ICS 624 Spring 2011

Multi-Dimensional Clustering

Asst. Prof. Lipyeow Lim

Information & Computer Science Department

University of Hawaii at Manoa

Clustered vs Unclustered Index

- Suppose data records are stored in a Heap file.
  - To build clustered index, first sort the Heap file (with some free space on each page for future inserts).
  - Overflow pages may be needed for inserts. (Thus, order of data recs is `close to', but not identical to, the sort order.)

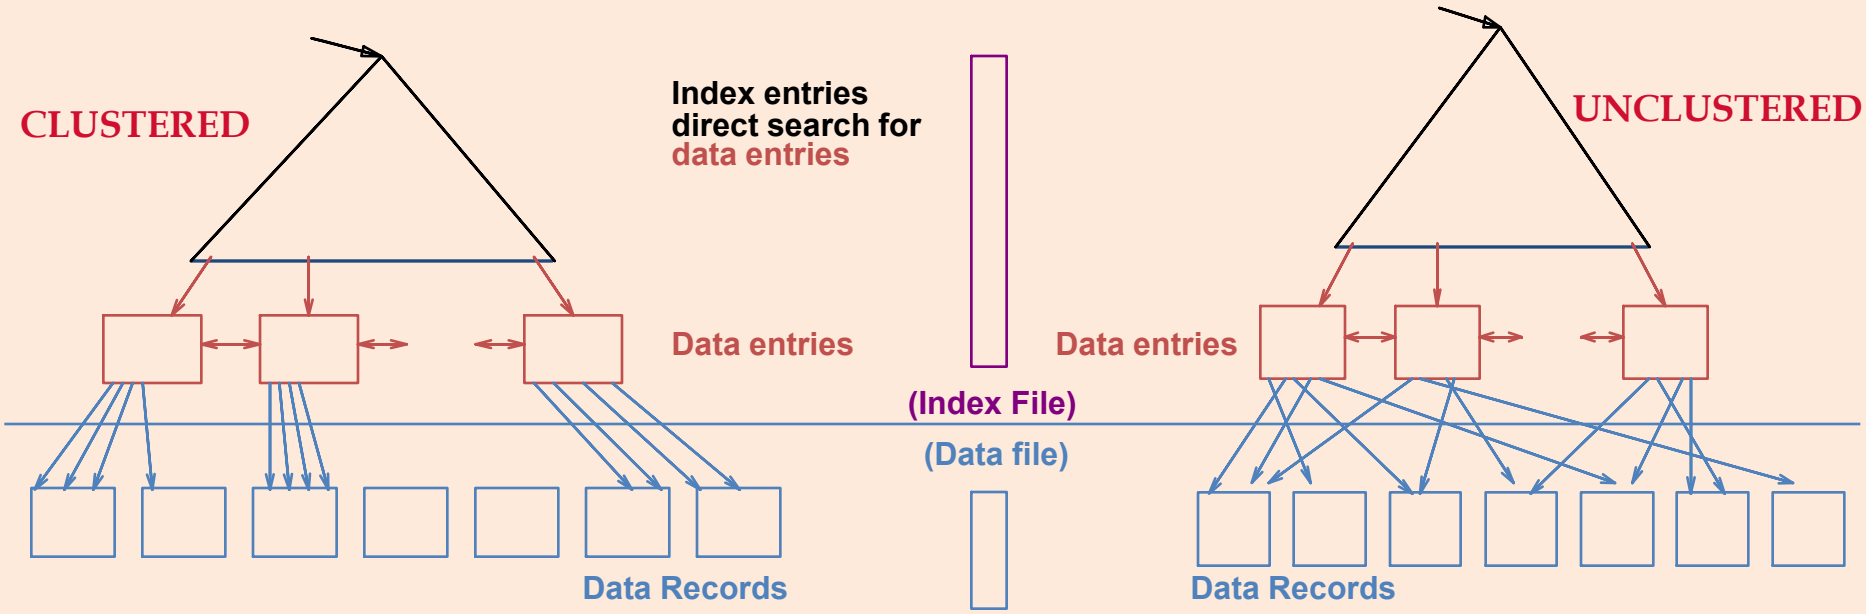

Accessing Data on Disk

- **Seek time**: time to move disk heads over the required track
- **Rotational delay**: time for desired sector to rotate under the disk head.
  - Assume uniform distribution, on average time for half a rotation
- **Transfer time**: time to actually read/write the data

