

Read/write factorizable programs

SIDDHARTH BHASKAR 

University of Copenhagen, Copenhagen, Denmark
(e-mail: bhask2sk@jmu.edu)

JAKOB GRUE SIMONSEN 

University of Copenhagen, Copenhagen, Denmark
(e-mail: simonsen@di.ku.dk)

Abstract

In the *cons-free* programming paradigm, we eschew constructors and program using only destructors. Cons-free programs in a simple first-order language with string data capture exactly P, the class of polynomial-time relations. By varying the underlying language and considering other data types, we can capture several other complexity classes. However, no cons-free programming language captures any *functional* complexity class for fundamental reasons. In this paper, we cleanly extend the cons-free paradigm to encompass functional complexity classes. Namely, we introduce programs with data that can either *only* be destructed or *only* be constructed, which we enforce by a type system on the program variables. We call the resulting programs *read/write-* (or *RW-*)factorizable, show that RW-factorizable string programs capture exactly the class FP of polynomial-time functions, and that tail-recursive RW-factorizable programs capture exactly the class FL of logarithmic-space functions. Finally, we state and solve the nontrivial problem of *syntactic composition* of two RW-factorizable programs.

1 Introduction

The primitive operations associated with certain data types can be separated into constructors and destructors. For example, over the natural numbers, the *successor* (or *increment*) function and the constant 0 are constructors, whereas the *predecessor* (or *decrement*) function and equals-zero test are destructors. Over strings on some fixed alphabet Σ , the *cons* function, which prepends a given character onto a given string, is a constructor, as is the constant naming the empty string. On the other hand, the *head* function, which isolates the first character of a given string, and the *tail* function, which deletes it, are destructors, as is the relation which tests whether a string is empty. Similarly, one can define constructors and destructors for trees, heaps, nested lists, and all sorts of other common data types.

A *cons-free* program over any one of these data types is a program in which no constructors occur. For example, consider the following cons-free natural number program, which decides whether its input is even:

```
even(n) = if zero(n) then true else if zero(n-1) then false, else even(n-2).
```

In other words, to check whether a given number is even, we return “true” if it is equal to zero, return “false” if it is equal to one, and otherwise subtract two and recurse. (Here, $n-2$ abbreviates $n-1-1$.) This program is clearly cons-free, as 0 and +1 do not occur within.

Cons-free programs clearly do not form a Turing complete language. Given their somewhat severe limitations, it is perhaps surprising that cons-free programs retain significant computational power: by a seminal result of Neil Jones (1999), cons-free string programs in a simple first-order functional language capture exactly the class P of polynomial-time relations. In other words, every language in P is decided by some cons-free program, and every language decided by a cons-free program is in P. Moreover, if we restrict our attention to *tail-recursive* programs, we capture the class L of logarithmic-space relations in exactly the same sense.

Cons-free programs, however, fail to capture the *functional* versions of any complexity class, e.g., the classes FP and FL of polynomial-time and logarithmic-space functions respectively. Indeed, they cannot even define the simplest function which increases the size of the input. To remedy this deficiency, we introduce the notion of an *read/write-factorizable program*. Such programs extend the destruct-only (or *read-only*) variables of cons-free programs with construct-only (or *write-only*) variables. For example, the following addition program is RW-factorizable:

$$\text{add}(n, m) = \text{if zero}(n) \text{ then } m \text{ else } \text{add}(n-1, m+1).$$

Here n is a read-only variable of type R , m is a write-only variable of type W , and $\text{add} : R \times W \rightarrow W$. Similarly, we can define a *string concatenation* function of type $R \times W \rightarrow W$ by

$$\text{cat}(x, w) = \text{if null}(x) \text{ then } w \text{ else } \text{cons}(\text{hd } x, \text{cat}(\text{tl } x, w)),$$

where x is a read-only variable of type R , w is a write-only variable of type W , $\text{null} : R \rightarrow 2$ tests whether read-only string is empty, $\text{cons} : \Sigma \times W \rightarrow W$ prepends a character in the alphabet Σ to a write-only string, $\text{hd} : R \rightarrow \Sigma$ is the head function, and $\text{tl} : R \rightarrow R$ is the tail function.

Consider the problem of programming a RW-factorizable *identity* function of type $R \rightarrow W$. Notice that the trivial program $\text{id}(x) = x$ cannot be consistently typed as $R \rightarrow W$; instead, we have to destruct the input and construct it again. Over strings, this would look something like $\text{id}(x) = \text{cat}(x, \text{nil})$. This shows that information can *flow* from R values to W values; however, the reverse is not possible.

Finally, let us try to define the base-2 exponential function $n \mapsto 2^n$. The program

$$\begin{aligned} \text{exp}(n) &= \text{if zero}(n) \text{ then } 1 \text{ else } \text{double}(\text{exp}(n-1)) \\ \text{double}(n) &= \text{add}(n, n) \end{aligned}$$

is *not* RW-factorizable, because $\text{double} : R \rightarrow W$, and exp cannot be well-typed, as its output must agree with both the input and output of double . However,

$$\text{leap}(n) = f(n, 0) \tag{1}$$

$$f(n, m) = \text{if zero}(n) \text{ then } m+1 \text{ else } f(n-1, f(n-1, m)) \tag{2}$$

is a perfectly legitimate RW-factorizable definition of $n \mapsto 2^n$, with $f : R \times W \rightarrow W$.

1.1 Our contribution

The central results of this paper are that

- RW-factorizable string programs of type $R \rightarrow W$ which are *non-nested* capture exactly the class FP of polynomial-time functions, and
- RW-factorizable strings programs of type $R \rightarrow W$ which are *tail-recursive* capture exactly the class FL of logarithmic-space functions.

In other words, the passage from cons-free to RW-factorizable programs is exactly what allows us to extend the results of Jones (1999) from relational to functional classes. The purpose of the *non-nested* stipulation is to exclude nested recursive programs like `leap` from line (1), whose outputs can have length exponential in the length of the inputs.¹

A foundational lesson of implicit computational complexity going back to Bellantoni & Cook (1992) is that any Turing-complete programming language must allow for the *same* piece of data to be *both read and written*. Our results illustrate this phenomenon in a particularly clear way.²

A natural question to ask about RW-factorizable programs is how to (syntactically) compose them, i.e., produce a program which computes the sequential composition of two given programs. Naive attempts fail: you cannot stick an output of type W directly into an input of type R . So the problem is nontrivial; however, we solve it at the end of this paper.

Both the capturing and composition results utilize the same technical device, namely *bit-length computability* of a function, which is introduced in this paper. This essentially means computability by two cons-free programs, one computing the bits, the other the length, of a given function on each input. Our results suggest that this idea might have some “legs” and be worthy of further study in its own right.

1.2 Organization of this paper

In Section 2, we review some technical background. We then introduce RW-factorizable programs over strings, our main object of study, in Section 3. In Section 4, we introduce bit-length programs, a dialect of cons-free programs, and define bit-length computability in Section 5. That section also states that FL and FP are captured by pairs of bit-length programs; a sketch of the proofs is postponed to Appendix B.

Sections 6 defines a compiling function from RW-factorizable to (pairs of) bit-length programs; Section 7 defines a compiling function in the other direction. These two sections establish the equivalence of the two models of computation and thus prove our desired capturing results for RW-factorizable programs. Section 8 treats syntactic composition of bit-length programs and hence of RW-factorizable programs by proxy. Finally in Section 9, we discuss further directions for research.

The beating technical heart of this paper consists of four program transformations in Section 6, two in Section 7, and one in Section 8. Some of these are simple and others more intricate, but each transformation implements a conceptually simple idea which we

¹ More precisely, we forbid nested recursive calls of type W , like $f(\dots f(\dots))$ in `leap`, but allow it for other types.

² Thanks to an anonymous referee who pointed this out.

outline at the top of its respective section. We advise that the reader read these first before diving into the details.

We develop most transformations according to a common template: first describing a transformation of types, then of variables and values, then of terms and programs, and finally proving correctness. In the interests of space, we postpone all the proofs of correctness to Appendix A.

1.3 Related work

The present paper falls squarely into a long tradition of characterizing complexity classes by function algebras, programming languages, and related models of computation, a topic known as *implicit computational complexity* (ICC). ICC has been an area of interest since, at least, the 1960s (Cobham, 1965), with a resurgence in the last 30 years since the groundbreaking work of Bellantoni & Cook (1992). There are multiple approaches to ICC, see, e.g., the survey (Hofmann, 2000) for early developments pre-2000, and Dal Lago (2022) for a recent survey about methods involving higher-order programs. Recently, implicit characterizations have been furnished for complexity classes using different modes of computation, or using different modes of computation as a means of characterizing standard complexity classes; in particular, this includes probabilistic computation (Lago & Toldin, 2015; Lago et al., 2021), reversible computation (Kristiansen, 2022), parallel computation (Baillot & Ghyselen, 2022), and higher-order complexity (Hainry et al., 2022).

Many of these results are obtained by imposing a type system on a base programming language, as we do in this paper. Of the multitude of different flavors of implicit characterizations, we briefly review the most well-known ones with connections to our work:

- *Data ramification*. By this, we mean factoring the base data into two or more copies, and restricting how we can access or modify data in each copy. (Our approach of RW-factorization falls into this broad category.) The most common of these is the *normal/safe* factorization, which underlies many of the foundational papers in the field, e.g., Bellantoni & Cook (1992), Leivant (1995), and has a wide variety of applications. Our work is distinguished from these by our use of a general-purpose programming language with flexible recursive definitions, rather than function algebras based off of primitive recursion.

There are other instances of data ramification, for example the *secure information flow* of Marion (2011), which partitions data into “higher” and “lower” security tiers. This bears resemblance to *W*- and *R*-data respectively, but ultimately yields different results: higher-security data does allow a limited amount of destruction, and FP is characterized by a class of tail-recursive programs. Still, this is perhaps the characterization closest to our own in spirit.

- *Linearity*. By linearity we mean (very roughly) restricting the reuse of variables or primitive operations, e.g., controlling the number of times a constructor can occur in the right-hand side of a recursive definition. Approaches to ICC based on linear logic include proof-theoretic approaches like light linear logic (Girard, 1998) and soft linear logic (Lafont, 2004), as well as the type-based approach of

non-size-increasing computation (Hofmann, 2002). From a practical perspective, they seem to yield more intensionally expressive languages than safe primitive recursion and have been applied to, e.g., cryptographic protocols (Baillot et al., 2019).

There does not seem to be any real technical connection between our work and this approach, except perhaps that our “non-nested” restriction of RW-factorizable programs carries a whiff of a linear typing discipline.

- *Restricted termination.* This is (roughly) the study of measures or bounds on programs or computation states, such that only those programs having execution that terminates within the bound, are part of the characterization. This is a flourishing field, both for traditional programs (Bonfante et al., 2011; Aubert et al., 2022) and for non-deterministic models such as term rewriting systems (Avanzini et al., 2011, 2015; Avanzini & Moser, 2016).

The approach of the present paper does not use restricted termination in any technical sense, and the characterizations we provide are not reliant on termination; there is, however, a tenuous connection in the sense that we use programs with very restricted ability to construct data structures, and this might be amenable to analysis by methods devised for ICC using restricted termination.

- *Cons-free programs.* This is the particular line of inquiry which the present paper belongs to. Jones (1999) was the first to observe that eschewing constructors yields a simple method of capturing polynomial-time and logarithmic-space relations. This paradigm has since been extended to obtain many capturing results in a variety of contexts, e.g., higher types (Jones, 2001), non-determinism (Bonfante, 2006), simultaneous higher types and non-determinism (Kop & Simonsen, 2017), tree-like data (Ben-Amram & Petersen, 1998), term rewriting systems (de Carvalho & Simonsen, 2014; Czajka, 2018), and cons-free time complexity classes (Bhaskar et al., 2022; Jones et al., 2020). The present paper is the first to consider computability of function classes.

2 Preliminaries

We assume a basic familiarity with syntax and semantics of first-order functional programming languages, Turing machines, the Turing machine complexity classes P (polynomial time) and L (logarithmic space), and their functional variants FP and FL. Note that when defining a functional space class, the space bound only applies to the *work* tapes and not the output tape. Hence, the length of the output of function in FL may in general grow polynomially in the length of the input.

Following the set-theoretic convention, we identify any natural number n with its set of predecessors $\{0, 1, \dots, n-1\}$ and we denote the set of all natural numbers by ω . Therefore, $1 = \{0\}$, $2 = \{0, 1\}$, and 1^n and 2^n are the sets of unary and binary strings of length n , respectively. We also identify 0 and 1 with \perp (*false*) and \top (*true*), respectively, so 2 is also the set of boolean values.

We denote the set of all finite unary and binary strings by 1^* and 2^* , respectively, following the convention in computer science.³ The *empty string* is denoted ε ; its underlying

³ In the set-theoretic convention these are $1^{<\omega}$ and $2^{<\omega}$ respectively.

alphabet must be inferred from context. The *head* of a string is its first character, and its *tail* is the remainder. The head and tail of ε are undefined. If x is a string, then $|x|$ is its *length*.

We typically use lowercase Latin letters such as m, n, u, v, w, x, y for variables that range over “type-0 objects” such as natural numbers and strings. By, e.g., a “ n -ary relation on strings,” we mean a subset of $(2^*)^n$. We use capital Latin letters, e.g., P, Q, X, Y , for variable that range over relations, and f, g, h for variables that range over functions. We use lowercase Greek letters for variables that range over types. We use typewriter script for program syntax.

A *partial function* $f : X \rightarrow Y$ is a function $f : X' \rightarrow Y$ for some $X' \subseteq X$; X' is the *domain of convergence* of f . If x is in the domain of convergence of f , we write $f(x) \downarrow$, otherwise $f(x) \uparrow$, and say that f “converges” or “diverges” on x , respectively. Partial functions converge strictly, meaning that $f(g(x)) \downarrow$ implies, in particular, that $g(x) \downarrow$. For two partial functions $f, f' : X \rightarrow Y$, we say that $f' \sqsubseteq f$ in case $f'(x) = y \implies f(x) = y$ for all $(x, y) \in X \times Y$. Finally, note that by “partial function” we do mean a possibly total function, otherwise we will say *properly partial*.

The semantics of a program p is the partial function denoted by $\llbracket p \rrbracket$. A program p *computes* a partial function f in case $\llbracket p \rrbracket$ is identical to f as a partial function. A program p *accepts* a set X in case X is the domain of convergence of $\llbracket p \rrbracket$. Finally, a boolean-valued program p *decides* a set X if it computes the characteristic function of X .

A partial function computed by a program may in general have a dependent type. Given a type α and family of types $\beta(x)$ indexed by $x : \alpha$, we will refer to the *dependent sum type* $\sum_{x:\alpha} \beta(x)$ and the *dependent product type* $\prod_{x:\alpha} \beta(x)$. (When β does not depend on x , these reduce to $\alpha \times \beta$ and $\alpha \rightarrow \beta$ respectively.) The language of dependent types provides an elegant formulation for *counting modules*, but we do not use the machinery of dependent types in any essential way.

A note on functional complexity. Suppose that $f : 2^* \rightarrow 2^*$ is contained in the class FP of polynomial-time computable functions. Consider the function $f_\ell : 2^* \rightarrow 1^*$ defined by $f_\ell(x) = |f(x)|$, i.e., the length of $f(x)$. Then, f_ℓ is also in FP; take, for example, any Turing machine computing f in polynomial time and identify all the characters of its output alphabet. Similarly, consider the function $f_b : 2^* \times 1^* \rightarrow 2$ defined by $f_b(x, i) = (f(x))_i$, meaning the i^{th} -bit of $f(x)$, for any $x \in 2^*$ and $i < |f(x)|$. This relation is computable in polynomial time as well: on input (x, i) , write $f(x)$ on a work tape and extract the i th-bit.⁴

All this is to say, roughly speaking, that if a function of type $2^* \rightarrow 2^*$ is computable in polynomial time, then so is its *length* function and its *bits* relation. We note that the converse implication is valid too. In other words, if f_ℓ and f_b are both computable in polynomial time, then we can compute f in polynomial time by concatenating the bits $f_b(x, i)$ for each $i < f_\ell(x)$.

In other words, we have reduced the notion of polynomial-time computability of functions of type $2^* \rightarrow 2^*$ to polynomial-time computability of relations and of functions $2^* \rightarrow 1^*$. What do we gain from this? We get *some* characterization of polynomial-time

⁴ There is some slight imprecision here, since we haven't said what $f_b(x, i)$ is if $i \geq |f(x)|$. It doesn't matter: we view $f_b(x, i) = (f(x))_i$ as a *partial* specification. If $f \in \text{FP}$, there is *some* such $f_b \in \text{FP}$; if there is some such $f_b \in \text{FP}$ and $f_\ell \in \text{FP}$, then $f \in \text{FP}$.

computability of functions $2^* \rightarrow 2^*$ in terms of cons-free programs. (Cons-free programs can handle data of type 1^* with *counting modules*, a device for reckoning with natural number quantities bounded by a polynomial in the length of the input.) So the polynomial-time computability of $f : 2^* \rightarrow 2^*$ may be witnessed by two cons-free programs: one with string input and counting module output computing the length and another with string and counting module input and boolean outputs, computing the bits. We shall elaborate on this later.

