

Datenbanken II A: DB-Entwurf

Chapter 7: Logical Design III

Prof. Dr. Stefan Brass

Martin-Luther-Universität Halle-Wittenberg

Wintersemester 2024/25

<http://www.informatik.uni-halle.de/~brass/dd24/>

Objectives

After completing this chapter, you should be able to:

- explain the steps in which a database schema is developed with Oracle Designer and name the tools that are used in this process.
- write a short paragraph about the Database Design Transformer of Oracle Designer: What it can do and what its limitations are.
- read Server Model Diagrams in Oracle Designer.

DB Design Transformer (7)

- For many-to-many relationships, a relationship table is generated:

```
COURSES_STUDENTS(CRS_CRN→COURSES,  
                  STUD_SSN→STUDENTS).
```

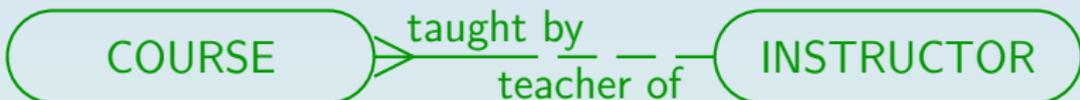
- The name of the table for the relationship is composed out of the plural forms of both entity types.

Probably it would have been nicer if the relationship names were used in some way. It is possible to change generated table and column names later in the Design Editor. Also, I would have preferred “STUDENTS_COURSES”, with the “from” side of the relationship first. But the Database Design Transformer always uses the alphabetic sequence.

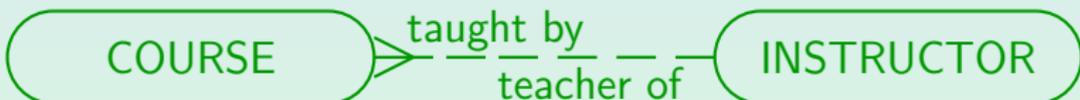
Restrictions (3)

Exactly Translated Relationship Types:

- Many-to-one, mandatory-to-optional:

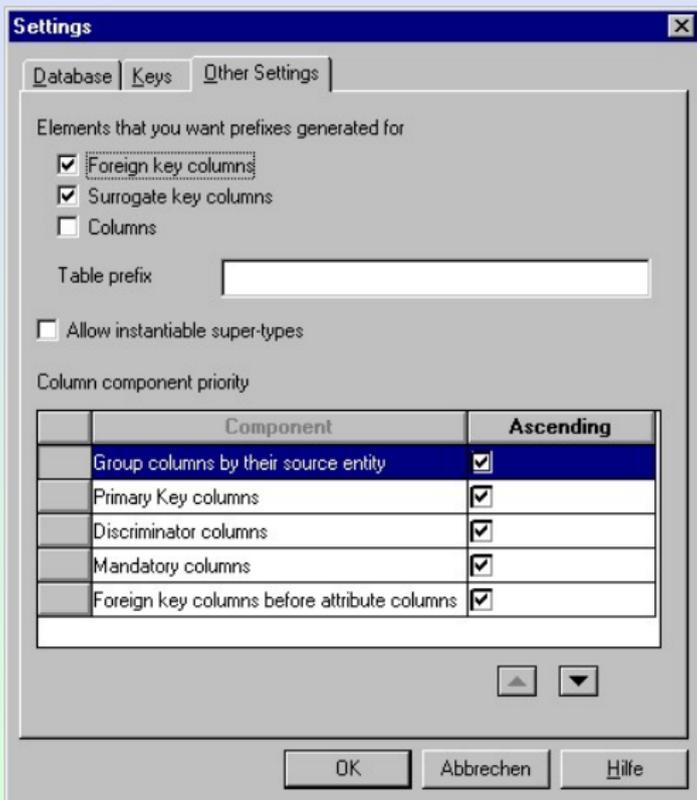


- Many-to-one, optional-to-optional:



