

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

Yannick Forster, Fabian Kunze, Marc Roth

POPL 2020
January 22

SAARLAND
UNIVERSITY



COMPUTER SCIENCE



SIC Saarland Informatics
Campus

The strong invariance thesis

Slot and Van Emde Boas (1984):

Reasonable machines
can simulate each other with a polynomial overhead in time and a
constant-factor overhead in space

The strong invariance thesis

Slot and Van Emde Boas (1984):

Turing machines and RAM machines
can simulate each other with a polynomial overhead in time and a
constant-factor overhead in space

The strong invariance thesis

This paper:

Turing machines and the weak call-by-value λ -calculus
can simulate each other with a polynomial overhead in time and a
constant-factor overhead in space

The strong invariance thesis

This paper:

Turing machines and the weak call-by-value λ -calculus
can simulate each other with a polynomial overhead in time and a
constant-factor overhead in space
with respect to the natural measures,
i.e. β -steps and size of largest term in the computation

The strong invariance thesis

This paper:

Turing machines and the weak call-by-value λ -calculus
can simulate each other with a polynomial overhead in time and a
constant-factor overhead in space
with respect to the natural measures,
i.e. β -steps and size of largest term in the computation

Standard complexity classes like P, NP, PSPACE, EXP
can be defined in terms of the λ -calculus¹

¹our result does not cover sublinear space or time

The Weak Call-by-Value λ -calculus L (Forster & Smolka, 2017)

$s, t, u, v : \text{Ter} ::= n \mid st \mid \lambda s$ where $n : \mathbb{N}$.

- Weak evaluation: don't reduce below abstractions
- Call-by-value evaluation: reduce arguments first
- Deterministic

Define (tree-)size of terms:

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

For $s = s_0 \succ s_1 \succ \dots \succ s_k = \lambda x.u$ define

- Time measure of s : k
- Space measure of s : $\max_{i=0}^k \|s_i\|$

The Weak Call-by-Value λ -calculus L (Forster & Smolka, 2017)

$s, t, u, v : \text{Ter} ::= n \mid st \mid \lambda s$ where $n : \mathbb{N}$.

- Weak evaluation: don't reduce below abstractions
- Call-by-value evaluation: reduce arguments first
- Deterministic

Define (tree-)size of terms:

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

For $s = s_0 \succ s_1 \succ \dots \succ s_k = \lambda x.u$ define

- Time measure of s : k
- Space measure of s : $\max_{i=0}^k \|s_i\|$

Measures are mapping terms to numbers!

The easy direction

Theorem

The weak call-by-value λ -calculus can simulate Turing machines with a polynomial overhead in time and a constant-factor overhead in space.

Follows by Accattoli / Dal Lago '17,
result for time mechanised in Forster / Kunze '19.

What's known for the harder direction?

- 1996: Lawall and Mairson: “total ink used” and “maximum ink used” are reasonable measures for the full λ -calculus
- 2008: Dal Lago and Martini: β -steps and accounting for the size of β -redexes are a reasonable time measure for the weak call-by-value λ -calculus
- 2016: Accattoli and Dal Lago: (leftmost-outermost) β -steps are a reasonable time measure for the full λ -calculus

Why isn't that enough?

To allow for e.g. mechanised complexity theoretic results we want

- **compositional** measures
- *for a compositional model*
- *for both time and space*

Why isn't that enough?

To allow for e.g. mechanised complexity theoretic results we want

- **compositional** measures

“total ink used” for time by Lawall/Mairson is not enough

- *for a* **compositional model**

- *for* **both** *time and space*

Why isn't that enough?

To allow for e.g. mechanised complexity theoretic results we want

- **compositional** measures

“total ink used” for time by Lawall/Mairson is not enough

- *for a* **compositional model**

Turing machines are not enough

- *for* **both** *time and space*

Why isn't that enough?

To allow for e.g. mechanised complexity theoretic results we want

- **compositional** measures

“total ink used” for time by Lawall/Mairson is not enough

- *for a* **compositional model**

Turing machines are not enough

- *for* **both** *time and space*

time complexity results by Accattoli/Dal Lago/Martini are not enough

Wait a moment!

There's a λ -term s_E with

$$\forall n : \mathbb{N}. s_E \bar{n} \succ^* \lambda xy.x$$

in $\mathcal{O}(n)$ steps but with $\mathcal{O}(2^n)$ space.

Wait a moment!

There's a λ -term s_E with

$$\forall n : \mathbb{N}. s_E \bar{n} \succ^* \lambda xy.x$$

in $\mathcal{O}(n)$ steps but with $\mathcal{O}(2^n)$ space.

