

The Time and Space of Interaction

Beniamino Accattoli¹ Ugo Dal Lago² Gabriele Vanoni²

¹INRIA & LIX, École Polytechnique

²Università di Bologna & INRIA

CHoCoLa, 22 April, 2021

Work appeared or to appear in PPDP 2020, POPL 2021, LICS 2021

Complexity and the λ -calculus

The Simplest Setting: Closed Call-by-Name λ -calculus

We set our analysis in the simplest higher-order Turing-complete framework: pure untyped λ -calculus with weak head reduction. Abstractions are the *only* values.

The Simplest Setting: Closed Call-by-Name λ -calculus

We set our analysis in the simplest higher-order Turing-complete framework: pure untyped λ -calculus with weak head reduction. Abstractions are the *only* values.

Closed CbN λ -calculus

Terms $t, u ::= x \mid \lambda x.t \mid tu$

Context $C ::= \langle \cdot \rangle \mid Ct$

Reduction $C\langle(\lambda x.t)u\rangle \rightarrow C\langle t\{x \leftarrow u\}\rangle$

The Simplest Setting: Closed Call-by-Name λ -calculus

We set our analysis in the simplest higher-order Turing-complete framework: pure untyped λ -calculus with weak head reduction. Abstractions are the *only* values.

Closed CbN λ -calculus

Terms $t, u ::= x \mid \lambda x.t \mid tu$

Context $C ::= \langle \cdot \rangle \mid Ct$

Reduction $C\langle(\lambda x.t)u\rangle \rightarrow C\langle t\{x \leftarrow u\}\rangle$

Of course, λ -terms may **diverge**, e.g. $\Omega = (\lambda x.xx)(\lambda x.xx) \rightarrow \Omega$.

Reasonable machines simulate each other with polynomially bounded overhead in time and constant factor overhead in space.

Slot and van Emde Boas

Reasonable machines simulate each other with polynomially bounded overhead in time and constant factor overhead in space.

Slot and van Emde Boas

Turing machines and RAMs are reasonable with the usual cost models.

Reasonable machines simulate each other with polynomially bounded overhead in time and constant factor overhead in space.

Slot and van Emde Boas

Turing machines and RAMs are reasonable with the usual cost models.

Can we implement the λ -calculus in a **reasonable** way?

Polynomial Overhead Time Simulations

We consider the number of β -steps as the time cost model for the λ -calculus.

We consider the number of β -steps as the time cost model for the λ -calculus.

- The direction TM \rightarrow λ -calculus is **easy**. [Dal Lago and Accattoli 2017]

We consider the number of β -steps as the time cost model for the λ -calculus.

- The direction TM \rightarrow λ -calculus is **easy**. [Dal Lago and Accattoli 2017]
- The direction λ -calculus \rightarrow TM is more delicate, because of **size explosion**.

We consider the number of β -steps as the time cost model for the λ -calculus.

- The direction TM \rightarrow λ -calculus is **easy**. [Dal Lago and Accattoli 2017]
- The direction λ -calculus \rightarrow TM is more delicate, because of **size explosion**.

Size Explosion

There exists a family of λ -terms t_n such that:

We consider the number of β -steps as the time cost model for the λ -calculus.

- The direction TM \rightarrow λ -calculus is **easy**. [Dal Lago and Accattoli 2017]
- The direction λ -calculus \rightarrow TM is more delicate, because of **size explosion**.

Size Explosion

There exists a family of λ -terms t_n such that:

$$|t_n| = \Theta(n)$$

We consider the number of β -steps as the time cost model for the λ -calculus.

- The direction TM \rightarrow λ -calculus is **easy**. [Dal Lago and Accattoli 2017]
- The direction λ -calculus \rightarrow TM is more delicate, because of **size explosion**.

Size Explosion

There exists a family of λ -terms t_n such that:

$$|t_n| = \Theta(n) \quad t_n \rightarrow_{\beta}^n \text{whnf}(t_n)$$

We consider the number of β -steps as the time cost model for the λ -calculus.

- The direction TM \rightarrow λ -calculus is **easy**. [Dal Lago and Accattoli 2017]
- The direction λ -calculus \rightarrow TM is more delicate, because of **size explosion**.

Size Explosion

There exists a family of λ -terms t_n such that:

$$|t_n| = \Theta(n) \quad t_n \rightarrow_{\beta}^n \text{whnf}(t_n) \quad |\text{whnf}(t_n)| = \Theta(2^n)$$

We consider the number of β -steps as the time cost model for the λ -calculus.

