


# Column-Stores vs. Row-Stores: How Different Are They Really?

SK Lee & Lucas Yoon





# Row-Stores


## Customer


ID	Name	Country	City
28495072	Jack Hamilton	USA	Seattle
78239842	Sarah Wilson	Finland	Helsinki
19389562	Jason Huang	China	Shanghai

```
INSERT INTO CUSTOMER(ID, Name,  
Country, City)  
VALUES(37583719, 'Sam Kim',  
'South Korea', 'Seoul')
```

## Customer

ID	Name	Country	City
28495072	Jack Hamilton	USA	Seattle
78239842	Sarah Wilson	Finland	Helsinki
19389562	Jason Huang	China	Shanghai
37583719	Sam Kim	South Korea	Seoul





# Column-Stores

## Customer


ID	Name	Country	City
28495072	Jack Hamilton	USA	Seattle
78239842	Sarah Wilson	Finland	Helsinki
19389562	Jason Huang	China	Shanghai

## Customer

ID	Name
28495072	Jack Hamilton
78239842	Sarah Wilson
19389562	Jason Huang

ID	Country
28495072	USA
78239842	Finland
19389562	China

ID	City
28495072	Seattle
78239842	Helsinki
19389562	Shanghai



# Index

1.

Introduction

2.

Experimental Setup

3.

Row-Orientation Execution

4.

Column-Orientation Execution


5.

Results

6.




What did We Learn?





# What is the problem this paper is solving?

1. Misconception
2. Lack of Systematic Comparison
3. Column Store Optimizations
4. Row Store Optimizations
5. Architectural Advantages




