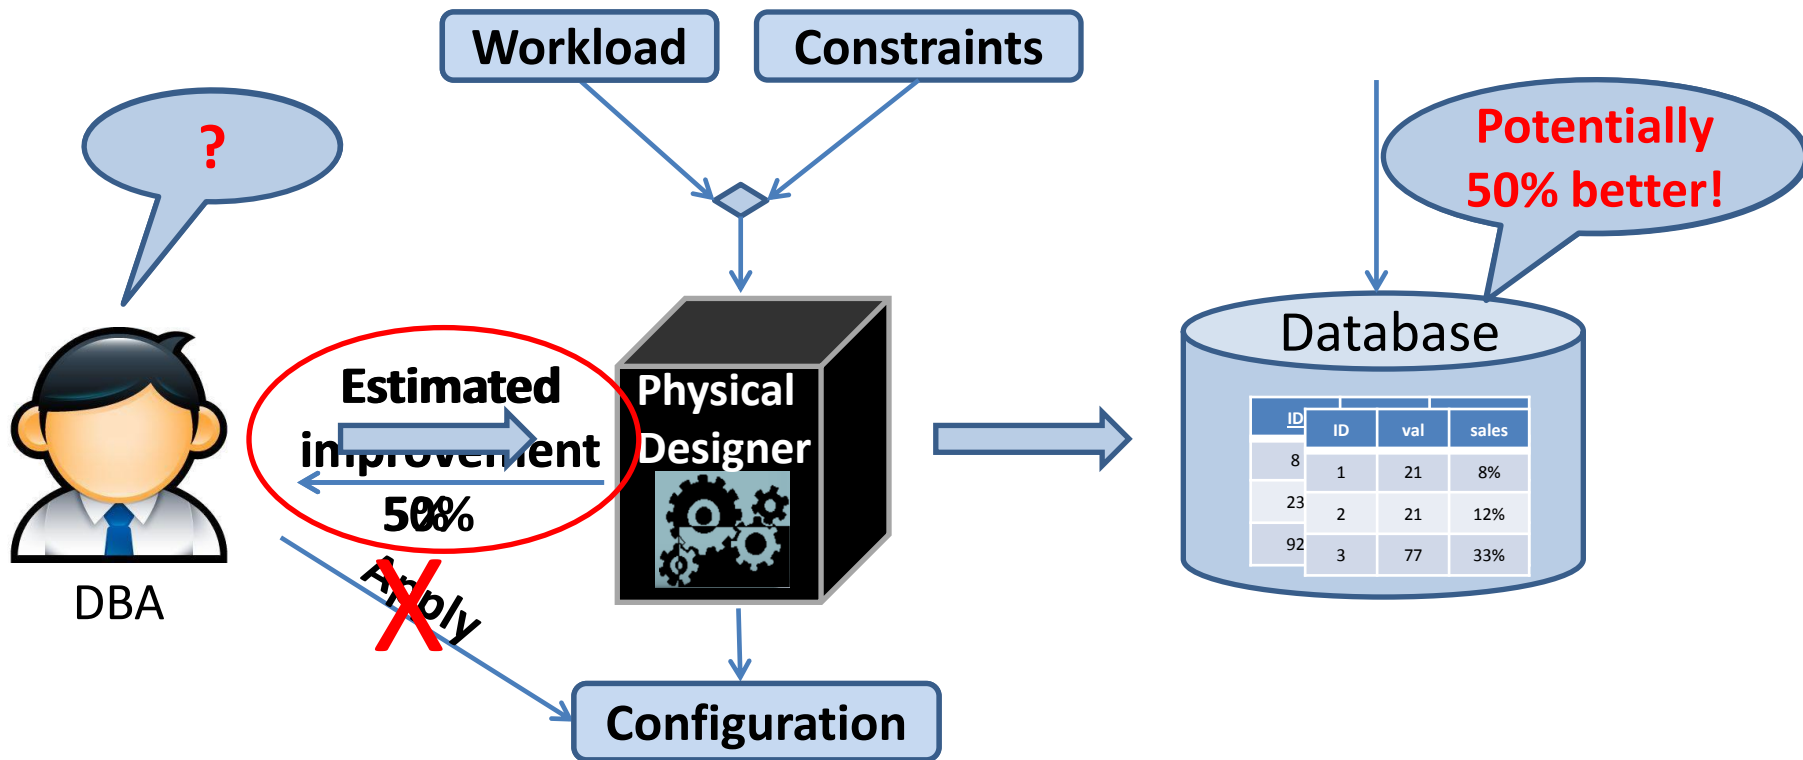


Automated Physical Designers: What You See is What You Get?

