The omniORB2 version 2.8 User's Guide

Sai-Lai Lo

(*email: slo@uk.research.att.com*)

David Riddoch

(*email: djr@uk.research.att.com*)

AT&T Laboratories Cambridge

22nd Sept, 1999

Changes and Additions, 22nd Sept 1999

- chapter 2 - added description on BOA and ORB shutdown.

- Chapter 4 - updates to command-line options

- Chapter 7 - updates to idle connection shutdown. Added desciption on the new remote call timeout feature.

Changes and Additions, 1st July 1999

- Chapter 1 - updates to features and setup information

- Chapter 3 - added notes on incompatibility with pre-omniORB 2.8.0 releases.

- Chapter 4 - updates to command-line options and configuration variables.

- Chapter 5 - new command-line option

- Chapter 9 - updates on the new Any and TypeCode behaviour and added notes on incompatibility with pre-omniORB 2.8.0 releases.

- Chapter 10 - added note to indicate that the implementation has not been updated to CORBA 2.3 yet.

- Chapter 11 - added note on the change in the default value of `CORBA::diiThrowsSysExceptions`.

- Chapter 12 - added note to indicate that the implementation has not been updated to CORBA 2.3 yet. Added note on new way to pass system exception back in the `invoke()` method.

Changes and Additions, 8th June 1999

- Chapter 9 - correction, recursive TypeCodes are supported

- Chapter 11 - correction, usage of CORBA::ORB::create_foo()

Changes and Additions, 22nd February 1999

- Chapter 11 - expanded and updated

Changes and Additions, 30 Sept 1998

- Chapter 1 - updates to features and setup information

- Chapter 2 - new example on tie implementation templates

- Chapter 4 - updates to command-line options and a new section on initial object reference bootstrapping.

- Chapter 5 - new section on-demand object loading

- Chapter 7 - new runtime configuration variable `omniORB::maxTcpConnectionPerServer`.

- New chapter - Dynamic Management of Any Values

Contents

Chapter 1

Introduction

OmniORB2 is an Object Request Broker (ORB) that implements the 2.3 specification of the Common Object Request Broker Architecture (CORBA) [OMG99a][1]. It has passed the Open Group CORBA compliant testsuite and is one of the three ORBs to have been granted the CORBA brand in June 1999[2].

This user guide tells you how to use omniORB2 to develop CORBA applications. It assumes a basic understanding of CORBA.

In this chapter, we give an overview of the main features of omniORB2 and what you need to do to setup your environment to run omniORB2.

1.1 Features

1.1.1 CORBA 2 compliant

OmniORB2 implements the Internet Inter-ORB Protocol (IIOP). This protocol provides omniORB2 the means of achieving interoperability with the ORBs implemented by other vendors. In fact, this is the native protocol used by omniORB2 for the communication amongst its objects residing in different address spaces. Moreover, the IDL to C++ language mapping provided by omniORB2 conforms to the latest revision of the CORBA specification. Type Any and TypeCode are now supported (introduced in version 2.5.0). DynAny is supported since 2.6.0. The Dynamic Invocation Interface and Dynamic Skeleton Interface are supported since 2.7.0. The C++ mapping has been updated to the CORBA 2.3 specification since 2.8.0.

1.1.2 Multithreading

OmniORB2 is fully multithreaded. To achieve low IIOP call overhead, unnecessary call-multiplexing is eliminated. At any time, there is at most one call in-flight in each communication channel between two address spaces. To do so without limiting the level of concurrency, new channels connecting the two address spaces are created on demand and cached when there are more concurrent calls in progress. Each channel

---

[1]Most of the 2.3 features have been implemented. The features still missing in this release is listed in 1.1.4. Where possible, backward compatibility has been maintained up to specification 2.0.

[2]More information can be found at *http://www.opengroup.org/press/7jun99_b.htm*

is served by a dedicated thread. This arrangement provides maximal concurrency and eliminates any thread switching in either of the address spaces to process a call. Furthermore, to maximise the throughput in processing large call arguments, large data elements are sent as soon as they are processed while the other arguments are being marshalled.

1.1.3 Portability

At AT&T Laboratories, the ability to target a single source tree to multiple platforms is very important. This is difficult to achieve if the IDL to C++ mapping for these platforms are different. We avoid this problem by making sure that only one IDL to C++ mapping is used. We run several flavours of Unices, Windows NT, Windows 95 and our in-house developed systems for our own hardware. OmniORB2 have been ported to all these platforms. **The IDL to C++ mapping for these targets are all the same**.

OmniORB2 uses real C++ exceptions and nested classes. We stay with the CORBA specification's standard mapping as much as possible and do not use the alternative mappings for C++ dialects. The only exception is the mapping of **modules**.

Starting with 2.6.0, the code generated by the IDL compiler of omniORB2 can be compiled using **C++ classes** or **namespaces** to represent IDL **modules** depending on the availability of namespace support in the compiler.

OmniORB2 relies on the native thread libraries to provide the multithreading capability. A small class library (omnithread [Richardson96a]) is used to encapsulated the (possibly different) APIs of the native thread libraries. In the application code, it is recommended but not mandatory to use this class library for thread management. It should be easy to port omnithread to any platform that either supports the POSIX thread standard or has a thread package that supports similar capabilities.

1.1.4 Missing features

OmniORB2 is not (yet) a complete implementation of the CORBA 2.3 core. The following is a list of the missing features.

- The BOA only support the persistent server activation policy. Other dynamic activation and deactivation policies are not supported.

- OmniORB2 does not has its own Interface Repository. However, it can act as a client to an IR.

- The POA is currently under development.

- The IDL types wchar, wstring, fixed, valuetype and native are not supported in this release.

These features may be implemented in the short to medium term. It is best to check out the latest status on the omniORB2 home page (`http://www.uk.research.att.com/omniORB/omniORB.html`).

1.2   Setting Up Your Environment

At AT&T Laboratories Cambridge, you should use the OMNI Development Environment (ODE) [Richardson96b] and the OMNI tree version 5.0 or above to compile your programs. If this is the case, there is no extra setup you have to do other than those described in the ODE documentation.

If you are running omniORB2 at other sites, you (or your system administrator) should install omniORB2 by following the instructions in the installation notes.

- On Unix platforms, the omniORB2 runtime looks for the environment variable `OMNIORB_CONFIG`. If this variable is defined, it contains the pathname of the omniORB2 configuration file. If the variable is not set, omniORB2 will use the compiled-in pathname to locate the file.

- On ARM/ATMos, the omniORB2 runtime looks for configuration information in the file `omniORB.cfg`.

- On Win32 platforms (Windows NT, Windows '95), omniORB2 first checks the environment variable (`OMNIORB_CONFIG`) to obtain the pathname of the configuration file. If this is not set, it then attempts to obtain configuration data in the system registry. It searches for the data under the key `HKEY_LOCAL_MACHINE\SOFTWARE\ORL\omniORB\2.0`

The configuration file is used to obtain an object reference for the COSS Naming Service. The entry in the configuration file should be specified in the following form:

`NAMESERVICE <stringified IOR for the COSS Naming Service>`

Comments in the configuration file should be prefixed with a #.

On Win32 platforms, the stringified IOR can be placed in the system registry, in the (string) value `NAMESERVICE`, under the key `HKEY_LOCAL_MACHINE\SOFTWARE\ORL\omniORB\2.0`.

Since 2.6.0, two other entries are supported:

```
ORBInitialHost <hostname string>
ORBInitialPort <port number (1-65535)>
```

The corresponding entries under the Win32 system registry is the key with name `ORBInitialHost` and `ORBInitialPort`.

The two entries provide information to the ORB to locate a bootstrap service at runtime. The bootstrap service is able to return the initial object reference for the COSS Naming Service and others. This is now the recommended way to configure omniORB2. More details are provided in section 4.6.

1.3   Platform specific variables

To compile omniORB2 programs correctly, several C++ preprocessor defines **must** be specified to identify the target platform.

Platform	CPP defines
Sun Solaris 2.5	`__sparc__  __sunos__  __OSVERSION__=5`
Digital Unix 3.2	`__alpha__  __osf1__    __OSVERSION__=3`
HPUX 10.x	`__hppa__  __hpux__    __OSVERSION__=10`
HPUX 11.x	`__hppa__  __hpux__    __OSVERSION__=11`
IBM AIX 4.x	`__aix__  __powerpc__    __OSVERSION__=4`
Linux 2.0 (x86)	`__x86__      __linux__  __OSVERSION__=2`
Linux 2.0 (alpha)	`__alpha__      __linux__  __OSVERSION__=2`
Linux 2.0 (powerpc)	`__powerpc__      __linux__  __OSVERSION__=2`
Windows/NT 3.5	`__x86__      __NT__      __OSVERSION__=3  __WIN32__`
Windows/NT 4.0	`__x86__      __NT__      __OSVERSION__=4  __WIN32__`
Windows/95	`__x86__  __WIN32__`
OpenVMS 6.x (alpha)	`__alpha__      __vms      __OSVERSION__=6`
OpenVMS 6.x (vax)	`__vax__      __vms      __OSVERSION__=6`
SGI Irix 6.x	`__mips__      __irix__      __OSVERSION__=6`
Reliant Unix 5.43	`__mips__      __SINIX__      __OSVERSION__=5`
ATMos 4.0	`__arm__      __atmos__  __OSVERSION__=4`
NextStep 3.x	`__m68k__  __nextstep__    __OSVERSION__=3`
Unixware 7	`__x86__      __uw7__      __OSVERSION__=5`

The preprocessor defines for new platform ports not listed above can be found in the corresponding platform configuration files. For instance, the platform configuration file for Sun Solaris 2.6 is in `mk/platforms/sun4_sosV_5.6.mk`. The preprocess defines to identify a platform is the value of the make variable `IMPORT_CPPFLAGS`.

In a single source multi-target environment, you can put the preprocessor defines as the command-line arguments for the compiler. Alternately, you could create a sitedef.h file in the same directory as `omniORB2/CORBA.h`. Write into the file the appropriate set of preprocessor defines and add `#include <omniORB2/sitedef.h>` at the beginning of `omniORB2/CORBA_sysdep.h`.

Chapter 2

The Basics

In this chapter, we go through three examples to illustrate the practical steps to use omniORB2. By going through the source code of each example, the essential concepts and APIs are introduced. If you have no previous experience with using CORBA, you should study this chapter in detail. There are pointers to other essential documents you should be familiar with.

If you have experience with using other ORBs, you should still go through this chapter because it provides important information about the features and APIs that are necessarily omniORB2 specific. For instance, the object implementation skeleton is covered in section 2.4.2.

2.1   The Echo Object Example

Our example is an object which has only one method. The method simply echos the argument string. We have to:

1. define the object interface in IDL;

2. use the IDL compiler to generate the stub code[1];

3. provide the object implementation;

4. write the client code.

The source code of this example is included in the last section of this chapter. A makefile written to be used under the OMNI Development Environment (ODE) [Richardson96b] is also included.

2.2   Specifying the Echo interface in IDL

We define an object interface, called Echo, as follows:

---

[1]The stub code is the C++ code that provides the object mapping as defined in the CORBA 2.0 specification.

```
interface Echo {
    string echoString(in string mesg);
};
```

If you are new to IDL, you can learn about its syntax in Chapter 3 of the CORBA specification 2.0 [OMG99a].

For the moment, you only need to know that the interface consists of a single operation, echoString, which takes a string as an argument and returns a copy of the same string.

The interface is written in a file, called `echo.idl`. If you are using ODE, all IDL files should have the same extension- `.idl` and should be placed in the `idl` directory of your export tree. This is done so that the stub code will be generated automatically and kept up-to-date with your IDL file.

For simplicity, the interface is defined in the global IDL namespace. This practice should be avoided for the sake of object reusability. If every CORBA developer defines their interfaces in the global IDL namespace, there is a danger of name clashes between two independently defined interfaces. Therefore, it is better to qualify your interfaces by defining them inside `module` names. Of course, this does not eliminate the chance of a name clash unless some form of naming convention is agreed globally. Nevertheless, a well-chosen module name can help a lot.

2.3   Generating the C++ stubs

From the IDL file, we use the IDL compiler to produce the C++ mapping of the interface. The IDL compiler for omniORB2 is called `omniidl2`. Given the IDL file, `omniidl2` produces two stub files: a C++ header file and a C++ source file. For example, from the file `echo.idl`, the following files are produced:

- `echo.hh`

- `echoSK.cc`

If you are using ODE, you don't need to invoke omniidl2 explicitly. In the example file `dir.mk`, we have the following line:

```
CORBA_INTERFACES = echo
```

That is all we need to instruct ODE to generate the stubs. Remember, you won't find the stubs in your working directory because all stubs are written into the `stub` directory at the top level of your build tree.

2.4   A Quick Tour of the C++ stubs

The C++ stubs conform to the mapping defined in the CORBA 2.0 specification (chapter 16-18). It is important to understand the mapping before you start writing any serious CORBA applications.

Before going any further, it is worth knowing what the mapping looks like.

2.4.1   Object Reference

The use of an object interface denotes an object reference. For the example interface Echo, the C++ mapping for its object reference is `Echo_ptr`. The type is defined in echo.hh. The relevant section of the code is reproduced below:

```
class Echo;
typedef Echo* Echo_ptr;

class Echo : public virtual omniObject, public vir-
tual CORBA::Object {
public:

  virtual char *  echoString ( const char *  mesg ) = 0;
  static Echo_ptr _nil();
  static Echo_ptr _duplicate(Echo_ptr);
  static Echo_ptr _narrow(CORBA::Object_ptr);

  ... // methods generated for internal use
};
```

In a compliant application, the operations defined in an object interface should **only** be invoked via an object reference. This is done by using arrow ("→") on an object reference. For example, the call to the operation `echoString` would be written as `obj→echoString(mesg)`.

It should be noted that the concrete type of an object reference is opaque, i.e. you must not make any assumption about how an object reference is implemented. In our example, even though `Echo_ptr` is implemented as a pointer to the class `Echo`, it should not be used as a C++ pointer, i.e. conversion to void*, arithmetic operations, and relational operations, including test for equality using **operation**== must not be performed on the type.

In addition to `echoString`, the mapping also defines three static member functions in the class Echo: `_nil`, `_duplicate`, and `_narrow`. Note that these are operations on an object reference.

The `_nil` function returns a nil object reference of the Echo interface. The following call is guaranteed to return TRUE:

```
CORBA::Boolean true_result = CORBA::is_nil(Echo::_nil());
```

Remember, `CORBA::is_nil()` is the only compliant way to check if an object reference is nil. You should not use the equality operator==.

The `_duplicate` function returns a new object reference of the Echo interface. The new object reference can be used interchangeably with the old object reference to perform an operation on the same object.

All CORBA objects inherit from the generic object `CORBA::Object`. `CORBA::Object_ptr` is the object reference for `CORBA::Object`. Any object reference is therefore conceptually inherited from `CORBA::Object_ptr`. In other words, an object reference such as `Echo_ptr` can be used in places where a `CORBA::Object_ptr` is expected.

The `_narrow` function takes an argument of the type `CORBA::Object_ptr` and returns a new object reference of the Echo interface. If the actual (runtime) type of the argument object reference can be widened to `Echo_ptr`, `_narrow` will return a valid object reference. Otherwise it will return a nil object reference.

To indicate that an object reference will no longer be accessed, you can call the `CORBA::release` operation. Its signature is as follows:

```
class CORBA {
    static void release(CORBA::Object_ptr obj);
    ... // other methods
};
```

You should not use an object reference once you have called `CORBA::release`. This is because the associated resources may have been deallocated. Notice that we are referring to the resources associated with the object reference and **not the object implementation**. Here is a concrete example, if the implementation of an object resides in a different address space, then a call to `CORBA::release` will only caused the resources associated with the object reference in the current address space to be deallocated. The object implementation in the other address space is unaffected.

As described above, the equality operator== should not be used on object references. To test if two object references are equivalent, the member function `_is_equivalent` of the generic object `CORBA::Object` can be used. Here is an example of its usage:

```
Echo_ptr A;
...              // initialised A to a valid object reference
Echo_ptr B = A;
CORBA::Boolean true_result = A->_is_equivalent(B);
// Note: the above call is guaranteed to be TRUE
```

You have now been introduced to most of the operations that can be invoked via `Echo_ptr`. The generic object `CORBA::Object` provides a few more operations and all of them can be invoked via `Echo_ptr`. These operations deal mainly with CORBA's dynamic interfaces. You do not have to understand them in order to use the C++ mapping provided via the stubs. For details, please read the CORBA specification [OMG99a] chapter 17.

Since object references must be released explicitly, their usage is prone to error and can lead to memory leakage. The mapping defines the **object reference variable** type to make life easier. In our example, the variable type `Echo_var` is defined[2].

The `Echo_var` is more convenient to use because it will automatically release its object reference when it is deallocated or when assigned a new object reference. For many operations, mixing data of type `Echo_var` and `Echo_ptr` is possible without any explicit operations or castings [3]. For instance, the operation `echoString` can be called using the arrow ("→") on a `Echo_var`, as one can do with a `Echo_ptr`.

The usage of `Echo_var` is illustrated below:

```
Echo_var a;
Echo_ptr p = ...      // somehow obtain an object reference

a = p;                // a assumes ownership of p, must not use p anymore

Echo_var b = a;       // implicit _duplicate

p = ...               // somehow obtain another object reference

a = Echo::_duplicate(p);   // release old object reference
                           // a now holds a copy of p.
```

2.4.2 Object Implementation

Unlike the client side of an object, i.e. the use of object references, the CORBA specification 2.0 deliberately leave many of the necessary functionalities to implement an object unspecified. As a consequence, it is very unlikely the implementation code of an object on top of two different ORBs can be identical. However, most of the code are expected to be portable. In particular, the body of an operation implementation can normally be ported with no or little modification.

OmniORB2 uses C++ inheritance to provide the skeleton code for object implementation. For each object interface, a skeleton class is generated. In our example, the skeleton class `_sk_Echo` is generated for the Echo IDL interface. An object implementation can be written by creating an implementation class that derives from the skeleton class.

The skeleton class `_sk_Echo` is defined in `echo.hh`. The relevant section of the code is reproduced below.

```
class _sk_Echo :  public virtual Echo {
public:
  _sk_Echo(const omniORB::objectKey& k);
  virtual char *  echoString ( const char *  mesg ) = 0;
  Echo_ptr        _this();
```

---

[2]In omniORB2, all object reference variable types are instantiated from the template type _CORBA_ObjRef_Var.

[3]However, the implementation of the type conversion operator() between Echo_var and Echo_ptr varies slightly among different C++ compilers, you may need to do an explicit casting when the compiler complains about the conversion being ambiguous.

```
  void             _obj_is_ready(BOA_ptr);
  void             _dispose();
  BOA_ptr          _boa();
  omniORB::objectKey _key();
  ... // methods generated for internal use
};
```

The code fragment shows the only member functions that can be used in the object implementation code. Other member functions are generated for internal use only. **Unless specified otherwise, the description below is omniORB2 specific.** The functions are:

echoString it is through this abstract function that an implementation class provides the implementation of the echoString operation. Notice that its signature is the same as the echoString function that can be invoked via the Echo_ptr object reference. **The signature of this function is specified by the CORBA specification**.

_this this function returns an object reference for the target object. The returned value must be deallocated via CORBA::release. See 2.7 for an example of how this function is used.

_obj_is_ready this function tells the Basic Object Adaptor[4] (BOA) that the object is ready to serve. Until this function is called, the BOA would not serve any incoming calls to this object. See 2.7 for an example of how this function is used.

_dispose this function tells the BOA to dispose of the object. The BOA will stop serving incoming calls of this object and remove any resources associated with it. See 2.7 for an example of how this function is used.

_boa this function returns a reference to the BOA that serves this object.

_key this function returns the key that the ORB used to identify this object. The type omniORB::objectKey is opaque to application code. The function omniORB::keyToOctetSequence can be used to convert the key to a sequence of octets.

2.5   Writing the object implementation

You define an implementation class to provide the object implementation. There is little constraint on how you design your implementation class except that it has to inherit from the stubs' skeleton class and to implement all the abstract functions defined in the skeleton class. Each of these abstract functions corresponds to an operation of the interface. They are hooks for the ORB to perform upcalls to your implementation.

Here is a simple implementation of the Echo object.

---

[4]The interface of a BOA is described in chapter 8 of the CORBA specification.

```
class Echo_i : public virtual _sk_Echo {
public:
  Echo_i() {}
  virtual ~Echo_i() {}
  virtual char * echoString(const char *mesg);
};


char *
Echo_i::echoString(const char *mesg) {
  char *p = CORBA::string_dup(mesg);
  return p;
}
```

There are three points to note here:

Storage Responsibilities A string, which is used as an IN argument and the return value of `echoString`, is a variable size data type. Other examples of variable size data types include sequences, type "any", etc. For these data types, you must be clear about who's responsibility to allocate and release their associated storage. As a rule of thumb, the client (or the caller to the implementation functions) owns the storage of all IN arguments, the object implementation (or the callee) must copy the data if it wants to retain a copy. For OUT arguments and return values, the object implementation allocates the storage and passes the ownership to the client. The client must release the storage when the variables will no longer be used. For details, please refer to Table 24-27 of the CORBA specification.

Multi-threading As omniORB2 is fully multithreaded, multiple threads may perform the same upcall to your implementation concurrently. It is up to your implementation to synchronise the threads' accesses to shared data. In our simple example, we have no shared data to protect so no thread synchronisation is necessary.

Instantiation You must not instantiate an implementation as automatic variables. Instead, you should always instantiate an implementation using the new operator, i.e. its storage is allocated on the heap. The reason behind this restriction will become clear in section 2.7.

2.6   Writing the client

Here is an example of how a `Echo_ptr` object reference is used.

```
void
hello(CORBA::Object_ptr obj)
{
  Echo_var e = Echo::_narrow(obj);                    // line 1

  if (CORBA::is_nil(e)) {                             // line 2
    cerr << "hello: cannot invoke on a nil object refer-
ence.\n" << endl;
    return;
```

```
  }

  CORBA::String_var src = (const char*) "Hello!";  // line 3
  CORBA::String_var dest;                          // line 4

  dest = e->echoString(src);                       // line 5

  cerr << "I said,\"" << src << "\"."
       << " The Object said,\"" << dest <<"\"" << endl;
}
```

Briefly, the function `hello` accepts a generic object reference. The object reference (`obj`) is narrowed to `Echo_ptr`. If the object reference returned by `Echo::_narrow` is not nil, the operation `echoString` is invoked. Finally, both the argument to and the return value of `echoString` are printed to cerr.

The example also illustrates how T_var types are used. As it was explained in the previous section, T_var types take care of storage allocation and release automatically when variables of the type are assigned to or when the variables go out of scope.

In line 1, the variable `e` takes over the storage responsibility of the object reference returned by `Echo::_narrow`. The object reference is released by the destructor of `e`. It is called automatically when the function returns. Line 2 and 5 shows how a `Echo_var` variable is used. As said earlier, `Echo_var` type can be used interchangeably with `Echo_ptr` type.

The argument and the return value of `echoString` are stored in `CORBA::String_var` variable `src` and `dest` respectively. The strings managed by the variables are deallocated by the destructor of `CORBA::String_var`. It is called automatically when the function returns. Line 5 shows how `CORBA::String_var` variables are used. They can be used in place of a string (for which the mapping is `char*`)[5]. As used in line 3, assigning a constant string (`const char*`) to a `CORBA::String_var` causes the string to be copied. On the otherhand, assigning a `char*` to a `CORBA::String_var`, as used in line 5, causes the latter to assume the ownership of the string[6].

Under the C++ mapping, T_var types are provided for all the non-basic data types. It is obvious that one should use automatic variables whenever possible both to avoid memory leak and to maximise performance. However, when one has to allocate data items on the heap, it is a good practice to use the T_var types to manage the heap storage.

2.7   Example 1 - Colocated Client and Implementation

Having introduced the client and the object implementation, we can now describe how to link up the two via the ORB. In this section, we describe an example in which both the client and the object implementation are in the same address space. In the next two sections, we shall describe the case where the two are in different address spaces.

---

[5]A conversion operator() of CORBA::String_var converts a CORBA::String_var to a char*.

[6]Please refer to the CORBA specification 16.7 for the details of the String_var mapping. Other T_var types are also covered in chapter 16.

