

A Test Framework for CORBA Interoperability

Mang Li¹, Arno Puder², Ina Schieferdecker¹

¹GMD FOKUS

Competence Center for Testing, Interoperability and Performance
Kaiserin-Augusta-Allee 31, D-10589 Berlin, Germany
{m.li, Schieferdecker}@fokus.gmd.de

²AT&T Labs Research

75 Willow Road
Menlo Park, CA 94025, USA
arno@research.att.com

Abstract

This paper presents a new approach to interoperability testing of CORBA implementations. The ORB Under Test (OUT) is embedded in the test system in such a way, that the behavior of either the client side or the server side of the OUT is controllable and observable. Components of the test system emulate the peer entities of the OUT and the applications based on the OUT, in order to get direct access both to the GIOP interface and to the CORBA APIs.

A specification framework for the test system is proposed. It separates the concerns of test data and test behavior to facilitate modularity and dynamic of the test system. The description languages CORBA IDL, ODMG ODL and TTCN-3 are used in the framework.

Keywords

CORBA, conformance testing, interoperability, ODMG-ODL, TTCN-3

1. Introduction

Applications increasingly cross boundaries of technological, administrative and political domains. Fields like E-commerce or Business-to-Business increase the need for cross-platform development of distributed applications. Because of the heterogeneity of these environments, it is not possible to enforce or even assume one single technology. Heterogeneity exists at different levels; ranging from

different network technologies, different operating systems and programming languages. Middleware platforms are proven to be an adequate means to cope with heterogeneous environments. A middleware platform spreads out like a table cloth in a heterogeneous environment, offering a well-known Application Programming Interface (API).

The Common Object Request Broker Architecture (CORBA) [7] provides a framework for the development of distributed object-oriented applications in heterogeneous environments. CORBA is published as a set of specifications that describe the behavior of a middleware platform. The specification does not prescribe implementation details.

Different vendors offering CORBA platforms can focus on different market segments and choose appropriate technical details for their implementations. As a consequence, there exist multiple CORBA implementations, ranging from commercial products to open source versions.

Deriving an implementation from a specification raises the question of conformance, i.e., if the implementation conforms to the specification. Conformance guarantees that the behavior is in accordance with the specification. It is the key to achieve interoperability and easy portability of applications between CORBA platforms. Without conformance, the very philosophy of CORBA is compromised which aims for the deployment of applications in heterogeneous environments.

This paper focuses on conformance testing of CORBA implementations, especially on the interoperability aspect

of CORBA. Section 2 gives a brief overview of CORBA and explains the difference between interoperability and portability. It is argued that interoperability plays a more crucial role than portability. Section 3 first discusses several different approaches to test interoperability with their respective drawbacks. It is followed by a proposal for a new test configuration. Section 4 then introduces a test framework based on the configuration discussed in the previous section. Section 5 finally concludes this paper with an outlook for future work.

2. CORBA and GIOP

In this section we first give a brief overview of CORBA. Our discussion focuses on the key concepts portability and interoperability of applications.

CORBA as a middleware platform provides the specification of two different kinds of interfaces which we call *horizontal* and *vertical interfaces* (see Figure 1). The horizontal interface separates the application from the middleware. The CORBA specification defines the horizontal interface amongst others through the Interface Definition Language (IDL), the Dynamic Invocation Interface (DII), the Dynamic Skeleton Interface (DSI) and the Portable Object Adapter (POA).




Figure 1 Vertical and horizontal interfaces

The vertical interface resides between two instances of a CORBA implementation. The vertical interface generally is defined through a protocol consisting of Protocol Data Units (PDUs) and marshalling conventions amongst others. The CORBA specification defines the vertical interface through the General Inter-ORB Protocol (GIOP), see [11]. GIOP defines an abstract framework for the vertical interface while making certain assumptions of the underlying transport layer. The Internet Inter-ORB Protocol (IIOP) is the application of GIOP using TCP as a transport layer, which is also defined in the CORBA specification. For the scope of this paper we will limit the discussion to GIOP.

The application programmer generally is unaware of the details of GIOP, while a detailed knowledge of the horizontal interface is necessary for building distributed applications. The standardization of the horizontal interface enables *portability* while the specification of the vertical

interface guarantees *interoperability* of applications. Horizontal and vertical interface are not independent of each other. E.g., data types definable based on the IDL at the horizontal interface must be marshalled at the vertical interface.

