



HAL
open science

Simply typed convertibility is TOWER-complete even for safe lambda-terms

Lê Thành Dũng Nguyễn

► **To cite this version:**

Lê Thành Dũng Nguyễn. Simply typed convertibility is TOWER-complete even for safe lambda-terms. Logical Methods in Computer Science, 2024, 20 (3), pp.21:1-21:15. 10.46298/LMCS-20(3:21)2024 . hal-04688374

HAL Id: hal-04688374

<https://hal.science/hal-04688374>

Submitted on 5 Sep 2024


HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

SIMPLY TYPED CONVERTIBILITY IS TOWER-COMPLETE EVEN FOR SAFE LAMBDA-TERMS

LÊ THÀNH DŨNG (TITO) NGUYỄN 

Laboratoire de l’informatique du parallélisme (LIP), École normale supérieure de Lyon, France
e-mail address: nltld@nguyentito.eu

ABSTRACT. We consider the following decision problem: given two simply typed λ -terms, are they β -convertible? Equivalently, do they have the same normal form? It is famously non-elementary, but the precise complexity – namely TOWER-complete – is lesser known. One goal of this short paper is to popularize this fact.

Our original contribution is to show that the problem stays TOWER-complete when the two input terms belong to Blum and Ong’s safe λ -calculus, a fragment of the simply typed λ -calculus arising from the study of higher-order recursion schemes. Previously, the best known lower bound for this safe β -convertibility problem was PSPACE-hardness. Our proof proceeds by reduction from the star-free expression equivalence problem, taking inspiration from the author’s work with Pradic on “implicit automata in typed λ -calculi”.

These results also hold for $\beta\eta$ -convertibility.

1. INTRODUCTION

Consider the following *simply typed β -convertibility* problem:

- *Input*: two simply typed λ -terms t and u of the same type – this involves some choices of representation for such terms, that will be discussed later (§1.3).
- *Output*: does $t =_{\beta} u$ hold? Equivalently – since β -reduction is confluent and terminating in the simply typed λ -calculus – do t and u have the same normal form?

This is, of course, a fundamental decision problem concerning the λ -calculus. For example, the special case when $t : \mathbf{Bool} = o \rightarrow o \rightarrow o$ and $u = \mathbf{true} = \lambda x^o. \lambda y^o. x$ amounts to *evaluating* the result of the program t . (Here, o is the single base type that we work with: our grammar of simple types is $A, B ::= o \mid A \rightarrow B$.) Also, proof assistants based on dependent types often need to check whether two terms are *definitionally equal* – which is nothing more than the generalization of the above problem to more sophisticated type theories. For these reasons, there have been many works on deciding convertibility efficiently, such as [Con20].

Key words and phrases: non-elementary complexity, safe λ -calculus.

Thanks to Anupam Das, Damiano Mazza and Noam Zeilberger for many instructive discussions on the topic of complexity of normalization for λ -calculi.

The author was supported by the LambdaComb project (ANR-21-CE48-0017) while working at École Polytechnique and by the LABEX MILYON (ANR-10-LABX-0070) of Université de Lyon, within the program “France 2030” operated by the French National Research Agency (ANR).

To find out whether $t =_{\beta} u$, a naive approach is to normalize both t and u and then compare their normal forms. However, it is not always the best course of action: see, for instance, the practical remarks of [Asp17, Section 4] or the use of denotational semantics to decide convertibility in [Ter12].

1.1. Complexity of β -convertibility. A well-known result of Statman [Sta79] is that simply typed β -convertibility is *not* in the complexity class ELEMENTARY, whose definition we recall now. Let $2_0^n = n$ and $2_{k+1}^n = 2^{2_k^n}$. A function is in the class k -FEXPTIME when it can be computed in time bounded by $2_k^{P(\text{input size})}$ for some polynomial P ; and the class FELEMENTARY of *elementary recursive functions* is the union of all the k -FEXPTIME classes for $k \in \mathbb{N}$. By ELEMENTARY, we refer to the class of decision problems that, when seen as functions returning a boolean, belong to FELEMENTARY.

The statement of Statman’s theorem can be strengthened as follows. Let $\text{tower}(n) = 2_1^n$. The complexity class TOWER consists of the decision problems that can be solved in time $\text{tower}(2_k^{P(\text{input size})})$ for some polynomial P and some $k \in \mathbb{N}$. A problem is TOWER-hard when every problem in TOWER reduces to it via many-one FELEMENTARY reductions. By the time hierarchy theorem, a TOWER-hard problem cannot be in ELEMENTARY.

Theorem 1.1. *Simply typed β -convertibility is TOWER-complete (that is, it belongs to the class TOWER, and is at the same time TOWER-hard).*

The hardness part is actually implicit in Statman’s proof [Sta79]: it provides a reduction from a problem which is essentially the same as the TOWER-complete Higher-Order Quantified Boolean Satisfiability (HOSAT) problem. For a more detailed discussion of this point, see the “related work” paragraph of [CHHM22, Section 3]; let us also mention that Mairson [Mai92] has given a simpler reduction from HOSAT to simply typed β -convertibility. In fact, in his relatively recent article [Sch16] where the complexity class TOWER is explicitly introduced for the first time, Schmitz observes that many non-elementary lower bounds in the literature can be read as proofs of TOWER-hardness – and he mentions the example of simply typed β -convertibility [Sch16, §3.1.2].

As for membership in TOWER, it is known to be the “easy part” of Theorem 1.1. In Section 2, we recall some folklore ideas about normalization in the simply typed λ -calculus, and observe that they yield a TOWER algorithm for β -convertibility. Section 2 also discusses how the complexity of normalizing a λ -term is controlled by the *degree* and the *order* of the types that it contains – defined respectively as $\text{deg}(o) = \text{ord}(o) = 0$ and

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

In other words, the degree of a type is the height of its syntax tree, while the order is the nesting depth of function arrows to the left.

1.2. The case of the safe λ -calculus. Coincidentally, the order of a simple type is also central to the definition of *safe* fragment of the simply typed λ -calculus, even though the motivation that led Blum and Ong [BO09, Blu09] to introduce this fragment were unrelated to computational complexity. (We discuss this background in Section 3.5.2.)

Let us recall briefly the definition of the safe λ -calculus. Let t be a simply typed λ -term in “Church style” – i.e., the variables have type annotations ensuring that the type of each subterm of t is uniquely determined. Then t is *unsafe* if it contains a subterm $u : A$ that

- contains some free variable $x : B$ with $\text{ord}(B) < \text{ord}(A)$
- while being in “unsafe position”, that is:
 - either u equals t itself,
 - or u is in argument position: $t = C[v u]$ for some term v and one-hole context $C[\cdot]$,
 - or u is applied to some argument ($t = C[u v]$) but u itself is not an application.

A safe λ -term is, obviously, a simply typed λ -term that is not unsafe. For example:

$$\lambda f^{(o \rightarrow o) \rightarrow o}. f (\lambda x^o. f (\underbrace{\lambda y^o. x}_{\text{argument subterm of type } (o \rightarrow o) \rightsquigarrow \text{order } 1})) \text{ is unsafe} \quad \lambda f^{(o \rightarrow o) \rightarrow o}. f (\lambda x^o. f (\lambda y^o. y)) \text{ is safe}$$

free variable of type $o \rightsquigarrow$ order 0

The precise complexity of deciding β -convertibility on safe λ -terms has been an open problem until now. In their paper introducing the safe λ -calculus, Blum and Ong prove that it is PSPACE-hard [BO09, Section 3]. This lower bound is far below TOWER, but they argue that both Statman’s proof [Sta79] of the hardness part in Theorem 1.1 and its simplification by Mairson [Mai92] fundamentally require unsafe terms. To quote [BO09, §3.1], this “does not rule out the possibility that another non-elementary problem is encodable in the safe lambda calculus”. This is precisely what our original contribution here is about:

Theorem 1.2. *β -convertibility in the safe λ -calculus is TOWER-complete.*

Membership in TOWER follows, of course, from the “easy part” of Theorem 1.1, which is the topic of Section 2. We establish TOWER-hardness in Section 3 even for the special case of “homogeneous long-safe” λ -terms of type `Bool` (Theorem 3.1), by reduction from a standard problem in automata theory: *star-free expression equivalence*. As a corollary, this provides an alternative proof of the “hard part” of Theorem 1.1.

