

Adaptive Physical Design for Curated Archives

Tanu Malik
Purdue University

Xiaodan Wang, Randal Burns
Johns Hopkins University

Debarata Dash
Carnegie Mellon University

Amitabh Chaudhary
University of Notre Dame

Anastasia Ailamaki
Ecole Polytechnique Fédérale de Lausanne

Introduction

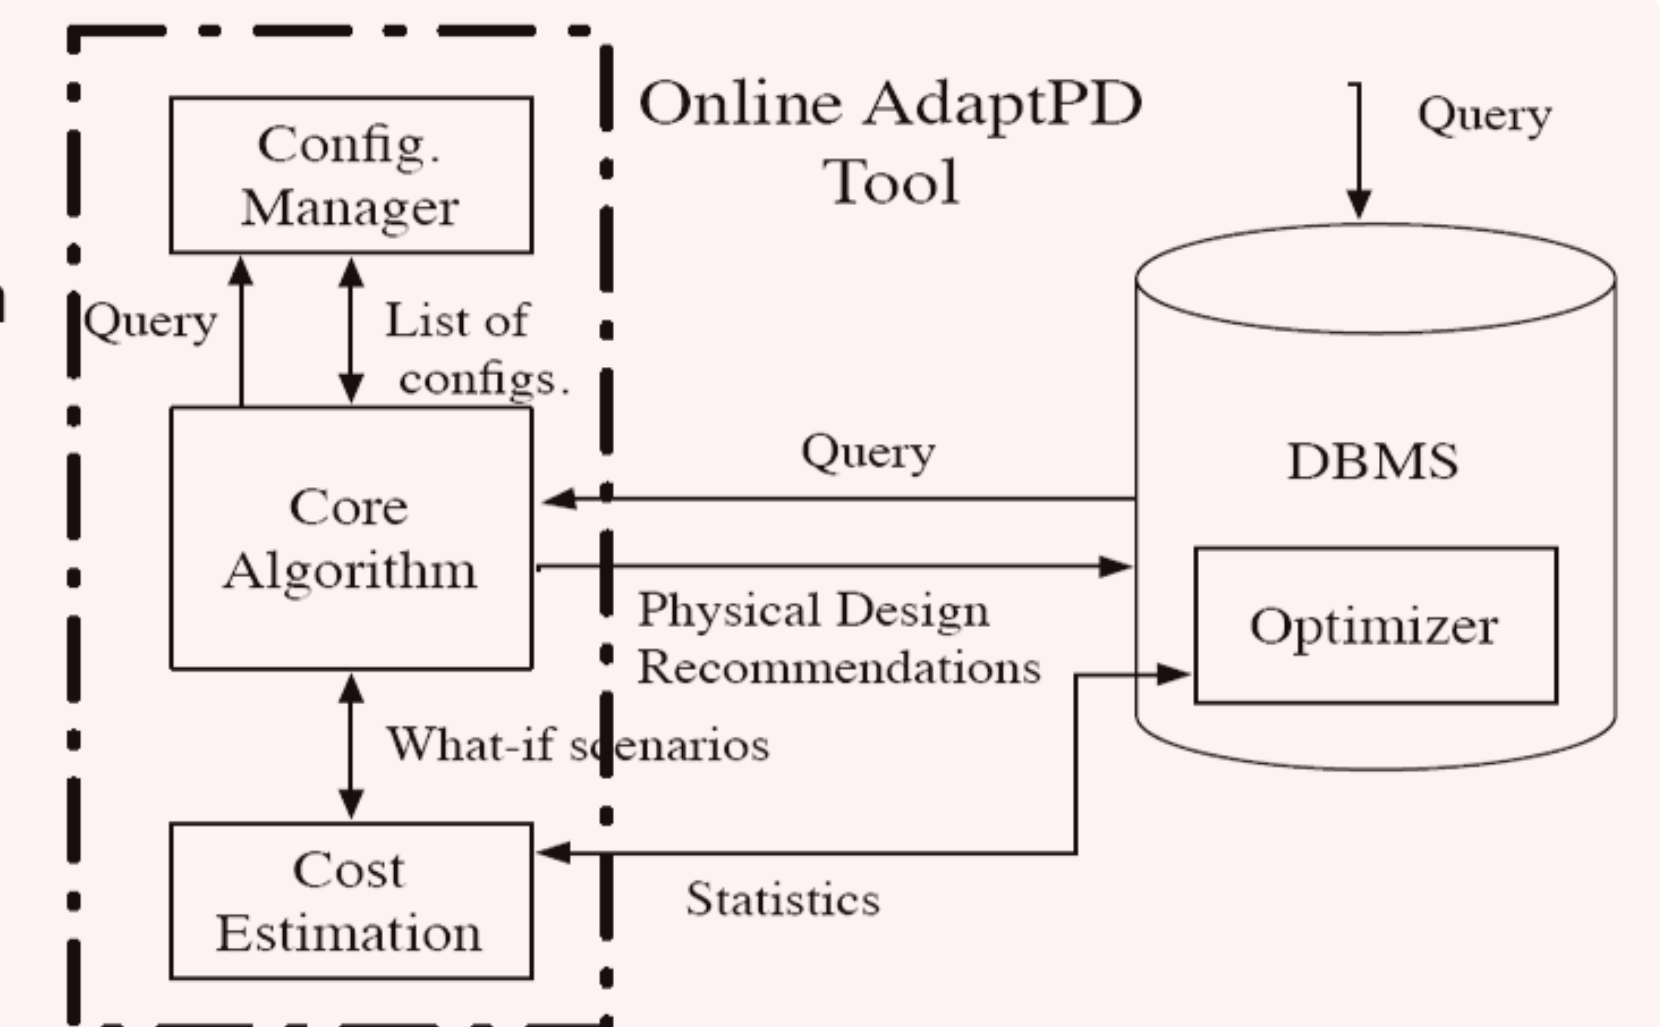
- Libraries are the new data centers that curate large scientific data such as the SDSS databases [1].
- It is crucial for curated databases to deliver optimal performance and to reduce DBMS' total cost of ownership.
- Automated physical design tools help achieve this goal.
- Current tools are partially automated—they do not adapt physical design continuously to changes in the workload.
- We present *AdaptPD*, a tool that quickly detects and adapts to changes in the workload.

