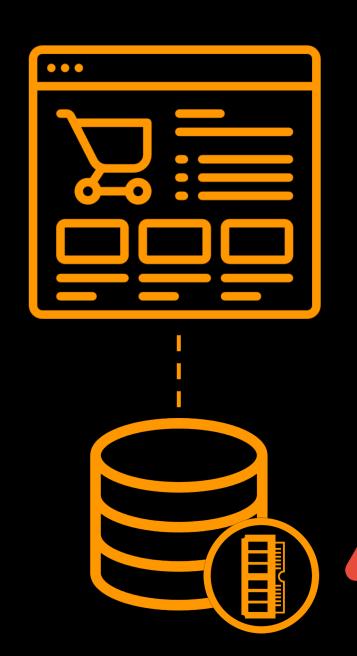


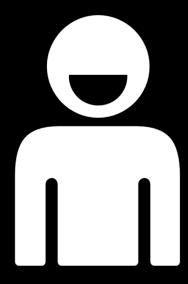
Reducing the Storage Overhead of Main-Memory OLTP Databases with Hybrid Indexes

Pavid G. Andersen

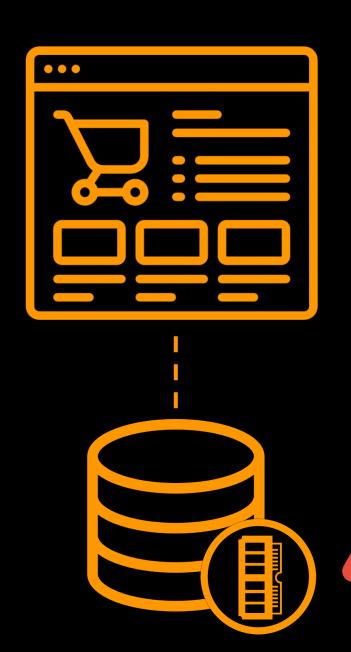
Huanchen Zhang







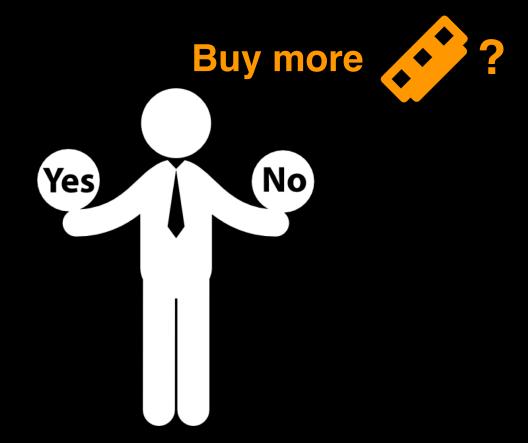




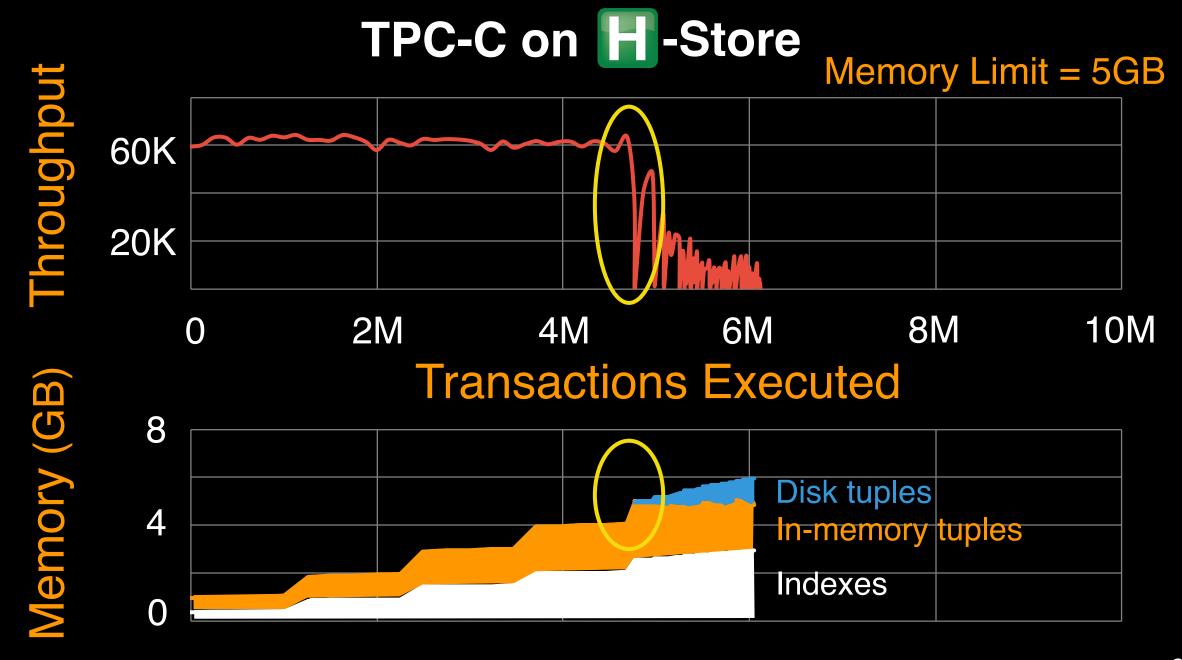








You are running out of memory





The better way: Use memory more efficiently



Indexes are LARGE

Hybrid Index Benchmark % space for index 58% → 34% TPC-C **55%** → 41% Voter 34% -> 18% Articles

