Java Bytecode Verification for @NonNull Types

Chris Male, David J. Pearce, Alex Potanin, and Constantine Dymnikov

Victoria University of Wellington, NZ
{malechri,djp,alex,dymnikkost}@mcs.vuw.ac.nz

Abstract. Java's annotation mechanism allows us to extend its type system with non-null types. However, checking such types cannot be done using the existing bytecode verification algorithm. We extend this algorithm to verify non-null types using a novel technique that identifies aliasing relationships between local variables and stack locations in the JVM. We formalise this for a subset of Java Bytecode and report on experiences using our implementation.

1 Introduction

`NullPointerExceptions` are a common error arising in Java programs when references holding `null` are dereferenced. Java 1.5 allows us to annotate types and, hence, to extend the type system with `@NonNull` types. An important step in the enforcement of such types is the bytecode verifier which must efficiently determine whether or not non-null types are used soundly. The standard bytecode verifier uses a dataflow analysis which is insufficient for this task. To address this, we present a novel, lightweight dataflow analysis ideally suited to the problem of verifying non-null types.

Java Bytecodes have access to a fixed size local variable array and stack [19]. These act much like machine registers in that they have no fixed type associated with them; rather, they can have different types at different program points. To address this, the standard bytecode verifier automatically infers the types of local variables and stack locations at each point within the program. The following illustrates a simple program, and the inferred types that hold immediately before each instruction:

```
static int f(Integer);     locals    stack
 0:   aload_0              [Integer] []
 1:   ifnull 8             [Integer] [Integer]
 4:   aload_0              [Integer] []
 5:   invokevirtual ...    [Integer] [Integer]
 8:   return               [Integer] []
```

Here, there is one local variable at index 0. On method entry, this is initialised with the `Integer` parameter. The `aload_0` instruction loads the local variable at index 0 onto the stack, and the `Integer` type is inferred for that stack location as a result.

A bytecode verifier for non-null types must infer that the value loaded onto the stack immediately before the `invokevirtual` method call cannot be `null`, as this is the call's receiver. The challenge here is that `ifnull` compares the top of the stack against `null`, but then discards this value. Thus, the bytecode verifier must be aware that, at

L. Hendren (Ed.): CC 2008, LNCS 4959, pp. 229–244, 2008.

that exact moment, the top of the stack and local 0 are aliases. The algorithm used by the standard bytecode verifier is unable to do this. Therefore, we extend this algorithm to maintain information about such aliases, and we refer to this technique as *type aliasing*. More specifically, this paper makes the following contributions:

- We formalise our non-null bytecode verifier for a subset of Java Bytecode.
- We detail an implementation of our system for Java Bytecode.
- We report on our experiences with using our system on real-world programs.

While there has already been considerable work on non-null types (e.g. [25,10,16,3,8]), none has directly addressed the problem of bytecode verification. While these existing techniques could be used for this purpose, they operate on higher-level program representations and must first translate bytecode into their representation. This introduces unnecessary overhead that is undesirable for the (performance critical) bytecode verifier. Our technique operates on bytecode directly, thus eliminating this inefficiency.

2 Preliminaries

We extend Java types to allow references to be declared as non-null and for arrays to hold non-null elements (in §5 we extend this to Java Generics). For example:

```
Vector v1;
@NonNull Vector v2;
@NonNull Integer @NonNull [] a1;
```

Here, `v1` is a *nullable* reference (one which may be `null`), while `v2` is a non-null reference (one which may not be `null`); similarly, `a1` is a non-null reference to an array holding non-null elements. When annotating arrays, the leftmost annotation associates with the element type, whilst that just before the braces associates with the array reference type. We formalise a cut-down version of the non-null types supported by our system using the following grammar:

$$\begin{array}{lll} \alpha & ::= & \texttt{@NonNull} \mid \epsilon \\ T & ::= & T\ \alpha\ \texttt{[]} \mid \alpha\ C \mid \mathit{null} \mid \bot \end{array}$$

Here, the special *null* type is given to the `null` value, ϵ denotes the absence of a `@NonNull` annotation, C denotes a class name (e.g. `Integer`) and $\bot$ is given to locations which hold no value (e.g. they are uninitialised, in deadcode, etc).

An important question is how our system deals with subtyping. For example, we require all array element types be identical between subtypes[1]. A formal definition of the subtype relation for our simplified non-null type language is given in Figure 1. An important property of our subtype relation is that it forms a *complete lattice* (i.e. that every pair of types T_1, T_2 has a unique least upper bound, $T_1 \sqcup T_2$, and a unique greatest lower bound, $T_1 \sqcap T_2$). This helps ensure termination of our non-null verification algorithm.

[1] While this contrasts slightly with Java's treatment of arrays, we cannot do better without adding runtime non-null type information to arrays.

$$\frac{}{\texttt{@NonNull} \leq \epsilon} \qquad \frac{\alpha_1 \leq \alpha_2 \quad C \texttt{ extends } B}{\alpha_1\ C\ \leq \alpha_2\ B}$$

$$\frac{\alpha_1 \leq \alpha_2}{T_1\ \alpha_1\ \texttt{[]} \leq T_1\ \alpha_2\ \texttt{[]}} \qquad \frac{}{T_1\ \alpha\ \texttt{[]} \leq \alpha\ \texttt{java.lang.Object}}$$

$$\frac{}{\bot \leq T\ \alpha\ \texttt{[]}} \quad \frac{}{\bot \leq \alpha\ C} \quad \frac{}{\bot \leq \mathit{null}} \qquad \frac{}{\mathit{null} \leq T_1\ \texttt{[]}} \quad \frac{}{\mathit{null} \leq C}$$

Fig. 1. Subtyping rules for non-null Java types. We assume reflexivity and transitivity, that `java.lang.Object` is the root of the class hierarchy and, hence, is also $\top$.

A well-known problem, however, is that Java's subtype relation does not form a complete lattice [17]. This arises because two classes can share the same super-class and implement the same interfaces; thus, they may not have a unique least upper bound. To resolve this, we adopt the standard solution of ignoring interfaces entirely and, instead, treating interfaces as type `java.lang.Object`. This works because Java supports only single inheritance between classes. This is the approach taken in Sun's Java Bytecode verifier and, hence, our system is no less general than it.

