

Abstract

In this work, we present examples for metaprogramming using `thisContext` and mixins and their implementation in Squeak/Pharo, both of which are dialects of Smalltalk.

1 Reflection and Metaprogramming in Smalltalk

Metaprogramming is a program's ability to reason about itself. In an abstract sense, this means that programs can be treated as ordinary data and, therefore, be analyzed and changed at runtime. In this section, we show an example for Smalltalk's `thisContext`, which is a mechanism to control the bytecode interpreter.

1.1 Stack Machine Model

Smalltalk source code is compiled to bytecode by a compiler written in Smalltalk itself. The bytecode is then executed by a virtual machine. Smalltalk's bytecode follows the stack machine model. When a message send should be performed, the compiler generates bytecode that pushes the receiver and all arguments onto the stack, followed by a `send:` instruction with the selector.

1.2 Accessing Stack Frames

Smalltalk provides a special keyword `thisContext` evaluating to an object representing the current stack frame. That object has methods for accessing and modifying stack variables (e.g., local variables), modifying the program counter and jumping to an arbitrary address within the executing method, and accessing the sender's stack frame. `thisContext` is an instance of class `MethodContext` which lets programmers effectively control the bytecode interpreter within the guest language. The compiler translates `thisContext` to a bytecode instruction that generates an instance of `MethodContext` and pushes it onto the stack.

Protocol The following list gives an overview of some interesting methods provided by `MethodContext`¹.

- `jump::` Modifies the program counter to perform a relative jump in the bytecode of the current method.
- `method:` Returns the `CompiledMethod` object of the executing method. This object contains meta information about the method and its bytecode.
- `tempAt::` Accesses a temporary variable of this stack frame using its index.
- `tempAt:put::` Sets a temporary variable of this stack frame using its index.
- `pc:` Returns the current program counter, an index into the bytecode of the method of this stack frame.
- `pc::` Sets the current program counter.
- `push::` Pushes a value onto the stack, effectively increasing the size of this stack frame.
- `receiver:` Returns the receiver of this stack frame.
- `return::` Causes the current stack frame to return with a certain value.
- `sender:` Returns the stack frame below this frame.
- `swapReceiver::` Sets the receiver of this stack frame.

¹Some methods are stored in a superclass of `MethodContext`, but we only mention `MethodContext` in this work.

Optimizations In the presence of a just-in-time compiler, a virtual machine might never generate a `MethodContext` object unless the programmer uses `thisContext`. Some of the methods shown below first try to use a primitive in the virtual machine (for performance reasons) and execute the shown code only as a fallback, in case the virtual machine does not provide an implementation for the primitive. Nevertheless, the programmer can change these methods at any time to force that the Smalltalk code is executed instead of the implementation in the virtual machine.

1.3 Example: Exceptions Implemented in the Guest Language

The mechanism for traversing the method stack and controlling the bytecode interpreter can be used to implement exception handling in the guest language, such that the underlying virtual machine does not have to be aware of that mechanism. The benefit of this implementation approach is that more functionality can be implemented in high-level Smalltalk code, resulting in a *smaller* virtual machine. This is not only a *cleaner* approach with respect to architectural design², but also us to use the guest language as a playground for new language features, since it is no longer necessary to recompile the virtual machine or to even restart the running Smalltalk system for a language modification to take effect.

Exception Handling Exceptions are raised in Smalltalk by sending the message `signal` to an exception object. An exception is caught by the first exception handler contained in a stack frame. The message `on:do:` can be sent to a block closure containing the code throwing an exception with the type of exception and an exception handler block as arguments.

```

1 HttpRequest»titechWebsiteContent
2   [ ↑ self httpGet: 'http://www.titech.ac.jp/' ]
3     on: Exception
4     do: [ :e | ↑ '<b>Unable to send HTTP GET request' ]
5
6 HttpRequest»httpGet: aURL
7   "... "
8   timedOut ifTrue: [ Exception new signal ].
9   "... "
```

If we run the execute `titechWebsiteContent` in the previous code snippet and the network host cannot be reached, `httpGet:` will `signal` (throw) an exception which will be handled by the exception handler block in the first method.

Implementation of Exception Handling This paragraph describes the implementation of exception handling in Squeak/Smalltalk using `thisContext`. The following source code snippets are taken from a Squeak 5.0 image and simplified. A number of subtle details are omitted such as checking if an exception handler should handle a certain type of exception.

