

Patterns of Component and Language Integration

Uwe Zdun
New Media Lab
Department of Information Science
Vienna University of Economics and
BA
Austria
zdun@acm.org

Integration is an important concern in many software systems. In this chapter, we present a number of patterns that are used to improve the integration of a system with components or code that is written in different languages than the system itself. Component integration is necessary when (foreign) components should be used within a system. The challenge of integrating components into a system is that heterogeneous kinds of components exist, perhaps without distinct interfaces or other component boundaries. The task of the component integration code is to provide suitable, stable invocation interfaces and to compose the components with the system. Sometimes, however, invocation and composition of components is not enough, but deeper language integration is required. Examples of what might be needed are automatic type conversions between languages, preserving the destruction order of the other language, automatic forwarding of invocations into the other language, and runtime integration of foreign language components. The patterns, presented in this chapter, are successful solutions in these areas. Each of these patterns plays an important role in different integration architectures. We will give examples from the areas of distributed systems, reengineering projects, scripting languages, and aspect-oriented programming systems. There are many other fields where the patterns are used as well.

Introduction

Software Integration

Integration of software systems is needed in many situations. Just consider two enterprise applications that need to be integrated after a fusion of the companies took place. Or, consider a reengineering project in which a terminal-based legacy system should be re-engineered to the Web – and thus must be integrated with a modern application server. The capability of the system's architecture to integrate foreign components and be integrated with other systems is thus

considered an important quality attribute [Bass+1998]. There are many integration concerns to be considered when integrating software architectures, including:

- integration of data, data models, and queries,
- integration of various software components,
- integration of various programming languages,
- integration of distributed systems,
- integration of system with legacy systems, and
- integration of different software paradigms such as the object-oriented and procedural paradigm.

In this chapter, we present patterns that describe solutions in the area of software component and language integration. These patterns are not a complete pattern language for the broad field of software architecture integration – rather, they are an excerpt working in the area of component, language, and paradigm integration mainly. There are also many patterns from other pattern collections that interact closely with the patterns presented in this chapter. (These other patterns are referenced in the pattern texts below).

The patterns presented in this chapter are applied where simple, ad hoc integration solutions fail. Thus the main audiences of the patterns are experienced software architects and developers.

Pattern Language Outline

In this section, we provide short thumbnails for the patterns and important external patterns, as well as a pattern map as an overview.

The following patterns are presented in this chapter:

- A **COMMAND LANGUAGE** is a language that is implemented within another language (called “host language”). It uses **COMMANDS** [Gamma+1995] implemented in that host language as implementations of its language elements. The **COMMAND LANGUAGE** lets the **COMMANDS** be assembled in scripts, and it provides an **INTERPRETER** [Gamma+1995] or another interpretation mechanism such as an on-the-fly compiler for these scripts.
- A **COMPONENT WRAPPER** is an object implemented as part of a component that represents another component. The wrapped component is imported into the wrapping component, and the **COMPONENT WRAPPER** handles all details of the import. For the integration it utilizes other patterns, such as **WRAPPER FACADE** [Schmidt+2000] or **PROXY** [Gamma+1995]. The **COMPONENT WRAPPER** is an accessible white-box representing the component within another component’s scope, where adaptations such as interface adaptations or version changes can take place.
- An **OBJECT SYSTEM LAYER** [Goedicke+2000] implements an object systems as a **LAYER** [Buschmann+1996] in another system. This way we can extend systems written in non-object-oriented languages with object-oriented concepts, integrate different object concepts, and extend existing object concepts. A part of an **OBJECT SYSTEM LAYER** is a **MESSAGE REDIRECTOR** [Goedicke+2001] that forms a **FACADE** [Gamma+1995] to the object system; all invocations into the **OBJECT SYSTEM LAYER** from the rest of the system are sent to its **MESSAGE REDIRECTOR**

and dispatched here to the objects in the object system.

- An AUTOMATIC TYPE CONVERTER converts types at runtime from one type to another. There are two main variants of the pattern: one-to-one converters between all type bindings and converters utilizing a canonical format to/from which all supported types can be converted. AUTOMATIC TYPE CONVERTER can, for instance, be used in a COMMAND LANGUAGE to convert types to/from the host language, in SPLIT OBJECTS to convert between the split object halves, and in COMPONENT WRAPPERS to convert invocations to the types of the wrapped component and vice versa.
- A SPLIT OBJECT is an object that is split across two languages. The SPLIT OBJECT is treated logically like a single instance, but it is implemented by two physical entities, one in each language. That is, both halves can forward invocations to the other half, access the inner state of the other half, and usually – to a certain extent – the user-defined class hierarchies of each half is mimicked by the respective other half.

There are a number of external patterns, documented elsewhere, that play an important role for the patterns described in this chapter. We want to explain some of these patterns briefly:

- A COMMAND [Gamma+1995] encapsulates an invocation to an object and provides a generic (abstract) invocation interface. COMMANDS alone only allow for adaptation by changing the association link to a COMMAND. In the COMMAND LANGUAGE pattern the COMMAND abstraction is extended with a binding of COMMANDS to language elements and an interpretation mechanism for that language.
- The important interpretation step in a COMMAND LANGUAGE can be implemented using another pattern from [Gamma+1995], the INTERPRETER pattern. In general an INTERPRETER defines a representation for a grammar along with an interpretation mechanism to interpret the language.
- The pattern INTERFACE DESCRIPTION [Völter+2004] describes how to specify interfaces of an object so that external clients of the object can reliably access the object. The pattern is originally described in the context of distributed systems: in this context an INTERFACE DESCRIPTION specifies the interface of a remote object for remote clients, but can also be used in other scenarios. Generating type conversion code from an INTERFACE DESCRIPTION is the most important alternative solution to using an AUTOMATIC TYPE CONVERTER. AUTOMATIC TYPE CONVERTERS can also be applied according to an INTERFACE DESCRIPTION.
- A MESSAGE REDIRECTOR [Goedicke+2001] receives symbolic invocations – for instance provided in form of strings – and redirects these invocations to objects implementing the invocation behavior. This way, for instance, a dynamic object dispatch mechanism can be implemented. A MESSAGE REDIRECTOR can also be used to integrate some non-object-oriented invocation interface with an object system. The MESSAGE REDIRECTOR pattern is an inherent part of an OBJECT SYSTEM LAYER: it is used to dispatch the invocations of OBJECT SYSTEM LAYER to its objects.

Figure 1.1 shows an overview of the patterns described in this chapter and the relationships explained above. In the pattern map the most important relationships of the patterns are represented by labeled arrows.

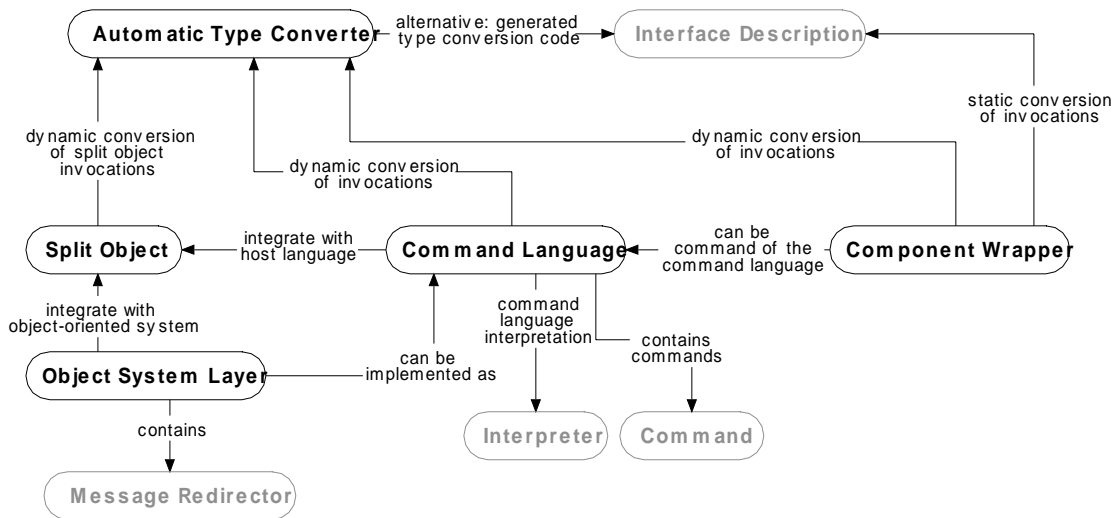


Figure 1.1
Pattern Interactions in the Pattern Language

Patterns for Component and Language Integration

In this section, we present the five patterns for component and language integration, summarized above. Please note that we first describe the patterns with only simple example scenarios and a few known uses. More detailed known uses and technology projections can be found in the next section, after the patterns have been described. These known uses and technology projections are used as examples for the patterns, instead of giving an example for each pattern on its own, because this way we can better describe the interactions of the patterns.

COMMAND LANGUAGE

Context

Some code needs to be executed by other code and needs to be exchangeable at runtime. In static languages you cannot simply exchange code. In such situations, COMMANDS provide a standard solution used in object-oriented and procedural systems: they encapsulate the implementation of an invocation and provide a generic, abstract invocation interface.

Problem

Using many COMMANDS without further support can be cumbersome in some cases, where the COMMANDS have to be composed in various different ways. Such situations occur frequently, when using COMMANDS for composing components or configuring a system, because in these situations a large number of different COMMANDS need to be composed. Runtime composition of the COMMANDS is not possible if the composition solution is hard-coded in static, compiled languages such as C, C++, or Java.

Scenario

Consider the following scenario: a system is built from a number of C components. Different customers require different components running in different configurations. The system needs to be composed and configured flexibly according to these customers' requirements. Rapid configuration to the customers' demands is needed. This is hard to achieve, however, because changing the system means programming in C, recompilation, and restart of the system. COMMANDS can be used to avoid this problem up to a certain degree, because they provide a way to encapsulate the composition and configuration code and exchange it at runtime. However, the composition and configuration of components is still hard-coded in the system, and COMMAND-based composition is rather cumbersome.

Forces

Composition in static, compiled languages, such as C, C++, or Java, is inflexible for situations where runtime behavior modification is required, because programming, recompilation, and restart are required for each change.

It is hard to implement simple, domain-specific languages, for instance for rapid configuration of software systems written in those languages, because this requires writing a parser, preprocessor, and interpreter for that language. Even if an existing parser, as for instance, an XML parser can be reused, still the mapping of language elements to implementations needs to be programmed by hand. For simple configuration options this is easily done, but for more complex languages that can handle behavior definitions, this can quickly result in a complex task.

The code of multiple, consecutive COMMAND invocations might be hard to read, just consider the following simple example:

```
if (expressionCommand.execute().toBoolean()) {
    result = command1.execute();
    command2.value = result;
    command2.execute();
}
```

Solution

Express COMMAND composition in a COMMAND LANGUAGE, instead of calling the COMMANDS directly using an API. Each COMMAND is accessed with a unique command name. The host language, in which COMMANDS are implemented, embeds the COMMAND LANGUAGE. In the host language, THE COMMAND LANGUAGE'S INTERPRETER or compiler is invoked at runtime to evaluate the COMMANDS expressed in the COMMAND LANGUAGE. Thus COMMANDS can be composed freely using COMMAND LANGUAGE scripts, even at runtime.

Discussion

A part of the COMMAND LANGUAGE is either an INTERPRETER or an on-the-fly compiler. Note that often byte-code compilers are used for COMMAND LANGUAGES. A byte-code compiler can also be part of an INTERPRETER. From a language user's perspective all these alternatives look the same: the user operates in a language that is interpreted at runtime.

