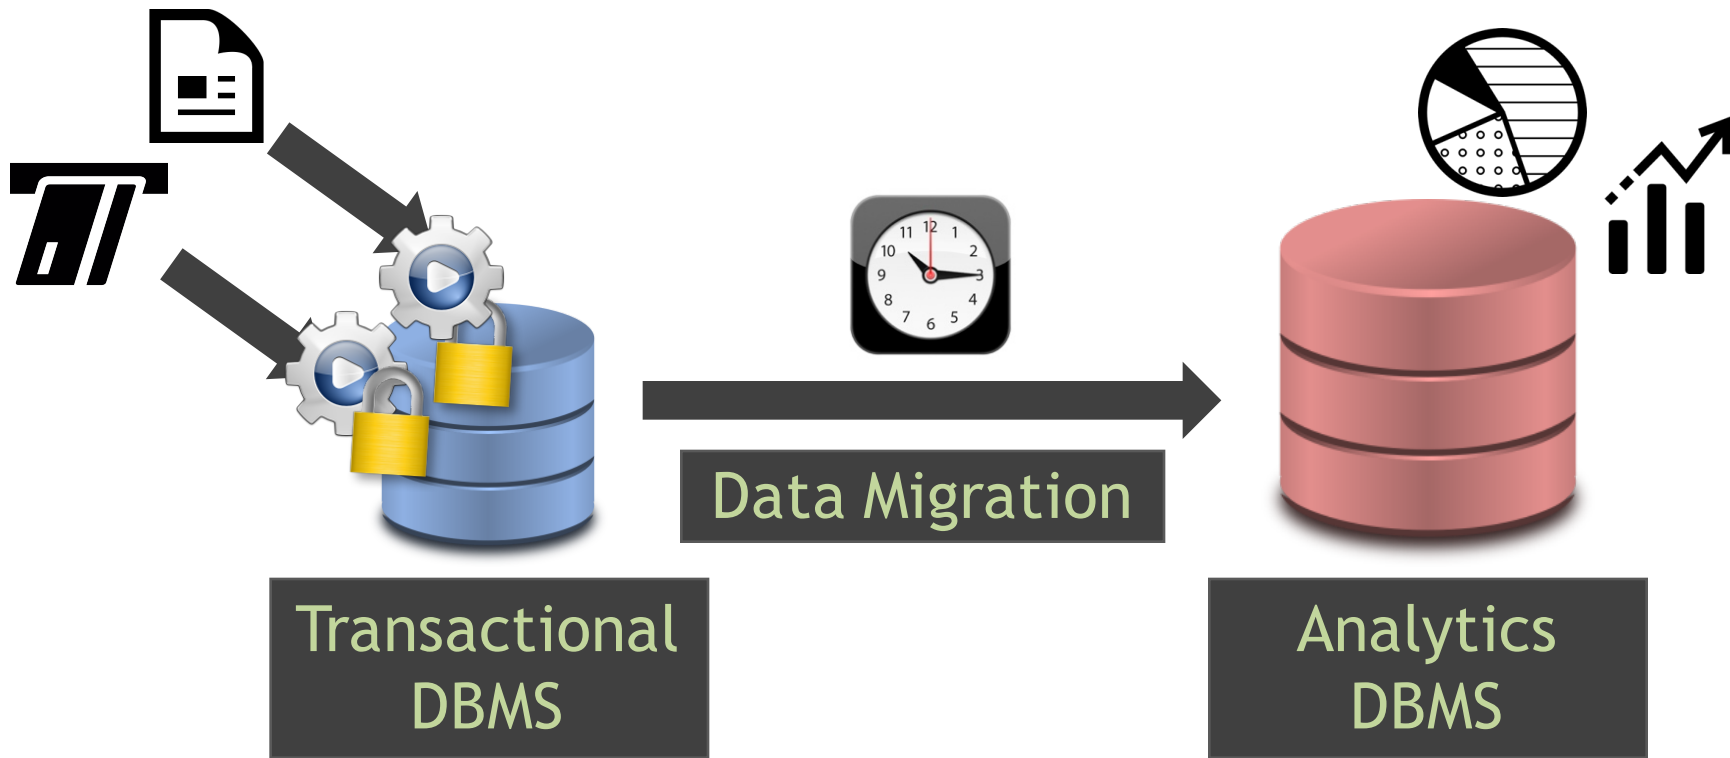


BRIDGING THE ARCHIPELAGO BETWEEN ROW-STORES AND COLUMN-STORES FOR HYBRID WORKLOADS

Joy Arulraj, Prashanth Menon, and Andy Pavlo
Carnegie Mellon University

SIGMOD 2016

REAL-TIME ANALYTICS



Transactional
DBMS

Analytics
DBMS

TRANSACTIONAL VS ANALYTICS DBMS



	TRANSACTIONAL DBMS	ANALYTICS DBMS
Storage Layout	Row-store	Column-store
Query Processing	Tuple-at-a-time	Attribute-at-a-time
Write-Intensive Workloads	✓	✗
Read-Intensive Workloads	✗	✓

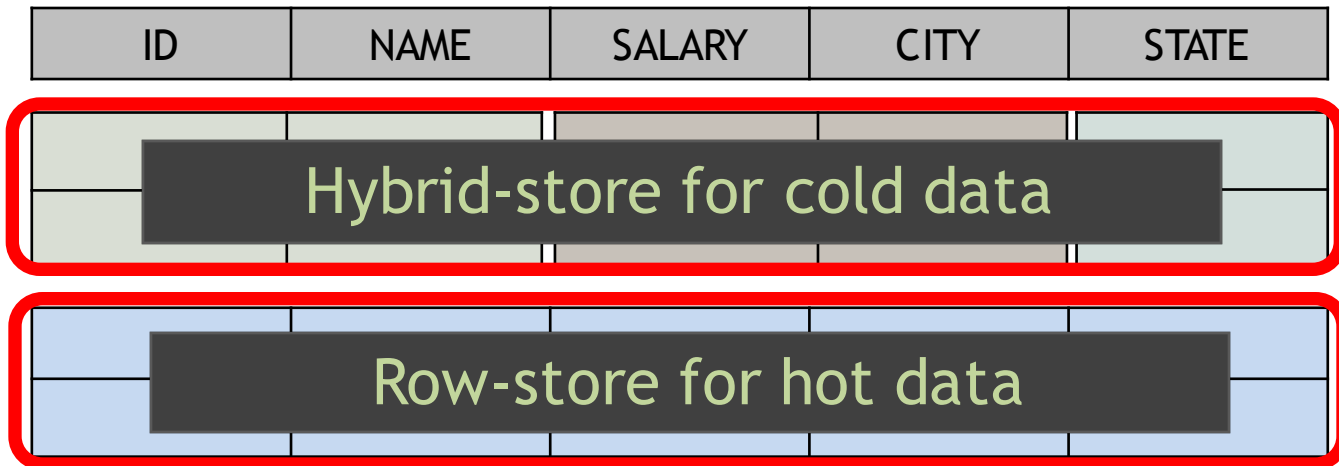
LIMITATIONS OF BIFURCATED SETUP

- ✗ Time taken to transfer data to analytics DBMS
- ✗ Overhead of maintaining two separate DBMSs

How can we do real time analytics on fresh data with lower overhead ?

