Predicting Query Execution time for JIT Compiled Database Engines

Konstantinos Chasialis, Srinivas Karthik,
Bikash Chandra, Anastasia Ailamaki

DATA-INTENSIVE APPLICATIONS AND SYSTEMS

Motivation & Problem

Use Cases

- Query admission
- Improve query scheduling
- Better resource allocation

JIT Analytical Engines

- Typically, majority spent on memory access and branch misprediction for in-memory systems

Prior Works

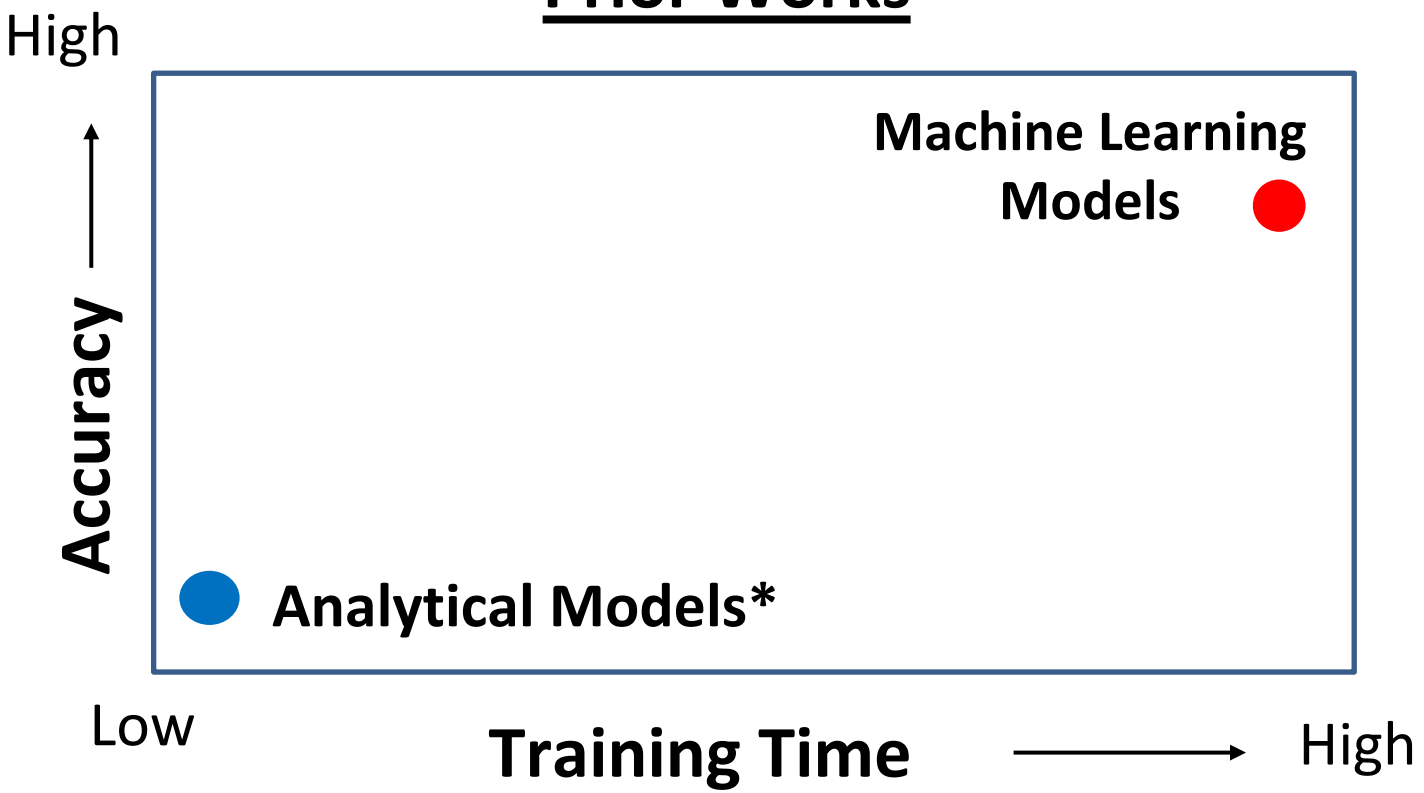

* Manegold et al. Generic database cost models for hierarchical memory systems

Quick and near accurate prediction of execution time is highly desirable

Baseline Analytical Model (Manegold et al)