Almost all elements of a COMMAND LANGUAGE are COMMANDS, what means that a COMMAND LANGUAGE'S syntax and grammar rules are usually quite simple. There are some additional syntax elements, as for instance grouping of instructions in blocks, substitutions, and operators:

- Grouping of instructions in blocks is required to build larger instruction sets, such as method bodies, lists of arguments, and so forth. Most often brackets are used to delimit grouping in blocks.
- Substitutions are required in a COMMAND LANGUAGE to hand over the results of one COMMAND to another. Other substitutions than COMMAND substitution are also possible, e.g. variable substitution or character substitution.
- Assignment operators are an alternative to COMMAND substitutions to hand over the results of one COMMAND to another. Also other operators can be implemented.
- The end of instructions has to be found during parsing. Usually either line ends or special characters (like “;”) are used to mark an instruction end. A grouping can also mark an instruction end.

These COMMAND LANGUAGE elements are usually evaluated by the INTERPRETER before the COMMANDS are evaluated. In general, the INTERPRETER evaluates one instruction after another according to grouping and instruction end rules. Then it performs substitutions, if necessary. Finally it finds the COMMAND in the instruction and invokes it with the rest of the instruction as arguments.

Note that all further syntactic elements, as well as the semantics of the COMMAND LANGUAGE, must be well-defined, as in any other programming language. Finding the optimal syntax and semantics for a particular COMMAND LANGUAGE is beyond this pattern description because these issues depend on the concrete, domain-specific tasks the COMMAND LANGUAGE is used for. When the COMMAND LANGUAGE pattern is chosen as foundation of an integration architecture there is one important rule, however: before building a COMMAND LANGUAGE from scratch, one should always consider to reuse an existing language that supports COMMAND execution, for instance, an extensible scripting language or an existing domain-specific language. It is often much easier to adapt an existing language infrastructure to the concrete requirements of a software system than building a language implementation from scratch.

Consider again the above simple example. Using a COMMAND LANGUAGE we can provide the dynamic expression by variable substitution (with '\$') and pass the result of `command1` as an argument to `command2` (with '[...]'). These changes shorten the resulting code and make it much more readable:

```
if {$expression} {  
  command2 [command1]  
}
```

COMMAND LANGUAGE code is typically expressed as strings of the host language, in which the COMMAND LANGUAGE is implemented. From within the host language, code can be evaluated in the COMMAND LANGUAGE, and the results of these evaluations can be obtained. A central strength of the COMMAND LANGUAGE is that scripts can be (dynamically) assembled in strings and (also dynamically) evaluated later in time. In other words, we can write dynamically assembled and evaluated programs in host languages that do not support this feature natively. For instance, we can put the above script into a host language string, change this string at runtime, and evaluate it later in the program.

```
StringBuffer script =  
  new StringBuffer("if {$expression} {command2 [command1]}");  
...  
// later change the script (here: append logging code)  
script.append("\nputs \"evaluating expression: \"$expression\"");  
...  
// later evaluate the script  
interpreter.eval(script.toString());
```

COMMANDS of a COMMAND LANGUAGE are typically implemented in the host language and bound to command names. The command names and implementations are registered in the INTERPRETER of the COMMAND LANGUAGE. When the INTERPRETER evaluates a COMMAND, it looks up the implementation, registered for a command name, and invokes it.

Usually all language elements of the COMMAND LANGUAGE are implemented in the same way as user-defined COMMANDS. For instance, control statements, such as `if` or `while`, are implemented as pre-defined COMMANDS. An `if`-command can be realized by first evaluating the expression given as first argument. If this expression returns true, the “then” part of the `if`-statement is evaluated in the INTERPRETER. An example is shown in Figure 1.2.

Sometimes interaction between code in the COMMAND LANGUAGE and the host language is necessary. Defining COMMANDS and using `eval`, as in the examples above, is one simple way of language integration. “Deeper” language integration is supported by the pattern SPLIT OBJECT. For all kinds of interaction, type conversion between the COMMAND LANGUAGE and the host language is necessary. This can be done by letting the user deal with type casting manually, or it can be automated using the pattern AUTOMATIC TYPE CONVERTER.

The PIGGYBACK pattern [Spinellis2001] uses the capabilities of a host language as a hosting base for a new language, such as a domain-specific language (DSL). The COMMAND LANGUAGE pattern follows the PIGGYBACK structure and uses COMMANDS to realize the definition of the new language. The LANGUAGE EXTENSION pattern, also described in [Spinellis2001], is used to add new features to an existing language. Thus it is a general alternative to defining a new language following the PIGGYBACK pattern, as described by the COMMAND LANGUAGE pattern.

Scenario Resolved

Consider again the following scenario: when using a COMMAND LANGUAGE for integrating multiple C components, we can flexibly assemble these components and change the COMMAND LANGUAGE scripts even at runtime. Once a C component is integrated in the COMMAND LANGUAGE, it can be used without C programming or recompilation. For instance, the three COMMANDS in the example in Figure 1.2 are implemented in the host language and can be freely composed using scripts. These scripts can be provided as stand-alone scripts to the COMMAND LANGUAGE'S INTERPRETER, read from files for instance. Alternatively, they can be evaluated as strings of the host language using the host language API of the COMMAND LANGUAGE.

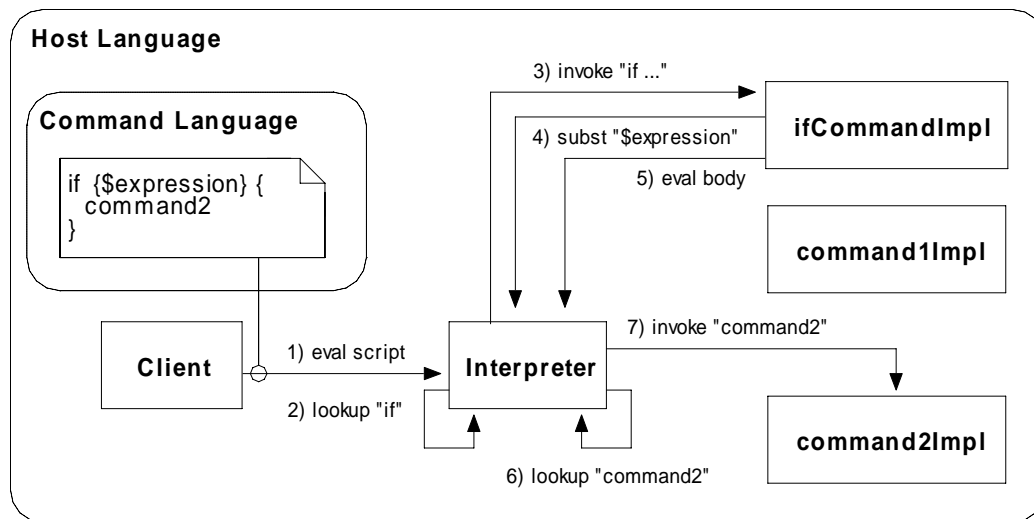


Figure 1.2

A script gets executed by the command language interpreter

Consequences

COMMAND LANGUAGE composition is very flexible because COMMAND LANGUAGE scripts can be evaluated and changed at runtime. COMMAND LANGUAGES are thus well suited for rapid prototyping and for implementing domain-specific languages. Any data can be treated as program code of a COMMAND LANGUAGE.

A COMMAND LANGUAGE is primarily designed for being easily extensible with COMMAND implementations. Thus it is simple to extend the COMMAND LANGUAGE with new components, typically written in another language. This is usually done by registering COMMANDS with command names in the COMMAND LANGUAGE'S INTERPRETER. In many COMMAND LANGUAGES solutions exist for automating the integration of components into the COMMAND LANGUAGES as new COMMANDS. How this is realized depends on the language features of the host language in which the COMMAND LANGUAGE is implemented. For instance, a wrapper generator can be used. Or, if supported, reflection can be used for performing a runtime lookup of the COMMAND implementation.

COMMAND composition in the host language is very efficient because COMMAND execution only requires one additional invocation in the host language. In contrast, a COMMAND LANGUAGE means a more severe overhead because it requires an additional dispatch of each invocation in the COMMAND LANGUAGE, as well as the interpretation of COMMAND LANGUAGE scripts.

Because the two languages have quite different tasks, each language can be designed for its particular task: the host language for writing efficient and reliable system components, and the COMMAND LANGUAGE for flexible component composition and other integration tasks.

When using a COMMAND LANGUAGE together with a host languages, developers have to learn two languages. Often, this is not a great problem: many COMMAND LANGUAGES are easy to learn because they have a simple syntax and only a limited number of language elements is required for COMMAND composition and configuration.

Implementing a COMMAND LANGUAGE from scratch might be a substantial effort, depending on the language features required. Thus reusing an existing language infrastructure (such as an existing scripting language or domain-specific language) should be considered before implementing the language from scratch.

Using a COMMAND LANGUAGE consumes resources. Invocations that are sent through the COMMAND LANGUAGE have a weaker performance than directly invoked COMMANDS. The language implementation requires additional memory, too.

Known Uses

Some known uses of the pattern are:

- Most scripting languages, such as Tcl, Python, and Perl are implemented as COMMAND LANGUAGES. That is, the language elements of the scripting language are implemented in another language; the languages named above are implemented in C. The INTERPRETER of the scripting language maps the command names in the scripts to the respective implementations in C. The scripting languages can thus be extended with new language elements implemented either in C or in the scripting language.
- A domain-specific language is a programming language tailored specifically for an application domain rather than providing a general purpose language. Many interpreted domain-specific languages are implemented as COMMAND LANGUAGES. This is because COMMAND LANGUAGES are easy to implement, easy to extend, or simply because another COMMAND LANGUAGE'S INTERPRETER is used for implementing the domain-specific language.
- The aspect-oriented framework of JAC [Pawlak+2001] implements a COMMAND LANGUAGE for aspect composition and configuration. This way, operations of an aspect component can be provided as command implementations and invoked from the configuration file. Each aspect can define its own "little" configuration language. This aspect can be configured with custom parameters and behavior without the need for programming and re-compilation.

COMPONENT WRAPPER

Context

A system is composed from a number of black-box components. Component configurations and other component variations are handled by parameterization using the components' interfaces.

Problem

You require a flexible composition of the system's components, but this is difficult, if component variations can only be introduced using parameterizations or exchange of the component. This is because parameters have to be foreseen by the component developers and thus are not well suited for dealing with unanticipated changes. Exchanging the component often means re-writing the component from scratch. Somehow we need to deal with unanticipated changes for a component without sacrificing the black-box property.

Scenario

Consider the following scenario: a company has two related products, one implemented in C and one in C++. Both are monoliths, each consisting of a number of hard-wired subsystems. Consider the systems should be evolved into a product line, consisting of a number of products. It should be possible to flexibly assemble products according to customer demands. That is, it should be possible to compose new products from the existing system components. Often composition requires little extensions and changes to the existing components. The product line architecture should provide some concept for integrating components into products and for extending individual components in the context of a product with new state or behavior.

Forces

If the component's source code is available, one way to cope with unanticipated changes is to change the component itself. This is problematic because the black-box property of the component should not be violated: black-box reuse eases understandability because the component can be reused without intimate knowledge of the component's internals. The component's internal implementation can even be replaced by another implementation, allowing us to get independent from implementation details. These are central reasons for using the component-based approach – and thus should not be sacrificed without need.

