

The Weak Call-By-Value λ -Calculus is Reasonable for Both Time and Space

Yannick Forster

Saarland University, Saarland Informatics Campus (SIC), Saarbrücken, Germany
forster@ps.uni-saarland.de

Fabian Kunze

Saarland University, Saarland Informatics Campus (SIC), Saarbrücken, Germany
kunze@ps.uni-saarland.de

Marc Roth 

Cluster of Excellence (MMCI), Saarland Informatics Campus (SIC), Saarbrücken, Germany
mroth@mmci.uni-saarland.de

Abstract

We study the weak call-by-value λ -calculus as a model for computational complexity theory and establish the natural measures for time and space – the number of beta-reductions and the size of the largest term in a computation – as reasonable measures with respect to the invariance thesis of Slot and van Emde Boas [STOC 84]. More precisely, we show that, using those measures, Turing machines and the weak call-by-value λ -calculus can simulate each other within a polynomial overhead in time and a constant factor overhead in space for all computations that terminate in (encodings) of “*true*” or “*false*”. We consider this result as a solution to the long-standing open problem, explicitly posed by Accattoli [ENTCS 18], of whether the natural measures for time and space of the λ -calculus are reasonable, at least in case of weak call-by-value evaluation.

Our proof relies on a hybrid of two simulation strategies of reductions in the weak call-by-value λ -calculus by Turing machines, both of which are insufficient if taken alone. The first strategy is the most naive one in the sense that a reduction sequence is simulated precisely as given by the reduction rules; in particular, all substitutions are executed immediately. This simulation runs within a constant overhead in space, but the overhead in time might be exponential. The second strategy is heap-based and relies on structure sharing, similar to existing compilers of eager functional languages. This strategy only has a polynomial overhead in time, but the space consumption might require an additional factor of $\log n$, which is essentially due to the size of the pointers required for this strategy. Our main contribution is the construction and verification of a space-aware interleaving of the two strategies, which is shown to yield both a constant overhead in space *and* a polynomial overhead in time.

2012 ACM Subject Classification Theory of computation \rightarrow Computational complexity and cryptography; Mathematics of computing \rightarrow Lambda calculus

Keywords and phrases invariance thesis, lambda calculus, weak call-by-value reduction, time and space complexity, abstract machines

Supplement Material A Coq formalisation of the results in Section 2 and Section 3 is available at <https://ps.uni-saarland.de/extras/wcbv-reasonable>.

1 Introduction

Turing machines are the de-facto foundation of modern computability and complexity theory, in part due to the conceptual simplicity of their definition. However, this simplicity is also one of the biggest disadvantages: When it comes to detailed or formal reasoning, Turing machines soon become impossible to treat, because they lack compositionality and heavy logical machinery has to be used to reason about them. This is best reflected by the fact that modern day researchers in computability and complexity theory usually have not faced

explicit Turing machines since their undergraduate studies. Instead, it is common to rely on pseudo code or mere algorithmic descriptions. For computability theory, other models of computation like RAM machines, recursive functions or variants of the λ -calculus can be used if details are of interest, because the notion of computation is invariant under changing the model. Especially the λ -calculus shines in this aspect, because tree-like inductive datatypes can be directly encoded and equational reasoning is accessible to verify the correctness of programs, which even makes the λ -calculus feasible as a model to formalise computability theory in proof assistants [19, 12]. However, this notion of *invariance* does not suffice for complexity theory. As stated by Slot and van Emde Boas [23]:

“Reasonable” machines can simulate each other within a polynomially bounded overhead in time and a constant factor overhead in space.

If only reasonable machines are considered, this invariance thesis makes complexity classes robust under changing the model of computation. Until now, only sequential models of computation have been shown to fulfil this strong notion of invariance with natural complexity measures for time and space [10]. The time and space complexity measures known to be reasonable for the full λ -calculus are “total ink used” and “maximum ink used” in a computation [18]. While the notion for space is natural, the notion for times is very unnatural and of no real interest. Other measures rely on concrete implementations, giving no satisfying answer to the question whether the λ -calculus can be considered reasonable.

Dal Lago and Martini [16] gave a preliminary result in 2008 for the weak call-by-value λ -calculus and showed that counting β -steps while taking the size of β -redices into account is a reasonable measure for time. In 2014 Accattoli and Dal Lago [6] showed that counting (leftmost-outermost) β -steps makes the full λ -calculus reasonable for time, starting a long line of research regarding measures for and implementations of the λ -calculus (see e.g. [1]). Whether the natural measure for space, i.e. the size of the largest term in a computation, can be used together with the number of β -steps or how it has to be adapted is a long-standing open problem.

We solve this problem for the deterministic weak call-by-value λ -calculus (which we call **L**) and show that the size of the largest intermediate term in a reduction makes **L** a reasonable machine in the strong sense. We consider our solution more than just a partial solution on the way to answering the question for the full λ -calculus in several aspects: First, weak call-by-value evaluation is the standard model of eager functional programming languages. Second, **L** is already Turing-complete and one does not gain more power for the implementation or verification of algorithms by the strong reduction allowed in the full λ -calculus. Third, from the complexity-theoretic point of view, the problem is solved: A certain form of the λ -calculus can be used to spell out arguments. However, from an implementation point of view, many questions remain open: Our simulation uses the freedom given by “polynomial overhead” and should not be seen as a proposal for a uniform, canonical implementation, which is still very much desirable.

In what follows, we explain how to simulate **L** on Turing machines with polynomial overhead in time and linear overhead in space, based on the natural measures, and vice-versa.

► **Definition 1.** *For a closed term s that reduces to a normal form $\lambda x.u$*

$$s = s_0 \succ s_1 \succ \cdots \succ s_k = \lambda x.u$$

we define the time consumption of the computation to be $\|s\|_T = k$ and the space consumption to be $\|s\|_S = \max_{i=0}^k \|s_i\|$ where $\|s\|$ is the size of s .

Result and Discussion We prove that the weak call-by-value λ -calculus is a reasonable machine with respect to the natural time and space measures defined above. For the formal statement we fix a finite alphabet Σ and say that a function $f : \Sigma^* \rightarrow \Sigma^*$ is *computable* by L in time \mathcal{T} and space \mathcal{S} if there exists an L -term s_f such that for all $x \in \Sigma^*$ we have that

$$s_f \ulcorner x \urcorner \succ^* \ulcorner f(x) \urcorner \quad \text{and} \quad \|s_f \ulcorner x \urcorner\|_{\mathcal{T}} \leq \mathcal{T}(|x|) \quad \text{and} \quad \|s_f \ulcorner x \urcorner\|_{\mathcal{S}} \leq \mathcal{S}(|x|).$$

Here $\ulcorner \cdot \urcorner$ is an encoding of strings over Σ .

► **Theorem 2.** *Let Σ be a finite alphabet such that $\{\mathbf{true}, \mathbf{false}\} \subseteq \Sigma$ and let $f : \Sigma^* \rightarrow \{\mathbf{true}, \mathbf{false}\}$ be a function. Furthermore, let $\mathcal{T}, \mathcal{S} \in \Omega(n)$.*

1. *If f is L -computable in time \mathcal{T} and space \mathcal{S} , then f is computable by a Turing machine in time $\mathcal{O}(\text{poly}(\mathcal{T}(n)))$ and space $\mathcal{O}(\mathcal{S}(n))$.*
2. *If f is computable by a Turing machine in time \mathcal{T} and space \mathcal{S} , then f is L -computable in time $\mathcal{O}(\text{poly}(\mathcal{T}(n)))$ and space $\mathcal{O}(\mathcal{S}(n))$.*

The conditions $\mathcal{T} \in \Omega(n)$ and $\mathcal{S} \in \Omega(n)$ state that we do not consider sublinear time and space. Furthermore, the restriction of f to $\{\mathbf{true}, \mathbf{false}\}$ can be seen as a restriction to characteristic functions, which is sufficient for the complexity theory of decidability problems.

To the best of our knowledge this is the first proof of the invariance thesis including time and space for a fragment of the λ -calculus using the natural complexity measures.

At this point the reader might have the following objection: A well-known problem in the λ -calculus is the issue of size explosion. There exist terms that reduce to a term of size $\Omega(2^n)$ with only $\mathcal{O}(n)$ beta reductions. Let us adapt an example from [16]: Given a natural number n , we write \bar{n} for its Church-encoding, defined as

$$\bar{n} := \lambda f x. \underbrace{(f(f(f \dots (f x) \dots)))}_{n \text{ times}}$$

and define the Church encoding of the Boolean \mathbf{true} as $\overline{\mathbf{true}} := \lambda x y. x$. Next we define the term $s_E := \lambda x. \overline{\mathbf{true}} \overline{\mathbf{true}} (x \bar{2}(\lambda x. x))$. Note that the application $\bar{n} \bar{2}$ reduces to a normal form t_n of size $\Omega(2^n)$, extensionally equivalent to the Church-exponentiation $\bar{2}^n$. The term s_E thus encodes a function, computing an exponentially big intermediate result, discarding it and returning $\overline{\mathbf{true}}$. Formally, we have

$$s_E \bar{n} \succ^4 (\lambda y. \overline{\mathbf{true}}) \underbrace{(\bar{2}(\bar{2} \dots (\bar{2}(\lambda x. x))))}_{n \text{ times}} \succ \dots \succ (\lambda y. \overline{\mathbf{true}}) t_n \succ \overline{\mathbf{true}}$$

for a term t_n with $\|t_n\| \in \Omega(2^n)$. Now $\|s_E \bar{n}\|_{\mathcal{T}} \in \Theta(n)$, i.e., $s_E \bar{n}$ reduces to $\overline{\mathbf{true}}$ in about n beta reductions. Moreover, $\|s_E \bar{n}\|_{\mathcal{S}} \geq \|(\lambda y. \overline{\mathbf{true}}) t_n\| \in \Omega(2^n)$, i.e., the largest term in the reduction is of exponential size. While it might seem counterintuitive that a reasonable machine allows a computation that requires much more space than time, which is impossible for Turing machines, we consider this as one of the major insights of this work: such computations returning values of bounded size can always be optimised, i.e. the size explosion is unnecessary to compute the result of the term. Of course, any reduction sequence that *ends* in a term of exponential size cannot be simulated in less space if the result has to be written down explicitly. However, in complexity theory of decision problems, the functions that matter in the end are characteristic functions which map to \mathbf{true} and \mathbf{false} . We make this more precise: For any size-exploding term s_E , we can use Theorem 2 twice to obtain a term \widehat{s}_E computing the same function with polynomial space usage. By Theorem 2 (1) there is a Turing machine that on an encoding of \bar{n} simulates $s_E \bar{n}$ with time and space

complexity of $\mathcal{O}(n^c)$ (the latter since Turing machines can not use more space than time). By Theorem 2 (2), there is a term \widehat{s}_E s.t. $\widehat{s}_E \bar{n}$ has the same normal form as $s_E \bar{n}$ – but with space complexity $\mathcal{O}(n^c)$, since the overhead in space is constant-factor.

