UNDER THE HOOD: THE JAVA VIRTUAL MACHINE II

CS2110 Fall 2010 Lecture 25

Pick up where we left off

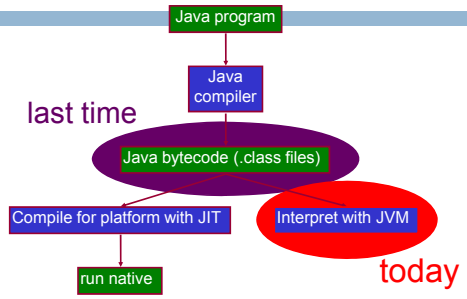

Today

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

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

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

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

Stack Frame of a Method

7

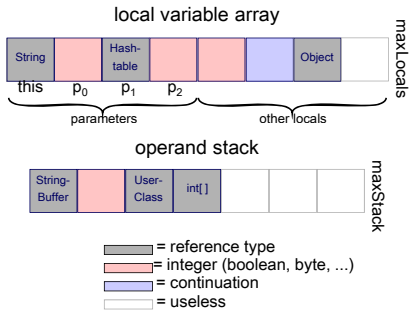

Instance Method Dispatch

8

```
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
```

Exception Handling

9

- 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

Exception Handling

10

- 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

Exception Table Entry

11

`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

Example

12

```
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!");
}
```

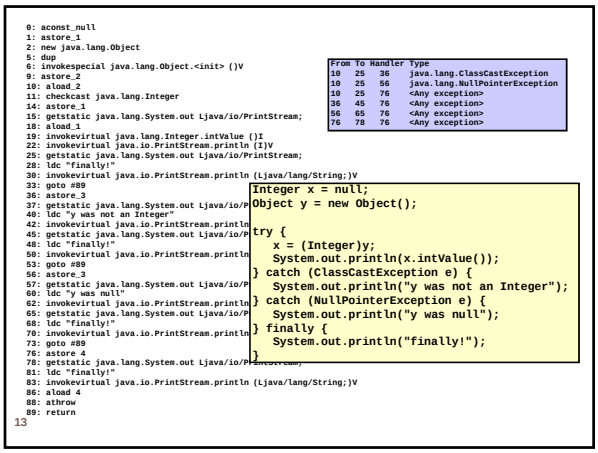
0: aconst_null
1: astore_1
2: new java.lang.Object
5: dup
6: invokespecial java.lang.Object.<init> ()V
9: astore_2
10: aload_2
11: checkcast java.lang.Integer
14: astore_1
15: getstatic java.lang.System.out Ljava/io/PrintStream;
18: aload_1
19: invokevirtual java.lang.Integer.intValue ()I
22: invokevirtual 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/P
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println
45: getstatic java.lang.System.out Ljava/io/P
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/P
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/P
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore 4
78: getstatic java.lang.System.out Ljava/io/PrintStream;
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
86: aload 4
88: athrow
89: return
From To Handler Type
10 25 36 java.lang.ClassCastException
10 25 56 java.lang.NullPointerException
10 25 76 <Any exception>
36 45 76 <Any exception>
56 65 76 <Any exception>
76 78 76 <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!");
}
13

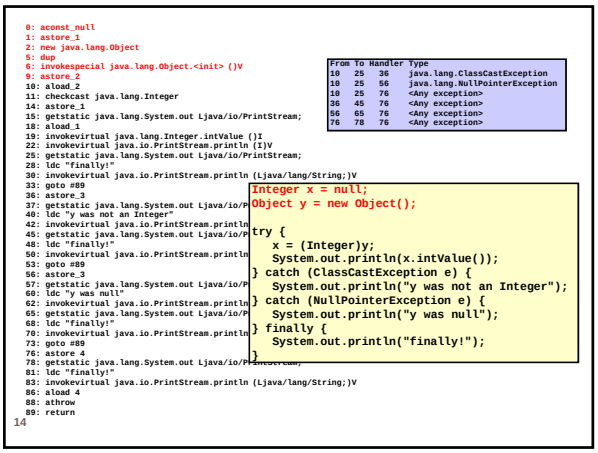
0: aconst_null
1: astore_1
2: new java.lang.Object
5: dup
6: invokespecial java.lang.Object.<init> ()V
9: astore_2
10: aload_2
11: checkcast java.lang.Integer
14: astore_1
15: getstatic java.lang.System.out Ljava/io/PrintStream;
18: aload_1
19: invokevirtual java.lang.Integer.intValue ()I
22: invokevirtual 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/P
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println
45: getstatic java.lang.System.out Ljava/io/P
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/P
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/P
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore 4
78: getstatic java.lang.System.out Ljava/io/PrintStream;
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
86: aload 4
88: athrow
89: return
From To Handler Type
10 25 36 java.lang.ClassCastException
10 25 56 java.lang.NullPointerException
10 25 76 <Any exception>
36 45 76 <Any exception>
56 65 76 <Any exception>
76 78 76 <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!");
}
14

