An SQL-based approach to Physics Analysis

M. Limper

CERNopenlab

Intro

***The challenge*: stress a database engine by forcing it to do physics analysis:**

- Store analysis data in relational database
- Complicated SQL queries
- Calls to external C++ libraries
- Let the database take care of parallelism

Analysis data in a relational database

- Test with root-ntuples (D3PDs) produced for ATLAS top-physics group
  - Sub-set of 7.2 million events (27 ntuples)
  - ~4000 variables ("branches") stored in event-tree

- DB design uses different tables for different physics-objects
  - Many columns per table

DATA12_8TEV			
Table name	***columns***	***M rows***	***size in GB***
photon	216	89.9	**114.4**
electron	340	49.5	**94.6**
jet	171	26.8	**26.3**
muon	251	7.7	**14.2**
primary_vertex	25	89.5	**11.9**
EF (trigger)	490	7.2	**7.9**
MET_RefFinal	62	6.6	**2.3**
eventData	52	7.2	**1.4**
	1607		**272.9**

Physics Analysis in SQL

Make temporary tables using the WITH-AS statement:

WITH goodmuons AS (SELECT … FROM muon WHERE pt>25.)

JOIN statements on the RunNumber,EventNumber put information from the different selections together:

SELECT … FROM good_muons INNER JOIN good_bjets USING (RunNumber,EventNumber) WHERE goodmuons.N=2 AND goodbjets.N=2

- ✓ ***Simple calculations were written in (PL/)SQL***
- ✓ ***Code from external C++ libraries was used for more complicated calculations***

Test setup

5 nodes with 40 cores total
5 disk arrays each with 12 disks
max. I/O reads of ~2500 MB/s

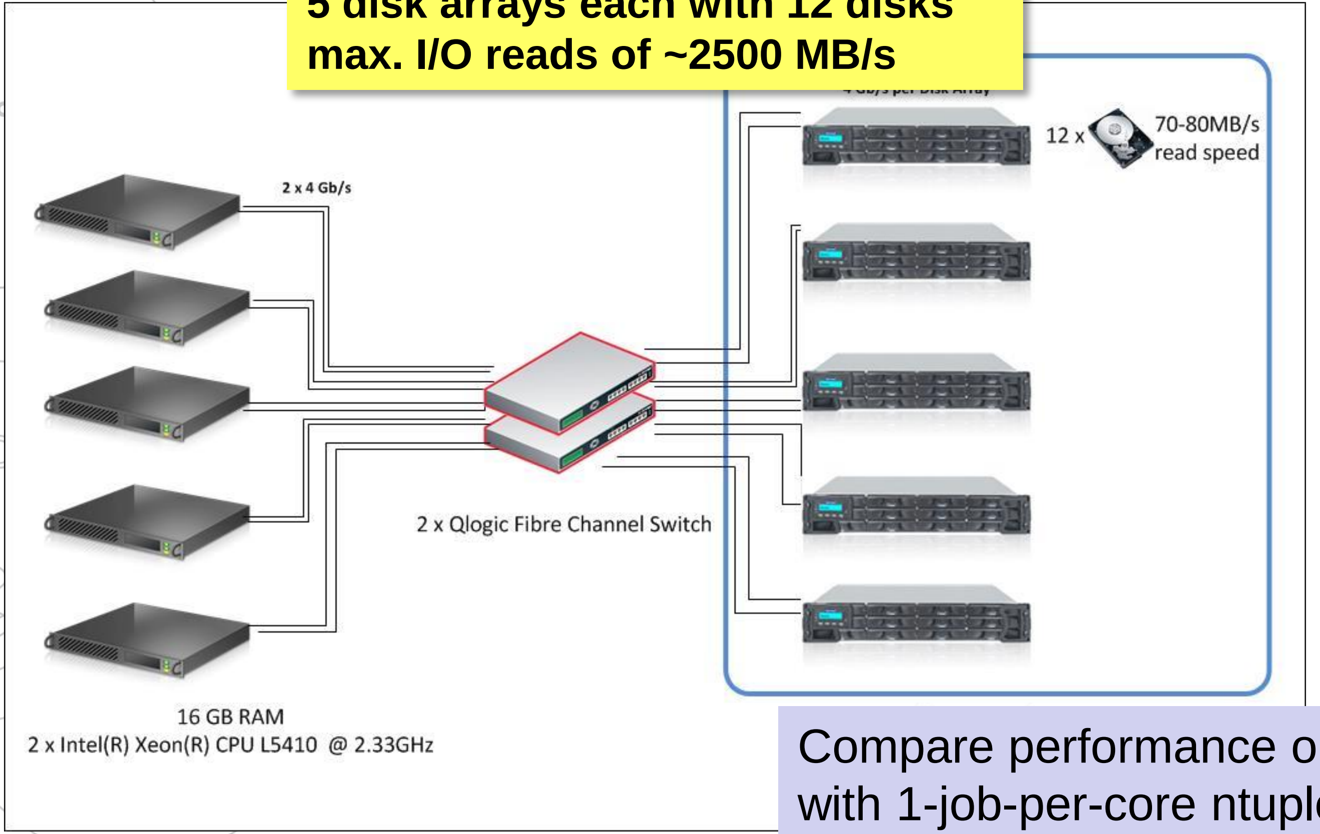

Compare performance on Oracle RAC with 1-job-per-core ntuple-analysis

Benchmark 1.

