



## **DataDirect Connect ODBC for dBASE**

Copyright 1998 INTERSOLV Inc. All rights reserved. INTERSOLV and DataDirect are registered trademarks of INTERSOLV, Inc. Connect ODBC is a trademark of INTERSOLV, Inc. Other company or product names mentioned herein may be trademarks or registered trademarks of their respective companies.

## About Connect ODBC for dBASE

Connect ODBC for dBASE (the “dBASE driver”) supports

- dBASE III, dBASE IV, and dBASE V files
- Clipper files
- FoxPro 2.5 and FoxPro 3.0 database containers (DBC)

The dBASE driver runs the SQL statements directly on dBASE-compatible files. You do not need to own the dBASE product to access these files.

The driver file name is IVDBF $nn$ .DLL, where  $nn$  denotes the revision level.

## Configuring Data Sources

To configure a dBASE data source, follow these steps.

- 1 Start the ODBC Administrator to display a list of data sources.
- 2 If you are configuring an existing data source, select the data source name and click **Configure** to display the [ODBC dBASE Driver Setup](#) dialog box.  
If you are configuring a new data source, click **Add** to display a list of installed drivers. Select the dBASE driver and click **Finish** to display the [ODBC dBASE Driver Setup](#) dialog box.
- 3 Specify a data source name and database directory. Specifying a database description and create type is optional. Click **Apply**.
- 4 Click the [Advanced tab](#) to configure optional data source settings, such as locking and file caching information. Click **Apply**.
- 5 Click **Translate** to display the Select Translator dialog box, which lists the translators specified in the ODBC Translators section of the system information. INTERSOLV provides a translator named INTERSOLV OEM ANSI that translates your data from the IBM PC character set to the ANSI character set.  
Select a translator, then click **OK** to close this dialog box and perform the translation.
- 6 If you use index files that have different names than their corresponding data files and you have not defined this association, click the **Define** tab. Follow the steps described in [Defining Index Attributes](#).
- 7 Click **OK** or **Cancel**. If you click **OK**, the values you have specified become the defaults when you connect to the data source. You can change these defaults by using this procedure to reconfigure your data source. You can override these defaults by [connecting to the data source using a connection string](#) with alternate values.

## Configuring Data Sources for FoxPro

To configure a data source for FoxPro 3.0 database containers, do the following:

- 1 Start the ODBC Administrator to display a list of data sources.
- 2 If you are configuring an existing data source, select the data source name and click **Configure** to display the [ODBC FoxPro Database Driver Setup](#) dialog box.  
If you are configuring a new data source, click **Add** to display a list of installed drivers. Select FoxPro 3.0 DBC and click **Finish** to display the [ODBC FoxPro Database Driver Setup](#) dialog box.
- 3 Specify a data source name and database directory. Specifying a database description is optional. You cannot change the Create Type for the FoxPro 3.0 DBC driver. Click **Apply**.
- 4 Click the [Advanced tab](#) to configure optional data source settings, such as locking and file caching information. Click **Apply**.
- 5 Click **Translate** to display the Select Translator dialog box, which lists the translators specified in the ODBC Translators section of the system information. INTERSOLV provides a translator named INTERSOLV OEM ANSI that translates your data from the IBM PC character set to the ANSI character set.  
Select a translator, then click **OK** to close this dialog box and perform the translation.
- 6 Click **Select** to display the Select FoxPro Database Container dialog box, where you can browse the available FoxPro 3.0 database containers.
- 7 If you use index files that have different names than their corresponding data files and you have not defined this association, click the **Define** tab on the dBASE Driver Setup. Follow the steps described in [Defining Index Attributes](#).
- 8 Click **OK** or **Cancel**. If you click **OK**, the values you have specified become the defaults when you connect to the data source. You can change these defaults by using this procedure to reconfigure your data source. You can override these defaults by [connecting to the data source using a connection string](#) with alternate values.

## Defining Index Attributes

The Define tab of the ODBC dBASE Driver Setup dialog box allows you to define the attributes of index files. With the exception of Clipper and dBASE III, the family of databases that includes dBASE and FoxPro uses a multiple index file associated with a particular table (database file). This index file has a .MDX or .CDX extension and is automatically maintained by the driver. Tags within this index can be marked as unique.

Clipper and dBASE III use single index files that are not automatically associated with a particular table. You can choose to have the driver maintain an index and choose whether or not the index is unique.

