

**Course:** MCTE 630 – Database Systems  
**Subject:** Homework #3  
**Task:** Problems from text:  
The Concepts of Database Management  
Pratt & Adamski (1997)  
**Professor:** Dr. Maxine Cohen  
**Student:** Leanne C. Boyd  
**Usercode:** boydl  
**Email:** boydl@scis.acast.nova.edu  
**Due date:** week of June 15, 1998

### HOMEWORK #3

---

#### CONCEPTS OF DATABASE MANAGEMENT:

**Chapter 7 # 2, 4, 10, 14, 19 (answer in terms of work or school), 20**

---

- 2) **What is the purpose of the catalog? What types of information usually are found in the catalogs that accompany microcomputer DBMSs? What additional types of information often are found in the catalogs that accompany mainframe DBMSs?**
- a) The catalog (which MUST be furnished by the DBMS) is where the user can access descriptions of the data items that are stored. It literally contains everything the DBMS contains, so that any questions imaginable can be researched, queried, answered, etc. This is CRUCIAL for people like those in charge of the dbase, or those writing programs to access the dbase, as they must know very quickly what the dbase “looks like.”
  - b) Types of information found in microcomputer DBMS catalogs: These are more limited than those for mainframe DBMSs, but usually will provide information on SOME of the comprehensive list shown below for mainframes. They will usually tell you about tables and fields and their descriptions, and info on relationships. They are not as extensive as the mainframe Super Catalog (see item c), next).
  - c) Mainframe DBMS and additional information: They often offer what is called a data dictionary, which is a separate entity and is somewhat of a Super Catalog. It will answer just about everything concerning the dbase, such as: tables and fields and their names; all characteristics of those tables and fields; possible values for the fields; restrictions; what the meanings of the fields are, as in a description of acronyms used; relationships and the meanings of those relationships; description of varying factors for making relationships and the necessity of certain relationships, or not; programs that access the database to retrieve, update, etc. and the limitations of those retrievals, updates, change abilities and the descriptions of how each can affect the dbase.

---

**4) Describe a situation, other than the one used in the text, in which uncontrolled shared update would produce incorrect results.**

I have worked several times for large corporations where this could, indeed, produce massive problems. One situation was at Rocky Flats (thermonuclear plant in Denver), where “mistakes” such as those described here could produce heinous results! My current position at Lucent Technologies could produce hundreds of similar problems, if there weren’t such a safeguard approach on the network. In both company examples, with such large databases, covering immense amounts of data, there has been a very detailed safeguard procedure against the loss of stored items by overwriting or deletion by user or by the system. Both systems allow use of common documents, but by seeking permission of the DBMS to download to your own hard disk. All the while, both systems also show WHO ELSE is currently working on the same item. Both systems worked with a system of work being produced IN VERSIONS of the document. In Lucent’s case, where the changes may be made in large online documents, this is especially helpful, as it tracks the history of the document. In no case are the changes permanent until approved by someone higher in administration. But, the key to both systems seems to be that the computer interface LETS YOU KNOW WHO ELSE IS WORKING ON WHAT YOU ARE WORKING ON. You aren’t sitting in the dark, thinking it’s just you and the dbase.

---

**10) What is meant by recovery? What facilities typically are provided by microcomputer DBMSs to handle backup and recovery? What main feature is lacking in such facilities? What problems can this cause for users?**

- a) Recovery is the process of returning the database to a correct state, after some sort of disaster that has altered or destroyed data. This might include incorrect input by users, update programs not functioning, and any sort of hardware or electrical malfunction. You literally must have a way to RECOVER the database, partially or entirely.
- b) The facilities for backup and recovery for microcomputer DBMSs is rather limited. They allow you a way to make a simple backup of all material. In the case of disaster, you then can copy over the mistake with the last saved working backup. Any successful input BEFORE your disaster is lost, along with eliminating the disaster.
- c) The main feature lacking is the JOURNAL or LOG that is found in mainframes, that keeps a record of all the updates to the dbase. This allows for going back and retrieving successful efforts and finding the disaster itself. But microcomputers don’t have this.
- d) This causes problems for a user, as the only real good alternative is to make very frequent backups, so that only a short time period is lost if a disaster should occur. This causes more work for the user, and work time eaten up just in backup efforts. Extra logic is needed to write the backup programs, as well as space for memory needed, etc. When the user makes a backup, more computer facilities including memory is needed, which slows the system down. Based on these considerations, it’s up to each user to weigh the losses against the benefits of frequent backups.

---

**14) How do views relate to security?**

When a DBMS allows the use of views, each user has their own view and not access to any other views. So, that information is literally invisible to that particular user. This creates a form of security in that when the data is not accessible, it is more secure against attack or accident.

