



A low Overhead Per Object Write Barrier --- for the Cog VM

Clément Béra





Introduction

- The Cog VM is the standard VM for:
 - Pharo
 - Squeak
 - Newspeak
 - Cuis



Introduction

- Working runtime optimizer for Cog's JIT
- Problem with literal mutability




Problem

- Is it possible to mark any object as read-only ?
- Smalltalk code to handle mutation failure
- Overhead



Terminology

- Discussion on VM mailing-list
 - Immutable: state cannot change after object's initialization
 - Write barrier or read-only object



Use-cases

- Modification tracker
- Read-only literals
 - Compiler optimizations
 - Inconsistent literal modifications
- Others...



This paper

- NOT about framework built using read-only objects
- Implementation details to limit the overhead




Feature

- Any object can be marked as read-only, except:
 - Immediate objects
 - Context instances
 - Objects related to Process scheduling
 - Objects internal to the runtime




APIs

- Object >> isReadOnlyObject
- Object >> setIsReadOnlyObject:
 - Object >> beWritableObject
 - Object >> beReadOnlyObject



Read-only object

- Instance variable store fail
- Primitives mutating a read-only object fail



IV store failure

```
| message |  
message := Message selector: #foo.  
message beReadOnlyObject.  
message setSelector: #bar.  
message
```

- Instance variable is not set.
- A call-back is sent:

```
message attemptToAssign: value withIndex: index
```




Primitive failure

```
| array |  
array := Array with: 1.  
array beReadOnlyObject.  
array at: 1 put: 2.  
array
```


- First value of array is not set



Primitive error code

```
instVarAt: index put: anObject  
  <primitive: 174 error: ec>  
self primitiveFailed
```

- new error code: `#'no modification'`



Other details

- Support flags


`Smalltalk vm supportsWriteBarrier`

- Mirror primitives
 - `Object >> object:setIsReadOnlyObject:`




VM compilation option

- VM C compiler flag
- The VM can be compiled with or without the feature.



Implementation

- Object representation
- Interpreter support
- JIT support



Implementation

- Most critical part:
 - How to keep IV store efficient ?
 - Machine code generated by the JIT
- Discussed in the paper...



IV Store details

x86 Assembly	Meaning
<code>movl -12(%ebp), %edx</code>	Load the receiver in %edx.
<code>popl %edi</code>	Load the value to store in %edi.
<code>movl %edi, %ds:0x8(%edx)</code>	Perform the store in the first instance variable using both registers (%edx and %edi)
<code>testl 0x00000003, %edi</code>	If the value to store is immediate, jump after the store check.
<code>jnz after_store</code>	
<code>movl 0x00040088, %eax</code>	Jump after the store check if the receiver is young: compare the young object space limit with receiver address
<code>cmpl %eax, %edx</code>	
<code>jb after_store</code>	
<code>cmpl %eax, %edi</code>	If the value to store is an old object, jump after the store check.
<code>jnb after_store</code>	
<code>movzbl %ds:0x3(%edx), %eax</code>	If the receiver is already in the remembered table, jump after the store check.
<code>testb 0x20, %al</code>	
<code>jnz after_store</code>	
<code>call store_check_trampoline</code>	Calls the store check trampoline.
<code>after_store:</code>	Code following the store.

- Wanted
- to show it,




x86 Assembly	Meaning
<code>movl -12(%ebp), %edx</code>	Load the receiver in %edx.
<code>popl %ecx</code>	Load the value to store in %ecx.
<code>movl %ds:(%edx), %eax</code>	If the receiver is read-only, jump to the store trampoline.
<code>testl 0x00800000, %eax</code>	
<code>jnz store_trampoline</code>	
<code>movl %ecx, %ds:0x8(%edx)</code>	Perform the store in the first instance variable using both registers (%edx and %ecx)
<code>testb 0x03, %cl</code>	If the value to store is immediate, jump after the store check.
<code>jnz after_store</code>	
<code>movl 0x00040088, %eax</code>	If the receiver is a young object, jump after the store check.
<code>cmpl %eax, %edx</code>	
<code>jb after_store</code>	
<code>cmpl %eax, %ecx</code>	If the value to store is an old object, jump after the store check.
<code>jnb after_store</code>	
<code>movzbl %ds:0x3(%edx), %eax</code>	
<code>testb 0x20, %al</code>	If the receiver is already in the remembered table, jump after the store check.
<code>jnz after_store</code>	
<code>store_trampoline:</code>	Calls the store check trampoline.
<code>call store_trampoline</code>	
<code>movl -12(%ebp), %edx</code>	Restore the receiver (to keep its register live).
<code>after_store:</code>	Code following the store.



Evaluation: Slow-down


- Binary trees
 - IV Store intensive
 - No significant difference



Evaluation: Slow-down


- Pathological case: setter

```
left: leftChild right: rightChild item: anItem  
left := leftChild.  
right := rightChild.  
item := anItem
```

Evaluation: Slow-down

- At writing time, setter overhead was 17%
- Stack frame creation problem
 - Two path compilation
 - Now faster than before



Conclusion

- New feature: read-only object
- Overhead is very limited