

AUTOMATED TERMINATION ANALYSIS OF JAVA BYTECODE BY TERM REWRITING

C. OTTO AND M. BROCKSCHMIDT AND C. VON ESSEN AND J. GIESL

LuFG Informatik 2, RWTH Aachen University, Germany

ABSTRACT. We present an automated approach to prove termination of **Java Bytecode (JBC)** programs by automatically transforming them to term rewrite systems (TRSs). In this way, the numerous techniques and tools developed for TRS termination can now be used for imperative object-oriented languages like **Java**, which can be compiled into **JBC**.

1. Introduction

Termination of TRSs and logic programs has been studied for decades. But as imperative programs dominate in practice, recently many results on termination of imperative programs were developed as well (e.g., [2, 3, 4, 5, 12]). Our goal is to re-use the wealth of techniques and tools from TRS termination when tackling imperative object-oriented programs. Similar TRS-based approaches have already proved successful for termination analysis of **Prolog** and **Haskell** [10, 17]. A first approach to prove termination of imperative programs by transforming them to TRSs was presented in [7]. However, [7] only analyzes a toy programming language without heap, whereas our goal is to analyze **JBC** programs.

JBC [14] is an assembly-like object-oriented language designed as intermediate format for the execution of **Java** [11] programs by a **Java Virtual Machine (JVM)**. Moreover, **JBC** is a common compilation target for many other languages besides **Java**. While there exist several static analysis techniques for **JBC**, we are only aware of two other automated methods to analyze termination of **JBC**, implemented in the tools **COSTA** [1] and **Julia** [19]. They transform **JBC** into a constraint logic program by abstracting every object of a dynamic data type to an integer denoting its path-length (i.e., the maximal length of the path of pointers that can be obtained by following the fields of objects). For example, consider a data structure **IntList** with the field **value** for the first list element and the field **next** which points to the next list element. Now an object of type **IntList** representing the list [0,1,2] would be abstracted to its length 3, but one would disregard the values of the list elements. While this fixed mapping from data objects to integers leads to a very efficient analysis, it also restricts the power of these methods. In contrast, in our approach we represent data objects not by integers, but by *terms*. To this end, we introduce a function symbol for every class. So the **IntList** object above is represented by a term like

Key words and phrases: Java Bytecode, termination, term rewriting.

Supported by the DFG grant GI 274/5-2 and by the G.I.F. grant 966-116.6.



`IntList(0, IntList(1, IntList(2, null)))`, which keeps the whole information of the data object.

So compared to [1, 19] and to direct termination analysis of imperative programs, rewrite techniques¹ have the advantage that they are very powerful for algorithms on user-defined data structures, since they can automatically generate suitable well-founded orders comparing arbitrary forms of terms. Moreover, by using TRSs with built-in integers [8], rewrite techniques are also powerful for algorithms on pre-defined data types like integers.

Inspired by our approach for termination of **Haskell** [10], in this paper we present a method to translate **JBC** programs to TRSs. More precisely, in Sect. 2 we show how to automatically construct a *termination graph* representing all execution paths of the **JBC** program. Similar graphs are also used in program optimization techniques, e.g. [18]. While we perform considerably less abstraction than [1, 19], we also apply a suitable abstract interpretation [6] in order to obtain finite representations for all possible forms of the heap at a certain state. In contrast to *control flow graphs*, the nodes of the termination graph contain not just the current program position, but also detailed information on the values of the variables and on the content of the heap. Thus, the termination graph usually has several nodes which represent the same program position, but where the values of the variables and the heap are different. This is caused by different runs through the program code. The termination graph takes care of all aliasing, sharing, and cyclicity effects in the **JBC** program. This is needed in order to express these effects in a TRS afterwards. Then, a TRS is generated from the termination graph such that termination of the TRS implies termination of the original **JBC** program (Sect. 3). The resulting TRSs can be handled by existing TRS termination techniques and tools.

As described in Sect. 4, we implemented the transformation in our tool **AProVE** [9]. In the first *International Termination Competition* on automated termination analysis of **JBC**, **AProVE** achieved competitive results compared to **Julia** and **COSTA**. So this paper shows for the first time that rewriting techniques can indeed be successfully used for termination of imperative object-oriented languages like **Java**.

2. From **JBC** to Termination Graphs

To obtain a finite representation of all execution paths, we evaluate the **JBC** program symbolically, resulting in a *termination graph*. Afterwards, this graph is used to generate a TRS suitable for termination analysis. Sect. 2.1 introduces the abstract states used in termination graphs. Then Sect. 2.2 illustrates the construction of termination graphs for simple programs and Sect. 2.3 extends it to programs with complex forms of sharing.

2.1. Representing States of the JVM

We define *abstract states* which represent *sets* of concrete **JVM** states, using a formalization which is especially suitable for a translation into TRSs (see e.g. [13] for related formalizations). Our approach is restricted to verified sequential **JBC** programs without recursion.

¹Of course, one could also use a transformation similar to ours where **JBC** is transformed to (constraint) logic programs, but where data objects are also represented by terms instead of integers. In principle, such an approach would be as powerful as ours, provided that one uses sufficiently powerful underlying techniques for automated termination analysis of logic programs. However, since some of the most powerful current termination analyzers for logic programs are based on term rewriting [15, 17], it seems more natural to transform **JBC** to term rewriting directly.

To simplify the presentation in the paper, we only consider program runs involving a single method, and exclude floating point arithmetic, arrays, exceptions, and static class fields. However, our approach can easily be extended to such constructs and to arbitrary many non-recursive methods. For the latter, we represent the frames of the call stack individually and simply “inline” the code of invoked methods. Indeed, our implementation also handles programs with several methods including floats, arrays, exceptions, and static fields.

Definition 2.1. The set of abstract states is $\text{STATES} = \text{PROGPOS} \times \text{LOCVAR} \times \text{OPSTACK} \times \text{HEAP}$.

The first component of a state corresponds to the program counter. We represent it by the next program instruction to be executed (e.g., by a **JBC** instruction like “`ifnull 8`”).

The second component is an array of the local variables which have a defined value at the current program position, represented by a partial function $\text{LOCVAR} = \mathbb{N} \rightarrow \text{REFERENCES}$. Here, **REFERENCES** are addresses in the heap. So in our representation, we do not store primitive values directly, but indirectly using references to the heap. This enables us to retain equality information for two otherwise unknown primitive values. Moreover, we require `null` \in **REFERENCES** to represent the `null` reference. To ease readability, in examples we usually denote local variables by names instead of numbers. Thus, “`o : o1, l : o2`” denotes an array where the 0-th local variable `o` references the address o_1 in the heap and the 1-st local variable `l` references the address o_2 in the heap. Of course, different local variables can point to the same address (e.g., in “`o : o1, l : o2, c : o1`”, `o` and `c` refer to the same object).

The third component is the operand stack that **JBC** instructions operate on. It will be filled with intermediate values such as operands of arithmetic operations when evaluating the bytecode. We represent it by a partial function $\text{OPSTACK} = \mathbb{N} \rightarrow \text{REFERENCES}$. The empty operand stack is denoted by “ ε ” and “ i_1, i_2 ” denotes a stack with top element i_2 .

| |
|--|
| <pre>ifnull 8 o:o₁, l:o₂ o₁ o₁ = Int(val = i₁) i₁ = (-∞, ∞) o₂ = Int(?)</pre> |
|--|

Figure 1: An abstract JVM state

To depict abstract states in examples, we write the first three components in the first line and separate them by “|”. The fourth **HEAP** component is written in the lines below, cf. Fig. 1. It describes the values of **REFERENCES**. We represent the **HEAP** by a partial

function $\text{HEAP} : \text{REFERENCES} \rightarrow \text{INTEGERS} \cup \text{INSTANCES} \cup \text{UNKNOWN}$.

The values in $\text{UNKNOWN} = \text{CLASSNAMES} \times \{?\}$ represent tree-shaped (and thus acyclic) objects for which we have no information except their type. **CLASSNAMES** contains the names of all classes and interfaces of the program. So for a class `Int`, “ $o_2 = \text{Int}(?)$ ” means that the object at address o_2 is `null` or an instance of type `Int` (or a subtype of `Int`).

We represent integers as possibly unbounded intervals, i.e. $\text{INTEGERS} = \{\{x \in \mathbb{Z} \mid a \leq x \leq b\} \mid a \in \mathbb{Z} \cup \{-\infty\}, b \in \mathbb{Z} \cup \{\infty\}, a \leq b\}$. So $i_1 = (-\infty, \infty)$ means that any integer can be at the address i_1 . Since current TRS termination tools cannot handle 32-bit `int`-numbers as in **JBC**, we treat `int` as the infinite set of all integers, i.e., we cannot handle problems related to overflows. Note that in **JBC**, `int` is also used for Boolean values.