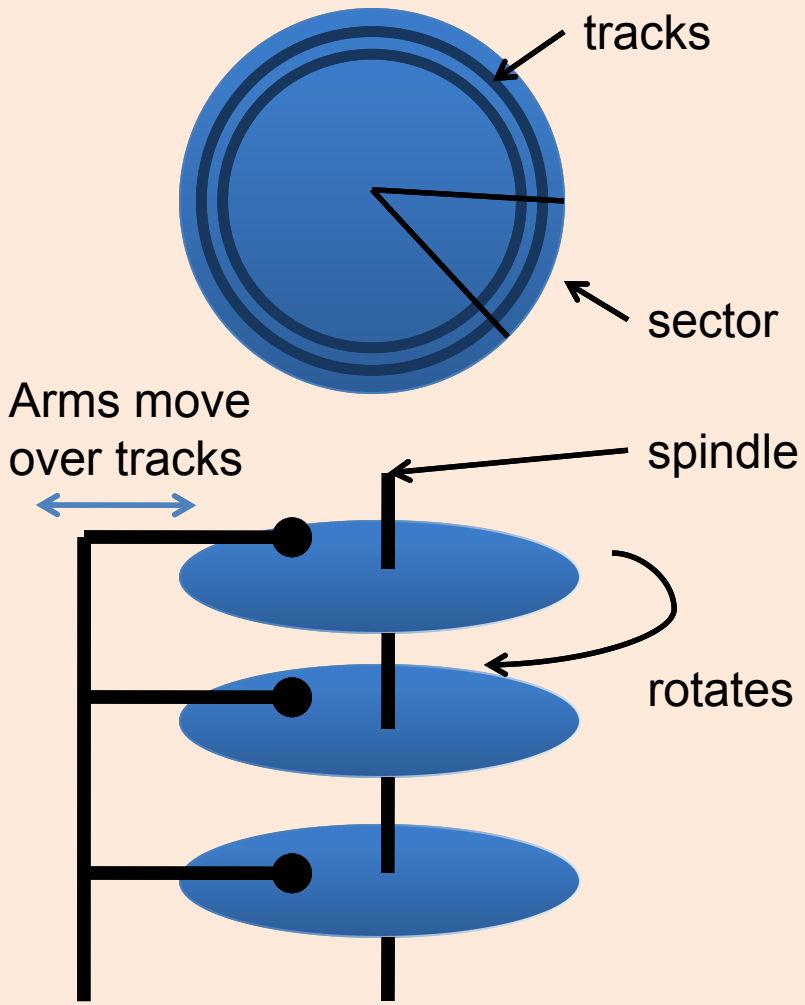

Example: Barracuda 1TB HDD (ST31000528AS)

- What is the average time to read 2048 bytes of data ?

= Seek time + rotational latency + transfer time

= 8.5 msec + 4.16 msec + (2048 / 512) / 63 * (60 000 msec / 7200 rpm)

= 8.5 + 4.16 + 0.265

cylinders	121601
Bytes/cylinder	16065*512
Blocks/cylinder	8029
Sectors/track	63
Heads	255
Sprindle Speed	7200 rpm
Average Latency	4.16 msec
Random read seek time	< 8.5 msec
Random read Write time	< 9.5 msec

Motivating Examples

```
SELECT  Sum(S.sales)
FROM  Sales S
WHERE S.yearOd=1997
```

```
SELECT  Sum(S.sales)
FROM  Sales S
WHERE S.region='Canada'
```

```
SELECT  Sum(S.sales)
FROM  Sales S
WHERE S.region='Mexico' AND
      S.itemID=2 AND
      S.yearOd=1997
```

- Suppose
  - Sales table is sorted by year
  - Index on region
  - Index on itemID
  - Index on yearOd
  - Index entries are <key,rids>

Multi-Dimensional Clustering (MDC)

- DB2 v8 and above.
- Physical layout mimics a multi-dimensional cube
- Associates a physical region called blocks for each unique combination of dimension attribute values.
- These blocks are the units of addressability for our clusters.
- A block index that addresses these blocks.

