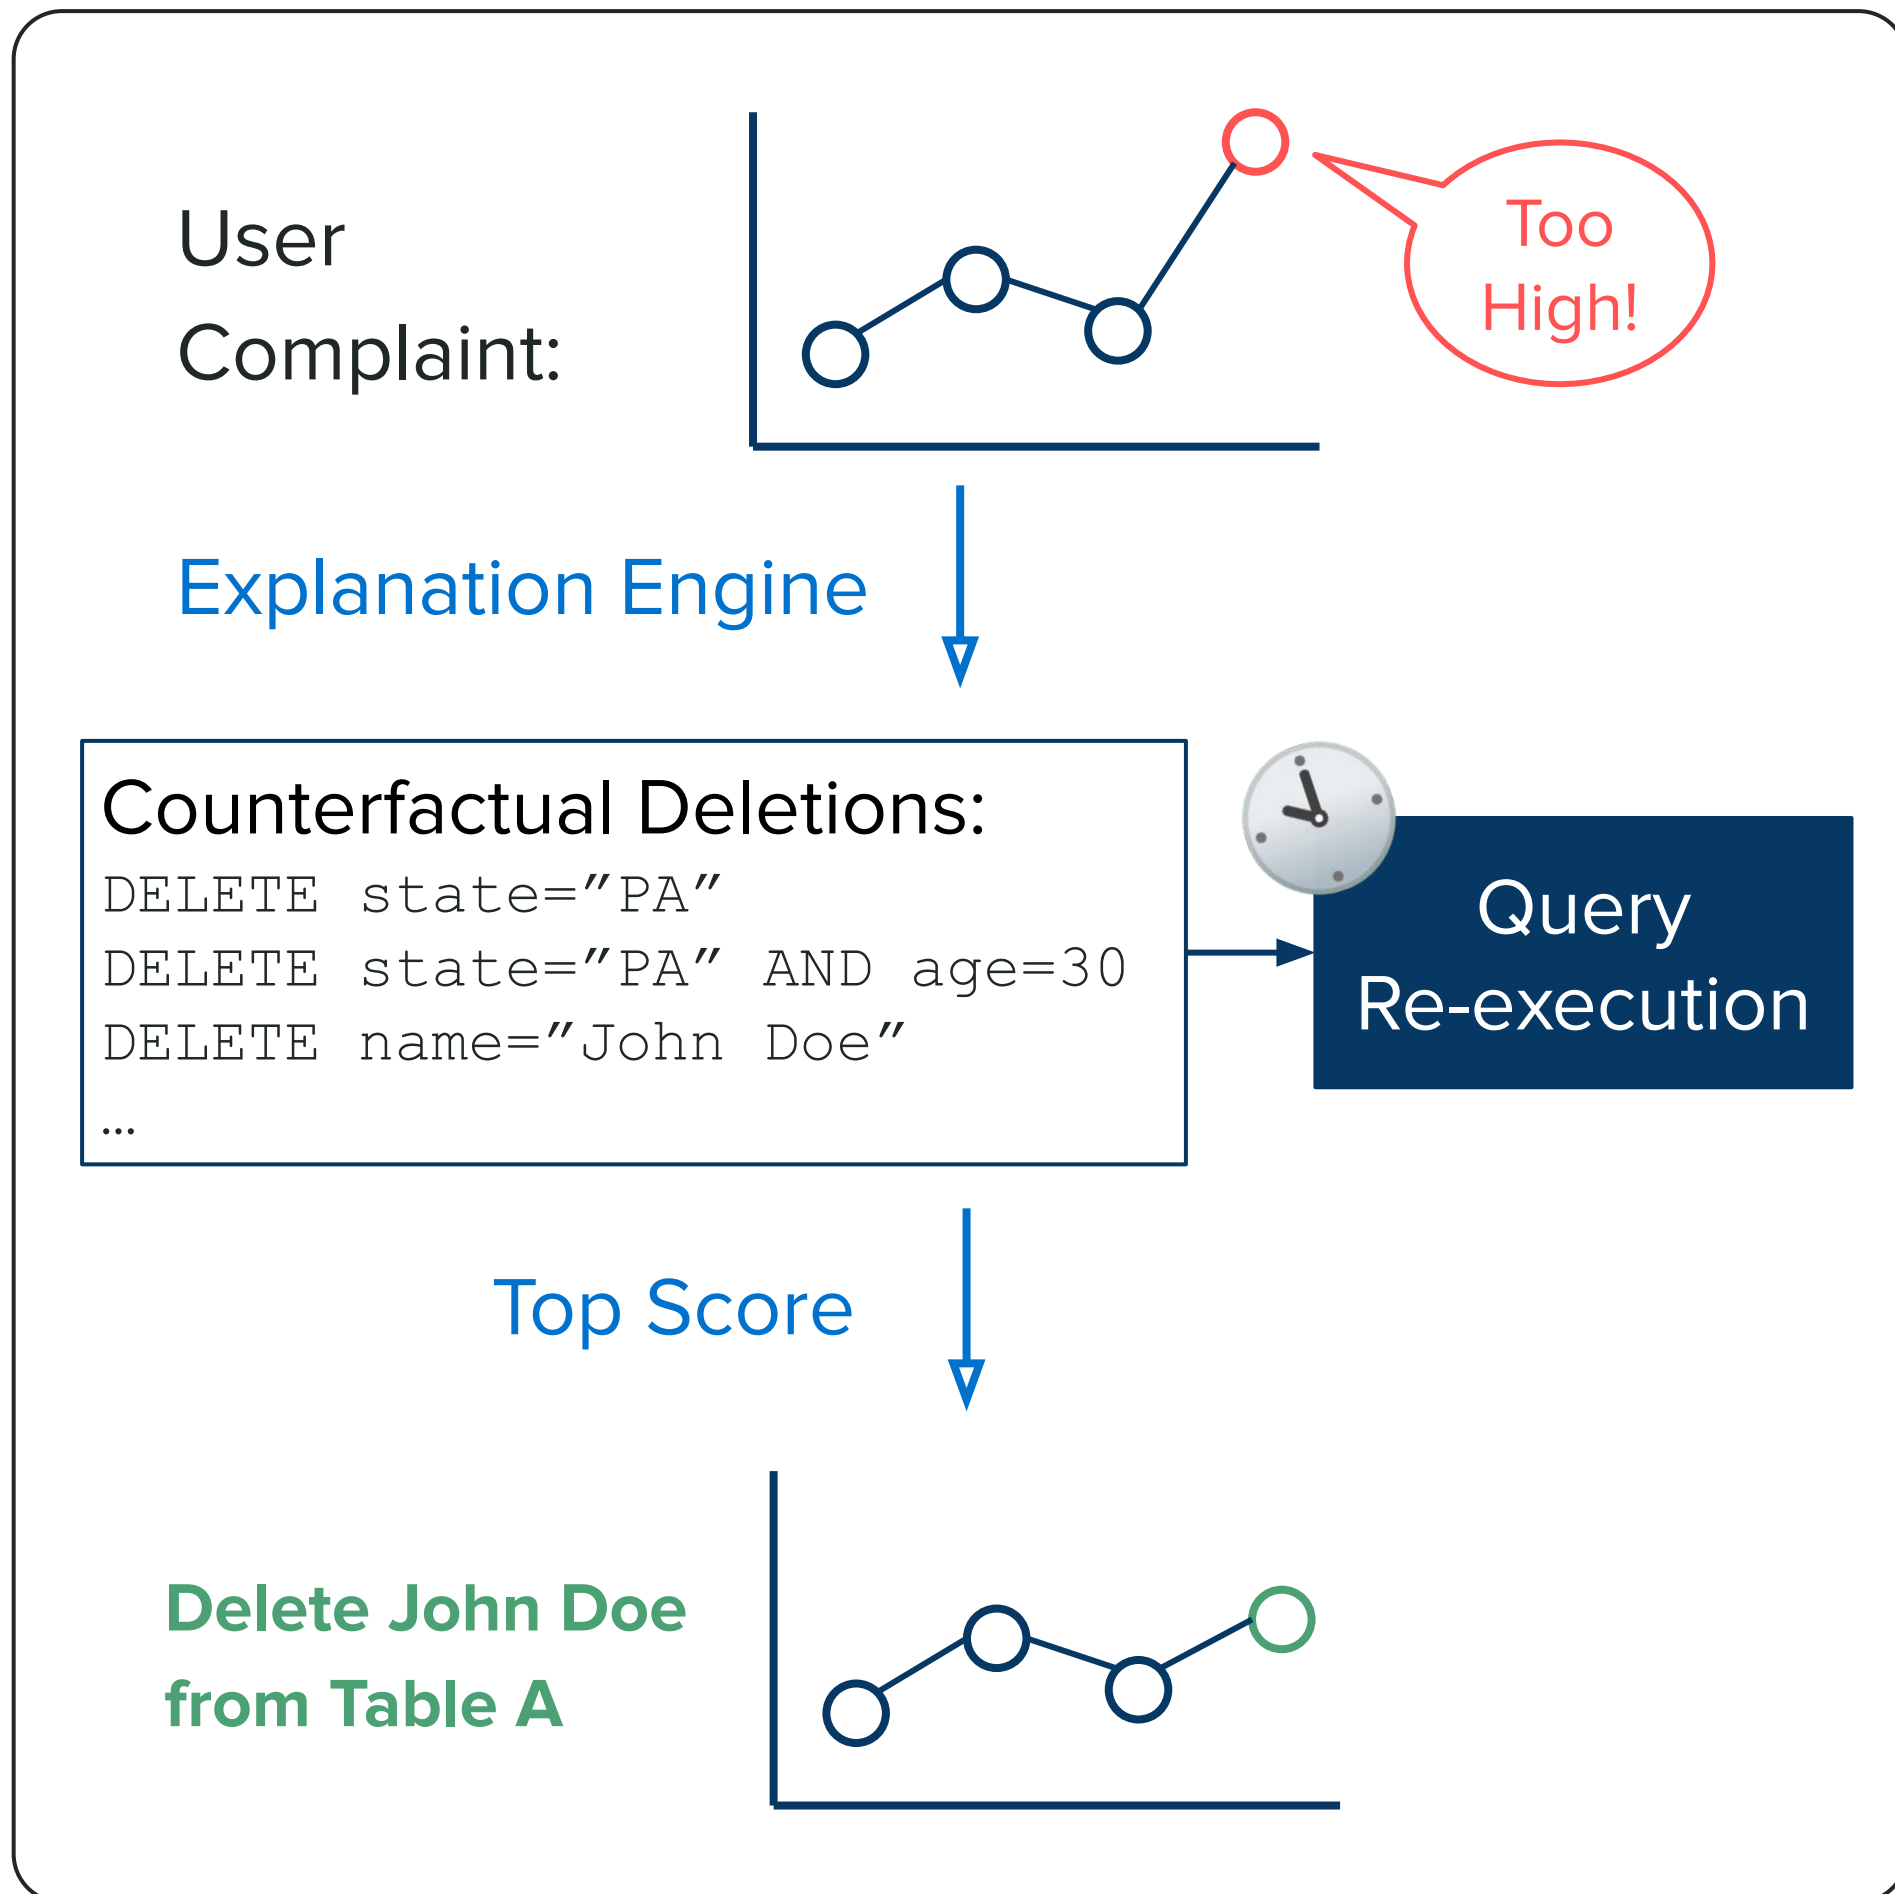


Accelerating Query Explanations Using Fine-Grained Provenance

Alex Yao, Columbia University

Background & Motivation



Problem Statement

Explanation engines are bottlenecked by **Query Re-execution** ⌚, which naively needs to run the original query once for each proposed deletion.

