

The Metacommand Pattern

(c) 2000 Markus Voelter, voelter@acm.org Version 1.2, 28.05.2000

INTENT

Metacommand enhances the Command pattern by providing a facility to enhance or modify the common behaviour of the command classes used in a system at any time without modifying the command classes themselves.

CONTEXT AND MOTIVATION

As described by the Gang-of-Four in [1], you can use the Command pattern to

Encapsulate a request as an object, thereby letting you parameterize clients with different requests, queue or log requests, and support undoable operations.

However, using Command requires that you know these additional features when you implement the command classes, i.e. at design/implementation time. It is not possible to dynamically add or remove features. For example, you might want to add logging facilities to all command classes in a system if debugging becomes necessary or your customer requires logging for security reasons. Or you might want to add a permissions check before executing a command, executing the command only if the check permits the execution.

The Metacommand pattern provides this additional flexibility by allowing you to add these features dynamically at runtime, without modifying the command classes themselves. The last point is especially important, because in general, there are quite many command classes in a system.

The pattern resolves the following forces:

- The common behaviour of your command classes should be kept flexible, enabling you to change common behaviour without modifying the command classes themselves. This should be achieved without requiring big changes to the programming model of the original Command pattern.
- This additional flexibility should ideally come without decreasing performance.
- Changing the common behaviour should not require changes allthrough the system. Changes should be localized.
- Retrofitting an existing system with these additional features should be simple. The necessary changes should be systematic to allow automatic conversion.

PROBLEM

How do you create a system that allows you to easily modify the common behaviour of command objects, even at runtime?

SOLUTION

Wrap the created command object with a *MetaCommand* object upon creation. Use a system-wide creation policy to determine which concrete *MetaCommand* class is used to wrap the commands. Use a factory to take care of wrapping. When a command is executed, the *MetaCommand* is executed instead. It adds its own behaviour before in turn executing the contained command.

Structure

Just as in the original Command pattern, the *Command* class is an abstract base class for the specific commands that you use in the application. It has an abstract *execute()* operation that gets called by the command-executing program entity, no matter what the concrete command class is.

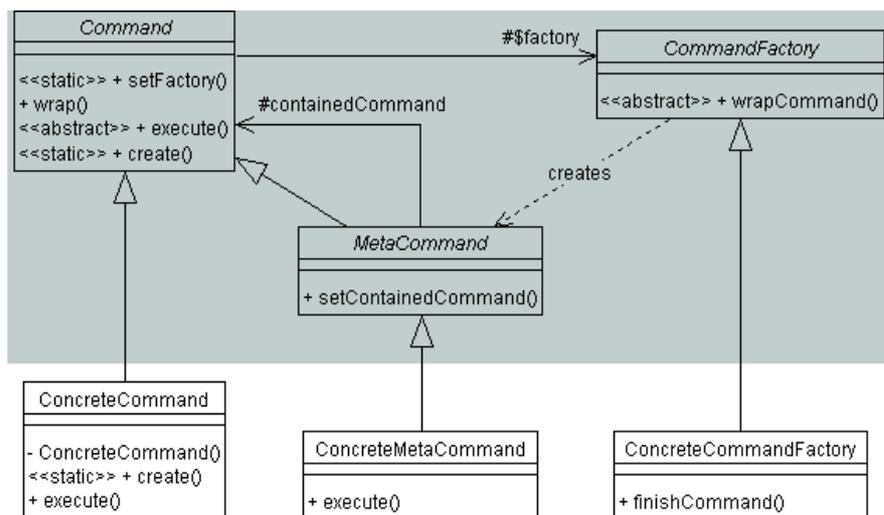


Illustration 1: Class diagram of the MetaCommand pattern. The classes in the grey box form the abstract base of the pattern, whereas the classes in the white part have to be created specifically for an application.

In addition, the *Command* class contains a static reference to a well-known *CommandFactory* object (according to the Abstract Factory pattern, [1]), which is responsible for automatically wrapping a *ConcreteMetaCommand* object around a *ConcreteCommand* upon creation. *MetaCommand* also inherits from *Command*, so that it can be used wherever a *Command* is expected. An application defines several *ConcreteMetaCommand* classes, the *ConcreteCommandFactory* determines, which of these classes will be used to wrap a specific *ConcreteCommand* object. You have to make sure that a *ConcreteMetaCommand* invokes the contained *ConcreteCommand*'s *execute()* operation during its own execution.

