



HARVARD

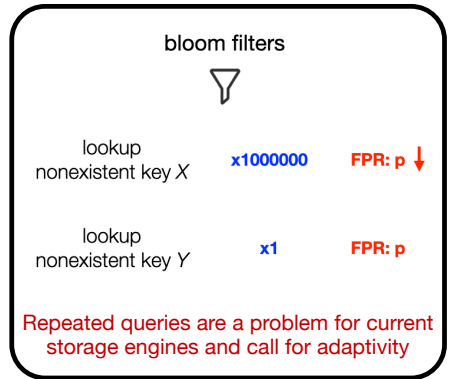

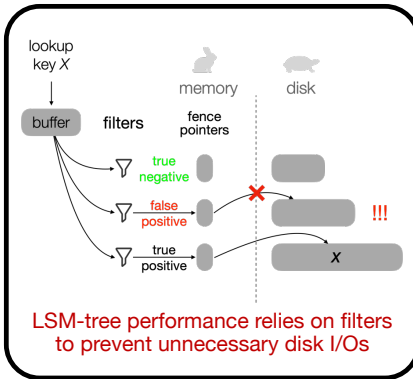
John A. Paulson
School of Engineering
and Applied Sciences

Workload-Adaptive Filtering in Storage Engines

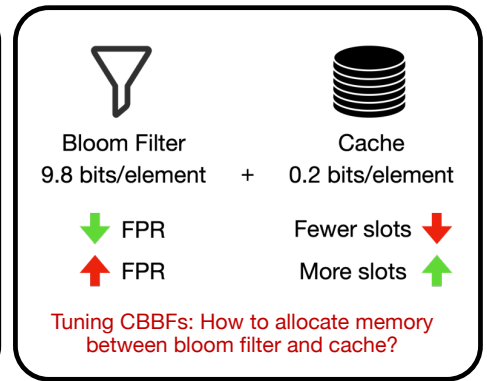
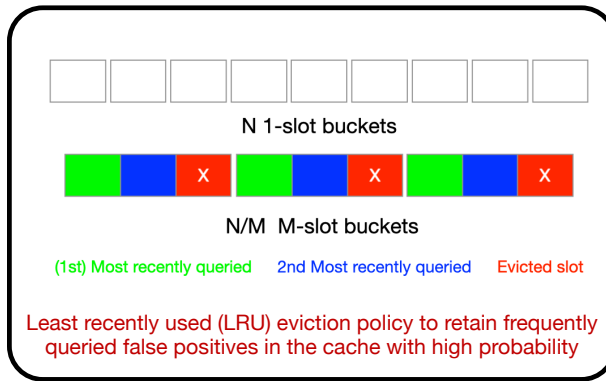
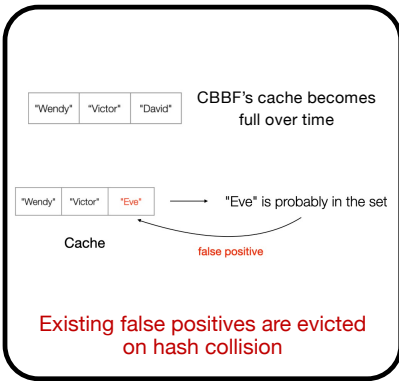
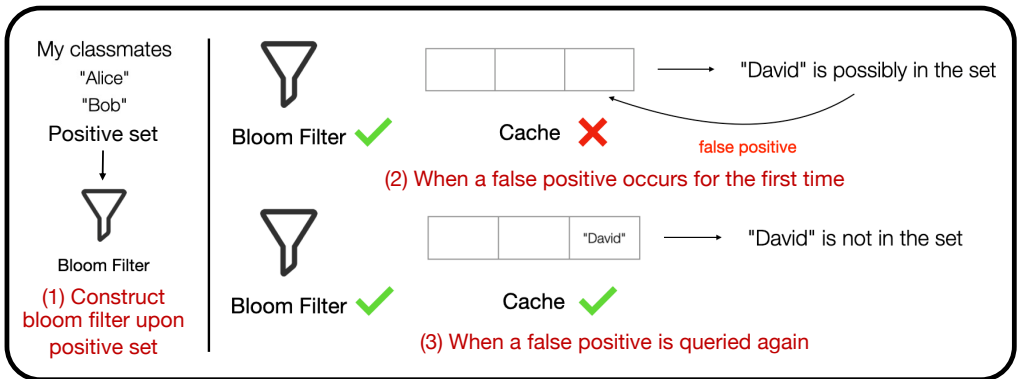
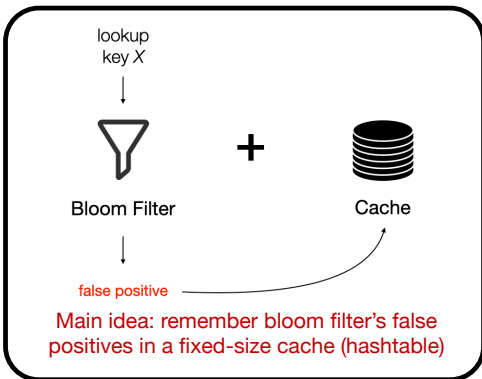
Joshua Pan | Advised by Brian Henschel and Stratos Idreos

The Problem: LSM-tree storage engines suffer with empty queries and skew workloads

LSM-tree based storage engines are everywhere



Our Contribution: Cache-Backed Bloom Filters (CBBFs)



End-to-End Storage Engine Throughput Improvement