To represent **INSTANCES** (i.e., objects) of some class, we describe the values of their fields, i.e., $\text{INSTANCES} = \text{CLASSNAMES} \times (\text{FIELDIDENTIFIERS} \rightarrow \text{REFERENCES})$. To prevent ambiguities, in general the **FIELDIDENTIFIERS** also contain the respective class names. So if the class `Int` has the field `val` of type `int`, then “ $o_1 = \text{Int}(\text{val} = i_1)$ ” means that at the address o_1 , there is an instance of class `Int` and its field `val` references the address i_1 in the heap. Note that all sharing and aliasing must be explicitly represented in the abstract state. So since the state in Fig. 1 contains no sharing information for o_1 and o_2 , o_1 and the references reachable from o_1 are disjoint from o_2 and from the references reachable from o_2 .

| | |
|--|---|
| <pre> 00: aload_0 // load orig to opstack 01: ifnull 8 // jump to line 8 if top // of opstack is null 04: aload_1 // load limit 05: ifnonnull 9 // jump if not null 08: return 09: aload_0 // load orig 10: astore_2 // store into copy 11: aload_0 // load orig 12: getfield val // load field val 15: aload_1 // load limit 16: getfield val // load field val 19: if_icmpge 35 // jump if // orig.val >= limit.val 22: aload_2 // load copy 23: aload_2 // load copy 24: getfield val // load field val 27: iconst_1 // load constant 1 28: iadd // add copy.val and 1 29: putfield val // store into copy.val 32: goto 11 35: return </pre> | <pre> public class Int { // only wrap a primitive int private int val; // count up to the value // in "limit" public static void count(Int orig, Int limit) { if (orig == null limit == null) { return; } // introduce sharing Int copy = orig; while (orig.val < limit.val) { copy.val++; } } } </pre> |
| (a) Java Bytecode | (b) Java Source Code |

Figure 2: Example using aliasing and an integer counting upwards

2.2. Termination Graphs for Simple Programs

We now introduce the *termination graph* using a simple example. In Fig. 2(a) we present the analyzed **JBC** program and Fig. 2(b) shows the corresponding **Java** source code.

We create the termination graph using the states of a run of our abstract virtual machine as nodes, starting in a suitable general state. In our example, we want to know if *all* calls of the method `count` with two distinct arbitrary `Int` objects (or `null`) as arguments terminate. Here it is important to handle the aliasing of the variables `copy` and `orig`.

In Fig. 3, node *A* contains the start state. For the local variables `orig` and `limit` (abbreviated *o* and *l*), we only know their type and we know that they do not share any part of the heap. The first **JBC** instruction `aload_0` loads the value of the 0-th local variable (the argument `orig`) on the operand stack. The variable `orig` references some address o_1 in the heap, but we do not need concrete information about o_1 for this instruction. The resulting new state *B* is connected to *A* by an *evaluation edge*.

To evaluate the `ifnull` instruction, we need to know if the reference on top of the operand stack is `null`. This is not yet known for o_1 . We *refine* the information and create successor nodes *C* and *D* for all possible cases (i.e., for $o_1 == \text{null}$, and for `Int` and all its non-abstract subclasses). In *C*, o_1 is `null`, and in *D* it is an instance of `Int` (`Int` has no proper subtypes). In *D*, the field values are new references in the heap. So instead of “ $o_1 = \text{Int}(?)$ ”, we now have “ $o_1 = \text{Int}(\text{val} = i_1)$ ”. Note that while “ $o_1 = \text{Int}(?)$ ” in node *B* means that if o_1 is not `null`, then it has type `Int` or a subtype of it, “ $o_1 = \text{Int}(\text{val} = i_1)$ ” in node *D* means that o_1 ’s type is exactly `Int` and not a proper subtype. We have no information about the value at i_1 . Therefore, i_1 gets the most general value for `INTEGERS`, i.e., $i_1 = (-\infty, \infty)$. *C* and *D* are connected to *B* by *refinement edges*.

Now we can evaluate the instruction both for *C* and *D*, leading to *E* and *F*. Evaluation stops in *E*, while for *F*, the same procedure is repeated for the argument `limit`, leading to node *G* (among others) after several steps, indicated by a dotted arrow. Note the aliasing between `copy` and `orig`, since both reference the same object at the address o_1 .

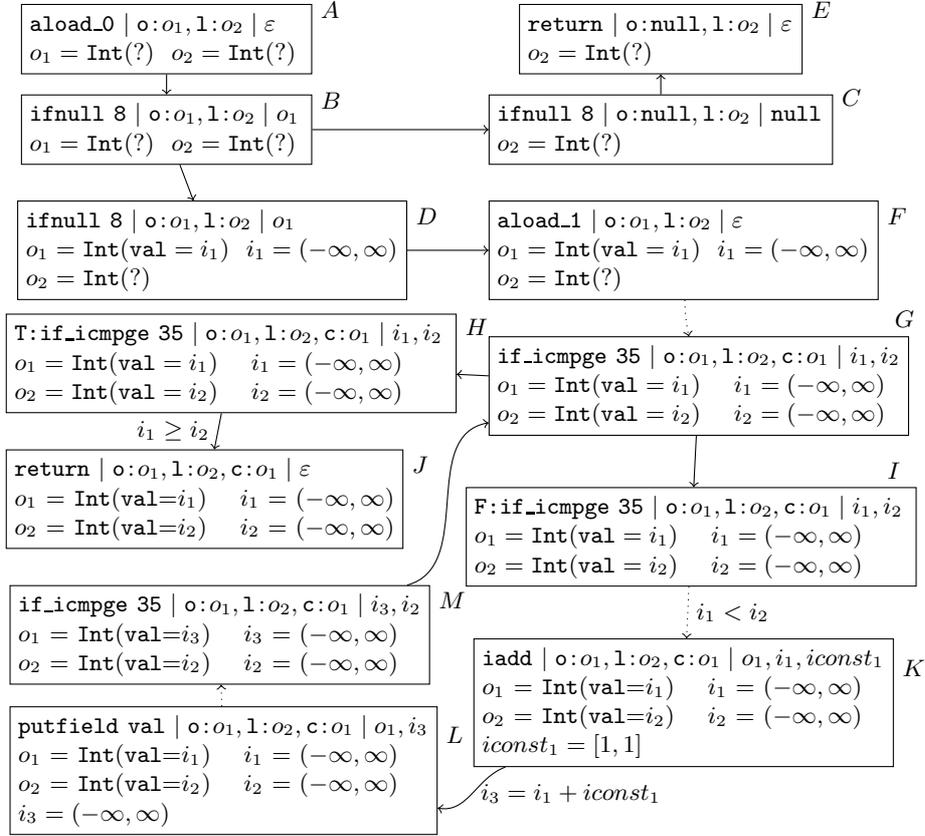


Figure 3: Termination graph for count

In G , we have already evaluated the two “getfield val” instructions and have pushed the two integer values on the operand stack. Now `if_icmpge` requires us to compare the unknown integers at i_1 and i_2 . If we had to compare i_1 with a fixed number like 0, we could refine the information about i_1 and i_2 and create two successor nodes with $i_1 = (-\infty, -1]$ and $i_1 = [0, \infty)$. But “ $i_1 \geq i_2$ ” is not expressible in our abstract states. Here, we split according to both possible values of the condition (depicted using the labels “T” and “F”, respectively). This leads to the nodes H and I which are connected to G by *split edges*.

We can evaluate the condition in H to **true** and label the resulting evaluation edge to J by this condition. We will use these labels when constructing a TRS from the termination graph. J marks the program end and thus, it remains a leaf of the graph.

In I , we can evaluate the condition to **false** and label the next edge by the converse of the condition. After evaluating the next four instructions we reach node K . On the top positions of the operand stack, there are two integer variables (where the topmost variable has the value 1). The instruction `iadd` adds these two variables resulting in a new integer variable i_3 . The relation between i_3 , i_1 , and $iconst_1$ is added as a label on the evaluation edge to the new node L . This label will again be used in the TRS construction.

From L on, we evaluate instructions until we again arrive at the instruction `if_icmpge` in node M . It turns out that M is an *instance* of the previous node G . Hence, we can connect M with G by an *instantiation edge*. The reason is that every concrete state which would be described by the abstract state M could also be described by the state G .

One has to expand termination graphs until all leaves correspond to program ends. Hence, our graph is now completed. By using appropriate generalization steps (which transform nodes into more general ones), one can always obtain a finite termination graph. To this end, one essentially executes the program symbolically until one reaches some position in the program for the second time. Then, a new state is created that is a generalization of both original states and one introduces instantiation edges from the two original states to the new generalized state. Of course, in our implementation we apply suitable heuristics to ensure that one only performs finitely many such generalization steps and to guarantee that the construction always terminates with a finite termination graph.

