

Factorized Databases: Past and Future Past

Dan Olteanu

Computer Science Department University of Oxford
dan.olteanu@cs.ox.ac.uk

Abstract. In this talk I will overview the FDB project at Oxford on succinct, lossless representations of relational data that I call factorized databases. I will first present a characterization of the succinctness of results to conjunctive queries and how factorizations can speed up query processing. I will then comment on how this succinctness characterization relates to seemingly disparate results on: readability of provenance polynomials, representation systems for incomplete information, one-pass query evaluation using finite cursor machines, tractability in probabilistic databases, and parallel query evaluation with one synchronization step. I will conclude with two near-future projects that brought me back to factorized data representations: scalable machine learning over relational data and distributed database systems with low communication cost.