

Abstracting an operational semantics to finite automata

Nadezhda Baklanova, Wilmer Ricciotti, Jan-Georg Smaus, Martin Strecker

IRIT (Institut de Recherche en Informatique de Toulouse)
Université de Toulouse, France
firstname.lastname@irit.fr **

Abstract. There is an apparent similarity between the descriptions of small-step operational semantics of imperative programs and the semantics of finite automata, so defining an abstraction mapping from semantics to automata and proving a simulation property seems to be easy. This paper aims at identifying the reasons why simple proofs break, among them artifacts in the semantics that lead to stuttering steps in the simulation. We then present a semantics based on the zipper data structure, with a direct interpretation of evaluation as navigation in the syntax tree. The abstraction function is then defined by equivalence class construction.

Keywords: Programming language semantics; Abstraction; Finite Automata; Formal Methods; Verification

Key Terms: FormalMethod, VerificationProcess

1 Introduction

Among the formalisms employed to describe the semantics of transition systems, two particularly popular choices are abstract machines and structural operational semantics (SOS). Abstract machines are widely used for modeling and verifying dynamic systems, e.g. finite automata, Büchi automata or timed automata [9,4,1]. An abstract machine can be represented as a directed graph with transition semantics between nodes. The transition semantics is defined by moving a pointer to a current node. Automata are a popular tool for modeling dynamic systems due to the simplicity of the verification of automata systems, which can be carried out in a fully automated way, something that is not generally possible for Turing-complete systems.

This kind of semantics is often extended by adding a background state composed of a set of variables with their values: this is the case of timed automata, which use background clock variables [2]. The UPPAAL model checker for timed

* N. Baklanova and M. Strecker were partially supported by the project *Verisync* (ANR-10-BLAN-0310).

** W. Ricciotti and J.-G. Smaus are supported by the project AJITPROP of the Fondation Airbus.

automata extends the notion of background state even further by adding integer and Boolean variables to the state [7] which, however, do not increase the computational power of such timed automata but make them more convenient to use.

Another formalism for modeling transition systems is structural semantics (“small-step”, contrary to “big-step” semantics which is much easier to handle but which is inappropriate for a concurrent setting), which uses a set of reduction rules for simplifying a program expression. It has been described in detail in [14] and used, for example, for the Jinja project developing a formal model of the Java language [10]. An appropriate semantic rule for reduction is selected based on the expression pattern and on values of some variables in a state. As a result of reduction the expression and the state are updated.

$$\frac{s' = s(v \mapsto \text{eval } \text{expr } s)}{(\text{Assign } v \text{ expr}, s) \rightarrow (\text{Unit}, s')} \quad [\text{ASSIGNMENT}]$$

This kind of rules is intuitive; however, the proofs involving them require induction over the expression structure. A different approach to writing a structural semantics was described in [3,12] for the CMinor language. It uses a notion of continuation which represents an expression as a control stack and deals with separate parts of the control stack consecutively.

$$(\text{Seq } e1 \ e2 \cdot \kappa, s) \rightarrow (e1 \cdot e2 \cdot \kappa, s) \quad (\text{Empty} \cdot \kappa, s) \rightarrow (\kappa, s)$$

Here the “.” operator designates concatenation of control stacks. The semantics of continuations does not need induction over the expression, something which makes proof easier; however it requires more auxiliary steps for maintaining the control stack which do not have direct correspondance in the modeled language.

For modeling non-local transfer of control, Krebbers and Wiedijk [11] present a semantics using (non-recursive) “statement contexts”. These are combined with the above-mentioned continuation stacks. The resulting semantics is situated mid-way between [3] and the semantics proposed below.

The present paper describes an approach to translation from structural operational semantics to finite automata extended with background state. All the considered automata are an extension of Büchi automata with background state, i.e. they have a finite number of nodes and edges but can produce an infinite trace. The reason of our interest in abstracting from structural semantics to Büchi automata is our work in progress [6]. We are working on a static analysis algorithm for finding possible resource sharing conflicts in multithreaded Java programs. For this purpose we annotate Java programs with timing information and then translate them to a network of timed automata which is later model checked. The whole translation is formally verified. One of the steps of the translation procedure includes switching from structural operational semantics of a Java-like language to automata semantics. During this step we discovered some problems which we will describe in the next section. The solutions we propose

extend well beyond the problem of abstracting a structured language to an automaton. It can also be used for compiler verification, which usually is cluttered up with arithmetic address calculation that can be avoided in our approach.