Creation of *ConcreteCommand* objects and the wrapping with a *ConcreteMetaCommand* is handled by one or more static *create()* operations (Factory Methods, see [1]) in the

ConcreteCommand classes. Please note that these operations formally return *Command*, not the *ConcreteCommand* in which they are declared. This is because the concrete class of the returned object depends on the factory's decision which *ConcreteMetaCommand* is wrapped around the command. (Note: You cannot use a constructor here, because a constructor always returns an instance of the class in which it is declared. A factory method like the *create()* operation is more flexible in this respect. That's why a factory method is also termed a virtual constructor). The create operation will usually look something like this (Java example)

```

01 class ConcreteCommand extends Command {
02     public static Command create( String anArgument ) {
03         ConcreteCommand c = new ConcreteCommand();
04         // process arguments with c
05         c.anArgument = anArgument;
06         // call wrap to wrap the created ConcreteCommand
07         return c.wrap();
08     }
09     public void execute() ....
10 }

```

The important part of this code snippet is line 7: After initialization, the *wrap()* operation is called on the new object before it is returned. This creates a *ConcreteMetaCommand* and wraps the *ConcreteCommand* (*c* in the example) with it. The operation returns this newly created *ConcreteMetaCommand*. Details follow below.

Participants

<i>Command</i>	Abstract base class for application-specific concrete command classes. It contains the abstract declaration for the <i>execute()</i> operation.
<i>MetaCommand</i>	Abstract base class for application specific <i>MetaCommand</i> classes. Wraps around (contains) a <i>ConcreteCommand</i> object.
<i>CommandFactory</i>	Responsible for setting up the wrapping between the <i>ConcreteCommand</i> and the <i>ConcreteMetaCommand</i> .
<i>ConcreteMetaCommand</i>	An application-specific <i>MetaCommand</i> class. Executes the <i>ConcreteCommand</i> during its own execution. It contains code common to all <i>ConcreteCommands</i> .
<i>ConcreteCommandFactory</i>	A concrete implementation of <i>CommandFactory</i> . Responsible for wrapping an instance of a <i>ConcreteMetaCommand</i> around a <i>ConcreteCommand</i> .
<i>ConcreteCommand</i>	An application-specific command class. Defines one or more <i>create()</i> operations, which work as described above.

Interactions

Initialization and command creation

The abstract class *Command* contains a static reference to an instance of a *CommandFactory* object which must be set at the beginning of a program by calling the *Command* class's static *setFactory()* operation with an instance of a *ConcreteCommandFactory* as the only parameter.

To create a *ConcreteCommand* object, you have to call the static *create()* operation on the *ConcreteCommand* class, of which you need an instance. As demonstrated above, this operation then creates an instance of its own class, processes all its arguments, and then calls *wrap()* on the newly created object.

The *wrap()* operation calls the *wrapCommand()* operation on the well-known factory that you have set before, passing the *ConcreteCommand* as a parameter. The factory determines which *ConcreteMetaCommand* should be used, creates an instance of it, and then sets the passed *ConcreteCommand* to be the contained command of the newly created *ConcreteMetaCommand*. This metacommmand (*mc*) is then returned to the client application. From the client's view, this looks like the following:

```
Command cmd = MyCommand.create();
```

Note that the type of the variable storing the returned object (*cmd*) is *Command*, not *MyCommand* because the concrete type of the returned object depends on the factory.

The following diagram illustrates the previously described processes:

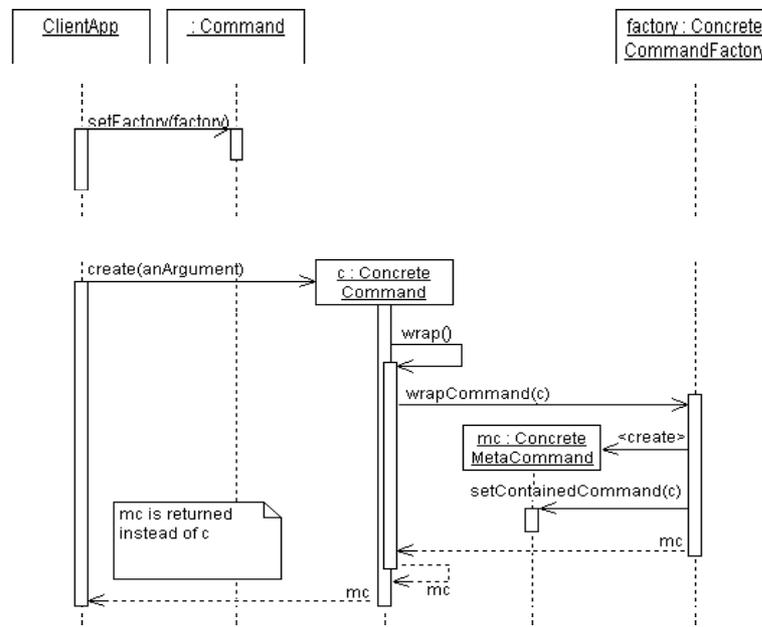


Illustration 2: Setting up the factory and creating a *ConcreteCommand* object.

