

Module_5

SQL & ER

CS 3410 DB

Introduction and Background

- In the modern world, corporations find themselves faced with a common problem: the lack of efficiency with record keeping. This, amongst other reasons, is why we are here today to introduce to you the language of data known as SQL.
 - SQL originated in the 1970's when Dr. Edgar F. Codd created the relational model of database management.
 - A special characteristic of SQL is that it is a 4th generation language meaning it falls more closely in line with that of the English language and therefore is more easily understood.
 - The first ANSI SQL language standards were published in 1986 and have been consistently updated since (the most notable of these changes being the one featured in 1992)

How has SQL changed over the years?

- In 1992 there was a significant change featuring a structural overhaul which would now entail 3 levels: Entry, Intermediate, and Expert
- SQL now has 4 major components which enable its function:
 - Query Dispatcher
 - Optimization Engines
 - Classic Query Engines
 - SQL Query Engine

Through the use of these components, a user can interact with the system by using commands which SQL specifies. This includes but is not limited to CREATE, DELETE, DROP, SELECT, and INSERT.

History of SQL

- SQL stands for Structured Query Language, and it is a non-procedural language used in an RDBMS.
- It is based on relational algebra and calculus expressions that can evaluate truth on and over relations.
- SQL became a particularly popular choice within databases when Chamberlin and Boyce, the creators of SQL, implemented relational operators into the language so as to compactly represent complex queries.
- SQL was also special in that it did not require the user to specify how to reach the data they wished to access with their query.
- SQL's unique capabilities: ease of use, high accessibility, and good readability allowed it to rise to common use, and as it stands now, there are no good alternatives to SQL.
- SQL is considered the standard language for relational database's by ANSI and ISO.

Unary Operators

- According to Ted Codd, a database is a set of relations that are often represented by tables that contain rows and columns
- Unary operators act on a single instance of a relation (one table)
- Project, Select, and Rename are the most popular and useful unary operators to know
- Projection will give the user the specified column(s) that they wish to extract from the tables
- Selection will give the user the rows that satisfy conditions they set for the data within the table
- Rename will allow the user to change the names of relations, entities, and attributes
- Composing these operators is possible, but it can cause errors if done incorrectly
- Examples follow on the next slide

Examples of Unary Operators

Name	ID	Class
Cam	1414	Bio
Ak	2976	Lit

Projection query:

```
SELECT Name, Class
```

```
FROM Table
```

This query will extract the Name Column and Class Column from the specified table named Table. This is an example of projection because we are obtaining desired columns and deleting undesirable ones.

Selection query:

```
SELECT *  
FROM Table  
WHERE class = 'Bio'
```

This query will select all of the columns in the table to evaluate. After obtaining these columns, it will only return the rows where the student's class is Bio. This is selection because we are extracting desired rows while removing ones that do not meet our requirements.



Binary Operators

- Oftentimes, databases will have multiple relation instances, multiple tables
- Chamberlin and Boyce also devised binary operators: join, union, intersect, and set difference
- Join will combine two relations based on a predicate (JOIN)
- Union will combine two relations that are compatible (UNION)
- Intersect will find and return the unique rows from each relation (INTERSECT)
- Set difference will return the unique rows in only the first relation (MINUS)
- It is important to remember that an error can be thrown if the queried relations are not compatible
- If relation 1 has a name and class while relation 2 has a name and ID, attempting to union or set difference the two relations will throw an error

Examples of Binary Operators

Name	ID	Class
Cam	1414	Bio
Ak	2976	Lit

Example of UNION

```
SELECT Name
FROM TableOne
UNION
SELECT Name
FROM TableTwo
```

This query will first get the Name column from each of the tables. Afterwards, it will join the Name columns of the two tables to get the resulting table. Note that Name is the only column that will be affected by the union operator, so the queried relations compatible.

Name	Range	Class
Shin	99	Bio
Navi	180	Lit

Example of Set Difference

```
SELECT Name
FROM TableOne
MINUS
SELECT Name
FROM TableTwo
```

This query will minus the Name columns of the two tables to get the resulting table. Note that Name is the only column that will be affected by the union operator, so the queried relations are compatible.



KENNESAW STATE
UNIVERSITY

SQL Triggers

In 2004, SQL released the update 5.0.2 which introduced triggers among other features.

Triggers were designed to make queries much more powerful. With a trigger, one can set a condition that needs to be met in order to perform certain actions.

Instead of writing the same query for each instance, SQL now has listeners to watch out for certain criteria.

```
CREATE TRIGGER TrigName [BEFORE|AFTER] [INSERT|UPDATE|DELETE] ON tableName
```

PHP

It was initially released in 1994 and has since been expanded upon and is now the most popular scripting language for MySQL databases.

Its relevance comes from being able to write scripts within the database to run queries.