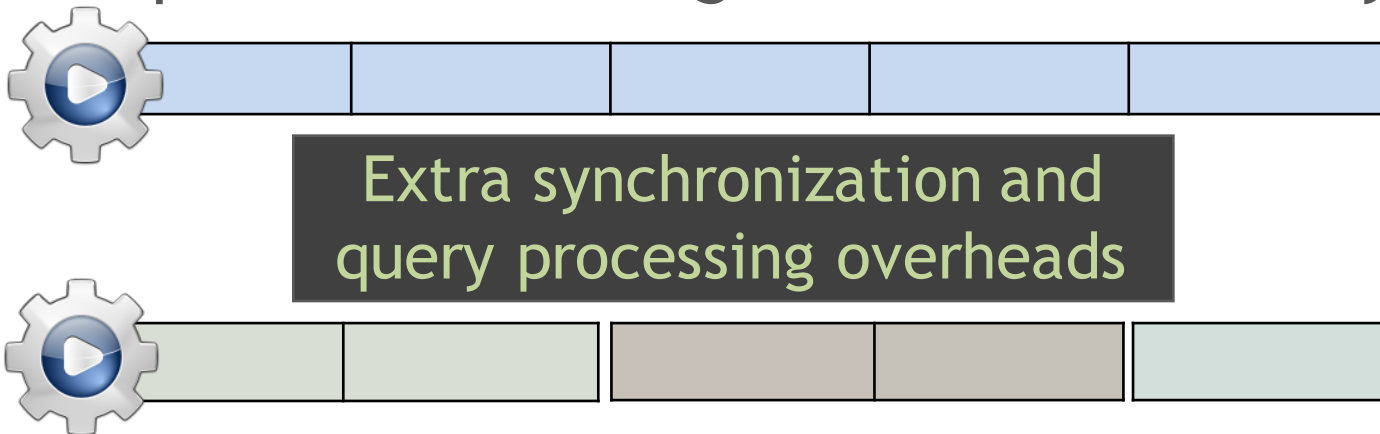
HYBRID DATABASE SYSTEM

- ❖ Hybrid storage layouts based on workload
 - *Prior work: Data Morphing, HYRISE, HANA, H2O*



HANDLING DIFFERENT LAYOUTS

X Multiple execution engines for different layouts



Can we design a single execution engine that can process tuples stored in different layouts ?

