

# Developing Distributed Systems

Barry Cornelius

Computing Services, University of Oxford

Date: 29th March 1998

<http://users.ox.ac.uk/~barry/papers/>

<mailto:barry.cornelius@oucs.ox.ac.uk>

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>RMI</b>	<b>1</b>
<b>3</b>	<b>CORBA</b>	<b>4</b>
<b>4</b>	<b>JDBC</b>	<b>12</b>
<b>5</b>	<b>Signed Applets</b>	<b>16</b>
<b>6</b>	<b>Three Tier Systems</b>	<b>19</b>
<b>7</b>	<b>Other Resources</b>	<b>24</b>

---

## 1 Introduction

Client-server computing has traditionally been done using *sockets*. A socket is a point of connection within a program to the Internet, and a program can start to communicate when it arranges for the socket to bind to a port of the computer that is running the program.

The Java classes `java.net.DatagramSocket`, `java.net.ServerSocket` and `java.net.Socket` provide easy-to-use ways of writing programs that use sockets. However, in Java, client-server computing can be performed without having to work at the level of sockets. This is because Java has several application programming interfaces (APIs) that offer higher levels of abstraction for client-server computing.

In this document, we look at three APIs: Remote Method Invocation, Common Object Request Broker Architecture and Java DataBase Connectivity. The latter is a way in which a Java program can use SQL to access database servers, whereas the other two APIs allow us to rise up above the socket-level details of client-server computing when we want to access objects distributed across different computers.

Although Java applications can use these APIs to attempt to communicate with any computer on the Internet, Java applets are normally restricted to talking to the site from which the applet was downloaded. This hinders the use of applets in client-server computing. This document explains how signed applets can be used to overcome these security restrictions.

In the past, client-server computing has usually involved just two computers talking to one another. Often a lot of the processing is going on in the client, which often means that the client is a big program. More recently, systems have been produced that use an intermediate layer: the client just provides a user interface whereas the intermediate layer contains all the complicated application logic. The final topic of this document introduces the idea of Three Tier Systems.

## 2 RMI

### 2.1 What is RMI?

RMI is *Remote Method Invocation*. With RMI, you can access the methods of remote objects. In more detail, on one computer, the *server*, you run:

- the RMI-Registry program;

and:


- a Java application that registers one or more Java objects with the RMI-Registry.

And, on the other computer, the *client* computer, you run:

- a Java application that accesses the methods of the objects of the application running on the server.

### 2.2 How do you use RMI?

Suppose that we want an object that looks after an integral value. We want the ability for any Java application/applet running on any computer to get the latest value, or to set it to some new value:



To use RMI, we need an interface that defines the services that will be available from the object:

```

0001: package RMI.Count; // CountIF.java
0002: public interface CountIF extends java.rmi.Remote {
0003:     int getCount() throws java.rmi.RemoteException;
0004:     void setCount(int vCount) throws java.rmi.RemoteException;
0005: }
  
```

We now need to define what happens when `getCount` and `setCount` get called. We do this by defining a class that implements the interface:

```

0006: package RMI.Count; // CountServant.java
0007: import java.rmi.RemoteException;
0008: import java.rmi.server.UnicastRemoteObject;
0009: public class CountServant extends UnicastRemoteObject implements CountIF {
0010:     public CountServant() throws RemoteException {
0011:         super();
0012:         System.out.println("CountServant: constructor called");
0013:         iCount = 0;
0014:     }
0015:     public synchronized int getCount() throws RemoteException {
0016:         System.out.println("CountServant: getCount called");
0017:         return iCount;
0018:     }
0019:     public synchronized void setCount(int vCount) throws RemoteException {
0020:         System.out.println("CountServant: setCount called with " + vCount);
0021:         iCount = vCount;
0022:     }
0023:     private int iCount;
0024: }
  
```

On the server computer, we need a Java application that creates an object that is of the `CountServant` class. The `CountServer` program creates an object called `realCount` and then registers the object with the RMI-Registry. It does this by using the `rebind` method from the class `java.rmi.Naming` passing to it the name `first Count` that will allow any client to refer to the object:

```

0025: package RMI.Count; // CountServer.java
0026: import java.rmi.Naming;
0027: import java.rmi.RMISecurityManager;
0028: public class CountServer {
0029:     public static void main(String[] args) {
0030:         System.setSecurityManager(new RMISecurityManager());
0031:         try {
0032:             CountServant realCount = new CountServant();
0033:             Naming.rebind("first Count", realCount);
0034:         }
0035:         catch(Exception rException) {
0036:             rException.printStackTrace();
0037:         }
0038:     }
0039: }
  
```

Finally, on the client computer, we want a Java application that accesses the methods of this object:

```

0040: package RMI.Count; // CountClient.java
0041: import java.io.DataInputStream;
0042: import java.rmi.Naming;
0043: import java.rmi.RMISecurityManager;
0044: public class CountClient {
0045:     public static void main(String[] args) {
0046:         DataInputStream stdin = new DataInputStream(System.in);
0047:         System.setSecurityManager(new RMISecurityManager());
0048:         try {
0049:             CountIF stubCount =
0050:                 (CountIF) Naming.lookup("rmi://" + args[0] + "/first Count");
0051:             System.out.println("before loop");
0052:             while (true) {
0053:                 String stdinLine = stdin.readLine();
0054:                 if ( stdinLine.equals("quit") ) break;
0055:                 if ( stdinLine.equals("show") )
0056:                     System.out.println(stubCount.getCount());
0057:                 else {
  
```

```

0058:         int intVal = Integer.parseInt(stdinLine);
0059:         stubCount.setCount(intVal);
0060:     }
0061: }
0062:     System.out.println("after loop");
0063: }
0064: catch(Exception rException) {
0065:     rException.printStackTrace();
0066: }
0067: }
0068: }

```

The CountClient program uses the lookup method from the class java.rmi.Naming to locate the remote object:

```

0049:         CountIF stubCount =
0050:             (CountIF) Naming.lookup("rmi://" + args[0] + "/first Count");


```

It needs to be passed a string like "rmi://hercules.dur.ac.uk/first Count". This string looks like a URL: it specifies the name of the server and the name of the object. In the CountClient program, the name of the server is actually obtained from args[0], the first argument on the command line used to run the program.

The call of lookup communicates with the RMI-Registry on the server. It returns a value that can be used as an object of the interface CountIF. So, we can use the reference variable stubCount to refer to the remote object. In fact, the CountClient program contains a call of both stubCount.setCount and stubCount.getCount.

## 2.3 How does RMI work?

The client program does not have immediate access to the object on the server. Instead, a lot of magic takes place:



The value returned by the call of lookup in the CountClient program is a reference to an object that is part of the client program. This object is a *stub* that acts as a proxy for the remote object. When the client program calls a method, say, it calls stubCount.setCount(42), then it is actually calling a method called setCount of the stub. This method knows nothing about how to set the remote object's value to 42. Instead, what it does is to transmit the required operation to the server: it does this by arranging for the method and its parameters to be *marshalled* into a stream of bytes. This marshalling makes use of JDK 1.1.x's *object serialization*.

Within the Java application running on the server, there is the real object. However, besides this real object there is also a *skeleton* object. The skeleton has a method that arranges for any incoming requests to be unmarshalled. By this means, it finds out what method is to be called and what arguments are to be supplied to the call. It then calls the appropriate method of the real object with the appropriate arguments.

If the method being called returns some result, then the skeleton's method will get the result, serialize the result, and pass the bytes back to the client computer. The method in the stub will receive these bytes, de-serialize them, and return a value to the caller.

So besides the code that you write for the client and the server, you will also need code for the stub and the skeleton. Fortunately, RMI comes with a tool (called rmi c) that will automatically generate the stub and skeleton files for you.

## 2.4 What steps are needed to run the above example?

First, compile the server program with a JDK 1.1.x compiler, e.g.:

```

setenv CLASSPATH /users/dcl0bjc/public_html/Java
javac CountServer.java

```

This javac command will compile CountIF.java, CountServant.java and CountServer.java to produce the files CountIF.class, CountServant.class and CountServer.class.

At a Unix/MS-DOS prompt, run RMI's rmi c command on the file CountServant.class:

```

rmi c RMI.Count.CountServant

```

This command will create the files: CountServant\_Stub.class and CountServant\_Skel.class. These files contain the bytecodes of the stub and skeleton.

Now compile the client program:

```
javac CountClient.java
```

Before running the server, we need to start the RMI-Registry. We can do this using RMI's `rmiregistry` command. On Unix, type:

```
rmiregistry &
```

whereas on Windows 95/NT, type:

```
start rmiregistry
```

The server can then be started using a command like:

```
java -Djava.rmi.server.codebase=http://hercules.dur.ac.uk/~dcl0bjc/Java/ \  
    RMI.Count.CountServer &
```

This should output:

```
CountServant: constructor called
```

Finally, run the client program:

```
setenv CLASSPATH /users/dcl0bjc/public_html/Java  
java RMI.Count.CountClient hercules.dur.ac.uk
```

The client program continually reads in lines from the keyboard. If the line contains the string `quit`, the client program is terminated; if it contains `show`, the client program displays the current value of the remote object; otherwise, the client program expects the line to contain an integral value and it instructs the server to set the remote object to this value.

## 2.5 Further information about RMI

The texts of the above programs are at <http://www.dur.ac.uk/barry.cornelius/papers/RMI/Count/>.

Sun's main RMI page is at <http://java.sun.com/products/jdk/rmi/>. On that page, you will find links to some useful pages entitled *Getting Started*, *Documentation*, *Examples*, *Specification*, *White Paper* and *FAQ*. You will also find a link which enables you to download RMI for JDK 1.0.2.

Users of RMI can discuss issues and tips with other users via the `rmi-users@javasoft.com` mailing list. You can subscribe to this list by sending an e-mail message containing the single line `subscribe rmi-users` to the address `listserv@javasoft.com`. Sun's main RMI page also contains a link to an archive of the `rmi-users` mailing list.

