

IDL*flex*: **A flexible and generic compiler for CORBA IDL**

Hans Reiser, Martin Steckermeier, Franz J. Hauck
Informatik 4, University of Erlangen-Nürnberg
{reiser, mstecker, hauck}@informatik.uni-erlangen.de

Abstract. For the development of CORBA applications, an IDL compiler is needed that generates code for communication stubs, helper classes and implementation skeletons. For each IDL language mapping, for every version of a particular language mapping, and for every CORBA implementation, the generated code has to be different. Traditionally the code generation is hard-wired into a compiler, thus many different compilers have to be programmed. *IDLflex* is a generic IDL compiler which is able to generate arbitrary code for arbitrary languages. Only a mapping program written in an XML-based mapping language and a language-specific utility class have to be provided. Thus, *IDLflex* can be adapted to another target programming language, to another mapping or to another ORB implementation in a very fast way. Furthermore, *IDLflex* allows to easily integrate additional functionality into a CORBA-based system, as it was done within the *AspectIX* middleware project.

1 Introduction

For the development of distributed applications, middleware platforms are used which support communication between the distributed components of an application. One such middleware is the Common Object Request Broker Architecture (CORBA) [5]. CORBA is an architecture that is realized by a particular CORBA implementation. Traditionally there are many vendors distributing different CORBA implementations.

CORBA allows the development of applications in terms of distributed objects, which can be written in various programming languages. To hide an object's implementation language from its clients, CORBA provides an Interface Definition Language (IDL) which is used to define the interfaces of distributed objects. For each programming language, a language-mapping standard defines how IDL types are mapped to language types, how parameters are passed, how the interface of an object looks like, how the client can use object references, etc.

Communication in CORBA is defined in terms of remote object invocations. Clients can hold object references and invoke methods at the referred object. The invocation of a method will finally invoke a method at the possibly remote server object. For the implementation of remote invocations, CORBA uses a client-side stub. The stub code depends on the corresponding IDL-based interface definition, the language mapping and on the particular communication mechanisms used by a CORBA implementation. The stub code is automatically generated by a code-generator tool, usually called an IDL compiler.

First, an application developer describes the interface of a distributed object in IDL. Second, he runs the IDL compiler for the chosen ORB implementation and target programming language. The IDL compiler generates a so-called skeleton which is a code frame for the object implementation. The application developer has to insert the actual object code into the skeleton in order to create the object implementation. For the client side of a remote object, the IDL compiler generates code for the necessary stub objects. However, the chosen programming language for stub objects may be different from the programming language for the object implementation, as CORBA supports heterogeneous programming environments; different IDL compilers may be used to generate either code. Beside skeleton and stub code, an IDL compiler generates code for auxiliary programming elements that are prescribed by the corresponding language mapping (e.g., the Java language mapping defines holder and helper classes for certain types [4]).

Thus, an IDL compiler is an important part of each particular CORBA implementation. Usually, such compilers are not very flexible, because the code generation is hard-wired into the compiler. For each supported programming language there needs to be a different compiler. With every new version of a language mapping the compiler code has to be adapted. Different vendors may have to use different compilers because the communication code inside of stub objects is usually vendor-dependent.

IDLflex is a novel approach to IDL compilers. The IDL-specific part is hard-wired into the compiler code whereas the mapping-specific code is generic and can be externally programmed and influenced in a very broad way. Thus, the compiler can run different mapping programs without any need to change the compiler itself. As the configuration of *IDLflex* is very easily changed, *IDLflex* is also very appropriate if ORB developers have to experiment with different generated code. The driving force behind *IDLflex* has been the *AspectIX* middleware project. *AspectIX* extends CORBA by generic quality-of-service support, and the implementation of *AspectIX* needs completely different stub objects than standard CORBA. So, *IDLflex* helped to design the stub objects for *AspectIX* and became the *AspectIX* IDL compiler.

This paper is organized as follows: Section 2 introduces the implementation of *IDLflex*. In Section 3, we outline the advantages of *IDLflex* demonstrated by examples of the IDL-to-Java language mapping. Section 4 compares *IDLflex* to related work. Section 5 will give our conclusions. Finally, in an appendix we will briefly define *IDLflex*'s mapping and configuration language.

