Virtual Lightweight Snapshots for Consistent Analytics in NoSQL Stores

Fernando Chirigati *NYU Tandon*
Jérôme Siméon *IBM Watson Research*
Martin Hirzel *IBM Watson Research*
Juliana Freire *NYU Tandon*

IBM Research

NYU TANDON SCHOOL OF ENGINEERING

ODM Insights

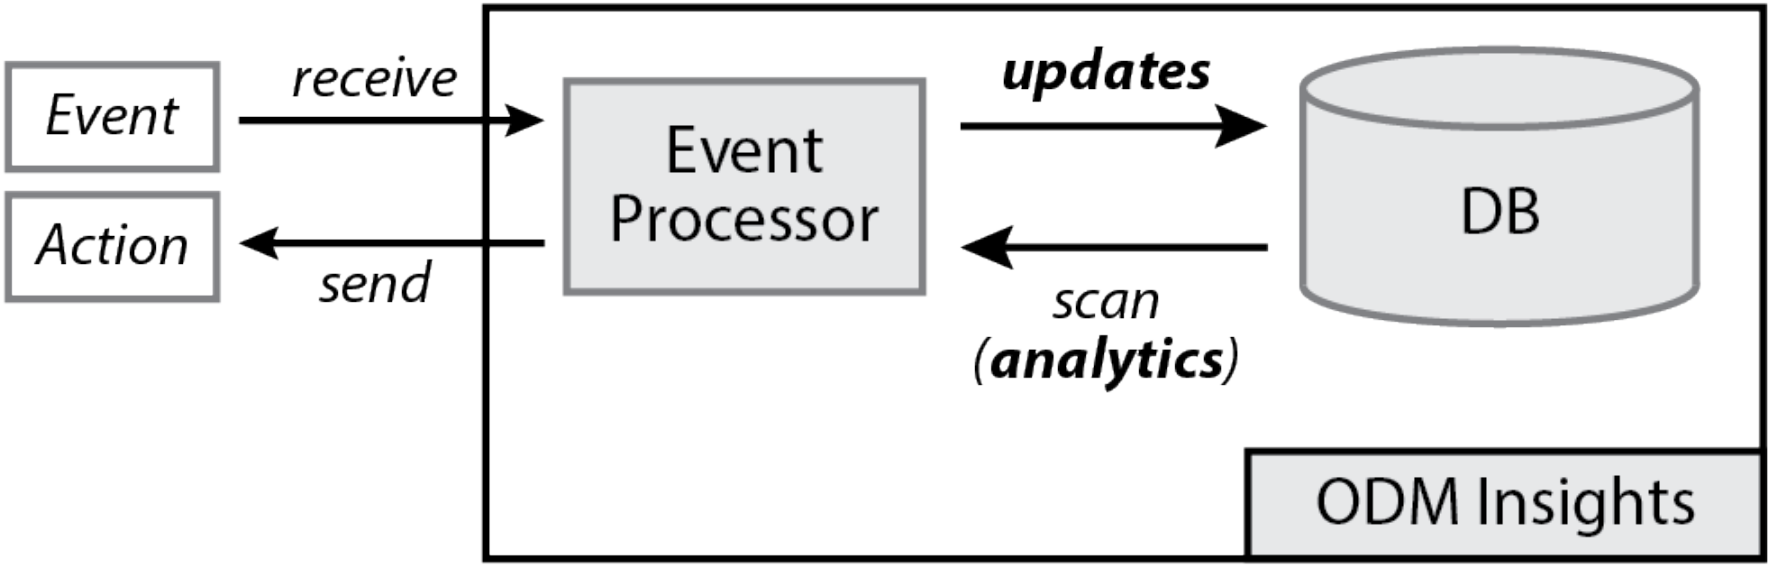

ODM Insights

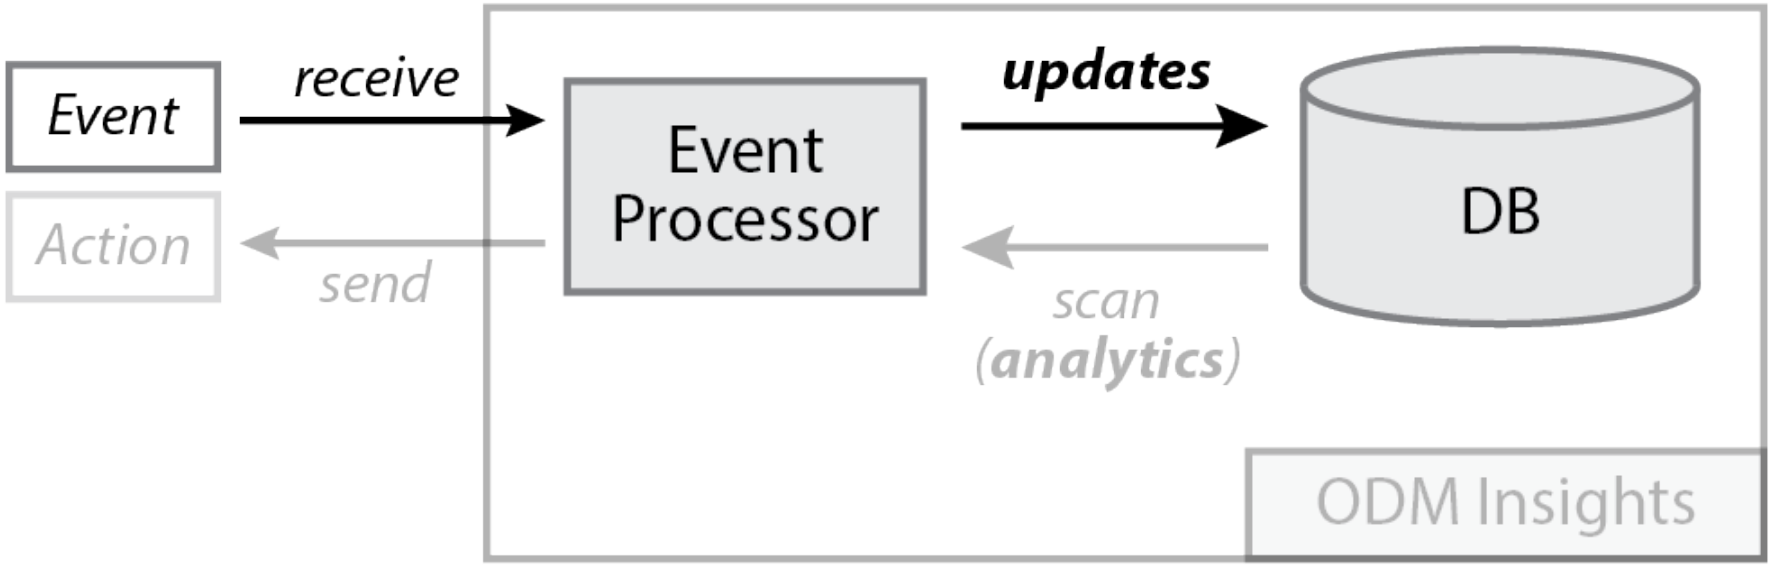

ODM Insights

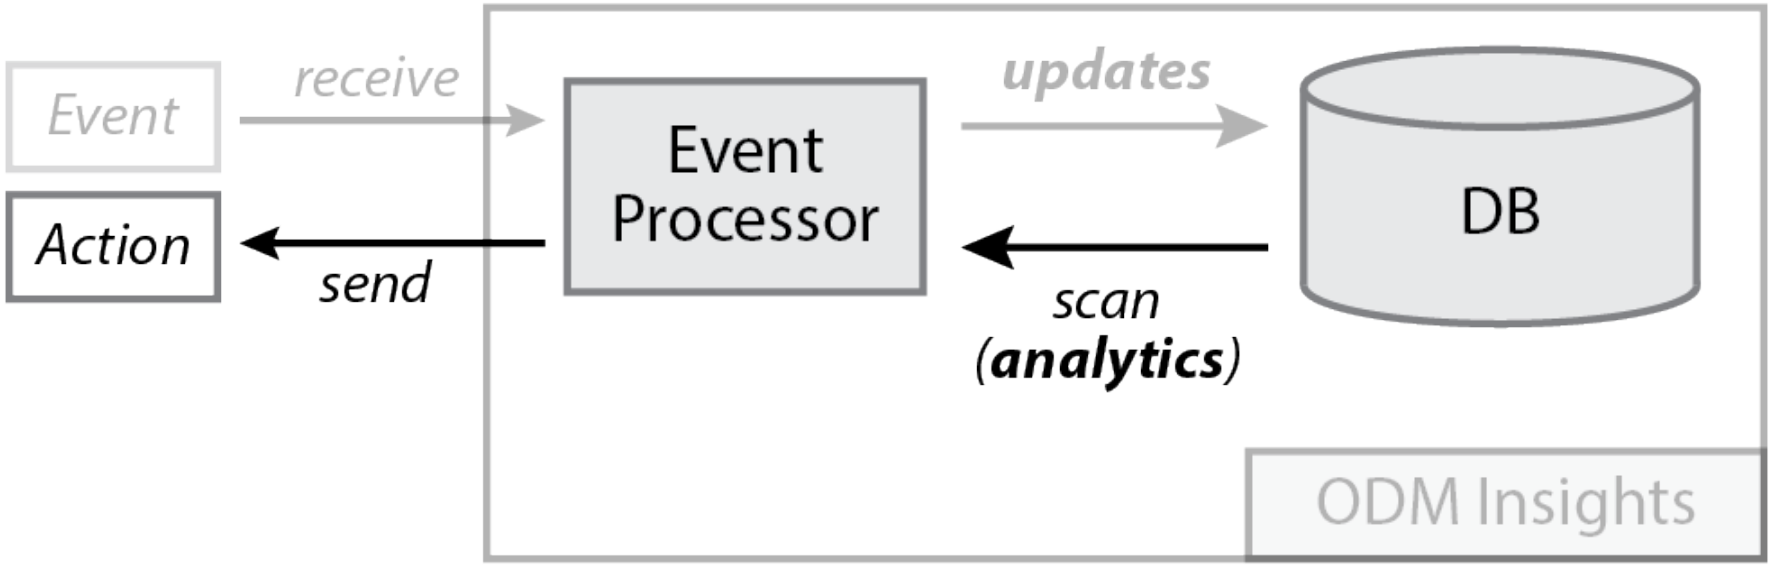

Air Traffic Control

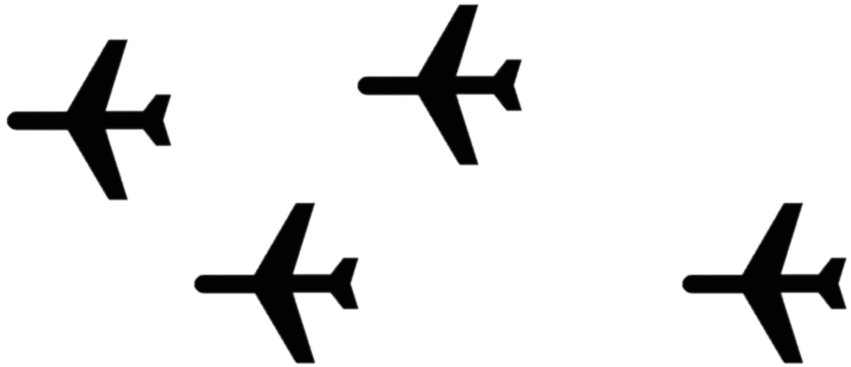

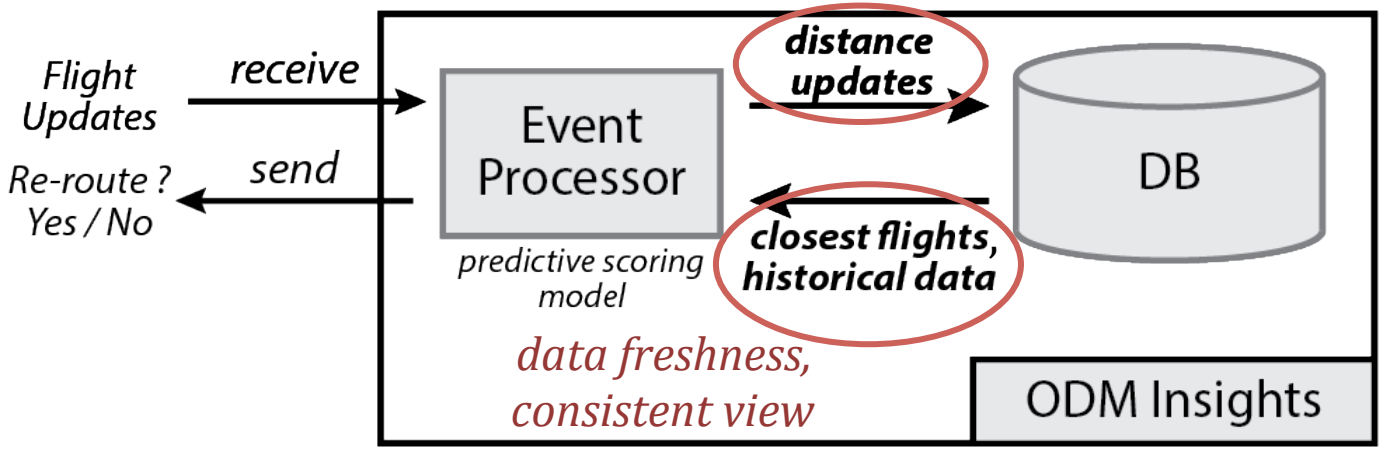

Updates

fast and numerous
(fast data)

X

Analytics

long and expensive
(big data)

Same system (fresh analytics) !

Impact on update throughput and latency

Same system, but lower level of isolation !

Inconsistent results

throughput / latency

consistent analytics

NoSQL Store

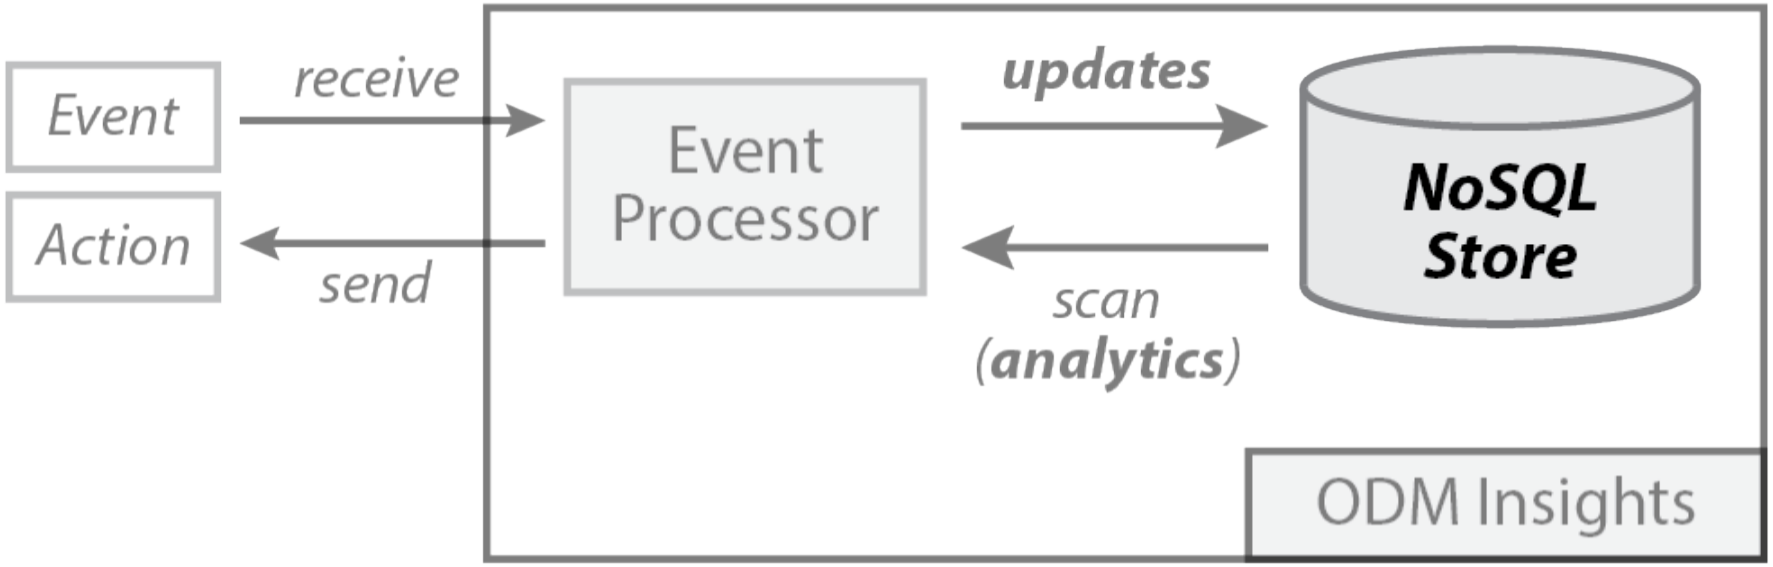

"read uncommitted"

Problem Statement

How to guarantee analytics' results to be ***fresh*** and ***consistent*** without impacting update latency and throughput on MongoDB-like NoSQL stores?

- *No support for ACID transactions*
- *No database versioning*
- *Disk-resident data*
- *Non-distributed data*

Previous Work

- ***Transient Versioning Techniques***
  [Bober and Carey, ICDE'92] [Mohan, et al., SIGMOD'92]
  [Chan, et al., SIGMOD'82] [Weihl, IEEE TSE'87] [Rastogi, et al., VLDB'97]

- ***Snapshot Isolation Techniques***
  [Padhye and Tripathi, CLOUD'12]

- ***Combined OLAP and OLTP Workloads***
  HyPer, SAP HANA, Hekaton, HYRISE, R-Store

- ACID Transactions
- Main-Memory Databases
- Native Database Versioning

Virtual Lightweight Snapshots
VLS

Inspiration:
Concurrent Garbage Collection

Memory Management	Databases
Object Pointers	Records in a Table
Traversal by Garbage Collector	Analytics / Scan
Changes in Object Pointers	Updates

Yuasa's algorithm:
Snapshot-at-the-beginning algorithm

