

Polymorphic Type Inference for the JNI *

Michael Furr
University of Maryland, College Park
furr@cs.umd.edu

Jeffrey S. Foster
University of Maryland, College Park
jfoster@cs.umd.edu

January 19, 2006

Abstract

We present a multi-lingual type inference system for checking type safety of programs that use the Java Native Interface (JNI). The JNI uses specially-formatted strings to represent class and field names as well as method signatures, and so our type system tracks the flow of string constants through the program. Our system embeds string variables in types, and as those variables are resolved to string constants during inference they are replaced with the structured types the constants represent. This restricted form of dependent types allows us to directly assign type signatures to each of the more than 200 functions in the JNI. Moreover, it allows us to infer types for user-defined functions that are parameterized by Java type strings, which we have found to be common practice. Our inference system allows such functions to be treated polymorphically by using instantiation constraints, solved with semi-unification, at function calls. Finally, we have implemented our system and applied it to a small set of benchmarks. Although semi-unification is undecidable, we found our system to be scalable and effective in practice. We discovered 155 errors 36 cases of suspicious programming practices in our benchmarks.

1 Introduction

Foreign function interfaces (FFIs) allow programs to call functions written in other languages. FFIs are important because they let new languages access system libraries and legacy code, but using FFIs correctly is difficult, as there is usually little or no compile-time consistency checking between native and foreign code. As a result, programs that use FFIs may produce run-time typing errors or even violate type safety, thereby causing program crashes or data corruption.

In this paper we develop a multi-lingual type inference system that checks for type safe usage of Java’s FFI to C, called the Java Native Interface (JNI).¹ In the JNI, most of the work is done in C “glue code,” which translates data between the two languages and then in turn invokes other routines, often in system or user libraries. Java primitives can be accessed directly by C glue code because they have the same representations as C primitives, e.g., a Java integer can be given the C type `int`. However all Java objects, no matter what class they are from, are assigned a single opaque type `jobject` by the JNI. Since `jobject` contains no further Java type information, it is easy for a programmer to use a Java object at a wrong type without any compile-time warnings.

We present a new type inference system that embeds Java type information in C types. When programmers manipulate Java objects, they use JNI functions that take as arguments specially-formatted strings representing class names, field names, and method signatures. Our inference system tracks the values of C string constants through the code and translates them into Java type annotations on the C `jobject` type, which is a simple form of dependent types.

We have found that JNI functions are often called with parameters that are not constants—in particular, programmers often write “wrapper” functions that group together common sequences of JNI operations, and

*This research was supported in part by NSF CCF-0346982 and CCF-0430118

¹The JNI also contains support for C++, but our system only analyzes C code.

the Java types used by these functions depend on the strings passed in by callers. Thus a novel feature of our system is its ability to represent partially-specified Java classes in which class, field, and method names and type information may depend on string variables in the program. During type inference these variables are resolved to constants and replaced with the structured types they represent. This allows us to handle wrapper functions and to directly assign type signatures to the more than 200 functions in the JNI.

Our system also supports polymorphism for functions that are parameterized by string variables and Java types. Our type inference algorithm generates instantiation constraints [28] at calls to polymorphic functions, and we present an algorithm based on semi-unification [19, 10] for solving the constraints.

We have implemented our system and applied it to a small set of benchmarks. Although semi-unification is undecidable, we found our algorithm to be both scalable and effective in practice. In our experiments, we found 155 errors and 36 suspicious but non-fatal programming mistakes in our benchmarks, suggesting that programmers do misuse the JNI and that multi-lingual type inference can reveal many mistakes.

In summary, the main contributions of this paper are as follows:

- We develop a multi-lingual type inference system that embeds Java types inside of C. Our system uses a simple form of dependent types so that Java object types may depend on the values of C string constants and variables.
- We present a constraint-based type inference algorithm for inferring multi-lingual types for C glue code. During inference, as string variables are resolved to constant strings they are replaced with the structured Java types they represent. Our system uses instantiation constraints to model polymorphism for JNI functions and user-defined wrapper functions.
- We describe an implementation of our type inference system. We have applied our implementation to a small set of benchmarks, and as a result, we found a number of bugs in our benchmark suite.

This work complements and extends our previous work on the OCaml-to-C FFI [13]. Our previous system was monomorphic and worked by tracking integers and memory offsets into the OCaml heap. Our previous system also did not model objects, which clearly limits its applicability to the JNI. Our new system can model accesses to Java objects using string constants and variables, and performs parametric polymorphic type inference.

2 Background

In this section we motivate our system by briefly describing the JNI [24]. To call a C function from Java, the programmer first declares a Java method with the `native` keyword and no body. When this native method is invoked, the Java runtime finds and invokes the correspondingly-named C function. Since Java and C share the same representation for primitive types such as integers and floating point numbers, C glue code requires no special support to manipulate them. In contrast, Java objects, such as instances of `Object`, `Class`, or `Int[]`, are all represented with a single opaque C type `jobject` (often an alias of `void *`), and glue code invokes functions in the JNI to manipulate `jobjects`. For example, to get the object `Point.class`, which represents the class `Point`, a programmer might write the following code²:

```
jobject pointClass = FindClass("java/awt/Point");
```

Here the `FindClass` function looks up a class by name. The resulting object `pointClass` is used to access fields and methods in and create new instances of class `Point`. For example, for a field access, the programmer next writes

```
jfieldID fid = GetFieldID(pointClass, "x", "I")
```

²The JNI functions mentioned in this section actually take an additional parameter, mentioned in Section 4.

```

int my_getIntField(jobject obj, char *field) {
    jobject cls = GetObjectClass(obj);
    jfieldID fid = GetFieldID(cls, field, "I");
    return GetIntField(obj, fid);
}

```

Figure 1: JNI Wrapper Function Example

After this call, `fid` contains a representation of the location of the field `x` with type `I` (a Java `int`) in class `Point`. This last parameter is a terse encoding of Java types called a *field descriptor*. Other examples are `F` for float, `[I` for array of integers, and `Ljava/lang/String;` for class `String`. Notice this is a slightly different encoding of class names than used by `FindClass`. Our implementation enforces this difference, but we omit it from the formal system for simplicity.

Finally, to read this field from a `Point` object `p`, the programmer writes

```

jobject p = ...;
int x = GetIntField(p, fid);

```

The function `GetIntField` returns an `int`, and there is one such function for each primitive type and one function `GetObjectField` for objects.

Thus we can see that a simple field access that would be written `int y = p.x` in Java requires three JNI calls, each of which corresponds to one internal step of the JVM: getting the type of the object, finding the offset of the field, and retrieving its contents. And while a Java compiler only accepts the code `y = p.x` if it is type correct, errors in C glue code, such as typos in `java/awt/Point`, `x`, or `I`, will produce a run time error. There also are several other places where mistakes could hide. For example, the programmer must be careful to maintain the dependence between the type of `x` and the call to `GetIntField`. If the type of `x` were changed to `float`, then the call must also be changed, to `GetFloatField`, something that is easy to overlook. Moreover, since `pointClass` and `p` both have type `jobject`, either could be passed where the other is expected with no C compiler warning, which we have seen happen in our benchmarks.

Invoking a Java method is similar to extracting a field, except the programmer calls `GetMethodID`, which accepts a string that encodes a list of parameter types. We omit the details due to lack of space.

One common pattern we have seen in JNI code is wrapper functions that specialize JNI routines to particular classes, fields, or methods. For example, Fig. 1 contains a function `my_getIntField` that extracts an integer field from an object. This routine invokes the JNI function `GetObjectClass`, which returns an object representing the class of its argument (as opposed to `FindClass`, which looks up a class by name). Calling `my_getIntField` is safe if the first parameter has an integer field whose name is given by the second parameter. Thus this function is parameterized by the object and by the name of the field, but not its type. Since this wrapper function might be called multiple times with different objects and different field names, we need to perform polymorphic type inference, not only in the types of objects but also in the values of string parameters.

3 Type System

In this section, we describe our multi-lingual type inference system. The input to our system is a collection of Java classes and a C program. Our type inference system analyzes the C program and generates *constraints* that describe how it uses Java objects. We also add constraints based on the type signatures of native methods and the actual class definitions from Java. We then solve the constraints to determine whether the C code uses Java objects consistently with the way they are declared, and we report any discrepancies as type errors.

Extracting type information from Java class files is relatively straightforward. Thus most of the work in our system occurs in the C analysis and constraint resolution phases, so that is the focus of this section.

$$\begin{aligned}
ct & ::= \text{void} \mid \text{int} \mid \text{str}\{s\} \mid (ct \times \dots \times ct) \rightarrow ct \\
& \quad \mid jt \text{ jobject} \mid (f, o) \text{ jfieldID} \mid (m, o) \text{ jmethodID} \\
jt & ::= \alpha \mid \text{JVoid} \mid \text{JInt} \mid o \mid jt \text{ JClass} \mid \text{JTSr}\{s\} \\
o & ::= \{s; F; M\} \\
s & ::= \nu \mid \text{"Str"} \\
f & ::= s : jt \\
F & ::= \phi \mid \emptyset \mid \langle f; \dots ; f \rangle \circ F \\
m & ::= s : (jt \times \dots \times jt) \rightarrow jt \\
M & ::= \mu \mid \emptyset \mid \langle m; \dots ; m \rangle \circ M
\end{aligned}$$

Figure 2: Type Language

3.1 Multi-Lingual Types

To check that C glue code uses Java types correctly, we must reconstruct the Java types for C variables that are declared with types like `jobject`, `jfieldID`, and `jmethodID`. Before giving our type language formally, consider the function `my_getIntField` in Fig. 1. Our system will give this function the following type, which corresponds to the informal description at the end of Section 2:

$$\text{my_getIntField} : \{\nu; \langle \nu_{field} : \text{JInt} \rangle \circ \phi; \mu\} \text{ jobject} \times \text{str}\{\nu_{field}\} \rightarrow \text{JInt jobject}$$

This is the type of a function that takes two parameters and returns a Java integer (which has the same representation as a C `int`). The second parameter is a C string whose contents are represented by the variable ν_{field} . The first parameter is a `jobject` that contains a field with name ν_{field} and Java type `int`. Here we have embedded Java type information into the bare C type in order to perform checking. The type of `my_getIntField` places no constraints on the class name ν of the first parameter or its other fields ϕ and methods μ . In order to infer this type, we need to track intermediate information about `cls` and `fid` as well. A detailed explanation of how this type is inferred is given in Section 3.4.