The contents of the paper has been entirely formalized in the Isabelle proof assistant [13]. We have not insisted on any Isabelle-specific features, therefore this formalization can be rewritten using other proof assistants. The full Isabelle formal development can be found on the web [5].

2 Problem Statement

We have identified the following as the main problems when trying to prove the correctness of the translation between a programming language semantics and its abstraction to automata:

1. Preservation of execution context: an abstract machine always sees all the available nodes while a reduced expression loses the information about previous reductions.
2. Semantic artifacts: some reduction rules are necessary for the functionality of the semantics, but may be missing in the modeled language. Additionally, the rules can produce expressions which do not occur in the original language.

These problems occur independently of variations in the presentation of semantic rules [14] adopted in the literature, such as [10] (recursive evaluation of sub-statements) or [3,12] (continuation-style).

We will describe these two problems in detail, and later our approach to their solution, in the context of a minimalistic programming language which only manipulates Boolean values (a *Null* value is also added to account for errors):

datatype *val* = *Bool bool* | *Null*

The language can be extended in a rather straightforward way to more complex expressions. In this language, expressions are either values or variables:

datatype *expr* = *Val val* | *Var vname*

The statements are those of a small imperative language:

datatype *stmt* =
 Empty — no-op
 | *Assign vname val* — assignment: *var := val*
 | *Seq stmt stmt* — sequence: *c₁; c₂*
 | *Cond expr stmt stmt* — conditional: if *e* then *c₁* else *c₂*
 | *While expr stmt* — loop: while *e* do *c*

2.1 Preservation of execution context

Problem 1 concerns the loss of an execution context through expression reductions which is a design feature of structural semantics. Let us consider a simple example.

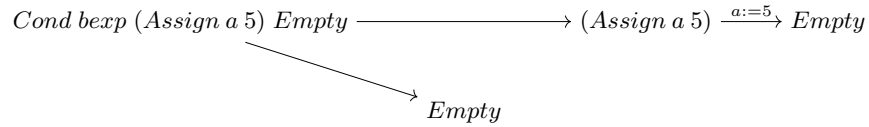
Assume we have a structural semantics for our minimal imperative language (some rules of a traditional presentation are shown in Figure 1): we want to translate a program written in this language into an abstract machine. Assume that the states of variable values have the same representation in the two systems: this means we only need to translate the program expression into a directed graph with different nodes corresponding to different expressions obtained by reductions of the initial program expression.

$$\begin{array}{c}
 \frac{s' = s(v \mapsto \text{eval } \text{expr } s)}{(\text{Assign } v \text{ expr}, s) \rightarrow (\text{Empty}, s')} \quad [\text{ASSIGN}] \\
 \\
 \frac{\text{eval } \text{bexp } s = \text{True}}{(\text{Cond } \text{bexp } e1 \ e2, s) \rightarrow (e1, s)} \quad [\text{CONDT}] \quad \frac{\text{eval } \text{bexp } s = \text{False}}{(\text{Cond } \text{bexp } e1 \ e2, s) \rightarrow (e2, s)} \quad [\text{CONDF}]
 \end{array}$$

Fig. 1. Semantic rules for the minimal imperative language.

On the abstract machine level the *Assign* statements would be represented as two-state automata, and the *Cond* as a node with two outgoing edges directed to the automata for the bodies of its branches.

Consider a small program in this language *Cond bexp (Assign a 5) Empty* and its execution flow.



The execution can select any of the two branches depending on the *bexp* value. There are two different *Empty* expressions appearing as results of two different reductions. The corresponding abstract machine would be a natural graph representation for a condition statement with two branches (Figure 2).

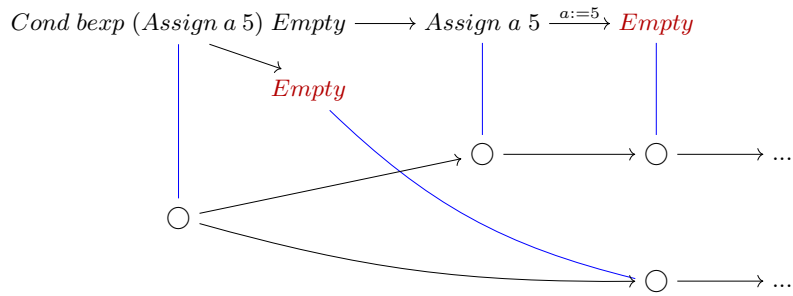


Fig. 2. The execution flow and the corresponding abstract machine for the program *Cond bexp (Assign a 5) Empty*.

