

Costs for Hash Join

- Assuming all tuples of both R and S can be place on main memory
 - Consider disk base hash join later
- We have to consider the following costs
 - Applying Hash Function: $\{R\} + \{S\}$
 - Building the Hash Table: $\{R\}$
 - Comparisons: $\{S\}$

Footnote: 2019/7/15 Advance Data Engineering (©H.Yokota) 151

Hash Join with Disk Accesses

- Three Algorithms
 - Simple Hash Join
 - GRACE Hash Join
 - Hybrid Hash Join
- Assumptions for cost estimation
 - Memory size for hash table: $|M|$
 - The distribution of attribute value is flat
 - $|R| > |S|$

Footnote: 2019/7/15 Advance Data Engineering (©H.Yokota) 152

Simple Hash Join

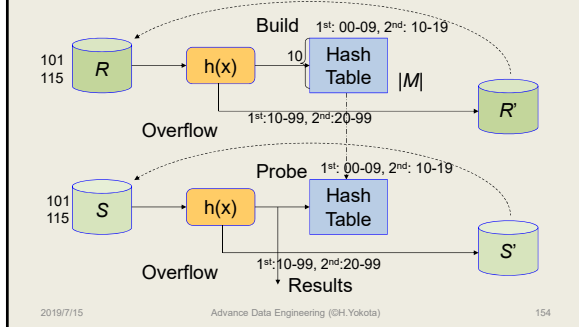
- Restore tuples that cannot be placed in main memory hash table
 - Repeat until no overflow
 - Both R and S (because of the identical hash table)
- The number of repeat (Expectation):
 - $L = |R|/|M| (\geq 1)$

2019/7/15

Advance Data Engineering (©H.Yokota)

153

Illustration of Simple Hash Join



2019/7/15

Advance Data Engineering (©H.Yokota)

154

Example of Overflow Control

- Suppose an attribute storing integer value
- Cardinality of R is 100
- The hash table $|M|$ can keep 10 tuples of R
- Apply hash function $h(X) = \text{mod}(X, 100)$
 - It generates two-digit numbers
 - Lower digit is used to choose an entry of the hash table
 - Higher digit is used to decide overflow tuples
- The first iteration: tuples of R , having higher digit of hash value is 0, are used to build the hash table, and others become overflow
- The second iteration: tuples having higher digit $\neq 1$ becomes overflow
- And so on
- Iteration count becomes 10 or more

2019/7/15

Advance Data Engineering (©H.Yokota)

155

Cost for Simple Hash Join (1)

- In the first loop
 - Applying Hash Function: $\{R\} + \{S\}$
 - Comparison: $\{S\} / L$
 - Disk I/O for read: $|R| + |S|$
 - Disk I/O for write: $|R| - |M| + |S| - |S| / L$
 - or $(|R| + |S|) - (|R| + |S|) / L$
 - from $|M| = |R| / L$

2019/7/15

Advance Data Engineering (©H.Yokota)

156

Cost for Simple Hash Join (2)

- In the i -th loop ($i = 2, \dots, L$)
 - Applying Hash Function:
 $\{R\} - (i - 1) \times \{R\} / L + \{S\} - (i - 1) \times \{S\} / L$
 - Comparison: $\{S\} / L$
 - Disk I/O for read:
 $(|R| + |S|) - (i - 1) \times (|R| + |S|) / L$
 - Disk I/O for write:
 $(|R| + |S|) - i \times (|R| + |S|) / L$

2019/7/15

Advance Data Engineering (©H.Yokota)

157

Cost for Simple Hash Join (3)

- Summation from the first to L -th (last) loop
 - Applying Hash Function:

$$L \times (\{R\} + \{S\}) - (L \times (L - 1) / 2 \times (\{R\} / L + \{S\} / L))$$

$$= (L + 1) / 2 \times (\{R\} + \{S\})$$

$$= (|R| + |M|) \times (\{R\} + \{S\}) / 2 |M|$$
 - Comparison: $(\{S\} / L) \times L = \{S\}$
 - Disk I/O :