In the following, we argue that with respect to conformance to the CORBA specification interoperability plays a more important role than portability. As shown in Figure 1, the vertical interface decouples technological domains. The vendors of ORB1 and ORB2 can target specific markets. While a specific technological domain may require the modification of the horizontal interface, the vertical interface should not be modified in order to retain interoperability.

One example of a technological domain are embedded systems. The CORBA specification offers solutions for these environments through its MinimumCORBA specification [7]. MinimumCORBA is a true subset of the full CORBA specification in order to enable support for CORBA in resource limited environments. MinimumCORBA removes certain elements from the horizontal interface (e.g., DII and DSI) but explicitly places no restrictions on the vertical interface in order not to compromise interoperability with other CORBA platforms.

MinimumCORBA is therefore a good example where the CORBA specification itself limits portability of applications but retains full interoperability. No matter if changes to the horizontal interface compromises portability of applications, violating the interoperability is against the very core of CORBA's philosophy. In terms of conformance testing we therefore focus on interoperability.

3. Test configuration

In this section we discuss the configuration for the conformance testing of the vertical GIOP interface. First, we briefly review some related approaches and discuss their drawbacks. Then, we introduce our proposal that aims to cope with the requirements raised in the previous section.

We will use the acronym OUT to denote the *ORB Under Test*.

3.1. Related approaches

Three approaches are introduced in this section, which are used by most of the available CORBA interoperability tests.

3.1.1. OUT on client and server side

In this configuration, the OUT is used on both the client and the server side (i.e., ORB1 and ORB2 in Figure 1). The test suite is running on top of the OUT, invoking operations with actual parameters. The test suite makes sure that the operation reaches the server and that all parameters are

transmitted correctly. This approach only guarantees the interoperability with respect to one ORB, namely the OUT. While this approach makes a statement about the possibility to build a distributed application based on the OUT, it cannot make any statements with respect to interoperability with other ORBs. In fact, the OUT could choose to implement a proprietary protocol that is not based on GIOP. This would make it impossible to interoperate with a different vendors ORB.

3.1.2. Passive testing of vertical interface

One way to assure that the protocol at the vertical interface is indeed based on GIOP is to inspect the network traffic between two instances of the OUT as the test proceeds. This approach was chosen for the conformance test suite developed by The Open Group [10]. Here a protocol analyzer sits between the client and the server OUT and ensures that the encoding of parameters and PDUs conform with GIOP.

While this approach can certify that the traffic between two instances of the OUT is conformant to GIOP, it does little to enforce all aspects of interoperability. The shortcomings have to do with the fact that GIOP allows certain implementation freedoms that cannot be controlled at the horizontal interface. An example for this is the endianness of the transmitted data. In GIOP, the sender can choose the endianness (little or big endian) encoding of the data to be transmitted. It is the receiver's responsibility to convert the received data into its preferred format. But GIOP does not prescribe a policy that is binding for the sender when to use which endianness. If the OUT decides to use little endian encodings for all outgoing PDUs, the observed traffic at the vertical interface will be conformant to GIOP, but the OUT will never have to demonstrate its fitness to cope with big endian encodings on the receiving side.

Since the OUT can only be controlled at its stub/skeleton interfaces, there is no way to influence when which kind of encoding will be used. There are other examples where GIOP does not prescribe specific policies that cannot be controlled at the horizontal interfaces. Examples are the padding during the alignment of data, encoding of union default members and encapsulations of embedded encodings.

3.1.3. Reference ORB implementation

To circumvent the problem mentioned in the last section, one might be tempted to use a reference ORB against the OUT. But doing this yields the same problems mentioned in the previous section, as the reference ORB will choose the native endianness for the encoding of the data, just as the OUT.

Also the approach based on the request-reply paradigm

(see section Section 3.1.2) may use a reference ORB. In fact, a reference implementation in a certain form is always used in testing. It must behave as an ORB implementation to provide the functionality as much as the tests require. The existence of a reference ORB is in many cases only an assumption. Important is to use an adequate method and granularity so that failures can be resolved easily.

3.2. Proposal for new test configuration

The test configuration we propose in the following is to allow direct access to the GIOP interface and active control of messages exchanged with the OUT. In addition, when the test control involves the horizontal interface, its use should not restrict testing of the GIOP interface. For instance, the existence of DII and DSI should not be required to allow testing of MinimumCORBA implementations.