1.3. Some pedantic points. Traditionally, the simply typed λ -calculus (and its safe fragment) can be presented in two ways (cf. [KSW16]). We already mentioned the intrinsically typed “Church style”. By contrast, in “Curry style”, a λ -term is given without type annotations, and may satisfy several different typing judgments (e.g. $\vdash \lambda xy. x : A \rightarrow B \rightarrow A$ for all simple types A and B). We might therefore consider two extreme choices when specifying the input for the simply typed β -convertibility problem, that might *a priori* make a difference regarding computational complexity:

- (1) All variables in the input terms are fully annotated with their types, and these annotations are counted in the size of the input.
- (2) There are no types in the input at all: we are given two untyped terms t and u , with the promise that there exists some simple type A such that $\vdash t : A$ and $\vdash u : A$.

The mathematical equivalence of these two questions already requires some care. It is not the case in general for type theories that type erasure is injective modulo conversion: for example, in the polymorphic λ -calculus (System F), there are Church-style λ -terms of the same type that are not convertible even though the underlying untyped terms are β -convertible.¹ However, this does not happen in the simply typed λ -calculus, because:

Fact 1.3. A Church-style simply typed λ -term *in normal form*² can be reconstructed uniquely from its type and its underlying untyped term.

¹In System F, the latter condition is called “Strachey equivalence” [PA93, §2.2], see also [JT18].

²The normal form assumption is required to exclude counter-examples such as $(\lambda xy. y) (\lambda z. z) : A \rightarrow A$ where, even when a specific A is fixed, x can be given any type of the form $B \rightarrow B$.

Complexity-wise, (1) is easier than (2): simply erasing all types in the input gives a trivial reduction. Conversely, type inference for the simply typed λ -calculus can be performed in linear time, using first-order unification (see e.g. [Wan87]) – this also works for inferring a common simple type for a pair (t, u) of λ -terms, by adding an equation to the unification problem. A solution to this unification problem – of size $O(|t| + |u|)$ – also gives us the type annotations to add to t and u ; we can ask a unification algorithm to return a unifier that has linear size using a representation of the syntax trees as directed acyclic graphs (in order to share subtrees), cf. [BN98, Sections 4.8 and 4.9]. Note that this guarantees that $\text{ord}(A) \leq \text{deg}(A) = O(|t| + |u|)$ for any type A among the computed annotations.

Thus, there is a linear time reduction from (2) to the variant (1') of (1) where such shared representations are allowed for the types. In turn, (1') reduces to (1) in exponential time by unfolding the syntax trees of the types. Since $1\text{-FEXPTIME} \subset \text{FELEMENTARY}$, this does not make a difference regarding membership in or hardness for TOWER.

1.4. The η rule. Theorems 1.1 and 1.2 also hold for $\beta\eta$ -conversion.

For the hardness part, this is because we prove it for the problem restricted to λ -terms of type `Bool`, and $\beta\eta$ -convertibility coincides with β -convertibility at type `Bool`. (Indeed, for every $t : \text{Bool}$, either $t =_{\beta} \text{true}$ or $t =_{\beta} \text{false} = \lambda x^o. \lambda y^o. y$, and $\text{true} \neq_{\beta\eta} \text{false}$.)

For the complexity upper bound, note that Section 2 actually gives an algorithm for computing the β -normal form t' of a simply typed λ -term t . Once we have t' , applying η -reductions takes time $|t'|^{O(1)}$ until a $\beta\eta$ -normal form³ is reached. And testing whether two terms are $\beta\eta$ -convertible can be done by comparing their $\beta\eta$ -normal forms for equality⁴, thanks to the confluence of $\beta\eta$ -reduction on simply typed λ -terms of a fixed type – a fact shown in [Geu92] which is a bit more subtle than one could expect. This also means that we can invoke Fact 1.3 again to conclude that the Church-style and Curry-style variants of the $\beta\eta$ -convertibility problem are equivalent.

But this does not work anymore when we add sum types endowed with the “strong” η rule: in this setting, “there can be no λ -calculus with sums *à la* Curry” [MS15, §I(a)] because untyped $\beta\eta$ -conversion is inconsistent (that is, it equates all λ -terms). In fact, decidability of $\beta\eta$ -convertibility for a simply typed λ -calculus with sums and the empty type, using typed conversion rules, has only been proved relatively recently [Sch17]. As future work, it could be interesting to know whether this problem is still in TOWER.

2. SIMPLY TYPED β -CONVERTIBILITY IS IN TOWER

2.1. A simple TOWER normalization algorithm using the degree. Let us first define the *type of a redex* $(\lambda x^A. t) u$ in a Church-style simply typed λ -term as $A \rightarrow B$ where B is the type of t (so $A \rightarrow B$ is also the type of $\lambda x^A. t$). The *degree of a redex* is the degree of its type. The idea of using redex degrees to get a weak normalization proof for the simply typed λ -calculus goes back at least to Turing, see [BS23, Section 1] for an explanation and historical discussion. Here, we use a variation on this idea based on the notion of *parallel reduction*, which comes from the Tait/Martin-Löf proof of confluence for untyped β -reduction.

³Also known as a β -normal η -short form. In many other contexts, β -normal η -long forms are more useful, but here η -reduction is more convenient than η -expansion because it gives us a terminating rewriting system.

⁴In this paper, we always consider λ -terms modulo α -renaming, so an equality test that takes concrete syntax trees as input has to take α -renaming into account.

Definition 2.1 [Tak95, p. 121]. The parallel reduction t^* of a λ -term t is, inductively:

$$x^* = x \quad (\lambda x. t)^* = \lambda x. t^* \quad (t u)^* = \begin{cases} (t')^* \{x := u^*\} & \text{when } t = \lambda x. t' \text{ for some } t' \\ t^* u^* & \text{otherwise} \end{cases}$$

This definition on untyped λ -terms can be extended to Church-style simply typed λ -terms in the obvious way.

It is immediate that $t \longrightarrow_{\beta}^* t^*$. Furthermore:

Proposition 2.2 [AL13, Appendix]. *Let t be a simply typed λ -term whose redexes have degree at most $d \in \mathbb{N}$. Then:*

- The redexes of t^* have degree at most $d - 1$.
- As a consequence, the normal form of t can be reached in d parallel reductions. In other words, let $t^{*0} = t$ and $t^{*(n+1)} = (t^{*n})^*$; then t^{*d} is the normal form of t .

This gives us a simple algorithm to compute the normal form of a simply typed λ -term. In fact, since parallel reduction does not inspect types, we may also apply it to simply *typable* terms, to solve the Curry-style version of the β -convertibility problem (cf. §1.3).

We consider the following measures on an untyped λ -term t to get a complexity bound:

- $\text{height}(t)$ denotes the height of its syntax tree;
- its *size* $|t|$ is the number of nodes of its syntax tree, including the leaves (variables):

$$|x| = 1 \quad |\lambda x. t| = 1 + |t| \quad |t u| = 1 + |t| + |u|$$

Fact 2.3. For any untyped λ -term t , we have $\text{height}(t^*) \leq |t| \leq 2^{\text{height}(t)}$.

