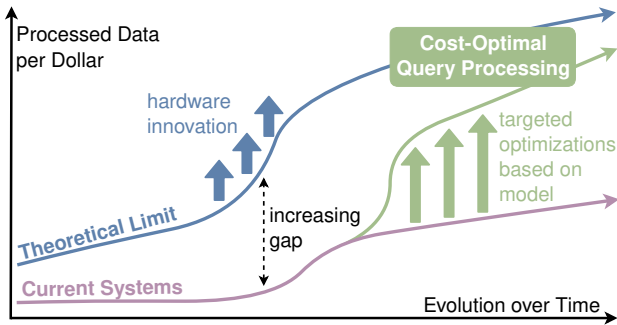


# Towards Cost-Optimal Query Processing in the Cloud

Viktor Leis, Maximilian Kuschewski

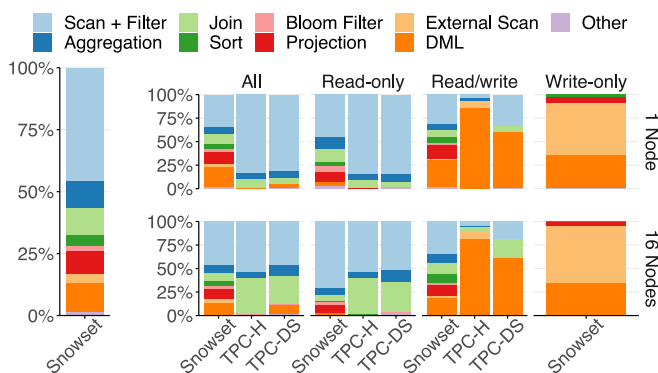
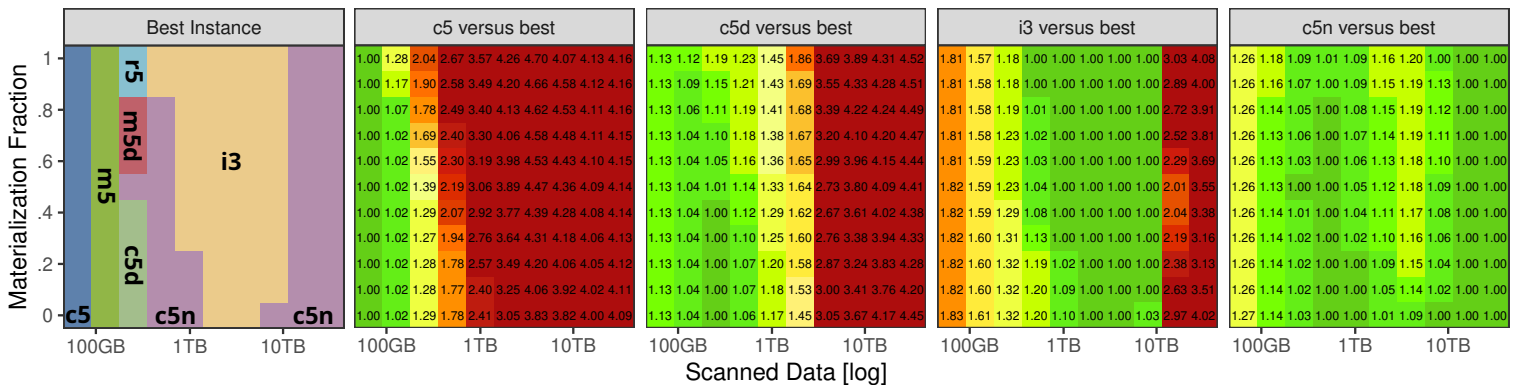
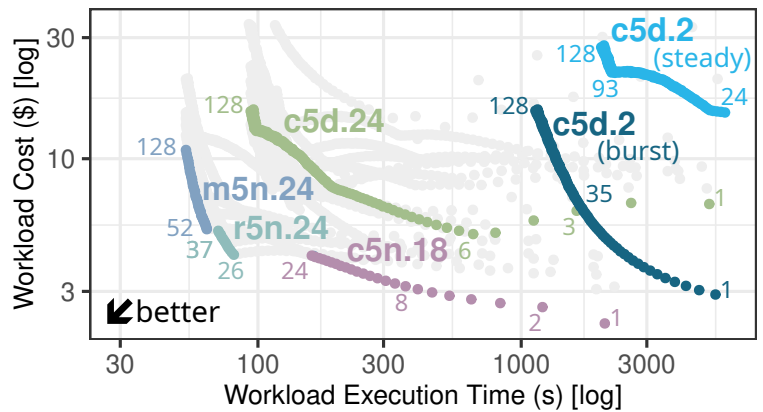


## Our Vision

- cloud analytics as a commodity
- workload cost as the primary metric
- model points way to cost-optimality
- per-workload hardware choice

## Model Predictions

- instance choice has a huge impact
- fast networking is a game changer
- scale up first, then scale out
- current systems are inefficient



## Snowset Operators

- almost 50% is spent in scan
- DML is important and complex
- no sort in TPC-\*
- aggregation is equal to join time