The test configuration is depicted in Figure 2. The OUT is highlighted by a shadowed box. Although a client-server paradigm is also used, the test components reside on different interfaces. Namely, the test client communicates with the OUT over the vertical interface, i.e. the GIOP interface, while the test server has access to the OUT over the adaptation layer and the horizontal interface which is in this case the skeleton interface.




Figure 2 Test configuration for server side testing

The test client is not a CORBA client because it does not use the standard stub interface. It is a test component that emulates the behavior of the counter part of the OUT. It must not be an implementation of a protocol engine. More important is that it supports the observation and assessment of the functionality of the OUT. In this sense, the test client is an active test component.

The test server uses the skeleton interface, is therefore a CORBA application. As a test component, the test server provides control of the OUT, e.g. acceptance of a request issued by the test client. It is also an active test component as it generates replies that are communicated via the OUT to the test client.

The adaptation layer is introduced for the customization of interfaces used by the test server. For example, in case that DII and DSI are used, the adaptation layer maps them

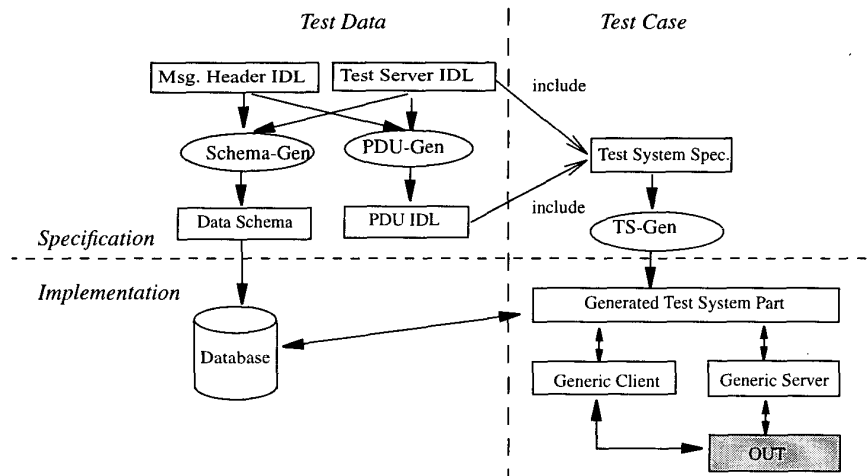


Figure 3 Test specification and implementation

to static interfaces of a MinimumCORBA ORB, while for a full CORBA ORB it delivers the calls one-to-one to its upper and lower layers.

The activities of the test client and test server need to be coordinated, e.g. for the timing of their setup, for parameterization of messages or operation invocations, or for the assignment of final test results. This can be done either manually or automatically. For the latter, a test manager can be used. It is less error-prone and more efficient.

The configuration shown in Figure 2 is used for testing the server side ORB. For the client side ORB, a second test configuration needs to be mentioned, in which the locations of the test client and the test server are swapped. The stub interface takes the place of the skeleton interface. With both test configurations we are able to verify the complete functionality of GIOP.

In the following, we continue to take the first configuration as the example but consider the testing methods so that they are applicable in both cases.

4. Specification of tests

After the test configuration is determined, its refinement in terms of test specification is considered. We propose to use high-level description in formal notations, in order to:

- allow unambiguous definition of test purposes;
- make test cases readable;
- reuse existing high-level data definitions;
- be abstract of target programming languages;
- enable high quality and efficient test development by automated code generation.

In particular, for adequate testing of GIOP traffic in terms of data transportation, we separate the focus of test specification into two major parts: Test Data and Test Case, as represented by Figure 3. It improves the modularity of the test system. Furthermore, a specification-based implementation concept is proposed. In this approach, test data covers the data structure used by the test case specification, and the data schema of the database used by the run-time of the test case implementation. The latter is composed of three parts. One part is generated from the Test System Specification. The other two parts, called Generic Client and Generic Server, correspond to the adaptation layer represented in Figure 2.

In this paper, we focus on the level of specification only. Some ideas to the test implementation can be found in Section 5.

In Section 4.1, we first elaborate the definition of the GIOP PDUs used by the test cases in CORBA IDL (abbr. IDL only). It is generated from the IDL definitions of the test server and the IDL definitions of the modules GIOP and IOP in the CORBA specification.