To define index file attributes:

- 1 Display the ODBC dBASE Driver Setup dialog box through the ODBC Administrator, then click the **Define** tab.

- 3 Click **Browse** to select the directory name and filename that contains the data file.

The upper section of the Define tab displays the directory name and the table containing the database information. The lower section of the tab displays the index information for that table.

- 4 For a displayed multiple index file, mark all tags in the Unique Tags list box that you want to be unique.

To define the attributes of a single index file, select the file from the Single Index File drop-down list. Check Maintain Index to associate this single index file with the selected table. Check Keep Unique to specify that the single index file is unique.

- 5 Click **Apply** or **OK** to save this information, or click **Cancel**. If you click **OK**, the values you have specified become the defaults when you connect to the data source. You can change these defaults by using this procedure to reconfigure your data source. You can override these defaults by connecting to the data source using a connection string with alternate values.

## Connecting to a Data Source Using a Connection String

If your application requires a connection string to connect to a data source, you must specify the data source name that tells the driver which section of the system information to use for the default connection information. Optionally, you may specify *attribute=value* pairs in the connection string to override the default values stored in the system information. These values are not written to the system information.

You can specify either long or short names in the connection string. The connection string has the form:

```
DSN=data_source_name[;attribute=value[;attribute=value]...]
```

An example of a connection string for dBASE is:

```
DSN=DBASE FILES;LCK=NONE;IS=0
```

The paragraphs that follow give the long and short names for each attribute, as well as a description. The defaults listed are initial defaults that apply when no value is specified in either the connection string or in the data source definition in the system information. If you specified a value for the attribute when configuring the data source, that value is your default.

**ApplicationUsingThreads (AUT):** ApplicationUsingThreads={0 | 1}. Ensures that the driver works with multi-threaded applications. The default is 1, which makes the driver thread-safe. When using the driver with single-threaded applications, you may set this option to 0 to avoid additional processing required for ODBC thread safety standards.

**CacheSize (CSZ):** The number of 64 KB blocks the driver uses to cache database records. The higher the number, the better the performance. The maximum number of blocks you can set depends on the system memory available. If the cache size is greater than 0, when browsing backwards, you will not be able to see updates made by other users until you run the Select statement again. The default is 4.

**CreateType (CT):** CreateType={dBASE3 | dBASE4 | dBASE5 | Clipper | FoxPro1 | FoxPro25 | FoxPro30}. Type of table or index to be created on a Create Table or Create Index statement. dBASE5 is the initial default.

**Note:** If you are using the FoxPro 3.0 DBC driver, you cannot change the CreateType attribute.

**Database (DB):** Directory in which the dBASE files are stored.

**DataFileExtension (DFE):** String of three or fewer characters that specifies the file extension to use for data files. The default DataFileExtension value is DBF. This value cannot be an extension the driver already uses, such as MDX or CDX. The DataFileExtension value is used for all Create Table statements. Sending a Create Table using an extension other than the DataFileExtension value causes an error.

In other SQL statements such as Select or Insert, users can specify an extension other than the DataFileExtension value. The DataFileExtension value is used when no extension is specified.

**DataSourceName (DSN):** String that identifies a dBASE data source configuration in the system information. Examples include "Accounting" or "dBASE Files."

**FileOpenCache (FOC):** Maximum number of unused file opens to cache. For example, when FileOpenCache=4, and a user opens and closes four files, the files are not actually closed. The driver keeps them open so that if another query uses one of these files, the driver does not have to perform another open, which is expensive. The advantage of using file open caching is increased performance. The disadvantage is that a user who tries to open the file exclusively may get a locking conflict even though no one appears to have the file open. The initial default is 0.

**IntISort (IS):** IntISort={0 | 1}. This attribute determines the order that records are retrieved when you issue a Select statement with an Order By clause. If IntISort=0 (the initial default), the driver uses the ASCII sort order. This order sorts items alphabetically with uppercase letters preceding lowercase letters. For example, "A, b, C" would be sorted as "A, C, b."

If IntISort=1, the driver uses the international sort order as defined by your operating system. This order is always alphabetic, regardless of case; the letters from the previous example would be sorted as "A, b, C." See your operating system documentation concerning the sorting of accented characters.