During the simple generation of an abstract machine from a program expression the two *Empty* statements cannot be distinguished although they should be mapped into two different nodes in the graph. We need to add more information about the context into the translation, and it can be done by different ways.

A straightforward solution would be to add some information in order to distinguish between the two *Empty* expressions. If we add unique identifiers to each subexpression of the program, they will allow to know exactly which subexpression we are translating (Figure 3). The advantage of this approach is its simplicity, however, it requires additional functions and proofs for identifier management.

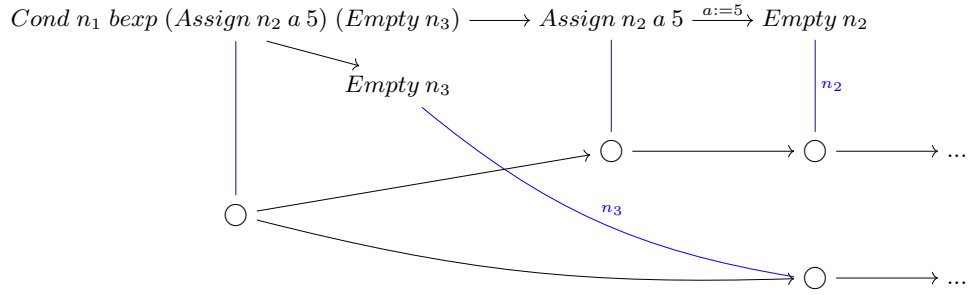


Fig. 3. The execution flow and the corresponding abstract machine for the program with subexpression identifiers $Cond\ n_1\ bexp\ (Assign\ n_2\ a\ 5)\ (Empty\ n_3)$.

Another solution for the problem proposed in this paper involves usage of a special data structure to keep the context of the translation. There are known examples of translations from subexpression-based semantics [10] and continuation-based semantics [12] to abstract machines. However, all these translations do not address the problem of context preservation during the translation.

2.2 Semantic artifacts

The second problem appears because of the double functionality of the *Empty* expression: it is used to define an empty operator which does nothing as well as the final expression for reductions which cannot be further reduced. The typical semantic rules for a sequence of expressions look as shown on Figure 4.

$$\frac{(e1, s) \rightarrow (e1', s')}{(Seq\ e1\ e2, s) \rightarrow (Seq\ e1'\ e2, s')} \text{ [SEQ1]} \quad \frac{}{(Seq\ Empty\ e2, s) \rightarrow (e2, s)} \text{ [SEQ2]}$$

Fig. 4. Semantic rules for the sequence of two expressions.

Here the *Empty* expression means that the first expression in the sequence has been reduced up to the end, and we can start reducing the second expression. However, any imperative language translated to an assembly language would not have an additional operator between the two pieces of code corresponding to the first and the second expressions. The rule SEQ2 must be marked as a silent transition when translated to an automaton, or the semantic rules have to be changed.

3 Zipper-based semantics of imperative programs

3.1 The zipper data structure

Our plan is to propose an alternative technique to formalize operational semantics that will make it easier to preserve the execution context during the translation to an automata-based formalism. Our technique is built around a zipper data structure, whose purpose is to identify a location in a tree (in our case: a *stmt*) by the subtree below the location and the rest of the tree (in our case: of type *stmt-path*). In order to allow for an easy navigation, the rest of the tree is turned inside-out so that it is possible to reach the root of the tree by following the backwards pointers. The following definition is a straightforward adaptation of the zipper for binary trees discussed in [8] to the *stmt* data type:

```
datatype stmt-path =
  PTop
| PSeqLeft stmt-path stmt          | PSeqRight stmt stmt-path
| PCondLeft expr stmt-path stmt | PCondRight expr stmt stmt-path
| PWhile expr stmt-path
```

Here, *PTop* represents the root of the original tree, and for each constructor of *stmt* and each of its sub-*stmts*, there is a “hole” of type *stmt-path* where a subtree can be fitted in. A location in a tree is then a combination of a *stmt* and a *stmt-path*:

```
datatype stmt-location = Loc stmt stmt-path
```