The changes to RMI that are planned for JDK 1.2 are documented at

<http://java.sun.com/products/jdk/1.2/docs/guide/rmi/>.

The Cetus Links on RMI are at [http://www.parallax.co.uk/cetus/oo\\_javabeans.html](http://www.parallax.co.uk/cetus/oo_javabeans.html).

## 3 CORBA

### 3.1 What is CORBA?

CORBA is the *Common Object Request Broker Architecture*. It is an ongoing effort by the Object Management Group (OMG) to define interfaces for interoperable software. Work on CORBA started several years ago, and there are now over 700 members of OMG including virtually all of the computer manufacturers and most of the major software vendors.

If you are using CORBA, communications between two computers is performed using an *Object Request Broker* (ORB). An object in a program running on some computer (the *client*) can use the ORB to access the public attributes of another object in some other program perhaps on a different computer (the *server*) that is also using an ORB. The client object can:

- *either* hard-code the name of the computer and the name of the object that it wants to contact (*Static Method Invocation*);
- *or* it can use the ORB to discover which objects have the required interface (*Dynamic Invocation Interface*).

In this document, we will only be using Static Method Invocation.

Although an ORB can be running on a single computer, it is more likely to be used to aid communication between one or more client computers and a server computer. If an ORB is unable to satisfy a request from a client, then, if it is connected to other ORBs (elsewhere in the world), the request may be satisfied by an object known to a remote ORB.

One of the key aspects of CORBA is the *Interface Definition Language (IDL)*. This is the language that is used to describe the interface that an object has. With RMI, we had to provide the interface in Java: whereas, with CORBA, the interface has to be provided in IDL. It is a language that looks something like C, C++ or Java. For example, it has function declarations and interfaces; it has types such as boolean, char, string, short, long, float and double.

Having produced an IDL file that describes the interface that an object has, we also need some code for the server that creates an object that implements that interface, and we will also need codes for clients that wish to access the object. These codes need to be written in *real* programming languages. In order for an interface written in IDL to be usable with existing languages, there needs to be a mapping that says what each IDL feature is in each language. For example, the IDL-to-Java mapping defines that an IDL long is an int in Java, an IDL string is a String in Java, ... . Currently, there are bindings for Ada, C, C++, Java and Smalltalk.


With RMI, both the client and the server have to be Java applications/applets (running in some Java environment). However, with CORBA, the client and server programs may be in different languages. Each program can run on any operating system written in any language so long as you have a compiler to translate the interface from IDL into that language on that operating system.

There are many vendors providing ORBs. In this document, we will look at two products: *JavaIDL* from Sun Microsystems and *OmniBroker* from Object Oriented Concepts. Some other ORBs are listed in the *Further information about some other ORBs* section given later.

### 3.2 Using Sun's JavaIDL

At the present time, Sun's JavaIDL is not part of the JDK: however, it can be downloaded from Sun's JavaIDL page at <http://java.sun.com/products/jdk/idl/> and used with JDK 1.0.2 or JDK 1.1.x. JavaIDL will form part of JDK 1.2, and a beta release of JDK 1.2 is available from <http://java.sun.com/products/jdk/1.2/>. A version of JavaIDL is also supplied with Sun's Solaris 2.6 operating system.

The use of JavaIDL will be demonstrated by using the same example as was used for RMI:



As before, we will need to provide an interface for our object: this time it is provided in IDL:

```
0069: module CountApp {
0070:     interface Count {
0071:         long getCount();
0072:         void setCount(in long vCount);
0073:     };
0074: };
```

When JavaIDL's `idltojava` command is used to translate this file into Java, it will also produce a class called `_CountImplBase` in a subdirectory called `CountApp`. We can implement the above interface by extending the class `_CountImplBase`:

```
0075: package IDL.Count; // CountServant.java
0076: import IDL.Count.CountApp._CountImplBase;
0077: public class CountServant extends _CountImplBase {
0078:     public CountServant() {
0079:         super();
0080:         System.out.println("CountServant: constructor called");
0081:         iCount = 0;
0082:     }
0083:     public synchronized int getCount() {
0084:         System.out.println("CountServant: getCount called");
0085:         return iCount;
0086:     }
0087:     public synchronized void setCount(int vCount) {
0088:         System.out.println("CountServant: setCount called with " + vCount);
0089:         iCount = vCount;
0090:     }
0091:     private int iCount;
0092: }
```

It is on the server that we have an object that is of the class CountServant. Here is a Java application that creates an object of this class and then asks the ORB's *Naming Service* to bind the name second Count to the object:

```

0093: package IDL.Count;                                     // CountServer.java
0094: import org.omg.CosNaming.NameComponent;
0095: import org.omg.CosNaming.NamingContext;
0096: import org.omg.CosNaming.NamingContextHelper;
0097: import org.omg.CORBA.ORB;
0098: public class CountServer {
0099:     public static void main(String[] args) {
0100:         try{
0101:             // create and initialize the ORB
0102:             ORB orb = ORB.init(args, null);
0103:             // create a CountServant object and register it with the ORB
0104:             CountServant realCount = new CountServant();
0105:             orb.connect(realCount);
0106:             // get the root naming context
0107:             org.omg.CORBA.Object objRef =
0108:                 orb.resolve_initial_references("NameService");
0109:             NamingContext ncRef = NamingContextHelper.narrow(objRef);
0110:             // bind the Object Reference in Naming
0111:             NameComponent nc = new NameComponent("second Count", "");
0112:             NameComponent path[] = {nc};
0113:             ncRef.rebind(path, realCount);
0114:             // wait for invocations from clients
0115:             java.lang.Object sync = new java.lang.Object();
0116:             synchronized (sync) {
0117:                 sync.wait();
0118:             }
0119:         }
0120:         catch (Exception rException) {
0121:             rException.printStackTrace();
0122:         }
0123:     }
0124: }

```

Finally, here is a client program. This is the program that accesses the methods of the object on the server:

```

0125: package IDL.Count;                                     // CountClient.java
0126: import IDL.Count.CountApp.Count;
0127: import IDL.Count.CountApp.CountHelper;
0128: import java.io.DataInputStream;
0129: import org.omg.CosNaming.NameComponent;
0130: import org.omg.CosNaming.NamingContext;
0131: import org.omg.CosNaming.NamingContextHelper;
0132: import org.omg.CORBA.ORB;
0133: public class CountClient {
0134:     public static void main(String[] args) {
0135:         DataInputStream stdin = new DataInputStream(System.in);
0136:         try{
0137:             // create and initialize the ORB
0138:             ORB orb = ORB.init(args, null);
0139:             // get the root naming context
0140:             org.omg.CORBA.Object objRef =
0141:                 orb.resolve_initial_references("NameService");
0142:             NamingContext ncRef = NamingContextHelper.narrow(objRef);
0143:             // resolve the Object Reference in Naming
0144:             NameComponent nc = new NameComponent("second Count", "");
0145:             NameComponent path[] = {nc};
0146:             Count stubCount = CountHelper.narrow(ncRef.resolve(path));
0147:             // enter a loop which processes requests from the keyboard
0148:             System.out.println("before loop");
0149:             while (true) {
0150:                 String stdinLine = stdin.readLine();
0151:                 if ( stdinLine.equals("quit") ) break;
0152:                 if ( stdinLine.equals("show") )
0153:                     System.out.println(stubCount.getCount());
0154:                 else {
0155:                     int intVal = Integer.parseInt(stdinLine);
0156:                     stubCount.setCount(intVal);
0157:                 }
0158:             }
0159:             System.out.println("after loop");
0160:         }
0161:         catch (Exception rException) {
0162:             rException.printStackTrace();
0163:         }
0164:     }
0165: }

```

```

0125: package IDL.Count;                                     // CountClient.java
0126: import IDL.Count.CountApp.Count;
0127: import IDL.Count.CountApp.CountHelper;
0128: import java.io.DataInputStream;
0129: import org.omg.CosNaming.NameComponent;

```

```

0130: import org.omg.CosNaming.NamingContext;
0131: import org.omg.CosNaming.NamingContextHelper;
0132: import org.omg.CORBA.ORB;

0133: public class CountClient {
0134:     public static void main(String[] args) {
0135:         DataInputStream stdin = new DataInputStream(System.in);
0136:         try{
0137:             // create and initialize the ORB
0138:             ORB orb = ORB.init(args, null);
0139:             // get the root naming context
0140:             org.omg.CORBA.Object objRef =
0141:                 orb.resolve_initial_references("NameService");
0142:             NamingContext ncRef = NamingContextHelper.narrow(objRef);
0143:             // resolve the Object Reference in Naming
0144:             NameComponent nc = new NameComponent("second Count", "");
0145:             NameComponent path[] = {nc};
0146:             Count stubCount = CountHelper.narrow(ncRef.resolve(path));
0147:             // enter a loop which processes requests from the keyboard
0148:             System.out.println("before loop");
0149:             while (true) {
0150:                 String stdinLine = stdin.readLine();
0151:                 if ( stdinLine.equals("quit") ) break;
0152:                 if ( stdinLine.equals("show") )
0153:                     System.out.println(stubCount.getCount());
0154:                 else {
0155:                     int intVal = Integer.parseInt(stdinLine);
0156:                     stubCount.setCount(intVal);
0157:                 }
0158:             }
0159:             System.out.println("after loop");
0160:         }
0161:         catch (Exception rException) {
0162:             rException.printStackTrace();
0163:         }
0164:     }

```

### 3.3 What steps are needed to run the above example?

First, you need to compile the interface with JavaIDL's `idltojava` command.

```

setenv JAVAIDL_HOME /users/hplives/JavaIDL-1.1EA
setenv PATH $JAVAIDL_HOME/bin:$PATH
idltojava -fclient -fserver -p IDL.Count -j ../.. Count.idl

```

Because the `Count.idl` file has a module called `CountApp`, this command will create a directory called `CountApp` in which it puts the files that implement the stub and skeleton objects.