2 Implementation

The internal structure of *IDLflex*'s implementation is outlined in Fig. 1. The two parts on the left side are responsible for reading in the IDL description and generating an internal representation of one or more interfaces that shall be processed by *IDLflex*. The third part, on the right side, will process the IDL descriptions and generate code.

IDLflex is entirely written in Java and thus portable to every platform that provides a Java Virtual Machine.

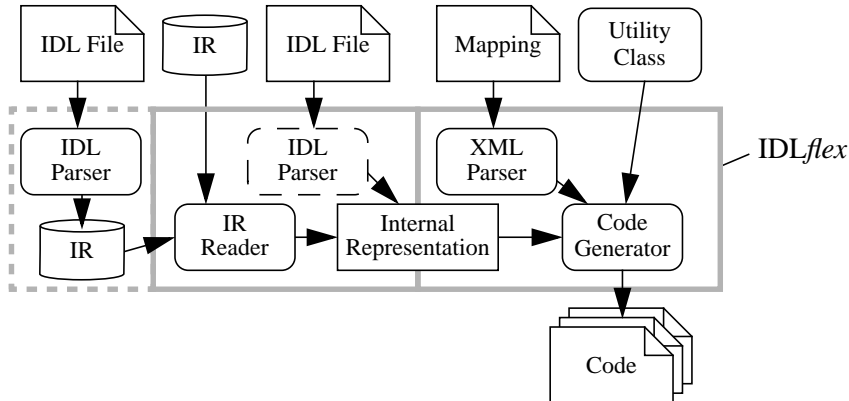


Fig. 1. Internal structure of *IDLflex*.

2.1 Reading IDL

As a first step, the IDL descriptions that shall be processed by *IDLflex* have to be read into the compiler. An internal data representation of an IDL interface has to be generated. This internal representation will finally serve as the basis for code generation.

In our first version of *IDLflex*, we retrieve IDL descriptions from an interface repository. The interface repository (IR) stores interface descriptions in form of CORBA objects. The types and the behavior of those objects are defined by the CORBA standard. As an interface repository is necessary for other CORBA mechanisms like type checking and the Dynamic Invocation Interface (DII), it is usually part of any CORBA implementation. A reader component inside of *IDLflex* can read an existing interface repository, retrieve the IDL information, and convert it into the internal representation needed for code generation. So, the IDL description of an interface first has to be loaded into an interface repository and then be processed by *IDLflex*.

As this is cumbersome in certain cases, we incorporated an existing interface repository including a corresponding IDL parser into *IDLflex* (see the left hand, dashed part in Fig. 1). The parser can read files containing IDL descriptions and feed them into the interface repository, which finally is read by the IR-reader component. In our implementation we used the interface repository and IDL parser of the open source CORBA implementation *JavaORB 2.2*¹.

Reading from an interface repository has the disadvantage that it can only process IDL descriptions that are usually stored there. For example, the interfaces of CORBA pseudo objects are not stored in an interface repository. Pseudo objects are used within the CORBA standard to make internal mechanisms and services appear as CORBA objects. The indirect reading of IDL descriptions via an internal repository also wastes resources. Thus, the next version of *IDLflex* will have an additional internal IDL parser that can directly read IDL files and convert them into the internal representation (see the dashed component in the middle part of Fig. 1).

1. This ORB is no longer supported by its developers. However the *JavaORB* technology migrated to *OpenORB*, also an open source project.

The internal data representation is similar to the representation inside of an interface repository. However, certain things have been simplified for the later processing. As an alternative we could have decided to directly generate code from the objects inside of an interface repository. We ruled this out due to the complex interface to the repository objects and for efficiency reasons.

In the internal IDL representation, every IDL item is represented by a Java object. So, IDL modules, interfaces, operations, attributes, exceptions, etc. are represented each by a particular object. All these objects inherit from a common base class called `IDLObject`, which provides generic access to retrieve all the necessary information (like kind of IDL item, name, type, etc.) and to de-reference child objects (like members of an interface or parameters of a method). The `IDLObject` interface is designed to be easily usable from the XML-based mapping language. A simple IDL interface definition of a `sumserver` with one method named `add()` is shown in Fig. 2.

```

interface sumserver
{
    void add( in long s1, in long s2, out long res );
}

```

Fig. 2. A simple IDL interface.