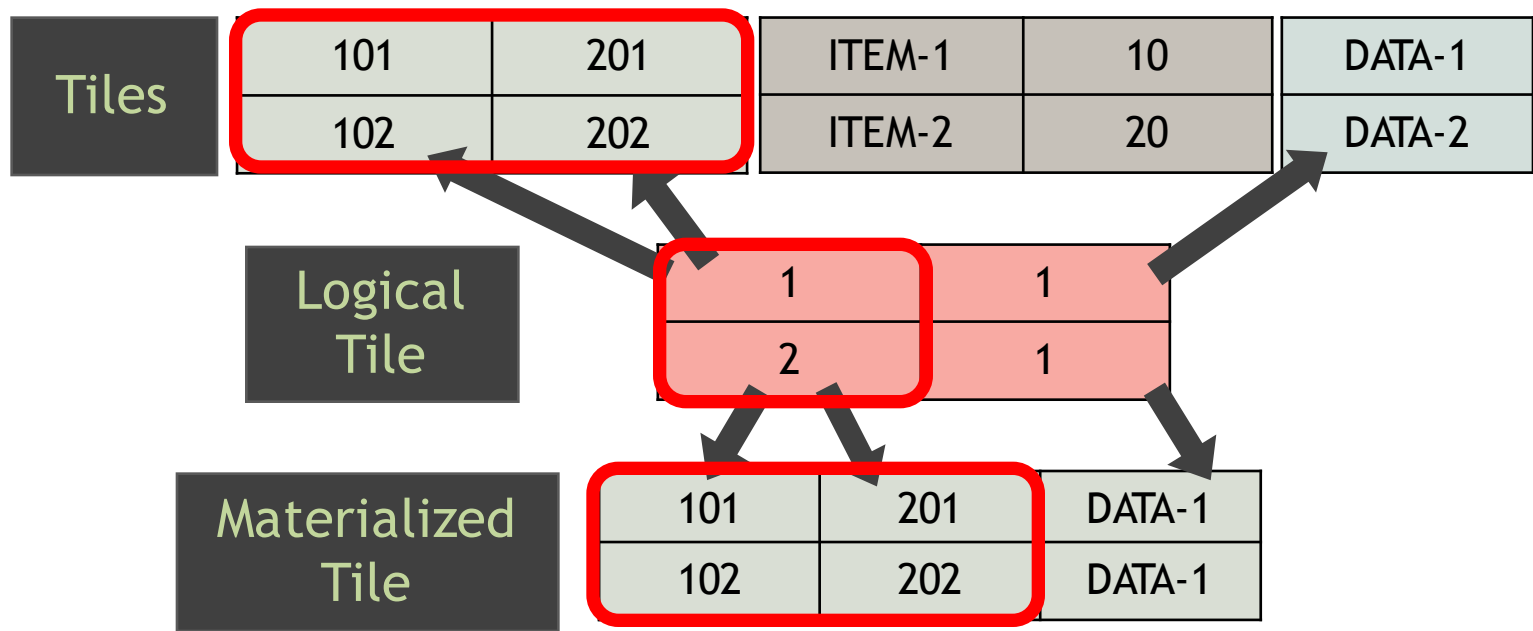
KEY IDEA: LOGICAL TILE

- ✓ Abstract the hybrid layout from the execution engine