**LockCompatibility (LCOMP):** LockCompatibility={Q+E | Q+EVirtual | dBASE | Clipper | Fox}. Locking schemes to be used in your dBASE tables.

- LockCompatibility=Q+E specifies that locks be placed on the actual bytes occupied by the record. Only applications that use the dBASE driver can read and write to the database. Other applications are locked out of the table completely (they cannot even read other records). This locking is compatible with earlier versions of Q+E products.

- LockCompatibility=Q+EVirtual specifies that locks be placed on bytes beyond the physical end-of-file. Q+EVirtual is the same as Q+E except that other applications can open the table and read the data.

The advantage of using a Q+E locking scheme over dBASE locking is that, on Inserts and Updates, Q+E locks only individual index tags, while dBASE locks the entire index.

- LockCompatibility=dBASE specifies Borland-compatible locking. This is the initial default.
- LockCompatibility=Clipper specifies Clipper-compatible locking.
- LockCompatibility=Fox specifies FoxPro-compatible locking.

If you are accessing a table with an application that uses the dBASE driver, your locking scheme does not have to match the Create Type. However, if you are accessing a table with two applications, and only one uses the dBASE driver, set your locking scheme to match the other application. For example, you don't have to set LCOMP=Fox to work with a FoxPro table. But if you are using a FoxPro application simultaneously with an application using the dBASE driver on the same set of tables, set LCOMP=Fox to ensure that your data does not get corrupted.

**Locking (LCK):** Locking={NONE | RECORD | FILE} Level of locking for the database tables.

Locking=NONE offers the best performance but is intended only for single-user environments..

Locking=RECORD locks only the records affected by the statement. This is the initial default.

Locking=FILE locks all of the records in the table.

**ModifySQL (MS):** ModifySQL={0 | 1}. This attribute is provided for backward compatibility with earlier versions of Q+E products. It determines whether the driver modifies SQL statements to conform to ODBC specifications or passes the SQL statement directly to dBASE. Specify ModifySQL=1 to have the driver modify the SQL statement to conform to ODBC specifications. Specify ModifySQL=0 to have the driver understand SQL dialects found in earlier drivers. The default is 1.

**UltraSafeCommit (USF):** UltraSafeCommit={0 | 1}. A number that specifies when the driver flushes the file cache. If UltraSafeCommit=1, the driver updates directory entries after each Commit. This decreases performance. If UltraSafeCommit=0 (the default), the driver updates the directory entry when the file is closed. In this case, a machine "crash" before closing the file causes newly inserted records to be lost.

**UseLongNames (ULN):** UseLongNames={0 | 1}. This attribute specifies whether the driver uses long filenames as table names. The default is 0, do not use long filenames. If UseLongNames=1, the driver uses long filenames. The maximum table name length is specific to the environment in which you are running (for example, in Windows 9x, the maximum table name length is 128).

**UseLongQualifiers (ULQ):** UseLongQualifiers={0 | 1}. This attribute specifies whether the driver uses long pathnames as table qualifiers. With UseLongQualifiers set to 1, pathnames can be up to 255 characters. The default is 0; maximum pathname length is 128 characters.

## Data Types

The following table shows how dBASE data types map to the standard ODBC data types. These dBASE data types can be used in a Create Table statement (described in the General Help on DataDirect Connect ODBC).

**Note:** A few products can create dBASE files with numbers that do not conform to the precision and scale of the Number column. For example, these products can store 100000 in a column declared as NUMBER(5,2). When this occurs, the dBASE driver displays error 1244, "Unsupported decimal format." To remedy this situation, multiply the nonconforming column by 1, which converts it to the Float data type. For example,

```
SELECT BADCOL * 1 FROM BADFILE
```

BADCOL \* 1 is evaluated as an expression and is returned as a float value.