(We have not managed to find this statement in the literature, but it is probably not new.)

Proof. The first inequality can be proved by structural induction on Definition 2.1, using

$$\text{height}((t')^* \{x := u^*\}) \leq \text{height}((t')^*) + \text{height}(u^*)$$

while the second inequality is due to the fact that each node in the syntax tree of a λ -term has at most two children. \square

As a consequence, $|t^*| \leq 2^{|t|}$ for any λ -term t . Putting all this together, we get:

Theorem 2.4. *There is an algorithm that takes as input a simply typable λ -term t and returns its normal form in time at most $2_d^{P(|t|)}$, where d is the maximum redex degree in some typing of t (not given as input) and P is some polynomial independent from t and d .*

Proof. The algorithm iterates parallel reductions until it reaches a normal form. This takes at most d steps; for each step $n \in \{1, \dots, d\}$, computing t^{*n} can be done by feeding the input $t^{*(n-1)}$ – which has size at most $2_{n-1}^{|t|}$ – to a naive 1-FEXPTIME procedure. \square

Given two simply typable λ -terms, we can compute their normal forms and compare them to decide whether they are β -convertible. Together with the considerations from §1.3, this establishes the “membership in TOWER” part of Theorem 1.1. (Again, we do not make any claim to originality concerning the material in this section; it has been included for expository purposes.)

2.2. Finer bounds with the order. Actually, a TOWER algorithm for normalization could also be obtained just by applying successive (non-parallel) β -reductions. The complexity then depends on the length of reduction sequences for simply typed λ -terms, for which bounds are known [Sch91, Bec01]. But note that these bounds are stated using the *order* of the types in a λ -term, not their degree!

Remark 2.5. The paper [Bec01] uses different terminology: “level” refers to what we call the order, while “degree” refers to the maximum order of the types appearing in a term. See also [AKST19] for a probabilistic analysis of β -reduction length in the simply typed λ -calculus, and [Sin22, Section 4.4] for something similar in the linear λ -calculus.

For any simple type A , we have $\text{ord}(A) \leq \text{deg}(A)$ and the gap can be arbitrarily large:

$$\text{deg}(\underbrace{o \rightarrow \cdots \rightarrow o}_{n \text{ times}} \rightarrow o) = n \quad \text{ord}(\underbrace{o \rightarrow \cdots \rightarrow o}_{n \text{ times}} \rightarrow o) = 1$$

While using the degree seems to be the easier way to devise a normalization algorithm together with a proof of a TOWER complexity bound, it seems that the order is the parameter that truly controls complexity. As further evidence towards this (cf. [Ngu21, §1.3.6] for a longer discussion of these three items):

- the normalization problem for simply typed λ -terms of type `Bool` containing only subterms whose types have *order at most* $2k + 2$ is k -EXPTIME-complete [Ter12];
- in fact, k -EXPTIME can be characterized as the predicates expressed by simply typed λ -terms of *order at most* $2k + 4$ using a certain input-output convention [HK96], refining an earlier characterization of ELEMENTARY in the simply typed λ -calculus with no constraint on the order [HKM96];
- the “call-by-name translation” ($A \rightarrow B = !A \multimap B$) from the simply typed λ -calculus to propositional linear logic (LL) maps the order to the nesting depth of the exponential modality ‘!’, and ‘!’ is generally considered to be the main source of complexity in LL (see [GRV09] for a connection between the exponential depth in propositional LL and the kind of stratification used in LL-based implicit computational complexity).

3. SAFE β -CONVERTIBILITY IS TOWER-HARD

As announced in Section 1.2 of the introduction, our goal here is to show:

Theorem 3.1. *Given a homogeneous long-safe λ -term t of type `Bool`, it is TOWER-hard to decide whether $t =_{\beta} \text{true}$.*

Recall that the type `Bool` is defined as $o \rightarrow o \rightarrow o$. It has exactly two inhabitants in normal form: `true` = $\lambda xy. x$ and `false` = $\lambda xy. y$. They are both safe [BO09, Remark 2.5(ii)], and even homogeneous long-safe (as defined in the next subsection).

3.1. Homogeneity and long-safety. Before proceeding with the proof of Theorem 3.1, let us explain the additional restrictions – beyond the safety condition recalled in Section 1.2 – imposed on the input terms. We chose to prove hardness in presence of these restrictions because they are natural in the context of higher-order recursion schemes (cf. §3.5.2).

First, *homogeneity* is a simple syntactic criterion on types.

Definition 3.2. A simple type $A_1 \rightarrow \cdots \rightarrow A_n \rightarrow o$ (for $n \geq 0$) is homogeneous when each of the A_i is itself homogeneous ($i \in \{1, \dots, n\}$) and $\text{ord}(A_1) \geq \cdots \geq \text{ord}(A_n)$.

The second restriction is *long-safety*. In Church style, it is the same as the notion of safety in §1.2 except more subterms of a simply typed λ -term are considered to be in “unsafe position”: if u is not a λ -abstraction, then it is in unsafe position in $t = C[\lambda x. u]$. Equivalently, in Curry style:

Definition 3.3 [BO09, Def. 1.21]. The typing rules of the long-safe λ -calculus are

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

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

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

where $\text{inford}(\Theta) = \inf_{(y:C) \in \Theta} \text{ord}(C)$ with the usual convention $\text{inf}(\emptyset) = +\infty$.

For the rest of Section 3, we only work with Curry-style typed λ -terms. This is because it allows our constructions to be carried out in polynomial time. If we had to build Church-style terms, writing out the type annotations would take exponential time, as explained in §1.3 (though this is harmless in a TOWER-hardness proof; the point is just to have a more precise statement).

Definition 3.4. For a λ -term t , a simple type A and a context Γ , we write $\Gamma \vdash_{\text{hls}} t : A$ when there is a derivation *using only homogeneous types* of $\Gamma \vdash t : A$ in the above type system. In that case we say that t is *homogeneous long-safe*.

3.2. Star-free expressions. As we said in the introduction, our TOWER-hardness proof proceeds by a reduction from an automata-theoretic problem, which we recall now.

Definition 3.5. A *star-free expression* over a finite alphabet Σ is a regular expression without the Kleene star, but with complementation, defining a language $\llbracket E \rrbracket \subseteq \Sigma^*$:

$$E, E' ::= \emptyset \mid \overbrace{\varepsilon}^{\text{empty string}} \mid \underbrace{a}_{\text{letter } a \in \Sigma} \mid E \cup E' \mid \overbrace{E \cdot E'}^{\text{concatenation}} \mid \underbrace{\neg E}_{\text{complement}}$$

The *star-free equivalence problem* consists in deciding whether two star-free expressions (over the same alphabet) denote the same language: given E and E' , does $\llbracket E \rrbracket = \llbracket E' \rrbracket$ hold?