Renata Borovica, Ioannis Alagiannis,
Anastasia Ailamaki

Physical designers' estimates

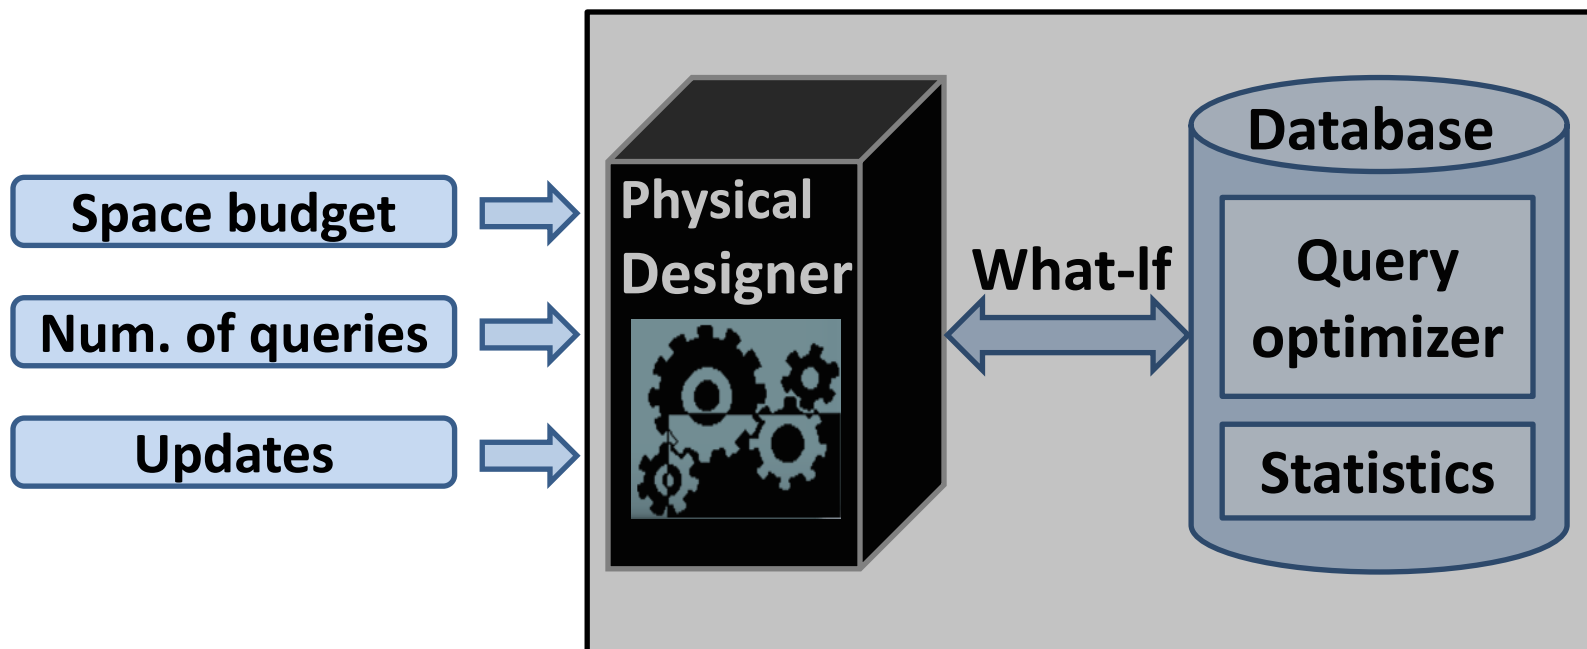


Overestimates = User frustration

Underestimates = Missing opportunities

Our approach

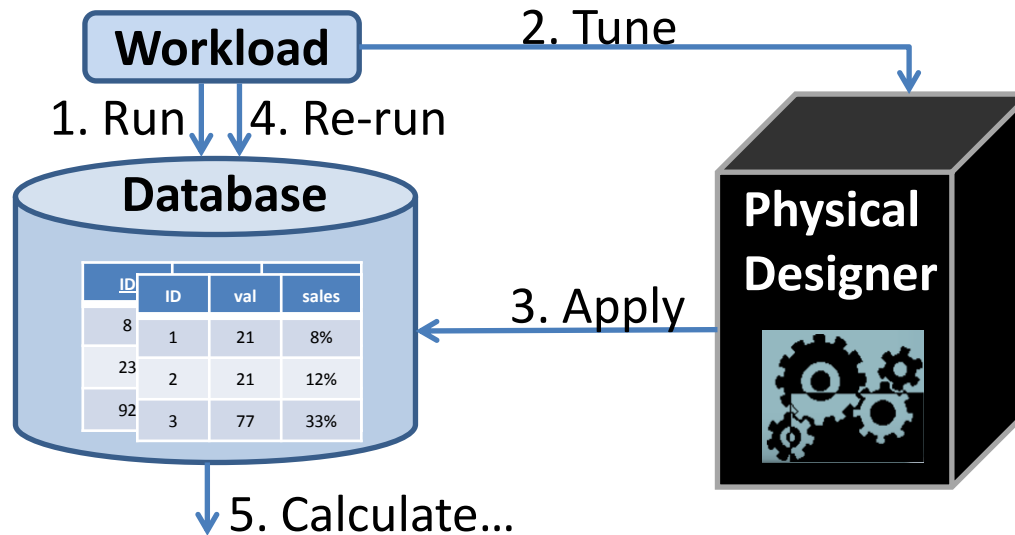
Compare existing physical database designers in terms of predictability (actual vs. estimated improvement).



Experimental setup

- Hardware
 - 2 x 4-core AMD 2.7GHz, 32 GB, Win 2008 R2
 - I/O: 2 x 750 GB SATA 7200rpm, RAID 0, 90 MB/s
- Commercial DBMS
 - System A, System B, System C
 - Buffer pool 20% of DB size, cold runs, updated statistics
- Workloads
 - TPC-H: SF (10 and 100), 17 queries
 - NREF: 6.7GB, 200 queries

