

C and C++ 8. JNI and STL

Alan Mycroft

University of Cambridge
(heavily based on previous years' notes – thanks to Alastair Beresford and Andrew Moore)

Michaelmas Term 2012–2013

1 / 40

JNI — Java Native Interface

Java Native Interface (JNI) is the Java interface to non-Java code.

- ▶ interoperate with applications and libraries written in other programming languages (e.g., C, C++, assembly)

Examples we will look at:

- ▶ Embedding C in Java
- ▶ Using Java from C

Why use JNI?

- ▶ Low-level specifics not provided by Java
- ▶ Performance (java is not famous for being fast)
- ▶ Legacy code wrapper

2 / 40

Justification

Pro:

- ▶ Reuse: allow access to useful native code
- ▶ Efficiency: use best language for the right task

Cons:

- ▶ Applets: mobility makes this painful
- ▶ Portability: native methods aren't portable
- ▶ Extra work: javah, creating shared library

3 / 40

Embedding C in Java

1. Declare the method using the keyword native (without implementation)
2. (Make sure Java loads the required library)
3. Run javah to generate names & headers
4. Implement the method in C
5. Compile as a shared library

```
1 class HelloWorld
2 {
3     public native void displayHelloWorld();
4     static
5     {
6         System.loadLibrary("hello");
7     }
8     public static void main(String[] args)
9     {
10         new HelloWorld().displayHelloWorld();
11     }
12 }
```

4 / 40

Generate JNI Header

Compile HelloWorld.java

```
$ javac HelloWorld.java
```

Create HelloWorld.h

```
$ javah HelloWorld
```

5 / 40

HelloWorld.h

```
1 #include <jni.h>
2 /* Header for class HelloWorld */
3 #ifndef _Included_HelloWorld
4 #define _Included_HelloWorld
5 #ifdef __cplusplus
6 extern "C" {
7 #endif
8 /*
9  * Class:      HelloWorld
10  * Method:    displayHelloWorld
11  * Signature: ()V
12  */
13 JNIEXPORT void JNICALL Java_HelloWorld_displayHelloWorld
14 (JNIEnv *, jobject);
15
16 #ifdef __cplusplus
17 }
18 #endif
19 #endif
```

6 / 40

HelloWorldImp.c

```
1 #include <jni.h>
2 #include "HelloWorld.h"
3 #include <stdio.h>
4
5 JNIEXPORT void JNICALL
6 Java_HelloWorld_displayHelloWorld(JNIEnv *env, jobject obj)
7 {
8     printf("Hello world!\n");
9     return;
10 }
```

7 / 40

Create a Shared Library

```
1 class HelloWorld
2 {
3     . . .
4         System.loadLibrary("hello");
5     . . .
6 }
```

Compile the native code into a shared library:

(on MCS machines, other OS may require the -I option)

```
cc -shared HelloWorldImpl.c -o libhello.so
```

8 / 40

Run the Program

```
$ java HelloWorld
```

Sometimes:

```
Hello World!
```

Other times:

```
Exception in thread "main" java.lang.UnsatisfiedLinkError: no hello
  at java.lang.ClassLoader.loadLibrary(ClassLoader.java:1698)
  at java.lang.Runtime.loadLibrary0(Runtime.java:840)
  at java.lang.System.loadLibrary(System.java:1047)
  at HelloWorld.<clinit>(HelloWorld.java:6)
Could not find the main class: HelloWorld. Program will exit.
```

Problems? set your LD_LIBRARY_PATH:
export LD_LIBRARY_PATH=.

9 / 40

Primitive Types and Native Equivalents

Some examples:

```
typedef unsigned char  jboolean; // 8, unsigned
typedef unsigned short jchar; // 16, unsigned
typedef short          jshort; // 16
typedef float          jfloat; // 32
typedef double         jdouble; // 64
```

Each Java language element must have a corresponding native equivalent

- ▶ Platform-specific implementation
- ▶ Generic programming interface

10 / 40

The JNIEnv *env argument

- ▶ passed to every native method as the first argument.
- ▶ contains function entries used by native code.
- ▶ organised like a C++ virtual function table.

