

Tiresias:

Enabling Predictive Autonomous Storage and Indexing

Michael Abebe mtabebe@uwaterloo.ca

Horatiu Lazu

September 2022

Khuzaima Daudjee

tiny.cc/tiresias



UNIVERSITY OF
WATERLOO

DBMS storage & indexing choices
have trade-offs based on workload

Index Choice

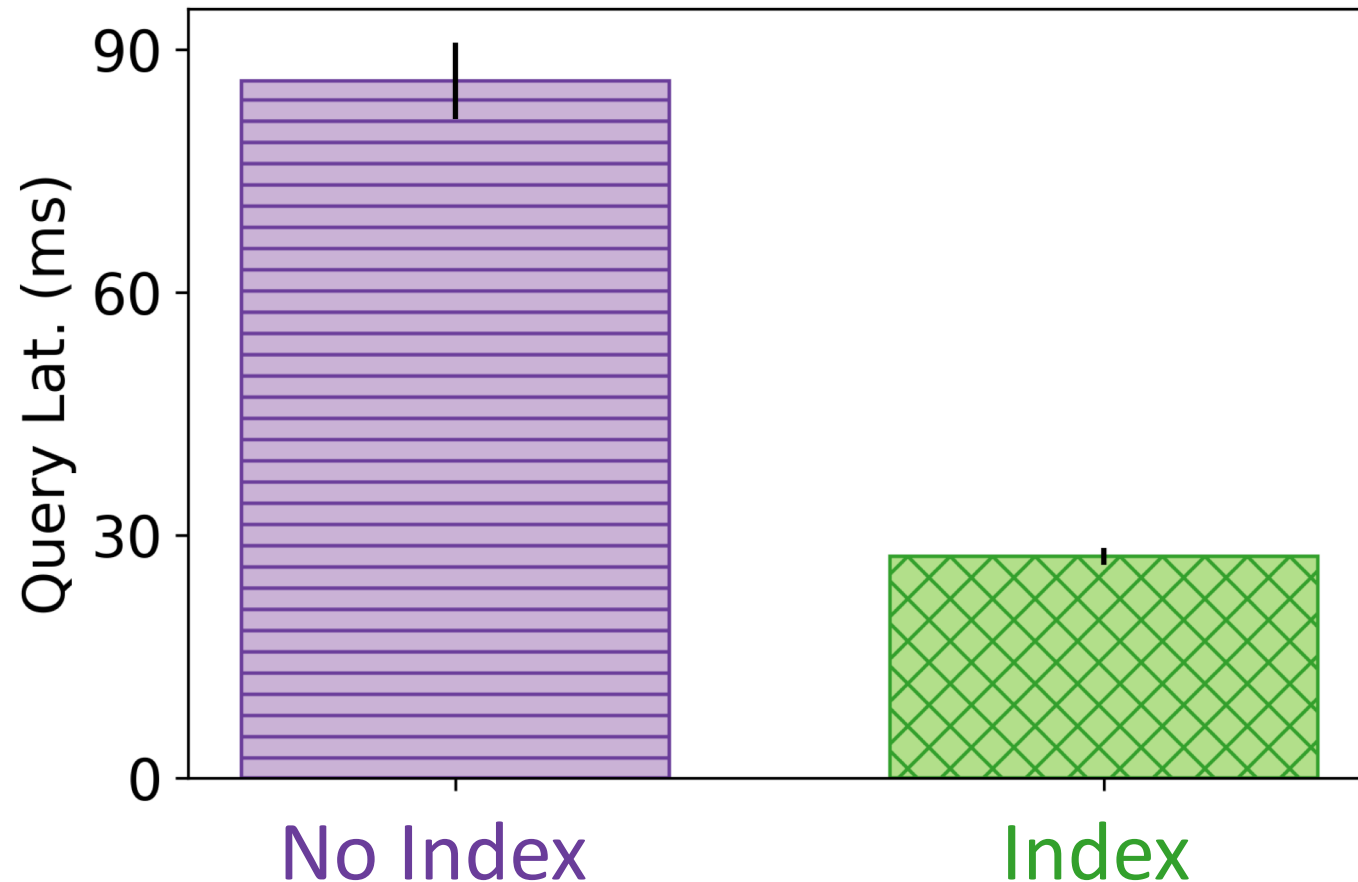
Order	Book	User	Qnt	\$
100	HP	Alice	7	\$15
101	Drac	Ted	90	\$3
102	SQL	Geoff	1	\$100
103	HP	Jean	3	\$15

SELECT * WHERE
QNT > 15

Should we index?

Index Choice

Lower is better



SELECT * WHERE
QNT > 15

Should we index?

Yes

Index Choice

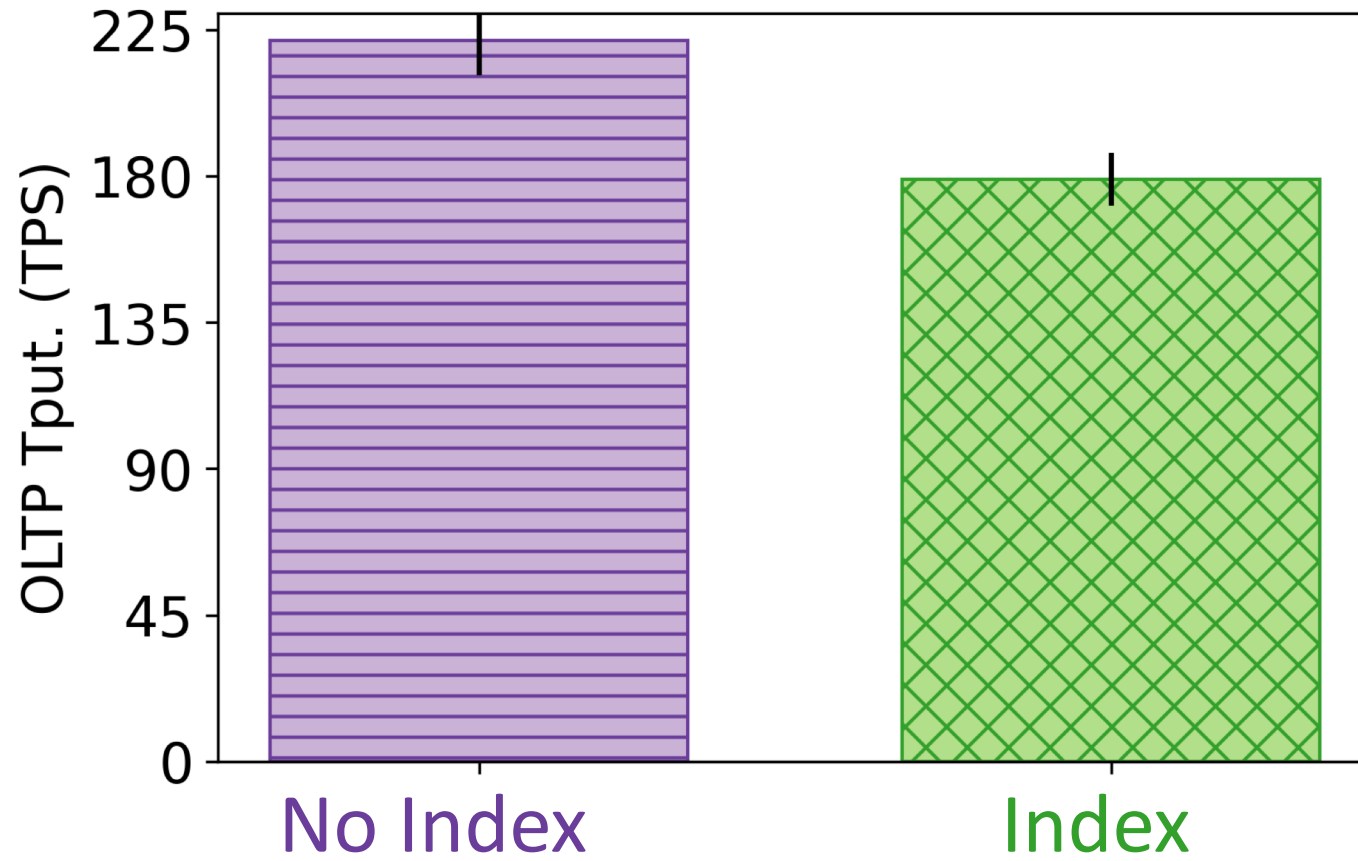
Order	Book	User	Qnt	\$
100	HP	Alice	7	\$15
101	Drac	Ted	90	\$3
102	SQL	Geoff	1	\$100
103	HP	Jean	3	\$15
104	LoTR	John	2	\$21

