

JNI – C++ integration made easy

Evgeniy Gabrilovich
gabr@acm.org

Lev Finkelstein
lev@zapper.com

Abstract

The Java Native Interface (JNI) [1] provides interoperation between Java code running on a Java Virtual Machine and code written in other programming languages (e.g., C++ or assembly). The JNI is useful when existing libraries need to be integrated into Java code, or when portions of the code are implemented in other languages for improved performance. The Java Native Interface is extremely flexible, allowing Java methods to invoke native methods and vice versa, as well as allowing native functions to manipulate Java objects. However, this flexibility comes at the expense of extra effort for the native language programmer, who has to explicitly specify how to connect to various Java objects (and later to disconnect from them, to avoid resource leak). We suggest a template-based framework that relieves the C++ programmer from most of this burden. In particular, the proposed technique provides automatic selection of the right functions to access Java objects based on their types, automatic release of previously acquired resources when they are no longer necessary, and overall simpler interface through grouping of auxiliary functions.

Introduction

The Java Native Interface¹ is a powerful framework for seamless integration between Java and other programming languages (called “native languages” in the JNI terminology). A common case of using the JNI is when a system architect wants to benefit from both worlds, implementing communication protocols in Java and computationally expensive algorithmic parts in C++ (the latter are usually compiled into a dynamic library, which is then invoked from the Java code). The JNI renders native applications with much of the functionality of Java, allowing them to call Java methods, access and modify Java variables, manipulate Java exceptions, ensure thread-safety through Java thread synchronization mechanisms, and ultimately to directly invoke the Java Virtual Machine.

This functional wealth is provided through a rather complex interface between the native code and the Java programming environment. For example, accessing an integer Java member variable from the native code is a multi-step process, which involves first querying the class type whose variable is to be retrieved (using the JNI function `GetObjectClass`), obtaining the field identifier (using `GetFieldID`), and finally retrieving the field per se (using `GetIntField`). The latter function is a representative of a set of functions (`Get<type>Field`), each corresponding to a different variable type; thus, accessing a variable requires explicit specification of the appropriate function. To streamline these steps, we suggest using a template (parameterized with a variable type), which encapsulates all these low-level operations, giving the programmer easy-to-use assignment/modification functionality similar to that of native C++ variables. As far as

¹ See the sidebar.

data types are concerned, we develop a complete set of templates for primitive types (jint, jfloat etc.), as well as outline a generic approach for accessing any user-defined type.

Observe also that there are a number of JNI API functions, which are complementary in their nature. For example, if a native function has a parameter of type jstring, it should first convert this parameter into a conventional C++ string of chars using GetStringChars (or GetStringUTFChars), and subsequently explicitly release this temporary representation using the corresponding function ReleaseStringChars (or ReleaseStringUTFChars). This mode of operation could be greatly simplified by implementing a proxy-style smart pointer, which realizes the “construction as resource acquisition” idiom [2]. This way, a Java string is transparently converted on access into a C++ string (char*), and when it later goes out of scope, the smart pointer destructor releases any temporary resources used. The proxy also provides proper handling of the transfer of ownership, similarly to the C++ standard auto_ptr construct [3]. We also suggest analogous treatment for Java arrays, featuring automatic selection of the access function based on the actual element type (e.g., Get<Int>ArrayElements), releasing array elements in the destructor of this “smart container”, and conventional access to array elements with operator[]. The “construction as resource acquisition” idiom may also be applied to function pairs such as NewGlobalRef / DeleteGlobalRef used for reservation and release of global references (respectively), MonitorEnter / MonitorExit used for protecting critical sections of code etc.

To make the discussion concrete, we start with developing a running example which solves a toy problem. We first show the native code for solving this task without using the framework (“before”), and then show the desired code (“after”), streamlined using the proposed framework. In a subsequent section we develop the JNI encapsulation framework step by step. The article ends with a larger-scale example.