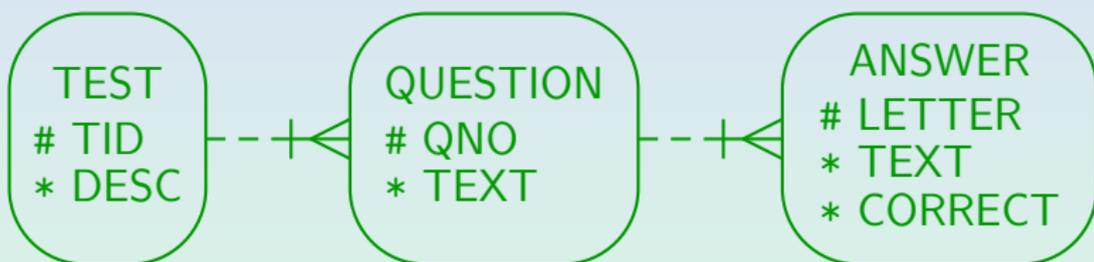
- Many-to-many, optional-to-optional:





Translation of Weak Entities

- The DB Design Transformer translates a hierarchy of weak entities as expected:



```

TESTS(TID, TEST_DESC) -- DESC is a reserved word
QUESTIONS(TEST_TID→TESTS, QNO, TEXT)
ANSWERS((QUEST_TEST_TID, QUEST_QNO)→QUESTIONS,
        LETTER, TEXT, CORRECT)
  
```


Translation of Subtypes (9)

- Result of Method 4 (“Arc Approach”):

COMMITTEES(CNAME)

COMMITTEES_FACULTY(COM_CNAME→COMMITTEES,
FAC_1_FAC_ID→FACULTY)

COURSES(CRN, TITLE, INST_NAME→INSTRUCTORS)

EXTERNAL(ADDRESS, EXT_ID)

FACULTY(TENURED, FAC_ID)

INSTRUCTORS(NAME, EMAIL, EXT_EXT_ID^o→EXTERNAL,
 FAC_FAC_ID^o→FACULTY)

- The foreign keys in **INSTRUCTORS** are connected with an (optional) arc and marked as non-transferable.

Propagating Changes (1)

- It is probably best to start the DB Design Transformer only when one is finished with the ER-design.
- If one has already changed the relational schema, and then changes the ER-schema and runs the DB Design Transformer again, it is a difficult problem to merge both changes into one version.
- In general, it is important that the ER-Schema and the relational schema remain in sync — otherwise the ER-schema loses its value as a documentation for the created tables.

Propagating Changes (2)

- Of course, if one has not yet worked on the relational schema, one can simply delete it and run the DB Design Transformer again.

Actually, it is not so simple to delete table definitions from the repository since they might be referenced in foreign keys. One must delete the foreign keys first. If one wants to delete all table definitions, one can click on the first, shift-click on the last, and then press the delete key. This will give an error message if a table is deleted that is still referenced by a foreign key. However, in Designer 6i (not Designer 6.0), one can choose to continue. After this is done, one simply presses “delete” again to remove the remaining tables (more runs might be needed, but if there are no cyclic foreign keys, finally all tables are deleted). In case of cyclic references, one must first delete at least one foreign key in the cycle before one can start to delete the tables.

Propagating Changes (3)

- Deleting the entire relational schema and running the DB Design Transformer again is the only completely automatic way that is guaranteed to keep both schemas in sync.
- The DB Design Transformer will never
 - remove existing tables (from a previous run) even if the corresponding entity type was deleted in the meantime,
 - remove columns from tables when the corresponding attribute was deleted.

Propagating Changes (4)

- The reason is probably that for denormalization, one could add columns and tables to the relational schema which are not present in the ER-schema.

This should be a big exception, only if the performance requirements cannot be met with a good schema. But in earlier times it was done quite often (programmer time was cheap compared to hardware).

- The DB Design Transformer protects this work.

The real reason probably is that in order to propagate deletions from the ER-schema to the relational schema, one must keep information about deleted schema elements. Also, the DB Design Transformer can be applied to a subset of the entity types. If one wants to delete tables, transforming the subset consisting of all entity types would be different from transforming the entire schema.