The interface `sumserver` is internally represented by an object of type `InterfaceObj`. Such an object contains references to other objects representing either members of the interface and inherited interfaces. In this example, just one reference is contained, a reference to an object of type `OperationObj` representing method `add()`. The operation object in turn contains references to objects representing its return type and its parameters. Parameter objects contain references to objects representing their types. The complete representation of the example interface is shown in Fig. 3. The internal objects also contain attributes to distinguish abstract and local interfaces as well as `in`, `inout` and `out` parameters, etc.




Fig. 3. The internal representation of interface `sumserver`.

2.2 Code Generation

As the code generation shall be generic there is no built-in code for code generation. Instead *IDLflex* provides an interpreter for a simple programming language. In this language the specific code generation depending on the IDL description fed into *IDLflex* can be programmed. This program is called a mapping.

The programming language to describe mappings is XML-based. The language defines statements as XML tags whereas ordinary text is usually directly written to the output code. There are three types of statements that are supported: output control, control flow and access to the internal data representation.

The output control allows to specify the names of output files where the generated code will be written. Usually a CORBA language mapping specification defines, which files have to be generated. With the output control, the mapping program can generate code for and switch between multiple files. Furthermore, the output control allows to write code to internal string variables, which can later be inserted at arbitrary places by the mapping program. This is useful for code that has to be inserted multiple times.

Control flow statements allow to group mapping code into so-called components that can be invoked like sub-routines. Thus, common mapping code can be extracted and put into an own component. This serves for modularity and code reuse. Parameters can be passed to components by filling the above mentioned string variables. Furthermore, conditions can be tested at run-time and according to such a condition multiple control flows can be chosen in the mapping program. Often the same mapping has to be done for a set of objects of the internal representation (e.g., for all members of an interface). A special statement allows to iterate over such sets.

Another statement allows to retrieve information from an object of the internal representation. At any time the flow of control inside of the mapping program is implicitly associated with one of the objects of the data representation. This is similar to the implicit *this* pointer in object-oriented languages. The mapping program starts with the root object of the representation, representing an IDL module containing other modules and interfaces. With the access statement (XML tag `GET`), it is possible to retrieve the name, type and other information from the representation object. With different values for an XML attribute of tag `GET` the actual information is named. Access to other objects of the internal representation is possible by explicit de-referencing which is allowed for certain cases. Implicit de-referencing takes place when the iteration operator is used. Inside of the loop, the implicit representation object is always switched to another member of the iteration set, e.g., to the next member of an interface or to the next base interface.

As some mapping mechanisms need complex algorithms it is not reasonable to program these algorithms using the statements of our mapping language. Therefore, those algorithms can be provided in form of a utility class. Some defined methods of the utility class can intercept attribute values of the `GET` tag and thus replace this XML tag by arbitrary strings that are put instead into the output code. As an example, the utility class for the Java language mapping will provide the correct file name for an IDL interface as prescribed by the Java language mapping. Another example is the conversion of IDL names to language names. This procedure has to especially handle language keywords and other reserved names.

The utility class is usually used to encapsulate language-mapping-specific algorithms. As those algorithms are defined in the mapping standard there is hardly any need to change the utility class if it is already available for a certain target language. The utility class is dynamically bound to *IDLflex*. The class name is defined in the mapping program. The mapping program that is to be executed by *IDLflex* can be configured as

a command-line parameter. The mapping program is parsed by an XML parser and read into memory as a DOM element tree. *IDLflex* uses the *Apache Xerxes* XML library. The mapping-program interpreter operates on the DOM elements. With the execution of the mapping statements in the mapping program, output code is generated. After completion of the mapping program *IDLflex* terminates.