Running example

Our toy problem defines a Java class JniExample with an integer, a static String and an integer array fields. Function main calls a native function implemented in C++, that accesses the Java fields, prints their original (default) values, and ultimately modifies them. Listing 1 shows (a fragment of) the Java class whose variables need to be accessed from C++ code.

Listing 1. A Java class for the running example (excerpt from JniExample.java):

```
public class JniExample {
    public int intField = 17;                // integer field
    public static String stringField = "Hello, world!"; // static String field
    public int[] intArray = new int[2];    // integer array
    ...
    private static native void native_call(JniExample x); // sample native call
    ...
}
```

Listing 2 shows sample implementation of native C++ code that modifies the above variables, implemented using the original JNI [1]. Observe that on average 3 to 4 preparatory operations are necessary to access a Java field from the native code. The

revised code in Listing 3 reduces this overhead to one constructor invocation per access. Note also that assignment to Java fields becomes more intuitive too.

Listing 2. Sample native code that uses the original JNI (excerpt from `jni_example_org.cpp`):

```
JNIEXPORT void JNICALL Java_JniExample_org_1native_1call (JNIEnv *env, jclass clazz, jobject obj) {
    // Lookup the integer field ('intField') in 'obj' and get its value
    jfieldID intFieldId = env->GetFieldID(clazz, "intField", "I");
    if (id == 0) throw JNIException("Field not found");
    int intFieldVal = env->GetIntField(obj, intFieldId);

    // Lookup the static String field ('stringField') in 'obj', then convert the Java string representation
    // to C++ type 'const char *'.
    jfieldID strFieldId = env->GetStaticFieldID(clazz, "stringField", "Ljava/lang/String;");
    if (id == 0) throw JNIException("Field not found");
    jstring jstr = (jstring) env->GetStaticObjectField(clazz, strFieldId);
    const char *str = (jstr == 0) ? 0 : env->GetStringUTFChars(jstr, 0);

    // Lookup the integer array field ('intArray') in 'obj', then convert it to C++ type 'jint *'.
    jfieldID arrFieldId = env->GetFieldID(clazz, "intArray", "[I");
    if (id == 0) throw JNIException("Field not found");
    jintArray jarr = (jintArray) env->GetObjectField(obj, arrFieldId);
    jint *arr = env->GetIntArrayElements(jarr, 0);

    // Set new values
    env->SetIntField(obj, intFieldId, 0);
    arr[0] = 0; arr[1] = 0;
    env->SetStaticObjectField(clazz, strFieldId, env->NewStringUTF("Good-bye, world!"));

    // Explicitly release resources
    env->ReleaseIntArrayElements(jarr, arr, 0);
    env->ReleaseStringUTFChars(jstr, str);
}
```

Listing 3. Sample native code that uses the JNI encapsulation framework (excerpt from `jni_example.cpp`):

```
JNIEXPORT void JNICALL Java_JniExample_native_1call (JNIEnv *env, jclass clazz, jobject obj) {
    // Lookup the Java fields in 'obj'
    JNIField<jint> intField(env, obj, "intField");
    JNIStringUTFChars str(env, "JniExample", "stringField");
    JNIArray<jint> arr(env, obj, "intArray");

    // Set new values
    intField = 0;
    arr[0] = 0; arr[1] = 0;
    JNIStaticField<jstring>(env, obj, "stringField") = env->NewStringUTF("Good-bye, world!");

    // Destructors for 'arr' and 'str' are invoked automatically
}
```

Template-based encapsulation of the JNI

In this section we evolve the JNI encapsulation framework. First, we discuss access to scalar variables (of both primitive and user-defined types). We then develop a generic

resource management scheme which underlies the implementation of containers (arrays and strings). Finally, we apply this scheme to advanced JNI features such as monitors and global references.

