

# A Mechanised Proof of the Time Invariance Thesis for the Weak Call-By-Value $\lambda$ -Calculus

Yannick Forster ✉ 

Saarland University, Saarland Informatics Campus, Saarbrücken, Germany

Fabian Kunze ✉

Saarland University, Saarland Informatics Campus, Saarbrücken, Germany

Gert Smolka ✉

Saarland University, Saarland Informatics Campus, Saarbrücken, Germany

Maximilian Wuttke ✉

Saarland University, Saarland Informatics Campus, Saarbrücken, Germany

---

## Abstract

---

The weak call-by-value  $\lambda$ -calculus  $L$  and Turing machines can simulate each other with a polynomial overhead in time. This time invariance thesis for  $L$ , where the number of  $\beta$ -reductions of a computation is taken as its time complexity, is the culmination of a 25-years line of research, combining work by Blelloch, Greiner, Dal Lago, Martini, Accattoli, Forster, Kunze, Roth, and Smolka. The present paper presents a mechanised proof of the time invariance thesis for  $L$ , constituting the first mechanised equivalence proof between two standard models of computation covering time complexity.

The mechanisation builds on an existing framework for the extraction of Coq functions to  $L$  and contributes a novel Hoare logic framework for the verification of Turing machines.

The mechanised proof of the time invariance thesis establishes  $L$  as model for future developments of mechanised computational complexity theory regarding time. It can also be seen as a non-trivial but elementary case study of time-complexity-preserving translations between a functional language and a sequential machine model. As a by-product, we obtain a mechanised many-one equivalence proof of the halting problems for  $L$  and Turing machines, which we contribute to the Coq Library of Undecidability Proofs.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Computability; Theory of computation  $\rightarrow$  Type theory

**Keywords and phrases** formalizations of computational models, computability theory, Coq, time complexity, Turing machines, lambda calculus, Hoare logic

**Digital Object Identifier** 10.4230/LIPIcs.ITP.2021.19

**Supplementary Material** *Software (Code Repository)*: <https://github.com/uds-psl/time-invariance-thesis-for-L>; archived at [swh:1:dir:cc13c3c6279985ae15903a24e23bd5d356b6f435](https://swh.io/1/dir/cc13c3c6279985ae15903a24e23bd5d356b6f435)

**Acknowledgements** We want to thank Lennard Gäher and Dominik Kirst for valuable feedback on drafts of this paper.

## 1 Introduction

Computability theory – i.e. the study of which problems can be solved computationally – is invariant under the chosen model of computation: Any Turing-complete model does the job. Similarly, but less generally, computational complexity theory – i.e. the study of how efficiently problems can be solved computationally – is invariant under the chosen model: Both Turing machines and RAM machines are used and can be exchanged as long as only complexity classes closed under polynomial time reductions are of interest. This practice is based on the *invariance thesis*, introduced by Slot and van Emde Boas as “all reasonable models of computation simulate each other with a polynomially bounded overhead in time and a constant factor overhead in space” [23]. For the present paper, we dub the first half concerning time complexity the *time invariance thesis*.



© Yannick Forster, Fabian Kunze, Gert Smolka, and Maximilian Wuttke; licensed under Creative Commons License CC-BY 4.0

12th International Conference on Interactive Theorem Proving (ITP 2021).

Editors: Liron Cohen and Cezary Kaliszyk; Article No. 19; pp. 19:1–19:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

While RAM machines and Turing machines are reasonable in this sense, the question how and whether the (untyped)  $\lambda$ -calculus [4] is reasonable was a long-standing open question. This may sound especially surprising given the fact that the  $\lambda$ -calculus was one of the first models for computability to be proven Turing-complete, by Turing himself [27]. On the other hand, the  $\lambda$ -calculus is indeed more complicated than sequential models: Terms are trees and contain binders, computation is non-local, multiple potentially non-equivalent reduction strategies can be chosen, etc. Maybe most crucially, it is not obvious what a reasonable time complexity measure is: The number of  $\beta$ -steps in a computation? Or does one have to account for (the size of)  $\beta$ -redexes? And even more unclear: What is the space complexity of a  $\lambda$ -calculus computation?

More or less independently of these questions, one direction of the time invariance thesis is easy to prove: The  $\lambda$ -calculus can simulate Turing machines.

In this paper, we focus on the weak call-by-value  $\lambda$ -calculus introduced by Plotkin [21], in the concrete variant L introduced by Forster and Smolka [12]. In L, only abstractions are values, and we take the number of  $\beta$ -steps as time complexity measure of a computation [7]. L is similar to an ML-like functional programming language: Reductions in abstractions are not allowed, and a  $\beta$ -reduction only applies when both sides of an application are a value.

A short timeline of time complexity for L and the full  $\lambda$ -calculus looks as follows:

- 1995** Blelloch and Greiner [3] prove the time invariance thesis for the weak call-by-value  $\lambda$ -calculus w.r.t. RAM machines, with the number of  $\beta$ -steps as time complexity measure.
- 2008** Dal Lago and Martini [17] prove the time invariance thesis for the weak call-by-value  $\lambda$ -calculus w.r.t. Turing machines, with the sum over all differences of the size of the redex and the size of the reduct as time complexity measure.
- 2015** Accattoli and Dal Lago [1] prove the time invariance thesis for the full  $\lambda$ -calculus w.r.t. RAM machines, with the number of  $\beta$ -steps as time complexity measure.
- 2020** Forster, Kunze, and Roth [7] prove the invariance thesis for L w.r.t. Turing machines, with the number of  $\beta$ -steps as time complexity measure and the maximum of the sizes of all intermediate terms as space complexity measure.

All of these results come with different levels of formality in their proofs. Blelloch and Greiner [3] only *state* their result, but do not give any details of a proof, much less provide details how the (RAM)-machine carrying out the simulation might look like. Dal Lago and Martini [17] give a *proof sketch*: They informally describe what a Turing machine simulating L has to do, and for instance how many tapes it has, but do not give invariants for inductions. The other, easier simulation direction is discussed in similar detail. In a later note, Dal Lago and Accattoli [16] give a *formal* proof on paper that the  $\lambda$ -calculus – independent from the concrete variant used – can simulate Turing machines with linear overhead. The proof of the time invariance thesis for the full  $\lambda$ -calculus by Accattoli and Dal Lago [1] is based on a careful and *formal* study of explicit sharing, but again the implementation in terms of machines is left out. This technique is omnipresent in theoretical computer science and mathematics: The folklore parts of results are only sketched or are entirely left out, while the interesting parts are formalised and detailed proofs are given.

When mechanising a result in an interactive theorem prover, both aspects form their own challenges: For folklore results, first a proof has to be found, then formalised, then mechanised, and each individual step can prove challenging. For a formal proof, only missing details of an argument have to be recovered, which can still impose challenges.

The proof of the invariance thesis for L [7] is accompanied by a mechanisation for one part of the novel contribution, namely two stack machine semantics for L. Since it is folklore that concrete algorithms can be implemented on Turing machines, this part is only sketched.

In the present paper, we give a mechanised proof of the time invariance thesis for L in Coq [26], providing all details left out in existing proofs. We mechanise two variants of the time invariance thesis for L. The first provides a Turing machine  $M_{\text{sim}}$ , simulating L based on a stack machine for L with a heap, and vice versa a term  $s_{\text{sim}}$  simulating Turing machines:

► **Theorem 1** (Time Invariance Thesis for L w.r.t. Simulation). *There are a Turing machine  $M_{\text{sim}}$ , an L-term  $s_{\text{sim}}$ , and encoding relations of L-terms and heaps on tapes, and Turing machines and tapes as L-terms s.t.*

1.  $M_{\text{sim}}$  simulates L with polynomial overhead, i.e. there is a polynomial  $p$  s.t. for all closed  $s$ :
  - a. Whenever  $s$  terminates in  $i$  steps with value  $v$ ,  $M_{\text{sim}}$  run on an input tape encoding  $s$  terminates in  $p(i, \|s\|)$  steps with an encoding of a heap containing  $v$  on its output tape, where  $\|s\|$  is the size of  $s$ .
  - b. If  $M_{\text{sim}}$  terminates on the encoding  $s$ , then  $s$  terminates.
2.  $s_{\text{sim}}$  simulates TMs with linear overhead, i.e. there is a constant  $c$  s.t. for all  $M$ :
  - a. Whenever  $M : \text{TM}_{\Sigma}^n$  terminates on tapes  $\mathbf{t}$  in  $i$  steps with result tapes  $\mathbf{t}'$ ,  $s_{\text{sim}}$  run on an encoding of both  $M$  and  $\mathbf{t}$  evaluates to the encoding of  $\mathbf{t}'$  in  $c \cdot i \cdot \|M\| + c$  steps.
  - b. If  $s_{\text{sim}}$  terminates on the encodings  $M$  and  $\mathbf{t}$ , then  $M$  terminates on  $\mathbf{t}$ .

As a by-product, this theorem yields the first mechanised proof that the halting problems for L and Turing machines are many-one equivalent, which we contribute as a corner stone to the Coq Library of Undecidability Proofs [11].

Two subtleties are important to point out: First, an L-computation does not have explicit input and output. Any L-term  $s_0$  can compute on its own. Input can be realised by application to an (encoded) value  $(s_0\bar{n})$  and the result of the computation can be considered its output. In contrast, a Turing machine  $M$  can not compute without the value of its tapes specified, unless one explicitly defines the canonical value of tapes to be e.g. empty. Secondly, the theorem depends on notions of encoding an L-term on a tape, of encoding heaps on a tape, and of encoding a Turing machine and tapes as L-terms. The choice of such encodings is not canonical, and they have to fulfil certain properties for the theorem to be meaningful. For instance, the unfolding of a heap to an L-term on a Turing machine has to be (at most) polynomial in time, otherwise the theorem is meaningless.

As is well-known, computation in L and the  $\lambda$ -calculus in general is potentially subject to so-called “size explosion”, i.e. the size of a result of a computation can be exponentially larger than both the size of the input and the number of steps. Thus, in the above theorem, unfolding the heap containing  $v$  will be polynomial in the size of  $v$  and  $i$ , but the size of  $v$  might be exponential in both the size of  $s$  and  $i$ . The following version of the time invariance thesis abstracts away from such subtleties by only considering the computability of  $k$ -ary relations on boolean strings, while still being as transparent as possible:

► **Theorem 2** (Time Invariance Thesis for L w.r.t. Computability). *Let  $R \subseteq (\mathbb{L}\mathbb{B})^k \times \mathbb{L}\mathbb{B}$ . Then*

1. *If  $R$  is L-computable with time complexity function  $\tau_R$ , then there is a polynomial  $p$  s.t.  $R$  is TM-computable with time complexity function  $(m, n_1, \dots, n_k) \mapsto p(m, n_1, \dots, n_k, \tau_R(n_1, \dots, n_k))$ .*
2. *If  $R$  is TM-computable with time complexity function  $\tau_R$ , then there is a constant  $c$  s.t.  $R$  is L-computable with time complexity function  $(m, n_1, \dots, n_k) \mapsto c \cdot m \cdot n_1 \cdots n_k \cdot \tau_R(n_1, \dots, n_k) + c$ .*

We accomodate for size explosion by making the time complexity function of a relation  $R$  also depend on the size  $m$  of the output. Of course, for Turing machines  $m$  is always bounded linearly by  $n_1, \dots, n_k$  and  $i$ .

Two corollaries are immediate: If the output of a relation  $R$  is polynomially bounded by its input,  $R$  is polynomially TM-computable if and only if it is polynomially L-computable. In particular, the complexity class P (i.e. PTIME) agrees for both L and Turing machines.

