

The safe λ -calculus is a fragment of the simply typed λ -calculus introduced in [BO09], motivated by the theory of higher-order recursion schemes. It has some interesting properties, e.g. there is a normalization procedure that never needs to perform any α -renaming (i.e. uses capture-permitting substitution).

Determining the precise complexity of deciding whether two safe λ -terms are $\beta\eta$ -convertible has been an open problem until now (to my knowledge). In [BO09, §3], a PSPACE-hardness result is established. This is far below the non-elementary lower bound for simply typed $\beta\eta$ -convertibility due to Statman, but it is argued that both Statman’s proof and its simplification by Mairson fundamentally require unsafe terms. To quote [BO09], this “does not rule out the possibility that another non-elementary problem is encodable in the safe lambda calculus”.

We provide such an encoding here: by reduction from the star-free expression equivalence problem, we show that safe $\beta\eta$ -convertibility is TOWER-complete in the sense of [Sch16]. (We only provide a proof of TOWER-hardness, membership holds because it works for STLC already. [TODO: actually this proves that $\beta\eta$ in STLC is TOWER-c which is not explicitly written anywhere in the published literature])

First, let us recall the definitions of our objects of interest:

Definition 1 (slight simplification of [BO09]). As usual, the *order* of a simple type is the nesting depth of function arrows to the left:

$$\text{ord}(o) = 0 \quad \text{ord}(A \rightarrow B) = \max(\text{ord}(A) + 1, \text{ord}(B))$$

The typing rules of the safe λ -calculus are

$$\frac{}{x : A \vdash x : A} \quad \frac{\Theta \vdash t : A}{\Theta' \vdash t : A} (\Theta \subset \Theta')$$

$$\frac{\Theta \vdash t : A_1 \rightarrow \dots \rightarrow A_n \rightarrow B \quad \Theta \vdash u_1 : A_1 \dots \Theta \vdash u_n : A_n}{\Theta \vdash t u_1 \dots u_n : B} \text{ when } \text{ord}(B) \leq \inf_{(y:C) \in \Theta} \text{ord}(C)$$

$$\frac{\Theta, x_1 : A_1, \dots, x_n : A_n \vdash t : B}{\Theta \vdash \lambda x_1. \dots \lambda x_n. t : A_1 \rightarrow \dots \rightarrow A_n \rightarrow B} \text{ when } \text{ord}(A_1 \rightarrow \dots \rightarrow A_n \rightarrow B) \leq \inf_{(y:C) \in \Theta} \text{ord}(C)$$

with the usual convention $\inf(\emptyset) = +\infty$.

Definition 2. *Star-free expressions* are regular expressions without the Kleene star, but with complementation:

$$E, F ::= \emptyset \mid \varepsilon \mid c \mid E \cup F \mid E \cdot F \mid E^c$$

(for c in an alphabet Σ). The *star-free equivalence problem* consists in deciding whether two star-free expressions denote the same language (subset of Σ^*).

Definition 3 (classical). The types $\mathbf{Str}_\Sigma = (o \rightarrow o) \rightarrow \dots \rightarrow (o \rightarrow o) \rightarrow o \rightarrow o$ (with $|\Sigma|$ occurrences of $(o \rightarrow o)$) and $\mathbf{Bool} = o \rightarrow o \rightarrow o$ are the so-called *Church encodings* of strings and booleans. In the case of a unary alphabet, we set $\mathbf{Nat} = \mathbf{Str}_{\{c\}} = (o \rightarrow o) \rightarrow o \rightarrow o$: it is the type of Church *numerals*.

The Church encoding of a word $w = w_1 \dots w_n$ over an ordered alphabet $\Sigma = \{c_1, \dots, c_n\}$ is $\bar{w} = \lambda f_{c_1}. \dots \lambda f_{c_n}. \lambda x. f_{w[1]} (\dots (f_{w[n]} x)) : \mathbf{Str}_\Sigma$. The Church encoding of $n \in \mathbb{N}$ is $\bar{n} = \bar{c} \dots \bar{c} : \mathbf{Nat}$ with n times c .

