

Under the Hood: The Java Virtual Machine, Part II

Announcements

- A5 due date postponed!
- New due date Monday Dec 8, 11:59pm

2

CS2110 F08 Quiz 3

Answer briefly the following questions.

1. Explain the difference between shadowing and overriding.
2. Explain the difference between interfaces and abstract classes.
3. Explain the difference between static and dynamic types.

3

CS2110 F08 Quiz 3

Answer briefly the following questions.

1. Explain the difference between shadowing and overriding.

Shadowing occurs when a field of a class has the same name and type as a field of a superclass, or when a local variable of a class has the same name and type as a field of that class. It is useful only in limited circumstances (e.g. initializing a field with a parameter). Overriding occurs when a method of a class has the same name and signature as a method of a superclass, and is one of the most useful features of object oriented languages.

2. Explain the difference between interfaces and abstract classes.

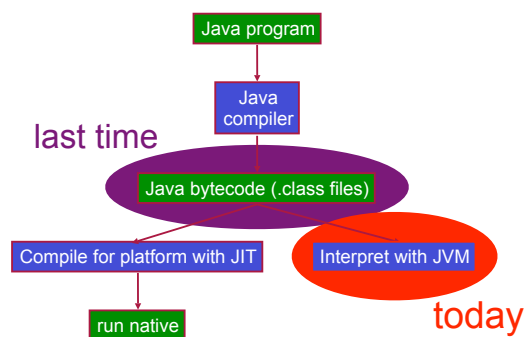
Interfaces may contain only constants and abstract methods (method signatures but no implementations). Abstract classes may contain fields and both abstract and concrete methods. Neither interfaces nor abstract classes can be directly instantiated with new.

3. Explain the difference between static and dynamic types.

Expressions have static types. The type of any expression is known to the Java compiler at compile time, before the program is run. Objects have dynamic types. Dynamic types are only known at runtime. An object receives its dynamic type when it is created with new.

4

last time



5

Today

- Class file format
- Class loading and initialization
- Object initialization
- Method dispatch
- Exception handling
- Java security model
 - Bytecode verification
 - Stack inspection

6

Instance Method Dispatch

`x.foo(...)`

- compiles to `invokevirtual`
- Every loaded class knows its superclass
 - name of superclass is in the constant pool
 - like a parent pointer in the class hierarchy
- bytecode evaluates arguments of `x.foo(...)`, pushes them on the stack
- Object `x` is always the first argument

7

Instance Method Dispatch

`invokevirtual foo(...)`

- Name and type of `foo(...)` are arguments to `invokevirtual` (indices into constant pool)
- JVM retrieves them from constant pool
- Gets the dynamic (runtime) type of `x`
- Follows parent pointers until finds `foo(...)` in one of those classes – gets bytecode from code attribute

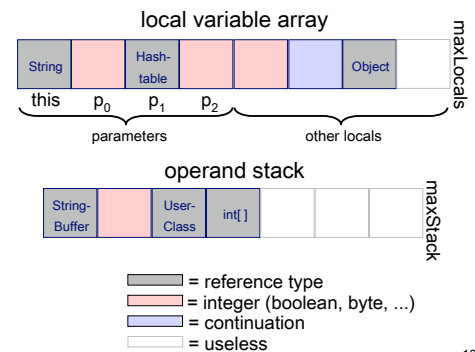
8

Instance Method Dispatch

- Creates a new *stack frame* on runtime stack around arguments already there
- Allocates space in stack frame for locals and operand stack
- Prepares locals (int=0, ref=null), empty stack
- Starts executing bytecode of the method
- When returns, pops stack frame, resumes in calling method after the `invokevirtual` instruction

9

Stack Frame of a Method



10

Instance Method Dispatch

```
byte[] data;
void getData() {
    String x = "Hello world";
    data = x.getBytes();
}
```

```
Code(maxStack = 2, maxLocals = 2, codeLength = 12)
0: ldc "Hello world"
2: astore_1
3: aload_0 //object of which getData is a method
4: aload_1
5: invokevirtual java.lang.String.getBytes () [B
8: putfield A.data [B
11: return
```