In C: (*env)->(env,foo);

In C++: env->(foo);

Through JNI interface pointer, native code can:

- ▶ operate on classes (loading, checking, etc.)
- ▶ catch and throw exceptions
- ▶ access fields
- ▶ call methods
- ▶ manipulate strings and arrays
- ▶ enter/exit monitors

11 / 40

Accessing Java Strings

The jstring type is not equivalent to the C string type

```
1 /* Wrong way */
2 JNIEXPORT jstring JNICALL Java_Prompt_getLine(JNIEnv *env, \
3 jobject obj, jstring prompt)
4 {
5     printf("%s", prompt); ...
6 }
7 /* Right way */
8 JNIEXPORT jstring JNICALL Java_Prompt_getLine(JNIEnv *env, \
9 jobject obj, jstring prompt)
10 {
11     char buf[128];
12     const char *str = (*env)->GetStringUTFChars(env, prompt, 0);
13     printf("%s", str);
14     /* release the memory allocated for the string operation */
15     (*env)->ReleaseStringUTFChars(env, prompt, str); ...
16 }
```

12 / 40

Accessing Java Array

```
1 /* Wrong way */
2 JNIEXPORT jint JNICALL Java_IntArray_sumArray(JNIEnv *env, \
3 jobject obj, jintArray arr) {
4     int i, sum = 0;
5     for (i=0; i<10; i++) {
6         sum += arr[i];
7     } ...
8 /* Right way */
9 JNIEXPORT jint JNICALL Java_IntArray_sumArray(JNIEnv *env, \
10 jobject obj, jintArray arr) {
11     int i, sum = 0;
12     /* 1. obtain the length of the array */
13     jsize len = (*env)->GetArrayLength(env, arr);
14     /* 2. obtain a pointer to the elements of the array */
15     jint *body = (*env)->GetIntArrayElements(env, arr, 0);
16     /* 3. operate on each individual elements */
17     for (i=0; i<len; i++) { sum += body[i]; }
18     /* 4. release the memory allocated for array */
19     (*env)->ReleaseIntArrayElements(env, arr, body, 0);
```

13 / 40

Accessing Java Member Variables

```
1 class FieldAccess {
2     static int si; /* signature is "si" */
3     String s; /* signature is "Ljava/lang/String;";
4 } /* run javap -s -p FieldAccess to get the signature */
5
6 /* 1. get the field ID */
7 fid = (*env)->GetStaticFieldID(env, cls, "si", "I");
8 /* 2. find the field variable */
9 si = (*env)->GetStaticIntField(env, cls, fid);
10 /* 3. perform operation on the primitive*/
11 (*env)->SetStaticIntField(env, cls, fid, 200);
12
13 /* 1. get the field ID */
14 fid = (*env)->GetFieldID(env, cls, "s", "Ljava/lang/String;");
15 /* 2. find the field variable */
16 jstr = (*env)->GetObjectField(env, obj, fid);
17 /* 3. perform operation on the object */
18 jstr = (*env)->NewStringUTF(env, "123");
19 (*env)->SetObjectField(env, obj, fid, jstr);
```

14 / 40

Calling a Java Method

1. Find the class of the object
Call `GetObjectClass`
2. Find the method ID of the object
Call `GetMethodID`, which performs a look-up for the Java method in a given class
3. Call the method
JNI provides an API for each method-type
e.g., `CallVoidMethod()`, etc.
You pass the object, method ID, and the actual arguments to the method (e.g., `CallVoidMethod`)

```
jclass cls = (*env)->GetObjectClass(env, obj);
jmethodID mid = (*env)->GetMethodID(env, cls, 'hello', '(I)V');
(*env)->CallVoidMethod(env, obj, mid, parm1);
```

15 / 40

Garbage Collection & Thread Issues

- ▶ Arrays and (explicit) global objects are *pinned* and must be explicitly released
- ▶ Everything else is released upon the native method returning

(JNIEnv *) is *only* valid in the current thread

- ▶ Interface pointer are not valid to pass between threads
- ▶ Local references are not valid to pass between threads
- ▶ Thread access to global variables requires locking