The observations we have made about polynomial time apply respectively to *logarithmic space* and *tail-recursive* cons-free programs. The proof is slightly different: notice that if $f : 2^* \rightarrow 2^*$ is in FL, then we cannot compute $f_b(x, i)$ by writing $f(x)$ on a work tape, because it will be too long, in general. However, there is a well-known trick in space complexity to circumvent this, viz., replacing the write-only output tape by a work tape that records only the *position* of the head. Similarly, when computing $f(x)$ using f_b and f_ℓ , we compute $f_\ell(x)$ *in binary*, which compresses it enough to stick it on a work tape.

We compile these observations into an official and easily referenced theorem.

Theorem 1. *A function $f : 2^* \rightarrow 2^*$ is computable in polynomial time, respectively, logarithmic space, iff there are functions $f_\ell : 2^* \rightarrow 1^*$ and $f_b : 2^* \times 1^* \rightarrow 2$ computing the length and bits of f which are computable in polynomial time, respectively logarithmic space.*

3 RW-factorizable programs

Let us work over the set 2^* of binary strings. Consider the following set of string primitives⁵

- hd , denoting the head function, of type $2^* \rightarrow 2$,
- tl , the tail function, of type $2^* \rightarrow 2$
- null , the empty test, of type $2^* \rightarrow 2$
- nil , a constant naming the empty string, of type 2^* , and
- cons , a binary function prepending a given character onto a given string, of type $2 \times 2^* \rightarrow 2$.

Now, consider two separate copies of 2^* , viz., R , whose strings are read-only, and W , whose strings are write-only. Then, we may retype these primitives, replacing each occurrence of 2^* by either R or W as follows:

$$\text{hd} : R \rightarrow 2, \quad \text{tl} : R \rightarrow R, \quad \text{null} : R \rightarrow 2, \quad \text{nil} : W, \quad \text{cons} : 2 \times W \rightarrow W.$$

Note that this typing is consistent with strings in R being “read-only” and strings in W being “write-only.”

From these primitives, we construct a simple first-order programming language. First we define the types, then terms, of our programming language, then we define the programs themselves; finally, we equip these programs with an environment-based big-step semantics.

⁵ Here we conflate program syntax, e.g., hd , with the function it denotes. We trust that this imprecision causes no confusion. The symbol \rightarrow occurs in hd and tl because they are undefined on the empty string.

| | | | | |
|---|----------------------------------|--|------------------------------------|---------------------------|
| $\frac{x \in \text{Var}_\alpha}{x : \alpha}$ | $\frac{}{w : W}$ | $\frac{}{\text{true} : 2}$ | $\frac{}{\text{false} : 2}$ | $\frac{}{\text{nil} : W}$ |
| $\frac{T : 2 \quad S : W}{\text{cons}(T, S) : W}$ | $\frac{T : R}{\text{hd}(T) : 2}$ | $\frac{T : R}{\text{tl}(T) : R}$ | $\frac{T : R}{\text{null}(T) : 2}$ | |
| $\frac{T_0 : \alpha_0 \quad \dots \quad T_{n-1} : \alpha_{n-1}}{T_0 \oplus \dots \oplus T_{n-1} : \alpha_0 \times \dots \times \alpha_{n-1}}$ | | $\frac{T : \tau_0 \times \dots \times \tau_{n-1} \quad i \leq j < n}{T[i, j] : \tau_i \times \dots \times \tau_j}$ | | |
| $\frac{T_0 : 2 \quad T_1 : \alpha \quad T_2 : \alpha}{\text{if } T_0 \text{ then } T_1 \text{ else } T_2 : \alpha}$ | | | | |
| $\frac{T : \beta \quad f \in \text{RFsymb}_{\beta \rightarrow \alpha}}{f(T) : \alpha}$ | | $\frac{T : \beta \quad f \in \text{RFsymb}_{\beta \rightarrow W}}{f(T) : W}$ | | |
| $\frac{T : \beta \quad S : W \quad f \in \text{RFsymb}_{\beta \times W \rightarrow W}}{f(T, S) : W}$ | | | | |

Fig. 1. RW-terms. Letters α, β , and each α_i range over product types and each τ_i ranges over the atomic types 2 and R .

Definition 1. The three *atomic types* are 2, R , and W , denoting the type of booleans, read-only strings, and write-only strings, respectively. A *product type* is any expression of the form $\tau_0 \times \dots \times \tau_{n-1}$, for $n \geq 0$, where each τ_i is either 2 or R . (W is excluded from product formation.) When $n = 0$, the product is empty. We extend the type constructor \times to apply to product types by

$$(\tau_0 \times \dots \times \tau_{n-1}) \times (\tau_n \times \dots \times \tau_{m-1}) = \tau_0 \times \dots \times \tau_{m-1}.$$

A *function type symbol* is an expression of one of the following forms

$$\beta \rightarrow \alpha, \quad \beta \rightarrow W, \quad \beta \times W \rightarrow W,$$

where β and α are product types, and α is nonempty.

Definition 2. For each product type α , fix an infinite set Var_α of variables of type α . For each product type β , fix infinite sets $\text{RFsymb}_{\beta \rightarrow W}$ and $\text{RFsymb}_{\beta \times W \rightarrow W}$ of function symbols of type $\beta \rightarrow W$ and $\beta \times W \rightarrow W$, respectively; similarly define $\text{RFsymb}_{\beta \rightarrow \alpha}$ for each pair of product types (α, β) .

By a *variable*, we mean any member of any Var_α , or the symbol w , which is the sole variable of type W . By a *recursive function symbol*, we mean any member of one of the sets $\text{RFsymb}_{\beta \rightarrow W}$ or $\text{RFsymb}_{\beta \times W \rightarrow W}$, for any β, W .

Definition 3. An *RW term* is any expression that can be derived according to the inference rules in Figure 1.

It is straightforward to show that any RW term has a unique derivation from these inference rules, and that if an RW term contains any subterm of type W , it must itself have type W .

Notice that we restrict the way that W terms can occur: there is only one variable of type W and recursive function symbols of output type W have at most one W input. This restriction, which we may refer to as the W -thinness of RW terms, does not ultimately limit the expressive power of the resulting language. Roughly speaking, this is because a W datum is a sort of black box: no information can flow out of it and it cannot affect the *shape* of a computation. Think of it as a write-only output tape of a Turing machine: a single write-only output tape suffices.

Definition 4. We identify a few important properties of RW terms.

- An RW term is *explicit* if it contains no occurrence of a recursive function symbol.
- An RW term is *cons-free* if it is not of type W ; equivalently, if it does not contain any subterm of type W .
- An RW term is *non-nested* in case it contains no occurrence of a recursion function symbol of type W inside another occurrence of a recursive function symbol of type W . More precisely, in any application of the bottom-most rule of Figure 1, the second hypothesis S must be explicit.
- An RW term is *tail-recursive* in case there is no occurrence of a recursive function symbol inside any other occurrence of a recursive function symbol, or any primitive call, or in the `if` clause of any `if / then / else` term. More precisely, in any application of a rule in the second row of Figure 1, or any application of the three recursive function calls at the bottom of 1, the terms in the hypotheses must be explicit; when forming `if T_0 then T_1 else T_2` , T_0 must be explicit.

Remark 1. Notice that every tail-recursive term is also non-nested. Tail recursive RW-factorizable programs will end up capturing FL and non-nested programs will end up capturing FP; the program `leap` of Section 1 is witness to the necessity of this restriction.

Notice also that to be non-nested we only forbid nested recursive calls of output type W . (Surprisingly, we can always eliminate nested recursive calls in the purely cons-free part of a program, though we will not need this result in the present paper.)

Remark 2. We will actually silently use a slightly more general formulation of tail recursion. Notice that tail recursion affords more flexibility to calls to the primitives over recursive calls. For example, it is perfectly legal to write phrases such as `hd(tl(...))` but not `hd(f(...))`, in a tail-recursive program.

However, consider a transformation $T \mapsto T^\circ$ on terms that uses the recursive functions of a previously defined tail-recursive program p^\dagger . In arguing that $T \mapsto T^\circ$ preserves tail recursion, it suffices to treat calls to the recursive functions of p^\dagger like calls to the primitives as opposed to other recursive calls.⁶

⁶ Suppose that a function f is computed by a tail-recursive program from the primitives Φ and that a function g is computed by a tail-recursive program from the primitives $\Phi \cup \{f\}$. Then, g can be computed by a tail-recursive program directly from Φ , but the simple transformation replacing each call to f by the Φ program computing it may not preserve tail recursion. So there is a theorem here, but an ancient one, almost certainly folklore. This remark appeals to this theorem, basically saying that to show that g is Φ -tail-recursive, it suffices to show that it's $\Phi \cup \{f\}$ -tail-recursive. Cf. x2A.1 of Moschovakis (2018).

Definition 5. An *RW-factorizable program* consists of a finite list of distinct recursive function symbols (f_0, \dots, f_k) and a definition for each one, which takes one of the following two forms:

$$f_i(x_i, w) = T_i \text{ or } f_i(x_i) = T_i,$$

where T_i is a RW-term, x_i is a variable, and every recursive function symbol that occurs in T_i must be one of (f_0, \dots, f_k) . Furthermore

- If the definition of f_i is of the form $f_i(x_i, w) = T_i$, then $f_i(x_i, w)$ must be a well-formed RW term of the same type as T_i , and the only variables that may occur in T_i are x_i and w .
- If the definition of f_i is of the form $f_i(x_i) = T_i$, then $f_i(x_i)$ must be a well-formed RW term of the same type as T_i , and the only variable that may occur in T_i is x_i .

The *head (term)* of a program is the left-hand side of its first line, e.g., $f_0(x_0)$ or $f_0(x_0, w)$. A program is *cons-free* in case all terms occurring within are cons-free.

Remark 3. Any RW-factorizable program can be split into a “purely cons-free part” and a “ W -part.” The purely cons-free part consists of all recursive function symbols of output type other than W and their definitions. This part is a self-contained “sub-program”: it does not contain any recursive function calls of type W , and its semantics is independent of the rest of the program. The W part of the program consists of all recursive function symbols of output type W and their definitions. It lies on top of the purely cons-free part, using it as a black box.

We equip programs with a standard, call-by-value, environment-based big-step semantics. By a *value*, we mean an element of the domain of some type, and by an *environment*, we mean a finite function mapping variables to values.

Definition 6. For a program p , term T , value v of the same type as T , and environment ρ , we define the relation $\rho \vdash_p T \rightarrow v$ according to the inference rules in Figure 2. We typically suppress the subscript p from \vdash for legibility.

Definition 7. Given a program p , define $\llbracket p \rrbracket(x) = v$ iff there is a derivation of $[x = x] \vdash_p f_0(x) \rightarrow v$, where $f_0(x)$ is the head of p . (If the head of p is of the form $f_0(x, w)$, then $\llbracket p \rrbracket(x, w) = v$ iff there is a derivation of $[x = x, w = w] \vdash_p f_0(x, w) \rightarrow v$.) Since p is deterministic, $\llbracket p \rrbracket$ is easily seen to be a partial function.

Note that the relation \vdash_p is independent of the head of p ; it is only when defining $\llbracket p \rrbracket$ that we care what the head is.⁷ Aside from specifying the head, neither \vdash_p nor $\llbracket p \rrbracket$ is sensitive to the particular order of the recursive function symbols in a program; it is only because of overwhelming programming intuition that we say *a list* instead of *a set* of function symbols.

⁷ A more precise treatment might distinguish *headless programs* from programs but we find this both slightly morbid and excessively pedantic.

| | |
|---|---|
| $(x \text{ a variable in some } \text{Var}_\alpha) \frac{\rho(x) = v}{\rho \vdash x \rightarrow v}$ | $\frac{\rho(\bar{w}) = w}{\rho \vdash \bar{w} \rightarrow w}$ |
| $\frac{}{\rho \vdash \text{true} \rightarrow \top}$ | $\frac{}{\rho \vdash \text{f} \rightarrow \perp}$ |
| $\frac{}{\rho \vdash \text{nil} \rightarrow \varepsilon}$ | |
| $(bw' = w) \frac{\rho \vdash T \rightarrow b \quad \rho \vdash S \rightarrow w'}{\text{cons}(T, S) \rightarrow w}$ | $((\exists r' \in 2^*) br' = r) \frac{\rho \vdash T \rightarrow r}{\text{hd}(T) \rightarrow b}$ |
| $((\exists b \in 2) br = r') \frac{\rho \vdash T \rightarrow r'}{\text{tl}(T) \rightarrow r}$ | $(r = \varepsilon) \frac{\rho \vdash T \rightarrow r}{\text{null}(T) \rightarrow \top}$ |
| $(r \neq \varepsilon) \frac{\rho \vdash T \rightarrow r}{\text{null}(T) \rightarrow \perp}$ | |
| $\frac{\rho \vdash T_0 \rightarrow v_0 \quad \dots \quad \rho \vdash T_{n-1} \rightarrow v_{n-1}}{\rho \vdash T_0 \oplus \dots \oplus T_{n-1} \rightarrow v_0 \circ \dots \circ v_{n-1}}$ | $(i < n) \frac{\rho \vdash T \rightarrow (a_0, \dots, a_{n-1})}{\rho \vdash T[i] \rightarrow a_i}$ |
| $\frac{\rho \vdash T_0 \rightarrow \text{true} \quad \rho \vdash T_1 \rightarrow u}{\rho \vdash \text{if } T_0 \text{ then } T_1 \text{ else } T_2 \rightarrow u}$ | $\frac{\rho \vdash T_0 \rightarrow \text{false} \quad \rho \vdash T_2 \rightarrow u}{\rho \vdash \text{if } T_0 \text{ then } T_1 \text{ else } T_2 \rightarrow u}$ |
| $\frac{\rho \vdash T \rightarrow v' \quad [x \mapsto v'] \vdash T^{\text{f}} \rightarrow v}{\rho \vdash \text{f}(T) \rightarrow v}$ | $\frac{\rho \vdash T \rightarrow v' \quad \rho \vdash S \rightarrow w \quad [x \mapsto v', \bar{w} \mapsto w] \vdash T^{\text{f}} \rightarrow v}{\rho \vdash \text{f}(T, S) \rightarrow v}$ |

Fig. 2. Program Semantics. In the bottom two rules, $\text{f}(x) = T^{\text{f}}$ or $\text{f}(x, \bar{w}) = T^{\text{f}}$ is the recursive definition of f in p . Also, w ranges over values of type W ; r and r' range over values of type R ; b ranges over values of type 2 ; the a_i range over values of type 2 and R ; v, v', v_i range over values of any product type; and u ranges over values of any type. The value $v_0 \circ \dots \circ v_{n-1}$ denotes concatenation of the constituent tuples. For example if $v_0 = (a_0, a_1)$ and $v_1 = (a_2, a_3, a_4)$ then $v_0 \circ v_1 = (a_0, a_1, a_2, a_3, a_4)$. In general we will not carefully maintain this correspondence between variable names and their types.

In subsequent sections, we will define programs by *extending* previously defined programs with additional recursive function symbols and their definitions. This simply means: take the old program, forget what its head is, add new lines to the program—we don't care in what order—and pick a new head according to the definition at hand. If q is the old program and p the new one, notice that \vdash_p extends \vdash_q .

4 Bit-length programs

In this section, we construct a dialect of cons-free programs and use them to define function computability in the manner sketched in Section 2. In other words, two programs are needed to witness the computability of a function $f : 2^* \rightarrow 2^*$ in polynomial time: one that computes the *bits* of f and the other that computes the *length*. Hence, we call these programs (rather unimaginatively) *bit-length* (or *BL*) programs.

The bit-length dialect differs from the standard cons-free language in two ways. The first difference is to replace R -data with string indices as follows. Imagine a cons-free program with a single string input. Then any string constructed during computation will be a suffix of that input. Instead, we might as well consider indices of the input string. Instead of querying the head of the various string suffixes, we query the *bit* of the input string at the given index.

Table 1. A correspondence between languages in the present paper on the left and Jones (1999) on the right. A bit-length program is C -free if it contains no counting modules

| | |
|--|----------------------|
| bit-length programs | $CM^{rec-poly}$ |
| C -free bit-length programs | $CM^{rec\setminus+}$ |
| tail-recursive bit-length programs | CM^{poly} |
| C -free tail-recursive bit-length programs | CM^+ |

The second difference is that counting modules—a type of natural numbers with magnitude bounded polynomially in the length of the input—are treated as a separate data type, a dependent type which varies with the length of the input string, as opposed to syntactic sugar.

Each of these modifications confers an advantage. Replacing R values by string indices lends a sort of extensional compositionality to bit-length programs that RW-factorizable programs lack. Meaning, the pair of programs computing a function transforms the same type of information about the input (its bits and length, given by the index primitives) into the same type of information about the output (its bits and length, given by the two programs). Moreover, a careful treatment of counting modules allows for cleaner representation strings in 1^* . (The relevance of unary strings to computation of functions was discussed in Section 2.)

