Lec 24 OLAP, Distributed

OLAP : long running
complex } queries
explaratory

Common workflow

OLTP DB accumulate data → Extract Transform Load → OLAP DB

" ETL", sometimes ppl do ELT

Decision Support Sys — analyse the data to inform decisions or even create ML dataset

Schemas

▷ Star Schema

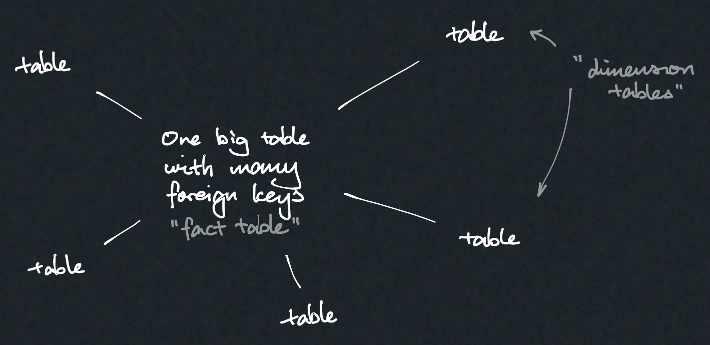

▷ Snowflake Schema

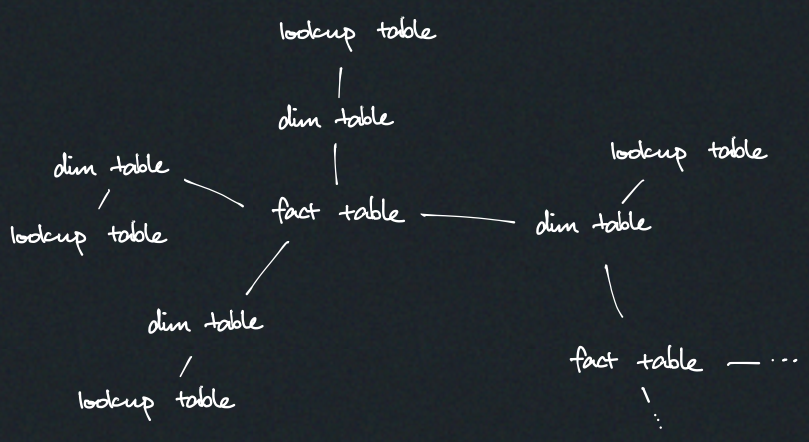

Dist. Query Execution

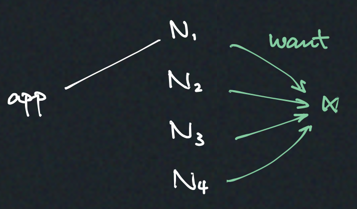

but :
- data might not fit on one node
- want parallel

Approaches : → Push query to data
  └ Send query to where the data is
→ Pull data to query
  └ Bring needed data to processor running the query

Optimisation : - Try filter before sending data
- Cache data from network in buf pool

If crash ... - Run again ...?
- Save intermediate result

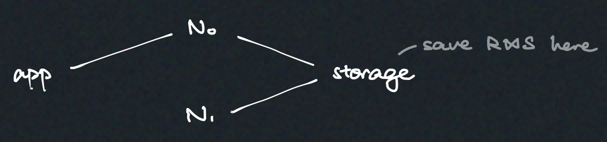

Query Planning

Need to consider data movement when optimising

Approaches : → Send physical operators ← more common
→ Send SQL queries ← works

Join

select * from R join S on R.id = S.id , 2 nodes

1. If S replicated : do local join and concat result
2. If S, R partitioned by same ranges : again do local join
3. If S, R partitioned differently and R part. by join key
   broadcast to case 1 to have S copied on both
4. Else — neither S, R part. by join key
   split up networking to go to case 2 (slide 25)
   like hash join, just with hashed partitions on nodes

▷ Semi-Join

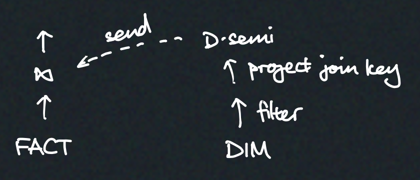

Cloud Systems

Approaches : → Managed DBMS — control hardware, install DB software , provide API
→ Cloud Native — build the software system specifically for serving
→ Serverless — not having hardware, just storage, but have on-demand compute node

Data Lakes

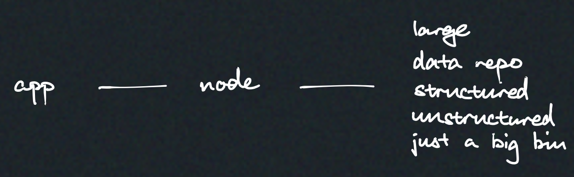

Misc

* Modularality — catalogue module
  optimiser ..
  exe. engine ..
  data format
  └ Parquet
  └ Arrow
  └ HDF5
  ⋮

Plug and play DB components !