



**Hasso
Plattner
Institut**


IT Systems Engineering | Universität Potsdam

Topics

1) Partitioned Bitvector

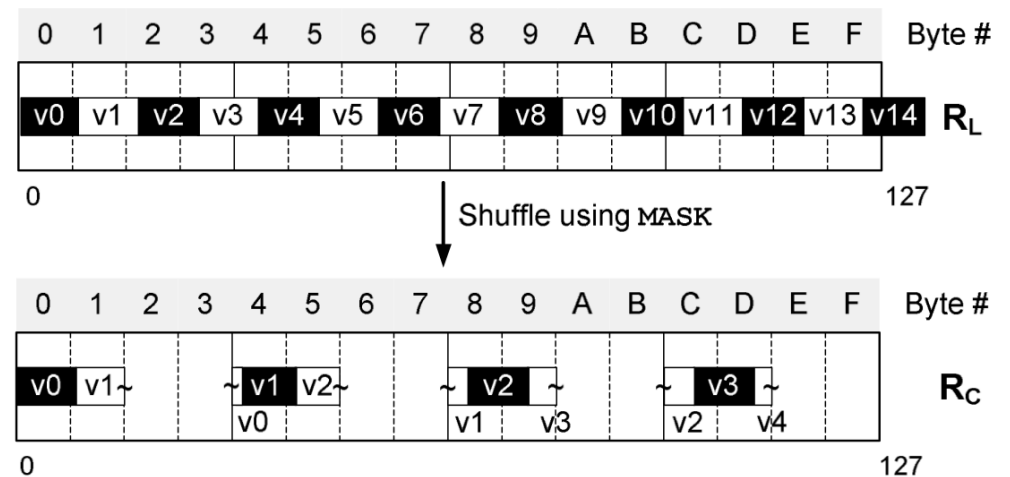



2

Delta Dictionary		Delta Partition (Compressed)	
U_D			
bravo	00	00	bravo
charlie	01	01	charlie
golf	10	10	golf
young	11	01	charlie
		11	young



2) Vertical Bitvector

3



3) Unsorted Dictionaries

4

Delta Dictionary

U'_D

bravo	00
charlie	01
golf	10
young	11

Delta Partition
(Compressed)

00	bravo
01	charlie
10	golf
01	charlie
11	young

SELECT name FROM table WHERE name > 'charlie'