Now compile the server program with a JDK 1.0.2 or JDK 1.1.x compiler, e.g.:

```

setenv CLASSPATH $JAVAIDL_HOME/lib/classes.zip:../..
javac CountServer.java

```

This `javac` command will compile `CountServant.java` and `CountServer.java` and some of the files in the `CountApp` directory.

Before running the server, we need to start JavaIDL's Naming Service program. We can do this using JavaIDL's `nameserv` command:

```
nameserv -ORBInitialPort 1050 &
```

This command will produce output like:

```

Initial Naming Context:
IOR:000000000000002849444c3a6f6d672e6f72672f436f734e616d696e672f4e616d696e67436
f6e746578743a312e30000000001000000000000003400010000000000968657263756c657300
00805e0000001cafabcafe00000023448a505000000000000008000000000000000000000
TransientNameServer: setting port for initial object references to: 1050

```

The server can then be started using a command like:

```
java IDL.Count.CountServer -ORBInitialPort 1050 &
```

This should output:

```
CountServant: constructor called
```

Finally, compile the client program with a JDK 1.0.2 or JDK 1.1.x compiler:

```
javac CountClient.java
```

and run it:

```
java IDL.Count.CountClient -ORBInitialHost hercules.dur.ac.uk -ORBInitialPort 1050
```

The client program continually reads in lines from the keyboard. If the line contains the string quit, the client program is terminated; if it contains show, the client program displays the current value of the remote object; otherwise, the client program expects the line to contain an integral value and it instructs the server to set the remote object to this value.

### 3.4 Further information about JavaIDL

The texts of the above programs are at <http://www.dur.ac.uk/barry.cornelius/papers/IDL/Count/>.

Sun's main JavaIDL page is at <http://java.sun.com/products/jdk/idl/>. On that page, you will find links to some useful pages entitled *Documentation*, *Examples*, *Download/EarlyAccess* and *Other References*. In particular, the *Download/EarlyAccess* link enables you to download the Early Access Release of JavaIDL for use with JDK 1.0.2 or JDK 1.1.x.

The above code is derived from Sun's Hello example which is subject to the copyright and licensing information given at <http://java.sun.com/products/jdk/idl/docs/examplelicense.html>.

The details of the JavaIDL being released with JDK 1.2 are documented at <http://java.sun.com/products/jdk/1.2/docs/guide/idl/>.

In many ways, RMI and CORBA are competing technologies. Sun's position about this is explained at <http://java.sun.com/features/1997/nov/rmi.html>.

### 3.5 Using an object's IOR instead of the Naming Service

Earlier, when we started nameserv (JavaIDL's Naming Service program), we saw that it output a line containing the characters IOR: followed by a very long string of hexadecimal characters. (In the output given earlier, the single line containing this string wraps round on to three lines.) This is an example of an *Interoperable Object Reference* (IOR). In this case, it is the IOR of the object that is the Naming Service.

Each object that is known to an ORB has an IOR that is unique to the object. In fact, all you need to communicate with a remote object is to know its IOR. So, if you have got the IORs of all the remote objects you want to access, then there is no need for you to run a Naming Service.

The API has two methods that help us work with IORs: they are called `object_to_string` and `string_to_object`. In the `CountServer` program, we currently have statements that register `realCount` with the Naming Service:

```
0104:         CountServant realCount = new CountServant();
0105:         orb.connect(realCount);
0106:         // get the root naming context
0107:         org.omg.CORBA.Object objRef =
0108:             orb.resolve_initial_references("NameService");
0109:         NamingContext ncRef = NamingContextHelper.narrow(objRef);
0110:         // bind the Object Reference in Naming
0111:         NameComponent nc = new NameComponent("second Count", "");
0112:         NameComponent path[] = {nc};
0113:         ncRef.rebind(path, realCount);
```

These can be replaced by statements that write the IOR of the `realCount` variable to a file called `Count.ior`:

```
0176:         CountServant realCount = new CountServant();
0177:         orb.connect(realCount);
0178:         String realCountString = orb.object_to_string(realCount);
0179:         PrintStream tPrintStream =
0180:             new PrintStream(new FileOutputStream("Count.ior"));
0181:         tPrintStream.println(realCountString);
0182:         tPrintStream.close();
```

And, in the `CountClient` program, the statements that look up the remote object in the Naming Service:

```
0139:         // get the root naming context
0140:         org.omg.CORBA.Object objRef =
0141:             orb.resolve_initial_references("NameService");
0142:         NamingContext ncRef = NamingContextHelper.narrow(objRef);
0143:         // resolve the Object Reference in Naming
0144:         NameComponent nc = new NameComponent("second Count", "");
0145:         NameComponent path[] = {nc};
0146:         Count stubCount = CountHelper.narrow(ncRef.resolve(path));
```

can be replaced by statements that read the IOR of the `realCount` variable from the file `Count.ior`:



```

0206:         DataInputStream tDataInputStream =
0207:             new DataInputStream(new FileInputStream("Count.ior"));
0208:         String realCountString = tDataInputStream.readLine();
0209:         Count stubCount =
0210:             CountHelper.narrow(orb.string_to_object(realCountString));

```

The complete texts of the programs that do not use a Naming Service are at

<http://www.dur.ac.uk/barry.cornelius/papers/IDL/CountNoNS/>.

Using the IOR of a remote object may be a useful technique if you are only communicating with one object. However, it becomes more cumbersome if several objects are involved because you will have to look after several IORs (one for each object).

### 3.6 Communicating with a Naming Service by using its IOR

It is also possible for the server program and/or the client program to use an IOR to contact the Naming Service rather than contact it through the information given by `ORBInitialHost` and `ORBInitialPort`. If we start the Naming Service by the command:

```
nameserv -ORBInitialPort 1050 | tee nameserv.ior &
```

the lines that `nameserv` outputs will be captured in the file `nameserv.ior`. Then instead of using:

```

java IDL.Count.CountServer -ORBInitialPort 1050 &
java IDL.Count.CountClient -ORBInitialHost hercules.dur.ac.uk -ORBInitialPort 1050

```

to tell `CountServer` and `CountClient` where the Naming Service is, we can pass the IOR of the Naming Service as the first argument of the command line being used to run these programs:

```

java IDL.CountNSIOR.CountServer `grep IOR nameserv.ior` &
java IDL.CountNSIOR.CountClient `grep IOR nameserv.ior`

```

In order to get at this IOR, the statement of `CountServer` that says:

```

0107:         org.omg.CORBA.Object objRef =
0108:             orb.resolve_initial_references("NameService");

```

has to be changed to:

```
0244:         org.omg.CORBA.Object objRef = orb.string_to_object(args[0]);
```

A similar change has to be made to the `CountClient` program.

The complete texts of the programs where an IOR is used to contact the Naming Service are at <http://www.dur.ac.uk/barry.cornelius/papers/IDL/CountNSIOR/>.

### 3.7 Using OmniBroker instead of JavaIDL

Object Oriented Concepts has an ORB called *OmniBroker*. This product supports both C++ and Java and is free for non-commercial purposes: <http://www.ooc.com/ob/>.

*OmniBroker* comes with a manual that includes a *Getting started* chapter. Here the authors introduce the idea of communicating with a remote object through its IOR, i.e., this is the technique explained above in the section entitled *Using an object's IOR instead of the Naming Service*. Elsewhere in the manual, they explain how to use *OmniBroker's* implementation of the Naming Service. It is run using *OmniBroker's* `nsserv` program.

In the previous section, we gave programs where remote objects were located using a Naming Service and we communicated with the Naming Service by using its IOR. If we want to use these programs with *OmniBroker* (instead of *JavaIDL*), then two changes have to be made to the `CountServer` program. First change:

```
0234: import org.omg.CORBA.ORB;
```

to:

```

0265: import org.omg.CORBA.ORB;
0266: import org.omg.CORBA.BOA;

```

and also change:

```
0239:         ORB orb = ORB.init(args, null);
```

to:

```
0271:         ORB orb = ORB.init(args, null);
0272:         BOA boa = orb.BOA_init(args, null);
```

The complete texts of the programs where OmniBroker is used instead of JavaIDL are at <http://www.dur.ac.uk/barry.cornelius/papers/OB/CountNSIOR/>.

Here are the steps needed to use these programs with OmniBroker. First, you need to compile the interface with OmniBroker's `jidl` command.

```
setenv PATH /usr/local/utills/OB/current/bin:$PATH
jidl --package OB.CountNSIOR --output-dir ../../ Count.idl
```

Because the `Count.idl` file has a module called `CountApp`, this command will create a directory called `CountApp` in which it puts the files that implement the stub and skeleton objects.

Now compile the server program with a JDK 1.0.2 or JDK 1.1.x compiler, e.g.:

```
setenv CLASSPATH /usr/local/utills/OB/current/lib/OB.jar
setenv CLASSPATH $CLASSPATH:/usr/local/utills/OB/current/lib/CosNaming.jar:../../
javac CountServer.java
```

This `javac` command will compile `CountServant.java` and `CountServer.java` and some of the files in the `CountApp` directory.

Before running the server, we need to start OmniBroker's Naming Service program. We can do this using OmniBroker's `nsserv` command:

```
nsserv --ior | tee nsserv.ior &
```

This command line will output the IOR of the Naming Service both to the screen and to the file `nsserv.ior`.