For instance, over the alphabet $\Sigma = \{a, b, c\}$, we have the equivalence

$$\llbracket a \cdot \neg \emptyset \cup b \cdot \neg \emptyset \rrbracket = \llbracket \neg(\varepsilon \cup c \cdot \neg \emptyset) \rrbracket = \{w \in \Sigma^* \mid w \text{ starts with an } a \text{ or a } b\}$$

Schmitz’s paper introducing TOWER presents star-free equivalence as a typical example of a TOWER-complete problem [Sch16, §3.1] (rephrasing a hardness result proved by Meyer and Stockmeyer, cf. [Sto74, §4.2]). Furthermore, equivalence reduces to emptiness: given two expressions E and F , the language denoted by $\neg(\neg E \cup F) \cup \neg(E \cup \neg F)$ is empty if and only if E and F are equivalent. And to test whether $\llbracket E \rrbracket = \emptyset$, it is well known that it suffices to examine words of length at most $\text{tower}(|E| - 1)$:

Proposition 3.6 (cf. e.g. [Sto74, proof of Prop. 4.25]). *If a star-free expression of size $n + 1$ denotes a non-empty language L , then L contains a word of length at most $\text{tower}(n)$.*

(Further remarks on the above proposition and on the complexity of star-free equivalence are given at the end in Section 3.5.1.) Altogether, this discussion leads us to conclude that:

Lemma 3.7. *The following problem is TOWER-hard:*

- Input: a star-free expression E and a natural number n given in unary.
- Output: whether or not $\llbracket E \rrbracket$ contains a word of length at most $\text{tower}(n)$.

3.3. From star-free expressions to safe λ -terms. We shall represent (star-free) languages as functions from strings to booleans. We already have a coding of booleans, that is involved in the statement of Theorem 3.1. For strings, we use the standard *Church encodings*.

Definition 3.8 [BB85, §4]. The type $\mathbf{Str}_\Sigma = \overbrace{(o \rightarrow o) \rightarrow \cdots \rightarrow (o \rightarrow o)}^{|\Sigma| \text{ times}} \rightarrow o \rightarrow o$ is the type of so-called *Church encodings* of strings over the finite alphabet Σ . In the case of a unary alphabet, the type of *Church numerals* is $\mathbf{Nat} = \mathbf{Str}_{\{1\}} = (o \rightarrow o) \rightarrow o \rightarrow o$.

The Church encoding of a word w over an ordered alphabet $\Sigma = \{c_1, \dots, c_n\}$ (resp. of a number $n \in \mathbb{N}$) is $\bar{w} = \lambda f_{c_1} \dots f_{c_n} x. \underbrace{f_{w[1]} (\dots (f_{w[n]} x))}_{w[i] \text{ refers to the } i\text{-th letter of } w} : \mathbf{Str}_\Sigma$ (resp. $\bar{n} = \underbrace{\overline{1 \dots 1}}_{n \text{ times}} : \mathbf{Nat}$).

For any finite alphabet Σ and string $w \in \Sigma^*$, one can check that the type \mathbf{Str}_Σ is homogeneous and that $\vdash_{\text{hls}} \bar{w} : \mathbf{Str}_\Sigma$ (the safety of \bar{w} is already used in [BO09, §2.2]). Next, to encode languages, a fundamental remark is that if A and B are homogeneous simple types and $\vdash_{\text{hls}} t : A$, then $\vdash_{\text{hls}} t : A[B]$, where:

Notation 3.9. $A[B]$ denotes the substitution of all occurrences of the base type o in the simple type A by the simple type B .

This allows us to use a λ -term t such that $\vdash_{\text{hls}} t : \mathbf{Str}_\Sigma[B] \rightarrow \mathbf{Bool}$ to represent the language $\{w \in \Sigma^* \mid t \bar{w} =_\beta \mathbf{true}\}$ – this makes sense since $\vdash_{\text{hls}} t \bar{w} : \mathbf{Bool}$. The key lemma in our proof of Theorem 3.1 is that star-free expressions can be efficiently translated to such representations by (Curry-style) homogeneous safe λ -terms.

Lemma 3.10. *A star-free expression E over a finite alphabet Σ can be turned in polynomial time into an λ -term t_E such that*

- $\vdash_{\text{hls}} t_E : \mathbf{Str}_\Sigma[A_E] \rightarrow \mathbf{Bool}$ for some simple type A_E ;
- for any $w \in \Sigma^*$, we have $t_E \bar{w} =_\beta \mathbf{true}$ if and only if $w \in \llbracket E \rrbracket$.

(The type A_E is not part of the output of the polynomial-time algorithm taking E and Σ as input. In fact, the size of A_E may be exponential in the size of E .)

To prove Lemma 3.10, we first note that boolean operations can be implemented by

$$\mathbf{and} = \lambda b_1 b_2 x y. b_1 (b_2 x y) y \quad \mathbf{or} = \lambda b_1 b_2 x y. b_1 x (b_2 x y) \quad \mathbf{not} = \lambda b x y. b y x$$

These λ -terms satisfy $\vdash_{\text{hls}} \mathbf{and}, \mathbf{or} : \mathbf{Bool} \rightarrow \mathbf{Bool} \rightarrow \mathbf{Bool}$ and $\vdash_{\text{hls}} \mathbf{not} : \mathbf{Bool} \rightarrow \mathbf{Bool}$. Furthermore, we can implement concatenation on strings (cf. [BO09, Theorem 2.8]) with

$$\mathbf{cat} = \lambda s_1 s_2 f_1 \dots f_{|\Sigma|} x. s_1 f_1 \dots f_{|\Sigma|} (s_2 f_1 \dots f_{|\Sigma|} x)$$

which is a homogeneous long-safe λ -term of type $\mathbf{Str}_\Sigma \rightarrow \mathbf{Str}_\Sigma \rightarrow \mathbf{Str}_\Sigma$. This generalizes to k -fold concatenation: $\vdash_{\text{hls}} \mathbf{cat}_k : \underbrace{\mathbf{Str}_\Sigma \rightarrow \cdots \rightarrow \mathbf{Str}_\Sigma}_{k \text{ times}} \rightarrow \mathbf{Str}_\Sigma$.

The next step in our proof is the following lemma, which is reused in Section 3.4.

Lemma 3.11. *For every finite alphabet Σ , one can build in time $|\Sigma|^{O(1)}$ a λ -term \mathbf{any}_Σ such that, for any homogeneous type A , there is some $F_{\mathbf{any}}(A)$ such that*

$$\vdash_{\text{hls}} \mathbf{any}_\Sigma : \mathbf{Str}_{\Sigma \cup \{\#\}}[F_{\mathbf{any}}(A)] \rightarrow (\mathbf{Str}_\Sigma[A] \rightarrow \mathbf{Bool}) \rightarrow \mathbf{Bool}$$

and for every list of words $w_0, \dots, w_n \in \Sigma^*$ and every λ -term $t : \mathbf{Str}_\Sigma[A] \rightarrow \mathbf{Bool}$,

$$\mathbf{any}_\Sigma \overline{w_0 \# \dots \# w_n} t =_\beta \mathbf{true} \iff \exists i \in \{0, \dots, n\}. t \overline{w_i} =_\beta \mathbf{true}$$

Proof. Let $\Sigma = \{c_1, \dots, c_{|\Sigma|}\}$. Take $F_{\mathbf{any}}(A) = \mathbf{Str}_\Sigma[A] \rightarrow \mathbf{Bool}$ and define \mathbf{any}_Σ as

$$\lambda s. s \overbrace{u_1 \dots u_{|\Sigma|}}^{\text{correspond to letters in } \Sigma} \underbrace{(\lambda f x. \text{or } (p x) (f \bar{\varepsilon})) p \bar{\varepsilon}}_{\text{corresponds to } \#} \quad \text{where } u_i = \lambda f x. f (\text{cat } x \bar{c}_i)$$

with $x, \bar{\varepsilon} : \mathbf{Str}_\Sigma[A]$, $p, f : \mathbf{Str}_\Sigma[A] \rightarrow \mathbf{Bool}$, $u_i : (\mathbf{Str}_\Sigma[A] \rightarrow \mathbf{Bool}) \rightarrow \mathbf{Str}_\Sigma[A] \rightarrow \mathbf{Bool}$ and $s : \mathbf{Str}_{\Sigma \cup \{\#\}}[\mathbf{Str}_\Sigma[A] \rightarrow \mathbf{Bool}]$ in the typing derivation for \mathbf{any}_Σ .