Note that the theorem still depends on encodings, namely on the precise definitions of a relation  $R$  being TM- and L-computable. However, the two notions can be assessed independently: L-computability only depends on how to encode boolean strings ( $\mathbb{L}\mathbb{B}$ ), and similar for TM-computability. Unsurprisingly, both encodings are very simple, and it is obvious that e.g. equality checking works in linear time.

The theorems are furthermore equivalent, in the sense that one can prove one from the other without additional complex simulations: We obtain our mechanised proof of the time invariance thesis w.r.t. computability from the time invariance thesis w.r.t. simulation, by defining terms and machines dynamically exchanging between the encodings of  $\mathbb{L}\mathbb{B}$  on Turing machines,  $\mathbb{L}\mathbb{B}$  in L, L-terms on Turing machines, and Turing machines as L-terms. The other direction would be possible by verifying a universal Turing machine with polynomial overhead in time for Turing machine computation, and similarly a universal L-term.

We avoid non-polynomial overhead for terms exhibiting size explosion by relying on heaps. However, this means that the space overhead is linear-logarithmic rather than constant factor, due to terms exhibiting pointer explosion. A mechanised proof of the time and space invariance thesis for L is left for future work.

The present work is based on the certifying extraction framework for L by Forster and Kunze [6] and the Turing machine verification framework by Forster, Kunze, and Wuttke [9].

The certifying extraction framework allows extracting (by definition total) Coq functions on first-order types to L and automatically proves correctness. The user can provide a time complexity function, which is then automatically verified as well. We use an L-computability proof of Turing machine transitions, which is already a case study of the framework.

The Turing machine verification framework allows giving algorithms in the style of a register-based while-language, and a corresponding machine is automatically constructed behind the scenes. Separate correctness and verification proofs are then inclusion proofs between the automatically derived and the user-given relations for the constructed machine.

**Contribution.** The main contribution of the paper are mechanised proofs of the time invariance thesis for L w.r.t. both simulation and computability. As a by-product, we obtain that the halting problems for Turing machines and L are many-one equivalent, a result we contribute to the Coq Library of Undecidability Proofs [11]. We also contribute a Hoare logic framework for the verification of correctness, time, and space complexity of Turing machines.

**Outline.** The paper is split into four parts: First, Section 3 introduces the weak call-by-value  $\lambda$ -calculus L with small-step, big-step, and stack machine semantics, and Section 4 introduces Turing machines and the Turing machine verification framework from [9]. Secondly, we explain the simulation of L computations on Turing machines with polynomial overhead in Section 5, and how to obtain a simulator  $s_{\text{sim}}$  from the L-computability of Turing machine steps [6] in Section 6, which yields a mechanised proof of the time invariance thesis for L w.r.t. simulation. Thirdly, we explain how to prove that TM-computable relations  $R \subseteq (\mathbb{L}\mathbb{B})^k \times \mathbb{L}\mathbb{B}$  are L-computable with polynomial overhead in Section 7. Lastly, we introduce a novel Hoare logic verification framework for Turing machines in Section 8, which we use to prove that L-computable relations  $R \subseteq (\mathbb{L}\mathbb{B})^k \times \mathbb{L}\mathbb{B}$  are TM-computable with polynomial overhead in Section 9, which yields a mechanised proof of the time invariance thesis for L w.r.t. computability.

**Related work.** Considering only the simulation of models of computation without time complexity, we are aware of Xu et al. [29] mechanising the equivalence of register machines and Turing machines, Forster and Larchey-Wendling [10] with a mechanised compilation of register machines to binary stack machines, and Larchey-Wendling and Forster [18], proving that the halting problems of register machines and  $\mu$ -recursive functions, and the solvability of diophantine equations are many-one equivalent. In unpublished work, Pous [22] mechanises an equivalence proof between counter machines and partial recursive functions in Coq.

## 2 Notations and Definitions

We use the following inductive types:

$$\begin{array}{llll}
\mathbb{1} ::= () & \text{(unit)} & o : \mathbb{O}A ::= \text{None} \mid \text{Some } a & \text{(options)} \\
n : \mathbb{N} ::= 0 \mid Sn & \text{(natural numbers)} & l : \mathbb{L}A ::= [] \mid a :: l & \text{(lists)} \\
b : \mathbb{B} ::= \text{true} \mid \text{false} & \text{(booleans)} & A + B ::= \text{inl } a \mid \text{inr } b & \text{(sums)} \\
& & A \times B ::= (a, b) & \text{(pairs)}
\end{array}$$

We use the notation **if  $a$  is  $s$  then  $b_1$  else  $b_2$**  for inline case analysis, which evaluates to  $b_1$  if  $a$  is of the shape  $s$  (e.g. for a non-zero number  $s := Sn$  or for a list  $s := []$  or  $s := x :: l$ ), and to  $b_2$  otherwise.

We use the functions  $\text{map} : (A \rightarrow B) \rightarrow \mathbb{L}A \rightarrow \mathbb{L}B$  and  $\text{map}_2 : (A \rightarrow B \rightarrow C) \rightarrow \mathbb{L}A \rightarrow \mathbb{L}B \rightarrow \mathbb{L}C$ , defined as follows

$$\begin{array}{ll}
\text{map } f (a :: l) := fa :: \text{map } f l & \text{map}_2 f (a :: l_1) (b :: l_2) := fab :: \text{map}_2 f l_1 l_2 \\
\text{map } f [] := [] & \text{map}_2 f l_1 l_2 := []
\end{array}$$

We use  $i$ ,  $n$ ,  $k$ , and  $m$  as letters for numbers, but try to be consistent in their use: numbers of steps are  $i$ , number of tapes are  $n$ , the arity of input and relations is  $k$ , the size of inputs are  $n_1, \dots, n_k$  and the size of the output is  $m$ .

We write  $\mathbb{P}$  for the type of propositions.  $R \subseteq A \times B$  is short for  $R : A \rightarrow B \rightarrow \mathbb{P}$ , and  $P \subseteq B$  is short for  $P : B \rightarrow \mathbb{P}$ . In particular, if  $R \subseteq A \times B$  and  $a : A$ , then  $Ra \subseteq B$ .

$A^k$  is the type of vectors of length  $k$  with elements in  $A$ . We use bold letters ( $\mathbf{t}, \mathbf{u}, \dots$ ) for vectors and reuse list functions such as  $\text{map}_2$  for vectors. The type  $\text{Fin}_n$  is inductively defined to have exactly  $n$  elements. We write the elements of the type  $\text{Fin}_{5n}$  as  $0, \dots, n$ .

A retraction  $X \hookrightarrow Y$  consists of a function  $I : X \rightarrow Y$  and an inverse function  $R : Y \rightarrow \mathbb{O}X$  s.t.  $\forall x. R(Ix) = \text{Some } x$ .

## 3 The call-by-value $\lambda$ -calculus L

The call-by-value  $\lambda$ -calculus was introduced by Plotkin [21] as variant of Church's  $\lambda$ -calculus [5]. The concrete variant of the call-by-value  $\lambda$ -calculus we present here is called L [12]. We define syntax and semantics for L in Section 3.1 and introduce a stack machine with a heap in 3.2.

### 3.1 Syntax, small-step, and big-step semantics

We define the syntax of L using de Bruijn indices as the syntax of the full  $\lambda$ -calculus, i.e. as variables, applications, or abstractions. The size  $\|s\|$  counts the number of constructors in  $s$  with unary encoded de Bruijn indices.

$$s, t, u : \text{tm}_L ::= n : \mathbb{N} \mid st \mid \lambda s \quad \|n\| := n + 1 \quad \|st\| := \|s\| + \|t\| + 1 \quad \|\lambda s\| := 1 + \|s\|$$

We use names for concrete terms on paper, e.g. write  $(\lambda xy.xx)(\lambda z.z)$  for  $(\lambda\lambda 11)(\lambda 0)$ .

We define a simple substitution operation  $s_t^n$ , agreeing with a more standard parallel substitution operation when the term substituted with is closed:

$$n_u^m := \text{if } n = m \text{ then } u \text{ else } n \quad (st)_u^n := s_u^n t_u^n \quad (\lambda s)_u^n := \lambda(s_u^{S_n})$$

Formally, we say that a term  $s$  is a *closed term* if  $\forall nu. s_u^n = s$ .

We now define a small-step relation  $\succ$ , its  $n$ -step repetition  $\succ^n$ , and an inductive characterisation of weak, call-by-value big-step evaluation  $s \triangleright v$ :

$$\frac{}{(\lambda s)(\lambda t) \succ s_{\lambda t}^0} \quad \frac{s \succ s' \quad t \succ t'}{st \succ s't'} \quad \frac{}{s \succ^0 s} \quad \frac{s \succ s' \quad s' \succ^n t}{s \succ^{S^i} t} \quad \frac{}{\lambda s \triangleright^0 \lambda s} \quad \frac{s \triangleright^{i_1} \lambda u \quad t \triangleright^{i_2} t' \quad u_{t'}^0 \triangleright^{i_3} v}{st \triangleright^{1+i_1+i_2+i_3} v}$$

Note that we have for example  $(\lambda xy.xx)(\lambda z.z) \triangleright^1 \lambda y.(\lambda z.z)(\lambda z.z)$ . Evaluation is called weak because the bodies of abstractions are not evaluated and call-by-value because arguments are evaluated before a function is called. Evaluation agrees with evaluation in Plotkin's calculus [21] on closed terms, but does not treat variables as values.

► **Lemma 3.** *If  $s$  is closed,  $s \triangleright^i t$  if and only if  $s \succ^i t$  and  $t$  is an abstraction.*

Weak call-by-value reduction is uniformly confluent [20], meaning for a terminating term  $s$ , we can talk about *the* time complexity of  $s$  without fixing a reduction path.

► **Fact 4.** *If  $s \succ t_1$  and  $s \succ t_2$ , then  $t_1 = t_2$  or  $\exists u. t_1 \succ u \wedge t_2 \succ u$ .*

► **Corollary 5.** *If  $s \triangleright^{n_1} t_1$  and  $s \triangleright^{n_2} t_2$ , then  $n_1 = n_2$  and  $t_1 = t_2$ .*

The L *halting problem* is defined as  $\text{Halt}_L(s : \text{tm}_L, H : \text{closed}s) := \exists t. s \triangleright t$ .

To define L-computability we introduce so-called Scott encodings [14, 19] for  $\mathbb{B}$ ,  $\mathbb{N}$ , and  $\mathbb{L}$ , which internalise the case-analysis behaviour of the respective types.

$$\begin{array}{lll} \overline{\text{true}} := \lambda xy.x & \bar{0} := \lambda xy.x & \bar{\square} := \lambda xy.x \\ \overline{\text{false}} := \lambda xy.y & \overline{Sn} := \lambda xy.y\bar{n} & \bar{b} :: \bar{l} := \lambda xy.y \bar{b} \bar{l} \end{array}$$

A relation  $R \subseteq \mathbb{L}\mathbb{B}^k \times \mathbb{L}\mathbb{B}$  is L-*computable* with time complexity relation  $\tau \subseteq \mathbb{N}^{1+k} \times \mathbb{N}$  if

$$\begin{aligned} \exists s : \text{tm}_L. \forall l_1 \dots l_k. (\forall l. R(l_1, \dots, l_k)l \rightarrow \exists c \leq \tau(|l|, |l_1|, \dots, |l_k|). s \bar{l}_1 \dots \bar{l}_k \triangleright^c \bar{l}) \wedge \\ \forall t. s \bar{l}_1 \dots \bar{l}_k \triangleright t \rightarrow \exists l. R(l_1, \dots, l_k)\bar{l} \end{aligned}$$

Note that time complexity is a relation rather than a function. We will only consider functional time complexity relations, and write them on paper as if they were functions. However, for undecidable relations  $R$  (e.g. expressing halting problems necessary for deducing a proof of the time invariance thesis w.r.t. simulation from the one w.r.t. computability), it is crucial that  $\tau$  is a relation: one could use a complexity *function*  $\tau : \mathbb{N}^{1+k} \rightarrow \mathbb{N}$  to write a Coq decision function checking whether  $\exists l. R(l_1, \dots, l_k)l$ . Since we know that such a decision is impossible in the constructive type theory of Coq, time complexity functions can not be defined for undecidable problems.