Our formal type grammar is given in Fig. 2. C types `ct` include `void`, integers, string types `str{s}` (corresponding to C types `char *`), and function types. In type `str{s}`, the string s may be either a known constant string `"Str"` or a type variable ν that is later resolved to a constant string. For example, in Fig. 1, `field` is given type `str{\nu_{field}}`, and ν_{field} is then used as a field name in the type of `my_getIntField`.

Our type language embeds a Java type jt in the C type `jobject`. In order to model the various ways the JNI can be used to access objects, we need a richer representation of Java types than just simple class names. For example, the wrapper function in Fig. 1 may be safely called with mutually incompatible classes as long as they all have the appropriate integer field. Thus our Java types include type variables α , the primitives `JVoid` and `JInt`, and *object descriptions* o of the form $\{s; F; M\}$, which represents an object whose class is named s with *field set* F and *method set* M . Since our inference system may discover the fields and methods of an object incrementally, we allow these sets to grow with the composition operator \circ . A set is *closed* if it is composed with \emptyset , and it is *open* if it is composed with a variable ϕ or μ . Since we never know just from C code whether we have accessed all the fields and methods of a class, field and method sets become closed only when unified with known Java types.

Because instances of Java’s `Class` class are essential for using the JNI, we distinguish them from other objects with the type $jt \text{ JClass}$. For example, in Fig. 1 the variable `cls` is given type $\{\nu; \langle \nu_{field} : \text{JInt} \rangle \circ \phi; \mu\} \text{ JClass}$, meaning it is the class of `obj`. Lastly, Java objects whose types are described by string s have type `JTSr{s}`. During inference when the value of s is determined, `JTSr{s}` will be replaced by the appropriate type. For example, initially the result type of `my_getIntField` is determined to be `JTSr{"I"}`, which is immediately replaced by `JInt`.

The types $(f, o) \text{ jfieldID}$ and $(m, o) \text{ jmethodID}$ represent intermediate JNI values for extracting field f or method m . The name of a field or method is a string s . For example, `fid` in Fig. 1 has type $(\nu_{field} : \text{JInt}, o)$ where o is our representation of `obj`. We include o so that we can check that this field identifier is used with an object of the correct type.

Given this type grammar, we can precisely describe the types of the JNI functions. Fig. 3 gives the types for the functions we have seen so far in this paper, plus two additional functions for invoking methods. For

$$\begin{aligned}
\text{FindClass} &: (\text{str}\{\nu\}) \rightarrow \text{JTStr}\{\nu\} \text{JClass } \text{jobject} \\
\text{GetObjectClass} &: (\{\nu; \phi; \mu\} \text{jobject}) \rightarrow \{\nu; \phi; \mu\} \text{JClass } \text{jobject} \\
\text{GetFieldID} &: (o \text{JClass } \text{jobject} \times \text{str}\{\nu_2\} \times \text{str}\{\nu_3\}) \rightarrow (f, o) \text{jfieldID} \\
&\quad \text{where } o = \{\nu_1; \langle f \rangle \circ \phi; \mu\} \text{ and } f = \nu_2 : \text{JTStr}\{\nu_3\} \\
\text{GetIntField} &: (o \text{jobject} \times (f, o) \text{jfieldID}) \rightarrow \text{JInt } \text{jobject} \\
&\quad \text{where } o = \{\nu_1; \langle f \rangle \circ \phi; \mu\} \text{ and } f = \nu_2 : \text{JInt} \\
\text{GetObjectField} &: (o \text{jobject} \times (f, o) \text{jfieldID}) \rightarrow \{\nu_3; \phi'; \mu'\} \text{jobject} \\
&\quad \text{where } o = \{\nu_1; \langle f \rangle \circ \phi; \mu\} \text{ and } f = \nu_2 : \{\nu_3; \phi'; \mu'\} \\
\text{GetMethodID} &: (o \text{JClass } \text{jobject} \times \text{str}\{\nu_2\} \times \text{str}\{\nu_3\}) \rightarrow (m, o) \text{jmethodID} \\
&\quad \text{where } o = \{\nu_1; \phi; \langle m \rangle \circ \mu\} \text{ and } m = \nu_2 : \text{JTStr}\{\nu_3\} \\
\text{CallIntMethod} &: (o \text{JClass } \text{jobject} \times (m, o) \text{jmethodID} \times \\
&\quad \alpha_1 \text{jobject} \times \dots \times \alpha_n \text{jobject}) \rightarrow \text{JInt } \text{jobject} \\
&\quad \text{where } o = \{\nu_1; \phi; \langle m \rangle \circ \mu\} \text{ and } m = \nu_2 : (\alpha_1 \times \dots \times \alpha_n) \rightarrow \text{JInt}
\end{aligned}$$

Figure 3: Sample JNI Type Signatures. All unconstrained variables are quantified.

instance, `FindClass` is a function that takes a string ν and returns an instance of `Class` that depends on the string represented by ν . The function `GetIntField` takes an object with a field f and a `jfieldID` describing f , and returns an integer. Notice that the object type o is open (expandable), because it may have other fields in addition to f . As we discussed in Section 2, it is important that these functions be polymorphic. In the type signatures in Fig. 3, any unconstrained variables are implicitly quantified. In the next section we describe how we use constraints to represent polymorphic function instantiation.

3.2 Constraint Generation

The core of our system is a type inference algorithm that traverses C source code and infers types in Fig. 2. Due to lack of space, we omit explicit typing rules for the source language, since they are mostly standard. Each C term of type `jobject` is initially assigned type α `jobject` for fresh α , and similarly for `jfieldID` and `jmethodID`. As we traverse the program, we generate *unification constraints* of the form $a = b$ (where a and b range over `ct`, `jt`, etc) whenever two terms must have the same type. For example, for the code in Fig. 1, we unify the type of `cls` with the return type of `GetObjectClass`. Since the only way to manipulate `jobject` types is through the JNI, in general our analysis only generates constraints at assignment, function definition, and function application. Because we use unification for `jt` types rather than subtyping, our inference algorithm could fail to unify two Java objects where one object is a subtype of the other. However, we did not find this to be a problem in practice (see Section 4).

To support polymorphism, we use *instantiation constraints* of the form $a \preceq_i b$ to model substitutions for quantified variables [19, 10]. Formally, we use the following two rules for generalization and instantiation:

$$\begin{array}{c}
\text{(LET)} \\
\frac{\Gamma \vdash e_1 : t_1 \quad \Gamma[f \mapsto (t_1, fv(\Gamma))] \vdash e_2 : t_2}{\Gamma \vdash \text{let } f = e_1 \text{ in } e_2 : t_2} \\
\text{(INST)} \\
\frac{\Gamma(f) = (t, \vec{\alpha}) \quad t \preceq_i \beta \quad \vec{\alpha} \preceq_i \vec{\alpha} \quad \beta \text{ fresh}}{\Gamma \vdash f_i : \beta}
\end{array}$$

In (Let), we represent a polymorphic type as a pair $(t, \vec{\alpha})$, where t is the base type and $\vec{\alpha}$ is the set of variables that may *not* be quantified. Then in (Inst), we generate an instantiation constraint $t \preceq_i \beta$. This constraint requires that there exist a substitution S_i such that $S_i(t) = \beta$. Our type rule also demands that

$\bar{\alpha} \preceq_i \bar{\alpha}$, i.e., S_i does not instantiate the free variables of the environment. As part of the constraint solving process (Section 3.3), we will implicitly find these substitutions S_i . The main advantage to this notation for polymorphism is that it allows us to traverse the source code in any order. In particular, we can generate instantiation constraints for a call to function f before we have seen a definition of f . A full discussion of these rules is beyond the scope of this paper and can be found elsewhere [19].

We have formalized a checking version of our type system in terms of lambda-calculus extended with strings, let-bound polymorphism, and primitives representing Java objects. We also include as primitive a set of JNI functions for operating on Java objects, and the small-step semantics for the language includes a reduction rule for each of these functions.

Theorem 1 (Soundness) *If $\Gamma \vdash e : ct$, then there exists a value v such that $e \rightarrow^* v$ and $\Gamma \vdash v : ct$.*

A proof of soundness can be found in Appendix B.

In our implementation, we do not keep track of the set $fv(\Gamma)$ for functions. Since C does not have nested functions, we simply issue warnings at any uses of global variables of type `object`. In general we have found that programmers use few globals when interacting with the JNI. We do not issue warnings for global `char *` types since these are common in C programs and would generate a very high rate of false positives.

3.3 Constraint Solving

Our type inference algorithm generates a set of constraints that we need to solve in order to decide if the program is well typed. To solve the constraints, we use a variant of Föhndrich et al’s worklist algorithm for semi-unification [10].

To simplify our constraint resolution rules below, whenever we refer to field and method sets we always assume they have been flattened so that they are composed with either \emptyset or a variable. During the process of unification, unknown strings ν will be replaced by known constant strings `str{Str}`. As this happens, we need to ensure that our object types still make sense. In particular, Java classes may not contain two fields with the same name but different types.³ Thus we also impose a well-formedness constraint on all field sets: any two fields with the same name in a field set must have the same type. Formally, for a field $f = s : jt$ we define $fname(f) = s$ and $ftype(f) = jt$. Then a field set $\langle f_1; \dots; f_n \rangle$ is well-formed if $fname(f_i) = fname(f_j) \Rightarrow ftype(f_i) = ftype(f_j)$ for all i, j , and we require that all field sets are well-formed. However, unlike fields, methods may be overloaded in Java. Therefore the above well-formed condition does not apply to methods.