3 Non-null Type Verification

Our non-null type verification algorithm infers the nullness of local variables at each point within a method. We assume method parameters, return types and fields are already annotated with `@NonNull`. Our algorithm is intraprocedural; that is, it concentrates on verifying each method in isolation, rather than the whole program together. The algorithm constructs an abstract representation of each method's execution; if this is possible, the method is type safe and cannot throw a `NullPointerException`. The abstract representation of a method mirrors the control-flow graph (CFG); its nodes contain an abstract representation of the program store, called an *abstract store*, giving the types of local variables and stack locations at that point.

We now formalise this construction process for methods. Constructors are ignored for simplicity and discussed informally in §5. Also, while the full Java Bytecode instruction set is supported, only a subset is considered here for brevity.

3.1 Abstract Store

In the Java Virtual Machine (JVM), each method has a fixed-size local variable array (for storing local variables) and a stack of known maximum depth (for storing temporary values). Our system models this using an abstract store, which we formalise as (Σ, Γ, κ), where Σ is the *abstract meta-heap*, Γ is the *abstract location array* and κ is the *stack pointer* which identifies the first free location on the stack. Here, Γ maps *abstract locations* to *type references*. These abstract locations are labelled $0, \ldots, n-1$, with the first m locations representing the local variable array, and the remainder representing the stack (hence, $n-m$ is the maximum stack size and $\kappa \leq n$). A type reference is a reference to a *type object* which, in turn, can be thought of as a non-null type with identity. Thus, we can have two distinct type objects representing the same non-null type. Crucially, this types-as-references approach allows two abstract locations to be

type aliases; that is, refer to the same type object. For example, in the following abstract store, locations 0 and 2 are type aliases:

$$\Sigma = \{r_1 \mapsto \texttt{@NonNull Integer}, r_2 \mapsto \texttt{String}\}, \Gamma = \{0 \mapsto r_1, 1 \mapsto r_2, 2 \mapsto r_1\}, \kappa = 3$$

Here, the abstract meta-heap, Σ, maps type references to non-null types. It's called a *meta-heap* as Σ does not abstract the program heap; rather it is an internal structure used only to enable type aliasing.

Definition 1. *An abstract store* (Σ, Γ, κ) *is well-formed iff* $dom(\Gamma) = \{0, \ldots, n-1\}$ *for some* n, $ran(\Gamma) \subseteq dom(\Sigma)$ *and* $0 \leq \kappa \leq n$.

3.2 Abstract Semantics

The effect of a bytecode instruction is given by its *abstract semantics*, which we describe using transition rules. These summarise the abstract store immediately after the instruction in terms of the abstract store immediately before it; any necessary constraints on the abstract store immediately before the instruction are also identified.

The abstract semantics for the bytecode instructions considered in our formalism are given in Figure 2. Here, $\Gamma[r_1/r_2]$ generates an abstract store from Γ where all abstract locations holding r_1 now hold r_2. Several helper functions are used: **fieldT**(`O`, `N`), returns the type of field `N` in class `O`; **methodT**(`O`, `M`) returns the type of method `M` in class `O`; **thisMethT**() gives the current method's type; finally, **validNewT**(T_1) holds if $T_1 \neq$ `@NonNull` T_2 α `[]` for any T_2. The latter prevents creation of arrays holding `@NonNull` elements, as Java always initialises array elements with `null` (see §5).

A useful illustration of our abstract semantics is the `arrayload` bytecode. This requires the array index on top of the stack, followed by the array reference itself; these are popped off the stack and the indexed element is loaded back on. Looking at the `arrayload` rule, we see κ decreases by one, indicating the net effect is one less element on the stack. The notation $\Gamma[\kappa - 2 \mapsto r]$ indicates the abstract store is updated so that abstract location $\kappa - 2$ now holds type reference r; thus, r has been pushed onto the stack and represents the loaded array element. The reference on top of the stack is ignored since this represents the actual index value, and is of no concern. The constraint $r \notin \Sigma$ ensures r references a *fresh* type object; such constraints are used to ensure an abstract location is not type aliased with any other. Another constraint ensures the array reference is non-null, thus protecting against a `NullPointerException`.

Considering the remaining rules from Figure 2, the main interest lies with `ifceq`. There is one rule for each of the true/false branches. The true branch uses the greatest lower bound operator, $T_1 \sqcap T_2$ (recall §2). This creates a single type object which is substituted for both operands to create a type aliasing relationship. For the false branch, a special *difference* operator, $T_1 - T_2$, is employed which is similar to set difference. For example, the set of possible values for a variable `o` of type `Object` includes all instances of `Object` (and its subtypes), as well as `null`; after a comparison `o!=null`, `null` is removed from this set. Thus, it is defined as follows:

$$\frac{}{\texttt{store}\ i : \Sigma, \Gamma, \kappa \longrightarrow \Sigma, \Gamma[i \mapsto \Gamma(\kappa-1)], \kappa-1} \qquad \frac{}{\texttt{load}\ i : \Sigma, \Gamma, \kappa \longrightarrow \Sigma, \Gamma[\kappa \mapsto \Gamma(i)], \kappa+1}$$

$$\frac{r \notin \Sigma \quad \Sigma' = \Sigma \cup \{r \mapsto \mathit{null}\}}{\texttt{loadnull} : \Sigma, \Gamma, \kappa \longrightarrow \Sigma', \Gamma[\kappa \mapsto r], \kappa+1} \qquad \frac{\begin{array}{c}\text{validNewT}(T) \\ r \notin \Sigma \quad \Sigma' = \Sigma \cup \{r \mapsto \texttt{@NonNull}\ T\}\end{array}}{\texttt{new}\ T : \Sigma, \Gamma, \kappa \longrightarrow \Sigma', \Gamma[\kappa \mapsto r], \kappa+1}$$

$$\frac{\begin{array}{c}\Sigma(\Gamma(\kappa-2)) = T\ \texttt{@NonNull}\ \texttt{[]} \\ r \notin \Sigma \quad \Sigma' = \Sigma \cup \{r \mapsto T\}\end{array}}{\texttt{arrayload} : \Sigma, \Gamma, \kappa \longrightarrow \Sigma', \Gamma[\kappa-2 \mapsto r], \kappa-1} \qquad \frac{\begin{array}{c}\Sigma(\Gamma(\kappa-1)) = T_1 \quad T_1 \leq T_2 \\ \Sigma(\Gamma(\kappa-3)) = T_2\ \texttt{@NonNull}\ \texttt{[]}\end{array}}{\texttt{arraystore} : \Sigma, \Gamma, \kappa \longrightarrow \Sigma, \Gamma, \kappa-3}$$