### 3.2 Stack machine semantics

The big-step semantics allows for a compact definition, but is not ideal for implementations of L. To prepare for a simulation of L on Turing machines we introduce a stack machine for L, utilising references to a heap instead of substitution, similar to the heap machine by Kunze et al. [15]. In contrast to the results there, we give a direct correctness proof instead of a step-wise refinement via several machines. Our machine is also similar to the heap machine by Forster et al. [7], but with fewer reduction rules simplifying verification.

Instead of terms, we will work with programs  $P, Q : \text{Pro} := \mathbb{L} \text{Com}$ , which are lists of commands. Commands are reference, application, abstraction, or return tokens:

$$c : \text{Com} ::= \text{ref } n \mid \text{app} \mid \text{lam} \mid \text{ret}$$

The *compilation function*  $\gamma : \text{Ter} \rightarrow \text{Pro}$  compiles terms to programs:

$$\gamma n := [\text{ref } n] \quad \gamma(st) := \gamma s \# \gamma t \# [\text{app}] \quad \gamma(\lambda s) := \text{lam} :: \gamma s \# [\text{ret}]$$

We have  $\gamma((\lambda xy.xx)(\lambda z.z)) = [\text{lam}; \text{lam}; \text{ref } 1; \text{ref } 1; \text{app}; \text{ret}; \text{ret}; \text{lam}; \text{ref } 0; \text{ret}; \text{app}]$ .

Compiled abstractions start with the token `lam` and end with `ret`. We can thus define a function  $\phi P : \mathbb{O}(\text{Pro} \times \text{Pro})$  that extracts the body of an abstraction by matching the tokens like parentheses. We define  $\phi P := \phi_{0, []} P$ , where  $\phi_{k,Q} P$  is an auxiliary function storing the number  $k$  of unmatched `lam` and the processed prefix  $Q$ :

$$\begin{aligned} \phi_{k,Q} [] &:= \text{None} & \phi_{0,Q} (\text{ret} :: P) &:= \text{Some}(Q, P) & \phi_{S_k,Q} (\text{ret} :: P) &:= \phi_{k,Q \# [\text{ret}]} P \\ \phi_{k,Q} (\text{lam} :: P) &:= \phi_{S_k,Q \# [\text{lam}]} P & \phi_{k,Q} (c :: P) &:= \phi_{k,Q \# [c]} P \text{ if } c = \text{ref } n \text{ or } \text{app} \end{aligned}$$

The *states of the heap machine* are tuples  $T, V, H$ . The *control stack*  $T$  and the *value stack*  $V$  are lists of *closures*  $g : \text{Clos} := \text{Pro} \times \mathbb{N}$ . A closure  $(P, a)$  denotes an open program, where the reference 0 in  $P$  has to be looked up at address  $a : \mathbb{N}$  in the heap when evaluating.

The *heap*  $H$  is a linked list of heap entries  $e : \text{Entry} := \mathbb{O}(\text{Clos} \times \mathbb{N})$ , i.e. an entry is either empty, or contains the head of the list and the address of its tail. Given a heap  $H$  and an address  $a$ ,  $H[a] : \mathbb{O} \text{Entry}$  denotes the  $a$ -th element of  $H$ . We define  $H[a, n]$  to be the  $n$ -th entry on the heap starting at address  $a$  as follows:

$$H[a, n] := \text{if } H[a] \text{ is } \text{Some}(\text{Some}(g, b)) \text{ then if } n \text{ is } S_n \text{ then } H[b, n] \text{ else } \text{Some } g \text{ else } \text{None}$$

We can now define the small-step semantics of the stack machine for  $L$ :

$$\begin{aligned} (\text{lam} :: P, a) :: T, V, H &\rightsquigarrow (P', a) ::_{\text{tc}} T, (Q, a) :: V, H & \text{if } \phi P = \text{Some}(Q, P') \\ (\text{ref } n :: P, a) :: T, V, H &\rightsquigarrow (P, a) ::_{\text{tc}} T, g :: V, H & \text{if } H[a, n] = \text{Some } g \\ (\text{app} :: P, a) :: T, g :: (Q, b) :: V, H &\rightsquigarrow (Q, |H|) :: (P, a) ::_{\text{tc}} T, V, H \# [\text{Some}(g, b)] \end{aligned}$$

Here,  $(P, a) ::_{\text{tc}} T := \text{if } P \text{ is } [] \text{ then } T \text{ else } (P, a) :: T$ .

In the abstraction rule, the complete abstraction is parsed via  $\phi$  and its body put on the value stack. In principle  $::$  instead of  $::_{\text{tc}}$  could be used to obtain a correct machine, however the time complexity of this machine is easier to verify using this optimising operation.

Similarly, in the reference rule, the body of the abstraction corresponding to the variable  $n$  is looked up in the heap starting at address  $a$  and the result is put on the value stack.

In the application rule, the machine takes the closure of the called function  $(Q, b)$  and its argument  $g$  from the value stack. The address  $b$  is bound to  $g$  in the heap, the entry being appended to  $H$ , thus obtaining address  $|H|$ . The machine continues evaluating the body  $Q$ , where the value for reference 0 can be looked up at address  $|H|$ , where it was just placed.

Given a closed term  $s$ , the initial state of the machine is  $([(\gamma s, 0)], [], [])$ , i.e. an empty value stack, an empty heap, and the closure  $(\gamma s, 0)$  on the task stack. In fact, for a closed term, 0 can be replaced by any address since there are no free references. An example run of the machine for  $(\lambda xy.xx)(\lambda z.z)$  can be found in Figure 1, using  $a$  as start address.

To state the correctness of the stack machine we need to define an unfolding operation  $\text{unf}_H(P, a)$ . We will use functional notation for unfolding on paper, but define it as a functional relation in Coq, since the function is not structurally recursive, and not even terminating on cyclic heaps. The function  $\text{unf} : \text{Pro} \rightarrow \text{tm}_L$  unfolds programs from the value

$$\begin{aligned}
 & \text{unf}_{\text{Some}(([\text{ref } 0], a), a)}([\text{ref } 1; \text{ref } 1; \text{app}], 0) \\
 &= \text{unf}_{\text{Some}(([\text{ref } 0], a), a), 0, 0}(\text{unf}[\text{ref } 1; \text{ref } 1; \text{app}]) \\
 &= \text{unf}_{\text{Some}(([\text{ref } 0], a), a), 0, 0}(\lambda 11) \\
 &= \lambda \text{unf}_{\text{Some}(([\text{ref } 0], a), a), 0, 1}(11) \\
 &= \lambda(\text{unf}_{\text{Some}(([\text{ref } 0], a), a), 0, 1^1})(\text{unf}_{\text{Some}(([\text{ref } 0], a), a), 0, 1^1}) \\
 &= \lambda(\text{unf}_{\text{Some}(([\text{ref } 0], a), a), 0, 0}(\text{unf}[\text{ref } 0]))) \\
 &= \lambda(\lambda 0)(\lambda 0) = \lambda y.(\lambda z.z)(\lambda z.z)
 \end{aligned}$$

$$\begin{aligned}
 & [([\text{lam}; \text{lam}; \text{ref } 1; \text{ref } 1; \text{app}; \text{ret}; \text{ret}; \text{lam}; \text{ref } 0; \text{ret}; \text{app}], a), [], [] \\
 & \rightsquigarrow [([\text{lam}; \text{ref } 0; \text{ret}; \text{app}], a), [([\text{lam}; \text{ref } 1; \text{ref } 1; \text{app}; \text{ret}], a)], [] \\
 & \rightsquigarrow [[[\text{app}], a], [[[\text{ref } 0], a], ([\text{lam}; \text{ref } 1; \text{ref } 1; \text{app}; \text{ret}], a)], [] \\
 & \rightsquigarrow [([\text{lam}; \text{ref } 1; \text{ref } 1; \text{app}; \text{ret}], 0), [], [\text{Some}(([\text{ref } 0], a), a) \\
 & \rightsquigarrow [], [[[\text{ref } 1; \text{ref } 1; \text{app}], 0)], [\text{Some}(([\text{ref } 0], a), a)
 \end{aligned}$$

■ **Figure 1** Example of the execution for term  $(\lambda xy. xx)(\lambda z.z)$ .

stack into a term by inverting  $\gamma$ . It adds  $\lambda$  to the result, since only the bodies of abstractions are saved on the value stack. The function  $\text{unf}_{H,a,k} : \text{tm}_L \rightarrow \text{tm}_L$  substitutes free variables  $n \geq k$  in a term by  $H[a, n - k]$ . Finally,  $\text{unf}_H : \text{Pro} \times \mathbb{N} \rightarrow \text{tm}_L$  unfolds a result using the two previous functions. The example from above is continued in Figure 1.

$$\begin{aligned}
 \text{unf } P &:= \lambda t \quad (\text{if } \gamma t = P) & \text{unf}_{H,a,k} n &:= n \quad (\text{if } n < k) \\
 \text{unf}_{H,a,k} n &:= \text{unf}_{H,b,0}(\text{unf } P) \quad (\text{if } n \geq k \text{ and } H[a, n - k] = \text{Some}(P, b)) \\
 \text{unf}_{H,a,k}(st) &:= (\text{unf}_{H,a,k}s)(\text{unf}_{H,a,k}s) & \text{unf}_{H,a,k}(\lambda s) &:= \lambda(\text{unf}_{H,a,k}s) & \text{unf}_H(P, a) &:= \text{unf}_{H,a,0}(\text{unf } P)
 \end{aligned}$$

We define the size of the components of a stack machine as follows:

$$\begin{aligned}
 \|\text{ref } n\| &:= \|n\| + 1 & \|\text{app}\| &:= \|\text{lam}\| := \|\text{ret}\| := 1 & \|n : \mathbb{N}\| &:= n + 1 \\
 \|(a, b)\| &:= \|a\| + \|b\| + 1 & \|[]\| &:= 1 & \|x :: l\| &:= \|x\| + \|l\|
 \end{aligned}$$

The final correctness theorem then reads:

► **Theorem 6.** *Let  $s$  be closed.*

1. *If  $s \triangleright^i v$  then  $([(\gamma s, 0)], [], []) \rightsquigarrow^{3 \cdot i + 1} ([[], [(P, a)], H)$  for some  $P$ ,  $a$  and  $H$  s.t.  $\text{unf}_H(P, a) = v$ .*
2. *If  $([(\gamma s, 0)], [], []) \rightsquigarrow^i (T, V, H)$  then  $\|(T, V, H)\| \leq c \cdot (i + 1) \cdot (i + \|s\| + 1)$  for some  $c$  and if furthermore  $\neg \exists \sigma. (T, V, H) \rightsquigarrow \sigma$ , then  $T = [], V = [(P, a)]$ , and  $s \triangleright v$  for some  $P$ ,  $a$ , and  $v$  s.t.  $\text{unf}_H(P, a) = v$  is defined.*

### 3.3 Mechanisation in Coq

The weak call-by-value  $\lambda$ -calculus L is a sweet spot for the mechanisation of computability and complexity theory – but only since it is engineered to be one. In principle, there is an abundance of options which  $\lambda$ -calculus to choose: call-by-value or call-by-name, weak or strong, variables are values or not, de Bruijn encoding with simple substitution, or with parallel substitution, or locally nameless, or parametric higher-order abstract syntax.