Air Traffic Control

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	FG752	835.87
3	AA758	190.45
4	TT995	238.60
5	DL992	367.21
6	KA221	1123.56
7	KA802	2192.31
8	AA321	194.10
9	DL293	2490.50

Analytics **A** finds the closest flight to JFK

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	0
1	DL635	103.20	0
2	FG752	835.87	0
3	AA758	190.45	0
4	TT995	238.60	0
5	DL992	367.21	0
6	KA221	1123.56	0
7	KA802	2192.31	0
8	AA321	194.10	0
9	DL293	2490.50	0

Remset

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	0
1	DL635	103.20	0
2	FG752	835.87	0
3	AA758	190.45	0
4	TT995	238.60	0
5	DL992	367.21	0
6	KA221	1123.56	0
7	KA802	2192.31	0
8	AA321	194.10	0
9	DL293	2490.50	0

stable version of the table for analytics A

Closest flight ➔ **DL635**

Remset

Air Traffic Control

A

ID	Flight	Distance	
0	AA123	234.00	0
1	DL635	103.20	0
2	FG752	835.87	0
3	AA758	190.45	0
4	TT995	238.60	0
5	DL992	367.21	0
6	KA221	1123.56	0
7	KA802	2192.31	0
8	AA321	194.10	0
9	DL293	2490.50	0

Closest flight ➔ **DL635**

Closest flight ➔ Null

Remset

Air Traffic Control

A

ID	Flight	Distance	
0	AA123	234.00	**1**
1	DL635	103.20	0
2	FG752	835.87	0
3	AA758	190.45	0
4	TT995	238.60	0
5	DL992	367.21	0
6	KA221	1123.56	0
7	KA802	2192.31	0
8	AA321	194.10	0
9	DL293	2490.50	0

Closest flight ➔ **DL635**

Closest flight ➔ **AA123**

234.00 miles

Remset

Air Traffic Control

A

ID	Flight	Distance	
0	AA123	234.00	**1**
1	DL635	103.20	0
2	FG752	835.87	0
3	AA758	190.45	0
4	TT995	238.60	0
5	DL992	367.21	0
6	KA221	1123.56	0
7	KA802	2192.31	0
8	AA321	194.10	0
9	DL293	2490.50	0

Closest flight ➔ **DL635**

Closest flight ➔ **AA123**

234.00 miles

Remset

Air Traffic Control

A

ID	Flight	Distance	
0	AA123	234.00	**1**
1	DL635	103.20	**1**
2	FG752	835.87	0
3	AA758	190.45	0
4	TT995	238.60	0
5	DL992	367.21	0
6	KA221	1123.56	0
7	KA802	2192.31	0
8	AA321	194.10	0
9	DL293	2490.50	0