The method `Exception»signal` uses the method `nestHandlerContext` to find a stack frame with an exception handler. To detect such a stack frame, the method `on:do:` starts with a primitive call. That primitive is not implemented in the virtual machine, i.e., when that method is executed, the primitive fails immediately and the code after the primitive statement is executed. The primitive only acts as a method marker (annotation). Note that in method `signal` the keyword `thisContext` refers to the stack frame executing the `signal` method. There is a stack frame calling that method on the stack below that frame, and we assume that there is also a stack frame containing an exception handler somewhere below that frame.

```

1 Exception»signal
```

²In the best case, we would like to have all functionality implemented in the guest language itself.

```

2  "Ask ContextHandlers in the sender chain to handle this signal. The default is to execute
   and return my defaultAction."
3  ↑ thisContext nextHandlerContext handleSignal: self
4
5  BlockClosure»on: exception do: handlerAction
6  "Evaluate the receiver in the scope of an exception handler."
7  <primitive: 199> "just a marker, fail and execute the following"
8  ↑ self value

```

The method `nextHandlerContext` finds the next stack frame containing an exception handler by iterating through all stack frames until one stack frame is marked, i.e., it contains a primitive call with number 199. A different approach could check if the method of a stack frame is the compiled method object `BlockClosure»on:do:`, but checking the primitive number might be more efficient.

```

1  MethodContext»nextHandlerContext
2  "Return the next handler marked context, returning nil if there is none. Search starts with
   self and proceeds up to nil."
3  | ctx |
4  ctx := self.
5  [ ctx isHandlerContext ifTrue:[ ↑ ctx ].
6  (ctx := ctx sender) == nil ] whileFalse.
7  ↑ nil
8
9  MethodContext»isHandlerContext
10 "Is this a context for a method that is marked?"
11 ↑ method primitive = 199

```

Once a stack frame handling exceptions was found, the handler must be executed and the method containing the handler must return. Notice that the `return:` message send in the following source code snippet is not a regular method return statement but a method defined on class `MethodContext`. It causes that method to return with a certain result, regardless of where the program counter points to. This automatically terminates all stack frames on top of that frame. The method `tempAt:` is used to retrieve the second temporary variable, which is the second argument to `on:do:` (exception handler block). The method `cull:` tries to execute the exception handler block with the exception object as argument or without any arguments in case the block does not take any arguments.

```

1  MethodContext»handleSignal: exception
2  self return: ((self tempAt: 2) cull: exception)

```

Implementation without Metaprogramming Exception handling is typically implemented in the virtual machine/interpreter. Stack frames can be marked as exception handlers by setting a flag similarly to the primitive call in the example above. Raising an exception translates to a primitive call or a special bytecode instruction, upon which the virtual machine traverses the stack of method frames until it finds one that is marked. This mechanism is very similar to the mechanism described above. Smalltalk is special in a sense that stack frames are guest language object and accessible and modifiable in the guest language.

It is not obvious how to implement exception handling in the guest language without using metaprogramming. One very tedious approach would have every method return a tuple of the actual return value and an optional exception object. As soon as a method call returns, the sender first checks if the tuple contains an exception object and, if so, returns immediately with that exception object as well. Otherwise, it proceeds with the execution, possibly using the actual return value of the called method. This mechanism requires modifying every return statement and every method call, but it might be possible to do this transformation automatically using macros.

2 Mixins in Smalltalk

A mixin is an abstract subclass that can be applied to multiple (super)classes. Mixins are typically used to share methods that are common to multiple classes, such that the source code does not have to be duplicated. Most Smalltalk dialects do not support mixins out of the box³, but it is easy to implement rudimentary mixin functionality in Squeak. The last part of this section describes how to do that.

2.1 Protocol

Classes are defined in Squeak using a message send to the superclass. The following snippet defines a subclass of `Object`.

```
1 Object subclass: #NewClass
2   instanceVariableNames: 'foo bar'
3   classVariableNames: 'qux'
```

The following listing shows how to apply three mixins during class definition. A mixin is an ordinary class but not meant to be instantiated.

```
1 Object subclass: #NewClass
2   instanceVariableNames: 'foo bar'
3   classVariableNames: 'qux'
4   mixins: { Mixin1. Mixin2. Mixin3 }
```

2.2 Example: Comparable Mixin

In this example, we assume that an application needs two classes `Time` and `Money`. Since our application should be deployed in an international environment, class `Time` must be aware of time zones and class `Money` should support multiple currencies. For that reason, we do not want to use temporal and numeric classes provided by the execution environment.

The following listing shows how these two classes are defined.

```
1 Object subclass: #Time
2   instanceVariableNames: 'hour minute second timeZone'
3   classVariableNames: ".
4
5 Object subclass: #Money
6   instanceVariableNames: 'amount currency'
7   classVariableNames: ".
```

A frequent operation in our application involves comparing instances of `Time` and instances of `Money`. Therefore, both classes should understand the methods `<`, `<=`, `>`, `>=`, `=`, and `≠` (inequality). We first present an implementation without mixins and then an implementation with mixins.