The code for this example is reproduced below:

```
int
main(int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");   // line 1
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA"); // line 2

  Echo_i *myobj = new Echo_i();                                 // line 3
  myobj->_obj_is_ready(boa);                                    // line 4

  boa->impl_is_ready(0,1);                                      // line 5

  Echo_ptr myobjRef = myobj->_this();                           // line 6
  hello(myobjRef);                                              // line 7
  CORBA::release(myobjRef);                                     // line 8

  myobj->_dispose();                                            // line 9

  boa->destory();                                               // line 10

  orb->NP_destory();                                            // line 11

  return 0;
}
```

The example illustrates several important interactions among the ORB, the object implementation and the client. Here are the details:

2.7.1   ORB/BOA initialisation

line 1 The ORB is initialised by calling the CORBA::ORB_init function. The function uses the 3rd argument to determine which ORB should be returned. To use omniORB2, this argument must either be "omniORB2" or NULL. If it is NULL, there must be an argument, -ORBid "omniORB2", in argv. Like any command-line arguments understood by the ORB, it will be removed from argv when CORBA::ORB_init returns. Therefore, an application is not required to handle any command-line arguments it does not understand. If the ORB identifier is not "omniORB2", the initialisation will fail and a nil ORB_ptr will be returned. If supplied, omniORB2 also reads the configuration file omniORB.cfg. Among other things, the file provides a list of initial object references. One example of these object references is the naming service. Its use will be discussed in section 2.9.1. If any error occurs during the processing of the configuration file, the system exception CORBA::INITIALIZE is raised.

line 2 The BOA is initialised by calling the ORB's BOA_init. The 3rd argument must either be "omniORB2_BOA" or NULL. If it is NULL, then argv must contain an argument, -BOAid "omniORB2_BOA". If the BOA identifier is not "omniORB2_BOA", the initialisation will fail and a nil BOA_ptr will be returned. Like ORB_init, any command-line arguments understood by BOA_init will be removed from argv.

2.7.2   ORB/BOA shutdown

line 10 Shutdown the BOA permanently. This call causes the ORB to release all its BOA resources, e.g. worker threads.

line 11 Shutdown the ORB permanently. This call causes the ORB to release its internal resources, e.g. internal threads.

These two calls are particularly important when writing a CORBA DLL on Windows NT that is to be used from ActiveX. If these calls are absent, the application will hang when the CORBA DLL is unloaded.

2.7.3   Object initialisation

line 3 An instance of the Echo object is initialised using the `new` operator.

line 4 The object's `_obj_is_ready` is called. This function informs the BOA that this object is ready to serve. Until this function is called, the BOA will not accept any invocation on the object and will not perform any upcall to the object.

line 5 The BOA's `impl_is_ready` is called. This function tells the BOA the implementation is ready. After this call, the BOA will accept IIOP requests from other address spaces. There are 2 points to note here:

1. `boa→impl_is_ready` can be called any time after `BOA_init` is called (line 2). In other words, object instances can be initialised and advertised to the BOA before or after this function is called.

2. The 2nd argument[7] to `impl_is_ready` tells the ORB whether this call should be non-blocking. The default value of this argument is FALSE(0) and the call will block indefinitely within the ORB. If there are more things the main thread should do after it calls `impl_is_ready`, as it is the case in this example, the non-blocking option (TRUE=1) should be specified. Whether the main thread blocks in this call or not, the ORB is not affected because its functions are provided by other threads spawned internally. Notice that the signature of `impl_is_ready` in the CORBA specification does not have the 2nd argument[8]. Therefore, calling `impl_is_ready` with the non-blocking option is omniORB2 specific.

2.7.4   Client invocation

line 6 The object reference is obtained from the implementation by calling `_this`. Like any object reference, the return value of `_this` must be released by `CORBA::release` when it is no longer needed.

line 7 Call `hello` with this object reference. The argument is widened implicitly to the generic object reference `CORBA::Object_ptr`.

---

[7]The 1st argument is a pointer to the implementation definition and is always ignored by omniORB2.
[8]The CORBA specification does not specify when `impl_is_ready` should return. Many ORB vendors choose to implement `impl_is_ready` as blocking until a certain time-out value is exceeded. In a single threaded implementation this is necessary to give the ORB the time to serve incoming requests.

line 8 Release the object reference.

One of the important characteristic of an object reference is that it is completely location transparent. A client can invoke on the object using its object reference without any need to know whether the object is colocated in the same address space or resided in a different address space.

In case of colocated client and object implementation, omniORB2 is able to short-circuit the client calls to direct calls on the implementation methods. The cost of an invocation is reduced to that of a function call. This optimisation is applicable **not only** to object references returned by the _this function but to any object references that are passed around within the same address space or received from other address spaces via IIOP calls.

2.7.5 Object disposal

line 9 To dispose of an object implementation and release all the resources associated with it, the _dispose function is called. In fact, this is the **only** clean way to get rid of an object implementation. Even though the object is created using the new operator in the application code, the application should never call the delete operator on the object directly.

Once an application calls _dispose on an object implementation, the pointer to the object should not be used any more. At the time the _dispose call is made, there may be other threads invoking on the object, omniORB2 ensures that all these calls are completed before removing the object from its internal tables and releasing the resources associated with it. The storage associated with the object is released by omniORB2 using the delete operator. This is why all object implementation should be initialised using the new operator (section 2.5).

The disposal of an object implementation by omniORB2 may also be deferred when **colocated** clients continue to hold on to copies of the object's reference[9]. This behavior is to prevent the short-circuited calls from the clients to fail unpredictably.

To summarise, an application can make no assumption as to when the object is disposed by omniORB2 after the _dispose call returns. If it is necessary to have better control on when to stop serving incoming requests, the work should be done by the object implementation itself, such as by keeping track of the current serving state.

2.8 Example 2 - Different Address Spaces

In this example, the client and the object implementation reside in two different address spaces. The code of this example is almost the same as the previous example. The only difference is the extra work need to be done to pass the object reference from the object implementation to the client.

---

[9]Object references held by clients in other address spaces will not prevent the object implementation from being disposed of. If these clients invoke on the object after it is disposed, the system exception INV_OBJREF is raised.

The simplest (and quite primitive) way to pass an object reference between two address spaces is to produce a stringified version of the object reference and to pass this string to the client as a command-line argument. The string is then converted by the client into a proper object reference. This method is used in this example. In the next example, we shall introduce a better way of passing the object reference using the COS Naming Service.

2.8.1   Object Implementation: Generating a Stringified Object Reference

The `main` function of the object implementation side is reproduced below. The full listing (`eg2_impl.cc`) can be found at the end of this chapter.

```
int
main(int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA");

  Echo_i *myobj = new Echo_i();
  myobj->_obj_is_ready(boa);

  {
    Echo_var myobjRef = myobj->_this();
    CORBA::String_var p;

    p = orb->object_to_string(myobjRef);            //line 1

    cerr << "'" << (char*)p << "'" << endl;
  }

  boa->impl_is_ready();    // block here indefinitely
                           // See the explanation in example 1
  return 0;
}
```

The stringified object reference is obtained by calling the ORB's function `_object_to_string` (line 1). This is a sequence starting with the signature "IOR:" and followed by a hexadecimal string. All CORBA 2.0 compliant ORBs are able to convert the string into its internal representation of a so-called Interoperable Object Reference (IOR). The IOR contains the location information and a key to uniquely identify the object implementation in its own address space[10]. From the IOR, an object reference can be constructed.

2.8.2   Client: Using a Stringified Object Reference

The stringified object reference is passed to the client as a command-line argument. The client uses the ORB's function `string_to_object` to convert the string into a generic object reference (`CORBA::Object_ptr`). The relevant section of the code is

---

[10]Notice that the object key is not globally unique across address spaces.

reproduced below. The full listing (`eg2_clt.cc`) can be found at the end of this chapter.

```
try {
  CORBA::Object_var obj = orb->string_to_object(argv[1]);
  hello(obj);
}
catch(CORBA::COMM_FAILURE& ex) {
  ... // code to handle communication failure
}
```

2.8.3 Catching System Exceptions

When omniORB2 detects an error condition, it may raise a system exception. The CORBA specification defines a series of exceptions covering most of the error conditions that an ORB may encounter. The client may choose to catch these exceptions and recover from the error condition[11]. For instance, the code fragment, shown in section 2.8.2, catches the system exception COMM_FAILURE which indicates that communication with the object implementation in another address space has failed.

All system exceptions inherit from the class `CORBA::SystemException`. With compilers that support RTTI[12][13], a single catch `CORBA::SystemException` will catch all the different system exceptions thrown by omniORB2.

When omniORB2 detects an internal inconsistency that is most likely to be caused by a bug in the runtime, it raises the exception `omniORB::fatalException`. When this exception is raised, it is not sensible to proceed with any operation that involves the ORB's runtime. It is best to exit the program immediately. The exception structure carries by `omniORB::fatalException` contains the exact location (the file name and the line number) where the exception is raised. You are strongly encourage to file a bug report and point out the location.

2.8.4 Lifetime of an Object Implementation

It may be obvious but it has to stated that an object implementation exists only for the duration of the process's lifetime. When the same program is run again, a different instance of the object implementation is created. More significantly, **the IOR, and hence the object reference, of this instance is different from that of the previous run**.

For instance, if you look at the stringified object reference produced by the program `eg2_impl` in different runs, they are all different. The implication is that you cannot store away the stringified object reference and expect to be able to use it again later when the original program run has terminated.

For system services and other applications, it may be desirable to have "persistent" object implementations. The objects are "persistent" in the sense that they can

---

[11]If a system exception is not caught, the C++ runtime will call the `terminate` function. This function is defaulted to abort the whole process and on some system will cause a core file to be produced.

[12]Run Time Type Identification

[13]A noticeable exception is the GNU C++ compiler (version 2.7.2). It doesn't support RTTI unless the compilation flag -frtti is specified. The omniORB2 runtime is not compiled with the -frtti flag. It is said that RTTI will be properly supported in the upcoming version 2.8.

be contacted using the same IOR when they are instantiated in different program runs. To provide this functionality, omniORB2 needs to be provided with two pieces of information: the (network) location and the object key. The details of how this can be done will be described in the later part of this manual.

Alternatively, an indirection from textual pathnames to object references can be used. Applications can register object implementations at runtime to a naming service and bind them to fixed pathnames. Clients can bind to the object implementations at runtime by asking the naming service to resolve the pathnames to the object references. CORBA defines a naming service, which is a component of the Common Object Services (COS) [OMG99b], that can be used for this purpose. The next section describes an example of how to use the COS Naming Service.

2.9   Example 3 - Using the COS Naming Service

In this example, the object implementation uses the COS Naming Service [OMG99b] to pass on the object reference to the client. This method is by-far more practical than using stringified object references. The full listing of the object implementation (`eg3_impl.cc`) and the client (`eg3_clt.cc`) can be found at the end of this chapter.

The object reference is bound to the pathname "**test/Echo**"[14]. The pathname consists of the context **test** and the object name **Echo**. Both the context and the object name has an attribute **kind**. This attribute is a string that is intended to be used to describe the name in a syntax-independent way. The naming service does not interpret, assign, or manage these values. However both the name and the kind attribute must match for a name lookup to succeed. In this example, the **kind** values for **test** and **Echo** are chosen to be "my_context" and "Object" respectively. This is an arbitrary choice for there is no standardised set of kind values.

2.9.1   Obtaining the Root Context Object Reference

The initial contact with the Naming Service can be established via what we called the **root** context. The object reference to the root context is provided by the ORB and can be obtained by calling `resolve_initial_references`. The following code fragment shows how it is used:

```
CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");

CORBA::Object_var initServ;
initServ = orb->resolve_initial_references("NameService");

CosNaming::NamingContext_var rootContext;
rootContext = CosNaming::NamingContext::_narrow(initServ);
```

---

[14]A pathname, or in the Naming Service's terminology- a *compound name*, is a sequence of textual names. Each name component except the last one is bound to a naming context. A naming context is analogous to a directory in a filing system, it can contain names of object references or other naming contexts. The last name component is bound to an object reference. Note: '/' is purely a notation to separate two components in the pathname. It does not appear in the *compound name* that is registered with the Naming Service.

Remember, omniORB2 constructs its internal list of initial references at initialisation time using the information provided in the configuration file `omniORB.cfg`. If this file is not present, the internal list will be empty and `resolve_initial_references` will raise a CORBA::ORB::InvalidName exception.

2.9.2 The Naming Service Interface

It is beyond the scope of this chapter to describe in detail the Naming Service interface. You should consult the CORBAservices specification [OMG99b] (chapter 3). The code listed in `eg3_impl.cc` and `eg3_clt.cc` are good examples of how the service can be used. Please spend time to study the examples carefully.

2.10 Example 4 - Using tie implementation templates

Since 2.6.0, a new command-line option (`-t`) has been added to `omniidl2`. When this flag is specified, `omniidl2` generates an extra template class for each interface. This template class can be used to tie a C++ class to the skeleton class of the interface.

The source code in `eg3_tieimpl.cc` at the end of this chapter illustrates how the template class can be used. The code is almost identical to `eg3_impl.cc` with only a few changes.

Firstly, the implementation class `Echo_i` does not inherit from any stub classes. This is the main benefit of using the template class because there are applications in which it is difficult to require every implementation class to subclass from CORBA classes.

Secondly, the instantiation of a CORBA object now involves creating an instance of the implementation class **and** an instance of the template. Here is the relevant code fragment:

```
class Echo_i { ... };

Echo_i *myimpl = new Echo_i();
_tie_Echo<Echo_i,1> *myobj = new _tie_Echo<Echo_i,1>(myimpl);
myobj->_obj_is_ready(boa);
```

For interface `Echo`, the name of its tie implementation template is `_tie_Echo`. The first template parameter is the implementation class that contains an implementation of each of the operations defined in the interface. The second template parameter is a boolean flag. When the flag is TRUE (1), as it is in this example, the ORB would call `delete` on the implementation object (`myimpl`) when `_dispose` is invoked on `myobj`. When the flag is FALSE (0), `delete` would not be called. Instantiating this template with the flag set to FALSE is useful when the same implementation class is used to implement multiple interfaces. In this situation, the same implementation would be used as the argument to a number of tie template instantiations. Provided that only one of the instantiation has the flag set to TRUE, the object implementation would not be deleted more than once!

2.11   Source Listing

2.11.1   echo_i.cc

```
// echo_i.cc - This source code demonstrates an implmentation of the
//             object interface Echo. It is part of the three examples
//             used in Chapter 2 "The Basics" of the om-
niORB2 user guide.
//
#include <string.h>
#include "echo.hh"

class Echo_i : public virtual _sk_Echo {
public:
  Echo_i() {}
  virtual ~Echo_i() {}
  virtual char * echoString(const char *mesg);
};

char *
Echo_i::echoString(const char *mesg) {
  char *p = CORBA::string_dup(mesg);
  return p;
}
```

2.11.2   greeting.cc

```
// greeting.cc - This source code demonstrates the use of an object
//               reference by a client to perform an operation on an
//               object. It is part of the three examples used
//               in Chapter 2 "The Basics" of the om-
niORB2 user guide.
//
#include <iostream.h>
#include "echo.hh"

void
hello(CORBA::Object_ptr obj)
{
  Echo_var e = Echo::_narrow(obj);

  if (CORBA::is_nil(e)) {
    cerr << "hello: cannot invoke on a nil object refer-
ence.\n" << endl;
    return;
  }

  CORBA::String_var src = (const char*) "Hello!";
  CORBA::String_var dest;

  dest = e->echoString(src);

  cerr << "I said,\"" << src << "\"."
       << " The Object said,\"" << dest <<"\"" << endl;
}
```

2.11.3   eg1.cc

```
// eg1.cc - This is the source code of example 1 used in Chapter 2
//          "The Basics" of the omniORB2 user guide.
//
//          In this example, both the object implementation and the
//          client are in the same process.
//
// Usage: eg1
//
#include <iostream.h>
#include "echo.hh"

#include "echo_i.cc"
#include "greeting.cc"

int
main(int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA");

  Echo_i *myobj = new Echo_i();
  // Note: all implementation objects must be instantiated on the
  // heap using the new operator.

  myobj->_obj_is_ready(boa);
  // Tell the BOA the object is ready to serve.
  // This call is omniORB2 specific.
  //
  // This call is equivalent to the following call sequence:
  //     Echo_ptr myobjRef = myobj->_this();
  //     boa->obj_is_ready(myobjRef);
  //     CORBA::release(myobjRef);

  boa->impl_is_ready(0,1);
  // Tell the BOA we are ready and to return immediately once it has
  // done its stuff. It is omniORB2 specific to call impl_is_ready()
  // with the extra 2nd argument- CORBA::Boolean NonBlocking,
  // which is set to TRUE (1) in this case.

  Echo_ptr myobjRef = myobj->_this();
  // Obtain an object reference.
  // Note: always use _this() to obtain an object reference from the
  //       object implementation.

  hello(myobjRef);

  CORBA::release(myobjRef);
  // Dispose of the object reference.

  myobj->_dispose();
  // Dispose of the object implementation.
```

```
  // This call is omniORB2 specific.
  // Note: *never* call the delete operator or the dtor of the object
  //       directly because the BOA needs to be informed.
  //
  // This call is equivalent to the following call sequence:
  //      Echo_ptr myobjRef = myobj->_this();
  //      boa->dispose(myobjRef);
  //      CORBA::release(myobjRef);

  boa->destory();

  orb->NP_destory();

  return 0;
}
```

2.11.4   eg2_impl.cc

```
// eg2_impl.cc - This is the source code of example 2 used in Chap-
ter 2
//                "The Basics" of the omniORB2 user guide.
//
//                This is the object implementation.
//
// Usage: eg2_impl
//
//        On startup, the object reference is printed to cerr as a
//        stringified IOR. This string should be used as the argu-
ment to
//        eg2_clt.
//
#include <iostream.h>
#include "omnithread.h"
#include "echo.hh"

#include "echo_i.cc"

int
main(int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA");

  Echo_i *myobj = new Echo_i();
  myobj->_obj_is_ready(boa);

  {
    Echo_var myobjRef = myobj->_this();
    CORBA::String_var p = orb->object_to_string(myobjRef);
    cerr << "'" << (char*)p << "'" << endl;
  }

  boa->impl_is_ready();
  // Tell the BOA we are ready. The BOA's default be-
haviour is to block
  // on this call indefinitely.

  // Call boa->impl_shutdown() from another thread would unblock the
  // main thread from impl_is_ready().
  //
  // To properly shutdown the BOA and the ORB, add the follow-
ing calls
  // after impl_is_ready() returns.
  //
  // boa->destroy();
  // orb->NP_destory();

  return 0;
}
```

2.11.5 eg2_clt.cc

```cpp
// eg2_clt.cc - This is the source code of example 2 used in Chap-
ter 2
//             "The Basics" of the omniORB2 user guide.
//
//             This is the client. The object refer-
ence is given as a
//             stringified IOR on the command line.
//
// Usage: eg2_clt <object reference>
//
#include <iostream.h>
#include "echo.hh"

#include "greeting.cc"

extern void hello(CORBA::Object_ptr obj);

int
main (int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA");

  if (argc < 2) {
    cerr << "usage: eg2_clt <object reference>" << endl;
    return 1;
  }

  try {
    CORBA::Object_var obj = orb->string_to_object(argv[1]);
    hello(obj);
  }
  catch(CORBA::COMM_FAILURE& ex) {
    cerr << "Caught system exception COMM_FAILURE, unable to con-
tact the "
         << "object." << endl;
  }
  catch(omniORB::fatalException& ex) {
    cerr << "Caught omniORB2 fatalException. This indi-
cates a bug is caught "
         << "within omniORB2.\nPlease send a bug report.\n"
         << "The exception was thrown in file: " << ex.file() << "\n"
         << "                                  line: " << ex.line() << "\n"
         << "The error message is: " << ex.errmsg() << endl;
  }
  catch(...) {
    cerr << "Caught a system exception." << endl;
  }

  orb->NP_destory();
```

```
  return 0;
}
```

2.11.6 eg3_impl.cc

```
// eg3_impl.cc - This is the source code of example 3 used in Chap-
ter 2
//               "The Basics" of the omniORB2 user guide.
//
//               This is the object implementation.
//
// Usage: eg3_impl
//
//         On startup, the object reference is registered with the
//         COS naming service. The client uses the naming service to
//         locate this object.
//
//         The name which the object is bound to is as follows:
//               root  [context]
//                |
//               text  [context] kind [my_context]
//                |
//               Echo  [object]  kind [Object]
//

#include <iostream.h>
#include "omnithread.h"
#include "echo.hh"

#include "echo_i.cc"

static CORBA::Boolean bindObjectToName(CORBA::ORB_ptr,CORBA::Object_ptr);

int
main(int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA");

  Echo_i *myobj = new Echo_i();
  myobj->_obj_is_ready(boa);

  {
    Echo_var myobjRef = myobj->_this();
    if (!bindObjectToName(orb,myobjRef)) {
      return 1;
    }
  }

  boa->impl_is_ready();
  // Tell the BOA we are ready. The BOA's default be-
haviour is to block
  // on this call indefinitely.

  // Call boa->impl_shutdown() from another thread would unblock the
  // main thread from impl_is_ready().
```

```
  //
  // To properly shutdown the BOA and the ORB, add the follow-
ing calls
  // after impl_is_ready() returns.
  //
  // boa->destroy();
  // orb->NP_destory();

  return 0;
}


static
CORBA::Boolean
bindObjectToName(CORBA::ORB_ptr orb,CORBA::Object_ptr obj)
{
  CosNaming::NamingContext_var rootContext;

  try {
    // Obtain a reference to the root context of the Name service:
    CORBA::Object_var initServ;
    initServ = orb->resolve_initial_references("NameService");

    // Narrow the object returned by resolve_initial_references()
    // to a CosNaming::NamingContext object:
    rootContext = CosNaming::NamingContext::_narrow(initServ);
    if (CORBA::is_nil(rootContext))
      {
        cerr << "Failed to narrow naming context." << endl;
        return 0;
      }
  }
  catch(CORBA::ORB::InvalidName& ex) {
    cerr << "Service required is invalid [does not exist]." << endl;
    return 0;
  }


  try {
    // Bind a context called "test" to the root context:

    CosNaming::Name contextName;
    contextName.length(1);
    contextName[0].id   = (const char*) "test";     // string copied
    contextName[0].kind = (const char*) "my_context"; // string copied
    // Note on kind: The kind field is used to indicate the type
    // of the object. This is to avoid conventions such as that used
    // by files (name.type -- e.g. test.ps = postscript etc.)

    CosNaming::NamingContext_var testContext;
    try {
      // Bind the context to root, and assign testContext to it:
      testContext = rootContext->bind_new_context(contextName);
```

```
    }
    catch(CosNaming::NamingContext::AlreadyBound& ex) {
      // If the context already exists, this excep-
tion will be raised.
      // In this case, just resolve the name and assign testContext
      // to the object returned:
      CORBA::Object_var tmpobj;
      tmpobj = rootContext->resolve(contextName);
      testContext = CosNaming::NamingContext::_narrow(tmpobj);
      if (CORBA::is_nil(testContext)) {
        cerr << "Failed to narrow naming context." << endl;
        return 0;
      }
    }

    // Bind the object (obj) to testContext, naming it Echo:
    CosNaming::Name objectName;
    objectName.length(1);
    objectName[0].id   = (const char*) "Echo";   // string copied
    objectName[0].kind = (const char*) "Object"; // string copied


    // Bind obj with name Echo to the testContext:
    try {
      testContext->bind(objectName,obj);
    }
    catch(CosNaming::NamingContext::AlreadyBound& ex) {
      testContext->rebind(objectName,obj);
    }
    // Note: Using rebind() will overwrite any Object previ-
ously bound
    //       to /test/Echo with obj.
    //       Alternatively, bind() can be used, which will raise a
    //       CosNaming::NamingContext::AlreadyBound excep-
tion if the name
    //       supplied is already bound to an object.

    // Amendment: When using OrbixNames, it is neces-
sary to first try bind
    // and then rebind, as rebind on it's own will throw a Not-
Foundexception if
    // the Name has not already been bound. [This is incorrect be-
haviour -
    // it should just bind].
  }
  catch (CORBA::COMM_FAILURE& ex) {
    cerr << "Caught system exception COMM_FAILURE, unable to con-
tact the "
         << "naming service." << endl;
    return 0;
  }
  catch (omniORB::fatalException& ex) {
    throw;
```

```
  }
  catch (...) {
    cerr << "Caught a system exception while using the naming ser-
vice."<< endl;
    return 0;
  }
  return 1;
}
```