Given a location in a tree, the function *reconstruct* reconstructs the original tree $reconstruct :: stmt \Rightarrow stmt-path \Rightarrow stmt$, and *reconstruct-loc* ($Loc\ c\ sp$) = *reconstruct* *c* *sp* does the same for a location.

```
fun reconstruct :: stmt  $\Rightarrow$  stmt-path  $\Rightarrow$  stmt where
  reconstruct c PTop = c
| reconstruct c (PSeqLeft sp c2) = reconstruct (Seq c c2) sp
| reconstruct c (PSeqRight c1 sp) = reconstruct (Seq c1 c) sp
| reconstruct c (PCondLeft e sp c2) = reconstruct (Cond e c c2) sp
| reconstruct c (PCondRight e c1 sp) = reconstruct (Cond e c1 c) sp
| reconstruct c (PWhile e sp) = reconstruct (While e c) sp
```

```
fun reconstruct-loc :: stmt-location  $\Rightarrow$  stmt where
  reconstruct-loc (Loc c sp) = reconstruct c sp
```

3.2 Semantics

Our semantics is a small-step operational semantics describing the effect of the execution a program on a certain program state. For each variable, the state yields *Some* value associated with the variable, or *None* if the variable is unassigned. More formally, the state is a mapping $vname \Rightarrow val\ option$. Defining the evaluation of an expression in a state is then standard.

Before commenting the rules of our semantics, let us discuss which kind of structure we are manipulating. The semantics essentially consists in moving around a pointer within the syntax tree. As explained in Section 3.1, a position in the syntax tree is given by a *stmt-location*. However, during the traversal of the syntax tree, we visit each position at least twice (and possibly several times, for example in a loop): before executing the corresponding statement, and after finishing the execution. We therefore add a Boolean flag, where *True* is a marker for “before” and *False* for “after” execution.

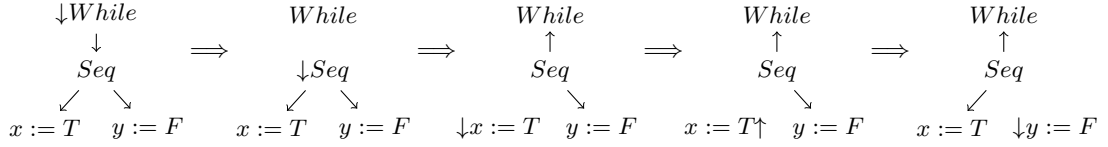


Fig. 5. Example of execution of small-step semantics

As an example, consider the execution sequence depicted in Figure 5 (with assignments written in a more readable concrete syntax), consisting of the initial steps of the execution of the program $While\ (e,\ Seq(x := T,\ y := F))$. The before (resp. after) marker is indicated by a downward arrow before (resp. an upward arrow behind) the current statement. The condition of the loop is omitted because it is irrelevant here. The middle configuration would be coded as $((Loc\ (x := T)\ (PSeqLeft\ (PWhile\ e\ PTop)\ (y := F))),\ True)$.

Altogether, we obtain a syntactic configuration (*synt-config*) which combines the location and the Boolean flag. The semantic configuration (*sem-config*) manipulated by the semantics adjoins the *state*, as defined previously.

type-synonym $synt\ config = stmt\ location \times bool$

type-synonym $sem\ config = synt\ config \times state$

The rules of the small-step semantics of Figure 7 fall into two categories: before execution of a statement s (of the form $((l,\ True), s)$) and after execution (of the form $((l,\ False), s)$); there is only one rule of this latter kind: SFALSE.

Let us comment on the rules in detail:

- SEMPTY executes the *Empty* statement just by swapping the Boolean flag.
- SASSIGN is similar, but it also updates the state for the assigned variable.

fun *next-loc* :: *stmt* \Rightarrow *stmt-path* \Rightarrow (*stmt-location* \times *bool*) **where**
next-loc *c* *P*Top = (*Loc* *c* *P*Top, *False*)
| *next-loc* *c* (*P*SeqLeft *sp* *c*₂) = (*Loc* *c*₂ (*P*SeqRight *c* *sp*), *True*)
| *next-loc* *c* (*P*SeqRight *c*₁ *sp*) = (*Loc* (*Seq* *c*₁ *c*) *sp*, *False*)
| *next-loc* *c* (*P*CondLeft *e* *sp* *c*₂) = (*Loc* (*Cond* *e* *c* *c*₂) *sp*, *False*)
| *next-loc* *c* (*P*CondRight *e* *c*₁ *sp*) = (*Loc* (*Cond* *e* *c*₁ *c*) *sp*, *False*)
| *next-loc* *c* (*P*While *e* *sp*) = (*Loc* (*While* *e* *c*) *sp*, *True*)