Components – especially third-party components – change over the time. Even though the black-box abstraction aims at stable interfaces, often the interfaces change. The client's developer should be able to cope with these changes without a necessity to seek for component accesses through the whole client's source code.

In many situations, one and the same client version has to deal with several component versions. Then adaptations to interfaces and to relevant internal implementation details have to take place within the same client implementation. With techniques such as preprocessor directives or if/switch control structures this may result in hard to read code.

Usually a client of a black-box component should be independent of the internal implementation of the component. But sometimes implementation details of the component implementation, such as access to shared resources, performance issues, and memory consumption, do matter and need to be accessed somehow.

Solution

Let the access point to a component, integrated into a system, be a first-class object of the programming language. Indirection by a COMPONENT WRAPPER object gives the application a central, white-box access point to the component. Here, the component access can be customized without interfering with the client or the component implementation. Because all components are integrated in the same way, a variation point for white-box extension by component's clients is provided for each black-box component in a system.

Discussion

The COMPONENT WRAPPER provides access to a component using the object abstraction of the component's client. A simple case is that both the component implementation and the component client use the same object abstraction; however, this is not always the case. In fact, a major goal of the COMPONENT WRAPPER pattern is to integrate components, written in different paradigms and languages, using the same integration style. To reach this goal, the COMPONENT WRAPPER pattern can use some other patterns within its internal implementation (there are many other integration styles though):

- In the simple case that another object-oriented component is integrated, the COMPONENT WRAPPER can use the PROXY pattern [Gamma+1995]. A PROXY is a placeholder object for another object offering the same interface. Note that a COMPONENT WRAPPER interface can also be different to the component's interface or only provide an excerpt of it, as for instance an "exported" interface. If two different object-oriented languages with slightly different object and type concepts are integrated, adaptations and type conversions can take place on the COMPONENT WRAPPER.
- When we want to integrate a procedural paradigm component into an object-oriented implementation, we apply the WRAPPER FACADE pattern [Schmidt+2000]. It uses one or more objects to represent a procedural component. The wrapper object forwards object-oriented messages to the procedures (or functions).
- If a COMMAND LANGUAGE is also used, the COMPONENT WRAPPERS can be exposed as COMMANDS of the COMMAND LANGUAGE. As a consequence, COMPONENT WRAPPERS can then be flexibly assembled using COMMAND LANGUAGE scripts. That is, two variability mechanisms are combined here: the COMPONENT WRAPPER allows for variation of the individual components and their integration in the system; the COMMAND LANGUAGE scripts allow for flexibly composing multiple COMPONENT WRAPPERS in the system. The COMPONENT WRAPPER is exposing the component's export interface as operations of the COMMAND, one for each exported functionality of the component. The arguments of the COMMAND in the COMMAND LANGUAGE script are used to dispatch the correct COMPONENT WRAPPER operation.

In each of the three variants, the COMPONENT WRAPPER can expose the same interface as the wrapped component (then it is a PROXY), or a different one. When changes to types of the interface are necessary during an invocation by a COMPONENT WRAPPER, the COMPONENT

WRAPPER can either contain the conversion code or use an AUTOMATIC TYPE CONVERTER for the conversion.

COMPONENT WRAPPERS are used to provide a uniform way of wrapping black-box components. The component client does not know about the style of wrapping or other implementation details. Moreover, the indirection to a COMPONENT WRAPPER object provides an accessible white-box on client-side. This placeholder provides a central access point to the component. DECORATORS [Gamma+1995], ADAPTERS [Gamma+1995], MESSAGE INTERCEPTORS [Zdun2003b], and other kinds of extensions, variations, or customizations optionally can be added to the COMPONENT WRAPPER. Only changes that necessarily have to interfere with the component's internals cannot be handled this way.

COMPONENT WRAPPERS can easily be generated, if the component's interface is electronically documented with an INTERFACE DESCRIPTION [Völter+2004]. The default implementation of such generated COMPONENT WRAPPERS just forwards the invocation to the component, but the client is able to extend the generated code. Using techniques, such subclassing, polymorphism, or MESSAGE INTERCEPTORS [Zdun2003b], we can avoid that re-generated code overwrites client extensions of the COMPONENT WRAPPER.

Using COMPONENT WRAPPERS we can modularize the aspect "integration of another component". Thus COMPONENT WRAPPERS are well suited to be applied together with an aspect-oriented approach [Kiczales +1997]. That is, an aspect composition framework can be used to weave the aspect represented by the COMPONENT WRAPPER into the system.

Scenario Resolved

Consider again the following scenario: two monolithic products should be evolved into a product line architecture. This can be done in an incremental fashion using COMPONENT WRAPPERS: they are used to decouple the components of the legacy systems step by step. Figure 1.3 shows that COMPONENT WRAPPERS are implemented as part of the product line, and instantiated by a new product. There is one COMPONENT WRAPPER for each of the export interfaces of the products. For the C product the pattern WRAPPER FACADE is used, for the C++ product we use a PROXY. In the example we can see that one COMPONENT WRAPPER performs an interface adaptation; the PROXY only forwards the invocation. On the COMPONENT WRAPPERS individual products, composed from the existing components, can add behavior or state to the components.

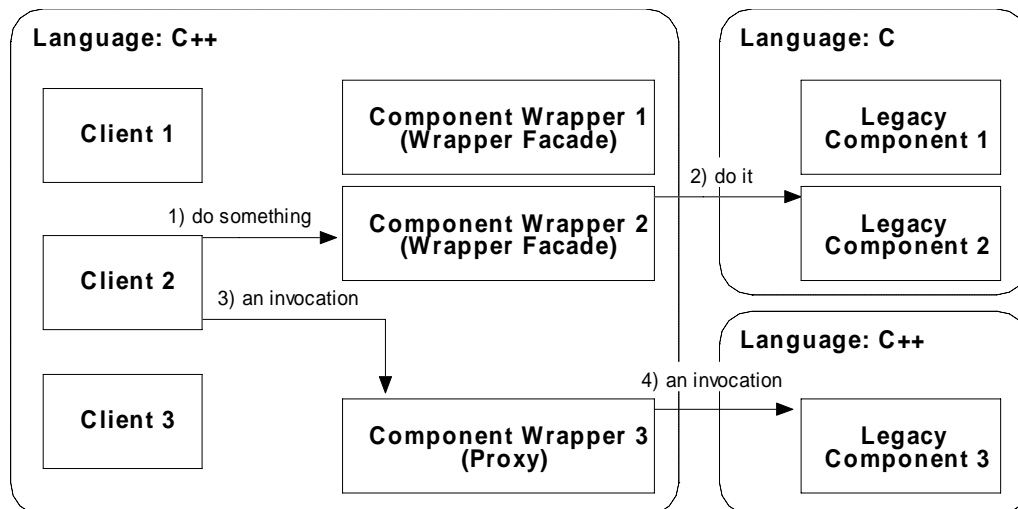


Figure 1.3

A number of component wrappers are integrated into a C++/C system: some are wrapper facades, some are proxies

Consequences

Using COMPONENT WRAPPERS customizability of black-box components is enhanced. Different styles of wrapping are conceptually integrated from a component client's perspective. Application developers can flexibly adapt and change component interfaces at a central access point. For instance, slight version changes in component version can be hidden by the COMPONENT WRAPPERS so that clients do not have to contain the adaptation code (such as #IFDEF directives). In other words, COMPONENT WRAPPERS reduce tangled code caused by component integration, composition, and configuration.

Using COMPONENT WRAPPERS an external component is used as an internal, first-class entity of the object system. That is, once integrated, all external and internal components look the same for component clients. The component use is decoupled from the internal realization. Therefore, the component itself is exchangeable with another implementation.

The indirections to the COMPONENT WRAPPER reduce performance slightly. More classes and flexibility hooks can result in a higher complexity of the application and thus reduce understandability. COMPONENT WRAPPERS introduce additional code. Thus there is more code to test and understand. Due to the indirection and the additional code, system's using COMPONENT WRAPPERS might be slightly harder to debug.

Known Uses

Some known uses of the pattern are:

- In [Goedicke+2002] we present a case study of a reengineering project for a document archive system implemented in C. We use COMPONENT WRAPPERS for migrating the existing C implementation into a component architecture in a stepwise manner.
- SWIG [Swig2003] is a wrapper generator that can generate COMPONENT

WRAPPERS for C and C++ components in various other languages. SWIG uses language bindings and header files as INTERFACE DESCRIPTIONS.

- VCF [OGJ03] integrates components from various Java component models, namely COM+, CORBA, EJBs, and Java Beans. A FACADE provides a common denominator interface of the technologies to clients. For integration of the component models, the approach uses a plug-in architecture that is based on a number of COMPONENT WRAPPERS.

Object System Layer

Context

You want to design with object-oriented techniques and use the benefits of advanced object-oriented concepts. Parts of a system implementation are realized in a language that does not support object-orientation, such as C, Tcl, or Cobol. Or the implementation is realized in an object system that does not support desired object-oriented techniques. For instance, C++ or Object Cobol do not implement reflection or interception techniques. Or, foreign object concepts are used, such as those of middleware system like CORBA.

Problem

Suppose you are faced with target programming languages that are non-object-oriented, or with legacy systems that cannot quickly be rewritten, or with target object systems that are not powerful enough or not properly integrated with other used object systems. But the target language is chosen for important technical or social reasons, such as integrating with legacy software, reusing knowledge of existing developers, and customer demands, so it cannot be changed. Somehow the object-oriented and non-object-oriented concepts need to be integrated in one design and implementation model.

Scenario

Consider the following scenario: legacy applications are often written in procedural languages, such as C or Cobol. A complete migration of a system to a new object-oriented language often makes the integration with the existing legacy components difficult and forces a reimplementing of several well functioning parts. The costs of such an evolution are often too high to consider the complete migration to object technology at all. Most often the costs are even very hard to estimate in advance.

Forces

There are a number of situations in which we require some integration with an object system, as for instance:

- Many modern systems are based on object-oriented concepts, but still legacy systems offering no support for object-orientation need to be supported.
- When two or more different object-oriented languages need to be integrated, we can use COMPONENT WRAPPERS for integration. However, we should avoid implementing this integration for each individual COMPONENT WRAPPER on its own.
- The need for integration of object concepts might also occur in only one object-oriented language when foreign object concepts need to be integrated. For instance the CORBA or COM object systems are not exactly the same as the object systems of object-oriented languages, such as C++ or Java, in which they are used. This problem arises for all kinds of technologies that introduce object concepts, as for instance distributed object systems, application servers, transaction monitors, and relational and object-oriented databases.
- Many applications use an object system of a mainstream programming language,

such as C++, Java, or Object Cobol. Those object systems only implement standard object-oriented techniques. Thus extensions of object concepts, such as interception techniques, reflection and introspection, role concepts, or aspect-orientation are not provided natively.

Manual integration of object systems often means that developers are performing the integration over and over again, which is cumbersome and error-prone. Instead the integration of object systems should be reusable.

Solution

Build or use an object system as a language extension in the host language and then implement the design on top of this OBJECT SYSTEM LAYER. Provide a well-defined interface to components that are non-object-oriented or implemented in other object systems. Make these components accessible through the OBJECT SYSTEM LAYER, and then the components can be treated as black-boxes. The OBJECT SYSTEM LAYER acts as a layer of indirection for applying changes centrally.