$$\frac{\begin{array}{c}\Sigma(\Gamma(\kappa-1)) = \texttt{@NonNull}\ C \\ T = \text{fieldT}(\texttt{O}, \texttt{N}) \\ r \notin \Sigma \quad \Sigma' = \Sigma \cup \{r \mapsto T\}\end{array}}{\texttt{getfield O.N} : \Sigma, \Gamma, \kappa \longrightarrow \Sigma', \Gamma[\kappa-1 \mapsto r], \kappa} \qquad \frac{\begin{array}{c}\Sigma(\Gamma(\kappa-1)) = T_1 \\ \Sigma(\Gamma(\kappa-2)) = \texttt{@NonNull}\ C \\ T_2 = \text{fieldT}(\texttt{O}, \texttt{N}) \quad T_1 \leq T_2\end{array}}{\texttt{putfield O.N} : \Sigma, \Gamma, \kappa \longrightarrow \Sigma, \Gamma, \kappa-2}$$

$$\frac{\begin{array}{c}(P_1, \ldots, P_n) \rightarrow T_r = \text{methodT}(\texttt{O}, \texttt{M}) \\ \Sigma(\Gamma(\kappa-n)), \ldots, \Sigma(\Gamma(\kappa-1)) = T_1, \ldots, T_n \\ \Sigma(\Gamma(\kappa-(n+1))) = \texttt{@NonNull}\ C \\ T_1 \leq P_1, \ldots, T_n \leq P_n \\ r \notin \Sigma \quad \Sigma' = \Sigma \cup \{r \mapsto T_r\} \quad \kappa' = \kappa - n\end{array}}{\texttt{invoke O.M} : \Sigma, \Gamma, \kappa \longrightarrow \Sigma', \Gamma[\kappa'-1 \mapsto r], \kappa'} \qquad \frac{\begin{array}{c}(P_1, \ldots, P_n) \rightarrow T_r = \text{thisMethT}() \\ \Sigma(\Gamma(\kappa-1)) = T \quad T \leq T_r\end{array}}{\texttt{return} : \Sigma, \Gamma, \kappa \longrightarrow \emptyset, \emptyset, 0}$$

$$\frac{\begin{array}{c}r_1 = \Gamma(\kappa-2) \quad r_2 = \Gamma(\kappa-1) \\ \Sigma(r_1) = T_1 \quad \Sigma(r_2) = T_2 \quad r_3 \notin \Sigma \\ \Sigma' = \Sigma \cup \{r_3 \mapsto T_1 \sqcap T_2\} \quad \kappa' = \kappa - 2\end{array}}{\texttt{ifceq} : \Sigma, \Gamma, \kappa \xrightarrow{\mathit{true}} \Sigma', \Gamma[r_1/r_3, r_2/r_3], \kappa'} \qquad \frac{\begin{array}{c}r_1 = \Gamma(\kappa-2) \quad r_2 = \Gamma(\kappa-1) \\ \Sigma(r_1) = T_1 \quad \Sigma(r_2) = T_2 \quad r_3, r_4 \notin \Sigma \\ \Sigma' = \Sigma \cup \{r_3 \mapsto T_1 - T_2, r_4 \mapsto T_2 - T_1\}\end{array}}{\texttt{ifceq} : \Sigma, \Gamma, \kappa \xrightarrow{\mathit{false}} \Sigma', \Gamma[r_1/r_3, r_2/r_4], \kappa-2}$$

Fig. 2. Abstract semantics for Java Bytecodes considered. Note, `ifceq` stands for `if_cmpeq`.

Definition 2. *$T_1 - T_2$ is* `@NonNull` *T, if $T_1 = \alpha\ T \wedge T_2 = \mathit{null}$, and T_1 otherwise.*

The semantics for the `return` bytecode indicate that: firstly, we always expect a return value (for simplicity); and, secondly, no bytecode can follow it in the CFG.

Finally, the Java Bytecodes not considered in Figure 2 include all arithmetic operations (e.g. `iadd`, `imul`, etc), stack manipulators (e.g. `pop`, `dup`, etc), other branching primitives (e.g. `ifnonull`, `tableswitch`, etc), synchronisation primitives (e.g. `monitorenter`, etc) and other miscellaneous ones (e.g. `instanceof`, `checkcast`, `athrow` and `arraylength`). It is easy enough to see how our abstract semantics extends to these and our implementation (see §5) supports them all.

3.3 An Example

Figure 3 illustrates the bytecode instructions for a simple method and its corresponding abstract representation. When a method is called, the local variable array is initialised with the values of the incoming parameters, starting from 0 and using as many as necessary; for instance methods, the first parameter is always the `this` reference. Thus, the first abstract location of the first store in Figure 3 has type `Test`; the remainder have

