[07/09/2011] CS 344 Notes Mayank Takker (08012325)

An overview of what will be covered in course in the rest of the semester was given. Data can be of three types: structured, unstructured and semi-structured. Databases are used to store structured data. Till now we have covered database designing - *ER and relational model, conversion from one model to another and normalization*. We have also seen how SQL is used to define schemas, update them and query for data. During all this, the DBMS was seen as a blackbox which stores stores schemas (and data) and replies to SQL queries. Now we plan to open up this blackbox and see how data is internally stored and what algorithms are used to answer queries.

Algorithms, data structures, computer architecture and OS are the pre-requisite to database course. The instructor gave a few examples of how each of these courses are relevant to the study of databases.

Computer Architecture: Though registers and cache are the fastest memory available, they are also the costliest and hence secondary storage like hard drives are used to store large databases. When data is fetched from secondary memory, it is cached in primary memory for performance gain. Since not all data can be stored in primary memory, various replacement policies like LRU, MRU, LFU, etc are used. Databases are designed keeping in mind the constraints imposed by architecture and also to optimally utilize the available resources.

Algorithms: A few searching algorithms and their relevance according to the context was briefly discussed. Searching in a tree can take at most linear time ('*n*' is the no of nodes), but requires logarithmic time in a balanced binary search treee. Similarly given an array whether we use binary search or linear search also depends on the situation. Binary search takes logarithmic time but requires array to be sorted (which itself requires *O(n.logn)* time). In case we need to search only once, we would prefer linear search over sorting the array and then using binary search.

Data Structures: Indexing is used to speed up search and answer conditional queries. Questions like - *Which attribute or set of attributes to use for indexing? What type of indexing – tree/bit/hash based to use? What data structure to use to store data?* will be answered in the later half of the database course. [Refer: http://en.wikipedia.org/wiki/B-tree]

Operating Systems: Several problems encountered in databases such as *atomicity* of transactions and managing *concurrent transactions* are also seen in OS. When we execute an update, we want that either the whole update operation should happen or nothing should happen. If certain steps of a transaction have been executed and the next step fails, there must be a mechanism to rollback ie undo the executed steps and flag an error. During concurrent transactions we need to avoid race conditions and deadlocks. Mechanisms similar to the ones used in OS are also used in DB to tackle these problems.