

Physical DB Issues, Indexes, Query Optimisation

Database Systems Lecture 13
Natasha Alechina

In This Lecture

- Physical DB Issues
 - RAID arrays for recovery and speed
 - Indexes and query efficiency
- Query optimisation
 - Query trees
- For more information
 - Connolly and Begg chapter 21 and appendix C.5

Physical Design


- Design so far
 - E/R modelling helps find the requirements of a database
 - Normalisation helps to refine a design by removing data redundancy
- Physical design
 - Concerned with storing and accessing the data
 - How to deal with media failures
 - How to access information efficiently

RAID Arrays

- RAID - redundant array of independent (inexpensive) disks
 - Storing information across more than one physical disk
 - Speed - can access more than one disk
 - Robustness - if one disk fails it is OK
- RAID techniques
 - Mirroring - multiple copies of a file are stored on separate disks
 - Striping - parts of a file are stored on each disk
 - Different levels (RAID 0, RAID 1...)


RAID Level 0

- Files are split across several disks
 - For a system with n disks, each file is split into n parts, one part stored on each disk
 - Improves speed, but no redundancy



RAID Level 1

- As RAID 0 but with redundancy
 - Files are split over multiple disks
 - Each disk is mirrored
 - For n disks, split files into n/2 parts, each stored on 2 disks
 - Improves speed, has redundancy, but needs lots of disks



Parity Checking

- We can use parity checking to reduce the number of disks

- Parity - for a set of data in binary form we count the number of 1s for each bit across the data

- If this is even the parity is 0, if odd then it is 1

```

1 0 1 1 0 0 1 1
0 0 1 1 0 0 1 1
1 0 1 0 1 0 0 1
0 1 1 0 1 1 1 0
-----
0 1 0 0 0 1 1 1
    
```

Recovery With Parity

- If one of our pieces of data is lost we can recover it


- Just compute it as the parity of the remaining data and our original parity information

```

1 0 1 1 0 0 1 1
0 0 1 1 0 0 1 1
0 1 1 0 1 1 1 0
-----
0 1 0 0 0 1 1 1
    
```

RAID Level 3

- Data is striped over disks, and a parity disk for redundancy
- For n disks, we split the data in n-1 parts
- Each part is stored on a disk
- The final disk stores parity information



Other RAID Issues

- Other RAID levels consider
 - How to split data between disks
 - Whether to store parity information on one disk, or spread across several
 - How to deal with multiple disk failures
- Considerations with RAID systems
 - Cost of disks
 - Do you need speed or redundancy?
 - How reliable are the individual disks?
 - 'Hot swapping'
 - Is the disk the weak point anyway?

Indexes

- Indexes are to do with ordering data
 - The relational model says that order doesn't matter
 - From a practical point of view it is very important
- Types of indexes
 - Primary or clustered indexes affect the order that the data is stored in a file
 - Secondary indexes give a look-up table into the file
 - Only one primary index, but many secondary ones

Index Example

- A telephone book
 - You store people's addresses and phone numbers
 - Usually you have a name and want the number
 - Sometimes you have a number and want the name
- Indexes
 - A clustered index can be made on name
 - A secondary index can be made on number

Index Example

As a Table

Name	Number
John	925 1229
Mary	925 8923
Jane	925 8501
Mark	875 1209

Order does not really concern us here

As a File

Jane, 9258501	8751209
John, 9251229	9251229
Mark, 8751209	9258501
Mary, 9258923	9258923

Most of the time we look up numbers by name, so we sort the file by name

Secondary Index


Sometimes we look up names by number, so we index number

Choosing Indexes

- You can only have one primary index
 - The most frequently looked-up value is often the best choice
 - Some DBMSs assume the primary key is the primary index, as it is usually used to refer to rows
- Don't create too many indexes
 - They can speed up queries, but they slow down inserts, updates and deletes
 - Whenever the data is changed, the index may need to change


Index Example

- A product database, which we want to search by keyword
 - Each product can have many keywords
 - The same keyword can be associated with many products



Index Example

- To search the products given a keyWord value
 - We look up the keyWord in Keywords to find its keyID
 - We look up that keyID in WordLink to find the related prodIDs
 - We look up those prodIDs in Products to find more information about them



Creating Indexes

- In SQL we use **CREATE INDEX:**

```
CREATE INDEX
<index name>
ON <table>
(<columns>)
```
- Example:


```
CREATE INDEX keyIndex
ON Keywords (keyWord)
CREATE INDEX linkIndex
ON WordLink(keyID)
CREATE INDEX prodIndex
ON Products (prodID)
```

Query Processing

- Once a database is designed and made we can query it
 - A query language (such as SQL) is used to do this
 - The query goes through several stages to be executed
- Three main stages
 - Parsing and translation - the query is put into an internal form
 - Optimisation - changes are made for efficiency
 - Evaluation - the optimised query is applied to the DB