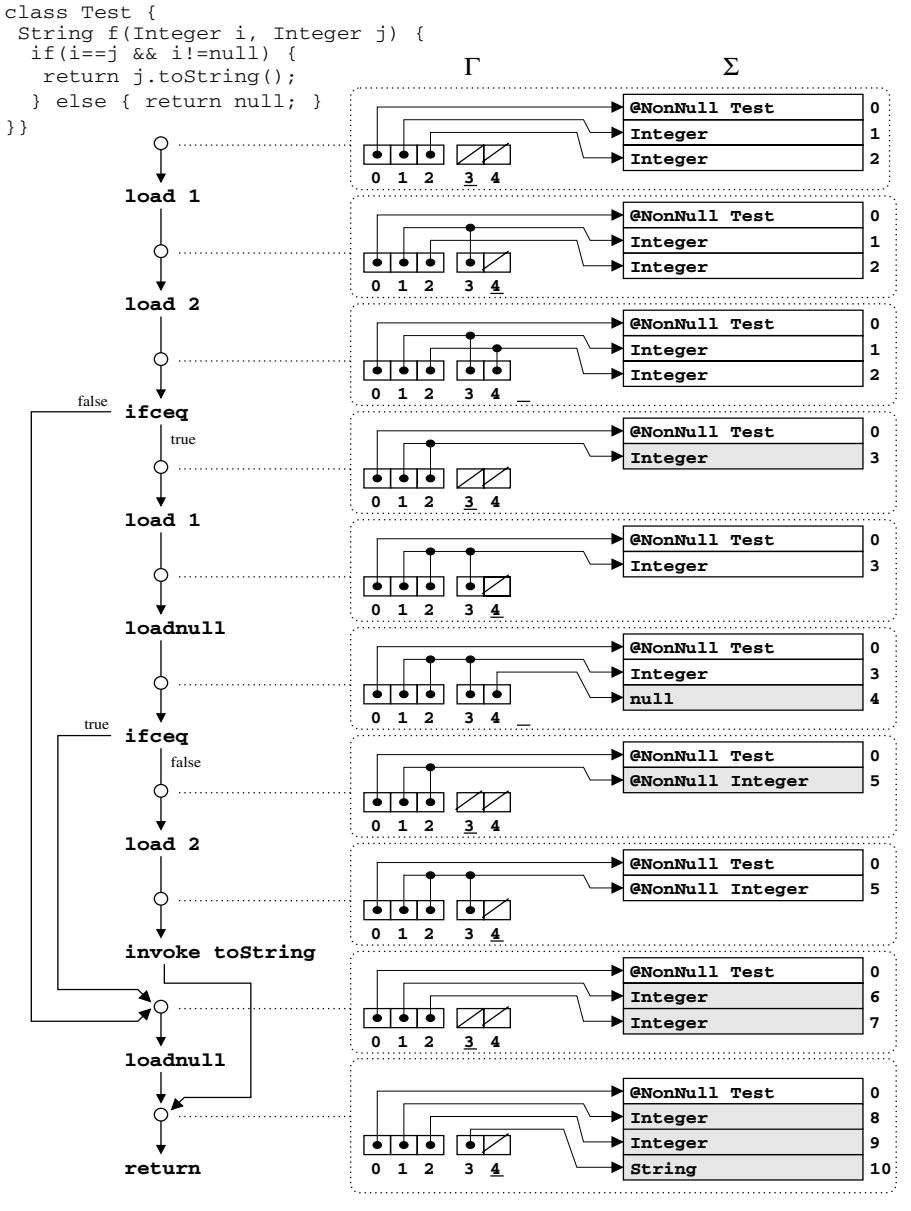

Fig. 3. Bytecode representation of a simple Java Method (source given above) and the state of the abstract store, (Σ, Γ, κ), going into each instruction. The value of κ is indicated by the underlined abstract location; when the stack is full, this points past the last location. The type objects in Σ are given a unique identifier to help distinguish new objects from old ones; we assume unreferenced type objects are immediately garbage collected, which is reflected in the identifiers becoming non-contiguous. Type aliases are indicated by references which are "joined". For example, the second abstract store reflects the state immediately after the `load 1` instruction, where locations 1 and 3 are type aliases.

nullable type `Integer`, with each referring to a unique type object (since we must conservatively assume parameters are not aliased on entry).

In Figure 3, the effect of each instruction is reflected in the changes between the abstract stores before and after it. Of note are the two `ifceq` instructions: the first establishes a type aliasing relationship between locations 1 and 2 (on the true branch); the second causes a retyping of location 1 to `@NonNull Integer` (on the false branch) which also retypes location 2 through type aliasing. Thus, at the `invoke` instruction, the top of the stack (which represents the receiver reference) holds `@NonNull Integer`, indicating it will not throw a `NullPointerException`.

We now consider what happens at join points in the CFG. The `return` instruction in Figure 3 is a good illustration, since two distinct paths reach it and each has its own abstract store. These must be combined to summarise all possible program stores at that point. In Figure 3, the store coming out of the `invoke` instruction has a type aliasing relationship, whereas that coming out of the `loadnull` instruction does not; also, in the former, location 2 has type `@NonNull Integer`, whilst the latter gives it nullable type `Integer`. This information must be combined conservatively. Since location 2 can hold `null` on at least one incoming path, it can clearly hold `null` at the join point. Hence, the least conservative type for location 2 is `Integer`. Likewise, if a type alias relationship does not hold on all incoming paths, we cannot assume it holds at the join. We formalise this notion of conservatism as a subtype relation:

Definition 3. *Let* $S_1 = (\Sigma_1, \Gamma_1, \kappa), S_2 = (\Sigma_2, \Gamma_2, \kappa)$ *be well-formed abstract stores. Then* $S_1 \leq S_2$ *iff* $\forall x, y \in \{0 \ldots \kappa\}\big[\Sigma_1(\Gamma_1(x)) \leq \Sigma_2(\Gamma_2(x)) \wedge (\Gamma_2(x) = \Gamma_2(y) \implies \Gamma_1(x) = \Gamma_1(y))\big]$.

Note, Definition 3 requires κ be identical on each incoming store; this reflects a standard requirement of Java Bytecode. Now, to construct the abstract store at a join point, our verification system finds the least upper bound, $\sqcup$, of incoming abstract stores — this is the least conservative information obtainable. We formalise this as follows:

Definition 4. *Let* $G = (V, E)$ *be the control-flow graph for a method* M*. Then, the dataflow equations for* M *are given by* $S_M(y) = \bigsqcup_{x \xrightarrow{l} y \in E} f(I(x), S_M(x), l)$.

Here, the *transfer function*, f, is defined by the abstract semantics of Figure 2, $I(x)$ gives the bytecode at node x, and the edge label, l, distinguishes the true/false branches for `ifceq`. Thus, $S_M(y)$ gives the abstract store going into y. Finally, the dataflow equations can be solved as usual by iterating to a fixed point using a *worklist algorithm*.

4 Soundness

We now demonstrate that our algorithm *terminates* and is *correct*; that is, if a method passes our verification process, then it cannot throw a `NullPointerException`.

Several previous works have formalised Java Bytecode and shown the standard verification algorithm is correct (e.g. [14,17]). Our system essentially operates in an identical fashion to the standard verifier, except that it additionally maintains type aliases and propagates `@NonNull` annotations. Indeed, our abstract semantics of Figure 2 would be identical to previous work (e.g. [17]) if we removed the requirement for `@NonNull`

types at dereference sites and prohibited type aliasing relationships. Thus, we leverage upon these existing works to simplify our proof by restricting attention to those details particular to our system.

An important issue regarding our formalism is that it applies only to *methods*, not *constructors*. The reason for this is detailed in §5. Therefore, in the following, we assume all fields annotated with `@NonNull` are correctly initialised.