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

Remset

Air Traffic Control

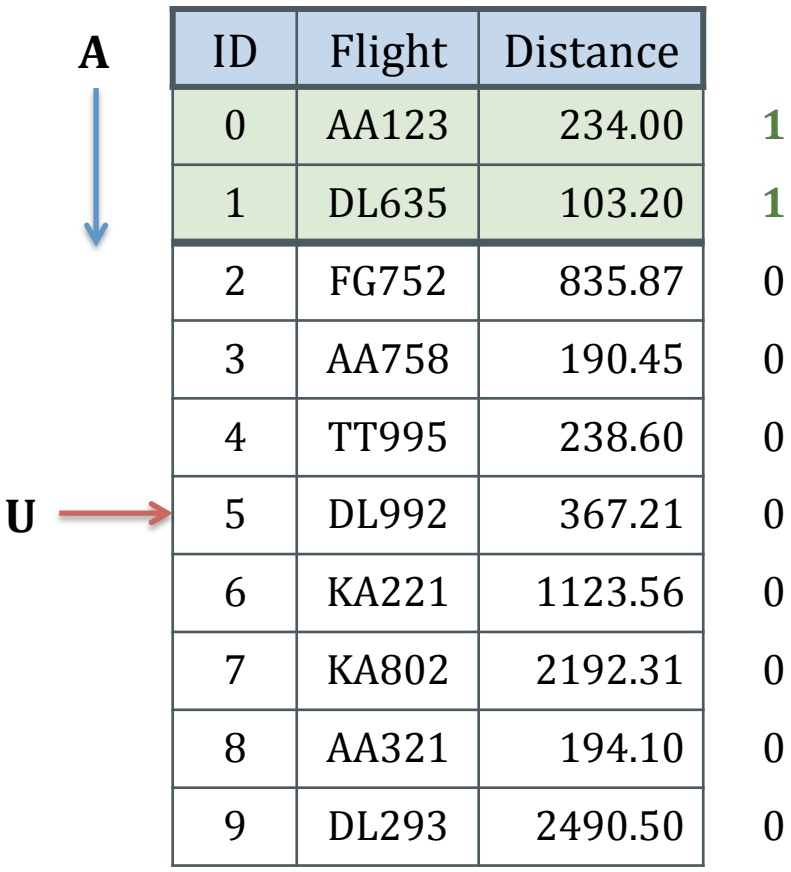

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

Remset

Air Traffic Control

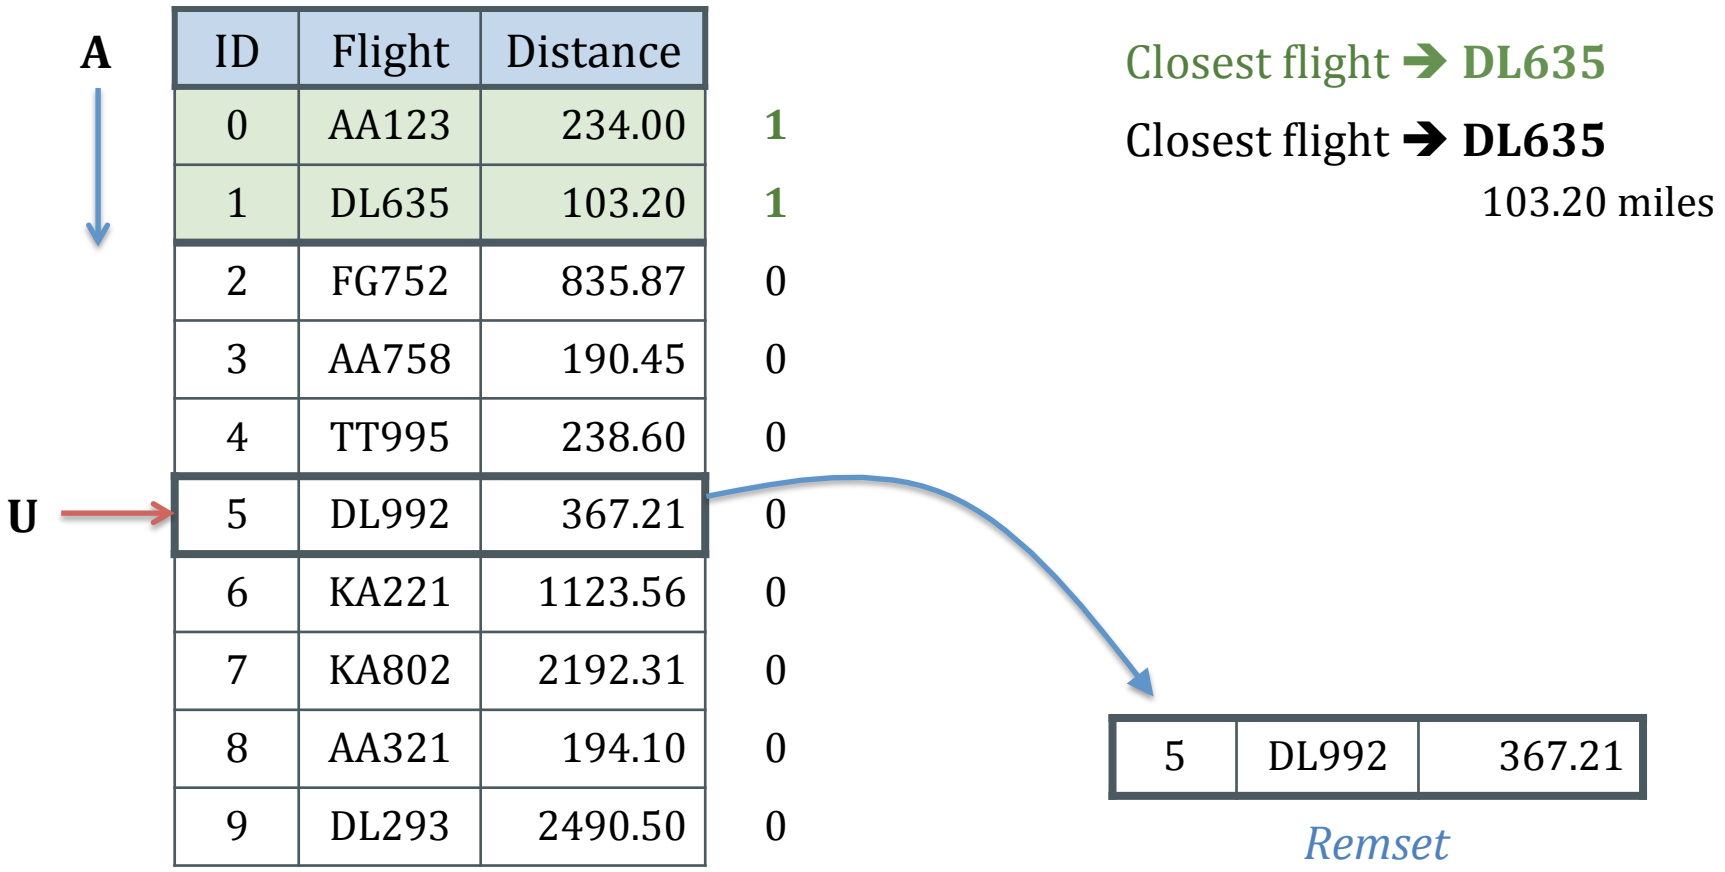

Air Traffic Control

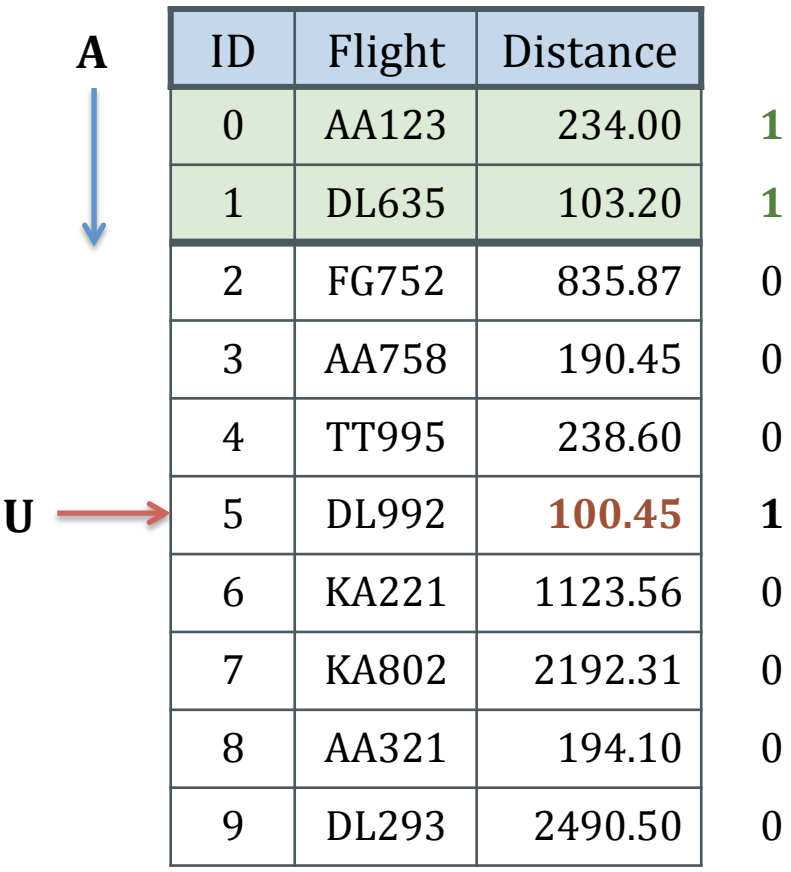

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21

Remset

Air Traffic Control

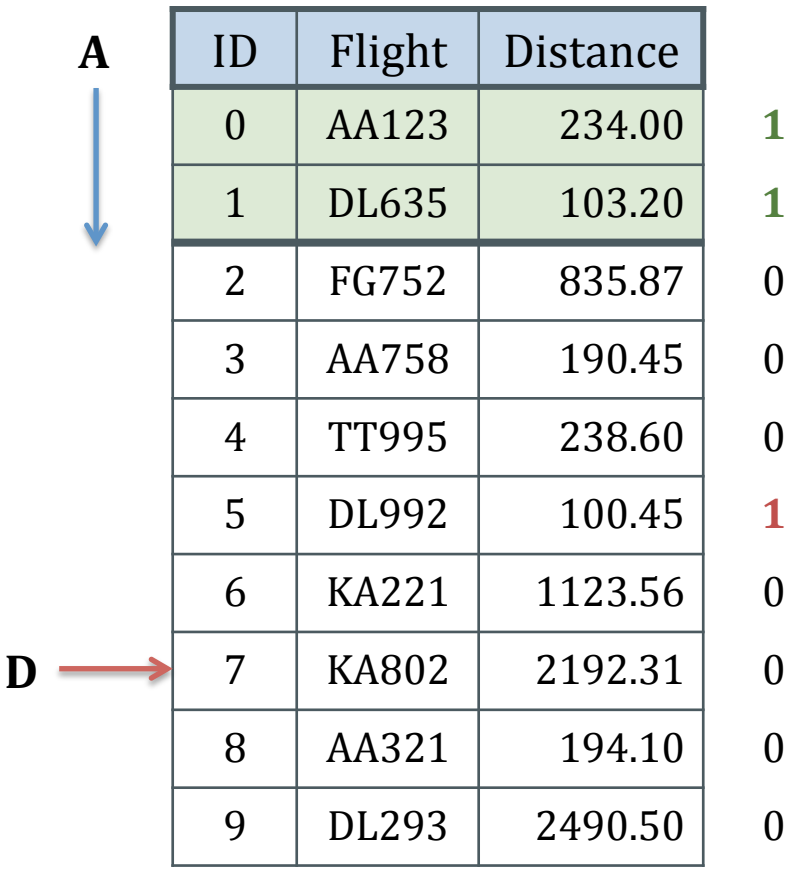

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	103.20	1
2	FG752	835.87	0
3	AA758	190.45	0
4	TT995	238.60	0
5	DL992	100.45	1
6	KA221	1123.56	0
7	KA802	2192.31	0
8	AA321	194.10	0
9	DL293	2490.50	0