The server can then be started using a command like:

```
java OB.CountNSIOR.CountServer `cat nsserv.ior` &
```

This should output:

```
CountServant: constructor called
```

Finally, compile the client program with a JDK 1.0.2 or JDK 1.1.x compiler:

```
javac CountClient.java
```

and run it:

```
java OB.CountNSIOR.CountClient `cat nsserv.ior`
```

### 3.8 Using a mix of JavaIDL and OmniBroker

It is possible to mix the use of vendors' ORBs, e.g., to use OmniBroker to create the server and JavaIDL to create the client. Suppose OmniBroker's Naming Service and a server created using OmniBroker have been started by the commands:

```
nsserv --ior | tee nsserv.ior &
java OB.CountNSIOR.CountServer `cat nsserv.ior` &
```

then we can communicate with this server from a client that has been created using JavaIDL:

```
java IDL.CountNSIOR.CountClient `cat ../../OB/CountNSIOR/nsserv.ior`
```

It is also possible to use these the opposite way round, i.e., to use JavaIDL to create the server and OmniBroker to create the client. Suppose JavaIDL's Naming Service and a server created using JavaIDL have been started by the commands:

```
nameserv -ORBInitialPort 1050 | tee nameserv.ior &
java IDL.CountNSIOR.CountServer `grep IOR nameserv.ior` &
```

then we can communicate with this server from a client that has been created using OmniBroker:

```
java OB.CountNSIOR.CountClient `grep IOR ../../IDL/CountNSIOR/nameserv.ior`
```

### 3.9 Using OmniBroker for a client written in C++

Because OmniBroker supports C++ as well as Java, the client and/or the server can be written in C++ instead of Java. For example, here is the client program that we have been using written in C++ instead of Java:

```
0294: #include <iostream.h>
0295: #include <stdlib.h>
0296: #include <OB/CORBA.h>
0297: #include <OB/CosNaming.h>
0298: #include <Count.h>
0299: int main(int argc, char* argv[], char*[] ) {
0300:     try {
0301:         CORBA_ORB_var orb = CORBA_ORB_init(argc, argv);
0302:         CORBA_BOA_var boa = orb -> BOA_init(argc, argv);
0303:         CORBA_Object_var objRef = orb -> string_to_object(argv[1]);
0304:         CosNaming_NamingContext_var ncRef =
0305:             CosNaming_NamingContext::_narrow(objRef);
0306:         CosNaming_Name path;
0307:         path.length(1);
0308:         path[0].id = CORBA_string_dup("second Count");
0309:         path[0].kind = CORBA_string_dup("");
0310:         CountApp_Count_var stubCount =
0311:             CountApp_Count::_narrow(ncRef -> resolve(path));
0312:         cout << "before loop" << endl;
0313:         while ( true ) {
0314:             char stdinLine[100];
0315:             cin >> stdinLine;
0316:             if ( stdinLine[0] == 'q' ) break;
0317:             if ( stdinLine[0] == 's' )
0318:                 cout << stubCount -> getCount() << endl;
0319:             else {
0320:                 int intVal = atoi(stdinLine);
0321:                 stubCount -> setCount(intVal);
0322:             }
0323:         }
0324:         cout << "after loop" << endl;
0325:     }
0326:     catch ( ... ) {
0327:         cout << "whoops" << endl;
0328:     }
0329: }
```

To support C++, OmniBroker comes with a program called `idl` that can be used to process the interface in the file `Count.idl`:

```
setenv PATH /usr/local/utills/OB/current/bin:$PATH
idl Count.idl
```

This will create the files `Count_skel.h`, `Count_skel.cpp`, `Count.h` and `Count.cpp` in the current directory. The client program can now be compiled using a C++ compiler, such as the GNU C++ compiler:

```
g++ -c --no-implicit-templates -fhandle-exceptions \
-I. -I/usr/local/utills/OB/current/include CountClient.cpp
g++ -c --no-implicit-templates -fhandle-exceptions \
-I. -I/usr/local/utills/OB/current/include Count.cpp
```

This will produce the files `CountClient.o` and `Count.o`. These can be linked to produce a binary in the file `CountClient` by the command:

```
g++ -o CountClient CountClient.o Count.o -L/usr/local/utills/OB/current/lib \
-lOB -lCosNaming -lIDL -lsocket -lnsl
```

Finally, the C++ version of the client program can be run, e.g., by using the command:

```
CountClient `grep IOR ../../IDL/CountNSIOR/nameserv.ior`
```

The text of the above program is at <http://www.dur.ac.uk/barry.cornelius/papers/OB/CountNSIORPP/>.

### 3.10 Further information about some other ORBs

Here are a few details about some other ORBs:

- Gerald Brose (of Freie Universitaet Berlin) has developed an ORB for Java called *JacORB*. It is written in Java and is free under the GNU General Public Licence:  
<http://www.inf.fu-berlin.de/brose/jacorb/>.

- IONA Technologies has an ORB called *Orbix* that has bindings for C++, Java, Smalltalk and Ada95. For C++ see <http://www.iona.com/Products/Orbix/> and for Java see <http://www.iona.com/Products/Orbix/OrbixWeb/>. The latter has free evaluation.
- The Oracle and Olivetti Research Laboratory (ORL) has an ORB called *OmniORB2* which is free under the GNU General Public Licence but which currently only supports C++: <http://www.orl.co.uk/omniORB/>.
- Visigenic has an ORB called *VisiBroker* which supports both C++ and Java and can be downloaded free for evaluation: <http://www.visigenic.com/>.

The Cetus Links on ORBs are at [http://www.parallax.co.uk/cetus/oo\\_object\\_request\\_brokers.html](http://www.parallax.co.uk/cetus/oo_object_request_brokers.html).

### 3.11 Further information about CORBA

OMG has a lot of information about CORBA at <http://www.omg.org/>.

The CORBA FAQ is at <http://www.cerfnet.com/mpcline/Corba-FAQ/>.

The Advanced Computing Laboratory at the Los Alamos National Laboratory has a page of resources about CORBA and the OMG at <http://www.acl.lanl.gov/CORBA>.

The Cetus Links on CORBA are at [http://www.parallax.co.uk/cetus/oo\\_corba.html](http://www.parallax.co.uk/cetus/oo_corba.html).

Netscape has a CORBA White Paper <http://developer.netscape.com/library/wpapers/corba/>, and other CORBA documents at <http://developer.netscape.com/library/documentation/corba.html>.

## 4 JDBC

### 4.1 What is JDBC?

The *Structured Query Language* (SQL) is a language for use with relational databases. One of the most used SQL statements is the SELECT statement: it is used to inspect a table of information. The statement:

```
SELECT last, first FROM ages WHERE age >= 21
```

would look at a table called *ages* and produce a new table containing the surnames and first names of the people having an age of at least 21 years.

Although SQL has been evolving since the early 1980's, it was in 1990 that the SQL Access Group defined the *Call Level Interface* (CLI) as a standard for accessing databases. To implement it, you need a *driver* that can translate a CLI call into the language used to access a particular database. Microsoft's Open Database Connectivity (ODBC), an API for Microsoft Windows that implements an extended version of the CLI, was first released in 1992. Most database vendors (including CA/Ingres, IBM, Informix, Oracle and Sybase as well as Microsoft) now support the ODBC API.

ODBC is an API written for the programming language C and although the Java Native Interface (JNI) allows a Java program to call a C function, there are many disadvantages with this approach:

- a Java program calling C code is one which is not so secure, robust and portable;
- calling C code is not usually possible from a Java applet because of security restrictions.

So, instead Java has its own API for submitting SQL statements to a database server: it is JDBC (sometimes called the *Java DataBase Connectivity*).

The JDBC API forms part of JDK 1.1.x: it is in the `java.sql` package. It was not part of JDK 1.0.2: however, if you have a JDK 1.0.2 program, a version of the JDBC API can be downloaded for use with JDK 1.0.2.

The JDBC API is implemented via a driver manager that can support multiple drivers connecting to different databases. A JDBC driver can either be written entirely in Java so that it can be downloaded as part of an applet or it can be implemented using native methods to bridge to existing database access libraries.

At the WWW page <http://java.sun.com/products/jdbc/jdbc.drivers.html>, Sun gives details of the drivers that are available from a number of vendors. On this page, they distinguish between four types of JDBC drivers:

1. 'The *JDBC-ODBC bridge* provides JDBC access via most ODBC drivers. Note that some ODBC binary code and in many cases database client code must be loaded on each client machine that uses this driver, so this kind of driver is most appropriate on a corporate network, or for application server code written in Java in a 3-tier architecture.'
2. 'A *native-API partly-Java driver* converts JDBC calls into calls on the client API for Oracle, Sybase, Informix, DB2, or other DBMS. Note that, like the bridge driver, this style of driver requires that some binary code be loaded on each client machine.'

3. 'A *net-protocol all-Java driver* translates JDBC calls into a DBMS-independent net protocol which is then translated to a DBMS protocol by a server. This net server middleware is able to connect all its Java clients to many different databases. The specific protocol used depends on the vendor. In general, this is the most flexible JDBC alternative. ...'
4. 'A *native-protocol all-Java driver* converts JDBC calls into the network protocol used by DBMSs directly. This allows a direct call from the client machine to the DBMS server and is a practical solution for Intranet access. Since many of these protocols are proprietary the database vendors themselves will be the primary source for this style of driver. Several database vendors have these in progress.'

## 4.2 A Type 4 driver for mSQL

Hughes Technologies has produced *mSQL* (or Mini SQL). This is a 'light-weight relational database management system that has been designed to provide rapid access to data sets with as little system overhead as possible'. The system is comprised of a database server that will run on most Unix platforms together with various tools that allow a user or a client application to communicate with the server. Use of *mSQL* in any commercial environment requires the purchase of a license from Hughes Technologies. However, free licenses are provided to organisations such as Universities, schools and registered charities. For more details about *mSQL*, see <http://www.Hughes.com.au>.

George Reese, a software engineer from Minneapolis, provides a free JDBC driver for *mSQL*. This can be downloaded from the WWW page <http://www.imaginary.com/Java/>.

## 4.3 A Type 1 driver for Microsoft Access

An ODBC driver for use with Microsoft Windows 95/NT products, such as Access and SQL Server, can be obtained from <http://support.microsoft.com/support/kb/articles/Q159/6/74.asp>. Once you have an ODBC driver, then the easiest method to use is the JDBC-ODBC bridge. This is provided as part of JDK 1.1.x.