---

**19) How well does your school's DBMS fulfill the functions of a DBMS described in this chapter? Which functions are supported fully, which are supported partially, and which are not supported at all?**

The DBMS at Nova seems to be fairly full functioning. Much of this very question has been addressed at length in our Forums. One of the major problems with Nova's computing, including the DBMS, is its susceptibility to the problems of not enough size and power to meet the needs of so many students utilizing at the same time. For my limited knowledge in this arena, however, Nova's system allows changes in an efficient manner, in creating, altering, changing, dropping, etc. tables. It allows index making and deletion. It allows easy access to services and also edit/query possibilities. It has somewhat of an easy ability to switch between programs, although our SQL software has a very basic, irritating, non-keyboard approach to editing. It has a help system, which I thought was very unclear to beginning students, but it IS there. Finally, I think that the system is very powerful, but the limited scope of this class didn't allow much knowledge of that, as our newly acquired set of skills don't allow much exploration. Overall, the system seems to be excellent, but with severe limitations because of either power or server size -- or whatever causes massive delays and cutoffs. I don't feel that our short trip through this subject allows enough information to know if a certain function is only partial or total on the Nova system. It met the needs of this class quite well, it seemed.

---

**20) Many computer magazines and a number of Web sites present comparisons of several DBMSs. Find one such DBMS comparison and compare the functions in this chapter to the listed features and functions in the comparison. Which functions from this chapter are included in the comparison, which functions are missing from the comparison, and what additional functions are included in the comparison?**

Much information was indeed to be found on comparison of DBMSs. One IMMEDIATE reaction I had to finding such websites was that much of this information is very, very OLD, with web pages covering information from 1996, backwards. Not knowing enough about DBMSs in general, I don't know if this is typical, and that the systems described could have remained much the same over the years. The article is found at:

<http://dbmsmag.mfi.com/9611d52.html>

One such site, *DBMS Server Comparison Supplement*, November 1996, has an article entitled, "An Analysis Of The Strengths And Weaknesses Of The Big Six Database Servers." In reading the article, I found it covered DBMSs and Relational Data Models, Database Objects, Queries, Connectivity and Distribution, Connectivity and Distribution, Replication, and Internet use.

Their comparison table includes: ORACLE7 VERSION 7.3, SYBASE SQL SERVER 11, INFORMIX-ONLINE 7.2, MICROSOFT SQL SERVER 6.5, IBM DB2 2.1.1, and CA-OPENINGRES 1.2.

From there, they compare several hundred criteria, including: domains, referential WHERE clauses (where NONE of them do this), updateable views, user-defined data types, table structures, index structures, stored procedures such as Language (PL/SQL, Transact-SQL, etc.), events, database administration tools and security, partial backup and recovery, Internet connectivity and use, distribution, Gateways, and much, much more.

In fact, quickly perusing several articles just like this one led me back to Forum discussions, where, at our beginning level, much of this is still a foreign language. You need to be more than a beginner to even be able to decipher much of the information on pages such as this article, where the “help” involved is for persons with a high understanding of DBMSs.

Meanwhile, MOST everything covered in our textbooks appear in one form or another in these web pages. They also include much more information than was included in our texts. Our system at Nova may indeed have some real upscale capabilities along the lines of what I saw in the articles I researched. However, I am at too basic of a level to even KNOW that if it is staring at me from the printed or electronic page!!

Personally, I believe this particular question belongs in a book for higher-level students. Secondly, it seems appropriate for a term paper, rather than a short answer in a set of problems. This, I feel, is part of the inconsistencies of these textbooks, and especially for use in such a beginning course. We’ve learned some basics, and a question like this immediately throws the student back into the realm of feeling very stupid, like nothing at all has been learned. At the very least, it will give just a short view into the magnitude of the subject, and will intimidate some students into NEVER looking further into DBMS worlds.

---

## CONCEPTS OF DATABASE MANAGEMENT:

**Chapter 8 Interview a DBA in your organization. Write up his/her job responsibilities, training, security issues, etc.**

---

I have just begun my position as part of the Web Development team at Lucent Technologies in Denver. I have not had the opportunity to meet any of the upper-level DBAs for Lucent. However, I have met with one individual quite frequently in setting up all my accesses at work, who is part of the extended DBA team, and deals with “local” issues, one at a time, per individual personnel. Troy is the person to call for any computing need, from setup and access, to email setup and coordination, to allocating passwords for general and specific access into the DBMS. In meeting Troy, much of our textbook studies came into focus.

Troy’s position requires much knowledge of the extended database at Lucent. He said that there wasn’t a lot of the DBMS that was inaccessible to him, as he deals with such a variety of problem-solving issues. He said he needs *two of himself* to keep up with daily demands. Part of what Troy does is to access the dbase in setting up, altering, or moving accesses for individual employees. In his discussion, it appeared

that what seems seamless to you and me, actually requires much in the way of clearances, queries, and altering of the database, to make it happen!! He is constantly and consistently concerned about security issues, which became apparent even in the short time he changed the parameters of my individual computer in connection to the network. I was able to see his frustration when the DBMS didn't perform as expected with his first set of queries. On his second visit, perhaps more information had been updated by another department for the technical writer who had left Lucent, and whose machine I inherited. The changes Troy then queried and committed, were successful.