Fig. 6. Finding the next location

$$\frac{}{((\text{Loc Empty } sp, \text{True}), s) \rightarrow ((\text{Loc Empty } sp, \text{False}), s)} \text{[SEMPY]}$$

$$\frac{}{((\text{Loc (Assign } vr \ vl) \ sp, \text{True}), s) \rightarrow ((\text{Loc (Assign } vr \ vl) \ sp, \text{False}), s(vr \mapsto vl))} \text{[SASSIGN]}$$

$$\frac{}{((\text{Loc (Seq } c_1 \ c_2) \ sp, \text{True}), s) \rightarrow ((\text{Loc } c_1 \ (\text{PSeqLeft } sp \ c_2), \text{True}), s)} \text{[SSEQ]}$$

$$\frac{\text{eval } e \ s = \text{Bool True}}{((\text{Loc (Cond } e \ c_1 \ c_2) \ sp, \text{True}), s) \rightarrow ((\text{Loc } c_1 \ (\text{PCondLeft } e \ sp \ c_2), \text{True}), s)} \text{[SCONDT]}$$

$$\frac{\text{eval } e \ s = \text{Bool False}}{((\text{Loc (Cond } e \ c_1 \ c_2) \ sp, \text{True}), s) \rightarrow ((\text{Loc } c_2 \ (\text{PCondRight } e \ c_1 \ sp), \text{True}), s)} \text{[SCONDF]}$$

$$\frac{\text{eval } e \ s = \text{Bool True}}{((\text{Loc (While } e \ c) \ sp, \text{True}), s) \rightarrow ((\text{Loc } c \ (\text{PWhile } e \ sp), \text{True}), s)} \text{[SWHILET]}$$

$$\frac{\text{eval } e \ s = \text{Bool False}}{((\text{Loc (While } e \ c) \ sp, \text{True}), s) \rightarrow ((\text{Loc (While } e \ c) \ sp, \text{False}), s)} \text{[SWHILEF]}$$

$$\frac{\text{sp} \neq \text{PTop}}{((\text{Loc } c \ sp, \text{False}), s) \rightarrow (\text{next-loc } c \ sp, s)} \text{[SFALSE]}$$

Fig. 7. Small-step operational semantics

- SSEQ moves the pointer to the substatement c_1 , pushing the substatement c_2 as continuation to the statement path.
- SCONDT and SCONDF move to the *then*- respectively *else*- branch of the conditional, depending on the value of the condition.
- SWHILET moves to the body of the loop.
- SWHILEF declares the execution of the loop as terminated, by setting the Boolean flag to *False*.

- SFALSE comes into play when execution of the current statement is finished. We then move to the next location, provided we have not already reached the root of the syntax tree and the whole program terminates.

The move to the next relevant location is accomplished by function *next-loc* (Figure 6) which intuitively works as follows: upon conclusion of the first substatement in a sequence, we move to the second substatement. When finishing the body of a loop, we move back to the beginning of the loop. In all other cases, we move up the syntax tree, waiting for rule SFALSE to relaunch the function.

4 Target language: Automata

4.1 Syntax

As usual, our automata are a collection of nodes and edges, with a distinguished initial state. In this general definition, we will keep the node type *'n* abstract. It will later be instantiated to *synt-config*. An edge connects two nodes; moving along an edge may trigger an assignment to a variable (*AssAct*), or have no effect at all (*NoAct*).

An automaton *'n ta* is a record consisting of a set of *nodes*, a set of *edges* and an initial node *init-s*. An edge has a *source* node, an *action* and a destination node *dest*. Components of a record are written between (| ... |).

4.2 Semantics

An automaton state is a node, together with a *state* as in Section 3.2.

type-synonym *'n ta-state* = *'n * state*

Executing a step of an automaton in an automaton state (*l, s*) consists of selecting an edge starting in node *l*, moving to the target of the edge and executing its action. Automata are non-deterministic; in this simplified model, we have no guards for selecting edges.

$$\frac{e \in \text{set}(\text{edges } aut) \quad l = \text{source } e \quad l' = \text{dest } e \quad s' = \text{action-effect}(\text{action } e) s}{aut \vdash (l, s) \rightarrow (l', s')} \text{ [ACTION]}$$

5 Automata construction

The principle of abstracting a statement to an automaton is simple; the novelty resides in the way the automaton is generated via the zipper structure: as nodes, we choose the locations of the statements (with their Boolean flags), and as edges all possible transitions of the semantics.