Contents

- 1 Database Design Transformer
- 2 Server Model Diagrams**
- 3 DB Admin.
- 4 DDL Generation

Design Editor (1)

- The relational schema generated by the database design transformer often still needs some work:
 - The names of the “relationship table” for many-to-many relationships often must be changed.
 - Column names and the sequence of columns within a table might need changes.
 - Often, some constraints are missing.

The DB Design transformer only generates keys, foreign keys, **NOT NULL**, and **CHECK** constraints for enumeration types or ranges automatically. Keys for one-to-one relationships are missing, as well as **CHECK**-constraints for subtypes, arcs, and other **CHECK**-constraints.

Design Editor (2)

- Manual work on the relational schema, continued:
 - For some foreign keys, one might have to select **“ON DELETE CASCADE”** etc.

A default can be specified in the settings of the DB design transformer, but it might be useful to consider each case individually. E.g. for weak entities **“ON DELETE CASCADE”** is probably right.

- **Warning:** Many students submitted the result of the DB Design Transformer as logical schema without ever reading it. They all lost points.

E.g. some generated column names are really ugly. Once application program development has started, it is difficult to change column or table names.

Design Editor (3)

- In addition, information necessary for the generation of application programs must be collected, e.g.
 - display title of the form generated for a table,
 - labels of input fields for columns,
 - field type (text, radio buttons, etc.),
 - field width,
 - help text,
 - display format (e.g. for date values),
 - columns that are not displayed.

Design Editor (4)

- After the logical design is finished, the following things must be defined:
 - Views.
 - Possibly triggers, stored procedures.
 - Users, table owners, access rights.
 - Physical design information.

E.g. indexes, storage parameters for tables, distribution of tables over disks/tablespaces, etc. It is quite likely that the physical design will need to change when it turns out that the assumptions about the system load were not quite right. However, changing it after the data was loaded can be quite a lot of work.

Design Editor (5)

- Although this information can be edited directly with the Repository Object Navigator, Oracle offers a special tool for all this work: The Design Editor.
- The Design Editor consists of four distinct tools:
 - Server Model (Relational Database Schema)
 - Modules (Application Programs)
 - DB Administration (Users, Tablespaces, etc.)
 - Distribution (for Distributed Databases)

Design Editor (6)

- Later, first-cut application programs (for Oracle Developer Forms, Visual Basic, etc.) will be generated from the “module definitions”.
- However, the module definitions contain only a link to the table name. The details such as the display width of input fields are defined in the server model (attached to tables).

Of course, some things such as the exact position of the input fields on the form cannot be generated, and must be later edited with the programming tool itself (e.g. Oracle Developer Forms).

Design Editor (7)

- One window of the Design Editor is the Server Model Navigator.
- It looks very similar to the Repository Object Navigator, but shows only objects that are part of the relational schema.

A student thought that she could remove the relational schema (for a fresh run of the DB Design Transformer) by selecting the application system name at the top of the Server Model Navigator window and pressing “Delete”. This removed her entire application system, not only the part shown in the window. For safety, export your design data at least once a day (with the Repository Object Navigator: “**Application→Export**”) and copy them on a floppy disk.

Design Editor (8)

- The Design Editor uses normally wizards/tabbed dialog boxes instead of the simple property palette in the Repository Object Navigator.

One can get also a property palette window under “**Tools**→**Property Palette**”.

- The Design Editor also contains a tool to put information about the schema of an existing relational database in the repository.

The “Design Capture Utility” (“**Generate**→**Capture Design of**→**Server Model**”) can read the information from the data dictionary of an Oracle Database, from a file with SQL DDL (Create Table) commands, or via the ODBC interface.

Design Editor (9)

- The Design Editor has also a “Server Model Guide” which shows a tree of all server model object types:
 - Domains
 - Tables (Indexes, Triggers, Constraints)
 - Constraints: Primary Keys, Foreign Keys, Unique Keys, Check.
 - Sequences
 - Advanced (Views, Snapshots, Clusters)
 - PL/SQL
 - Oracle8 (Collection Types, Object Types, Object Tables, Object Views).