During constraint solving, our system may unify a string variable ν with a constant string “Str”. When this occurs, our system uses the *Eval* function shown below to convert a type `JTStr{s}` into the Java type it represents:

$$\begin{aligned}
 Eval(\text{JTStr}\{\text{“V”}\}) &\Rightarrow \text{JVoid} \\
 Eval(\text{JTStr}\{\text{“I”}\}) &\Rightarrow \text{JInt} \\
 Eval(\text{JTStr}\{\text{“Ljava/lang/String;”}\}) &\Rightarrow \{\text{“java/lang/String”;} \dots; \dots\} \\
 Eval(\text{JTStr}\{\text{“Ljava/awt/Point;”}\}) &\Rightarrow \{\text{“java/awt/Point”;} \dots; \dots\} \\
 &\vdots
 \end{aligned}$$

We use a similar function to convert the string representation of a method signature into a list of Java types.

We express constraint solving in terms of rewrite rules, shown in Fig. 4. Given a set of constraints C , we apply these rules exhaustively, replacing the left-hand side with the right-hand side until we reach a fixpoint. Technically because we use semi-unification this algorithm may not terminate, but we have not found a case of this in practice. The complete list of rewrite rules is long and mostly standard, and so Fig. 4 contains only the interesting cases. The exhaustive set of rules may be found in Appendix A.

In Fig. 4, the (Closure) rule unifies two terms b and c when they are both instantiations of the same variable a at the same instantiation site. Intuitively, this rule enforces the property that substitution S_i must replace variable a consistently [19]. The rule (JTStr Ineq) applies the usual semi-unification rule for constructed types. Since the substitution S_i renames the left-hand side to yield the right-hand side in a

³JNI operations do not walk the Java type hierarchy, and so an overloaded field via inheritance is not possible.

(Closure)	$C \cup \{a \preceq_i b\} \cup \{a \preceq_i c\} \Rightarrow C \cup \{a \preceq_i b\} \cup \{b = c\}$	
(JTStr Ineq)	$C \cup \{ \text{JTStr}\{s\} \preceq_i jt \} \Rightarrow C \cup \{jt = \text{JTStr}\{\nu\}\} \cup \{s \preceq_i \nu\}$	ν fresh
(FieldSet EmptyErr1)	$C \cup \{\emptyset \preceq_i \langle f; \dots; f \rangle\} \Rightarrow \text{error}$	
(FieldSet EmptyErr2)	$C \cup \{\langle f; \dots; f \rangle \preceq_i \emptyset\} \Rightarrow \text{error}$	
(FieldSet InEq)	$C \cup \{\langle f_1; \dots; f_n \rangle \circ F \preceq_i \langle f'_1; \dots; f'_m \rangle \circ F'\} \Rightarrow$	$C \cup \{ \text{ftype}(f_k) \preceq_i \text{ftype}(f'_j) \mid \text{fname}(f_k) = \text{fname}(f'_j) \}$ $\cup \{ \langle f_k \mid \exists j. \text{fname}(f_k) = \text{fname}(f'_j) \rangle \preceq_i F' \}$ $\cup \{ F \preceq_i \langle f'_j \mid \exists i. \text{fname}(f_k) = \text{fname}(f'_j) \rangle \}$
(Str Sub)	$C \cup \{\nu_1 = \nu_2\} \Rightarrow C[\nu_1 \mapsto \nu_2]$	
(Str Resolve)	$C \cup \{\nu = \text{"Str''"}\} \Rightarrow C[\nu \mapsto \text{"Str''"}][\text{JTStr}\{\nu\} \mapsto \text{Eval}(\text{"Str''"})]$	
(Str Eq)	$C \cup \{\text{"Str}_1" = \text{"Str}_2"\} \Rightarrow C$	$\text{Str}_1 = \text{Str}_2$
(Str Neq)	$C \cup \{\text{"Str}_1" = \text{"Str}_2"\} \Rightarrow \text{error}$	$\text{Str}_1 \neq \text{Str}_2$
(JTStr Sub)	$C \cup \{\alpha = \text{JTStr}\{s\}\} \Rightarrow C[\alpha \mapsto \text{JTStr}\{s\}]$	
(JTStr Void)	$C \cup \{\text{JVoid} = \text{JTStr}\{s\}\} \Rightarrow C \cup \{s = \text{"V"}\}$	
(JTStr Obj)	$C \cup \{\{\nu; F; M\} = \text{JTStr}\{s\}\} \Rightarrow C \cup \{s = \nu\}$	

Figure 4: Selected Constraint Rewrite Rules

constraint \preceq_i , the right-hand side must have the same shape. Thus in (JTStr Ineq), we unify jt with $\text{JTStr}\{\nu\}$ where ν is fresh and propagate the semi-unification constraint to s and ν .

The next set of rules is for field sets. In (FieldSet EmptyErr1) and (FieldSet EmptyErr2), we forbid the empty set of fields from instantiation to or from a non-empty set. In (FieldSet InEq), we match up the fields of two non-empty field sets. We directly propagate instantiation constraints to fields f_k and f'_j with the same field name or variable. We then make an instantiation constraint to F' from the remaining fields of the left side for which there does not exist a field with the same name on the right-hand side. Recall that we assume the sets have been flattened, so that F' is either a variable or \emptyset . We then generate the analogous constraint for F and the right-hand side. Notice that this process may result in a field set with multiple fields with variables ν for names. Our implicit well-formedness requirement from Section 3 will handle the case where these variables are later unified with known strings.

We omit the rules for method sets due to lack of space. These rules are similar to the field set rules except for one important detail. Suppose we have an open method set $\langle x : \alpha \rightarrow \alpha \rangle \circ \mu$ which is unified with a closed method set $\langle x : \text{JInt} \rightarrow \text{JInt}; x : \text{JVoid} \rightarrow \text{JVoid} \rangle \circ \emptyset$. Since the method x is overloaded, we do not know if α should unify with JInt or JVoid . Therefore, if this occurs during unification, our tool emits a warning and removes the constraint.

The next three rules handle strings. (Str Sub) replaces one string variable by another (we write $C[a \mapsto b]$ to mean substitute all instance of a with b in set C). (Str Resolve) uses Eval to replace all occurrences of $\text{JTStr}\{\nu\}$ with their appropriate representations. (Str Eq) and (Str Neq) test string constants for equality.

Because $\text{JTStr}\{\}$ encodes its type as a string, we use a slightly different rewrite rule for this case. In rule (JTStr Sub), if a $\text{JTStr}\{\}$ type is unified with a type variable, then the variable is replaced as normal. However, if a $\text{JTStr}\{s\}$ type is unified with a void type as in (JTStr Void), then we add the constraint that the $s = \text{"V"}$, since $\text{Eval}(s)$ must produce a void type. We use a similar rule for integers. Similarly, the rule (JTStr Obj) adds the constraint that s must have the same name as the object it unifies with.

The remainder of the rewrite rules are standard, and all either replace type variables with terms or match up type constructors and propagate constraints recursively to the constructor parameters.

3.4 Example

In this section, we demonstrate our inference system on the `my_getIntField` function from Fig. 1. Initially our analysis assigns each parameter and local variable of this function a fresh type:

$$\begin{array}{ll} \text{obj} & : \alpha_{obj} \text{ jobject} & \text{cls} & : \alpha_{cls} \text{ jobject} \\ \text{field} & : \text{str}\{\nu_{field}\} & \text{fid} & : (\{\nu_{fid}; \phi_{fid}; \mu_{fid}\} \text{ jfieldID}) \end{array}$$

The first line of the function calls the `GetObjectClass` function (call this instantiation site 1). After looking up its type in the environment (shown in Fig. 3 with quantified type variables ν , ϕ , and μ), we add the following constraints:

$$\begin{array}{ll} \{\nu; \phi; \mu\} \text{ jobject} & \preceq_1 \alpha_{obj} \text{ jobject} \\ \{\nu; \phi; \mu\} \text{ jclass} & \preceq_1 \alpha_{cls} \text{ jobject} \end{array}$$

Applying our constraint rewrite rules yields the following new constraints:

$$\begin{array}{ll} \alpha_{obj} = \{\nu_2; \phi_2; \mu_2\} & \alpha_{cls} = \{\nu_3; \phi_3; \mu_3\} \text{ jclass} \\ \nu \preceq_1 \nu_2 & \nu \preceq_1 \nu_3 \\ \phi \preceq_1 \phi_2 & \phi \preceq_1 \phi_3 \\ \mu \preceq_1 \mu_2 & \mu \preceq_1 \mu_3 \\ \nu_2, \phi_2, \mu_2 \text{ fresh} & \nu_3, \phi_3, \mu_3 \text{ fresh} \end{array}$$

Then rule (Closure) in Fig. 4 generates the constraints $\nu_2 = \nu_3$, $\phi_2 = \phi_3$, and $\mu_2 = \mu_3$ to require that the substitution corresponding to this call is consistent. Next, `my_getIntField` calls `GetFieldID`, and after applying our constraint rules, we discover (among other things):

$$\phi_2 = \langle \nu_{field} : \text{JTStr}\{\text{"I"}\} \rangle \circ \phi_4 \quad \phi_4 \text{ fresh}$$

Now since we have a `JTStr` with a known string, our resolution rules call `Eval` to replace it, yielding

$$\phi_2 = \langle \nu_{field} : \text{JInt} \rangle \circ \phi_4$$

The last call to `GetIntField` generates several new constraints, but they do not affect the types. Thus after the function has been analyzed, it is given the type

$$\text{my_getIntField} : \{\nu_1; \langle \nu_{field} : \text{JInt} \rangle \circ \phi_4; \mu\} \text{ jobject} \times \text{str}\{\nu_{field}\} \rightarrow \text{JInt} \text{ jobject}$$

In other words, this function accepts any object as the first parameter as long as it has an integer field whose name is given by the second parameter, exactly as intended.

4 Implementation and Experiments

We have implemented the inference system described in Section 3 in the form of two tools used in sequence during the build process. The first tool is a light-weight Java compiler wrapper. The wrapper intercepts calls to `javac` and records the class path so that the second tool can retrieve class files automatically. The wrapper itself does not perform any code analysis. The second tool applies our type inference algorithm to C code and issues warnings whenever it finds a type error. Our tool uses CIL [26] to parse C source code and the OCaml JavaLib [7] to extract Java type information from compiled class files.

