Robust LSM-Trees Tuning for Workload and Resources Uncertainty

Authors: Andy Huynh, Harshal Chaudhari,
Evimaria Terzi, Manos Athanassoulis

IBM Mentor: Josh Berkus

Tuning Data Systems

Workloads Resources Database Performance

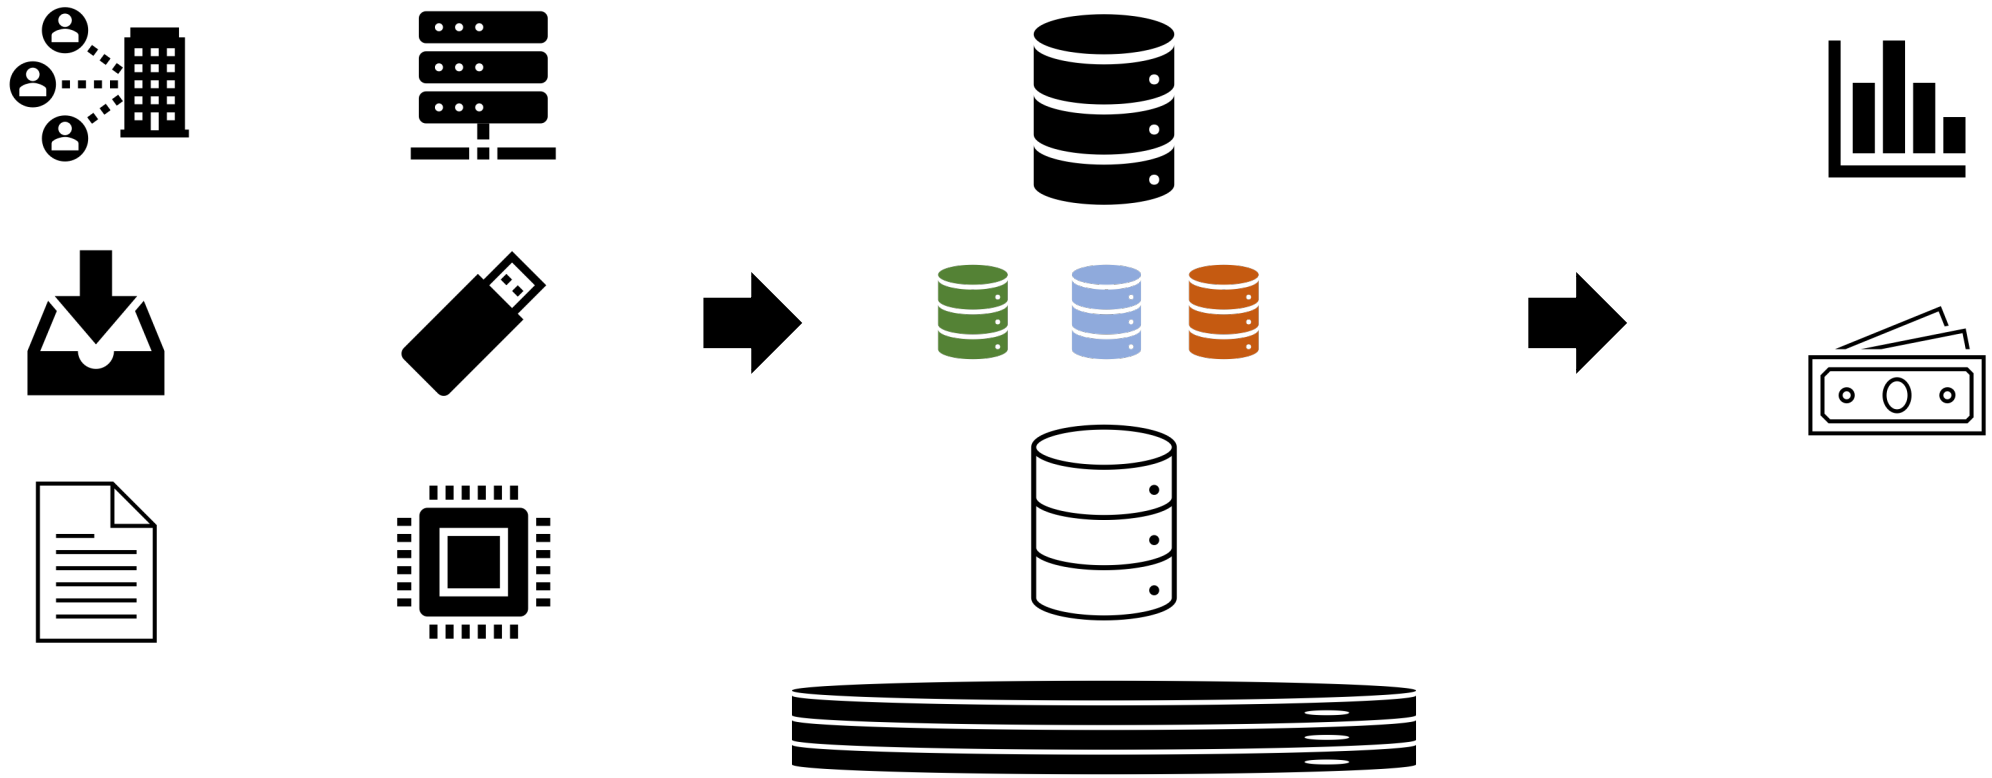

Tuning Data Systems

Workloads

Resources

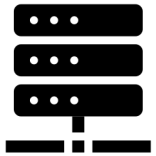

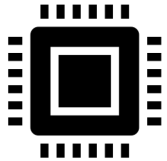

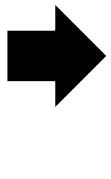

Production

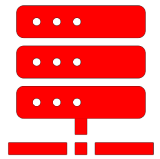

How often do our assumptions hold up?

How often do we have to retune our systems?

