

Tuple Reconstruction

Prof. Hasso Plattner

Tuple Reconstruction (I)

Accessing a record in a row store

Table: world_population


	First Name	Last Name	Gender	Country	City	Birthday
Row 1						
Row 2						
Row 3						
Row 4						
...						
Row 8 x 10 ⁹						

- All attributes are stored consecutively
- 200 byte → 4 cache accesses à 64 byte → **256 byte**
- Read with 4MB/ms/core → **≈ 0.064 μs** with 1 core



Tuple Reconstruction (2)

Virtual record IDs



- All attributes are stored in separate columns
- Implicit record IDs are used to reconstruct rows

HPI Hasso Plattner Institut
IT Systems Engineering | Universität Potsdam

Table: world_population

First Name	Last Name	Gender	Country	City	Birthday
1	1	2	1	2	1
2	2	3	2	3	2
3	3	...	3	4	3
4	4	8x10 ⁹	4	5	4
5	5		5	6	5
6	6		6	7	6
7	7		7	8	7
8	8		8	9	8
9	9		9	10	9
10	10		10	11	10
11	11		11	12	11
12	12		12	13	12
13	13		13	14	13
14	14		14	15	14
15	15		15	16	15
16	16		16	17	16
17	17		17	18	17
18	18		18	19	18
19	19		19	20	19
20	20		20	21	20
21	21		21	22	21
22	22		22	23	22
23	23		23	24	23
24	24		24	25	24
25	25		25	26	25
26	26		26	27	26
27	27		27	28	27
28	28		28	29	28
29	29		29	30	29
30	30		30	31	30
31	31		31	32	31
32	32		32	33	32
33	33		33	34	33
34	34		34	35	34
35	35		35	36	35
...	...		...	...	...
8x10 ⁹	8x10 ⁹				8x10 ⁹

Legend:  ≈ 1 cache access ≈ 64 byte

- 1 cache access for each attribute
- 6 cache accesses à 64 byte
→ **384 byte**
- Read with 4MB/ms/core
→ **$\approx 0.096 \mu\text{s}$**
with 1 core