## 4.4 How do you use JDBC?


A Java program that wants to use JDBC to access a database server will normally consists of the following steps:

1. Use the `getConnection` method (of the class `java.sql.DriverManager`) to establish a connection to a database.
2. Use the `createStatement`, `prepareStatement` or `prepareCall` method.
3. For the `prepareStatement` and `prepareCall` methods, use calls of `setXXX` methods to establish any parameters that are needed by the statement.
4. Use `execute`, `executeQuery` or `executeUpdate` to execute the statement.
5. If the statement is a query, process the result set that is returned.
6. Close the statement.
7. Do steps 2 to 6 for each statement you wish to process.
8. Close the connection that was established by `getConnection`.


By default, all the SQL statements will automatically be *committed*: this is called *autocommit* mode. However, you can disable this and execute calls of `commit` or `rollback` yourself.

## 4.5 A Java application that uses JDBC

Here is a program that is divided into the above steps. The program is written so that it accesses either a database called `bjc1` that is attached to an *mSQL* database server running on a computer called `hercules`:



or it uses the JDBC-ODBC bridge to access a database called `cestria`:



The choice depends on a value passed on the command line used to start the program:

```

0330: import java.sql.Connection; // BothTion.java
0331: import java.io.DataInputStream;
0332: import java.sql.DriverManager;
0333: import java.sql.ResultSet;
0334: import java.sql.ResultSetMetaData;
0335: import java.sql.SQLException;
0336: import java.sql.SQLWarning;
0337: import java.sql.Statement;
0338: public class BothTion {
0339:     public static void main(String[] args) {
0340:         try {
0341:             DataInputStream stdin = new DataInputStream(System.in);
0342:             Connection tConnection = null;
0343:             String tURLString;
0344:             if ( args[0].equals("bjc1") ) {
0345:                 Class.forName("COM.imaginary.sql.mssql.MssqlDriver");
0346:                 tURLString = "jdbc:mssql://hercules.dur.ac.uk:4333/bjc1";
0347:             }
0348:             else {
0349:                 Class.forName("sun.jdbc.odbc.JdbcOdbcDriver");
0350:                 tURLString = "jdbc:odbc:cestria";
0351:             }
0352:             boolean tStillTrying = true;
0353:             int tNumberOfTries = 0;
0354:             while (tStillTrying) {
0355:                 try {
0356:                     tConnection = DriverManager.getConnection(tURLString, "", "");
0357:                     tStillTrying = false;
0358:                 }
0359:                 catch (Exception rException) {
0360:                     tNumberOfTries++;
0361:                     tStillTrying = (tNumberOfTries > 20);
0362:                 }
0363:             }
0364:             if ( tConnection == null ) System.out.println("tConnection is null");
0365:             check(tConnection.getWarnings());
0366:             Statement tStatement = tConnection.createStatement();
0367:             boolean tMoreSQLRequests = true;
0368:             while (tMoreSQLRequests) {
0369:                 System.out.print("> ");
0370:                 String tSQLString = stdin.readLine();
0371:                 try {
0372:                     if ( tSQLString.startsWith("select") ) {
0373:                         ResultSet tResultSet = tStatement.executeQuery(tSQLString);
0374:                         showResultSet(tResultSet);
0375:                         tResultSet.close();
0376:                     }
0377:                     else if ( tSQLString.startsWith("quit") ) {
0378:                         tMoreSQLRequests = false;
0379:                     }
0380:                     else {
0381:                         // assume line contains a SQL insert, update or delete
0382:                         int tNumberOfRows = tStatement.executeUpdate(tSQLString);
0383:                     }
0384:                 }
0385:                 catch (Exception rException) {
0386:                     outputException(rException);
0387:                 }
0388:             }
0389:             tStatement.close();
0390:             tConnection.close();
0391:         }
0392:         catch (Exception rException) {
0393:             outputException(rException);
0394:         }
0395:     }
0396:     private static void showResultSet(ResultSet rResultSet) throws SQLException {
0397:         // Get the ResultSetMetaData. This will be used for
0398:         // the column headings
0399:         ResultSetMetaData tResultSetMetaData = rResultSet.getMetaData();
0400:         // Get the number of columns in the result set
0401:         int tNumCols = tResultSetMetaData.getColumnCount();
  
```

```

0402:         for (int tColNum=1; tColNum<=tNumCols; tColNum++) {
0403:             if (tColNum>1) System.out.print(", ");
0404:             System.out.print(tResultSetMetaData.getColumnLabel(tColNum));
0405:         }
0406:         System.out.println();
0407:         // Display data, fetching until end of the result set
0408:         boolean tMoreRows = rResultSet.next();
0409:         while (tMoreRows) {
0410:             // Loop through each column, getting the
0411:             // column data and displaying
0412:             for (int tColNum=1; tColNum<=tNumCols; tColNum++) {
0413:                 if (tColNum>1) System.out.print(", ");
0414:                 System.out.print(rResultSet.getString(tColNum));
0415:             }
0416:             System.out.println();
0417:             // Fetch the next result set row
0418:             tMoreRows = rResultSet.next();
0419:         }
0420:     }
0421:     private static boolean check(SQLWarning rSQLWarning) throws SQLException {
0422:         boolean tReturnCode = false;
0423:         if (rSQLWarning != null) {
0424:             System.out.println("\n *** Warning ***");
0425:             tReturnCode = true;
0426:             while (rSQLWarning != null) {
0427:                 System.out.println("SQLState: " + rSQLWarning.getSQLState());
0428:                 System.out.println("Message: " + rSQLWarning.getMessage());
0429:                 System.out.println("Vendor: " + rSQLWarning.getErrorCode());
0430:                 System.out.println();
0431:                 rSQLWarning = rSQLWarning.getNextWarning();
0432:             }
0433:         }
0434:         return tReturnCode;
0435:     }
0436:     private static void outputException(Exception rException) {
0437:         if ( rException instanceof SQLException ) {
0438:             SQLException tSQLException = (SQLException) rException;
0439:             System.out.println("\n*** SQLException caught ***");
0440:             while (tSQLException != null) {
0441:                 System.out.println("SQLState: " + tSQLException.getSQLState());
0442:                 System.out.println("Message: " + tSQLException.getMessage());
0443:                 System.out.println("Vendor: " + tSQLException.getErrorCode());
0444:                 tSQLException = tSQLException.getNextException();
0445:                 System.out.println();
0446:             }
0447:         }
0448:         else {
0449:             rException.printStackTrace();
0450:         }
0451:     }
0452: }

```

## 4.6 What steps are needed to run the above example?

First, compile the program with a JDK 1.1.x compiler:

```
javac BothTion.java
```

At runtime, the program either loads the class `COM.imaginary.sql.mssql.MssqlDriver` or it loads the class `sun.jdbc.odbc.JdbcOdbcDriver`. The latter (which is the JDBC-ODBC Bridge) forms part of the JDK 1.1.x distribution whereas the class files for George Reese's mSQL JDBC driver have been put in the directory `/users/dcl0bjc/public_html/Java/COM/imaginary/sql/mssql`. So at runtime the `CLASSPATH` variable needs to be set to ensure that the Java interpreter finds the `MssqlDriver` class:

```
setenv CLASSPATH /users/dcl0bjc/public_html/Java:.
```

When we run the `BothTion` program, an argument on the command line is used to indicate which database we want the program to access:

```
java BothTion bjcl
```

## 4.7 Further information about JDBC

The texts of the above programs are at <http://www.dur.ac.uk/barry.cornelius/papers/JDBC/BothTion/>.

Sun has a JDBC page at <http://java.sun.com/products/jdk/1.1/docs/guide/jdbc/>. A useful guide to JDBC is at <http://java.sun.com/products/jdk/1.1/docs/guide/jdbc/getstart/introTOC.doc.html>.

Sun's main JDBC page at <http://java.sun.com/products/jdbc/> includes details about vendors who have JDBC drivers, an FAQ, and a link to download a JDBC for use with JDK 1.0.2.

Another JDBC FAQ is at <http://japanese.yoyoweb.com/JDBC/FAQ.txt>.

As mentioned above, details about mSQL can be obtained from <http://www.Hughes.com.au> and details about the mSQL-JDBC driver can be obtained from <http://www.imaginary.com/Java/>. This latter page also contains a link to some documentation for the driver, details about the `msql-jdbc` mailing list, and some details about George Reese's book *Database Programming with JDBC and Java*. This book is published by O'Reilly; its ISBN is ISBN 1-56592-270-0.

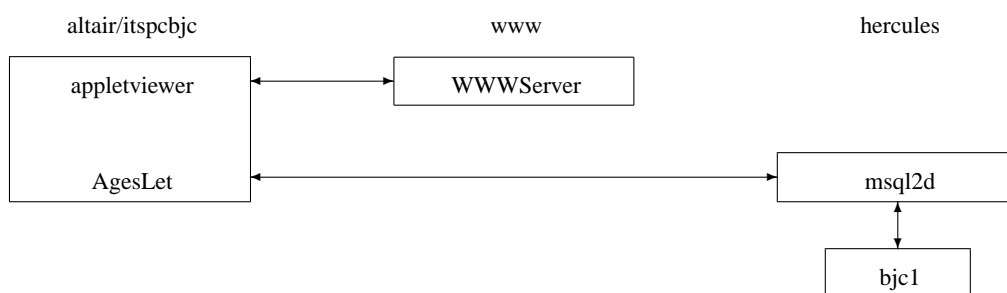
Netscape has a technical note about writing database applications in Java at <http://developer.netscape.com/library/technote/database/web/>.

## 5 Signed Applets

### 5.1 Overview

So far we have seen three different techniques (RMI, CORBA and JDBC): in each case, a Java *application* running on one computer has contacted another computer. When a WWW user wishes to download a Java *applet* from a WWW server, there is no problem if that applet wants to access a program running on the same computer as the WWW server. For example, it is reasonably easy to provide a WWW page that runs a Java applet which uses JDBC to connect to a database server (e.g., the mSQL server) if the database server is on the same computer as the WWW server that delivers the WWW page.

However, if the WWW server and the database server are on different computers:



Java's security restrictions are such that the applet is not allowed access to the database server unless the applet is a *trusted applet*. An applet is *trusted* if the owner of the applet has added a *digital signature* to the applet and the visitor to the WWW page has signified to their WWW browser that they trust the person described in the digital signature.