4.1 Termination

Demonstrating termination amounts to showing the dataflow equations always have a *least fixed-point*. This requires the transfer function, f, is monotonic and that our subtyping relation is a *join-semilattice* (i.e. any two abstract stores always have a unique least upper bound). These are addressed by Lemmas 1 and 2.

Strictly speaking, Definition 3 does not define a join-semilattice over abstract stores, since two stores may not have a unique least upper bound. For example, consider:

$$S_1 = (\{r_1 \mapsto \texttt{Integer}, r_2 \mapsto \texttt{Float}\}, \{0 \mapsto r_1, 1 \mapsto r_1, 2 \mapsto r_2\}, 3)$$
$$S_2 = (\{r_1 \mapsto \texttt{Integer}, r_2 \mapsto \texttt{Float}\}, \{0 \mapsto r_2, 1 \mapsto r_2, 2 \mapsto r_1\}, 3)$$

Then, the following are minimal upper bounds of S_1 and S_2:

$$S_3 = (\{r_1 \mapsto \texttt{Number}, r_2 \mapsto \texttt{Number}\}, \{0 \mapsto r_1, 1 \mapsto r_1, 2 \mapsto r_2\}, 3)$$
$$S_4 = (\{r_1 \mapsto \texttt{Number}, r_2 \mapsto \texttt{Number}\}, \{0 \mapsto r_2, 1 \mapsto r_2, 2 \mapsto r_1\}, 3)$$

Here, $S_3 \leq S_4$, $S_4 \leq S_3$, $\{S_1, S_2\} \leq \{S_3, S_4\}$ and $\neg\exists S.[\{S_1, S_2\} \leq S \leq \{S_3, S_4\}]$. Hence, there is no unique least upper bound of S_1 and S_2. Such situations arise in our implementation as type objects are Java Objects and, hence, $r_1 \neq r_2$ simply means different object addresses. Now, while S_3 and S_4 are distinct, they are also *equivalent*:

Definition 5. *Let $S_1 = (\Sigma_1, \Gamma_1, \kappa), S_2 = (\Sigma_2, \Gamma_2, \kappa)$, then S_1 and S_2 are equivalent, written $S_1 \equiv S_2$, iff $S_1 \leq S_2$ and $S_1 \geq S_2$.*

Lemma 1. *Let $S_1 = (\Sigma_1, \Gamma_1, \kappa), S_2 = (\Sigma_2, \Gamma_2, \kappa)$ with dom(Γ_1) = dom(Γ_2). If U is the set of minimal upper bounds of S_1 and S_2, then $U \neq \emptyset$ and $\forall x, y \in U.[x \equiv y]$.*

Proof. See companion Technical Report [20].

Lemma 2. *The dataflow equations from Definition 4 are monotonic.*

Proof. By case analysis on the instructions of Figure 2. See companion Technical Report [20].

4.2 Correctness

We now show the type aliasing information maintained is correct (Lemma 3), and that any location with `@NonNull` type cannot hold `null` (Lemma 4). This yields an overall correctness result for the subset of Java Bytecode we have formalised (Theorem 1).

Definition 6. *A Java method is considered to be valid if it passes the standard JVM verification process [19].*

The consequences of Definition 6 include: all conventional types (i.e. ignoring non-null types) are used safely; stack sizes are always the same at the meet points; method and field lookups always resolve; etc.

Lemma 3. *Let $S_M = (\Sigma, \Gamma, \kappa)$ be the abstract store for an instruction in a valid method M. If $\{l_1 \mapsto r, l_2 \mapsto r\} \subseteq \Gamma$, then the local array/stack locations represented by l_1, l_2 refer to the same object or array immediately before that instruction in any execution trace of M.*

Proof. By case analysis on the different instruction types of Figure 2 and the notion of conservatism from Definition 3. See companion Technical Report [20]. □

Lemma 4. *Let $S_M = (\Sigma, \Gamma, \kappa)$ be the abstract store for an instruction in a valid method M. Assume the parameters of M, the fields accessed by M and the return value of all methods invoked by M respect their declared non-null type. Then, if $\{l \mapsto r\} \subseteq \Gamma \wedge \{r \mapsto$ `@NonNull` $T\} \subseteq \Sigma$, the local array/stack location represented by l does not hold* `null` *immediately before that instruction in any execution trace of M.*

Proof. Again, by case analysis on the different instruction types of Figure 2, the notion of conservatism from Definition 3 and Lemma 3. See companion Technical Report [20]. □

Theorem 1. *If our abstract representation can be correctly constructed for all methods in a Java Bytecode program, then no method will throw a* `NullPointerException`*, assuming all fields are correctly initialised.*

Proof. By induction on the call sequence, starting from `main(String[])`. Using Lemma 4, we formulate an inductive hypothesis stating, for a method M, that if the arguments to M respect their non-null types, so do the return value of M, the arguments to any calls made by M, and any assignments to fields / array elements made by M. See companion Technical Report [20]. □

5 Implementation

We have implemented our system on top of Java Bytecode and we now discuss many aspects not covered by our discussion so far.

Constructors. In Java, a field is assigned `null` before it is initialised in a constructor [10]. Thus, a field with non-null type will temporarily hold `null` inside a constructor. Figure 4 highlights the problem. We must ensure such fields are properly initialised, and must restrict access prior to this occurring. Two mechanisms are used to do this:

1. A simple dataflow analysis is used to ensure that all non-null (instance) fields in a class declaration are initialised by that class's constructor.
2. Following [10], we use a secondary type annotation, `@Raw`, for references to indicate the object referred to may not be initialised. Reads from fields through these return nullable types. The `this` reference in a constructor is implicitly typed `@Raw` and `@Raw` is strictly a supertype of a normal reference.

Inheritance. When a method overrides another via inheritance our tool checks that `@NonNull` types are properly preserved. As usual, types in the parameter position are *contravariant* with inheritance, whilst those in the return position are *covariant*.

```
class Parent {
 Parent() { doBadStuff(); } // error #1, f1 not initialised yet!
 int doBadStuff() { return 0; }
}
class Child extends Parent {
 @NonNull String f1; @NonNull String f2;
 Child() {
  doBadStuff();    // error #2, f1 not initialised before call!
  f1 = "Hello World";
 }                 // error #3, f2 not initialised yet!
 int doBadStuff() { return f1.length(); }
}}
```

Fig. 4. Illustrating three distinct problems with constructors and default values. Error #3 arises as all `@NonNull` fields must be initialised! Error #2 arises as a method is called on `this` before all `@NonNull` fields are initialised. Error #1 arises as, when the `Child`'s constructor is called, it calls the `Parent`'s constructor. This, in turn, calls `doBadStuff()` which dynamically dispatches to the `Child`'s implementation. However, field `f1` has not yet been initialised!