This eases the overhead on the front end making the entire service more efficient.

```
$sql =INSERT INTO data VALUES('ashok','you are awesome')";
```

PL/SQL

PL/SQL is another scripting language released in 1989 by Oracle.

Specifically designed with SQL in mind, PL/SQL provides an incredibly straightforward way to incorporate automation in your database.

Moving automation and triggers onto the database leaves less work for the front end of the service as less information needs to travel back and forth from client to server thus reducing network traffic.

PL/SQL Cont.

PL/SQL is designed to be written in blocks.

In this example we see a simple way to incorporate an SQL statement in a PL/SQL block.

PL/SQL also has its own built-in functions like DBMS_OUTPUT that do not require SQL to write.

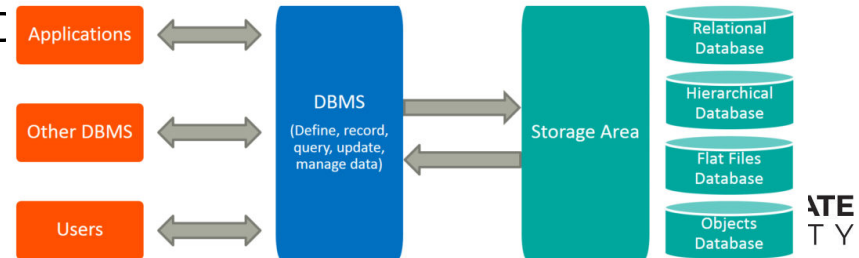
```
DECLARE
    l_name  employees.last_name%TYPE;
BEGIN
    SELECT last_name
       INTO l_name
      FROM employees
     WHERE employee_id = 138;

    DBMS_OUTPUT.put_line (l_name);
END;
```

Database Management Systems

- The most popular DB management systems today are:
 - Oracle – Ranked the highest out of the ones listed due to its portability and functionality
 - Microsoft SQL - Relational DB management system developed by microsoft (useful for those involved in the Microsoft stack)
 - MySQL – Another Relational DB management system which is open source and made by Allan Larsson
 - MongoDB – A document-oriented [

Database Management System



Defining a Database: Data Definition Language

- There are 4 major categories which are used to define a database:
 - Data Definition Language – This is what's used to define the data, database objects and schema within the database. This is done through the use of command such as CREATE, DROP, ALTER, TRUNCATE, COMMENT, and RENAME. The main utility for these commands would be for large generic changes to your DB, or for startup of your DB.
 - As you can see before we added a salary column by using the "ALTER TABLE

EmployeeName	EmployeeID	Position
John	123	Engineer
Jeff	456	Accountant
Jason	789	QA

EmployeeName	EmployeeID	Position	Salary
John	123	Engineer	
Jeff	456	Accountant	
Jason	789	QA	



KENNESAW STATE
UNIVERSITY

Data Query Language and Data Control Language

- DQL - This category of language is actually used to perform manipulations to the data within schemas rather than the schemas themselves as mentioned above. The primary example of a DQL is the select command which selects a set of data which you want to perform your operation on and does this by retrieving the data from the database (Varshini, D., 2019, August 26)
- DML - This category is mostly just concerning the rights and permissions of other users on the network database and has commands which manage these things. These commands such as GRANT and REVOKE simply grant or revoke permissions to a specified user on the database system. This is important for both data security and data integrity. It ensures that data can only be seen by people who must see it, and it also ensures data cannot be changed by someone who does not have an understanding of what is happening in the database. Otherwise, information can be incorrect or even leaked.

Data Manipulation Language

- This is different from DQL because DQL only takes into consideration data from the schema object. It is also worth noting that this category includes most of the commonly used SQL statements such as INSERT, DELETE, and UPDATE which each has a pretty intuitive utility corresponding to its name. These allow users to actually modify data in small increments, rather than make sweeping changes that would be seen in DDL.

- Example:**

Employee Name	Employee ID	Position	Salary

After Command:
INSERT INTO
Employee(EmployeeName,EmployeeID,Position,Salary)

VALUES ("Jonah", 159,
Janitor, 1000000)

EmployeeName	EmployeeID	Position	Salary
Jonah	159	Janitor	1000000



KENNESAW STATE
UNIVERSITY

5.3 Write Single Table Queries Using SQL

- Query – A call for a specific set, group or combination of data. To query in a database one must use a language that the database understands.
- Table – Objects within a database that include some, or all, of the data from the database. Data is organized into rows and columns.
- Rows – Represent unique item
- Columns – Represent different attributes the item contains

5.3 Student Table

Student (Kashefi 2020)

StudentName	StudentID	ClassID	Grade
John	123	987	90
Jeff	456	654	83
Jason	789	321	97



5.3 Continued

- The example table labeled Students illustrates the use of rows and columns. The first column holds the student's name which will be used to represent unique items in the rows.
- The columns in each row contain unique information specific to that row. This makes each row unique with its own unique data.
- We use SQL commands to query the table such as SELECT and FROM commands.
- SELECT allows us to choose the columns we will be querying from.
- FROM allows us to choose which table to query.