In order for you to add a digital signature to an applet you need to know about JDK's `javakey` command in order to generate keys, certificates and signed applets. Furthermore, in order to use `javakey`, your applet's class must be stored in a *Java Archive*, and so you will also have to know about JDK's `jar` command (in order to create a Java Archive). Although the `jar` and `javakey` commands of JDK 1.1.x enable you to generate signed applets, these signed applets are not understood by Netscape Navigator 4.03. So in order for a visitor to your WWW page to run your signed applet, he/she will have to use the appletviewer that comes with JDK 1.1.x. In a later version of the JDK (maybe as late as JDK 1.2), the security aspects of the JDK will allow those used by Netscape Navigator 4.03 and so hopefully the two will interoperate.

### 5.2 Creating a Java Archive containing a Java Applet

We will use the following Java applet `AgesLet.java` to communicate with an mSQL database server using the JDBC API:

```
0453: import java.applet.Applet; // AgesLet.java
0454: import java.awt.Button;
0455: import java.sql.Connection;
0456: import java.sql.DriverManager;
0457: import java.awt.Event;
0458: import java.sql.ResultSet;
0459: import java.sql.ResultSetMetaData;
0460: import java.sql.SQLException;
0461: import java.sql.SQLWarning;
0462: import java.sql.Statement;
0463: import java.awt.TextArea;
0464: import java.awt.TextField;
0465: public class AgesLet extends Applet {
0466:     public void init() {
0467:         add(iQueryTextField);
0468:         add(iQueryButton);
0469:         add(iResultsTextArea);
0470:         try {
0471:             Class.forName("COM.imaginary.sql.msql.MsqlDriver");
0472:             String tURLString = "jdbc:msql://hercules.dur.ac.uk:4333/bjc1";
0473:             boolean tStillTrying = true;
```



```

0474:         int tNumberOfTries = 0;
0475:         while (tStillTrying) {
0476:             try {
0477:                 iConnection = DriverManager.getConnection(tURLString, "", "");
0478:                 tStillTrying = false;
0479:             }
0480:             catch (Exception rException) {
0481:                 tNumberOfTries++;
0482:                 tStillTrying = (tNumberOfTries > 20);
0483:             }
0484:         }
0485:         if ( iConnection == null) iResultsTextArea.append("is null\n");
0486:         check(iConnection.getWarnings());
0487:     }
0488:     catch (Exception rException) {
0489:         outputException(rException);
0490:     }
0491: }
0492: public void destroy() {
0493:     try {
0494:         iConnection.close();
0495:     }
0496:     catch (Exception rException) {
0497:         outputException(rException);
0498:     }
0499: }
0500: public boolean handleEvent(Event rEvent) {
0501:     if (rEvent.target.equals(iQueryButton) && rEvent.id==Event.ACTION_EVENT) {
0502:         try {
0503:             String tSQLString = iQueryTextField.getText();
0504:             Statement tStatement = iConnection.createStatement();
0505:             if ( tSQLString.startsWith("select") ) {
0506:                 ResultSet tResultSet = tStatement.executeQuery(tSQLString);
0507:                 displayResultSet(tResultSet);
0508:                 tResultSet.close();
0509:             }
0510:             else {
0511:                 int tNumberOfRows = tStatement.executeUpdate(tSQLString);
0512:             }
0513:             tStatement.close();
0514:         }
0515:         catch (Exception rException) {
0516:             outputException(rException);
0517:         }
0518:         iResultsTextArea.append("-----\n");
0519:         return true;
0520:     }
0521:     return false;
0522: }
0523: private TextField iQueryTextField = new TextField(80);
0524: private Button iQueryButton = new Button("click here");
0525: private TextArea iResultsTextArea = new TextArea(20, 80);
0526: private Connection iConnection = null;
0527: private void displayResultSet(ResultSet rResultSet) throws SQLException {
0528:     ResultSetMetaData tResultSetMetaData = rResultSet.getMetaData();
0529:     int tNumCols = tResultSetMetaData.getColumnCount();
0530:     boolean tMoreRows = rResultSet.next();
0531:     while (tMoreRows) {
0532:         for (int tColNum=1; tColNum<=tNumCols; tColNum++) {
0533:             iResultsTextArea.append(rResultSet.getString(tColNum));
0534:             iResultsTextArea.append(" ");
0535:         }
0536:         iResultsTextArea.append("\n");
0537:         tMoreRows = rResultSet.next();
0538:     }
0539: }
0540: private static boolean check(SQLWarning rSQLWarning) throws SQLException {

```

... same as the check method of BothTion.java (see Section 4.5)

```

0554:     }
0555:     private static void outputException(Exception rException) {

```

... same as the outputException method of BothTion.java (see Section 4.5)

```

0570:     }
0571: }

```

This file can be compiled using the command:

```
javac AgesLet.java
```

The compiler produces a file called AgesLet.class.

You will need to put the classes/COM directory of the mSQL-JDBC 1.0a4 distribution into the current directory. On a Unix system, you can do this by a command like:

```
cp -pr /usr/local/utils/mSQL/JDBC/1.0a4/classes/COM .
```

This creates a subdirectory called COM in the current directory.

You now need to produce a *Java Archive*. This is Java's way of packaging up a lot of files to produce one file. It is done using JDK 1.1.x's jar command at the Unix/MS-DOS prompt:

```
jar cf AgesLet.jar AgesLet.class COM
```

This command will produce a *jar file* in the file AgesLet.jar.

### 5.3 Adding a Digital Signature to a Java Archive

You now need to digitally sign the jar file. This is complicated. First, you need to use JDK 1.1.x's `javakey` command to generate an *identity*. For example, suppose you want to use the identity Barry. You can do this by typing the Unix/MS-DOS command:

```
javakey -cs "Barry" true
```

On a Unix system, this command creates a file called `identitydb.obj` in your home directory, whereas on Microsoft's Windows 95/NT it will create the file in the top directory of your JDK's directory tree. Whenever you use the `javakey` to make a change, you can check what has happened using:

```
javakey -ld
```

You can now create a pair of *keys* (a *private key* and a *public key*) to be used by this identity. This can be done using the command:

```
javakey -gk "Barry" DSA 1024
```

You now use a text editor to create a *certificate directive file*. Suppose the file `Barry.cdirective` contains the lines:

```
issuer.name=Barry
issuer.cert=1
subject.name=Barry
subject.real.name=Barry Cornelius
subject.org.unit=IT Service
subject.org=University of Durham
subject.country=UK
start.date=1 Nov 1997
end.date=30 Nov 1997
serial.number=1
out.file=Barry.x509
```

You can then issue the command:

```
javakey -gc Barry.cdirective
```

This alters the file `identitydb.obj` and also creates the file `Barry.x509`.

You are now (at long last) in a position to sign your jar file. First, you need to use a text editor to create a *signature directive file*. Suppose the file `Barry.sdirective` contains the lines:

```
signer=Barry
cert=1
chain=0
signature.file=Barry
```

Then type the command:

```
javakey -gs Barry.sdirective AgesLet.jar
```

This creates a signed version of your jar file in a file called `AgesLet.jar.sig`.

### 5.4 Creating a WWW page that loads the Java Archive

You can then produce an appropriate HTML file in a file called `AgesLet.html`:

```
0572: <!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 3.2//EN">
0573: <HTML><HEAD><TITLE> The AgesLet Program </TITLE></HEAD>
0574: <BODY>
0575: <APPLET ARCHIVE="AgesLet.jar.sig" CODE="AgesLet.class" WIDTH=700 HEIGHT=700>
0576: </APPLET>
0577: </BODY></HTML>
```

Note that the `APPLET` in this HTML file has both an `ARCHIVE` and a `CODE` section. When a WWW browser accesses this page, it will download all the classes in the Java Archive `AgesLet.jar.sig` and then it will start interpreting the class that came from the file `AgesLet.class`.

## 5.5 Accessing your WWW page from a WWW browser

The author of this applet can now run it using the appletviewer that comes with JDK 1.1.x, e.g., by typing the Unix/MS-DOS command:

```
appletviewer http://www.dur.ac.uk/~dcl0bjc/Java/Signed/Ages/AgesLet.html
```

If another user wants to execute the above appletviewer command perhaps somewhere else in the world, then first they will have to have a copy of your x509 certificate file (that is in the file `Barry.x509`) and then they will have to execute the Unix/MS-DOS command:

```
javakey -c "Barry" true
```

to introduce Barry as an identity and then execute:

```
javakey -ic "Barry" Barry.x509
```

to add Barry's certificate to their `identitydb.obj` file. Having done that, appletviewer will execute any applet that have been signed by Barry.

## 5.6 Further information about Signed Applets

The texts of the above programs are at <http://www.dur.ac.uk/barry.cornelius/papers/Signed/Ages/>.

JDK 1.1.x's appletviewer must be used instead of a browser such as Netscape Navigator. This is because JDK 1.1.x's certificates are X509v1 certificates for DSA whereas Netscape's Navigator uses X509v3. However, Sun has said that the next release of the JDK will be able to read X509v3 certificates, and hopefully 'it will be a straightforward matter for the Java-enabled browsers to use the JDK certificate APIs to import and parse the JDK signing certificates'. In the next release, Sun will replace the `javakey` tool by `keytool`, `jarsigner` and `policytool`, and the security database used by these new tools will be incompatible with the database in the `identitydb.obj` file. Although these changes may take place in JDK 1.1.6 (possibly due in December 1997), they are more likely to occur in JDK 1.2 (possibly due in Q2 1998).

Sun has one page about security at <http://java.sun.com/products/jdk/1.1/docs/guide/security/>. However, their main security page is at <http://java.sun.com/security/>. There is a guide to using `javakey` at <http://java.sun.com/security/usingJavakey.html> and recommendations on its use at <http://java.sun.com/security/policy.html>. The details about the replacement of `javakey` by `keytool`, `jarsigner` and `policytool` is given at <http://java.sun.com/products/jdk/1.2/docs/guide/security/>.