2.11.7 eg3_clt.cc

```
// eg3_clt.cc - This is the source code of example 3 used in Chap-
ter 2
//              "The Basics" of the omniORB2 user guide.
//
//              This is the client. It uses the COSS naming service
//              to obtain the object reference.
//
// Usage: eg3_clt
//
//
//        On startup, the client lookup the object refer-
ence from the
//        COS naming service.
//
//        The name which the object is bound to is as follows:
//              root  [context]
//                |
//              text  [context] kind [my_context]
//                |
//              Echo  [object]  kind [Object]
//

#include <iostream.h>
#include "echo.hh"

#include "greeting.cc"

extern void hello(CORBA::Object_ptr obj);

static CORBA::Object_ptr getObjectReference(CORBA::ORB_ptr orb);

int
main (int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA");

  try {
    CORBA::Object_var obj = getObjectReference(orb);
    hello(obj);
  }
  catch(CORBA::COMM_FAILURE& ex) {
    cerr << "Caught system exception COMM_FAILURE, unable to con-
tact the "
         << "object." << endl;
  }
  catch(omniORB::fatalException& ex) {
    cerr << "Caught omniORB2 fatalException. This indi-
cates a bug is caught "
         << "within omniORB2.\nPlease send a bug report.\n"
         << "The exception was thrown in file: " << ex.file() << "\n"
```

```
          << "                                    line: " << ex.line() << "\n"
          << "The error message is: " << ex.errmsg() << endl;
  }
  catch(...) {
    cerr << "Caught a system exception." << endl;
  }

  orb->NP_destory();

  return 0;
}

static
CORBA::Object_ptr
getObjectReference(CORBA::ORB_ptr orb)
{
  CosNaming::NamingContext_var rootContext;

  try {
    // Obtain a reference to the root context of the Name service:
    CORBA::Object_var initServ;
    initServ = orb->resolve_initial_references("NameService");

    // Narrow the object returned by resolve_initial_references()
    // to a CosNaming::NamingContext object:
    rootContext = CosNaming::NamingContext::_narrow(initServ);
    if (CORBA::is_nil(rootContext))
      {
        cerr << "Failed to narrow naming context." << endl;
        return CORBA::Object::_nil();
      }
  }
  catch(CORBA::ORB::InvalidName& ex) {
    cerr << "Service required is invalid [does not exist]." << endl;
    return CORBA::Object::_nil();
  }


  // Create a name object, containing the name test/context:
  CosNaming::Name name;
  name.length(2);

  name[0].id   = (const char*) "test";       // string copied
  name[0].kind = (const char*) "my_context"; // string copied
  name[1].id   = (const char*) "Echo";
  name[1].kind = (const char*) "Object";
  // Note on kind: The kind field is used to indicate the type
  // of the object. This is to avoid conventions such as that used
  // by files (name.type -- e.g. test.ps = postscript etc.)


  CORBA::Object_ptr obj;
  try {
```

```
    // Resolve the name to an object reference, and as-
sign the reference
    // returned to a CORBA::Object:
    obj = rootContext->resolve(name);
  }
  catch(CosNaming::NamingContext::NotFound& ex)
    {
      // This exception is thrown if any of the components of the
      // path [contexts or the object] aren't found:
      cerr << "Context not found." << endl;
      return CORBA::Object::_nil();
    }
  catch (CORBA::COMM_FAILURE& ex) {
    cerr << "Caught system exception COMM_FAILURE, unable to con-
tact the "
         << "naming service." << endl;
    return CORBA::Object::_nil();
  }
  catch(omniORB::fatalException& ex) {
    throw;
  }
  catch (...) {
    cerr << "Caught a system exception while using the naming ser-
vice."<< endl;
    return CORBA::Object::_nil();
  }
  return obj;
}
```

2.11.8   eg3_tieimpl.cc

```
// eg3_tieimpl.cc - This example is similar to eg3_impl.cc ex-
cept that
//                  the tie implementation skeleton is used.
//
//               This is the object implementation.
//
// Usage: eg3_tieimpl
//
//        On startup, the object reference is registered with the
//        COS naming service. The client uses the naming service to
//        locate this object.
//
//        The name which the object is bound to is as follows:
//              root  [context]
//                 |
//              text  [context] kind [my_context]
//                 |
//              Echo  [object]  kind [Object]
//

#include <iostream.h>
#include "omnithread.h"
#include "echo.hh"


// Real implementation, notice that it does not in-
herit from any stub class
class Echo_i {
public:
  Echo_i() {}
  virtual ~Echo_i() {}
  virtual char * echoString(const char *mesg);
};

char *
Echo_i::echoString(const char *mesg) {
  char *p = CORBA::string_dup(mesg);
  return p;
}


static CORBA::Boolean bindObjectToName(CORBA::ORB_ptr,CORBA::Object_ptr);

int
main(int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA");

  Echo_i *myimpl = new Echo_i();
  _tie_Echo<Echo_i,1> *myobj = new _tie_Echo<Echo_i,1>(myimpl);
```

```
  myobj->_obj_is_ready(boa);

  {
    Echo_var myobjRef = myobj->_this();
    if (!bindObjectToName(orb,myobjRef)) {
      return 1;
    }
  }

  boa->impl_is_ready();
  // Tell the BOA we are ready. The BOA's default be-
haviour is to block
  // on this call indefinitely.

  return 0;
}


static
CORBA::Boolean
bindObjectToName(CORBA::ORB_ptr orb,CORBA::Object_ptr obj)
{
  CosNaming::NamingContext_var rootContext;

  try {
    // Obtain a reference to the root context of the Name service:
    CORBA::Object_var initServ;
    initServ = orb->resolve_initial_references("NameService");

    // Narrow the object returned by resolve_initial_references()
    // to a CosNaming::NamingContext object:
    rootContext = CosNaming::NamingContext::_narrow(initServ);
    if (CORBA::is_nil(rootContext))
      {
        cerr << "Failed to narrow naming context." << endl;
        return 0;
      }
  }
  catch(CORBA::ORB::InvalidName& ex) {
    cerr << "Service required is invalid [does not exist]." << endl;
    return 0;
  }


  try {
    // Bind a context called "test" to the root context:

    CosNaming::Name contextName;
    contextName.length(1);
    contextName[0].id   = (const char*) "test";     // string copied
    contextName[0].kind = (const char*) "my_context"; // string copied
    // Note on kind: The kind field is used to indicate the type
    // of the object. This is to avoid conventions such as that used
```

```
     // by files (name.type -- e.g. test.ps = postscript etc.)

     CosNaming::NamingContext_var testContext;
     try {
       // Bind the context to root, and assign testContext to it:
       testContext = rootContext->bind_new_context(contextName);
     }
     catch(CosNaming::NamingContext::AlreadyBound& ex) {
       // If the context already exists, this excep-
tion will be raised.
       // In this case, just resolve the name and assign testContext
       // to the object returned:
       CORBA::Object_var tmpobj;
       tmpobj = rootContext->resolve(contextName);
       testContext = CosNaming::NamingContext::_narrow(tmpobj);
       if (CORBA::is_nil(testContext)) {
         cerr << "Failed to narrow naming context." << endl;
         return 0;
       }
     }

     // Bind the object (obj) to testContext, naming it Echo:
     CosNaming::Name objectName;
     objectName.length(1);
     objectName[0].id   = (const char*) "Echo";   // string copied
     objectName[0].kind = (const char*) "Object"; // string copied


     // Bind obj with name Echo to the testContext:
     try {
       testContext->bind(objectName,obj);
     }
     catch(CosNaming::NamingContext::AlreadyBound& ex) {
       testContext->rebind(objectName,obj);
     }
     // Note: Using rebind() will overwrite any Object previ-
ously bound
     //       to /test/Echo with obj.
     //       Alternatively, bind() can be used, which will raise a
     //       CosNaming::NamingContext::AlreadyBound excep-
tion if the name
     //       supplied is already bound to an object.

     // Amendment: When using OrbixNames, it is neces-
sary to first try bind
     // and then rebind, as rebind on it's own will throw a Not-
Foundexception if
     // the Name has not already been bound. [This is incorrect be-
haviour -
     // it should just bind].
   }
   catch (CORBA::COMM_FAILURE& ex) {
     cerr << "Caught system exception COMM_FAILURE, unable to con-
```

```
tact the "
        << "naming service." << endl;
    return 0;
  }
  catch (omniORB::fatalException& ex) {
    throw;
  }
  catch (...) {
    cerr << "Caught a system exception while using the naming ser-
vice."<< endl;
    return 0;
  }
  return 1;
}
```

2.11.9 dir.mk

```
# dir.mk - This is the makefile to compile all the example pro-
grams used in
#          Chapter 2 The Basics.
#          This makefile is to be used under the OMNI develop-
ment environment.
#
CXXSRCS = greeting.cc eg1.cc \
          eg2_impl.cc eg2_clt.cc \
          eg3_impl.cc eg3_clt.cc

DIR_CPPFLAGS = $(CORBA_CPPFLAGS)

CORBA_INTERFACES = echo

eg1        = $(patsubst %,$(BinPattern),eg1)
eg2_impl   = $(patsubst %,$(BinPattern),eg2_impl)
eg2_clt    = $(patsubst %,$(BinPattern),eg2_clt)
eg3_impl   = $(patsubst %,$(BinPattern),eg3_impl)
eg3_clt    = $(patsubst %,$(BinPattern),eg3_clt)

all:: $(eg1) $(eg2_impl) $(eg2_clt)  $(eg3_impl) $(eg3_clt)

clean::
        $(RM) $(eg1) $(eg2_impl) $(eg2_clt) $(eg3_impl) $(eg3_clt)

$(eg1): eg1.o $(CORBA_STUB_OBJS) $(CORBA_LIB_DEPEND)
        @(libs="$(CORBA_LIB)"; $(CXXExecutable))

$(eg2_impl): eg2_impl.o $(CORBA_STUB_OBJS) $(CORBA_LIB_DEPEND)
        @(libs="$(CORBA_LIB)"; $(CXXExecutable))

$(eg2_clt): eg2_clt.o $(CORBA_STUB_OBJS) $(CORBA_LIB_DEPEND)
        @(libs="$(CORBA_LIB)"; $(CXXExecutable))

$(eg3_impl): eg3_impl.o $(CORBA_STUB_OBJS) $(CORBA_LIB_DEPEND)
        @(libs="$(CORBA_LIB)"; $(CXXExecutable))

$(eg3_clt): eg3_clt.o $(CORBA_STUB_OBJS) $(CORBA_LIB_DEPEND)
        @(libs="$(CORBA_LIB)"; $(CXXExecutable))
```

Chapter 3

IDL to C++ Language Mapping

Now that you are familiar with the basics, it is important to familiar yourselves with the IDL to C++ language. The mapping is described in detail in [OMG99a]. If you have not done so, you should obtain a copy of the document and use that as the programming guide to omniORB2.

The specification is not an easy read. The alternative is to use one of the books on CORBA programming that has begun to appear. For instance, the "Advanced CORBA Programming with C++" by Michi Henning and Steve Vinoski includes many example code bits to illustrate how to use the CORBA 2.3 C++ mapping.

3.1 Incompatibilities with pre-2.8.0 releases

Before 2.8.0, omniORB2 implements the CORBA 2.0 C++ mapping. Since 2.8.0, the mapping has been updated to CORBA 2.3. Unfortunately, to comply with the CORBA 2.3 specification, it is necessary to change the semantics of a few APIs in a way that is incompatible with older omniORB2 releases. The incompatible changes are limited to:

1. the extraction of string, object reference and typecode from an Any.

2. the DII interface now defaults to report a system exception by raising a C++ exception instead of returning the exception as an environment value.

The changes are minor and requires minimal changes to the application source code. **However, it is not possible to detect the old usage at compile time. In particular, unmodified code that use the affected Any extraction operators would most certainly cause runtime errors to occur.**

To smooth the transition from the old usage to the new, an ORB configuration variable `omniORB::omniORB_27_CompatibleAnyExtraction` can be set to revert the any extraction operators to the old semantics. More information can be found in chapter 9.

Chapter 4

The omniORB2 API

In this chapter, we introduce the omniORB2 API. The purpose of this API is to provide access points to omniORB2 specific functionalities that are not covered by the CORBA specification. Obviously, if you use this API in your application, that part of your code is not going to be portable to run unchanged on other vendors' ORBs. To make it easier to identify omniORB2 dependent code, this API is defined under the name space "omniORB"[1].

4.1   ORB and BOA initialisation options

`CORBA::ORB_init` accepts the following command-line arguments:

`-ORBid ''omniORB2''` The identifier supplied must be "omniORB2".

`-ORBtraceLevel <level>` See section 4.2.

`-ORBserverName <string>` See section 4.3.

`-ORBtcAliasExpand <0 or 1>` See section 9.2.

`-ORBgiopMaxMsgSize <size in bytes>` See section 4.5.

`-ORBInitialHost <string>` See section 4.6.

`-ORBInitialPort <1-65535>` See section 4.6.

`-ORBdiiThrowsSysExceptions <0 or 1>` See section 11.4.

`-ORBinConScanPeriod <0-max integer>` See section 7.3.

`-ORBoutConScanPeriod <0-max integer>` See section 7.3.

`-ORBclientCallTimeOutPeriod <0-max integer>` See section 7.3.

`-ORBserverCallTimeOutPeriod <0-max integer>` See section 7.3.

`-ORBscanGranularity <0-max integer>` See section 7.3.

---

[1]omniORB is a class name if the C++ compiler does not support the namespace keyword.

`-ORBverifyObjectExistsAndType <0 or 1>` See section 4.7.

`-ORBlcdMode` See secion 4.7

`-ORBabortOnInternalError <0 or 1>` See section 4.8.

`-ORBhelp` List all ORB command line options.

   `BOA_init` accepts the following command-line arguments:

`-BOAid ''omniORB2_BOA''` The identifier supplied must be "omniORB2_BOA".

`-BOAiiop_port <port number>` This option tells the BOA which TCP/IP port to use to accept IIOP calls. If this option is not specified, the BOA will use an arbitrary port assigned by the operating system.

`-BOAno_bootstrap_agent` See section 4.6.

`-BOAiiop_name_port <hostname[:port number]>` This options tells the BOA the hostname and optionally the port number to use. See below for details.

`-BOAhelp` List all BOA command line options.

   By default, the BOA can work out the IP address of the host machine. This address is recorded in the object references of the local objects. However, when the host has multiple network interfaces and multiple IP addresses, it may be desirable for the application to control what address the BOA should use. This can be done by defining the environment variable `OMNIORB_USEHOSTNAME` to contain the preferred host name or IP address in dot-numeric form. Alternatively, the same can be acheived using the `-BOAiiop_name_port` option.

   As defined in the CORBA specification, any command-line arguments understood by the ORB/BOA will be removed from `argv` when the initialisation functions return. Therefore, an application is not required to handle any command-line arguments it does not understand.

4.2   Run-time Tracing and Diagnostic Messages

OmniORB2 uses the C++ iostream `cerr` to output any tracing and diagnostic messages. Some or all of these messages can be turned-on/off by setting the variable `omniORB::traceLevel`. The type definition of the variable is:

```
CORBA::ULong omniORB::traceLevel = 1;  // The default value is 1
```

   At the moment, the following trace levels are defined:

level 0 turn off all tracing and informational messages

level 1 informational messages only

level 2 the above plus configuration information

level 5 the above plus notifications when server threads are created or communication endpoints are shutdown

level 10-20 the above plus execution traces

level 25 the above plus hex dump of all data sent and received by the ORB via its network connections.

The variable can be changed by assignment inside your applications. It can also be changed by specifying the command-line option: `-ORBtraceLevel <level>`. For instance:

```
$ eg2_impl -ORBtraceLevel 5
```

4.3 Server Name

Applications can optionally specified a name to identify the server process. At the moment, this name is only used by the host-based access control module. See section 7.5 for details.

The name is stored in the variable `omniORB::serverName`.

```
CORBA::String_var omniORB::serverName;
```

The variable can be changed by assignment inside your applications. It can also be changed by specifying the command-line option: `-ORBserverName <string>`.

4.4 Object Keys

OmniORB2 uses a data type `omniORB::objectKey` to uniquely identify each object implementation. This is an opaque data type and can only be manipulated by the following functions:

```
void omniORB::generateNewKey(omniORB::objectKey &k);
```

`omniORB::generateNewKey` returns a new `objectKey`. The return value is guaranteed to be unique among the keys generated during this program run. On the platforms that have a realtime clock and unique process identifiers, a stronger assertion can be made, i.e. the keys are guaranteed to be unique among all keys ever generated on the same machine.

```
const unsigned int omniORB::hash_table_size;
int omniORB::hash(omniORB::objectKey& k);
```

`omniORB::hash` returns the hash value of an `objectKey`. The value returned by this function is always between 0 and `omniORB:hash_table_size - 1` inclusively.

```
omniORB::objectKey omniORB::nullkey();
```

omniORB::nullkey always returns the same objectKey value. This key is guaranteed to hash to 0.

```
int operator==(const omniORB::objectKey &k1,const om-
niORB::objectKey &k2);
int operator!=(const omniORB::objectKey &k1,const om-
niORB::objectKey &k2);
```

ObjectKeys can be tested for equality using the overloaded operator== and operator!=.

```
omniORB::seqOctets*
omniORB::keyToOctetSequence(const omniORB::objectKey &k1);

omniORB::objectKey
omniORB::octetSequenceToKey(const omniORB::seqOctets& seq);
```

omniORB::keyToOctetSequence takes an objectKey and returns its externalised representation in the form of a sequence of octets. The same sequence can be converted back to an objectKey using omniORB::octetSequenceToKey. If the supplied sequence is not an objectKey, omniORB::octetSequenceToKey raises a CORBA::MARSHAL exception.

4.5   GIOP Message Size

omniORB2 sets a limit on the GIOP message size that can be sent or received. The value can be obtained by calling:

```
size_t omniORB::MaxMessageSize();
```

and can be changed by:

```
void omniORB::MaxMessageSize(size_t newvalue);
```

or by the command-line option -ORBgiopMaxMsgSize.

The exact value is somewhat arbitrary. The reason such a limit exists is to provide some way to protect the server side from resource exhaustion. Think about the case when the server receives a rogue GIOP(IIOP) request message that contains a sequence length field set to 2**31. With a reasonable message size limit, the server can reject this rogue message straight away.

4.6   Initial Object Reference Bootstrapping

Starting from 2.6.0, a new mechanism is available for the ORB runtime to obtain the initial object references to CORBA services. Previously, it is necessary to write the IOR string of these services in the configuration file (section 1.2). Now the object references can be obtained from a bootstrap service. The bootstrap service is a special object with the object key 'INIT' and the following interface[2].:

---

[2]This interface is first defined by Sun's NEO and is in used in Sun's JavaIDL

```
// IDL
module CORBA {
   interface InitialReferences {
        Object get(in ORB::ObjectId id);
        // returns the initial object reference of the service
        // identified by <id>. For example the id for the COSS
        // Naming service is "NameService".

        ORB::ObjectIdList list();
        // returns the list of service id that this agent knows of
        // their initial object reference.
   };
};
```

By default, all omniORB2 servers, i.e. those applications with the BOA initialised, contains an instance of this object and is able to responds to remote invocations. To prevent the ORB from instantiating this object, the command-line option `-BOAno_bootstrap_agent` should be specified.

In particular, the Naming Service omniNames is able to respond to a query through this interface and returns the object reference of its root context. In effect, the bootstrap agent provides a level of indirection. All omniORB2 clients still have to be supplied with the address of the bootstrap agent. However the information is much easier to specify than a stringified IOR! Another advantage of this approach is that it is completely compatiable with JavaIDL. This makes it possible for a client written in JavaIDL to share with a omniORB2 server the same Naming Service.

The address of the bootstrap agent is given by the `ORBInitialHost` and `ORBInitialPort` parameter in the omniORB configuration file (section 1.2). The parameters can also be specified as command-line options (section 4.1). `ORBInitialHost` is the host name and `ORBInitialPort` is the TCP/IP port number. The parameter `ORBInitialPort` is optional. If it is not specified, port number 900 will be used.

During initialisation, the ORB reads the parameters in the omniORB configuration file. If the parameter `NAMESERVICE` is specified, the stringified IOR would be used as the object reference of the root naming context. If the parameter is absent and the parameter `ORBInitialHost` is present, the ORB would contact the bootstrap agent at the address specified to obtain the root naming context when the application calls `resolve_initial_reference()`. If neither the `NAMESERVICE` nor `ORBInitialHost` is present, a call to `resolve_initial_reference()` returns a nil object. Finally, the command line argument `-ORBInitialHost` overrides any parameters in the configuration file. The ORB would always contact the bootstrap agent at the address specified to obtain the root naming context.

Now we are ready to describe a simple way to set up omniNames.

1. Start omniNames for the first time on a machine (e.g. wobble):

   ```
   $ omniNames -start 1234
   ```

2. Add to omniORB.cfg:

   ```
   ORBInitialHost wobble
   ORBInitialPort 1234
   ```

3. All omniORB2 applications will now be able to contact omniNames.

Alternatively, the command line options can be used, for example:
```
$ eg3_impl -ORBInitialHost wobble -ORBInitialPort 1234 &
$ eg3_clt -ORBInitialHost wobble -ORBInitialPort 1234
```

4.7   GIOP Lowest Common Denominator Mode

Sometimes, to cope with bugs in another ORB, it is necessary to disable various GIOP and IIOP features in order to achieve interoperability. If the command line option `-ORBlcdMode` or the function `omniORB::enableLcdMode()` is called, the ORB enters the so-called "lowest common denominator mode". It bends over backwards to cope with bugs in the ORB at the other end. This is purely a transitional measure. The long term solution is to report the bugs to the other vendors and ask them to fix them expediently.

By default, omniORB2 uses GIOP LOCATE_REQUEST message to verify the existence of an object prior to the first invocation. If another vendor's ORB is known not to be able to handle this GIOP message, set the variable `omniORB::verifyObjectExistsAndType` to 0 would disable this feature, and hence achieve interoperability. The option can also be set via the command line option `-ORBverifyObjectExistsAndType`.

4.8   Trapping omniORB2 Internal Errors

```
class fatalException {
public:
    const char *file() const;
    int line() const;
    const char *errmsg() const;
};
```

When omniORB2 detects an internal inconsistency that is most likely to be caused by a bug in the runtime, it raises the exception `omniORB::fatalException`. When this exception is raised, it is not sensible to proceed with any operation that involves the ORB's runtime. It is best to exit the program immediately. The exception structure carries by `omniORB::fatalException` contains the exact location (the file name and the line number) where the exception is raised. You are strongly encourage to file a bug report and point out the location.

It may help to cause a core-dump and look at the stack trace to locate where the exception was thrown. This can be done by setting the variable `omniORB::abortOnInternalError` to 1. The variable can also be set via the command line option `-ORBabortOnInternalError`.

Chapter 5

The Basic Object Adaptor (BOA)

This chapter describes the BOA implementation in omniORB2. The CORBA speci-
fication defines the Basic Object Adaptor as the entity that mediates between object
implementations and the ORB. Unfortunately, the BOA specification is incomplete
and does not address the multi-threading issues appropriately. The end result is that
different ORB vendors implement different extensions to their BOAs. Worse, the im-
plementation of the operations defined in the specification are different in different
ORBs. Recently, a new Object Adaptor specification (the Portable Object Adaptor-
POA) has been adopted and will replace the BOA as the standard Object Adaptor in
CORBA. The new specification recognises the compatibility problems of BOA and rec-
ommends that all BOAs should be considered propriety extensions. OmniORB2 will
support POA in future releases. Until then, you have to use the BOA to attach object
implementations to the ORB.

The rest of this chapter describes the interface of the BOA in detail. It is important
to recognise that the interface described below is omniORB2 specific and hence the
code using this interface is unlikely to be portable to other ORBs.

Unless it is stated otherwise, the term "object" will be used below to refer to object
implementations. This should not be confused with "object references" which are
handles held by clients.

5.1   BOA Initialisation

It takes two steps to put the BOA into service. The BOA has to be initialised using
`BOA_init` and activated using `impl_is_ready`.
`BOA_init` is a member of the `CORBA::ORB` class. Its signature is:

```
BOA_ptr BOA_init(int & argc,
                 char ** argv,
                 const char * boa_identifier);
```

Typically, it is used in the startup code as follows:

```
CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");   // line 1
CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA"); // line 2
```

47

The `argv` parameters may contain BOA options. These options will be removed from the `argv` list when `BOA_init` returns. Other parameters in `argv` will remain. The supported options are:

-BOAiiop_port `<port number (0-65535)>` Use the port number to receive IIOP requests. This option can be specified multiple times in the command line and the BOA would be initialised to listen on all of the ports.

-BOAid `<id (string)>` If this option is used the id must be "omniORB2_BOA".

-BOAiiop_name_port `<hostname[:port number]>` Similar to `-BOAiiop_port`, this options tells the BOA the hostname and optionally the port number to use.

If the third argument of `BOA_init` is non-nil, it must be the string constant "omniORB2_BOA". If the argument is nil, -BOAid must be present in `argv`.

If there is any problem in the initialisation process, a `CORBA::INITIALIZE` exception would be raised.

To register an object with the BOA, the method `_obj_is_ready` should be called with the return value of `BOA_init` as the argument.