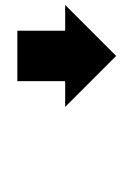

Performance

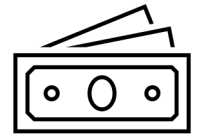

Outline

LSM Trees: A Review

Tuning Problem

Modeling and Deploying LSM Trees

Outline

LSM Trees: A Review

Tuning Problem

Modeling and Deploying LSM Trees

LSM Trees : A Review

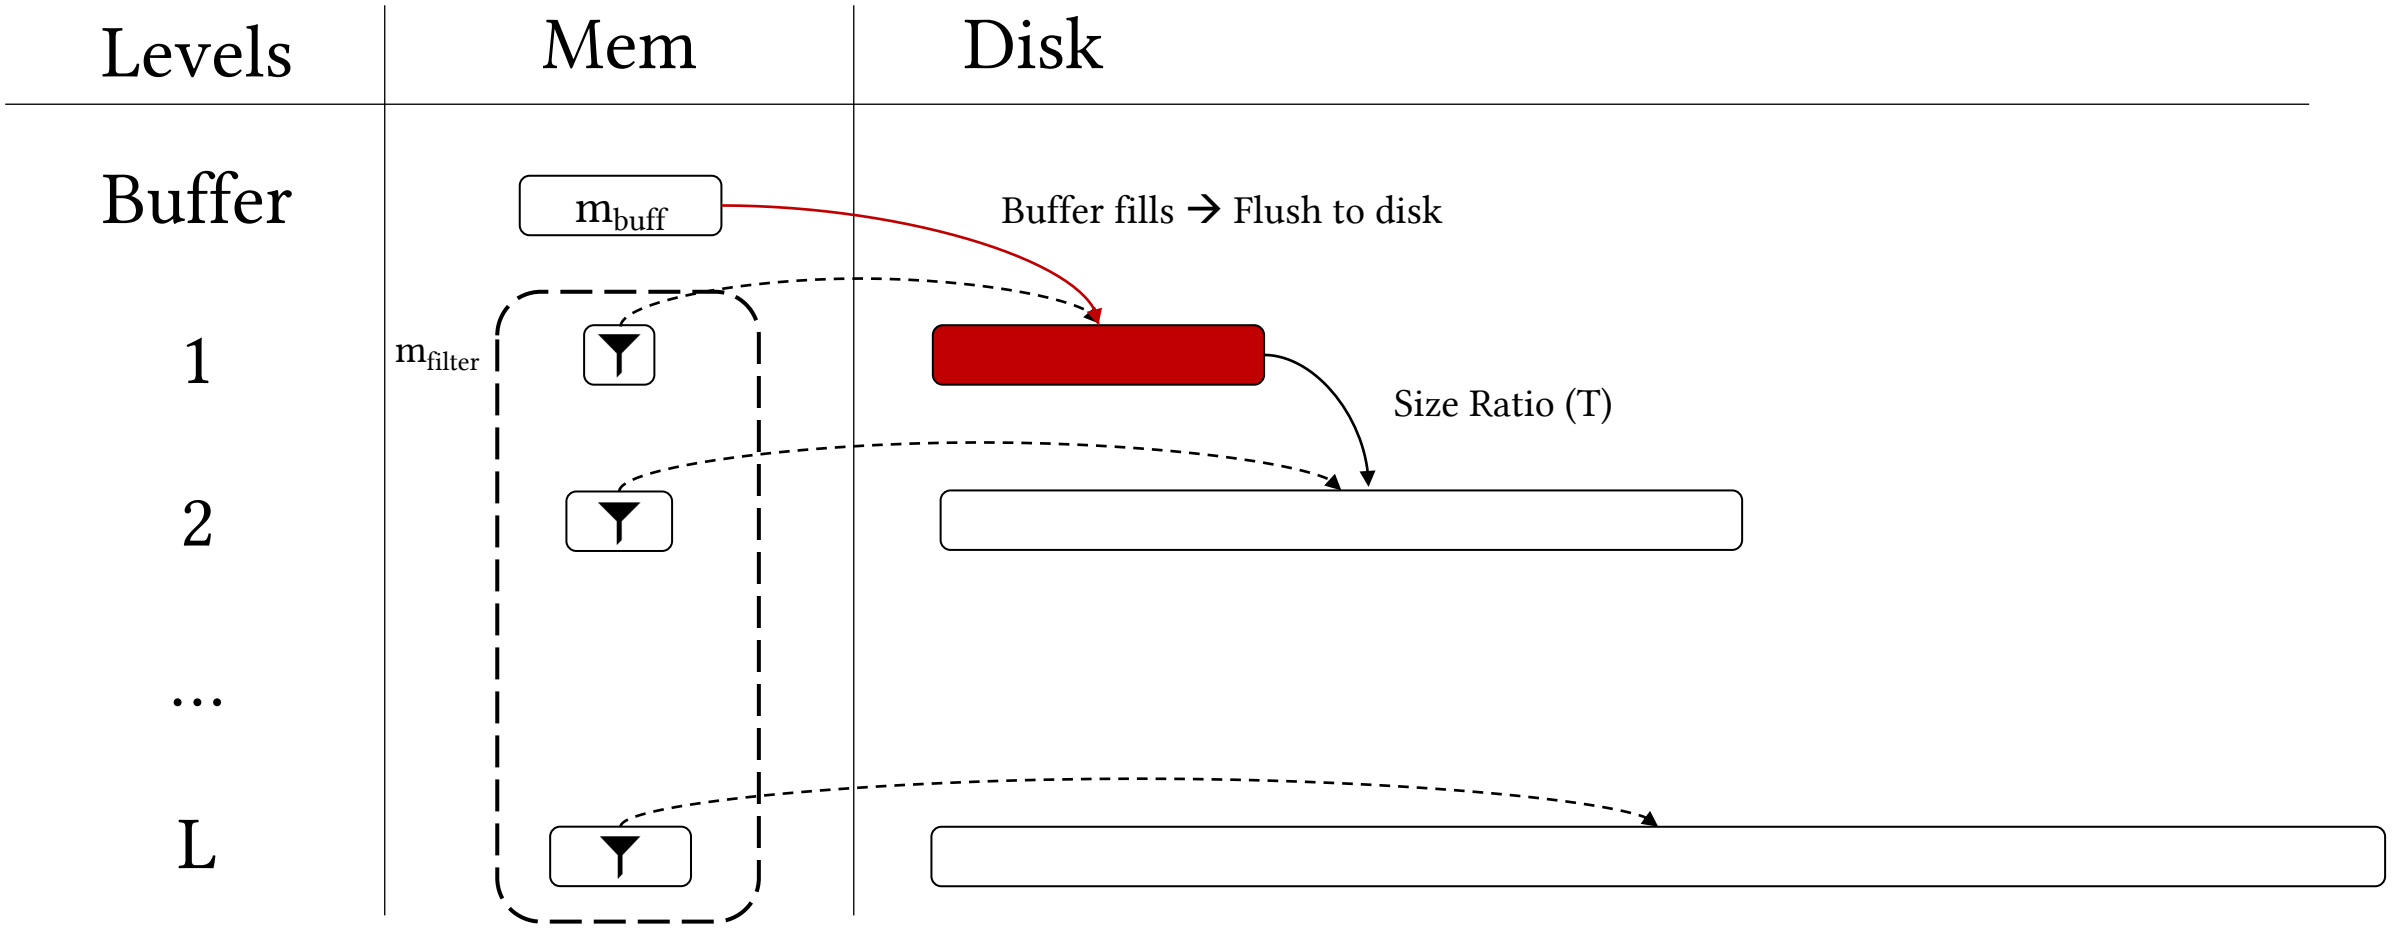

LSM Trees: Compactions

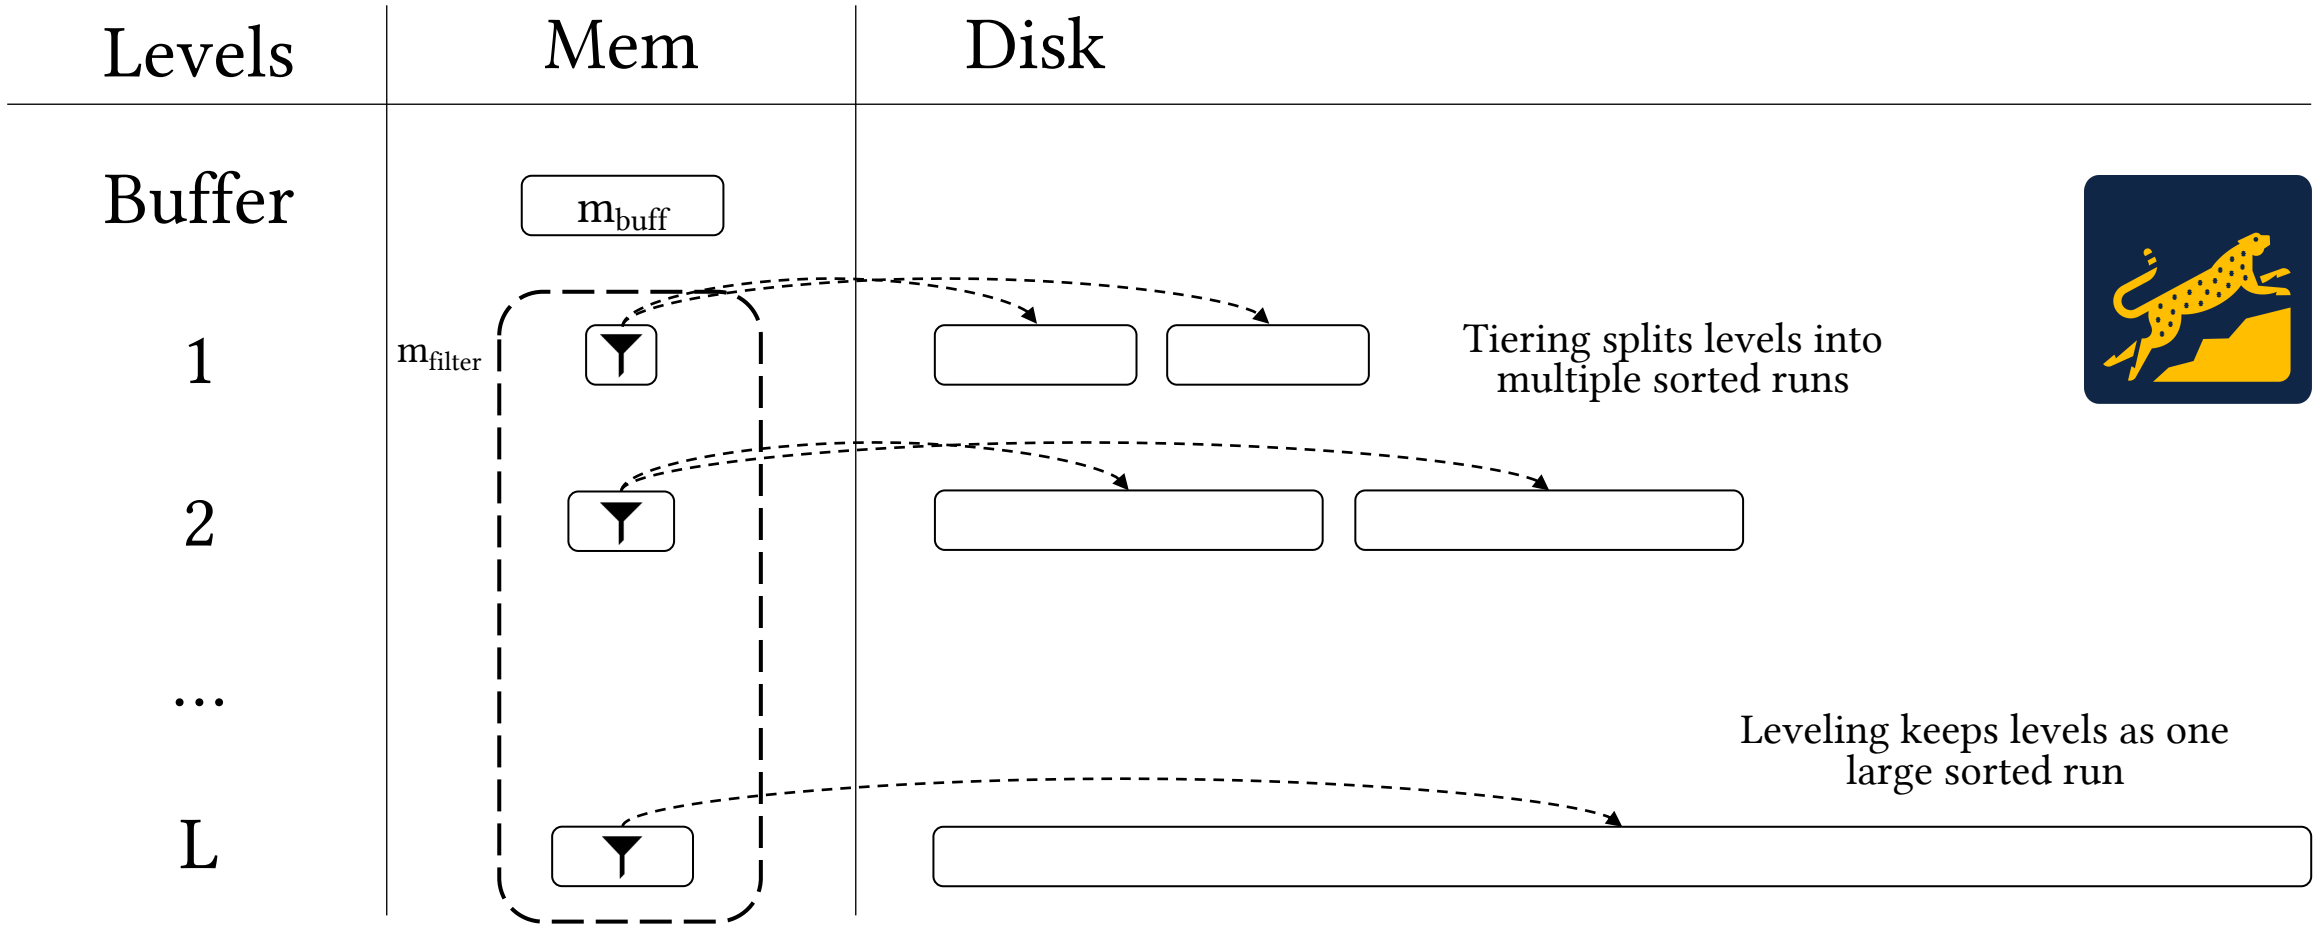

Outline

LSM Trees: A Review

Tuning Problem

Modeling and Deploying LSM Trees

Tuning Problem

$\boldsymbol{w}$: Workload

D : LSM Tree Design

ϕ : Fixed System Parameters

C : Cost (Latency)