INSERT (104, LoTR,
John, 2, \$21)

Should we index?

Index Choice

Higher is better



INSERT (104, LoTR,
John, 2, \$21)

Should we index?

No


Storage Choice

Order	Book	User	Qnt	\$
100	HP	Alice	7	\$15
101	Drac	Ted	90	\$3
102	SQL	Geoff	1	\$100
103	HP	Jean	3	\$15
104	LoTR	John	2	\$21

INSERT (104, LoTR,
John, 2, \$21)

Storage Choice

Higher is better



INSERT (104, LoTR,
John, 2, \$21)

How should we store?

Row


Storage Choice

Order	Book	User	Qnt	\$
100	HP	Alice	7	\$15
101	Drac	Ted	90	\$3
102	SQL	Geoff	1	\$100
103	HP	Jean	3	\$15
104	LoTR	John	2	\$21

**SELECT * WHERE
QNT > 15**

Storage Choice

Lower is better




SELECT * WHERE
QNT > 15

How should we store?
Column

DBMS storage & indexing choices
have trade-offs based on workload
so choices should be adaptive!


Workloads Change



Tiresias


Predict upcoming accesses and latency under different storage and indexing choices

Tiresias



Model workload *access* history
Model *access latencies*

DBMS

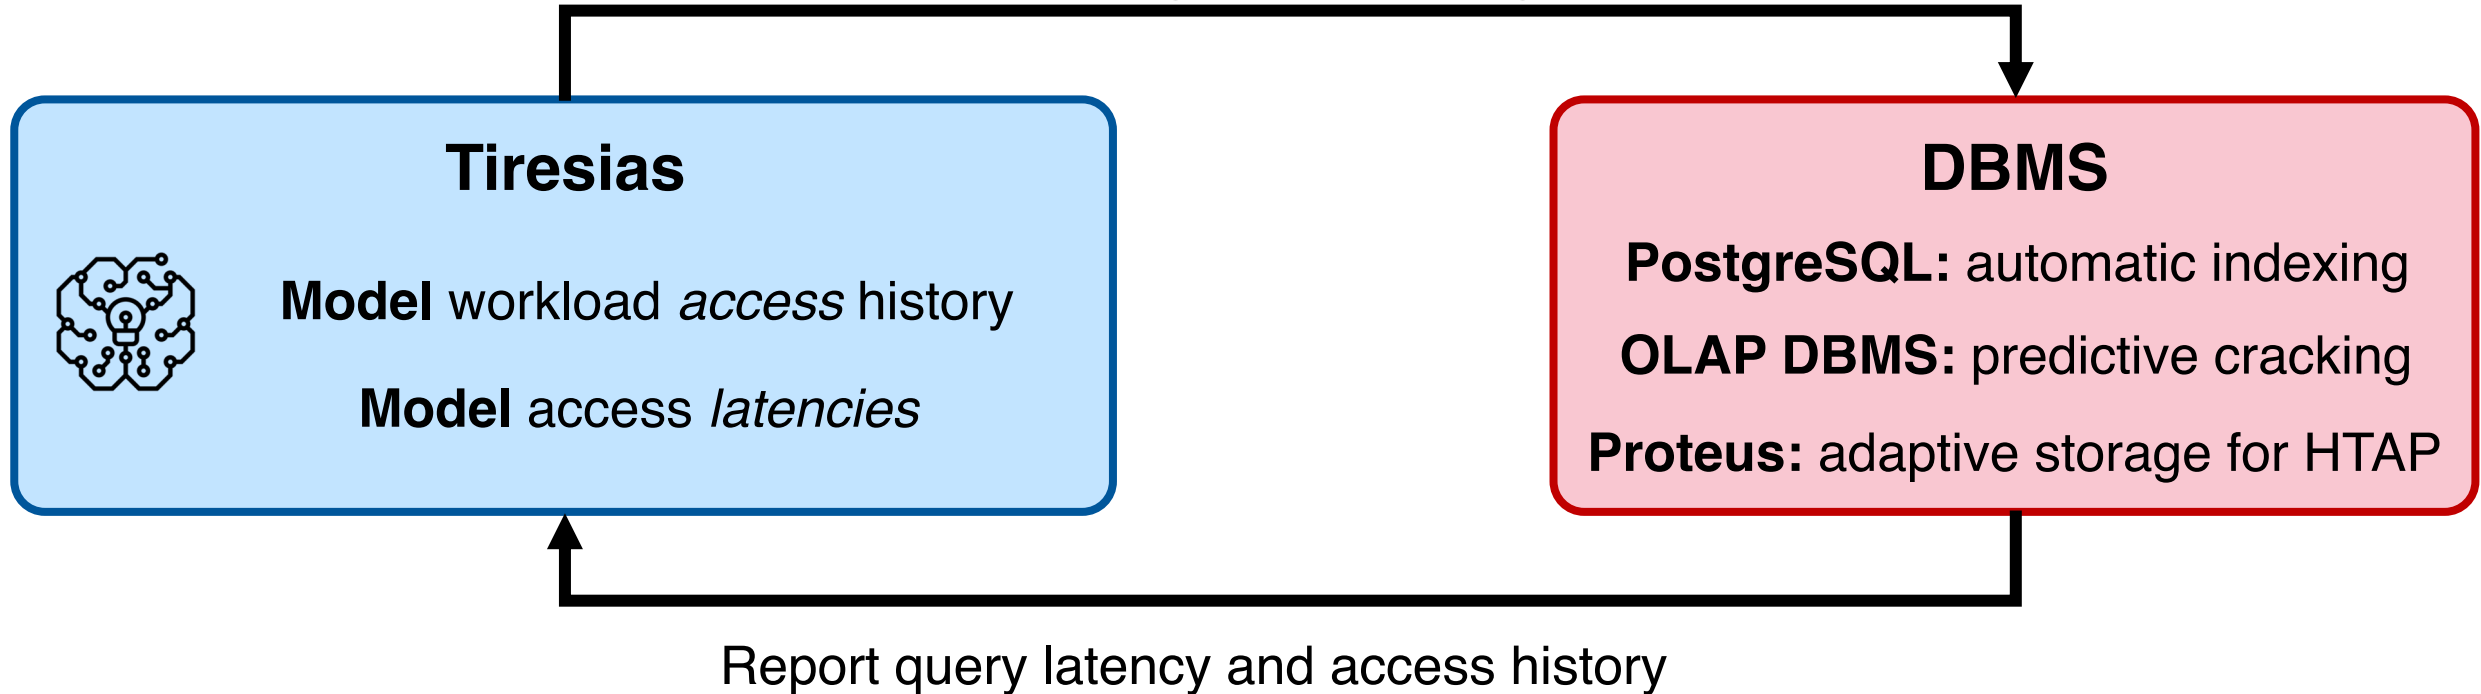


Store and index data
Change data storage and indexing

Report query latency and access history


Tiresias

Predict upcoming accesses and latency under different storage and indexing choices





Tiresias

Predict upcoming accesses and latency under different storage and indexing choices




Report query latency and access history


Access Arrival Patterns




Access Arrival Patterns




Learning Access Arrival Patterns




Learning Access Arrival Patterns



=



+



Periodicity

Avg. over user
defined intervals

Trend

Avg. over
interval

SPAR


**Hybrid-
Ensemble**

Recurrent
Neural Network

Linear
Regression


Access Arrival Patterns - Results

Hybrid-Ensemble

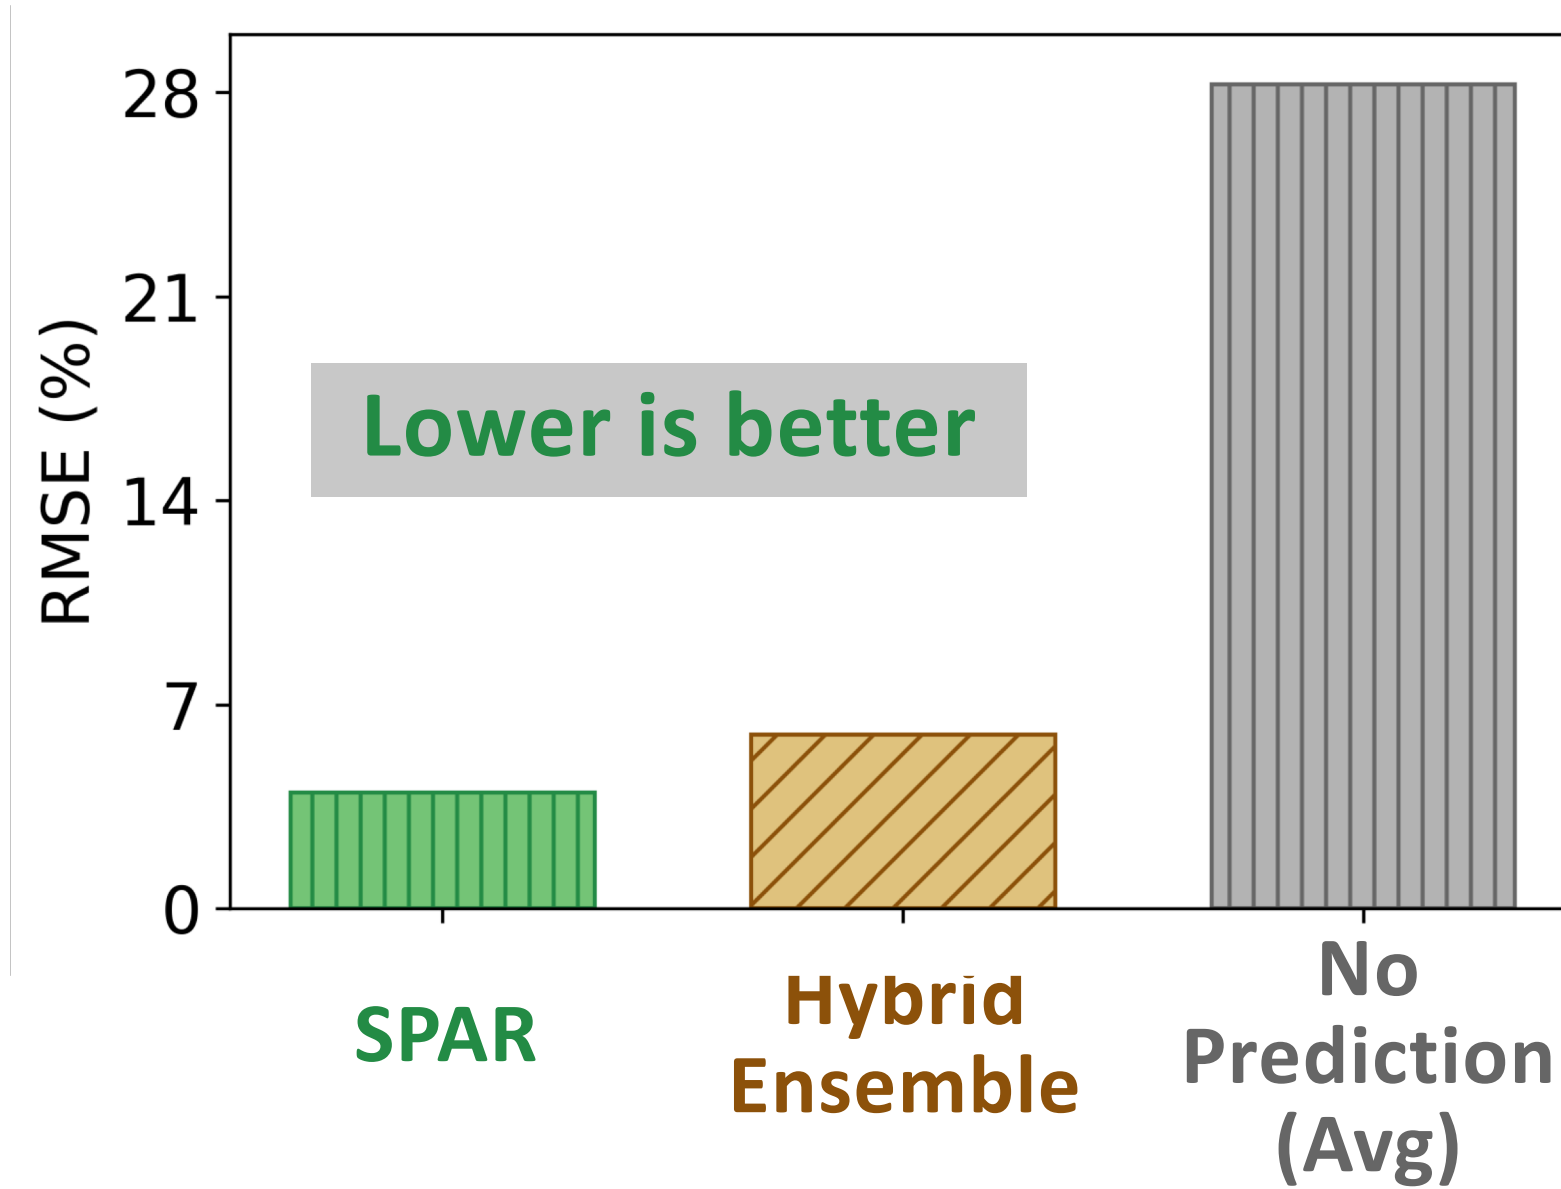


Access Arrival Patterns - Results

SPAR




Access Arrival Patterns - Accuracy



Access Arrival Patterns - Results

SPAR (Mis-Configured)




Incorrect definition of period **decreases** accuracy!

