

Question #1: What are the main issues considered during physical database design?

Respondant: The main issues considered during the physical database design are the storage format for each attribute (choosing the data type), the grouping of attributes into physical records, arranging records into file organizations, selecting structures for storing and connecting files to make retrieving data efficient, preparing strategies for handling queries against the database that will optimize performance

Respondant: How to make processing of physical files and database efficient to minimize the time by user to interact with the information system

How to use the space efficiently.

Controlling data integrity

Using and selecting index

Respondant: Choosing the data type for each attribute from the logical data model to minimize storage space and to maximize data integrity.

Grouping attributes from the logical data model into physical records.

Arranging similarly structured records in secondary memory so that individual and groups of records can be stored, retrieved and updated rapidly.

protecting and recovering data after errors are found.

Selecting structures for storing and connecting files to make retrieving related data more efficient.

preparing strategies for handling queries against the database that will optimize performance and take advantage of the file organizations and indexes which have been specified.

Respondant: 1. The first step to take in physical database design is to estimate the size and usage patterns of the database.(Data volume and access frequencies)

2. Choosing data type for each attribute from the logical data model.

3. To minimize storage space and to maximize data integrity.

4. Selecting indexes for storing and connecting files to make retrieving related data more efficient.

5. Preparing strategies for handling queries against the database.

Respondant: The main issues considered during physical database design is how data will be processed and stored in the most effective way to insure database integrity, security, and recoverability.

Respondant: It is important to minimize time required by users to interact with the system. Speed and space need to be considered. The storage format that is chosen must minimize storage space and maximize data integrity. Another consideration is grouping attributes from the logical data model into physical records. You need to look at data volume and usage statistics. Have strategies prepared for handling queries. Need to optimize performance.

Respondant: Database integrity and performance are important issues to consider. The following decisions effect these issues during physical database design.

1. Choosing the storage format for each attribute from the logical data model.

2. Grouping attributes from the logical data model into physical records.
3. Arranging similar structured records in secondary memory so that individual and groups of records can be stored, retrieved, and updated rapidly.
3. Selecting structures for storing and connecting files to make retrieving related data more efficient.
5. Preparing strategies for handling queries against the database that will optimize performance and take advantage of the file organizations and indexes that you have specified.

Respondant: Well the goal is to design in such a way to provide processing efficiency while ensuring database integrity, security, and recoverability. Information includes the following: 1) Normalized relations, including volume estimates. 2) Definitions of each attribute 3) Description of where and when data are used: entered, retrieved, deleted and updated. 4) Expectations or requirements for response time and data security, backup, recovery, retention and integrity. 5) Descriptions of the technologies (database management systems) used for implementing the database.

Respondant: choosing the storage format, grouping attributes from the logical to the physical model, selecting structures for storing and connecting files to make retrieving data more efficient.

Respondant: The main issues of developing a physical database design is to provide good performance and insure database integrity, security, and recoverability.

Respondant: Choosing the storage format for each attribute from the logical data model. Grouping attributes from the logical data model into physical records. Arranging similarly structured records in secondary memory so that individual and groups of records can be stored, retrieved, and updated rapidly. Selecting structures for storing and connecting files to make retrieving related data more efficient. Preparing strategies for handling queries against the database that will optimize performance and take advantage of the file organizations and indexes that you have specified.

Respondant: Issues include decisions about storing the data most efficient way and detailed specification of attributes.

Respondant: Some important physical database design issues include: choosing a storage format to minimize storage space and maximize data integrity, grouping attributes from logical model into physical records, arranging data to improve speed for using the database, selecting structures to make retrieving data easier, and setting up strategies to handle likely queries.

Respondant: Type of storage format for each attribute from the logical data model. How attributes are grouped. Storing, protecting and recovering data after errors are found. Selecting structures called indexes and database architectures for storing and connecting files to make retrieving data more efficient. Overall database efficiency, handling queries and file organization.

Respondant: 1. choosing the storage format to minimize storage space and maximize data integrity.

2. make physical records out of the attributes from the logical model.
3. arrangement of records into another form of memory so that they can be accessed easily.
4. giving structure to store and connect files making retrieval easy.
5. develop an efficient strategy to handle the queries to the database.

Respondant: The main issues are how to store data from the logical model, how to group attributes from the logical design, how to organize, protect and recover files, how to structure the database for storing and connecting files, and how to design queries so as to use the most of the advantages gained by the physical design.

Respondant: 1. Choosing the storage format for each attribute from the logical data model.
 2. Grouping attributes from the logical data model into physical records.
 3. Arranging similarly structured records in secondary memory so that individual and groups of records can be stored, retrieved, and updated rapidly.
 4. Select structures for storing and connecting files to make retrieving related data more efficient.
 5. Prepare strategies for handling queries against the database that will optimize performance and take advantage of the file organizations and specified indexes.