`BOA_init` is thread-safe. It can be called multiple times and the same `BOA_ptr` will be returned. However, only the `argv` in the first call will be scanned, the argument is ignored in subsequent calls.

`BOA_init` returns a pseudo object of type `CORBA::BOA_ptr`.   Similar to `CORBA::Object_ptr`, the pointer can be managed using `CORBA::BOA_var`, `BOA::_duplicate` and `CORBA::release`.   The pointer can be tested using `CORBA::is_nil` which returns true if the pointer is equivalent to the return value of `BOA::_nil`.

After `BOA_init` is called, objects can be registered. However, incoming IIOP requests would not be despatched until `impl_is_ready` is called.

```
class BOA {
public:
    impl_is_ready(CORBA::ImplementationDef_ptr p = 0,
                  CORBA::Boolean NonBlocking = 0);
};
```

One of the common pitfall in using the BOA is to forget to call impl_is_ready. Until this call returns, there is no thread listening on the port from which IIOP requests are received. The remote client may hang because of this.

When `impl_is_ready` is called with no argument. The calling thread would be blocked indefinitely in the function until `impl_shutdown` (see below) is called. The thread that is calling `impl_is_ready` is not used by the BOA to perform its internal functions. The BOA has its own set of threads to process incoming requests and general housekeeping. Therefore, it is not necessary to have a thread blocked in the call if it can be put into use elsewhere. For example, the main thread may call impl_is_ready once in non-blocking mode (see below) and then enter the event loop to handle the GUI frontend.

If non-blocking behaviour is needed, the `NonBlocking` argument should be set to 1. For instance, if you creates a callback object, you might call impl_is_ready in

non-blocking mode to tell the BOA to start receiving IIOP requests before sending the callback object to the remote object. The first argument `ImplementationDef_ptr` is ignored by the BOA. Just set the argument to nil.

`impl_is_ready` is thread safe and can be called multiple times. Multiple threads can be blocked in `impl_is_ready`.

5.2 Object Registration

Once the BOA is initialised, objects can be registered. The purpose of object registration is to let the BOA know of the existence of the object and to dispatch requests for the object as upcalls into the object.

To register an object, the `_obj_is_ready` function should be called. `_obj_is_ready` is a member function of the implementation skeleton class. The function should be called only once for each object. The call should be made only after the object is fully initialised.

The member function `obj_is_ready` of the BOA may also be used to register an object. However, this function has been superseded by `_obj_is_ready` and should not be used in new application code.

5.3 Object Disposal

Once an object is registered, it is under the management of the BOA. To remove the object from the BOA and to delete it (when it is safe to do so), the `_dispose` function should be called. `_dispose` is a member function of the implementation skeleton class. The function should be called only once for each object.

Notice the asymmetry in object instantiation and destruction. To instantiate an object, the application code has to call the **new** operator. To remove the object, the application should never call the delete operator on the object directly.

At the time the `_dispose` call is made, there may be other threads invoking on the object, the BOA ensures that all these calls are completed before removing the object from its internal tables and calling the **delete** operator.

Internally, the BOA keeps a reference count on each object. Initially, the reference count is 0. After a call to `_obj_is_ready`, the reference count is 1. The BOA increases the reference count by 1 before an upcall into the object is made. The count is decreased by 1 when the upcall returns. `_dispose` decreases the reference count by 1, if the reference count is 0, the delete operator is called. If the count is non-zero, the object is marked as disposed. The object will be deleted when the reference count eventually goes to zero.

The reference count is also increased by 1 for each object reference held in the same address space. Hence, the **delete** operator will not be called when there are outstanding object references in the same address space. To ensure that an object is deleted, all its object references in the same address space should be released using `CORBA::release`.

Unlike colocated object references, references held by clients in other address spaces would not prevent the deletion of objects. If these clients invoke on the object

after it is disposed, the system exception INV_OBJREF would be raised. The difference in semantics is an undesirable side-effect of the current BOA implementation. In future, colocated references will have the same semantics as remote references, i.e. their presence will not delay the deletion of the objects.

Instead of _dispose, it may be useful to have a method to deactivate the object but not deleting it. This feature is not supported in the current BOA implementation.

5.4   BOA Shutdown

The BOA can be withdrawn from service using member functions impl_shutdown and destroy.

```
class BOA {
public:
   void impl_shutdown();
   void destroy();
};
```

impl_shutdown and destroy are the inverse of impl_is_ready and BOA_init respectively.

impl_shutdown deactivates the BOA. When the call returns, all the internal threads and network connections will be shutdown. Any thread blocking in impl_is_ready would be unblocked. After the call, no request from other address spaces will be processed. In other words, the BOA will be in the same state as it was in before impl_is_ready was called. For example, a remote client may hang if it tries to connect to the server after impl_shutdown was called because no thread is listening on the IIOP port.

impl_shutdown does not wait for incoming requests to complete before it closes the network connections. The remote clients will see the network connections shutdown and the replies may not reach them even if the upcalls have been completed. Therefore, if the application is to define an operation in an IDL interface to shutdown the BOA, the operation should be defined as an oneway operation.

impl_shutdown is thread-safe and can be called multiple times. The call is silently ignored if the BOA has already been shutdown. After impl_shutdown is called, the BOA can be reactivated by another call to impl_is_ready.

It should be noted that impl_shutdown does not affect outgoing network connections. That is, clients in the same address space will still be able to make calls to objects in other address spaces.

While remote requests are not delivered after impl_shutdown is called, the current implementation does not stop colocated clients from calling the objects. In future, colocated clients will exhibit the same behaviour as remote clients.

destroy permanently removed the BOA. This function will call impl_shutdown implicitly if it has not been called. When this call returns, the IIOP port(s) held by the BOA will be freed. Remote clients will see their requests refused by the operating system when they try to open a connection to the IIOP port(s).

After destroy is called, the BOA should not be used. If there is any objects still registered with the BOA, the objects should not be invoked afterwards. The objects

are not disposed. Invoking on the objects after `destroy` would result in undefined behaviour. Initialisation of another BOA using `BOA_init` is not supported. The behaviour of `BOA_init` after this call is undefined.

5.5 Unsupported functions

The following member functions are not implemented. Calling these functions do not have any effect.

- `Object_ptr create(...)`

- `ReferenceData* get_id(Object_ptr)`

- `Principal_ptr get_principal(Object_ptr,Environment_ptr)`

- `void change_implementation(Object_ptr, Implementation-Def_ptr)`

- `void deactivate_impl(ImplementationDef_ptr)`

- `void deactivate_obj(Object_ptr)`

5.6 Loading Objects On Demand

Since 2.5.0, there is limited support for loading objects on demand. An application can register a handler for loading objects dynamically. The handler should have the signature `omniORB::loader::mapKeyToObject_t`:

```
namespace omniORB {
  ...
  class loader {
  public:
    typedef CORBA::Object_ptr (*mapKeyToObject_t) (const objec-
tKey& key);
    static void set(mapKeyToObject_t NewKeyToObject);
  };
};
```

When the ORB cannot locate the target object in this address space, it calls the handler with the object key of the target. The handler is expected to instantiate the object, either in this address space or in another address space, and returns the object reference to the newly instantiated object. The ORB will then reply with a LOCATION_FORWARD message to instruct the client to retry using the object reference returned by the handler. When the handler returns, the ORB assumes ownership of the returned value. It will call CORBA::release() on the returned value when it has finished with it.

The handler may be called concurrently by multi-threads. Hence it must be thread-safe.

If the handler cannot load the target object, it should return CORBA::Object::_nil(). The object will be treated as non-existing.

The application registers the handler with the ORB at runtime using omniORB::loader::set(). This function is not thread-safe. Calling this function again will replace the old handler with the new one.

Chapter 6

Interface Type Checking

This chapter describes the mechanism used by omniORB2 to ensure type safety when object references are exchanged across the network. This mechanism is handled completely within the ORB. There is no programming interface visible at the application level. However, for the sake of diagnosing the problem when there is a type violation, it is useful to understand the underlying mechanism in order to interpret the error conditions reported by the ORB.

6.1   Introduction

In GIOP/IIOP, an object reference is encoded as an Interoperable Object Reference (IOR) when it is sent across a network connection. The IOR contains a Repository ID (REPOID) and one or more communication profiles. The communication profiles describe where and how the object can be contacted. The REPOID is a string which uniquely identifies the IDL interface of the object.

Unless the **ID** pragma is specified in the IDL, the ORB generates the REPOID string in the so-called OMG IDL Format[1]. For instance, the REPOID for the `Echo` interface used in the examples of chapter 2 is `IDL:Echo:1.0`.

When interface inheritance is used in the IDL, the ORB always sends the REPOID of the most derived interface. For example:

```
// IDL
   interface A {
     ...
   };
   interface B : A {
     ...
   };
   interface C {
     void op(in A arg);
   };

// C++
   C_ptr server;
```

---

[1]For further details of the repository ID formats, see section 6.6 in the CORBA specification.

```
    B_ptr objB;
    A_ptr objA = objB;
    server->op(objA);          // Send B as A
```

In the example, the operation C::op accepts an object reference of type A. The real type of the reference passed to C::op is B, which inherits from A. In this case, the REPOID of B, and not that of A, is sent across the network.

The GIOP/IIOP specification allows an ORB to send a null string in the REPOID field of an IOR. It is up to the receiving end to work out the real type of the object. OmniORB2 never sends out null strings as REPOID. However, it may receive null REPOID from other ORBs. In that case, it will use the mechanism described below to ensure type safety.

6.2   Basic Interface Type Checking

The ORB is provided with the interface information by the stubs via the proxyObject-Factory class. For an interface A, the stub of A contains a A_proxyObjectFactory class. This class is derived from the proxyObjectFactory class. The proxyObjectFactory is an abstract class which contains 3 virtual functions.

```
class proxyObjectFactory {
public:

  virtual const char *irRepoId() const = 0;

  virtual _CORBA_Boolean  is_a(const char *base_repoId) const = 0;

  virtual CORBA::Object_ptr newProxyObject(Rope *r,
                                           CORBA::Octet *key,
                                           size_t keysize,
                                           IOP::TaggedProfileList
                                           *profiles,
                                           CORBA::Boolean re-
lease) = 0;

};
```

- `irRepoId` returns the REPOID of the interface.

- `is_a` returns true(1) if the argument is the REPOID of the interface itself or it is that of its base interfaces.

- `newProxyObject` returns an object reference based on the information supplied in the arguments.

A single instance of every *_proxyObjectFactory is instantiated at runtime. The instances are entered into a list inside the ORB. The list constitutes all the interface information known to the ORB.

When the ORB receives an IOR from the network, it unmarshals and extracts the REPOID from the IOR. At this point, the ORB has two pieces of information in hand:

1. The REPOID of the object reference received from the network.

2. The REPOID the ORB is expecting. This comes from the unmarshal function that tells the ORB to receive the object reference.

Using the REPOID received, the ORB searches its proxyObjectFactory list for an exact match. If there is an exact match, all is well because the runtime can use the `is_a` method of the proxyFactory to check if the expected REPOID is the same as the received REPOID or if it is that of its base interfaces. If the answer is positive, the IOR passes the type checking test and the ORB can proceed to create an object reference in its own address space to represent the IOR.

However, the ORB may fail to find a match in its proxyObjectFactory list. This means that the ORB has no local knowledge of the REPOID. There are three possible causes:

1. The remote end is another ORB and it sends a null string as the REPOID.

2. The ORB is expecting an object reference of interface A. The remote end sends the REPOID of B which is an interface that inherits from A. The stubs of A is linked into the executable but the stubs of B is not.

3. The remote end has sent a duff IOR.

To handle this situation, the ORB must find out the type information dynamically. This is explained in the next section.

6.3 Interface Inheritance

When the ORB receives an IOR of interface type B when it expects the type to be A, it must find out if B inherits from A. When the ORB has no local knowledge of the type B, it must work out the type of B dynamically.

The CORBA specification defines an Interface Repository (IR) from which IDL interfaces can be queried dynamically. In the above situation, the ORB could contact the IR to find out the type of B. However, this approach assumes that an IR is always available and contains the up-to-date information of all the interfaces used in the domain. This assumption may not be valid in many applications.

An alternative is to use the `_is_a` operation to work out the actual type of an object. This approach is simpler and more robust than the previous one because no 3rd party is involved.

```
class Object{
    CORBA::Boolean _is_a(const char* type_id);
};
```

The ⌐is⌐a operation is part of the CORBA::Object interface and must be implemented by every object. The input argument is a REPOID. The function returns true(1) if the object is really an instance of that type, including if that type is a base type of the most derived type of that object.

In the situation above, the ORB would invoke the ⌐is⌐a operation on the object and ask if the object is of type A **before** it processes any application invocation on the object.

Notice that the ⌐is⌐a call is **not** performed when the IOR is unmarshalled. It is performed just prior to the first application invocation on the object. This leads to some interesting failure mode if B reports that it is not an A. Consider the following example:

```
\\ IDL
   interface A { ... };
   interface B : A { ... };
   interface D { ... };
   interface C {
     A      op1();
     Object op2();
   };

\\ C++

   C_ptr objC;
   A_ptr objA;
   CORBA::Object_ptr objR;

   objA =  objC->op1();              // line 1
   (void) objA->_non_existent();    // line 2

   objR =  objC->op2();              // line 3
   objA =  A::_narrow(objR);        // line 4
```

If the stubs of A,B,C,D are linked into the executable and:

Case 1 C::op1 and C::op2 returns a B. Line 1-4 complete successful. The remote object is only contacted at line 2.

Case 2 C::op1 and C::op2 returns a D. This condition only occurs if the runtime of the remote end is buggy. The ORB raises a CORBA::Marshal exception at line 1 because it knows it has received an interface of the wrong type.

If only the stub of A is linked into the executable and:

Case 1 C::op1 and C::op2 returns a B. Line 1-4 completes successful. When line 2 and 4 is executed, the object is contacted to ask if it is a A.

Case 2 C::op1 and C::op2 returns a D. This condition only occurs if the runtime of the remote end is buggy. Line 1 completes and no exception is raised. At line 2, the

object is contacted to ask if it is a A. If the answer is no, a CORBA::INV_OBJREF exception is raised. The application will also see a CORBA::INV_OBJREF at line 4.

Chapter 7

Connection Management

This chapter describes how omniORB2 manages network connections.

7.1 Background

In CORBA, the ORB is the "middleware" that allows a client to invoke an operation on an object without regard to its implementation or location. In order to invoke an operation on an object, a client needs to "bind" to the object by acquiring its object reference. Such a reference may be obtained as the result of an operation on another object (such as a naming service) or by conversion from a stringified representation previously generated by the same ORB. If the object is in a different address space, the binding process involves the ORB building a proxy object in the client's address space. The ORB arranges for invocations on the proxy object to be transparently mapped to equivalent invocations on the implementation object.

For the sake of interoperability, CORBA mandates that all ORBs should support IIOP as the means to communicate remote invocations over a TCP/IP connection. IIOP is asymmetric with respect to the roles of the parties at the two ends of a connection. At one end is the client which can only initiate remote invocations. At the other end is the server which can only receive remote invocations.

Notice that in CORBA, as in most distributed systems, remote bindings are established implicitly without application intervention. This provides the illusion that all objects are local, a property known as "location transparency". CORBA does not specify when such bindings should be established or how they should be multiplexed over the underlying network connections. Instead, ORBs are free to implement implicit binding by a variety of means.

The rest of this chapter describes how omniORB2 manages network connections and the programming interface to fine tune the management policy.

7.2 The Model

OmniORB2 is designed from the ground up to be fully multi-threaded. The objective is to maximise the degree of concurrency and at the same time eliminate any unnecessary thread overhead. Another objective is to minimise the interference by

the activities of other threads on the progress of a remote invocation. In other words, thread "cross-talk" should be minimised within the ORB. To achieve these objectives, the degree of multiplexing at every level is kept to a minimum.

On the client side of a connection, the thread that invokes on a proxy object drives the IIOP protocol directly and blocks on the connection to receive the reply. On the server side, a dedicated thread blocks on the connection. When it receives a request, it performs the up-call to the object and sends the reply when the upcall returns. There is no thread switching along the call chain.

With this design, there is at most one call in-flight at any time in a connection. If there is only one connection, concurrent invocations to the same remote address space would have to be serialised. To eliminate this limitation, omniORB2 implements a dynamic policy- multiple connections to the same remote address space are created on demand and cached when there are concurrent invocations in progress.

To be more precise, a network connection to another address space is only established when a remote invocation is about to be made. Therefore, there may be one or more object references in one address space that refers to objects in a different address space but unless the application invokes on these objects, no network connection is made. The maximum number of connections opened to another address space is 5 by default. Since 2.6.0, this parameter can be changed by setting the variable `om-niORB::maxTcpConnectionPerServer` **before** calling `ORB_init`.

It is wasteful to leave a connection opened when it has been left unused for a considerable time. Too many idle connections could block out new connections to a server when it runs out of spare communication channels. For example, most unix platforms has a limit on the number of file handles a process can open. 64 is the usual default limit. The value can be increased to a maximum of a thousand or more by changing the "ulimit" in the shell.

7.3   Idle Connection Shutdown and Remote Call Timeout

Inside the ORB, a thread is dedicated to scan for idle connections. The thread looks after both the outgoing connections and the incoming connections.

When a connection is idled for a period of time, the connection is shutdown. Similarly, if a remote call has not completed within a defined period of time, the connection is shutdown and the ORB will return `COMM_FAILURE` to the client.

How often the internal thread scan the connections is determined by the value of the *scan granularity*.   This value is default to 5 seconds and can be changed using the command-line option `-ORBscanGranularity` or using the `om-niORB::scanGranularity()` call. Notice that this value determines the precision the ORB is able to keep to the value of the idle connection or remote call timeout.

How long the ORB will wait before it shuts down an idle connection is determined by the idleConnectionPeriods. There are separate values for incoming and outgoing connections. The default values are 180 and 120 seconds for incoming and outgoing connections respectively. These values can be changed using the command-line options `-ORBinConScanPeriod` and `-ORBoutConScanPeriod`. They can also be controlled by the `omniORB::idleConnectionScanPeriod()` call.

Similary, how long the ORB will wait for a remote call to complete is determined by

the parameter clientCallTimeOutPeriod for the client side and the serverCallTimeOut-
Peroid for the server side. The default values are 60 and 90 seconds for the client and
the server side respectively. These values can be changed using the command-line op-
tion `-ORBclientCallTimeOutPeriod` and `-ORBserverCallTimeOutPeriod`.
They can also be controlled by the `omniORB::callTimeOutPeriod()` call.

```
class omniORB {

public:

  static void scanGranularity(CORBA::ULong sec);

  static CORBA::ULong scanGranularity();

  enum idleConnType { idleIncoming, idleOutgoing };

  static void idleConnectionScanPeriod(idleConnType direction,
                                       CORBA::ULong sec);

  static CORBA::ULong idleConnectionScanPeriod(idleConnType direction);

  enum callTimeOutType { clientSide, serverSide };

  static void callTimeOutPeriod(callTimeOutType direction,
                      CORBA::ULong sec);

  static CORBA::ULong callTimeOutPeriod(callTimeOutType direction);

};
```

The scan can be disabled completely by setting the scan granularity to 0.

7.4 Interoperability Considerations

The IIOP specification allows both the client and the server to shutdown a connec-
tion unilaterally. When one end is about to shutdown a connection, it should send a
closeConnection message to the other end. It should also make sure that the message
will reach the other end before it proceeds to shutdown the connection.

The client should distinguish between an orderly and an abnormal connection
shutdown. When a client receives a closeConnection message before the connection is
closed, the condition is an orderly shutdown. If the message is not received, the con-
dition is an abnormal shutdown. In an abnormal shutdown, the ORB should raise a
COMM_FAILURE exception whereas in an orderly shutdown, the ORB should **not** raise
an exception and should try to re-establish a new connection transparently.

OmniORB2 implements this semantics completely. However, it is known that
some ORBs are not (yet) able to distinguish between an orderly and an abnormal shut-
down. Usually this is manifested as the client in these ORBs seeing a COMM_FAILURE
occasionally when connected to an omniORB2 server. The workaround is either to

catch the exception in the application code and retries or to turn off the idle connection shutdown inside the omniORB2 server.

7.5   Connection Acceptance

OmniORB2 provides the hook to implement a connection acceptance policy. Inside the ORB runtime, a thread is dedicated to receive new connections. When the thread is given the handle of a new connection by the operating system, it calls the policy module to decide if the connection can be accepted. If the answer is yes, the ORB will start serving requests coming in from that connection. Otherwise, the connection is shutdown immediately.

There can be a number of policy module implementations. The basic one is a dummy module which just accepts every connection.

In addition, a host-based access control module is available on unix platforms. The module uses the IP address of the client to decide if the connection can be accepted. The module is implemented using *tcp_wrappers 7.6*. The access control policy can be defined as rules in two access control files: `hosts.allow` and `hosts.deny`. The syntax of the rules is described in the manual page `hosts_access(5)` which can be found in appendix A. The syntax defines a simple access control language that is based on client (host name/address, user name), and server (process name, host name/address) patterns. When searching for a match on the server process name, the ORB uses the value of `omniORB::serverName`. `ORB_init` uses the argument `argv[0]` to set the default value of this variable. This can be overridden by the application by passing the option: `-ORBserverName <string>` to `ORB_init`.

The default location of the access control files is `/etc`. This can be overridden by the extra options in `omniORB.cfg`. For instance:

```
# omniORB configuration file - extra options
#

GATEKEEPER_ALLOWFILE    /project/omni/var/hosts.allow

GATEKEEPER_DENYFILE     /project/omni/var/hosts.deny
```

As each policy module is implemented as a separate library, the choice of policy module is determined at program linkage time.

For instance, if the host-based access control module is in use:

```
% eg1 -ORBtraceLevel 2
omniORB2 gateKeeper is tcpwrapGK 1.0 - based on tcp_wrappers_7.6
I said,"Hello!". The Object said,"Hello!"
```

Whereas if the dummy module is in use:

```
% eg1 -ORBtraceLevel 2
omniORB2 gateKeeper is not installed. All incoming are accepted.
I said,"Hello!". The Object said,"Hello!"
```

Chapter 8

Proxy Objects

When a client acquires a reference to an object in another address space, omniORB2 creates a local representation of the object and returns a pointer to this object as its object reference. The local representation is known as the proxy object.

The proxy object maps each IDL operation into a method to deliver invocations to the remote object. The method implements argument marshalling using the ORB runtime. When the ORB runtime detects an error condition, it may raise a system exception. These exceptions will normally be propagated by the proxy object to the application code. However, there may be applications that prefer to have the system exceptions trapped in the proxy object. For these applications, it is possible to install exception handlers for individual proxy object or all proxy objects. The API to do this will be explained in this chapter.

As described in section 6.2, proxy objects are created by instances of the proxyObjectFactory class. For each IDL interface A, the stubs of A contains a derived class of proxyObjectFactory (A_proxyObjectFactory). This derived class is responsible for creating proxy objects for A. This process is completely transparent to the application. However, there may be applications that require greater control on the creation of proxy objects or even want to change the behavior of the proxy objects. To cater for this requirement, applications can override the default proxyObjectFactories and install their own versions of proxyObjectFactories. The way to do this will be explained in this chapter.

8.1   System Exception Handlers

By default, all system exceptions, with the exception of CORBA::TRANSIENT, are propagated by the proxy objects to the application code. Some applications may prefer to trap these exceptions within the proxy objects so that the application logic does not have to deal with the error condition. For example, when a CORBA::COMM_FAILURE is received, an application may just want to retry the invocation until it finally succeeds. This approach is useful for objects that are persistent and their operations are idempotent.

OmniORB2 provides a set of functions to install exception handlers. Once they are installed, proxy objects will call these handlers when the target system exceptions are raised by the ORB runtime. Exception handlers can be installed for

CORBA::TRANSIENT, CORBA::COMM␣FAILURE and CORBA::SystemException. The last handler covers all system exceptions other than the two covered by the first two handlers. An exception handler can be installed for individual proxy object or it can be installed for all proxy objects in the address space.

8.1.1   CORBA::TRANSIENT handlers

When a CORBA::TRANSIENT exception is raised by the ORB runtime, the default behaviour of the proxy objects is to retry indefinitely until the operation succeeds. Successive retries will be delayed progressively by multiples of `om-niORB::defaultTransientRetryDelayIncrement`. The delay will be limited to a maximum specified by `omniORB::defaultTransientRetryDelayMaximum`. The unit of both values are in seconds.

The ORB runtime will raised CORBA::TRANSIENT under the following conditions:

1. When a **cached** network connection is broken while an invocation is in progress. The ORB will try to open a new connection at the next invocation.