- The direction TM \rightarrow λ -calculus is **easy**. [Dal Lago and Accattoli 2017]
- The direction λ -calculus \rightarrow TM is more delicate, because of **size explosion**.

Size Explosion

There exists a family of λ -terms t_n such that:

$$|t_n| = \Theta(n) \quad t_n \xrightarrow{\beta}^n \text{whnf}(t_n) \quad |\text{whnf}(t_n)| = \Theta(2^n)$$

A nice proof can be carried out through the use of **abstract machines** (e.g. Krivine's).

Constant Overhead Space Simulations

It is not clear what the **space cost model** for the λ -calculus should be.

It is not clear what the **space cost model** for the λ -calculus should be.

Considering the maximum size of a term in a computation, one would have $\text{SPACE} > \text{TIME}$, because of **size explosion**.

It is not clear what the **space cost model** for the λ -calculus should be.

Considering the maximum size of a term in a computation, one would have $\text{SPACE} > \text{TIME}$, because of **size explosion**.

Abstract machines **share** redex arguments, thus consuming $\text{SPACE} \simeq \text{TIME}$ (unless some garbage collection mechanism is used).

We want to be **parsimonious** in **space**.

We want to be **parsimonious** in **space**.

Ideas

- 1 No tracing of β -redexes

We want to be **parsimonious** in **space**.

Ideas

- 1 No tracing of β -redexes
⇒ space disentangled from time.

We want to be **parsimonious** in **space**.

Ideas

- 1 No tracing of β -redexes
⇒ space disentangled from time.
- 2 **Backtracking** to retrieve β -redexes

We want to be **parsimonious** in **space**.

Ideas

- 1 No tracing of β -redexes
⇒ space disentangled from time.
- 2 **Backtracking** to retrieve β -redexes
⇒ wasting time to gain space.

We want to be **parsimonious** in **space**.

Ideas

- 1 No tracing of β -redexes
⇒ space disentangled from time.
- 2 **Backtracking** to retrieve β -redexes
⇒ wasting time to gain space.
- 3 Computation is **local**

We want to be **parsimonious** in **space**.

Ideas

- 1 No tracing of β -redexes
⇒ space disentangled from time.
- 2 **Backtracking** to retrieve β -redexes
⇒ wasting time to gain space.
- 3 Computation is **local**
⇒ a **token** that travels inside the term.

We want to be **parsimonious** in **space**.

Ideas

- 1 No tracing of β -redexes
⇒ space disentangled from time.
- 2 **Backtracking** to retrieve β -redexes
⇒ wasting time to gain space.
- 3 Computation is **local**
⇒ a **token** that travels inside the term.
- 4 The code never changes.

The Interaction Abstract and Space Complexity

- **Mackie**: small runtime system for PCF.

- **Mackie**: small runtime system for PCF.
- **Ghica**: compilation of higher-order functions to digital circuits.


- **Mackie**: small runtime system for PCF.
- **Ghica**: compilation of higher-order functions to digital circuits.
- **Dal Lago** and **Schöpp**: functional programming in Logspace.

- **Mackie**: small runtime system for PCF.
- **Ghica**: compilation of higher-order functions to digital circuits.
- **Dal Lago** and **Schöpp**: functional programming in Logspace.
- **Aubert et al.**: logic programming in Logspace.

The position of the token inside the term t is represented via a pair (u, C) such that $C\langle u \rangle = t$.

The position of the token inside the term t is represented via a pair (u, C) such that $C\langle u \rangle = t$.


Example



The position of the token ■ in the term $\lambda x.t$ is represented by the position $(t, \lambda x.\langle \cdot \rangle)$.


No information is saved, traversing a β -redex $(\lambda x.t)u$. The token remains untouched.

No information is saved, traversing a β -redex $(\lambda x.t)u$. The token remains untouched.



The token ■ is at the root of the redex.


No information is saved, traversing a β -redex $(\lambda x.t)u$. The token remains untouched.



The token \blacksquare moves to the left-hand side of the application, changing color.

$$ru \mid C \mid \blacksquare \rightarrow_{\bullet 1} r \mid C \langle \langle \cdot \rangle u \rangle \mid \square$$

No information is saved, traversing a β -redex $(\lambda x.t)u$. The token remains untouched.