Tiresias

Predict upcoming accesses and latency under different storage and indexing choices

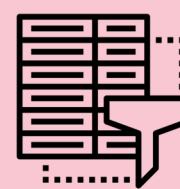
Tiresias



Model workload *access* history
Model *access latencies*

1
2

DBMS



Store and index data
Change data storage and indexing

3

Report query latency and access history

Predicting Latency

Transactions are composed of **physical operators**

SELECT book, **SUM**(qnt) **GROUP BY** book

Row layout

Logical Plan

Sorted column layout

Row scan

Scan & Project
book, qnt

Sequential col
scan

Hash aggregation
book, sum(qnt)

Aggregate
book, sum(qnt)

Sorted col aggregation
book, sum(qnt)

Predicting Latency

Transactions are composed of **physical operators**


Predict physical operator latency

Per layout with workload stats as parameters

Cardinality

Data Width

Est Selectivity



Predicted Latency

Seq col scan

Predicting Latency

Transactions are composed of **physical operators**


Predict physical operator latency

Per layout with workload stats as parameters

Cardinality

Data Width

Est Selectivity



Predicted Latency

Row scan

Predicting Latency

Transactions are composed of **physical operators**

Predict physical operator latency

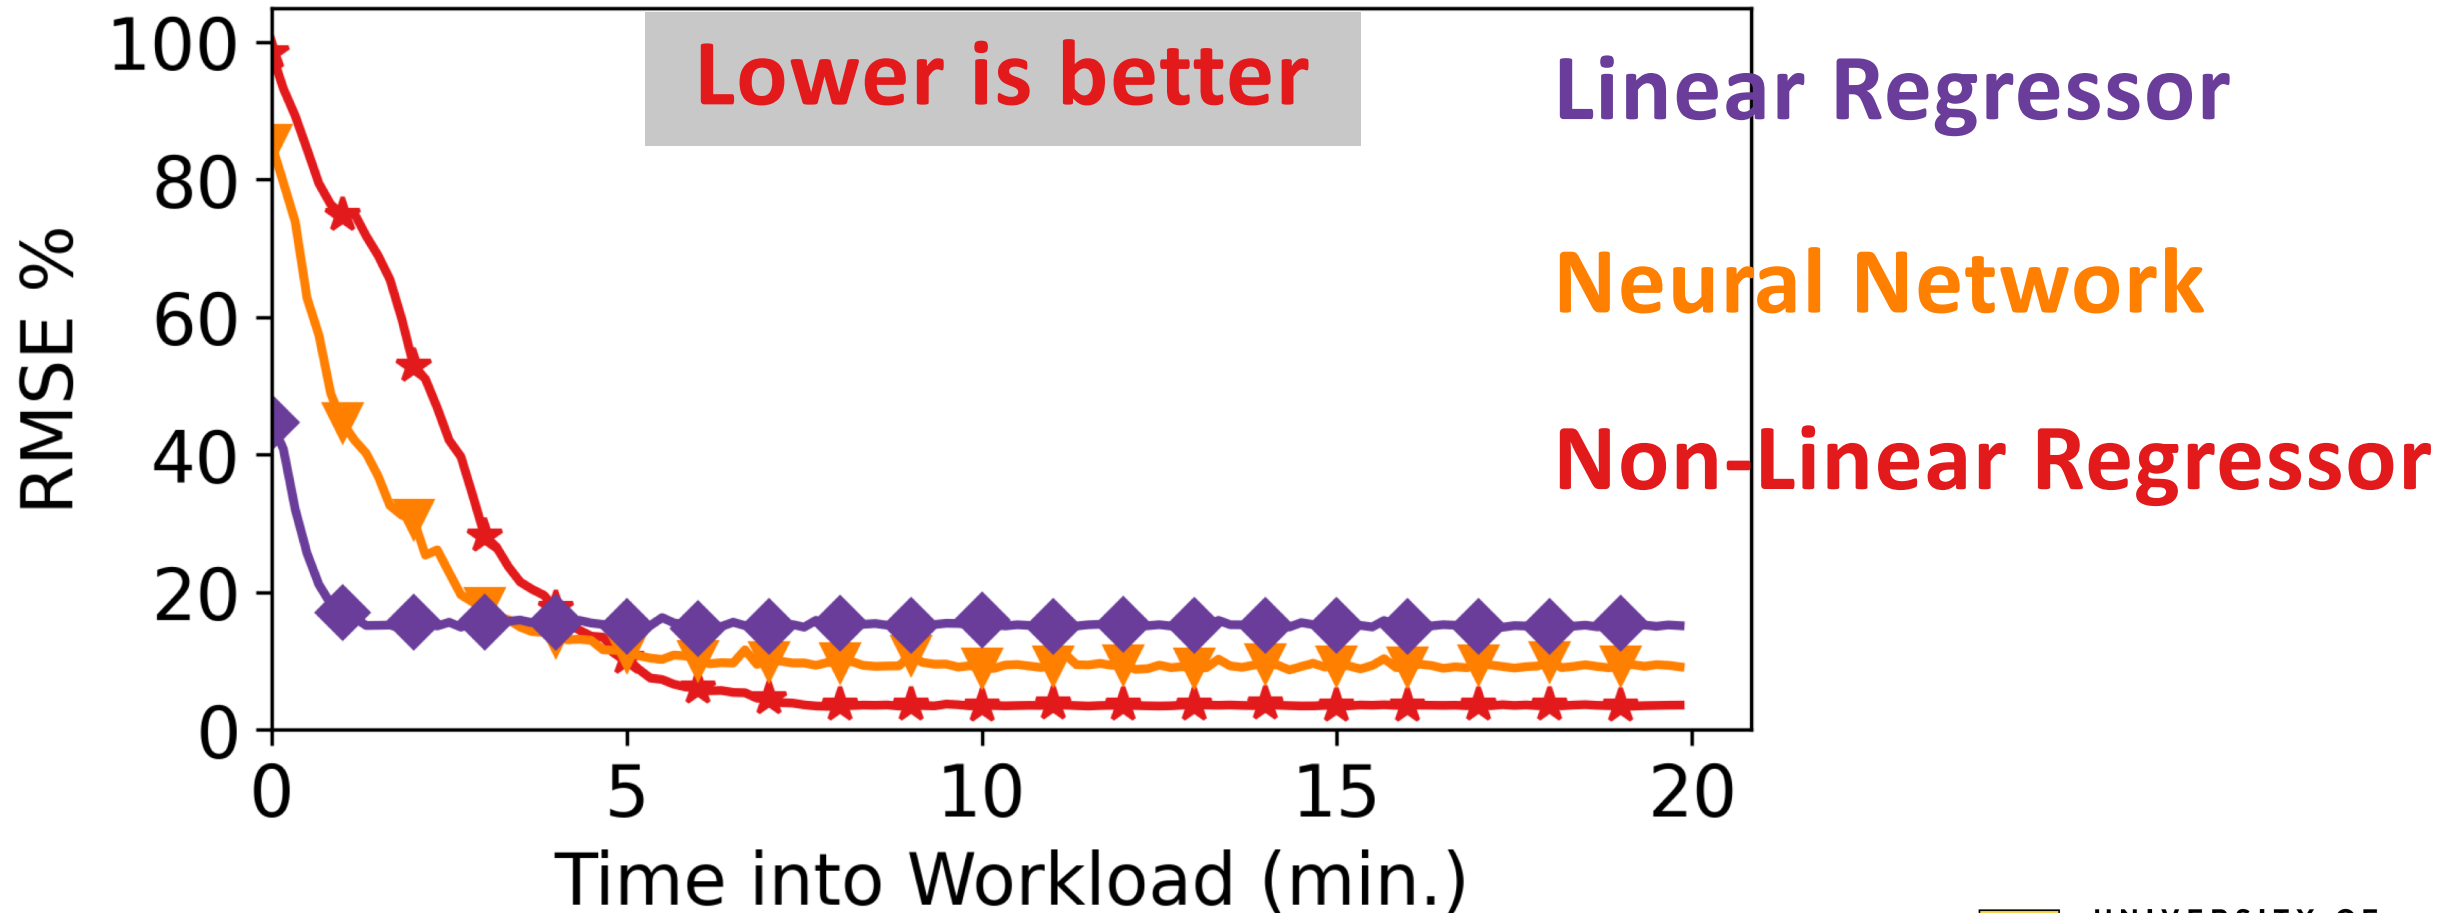
Per layout with workload stats as parameters