In Section 4.2, we proceed to the data schema that is built up from the same input, but using different transformation rules.

The specification of test cases includes the PDU data structure to have run-time access to the database that is derived from the data schema. It includes also the test server IDL for the description of the test component that resides on the CORBA APIs. This point is discussed in Section 4.3.

4.1. Data structure

As shown in Figure 3, the PDU IDL is generated from the test server IDL and the message header IDL. The




Figure 4 Request message structure

generator is called *PDU-Gen*. The test server IDL is designed by the test developer according to the test purposes, e.g. to verify the transport of different IDL typed values. The message header IDL consists of the modules *GIOP* and *IOp* that are defined in the CORBA specification.

The following is an excerpt that describes the structure of a *GIOP 1.2* message header:

```
// Msg. header IDL
module GIOP {
    struct Version {
        octet major;
        octet minor;
    };
    struct MessageHeader_1_2 {
        char magic[4];
        Version GIOP_Version;
        octet flags;
        octet message_type;
        unsigned long message_size;
    };
    struct RequestHeader_1_2 {
        // ...
    };
    // ...
};
```

For the current *GIOP* version, eight messages (PDUs) are defined: *Request*, *Reply*, *CancelRequest*, *LocateRequest*, *LocateReply*, *CloseConnection*, *MessageError* and *Fragment*. A *GIOP* PDU is composed of a general header, a specific header and a specific body. For example, a *Request* PDU consists of an instance of *MessageHeader_1_2*, followed by an instance of *RequestHeader_1_2*, followed by an instance of the *Request* body that consists of a sequence of the actual parameters that are defined in the signature in the test server IDL. The on-the-wire representation is determined by applying the *Common Data Representation* (CDR) rules to the sequence of those IDL definitions. Figure 4 gives a graphical depiction of the structure of a *Request* PDU according to *GIOP 1.2*.

Since large portions of a *GIOP* PDU are already described in IDL, we use IDL as the language to describe the PDU structure for test cases.

Consider the following example:

```
// Test server IDL
interface TestServer {
    void foo (in short x, inout long y,
             out double z);
    // ...
};
```

The operation *foo* has an input parameter of type *short*, an inout parameter of type *long* and an output parameter of type *double*. The body of the *Request*-PDU will contain the *short* and *long* parameters and the body of the *Reply* PDU will contain the *long* and *double* parameters respectively (an inout parameter is transferred in both directions between client and server). Using IDL, the structure of the bodies of the *Request* and *Reply* PDUs can be described as follows:

```
// IDL
struct foo_Request_body {
    short x;
    long y;
};
struct foo_Reply_body {
    long y;
    double z;
};
```

The CDR is used again to marshal the actual parameters for the *Request* and *Reply* bodies. *GIOP*-Replies make a distinction between normal and abnormal execution. The abnormal execution is indicated through an exception. This information can be coded in a similar way as shown above for normal execution.

For each of the operations of the test server interface, several structs are derived that each describes the content of a specific *GIOP* PDU. For example, the operation *foo* mentioned above yields the following structs:

```
// PDU IDL
struct foo_GIOP_Request {
    GIOP::MessageHeader_1_2 giop_header;
    GIOP::RequestHeader_1_2 request_header;
    foo_Request_body request_body;
};
struct foo_GIOP_Reply_normal {
    // ...
};
struct foo_GIOP_Reply_exception {
    // ...
};
```

Repeating this scheme, all operations of the test server interface can be translated into a set of IDL specifications that describe the logical structure of all the GIOP PDUs that will be exchanged between the test client and the OUT.

4.2. Data schema

The generated PDU structure definitions, as described in the previous section, are used by test cases to get access to the database at run-time. A data schema that is in-line with the IDL definitions eases both the specification of constraints and the implementation of data access.

However, the PDU IDL is not directly used to derive the data schema. Because some information in the input IDLs, which is relevant for the test data but irrelevant for the test case specification, is no more contained in the generated PDU IDL, e.g. the attribute of an operation (normal or oneway), or the relations between message bodies of the same operation. Therefore, the generator for the data schema *Schema-Gen* takes the original IDLs as input.

Further, using an IDL-related language would be beneficial, because only minimal mappings of language constructs are required. We propose to use the Object-Definition Language (ODL) defined by the Object Database Management Group (ODMG). ODMG 3.0 [1] is the most recent specification of this consortium. It defines an object database framework, containing the data schema language ODL and the Object Query Language (OQL).

