CORADD: Correlation Aware Database Designer for Materialized Views & Indexes

Hideaki Kimura 

With: George Huo  , Alex Rasin 

Sam Madden  , Stan Zdonik 

Intro: Data-warehouse (OLAP)

- Very Large Table
  - Million-Billion tuples
- Analytical Query
  - Not Very Selective Predicates
- Huge Market
  - Several $billion/year

Goal: *Speed Up OLAP Queries*

Problem: Secondary Index

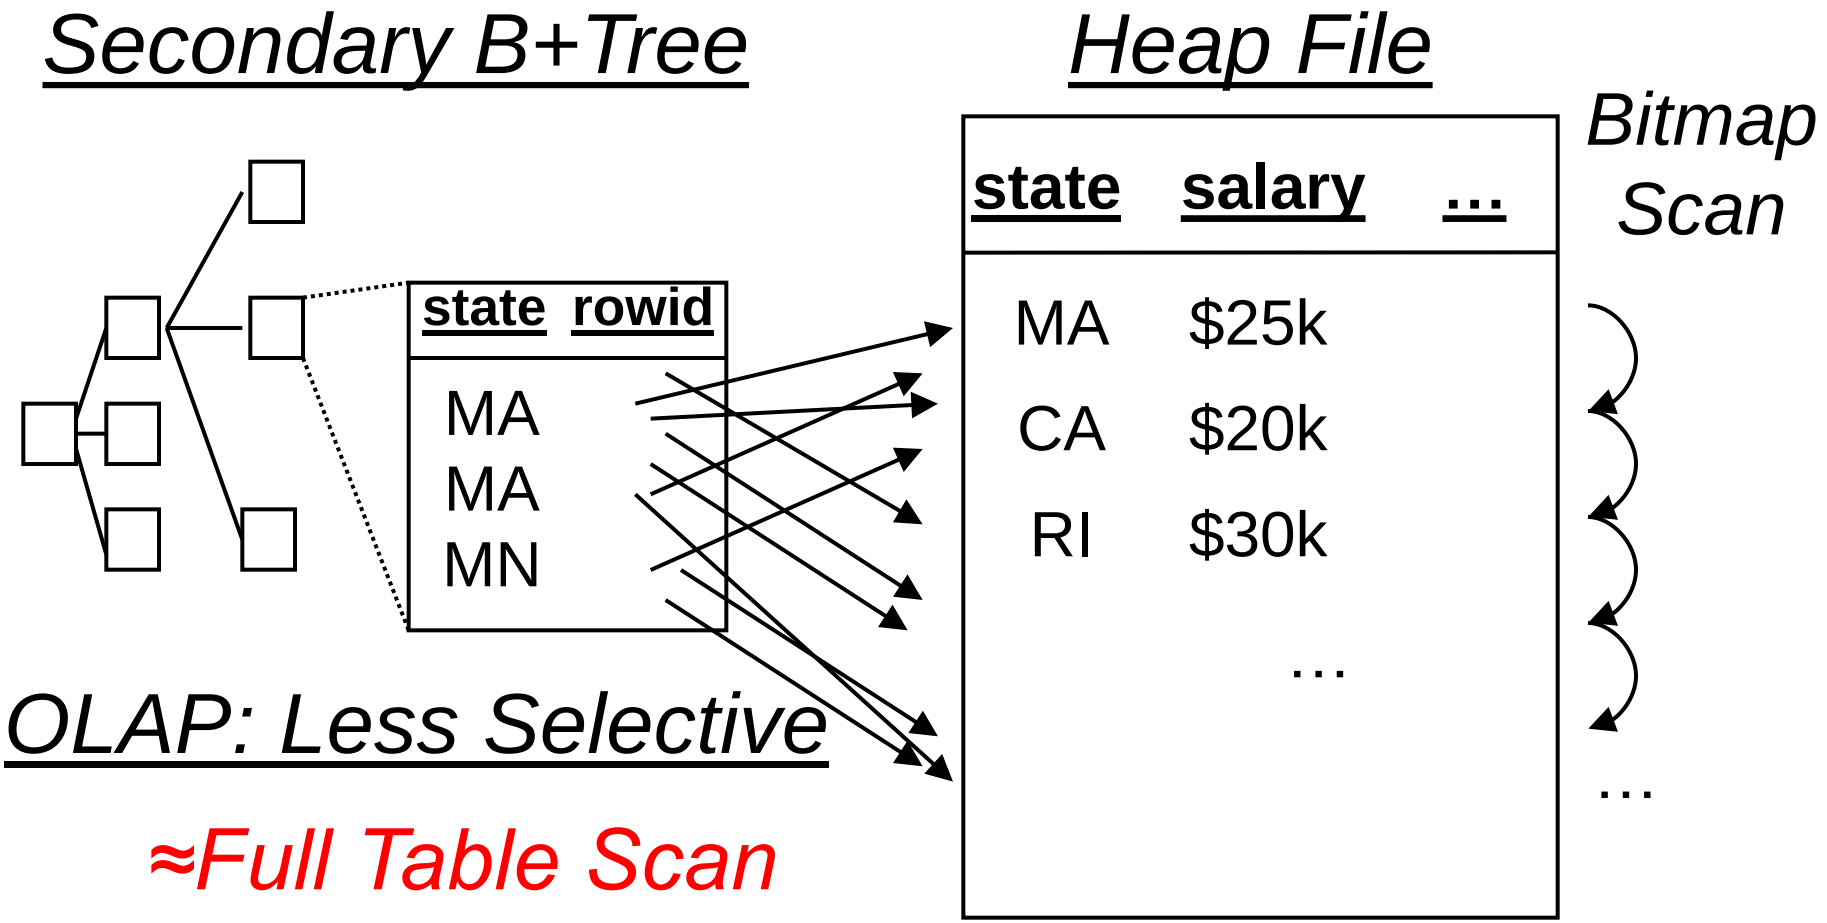

OLAP: Less Selective

≈Full Table Scan

Secondary Index is Too Slow

Problem: Costs of B+Tree

- Storage Cost
- Maintenance Cost
  - More Pressure on Bufferpool
  - More Eviction from Bufferpool = disk writes