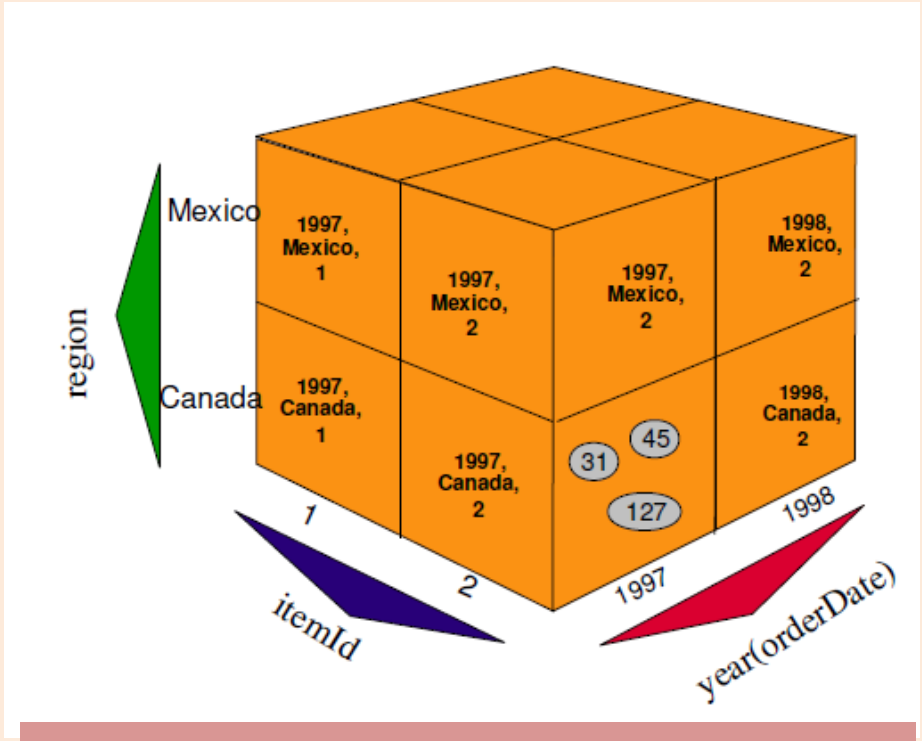


```
CREATE TABLE Sales(
int storeId, date orderDate, int region,
int itemId, float price, int yearOd
generated always as year(orderDate))
ORGANIZE BY DIMENSIONS
(region, yearOd, itemId)
```

See *Multi-Dimensional Clustering: A New Data Layout Scheme* in DB2. SIGMOD 2003: 637-641