2. When the proxy object has been redirected by a location forward message by the remote object to a new location and the object at the new location cannot be contacted. In addition to the CORBA::TRANSIENT exception, the proxy object also resets its internal state so that the next invocation will be directed at the original location of the remote object.

3. When the remote object reports CORBA::TRANSIENT.

Applications can override the default behaviour by installing their own exception handler. The API to do so is summarised below:

```
class omniORB {

public:

typedef CORBA::Boolean (*transientExceptionHandler_t)(void* cookie,
                                                CORBA::ULong n_retries,
                                                const CORBA::TRANSIENT& ex);

static void installTransientExceptionHandler(void* cookie,
                                                transientExceptionHan-
dler_t fn);

static void installTransientExceptionHandler(CORBA::Object_ptr obj,
                                                void* cookie,
                                                transientExceptionHan-
dler_t fn);

static CORBA::ULong defaultTransientRetryDelayIncrement;
static CORBA::ULong defaultTransientRetryDelayMaximum;
```

```
}
```

The overloaded functions `installTransientExceptionHandler` can be used to install the exception handlers for CORBA::TRANSIENT.

Two overloaded forms are available. The first form install an exception handler for all object references except for those which have an exception handler installed by the second form, which takes an addition argument to identify the target object reference. The argument `cookie` is an opaque pointer which will be passed on by the ORB when it calls the exception handler.

An exception handler will be called by proxy objects with three arguments. The `cookie` is the opaque pointer registered by `installTransientExceptionHandler`. The argument `n_retries` is the number of times the proxy has called this handler for the same invocation. The argument `ex` is the value of the exception caught. The exception handler is expected to do whatever is appropriate and returns a boolean value. If the return value is TRUE(1), the proxy object would retry the operation again. If the return value is FALSE(0), the CORBA::TRANSIENT exception would be propagated into the application code.

The following sample code installs a simple exception handler for all objects and for a specific object:

```
CORBA::Boolean my_transient_handler1 (void* cookie,
                                      CORBA::ULong retries,
                                      const CORBA::TRANSIENT& ex)
{
   cerr << ''transient handler 1 called.'' << endl;
   return 1;            // retry immediately.
}

CORBA::Boolean my_transient_handler2 (void* cookie,
                                      CORBA::ULong retries,
                                      const CORBA::TRANSIENT& ex)
{
   cerr << ''transient handler 2 called.'' << endl;
   return 1;            // retry immediately.
}


static Echo_ptr myobj;

void installhandlers()
{
   omniORB::installTransientExceptionHandler(0,my_transient_handler1);
   // All proxy objects will call my_transient_handler1 from now on.

   omniORB::installTransientExceptionHandler(myobj,0,my_transient_handler2);
   // The proxy object of myobj will call my_transient_handler2 from now on.
}
```

8.1.2   CORBA::COMM FAILURE

When the ORB runtime fails to establish a network connection to the remote object and none of the conditions listed above for raising a CORBA::TRANSIENT is applicable, it raises a CORBA::COMM FAILURE exception.

The default behaviour of the proxy objects is to propagate this exception to the application.

Applications can override the default behaviour by installing their own exception handlers. The API to do so is summarised below:

```
class omniORB {

public:

typedef CORBA::Boolean (*commFailureExceptionHan-
dler_t)(void* cookie,
                                          CORBA::ULong n_retries,
                                          const CORBA::COMM_FAILURE& ex);

static void installCommFailureExceptionHandler(void* cookie,
                                          commFailureException-
Handler_t fn);

static void installCommFailureExceptionHandler(CORBA::Object_ptr obj,
                                          void* cookie,
                                          commFailureExceptionHandler_t
                                          fn);
}
```

The functions are equivalent to their counterparts for CORBA::TRANSIENT.

8.1.3   CORBA::SystemException

To report an error condition, the ORB runtime may raise other SystemExceptions. If the exception is neither CORBA::TRANISENT nor CORBA::COMM FAILURE, the default behaviour of the proxy objects is to propagate this exception to the application.

Application can override the default behaviour by installing their own exception handlers. The API to do so is summarised below:

```
class omniORB {

public:

typedef CORBA::Boolean (*systemExceptionHandler_t)(void* cookie,
                                          CORBA::ULong n_retries,
                                          const CORBA::SystemException& ex);

static void installSystemExceptionHandler(void* cookie,
                                          systemExceptionHan-
dler_t fn);
```

```
static void installSystemExceptionHandler(CORBA::Object_ptr obj,
                                          void* cookie,
                                          systemExceptionHan-
dler_t fn);
}
```

The functions are equivalent to their counterparts for CORBA::TRANSIENT.

8.2 Proxy Object Factories

This section describes how an application can control the creation or change the behaviour of proxy objects.

8.2.1 Background

For each interface A, its stub contains a proxy factory class- A_proxyObjectFactory. This class is derived from CORBA::proxyObjectFactory and implements three virtual functions:

```
class A_proxyObjectFactory : public virtual CORBA::proxyObjectFactory {
public:

  virtual const char *irRepoId() const;

  virtual _CORBA_Boolean  is_a(const char *base_repoId) const;

  virtual CORBA::Object_ptr newProxyObject(Rope *r,
                                           CORBA::Octet *key,
                                           size_t keysize,
                                           IOP::TaggedProfileList
                                           *profiles,
                                           CORBA::Boolean release);

};
```

As described in chapter 6, the functions allow the ORB runtime to perform type checking. The function newProxyObject creates a proxy object for A based on its input arguments. The return value is a pointer to the class _proxy_A which is automatically re-casted into a CORBA::Object_ptr. _proxy_A implements the proxy object for A:

```
class _proxy_A :  public virtual A {
public:

  _proxy_A (Rope *r,
            CORBA::Octet *key,
            size_t keysize,IOP::TaggedProfileList *profiles,
```

```
              CORBA::Boolean release);
   virtual ~_proxy_A();

   // plus other internal functions.

};
```

The stub of A guarantees that exactly **one** instance of A_proxyObjectFactory is instantiated when an application is executed.   The constructor of A_proxyObjectFactory, via its base class proxyObjectFactory links the instance into the ORB's proxy factory list.

Newly instantiated proxy object factories are always entered at the front of the ORB's proxy factory list. Moreover, when the ORB searches for a match on the type, it always stops at the first match. In other words, when additional instances of A_proxyObjectFactory or derived classes of it are created, the last instantiation will override earlier instantiations to be the proxy factory selected to create proxy objects of A. This property can be used by an application to install its own proxy object factories.

8.2.2   An Example

Using the Echo example in chapter 2 as the basis, one can tell the ORB to use a modified proxy object class to create proxy objects. The steps involved are as follows:

8.2.2.1   Define a new proxy class

We define a new proxy class to cache the result of the last invocation of echoString.

```
class _new_proxy_Echo : public virtual _proxy_Echo {
public:
   _new_proxy_Echo (Rope *r,
                    CORBA::Octet *key,
                    size_t keysize,IOP::TaggedProfileList *profiles,
                    CORBA::Boolean release)
         : _proxy_Echo(r,key,keysize,profiles,release),
           omniObject(Echo_IntfRepoID,r,key,keysize,profiles,release)
   {
     // You have to look at the _proxy_Echo class and copy from its
     // ctor all the explicit ctor calls to its base member.
   }
   virtual ~_new_proxy_Echo() {}


   virtual char* echoString(const char* mesg) {
     //
     // Only calls the remote object if the argument is differ-
ent from the
     // last invocation.

     omni_mutex_lock sync(lock);
     if ((char*)last_arg) {
```

```
        if (strcmp(mesg,(char*)last_arg) == 0) {
            return CORBA::string_dup(last_result);
        }
      }
    }
    char* res = _proxy_Echo::echoString(mesg);
    last_arg = mesg;
    last_result = (const char*) res;
    return res;
  }

private:
  omni_mutex        lock;
  CORBA::String_var last_arg;
  CORBA::String_var last_result;
};
```

8.2.2.2 Define a new proxy factory class

Next, we define a new proxy factory class to instantiate _new_proxy_Echo as proxy objects for `Echo`.

```
class _new_Echo_proxyObjectFactory : public vir-
tual Echo_proxyObjectFactory
{
public:
   _new_Echo_proxyObjectFactory () {}
   virtual ~_new_Echo_proxyObjectFactory() {}

   // Only have to override newProxyObject
   virtual CORBA::Object_ptr newProxyObject(Rope *r,
                                            CORBA::Octet *key,
                                            size_t keysize,
                                            IOP::TaggedProfileList *pro-
files,
                                            CORBA::Boolean re-
lease) {
      _new_proxy_Echo *p = new _new_proxy_Echo(r,key,keysize,profiles,release);
      return p;
   }
};
```

Finally, we have to instantiate a single instance of the new proxy factory in the application code.

```
int main(int argc, char** argv)
{
   // Other initialisation steps

   _new_Echo_proxyObjectFactory* f =  new _new_Echo_proxyObjectFactory;

   // Use the new operator to instantiate the proxy fac-
tory and never
```

```
    // call the delete operator on this instance.

    // From this point onwards, _new_proxy_Echo will be used to create
    // proxy objects for Echo.

}
```

8.2.3   Further Considerations

Notice that the ORB may call `newProxyObject` multiple times to create proxy objects for the same remote object. In other words, the ORB does not guarantee that only one proxy object is created for each remote object. For applications that require this guarantee, it is necessary to check within `newProxyObject` whether a proxy object has already been created for the current request. If the argument `Rope* r` points to the same structure and the content of the sequence `CORBA::Octet* key` is the same, then an existing proxy object can be returned to satisfy the current request. Do not forget to call `CORBA::duplicate()` before returning the object reference.

`newProxyObject` may be called concurrently by different threads within the ORB. Needless to say, the function must be thread-safe.

Chapter 9

Type Any and TypeCode

The CORBA specification provides for a type that can hold the value of any OMG IDL type. This type is known as type Any. The OMG also specifies a pseudo-object, TypeCode, that can encode a description of any type specifiable in OMG IDL.

In this chapter, an example demonstrating the use of type Any is presented. This is followed by sections describing the behaviour of type Any and TypeCode in omniORB2. For further information on type Any, refer to the C++ Mapping section of the CORBA 2 specification [OMG99a], and for more information on TypeCode, refer to the Interface Repository chapter in the CORBA core section of the CORBA 2 specification.

WARNING: Since 2.8.0, omniORB2 has been updated to CORBA 2.3. In order to comply with the 2.3 specification, it is necessary to change the semantics of *the extraction of string, object reference and typecode from an Any.* The memory of the extracted values of these types now belong to the Any value. The storage is freed when the Any value is deallocated. Previously the extracted value is a copy and the application is responsible to release the storage. It is not possible to detect the old usage at compile time. In particular, unmodified code that uses the affected Any extraction operators would most certainly cause runtime errors to occur. To smooth the transition from the old usage to the new, an ORB configuration variable `om-niORB::omniORB_27_CompatibleAnyExtraction` can be set to revert the any extraction operators to the old semantics.

9.1 Example using type Any

Before going through this example, you should make sure that you have read and understood the examples in chapter 2. The source code for this example is included in the omniORB2 distribution, in the directory src/examples/anyExample. A listing of the source code is provided at the end of this chapter.

9.1.1 Type Any in IDL

Type Any allows one to delay the decision on the type used in an operation until run-time. To use type any in IDL, use the keyword `any`, as in the following example:

```
// IDL
```

```
interface anyExample {
  any testOp(in any mesg);
};
```

The operation `testOp()` in this example can now take any value expressible in OMG IDL as an argument, and can also return any type expressible in OMG IDL.

Type Any is mapped into C++ as the type `CORBA::Any`. When passed as an argument or as a result of an operation, the following rules apply:

In	InOut	Out	Return
const CORBA::Any&	CORBA::Any&	CORBA::Any*&	CORBA::Any*

So, the above IDL would map to the following C++

```
// C++

class anyExample_i : public virtual _sk_anyExample {
public:
  anyExample_i() { }
  virtual ~anyExample_i() { }
  virtual CORBA::Any* testOp(const CORBA::Any& a);
};
```

9.1.2   Inserting and Extracting Basic Types from an Any

The question now arises as to how values are inserted into and removed from an Any. This is achieved using two overloaded operators: `<<=` and `>>=` .

Two insert a value into an Any, the `<<=` operator is used, as in this example:

```
// C++

CORBA::Any an_any;
CORBA::Long l = 100;
an_any <<= l;
```

Note that the overloaded `<<=` operator has a return type of `void`.

To extract a value, the `>>=` operator is used, as in this example (where the Any contains a long):

```
// C++

CORBA::Long l;
an_any >>= l;

cout << "This is a long: " << l << endl;
```

The overloaded `>>=` operator returns a CORBA::Boolean. If an attempt is made to extract a value from an Any when it contains a different value (e.g. an attempt to extract a long from an Any containing a double), the overloaded `>>=` operator will return False; otherwise it will return True. Thus, a common tactic to extract values from an Any is as follows:

```
// C++


CORBA::Long l;
CORBA::Double d;
const char* str;      // From CORBA 2.3 onwards, uses const char*
                      // instead of char*.

if (an_any >>= l) {
    cout << "Long: " << l << endl;
}
else if (an_any >>= d) {
    cout << "Double: " << d << endl;
}
else if (an_any >>= str) {
    cout << "String: " << str << endl;
    // Since 2.8.0 the storage of the extracted string is still
    // owned by the any.
    // In pre-omniORB 2.8.0 releases, the string returned is a copy.
}
else {
    cout << "Unknown value." << endl;
}
```

9.1.3   Inserting and Extracting Constructed Types from an Any

It is also possible to insert and extract constructed types and object references from an Any. `omniidl2` will generate insertion and extraction operators for the constructed type. Note that it is necessary to specify the `-a` command-line flag when running omniidl2 in order to generate these operators. The following example illustrates the use of constructed types with type Any:

```
// IDL

struct testStruct {
  long l;
  short s;
};


interface anyExample {
  any testOp(in any mesg);
};
```

Upon compiling the above IDL with `omniidl2 -a`, the following overloaded operators are generated:

1. `void operator<<=(CORBA::Any&, const testStruct& )`

2. `void operator<<=(CORBA::Any&, testStruct* )`

3. `CORBA::Boolean operator>>=(const CORBA::Any&, const test-Struct*&)`

Operators of this form are generated for all constructed types, and for interfaces.
The first operator, *(1)* , copies the constructed type, and inserts it into the Any. The second operator, *(2)* , inserts the constructed type into the Any, and then manages it. Note that if the second operator is used, the Any consumes the constructed type, and the caller should not used the pointer to access the data after insertion. The following is an example of how to insert a value into an Any using operator *(1)* :

```
// C++

CORBA::Any an_any;

testStruct t;
t.l = 456;
t.s = 8;

an_any <<= t;
```

The third operator, *(3)* , is used to extract the constructed type from the Any, and can be used as follows:

```
const testStruct* tp;    // From CORBA 2.3 onwards, use
                         // const testStruct* instead of testStruct*

if (an_any >>= tp) {
    cout << "testStruct: l: " << tp->l << endl;
    cout << "            s: " << tp->s << endl;
}
else {
    cout << "Unknown value contained in Any." << endl;
}
```

As with basic types, if an attempt is made to extract a type from an Any that does not contain a value of that type, the extraction operator returns False. If the Any does contain that type, the extraction operator returns True. If the extraction is successful, the caller's pointer will point to memory managed by the Any. The caller must not delete or otherwise change this storage, and should not use this storage after the contents of the Any are replaced (either by insertion or assignment), or after the Any has

been destroyed. In particular, management of the pointer should not be assigned to a _var type.

WARNING!!! In pre-omniORB 2.8.0 releases, it is unclear in the CORBA specification whether or not object references should be managed by an Any. The omniORB2 implementation leaves management of an extracted object reference to the caller. Therefore, the programmer should release object references and TypeCodes that have been extracted from an Any. The same also applies to string extraction. CORBA 2.3 has clarified this issue and decreed that the management of an extracted object reference still belongs to the Any! Since 2.8.0, the omniORB2 implementation conforms to the CORBA 2.3 specification. For backward compatibility, the runtime variable `omniORB::omniORB_27_CompatibleAnyExtraction` can be set to 1 to get back the old behaviour. Notice that this should be used as a transitional measure and in the long run, applications should be written to use the new behaviour.

If the extraction fails, the caller's pointer will be set to point to null.

Note that there are special rules for inserting and extracting arrays (using _forany types), and for inserting and extracting booleans, octets, chars, and bounded strings. Please refer to the C++ Mapping chapter of the CORBA 2 specification [OMG99a] for further information.

9.2   Type Any in omniORB2

This section contains some notes on the use and behaviour of type Any in omniORB2.

Generating Insertion and Extraction Operators.   To generate type Any insertion and extraction operators for constructed types and interfaces, the `-a` command line flag should be specified when running `omniidl2`.

TypeCode comparison when extracting from an Any.   When an attempt is made to extract a type from an Any, the TypeCode of the type is checked for **equivalence** with the TypeCode of the type stored by the Any. The equivalent() test in the TypeCode interface is used for this purpose[1].

Examples:

```
// IDL 1

typedef double Double1;

struct Test1 {
Double1 a;
};
```

------

[1]In pre-omniORB 2.8.0 releases, omniORB2 performs an equality test and will ignore any alias Type-Codes (`tk_alias`) when making this comparison. The semantics is similar to the equivalent() test in the TypeCode interface of CORBA 2.3.

```
// IDL 2

typedef double Double2;

struct Test1 {
  Double2 a;
};
```

If an attempt is made to extract the type `Test1` defined in IDL 1 from an Any containing the `Test1` defined in IDL 2, this will succeed (and vice-versa), as the two types differ only by an alias.

Object references.   WARNING!! In pre-omniORB 2.8.0 releases, the type Any does not manage object reference types - it was unclear in the pre-2.3 CORBA specification whether this is required or not. Therefore, the programmer should release object references and pseudo-objects (such as TypeCode) that have been extracted from an Any. Type Any will, however, manage constructed types (as per the CORBA 2 specification) - constructed types extracted from an Any should not be deleted, as they will be deleted by the Any when it is destroyed. CORBA 2.3 has clarified this issue and decreed that the management of an extracted object reference still belongs to the Any! Since 2.8.0, the omniORB2 implementation conforms to the CORBA 2.3 specification. For backward compatibility, the runtime variable `om-niORB::omniORB_27_CompatibleAnyExtraction` can be set to 1 to get back the old behaviour.

Top-level aliases.   When a type is inserted into an Any, the Any stores both the value of the type and the TypeCode for that type. The treatment of top-level aliases from omniORB 2.8.0 onwards is different from pre-omniORB 2.8.0 releases.

In pre-omniORB 2.8.0 releases, if there are any top-level `tk_alias` TypeCodes in the TypeCode, they will be removed from the TypeCode stored in the Any. Note that this does not affect the `_tc_` TypeCode generated to represent the type (see section on TypeCode, below). This behaviour is necessary, as two types that differ only by a top-level alias can use the same insertion and extraction operators. If the `tk_alias` is not removed, one of the types could be transmitted with an incorrect `tk_alias` TypeCode. Example:

```
// IDL 3

typedef sequence<double> seqDouble1;
typedef sequence<double> seqDouble2;
typedef seqDouble2       seqDouble3;
```

If either seqDouble1 or seqDouble2 is inserted into an Any, the TypeCode stored will be for a `sequence<double>`, and not for an alias to a `sequence<double>`.

From omniORB 2.8.0 onwards, there are two changes. Firstly, in the example, seq-Double1 and seqDouble2 are now distinct types and therefore each has its own set of C++ operators for Any insertion and extraction. Secondly, the top level aliases are not removed. For example, if seqDouble3 is inserted into an Any, the insertion operator for seqDouble2 is invoked (because seqDouble3 is just a C++ typedef of seqDouble2). Therefore, the typecode in the Any would be that of seqDouble2. If this is not desirable, one can use the new member function 'void type(TypeCode ptr)' of the Any interface to explicitly set the typecode to the correct one.

Removing aliases from TypeCodes. Some ORBs (e.g. Orbix) will not accept Type-Codes containing `tk alias` TypeCodes. When using type Any while interoperating with these ORBs, it is necessary to remove `tk alias` TypeCodes from throughout the TypeCode representing a constructed type.

To remove all `tk alias` TypeCodes from TypeCodes stored in Anys, supply the `-ORBtcAliasExpand 1` command-line flag when running an omniORB2 executable. There will be some (small) performance penalty when inserting values into an Any.

Note that the `tc` TypeCodes generated for all constructed types will contain the complete TypeCode for the type (including any `tk alias` TypeCodes), regardless of whether the `-ORBtcAliasExpand` flag is set to 1 or not.

Recursive TypeCodes. omniORB2 does now (as of version 2.7) support recursive TypeCodes. This means that types such as the following can be inserted or extracted from an Any:

```
// IDL 4

struct Test4 {
  sequence<Test4> a;
};
```

Type-unsafe construction and insertion. If using the type-unsafe Any constructor, or the `CORBA::Any::replace()` member function, ensure that the value returned by the `CORBA::Any::value()` member function and the TypeCode returned by the `CORBA::Any::type()` member function are used as arguments to the constructor or function. Using other values or TypeCodes may result in a mismatch, and is undefined behaviour.

Note that a non-CORBA 2 function,

```
CORBA::ULong CORBA::Any::NP_length() const
```

is supplied. This member function returns the length of the value returned by the `CORBA::Any::value()` member function. It may be necessary to use this function if the Any's value is to be stored in a file.

Threads and type Any.   Inserting and extracting simultaneously from the same Any (in 2 different threads) is undefined behaviour.

Extracting simultaneously from the same Any (in 2 or more different threads) also leads to undefined behaviour. It was decided not to protect the Any with a mutex, as this condition should rarely arise, and adding a mutex would lead to performance penalties.

9.3   TypeCode in omniORB2

This section contains some notes on the use and behaviour of TypeCode in omniORB2

TypeCodes in IDL.   When using TypeCodes in IDL, note that they are defined in the CORBA scope. Therefore, CORBA::TypeCode should be used. Example:

```
// IDL 5

struct Test5 {
  long length;
  CORBA::TypeCode desc;
};
```

orb.idl   Inclusion of the file `orb.idl` in IDL using CORBA::TypeCode is optional. An empty orb.idl file is provided for compatibility purposes.

Generating TypeCodes for constructed types.   To generate a TypeCode for constructed types, specify the `-a` command-line flag when running omniidl2. This will generate a ‿tc‿ TypeCode describing the type, at the same scope as the type (as per the CORBA 2 specification). Example:

```
// IDL 6

struct Test6 {
 double a;
 sequence<long> b;
};
```

A TypeCode, ‿tc‿Test6, will be generated to describe the struct `Test6`. The operations defined in the TypeCode interface (see section 6.7.1 of the CORBA 2 specification [OMG99a]) can be used to query the TypeCode about the type it represents.

TypeCode equality.   The behaviour of `CORBA::TypeCode::equal()` member function from omniORB 2.8.0 onwards is different from pre-omniORB 2.8.0 releases. In summary, the pre-omniORB 2.8.0 is close to the semantics of the new `CORBA::TypeCode::equivalent()` member function. Details are as follows:

The `CORBA::TypeCode::equal()` member function will return true only if the two TypeCodes are **exactly** the same. `tk‿alias` TypeCodes are included in this comparison, unlike the comparison made when values are extracted from an Any (see section on Any, above).

In pre-omniORB 2.8.0 releases, equality test would ignore the optional fields when one of the fields in the two typecodes is empty. For example, if one of the Type-Codes being checked is a `tk_struct`, `tk_union`, `tk_enum`, or `tk_alias`, and has an empty repository ID parameter, then the repository ID parameter will be ignored when checking for equality. Similarly, if the `name` or `member_name` parameters of a TypeCode are empty strings, they will be ignored for equality checking purposes. This is because a CORBA 2 ORB does not have to include these parameters in a Type-Code (see the Interoperability section of the CORBA 2 specification [OMG99a]). Note that these (optional) parameters are included in TypeCodes generated by omniORB2.

Since CORBA 2.3, the issue of typecode equality has been clarified. There is now a new member `CORBA::TypeCode::equivalent()` which provides the semantics of the `CORBA::TypeCode::equal()` as implemented in pre-omniORB 2.8.0 releases. So from omniORB 2.8.0 onwards, the `CORBA::TypeCode::equal()` function has been changed to enforce strict equality. The pre-2.8.0 behaviour can be obtained with `equivalent()`.