For the implementation of L on Turing machines it is mainly the choice of a simple substitution operation based on de Bruijn indices that is crucial. Parallel substitution is considerably more complicated to define, since it is not structurally recursive *and* a priori uses functions, i.e. uncountable types, to represent substitutions. In contrast, simple substitutions are structurally recursive and only require natural numbers for their definition.

To simulate Turing machines in L, it is crucial that also a small step semantics is available: the non-determinism of  $\succ$ , i.e. that it applies to both sides of an application, allows (directed) equational reasoning in correctness proofs without any overhead.



## 4 Turing machines

Turing machines [27] are widely used in books on computability theory, are the standard model of computation for complexity theory, and are still considered by many to be the model of computation most convincingly capturing all “effectively calculable” functions. Despite their universal use, there is however no consensus on how to formally define Turing machines in the literature. There are a multitude of definitions which can all be proved equivalent. In Appendix A, we compare the definition we use [9] to the one by Hopcroft et al. [13]<sup>1</sup>

In Section 4.2, we give an overview of the Turing machine verification framework from [9].

### 4.1 Definition

We start by defining a tape over type  $\Sigma$  inductively using four constructors [2]:

$$\text{tp}_\Sigma ::= \text{niltp} \mid \text{leftof } r \text{ } rs \mid \text{midtp } ls \text{ } m \text{ } rs \mid \text{rightof } l \text{ } ls \quad \text{where } m, l, r : \Sigma \text{ and } ls, rs : \mathbb{L}\Sigma$$

The representation does enforce that a tape contains a continuous sequence of symbols, and that either a symbol or one of the ends of the tape is distinguished as head position. There are no explicit blank symbols, which allows for a unique representation of every tape and no well-formedness predicate is needed.

We define a type of moves  $\text{Move}$ , a function  $\text{mv} : \text{Move} \rightarrow \text{tp} \rightarrow \text{tp}$  applying a move to a tape, a function  $\text{wr} : \mathbb{O}\Sigma \rightarrow \text{tp} \rightarrow \text{tp}$  writing to a tape, and a function  $\text{curr} : \text{tp} \rightarrow \mathbb{O}\Sigma$  obtaining the current symbol of a tape in Figure 2.

We define multi-tape *Turing machines*  $M : \text{TM}_\Sigma^n$ , where  $n : \mathbb{N}$  is the number of tapes and the finite type  $\Sigma$  is the alphabet, as dependent pairs  $(Q, \delta, q_0, \text{halt})$  where  $Q$  is a finite type,  $\delta : Q \times (\mathbb{O}\Sigma)^n \rightarrow Q \times (\mathbb{O}\Sigma \times \text{Move})^n$ ,  $q_0 : Q$  is the starting state, and  $\text{halt} : Q \rightarrow \mathbb{B}$  indicates halting states. The definition of Turing machine evaluation  $M(q, t) \triangleright (q', t')$  and the *Turing machine halting problem*  $\text{Halt}_{\text{TM}}$  are defined in Figure 2.

A relation  $R \subseteq (\mathbb{L}\mathbb{B})^k \times (\mathbb{L}\mathbb{B})$  is *TM-computable* with time-complexity  $\tau \subseteq \mathbb{N}^{1+k} \times \mathbb{N}$  if

$$\begin{aligned} \exists n : \mathbb{N}. \exists \Sigma. \exists s \text{ bl} : \Sigma. s \neq \text{bl} \wedge \exists M : \text{TM}_\Sigma^{1+k+n}. \forall l_1 \dots l_k. \\ (\forall l. R(l_1, \dots, l_k) l \rightarrow \exists i \leq \tau(|l|, |l_1|, \dots, |l_k|). \exists q \mathbf{t}. M(q_0, [\text{niltp}, \bar{l}_1, \dots, \bar{l}_k, \text{niltp}, \dots, \text{niltp}]) \triangleright^i (q, \mathbf{t}) \wedge \mathbf{t}[0] = \bar{l}) \wedge \\ \forall q \mathbf{t} i. M(q_0, [\text{niltp}, \bar{n}_1, \dots, \bar{n}_k, \text{niltp}, \dots, \text{niltp}]) \triangleright^i (q, \mathbf{t}) \rightarrow \exists l. R(l_1, \dots, l_k) l \end{aligned}$$

with  $[\bar{x}_1, \dots, \bar{x}_n] := \text{midtp } [] \text{ bl } [\bar{x}_1, \dots, \bar{x}_n]$  and  $\overline{\text{true}} := s, \overline{\text{false}} := \text{bl}$ .

### 4.2 Verified programming of Turing machines

As presented, Turing machines are not compositional: There is no canonical way how to execute a 5-tape machine over alphabet  $\mathbb{B}$  after a 3-tape machine over  $\mathbb{O}(\mathbb{B} \times \mathbb{B})$ .

To allow for the composition of Turing machines and their verification, we first introduce labellings in order to abstract away from the state space. A *labelled Turing machine* over a type  $L$ , written  $M : \text{TM}_\Sigma^n(L)$ , is a dependent pair  $(M', \text{lab}_M)$  of a machine  $M' : \text{TM}_\Sigma^n$  and a labelling function  $\text{lab}_M : Q_{M'} \rightarrow L$ .

To prove the soundness of machines, we introduce *realisation*. A Turing machine  $M : \text{TM}_\Sigma^n(L)$  realises a relation  $R \subseteq \text{tp}_\Sigma^n \times (L \times \text{tp}_\Sigma^n)$  if

$$M \models R := \forall \mathbf{t} \mathbf{t}'. M(q_0, \mathbf{t}) \triangleright (q, \mathbf{t}') \rightarrow R \mathbf{t} (\text{lab}_M q, \mathbf{t}')$$

<sup>1</sup> chosen for instance by Wikipedia as reference definition

## 19:10 Mechanised Time Invariance Thesis for L

Move ::= L   N   R	
$\begin{aligned} \text{mv} &: \text{Move} \rightarrow \text{tp} \rightarrow \text{tp} \\ \text{mv L}(\text{rightof } l \text{ } rs) &:= \text{midtp } l \text{ } l \text{ } [] \\ \text{mv L}(\text{midtp } [] \text{ } m \text{ } rs) &:= \text{leftof } m \text{ } rs \\ \text{mv L}(\text{midtp } (l :: ls) \text{ } a \text{ } rs) &:= \text{midtp } l \text{ } l \text{ } (a :: rs) \\ \text{mv m } t &:= t \quad \text{in all other cases} \end{aligned}$	$\begin{aligned} \text{mv R}(\text{leftof } r \text{ } rs) &:= \text{midtp } [] \text{ } r \text{ } rs \\ \text{mv R}(\text{midtp } l \text{ } s \text{ } []) &:= \text{rightof } a \text{ } l \text{ } s \\ \text{mv R}(\text{midtp } l \text{ } s \text{ } a \text{ } (r :: rs)) &:= \text{midtp } (a :: l \text{ } s) \text{ } r \text{ } rs \end{aligned}$
$\begin{aligned} \text{wr} &: \mathbb{O}\Sigma \rightarrow \text{tp} \rightarrow \text{tp} \\ \text{wr None } t &:= t & \text{wr (Some } a \text{) niltp} &:= \text{midtp } [] \text{ } a \text{ } [] & \text{wr (Some } a \text{) (midtp } l \text{ } s \text{ } b \text{ } rs)} &:= \text{midtp } l \text{ } s \text{ } a \text{ } rs \\ \text{wr (Some } a \text{) (leftof } r \text{ } rs)} &:= \text{midtp } [] \text{ } a \text{ } (r :: rs) & \text{wr (Some } a \text{) (rightof } l \text{ } ls)} &:= \text{midtp } (l :: l \text{ } s) \text{ } a \text{ } [] \end{aligned}$	
$\begin{aligned} \text{curr} &: \text{tp} \rightarrow \mathbb{O}\Sigma \\ \text{curr}(\text{midtp } l \text{ } s \text{ } a \text{ } rs) &:= \text{Some } a & \text{curr } t &:= \text{None} \quad \text{otherwise} \end{aligned}$	
$\frac{\text{halt } q = \text{true}}{M(q, \mathbf{t}) \triangleright^i (q, \mathbf{t})} \qquad \frac{\text{halt } q = \text{false} \quad \delta(q, \text{map curr } \mathbf{t}) = (q', a) \quad M(q', \text{map}_2(\lambda(c, m) \text{ } t. \text{mv } m \text{ } (\text{wr } c \text{ } t))) \text{ } a \text{ } t) \triangleright^i (q'', \mathbf{t}')}{M(q, \mathbf{t}) \triangleright^{S^i} (q'', \mathbf{t}')}$	
$\text{Halt}_{\text{TM}_{\Sigma}^n}(M : \text{TM}_{\Sigma}^n, \mathbf{t} : \text{tp}_{\Sigma}^n) := \exists i q' \mathbf{t}'. M(q_0, \mathbf{t}) \triangleright^i (q', \mathbf{t}')$	
$\text{Halt}_{\text{TM}}(n : \mathbb{N}, \Sigma, M : \text{TM}_{\Sigma}^n, \mathbf{t} : \text{tp}_{\Sigma}^n) := \text{Halt}_{\text{TM}_{\Sigma}^n}(M, \mathbf{t})$	

■ **Figure 2** Definitions for Turing machines.

Dually, we introduce *termination*.  $M : \text{TM}_{\Sigma}^n(L)$  terminates in  $T \subseteq \text{tp}_{\Sigma}^n \times \mathbb{N}$  if

$$M \downarrow T := \forall \mathbf{t} i. T \mathbf{t} i \rightarrow \exists q \mathbf{t}'. M(\mathbf{t}) \triangleright^i (q, \mathbf{t}').$$

We call a machine *total* if  $\exists c. M \downarrow \lambda \mathbf{t} i. i \geq c$ , i.e. if it halts on any tape in at most  $c$  steps.

► **Fact 7.** *The introduced predicates are (anti-)monotonic:*

1. If  $M \models R'$  and  $\forall \mathbf{t} \ell \mathbf{t}'. R' \mathbf{t} (\ell, \mathbf{t}') \rightarrow R \mathbf{t} (\ell, \mathbf{t}')$ , then  $M \models R$ .
2. If  $M \downarrow T'$  and  $\forall \mathbf{t} i. T \mathbf{t} i \rightarrow T' \mathbf{t} i$ , then  $M \downarrow T$ .