Field access

Accessing a Java variable from C++ is a cumbersome process at the least. Our aim is to develop a technique for establishing correspondence between a C++ object and the Java variable, so that all low-level access operations become transparent to the programmer. For example, a C++ object corresponding to `intField` is a proxy² of type `JNIField<jint>`, created using the environment handle and the Java object passed via the JNI native call: `JNIField<jint> intField(env, obj, "intField")`. Then changing the variable value in C++ is as simple as that: `intField = 17`.

To this end, we use a template class `JNIField` (see Listing 4), whose three constructors present three ways to attach a C++ variable to a Java field. The first constructor receives an environment handle, a reference to the Java object whose field is to be accessed, the field name and type signature, and connects to the designated field. The second constructor allows creation of fields whose type signature may be deduced automatically (more on this below). Both constructors compute the field identifier (using implicit calls to the JNI functions `GetObjectClass` and `GetFieldId`), and store it together with the object handle. Template class `JNIFieldId`, which encapsulates the notion of a field identifier, is also defined in file `jni_field.h`. Occasionally, it might be necessary to access the same field in numerous Java objects of the same type (e.g., iterating over array elements). In such a case, the field identifier may only be computed once, and cached for subsequent reuse; this mode is supported by the third constructor. The assignment and casting operators provide easy access to the variable value.

Listing 4. `JNIField` template (excerpt from `jni_field.h`):

```
template<class NativeType> class JNIField {
    typedef JNIField<NativeType> _self;
    jobject _obj;                // Object that holds the field
    JNIFieldId<NativeType> _id;   // field ID

public:
    JNIField(JNIEnv *env, jobject obj, const char *name, const char *sig) :
        _obj(obj), _id(env, obj, name, sig) {}
    JNIField(JNIEnv *env, jobject obj, const char *name) :
        _obj(obj), _id(env, obj, name) {}
    JNIField(JNIFieldId<NativeType> id, jobject obj) : _obj(obj), _id(id) {}

    _self &operator= (const _self &rhs);
    _self &operator= (const NativeType &rhs);

    // Casting to NativeType
    operator NativeType() const { return _id.Get(_obj); }
};
```

² The name of the C++ variable is obviously arbitrary; we use the same variable name in C++ and Java merely for convenience.

An interesting aspect of the field-handling templates (JNIField, JNIFieldId) is their type inference capability, which is mostly hidden from the user. Observe that the JNI only supplies a set of Get<type>Field functions, each corresponding to a different variable type; thus, accessing a variable apparently requires explicit specification of the appropriate function. We circumvent this requirement by using the *template specialization* technique, namely, we preinstantiate the member functions of JNIFieldId for all the primitive types (i.e., JNIFieldId<jint>::operator=() is actually implemented using the SetIntField function etc.). This way, once the field is instantiated (for example, on an integer Java variable – JNIField<jint>), we no longer need to specify its type on every variable access.

This mechanism uses a series of basic declarations, where each primitive C++ type is associated with the corresponding Java basic type and Java array type, and is assigned the JNI type signature. The declarations are realized as C++ structs, which are later used by the compiler as a lookup table (among the rest, this relieves the user from having to remember signatures of various Java types; instead, those can be looked up whenever necessary). For more details, see the full implementation in files jni_field.h and jni_declarations.h.

To conclude the presentation of field access, file jni_field.h provides similar definitions for static fields (“class variables” in Java terminology). It defines a template class JNIStaticFieldId (also specialized for the primitive types), and a template class JNIStaticField, which has an additional data member of type jclass for keeping the appropriate class object.

Arrays and user-defined data type

The framework provides (in file jni_declarations.h) a set of declarations to facilitate JNI arrays. We actually build a lookup table, which the compiler consults for type inference. This table may be looked up either given a primitive type or an array type. For example, this way the compiler can automatically deduce that an array of jint is of type jintArray, and that an array of type jcharArray consists of jchars and has the signature “[C”. Listing 3 contains an example of how to connect from C++ code to a Java array of integers (“intArray”).