In [2], Accattoli writes “*Essentially one is assuming that space in the λ -Calculus is given by the maximum size of terms during evaluation, and since in sequential models time is greater or equal to space (because one needs a unit of time to use a unit of space), the time cost of the size exploding family must be at least exponential. The wrong hypothesis then is that space is the maximum size of the terms during evaluation. It is not yet clear what accounts for space in the λ -Calculus, however the study of time cost models made clear that space is not the size of the term.*” Our result implies that this conclusion does not apply for weak call-by-value evaluation in L . In this particular case, the wrong hypothesis is that even when simulating on Turing machines, where time is greater or equal to space, the space-measure of a λ -calculus term does not have to coincide with the actual resources used when simulating it and can thus be much larger than the time-measure. Furthermore, we point out that there exist reasonable sequential models that might consume asymptotically more space than time as illustrated in Appendix D in case of RAM machines.

Simulation Strategies In the previous paragraph we argued that our proposed cost and time measures for L are not inherently contradictory with respect to the invariance thesis. However, we did not provide explicit information of the simulations yet, which we are going to catch up on now. Note that a simulation of Turing machines in L for the second part of Theorem 2 regarding time has already been given by Accattoli and Dal Lago [5, 15]. We argue in Section 5 that their construction also only has a constant factor overhead in space and thus works for our purposes as well. The main part of the paper thus focuses on the simulation of L by Turing machines. Essentially, we rely on an interleaving of two different strategies for simulating a reduction in L with a Turing machine. Both strategies are formally introduced in Section 3; in what follows, we provide an intuitive overview.

The first one, which we call the *substitution-based strategy* simulates a reduction sequence naively as given by the reduction rules of L . In particular, all substitutions are executed immediately if a β -reduction is performed. However, we have already seen an example which shows this strategy to be insufficient: Consider again the term $s_E \bar{n}$ which reduces to $\overline{\text{true}}$ in $\Theta(n)$ beta reductions. If this reduction sequence is simulated naively, exponentially many substitutions have to be performed, and hence the time consumption of that machine would be exponential in n . At the same time, for this term, the strategy is valid if we would only care for space, because the space complexity of the L term is already exponential. In general, we show that any reduction sequence in L can be simulated with only a constant overhead in space by a Turing machine using the substitution-based strategy.

Solving the issue regarding the time consumption requires us to rely on the second simulation strategy which we call the *heap-based strategy*. Intuitively, we do not execute any substitution if a β -reduction is simulated. Instead we use *closures* and keep track of the values assigned to variables in an environment. These environments are stored on an explicit heap containing pointers and terms. This allows for *structure sharing*, similar to real-world execution of functional languages as well as to the strategy used in [6]. Indeed, applying this strategy to the reduction sequence $s_E \bar{n} \succ \dots \succ \overline{\text{true}}$ yields a polynomial number of steps in a simulation with a Turing machine.

At this point, one might be tempted to think that the heap-based strategy is strictly superior to the substitution-based strategy. However, there is one (major) catch: There exist reduction sequences of time and space linear in the input term size n , which yield an overhead

of factor $\log n$ in space when simulated using the heap-based strategy. The reason is that the number of heap entries is linear, which requires the pointers, i.e. the heap addresses, to grow in size. The following example illustrates this phenomenon: Let $N := (\lambda xy.xx)\overline{\text{true}}$.

$$s_P := \underbrace{N(\cdots(N\overline{\text{true}})\cdots)}_{n \text{ times}} \succ^n \underbrace{(\lambda y.\overline{\text{true}}\overline{\text{true}})(\cdots((\lambda y.\overline{\text{true}}\overline{\text{true}})\overline{\text{true}})\cdots)}_{n \text{ times}} \succ^{2^n} \overline{\text{true}}$$

Since s_P performs $3n$ beta reductions it needs $3n$ entries on the heap. The heap pointers then make the space consumption “explode” again. They are of size $\log n$ if binary numbers are used and n if unary numbers are used, resulting in an overall space consumption of $\Omega(n \log n)$ or $\Omega(n^2)$, both forming more than constant factor overhead. We call this problem *pointer explosion*, analogous to the discussed size explosion problem, and point out that both phenomena have already been identified and discussed by Slot and van Emde Boas [23] in their treatment of RAM simulations by Turing machines.

In our case of the weak call-by-value λ -calculus L , we have obtained two simulation strategies, each solving *one* of the problems: The substitution-based strategy works for space, but is insufficient for time on terms exhibiting size explosion (i.e. which have exponentially big intermediate terms). The heap-based strategy works for time, but is insufficient for space on terms exhibiting pointer explosion. The crucial observation is now that on terms exhibiting size explosion, i.e. reaching a term of size $\Omega(2^n)$ in n steps, pointer explosion is a non-issue: n pointers of size $\log n$ can easily be accommodated for in space $\mathcal{O}(2^n)$.

Since it is a-priori not decidable whether a term exhibits size explosion or pointer explosion, we interleave the execution of the two simulation strategies. We simulate the execution for every step number k repeatedly, and always try to run the substitution-based strategy first. If the size of intermediate terms becomes big enough to accommodate exploding pointers, we immediately abort and try the heap-based strategy for k steps instead. The heap-based strategy is thus guaranteed to not encounter the pointer explosion problem and the substitution-based strategy can not encounter the size explosion problem, because it is aborted beforehand. The details of this interleaving machine, which we consider as our main technical contribution, are given in Section 4.


Formalisation in Coq A technically demanding and error-prone part of our proof is analysing the exact complexity of the abstract machines involved (see Section 3).

Because our proofs rely on many simulation notions containing hard-to-check side conditions, we provide a formalisation of all results for the abstract machines, i.e. every theorem and definition needed for and including Section 3, in the proof assistant Coq [25].¹

A formalisation of the full results, including a formal verification of the Turing machines involved, is an ongoing and challenging project. We reported on the project previously in [11]; the current paper presents the finalised theoretical contributions.

2 Preliminaries

We adopt a notation closely related to type-theory, but the paper can be read with no background in type theory. All defined functions are always total. We use the type X_\perp to denote the type X enriched with a new element \perp . This allows us to view $X \rightarrow Y_\perp$ as the

¹ The code is hyperlinked with the PDF version of this document and can be accessed at <https://ps.uni-saarland.de/extras/wcbv-reasonable> or by clicking on the formalised statements and definitions, which are marked with a -symbol.

type of partial functions from X to Y . In the definition of such partial functions, left out cases are meant to default to \perp . Concerning lists $A, B : X^*$ over X , we use $[]$ for the empty list, write $x :: A$ to prepend an element to a list and write $[x_1, \dots, x_k]$ for a list built from the elements x_i . We write $A \# B : X^*$ for list concatenation, $|A| : \mathbb{N}$ for length and $A[n] : X_\perp$ for list lookup.

2.1 Call-by-value λ -calculus L

The call-by-value λ -calculus introduced by Plotkin [20] in his seminal paper is known to be a reasonable machine for time complexity [16]. In those works, abstractions and variables are treated as values, but β -reduction below binders is not allowed, i.e. reduction is *weak*. We use a deterministic version of the weak call-by-value λ -calculus we call L, originally introduced in [12]. We treat only abstractions as values but keep the weak behaviour of reduction. On closed terms, the number of steps to a normal form agrees with the number of steps needed in the version in [16]. We keep the definitions short and use the same notations as in [14], where more details can be found.

We define the syntax of the λ -calculus using a de Bruijn representation of terms [9]: $s, t, u, v : \text{Ter} ::= n \mid st \mid \lambda s$ where $n : \mathbb{N}$.

✦ **Definition 3.** We define a recursive function s_u^k providing a single-point, capturing substitution operation:

$$\begin{aligned} k_u^k &:= u & n_u^k &:= n & \text{if } n \neq k \\ (st)_u^k &:= (s_u^k)(t_u^k) & (\lambda s)_u^k &:= \lambda(s_u^{1+k}) \end{aligned}$$

We say that s is bounded by k if all free de Bruijn indices in s are lower than k . Consequently, s is a closed term iff it is bounded by 0.

✦ **Definition 4.** We define a deterministic inductive reduction relation $s \succ t$, which is weak, call-by-value and agreeing with the reduction relation in [20, 16] on closed terms:

$$\frac{}{(\lambda s)(\lambda t) \succ s_{\lambda t}^0} \quad \frac{s \succ s'}{st \succ s't} \quad \frac{t \succ t'}{(\lambda s)t \succ (\lambda s)t'}$$

Note that the only closed, irreducible terms are abstractions. We write $s \Downarrow_k^T t$ and $s \Downarrow_m^S t$ if $s \succ^k t$ for t being an abstraction and $m = \|s\|_S$ as defined in Definition 1.