Explanation: computing $\mathbf{any}_\Sigma \overline{w_0 \# \dots \# w_n} t$ performs a “right fold” over the input with an accumulator of type $\mathbf{Str}_\Sigma[A] \rightarrow \mathbf{Bool}$, representing a language $L(\sigma) \subseteq \Sigma^*$ where σ is the suffix processed up until this point of the computation. The idea is that

$$u \in L(v \# w_m \# \dots \# w_n) \iff t \overline{uv} =_\beta \mathbf{true} \vee \exists i \geq m. t \overline{w_i} =_\beta \mathbf{true}$$

This makes the condition that we want to check equivalent to $\varepsilon \in L(w_0 \# \dots \# w_n)$. \square

Proof of Lemma 3.10. By induction on the expression, leading to a recursive algorithm. We also want the algorithm to return $\text{ord}(A_E)$ on the input E , because this information on subexpressions is used in two cases to make sure that the term we build admits a homogeneous type. It will be clear that in every case, $\text{ord}(A_E)$ can be recursively computed and is bounded by $O(|E|)$.

- We take $A_\emptyset = o$ and $t_\emptyset = \lambda s. \mathbf{false} : \mathbf{Str}_\Sigma \rightarrow \mathbf{Bool}$.
- In the simply typed λ -calculus, testing for the empty word can be done with type $\mathbf{Str}_\Sigma \rightarrow \mathbf{Bool}$, but this is not possible with safe λ -terms as discussed in [BO09, Section 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 consists of a single letter, say, the first one in the alphabet Σ (call it c_1), we use $A_{c_1} = o \rightarrow o \rightarrow o \rightarrow o$ and define t_{c_1} to be a λ -term corresponding to a deterministic finite automaton that recognizes the language $\{c_1\}$, namely:

$$\begin{array}{l} \text{represents the transition } \delta(c_1, q_0)=q_1, \delta(c_1, q_1)=\delta(c_1, q_2)=q_2 \quad \text{accepting states} = \{q_1\} \\ \lambda sxy. s \underbrace{(\lambda qz_0z_1z_2. q z_1 z_2 z_2)}_{|\Sigma|-1 \text{ times}} \underbrace{(\lambda q. q_2) \dots (\lambda q. q_2)}_{\text{the 3 states of the DFA (of type } A_{c_1})} q_0 \overbrace{y x y} \quad \text{where } \underbrace{q_i = \lambda z_0z_1z_2. z_i} \end{array}$$

- Complementation is trivially implemented by post-composing with \mathbf{not} .
- To handle a union $E \cup F$, we first check whether $\text{ord}(A_E) \geq \text{ord}(A_F)$. If that is the case, we take $A_{E \cup F} = \mathbf{Str}_\Sigma[A_E] \rightarrow \mathbf{Str}_\Sigma[A_F] \rightarrow \mathbf{Bool}$ and

$$t_{E \cup F} = \lambda s. s \underbrace{\dots (\lambda kxy. k (\text{cat } x \bar{c}_i) (\text{cat } y \bar{c}_i)) \dots (\lambda xy. \text{or } (t_E x) (t_F y)) \bar{\varepsilon} \bar{\varepsilon}}_{\text{for each } c_i \in \Sigma}$$

with $k : \mathbf{Str}_\Sigma[A_E] \rightarrow \mathbf{Str}_\Sigma[A_F] \rightarrow \mathbf{Bool}$, $x : \mathbf{Str}_\Sigma[A_E]$ and $y : \mathbf{Str}_\Sigma[A_F]$ in the typing derivation; the two occurrences of $\bar{\varepsilon}$ at the end must be given the different types $\mathbf{Str}_\Sigma[A_E]$ and $\mathbf{Str}_\Sigma[A_F]$ respectively.

This can be seen as a continuation-passing-style transformation of the following procedure using pair types (which we do not have here): perform a “left fold” over the input string with an accumulator of type $\mathbf{Str}_\Sigma[A_E] \times \mathbf{Str}_\Sigma[A_F]$; for each input letter c , apply the function $(w, w') \in \Sigma^* \mapsto (wc, w'c)$, so that the result of the fold is two copies of the input string with different types; finally, check on (w, w') whether $w \in \llbracket E \rrbracket$ or $w' \in \llbracket F \rrbracket$.

When $\text{ord}(A_E) < \text{ord}(A_F)$, we take $t_{E \cup F} = t_{F \cup E}$. Observe that t_E and t_F appear only once in $t_{E \cup F}$ – this is true in every recursive case, ensuring that t_E has size $O(|E|)$.

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

First, we introduce a new letter $\square \notin \Sigma$ and build a λ -term $t'_{E,F}$ that computes the language $\llbracket (E\square)^*F \rrbracket$. When $\text{ord}(A_E) \geq \text{ord}(A_F)$, we get

$$\vdash_{\text{hls}} t'_{E,F} : \mathbf{Str}_{\Sigma \cup \{\square\}}[\mathbf{Str}_\Sigma[A_E] \rightarrow \mathbf{Str}_\Sigma[A_F] \rightarrow \text{Bool}] \rightarrow \text{Bool}$$

by taking $t'_{E,F} = \lambda s. s v_{c_1} \dots v_{c_{|\Sigma|}} v_\square u \bar{\varepsilon} \bar{\varepsilon}$ (for $\Sigma = \{c_1, \dots, c_{|\Sigma|}\}$) where

$$v_c = \lambda fxy. f (\text{cat } x \bar{c}) (\text{cat } y \bar{c}) \quad v_\square = \lambda fxy. \text{and } (t_E x) (f \bar{\varepsilon} \bar{\varepsilon}) \quad u = \lambda xy. t_F y$$

Similarly to the proof of Lemma 3.11, $t'_{E,F}$ performs a “right fold” on its input string (over the alphabet $\Sigma \cup \{\square\}$), whose accumulator $f : \mathbf{Str}_\Sigma[A_E] \rightarrow \mathbf{Str}_\Sigma[A_F] \rightarrow \text{Bool}$ represents a language (over Σ) – except that, analogously to the case $E \cup F$ above, this representation must be fed two copies $x : \mathbf{Str}_\Sigma[A_E]$ and $y : \mathbf{Str}_\Sigma[A_F]$ of the same string. We get

$$\{w \in \Sigma^* \mid w\sigma \in \llbracket (E\square)^*F \rrbracket\}$$

after reading the input suffix $\sigma \in (\Sigma \cup \{\square\})^*$; therefore, once the whole input string has been traversed, we just need to check that the final accumulator contains ε .

If $\text{ord}(A_E) < \text{ord}(A_F)$, to ensure homogeneity, we replace in the definition of $t'_{E,F}$

$$v_\square \text{ by } \lambda kxy. \text{and } (t_E y) (k \bar{\varepsilon} \bar{\varepsilon}) \quad \text{and } u \text{ by } \lambda xy. t_F x$$

leading to $\vdash_{\text{hls}} t'_{E,F} : \mathbf{Str}_{\Sigma \cup \{\square\}}[\mathbf{Str}_\Sigma[A_F] \rightarrow \mathbf{Str}_\Sigma[A_E] \rightarrow \text{Bool}] \rightarrow \text{Bool}$.