Note that \mathbf{Bool} only has two closed β -normal inhabitants, $\mathbf{true} = \lambda xy. x$ and $\mathbf{false} = \lambda xy. y$. We will want to prove the following:

Theorem 4. *Given a safe λ -term t , it is TOWER-hard to decide whether $t =_{\beta\eta} \mathbf{true}$.*

This will proceed by reduction from star-free equivalence, which was given in [Sch16, §3.1] as an example of a TOWER-complete problem. Any elementary complexity reduction will do, in fact we will even have a polynomial time reduction (or exponential time if we insist on writing out the full type derivation for the term we build).

First, note that star-free equivalence reduces to star-free *emptiness*: given two expressions E and F , the language denoted by $(E^c \cup F)^c \cup (E \cup F^c)^c$ is empty if and only if E and F are equivalent. This new problem can be solved by a kind of bounded search:

Lemma 5 (classical). *Suppose that the expression E of size n denotes a non-empty language. This language then contains a word of length at most $\text{tower}(n)$.*

(Recall that $\text{tower}(n+1) = 2^{\text{tower}(n)}$.)

Proof sketch. We translate E to an equivalent nondeterministic finite automaton (NFA), whose number of states bounds the length of a shortest word (indeed such a word has an accepting run that visits each state at most once). This can be done by induction of E , using any standard construction on NFA for union and concatenation; the costliest operation is complementation, for which we use determinization, inducing a single exponential state blowup. \square

Remark 6. If we wanted to inductively build *deterministic* finite automata, then the exponential blowup would occur on the treatment of concatenation. In fact it is a well-understood phenomenon that the source of complexity lies in the *alternations* of concatenation and complementation (cf. dot-depth / Straubing-Thérien hierarchy).

Thus, we only have to test E against all inputs of length at most $\text{tower}(n)$. We do so using the following lemmas (notation: $A[B] = A[B/o]$; note that any $t : A$ can be typecast to a term of type $A[B]$):

Lemma 7 (see [BO09, Remark 2.5(iii)]). *One can build a safe λ -term that is $\beta\eta$ -convertible to $\text{tower}(n) : \mathbf{Nat}$ in polynomial time.*

Lemma 8. *For any finite alphabet Σ and $\# \notin \Sigma$, there exists a safe λ -term $\mathbf{enum} : \mathbf{Nat}[A_{\mathbf{enum}}] \rightarrow \mathbf{Str}_{\Sigma \cup \{\#\}}$ such that for any $m \in \mathbb{N}$, $\mathbf{enum} \bar{m}$ reduces to the encoding of a $\#$ -separated list containing all words in Σ^* of length at most m .*

Lemma 9. *There exists a simply typed λ -term $\mathbf{exists} : (\mathbf{Str}_\Sigma[\alpha] \rightarrow \mathbf{Bool}) \rightarrow \mathbf{Str}_{\Sigma \cup \{\#\}}[F_{\mathbf{exists}}(\alpha)] \rightarrow \mathbf{Bool}$, where α is a type variable, whose instantiations $\alpha = A$ are all safe and test whether some word (over Σ) in a given list of words (encoded using the separator $\#$) satisfies a given predicate.*

Lemma 10. *E can be turned in polynomial time into an equivalent safe λ -term $t_E : \mathbf{Str}_\Sigma[A_E] \rightarrow \mathbf{Bool}$.*

We conclude by composing the four terms built in the above lemmas and $\mathbf{not} : \mathbf{Bool} \rightarrow \mathbf{Bool}$; this is allowed because safe terms of type $A[T] \rightarrow B$ and $B[U] \rightarrow C$ can be composed into a safe term of type $A[T[U]] \rightarrow C$ whenever $\text{ord}(B) \leq \text{ord}(A[T])$.