$$D^* = \text{argmin}_D \, C(\boldsymbol{w}, D, \phi)$$

Uncertainty Regions

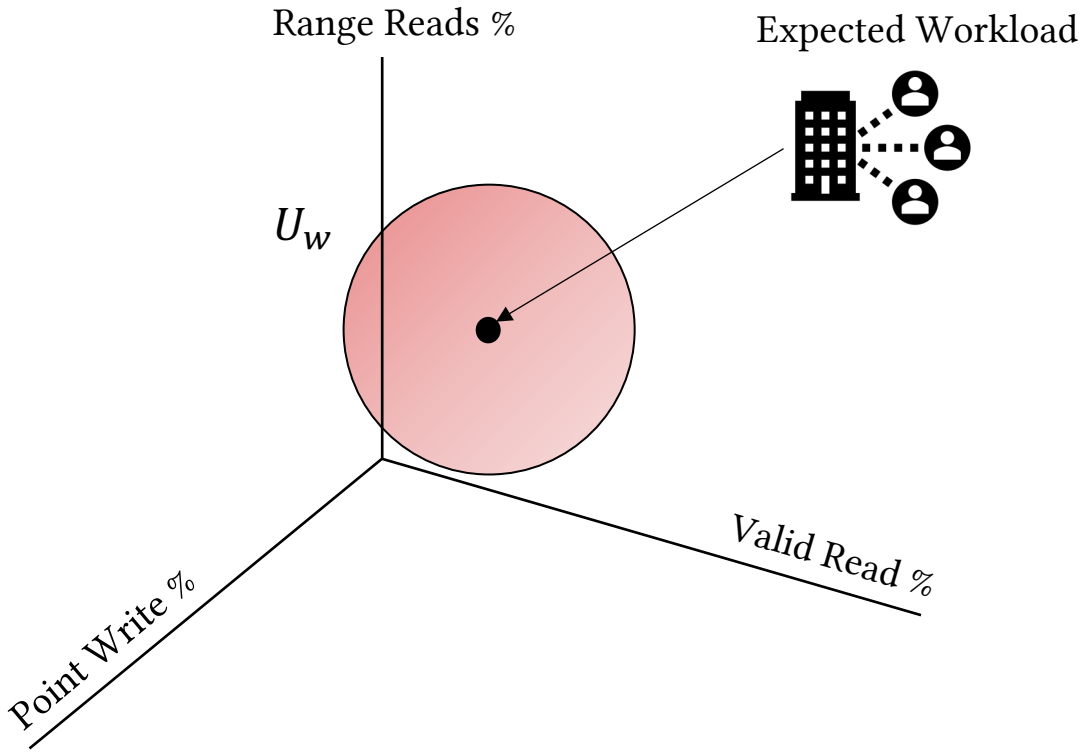

Cloud resource variance

Workload drift

Multi-tenant nodes

Robust Tuning Problem

$\boldsymbol{w}$: Workload

D : LSM Tree Design

ϕ : Fixed System Parameters

C : Cost (Latency)