To define “*instance*” formally, we first define all positions π of references in a state s , where $s|_\pi$ denotes the reference at position π . A position π is a sequence starting with LV_n or OS_n for some $n \in \mathbb{N}$ (indicating the n -th reference in the local variable array or in the operand stack), followed by zero or more `FIELDIDENTIFIERS`.

Definition 2.2 (position, SPOS). Let $s = (pp, l, op, h) \in \text{STATES}$. Then $\text{SPOS}(s)$ is the smallest set such that one of the following holds for all $\pi \in \text{SPOS}(s)$:

- $\pi = LV_n$ for some $n \in \mathbb{N}$ where $l(n)$ is defined. Then $s|_\pi$ is $l(n)$.
- $\pi = OS_n$ for some $n \in \mathbb{N}$ where $op(n)$ is defined. Then $s|_\pi$ is $op(n)$.
- $\pi = \pi'v$ for some $v \in \text{FIELDIDENTIFIERS}$ and some $\pi' \in \text{SPOS}(s)$ where $h(s|_{\pi'}) = (c, f) \in \text{INSTANCES}$ and where $f(v)$ is defined. Then $s|_\pi$ is $f(v)$.

As an example, consider the state s depicted in node G of Fig. 3. Here we have three local variables and two elements on the operand stack. Thus, $\text{SPOS}(s)$ contains $LV_0, LV_1, LV_2, OS_0, OS_1$, where $s|_{LV_0} = s|_{LV_2} = o_1$, $s|_{LV_1} = o_2$, $s|_{OS_0} = i_1$, and $s|_{OS_1} = i_2$. If h is the heap of that state, then $h(o_1) = (\text{Int}, f_1) \in \text{INSTANCES}$, where $f_1(\text{val}) = i_1$. Hence, “ $LV_0 \text{ val}$ ” is also a position in $\text{SPOS}(s)$ and $s|_{LV_0 \text{ val}} = i_1$. The remaining elements of $\text{SPOS}(s)$ are “ $LV_2 \text{ val}$ ” and “ $LV_1 \text{ val}$ ”, where $s|_{LV_2 \text{ val}} = i_1$ and $s|_{LV_1 \text{ val}} = i_2$.

Intuitively, a state s' is an instance of a state s if they correspond to the same program position and whenever there is a reference $s'|_\pi$, then either the values represented by $s'|_\pi$ in the heap of s' are a subset of the values represented by $s|_\pi$ in the heap of s or else, π is no position in s . Moreover, shared parts of the heap in s' must also be shared in s . Note that since s and s' correspond to the same position in a *verified JBC* program, s and s' have the same number of local variables and their operand stacks have the same size.

Definition 2.3 (Instance). We say that $s' = (pp', l', op', h')$ is an *instance* of state $s = (pp, l, op, h)$ (denoted $s' \sqsubseteq s$) iff $pp = pp'$, and for all $\pi, \pi' \in \text{SPOS}(s')$:

- (a) if $s'|_\pi = s'|_{\pi'}$ and $h'(s'|_\pi) \in \text{INSTANCES} \cup \text{UNKNOWN}$, then $\pi, \pi' \in \text{SPOS}(s)$ and $s|_\pi = s|_{\pi'}$
- (b) if $s'|_\pi \neq s'|_{\pi'}$ and $\pi, \pi' \in \text{SPOS}(s)$, then $s|_\pi \neq s|_{\pi'}$
- (c) if $h'(s'|_\pi) \in \text{INTEGERS}$ and $\pi \in \text{SPOS}(s)$, then $h(s|_\pi) \in \text{INTEGERS}$ and $h'(s'|_\pi) \subseteq h(s|_\pi)$
- (d) if $s'|_\pi = \text{null}$ and $\pi \in \text{SPOS}(s)$, then $s|_\pi = \text{null}$ or $h(s|_\pi) = (c, ?) \in \text{UNKNOWN}$
- (e) if $h'(s'|_\pi) = (c', ?)$ and $\pi \in \text{SPOS}(s)$, then $h(s|_\pi) = (c, ?)$ where c' is c or a subtype of c
- (f) if $h'(s'|_\pi) = (c', f') \in \text{INSTANCES}$ and $\pi \in \text{SPOS}(s)$, then $h(s|_\pi) = (c', f) \in \text{INSTANCES}$ or $h(s|_\pi) = (c, ?)$, where c' must be c or a subtype of c .

The state s' in node M of Fig. 3 is an instance of the state s in node G . Clearly, they both refer to the same program position. It remains to examine the references reachable in s' . We have $\text{SPOS}(s') = \text{SPOS}(s) = \{LV_0, LV_1, LV_2, OS_0, OS_1, LV_0 \text{ val}, LV_1 \text{ val}, LV_2 \text{ val}\}$. It is easy to check that the conditions of Def. 2.3 are satisfied for all these positions π . We illustrate this for $\pi = LV_0 \text{ val}$. Here, $s'|_\pi = i_3$ and if h' is the heap of s' , then $h'(i_3) =$

$(-\infty, \infty)$. Similarly, $s|_\pi = i_1$ and if h is the heap of s , then $h(i_1) = (-\infty, \infty)$. Here, s' and s are in fact equivalent, since M is an instance of G and G is an instance of M .

| |
|--|
| <pre>if_icmpge35 o:o1, l:o2, c:o1 i1, i2 o1 = Int(val=i1) i1 = [1, 1] o2 = Int(val=i2) i2 = [10000, 10000]</pre> |
|--|

As remarked before, abstract states describe sets of *concrete states* like the one in Fig. 4, which is an instance of G and M . Here, the values for i_1 and i_2 are proper integers instead of intervals.

Figure 4: A concrete state

Definition 2.4 (Concrete state). A state $s = (pp, l, op, h)$ is *concrete* if for all $\pi \in \text{SPOS}(s)$:

- $h(s|_\pi) \notin \text{UNKNOWN}$ and
- if $h(s|_\pi) \in \text{INTEGERS}$, then $h(s|_\pi)$ is just a singleton interval $[i, i]$ for some $i \in \mathbb{Z}$

A concrete state has no proper instances (i.e., if s is concrete and $s' \sqsubseteq s$, then $s \sqsubseteq s'$). Concrete states that are not a program end can always be evaluated and have exactly one (concrete) successor state. For Fig. 4, since i_1 's value is not greater or equal than i_2 's, the successor state corresponds to the instruction “`aload_2`”, with the same local variables and empty operand stack. Such a sequence of concrete states, obtained by **JBC** evaluation, is called a *computation sequence*. Our construction of termination graphs ensures that

if s is an abstract state in the termination graph and there is a concrete state $t \sqsubseteq s$ where t evaluates to the concrete state t' , then the termination graph contains a path from s to a state s' with $t' \sqsubseteq s'$. (2.1)

To see why (2.1) holds, note that in the termination graph, s is first refined to a state \bar{s} with $t \sqsubseteq \bar{s}$. So there is a path from s to \bar{s} , and in the state \bar{s} , all concrete information needed for an actual evaluation according to the **JBC** specification [14] is available. Note that “evaluation edges” in the termination graph are defined by exactly following the specification of **JBC** in [14]. Thus, there is an evaluation edge from \bar{s} to s' , where $t' \sqsubseteq s'$.

The computation sequence from Fig. 4 to its concrete successor corresponds to the path from node M or G to I 's successor. Paths in the graph that correspond to computation sequences are called *computation paths*. Our goal is to show that all these paths are finite.

Definition 2.5 (Graph termination). A finite or infinite path $s_1^1, \dots, s_1^{n_1}, s_2^1, \dots, s_2^{n_2}, \dots$ through the termination graph is called a *computation path* iff there is a computation sequence t_1, t_2, \dots of concrete states where $t_i \sqsubseteq s_i^1$ for all i . A termination graph is called *terminating* iff it has no infinite computation path. Note that due to (2.1), if the termination graph is terminating, then the original **JBC** program is also terminating for all concrete states t where $t \sqsubseteq s$ for some abstract state s in the termination graph.

2.3. Termination Graphs for Complex Programs

Now we discuss sharing problems in complex programs with recursive data types. In Fig. 5, `flatten` takes a list of binary trees whose nodes are labeled by integers. It performs a depth-first run through all trees and returns the list of all numbers in these trees. It terminates because each loop iteration decreases the total number of all nodes in the trees of `list`, even though `list`'s length may increase. Note that `list` and `cur` share part of the heap.