But $P \subseteq PSPACE$?!

$P \subseteq PSPACE$

Let M be a machine computing a function $\mathbb{N} \rightarrow \mathbb{B}$ and $x : \mathbb{N}$

Theorem (for TMs, free)

$$\mathcal{S}(M, x) \leq \mathcal{T}(M, x)$$

$P \subseteq PSPACE$

Let M be a machine computing a function $\mathbb{N} \rightarrow \mathbb{B}$ and $x : \mathbb{N}$

Theorem (for TMs, free)

$$\mathcal{S}(M, x) \leq \mathcal{T}(M, x)$$

Theorem (for RAM machines, Slot & van Emde Boas, hard)

$$\mathcal{S}(M, x) \leq p(\mathcal{T}(M, x)) \quad \text{for a polynomial } p$$

$P \subseteq PSPACE$

Let M be a machine computing a function $\mathbb{N} \rightarrow \mathbb{B}$ and $x : \mathbb{N}$

Theorem (for TMs, free)

$$\mathcal{S}(M, x) \leq \mathcal{T}(M, x)$$

Theorem (for RAM machines, Slot & van Emde Boas, hard)

$$\mathcal{S}(M, x) \leq p(\mathcal{T}(M, x)) \quad \text{for a polynomial } p$$

Theorem (for L, follows from our result)

$$\begin{aligned} \exists M'. \quad & M' \text{ is ext. equiv. to } M \text{ and} \\ \mathcal{S}(M', x) & \leq p_1(\mathcal{T}(M', x)) \leq p_2(\mathcal{T}(M, x)) \\ & \text{for polynomials } p_1, p_2 \end{aligned}$$

Terms can exhibit size explosion,
but decision functions can be optimised
to not explode.

Theorem (Accattoli & Dal Lago, 2016)

There is an algorithm which takes as input a λ -term t and which, in time polynomial in the number of left-most outermost β -steps of t and the size of t

outputs an LSC term u such that the unfolding of u is the normal form of t .

Theorem (This paper, 2020)

*There is an algorithm which takes as input a **closed L-term** t and which,*

Theorem (Accattoli & Dal Lago, 2016)

There is an algorithm which takes as input a λ -term t and which, in time polynomial in the number of left-most outermost β -steps of t and the size of t

outputs an LSC term u such that the unfolding of u is the normal form of t .

Theorem (This paper, 2020)

*There is an algorithm which takes as input a **closed L-term** t and which, in time polynomial in the number of β -steps of t and the size of t*

Theorem (Accattoli & Dal Lago, 2016)

There is an algorithm which takes as input a λ -term t and which, in time polynomial in the number of left-most outermost β -steps of t and the size of t

outputs an LSC term u such that the unfolding of u is the normal form of t .

Theorem (This paper, 2020)

*There is an algorithm which takes as input a **closed L-term** t and which, in time polynomial in the number of β -steps of t and the size of t
and space linear in the size of the largest term in the reduction*

Theorem (Accattoli & Dal Lago, 2016)

There is an algorithm which takes as input a λ -term t and which, in time polynomial in the number of left-most outermost β -steps of t and the size of t

outputs an LSC term u such that the unfolding of u is the normal form of t .

Theorem (This paper, 2020)

*There is an algorithm which takes as input a **closed L-term** t and which, in time polynomial in the number of β -steps of t and the size of t*

and space linear in the size of the largest term in the reduction

*outputs a **heap containing a term** u such that the unfolding of u is the normal form of t .*

Theorem (Accattoli & Dal Lago, 2016)

There is an algorithm which takes as input a λ -term t and which, in time polynomial in the number of left-most outermost β -steps of t and the size of t

outputs an LSC term u such that the unfolding of u is the normal form of t .

Theorem (This paper, 2020)

*There is an algorithm which takes as input a **closed L-term** t and which, in time polynomial in the number of β -steps of t and the size of t
and space linear in the size of the largest term in the reduction
*outputs a **heap containing a term** u such that the unfolding of u is the normal form of t .**

Unfolding takes time polynomial in number of β -steps and size of result.

For decision functions, evaluation is polynomial and P, NP, PSPACE, EXP can be defined in terms of L ...

Theorem (This paper, 2020)

There is an algorithm which takes as input a closed L-term t and which, in time polynomial in the number of β -steps of t and the size of t and space linear in the size of the largest term in the reduction outputs a heap containing a term u such that the unfolding of u is the normal form of t .

Unfolding takes time polynomial in number of β -steps and size of result.