Our Contributions

- The hybrid index architecture
- The Dual-Stage Transformation
- Applied to 4 index structures
 - B+tree Skip List
 - Masstree Adaptive Radix Tree (ART)

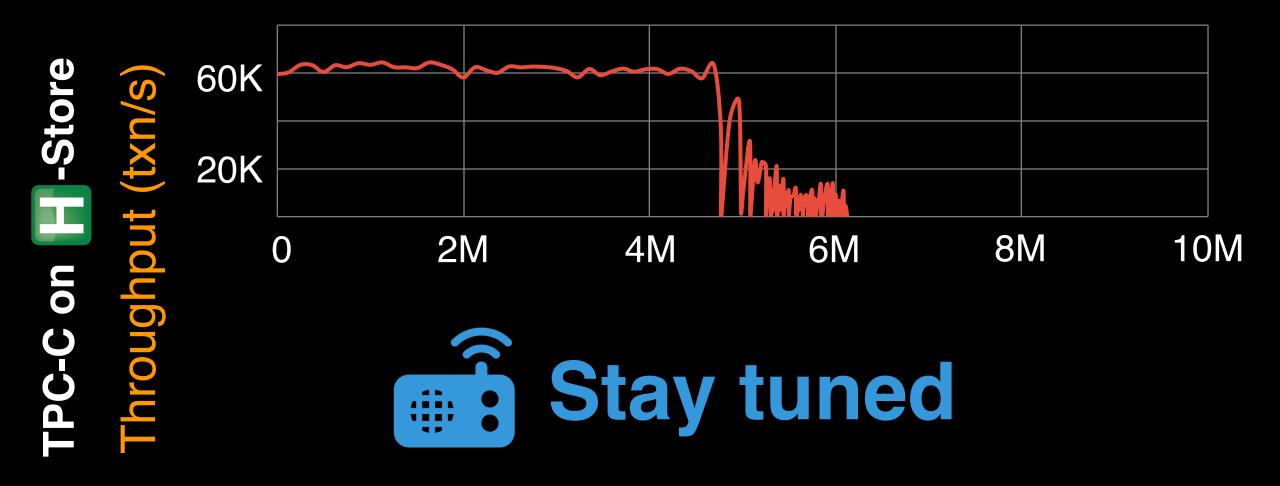
Performance

Space



30 - 70%

Did we solve this problem?



Transactions Executed

How do hybrid indexes achieve memory savings?