Current Approaches

- Current physical design tools provide partial automation [2][3]:
 - Require an explicit specification of a representative workloads for tuning
 - Do not provide relative benefits of implementing the recommendations

AdaptPD

- Identifies when physical design changes are required
- Accounts for the cost of reorganization
- Efficient cost estimation for query and transition costs



Problem

Online Problem


- Given finite sequence of queries $\sigma = q_1, \dots, q_n$
- $S = \{S_1, \dots, S_N\}$ be the set of configurations
- Function $d: S \times S \rightarrow \mathbb{R}^+$ measures transition cost between two configurations
- In AdaptPD, the objective is to find sequence $S = (S_0, S_1, \dots, S_n)$ of configurations which minimizes:

$$\text{cost}(\sigma, S) = \sum_{i=1}^n q_i(S_i) + \sum_{i=0}^{n-1} d(S_i, S_{i+1})$$

Core Algorithm

Core algorithm is based on two insights:

- When $N=2$ and $d(\cdot)$ is symmetric, the AdaptPD problem is similar to the Online Ski Rental Problem [4]
 - Costs no more than four times the optimal
- With large N and $d(\cdot)$ is asymmetric, we can transform directed S into undirected H where costs are symmetric
 - Does not increase cost more than $8(N-1)\rho$ optimal
 - ρ is a factor that measures amount of asymmetry ($\rho=10$ in the Figure)