16 / 40

Synchronisation

- ▶ Synchronize is available as a C call
- ▶ Wait and Notify calls through JNIEnv do work and are safe to use

```
In Java:
synchronized(obj)
{ ...
  /* synchronized block */
... }

In C:
(*env)->MonitorEnter(env,obj);
/* synchronized block */
(*env)->MonitorExit(env,obj);
```

17 / 40

Embedding a JVM in C

- ▶ JDK ships JVM as a shared library
- ▶ This is handled as a special function call from C
- ▶ Call does not return
- ▶ The call provides a pointer for access from native code

18 / 40

JVMinC.c

```
1 #include <stdlib.h>
2 #include <jni.h>
3 int main(int argc, char *argv[]) {
4     JNIEnv *env;
5     JavaVM *jvm;
6     jint res;
7     jclass cls;
8     jmethodID mid;
9     jstring jstr;
10    jobjectArray args;
11    JavaVMInitArgs vm_args;
12    JavaVMOption options[4];
13    /* disable JIT */
14    options[0].optionString = "-Djava.compiler=NONE";
15    /* set user class and native library path */
16    options[1].optionString = "-Djava.class.path=";
17    options[2].optionString = "-Djava.library.path=";
18    /* print JNI-related messages */
19    options[3].optionString = "-verbose:jni";
```

19 / 40

JVMinC.c (cont.)

```
1     vm_args.version = JNI_VERSION_1_4;
2     vm_args.options = options;
3     vm_args.nOptions = 4;
4     vm_args.ignoreUnrecognized = 1;
5     /* spawn JVM, find class and the class entry-point */
6     res = JNI_CreateJavaVM(&jvm, (void **)&env, &vm_args);
7     cls = (*env)->FindClass(env, "JVMinCTest");
8     mid = (*env)->GetStaticMethodID(env, cls, "main", \
9         "(Ljava/lang/String;)V");
10    /* create a valid string to pass */
11    jstr = (*env)->NewStringUTF(env, " from C!");
12    args = (*env)->NewObjectArray(env, 1, \
13        (*env)->FindClass(env, "java/lang/String"), jstr);
14    /* Run and cleanup */
15    (*env)->CallStaticVoidMethod(env, cls, mid, args);
16    (*jvm)->DestroyJavaVM(jvm);
17    return 0;
18 }
```

20 / 40

Compiling Embedded JVM and C

```
1 public class JVMinCTest
2 {
3     public static void main(String [] args)
4     {
5         System.out.println("Hello from JVMinCTest");
6         System.out.println("String passed from C: " + args[0]);
7     }
8 }
```

To compile

```
$ javac JVMinCTest
$ cc JVMinC.c -o JVMinC -L
/usr/lib/jvm/java-1.6.0-openjdk-1.6.0/jre/lib/i386/client/
-ljvm
```

Running:

```
$ ./JVMinCTest Hello from JVMinCTest
String passed from C: from C!
$
```

21 / 40

Standard Template Library (STL)

Alexander Stepanov, designer of the Standard Template Library says:

"STL was designed with four fundamental ideas in mind:

- ▶ Abstractness
- ▶ Efficiency
- ▶ Von Neumann computational model
- ▶ Value semantics"

It's an example of generic programming; in other words reusable or "widely adaptable, but still efficient" code

22 / 40

Additional references

- ▶ Musser et al (2001). STL Tutorial and Reference Guide (Second Edition). Addison-Wesley.
- ▶ <http://gcc.gnu.org/onlinedocs/libstdc++/documentation.html>

23 / 40

Advantages of generic programming

- ▶ Traditional container libraries place algorithms as member functions of classes
 - ▶ Consider, for example, "test".substring(1,2); in Java
- ▶ So if you have m container types and n algorithms, that's nm pieces of code to write, test and document
- ▶ Also, a programmer may have to copy values between container types to execute an algorithm
- ▶ The STL does not make algorithms member functions of classes, but uses meta programming to allow programmers to link containers and algorithms in a more flexible way
- ▶ This means the library writer only has to produce $n + m$ pieces of code
- ▶ The STL, unsurprisingly, uses templates to do this

24 / 40

Plugging together storage and algorithms

Basic idea:

- ▶ define useful data storage components, called containers, to store a set of objects
- ▶ define a generic set of access methods, called iterators, to manipulate the values stored in containers of any type
- ▶ define a set of algorithms which use containers for storage, but only access data held in them through iterators

The time and space complexity of containers and algorithms is specified in the STL standard

25 / 40

A simple example

```
1 #include <iostream>
2 #include <vector> //vector<T> template
3 #include <numeric> //required for accumulate
4
5 int main() {
6     int i[] = {1,2,3,4,5};
7     std::vector<int> vi(&i[0],&i[5]);
8
9     std::vector<int>::iterator viter;
10
11    for(viter=vi.begin(); viter < vi.end(); ++viter)
12        std::cout << *viter << std::endl;
13
14    std::cout << accumulate(vi.begin(),vi.end(),0) << std::endl;
15 }
```

26 / 40

Containers

- ▶ The STL uses containers to store collections of objects
- ▶ Each container allows the programmer to store multiple objects of the same type
- ▶ Containers differ in a variety of ways:
 - ▶ memory efficiency
 - ▶ access time to arbitrary elements
 - ▶ arbitrary insertion cost
 - ▶ append and prepend cost
 - ▶ deletion cost
 - ▶ ...
- ▶ The STL specifies bounds on these costs for each container type.

27 / 40

Containers

- ▶ Container examples for storing sequences:
 - ▶ `vector<T>`
 - ▶ `deque<T>`
 - ▶ `list<T>`
- ▶ Container examples for storing associations:
 - ▶ `set<Key>`
 - ▶ `multiset<Key>`
 - ▶ `map<Key,T>`
 - ▶ `multimap<Key, T>`

28 / 40

Using containers

```
1 #include <string>
2 #include <map>
3 #include <iostream>
4
5 int main() {
6
7     std::map<std::string,std::pair<int,int> > born_award;
8
9     born_award["Perlis"] = std::pair<int,int>(1922,1966);
10    born_award["Wilkes"] = std::pair<int,int>(1913,1967);
11    born_award["Hamming"] = std::pair<int,int>(1915,1968);
12    //Turing Award winners (from Wikipedia)
13
14    std::cout << born_award["Wilkes"].first << std::endl;
15
16    return 0;
17 }
```

29 / 40

std::string

- ▶ Built-in arrays and the `std::string` hold elements and can be considered as containers in most cases
- ▶ You can't call `.begin()` on an array however!
- ▶ Strings are designed to interact well with C char arrays
- ▶ String assignments, like containers, have value semantics:

```
1 #include <iostream>
2 #include <string>
3
4 int main() {
5     char s[] = "A string ";
6     std::string str1 = s, str2 = str1;
7
8     str1[0]='a', str2[0]='B';
9     std::cout << s << str1 << str2 << std::endl;
10    return 0;
11 }
```

30 / 40

Iterators

- ▶ Containers support iterators, which allow access to values stored in a container
- ▶ Iterators have similar semantics to pointers
 - ▶ A compiler may represent an iterator as a pointer at run-time
- ▶ There are a number of different types of iterator
- ▶ Each container supports a subset of possible iterator operations
- ▶ Containers have a concept of a `beginning` and `end`

31 / 40

Iterator types

Iterator type	Supported operators
Input	<code>== != ++ *(read only)</code>
Output	<code>== != ++ *(write only)</code>
Forward	<code>== != ++ *</code>
Bidirectional	<code>== != ++ * --</code>
Random Access	<code>== != ++ * -- + - += -= < > <= >=</code>

- ▶ Notice that, with the exception of input and output iterators, the relationship is hierarchical
- ▶ Whilst iterators are organised logically in a hierarchy, they do not do so formally through inheritance!
- ▶ There are also `const` iterators which prohibit writing to ref'd objects

32 / 40

Adaptors

- ▶ An adaptor modifies the interface of another component
- ▶ For example the `reverse_iterator` modifies the behaviour of an iterator

```
1 #include <vector>
2 #include <iostream>
3
4 int main() {
5     int i[] = {1,3,2,2,3,5};
6     std::vector<int> v(&i[0],&i[6]);
7
8     for (std::vector<int>::reverse_iterator i = v.rbegin();
9          i != v.rend(); ++i)
10         std::cout << *i << std::endl;
11
12     return 0;
13 }
```