Design Editor (10)

- When one selects an object type in the map, all objects of that type are shown. One can create, edit, or delete an object of the selected type.
- Basically, this is the same functionality as the “Server Model Navigator” which is also part of the Design Editor. Only the user interface is a bit different.

One can also choose that the two tools are linked: When an object is selected in the Server Model Guide, it is automatically also selected in the Server Model Navigator. The Server Model Guide gives more advice what to do in which sequence and sometimes has links to documentation.

Server Model Diagrams (1)

- The “Server Model” part of the Design Editor has a graphical interface showing tables and their foreign key connections in “Server Model Diagrams”.

The easiest way to create a diagram is to expand the “Relational Table Definitions” in the “Server Model Navigator” on the left, then to select the tables that should appear on the diagram (e.g. click on the first table and shift-click on the last) and then to select

“File→New→Server Model Diagram”.

- These diagrams are quite similar to ER-Diagrams.

However, the orientation on ER-diagrams is simpler. Server model diagrams are overloaded with information, table boxes are larger than entity boxes. Also many-to-many relationships are now shown as tables of their own, and foreign key columns do not appear on ER-diagrams.

Server Model Diagrams (3)

- Tables are shown as boxes with three sections:
 - The first section contains the table name and a number of buttons.

Buttons: “Database Triggers”, “Indexes”, “Database Synonyms”, “Primary Key”, “Unique Keys”, “Check Constraints”, “Foreign Keys”. A dimmed button means that the table has no object of that type. The button left to the table name is unusable.
 - The second section lists the columns (→ below).
 - The third section shows additional information as selected by the buttons in the first section.

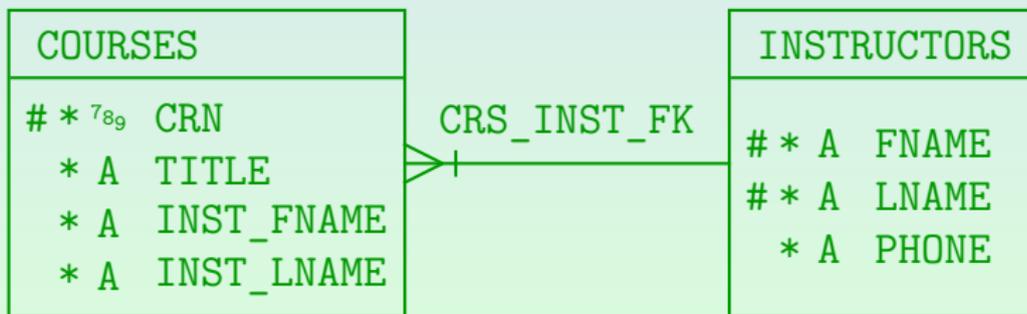
If one selects “View→Track Associations” and clicks on e.g. an index, the corresponding columns are shown inverted above.

Server Model Diagrams (4)

- The second section of the table box contains one row per column with the following information:
 - “#”: member of the primary key.
 - “*”: mandatory column (not null),
“o”: optional column.
 - “☰”: enumeration type value list.
 - “A”: character/string data type,
“789”: numeric data type.
 - “ $1\overline{2}_3$ ”: sequence (unique number generator).
 - “☰☰”: column belongs to domain.

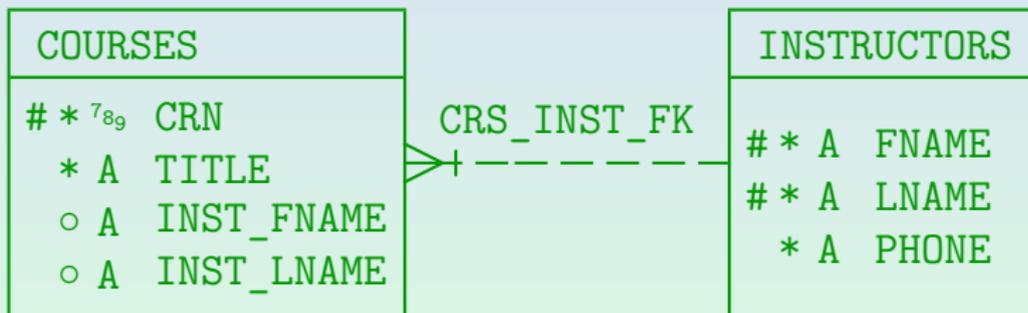
Server Model Diagrams (5)