Discussion

An OBJECT SYSTEM LAYER requires some dispatch mechanism that translates invocations in the host language into invocations to the objects of the OBJECT SYSTEM LAYER. This can be done using a MESSAGE REDIRECTOR [Goedicke+2001]. It acts as a FACADE [Gamma+1995] to the whole OBJECT SYSTEM LAYER and is not bypassed. For instance, the MESSAGE REDIRECTOR invocations might be string-based messages, which are translated into the respective invocations.

The OBJECT SYSTEM LAYER can be part of a COMMAND LANGUAGE, if the COMMAND LANGUAGE implements an object system. Then the OBJECT SYSTEM LAYER'S MESSAGE REDIRECTOR is part of the COMMAND LANGUAGE'S INTERPRETER. Once a script is parsed, the INTERPRETER uses a MESSAGE REDIRECTOR to indirect the invocation to the implementation of the COMMAND LANGUAGE element. A typical integration of the patterns OBJECT SYSTEM LAYER and COMMAND LANGUAGE uses one object per COMMAND, and the command name is used as an object ID for the OBJECT SYSTEM LAYER object.

The SPLIT OBJECT pattern is used, if the OBJECT SYSTEM LAYER needs to be integrated with a host language object system. Using COMPONENT WRAPPERS foreign paradigm elements, such as procedural APIs, can be integrated into the OBJECT SYSTEM LAYER'S object system.

Besides the conceptual integration with host language concepts, an OBJECT SYSTEM LAYER also must deal with other potential mismatches of the host language and the OBJECT SYSTEM LAYER'S model. For instance, mismatches in the threading models or mismatches in transaction contexts must be resolved.

The OBJECT SYSTEM LAYER can be used to introduce extensions of the object concept of an object-oriented language. Typical examples are the introduction of runtime dynamic object and class concepts such as TYPE OBJECTS [Johnson+1998], ADAPTIVE OBJECT MODEL [Yoder+2002], role concepts, or MESSAGE INTERCEPTORS [Zdun2003b].

Many technologies integrated into a system, implement an OBJECT SYSTEM LAYER, as for instance distributed object systems, application servers, transaction monitors, and databases. If needed, an OBJECT SYSTEM LAYER can also be introduced as a superset of (a part of) two or more object concepts and used for integrating them. For instance, the OBJECT SYSTEM LAYERS of technologies

that are implemented in multiple languages, such as distributed object frameworks, are primarily used for integrating the object concepts of these languages.

Scenario Resolved

Consider again the following scenario: procedural components should be composed using object-oriented concepts. An OBJECT SYSTEM LAYER can be used for this purpose, as depicted in Figure 1.4. This has the benefit that we do not have to translate the system into an object-oriented language completely. Instead existing components can be incrementally evolved; others can stay in the procedural language and be integrated into the OBJECT SYSTEM LAYER using COMPONENT WRAPPERS.

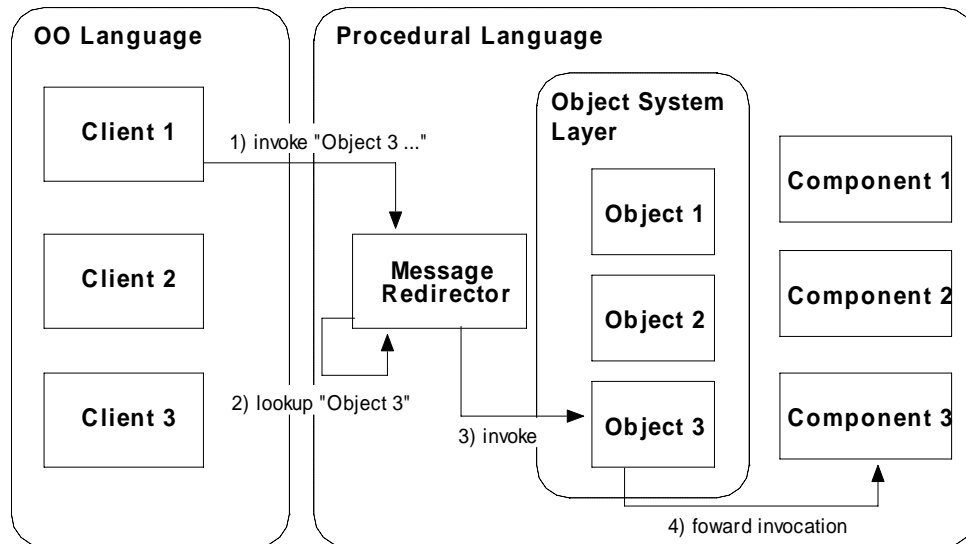


Figure 1.4

An object system layer is embedded in a procedural language

Consequences

When using an OBJECT SYSTEM LAYER, the flexibility of the application can be improved, since the OBJECT SYSTEM LAYER allows us to introduce variation points and to implement high-level language constructs, such as interception, reflection, and adaptation techniques.

The complexity of the application can be reduced, because – in scenarios where they are really needed – more high-level concepts, such as roles, interceptors, or aspects can simplify the program code and avoid complex workarounds.

Complexity of the application can also rise, if the client has to maintain the OBJECT SYSTEM LAYER. Then issues such as garbage collection, object destruction, relationship dynamics, reflection, etc. have to be programmed by hand. But this problem can be avoided by using an existing OBJECT SYSTEM LAYER as a library.

Performance can be decreased due to additional indirections.

The OBJECT SYSTEM LAYER'S conventions and interfaces have to be learned by the developers in addition to the system's APIs.

Known Uses

Some known uses of the pattern are:

- Object-oriented scripting languages implemented in C, such as Python or XOTcl [Neumann+2000], provide an OBJECT SYSTEM LAYER for C. These languages provide a C API – that means the object models can be accessed and used from C for using object-oriented abstractions in C programs.
- The Redland RDF library [Beckett2004] implements a simple OBJECT SYSTEM LAYER for C to use object-oriented abstractions for RDF nodes, models, streams, and other elements of the library. The library implements classes, object IDs, constructors, and destructors.
- Procedural implementations of object-oriented middleware, such as C or Tcl implementations of CORBA or Web services, provide the object abstractions of the middleware in the procedural language. Thus in the distributed context these implementations provide an OBJECT SYSTEM LAYER.

AUTOMATIC TYPE CONVERTER

Context

Languages or systems supporting different types are used in one system.

Problem

When two different type systems need to be integrated, it is necessary to convert corresponding types. An INTERFACE DESCRIPTION can be used to describe the type differences, and the corresponding type conversion code can be generated. Conversion code for each particular invocation can only be generated, however, if the signatures of all operations are known during generation. Sometimes operation signatures are not known before runtime. In many integration situations it is also necessary that it is easy to add new type bindings rapidly.

Scenario

Consider the following scenarios:

- In Web services frameworks remote invocations and other messages are transferred using SOAP [Box+2000] as a XML-based payload format. In SOAP, invocations containing parameter and return types are encoded in string format. To be interpreted by a programming language the SOAP messages must be converted to programming language native formats. From AN INTERFACE DESCRIPTION we can generate client proxy and server stubs that perform the type conversion. This, however, does only work, if the types are known before runtime; for dynamic deployment of services we basically have two choices: we can generate a server stub while the system runs or and perform type conversions at runtime.
- A COMMAND LANGUAGE is embedded in a statically typed language. The COMMAND LANGUAGE is given in form of scripts that are interpreted at runtime. That is, all types to be accessed by the COMMAND LANGUAGE need to be dynamically convertible to and from strings. To ease type integration, it should be simple to integrate new types.

Forces

Different languages support different types that cannot be exchanged directly. This problem also arises for other systems supporting types – not only for programming languages. For instance, most distributed object frameworks support their own type system to integrate across languages or platforms. Systems supporting dynamic types – for instance using the patterns TYPE OBJECT [Johnson+1998] or ADAPTIVE OBJECT MODEL [Yoder+2002] – require some way to cast the dynamic types into other dynamic types or static types of the host language. Some languages and typed systems only support canonical types, such as strings. To integrate them with a strongly typed system, we need to convert the types to the canonical format and vice versa.

When an INTERFACE DESCRIPTION of all operation signatures exists, an option is to generate the type conversion code. This option is for instance often chosen in the proxies and stubs generated by many distributed object frameworks. However, this does not work for more dynamic or complex type conversion situations. For instance, generation is not as well suited for dynamic

typing or runtime deployment because here the operation signatures are not known before runtime of the system. Note that runtime generation of conversion code is possible – it is just more complex than writing simple wrappers.

Also, in many cases, it is necessary to provide a very simple kind of type conversion because it should be easily extensible: only the type binding and the conversion code should be needed to be specified.

Solution

Provide an AUTOMATIC TYPE CONVERTER as a means to convert an instance of a particular type supported by the AUTOMATIC TYPE CONVERTER to a particular target type, and vice versa. The AUTOMATIC TYPE CONVERTER is either extensible with new type bindings for one-to-one transformations, or converts every type to a canonical format, such as strings, and can convert the canonical format to any supported type. If no user-defined type binding is found, the AUTOMATIC TYPE CONVERTER should execute a default conversion routine.

Discussion

In general, an AUTOMATIC TYPE CONVERTER should be able to deal with all kinds of conversions that occur at runtime because manual intervention is not possible at runtime and a conversion that fails leads to a runtime error. That is, all kinds of AUTOMATIC TYPE CONVERTERS should always provide a default conversion routine that takes place when no custom type binding is defined. Hooks for type conversion are provided so that users can extend the type converter with additional conversions.

There are two main variants of AUTOMATIC TYPE CONVERTERS:

- There might be a canonical format, such as strings, `void*` in C/C++, `Object` in Java, supported by the AUTOMATIC TYPE CONVERTER and a number of other types. The AUTOMATIC TYPE CONVERTER can convert any type to the canonical format and the canonical format to any type. This way, any type can be converted into any other type by a two-step conversion. Note that it is the developer's responsibility to provide sensible canonical representations of the various types which allow for conversion without information loss and which are efficient. Also, it is the developers burden to request only sensible conversions. For instance, converting a whole database into a in-memory string representation is problematic; it might be better to convert a database handle only. The system must be designed in a way so that continuous back-and-forth conversions are avoided, where possible.
- An alternative to using a canonical format is to use direct, one-to-one conversion. This is usually faster because only one conversion is required. The drawback of this variant is that we might have to write more converters: if each type can be mapped to each other type, we have to write $N * (N - 1)$ converters, instead of $2 * N$ converters needed to support a canonical format. Therefore, one-to-one conversion is often chosen if there is only one target type for each type. A second reason for using one-to-one conversion is that specialties of the individual types must be considered during conversion, for instance, because information would get lost during conversion to a canonical format. Finally, a third reason for choosing one-to-one conversion might be the better performance because conversion to/from the intermediate canonical representation can be avoided.

An AUTOMATIC TYPE CONVERTER can also be used for more than two types. Thus, for instance, it can also be used to implement polymorphism for TYPE OBJECTS [Johnson+1998] or other dynamic typing solutions. The AUTOMATIC TYPE CONVERTER can convert an object to all types, supported by its classes, and delivers the casted type, as requested by the client. This is for instance useful if an AUTOMATIC TYPE CONVERTER is combined with an OBJECT SYSTEM LAYER or SPLIT OBJECTS.

A COMMAND LANGUAGE implemented in a host language usually requires an AUTOMATIC TYPE CONVERTER. In this scenario, the canonical format variant is often supported because the COMMAND LANGUAGE scripts are represented as strings, and these strings can be used as a canonical format. Thus the AUTOMATIC TYPE CONVERTER is used as a mechanism to exchange data, specify operation parameters, and receive operation results between a COMMAND LANGUAGE and its host language.