Our implementation contains some additional features not discussed in the formal system. In addition to the type `jobject`, the JNI contains a number of typedefs (aliases) for other object types, such as `jstring` for Java `Strings`. These are all aliases of `jobject`, and so their use is not required by the JNI, and they do not result in any more checking by the C compiler. Our system does not require their use either, but since they are a form of documentation we enforce their intended meaning, e.g., values of type `jstring` are assigned a type corresponding to `String`. We found 3 examples in our benchmarks where programmers used the wrong alias. The JNI also defines types `jvoid` and `jint`, representing Java voids and integers, as aliases

Program	C LOC	Java LOC	Time (s)	Errs	Warnings	False Pos	Impr
libgconf-java-2.10.1	1119	670	2.4	0	0	10	0
libglade-java-2.10.1	149	1022	6.9	0	0	0	1
libgnome-java-2.10.1	5606	5135	17.4	45	0	0	1
libgtk-java-2.6.2	27095	32395	36.3	74	8	34	18
libgtkhtml-java-2.6.0	455	729	2.9	27	0	0	0
libgtkmozembed-java-1.7.0	166	498	3.3	0	0	0	0
libvte-java-0.11.11	437	184	2.5	0	26	0	0
jnetfilter	1113	1599	17.3	9	0	0	0
libreadline-java-0.8.0	1459	324	2.2	0	0	0	1
pgpjava	10136	123	2.7	0	1	0	1
posix1.0	978	293	1.8	0	1	0	0
Total				155	36	44	22

Figure 5: Experimental Results

of the C types `void` and `int`, respectively. Our system does not distinguish between the C name and its `j-`prefixed counterpart.

Rather than being called directly, JNI functions are actually stored in a table that is passed as an extra argument (usually named `env`) to every C function called from Java, and this table is in turned passed to every JNI function. For example, the `FindClass` function is actually called with `(*env)->FindClass(env, ...)`. Our system extracts FFI function names via syntactic pattern matching, and we assume that the table is the same everywhere. Function pointers that are not part of the JNI are not handled by our system, and we do not generate any constraints when they are used in a program.

The JNI functions for invoking Java methods must take a variable number of arguments, since they may be used to invoke methods with any number of parameters. Our system handles the commonly-used interface, which is JNI functions declared to be `varargs` using the `...` convention in C. However, the JNI provides two other calling mechanisms that we do not model: Passing arguments as an array, and passing arguments using the special `va_list` structure. We issue warnings if either is used.

Although our type system is flow-insensitive, we treat the types of local variables flow-sensitively. Each assignment updates the type of a variable in the environment, and we add a unification constraint to variables of the same name at join points in the control flow graph. See [13] for details.

Lastly, our implementation models strings in a very simple way to match how they are used in practice in C glue code. We currently ignore string operations like `strcat` or destructive updates via array operations. We also assume that strings are always initialized before they are used, since most compilers produce a warning when this is not the case.

We ran our analysis tool on a suite of 11 benchmarks that use the JNI. Fig. 5 shows our results. All benchmarks except `pgpjava` are glue code libraries that connect Java to an external C library. The first 7 programs are taken from the Java-Gnome project [22], and the remaining programs are unrelated. For each program, Fig. 5 lists the number of lines of C and Java code, the analysis time in seconds (3 run avg.), and the number of messages reported by our tool, divided manually into four categories as described below. The running time includes the C code analysis (including extracting Java types from class files) but not the parsing of C code. The measurements were performed on a 733 MHz Pentium III machine with 1GB of RAM.

Our tool reported 155 errors, which are programming mistakes that may cause a program to crash or to emit an unexpected exception. Surprisingly, the most common error was declaring a C function with the wrong arity, which accounted for 68 errors (30 in `libgtk` and 38 in `libgnome`). All C functions called from Java must start with one parameter for the JNI environment and a second parameter for the invoking object or class. In many cases the second parameter was omitted in the call, and hence any subsequent arguments would be retrieved from the wrong stack location, which could easily cause a program crash.

56 of the errors were due to mistakes made during a software rewrite. Programs that use the JNI typically use one of two techniques to pass C pointers (e.g., GUI window objects) through Java: they either pass the pointer directly as an integer, or they embed the pointer as an integer field inside a Java object. Several of the libraries in the Java-Gnome project appear to be switching from the integer technique to the object technique, which requires changing Java declarations in parallel with C declarations, an error-prone process. Our tool detected many cases when a Java native method specified an `Object` parameter but the corresponding C function specified an integer parameter, or vice-versa. This accounted for 4 errors in `libgnome`, 25 in `libgtk`, and 27 in `libgtkhtml`.

Type mismatches accounted for 17 of the remaining errors. 6 errors occurred because a native Java method was declared with a `String` argument, but the C code took a byte array argument. In general Java strings must be translated to C strings using special JNI functions, and hence this is a type error. Another type error occurred because one C function passed a (non-array) Java object to another C function expecting a Java array. Since both of these are represented with the type `jobject` in C, the C compiler did not catch this error.

Finally, 14 errors were due to incorrect namings. 11 of these errors (9 in `jnetfilter` and 2 in `gtk-java-2.6.2`) were caused by calls to `FindClass` with an incorrect string. Ironically, all 9 `jnetfilter` errors occurred in code that was supposed to construct a Java exception to throw—but since the string did not properly identify the exception class, the JVM would throw a `ClassNotFoundException` exception instead. The remaining 3 errors were due to giving incorrect names to C functions corresponding to Java native methods; such functions must be given long names following a particular naming scheme, and it is easy to get this wrong.

Most of the errors we found are easy to trigger with a small amount of code. In cases such as incorrectly-named function, errors would likely be immediately apparent as soon as the native method is called. Thus clearly many of the errors are in code that has not been tested very much, most likely the parts of libraries that have not yet been used by Java programmers.

Our tool also produced 36 warnings, which are suspicious programming practices that do not actually cause run-time errors. One warning arose when a programmer called the function `FindClass` with a field descriptor of the form `Ljava/lang/String`; rather than a fully qualified class name `java/lang/String`. Technically this is an error [24], but the Sun JVM we tested allows both versions, so we only consider this a warning.

Another example that accounted for 2 warnings in our benchmarks was due to incorrect type declarations but correct usage. For example, one C function was declared to return a byte array, but it corresponded to a native method that was declared to return a `String`. As we mentioned earlier this is an error, except that in this case the C function actually returned an object of type `String` and so behaved correctly at run time.

Finally, 33 warnings were due to the declaration of C functions that appear to implement a specific native method (because they have mangled names), but do not correspond to any native Java method. In many cases there was a native method in the Java code, but it had been commented out or moved without deleting the C code. This will not cause any run-time errors, but it seems helpful to notify the programmer about this dead code.

Our tool also produced 44 false positives, which are warnings that correspond to code that is actually correct. All of the false positives were caused by the use of subtyping inside of C code, which our analysis does not model precisely because it uses unification. Our tool also produced 22 imprecision warnings, which are cases where the analysis had insufficient information about the value of a string. 16 of the warnings were due to unification failures with partially specified methods that could not be resolved because of Java's overloading rules. The other 6 warnings occurred when the programmer called a Java method by passing arguments via an array. Since our analysis does not model arrays, we conservatively produce a warning at this point.

5 Related Work

In prior work, we presented a system for inferring types for programs that use the OCaml-to-C FFI [13]. Our new work on the JNI differs in several ways: To infer types for the JNI, we need to track string values through

in the source code, whereas for the OCaml FFI we mainly needed to track integers and pointer offsets. Our new system can directly assign polymorphic types to JNI functions and to user-written wrapper functions, in contrast to our previous system which did not track sufficient interprocedural information to allow this and was not polymorphic. Our new system also includes support for indexing into records using strings for field names and method signatures, and those strings may not be known until constraint resolution. Our previous system did not support objects.

Recently several researchers have developed sophisticated techniques to track strings in programs [8, 9, 14]. One goal of these systems is to check that dynamically-generated SQL queries are well-formed. For purposes of checking clients of the JNI, we found that simple tracking of string constants was sufficient.

There are many different foreign function interfaces with various design tradeoffs [5, 11, 12, 20, 23, 24]. We believe that the techniques we develop in this paper and in our previous work [13] can be adapted to many FFIs.

Trifonov and Shao [29] develop a technique for ensuring that code fragments in languages with different runtime resource requirements have access to necessary resources while preserving semantics, which is different than our goal of type safety.

An alternative to using FFIs directly is to use automatic interface generators to produce glue code. SWIG [3] generates glue code based on an interface specification file. Exu [6] provides programmers with a light-weight system for automatically generating JNI-to-C++ glue code for the common cases. Mockingbird [1] is a system for automatically finding matchings between two types written in different languages and generating the appropriate glue code. Our benchmark suite contained custom glue code that was generated by hand.

In addition to the JNI, there has been significant work on other approaches to object-oriented language interoperation, such as the commercial solutions COM [15], SOM [17] and CORBA [27]. Barret [2] proposes the PolySPIN system as an alternative to CORBA. All of these systems check for errors mainly at runtime (though in some cases interface generators can be used to provide some compile-time checking). The Microsoft common-language runtime (CLR) [18, 25] provides interoperation by being the target of compilers for multiple different languages, and the CLR includes a strong static type system. However, the type system only checks CLR code, and not unmanaged code that it may invoke.

Grechanik et al [16] present a framework called ROOF that allows many different languages to interact using a common syntax. This system includes both run-time and static type checking for code that uses ROOF. It is unclear whether ROOF supports polymorphism and whether it can infer types for glue code in isolation.

6 Acknowledgments

We wish to thank Michael Hicks for his helpful comments on earlier versions of this paper.

7 Conclusion

We have presented a multi-lingual type inference system for checking that programs use the JNI safely. Our system tracks the values of C strings to determine what names and type descriptors are passed to JNI functions. Thus we are able to infer what type assumptions C glue code makes about Java objects and check that they are consistent with the actual Java class definitions. Our system treats wrapper functions and JNI functions polymorphically, allowing them to be parametric even in string arguments. Using an implementation of our system, we found many errors and suspicious practices in our suite of benchmarks. Our results suggest that our static checking system can be an important part of ensuring that the JNI is used correctly, and we believe that the same ideas can be applied to other object-oriented FFIs.