# Why is it Challenging?



## Challenge 1

- Architectural Differences

## Challenge 2

- Balancing Performance





# Why is it important?




**Informed System Design**



**Hybrid Solutions**



**Future Development**




# Solution Description



**Compare Performance**



**Dissect Optimizations**

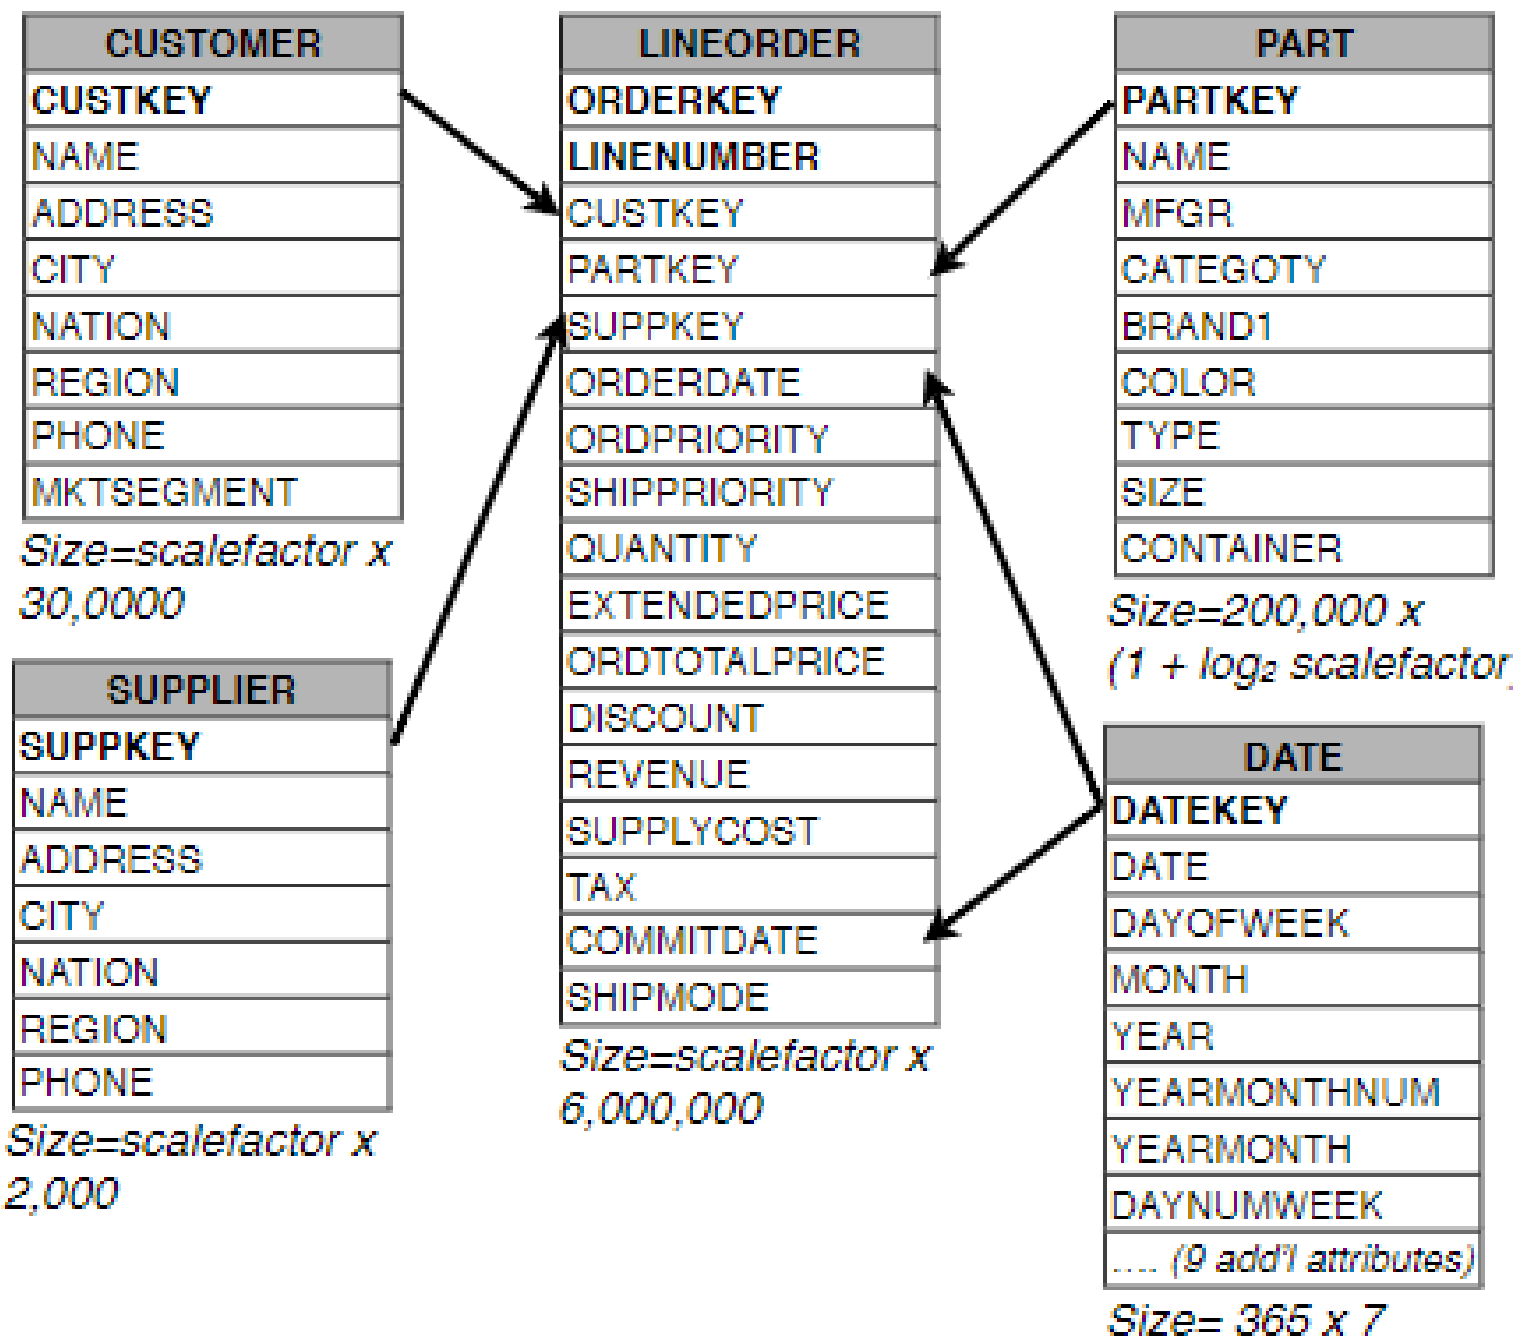


**Highlight Implications**





# Star Schema Benchmark




Flight 1: 1 dimension + DISCOUNT & QUANTITY

Flight 2: 2 dimensions & calculate revenue of product + region

Flight 3: 3 dimensions & analyze revenue in REGION + time

Flight 4: 3 dimensions & profit by YEAR, NATION, REGION






# Row-Stores

Unnecessary reads due to row-by-row
Easy read and write
Expensive and low compression rate
Best suited for transaction system

# Column-Stores

Read only relevant data
Slower read and write
Efficient when going through entire data set and high compression rate
Best suited for analytical system





# Row-Stores



## Customer


ID	Name	Country	City
28495072	Jack Hamilton	USA	Seattle
78239842	Sarah Wilson	Finland	Helsinki
19389562	Jason Huang	China	Shanghai

```
INSERT INTO CUSTOMER(ID, Name,  
Country, City)  
VALUES(37583719, 'Sam Kim',  
'South Korea', 'Seoul')
```

## Customer

ID	Name	Country	City
28495072	Jack Hamilton	USA	Seattle
78239842	Sarah Wilson	Finland	Helsinki
19389562	Jason Huang	China	Shanghai
37583719	Sam Kim	South Korea	Seoul





# Column-Stores

## Customer



ID	Name	Country	City
28495072	Jack Hamilton	USA	Seattle
78239842	Sarah Wilson	Finland	Helsinki
19389562	Jason Huang	China	Shanghai

## Customer

ID	Name
28495072	Jack Hamilton
78239842	Sarah Wilson
19389562	Jason Huang

ID	Country
28495072	USA
78239842	Finland
19389562	China



ID	City
28495072	Seattle
78239842	Helsinki
19389562	Shanghai





# Row-orientation execution

1. Vertical partitioning
2. Index-only plans
3. Materialized views





# Vertical partitioning

## Employees Table

Name	Age	department
Bob	29	Marketing
Hugh	30	Marketing
John	32	Business



Bob		1
Hugh		2
John		3


**NameTable**

29		1
30		2
32		3

**AgeTable**

Marketing		1
Marketing		2
Business		3

**DeptTable**







# Vertical partitioning

## Employees Table

NameTable	AgeTable	DeptTable
Bob   1	29   1	Marketing   1
Hugh   2	30   2	Marketing   2
John   3	32   3	Business   3

```
SELECT *  
FROM Employees  
WHERE name = "Bob"
```

```
SELECT name  
FROM Employees  
WHERE name = "Bob"
```







# Vertical partitioning

## Pros:

1. Can emulate column stores

## Cons:

1. It will need a join when different column data needs to be used
2. It requires the position attribute which wastes disc space.