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21

Remset

Air Traffic Control

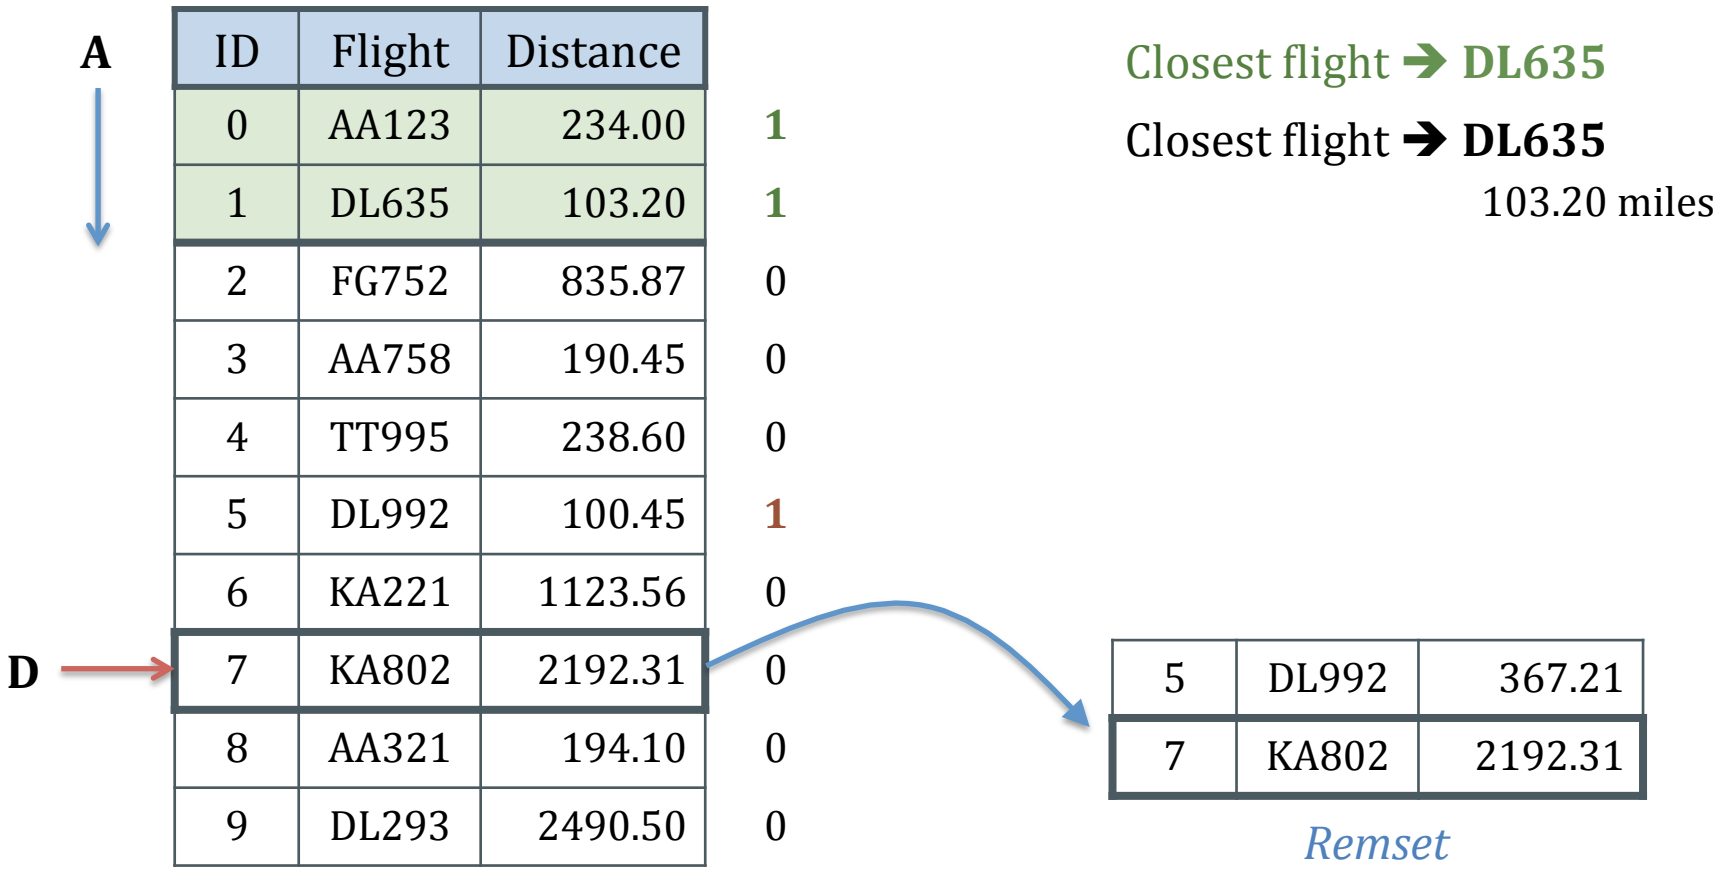

Air Traffic Control

A

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	103.20	1
2	FG752	835.87	0
3	AA758	190.45	0
4	TT995	238.60	0
5	DL992	100.45	1
6	KA221	1123.56	0
8	AA321	194.10	0
9	DL293	2490.50	0

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21
7	KA802	2192.31

Remset

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	103.20	1
2	FG752	835.87	0
3	AA758	190.45	0
4	TT995	238.60	0
5	DL992	100.45	1
6	KA221	1123.56	0
8	AA321	194.10	0
9	DL293	2490.50	0
10	AA555	3489.66	1

I →

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21
7	KA802	2192.31

Remset

Air Traffic Control

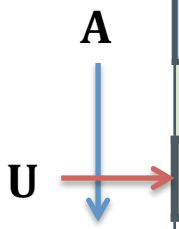

	ID	Flight	Distance	
	0	AA123	234.00	1
U →	1	DL635	103.20	1
	2	FG752	835.87	0
	3	AA758	190.45	0
	4	TT995	238.60	0
	5	DL992	100.45	1
	6	KA221	1123.56	0
	8	AA321	194.10	0
	9	DL293	2490.50	0
U →	10	AA555	3489.66	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21
7	KA802	2192.31

Remset

Air Traffic Control

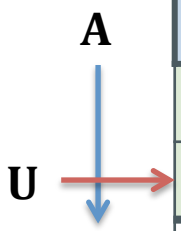

	ID	Flight	Distance	
	0	AA123	234.00	1
U →	1	DL635	**90.34**	1
	2	FG752	835.87	0
	3	AA758	190.45	0
	4	TT995	238.60	0
	5	DL992	100.45	1
	6	KA221	1123.56	0
	8	AA321	194.10	0
	9	DL293	2490.50	0
U →	10	AA555	**3290.21**	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21
7	KA802	2192.31

Remset

Air Traffic Control

A

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	90.34	1
2	FG752	835.87	1
3	AA758	190.45	1
4	TT995	238.60	1
5	DL992	100.45	1
6	KA221	1123.56	0
8	AA321	194.10	0
9	DL293	2490.50	0
10	AA555	3290.21	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21
7	KA802	2192.31

Remset

Air Traffic Control

A

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	90.34	1
2	FG752	835.87	1
3	AA758	190.45	1
4	TT995	238.60	1
5	DL992	100.45	1
6	KA221	1123.56	0
8	AA321	194.10	0
9	DL293	2490.50	0
10	AA555	3290.21	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21
7	KA802	2192.31

Remset

Air Traffic Control

A

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	90.34	1
2	FG752	835.87	1
3	AA758	190.45	1
4	TT995	238.60	1
5	DL992	100.45	1
6	KA221	1123.56	0
8	AA321	194.10	0
9	DL293	2490.50	0
10	AA555	3290.21	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21
7	KA802	2192.31

Remset

Air Traffic Control

A

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	90.34	1
2	FG752	835.87	1
3	AA758	190.45	1
4	TT995	238.60	1
5	DL992	100.45	1
6	KA221	1123.56	1
8	AA321	194.10	1
9	DL293	2490.50	1
10	AA555	3290.21	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21
7	KA802	2192.31

Remset

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	90.34	1
2	FG752	835.87	1
3	AA758	190.45	1
4	TT995	238.60	1
5	DL992	100.45	1
6	KA221	1123.56	1
8	AA321	194.10	1
9	DL293	2490.50	1
10	AA555	3290.21	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

5	DL992	367.21
7	KA802	2192.31

Remset

record removed

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	90.34	1
2	FG752	835.87	1
3	AA758	190.45	1
4	TT995	238.60	1
5	DL992	100.45	1
6	KA221	1123.56	1
8	AA321	194.10	1
9	DL293	2490.50	1
10	AA555	3290.21	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

A

5	DL992	367.21
7	KA802	2192.31

Remset

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	90.34	1
2	FG752	835.87	1
3	AA758	190.45	1
4	TT995	238.60	1
5	DL992	100.45	1
6	KA221	1123.56	1
8	AA321	194.10	1
9	DL293	2490.50	1
10	AA555	3290.21	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

A

7	KA802	2192.31

Remset

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	90.34	1
2	FG752	835.87	1
3	AA758	190.45	1
4	TT995	238.60	1
5	DL992	100.45	1
6	KA221	1123.56	1
8	AA321	194.10	1
9	DL293	2490.50	1
10	AA555	3290.21	1

Closest flight ➔ **DL635**

Closest flight ➔ **DL635**

103.20 miles

Remset

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	**1**
1	DL635	90.34	**1**
2	FG752	835.87	**1**
3	AA758	190.45	**1**
4	TT995	238.60	**1**
5	DL992	100.45	**1**
6	KA221	1123.56	**1**
8	AA321	194.10	**1**
9	DL293	2490.50	**1**
10	AA555	3290.21	**1**

Remset

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	**1**
1	DL635	90.34	**1**
2	FG752	835.87	**1**
3	AA758	190.45	**1**
4	TT995	238.60	**1**
5	DL992	100.45	**1**
6	KA221	1123.56	**1**
8	AA321	194.10	**1**
9	DL293	2490.50	**1**
10	AA555	3290.21	**1**

Before **A** scans table:

0 ➔ stable

1 ➔ read / changed

Remset

Air Traffic Control

ID	Flight	Distance	
0	AA123	234.00	**1**
1	DL635	90.34	**1**
2	FG752	835.87	**1**
3	AA758	190.45	**1**
4	TT995	238.60	**1**
5	DL992	100.45	**1**
6	KA221	1123.56	**1**
8	AA321	194.10	**1**
9	DL293	2490.50	**1**
10	AA555	3290.21	**1**

After **A** scans table:

0 ➔ read / changed

1 ➔ stable

Remset

Virtual Lightweight Snapshots (***VLS***)

ID	Flight	Distance
0	AA123	234.00
1	DL635	90.34
2	FG752	835.87
3	AA758	190.45
4	TT995	238.60
5	DL992	100.45
6	KA221	1123.56
8	AA321	194.10
9	DL293	2490.50
10	AA555	3290.21

Table

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	FG752	835.87
3	AA758	190.45
4	TT995	238.60
5	DL992	367.21
6	KA221	1123.56
7	KA802	2192.31
8	AA321	194.10
9	DL293	2490.50

Virtual Lightweight Snapshot for **A**

5	DL992	367.21
7	KA802	2192.31

Remset

Problem solved ?