<b>DBASE</b>	<b>ODBC</b>
Binary1	SQL_LONGVARBINARY
Char2	SQL_CHAR
Date	SQL_TYPE_DATE
Float3	SQL_DECIMAL
General4	SQL_LONGVARBINARY
Logical	SQL_BIT
Memo	SQL_LONGVARCHAR
Numeric	SQL_DECIMAL
Picture5	SQL_LONGVARBINARY

1 dBASE V only

2 254 characters maximum (1024 for Clipper)

3 dBASE IV and V only

4 FoxPro 2.5 and dBASE V only

5 FoxPro and Clipper

### FoxPro 3.0 Data Types

Data types in FoxPro 3.0 tables and database containers map to the ODBC data types as follows:

<b>FoxPro 3.0</b>	<b>ODBC</b>
Character (binary)	SQL_CHAR
Currency	SQL_DOUBLE
Datetime	SQL_TYPE_TIMESTAMP
Double	SQL_DOUBLE
Integer	SQL_INTEGER
Memo (binary)	SQL_LONGVARBINARY

## SQL Statements for dBASE



If you are using the dBASE driver, you can use SQL statements (described in *SQL for Flat-File Drivers* in the General Help on DataDirect Connect ODBC) to read, insert, update, and delete records from a database, create new tables, and drop existing tables. In addition, you can use SQL statements and clauses that are specific to the dBASE driver. The dBASE-specific SQL information is described in the following topics.

[Select Statement](#)

[SQL Statements for FoxPro 3.0 Database Containers](#)

[Alter Table Statement](#)

[Create Index Statement](#)

[Drop Index Statement](#)

[Pack Statement](#)

## Select Statement

You use a SQL Select statement to specify the columns and records to be read. dBASE Select statements support all of the Select statement clauses (described in the General Help on DataDirect Connect ODBC).

### Column Names

The maximum length of a column name is 10 characters. A column name can contain alphanumeric characters and the hyphen character (-). The first character must be a letter (a through z).

### Rowid Pseudo-Column

Each dBASE record contains a special column named Rowid. This field contains a unique number that indicates the record's sequence in the database. For example, a table that contains 50 records has Rowid values from 1 to 50 (if no records are marked deleted). You can use Rowid in Where and Select clauses.

Rowid is particularly useful when you are updating records.

```
SELECT last_name, first_name, salary, rowid FROM emp
```

Then you can use the Rowid of the record that you want to update to ensure that you are updating the correct record and no other. For example,

```
UPDATE emp set salary = 40000 FROM emp WHERE rowid=21
```

The fastest way of updating a single row is to use a Where clause with the Rowid. You cannot update the Rowid column.

## SQL Statements for FoxPro 3.0 Database Containers

The FoxPro 3.0 DBC driver supports four additional SQL statements.

To create a new FoxPro 3.0 database container, use:

```
CREATE DATABASE database_name
```

To add an existing table to the database container, use:

```
ADD TABLE table_name
```

To remove a table from the database container (not delete the table, but unlink it from the database container), use:

```
REMOVE TABLE table_name
```

To set the current database container to an existing database container, use:

```
USE database_name
```

To add or delete columns from a table in a database container, use the [Alter Table statement](#).

## Alter Table Statement

The dBASE driver supports the Alter Table statement to add one or more columns to the table, or to delete (drop) a single column.

The Alter Table statement has the form:

```
ALTER TABLE table_name {ADD column_name data_type | ADD (column_name data_type [, column_name data_type] . . . ) | DROP [COLUMN] column_name}
```

*table\_name* is the name of the table for which you are adding or dropping columns.

*column\_name* assigns a name to the column you are adding or specifies the column you are dropping.

*data\_type* specifies the native data type of each column you add. See [dBASE Data Types](#) for more information.

For example, to add two columns to the emp table,

```
ALTER TABLE emp {ADD startdate date, dept char (10)}
```

You cannot add and drop columns in a single statement, and you can drop only one column at a time. For example, to drop a column,

```
ALTER TABLE emp DROP startdate
```

The Alter Table statement fails when you attempt to drop a column upon which other objects, such as indexes or views, are dependent.

## Create Index Statement

The type of index you create is determined by the value of the CreateType attribute, which you set in the setup dialog box or as a connection string option. The index can be:

- dBASE III (.NDX)
- dBASE IV (.MDX)
- Clipper (.NTX)
- FoxPro (.CDX)

The syntax for creating an index is:

```
CREATE [UNIQUE] INDEX index_name ON base_table_name
(field_name [ASC | DESC] [,field_name [ASC | DESC]] ...)
```

*index\_name* is the name of the index file. For FoxPro 2.5 and dBASE IV, this is a tag, which is required to identify the indexes in an index file. Each index for a table must have a unique name.