We note that the original paper (Jones, 1999) contains a language called CM for *counter machine*, which was used as a technical tool in the proofs of the main results of that paper. A variant, $CM^{rec-poly}$, is basically identical to our bit-length language. However, our presentation is different enough to justify a separate treatment. For one, CM and its variants are imperative instead of functional. Secondly, indices and counting modules are conflated in CM, whereas we treat them as different types. A glossary between the present paper and Jones (1999) can be found in Table 1.

As for RW programs, we explain bit-length programs first by discussing types, then terms, then programs, and finally semantics.

Definition 8. For any $n \in \omega$, the *counting module* $C(n)$ is a data type whose domain is the set $n = \{0, 1, \dots, n - 1\}$, and whose primitives consist of constants naming the maximum ($n - 1$), minimum (0), and 1, plus operations of addition, subtraction, and comparison and equality tests. Addition and subtraction “top out” and “bottom out” at the maximum and minimum element, respectively.⁸

Definition 9. For any string $x \in 2^*$, the set of indices $I(x)$ is a data type whose domain is $|x| + 1 = \{0, 1, \dots, |x|\}$, and whose primitives consist of constants naming the maximum ($|x|$), minimum (0), predecessor (which decrements the index), equality to zero, and *bit* (which returns the bit of the input at the given index).

Remark 4. We index strings such that the leftmost character has index the length of the string, and the rightmost character has index 1. That is, for any string $x, x = x_{|x|}x_{|x|-1} \dots x_2x_1$. The bit x_0 is undefined. This funny indexing makes the compiling

⁸ If the sum of two numbers is greater than the maximum, for example, then the result is simply the maximum.

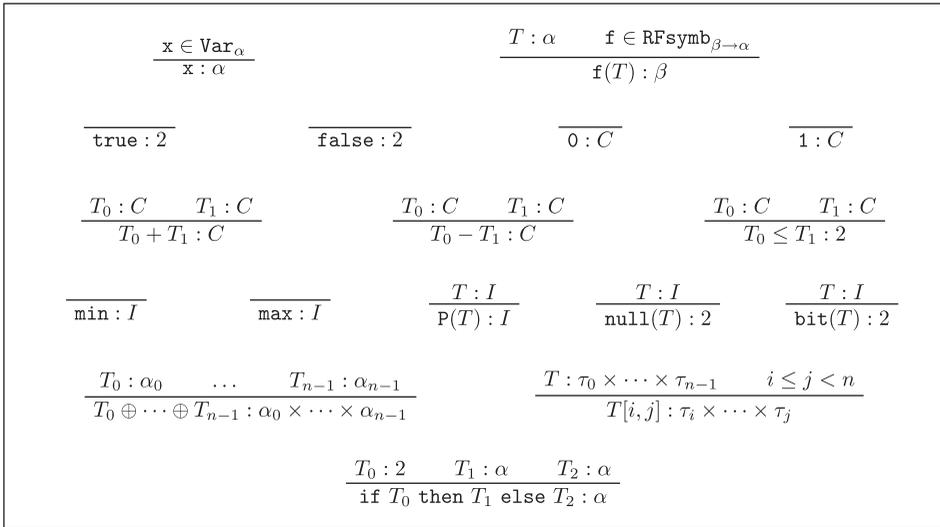


Fig. 3. Bit/length terms. Here α ranges over product types and τ over atomic types.

between bit-length and cons-free programs cleaner, as identifying indices with suffixes identifies `tl` with the predecessor function.

Notice that C and I are *dependent types*; namely, they are families of types indexed by some other value (natural numbers or binary strings). Now we introduce program syntax.

Definition 10. There are three *atomic type symbols*, 2 , I , and C , denoting the type of booleans, indices, and counting modules respectively. A *BL-type symbol* is any expression of the form $\tau_0 \times \dots \times \tau_{n-1}$, for $n \geq 0$, where each τ_i is an atomic type symbol. When $n = 0$, the product is empty, and when $n = 1$, we recover the atomic type symbols. A *BL-function type symbol* is an expression of the form $\beta \rightarrow \alpha$, where β and α are product types, and α is nonempty.

Definition 11. For a BL-type symbol α , we write $\alpha(x, n)$ to denote the type obtained by specializing each coordinate of α to x or n as appropriate. (Similarly for BL-function types.)

Remark 5. In the context of a single computation of a bit-length program, x and n will be fixed. Therefore, it is fine to type variables by type *symbols*, like 2 , I , C , $2 \times I \times C$, etc., since we do not have multiple instantiations of these types in a single computation. In contrast, *values* have types like $I(x)$, $I(x) \times C(n)$, etc., where x is a string and n a number.

Definition 12. Fix an infinite set Var_α of variables of type α , for each BL-type α , and an infinite set $\text{RFsymb}_{\beta \rightarrow \alpha}$ of BL-function symbols of type $\beta \rightarrow \alpha$, for each function type $\beta \rightarrow \alpha$. Then, a *bit-length term* is any expression which can be derived according to the inference rules in Figure 3.

| | |
|---|--|
| $\frac{\rho(x) = v}{\rho \vdash x \rightarrow v}$ | $\frac{\rho \vdash T \rightarrow v' \quad [x^f = v'] \vdash T^f \rightarrow v}{\rho \vdash f(T) \rightarrow v}$ |
| $\overline{\rho \vdash \text{true} \rightarrow \top}$ | $\overline{\rho \vdash f \rightarrow \perp}$ |
| $\overline{\rho \vdash 0 \rightarrow 0}$ | $\overline{\rho \vdash 1 \rightarrow 1}$ |
| $\frac{\rho \vdash T_0 \rightarrow c \quad \rho \vdash T_1 \rightarrow d}{\rho \vdash T_0 + T_1 \rightarrow \min\{c + d, n\}}$ | $\frac{\rho \vdash T_0 \rightarrow c \quad \rho \vdash T_1 \rightarrow d}{\rho \vdash T_0 - T_1 \rightarrow \max\{c - d, 0\}}$ |
| (if $c \leq d$) $\frac{\rho \vdash T_0 \rightarrow c \quad \rho \vdash T_1 \rightarrow d}{\rho \vdash T_0 \leq T_1 \rightarrow \top}$ | (if $c < d$) $\frac{\rho \vdash T_0 \rightarrow c \quad \rho \vdash T_1 \rightarrow d}{\rho \vdash T_0 \leq T_1 \rightarrow \perp}$ |
| $\overline{\rho \vdash \min \rightarrow 0}$ | $\overline{\rho \vdash \max \rightarrow x }$ |
| (if $i = 0$) $\frac{\rho \vdash T \rightarrow i}{\rho \vdash \text{null}(T) \rightarrow \top}$ | (if $i \neq 0$) $\frac{\rho \vdash T \rightarrow i}{\rho \vdash \text{null}(T) \rightarrow \perp}$ |
| $\frac{\rho \vdash T \rightarrow i}{\rho \vdash \text{bit}(T) \rightarrow x_i}$ | |
| $\frac{\rho \vdash T_0 \rightarrow v_0 \quad \dots \quad \rho \vdash T_{n-1} \rightarrow v_{n-1}}{\rho \vdash T_0 \oplus \dots \oplus T_{n-1} \rightarrow v_0 \circ \dots \circ v_{n-1}}$ | ($i < n$) $\frac{\rho \vdash T \rightarrow (a_0, \dots, a_{n-1})}{\rho \vdash T[i] \rightarrow a_i}$ |
| $\frac{\rho \vdash T_0 \rightarrow \top \quad \rho \vdash T_1 \rightarrow v}{\rho \vdash \text{if } T_0 \text{ then } T_1 \text{ else } T_2 \rightarrow v}$ | $\frac{\rho \vdash T_0 \rightarrow \perp \quad \rho \vdash T_2 \rightarrow v}{\rho \vdash \text{if } T_0 \text{ then } T_1 \text{ else } T_2 \rightarrow v}$ |

Fig. 4. Semantics for bit-length programs. Dependence of \vdash on x, n , and ρ is suppressed for legibility. $f(x^f) = T^f$ is the recursive definition of f in p and x_i is the i -th bit of $x = x_{|x|} \dots x_1$. Variables c and d range over $C(n)$, i ranges over $I(x)$, the a_i range over values in any atomic type, and v, v' and the v_i range over values of any type. In general we will not carefully maintain this correspondence between variables and their types. As in RW-factorizable programs, note that while we can form larger tuples from smaller proper tuples, we can only decompose tuples into atomic types.

Definition 13. A bit-length term is *explicit* in case it contains no occurrence of a recursive function symbol. It is *tail-recursive* in case no recursive function symbol occurs inside any other recursive call, primitive call, or if clause. (We are not concerned with non-nested bit-length terms.)

Definition 14. A *bit-length program* is a finite list of lines of the form $f_i(x_i) = T_i$, for $0 \leq i \leq k$, where x_i is a variable whose type agrees with the domain of f , and the type of T_i agrees with the codomain of f_i . In addition, the only variable that occurs in T_i must be x_i , and the only recursive function symbols that occur in T_i must be one of (f_0, \dots, f_k) .

Definition 15. For any bit-length program p , input string $x \in 2^*$, seed $n \in \omega$, term T of type α , environment ρ binding the free variables of T , and value v of type $\alpha(x, n)$, we define the relation

$$x, n, \rho \vdash_p T \rightarrow v$$

according to the inference rules in Figure 4. Note that there is at most one v satisfying the above relation for each x, n, ρ, p , and T , which means that p is *deterministic*.

Remark 6. Note the extra two parameters x and n in the definition of semantics of bit-length programs. These may be viewed as analogous to a global variable if you're a programmer or the structure on the left-hand side of the satisfiability relation if you're a logician. Without them, the function symbols `max`, `+`, and `bit` are ill defined.

Definition 16. For any program p and function $\lambda : \omega \rightarrow \omega$, we define the relation $\llbracket p \rrbracket_\lambda(x, y) = w$ by

$$x, \lambda(|x|), [x_0 = y] \vdash_p \mathbf{f}_0(x_0) \rightarrow w,$$

where $\mathbf{f}_0(x_0) = T_0$ is the head of p . By the determinism of p , $\llbracket p \rrbracket_\lambda$ is a partial function.

Note that the string variable x becomes the first argument of $\llbracket p \rrbracket_\lambda$. Therefore, if we want to compute a function which has a single string input by a bit-length program, the head of that program must be nullary, i.e., have zero inputs.

Definition 17. Fix a string x , natural number n , and program p , and let $\rho \vdash T \rightarrow v$ mean $x, n, \rho \vdash_p T \rightarrow v$. A *collision* is an occurrence of an inference rule of the form

$$\frac{\rho \vdash T_0 \rightarrow v \quad \rho \vdash T_1 \rightarrow w}{\rho \vdash T_0 + T_1 \rightarrow \min\{v + w, n\}}$$

such that $\min\{v + w, n\} = n$, in a derivation formed from the inference rules in Figure 4. We say a derivation of $\rho \vdash T \rightarrow v$ is *collision-free* in case it contains no collisions.

Informally, a collision occurs when an addition operation attempts to overstep (or even meet) the maximum integer bound n in a derivation. The nice thing about collision-free derivations is that they are oblivious to this bound.

Lemma 1. *If $x, n, \rho \vdash_p T \rightarrow v$ by a collision-free derivation and $n' \geq n$, then $x, n', \rho \vdash_p T \rightarrow v$.*

Proof Notice that in the inference rules of Figure 4, a collision is the only instance in which the conclusion depends on n . Hence if we take a collision-free derivation and increase n , it is still a valid derivation. ■

Remark 7. As a consequence of Lemma 1, if $\llbracket p \rrbracket_\lambda(x, y) = w$ by a collision-free derivation and $\lambda(n) \leq \lambda'(n)$ for all $n \in \omega$, then $\llbracket p \rrbracket_{\lambda'}(x, y) = w$.

Finally, let us discuss the *type* of the partial function computed $\llbracket p \rrbracket_\lambda$ by a bit-length program p , which will in general be a dependent product type. Suppose the recursive function symbol \mathbf{f}_0 in the head of program p has type $\beta \rightarrow \alpha$. Then, the type of $\llbracket p \rrbracket_\lambda$ is

$$\prod_{x:2^*} (\beta(x, \lambda(|x|)) \rightarrow \alpha(x, \lambda(|x|))).$$

We identify two important special cases. When β is the empty tuple and $\alpha = C$, then

$$\llbracket p \rrbracket_\lambda : \prod_{x:2^*} C(\lambda(|x|)).$$

On the other hand, when $\alpha = 2$ and $\beta = C$, then

$$\llbracket p \rrbracket_\lambda : \sum_{x:2^*} C(\lambda(|x|)) \rightarrow 2.$$

(Note that when both β is empty and $\alpha = 2$, then $\llbracket p \rrbracket_\lambda : 2^* \rightarrow 2$.)

5 Bit-length computability of functions

The following theorem is a restatement of the theorems $\text{TM}^{\text{ptime}} \equiv \text{CM}^{\text{rec-poly}}$ and $\text{TM}^{\text{logspace}} \equiv \text{CM}^{\text{poly}}$ from Jones (1999).

Theorem. *Let f be any function of type $2^* \rightarrow 2$. Then, the following are equivalent:*

- *f is computable in polynomial time.*
- *There is a polynomially bounded function $\lambda : \omega \rightarrow \omega$ and bit-length program p such that $\llbracket p \rrbracket_\lambda(x) = f(x)$, without collisions, for any $x \in 2^*$.*

We get an analogous result by replacing FP with FL and “bit-length program” with “tail-recursive bit-length program” throughout.

We use the following extension of this theorem. The basic observation is that the simulations of Turing machines by programs can be augmented by counting modules that keep track of the *lengths* of tapes or what comes to the same thing, unary strings. So the simulation extends to Turing machines which return unary strings as output. We sketch the proof in Appendix B.

Theorem 2. *For any function $f : 2^* \rightarrow 1^*$, the following are equivalent:*

- *f is computable in polynomial time.*
- *There is a polynomially bounded function $\lambda : \omega \rightarrow \omega$ and a bit-length program p such that $\llbracket p \rrbracket_\lambda(x) = f(x)$, without collisions, for any $x \in 2^*$.*

For any function $f : 2^ \times 1^* \rightarrow 2$ and polynomially bounded function $\pi : \omega \rightarrow \omega$, the following are equivalent:*

- *There is a polynomial-time computable function $g : 2^* \times 1^* \rightarrow 2$ such that $g(x, y) = f(x, y)$ for any string $x \in 2^*$ and $y < \pi(|x|)$.*
- *There is a polynomially bounded function $\lambda : \omega \rightarrow \omega$ and a bit-length program p such that $\llbracket p \rrbracket_\lambda(x, y) = f(x, y)$, without collisions, for any $x \in 2^*$ and $y < \pi(|x|)$.*

Moreover, we get an analogous result by replacing FP with FL and “bit-length program” with “tail-recursive bit-length program” throughout.

Note that the recursive function symbol f_0 in the head of the program p should have type C if $f : 2^* \rightarrow 1^*$ and type $C \rightarrow 2$ if $f : 2^* \times 1^* \rightarrow 2$.

Combining Theorems 1 and 2, we are able to show the fundamental property of bit-length programs; namely that a function is computable in polynomial time iff there is a pair of bit-length programs computing the length and bits of f , respectively. Similarly, a function is computable in logarithmic space iff there is a pair of tail-recursive bit-length programs computing the length and bits of f respectively.

Theorem 3. *For any function $f : 2^* \rightarrow 2^*$, f is computable in polynomial time if and only if there is a polynomially bounded function $\lambda : \omega \rightarrow \omega$ such that $|f(x)| < \lambda(|x|)$ for all $x \in 2^*$ and a pair of bit-length programs p and q such that*

$$|f(x)| = \llbracket q \rrbracket_\lambda(x)$$

without collisions, for all $x \in 2^*$, and additionally for any $i < |f(x)|$,

$$(f(x))_i = \llbracket p \rrbracket_\lambda(x, i)$$

without collisions, where $(f(x))_i$ is the i th-bit of $f(x)$.

Furthermore, we get the analogous result for logarithmic space by requiring that p and q be tail-recursive.

Proof We go through the proof in the polynomial-time case; the proof for the logarithmic-space case is verbatim, replacing “program” by “tail-recursive program” throughout.

Suppose $f : 2^* \rightarrow 2^*$ is computable in polynomial time. By Theorem 1, there exist polynomial-time computable functions $f_\ell : 2^* \rightarrow 1^*$ and $f_b : 2^* \times 1^* \rightarrow 2$ such that, for all $x \in 2^*$ $f_\ell(x) = |f(x)|$, and additionally for each $i < |f(x)|$, $f_b(x, i) = (f(x))_i$. Let π be a polynomially bounded function such that $|f(x)| < \pi(|x|)$.