Parsing and Translation

- SQL is a good language for people
 - It is quite high level
 - It is non-procedural
- Relational algebra is better for machines
 - It can be reasoned about more easily
- Given an SQL statement we want to find an equivalent relational algebra expression
- This expression may be represented as a tree - the query tree

Relational Operators

- Product \times
 - Product finds all the combinations of one tuple from each of two relations
 - $R1 \times R2$ is equivalent to


```
SELECT DISTINCT *
FROM R1, R2
```
- Selection σ
 - Selection finds all those rows where some condition is true
 - $\sigma_{\text{cond}} R$ is equivalent to


```
SELECT DISTINCT *
FROM R
WHERE <cond>
```

Relational Operators

- Projection π
 - Projection chooses a set of attributes from a relation, removing any others
- $\pi_{A1,A2,\dots} R$ is equivalent to


```
SELECT DISTINCT
A1, A2, ...
FROM R
```
- Projection, selection and product are enough to express queries of the form

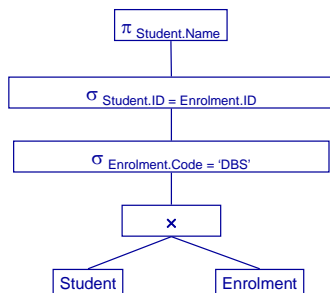

```
SELECT <cols>
FROM <table>
WHERE <cond>
```

SQL \rightarrow Relational Algebra

- SQL statement


```
SELECT Student.Name
FROM Student,
      Enrolment
WHERE
  Student.ID =
  Enrolment.ID
AND
  Enrolment.Code =
  'DBS'
```
- Relational Algebra
 - Take the product of Student and Enrolment
 - select tuples where the IDs are the same and the Code is DBS
 - project over Student.Name

Query Tree



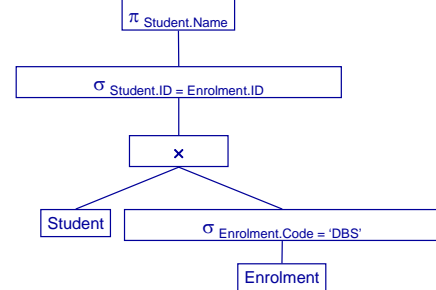
Optimisation

- There are often many ways to express the same query
- Some of these will be more efficient than others
- Need to find a good version
- Many ways to optimise queries
 - Changing the query tree to an equivalent but more efficient one
 - Choosing efficient implementations of each operator
 - Exploiting database statistics

Optimisation Example

- In our query tree before we have the steps
 - Take the product of Student and Enrolment
 - Then select those entries where the Enrolment.Code equals 'DBS'
- This is equivalent to
 - selecting those Enrolment entries with Code = 'DBS'
 - Then taking the product of the result of the selection operator with Student

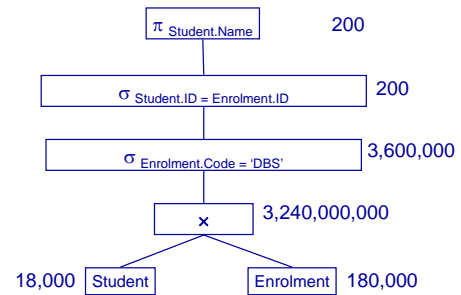
Optimised Query Tree



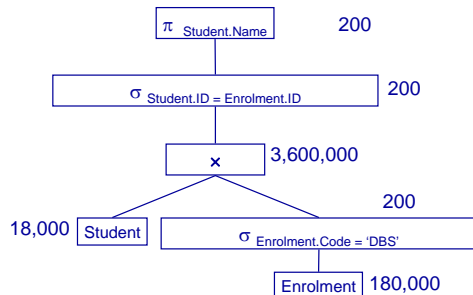
Optimisation Example

- To see the benefit of this, consider the following statistics
 - Nottingham has around 18,000 full time students
 - Each student is enrolled in at about 10 modules
 - Only 200 take DBS
- From these statistics we can compute the sizes of the relations produced by each operator in our query trees

Original Query Tree



Optimised Query Tree



Optimisation Example

- The original query tree produces an intermediate result with 3,240,000,000 entries
- The optimised version at worst has 3,600,000 entries
- A big improvement!
- There is much more to optimisation
 - In the example, the product and the second selection can be combined and implemented efficiently to avoid generating all Student-Enrolment combinations

Optimisation Example

- If we have an index on Student.ID we can find a student from their ID with a binary search
- For 18,000 students, this will take at most 15 operations
- For each Enrolment entry with Code 'DBS' we find the corresponding Student from the ID
- $200 \times 15 = 3,000$ operations to do *both* the product and the selection.

Next Lecture

- Database Security
 - Aspects of security
 - Access to databases
 - Privileges and views
- Database Integrity
 - View updating, Integrity constraints
- For more information
 - Connolly and Begg chapters 6 and 19