

Provenance-Directed Chase and Backchase

Alin Deutch, UCSD

The Chase&Backchase algorithm for rewriting queries using views is based on constructing a canonical rewriting candidate called a universal plan (during the chase phase), then chasing its exponentially many subqueries in search for minimal rewritings (during the backchase phase). We show that the backchase phase can be sped up significantly if we instrument the standard chase to maintain provenance information. The particular provenance flavor required is known as minimal why-provenance in the literature, and it can be computed by exploiting the analogy between a chase step execution and query evaluation.