We will use the following total machines we call *primitive machines*:

$$\begin{aligned} \text{Read } s : \text{TM}_{\Sigma}^1(\mathbb{O}(\Sigma)) \models \lambda \mathbf{t} (\ell, \mathbf{t}'). \ell = \text{curr } \mathbf{t}[0] \wedge \mathbf{t}' = \mathbf{t} & \quad \text{Write } s : \text{TM}_{\Sigma}^1(\mathbb{1}) \models \lambda \mathbf{t} \mathbf{t}'. \mathbf{t}'[0] = \text{wr } s \text{ } \mathbf{t}[0] \\ \text{Move } d : \text{TM}_{\Sigma}^1(\mathbb{1}) \models \lambda \mathbf{t} \mathbf{t}'. \mathbf{t}'[0] = \text{mv } d \text{ } \mathbf{t}[0] & \quad \text{Return } \ell : \text{TM}_{\Sigma}^n(L) \models \lambda \mathbf{t} (\ell', \mathbf{t}'). \mathbf{t}' = \mathbf{t} \wedge \ell' = \ell \end{aligned}$$

The last necessary tool now are *combinators* to compose machines. Given  $M : \text{TM}_{\Sigma}^n(L)$  and  $M'_{\ell} : \text{TM}_{\Sigma}^n(L')$  (for  $\ell : L$ ), we introduce the combinator  $\text{Switch } M \ M' : \text{TM}_{\Sigma}^n(L')$ , which executes  $M$  and, depending on the label  $\ell$  returned by  $M$ , executes  $M'_{\ell}$ .

$$\frac{M \models R \quad \forall (\ell : L). M'_{\ell} \models R'_{\ell}}{\text{Switch } M \ M' \models \lambda \mathbf{t}_0 (\ell', \mathbf{t}'). \exists \mathbf{t} (\ell : L). R \mathbf{t}_0 (\ell, \mathbf{t}) \wedge R'_{\ell} \mathbf{t} (\ell', \mathbf{t}')} \quad \frac{M \downarrow T \quad M \models R \quad \forall (\ell : L). M'_{\ell} \downarrow T'_{\ell}}{\text{Switch } M \ M' \downarrow \lambda \mathbf{t} k. \exists k_1 k_2. T \mathbf{t} k_1 \wedge 1 + k_1 + k_2 \leq k \wedge \forall \ell \mathbf{t}'. R \mathbf{t} (\ell, \mathbf{t}') \rightarrow T'_{\ell} \mathbf{t}' k_2}$$

Given a machine  $M : \text{TM}_{\Sigma}^n(\mathbb{O}(L))$ , we introduce the combinator  $\text{While } M : \text{TM}_{\Sigma}^n(L)$ , which loops  $M$  until the result label is  $\text{Some } \ell$ . The realisation relation for  $\text{While}$  is defined

inductively, whereas the termination relation is the co-inductively defined accessibility relation (the dashed line indicates coinduction).

$$\begin{array}{c}
\frac{M \vDash R}{\text{While } M \vDash \text{While}R R} \qquad \frac{R \ t \ (\text{Some } \ell, \ t')}{\text{While}R R \ t \ (\ell, \ t')} \qquad \frac{R \ t \ (\text{None}, \ t') \quad \text{While}R R \ t' \ (\ell, \ t'')}{\text{While}R R \ t \ (\ell, \ t'')} \\
\\
\frac{M \downarrow T \quad M \vDash R}{\text{While } M \downarrow \text{While}T R T} \qquad \frac{\begin{array}{c} T \ t \ k_1 \quad \forall t' \ l. R \ t \ (\text{Some } l, \ t') \rightarrow k_1 \leq k \\ \forall t'. R \ t \ (\text{None}, \ t') \rightarrow \exists k_2. \text{While}T R T \ t' \ k_2 \wedge 1 + k_1 + k_2 \leq k \end{array}}{\text{While}T R T \ t \ k}
\end{array}$$

Composing machines with the operator `Switch` only works for machines over the same alphabet and the same number of tapes. To remedy this situation we introduce lifting operations for alphabets and tapes, and also a relabelling operation `Relabel`.

Given a retraction  $f : \Sigma \rightarrow \Gamma$ , a default symbol  $d : \Sigma$ , and a machine  $M : \text{TM}_{\Sigma}^n$ , the *alphabet lift*  $\uparrow_{(f,d)} M : \text{TM}_{\Gamma}^n$  translates every read symbol via  $f^{-1}$ , passes it to  $M$ , and translates the symbol  $M$  writes via  $f$ . In case  $f^{-1}$  returns `None`,  $d$  is passed to  $M$ .

Given a retraction  $I : \text{Fin}_m \rightarrow \text{Fin}_n$  and a machine  $M : \text{TM}_{\Sigma}^m$ , the *tape lift*  $\uparrow_I M : \text{TM}_{\Sigma}^n$  replicates the behavior of  $M$  on tape  $i$  on tape  $Ii$ , and leaves all other tapes untouched.

We only show the canonical realisation and termination relations for the tape lift here:

$$\begin{aligned}
\uparrow_I R &:= \lambda t \ (\ell, \ t'). R \ (\text{select } I \ t) \ (\ell, \ \text{select } I \ t') \wedge \forall (i : \text{Fin}_n). i \notin I \rightarrow t'[i] = t[i] \\
\uparrow_I T &:= \lambda t \ k. T \ (\text{select } I \ t) \ k \quad \text{where } \text{select } I \ t := \text{map } (\lambda i. t[I(i)]) \ [0, \dots, m-1]
\end{aligned}$$

$$\frac{M : \text{TM}_{\Sigma}^m(L) \vDash R \quad I : \text{Fin}_m \hookrightarrow \text{Fin}_n}{\uparrow_I M : \text{TM}_{\Sigma}^n \vDash \uparrow_I R} \quad \frac{M \downarrow T \quad I : \text{Fin}_m \hookrightarrow \text{Fin}_n}{\uparrow_I M \downarrow \uparrow_I T}$$

Given a function  $r : L_1 \rightarrow L_2$  and a machine  $M : \text{TM}_{\Sigma}^n(L_1)$ , `Relabel`  $M \ r : \text{TM}_{\Sigma}^n(L_2)$  behaves like  $M$ , but returns label  $r\ell$  where  $M$  returned  $\ell$ .

The last important layer of abstraction introduces the treatment of tapes as registers, based on the notion  $t[i] \simeq v$  expressing that tape  $t[i]$  contains an encoded value  $v : V$ . A type  $V$  is a *TM-encodable type* on alphabet  $\Sigma$  if there is a (designated) injective function  $\varepsilon : V \rightarrow \mathbb{L}\Sigma$ . We define such designated encoding functions for several data types, e.g. booleans, tuples, and lists of encodable types. In Coq, the implementation of encodable types relies on type classes, such that users can define their own encoding functions.

We then define *tape containment*  $t \simeq_f v$ , where  $V$  is TM-encodable on alphabet  $\Sigma$ ,  $v : V$ ,  $t : \text{tp}_{\Gamma^+}$ , and  $f : \Sigma \rightarrow \Gamma$ .

$$\begin{aligned}
\Gamma^+ &::= \text{START} \mid \text{STOP} \mid \text{UNKNOWN} \mid (s : \Gamma) \\
t \simeq_f v &:= \exists ! s. t = \text{midtp } ls \text{ START } (\text{map } f \ (\varepsilon v) \ ++ \ [\text{STOP}])
\end{aligned}$$

Note that the position of the head is fixed in the definition of tape containment: The head must be located on the start symbol. By extending the alphabet  $\Sigma$  with the delimiting symbols `START` and `END`, values can effectively be copied from one tape to another tape. The symbol `UNKNOWN` is used as the canonical default symbol for the alphabet lift. Let  $M : \text{TM}_{\Sigma^+}^n$  and  $f : \Sigma \rightarrow \Gamma$ . Then  $\uparrow_{f^+} M : \text{TM}_{\Gamma^+}^n$  (with the canonically inferrable injection  $f^+ : \Sigma^+ \rightarrow \Gamma^+$ ) is an alphabet lift of  $M$ . In case the lifted machine reads a symbol that is not in the image of  $f$ ,  $\uparrow_{f^+} M$  behaves like if  $M$  reads `UNKNOWN`. However, this will by design not happen if the head is under a symbol of the encoding of a value.

*Void* tapes (written `isVoid`  $t$ ) do not contain values. The head of the tape is located at the right-most symbol:

$$\text{isVoid } t := \exists m \ ls. t = \text{midtp } ls \ m \ []$$

A void tape can be initialised with a value by writing the encoding of a value delimited by `START` and `END`.

## 5 Simulating L on Turing machines

To simulate L on Turing machines, we use the stack machine semantics, meaning we have to implement the relation  $\rightsquigarrow$  from Section 3.2 as multi-tape Turing machine  $\text{Step} : \text{TM}_{\Sigma}^5$ . The central components of  $\text{Step}$  are machines implementing the heap lookup operation  $H[a, n]$  and the parsing operation  $\phi$ . We will omit concrete implementations, but show the correctness and termination relations and briefly discuss the proof goals for such verifications. The machines will share an alphabet  $\Sigma$  consisting of 30 symbols, allowing to encode commands, programs, addresses, closures, task and value stack, heap entries, and heap.

The relations we display here are simplified in comparison to the actual relations in Coq w.r.t. two aspects: First, we omit the retractions  $f_X : \Sigma_X \rightarrow \Sigma$  when writing  $\simeq$ . Since as long as concrete  $f_X$  are fixed for every type  $X$  encodable on type  $\Sigma_X$ , their concrete definitions do not matter. Secondly, we omit the condition  $\text{isVoid } t$  both in the premise and conclusion of rules: Any unspecified tape is always implicitly void.

► **Fact 8.** *There is a machine  $\text{Lookup} : \text{TM}_{\Sigma}^5(\mathbb{B})$  and a  $c : \mathbb{N}$  s.t.*

1.  $\text{Lookup} \models \lambda t(\ell, t'). \forall H a n. t[0] \simeq H \rightarrow t[1] \simeq a \rightarrow t[2] \simeq n \rightarrow$   
 $\text{if } \ell \text{ then } \exists g. H[a, b] = \text{Some } g \wedge t'[0] \simeq H \wedge t'[3] \simeq g \text{ else } H[a, b] = \text{None}$
2.  $\text{Lookup} \downarrow \lambda ti. \exists H a n. t[0] \simeq H \wedge t[1] \simeq a \wedge t[2] \simeq n \wedge i \geq c \cdot (n + 1) \cdot (\|H\| + \max \|H\| \|a\|)$

**Proof.** The machine  $\text{Lookup}$  can be defined by using building blocks like  $\text{While}$  and  $\text{Switch}$ . Once the machine is defined, an inductive relation  $R$  s.t.  $\text{Lookup} \models R$  can be automatically inferred from the relations of the building blocks.

By Fact 7 it then suffices to prove  $R t(\ell, t') \rightarrow \forall H a n. t[0] \simeq H \rightarrow t[1] \simeq a \rightarrow t[2] \simeq n \rightarrow$   
 $\text{if } \ell \text{ then } \exists g. H[a, b] = \text{Some } g \wedge t'[0] \simeq H \wedge t'[3] \simeq g \text{ else } H[a, b] = \text{None}$  by induction on  $R$ .

The termination proof is dual. ◀

► **Fact 9.** *There is a machine  $\text{Parse} : \text{TM}_{\Sigma}^5(\mathbb{B})$  and a  $c : \mathbb{N}$  s.t.*

1.  $\text{Parse} \models \lambda t(\ell, t'). \forall P. t[0] \simeq P \rightarrow$   
 $\text{if } \ell \text{ then } \exists Q P'. \phi P = (Q, P') \wedge t'[0] \simeq P' \wedge t'[1] \simeq Q \text{ else } \phi P = \text{None}$
2.  $\text{Parse} \downarrow \lambda ti. \exists P. t[0] \simeq P \wedge i \geq \|c\| \cdot P^2$

► **Fact 10.** *There is a machine  $\text{Step} : \text{TM}_{\Sigma}^{11}(\mathbb{B})$  and a  $c : \mathbb{N}$  s.t.*

1.  $\text{Step} \models \lambda t(\ell, t'). \forall T V H. t[0] \simeq T \rightarrow t[1] \simeq V \rightarrow t[2] \simeq H \rightarrow$   
 $\text{if } \ell \text{ then } \exists T' V' H'. (T, V, H) \rightsquigarrow (T', V', H') \wedge t'[0] \simeq T' \wedge t'[1] \simeq V' \wedge t'[2] \simeq H'$   
 $\text{else } (\neg \exists \sigma. (T, V, H) \rightsquigarrow \sigma) \wedge T = [] \rightarrow t'[0] \simeq [] \wedge t'[1] \simeq V \wedge t'[2] \simeq H$
2.  $\text{Step} \models \lambda ti. t[0] \simeq T \wedge t[1] \simeq V \wedge t[2] \simeq H \wedge i \geq 1 + c \cdot \text{if } T \text{ is } (a, P) :: \_ \text{ then}$   
 $\|a\| + \|H\| + \|V\| + \|P\| \cdot (1 + \|H\| + \max a \|H\|_{\text{addr}} + \|P\|) \text{ else } 0$

This suffices to prove one direction of the time invariance thesis w.r.t. simulation:

► **Theorem 11.** *There is  $M_{\text{sim}} : \text{TM}_{\Sigma}^{11}$  and a polynomial  $p : \mathbb{N} \rightarrow \mathbb{N}$  s.t. for closed terms  $s$*

1. *If  $s \triangleright^i v$ , then there exist  $t, H, P$ , and a s.t.  $M_{\text{sim}}(\overline{[(\gamma s, 0)]}, \text{niltp}, \dots, \text{niltp}) \triangleright^{p(i, \|s\|)} t$  with  
 $t[1] = \overline{[]}, t[2] = \overline{(P, a)}, t[3] = \overline{H}$ , and  $\text{unf}_H(P, a) = v$ .*
2. *If  $M_{\text{sim}}(\overline{[(\gamma s, 0)]}, \text{niltp}, \dots, \text{niltp})$  terminates, so does  $s$ .*

**Proof.** Define  $M_{\text{sim}} := \text{While Step}$  and  $p i m := c \cdot (3i + 2)^2 \cdot (3i + 1 + m) \cdot m$  for some  $c$ . ◀

► **Corollary 12.**  $\text{Halt}_L \preceq \text{Halt}_{\text{TM}}$

A first version of the Coq verification of  $M_{\text{sim}}$  was also discussed in Wuttke's bachelor's thesis [28].