0— Static

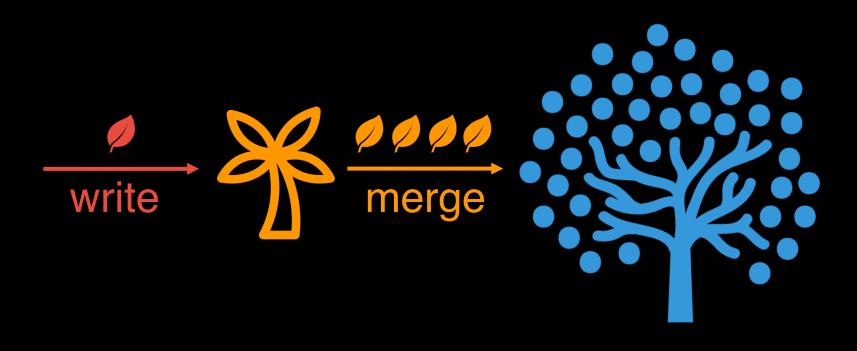
Hybrid Index: a dual-stage architecture





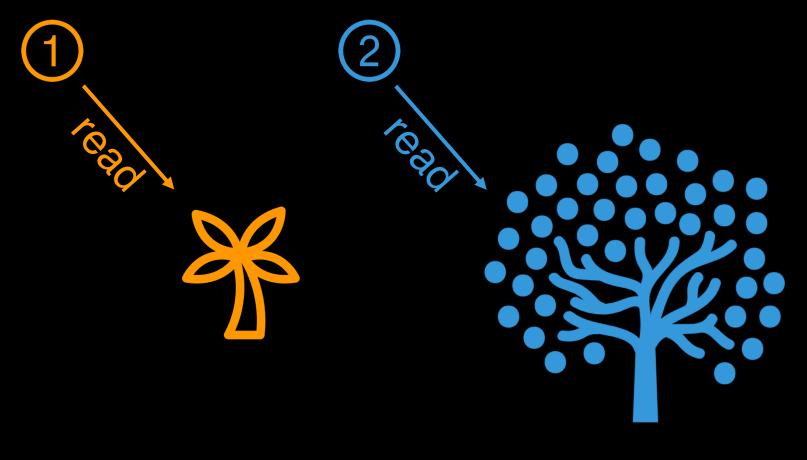
dynamic stage

Inserts are batched in the dynamic stage



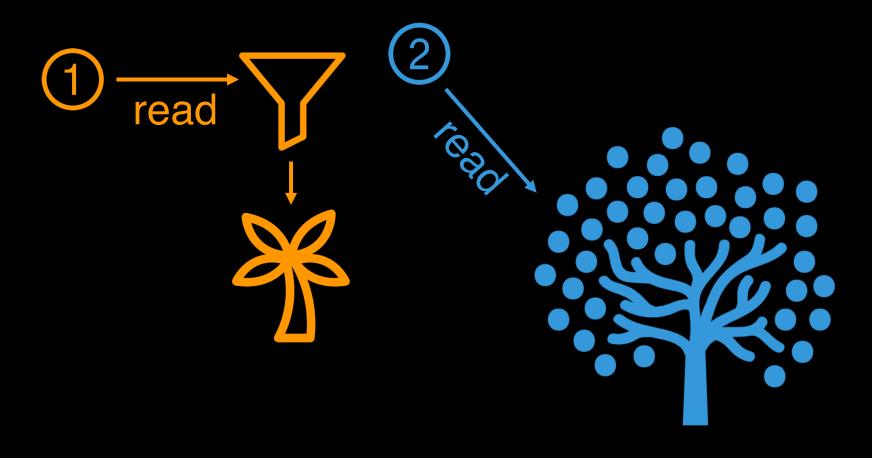
dynamic stage

Reads search the stages in order

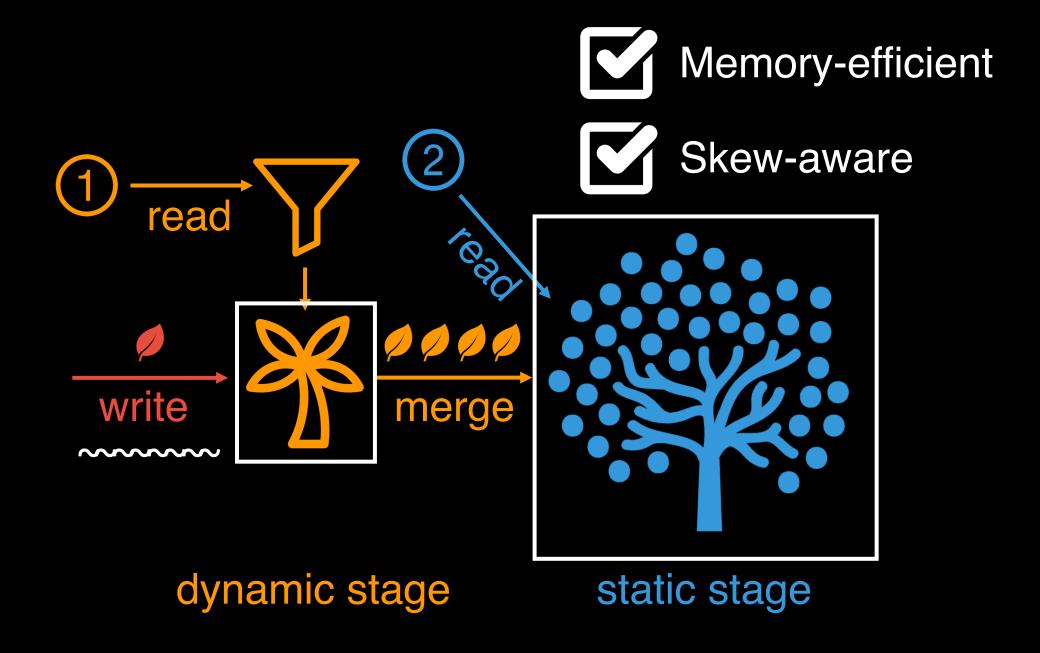


dynamic stage

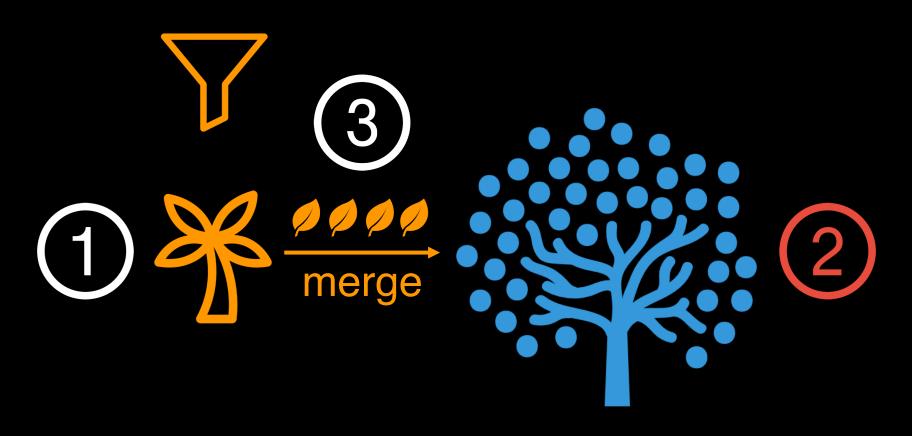
A Bloom filter improves read performance