The size of a term is defined with a unary encoding of indices in mind:

$$\|n\| := 1 + n \quad \|\lambda s\| := 1 + \|s\| \quad \|st\| := 1 + \|s\| + \|t\|$$

For a binary encoding, i.e. $\|n\| := 1 + \log_2 n$, we conjecture that the remainder of this paper can be adapted with no essential change.

2.2 Encoding Terms as Programs

Turing machines can not directly operate on tree-like data structures like L-terms. We encode terms as *programs* $P, Q, R : \text{Pro}$, which are lists of *commands*:

$$c : \text{Com} ::= \text{ret} \mid \text{var } n \mid \text{lam} \mid \text{app} \quad \text{with } n : \mathbb{N}$$

✦ **Definition 5.** The encoding function $\gamma : \text{Ter} \rightarrow \text{Pro}$ compiles terms to programs:

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

This encoding is similar to postfix notation, the additional command `lam` makes it easier to detect subprograms representing values when traversing the encoding.

✦ **Definition 6.** We write $P \gg s$, read as P represents s , to connect programs with values in L . This relation is defined with the single rule $\gamma t \gg \lambda t$.

To store the encoding of de Bruijn indices on tapes, we will use a unary encoding, motivating the following definition of the size of commands and programs:

$$\|\text{var } n\| := 1 + n \quad \|c\| := 1 \quad \text{if } c \text{ is no variable} \quad \|P\| := 1 + \sum_{c \in P} \|c\|$$

This size is compatible with term size, with factor 2 due to the two commands for abstractions:

✦ **Lemma 7.** $1 \leq \|s\| \leq \|\gamma s\| \leq 2\|s\| - 1$

The use of `lam` and `ret` to encode abstraction allows to define a function $\varphi P : (\text{Pro} \times \text{Pro})_{\perp}$ that extracts the body of an abstraction by matching `lam` with `ret` like parentheses. It uses an auxiliary function $\varphi_{k,Q} P$ that stores the number k of currently unmatched `lam` and the prefix Q already processed.

✦ **Definition 8.** $\varphi P := \varphi_{P,0} []$ with

$$\begin{aligned} \varphi_{0,Q}(\text{ret} :: P) &:= (Q, P) & \varphi_{1+k,Q}(\text{ret} :: P) &:= \varphi_{k,Q+[ret]} P \\ \varphi_{k,Q}(\text{lam} :: P) &:= \varphi_{1+k,Q+[lam]} P & \varphi_{k,Q}(c :: P) &:= \varphi_{k,Q+[c]} P \quad \text{if } c = \text{var } n \text{ or app} \end{aligned}$$

✦ **Lemma 9.** $\varphi(\gamma s + \text{ret} :: P) = (\gamma s, P)$

We define a substitution operation P_Q^k on programs, analogous to substitution on terms.

✦ **Definition 10** (Substitution on programs).

$$\begin{aligned} (\text{var } k :: P)_Q^k &:= Q :: P_Q^k & (\text{var } n :: P)_Q^k &:= \text{var } n :: P_Q^k \\ (\text{lam} :: P)_Q^k &:= \text{lam} :: P_Q^{S^k} & (\text{app} :: P)_Q^k &:= \text{app} :: P_Q^k \\ (\text{ret} :: P)_Q^0 &:= [\text{ret}] & (\text{ret} :: P)_Q^{S^k} &:= \text{ret} :: P_Q^k \\ []_Q^k &:= [] \end{aligned}$$

Term and program substitution are compatible:

✦ **Lemma 11.** $(\gamma s)_{\gamma t}^0 = \gamma(s_t^0)$

2.3 Closures and Heaps

To allow for structure sharing later, we introduce closures whose environments are stored in an explicitly modelled heap. Environments are stored as linked lists of closures on the heap, and closures $g : \text{HC} := \text{Pro} \times \text{HA}$ contain programs and pointers $a, b : \text{HA} := \mathbb{N}$ to the environment. We represent the heap $H : \text{Heap} := \text{HE}^*$ as a list of its cells $e : \text{HE} ::= (g, a)$ that store the head and the address of the tail. To interpret the linked list structure, we define a *lookup function* $H[a, n] : \text{HC}_{\perp}$ that returns the n -th entry of the list at address a , i.e. the value bound to the de Bruijn index n in the environment a :

$$H[a, n] := \text{if } n = 0 \text{ then } g \text{ else } H[b, n - 1] \quad \text{where } H[a] = (g, b)$$

The operation $\text{put He} : \text{Heap} \times \text{HA}$ puts a heap entry on the heap and returns the new heap and the address of the new element.

$$\text{put } H e := (H \# [e], 1 + |H|)$$

In our setting, we allocate at the end of the heap and have no need for garbage collection.

A closure (P, a) represents some term if the environment a contains bindings for all free variables in P . To make this more precise, we first define the unfolding of a term relative to some environment:

✦ **Definition 12** (Unfolding). *The unfolding $s\langle a, k \rangle \downarrow_H s'$ is inductively given by the rules*

$$\frac{n < k}{n\langle a, k \rangle \downarrow_H n} \qquad \frac{n \geq k \quad H[a, n - k] = (P, b) \quad P \gg s \quad s\langle b, 0 \rangle \downarrow_H s'}{n\langle a, k \rangle \downarrow_H s'}$$

$$\frac{s\langle a, 1 + k \rangle \downarrow_H s'}{\lambda s\langle a, k \rangle \downarrow_H \lambda s'} \qquad \frac{s\langle a, k \rangle \downarrow_H s' \quad t\langle a, k \rangle \downarrow_H t'}{st\langle a, k \rangle \downarrow_H s't'}$$

Intuitively, $s\langle a, k \rangle \downarrow_H s'$ holds if s' is obtained by recursively substituting all free variables in s by their values in the environment a . The index k is an artefact of the de Bruijn representation and denotes which variables in s are locally bound during the traversal of s .

The first rule states that bound variables are left unchanged. The second rule states that for free variables, the environment a binds n to some value s' that can be looked up in H . The third rule descends under an abstraction and thus one more variable is considered bound in s . The last rule descends under application.

✦ **Definition 13.** *The relation $g \gg_H s$, read as g represents s relative to H , is defined by the single rule*

$$\frac{P \gg t \quad t\langle a, 0 \rangle \downarrow_H s}{(P, a) \gg_H s}$$

3 Abstract Machines

In order to analyse the two mentioned strategies on a more semantic level than just as implementations on Turing machines we introduce two abstract machines implementing these strategies – based on substitutions and based on heaps. The machines are variants of the ones presented in [14]. Both machines will take $\mathcal{O}(\|s\|_T)$ abstract steps to evaluate a term s , but differ in the size of intermediate states and in the complexity of their respective implementations as Turing machines, which we construct in Section 4.

3.1 Substitution Machine

We define an abstract machine that uses substitution on programs. The implemented strategy is close to the small-step semantics for L. One important property is that the size of machine states during the machine run is linear in the size of the intermediate terms. Therefore, the substitution-based Turing machine will have constant factor overhead for space.

The abstract machine executes terms using two stacks of programs T and V called *task* and *value* stack. The task stack holds the parts of the program yet to be executed, and the value stack holds the already fully evaluated parts.

$$\begin{aligned}
& (\text{lam} :: P) :: T, V \succ P' ::_{\text{tc}} T, Q :: V && \text{if } \varphi P = (Q, P') \\
& (\text{app} :: P) :: T, Q :: R :: V \succ R_{\text{lam}::Q \text{+[ret]}}^0 :: P ::_{\text{tc}} T, V \\
& \text{where } P ::_{\text{tc}} T := \text{if } P = [] \text{ then } T \text{ else } P :: T
\end{aligned}$$

■ **Figure 1** Reduction rules of the substitution machine

$$\begin{aligned}
& (\text{var } n :: P, a) :: T, V, H \succ (P, a) :: T, g :: V, H && \text{if } H[a, n] = g \\
& (\text{lam} :: P, a) :: T, V, H \succ (P', a) :: T, (Q, a) :: V, H && \text{if } \varphi P = (Q, P') \\
& (\text{app} :: P, a) :: T, g :: (Q, b) :: V, H \succ (Q, b') :: (P, a) :: T, V, H' && \text{if } \text{put } H(g, b) = (H', b') \\
& ([], a) :: T, V, H \succ T, V, H
\end{aligned}$$

■ **Figure 2** Reduction rules of the heap machine

The semantics of the substitution machine is defined in Figure 1. The machine executes the first command of the topmost program of the task stack. In the *lambda rule*, the command `lam` marks the start of an abstraction. The sub-program corresponding to the body of the abstraction is moved to the value stack. In the *application rule*, the topmost values are applied to each other: The program R is instantiated with the argument Q to obtain a new task to be evaluated.

We need tail call optimisation $::_{\text{tc}}$ to guarantee that the size of the machine state is linear in the size of the represented term. Without it, the application rule could pile up return-tasks $P = []$ inside the task stack, invalidating Theorem 15. While we only need tail call optimisation for the application rule, adding it to the lambda rule as well streamlines proofs and allows us to avoid a rule to discard empty programs. The *initial state* τ_s for a term s is $\tau_s := ([\gamma s], [])$.

The machine evaluates L with a number of steps linear in the time-measure:

✦ **Theorem 14** (Substitution machine runtime). *If $s \Downarrow_k^T t$, then $\tau_s \succ^{3k+1} ([], [P])$ for some P with $P \gg t$.*

The size $\|T\|$ of T is defined to be just the sum of the sizes of the elements in T , and similar for V . The size of a state is defined by $\|(T, V)\| = \|T\| + \|V\|$. We write $\tau \succ_m^* \tau'$ for a sequence of machine reductions where the largest state has size m .