9.4   Source Listing

9.4.1   anyExample_impl.cc

```
// anyExample_impl.cc - This is the source code of the exam-
ple used in
//                      Chapter 9 "Type Any and Type-
Code" of the omniORB2
//                      users guide.
//
//              This is the object implementation.
//
// Usage: anyExample_impl
//
//          On startup, the object reference is registered with the
//          COS naming service. The client uses the naming service to
//          locate this object.
//
//          The name which the object is bound to is as follows:
//                  root   [context]
//                    |
//                  text   [context] kind [my_context]
//                    |
//                  anyExample  [object]  kind [Object]
//


#include <iostream.h>

#include "anyExample.hh"


static CORBA::Boolean bindObjectToName(CORBA::ORB_ptr,CORBA::Object_ptr);


class anyExample_i : public virtual _sk_anyExample {
public:
  anyExample_i() { }
  virtual ~anyExample_i() { }
  virtual CORBA::Any* testOp(const CORBA::Any& a);
};


CORBA::Any*
anyExample_i::testOp(const CORBA::Any& a) {

  cout << "Any received, containing: " << endl;

#ifndef NO_FLOAT
  CORBA::Double d;
#endif

  CORBA::Long l;
  const char* str;        // From CORBA 2.3 onwards, uses const char*
```

```
                            // instead of char*.

  const testStruct* tp;    // From CORBA 2.3 onwards, use
                           // const testStruct* in-
stead of testStruct*


  if (a >>= l) {
    cout << "Long: " << l << endl;
  }
#ifndef NO_FLOAT
  else if (a >>= d) {
    cout << "Double: " << d << endl;
  }
#endif
  else if (a >>= str) {
    cout << "String: " << str << endl;
    // Since 2.8.0 the storage of the extracted string is still
    // owned by the any.
    // In pre-omniORB 2.8.0 releases, the string returned is a copy.
  }
  else if (a >>= tp) {
    cout << "testStruct: l: " << tp->l << endl;
    cout << "             s: " << tp->s << endl;
  }
  else {
    cout << "Unknown value." << endl;
  }

  CORBA::Any* ap = new CORBA::Any;

  *ap <<= (CORBA::ULong) 314;

  cout << "Returning Any containing: ULong: 314\n" << endl;
  return ap;
}


int
main(int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA");

  anyExample_i *myobj = new anyExample_i();
  myobj->_obj_is_ready(boa);

  {
    anyExample_var myobjRef = myobj->_this();
    if (!bindObjectToName(orb,myobjRef)) {
      return 1;
    }
  }
```

```
  boa->impl_is_ready();
  // Tell the BOA we are ready. The BOA's default be-
haviour is to block
  // on this call indefinitely.

  return 0;
}


static
CORBA::Boolean
bindObjectToName(CORBA::ORB_ptr orb,CORBA::Object_ptr obj)
{
  CosNaming::NamingContext_var rootContext;

  try {
    // Obtain a reference to the root context of the Name service:
    CORBA::Object_var initServ;
    initServ = orb->resolve_initial_references("NameService");

    // Narrow the object returned by resolve_initial_references()
    // to a CosNaming::NamingContext object:
    rootContext = CosNaming::NamingContext::_narrow(initServ);
    if (CORBA::is_nil(rootContext))
      {
        cerr << "Failed to narrow naming context." << endl;
        return 0;
      }
  }
  catch(CORBA::ORB::InvalidName& ex) {
    cerr << "Service required is invalid [does not exist]." << endl;
    return 0;
  }


  try {
    // Bind a context called "test" to the root context:

    CosNaming::Name contextName;
    contextName.length(1);
    contextName[0].id   = (const char*) "test";     // string copied
    contextName[0].kind = (const char*) "my_context"; // string copied
    // Note on kind: The kind field is used to indicate the type
    // of the object. This is to avoid conventions such as that used
    // by files (name.type -- e.g. test.ps = postscript etc.)

    CosNaming::NamingContext_var testContext;
    try {
      // Bind the context to root, and assign testContext to it:
      testContext = rootContext->bind_new_context(contextName);
    }
    catch(CosNaming::NamingContext::AlreadyBound& ex) {
```

```cpp
      // If the context already exists, this excep-
tion will be raised.
      // In this case, just resolve the name and assign testContext
      // to the object returned:
      CORBA::Object_var tmpobj;
      tmpobj = rootContext->resolve(contextName);
      testContext = CosNaming::NamingContext::_narrow(tmpobj);
      if (CORBA::is_nil(testContext)) {
        cerr << "Failed to narrow naming context." << endl;
        return 0;
      }
    }

    // Bind the object (obj) to testContext, naming it anyExample:
    CosNaming::Name objectName;
    objectName.length(1);
    objectName[0].id   = (const char*) "anyExam-
ple";   // string copied
    objectName[0].kind = (const char*) "Object"; // string copied


    // Bind obj with name anyExample to the testContext:
    try {
      testContext->bind(objectName,obj);
    }
    catch(CosNaming::NamingContext::AlreadyBound& ex) {
      testContext->rebind(objectName,obj);
    }
    // Note: Using rebind() will overwrite any Object previ-
ously bound
    //       to /test/anyExample with obj.
    //       Alternatively, bind() can be used, which will raise a
    //       CosNaming::NamingContext::AlreadyBound excep-
tion if the name
    //       supplied is already bound to an object.
  }
  catch (CORBA::COMM_FAILURE& ex) {
    cerr << "Caught system exception COMM_FAILURE, unable to con-
tact the "
         << "naming service." << endl;
    return 0;
  }
  catch (omniORB::fatalException& ex) {
    throw;
  }
  catch (...) {
    cerr << "Caught a system exception while using the naming ser-
vice."<< endl;
    return 0;
  }
  return 1;
}
```

9.4.2   anyExample_clt.cc

```
// anyExample_clt.cc -  This is the source code of the exam-
ple used in
//                      Chapter 9 "Type Any and Type-
Code" of the omniORB2
//                      users guide.
//
//           This is the client. It uses the COSS naming service
//           to obtain the object reference.
//
// Usage: anyExample_clt
//
//
//         On startup, the client lookup the object refer-
ence from the
//         COS naming service.
//
//         The name which the object is bound to is as follows:
//              root  [context]
//                 |
//              text  [context] kind [my_context]
//                 |
//              anyExample  [object]  kind [Object]
//

#include <iostream.h>
#include "anyExample.hh"

static CORBA::Object_ptr getObjectReference(CORBA::ORB_ptr orb);
static void invokeOp(anyExample_ptr& tobj, const CORBA::Any& a);

int
main (int argc, char **argv)
{
  CORBA::ORB_ptr orb = CORBA::ORB_init(argc,argv,"omniORB2");
  CORBA::BOA_ptr boa = orb->BOA_init(argc,argv,"omniORB2_BOA");
  CORBA::Object_var obj;

  try {
    obj = getObjectReference(orb);
  }
  catch(CORBA::COMM_FAILURE& ex) {
    cerr << "Caught system exception COMM_FAILURE, unable to con-
tact the "
         << "object." << endl;
    return -1;
  }
  catch(omniORB::fatalException& ex) {
    cerr << "Caught omniORB2 fatalException. This indi-
cates a bug is caught "
         << "within omniORB2.\nPlease send a bug report.\n"
         << "The exception was thrown in file: " << ex.file() << "\n"
```

```
        << "                                line: " << ex.line() << "\n"
        << "The error message is: " << ex.errmsg() << endl;
    return -1;
  }
  catch(...) {
    cerr << "Caught a system exception." << endl;
    return -1;
  }


  anyExample_ptr tobj = anyExample::_narrow(obj);

  if (CORBA::is_nil(tobj)) {
    cerr << "Can't narrow object reference to type anyExam-
ple." << endl;
    return -1;
  }


  CORBA::Any a;

  try {
    // Sending Long
    CORBA::Long l = 100;
    a <<= l;
    cout << "Sending Any containing Long: " << l << endl;
    invokeOp(tobj,a);

    // Sending Double
#ifndef NO_FLOAT
    CORBA::Double d = 1.2345;
    a <<= d;
    cout << "Sending Any containing Double: " << d << endl;
    invokeOp(tobj,a);
#endif

    // Sending String
    const char* str = "Hello";
    a <<= str;
    cout << "Sending Any containing String: " << str << endl;
    invokeOp(tobj,a);

    // Sending testStruct  [Struct defined in IDL]
    testStruct t;
    t.l = 456;
    t.s = 8;
    a <<= t;
    cout << "Sending Any containing testStruct: l: " << t.l << endl;
    cout << "                                    s: " << t.s << endl;
    invokeOp(tobj,a);
  }
  catch(CORBA::COMM_FAILURE& ex) {
    cerr << "Caught system exception COMM_FAILURE, unable to con-
```

```
tact the "
         << "object." << endl;
    return -1;
  }
  catch(omniORB::fatalException& ex) {
    cerr << "Caught omniORB2 fatalException. This indi-
cates a bug is caught "
         << "within omniORB2.\nPlease send a bug report.\n"
         << "The exception was thrown in file: " << ex.file() << "\n"
         << "                                 line: " << ex.line() << "\n"
         << "The error message is: " << ex.errmsg() << endl;
    return -1;
  }
  catch(...) {
    cerr << "Caught a system exception." << endl;
    return -1;
  }

  return 0;
}

static
CORBA::Object_ptr
getObjectReference(CORBA::ORB_ptr orb)
{
  CosNaming::NamingContext_var rootContext;

  try {
    // Obtain a reference to the root context of the Name service:
    CORBA::Object_var initServ;
    initServ = orb->resolve_initial_references("NameService");

    // Narrow the object returned by resolve_initial_references()
    // to a CosNaming::NamingContext object:
    rootContext = CosNaming::NamingContext::_narrow(initServ);
    if (CORBA::is_nil(rootContext))
      {
        cerr << "Failed to narrow naming context." << endl;
        return CORBA::Object::_nil();
      }
  }
  catch(CORBA::ORB::InvalidName& ex) {
    cerr << "Service required is invalid [does not exist]." << endl;
    return CORBA::Object::_nil();
  }


  // Create a name object, containing the name test/context:
  CosNaming::Name name;
  name.length(2);

  name[0].id   = (const char*) "test";       // string copied
  name[0].kind = (const char*) "my_context"; // string copied
```

```cpp
  name[1].id   = (const char*) "anyExample";
  name[1].kind = (const char*) "Object";
  // Note on kind: The kind field is used to indicate the type
  // of the object. This is to avoid conventions such as that used
  // by files (name.type -- e.g. test.ps = postscript etc.)


  CORBA::Object_ptr obj;
  try {
    // Resolve the name to an object reference, and as-
sign the reference
    // returned to a CORBA::Object:
    obj = rootContext->resolve(name);
  }
  catch(CosNaming::NamingContext::NotFound& ex)
    {
      // This exception is thrown if any of the components of the
      // path [contexts or the object] aren't found:
      cerr << "Context not found." << endl;
      return CORBA::Object::_nil();
    }
  catch (CORBA::COMM_FAILURE& ex) {
    cerr << "Caught system exception COMM_FAILURE, unable to con-
tact the "
         << "naming service." << endl;
    return CORBA::Object::_nil();
  }
  catch(omniORB::fatalException& ex) {
    throw;
  }
  catch (...) {
    cerr << "Caught a system exception while using the naming ser-
vice."<< endl;
    return CORBA::Object::_nil();
  }
  return obj;
}


static void
invokeOp(anyExample_ptr& tobj, const CORBA::Any& a)
{
  CORBA::Any_var bp;

  cout << "Invoking operation." << endl;
  bp = tobj->testOp(a);

  cout << "Operation completed. Returned Any: ";
  CORBA::ULong ul;

  if (bp >>= ul) {
    cout << "ULong: " << ul << "\n" << endl;
  }
```

```
  else {
    cout << "Unknown value." << "\n" << endl;
  }
}
```

Chapter 10

Dynamic Management of Any Values

In CORBA specification 2.2, a new facility- **DynAny** is introduced. Previously, it is not possible to insert or extract constructed and other complex types from an **any** without using the stub code generated by an idl compiler for these types. This makes it impossible to write generic servers (bridges, event channels supporting filtering etc) because these servers can not have static knowledge of all the possible data types that they have to handle.

To fill this gap, the **DynAny** facility is defined to enable traversal of the data value associated with an **any** at runtime and extraction of its constituents. This facility also enables the construction of an **any** at runtime, without having static knowledge of its types.

This chapter explains how **DynAny** may be used. For completeness, you should also read the DynAny specification defined in Chapter 7 of the CORBA specification 2.2. Where possible, the implementation in omniORB2 adheres closely to the specification. However, there are areas in the specification that are ambiguous or lacking in details. A number of these issues are currently opened with the ORB revision task force. Until the issues are resolved, it is possible that a different implementation may choose to intepret the specification differently. This chapter provides clarifications to the specification, explains the interpretation used and offers some advice and warnings on potential portability problems.

Notice that the **DynAny** interface has been changed in CORBA 2.3, particularly with the addition of the support for the IDL type `valuetype`. Future releases of omniORB will be updated to implement the interface as defined in CORBA 2.3.

10.1   C++ mapping

```
// namespace CORBA

class ORB {
public:
  ...
```

89

```
  class InconsistentTypeCode : public UserException { ... };

  DynAny_ptr create_dyn_any(const Any& value);

  DynAny_ptr create_basic_dyn_any(TypeCode_ptr tc);

  DynStruct_ptr create_dyn_struct(TypeCode_ptr tc);

  DynSequence_ptr create_dyn_sequence(TypeCode_ptr tc);

  DynArray_ptr create_dyn_array(TypeCode_ptr tc);

  DynUnion_ptr create_dyn_union(TypeCode_ptr tc);

  DynEnum_ptr create_dyn_enum(TypeCode_ptr tc);

};

typedef DynAny* DynAny_ptr;
class DynAny_var { ... };

class DynAny {
public:

  class Invalid : public UserException { ... };
  class InvalidValue : public UserException { ... };
  class TypeMismatch : public UserException { ... };
  class InvalidSeq : public UserException { ... };

  typedef _CORBA_Unbounded_Sequence__Octet OctetSeq;

  TypeCode_ptr type() const;

  void assign(DynAny_ptr dyn_any) throw(Invalid,SystemException);
  void from_any(const Any& value) throw(Invalid,SystemException);
  Any* to_any() throw(Invalid,SystemException);
  void destroy();
  DynAny_ptr copy();

  DynAny_ptr current_component();
  Boolean next();
  Boolean seek(Long index);
  void rewind();

  void insert_boolean(Boolean value) throw(InvalidValue,SystemException);
  void insert_octet(Octet value) throw(InvalidValue,SystemException);
  void insert_char(Char value) throw(InvalidValue,SystemException);
  void insert_short(Short value) throw(InvalidValue,SystemException);
  void insert_ushort(UShort value) throw(InvalidValue,SystemException);
  void insert_long(Long value) throw(InvalidValue,SystemException);
  void insert_ulong(ULong value) throw(InvalidValue,SystemException);
  void insert_float(Float value) throw(InvalidValue,SystemException);
  void insert_double(Double value) throw(InvalidValue,SystemException);
```

```
   void insert_string(const char* value) throw(InvalidValue,SystemException);
   void insert_reference(Object_ptr v) throw(InvalidValue,SystemException);
   void insert_typecode(TypeCode_ptr v) throw(InvalidValue,SystemException);
   void insert_any(const Any& value) throw(InvalidValue,SystemException);

   Boolean get_boolean() throw(TypeMismatch,SystemException);
   Octet get_octet() throw(TypeMismatch,SystemException);
   Char get_char() throw(TypeMismatch,SystemException);
   Short get_short() throw(TypeMismatch,SystemException);
   UShort get_ushort() throw(TypeMismatch,SystemException);
   Long get_long() throw(TypeMismatch,SystemException);
   ULong get_ulong() throw(TypeMismatch,SystemException);
   Float get_float() throw(TypeMismatch,SystemException);
   Double get_double() throw(TypeMismatch,SystemException);
   char* get_string() throw(TypeMismatch,SystemException);
   Object_ptr get_reference() throw(TypeMismatch,SystemException);
   TypeCode_ptr get_typecode() throw(TypeMismatch,SystemException);
   Any* get_any() throw(TypeMismatch,SystemException);

   static DynAny_ptr _duplicate(DynAny_ptr);
   static DynAny_ptr _narrow(DynAny_ptr);
   static DynAny_ptr _nil();
};

// DynFixed is not supported.

typedef DynEnum* DynEnum_ptr;
class DynEnum_var { ... };

class DynEnum :  public DynAny {
public:

   char* value_as_string();
   void value_as_string(const char* value);
   ULong value_as_ulong();
   void value_as_ulong(ULong value);

   static DynEnum_ptr _duplicate(DynEnum_ptr);
   static DynEnum_ptr _narrow(DynAny_ptr);
   static DynEnum_ptr _nil();
};

typedef char* FieldName;
typedef String_var FieldName_var;

struct NameValuePair {
  String_member id;
  Any value;
};

typedef _CORBA_ConstrType_Variable_Var<NameValuePair> NameValuePair_var;
typedef _CORBA_Unbounded_Sequence<NameValuePair > NameValuePairSeq;
```

```
typedef DynStruct* DynStruct_ptr;
class DynStruct_var { ... };

class DynStruct :  public DynAny {
public:

  char*  current_member_name();
  TCKind current_member_kind();
  NameValuePairSeq* get_members();
  void set_members(const NameValuePairSeq& NVSeqVal)
                    throw(InvalidSeq,SystemException);

  static DynStruct_ptr _duplicate(DynStruct_ptr);
  static DynStruct_ptr _narrow(DynAny_ptr);
  static DynStruct_ptr _nil();
};

typedef DynUnion* DynUnion_ptr;
class DynUnion_var { ... };

class DynUnion :  public DynAny {
public:

  Boolean set_as_default();
  void set_as_default(Boolean value);
  DynAny_ptr discriminator();
  TCKind discriminator_kind();
  DynAny_ptr member();
  char*  member_name();
  void member_name(const char* value);
  TCKind member_kind();

  static DynUnion_ptr _duplicate(DynUnion_ptr);
  static DynUnion_ptr _narrow(DynAny_ptr);
  static DynUnion_ptr _nil();
};

typedef _CORBA_Unbounded_Sequence<Any > AnySeq;

typedef DynSequence* DynSequence_ptr;
class DynSequence_var { ... };

class DynSequence :  public DynAny {
public:

  ULong length();
  void length (ULong value);
  AnySeq* get_elements();
  void set_elements(const AnySeq& value) throw(InvalidValue,SystemException);

  static DynSequence_ptr _duplicate(DynSequence_ptr);
  static DynSequence_ptr _narrow(DynAny_ptr);
  static DynSequence_ptr _nil();
```

```
};

typedef DynArray* DynArray_ptr;
class DynArray_var { ... };

class DynArray : public DynAny {
public:

  AnySeq* get_elements();
  void set_elements(const AnySeq& value) throw(InvalidValue,SystemException);

  static DynArray_ptr _duplicate(DynArray_ptr);
  static DynArray_ptr _narrow(DynAny_ptr);
  static DynArray_ptr _nil();
};
```

10.2   The DynAny Interface

10.2.1   Example: extract data values from an Any

If an **any** contains a value of one of the basic data types, its value can be extracted using
the pre-defined operators in the Any interface. When the value is a struct or other
non-basic types, one can use the DynAny interface to extract its constituent values.

In this section, we use a struct as an example to illustrate how the DynAny inter-
face can be used.

The example struct is as follows:

```
// IDL

struct exampleStruct1 {
  string s;
  double d;
  long   l;
};
```

To **create** a DynAny from an Any value, one uses the `create_dyn_any` method:

```
// C++

CORBA::ORB_ptr orb;   // orb initialised by CORBA::ORB_init.

Any v;
...        // Initialise v to contain a value of type exampleStruct1.

CORBA::DynAny_var dv = orb->create_dyn_any(v);
```

Like CORBA object and pseudo object references, a `DynAny_ptr` can be managed
by a _var type (`DynAny_var`) which will release the `DynAny_ptr` automatically when
the variable goes out of scope.

10.2.1.1 Iterate through the components

Once the DynAny object is created, we can use the DynAny interface to extract the individual components in `exampleStruct1`. The DynAny interface provides a number of functions to extract and insert component values. These functions are defined to operate on the component identified by the **current component** pointer.

A **current component** pointer is an internal state of a DynAny object. When a DynAny object is created, the pointer is initialised to point to the first component of the any value.

The pointer can be advanced to the next component with the `next()` operation. The function returns FALSE (0) if there are no more components. Otherwise it returns TRUE (1). When the any value in the DynAny object contains only one component, the `next()` operation always returns FALSE(0).

Another way of adjusting the pointer is the `seek()` operation. The function returns FALSE (0) if there is no component at the specified index. Otherwise it returns TRUE (1). The index value of the first component is zero. Therefore, a `seek(0)` call rewinds the pointer to the first component, this is also equivalent to a call to the `rewind()` operation.

For completeness, we should also mention here the `current_component()` operation. This operation causes the DynAny object to return a reference to another DynAny object that can be used to access the current component. It is possible that the current component pointer is not pointing to a valid component, for instance, the `next()` operation has been invoked and there is no more component. Under this circumstance, the `current_component()` operation returns a nil DynAny object reference[1]. For components which are just basic data types, calling `current_component()` is an overkill because we can just use the basic type extraction and insertion functions directly.

10.2.1.2 Extract basic type components

In our example, the component values can be extracted as follows:

```
CORBA::String_var s = dv->get_string();
CORBA::Double     d = dv->get_double();
CORBA::Long       l = dv->get_long();
```

Each get basic type operation has the side-effect of advancing the current component pointer. For instance:

```
CORBA::String_var s = dv->get_string();
```

is equivalent to:

```
CORBA::DynAny_var temp = dv->current_component();
CORBA::String_var s = temp->get_string();
dv->next();
```

---

[1]Testing a nil DynAny object with CORBA::is_nil() returns TRUE(1). The CORBA 2.2 specification does not specify what is the return value of this function when the current component pointer is invalid. To ensure portability, it is best to avoid calling `current_component()` under this condition.

The get operations ensure that the current component is of the same type as requested. Otherwise, the object throws a TypeMismatch exception. If the current component pointer is invalid when a get operation is called, the object also throws a TypeMismatch exception[2].

To repeatedly access the components, one can use the `rewind()` or `seek()` operation to manipulate the current component pointer. For instance, to access the `d` member in `exampleStruct1` directly:

```
dv->seek(1);            // position current component to member d.
CORBA::Double    d = dv->get_double();
```

10.2.1.3   Extract complex components

When a component is not one of the basic data types, it is not possible to extract its value using the get operations. Instead, a DynAny object has to be created from which the component is accessed.

Consider this example:

```
// IDL

struct exampleStruct2 {
  string m1;
  exampleStruct1 m2;
};
```

In order to extract the data members within `m2` (of type `exampleStruct1`), we use `current_component()` as follows:

```
// C++

CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.
Any v;
...         // Initialise v to contain a value of type exampleStruct2.

CORBA::DynAny_var dv = orb->create_dyn_any(v);

CORBA::String_var m1 = dv->get_string();  // extract member m1
CORBA::DynAny_var dm = dv->current_component(); // DynAny refer-
ence to m2
CORBA::String_var s = dm->get_string();   // m2.s
CORBA::Double    d = dm->get_double();   // m2.d
CORBA::Long      l = dm->get_long();     // m2.l
```

10.2.1.4   Clean-up

Now we finish off this example with a description on destroying DynAny objects. There are two points to remember:

---

[2]The CORBA 2.2 specification does not define the behavior of this error condition. To ensure portability, it is best to avoid calling the get operations when the current component pointer is known to be invalid.

1. A DynAny reference (`DynAny_ptr`) is like any CORBA object or psuedo object reference and should be handled in the same way. In particular, one has to call the `CORBA::release` operation to indicate that a DynAny reference will no longer be accessed. In the example, this is done automatically by `DynAny_var`.

2. A DynAny object and its references are separate entities, just as a CORBA object implementation and its object references are different entities. While `CORBA::release` will release any resource associated with a `DynAny_ptr`, one has to separately destroy the DynAny object to avoid any memory leak. This is done by calling the `destroy()` operation.

In the example, the DynAny object can be destroyed as follows:

```
// C++
...
CORBA::DynAny_var dv = orb->create_dyn_any(v);
...
dv->destroy();

// From now on, one should not invoke any operation in dv. Oth-
erwise the
// behavior is undefined.
```

10.2.2   Example: insert data values into an Any

Using the DynAny interface, one can create an Any value from scratch. In this example, we are going to create an Any containing the value of the `exampleStruct1` type.