dynamic stage



The Dual-Stage Transformation



dynamic stage

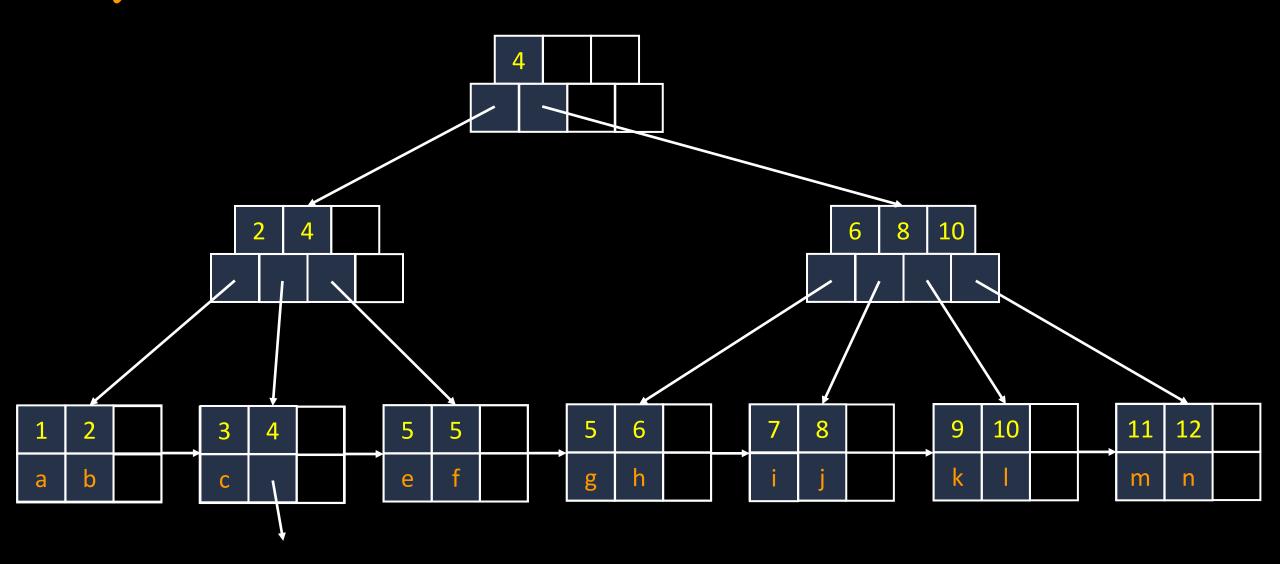
The Dynamic-to-Static Rules



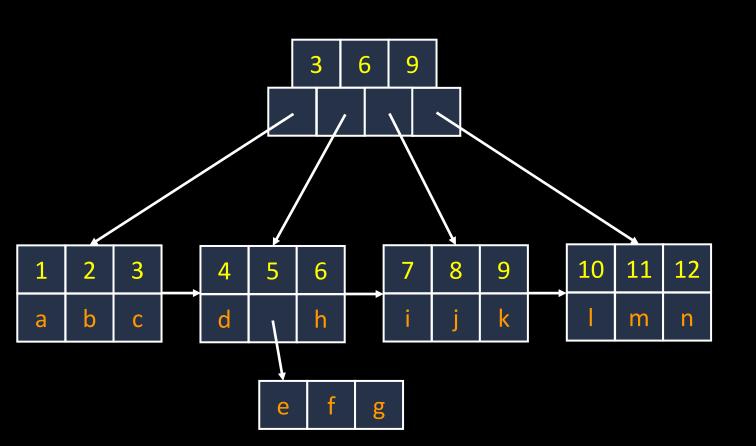




Compaction: minimize # of memory blocks

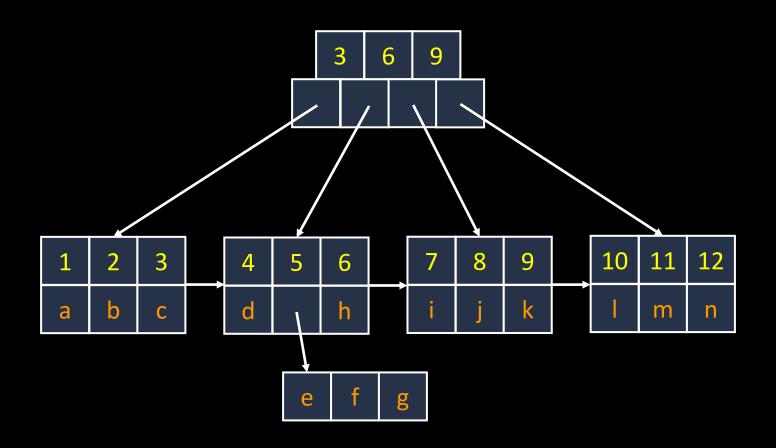


Compaction: minimize # of memory blocks