## 6 Simulating Turing machines in L

To simulate Turing machines in L, we first give an alternative, executable semantics for Turing machines based on iteration of a step-function. We then recap the central ingredients of the certifying extraction framework [6], which we use to extract the step-function to L. And lastly, we implement a (potentially non-terminating) iteration combinator in L.

We define a function  $\text{nxt}_M : Q_M \times \text{tp}_\Sigma^n \rightarrow Q_M \times \text{tp}_\Sigma^n + \text{tp}_\Sigma^n$  and a polymorphic function  $\text{loop} : (X \rightarrow X + Y) \rightarrow X \rightarrow \mathbb{N} \rightarrow \mathbb{O}Y$  as follows:

$$\text{nxt}_M(q, t) := \text{if halt } q \text{ then } t \text{ else let } (q', a) := \delta_M(q, \text{curr } t) \text{ in } (q', \text{map}_2(\lambda(c, m)t. \text{mv } m(\text{wr } ct)) a t)$$

$$\text{loop } f x 0 := \text{None} \quad \text{loop } f x (Sn) := \text{loop } f x' n \text{ (if } fx = \text{inl } x')$$

$$\text{loop } f x (Sn) := \text{Some } y \text{ (if } fx = \text{inr } y)$$

► **Fact 13.**  $\text{loop } \text{nxt}_M(q_0, t)(Si) = \text{Some } t' \leftrightarrow \exists q'. M(q, t) \triangleright^i (q', t')$

The certifying extraction framework [6] automatically extracts L-terms  $s_f$  for functions  $f$  and proves that  $s_f$  computes  $f$ . Additionally, one can pass a time complexity function  $\tau_f$  and is then presented with proving certain recurrence equations for  $\tau_f$ . Since we do not implement higher-order functions, we can give a simplified account of the framework here.

Central in the framework are Scott encodings, which are used to encode elements of arbitrary, first-order types as L-terms. The idea behind Scott encodings is that case analysis is by application: For instance, **if**  $b$  **then**  $a_1$  **else**  $a_2$  corresponds to the L-term  $\bar{b} s_{a_1} s_{a_2}$ , where  $s_{a_1}$  and  $s_{a_2}$  compute  $a_1$  and  $a_2$  respectively.

In general, for types  $A_1, \dots, A_n, B$  with a Scott encoding, a function  $f : A_1 \rightarrow \dots A_n \rightarrow B$ , is computed by a term  $s_f$  with time complexity function  $\tau_f : A_1 \rightarrow \dots \rightarrow A_n \rightarrow \mathbb{N}$  if

$$\forall a_1, \dots, a_n. \exists k \leq \tau_f a_1 \dots a_n. s_f \bar{a}_1 \dots \bar{a}_n \triangleright^k \overline{f a_1 \dots a_n}$$

The framework also supports higher-order functions and currying, but we omit those features complicating e.g. the definition of time complexity since we do not rely on them.

The certifying extraction framework comes with a library of computability proofs including time complexity, covering natural numbers, list, and vectors. Furthermore, one can give a general computability proof for functions with listable domain. I.e. if there is  $[x_1, \dots, x_n] : \mathbb{L}X$  s.t.  $\forall x : X. x \in [x_1, \dots, x_n]$  and  $f : X \rightarrow Y$ , then there is  $s_f$  and a constant  $c$  s.t.  $\forall x. \exists i \leq c \cdot n. s_f \bar{x} \triangleright^i \overline{fx}$ .

Since  $\delta_M$  has a listable domain, we can use the certifying extraction framework to extract  $\text{nxt}$  for every  $M$ :

► **Fact 14.** *Let  $M : \text{TM}_\Sigma^n$ . There is  $\text{nxt}_M : \text{tm}_L$  and  $C_{\text{nxt}} : \mathbb{N}$  s.t.  $s_{\text{nxt}_M}(\bar{q}, \bar{t}) \triangleright^{C_{\text{nxt}}} \overline{\text{nxt}(q, t)}$ .*

**Proof.** The framework generates the proof obligations  $52 \cdot |Q_M|^2 + 56 \leq C_{\text{nxt}}$ , and  $52 \cdot |Q_M|^2 + 130 \cdot n + 216 \leq C_{\text{nxt}}$ , where  $|Q_M|$  is the number of states of  $M$ . If  $C$  is picked large enough before running extraction, the obligations can be discharged automatically by the tactic `solverec` provided by the framework. ◀

We now define a term  $s_{\text{loop}}$  which expects  $f$  and  $x$  and loops  $f$  until a value  $y$  is found, or indefinitely if not. Since the extraction framework only covers total functions, we have to manually implement  $s_{\text{loop}}$ . To do so, we rely on a recursion combinator  $\rho$  [12], also employed by the framework to use recursion:

## 19:14 Mechanised Time Invariance Thesis for L

► **Fact 15.** *If  $s$  and  $s'$  are closed abstractions,  $\rho s s' \succ^3 s(\rho s)s'$ .*

We then define  $s_{\text{loop}}$  with time complexity function  $\tau_{\text{loop}} : (X \rightarrow \mathbb{N}) \rightarrow X \rightarrow \mathbb{N} \rightarrow \mathbb{N}$ :

$$\begin{aligned} s_{\text{loop}} &:= \rho(\lambda r f x. f x(\lambda x' z. r f x')(\lambda y z. y)(\lambda z. z)) \\ \tau_{\text{loop}} \tau_f x 0 &:= 0 \quad \tau_{\text{loop}} \tau_f x (Si) := \tau_f x + 11 + \text{if } f x \text{ is inl } x' \text{ then } \tau_{\text{loop}} \tau_f x i \text{ else } 0 \end{aligned}$$

► **Lemma 16.** *Let  $f$  be computable by  $s_f$  with time complexity function  $\tau_f$ , i.e.  $\forall x. \exists i \leq \tau_f x. s_f \bar{x} \triangleright^i \bar{f}x$ . Then*

1. *If  $\text{loop } f x i = \text{Some } y$ , then  $\exists k \leq \tau_{\text{loop}} \tau_f x i. s_{\text{loop}} s_f \bar{x} i \triangleright^k \bar{y}$ .*
2. *If  $s_{\text{loop}} s_f \bar{x}$  terminates, there exist  $i$  and  $y$  s.t.  $\text{loop } f x i = \text{Some } y$ .*

This suffices to prove one direction of the time invariance thesis w.r.t. simulation:

► **Theorem 17.** *There is  $s_{\text{sim}} : \text{tm}_L$  s.t. for all  $M : \text{TM}_{\Sigma}^n$  there is  $C : \mathbb{N}$  s.t. for all  $\mathbf{t} : \text{tp}_{\Sigma}^n$*

1. *If  $M(q_0, \mathbf{t}) \triangleright^i (q, \mathbf{t}')$ , then  $\exists j \leq C \cdot i + C. s_{\text{sim}} s_{\text{next}_M} \bar{\mathbf{t}} \triangleright^j \bar{\mathbf{t}}'$ .*
2. *If  $s_{\text{sim}} s_{\text{next}_M} \bar{\mathbf{t}} \triangleright v$ , then  $\exists q \mathbf{t}'. M(q_0, \mathbf{t}) \triangleright (q, \mathbf{t}')$ .*

## 7 TM-computable relations over $\mathbb{L}\mathbb{B}$ are L-computable

Let  $R \subseteq (\mathbb{L}\mathbb{B})^k \times \mathbb{L}\mathbb{B}$  be computable by  $M : \text{TM}_{\Sigma}^n$ . We define an L-term computing  $R$  by taking  $l_1, \dots, l_k$  as input, converting them to their respective TM-encoding, and then running  $M$  with the help of  $s_{\text{sim}}$ . Step-by-step,  $s$  has to:

1. Expect input in the form  $s \overline{l_1 \dots l_k}$ ,
2. for every  $1 \leq i \leq k$  compute  $\text{midtp } [] \text{ bl } l_i$ , i.e. the L-encoding of the TM-encoding of  $l_i$ .
3. run the simulation  $s_{\text{sim}} s_{\text{next}_M} [\text{niltp}, t_1, \dots, t_k, \text{niltp}, \dots, \text{niltp}]$ .
4. this computation will (if it terminates) terminate with a value  $\overline{(\text{midtp } [] \text{ bl } t'_1 \dots t'_n)}$ ,
5. meaning  $s$  has to output  $\bar{l}$ .

Three challenges arise: the term  $s$  has to be defined parametric in  $k$ , the L-encoding of the lists  $l_1, \dots, l_k$  has to be converted to the L-encoding  $[\text{niltp}, t_1, \dots, t_k, \text{niltp}, \dots, \text{niltp}]$ , and the L encoding of a result  $\mathbf{t}'$  has to be analysed, and the TM-encoding of a list  $l$  contained  $\mathbf{t}'[0]$  has to be converted to the L-encoding of  $l$ .

For the first task, we implement  $k$ -ary substitutions and combinators.

► **Fact 18.** *One can define functions  $s_{\mathbf{u}}^n : \text{tm}_L$  where  $s : \text{tm}_L, n : \mathbb{N}, \mathbf{u} : \text{tm}_L^k$  and*

$$\lambda_k : \text{tm}_L \rightarrow \text{tm}_L \quad \text{app}_k : \text{tm}_L \rightarrow \text{tm}_L^k \rightarrow \text{tm}_L \quad \text{vars}_k : \text{tm}_L^k$$

such that the following hold:

1.  $\text{vars}_{S_k} = k :: \text{vars}_k$ ,
2.  $(\text{app}_k s(s_1, \dots, s_k))_{\mathbf{u}}^n = \text{app}_k(s_{\mathbf{u}}^n)((t_1)_{\mathbf{u}}^n, \dots, (t_k)_{\mathbf{u}}^n)$ ,
3. *if all elements of  $\mathbf{u}$  are closed abstractions, then  $\text{app}_k(\lambda_k s)\mathbf{u} \succ^k s_{\mathbf{u}}^0$ .*

The second and third tasks can again be done by extraction.