The framework can also be easily extended for custom (non-primitive) types, so that they can be used by the compiler for automatic type inference. File jni_declarations.h exemplifies this feature with the declarations for String data type (struct StringDeclarations). Such declarations are immediately available for the compiler to utilize. For example, to assign a new value to the string field of the running example (see Listing 1) from the C++ code, we use the following definition:

```
JNIStaticField<jstring>(env, obj, "stringField") = env->NewStringUTF("Good-bye, world!");
```

Specifically observe that we do not designate explicitly the corresponding type signature, as the compiler infers it automatically.

Resource management

The JNI encapsulation framework features a generic resource management mechanism, which (when adapted properly) greatly simplifies various JNI use cases. Our resource

management approach is based on the C++ "construction as resource acquisition" idiom [2]: resources are allocated in the constructor, and are released in the destructor of dedicated auxiliary objects. This idiom is implemented using the Proxy pattern [4], with functionality similar to that of `auto_ptr` template of the C++ Standard Library [3]. This way, Java resource acquisition is performed on construction of the corresponding C++ object, transparently for the programmer. Resource deallocation (release) is also performed automatically, through the C++ object destructor. Customized resource management (e.g., the `isCopy` parameter for allocation and the `mode` parameter for deallocation) is also supported.

Strings

The JNI allows two kinds of strings, namely, using regular (UTF-8) or wide (Unicode) characters. Class `JNIStringUTFChars` provides an interface to Java UTF-8 strings of type `jstring`. Such a string can be accessed as raw `const char*`, or converted to C++ `std::string` for convenience³. Since Java strings cannot be modified, class `JNIStringUTFChars` only provides the `const` version of `operator[]`. Function `length` returns the length of the C++ string (using the JNI function `GetStringUTFLength`). File `jni_resource.h`, which contains the complete implementation, also defines class `JNIStringChars` for accessing Java strings of wide (Unicode) characters.

As an example, consider the following code that attaches a C++ variable to the static string field of the running example (see Listing 1), using the names of the host class and the string field: `JNIStringUTFChars str(env, "JniExample", "stringField")`. The string value can then be printed simply using `cout << str.get()`.

To replace the value of a Java string, we should first create a new string that can later "survive" in the Java environment. In the code fragment below, a C++ variable is instantiated on the static string field of the running example (see Listing 1), and then assigned a brand new Java string created with the JNI function `NewStringUTF`:

```
JNIStaticField<jstring>(env, obj, "stringField") = env->NewStringUTF("Good-bye, world!");
```

Note that there is no need to worry about eventually releasing the memory occupied by this newly created string – this is performed by Java garbage collector.

Arrays

Arrays feature most of the functionality presented till now. In particular, they provide automatic acquisition and release of elements in the constructor and destructor respectively (the template specialization trick is used here again to preinstantiate the array template for all primitive types, so that the appropriate `Get<Type>ArrayElements / Release<Type>ArrayElements` functions are automatically selected based on the context). Function `size` uses the JNI facility `GetArrayLength` to determine the number of array elements. Two versions of `operator[]` (regular and `const`) are provided to access the individual elements. File `jni_resource.h` contains the implementation.

³ Note that conversion function `asString` physically copies the characters.

To access the integer array field of the running example (see Listing 1), we instantiate a corresponding C++ variable as follows: `JNIArray<jint> arr(env, obj, "intArray")`. Subsequent access of the array elements is straightforward: `arr[0] = 0`.

The default behavior of the `JNIArray` is to copy all the array elements back into Java environment once the C++ array goes out of scope (this is done by function `Release<Type>ArrayElements` invoked from the array destructor). When this behavior needs to be overridden, use member function `CustomRelease` to set the desired mode for `Release<Type>ArrayElements`.

Occasionally, it is not necessary to manipulate an entire Java array, which may be quite large. For cases when only a part of the array needs to be accessed, the JNI provides a pair of functions `Get<Type>ArrayRegion` / `Set<Type>ArrayRegion`. File `jni_utils.h` defines template functions `GetArrayRegion` / `SetArrayRegion` (preinstantiated at compile time for primitive types) that are capable of deducing the element type based on their parameters.