Linear Regressor

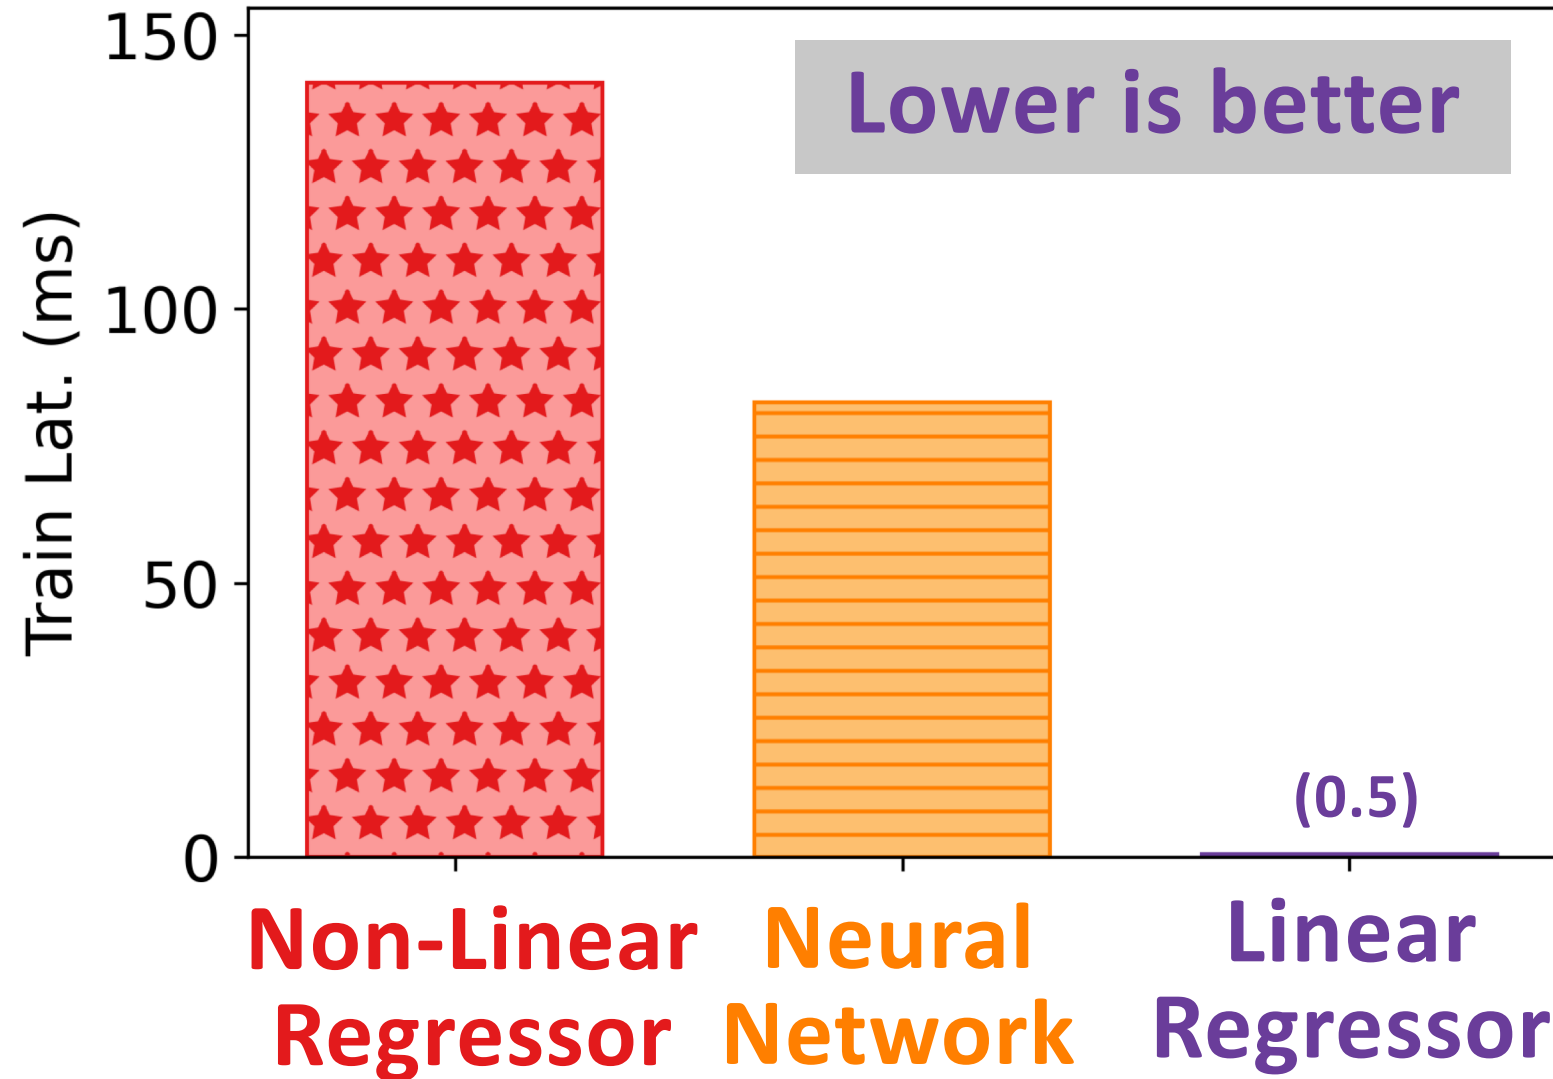
Neural Network

Non-Linear Regressor


Predicting Latency Accuracy



Predicting Latency - Training




Predicting Latency - Inference



Tiresias

Predict upcoming accesses and latency under different storage and indexing choices