2.3 Example

As a brief example of the code generation, Fig. 4 shows three components of the mapping program for the IDL-to-Java language mapping. The first component, named `OperationCompiler` is called whenever an operation has to be mapped to a Java method declaration. This is necessary both for stub objects and for skeletons. We assume that the output file is already selected by the calling component. Ordinary text, like the word `public`, is then directly written to the selected output file. The next `GET` tag de-references the return type from the implicit `OperationObj-Object` and retrieves the Java type declaration. The following `GET` tag retrieves the name of the implicit operation object. Then a loop is executed iterating over all parameter objects stored in the implicit operation object. Within the loop each parameter object once becomes the implicit representation object for the loop body. Inside of the loop another component for the parameter type definition is executed followed by the name of the parameter. Finally, all defined exceptions are processed in another loop by calling yet another component.

```

<!-- Processing an operation -->
<COMPONENT NAME="OperationCompiler">
  public <GET OBJ="RETURN" T="JAVA:TYPE:decl"/> <GET T="JAVA:TYPE:name"/>{
    <ITERATE NAME="PARAM">
      <CALL NAME="ParametersDefCompiler"/> <GET T="JAVA:TYPE:name"/>
    </ITERATE> )
    <ITERATE NAME="EXCEPT"><CALL NAME="ThrowsDefCompiler"/> </ITERATE>
  </COMPONENT>

<!-- Processing the parameters -->
<COMPONENT NAME="ParametersDefCompiler">
  <IF ATTR="LOOP:NotFirst">, </IF>
  <IF ATTR="IDL:inarg"><GET OBJ="BASE" T="JAVA:TYPE:decl"/>
  <ELSE/><GET OBJ="BASE" T="JAVA:TYPE:holder"/>
  </IF>
</COMPONENT>

<!-- Processing thrown exceptions -->
<COMPONENT NAME="ThrowsDefCompiler">
  <IF ATTR="LOOP:First">throws <ELSE/>, </IF><GET T="JAVA:TYPE:decl"/>
</COMPONENT>

```

Fig. 4. Mapping components for operations in the IDL-to-Java mapping.

The definition of the `ParametersDefCompiler` component shows two conditional statements of the mapping language. With `LOOP:NotFirst` a conditional statement can retrieve whether it belongs to the first run of a loop or not. The second statement checks whether the implicit parameter object is an in parameter or not.

Executing the component `OperationCompiler` on the internal operation object named `add` of our `sumserver` interface (see Fig. 2 and 3) leads to the output shown in

Fig. 5. For the out parameter `res` a holder type has to be inserted. The IDL type `long` is mapped to Java's type `int`. This mapping happens in the utility class which processes the request `JAVA:TYPE:decl`.

The output text is post-processed by a simple formatter that prunes unnecessary white space characters which are introduced as delimiters for mapping statements. This makes the output code more readable.

```
public void add( int s1, int s2, org.omg.CORBA.IntHolder res )
```

Fig. 5. Result of the `OperationCompiler` for operation `add`.

3 Advantages

The advantage of *IDLflex* lies in its flexibility. Existing mapping programs can be adapted to the needs of a certain CORBA implementation. One example is the generation of client-side stub objects. Stub objects have to look like the actual object but forward invocation request to the remote object on another host. Therefore, stub objects have to use communication mechanisms to contact the remote host, or more precisely, the remote object adaptor responsible for the remote object. These communication mechanisms are often vendor-specific; every CORBA implementation may use different, sometimes even incompatible, communication mechanisms. In this case, for each CORBA implementation just another, slightly different mapping program has to be provided.

In the Java mapping, communication mechanisms are entirely encapsulated outside of the stub object, so that stub objects of all vendors could look similar. However, the Java language mapping allows two different implementations of stub objects: one based on the dynamic invocation interface (DII) and another one based on output streams. With *IDLflex* one could just use different mapping programs for each implementation. Even better, a mapping program for *IDLflex* can contain both variants and the variant is chosen by a command line option at execution time of *IDLflex*. Fig. 6 shows the mapping code for the selection. A command-line parameter named `-DDIISub` selects the mapping for DII-based communication. Otherwise the stream-based mapping is selected. The two mapping variants are each defined in an own mapping component.

```
<IF ATTR="DEF:DIIStub">
  <CALL COMP="DIIStubCompiler"/>
<ELSE/>
  <CALL COMP="StreamStubCompiler"/>
</IF>
```

Fig. 6. Selection of two mapping variants.

IDLflex has also been used for creating stubs, and object references respectively, of the *AspectIX* middleware platform [3, 2]. *AspectIX* is a CORBA-based system, but integrates general quality-of-service requirements into distributed objects. To that end, client-side stub objects can carry object-specific code in *AspectIX*, because quality-of-service implementation often need not only code at the server- or object-side but also at the client side, e.g., for maintaining a communication link based on special protocols.