There is a section about keys, certificates and digital signatures in Sun's *Tutorial on Java*: <http://java.sun.com/docs/books/tutorial/security1.1/>.

Netscape's WWW page for signed objects is at

<http://developer.netscape.com/library/documentation/signedobj/>.

## 6 Three Tier Systems

### 6.1 Two Tier and Three Tier Systems

In the past, many distributed systems have been built that are *two tier*. This is a system that has a client program communicating with a server program. The examples of RMI, JavaIDL and JDBC given earlier are simple two-tier systems. Typically, the user interface is in the client and the server is there to process data, e.g., to manipulate a database. The complicated application logic (or business logic) is often put in the client (or in the server). More recently, *three tier* systems have come into favour. This is where an additional tier is interposed: it is there to handle the application logic leaving the client just to handle the user interface.

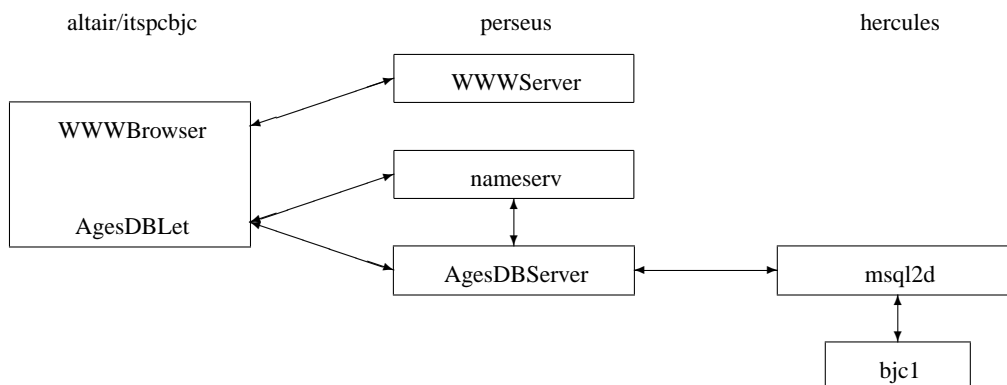
There are a number of advantages to this approach. For example:

- The code running on the client is small.
- When different client programs are needed to handle the differences between different client platforms, the amount of code that has to be provided is small because it only has to handle the user interface.
- It is possible to write the code running on the middle tier so that it handles requests from more than one client. In this way, requests can be more efficiently handled.

### 6.2 An example of a Three Tier System

We will now look at an unsophisticated three tier system. It will use a WWW browser on a computer called `altair` that downloads a Java applet from a WWW server running on `perseus`. The Java applet will display some buttons and textfields. A click on a button will cause the Java applet to send a request (using JavaIDL) to a

Java application that is also running on perseus. This Java application will send (using JDBC) an appropriate SQL statement to a database server running on a computer called hercules:



### 6.3 The code for this Three Tier System

When the following WWW page is visited, the WWW browser will download all the classes in the Java Archive AgesDBLet.jar and then it will start interpreting the class that came from the file AgesDBLet.class passing to it the two parameters mentioned in the PARAM tags:

```

0578: <!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 3.2//EN">
0579: <HTML><HEAD><TITLE> The AgesDBLet Program </TITLE></HEAD>
0580: <BODY>
0581: <APPLET ARCHIVE=AgesDBLet.jar CODE="AgesDBLet.class" WIDTH=700 HEIGHT=700>
0582: <PARAM NAME=org.omg.CORBA.ORBInitialHost VALUE="perseus.dur.ac.uk">
0583: <PARAM NAME=org.omg.CORBA.ORBInitialPort VALUE=1050>
0584: </APPLET>
0585: </BODY></HTML>
  
```

The Java applet is going to communicate using JavaIDL with an object that is part of a Java application running on perseus. Suppose the object has the following interface (given in IDL):

```

0586: module AgesDBApp {
0587:     interface AgesDB {
0588:         void getPerson(in string rLast, out string rFirst, out string rAge);
0589:         void addPerson(in string rLast, in string rFirst, in string rAge);
0590:         void removePerson(in string rLast);
0591:     };
0592: };
  
```

Here is the text of the Java applet AgesDBLet.java:

```

0593: import AgesDBApp.AgesDB; // AgesDBLet.java
0594: import AgesDBApp.AgesDBHelper;
0595: import java.applet.Applet;
0596: import java.awt.Button;
0597: import java.awt.Event;
0598: import org.omg.CosNaming.NameComponent;
0599: import org.omg.CosNaming.NamingContext;
0600: import org.omg.CosNaming.NamingContextHelper;
0601: import org.omg.CORBA.ORB;
0602: import org.omg.CORBA.StringHolder;
0603: import java.awt.TextField;
0604: public class AgesDBLet extends Applet {
0605:     public void init() {
0606:         // establish connection to the remote object through the iAgesDB variable
0607:         try {
0608:             ORB orb = ORB.init(this, null);
0609:             org.omg.CORBA.Object objRef =
0610:                 orb.resolve_initial_references("NameService");
0611:             NamingContext ncRef = NamingContextHelper.narrow(objRef);
0612:             NameComponent nc = new NameComponent("AgesDB", "");
0613:             NameComponent path[] = {nc};
0614:             iAgesDB = AgesDBHelper.narrow(ncRef.resolve(path));
0615:         }
0616:         catch (Exception rException) {
0617:             rException.printStackTrace();
0618:         }
0619:         // add the three buttons to the applet
0620:         add(iGetButton);
0621:         add(iAddButton);
0622:         add(iRemoveButton);
0623:         // add the three textfields to the applet
0624:         add(iLastTextField);
0625:         add(iFirstTextField);
  
```

```

0626:         add(iAgeTextField);
0627:     }
0628:     public boolean handleEvent(Event rEvent) {
0629:         try {
0630:             if (rEvent.id == Event.ACTION_EVENT) {
0631:                 if (rEvent.target.equals(iGetButton)) {
0632:                     // In the interface, getPerson has two 'out' parameters. However,
0633:                     // Java only has 'in'. The class org.omg.CORBA.StringHolder
0634:                     // can be used to provide a wrapper object for a string value.
0635:                     StringHolder tFirstHolder = new StringHolder();
0636:                     StringHolder tAgeHolder = new StringHolder();
0637:                     iAgesDB.getPerson(iLastTextField.getText(),
0638:                                     tFirstHolder, tAgeHolder);
0639:                     iFirstTextField.setText(tFirstHolder.value);
0640:                     iAgeTextField.setText(tAgeHolder.value);
0641:                     return true;
0642:                 }
0643:                 if (rEvent.target.equals(iAddButton)) {
0644:                     iAgesDB.addPerson(iLastTextField.getText(),
0645:                                       iFirstTextField.getText(), iAgeTextField.getText());
0646:                     return true;
0647:                 }
0648:                 if (rEvent.target.equals(iRemoveButton)) {
0649:                     iAgesDB.removePerson(iLastTextField.getText());
0650:                     return true;
0651:                 }
0652:             }
0653:             return false;
0654:         }
0655:         catch (Exception rException) {
0656:             rException.printStackTrace();
0657:             return false;
0658:         }
0659:     }
0660:     private AgesDB iAgesDB = null;
0661:     private Button iGetButton = new Button("Get");
0662:     private Button iAddButton = new Button("Add");
0663:     private Button iRemoveButton = new Button("Remove");
0664:     private TextField iLastTextField = new TextField(30);
0665:     private TextField iFirstTextField = new TextField(30);
0666:     private TextField iAgeTextField = new TextField(4);
0667: }

```

In the Java application on perseus, an object of the following class implements the interface given earlier (in IDL) by extending the class `_AgesDBImplBase`:

```

0668: import AgesDBApp._AgesDBImplBase; // AgesDBServant.java
0669: import java.sql.Connection;
0670: import java.sql.DriverManager;
0671: import java.sql.ResultSet;
0672: import java.sql.ResultSetMetaData;
0673: import java.sql.SQLException;
0674: import java.sql.SQLWarning;
0675: import java.sql.Statement;
0676: import org.omg.CORBA.StringHolder;
0677: public class AgesDBServant extends _AgesDBImplBase {
0678:     public AgesDBServant() {
0679:         super();
0680:         System.out.println("AgesDBServant: constructor start");
0681:         try {
0682:             Class.forName("COM.imaginary.sql.mssql.MssqlDriver");
0683:             String tURLString = "jdbc:mssql://hercules.dur.ac.uk:4333/bjc1";
0684:             boolean tStillTrying = true;
0685:             int tNumberOfTries = 0;
0686:             while (tStillTrying) {
0687:                 try {
0688:                     iConnection = DriverManager.getConnection(tURLString, "", "");
0689:                     tStillTrying = false;
0690:                 }
0691:                 catch (Exception rException) {
0692:                     tNumberOfTries++;
0693:                     tStillTrying = (tNumberOfTries > 20);
0694:                 }
0695:             }
0696:             if ( iConnection == null ) System.out.println("iConnection is null");
0697:             check(iConnection.getWarnings());
0698:         }
0699:         catch (Exception rException) {
0700:             outputException(rException);
0701:         }
0702:         System.out.println("AgesDBServant: constructor finish");
0703:     }
0704:     public synchronized void getPerson(String rLast,
0705:                                       StringHolder rFirst, StringHolder rAge) {
0706:         System.out.println("AgesDBServant: getPerson start");
0707:         rFirst.value = "";