5.3 Commands

- `SELECT * FROM Student`, allows entire table to be asked to be queried. The asterisk is shorthand way to ask for every column.
- `SELECT StudentName, Grade FROM Student`, This will tell us the name and grade for every student.
- `SELECT StudentName, Grade FROM Student WHERE Grade < 95`, `WHERE` will only return a row if the information in one of the specific columns fits a specific condition of the query.

5.4 Establish Referential Integrity using SQL

- Referential integrity - The accuracy and consistency of data in a relationship. Referential integrity requires that a foreign key references a primary key.
- Primary key - The key or specific column in a parent table
- Foreign key - The key in a child table that references the primary key

5.4 Child and Parent Tables

Child Table

Department	
Employee ID	Department
6789	Marketing
5632	R&D

Parent Table

Employee		
Employee ID	Age	Salary
6789	25	56000
5632	29	83000

5.4 Continued

- Primary key is Employee ID and it is referenced in the Child table, making it a foreign key in this table.
- As a result of this relationship a user will be prevented from:
- Adding information in the child table if the same information is not also in the parent table
- Changing data in the primary table that result in parentless keys in the child table(orphaned)
- Deleting records from the parent table if they exist in the child table.

5.4 Conclusion

- A lack of Referential integrity can result in records being lost and/or inaccurate or confusing.
- Therefore it is imperative to have Referential integrity or else this lack of can have negative repercussions for those that use database systems.

Resources

- Raza, M. (2018, August 29). What is a DBMS? Database Management Systems Explained. Retrieved from <https://www.bmc.com/blogs/dbms-database-management-systems/>

Introduction

- Physical database design is the process of transforming logical data models into physical data models.
- Conceptual data modeling is a map of concepts and their relationships used for databases.
- The relational model is one of the most commonly used models in contemporary database applications.
- The principles of logical database design for the relational model apply to many other logical models as well.

Background

- The first database was developed in the 1960s when computers were mostly used for private organizations.
- The two most popular data models during the 1960s were the CODASYL (Conference/Committee on Data Systems Languages) and IMS (Information Management System).
- These database systems began to change during the 1970s when E.F Codd published a paper on his revolutionary ideas about the relational model/database.

Background Continued...

- The ER(Entity-Relationship) Diagram was developed in 1976. By P. Chen. The ER module allows developers to focus less on the logic table structure and more on research data application.
- In 1980, SQL(Structured Query Language) became the standard query language among databases.
- Before 1980, Government organizations were the first ones to invest heavily into security of data.
- Today, research is still being conducted on database security and continues to evolve year to year.

Logical Design vs Physical Design

- The logical design is made up of several characteristics such as: Entity, Relationship, Attribute(s), and a Unique Identifier.
- The physical design consists of a table, foreign key, column(s), and a primary key.

[4.2.1] Physical Database Design Process

- The design of a physical database design is heavily influenced on integrity and performance.
- According to Adrienne Watt, database design starts with a conceptual data model and produces a specification of a logical schema.
- The database design process is initialized from the logical data model that will be used in the database design and can be represented as an E-R(Entity-Relationship) diagram.
- Physical database design is concerned with the design of fields. A field is the smallest unit of application data recognized by system software.

[4.2.2] Data Partitioning

- Partitioning is a concept in databases in which very large tables and data are partitioned into smaller, individual tables, and queries.
- Horizontal partitioning is the classification of the rows based on common characteristics into several, separate tables.
- In range partitioning, each partitioned portion is characterized by a range of values for one or multiple columns such as IDs, or dates.
- Hash partitioning is the spreading of data in even partitions autonomous of the key value.
- List partitioning is a technique where a list of distinct values is defined as the partitioning key in the characterization for each partition

Pros of Data Partitioning

- Partitioning is practical and helps manage the table because partitioning helps identify the area where maintenance is needed and saves storage space.
- Partitioning is also secure as only the relevant and necessary data can be specifically chosen and accessed by the user.
- Backing up and securing files is easier due to their smaller size and if one file is corrupted, the other is still accessible.
- Partitioning also helps with balancing the load. The partitioned files can be designated to different storage locations which reduces conflict and maximizes performance.

Cons of Data Partitioning

- Partitioning is inconsistent with the access speed. Due to all partitions not being identical, the access speeds tend to differ.
- Due to the complex nature of partitions, the code required to program will need to be more complex, and challenging.
- Partition takes up excess storage space and time.

