

OctopusDB: A Database System for Managing Multimodal Data

Alekh Jindal, Jens Dittrich
Computer Science, Saarland University



INTRODUCTION

Research Goal

To build a one-size-fits-all database system which:


- Is able to efficiently handle different types of multimodal data
- Caters well to existing and newer data management use cases
- Adapts automatically to initial as well as changing workload
- Provides overall improved performance
- Has lower cost
- Is easier to maintain

Research Challenges

- Different types of multimodal data in a single system
- Different storage layouts under a single umbrella
- Online adaption of storage layouts
- Storage layout optimization algorithms
- Picking right layout for query processing
- Update strategies for heterogenous layouts
- Compression schemes for different storage layouts
- Simplicity versus Optimization

METHODS

Architecture



Storage Views

- Arbitrary physical representations of data
- Examples: Row, Column, Partitioned, Index, Adaptive etc.

Storage View Selection

Single optimization problem for:

- Query optimization
- View maintenance
- Index selection
- Store selection


Holistic Storage View Optimizer

- Storage View Selection
- Storage View Update Propagation
- Still, Storage Views can have specialized techniques within


RESULTS

Use-case: Flight Booking System

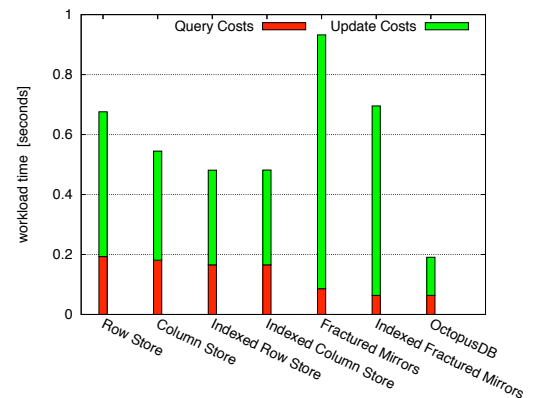
Query Costs



Update Costs



Comparison with state-of-the-art



REFERENCES

Alekh Jindal, Supervised by Prof. Jens Dittrich
The Mimicking Octopus: Towards a one-size-fits-all Database Architecture, VLDB 2010 PhD Workshop