Executing the command

Just as in the conventional Command pattern, you execute a command by invoking its *execute()* operation. But because the *create()* operation returned the *ConcreteMetaCommand*, the invocation reaches the *ConcreteMetaCommand's execute()*, which in turn calls the *ConcreteCommand's execute()* after adding its own behaviour. This is shown in the diagram below

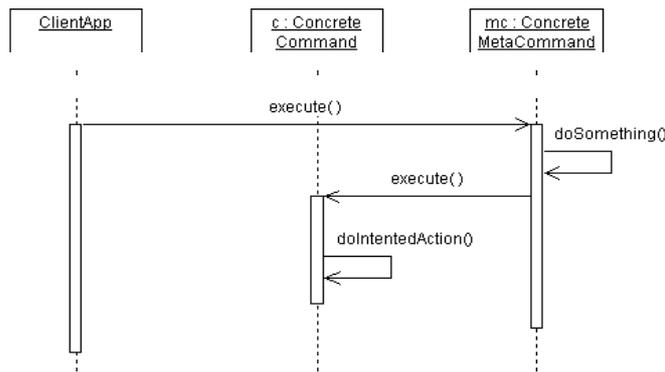


Illustration 3: Executing a command in fact executes the *ConcreteMetaCommand*, which, after executing its own code, calls the *ConcreteCommand*'s execute operation.

CONSEQUENCES

- You can change the common behaviour of all commands by exchanging the *ConcreteCommandFactory*. No change to the *ConcreteCommand* classes is necessary. The change is localized to the factory. The programming model is not very different from the original command pattern, you only have to use the *create()* operation instead of the constructor (and provide the *CommandFactories* and *MetaCommands*).
- This additional flexibility has introduced an additional level of indirection (A *ConcreteCommand*'s *execute()* is called through *ConcreteMetaCommand*'s *execute()*). This might impose a slight performance overhead. The same is true for the creation process, when two objects must be created instead of just one.
- You can easily retrofit an existing application with this increased flexibility. The only required change is where the command objects are created, and some „infrastructure“: the factory class, and the *create()* operations. However, the handling of created command objects and their execution remains the same.
- The pattern relies on the wrapping *MetaCommand* object to call the *execute()* operation of the contained *Command*. However, this cannot easily be enforced. In some cases this is intended (e.g. not executing a *Command* in case of a failed permissions check in the *MetaCommand*). On the other hand, it is possible to actually modify the behaviour of *Command* classes, and not just to add behaviour. This might be a problem in certified / well tested / safety relevant systems.

EXAMPLES AND SAMPLE CODE

The simplest case

Let's begin with a very simple implementation of the *Command* class. Java is used for the examples.

```

public abstract class Command {
    protected static CommandFactory factory;

    public static void setFactory(CommandFactory cf) {
        factory = cf;
    }
}
  
```

```

public Command wrap() {
    if ( factory == null ) return this;
    return factory.wrapCommand( this );
}

public abstract void execute();
}

```

The Command class above has the important property, that whenever no factory is set, the *wrap()* operation returns *this*, i.e. it returns the object on which it was called. So whenever no factory is set, the system behaves just like the standard Command pattern (without any performance overhead). The check whether the *factory* is *null* is implemented very efficiently in Java, and imposes nearly no performance overhead.

Logging

The above example might be the initial implementation for a software system. Later, when the system is deployed, you might want to log all executed commands. You can achieve this by executing the following three-step process:

1. A suitable *MetaCommand* class is created. It is called *LogMetaCommand* in the example and could look something like the following:

```

public class LogMetaCommand extends MetaCommand {
    public void execute() {
        System.out.println( (new Date()).toString() +
            " executing : "+containedCommand );
        containedCommand.execute();
    }
}

```

2. Create a factory that wraps the *ConcreteCommands* with a *LogMetaCommand*. This is also very simple and straightforward:

```

public class LogCommandFactory extends CommandFactory {
    public Command wrapCommand(Command c) {
        LogMetaCommand lcmd = new LogMetaCommand();
        lcmd.setContainedCommand( c );
        return lcmd;
    }
}

```

3. Set the factory object in the *Command* class:

```

public class Test {
    public Test() {
        Command.setFactory( new LogCommandFactory() );
        Command t = TestCommand.create( "Hallo" );
        t.execute();
    }
    public static void main(String[] args) {
        Test test = new Test();
    }
}

```

The above piece of code has the consequence, that every command execution in the system gets logged. This is true for command classes already used in the system, as well as for those, that are introduced later.

