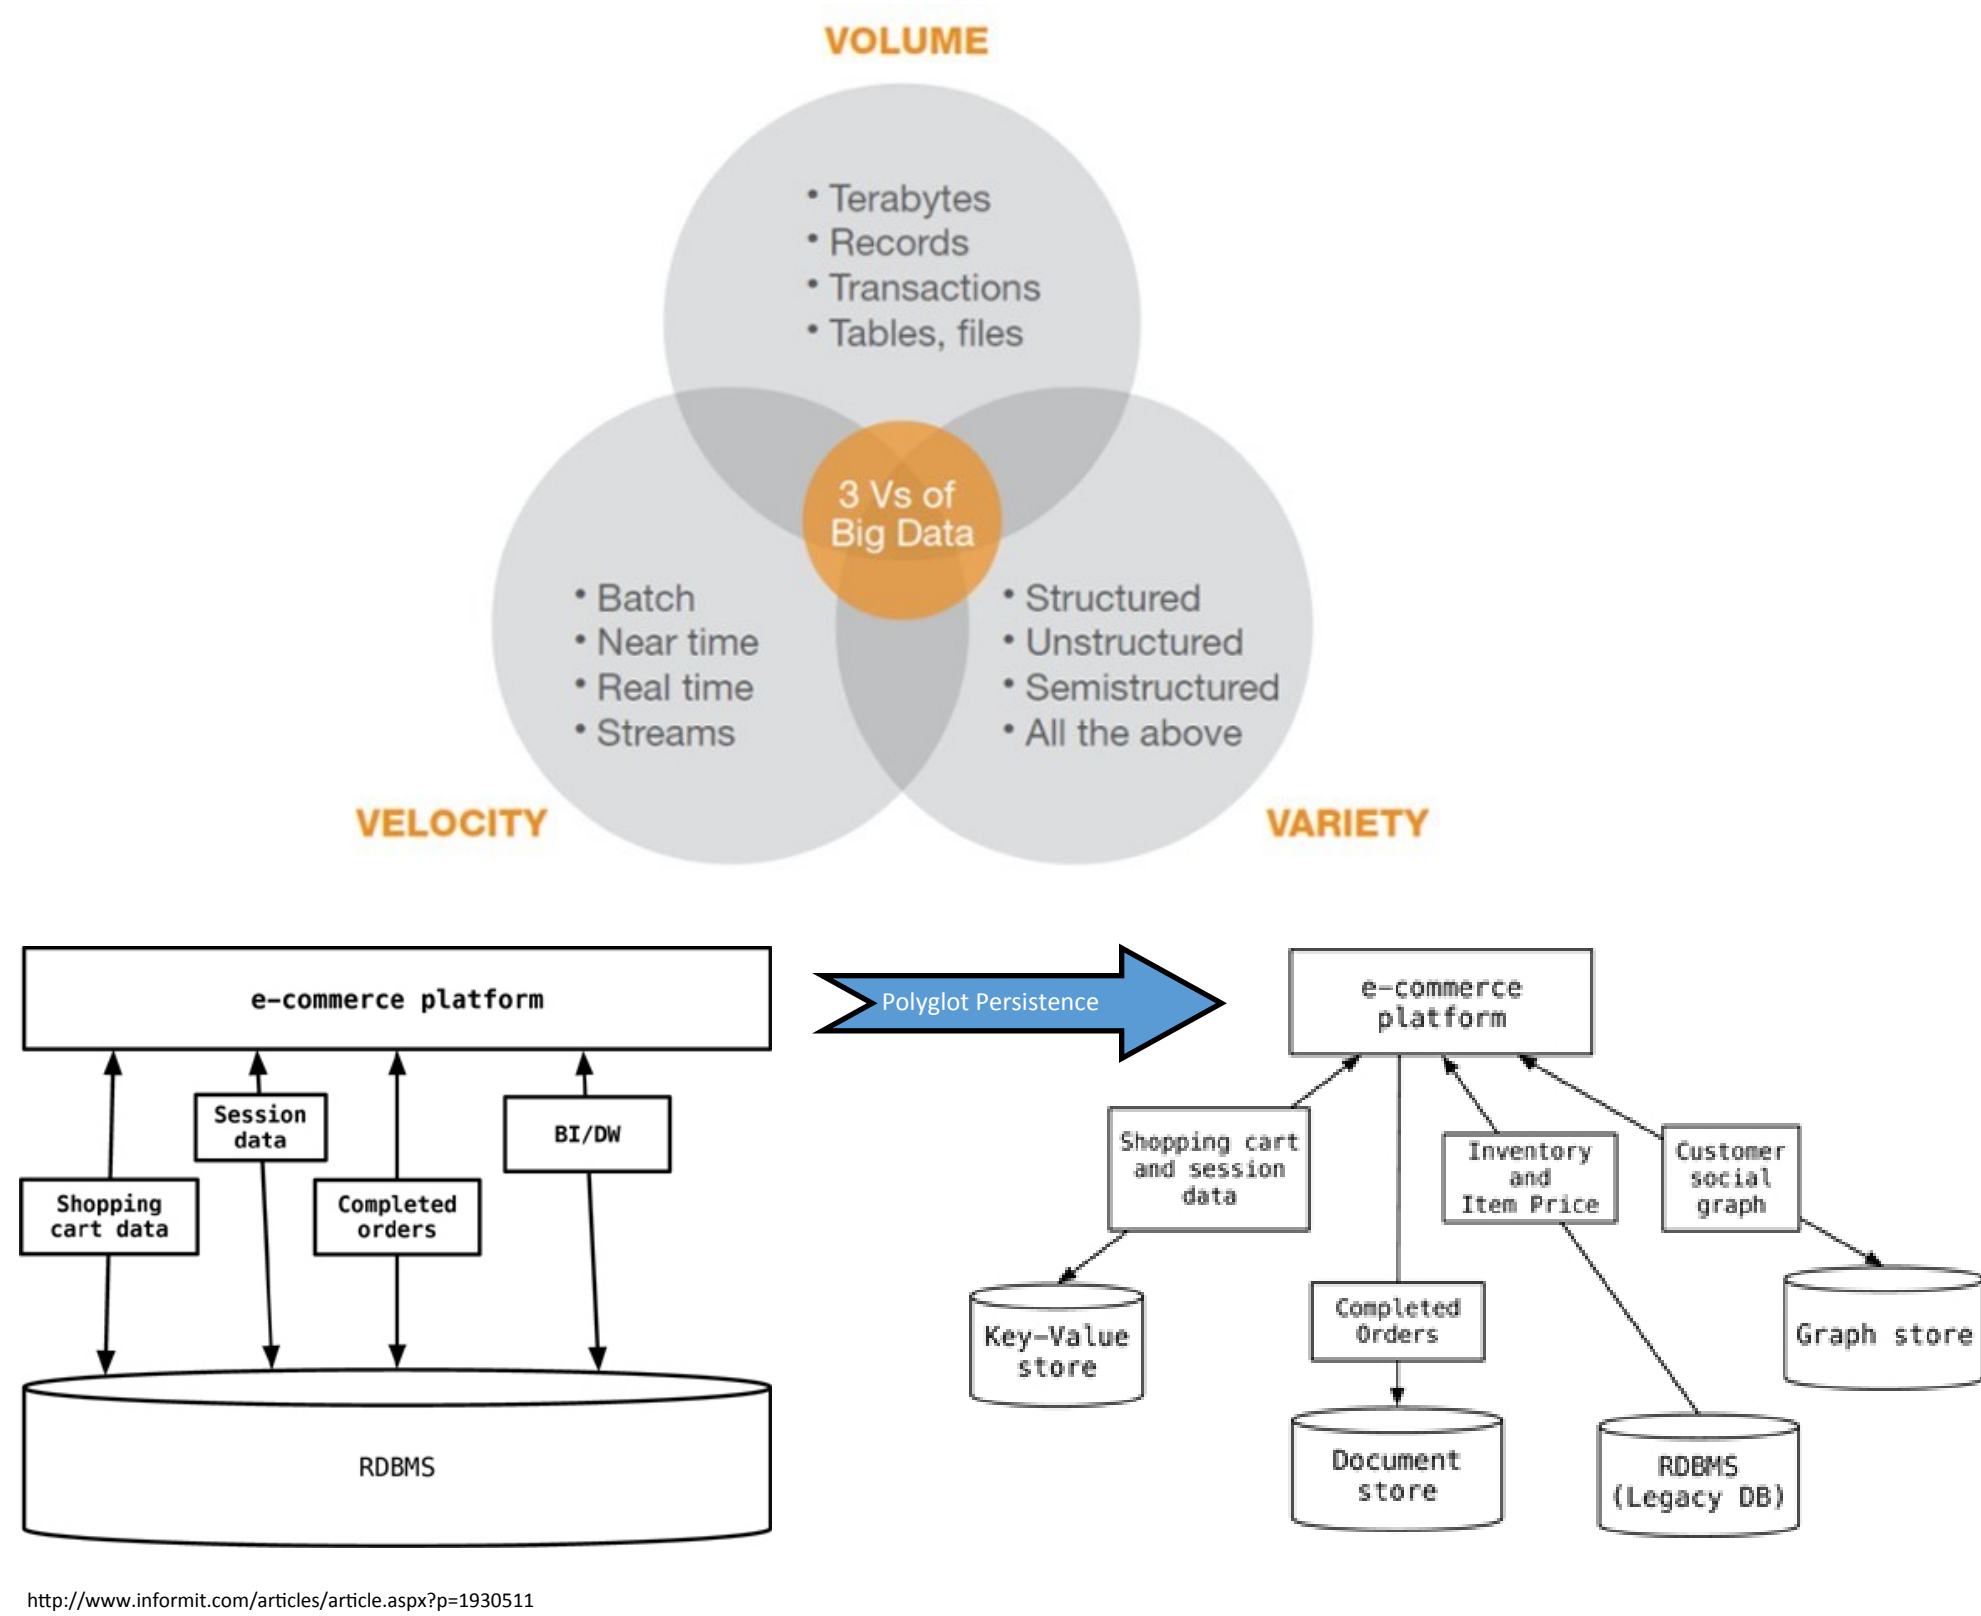


Optimizing Analytic Data Flows in Polyglot Persistence

Rana Faisal Munir, Alberto Abelló, Oscar Romero
 Universitat Politècnica de Catalunya, BarcelonaTech
 [fmunir | aabello | oromero]@essi.upc.edu

Wolfgang Lehner, Maik Thiele
 Technische Universität Dresden
 [wolfgang.lehner | maik.thiele]@tu-dresden.de



Challenges

- ⇒ Perform data analysis in polyglot persistence
- ⇒ Different query processing capabilities
 - ⇒ Joins
 - ⇒ Range queries
 - ⇒ Secondary Indexes
 - ⇒ Aggregation
- ⇒ Optimization of analytic data flows
- ⇒ Selection of store and storage format for intermediate results