For decision functions, evaluation is polynomial and P, NP, PSPACE, EXP can be defined in terms of L ...

... but sublinear time or space is not covered

Unfolding takes time polynomial in number of β -steps and size of result.

Theorem (This paper, 2020)

There is an algorithm which takes as input a closed L-term t and which, in time polynomial in the number of β -steps of t

and the size of t

and space linear in the size of the largest term in the reduction

outputs a heap containing a term u such that the unfolding of u is the normal form of t .

Size explosion

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

Exponentiation: $\bar{n} \bar{m} \succ^* \overline{m^n}$

Church booleans: $\overline{true} := \lambda x y. x$

Define: $s_E := \lambda x. \overline{true} \overline{true} (x \bar{2}(\lambda x. x))$

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

with $\|t_n\| \in \Omega(2^n)$.

Easy theorem I

Theorem

Turing machines can simulate L with a constant-factor overhead in space using a naive substitution-based strategy.

Easy theorem I

Theorem

Turing machines can simulate L with a constant-factor overhead in space using a naive substitution-based strategy.

For time, this strategy is exponentially wrong.

Easy theorem II

Theorem

Turing machines can simulate L with a polynomial overhead in time using a heap-based strategy.

Easy theorem II

Theorem

Turing machines can simulate L with a polynomial overhead in time using a heap-based strategy.

For space, this strategy has a logarithmic factor overhead because pointers can become too large

Pointer explosion

$$N := (\lambda xy.xx)\overline{true}.$$

$$\begin{aligned} s_P &:= \underbrace{N(\dots(N\overline{true})\dots)}_{n \text{ times}} \\ &\succ^n \underbrace{(\lambda y.\overline{true\ true})(\dots((\lambda y.\overline{true\ true})\overline{true})\dots)}_{n \text{ times}} \\ &\succ^{2n} \overline{true} \end{aligned}$$

Pointer explosion

$$N := (\lambda xy.xx)\overline{true}.$$

$$\begin{aligned} s_P &:= \underbrace{N(\dots(N\overline{true})\dots)}_{n \text{ times}} \\ &\succ^n \underbrace{(\lambda y.\overline{true\ true})(\dots((\lambda y.\overline{true\ true})\overline{true})\dots)}_{n \text{ times}} \\ &\succ^{2n} \overline{true} \end{aligned}$$

$3n$ beta reductions \rightsquigarrow $3n$ entries on the heap

Pointer explosion

$$N := (\lambda xy.xx)\overline{true}.$$

$$\begin{aligned} s_P &:= \underbrace{N(\dots(N\overline{true})\dots)}_{n \text{ times}} \\ &\succ^n \underbrace{(\lambda y.\overline{true\ true})(\dots((\lambda y.\overline{true\ true})\overline{true})\dots)}_{n \text{ times}} \\ &\succ^{2n} \overline{true} \end{aligned}$$

$3n$ beta reductions \rightsquigarrow $3n$ entries on the heap

Heap pointers are of size $\log n$, space consumption is $\Omega(n \log n)$

Substitution-based strategy is ok for space, wrong for time
on size-exploding terms

Heap-based strategy is ok for time, wrong for space
due to pointer explosion

Substitution-based strategy is ok for space, wrong for time
on size-exploding terms

Heap-based strategy is ok for time, wrong for space
due to pointer explosion

Size-exploding terms do not exhibit pointer
explosion!

Substitution-based strategy is ok for space, wrong for time
on size-exploding terms

Heap-based strategy is ok for time, wrong for space
due to pointer explosion

Size-exploding terms do not exhibit pointer explosion!

(There's always enough space for the pointers)

Substitution-based strategy is ok for space, wrong for time
on size-exploding terms

Heap-based strategy is ok for time, wrong for space
due to pointer explosion

Solution:

Interleave both strategies.

For each k , run substitution-based strategy for k steps,
if size explodes run the heap-based strategy instead

In the paper

- Substitution-based stack machine verified in Coq w.r.t. time and space
- Heap-based stack machine verified in Coq w.r.t. time and space
- Sketch of a Turing machine simulating heap-based stack machine
- Sketch of a Turing machine simulating substitution-based stack machine while checking space consumption
- Detailed proof of interleaving strategy

Conclusion

The natural measures for the weak call-by-value λ -calculus

- are compositional
- for a compositional model of computation
- cover both time and space
- can be used to define standard complexity classes

Conclusion

The natural measures for the weak call-by-value λ -calculus

- are compositional
- for a compositional model of computation
- cover both time and space
- can be used to define standard complexity classes
- are feasible to use in mechanisations

Future Work

- mechanise basic complexity theory in Coq based on extraction

Future Work

- mechanise basic complexity theory in Coq based on extraction
- extension to the full λ -calculus
- extension to sublinear time and space
- prove $P \subseteq PSPACE$ without reference to sequential models
- space measure without size explosion

Future Work

- mechanise basic complexity theory in Coq based on extraction
- extension to the full λ -calculus
- extension to sublinear time and space
- prove $P \subseteq PSPACE$ without reference to sequential models
- space measure without size explosion

Questions?