► **Fact 19.** *There is a closed abstraction  $s_{\text{prep}}$  s.t.  $s_{\text{prep}} \overline{(l_1, \dots, l_k)} \triangleright [\text{niltp}, t_1, \dots, t_k, \text{niltp}, \dots, \text{niltp}]$ , where  $t_i := \text{midtp } [] \text{ bl } [\bar{x}_1, \dots, \bar{x}_n]$  for  $l_i = [x_1, \dots, x_n]$ .*

► **Fact 20.** *There is a closed abstraction  $s_{\text{unenc}_{\text{TM}}}$  s.t. if  $\mathbf{t}[0] = \text{midtp } [] \text{ bl } [\bar{x}_1, \dots, \bar{x}_m]$  and  $l = [x_1, \dots, x_n]$  we have  $s_{\text{unenc}_{\text{TM}}} \bar{\mathbf{t}} \triangleright \text{Some } \bar{l}$  and  $s_{\text{unenc}_{\text{TM}}} \bar{\mathbf{t}} \triangleright \text{None}$  otherwise.*

► **Theorem 21.** *If  $R \subseteq (\mathbb{L}\mathbb{B})^k \times \mathbb{L}\mathbb{B}$  is TM-computable by a machine  $M$  with time complexity relation  $\tau$  there are a term  $s_M$  and a constant  $c$  s.t.  $s_M$  computes  $R$  with time complexity relation  $(m, n_1, \dots, n_k) \mapsto c \cdot m \cdot n_1 \cdots n_k \cdot \tau(m, n_1, \dots, n_k) + c$ .*

**Proof.** Define  $s_M := \lambda k. s_{\text{unenc}_{\text{TM}}}(s_{\text{sim}} s_{\text{next}_M} (s_{\text{prep}}(s_{\text{cons}} k(\dots (s_{\text{cons}} 0 \bar{[]})))))(\lambda x x) \bar{[]}$ . ◀

## 8 Hoare logic verification framework for Turing machines

Although the relational verification approach is quite powerful, it suffers from two crucial problems. First, realisation and termination are inherently separated proof goals, even for total machines, which form the majority of machines we are interested in. The premises of the correctness and termination relations are usually almost the same, thus manipulation of premises is duplicated. Secondly, in an interactive proof the proof context can grow very large. When two machines with say 9 tapes are sequentially composed, the proof context contains 9 assumptions on the initial tapes, 9 assumptions for the intermediate tapes between the two machines, and 9 final propositions on the tapes. Many of these assumptions are equalities. The Turing machine verification framework has naive tactics to simplify such proof goals by substitution and rewriting, which becomes quadratically slower the longer a realisation proofs takes.

We propose a verification framework based on Hoare logic, solving both problems. Correctness and termination can be proved at once, eliminating the need for repeated manipulation of premises. Furthermore, at each step in the verification of an  $n$  tape TM, the user sees exactly  $n$  specifications  $S$  for the  $n$  tapes, plus optionally a custom invariant depending on both the tapes and the label.

We here give a high-level overview. More details are in the separate Appendix B [8].

The Hoare logic is built as a new layer of abstraction for the relational framework. Given  $M : \text{TM}_{\Sigma}^n(L)$ , a predicate  $P \subseteq \text{tp}_{\Sigma}^n$  and a relation  $Q \subseteq L \times \text{tp}_{\Sigma}^n$ , we write *weak* Hoare triples:

$$\models \{P\} M \{Q\} := M \models (\lambda \mathbf{t} (\ell, \mathbf{t}'). P \mathbf{t} \rightarrow Q (\ell, \mathbf{t}'))$$

To state that the machine is functionally correct and terminates in a certain time, we use *total* Hoare triples. In addition to the relation of the weak Hoare triple, a total Hoare triple asserts that the machine terminates in (at most)  $i$  steps if the precondition is satisfied. Thus, we avoid spelling out the precondition again.

$$\models \{P\}^i M \{Q\} := (\models \{P\} M \{Q\}) \wedge M \downarrow (\lambda \mathbf{t} i'. P \mathbf{t} \wedge i \leq i')$$

Similar to the old verification framework, every building block like `While` and `Switch` comes with an associated Hoare triple, shown in Figure 1 in the separate Appendix B [8]. Using these rules, an interactive verification is akin to a symbolic execution of machines with explicitly annotated invariants.

The Hoare triples of user-defined machines  $M : \text{TM}_{\Sigma}^n$  exclusively have triples with both pre and post conditions of the form  $S_1 \wedge \dots \wedge S_n \wedge I$  where  $S_i$  for  $1 \leq i \leq n$  are either  $\mathbf{t}[i] \simeq x$  for some  $x$ , `isVoid`  $\mathbf{t}[i]$ , or  $\mathbf{t}[i] = t_0$  for some fixed  $t_0$ , and  $I$  is a custom, user-chosen invariant. Thus, the following specification for a binary machine

$$\begin{aligned} M \models \lambda \mathbf{t} (\ell, \mathbf{t}'). \forall (x : X) (y : Y). P_X x \rightarrow P_Y y \rightarrow \\ \mathbf{t}[0] \simeq x \rightarrow \mathbf{t}[1] \simeq y \rightarrow \mathbf{t}'[0] \simeq f(x, y) \wedge \mathbf{t}'[1] \simeq g(x, y) \\ M \downarrow \lambda \mathbf{t} i. \exists (x : X) (y : Y). P_X x \wedge P_Y y \wedge \mathbf{t}[0] \simeq x \wedge \mathbf{t}[1] \simeq y \wedge i \geq \tau xy \end{aligned}$$

can be compactly restated using the following triple:

$$\forall xy. P_X x \rightarrow P_Y y \rightarrow \models \{\lambda \mathbf{t}. \mathbf{t}[0] \simeq x \wedge \mathbf{t}[1] \simeq y\}^{\tau xy} M \{\lambda (\ell, \mathbf{t}'). \mathbf{t}'[0] \simeq f(x, y) \wedge \mathbf{t}'[1] \simeq g(x, y)\}$$

A typical workflow for the verification of total machines then looks as follows: First, the user defines Hoare triples for their machines, following the shape  $S_1 \wedge \dots \wedge S_n \wedge I$ . Secondly, they perform the correctness proof. Thirdly, they define the time complexity function of the machine, and add the running time to the triple in the proved lemma. Fourthly, they replay the proof script and in the very last step of the verification, show that the accumulated running time is indeed bounded by the time complexity function.

## 9 L-computable relations over $\mathbb{L}\mathbb{B}$ are TM-computable

To convert a term  $s$  computing a relation  $R$ , we follow the same strategy as in Section 7. We exemplarily show the specification of the machine  $M_{\text{conv}}$  corresponding to  $s_{\text{unencTM}}$ :

► **Fact 22.** *There is a machine  $M_{\text{unencL}} : \text{TM}_{\Sigma}^4$  and a constant  $c$  s.t. for all  $l : \mathbb{L}\mathbb{B}$  we have  $\models \{P_l\}^{c \cdot |l|} M_{\text{unencL}} \{Q_l\}$  where  $P_l := \lambda t. \mathbf{t}[0] \simeq \bar{l} \wedge \mathbf{t}[1] = \text{niltp} \wedge \text{isVoid } \mathbf{t}[2] \wedge \text{isVoid } \mathbf{t}[3]$  and  $Q_l := \lambda(\_, t'). \mathbf{t}[1] = \text{midtp} [] \text{ bl } \bar{l} \wedge \text{isVoid } \mathbf{t}[2] \wedge \text{isVoid } \mathbf{t}[3]$ .*

► **Theorem 23.** *If  $R \subseteq (\mathbb{L}\mathbb{B})^k \times \mathbb{L}\mathbb{B}$  is L-computable by a closed term  $s$  with time complexity relation  $\tau$  there is a polynomial  $p$  and a Turing machine  $M_s$  s.t.  $M_s$  computes  $R$  with time complexity relation  $(m, n_1, \dots, n_k) \mapsto p(m, n_1, \dots, n_k, \tau_R(m, n_1, \dots, n_k))$ .*

**Proof.** The machine  $M_s$  is defined parametric in  $s$  and by recursion on  $k$ , i.e. we define a different machine for each  $k$ . First,  $M_s$  reads its input and writes the task stack  $[(\gamma(s \bar{n}_1 \bar{n}_k), 0)]$  to an auxiliary tape. As a total subroutine, we verify this part using the Hoare framework. Then  $M_s$  runs  $M_{\text{sim}}$ . If the computation terminates, the resulting heap will contain a term  $\bar{l}$  for a list  $l : \mathbb{L}\mathbb{B}$ . Thus,  $M_s$  runs a machine implementing  $\text{unf}$  and lastly runs  $M_{\text{unencL}}$ .

The size of the heap after  $i$  steps is in  $\mathcal{O}(i \cdot (i + N))$ , where  $N := n_1 + \dots + n_k$ . For one reduction step,  $N$  variables might have to be looked up, resulting in a runtime of  $\mathcal{O}(i \cdot (i + N + 1) \cdot (N + 1))$  per step. Unfolding a heap takes  $\mathcal{O}((m + 1) \cdot (H + N + 1))$ , where  $H$  is the size of the heap. In total, we can thus define the polynomial  $p$  cubic in the number of steps  $i$ , quadratic in  $N$  and linear in the size of the output  $m$  as follows, where  $c$  is a constant:

$$(m, n_1, \dots, n_k, i) \mapsto c \cdot (i + 1) \cdot (i + n_1 + \dots + n_k + 1) \cdot ((n_1 + \dots + n_k + 1) \cdot (i + 1) + m). \quad \blacktriangleleft$$

## 10 Discussion

We have presented the first mechanised proof of an instance of the time invariance thesis, connecting L with Turing machines with a polynomial overhead in time. We prove two variants, respectively concerned with simulation of one model of computation on the other, and with the computability of relations on boolean strings. In total, including dependencies, our development consists of 30.000 LoC and takes about 18 minutes to compile on an Intel Core i7-6600U CPU @ 2.60GHz machine. The novel code for this paper still constitutes 7800 LoC, with about 40% specification and 60% proofs.

It is folklore that the two variants are equivalent. For instance, in [1, 7, 23] more emphasis is put on the simulation variant. In the interactive theorem proving community it is however well-known that “folklore” is not equivalent to “easy to mechanise”. In the setting of computability and complexity theory, where one has to deal with models of computation, this fact is amplified: proofs on paper virtually never provide concrete implementations in a model of computation, but focus on the high-level invariants of a proof. This is based on the (again folklore) fact that algorithms can be implemented in models of computation provided enough time and strength to sustain the tedium to do so. However, the proof engineering time needed to show the correctness of an algorithm (as e.g. a Coq function) and to show the correctness of its implementation (as e.g. a Turing machine) are completely independent. Correctness proofs of algorithms depend on the intricacy of invariants. Correctness proofs of implementations depend on the length of the code of a Turing machine, and the size of the gap between specification language and Turing machines. If the specification is a first-order, tail-recursive Coq function, the length of the Turing machine is the main factor.

We however hope that future mechanised proofs of results in complexity theory and of the time invariance thesis for other models of computation can profit from our development.



For us, simulating Turing machines on  $L$  and simulating  $L$  on Turing machines was the only possible way. Now that  $L$  is shown reasonable for time complexity, future proofs of the time invariance thesis can choose which reasonable model to use for each direction of the proof. For the time invariance thesis for a model of computability  $C$ , one can show that  $C$  can simulate Turing machines, which are structurally simple and where computation is local. In the other direction, one can show that  $C$  can be simulated in  $L$ , which has rich structure and supports non-local computation on almost arbitrary (first-order) data-structures.

For our concrete cases of Turing machines and  $L$ , powerful abstraction layers and verification frameworks are necessary for mechanisation. We would assess that a “manual” approach to any invariance thesis involving Turing machines is completely unfeasible. The Turing machine verification framework [9] and the certifying extraction framework [6] proved very valuable in this regard. While having similar goals, they use different mathematical approaches to both verification and time complexity analysis.