The maximal machine state size when evaluating s is asymptotically as large as $\|s\|_S$.

✦ **Theorem 15** (Substitution machine state size). *If $s \Downarrow_m^S t$, then $\tau_s \succ_{m'}^* ([], [P])$ for some P and m' with $P \gg t$ and $m \leq m' \leq 2m$.*

3.2 Heap Machine

This machine uses the heap described in Section 2.3 to enable sharing of environments. One important feature of this machine is that the size of intermediate states does only depend on the size of the initial term s and the number of machine steps, but not on $\|s\|_S$.

The machine is defined in Figure 2. Its task and value stacks contain closures. The *variable rule* loads the value bound to a variable to the value stack. The *lambda rule* copies a subprogram representing an abstraction to the value stack. The *application rule* calls the subprogram Q after adding the value g as argument to the environment of Q . The *return rule* drops finished tasks. The use of closures instead of programs allows an explicit variable rule instead of program-level substitution. The *initial state* σ_s for a closed term s is $([\gamma s, 0], [], [])$ as $s \langle 0, 0 \rangle \Downarrow_H s$ by Lemma 30.

The machine evaluates L with a number of steps linear in the time-measure:

✦ **Theorem 16** (Heap machine runtime). *If $s \Downarrow_k^T t$ and s is closed, then $\sigma_s \succ^{4k+2} ([], [g], H)$ for some g, H with $g \gg_H t$.*

We define the size of a closure as $\|(P, a)\| := \|P\| + a$ and the size of a heap entry to be $\|(g, a)\| := \|g\| + a$. The size of a state is the sum of the sizes of all elements in T , V and H .

The size of the k -th state starting from σ_s is a polynomial in k and $\|s\|$:

✦ **Theorem 17** (Heap machine state size). *If $\sigma_s \succ^k \sigma$, then $\|\sigma\| \leq (k + 1)(3k + 4\|s\|)$*

4 Simulating L with Turing Machines

We now sketch how to construct the Turing machine that simulates L with polynomially bounded overhead time and constant factor overhead in space. The considered Turing machines will operate on various kinds of data (e.g. natural numbers, programs, heap closures, heap entries, heaps, ...). For programs, we use a symbol for each of the four constructors and a fifth symbol to encode de Bruijn indices in unary. All other natural numbers will also be encoded in unary, unless explicitly stated. The encoding of the further structures on tapes is straightforward.

4.1 The Substitution-based Turing Machine Simulating L

We construct a Turing machine M_{subst} that executes the substitution-based strategy from Section 3.1 for k steps, where k is an input. The Turing machine takes an additional input m and aborts if the abstract machine would reach a state of size greater m .

► **Theorem 18.** *There is a Turing machine M_{subst} that, given two binary numbers k, m and a term s , halts in time $\mathcal{O}(k \cdot \text{poly}(\min(m, \|s\|_S)))$ and space $\mathcal{O}(\min(m, \|s\|_S) + \log m + \log k)$. Either the machine outputs a term t , then s has normal form t and $m \geq \|s\|_S$ and $k \geq 3 \cdot \|s\|_T + 1$. Or it halts in one of two other final states: Either a state named space bound reached, implying that $m \leq 2 \cdot \|s\|_S$ holds, or in a state named space bound not reached, implying that $k < 3 \cdot \|s\|_T + 1$ holds.*

Furthermore, by Theorem 15, the machine can only approximate the size of the 'current' term up to a factor of 2, which further complicates the theorem.

Proof. The Turing machine can be constructed by iterating the rules of the abstract substitution machine from Figure 1 on the initial state τ_s . The machine has to keep track of the size of the abstract machine state, even *during* the execution of the substitution: As soon as the size of the next state to be computed is known to exceed m , it aborts *before* consuming more than $\Theta(m)$ space. This is necessary because the result of a substitution P_Q^0 with $\|P\| + \|Q\| \in \mathcal{O}(m)$ could have quadratic size $\mathcal{O}(m^2)$, e.g. if P applies the variable 0 to itself m times and Q has size m as well. The function φP can be implemented via the tail-recursive $\varphi_{k,Q} P$, which takes space and time $\mathcal{O}(\|Q\| + k + \|P\|)$, as it just traverses P and accumulates the result. The argument k during the run is bound by $\|s\|_S$. Then the size of all intermediate states and the overall space consumption follow from Theorem 15. The existence of the result for large enough k follows in combination with Theorem 14. ◀

The precise specification of the machine is subtle: Intuitively, the machine state size is as large as the 'current' term, but we don't know if a state larger m is reached in the first k

steps. Therefore, we don't specify which of the last two cases occurs if *both* bounds on k and m are exceeded.

If s diverges, Theorem 18 states that M_{subst} can only halt in the two special final states (with $\|s\|_T = \infty$ for diverging terms s).

4.2 The Heap-based Turing Machine Simulating L

We construct a Turing machine executing the heap-based strategy from Sect. 3.2 for k steps:

► **Theorem 19.** *There is a Turing machine M_{heap} that, given a number k and a closed term s , halts in time $\mathcal{O}(\text{poly}(\|s\|, k))$ and space $\mathcal{O}(\|s\| \cdot \text{poly}(k))$. If s has a normal form t and $k \geq 4 \cdot \|s\|_T + 2$, it computes a heap H and a closure g such that $g \gg_H t$. Otherwise, it halts in a distinguished final state (denoting ‘failure’).*

Proof. The Turing machine can be constructed by iterating the rule of the abstract substitution machine on the initial state σ_s . We already argued on the runtime of φ for Theorem 18. And $H[a, n]$ can be computed by iterating over H for at most n times. So each abstract step $(T, V, H) \succ (T', V', H')$ can be implemented in time $\mathcal{O}(\text{poly}(\|(T, V, H)\|))$ and space $\mathcal{O}(\max(\|(T, V, H)\|, \|(T', V', H')\|))$. The space consumption of all involved operations in Figure 2 is bounded by their input or output. Using Theorem 17, the size of all intermediate (T, V, H) can be bound by k and $\|s\|$ to derive the claimed resource bounds. The successful computation of g and H for large enough k follows with Theorem 16. ◀

4.3 The Combined Turing Machine Simulating L

We now combine the machines from the last two sections to execute the heap-machine only if we know that its space consumption is bounded by the space measure of the simulated term:

► **Theorem 20.** *There is a Turing machine M_L that, given a closed term s that has a normal form t , computes a heap H and a closure g such that $g \gg_H t$ in time $\mathcal{O}(\text{poly}(\|s\|, \|s\|_T))$ and space $\mathcal{O}(\|s\|_S)$.*

Proof. Let p be the polynomial such that the machine from Theorem 19 runs in space $\mathcal{O}(\|s\| \cdot p(k))$. Then the combined machine executes the following algorithm:

1. Initialise $k := 0$ (in binary)
2. Compute $m := \|s\| \cdot p(k)$ (in binary)
3. Run M_{subst} on s , k and m .
 - If M_{subst} computes the normal form t , output $(\gamma t, 0)$ and an empty heap $[]$ and halt.
 - If M_{subst} halts with *space bound not reached*, set $k := k + 1$ and go to **2**.
 - If M_{subst} halts with *space bound reached*, continue at **4**.
4. Run M_{heap} on s and k .
 - If this computed a closure and a heap representing t , output that and halt.
 - Otherwise, set $k := k + 1$ and go to **2**.

First, we show that if this machine halts, its output is a closure-heap pair representing the normal form t of s : If the machine halts during **3**, the output is a representation of the normal form by Theorem 18 and Lemma 30. If it halts during **4**, it does so by Theorem 19.

Second, we analyse termination and the time complexity of this machine. As intermediate step, we analyse the run time for a fixed k . Step **2** takes time $\mathcal{O}(\text{poly}(\|s\|, k))$, and the size

of s can be computed from its encoding in straightforward fashion. Using Theorem 18, Step **3** takes time

$$\begin{aligned} \mathcal{O}(k \cdot \text{poly}(\min(m, \|s\|_S))) &\subseteq \mathcal{O}(k \cdot \text{poly}(m)) \\ &= \mathcal{O}(k \cdot \text{poly}(\|s\| \cdot p(k))) \\ &\subseteq \mathcal{O}(k \cdot \text{poly}(\|s\|, k)) && p \text{ is a polynomial} \\ &\subseteq \mathcal{O}(\text{poly}(\|s\|, k)) \end{aligned}$$

If Step **4** is executed, this takes time $\mathcal{O}(\text{poly}(\|s\|, k))$ by Theorem 19. This means for arbitrary k , one iteration of the described algorithm can be computed in time $\mathcal{O}(\text{poly}(\|s\|, k))$.

The algorithm will eventually halt: We consider $k = 4\|s\|_T + 2$, which is larger than the two values required in Theorem 18 and Theorem 19: By Theorem 18, the machine does halt during Step **3**, unless $m < \|s\|_S$. In the latter case, **4** is tried. Then, by Theorem 19, as k is large enough, we have that M_{heap} indeed halts with a closure-heap pair.