Not quite yet…

Challenge:
Support for Multiple Scans / Analytics

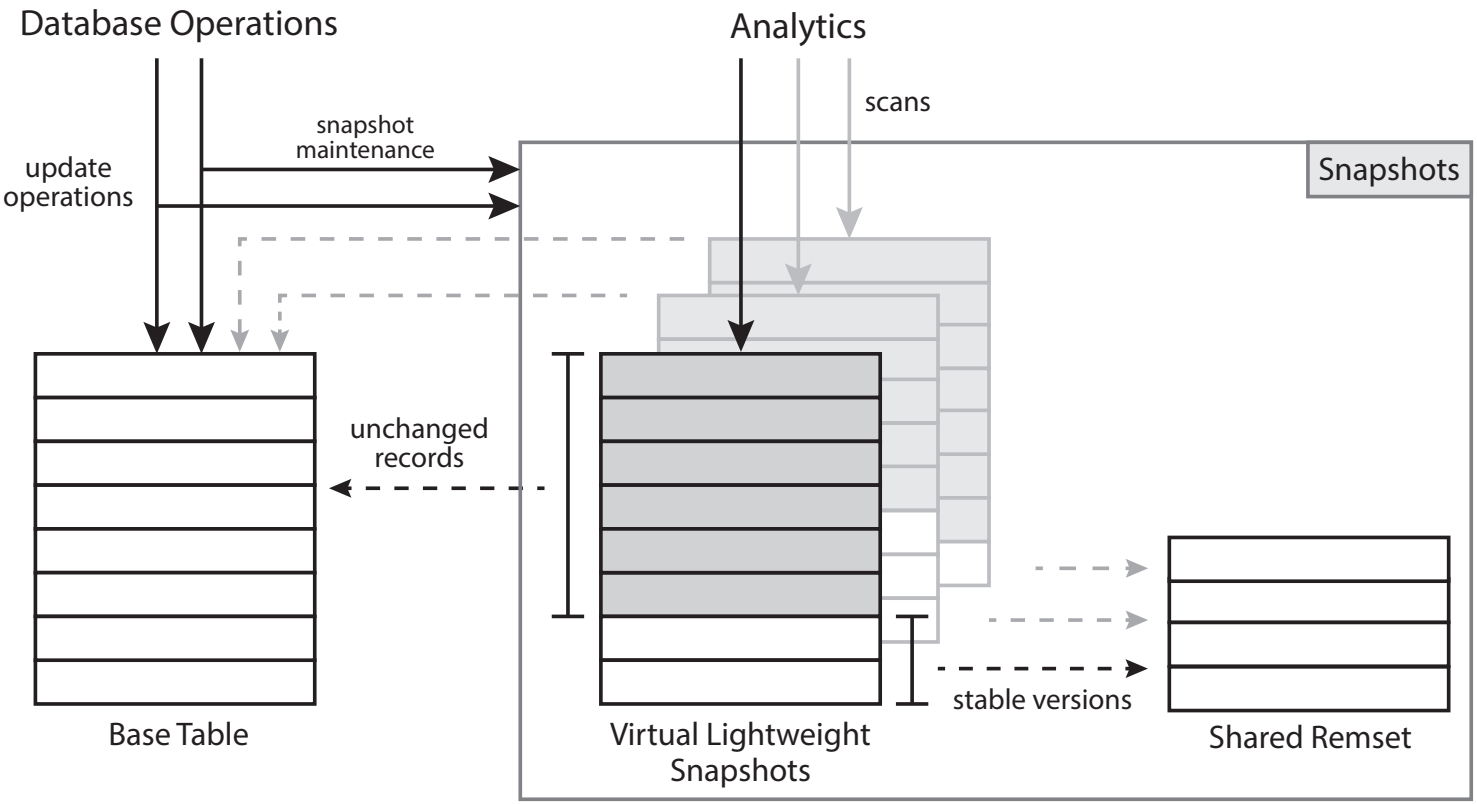
Database Operations
Analytics
scans
snapshot
maintenance
update
operations
Snapshots
unchanged
records
stable versions
Base Table
Virtual Lightweight
Snapshots
Shared Remset

Support for Multiple Scans

N = 3

Active Mask

1	1	1

Stable Mask

0	0	0

ID	Flight	Distance	Status Mask		
0	AA123	234.00	0	0	0
1	DL635	90.34	0	0	0
2	FG752	835.87	0	0	0
3	AA758	190.45	0	0	0
4	TT995	238.60	0	0	0
5	DL992	100.45	0	0	0
6	KA221	1123.56	0	0	0
8	AA321	194.10	0	0	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 1 1

ID	Flight	Distance	Status Mask
0	AA123	234.00	0 0 0
1	DL635	90.34	0 0 0
2	FG752	835.87	0 0 0
3	AA758	190.45	0 0 0
4	TT995	238.60	0 0 0
5	DL992	100.45	0 0 0
6	KA221	1123.56	0 0 0
8	AA321	194.10	0 0 0

Stable Mask

0 0 0

Remset Status Mask

8	AA321	194.10	1 1 0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 1 1

ID	Flight	Distance	Status Mask		
0	AA123	234.00	0	0	0
1	DL635	90.34	0	0	0
2	FG752	835.87	0	0	0
3	AA758	190.45	0	0	0
4	TT995	238.60	0	0	0
5	DL992	100.45	0	0	0
6	KA221	1123.56	0	0	0
8	AA321	194.10	0	0	0

Stable Mask

0 0 0

8	AA321	194.10	1	1	1

Remset

Support for Multiple Scans

N = 3

Active Mask

1 1 1

ID	Flight	Distance	Status Mask
0	AA123	234.00	0 0 0
1	DL635	90.34	0 0 0
2	FG752	835.87	0 0 0
3	AA758	190.45	0 0 0
4	TT995	238.60	0 0 0
5	DL992	100.45	0 0 0
6	KA221	1123.56	0 0 0
8	AA321	194.10	0 0 0

Stable Mask

0 0 0

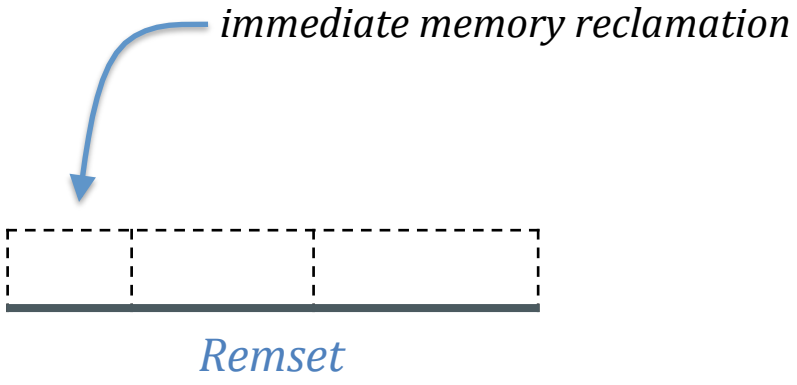

Bit Operations

Remset Status Mask ← (Stable Mask XOR Status Mask) OR Active Mask
UPDATE / DELETE

Status Mask ← NOT (Stable Mask XOR Active Mask)
UPDATE / INSERT

Status Mask ← NOT (NOT (Status Mask) XOR Scan Mask)
SCAN

Status ← (Stable Mask XOR Status Mask) AND Scan Mask
SCAN

Stable Mask ← NOT (NOT (Stable Mask) XOR Scan Mask)
SCAN

Active Mask ← Active Mask OR Scan Mask
SCAN

Challenge:
Support for Index Scans

1) Spurious Records

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	DL992	367.21
3	KA221	1123.56
4	KA802	2192.31
5	AA321	194.10
6	DL293	2490.50

*Flights that are at most **400** miles away?*

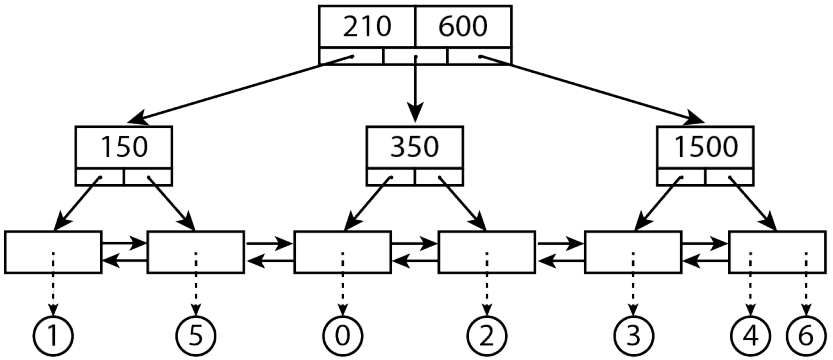

B-Tree Index

1) Spurious Records

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	DL992	367.21
3	KA221	1123.56
4	KA802	2192.31
5	AA321	194.10
6	DL293	2490.50

U

6	DL293	2490.50

spurious record

Remset

*Flights that are at most **400** miles away?*

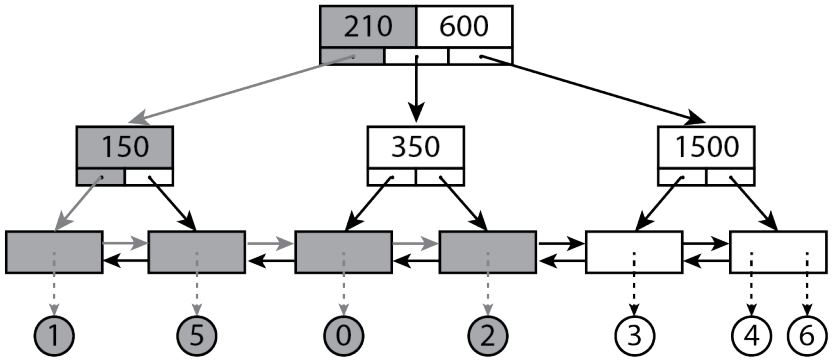

B-Tree Index

Solution:

record set = {1,5,0,2}

2) Consistency of Bit Masks

ID	Flight	Distance	
0	AA123	234.00	1
1	DL635	103.20	1
2	DL992	367.21	1
3	KA221	1123.56	**0**
4	KA802	2192.31	**0**
5	AA321	194.10	1
6	DL293	2490.50	**0**