The token \blacksquare moves to the body of the abstraction, changing color again.

$$\lambda x.t \mid C \mid \square \rightarrow_{\bullet 2} t \mid C\langle \lambda x.\langle \cdot \rangle \rangle \mid \blacksquare$$

Quering Variables and their Arguments

Computation in the λ -calculus is done by substituting variables for arguments.

Quering Variables and their Arguments


Computation in the λ -calculus is done by substituting variables for arguments.

Our machine first looks for variables, in **red** mode, going deep inside the term.

Querying Variables and their Arguments

Computation in the λ -calculus is done by substituting variables for arguments.

Our machine first looks for variables, in **red** mode, going deep inside the term.




The token ■ is pointing a variable.

Querying Variables and their Arguments

Computation in the λ -calculus is done by substituting variables for arguments.

Our machine first looks for variables, in **red** mode, going deep inside the term.




The token **■** moves locally to the binder, changing its mode to **blue**, *i.e.* the machine is now looking for the argument of the variable.

$$x \mid C \langle \lambda x . D \rangle \mid \blacksquare \rightarrow_{\text{var}} \lambda x . D \langle x \rangle \mid C \mid \blacksquare$$


When an argument is found, it has to be evaluated.

When an argument is found, it has to be evaluated.



The token ■, *i.e.* looking for an argument is on the left-hand side of an application.

When an argument is found, it has to be evaluated.



The token moves to the argument, changing its mode to **red**, *i.e.* looking for the head variable.

$$t \mid C\langle\langle\cdot\rangle u\rangle \mid \blacksquare \rightarrow_{\text{arg}} u \mid C\langle t\langle\cdot\rangle\rangle \mid \blacksquare$$

Hiding Some Details

There is also a **backtracking** mechanism (not explained here).

There is also a **backtracking** mechanism (not explained here).

It uses position **enriched** with history/log.

There is also a **backtracking** mechanism (not explained here).

It uses position **enriched** with history/log.

The token is actually made of two data structures, tape and log.

There is also a **backtracking** mechanism (not explained here).

It uses position **enriched** with history/log.

The token is actually made of two data structures, tape and log.

Implementation Theorem

The λ -term t has weak head normal form iff the IAM terminates on $(\underline{t}, \langle \cdot \rangle, \epsilon, \epsilon)$.

Types and Complexity

Non-Idempotent Intersection Types

Sequence Types for the CbN λ -Calculus

Linear Types $A ::= \star \mid S \rightarrow A$

Sequence Types $S ::= [A_1 \cdots A_n] \quad n \geq 0$

Sequence Types for the CbN λ -Calculus

Linear Types $A ::= \star \mid S \rightarrow A$

Sequence Types $S ::= [A_1 \cdots A_n] \quad n \geq 0$

Correctness and Completeness

$\vdash t : \star$ if and only if $t \Downarrow$.

A Bridge Between Syntax and Semantics

Sequence types are nothing but another way of formulating operational/denotational **semantics** of the λ -calculus in a Linear Logic fashion. There are many other examples.

A Bridge Between Syntax and Semantics

Sequence types are nothing but another way of formulating operational/denotational **semantics** of the λ -calculus in a Linear Logic fashion. There are many other examples.

Operational Semantics

A Bridge Between Syntax and Semantics

Sequence types are nothing but another way of formulating operational/denotational **semantics** of the λ -calculus in a Linear Logic fashion. There are many other examples.

Operational Semantics

- Linear Substitution Calculus (*i.e. proof net reduction*)

A Bridge Between Syntax and Semantics

Sequence types are nothing but another way of formulating operational/denotational **semantics** of the λ -calculus in a Linear Logic fashion. There are many other examples.

Operational Semantics

- Linear Substitution Calculus (*i.e. proof net reduction*)
- Krivine's Abstract Machine

A Bridge Between Syntax and Semantics

Sequence types are nothing but another way of formulating operational/denotational **semantics** of the λ -calculus in a Linear Logic fashion. There are many other examples.

Operational Semantics

- Linear Substitution Calculus (*i.e. proof net reduction*)
- Krivine's Abstract Machine
- Interaction Abstract Machine

A Bridge Between Syntax and Semantics

Sequence types are nothing but another way of formulating operational/denotational **semantics** of the λ -calculus in a Linear Logic fashion. There are many other examples.