33 / 40

Generic algorithms

- ▶ Generic algorithms make use of iterators to access data in a container
- ▶ This means an algorithm need only be written once, yet it can function on containers of many different types
- ▶ When implementing an algorithm, the library writer tries to use the most restrictive form of iterator, where practical
- ▶ Some algorithms (e.g. `sort`) cannot be written efficiently using anything other than random access iterators
- ▶ Other algorithms (e.g. `find`) can be written efficiently using only input iterators
- ▶ Lesson: use common sense when deciding what types of iterator to support
- ▶ Lesson: if a container type doesn't support the algorithm you want, you are probably using the wrong container type!

34 / 40

Algorithm example

- ▶ Algorithms usually take a `start` and `finish` iterator and assume the valid range is `start` to `finish-1`; if this isn't true the result is undefined

Here is an example routine `search` to find the first element of a storage container which contains the value `element`:

```
1 //search: similar to std::find
2 template<class I, class T> I search(I start, I finish, T element)
3     while (*start != element && start != finish)
4         ++start;
5     return start;
6 }
```

35 / 40

Searching over multiple containers

```
1 #include "example23.hh"
2
3 #include "example23a.cc"
4
5 int main() {
6     char s[] = "The quick brown fox jumps over the lazy dog";
7     std::cout << search(&s[0], &s[strlen(s)], 'd') << std::endl;
8
9     int i[] = {1,2,3,4,5};
10    std::vector<int> v(&i[0], &i[5]);
11    std::cout << search(v.begin(), v.end(), 3) - v.begin()
12                << std::endl;
13
14    std::list<int> l(&i[0], &i[5]);
15    std::cout << (search(l.begin(), l.end(), 4) != l.end())
16                << std::endl;
17
18    return 0;
19 }
```

36 / 40

Heterogeneity of iterators

```
1 #include "example24.hh"
2
3 int main() {
4     char one[] = {1,2,3,4,5};
5     int two[] = {0,2,4,6,8};
6     std::list<int> l (&two[0], &two[5]);
7     std::deque<long> d(10);
8
9     std::merge(&one[0], &one[5], l.begin(), l.end(), d.begin());
10
11    for (std::deque<long>::iterator i = d.begin(); i != d.end(); ++i)
12        std::cout << *i << " ";
13    std::cout << std::endl;
14
15    return 0;
16 }
```

We use `merge` to merge different element types from different containers.

37 / 40

Function objects

- ▶ C++ allows the function call "`()`" to be overloaded
- ▶ This is useful if we want to pass functions as parameters in the STL
- ▶ More flexible than function pointers, since we can store per-instance object state inside the function
- ▶ Example:

```
1 struct binaccum {
2     int operator()(int x, int y) const {return 2*x + y;}
3 };
```

38 / 40

Higher-order functions in C++

- ▶ In ML we can write: `foldl (fn (y,x) => 2*x+y) 0 [1,1,0]`;
- ▶ Or in Python: `reduce(lambda x,y: 2*x+y, [1,1,0])`
- ▶ Or in C++:

```
1 #include<iostream>
2 #include<numeric>
3 #include<vector>
4
5 #include "example27a.cc"
6
7 int main() { //equivalent to foldl
8
9     bool binary[] = {true,true,false};
10    std::cout << std::accumulate(&binary[0], &binary[3], 0, binaccum())
11                << std::endl; //output: 6
12
13    return 0;
14 }
```

39 / 40

Higher-order functions in C++

- ▶ By using reverse iterators, we can also get `foldr`:

```
1 #include<iostream>
2 #include<numeric>
3 #include<vector>
4
5 #include "example27a.cc"
6
7 int main() { //equivalent to foldr
8
9     bool binary[] = {true,true,false};
10    std::vector<bool> v(&binary[0], &binary[3]);
11
12    std::cout << std::accumulate(v.rbegin(), v.rend(), 0, binaccum());
13    std::cout << std::endl; //output: 3
14
15    return 0;
16 }
```

40 / 40