His job requires coordination with the security and badging department, the internal phone communications department, Human Resources, and probably dozens of other departments that he didn't mention.

The evidence of a well-organized DBMS and a well-placed, thoughtful and caring DBA has been very evident in dealing with Troy. The DBA is a tremendous asset to the company when the jobs have been done correctly. I can see, from working with Troy in getting my "spot" comfortable in the Lucent environment, that anyone with less technical skills would NOT be an asset. I can also see that the DBA requires much in personable attitudes ... people skills. In fact, I would say that at least 50% or more of Troy's job description would have to do with those people skills, rather than just the technical skill set of accessing the DBMS and making necessary, daily alterations to such a huge collection of data.

Even though Troy's job does not seem to require accessing the massive number of remote Lucent systems, it does appear that his correct manipulation of local data would affect the success rate of the Lucent plants in New Jersey, Georgia, Louisiana, Ohio, and other parts of the world. It is as far-reaching a position, in many ways, as the one I have, where locally planned and created learning websites will affect personnel and customers, worldwide.

---

## CONCEPTS OF DATABASE MANAGEMENT:

### Chapter 9 #4, 10, 12, 34 (answer in terms of your institution or workplace)

---

#### 4) What is location transparency?

One of the most desirable objectives of a distributed system is to make the user feel that ALL of the information is at a local level, and that they are not accessing data that is very far away, distance-wise. The user should be able to access from a REMOTE site as easily as from the LOCAL site. The term, LOCATION TRANSPARENCY, describes this situation, where the user feels that all data is together, on-site, and immediately accessible, regardless of the actual location of a remote system. Literally, the location and the traditional meaning of location, distance, etc. – becomes an invisible, non-pertinent fact in utilizing the database.

---

#### 10) Why is query processing more complex in a distributed environment?

In a distributed environment, much of the data is being accessed via a REMOTE system. The additional time needed for making the query and waiting for the results can bring back more than just irritation. It takes much less time to access a local disk than to send messages between sites.

Additionally, sometimes the remote site will not have the information requested, and a new request to another site will take up even more time. If your system is limited to a “one-record-at-a-time” orientation, the communication factor on a large query could take an immense amount of time. The problem becomes much less if your system is able to perform requests for sets of records instead of just one at a time.

**12) Describe the principle of two-phase commit. How does it work? Why is it necessary?**

In a two-phase commit form of recovery measure in a distributed environment, you have several sites needing to backup the database in a cooperative and consistent manner. At Site #1, you will have a coordinator of the process, who will send messages to all other sites, advising of the upcoming backup, requesting they acquire the necessary locks and prepare for the update. The sites are required to do so, and then send back messages stating they are ready. Any reason for aborting on either end requires further messages and a new time set by yet more messages for a new update time. When all replies are positive, the update may proceed. The success of this method lies with the capability of the coordinator in overseeing the process, and wading through many messages when an update needs to be aborted. This means that the remote sites are not as independent as might be wished, as they rely heavily on the coordinating site for backup of data. As with any DBMS, an efficient backup system is absolutely necessary, in spite of any shortcomings or irritations of a bulky method – in case of dbase disaster.

**34) If your school is using the Internet and/or an Intranet to manipulate data stored in a database, investigate how this is being accomplished (which software and database entities are used). If not, what are your school’s future plans in this arena?**

Again, I think this question may be more well-suited to students with more experience. I also think that proper investigation of this would lead to many students requiring much information from already over-loaded technical personnel – leading to less efficiency than we already have at Nova, in the administration of computer issues.

Nova utilizes the Internet in every course I’ve taken. I would surmise that there is Internet use for manipulating data in a dbase, simply because there are many indicators of it ... such as accessing information on courses, the listings in the library system and of student dissertations. I have become somewhat familiar with the UNIX system, the Oracle capacities (in this SQL course), the proprietary software and platform used for the Moderator and the ESET system. I’ve utilized the ECR and other telnet-based programs, which is part of the UNIX setup at Nova.

Nova’s current plans (I don’t have a clue) SHOULD include massive reconstruction or upgrade in Moderator and ESET. This would be a huge help to students. Nova’s plans SHOULD include a massive upgrade to the server systems in size and power, in order to give quality service to ever-growing numbers of students. I have heard from teachers that there have been plans to set up capacities for MS NetMeeting, although I don’t know if that is school-wide or more local in our department. Finally, if courses are given in Multimedia, which they are – then these systems need to have immediate plans for updating equipment and software, in order to be able to deliver projects without harming this overloaded and antique system we seem to have.

In other words, I do not know if the information requested is available without bothering Nova personnel.



*I diddit, I geddit – I guess it wasn't THAT bad !! 😊*