Then, the λ -term that we want is $t_{E \cdot F} = \lambda s. \overbrace{\text{any}_{\Sigma \cup \{\square\}}}$ $\overbrace{(\text{split}_\Sigma s)}$ $t'_{E,F}$ where the role of $t'_{E,F}$ is to distinguish, among the strings containing exactly one ‘ \square ’, those that belong to $\llbracket E\square F \rrbracket$ – indeed, $\llbracket (E\square)^*F \rrbracket \cap \Sigma^*\square\Sigma^* = \llbracket E\square F \rrbracket$. \square

We isolate the following part of the above proof to serve as an example of independent interest (cf. §3.5.3) of a string-to-string function expressed in the safe λ -calculus.

Lemma 3.12. *For every finite alphabet Σ , one can build in time $|\Sigma|^{O(1)}$ a λ -term split_Σ such that, for some type A_{split} ,*

$$\vdash_{\text{hls}} \text{split}_\Sigma : \mathbf{Str}_\Sigma[A_{\text{split}}] \rightarrow \mathbf{Str}_\Gamma \quad \text{where } \Gamma = \Sigma \cup \{\square, \#\}$$

and for every list of letters $a_1, \dots, a_n \in \Sigma$,

$$\text{split}_\Sigma \overline{a_1 \dots a_n} =_\beta \overline{\square a_1 \dots a_n \# a_1 \square a_2 \dots a_n \# \dots \# a_1 \dots a_{n-1} \square a_n \# a_1 \dots a_n \square}$$

Proof. As in the case of unions in the proof of Lemma 3.10, the idea is to start from a left-to-right procedure using a pair type, and then apply a continuation-passing-style transformation. For $p = a_1 \dots a_m \in \Sigma^*$ (where $a_i \in \Sigma$), let us consider the maps

$$\begin{array}{ll} X_p : \mathbb{N} \rightarrow \Sigma^* & F_{a_1 \dots a_m} : \Gamma^* \rightarrow \Gamma^* \\ n \mapsto p & y \mapsto \square a_1 \dots a_m y a_1 \square a_2 \dots a_m y \dots y a_1 \dots a_{m-1} \square a_m y \end{array}$$

In particular, $F_\varepsilon(y) = \varepsilon$ and $F_a(y) = \square ay$ for $a \in \Sigma$. The maps X_p are constant; the role of their dummy argument is to increase the order of an associated variable in the λ -term defined below so that it satisfies the safety condition – $\text{ord}(\text{Nat} \rightarrow \text{Str}_\Gamma) = 3 > 2 = \text{ord}(\text{Str}_\Gamma)$.

For $p \in \Sigma^*$ and $c \in \Sigma$, we have the inductive equations

$$X_{pc}(n) = X_p(0) \cdot c \quad F_{pc}(y) = F_p(cy) \cdot X_p(0) \cdot \square c \cdot y$$

which we translate into the λ -term

$$v_c = \lambda k x f. k (\lambda n. \text{cat} (x \bar{0}) \bar{c}) \underbrace{(\lambda y. \text{cat}_4 (f (\text{cat} \bar{c} y)) (x \bar{0}) \overline{\square c} y)}_{x \text{ appears free in this subterm of order } 3 \rightsquigarrow \text{the } \text{Nat} \rightarrow \text{Str}_\Gamma \text{ trick ensures (long-)safety}}$$

with $k : (\text{Nat} \rightarrow \text{Str}_\Gamma) \rightarrow (\text{Str}_\Gamma \rightarrow \text{Str}_\Gamma) \rightarrow \text{Str}_\Gamma$, $x : \text{Nat} \rightarrow \text{Str}_\Gamma$, $f : \text{Str}_\Gamma \rightarrow \text{Str}_\Gamma$, $n : \text{Nat}$ and $y : \text{Str}_\Gamma$. Finally, we take

$$\text{split} = \lambda s. s v_{c_1} \dots v_{c_{|\Sigma|}} (\lambda x f. \text{cat}_3 (f \overline{\#}) (x \bar{0}) \overline{\square}) (\lambda n. \bar{\varepsilon}) (\lambda y. \bar{\varepsilon})$$

– thus, A_{split} is the aforementioned type of k . \square

3.4. Proof of Theorem 3.1. We now have all the ingredients to reduce the TOWER-hard decision problem of Lemma 3.7 to safe β -convertibility. This reduction is performed by the following lemma, which therefore immediately entails our main theorem.

Lemma 3.13. *Given a star-free expression E over a finite alphabet Σ , and $n \in \mathbb{N}$ written in unary, one can build in polynomial time a λ -term b_E such that $\vdash_{\text{hls}} b_E : \text{Bool}$ and*

$$b_E =_\beta \text{true} \iff [E] \cap \Sigma^{\leq \text{tower}(n)} \neq \emptyset$$

Proof. The key is to build a λ -term enum_Σ such that:

- $\vdash_{\text{hls}} \text{enum}_\Sigma : \text{Nat}[A_{\text{enum}}] \rightarrow \text{Str}_{\Sigma \cup \{\#\}}$ for some type A_{enum} ;
- for any $m \in \mathbb{N}$, the application $\text{enum} \bar{m}$ reduces to the Church encoding of a $\#$ -separated list containing all words in Σ^* of length at most m .

We define it for $\Sigma = \{c_1, \dots, c_{|\Sigma|}\}$ in such a way that $\text{enum}_\Sigma \bar{N} =_\beta \overline{f_N(\#)}$ for $N \in \mathbb{N}$, where the sequence of functions $f_n : (\Sigma \cup \{\#\})^* \rightarrow (\Sigma \cup \{\#\})^*$ is inductively defined by

$$f_0(x) = x \quad f_{n+1}(x) = x \cdot f_n(c_1 x) \cdot \dots \cdot f_n(c_{|\Sigma|} x)$$

e.g. for $\Sigma = \{0, 1\}$ we have $f_1(x) = x0x1x$ and $f_2(x) = x0x00x10x1x01x11x$. Observe that the string $f_2(\#)$ is a $\#$ -separated representation of the list $[\varepsilon, 0, 00, 10, 1, 01, 11, \varepsilon]$ that indeed contains all words of length at most 2 over Σ (with one redundancy). To implement this, take $A_{\text{enum}} = \text{Str}_\Sigma \rightarrow \text{Str}_\Sigma$ and

$$\text{enum}_\Sigma = \lambda n. n (\lambda f s. \text{cat}_{|\Sigma|+1} s (f (\text{cat} \bar{c}_1 s)) \dots (f (\text{cat} \bar{c}_{|\Sigma|} s))) (\lambda y. y) \overline{\#}$$

Once we have that, we can take $b_E = \underbrace{\text{any}_\Sigma}_{\text{Lemma 3.11}} (\text{enum}_\Sigma \text{tow}_n) \underbrace{t_E}_{\text{Lemma 3.10}}$ where $\text{tow}_n = \overbrace{\bar{2} \dots \bar{2}}^{n \text{ times}}$ satisfies $\vdash_{\text{hls}} \text{tow}_n : \text{Nat}$ and is β -convertible to $\text{tower}(n)$ (cf. [BO09, Remark 3.5(iii)]). \square

3.5. Final remarks. To conclude this paper, we discuss some topics related to Theorem 3.1 and its proof.

3.5.1. *On the complexity of star-free expression equivalence.* The equivalence problem for usual regular expressions, with Kleene star but no complementation, is “merely” PSPACE-complete [Sto74, Theorem 4.13] – see also [HK11, Theorem 15]. One could therefore think that the main source of complexity in the case of star-free expressions is *complementation*.

This can be witnessed in the proof idea for Proposition 3.6. Translate E to an equivalent nondeterministic finite automaton (NFA), whose number of states bounds the length of a shortest word in the language (since the shortest accepting runs visit each state at most once). This can be done by induction on E , using any standard construction on NFA for union and concatenation; the costliest operation is complementation, which uses determinization, inducing a single exponential state blowup.