Careful naive substitution

There is a Turing machine M_{subst} that, given

- a term s
- a binary number k indicating the number of β -steps
- a binary number m indicating the maximum space to use

halts in time $\mathcal{O}(k \cdot \text{poly}(\min(m, \|s\|_S)))$

and space $\mathcal{O}(\min(m, \|s\|_S) + \log m + \log k)$

s.t. one of the following holds:

- The machine outputs a term t , then s has normal form t and $m \geq \|s\|_S$ and $k \geq \|s\|_T$.
- The machine halts in a state named *space bound reached* and $m \leq \|s\|_S$.
- The machine halts in a state named *space bound not reached* and $k < \|s\|_T$.

Hybrid strategy

- 1 Initialise $k := 0$ (in binary)
- 2 Compute $m := \|s\| \cdot p(k)$
- 3 Run M_{subst} on s for k β -steps with space bound m :
 - ▶ If M_{subst} computes a normal form, output it and halt.
 - ▶ If M_{subst} halts in *space bound not reached*, set $k := k + 1$ and go to (2).
 - ▶ If M_{subst} halts in *space bound reached*, continue at (4).
- 4 Run M_{heap} on s for k β -steps.
 - ▶ If this computes a normal form, output it and halt.
 - ▶ Otherwise, set $k := k + 1$ and go to (2).

Hybrid strategy

- 1 Initialise $k := 0$ (in binary)
- 2 Compute $m := \|s\| \cdot p(k)$
- 3 Run M_{subst} on s for k β -steps with space bound m :
 - ▶ If M_{subst} computes a normal form, output it and halt.
 - ▶ If M_{subst} halts in *space bound not reached*, set $k := k + 1$ and go to (2).
 - ▶ If M_{subst} halts in *space bound reached*, continue at (4).
- 4 Run M_{heap} on s for k β -steps.
 - ▶ If this computes a normal form, output it and halt.
 - ▶ Otherwise, set $k := k + 1$ and go to (2).

Lemma

$\log \|s\|_T \in \mathcal{O}(\|s\|_S)$ – *there's always enough space to count steps in binary*

Turing machines as mechanised model of computation

Turing machines are

- the de-facto standard model for computational complexity theory
- simple to define
- easy to understand

Turing machines as mechanised model of computation

Turing machines are

- the de-facto standard model for computational complexity theory
- simple to define
- easy to understand

“Turing machines as model of computation are *inherently infeasible* for the formalisation of any computability or complexity theoretic result.”

Verified Programming of Turing Machines in Coq

Yannick Forster
Saarland University
Saarbrücken, Germany
forster@ps.uni-saarland.de

Fabian Kunze
Saarland University
Saarbrücken, Germany
kunze@ps.uni-saarland.de

Maximilian Wuttke
Saarland University
Saarbrücken, Germany
sflmaxwutt@stud.uni-saarland.de

Abstract

We present a framework for the verified programming of multi-tape Turing machines in Coq. Improving on prior work by Asperti and Ricciotti in MatTA, we implement multiple layers of abstraction. The highest layer allows a user to implement nontrivial algorithms as Turing machines and verify their correctness, as well as time and space complexity compositionally. The user can do so without ever mentioning states, symbols on tapes or transition functions: They write programs in an imperative language with registers containing values of encodable data types, and our framework constructs corresponding Turing machines.

As case studies, we verify a translation from multi-tape to

not clear at all how to compose a two-tape Turing machine with a three-tape Turing machine that works on a different alphabet. Therefore, it is common to rely on pseudo code prose describing the intended behaviour. The exact implementation as well as its correctness or resource analysis left as an exercise to the reader. In a mechanised proof, all details cannot be left out. Luckily, it is possible to hide all details behind suitable abstractions.

We present a framework that allows the user to have the cake eat it too when it comes to mechanising computation terms of Turing machines. Algorithms are stated in the syntax of a register based imperative language; a corresponding Turing machine is automatically constructed behind the scene.