References

- [1] J. Auerbach, C. Barton, M. Chu-Carroll, and M. Raghavachari. Mockingbird: Flexible stub compilation from paris of declarations. In *Proceedings of the 19th International Conference on Distributed Computing Systems*, pages 393–402, Austin, TX, USA, 1999. IEEE Computer Society.
- [2] D. J. Barrett. *Polylingual Systems: An Approach to Seamless Interoperability*. PhD thesis, University of Massachusetts Amherst, Feb. 1998.
- [3] D. M. Beazley. SWIG: An easy to use tool for integrating scripting languages with C and C++. In *USENIX Fourth Annual Tcl/Tk Workshop*, July 1996.
- [4] N. Benton and A. Kennedy, editors. *BABEL'01: First International Workshop on Multi-Language Infrastructure and Interoperability*, volume 59 of *Electronic Notes in Theoretical Computer Science*, Firenze, Italy, Sept. 2001. <http://www.elsevier.nl/locate/entcs/volume59.html>.
- [5] M. Blume. No-Longer-Foreign: Teaching an ML compiler to speak C “natively”. In Benton and Kennedy [4]. <http://www.elsevier.nl/locate/entcs/volume59.html>.
- [6] J. F. Bubba, A. Kaplan, and J. C. Wileden. The Exu Approach to Safe, Transparent and Lightweight Interoperability. In *25th International Computer Software and Applications Conference (COMPSAC 2001)*, pages 393–400, Chicago, IL, USA, Oct. 2001. IEEE Computer Society.
- [7] N. Cannasse. Ocaml javalib. <http://team.motion-twin.com/ncannasse/javaLib/>.
- [8] A. S. Christensen, A. Møller, and M. I. Schwartzbach. Precise Analysis of String Expressions. In R. Cousot, editor, *Static Analysis, 10th International Symposium*, volume 2694 of *Lecture Notes in Computer Science*, pages 1–18, San Diego, CA, USA, June 2003. Springer-Verlag.
- [9] R. DeLine and M. Fähndrich. The Fugue Protocol Checker: Is your software Baroque? Technical Report MSR-TR-2004-07, Microsoft Research, Jan. 2004.
- [10] M. Fähndrich, J. Rehof, and M. Das. Scalable Context-Sensitive Flow Analysis using Instantiation Constraints. In *Proceedings of the 2000 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 253–263, Vancouver B.C., Canada, June 2000.
- [11] S. Finne, D. Leijen, E. Meijer, and S. P. Jones. Calling hell from heaven and heaven from hell. In *Proceedings of the Fourth ACM SIGPLAN International Conference on Functional Programming*, pages 114–125, Paris, France, Sept. 1999.
- [12] K. Fisher, R. Pucella, and J. Reppy. A framework for interoperability. In Benton and Kennedy [4]. <http://www.elsevier.nl/locate/entcs/volume59.html>.
- [13] M. Furr and J. S. Foster. Checking Type Safety of Foreign Function Calls. In *Proceedings of the 2005 ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 62–72, Chicago, Illinois, June 2005.
- [14] C. Gould, Z. Su, and P. Devanbu. Static Checking of Dynamically Generated Queries in Database Applications. In ICSE'04 [21], pages 645–654.
- [15] D. N. Gray, J. Hotchkiss, S. LaForge, A. Shalit, and T. Weinberg. Modern Languages and Microsoft's Component Object Model. *cacm*, 41(5):55–65, May 1998.
- [16] M. Grechanik, D. Batory, and D. E. Perry. Design of large-scale polylingual systems. In ICSE'04 [21], pages 357–366.
- [17] J. Hamilton. Interlanguage Object Sharing with SOM. In *Proceedings of the Usenix 1996 Annual Technical Conference*, San Diego, California, Jan. 1996.

- [18] J. Hamilton. Language Integration in the Common Language Runtime. *ACM SIGPLAN Notices*, 38(2):19–28, Feb. 2003.
- [19] F. Henglein. Type Inference with Polymorphic Recursion. *ACM Transactions on Programming Languages and Systems*, 15(2):253–289, Apr. 1993.
- [20] L. Huelsbergen. A Portable C Interface for Standard ML of New Jersey. <http://www.smlnj.org/doc/SMLNJ-C/smlnj-c.ps>, 1996.
- [21] *Proceedings of the 26th International Conference on Software Engineering*, Edinburgh, Scotland, UK, May 2004.
- [22] Java-Gnome Developers. Java bindings for the gnome and gtk libraries. <http://java-gnome.sourceforge.net>.
- [23] X. Leroy. The Objective Caml system, Aug. 2004. Release 3.08, <http://caml.inria.fr/distrib/ocaml-3.08/ocaml-3.08-refman.pdf>.
- [24] S. Liang. *The Java Native Interface: Programmer’s Guide and Specification*. Addison-Wesley, 1999.
- [25] E. Meijer, N. Perry, and A. van Yzendoorn. Scripting .NET using Mondrian. In J. L. Knudsen, editor, *ECOOP 2001 - Object-Oriented Programming, 15th European Conference*, volume 2072 of *Lecture Notes in Computer Science*, pages 150–164, Budapest, Hungary, June 2001. Springer-Verlag.
- [26] G. C. Necula, S. McPeak, S. P. Rahul, and W. Weimer. CIL: Intermediate Language and Tools for Analysis and Transformation of C Programs. In R. N. Horspool, editor, *Compiler Construction, 11th International Conference*, volume 2304 of *Lecture Notes in Computer Science*, pages 213–228, Grenoble, France, Apr. 2002. Springer-Verlag.
- [27] Object Management Group. *Common Object Request Broker Architecture: Core Specification*, version 3.0.3 edition, Mar. 2004.
- [28] J. Rehof and M. Fähndrich. Type-Based Flow Analysis: From Polymorphic Subtyping to CFL-Reachability. In *Proceedings of the 28th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 54–66, London, United Kingdom, Jan. 2001.
- [29] V. Trifonov and Z. Shao. Safe and Principled Language Interoperation. In D. Swierstra, editor, *8th European Symposium on Programming*, volume 1576 of *Lecture Notes in Computer Science*, pages 128–146, Amsterdam, The Netherlands, Mar. 1999. Springer-Verlag.

A Complete Constraint Resolution Rules

In this section, we provide our complete rewrite rules for our system. Given our constraint grammar:

$$\text{op} ::= = \mid \preceq_i$$

we write $a \text{ op } b$ for the case when either $a = b$ or $a \preceq_i b$ is present. In order to unify methods, we require that two methods are strictly equal, i.e., no unification or substitution is allowed. We write this relation as $=_{eq}$ and define it formally as

$$\begin{array}{lcl}
 \alpha & =_{eq} & \alpha \\
 \nu & =_{eq} & \nu \\
 \phi & =_{eq} & \phi \\
 \mu & =_{eq} & \mu \\
 \text{JVoid} & =_{eq} & \text{JVoid} \\
 \text{JInt} & =_{eq} & \text{JInt} \\
 \emptyset & =_{eq} & \emptyset \\
 jt_1 =_{eq} jt_2 & \Rightarrow & jt_1 \text{ JClass} =_{eq} jt_2 \text{ JClass} \\
 s_1 =_{eq} s_2 \wedge F_1 =_{eq} F_2 \wedge M_1 =_{eq} M_2 & \Rightarrow & \{s_1; F_1; M_1\} =_{eq} \{s_2; F_2; M_2\} \\
 \begin{array}{l} f_1 =_{eq} f_{n+1} \wedge \dots \wedge f_n =_{eq} f_{n+n} \\ \wedge F_1 =_{eq} F_2 \end{array} & \Rightarrow & \langle f_1; \dots; f_n \rangle \circ F_1 =_{eq} \langle f_{n+1}; \dots; f_{n+n} \rangle \circ F_2 \\
 s_1 =_{eq} s_2 \wedge jt_1 =_{eq} jt_2 & \Rightarrow & s_1 : jt_1 =_{eq} s_2 : jt_2 \\
 \begin{array}{l} m_1 =_{eq} m_{n+1} \wedge \dots \wedge m_n =_{eq} m_{n+n} \\ \wedge M_1 =_{eq} M_2 \end{array} & \Rightarrow & \langle m_1; \dots; m_n \rangle \circ M_1 =_{eq} \langle m_{n+1}; \dots; m_{n+n} \rangle \circ M_2 \\
 \begin{array}{l} s_1 =_{eq} s_2 \wedge jt_1 =_{eq} jt_{n+1} \wedge \dots \\ \wedge jt_n =_{eq} jt_{n+n} \end{array} & \Rightarrow & \begin{array}{l} s_1 : (jt_1 \times \dots \times jt_n) =_{eq} \\ s_2 : (jt_{n+1} \times \dots \times jt_{n+n}) \end{array}
 \end{array}$$

When comparing two field or method sets, we implicitly allow their elements to be permuted so that relation does not require a specific ordering of their elements.