But in our translation from star-free expressions to λ -terms (Lemma 3.10), it turns out that complementation is trivial while concatenation increases the order of the types. Relatedly, in *deterministic* finite automata, complementation can be done without changing the set of states, while concatenation may need an exponential blowup of the number of states. So the complexity of the problem is in fact rooted in the *alternation* between complementation and concatenation – which leads to the “dot-depth hierarchy” [Pin17].

3.5.2. *Safety in higher-order recursion schemes.* Blum and Ong’s main motivation for introducing the safe λ -calculus [BO09] was unrelated to complexity: they wanted to transpose to the λ -calculus the notion of safety on higher-order recursion schemes (HORS) [KNU02] – and indeed, the correspondence between HORS and $\lambda\mathbf{Y}$ -terms (i.e. terms in the simply typed λ -calculus extended with fixed-point combinators $\mathbf{Y}_A : (A \rightarrow A) \rightarrow A$) relates safe HORS with safe $\lambda\mathbf{Y}$ -terms [SW16]. Even without \mathbf{Y} , the safe λ -calculus has some interesting properties: for example, when substituting a safe λ -term into another, no spurious variable capture can happen [BO09, Lemma 1.10] – this leads to a way to normalize long-safe λ -terms by rewriting without α -renaming, but it is a bit subtle, see [FMvO23, Section 5.2]. A universal algebra perspective on the safe λ -calculus is also sketched in [Sal15, §2.3].

The homogeneity condition also comes from the study of HORS, cf. [Par18]. In fact, the homogeneous long-safe λ -calculus is equivalent (cf. [Blu09, Remark 3.53]) to an earlier attempt [dM06, Section 2.4.2] to design a safe λ -calculus inspired by HORS.

3.5.3. *Implicit automata in typed λ -calculi and transducer theory.* Lemma 3.10 tells us that every language of the form $\llbracket E \rrbracket$ for some star-free expression can be expressed by some homogeneous long-safe λ -term of type $\text{Str}_\Sigma[A] \rightarrow \text{Bool}$. These *star-free languages* are a fundamental and well-studied (cf. [Str18]) subclass of *regular languages*. Hillebrand and Kanellakis have shown that the languages computed by simply typed λ -terms of type $\text{Str}_\Sigma[A] \rightarrow \text{Bool}$ are, in fact, exactly the regular languages [HK96, Theorem 3.4]⁵; they have a translation of deterministic finite automata (rather than regular expressions) into the simply typed λ -calculus that actually produces safe λ -terms. In fact, we used this translation of DFA in the proof of Lemma 3.10, for the “single letter” case of the induction.

By using an affine and non-commutative type system, one can get a variant of Hillebrand and Kanellakis’s theorem characterizing star-free languages instead [NP20].⁶ The proof does not translate star-free expressions; instead, to encode star-free languages, it goes through

⁵This paper [HK96] is the same that was already mentioned in §2.2.

⁶See also [Ngu21, Chapter 7] for a variant of this result (with linear instead of affine types).

an algebraic characterization. Nevertheless, these works were the main inspirations for our proof strategy for Theorem 3.1.

Further works [Ngu21, PP24] in this “implicit automata in typed λ -calculi” research programme have focused on characterizations of string-to-string (or tree-to-tree) functions computed by *transducers*, i.e. automata with output. For reasons that are close to the aforementioned connection with HORS, as discussed in [Ngu21, §1.4.1], definability in the safe λ -calculus characterizes an important class of functions from 1980s transducer theory. The functions defined by the safe λ -terms split_Σ and enum_Σ (from Lemmas 3.12 and 3.13 respectively) therefore belong to this class. The “split” function is also a typical example of the more recently introduced *polyregular functions*, cf. [Boj22, Kie24] – in fact, it is a slight variation on the “squaring with underlining” function of [Boj22, Example 3].

REFERENCES

- [AKST19] Kazuyuki Asada, Naoki Kobayashi, Ryoma Sin’ya, and Takeshi Tsukada. Almost every simply typed lambda-term has a long beta-reduction sequence. *Logical Methods in Computer Science*, 15(1), 2019. doi:10.23638/LMCS-15(1:16)2019.
- [AL13] Andrea Asperti and Jean-Jacques Lévy. The cost of usage in the lambda-calculus. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28, 2013*, pages 293–300. IEEE Computer Society, 2013. doi:10.1109/LICS.2013.35.
- [Asp17] Andrea Asperti. About the efficient reduction of lambda terms. *CoRR*, abs/1701.04240, 2017. arXiv:1701.04240.
- [BB85] Corrado Böhm and Alessandro Berarducci. Automatic synthesis of typed λ -programs on term algebras. *Theoretical Computer Science*, 39:135–154, January 1985. doi:10.1016/0304-3975(85)90135-5.
- [Bec01] Arnold Beckmann. Exact bounds for lengths of reductions in typed lambda-calculus. *Journal of Symbolic Logic*, 66(3):1277–1285, 2001. doi:10.2307/2695106.
- [Blu09] William Blum. *The safe lambda calculus*. PhD thesis, University of Oxford, 2009. URL: <https://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.504329>.
- [BN98] Franz Baader and Tobias Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998. doi:10.1017/CB09781139172752.
- [BO09] William Blum and C.-H. Luke Ong. The Safe Lambda Calculus. *Logical Methods in Computer Science*, 5(1), February 2009. doi:10.2168/LMCS-5(1:3)2009.
- [Boj22] Mikołaj Bojańczyk. Transducers of polynomial growth. In Christel Baier and Dana Fisman, editors, *LICS ’22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, Haifa, Israel, August 2 - 5, 2022*, pages 1:1–1:27. ACM, 2022. doi:10.1145/3531130.3533326.
- [BS23] Pablo Barenbaum and Cristian Sottile. Two decreasing measures for simply typed λ -terms. In Marco Gaboardi and Femke van Raamsdonk, editors, *8th International Conference on Formal Structures for Computation and Deduction, FSCD 2023, July 3-6, 2023, Rome, Italy*, volume 260 of *LIPICs*, pages 11:1–11:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.FSCD.2023.11.
- [CHHM22] Dmitry Chistikov, Christoph Haase, Zahra Hadizadeh, and Alessio Mansutti. Higher-Order Quantified Boolean Satisfiability. In Stefan Szeider, Robert Ganian, and Alexandra Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 33:1–33:15. Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPICs.MFCS.2022.33.
- [Con20] Andrea Condoluci. *Beta-Conversion, Efficiently*. PhD thesis, Alma Mater Studiorum – Università di Bologna, April 2020. URL: <http://amsdottorato.unibo.it/9444/>.
- [dM06] Jolie G. de Miranda. *Structures generated by higher-order grammars and the safety constraint*. PhD thesis, University of Oxford, 2006. URL: <http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.442397>.