11

Exception Handling

- Each method has an *exception handler table* (possibly empty)
- Compiled from `try/catch/finally`
- An exception handler is just a designated block of code
- When an exception is thrown, JVM searches the exception table for an appropriate handler that is in effect
- **finally** clause is executed last

12

Exception Handling

- Finds an exception handler → empties stack, pushes exception object, executes handler
- No handler → pops runtime stack, returns exceptionally to calling routine
- **finally** clause is always executed, no matter what

13

Exception Table Entry

startRange	start of range handler is in effect
endRange	end of range handler is in effect
handlerEntry	entry point of exception handler
catchType	exception handled

- **startRange** → **endRange** give interval of instructions in which handler is in effect
- **catchType** is any subclass of **Throwable** (which is a superclass of **Exception**) -- any subclass of **catchType** can be handled by this handler

14

Example

```
Integer x = null;
Object y = new Object();

try {
    x = (Integer)y;
    System.out.println(x.intValue());
} catch (ClassCastException e) {
    System.out.println("y was not an Integer");
} catch (NullPointerException e) {
    System.out.println("y was null");
} finally {
    System.out.println("finally!");
}
```

15

```
0: aconst null
1: astore_1
2: new java.lang.Object
5: dup
6: invokeSpecial java.lang.Object.<init> ()V
9: astore_2
10: aload_2
11: checkCast java.lang.Integer
14: astore_1
15: getstatic java.lang.System.out Ljava/io/PrintStream;
18: aload_1
19: invokevirtual java.lang.Integer.intValue ()I
22: invokevirtual java.io.PrintStream.println (I)V
25: getstatic java.lang.System.out Ljava/io/PrintStream;
28: ldc "finally!"
30: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
33: goto #89
36: astore_3
37: getstatic java.lang.System.out Ljava/io/
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println (
45: getstatic java.lang.System.out Ljava/io/
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore_4
78: getstatic java.lang.System.out Ljava/io/
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println
86: aload 4
88: athrow
89: return
```

From	To	Handler Type
10	25	java.lang.ClassCastException
10	25	java.lang.NullPointerException
10	25	<Any exception>
36	45	<Any exception>
56	65	<Any exception>
76	78	<Any exception>

```
Integer x = null;
Object y = new Object();

try {
    x = (Integer)y;
    System.out.println(x.intValue());
} catch (ClassCastException e) {
    System.out.println("y was not an Integer");
} catch (NullPointerException e) {
    System.out.println("y was null");
} finally {
    System.out.println("finally!");
}
```

16

```
0: aconst null
1: astore_1
2: new java.lang.Object
5: dup
6: invokeSpecial java.lang.Object.<init> ()V
9: astore_2
10: aload_2
11: checkCast java.lang.Integer
14: astore_1
15: getstatic java.lang.System.out Ljava/io/PrintStream;
18: aload_1
19: invokevirtual java.lang.Integer.intValue ()I
22: invokevirtual java.io.PrintStream.println (I)V
25: getstatic java.lang.System.out Ljava/io/PrintStream;
28: ldc "finally!"
30: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
33: goto #89
36: astore_3
37: getstatic java.lang.System.out Ljava/io/
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println (
45: getstatic java.lang.System.out Ljava/io/
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore_4
78: getstatic java.lang.System.out Ljava/io/
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println
86: aload 4
88: athrow
89: return
```

From	To	Handler Type
10	25	java.lang.ClassCastException
10	25	java.lang.NullPointerException
10	25	<Any exception>
36	45	<Any exception>
56	65	<Any exception>
76	78	<Any exception>

```
Integer x = null;
Object y = new Object();

try {
    x = (Integer)y;
    System.out.println(x.intValue());
} catch (ClassCastException e) {
    System.out.println("y was not an Integer");
} catch (NullPointerException e) {
    System.out.println("y was null");
} finally {
    System.out.println("finally!");
}
```

