DIAS: Data-Intensive Applications and Systems Laboratory
School of Computer and Communication Sciences
Ecole Polytechnique Federale de Lausanne
Building BC, Station 14
CH-1015 Lausanne
URL: http://dias.epfl.ch

Physical Database Designers: A Comparison of Robustness

(Fall Semester Sep-Dec 2010)

Keywords: Physical Design, Robustness

Problem: Database Management Systems (DBMSs) have become extremely popular the last decades and they support a large number of different applications. Those applications along with the explosion of available data have increased the importance of physical design since the selection of the correct physical structures (e.g. indexes, partitions and materialized views) may improve the query execution performance by orders of magnitude. Commercial DBMS vendors have recognized the problem and they offer automated physical design tools in their products [1][9]. Those tools use what-if interfaces [5] to simulate the presence of different physical structures and recommend physical designs that minimize the estimated execution time for a given workload. However, there is not a clear method of comparing the quality of the proposed solutions [2]. Additionally, considering only the estimated execution time can lead to physical designs that are not robust due to induced error from the query optimizer's estimations and changes in the workload [4].

Project: In this project, a method of evaluating the robustness of physical designers' recommendations is proposed. The robustness of different physical designs is measured using template queries from the standard TPC-H benchmark. Experiments will be performed using DB2, MS SQL Server and PostgreSQL. The study plan for this project is the following:

1. Literature review of relevant papers in the area of Automated Physical Design.
2. Configuration of different DBMSs.

3. Benchmarking and analysis using TPC-H [7] and SDSS [8] workloads.
4. Preparation and presentation of a short report based on the experimental results.

Supervisor:	Prof. Anastasia Ailamaki, anastasia.ailamaki@epfl.ch
Responsible collaborator:	Ioannis Alagiannis, ioannis.alagiannis@epfl.ch
Duration:	4 months

References

1. S. Agrawal et al. Database Tuning Advisor for Microsoft SQL Server 2005. In Proceedings of the International Conference on Very Large Databases (VLDB), 2004.
2. S. Chaudhuri and V. Narasayya, "Self-Tuning Database Systems: A Decade of Progress," in Proceedings of the 33rd International Conference on Very Large Databases, Vienna, Austria, 2007
3. D. Dash, A. Ailamaki. CoPhy: Automated Physical Design with Quality Guarantees. Technical Report CMUCS-10-109
4. K. El Gebaly and A. Aboulnaga, "Robustness in Automatic Physical Database Design", EDBT '08, Nantes, France, pages 145-156, March 2008.
5. S. Finkelstein,M. Schkolnick, P. Tiberio "Physical database design for relational databases" ACM ToDS. 1988
6. S. Papadomanolakis, D. Dash, A. Ailamaki. Efficient Use of the Query Optimizer for Automated Physical Design. VLDB 2007.
7. TPC Benchmark H. Decision Support. http://www.tpc.org
8. http://www.sdss.org/
9. D. C. Zilio, J. Rao, et al. DB2 Design Advisor: Integrated Automatic Physical Database Design. In Proceedings of the 30th Intl. Conf. on Very Large Data Bases, pages 1087–1097, 2004.