4) Default Value Index

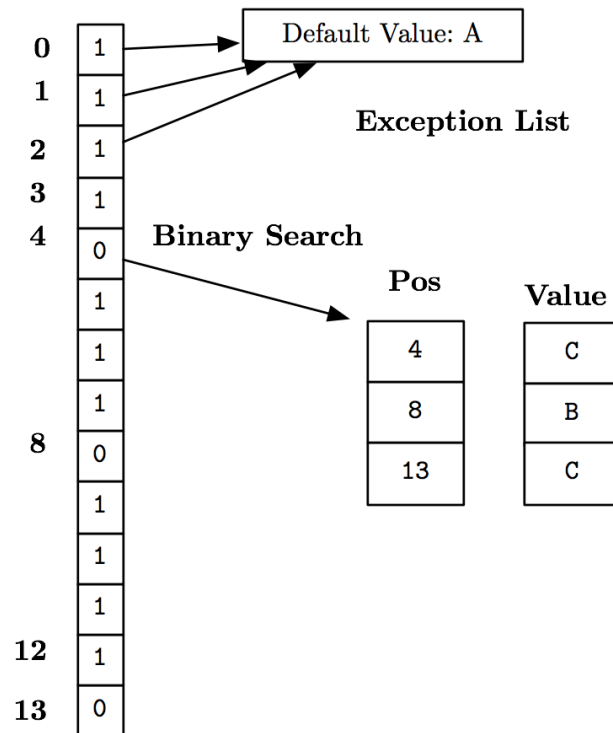
5

Default-Value Bitvector

Logical View



A
A
A
A
C
A
A
A
A
A
B
A
A
A
A
C

Storage Layout

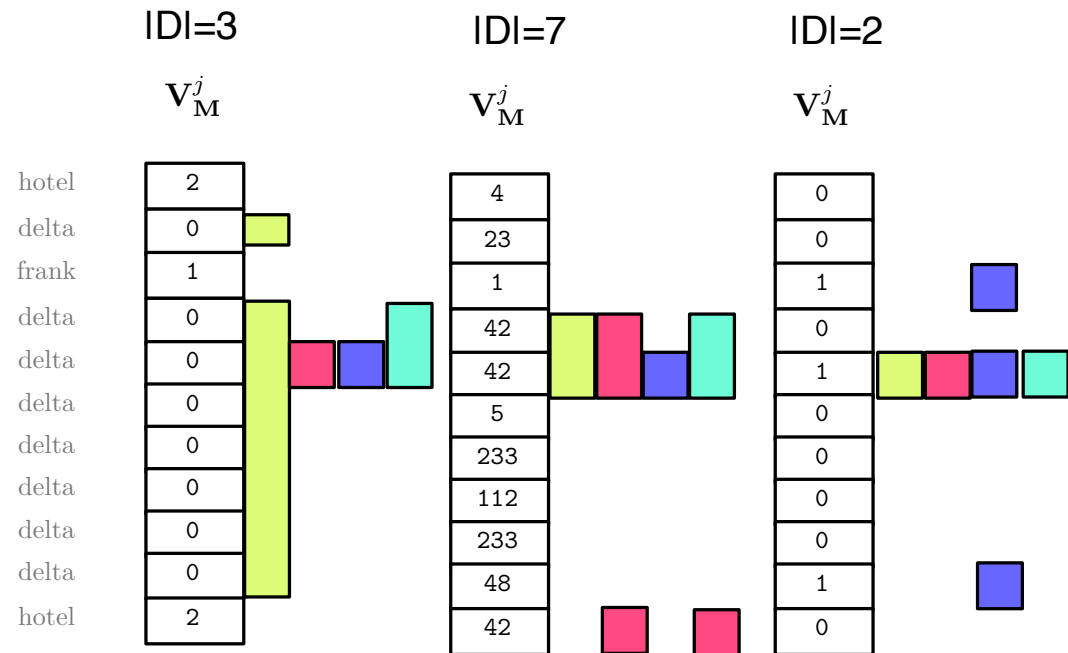


5) Multi Column Indices

6



- ABC
- BCA
- CBA
- BAC



Select * from T where colA='delta' and colB='42' and colC=1

6) Multi Column Join Algorithms

7

When joining tables with multiple key columns, column stores need to materialize all keys or need an uncompressed tree structure with actual values not value-ids.

- How can column stores perform fast and optimized joins over multiple columns
- Avoid complete materialization, leverage dictionary compression
- Special index structures might be used