UNIQUE means that the driver creates an ANSI-style unique index over the column and ensures uniqueness of the keys. Use of unique indexes improves performance. ANSI-style unique indexes are different from dBASE-style unique indexes. With ANSI-style unique indexes, you receive an error message when you try to insert a duplicate value into an indexed field. With dBASE-style unique indexes, you do not see an error message when you insert a duplicate value into an indexed field. This is because only one key is inserted in the index file.

*base\_table\_name* is the name of the database file whose index is to be created. The .DBF extension is not required, the driver automatically adds it if it is not present. By default, dBASE IV index files are named *base\_table\_name*. MDX and FoxPro 2.5 indexes are named *base\_table\_name*.CDX.

*field\_name* is a name of a column in the dBASE table. You can substitute a valid dBASE-style index expression for the list of field names.

Asc tells dBASE to create the index in ascending order. Desc tells dBASE to create the index in descending order. By default, indexes are created in ascending order. You cannot specify both Asc and Desc orders within a single Create Index statement. For example, the following statement is invalid:

```
CREATE INDEX emp_i ON emp (last_name ASC, emp_id DESC)
```

The following table shows the attributes of the different index files supported by the dBASE driver. For each type supported, it provides the following details:

- Whether dBASE-style unique indexes are supported
- Whether descending order is supported
- The maximum size supported for key columns
- The maximum size supported for the column specification in the Create Index statement
- Whether production/structural indexes are supported

Create Type/Ext.	dBASE UNIQUE	DESC	Max Size of Key Column	Max Size of Column Spec.	Production/ Structural Indexes	Supports FOR Expressions
dBASE III .NDX	Yes	No	100	219	No	No
Clipper .NTX	Yes	Yes	250	255	No	Yes
dBASE IV and V .MDX	Yes	Yes	100	220	Yes	Yes
FoxPro .IDX*	Yes	Yes	240	255	No	Yes
FoxPro .CDX	Yes	Yes	240	255	Yes	Yes

\* Compact IDX indexes have the same internal structure as a tag in a CDX file. These indexes can be created if the IDX extension is included with the index name in the Create Index statement.

## Drop Index Statement

The syntax for dropping a dBASE index is as follows:

```
DROP INDEX table_name.index_name
```

*table\_name* is the name of the dBASE file without the extension.

For FoxPro 2.5 and dBASE IV, *index\_name* is the tag. Otherwise, *index\_name* is the name of the index file without the extension.

To drop the index EMPHIRE.NDX, issue the following statement:

```
DROP INDEX emp.emphire
```

## Pack Statement

When records are deleted from a dBASE file, they are not removed from the file. Instead, they are marked as having been deleted. Also, when memo fields are updated, space may be wasted in the files. To remove the deleted records and free the unused space from updated memo fields, you must use the Pack statement. It has the following form:

```
PACK filename
```

*filename* is the name of the dBASE file to be packed. The .DBF extension is not required; the driver automatically adds the extension if it is not present. For example,

```
PACK emp
```

You cannot pack a file that is opened by another user, and you cannot use the Pack statement in manual commit mode.

For the specified file, the Pack statement does the following:

- Removes all deleted records from the file
- Removes the entries for all deleted records from .CDX and .MDX files having the same name as the file
- Compresses unused space in the memo (.DBT or .FPT) file

## Locking

With the dBASE driver, you can build and run applications that share dBASE database files on a network. Whenever more than one user is running an application that accesses a shared database file, the applications should lock the records that are being changed. Locking a record prevents other users from locking, updating, or deleting the record.

### Levels of Database Locking

The dBASE driver supports three levels of database locking: None, Record, and File. You can set these levels in

- The connection string (LCK=)
- The setup dialog box

No locking offers the best performance but is intended only for single-user environments.

With record or file locking, the system locks the database tables during Insert, Update, Delete, or Select...For Update statements. The locks are released when the user commits the transaction. The locks prevent other users from modifying the locked objects, but they do not lock out readers.

With record locking, only records affected by the statement are locked. Record locking provides better concurrency with other users who also want to modify the table.

With file locking, all the records in the table are locked. File locking has lower overhead and may work better if records are modified infrequently, if records are modified primarily by one user, or if a large number of records are modified.

### Using Locks on Local Files

If you use database locking and are accessing files locally (not on a network), run the DOS utility SHARE.EXE before running Windows. If you add SHARE.EXE to your AUTOEXEC.BAT file, it runs automatically each time you boot your computer.