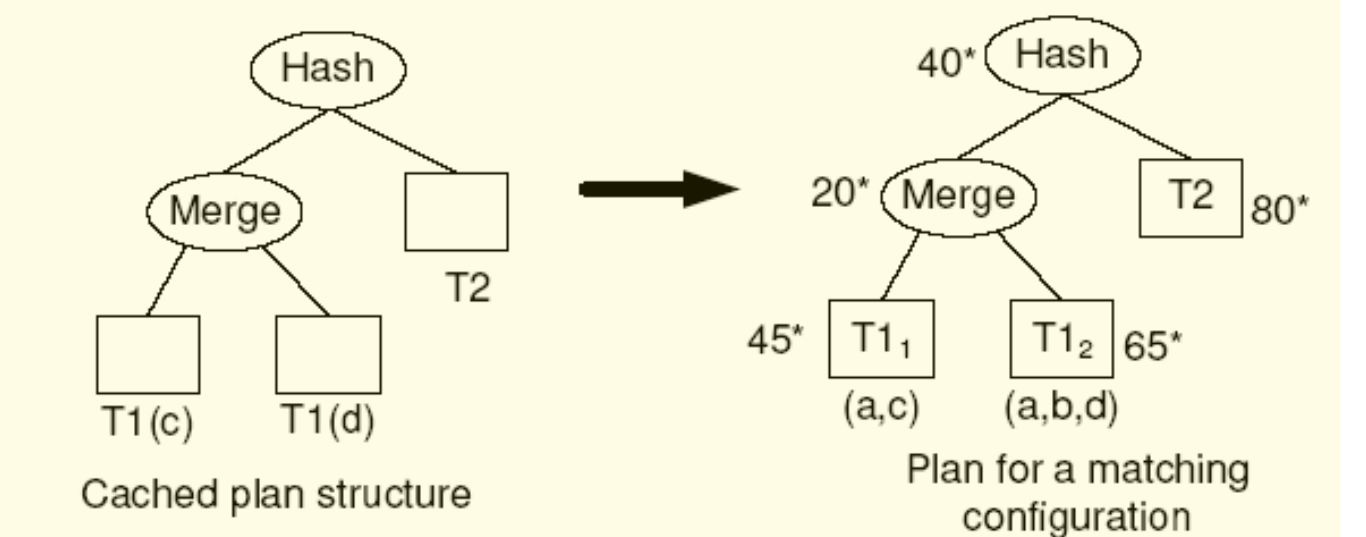


Cost Estimation

Query Cost Estimation

Key idea is to invoke optimizer less often by caching and reusing query plans that remain invariant, which is defined as:

- Having the same # of partitions
- Division of predicate columns are the same
- Page size distributions are similar



Transition Cost Estimation

Key idea is to Capture cost of frequent transitions using bulk-insert for repartitioning. Bulk-insert is the most efficient method to copy data.