8) Cache-Optimized Parallel Join

9

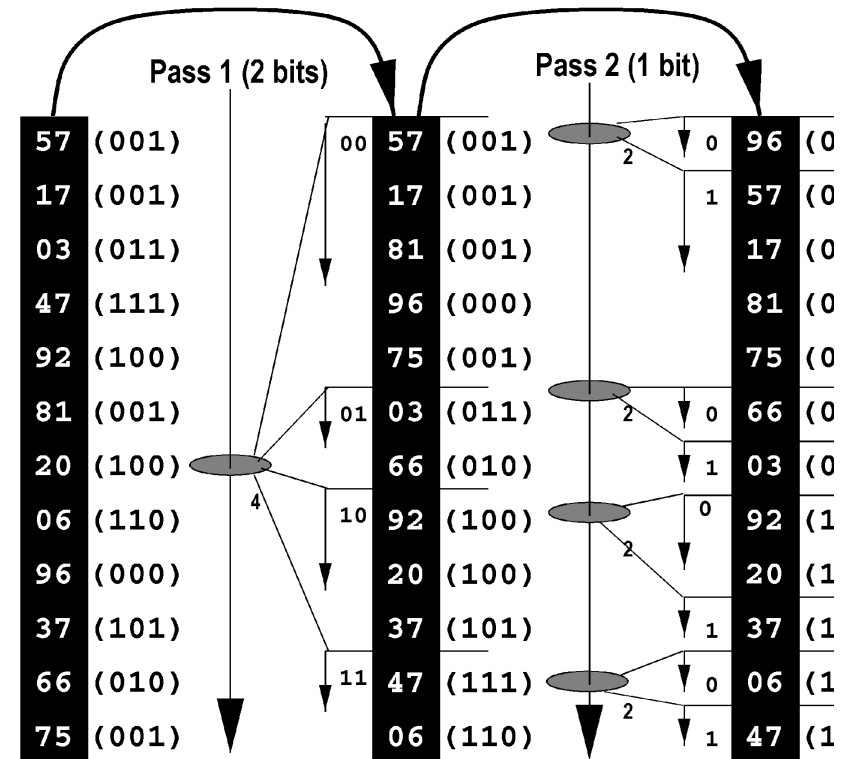



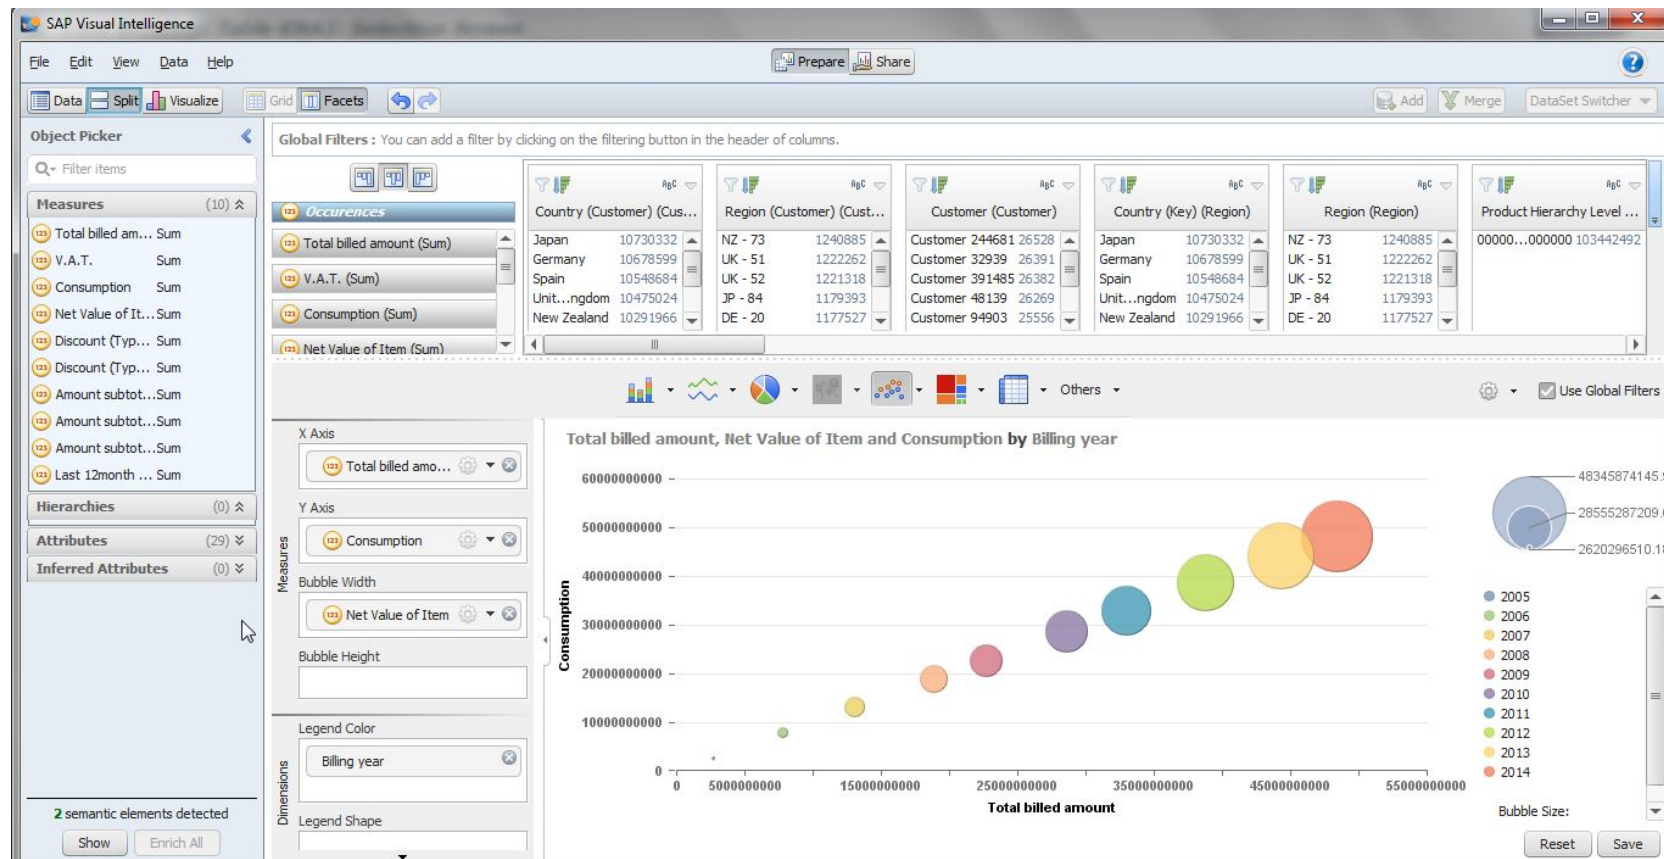
Fig. 9. Two-pass/3-bit radix cluster (lower bits indicated by parentheses).

straightforward clustering algorithm.

From [1] S. Manegold, P. Boncz, M. Kersten, S. Manegold, P. Boncz, and M. Kersten, "Optimizing main-memory join on modern hardware," IEEE Transactions on Knowledge and Data Eng, 2002.