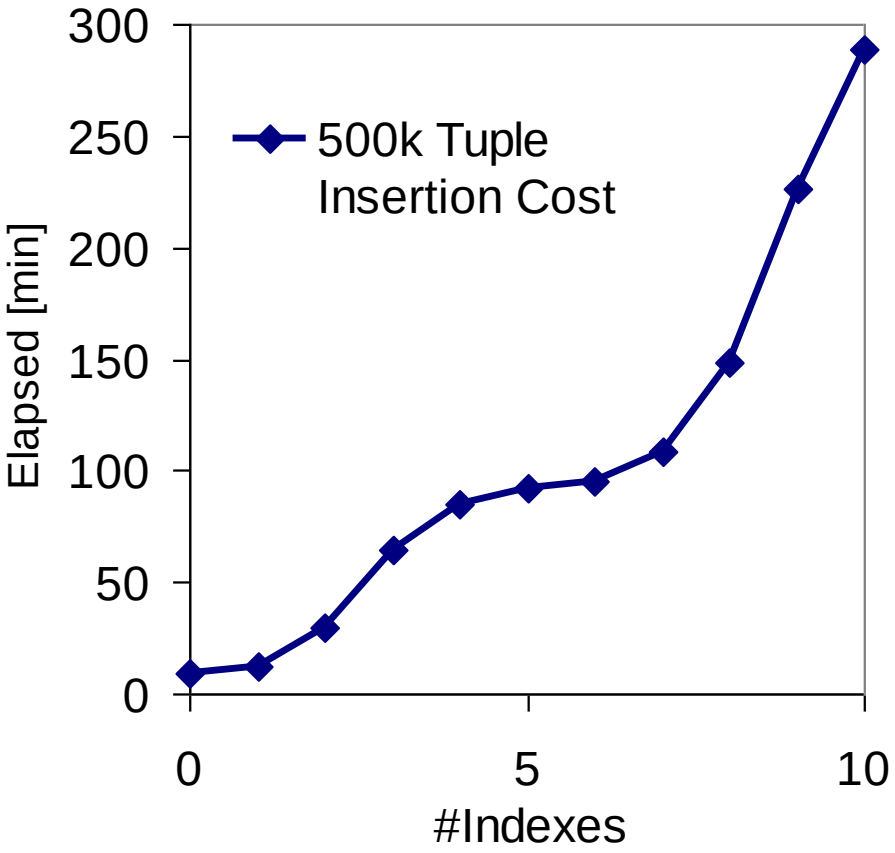

Secondary Index is Too Large

Talk Overview

Idea:

Exploit Correlations between Primary and Secondary Indexes

- Correlation and its potentials
- Correlation-Aware Database Designer
- Experiments

What is Correlation?

ORDERKEY	COMMITDATE	SHIPDATE	PRICE
123	1999/04/01	1999/04/06	$1223
124	1995/09/20	1995/09/19	$210
125	1997/01/01	1997/01/11	$6543
…			

COMMITDATE SHIPDATE

1999/04/01 —*co-occur*→ 1999/03/01 ~ 1999/05/01

1999/04/02 —*co-occur*→ 1999/03/02 ~ 1999/05/02

COMMITDATE is correlated with SHIPDATE, but not with ORDERKEY or PRICE

Correlation and Query Runtime

```
SELECT SUM(PRICE) FROM LINEITEM
WHERE COMMITDATE = 1995/01/01
```

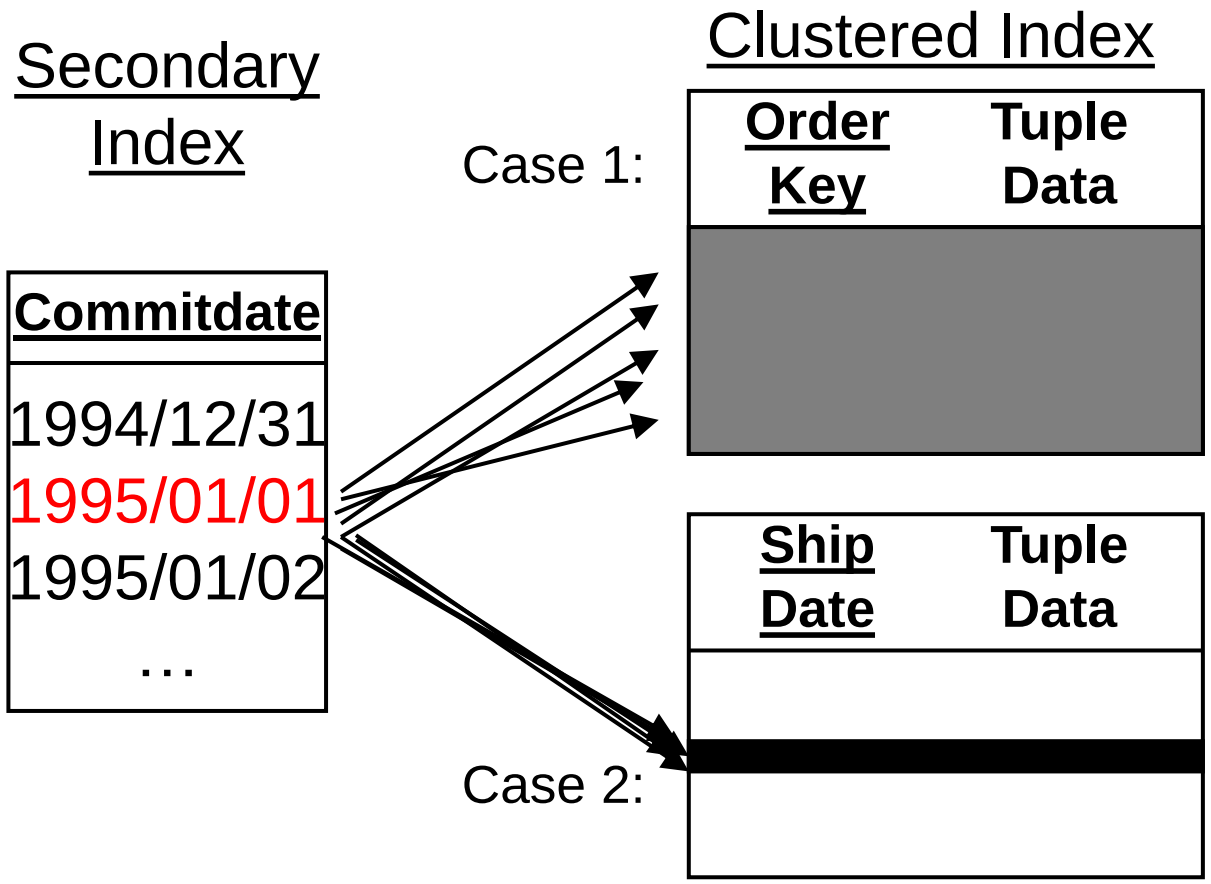

Correlation and Compact Index

Secondary indexes can be *orders of magnitude smaller* with correlation.