Consider the three states A , B , and C in Fig. 6. A is the state of our abstract **JVM** when it first reaches the loop condition “`cur != null`” (where `list`, `cur`, and `result` are abbreviated by `l`, `c`, and `r`). After one execution of the loop body, one obtains the state B if `tree` is `null` and C otherwise. Note that local variables declared in the loop body are no longer defined at the loop condition, and hence, they do not occur in A , B , or C .

```

public class Flatten {
    public static IntList flatten(TreeList list) {
        TreeList cur = list;
        IntList result = null;
        while (cur != null) {
            Tree tree = cur.value;
            if (tree != null) {
                IntList oldIntList = result;
                result = new IntList();
                result.value = tree.value;
                result.next = oldIntList;
                TreeList oldCur = cur;
                cur = new TreeList();
                cur.value = tree.left;
                cur.next = oldCur;
                oldCur.value = tree.right;
            } else cur = cur.next;
        }
        return result;
    }
}

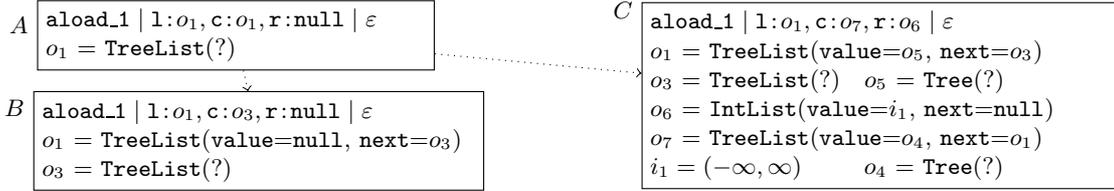
public class Tree {
    int value;
    Tree left;
    Tree right;
}

public class TreeList {
    Tree value;
    TreeList next;
}

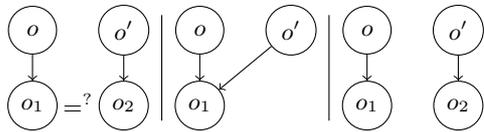
public class IntList {
    int value;
    IntList next;
}

```

Figure 5: Example converting a list of binary trees to a list of integers

Figure 6: Three states of the termination graph of `flatten`

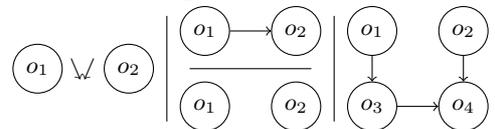
If one continued the evaluation like this, one would obtain an infinite tree, since one never reaches any state which is an instance of a previous state. (In particular, B and C are no instances of A .) Hence, to obtain *finite* graphs, one sometimes has to *generalize* states. Thus, we want to create a new general state S such that A , B , and C are instances of S . Note that in S , l and c cannot point to different references with UNKNOWN values, since then S would only represent states where l and c are tree-shaped and not sharing. However, l and c point to the *same* object in A , one can reach $c:o_3$ from $l:o_1$ in B (i.e., l joins c , since a field value of o_1 is o_3), and one can reach $l:o_1$ from $c:o_7$ in C . To express such sharing information in general states, we extend states by *annotations*.

Figure 7: “ $=?$ ” annotation

In Fig. 7, the leftmost picture depicts a heap where an instance referenced by o has a field value o_1 and o' has a field value o_2 . The annotation “ $o_1 =? o_2$ ” means that o_1 and o_2 could be equal. Here the value of at least one of o_1 and o_2 must be UNKNOWN. So both the second and the third shape in Fig. 7 are instances of the first. In the second shape, o_1 and o_2 are equal and all occurrences of o_2 can be replaced by o_1 (or vice versa). In the third shape, o_1 and o_2 are not the same and thus, the annotation has been removed.

So the $=?$ annotation covers both the equality of l and c in state A and their non-equality in states B and C . To represent states where l and c *may* join, we use the annotation “ \surd ”. We say that a

In Fig. 7, the leftmost picture depicts a heap where an instance referenced by o has a field value o_1 and o' has a field value o_2 . The annotation “ $o_1 =? o_2$ ” means that o_1 and o_2 could be equal. Here the value of at least one of o_1 and o_2 must be UNKNOWN. So both the second and the third shape in Fig. 7 are instances of the first. In the second shape, o_1 and o_2 are equal and all occurrences of o_2 can be replaced by o_1 (or vice versa). In the third shape, o_1 and o_2 are not the same and thus, the annotation has been removed.

Figure 8: “ \surd ” annotation

reference o' is a *direct successor* of a reference o (denoted $o \rightarrow o'$) iff the object at address o has a field whose value is o' . As an example, consider state B in Fig. 6, where $o_1 \rightarrow o_3$ holds. Then the annotation “ $o_1 \Downarrow o_2$ ” means that if o_1 is UNKNOWN, then there could be an object o with $o_1 \rightarrow^+ o$ and $o_2 \rightarrow^* o$, i.e., o is a proper successor of o_1 and a (possibly non-proper) successor of o_2 . Note that \Downarrow is symmetric,² so $o_1 \Downarrow o_2$ also means that if o_2 is UNKNOWN, then there could be an object o' with $o_1 \rightarrow^* o'$ and $o_2 \rightarrow^+ o'$. The shapes 2-4 in Fig. 8 visualize three possible instances of the state with annotation “ $o_1 \Downarrow o_2$ ”. Note that a state in which o_1 and o_2 do not share is also an instance.

We can now create a state S (see Fig. 9) such that $A, B, C \sqsubseteq S$. The annotations state that l and c may be equal (as in A), that l may join c (as in B), or c may join l (as in C).

So to obtain a finite termination graph, after reaching A , we generalize it to a new node S connected by an instantiation edge. As seen in D , we introduce new forms of refinement edges to refine a state with the annotation “ $o_1 =^? o_7$ ” into the two instances where $o_1 = o_7$ and where $o_1 \neq o_7$. For $o_1 = o_7$, we reach B' and C' which are like B and C but now r points to a list ending with o_6 instead of `null`. The nodes B' and C' are connected back to S with instantiation edges. For $o_1 \neq o_7$, due to `c != null`, we first refine the information about o_7 , and obtain $o_7 = \text{TreeList}(\text{value} = o_8, \text{next} = o_9)$. Note that “ \Downarrow ” annotations have to be updated during refinements. If we have the annotation “ $o_1 \Downarrow o_7$ ” and if one refines o_7 by introducing references like o_8, o_9 for its non-primitive fields, then we have to add corresponding annotations such as “ $o_1 =^? o_9$ ” for all field references like o_9 whose types correspond to the type of o_1 . Moreover, we add “ \Downarrow ” annotations for all non-primitive field references (i.e., “ $o_1 \Downarrow o_8$ ” and “ $o_1 \Downarrow o_9$ ”). If after this refinement neither o_1 nor o_7 were UNKNOWN, we would delete the annotation $o_1 \Downarrow o_7$ since it has no effect anymore.

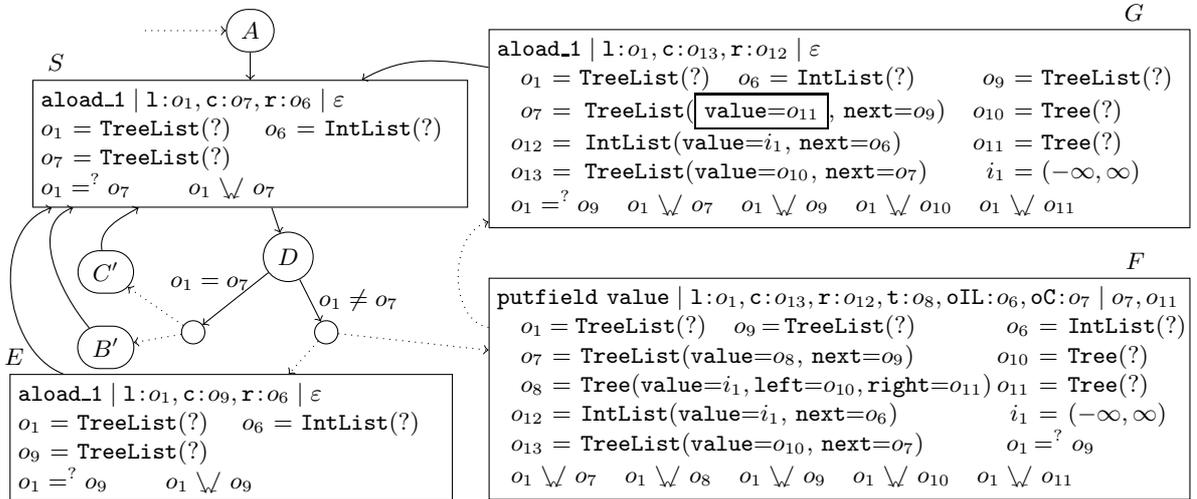


Figure 9: Termination graph for flatten

Now we use a refinement that corresponds to the case analysis whether `tree` is `null`. For `tree == null`, after one loop iteration we reach node E which is again an instance of S . Here, the local variable `tree` is no longer visible.