Report query latency and access history

Storage and Index Changes

Expected Benefit of Change

Predicted access **latency** under:
current & proposed storage/indexing choice

Weighted by **likelihood** of access

Cost incurred to perform change

Tiresias End-to-End Evaluation

DBMS

PostgreSQL: automatic indexing

OLAP DBMS: predictive cracking

Proteus: adaptive storage for HTAP

Adapt with Tiresias

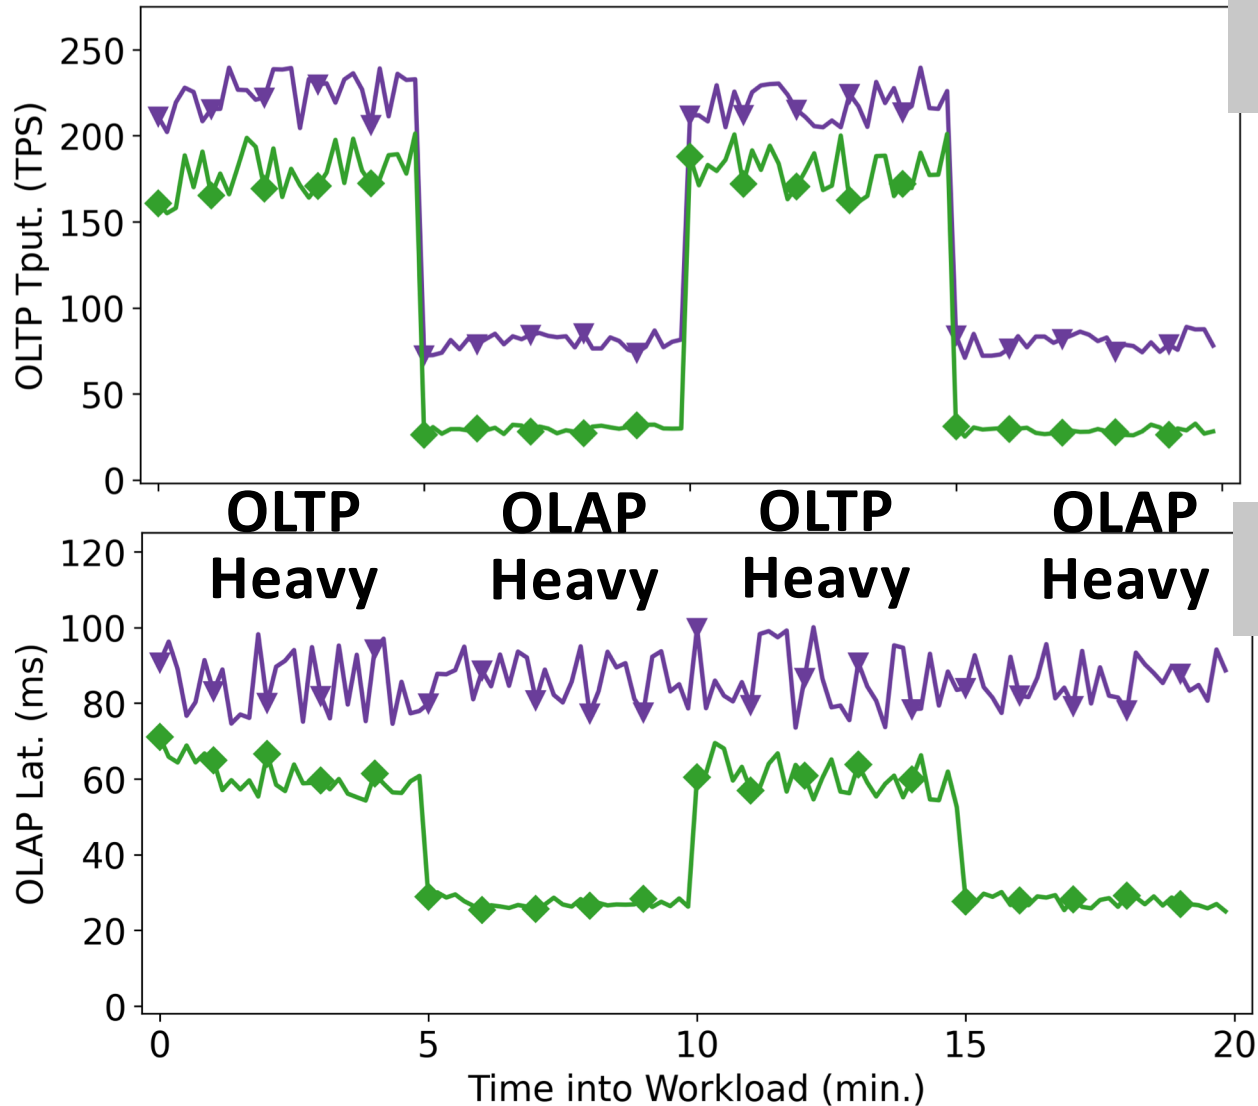
Static choices
without Tiresias

Shifting Workloads

Skew (hotspots)

Mix (OLTP or OLAP heavy)

