes.

Use the vector-at-a-time iterator model.



DECOMPOSITION STORAGE MODEL (DSM)

1970s: Cantor DBMS **1980s:** DSM Proposal **1990s:** SybaseIQ (in-memory only) 2000s: Vertica, Vectorwise, MonetDB 2010s: "The Big Three" Cloudera Impala, Amazon Redshift, SAP HANA, MemSQL, Clickhouse, LinkedIn Pinot, and most OLAP systems



DSM: TUPLE IDENTIFICATION

Choice #1: Fixed-length Offsets

 \rightarrow Each value is the same length for an attribute.

Choice #2: Embedded Tuple Ids

 \rightarrow Each value is stored with its tuple id in a column.



Embedded Ids





DSM: QUERY PROCESSING

Late Materialization Columnar Compression Block/Vectorized Processing Model



DECOMPOSITION STORAGE MODEL (DSM)

Advantages

- \rightarrow Reduces the amount wasted work because the DBMS only reads the data that it needs.
- \rightarrow Better compression.

Disadvantages

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



OBSERVATION

Data is "hot" when first entered into database \rightarrow A newly inserted tuple is more likely to be updated again the near future.

As a tuple ages, it is updated less frequently. \rightarrow At some point, a tuple is only accessed in read-only queries along with other tuples.

What if we want to use this data to make decisions that affect new txns?

HYBRID STORAGE MODEL

Single logical database instance that uses different storage models for hot and cold data.

Store new data in NSM for fast OLTP Migrate data to DSM for more efficient OLAP



HYBRID STORAGE MODEL

Choice #1: Separate Execution Engines

 \rightarrow Use separate execution engines that are optimized for either NSM or DSM databases.

Choice #2: Single, Flexible Architecture

 \rightarrow Use single execution engine that is able to efficiently operate on both NSM and DSM databases.



SEPARATE EXECUTION ENGINES

Run separate "internal" DBMSs that each only operate on DSM or NSM data.

- \rightarrow Need to combine query results from both engines to appear as a single logical database to the application.
- \rightarrow Have to use a synchronization method (e.g., 2PC) if a txn spans execution engines.

Two approaches to do this:

- \rightarrow **Fractured Mirrors** (Oracle, IBM)
- \rightarrow **<u>Delta Store</u>** (SAP HANA)



FRACTURED MIRRORS

Store a second copy of the database in a DSM layout that is automatically updated.
→ All updates are first entered in NSM then eventually copied into DSM mirror.



DELTA STORE

Stage updates to the database in an NSM table. A background thread migrates updates from delta store and applies them to DSM data.





CATEGORIZING DATA

Choice #1: Manual Approach

 \rightarrow DBA specifies what tables should be stored as DSM.

Choice #2: Off-line Approach

 \rightarrow DBMS monitors access logs offline and then makes decision about what data to move to DSM.

Choice #3: On-line Approach

 \rightarrow DBMS tracks access patterns at runtime and then makes decision about what data to move to DSM.



PELOTON ADAPTIVE STORAGE

Employ a single execution engine architecture that is able to operate on both NSM and DSM data. \rightarrow Don't need to store two copies of the database. \rightarrow Don't need to sync multiple database segments.

Note that a DBMS can still use the delta-store approach with this single-engine architecture.

BRIDGING THE ARCHIPELAGO BETWEEN ROW-STORES AND COLUMN-STORES FOR HYBRID WORKLOADS SIGMOD 2016



PELOTON ADAPTIVE STORAGE



TILE ARCHITECTURE



TILE ARCHITECTURE



DATABASE GROUP

TILE ARCHITECTURE



SELECT AVG(B) FROM AndySux WHERE C = "yyy"В 2 1 2 3

TILE ARCHITECTURE







SYSTEM CATALOGS

Almost every DBMS stores their a database's catalog in itself.

- \rightarrow Wrap object abstraction around tuples.
- \rightarrow Specialized code for "bootstrapping" catalog tables.

The entire DBMS should be aware of transactions in order to automatically provide ACID guarantees for DDL commands and concurrent txns.



SCHEMA CHANGES

ADD COLUMN:

- \rightarrow **NSM**: Copy tuples into new region in memory.
- \rightarrow **DSM**: Just create the new column segment

DROP COLUMN:

- \rightarrow NSM #1: Copy tuples into new region of memory.
- → NSM #2: Mark column as "deprecated", clean up later.
- \rightarrow **DSM**: Just drop the column and free memory.

CHANGE COLUMN:

 \rightarrow Check whether the conversion is allowed to happen. Depends on default values.

INDEXES

CREATE INDEX:

- \rightarrow Scan the entire table and populate the index.
- \rightarrow Have to record changes made by txns that modified the table while another txn was building the index.
- \rightarrow When the scan completes, lock the table and resolve changes that were missed after the scan started.

DROP INDEX:

- \rightarrow Just drop the index logically from the catalog.
- \rightarrow It only becomes "invisible" when the txn that dropped it commits. All existing txns will still have to update it.



SEQUENCES

Typically stored in the catalog. Used for maintaining a global counter \rightarrow Also called "auto-increment" or "serial" keys

Sequences are not maintained with the same isolation protection as regular catalog entries.

- \rightarrow Rolling back a txn that incremented a sequence does not rollback the change to that sequence.
- \rightarrow All **INSERT** queries would incur write-write conflicts.



PARTING THOUGHTS

We abandoned the hybrid storage model

- \rightarrow Significant engineering overhead.
- \rightarrow Delta version storage + column store is almost equivalent.

Catalogs are hard.