[4.3.1] Describe three types of file organization

File Organization

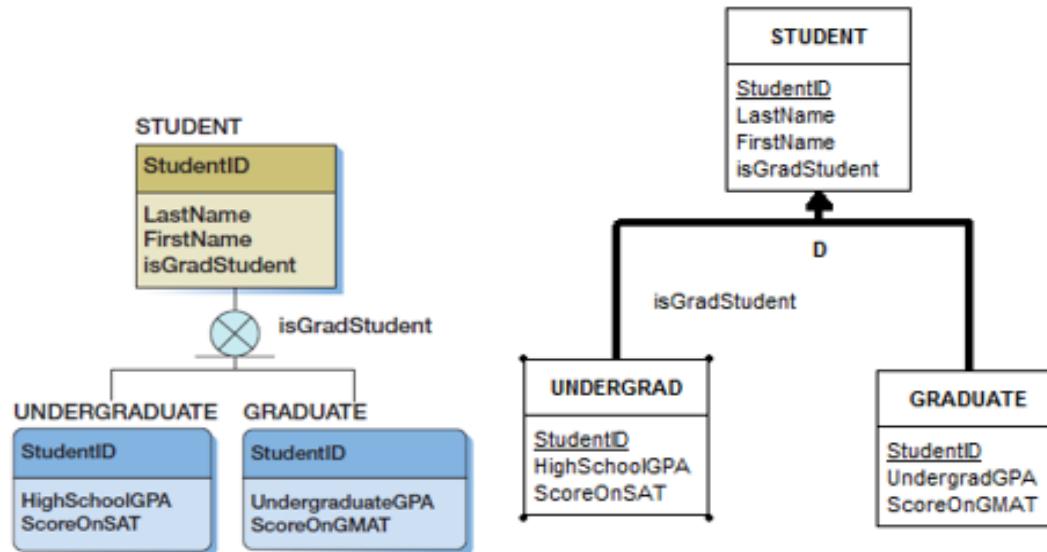
- According to (Venkataraman, R., Topi, H. 2011) a “*file organization* is a technique for arranging the records of a file on secondary storage devices.”
- With modern relational DBMS it is not necessary to design file organizations, but you are to be allowed to select an organization and its parameters for a table or physical file.
- In choosing a file organization for a particular file in a database consider seven important factors: Fast data retrieval, high throughput for processing data input and maintenance transactions, Efficient use of storage space, Protection from failures or data loss, minimizing need for reorganization, accommodating growth and security from unauthorized use





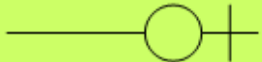
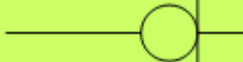
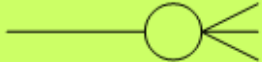





SQL and ER

- It is very important to remember that a database is a model of a user's view of the world. The only question is "How well does it fit the mental models of the people who are going to use the database system?"
- It is up to database administrators to create a SQL and Entity-Relationship (E-R) platform and provide the needs to fit user requirements.
- ER-Assistant provides relationships that are expressed using a different notation. It is easy to use but the entity boxes cannot be resized, leading to text length limitations.




- Erwin uses **solid vs. dashed lines** for **M:N relationships**; this can only be specified on “**children**” in a relationship.
- These features mean that nearly any E-R diagram created using Erwin will be incorrect for this text. Figure_1 below illustrates the differences between the notation used in the text and the notation used by Erwin.



Symbol Used In Database Concepts	ERWin Symbol	Meaning
		One - Mandatory
		Many - Mandatory
		One - Optional
		Many - Optional
		Exclusive Subtype
		Inclusive Subtype



Database Drawing such as Dia and Visio 2016 display names and roles of relationships and do not distinguish between weak and strong entities (no rounded corners). Connecting lines are solid:




Information Engineering (IE) began with the work of Clive data modeling tools which is one of the most popular notations for *database* design using with a *crow's* foot.

Give examples of M-M, M-O, O-M, and O-O relationships (other than those presented in this chapter). Draw two E-R diagrams for each of your examples: one using the traditional diamond notation and one using **IE Crow's Foot** notation.

In the Real Estate Agency example in question 5.5, each AGENT must use an agency car when on agency business. Further, to keep costs down the agency keeps exactly enough cars for the agents. Therefore, each AGENT must have a CAR, and each CAR must be assigned to an AGENT. This is an M-M relationship.



In the Real Estate Agency example in question 5.5, each CLIENT must be assigned to an AGENT, but there may be AGENTs who currently have no CLIENTs. This is an M-O (same as O-M, but seen reversed) relationship.



