




DuckDB in Science Meetup 2025 (London)

BOSTON
UNIVERSITY

Elevating Bitmap Indexing in OLAP DBMSs

CUBIT, RABIT, and ...

Junchang (Jason) Wang and Manos Athanassoulis



Bitmap Index Basics

Column A	A=10	A=20	A=30
30	0	0	1
20	0	1	0
30	0	0	1
10	1	0	0
20	0	1	0
10	1	0	0
30	0	0	1
20	0	1	0

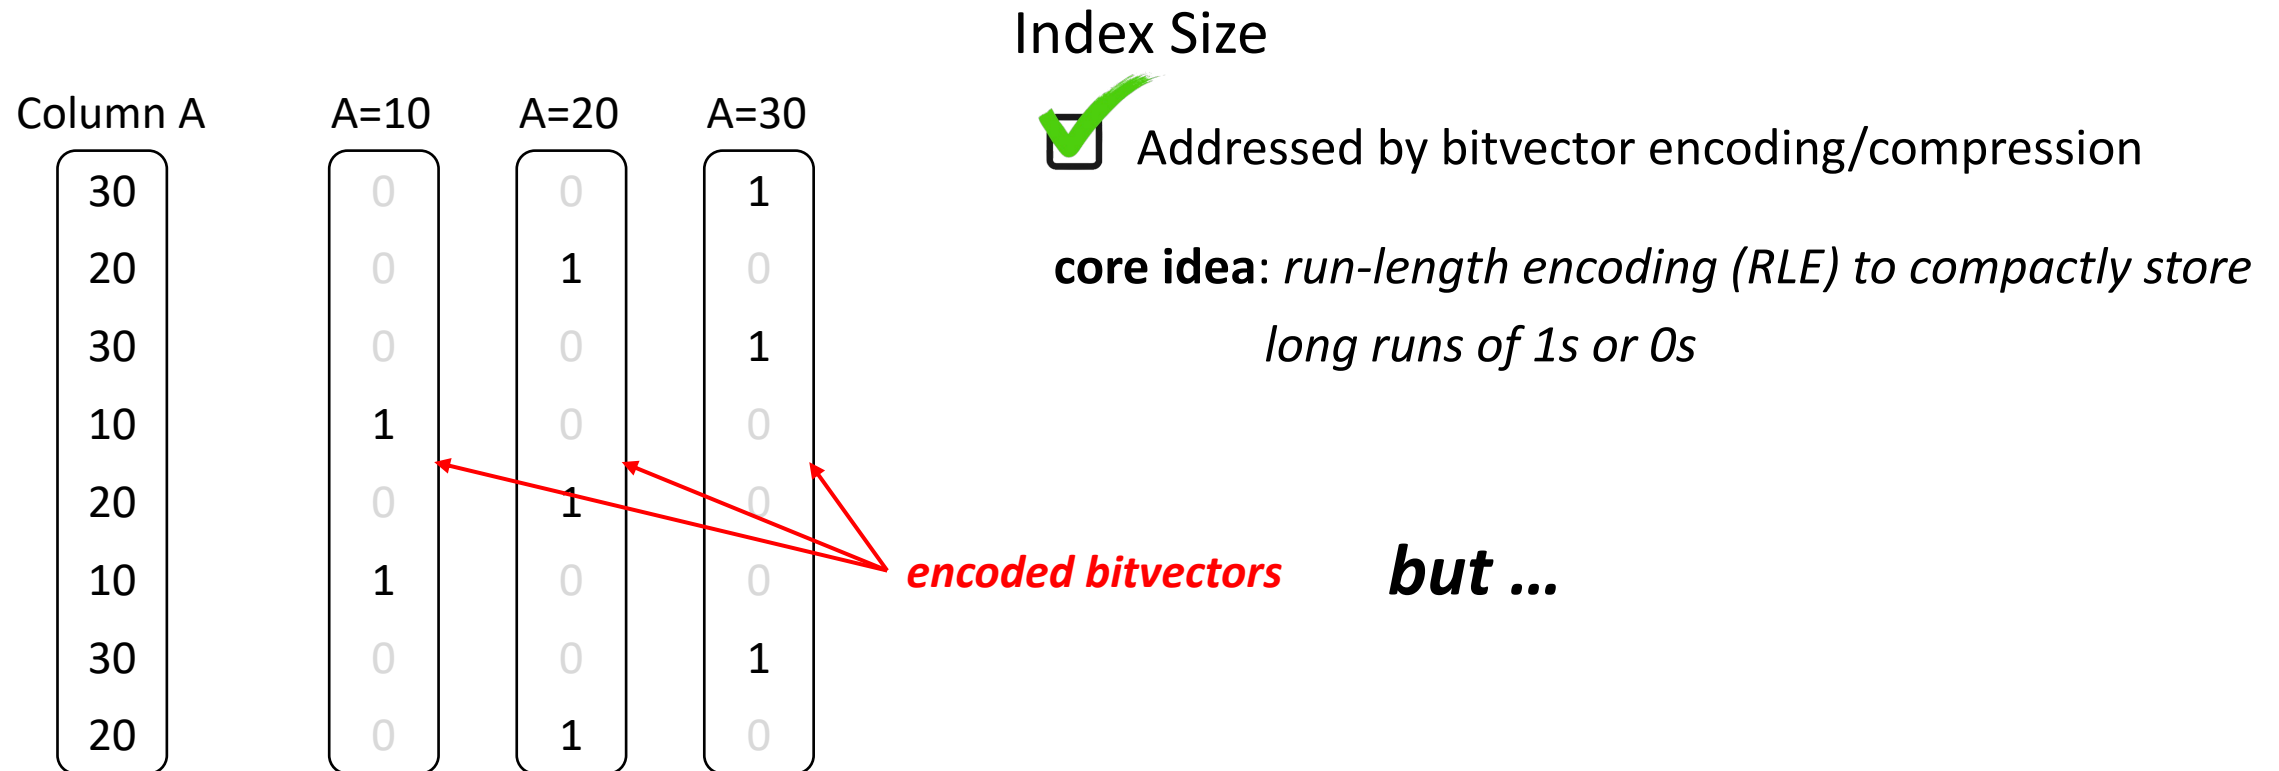
Specialized indexing

- ✓ Query result is readily available
- ✓ Compact representation of selective query result

Bitvectors

- ✓ 0s are highly compressible
- ✓ Efficient bitwise operations (e.g., OR/AND)
- ✓ Sequential access minimizes cache and TLB misses

Bitmap Indexing Limitations



☒ Updating encoded bitvectors is **very** inefficient

CUBIT's Goal


Update-friendly Bitmap Indexes on Multicore Systems

- **Queries are wait-free** with guaranteed completion
- **Updates are lock-free** and avoid blocking each other

Concurrent Updatable BITmap Indexing (CUBIT)



Out-of-place update mitigates synchronization complexity



Multi-versioning enables parallel execution of queries and updates



Bitvector Segmentation exploits parallelism on multicore systems

CUBIT

Goals

Experiments



Out-of-place update

Mitigate synchronization


Synthetic data

n : # tuples

d : # domain values (cardinality)

q : # queries

u : % updates in the workload



Multi-versioning

Parallel queries and updates

Prototype C++ implementation of CUBIT, UpBit, UCB, and In-place using FastBit



Bitvector Segmentation

Parallelism on multicores


Integrated into a row-store prototype system (DBx1000) and a column-store system (DuckDB)

Integrating CUBIT into DuckDB

- TPC-H SF 10 dataset
- DuckDB version 1.0
- Single-core execution to focus on executor logic