Permissions

Often, you need to retrofit a system with security features, i.e. that certain commands may only be executed by certain users (or groups of users). Once again, Metaccommand provides a simple solution for this problem by using *MetaCommands* that check the permissions before executing the contained command. The check itself could be delegated to another class:

```
public class PermissionMetaCommand extends MetaCommand {
    public void execute() {
        if ( PermissionManager.instance().allows( containedCommand )
            containedCommand.execute();
        else showMessage( „Required permissions not available!“ );
    }
}
```

Queueing

In some applications, you will find that it makes sense to queue some commands for asynchronous execution. You can use Metaccommand to implement such a queueing facility. Two building blocks are necessary here. A *QueueingMetaCommand* and a *QueueFactory*.

Upon execution, the *QueueingMetaCommand* puts the contained *Command* into a queue which is then processed asynchronously by another thread. The *QueueFactory* determines whether the command should be queued or not. If it should not be queued, then the factory does nothing and returns the command itself. If, on the other hand, the command should be queued, it wraps the command with a *QueueingMetaCommand*.

Note: If the time between enqueueing the commands and their execution is long, and if the factory is replaced during this timespan, then the already enqueuee (and therefore wrapped) command objects will not change their behaviour to reflect the changes imposed by the new factory.

ANOTHER SOLUTION – AND WHY IT DOES NOT WORK

Another solution that comes to mind immediately is simply to provide the *Command* classes with a *before()* and *after()* method. It would be the responsibility of the *ConcreteCommand* classes *execute()* methods to call them. This does not work because:

- You cannot give different groups of *ConcreteCommand* classes different common behaviour in the *before()* and *after()* methods. To do this, you would have to insert an additional layer of abstract command classes which implement common behaviour in the *before()* and *after()* methods. But because you usually do not know the grouping of the *ConcreteCommands* from the beginning, you would have to change the class hierarchy constantly, every time, you want to change the grouping.

- You cannot change anything at runtime, except by using state dependent case statements in the command classes *before()* and *after()* methods.
- You cannot implement permissions and queueing cleanly this way because the result of the *before()* method does not influence the behaviour of the main *execute()* method. You could create such a behaviour by using exceptions somehow, but this is not a very practical solution. (Example, Queueing: The *before()* method could place itself (*this*) into a queue, then set a flag, which is tested by the *execute()* method after its *before()* call. If set, the method returns. When *execute()* is once again called, this time by the thread that processes the queue, *before()* must not be called again, instead the core of the *execute()* operation must be run. This can also be achieved by using a couple of flags. But all in all, this is not a very elegant solution).
- Composability (as described in the next section) is not possible.

VARIANTS AND IMPROVEMENTS

Multiple Layers

You can compose multiple layers of *MetaCommands*. If your system contains a *LogMetaCommand* class, a *QueueingMetaCommand* class and a *PermissionMetaCommand* class, then you can easily build systems that have all features just by creating another factory:

```
public class LQPCommandFactory extends CommandFactory {
    public Command wrapCommand(Command c) {
        LogMetaCommand lcmd = new LogMetaCommand();
        lcmd.setContainedCommand( c );
        QueueingMetaCommand qcmd = new QueueingMetaCommand();
        qcmd.setContainedCommand( lcmd );
        PermissionMetaCommand pcmd = new PermissionMetaCommand();
        pcmd.setContainedCommand( qcmd );
        return pcmd;
    }
}
```

The result of the work of this factory is the following containment structure:

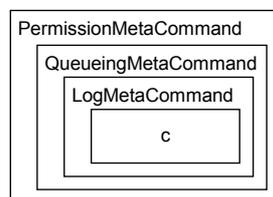


Illustration 4: Containment structure if multiple *MetaCommands* are wrapped.

Upon execution, the commands are executed “outside-in”. First, the *PermissionMetaCommand* checks whether the command is allowed to be executed. If so, the command is put in a queue. Upon execution, a log message is created and then, finally, the command itself gets executed. Note, that the order wrap order of the commands is significant: If, for example, permission checking is contained inside of logging, the command will be logged although perhaps it will not be executed. The

same is true for queueing and permissions. There is no use of checking a permission after it has been inserted into the queue (because they will be executed asynchronously).