Respondant: The main issues considered are: data processing efficiency, normalized relations-including volume estimates, definitions of each attributes, descriptions of data changes, security and integrity requirements, descriptions of technologies used, storage format, grouping of physical records, selecting structures, and strategies for handling queries.

Respondant: The main issues considered during physical database design are those that are not already implicit to the organization's technologies. The main goal is data processing efficiency. Some decisions made in this process are:

1. selecting the data type for each attribute.
2. grouping the attributes into records.
3. creating similar records in secondary memory.
4. selecting structures for storing files.
5. strategies for handling queries.

Respondant: The main issues for physical database design are:

1. Choosing the storage format to minimize storage space and maximize integrity.
2. Grouping attributes from the logical data model into actual records.
3. Arrange similar records in secondary memory so records can be stored, retrieved, and updated.
4. Selecting structures for storing and retrieving more efficient.
5. Creating optimal efficiency database.

Respondant: There's five of them outlined on page 255: 1) Choosing the data type for each attribute from the logical data model; 2) Grouping attributes from the logical data model into physical records; 3) Arranging similarly structured records in secondary memory so that individuals and groups of records can be stored, retrieved, and updated rapidly; 4) Selecting structures for storing and connecting files to make retrieving related data more efficient; 5)

Preparing strategies for handling queries against the database that will optimize performance and take advantage of the file organizations and indexes

Respondant: 1)Choosing the storage format for each attribute from the logical model to minimize storage space and maximized data integrity.
2)Group attributes into physical records.
3)Arranging similarly structured records in secondary memory.
4)Selecting structures for storing and connecting files to make retrieving related data more efficient.
5)Preparing strategies for handling queries against the database that will optimize performance and take advantage of the file organizations and indexes.

Respondant: The main issues considered during physical database design are:
-obtaining data processing efficiency
-data volume and usage analysis
-choosing data types
-controlling data integrity
-using and selecting indexes

Respondant: The main issues considered during physical database design are:
- the storage space which is limited
-the controls of data integrity
- the structure of the database which has to be efficient enough to enable a good storage
-easy accesses and manipulations to the data
(that is the responsiveness of the system is reasonable)
-the type of the data stored

Respondant: The main issue involved in designing a physical database is ways to minimize the time needed to interact with the information system, without taking much consideration on space efficiency. Key decisions that need to be made include: choosing the appropriate storage format, group attributes that came from logical data model into physical records, further arranging similar structured records in a secondary memory for rapid access, selecting structures, and finally the needed strategies for improving queries method.

Respondant: Normalizing relations, defining each attribute, describing where and when data is used, expectations and requirements for response time and data security, and descriptions of the technologies used for implementing the database, are the main issues in physical database design.

Question #2: How does the Garden.com article relate to issues of physical database design? What physical database design issues are mentioned?

Respondant:

The Garden.com article relates to the issues of physical database design through performance. The man interviewed commented that their performance was slow and that increasing the performance of the database was a major factor for success. The physical database design will determine the performance of the overall system. Issues mentioned are data volume and usage analysis. They mention that All of the files were on a single web server. Users must have great accuracy or else they will not have access to the desired information.

Respondant: System performance.
Controlling data integrity
Minimize the response time.

Respondant: The company struggles with database performance issues and hiring people who can understand its business", "One of the things on our Web site that the customer service reps use is the ability to look up users. It's certainly not done using indexes, and that's by default..." these problems are related to the physical database design.

Relative physical database design issue:

1. Selecting structures for storing and connecting files to make retrieving related data more efficient. "It's certainly not done using indexes, and that's by default. When we started, the database was fairly small, and it seemed to make a lot of sense"
2. Arranging similarly structured records in secondary memory so that individual and groups of records can be stored, retrieved and updated rapidly.
protecting and recovering data after errors are found. "I have 300,000 cache files on every single one of my Web servers, all storing this data locally

Respondant: The databases were very slow, and it affected the company to get performance. With the ever-growing of the company, the databases of Garden.com became much larger than beginning, and the frequency-of-use changed. At the same time, retrieving data from the databases became difficult, since the indexes of files should be improved.

Issues mentioned:

1. Data volume and access frequencies has changed.
2. using indexes to speed up the process of retrieving data.

Respondant: Without the proper considerations when developing the physical database design, you could run into major problems in the future as Garden.com has seemed to. The article emphasizes the importance of creating a robust physical database design. Issues on indexing and performance optimization were mentioned. They specifically stressed performance as the key to an effective database. In other words, how the database is set up to handle the large flow of data coming in and how to get that information out in the most fastest and easiest way possible without taking too much time, space in the system, and effort.

