Lazy users and agile databases: toward timely, predictable and cheap data analytics

ecocloud, an EPFL research center

Renata Borovica-Gajic and Anastasia Ailamaki

The Need for Agile Approach in building DBMS

What every business analyst wants

Predictably fast (i.e., interactive) querying with low initial investment

Why timely and interactive?

Revenue depends upon ability to make decisions at real-time

Why predictable?

When an analyst submits a query, s/he does not know whether to :

- Wait for the response
- Go for coffee/dinner
- Go home for the night

[Guy Lohman, Blink]

Why cost-efficient?

Infrastructure is expensive

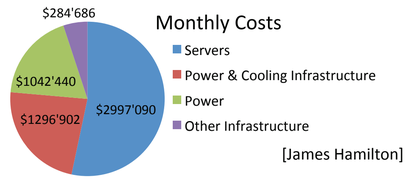

[James Hamilton]

Companies spend $0.30/GB/Month to store their data

Challenges for DBMS

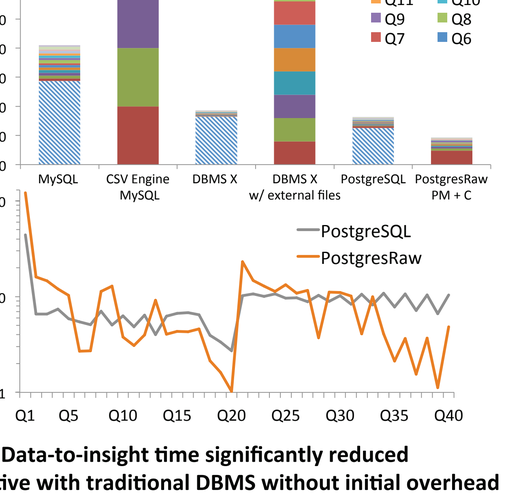

With ever bigger data and more complex ad-hoc workloads, the business analyst's goal is becoming unattainable

Timely, interactive, workload-driven data exploration

Time to first insight

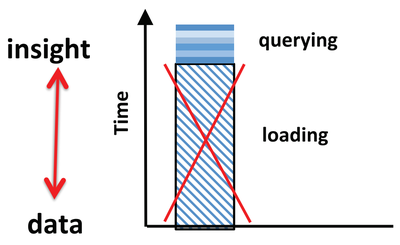

Long data-to-insight time due to high startup cost

Issues for data exploration:

- Why wait for loading and tuning?
- Why load the entire dataset?
- Hire DB expert to tune the system?

Skipping data loading in favor of raw data processing

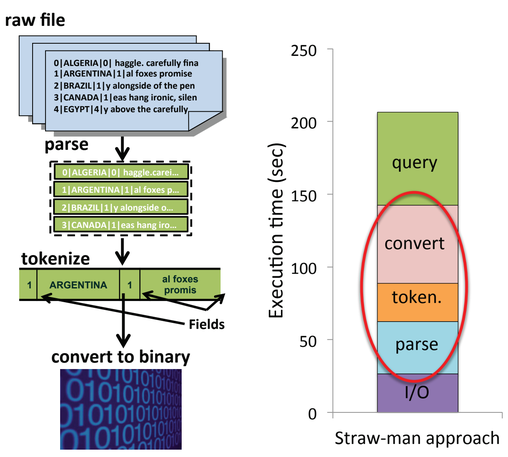

Too slow with current technology

Tuning as a side effect of workload execution in PostgresRaw

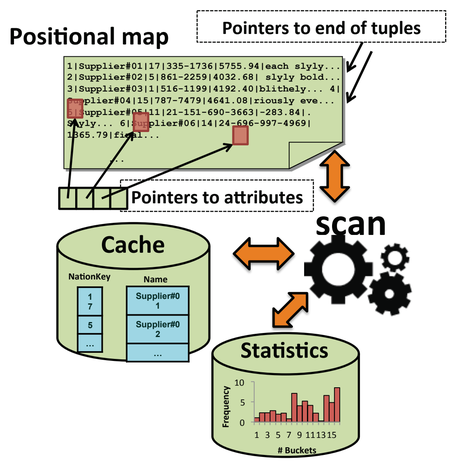

Seamlessly adjusting to queries
Progressively cheaper access

PostgresRaw in Action

Data-to-insight time significantly reduced
Competitive with traditional DBMS without initial overhead