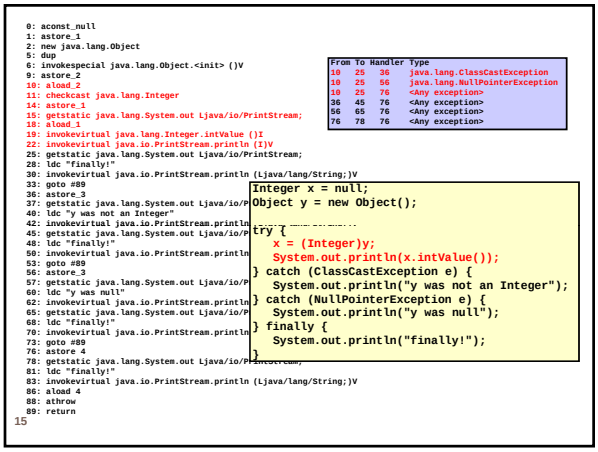
0: aconst_null
1: astore_1
2: new java.lang.Object
5: dup
6: invokespecial java.lang.Object.<init> ()V
9: astore_2
10: aload_2
11: checkcast java.lang.Integer
14: astore_1
15: getstatic java.lang.System.out Ljava/io/PrintStream;
18: aload_1
19: invokevirtual java.lang.Integer.intValue ()I
22: invokevirtual 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/P
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println
45: getstatic java.lang.System.out Ljava/io/P
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/P
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/P
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore 4
78: getstatic java.lang.System.out Ljava/io/PrintStream;
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
86: aload 4
88: athrow
89: return
From To Handler Type
10 25 36 java.lang.ClassCastException
10 25 56 java.lang.NullPointerException
10 25 76 <Any exception>
36 45 76 <Any exception>
56 65 76 <Any exception>
76 78 76 <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!");
}
15

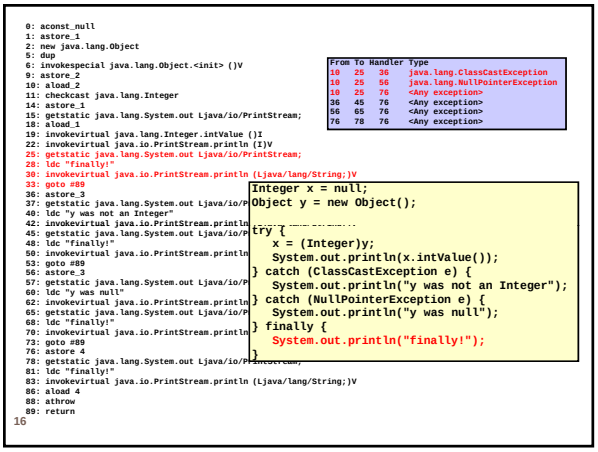
0: aconst_null
1: astore_1
2: new java.lang.Object
5: dup
6: invokespecial java.lang.Object.<init> ()V
9: astore_2
10: aload_2
11: checkcast java.lang.Integer
14: astore_1
15: getstatic java.lang.System.out Ljava/io/PrintStream;
18: aload_1
19: invokevirtual java.lang.Integer.intValue ()I
22: invokevirtual 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/P
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println
45: getstatic java.lang.System.out Ljava/io/P
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/P
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/P
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore 4
78: getstatic java.lang.System.out Ljava/io/PrintStream;
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
86: aload 4
88: athrow
89: return
From To Handler Type
10 25 36 java.lang.ClassCastException
10 25 56 java.lang.NullPointerException
10 25 76 <Any exception>
36 45 76 <Any exception>
56 65 76 <Any exception>
76 78 76 <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