²Since both “ $=^?$ ” and “ \Downarrow ” are symmetric, we do not distinguish between “ $o_1 =^? o_2$ ” and “ $o_2 =^? o_1$ ” and we also do not distinguish between “ $o_1 \Downarrow o_2$ ” and “ $o_2 \Downarrow o_1$ ”.

For `tree != null`, the graph shows nodes F and G . In F we need to evaluate a `putfield` instruction (corresponding to “`oldCur.value = tree.right`”), i.e., we have to put the object at address o_{11} to the field `value` of the object at address o_7 . The effect of this operation can be seen in the box in state G , where the value of the object at o_7 was changed from o_8 to o_{11} . In G (which again corresponds to the loop condition), we removed the reference o_8 since it is no longer accessible from the local variables or the operand stack.

In contrast to other evaluation steps, such `putfield` instructions can give rise to additional annotations, since objects that already shared parts of the heap with o_7 now may also share parts of the heap with o_{11} . We say that a reference o *reaches* a reference o' iff there is a successor r of o (i.e., $o \rightarrow^* r$) such that $r = o'$ or $r = ? o'$ or $r \Downarrow o'$. So in our example, o_{11} reaches just o_{11} and o_1 . Now if we write o_{11} to a field of o_7 , then for all references o with $o \Downarrow o_7$, we have to add the annotation $o \Downarrow o'$ for all o' where o_{11} reaches o' . Hence, in our example, we would have to add $o_1 \Downarrow o_{11}$ (which is already present in the state) and $o_1 \Downarrow o_1$. However, the annotation $o_1 \Downarrow o_1$ has no effect, since by “ $o_1 = \text{TreeList}(?)$ ”, we know that o_1 only represents tree-shaped objects. Therefore, we can immediately drop $o_1 \Downarrow o_1$ from the state. Concrete non-tree-shaped objects can of course be represented easily (e.g., “ $o = \text{TreeList}(\text{value} = o', \text{next} = o)$ ”). But to represent an *arbitrary* possibly non-tree-shaped object o , we use a special annotation (depicted “ $o!$ ”).³

Definition 2.6 (Instance & annotations). We extend Def. 2.3 to $s, s' = (pp', l', op', h') \in \text{STATES}$ possibly containing annotations. Now $s' \sqsubseteq s$ holds iff for all $\pi, \pi' \in \text{SPOS}(s')$, the following conditions are satisfied in addition to Def. 2.3 (b)-(f). Here, let τ resp. τ' be the maximal prefix of π resp. π' such that both $\tau, \tau' \in \text{SPOS}(s)$.

- (a) if $s'|_\pi = s'|_{\pi'}$ where $h'(s'|_\pi) \in \text{INSTANCES} \cup \text{UNKNOWN}$, and if $\pi, \pi' \in \text{SPOS}(s)$,⁴
then $s|_\pi = s|_{\pi'}$ or $s|_\pi = ? s|_{\pi'}$
- (b) if $s'|_\pi = ? s'|_{\pi'}$ and $\pi, \pi' \in \text{SPOS}(s)$, then $s|_\pi = ? s|_{\pi'}$
- (c) if $(s'|_\pi = s'|_{\pi'})$ where $h'(s'|_\pi) \in \text{INSTANCES} \cup \text{UNKNOWN}$, or $s'|_\pi = ? s'|_{\pi'}$
and π or $\pi' \notin \text{SPOS}(s)$ with $\pi \neq \pi'$, then $s|_\tau \Downarrow s|_{\tau'}$
- (d) if $s'|_\pi \Downarrow s'|_{\pi'}$, then $s|_\tau \Downarrow s|_{\tau'}$
- (e) if $s'|_\pi!$ holds, then $s|_\tau!$
- (f) if there exist (possibly empty) sequences $\rho \neq \rho'$ of `FIELDIDENTIFIERS` without common prefix, where $s'|_{\pi\rho} = s'|_{\pi\rho'}$, $h'(s'|_{\pi\rho}) \in \text{INSTANCES} \cup \text{UNKNOWN}$,
and $(\pi\rho$ or $\pi\rho' \notin \text{SPOS}(s)$ or $s|_{\pi\rho} = ? s|_{\pi\rho'})$, then $s|_\tau!$

3. From Termination Graphs to TRSs

Now we transform termination graphs into *integer term rewrite systems (ITRSs)*. These are TRSs where the Booleans \mathbb{B} , the integers \mathbb{Z} , and their built-in operations $\text{ArithOp} = \{+, -, *, /, \%, \ll, \gg, \ggg, \hat{\ }, \&, |\}$ and $\text{RelOp} = \{>, \geq, <, \leq, ==, \neq\}$ are pre-defined by an infinite set of variable-free rules \mathcal{PD} . For example, \mathcal{PD} contains $1+2 \rightarrow 3$ and $5 < 4 \rightarrow \text{false}$. As shown in [8], TRS termination techniques can easily be adapted to ITRSs as well.

³Such annotations can also result from `putfield` operations which write a reference o_2 to a field \mathbf{f} of o_1 . If o_1 already reached o_2 before through some field $\mathbf{g} \neq \mathbf{f}$, then we add “ $o_1!$ ”, since o_1 is no longer a tree. Even worse, if o_2 reached o_1 before, then `putfield` creates a cyclic object and we add “ $o_1!$ ” and “ $o_2!$ ”.

⁴In contrast to Def. 2.3(a), here one may allow that π or $\pi' \notin \text{SPOS}(s)$. This case is handled in (c).

Definition 3.1 (ITRS [8]). An *ITRS* is a finite conditional TRS with rules “ $\ell \rightarrow r \mid b$ ”. Here ℓ, r, b are terms, where $\ell \notin \mathbb{B} \cup \mathbb{Z}$ and ℓ contains no symbol from $\mathcal{A}rithOp \cup \mathcal{R}elOp$. However, b and r may contain extra variables not occurring in ℓ . We often omit the condition b if b is true. The *rewrite relation* $\hookrightarrow_{\mathcal{R}}$ of an ITRS \mathcal{R} is the smallest relation where $t_1 \hookrightarrow_{\mathcal{R}} t_2$ iff there is a rule $\ell \rightarrow r \mid b$ from $\mathcal{R} \cup \mathcal{PD}$ such that $t_1|_p = \ell\sigma$, $b\sigma \hookrightarrow_{\mathcal{R}}^* \text{true}$, and $t_2 = t_1[r\sigma]_p$. Here, $\ell\sigma$ must not have instances of left-hand sides of rules as proper subterms, and σ must be *normal* (i.e., $\sigma(y)$ is in normal form also for variables y occurring only in b or r). Thus, the rewrite relation $\hookrightarrow_{\mathcal{R}}$ corresponds to an *innermost* evaluation strategy.

So if \mathcal{R} contains the rule “ $f(x) \rightarrow g(x, y) \mid x > 2$ ”, then $f(1 + 2) \hookrightarrow_{\mathcal{R}} f(3) \hookrightarrow_{\mathcal{R}} g(3, 27)$. Hence, extra variables in conditions or right-hand sides of rules stand for arbitrary values.

We first show how to transform a reference o in a state s into a term $\text{tr}(s, o)$. References pointing to concrete integers like $\text{iconst}_1 = [1, 1]$ in state K of Fig. 3 are transformed into the corresponding integer constant `1`. The reference `null` is transformed into the constant `null`. References pointing to instances will be transformed by a refined transformation $\text{ti}(s, o)$ in a more subtle way in order to take their types and the values of their fields into account. Finally, any other reference o is transformed into a variable (which we also call o). So $i_1 = (-\infty, \infty)$ in state K of Fig. 3 is transformed to the variable i_1 .

Definition 3.2 (Transforming references). Let $s = (pp, l, op, h) \in \text{STATES}$, $o \in \text{REFERENCES}$.

$$\text{tr}(s, o) = \begin{cases} i & \text{if } h(o) = [i, i], \text{ where } i \in \mathbb{Z} \\ \text{null} & \text{if } o = \text{null} \\ \text{ti}(s, o) & \text{if } h(o) \in \text{INSTANCES} \\ o & \text{otherwise} \end{cases}$$

The main advantage of our approach becomes obvious when transforming instances (i.e., data objects) into terms. The reason is that data objects essentially *are* terms and we simply keep their structure when transforming them. So for any object, we use the name of its class as a function symbol. The arguments of the function symbol correspond to the fields of the class. As an example, consider o_{13} in state F of Fig. 9. This data object is transformed to the term $\text{TreeList}(o_{10}, \text{TreeList}(\text{Tree}(i_1, o_{10}, o_{11}), o_9))$.