# Index-only plans

Name	Age	department
Bob	29	Marketing
Hugh	30	Marketing
John	32	Business







# Index-only plans

## Pros:

1. Minimizes space and performance overhead associated with vertical partitioning



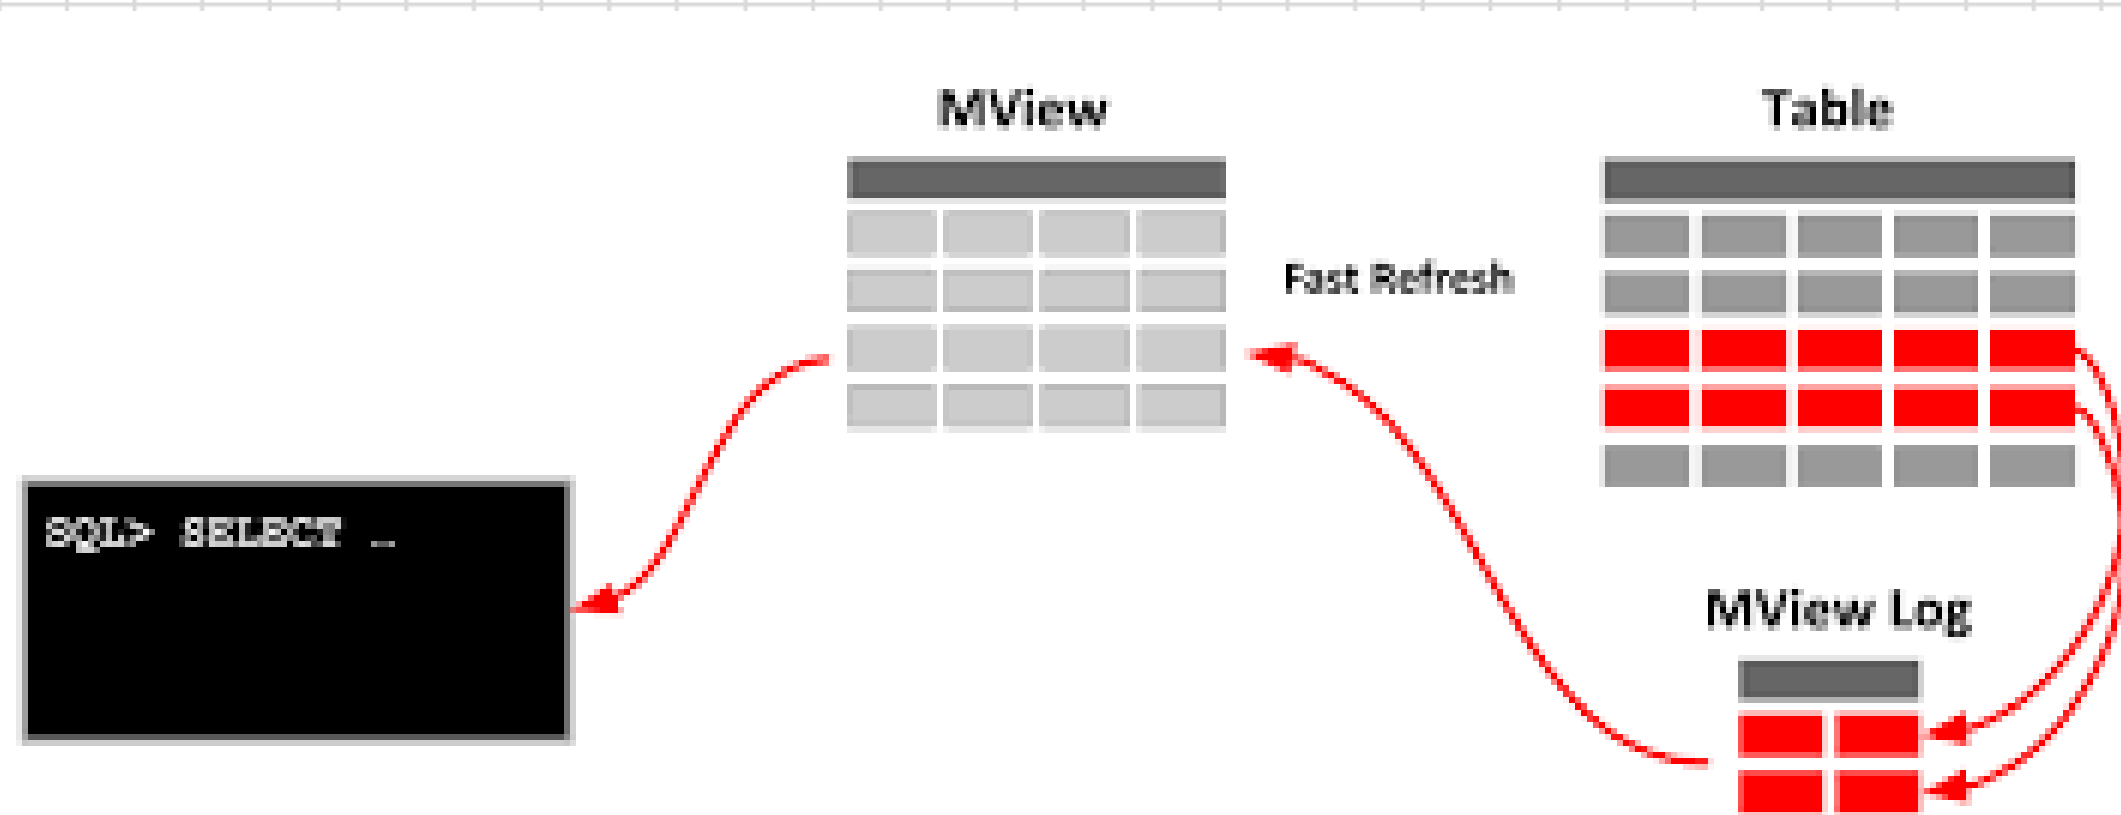
## Cons:

1. Expensive writes
2. Indexes still take up space






# Materialized views



# Materialized views

Name	Age	department
Bob	29	Marketing
Hugh	30	Marketing
John	32	Business

Save view



Age	department
29	Marketing
30	Marketing
32	Business

```
SELECT AVG(Age)  
FROM Employees  
WHERE department = "Marketing"
```





# Materialized Views

## Pros:

1. Guarantees improved performance

## Cons:



1. The query workload has to be known in advance to take advantage






# Column-orientation execution


1. Compression
2. Late materialization
3. Block iteration
4. Invisible join



# Compression



# Late materialization



Pos(T.A)	Pos(S.A)
0	1
1	0

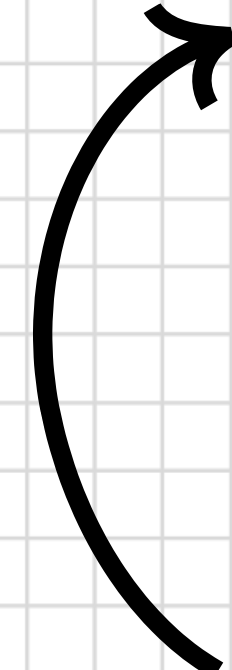
T.B	S.C
7	4
3	3



# Block iteration

ID	Name	Country	City
28495072	Jack Hamilton	USA	Seattle
78239842	Sarah Wilson	Finland	Helsinki
19389562	Jason Huang	China	Shanghai

Block



City
Seattle
Helsinki
Shanghai

# Invisible join

Apply region = 'Asia' on Customer table

custkey	region	nation	...
1	Asia	China	...
2	Europe	France	...
3	Asia	India	...

Hash table  
with keys  
1 and 3

Apply region = 'Asia' on Supplier table

supkey	region	nation	...
1	Asia	Russia	...
2	Europe	Spain	...