An AUTOMATIC TYPE CONVERTER is a flexible alternative to type conversion code generated from an INTERFACE DESCRIPTION. In many cases, generated type conversion code uses the AUTOMATIC TYPE CONVERTER internally. An AUTOMATIC TYPE CONVERTER can also use INTERFACE DESCRIPTIONS as an aid in the type conversion decision.

The pattern CONTENT CONVERTER [Vogel+2002] is a similar pattern describing content integration – it also can be used with a canonical format or using one-to-one integration. If the content of a CONTENT CONVERTER is typed, an AUTOMATIC TYPE CONVERTER can be used as part of a CONTENT CONVERTER used for the type conversion task.

Scenario Resolved

Consider again the following scenarios:

- In a Web services framework we usually can map every programming language type to one type in the INTERFACE DESCRIPTION. Thus here we can use a one-to-one converter between INTERFACE DESCRIPTION type, transmitted over the network, and the type of the programming language.
- A COMMAND LANGUAGE needs to be integrated with a statically typed language. The COMMAND LANGUAGE uses scripts consisting of strings. Thus, for each host language type that we want to expose, we need a converter that converts the host language type to strings and from strings. In Figure 1.5 a solution using a string representation as a canonical format is sketched. Each AUTOMATIC TYPE CONVERTER object internally has two representations: a string representation and an internal representation containing the target type of the statically typed host language. The AUTOMATIC TYPE CONVERTER class provides an operation for converting the internal representation to strings, an operation for converting strings to the internal representation, getter/setter operations for both representations, and operations for discarding one of the two representations. At each point in time, at least one of the two representations must be valid.

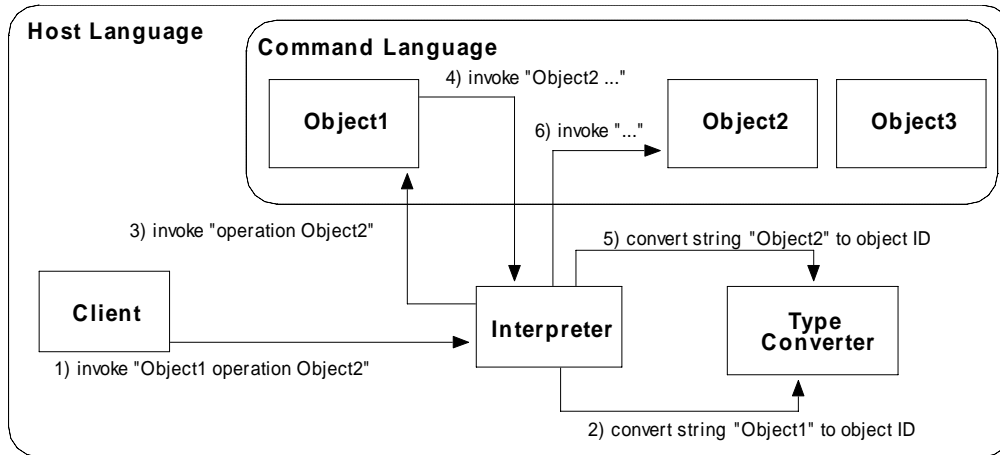


Figure 1.5

A host language client invokes an command language object and the interpreter performs automatic type conversions

Consequences

An AUTOMATIC TYPE CONVERTER is in many scenarios slower than a solution generated from an INTERFACE DESCRIPTION because it performs the conversion at runtime, whereas the generated solution only performs a single cast. The benefit of an AUTOMATIC TYPE CONVERTER is that it is more flexible than generated solutions: this way dynamic typing can be supported. Type mappings can be changed at runtime. Also, the type mappings in an AUTOMATIC TYPE CONVERTER are easy to extend.

Known Uses

Some known uses of the pattern are:

- The scripting language Tcl uses an AUTOMATIC TYPE CONVERTER for converting the string-based scripts to native types in C and vice versa. Strings are used as a canonical format.
- IONA Artix is a Web services framework that supports conversion between different protocols (HTTP, IIOP, Websphere MQ, Tuxedo, etc.) and payload formats (SOAP, Tuxedo's FML, GIOP, etc.). To increase conversion performance, Artix uses one-to-one converters. An AUTOMATIC TYPE CONVERTER for each converter is needed to recognize the types in one payload format and translate them into the target format.
- Apache Axis is a Web services framework that uses AUTOMATIC TYPE CONVERTERS for one-to-one conversion between programming language types in Java and the payload format SOAP.

SPLIT OBJECT

Context

Two languages need to be used together in one system.

Problem

To integrate two languages one can write simple wrappers that allow for invocations in the other language. However, pure wrapping poses some problems in more complex language integration situations. A wrapper provides only a “shallow” interface into a system, and it does not reflect further semantics of the two languages. Examples of such semantics are class hierarchies or delegation relationships. Further, a wrapper does not allow one to introspect the system’s structure. The logical object identity between wrapper and its wrappees is not explicit. Complex wrappers that are implemented manually are hard to maintain.

Scenario

Consider the following scenario: one language embeds another language, such as a COMMAND LANGUAGE that is embedded in a host language for configuring host language objects. That is, the COMMAND LANGUAGE requires access to its host language, and vice versa. Here “access” means, for instance, performing lookups, invocations, creations, and destructions of objects and methods.

Forces

A main driving force for considering deeper language integration than COMPONENT WRAPPERS is that it is cumbersome to write wrappers for each and every element that requires language integration. To a certain degree, wrappers can be generated, but still we would require custom code for each language integration situation. Instead it should be possible to reuse language integration code. Automation and reuse would also foster maintainability: in general, complex wrapping structures can quickly get quite complex and thus hard to maintain. Especially scattering wrapping code across the functional code should be avoided.

Simple wrappers are usually realized by a single indirection, and thus they are relatively fast and consume only little extra memory. Deeper language integration needs to be compared to simpler wrapper solutions with regard to resource consumption.

Solution

A SPLIT OBJECT is an object that physically exists as an instance in two languages, but logically it is treated like one, single instance. Both SPLIT OBJECT halves can delegate invocations to the other half. The SPLIT OBJECT halves mimic the user-defined class hierarchy of the counterpart, variables are automatically traced or shared, and methods can be wrapped. Depending on the language features of the two languages, these functionalities can either be implemented by extending the language’s dispatch process, using reflection, or using program generation.

Discussion

There are different ways how a SPLIT OBJECT half can communicate with its counterpart in the other language:

- A method can be provided, for instance by a superclass of all SPLIT OBJECTS, that allows for sending an invocation to the other half. The invocation parameters are usually provided as arguments using a generic type (such as strings or `Object` in Java) and converted using an AUTOMATIC TYPE CONVERTER. Depending on the language features, the method implementation on the other split object half can be looked up using reflection, using a dynamic dispatch mechanism, using a registration in a method table (provided by the developer), or using a generated method table.
- A wrapper generator can generate a method on a SPLIT OBJECT half for each exported method of the other SPLIT OBJECT half. This generated method is a simple forwarder to the other split object half. The invocation parameters can be provided using the correct types, required by the counterpart, if possible, or an AUTOMATIC TYPE CONVERTER is needed for conversion. This variant is very efficient because it requires no dynamic lookup.
- If the language’s dispatch mechanism can be extended and method dispatch is performed at runtime, we can implement an automatic forwarding mechanism. Whenever a method cannot be found for the SPLIT OBJECT half, the “method not found” error is not raised directly, but it is first tried to find the method on the other SPLIT OBJECT half. If this succeeds, the method is invoked there, otherwise the error is raised. In this variant, the same means for method lookup as in the first alternative can be used.
- In dynamic languages (e.g. a COMMAND LANGUAGE), wrapper methods can also be generated at runtime. Again, these methods are simple forwarders, as in the second alternative, but the means for method lookup, described in the first alternative, are used.

Note that only in the first variant clients recognize that the object does not itself implement the method; in the following three variants the client can use the same invocation style as used for methods implemented in the same language.

SPLIT OBJECTS assume a language integration structure following the PIGGYBACK pattern [Spinellis2001]. This pattern generally describes how to use the capabilities of a host language as a hosting base for a new language, such as a domain-specific language (DSL) or a COMMAND LANGUAGE.

In many cases, split object halves will mimic the user-defined class hierarchy of the other half. For instance, when you generate forwarder methods to the other split object half, you face the problem of name clashes: methods on different classes of one object might have the same name. This problem does not arise, if every class for which a forwarder method is generated, exists within both languages. For this to work, to a certain extent, we need to integrate the class concepts of the two languages. For instance, if one language supports multiple inheritance, and the other not, we need to simulate multiple inheritance in the other language. To mimic a class hierarchy in a procedural language or integrate different class concepts in two object-oriented languages, you might consider implementing a simple OBJECT SYSTEM LAYER.

From the COMMAND LANGUAGE we can automatically forward invocations into the host language. In turn, ordinary host language invocations bypass the SPLIT OBJECT in the COMMAND LANGUAGE. Thus, from within the host language, we have to use the INTERPRETER’S `eval` method to access a SPLIT OBJECT. To avoid this additional invocation style, the pattern HOOK

INJECTOR [Zdun2003b] can be used to replace host language invocations with indirections to SPLIT OBJECTS.

Scenario Resolved

Consider again the following scenario: when integrating a COMMAND LANGUAGE with a host language, we can provide a method for invoking the COMMAND LANGUAGE from the host language. The COMMAND LANGUAGE object, in turn, provides an automatic forwarding mechanism to invoke the host language half of a SPLIT OBJECT. As an example, Figure 1.6 shows a host language client that needs to invoke a host language object `Object1`. This object is a SPLIT OBJECT: instead of invoking this object directly, the counterpart in the COMMAND LANGUAGE is invoked first, which forwards the invocation back into the host language. This way the COMMAND LANGUAGE can intercept the invocation.

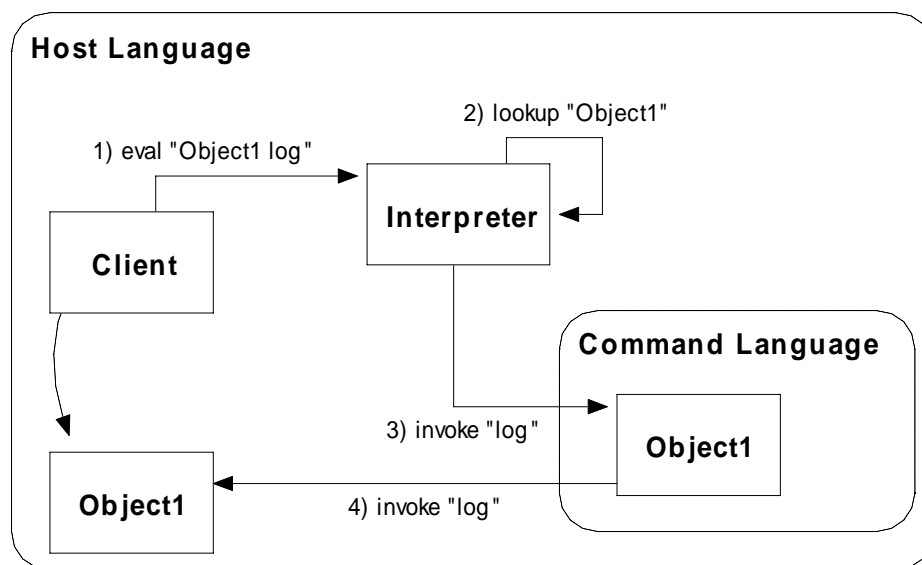


Figure 1.6

Invoking a split object through the command language

Consequences

SPLIT OBJECTS can be used to deeply integrate two object systems. Concepts realized in one object system can be used from within the other object system. The language integration code in the SPLIT OBJECT implementation can be reused. Thus language integration is easier to apply and the code is easier to maintain than integration code based on COMPONENT WRAPPERS only. SPLIT OBJECTS can be used as a model to automate wrapper generation for two languages.