The E-R Crow's Foot model above is based on the model in question C but adds the entity ACTOR. Since there are no additional attributes needed, this can be modeled as an N:M relationship. The data for the new parts of the model are contained in the following table:

RELATIONSHIP			CARDINALITY [Blue = Inferable]	
PARENT	CHILD	TYPE	MAX	MIN
ACTOR	MOVIE	Strong	N:M	O-O

Sequential & Indexed File Organization

- **Sequential File Organization**
- **In a sequential file organization, the records in the file are stored in sequence according to a primary key value.**
- **To locate a particular record, a program must normally scan the file from the beginning until the desired record is located. A common example of a sequential file is the alphabetical list of persons in the white pages of a telephone directory.**
- **A comparison of the capabilities of sequential files with the other two types of files can be seen in figure 1.3. “Because of their inflexibility, sequential files are not used in a database but may be used for files that back up data from a database.” (Venkataraman, R., Topi, H. 2011)**

Indexed File Organization

- The records are stored either sequentially or not sequential, and an index is created that allows the application software to locate individuals. “A card catalog in a library, an *index* is a table that is used to determine in a file the location of records that satisfy some condition.”(Venkataraman, R., Topi, H. 2011) Each index entry matches a key value with one or more records.
- An index can point to unique records or to potentially more than one record. According to (Venkataraman, R., Topi, H. 2011) “an index that allows each entry to point to more than one record is called a *secondary key index*.” Secondary key indexes are important for supporting many reporting requirements and for providing rapid ad hoc data retrieval.
- An example would be an index on the ProductFinish column of a Product table. Because indexes are extensively used with relational DBMSs, and the choice of what index and how to store the index entries matters greatly in database processing performance.

Hash File Organization

- To determine or compute the address of a record within a file is to a hash file organization can be used and implemented as algorithm or function.
- A hash algorithm is an algorithm that takes an input of random size and proceeds to transform the input such that the hash result is an output of fixed length.
- Once the output is determined or computed, the hash result is irreversible, meaning that the algorithm can only process data in one-way.
- The use of hashing algorithms is commonly found in databases for practically any website that requires a password to login to an account and is illustrated in Figure 1.4.

Clustering File Organization:

- Defining a table to be in only one cluster reduces retrieval time for only those tables stored in the same cluster.
- This technique of file organization is known as clustering files and is illustrated in Figure 1.5.

Hash File Organization

```
CREATE CLUSTER Ordering (CustomerID CHAR(25));
```

The term Ordering names the cluster space; the attribute CustomerID specifies the attribute with common values.

Then tables are assigned to the cluster when the tables are created, as in the following example:

```
CREATE TABLE Customer_T(  
    CustomerID          VARCHAR2(25) NOT NULL,  
    CustomerAddress      VARCHAR2(15)  
)  
    CLUSTER Ordering (CustomerID);  
CREATE TABLE Order_T (  
    OrderID              VARCHAR2(20) NOT NULL,  
    CustomerID            VARCHAR2(25) NOT NULL,  
    OrderDate             DATE  
)  
    CLUSTER Ordering (CustomerID);
```



Security

- *Database files are stored in a proprietary format by the database which allows for access controls over the files.*
- *A useful procedures to consider are backups to ensure that stored data may be retrieved in the event that data may be compromised.*
- *Another technique employs the utilization of encryption to encrypt data contained within files and allow for only programs with access to decrypt the encrypted files to read them.*
- *Encryption involves two methods of encryption: symmetric encryption and asymmetric encryption. Symmetric encryption makes use of a single key for all parties communicating and is used for both encrypting data and decrypting data.*
- *Asymmetric encryption makes use of two keys for all parties communicating where the first key is used for encryption and the second key is used for decryption.*