Proof of Lemma 8. Let $\Sigma = \{c_1, \dots, c_{|\Sigma|}\}$. Take $A_{\mathbf{enum}} = \mathbf{Str}_\Sigma \rightarrow \mathbf{Str}_\Sigma$ and

$$\mathbf{enum} = \lambda x. x (\lambda f. \lambda s. \mathbf{conc}_{|\Sigma|+1} (f \overline{\#}) (f (\mathbf{conc} \overline{c_1} s)) \dots (f (\mathbf{conc} \overline{c_{|\Sigma|}} s))) (\lambda y. y) \overline{\#}$$

where $\mathbf{conc}_k : \mathbf{Str} \rightarrow \dots \rightarrow \mathbf{Str} \rightarrow \mathbf{Str}$ is the concatenation of k strings, which is safely λ -definable according to [BO09, Theorem 2.8], and $\mathbf{conc} = \mathbf{conc}_2$. \square

Proof of Lemma 9. Let $\Sigma = \{c_1, \dots, c_{|\Sigma|}\}$. Take $F_{\mathbf{exists}}(\alpha) = \mathbf{Str}_\Sigma[\alpha] \rightarrow \mathbf{Bool}$ and $\mathbf{exists} = \lambda p. \lambda s. s u_1 \dots u_{|\Sigma|} v p \overline{\varepsilon}$ where

$$u_i = \lambda f. \lambda x. f (\mathbf{conc} x \overline{c_i}) \quad u_{\#} = \lambda f. \lambda x. \mathbf{or} (p x) (f \overline{\varepsilon})$$

where $\mathbf{or} : \mathbf{Bool} \rightarrow \mathbf{Bool} \rightarrow \mathbf{Bool}$ can be safely defined following the recipe of [BO09, Remark 2.5(ii)]. (Explanation of this code: the accumulator of the “right fold” contains a function that takes the maximal $\#$ -free factor strictly before the current position, and tells you whether all blocks up to a certain point satisfy the predicate p .) Note that all instantiations $\alpha = A$ of this term are safe. \square

Proof of Lemma 10. By induction on the expression.

- We take $A_\emptyset = o$ and $t_\emptyset = \lambda s. \mathbf{false} : \mathbf{Str}_\Sigma \rightarrow \mathbf{Bool}$.
- In the unsafe λ -calculus, testing for the empty word could be done with type $\mathbf{Str}_\Sigma \rightarrow \mathbf{Bool}$, but here this is not possible anymore as discussed in [BO09, §2]. We use instead $A_\varepsilon = \mathbf{Bool}$ and

$$t_\varepsilon = \lambda s. s (\lambda x. \mathbf{false}) \dots (\lambda x. \mathbf{false}) \mathbf{true}$$

- To test whether the word contains a single letter, say, the first one in the alphabet Σ (call it c_1), we use $A_{c_1} = \mathbf{Bool} \rightarrow \mathbf{Bool} \rightarrow \mathbf{Bool}$ and

$$t_{c_1} = \lambda s. s (\lambda fxy. f \mathbf{true} (\mathbf{and} (\mathbf{not} x) y)) t' \dots t' \mathbf{and} \mathbf{false} \mathbf{true}$$

where $t' = \lambda fxy. f x \mathbf{false}$, and \mathbf{and} is defined from \mathbf{or} and \mathbf{not} .

- Complementation is implemented by post-composing with \mathbf{not} .

- To handle a union, we take $A_{E \cup F} = \mathbf{Str}[A_E] \rightarrow \mathbf{Str}[A_F] \rightarrow \mathbf{Bool}$ and

$$t_{E \cup F} = \lambda s. s (\lambda fxy. f (\mathbf{conc} \bar{c}_1 x) (\mathbf{conc} \bar{c}_1 y)) \dots (\lambda xy. \mathbf{or} (t_E x) (t_F y)) \bar{\varepsilon} \bar{\varepsilon}$$

Observe that t_E and t_F appear only once; otherwise, our construction would not run in polynomial time.