Operational Semantics

- Linear Substitution Calculus (*i.e. proof net reduction*)
- Krivine's Abstract Machine
- Interaction Abstract Machine
- Polyadic approximations

A Bridge Between Syntax and Semantics

Sequence types are nothing but another way of formulating operational/denotational **semantics** of the λ -calculus in a Linear Logic fashion. There are many other examples.

Operational Semantics

- Linear Substitution Calculus (*i.e. proof net reduction*)
- Krivine's Abstract Machine
- Interaction Abstract Machine
- Polyadic approximations

Denotational Semantics

A Bridge Between Syntax and Semantics

Sequence types are nothing but another way of formulating operational/denotational **semantics** of the λ -calculus in a Linear Logic fashion. There are many other examples.

Operational Semantics

- Linear Substitution Calculus (*i.e. proof net reduction*)
- Krivine's Abstract Machine
- Interaction Abstract Machine
- Polyadic approximations

Denotational Semantics

- **Relational model**

A Bridge Between Syntax and Semantics

Sequence types are nothing but another way of formulating operational/denotational **semantics** of the λ -calculus in a Linear Logic fashion. There are many other examples.

Operational Semantics

- Linear Substitution Calculus (*i.e. proof net reduction*)
- Krivine's Abstract Machine
- Interaction Abstract Machine
- Polyadic approximations

Denotational Semantics

- **Relational model**
- Game semantics

$$\frac{}{x : [A] \vdash x : A} \text{T-Var} \qquad \frac{\Gamma, x : S \vdash t : A}{\Gamma \vdash \lambda x.t : S \rightarrow A} \text{T-}\lambda \qquad \frac{}{\Gamma \vdash \lambda x.t : \star} \text{T-}\lambda_{\star}$$

$$\frac{\Gamma \vdash t : [A'_1, \dots, A'_n] \rightarrow A \quad [\Delta_i \vdash u : A'_i]_{i \in [1, \dots, n]}}{\Gamma \uplus \sum_{i \in [1, \dots, n]} \Delta_i \vdash tu : A} \text{T-}\odot$$

A Machine Built on Type Derivations

Given a type derivation ending in $\vdash t : \star$, we can consider a **token machine** that starts from the root occurrence of \star .

Given a type derivation ending in $\vdash t : \star$, we can consider a **token machine** that starts from the root occurrence of \star .

One can move in the *natural* way inside the derivation, following that occurrence of \star .

Given a type derivation ending in $\vdash t : \star$, we can consider a **token machine** that starts from the root occurrence of \star .

One can move in the *natural* way inside the derivation, following that occurrence of \star .

The obtained machine is strongly bisimilar to the **Interaction Abstract Machine**.

Given a type derivation ending in $\vdash t : \star$, we can consider a **token machine** that starts from the root occurrence of \star .

One can move in the *natural* way inside the derivation, following that occurrence of \star .

The obtained machine is strongly bisimilar to the **Interaction Abstract Machine**.

The **multiplicative** part is handled by **types**.

Given a type derivation ending in $\vdash t : \star$, we can consider a **token machine** that starts from the root occurrence of \star .

One can move in the *natural* way inside the derivation, following that occurrence of \star .

The obtained machine is strongly bisimilar to the **Interaction Abstract Machine**.

The **multiplicative** part is handled by **types**.

The **exponential** part is **no more needed**: in type derivations each subterm has already been copied/discarded. This way, each copy is used only once, in a **linear** way.

$$\begin{array}{c}
 \frac{x : [[\star] \rightarrow \star] \vdash x : [\star_{\downarrow 16}] \rightarrow \star_{\uparrow 6} \quad y : [\star] \vdash y : \star_{\uparrow 17}}{y : [\star], x : [[\star] \rightarrow \star] \vdash xy : \star_{\uparrow 5}} \\
 \frac{y : [\star] \vdash \lambda x. xy : [[\star_{\uparrow 15}] \rightarrow \star_{\downarrow 7}] \rightarrow \star_{\uparrow 4}}{\vdash \lambda y. \lambda x. xy : [\star_{\downarrow 18}] \rightarrow [[\star_{\uparrow 14}] \rightarrow \star_{\downarrow 8}] \rightarrow \star_{\uparrow 3}} \quad \frac{}{\vdash I : \star_{\uparrow 19}} \quad \frac{z : [\star] \vdash z : \star_{\uparrow 11}}{\vdash \lambda z. z : [\star_{\downarrow 12}] \rightarrow \star_{\uparrow 10}} \\
 \hline
 \vdash (\lambda y. \lambda x. xy)I : [[\star_{\uparrow 13}] \rightarrow \star_{\downarrow 9}] \rightarrow \star_{\uparrow 2} \\
 \hline
 \vdash (\lambda y. \lambda x. xy)I(\lambda z. z) : \star_{\uparrow 1}
 \end{array}$$

Computing the Run Time

We observe the following facts:

We observe the following facts:

- there is only **one** initial state,

We observe the following facts:

- there is only **one** initial state,
- there are **no cycles**.