Without Mixins The following source code snippet shows how to implement `Time` without mixins. `Time` has a method `gmtTime` which returns the time in seconds according to the GMT time zone. Methods for comparing two `Time` instances compare this value.

```
1 Time>>gmtTime
2   ↑ (self hour * 3600 + self minute * 60 + self second - self timeZone gmtDifference * 3600) ←
   \ \ 86400
```

³Newspeak is similar to Smalltalk and supports Mixins.

```

3
4 Time>< other
5   ↑ self gmtTime < other gmtTime
6
7 Time><= other
8   ↑ self gmtTime <= other gmtTime
9
10 Time>> other
11  ↑ self gmtTime > other gmtTime
12
13 Time>>= other
14  ↑ self gmtTime >= other gmtTime
15
16 Time>= other
17  ↑ self gmtTime = other gmtTime
18
19 Time>~ other
20  ↑ self gmtTime ~ other gmtTime

```

The following source code snippet shows how to implement `Money` without mixins. `Money` has a method `toUSD` which returns the amount in US dollars according to current exchange rate. Methods for comparing two `Money` instances compare this value. Note that methods for comparing instances are similar in `Time` and `Money`. In the next paragraph, we will get rid of this code duplication using mixins.

```

1 Money>toUSD
2   ↑ self amount * (WebRequest queryRate: self currency to: 'USD').
3
4 Money>< other
5   ↑ self toUSD < other toUSD
6
7 Money><= other
8   ↑ self toUSD <= other toUSD
9
10 Money>> other
11  ↑ self toUSD > other toUSD
12
13 Money>>= other
14  ↑ self toUSD >= other toUSD
15
16 Money>= other
17  ↑ self toUSD = other toUSD
18
19 Money>~ other
20  ↑ self toUSD ~ other toUSD

```

With Mixins We first define a mixin `Comparable` providing methods for comparing instances of any kind of class, given that the class provides implementations for `=` and `>`. Based on these two methods, the remaining for methods can be implemented as follows.

```

1 Object subclass: #Comparable
2   instanceVariableNames: "

```

```

3   classVariableNames: ".
4
5   Comparable»= other
6     self subclassResponsibility.
7
8   Comparable»> other
9     self subclassResponsibility.
10
11  Comparable»< other
12    ↑ (self = other | (self > other)) not
13
14  Comparable»<= other
15    ↑ (self > other) not
16
17  Comparable»>= other
18    ↑ self = other | (self > other)
19
20  Comparable»~ other
21    ↑ (self = other) not

```

We now define `Time` and `Money` as subclasses of `Object` with the mixin `Comparable`. This means that these classes are subclasses of a mixin application of `Comparable`, which is a subclass of `Object`. Notice that we only have to implement the methods `>` and `=` along with the converter methods `toUSD` and `gmtTime`, removing the code duplication partly.

```

1  Object subclass: #Time
2    instanceVariableNames: 'hour minute second timeZone'
3    classVariableNames: ".
4    mixins: { Comparable }.
5
6  Object subclass: #Money
7    instanceVariableNames: 'amount currency'
8    classVariableNames: ".
9    mixins: { Comparable }.
10
11  Time»> other
12    ↑ self gmtTime > other gmtTime
13
14  Time»= other
15    ↑ self gmtTime = other gmtTime
16
17  Money»> other
18    ↑ self toUSD > other toUSD
19
20  Money»= other
21    ↑ self toUSD = other toUSD

```

2.3 Implementation of Mixins

The following listing describes how mixins can be implemented in Squeak using metaprogramming. We provide a new method that takes an additional collection of mixins during subclassing. For every mixin, the algorithm generates a new subclass with the instance/class variables of the mixin and adds the methods of the mixin to that subclass. This is possible because Smalltalk allows generating new classes and adding new methods from source code while a program is running. Class-side methods and instance variables are added to the meta class object, which can be obtained by sending the message `class` to the class object⁴.

```

1 Class class»subclass: name instanceVariableNames: instVarNames classVariableNames: classVars
   mixins: mixinClasses
2 | result |
3 result := self
4 mixinClasses withIndexDo: [ :cls :idx |
5     result := result subclass: name, '_', idx asString
6         instanceVariableNames: cls instVarNames
7         classVariableNames: cls classVarNames.
8     cls methodDict do: [ :sel :meth | result compile: meth getSource ].
9     cls class methodDict do: [ :sel :meth | result class compile: meth getSource ] ].
10 ↑ result subclass: name
11     instanceVariableNames: instVarNames
12     classVariableNames: classVars

```

⁴For that reason we write `Classname class»methodName` to denote class-side methods.