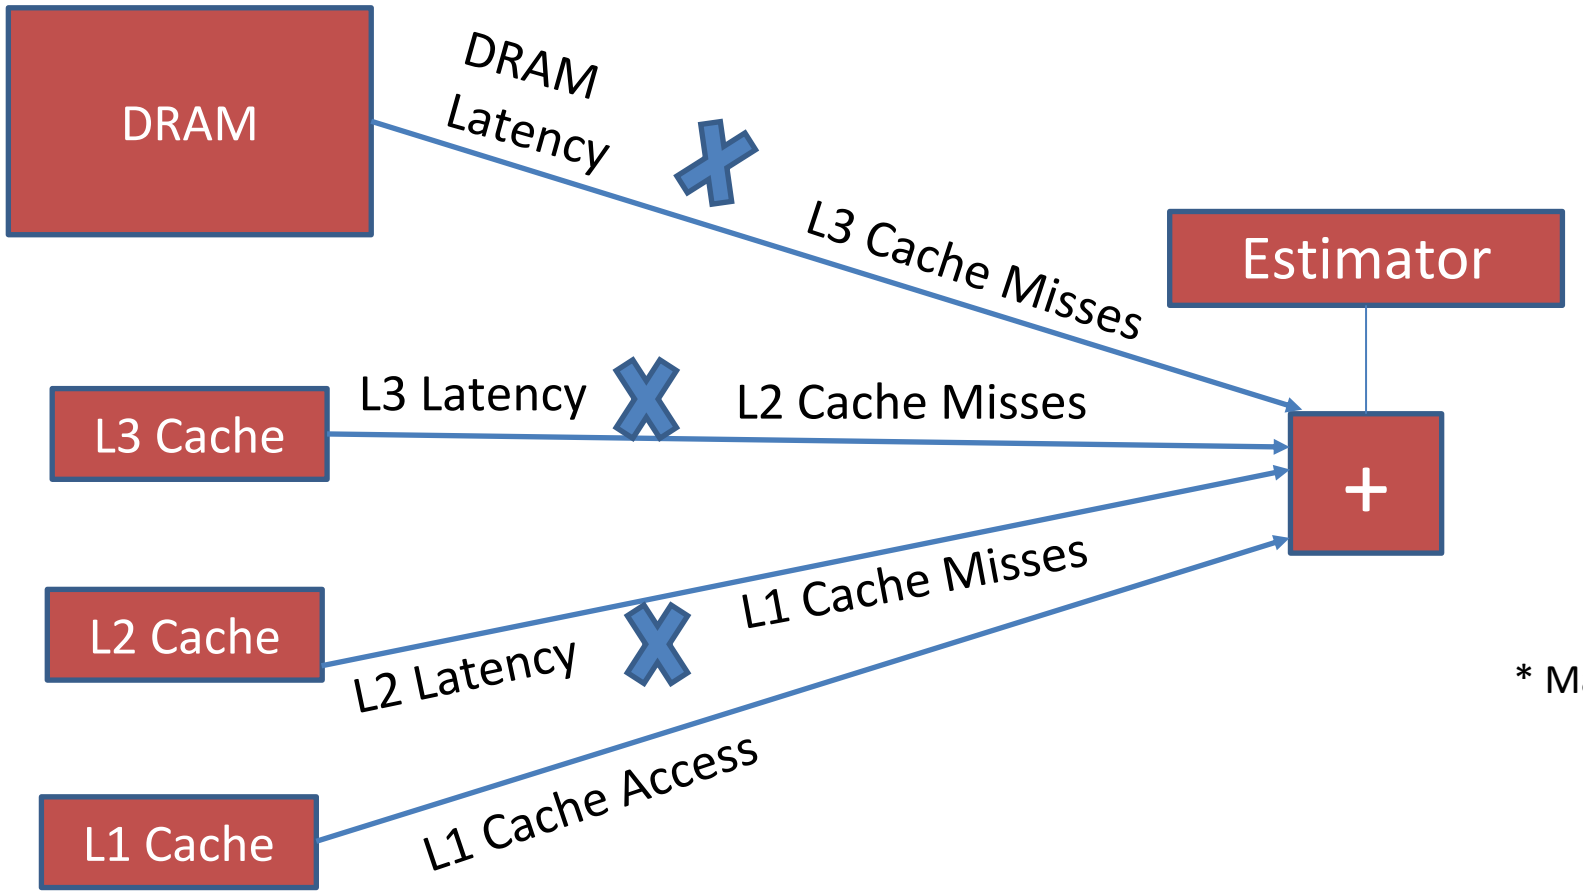

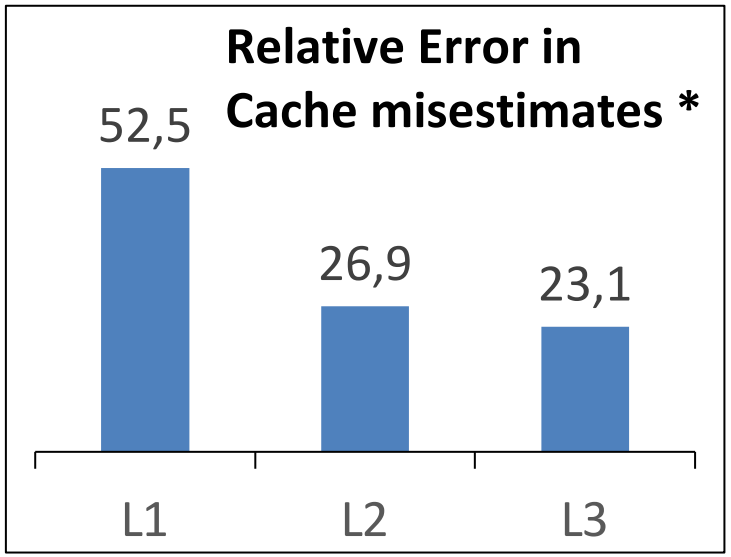

* Manegold et al. Generic database cost models for hierarchical memory systems

Analytical models have low accuracy for runtime parameters