- Foreign keys are shown as lines between the tables and use symbols similar to “one-to-many” relationships (but beware of the differences).
- Mandatory foreign keys (i.e. foreign keys that must be not null) are shown as solid lines:



Server Model Diagrams (6)

- Optional foreign keys (i.e. foreign keys that can be null) are shown as dashed lines:



- The entire line is either solid or dashed, there are no longer two halves.

Server Model Diagrams (7)

- Corresponding to the one-to-many relationship, the “crows foot” is on the side with the foreign key.

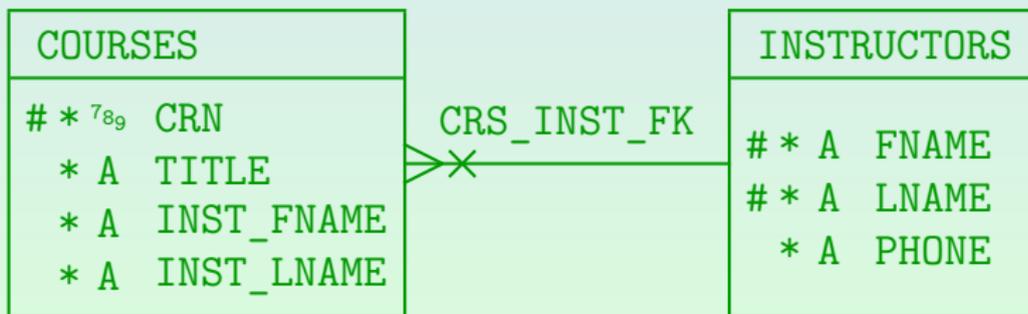
It can also be seen as indicating the direction of the pointer, although a real arrowhead would be on the opposite side.

- The names of the foreign keys are often not helpful, but take space on the diagram.

With “Options→Show/Hide” one can determine what is shown on the diagram. Removing the check mark from “Text” of “Associations” hides the foreign key names. One can specify which kinds of columns are shown, e.g. hide the foreign key columns on the diagram. One can also select which of the column type symbols are shown.

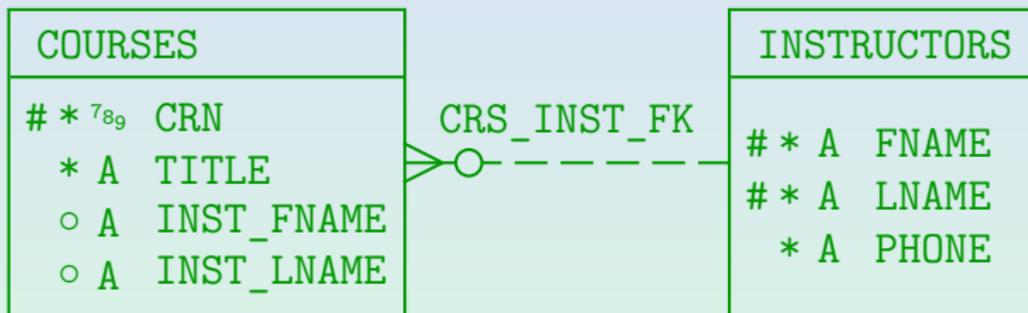
Server Model Diagrams (8)

- The small vertical bar near the crow's foot means that deletions do not cascade ("restricted", one cannot delete an instructor that teaches courses).
- If one selects "ON DELETE CASCADE", the line is crossed with an "x":



Server Model Diagrams (9)