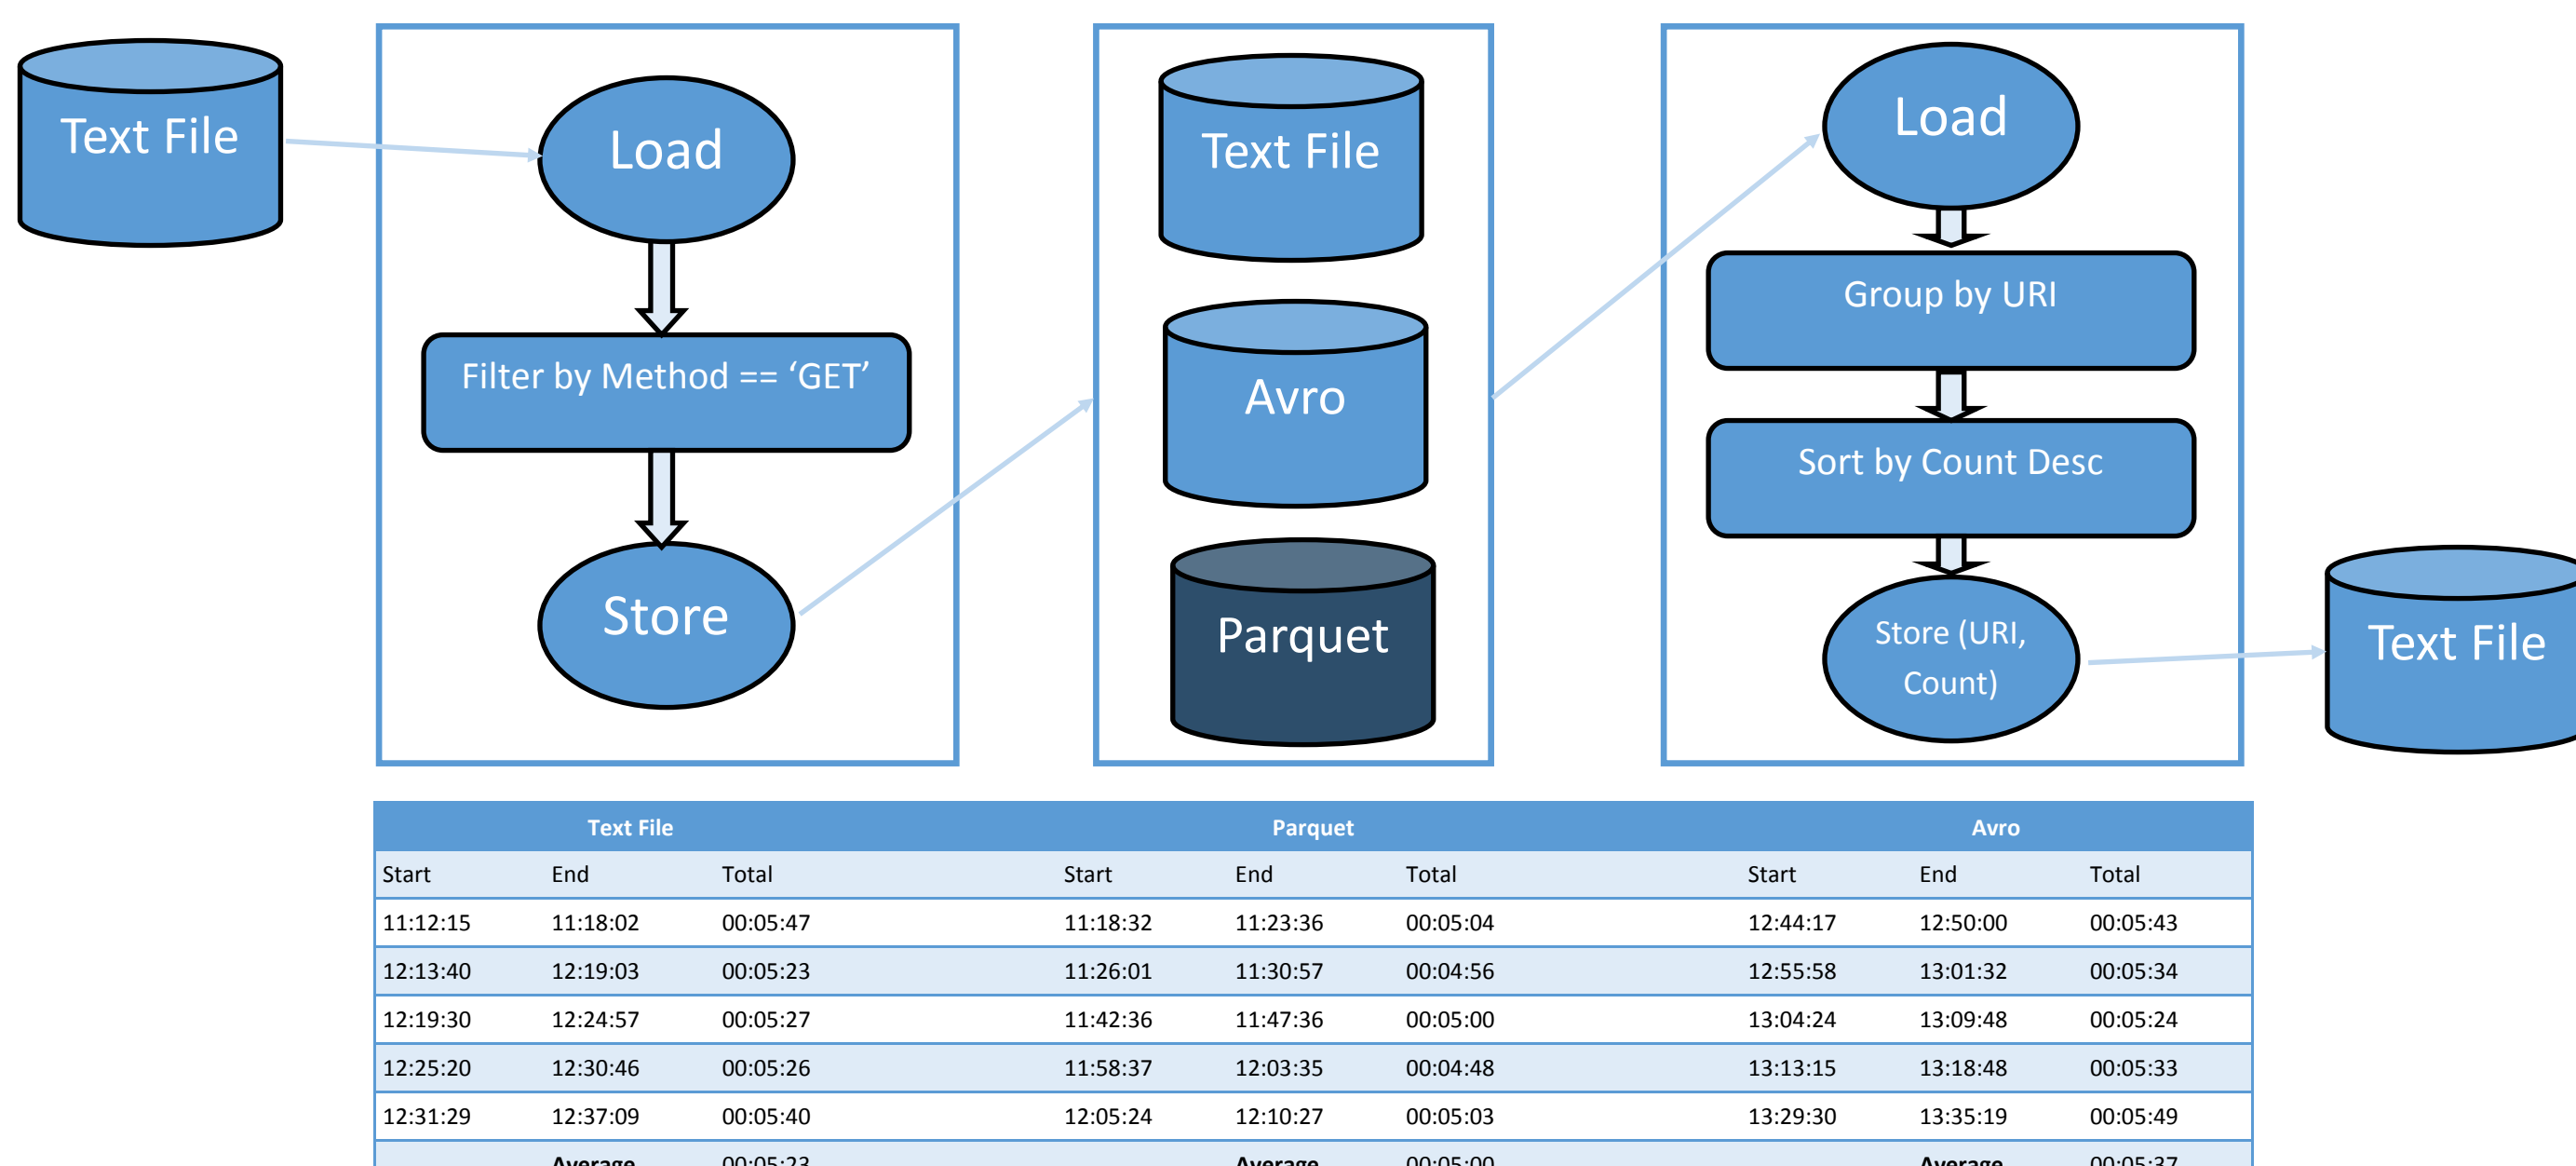
Motivation

- ⇒ Iterative workflows in real workloads
- ⇒ 80% of data re-accesses occur on the range of minutes to hours (VLDB 2012)
- ⇒ Workflows of different users share computation
- ⇒ Reuse of shared computation save storage and computation cost
- ⇒ Selection of data store based on access characteristics of intermediate results gives speedup

Our Research Objectives

- ⇒ Better utilization of the capabilities of data stores
 - ⇒ Fully utilize the features of data stores
 - ⇒ Reuse of intermediate results
 - ⇒ Data flows produce intermediate results
 - ⇒ Results can be materialized for future reuse
- ⇒ Selection of store for intermediate results
 - ⇒ Based on access characteristics of intermediate results

Evaluation



Existing Frameworks with Polyglot Persistence Support

Features	Apache Spark	Apache Drill	SQL++
Data Structure	RDD	JSON	Similar to JSON
Query Language	SparkSQL	Similar to ANSI SQL	Similar to SQL
Architecture	Connectors for different data stores	Connectors for different data stores	Mediate-Wrapper
Query Optimization	Apache Catalyst	Apache Calcite	No
Local Query Optimization	Predicates Push Down and Column Pruning	Predicates Push Down and Column Pruning	No
Scalable	Yes	Yes	No
Aggregation Framework	Yes	Yes	No
Data Locality	Yes	Yes	No
Company	Databricks	MapR	UC San Diego, USA