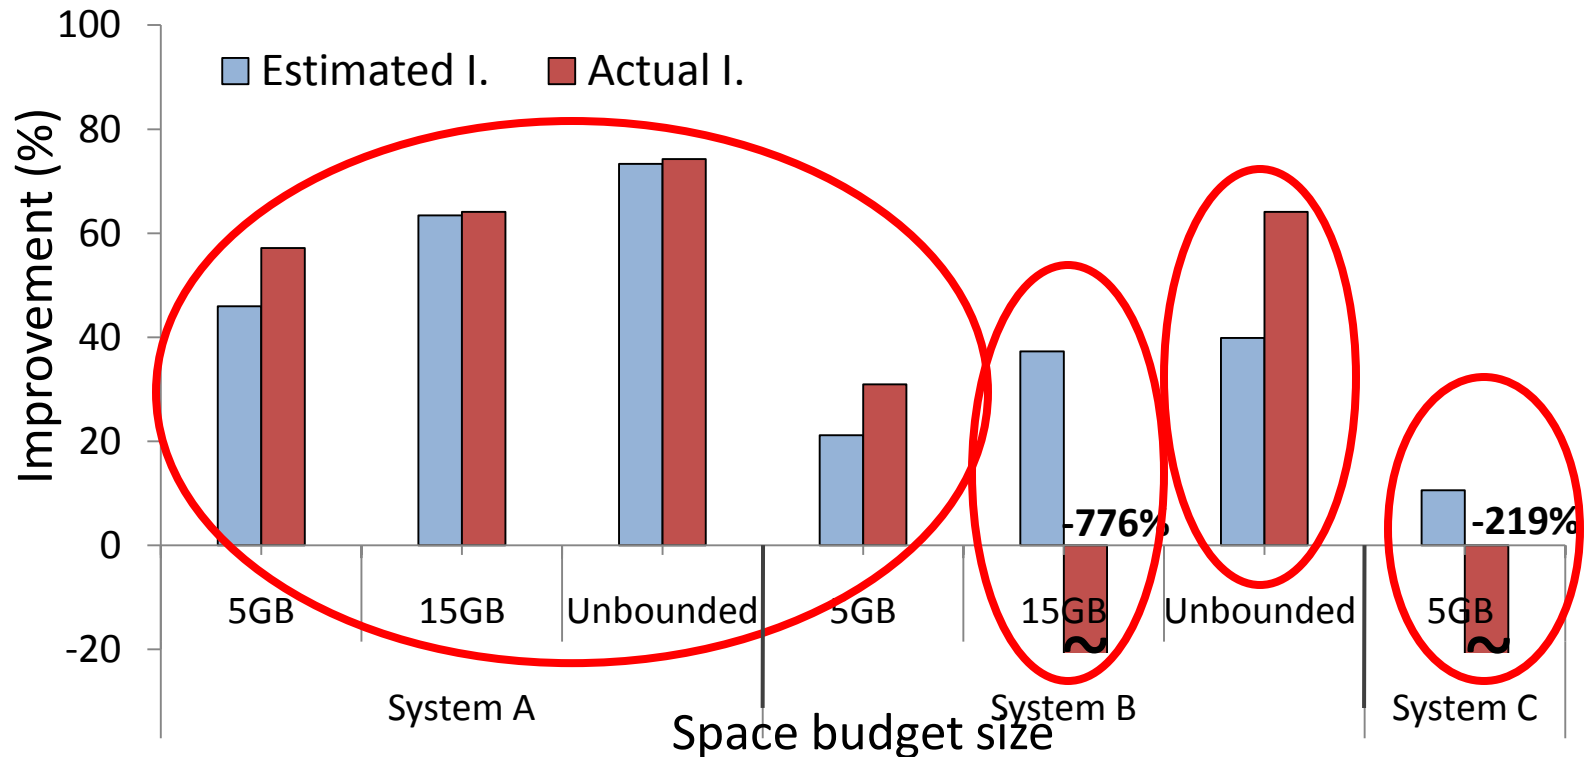
Experimental methodology



Metric	Label	Formula
Actual improvement (%)	I_A	$\left(1 - \frac{\text{Tuned_time}}{\text{Original_time}}\right) \times 100$
Estimated tuned time (sec)	E_{TT}	$\text{Original_time} - \frac{\text{Estimated_improvement} \times \text{Original_time}}{100}$
Relative estimation error (%)	R_{EE}	$\frac{E_{TT} - \text{Tuned_time}}{\text{Tuned_time}} \times 100$