Predictable query performance with adaptive, data-driven access operators

State of affairs in database systems

Setting: TPC-H, SF10, DBMS-X, Tuning tool 5GB space

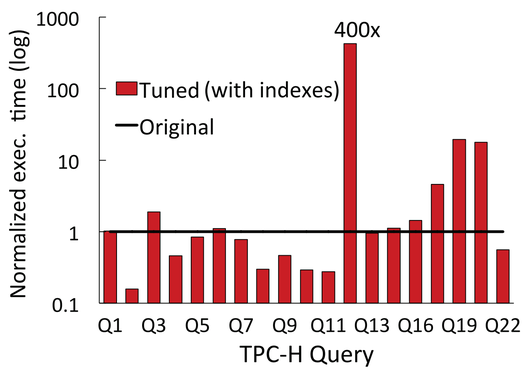

Degradation due to sub-optimal access path choices

Quest for robust & predictable execution

Statistics - unreliable advisor
Reoptimization - risky

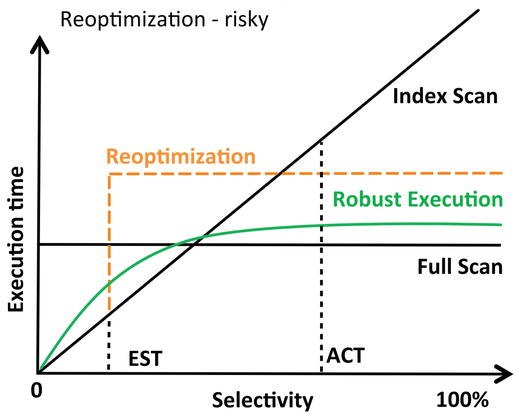

Robust = Near-optimal execution for all inputs

Region Snooping with Smooth Scan

INSIGHT: **Morph** between Index and Sequential Scan based on observed result distribution

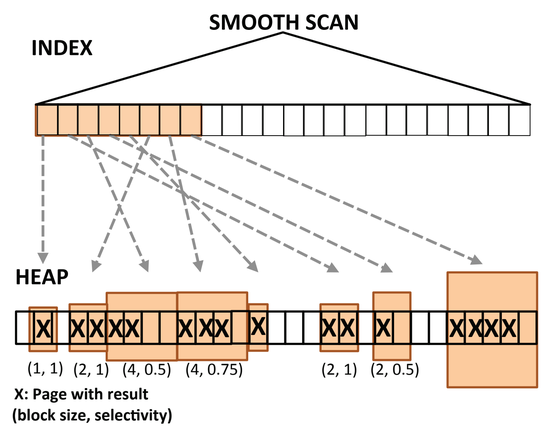

Data driven adaptation = near-optimal performance

Smooth Scan in Action

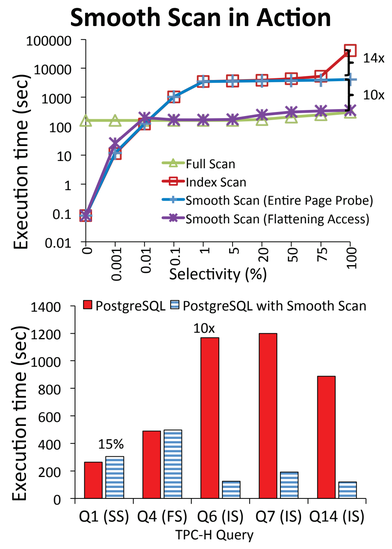

Robust & predictable performance

Data analytics for a penny

Database storage tiering

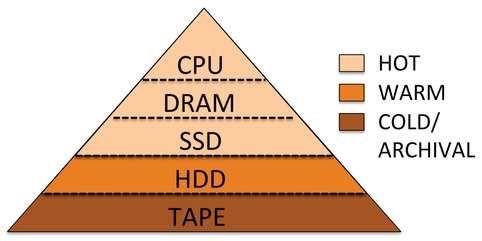

Data waterfall reduces capital and operational expenses

Cold storage in the tiering hierarchy

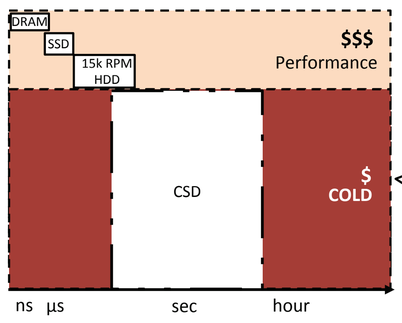

Can cold storage subsume archival and capacity tiers?