$$d(c_i, c_j) = \sum_{t \in \Delta_{ij}} BCP(t)$$

Config. Manager

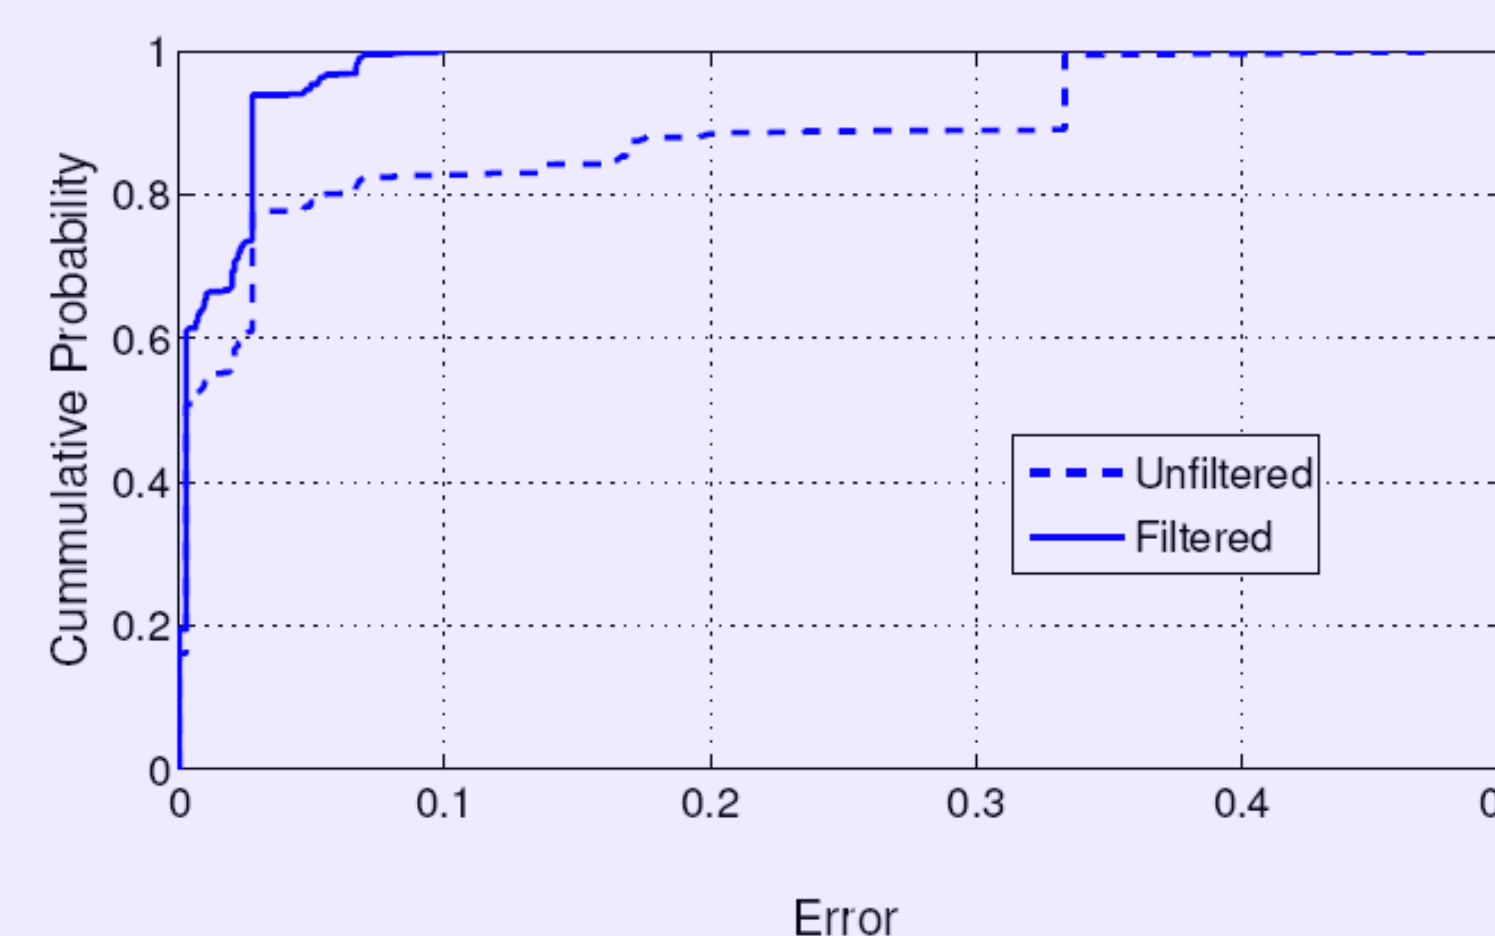
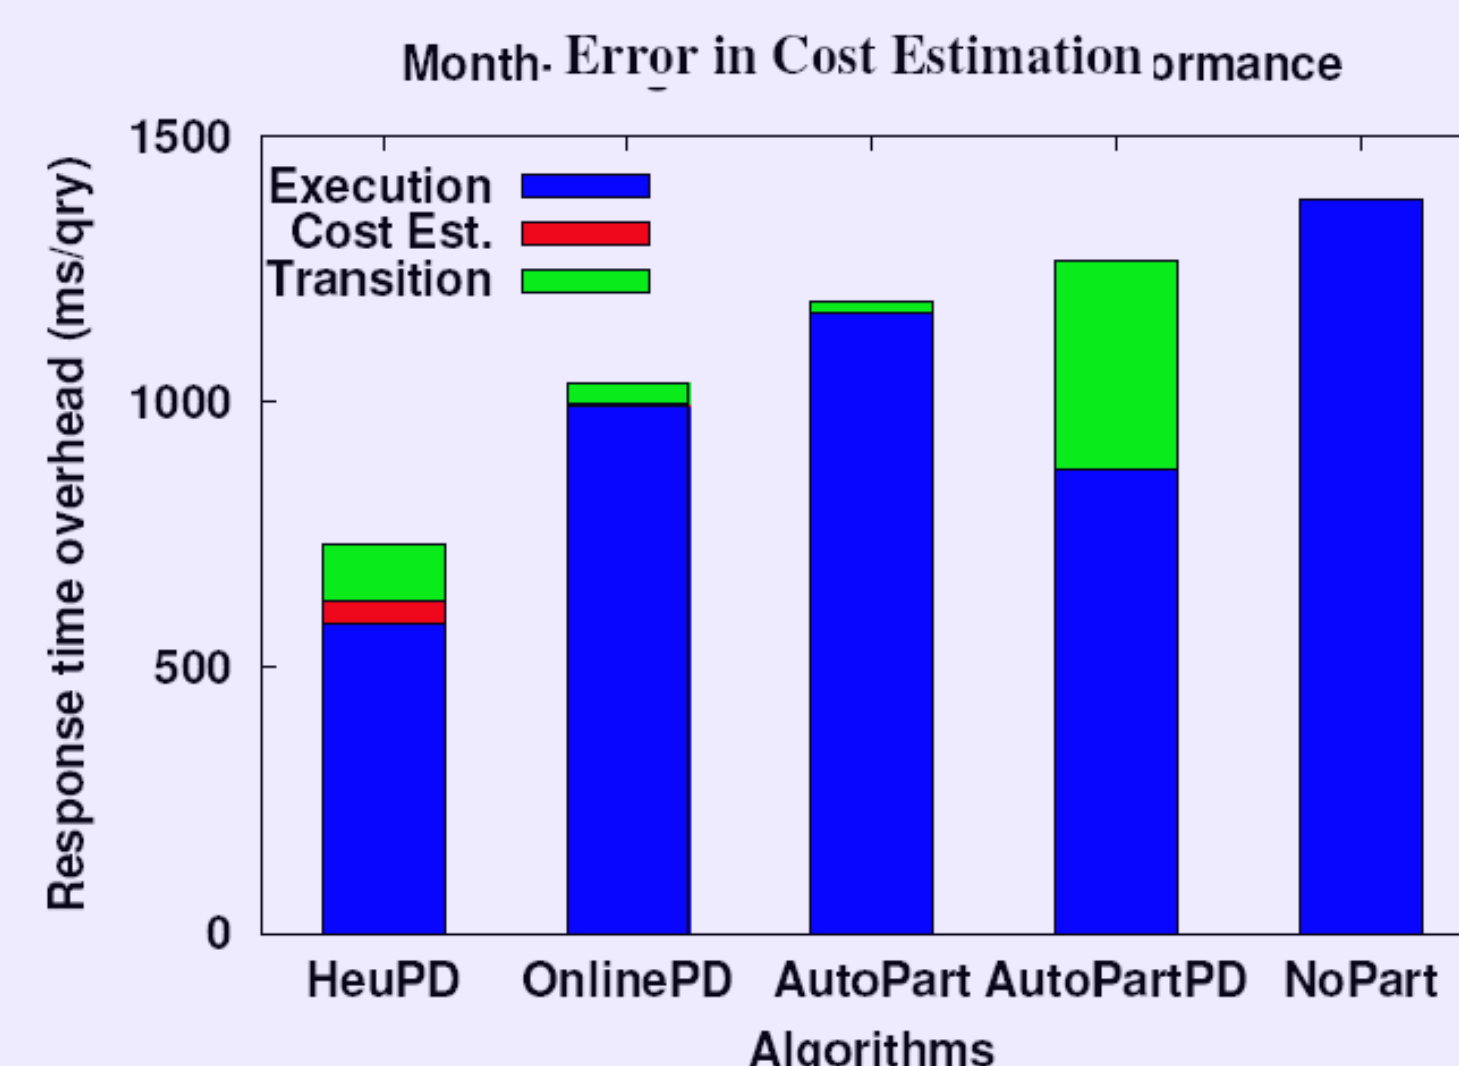
- A prototype is a set of attributes a query accesses
- Candidate configs are determined from recurring query prototypes in the workload
- The following queries form the same prototype:

```
SELECT objID, ra, dec FROM PhotoPrimary
WHERE dec between 2.25 and 2.75
```

```
SELECT top 1 ra, dec FROM PhotoPrimary
WHERE objID = 5877311875315
```

