


Database Systems on Modern CPU Architectures

Adrian Riedl, Philipp Fent


Tuesday 18th July, 2023

Chair of Data Science and Engineering
TUM School of Computation, Information and Technology
Technical University of Munich




Insert the following tuples into an extendible hash table whose buckets can hold up to two entries.

ID	Name	hash
10	Müller	010100
25	Meier	100110
30	Schmidt	011110
18	Krause	010010
40	Schulz	000101
45	Kaufmann	101101




Insert the following tuples into an extendible hash table whose buckets can hold up to two entries.

ID	Name	hash
10	Müller	010100
25	Meier	100110
30	Schmidt	011110
18	Krause	010010
40	Schulz	000101
45	Kaufmann	101101




Insert the following tuples into an extendible hash table whose buckets can hold up to two entries.

ID	Name	hash
10	Müller	010100
25	Meier	100110
30	Schmidt	011110
18	Krause	010010
40	Schulz	000101
45	Kaufmann	101101




Insert the following tuples into an extendible hash table whose buckets can hold up to two entries.

ID	Name	hash
10	Müller	010100
25	Meier	100110
30	Schmidt	011110
18	Krause	010010
40	Schulz	000101
45	Kaufmann	101101




Insert the following tuples into an extendible hash table whose buckets can hold up to two entries.

ID	Name	hash
10	Müller	010100
25	Meier	100110
30	Schmidt	011110
18	Krause	010010
40	Schulz	000101
45	Kaufmann	101101



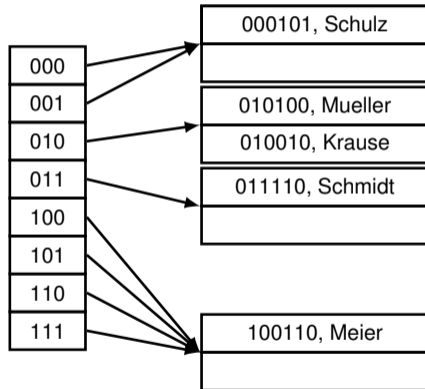
Insert the following tuples into an extendible hash table whose buckets can hold up to two entries.

ID	Name	hash
10	Müller	010100
25	Meier	100110
30	Schmidt	011110
18	Krause	010010
40	Schulz	000101
45	Kaufmann	101101




Insert the following tuples into an extendible hash table whose buckets can hold up to two entries.

ID	Name	hash
10	Müller	010100
25	Meier	100110
30	Schmidt	011110
18	Krause	010010
40	Schulz	000101
45	Kaufmann	101101

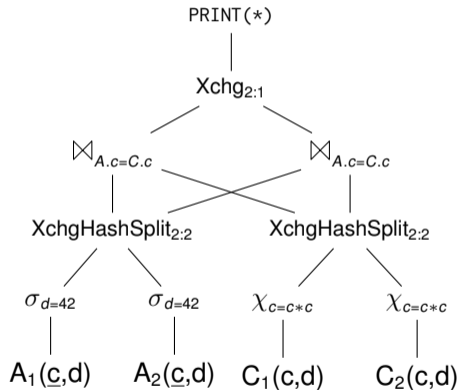



Insert the following tuples into an extendible hash table whose buckets can hold up to two entries.

ID	Name	hash
10	Müller	010100
25	Meier	100110
30	Schmidt	011110
18	Krause	010010
40	Schulz	000101
45	Kaufmann	101101



Parallelize the given query plan by introducing $XchgHashSplit(n:m)$ and $Xchg(n:m)$ operators for **2 threads** as necessary.



Questions?