Summing up the run time of each iteration, we have that the machine terminates in time

$$\mathcal{O}\left(\sum_{k=0}^{4\|s\|_T+2} (\text{poly}(\|s\|, k))\right) \subseteq \mathcal{O}(\|s\|_T \cdot (\text{poly}(\|s\|, \|s\|_T))) \subseteq \mathcal{O}(\text{poly}(\|s\|, \|s\|_T))$$

Third, we analyse the space complexity of this machine. Again, we first analyse one iteration for a fixed k . Step **2** takes space $\mathcal{O}(\log(m))$, since we use binary numbers. By Theorem 18, Step **3** takes space $\mathcal{O}(\min(m, \|s\|_S) + \log m + \log k) \subseteq \mathcal{O}(\|s\|_S + \log m + \log k)$. If Step **4** is executed, then $m < \|s\|_S$. By Theorem 19, this step runs in space $\mathcal{O}(m) \subseteq \mathcal{O}(\|s\|_S)$. So, we can compute the space consumption of a single iteration as:

$$\begin{aligned} &\mathcal{O}(\log m + \log k + \|s\|_S) \\ &= \mathcal{O}(\log(\|s\| \cdot p(k)) + \log k + \|s\|_S) && \text{definition } m \\ &\subseteq \mathcal{O}(\log \|s\| + \log(p(k)) + \log k + \|s\|_S) \\ &= \mathcal{O}(\log(p(k)) + \log k + \|s\|_S) && \text{as } \|s\| \leq \|s\|_S \\ &\subseteq \mathcal{O}(\log k + \|s\|_S) && \log(p(k)) \in \mathcal{O}(\log k) \text{ as } p \text{ polynomial} \end{aligned}$$

Overall, we have that the whole machine runs in space (the last equation is by Theorem 21):

$$\mathcal{O}\left(\max_{0 \leq k \leq 4\|s\|_T+2} (\log k + \|s\|_S)\right) \subseteq \mathcal{O}(\log \|s\|_T + \|s\|_S) = \mathcal{O}(\|s\|_S) \quad \blacktriangleleft$$

Note that the machine only terminates for terminating terms, making this a full simulation also for diverging terms. For terms with $\|s\|_T \notin \mathcal{O}(\|s\|_S)$ it is crucial that the machine tracks the step number k in binary, because it would need $\Omega\|s\|_T$ space otherwise. This suffices due to the following theorem, which is proved in the appendix:

► **Theorem 21.** $\log \|s\|_T \in \mathcal{O}(\|s\|_S)$.

The simulation of L on Turing machines computes normal form as pair of closure and heap, as defined in Definition 13. It is possible to unfold this heap into a program:

► **Lemma 22.** *There is a machine M_{unf} that, given a heap H and a closure g that represent s , i.e. $g \gg_H s$, computes s (explicitly encoded as γs) in time $\mathcal{O}(\text{poly}(\|s\|, \|H\|, \|g\|))$ and space $\mathcal{O}(\|s\| \cdot (\|g\| + \|H\|))$.*

5 Simulating Turing Machines in L

The remaining direction of the proof of the strong invariance thesis requires us to prove that Turing machines can be simulated with L consuming only a constant overhead in space and a polynomial overhead in time with respect to our measures $\|\cdot\|_S$ and $\|\cdot\|_T$.

Accattoli and Dal Lago [5] show that counting head-reductions is an invariant time measure. In the associated technical report, they give a linear simulation of Turing machines in the deterministic λ -calculus, a fragment of the λ -calculus where all weak evaluation strategies coincide. Although they treat variables as values, reduction in L also coincides, because all considered terms are closed. The construction uses standard Scott encodings $\ulcorner x \urcorner$ for strings x and is explained in all detail in [15], spelling out all intermediate terms during simulation explicitly.

It turns out that this construction also only has a constant factor overhead in space w.r.t our measure $\|\cdot\|_S$. This can easily be verified by checking all intermediate terms spelled out in the proofs of [15]. One has to take care that a linear amount of steps (i.e. all steps annotated with $\mathcal{O}(\cdot)$ or $\Theta(\cdot)$ instead of constants) does not introduce a super-linear space overhead. This is the case, because all such sequences of steps only use substitutions where the substituted variable occurs at most once, effectively decreasing the term size. Note that since names in the simulation are all distinct, the translation to de Bruijn indices has no overhead. Thus the simulation is linear in time *and space*:

► **Theorem 23.** *Let $f : \Sigma^* \rightarrow \Sigma^*$ be a function that is computable by a Turing machine \mathcal{M} in time \mathcal{T} and in space \mathcal{S} . Then there exists an L-term $\overline{\mathcal{M}}$ such that for every $x \in \Sigma^*$ we have that*

1. $\overline{\mathcal{M}} \ulcorner x \urcorner \succ^* \ulcorner f(x) \urcorner$,
2. $\|\overline{\mathcal{M}} \ulcorner x \urcorner\|_S \in \mathcal{O}(|x| + \mathcal{S}(|x|))$, and
3. $\|\overline{\mathcal{M}} \ulcorner x \urcorner\|_T \in \mathcal{O}(|x| + \mathcal{T}(|x|))$.

Proof. Take $\overline{\mathcal{M}}$ as in Theorem 5.5. in [15]. ◀

6 The Weak Call-By-Value λ -Calculus is Reasonable

We explain how existing simulations of Turing machine in the λ -calculus already have polynomial time and constant factor space overhead in 5. With both the simulations, we are now able to show Theorem 2, that is, the invariance thesis for the weak call-by-value λ -calculus.

Proof. Let Σ be a finite alphabet such that $\{\mathbf{true}, \mathbf{false}\} \subseteq \Sigma$ and let $f : \Sigma^* \rightarrow \{\mathbf{true}, \mathbf{false}\}$ be a function. Furthermore, let $b = \max\{\|\ulcorner \mathbf{true} \urcorner\|, \|\ulcorner \mathbf{false} \urcorner\|\}$ and $\mathcal{T}, \mathcal{S} \in \Omega(n)$. Note that b is a constant only depending on the fixed alphabet Σ .

For the first direction, we assume that f is L-computable in time \mathcal{T} and space \mathcal{S} . By definition, there is hence a term s_f such that for all $x \in \Sigma^*$ we have that

$$s_f \ulcorner x \urcorner \succ^* \ulcorner f(x) \urcorner \quad \text{and} \quad \|s_f \ulcorner x \urcorner\|_T \leq \mathcal{T}(|x|) \quad \text{and} \quad \|s_f \ulcorner x \urcorner\|_S \leq \mathcal{S}(|x|).$$

We construct a Turing machine M_f as follows. On input x , M_f executes M_L on the (closed) term $s := s_f \ulcorner x \urcorner$, which computes a heap H and a closure g such that $g \ggg_H \ulcorner f(x) \urcorner$ in time $\mathcal{O}(\text{poly}(\|s\|, \|s\|_T))$ and space $\mathcal{O}(\|s\|_S)$, by Theorem 20 – note that s_f as well as M_L are hard-coded in M_f . We observe that

$$\|g\| + \|H\| \in \mathcal{O}(\text{poly}(\|s\|, \|s\|_T)) \quad \text{and} \quad \|g\| + \|H\| \in \mathcal{O}(\|s\|_S), \quad (1)$$

where the former holds as writing down g and H cannot take more time than the overall running time bound $\mathcal{O}(\text{poly}(\|s\|, \|s\|_{\mathcal{T}}))$ and the latter is due to the space bound $\mathcal{O}(\|s\|_{\mathcal{S}})$ of M_f . After that, M_f executes M_{unf} on H and g which yields $\lceil f(x) \rceil$ and finally, depending on whether $\lceil f(x) \rceil = \text{true}$ or $\lceil f(x) \rceil = \text{false}$, M_f outputs `true` or `false` accordingly. By Lemma 22, the final steps take time $\mathcal{O}(\text{poly}(b, \|H\|, \|g\|))$ and space $\mathcal{O}(b \cdot (\|g\| + \|H\|))$. Now the final time consumption is given by

$$\mathcal{O}(\text{poly}(\|s\|, \|s\|_{\mathcal{T}}) + \text{poly}(b, \|H\|, \|g\|)) \leq \mathcal{O}(\text{poly}(\|s\|, \|s\|_{\mathcal{T}}) + \text{poly}(b, \|s\|, \|s\|_{\mathcal{T}})) \quad (2)$$

$$\leq \mathcal{O}(\text{poly}(|x|, \mathcal{T}(|x|))) \quad (3)$$

$$\leq \mathcal{O}(\text{poly}(\mathcal{T}(|x|))), \quad (4)$$

where (2) is due to Equation (1), (3) holds as $\|s_f\|$ and b are constants and (4) follows from the fact that $\mathcal{T} \in \Omega(n)$. The overall space consumption is bounded by

$$\mathcal{O}(\|s\|_{\mathcal{S}} + b \cdot (\|H\| + \|g\|)) \leq \mathcal{O}((b+1) \cdot \|s\|_{\mathcal{S}}) \quad (5)$$

$$\leq \mathcal{O}(\mathcal{S}(|x|)), \quad (6)$$

where (5) is due to Equation (1) and (6) holds as b is a constant.

For the converse direction, we assume that f can be computed by a Turing machine \mathcal{M} in time \mathcal{T} and space \mathcal{S} . We invoke Theorem 23 to obtain a term $\overline{\mathcal{M}}$ which shows that f is L-computable in space $\mathcal{O}(|x| + \mathcal{S}(|x|))$, and time $\mathcal{O}(|x| + \mathcal{T}(|x|))$. We conclude the proof by observing that $\mathcal{O}(|x| + \mathcal{S}(|x|)) = \mathcal{O}(\mathcal{S}(|x|))$ and $\mathcal{O}(|x| + \mathcal{T}(|x|)) = \mathcal{O}(\mathcal{T}(|x|))$ as both, \mathcal{S} and \mathcal{T} are contained in $\Omega(n)$. \blacktriangleleft

7 Related and Future Work

We have already mentioned the recent long line of work by Accattoli, Dal Lago, Sacerdoti Coen, Guerrini and Martini (for an overview see [2]) analysing reasonable time measures and implementations of several λ -calculi.

Type systems for call-by-name and call-by-value λ -calculi can be used to logically characterise complexity classes (P [7], LOGSPACE [22], PSPACE [13]). Connecting these insights with our measures would make it even more feasible to use L as a formal basis for complexity theory, which we plan to do as future work, building on existing formalisations of computability theory [12].