By Theorem 2 applied to f_ℓ , there is a polynomially bounded function λ_1 and a bit-length program q such that $\llbracket q \rrbracket_{\lambda_1}(x) = f_\ell(x)$, without collisions, for any string x . By Theorem 2 applied to f_b , there is a polynomially bounded function λ_2 and a bit-length program p such that $\llbracket p \rrbracket_{\lambda_2}(x, i) = f_b(x, i)$, without collisions, for any string x and $i < \pi(|x|)$. Let λ be a polynomially bounded function dominating both λ_1 and λ_2 . By Remark 7 concerning collision-free computation, we can replace λ_1 and λ_2 by λ in the statements above. By definition of f_ℓ , $\llbracket q \rrbracket_\lambda(x) = |f(x)|$, for any string x . By definition of f_b and since $\pi(|x|)$ dominates $|f(x)|$, $\llbracket p \rrbracket_\lambda(x, i) = (f(x))_i$ for any string x and $i < |f(x)|$. This concludes the forward direction.

In the other direction, suppose we have a bound λ and programs p and q satisfying the desired properties. Then by Theorem 2, the function $x \mapsto |f(x)| : 2^* \rightarrow 1^*$ is computable in polynomial time. Fix a polynomially bounded function π such that $|f(x)| < \pi(|x|)$. Then by Theorem 2 again, there is a function $g : 2^* \times 1^* \rightarrow 2$, computable in polynomial time, such that $g(x, i) = \llbracket p \rrbracket_\lambda(x, i)$ for every string x and $i < \pi(|x|)$. (In particular, $g(x, i) = (f(x))_i$ for every string x and $i < |f(x)|$.) Finally, by applying Theorem 1 to $x \mapsto |f(x)|$ and g , we conclude that f is computable in polynomial time. ■

Theorem 3 suggests the following definition.

Definition 18. Let $f : 2^* \rightarrow 2^*$, $\lambda : \omega \rightarrow \omega$, p be a bit-length program whose head recursive function symbol has type $C \rightarrow 2$, and q be a bit-length program whose head recursive function symbol has type C .

Then we say (λ, p, q) *properly computes* f in case λ is increasing, $|f(x)| < \lambda(|x|)$ for all $x \in 2^*$, $|f(x)| = \llbracket q \rrbracket_\lambda(x)$, without collisions, for all $x \in 2^*$, and additionally for any $i < |f(x)|$, $(f(x))_i = \llbracket p \rrbracket_\lambda(x, i)$ without collisions.

Then, Theorem 3 can be succinctly restated as follows: membership in FP is equivalent to being properly computed by some triple (λ, p, q) with polynomially bounded λ , and membership in FL is equivalent to being properly computed by some triple (λ, p, q) with polynomially bounded λ and p and q tail-recursive.

6 Compiling RW-factorizable to bit-length programs

We now show how to compile an RW-factorizable program into a pair of bit-length programs that properly computes the same function. This consists of four program transformations:

1. The “dagger” transformation \dagger , that translates the RW-terms without W into the bit-length variant without counting modules. (So, substrings of the input string are reinterpreted as indices, `hd` is reinterpreted as `bit`, `tl` is reinterpreted as `P`, etc.)
2. The “diamond” transformation \diamond that takes an RW-term T of type W into a bit-length term T^\diamond of type 2, which detects whether T denotes a string built up from `w` or built up from `nil`.
3. A “length” transformation ℓ that takes an RW-term T of type W into a bit-length term T^ℓ of type C computing the length of T .
4. A “bit” transformation b that takes an RW-term T of type W into a bit-length term T^b of type 2, which has an extra variable `c` and computes the bit of T at `c`.

Each of these transformations requires a translation of types, then terms and values, then programs, and finally a proof of correctness.⁹ Moreover, we need to observe that each transformation preserves tail recursion. For that reason, this section is long, even though the main idea of each translation is easy to intuit:

- The dagger transformation is straightforward, essentially amounting to a “renaming” of primitives.
- Operationally, when we evaluate a RW-term T of type W into a value v according to an environment ρ , v is constructed by starting with either the empty string ε or the string $w = \rho(\mathbf{w})$ and `cons-ing` various bits in front. The term T^\diamond detects which of ε or w v is “built up” from.
- Suppose p is an RW-factorizable program that contains a recursive function symbol \mathbf{f} of type $\beta \times W \rightarrow W$. Suppose that $\rho \vdash_p \mathbf{f}(T, S) \rightarrow v$, $\rho \vdash_p \mathbf{f}(T, \mathbf{nil}) \rightarrow u$ and $\rho \vdash_p S \rightarrow w$. Then, there are two possibilities. Either $v = uw$, in which case

⁹ The map on types is only explicitly defined for the dagger transformation; the remaining three map the type W to 2, C , and 2, respectively.

- $|v| = |u| + |w|$, and
 - $v_i = w_i$ (if $i \leq |w|$), and otherwise is $u_{i-|w|}$.
- Otherwise $v = u$, in which case $|v| = |u|$ and $v_i = u_i$. Which case we are in is given to us by the diamond transformation.

The rest of the section essentially consists of making these intuitions formal.

6.1 The dagger transformation

The types in an RW-factorizable program are 2 , R , and W . The types in a bit-length program are 2 , I , and C . We first define a map from W -free types to bit-length types and then extend it to variables, recursive function symbols, terms, and programs.

Definition 19. Let $2^\dagger = 2$ and $R^\dagger = I$, and extend this to product and function types coordinate-wise.

Definition 20. For RW variables x not of type W and function symbols f not containing any type W , let $x \mapsto x^\dagger$ and $f \mapsto f^\dagger$ be a map from RW-factorizable variables and function symbols into bit-length variables and function symbols, so that for example, if x is of type τ , then x^\dagger is of type τ^\dagger . (We may assume these maps are injective.)

Definition 21. Define a map $T \mapsto T^\dagger$ from W -free RW-factorizable terms to bit-length terms by:

- If $T \equiv \text{true}$ or $T \equiv \text{false}$, then $T^\dagger \equiv T$.
- If $T \equiv x$, then $T^\dagger \equiv x^\dagger$.
- If $T \equiv \text{hd}(S)$ then $T^\dagger \equiv \text{bit}(S^\dagger)$.
- If $T \equiv \text{tl}(S)$ then $T^\dagger \equiv P(S^\dagger)$.
- If $T \equiv \text{null}(S)$ then $T^\dagger \equiv \text{null}(S^\dagger)$.
- If $T \equiv T_0 \oplus \dots \oplus T_{n-1}$ then $T^\dagger \equiv T_0^\dagger \oplus \dots \oplus T_{n-1}^\dagger$.
- If $T \equiv S[i]$ then $T^\dagger \equiv S^\dagger[i]$.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, then $T^\dagger \equiv \text{if } T_0^\dagger \text{ then } T_1^\dagger \text{ else } T_2^\dagger$.
- If $T \equiv f(S)$ then $T^\dagger \equiv f^\dagger(S^\dagger)$.

Notice that this transformation trivially preserves tail recursion.

Definition 22. For a cons-free program $p = (f_i(x_i) = T_i)_{0 \leq i \leq k}$, let the bit-length program p^\dagger be $(f_i^\dagger(x_i^\dagger) = T_i^\dagger)_{0 \leq i \leq k}$. Note that p^\dagger is a well-formed program which is tail-recursive if p is.

Finally, we extend \dagger to a map on values. Note that the type of the map $y \mapsto y^\dagger : R \rightarrow I(x)$ depends on an ‘‘ambient string’’ x of which y is a suffix.

Definition 23. For any $b \in 2$, let $b^\dagger = b$. For any string $x \in 2^*$ and any nonempty suffix y of x , let $y^\dagger = |y|$ (as a member of $I(x)$). Furthermore,

- Extend the map to tuples of data coordinate-wise, i.e., if $v = (v_0, \dots, v_{n-1})$ is a decomposition of v into atomic types, then $v^\dagger = (v_0^\dagger, \dots, v_{n-1}^\dagger)$. Note that $(u \circ v)^\dagger = u^\dagger \circ v^\dagger$, where \circ denotes concatenation.
- For an environment ρ , define the environment ρ^\dagger by $\rho(x) = v \iff \rho^\dagger(x^\dagger) = v^\dagger$. (By this, we mean that the domain of ρ^\dagger is the \dagger -image of the domain of ρ .)

Next, we prove correctness of the dagger translation.

Definition 24. Let ρ be an RW environment, i.e., a partial, finite map from RW variables to RW data. For a string $x \in 2^*$, we say that ρ is an x -environment in case for every variable x of type R , $\rho(x)$ is a suffix of x , and similarly, for every variable x of product type $\tau_0 \times \dots \times \tau_{n-1}$, $\rho(x)[i]$ is a suffix of x whenever $\tau_i = R$.

Remark 8. Suppose p is RW-factorizable program, ρ is an x -environment, T is a cons-free term, and $\rho \vdash_p T \rightarrow v$. Then, every R value that appears in the derivation of $\rho \vdash T \rightarrow v$ must also be a suffix of x .

The proof of the next lemma is postponed to Appendix A.

Lemma 2. Suppose p is a cons-free program, x is a string, ρ is an x -environment. Then,

$$\rho \vdash_p T \rightarrow v \implies x, \rho^\dagger \vdash_{\rho^\dagger} T^\dagger \rightarrow v^\dagger.$$

Example. Consider the program p defined by

```

evt(y) = if null(y) then true else if hd(y) then odt(tl y) else evt(tl y)
odt(y) = if null(y) then false, else if hd(y) then evt(tl y) else odt(tl y).
    
```

Then, `evt` and `odt` compute whether the number of occurrences of the character `true` in the string y is even or odd, respectively. The program p^\dagger is defined by

```

evt^\dagger(y^\dagger) = if null(y^\dagger) then true else if bit(y^\dagger) then odt^\dagger(P(y^\dagger)) else evt^\dagger(P(y^\dagger))
odt^\dagger(y^\dagger) = if null(y^\dagger) then false, else if bit(y^\dagger) then evt^\dagger(P(y^\dagger)) else odt^\dagger(P(y^\dagger)).
    
```

Notice that evt^\dagger and odt^\dagger compute whether the number of occurrences of `true` with index at most y^\dagger in the input string is even or odd respectively. Thus, if y is bound to a suffix y of some string x , then `evt` and `odt` behave the same on input y as evt^\dagger and odt^\dagger behave on input y^\dagger (the index $|y|$) with global input x .

6.2 The diamond transformation

These next three transformations each extend the dagger transformation from pure cons-free terms to terms of type W in different ways. In the diamond transformation, we will transform terms of type W into terms of type 2, the idea being that the transformed term will encode which of `nil` or `w` the original term is built up from.

Definition 25. For each recursive function symbol $f : \beta \times W \rightarrow W$, let f^\diamond be a recursive function symbol of type $\beta^\dagger \rightarrow 2$. (We may assume that the map $f \mapsto f^\diamond$ is injective.)

Definition 26. We define a transformation $T \mapsto T^\diamond$ from RW terms of type W to bit-length terms of type 2 as follows:

- If $T \equiv \mathbf{w}$, then $T^\diamond \equiv \text{false}$.
- If $T \equiv \text{nil}$, then $T^\diamond \equiv \text{true}$.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$ then $T^\diamond \equiv \text{if } T_0^\dagger \text{ then } T_1^\diamond \text{ else } T_2^\diamond$.
- If $T \equiv \mathbf{f}(T')$ for some $\mathbf{f} : \beta \rightarrow W$, then $T^\diamond \equiv \text{true}$.
- If $T \equiv \mathbf{f}(T', S)$ for some $\mathbf{f} : \beta \times W \rightarrow W$, then $T^\diamond \equiv \text{if } S^\diamond \text{ then true else } \mathbf{f}^\diamond((T')^\dagger)$
- If $T \equiv \text{cons}(T', S)$ then $T^\diamond \equiv S^\diamond$.

Remark 9. This transformation transforms tail-recursive terms into tail-recursive terms. It never produces a primitive call, and the only things it puts inside recursive calls are outputs of the \dagger transformation. The only way it possibly produces a non-tail-recursive term is in the case $T \equiv \mathbf{f}(T', S)$, when S^\diamond goes in an if clause. But if T is tail-recursive, S is explicit, so S^\diamond is explicit, hence contains no recursive function calls.

Definition 27. Given an RW-factorizable program p of type $R \rightarrow W$, we define the bit-length program p^\diamond by extending p^\dagger as follows. For every function symbol \mathbf{f}_i of type $\beta \times W \rightarrow W$ with definition $\mathbf{f}_i(x_i, \mathbf{w}) = T_i$, we add a new line $\mathbf{f}_i^\diamond(x_i^\dagger) = T_i^\diamond$. Note that p^\diamond is tail-recursive if p is.

We ignore function symbols of type $\beta \rightarrow W$. Also, we do not bother to specify a head, because we will only need \vdash_{p^\diamond} , not $\llbracket p^\diamond \rrbracket$.

Note that p^\diamond is a well-formed program: its recursive function symbols consist exactly of \mathbf{f}_i^\dagger for \mathbf{f}_i of type non- W and \mathbf{f}_i^\diamond for \mathbf{f}_i of type $\beta \times W \rightarrow W$. For the latter, the term $\mathbf{f}_i^\diamond(x_i^\dagger)$ is well-formed and of the same type as T_i^\diamond . The only variable that may occur in T_i^\diamond is x_i^\dagger . Notice as well that p^\diamond contains no terms of type C . Hence, when specifying a semantics, we do not need to give a natural number bound n on the upper end of the counting module.

Lemma 3. *Suppose that p is an RW-factorizable program, T is an RW term, $x, v, v', w \in 2^*$, ρ is an x -environment, and \mathbf{w} is not bound by ρ .*

Suppose that $\rho, [\mathbf{w} = w] \vdash_p T \rightarrow v$ and $\rho, [\mathbf{w} = \varepsilon] \vdash_p T \rightarrow v'$.¹⁰ Then either

- $x, \rho^\dagger \vdash_{p^\diamond} T^\diamond \rightarrow \top$ and $v = v'$, or
- $x, \rho^\dagger \vdash_{p^\diamond} T^\diamond \rightarrow \perp$ and $v = v'w$.

(The proof is postponed to Appendix A.)

In particular, Lemma 3 implies that for any x , x -environment ρ , w , and term T , if there exists a v such that $\rho, [\mathbf{w} = w] \vdash T \rightarrow v$, then there exists a boolean b such that $x, \rho^\dagger \vdash_{p^\diamond} T^\diamond \rightarrow b$. We will silently use this fact in the remainder of this paper.

¹⁰ Note that the convergence of the term T only depends on the environment ρ and not on the binding of the variable \mathbf{w} .

Example. Consider the program *id* defined by

$$\begin{aligned} \text{id}(x) &= \text{cat}(x, \text{nil}) \\ \text{cat}(x, w) &= \text{if null}(x) \text{ then } w \text{ else cons}(\text{hd } x, \text{cat}(\text{tl } x, w)), \end{aligned}$$

which computes the identity function of type $R \rightarrow W$. Both *id* and *cat* have output type W ; note that *id* always builds its output from *nil* and *cat* always builds its output from the input *w*. Then,

$$\begin{aligned} \text{id}^\diamond(x^\dagger) &= \text{true} \\ \text{cat}^\diamond(x^\dagger) &= \text{if null}(x^\dagger) \text{ then false, else if false, then true else cat}^\diamond(\text{P}(x^\dagger)). \end{aligned}$$

We can see that id^\diamond and cat^\diamond compute the always-true and always-false functions, respectively, as they ought to. Note that the dagger and diamond transformations produce C -free BL-terms. By contrast, the subsequent two transformations make crucial use of counting modules.

6.3 The length transformation

In the length transformation, we will be transforming terms of type W into terms of type C . Since the length of any W output of an RW-factorizable program is bounded by a polynomial in the length of the inputs, it makes sense to try to capture the length within a counting module.

Definition 28. For each recursive function symbol f of type $\beta \times W \rightarrow W$ or $\beta \rightarrow W$, let f^ℓ be a recursive function symbol of type $\beta^\dagger \rightarrow C$. (We may assume that the map $f \mapsto f^\ell$ is injective.)

Definition 29. We define a transformation $T \mapsto T^\ell$ from RW-terms of type W to bit-length terms of type C as follows:

- If $T \equiv w$ or $T \equiv \text{nil}$, then $T^\ell \equiv 0$.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, then $T^\ell \equiv \text{if } T_0^\dagger \text{ then } T_1^\ell \text{ else } T_2^\ell$.
- If $T \equiv f(T')$, then $T^\ell \equiv f^\ell((T')^\dagger)$.
- If $T \equiv f(T', S)$, where S has type W , then

$$T^\ell \equiv \text{if } f^\diamond((T')^\dagger) \text{ then } f^\ell((T')^\dagger) \text{ else } f^\ell((T')^\dagger) + S^\ell.$$

- If $T \equiv \text{cons}(T', S)$, then $T^\ell \equiv 1 + S^\ell$.