Comparison of File Organization

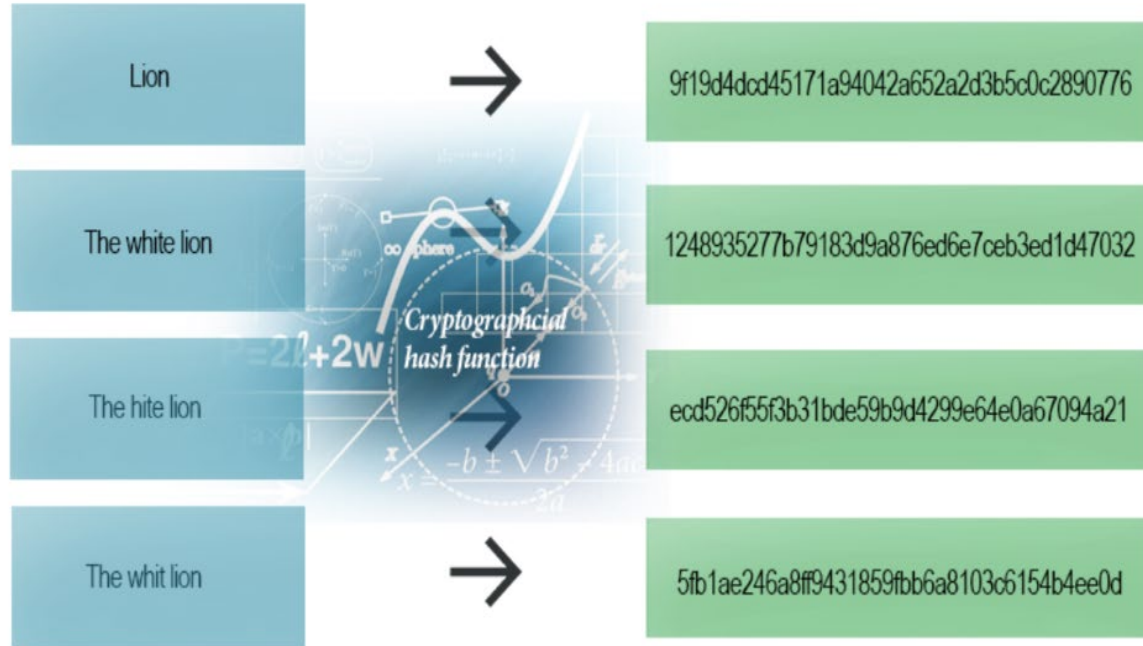
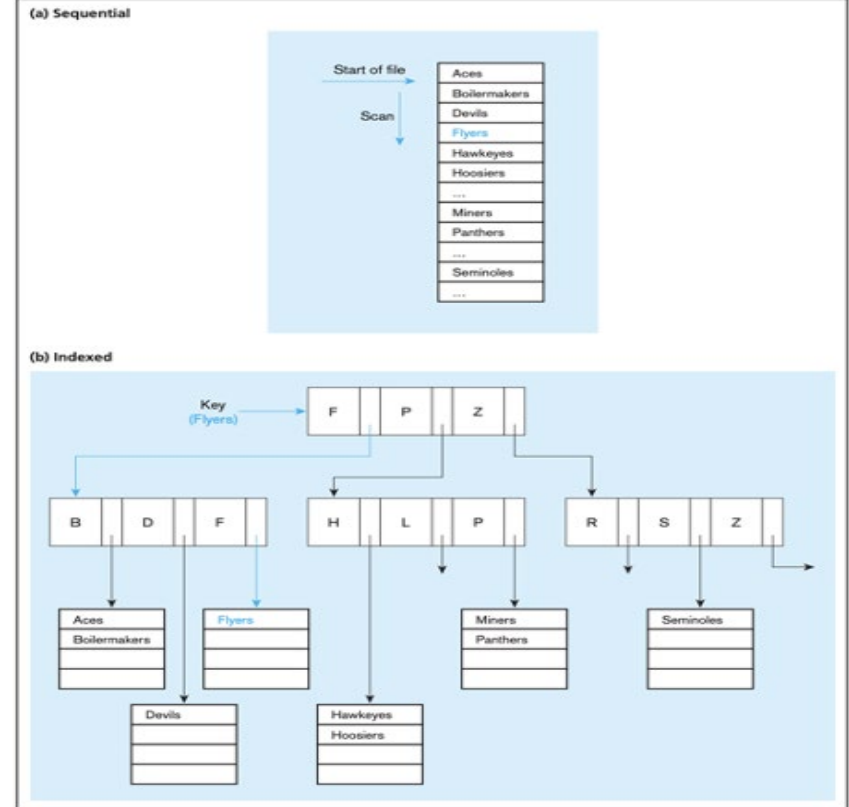


Figure 1.4. Hashing Algorithm, by jscribler, 2020, <https://blog.jscribler.com/hashing-algorithms/>. Copyright 2020 by jscribler

Comparison of File Organization

Fig 1.3 Modern Database Management 10th edition. ((Venkataraman, R., Topi, H. 2011))

FIGURE 5-7 Comparison of file organizations



Comparative Features of different File Organization


TABLE 5-3 Comparative Features of Different File Organizations

Factor	File Organization		
	Sequential	Indexed	Hashed
Storage space	No wasted space	No wasted space for data but extra space for index	Extra space may be needed to allow for addition and deletion of records after the initial set of records is loaded
Sequential retrieval on primary key	Very fast	Moderately fast	Impractical, unless using a hash index
Random retrieval on primary key	Impractical	Moderately fast	Very fast
Multiple-key retrieval	Possible but requires scanning whole file	Very fast with multiple indexes	Not possible unless using a hash index
Deleting records	Can create wasted space or require reorganizing	If space can be dynamically allocated, this is easy but requires maintenance of indexes	Very easy
Adding new records	Requires rewriting a file	If space can be dynamically allocated, this is easy but requires maintenance of indexes	Very easy, but multiple keys with the same address require extra work
Updating records	Usually requires rewriting a file	Easy but requires maintenance of indexes	Very easy

[4.4.1] Translate a database model into efficient structures

- Database manipulations demand the location of a row or a collection of rows that satisfies a condition.
- Searching for data can be quite the laborious task, given the magnitude of a database. Hence, using indexes can vastly increase the speed of the process and reduce the time and work.
- The usage and definition of indexes are a crucial spoke on the wheel of physical database design. Indexes are defined as either a primary key, secondary key, or both. It is ordinary to define an index for the primary key of a table.
- The index is formed of two columns: one column for the key and the other column for the address of the record that consists of the key value. In the case of a primary key, the index will only have one entry for each key value.