Simplified Higgs+Z: compare simple root-macro with SQL-query returning same results

- In both cases limited by iowait !
- I/O reads for root-ntuple analysis 10x less than for DB

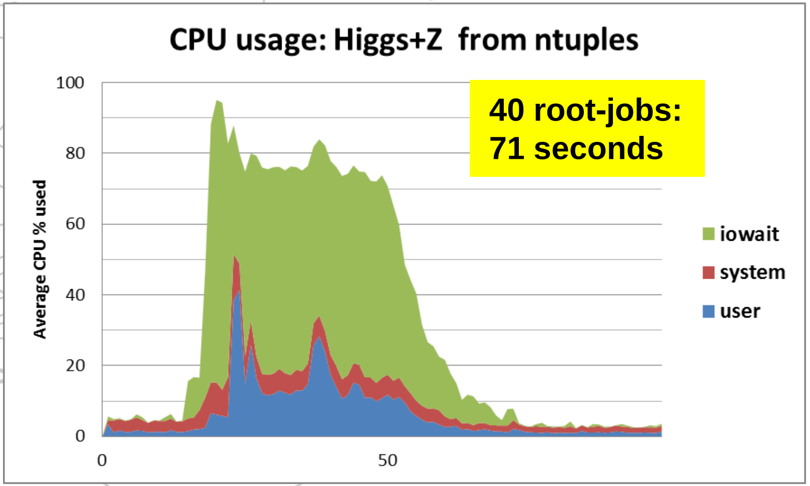

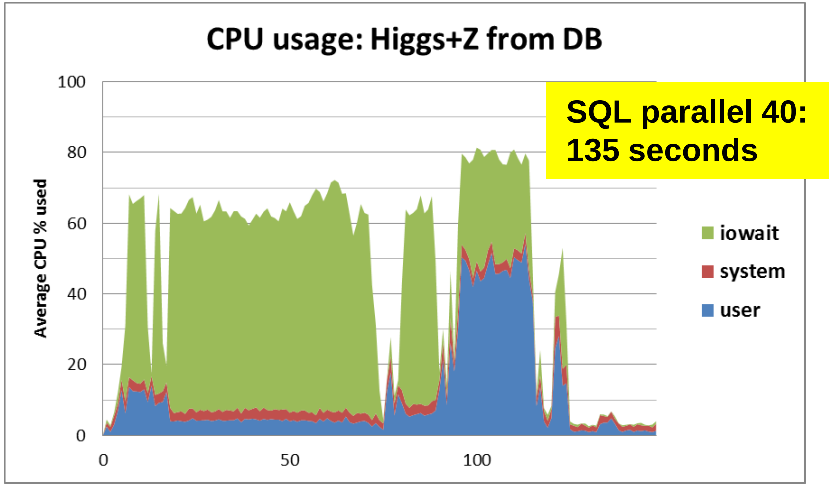

Benchmark 2.

Ttbar cutflow: compare existing ‘root-core’ packages with modified version that constructs SQL-query

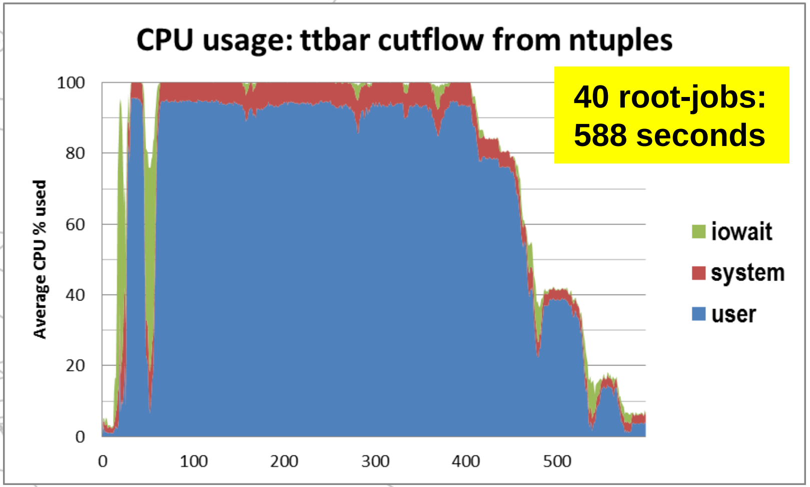

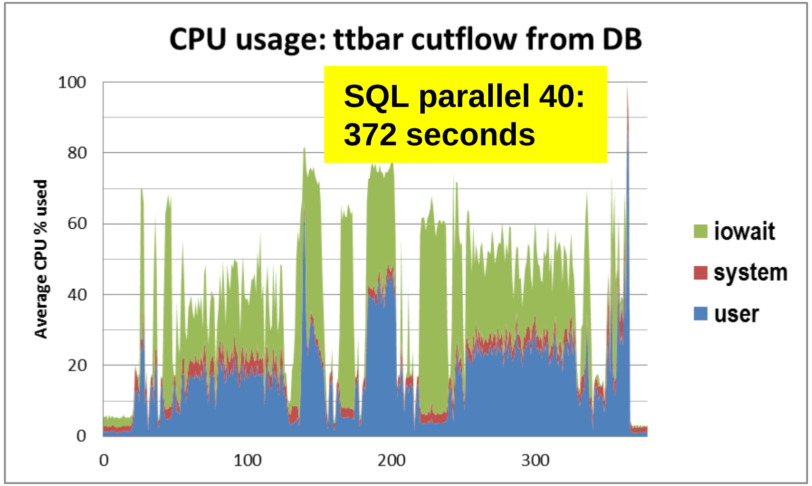

Conclusion

SQL-based physics analysis using data stored in a relational database could reproduce results from root-ntuple analysis'

- Database takes care of parallelism
- Row-based storage in combination with wide tables limits performance by the I/O read speed of the system