

CSE 544

Principles of Database Management Systems

Magdalena Balazinska


Winter 2009

Lecture 15 - Data Warehousing: CStore



References

- **Column-Stores vs Row-Stores: How Different Are They Really?** D. Abadi, S. Madden, and N. Hachem. SIGMOD'08
- Sections on the final: 1, 4, 5.1-5.3, and 6

From Row-Store to Column-Store



Rows stored
contiguously on disk
(+ tuples headers)



Columns stored
contiguously on disk
(no headers needed)₃

More Detailed Example

Row-based
(4 pages)

Page {

A	1
A	2
A	2
A	2
B	2
B	4
C	4
C	4

Column-based
(4 pages)

A	1
A	2
A	2
A	2
B	2
B	4
C	4
C	4

} Page

C-Store also
avoids large
tuple headers


Column-Store Optimizations

- **Late tuple materialization** (3X improvement)
 - Process individual columns as long as possible
 - Merge columns into complete tuples as late as possible
- **Block iteration** (1.5X)
 - Pass blocks of values between ops instead of individual tuples
- **Compression**: e.g., run-length encoding of columns (10X)
- **Invisible joins** (1.5X)


Late Tuple Materialization

Ex: SELECT R.b from R where R.a=X and R.d=Y

Early materialization



Late materialization



Compression Example

Row-based
(4 pages)

Page {

A	1
A	2
A	2
A	2
B	2
B	4
C	4
C	4

Column-based
(4 pages)

A	1
A	2
A	2
A	2
B	2
B	4
C	4
C	4

Compressed
(2 pages)

4XA	1X1
2XB	4X2
2XC	5X4

Page

Simulating a Column-Store DBMS in a Row-Store DBMS

- Vertical partitioning
 - Two-column tables: (key, attribute)
- Index-only plans
 - Create a B+ tree index on each attribute
 - Answer queries using indexes only, without reading actual data
- Materialized views
 - Each view contains a subset of columns

Performance Comparison

- See Section 6 in the paper

Conclusion

- Column-store DBMS outperforms row-store DBMS
 - Measured on a data warehousing benchmark (SSBM)
- Late materialization and compression are key factors
- Difficult to simulate a column-store in a row-store
 - Tuple overheads cause data blow-up
 - Column joins are expensive
 - Hard to get the DBMS to “do the right thing” (e.g., index plans)
- Not the end of the story, however, ... see CIDR'09 paper