Remark 10. This transformation does *not* preserve tail recursion, the problem being $f^\ell((T')^\dagger) + S^\ell$, where the recursive function symbol f^ℓ occurs within the primitive $+$. This can be fixed in a couple of ways. For one, we can observe that if T is tail-recursive then T^ℓ is *linear recursive*, which can always be transformed into an equivalent tail-recursive term (Greibach, 1975).

But there is an easier fix in this case: just give each f^ℓ an extra input of type c , the idea being that the new $f^\ell(T, c)$ is the old $f^\ell(T) + c$. We have to change each recursive definition $f^\ell(x^\dagger) = T^\ell$ to $f^\ell(x^\dagger, c) = T^\ell + c$. Then $f^\ell((T')^\dagger) + S^\ell$ becomes $f^\ell((T')^\dagger, S^\ell)$, and this solves our problem.

Definition 30. Given an RW-factorizable program p of type $R \rightarrow W$, we define the bit-length program p^ℓ by extending p^\diamond by adding the following lines for each recursive function symbol of output type W :

| | |
|--|------------------------------------|
| For each line in p of the following form, | add the line |
| $f_i(x_i) = T_i (f : \beta \rightarrow W)$ | $f_i^\ell(x_i^\dagger) = T_i^\ell$ |
| $f_i(x_i, w) = T_i (f : \beta \times W \rightarrow W)$ | $f_i^\ell(x_i^\dagger) = T_i^\ell$ |

Finally, add a new head $h = f_0^\ell(\text{max})$. The resulting program is p^ℓ ; it is tail-recursive if p is.

Let us check that p^ℓ is a well-defined bit-length program. In addition to p^\diamond , it contains recursive function symbols f_i^ℓ for f_i of output type W . If the only variables that may occur in T_i are x_i and w , then the only variable that may occur in T_i^ℓ is x_i^\dagger , and the only recursive function symbols that may occur are the recursive function symbols listed above. Finally, the terms in each line are well-formed and the types of each recursive functions symbol and its definition agree. Finally, notice that since p has type $R \rightarrow W$, p^ℓ has type C .

Definition 31. For a value v of type W , let $v^\ell = |v|$.

Lemma 4. For every RW-program p of type $R \rightarrow W$, RW term T from p of type W , strings x, v , and w , natural number $n > |v|$, and x -environment ρ binding w to w :

$$\rho \vdash_p T \rightarrow v \implies x, n, \rho^\dagger \vdash_{p^\ell} T^\ell \rightarrow v^\ell - \delta,$$

where $\delta = 0$ if $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$ or $\delta = |w|$ if $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \perp$. Moreover, the derivation of the right-hand side is collision-free.

(The proof is postponed to Appendix A.)

Theorem 4. Suppose that p is an RW-factorizable program of type $R \rightarrow W$ and suppose that $\lambda : \omega \rightarrow \omega$ satisfies

$$(\forall x, w \in 2^*) \llbracket p \rrbracket(x) = w \implies |w| < \lambda(|x|).$$

Then,

$$(\forall x, w \in 2^*) \llbracket p \rrbracket(x) = w \implies \llbracket p^\ell \rrbracket_\lambda(x) = |w|;$$

moreover, without collisions.

Proof Fix x and suppose that $\llbracket p \rrbracket(x) = w$. Suppose $f_0(x_0)$ is the head of p . Then, $[x_0 = x] \vdash_p f_0(x_0) \rightarrow w$, so by Lemma 4, since $\lambda(|x|) > w$,

$$x, \lambda(|x|), [x_0^\dagger = x^\dagger] \vdash_{p^\ell} f_0^\ell(x_0^\dagger) \rightarrow |w|$$

by a collision-free derivation. Therefore, since $x^\dagger = |x|$ (as a member of $I(x)$), and since $|x|$ is the denotation of max , we have

$$x, \lambda(|x|) \vdash_{p^\ell} f_0^\ell(\text{max}) \rightarrow |w|.$$

Since $f_0(\max)$ is precisely the head of p^ℓ , we have $\llbracket p^\ell \rrbracket_\lambda(x) = |w|$, which is what we wanted to prove. ■

Examples. Consider again the program *id* given by

$$\begin{aligned} \text{id}(x) &= \text{cat}(x, \text{nil}) \\ \text{cat}(x, w) &= \text{if } \text{null}(x) \text{ then } w \text{ else } \text{cons}(\text{hd } x, \text{cat}(\text{tl } x, w)), \end{aligned}$$

Then,

$$\begin{aligned} \text{id}^\ell(x^\dagger) &= \text{if } \text{cat}^\diamond(x^\dagger) \text{ then } \text{cat}^\ell(x^\dagger) \text{ else } \text{cat}^\ell(x^\dagger) + 0 \\ \text{cat}^\ell(x^\dagger) &= \text{if } \text{null}(x^\dagger) \text{ then } 0 \text{ else } 1 + (\text{if } \text{cat}^\diamond(P(x^\dagger)) \text{ then } \text{cat}^\ell(P(x^\dagger)) \text{ else } \\ &\quad \text{cat}^\ell(P(x^\dagger)) + 0), \end{aligned}$$

which is semantically equivalent to

$$\begin{aligned} \text{id}^\ell(x^\dagger) &= \text{cat}^\ell(x^\dagger) \\ \text{cat}^\ell(x^\dagger) &= \text{if } \text{null}(x^\dagger) \text{ then } 0 \text{ else } 1 + \text{cat}^\ell(P(x^\dagger)). \end{aligned}$$

It's now easy to see that for sufficiently large λ , $\llbracket \text{id}^\ell \rrbracket_\lambda(x) = |x|$, as it ought to. (Recall that the head of id^ℓ is $\text{id}^\ell(\max)$.)

Consider the following program *leap* based off of modifying the program from line (1) in Section 1 to string data:

$$\begin{aligned} \text{leap}(x) &= f(x, \text{nil}) \\ f(x, w) &= \text{if } \text{null}(x) \text{ then } \text{cons}(\text{true}, w) \text{ else } f(\text{tl } x, f(\text{tl } x, w)). \end{aligned}$$

Then for any string x , $\llbracket \text{leap} \rrbracket(x)$ is a $2^{|x|}$ -length string of only the character `true`. Now,

$$\begin{aligned} \text{leap}^\ell(x^\dagger) &= \text{if } f^\diamond(x^\dagger) \text{ then } f^\ell(x^\dagger) \text{ else } f^\ell(x^\dagger) + 0 \\ f^\ell(x^\dagger) &= \text{if } \text{null}(x^\dagger) \text{ then } 1 + 0 \text{ else} \\ &\quad \text{if } f^\diamond(Px^\dagger) \text{ then } f^\ell(Px^\dagger) \text{ else } f^\ell(Px^\dagger) + \\ &\quad (\text{if } f^\diamond(Px^\dagger) \text{ then } f^\ell(Px^\dagger) \text{ else } f^\ell(Px^\dagger) + 0). \end{aligned}$$

This is semantically equivalent to

$$\begin{aligned} \text{leap}^\ell(x^\dagger) &= f^\ell(x^\dagger) \\ f^\ell(x^\dagger) &= \text{if } \text{null}(x^\dagger) \text{ then } 1 \text{ else } f^\ell(Px^\dagger) + f^\ell(Px^\dagger), \end{aligned}$$

which we can see correctly computes the base-2 exponential of the length of the input string (given a sufficiently large counting module).

6.4 The bit transformation

In the bit transformation, we will transform terms of type W into terms of type 2, with an additional type- C input. The idea is that the transformed term encodes the *bits* of the original W -term.

Definition 32. For each function symbol f of type $\beta \times W \rightarrow W$ or $\beta \rightarrow W$, let f^b be a recursive function symbol of type $\beta^\dagger \times C \rightarrow 2$. (We may assume that the map $f \mapsto f^b$ is injective.)

Definition 33. We define a transformation $T \mapsto T^b$ from RW terms of type W to bit-length terms of type 2 as follows. Here c is a fixed variable of type C .

- If $T \equiv w$ or $T \equiv \text{nil}$, then $T^b \equiv \text{false}$.¹¹
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, then $T^b \equiv \text{if } T_0^\dagger \text{ then } T_1^b \text{ else } T_2^b$.
- If $T \equiv f(T')$, then $T^b \equiv f^b((T')^\dagger, c)$
- If $T \equiv f(T', S)$, where S has type W , then

$$T^b \equiv \text{if } f^\diamond((T')^\dagger) \text{ then } f^b((T')^\dagger, c) \text{ else if } c \leq S^\ell \text{ then } S^b \text{ else } f^b((T')^\dagger, c - S^\ell).$$
- If $T \equiv \text{cons}(T', S)$, then $T^b \equiv \text{if } c \leq S^\ell \text{ then } S^b \text{ else } (T')^\dagger$,

Notice that the only variables that may occur in T^b are c and x^\dagger for x which occur in T , and the only recursive function symbols that may occur are f^b for f of output type W , and f^\dagger for f of output type non- W , which occur in T .

Remark 11. Not only does this transformation preserve tail-recursive terms, but every term it produces is tail-recursive (as long as we regard previously defined subterms like S^ℓ and $f^\diamond((T')^\dagger)$ as explicit, as usual). No matter what T is, no f^b will occur inside anything other than a then- or else-clause in T^b .

Definition 34. Given an RW-factorizable program p of type $R \rightarrow W$, we define the bit-length program p^b by extending p^ℓ as follows

| For each line of the following form, | add the following |
|--|---------------------------------|
| $f_i(x_i) = T_i (f_i : \beta \rightarrow W)$ | $f_i^b(x_i^\dagger, c) = T_i^b$ |
| $f_i(x_i, w) = T_i (f_i : \beta \times W \rightarrow W)$ | $f_i^b(x_i^\dagger, c) = T_i^b$ |

Finally, add a new head $h(c) = f_0^b(\text{max}, c)$. The resulting program is p^b . It is tail-recursive if p is.

Let us check that for any RW-factorizable program p of type $R \rightarrow W$, p^b is a well-defined bit-length program. It contains the recursive function symbols f^b for f in p of type W , and f^\dagger for f in p of type non- W . Since the only variable that appears in T is x_i , the only variables that may appear in T^b are x_i^\dagger and c . Finally, the types of $f_i^b(x_i^\dagger, c)$ and T_i agree.

Notice that p^b contains p^ℓ , and hence p^\diamond and p^\dagger , as “subprograms.” Hence, the semantics of p^b extends the semantics of the previous programs. We now prove its correctness; the proof is postponed to Appendix A.

Lemma 5. For every RW program p of type $R \rightarrow W$, RW term T from p of type W , strings x, w and v , natural numbers $n > |v|$ and $1 \leq c \leq |v| - \delta$, and x -environment ρ binding w to w ,

¹¹ This definition will not be important in practice; rather, it is purely so that the transformation is always well-defined. We could have equally well set $T^b \equiv \text{true}$.

$$\rho \vdash_p T \rightarrow v \implies x, n, \rho^\dagger, [c = c] \vdash_{p^b} T^b \rightarrow v_{c+\delta},$$

where $\delta = 0$ if $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$ or $\delta = |w|$ if $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \perp$.¹² Moreover, the computation on the right-hand side is without collisions.

Theorem 5. Suppose that p is an RW-factorizable program of type $R \rightarrow W$ and suppose that $\lambda : \omega \rightarrow \omega$ satisfies

$$(\forall x, w \in 2^*) \llbracket p \rrbracket(x) = w \implies |w| < \lambda(|x|).$$

Then,

$$(\forall x, w \in 2^*) \llbracket p \rrbracket(x) = w \implies (\forall c \leq |w|) \llbracket p^b \rrbracket_\lambda(x, c) = w_c,$$

without collisions.

Proof Fix x, v , and c such that $\llbracket p \rrbracket(x) = v$ and $1 \leq c \leq |v|$. Suppose that $f_0(x_0)$ is the head of p , so that $[x_0 = x] \vdash_p f_0(x_0) \rightarrow v$. Therefore, $[x_0 = x, w = \varepsilon] \vdash_p f_0(x_0) \rightarrow v$. By Lemma 5, since $\lambda(|x|) > v$,

$$x, \lambda(|x|), [x_0^\dagger = x^\dagger, c = c] \vdash_{p^b} f_0^b(x_0^\dagger, c) \rightarrow v_c,$$

without collisions. (Notice that $\delta = 0$ as w is bound to ε .)

Since $x, \lambda(|x|) \vdash_{p^b} \max \rightarrow |x|$ and $|x| = x^\dagger$, we have

$$x, \lambda(|x|), [c = c] \vdash_{p^b} f_0^b(\max, c) \rightarrow v_c.$$

But since the head of $h(c)$ is defined to be $f_0^b(\max, c)$, this implies $\llbracket p^b \rrbracket_\lambda(x, c) = v_c$, which is exactly what we wanted to show. ■

Theorems 4 and 5 immediately imply the following result, the statement of correctness for the translation from RW-factorizable to bit-length programs.

Theorem 6. Let $f : 2^* \rightarrow 2^*$ be a function, $\lambda : \omega \rightarrow \omega$ satisfy $\lambda(|x|) > |f(x)|$, and p be an RW-factorizable program of type $R \rightarrow W$ computing f . Then (λ, p^ℓ, p^b) properly computes f .

Example. Given the program id defined by

$$\begin{aligned} id(x) &= \text{cat}(x, \text{nil}) \\ \text{cat}(x, w) &= \text{if } \text{null}(x) \text{ then } w \text{ else } \text{cons}(\text{hd } x, \text{cat}(\text{tl } x, w)), \end{aligned}$$

id^b is defined by

$$\begin{aligned} id^b(x^\dagger, c) &= \text{if } \text{cat}^\diamond(x^\dagger) \text{ then } \text{cat}^b(x^\dagger, c) \text{ else if } c \leq 0 \text{ then false,} \\ &\quad \text{else } \text{cat}^b(x^\dagger, c - 0) \\ \text{cat}^b(x^\dagger, c) &= \text{if } \text{null}(x^\dagger) \text{ then false, else if } c \leq S^\ell \text{ then } S^b \text{ else } \text{bit}(x^\dagger), \end{aligned}$$

¹² Recall that $v_{c+\delta}$ refers to a particular bit of v , where $v = v_{|v|} \dots v_2 v_1$.

where $S \equiv \text{cat}(\text{tl } x, w)$. We know that S^ℓ correctly computes the length of $\text{tl}(x)$, which is (abusing notation) $x^\dagger - 1$. Moreover,

$S^b \equiv \text{if } \text{cat}^\diamond(\text{Px}^\dagger) \text{ then } \text{cat}^b(\text{Px}^\dagger, c) \text{ else if } c \leq 0 \text{ then false, else } \text{cat}^b(\text{Px}^\dagger, c - 0)$.

Recall that cat^\diamond computes the always-false function. Combining this all into semantically equivalent and legible pseudocode, we get

$$\begin{aligned} \text{id}^b(x^\dagger, c) &= \text{if } c \leq 0 \text{ then false, else } \text{cat}^b(x^\dagger, c) \\ \text{cat}^b(x^\dagger, c) &= \text{if } \text{null}(x^\dagger) \vee c \leq 0 \text{ then false, else} \\ &\quad \text{if } c \leq x^\dagger - 1 \text{ then } \text{cat}(\text{Px}^\dagger, c) \text{ else } \text{bit}(x^\dagger). \end{aligned}$$

We can see that $\text{cat}^b(x^\dagger, c)$ computes the bit indexed by c : operationally, it decrements the index x^\dagger until it is equal to the counting module c , then spits out the bit indexed by x^\dagger . Hence, $\llbracket \text{id}^b \rrbracket_\lambda(x, c)$ computes the bit x_c for sufficiently large λ and appropriate values of c .

Observe that for sufficiently large λ , $(\lambda, \text{id}^\ell, \text{id}^b)$ correctly bit-length computes the identity function, as $\llbracket \text{id}^\ell \rrbracket_\lambda(x) = |x|$ and $\llbracket \text{id}^b \rrbracket_\lambda(x, c) = x_c$.

7 Compiling BL- to RW-factorizable programs

In this section, we show how to compile a pair of bit-length programs into a single RW-factorizable program, thus establishing extensional equivalence for these two notions of computability, at least for total functions. The core of this section consists of two transformations:

- The more significant of these is a transformation \ddagger , which acts as a sort of inverse to \dagger : it eliminates indices in favor of string suffixes, i.e., R -data, and counting modules in favor of tuples of R -data. (More precisely, this is a *family* of transformations parameterized by how many R values we need to encode a single counting module.) This process transforms every bit-length program into a cons-free program with a “global input variable.”
- The simpler transformation takes a cons-free program with a global input variable into one without. We do this in most naive way possible, simply by passing the global input as an additional parameter to every recursive call.

The former transformation uses a well-known trick of cons-free programming, namely that we can simulate a counting module of size polynomial in the input length by a tuple of suffixes of the input. If we just think of a suffix of the input as encoding its length (i.e., forgetting about its bits), then we can identify it with some single digit $\{0, 1, \dots, n\}$, where n is the length of the input. Therefore, we can identify a k -tuple of suffixes with some k -digit number in base $n + 1$, which we can identify in turn with some number less than $(n + 1)^k$. To implement the full data type of a counting module, it simply suffices to implement the fixed-width arithmetic and comparison operations using cons-free programs. In this way, we can replace polynomially bounded counting modules with fixed-width tuples of R values.