- Correlation Map (CM) [VLDB'09]
  Stores soft functional dependencies
- IBM DB2 BHUNT
  Stores algebraic relationships b/w attributes
- SQLServer Datetime Correlation Optimization
  Stores co-occurring datetime pairs

Correlation Map (CM)

SELECT AVG(Sal.) FROM People
WHERE City = Boston

***People* Table**
(Clustered on State)

PID	Sal.	City	State
1	$25K	Boston	MA
2	$90K	Cambr.	MA
3	$80K	Jackson	MN
4	$30K	Boston	NH
5	$50K	Manch.	NH
6	$20K	Toledo	OH

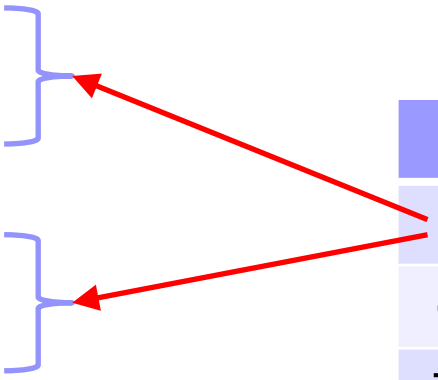

Correlation Map
(City -> State)

City	State
Boston	{MA,NH}
Cambr.	{MA}
Jackson	{MN}
…	…

1,000x Smaller Than Conventional B+Tree

Correlation in DB Design

- Good News

  Faster and Smaller (Maintainable) Secondary Index

- Bad News

  Too Complex for Human DBAs to Pick Primary/Secondary Indexes

Let's Automate Designing

Problem Statement

Automatic Database Designer

- Input: Queries, Space Budget
- Output: Database Objects
  - Materialized Views (MVs)
  - Clustered Indexes
  - Secondary Indexes
- Constraint: Size(Objects) ≤ Budget
- Objective: *min* TotalRuntime (Queries)

Automatic Database Designer

- Plethora of work In Industry/Academia
- Key Component in Today's DBMS
  - Microsoft SQLServer Database Tuning Advisor
  - IBM DB2 Design Advisor

Still, No Existing Database Designer Exploits or is aware of Correlation

CORADD Framework

- MV Candidate Generator
  - Projection (Columns)
  - Clustered Index
- Correlation Aware Cost Model
- ILP/ILP Feedback
- CM Advisor

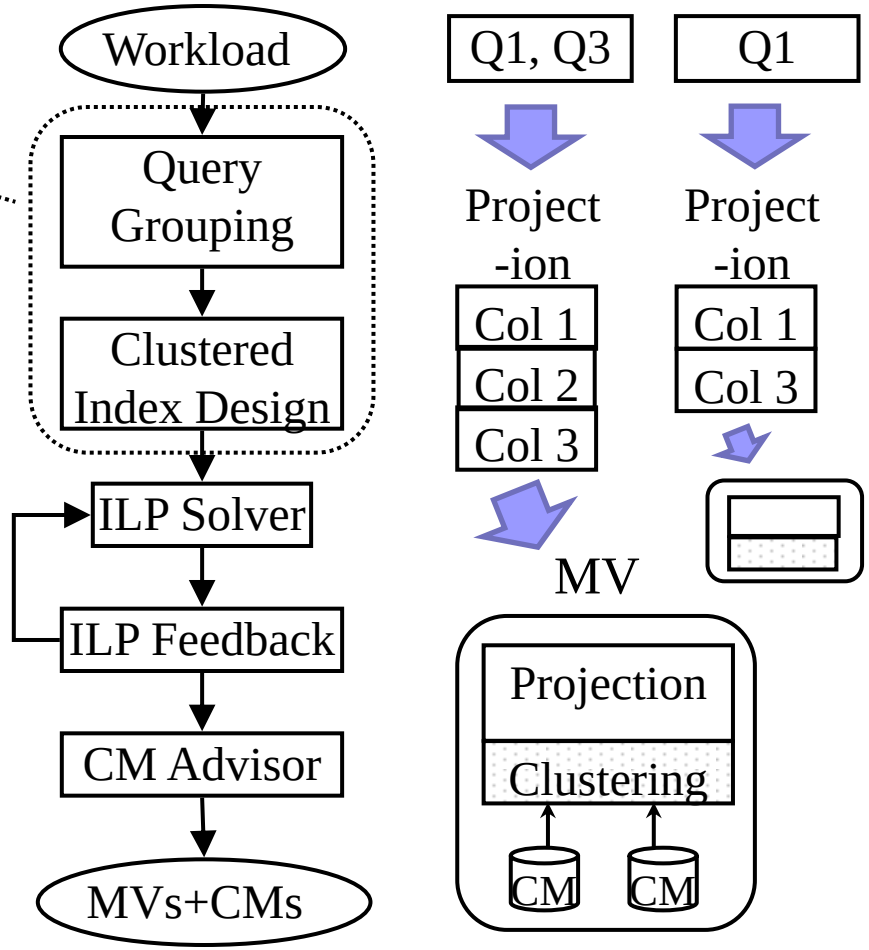

MV Candidate Generator

- MV = Projection + Clustered Index
- Projection Design
  - Choose Queries To Serve
  - Maximize *Similarity* Of Queries (Predicates)
- Clustered Index Design
  - Choose Clustered Index On Projection
  - Maximize *Correlation* with Secondary Indexes (CMs)

Projection: Query Grouping

- Selectivity Vector

Query	*year*	*yearmonth*
Q1	0.15 →	0.15
Q2	0.15 ←	0.013

Predicates

Correlated

year = '95

yearmonth = '94Dec

- Selectivity Propagation

- k-means to Group Similar Queries

$$dist\left(\vec{v_1}, \vec{v_2}\right) \equiv \sqrt{\sum_{attr} \left(v_1[attr] - v_2[attr]\right)^2}$$