Firstly, we have to create a DynAny to store the value using one of the `create_dyn` functions. Because `exampleStruct1` is a struct type, we use the `create_dyn_struct()` operation.

```
// C++

CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.

// create the TypeCode for exampleStruct.
StructMemberSeq tc_members;
tc_members.length(3);
tc_members[0].name = (const char*)"s";
tc_members[0].type = CORBA::TypeCode::_duplicate(CORBA::_tc_string);
tc_members[0].type_def = CORBA::IDLType::_nil();
tc_members[1].name = (const char*)"d";
tc_members[1].type = CORBA::TypeCode::_duplicate(CORBA::_tc_double);
tc_members[1].type_def = CORBA::IDLType::_nil();
tc_members[2].name = (const char*)"l";
tc_members[2].type = CORBA::TypeCode::_duplicate(CORBA::_tc_long);
tc_members[2].type_def = CORBA::IDLType::_nil();
CORBA::TypeCode_var tc = orb->create_struct_tc("IDL:exampleStruct1:1.0",
                                               "exampleStruct1",
```

```
                                            tc_members);

    // create the DynAny object to represent the any value
    CORBA::DynAny_var dv = orb->create_dyn_struct(tc);
```

10.2.2.1   Insert basic type components

Once the DynAny object is created, we can use the DynAny interface to insert the components. The DynAny interface provides a number of insert operations to insert basic types into the any value. In our example, the component values can be inserted as follows:

```
    CORBA::String_var s = (const char*)"Hello";
    CORBA::Double     d = 3.1416;
    CORBA::Long       l = 1;

    dv->insert_string(s);
    dv->insert_double(d);
    dv->insert_long(l);
```

Each insert basic type operation has the side-effect of advancing the current component pointer. For instance:

```
    dv->insert_string(s);
```

is equivalent to:

```
    CORBA::DynAny_var temp = dv->current_component();
    temp->insert_string(s);
    dv->next();
```

The insert operations ensure that the current component is of the same type as the inserted value. Otherwise, the object throws a InvalidValue exception. If the current component pointer is invalid when an insert operation is called, the object also throws a InvalidValue exception[3].

Sometimes, one may just want to modify one component in an Any value. For instance, one may just want to change the value of the double member in `exampleStruct1`. This can be done as follows:

```
    // C++

    CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.

    Any v;
    ...        // Initialise v to contain a value of type exampleStruct1.

    CORBA::Double d = 6.28;
```

---

[3]The CORBA 2.2 specification does not define the behavior of this error condition. To ensure portability, it is best to avoid calling the insert operations when the current component pointer is known to be invalid.

```
CORBA::DynAny_var dv = orb->create_dyn_any(v);

dv->seek(1);
dv->insert_double(d);     // Change the value of the member d.
```

Finally, the any value can be obtained from the DynAny object using the to_any() operation:

```
CORBA::Any_var v = dv->to_any();     // Obtain the any value.
```

10.2.2.2   Insert complex components

When a component is not one of the basic data types, it is not possible to insert its value using the insert operations. Instead, a DynAny object has to be created through which the component can be inserted.

In our example, one can insert component values into exampleStruct2 as follows:

```
// C++

CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.

CORBA::TypeCode_var tc;
// create the TypeCode for exampleStruct2.
...
// create the DynAny object to represent the any value
CORBA::DynAny_var dv = orb->create_dyn_struct(tc);

CORBA::String_var m1  = (const char*)"Greetings";
CORBA::String_var m2s = (const char*)"Hello";
CORBA::Double     m2d = 3.1416;
CORBA::Long       m2l = 1;

dv->insert_string(m1);   // insert member m1
CORBA::DynAny_var dm = dv->current_component(); // DynAny refer-
ence to m2
dm->insert_string(m2s);  // insert member m2.s
dm->insert_double(m2d);  // insert member m2.d
dm->insert_long(m2l);    // insert member m2.l

CORBA::Any_var v = dv->to_any();  // obtain the any value

dv->destroy();            // destroy the DynAny object.
                          // No operation should be invoked on dv
                          // from this point on ex-
cept CORBA::release.
```

In addition to the DynAny interface, a number of derived interfaces are defined. These interfaces are specialisation of the DynAny interface to facilitate the handling of any values containing non-basic types: struct, sequence, array, enum and union[4]. The next few sections will provide more details on these interfaces.

---

[4]In the CORBA 2.2 specification, the DynFixed interface is defined to handle the fixed data type. This is not supported in this implementation.

10.3 The DynStruct Interface

When a DynAny object is created through the `create_dyn_any()` operation and the any value contains a struct type, a DynStruct object is created. The DynAny reference returned can be narrowed to a DynStruct reference using the `CORBA::DynStruct::_narrow()` operation.

In the previous example, the components are extracted using the get operations. Alternatively, the DynStruct interface provides an addition operation (`get_members()`) to return all the components in a single call. The returned value is a sequence of name value pair. The member name is given in the name field and its value is returned as an Any value. For example, an alternative way to extract the components in the previous example is as follows:

```
// C++

CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.
Any v;
...        // Initialise v to contain a value of type exampleStruct1.
CORBA::DynAny_var dv = orb->create_dyn_any(v);

CORBA::DynStruct_var ds = CORBA::DynStruct::_narrow(dv);

CORBA::NameValuePairSeq* sq = ds->get_members();

char*       s;
CORBA::Double d;
CORBA::Long   l;

(*sq)[0].value >>= s;        // 1st element contains member s
(*sq)[1].value >>= d;        // 2nd element contains member d
(*sq)[2].value >>= l;        // 3rd element contains member l
```

Similarly, the DynStruct interface provides an addition operation (`set_members()`) to insert all the components in a single call. The following is an alternative way to insert the components of the type `exampleStruct1` into an Any value:

```
// C++

CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.

CORBA::TypeCode_var tc;
// create the TypeCode for exampleStruct1.
...
// create the DynAny object to represent the any value
CORBA::DynAny_var dv = orb->create_dyn_struct(tc);

CORBA::String_var s = (const char*)"Hello";
CORBA::Double    d = 3.1416;
CORBA::Long      l = 1;
```

```
CORBA::NameValuePairSeq sq;
sq.length(3);
sq[0].id = (const char*)"s";
sq[0].value <<= CORBA::Any::from_string(s,0);
                                 // 1st element contains member s
sq[1].id = (const char*)"d";
sq[1].value <<= d;               // 2nd element contains member d
sq[2].id = (const char*)"l";
sq[2].value <<= l;               // 3rd element contains member l

dv->set_members(sq);
```

Notice that the name-value pairs in the argument to set_members() must match the members of the struct exactly or the object would throw the InvalidSeq exception.

In addition to the current_component() operation, the DynStruct interface provides two operations: current_member_name() and current_member_kind(), to return information about the current component.

10.4   The DynSequence Interface

Like struct values, sequence values can be traversed using the operations introduced in section 10.2. The first sequence element can be accessed as the first DynAny component, the second sequence element as the second DynAny component and so on.

To extract component values from an Any containing a sequence, the length of the sequence can be obtained using the get length operation in the DynSequence interface. Here is an example to extract the components of a sequence of long:

```
// C++

CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.
Any v;
...         // Initialise v to contain a value of a se-
quence of long
CORBA::DynAny_var dv = orb->create_dyn_any(v);

CORBA::DynSequence_var ds = CORBA::DynSequence::_narrow(dv);
CORBA::ULong len = ds->length();     // ex-
tract the length of the sequence
CORBA::ULong index;
for (index = 0; index < len; index++) {
  CORBA::Long v = ds->get_long();
  cerr << "[" << index << "] = " << v << endl;
}
```

Conversely, the set length operation is provided to set the length of the sequence. Here is an example to insert the components of a sequence of long:

```
// C++

CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.
```

```
CORBA::TypeCode_var tc;
// create the TypeCode for a sequence of long.
...
// create the DynAny object to represent the any value
CORBA::DynSequence_var ds = orb->create_dyn_sequence(tc);

CORBA::ULong len = 3;

ds->length(len);              // set the length of the sequence

CORBA::ULong index;
for (index = 0; index < len; index++) {
  ds->insert_long(index);    // insert a sequence element
}
```

Similar to the DynStruct interface, the get_elements() operation is provided to return all the sequence elements and the set_elements() operation is provided to insert all the sequence elements.

10.5   The DynArray Interface

Array values are handled by the DynArray interface. The DynArray interface is the same as the DynSequence interface except that the former does not provide the set length and get length operations.

10.6   The DynEnum Interface

Enum values are handled by the DynEnum interface. A DynEnum object contains a single component which is the enum value. This value cannot be extracted or inserted using the get and insert operations of the DynAny interface. Instead, two pairs of operations are provided to handle this value.

The value_as_string operations allow the enum value to be extracted or inserted as a string. The value_as_ulong operations allow the enum value to be extracted or inserted as an unsigned long.

10.7   The DynUnion Interface

Union values are handled by the DynUnion interface. Unfortunately, the CORBA 2.2 specification does not define the DynUnion interface in suffcient details to nail down its intended usage[5]. In this section, we try to fill in the gaps and describe a sensible way to use the DynUnion interface. Where necessary, the semantics of the operations is clarified. It is possible that the behavior of this interface in another ORB is different from this implmentation. Where appropriate, we give warnings on usage that might cause problems with portability.

---

[5]This interface is currently an open issue with the ORB revision task force.

In relation to the current component pointer (10.2.1.1), a DynUnion object contains two components. The first component (with the index value equals 0) is the discriminator value, the second one is the member value. Therefore, one can use the `seek()` and `current_component()` operations to obtain a reference to the DynAny objects that handle the two components. However, it is better to use the operations defined in the DynUnion interface to manipulate these components as the semantics of the operations is easier to understand.

10.7.1   Three Categories of Union

Before we continue, it is important to understand that unions can be classified into the following categories:

1. One that has a default branch defined in the IDL. This will be called **explicit default union** in the rest of this section.

2. One that has no default branch and not all the possible values of the discriminator type are covered by the branch labels in the IDL. This will be called **implicit default union**.

3. One that has no default branch but all the possible values of the discriminator type are covered. This will be called **no default union**.

Of the three categories, the implicit default union is interesting because by definition if the discriminator value is not equal to any of the branch labels, the union has **no** member. That is, the union value consists solely of the discriminator value.

10.7.2   Example: extract data values from a union

10.7.2.1   Explicit default union

Consider a union of the following type:

```
// IDL

union exampleUnion1 switch(boolean) {
case TRUE: long l;
default:   double d;
};
```

The most straightforward way to extract the member value is as follows:

```
// C++

CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.

Any v;
...        // Initialise v to contain a value of type exampleUnion1.

CORBA::DynAny_var dv = orb->create_dyn_any(v);
CORBA::DynUnion_var du = CORBA::DynUnion::_narrow(dv);
```

```
CORBA::String_var di = du->member_name();
CORBA::DynAny_var dm = du->member();

if (strcmp((const char*)di,"l") == 0) {
  // branch label is TRUE
  CORBA::Long v = dm->get_long();
  cerr << "l = " << v << endl;
}

if (strcmp((const char*)di,"d") == 0) {
  // Is default branch
  CORBA::Double v = dm->get_double();
  cerr << "d = " << v << endl;
}
```

In the example, the operation member_name() is used to determine which branch the union has been instantiated. The operation member() is used to obtain a reference to the DynAny object that handles the member.

Alternatively, the branch can be determined by reading the discriminator value:

```
// C++

CORBA::DynAny_var di = du->discriminator();
CORBA::DynAny_var dm = du->member();

CORBA::Boolean di_v = di->get_boolean();

switch (di_v) {
case 1:
  CORBA::Long v = dm->get_long();
  cerr << "l = " << v << endl;
  break;
default:
  CORBA::Double v = dm->get_double();
  cerr << "d = " << v << endl;
}
```

The operation discriminator() is used to obtain the value of the discriminator. Finally, the third way to determine the branch is to test if the default is selected:

```
// C++

switch (dv->set_as_default()) {
case 1:
  CORBA::Double v = dm->get_double();
  cerr << "d = " << v << endl;
  break;
default:
  CORBA::Long v = dm->get_long();
  cerr << "l = " << v << endl;
}
```

The operation `set_as_default()` returns TRUE (1) if the discriminator has been assigned a valid default value.

10.7.2.2   Implicit default union

Consider a union of the following type:

```
// IDL

union exampleUnion2 switch(long) {
case 1: long l;
case 2: double d;
};
```

This example is similar to the previous one but there is no default branch. The description above also applies to this example. However, the discriminator may be set to neither 1 nor 2. Under this condition, the implicit default is selected and the union value contains the discriminator only!

When the discriminator contains an implicit default value, one might ask what is the value returned by the `member_name()` and `member()` operation. Since there is no member in the union value, omniORB2 returns a null string and a nil DynAny reference respectively. This behavior is not specified in the CORBA 2.2 specification. To ensure that your application is portable, it is best to avoid calling these operations when the DynUnion object might contain an implicit default value.

10.7.2.3   No default union

This is the last union category. For instance:

```
// IDL

union exampleUnion3 switch(boolean) {
case TRUE: long l;
case FALSE: double d;
};
```

In this example, all the possible values of the discriminator are used as union labels. There is no default branch. The only difference between this category and the explicit default union is that the `set_as_default()` operation always returns FALSE (0).

10.7.3   Example: insert data values into a union

Writing into a union involves selecting the union branch with the appropriate discriminator value and then writing the member value. There are three ways to set the discriminator value:

1. Use the `member_name()` write operation to specify the union branch by specifying the union member directly. This operation has the side effect of setting the discriminator to the label value of the branch.

2. Write the label value of a union branch into the DynAny object that handles the discriminator.

3. If the union has a default branch, either explicitly or implicitly, use the set_as_default() write operation to set the discriminator to a valid default value.

 The following example shows the three ways of writing into a union:

```
// C++

CORBA::ORB_ptr orb;  // orb initialised by CORBA::ORB_init.

CORBA::TypeCode_var tc;
// create the TypeCode for exampleUnion1.
...
// create the DynAny object to represent the any value
CORBA::DynUnion_var dv = orb->create_dyn_union(tc);

CORBA::Any_var v;
DynAny_ptr dm;

// Use member_name to select the union branch
dv->member_name("l");
dm = dv->member();
dm->insert_long(10);
v = dv->to_any();            // transfer to an Any
CORBA::release(dm);

// Setting the discriminator value to select the union branch
CORBA::DynAny_var di = dv->discriminator();
di->insert_boolean(1);     // set discriminator to label TRUE
dm = dv->member();
dm->insert_long(20);
v = dv->to_any();            // transfer to an Any
CORBA::release(dm);

// Use set_as_default to select the default union branch
dv->set_as_default(1);
dm = dv->member();
dm->insert_double(3.14);
v = dv->to_any();            // transfer to an Any
CORBA::release(dm);

dv->destroy();
```

10.7.3.1   Ambiguous usage

1. When the discriminator is set to a different value, a different member branch is selected. Suppose the application has previously obtained a DynAny reference to a union member when it changes the discriminator value. As a result of the

value change, the union is now instantiated to another union branch, i.e. a call to the `member()` operation will now return a reference to a different DynAny object. If the application continues to access the DynAny object of the old union member, the behavior of the ORB under this condition is not defined by the CORBA 2.2 specification. With omniORB2, the DynAny object of the old union member is detached from the union when a new union branch is selected. Therefore reading or writing this object will not have any relation to the current value of the union. To avoid this ambiguity, the reference to the old union member should be released before a different union branch is selected.

2. The write operation `set_as_default()` takes a boolean argument. It is ambiguous to call this function with the argument set to FALSE (0). With omniORB2, such a call will be silently ignored.

3. It is also ambiguous to pass the value TRUE (1) to the `set_as_default()` operation when the union is a no default union (10.7.1). With omniORB2, such a call will be silently ignored.

4. When the discriminator value is not set, calling the `member()` operation is ambiguous. With omniORB2, such a call will return a nil DynAny reference. Similarly, a call to the `member_kind()` operation under this condition will return `tk_null`.

To ensure portability, it is best to avoid using the DynUnion interface and not to rely on the ORB to behave as omniORB2 does under these ambiguous conditions.

10.8   Duplicate DynAny References

Like any CORBA object and psuedo object references, a DynAny reference can be duplicated using the `_duplicate()` operations. When an application has obtained multiple DynAny references to the same DynAny object, it should be noted that a change made to the object by invoking on one reference is also visible through the other references. In particular, if a call through one reference has caused the current component pointer to be changed, subsequent calls through other references will operate on the new current component pointer.

10.9   Other Operations

The following is a short summary of the other operations in the DynAny interface which have not been covered in previous sections:

`assign()` initialises a DynAny object with another DynAny object. The two objects must have the same typecode.

`from_any()` initialises a DynAny object from the value in an any. The typecode in the two objects must be the same.

`copy()` creates a new DynAny object whose value is a deep copy of the current object.

`type()` returns the typecode associated with the DynAny object.

Chapter 11

The Dynamic Invocation Interface

The Dynamic Invocation Interface (or DII) allows applications to invoke operations on CORBA objects about which they have no static information. That is to say the application has not been linked with stub code which performs the remote operation invocation. Thus using the DII applications may invoke operations on *any* CORBA object, possibly determining the object's interface dynamically by using an Interface Repository.

This chapter presents an overview of the Dynamic Invocation Interface. An toy example use of the DII can be found in the omniORB2 distribution in the `<top>/src/examples/dii` directory. The DII makes extensive use of the type Any, so ensure that you have read chapter 9. For more information refer to the Dynamic Invocation Interface and C++ Mapping sections of the CORBA 2 specification [OMG99a].

11.1 Overview

To invoke an operation on a CORBA object an application needs an object reference, the name of the operation and a list of the parameters. In addition the application must know whether the operation is one-way, what user-defined exceptions it may throw, any user-context strings which must be supplied, a 'context' to take these values from and the type of the returned value. This information is given by the IDL interface declaration, and so is normally made available to the application via the stub code. In the DII this information is encapsulated in the `CORBA::Request` pseudo-object.

To perform an operation invocation the application must obtain an instance of a `Request` object, supply the information listed above and call one of the methods to actually make the invocation. If the invocation causes an exception to be thrown then this may be retrieved and inspected, or the return value on success.

11.2 Pseudo Objects

The DII defines a number of psuedo-object types, all defined in the CORBA namespace. These objects behave in many ways like CORBA objects. They should only be accessed by reference (through `foo_ptr` or `foo_var`), may not be instantiated directly

and should be released by calling `CORBA::release()`[1]. A nil reference should only
be represented by `foo::_nil()`.

These pseudo objects, although defined in pseudo-IDL in the specification do not
follow the normal mapping for CORBA objects. In particular the memory manage-
ment rules are different - see the CORBA 2 specification [OMG99a] for more details.
New instances of these objects may only be created by the ORB. A number of methods
are defined in `CORBA::ORB` to do this.

11.2.1   Request

A `Request` encapsulates a single operation invocation. It may *not* be re-used - even
for another call with the same arguments.

```
class Request {
public:
  virtual Object_ptr        target() const;
  virtual const char*       operation() const;
  virtual NVList_ptr        arguments();
  virtual NamedValue_ptr    result();
  virtual Environment_ptr   env();
  virtual ExceptionList_ptr exceptions();
  virtual ContextList_ptr   contexts();
  virtual Context_ptr       ctxt() const;
  virtual void              ctx(Context_ptr);

  virtual Any& add_in_arg();
  virtual Any& add_in_arg(const char* name);
  virtual Any& add_inout_arg();
  virtual Any& add_inout_arg(const char* name);
  virtual Any& add_out_arg();
  virtual Any& add_out_arg(const char* name);

  virtual void set_return_type(TypeCode_ptr tc);
  virtual Any& return_value();

  virtual Status  invoke();
  virtual Status  send_oneway();
  virtual Status  send_deferred();
  virtual Status  get_response();
  virtual Boolean poll_response();

  static Request_ptr _duplicate(Request_ptr);
  static Request_ptr _nil();
};
```

11.2.2   NamedValue

A pair consisting of a string and a value - encapsulated in an Any. The name is op-
tional. This type is used to encapsulate parameters and returned values.

---

[1]if not managed by a _var type.

```
class NamedValue {
public:
  virtual const char* name() const;
  // Retains ownership of return value.

  virtual Any* value() const;
  // Retains ownership of return value.

  virtual Flags flags() const;

  static NamedValue_ptr _duplicate(NamedValue_ptr);
  static NamedValue_ptr _nil();
};
```

11.2.3 NVList

A list of `NamedValue` objects.

```
class NVList {
public:
  virtual ULong count() const;
  virtual NamedValue_ptr add(Flags);
  virtual NamedValue_ptr add_item(const char*, Flags);
  virtual NamedValue_ptr add_value(const char*, const Any&, Flags);
  virtual NamedValue_ptr add_item_consume(char*,Flags);
  virtual NamedValue_ptr add_value_consume(char*, Any*, Flags);
  virtual NamedValue_ptr item(ULong index);
  virtual Status remove (ULong);

  static NVList_ptr _duplicate(NVList_ptr);
  static NVList_ptr _nil();
};
```

11.2.4 Context

Represents a set of context strings. User contexts are not supported by the omniidl2 IDL compiler - and so cannot be used with statically defined operations. However they are supported in the DII.

```
class Context {
public:
  virtual const char* context_name() const;
  virtual CORBA::Context_ptr parent() const;
  virtual CORBA::Status create_child(const char*, Context_out);
  virtual CORBA::Status set_one_value(const char*, const CORBA::Any&);
  virtual CORBA::Status set_values(CORBA::NVList_ptr);
  virtual CORBA::Status delete_values(const char*);
  virtual CORBA::Status get_values(const char* start_scope,
                                   CORBA::Flags op_flags,
                                   const char* pattern,
                                   CORBA::NVList_out values);
  // Throws BAD_CONTEXT if <start_scope> is not found.
```

```
  // Returns a nil NVList in <values> if no matches are found.

  static Context_ptr _duplicate(Context_ptr);
  static Context_ptr _nil();
};
```

11.2.5   ContextList

A `ContextList` is a list of strings, and is used to specify which strings from the 'context' should be sent with an operation.

```
class ContextList {
public:
  virtual ULong count() const;
  virtual void add(const char* ctxt);
  virtual void add_consume(char* ctxt);
  // consumes ctxt

  virtual const char* item(ULong index);
  // retains ownership of return value

  virtual Status remove(ULong index);

  static ContextList_ptr _duplicate(ContextList_ptr);
  static ContextList_ptr _nil();
};
```

11.2.6   ExceptionList

`ExceptionList`s contain a list of TypeCodes - and are used to specify which user-defined exceptions an operation may throw.

```
class ExceptionList {
public:
  virtual ULong count() const;
  virtual void add(TypeCode_ptr tc);
  virtual void add_consume(TypeCode_ptr tc);
  // Consumes <tc>.

  virtual TypeCode_ptr item(ULong index);
  // Retains ownership of return value.

  virtual Status remove(ULong index);

  static ExceptionList_ptr _duplicate(ExceptionList_ptr);
  static ExceptionList_ptr _nil();
};
```

11.2.7   UnknownUserException

When a user-defined exception is thrown by an operation it is unmarshalled into a value of type Any. This is encapsulated in an `UnknownUserException`. This type

follows all the usual rules for user-defined exceptions - it is not a pseudo object, and its resources may be released by using `delete`.

```
class UnknownUserException : public UserException {
public:
  UnknownUserException(Any* ex);
  // Consumes <ex> which MUST be a UserException.

  virtual ~UnknownUserException();

  Any& exception();

  virtual void _raise();
  static const UnknownUserException* _downcast(const Exception*);
  static UnknownUserException* _downcast(Exception*);
  static UnknownUserException* _narrow(Exception*);
  // _narrow is a deprecated function from CORBA 2.2,
  // use _downcast instead.
};
```

11.2.8   Environment

An `Environment` is used to hold an instance of a system exception or an `UnknownUserException`.

```
class Environment {
  virtual void exception(Exception*);
  virtual Exception* exception() const;
  virtual void clear();

  static Environment_ptr _duplicate(Environment_ptr);
  static Environment_ptr _nil();
};
```

11.3   Creating Requests

`CORBA::Object` defines three methods which may be used to create a `Request` object which may be used to perform a single operation invocation on that object:

```
class Object {
  ...
  Status _create_request(Context_ptr ctx,
                         const char* operation,
                         NVList_ptr arg_list,
                         NamedValue_ptr result,
                         Request_out request,
                         Flags req_flags);

  Status _create_request(Context_ptr ctx,
                         const char* operation,
                         NVList_ptr arg_list,
```

```
                    NamedValue_ptr result,
                    ExceptionList_ptr exceptions,
                    ContextList_ptr ctxlist,
                    Request_out request,
                    Flags req_flags);

  Request_ptr _request(const char* operation);
  ...
};
```

operation is the name of the operation - which is the same as the name given in IDL. To access attributes the name should be prefixed by _get_ or _set_

In the first two cases above the list of parameters may be supplied. If the parameters are not supplied in these cases, or _request() is used then the parameters (if any) may be specified using the add_*_arg() methods on the Request. You must use one method or the other - not a mixture of the two. For *in/inout* arguments the value must be initialised, for *out* arguments only the type need be given. Similarly the type of the result may be specified by passing a NamedValue which contains an Any which has been initialised to contain a value of that type, or it may be specified using the set_return_type() method of Request.