ODL provides *interface* and *class* types for objects. Interface types are used for generalization. Class types are used to directly instantiate objects. Objects have identifiers and names to refer to each other. ODL provides notations to specify relationships between objects. In particular, the referential integrity is guaranteed.

ODL is designed to be a super-set of IDL. Therefore it supports most of the basic types of IDL. In addition, collection objects are supported by the so-called type generators: *set*, *bag*, *list*, *array* and *dictionary*. Elements of collection objects are of the same type. Mostly used are *set* objects - collections of unordered, non-duplicated elements, and *list* objects - ordered collections of objects.

The data schema specification using ODL is guided by the following goals:

- Represent the complete information contained in the original test server IDL and the message header IDL.
- Take care of aggregation relations and other dependence relations between data.
- Build data objects in a modular and hierarchical manner to allow flexible use.

As the GIOP interface provides the transparent communication between CORBA clients and servers, it is

adequate to store values needed both by the test client and the test server in PDU structures.

The database does not store complete PDUs, such as *foo_GIOP_Request*, but data objects in a hierarchical structure that allows flexible combination of PDU fields. It reduces redundancy of storage space. The top-level structures are general message headers, specific message headers and specific message bodies. Using constraints, PDUs can be constructed on-the-fly.

In IDL, a PDU part, e.g. *MessageHeader_1_2*, is defined by a *struct* type. Its fields are represented by *struct* members. There is an aggregation or composition relation between a *struct* type and its members. Each member may have different values. From testing perspective, these values may be valid, i.e. values that in general lead to normal behavior, or invalid, i.e. values that in general cause exceptional or forbidden behavior. In addition, there can be also relations between members of a structured type. Therefore, we propose an IDL to ODL mapping described in the following.

An IDL *struct* type is mapped to an ODL *class* type. Members of the *struct* type are mapped to attributes of the class type. This is to protect the aggregation relation between a *struct* and its members. Below is a fraction of translated ODL definition for the message header IDL in Section 4.1 (see also Figure 5).




Figure 5 Aggregation relation

```

// ODL
module ODMG_GIOP {
    interface DataObj {
        octet flag;
    }
    class MessageHeader_1_2 : DataObj {
        MessageHeader_magic magic;
        Version GIOP_version;
        MessageHeader_flags flags;
        MessageHeader_message_type
        message_type;
        MessageHeader_message_size
        message_size;
    };
}
  
```

```

class Version : DataObj {
    Version_major major;
    Version_minor minor;
};
class Version_major : DataObj {
    list<Version_major_value> valuelist;
};
class Version_major_value : DataObj {
    octet value;
}
class RequestHeader_1_2 : DataObj {
    //...
};
//...
};

```

To allow referencing of individual values, for each basic type member of an IDL struct, two ODL class types are derived, e.g. `Version_major` and `Version_major_value`. The first class contains a list type attribute named `valuelist` to maintain alternative values of a PDU field by object references. The second class is to store a concrete single value of the PDU field. It has a `value` attribute of the type that corresponds to the original IDL type, e.g. `octet` for `Version_major_value`.

For data usage, an interface named `DataObj` is defined that is inherited by every class type derived from IDL definitions. It contains an attribute `flag` of type `octet`, which can be used to denote *valid/invalid* values for simple fields, or *normal/exceptional* values for message bodies.

It is important that each value instance is accessible by referencing object instances for building dependence relations between data, as discussed below.

Figure 6 illustrates a simplified view on the dependence relations (represented by arrows with dotted lines) in a Request message. The complete structure is depicted in Figure 4.

Within the Request header, the values of fields `Operation`, `Resp_flag` and `IOR` of `Target` (actually the type id of the operation interface contained in the `IOR`) are dependent from each other. The dependence is determined by the signature of the operation defined in the test server IDL. Also, the relation between the `Operation` value and the corresponding Request body is defined by the operation signature, e.g. `opl` and `opl_Req_body`.

On the other hand, the relations between message bodies for a given operation, are determined by the semantics of the test server interfaces.

In ODL, the identified dependence relations are represented by `relationship` properties. For Figure 6, we may have the following ODL definitions, in which the class `ReqHdr_operation_value` and the class `ReqHdr_resp_flag_value` reference to each other by bi-directional relationship properties. A pre-defined interface `RequestBody` is used by specific Request bodies to refer to corresponding operation names.