$$\begin{array}{l}
C \cup \{a \preceq_i b\} \cup \{a \preceq_i c\} \Rightarrow C \cup \{a \preceq_i b\} \cup \{b = c\} \\
C \cup \{\text{void op void}\} \Rightarrow C \\
C \cup \{\text{JVoid op void}\} \Rightarrow C \\
C \cup \{\text{void op JVoid}\} \Rightarrow C \\
C \cup \{\text{JVoid op JVoid}\} \Rightarrow C \\
C \cup \{\text{int op int}\} \Rightarrow C \\
C \cup \{\text{JInt op int}\} \Rightarrow C \\
C \cup \{\text{int op JInt}\} \Rightarrow C \\
C \cup \{\text{JInt op JInt}\} \Rightarrow C \\
C \cup \{\text{str}\{s_1\} \text{ op str}\{s_2\}\} \Rightarrow C \cup \{s_1 \text{ op } s_2\} \\
C \cup \{jt_1 \text{ jobject op } jt_2 \text{ jobject}\} \Rightarrow C \cup \{jt_1 \text{ op } jt_2\} \\
C \cup \{(f_1, o_1) \text{ jfieldID op } (f_2, o_2) \text{ jfieldID}\} \Rightarrow C \cup \{f_1 \text{ op } f_2\} \cup \{o_1 \text{ op } o_2\} \\
C \cup \{(m_1, o_1) \text{ jmethodID op } (m_2, o_2) \text{ jmethodID}\} \Rightarrow C \cup \{m_1 \text{ op } m_2\} \cup \{o_1 \text{ op } o_2\} \\
C \cup \{\alpha = jt\} \Rightarrow C[\alpha \mapsto jt] \\
C \cup \{\{s_1; F_1; M_1\} = \{s_1; F_1; M_1\}\} \Rightarrow C \cup \{s_1 = s_2\} \cup \{F_1 = F_2\} \cup \{M_1 = M_2\} \\
C \cup \{\{s; F; M\} \preceq_i \alpha\} \Rightarrow C \cup \{\alpha = \{\nu; \phi; \mu\}\} \cup \{s \preceq_i \nu\} \cup \{F \preceq_i \phi\} \cup \{M \preceq_i \mu\} \\
C \cup \{jt_1 \text{ JClass op } jt_2 \text{ JClass}\} \Rightarrow C \cup \{jt_1 \text{ op } jt_2\} \\
C \cup \{jt \text{ JClass} = \text{JTStr}\{s\}\} \Rightarrow C \cup \{s = \text{"java.lang.Class"}\} \\
C \cup \{jt \text{ JClass} = \{s; F; M\}\} \Rightarrow C \cup \{s = \text{"java.lang.Class"}\} \\
C \cup \{\text{JTStr}\{s\} \preceq_i jt\} \Rightarrow C \cup \{jt = \text{JTStr}\{\nu\}\} \cup \{s \preceq_i \nu\} \quad \nu \text{ fresh} \\
C \cup \{\text{JVoid} = \text{JTStr}\{s\}\} \Rightarrow C \cup \{s = \text{JTStr}\{\text{"V"}\}\} \\
C \cup \{\text{JInt} = \text{JTStr}\{s\}\} \Rightarrow C \cup \{s = \text{JTStr}\{\text{"I"}\}\} \\
C \cup \{\{\nu; F; M\} = \text{JTStr}\{s\}\} \Rightarrow C \cup \{s = \nu\} \\
C \cup \{\nu_1 = \nu_2\} \Rightarrow C[\nu_1 \mapsto \nu_2] \\
C \cup \{\nu = \text{"Str"}\} \Rightarrow C[\nu \mapsto \text{"Str"}][\text{JTStr}\{\nu\} \mapsto \text{Eval}(\text{"Str"})] \\
C \cup \{\text{"Str"}'_1 = \text{"Str"}'_2\} \Rightarrow C \quad \text{if } Str_1 = Str_2
\end{array}$$

Figure 6: Complete Constraint Rewrite Rules Part 1

$$\begin{aligned}
C \cup \{\emptyset \text{ op } \emptyset\} &\Rightarrow C \\
C \cup \{\phi = \phi'\} &\Rightarrow C[\phi \mapsto \phi'] \\
C \cup \{\mu = \mu'\} &\Rightarrow C[\mu \mapsto \mu'] \\
C \cup \{\phi = \langle f_1; \dots; f_n \rangle \circ F\} &\Rightarrow C[\phi \mapsto \langle f_1; \dots; f_n \rangle \circ F] \\
C \cup \{\langle f_1; \dots; f_n \rangle \circ F \preceq_i \phi\} &\Rightarrow C \cup \{\phi = \langle \nu_1 : \alpha_1; \dots; \nu_n : \alpha_n \rangle \circ \phi'\} \\
&\quad \cup \{f_j \preceq_i \nu_j : \alpha_j\} \cup F \preceq_i \phi' \\
C \cup \left\{ \begin{array}{l} C \cup \langle f_1; \dots; f_n \rangle \circ F \text{ op} \\ \langle f'_1; \dots; f'_m \rangle \circ F' \end{array} \right\} &\Rightarrow C \cup \{f_{\text{type}}(f_k) \text{ op } f_{\text{type}}(f'_j) \mid f_{\text{name}}(f_k) = f_{\text{name}}(f'_j)\} \\
&\quad \cup \{\langle f_k \mid \exists j. f_{\text{name}}(f_k) = f_{\text{name}}(f'_j) \rangle \text{ op } F'\} \\
&\quad \cup \{F \text{ op } \langle f'_j \mid \exists i. f_{\text{name}}(f_k) = f_{\text{name}}(f'_j) \rangle\} \\
C \cup \{\mu = \langle m_1; \dots; m_n \rangle \circ M\} &\Rightarrow C[\mu \mapsto \langle m_1; \dots; m_n \rangle \circ M] \\
C \cup \{\langle m_1; \dots; m_n \rangle \circ M \preceq_i \mu\} &\Rightarrow C \cup \{\mu = \langle m'_1; \dots; m'_n \rangle \circ \mu'\} \cup \{m_j \preceq_i m'_j\} \cup M \preceq_i \mu' \\
C \cup \left\{ \begin{array}{l} C \cup \langle m_1; \dots; m_n \rangle \circ M \preceq_i \\ \langle m_{n+1}; \dots; m_p \rangle \circ M' \end{array} \right\} &\Rightarrow C \cup \{\langle m'_1; \dots; m'_n \rangle \circ \mu = \langle m_{n+1}; \dots; m_p \rangle \circ M'\} \\
&\quad \cup \{M \preceq_i M'\} \cup \{m_1 \preceq_i m'_1\} \dots \cup \{m_n \preceq_i m'_n\} \\
C \cup \left\{ \begin{array}{l} \langle m_1; \dots; m_n \rangle \circ \mu_1 = \\ \langle m'_1; \dots; m'_p \rangle \circ \mu_2 \end{array} \right\} &\Rightarrow C[\mu_2 \mapsto \langle m_1; \dots; m_n \rangle \circ \mu'_2][\mu \mapsto \langle m'_1; \dots; m'_p \rangle \circ \mu'_1] \\
C \cup \left\{ \begin{array}{l} \langle m_1; \dots; m_n \rangle \circ \mu = \\ \langle m'_1; \dots; m'_p \rangle \circ \emptyset \end{array} \right\} &\Rightarrow C[\mu \mapsto \emptyset] \cup \{\langle m_1; \dots; m_n \rangle \circ \emptyset = \langle m'_1; \dots; m'_p \rangle \circ \emptyset\} \\
C \cup \left\{ \begin{array}{l} \langle m_1; \dots; m_n \rangle \circ \emptyset = \\ \langle m'_1; \dots; m'_p \rangle \circ \emptyset \end{array} \right\} &\Rightarrow C \text{ if } \langle m_1; \dots; m_n \rangle =_{eq} \langle m'_1; \dots; m'_p \rangle \\
C \cup \{\cdot \text{ op } \cdot\} &\Rightarrow \text{error}
\end{aligned}$$

Figure 7: Complete Constraint Rewrite Rules Part 2

B Soundness

We illustrate soundness for our inference system on the following variant of lambda calculus (we use this language rather than C because it is easier to reason about):

$$\begin{aligned}
e &::= v \mid \delta(e, \dots, e) \mid e e \mid \text{let } f = e \text{ in } e \\
v &::= x \mid f_i \mid \lambda x. e \mid \text{“Str”} \\
&\quad \mid \text{class(“Class”)} \mid \text{object(“Class”)} \mid \text{fid(“Name”, “FClass”, “OClass”)} \\
\delta &::= \text{FindClass} \mid \text{FindObjectClass} \mid \text{GetFieldID} \mid \text{GetObjectField}
\end{aligned}$$

This language is the lambda calculus extended with primitive strings “Str”. We have also added representations of Java classes, Java objects and Java field IDs. Since we are only concerned about type correctness, and since Java objects cannot be directly manipulated, we model all instances of *Class* with the value *object*(“Class”). Similarly, we model a field ID for a field *Name* of type *FClass* inside of *OClass* with the value *fid*(“Name”, “FClass”, “OClass”). We omit methods and method IDs, since they are similar to fields. We add primitive JNI functions given by δ . We also add let-bound variables f whose occurrences are labeled with an index i . As is standard, let introduces polymorphism.

Fig. 8 contains the type checking rules for our language. We overload *Eval* to work on strings; we ignore void and int. The JNI functions in δ are given types as before. We write $T(\delta)$ for the type of JNI function δ . Notice that unlike our implementation we do not treat occurrences of δ as polymorphic; instead in this language polymorphism can be achieved by writing wrappers around the calls to δ . In finding the free variables of the environment, we ignore the variables in $T(\delta)$, however.

We next give a small-step semantics. We define reduction contexts as usual:

$$\begin{aligned}
C &::= [] \mid \lambda x. C \mid C e \mid e C \\
&\quad \mid \delta(C, e, \dots, e) \mid \delta(e, C, e, \dots, e) \mid \dots
\end{aligned}$$

Fig. 9 gives our operational semantics. The rule β is standard. The remaining rules δ describe the behavior of JNI functions. For example, **FindClass** converts a string into a Java class object.

Next we prove preservation and progress for this language.

Lemma 1 (Substitution) *If $\Gamma[x \mapsto ct'] \vdash e : ct$ and $\Gamma \vdash e' : ct'$ then $\Gamma \vdash e[x \mapsto e'] : ct$.*

Proof: By induction on the structure of e .

Case $\delta(e_1, \dots, e_n)$. We have

$$\frac{\text{DELTA} \quad \Gamma[x \mapsto ct'] \vdash e_i : ct_i \quad T(\delta) = ct_1 \times \dots \times ct_n \rightarrow ct}{\Gamma[x \mapsto ct'] \vdash \delta(e_1, \dots, e_n) : ct}$$

By induction, $\Gamma \vdash e_i[x \mapsto e'] : ct_i$. But then by (DELTA) we have $\Gamma \vdash (\delta(e_1, \dots, e_n))[x \mapsto e'] : ct$.