When using _create_request(), the management of any pseudo-object references passed in remains the responsibility of the application. That is, the values are not consumed - and must be released using CORBA::release(). The CORBA specification is unclear about when these values may be released, so to be sure of portability do not release them until after the request has been released. Values which are not needed need not be supplied - so if no parameters are specified then it defaults to an empty parameter list. If no result type is specified then it defaults to void. A Context need only be given if a non-empty ContextList is specified. The req_flags argument is not used in the C++ mapping.

11.3.1   Examples

An operation might be specified in IDL as:

```
short anOpn(in string a);
```

An operation invocation may be created as follows:

```
CORBA::ORB_var orb = CORBA::ORB_init(argc, argv, "omniORB2");
...
CORBA::NVList_var args;
orb->create_list(1, args);
*(args->add(CORBA::ARG_IN)->value()) <<= (const char*) "Hello World!";

CORBA::NamedValue_var result;
orb->create_named_value(result);
result->value()->replace(CORBA::_tc_short, 0);

CORBA::Request_var req = obj->_create_request(CORBA::Context::_nil(),
                                   "anOpn", args, result, 0);
```

or alternatively and much more concisely:

```
CORBA::Request_var req = obj->_request("anOpn");
req->add_in_arg() <<= (const char*) "Hello World!";
req->set_return_type(CORBA::_tc_short);
```

11.4   Invoking Operations

Once the `Request` object has been properly constructed the operation may be invoked by calling one of the following methods on the request object:

invoke()   blocks until the request has completed. The application should then test to see if an exception was raised. Since the CORBA spec is not clear about whether or not system exceptions should be thrown from this method, a runtime configuration variable is supplied so that you can specify the behavior:

```
namespace omniORB {
  ...
  CORBA::Boolean diiThrowsSysExceptions;
  ...
};
```

If this is FALSE, and the application should call the `env()` method of the request to retrieve an exception (it returns 0 (nil) if no exception was generated). If it is TRUE then system exceptions will be thrown out of `invoke()`. User-defined exceptions are always passed via `env()`, which will return a pointer to a `CORBA::UnknownUserException`. The application can determine which type of exception was returned by `env()` by calling the `_narrow()` method defined for each exception type.

WARNING!! In pre-omniORB 2.8.0 releases, the default value of `diiThrowsSysExceptions` is FALSE. From omniORB 2.8.0 onwards, the default value is TRUE.

After determining that no exception was thrown the application may retrieve any returned values by calling `return_value()` and `arguments()`.

send_oneway()   has the same semantics as a *oneway* IDL operation. It is important to note that oneway operations have at-most-once semantics, and it is not guaranteed that they will not block. Any operation may be invoked 'oneway' using the DII, even if it was not declared as 'oneway' in IDL. A system exception may be generated, in which case it will either be thrown or may be retrieved using `env()` depending on `diiThrowsSysExceptions` as above.

send_deferred()   initiates the invocation, and then returns without waiting for the result. At some point in the future the application must retrieve the result of the operation - but other than testing for completion of the operation the application must not call any of the request's methods in the meantime.

- `get_response()` blocks until the reply is received.

- `poll_response()` returns TRUE if the reply has been received, and FALSE if not. It does not block.

Once `poll_response()` has returned TRUE, or `get_response()` has been called and returned, the application may test for an exception and retrieve returned values as above. If `diiThrowsSysExceptions` is true, then a system exception may be thrown from `get_response()`. From omniORB 2.8.0 onwards, `poll_response()` will raise a system exception if one has occured during the invocation. Previously, `poll_response()` will not raise an exception, so if polling, the application must also call another method to give the request an opportunity to raise the exception. This can be one of the methods to retrieve values from the request, or `get_response()`.

11.5   Multiple Requests

The following methods are provided by the ORB to enable multiple requests to be invoked asynchronously.

```
namespace CORBA {
  ...
  class ORB {
  public:
    ...
    Status send_multiple_requests_oneway(const RequestSeq&);
    Status send_multiple_requests_deferred(const RequestSeq&);
    Boolean poll_next_response();
    Status get_next_response(Request_out);
    ...
  };
  ...
};
```

send_multiple_requests_oneway()   is used to invoke a number of oneway requests. An attempt will be made to invoke each of the requests, even if one or more of the early requests fails. The application may check for failure of any of the requests by testing the request's `env()` method. System exceptions are never raised by this method.

send_multiple_requests_deferred()   will initiate an invocation of each of the given requests, and return without waiting for the reply. At some point in the future the application must retrieve the reply by calling `get_next_response()`, which returns a completed request. If no requests have yet completed it will block. This method never throws exceptions - the request's `env()` method must be used to determine if an exception was generated. If not then any returned values may then be queried.

   `poll_next_response()` returns TRUE if there are any completed requests, and FALSE otherwise, without blocking. If this returns true then the next call to `get_next_response()` will not block. However, if another thread may also be calling `get_next_response()` then it could retrieve the completed message first - in which case this thread might block.

There are no guarantee as to the order in which replies will be received. If multiple threads are using this interface then it is not even guaranteed that a thread will receive replies to the requests it sent. Any thread may receive replies to requests sent by any other thread. It is legal to call `get_next_response()` even if no requests have yet been invoked - in which case the calling thread blocks until another thread invokes a request and the reply is received.

Chapter 12

The Dynamic Skeleton Interface

The Dynamic Skeleton Interface (or DSI) allows applications to provide implementations of the operations on CORBA objects without static knowledge of the object's interface. It is the server-side equivalent of the Dynamic Invocation Interface.

This chapter presents the Dynamic Skeleton Interface and explains how to use it. An toy example use of the DSI can be found in the omniORB2 distribution in the `<top>/src/examples/dsi` directory. For further information refer to the Dynamic Skeleton Interface and C++ Mapping sections of the CORBA 2 specification [OMG99a].

The DSI interface has changed in CORBA 2.3. The implementation described below conforms to CORBA 2.1 or 2.2.

12.1   Overview

When an ORB receives an invocation request, the information includes the object reference and the name of the operation. Typically this information is used by the ORB to select an instance of an object and call into the implementation of the operation (which knows how to unmarshal the parameters etc.). The Dynamic Skeleton Interface however makes this information directly available to the application - so that it can implement the operation (or pass it on to another server) without static knowledge of the interface. In fact it is not even necessary for the server to always implement the same interface on any particular object!

To provide an implementation for one or more objects an application must subclass `DynamicImplementation` and override the method `invoke()`. An instance of this class is registered with the BOA and is assigned an object reference (see below). When the ORB receives a request for that object the `invoke()` method is called and will be passed a `ServerRequest` object which provides:

- the operation name

- context strings

- access to the parameters

- a way to set the returned values

119

- a way to throw user-defined exceptions.

12.2   DSI Types

12.2.1   DynamicImplementation

This class must be sub-classed by the application to provide an implementation for
DSI objects. The method `invoke()` will be called for each operation invocation.

```
namespace CORBA {
  ...
  class BOA {
    ...
    class DynamicImplementation {
    public:
      DynamicImplementation();
      virtual ~DynamicImplementation();

      virtual void invoke(ServerRequest_ptr request,
                          Environment& env) throw() = 0;

    protected:
      Object_ptr _this();
      // Must only be called from within invoke(). Caller must release
      // the reference returned.

      BOA_ptr _boa();
      // Must only be called from within invoke(). Caller must NOT
      // release the reference returned.
    };
    ...
  };
  ...
};
```

12.2.2   ServerRequest

A `ServerRequest` object provides the interface between a dynamic implementation
and the ORB.

```
namespace CORBA {
  ...
  class ServerRequest {
  public:
    virtual const char*      op_name();
    virtual OperationDef_ptr op_def();
    virtual Context_ptr      ctx();
    virtual void             params(NVList_ptr parameters);
    virtual void             result(Any* value);
    virtual void             exception(Any* value);
```

```
     static ServerRequest_ptr _duplicate(ServerRequest_ptr);
     static ServerRequest_ptr _nil();
  };
  ...
};
```

12.3 Creating Dynamic Implementations

The application must override the `invoke()` method of `DynamicImplementation` to provide an implementation for DSI objects. This method must behave as follows:

- It may be called concurrently by multiple threads of execution, and so must be thread-safe.

- It may not throw any exceptions. User-defined exceptions are passed in a value of type Any and given to `ServerRequest::exception()`. In pre-omniORB 2.8.0 releases, system exceptions may be passed via the `env` parameter. From omniORB 2.8.0 onwards, system exceptions can be inserted into an Any and given to `ServerRequest::exception()` in the same way as user-defined exceptions. Since the methods provided by `ServerRequest` may throw system exceptions, the application must catch any such exception, passed it via `ServerRequest::exception()`.

- The operations on the `ServerRequest` object must be carried out in the correct order, as described below.

12.3.1 Operations on the ServerRequest

`op_name()` will return the name of the operation, and may be called at any time. For attribute access the operation name is the IDL name of the attribute, prefixed by `_get_` or `_set_` If the operation name is not recognised a `CORBA::BAD_OPERATION` exception should be passed back through `env`. This will allow the ORB to then see if it is one of the standard object operations.

Firstly `params()` must be called passing a `CORBA::NVList`[1] which must be initialised to contain the type and mode of the parameters. The ORB consumes this value and will release it when the operation is complete. At this point any *in/inout* arguments will be unmarshalled, and when this operation returns their values will be in the `NVList`. The application may set the value of *inout/out* arguments by modifying this parameter list.

If the operation has user-context information, then `ctx()` must be called after `params()` to retrieve it.

`result()` must then be called exactly once if the operation has a non-void return value (unless an exception is thrown). The value passed should be an Any allocated with `new`, and will be freed by the ORB.

At any point in the above sequence `exception()` may be called to set a user-defined exception or a system exception. If this happens then no further operations

---

[1]obtained by calling `CORBA::ORB::create_list()`

should be invoked on the `ServerRequest` object, and the `invoke()` method should return.

Within the `invoke()` method `_this()` and `_boa()` may be called to obtain the object reference and BOA reference respectively. These methods may not be used at any other time.

12.4   Registering Dynamic Objects

To use a `DynamicImplementation` a CORBA object must be created and associated with the implementation. The way in which this is done is not defined by the CORBA 2.0 specification, so the following method is omniORB2 specific:

```
namespace CORBA {
  ...
  class BOA {
    ...
    Object_ptr create_dynamic_object(DynamicImplementation_ptr dir,
                                      const char* intfRepoId);
    ...
  };
  ...
};
```

Ownership of the `DynamicImplementation` object is taken over by the ORB, and it will be deleted when associated the object is destroyed. The returned object may then be entered into the object table with a call to `BOA::obj_is_ready()` as usual, and will then start accepting operation invocations.

For some applications it will not be possible to register all DSI objects in advance of invocations arriving. In this case DSI objects can be created on demand in the same way as normal objects - see section 5.6.

12.5   Example

This implementation of `DynamicImplementation::invoke()` is taken from an example which can be found in the omniORB2 distribution. The operation "echoString" is declared in IDL as:

```
string echoString(in string mesg);
```

Here is the Dynamic Implementation Routine:

```
void
MyDynImpl::invoke(CORBA::ServerRequest_ptr re-
quest, CORBA::Environment& env)
  throw()
{
  try {
    if( strcmp(request->op_name(), "echoString") )
      throw CORBA::BAD_OPERATION(0, CORBA::COMPLETED_NO);
```

```
    CORBA::NVList_ptr args;
    orb->create_list(0, args);
    CORBA::Any a;
    a.replace(CORBA::_tc_string, 0);
    args->add_value("", a, CORBA::ARG_IN);

    request->params(args);

    CORBA::Any& input_any = *(args->item(0)->value());
    CORBA::String_var input;
    input_any >>= input.out();

    CORBA::Any* result = new CORBA::Any();
    *result <<= CORBA::Any::from_string(input._retn(), 0);
    request->result(result);
  }
  catch(CORBA::Exception& ex) {
    CORBA::Any* v = new CORBA::Any;
    ::operator<<=(*v,ex);
    request->exception(v);
  }
  // In pre-omniORB 2.8.0, one has to do this:
  //   catch(CORBA::SystemException& ex){
  //   env.exception(CORBA::Exception::_duplicate(&ex));
  //   }
  catch(...){
    cout << "echo_dsiimpl: MyDynImpl::invoke - caught an"
         " unknown exception." << endl;
    env.exception(new CORBA::UNKNOWN(0, CORBA::COMPLETED_NO));
  }
}
```

Appendix A

hosts_access(5)

DESCRIPTION

This manual page describes a simple access control language that is based on client (host name/address, user name), and server (process name, host name/address) patterns. Examples are given at the end. The impatient reader is encouraged to skip to the EXAMPLES section for a quick introduction.

   An extended version of the access control language is described in the hosts_options(5) document. The extensions are turned on at program build time by building with -DPROCESS_OPTIONS.

   In the following text, *daemon* is the the process name of a network daemon process, and *client* is the name and/or address of a host requesting service. Network daemon process names are specified in the inetd configuration file.

ACCESS CONTROL FILES

The access control software consults two files. The search stops at the first match:

- Access will be granted when a (daemon,client) pair matches an entry in the `/etc/hosts.allow` file.

- Otherwise, access will be denied when a (daemon,client) pair matches an entry in the `/etc/hosts.deny` file.

- Otherwise, access will be granted.

   A non-existing access control file is treated as if it were an empty file. Thus, access control can be turned off by providing no access control files.

ACCESS CONTROL RULES

Each access control file consists of zero or more lines of text. These lines are processed in order of appearance. The search terminates when a match is found.

- A newline character is ignored when it is preceded by a backslash character. This permits you to break up long lines so that they are easier to edit.

- Blank lines or lines that begin with a # character are ignored. This permits you to insert comments and whitespace so that the tables are easier to read.

- All other lines should satisfy the following format, things between [] being optional: `daemon_list  :   client_list [  :   shell_command ]`

`daemon_list` is a list of one or more daemon process names (argv[0] values) or wildcards (see below).

`client_list` is a list of one or more host names, host addresses, patterns or wildcards (see below) that will be matched against the client host name or address.

The more complex forms `daemon@host` and `user@host` are explained in the sections on server endpoint patterns and on client username lookups, respectively.

List elements should be separated by blanks and/or commas.

With the exception of NIS (YP) netgroup lookups, all access control checks are case insensitive.

PATTERNS

The access control language implements the following patterns:

- A string that begins with a . character. A host name is matched if the last components of its name match the specified pattern. For example, the pattern `.tue.nl` matches the host name `wzv.win.tue.nl`.

- A string that ends with a . character. A host address is matched if its first numeric fields match the given string. For example, the pattern `131.155.` matches the address of (almost) every host on the Eindhoven University network (`131.155.x.x`).

- A string that begins with an character is treated as an NIS (formerly YP) netgroup name. A host name is matched if it is a host member of the specified netgroup. Netgroup matches are not supported for daemon process names or for client user names.

- An expression of the form `n.n.n.n/m.m.m.m` is interpreted as a "net/mask" pair. A host address is matched if "net" is equal to the bitwise AND of the address and the "mask". For example, the net/mask pattern `131.155.72.0/255.255.254.0` matches every address in the range `131.155.72.0` through `131.155.73.255`.

WILDCARDS

The access control language supports explicit wildcards:

`ALL` The universal wildcard, always matches.

`LOCAL` Matches any host whose name does not contain a dot character.

UNKNOWN Matches any user whose name is unknown, and matches any host whose name or address are unknown. This pattern should be used with care: host names may be unavailable due to temporary name server problems. A network address will be unavailable when the software cannot figure out what type of network it is talking to.

KNOWN Matches any user whose name is known, and matches any host whose name and address are known. This pattern should be used with care: host names may be unavailable due to temporary name server problems. A network address will be unavailable when the software cannot figure out what type of network it is talking to.

PARANOID Matches any host whose name does not match its address. When tcpd is built with -DPARANOID (default mode), it drops requests from such clients even before looking at the access control tables. Build without -DPARANOID when you want more control over such requests.

OPERATORS

EXCEPT Intended use is of the form: `list_1 EXCEPT list_2`; this construct matches anything that matches `list_1` unless it matches `list_2`. The EXCEPT operator can be used in `daemon_lists` and in `client_lists`. The EXCEPT operator can be nested: if the control language would permit the use of parentheses, `a EXCEPT b EXCEPT c` would parse as `(a EXCEPT (b EXCEPT c))`.

SHELL COMMANDS

If the first-matched access control rule contains a shell command, that command is subjected to `%<letter>` substitutions (see next section). The result is executed by a /bin/sh child process with standard input, output and error connected to /dev/null. Specify an `&` at the end of the command if you do not want to wait until it has completed.

Shell commands should not rely on the PATH setting of the inetd. Instead, they should use absolute path names, or they should begin with an explicit PATH=whatever statement.

The hosts_options(5) document describes an alternative language that uses the shell command field in a different and incompatible way.

% EXPANSIONS

The following expansions are available within shell commands:

%a (%A) The client (server) host address.

%c Client information: user@host, user@address, a host name, or just an address, depending on how much information is available.

%d The daemon process name (argv[0] value).

`%h` (`%H`) The client (server) host name or address, if the host name is unavailable.

`%n` (`%N`) The client (server) host name (or "unknown" or "paranoid").

    `%p` The daemon process id.

    `%s` Server information: daemon@host, daemon@address, or just a daemon name, depending on how much information is available.

    `%u` The client user name (or "unknown").

    `%%` Expands to a single `%` character.

Characters in % expansions that may confuse the shell are replaced by underscores.

SERVER ENDPOINT PATTERNS

In order to distinguish clients by the network address that they connect to, use patterns of the form:

    `process_name@host_pattern :  client_list ...`

Patterns like these can be used when the machine has different internet addresses with different internet hostnames. Service providers can use this facility to offer FTP, GOPHER or WWW archives with internet names that may even belong to different organisations. See also the "twist" option in the hosts_options(5) document. Some systems (Solaris, FreeBSD) can have more than one internet address on one physical interface; with other systems you may have to resort to SLIP or PPP pseudo interfaces that live in a dedicated network address space. .sp The `host_pattern` obeys the same syntax rules as host names and addresses in `client_list` context. Usually, server endpoint information is available only with connection-oriented services.

CLIENT USERNAME LOOKUP

When the client host supports the RFC 931 protocol or one of its descendants (TAP, IDENT, RFC 1413) the wrapper programs can retrieve additional information about the owner of a connection. Client username information, when available, is logged together with the client host name, and can be used to match patterns like:

    `daemon_list :  ...  user_pattern@host_pattern ...`

The daemon wrappers can be configured at compile time to perform rule-driven username lookups (default) or to always interrogate the client host. In the case of rule-driven username lookups, the above rule would cause username lookup only when both the `daemon_list` and the `host_pattern` match.

A user pattern has the same syntax as a daemon process pattern, so the same wildcards apply (netgroup membership is not supported). One should not get carried away with username lookups, though.

- The client username information cannot be trusted when it is needed most, i.e. when the client system has been compromised. In general, ALL and (UN)KNOWN are the only user name patterns that make sense.

- Username lookups are possible only with TCP-based services, and only when the client host runs a suitable daemon; in all other cases the result is "unknown".

- A well-known UNIX kernel bug may cause loss of service when username lookups are blocked by a firewall. The wrapper README document describes a procedure to find out if your kernel has this bug.

- Username lookups may cause noticeable delays for non-UNIX users. The default timeout for username lookups is 10 seconds: too short to cope with slow networks, but long enough to irritate PC users.

Selective username lookups can alleviate the last problem. For example, a rule like:

```
daemon_list :   @pcnetgroup ALL@ALL
```

would match members of the pc netgroup without doing username lookups, but would perform username lookups with all other systems.

DETECTING ADDRESS SPOOFING ATTACKS

A flaw in the sequence number generator of many TCP/IP implementations allows intruders to easily impersonate trusted hosts and to break in via, for example, the remote shell service. The IDENT (RFC931 etc.) service can be used to detect such and other host address spoofing attacks.

Before accepting a client request, the wrappers can use the IDENT service to find out that the client did not send the request at all. When the client host provides IDENT service, a negative IDENT lookup result (the client matches `UNKNOWN@host`) is strong evidence of a host spoofing attack.

A positive IDENT lookup result (the client matches `KNOWN@host`) is less trustworthy. It is possible for an intruder to spoof both the client connection and the IDENT lookup, although doing so is much harder than spoofing just a client connection. It may also be that the client's IDENT server is lying.

Note: IDENT lookups don't work with UDP services.

EXAMPLES

The language is flexible enough that different types of access control policy can be expressed with a minimum of fuss. Although the language uses two access control tables, the most common policies can be implemented with one of the tables being trivial or even empty.

When reading the examples below it is important to realise that the allow table is scanned before the deny table, that the search terminates when a match is found, and that access is granted when no match is found at all.

The examples use host and domain names. They can be improved by including address and/or network/netmask information, to reduce the impact of temporary name server lookup failures.

MOSTLY CLOSED

In this case, access is denied by default. Only explicitly authorised hosts are permitted access.

The default policy (no access) is implemented with a trivial deny file:

```
/etc/hosts.deny:
    ALL: ALL
```

This denies all service to all hosts, unless they are permitted access by entries in the allow file.

The explicitly authorised hosts are listed in the allow file. For example:

```
/etc/hosts.allow:
   ALL: LOCAL @some_netgroup
   ALL: .foobar.edu EXCEPT terminalserver.foobar.edu
```

The first rule permits access from hosts in the local domain (no . in the host name) and from members of the `some_netgroup` netgroup. The second rule permits access from all hosts in the `foobar.edu` domain (notice the leading dot), with the exception of `terminalserver.foobar.edu`.

MOSTLY OPEN

Here, access is granted by default; only explicitly specified hosts are refused service.

The default policy (access granted) makes the allow file redundant so that it can be omitted. The explicitly non-authorised hosts are listed in the deny file. For example:

```
/etc/hosts.deny:
   ALL: some.host.name, .some.domain
   ALL EXCEPT in.fingerd: other.host.name, .other.domain
```

The first rule denies some hosts and domains all services; the second rule still permits finger requests from other hosts and domains.

BOOBY TRAPS

The next example permits tftp requests from hosts in the local domain (notice the leading dot). Requests from any other hosts are denied. Instead of the requested file, a finger probe is sent to the offending host. The result is mailed to the superuser.

```
/etc/hosts.allow:
   in.tftpd: LOCAL, .my.domain

/etc/hosts.deny:
   in.tftpd: ALL: (/some/where/safe\_finger -l @%h | \
       /usr/ucb/mail -s %d-%h root) &
```

The `safe_finger` command comes with the tcpd wrapper and should be installed in a suitable place. It limits possible damage from data sent by the remote finger server. It gives better protection than the standard finger command.

The expansion of the %h (client host) and %d (service name) sequences is described in the section on shell commands.

Warning: do not booby-trap your finger daemon, unless you are prepared for infinite finger loops.

On network firewall systems this trick can be carried even further. The typical network firewall only provides a limited set of services to the outer world. All other services can be "bugged" just like the above tftp example. The result is an excellent early-warning system.

DIAGNOSTICS

An error is reported when a syntax error is found in a host access control rule; when the length of an access control rule exceeds the capacity of an internal buffer; when an access control rule is not terminated by a newline character; when the result of expansion would overflow an internal buffer; when a system call fails that shouldnt́. All problems are reported via the syslog daemon.

FILES

`/etc/hosts.allow`, (daemon,client) pairs that are granted access.
`/etc/hosts.deny`, (daemon,client) pairs that are denied access.

SEE ALSO

tcpd(8) tcp/ip daemon wrapper program.
tcpdchk(8), tcpdmatch(8), test programs.

BUGS

If a name server lookup times out, the host name will not be available to the access control software, even though the host is registered.

Domain name server lookups are case insensitive; NIS (formerly YP) netgroup lookups are case sensitive.

AUTHOR

Wietse Venema (wietse@wzv.win.tue.nl)
Department of Mathematics and Computing Science
Eindhoven University of Technology
Den Dolech 2, P.O. Box 513,
5600 MB Eindhoven, The Netherlands

Bibliography

[OMG99a] *The Common Object Request Broker: Architecture and Specification*, Revision 2.3, OMG, Final publication expected in 1999.

[OMG99b] *CORBAservices: Common Object Services Specification*, OMG, Updated July 1996.

[Richardson96a] *The OMNI Thread Abstraction*, Tristan Richardson, ORL, 22 October 1996.

[Richardson96b] *The OMNI Development Environment Version 4.0*, Tristan Richardson, ORL, 5 November 1996.