Hash table  
with key 1

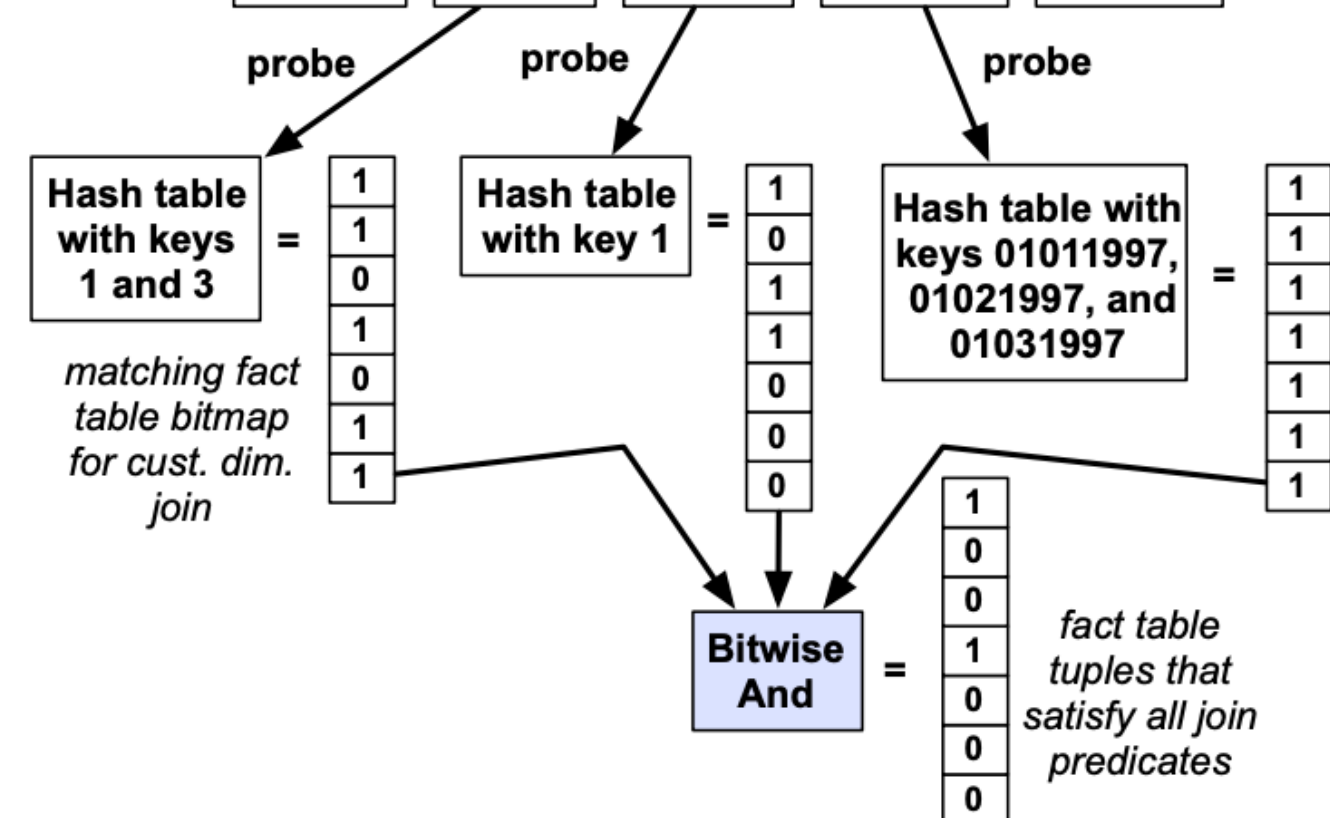
Apply year in [1992,1997] on Date table

dateid	year	...
01011997	1997	...
01021997	1997	...
01031997	1997	...