The client-side part of an *AspectIX* object is called a fragment. It consists of a generic part that can easily be generated by *IDLflex* by adapting the stub mapping of the standard Java mapping. The generic part forwards invocations to an object-specific part, called the fragment implementation. For that implementation, *IDLflex* can generate skeletons, similar to servant skeletons in plain CORBA. For the *AspectIX* project, *IDLflex* proved to be an efficient platform for experimenting with different code generation. Changes in the output code were very easy and the turn-around time for testing the newly designed output code, the new mapping program respectively, was very small.

4 Related Work

The only comparable project is *Flick* (Flexible IDL Compiler Kit) [1]. *Flick* is very flexible as it can deal with CORBA IDL, Sun RPC, Mach Messages and other interface description languages. Multiple back-end modules of the compiler kit allow to generate code for different languages and mappings. The back-end modules are explicitly programmed for a specific mapping. However, there is one back-end that allows to include a SGML-based mapping description into the code generation.

While *Flick* is a very complex system which urges the programmer to write C++ code for complex mapping changes, *IDLflex* has a lean design and a powerful but still simple mapping language allowing the description of CORBA IDL language mappings.

5 Conclusion

IDLflex is a flexible IDL compiler as it allows to specify the code generation for a specific language mapping including vendor-specific code by programming the output in an XML-based mapping language. Language-specific mapping constructs can be encapsulated in a utility class. However, we anticipate that the utility class is hardly ever changed for a specific language mapping, e.g., for Java. So the same IDL compiler can be used with different mapping programs for producing code for different languages and CORBA implementations.

IDLflex has been used to generate special object references of the *AspectIX* ORB architecture. Especially for experimenting with different *AspectIX* implementations, *IDLflex* was very helpful, because just the mapping file had to be changed and there was no need to adapt the IDL compiler.

6 Availability

The *IDLflex* software is available to the public and can be downloaded under the URL <http://www.aspectix.org/IDLflex/>. *IDLflex* has been put under the terms of the GNU Public License (GPL). The current version, as of July 2001, is beta software, and lacks some documentation. Currently we have a mapping program for the CORBA 2.3 Java language mapping. This mapping does not yet consider value boxes.

References

- [1] E. Eide, J. L. Simister, T. Stack, J. Lepreau: Flexible IDL compilation for complex communication patterns. *Scientific Progr.*, 7(3, 4), 1999, pp. 275-287.
- [2] F. J. Hauck, U. Becker, M. Geier, E. Meier, U. Rasthofer, M. Steckermeier: AspectIX: a quality-aware, object-based middleware architecture. *Proc. of the DAIS 2001 Conf.*, to appear 2001.
- [3] F. J. Hauck, E. Meier, U. Becker, M. Geier, U. Rasthofer, M. Steckermeier: A middleware architecture for scalable, QoS-aware and self-organizing global services. *Proc. of the USM '00 Conf.* (Munich, Sept. 12–14, 2000), LNCS 1890, Springer, 2000.
- [4] Object Mgmt. Group, OMG: *Java language to IDL mapping specification*. OMG Doc. formal/99-07-59, June 1999.
- [5] Object Mgmt. Group, OMG: *The Common Object Request Broker, architecture and specification*. Ver. 2.4.2, OMG Doc. formal/01-02-33. Framingham, Mass., Feb. 2001.

Appendix: The IDLflex mapping language

This section describes the XML elements of the mapping language, their attributes, and their semantics.