A complication is the fact that our language does not accommodate nested tuples. For example, suppose that we had to replace every counting module of type C by three R values

of type $R \times R \times R$. Then if we had a tuple of type $I \times C \times 2$, it *should* be replaced by a tuple of type $(R, R^3, 2)$, but since we don't have nested tuples, we are required to "flatten it out" into a tuple of type $(R, R, R, R, 2)$. This makes indexing harder: to get the element of R^3 encoding the counting module, we have to extract the middle three coordinates.

7.1 Programs with global input

Bit-length programs have a global input value, which is the x in the judgment $x, n, \rho \vdash_p T \rightarrow v$. RW-factorizable programs, and cons-free programs in particular, do not. It will be convenient to first transform bit-length programs into cons-free programs with some global input variable `in`, and then eliminate it, instead of trying to cram both in one transformation.

Definition 35. A RW term *term with global input* is obtained by extending the formation rules of Figure 1 by the additional axiom $\text{in} : R$. A RW-factorizable program with global input is defined like an ordinary RW-factorizable program, except that the terms T_i are RW-terms with global input.

The semantics relation \vdash has the form $x, \rho \vdash T \rightarrow v$ (note the additional x argument on the left-hand side) and is defined by extending the rules of Figure 2 by the axiom $x, \rho \vdash \text{in} \rightarrow x$. If $\mathfrak{f}_0(x_0)$ is the head of a program p , then we define

$$\llbracket p \rrbracket(x, y) = v \iff x, [x_0 = y] \vdash_p \mathfrak{f}_0(x_0) \rightarrow v.$$

An RW-factorizable program with global input is *cons-free* if it contains no term of type W .

We can easily transform a RW program with global input into one without. The basic idea is to pass the global input explicitly into each recursive function symbol as an extra argument which never gets modified. The remainder of this subsection consists of making this intuition formal. Moreover, we will restrict our attention to cons-free programs, because that is the only case we need.

Definition 36. For any RW product term α , let α_R be $\alpha \times R$. For any RW function term $\rho = \beta \rightarrow \alpha$, let ρ_R be $\beta_R \rightarrow \alpha$. Define an injection $\mathbf{x} \mapsto \mathbf{x}_R, \mathfrak{f} \mapsto \mathfrak{f}_R$ from variables and recursive function symbols of type α and ρ to type α_R and ρ_R , respectively.

The idea of the next transformation is we replace the variable \mathbf{x} with the variable \mathbf{x}_R . The last coordinate of \mathbf{x}_R stores the "global input" which gets passed around to all the recursive functions in the program, and the rest of the variable stores the "original variable" \mathbf{x} .

Definition 37. Suppose T is a term in which only the variable \mathbf{x} may occur. (Call this an *x-term* for brevity.) Let n be the length of the type of \mathbf{x} . Define the map $T \mapsto T_R^{\mathbf{x}}$ from cons-free terms with global input to cons-free terms as follows (for brevity, we omit the superscript \mathbf{x}):

- If $T \equiv \mathbf{x}$, then $T_R \equiv \mathbf{x}_R[0, n - 1]$.
- If $T \equiv \text{in}$, then $T_R \equiv \mathbf{x}_R[n]$.

- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2 \text{ then } T_R \equiv \text{if } (T_0)_R \text{ then } (T_1)_R \text{ else } (T_2)_R$.
- If $T \equiv \text{hd}(S)$, $\text{tl}(S)$, or $\text{null}(S)$, then $T_R \equiv \text{hd}(S_R)$, $\text{tl}(S_R)$, or $\text{null}(S_R)$ respectively.
- If $T \equiv T_0 \oplus \dots \oplus T_{n-1}$, then $T_R \equiv (T_0)_R \oplus \dots \oplus (T_{n-1})_R$.
- If $T \equiv S[i, j]$, then $T_R \equiv S_R[i, j]$.
- If $T \equiv \mathbf{f}(S)$, then $T_R \equiv \mathbf{f}_R(S_R \oplus \mathbf{x}_R[n])$.

Notice that for any RW-term with global input T containing only \mathbf{x} , then $T_R^{\mathbf{x}}$ is a well-formed RW-term of the same type of T . Moreover, \mathbf{x}_R is the only variable which may occur in $T_R^{\mathbf{x}}$. This transformation is easily seen to preserve tail recursion.

Definition 38. Suppose the cons-free program with global input p consists of the lines $\mathbf{f}_i(\mathbf{x}_i) = T_i$ for $0 \leq i \leq k$. Then define the program p_R to consist of lines $(\mathbf{f}_i)_R((\mathbf{x}_i)_R) = (T_i)_R^{\mathbf{x}_i}$ for $0 \leq i \leq k$.

Note that p_R is a well-defined program. It contains the recursive function symbols $(\mathbf{f}_i)_R$ for each \mathbf{f}_i from p . The type of $(\mathbf{x}_i)_R$ is the input type of $(\mathbf{f}_i)_R$. The type of $(T_i)_R^{\mathbf{x}_i}$ is the type of T_i , which is the output type of \mathbf{f}_i , which is the output type of $(\mathbf{f}_i)_R$. The only variable that may occur in $(T_i)_R$ is $(\mathbf{x}_i)_R$, and the only recursive function symbols that may occur are among $((\mathbf{f}_0)_R, \dots, (\mathbf{f}_k)_R)$.

In the next lemma, recall that if (u_0, \dots, u_{n-1}) is a decomposition of the value u into atomic values, then by $u \circ x$ we mean (u_0, \dots, u_{n-1}, x) . Its proof is postponed to Appendix A.

Lemma 6. For any cons-free program with global input p , values x u , and v , variable \mathbf{x} , and \mathbf{x} -term T ,

$$x, [\mathbf{x} = u] \vdash T \rightarrow v \implies [\mathbf{x}_R = u \circ x] \vdash T_R^{\mathbf{x}} \rightarrow v.$$

Hence,

Theorem 7. For any cons-free program with global input p with nullary input, string x and value v , if $\llbracket p \rrbracket(x) = v$, then $\llbracket p_R \rrbracket(x) = v$.

Proof Let $\mathbf{f}_0(\mathbf{x})$ be the head of p . Since p has nullary input, \mathbf{x} is a variable of type the empty product, and \mathbf{x}_R is a variable of type R . If $\llbracket p \rrbracket(x) = v$, then $x \vdash \mathbf{f}_0 \rightarrow v$. By Lemma 6, $[\mathbf{x}_R = x] \vdash (\mathbf{f}_0)_R((\mathbf{x}_0)_R) \rightarrow v$. But since $(\mathbf{f}_0)_R((\mathbf{x}_0)_R)$ is the head of p_R , $\llbracket p_R \rrbracket(x) = v$. ■

7.2 Eliminating indices and counting modules

In this subsection, we show how to compile any bit-length program into a cons-free program with global input. We fix a natural number $k \geq 1$ which is the number of copies of R we want to replace every copy of C by. The transformation \ddagger defined in this section should be understood as parameterized by this number k . First we will define a map on types, variables, and terms, then values, then programs, and we conclude with a proof of correctness.

The main idea: Bit-length programs contain two types of data that RW-factorizable programs do not: indices and counting modules. We replace indices by single R values and counting modules by k -tuples of R values.

The former correspondence is straightforward: if the input string x has length n , then its bits are $x_n \dots x_2x_1$. R values range over suffixes of x . We encode the index i by the R value $x_i \dots x_2x_1$. The “dummy index” 0 is encoded by the empty string. With respect to this encoding, the index primitives `bit`, `P`, `null`, and `max` correspond to `hd`, `tl`, `null`, and `in` (the global input variable) respectively. The only index primitive that needs to be programmed is `min`, which can be replaced by `f(in)`, where

$$f(x) \equiv \text{if null}(x) \text{ then } x \text{ else } f(\text{tl}(x)).$$

Let zero_l be the term `f(in)`.

If we ignore the characters in an R -string, it simply encodes a *length* ranging from 0 to $n + 1$; i.e., a single digit base $n + 1$. Hence, a k -tuple of R -data encodes a k -digit number base $n + 1$, or alternatively, as a natural number less than $(n + 1)^k$. Therefore, k -tuples of R data can be used to simulate counting modules which are polynomially bounded in the length of the input.¹³ The counting modules primitives `+`, `-`, and `≤` can be replaced by cons-free programs `add`, `minus`, and `less`, which simulate the corresponding primitives on k -tuples of R values by mimicking any common algorithm for fixed-width arithmetic.¹⁴ The cons-free constants 0 and 1 can similarly be replaced with programs `zero` and `one`, which compute the constant- $(0, \dots, 0, 0)$ and constant- $(0, \dots, 0, 1)$ functions, respectively.¹⁵

We assume the existence of these five programs without constructing them and note that they can always be made tail-recursive. In fact, all the recursion we need can be collected into a few important subroutines. Thinking of R values as single digits, we can name the largest digit (by the global input variable `in`) and we can decrement digits (by `tl`). Using these primitives, we can define simple tail-recursive subroutines incrementing a digit, comparing two digits, and naming the digit 0. Then, the programs `add`, `minus`, `less`, `zero`, and `one` are *explicit* in these subroutines; i.e., can be defined in terms of them using no additional recursion.

Definition 39. Let $2^\ddagger = 2$, $I^\ddagger = R$, and $C^\ddagger = R^k$. For every bit-length product type α , define the RW product type α^\ddagger by replacing every copy of I by R , replacing every copy of C by R^k , and flattening. For example:

$$\begin{aligned} C^\ddagger &= R^k \\ (2 \times I \times C)^\ddagger &= 2 \times R^{k+1} \\ (2 \times C \times I \times 2 \times C \times C)^\ddagger &= 2 \times R^{k+1} \times 2 \times R^{2k}. \end{aligned}$$

Extend this to a map on function types by $(\beta \rightarrow \alpha)^\ddagger = \beta^\ddagger \rightarrow \alpha^\ddagger$.

¹³ Indeed, in the original formulation of Jones (1999), counting modules were introduced as syntactic sugar for fixed-width tuples of R data.

¹⁴ When a sum or difference would overflow its bounds of $(n + 1)^k - 1$ or 0, add “tops out” at $(n + 1)^k - 1$ and minus “bottoms out” at 0, respectively, thus mimicking the behavior of `+` and `-` on counting modules.

¹⁵ When we say, e.g., constant $(0, \dots, 0, 0)$ -function, the values in the tuple are understood to be R -data.

Finally, for each product type α , define the map $s_\alpha : |\alpha^\dagger| \rightarrow |\alpha|$ by mapping each coordinate of α^\dagger to the coordinate “which it comes from” in α . For example, $s_C : k \rightarrow 1$ is defined by $s_C(i) = 0$ for all $i \in k$; $s_{2 \times I \times C} : k + 2 \rightarrow 3$ and $s_{2 \times I \times C}(i)$ is 0 if $i = 0$, 1 if $i = 1$, and 2 otherwise; and

$$s_{2 \times C \times I \times 2 \times C \times C} : 3k + 2 \rightarrow 6$$

is defined by

$$s_{2 \times C \times I \times 2 \times C \times C}(i) = \begin{cases} 0 & \text{if } i = 0 \\ 1 & \text{if } 1 \leq i \leq k \\ 2 & \text{if } i = k + 1 \\ 3 & \text{if } i = k + 2 \\ 4 & \text{if } k + 3 \leq i \leq 2k + 2 \\ 5 & \text{otherwise} \end{cases}$$

Notice that each such map s_α is a monotone (non-decreasing) surjection.

Definition 40. Fix an injection $x \mapsto x^\dagger$ from bit-length variables to RW variables such that if x has type α , x^\dagger has type α^\dagger . Fix an injection $f \mapsto f^\dagger$ from bit-length recursive function symbols of type ρ to RW function symbols of type ρ^\dagger , for every bit-length function type ρ .

In the next definition, `in` is a fixed variable of type R naming the global input, `zero` and `one` are fixed function symbols of type C^\dagger , `add` and `minus` are fixed function symbols of type $(C \times C \rightarrow C)^\dagger$, and `less` is a fixed function symbol of type $(C \times C \rightarrow 2)^\dagger$.

Definition 41. Define a map $T \mapsto T^\dagger$ from bit-length terms to cons-free terms with global input by:

- If $T \equiv \text{true}$ or $T \equiv \text{false}$, then $T^\dagger \equiv T$.
- If $T \equiv x$, a variable of type α , then $T^\dagger \equiv x^\dagger$.
- If $T \equiv 0$ then $T^\dagger \equiv \text{zero}$.
- If $T \equiv 1$ then $T^\dagger \equiv \text{one}$.
- If $T \equiv T_0 + T_1$ then $T^\dagger \equiv \text{add}(T_0^\dagger, T_1^\dagger)$.
- If $T \equiv T_0 - T_1$ then $T^\dagger \equiv \text{minus}(T_0^\dagger, T_1^\dagger)$.
- If $T \equiv T_0 \leq T_1$, then $T^\dagger \equiv \text{less}(T_0^\dagger, T_1^\dagger)$.
- If $T \equiv \text{min}$ then $T^\dagger \equiv \text{zero}_I$.
- If $T \equiv \text{max}$ then $T^\dagger \equiv \text{in}$.
- If $T \equiv P(S)$ then $T^\dagger \equiv \text{tl}(S^\dagger)$.
- If $T \equiv \text{null}(S)$ then $T^\dagger \equiv \text{null}(S^\dagger)$.
- If $T \equiv \text{bit}(S)$ then $T^\dagger \equiv \text{hd}(S^\dagger)$.
- If $T \equiv T_0 \oplus \dots \oplus T_{n-1}$ then $T^\dagger \equiv T_0^\dagger \oplus \dots \oplus T_{n-1}^\dagger$.
- If $T \equiv S[i, j]$, let m be the length of S and n the length of S^\dagger . Then, $T^\dagger \equiv S^\dagger[l, j]$, where the interval $[l, j] \subseteq m$ is the s_α -pre-image of $[i, j] \subseteq n$.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$ then $T^\dagger \equiv \text{if } T_0^\dagger \text{ then } T_1^\dagger \text{ else } T_2^\dagger$.
- If $T \equiv f(S)$, then $T \equiv f^\dagger(S^\dagger)$.

Now we define a map on values. For this definition, recall the bijection

$$f : \underbrace{n \times \dots \times n}_{k \text{ copies}} \rightarrow n^k$$

defined by $(d_0, \dots, d_{k-1}) \mapsto \sum_{i < k} d_i n^i$. Intuitively, this identifies k -digit numerals base n with the numbers they denote.

Remark 12. As long as zero, one, add, minus, and less are implemented by tail-recursive programs, this transformation is easily seen to preserve tail recursion.

Definition 42. For any string $x \in 2^*$, define the map $v \mapsto v^\ddagger$ as follows. If $v \in 2$, then $v_x^\ddagger = v$. If $v \in I(x)$, then v_x^\ddagger is the suffix of x of length $|v|$. If $v \in C((|x| + 1)^k)$, then v_x^\ddagger is the unique k -tuple of suffixes (s_0, \dots, s_{k-1}) of x such that

$$v = \sum_{i < k} |s_i| (|x| + 1)^i.$$

If (v_0, \dots, v_{n-1}) is a decomposition of v into atomic types, then $v^\ddagger = v_0^\ddagger \circ \dots \circ v_{n-1}^\ddagger$. Notice that for any values u and v , $(u \circ v)^\ddagger = u^\ddagger \circ v^\ddagger$.

Definition 43. Given a bit-length program $p = (\mathbf{f}_i(x_i) = T_i)$, the cons-free program with global input p^\ddagger is defined to be $(\mathbf{f}_i^\ddagger(x_i^\ddagger) = T_i^\ddagger)$.

This program is well-defined: if T_i is an x_i -term, then T_i^\ddagger is an x_i^\ddagger -term. The only recursive function symbols that may occur in T_i^\ddagger , besides the previously defined zero_I , zero, one, add, minus and less, are the \mathbf{f}_i^\ddagger . Moreover, it is tail-recursive if p is.

The proof of the next result is postponed to Appendix A.

Lemma 7. For every bit-length program p , bit-length term T from p , string x , x -environment ρ , and value v ,

$$x, (|x| + 1)^k, \rho \vdash_p T \rightarrow v \implies x, \rho^\ddagger \vdash_{p^\ddagger} T^\ddagger \rightarrow v^\ddagger.$$

Lemma 7 and Theorem 7 have the following corollary. Let p_R^\ddagger be $(p^\ddagger)_R$.

Theorem 8. For any bit-length program p and polynomially bounded $\lambda : \omega \rightarrow \omega$, string x and value v , if $\llbracket p \rrbracket_\lambda(x) = v$ without collisions, then $\llbracket p_R^\ddagger \rrbracket(x) = v^\ddagger$.