*Flights that are at most **400** miles away?*

Solution:

*Background thread to flip all bits **after** scan finishes*

3) Physical Structure of Index

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	DL992	367.21
3	KA221	1123.56
4	KA802	2192.31
5	AA321	194.10
6	DL293	2490.50

U →

*Flights that are at most **400** miles away?*

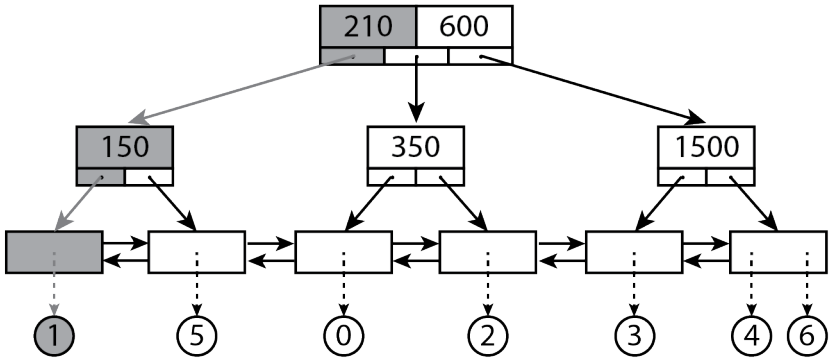

B-Tree Index

3) Physical Structure of Index

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	DL992	**100.45**
3	KA221	1123.56
4	KA802	2192.31
5	AA321	194.10
6	DL293	2490.50

U →

*Flights that are at most **400** miles away?*

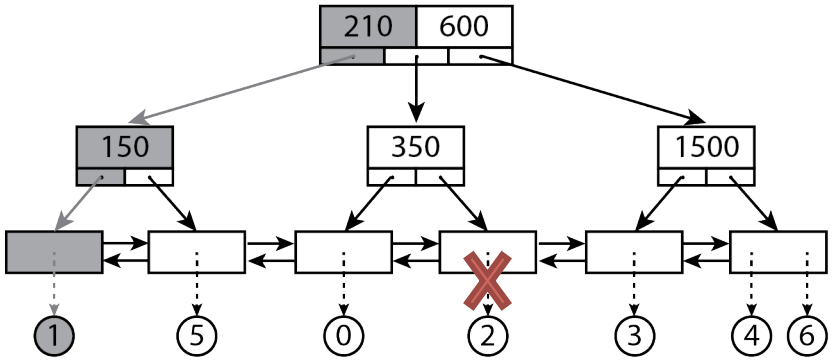

B-Tree Index

3) Physical Structure of Index

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	DL992	**100.45**
3	KA221	1123.56
4	KA802	2192.31
5	AA321	194.10
6	DL293	2490.50

U →

*Flights that are at most **400** miles away?*

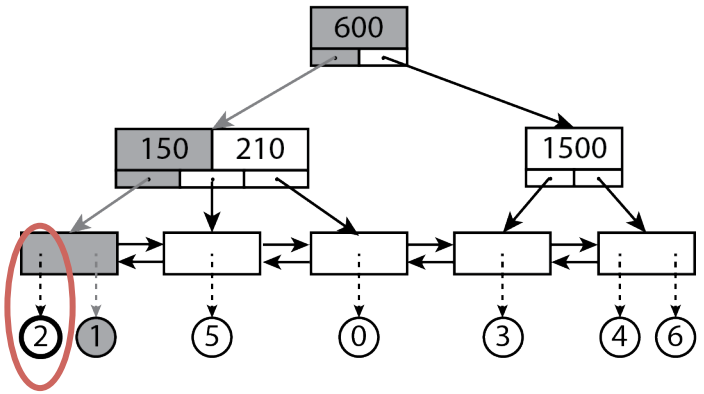

B-Tree Index

3) Physical Structure of Index

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	**DL992**	**100.45**
3	KA221	1123.56
4	KA802	2192.31
5	AA321	194.10
6	DL293	2490.50

*Flights that are at most **400** miles away?*

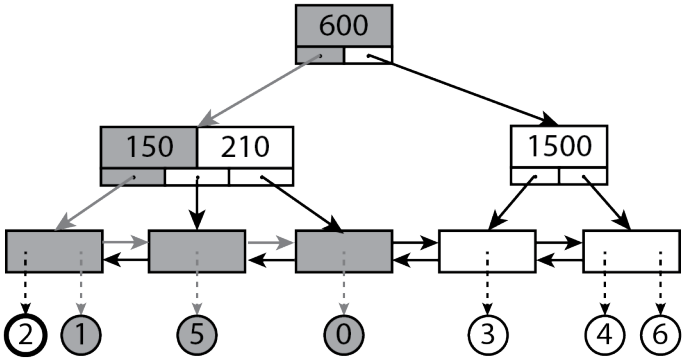

B-Tree Index

record set = {1,5,0}

Solution:

- *Removal of record pointers is **postponed***
- *Masks for record pointers*

Bit Operations

Index Status Mask ← NOT (Index Active Mask)
INSERT

Index Status Mask ← Index Status Mask OR Index Active Mask
DELETE

Index Active Mask ← Index Active Mask OR Scan Mask
INDEX SCAN

Index Status Mask ← Index Status Mask OR Scan Mask
MEMORY RECLAMATION

Please check the details in the paper!

MongoDB-VLS:

https://github.com/ViDA-NYU/mongodb-vls

- Built-in implementation (MongoDB 2.5)
- Modifications to
  - *insert*, *update*, and *delete* operations
  - *scan* operations
  - *B-Tree index*
- Masks and *remsets* maintained by the system

Experimental Evaluation

vs. MongoDB-VLS

Goals:
1) Analyze the impact of VLS on update throughput, latency, and analytics completion time
2) Examine the costs for maintaining the snapshots

Benchmark: Yahoo! Cloud System Benchmark (YCSB) [1]

[1] *B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears, "**Benchmarking Cloud Serving Systems with YCSB**," in SoCC'10, 2010, pp. 143–154.*

Overview of Results

- Query Execution Time
  - Time increases by at most 10% (full and index scans)
  - In absolute numbers, for many cases, less than 1s
- Update Throughput and Latency
  - Overheads of at most 15%
  - Error bars overlap: overheads are often not statistically significant
- *Remset* Sharing
  - VLS leads to storage savings of up to 60 times
- Size of Bit Masks
  - Setting the size as the register width gives the best performance

VLS leads to small overheads while enabling consistent results !

Please check our full experiments and results in the paper!

Summary

- Virtual Lightweight Snapshots (VLS) is a snapshotting technique for NoSQL stores
  - It brings consistency for *analytics* in the presence of concurrent *updates*
  - It depends on neither a native database versioning nor a transaction manager
- VLS brings consistency to analytics at a *low cost*
  - Bit mask operations are very efficient
  - Overheads in query execution time, update latency, and update throughput are negligible
  - Remset sharing approach is key to memory efficiency
- Future work
  - MongoDB-VLS in ODM Insights
  - Comparison with WiredTiger (MongoDB 3.0+)
  - VLS for distributed data
  - Compression of remset records

Acknowledgements

- Funding
  - Google Faculty Award
  - IBM Faculty Award
  - Moore-Sloan Data Science Environment at NYU
  - DARPA
  - National Science Foundation
- Insightful comments from reviewers

Thanks! Questions?

Fernando Chirigati
fchirigati@nyu.edu

Code available here: https://github.com/ViDA-NYU/mongodb-vls

Additional Slides

Example: Air Traffic Control

Redirecting Flights

when an airport closed event occurs, called 'airport closed' ← EVENT

```
definitions
  set 'closest flights' to all flights in the inbound flights of 'the airport'
    where the location of each flight is within 300 miles range of the
    location of 'the airport';

if
  the propensity that 'the airport' will stay closed at least 1 hour,
  given the closing reason of 'airport closed' and the closing time of
  'airport closed', is more than 0.75

then
  for each flight in 'closest flights':
  - set 'the new arrival airport' to the closest airport from the location of
    this flight;
  - set airport of this flight to 'the new arrival airport';
```

→ RULE

Example: Air Traffic Control

Redirecting Flights

ANALYTICS TO FIND CLOSEST FLIGHTS

```
when an airport closed event occurs, called 'airport closed'
definitions
  set 'closest flights' to all flights in the inbound flights of 'the airport'
    where the location of each flight is within 300 miles range of the
    location of 'the airport';
if
  the propensity that 'the airport' will stay closed at least 1 hour,
  given the closing reason of 'airport closed' and the closing time of
  'airport closed', is more than 0.75
then
  for each flight in 'closest flights':
  - set 'the new arrival airport' to the closest airport from the location of
    this flight;
  - set airport of this flight to 'the new arrival airport';
```

UPDATE

Example: Air Traffic Control

Redirecting Flights

```
when an airport closed event occurs, called 'airport closed'

definitions
  set 'closest flights' to all flights in the inbound flights of 'the airport'
    where the location of each flight is within 300 miles range of the
    location of 'the airport';

if
  the propensity that 'the airport' will stay closed at least 1 hour,
  given the closing reason of 'airport closed' and the closing time of
  'airport closed', is more than 0.75

then
  for each flight in 'closest flights':
  - set 'the new arrival airport' to the close
    this flight;
  - set airport of this flight to 'the new arrival airport';
```

ANALYTICS THAT INVOKES A PREDICTIVE SCORING MODEL

Updates

fast and numerous
(fast data)

X

Analytics

long and expensive
(big data)

Different systems: OLTP and OLAP !

No data freshness for analytics

Same system !

Impact on update throughput and latency

Same system and lower level of isolation !