Performance Evaluation

- Analyze query execution cost, cost estimation overhead, and transition cost for month-long, 1.4 million trace from SDSS.
- Compare $8(N-1)\rho$ -competitive alg (*OnlinePD*) with a highly tuned heuristic alg (*HeuPD*), an offline partitioning tool (*AutoPart*), and no partitioning (*NoPart*).
- Query cost estimation exhibit less than 5% error for 94% of the queries.



Conclusion

- *AdaptPD* automates vital DBA tasks such as:
 - Estimating “when” to tune the current physical design
 - Finding representative workloads
- *AdaptPD* gives DBA a quantitative reason for choosing the recommended design.
- We are extending the tool to include automation of index design.

- [1] <http://www.sdss.org>: The Sloan Digital Sky Survey
- [2] Papadomanolakis, @., Ailamaki, A.: AutoPart: Automating Schema Design for Large Scientific Databases Using Data Partitioning. In: SSDBM. (2004)
- [3] Agrawal, S., Narasayya, V.R., Yang, B.: Integrating Vertical and Horizontal Partitioning Into Automated Physical Database Design. In: SIGMOD. (2004)
- [4] Borodin, A., El-Yaniv, R.: Online Computation and Competitive Analysis. Cambridge University Press, New York, NY, USA (1998)