JIT Prediction: Analytical models + JIT Calibration

$$Tmem = M_1 * l_2 + M_2 * l_3 + M_3 * l_4$$

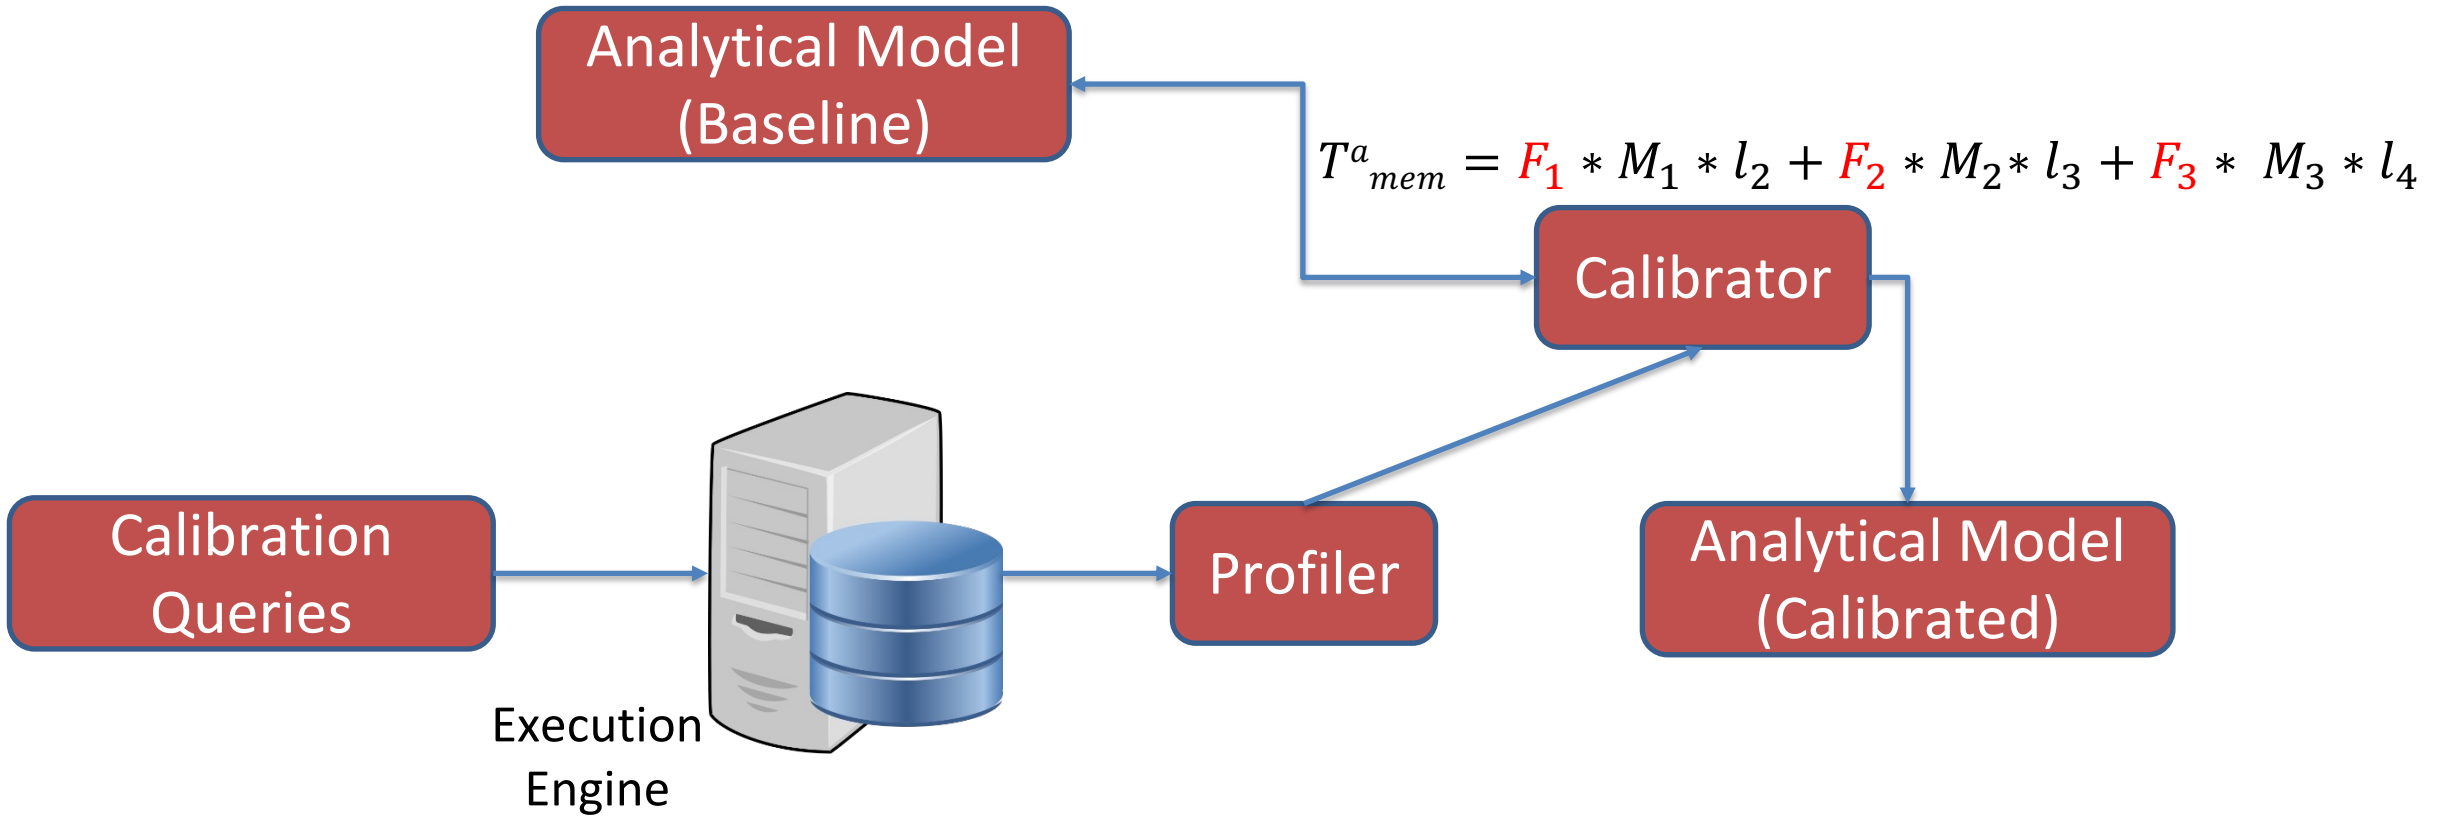

Low overhead calibration can significantly improve accuracy

Experimental Results

Setup

Intel(R) Xeon(R) Gold 5118
L1: 32KB
L2: 1MB
L3: 16MB (non-inlusive)
RAM: 360 GB
Database: SSBM SF100
Proteus – engine based on JIT compilation