Working with two languages might make debugging more difficult because you have align the debugging information in the two languages and maybe even integrate debugging tools. On the other hand, SPLIT OBJECTS can even be used to provide debugging information for other languages (see [Zdun2004a] for examples).

SPLIT OBJECTS pose a memory and performance overhead. This overhead can be minimized by not using dynamic lookup and dispatch mechanisms, but generative techniques instead.

Generated code is hard to change at runtime (if this is possible at all); that is, using generative techniques means less flexibility.

Known Uses

Some known uses of the pattern are:

- The network simulator [UCB2000a] and Open Mash [OpenMash2000] are two projects that are implemented in C++ and use OTcl [Wetherall+1995] for configuring and customizing the C++ applications. The integration of the two languages is performed by a SPLIT OBJECT solution called Tclcl [UCB2000b].
- In [Zdun2004a] it is demonstrated how to use split objects for dynamically performing maintenance tasks, such as component testing, feature tracing, and variation and configuration management. For instance, the language Frag [Zdu2003a] is used to configure and compose AspectJ [Kiczales+2001] aspects dynamically to perform these tasks in Java.
- The Redland RDF library [Beckett2004] contains a C-based OBJECT SYSTEM LAYER. The Perl interface to that system uses a hand-build SPLIT OBJECT solution: when a Perl wrapper is created, the respective C-based object is created as well – thus object IDs are mapped one to one. Similarly, the C destructors are automatically invoked in the correct order upon destruction of the Perl wrapper object.

Known Uses and Technology Projections

In this section we discuss some known uses and examples for the patterns in an integrated manner. We concentrate on a few interesting technology projections from the areas of distributed object frameworks, reengineering, scripting languages, configuration and composition of components, and dynamic aspect configurations. There are many other systems in each of these areas using some of the patterns, described in this chapter.

Reengineering a Document Archive System

In [Goedicke+2002] we present a case study of reengineering a large-scale documented archive system, implemented in C, to a component-oriented architecture. An important concern was incremental evolution of the system. Here, the pattern COMPONENT WRAPPER played a central role: the idea was to start off with simplistic wrappers, in particular WRAPPER FACADES for the original procedural components – these were more or less similar to the original subsystems of the system. The COMPONENT WRAPPERS provided a means to unify wrapping styles and handle extensions or adaptations. For instance, they allowed us to deal with situations in which different component versions had to be supported, interface changes had to be handled, and components with other wrapper styles had to be integrated.

To objectify the system, we applied the pattern OBJECT SYSTEM LAYER. A simple object system was provided with which the components can be treated in an object-oriented manner and easily integrated with the object systems of distributed object systems (in [Goedicke+2002]: CORBA).

New requirements impose changes on the wrappers constantly. Such extensions or adaptations can be transparently handled by MESSAGE INTERCEPTORS [Zdun2003b]. Here, the pattern OBJECT SYSTEM LAYER is beneficial again, because the MESSAGE REDIRECTOR of the OBJECT SYSTEM LAYER can be used to apply the MESSAGE INTERCEPTORS. It can also trigger an AUTOMATIC TYPE CONVERTER, where necessary.

Hardware selection for an MHP product line

The TPMHP project [Goedicke+2004] is concerned with developing a product line of software components for the Multimedia Home Platform (MHP) [ETSI2001]. The MHP specification is a generic set of APIs for digital content broadcast applications, interaction via a return channel, and Internet access on an MHP terminal. Typical MHP terminals are for instance digital set-top boxes, integrated digital TV sets, and multimedia PCs. The MHP standard defines a client-side software layer for MHP terminal implementations, including the platform architecture, an embedded Java virtual machine implementation running DVB-J applications, broadcast channel protocols, interaction channel protocols, content formats, application management, security aspects, a graphics model, and a GUI framework.

Even though there are just a few MHP set-top boxes available right now, a great diversity of products can be foreseen in this field. In addition, there will be a great number of supporting systems, such as input devices. The whole variety of different TV systems and PCs also might be used as an MHP client platform.

Software products that should run on a MHP platform need to consider multiple kinds of variations regarding the hardware environment. As MHP applications are broadcasted, the earliest possible binding time for the variation points is during startup of the application on the client platform.

Let us consider a few typical hardware selection problems. Different input devices offer various input events, and the MHP application needs to dynamically configure itself to these hardware elements. Graphical elements in the MHP API include a few graphic functions of Java, as well as the more advanced GUI elements of the HAVI standard [HAVI2001]. It depends on the type of set-top box if HAVI is supported or GUI elements have to be painted using the graphic functions of Java. Different back-channels can be identified at runtime and may in some situations be handled differently.

In pure Java these issues are typically handled by loose associations with objects. This, however, does not provide a central variation point management architecture that, for instance, supports querying various hardware options or reconfiguring them at runtime with a common model.

The solution, proposed in [Goedicke+2004], is to use a COMMAND LANGUAGE to handle the composition and configuration tasks for hardware-related software components. We use scripts like the following example for behavioral compositions and configurations:

```
...
MHPButton create OK -set parent OKBOX
MHPButton create CANCEL -set parent OKBOX
if {$earlyVisible == 1} {
    CANCEL setVisible
    OK setVisible
}
...
```

In the script two button components are created and composed with another component OKBOX. Then the components are configured conditionally: if the variable `earlyVisible` is true, the configuration option `visible` is set on both components.

Scripts, like this example, can be rapidly composed and changed, and they can be stored in a central repository. COMPONENT WRAPPERS for the respective GUI components (such as the MHPButton in the example above) are implemented in Java. Both types, HAVI and Graphics, support a common interface. In the COMPONENT WRAPPERS, the common interface is mapped to the HAVI API or the Java-based Graphics API respectively:

```
class HAVI_Button implements MHP_Button {
    HtextButton button;
    // use HAVI HtextButton to draw a button
    ...
}
class Graphics_Button implements MHP_Button {
    // use Graphics API to draw a button
    ...
}
```

The COMPONENT WRAPPERS are integrated into the COMMAND LANGUAGE as SPLIT OBJECTS. Each SPLIT OBJECTS realizes one COMMAND in the COMMAND LANGUAGE. At application startup

we can select which of the two COMPONENT WRAPPER sets is bound to the COMMANDS. That is, we can use the COMPONENT WRAPPER indirection to select between the HAVI implementation and the Graphics implementation at runtime. The same COMMAND LANGUAGE scripts can be used for both implementations.

The resulting object structure of the example, if HAVI is selected, is shown in Figure 1.7.

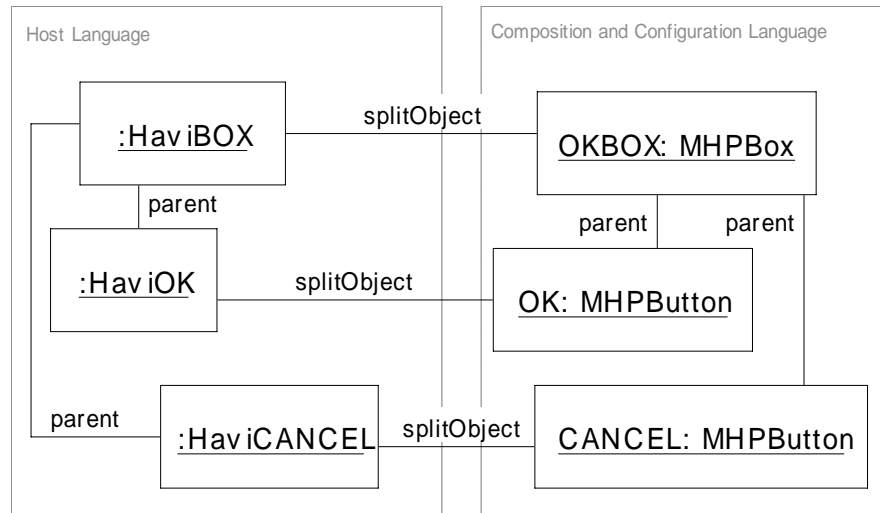


Figure 1.7

Object structure of the example

Let us compare the solution briefly to composition and configuration using hard-coded invocations in Java. The solution has many benefits: behavioral composition and configuration is not tangled in the business logic code. Thus the solution fosters understandability and maintainability. The developer does not have to care at all for language integration: this is done automatically by using the SPLIT OBJECT solution. Thus the solution is easy to use. The composition and configuration behavior can be changed at runtime. Thus the solution supports unanticipated software evolution. There are a few disadvantages as well: each composition and configuration object consumes memory, and invocations routed through two languages are slower than invocation performed in only one language. This is not a big drawback, as it only applies for composition and configuration code, which is not a large part of the system and does not affect performance-critical parts of the system. A second drawback is that more learning effort is required: developers must learn the Java MHP system plus the configuration options and the COMMAND LANGUAGE syntax. Overall, the benefits outweigh the drawbacks in this example.

TclCL and XOTcl/SWIG

TclCL [UCB2000b] is an integration of OTcl [Wetherall+1995] and C++. It utilizes SPLIT OBJECTS for language integration. TclCL is used in:

- Mash [OpenMash2000] is a streaming media toolkit for distributed collaboration applications based on the Internet Mbone tools and protocols.

- The Network Simulator (NS) [UCB2000a] supports network simulation including TCP, routing, multicast, network emulation, and animation.

In both projects, TclCL is applied to utilize the two language concept: the COMMAND LANGUAGE OTcl is used for high-level tasks, such as configuration, testing, and others, and C++ is used for the efficient implementation of the core tasks as C++ COMMANDS. The user-defined OTcl classes mimic the C++ class hierarchy. This is mainly done because the C++ variables /methods can be dynamically registered for OTcl in the C++ constructors of the respective classes. The OTcl counterpart remembers the capabilities of the C++ half, and thus can forward invocations that are implemented on the C++ half into the other language. OTcl is an OBJECT SYSTEM LAYER that adds dynamic object-oriented language concepts to the static languages C and C++.

XOTcl [Neumann+2000] is the successor of the language OTcl. We have implemented a C++/XOTcl binding using the wrapper generator SWIG [Swig2003]. This project has similar goals and a similar architecture as the TclCL solution explained above. The main difference is that in this solution the variable/method bindings are not registered in the C++ constructors, but are generated from an INTERFACE DESCRIPTION (the SWIG header file – which is similar to a C++ header file).

As Tcl extensions, both OTcl and XOTcl use strings in scripts to represent primitive data types. Internally, all primitive data types of XOTcl, such as ints, floats, doubles, booleans, strings, lists, and XOTcl objects are presented by so-called `Tcl_Objects`. These are AUTOMATIC TYPE CONVERTERS which contain data structures for a string representation and an internal representation for the respective type implemented in C. Also they contain conversion routines for each type to strings and for strings to each type – also implemented in C. `Tcl_Objects` use an abstract interface – using C function pointers – so that the Tcl INTERPRETER can automatically convert each type without knowing any more details of the types.