There is recent work in investigating strategies to evaluate open terms, for instance open call-by-value, which is reasonable for time [4, 3], but the question for space is open. On the more applied side, there is work on time and space profiling based on lazy graph reduction [21] in Haskell. More recent work uses a graph-based cost-semantics used for space-profiling [24], based on earlier measures in [8]. Moreover, computation in sub-linear space with an external memory has been studied [17], which we do not cover in this paper.

And finally, the full λ -calculus can be translated into weak call-by-value e.g. using a CPS translation. The longstanding question whether the natural time and space measures for the λ -calculus are reasonable remains open. We want to investigate whether our results can contribute to an answer.

References

- 1 Beniamino Accattoli. The Complexity of Abstract Machines. In *Proceedings Third International Workshop on Rewriting Techniques for Program Transformations and Evalu-*

- ation, *WPTE@FSCD 2016, Porto, Portugal, 23rd June 2016.*, pages 1–15, 2016. doi:10.4204/EPTCS.235.1.
- 2 Beniamino Accattoli. (In)Efficiency and Reasonable Cost Models. *Electr. Notes Theor. Comput. Sci.*, 338:23–43, 2018. doi:10.1016/j.entcs.2018.10.003.
 - 3 Beniamino Accattoli and Claudio Sacerdoti Coen. On the Relative Usefulness of Fireballs. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 141–155, 2015. doi:10.1109/LICS.2015.23.
 - 4 Beniamino Accattoli and Giulio Guerrieri. Implementing Open Call-by-Value. In *Fundamentals of Software Engineering - 7th International Conference, FSEN 2017, Tehran, Iran, April 26-28, 2017, Revised Selected Papers*, pages 1–19, 2017. doi:10.1007/978-3-319-68972-2_1.
 - 5 Beniamino Accattoli and Ugo Dal Lago. On the Invariance of the Unitary Cost Model for Head Reduction. In *23rd International Conference on Rewriting Techniques and Applications (RTA'12), RTA 2012, May 28 - June 2, 2012, Nagoya, Japan*, pages 22–37, 2012. doi:10.4230/LIPIcs.RTA.2012.22.
 - 6 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.
 - 7 Andrea Asperti and Luca Roversi. Intuitionistic Light Affine Logic. *ACM Trans. Comput. Log.*, 3(1):137–175, 2002. doi:10.1145/504077.504081.
 - 8 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.
 - 9 Nicolaas Govert De Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. In *Indagationes Mathematicae (Proceedings)*, volume 75, pages 381–392. Elsevier, 1972.
 - 10 Nachum Dershowitz and Evgenia Falkovich-Derzhavetz. The invariance thesis. *Logical Methods in Computer Science (to appear)*, 2015. URL: <http://www.cs.tau.ac.il/~nachumd/papers/InvarianceThesis.pdf>.
 - 11 Yannick Forster, Fabian Kunze, and Marc Roth. The strong invariance thesis for a λ -calculus. *Workshop on Syntax and Semantics of Low-Level Languages (LOLA)*, 2017.
 - 12 Yannick Forster and Gert Smolka. Weak Call-by-Value Lambda Calculus as a Model of Computation in Coq. In *Interactive Theorem Proving - 8th International Conference, ITP 2017, Brasilia, Brazil, September 26-29, 2017, Proceedings*, pages 189–206, 2017. doi:10.1007/978-3-319-66107-0_13.
 - 13 Marco Gaboardi, Jean-Yves Marion, and Simona Ronchi Della Rocca. A logical account of PSPACE. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, pages 121–131, 2008. doi:10.1145/1328438.1328456.
 - 14 Fabian Kunze, Gert Smolka, and Yannick Forster. Formal Small-Step Verification of a Call-by-Value Lambda Calculus Machine. In *Programming Languages and Systems - 16th Asian Symposium, APLAS 2018, Wellington, New Zealand, December 2-6, 2018, Proceedings*, pages 264–283, 2018. doi:10.1007/978-3-030-02768-1_15.
 - 15 Ugo Dal Lago and Beniamino Accattoli. Encoding Turing Machines into the Deterministic Lambda-Calculus. *CoRR*, abs/1711.10078, 2017. URL: <http://arxiv.org/abs/1711.10078>, arXiv:1711.10078.
 - 16 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.
 - 17 Ugo Dal Lago and Ulrich Schöpp. Functional Programming in Sublinear Space. In *Programming Languages and Systems, 19th European Symposium on Programming, ESOP 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings*, pages 205–225, 2010. doi:10.1007/978-3-642-11957-6_12.

- 18 Julia L. Lawall and Harry G. Mairson. Optimality and Inefficiency: What Isn't a Cost Model of the Lambda Calculus? In *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming, ICFP 1996, Philadelphia, Pennsylvania, USA, May 24-26, 1996.*, pages 92–101, 1996. doi:10.1145/232627.232639.
- 19 Michael Norrish. Mechanised Computability Theory. In *Interactive Theorem Proving - Second International Conference, ITP 2011, Berg en Dal, The Netherlands, August 22-25, 2011. Proceedings*, pages 297–311, 2011. doi:10.1007/978-3-642-22863-6_22.
- 20 Gordon D. Plotkin. Call-by-Name, Call-by-Value and the lambda-Calculus. *Theor. Comput. Sci.*, 1(2):125–159, 1975. doi:10.1016/0304-3975(75)90017-1.
- 21 Patrick M. Sansom and Simon L. Peyton Jones. Time and Space Profiling for Non-Strict Higher-Order Functional Languages. In *Conference Record of POPL'95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23-25, 1995*, pages 355–366, 1995. doi:10.1145/199448.199531.
- 22 Ulrich Schöpp. Space-Efficient Computation by Interaction. In *Computer Science Logic, 20th International Workshop, CSL 2006, 15th Annual Conference of the EACSL, Szeged, Hungary, September 25-29, 2006, Proceedings*, pages 606–621, 2006. doi:10.1007/11874683_40.
- 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 Daniel Spoonhower, Guy E. Blelloch, Robert Harper, and Phillip B. Gibbons. Space profiling for parallel functional programs. *J. Funct. Program.*, 20(5-6):417–461, 2008. doi:10.1017/S0956796810000146.
- 25 The Coq Proof Assistant. <http://coq.inria.fr>, 2018.

A Big-step characterisation of reduction

While the characterisation of our time and space measure in terms of a normalising reduction $s_0 \succ \dots \succ s_k$ are intuitive, a big-step characterisation allows for easy, inductive analyses of the abstract machines evaluating L in Section 3.

✦ **Definition 24** (Time Measure).

$$\frac{}{\lambda s \Downarrow_0^T \lambda s} \quad \frac{s \Downarrow_{k_1}^T \lambda s' \quad t \Downarrow_{k_2}^T \lambda t' \quad s' \lambda t' \Downarrow_{k_3}^T u}{st \Downarrow_{k_1+k_2+1+k_3}^T u}$$

✦ **Definition 25** (Space Measures).

$$\frac{\lambda s \Downarrow_{|\lambda s|}^S \lambda s}{s \Downarrow_{m_1}^S \lambda s' \quad t \Downarrow_{m_2}^S \lambda t' \quad s' \lambda t' \Downarrow_{m_3}^S u \quad m = \max(1 + m_1 + |t|, 1 + |\lambda s'| + m_2, m_3)}{st \Downarrow_m^S u}$$

In the second rule, each of the three recursive assumptions could contain the largest subterm, so we take the maximum of each m_i , while accounting for the size of the remaining part of the term, e.g. during the reduction of s in st , the t and the application itself contribute to $1 + \|t\|$ additional size by definition of the term siz .

The following lemmas allow us to use the two characterisations interchangeably:

✦ **Lemma 26.** $s \Downarrow_k^T t$ iff $s \succ^k t$ and t is an abstraction.

Note that especially if $\|s\|_T = k$, then $s \Downarrow_k^T t$ for some t .

We write $s \succ_m^* t$ if s reduces to t where the largest intermediate term has size m .

✦ **Lemma 27.** $s \Downarrow_m^S t$ iff $s \succ_m^* t$ and t is an abstraction.

Note that especially, if $\|s\|_S = m$, then $s \Downarrow_m^S t$ for some t .

B Technical Definitions and Lemmas

✦ **Lemma 28.** γ is injective.

Adding a value t' to an environment a results in substitution in the unfolded term:

✦ **Lemma 29.** If $H[a'] = (g, a)$ with $g \gg_H t'$ and $s\langle a, 1 \rangle \Downarrow_H s'$, then $s\langle a', 0 \rangle \Downarrow_H s' \lambda t'$.

Unfolding only changes de Bruijn indices starting at k :

✦ **Lemma 30.** If s is bounded by k , then $s\langle a, k \rangle \Downarrow_H s$.

Proof. Induction on $s < k$. ◀

In particular, closed terms are invariant under unfolding.

The unfolding relation only holds if all de Bruijn indices up to k are bound in a .

✦ **Lemma 31.** If $s\langle a, k \rangle \Downarrow_H s'$, then $s' < k$

Proof. Induction on $s\langle a, k \rangle \Downarrow_H s'$. ◀

In particular for $k = 0$, unfolding results in closed terms.