Proof Choose k such that $\lambda(n) < (n + 1)^k$ for all n and define $\mu : \omega \rightarrow \omega$ by $\mu(n) = (n + 1)^k$. Suppose that $\llbracket p \rrbracket_\lambda(x) = v$ without collisions. Then, $\llbracket p \rrbracket_\mu(x) = v$ by monotonicity of collision-free computation. By Lemma 7 and the definition of p^\ddagger , $\llbracket p^\ddagger \rrbracket(x) = v^\ddagger$. By Theorem 7, $\llbracket p_R^\ddagger \rrbracket(x) = v^\ddagger$. ■

Example. The following program of type $C \times C \rightarrow C$ multiplies its two inputs. (Or rather, the two coordinates of its single input—recall these are $c[0]$ and $c[1]$.)

$$f(c) = \text{if } c[0] \leq 0 \text{ then } 0 \text{ else } c[1] + f(c[0] - 1, c[1]).$$

Suppose that $k = 2$, so we replace each occurrence of C by two copies of R . Then, the type of c^\ddagger is R^4 , where the former and latter two copies of R correspond to $c[0]$ and $c[1]$, respectively. Hence,

$$f^\ddagger(c^\ddagger) = \text{if less}(c^\ddagger[0, 1], \text{zero}) \text{ then zero else add}(c^\ddagger[2, 3], f^\ddagger(\text{minus}(c^\ddagger[0, 1], \text{one}), c^\ddagger[2, 3])).$$

7.3 Building type- W output

So far we have shown that we can convert any bit-length program p into an equivalent cons-free RW-factorizable program p_R^\ddagger . Given a polynomially bounded function λ and bit-length programs p and q such that (λ, p, q) properly computes a function f , how do we use the cons-free programs q_R^\ddagger and p_R^\ddagger to form an RW-factorizable program computing f ? (As above, \ddagger is dependent on a parameter k , which we choose large enough so that $(n + 1)^k$ dominates λ .)

As usual, the basic idea is straightforward. We compute the length of the output string using q_R^\ddagger . We iterate through indices of the output string less than q_R^\ddagger and compute the corresponding bit of the output string using p_R^\ddagger . For each of these computed bits, we cons them on to a variable w , which “accumulates” the eventual output.

If we were to write this in a legible but informal imperative pseudocode, it would look like this:

```
w = nil
for 1 ≤ c ≤ ⌊q_R^\ddagger⌋(x)
    w = cons(⌊p_R^\ddagger⌋(x, c), w)
return w
```

It is easy to see that this program computes $f(x)$ on input x , as it outputs a string of length $\llbracket q_R^\ddagger \rrbracket(x)$ whose i -th bit is $\llbracket p_R^\ddagger \rrbracket(x, i)$. Written slightly more carefully, we get this:

$$\begin{aligned} \text{out}(x) &= f(x, \text{one}, \text{nil}) \\ f(x, c, w) &= \text{if } c = h_q(x) \text{ then } w \text{ else } f(x, \text{add}(c, \text{one}), \text{cons}(h_p(x, c), w)), \end{aligned}$$

where h_p and h_q denote the heads of the programs p_R^\ddagger and q_R^\ddagger , respectively. But even here, we have sacrificed some precision for legibility: for example, we must replace $=$ by two occurrences of `less`, and a term like $f(x, 0, \text{nil})$ must be understood as $f(x \oplus 0, \text{nil})$.

Thus, we have transformed (λ, p, q) into an RW-factorizable program computing the same function. Moreover, note that it is non-nested, and tail-recursive if p and q are.¹⁶ Combined with Theorem 6, we get the following translation result:

Theorem 9. *For every function $f : 2^* \rightarrow 2^*$ whose length is polynomially bounded, there is a non-nested RW-factorizable program of type $R \rightarrow W$ computing f iff there is a polynomially bounded function λ and bit-length programs p and q such that (λ, p, q) properly computes f .*

¹⁶ Recall that “non-nested” only applied to recursive functions of output type W .

Similarly, there is a tail-recursive RW-factorizable program of type $R \rightarrow W$ computing f iff there is a polynomially bounded function λ and tail-recursive bit-length programs p and q such that (λ, p, q) properly computes f .

Theorems 9 and 2 have the following immediate consequence, which is the central result of this paper:

Theorem 10. For any function $f : 2^* \rightarrow 2^*$, $f \in \text{FP}$ iff f is computed by a non-nested RW-factorizable program of type $R \rightarrow W$, and $f \in \text{FL}$ iff f is computed by a tail-recursive RW-factorizable program of type $R \rightarrow W$.

A few comments are in order. The proof of this theorem gives us a rather strong normal form for non-nested RW-factorizable programs. Namely, each such program can be written with a single `for` loop on top of purely `cons-free` subroutines. We know that those `cons-free` subroutines can be made non-nested in general, and can furthermore be made tail-recursive if the original program was tail-recursive.

It remains an open question which complexity class general (i.e., possibly nested) RW-factorizable programs capture, but it seems to be intermediate between the class of polynomial time and polynomial space functions. (Recall that the length of the output of the latter functions can grow exponentially in the length of the input.) A plausible candidate is functions whose length is computable in polynomial space, but whose bits are computable in polynomial time.

8 Composing bit-length programs

Finally, we turn to the problem of *syntactic composition*: producing a program r such that $\llbracket r \rrbracket = \llbracket p \rrbracket \circ \llbracket q \rrbracket$ given p and q . In most programming languages this is trivial: we combine programs p and q and add a new head $h_p(h_q(x))$, where h_p and h_q are the recursive function symbols in the heads of p and q respectively. However, this does not work for RW-factorizable programs of type $R \rightarrow W$, as the term $h_p(h_q(x))$ would be ill-typed.

However, since we can transform RW-factorizable programs to and from equivalent pairs of bit-length programs, it suffices to syntactically compose these.¹⁷ As we have remarked above, pairs of bit-length programs *are* compositional in a rough extensional sense: we are given bit- and length-access to the input, and we compute bit- and length-access to the output.

More precisely, suppose we have a triple (λ_f, p_f, q_f) that properly computes a function f and a triple (λ_g, p_g, q_g) that properly computes g . We want to transform p_f and q_f so that they output the bits and length of $f(g(x))$, as opposed to $f(x)$. Namely, we must “re-interpret” the primitives in p_f and q_f so that it returns the bits of $g(x)$ instead of bits of x , and the constant `max` so that it returns the maximum index of $g(x)$ instead of the maximum index of x . Luckily this is exactly what we have in p_g and q_g .

¹⁷ To be completely syntactic when we transform an RW-factorizable program into an equivalent triple (λ, p, q) , we must provide a finite description for λ as well. This is not hard: if f is computable by an RW-factorizable program, then its length on inputs of length n is bounded by $a(n+1)^b$, where a is something like the number of distinct terms that occur in the program and b is the maximum length of any tuple.

Finally, we have to find a suitable bound for the size of the counting module in the composed program. We assume λ_f and λ_g are increasing and dominate the identity function. Note that if $|f(x)| \leq \lambda_f(|x|)$ and $|g(x)| \leq \lambda_g(|x|)$, then $|f(g(x))| \leq (\lambda_f \circ \lambda_g)(|x|)$. Hence, we may take $\lambda_f \circ \lambda_g$.

We now present the formal development, which follows a familiar rhythm: a map on types, on values, terms, and programs, followed by a proof of that the program transformation behaves how we want it to. In this section, let us fix:

- functions $f, g : 2^* \rightarrow 2^*$,
- increasing functions $\lambda_f, \lambda_g : \omega \rightarrow \omega$, each dominating the identity function, such that $\lambda_f(|x|) > |f(x)|$ and $\lambda_g(|x|) > |g(x)|$ for all x ,
- a function $\mu : \omega \rightarrow \omega$ satisfying $\mu \geq \lambda_f \circ \lambda_g$,
- a triple (λ_f, p_f, q_f) properly computing f , and
- a triple (λ_g, p_g, q_g) properly computing g .

Note that the map defined in this section will be denoted by a g in the subscript, e.g., $T \mapsto T_g$. For types, terms, and programs, this g is purely formal; not so for values, where it actually depends on the partial function g .

Definition 44. Let $2_g = 2$, $I_g = C$, and $C_g = C$. Extend this map to a map $\alpha \mapsto \alpha_g$ on product types coordinate-wise, and thence to a map $(\beta \rightarrow \alpha) \mapsto (\beta_g \rightarrow \alpha_g)$ on function types.

Definition 45. For each product type α , fix an injection $x \mapsto x_g : \text{Var}_\alpha \rightarrow \text{Var}_{\alpha_g}$ and an injection $f \mapsto f_g : \text{RFSymb}_\rho \rightarrow \text{RFSymb}_{\rho_g}$.

For the next definition, notice that $|g(x)| \leq \lambda_g(|x|) \leq \lambda_f(\lambda_g(|x|)) \leq \mu(|x|)$ and $\lambda_f(|g(x)|) \leq \lambda_f(\lambda_g(|x|)) \leq \mu(|x|)$, by the assumption that λ_f is increasing and dominates the identity.

Definition 46. For each $x \in 2^*$, we define a map $v \mapsto v_g : I(g(x)) \rightarrow C(\mu(|x|))$ and $C(\lambda_f(|g(x)|)) \rightarrow C(\mu(|x|))$ by inclusion as initial segments of ω . Extend this to a map $v \mapsto v_g$ on product types coordinate-wise, fixing boolean values.

For the next definition and henceforth, let $h_p(c)$ and h_q be the heads of p_g and q_g . Notice that these are terms of type 2 and C respectively, and that c is a variable of type C .

Definition 47. Given a term T of type α , we define the term T_g of type α_g as follows:

1. If $T \equiv x$ (a variable) then $T_g \equiv x_g$.
2. If $T \equiv f(T')$, then $T_g \equiv f_g(T'_g)$.
3. If $T \equiv \text{true}, \text{false}, 0$, or 1 , then $T_g \equiv T$.
4. If $T \equiv T_0 + T_1, T_0 - T_1$, or $T_0 \leq T_1$, then $T_g \equiv (T_0)_g + (T_1)_g, (T_0)_g - (T_1)_g$, or $(T_0)_g \leq (T_1)_g$ respectively.
5. If $T \equiv \text{min}$, then $T_g \equiv 0$.
6. If $T \equiv \text{max}$, then $T_g \equiv h_q$.
7. If $T \equiv P(T')$ or $\text{null}(T')$, then $T_g \equiv T'_g - 1$, or $T'_g \leq 0$, respectively.
8. If $T \equiv \text{bit}(T')$, then $T_g \equiv h_p(T'_g)$.

- 9. If $T \equiv (T_0, \dots, T_{n-1})$, then $T_g \equiv ((T_0)_g, \dots, (T_{n-1})_g)$.
- 10. If $T \equiv T'[i]$, then $T_g \equiv (T'_g)[i]$.
- 11. If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, then $T_g \equiv \text{if } (T_0)_g \text{ then } (T_1)_g \text{ else } (T_2)_g$.

Definition 48. Given a program $p = (\mathbf{f}_i(\mathbf{x}_i) = T_i)_{0 \leq i \leq k}$, let program $p_{(g)}$ be defined by combining the programs p_g and q_g , then adding new lines $((\mathbf{f}_i)_g(\mathbf{x}_i)_g) = (T_i)_g)_{0 \leq i \leq k}$ on top. (The head of $p_{(g)}$ is the line $(\mathbf{f}_0)_g(\mathbf{x}_0)_g = (T_0)_g$, where $\mathbf{f}_0(\mathbf{x}_0) = T_0$ is the head of p .)

Lemma 8. For any program p , environment ρ , string x , term T , and value v ,

$$g(x), \lambda_f(|g(x)|), \rho \vdash_p T \rightarrow v \implies x, \mu(|x|), \rho_g \vdash_{p_{(g)}} T_g \rightarrow v_g,$$

without collisions, if the derivation of the left-hand side has no collisions.

(The proof is in Appendix A.)

Theorem 11. $(\mu, (p_f)_{(g)}, (q_f)_{(g)})$ properly computes $f \circ g$.

Proof Let $h'_p(c)$ and h'_q be the left-hand-sides of the heads of p_f and q_f respectively. Fix a string x . By definition of q_f ,

$$g(x), \lambda_f(|g(x)|) \vdash_{q_f} h'_q \rightarrow |f(g(x))|$$

without collisions. Hence by Lemma 8,

$$x, \mu(|x|) \vdash_{(q_f)_{(g)}} (h'_q)_g \rightarrow |f(g(x))|,$$

i.e., $\llbracket (q_f)_{(g)} \rrbracket(x) = |f(g(x))|$, without collisions.

By definition of p_f , for each $i < |f(g(x))|$,

$$g(x), \lambda_f(|g(x)|), [c = i] \vdash_{p_f} h'_p(c) \rightarrow (f(g(x)))_i$$

without collisions. Hence, by Lemma 8,

$$x, \mu(|x|), [c_g = i] \vdash_{(p_f)_g} (p'_f)_g(c_g) \rightarrow (f(g(x)))_i,$$

i.e., $\llbracket (p_f)_{(g)} \rrbracket(x, i) = (f(g(x)))_i$, without collisions. ■

Example. Let q_g be the program

$$\begin{aligned} h_q &= \mathbf{q}(\max) \\ \mathbf{q}(x) &= \text{if } \text{null}(x) \text{ then } 0 \text{ else } 1 + \mathbf{q}(\mathbf{P}(x)), \end{aligned}$$

and let p_g be the program

$$h_p(x) = \text{if } \text{bit}(x) \text{ then } \text{false}, \text{ else } \text{true}.$$

In other words, for sufficiently large λ , (λ, p_g, q_g) computes the function which flips the bits of the input string. Let us *square* this function, i.e., create an identical pair of programs (p_f, q_f) , distinguishing the function symbols of the latter by replacing, e.g., \mathbf{q} by \mathbf{q}' , and

then composing (p_f, q_f) with (p_g, q_g) . We get for the length component:

$$\begin{aligned} h'_q &= q'(h_q) \\ q'(x_g) &= \text{if } x_g \leq 0 \text{ then } 0 \text{ else } 1 + q'(x_g - 1) \\ h_q &= q(\text{max}) \\ q(x) &= \text{if null}(x) \text{ then } 0 \text{ else } 1 + q(P(x)), \end{aligned}$$

and for the bits component,

$$\begin{aligned} h'_p(x_g) &= \text{if } h_p(x_g) \text{ then false, else true} \\ h_p(x) &= \text{if bit}(x) \text{ then false, else true.} \end{aligned}$$

By inspection, we can see that this pair of programs computes the identity function, as it should.

9 Discussion and open questions

We have identified RW-factorizable programs as a simple extension of cons-free programs that captures functional polynomial time and logarithmic space, and we have also shown how to compose two such programs. We introduced the notion of bit-length computability as auxiliary machinery to help us show these results. Our work suggests a number of subsequent questions:

1. Higher orders, non-determinism, and other data types.

Cons-free programs have been studied in all of these contexts. Jones (2001) showed that deterministic cons-free programs at higher orders capture the deterministic exponential hierarchy, whose union is the class of elementary relations. Kop & Simonsen (2017) showed that whereas non-deterministic first-order cons-free programs are no more expressive than deterministic ones, even non-deterministic second-order programs capture the entire class of elementary relations. Ben-Amram & Petersen (1998) studied cons-free programs over tree data. Does the RW-factorizable paradigm extend smoothly to all of these situations?

2. RW-factorizable algorithms

Many algorithms seem to be RW-factorizable, in the sense that they have read-only input and write-only output (allowing for some flexibility in what we mean by “read” and “write”). For example, both selection sort and insertion sort seem to follow this paradigm: both algorithms have an input list which gets whittled down over the course of computation while the output list is simultaneously built up. Mergesort, on the other hand, is definitely *not* RW-factorizable: whereas *merge* can be consistently typed by $R \rightarrow W$, the recursive definition $\text{sort}(u) = \text{merge}(\text{sort}(u_0), \text{sort}(u_1))$ forces $\text{sort}(u)$ to be of type R and W simultaneously.

Are there general techniques or results that apply to RW-factorizable algorithms in general, for example, stronger lower bounds? More generally, can we consider other ways to separate construction from destruction that just variable typing? (e.g., in heapsort we separate them *chronologically*: the same piece of data is first constructed, then destructed.)

3. RW-factorizable programs as an indexing of polynomial-time functions.

If we fix an encoding of RW-factorizable programs by binary strings, we get an *indexing* of polynomial-time functions, i.e., an identification of strings with functions. Indexings of partial recursive functions are the central topic in recursion theory, but indexings of complexity classes (or *subrecursive indexings*) have received a little bit of attention in the context of proving, e.g., speedup phenomena for non-Turing complete languages (Constable & Borodin, 1972; Alton, 1980).¹⁸ Kozen (1978) has identified several simple axioms of subrecursive indexings with significant consequences, for example, Kleene's second recursion theorem. One of these axioms is the existence of a polynomial-time function which gives a code for $f \circ g$ given codes for f and g . It seems likely that the transformations in this paper can be computed in polynomial time, and that the natural indexing of FP given by RW-factorizable programs satisfies this and the other axioms.