### Limit on Number of Locks

There is a limit on the number of locks that can be placed on a file. If you are accessing a dBASE file from a server, the limit depends on the server (see your server documentation). If you are accessing a dBASE file locally, the limit depends on the buffer space allocated when SHARE.EXE was loaded (see your DOS documentation). If you are exceeding the number of locks available, you may want to switch to file locking.

### How Transactions Affect Record Locks

When an Update or Delete statement is run, the driver locks the records affected by that statement. The locks are released after the driver commits the changes. Under manual commit mode, the locks are held until the application commits the transaction. Under autocommit mode, the locks are held until the statement is run.

When a Select...For Update statement is run, the driver locks a record only when the record is fetched. If the record is updated, the driver holds the lock until the changes are committed. Otherwise, the lock is released when the next record is fetched.



## **Isolation and Lock Levels Supported**

dBASE supports isolation level 1. It supports both file- and record-level locking.

## **ODBC Conformance Level**

The API functions supported are listed in *ODBC Functions Supported*, found in the General Help on DataDirect Connect ODBC. The dBASE driver also supports the Level 2 function SQLSetPos, backward and random fetching in SQLExtendedFetch and SQLFetch Scroll, and the minimum SQL grammar.

## **Number of Connections and Statements Supported**

dBASE supports multiple connections and multiple statements per connection.

## ODBC dBASE Driver Setup Dialog Box

Use the ODBC dBASE Driver Setup dialog box to [create](#) new dBASE data sources or [configure](#) existing data sources.

**Data Source Name:** A string that identifies this dBASE data source configuration in the system information. Examples include "Accounting" or "dBASE Files."

**Description:** An optional long description of a data source name. For example, "My Accounting Database" or "dBASE files in C:\ACCOUNTS."

**Database Directory:** A path specification to the directory that contains the database files. If none is specified, the current working directory is used.

**Create Type:** The type of table or index to be created on a Create Table or Create Index statement. Select dBASE III, dBASE IV, dBASE V, Clipper, FoxPro1, FoxPro25, or FoxPro30. The default is dBASE V.

### Advanced Tab

Displays the [Advanced tab](#), where you can configure optional data source settings, such as locking and file caching.

### Define Tab

Displays the [Define tab](#), where you can define the attributes for index files that are associated with a dBASE/FoxPro file.

**OK**

**Cancel**

**Apply**

## ODBC FoxPro Database Driver Setup Dialog Box

Use the ODBC FoxPro Database Driver Setup dialog to [create](#) new data sources for FoxPro 3.0 database containers (DBC) or to [configure](#) existing data sources.

**Data Source Name:** A string that identifies this FoxPro data source configuration in the system information. Examples include "Accounting" or "FoxPro Files."

**Description:** An optional long description of a data source name. For example, "My Accounting Database" or "FoxPro files in C:\ACCOUNTS."

**Database Name:** A path specification to the directory that contains the database files. If none is specified, the current working directory is used.

**Create Type:** You cannot change the Create Type for the FoxPro 3.0 DBC driver.

### Advanced Tab

Displays the [Advanced tab](#), where you can configure optional data source settings, such as locking and file caching.

### Select Button

Displays the Select FoxPro Database Container dialog box, where you can browse the available FoxPro 3.0 database containers.

[OK](#)

[Cancel](#)

[Apply](#)

## Advanced Tab, ODBC dBASE/FoxPro Driver Setup Dialog Box

To configure optional settings for a dBASE/FoxPro data source, specify values as follows:

**Locking:** The level of locking for the database file (File, Record, or None). File locks all of the records in the table. Record (the default) locks only the records affected by the statement. None offers the best performance but is intended only for single-user environments.

**Lock Compatibility:** The locking scheme the driver uses when locking records. Select dBASE, Q+E, Q+EVirtual, Clipper, or Fox. The default is dBASE. These values determine locking support as follows:

- dBASE specifies Borland-compatible locking.
- Q+E specifies that locks be placed on the actual bytes occupied by the record. Only applications that use the dBASE driver can read and write to the database. Other applications are locked out of the table completely (they cannot even read other records). This locking is compatible with earlier versions of Q+E products.
- Q+EVirtual specifies that locks be placed on bytes beyond the physical end-of-file. Q+EVirtual is the same as Q+E except that other applications can open the table and read the data.