A heap is extended by another heap if the latter contains a superset of the entries:

✦ **Definition 32.** $H \subseteq H' := \forall a, H[a] \neq \perp \rightarrow H[a] = H'[a]$

✦ **Lemma 33.** *Heap extension $H \subseteq H'$ is transitive and reflexive.*

Heap extension does not change the result of certain operations:

✦ **Lemma 34.** *Assume $H \subseteq H'$*

1. *If $H[a, n] \neq \perp$, then $H[a, n] = H'[a, n]$.*
2. *If $s\langle a, k \rangle \downarrow_H s'$, then $s\langle a, k \rangle \downarrow_{H'} s'$.*
3. *If $g \gg_H s$, then $g \gg_{H'} s$.*

Proof. The first claim follow by induction on n . The second claim follows by induction on $s\langle a, k \rangle \downarrow_H s'$. The only interesting case is the one where $s = n \geq k$, which requires the first claim. The third claim follows from the second by definition of \gg_H . ◀

✦ **Lemma 35.** *Assume $\sigma_s \succ^k (T, V, H) = \sigma$ for some term s .*

1. $|T| + |V| \leq k + 1$
2. $|H| \leq k$
3. $\|P\| \leq \|s\|$ and $a \leq |H|$ for all $P/a \in T + V$
4. $\|P\| \leq \|s\|$ and $a \leq |H|$ and $b \leq |H|$ for all $((P, a), b) \in H$

C Proofs

Proof of Lemma 7:

Proof. Induction on s . ◀

Proof of Lemma 9:

Proof. The generalisation $\varphi_{k,Q}(\gamma s + P) = \varphi_{k,Q + \gamma s} P$ follows by induction on s . ◀

Proof of Lemma 11:

Proof. The generalisation $(\gamma s + P)_{\gamma t}^k = \gamma(s_t^k) + P_{\gamma t}^k$ holds by induction on s . ◀

Proof of Theorem 14:

Proof. We show the generalisation, *if $s \downarrow_k^T t$, then for all Q, T, V we have $((\gamma s + Q) :: T, V) \succ^{3k+1} (Q ::_{tc} T, P :: V)$ for some P with $P \gg t$, from which the claim follows for $Q = R = V = []$.*

Proof by induction on $s \downarrow_k^T t$ as defined in Definition 24.

In the case $\lambda s \downarrow_0^T \lambda s$, we have

$$\begin{aligned} ((\gamma(\lambda s) + Q) :: T, V) &= ((\text{lam} :: \gamma s :: \text{ret} + Q) :: T, V) \\ &\succ (Q ::_{tc} T, \gamma s :: V) \end{aligned}$$

Lemma 9

and $\gamma s \gg s$ holds by definition.

In the case $st \Downarrow_{k_1+k_2+1+k_3}^T u$ with all names as in Definition 24, we have

$$\begin{aligned}
((\gamma(st) + Q) :: T, V) &= ((\gamma s + \gamma t + \text{app} :: Q) :: T, V) \\
&\succ^{3k_1+1} ((\gamma t + \text{app} :: Q) ::_{tc} T, \gamma s' :: V) && \text{IH for } s \Downarrow_{k_1}^T \lambda s' \\
&\succ^{3k_2+1} ((\text{app} :: Q) ::_{tc} T, \gamma t' :: \gamma s' :: V) && \text{IH for } t \Downarrow_{k_2}^T \lambda t' \\
&\succ (((\gamma s')^0_{\gamma(\lambda t')} :: Q) ::_{tc} T, \gamma t' :: \gamma s' :: V) \\
&= (\gamma(s'^0_{\lambda t'}) :: Q) ::_{tc} T, V && \text{Lemma 11} \\
&\succ^{3k_3+1} (Q) ::_{tc} T, \gamma u :: V && \text{IH for } s'^0_{\lambda t'} \Downarrow_{k_3}^T \lambda u
\end{aligned}$$

Note that $::_{tc}$ is $::$ in the first two reductions.

The claim follows as $3(k_1 + k_2 + 1 + k_3) + 1 = (3k_1 + 1) + (3k_2 + 1) + 1 + (3k_3 + 1)$ and $\gamma u \gg \lambda u$ by definition. \blacktriangleleft

Proof of Theorem 15:

Proof. We show a generalisation, *if $s \Downarrow_m^S t$, then for all Q, T, V we have $((\gamma s + Q) :: T, V) \succ_{m'}^* (Q) ::_{tc} T, P :: V$ for some P, m' with $P \gg t$ and $m + \|P\| + \|V\| \leq m' \leq 2m + \|P\| + \|V\|$, from which the claim follows with $Q = T = V = []$.*

Proof by induction on $s \Downarrow_m^S t$ as defined in Definition 25. By definition of \Downarrow^S , this proof is very similar to the one for Theorem 14. The only difference is the needed equalities between the various space-measures m_i . Those are proven by tedious, but straightforward computations when using the facts that $\|P\| + \|T\| \leq \|P\| + \|T\| + 1$ and Lemma 11 and Lemma 7 and the fact that $s \Downarrow_m^S t$ implies $\|s\| \leq m \geq \|t\|$. \blacktriangleleft

Proof of Theorem 16:

Proof. We show a generalisation, *If $s \Downarrow_k^T t$ and $s_0 \langle a, 0 \rangle \Downarrow_H s$, then there are g and H' with $g \gg_{H'} t$ such that $((\gamma s_0 + P, a) :: T, V, H) \succ^{4k+1} ((P, a) :: T, g :: V, H')$ for any P, T, V , and $H \subseteq H'$. Here $H \subseteq H'$ is meant as in Definition 32. The original claim follows with $P = T = V = H = []$, Lemma 30 and the reduction rule for empty tasks.*

Proof by induction on $s \Downarrow_m^T t$ as in Definition 24. In the case of $\lambda s \Downarrow_0^T \lambda s$, a case distinction on $s_0 \langle a, 0 \rangle \Downarrow_H \lambda s$ yields two cases: s_0 is either a variable with a value bound in a , or s_0 is an abstraction.

In the case $s_0 = n$, we obtain Q, b, s_1 such that $H[a, n] = (Q, b)$ with $Q \gg s_1$ and $s_1 \langle b, 0 \rangle \Downarrow_H \lambda s$. The claim holds as $(Q, b) \gg_H \lambda s$ and $((\gamma n + P, a) :: T, V, H) = ((\text{var } n :: P, a) :: T, V, H) \succ ((P, a) :: T, (Q, b) :: V, H)$.

In the case $s_0 = \lambda s_1$, we have that $s_1 \langle a, 1 \rangle \Downarrow_H s$. The claim holds as $(\gamma s_1, a) \gg_H \lambda s$ and $((\gamma(\lambda s_1) + P, a) :: T, V, H) = ((\text{lam} :: \gamma s_1 + \text{ret} :: P, a) :: T, V, H) \succ ((P, a) :: T, (\gamma s_1, a) :: V, H)$.

In the other case of the induction, $st \Downarrow_{k_1+k_2+1+k_3}^T u$, we have $s \Downarrow_{k_1}^T \lambda s'$ and $t \Downarrow_{k_2}^T \lambda t'$ and $s'^0_{\lambda t'} \Downarrow_{k_3}^T u$ and an inductive hypothesis for each of those. We also have $s_0 \langle a, 0 \rangle \Downarrow_H st$. Now $s_0 = s_1 t_1$ must be an application. Note that even in the second rule of Definition 12, the definition of \gg on programs implies that the unfolded term would be an abstraction.

So we have $s_1 \langle a, 0 \rangle \Downarrow_H s$ and $t_1 \langle a, 0 \rangle \Downarrow_H t$ for some s_1, t_1 . We now construct the reduction of the machine using the inductive hypothesis. We will explain where the new objects in the

following reduction come from in the next paragraph.

$$(\gamma(s_1 t_1) \# P, a) :: T, V, H) = (\gamma s_1 \# \gamma t_1 \# \text{app} :: P, a) :: T, V, H) \quad (7)$$

$$\succ^{4k_1+1} ((\gamma t_1 \# \text{app} :: P, a) :: T, (\gamma s_2, a_2) :: V, H_1) \quad \text{IH} \quad (8)$$

$$\succ^{4k_2+1} ((\text{app} :: P, a) :: T, g_t :: (\gamma s_2, a_2) :: V, H_2) \quad \text{IH} \quad (9)$$

$$\succ ((\gamma s_2, a'_2) :: (P, a) :: T, V, H'_2) \quad (10)$$

$$\succ^{4k_3+1} ((\square, a'_2) :: (P, a) :: T, g_u :: V, H_3) \quad \text{IH} \quad (11)$$

$$\succ ((P, a) :: T, g_u :: V, H_3) \quad (12)$$

In this reduction, the inductive hypothesis for s_1 in (8) yields s_2, a_2 and H_1 such that $H \subseteq H_1$ and $s_2 \langle a_2, 1 \rangle \downarrow_{H_1} s'$. The inductive hypothesis on t_1 in (9) yields g_t and H_2 such that $H_1 \subseteq H_2$ and $g_t \gg_{H_2} \lambda t'$. In the step for beta reduction, (10), we have $(H'_2, a'_2) = \text{put}H_2(g_t, a_2)$ and $H_2 \subseteq H'_2$. With Lemma 29, this implies $s_2 \langle a'_2, 0 \rangle \downarrow_{H'_2} s''_{\lambda t'}$. This now allows the use of the third inductive hypothesis in (11), obtaining g_u and H_3 with $g_u \gg u$ and $H'_2 \subseteq H_3$. Note that we use Lemma 33 to transfer several properties along the changing heaps. Now the claim holds for g_u and H_3 . \blacktriangleleft

Proof of Theorem 17:

Proof. Follows from Lemma 35 \blacktriangleleft

Proof of Theorem 21:

Proof. The main insight is that for any given size, there are only exponentially many terms smaller than that size. As reduction is deterministic, a terminating term s can not contain the same intermediate term twice. This bounds $\|s\|_T$ by the number of terms with size smaller than $\|s\|_S$, i.e. $\|s\|_T \leq c^{\|s\|_S}$ for a constant c .