However, we also have to take the class hierarchy into account. Therefore, for any class c with n fields, let the corresponding function symbol now have arity $n + 1$. The arguments $2, \dots, n + 1$ correspond to the values of the fields declared in class c . The first argument represents the part of the object that corresponds to subclasses of c . As an example, consider a class `A` with a field `a` of type `int` and a class `B` which extends `A` and has a field `b` of type `int`. If x is a data object of type `A` where $x.a$ is `1`, then we now represent it by the term $\text{A}(\text{eoc}, 1)$. Here, the first argument of `A` is a constant `eoc` (for “end of class”) which indicates that the type of x is really `A` and not a subtype of `A`. If y is a data object of type `B` where $y.a$ is `2` and $y.b$ is `3`, then we represent it by the term $\text{A}(\text{B}(\text{eoc}, 3), 2)$. So the class hierarchy is represented by nesting the function symbols corresponding to the respective classes.

More precisely, since every class extends `java.lang.Object` (which has no fields), each such term now has the form `java.lang.Object(...)`. Hence, if we abbreviate the function symbol `java.lang.Object` by `jIO`, then for the `TreeList` object above, now the corresponding term is `jIO(TreeList(eoc, o10, jIO(TreeList(eoc, jIO(Tree(eoc, i1, o10, o11), o9))))))`.

Of course, we can only transform tree-shaped objects to terms. If $\pi \in \text{SPOS}(s)$ and if there is a non-empty sequence ρ of `FIELDIDENTIFIERS` such that $s|_{\pi} = s|_{\pi\rho}$, then $s|_{\pi}$ is called *cyclic* in s . If $s|_{\pi}$ is cyclic or marked by “!”, then $s|_{\pi}$ is called *special*. Every special

reference o is transformed into a variable o in order to represent an “arbitrary unknown” object. To define the transformation $\text{ti}(s, o)$ formally, we use an auxiliary transformation $\overline{\text{ti}}(s, o, c)$ which only considers the part of the class hierarchy starting with c .

Definition 3.3 (Transforming instances). We start the construction at the root of the class hierarchy (i.e., with `java.lang.Object`) and define $\text{ti}(s, o) = \overline{\text{ti}}(s, o, \text{java.lang.Object})$.

Let $s = (pp, l, op, h) \in \text{STATES}$ and let $h(o) = (c_o, f) \in \text{INSTANCES}$. Let $(c_1 = \text{java.lang.Object}, c_2, \dots, c_n = c_o)$ be ordered according to the class hierarchy, i.e., c_i is the direct superclass of c_{i+1} . We define the term $\overline{\text{ti}}(s, o, c_i)$ as follows:

$$\overline{\text{ti}}(s, o, c_i) = \begin{cases} o & \text{if } o \text{ is special} \\ c_o(\text{eoc}, \text{tr}(s, v_1), \dots, \text{tr}(s, v_m)) & \text{if } c_i = c_o, \text{fv}(c_o, f) = v_1, \dots, v_m \\ c_i(\overline{\text{ti}}(s, o, c_{i+1}), \text{tr}(s, v_1), \dots, \text{tr}(s, v_m)) & \text{if } c_i \neq c_o, \text{fv}(c_i, f) = v_1, \dots, v_m \end{cases}$$

For $c \in \text{CLASSNAMES}$, let $\mathbf{f1}, \dots, \mathbf{fm}$ be the fields declared in c in some fixed order. Then for $f : \text{FIELDIDENTIFIERS} \rightarrow \text{REFERENCES}$, $\text{fv}(c, f)$ gives the values $f(\mathbf{f1}), \dots, f(\mathbf{fm})$.

So for class **A** and **B** above, if $h(o) = (\mathbf{B}, f)$ where $f(\mathbf{a}) = [2, 2]$ and $f(\mathbf{b}) = [3, 3]$, then $\text{fv}(\mathbf{A}, f) = [2, 2]$, $\text{fv}(\mathbf{B}, f) = [3, 3]$ and thus, $\overline{\text{ti}}(s, o, \text{java.lang.Object}) = \text{jIO}(\mathbf{A}(\mathbf{B}(\text{eoc}, 3), 2))$.

To transform a whole state, we create the tuple of the terms that correspond to the references in the local variables and the operand stack. For example, state J in Fig. 3 is transformed to the tuple of the terms $\text{jIO}(\text{Int}(\text{eoc}, i_1))$, $\text{jIO}(\text{Int}(\text{eoc}, i_2))$, and $\text{jIO}(\text{Int}(\text{eoc}, i_1))$.

Definition 3.4 (Transforming states). Let $s = (pp, l, op, h) \in \text{STATES}$, let lv_0, \dots, lv_n and os_0, \dots, os_m be the references in l and op , respectively (i.e., $h(\text{LV}_i) = lv_i$ and $h(\text{OS}_i) = os_i$). We define the following mapping: $\text{ts}(s) = (\text{tr}(s, lv_0), \dots, \text{tr}(s, lv_n), \text{tr}(s, os_0), \dots, \text{tr}(s, os_m))$.

There is a connection between the instance relation on states and the matching relation on the corresponding terms. If s' is an instance of state s , then the terms in the transformation of s match the terms in the transformation of s' . Hence, if one generates rules matching the term representation of s , then these rules also match the term representation of s' .

Lemma 3.5. *Let $s' \sqsubseteq s$. Then there exists a substitution σ such that $\text{ts}(s)\sigma = \text{ts}(s')$.*⁵

Now we show how to build an ITRS from a termination graph such that termination of the ITRS implies termination of the graph, i.e., that there is no infinite computation path $s_1^1, \dots, s_1^{n_1}, s_2^1, \dots, s_2^{n_2}, \dots$. In other words, there should be no infinite computation sequence t_1, t_2, \dots of concrete states where $t_i \sqsubseteq s_i^1$ for all i .

For any abstract state s of the graph, we introduce a new function symbol \mathbf{f}_s . The arity of \mathbf{f}_s is the number of components in the tuple $\text{ts}(s)$. Our goal is to generate an ITRS \mathcal{R} such that $\mathbf{f}_{s_i^1}(\text{ts}(t_i)) \xrightarrow{+}_{\mathcal{R}} \mathbf{f}_{s_{i+1}^1}(\text{ts}(t_{i+1}))$ for all i . In other words, every computation path in the graph must be transformable into a rewrite sequence. Then each infinite computation path corresponds to an infinite rewrite sequence with \mathcal{R} .

To this end, we transform each edge in the termination graph into a rewrite rule. Let s, s' be two states connected by an edge e . If e is a split edge or an evaluation edge, then the corresponding rule should rewrite any instance of s to the corresponding instance of s' . Hence, we generate the rule $\mathbf{f}_s(\text{ts}(s)) \rightarrow \mathbf{f}_{s'}(\text{ts}(s'))$. For example, the edge from D to F in Fig. 3 results in the rule $\mathbf{f}_D(\text{jIO}(\text{Int}(\text{eoc}, i_1)), o_2, \text{jIO}(\text{Int}(\text{eoc}, i_1))) \rightarrow \mathbf{f}_F(\text{jIO}(\text{Int}(\text{eoc}, i_1)), o_2)$.

⁵For all proofs, we refer to [16].

If the evaluation involves checking some integer condition, we create a corresponding conditional rule. For example, the edge from H to J in Fig. 3 yields the rule

$$\begin{array}{l} f_H(\text{jIO}(\text{Int}(\text{eoc}, i_1)), \text{jIO}(\text{Int}(\text{eoc}, i_2)), \text{jIO}(\text{Int}(\text{eoc}, i_1)), i_1, i_2) \rightarrow \\ f_J(\text{jIO}(\text{Int}(\text{eoc}, i_1)), \text{jIO}(\text{Int}(\text{eoc}, i_2)), \text{jIO}(\text{Int}(\text{eoc}, i_1))) \quad | \quad i_1 \geq i_2 \end{array}$$

The only evaluation edges which do not result in the rule $f_s(\text{ts}(s)) \rightarrow f_{s'}(\text{ts}(s'))$ are evaluations of `putfield` instructions. If `putfield` writes to the field of an object at reference o , then this could modify all objects at references o' with $o \searrow o'$.⁶ Therefore, in the right-hand side of the rule corresponding to `putfield`, we do not transform the reference o' to the variable o' , but to a fresh variable $\overline{o'}$. As an example consider state F in Fig. 9, where we write to a field of o_7 , and we have the annotation $o_1 \searrow o_7$. In the resulting rule, we therefore have the variable o_1 on the left-hand side, but a fresh variable $\overline{o_1}$ on the right-hand side. The terms corresponding to o_7 on the left- and right-hand side of the resulting rule describe the update of its field precisely (i.e., $\text{jIO}(\text{Tree}(\text{eoc}, i_1, o_{10}, o_{11}))$ is replaced by o_{11}).