Figure 6 Dependence relations

```

// ODL
module ODMG_GIOP {
    interface RequestBody {
        relationship ReqHdr_operation_value op
        inverse ReqHdr_operation_value::req_bd;
        relationship ReplyBody_normal rpl
        inverse ReplyBody_normal::req;
        relationship ReplyBody_exception ex
        inverse ReplyBody_exception::req;
    };
    interface ReplyBody_normal { //... };
    interface ReplyBody_exception { //... };
    class RequestHeader_1_2 : DataObj {
        ReqHdr_request_id request_id;
        ReqHdr_resp_flags resp_flags;
        ReqHdr_reserved reserved;
        TargetAddress target;
        ReqHdr_operation operation;
        ODMG_IOP::ServiceContextList
            service_context;
    };
    class ReqHdr_resp_flag : DataObj {
        list<ReqHdr_resp_flag_value> valuelist;
    };
    class ReqHdr_resp_flag_value : DataObj {
        octet value;
        relationship ReqHdr_operation_value op
        inverse ReqHdr_operation_value::flg;
    };
    class ReqHdr_operation : DataObj {
        list<ReqHdr_operation_value> valuelist;
    };
}

```

```

class ReqHdr_operation_value : DataObj {
    string value;
    relationship ReqHdr_resp_flag_value flg
    inverse ReqHdr_resp_flag_value::op;
    relationship RequestBody req_bd
    inverse RequestBody::op;
    //...
};
class op1_Req_body : DataObj, RequestBody {
    //...
};
//...
};

```

4.3. Test cases

The test case specification is the second major part of a test specification. The test case behaviour describes the exchange of test events with the OUT. Since GIOP is a message-based protocol, messages are exchanged only as test events. In black-box testing, the basic principle of testing is that the test case sends a stimulus to the OUT, awaits in general a set of possible reactions from the OUT (including also unexpected ones), compares received reactions (including also the absence of a reaction) with the expected ones, and decides on the basis of this comparison how to proceed with the test cases or which verdict to assign (and to terminate the test case). The timing for the exchange of messages is controlled with timers: that a reaction from the OUT shall occur within a certain amount of time, that a reaction from the OUT shall be delayed by a certain amount of time, or that no reaction shall occur at all.

In the past, a methodology and framework for testing distributed systems, known as the *conformance testing methodology and framework* (CTMF) [5], was developed and internationally standardized. CTMF covers all aspects of testing distributed systems, such as test suite specification, test notation (TTCN - Tree and Tabular Combined Notation), test implementation and test execution. CTMF has been successfully applied to a wide range of systems, e.g. from e-mail systems and directory services to management and IN systems.

Currently, the third edition of TTCN (TTCN-3) [4] [4] [9] has been developed to address testing needs of modern telecom and datacom technologies and to widen the scope of its applicability. TTCN-3 is a text-based language for the specification of tests for reactive systems in general.

TTCN-3 is on syntactical (and methodological) level very different to previous TTCN versions. However, the main concepts of TTCN have been retained and improved, and new concepts have been included, so that TTCN-3 will be applicable for a broader class of systems. New concepts are, e.g. a test execution control program to describe relations between test cases such as sequences, repetitions and dependencies on test outcomes, dynamic concurrent test configurations, and test behaviour in asynchronous and

synchronous communication environments. Further improved concepts are, e.g. the integration of ASN.1, the module and grouping concepts to improve the test suite structure, and the test component concepts to describe concurrent test setups.

A test specification in TTCN-3 is included in a module, which declares all objects such as types, timers, test components, which are used to define the test cases and their execution. The `GIOP-Tests` module described below imports data type definitions in IDL for test server and GIOP PDUs as introduced in Section 4.1.

For the GIOP tests we use a test configuration with three test components (see also Figure 7): a test component for the vertical interface (`TestClient`), a test component for the horizontal interface (`TestServer`), and a master test component (MTC) for the overall control of the other two test components over the coordination points (CPs). The interface of the OUT is defined by a separate component type definition for OUT containing the vertical (VG) and horizontal (HG) interfaces.