- <IDLflex> Root element of the XML-based language. Attributes: ROOT describes the component that processes the root of the internal representation objects. UTILITY specifies the class name for the mapping-specific utility class. WRITER specifies the class name for a class post-processing the output.
- <COMPONENT> Element that specifies a mapping component. Attributes: NAME specifies the name of the component.
- <CALL> Element executes the mapping code of a referenced component. Attributes: OBJ references an internal representation object that is used within the component code; if omitted the internal object is not re-assigned. For allowed values for the OBJ tag see the GET tag. NAME is the name of the executed component.
- <ITERATE> This elements iterates over a list of internal objects. Attributes: NAME denotes the kind of objects used within the loop. The object set depends on the current internal object. Examples are: MEMBER for the members of modules, interfaces and enums. ALLMEMBERS for the members of interface including inherited members. ELEMENTS for the members of structs, unions and exceptions. UNIQUE for a list of all unique union members. BASE for the immediate base interfaces of interfaces; ALLBASE for all inherited interfaces. EXCEPT for the list of exceptions, PARAM for the list of parameters of an operation.
- <IF> Starts a conditional control flow. Attributes: TYPE checks the type of the current representation object. OBJ selects another representation object for consideration. For allowed values see the GET tag. ATTR checks a condition defined in the utility class. Allowed conditions are:
 - DEF:name Checks for a command line argument -D name
 - LOOP:First True if the first representation object of an ITERATE element
likewise defined: LOOP:NotFirst, LOOP>Last, LOOP:NotLast
 - HAVE:RETURN True if an operation object has a non-void return type
 - HAVE:name True if the named element list (allowed names see attribute SPEC of
ITERATE) is nonempty.
 - IDL:readonly True if attribute object is read-only
 - IDL:abstractif True if interface is abstract
 - IDL:inarg, outarg, inoutarg Exactly one is true for parameter objects
 - IDL:wide, bound True if string object is wide or bound

IDL:unique True if union member is unique
 IDL:isdefault True if union member is default
 IDL:needsdef True if implicit default label necessary for unions

Conditions may be inverted with the !-operator and concatenated by the or-operator |.

<SWITCH> Element denoting a conditional expression with several branches. It may contain an arbitrary number of CASE elements and one optional DEFAULT element. The allowed attributes of SWITCH are the same as of the IF-tag. The DEFAULT element has no attributes. Only the first CASE with a true condition is evaluated, the DEFAULT element is used if no true condition was found.

<FILE> Redirects the following output text to a specific file. Attributes: SPEC selects a file. Allowed names are defined by the mapping-specific utility class. For the Java mapping the same names as for FILE:name of attribute T of tag GET are allowed.

<SBOX> Redirects the following output test to an internal variable. Attributes: NAME of the SBOX variable.

<GET> This is the generic tag to access external information, e.g., from the internal representation objects. Attributes: T denotes what is to be retrieved. OBJ denotes the internal representation object to be used for retrieval; if omitted the current object is used. OBJ may have the following content:

BASE	The type object of an array, attribute, constant, parameter, sequence, union or struct member
CONTAINER	The object the current object is contained in as a member
DISCR	The type object of the union discriminator
RDISCR	Same as DISCR but typedefs are resolved
RESOLVE	Referred internal object for typedefs
RETURN	The return type of an operation object
SUPER	The internal object for the enum of an enum member; the union or struct object of a union or struct member

Attribute T may have the following values (partially defined by the Java utility class):

DEF:name	Retrieves content of command line option -Dname
IDL:name	Retrieves information from an internal object. name could be:
bound	Size of an array, sequence or string
defdiscr	Smallest discriminator value of union
digits/scale	Digits and scale of a fixed width number
discr	Label number of union discriminator value
id/name/fullname	ID, name and qualified name of an internal object
length	Size of an array or sequence
value	Value of an enum member
CONSTVAL	Java code for the value of a constant object
LIST:COUNT:name	Number of elements of object lists (see ITERATE for allowed name)
LOOP:Count	Total number of list elements for an ITERATE loop
LOOP:Index	Index of the current ITERATE loop (counted from 0)
DISCR:num/sym	Numeric or symbolic reference to union discriminator of a union member
SBOX:name	Retrieves the content of an SBOX variable
FILE:name	Retrieves the file name for JAVA:TYPE:name class of the Java mapping.
JAVA:PACKAGEDEF	Java package declaration for current object
JAVA:PACKAGENAME	Name of surrounding Java package
JAVA:TYPE:name	Conversion for the Java language mapping. Allowed values for name:
decl	Java-mapped type declaration of internal object
helper/holder	Name of mapped helper/holder class
name	Java-mapped name of internal object
operationif	Interface name of operation interface
stub/skeleton/tie	Stub/skeleton/tie Java class name
signatureif	Interface name of signature interface