Existing Approach 1: IVM

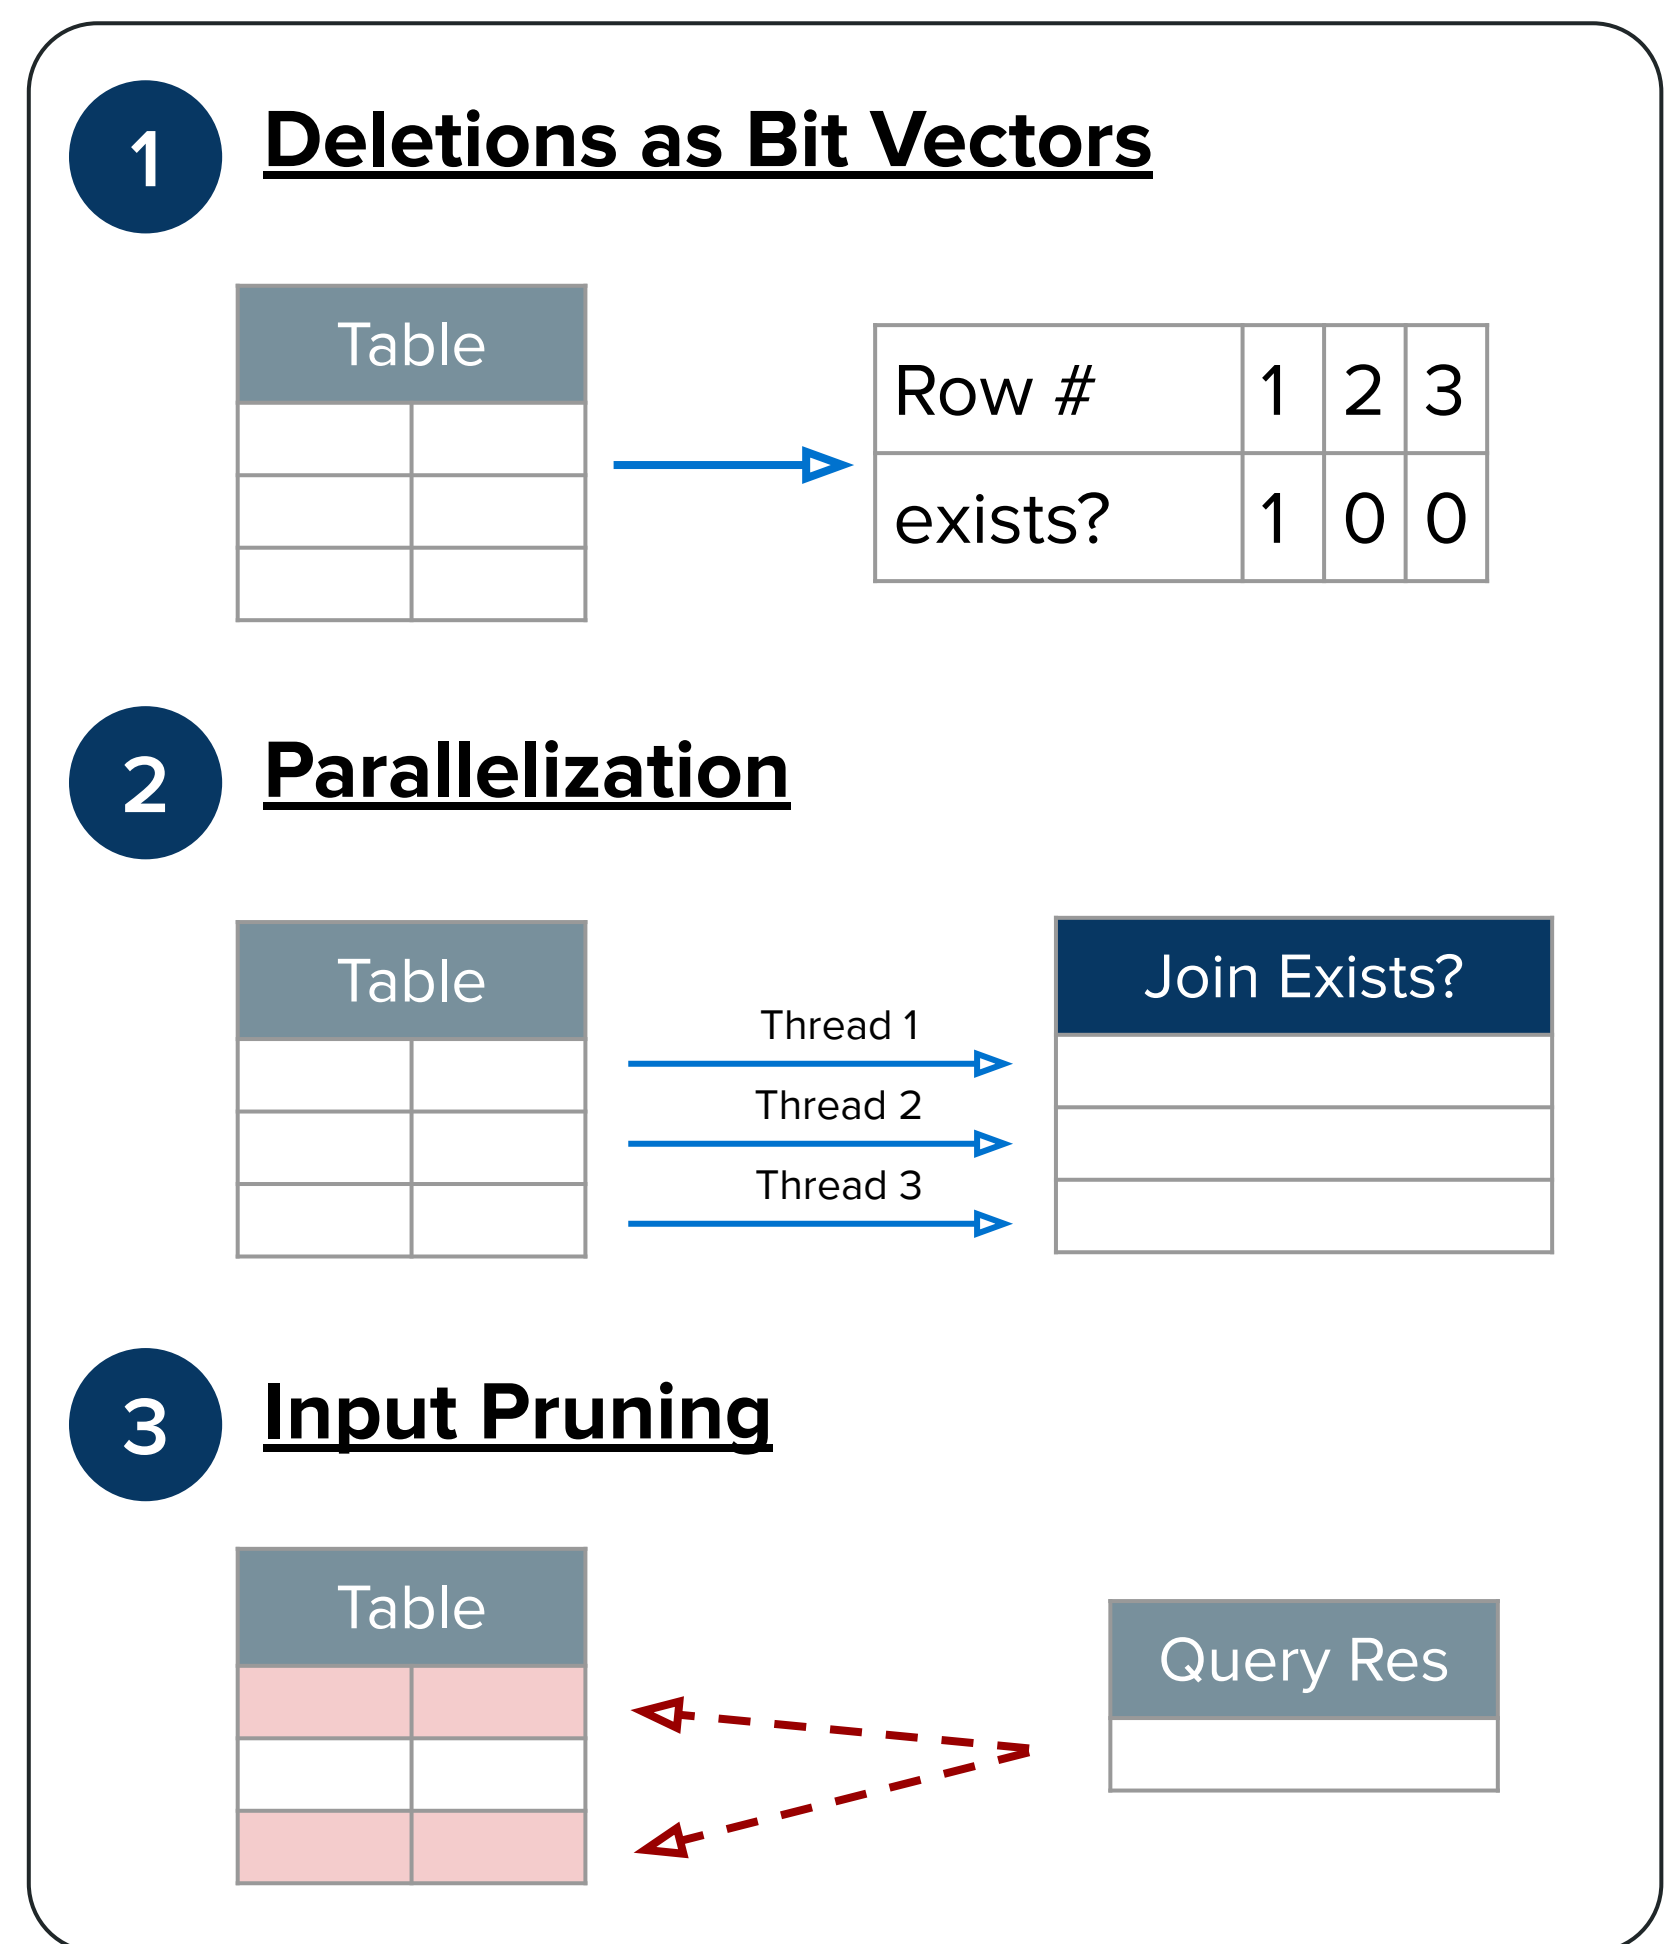
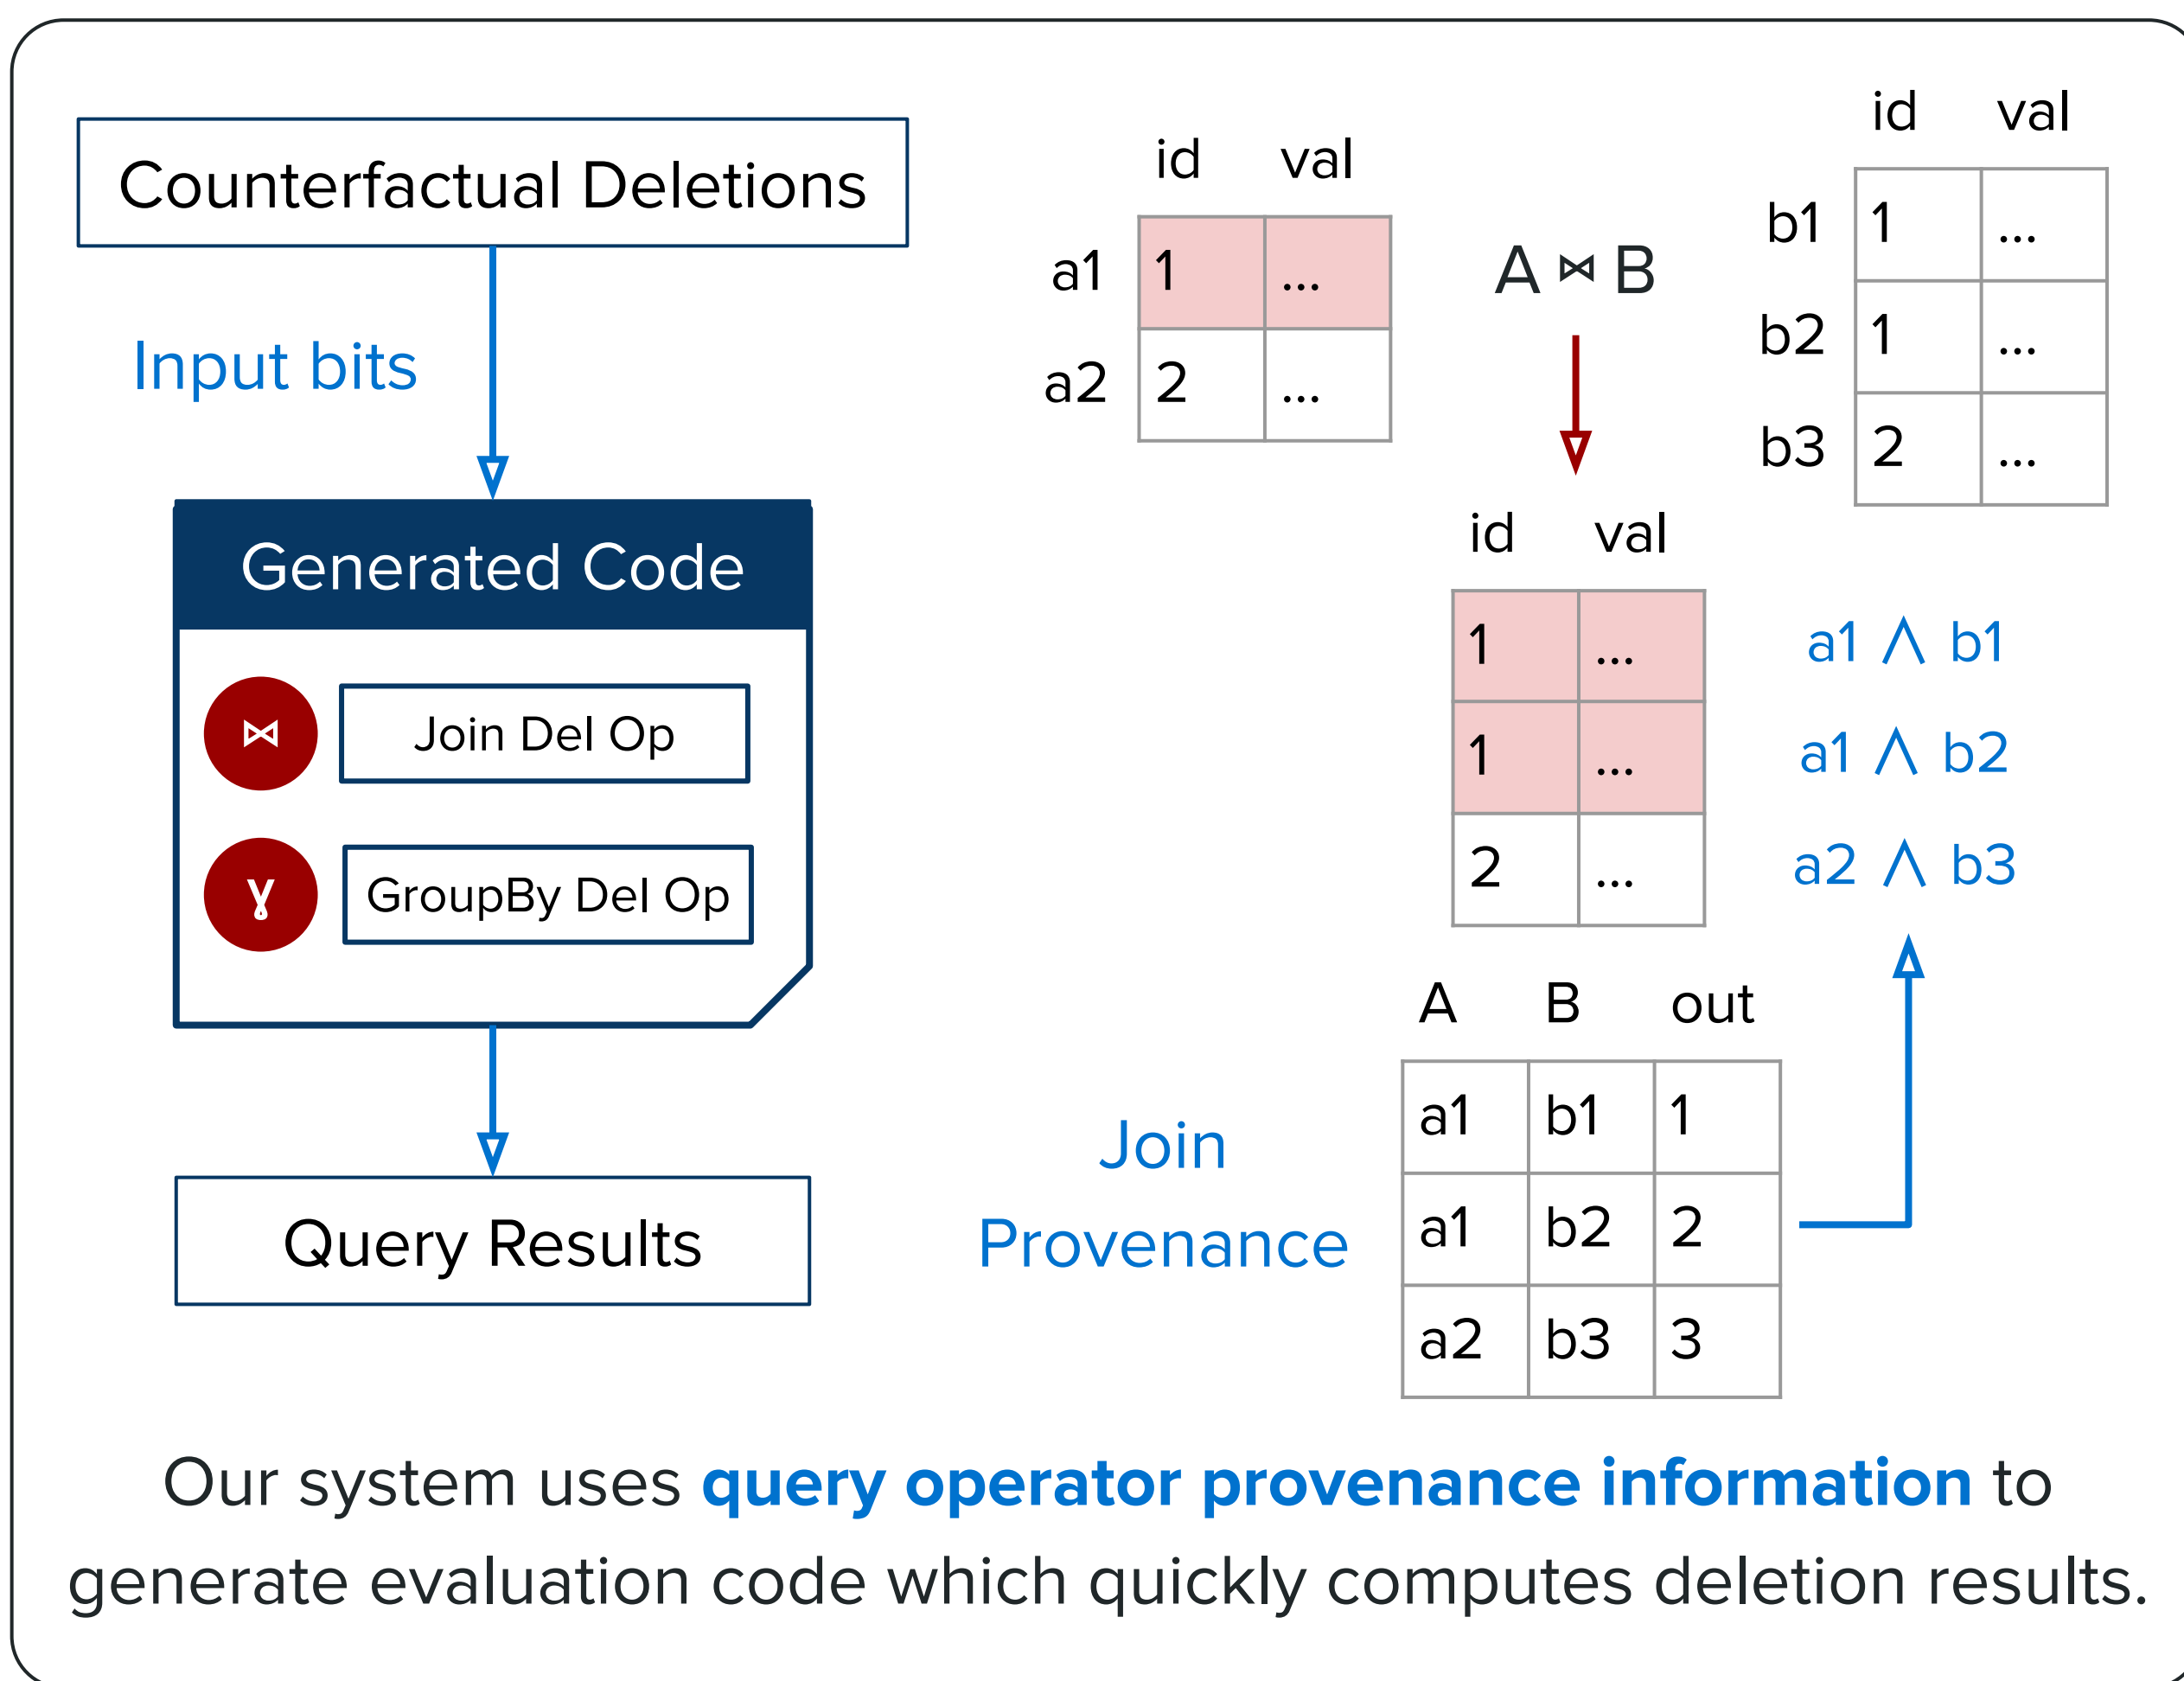
Poor performance over large deletions and complex queries

$$(A - \nabla A) \bowtie (B - \nabla B) = (A \bowtie B) - (\nabla A \bowtie B + A \bowtie \nabla B - \nabla A \bowtie \nabla B)$$

Existing Approach 2: Provenance Circuits

Conceptually great, but impractical, inefficient evaluation logic

Our Approach



Experiments & Results