Impact of space budget

Setting: TPC-H, SF10, Unlimited time




Improvement usually higher than estimated

Performance hurt in Systems B and C

Analyzing performance degradation

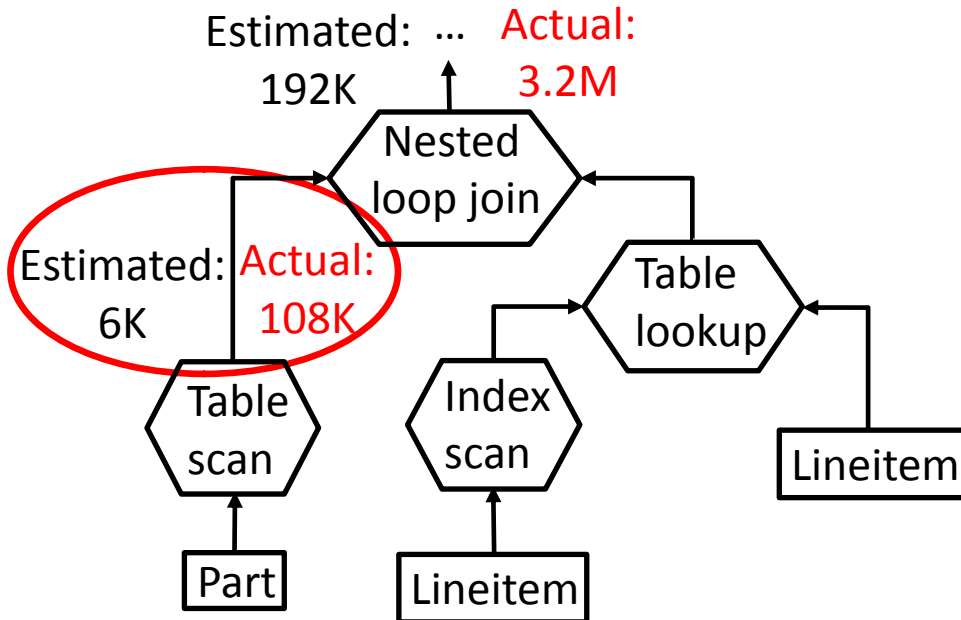
Setting: TPC-H, SF10, System B, Space budget 15 GB



Two queries prolonged workload execution 8 times


Cause for sub-optimal plans

- Cardinality errors



- Order of magnitude more tuples
- 75x longer execution time!