We observe the following facts:

- there is only **one** initial state,
- there are **no cycles**.

As a consequence:

We observe the following facts:

- there is only **one** initial state,
- there are **no cycles**.

As a consequence:

- 1 all states are reached exactly **once**, and thus

We observe the following facts:

- there is only **one** initial state,
- there are **no cycles**.

As a consequence:

- 1 all states are reached exactly **once**, and thus
- 2 run time = number of states = number of occurrences of \star in the derivation.

$$\|\star\| = 1 \quad \|S \rightarrow A\| = \|S\| + \|A\| \quad \|[A_1, \dots, A_n]\| = \sum_{1 \leq i \leq n} \|A_i\|$$

$$\frac{}{\vdash^0 \lambda x.t : \star} \quad \frac{\vdash^w}{\vdash^{w+\|S \rightarrow A\|} \lambda x.t : S \rightarrow A}$$

$$\frac{}{\vdash^{\|A\|} x : A} \quad \frac{\vdash^w \quad [\vdash^{v_i}]_{i \in [1, \dots, n]}}{\vdash^{w+\sum v_i+\|A\|} tu : A}$$

Higher-Order Types need Time

- **IAM (unreasonable?):** weights, thus run time, depend on types,

- **IAM (unreasonable?)**: weights, thus run time, depend on types,
- **KAM (time reasonable)**: weights, thus run time, do *not* depend on types.

- **IAM (unreasonable?)**: weights, thus run time, depend on types,
- **KAM (time reasonable)**: weights, thus run time, do *not* depend on types.

Indeed, give a term t such that $t \rightarrow^n \lambda x.u$:

- **IAM (unreasonable?)**: weights, thus run time, depend on types,
- **KAM (time reasonable)**: weights, thus run time, do *not* depend on types.

Indeed, give a term t such that $t \rightarrow^n \lambda x.u$:

- **types** of subterms of t can be **exponential** in $|t|$ and n ,

- **IAM (unreasonable?)**: weights, thus run time, depend on types,
- **KAM (time reasonable)**: weights, thus run time, do *not* depend on types.

Indeed, give a term t such that $t \rightarrow^n \lambda x.u$:

- **types** of subterms of t can be **exponential** in $|t|$ and n ,
- the **size** of the type derivation for t is **polynomial** in $|t|$ and n .

- **IAM (unreasonable?)**: weights, thus run time, depend on types,
- **KAM (time reasonable)**: weights, thus run time, do *not* depend on types.

Indeed, give a term t such that $t \rightarrow^n \lambda x.u$:

- **types** of subterms of t can be **exponential** in $|t|$ and n ,
- the **size** of the type derivation for t is **polynomial** in $|t|$ and n .

Ultimately, IAM inefficiency comes from the use of **higher-order** types.

- IAM run time can be characterized by intersection type derivations.

How Can Space Be Captured?

How Can Space Be Captured?

We need to refine the type system:

How Can Space Be Captured?

We need to refine the type system:

Sequence Types for the CbN λ -Calculus

Linear types $A, A' ::= \star \mid T \rightarrow A$

Tree types $T, T' ::= [G_1, \dots, G_n] \quad n \geq 0$

(Generic) Types $G, G' ::= A \mid T$

How Can Space Be Captured?