$$D^* = \text{argmin}_D \; C(\boldsymbol{w}, D, \phi)$$

U_w: Uncertainty Region of Workloads

$$D^* = \text{argmin}_D \; C(\hat{\boldsymbol{w}}, D, \phi)$$

$$s.t., \qquad \hat{\boldsymbol{w}} \in U_w$$

Robust Designs

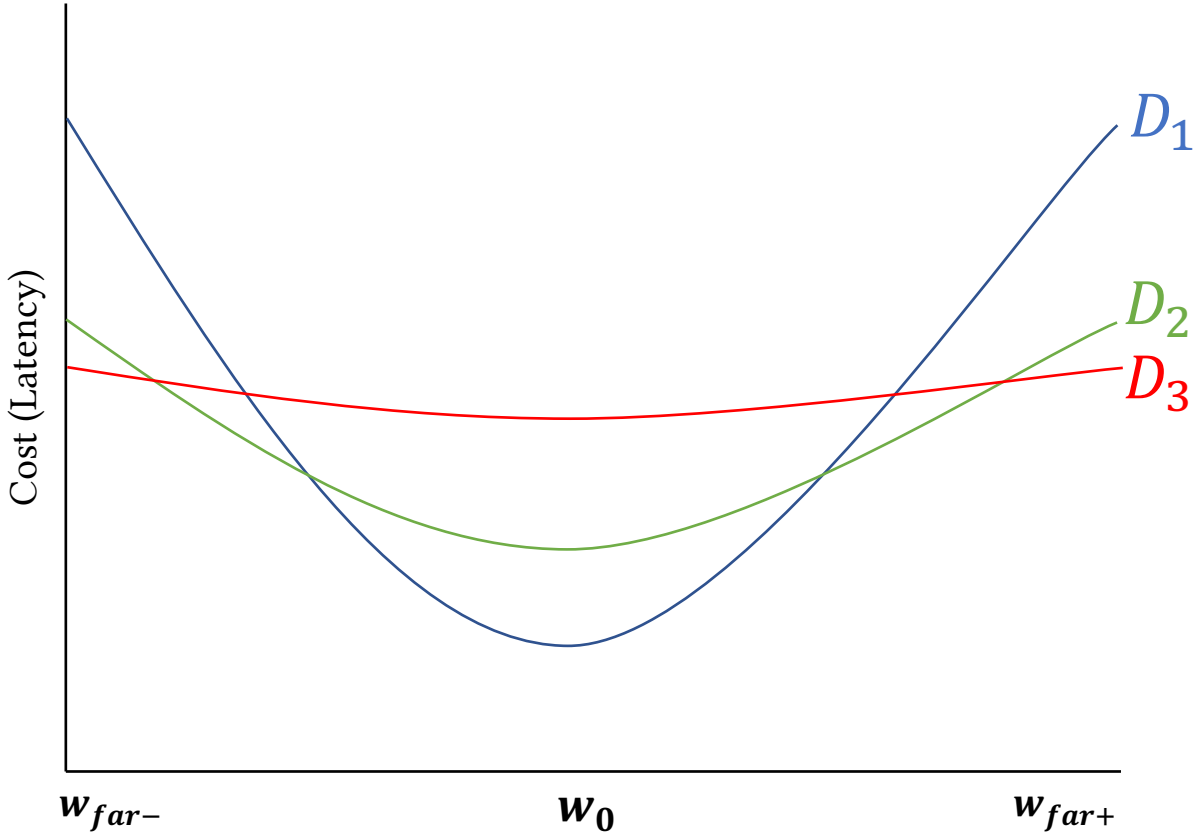

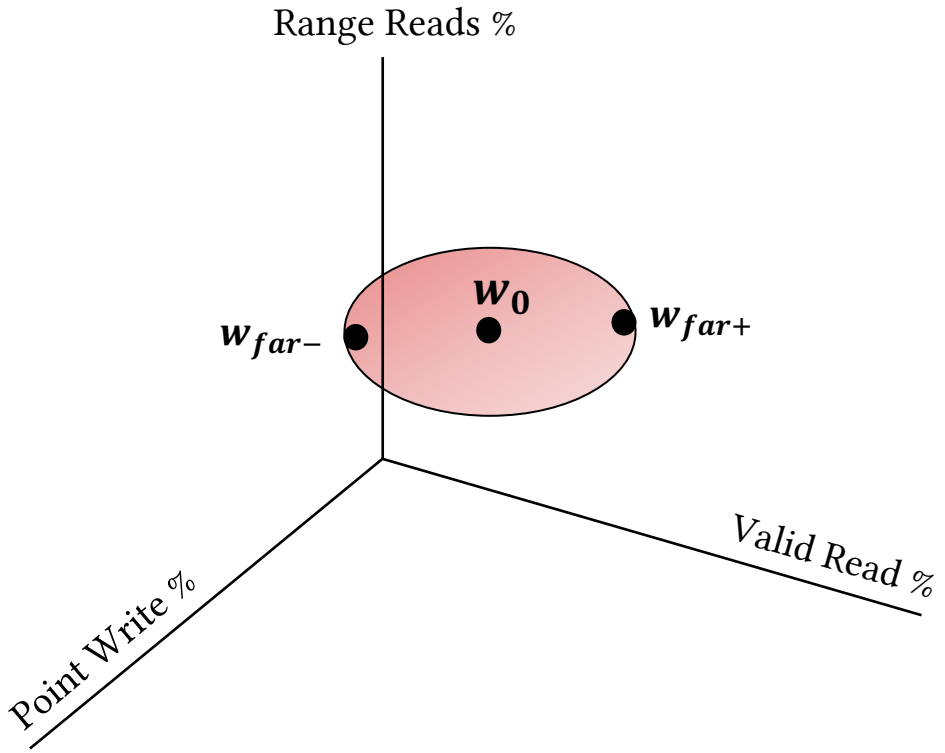

Outline

LSM Trees: A Review