17

```
0: aconst null
1: astore_1
2: new java.lang.Object
5: dup
6: invokeSpecial java.lang.Object.<init> ()V
9: astore_2
10: aload_2
11: checkCast java.lang.Integer
14: astore_1
15: getstatic java.lang.System.out Ljava/io/PrintStream;
18: aload_1
19: invokevirtual java.lang.Integer.intValue ()I
22: invokevirtual java.io.PrintStream.println (I)V
25: getstatic java.lang.System.out Ljava/io/PrintStream;
28: ldc "finally!"
30: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
33: goto #89
36: astore_3
37: getstatic java.lang.System.out Ljava/io/
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println (
45: getstatic java.lang.System.out Ljava/io/
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore_4
78: getstatic java.lang.System.out Ljava/io/
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println
86: aload 4
88: athrow
89: return
```

From	To	Handler Type
10	25	java.lang.ClassCastException
10	25	java.lang.NullPointerException
10	25	<Any exception>
36	45	<Any exception>
56	65	<Any exception>
76	78	<Any exception>

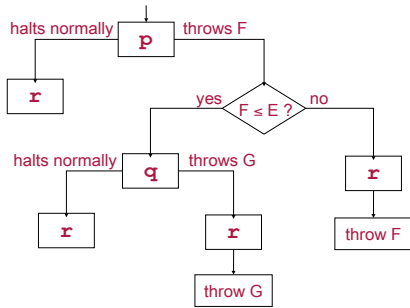
```
Integer x = null;
Object y = new Object();

try {
    x = (Integer)y;
    System.out.println(x.intValue());
} catch (ClassCastException e) {
    System.out.println("y was not an Integer");
} catch (NullPointerException e) {
    System.out.println("y was null");
} finally {
    System.out.println("finally!");
}
```

18

Try/Catch/Finally

`try {p} catch (E) {q} finally {r}`



25

Java Security Model

- Bytecode verification
 - Type safety
 - Private/protected/package/final annotations
 - Basis for the entire security model
 - Prevents circumvention of higher-level checks
- Secure class loading
 - Guards against substitution of malicious code for standard system classes
- Stack inspection
 - Mediates access to critical resources

26

Bytecode Verification

- Performed at load time
- Enforces type safety
 - All operations are well-typed (e.g., may not confuse refs and ints)
 - Array bounds
 - Operand stack overflow, underflow
 - Consistent state over all dataflow paths
- Private/protected/package/final annotations


27

Bytecode Verification

- A form of *dataflow analysis* or *abstract interpretation* performed at load time
- Annotate the program with information about the execution state at each point
- Guarantees that values are used correctly

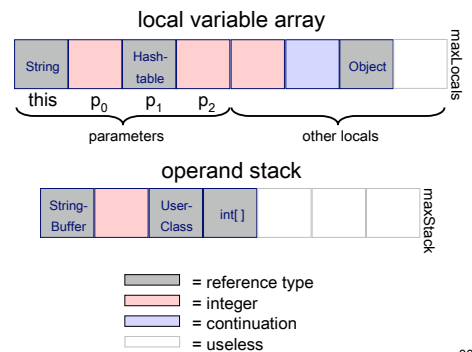
28

Types in the JVM




29

Typing of Java Bytecode



30

Example



Preconditions for safe execution:


- local 3 is an integer
- stack is not full

Effect:

