



Technische Universität Berlin
Fakultät IV Elektrotechnik und Informatik
Fachgebiet Regelungssysteme
Leitung: Prof. Dr.-Ing. Jörg Raisch
Skr. EN 11, Tel. 314-22999



Vortragsankündigung

Seminar Regelungssysteme LV 0430 L 654

Montag, 23. August 2010, 16:00 Uhr

Vortragsort: EN 223

Prof. Dr. rer. nat. Volker Markl

Fachgebiet Datenbanksysteme und Informationsmanagement, Technische Universität Berlin

“Learning in Query Optimization”

Database Systems let users specify queries in a declarative language like SQL. Most modern DBMS optimizers rely upon a cost model to choose the best query execution plan (QEP) for any given query. Cost estimates are heavily dependent upon the optimizer’s estimates for the number of rows that will result at each step of the QEP for complex queries involving many predicates and/or operations. These estimates, in turn, rely upon statistics on the database and modeling assumptions that may or may not be true for a given database. In my talk, I will present an overview of the research on learning in query optimization. I will introduce LEO, DB2’s LEarning Optimizer, as a comprehensive way to repair incorrect statistics and cardinality estimates of a query execution plan. By monitoring executed queries, LEO compares the optimizer’s estimates with actuals at each step in a QEP, and computes adjustments to cost estimates and statistics that may be used during the current and future query optimizations. LEO introduces a feedback loop to query optimization that enhances the available information on the database where the most queries have occurred, allowing the optimizer to actually learn from its past mistakes.

In the second part of the talk, I describe how the knowledge gleaned by LEO is exploited consistently in a query optimizer, by adjusting the optimizer’s model and by maximizing information entropy.