


Data Queries: The Basics


Row Store:

- Transactional, fast processing



Column Store:

- Analytical, fewer tables



Hybrid Systems




HTAP Systems:

- Finds the ideal layout between row and column processing
- Faces massive bookkeeping costs

Ephemeral Variables

Ephemeral Variables:

- RM uses to find the ideal layout
- Represents data as CPU addresses (minimized data movement)



Query Details

Row Size: 4 - 524288 bytes (2^n)


Row Count: 524288 - 4 bytes (2^{m-n})

Column Width: 1 byte

Table Size: 2 megabytes (2097152 bytes)

Selectivity: None

(These ones)



Execution:

- Access the first column only in tables of various dimensions
- Begins at 4 x 524288, changes by powers of 2


Relational Memory

Hot vs Cold:

- Hot uses data in Relational Buffer for faster variable representation
- Data is gathered through sequential queries
- Cold lacks this data

Execution Speed (Normalized to Row Store)

Effect of Type of Access on Row Size vs Runtime




Discussion

Important Factors:

- RM's execution speed is 30% faster than column-store
- Row-store's major spike in execution speed and miss frequency after 256 byte rows
- Exponential RM miss ratio growth

Miss Ratio in the 1st Cache

l1 refills/references



References

[1] Shahin Roozkhosh, Denis Hoornaert, Ju Hyung Mun, Tarikul Islam Papon, Ahmed Sanallah. Ulrich Drepper, Renato Mancuso, and Manos Athanassoulis. 2021. Relational Memory: Native In-Memory Accesses on Rows and Columns. CORR, abs/2109.14349. <https://arxiv.org/abs/2109.14349>

[2] "Hyperdata Computing." Hyperdatacomputing, 31 Mar. 2022, www.hyperdatacomputing.com/.

[3] Dittrich, Prof. Jens, et al. "OctopusDB." Information Systems Group - Prof. Jens Dittrich, bigdata.uni-saarland.de/projects/octopusdb.php.