To make this precise, we need some auxiliary functions. We first define a function *all-locations* of type $stmt \Rightarrow stmt\text{-path} \Rightarrow stmt\text{-location list}$ which gathers all locations in a statement, and a function *nodes-of-stmt-locations* which adds the Boolean flags.

As for the edges, the function *synt-step-image* yields all possible successor configurations for a given syntactic configuration. This is of course an over-approximation of the behavior of the semantics, since some of the source tree locations may be unreachable during execution.

```
fun synt-step-image :: synt-config  $\Rightarrow$  synt-config list where
  synt-step-image (Loc Empty sp, True) = [(Loc Empty sp, False)]
| synt-step-image (Loc (Assign vr vl) sp, True) = [(Loc (Assign vr vl) sp, False)]
| synt-step-image (Loc (Seq c1 c2) sp, True) = [(Loc c1 (PSeqLeft sp c2), True)]
| synt-step-image (Loc (Cond e c1 c2) sp, True) =
  [(Loc c1 (PCondLeft e sp c2), True), (Loc c2 (PCondRight e c1 sp), True)]
| synt-step-image (Loc (While e c) sp, True) =
  [(Loc c (PWhile e sp), True), (Loc (While e c) sp, False)]
| synt-step-image (Loc c sp, False) = (if sp = PTop then [] else [next-loc c sp])
```

Together with the following definitions:

```
fun action-of-synt-config :: synt-config  $\Rightarrow$  action where
  action-of-synt-config (Loc (Assign vn vl) sp, True) = AssAct vn vl
| action-of-synt-config (Loc c sp, b) = NoAct
```

```
definition edge-of-synt-config :: synt-config  $\Rightarrow$  synt-config edge list where
  edge-of-synt-config s =
  map( $\lambda$  t. ( $\downarrow$ source = s, action = action-of-synt-config s, dest = t))(synt-step-image s)
```

```
definition edges-of-nodes :: synt-config list  $\Rightarrow$  synt-config edge list where
  edges-of-nodes nds = concat (map edge-of-synt-config nds)
```

we can define the translation function from statements to automata:

```
fun stmt-to-ta :: stmt  $\Rightarrow$  synt-config ta where
  stmt-to-ta c =
  (let nds = nodes-of-stmt-locations (all-locations c PTop) in
  ( $\downarrow$  nodes = nds, edges = edges-of-nodes nds, init-s = ((Loc c PTop), True)  $\downarrow$ ))
```

6 Simulation Property

We recall that the nodes of the automaton generated by *stmt-to-ta* are labeled by configurations (location, Boolean flag) of the syntax tree. The simulation lemma (Lemma 1) holds for automata with appropriate closure properties: a successor configuration wrt. a transition of the semantics is also a label of the automaton (*nodes-closed*), and analogously for edges (*edges-closed*) or both nodes and edges (*synt-step-image-closed*).

The simulation statement is a typical commuting-diagram property: a step of the program semantics can be simulated by a step of the automaton semantics,

for corresponding program and automata states. For this correspondence, we use the notation \approx , even though it is just plain syntactic equality in our case.

Lemma 1 (Simulation property).

Assume that *synt-step-image-closed aut* and $((lc, b), s) \approx ((lca, ba), sa)$. If $((lc, b), s) \rightarrow ((lc', b'), s')$, then there exist lca', ba', sa' such that $(lca', ba') \in \text{set}(\text{nodes aut})$ and the automaton performs the same transition: $\text{aut} \vdash ((lca, ba), sa) \rightarrow ((lca', ba'), sa')$ and $((lc', b'), s') \approx ((lca', ba'), sa')$.

The proof is a simple induction over the transition relation of the program semantics and is almost fully automatic in the Isabelle proof assistant.

We now want to get rid of the precondition *synt-step-image-closed aut* in Lemma 1. The first subcase (edge closure), is easy to prove. Node closure is more difficult and requires the following key lemma:

Lemma 2.