Inconsistent results

throughput / latency

consistent analytics

Read Uncommitted

Closest flight?

t_0

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	FG752	835.87
3	AA758	190.45
4	TT995	238.60
5	DL992	367.21
6	KA221	1123.56
7	KA802	2192.31
8	AA321	194.10
9	DL293	2490.50

t_1

ID	Flight	Distance
0	AA123	234.00
1	DL635	**90.34**
2	FG752	835.87
3	AA758	190.45
4	TT995	238.60
5	DL992	**100.45**
6	KA221	1123.56
8	AA321	194.10
9	DL293	2490.50
10	**AA555**	**3290.21**

Read Uncommitted

Closest flight?

t_0

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	FG752	835.87
3	AA758	190.45
4	TT995	238.60
5	DL992	367.21
6	KA221	1123.56
7	KA802	2192.31
8	AA321	194.10
9	DL293	2490.50

Read Uncommitted

Closest flight?

t_0

ID	Flight	Distance
0	AA123	234.00
1	DL635	103.20
2	FG752	835.87
3	AA758	190.45
4	TT995	238.60
5	DL992	367.21
6	KA221	1123.56
7	KA802	2192.31
8	AA321	194.10
9	DL293	2490.50

t_1

ID	Flight	Distance
0	AA123	234.00
1	DL635	**90.34**
2	FG752	835.87
3	AA758	190.45
4	TT995	238.60
5	DL992	**100.45**
6	KA221	1123.56
8	AA321	194.10
9	DL293	2490.50
10	**AA555**	**3290.21**

Closest flight = DL992

Concurrent Garbage Collection

Yuasa's algorithm
(snapshot-at-the-beginning)

- A queue is initialized
- Every object receives a bit 0
- If an application modifies an object pointer
  - Bit is flipped to 1
  - Current value is placed in queue
- GC processes an object if bit is 0, and sets the bit to 1
- GC skips an object if bit is 1
- After looking at the graph, GC looks at the queue

Support for Multiple Scans

- Extending the bit representation
  - *N*-bit vectors for table records
  - *N*-bit vectors for *remset* records

- Need for additional bookkeeping
  - Allocating bits to incoming scans
  - Maintaining the meaning of the bits (stable *vs.* read/changed)
  - Performing memory reclamation of *remset* records

Support for Multiple Scans

N = 3

```
0: scan
1: no scan
```

Active Mask

1	1	1

ID	Flight	Distance	Status Mask
0	AA123	234.00	0 0 0
1	DL635	90.34	0 0 0
2	FG752	835.87	0 0 0
3	AA758	190.45	0 0 0
4	TT995	238.60	0 0 0
5	DL992	100.45	0 0 0
6	KA221	1123.56	0 0 0
8	AA321	194.10	0 0 0

Stable Mask

0	0	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

ID	Flight	Distance			
0	AA123	234.00	0	0	0
1	DL635	90.34	0	0	0
2	FG752	835.87	0	0	0
3	AA758	190.45	0	0	0
4	TT995	238.60	0	0	0
5	DL992	100.45	0	0	0
6	KA221	1123.56	0	0	0
8	AA321	194.10	0	0	0

Stable Mask

0 0 0

Remset

Support for Multiple Scans

N = 3

Active Mask

1	0	0

Stable Mask

0	0	0

	ID	Flight	Distance		A1	A2
	0	AA123	234.00	0	1	1
	1	DL635	90.34	0	1	1
	2	FG752	835.87	0	1	1
	3	AA758	190.45	0	1	1
U	4	TT995	238.60	0	1	0
	5	DL992	100.45	0	1	0
	6	KA221	1123.56	0	1	0
	8	AA321	194.10	0	0	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 0 0

ID	Flight	Distance		A1	A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	238.60	0	1	0
5	DL992	100.45	0	1	0
6	KA221	1123.56	0	1	0
8	AA321	194.10	0	0	0

U

4	TT995	238.60

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 0 0

	ID	Flight	Distance		A1	A2
	0	AA123	234.00	0	1	1
	1	DL635	90.34	0	1	1
	2	FG752	835.87	0	1	1
	3	AA758	190.45	0	1	1
U	4	TT995	238.60	0	1	0
	5	DL992	100.45	0	1	0
	6	KA221	1123.56	0	1	0
	8	AA321	194.10	0	0	0

0: read
1: skip

S[illegible]k

4	TT995	238.60	1 1 0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 0 0

ID	Flight	Distance		A1	A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	0
6	KA221	1123.56	0	1	0
8	AA321	194.10	0	0	0

U

4	TT995	238.60	1	1	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 | 0 0

Stable Mask

0 | 0 0

ID	Flight	Distance		A1	A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	**1**
5	DL992	100.45	0	1	0
6	KA221	1123.56	0	1	0
8	AA321	194.10	0	0	0

U

4	TT995	238.60	1	1	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 | 0 0

Stable Mask

0 | 0 0

ID	Flight	Distance		A1	A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	**1**
5	DL992	100.45	0	1	0
6	KA221	1123.56	0	1	0
8	AA321	194.10	0	0	0

U

4	TT995	238.60	1	1	0
8	AA321	194.10	1	0	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 0 0

ID	Flight	Distance		A1	A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	**1**
5	DL992	100.45	0	1	0
6	KA221	1123.56	0	1	0
8	AA321	**187.32**	0	**1**	**1**

U

4	TT995	238.60	1	1	0
8	AA321	194.10	1	0	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 0 0

ID	Flight	Distance		A1	A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

A1

4	TT995	238.60	1	1	0
8	AA321	194.10	1	0	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 0 0

ID	Flight	Distance		A1	A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

A1

4	TT995	238.60	1	1	0
8	AA321	194.10	1	0	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 0 0

ID	Flight	Distance		A1	A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

A1

4	TT995	238.60	1	1	0
8	AA321	194.10	1	0	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 0 0

ID	Flight	Distance		A1	A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

A1

4	TT995	238.60	1	1	0
8	AA321	194.10	1	**1**	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 1 0

ID	Flight	Distance			A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

A2

4	TT995	238.60	1	1	0
8	AA321	194.10	1	1	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

ID	Flight	Distance			A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

Stable Mask

0 1 0

A2

4	TT995	238.60	1	1	1
8	AA321	194.10	1	1	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

ID	Flight	Distance			A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

Stable Mask

0 1 0

record removed by **A2**

immediate memory reclamation

A2

8	AA321	194.10	1	1	0

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

Stable Mask

0 1 0

ID	Flight	Distance			A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

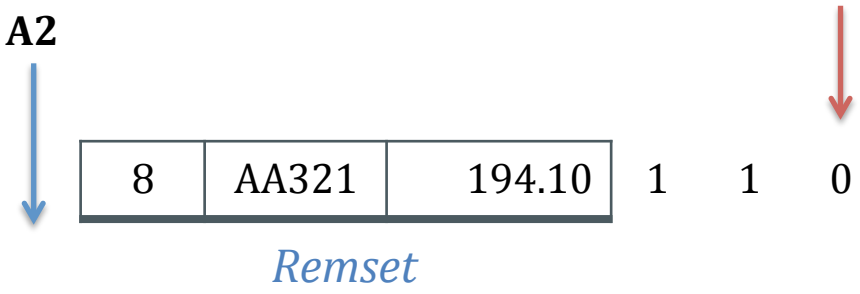

Support for Multiple Scans

N = 3

Active Mask

1 0 0

ID	Flight	Distance			A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

Stable Mask

0 1 0

A2

8	AA321	194.10	1	1	1

Remset

Support for Multiple Scans

N = 3

Active Mask

1 0 0

ID	Flight	Distance			A2
0	AA123	234.00	0	1	1
1	DL635	90.34	0	1	1
2	FG752	835.87	0	1	1
3	AA758	190.45	0	1	1
4	TT995	189.89	0	1	1
5	DL992	100.45	0	1	1
6	KA221	1123.56	0	1	1
8	AA321	187.32	0	1	1