Respondant: This article relates to physical database design by the CTO discussing the problems they have had in designing and using their database. One issue mentioned is

performance, making the database run quickly. Another thing they mention is that once a shipping label is printed, the tracking number is in the system. He also mentions that they need to implement some sort of monitoring of the system.

Respondant: Garden.com (similar to others) have issues with database performance. They

Respondant: Performance. Garden.com is having issues with the performance of its database. Physical database design helps increase this performance. The article mentions that Garden.com looked seriously at where and when data is retrieved, and updated with scripts. They also mentioned the use of different DBMS such as Oracle and Solaris.

Respondant: garden.com's database is slow and hard to use, this is due to the physical design of the database. the database can't take hits, it's slow, and it is hard to get the information you want from queries.

Respondant: The article talks about a few issues of the physical database design:

1. The database should be able to be used by inexperienced users.
2. The database should be able to access information quickly.
3. A database developer must have business background.

The physical design issues mentioned are performance and database integrity.

Respondant: Because of the physical design, the database is slow and hard to use. The design issues mentioned are how the database can't take hits, its slow speed, and the fact that it is hard to get the information you want from queries.

Respondant: The article discuss about maximazing the performance and mentions problems with database speed. With right physical design the databases are able to be faster and give better performance.

Article mentiones for example the use of cache.

Respondant: In this article, the CTO of Garden.com talks about some of the problems and key issues they have faced from having their whole business rely on a database. He talked about how he wanted the system to be very easy to use so that even an inexperienced developer could do "good things" to the system. He also mentioned that the use of indexes and how it was too slow for their sales people to look up users information. He also spent a lot of time going over the fact that cache was very important to his system and that the traditional database software was not strong enough for his company.

Respondant: They talk about structuring issues, the main one being performance. The database system must perform the task required quickly and efficiently. Monitoring the system is a very important issue because it's they only way you can get feed back about what's going on within your system. Having a system that is totally programmable is the way to because it allows the developer versitility and makes upgrade the system much faster.

Respondant: one issue that was discussed is that of choosing the data format to get the best performance. another design issue that is mentioned is preparing strategies to handle queries to optimize performance, like looking up users.

Respondant: Garden.com seeks performance, which is affected by several issues. When the database is not populated at all, it could be seen as performing well, but as the data in the system increases, the database may perform poor. CTO states that they made the most of cache usage, to increase the performance, and design the database such that it becomes easy to use. Most of the decisions made are related to the issues (a) how to structure the database for storing and connecting files, and (b) how to design queries so as to use the most of the advantages gained by the physical design.

Respondant: CTO of the Garden.com specified Performance as the most important issue regarding optimized database architecture. This relates to preparation of strategies for handling queries against the database that will optimize performance and take advantage of the file organizations and indexes specified. Efficient database structure is critical to database management systems and will dictate the performance of the database.

Respondant: One of the issues mentioned is performance. This relates to the efficiency issue of the physical design. Specifically, how to increase the speed of the database.

Another issue is selecting structure. The architecture used on the website is very old, and this is affecting the performance.

Another issue is the handling of the queries. Because of the many users data in the database, it can take up to 20 seconds to retrieve the data if it wasn't queried properly (ie, by clicking the 'Exact Match Only' button).

Lastly, the issue of selecting and describing appropriate technologies was also mentioned (eg. what type of programming language to use or what type of technologies to adopt next).

Respondant: The article opens stating that "the company struggles with database performance issues.." and then begins to go into detail about why this is the case and how the design of the physical database affects performance. A physical database design issue mentioned is how the company by default does not use indexes to organize the data.

Respondant: Garden.com is related to issues of physical design in that when you have a large database, you have to store information in logical senses that most information can be derived so the amount of data stored in the database can be minimized and made efficient.

The physical design issues mentioned is that Garden.com has 300,000 cache files on every single one of my Web servers, all storing this data locally. Customer service reps use ability to look up users. It's not done using indexes. When the database was fairly small, you didn't have to be that accurate. When the database gets big it takes lots of time to search in the web page. But you don't want your customers to wait for search.

Respondant: It sounds to me like the Garden.com CTO has a lot of trouble with his database, and it's probably not planned as well as it could be. But when you're trying to run a company based on a database, there are obviously many physical design issues that come up. The issues he mentioned were caching, database speed, use of Oracle or Solaris or Access, use of indexes,

64/128-bit systems, C vs Java, monitoring, user interface quality, linking with FedEx, data integrity, a whole bunch of issues.