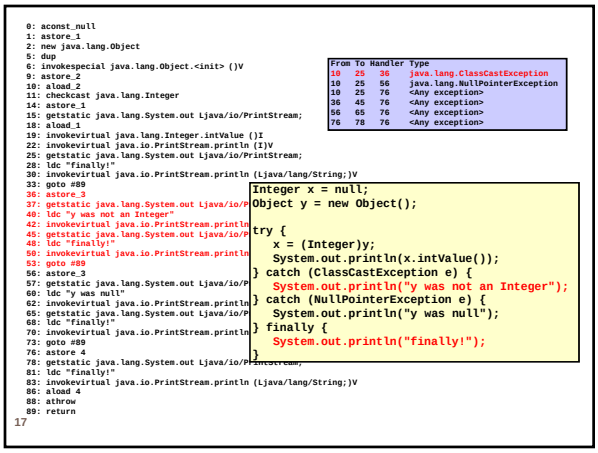
0: aconst_null
1: astore_1
2: new java.lang.Object
5: dup
6: invokespecial java.lang.Object.<init> ()V
9: astore_2
10: aload_2
11: checkcast java.lang.Integer
14: astore_1
15: getstatic java.lang.System.out Ljava/io/PrintStream;
18: aload_1
19: invokevirtual java.lang.Integer.intValue ()I
22: invokevirtual 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/P
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println
45: getstatic java.lang.System.out Ljava/io/P
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/P
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/P
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore 4
78: getstatic java.lang.System.out Ljava/io/PrintStream;
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
86: aload 4
88: athrow
89: return
From To Handler Type
10 25 36 java.lang.ClassCastException
10 25 56 java.lang.NullPointerException
10 25 76 <Any exception>
36 45 76 <Any exception>
56 65 76 <Any exception>
76 78 76 <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

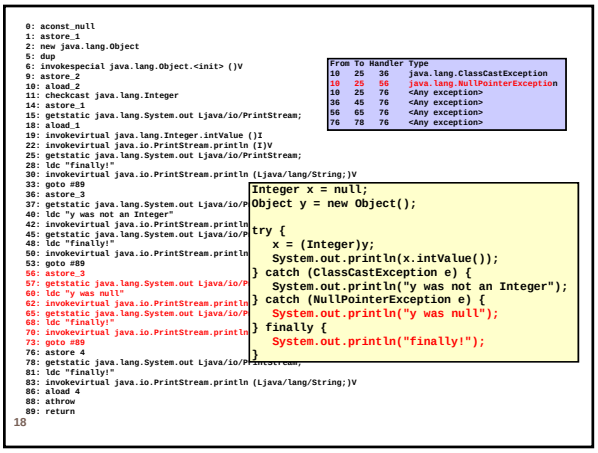
0: aconst_null
1: astore_1
2: new java.lang.Object
5: dup
6: invokespecial java.lang.Object.<init> ()V
9: astore_2
10: aload_2
11: checkcast java.lang.Integer
14: astore_1
15: getstatic java.lang.System.out Ljava/io/PrintStream;
18: aload_1
19: invokevirtual java.lang.Integer.intValue ()I
22: invokevirtual 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/P
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println
45: getstatic java.lang.System.out Ljava/io/P
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/P
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/P
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore 4
78: getstatic java.lang.System.out Ljava/io/PrintStream;
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
86: aload 4
88: athrow
89: return
From To Handler Type
10 25 36 java.lang.ClassCastException
10 25 56 java.lang.NullPointerException
10 25 76 <Any exception>
36 45 76 <Any exception>
56 65 76 <Any exception>
76 78 76 <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

```
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/P
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println
45: getstatic java.lang.System.out Ljava/io/P
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/P
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/P
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore_4
78: getstatic java.lang.System.out Ljava/io/PrintStream;
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
86: aload_4
88: athrow
89: return
```