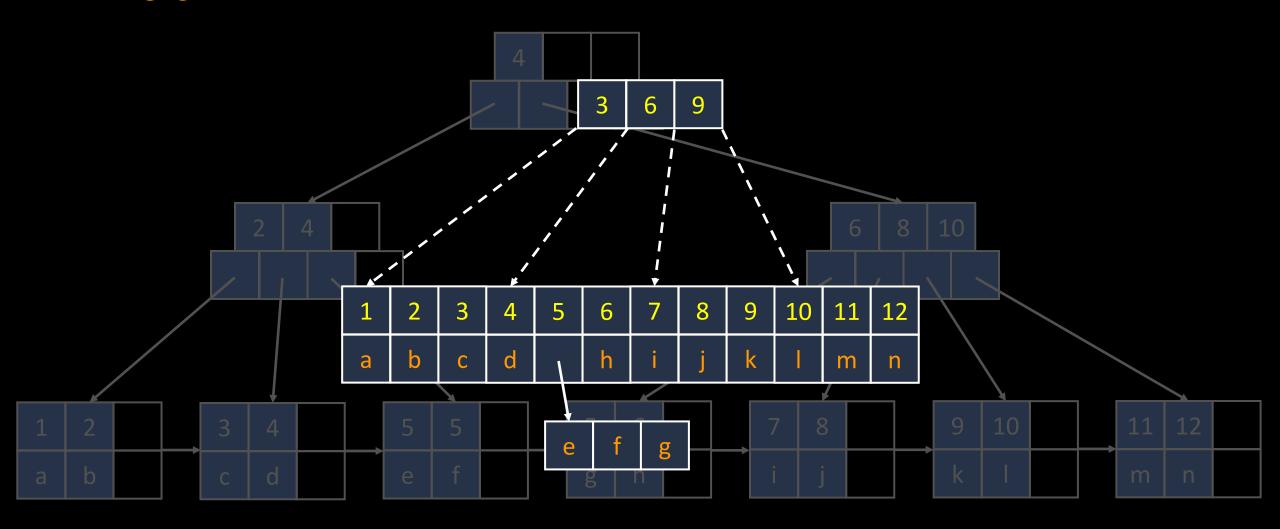




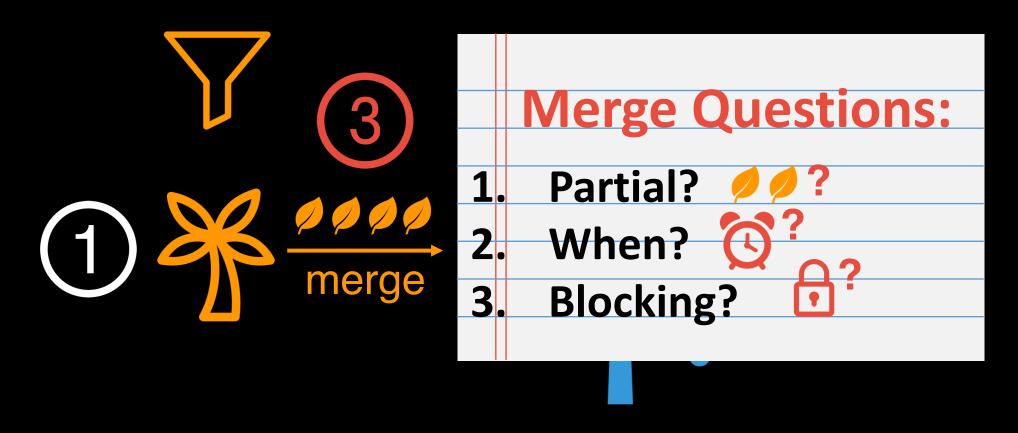
Reduction: minimize structural overhead



Reduction: minimize structural overhead

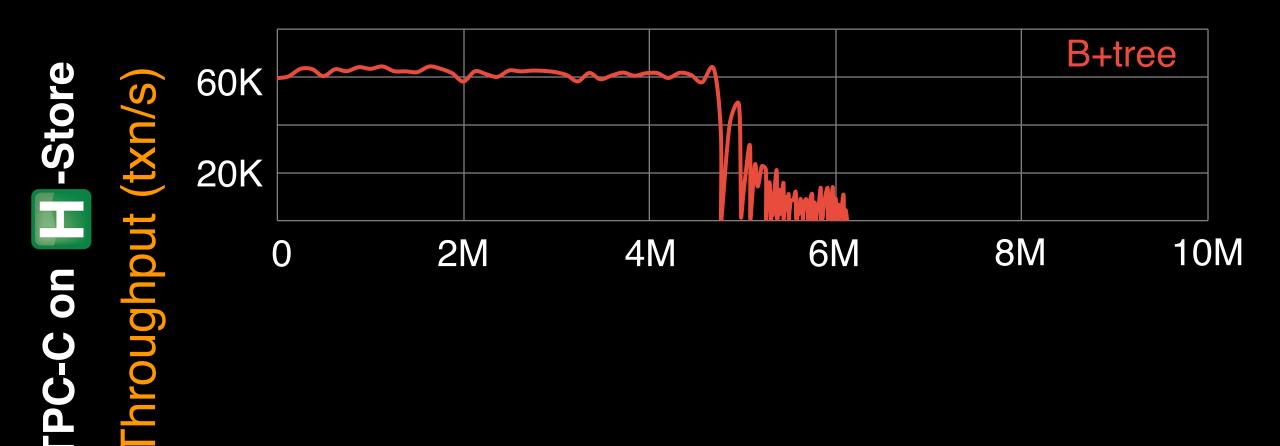


The Dual-Stage Transformation



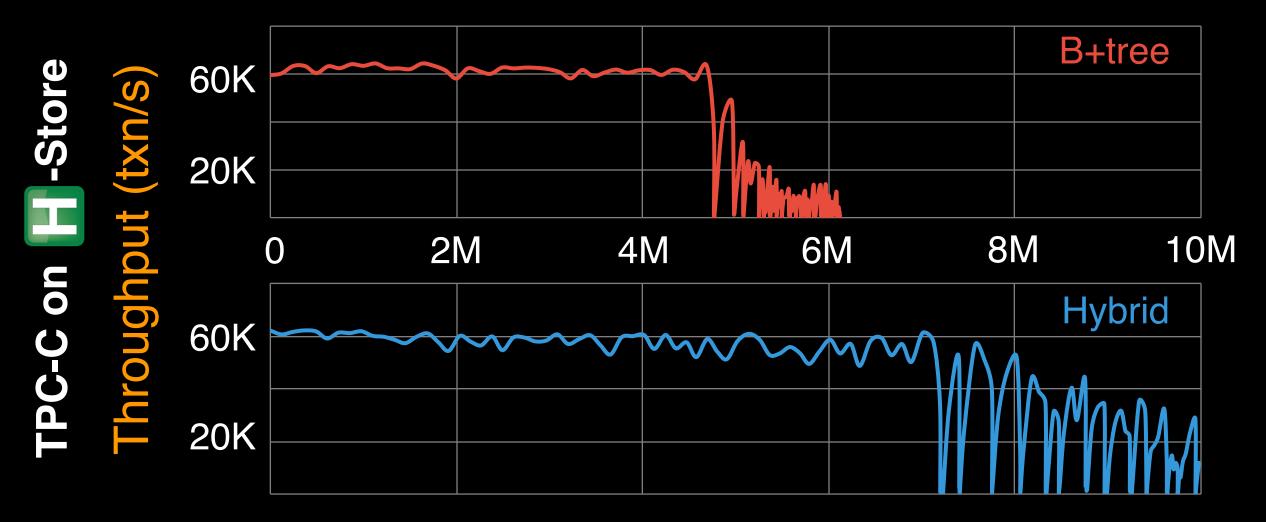
dynamic stage

Did we solve this problem?