Respondant: I could not find the article on Garden.com

Respondant: The Garden.com article relates to issues of physical database design because the main emphasis is on performance. This relates to the above mentioned--data processing efficiency. An example of a physical database design issue that is mentioned is the ability to look up users (in their case, it is not done by indexes).

Respondant: Garden.com wants to maintain a high level of performance for its information systems. But as the architecture of the system used is quite old, the company has to face some physical database design issues. The article mentioned these issues:

-the storage capacity is not sufficient enough any more thus the data are stored in many different places

-the language used is not appropriate any more; Andy Martin thinks that they should use JAVA now instead of C

- the controls of data integrity are not sufficient, they could be more automatic; thus Andy Martin said " and you can't just put the system in and hope that everyone's going to do it the right way". He added too: "Don't hit the database. Put everything in it, but don't hit it. It's really unbelievable that we had to do that."

The constraints could be verified automatically by the DBMS and not by Andy Martin for instance

-the architecture needs improvements: there is no index

-the responsiveness of the system is not always optimal: the manipulations or retrievals can be difficult and long sometimes.

Respondant: The article mentioned the problems the company are having with the architecture of their database. Garden.com also desires to make attempt on a faster user utilization on the database. The CTO talked about how on each of their webpage there are about 300,000 cache files being stored in the server. This promotes rapid user retrieval and storing of data. When he was complaining about the current database method of retrieving data, he was referring to the lack of efficiency in data processing as it takes longer if accuracy in using indexes is not high.

Respondant: The article about flowers blooming on the web relates to physical database design. Data security and describing when and where the data is used were mentioned in the article.

Question #3: What class on campus could you take if you wished to know more about physical database design issues?

Respondant: Physical Database design provides the technical specifications that programmers and others involved in information systems construction will use during the implementation phase. It seems that taking CSE 444 would give you the opportunity to learn more about the physical database design issues. CSE 444: Introduction to Database Systems
Catalog Description: Fundamental concepts, system organization, and implementation of database systems. Relational, hierarchical, and network data models; file organizations and data structures; query languages; query optimization; database design; concurrency control; security; issues involving distributed database systems.

Respondant: Database design

Respondant: data structure, software engineering

Respondant: Data structures
Database system

Respondant: Probably IS 300 or 480 (Database management classes)

Respondant: Information Systems 585 - Advanced Database Research.

Respondant: I S 460 Systems Analysis and Design I
I S 461 Systems Analysis and Design II
I S 480 Database Management

Respondant: IS 480 Database Management

Respondant: i would take is 480 Database Management.

Respondant: To learn more about physical database design issues, a student can take IS300 and IS480.

Respondant: One class that I heard about was 480 Database Management in the business school.

Respondant: Courses in data modeling - courses from computer science department.

Respondant: IS 480, Database Management deals with concepts of logical and physical database organization.

Respondant: IS 300

Respondant: IS 321: Database Management and Telecommunications for Accountants. this class includes processing architectures, database reliability, database recover, database security, database administration, etc.

Respondant: Looking from the Course Catalog for Autumn 2000, I guess the most related courses would be:

CSE 326 DATA STRUCTURES, and
CSE 444 INTRODUCTION TO DATABASE SYSTEMS,
which are both offered by the Computer and Systems Engineering Dept.

Respondant: I S 300 Introduction to Information Systems
I S 321 Database Management and telecommunications for Accountants
I S 480 Database Management
I S 545 Database Systems and Applications
I S 560 Information Systems Development

Respondant: I could take I S 461: Systems Analysis and Design II.
This is a second course in analysis and design of business information systems. Concentrates on design and implementation phases of systems development. Translation of logical system model into physical model, design of modules, file design, testing, implementation. Includes a project using third- and fourth-generation software development tools. Prerequisite: I S 460.

Respondant: IS 480: Database Management.
This course description from the School of Business website puts an emphasis on learning about physical design issues.

Respondant: I would choose to take C++ or java or MS access or Homepage classes. All of these computer programming classes emphasize on getting the functionality working while reducing redundancy.

Respondant: IS 461 - Systems Analysis and Design II

Respondant: Enterprise Data Integration
or Data Modeling

Respondant: If you wished to know more about physical database design issues, you could take this class or IS480 "Database Management".

Respondant: I could take CSE 544 "Principles of Databases Systems" to know more about physical database design issues. Indeed the following part of the catalog description, introduces some interesting points I could analyze in addition to what I am learning with INDE 499 B course regarding physical design:
"Principles of data storage and indexing. Query-execution methods and query optimization algorithms".

Respondant: IS 480: Database Management.

Respondant: IS 480