$$2L \times (|R| + |S|) - (L \times (L + 1)) / 2$$

$$+ L \times (L - 1) / 2 \times (|R| + |S|) / L$$

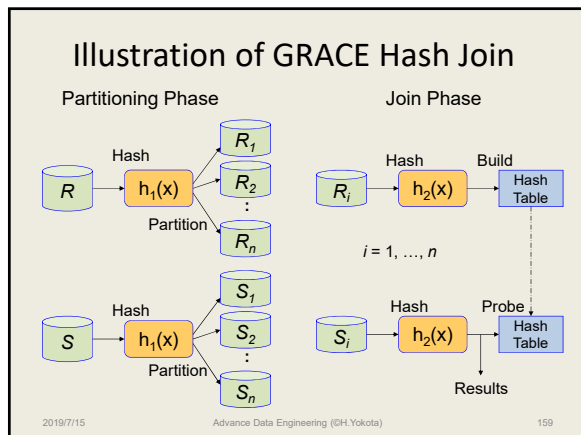
$$= L \times (|R| + |S|)$$

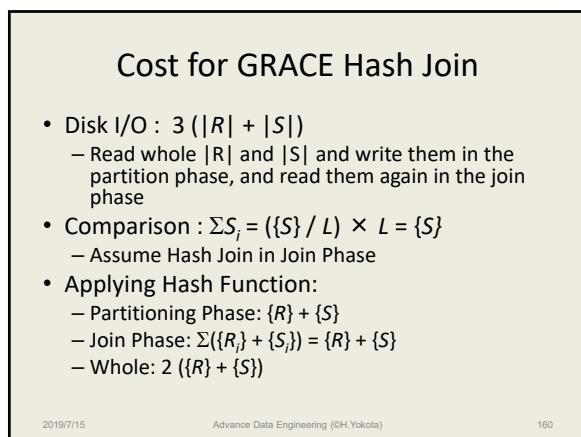
$$= (|R| \times (|R| + |S|)) / |M|$$

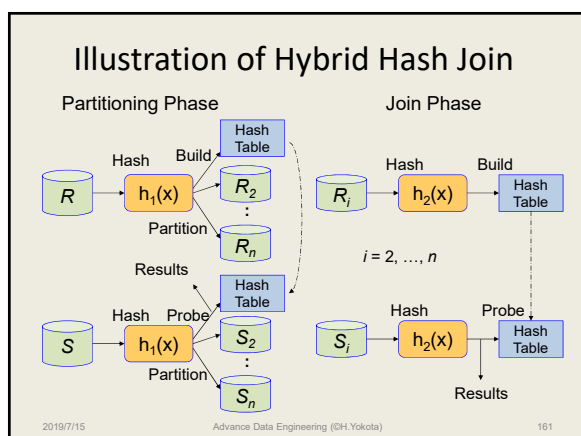
2019/7/15

Advance Data Engineering (©H.Yokota)

158







Cost for Hybrid Hash Join

- Disk I/O :
 - $3(|R| + |S|) - 2(|R| + |S|) / L$
 - $= (3 - 2|M| / |R|) \times (|R| + |S|)$
- Comparison and Applying Hash Function
 - The same as GRACE Hash Join

2019/7/15

Advance Data Engineering (©H.Yokota)

162

Comparison of Hash Join Algorithm

	Disk I/O	Applying Hash	Comparison
Simple	$\frac{ R }{ M }(R + S)$	$\frac{ R + M }{2 M }(R + S)$	$\{S\}$
GRACE	$3(R + S)$	$2(R + S)$	$\{S\}$
Hybrid	$(3 - 2\frac{ M }{ R })(R + S)$	$2(R + S)$	$\{S\}$

2019/7/15

Advance Data Engineering (©H.Yokota)

163

Optimization of Join (1)

- Star Query (Oracle 7)
 - Derive Cartesian Products among Dimension Tables
 - to avoid handling the huge Fact Table
 - Tradeoff between the cost for Cartesian Product and that for handling the Fact Table
 - The dimension tables can be filtered by conditions in advance
 - e.g. A4 Laptop, Aug., etc

2019/7/15

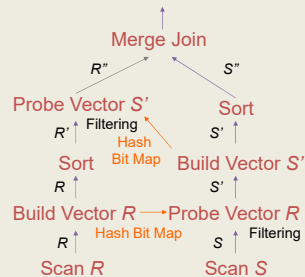
Advance Data Engineering (©H.Yokota)

164

Optimization of Join (2)

- Bit Vector Filtering

- Bloom Filtering
- Make each entry of the Hash Table 1 bit
 - Only Existence (Allow collisions)
- Bit Vector can be placed in main memory
- This filtering can apply both sort merge and hash join



2019/7/15

Advance Data Engineering (©H.Yokota)

165

Join Index (1)