Now, we show that the number of terms smaller than a certain size m is an exponential. We use the encoding γ to allow us to count linear strings (programs) instead of trees (terms):

$$\begin{aligned} & \#\{t \mid \|t\| \leq m\} \\ &= \#\{t \mid 2 \cdot \|t\| \leq 2 \cdot m\} \\ &\leq \#\{t \mid \|\gamma t\| \leq 2 \cdot m\} && \text{Lemma 7} \\ &= \#\{\gamma t \mid \|\gamma t\| \leq 2 \cdot m\} && \text{Lemma 28} \\ &\leq \#\{P \mid \|P\| \leq 2 \cdot m\} \\ &\leq 5^{2 \cdot m} \end{aligned}$$

In the last step, we use that $\#\{P \mid \|P\| \leq n\} \leq 5^{n-1}$ for all $n > 0$ by induction on n , where the intuition behind the 5 is that there are four different symbols with which a program can start, and that variables require a fifth symbol to encode the index in unary. Thus the claim holds for $c = 5^2$. \blacktriangleleft

Proof of Lemma 22:

Proof. We first consider a partial function f_{HPak} that computes the unfolding, but on programs instead of terms:

$$\begin{aligned}
f_H[]ak &:= [] \\
f_H(\text{app} :: P)ak &:= \text{app} :: f_HPak \\
f_H(\text{ret} :: P)a(1+k) &:= \text{ret} :: f_HPak \\
f_H(\text{lam} :: P)ak &:= \text{lam} :: f_HPa(1+k) \\
f_H(\text{var } n :: P)ak &:= \text{lam} :: f_HQb1 + \text{ret} :: f_HPak \quad \text{if } n \geq k \text{ and } H[a, n-k] = (Q, b) \\
f_H(\text{var } n :: P)ak &:= \text{var } n :: f_HPak \quad \text{if } n < k
\end{aligned}$$

For this set of equations, we can show *if* $k \langle s, a \rangle \downarrow_H s'$, *then* $f_H(\gamma s + Q)ak = \gamma s' + f_HQak$ by induction on s' . With $Q = []$, this means that $(P, a) \gg_H \lambda s'$ implies $f_HP a 1 = \gamma s'$. So f indeed computes the unfolding on programs.

Implementing f in Turing machines, we first note that during execution, all considered k , a and P are bound is bound by $\|g\| + \|H\|$, as all addresses come from g or H and k can not be larger than the largest program in g or H . In the equation using $H[\cdot, \cdot]$, an additional explicit stack is needed to remember P for after the recursive call on Q . This stack is bound in length by $\mathcal{O}(\|s\|)$, as every recursive call computes at least one symbol of the result. This means that the algorithm runs in space $\mathcal{O}(\|s\| \cdot (\|g\| + \|H\|))$.

Furthermore, each equation produces a symbol of the result, so the total number of calls on f is bound by $\|s\|$. Every equation, except the one where $H[\cdot, \cdot]$ occurs, performs a constant number of operations. This other equation needs to traverse H at most n times before recurring, where n is the largest de Bruijn index occurring. In total, this means that the algorithm runs in time $\mathcal{O}(\text{poly}(\|s\|, \|g\|, \|H\|))$. \blacktriangleleft

Proof of Lemma 26:

Proof. For the direction assuming $s \Downarrow_k^T t$, the claim follows by induction on \Downarrow^T using two compatibility lemmas of \succ^k with term-level application, the first being that $s \succ^k s'$ implies $st \succ^k s't$, and the second that $t \succ^k t'$ implies $(\lambda s)t \succ^k (\lambda s)t'$.

For the other direction, we first show

\triangleright **Claim 1.** If $s \succ s'$ and $s' \Downarrow_k^T t$, then $s \Downarrow_{1+k}^T t$.

This claim follows by induction in $s \succ s'$.

Now, assuming $s \succ^k \lambda t$, we can show $s \Downarrow_k^T t$ by induction on k using Claim 1 in the case where $k > 0$. \blacktriangleleft

Proof of Lemma 27:

Proof. For the direction assuming $s \Downarrow_m^S t$, the claim follows by induction on \Downarrow^S . In the inductive case, two compatibility lemmas of \succ_m^* with term-level application are helpful: the first being that $s \succ_m^* s'$ implies $st \succ_{1+m+\|t\|}^* s't$, and the second that $t \succ_m^* t'$ implies $(\lambda s)t \succ_{1+m+\|\lambda s\|}^* (\lambda s)t'$. Furthermore, the fact that $t \Downarrow_{m_2}^S \lambda t'$ implies $\|\lambda t'\| \leq m_2$ is needed.

For the other direction, we first show

\triangleright **Claim 2.** If $s \succ s'$ and $s' \Downarrow_m^S t$, then $s \Downarrow_{\max(\|s\|, m)}^S t$.

This claim follows by induction in $s \succ s'$. The equalities between expressions involving \max are tedious to check, but follow only using the inductive hypothesis, the definition of \Downarrow^S , the definition of the size of terms and that $s \Downarrow_m^S t$ implies $\|s\| \leq m \geq \|t\|$.

Now, assuming $s \succ_m^k \lambda t$, we can show $s \Downarrow_m^S \lambda t$ by induction on k . In the base case, $\lambda t \succ_m^0 \lambda t$ implies $m = \|\lambda t\|$ implies $\lambda t \Downarrow_m^S \lambda t$.

In the inductive case, $s \succ_m^{1+k} \lambda t$ implies a decomposition $s \succ s' \succ_{m'}^k \lambda t$ for some s', m' with $m = \max(\|s\|, m')$. Then the inductive hypothesis for k is $s' \Downarrow_{m'}^S \lambda t$, which together with Claim 2 implies $s \Downarrow_m^S \lambda t$. \blacktriangleleft

Proof of Lemma 28:

Proof. We define an inverse $\delta : \mathbb{N} \rightarrow \text{Pro} \rightarrow \text{Ter}^* \rightarrow \text{Ter}^*_\perp$ of γ :

$$\begin{aligned} \delta k(\text{var } n :: P)A &:= \delta kP(n :: P) & \delta k(\text{app} :: P)(t :: s :: A) &:= \delta kP(st :: A) \\ \delta k(\text{lam} :: P)A &:= \delta(1+k)PA & \delta(1+k)(\text{ret} :: P)(s :: A) &:= \delta kP(\lambda s. :: A) \\ \delta k[]A &:= A \end{aligned}$$

Now $\delta k(\gamma s \# P)A = \delta kP(s :: A)$ holds by induction on s . \blacktriangleleft

Proof of Lemma 29:

Proof. Let $H[a'] = (g, a)$ and $g \gg_H t'$. We show a generalisation: *If $s\langle a, 1+k \rangle \Downarrow_H s'$, then $s\langle a', k \rangle \Downarrow_H s'_{\text{t}'}$, by induction on $s\langle a, 1+k \rangle \Downarrow_H s'$.*

In the case $s = n < 1+k$ and $s' = n$, there are two subcases: Assuming $n < k$, $n\langle a', k \rangle \Downarrow_H n = s'_{\text{t}'}$ holds by definition. Otherwise, we have $n = k$. Since $H[a'] = (g, a)$, we have $H[a', n-k] = H[a', 0] = g$. With P, b such that $g = (P, b)$, we have $n\langle a', k \rangle \Downarrow_H n_{\text{t}'}^k = t'$ by the second rule since $(P, b) \gg_H t'$ implies $(P \gg t)$ and $(t, b) \gg_H t'$ for some t .

In the case $s = n \geq 1+k$, we have $H[a, n-(1+k)] = (P, b)$ and $P \gg u$ with $u\langle b, 0 \rangle \Downarrow_H s'$ for some P, b . As $H[a'] = (g, a)$, we have $H[a', n-k] = H[a, n-(1+k)] = (P, b)$. Therefore $n\langle a, k \rangle \Downarrow_H s'_{\text{t}'}^k = s'$ by the second rule, where the equality holds as s' is closed by Lemma 31.

In the other cases, i.e. application and abstraction, the claim follows by the inductive hypothesis and the definition of $\langle \cdot, k \rangle \Downarrow_H \cdot$. \blacktriangleleft

Proof of Lemma 33:

Proof. Transitivity and reflexivity follow from the same properties for equality. \blacktriangleleft

Proof of Lemma 35:

Proof. All claims follow by induction on k . The third claim uses that φP always returns a sublist of P . \blacktriangleleft

D RAM machines can consume more space than time

Turing machines can not consume more space than time, since it costs a time unit to allocate a new space unit. For RAM machines, this is different, as analysed by Slot and van Emde Boas [23].

The *time consumption* of a RAM computation is the number of steps; we denote it by \mathcal{T} . The *space consumption* is given by

$$\mathcal{S}_b = \sum_{i=0}^m \text{size}_b(i, \max(i)),$$

where m is the index of the highest address for which a register was accessed and $\max(i)$ is the maximal content of register $\mathcal{R}[i]$ during the computation. Furthermore,

$$\text{size}_b(i, x) := \begin{cases} 0 & \text{if } \mathcal{R}[i] \text{ is unused} \\ \log(x) + \log(i) & \text{otherwise} \end{cases}$$

Now, intuitively, \mathcal{S}_b is the sum of the maximum sizes of contents of used registers and the sizes of the addresses required to access those registers. It is known that, using \mathcal{S}_b as space measure and \mathcal{T} as time measure, RAM machines and Turing machines can simulate each other with a constant overhead in space and a polynomial overhead in time (see e.g. Section 1 in [23]). Next we consider the following RAM program P :

Input: x
 $a \leftarrow 1$;
for $i = 1$ **to** $|x|$ **do** (where $|x|$ is the length of x in binary)
 $a \leftarrow a + a$;
 $\mathcal{R}[a] \leftarrow 1$;
od;
Output: 1

Observe that P is similar to the size-exploding term s we have seen in the introduction. The time consumption of P is given by the following function; recall that the size of an input is given by its length.

$$n \mapsto 2n + 2 \in \mathcal{O}(n).$$

However, the space consumption is given by

$$n \mapsto \sum_{i=0}^{2^n} \text{size}_b(i, \max(i)) = \sum_{i=0}^n \text{size}_b(2^i, \max(2^i)) = \sum_{i=0}^n \log(2^i) + 1 \in \Omega(n^2).$$

Thus there are RAM machines that consume asymptotically more space than time, despite being a sequential model.