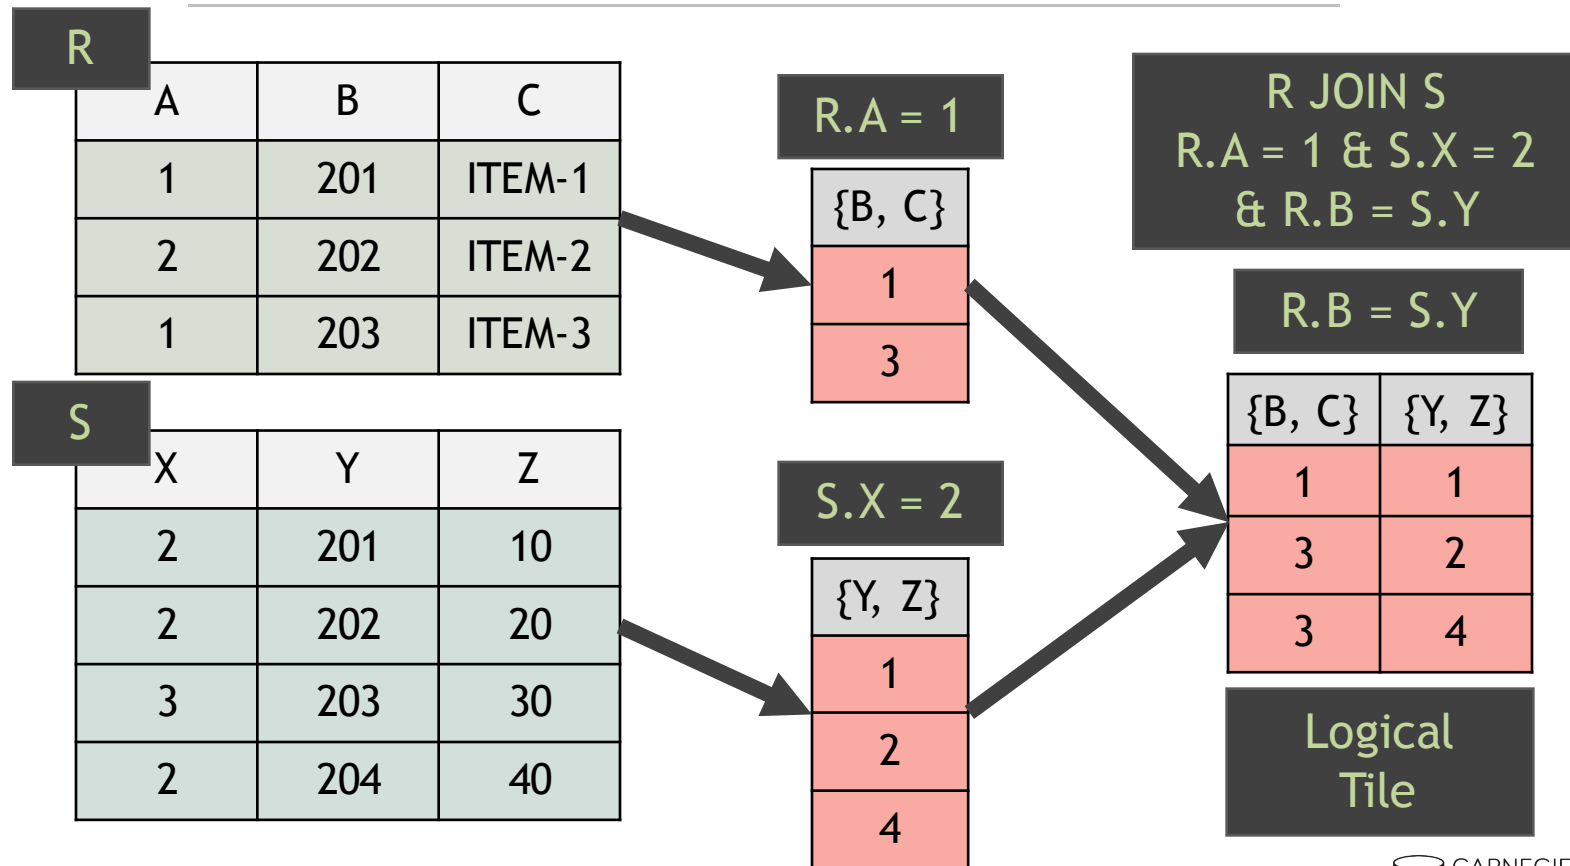
ID	NAME	SALARY	CITY	STATE
Tile #1		Tile #2		Tile #3