Structure of an Index in Database



Indexing in Databases, by
GeeksforGeeks, 2020,
<https://www.geeksforgeeks.org/indexing-in-databases-set-1/> Copyright
2020 by GeeksforGeeks.org

Using, Selecting, Creating & When to use Indexes

- **Using and selecting Indexes**

- Given the magnitude of a database, searching for data can be quite the laborious task. Hence, using indexes can vastly increase the speed of the process and reduce the time and work. The usage and definition of indexes are a crucial spoke on the wheel of physical database design. Indexes are defined as either a primary key, secondary key, or both. It is ordinary to define an index for the primary key of a table. The index is formed of two columns: one column for the key and the other column for the address of the record that consists of the key value. In the case of a primary key, the index will only have one entry for each key value.

- **Creating a unique index**

- The syntax to create a unique key index in SQL is "CREATE [UNIQUE] INDEX index_name ON table_name(column1, ... column_n);". The UNIQUE modifier specifies the values in the indexed columns. Creating a non unique key index is equivalent to a secondary key index. The term UNIQUE isn't used to create a secondary key index because values can be repeated.

- **When to use indexes**

- It is important to know when to use an index and which attributes to use when creating an index. Using indexes come at the price of performance. Performance is compromised when using indexes due to the overload for maintenance for insertions, deletions, and updating records. For this reason, indexes should be utilized mainly for data retrieval. According to the book, "Indexes should be used judiciously for databases that support transaction processing and other applications with heavy updating requirements, because the indexes impose additional overhead" (Hoffer, Venkataraman, & Topi, 2011). Here are some rules or conditions that suggest the use of indexes.

-

Rules & Conditions That Suggest The Use of Indexes

- 1. Indexes are a lot more efficient and practical for substantial tables.
- 2. Indexes are useful when there is a need to set out a unique index for the primary key.
- 3. Indexes are frequently used for columns that appear in WHERE modifiers of SQL commands.
- 4. Indexes should be used when for attributes referenced in ORDER BY and GROUP BY statements.
- 5. Indexes are convenient when there is diversity in the values of an attribute. For Oracle's standards, it is unproductive to use an index when an attribute has fewer than 30 values.
- 6. One point to keep in mind is to consider developing a compressed version of the values. Doing this will ensure that the index isn't slower to process.
- 7. If the index is used for finding the location of where the record will be stored, make sure the key of this index is a surrogate key to ensure the records will be fairly spread across the storage space.
- 8. Make sure to check the limit of indexes on the DBMS because some systems do not allow for more than 16 indexes.
- 9. Find a way to index attributes that have null values because rows with a null value won't be referenced.
-

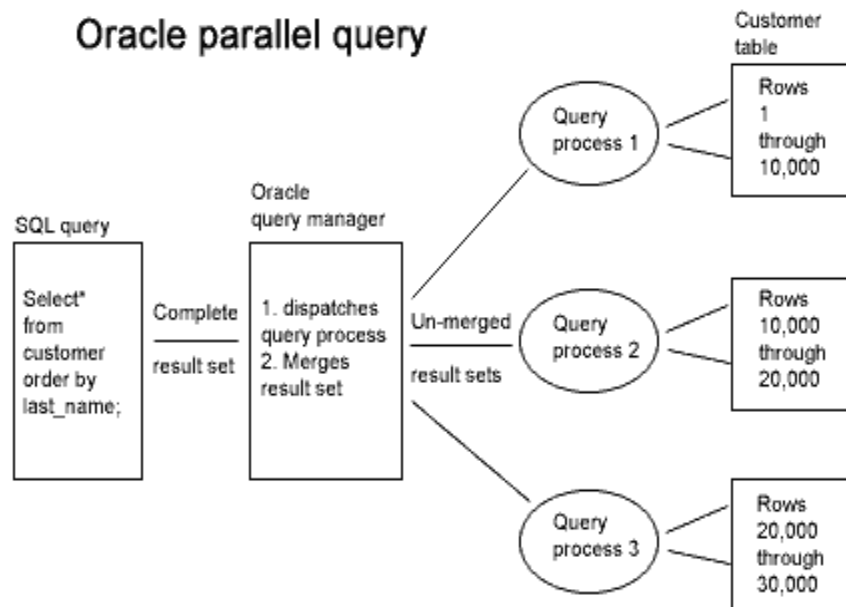
4.4.1 Designing a Database for optimal Query Performance

- Database processing can include adding, deleting and modifying a database along with method of retrieving data. The amount of work required to optimize query for performance heavily relies on DBMS.
- Architecture of modern computers has changed greatly over the years and the use of multiple processors in database servers has become standard. Symmetric multiprocessor (SMP) is commonly used in database servers to allow multiple processing. DBMS that use parallel query processing include planning on breaking up a query that can be processed in parallel by different processors.
- “Suppose you have an Order table with several million rows for which query performance has been slow. To ensure that subsequent scans of this table are performed in parallel, using at least three processors, you would alter the structure of the table with the SQL command:”(Hoffer, Venkataraman, & Topi, 2011)
- `ALTER TABLE ORDER_T PARALLEL 3; (Hoffer, Venkataraman, & Topi, 2011)`
- Schumacher reported, “on a test in which the time to perform a query was cut in half with parallel processing compared to using a normal table scan. Because an index is a table, indexes can also be given the parallel structure, so that scans of an index are also faster.” Schumacher also reported an example of parallel processing reducing the time of creating an index from seven minutes to five seconds.(Hoffer, Venkataraman, & Topi, 2011)

Designing a Database for optimal Query Performance (continued...)