Monitors

Monitors serve to ensure mutual exclusion of threads competing for a shared resource. To ensure resource integrity (“thread safety”), threads should request to enter a monitor at the beginning of the critical section, and leave it at the end of the section. We propose a resource management technique that uses an auxiliary automatic object of type `JNIMonitor`, whose constructor enters a monitor, and whose destructor leaves the monitor as soon as the object goes out of scope (see file `jni_resource.h` for implementation details). The constructor of `JNIMonitor` receives a handle to the object that constitutes a shared resource protected by this monitor. Listing 5 gives a sample code fragment that uses a monitor.

Listing 5. Sample code fragment that uses a monitor.

```
void sample_function(JNIEnv *env, jobject obj) {
...
  { // start the critical section block
    JNIMonitor mon(env, obj);
    // do the critical section stuff here
    ...
    // the destructor of 'mon' automatically exits
    // the underlying Java monitor
  }
...
}
```

Global references

It is sometimes necessary to obtain a reference to a Java object, so that it can be used across the function boundaries of C++ code. In such cases, a *global reference* to this object should be reserved (in contrast to most JNI functions that yield a *local* reference, which expires as soon as the current scope terminates). The last use case defined in file `jni_resource.h` implements an auxiliary template class `JNIGlobalRef`, whose constructor acquires and destructor releases a global reference to the specified object (see sample usage in Listing 6).

Using the code

Following the STL methodology, all the framework code resides in header files and is entirely template based, so clients do not need to compile their applications with any additional libraries. In fact, client applications only need to include the master file `jni_master.h`, which itself `#includes` all the other headers. The entire code of this article with complete Java – C++ integration examples can be obtained from the C/C++ Users Journal Web site at <http://www.cuj.com/code>.

A more elaborate example

We now apply the JNI encapsulation framework to a more substantial example. Suppose we have a Java application in which several concurrent threads generate (a predefined number of) objects with string IDs. The task is to collect these objects in a thread-safe way, and ultimately sort them by their IDs. Since earlier versions of the Java Development Kit (such as JDK 1.1) did not have the Collections Framework, it is quite natural to implement the sorting container in native C++ code using the STL. File `JniComplexExample.java` (not shown here for the sake of brevity) contains the Java part of this example, which uses the following native functions:

- `init_native_resources()` – initializes the native code data structures
- `clean_native_resources()` – releases the native resources
- `register_object()` – inserts a given object into the container
- `recall_objects()` – returns a sorted array of the collected objects

Listing 6 shows the native code for the container, implemented using the STL `multimap`. The container holds the (Java originated) objects by global references⁴, associated with their string IDs. The container is realized as a singleton object, which has two thread-safe access functions:

- `insert()` – collects a given object
- `exportAllObjects()` – returns all the collected objects as a vector (observe that this vector is inherently sorted, as the objects are extracted from a `multimap`).

To ensure code portability, thread-safety is implemented using JNI monitors. Critical sections start with monitor definition and last until the monitor is automatically destroyed as it goes out of scope.

Listing 6. Native code implementation of a thread-safe sorting container (excerpt from file `jni_complex_example.cpp`):

```
class SampleContainer {
    static SampleContainer *instance;

    typedef multimap<string, JNIGlobalRef< jobject > * > MapOfObjects;
    MapOfObjects mapOfObjects;           // the container implementation
    JNIGlobalRef< jobject > monitor;      // monitor (for critical sections)
    JNIEnv *_env;                        // the environment variable

    SampleContainer(JNIEnv *env) : _env(env), monitor(env, getMonitorObject(env)) {}
    ~SampleContainer();                  // the dtor purges all the map elements
};
```