```

```

0708:     rAge.value = "";
0709:     try {
0710:         Statement tStatement = iConnection.createStatement();
0711:         String tSQLString = "select * from ages where last = '" + rLast + "'";
0712:         System.out.println(tSQLString);
0713:         ResultSet tResultSet = tStatement.executeQuery(tSQLString);
0714:         if ( tResultSet.next() ) {
0715:             rFirst.value = tResultSet.getString(1);
0716:             rAge.value = tResultSet.getString(3);
0717:         }
0718:         tResultSet.close();
0719:         tStatement.close();
0720:     }
0721:     catch (Exception rException) {
0722:         outputException(rException);
0723:     }
0724:     System.out.println("AgesDBServant: getPerson finish");
0725: }
0726: public synchronized void addPerson(String rLast, String rFirst, String rAge){
0727:     System.out.println("AgesDBServant: addPerson start");
0728:     try {
0729:         Statement tStatement = iConnection.createStatement();
0730:         String tSQLString = "insert into ages values('"
0731:             + rFirst + "', '" + rLast + "', " + rAge + ")";
0732:         System.out.println(tSQLString);
0733:         int tNumberOfRows = tStatement.executeUpdate(tSQLString);
0734:         tStatement.close();
0735:     }
0736:     catch (Exception rException) {
0737:         outputException(rException);
0738:     }
0739:     System.out.println("AgesDBServant: addPerson finish");
0740: }

0741: public synchronized void removePerson(String rLast) {
0742:     System.out.println("AgesDBServant: removePerson start");
0743:     try {
0744:         Statement tStatement = iConnection.createStatement();
0745:         String tSQLString = "delete from ages where last = '" + rLast + "'";
0746:         System.out.println(tSQLString);
0747:         int tNumberOfRows = tStatement.executeUpdate(tSQLString);
0748:         tStatement.close();
0749:     }
0750:     catch (Exception rException) {
0751:         outputException(rException);
0752:     }
0753:     System.out.println("AgesDBServant: removePerson finish");
0754: }
0755: private Connection iConnection = null;
0756: private static boolean check(SQLWarning rSQLWarning) throws SQLException {

```

... same as the check method of BothTion.java (see Section 4.5)

```

0770:     }
0771:     private static void outputException(Exception rException) {

```

... same as the outputException method of BothTion.java (see Section 4.5)

```

0786:     }
0787: }

```

The Java application running on perseus creates an object of the above class and registers it:

```

0788: import org.omg.CosNaming.NameComponent; // AgesDBServer.java
0789: import org.omg.CosNaming.NamingContext;
0790: import org.omg.CosNaming.NamingContextHelper;
0791: import org.omg.CORBA.ORB;
0792: public class AgesDBServer {
0793:     public static void main(String[] args) {
0794:         try{
0795:             ORB orb = ORB.init(args, null);
0796:             AgesDBServant AgesDBRef = new AgesDBServant();
0797:             orb.connect(AgesDBRef);
0798:             org.omg.CORBA.Object objRef =
0799:                 orb.resolve_initial_references("NameService");
0800:             NamingContext ncRef = NamingContextHelper.narrow(objRef);
0801:             NameComponent nc = new NameComponent("AgesDB", "");
0802:             NameComponent path[] = {nc};
0803:             ncRef.rebind(path, AgesDBRef);
0804:             java.lang.Object sync = new java.lang.Object();
0805:             synchronized (sync) {
0806:                 sync.wait();
0807:             }

```

```

0808:     }
0809:     catch (Exception rException) {
0810:         rException.printStackTrace();
0811:     }
0812: }
0813: }

```

The above Java application uses JDBC to talk to a database server running on hercules.

## 6.4 What steps are needed to run this example?

First, compile the interface in the file AgesDB.idl by executing JavaIDL's idltojava command:

```

setenv JAVAIDL_HOME /users/hplives/JavaIDL-1.1EA
setenv PATH $JAVAIDL_HOME/bin:$PATH
idltojava -fclient -fserver AgesDB.idl

```

This command creates a directory called AgesDBApp for the stub and skeleton classes.

We need to compile AgesDBServer.java, the Java application of the middle tier. In order to compile this file, we will need access to the bytecodes of the classes of the JavaIDL API. Earlier, this was done by setting the CLASSPATH variable to include the file \$JAVAIDL\_HOME/lib/classes.zip. However, for a technical reason (to do with JavaStations), this file will be unzipped, to create two directories in the current directory called com and org. Having done that, we can compile the Java application of the middle tier with a JDK 1.1.x compiler:

```
javac AgesDBServer.java
```

Before running the server, we need to start nameserv (JavaIDL's Naming Service program):

```
nameserv -ORBInitialPort 1050 &
```

The server can then be started using the commands:

```

setenv CLASSPATH /users/dcl0bjc/public_html/Java:.
java AgesDBServer -ORBInitialPort 1050 &

```

Finally, compile the client program with a JDK 1.0.2 or JDK 1.1.x compiler:

```
javac AgesDBLet.java
```

This compilation will use the JavaIDL classes in the com and org directories that were created earlier. In order to simplify the loading of these classes whenever the Java applet is downloaded, we will use JDK's 1.1.x's jar command to build a Java Archive that contains all of these files:

```
jar cf AgesDBLet.jar AgesDBLet.class AgesDBApp org com
```

This Java applet can be run by anyone having a WWW browser that can handle Java Archives. Examples are Netscape's Navigator 3.01 (or later) and JDK 1.1.x's appletviewer:

```
appletviewer http://perseus.dur.ac.uk/~dcl0bjc/Java/ThreeTier/Ages/AgesDBLet.html
```

Because of an incompatibility, such applets do not work with Netscape's Navigator 4.0x. I have also used this applet as the applet behind one of the buttons on the desktop when using *HotJava Views* on a *JavaStation*.

## 6.5 Using a C++ program as a client of this Three Tier System

Here is a C++ program that can be used as the client of this Three Tier System:

```

0814: #include <iostream.h>
0815: #include <stdlib.h>
0816: #include <OB/CORBA.h>
0817: #include <OB/CosNaming.h>
0818: #include <AgesDB.h>
0819: int main(int argc, char* argv[], char*[]) {
0820:     try {
0821:         CORBA_ORB_var orb = CORBA_ORB_init(argc, argv);
0822:         CORBA_BOA_var boa = orb -> BOA_init(argc, argv);
0823:         CORBA_Object_var objRef = orb -> string_to_object(argv[1]);
0824:         CosNaming_NamingContext_var ncRef =
0825:             CosNaming_NamingContext::_narrow(objRef);
0826:         CosNaming_Name path;
0827:         path.length(1);
0828:         path[0].id = CORBA_string_dup("AgesDB");
0829:         path[0].kind = CORBA_string_dup("");

```

```

0830:     AgesDBApp_AgesDB_var stubAgesDB =
0831:         AgesDBApp_AgesDB::_narrow(ncRef -> resolve(path));
0832:     char stdinLine[100];
0833:     char in_last[10];
0834:     char in_first[10];
0835:     char in_age[10];
0836:     char *out_firstptr = 0;
0837:     char *out_ageptr = 0;
0838:     bool stillMoreRequests = true;
0839:     while ( stillMoreRequests ) {
0840:         cout << "agqr: "; cin >> stdinLine;
0841:         switch ( stdinLine[0] ) {
0842:             case 'a':
0843:                 cout << "last: "; cin >> in_last;
0844:                 cout << "first: "; cin >> in_first;
0845:                 cout << "age: "; cin >> in_age;
0846:                 stubAgesDB -> addPerson(in_last, in_first, in_age);
0847:                 break;
0848:             case 'g':
0849:                 cout << "last: "; cin >> in_last;
0850:                 stubAgesDB -> getPerson(in_last, out_firstptr, out_ageptr);
0851:                 cout << out_firstptr << " " << out_ageptr << endl;
0852:                 break;
0853:             case 'q':
0854:                 stillMoreRequests = false;
0855:                 break;
0856:             case 'r':
0857:                 cout << "last: "; cin >> in_last;
0858:                 stubAgesDB -> removePerson(in_last);
0859:                 break;
0860:             default:
0861:                 break;
0862:         }
0863:     }
0864: }
0865: catch ( ... ) {
0866:     cout << "whoops" << endl;
0867: }
0868: }

```

With the help of OmniBroker, this C++ program can be compiled and executed in a similar way to the C++ program given earlier.

## 6.6 Further information about Three Tier Systems

The texts of the above programs are at <http://www.dur.ac.uk/barry.cornelius/papers/ThreeTier/Ages/> and at <http://www.dur.ac.uk/barry.cornelius/papers/OB/ThreeTierPP/>.

Sun has a WWW page describing Three Tier Systems at

<http://java.sun.com/products/jdk/1.2/docs/guide/idl/jidlDistApp.html>.

For more about JavaStations and HotJava Views, see <http://www.sun.com/javastation/> and <http://java.sun.com/products/hotjavaviews/>.

## 7 Other Resources

One good book is *Client/Server Programming with Java and CORBA* (1997) by Robert Orfali and Dan Harkey published by Wiley (0-471-16351-1): <http://www.corbajava.engr.sjsu.edu/Progrmg.htm>. It contains extensive sections on both CORBA and JDBC, and the use of CORBA and JDBC to build multi-tier systems. They also compare CORBA with other methods: sockets, HTTP/CGI and RMI. The second edition of this book (0-471-24578-X) 'with 250 pages of new material' is due out in February 1998. Working with Jeri Edwards, Orfali and Harkey have three other books in this area: *Essential Client/Server Survival Guide* (2nd edition 1996), *Essential Distributed Objects Survival Guide* (1996) and *Instant CORBA* (1997). See <http://www.corbajava.engr.sjsu.edu/mbooks.htm>.

Prashant Sridharan's book *Advanced Java Networking* explores the topics covered in this document. Its ISBN is 0-13-749136-0 and it is published by Prentice-Hall: <http://www.prenhall.com/>.

Sams.net Publishing have several overlapping books on Java, and their *Unleashed* book keeps moving into new editions. Many of the chapters of *Java 1.1 Unleashed Third Edition* describe various aspects of client-server computing. Its ISBN is 1-57521-298-6. For more details, see: <http://merchant.superlibrary.com:8000/catalog/hg/PRODUCT/PAGE/15752/bud/1575212986.html>.

The text of this document is updated from time to time. The latest version is available at:

<http://www.dur.ac.uk/dcl0bjc/Java/>. This WWW page also contains other documents on Java including a tutorial on Java, an overview of the material of the tutorial and some notes comparing Java and C++.