- Cost model

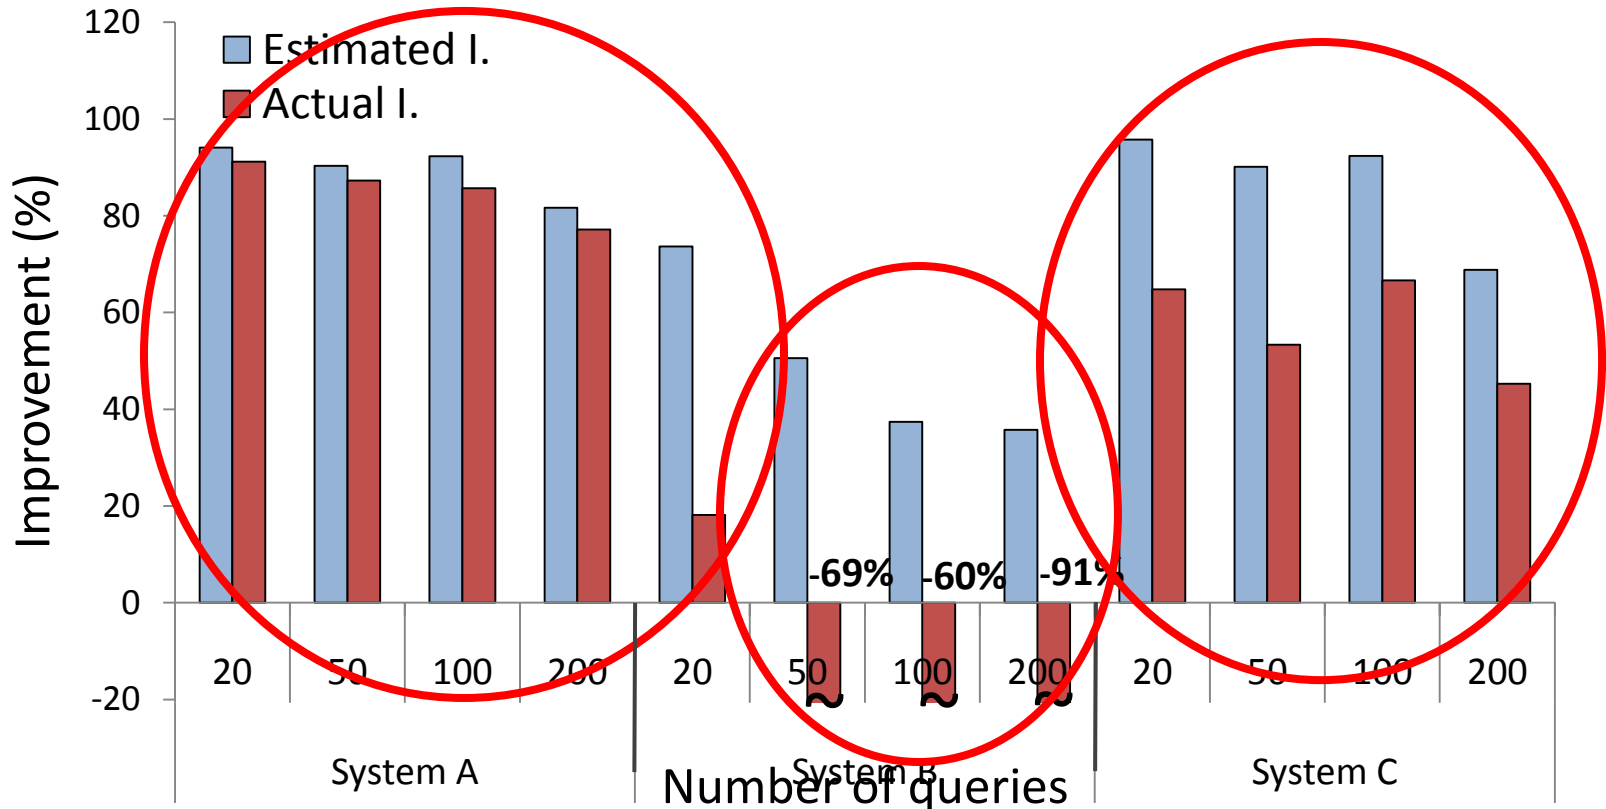


- Wrong decision of cost model
- 5x longer execution time!

Optimizer's mistakes -> mislead designer -> hurt predictability

Increasing number of queries

Setting: NREF, Space budget 20GB, Time budget 30min



Improvement lower than estimated

Wrong cardinalities hurt performance of System B!

Impact of updates

Setting: NREF, Space budget 20GB, Time budget 30min,
400 statements

Metric	System A	System B	System C
Estimated I. (%)	58.62	--	2.23
Actual I. (%)	-18.3	--	-8.13
Relative error (%)	65.02	--	9.58

Complex trade-off between improvement and maintenance

Summary

- System A
 - Relative error 2 – 46 % in read-only workloads
 - Performance hurt by 18% only in update-intensive workload
- System B
 - Relative error 14 – 92 % in read-only workloads
 - Performance hurt up to 8x after tuning
- System C
 - Relative error 42 – 87 % in read-only workloads
 - Performance hurt up to 2x after tuning

What you see is NOT what you get

Thank you!

<http://dias.epfl.ch>