⁴ Global references are required so that the Java garbage collector does not destroy the objects prematurely.


```

jobject getMonitorObject(JNIEnv *env) { // allocating the monitor object
    jclass objectClass(env, "java/lang/Object");
    jmethodID constructorId = env->GetMethodID(objectClass, "<init>", "()V");
    return env->NewObject(objectClass, constructorId);
}

public:
static SampleContainer *getInstance(JNIEnv *env = 0) {
    if (instance == 0) {
        if (env == 0) { // raise an exception flag in Java, then throw a C++ exception
            env->ThrowNew(jclass(env, "java/lang/Exception"), "Init failed");
            throw new JNIException("Init failed");
        }
        static JNIGlobalRef<jobject> initMonitor(env, getMonitorObject(env));

        // Double-checked locking is used to provide correct initialization
        JNIMonitor startCriticalSection(env, initMonitor);
        if (instance == 0)
            instance = new SampleContainer(env);
    }
    return instance;
}

static void clean() { delete instance; } // explicitly release the native resources

void insert(jobject obj) {
    JNIMonitor startCriticalSection(_env, monitor);

    JNIStringUTFChars str(_env, obj, "name"); // retrieve the object ID
    JNIGlobalRef<jobject> *ref = new JNIGlobalRef<jobject>(_env, obj);
    mapOfObjects.insert(make_pair(str.asString(), ref));
}

vector<JNIGlobalRef<jobject> *> exportAllObjects() {
    JNIMonitor startCriticalSection(_env, monitor);

    vector<JNIGlobalRef<jobject> *> result(mapOfObjects.size(), 0);
    MapOfObjects::iterator p; vector<JNIGlobalRef<jobject> *>::iterator q;
    for (p = mapOfObjects.begin(), q = result.begin();
        p != mapOfObjects.end(); p++, q++)
        *q = (*p).second;
    return result;
}
};

// Singleton instance
SampleContainer *SampleContainer::instance = 0;

/* Implementation of native calls (note that their prototypes are automatically generated) */

// init_native_resources()
JNIEXPORT void JNICALL
Java_JniComplexExample_init_1native_1resources (JNIEnv *env, jclass clazz)
{ SampleContainer::getInstance(env); }

// clean_native_resources()

```

```

JNIEXPORT void JNICALL
Java_JniComplexExample_clean_1native_1resources (JNIEnv *env, jclass clazz)
{ SampleContainer::clean(); }

//register_object()
JNIEXPORT void JNICALL
Java_JniComplexExample_register_1object (JNIEnv *env, jclass clazz, jobject obj)
{ SampleContainer::getInstance()->insert(obj); }

//recall_objects
JNIEXPORT jobjectArray JNICALL
Java_JniComplexExample_recall_1objects (JNIEnv *env, jclass clazz) {
    // obtain the vector of global references
    vector<JNIGlobalRef<jobject> *> allObjects = SampleContainer::getInstance()->exportAllObjects();
    // create an output array of type 'NameWithInfo'
    jclass objectClass(env, "NameWithInfo");
    jobjectArray result = env->NewObjectArray(allObjects.size(), objectClass, 0);
    // export the objects
    for (int i = 0; i < allObjects.size(); i++)
        env->SetObjectArrayElement(result, i, *allObjects[i]);
    return result;
}

```

Discussion

Let us review the properties of the solution we developed:

1. *Easier to use, more straightforward approach*
Whenever possible, the compiler infers variable types from the context. Java data structures are automatically exported to (and in order to save changes, are later imported from) the C++ code. Auxiliary technical operations are encapsulated in higher-level templates.
2. *Less error-prone API*
Fewer functions to call means fewer opportunities to err in successive function invocations. Also, it is now possible to perform various checks at compile time, instead of discovering problems much later as run-time errors.
3. *Proper resource management*
Resources are automatically deallocated when they are no longer necessary, thus preventing resource leaks, deadlocks, and starvation.
4. *Portability issues*
Java portability is preserved by using only ANSI-standard C++ and the STL [3].
5. *Compilation overhead*
A possible drawback of the suggested framework is the compile-time penalty it imposes, due to heavy use of the preprocessor and embedded templates. However, this overhead is limited to the compilation time, and does not propagate to the run-time. The code size increase is also negligible, since most of the templates only provide type definitions (and thus do not need run-time representation at all), and unused template instantiations are discarded by the code optimizer.