Tuning Problem

Modeling and Deploying LSM Trees

Modeling LSM Trees

Accurate modeling is **extremely** important

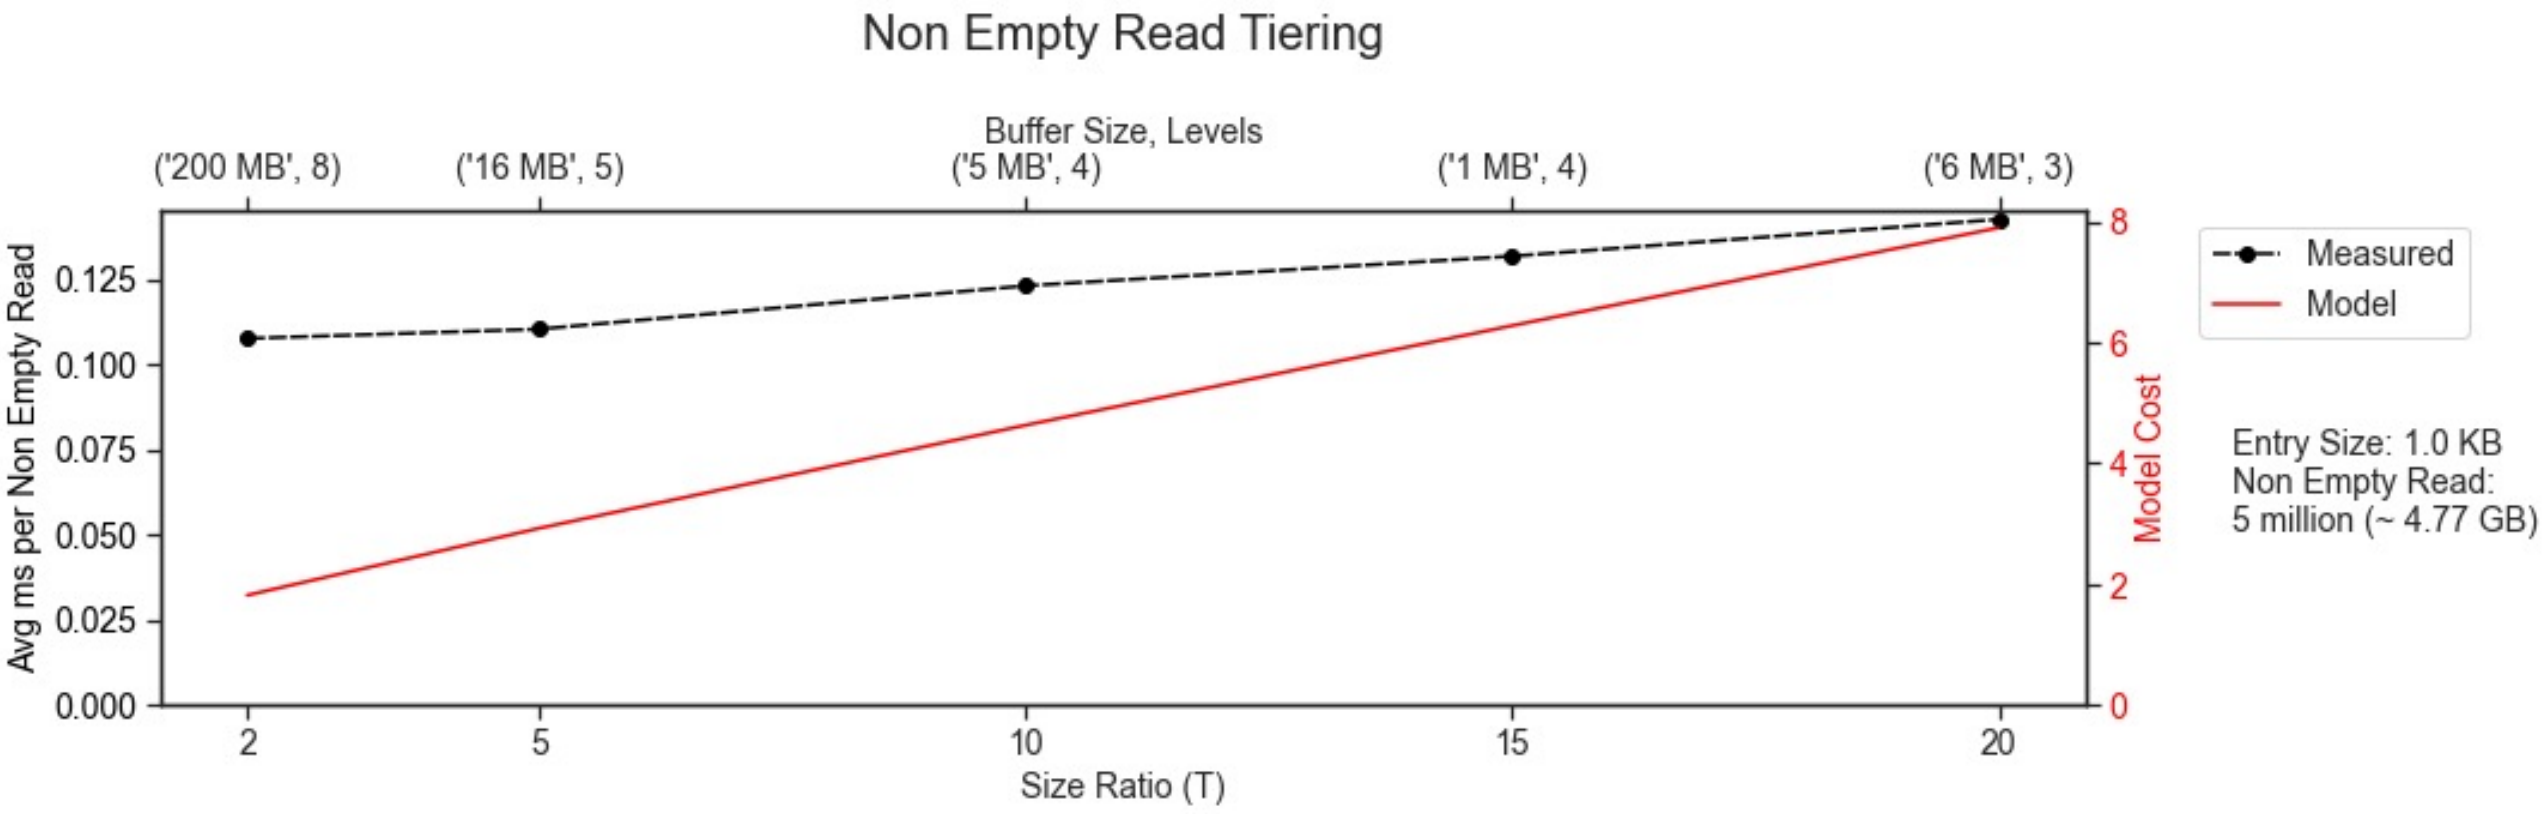

Modeling Reads

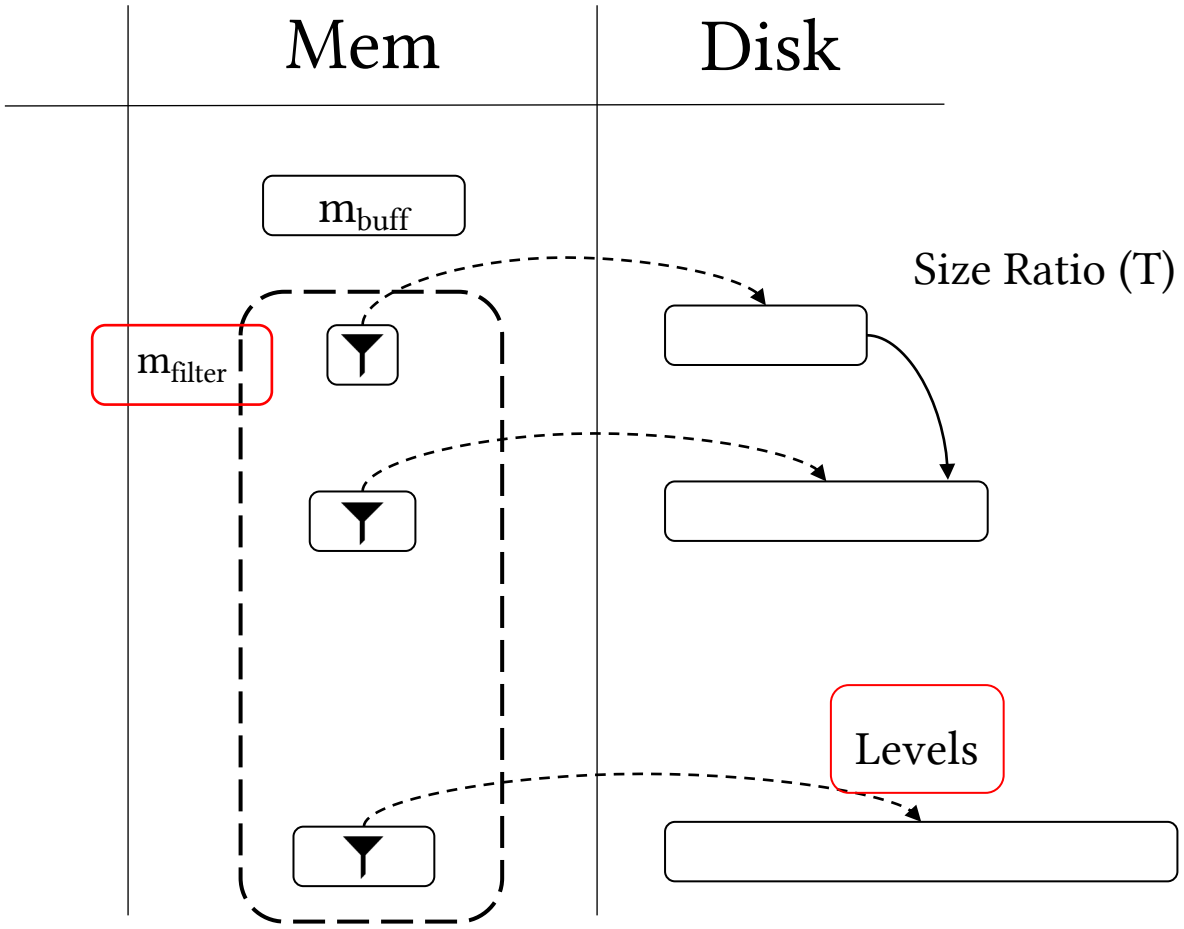

Bloom filters introduce a **false positive rate**

Reads can be split between **empty** and **non-empty** reads