Configuring Aspects using Split Objects

A more sophisticated SPLIT OBJECT solution [Zdun2004a] is implemented in Frag [Zdun2003a]. Frag is a Tcl variant that can be completely embedded in Java. It is implemented on top of Jacl [DeJong +2003], a Tcl INTERPRETER implemented in Java. Frag also exploits SPLIT OBJECTS for language integration, but in contrast to the OTcl and XOTcl solutions explained above, it does not use registration or generative programming for access of the other SPLIT OBJECT half, but Java Reflection. Frag thus can be embedded in Java as a COMMAND LANGUAGE and be dynamically connected to Java classes.

Within Frag, there is a Frag class `JavaClass` provided. This class is used for wrapping a Java class with a Frag object. When an object is derived from this class or any of its subclasses, a SPLIT OBJECT consisting of a Java half and a Frag half is created.

All SPLIT OBJECTS in Frag inherit from a class `Java`. The method `dispatcher` of the class `Java` is responsible for forwarding all invocations that cannot otherwise be dispatched to the Java half. `dispatcher` is automatically invoked as a default behavior, when a method cannot be dispatched in Frag. Both, wrapping a Java class and forwarding invocations, is implemented using Java reflection. Internally, primitive Java types are automatically converted to and from strings. Non-primitive Java types have to be used as SPLIT OBJECTS in order to be accessed from Frag.

Figure 1.8 shows the structure for an example Java class `AJavaClass` that has a SPLIT OBJECT in `Frag`.

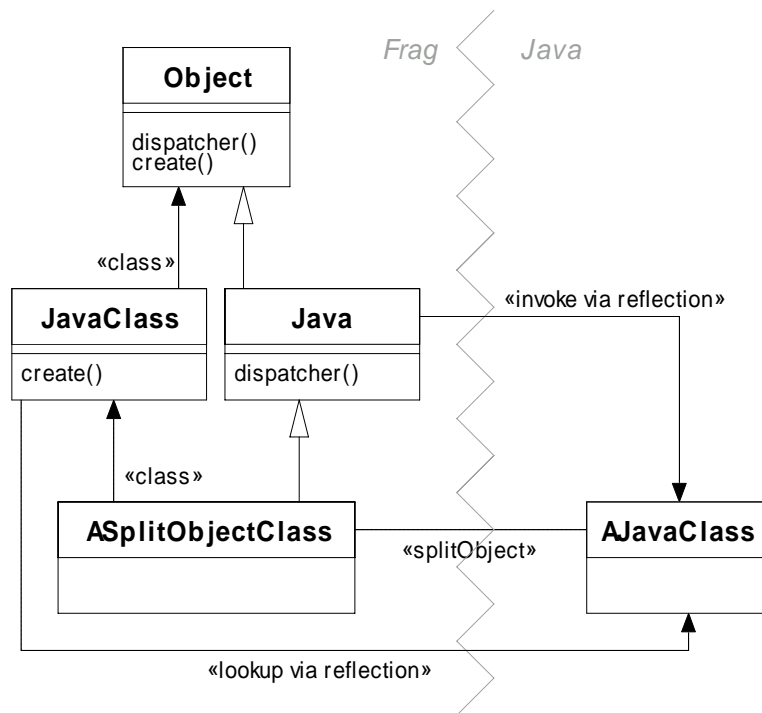


Figure 1.8

Frag Split object structures example

When the goal is to configure a Java application, typically the `Frag` script are used from within Java. However, Java invocations require the developer to know which objects are SPLIT OBJECTS. Consider a Java object `circle` is defined as a SPLIT OBJECT and the following Java invocation is performed:

```
circle.setRadius(2.0);
```

This Java invocation bypasses the `COMMAND LANGUAGE Frag`. What would be needed instead is an invocation sent through the `COMMAND LANGUAGE`:

```
frag.eval("circle setRadius 2.0");
```

The pattern `HOOK INJECTOR` [Zdun2003b] provides a solution. For different languages different tools exist that can perform static hook injections. For instance, for Java we can use `AspectJ`. The following aspect associates statically with a `Frag INTERPRETER`. It contains one advice that is invoked before the constructors of user-defined classes. It calls `makeSplitObject` to create a SPLIT OBJECT half in `Frag`. All other methods of the user-defined classes are intercepted by an `around` advice. The method `invokeSplitObject` sends the invocation to the SPLIT OBJECT half first, and if `next` is invoked, it is sent to the Java implementation as well.

```

abstract aspect FragSplitObject {
    static Frag frag;
    ...
    protected static void makeSplitObject(
        JoinPoint jp, Object o) {...}
    protected static Object invokeSplitObject(
        JoinPoint jp, Object o) {...}

    abstract pointcut splitObjectClasses(Object obj);
    pointcut theConstructors(Object obj):
        splitObjectClasses(obj) && execution(new(...));
    pointcut theMethods(Object obj):
        splitObjectClasses(obj) &&
        execution(* *(..)) &&
        !execution(String toString());
    before(Object obj): theConstructors(obj) {
        makeSplitObject(thisJoinPoint, obj);
    }
    Object around(Object obj) : theMethods(obj) &&
        !cflow(execution(* invokeSplitObject(..))) {
        return invokeSplitObject(thisJoinPoint, obj);
    }
}

```

Note that this aspect excerpt is simplified. The actual aspect implementation in the Frag distribution is more complex because it treats all Java primitive types by separate pointcuts and advices. (This terminology stems from AspectJ and is explained in detail in [Kiczales+2001]).

The aspect is defined as an abstract aspect. In concrete aspects the pointcut `splitObjectClasses` is refined. Here, the classes can be defined by the user to which the aspect is applied. For instance, we can apply the split objects to the classes `Circle` and `Square`:

```

public aspect FragSplitObjectCircleSquare
    extends FragSplitObject {
    pointcut splitObjectClasses(Object obj):
        this(obj) &&
        (within(Circle) || within(Square));
    ...
}

```

The default behavior of this aspect is that all invocations are sent through the `COMMAND LANGUAGE`, but the invocations are not altered. By re-defining the classes in the dynamic Frag language, one can dynamically compose the internals of the aspect. Thus, here we have applied the `SPLIT OBJECT` pattern for integrating AspectJ aspect with the Frag `COMMAND LANGUAGE`. That is, we can compose and configure aspect implementations dynamically at runtime. The same Frag code used for specifying a Java aspect can also be used to specify an aspect in another language (e.g. Tcl, C, or C++). We only have to use another aspect framework for weaving the aspect.

Semantic lookup service

In many distributed systems, servers offer services to clients. Time and again, services get added or removed. Clients need to know which services are currently offered by which server application, running at which location. To allow for efficient, inexpensive, and timely publication of services, the LOOKUP pattern is used as a service by most distributed object frameworks (see [Völter+2004]). Typical implementations of lookup services, such as naming or trading services, can be queried using a name (or a set of properties). As a result of the lookup, the ABSOLUTE OBJECT REFERENCE [Völter+2004] of a remote object matching the query is returned, or other location information such as an INTERFACE DESCRIPTION [Völter+2004] like a WSDL-file.

Today, we see many new requirements for lookup services arising that are due to new kinds of applications for these services and new lookup requirements. For instance, Web services frameworks and other emerging kinds of middleware, such as P2P systems, GRIDs, or spontaneous networks, require highly dynamic lookup repositories with services that can connect, disconnect, move, and change their properties dynamically. In many existing lookup implementations, however, the properties and relationships of elements in the lookup service are too simple or hard to extend. Also, more advanced lookup functionality might be needed, such as ontology-based queries, behavioral extension, query extension, load-balancing, and fault tolerance. This requires a high flexibility regarding inputs, processing, and outputs of the lookup service.

This section discusses a semantic lookup service [Zdun2004b] which is used to cope with these challenges. The basic idea is to make the lookup service itself an extensible Semantic Web repository and make it programmable at runtime. We offer the Redland RDF store [Beckett2004] remotely in which we can define resources and resource relationships according to different Semantic Web ontologies, such as RDFS, OWL, Dublin Core, or LOM. In contrast to most existing solutions in the Semantic Web area, we do not use a logic reasoner, but an OBJECT SYSTEM LAYER as part of a COMMAND LANGUAGE which provides a dynamic, imperative interface to the RDF repository.

As a solution to these problems, we use mobile code that is sent from the client to the server, and an INTERPRETER on server side for interpreting the lookup queries. In particular, this architecture consists of the following components:

- All metadata about the services is stored in an *RDF store* (the Redland RDF store [Beckett2004]), which stores semantic metadata about Web resources described in some ontology or schema.
- The RDF store is accessed on server side using a *script interpreter* (the XOTcl [Neumann+2000] INTERPRETER is used here). Queries are executed in a (safe) sub-interpreter that can read the data from the RDF store. Each script interpreter runs in its own thread of control to make the architecture scalable and efficient.
- A *remote query interface* is offered so that remote clients can send scripts to the lookup service. The lookup service evaluates the scripts using its embedded script interpreter.

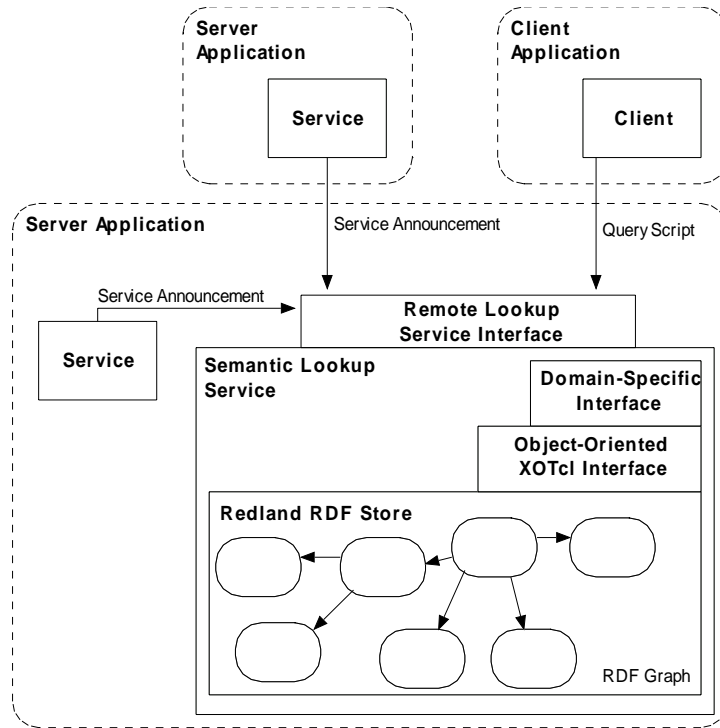


Figure 1.9

Architecture overview of the semantic lookup service

Figure 1.9 shows how these components are assembled to a remote service architecture, and how it is used by services and clients for lookups. Redland is implemented as a set of C libraries. XOTcl is used as an OBJECT SYSTEM LAYER, which provides an object-oriented interface to the Redland store. Scripts can use this interface for their queries. The class diagram in Figure 1.10 gives an overview of the object-oriented interface to the Redland RDF store library. Each class depicted in this figure is a COMPONENT WRAPPER that wraps one of the Redland C libraries.