Now let e be an instance edge from s to s' . Here we keep the information that we already have for the specialized state s (i.e., we keep $\text{ts}(s)$) and continue rewriting with the rules we already created for s' . So instead of $f_s(\text{ts}(s)) \rightarrow f_{s'}(\text{ts}(s'))$ we generate $f_s(\text{ts}(s)) \rightarrow f_{s'}(\text{ts}(s))$.

Finally, let e be a refinement edge from s to s' . So some abstract information in s is refined to more concrete information in s' (e.g., by refining $(c, ?)$ to `null`). These edges represent a case analysis and hence, some instances of s are also instances of s' , but others are no instances of s' . Note that by Lemma 3.5, if a state t is an instance of s' , then the term representation of s' matches t 's term representation. Hence, we can use pattern matching to perform the necessary case analysis. So instead of the rule $f_s(\text{ts}(s)) \rightarrow f_{s'}(\text{ts}(s'))$, we create a rule whose left-hand side only matches instances of s' , i.e., $f_s(\text{ts}(s')) \rightarrow f_{s'}(\text{ts}(s'))$. Consider for example the edge from B to C in Fig. 3. Any concrete state whose evaluation corresponds to this edge must have `null` at positions LV_0 and OS_0 . Thus, we create the rule $f_B(\text{null}, o_2, \text{null}) \rightarrow f_C(\text{null}, o_2, \text{null})$ which is only applicable to such states.

Recall that possibly cyclic data objects are translated to variables in Def. 3.3. Although variables are only instantiated by finite (non-cyclic) terms in term rewriting, our approach remains sound because states with possibly cyclic objects result in rules with extra variables on right-hand sides. For example, consider a simple list traversal algorithm. Here, we would have a state s where the local variable points to a reference o_1 with $o_1 = \text{IntList}(\text{value} = i_1, \text{next} = o_2)$ and in the successor state s' , the local variable would point to o_2 . Then, after refinement to $o_2 = \text{IntList}(\text{value} = i_2, \text{next} = o_3)$, there would be an instantiation edge back to s . For acyclic lists, this results in the rules $f_s(\text{jIO}(\text{IntList}(\text{eoc}, i_1, o_2))) \rightarrow f_{s'}(o_2)$, $f_{s'}(\text{jIO}(\text{IntList}(\text{eoc}, i_2, o_3))) \rightarrow f_{s''}(\text{jIO}(\text{IntList}(\text{eoc}, i_2, o_3)))$ and $f_{s''}(\text{jIO}(\text{IntList}(\text{eoc}, i_2, o_3))) \rightarrow f_s(\text{jIO}(\text{IntList}(\text{eoc}, i_2, o_3)))$ whose termination is easy to show. But if we had the annotations $o_1!$ and $o_2!$ in s , $o_2!$ in s' , and $o_2!$ and $o_3!$ in s'' , then we would obtain the rules $f_s(o_1) \rightarrow f_{s'}(o_2)$, $f_{s'}(o_2) \rightarrow f_{s''}(o_2)$ and $f_{s''}(o_2) \rightarrow f_s(o_2)$. So in the first rule, o_2 would be an extra variable representing an arbitrary list, and the resulting rules would not be terminating.

Definition 3.6 (Rewrite rules from termination graphs). Let there be an edge e from the state $s = (pp, l, op, h)$ to the state s' in a termination graph. Then we generate $\text{rule}(e)$:

- if e is an instance edge, then $\text{rule}(e) = f_s(\text{ts}(s)) \rightarrow f_{s'}(\text{ts}(s))$
- if e is a refinement edge, then $\text{rule}(e) = f_s(\text{ts}(s')) \rightarrow f_{s'}(\text{ts}(s'))$
- if e is an evaluation or split edge, we perform the following case analysis:

⁶Note that this is only possible if o' is UNKNOWN.

- if e is labeled by a statement of the form $o_1 = o_2 \oplus o_3$ where $\oplus \in \mathit{ArithOp}$, then $rule(e) = f_s(\text{ts}(s)) \rightarrow f_{s'}(\text{ts}(s'))\sigma$, where σ substitutes o_1 by $\text{tr}(s, o_2) \oplus \text{tr}(s, o_3)$
- if e is labeled by a condition $o_1 \oplus o_2$ where $\oplus \in \mathit{RelOp}$, then $rule(e) = f_s(\text{ts}(s)) \rightarrow f_{s'}(\text{ts}(s')) \mid \text{tr}(s, o_1) \oplus \text{tr}(s, o_2)$
- if pp is the instruction `putfield` writing to a field of reference o , then $rule(e) = f_s(\text{ts}(s)) \rightarrow f_{s'}(\text{ts}_o(s'))$, where $\text{ts}_o(s')$ is defined like $\text{ts}(s')$, but each reference o' with $o \searrow o'$ is transformed into a new fresh variable.
- for all other instructions, $rule(e) = f_s(\text{ts}(s)) \rightarrow f_{s'}(\text{ts}(s'))$

Our main theorem states that every computation path of the termination graph can be simulated by a rewrite sequence using the corresponding ITRS. Of course, the converse does not hold, i.e., our approach cannot be used to prove non-termination of **JBC** programs.

Theorem 3.7 (Proving termination of **JBC** by ITRSs). *If the ITRS corresponding to a termination graph is terminating, then the termination graph is terminating as well. Hence, then the original **JBC** program is also terminating for all concrete states t where $t \sqsubseteq s$ for some abstract state s in the termination graph.*

The resulting ITRSs are usually large, since they contain one rule for each edge of the termination graph. But since our ITRSs have a special form where the roots of all left- and right-hand sides are defined, where defined symbols do not occur below the roots, and where we only consider rewriting with normal substitutions, one can simplify the ITRSs substantially by *merging* their rules: Let \mathcal{R}_1 (resp. \mathcal{R}_2) be those rules in \mathcal{R} where the root of the right- (resp. left-)hand side is f . Then one can replace the rules $\mathcal{R}_1 \cup \mathcal{R}_2$ by the rules “ $\ell\sigma \rightarrow r'\sigma \mid b\sigma \ \&\& \ b'\sigma$ ” for all $\ell \rightarrow r \mid b \in \mathcal{R}_1$ and all $\ell' \rightarrow r' \mid b' \in \mathcal{R}_2$, where $\sigma = \text{mgu}(r, \ell')$. Of course, we also have to add rules for the Boolean conjunction “ $\&\&$ ”. Clearly, this process does not modify the termination behavior of \mathcal{R} . Moreover, it suffices to create rules only for those edges that occur in cycles of the termination graph.

With this simplification, we automatically obtain the following 1-rule ITRS for the `count` program. It increases the value i_1 of f_G 's first and third argument (corresponding to the value-field of `orig` and `copy`) as long as $i_1 < i_2$ (where i_2 is the value-field of `limit`).

$$\begin{array}{l} f_G(\text{jIO}(\text{Int}(\text{eoc}, i_1)), \text{jIO}(\text{Int}(\text{eoc}, i_2)), \text{jIO}(\text{Int}(\text{eoc}, i_1)), i_1, i_2) \rightarrow \\ f_G(\text{jIO}(\text{Int}(\text{eoc}, i_1 + 1)), \text{jIO}(\text{Int}(\text{eoc}, i_2)), \text{jIO}(\text{Int}(\text{eoc}, i_1 + 1)), i_1 + 1, i_2) \mid i_1 < i_2 \end{array}$$

For the `flatten` program we automatically obtain the following ITRS. To ease readability, we replaced every subterm “ $\text{jIO}(t)$ ” by just t , and we replaced “ $\text{TreeList}(\text{eoc}, v, n)$ ” by “ $\text{TL}(v, n)$ ”, “ $\text{Tree}(\text{eoc}, v, l, r)$ ” by “ $\text{T}(v, l, r)$ ”, and “ $\text{IntList}(\text{eoc}, v, n)$ ” by “ $\text{IL}(v, n)$ ”.