End-to-End PostgreSQL Indexes



Higher is better

No Index


Index

Lower is better

No Index

Index

End-to-End PostgreSQL Indexes



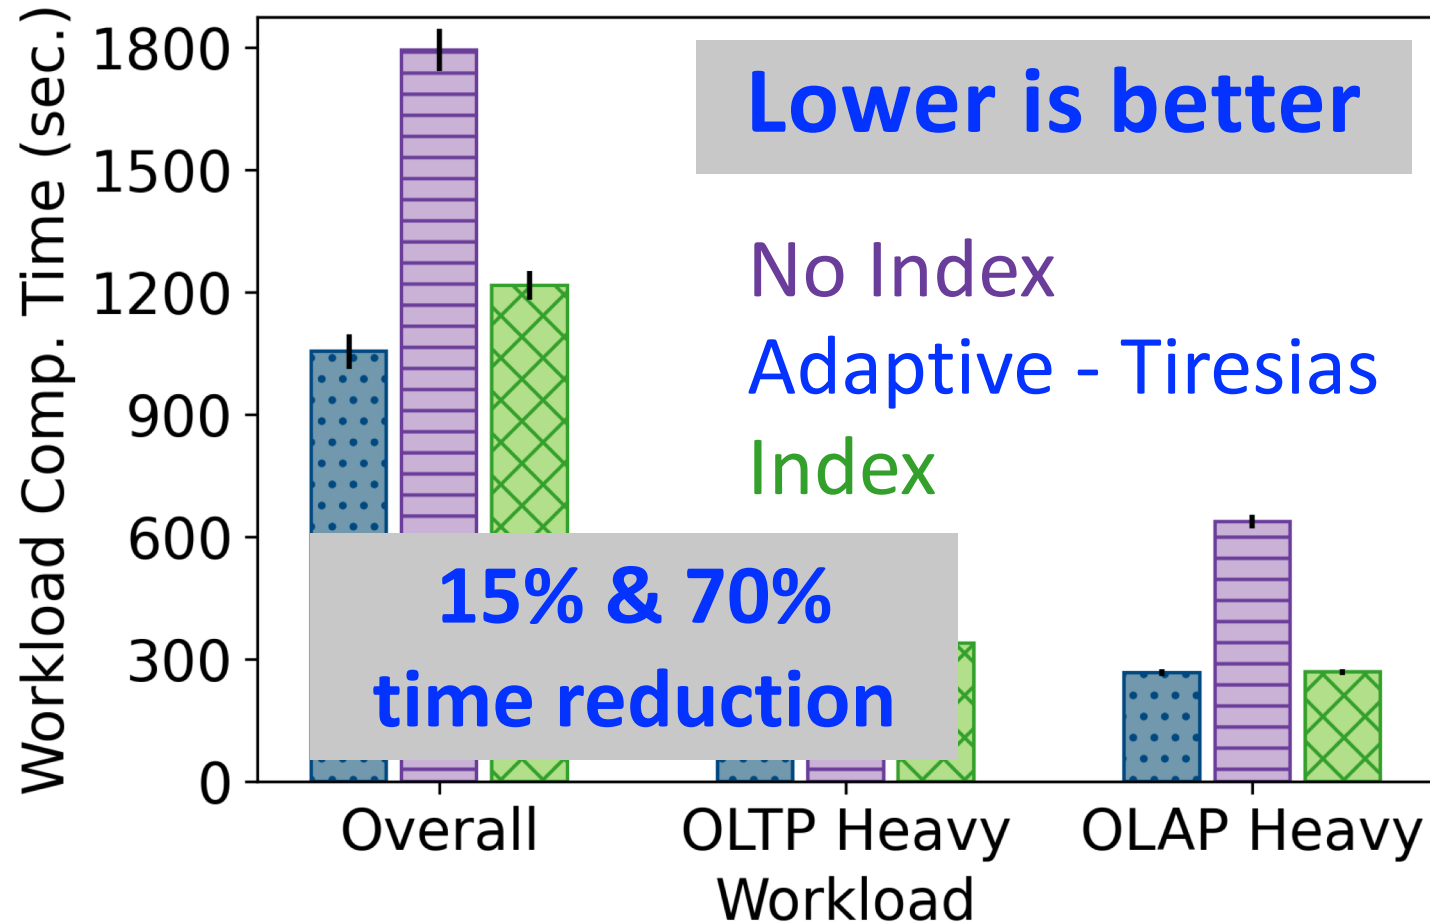
Higher is better

No Index
Adaptive - Tiresias
Index

Lower is better

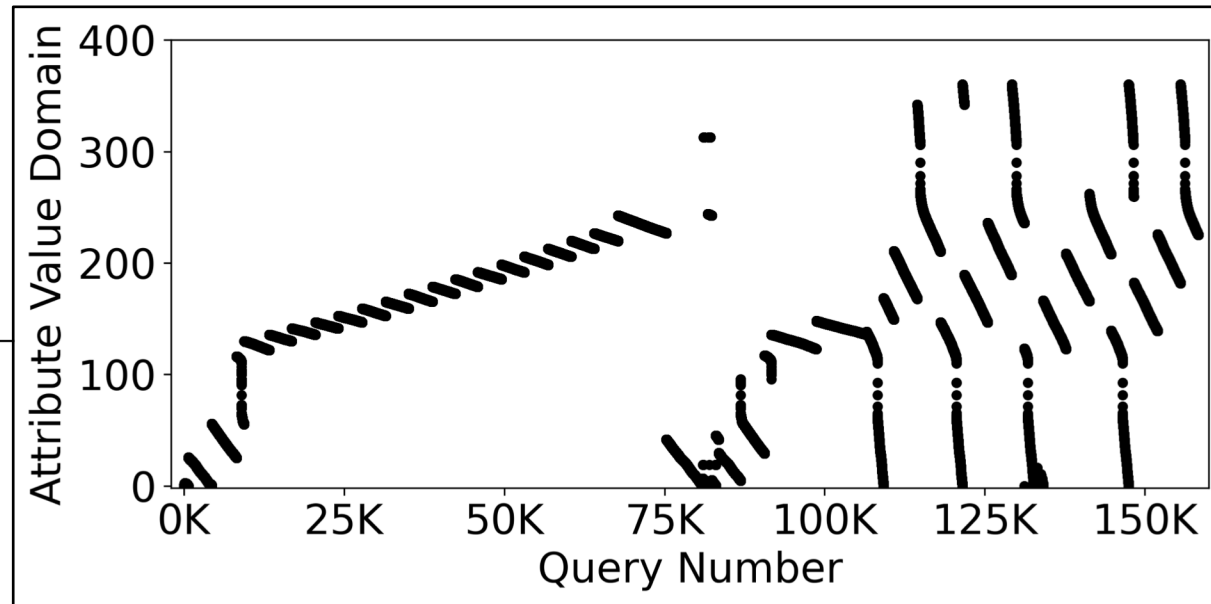
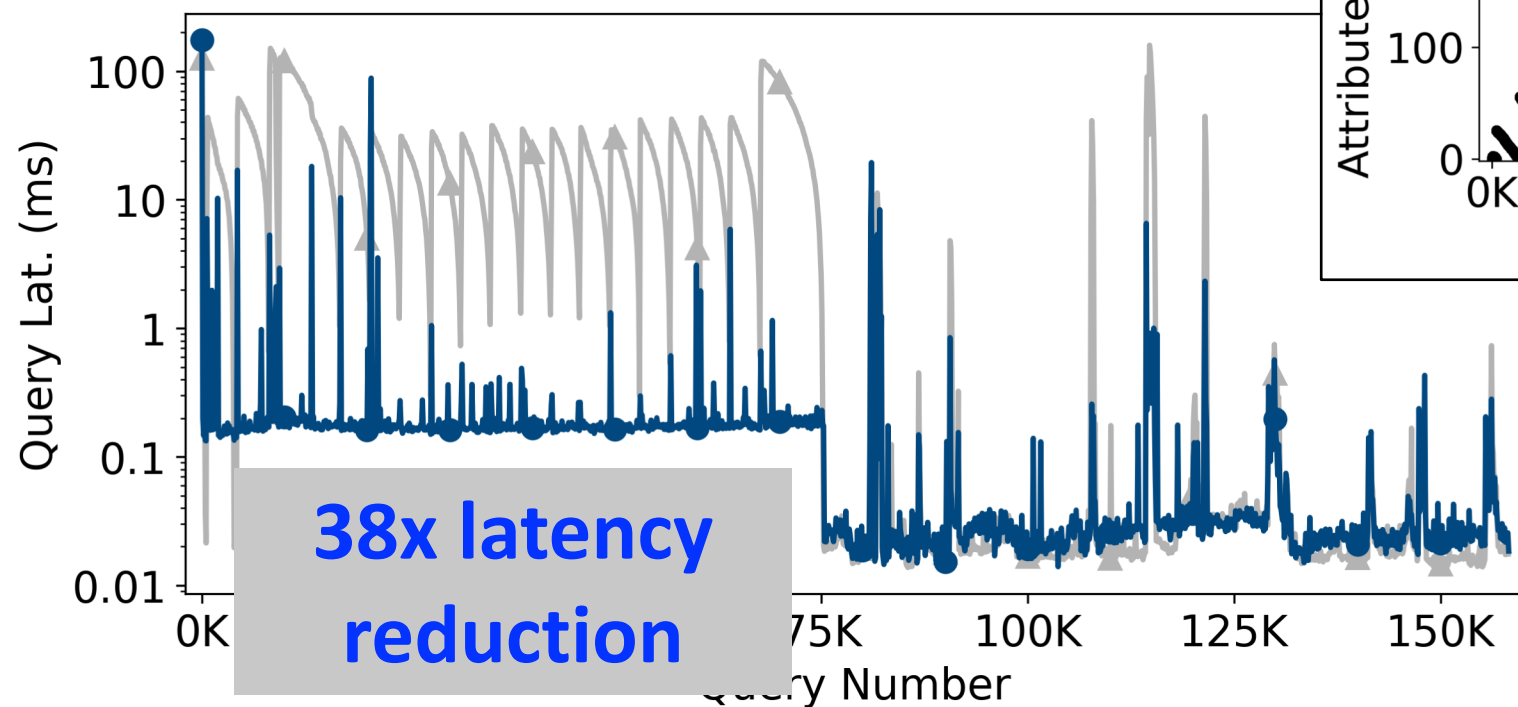
No Index
Adaptive - Tiresias
Index