Figure 7 Test component configuration

```

// TTCN-3
module GIOP-Tests {
    import all from GIOP language IDL;
    import all from TestServer language IDL;
    type component MTC {
        port Coord CP;
    }
    type component TestClient {
        timer T_Response:= ...;
        timer T_NoResponse:= ...;
        timer T_Wait:= ...;
        port Coord CP;
        port VerGIOP VG;
    }
    type component TestServer {
        timer T_Response:= ...;
        timer T_NoResponse:= ...;
        timer T_Wait:= ...;
        port Coord CP;
        port HorGIOP HG;
    }
    type component OUT {
        port HorGIOP HG;
        port VerGIOP VG;
    }
    ...
}

```


The master test component executes the test cases which are defined as the following example shows. The master test component uses the component type definition MTC (keyword `runs on`) and can be executed for OUT, which have interfaces according to the OUT component type definition (keyword `system`).

```
// TTCN-3
testcase ReqMessage_T1
(fRequest: foo_GIOP_Request,
 fRequestReply: foo_GIOP_Reply_normal, ...)
runs on MTC
system OUT
{
  var TestClient TC := TestClient.create;
  var TestServer TS := TestServer.s;
  connect(mtc:CP, TC:CP);
  connect(mtc:CP, TS:CP);
  map(TC:VG, system:VG);
  map(TS:HG, system:HG);
  activate(Default_1);
  TC.start(ReqMessage_T1_TC
    (fRequest, fRequestReply));
  TS.start(...);
  all.done;
  log ("successful termination");
  stop;
}
```

Initially, it creates the other two test components: TC of type `TestClient`, and TS of type `TestServer`; activates a default `Default_1` to capture unexpected events during test execution; starts TC and TS with test behaviour, e.g. `ReqMessage_T1_TC` for TC (see also below); and awaits the termination of the created test components (`all.done` is blocking until all test components have terminated) in order to collect their individual test verdicts and to assign the overall test verdict. Test verdicts are *pass* (if the observed behaviour validates the test purpose of a test case), *fail* (if the observed behaviour disproves the test purpose of a test case), *inconclusive* (if the observed behaviour leads neither clearly to pass or fail), or *error* (for the exceptional cases of run-time errors in the test system).

Verdicts are collected by a TTCN-3 specific mechanism: each test component (i.e. all three in our case) have a local test verdict. This verdict is treated according to the “get-only-worse” rule: whenever during the execution of this test component an inconclusive or fail is assigned, the overall verdict of the test cases is at most inconclusive or respectively fail.

The behaviour of the test components `TestClient` and `TestServer` are defined in terms of functions, for instance, `ReqMessage_T1_TC` for `TestClient` (see below). After activation of a default (in order to make the test behaviour robust for unexpected responses from the OUT), a request message `fRequest` is sent to the OUT via port `VerGIOP`. A timer is started in order to prevent from infinite waiting.

In an alternative statement (keyword `alt`), the various responses are expected at port `VerGIOP`. Only in the case that the response is correct, a pass will be assigned. In the other cases (i.e. wrong response, e.g. different to the template `fRequestReply`, or no response, e.g. timeout), a fail will be assigned.

```
// TTCN-3
function ReqMessage_T1_TC
(fRequest: foo_GIOP_Request,
 fRequestReply: foo_GIOP_Reply_normal)
runs on TestClient
{
  activate (Default_2());
  VerGIOP.send(fRequest);
  T_Response.start;
  alt {
    [] VerGIOP.receive(fRequestReply)
    { set.verdict(pass); ... }
    [] VerGIOP.receive
    { set.verdict(fail); stop; }
    [] T_Response.timeout()
    { set.verdict(fail); stop; }
  }
}
```

The control part is used to denote the successive execution of test cases, to define the parameterization of test cases (using the data access described in the previous section) and to make the execution of test cases dependent on the outcome of previous test cases.

While message types (the structure for a test event exchanged between test system and OUT) are included from imported PDU IDL definitions, message templates (the concrete values or value constraints for a test event) are specified as the following examples `foo_req_v` and `foo_rpl_v` show.

A template definition aligns to the PDU IDL definition. It may use concrete values or symbolic values. The use of symbolic values in templates extends the current TTCN-3 definition[4]. The symbolic value `any` refers to any available values for the appropriate field. `any` is further constrained by a flag which can be `valid` or `invalid`, `normal` or `exceptional` (see also Section 4.2).