$$f_S(\text{TL}(\text{null}, o_9), \text{TL}(\text{null}, o_9), o_6) \rightarrow f_S(\text{TL}(\text{null}, o_9), o_9, o_6) \quad (3.1)$$

$$f_S(\text{TL}(\text{T}(i_1, o_{10}, o_{11}), o_9), \text{TL}(\text{T}(i_1, o_{10}, o_{11}), o_9), o_6) \rightarrow f_S(\text{TL}(o_{11}, o_9), \text{TL}(o_{10}, \text{TL}(o_{11}, o_9)), \text{IL}(i_1, o_6)) \quad (3.2)$$

$$f_S(o_1, \text{TL}(\text{null}, o_9), o_6) \rightarrow f_S(o_1, o_9, o_6) \quad (3.3)$$

$$f_S(o_1, \text{TL}(\text{T}(i_1, o_{10}, o_{11}), o_9), o_6) \rightarrow f_S(\overline{o_1}, \text{TL}(o_{10}, \text{TL}(o_{11}, o_9)), \text{IL}(i_1, o_6)) \quad (3.4)$$

Rules (3.1) and (3.3) correspond to the cycles from S over B' and over E . Their difference is whether 1 and c point to the same object in S (i.e., whether the first two arguments of f_S in the left-hand side are identical). But both handle the case where the first tree in the list c (i.e., in f_S 's second argument) is `null`. Then this null-tree is simply removed from the list and the result r (i.e., the third argument of f_S) does not change. The rules (3.2) and (3.4) correspond to the cycles from S over C' and over G . Here, the list c has the form $\text{TL}(\text{T}(i_1, o_{10}, o_{11}), o_9)$. Hence, the value i_1 of the first tree in the list is stored in the result

list (which is modified from o_6 to $\text{IL}(i_1, o_6)$) and the list \mathbf{c} is modified to $\text{TL}(o_{10}, \text{TL}(o_{11}, o_9))$. So the length of the list increases, but the number of nodes in the list decreases.

These examples illustrate that the ITRSs resulting from our automatic transformation of **JBC** are often very readable and constitute a natural representation of the original algorithm as a rewrite system. Not surprisingly, existing TRS techniques can easily prove termination of the resulting rules. For example, termination of the above ITRS for **flatten** is easily proved using a straightforward polynomial interpretation and dependency pairs. In contrast, abstraction-based tools like **Julia** and **COSTA** fail on examples like **flatten**. In fact, **Julia** and **COSTA** also fail on the **count** example from Sect. 2.2.

4. Experiments and Conclusion

We introduced an approach to prove termination of **JBC** programs automatically by first transforming them to termination graphs. Then an integer TRS is generated from the termination graph and existing TRS tools can be used to show its termination.

We implemented our approach in the termination prover **AProVE** [9] and evaluated it on the 106 non-recursive **JBC** examples from the *termination problem data base (TPDB)* used in the *International Termination Competition*.⁷ In our experiments, we removed one controversial example (“**overflow**”) from the TPDB whose termination depends on the treatment of integer overflows and we added the two examples **count** and **flatten** from this paper. Of these 106 examples, 10 are known to be non-terminating. See <http://aprove.informatik.rwth-aachen.de/eval/JBC> for the origins of the individual examples. As in the competition, we ran **AProVE** and the tools **Julia** [19] and **COSTA** [1] with a time limit of 60 seconds on each example. “Success” gives the number of examples where termination was proved, “Failure” means that the proof failed in less than 60 seconds, “Timeout” gives the number of examples where the tool took longer than 60 seconds, and “Runtime” is the average time (in s) needed per example. Note that for those examples from this collection where **AProVE** resulted in a timeout, the tool would also fail when using a longer timeout.

| | all 106 non-recursive examples | | | |
|---------------|--------------------------------|---------|---------|---------|
| | Success | Failure | Timeout | Runtime |
| AProVE | 89 | 5 | 12 | 14.3 |
| Julia | 74 | 32 | 0 | 2.6 |
| COSTA | 60 | 46 | 0 | 3.4 |

Our experiments show that for the problems in the current example collection, our rewriting-based approach in **AProVE** currently yields the most precise results. The main reason is that we do not use a fixed abstraction from data objects to integers, but represent objects as terms. On the other hand, this also explains the larger runtimes of **AProVE** compared to **Julia** and **COSTA**. Still, our approach is efficient enough to solve most examples in reasonable time. Our method benefits substantially from the representation of objects as terms, since afterwards arbitrary TRS termination techniques can be used to prove termination of the algorithms. Of course, while the examples in the TPDB are challenging, they are still quite small. Future work will be concerned with the application and adaption of our approach in order to use it also for large examples and **Java** libraries.

In the current paper, we restricted ourselves to **JBC** programs without recursion, whereas the approaches of **Julia** and **COSTA** also work on recursive programs. Of course,

⁷See http://www.termination-portal.org/wiki/Termination_Competition.

an extension of our method to recursive programs is another main point for future work. Our experiments also confirm the results at the *International Termination Competition* in December 2009, where the first competition on termination of **JBC** programs took place. Here, the three tools above were run on a random selection of the examples from the TPDB with similar results. To experiment with our implementation via a web interface and for details about the above experiments, we refer to <http://aprove.informatik.rwth-aachen.de/eval/JBC>.

Acknowledgement

We are indebted to the **Julia**- and the **COSTA**-team for their help with the experiments. We thank the anonymous reviewers for their valuable comments.

References

- [1] E. Albert, P. Arenas, M. Codish, S. Genaim, G. Puebla, and D. Zanardini. Termination analysis of **Java Bytecode**. In *Proc. FMOODS '08*, LNCS 5051, pages 2–18, 2008.
- [2] J. Berdine, B. Cook, D. Distefano, and P. O’Hearn. Automatic termination proofs for programs with shape-shifting heaps. In *Proc. CAV '06*, LNCS 4144, pages 386–400, 2006.
- [3] A. R. Bradley, Z. Manna, and H. B. Sipma. Termination of polynomial programs. In *Proc. VMCAI '05*, LNCS 3385, pages 113–129, 2005.
- [4] M. Colón and H. Sipma. Practical methods for proving program termination. In *Proc. CAV '02*, LNCS 2404, pages 442–454, 2002.
- [5] B. Cook, A. Podelski, and A. Rybalchenko. Termination proofs for systems code. In *Proc. PLDI '06*, pages 415–426. ACM Press, 2006.
- [6] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proc. POPL '77*, pages 238–252. ACM Press, 1977.
- [7] S. Falke and D. Kapur. A term rewriting approach to the automated termination analysis of imperative programs. In *Proc. CADE '09*, LNAI 5663, pages 277–293, 2009.
- [8] C. Fuhs, J. Giesl, M. Plücker, P. Schneider-Kamp, and S. Falke. Proving termination of integer term rewriting. In *Proc. RTA '09*, LNCS 5595, pages 32–47, 2009.
- [9] J. Giesl, P. Schneider-Kamp, and R. Thiemann. **AProVE 1.2**: Automatic termination proofs in the dependency pair framework. In *Proc. IJCAR '06*, LNAI 4130, pages 281–286, 2006.
- [10] J. Giesl, S. Swiderski, P. Schneider-Kamp, and R. Thiemann. Automated termination analysis for **Haskell**: From term rewriting to programming languages. In *RTA '06*, LNCS 4098, pp. 297–312, 2006.
- [11] J. Gosling, B. Joy, G. Steele, and G. Bracha. *The Java Language Specification*. Addison Wesley, 2005.
- [12] S. Gulwani, K. Mehra, and T. Chilimbi. **SPEED**: Precise and efficient static estimation of program computational complexity. In *Proc. POPL '09*, pages 127–139. ACM Press, 2009.
- [13] G. Klein and T. Nipkow. A machine-checked model for a **Java**-like language, virtual machine, and compiler. *ACM Transactions on Programming Languages and Systems*, 28(4):619–695, 2006.
- [14] T. Lindholm and F. Yellin. *Java Virtual Machine Specification*. Prentice Hall, 1999.
- [15] M. T. Nguyen, D. De Schreye, J. Giesl, and P. Schneider-Kamp. **Polytool**: Polynomial interpretations as a basis for termination analysis of logic programs. *Theory and Practice of Logic Programming*, 2010. To appear. Available from <http://arxiv.org/pdf/0912.4360>.
- [16] C. Otto, M. Brockschmidt, C. von Essen, and J. Giesl. Automated termination analysis of **Java Bytecode** by term rewriting. Technical Report AIB-2010-08, RWTH Aachen, 2010. <http://aib.informatik.rwth-aachen.de>.
- [17] P. Schneider-Kamp, J. Giesl, A. Serebrenik, and R. Thiemann. Automated termination proofs for logic programs by term rewriting. *ACM Transactions on Computational Logic*, 11(1), Article 2, 2009.
- [18] M. H. Sørensen and R. Glück. An algorithm of generalization in positive supercompilation. In *Proc. ILPS '95*, pages 465–479. MIT Press, 1995.
- [19] F. Spoto, F. Mesnard, and É. Payet. A termination analyser for **Java Bytecode** based on path-length. *ACM Transactions on Programming Languages and Systems*, 32(3), Article 8, 2010.