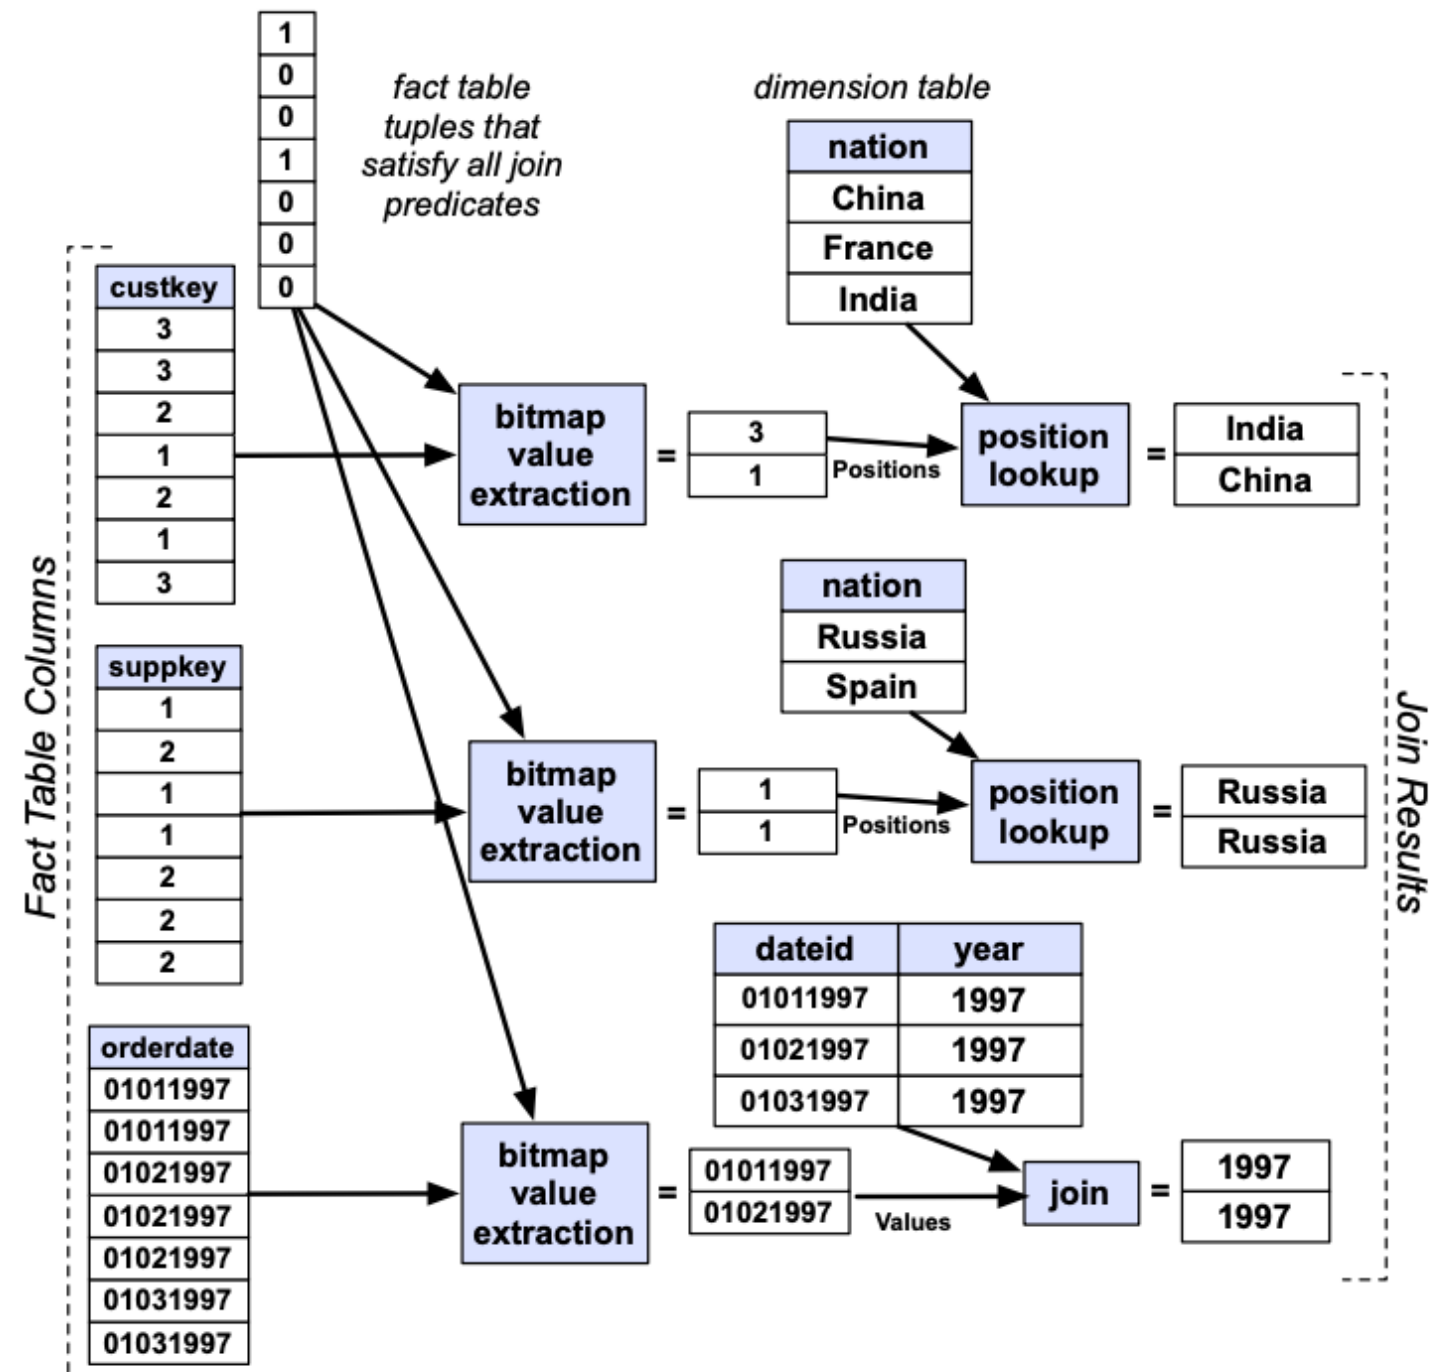
Hash table with  
keys 01011997,  
01021997, and  
01031997

Fact Table

orderkey	custkey	supkey	orderdate	revenue
1	3	1	01011997	43256
2	3	2	01011997	33333
3	2	1	01021997	12121
4	1	1	01021997	23233
5	2	2	01021997	45456
6	1	2	01031997	43251
7	3	2	01031997	34235



# Invisible join



# Experiments




Figure 5: Baseline performance of C-Store “CS” and System X “RS”, compared with materialized view cases on the same systems.