- Parallel processing not only improves the time of table scans but also can be used on joining tables, grouping query results, sorting, deleting, updating, and insertion.
- Often the designer creating the query has information to better optimize the query that the query module in the DBMS does not. In most relational DBMs the optimizer's plan for processing the query can be known by the designer before actually running the query. This is done through the command EXPLAIN or EXPLAIN PLAN which display all the information about the optimizer's plans to process the query.
- The query optimizer makes its decision on how to process the query by looking at data from each table such as average row length or the number of rows. You can submit multiple EXPLAIN commands with a query written in different ways to see if the optimizer predicts different performance.
- That allows you to find the best performance and submit that for actual processing. With some DBMs you can force the optimizer to take different steps or use resources other than what the optimizer thinks is the best performance. The clause (/**/) can be used to override what the query determines is the best way to process the query.
-

Oracle parallel query



KENNESAW STATE
UNIVERSITY

Extended Resources

Description & Links

- This is a video lecture from the University of Washington by Grey Hay, about the physical database design methodology. This video goes into extreme detail about physical database design from the ground up and how the methodology is implemented. https://www.youtube.com/watch?v=S98_8HalY5Q
- This video talks about the oracle database security in a broad approach, mainly focused in Europe. The video discusses important security topics from current database security laws, benefits, history, risks and many more topics. <https://www.youtube.com/watch?v=GXF3T4g2tJg>
- This video by Kimberly Tripp where she discusses why physical database design matters. Kimberly discusses the importance of good design and also how poor design and can lead to major performance issues. <https://www.youtube.com/watch?v=H-jPsp2QIT0>
- Lightstone, S., Nadeau, T., & Teorey, T. J. (2007). Physical Database Design : The Database Professional's Guide to Exploiting Indexes, Views, Storage, and More. Morgan Kaufmann.
- Erickson, J., & Siau, K. (2009). Advanced Principles for Improving Database Design, Systems Modeling and Software Development. IGI Global.
- Carmel-Gilfilen, C. (2013). Bridging security and good design: Understanding perceptions of expert and novice shoplifters. Security Journal, 26(1), 80–105. <https://doi.org/10.1057/sj.2011.34>
- Burtescu, E. (2009). Database Security - Attacks and Control Methods. Journal of Applied Quantitative Methods, 4(4), 449–454.

References

- CloudGirl, & CloudGirl. (2017, March 26). Data Partitioning: Vertical Partitioning, Horizontal Partitioning, and Hybrid Partitioning Project Update #3. Retrieved from <http://cloudgirl.tech/data-partitioning-vertical-horizontal-hybrid-partitioning/>
- DigitalOcean. (2019, October 28). Understanding Database Sharding. Retrieved from <https://www.digitalocean.com/community/tutorials/understanding-database-sharding>
- Database VLDB and Partitioning Guide. (2016, July 20). Retrieved from <https://docs.oracle.com/database/121/VLDBG/GUID-BE424ACC-F746-4CA8-973C-F578CF98FF10.htm#VLDBG00225>
- Lesov, P. (2010). Database Security: A Historical Perspective. ArXiv, abs/1004.4022.
- Quick Base. (n.d.). A Timeline of Database History. Retrieved from <https://www.quickbase.com/articles/timeline-of-database-history>
- Watt, Adrienne. (n.d.). Chapter 13 Database Development Process. Retrieved from <https://opentextbc.ca/dbdesign01/chapter/chapter-13-database-development-process/>
- What is Entity Relationship Diagram? (ERD). (2020). Retrieved from <https://www.visual-paradigm.com/guide/data-modeling/what-is-entity-relationship-diagram/>
- Hoffer, A. J., Venkataraman, R., Topi, H. (2011). *Modern Database Management 10e*[E-Reader Version]. Retrieved from <https://www.amazon.com/Modern-Database-Management-Jeffrey-Hoffer/dp/0136088392>
- Burleson Consulting. (2014, April 23). Oracle parallel query tips. Retrieved March 28, 2020, from http://www.dba-oracle.com/art_opq.htm
- Chung, C. (n.d.). Introduction to Hashing and its uses. 2BrightSparks. <https://www.2brightsparks.com/resources/articles/introduction-to-hashing-and-its-uses.html>
-



KENNESAW STATE
UNIVERSITY