The templates are used by external functions e.g. `get_fooRequestMessage` to get access to the database.

```
// TTCN-3
template foo_GIOP_Request foo_req_v
:= {
  giop_header any.valid,
  request_header any.valid,
  request_body.x any.normal
}
template foo_GIOP_Reply_normal foo_rpl_v
:= {
  giop_header any.valid,
  reply_header any.valid,
  reply_body_normal any.normal
}
```

```

external function get_fooRequestMessage
(in template foo_GIOP_Request foo_req_v,
 out foo_GIOP_Request curReq)
return Boolean;

external function get_fooReplyNormalMessage
(in template foo_GIOP_Reply_normal
 foo_rpl_v,
 in foo_GIOP_Request curReq,
 out foo_GIOP_Reply_normal curRpl)
return Boolean;

control {
  var foo_GIOP_Request curReq;
  var foo_GIOP_Reply_normal curRpl;
  while (get_fooRequestMessage
    (foo_req_v, curReq)) {
    get_fooReplyNormalMessage
    (foo_rpl_v, curReq, curRpl);
    execute(ReqMessage_T1
    (curReq, curRpl, ...), 20);
  }
}

```

The execution of the test case is defined in the `execute` statement. There, the concrete data is bound to the test case. The execution of the test case is limited to 20 seconds only. If the overall verdict cannot be assigned within that time, an error verdict will be assigned instead.

5. Conclusions

CORBA consists of a set of specifications that define the portability and interoperability of applications in a heterogeneous, distributed environment. With respect to cross-platform development we have argued that interoperability plays a more important role than portability, as it helps to decouple technological domains. The key to CORBA interoperability is the conformance of CORBA implementations. In this paper we have presented a new approach to test an ORB's conformance of its implementation of the interoperability specification. The test system accesses the ORB under test directly at its GIOP interface, which allows for greater flexibility when generating GIOP traffic. We have presented a framework that separates the specification from the implementation of the test system. ODMG-ODL is used to describe the data schema for the PDUs used during testing, while TTCN-3 is used to specify the behavior of the test cases.

Some aspects that have not been mentioned in this paper need further investigation. Among those are the population

of the ODMG-based database and the design of the generic client and server that act as the interface between the ORB under test and the test system. It is planned to implement the concepts presented in this paper. To underline the importance of conformance testing of interoperability, it is intended to release the code of the tests under an open source license. The tests will be contributed to the COST (CORBA Open Source Test Suite) [8] initiative currently organized by the OMG.

6. References

- [1] R.G.G. Cattell, et al., *The Object Database Standard: ODMG 3.0*, Morgan Kaufmann Publ., Inc., 2000.
- [2] R. Elmasri, S.B. Navathe, *Fundamentals of Database Systems*, Addison-Wesley, 2000.
- [3] ETSI (European Telecommunications Standards Institute) ES 201 873-1, *Methods for Testing and Specification; The Tree and Tabular Combined Notation version 3 (TTCN-3); Part 1: TTCN-3 Core Language*, Sophia-Antipolis, France, 2001.
- [4] J. Grabowski, A. Wiles, C. Willcock, D. Hogrefe, *On the Design of the New Testing Language TTCN-3*, Int. Conf. on Testing of Communication Systems (TestCom 2000), Ottawa, Canada, Aug. 2000.
- [5] ISO/IEC 9646, *Information Technology - Open Systems Interconnection - Conformance Testing Methodology and Framework*, 1997.
- [6] M. Li, A. Rennoch, I. Schieferdecker, D. Witaszek, O. Halabi, A. Vouffou, A. Yin, *Experience report on Conformance Tests for CORBA ORBs*, submitted to the Second Asia-Pacific Conference on Quality Software (APQS'01), Hong Kong, Dec. 2001.
- [7] Object Management Group, *Common Object Request Broker Architecture (CORBA)*, ver. 2.4, Feb. 2001.
- [8] Object Management Group, *COST (CORBA Open Source Test Suite)*, public document, <http://cost.omg.org>.
- [9] I. Schieferdecker, S. Pietsch, T. Vassiliou-Gioles, *Systematic Testing of Internet Protocols - First Experiences in Using TTCN-3 for SIP*, 5th. Africa Conf. on Communications (Afri-com'01), Capetown, May 2000.
- [10] The Open Group, *CORBA Verification Suite, User's Guide*, ver. 1.1.1, Sep. 1999.
- [11] A.R. William, T. Herron, P. Klinker, *IIOP Complete: Understanding CORBA and Middleware Interoperability*, Addison-Wesley, 1999.