Clustered Index Design: Split

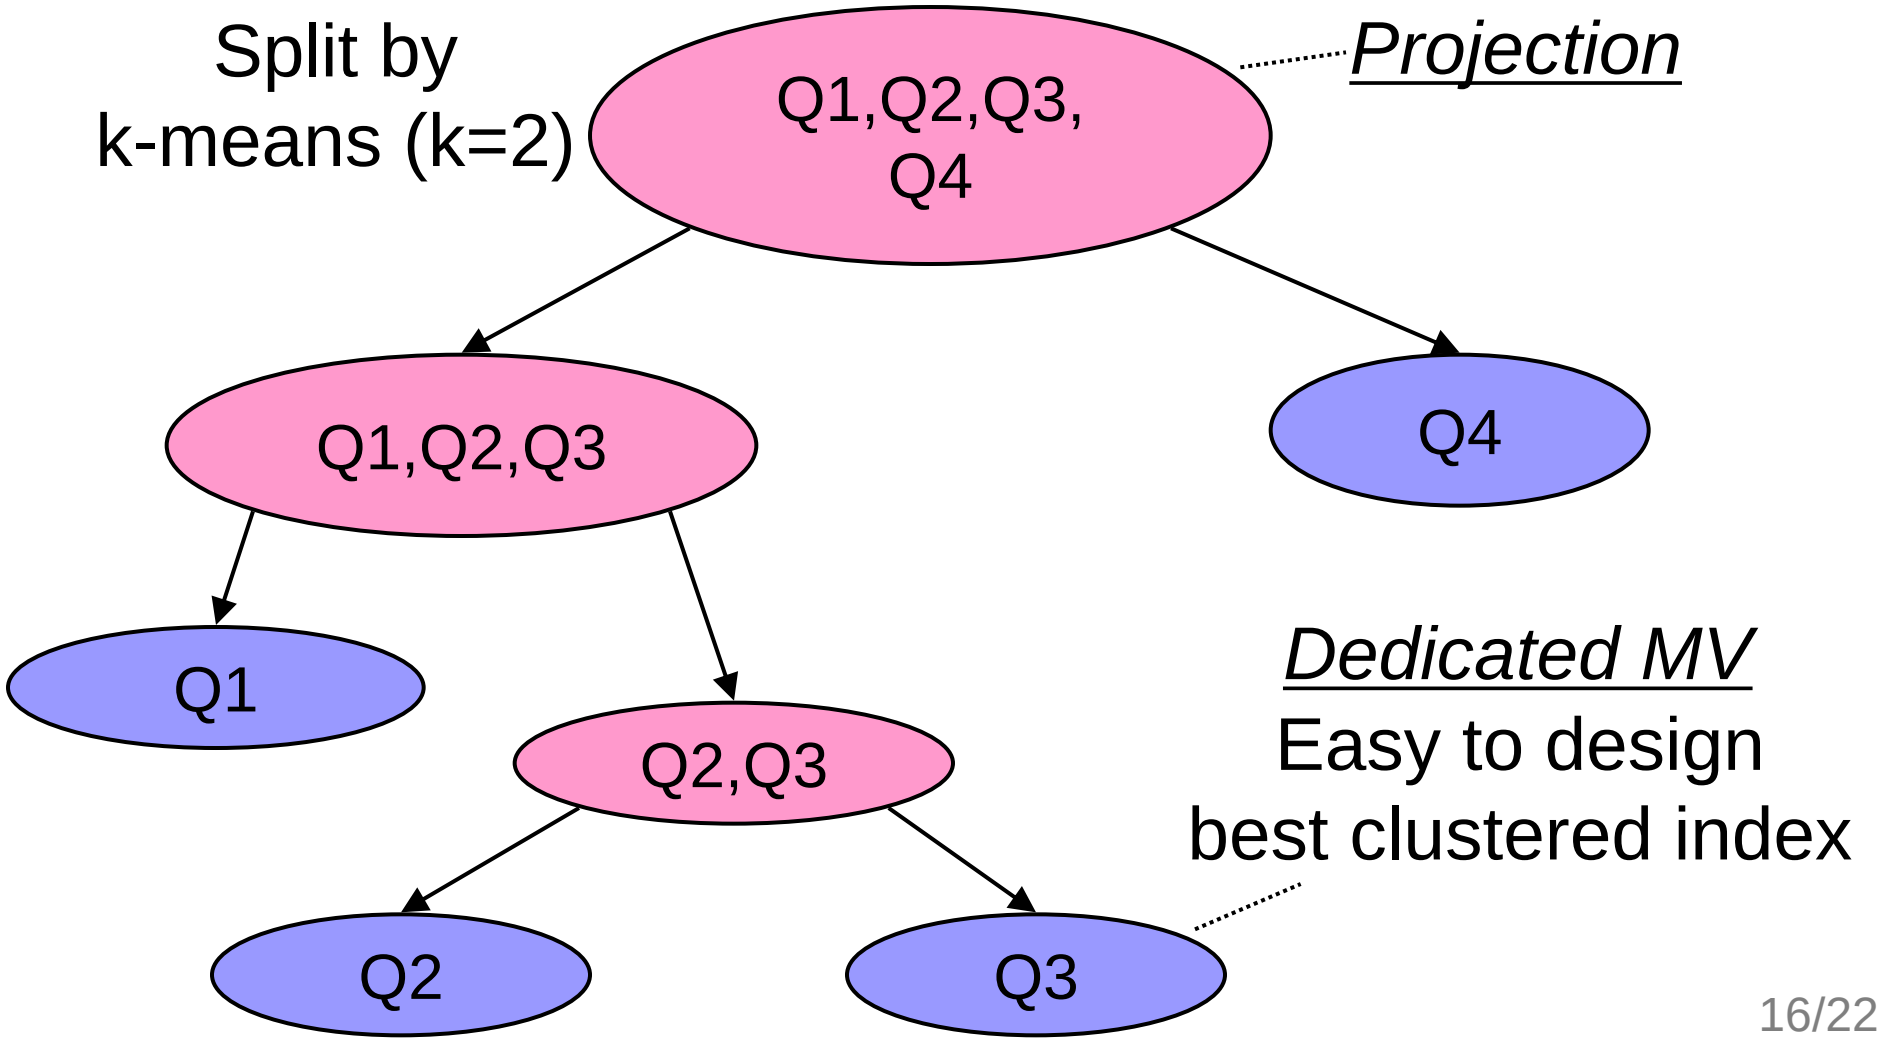

Clustered Index Design: Merge

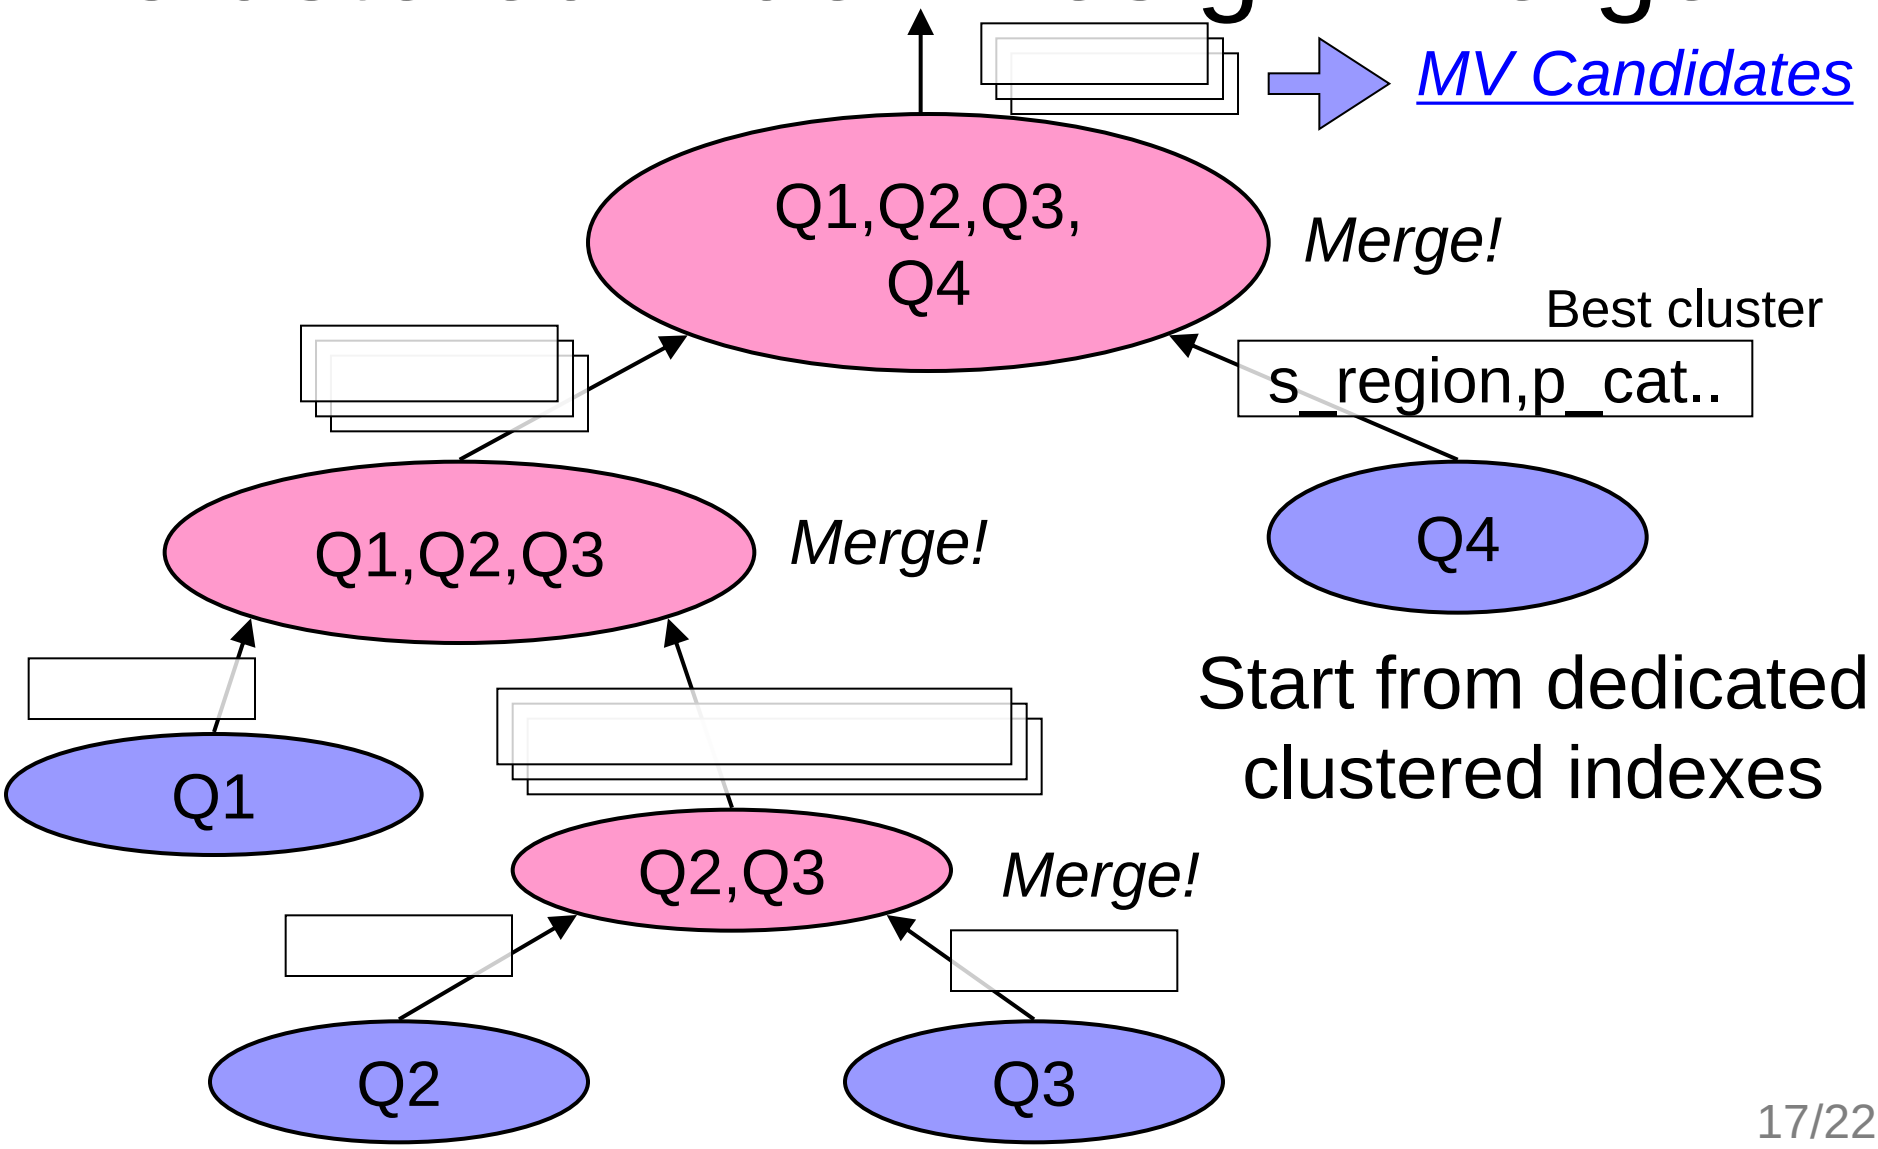

Experiments

- Environment
  - DBMS X on Windows Server Enterprise Ed.
  - Quad-core, 4GB RAM, 10k rpm HDD
- Dataset/Workload
  - OLAP Council APB-1 (45M tuples)
  - Star Schema Benchmark (24M tuples)
- Compared w/ DBMS X's own DB Designer