End-to-End PostgreSQL Indexes



End-to-End Cracking


Lower is better



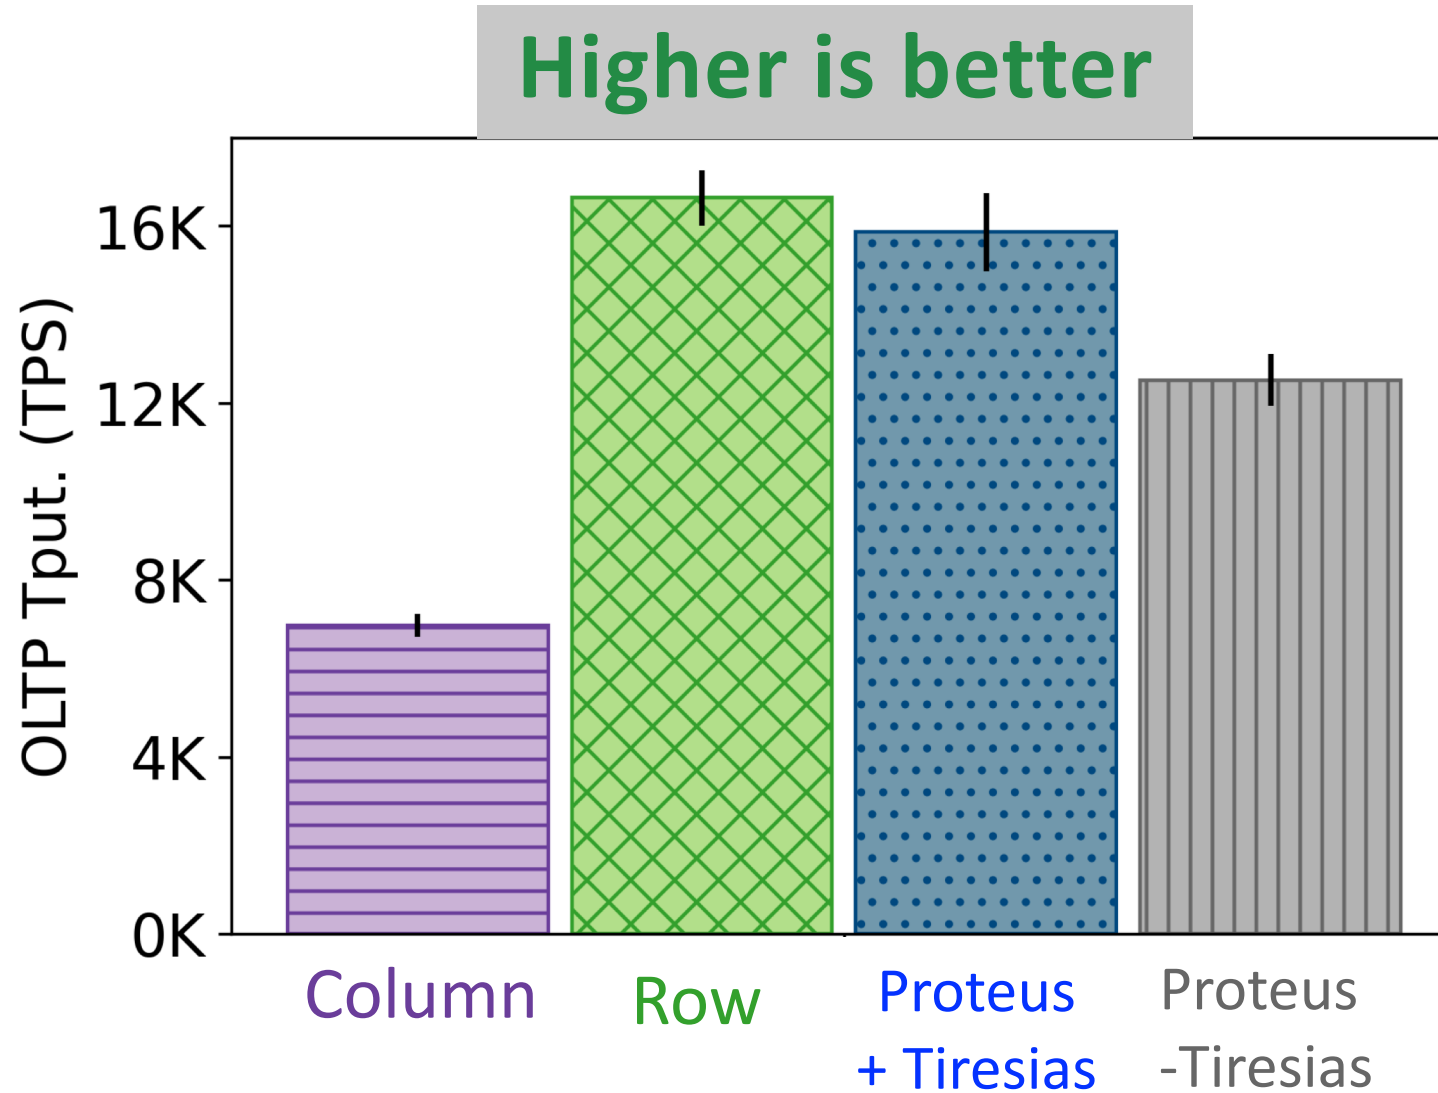
Cracking (No Tiresias)

Predictive Cracking: Tiresias


End-to-End Proteus



End-to-End Proteus



End-to-End Proteus



Tiresias Takeaways

tiny.cc/tiresias

Automatic adaptation of data storage and indexing

Generalizable API used in different DBMSs

Predict data accesses based on workload patterns

Learn access costs under different storage choices