

Weaving Relations for Cache Performance

Johannes Kern

Universität Tübingen

03.12.2010


Motivation I

Hard Disk ~~X~~ RAM



Motivation II

- **Key memory bottleneck:** Data cache misses
- 50-90% of total memory-related stall time



Remedy

- Create cache-conscious storage model
- Maximize spatial locality
- Maintain good I/O performance

Remedy

- Create cache-conscious storage model
- Maximize spatial locality
- Maintain good I/O performance

→ Partition Attributes Across (PAX)

Cache Evaluation

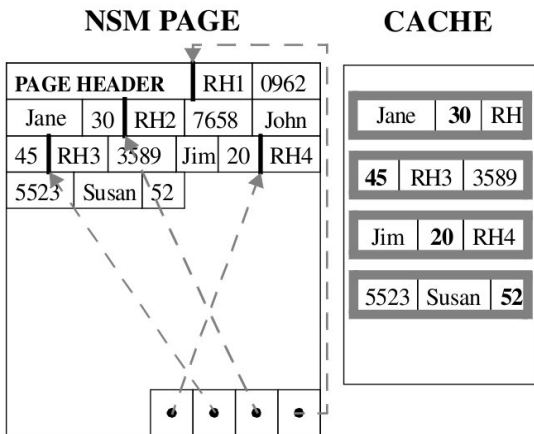
Relation R

SSN	name	age
0962	Jane	30
7658	John	45
3589	Jim	20
5523	Susan	52

Query

```
SELECT name
FROM R
WHERE age < 40;
```


N-ary Storage Model (NSM)



Inter-record spatial locality ✗

Low record reconstruction cost ✓


Decomposition Storage Model (DSM)




Inter-record spatial locality ✓

Low record reconstruction cost ✗

Partition Attributes Across (PAX)



CACHE



Inter-record spatial locality ✓

Low record reconstruction cost ✓

Evaluation

Characteristic	NSM	DSM	PAX
Inter-record spatial locality	✗	✓	✓
Low record reconstruction cost	✓	✗	✓

Setup

- Dell 6400 PII Xeon/MT
- Windows NT 4.0
- Pentium II Xeon 400Mhz
 - 16 kB split L1 cache
 - 512 kB unified L2 cache
 - 32 Byte cache blocks
- 512 MB main memory
- 100 Mhz system bus



Workload


Query

```
SELECT AVG(a)
FROM R
WHERE b > Lo
      AND b < Hi
```

- Eight 8 Byte attributes
- 1.2 million records




Record Reconstruction Costs



Memory Stall Time Analysis

- Cache misses
 - **NSM**: 1 per record
 - **PAX**: 1 every n records ($n = \frac{\text{cacheblocksize}}{\text{attributesize}}$)
- **PAX saves 75% L2 data cache misses compared to NSM**



Summary

- Data access to cache hierarchy major performance bottleneck
- NSM incurs negative effects on data cache performance
- DSM incurs high record reconstruction costs
- PAX combines the good attributes of NSM/DSM

PAX summary

PAX groups values for the same attribute together in minipages, combining inter-record spatial locality and high data cache performance with minimal record reconstruction cost at no extra storage overhead.