Cost Model Accuracy (APB-1)

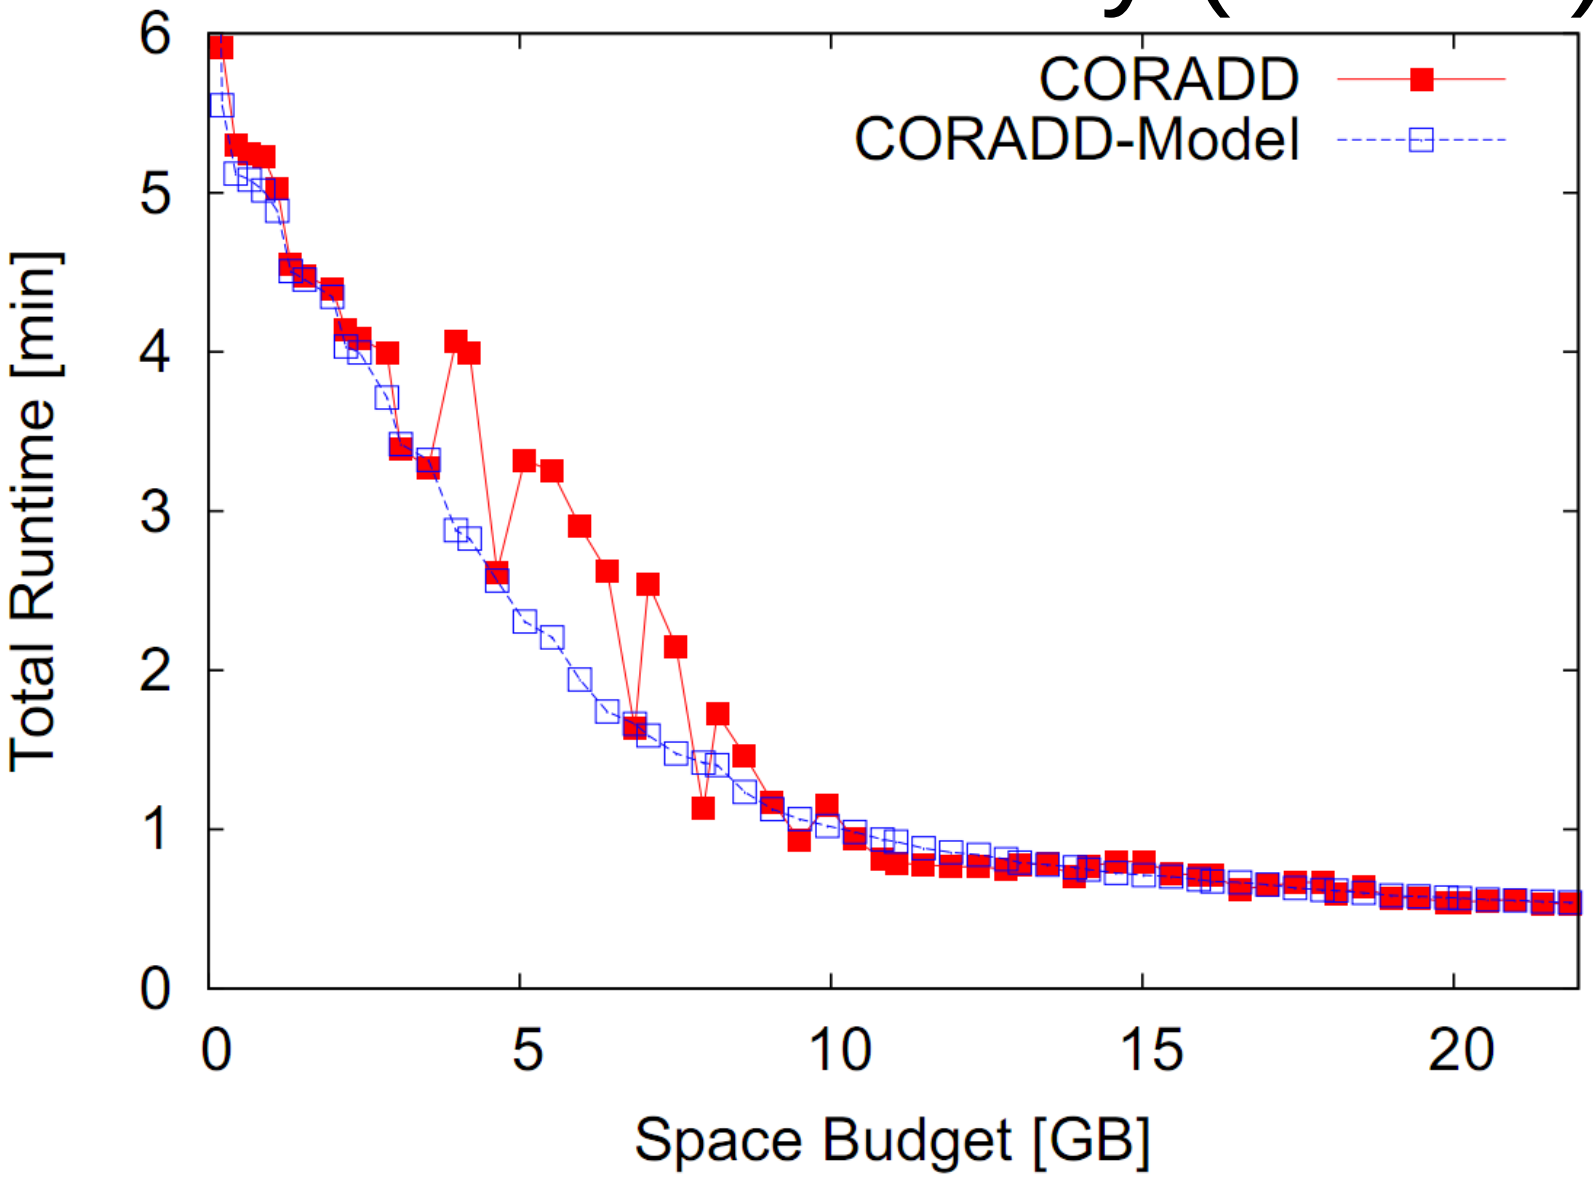

Real Runtimes on APB-1

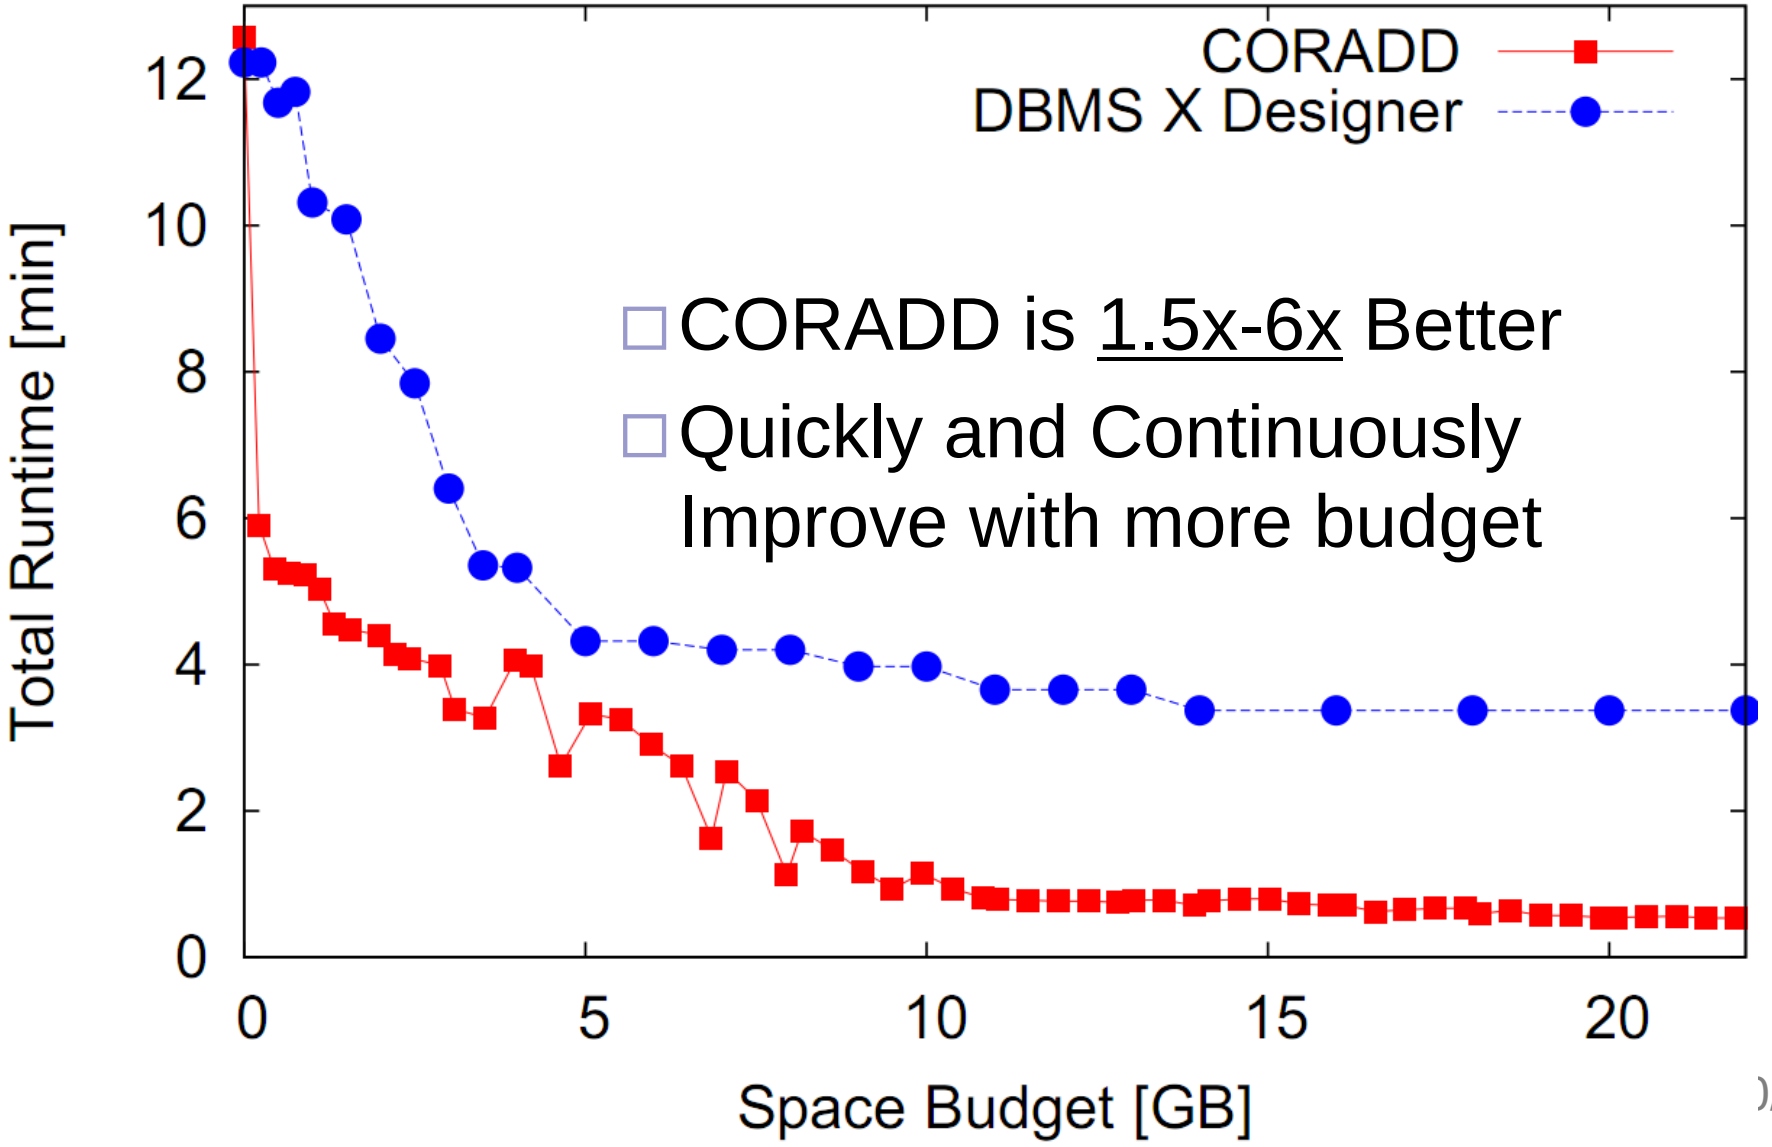

Real Runtimes on SSB