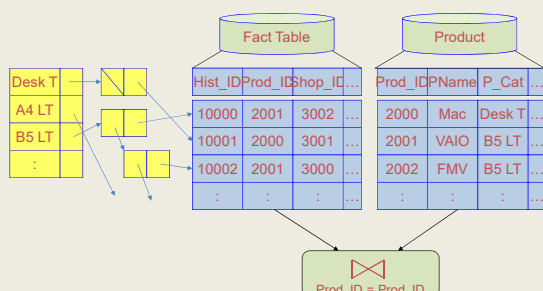
- Assuming join attributes for relations R and S are R.A and S.B, respectively, an index from the other attribute of R.C of the relation R is a Join Index.
 - Structure: Inverted File or B⁺-tree
 - Collisions for one entry: Inverted List or in Inverted File / B-tree

2019/7/15

Advance Data Engineering (©H.Yokota)

166

Join Index (2)




2019/7/15

Advance Data Engineering (©H.Yokota)

167

Multidimensional Join Index

- Join Index for combination of multiple attributes



2019/7/15

Advance Data Engineering (©H.Yokota)

168

Bitmap Index

- Toggle 0 or 1 by existence of a value of the attribute
- When the variety of values is small, it provides good space efficiency
 - Bitmap can be place in main memory
 - Bitmap can be calculated as an array
 - AND/OR operation can be used for filtering
- On the other hand, large variety of attribute values make the space efficiency worse.
 - combination with value index by segmentation is proposed

2019/7/15

Advance Data Engineering (©H.Yokota)

169

An Example of Bitmap Index

Condition: Prod_ID=2001 and Shop_ID=3001

Fact Table	Prod_ID Bitmap Index	Shop_ID Bitmap Index	Results
Hist_ID Prod_ID Shop_ID ...	2000 2001 2002 2003	3000 3001 3002 3003	
10000 2001 3002 ...	0 1 0 0	0 0 1 0	0
10001 2000 3001 ...	1 0 0 0	0 1 0 0	0
10002 2001 3000 ...	0 1 0 0	1 0 0 0	0
10003 2020 3002 ...	0 0 0 0	0 0 1 0	0
10004 2003 3010 ...	0 0 0 1	0 0 0 0	0
10002 2001 3001 ...	0 1 0 0	0 1 0 0	1
10003 2002 3011 ...	0 0 1 0	0 0 0 0	0
...

2019/7/15

Advance Data Engineering (©H.Yokota)

170

Semi-Join Operation

- A kind of eq-join, but attributes in the result relation belong to only one relation

Product ID	Product Type	Price		Product ID	Stocks
P001	Laptop PC	1,500	⋈	P001	10
P002	Laptop PC	2,000		P003	5
P003	Desktop PC	2,000			

Product ID	Product Type	Price
P001	Laptop PC	1,500
P003	Desktop PC	2,000

2019/7/15

Advance Data Engineering (©H.Yokota)

171

Anti-Semi-Join Operation

- Similar to semi-join, but generating tuples which does not much with another relation

Product ID	Product Type	Price		Product ID	Stocks
P001	Laptop PC	1,500	⋈ ⁻	P001	10
P002	Laptop PC	2,000		P003	5
P003	Desktop PC	2,000			

Product ID	Product Type	Price
P002	Laptop PC	2,000

2019/7/15

Advance Data Engineering (©H.Yokota)

172

Costs of Semi-Join & Anti-Semi-Join

- Cost of Semi-Join are basically same as ordinary join operations
 - For the case of hash join, the size of hash table can be reduced
- Implementation of anti-semi-join is also similar to semi-join operation except outputting unmached tuples instead of mached tuples
 - The cost of anti-semi-join is equal to semi join
- These operations are related to set operations
- Semi-Join is also related to distributed database

2019/7/15

Advance Data Engineering (©H.Yokota)

173

Implementation of Set Operations

- **Intersection** ($R \cap S$)
 - Apply the **semi-join** targeting all attributes as conditions
 - Costs is equal to one for a join operation
- **Difference** ($R - S$)
 - Apply the **anti-semi-join** targeting all attributes
 - Costs is also equal to one for a join operation
- **Union** ($R \cup S$)
 - Apply difference operation to derive $(R-S)$ and concatenate $(R-S)$ with S
 - Or concatenate R with S and eliminate duplications

2019/7/15

Advance Data Engineering (©H.Yokota)

174

Division \div

- Division produce a relation that consists of the set of tuples from R defined over the attribute C that match the combination of every tuple in S , where C is the set of attributes that are in R but not in S

R			S	
Product ID	Salesman		Product ID	
P001	John	\div	P001	
P001	Alice		P002	
P001	Bill		P003	
P002	John			
P002	Bill			
P003	John			
P003	Bill			
P004	Alice			
P004	Bill			

$R[\text{ProductID}] \div S[\text{ProductID}]$
Salesman who sold all products listed in relation S

Salesman
Bill
John

2019/7/15

Advance Data Engineering (©H.Yokota)