Field Retyping. Consider this method and its bytecode (recall local 0 holds `this`):

```
class Test {                0.  load 0
 Integer field;             2.  getfield Test.field
 void f() {                 5.  ifnull 16
  if(field != null) {       8.  load 0
   field.toString()         10. getfield Test.field
}}}                         13. invoke Integer.toString
                            16. return
```

The above is not type safe in our system as the non-nullness of the field is lost when it is reloaded. This is strictly correct, since the field's value may have been changed between loads (e.g. by another thread). We require this is resolved manually by adjusting the source to first store the field in a local variable (which is strictly thread local).

Generics. Our implementation supports Java Generics. For example, we denote a `Vector` containing non-null `Strings` with `Vector<@NonNull String>`. Extending the subtype relation of Figure 1 is straightforward and follows the conventions of Java Generics (i.e. prohibiting variance on generic parameters). Verifying methods which accept generic parameters is more challenging. To deal with this, we introduce a special type, $\top_i$, for each (distinct) generic type used in the method; here, $\top_i \leq$ `java.lang.Object` and $\top_i \not\leq \top_j$, for $i \neq j$. When checking a method `f(T x)`, the abstract location representing `x` is initialised to the type $\top_i$ used exclusively for representing the generic type `T`. The subtyping constraints ensure $\top_i$ can only flow into variables/return types declared with the same generic type `T`. However, an interesting problem arises with some existing library classes. For example:

```
class Hashtable<K,V> ... { ...
 V get(K key) { ...; return null; } }
```

Clearly, this class assumes `null` is a subtype of every type; unfortunately, this is not true in our case, since e.g. `null` $\not\leq$ `@NonNull String`. To resolve this, we prohibit instances of `Hashtable`/`HashMap` from having a non-null type in `V`'s position.

Casting + Arrays. We explicitly prevent the creation of arrays with non-null elements (e.g. `new @NonNull Integer[10]`), as Java always initialises array elements of reference type with `null`. Instead, we require an explicit *cast* to `@NonNull Integer[]` when the programmer knows the array has been fully initialised. Casts from nullable to non-null types are implemented as runtime checks which fail by throwing `ClassCastException`s. Their use weakens Theorem 1, since we are essentially trading `NullPointerException`s for `ClassCastException`s. While this is undesirable, it is analogous to the issue of downcasts in Object-Oriented Languages.

Instanceof. Our implementation extends the type aliasing technique to support retyping via `instanceof`. For example:

```
if(x instanceof String) { String y = (String) x; .. }
```

Here, our system retypes `x` to type `@NonNull String` on the true branch, rending the cast redundant (note, an `instanceof` test never passes on `null`).

Type Annotations. The Java Classfile format doesn't allow annotations on generic parameters or in the array type reference position. Therefore, we use a simple mechanism for encoding this information into a classfile. We expect future versions of Java will support such types directly and, indeed, work is already underway in this regard [9].

6 Case Studies

We have manually annotated and checked several real-world programs using our non-null type verifier. The largest practical hurdle was annotating Java's standard libraries. This task is enormous and we are far from completion. Indeed, finishing it by hand does not seem feasible; instead, we plan to develop (semi-)automatic procedures to help.

We now consider four real-world code bases which we have successfully annotated: the `java/lang` and `java/io` packages, the `jakarta-oro` text processing library and `javacc`, a well-known parser generator. Table 1 details these. Table 2 gives a breakdown of the annotations added, and the modifications needed for the program to type check. The most frequent modification, "Field Load Fix", was for the field retyping issue identified in §5. To resolve this, we manually added a local variable into which the field was loaded before the null check. Many of these fixes may represent real concurrency bugs, although a deeper analysis of each situation is needed to ascertain this. The next most common modification, "Context Fixes", were for situations where the programmer knew a reference could not hold `null`, but our system was unable to determine this. These were resolved by adding dummy null checks. Examples include:

Table 1. Details of our four benchmarks. Note, `java/lang` does not include subpackages.

benchmark	version	LOC	source
`java/lang` package	1.5.0	14K	`java.sun.com`
`java/io` package	1.5.0	10.6K	`java.sun.com`
`jakarta-oro`	2.0.8	8K	`jakarta.apache.org/oro`
`javacc`	3.2	28K	`javacc.dev.java.net`

Table 2. Breakdown of annotations added and related metrics. "Annotated Types" gives the total number of annotated parameter, return and field types against the total number of reference / array types in those positions. A breakdown according to position (i.e. parameter, return type or field) is also given. "Field Load Fixes" counts occurrences of the field retyping problem outlined in §5. "Context Fixes" counts the number of dummy null checks which had to be added. "Required Null Checks" counts the number of required null checks, versus the total number of dereference sites. Finally, "Required Casts" counts the number of required casts, versus the total number of casts.

	Annotated Types	Parameter Annotations	Return Annotations	Field Annotations
`java/lang`	931 / 1599	363 / 748	327 / 513	241 / 338
`java/io`	515 / 1056	322 / 672	96 / 200	97 / 184
`jakarta-oro`	413 / 539	273 / 320	85 / 108	55 / 111
`javacc`	420 / 576	199 / 278	53 / 65	168 / 233

	Field Load Fixes	Context Fixes	Other Fixes	Required Null Checks	Required Casts
`java/lang`	65	61	36	281 / 2550	51 / 96
`java/io`	59	82	21	207 / 2254	54 / 110
`jakarta-oro`	53	327	29	73 / 2014	29 / 33
`javacc`	109	137 (28)	74	287 / 5700	141 / 431

- `Thread.getThreadGroup()` returns `null` when the thread in question has stopped. But, `Thread.currentThread().getThreadGroup()` will return a non-null value, since the current thread cannot complete `getThreadGroup()` if it has stopped! This assumption was encountered in several places.
- Another difficult situation for our tool is when the nullness of a method's return value depends either on its parameters, or on the object's state. A typical example is illustrated in Figure 5. More complex scenarios were also encountered where, for example, an array was known to hold non-null values up to a given index.
- As outlined in §5, `Hashtable.get(K)` returns `null` if no item exists for the key. A programmer may know that, for specific keys, `get()` cannot return `null` and so can avoid unnecessary `null` check(s). The `javacc` benchmark used many `hashtables` and many context fixes were needed as a result. In Table 2, the number of "Context Fixes" for this particular problem are shown in brackets.