From	To	Handler	Type
10	25	36	java.lang.ClassCastException
10	25	56	java.lang.NullPointerException
10	25	76	<Any exception>
36	45	76	<Any exception>
56	65	76	<Any exception>
76	78	76	<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!");
}
```

19


```
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/P
40: ldc "y was not an Integer"
42: invokevirtual java.io.PrintStream.println
45: getstatic java.lang.System.out Ljava/io/P
48: ldc "finally!"
50: invokevirtual java.io.PrintStream.println
53: goto #89
56: astore_3
57: getstatic java.lang.System.out Ljava/io/P
60: ldc "y was null"
62: invokevirtual java.io.PrintStream.println
65: getstatic java.lang.System.out Ljava/io/P
68: ldc "finally!"
70: invokevirtual java.io.PrintStream.println
73: goto #89
76: astore_4
78: getstatic java.lang.System.out Ljava/io/PrintStream;
81: ldc "finally!"
83: invokevirtual java.io.PrintStream.println (Ljava/lang/String;)V
86: aload_4
88: athrow
89: return
```

From	To	Handler	Type
10	25	36	java.lang.ClassCastException
10	25	56	java.lang.NullPointerException
10	25	76	<Any exception>
36	45	76	<Any exception>
56	65	76	<Any exception>
76	78	76	<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!");
}
```


20

Try/Catch/Finally

21

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

- **r** is always executed, regardless of whether **p** and/or **q** halt normally or exceptionally
- If **p** throws an exception not caught by the catch clause, or if **q** throws an exception, that exception is *rethrown* upon normal termination of **r**

Try/Catch/Finally

22

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

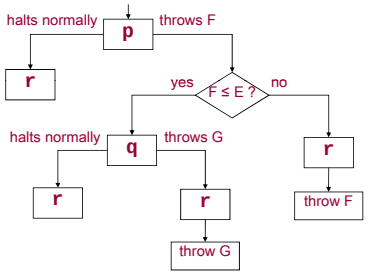

Java Security Model

23

- 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

Bytecode Verification

24

- 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

Bytecode Verification

25

- 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

Types in the JVM

26

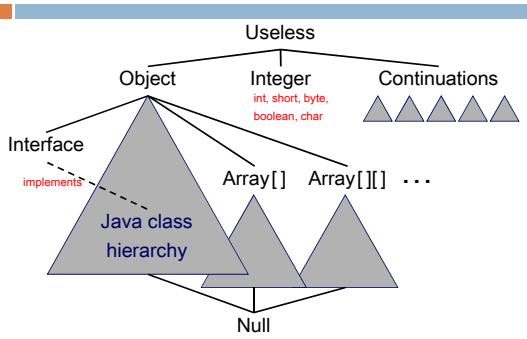

Typing of Java Bytecode

27

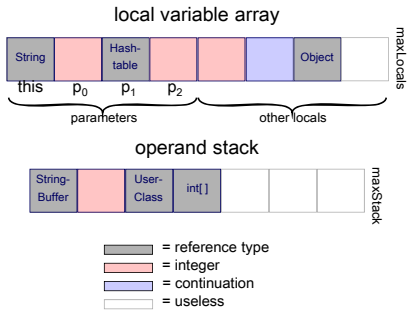

Example

28

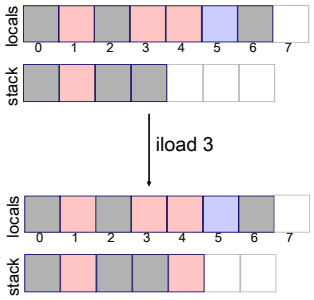

Preconditions for safe execution:

- local 3 is an integer
- stack is not full

Effect:

- push integer in local 3 on stack

Example

29

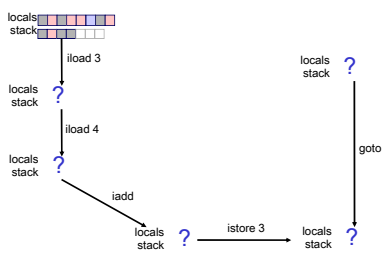

Example

30

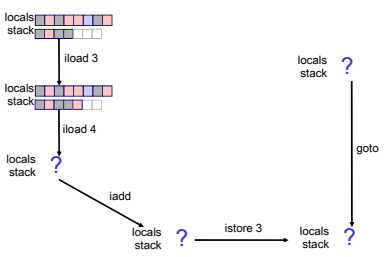

Example

31

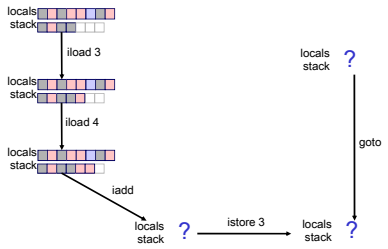

Example

32

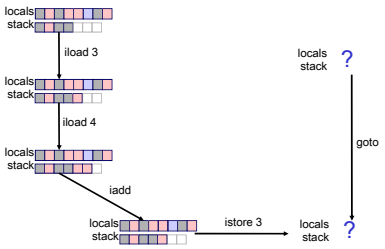

Example

33

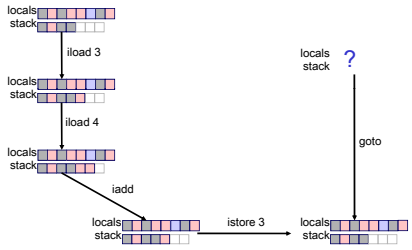

Example

34

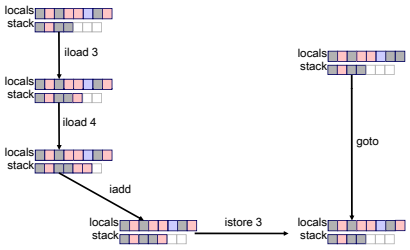

Example

35

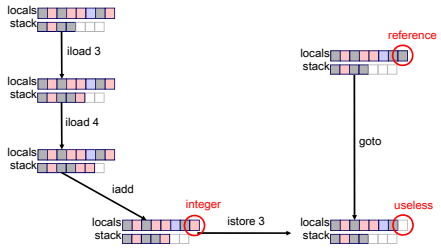

Example

36

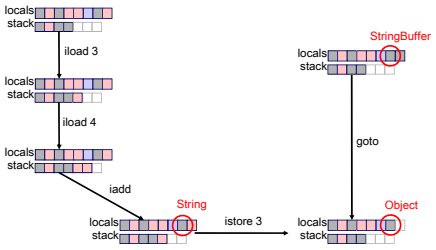

Mobile Code

37

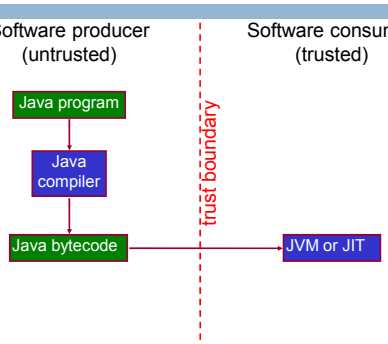

Mobile Code

38

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?*

Mobile Code

39

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

Mobile Code

40

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

Runtime Stack

41

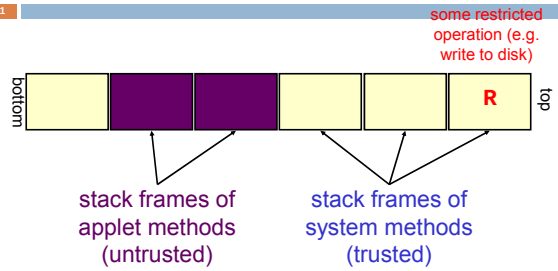

Runtime Stack

42

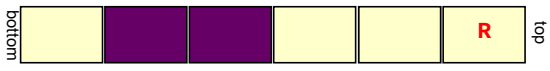

Maybe we want to disallow it
–the malicious applet may be trying to erase our disk
–it's calling system code to do that