- The remaining case, concatenation, is the most delicate:

- First, we introduce a new symbol $\square \notin \Sigma$ and build a term of type $\mathbf{Str}_{\Sigma \cup \{\square\}}[\mathbf{Str}_{\Sigma}[A_E] \rightarrow (\mathbf{Str}_{\Sigma}[A_F] \rightarrow \mathbf{Bool}) \rightarrow \mathbf{Bool}] \rightarrow \mathbf{Bool}$ that distinguishes, among the words of $\Sigma^* \square \Sigma^*$, those that belong to $E \square F$ (we do not care what it computes on the rest of $(\Sigma \cup \{\square\})^*$):

$$t_{E \square F} = \lambda s. s v_{c_1} \dots v_{c_{|\Sigma|}} v_{\square} (\lambda x f. f \bar{\varepsilon}) \bar{\varepsilon} (\lambda y. \mathbf{false})$$

where

$$v_c = \lambda kxf. k (\mathbf{conc} x \bar{c}) (\lambda y. f (\mathbf{conc} \bar{c} y))$$

$$v_{\square} = \lambda kxf. k \underbrace{\bar{\varepsilon} (\lambda y. \mathbf{and} (t_E x) (t_F y))}_{\text{underbraced subterm}}$$

Note that the underbraced subterm has type $\mathbf{Str}_{\Sigma}[A_F] \rightarrow \mathbf{Bool}$ and contains a free variable of type $\mathbf{Str}_{\Sigma}[A_E]$. The safety condition tells us that we need to have $\text{ord}(A_E) \geq \text{ord}(A_F) + 1$. We can always make sure that we are in such a situation: by composing t_E m times with a safe λ -term $\mathbf{copy} : \mathbf{Str}_{\Sigma}[\mathbf{Str}_{\Sigma}] \rightarrow \mathbf{Str}_{\Sigma}$ which realizes the identity function on Σ^* (its existence is left to the reader...) we can get a term $t_E^{(m)} : \mathbf{Str}_{\Sigma}[A_E^{(m)}] \rightarrow \mathbf{Bool}$ with $\text{ord}(A_E^{(m)}) = \text{ord}(A_E) + 2m$ that recognizes the same language.

- Then, we show that $123 \mapsto \square 123 \# 1 \square 23 \# 12 \square 3 \# 123 \square : \Sigma^* \rightarrow \Gamma^*$ where $\Gamma = \Sigma \cup \{\square, \#\}$ (a typical *polyregular function*) is defined by $\mathbf{split} : \mathbf{Str}_{\Sigma}[\mathbf{Str}_{\Gamma}[\mathbf{Str}_{\Gamma}]] \rightarrow (\mathbf{Str}_{\Gamma} \rightarrow \mathbf{Str}_{\Gamma}) \rightarrow \mathbf{Str}_{\Gamma} \rightarrow \mathbf{Str}_{\Gamma}$:

$$\mathbf{split} = \lambda s. s v'_{c_1} \dots v'_{c_{|\Sigma|}} (\lambda x f. \mathbf{conc}_3 (f \bar{\#}) (\mathbf{copy} x) \bar{\square}) \bar{\varepsilon} (\lambda y. \bar{\varepsilon})$$

$$v'_c = \lambda kxf. k (\mathbf{conc} x \bar{c}) (\lambda y. \mathbf{conc}_4 (f (\mathbf{conc} \bar{c} y)) (\mathbf{copy} x) \bar{\square} c y)$$

Thanks to our use of $\mathbf{Str}_{\Gamma}[\mathbf{Str}_{\Gamma}]$ and \mathbf{copy} , this λ -term is safe.

- Finally, the term that we want is $\lambda s. \mathbf{exists} t_{E \square F} (\mathbf{split} s)$ (here we use \mathbf{exists} from Lemma 9). \square

References

- [BO09] William Blum and C.-H. Luke Ong. The Safe Lambda Calculus. *Logical Methods in Computer Science*, 5(1), February 2009.
- [Sch16] Sylvain Schmitz. Complexity hierarchies beyond elementary. *ACM Transactions on Computation Theory*, 8(1):3:1–3:36, 2016.