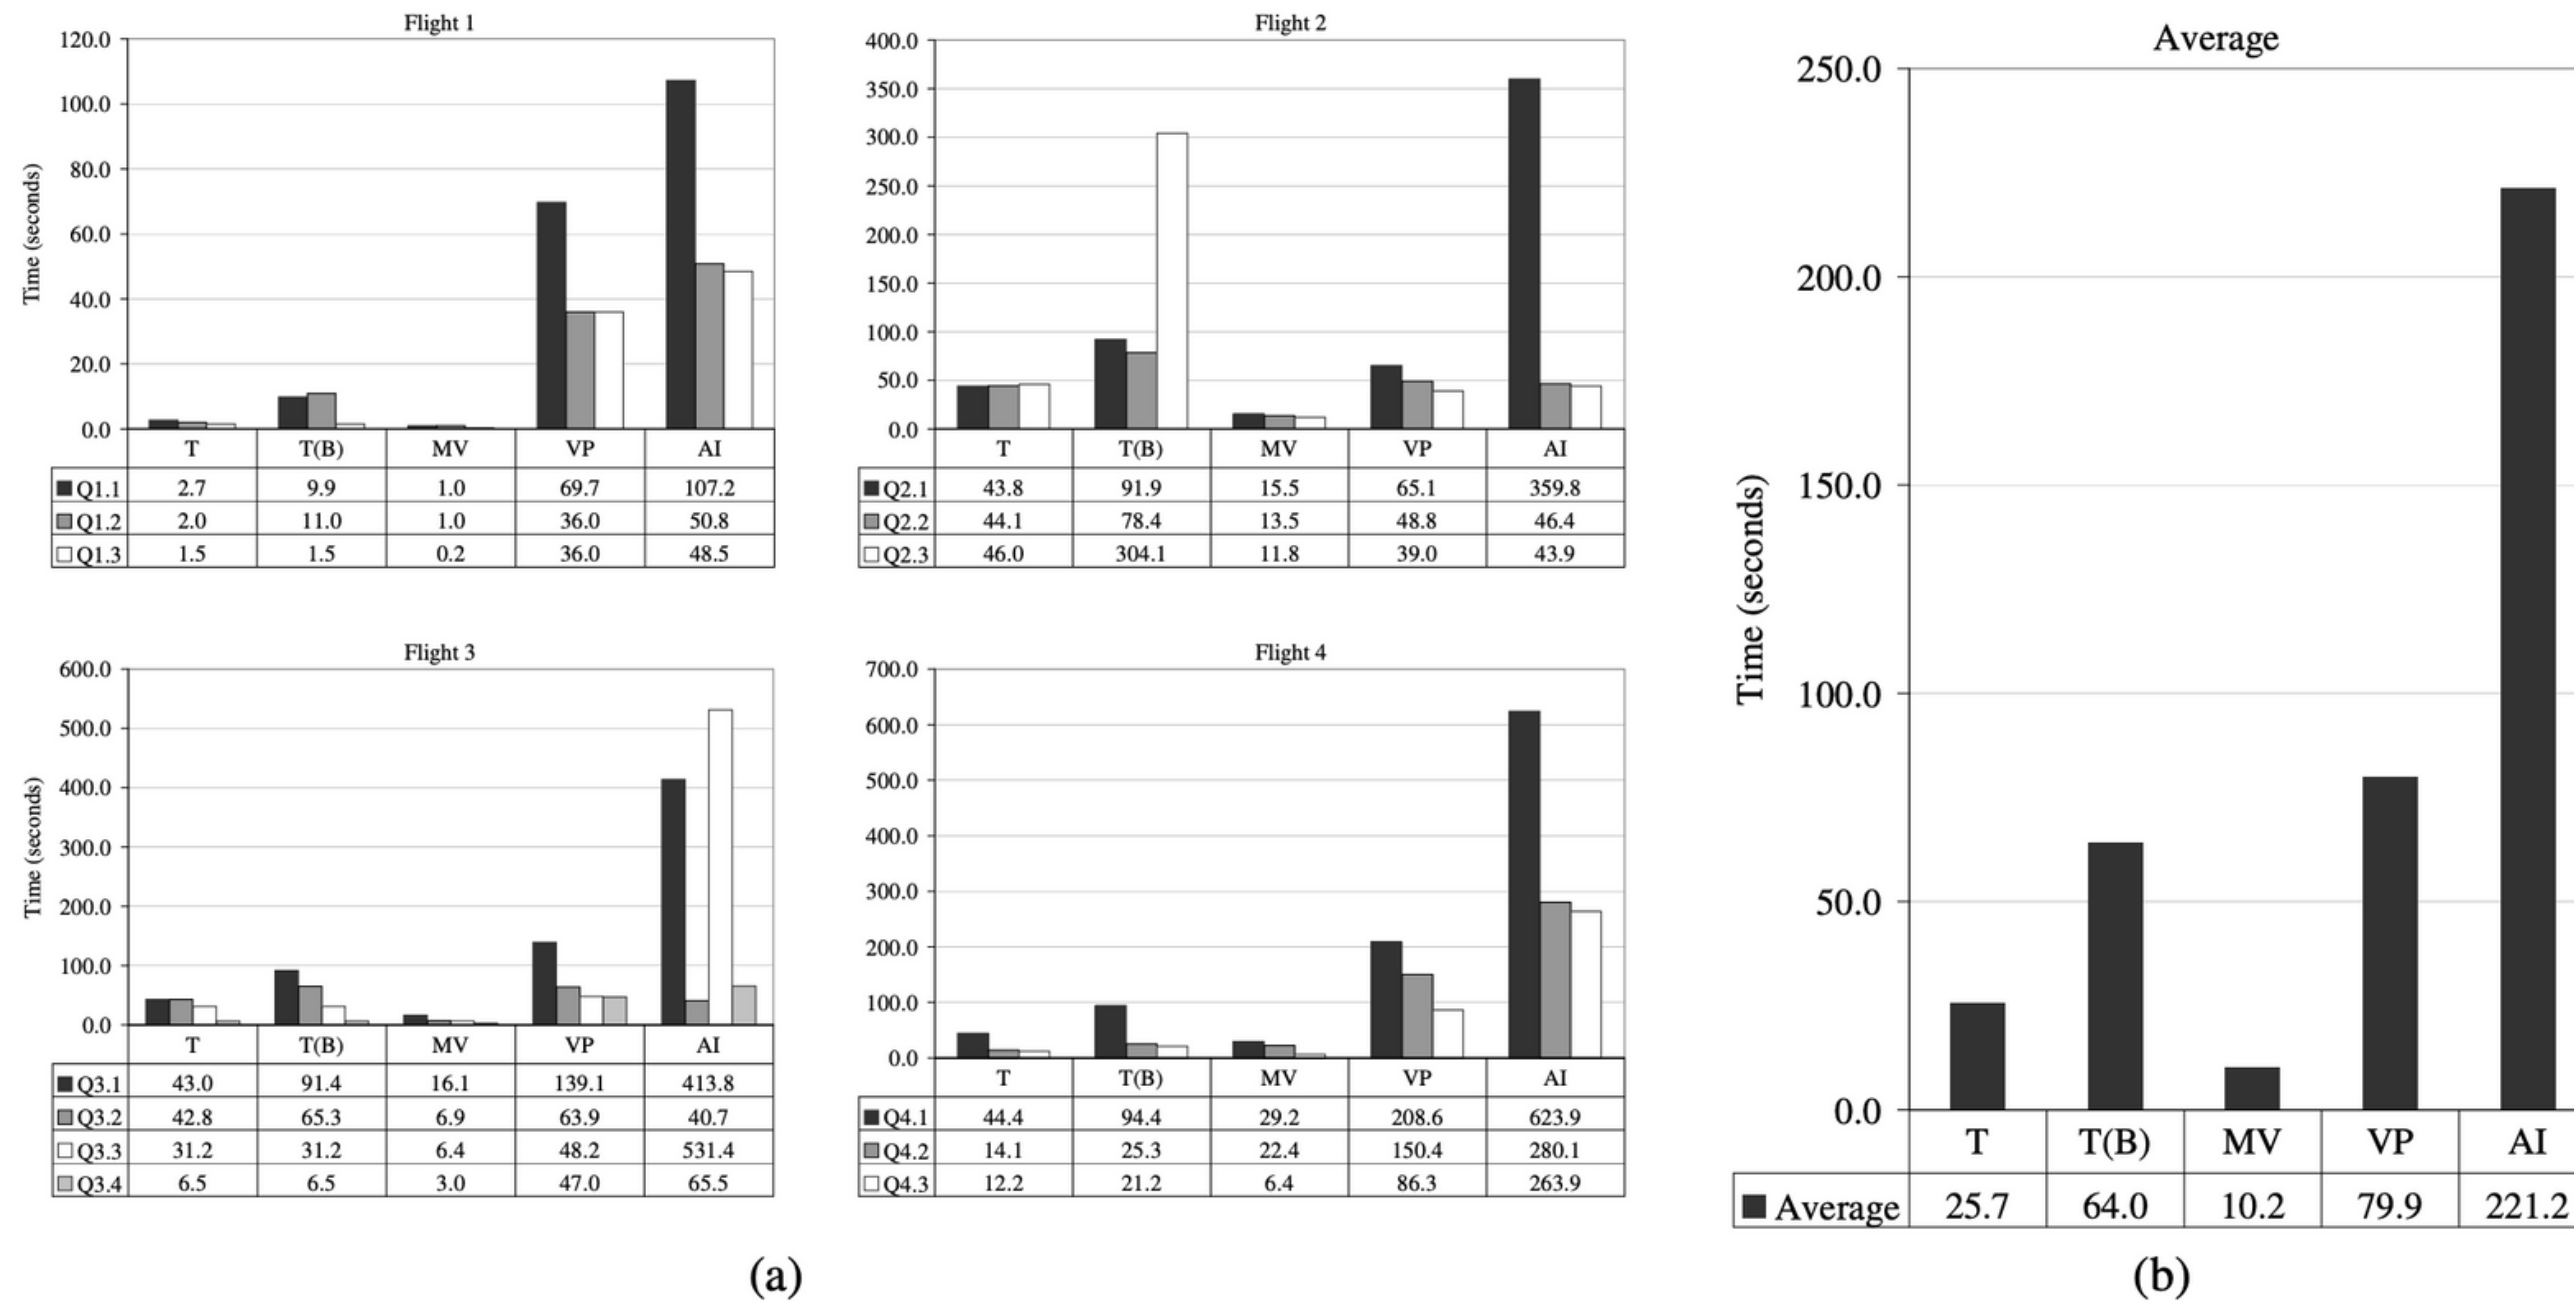
**RS: Traditional Row Store (System X)**