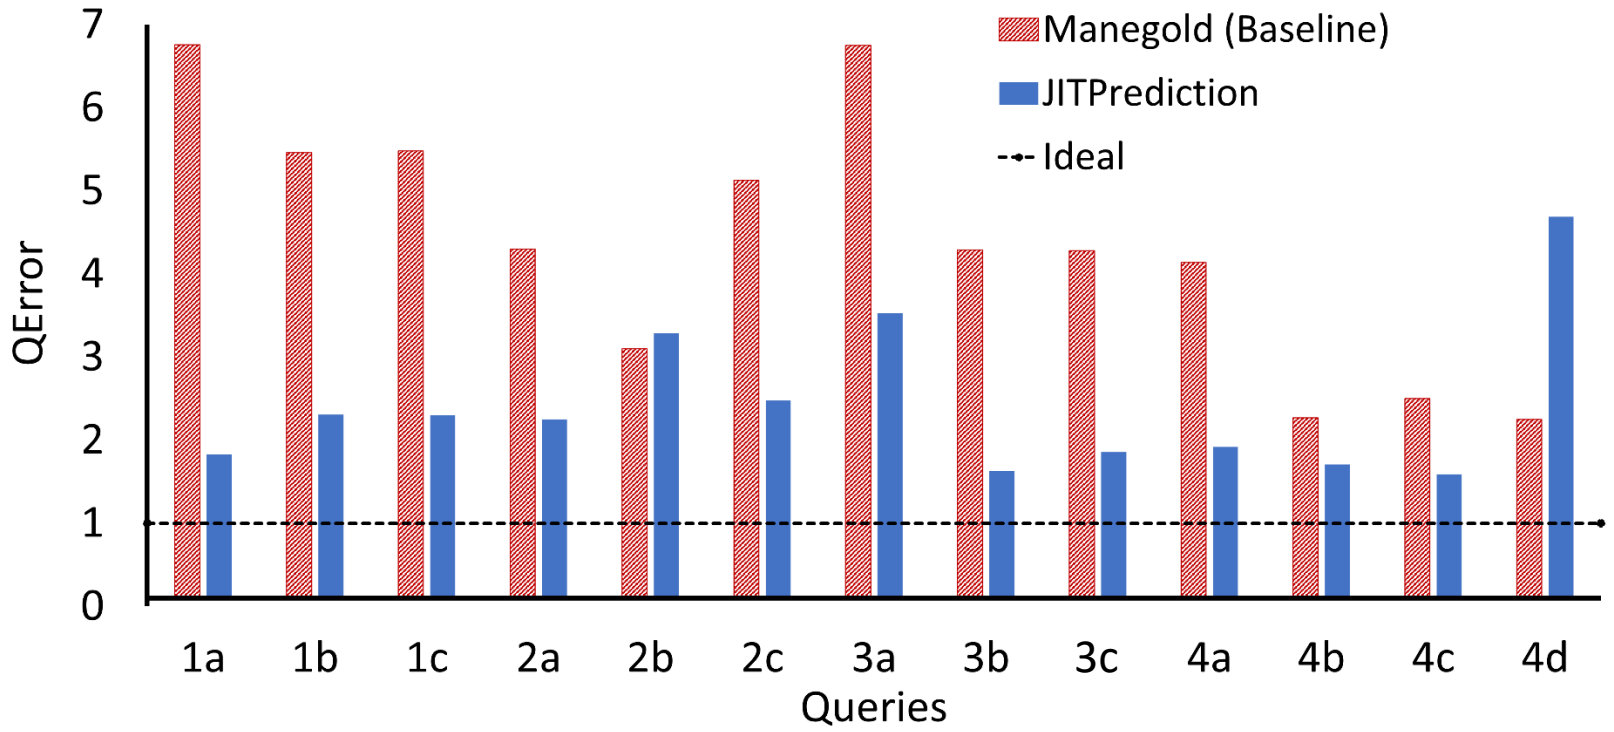

Accuracy: Improves over Analytical Model by 93%
Training Time: < 10% of the batch execution time

Hardware Parameters

JIT Calibration

Workload

Analytical Model

JIT Prediction

Low Training + Near Accurate Prediction

Thank You!