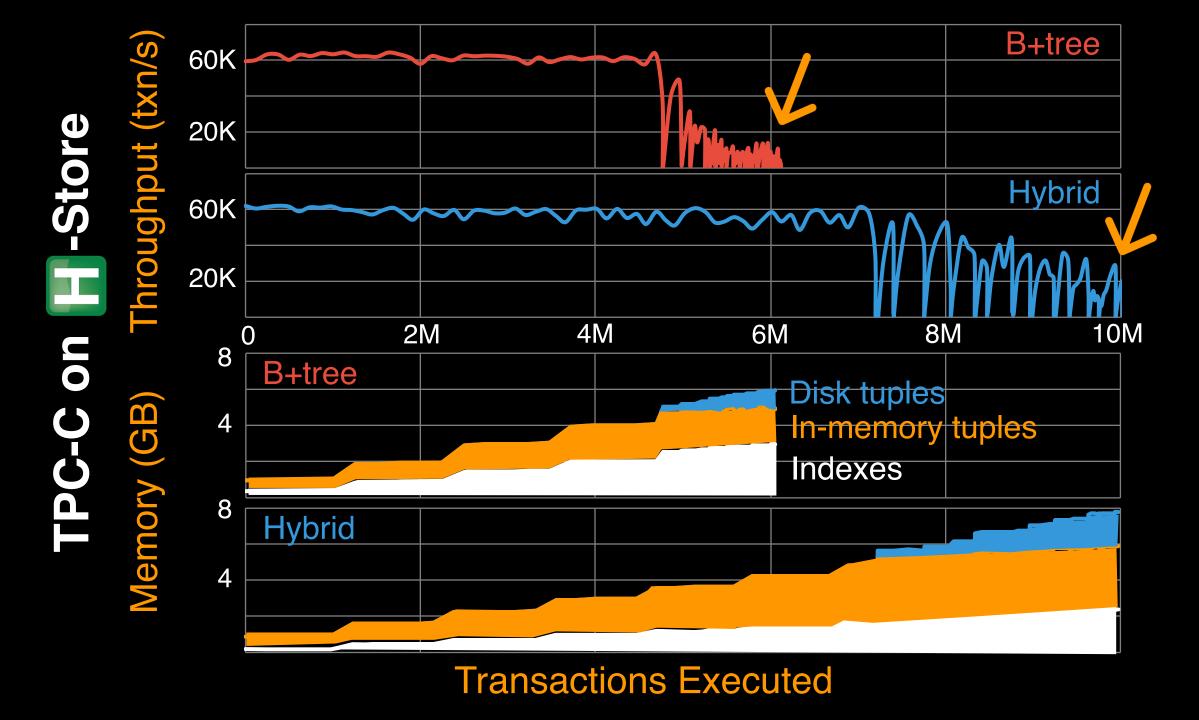


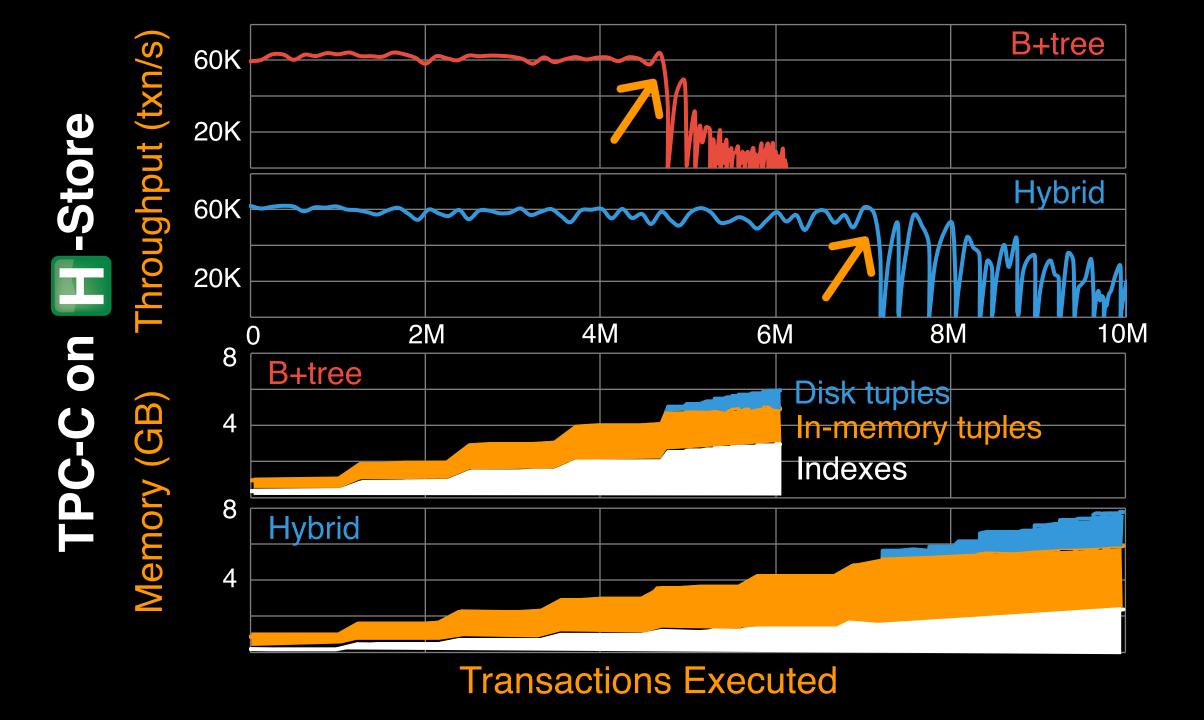