KEY IDEA: LOGICAL TILE

- ✓ Abstract the hybrid layout from the execution engine



SELECTION & JOIN OPERATORS




LAYOUT-AGNOSTIC EXECUTION ENGINE

- ❖ All operators produce and consume logical tiles
 - *Aggregation, Projection, etc.*

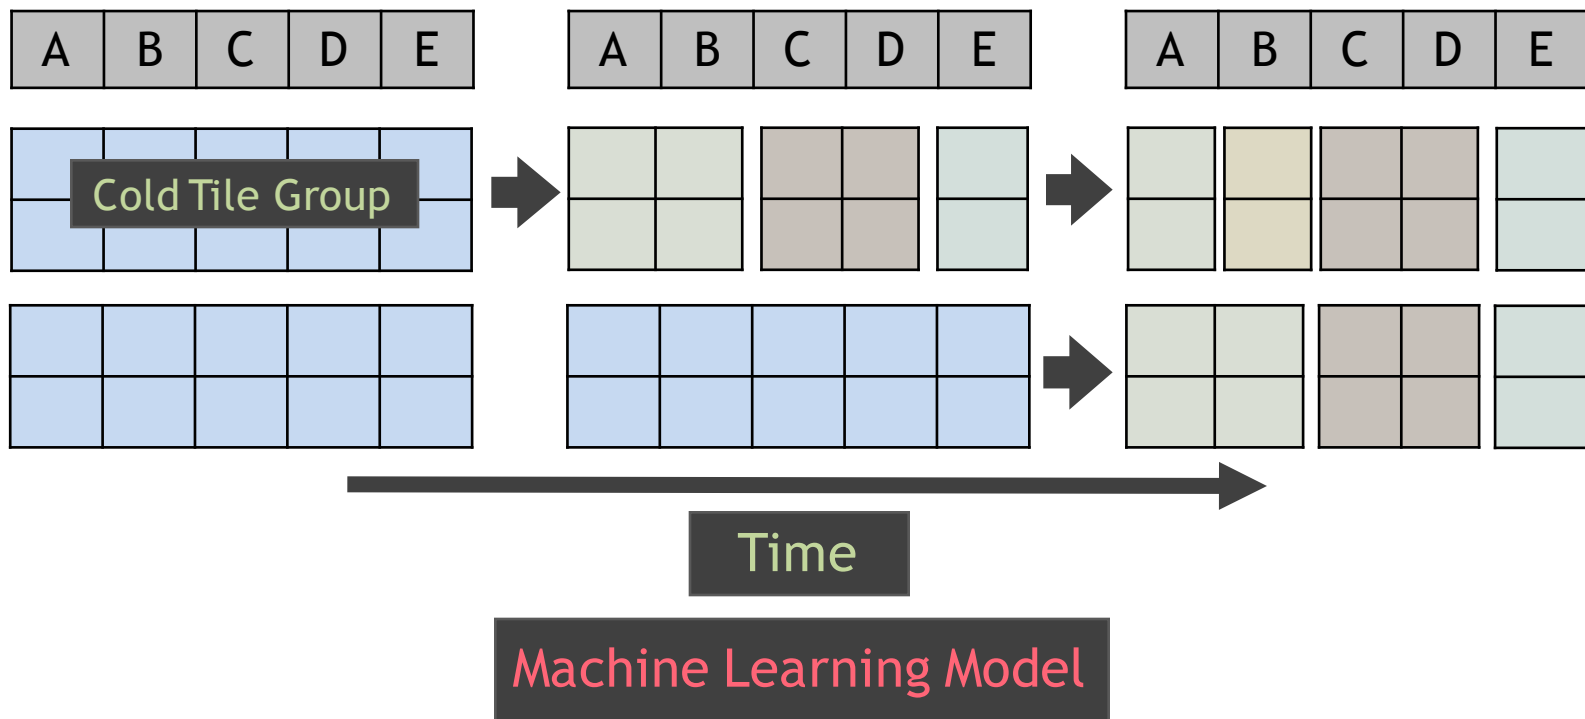
Logical tile algebra can be used to build a layout-agnostic engine for real-time analytics

DYNAMIC HYBRID WORKLOADS



How should the DBMS adapt the hybrid storage layout with workload shifts ?