The advantage of using a Q+E locking scheme over dBASE locking is that on Inserts and Updates, Q+E locks only individual index tags, while dBASE locks the entire index.

- Clipper specifies Clipper-compatible locking.
- Fox specifies FoxPro-compatible locking.

If you are accessing a table with an application that uses the dBASE driver, your locking scheme does not have to match the Create Type. However, if you are accessing a table with two applications, and only one uses the dBASE driver, set your locking scheme to match the other application. For example, you do not have to set this value to Fox to work with a FoxPro table. But if you are using a FoxPro application simultaneously with an application using the dBASE driver on the same set of tables, set this value to Fox to ensure that your data does not get corrupted.

**File Open Cache:** An integer value to specify the maximum number of unused file opens to cache. For example, the value 4 specifies that when a user opens and closes four tables, the tables are not actually closed. The driver keeps them open so that if another query uses one of these tables, the driver does not have to perform another open, which is expensive. The advantage of file open caching is increased performance. The disadvantage is that a user who specifies file locking on open may get a locking conflict even though no one appears to have the file open. The default is 0, which means no file open caching.

**Cache Size:** The number of 64 KB blocks the driver uses to cache database records. The higher the number, the better the performance. The maximum number of blocks you can set depends on the system memory available. If the cache size is greater than 0, when browsing backwards, you will not be able to see updates made by other users until you run the Select statement again. The default is 4.

**Data File Extension:** Specifies the file extension to use for data files. The default Data File Extension setting is DBF. The Data File Extension setting cannot be greater than three characters, and it cannot be one the driver already uses, such as MDX or CDX. The Data File Extension setting is used for all Create Table statements. Sending a Create Table using an extension other than the Data File Extension setting causes an error.

In other SQL statements such as Select or Insert, users can specify an extension other than the Data File Extension setting. The Data File Extension setting is used when no extension is specified.

**International Sort:** A setting to indicate the order in which records are retrieved when you issue a Select statement with an Order By clause. Clear this box to use ASCII sort order (the default setting). This order sorts items alphabetically with uppercase letters preceding lowercase letters. For example, "A, b, C" would be sorted as "A, C, b."

Select this box to use international sort order as defined by your operating system. This order is always alphabetic, regardless of case; the letters from the previous example would be sorted as "A, b, C." See your operating system documentation concerning the sorting of accented characters.

**Use Long Names:** Select this check box to use long filenames as table names. The maximum table

name length is specific to the environment in which you are running (for example, in Windows 9x, the maximum table name length is 128 characters).

**Use Long Qualifiers:** Select this check box to use long pathnames as table qualifiers. When you select this check box, pathnames can be up to 255 characters. The default length for pathnames is 128 characters.

**Application Using Threads:** A setting that ensures that the driver works with multi-threaded applications. You can clear this check box when using the driver with single-threaded applications. Turning off this setting avoids additional processing required for ODBC thread safety standards.

#### **Translate Button**

Displays the Select Translator dialog box, where you can translate your data from one character set to another. INTERSOLV provides a translator named INTERSOLV OEM ANSI that translates your data from the IBM PC character set to the ANSI character set.

**OK**

**Cancel**

**Apply**

## Define Tab, ODBC dBASE/FoxPro Driver Setup Dialog Box

Use the Define Tab to [define the attributes of index files](#) that are associated with a dBASE/FoxPro file.

**Directory:** Displays the name of the directory that contains the dBASE data file for which the index is defined.

**Table:** Displays the name of the database file for which the index is defined.

**Multiple Index File:** Displays any multiple index file (with a .CDX extension or .MCX extension) associated with the Table.

**Unique Tags:** Displays tags associated with the multiple index file. Select a tag to mark it as unique. You can select multiple tags.

**Single Index Files:** Displays a drop-down list of single index files. You select a file to define its attributes.

**Note:** The Single Index Files pane is active only if you have selected a dBASE III or Clipper table.

**Maintain Index:** If checked, associates the specified single index file with the selected Table.

**Keep Unique:** If checked, specifies that the single index file is unique.

**OK**

**Cancel**

**Apply**



## Apply Button

Writes the settings you have specified to the system information. These settings remain in effect until you change them in this dialog box. Clicking **Cancel** does not affect settings that have been applied.

**OK Button**

Writes the settings you have specified to the system information and closes the dialog box.

**Cancel Button**

Closes the dialog box without saving settings that have not been applied.

