

Using ODBC

- The main ideas are presented via a sequence of four annotated C programs.
- These slides provide only supporting information.

Context:

- These notes deal with the Microsoft Visual C++ programming environment.
- The notes deal primarily with the Windows NT operating system.
- Although the examples are intended to be generic, they have been tested only with the Microsoft Access database system.

References:

- On-line documentation for ODBC is available at the Microsoft web site. Follow the link on the course home page.

The following hardcopy references are provided only as information for those with insatiable appetites for knowledge.

These notes, together with the accompanying sample programs and lectures, should provide enough information to write reasonable ODBC-based applications.

- One may also purchase hardcopy of the Microsoft documentation and software from booksellers. (*Microsoft ODBC 3.0 Software Development Kit and Programmer's Reference, Second Edition* -- Two books plus a CD, 1997)
- A decent book for the eager is *ODBC 3.5 Developers Guide*, by Roger E. Sanders, McGraw-Hill, 1999.
- A great reference????? If one were available at a reasonable price, it would become part of the course literature.

Interfacing to the Operating System:

The ODBC manager of operating system must be configured before database access is possible.

- The first step is to install the required DBMS-specific drivers. For Microsoft Access, this has already been done.
- The next step is to tell the ODBC manager about the database(s) to be accessed. In both Windows 95/98 and NT, this is accomplished via the ODBC icon in the control panel, although the details are a bit different. This explanation will focus upon Windows NT.

Configuring a data source under Windows NT:

- Open the ODBC menu in the control panel.
- Click on the User DSN tab.
- Highlight MS Access 97 Database and click on Add.
- From the menu in the new window, select Microsoft Access Driver (*.mdb), and click on Finish.
- From the menu in the new window, type in a data source name (e.g., Company), and optionally, a description. The click on either Select or Create, depending upon whether you want to link to an existing database, or create a new blank one.
- In the new window, give the path to the database.
- “OK” away the pile of subwindows; the new database should appear under the top-level ODBC User DSN tab.

Remarks:

- Put your database(s) under the User DSN tab. This will insure that they will accompany you from machine to machine, and that they will not interfere with those of others. Do **not** put your databases under the System DSN or File DSN tabs.

Using the development environments:

- The appropriate environment to use on the departmental machines is Microsoft Visual C++ 6.0 Professional.
- Do not use Borland C++ 4.5 or 5.0, or Borland C++ Builder 1.0. They do not have ODBC support.
- It is not practical to use a Unix environment, as ODBC drivers are not installed.
- If you have an office in the department, Borland C++ Builder 4.0 Professional may be installed there. You may use it, although you are yourself responsible for ensuring that your final submission is compatible with the Microsoft environment. See the slides from the 1999 course for a list of nuances of the Borland environment.
- Microsoft Visual C++ will happily let you program in C, even in your main file. Just use your C program as the main file.
- Make sure that you create a console application, and not a windows application, unless the latter is what you really want.

- **Some basics:**

Identifiers:

- Most ODBC identifiers begin with `SQL` (note the capitalization). Thus, it is a very good idea to avoid using this sequence as the beginning of user-defined identifiers.

API calls:

- ODBC contains a large number of functions (around 80). They have names like `SQLAllocHandle`, and `SQLCloseCursor`. Only a few will be used in this course.
- All (most?) return a value of type `SQLRETURN`. This value is zero if the execution was normal, and nonzero if it was special.

Includes:

- To run ODBC API calls, the following two includes must be issued:

```
#include <sql.h>
#include <sqlext.h>
```

Variable types:

There are three classes of variables associated with ODBC.

1. Types to be used as declarations to C. These begin with `SQL`, and continue with a sequence of capital letters, without underscores. They are `#defined` within the header files to be certain C types. Here are some of the principal ones:

ODBC Type	C Type
<code>SQLCHAR</code>	<code>char</code>
<code>SQLSCHAR</code>	<code>signed char</code>
<code>SQLINTEGER</code>	<code>long int</code>
<code>SQLUINTEGER</code>	<code>unsigned long int</code>
<code>SQLSMALLINT</code>	<code>short int</code>
<code>SQLUSMALLINT</code>	<code>unsigned short int</code>
<code>SQLREAL</code>	<code>float</code>
<code>SQLDOUBLE</code> , <code>SQLFLOAT</code>	<code>double</code>
<code>SQLDATE</code>	<code>a large struct..</code>

There are also a number of special ones for date, time etc., which correspond to `structs` in C.