Runtime Stack

43

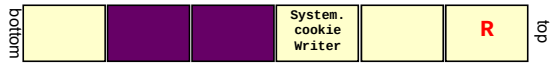

Or, maybe we want to allow it
- –it may just want to write a cookie
- –it called **`System.cookieWriter`**
- –**`System.cookieWriter`** knows it's ok

Runtime Stack

44

Maybe we want to allow it for another reason
- –all running methods are trusted

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

A) *Stack inspection!*

45

Stack Inspection

46

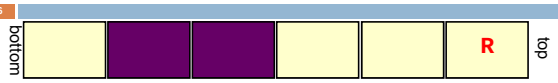

- An invocation of a trusted method, when calling another method, may either:
  - *permit* R on the stack above it
  - *forbid* R on the stack above it
  - *pass* permission from below (be transparent)
- An instantiation of an untrusted method must *forbid* R above it

Stack Inspection

47

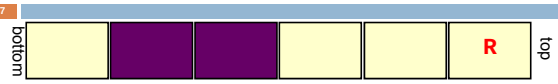

- When about to execute R, look down through the stack until we see either
  - a system method permitting R -- do it
  - a system method forbidding R -- don't do it
  - an untrusted method -- don't do it
- If we get all the way to the bottom, do it (IE, Sun JDK) or don't do it (Netscape)

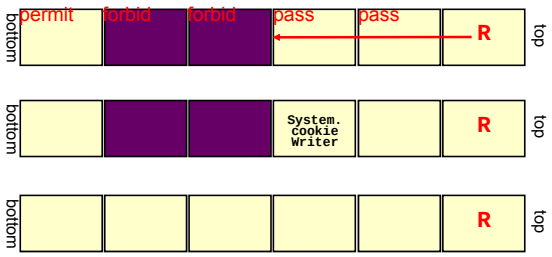

Case A: R is not executed

48

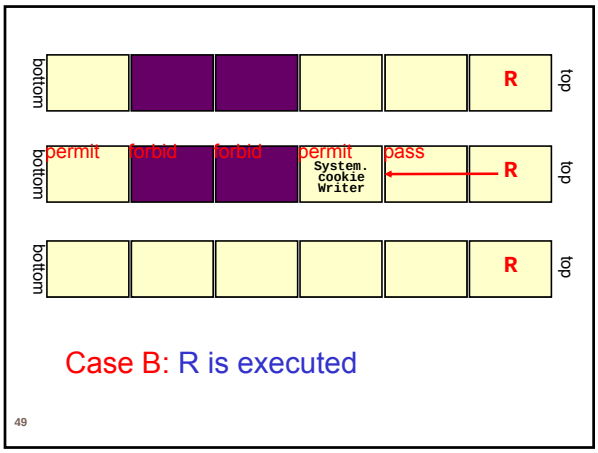

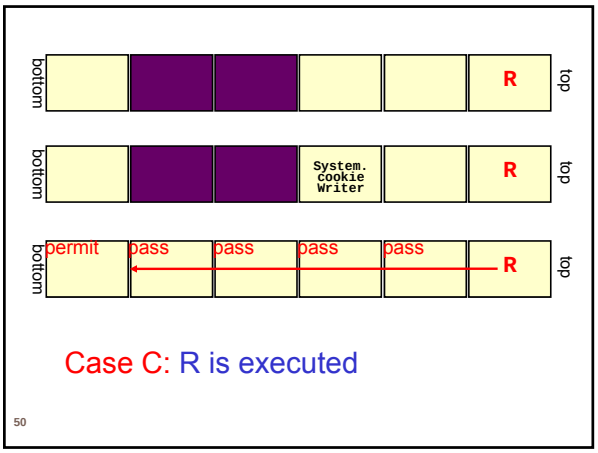

Conclusion

51

Java and the Java Virtual Machine:
Full of interesting ideas

Many systems have been built by taking an open source JVM and then somehow "doing surgery" on it. You can too!