9) Analytical User Interface for HYRISE

10



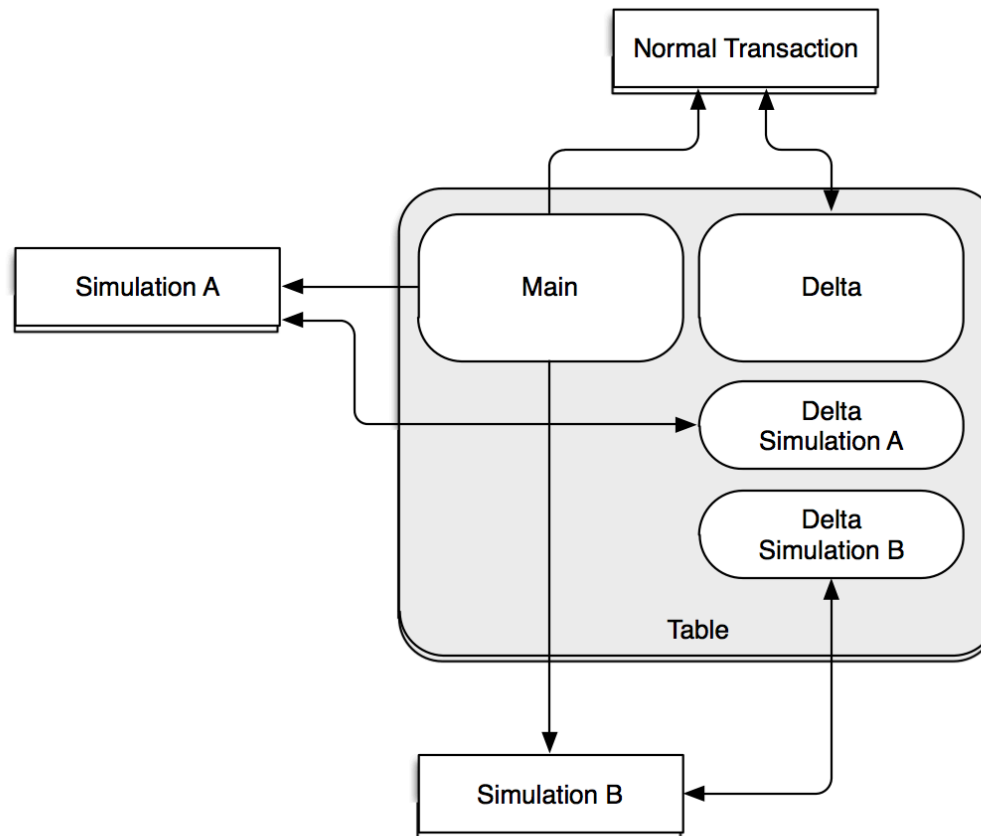
10) Set-based Query Language

11



11) Branching Deltas (Simulation)

12



12) Result Set Compression

13

Currently, result sets are transferred “as is” to the client, where it is simply read row-by-row.

- In scenarios where the connection between database and applications is limited (mobile, cloud), compressing the result set might lead to significant performance gains
- Standard compression algorithms (like gzip)
- Late materialization on client side

13) Memory Compression Algorithms in Hardware (PPC)

14

Modern PowerPC processors have hardware supported main memory compression.

- Main memory is separated into two areas – compressed and uncompressed
- When accessing memory area that is compressed, decompression is performed transparently
- Memory compression double available memory size

A first evaluation could investigate how well it compresses enterprise data, how performance is affected, and how a row store compares to a column store.

14) Query runtime prediction based on optimizer results

15

The idea is to use a cost-based query optimizer to predict query execution time:

- In the seminar, the student should validate whether the HANA query optimizer can be used to roughly predict the query execution time without actually running it
- The objective is to support developers to estimate the runtime of queries during development

Envisioned steps:

- Get to know the query optimizer of HANA
- Experimentally validate the correlation of estimated cost and actual run-time
- Develop a wrapper that estimates runtime for a given query

15) Execution model for co-processors

Based on the memory bandwidth limitation of database co-processors, multiple execution models are possible. Evaluation of the following different models: Co-Process, Procedural Language Architecture, GPU hosted data.

16) Co-processor integration into databases

Evaluation how a co-processor can be integrated with a database for application specific logic with SAP HANA (AFL) or HYRISE.

17) Automatic execution unit detection in a heterogeneous system

In a heterogeneous system, the capabilities of database (co-)processors differ with respects to bandwidth, parallelization, clock frequency. To determine the best suited execution unit an execution model needs to be created based on device benchmarks and a programs current workflow.

18) Optimized co-processor data structures

Database co-processors used for application specific tasks, such as prediction algorithms or scientific calculations, require the to have a different layout than database tables. The student working on this topic will evaluate which data structures are favorable for different co-processors and algorithms.

19) Automatic code optimization for co-processors

As a heterogeneous programming framework, OpenCL enables the developer to write a single program that can be executed on many different hardware platforms. Nevertheless, the code needs to be optimized for each device to run optimal. The goal will be to optimize the code written in OpenCL C for different devices automatically based on micro benchmarks.