Stable Mask

0 1 0

A2

Remset

Query Execution Time

- Without Concurrent Updates

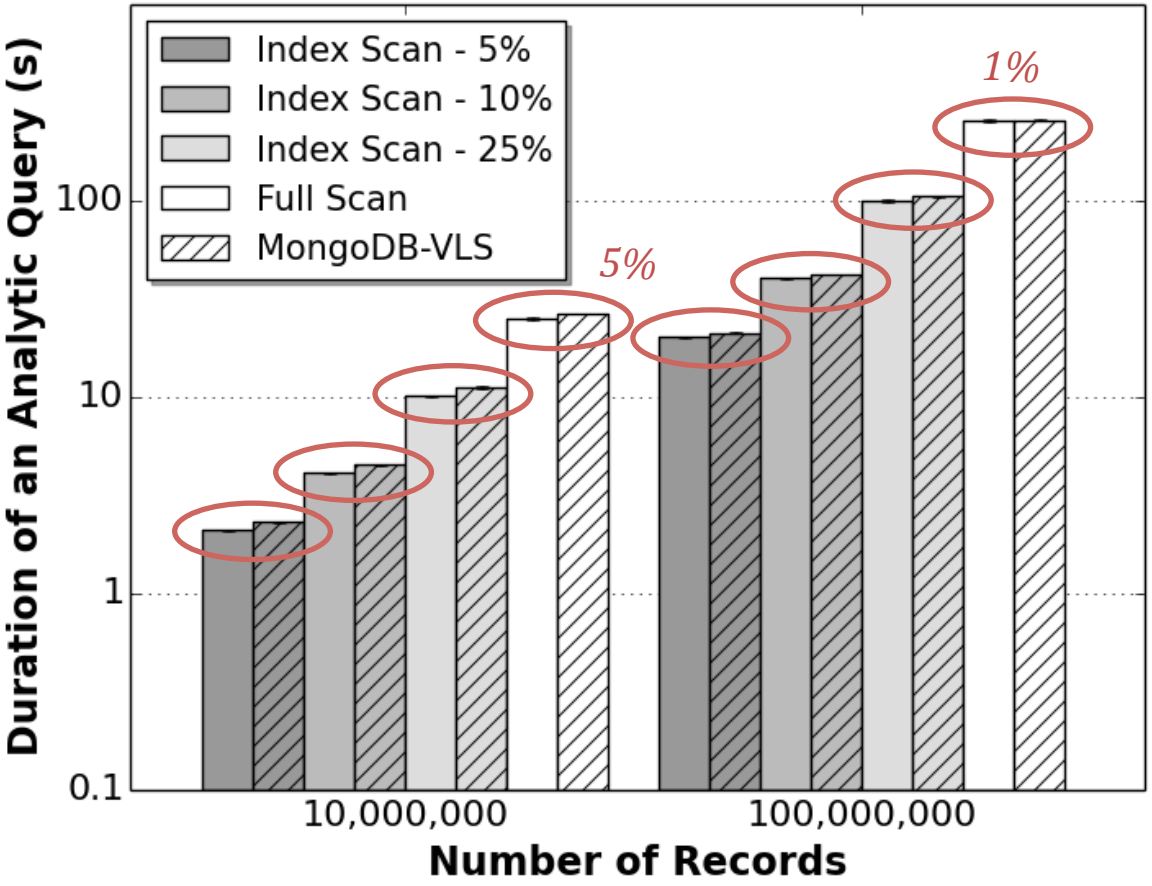

Query Execution Time

- With Concurrent Updates

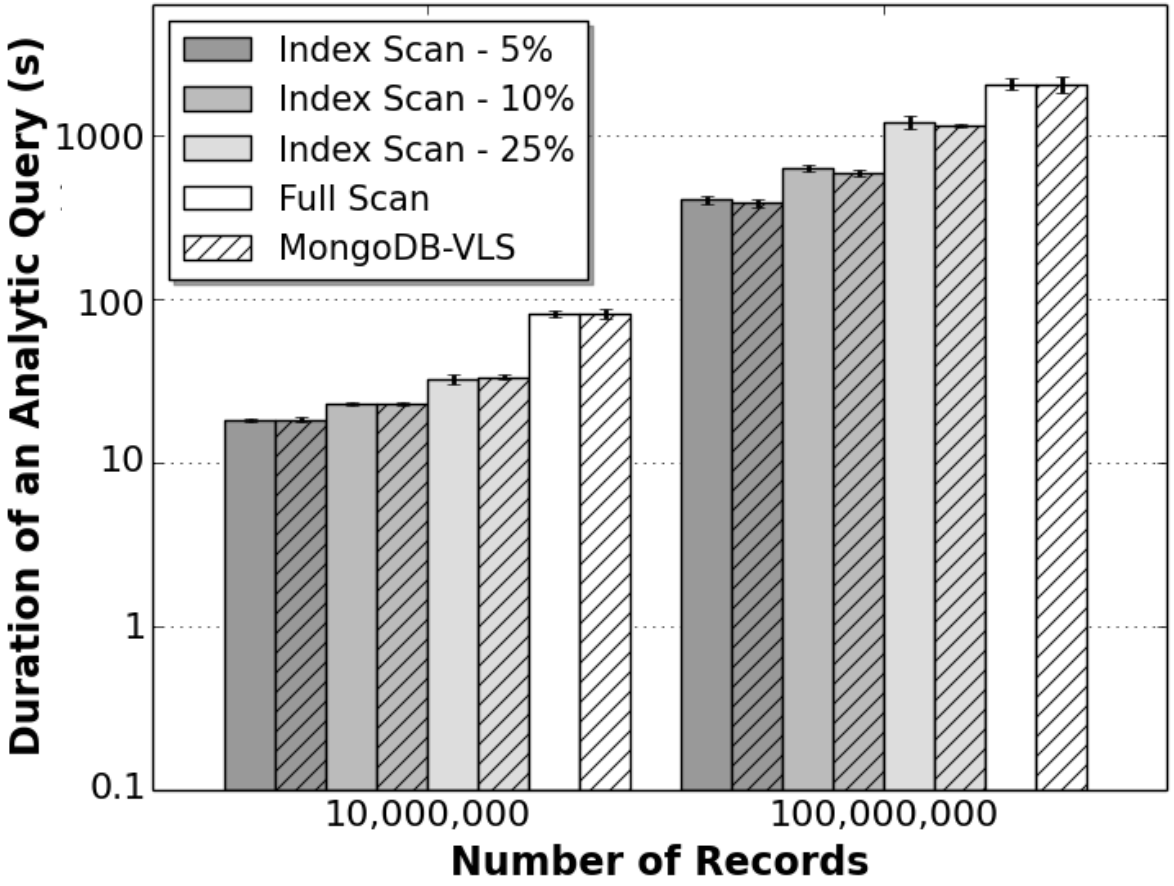

Update Throughput and Latency

- Full Scans

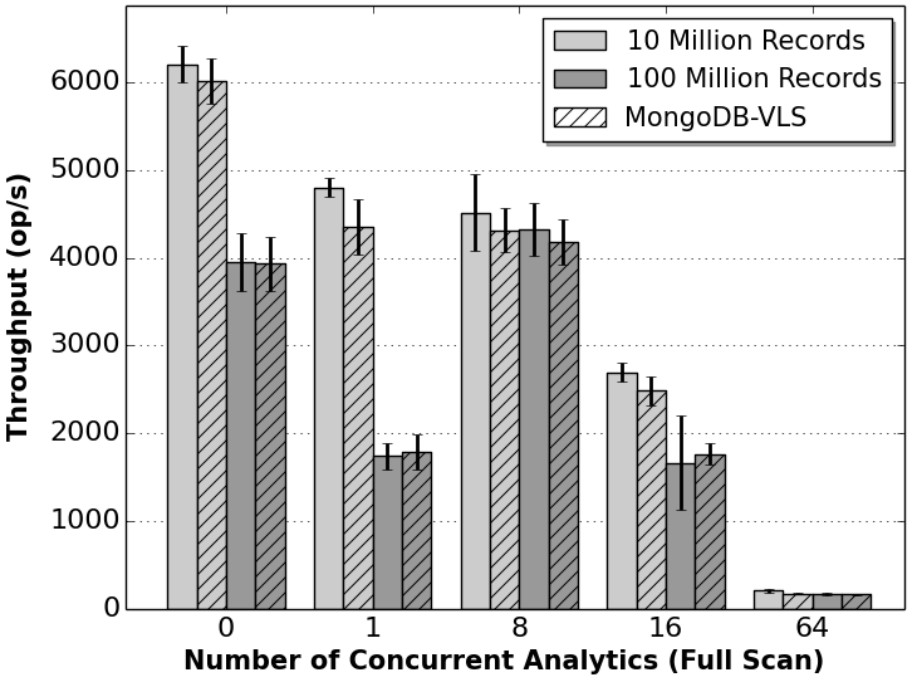

3% - 16%

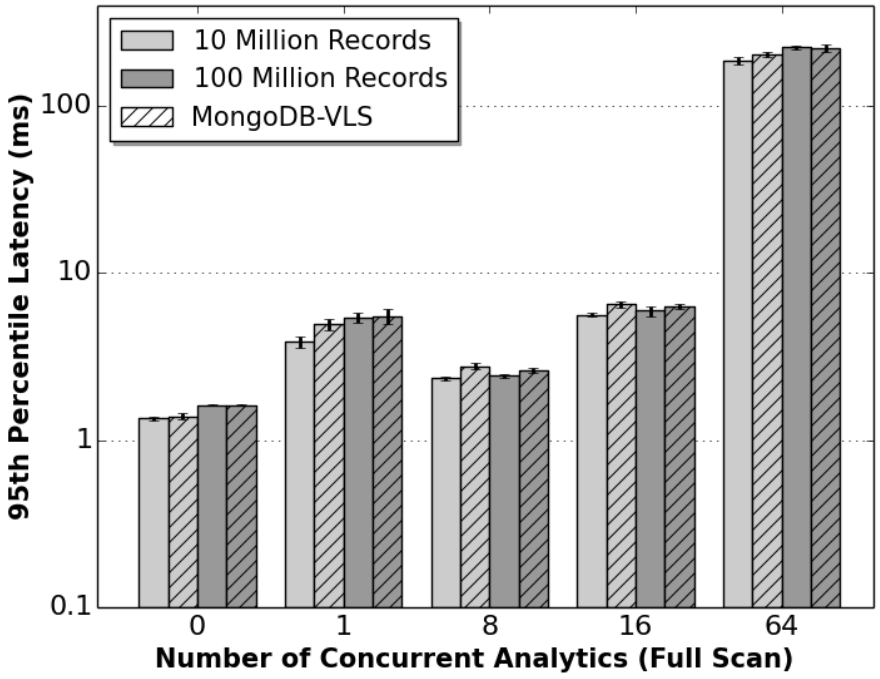

2% - 15%

Update Throughput and Latency

- Index Scans

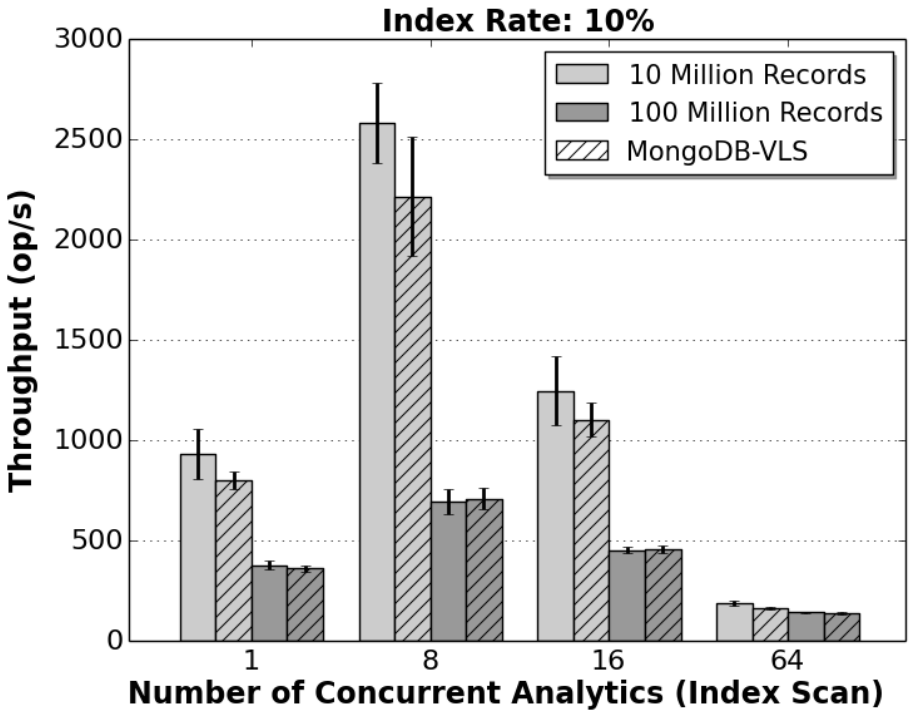

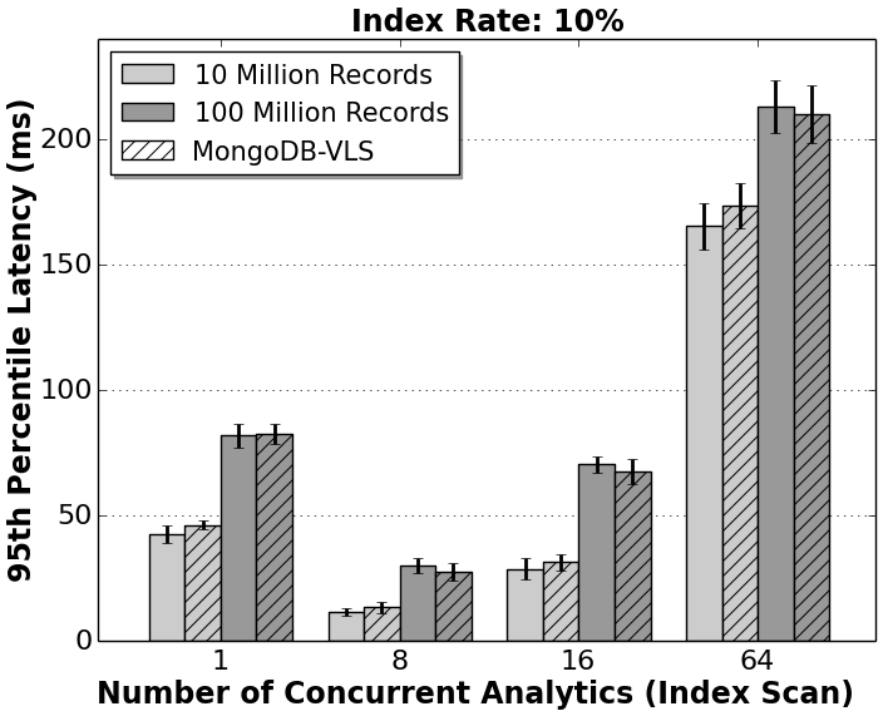

0.70% - 18%

0% - 7%

Remset Sharing