Transactions Executed

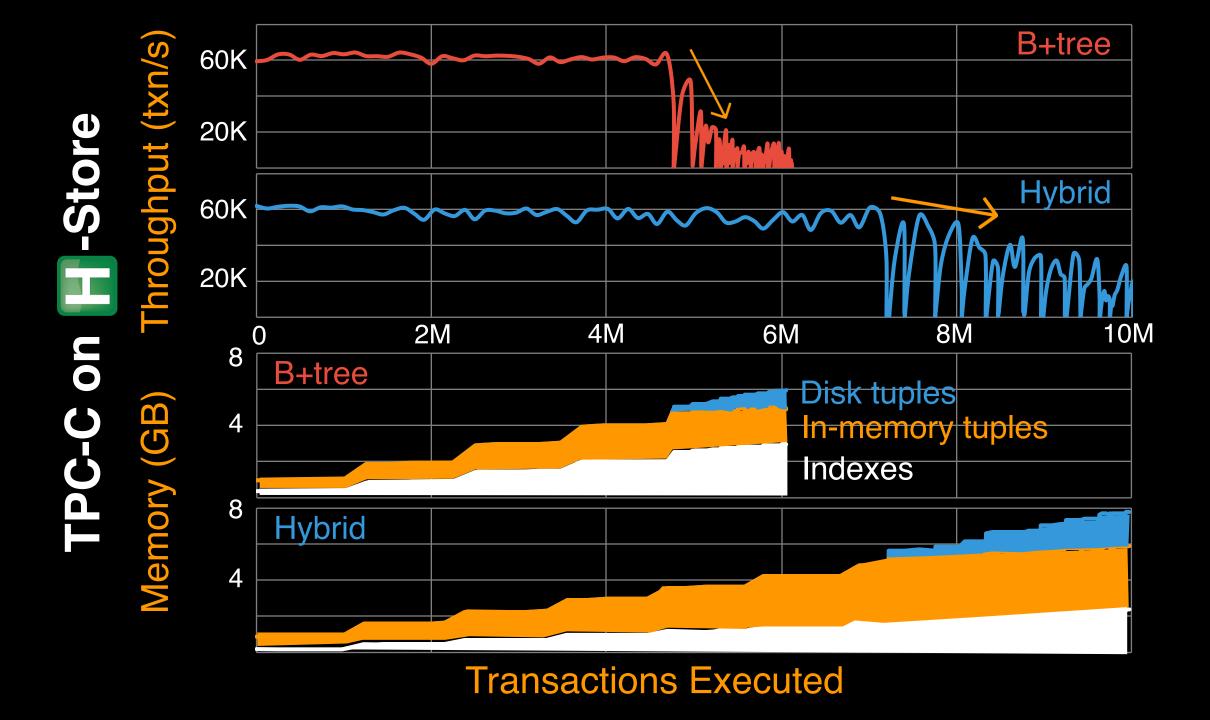
Yes, we improved the DBMS's capacity!