We need to refine the type system:

Sequence Types for the CbN λ -Calculus

Linear types $A, A' ::= \star \mid T \rightarrow A$

Tree types $T, T' ::= [G_1, \dots, G_n] \quad n \geq 0$

(Generic) Types $G, G' ::= A \mid T$

Space Type Measure

$$\|\star\| := 0 \quad \|T \rightarrow A\| := \max\{\|T\|, \|A\| + 1\}$$
$$\|[G_1, \dots, G_n]\| := \mathbf{X} + \max_i \{\|G_i\|\}$$

$$\frac{}{x : [A] \vdash x : A} \text{T-Var}$$

$$\frac{\Gamma, x : T \vdash t : A}{\Gamma \vdash \lambda x.t : T \rightarrow A} \text{T-}\lambda$$

$$\frac{\Gamma \vdash t : T \rightarrow A \quad \Delta \vdash u : T}{\Gamma \uplus \Delta \vdash tu : A} \text{T-}\odot$$

$$\frac{}{\Gamma \vdash \lambda x.t : \star} \text{T-}\lambda_{\star}$$

$$\frac{\Gamma_i \vdash t : G_i \quad 1 \leq i \leq n}{[\uplus_{i=1}^n \Gamma_i] \vdash t : [G_1, \dots, G_n]} \text{T-many}$$

$$\frac{}{\vdash t : [\cdot]} \text{T-none}$$

The Inference Rules - Weighted

$$\frac{}{x : [A] \vdash x : A} \text{ T-Var}$$

$$\frac{\Gamma, x : T \vdash t : A}{\Gamma \vdash \lambda x. t : T \rightarrow A} \text{ T-}\lambda$$

$$\frac{\Gamma \vdash t : T \rightarrow A \quad \Delta \vdash u : T}{\Gamma \uplus \Delta \vdash tu : A} \text{ T-}\odot$$

$$\frac{}{\Gamma \vdash \lambda x. t : \star} \text{ T-}\lambda_{\star}$$

$$\frac{\Gamma_i \vdash t : G_i \quad 1 \leq i \leq n}{[\uplus_{i=1}^n \Gamma_i] \vdash t : [G_1, \dots, G_n]} \text{ T-many}$$

$$\frac{}{\Gamma \vdash t : [\cdot]} \text{ T-none}$$

- IAM run time can be characterized by sequence type derivations.
- IAM space consumption can be characterized by tree type derivations.

We would like to prove that the IAM is invariant.

We would like to prove that the IAM is invariant.

About time, the difficult direction is $\text{TM} \rightarrow \text{IAM}$. We need to reason about the IAM execution of terms which are the image of *an* encoding of TM in the λ -calculus.

We would like to prove that the IAM is invariant.

About time, the difficult direction is $\text{TM} \rightarrow \text{IAM}$. We need to reason about the IAM execution of terms which are the image of *an* encoding of TM in the λ -calculus.

The *while loop* of Turing machines is typically implemented through the fixed point combinator $\Theta := \theta\theta$.

We would like to prove that the IAM is invariant.

About time, the difficult direction is $\text{TM} \rightarrow \text{IAM}$. We need to reason about the IAM execution of terms which are the image of *an* encoding of TM in the λ -calculus.

The *while loop* of Turing machines is typically implemented through the fixed point combinator $\Theta := \theta\theta$.

$$\mathbb{F}_{n+1}^{\bar{A}} := [\mathbb{Y}_{n+1}^{\bar{A}}, \dots, \mathbb{Y}_0^{\bar{A}}] \rightarrow A_{n+1}$$

$$\mathbb{Y}_{n+1}^{\bar{A}} := [A_n] \rightarrow A_{n+1}$$

$$\mathbb{X}_{n+1}^{\bar{A}} := [\mathbb{X}_n^{\bar{A}}, \dots, \mathbb{X}_0^{\bar{A}}] \rightarrow \mathbb{F}_{n+1}^{\bar{A}}$$

We would like to prove that the IAM is invariant.

About time, the difficult direction is $\text{TM} \rightarrow \text{IAM}$. We need to reason about the IAM execution of terms which are the image of *an* encoding of TM in the λ -calculus.

The *while loop* of Turing machines is typically implemented through the fixed point combinator $\Theta := \theta\theta$.

$$\begin{aligned}\mathbb{F}_{n+1}^{\bar{A}} &:= [\mathbb{Y}_{n+1}^{\bar{A}}, \dots, \mathbb{Y}_0^{\bar{A}}] \rightarrow A_{n+1} \\ \mathbb{Y}_{n+1}^{\bar{A}} &:= [A_n] \rightarrow A_{n+1} \\ \mathbb{X}_{n+1}^{\bar{A}} &:= [\mathbb{X}_n^{\bar{A}}, \dots, \mathbb{X}_0^{\bar{A}}] \rightarrow \mathbb{F}_{n+1}^{\bar{A}}\end{aligned}$$

Lemma

For each $n \geq 0$ and type list \bar{A} , $\vdash \theta : \mathbb{X}_n^{\bar{A}}$ and $\vdash \Theta : \mathbb{F}_n^{\bar{A}}$.