The "Other Fixes" category in Table 2 covers other miscellaneous modifications needed for the code to check. Figure 6 illustrates one such example. Most relate to the initialisation of fields. In particular, helper methods called from constructors which initialise fields are a problem. This is because our system checks each constructor initialises its fields, but does not account for those initialised in helper methods. To resolve this, we either inlined helper methods or initialised fields with dummy values before they were called.

The "Required Null Checks" counts the number of explicit null checks (as present in the original program's source), against the total number of dereference sites. Since, in the normal case, the JVM must check every dereference site, this ratio indicates the potential for speedup resulting from non-null types. Likewise, "Required Casts" counts

```
public void actionPerformed(@NonNull ActionEvent ae) { ...
  JFileChooser jfc = new JFileChooser(); ...
  int rval = jfc.showOpenDialog(null);
  if(rval == JFileChooser.APPROVE_OPTION) {
    File f = jfc.getSelectedFile();
    filePath.setText(f.getCanonicalPath());
  ...
```

Fig. 5. A common scenario where the nullness of a method's return type depends upon its context; in this case, if `rval==APPROVE_OPTION`, then `getSelectedFile()` won't return `null`. To resolve this, we must add a "dummy" check that `f!=null` before the method call.

```
public ThreadGroup(String name) {
  this(Thread.currentThread().getThreadGroup(), name);
  ...
```

Fig. 6. An interesting example from `java.lang.ThreadGroup`. The constructor invoked via the `this` call requires a non-null argument (and this is part of its Javadoc specification). Although `getThreadGroup()` can return `null`, it cannot here (as discussed previously). Our tool reports an error for this which cannot be resolved by inserting a dummy `null` check, since the `this` call must be the first statement of the constructor. Therefore, we either inline the constructor being called, or construct a helper method which can accept a `null` parameter.

the number of casts actually required, versus the total number present (recall from §5 that our tool automatically retypes local variables after `instanceof` tests, making numerous casts redundant.)

We were also interested in whether or not our system could help documentation. In fact, it turns out that of the 1101 public methods in `java/lang`, 83 were misdocumented. That is, the Javadoc failed to specify that a parameter must not be `null` when, according to our system, it needed to be. We believe this is actually pretty good, all things considered, and reflects the quality of documentation for `java/lang`. Interestingly, many of the problem cases were found in `java/lang/String`.

Finally, a comment regarding performance seems prudent, since we have elided performance results for brevity. In fact, the performance of our system is very competitive with the standard bytecode verifier. This is not surprising, since our system uses a very similar algorithm to the standard bytecode verifier, albeit extended with type aliasing.

7 Related Work

Several works have considered the problem of checking non-null types. Fähndrich and Leino investigated the constructor problem (see §5) and outlined a solution using raw types [10]. However, no mechanism for actually checking non-null types was presented. The FindBugs tool checks `@NonNull` annotations using a dataflow analysis that accounts for comparisons against `null` [16,15]. Their approach does not employ type aliasing and provides no guarantee that all potential errors will be reported. While this is reasonable for a lightweight software quality tool, it is not suitable for bytecode

verification. ESC/Java also checks non-null types and accounts for the effect of conditionals [11]. The tool supports type aliasing (to some extent), can check very subtle pieces of code and is strictly more precise than our system. However, it relies upon a theorem prover which employs numerous transformations and optimisations on the intermediate representation, as well as a complex back-tracking search procedure. This makes it rather unsuitable for bytecode verification, where efficiency is paramount.

Ekman *et al.* implemented a non-null checker within the JustAdd compiler [8]. This accounts for the effect of conditionals, but does not consider type aliasing as there is little need in their setting where a full AST is available. To apply their technique to Java Bytecode would require first reconstructing the AST to eliminate type aliasing between stack and local variable locations. This would add additional overhead to the bytecode verification process, compared to our more streamlined approach. Pominville *et al.* also discuss a non-null analysis that accounts for conditionals, but again does not consider type aliasing [25]. They present empirical data suggesting many internal null checks can be eliminated, and that this leads to a useful improvement in program performance.

Chalin *et al.* empirically studied the ratio of parameter, return and field declarations which are intended to be non-null, concluding that $2/3$ are [3]. To do this, they manually annotated existing code bases, and checked for correctness by testing and with ESC/-Java. JavaCOP provides an expressive language for writing type system extensions, such as non-null types [2]. This system cannot account for the effects of conditionals; however, as a work around, the tool allows assignment from a nullable variable `x` to a non-null variable if this is the first statement after a `x!=null` conditional.

CQual is a flow-sensitive qualifier inference algorithm which supports numerous type qualifiers, but does not account for conditionals at all [12,13]. Building on this is the work of Chin *et al.* which also supports numerous qualifiers, including `nonzero`, `unique` and `nonnull` [5,6]. Again, conditionals cannot be accounted for, which severely restricts the use of `nonnull`. The Java Modelling Language (JML) adds formal specifications to Java and supports non-null types [7]. However, JML is strictly a specification language, and requires separate tools (such as ESC/Java) for checking.

Related work also exists on type inference for Object-Oriented languages (e.g. [21,24,28]). These, almost exclusively, assume the original program is completely untyped and employ set constraints (see [1]) for inferring types. This proceeds across method calls, necessitating knowledge of the program's call graph (which must be approximated in languages with dynamic dispatch). Typically, a constraint graph representing the entire program is held in memory at once, making these approaches somewhat unsuited to separate compilation [21]. Such systems share a strong relationship with other constraint-based program analyses, such as *points-to* analysis (e.g. [18,26,22,23]).

Several works also use techniques similar to type aliasing, albeit in different settings. Smith *et al.* capture aliasing constraints between locations in the program store to provide safe object deallocation and imperative updates [27]; for example, when an object is deallocated the supplied reference and any aliases are retyped to *junk*. Chang *et al.* maintain a graph, called the *e-graph*, of aliasing relationships between elements from different abstract domains [4]; their least upper bound operator maintains a very similar

invariant to ours. Zhang *et al.* consider aliasing of constraint variables in the context of set-constraint solvers [29].