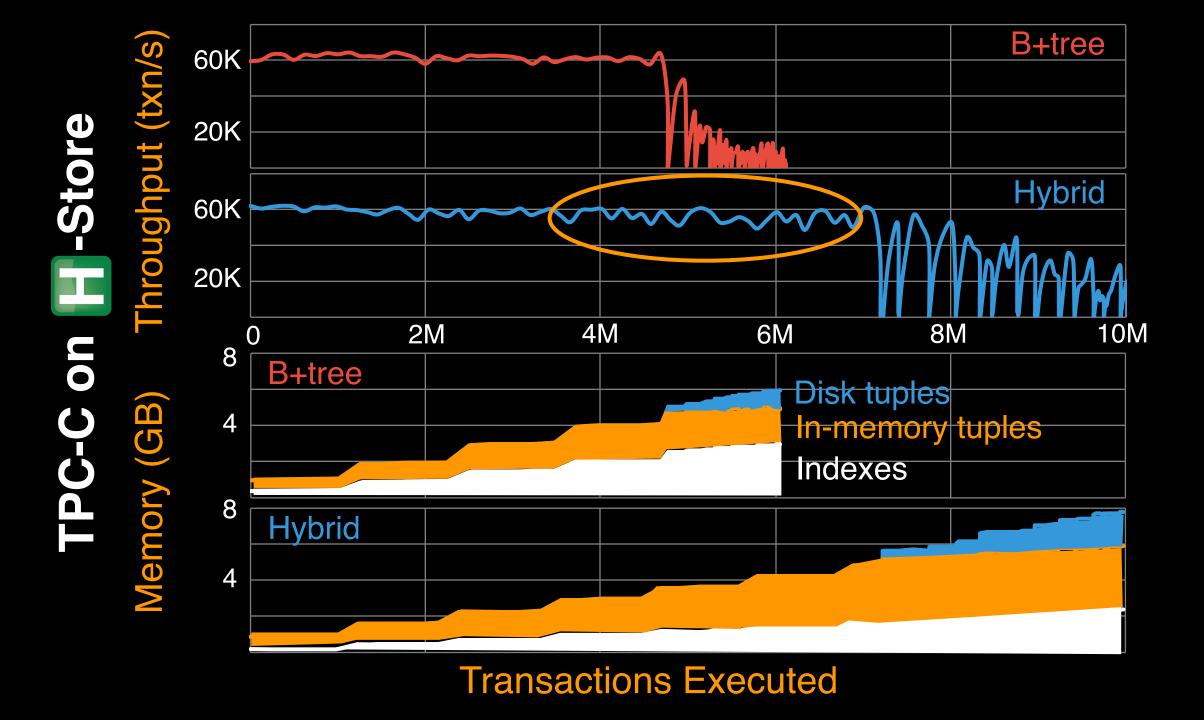


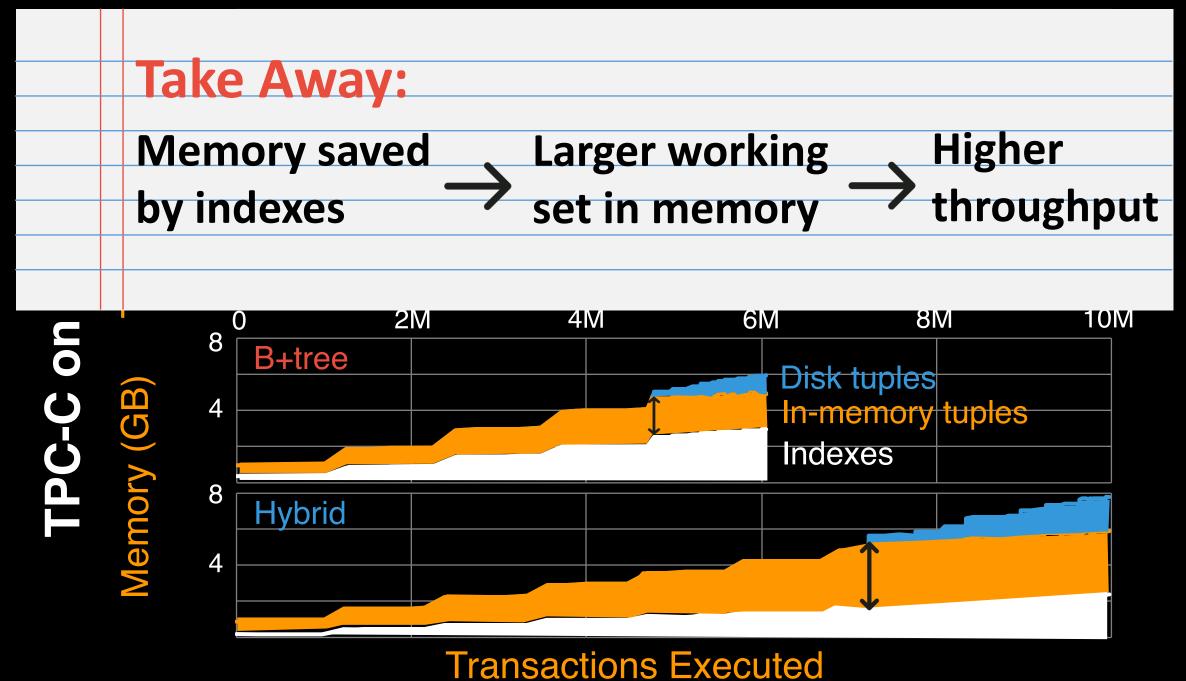
Transactions Executed











This is just the **BEGINNING**

Conclusions

- (1) The hybrid index architecture GENERAL
- (2) The Dual-Stage Transformation PRACTICAL
- (3) Applied to 4 index structures USEFUL
 - B+tree Skip List
 - Masstree Adaptive Radix Tree (ART)

Toll-Free Hotline:



1-844-88-CMUDB