Maintain CUBIT instances for table attributes under updates

- l_quantity (cardinality = 50)
- l_shipdate (cardinality = 2,526)
- l_orderkey (cardinality = 15M)
- ...



Implement CUBIT-based Scan, Aggregation, and Join executors

- See our paper for details

DuckDB makes it easy to integrate and experiment with CUBIT
CUBIT accelerates not only Scan, but also Aggregation and Join

To be continued ...

- CUBIT mainly focuses on point queries**, meaning
- it mainly adopts equality encoding scheme
 - each query reads only one bitvector

However,

- to support range queries** (e.g., "A ≤ 8"),
- **Extend updatability to encoding schemes beyond equality encoding**
 - **Support high-cardinality attributes**

← Must OR many bitvectors →

	E ₀	E ₁	E ₂	E ₃	E ₄	E ₅	E ₆	E ₇	E ₈	E ₉
1	3	0	0	0	1	0	0	0	0	0
2	2	0	0	1	0	0	0	0	0	0
3	1	0	1	0	0	0	0	0	0	0
4	2	0	0	1	0	0	0	0	0	0
5	8	0	0	0	0	0	0	0	1	0
6	2	0	0	1	0	0	0	0	0	0
7	9	0	0	0	0	0	0	0	0	1
8	0	1	0	0	0	0	0	0	0	0
9	7	0	0	0	0	0	0	1	0	0
10	5	0	0	0	0	1	0	0	0	0
11	6	0	0	0	0	0	1	0	0	0
12	4	0	0	0	1	0	0	0	0	0

Equality Encoding (EE)

Barely compressible
High update cost

	R ₀	R ₁	R ₂	R ₃	R ₄	R ₅	R ₆	R ₇	R ₈
0	0	0	0	1	1	1	1	1	1
0	0	1	1	1	1	1	1	1	1
0	1	1	1	1	1	1	1	1	1
0	0	1	1	1	1	1	1	1	1
0	0	0	0	0	0	0	0	0	1
0	0	1	1	1	1	1	1	1	1
0	0	0	0	0	0	0	0	0	0
1	1	1	1	1	1	1	1	1	1
0	0	0	0	0	0	0	1	1	1
0	0	0	0	0	0	1	1	1	1
0	0	0	0	0	1	1	1	1	1
0	0	0	0	1	1	1	1	1	1


Range Encoding (RE)

CUBIT is Update-Friendly for Multicore Processors

DuckDB enables rapid integration and test with CUBIT !



Out-of-place update



Multi-versioning



Bitvector Segmentation

Thanks!



<https://disc.bu.edu>

We are pushing bitmap toward general secondary indexes

Please contact if you are interested

Thank you!