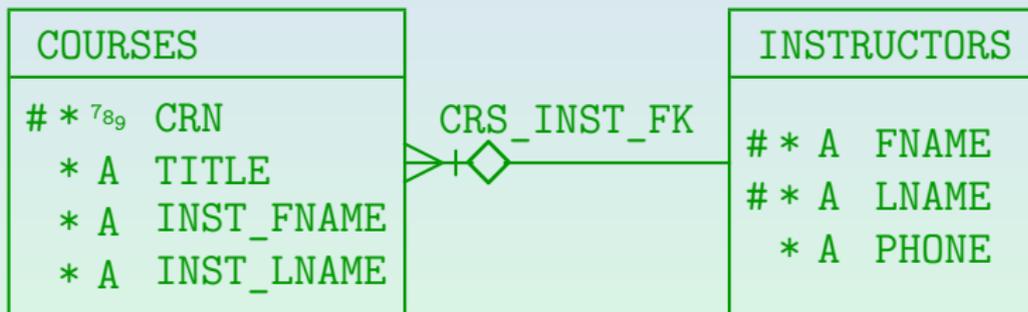
- If “ON DELETE SET NULL” is selected, a circle is used:



- A filled circle means “ON DELETE SET DEFAULT”.
- The cascade rule for “ON UPDATE” is not shown on the diagram.

Server Model Diagrams (10)

- One can also mark foreign keys as non updatable (corresponding to a non-transferable relationship):



- Foreign keys can be marked as mutually exclusive by means of arcs (as on ER-diagrams).

Server Model Diagrams (12)

- The properties dialog box for tables (“Edit Table”) has tabs “Name”, “Columns”, “Display”, “Controls”, “UI”.

Under “Display”, one can define which columns correspond to input fields in a form. Under “Controls” the type, size, etc. of these input fields is defined. Under “UI” (User Interface), more information about input fields is defined, e.g. a help text and a display format.

- Column and table names can be edited directly in the diagram, one does not have to go over the properties dialog box.

Server Model Diagrams (13)

- Tables also have the “Edit Text” dialog box, where one can define a description, notes, help text, and code for insert, update, delete, and locks.

One can open this dialog box from the menu that appears if one right-clicks on the table. This menu also permits to add columns, triggers, indexes, synonyms, keys, check constraints, foreign keys.

- The properties dialog box for foreign keys has tabs “Foreign Key Mandatory”, “Foreign Key Column”, “Cascade Rules”, “Validation”.

E.g. under “Validation” one can choose whether the constraint should be enforced on the server, the client, or both. One can also specify an error message and a table for exceptions (rows that violate it).

Server Model Diagrams (14)

- If one chooses to add e.g. a check constraint to a table, a wizard is opened that asks for the required information.

To edit it later, display it in the third part of the box (by clicking on the button for the object type) and click on the symbol in front of the name. Clicking on the name only permits to edit the name. Editing an existing check constraint etc. shows the same screens as the wizard, but now one can jump with tabs between them.

- Oracle Designer does not check the SQL syntax e.g. of **CHECK**-constraint definitions.

One can enter any text. Column names can be selected from a list. Of course, the exact SQL syntax depends on the DBMS.

Server Model Diagrams (15)

- The DB Design Transformer has already created indexes for foreign keys.

In addition, the DBMS automatically creates indexes for primary and alternate keys.

- As part of the physical design, one can add further indexes to a table.
- One can also add triggers (e.g. for enforcing complex constraints or logging changes to a table).

Server Model Diagrams (16)

- Server model diagrams can also contain other objects, such as

- Views (shown as grey-blue boxes).

In order to create a view, one can e.g. right-click on the background of a server model diagram. Alternatively there is also a symbol on the left toolbar. A wizard is started that asks the required information. One can select base tables (**FROM**) and columns (**SELECT**), and then any **WHERE** clause can be entered.

- Object types (shown as red boxes).

In Oracle, an object type is a generalization of a record/row type. One can create one or more tables over an object type. Object type and tables are connected with a line that ends in a diamond attached to the table.