Probabilistic model on expected number of I/Os

Modeling Writes

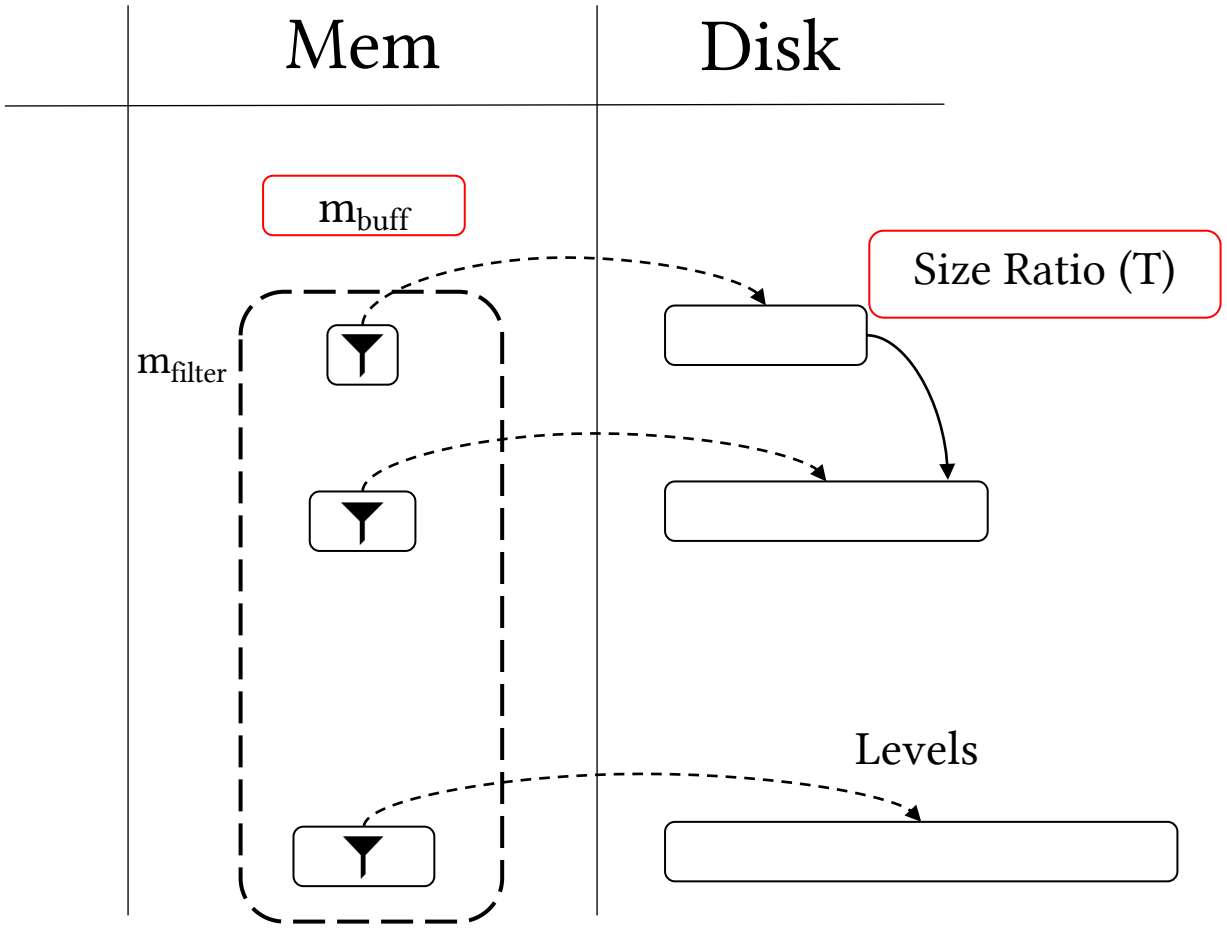

Write speeds depend heavily on **buffer size** relative to entries

Compaction policy plays a heavy role to frequency of writes

Amortized cost of write I/Os over a sustained period

Tuning Pipeline