8 Conclusion

We have presented a novel approach to the bytecode verification of non-null types. A key feature is that our system infers two kinds of information from conditionals: nullness information and type aliases. We have formalised this system for a subset of Java Bytecode, and proved soundness. Finally, we have detailed an implementation of our system and reported our experiences gained from using it. The tool itself is freely available from `http://www.mcs.vuw.ac.nz/~djp/JACK/`.

Acknowledgements. Thanks to Lindsay Groves, James Noble, Paul H.J. Kelly, Stephen Nelson, and Neil Leslie for many excellent comments on earlier drafts. This work is supported by the University Research Fund of Victoria University of Wellington.

References

1. Aiken, A.: Introduction to set constraint-based program analysis. Science of Computer Programming 35(2–3), 79–111 (1999)
2. Andreae, C., Noble, J., Markstrum, S., Millstein, T.: A framework for implementing pluggable type systems. In: OOPSLA, pp. 57–74. ACM Press, New York (2006)
3. Chalin, P., James, P.R.: Non-null references by default in Java: Alleviating the nullity annotation burden. In: Ernst, E. (ed.) ECOOP 2007. LNCS, vol. 4609, pp. 227–247. Springer, Heidelberg (2007)
4. Chang, B.-Y.E., Leino, K.R.M.: Abstract interpretation with alien expressions and heap structures. In: Cousot, R. (ed.) VMCAI 2005. LNCS, vol. 3385, pp. 147–163. Springer, Heidelberg (2005)
5. Chin, B., Markstrum, S., Millstein, T.: Semantic type qualifiers. In: PLDI, pp. 85–95. ACM Press, New York (2005)
6. Chin, B., Markstrum, S., Millstein, T., Palsberg, J.: Inference of user-defined type qualifiers and qualifier rules. In: Sestoft, P. (ed.) ESOP 2006. LNCS, vol. 3924, pp. 264–278. Springer, Heidelberg (2006)
7. Cielecki, M., Fulara, J., Jakubczyk, K., Jancewicz, L.: Propagation of JML non-null annotations in Java programs. In: PPPJ, pp. 135–140. ACM Press, New York (2006)
8. Ekman, T., Hedin, G.: Pluggable checking and inferencing of non-null types for Java. Journal of Object Technology 6(9), 455–475 (2007)
9. Ernst, M.: Annotations on Java types. Java Specification Request (JSR) 308 (2007)
10. Fähndrich, M., Leino, K.R.M.: Declaring and checking non-null types in an object-oriented language. In: OOPSLA, pp. 302–312. ACM Press, New York (2003)
11. Flanagan, C., Leino, K.R.M., Lillibridge, M., Nelson, G., Saxe, J.B., Stata, R.: Extended static checking for Java. In: Proc. PLDI, pp. 234–245. ACM Press, New York (2002)
12. Foster, J.S., Fähndrich, M., Aiken, A.: A theory of type qualifiers. In: Proc. PLDI, pp. 192–203. ACM Press, New York (1999)
13. Foster, J.S., Terauchi, T., Aiken, A.: Flow-sensitive type qualifiers. In: Proc. PLDI, pp. 1–12. ACM Press, New York (2002)
14. Goldberg, A.: A Specification of Java Loading and Bytecode Verification. In: Conference on Computer & Communications Security, pp. 49–58. ACM Press, New York (1998)

15. Hovemeyer, D., Pugh, W.: Finding more null pointer bugs, but not too many. In: Proc. PASTE, pp. 9–14. ACM Press, New York (2007)
16. Hovemeyer, D., Spacco, J., Pugh, W.: Evaluating and tuning a static analysis to find null pointer bugs. In: Proc. PASTE, pp. 13–19. ACM Press, New York (2005)
17. Leroy, X.: Java bytecode verification: algorithms and formalizations. Journal of Automated Reasoning 30(3/4), 235–269 (2003)
18. Lhoták, O., Hendren, L.J.: Context-sensitive points-to analysis: Is it worth it? In: Mycroft, A., Zeller, A. (eds.) CC 2006. LNCS, vol. 3923, pp. 47–64. Springer, Heidelberg (2006)
19. Lindholm, T., Yellin, F.: The Java Virtual Machine Specification, 2nd edn. The Java Series. Addison Wesley Longman, Inc., Amsterdam (1999)
20. Male, C., Pearce, D.J., Potanin, A., Dymnikov, C.: Java bytecode verification for @NonNull types. Technical report, Victoria University of Wellington (2007)
21. Palsberg, J., Schwartzbach, M.I.: Object-oriented type inference. In: Proc. OOPSLA, pp. 146–161. ACM Press, New York (1991)
22. Pearce, D.J., Kelly, P.H.J., Hankin, C.: Online cycle detection and difference propagation: Applications to pointer analysis. Software Quality Journal 12(4), 309–335 (2004)
23. Pearce, D.J., Kelly, P.H.J., Hankin, C.: Efficient field-sensitive pointer analysis for C. Transactions on Programming Languages and Systems 30(1) (2008)
24. Plevyak, J., Chien, A.A.: Precise concrete type inference for object-oriented languages. In: Proc. OOPSLA, pp. 324–340. ACM Press, New York (1994)
25. Pominville, P., Qian, F., Vallée-Rai, R., Hendren, L., Verbrugge, C.: A framework for optimizing Java using attributes. In: Wilhelm, R. (ed.) CC 2001. LNCS, vol. 2027, pp. 334–554. Springer, Heidelberg (2001)
26. Rountev, A., Milanova, A., Ryder, B.G.: Points-to analysis for Java using annotated constraints. In: Proc. OOPSLA, pp. 43–55. ACM Press, New York (2001)
27. Smith, F., Walker, D., Morrisett, G.: Alias types. In: Smolka, G. (ed.) ESOP 2000. LNCS, vol. 1782, pp. 366–381. Springer, Heidelberg (2000)
28. Wang, T., Smith, S.F.: Precise constraint-based type inference for Java. In: Knudsen, J.L. (ed.) ECOOP 2001. LNCS, vol. 2072, pp. 99–117. Springer, Heidelberg (2001)
29. Zhang, Y., Nielson, F.: A scalable inclusion constraint solver using unification. In: King, A. (ed.) LOPSTR 2007. LNCS, vol. 4915, Springer, Heidelberg (2007)