Typing Turing Machines

We would like to prove that the IAM is invariant.

About time, the difficult direction is $\text{TM} \rightarrow \text{IAM}$. We need to reason about the IAM execution of terms which are the image of *an* encoding of TM in the λ -calculus.

The *while loop* of Turing machines is typically implemented through the fixed point combinator $\Theta := \theta\theta$.

$$\begin{aligned}\mathbb{F}_{n+1}^{\bar{A}} &:= [\mathbb{Y}_{n+1}^{\bar{A}}, \dots, \mathbb{Y}_0^{\bar{A}}] \rightarrow A_{n+1} \\ \mathbb{Y}_{n+1}^{\bar{A}} &:= [A_n] \rightarrow A_{n+1} \\ \mathbb{X}_{n+1}^{\bar{A}} &:= [\mathbb{X}_n^{\bar{A}}, \dots, \mathbb{X}_0^{\bar{A}}] \rightarrow \mathbb{F}_{n+1}^{\bar{A}}\end{aligned}$$

Lemma

For each $n \geq 0$ and type list \bar{A} , $\vdash \theta : \mathbb{X}_n^{\bar{A}}$ and $\vdash \Theta : \mathbb{F}_n^{\bar{A}}$.

Observation

$$\|\mathbb{X}_n\| = \Omega(2^n).$$

- IAM run time can be characterized by sequence type derivations.
- IAM space consumption can be characterized by tree type derivations.
- Given a TM that halts in n steps, its encoding in the λ -calculus takes $\Omega(2^n)$ steps when evaluated on the IAM.

The fixed point operator can be analogously typed with tree types.

The fixed point operator can be analogously typed with tree types.

$$\begin{aligned} \Theta : \mathbb{F}_{n+1}^{\bar{A}} &:= \mathbb{T}_{n+1}^{\bar{A}} \rightarrow A_{n+1} \\ \text{where } \mathbb{T}_{n+1}^{\bar{A}} &:= [\mathbb{Y}_{n+1}^{\bar{A}}, [\mathbb{T}_n^{\bar{A}}]] \\ \text{and } \mathbb{Y}_{n+1}^{\bar{A}} &:= [A_n] \rightarrow A_{n+1} \end{aligned}$$

The fixed point operator can be analogously typed with tree types.

$$\begin{aligned}\Theta : \mathbb{F}_{n+1}^{\bar{A}} &:= \mathbb{T}_{n+1}^{\bar{A}} \rightarrow A_{n+1} \\ \text{where } \mathbb{T}_{n+1}^{\bar{A}} &:= [\mathbb{Y}_{n+1}^{\bar{A}}, [\mathbb{T}_n^{\bar{A}}]] \\ \text{and } \mathbb{Y}_{n+1}^{\bar{A}} &:= [A_n] \rightarrow A_{n+1}\end{aligned}$$

Lemma

For each $n \geq 0$ and type list \bar{A} , $\vdash \theta : \mathbb{F}_n^{\bar{A}}$.

The fixed point operator can be analogously typed with tree types.

$$\begin{aligned}\Theta : \mathbb{F}_{n+1}^{\bar{A}} &:= \mathbb{T}_{n+1}^{\bar{A}} \rightarrow A_{n+1} \\ \text{where } \mathbb{T}_{n+1}^{\bar{A}} &:= [\mathbb{Y}_{n+1}^{\bar{A}}, [\mathbb{T}_n^{\bar{A}}]] \\ \text{and } \mathbb{Y}_{n+1}^{\bar{A}} &:= [A_n] \rightarrow A_{n+1}\end{aligned}$$

Lemma

For each $n \geq 0$ and type list \bar{A} , $\vdash \theta : \mathbb{F}_n^{\bar{A}}$.

Observation

$$\|\mathbb{F}_n\| = \Omega(n).$$

- IAM run time can be characterized by sequence type derivations.
- IAM space consumption can be characterized by tree type derivations.
- Given a TM that halts in n steps, its encoding in the λ -calculus takes $\Omega(2^n)$ steps when evaluated on the IAM.
- Given a TM that halts in n steps, its encoding in the λ -calculus consumes $\Omega(n)$ space when evaluated on the IAM.

Analyzing the space/time behaviour of environment machines *with* garbage collection.