Repository Reports

- Again, there are many repository reports which can be printed for documenting the DB Design, e.g.:
 - Entity to Table Implementation
 - Table Definition
 - Column Definition
 - Columns in Domain
 - Constraint Definition
 - Database Trigger
 - Cluster Definition
 - Tables, Columns, and Foreign Key Derivations

Contents

- 1 Database Design Transformer
- 2 Server Model Diagrams
- 3 DB Admin.**
- 4 DDL Generation

Database Administration (1)

- The following information can be specified with this part of the Design Editor:
 - Database Name and connection information.
 - Access information: Users, Roles, Profiles.
 - Storage Information: Tablespace, Datafiles, Logfiles, Rollback Segments, Directories.
- In the Server Model view, the really physical information (like storage parameters) was not yet asked.

Also, in the server model, the tables do not yet belong to users (there is no such property). One must move from the Server Model Relational Table Definitions to Table Implementations under “DB Admin”.

Database Administration (2)

- One now can “implement” the tables under a user account in a database. A new wizard asks for a table from the server model and takes the designer through the physical options.

One gets this wizard e.g. by selecting a user in the Database Administrator Guide, then selecting “Tables” and clicking on “Create”. This does not mean that a table is created from scratch, one can select a table from the Server Model. Since often several tables have the same storage parameters, one can create named sets of such parameters (“Storage Definitions”) and assign to tables.

- Of course, the resulting data are still stored in the repository. The table is not yet really implemented.

Database Administration (3)

- The Database Administration part of the Design Editor is only a subset of the Repository Object Navigator (there are no new diagrams).

However, again wizards/tabbed dialog boxes are used instead of the property palette. And it has a “Database Administrator Guide” that shows the steps for specifying a database.

- The tool is similar to a graphical user interface for a DBA (but stores all information in the repository).
- From the collected information, database creation scripts can be generated.

Contents

- 1 Database Design Transformer
- 2 Server Model Diagrams
- 3 DB Admin.
- 4 DDL Generation**

DDL Generation (1)

- The Database Design Transformer stores the relational schema in the repository. It does not actually create the tables.
- The reason for this is that in most cases, some things must still be changed/added manually.
- Once one is satisfied with the relational schema, one can generate SQL DDL code containing e.g. **CREATE TABLE** statements.

DDL = Data Definition Language. The generation is done with the Design Editor: “**Generate→Generate Database from Server Model**”.

DDL Generation (2)

- Oracle Designer can create DDL code for different DBMS: ANSI 92, DB2, Oracle (different versions), RDB7, SQL Server, Sybase.
- The creation of tables etc. can be done as follows:
 - Files with DDL statements are created, these must be executed manually in the target DB.
 - If the target database is an Oracle Database, Oracle Designer can directly create the tables.
 - If the target DB supports ODBC connections, tables can also be directly created.

DDL Generation (3)

- One can select for which schema objects should be generated (e.g. only a subset of the tables).

This is done on the “Objects” tab. E.g. click on the double right arrow:
“Generate All”.

- What can be generated, depends on the DBMS chosen, e.g.:
 - “ANSI 92”: Only tables and views.
 - “SQL Server” Only domains, tables, and views.

Which is strange, since it has indexes.

References

- Teorey: Database Modeling & Design, 3rd Edition. Morgan Kaufmann, 1999, ISBN 1-55860-500-2, ca. \$32.
- Rauh/Stickel: Konzeptuelle Datenmodellierung (in German), Teubner, 1997.
- Graeme C. Simsion, Graham C. Witt: Data Modeling Essentials, 2nd Edition. Coriolis, 2001, ISBN 1-57610-872-4, 459 pages.
- Barker: CASE*Method, Entity Relationship Modelling. Addison-Wesley, 1990, ISBN 0-201-41696-4, ca. \$61.
- Koletzke/Dorsey: Oracle Designer Handbook, 2nd Edition. ORACLE Press, 1998, ISBN 0-07-882417-6, ca. \$40.
- A. Lulushi: Inside Oracle Designer/2000. Prentice Hall, 1998, ISBN 0-13-849753-2, ca. \$50.
- Oracle/Martin Wykes: Designer/2000, Release 2.1.1, Tutorial. Part No. Z23274-02, Oracle, 1998.
- Oracle Designer Model, Release 2.1.2 (Element Type List).
- Oracle Designer Online Help System.