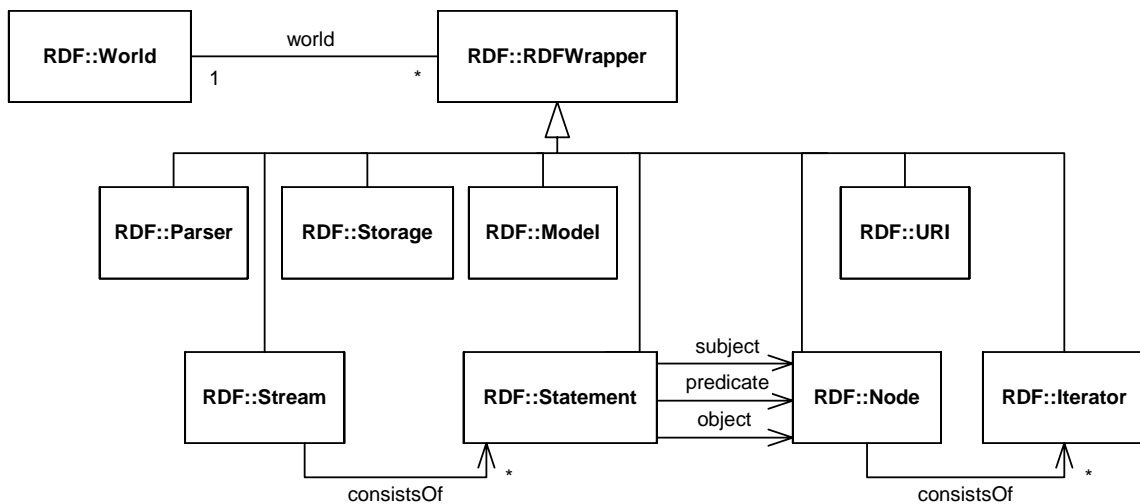


Figure 1.10

Object-oriented RDF store interface

Based on this infrastructure we can build the lookup and query models. In fact, the basic lookup model is just the object-oriented COMPONENT WRAPPER interface. The basic model can be used for simple RDF graphs with ABSOLUTE OBJECT REFERENCES, properties, and relationships. In the case of ordinary Web services, we simply announce the services with a service name, their WSDL INTERFACE DESCRIPTION, and other information like the service's URL. More sophisticated descriptions might contain properties provided as domain-specific ontologies. For instance, information in the WSDL-file (like operation names, supported protocols, etc.) can be reflected in the RDF graph as well – then this information can be included in queries.

The benefit of this architecture is that clients can construct complex queries in an imperative, object-oriented script language. For instance, the following simple query example produces a list of all the contents of all statements that have a creator (according to the Dublin Core ontology) and the value of the creator is “Jim”:

```
set results ""
Statement queryStatement
queryStatement predicate "http://purl.org/dc/elements/1.1/creator"
queryStatement object "Jim"
set stream [m findStatements queryStatement]
while {![$stream end]} {
  lappend results [[$stream current] toString]
  $stream proceed
}
set results
```

It is also easy to extend the server-side INTERPRETER with pre-defined query methods. Just consider the example above: Here, the client knows what to do with the information that someone is the creator of a Web service. Thus, if such queries are recurring, we can provide the implementation of query semantics on server side so that the client does not have to provide the query logic with each request, but can reuse the pre-defined methods instead. Such an interface must then be loaded into the server-side INTERPRETER before the query happens. Of course, this kind of server-side extension is only needed for convenience. Any kind of query can also be defined at client side and be sent with each query – thus the query model is fully extensible without the need for changes in the server.

Apache Axis

We want to discuss the Web services framework Apache Axis as an example of a distributed object framework. In many other distributed object frameworks similar concerns occur and similar solutions are applied.

The MARSHALLER [Völter+2004] of the Web services framework needs to convert the types of remote invocations into strings to be transported across the network using the SOAP protocol, and on server side it needs to convert the strings into the respective types again. In Axis the XML Schema Data (XSD) types are used, and WSDL files are used as an INTERFACE DESCRIPTION. Thus Axis only requires one-to-one conversion between programming language types and the

string representation of the XSD types. This is done using an AUTOMATIC TYPE CONVERTER. This way type conversion is flexibly extensible with new type converters.

In Axis a type converter consists of a factory for a serializer and de-serializer, as well as an implementation of serializer and de-serializer. Besides primitive types which can use a simple converter because type conversion is supported by Java, Axis can convert many more complex types, such as arrays, dates, calendars, DOM documents, object references, vectors, beans, and others. In a class containing the type bindings, these AUTOMATIC TYPE CONVERTERS are mapped to the XSD types. According to the INTERFACE DESCRIPTION, Axis can lookup the Java type converter corresponding to an XSD type and perform the conversion automatically.

A Web services framework usually supports heterogeneity of backends implementing the service. For instance, a service might be implemented as a Java class, a procedural library, an EJB component, a COM component, and many others. Axis provides a “provider” abstraction. One provider class for each type of backend service is provided. The provider can be dynamically configured for particular invocations, and the framework can be extended with new providers for new types. This is a variant of the pattern COMPONENT WRAPPER: each backend component type can be flexibly integrated, and we can manipulate the mapping by either providing new provider types or by configuring the provider objects.

A distributed object framework is an implementation of a canonical object system, often implemented in many languages. Thus it can be used as an OBJECT SYSTEM LAYER to provide objectified access to non-object-oriented languages, as well as a unified access to the object systems of the language supported. The object concepts of this OBJECT SYSTEM LAYER are usually described by a language for INTERFACE DESCRIPTIONS, such as WSDL for Web services and IDL for CORBA.

Conclusion

In this chapter, we have presented a number of patterns for software architecture integration from the specific areas of language integration and component integration. The patterns play a central role in a huge number of integration architectures. There are many external patterns and many patterns missing in this broad area yet in order to provide a full pattern language for this important field.

Acknowledgments

Many thanks to Charles Weir for his attention as a EuroPLoP 2004 shepherd and his constructive comments. Also, thanks to the participants of the EuroPLoP 2004 writers’ workshop who provided valuable feedback.

References

[Bass+1998] L. Bass, P. Clement, and R. Kazman. *Software Architecture in Practice*. Addison-Wesley, Reading, USA, 1998.

[Beckett2004] Dave Beckett. *Redland RDF application framework*. 2004. Available:
<http://www.redland.opensource.ac.uk/>

[Box+2000] D. Box, D. Ehnebuske, G. Kakivaya, A. Layman, N. Mendelsohn, H. F. Nielsen, S. Thatte, and D. Winer. *Simple object access protocol (SOAP) 1.1*. 2000. Available:
<http://www.w3.org/TR/SOAP/>

[Buschmann+1996] F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad, and M. Stal. *Pattern-oriented Software Architecture - A System of Patterns*. J. Wiley and Sons Ltd., 1996.

[DeJong +2003] M. DeJong and S. Redman. *Tcl Java Integration*. 2003. Available:
<http://www.tcl.tk/software/java/>

[ETSI2001] ETSI. *MHP specification 1.0.1*. ETSI standard TS101-812, October 2001.

[Gamma+1995] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns-Elements of Reusable Object-Oriented Software*. Reading, MA: Addison-Wesley, 1995.

[Goedicke+2000] M. Goedicke, G. Neumann, and U. Zdun. Object system layer. In *Proceedings of EuroPlop 2000*, Irsee, Germany, July 2000.

[Goedicke+2001] M. Goedicke, G. Neumann, and U. Zdun. Message Redirector. In *Proceedings of EuroPlop 2001*, Irsee, Germany, July 2001.

[Goedicke+2002] M. Goedicke and U. Zdun. Piecemeal Legacy Migrating with an Architectural Pattern Language: A Case Study. In: *Journal of Software Maintenance and Evolution: Research and Practice*, 14:1-30, 2002.

[Goedicke+2004] M. Goedicke, C. Koellmann, and U. Zdun. Designing runtime variation points in product line architectures: Three cases. In: *Science of Computer Programming*, 53(3):353-380, 2004.

[HAVI2001] HAVI. *HAVI specification 1.1.*, May 2001. Available:

<http://www.havi.org>

[Johnson+1998] R. Johnson and B. Woolf. Type object. In R. C. Martin, D. Riehle, and F. Buschmann, editors, *Pattern Languages of Program Design 3*. Addison-Wesley, 1998.

[Kiczales+2001] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. G. Griswold. Getting started with AspectJ. *Communications of the ACM*, October 2001.

[Kiczales +1997] G. Kiczales, J. Lamping, A. Mendhekar, C. Maeda, C. V. Lopes, J. M. Loingtier, and J. Irwin. Aspect-oriented programming. In *Proceedings of ECOOP'97*, Finland, June 1997. LCNS 1241, Springer-Verlag.

[Neumann+2000] G. Neumann and U. Zdun. XOTcl, an object-oriented scripting language. In *Proceedings of Tcl2k: The 7th USENIX Tcl/Tk Conference*, Austin, Texas, USA, February 2000.

[Oberleitner+2003] J. Oberleitner, T. Gschwind, and M. Jazayeri. Vienna component framework enabling composition across component models. In *Proceedings of the International Conference on Software Engineering (ICSE 2003)*, Portland, Oregon, USA, May 2003.

[OpenMash2000] Open Mash Consortium. *The Open Mash consortium*. 2000. Available:

<http://www.openmash.org>

[Pawlak+2001] R. Pawlak, L. Seinturier, L. Duchien, and G. Florin. JAC: a flexible framework for AOP in Java. In *Reflection 2001: Meta-level Architectures and Separation of Crosscutting Concerns*, Kyoto, Japan, Sep 2001.

[Schmidt+2000] D. C. Schmidt, M. Stal, H. Rohnert, and F. Buschmann. *Patterns for Concurrent and Distributed Objects*. Pattern-Oriented Software Architecture. J. Wiley and Sons Ltd., 2000.

[Spinellis2001] D. Spinellis. Notable design patterns for domain specific languages. *Journal of Systems and Software*, 56(1):91–99, February 2001.

[Swig2003] Swig Project. *Simplified wrapper and interface generator*. 2003. Available:

<http://www.swig.org/>

[UCB2000a] UCB Multicast Network Research Group. *Network simulator - ns (version 2)*. 2000.

Available:

<http://www.isi.edu/nsnam/ns/>

[UCB2000b] UCB Multicast Network Research Group. *Tclcl*. 2000. Available:

<http://www.isi.edu/nsnam/tclcl/>

[Völter+2004] M. Völter, M. Kircher, and U. Zdun. *Remoting Patterns - Foundations of Enterprise, Internet, and Realtime Distributed Object Middleware*. Wiley Series in Software Design Patterns. J. Wiley & Sons. 2004.

[Vogel+2002] O. Vogel and U. Zdun. Content conversion and generation on the web: A pattern language.

In *Proceedings of EuroPlop 2002*, Irsee, Germany, July 2002.

[Wetherall+1995] D. Wetherall and C. J. Lindblad. Extending Tcl for dynamic object-oriented programming. In *Proc. of the Tcl/Tk Workshop '95*, Toronto, July 1995.

[Yoder+2002] J.W. Yoder and R. Johnson. The Adaptive Object-Model Architectural Style. In: *3rd Working IEEE/IFIP Conference on Software Architecture (WICSA'2002)*, Montreal, Canada, August, 2002.

[Zdun2003a] U. Zdun. *Frag*. 2003. Available:
<http://frag.sourceforge.net/>

[Zdun2003b] U. Zdun. Patterns of tracing software structures and dependencies. In *Proceedings of EuroPlop 2003*, Irsee, Germany, June 2003.

[Zdun2004a] U. Zdun. Using split objects for maintenance and reengineering tasks. In *Proceedings of 8th European Conference on Software Maintenance and Reengineering (CSMR'04)*, Tampere, Finland, Mar 2004.

[Zdun2004b] U. Zdun. Semantic lookup in service-oriented architectures. In *Proceedings of Fourth International Workshop on Web-Oriented Software Technologies*, pages 101-110, Munich, Germany, July 2004.