Case let $f = e_1$ in e_2 .

$$\frac{\text{LET} \quad \Gamma[x \mapsto ct'] \vdash e_1 : ct_1 \quad \Gamma[x \mapsto ct', f \mapsto (ct_1, \vec{\alpha})] \vdash e_2 : ct_2 \quad fv(\Gamma[x \mapsto ct']) \subseteq \vec{\alpha}}{\Gamma[x \mapsto ct'] \vdash \text{let } f = e_1 \text{ in } e_2 : ct_2}$$

By induction, $\Gamma \vdash e_1[x \mapsto e'] : ct_1$. Also, observe that $fv(\Gamma) \subseteq fv(\Gamma[x \mapsto ct'])$ and thus $fv(\Gamma) \subseteq \vec{\alpha}$. Since we syntactically distinguish let- and lambda-bound variables, then $f[x \mapsto e'] = f$ always and therefore

$$\Gamma[x \mapsto ct', f \mapsto (ct_1, \vec{\alpha})] = \Gamma[f \mapsto (ct_1, \vec{\alpha}), x \mapsto ct']$$

Thus, by induction, $\Gamma[f \mapsto (ct_1, \vec{\alpha}), x \mapsto ct'] \vdash e_2 : ct_2$. Therefore by (Let), we have $\Gamma \vdash (\text{let } f = e_1 \text{ in } e_2)[x \mapsto e'] : ct_2$.

The other inductive cases are straightforward. □

Lemma 2 (Polymorphic Substitution) *If $\Gamma[f \mapsto (\mathbf{ct}', \vec{\alpha})] \vdash e : \mathbf{ct}$ and $\Gamma \vdash e' : \mathbf{ct}'$ and $fv(\Gamma) \subseteq \vec{\alpha}$ then $\Gamma \vdash e[x \mapsto e'] : \mathbf{ct}$.*

Proof: By induction on the structure of e . Most of the cases are straightforward.

Case g . If $g \neq f$ the result is trivial. Otherwise we have

$$\frac{\text{INST} \quad \Gamma(f) = (\mathbf{ct}, \vec{\alpha}) \quad \mathbf{ct} \preceq_i \mathbf{ct}' \quad \vec{\alpha} \preceq_i \vec{\alpha}}{\Gamma \vdash f_i : \mathbf{ct}'}$$

We have $f_i[f \mapsto e'] = e'$. Let S_i be the substitution defined by \preceq_i . Then we have $\Gamma \vdash e' : \mathbf{ct}'$ by assumption, and thus $S_i(\Gamma) \vdash e' : S_i(\mathbf{ct})$. But $fv(\Gamma) \subseteq \vec{\alpha}$, and $S_i(\vec{\alpha}) = \vec{\alpha}$. Further, by definition $S_i(\mathbf{ct}) = \mathbf{ct}'$. Thus we have $\Gamma \vdash f_i[f \mapsto e'] : \mathbf{ct}'$. \square

Lemma 3 (Preservation) *If $\Gamma \vdash e : \mathbf{ct}$ and $e \rightarrow e'$, then $\Gamma \vdash e' : \mathbf{ct}$.*

Proof: By induction on the structure of e . The inductive cases are straightforward. For (β) we use the substitution lemma, and for (let) we use the polymorphic substitution lemma.

Case $\delta(v_1, \dots, v_n)$. We proceed by case analysis on the reduction step.

If the reduction was $(\delta - FC)$, then we have

$$\frac{\text{DELTA} \quad \Gamma \vdash \text{“Class”} : \mathbf{str}\{\text{Class}\} \quad T(\delta) = \mathbf{str}\{\text{Class}\} \rightarrow \mathbf{ct} \quad \mathbf{ct} = \text{JStr}\{\text{“Class”}\} \text{JClass jobject}}{\Gamma \vdash \delta(\text{“Class”}) : \mathbf{ct}}$$

But since “Class” is a known value, we have $\text{JStr}\{\text{“Class”}\} = \text{Eval}(\text{Class})$. Thus $\mathbf{ct} = \text{Eval}(\text{Class}) \text{JClass jobject}$. But the reduction produces the value $\text{class}(\text{“Class”})$, which has the same type by rule (CLASS) .

If the reduction was $(\delta - FOC - 1)$, then we have

$$\frac{\text{DELTA} \quad \Gamma \vdash \text{object}(\text{“Class”}) : \text{Eval}(\text{Class}) \text{jobject} \quad T(\delta) = \{\nu; \phi\} \text{jobject} \rightarrow \mathbf{ct} \quad \mathbf{ct} = \{\nu; \phi\} \text{JClass jobject}}{\Gamma \vdash \delta(\text{object}(\text{“Class”})) : \mathbf{ct}}$$

Since $\text{object}(\text{“Class”})$ is a known value, we have $\{\nu; \phi\} = \text{Eval}(\text{Class})$. Therefore, $\mathbf{ct} = \text{Eval}(\text{Class}) \text{JClass jobject}$, which is the same type as the value $\text{class}(\text{“Class”})$, produced by $(\delta - FOC - 1)$, as given by the type rule (CLASS) .

If the reduction was $(\delta - FOC - 2)$, then we have

$$\frac{\text{DELTA} \quad \Gamma \vdash \text{class}(\text{“Class”}) : \text{Eval}(\text{Class}) \text{JClass jobject} \quad T(\delta) = \{\nu; \phi\} \text{jobject} \rightarrow \mathbf{ct} \quad \mathbf{ct} = \{\nu; \phi\} \text{JClass jobject}}{\Gamma \vdash \delta(\text{class}(\text{“Class”})) : \mathbf{ct}}$$

However, since $\{\nu; \phi\} = \text{Eval}(\text{Class}) \text{JClass}$, we must apply JAVACLASS . Therefore $\text{Class} = \text{“java.lang.Class”}$ and thus $\{\nu; \phi\} = \text{Eval}(\text{“java.lang.Class”})$. Therefore $\mathbf{ct} = \text{Eval}(\text{“java.lang.Class”}) \text{JClass jobject}$, which is exactly the type for $\text{class}(\text{“java.lang.Class”})$ produced by $(\delta - FOC - 2)$.

If the reduction was $(\delta - GFID)$, then we have

$$\frac{\text{DELTA} \quad \Gamma \vdash \text{class}(\text{“OClass”}) : \text{Eval}(\text{OClass}) \text{JClass jobject} \quad \Gamma \vdash \text{“Name”} : \mathbf{str}\{\text{Name}\} \quad \Gamma \vdash \text{“FClass”} : \mathbf{str}\{\text{FClass}\} \quad T(\delta) = \{\text{“OClass”}; \langle f \rangle \circ \phi\} \text{JClass jobject} \times \mathbf{str}\{\text{Name}\} \times \mathbf{str}\{\text{FClass}\} \rightarrow \mathbf{ct} \quad f = \text{“Name”} : \text{JStr}\{\text{“FClass”}\} \quad \mathbf{ct} = (f, \{\text{OClass}; \langle f \rangle \circ \phi\}) \text{jfieldID}}{\Gamma \vdash \delta(\text{class}(\text{“OClass”}), \text{“Name”}, \text{“FClass”}) : \mathbf{ct}}$$

By examining the type rules, we see that we can apply rule FID,

$$\text{FID} \frac{f = \text{"Name"} : (Eval(FClass)) \quad o = Eval(OCClass) \quad f \in o}{\Gamma \vdash fid(\text{"Name"}, \text{"FClass"}, \text{"OCClass"}) : (f, o) \text{jfieldID}}$$

Since the name of every object is unique (Eval is a 1-1 function) then $\{OCClass; \dots\}$ and $Eval(OCClass)$ must share the same fields and methods. Therefore since $\{OCClass; \langle f \rangle \circ \phi\}$, then $f \in Eval(OCClass)$. Thus the type given by (FID) is exactly the same type given above in DELTA.

If the reduction was $(\delta - GOF - 1)$, then we have

$$\text{DELTA} \frac{\Gamma \vdash object(\text{"Class"}) : Eval(Class) \quad \Gamma \vdash fid(\text{"Name"}, \text{"FClass"}, \text{"OCClass"}) : (f, o) \text{jfieldID} \quad \begin{array}{l} f = \text{"Name"} : (Eval(FClass)) \quad o = Eval(OCClass) \quad f \in o \\ T(\delta) = \{\text{"Class"}; \langle f \rangle \circ \phi\} \text{jobject} \times (f, \{\text{"Class"}; \langle f \rangle \circ \phi\}) \text{jfieldID} \rightarrow \text{ct} \\ \text{ct} = Eval(FClass) \text{jobject} \end{array}}{\Gamma \vdash \delta(object(\text{"Class"}), fid(\text{"Name"}, \text{"FClass"}, \text{"OCClass"})) : \text{ct}}$$

By examining the type rules, we see that the rule OBJ applies to the result of $(\delta - GOF - 1)$,

$$\text{OBJ} \frac{o = Eval(FClass)}{\Gamma \vdash object(\text{"FClass"}) : o \text{jobject}}$$

which is exactly the same type given to ct above in DELTA.

If the reduction was $(\delta - GOF - 2)$, then we have

$$\text{DELTA} \frac{\Gamma \vdash class(\text{"Class"}) : Eval(\text{"java.lang.Class"}) \text{jobject} \quad \Gamma \vdash fid(\text{"Name"}, \text{"FClass"}, \text{"java.lang.Class"}) : (f, o) \text{jfieldID} \quad \begin{array}{l} f = \text{"Name"} : (Eval(FClass)) \quad o = Eval(\text{"java.lang.Class"}) \quad f \in o \\ T(\delta) = \{OCClass; \langle f \rangle \circ \phi\} \text{jobject} \times (f, \{OCClass; \langle f \rangle \circ \phi\}) \text{jfieldID} \rightarrow \text{ct} \\ \text{ct} = Eval(FClass) \text{jobject} \quad OCClass = \text{"java.lang.Class"} \end{array}}{\Gamma \vdash \delta(class(\text{"Class"}), fid(\text{"Name"}, \text{"FClass"}, \text{"java.lang.Class"})) : \text{ct}}$$

By examining the type rules, we see that the rule OBJ applies to the result of $(\delta - GOF - 2)$,

$$\text{OBJ} \frac{o = Eval(FClass)}{\Gamma \vdash object(\text{"FClass"}) : o \text{jobject}}$$

which is exactly the same type given to ct above in DELTA. □

Lemma 4 (Progress) *If $\Gamma \vdash e : \text{ct}$, then either e is a value, or there exists an e' such that $e \rightarrow e'$.*

Proof: By induction on the structure of e . The inductive cases are straightforward. Therefore we present only those cases where expressions only contain values as sub-expressions.