Acknowledgments

The authors are thankful to Marc Briand, Herb Sutter and Jerry Schwarz for their constructive comments and suggestions.

References

- [1] Java Native Interface Specification
<http://java.sun.com/products/jdk/1.2/docs/guide/jni/spec/jniTOC.doc.html>
- [2] Stroustrup, B. "The C++ Programming Language", 3rd edition, Addison Wesley, 1990.
- [3] "Information Technology – Programming Languages – C++",
International Standard ISO/IEC 14882-1998(E).
- [4] Gamma, E., *et al.* "Design Patterns: Elements of Reusable Software Architecture",
Addison Wesley, 1995.

Sidebar: Java Native Interface in a nutshell

As set forth in [1], the *Java Native Interface (JNI)* enables Java code running on a Java Virtual Machine (JVM) to work in concert with code written in other programming languages (e.g., C, C++ or assembly, referred to as “native languages” due to their “nativeness” to the execution environment). The JNI can be handy to foster code reuse or to implement mission-critical parts of the code in a native language for superior efficiency. The JNI may also occasionally be useful to facilitate platform-dependent features not available through the standard Java class library. The Java Native Interface provides *bidirectional* cooperation between Java code and native code, so that Java methods may transparently invoke native routines, and vice versa. Additional functional wealth available to native applications includes manipulating Java objects, access to Java class information and run-time type checking facilities, dispatching Java exceptions, as well as convenient usage of Java thread synchronization mechanisms. Finally, the so-called *Invocation API* allows any native application to directly operate the Java Virtual Machine as a regular native object.

Every native function receives a *JNI interface pointer* through which it calls all the other JNI functions. For the sake of implementation flexibility, the interface pointer indirectly points to a table of pointers to JNI functions. Observe, therefore, that calling a native method always entails several dereference operations. Note also that the interface pointer is only valid in the current thread. This pointer is implicit in Java signatures of native methods, and constitutes the (explicit) first parameter in their native programming language. The JNI also prescribes the meaning of the second parameter to native methods. This parameter contains a reference to the host object (*this*) for *instance methods* (nonstatic functions), and a reference to the Java class object for *class methods* (static functions). Libraries of native functions are normally loaded dynamically at run-time, using the `System.loadLibrary` method. Name mangling conventions for native methods allow overloading, and are stipulated by the JNI Specification.

The JNI allows native code to access Java objects of both primitive types (e.g., `int`, `char`) and user-defined types. The JNI Specification associates each Java primitive type with an equivalent native type (for instance, the `jfloat` C++ native type corresponds to the Java `float`, and is implemented as a 32-bit variable). Native methods may receive Java objects as parameters. Retrieving data members of a compound parameter (including individual array entries), or creating new objects in the Java environment, is performed by calling appropriate JNI functions. The Java Virtual Machine keeps track of all the objects made

available to the native code, so that they do not get garbage-collected while in use. Calling Java methods and raising exceptions from the native code is also accomplished through the variety of JNI functions.

About the authors

Evgeniy Gabrilovich is an Algorithm Developer at Zapper Technologies Inc. He holds an M.Sc. degree in Computer Science from the Technion – Israel Institute of Technology. His interests involve Computational Linguistics, Information Retrieval, Artificial Intelligence, and Speech Processing. He can be contacted at gabr@acm.org.

Lev Finkelstein is an Algorithm Developer at Zapper Technologies Inc., and is a Ph.D. student in Computer Science at the Technion – Israel Institute of Technology. His interests include Artificial Intelligence, Machine Learning, Multi-agent systems, and Data Mining. He can be reached at lev@zapper.com.