175

Implementation of $R[X] \div S[X]$ (1)

- Sort Base Direct Method
 - Sort R by X as minor and C as major
 - Check X of R with S from top of the same value C

(X)	R	(C)		minor	major		S
Product ID		Salesman		Product ID	Salesman		Product ID
P001		John	Sort	P001	Alice	\div	P001
P001		Alice		P004	Alice		P002
P001		Bill		P001	Bill		P003
P002		John		P002	Bill		
P002		Bill		P003	Bill		
P003		John		P004	Bill		
P003		Bill		P001	John		
P004		Alice		P002	John		
P004		Bill		P003	John		

results

Salesman
Bill
John

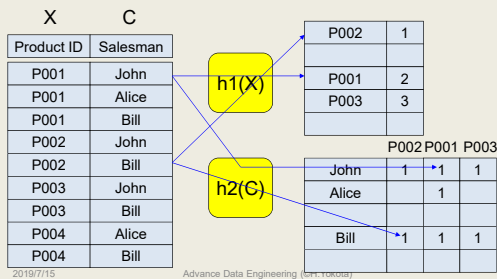
2019/7/15

Advance Data Engineering (©H.Yokota)

176

Implementation of $R[X] \div S[X]$ (2)

- Hash Base Direct Method
 - Prepare two hash functions $h1(X)$ and $h2(C)$
 - $h1(X)$ is used to derive entry # in the indicated $h2(C)$ and set 1



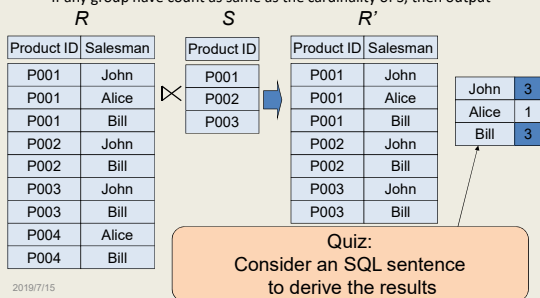
2019/7/15

Advance Data Engineering (©H. Yokota)

177

Implementation of $R[X] \div S[X]$ (3)

- Aggregate Function based Indirect Method
 - Apply semi-join in advance, then count tuples for each group of C
 - If any group have count as same as the cardinality of S, then output



2019/7/15

Cost for Group-by and Aggregation Operations (1)

- Nested Loop base
 - Scan all tuples to search tuples in the same group, and then calculate COUNT, SUM, AVG, MAX, MIN
 - $(\{R\} - 1) + (\{R\} - 2) + \dots + 1$
 - Comparison : $\{R\}(\{R\} - 1) / 2$
 - Write the result into disk at the last
 - Disk I/O: $|R|(|R| + 1) / 2$

2019/7/15

Advance Data Engineering (©H. Yokota)

179

Cost for Group-by and Aggregation Operations (2)

- Sort base
 - At first sort all tuples and scan from the top
 - Assume Merge Sort
 - Cost for sort and one scan at the last
 - Comparison : $\{R\}(\log\{R\}+1)$
 - Disk I/O : $|R|(2 \log |R| + 1)$
- Hash base
 - Divide tuples by a hash function
 - Apply the aggregate function for each hash bucket
 - If there is no hash collision
 - Applying Hash Function, Comparison: $\{R\}$
 - Disk I/O : $|R|$

2019/7/15

Advance Data Engineering (©H.Yokota)

180

Indexing for Aggregation Functions

- Bit-Sliced Index
 - Divide each bit of binary expression of an integer value
 - A list of each bit in tuples is treated as a bitmap index
 - Each bit-slice can be placed on main memory
 - Calculate SUM or AVG for each bit-slice and summarize them
 - Adopted by Sybase IQ, CCA Model 204

2019/7/15

Advance Data Engineering (©H.Yokota)

181

Indexing for Aggregation Functions

- Projection Index
 - Derive an attribute (Projection), and access by its location.
 - Reduce Disk I/O by reducing amount of data
 - Suited for complex aggregate functions on the attribute
 - Derive SUM for the results of some calculation
 - Adopted by Sybase IQ

2019/7/15

Advance Data Engineering (©H.Yokota)

182

Comparison on the AF Indices

Aggregate	Value-list	Bit-Sliced	Projection
Max, Min	Best	Slow	Slow
SUM, AVG	Not Bad	Best	Good
SUM(A1*(1-A2))	Very Slow	Very Slow	Best
Narrow Range	Best	Good	Good
Wide Range	Not Bad	Best	Good

2019/7/15

Advance Data Engineering (©H. Yokota)

183