**RS (MV): Row Store with Materialized Views**

**CS: Traditional Column Store (C-Store)**

**CS (Row-MV): Row Store with Materialized Views (C-Store)**

# Experiments



**Figure 6: (a) Performance numbers for different variants of the row-store by query flight. Here, T is traditional, T(B) is traditional (bitmap), MV is materialized views, VP is vertical partitioning, and AI is all indexes. (b) Average performance across all queries.**

# Experiments

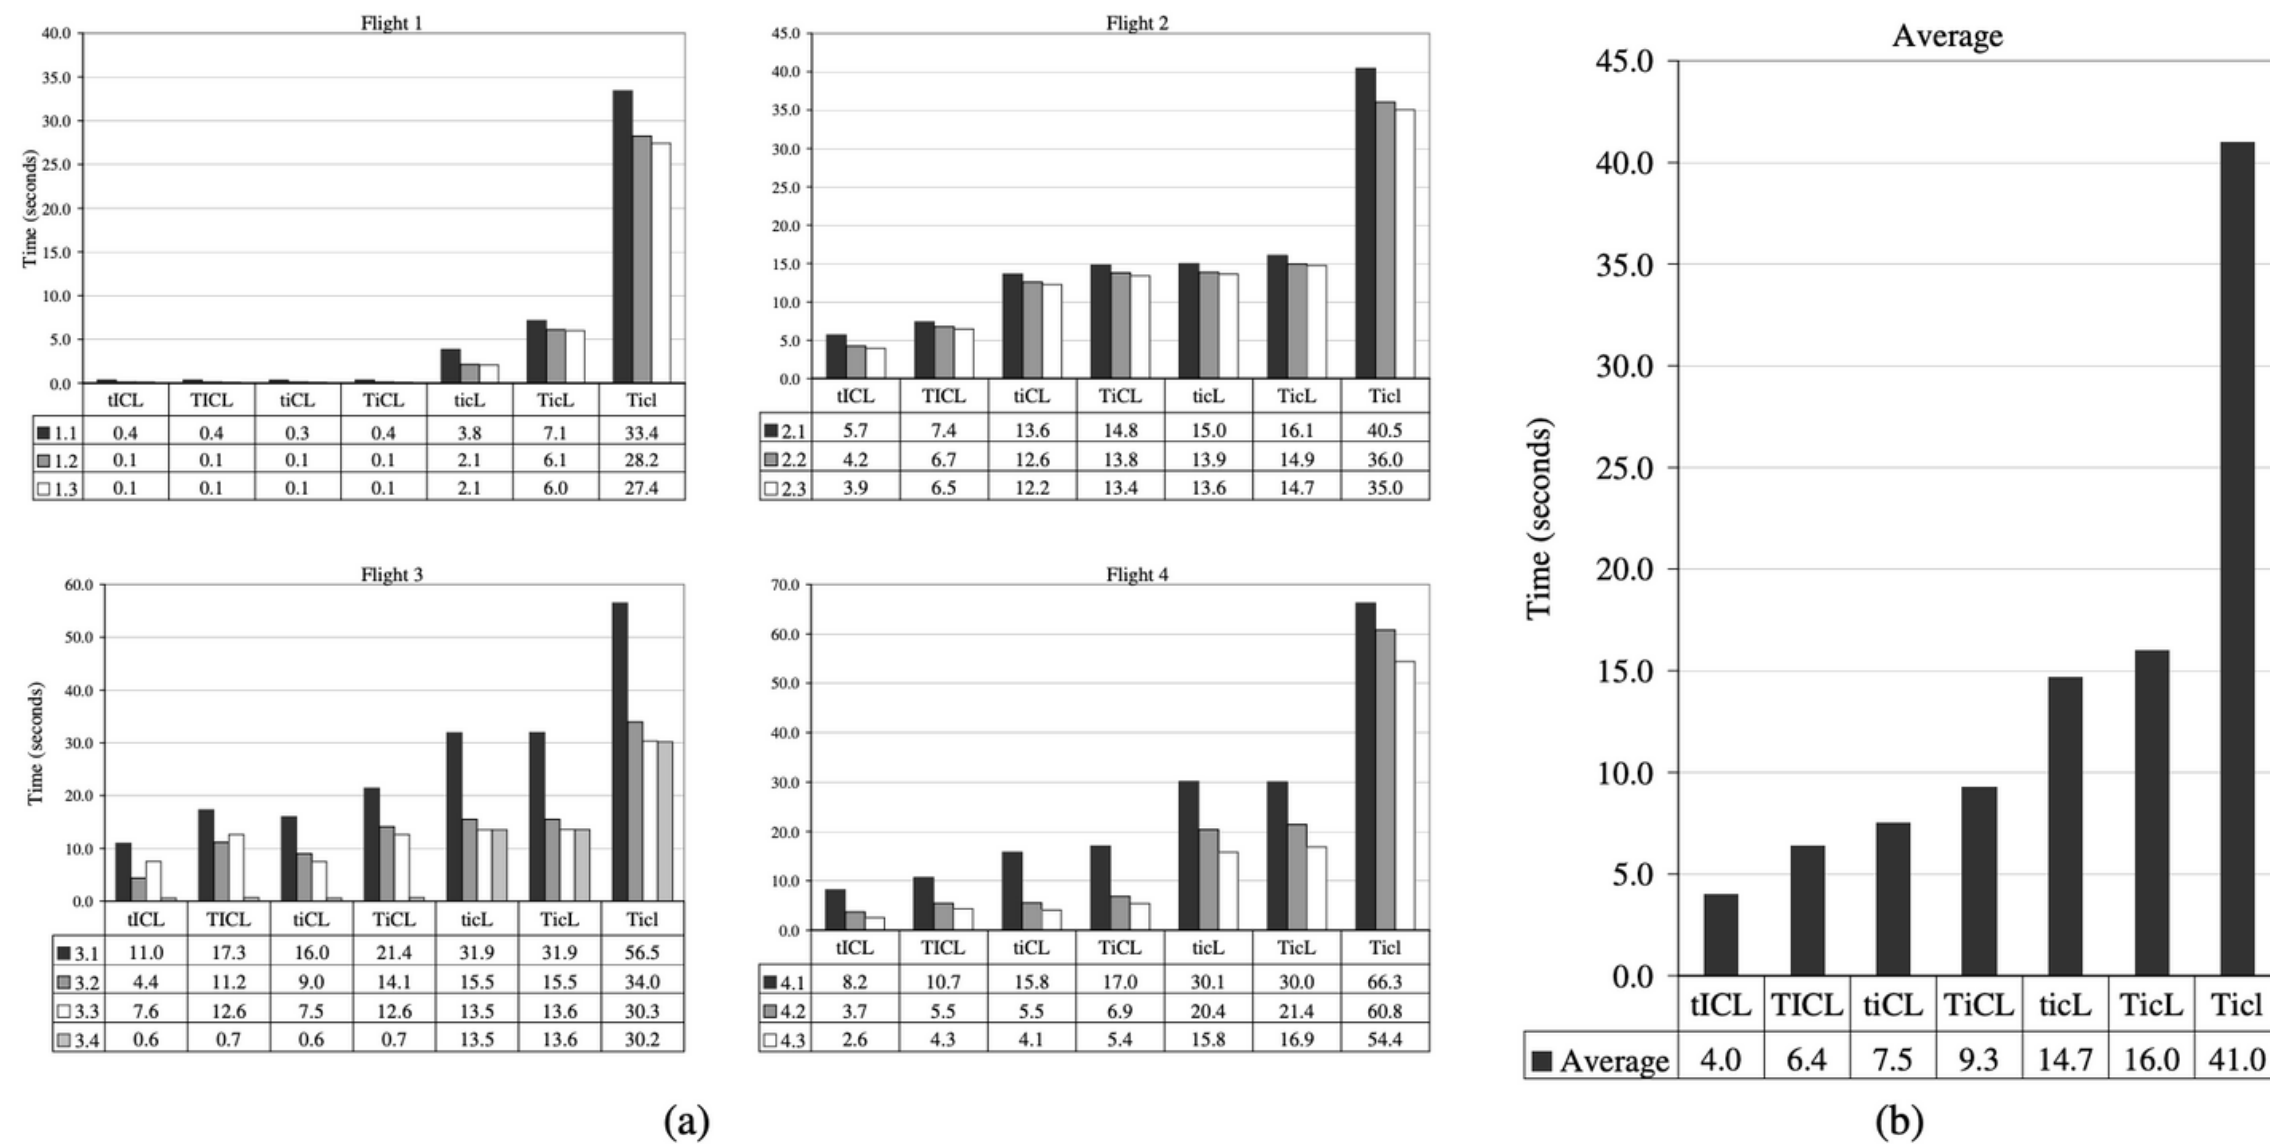



Figure 7: (a) Performance numbers for C-Store by query flight with various optimizations removed. The four letter code indicates the C-Store configuration: T=tuple-at-a-time processing, t=block processing; I=invisible join enabled, i=disabled; C=compression enabled, c=disabled; L=late materialization enabled, l=disabled. (b) Average performance numbers for C-Store across all queries.

# Experiments



**Figure 8: Comparison of performance of baseline C-Store on the original SSBM schema with a denormalized version of the schema. Denormalized columns are either not compressed (“PJ, No C”), dictionary compressed into integers (“PJ, Int C”), or compressed as much as possible (“PJ, Max C”).**

# Related Work

## Monet DB Monet DB/X100

- Pioneers of the design of modern column-oriented database system
- Contributed to superior CPU and cache performance and reduce I/O compared to row-store

## C-Store

- Optimization for direct operation on compressed data
- Dramatic performance improvements on warehouse workloads

## Fractured Mirrors

- Hybrid of row/column-stores
- Updates using row-store and reads using column-store

## Shore

- Halverson et al worked with Shore to compare against unmodified version
- “Super tuples” optimization made vertically partitioned database competitive to column-stores



# What Did We Learn?




**Differences in  
Architecture**



**Need for Changes  
in Row-Stores**




**Future  
Directions**



# In Class Discussion


**How can we further improve a row-store & column-store?**





# In Class Discussion

**Do you think row-store can follow the performance rate of column-store in the future?**





# Q&A