The definitions are found in the library file `sqltypes.h`. Consult this file or the ODBC documentation for complete information.

For types involved in API calls, these types, rather than the C types, should be used.

2. C data type encodings. These are not true data types, but rather numerical encodings of the types listed in the previous group. These numerical encodings are used as arguments to API function calls. The following table lists some of the principal types.

Encoded Type	ODBC Type
SQL_C_CHAR	SQLCHAR
SQL_S_TINYINT	SQLSCHAR
SQL_C_SLONG	SQLINTEGER
SQL_C_ULONG	SQLUINTEGER
SQL_C_SSHORT	SQLSMALLINT
SQL_C_USHORT	SQLUSMALLINT
SQL_C_FLOAT	SQLREAL
SQL_C_DOUBLE	SQLDOUBLE , SQLFLOAT
SQL_C_TYPE_DATE	SQLDATE

The definitions for these types are found in the file `sqlext.h`. Consult that file or the ODBC documentation for further information.

It is important to remember that these are **not** C-language data types. They cannot be used in type declarations!!!

3. SQL data types encodings. These provide an association between the types allowed in SQL declarations, and those of the programming language. They are used in arguments to API calls, but never in variable declarations in the program itself.

These are not true data types, but rather numerical encodings which correspond to the numerical encodings of the types in the previous list. They cannot be used in type declarations!!!

The following table gives some principal examples.

Encoded Type	SQL Type
SQL_CHAR	Char (n)
SQL_VARCHAR	Varchar (n)
SQL_SMALLINT	Smallint
SQL_INTEGER	Integer
SQL_REAL	Real
SQL_DECIMAL	Decimal (p , s)
SQL_TYPE_DATE	Date

The exact mapping between these types and those of the previous table is implementation dependent.

The definitions for these types are found in the file `sqlext.h`. Consult that file or the ODBC documentation for further information.

Handles:

- Handles are numerical values which are associated with certain items.
- Example: File handles are familiar in operating system programming.

In ODBC, there are four types of handles:

- **Environment handles:** In order to access a database via ODBC, an ODBC environment must be established. There is normally only one such environment per program.
- **Connection handles:** Just as one must have a file handle for every open file in an operating system, so too must one have a connection handle for every ODBC database which is opened.
- **Statement handles:** A statement handle is associated with an SQL statement which is to be issued to an ODBC database for execution.
- **Descriptor handles:** Descriptors are metadata which describe formats associated with SQL statements. They will not be studied in this course.

In ODBC 3.0 and higher:

- Handles are declared using the type `SQLHANDLE`.
- Handles are allocated using the function `SQLAllocHandle`.
- Handles are freed using the function `SQLFreeHandle`.

The slides show examples of these activities.

Remark: There are older, ODBC 2 data types and calls which deal with each of the first three types of handles (all except descriptor handles) separately.

- The types are `HENV`, `HDBC`, and, `HSTMT`.
- The allocation functions are `SQLAllocEnv`, `SQLAllocEnv`, and `SQLAllocStmt`.
- The freeing functions are `SQLFreeEnv`, `SQLFreeEnv`, and `SQLFreeStmt`.

Although most ODBC implementations are backwards compatible with these calls, their use is to be discouraged in new software. (Translation: Do not use them in your project!)

Other important general operations:

- Inform the system of the ODBC version in use:
`SQLSetEnvAttr`.
- Connect to a database identified by an allocated connection handle: `SQLConnect`.
- Disconnect from the database allocated to a connection handle: `SQLDisconnect`.
 - The handle remains available for connection to another database.

Special query operations:

- Prepare ("compile") an SQL statement for execution: `SQLPrepare`.
- Execute a compiled SQL statement: `SQLExecute`.

Note: The function `SQLExecDirect` combines the above two functions, and is appropriate in situations in which the SQL statement is executed only once.

- Bind an input parameter index in an SQL statement with a variable in the program: `SQLBindParameter`.
- Bind a column of a query result (output parameter) to a variable in the program. `SQLBindCol`.
- Fetch the next tuple from the result of a query: `SQLFetchTuple`.
- Close the cursor on a given query, so that the statement handle may be used to collect the results of a new query: `SQLCloseCursor`.

Other classes of API calls:

- **Catalog queries:** Find out which relations are in a given database, what the types of the columns are, what the constraints are, etc.
- **Optimization directives:** Handle large queries with efficient batch operations.
- **Error management:** If something goes wrong, find out what the problem is.

All in all, there are over 80 API calls in ODBC.