Multiple wrapped Factories

If you use many different combinations of wrapping *MetaCommands* in an application, it might be impractical to create a corresponding factory for each combination. Instead, the *CommandFactory* class can be modified to contain a reference to another factory. The factories could recursively call *wrapCommand()* on the referenced factory, thereby adding multiple layers of *MetaCommands* around the original *ConcreteCommand*.

Controlling the Factory

If there is no general common policy about which *MetaCommand* the factory should use, then it might make sense to add additional *create...(...)* operations on the *Command* classes that have other names or additional parameters. This makes it possible for the creator of a command to determine the behaviour of the factory at runtime, because the factory can use the supplied parameters to determine whether a *MetaCommand* will be set or not, and to pass additional parameters to the *MetaCommand*.

RELATED PATTERNS

The Metacommand pattern of course has a close relationship to the Command pattern – it is an extension. The Command pattern is described in the classic Gang of Four book [1]. In addition, Metacommand makes use of the Abstract Factory pattern and the Factory Method, which are also described in [1]. The containment structure of the commands (and of the factories, see VARIANTS AND IMPROVEMENTS) can be seen as an implementation of the Decorator pattern [1].

Another pattern that extends Command and addresses quite similar concerns is Peter Sommerlad's Command Processor pattern [2]. In the Command Processor pattern, commands are not executed directly. Instead, they are executed by a command processor. This component can be used to implement additional code, for logging, queueing etc. So, instead of wrapping *MetaCommand* objects around each *ConcreteCommand* to implement common behaviour, the Command Processor pattern locates this common code into the command processor component.

The decision which of the two patterns should be used, could be guided by the following observation:

- The Command Processor leaves the creation of commands unchanged and requires modifications at all locations where commands are executed. So whenever it is not possible (or feasible) to modify the creation of command objects, the Command Processor pattern should be used.
- The MetaCommand leaves the execution of commands unchanged but requires changes in the code at all locations, where commands are created. So whenever it is not possible (or feasible) to modify the execution of commands, the MetaCommand pattern should be used.

Because of the above, the Metacommand pattern as well as the Command Processor pattern can be used in refactorization projects, where it is necessary to introduce unanticipated features into a software system.

OTHER RELATED WORK

It is also interesting to have a look at the relationship to aspect-oriented programming (AOP, [3]). AOP supports the definition of cross-cutting concerns of a software system in a single entity called an *aspect*. Aspects can introduce new methods into (one or more) classes, and can add *before* and *after* code to existing methods. Modifying existing methods this way is called „advising a method“.

You can use this to create an aspect that advises the *execute()* operations of all *ConcreteCommand* objects, thereby achieving the same effect as the code within the *MetaCommands*. AOP requires a special tool (called an *aspect weaver* for AspectJ, an AOP implementation for Java, [4]). *MetaCommand* allows similar features to be implemented without the availability of aspects (but requiring more work!).

You could also view the *Metacommand* pattern as a way to dynamically change the class of the commands, because arbitrary behaviour can be added or removed at any time. Another way to achieve this is by using meta level programming, if it is available in the language in use. This is where the pattern borrows its name.

KNOWS USES

The pattern was used in the ThoughtPad application [5], a tool create topic maps. Here, the pattern has been applied after the program has been finished (it already used the conventional Command pattern for all user interface actions). As described above, only the creation of the commands had to be changed. The *MetaCommands* have been used to add permission checks to the program.

IBM's PRODIKOS project uses the *Metacommand* pattern in its user interface architecture. The system will later be integrated with Lotus Workflow, although it is not yet clear at the beginning of the project how this integration will look like. A single *metacommand* will be used for all command objects, that analyses the concrete command and then triggers the correct process in Lotus Workflow.

ACKNOWLEDGEMENTS

Many thanks to my shepherd Wolfgang Keller for giving many useful hints and comments on how to improve the paper. It has improved significantly! I'd also like to thank the participants of MATHEMA's design patterns courses (Feb. 14 - Feb. 21, 2000) for discussing this paper with me.

REFERENCES

- [1] Gamma, Helm, Johnson, Vlissides; *Design Patterns, elements of reusable software design*, Addison-Wesley 1996
- [2] Peter Sommerlad, *Command Processor*, in *Pattern Languages of Program Design 2*, Addison-Wesley 1996
- [3] Xerox PARC, *AOP homepage* at <http://www.parc.xerox.com/csl/projects/aop>
- [4] Xerox PARC, *AspectJ homepage* at <http://aspectj.org>
- [5] voelterSOFTWARE, *ThoughtPad homepage* at <http://www.voelter.de/thoughtpad>