KEY IDEA: DYNAMIC LAYOUT ADAPTATION

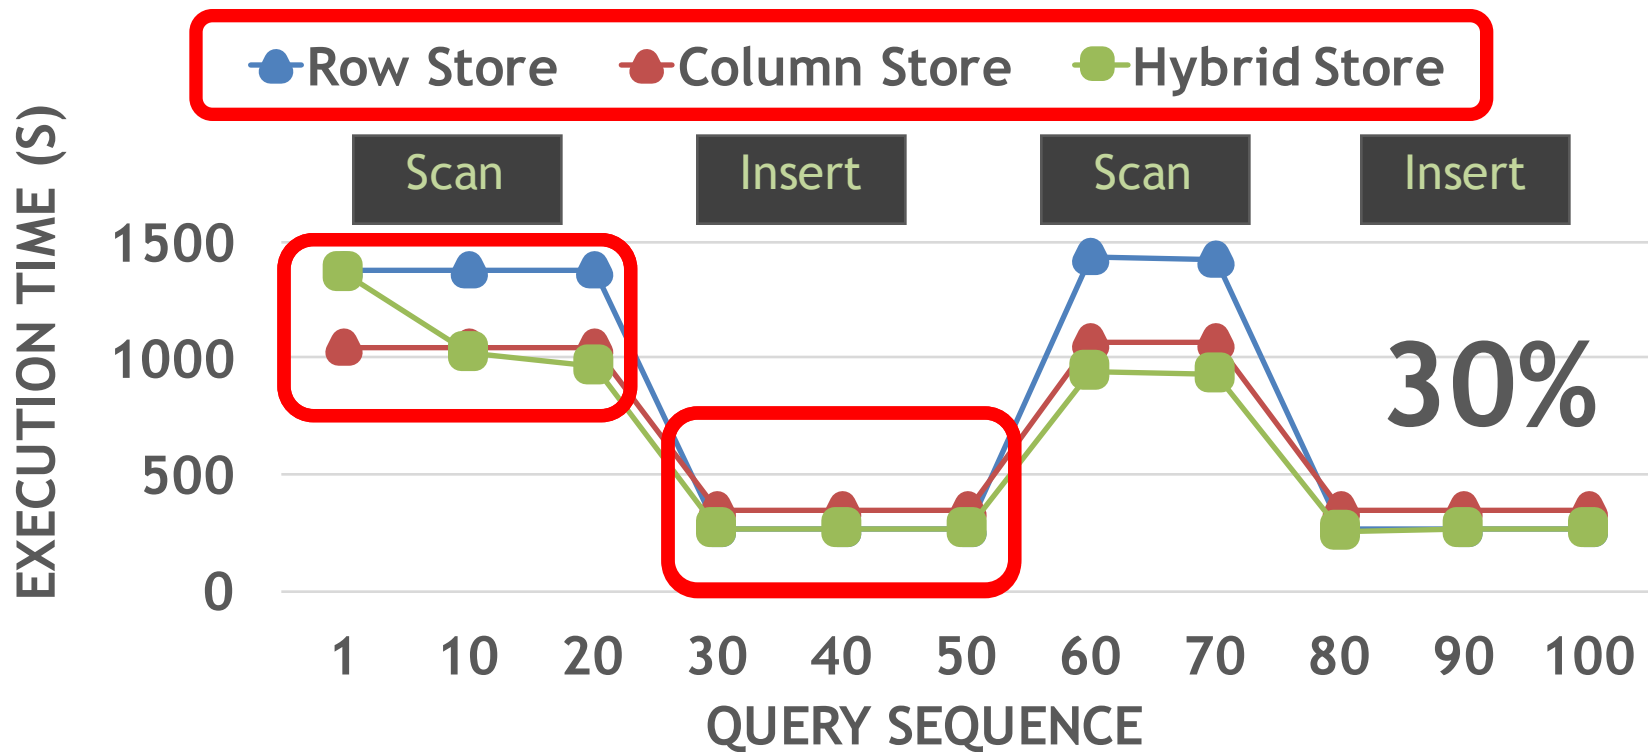


EXPERIMENTAL EVALUATION

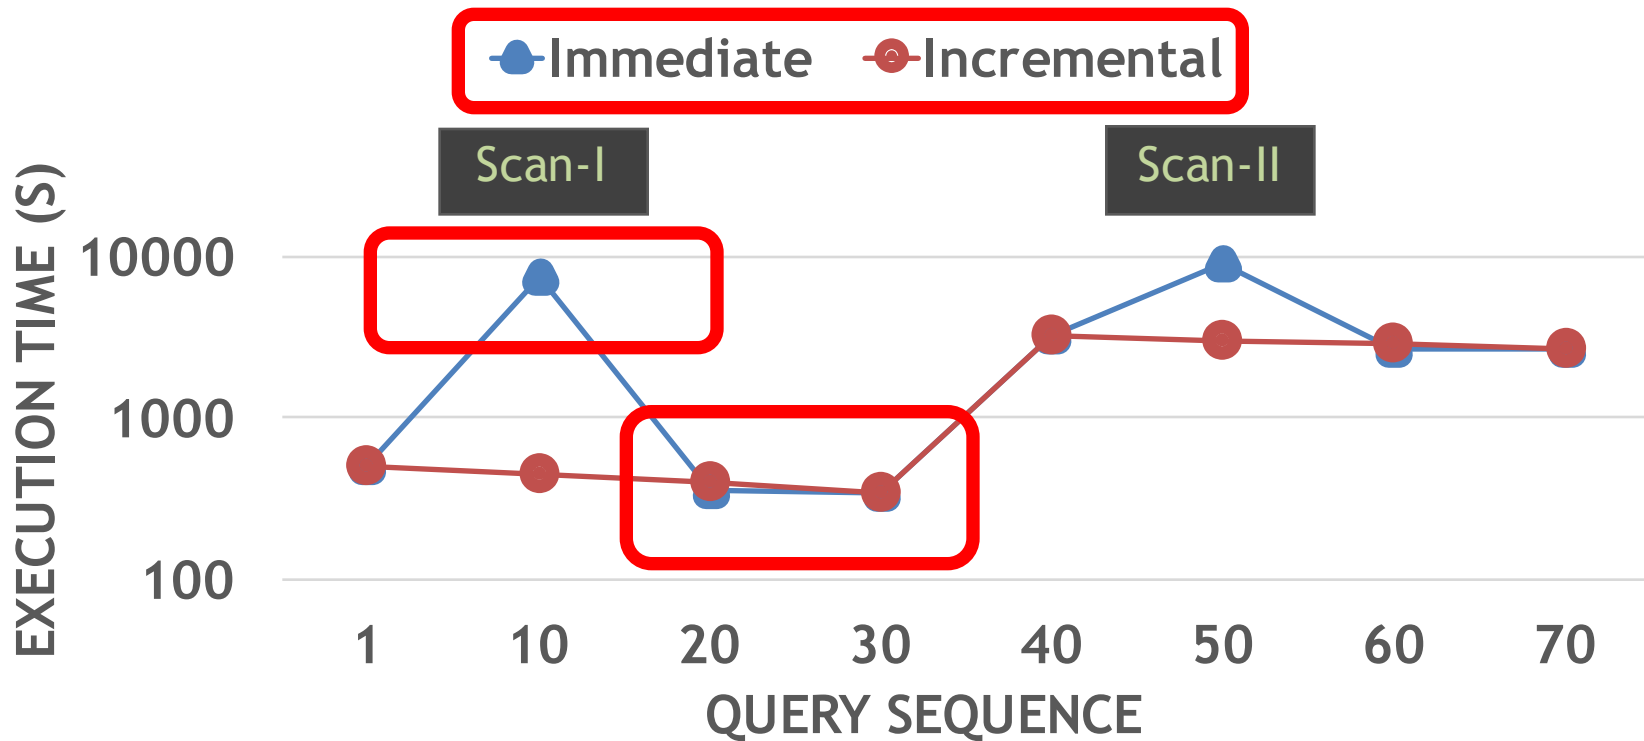
❖ ADAPT Benchmark

- *Dynamic hybrid workload*
- *Row-store, column-store, and hybrid-store*
- *Mixture of 1000 scan queries and 1M inserts*
- *Table with 10M tuples (~2GB)*

DYNAMIC HYBRID WORKLOAD



“INCREMENTAL” LAYOUT ADAPTATION




TAKEAWAYS

- ❖ Layout-agnostic execution engine
 - *Enables real-time analytics with lower overheads*
- ❖ Automatic layout reorganization module
 - *Obviates the need for manual layout tuning*




PELOTON

<http://pelotondb.org>



WE'RE ALWAYS LISTENING !
TOLL-FREE HOTLINE
1-844-88-CMUDB



END