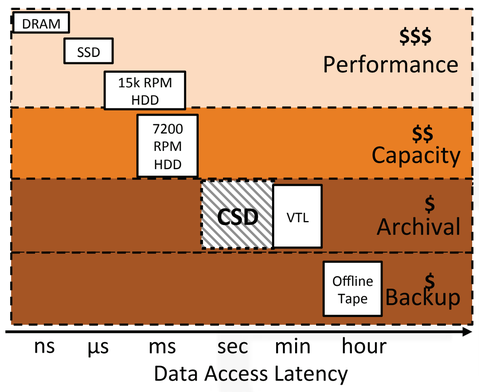

CSD = price of tapes, latency of HDD

CSD in capacity tier

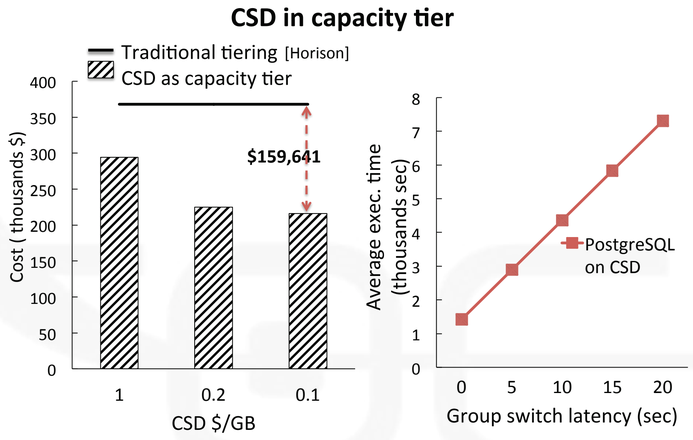

Major cost savings with CSD, BUT performance hit due to group switch latency

Skipper: adaptive CSD-driven execution

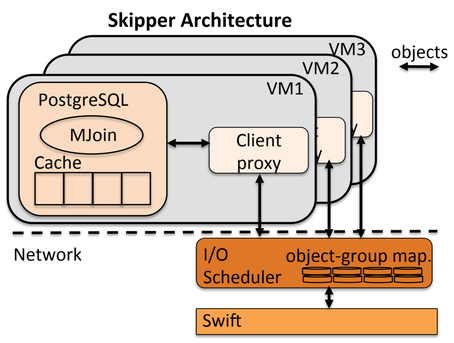

Support for out-of-order execution

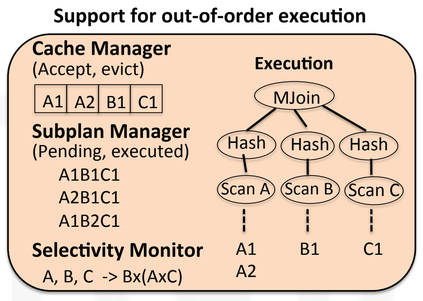

Out-of-order execution coupled with smart caching and I/O scheduling policies

Skipper in Action

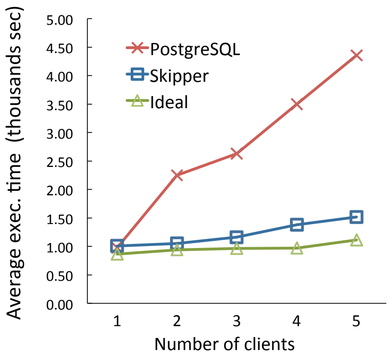

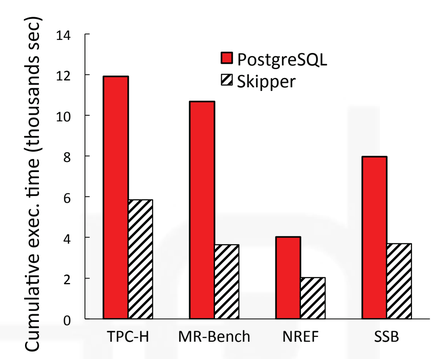

Skipper approximates ideal performance