MDC Storage Layout

Conventional Heap File

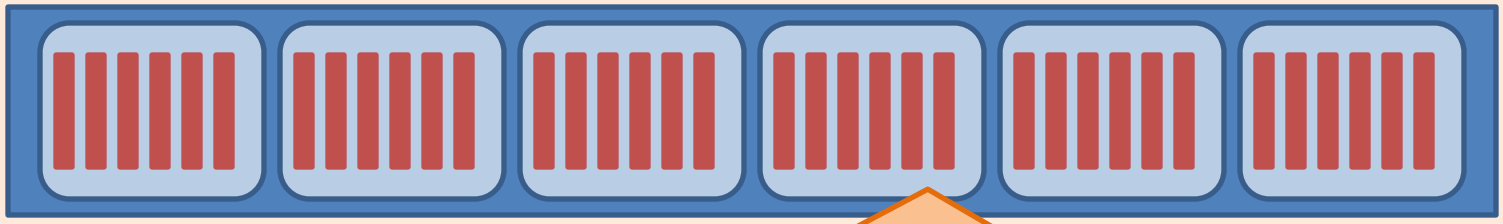

Tuples and pages sorted by eg. time

MDC Storage Layout

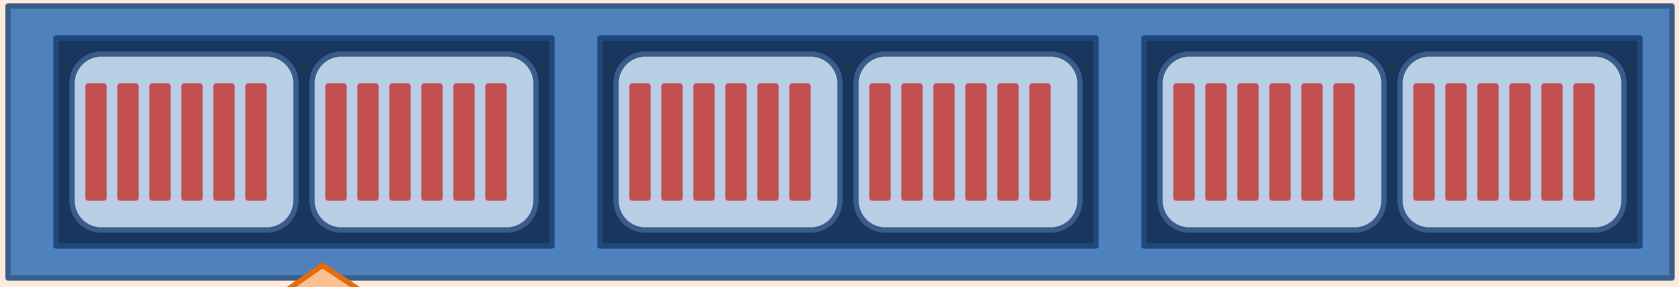

- Pages grouped into blocks by multi-dimensional values
- Blocks sorted by multi-dimensional values.

Dimension Block Index

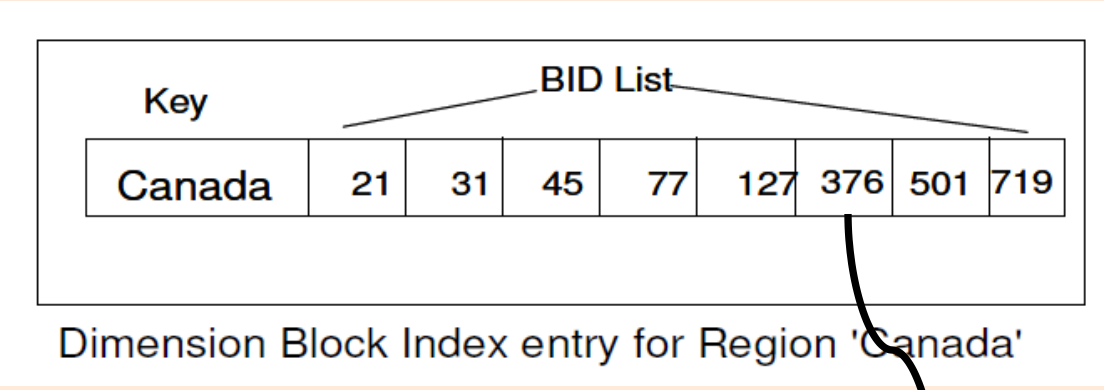

MDC Storage Layout

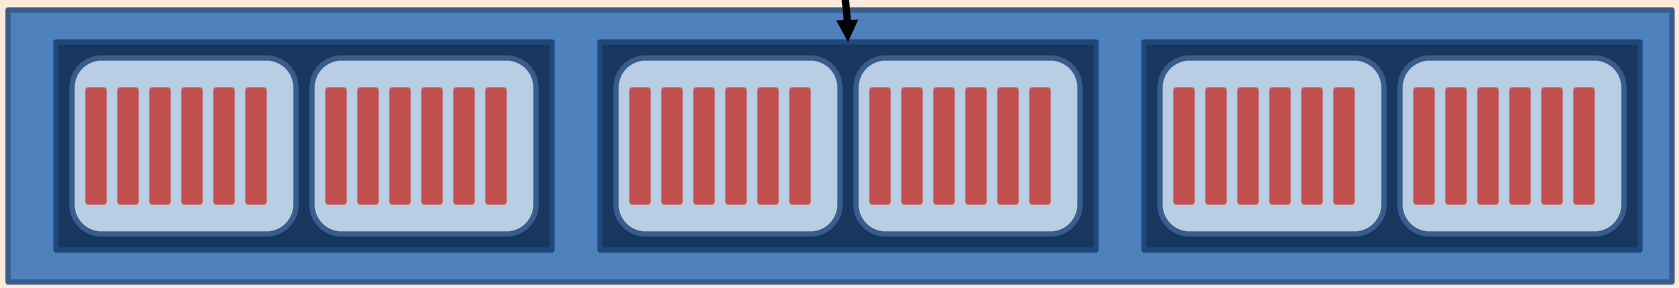

- Regular B-tree indexes except that data entries contain <key,bid list>
- One index for each dimension

SELECT Sum(S.sales)
FROM Sales S
WHERE S.region='Canada'

Issues & Questions

- Choosing the MDC dimensions
- Overhead in maintaining the additional level of indirection
- Maintaining the clustering with updates