First, the certifying  $L$ -extraction framework gives support to automatically prove the computability of a large subset of Coq functions. For time complexity, a user has to give a time complexity function as input, and the framework automatically generates equations the function has to fulfil and furthermore provides tactics to solve these equations. Finding a time complexity function can prove challenging, but an interactive approach where a wrong function is picked and a correct function is reverse-engineered from the equations works well. A priori, the framework does not support partial terms. The `Lsimpl` tactic used in the framework however can be used to normalise terms in manual verification.

Secondly, the Turing machine verification framework provides tools to verify the correctness, time and even space complexity of Turing machines, but the user has to implement those machines manually. For the implementation, the user can write Turing machines in the style of a register based while-language, for which a canonical realisation and termination relation are inferred automatically. Correctness and time complexity proofs w.r.t. user-defined relations are now simply inclusion proofs between the canonical and the user-defined relations. Due to the split of correctness and termination, the framework works well for the verification of partial machines. However, for total machines, termination and correctness are distinct proof goals and have to be proved separately, which leads to a mathematical duplication of similar proof goals, sometimes even to actual proof code duplication.

The novel Hoare logic verification framework we present remedies this situation: For total machines, only one proof subsuming both correctness and time complexity has to be carried out, while it still supports separate proof goals for partial machines. No canonical relations are used. The verification of Hoare triples is carried out directly on the implementation of the machine, using the proof rules for the used combinators and user-defined machines. We conjecture that the Hoare logic framework scales further and possibly far. Our simulation of  $L$  on Turing machines is reasonable on time, but if terms exhibit pointer explosion [7], the space overhead might not be constant factor. The proof of this stronger time and space invariance thesis for  $L$  [7] is considerably more complicated, but might now be in reach.

However, the assessment that “Turing machines as model of computation are inherently infeasible for the formalisation of any computability or complexity theoretic result” [9] still stands. Future developments of mechanised results in complexity theory can and should be based on  $L$  or similarly well-suited models. Our certifying extraction framework provides valuable help in doing so, but there are interesting extensions worth exploring in the future.

First, this framework does not cover space complexity, which would however be needed for a proof of the full invariance thesis for  $L$ . Secondly, exploring an automatic generation of time-complexity functions, where the user can then provide a (polynomial) upper bound

for the function in a second step might be interesting. Lastly, the framework requires user input for every single recursive function used, also for auxiliary functions, and the automatic verification of these is based on `Ltac` tactics which sometimes fail, and sometimes are slow. Here, an automatic generation of correctness proofs based on a meta-programming tool for Coq like MetaCoq [24, 25] might be a possible solution.

---

## References

- 1 Beniamino Accattoli and Ugo Dal Lago. (Leftmost-Outermost) Beta Reduction is Invariant, Indeed. *Logical Methods in Computer Science*, 12(1), 2016. doi:10.2168/LMCS-12(1:4)2016.
- 2 Andrea Asperti and Wilmer Ricciotti. Formalizing Turing Machines. In *Logic, Language, Information and Computation*, pages 1–25. Springer, 2012.
- 3 Guy E. Blelloch and John Greiner. Parallelism in Sequential Functional Languages. In *Proceedings of the seventh international conference on Functional programming languages and computer architecture, FPCA 1995, La Jolla, California, USA, June 25-28, 1995*, pages 226–237, 1995. doi:10.1145/224164.224210.
- 4 Alonzo Church. A set of postulates for the foundation of logic. *Annals of mathematics*, pages 346–366, 1932.
- 5 Alonzo Church. An unsolvable problem of elementary number theory. *American journal of mathematics*, 58(2):345–363, 1936.
- 6 Yannick Forster and Fabian Kunze. A Certifying Extraction with Time Bounds from Coq to Call-By-Value Lambda Calculus. In John Harrison, John O’Leary, and Andrew Tolmach, editors, *10th International Conference on Interactive Theorem Proving (ITP 2019)*, volume 141 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:19, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ITP.2019.17.
- 7 Yannick Forster, Fabian Kunze, and Marc Roth. The weak call-by-value  $\lambda$ -calculus is reasonable for both time and space. *Proceedings of the ACM on Programming Languages*, 4(POPL):1–23, 2019.
- 8 Yannick Forster, Fabian Kunze, Gert Smolka, and Maximilian Wuttke. Appendix b: A mechanised proof of the time invariance thesis for the weak call-by-value  $\lambda$ -calculus. Technical report, Saarland University, 2021. URL: [https://www.ps.uni-saarland.de/Publications/documents/AppendixB\\_ForsterEtAl\\_2021\\_Mechanised-Time-Invariance-L.pdf](https://www.ps.uni-saarland.de/Publications/documents/AppendixB_ForsterEtAl_2021_Mechanised-Time-Invariance-L.pdf).
- 9 Yannick Forster, Fabian Kunze, and Maximilian Wuttke. Verified programming of Turing machines in Coq. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 114–128, 2020.
- 10 Yannick Forster and Dominique Larchey-Wendling. Certified undecidability of intuitionistic linear logic via binary stack machines and minsky machines. In *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 104–117, 2019.
- 11 Yannick Forster, Dominique Larchey-Wendling, Andrej Dudenhefner, Edith Heiter, Dominik Kirst, Fabian Kunze, Gert Smolka, Simon Spies, Dominik Wehr, and Maximilian Wuttke. A Coq library of undecidable problems. In *The Sixth International Workshop on Coq for Programming Languages (CoqPL 2020)*, 2020. URL: <https://github.com/uds-psl/coq-library-undecidability>.
- 12 Yannick Forster and Gert Smolka. Weak call-by-value lambda calculus as a model of computation in Coq. In *International Conference on Interactive Theorem Proving*, pages 189–206. Springer, 2017.
- 13 John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation (3rd Edition)*. Addison-Wesley Longman Publishing Co., Inc., USA, 2006.

- 14 Jan Martin Jansen. *Programming in the  $\lambda$ -Calculus: From Church to Scott and Back*, pages 168–180. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. doi:10.1007/978-3-642-40355-2\_12.
- 15 Fabian Kunze, Gert Smolka, and Yannick Forster. Formal small-step verification of a call-by-value lambda calculus machine. In *Asian Symposium on Programming Languages and Systems*, pages 264–283. Springer, 2018.
- 16 Ugo Dal Lago and Beniamino Accattoli. Encoding Turing machines into the deterministic lambda-calculus. *arXiv preprint arXiv:1711.10078*, 2017.
- 17 Ugo Dal Lago and Simone Martini. The weak lambda calculus as a reasonable machine. *Theor. Comput. Sci.*, 398(1-3):32–50, 2008. doi:10.1016/j.tcs.2008.01.044.
- 18 Dominique Larchey-Wendling and Yannick Forster. Hilbert’s tenth problem in Coq. *arXiv preprint arXiv:2003.04604*, 2020.
- 19 Torben Æ. Mogensen. Efficient self-interpretation in lambda calculus. *JOURNAL OF FUNCTIONAL PROGRAMMING*, 2:345–364, 1994.
- 20 Joachim Niehren. Uniform confluence in concurrent computation. *Journal on Functional Programming*, 10(5):453–499, 2000.
- 21 Gordon D. Plotkin. Call-by-name, call-by-value and the  $\lambda$ -calculus. *Theoretical computer science*, 1(2):125–159, 1975.
- 22 Damien Pous. A certified compiler from recursive functions to minski machines. Technical report, Université Paris-Diderot, PPS, 2004. URL: <http://perso.ens-lyon.fr/damien.pous/recursive-minski/>.
- 23 Cees F. Slot and Peter van Emde Boas. On Tape Versus Core; An Application of Space Efficient Perfect Hash Functions to the Invariance of Space. In *Proceedings of the 16th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1984, Washington, DC, USA*, pages 391–400, 1984. doi:10.1145/800057.808705.
- 24 Matthieu Sozeau, Abhishek Anand, Simon Boulier, Cyril Cohen, Yannick Forster, Fabian Kunze, Gregory Malecha, Nicolas Tabareau, and Théo Winterhalter. The MetaCoq Project. *Journal of Automated Reasoning*, 2020. doi:10.1007/s10817-019-09540-0.
- 25 Matthieu Sozeau, Simon Boulier, Yannick Forster, Nicolas Tabareau, and Théo Winterhalter. Coq Coq correct! verification of type checking and erasure for Coq, in Coq. *Proceedings of the ACM on Programming Languages*, 4(POPL):1–28, 2019.
- 26 The Coq Development Team. The coq proof assistant, version 8.12.0, July 2020. doi:10.5281/zenodo.4021912.
- 27 Alan Mathison Turing. On computable numbers, with an application to the Entscheidungsproblem. *J. of Math*, 58(345-363):5, 1936.
- 28 Maximilian Wuttke. Verified programming of Turing machines in Coq. Bachelor’s thesis, Saarland University, 2018.
- 29 Jian Xu, Xingyuan Zhang, and Christian Urban. Mechanising Turing machines and computability theory in Isabelle/HOL. In *International Conference on Interactive Theorem Proving*, pages 147–162. Springer, 2013.

## **A** Definition of Turing machines

We compare the definition of Turing machines we use [9] to the one by Hopcroft, Motwani, and Ullman [13].

1. The logical system in [13] is classical set theory, whereas we work in constructive type theory. We discuss the impact of this foundational choice below.
2. The alphabet in [13] is separated into a set of input symbols  $\Sigma$  and a superset of tape symbols  $\Gamma$ , where we unify both into a single type.

3. The blank symbol in [13] is an explicit part of Turing machines as an element of  $\Gamma$ , but not of  $\Sigma$ . We do not specify blank symbols explicitly and instead leave it to a user to specify a blank symbol or even various blank symbols.
4. Tapes in [13] are not formally defined. It is only stated that tapes ‘extend infinitely to the left and right, initially hold[ing] [...] the input’. Instantaneous descriptions of Turing machines are formally defined as strings over  $\Gamma$  and  $Q$ , which contain ‘the portion of the tape between the leftmost and the rightmost non-blank, unless the head is to the left of the leftmost non-blank or to the right of the rightmost non-blank.’ In the latter case, the blanks between the head and the non-blank content are part of the string.
5. The transition function in [13] is a partial function, whereas ours is total. If the transition function is unspecified, the computation of the machine halts, whereas we have an explicit boolean halting function. In Coq’s type theory one requires classical logic and  $\text{AUC}_{\mathbb{N},\mathbb{N}}$  to compile Turing machines with a partial transition function into an equivalent definition with total transition functions. In general, any compilation of partial functions on finite types to total functions is non-computable and thus not definable in Coq’s type theory without axioms.
6. A machine in [13] always writes a symbol and always moves to the left or right. We allow to not write a symbol and to not move the head. The first is important regarding our definition of tapes (otherwise a fully empty `niltp` can never stay fully empty), whereas the second is a relatively arbitrary choice to allow more freedom in the definition of concrete machines.
7. Turing machines have an explicit set of accepting states in [13]. We do not add one, because our definition is not aimed at formalising computability theory directly. Instead, our more flexible definition of labels subsumes the binary notion of accepting states, but allows for more interesting constructions like the `Switch` and `MemWhile` machines.

Subtle difficulties might arise when defining notions of computability theory in classical set theory. For instance, defining Turing machines with a transition function which takes as input the whole tape is problematic: Nothing permits non-computable transition functions then, making arbitrary problems decidable by encoding the decision into the transition function. When imposing the transition function to be computable one obtains a circular dependency: The notion of computability is needed to define transition functions, transition functions are needed to define Turing machines, but Turing machines are needed to define the notion of computability. The well-known solution here is to define Turing machines as finite objects, i.e. let the transition function, although partial, work on a finite domain and codomain. In classical set theory one can then always show afterwards that this transition function is computable, since every function with finite domain and codomain is. In our type-theoretic setting, we can also show that the transition function is computable, provided it is, as we defined it, a total function. We did so in Section 6.