

# Project

Due Friday April 18, 2008

Checkpoint 3: Mon Mar 31, 2008

(updated schedule, executables for at least 2 interfaces)

Checkpoint 2: Mon Mar 17, 2008

(relational design, user interfaces and their features, development schedule)

Checkpoint 1: Wed Feb 27, 2008 (E-R diagram)

CSE 241 and 398

Database Systems

Spring 2008

**Goal:** The goal of this project is to provide a realistic experience in the conceptual design, logical design, implementation, operation, and maintenance of a relational database and associated applications. First, I shall describe the application, then the categories of requirements, and then some suggestions on how deeply you need to go in each category. A real project of this sort would require a substantial development team working for several months (or more). You will do this alone over several weeks. I have chosen to go with individual rather than group projects because the goal of this project is for you to gain a personal appreciation of the depth and breadth of issues that go into the design of a database application, rather than to have you specialize in just one aspect (and rely on others for the rest).

The project can go well beyond the minimal requirements I outline at the end. I encourage such extensions. They could turn into a senior design project or other independent work.

**Application description:** The application is an electronics vendor that operates both a Web site and a chain of many physical stores. Examples include Best Buy and Circuit City. To find out more about this application, think about any experiences you may have had making purchases both online and in-store, and browse their Web sites.

In our hypothetical company, it has been decided to redesign a major part of the database that underlies company operations. Unfortunately, the manager assigned to solicit database design proposals is not very computer literate and is unable to provide a very detailed specification at the technical level. Fortunately, you are able to do that.

Here are a few points to consider:

- There are many different products, grouped into a variety of (possibly overlapping) categories. Groupings can be by type of product (cameras, phones, etc.), by manufacturer

(Sony, Apple, etc.), or by other means (for example, a Gateway PC might be packaged with a Sony monitor and an HP printer and marketed as a package).

- Some customers have a contract with the company and bill their purchases to an account number. They are billed monthly. Other customers are infrequent customers and pay with a credit or debit card. Card information may be stored for online customers, but not for in-store customers.
- Online sales must be sent to a shipper. The company needs to store the tracking number for the shipping company so it can respond to customer inquiries.
- Inventory must be accurate both in stores and in warehouses used to replenish stores and to ship to online customers. When inventory is low, a reorder should be sent to the manufacturer and listed in the database. When goods arrive, inventory should be updated and reorders marked as having been filled.
- Sales data are important for corporate planning. Marketers may want to look at sales data by time period, product, product grouping, season, region (for stores), etc.

### **Client Requests:**

#### 1. E-R Model

- Construct an E-R diagram representing the conceptual design of the database.
- Be sure to identify primary keys, relationship cardinalities, etc.

#### 2. Relational Model

- After creating an initial relational design from your E-R design, refine it based on the principles of relational design (Chapter 7).
- Create the relations in Oracle database you used for earlier projects.
- Create indices (Chapter 12) and constraints (Chapter 8) as appropriate.
- If as you refine your design, you discover flaws in the E-R design, go back and change it (even if the earlier design passed the checkpoint.) Your final E-R design must be consistent with your relational design.

#### 3. Populate Relations

- Include enough data to make answers to your queries interesting and nontrivial for test purposes.
- You may find it helpful to write a program to generate test data.

#### 4. Queries: You should run a number of test queries to see that you have loaded your database in the way you intended. The queries listed below are those that your client (the manager from the package delivery company) wants turned in. They may provide further hints about database design, so think about them at the outset of the project.

- Assume the package shipped by USPS with tracking number 123456 is reported to have been destroyed in an accident. Find the contact information for the customer. Also, find the contents of that shipment and create a new shipment of replacement items.
  - Find the customer who has bought the most (by price) in the past year.
  - Find the top 2 products by dollar-amount sold in the past year.
  - Find the top 2 products by unit sales in the past year.
  - Find those products that are out-of-stock at every store in California.
  - Find those packages that were not delivered within the promised time.
  - Generate the bill for each customer for the past month.
5. Interfaces: There are several types of users who access the database. Each may need a special application
- The database administrator (you) may use SQL either via the command line or SQL Developer.
  - Customer service needs a lookup application to check inventory both locally and at nearby stores.
  - Online customers need an elegant Web interface. (depending on your skill with applets, they may or may not get this.)
  - Call center staff need an application that allows quick access to customer data and the ability quickly to enter phone orders.
  - The stocking clerks at the warehouses need an application to help them record incoming shipments and update inventory
  - The marketing department needs sales reports and may want to do special data mining and analysis. Because we do not cover these topics until late in the course, these features may be left to version 2, planned for the year after you graduate Lehigh.

These interfaces may, in some cases, use existing Oracle tools, in which case all you need to provide are instructions (e.g. run Query 17.SQL). In other cases, the target user of the interface is not someone we would expect to be fluent in SQL. Those interfaces can be built as

- Web applications using Java applets or a scripting language.
- A standalone Java application using Swing to create a GUI
- Other GUI development tools you may know (but be sure they are platform independent, see note below)

- Since this course is not a Java course, nor a GUI course, I will accept a simple command-line interface. In fact, even if you are expert in GUI development, you may want to start with a simple command-line and then upgrade later.
6. Concurrency: The company stock will go down rapidly if the database cannot support more than one customer at a time making a purchase. Be sure the Oracle transaction mechanism is providing the needed guarantees.

**What to turn in and how much I actually expect you to do:**

The checkpoints are meetings only; nothing needs to be turned in. Usually, I find the first checkpoint requires some discussion, while checkpoint 2 can often be handled quickly (even via email). Checkpoint 3 is mainly in place to ensure that the projects are on schedule (again, email will do). Please talk to me about questions at any point; don't wait for a checkpoint.

The final version of the project is to be turned in as a single zip file on blackboard. I will accept paper for the ER diagram since we are not covering drawing tools for these diagrams.

*I, as DBA, will have access to your database online and will use that if your submission leaves me with some unresolved questions.*

1. E-R diagram, plus any explanatory notes. At minimum you must include all the entity and relationship sets implied by this handout. You may go beyond the minimum. Remember that the manager who defined the specifications is not computer literate so the specifications should not be viewed as necessarily being precise and complete.
2. Relational schema. It is likely for many of you that your ER design will be sufficiently extensive that we agree that only a part of the resulting relational design will actually be implemented under Oracle on Edgar3. This is something on which we'll agree before Checkpoint 2. This is the point where we cut implementation effort and data-entry time to something realistic for the course time frame. *You may share data with your classmates but not design nor code (SQL, Java or anything else) provided you credit your source.*
3. Do NOT turn in a listing of all your data. I can see them online if I find it necessary.
4. A set of sample queries.
5. The code to implement the various interfaces. (By restricting your development to Java, your code should be platform independent. Please discuss any non-platform independent strategies before you implement them.) I will accept quite basic interfaces (command line with a modest command set), but encourage more elegant interfaces. Depending on the degree of sophistication you plan for each interface, we can agree to fewer than the 6 interfaces requested by the client.
6. A README file in the top-level folder that explains what is where, etc. Include usage instructions for the interfaces