- [FMvO23] Samuel Frontull, Georg Moser, and Vincent van Oostrom. α -Avoidance. In Marco Gaboardi and Femke van Raamsdonk, editors, *8th International Conference on Formal Structures for Computation and Deduction (FSCD 2023)*, volume 260 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:22, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.FSCD.2023.22.
- [Geu92] Herman Geuvers. The Church-Rosser property for beta-eta-reduction in typed lambda-calculi. In *Proceedings of the Seventh Annual Symposium on Logic in Computer Science (LICS '92)*, Santa Cruz, California, USA, June 22–25, 1992, pages 453–460. IEEE Computer Society, 1992. doi:10.1109/LICS.1992.185556.
- [GRV09] Marco Gaboardi, Luca Roversi, and Luca Vercelli. A By-Level Analysis of Multiplicative Exponential Linear Logic. In *Mathematical Foundations of Computer Science 2009*, Lecture Notes in Computer Science, pages 344–355. Springer, Berlin, Heidelberg, August 2009. doi:10.1007/978-3-642-03816-7_30.
- [HK96] Gerd G. Hillebrand and Paris C. Kanellakis. On the Expressive Power of Simply Typed and Let-Polymorphic Lambda Calculi. In *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science*, pages 253–263. IEEE Computer Society, 1996. doi:10.1109/LICS.1996.561337.
- [HK11] Markus Holzer and Martin Kutrib. The complexity of regular(-like) expressions. *International Journal of Foundations of Computer Science*, 22(7):1533–1548, 2011. doi:10.1142/S0129054111008866.
- [HKM96] Gerd G. Hillebrand, Paris C. Kanellakis, and Harry G. Mairson. Database Query Languages Embedded in the Typed Lambda Calculus. *Information and Computation*, 127(2):117–144, June 1996. doi:10.1006/inco.1996.0055.
- [JT18] Guilhem Jaber and Nikos Tzevelekos. A trace semantics for system F parametric polymorphism. In Christel Baier and Ugo Dal Lago, editors, *Foundations of Software Science and Computation Structures - 21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14–20, 2018, Proceedings*, volume 10803 of *Lecture Notes in Computer Science*, pages 20–38. Springer, 2018. doi:10.1007/978-3-319-89366-2_2.
- [Kie24] Sandra Kiefer. Polyregular functions – characterisations and refutations. In Joel D. Day and Florin Manea, editors, *Developments in Language Theory - 28th International Conference, DLT 2024, Göttingen, Germany, August 12–16, 2024, Proceedings*, volume 14791 of *Lecture Notes in Computer Science*, pages 13–21. Springer, 2024. doi:10.1007/978-3-031-66159-4_2.
- [KNU02] Teodor Knapik, Damian Niwiński, and Paweł Urzyczyn. Higher-order pushdown trees are easy. In Mogens Nielsen and Uffe Engberg, editors, *Foundations of Software Science and Computation Structures, 5th International Conference, FOSSACS 2002. Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2002 Grenoble, France, April 8–12, 2002, Proceedings*, volume 2303 of *Lecture Notes in Computer Science*, pages 205–222. Springer, 2002. doi:10.1007/3-540-45931-6_15.
- [KSW16] Fairouz Kamareddine, Jonathan P. Seldin, and J. B. Wells. Bridging Curry and Church’s typing style. *Journal of Applied Logic*, 18:42–70, 2016. doi:10.1016/j.jal.2016.05.008.
- [Mai92] Harry G. Mairson. A simple proof of a theorem of Statman. *Theoretical Computer Science*, 103(2):387–394, September 1992. doi:10.1016/0304-3975(92)90020-G.
- [MS15] Guillaume Munch-Maccagnoni and Gabriel Scherer. Polarised intermediate representation of lambda calculus with sums. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6–10, 2015*, pages 127–140. IEEE Computer Society, 2015. doi:10.1109/LICS.2015.22.
- [Ngu21] Lê Thành Dũng Nguyễn. *Implicit automata in linear logic and categorical transducer theory*. PhD thesis, Université Paris XIII (Sorbonne Paris Nord), December 2021. URL: <https://theses.hal.science/tel-04132636>.
- [NP20] Lê Thành Dũng Nguyễn and Cécilia Pradic. Implicit automata in typed λ -calculi I: aperiodicity in a non-commutative logic. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8–11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 135:1–135:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ICALP.2020.135.

- [PA93] Gordon D. Plotkin and Martín Abadi. A logic for parametric polymorphism. In Marc Bezem and Jan Friso Groote, editors, *Typed Lambda Calculi and Applications, International Conference on Typed Lambda Calculi and Applications, TLCA '93, Utrecht, The Netherlands, March 16-18, 1993, Proceedings*, volume 664 of *Lecture Notes in Computer Science*, pages 361–375. Springer, 1993. doi:10.1007/BFb0037118.
- [Par18] Paweł Parys. Homogeneity without loss of generality. In Hélène Kirchner, editor, *3rd International Conference on Formal Structures for Computation and Deduction, FSCD 2018, July 9-12, 2018, Oxford, UK*, volume 108 of *LIPICs*, pages 27:1–27:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICs.FSCD.2018.27.
- [Pin17] Jean-Éric Pin. The dot-depth hierarchy, 45 years later. In Stavros Konstantinidis, Nelma Moreira, Rogério Reis, and Jeffrey O. Shallit, editors, *The Role of Theory in Computer Science – Essays Dedicated to Janusz Brzozowski*, pages 177–202. World Scientific, 2017. doi:10.1142/9789813148208_0008.
- [PP24] Cécilia Pradic and Ian Price. Implicit automata in λ -calculi III: affine planar string-to-string functions, 2024. arXiv:2404.03985.
- [Sal15] Sylvain Salvati. *Lambda-calculus and formal language theory*. Habilitation à diriger des recherches, Université de Bordeaux, December 2015. URL: <https://theses.hal.science/tel-01253426>.
- [Sch91] Helmut Schwichtenberg. An upper bound for reduction sequences in the typed λ -calculus. *Archive for Mathematical Logic*, 30(5-6):405–408, 1991. doi:10.1007/BF01621476.
- [Sch16] Sylvain Schmitz. Complexity hierarchies beyond elementary. *ACM Transactions on Computation Theory*, 8(1):3:1–3:36, 2016. doi:10.1145/2858784.
- [Sch17] Gabriel Scherer. Deciding equivalence with sums and the empty type. In Giuseppe Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, pages 374–386. ACM, 2017. doi:10.1145/3009837.3009901.
- [Sin22] Alexandros Singh. *A unified approach to the combinatorics of the λ -calculus and maps: bijections and limit properties*. PhD thesis, Université Paris XIII (Sorbonne Paris Nord), 2022. URL: <https://theses.hal.science/tel-04069290>.
- [Sta79] Richard Statman. The typed λ -calculus is not elementary recursive. *Theoretical Computer Science*, 9(1):73–81, July 1979. doi:10.1016/0304-3975(79)90007-0.
- [Sto74] Larry J. Stockmeyer. *The complexity of decision problems in automata theory and logic*. PhD thesis, Massachusetts Institute of Technology, 1974. URL: <http://hdl.handle.net/1721.1/15540>.
- [Str18] Howard Straubing. First-order logic and aperiodic languages: a revisionist history. *ACM SIGLOG News*, 5(3):4–20, 2018. doi:10.1145/3242953.3242956.
- [SW16] Sylvain Salvati and Igor Walukiewicz. Simply typed fixpoint calculus and collapsible pushdown automata. *Mathematical Structures in Computer Science*, 26(7):1304–1350, October 2016. doi:10.1017/S0960129514000590.
- [Tak95] Masako Takahashi. Parallel reductions in lambda-calculus. *Information and Computation*, 118(1):120–127, 1995. doi:10.1006/inco.1995.1057.
- [Ter12] Kazushige Terui. Semantic Evaluation, Intersection Types and Complexity of Simply Typed Lambda Calculus. In *23rd International Conference on Rewriting Techniques and Applications (RTA '12)*, pages 323–338, 2012. doi:10.4230/LIPICs.RTA.2012.323.
- [Wan87] Mitchell Wand. A simple algorithm and proof for type inference. *Fundamenta Informaticae*, 10(2):115–121, 1987. doi:10.3233/FI-1987-10202.