If $lc \in \text{set}(\text{all-locations } c \text{ PTop})$ then $\text{set}(\text{map fst}(\text{synt-step-image}(lc, b))) \subseteq \text{set}(\text{all-locations } c \text{ PTop})$.

With this, we obtain the desired

Lemma 3 (Closure of automaton). *synt-step-image-closed (stmt-to-ta c)*

For the proofs, see [5].

Let us combine the previous results and write them more succinctly, by using the notation \rightarrow^* for the reflexive-transitive closure for the transition relations of the small-step semantics and the automaton. Whenever a state is reachable by executing a program c in its initial configuration, then a corresponding (\approx) state is reachable by running the automaton generated with function *stmt-to-ta*:

Theorem 1.

If $((\text{Loc } c \text{ PTop}, \text{True}), s) \rightarrow^* (cf', s')$ then $\exists cfa' sa'. \text{stmt-to-ta } c \vdash (\text{init-s}(\text{stmt-to-ta } c), s) \rightarrow^* (cfa', sa') \wedge (cf', s') \approx (cfa', sa')$.

Obviously, the initial configuration of the semantics and the automaton are in the simulation relation \approx , and for the inductive step, we use Lemma 1.

7 Conclusions

This paper has presented a new kind of small-step semantics for imperative programming languages, based on the zipper data structure. Our primary aim is to show that this semantics has decisive advantages for abstracting programming language semantics to automata. Even if the generated automata have a great number of silent transitions, these can be removed.

We are currently in the process of adopting this semantics in a larger formalization from Java to Timed Automata [6]. As most constructs (zipper data

structure, mapping to automata) are generic, we think that this kind of semantics could prove useful for similar formalizations with other source languages. The proofs (here carried out with the Isabelle proof assistant) have a pleasingly high degree of automation that are in sharp contrast with the index calculations that are usually required when naming automata states with numbers.

Renaming nodes from source tree locations to numbers is nevertheless easy to carry out, see the code snippet provided on the web page [5] of this paper. For these reasons, we think that the underlying ideas could also be useful in the context of compiler verification, when converting a structured source program to a flow graph with basic blocs, but before committing to numeric values of jump targets.

References

1. Rajeev Alur, Costas Courcoubetis, and David L. Dill. Model-checking for real-time systems. In *LICS*, pages 414–425. IEEE Computer Society, 1990.
2. Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:183–235, 1994.
3. Andrew W. Appel and Sandrine Blazy. Separation logic for small-step cminor. In *Theorem Proving in Higher Order Logics, 20th int. conf. TPHOLS*, pages 5–21. Springer, 2007.
4. Ch. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
5. Nadezhda Baklanova, Wilmer Ricciotti, Jan-Georg Smaus, and Martin Strecker. Abstracting an operational semantics to finite automata (formalization), 2014. https://bitbucket.org/Martin_Strecker/abstracting_op_sem_to_automata.
6. Nadezhda Baklanova and Martin Strecker. Abstraction and verification of properties of a Real-Time Java. In *Proc. ICTERI*, volume 347 of *Communications in Computer and Information Science*, pages 1–18. Springer, 2013.
7. Johan Bengtsson and Wang Yi. Timed automata: Semantics, algorithms and tools. In *Lectures on Concurrency and Petri Nets*, LNCS, pages 87–124. Springer, 2004.
8. Gérard Huet. Functional pearl: The zipper. *Journal of Functional Programming*, 7(5):549–554, September 1997.
9. Bakhadyr Khossainov and Anil Nerode. *Automata Theory and Its Applications*. Birkhauser Boston, 2001.
10. Gerwin Klein and Tobias Nipkow. A machine-checked model for a Java-like language, virtual machine, and compiler. *ACM Trans. Program. Lang. Syst.*, 28:619–695, July 2006.
11. Robbert Krebbers and Freek Wiedijk. Separation logic for non-local control flow and block scope variables. In Frank Pfenning, editor, *Foundations of Software Science and Computation Structures*, volume 7794 of *Lecture Notes in Computer Science*, pages 257–272. Springer Berlin Heidelberg, 2013.
12. Xavier Leroy. A formally verified compiler back-end. *Journal of Automated Reasoning* 43(4), 43(4), 2009.
13. Tobias Nipkow, Lawrence Paulson, and Markus Wenzel. *Isabelle/HOL. A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
14. Glynn Winskel. *The Formal Semantics of Programming Languages: An Introduction*. MIT Press, Cambridge, MA, USA, 1993.