Case $\delta(v_1, \dots, v_n)$. We proceed by case analysis on the JNI function from δ .

If we are checking `FindClass`, then we have

$$\text{DELTA} \frac{\Gamma \vdash v_1 : \text{str}\{Class\} \quad T(\text{FindClass}) = \text{str}\{Class\} \rightarrow \text{ct} \quad \text{ct} = \text{JTStr}\{\text{"Class"}\} \text{JClass} \text{jobject}}{\Gamma \vdash \text{FindClass}(v_1) : \text{ct}}$$

Then, by examining the type rules, the only value with type $\text{str}\{Class\}$ is the string “Class” by rule STRING. Therefore, $v_1 = \text{“Class”}$ and we can apply $(\delta - FC)$.

If we are checking `FindObjectClass`, then we have

$$\begin{array}{c} \text{DELTA} \\ \Gamma \vdash v_1 : \text{Eval}(Class) \text{ jobject} \\ \hline T(\text{FindObjectClass}) = \text{Eval}(Class) \text{ jobject} \rightarrow \text{ct} \quad \text{ct} = \text{Eval}(Class) \text{ JClass jobject} \\ \hline \Gamma \vdash \text{FindObjectClass}(v_1) : \text{ct} \end{array}$$

By examining the type rules, the only values with type $\text{Eval}(Class) \text{ jobject}$ are objects, $\text{object}(\text{“Class”})$ (via OBJ) and classes $\text{class}(\text{“Class”}) \text{ JClass}$ (via JAVACLASS). If $v_1 = \text{object}(\text{“Class”})$ then we can apply $(\delta - FOC - 1)$. Otherwise, $v_1 = \text{class}(\text{“Class”}) \text{ JClass}$ and we can apply $(\delta - FOC - 2)$.

If we are checking `GetFieldID`, then we have

$$\begin{array}{c} \text{DELTA} \\ \Gamma \vdash v_1 : o \text{ JClass jobject} \quad \Gamma \vdash v_2 : \text{str}\{Name\} \\ \Gamma \vdash v_3 : \text{str}\{FClass\} \quad T(\text{GetFieldID}) = (o \text{ JClass jobject} \times \text{str}\{Name\} \times \text{str}\{FClass\}) \rightarrow \text{ct} \\ \text{ct} = (f, o) \text{ jfieldID} \quad f = \text{“Name”} : \text{JTStr}\{FClass\} \quad o = \text{Eval}(OClass) \\ \hline \Gamma \vdash \text{GetFieldID}(v_1, v_2, v_3) : \text{ct} \end{array}$$

By examining the type rules, we see that by rule (CLASS), $v_1 = \text{class}(\text{“OClass”})$. Also, by rule (STRING), $v_2 = \text{“Name”}$ and $v_3 = \text{“FClass”}$. Since $\text{ct} = (f, o) \text{ jfieldID}$, then by rule (FID), we see that $f \in o$. Since class names are unique (Eval is a 1-1 function), it must also be the case that $\text{“Name”} : (\text{Eval}(FClass)) \in \text{Eval}(OClass)$. Therefore, we can apply $(\delta - GFID)$.

Finally, if we are checking `GetObjectField`, then we have

$$\begin{array}{c} \text{DELTA} \\ \Gamma \vdash v_1 : o \text{ jobject} \\ \Gamma \vdash v_2 : (f, o) \text{ jfieldID} \quad T(\text{GetObjectField}) = (o \text{ jobject} \times (f, o) \text{ jfieldID}) \rightarrow \text{ct} \\ \text{ct} = \text{Eval}(FClass) \quad f = \text{“Name”} : \text{Eval}(FClass) \quad o = \text{Eval}(OClass) \\ \hline \Gamma \vdash \text{GetObjectField}(v_1, v_2) : \text{ct} \end{array}$$

By examining the type rules, the only values with type $(f, o) \text{ jfieldID}$ are field identifiers. Therefore $v_2 = \text{fid}(\text{“Name”}, \text{“FClass”}, \text{“OClass”})$. Since $f \in o$ (by (FID)), then $\text{“Name”} : (\text{Eval}(FClass)) \in \text{Eval}(OClass)$ since class names are unique. By again looking at the type rules, the only values with type $o \text{ jobject}$ are objects, $\text{object}(\text{“Class”})$ (via OBJ) and classes $\text{class}(\text{“Class”}) \text{ JClass}$ (via JAVACLASS). If $v_1 = \text{object}(\text{“Class”})$ then $\Gamma \vdash v_1 : \text{Eval}(Class) \text{ jobject}$ and $Class = OClass$. Therefore, we can apply $(\delta - GOF - 1)$. Otherwise, $v_1 = \text{class}(\text{“Class”})$ and thus $OClass = \text{“java.lang.Class”}$ in which case we can apply $(\delta - GOF - 2)$.

Case $v_1 v_2$. The only type rule which applies to this form is (App), therefore since $\Gamma \vdash v_1 v_2 : \text{ct}$, there exists ct' such that $\Gamma \vdash v_1 : \text{ct}' \rightarrow \text{ct}$ and $\Gamma \vdash v_2 : \text{ct}'$. By examining the type rules again, we see that the only values with function types are lambda expressions. Therefore, v must be of the form $\lambda x.e$. Therefore we can apply the step (β) .

Case `let $f = v$ in e` . We can simply apply (let) in this case.

□

$$\begin{array}{c}
\text{STRING} \\
\hline
\Gamma \vdash \text{"Str"} : \text{str}\{\text{'Str'}\} \\
\\
\text{DELTA} \\
\frac{\Gamma \vdash e_i : \text{ct}_i \quad T(\delta) = \text{ct}_1 \times \dots \times \text{ct}_n \rightarrow \text{ct}}{\Gamma \vdash \delta(e_1, \dots, e_n) : \text{ct}} \\
\\
\text{VAR} \\
\hline
\Gamma \vdash x : \Gamma(x) \\
\\
\text{INST} \\
\frac{\Gamma(f) = (\text{ct}, \vec{\alpha}) \quad \text{ct} \preceq_i \text{ct}' \quad \vec{\alpha} \preceq_i \vec{\alpha}'}{\Gamma \vdash f_i : \text{ct}'} \\
\\
\text{LAM} \\
\frac{\Gamma[x \mapsto \text{ct}] \vdash e : \text{ct}'}{\Gamma \vdash \lambda x. e : \text{ct} \rightarrow \text{ct}'} \\
\\
\text{APP} \\
\frac{\Gamma \vdash e_1 : \text{ct} \rightarrow \text{ct}' \quad \Gamma \vdash e_2 : \text{ct}}{\Gamma \vdash e_1 e_2 : \text{ct}'} \\
\\
\text{LET} \\
\frac{\Gamma \vdash e_1 : \text{ct}_1 \quad \Gamma[f \mapsto (\text{ct}_1, \vec{\alpha})] \vdash e_2 : \text{ct}_2 \quad fv(\Gamma) \subseteq \vec{\alpha}}{\Gamma \vdash \text{let } f = e_1 \text{ in } e_2 : \text{ct}_2} \\
\\
\text{CLASS} \\
\frac{o = \text{Eval}(\text{Class})}{\Gamma \vdash \text{class}(\text{"Class"}) : o \text{ JClass jobject}} \\
\\
\text{JAVACLASS} \\
\frac{\Gamma \vdash e : o \text{ JClass jobject}}{\Gamma \vdash e : \{\text{"java.lang.Class"; ...}\} \text{ jobject}} \\
\\
\text{OBJ} \\
\frac{o = \text{Eval}(\text{Class})}{\Gamma \vdash \text{object}(\text{"Class"}) : o \text{ jobject}} \\
\\
\text{FID} \\
\frac{f = \text{"Name"} : (\text{Eval}(\text{FClass})) \quad o = \text{Eval}(\text{OClass}) \quad f \in o}{\Gamma \vdash \text{fid}(\text{"Name"}, \text{"FClass"}, \text{"OClass"}) : (f, o) \text{ jfieldID}}
\end{array}$$

Figure 8: Type Checking Rules

$$\begin{array}{l}
(\beta) \quad C[(\lambda x.e) v] \rightarrow C[e[x \mapsto v]] \\
(\text{let}) \quad C[\text{let } f = v \text{ in } e] \rightarrow C[e[f \mapsto v]] \\
(\delta - FC) \quad C[\text{FindClass}(\text{"Class"})] \rightarrow C[\text{class}(\text{"Class"})] \\
(\delta - FOC - 1) \quad C[\text{FindObjectClass}(\text{object}(\text{"Class"}))] \rightarrow C[\text{class}(\text{"Class"})] \\
(\delta - FOC - 2) \quad C[\text{FindObjectClass}(\text{class}(\text{"Class"}))] \rightarrow C[\text{class}(\text{"java.lang.Class"})] \\
(\delta - GFID) \quad C[\text{GetFieldID}(\text{class}(\text{"OClass"}), \text{"Name"}, \text{"FClass"})] \rightarrow \\
\quad C[\text{fid}(\text{"Name"}, \text{"FClass"}, \text{"OClass"})] \\
\quad \text{"Name"} : (\text{Eval}(\text{FClass})) \in \text{Eval}(\text{OClass}) \\
(\delta - GOF - 1) \quad C[\text{GetObjectField}(\text{object}(\text{"Class"}), \text{fid}(\text{"Name"}, \text{"FClass"}, \text{"OClass"}))] \rightarrow \\
\quad C[\text{object}(\text{"FClass"})] \\
\quad \text{Class} = \text{OClass} \wedge \text{"Name"} : (\text{Eval}(\text{FClass})) \in \text{Eval}(\text{OClass}) \\
(\delta - GOF - 2) \quad C[\text{GetObjectField}(\text{class}(\text{"Class"}), \text{fid}(\text{"Name"}, \text{"FClass"}, \text{"java.lang.Class"}))] \rightarrow \\
\quad C[\text{object}(\text{"FClass"})] \\
\quad \text{"Name"} : (\text{Eval}(\text{FClass})) \in \text{Eval}(\text{"java.lang.Class"})
\end{array}$$

Figure 9: Operational Semantics