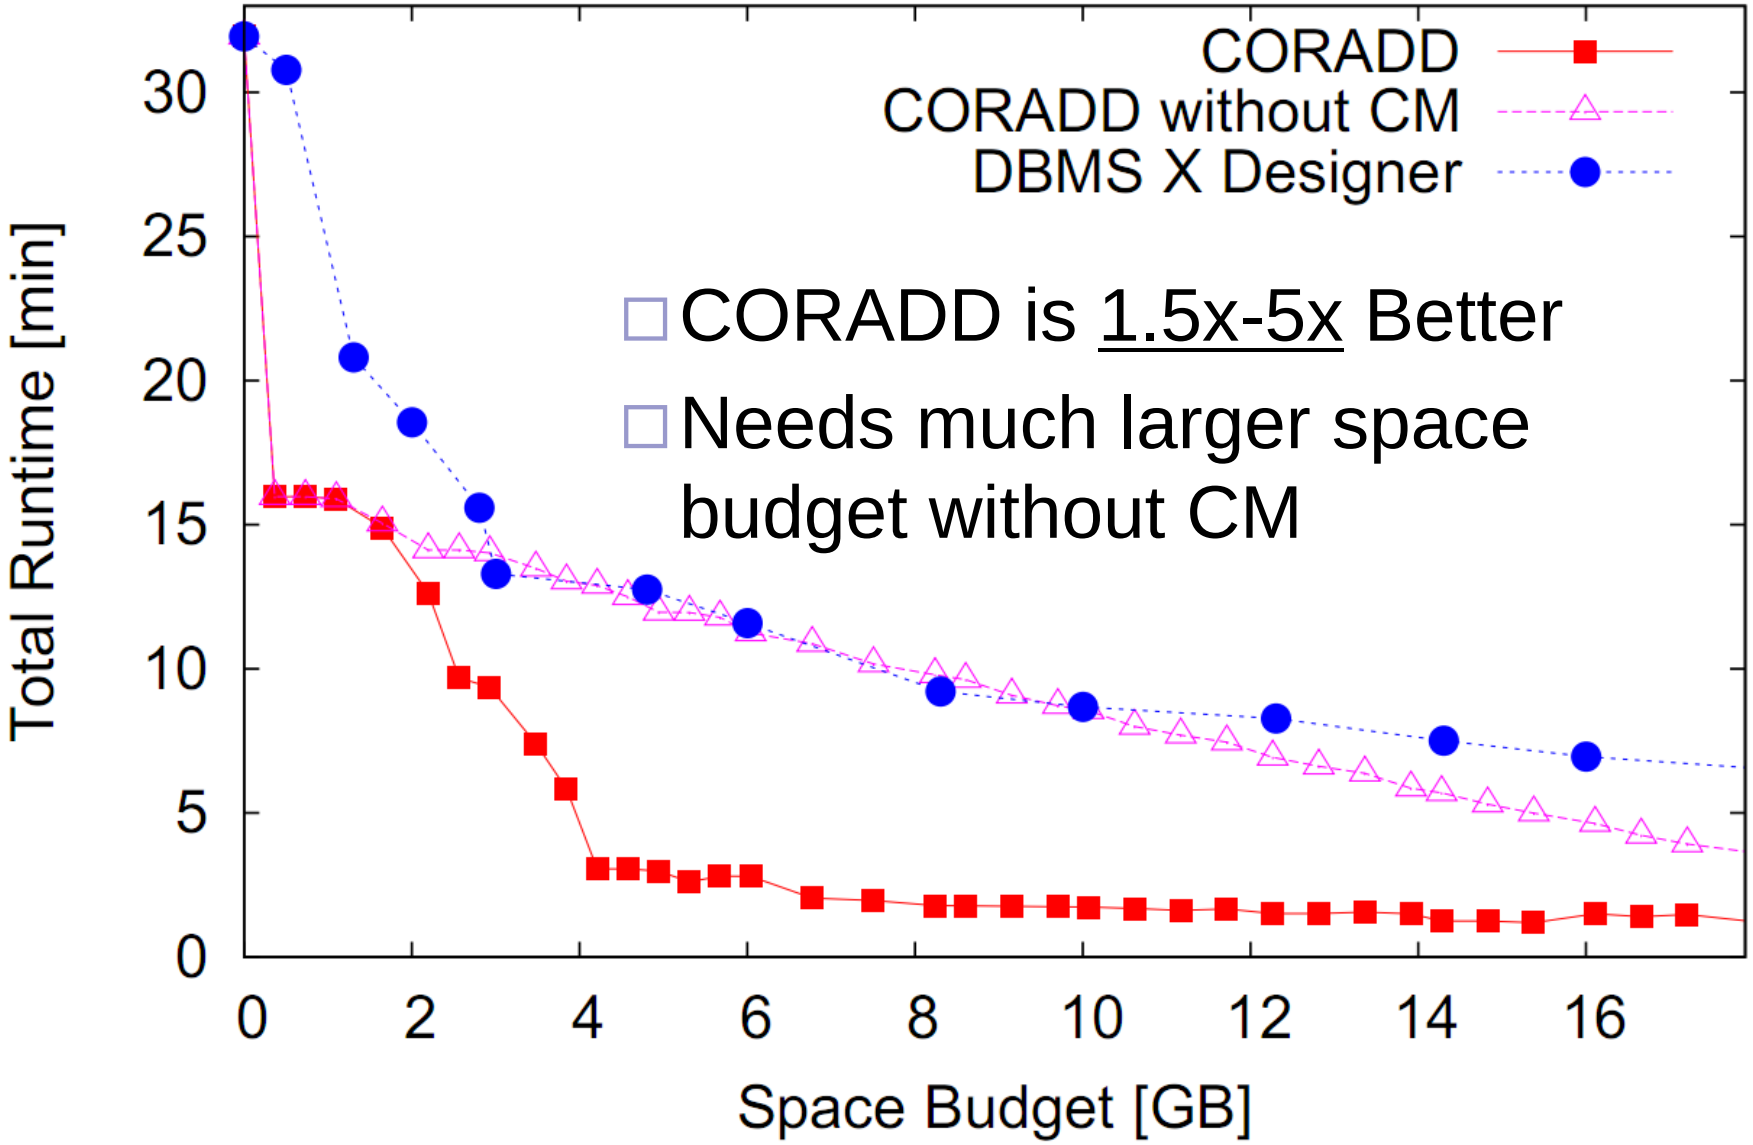

Conclusion

- Index Correlation
  - Correlation = Faster & Smaller Index
  - Choice of Clustered Index matters!
- Correlation-Aware Database Designer
  - Design MVs based on Selectivity Vector
  - Maximize Correlation between Clustered/Secondary Index
  - 5-6x Better Quality to State-of-the-art