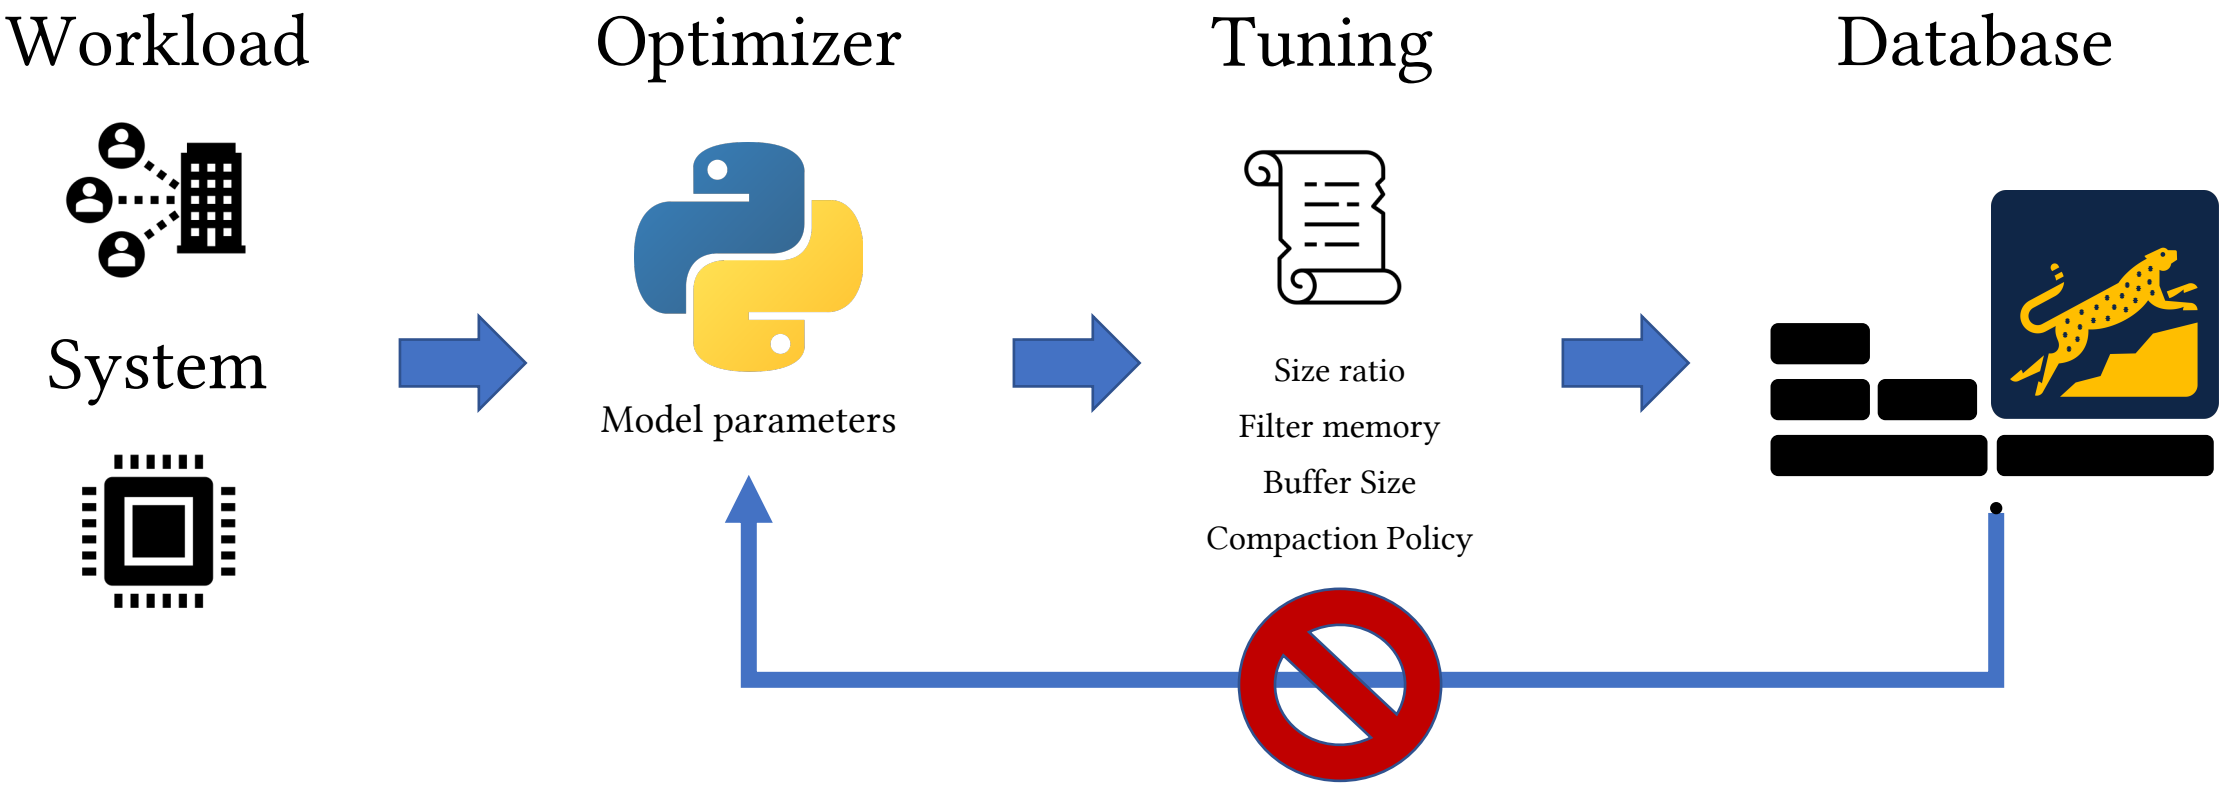

Closing Thoughts

LSM Trees operate in **uncertain** environments requiring **robust** tunings

Any tuning pipelines requires **accurate** models

Can we extend this design paradigm to other data structures?

Working towards promising results!

Questions? Contact me: ndhuynh@bu.edu, ndhuynh.com