- push integer in local 3 on stack


31

Example




32

Example




33

Example




34

Example




35

Example




36

Example




37

Example



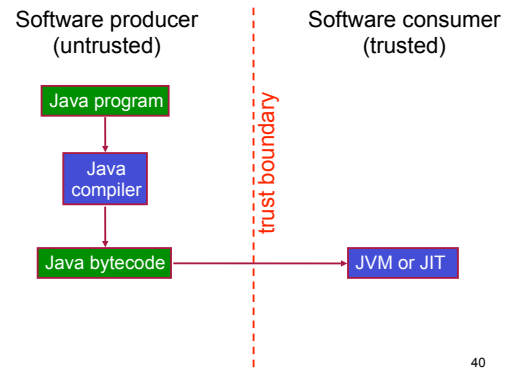
38

Example



39

Mobile Code



40

Mobile Code

Problem: mobile code is not trustworthy!

- We often have *trusted* and *untrusted* code running together in the same virtual machine
 - e.g., applets downloaded off the net and running in our browser
- Do not want untrusted code to perform critical operations (file I/O, net I/O, class loading, security management,...)
- *How do we prevent this?*

41

Mobile Code

Early approach: *signed applets*

- Not so great
 - everything is either trusted or untrusted, nothing in between
 - a signature can only *verify* an already existing relationship of trust, it cannot *create* trust
- Would like to allow untrusted code to interact with trusted code
 - just monitor its activity somehow

42


Mobile Code

Q) Why not just let trusted (system) code do anything it wants, even in the presence of untrusted code?

A) Because untrusted code calls system code to do stuff (file I/O, etc.) – system code could be operating on behalf of untrusted code

43

Runtime Stack



44

Runtime Stack




Maybe we want to disallow it

- the malicious applet may be trying to erase our disk
- it's calling system code to do that

45

Runtime Stack



Or, maybe we want to allow it

- it may just want to write a cookie
- it called `System.cookieWriter`
- `System.cookieWriter` knows it's ok

46


Runtime Stack



Maybe we want to allow it for another reason

- all running methods are trusted

47



Q) How do we tell the difference between these scenarios?

A) *Stack inspection!*

48

Stack Inspection



- An invocation of a trusted method, when calling another method, may either:
 - *permit* R on the stack above it
 - *forbid* R on the stack above it
 - *pass* permission from below (be transparent)
- An instantiation of an untrusted method must *forbid* R above it

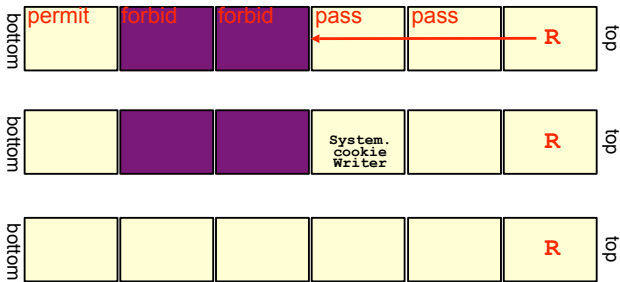
49

Stack Inspection



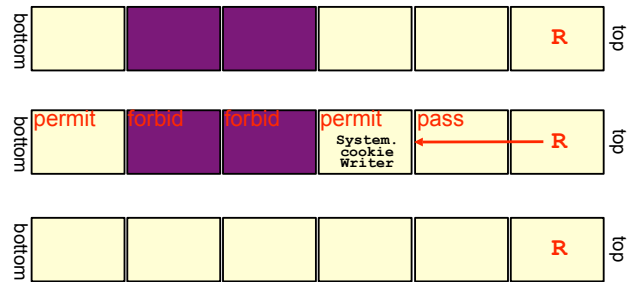
- When about to execute R, look down through the stack until we see either
 - a system method permitting R -- *do it*
 - a system method forbidding R -- *don't do it*
 - an untrusted method -- *don't do it*
- If we get all the way to the bottom, *do it* (IE, Sun JDK) or *don't do it* (Netscape)

50




Case A: R is not executed

51



Case B: R is executed

52



Case C: R is executed

53

Conclusion

Java and the Java Virtual Machine:
Lots of interesting ideas!

54