This suggests that we might look for speedup phenomena within the class of RW-factorizable programs itself. Given that cons-free running time corresponds to something like (the exponential of) circuit depth, such speedups might have complexity-theoretic consequences (Bhaskar et al., 2022).

Acknowledgements

We are immensely grateful to Andrzej Filinski, Neil Jones, and Jakob Nordström for their comments on this work. This paper is dedicated to the memory of Neil Jones.

Conflicts of interest

None.

References

- Alton, D. A. (1980) "Natural" programming languages and complexity measures for subrecursive programming languages: An abstract approach. In *Recursion Theory, its Generalisations and Applications*. London Mathematical Society Lecture Note Series. Cambridge University Press, pp. 248–285. doi:10.1007/CBO9780511629181.011.
- Aubert, C., Seiller, T., Rubiano, T. & Rusch, N. (2022) mwp-Analysis Improvement and Implementation: Realizing Implicit Computational Complexity. In 7th International Conference on Formal Structures for Computation and Deduction, FSCD 2022, August 2–5, 2022, Haifa, Israel, Felty, A. P. (ed). LIPIcs, vol. 228. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 26:1–26:23.
- Avanzini, M., Eguchi, N. & Moser, G. (2011) A path order for rewrite systems that compute exponential time functions. In Proceedings of the 22nd International Conference on Rewriting Techniques and Applications, RTA 2011, May 30–June 1, 2011, Novi Sad, Serbia, Schmidt-Schauß, M. (ed). LIPIcs, vol. 10. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 123–138.
- Avanzini, M., Eguchi, N. & Moser, G. (2015) A new order-theoretic characterisation of the polytime computable functions. *Theoret. Comput. Sci.* **585**, 3–24.

¹⁸ These are qualitatively different objects: non-pathological indexings of partial recursive functions always contain a *universal program*, whereas indexings of complexity classes never contain one.

- Avanzini, M. & Moser, G. (2016) A combination framework for complexity. *Inf. Comput.* **248**, 22–55.
- Baillot, P., Barthe, G. & Lago, U. D. (2019) Implicit computational complexity of subrecursive definitions and applications to cryptographic proofs. *J. Autom. Reason.* **63**(4), 813–855.
- Baillot, P. & Ghyselen, A. (2022) Types for complexity of parallel computation in Pi-calculus. *ACM Trans. Program. Lang. Syst.* **44**(3), 15:1–15:50.
- Bellantoni, S. & Cook, S. (1992) A new recursion-theoretic characterization of the polytime functions (extended abstract). In Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing (STOC '92). Association for Computing Machinery, pp. 283–293. isbn: 0897915119.
- Ben-Amram, A. M. & Petersen, H. (1998) CONS-free programs with tree input (extended abstract). In Automata, Languages and Programming, 25th International Colloquium, ICALP'98, Aalborg, Denmark, July 13–17, 1998, Proceedings, Larsen, K. G., Skyum, S. & Winskel, G. (eds). Lecture Notes in Computer Science, vol. 1443. Springer, pp. 271–282.
- Bhaskar, S., Kop, C. & Simonsen, J. G. (2022) Subclasses of Ptime Interpreted by Programming Languages. *Theory Comput. Syst.*
- Bonfante, G. (2006) Some programming languages for logspace and Ptime. In Algebraic Methodology and Software Technology, 11th International Conference, AMAST 2006, Kuressaare, Estonia, July 5–8, 2006, Proceedings, pp. 66–80.
- Bonfante, G., Marion, J.-Y. & Moyon, J.-Y. (2011) Quasi-interpretations a way to control resources. *Theoret. Comput. Sci.* **412**(25), 2776–2796.
- Cobham, A. (1965) The intrinsic computational difficulty of functions. In Logic, Methodology and Philosophy of Science: Proceedings of the 1964 International Congress (Studies in Logic and the Foundations of Mathematics), Bar-Hillel, Y. (ed). North-Holland Publishing, pp. 24–30.
- Constable, R. L. & Borodin, A. B. (1972) Subrecursive programming languages, part I. *J. ACM* **19**(3), 526–568.
- Czajka, L. (2018) Term rewriting characterisation of LOGSPACE for finite and infinite data. In 3rd International Conference on Formal Structures for Computation and Deduction, FSCD 2018, July 9–12, 2018, Oxford, UK, pp. 13:1–13:19.
- Dal Lago, U. (2022) Implicit computation complexity in higher-order programming languages: A Survey in Memory of Martin Hofmann. *Math. Struct. Comput. Sci.*, 1–17. doi:[10.1017/S0960129521000505](https://doi.org/10.1017/S0960129521000505).
- de Carvalho, D. & Simonsen, J. G. (2014) An implicit characterization of the polynomial-time decidable sets by cons-free rewriting. In Rewriting and Typed Lambda Calculi - Joint International Conference, RTA-TLCA 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 14–17, 2014. Proceedings, pp. 179–193.
- Girard, J.-Y. (1998) Light linear logic. *Inf. Comput.* **143**(2), 175–204.
- Greibach, S. A. (1975) *Theory of Program Structures: Schemes, Semantics, Verification*. Lecture Notes in Computer Science, vol. 36. Springer.
- Hainry, E., Kapron, B. M., Marion, J.-Y. & Péchoux, R. (2022) A tier-based typed programming language characterizing Feasible Functionals. *Log. Methods Comput. Sci.* **18**(1), 1–31.
- Hofmann, M. (2000) Programming languages capturing complexity classes. *SIGACT News* **31**(1), 31–42.
- Hofmann, M. (2002) The strength of non-size increasing computation. In Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL'02. Portland, Oregon: Association for Computing Machinery, pp. 260–269. isbn: 1581134509.
- Jones, N. D. (1999) LOGSPACE and PTIME characterized by programming languages. *Theoret. Comput. Sci.* **228**(1), 151–174.
- Jones, N. D. (2001) The expressive power of higher-order types or, life without CONS. *J. Funct. Program.* **11**(1), 5–94.
- Jones, N. D., Kop, C., Bhaskar, S. & Simonsen, J. G. (2020) Cons-free programs and complexity classes between LOGSPACE and PTIME. In Proceedings 8th International Workshop on Verification and Program Transformation and 7th Workshop on Horn Clauses for Verification and

- Synthesis, VPT/HCVS@ETAPS 2020, and 7th Workshop on Horn Clauses for Verification and Synthesis Dublin, Ireland, 25–26th April 2020, Fribourg, L. & Heizmann, M. (eds), vol. 320. EPTCS, pp. 65–79.
- Kop, C. & Simonsen, J. G. (2017) The power of non-determinism in higher-order implicit complexity - characterising complexity classes using non-deterministic cons-free programming. In Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22–29, 2017, Proceedings, Yang, H. (ed). Lecture Notes in Computer Science, vol. 10201. Springer, pp. 668–695.
- Kozen, D. (1978) Indexing of subrecursive classes. In Proceedings of the Tenth Annual ACM Symposium on Theory of Computing. STOC'78. San Diego, California, USA: Association for Computing Machinery, pp. 287–295.
- Kristiansen, L. (2022) Reversible computing and implicit computational complexity. *Sci. Comput. Program.* **213**, 102723.
- Lafont, Y. (2004) Soft linear logic and polynomial time. *Theor. Comput. Sci.* **318**(1–2), 163–180.
- Lago, U. D., Kahle, R. & Oitavem, I. (2021) A recursion-theoretic characterization of the probabilistic class PP. In 46th International Symposium on Mathematical Foundations of Computer Science, MFCS 2021, August 23–27, 2021, Tallinn, Estonia, Bonchi, F. & Puglisi, S. J. (eds). LIPIcs, vol. 202. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 35:1–35:12.
- Lago, U. D. & Toldin, P. P. (2015) A higher-order characterization of probabilistic polynomial time. *Inf. Comput.* **241**, 114–141.
- Leivant, D. (1995) Ramified recurrence and computational complexity I: Word recurrence and polytime. In *Feasible Mathematics II. Birkhäuser Boston*, pp. 320–343.
- Marion, J. (2011) A type system for complexity flow analysis. In 2011 IEEE 26th Annual Symposium on Logic in Computer Science (LICS 2011), pp. 123–132.
- Moschovakis, Y. N. (2018) *Abstract Recursion and Intrinsic Complexity*. Lecture Notes in Logic. Cambridge University Press.

Appendix A: Proofs of correctness

In this section, we collect all the proofs of correctness for transformations defined earlier.

Lemma (Lemma 2). *Suppose p is a cons-free program, x is a string, ρ is an x -environment. Then,*

$$\rho \vdash_p T \rightarrow v \implies x, \rho^\dagger \vdash_{p^\dagger} T^\dagger \rightarrow v^\dagger.$$

Proof The proof is by induction on the length of the derivation of $\rho \vdash T \rightarrow v$. Abbreviate $x, \rho^\dagger \vdash_{p^\dagger}$ by $\rho^\dagger \vdash$.

- If $T \equiv \text{true}$ or $T \equiv \text{false}$, then $v = \top$ or $v = \perp$, and $v^\dagger = v$ and $T^\dagger \equiv T$.
- If $T \equiv x$, then $T^\dagger \equiv x^\dagger$, and hence, $\rho^\dagger \vdash T^\dagger \rightarrow \rho(x^\dagger)$, which is v^\dagger .
- If $T \equiv \text{hd}(S)$, then let u satisfy $\rho \vdash S \rightarrow u$. Then by the suffix property, u is a suffix of x , so the head of u is the bit $x_{|u|}$ and $\rho \vdash T \rightarrow x_{|u|}$. By induction, $\rho^\dagger \vdash S \rightarrow u^\dagger$ so $\rho^\dagger \vdash \text{bit}(S^\dagger) \rightarrow x_{u^\dagger}$, but $u^\dagger = |u|$.
- If $T \equiv \text{tl}(S)$, then there exists a u and $b \in 2$ such that $v = bu$ and $\rho \vdash S \rightarrow u$. By induction, $\rho^\dagger \vdash S^\dagger \rightarrow |u|$. Hence, $\rho^\dagger \vdash P(S^\dagger) \rightarrow |u| - 1$.
- If $T \equiv \text{null}(S)$, then v is true or false depending on whether S is the empty string. Similarly, $\text{null}(S^\dagger)$ is true or false depending on whether S^\dagger is zero, the length of the empty string.

- If $T \equiv T_0 \oplus \dots \oplus T_{n-1}$, then there exist v_0, \dots, v_{n-1} such that $\rho \vdash T_i \rightarrow v_i$ for each $i < n$ and $v = v_0 \circ \dots \circ v_{n-1}$. By induction, $\rho^\dagger \vdash T_i^\dagger \rightarrow v_i^\dagger$ for each $i < n$. Since $v^\dagger = v_0^\dagger \circ \dots \circ v_{n-1}^\dagger$, $\rho^\dagger \vdash T_0^\dagger \oplus \dots \oplus T_{n-1}^\dagger \rightarrow v^\dagger$; hence, $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- If $T \equiv S[i]$, then there exists a value u such that $\rho \vdash S \rightarrow u$, (u_0, \dots, u_{n-1}) is a decomposition of u into atomic values, and $v = u_i$. By induction, $\rho^\dagger \vdash S^\dagger \rightarrow u^\dagger$, where $u^\dagger = (u_0^\dagger, \dots, u_{n-1}^\dagger)$. Thus, $\rho^\dagger \vdash S^\dagger[i] \rightarrow u_i^\dagger$, which says $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, assume $\rho \vdash T_0 \rightarrow \top$. (The case $\rho \vdash T_0 \rightarrow \perp$ is similar.) Then, $\rho \vdash T_1 \rightarrow v$. By induction, $\rho^\dagger \vdash T_0^\dagger \rightarrow \top$ and $\rho^\dagger \vdash T_1^\dagger \rightarrow v^\dagger$. Hence, $\rho^\dagger \vdash T \rightarrow v^\dagger$.
- If $T \equiv f(S)$, then let $f(x^f) = T^f$ be the definition of f in p . Then there exists some u such that $\rho \vdash S \rightarrow u$ and $[x^f = u] \vdash T^f \rightarrow v$. By induction, $\rho^\dagger \vdash S^\dagger \rightarrow u^\dagger$ and $[x^f = u^\dagger] \vdash (T^f)^\dagger \rightarrow v^\dagger$. Hence, $\rho^\dagger \vdash f^\dagger(S^\dagger) \rightarrow v^\dagger$, and $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$. ■

Lemma (Lemma 3). *Suppose that p is an RW-factorizable program, T is an RW-term, $x, v, v', w \in 2^*$, ρ is an x -environment, and w is not bound by ρ .*

Suppose that $\rho, [w = w] \vdash_p T \rightarrow v$ and $\rho, [w = \varepsilon] \vdash_p T \rightarrow v'$.¹⁹ Then either

- $x, \rho^\dagger \vdash_{\rho^\diamond} T^\diamond \rightarrow \top$ and $v = v'$, or
- $x, \rho^\dagger \vdash_{\rho^\diamond} T^\diamond \rightarrow \perp$ and $v = v'w$.

Proof Abbreviate $\rho, [w = w] \vdash_p T \rightarrow v$ by $[\rho, w] \vdash T \rightarrow v$, and similarly abbreviate $\rho, [w = \varepsilon] \vdash_p T \rightarrow v$ by $[\rho, \varepsilon] \vdash T \rightarrow v$. Abbreviate $x, \rho^\dagger \vdash_{\rho^\diamond} T^\diamond \rightarrow b$ by $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow b$. The proof proceeds by induction on the size of the derivation of $[\rho, w] \vdash T \rightarrow v$ and breaks up into cases depending on the form of T .

- If $T \equiv w$, then $v = w$, $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \perp$, and $v' = \varepsilon$. In this case, $v = w = \varepsilon w = v'w$, so we're done.
- If $T \equiv \text{nil}$, then $v = \varepsilon$, $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$, and $v' = \varepsilon$. In this case $v = \varepsilon = v'$, so we're done.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, first suppose that $\rho \vdash T_0 \rightarrow \top$. By Lemma 2, $\rho^\dagger \vdash_\diamond T_0^\dagger \rightarrow \top$. By definition of T , $[\rho, w] \vdash T_1 \rightarrow v$ and $[\rho, \varepsilon] \vdash T_1 \rightarrow v'$. By induction, either $\rho^\dagger \vdash_\diamond T_1^\diamond \rightarrow \top$ or $\rho^\dagger \vdash_\diamond T_1^\diamond \rightarrow \perp$.
Suppose that $\rho^\dagger \vdash_\diamond T_1^\diamond \rightarrow \top$. By induction $v = v'$. By definition of T^\diamond , $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$. Similarly, suppose that $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \perp$. By induction, $v = v'w$. By definition of T^\diamond , $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \perp$, and we're done.
(The case $\rho \vdash T_0 \rightarrow \perp$ is similar.)
- Suppose that $T \equiv f(T')$ for some $f : \beta \rightarrow W$. Then, the variable w does not occur in T , so $v' = v$. Since $T^\diamond \equiv \text{true}$, $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$, and we are done.
- Suppose that $T \equiv f(T', S)$ for some $f : \beta \times W \rightarrow W$. Suppose that $f(x^f, w) = T^f$ is the recursive definition of f in p . Then there exist t and s such that $\rho \vdash T' \rightarrow t$, $[\rho, w] \vdash S \rightarrow s$, and $[x^f = t, w = s] \vdash T^f \rightarrow v$, and there exists s' such that $[\rho, \varepsilon] \vdash S \rightarrow s'$ and $[x^f = t, w = s'] \vdash T^f \rightarrow v'$. By Lemma 2, $\rho^\dagger \vdash (T')^\dagger \rightarrow t^\dagger$. Since convergence is independent of the binding of w , there exists a v'' such that $[x^f = t, w = \varepsilon] \vdash T^f \rightarrow v''$.

¹⁹ Note that the convergence of the term T only depends on the environment ρ and not on the binding of the variable w .

Now

$$[x^f = t, w = s] \vdash T^f \rightarrow v \text{ and } [x^f = t, w = \varepsilon] \vdash T^f \rightarrow v''.$$

By induction,

$$([(x^f)^\dagger = t^\dagger] \vdash_\diamond (T^f)^\diamond \rightarrow \top \wedge v = v'') \vee ([(x^f)^\dagger = t^\dagger] \vdash_\diamond (T^f)^\diamond \rightarrow \perp \wedge v = v'').$$

Similarly,

$$[x^f = t, w = s'] \vdash T^f \rightarrow v' \text{ and } [x^f = t, w = \varepsilon] \vdash T^f \rightarrow v''.$$

By induction,

$$([(x^f)^\dagger = t^\dagger] \vdash_\diamond (T^f)^\diamond \rightarrow \top \wedge v' = v'') \vee ([(x^f)^\dagger = t^\dagger] \vdash_\diamond (T^f)^\diamond \rightarrow \perp \wedge v' = v''s').$$

First assume $[(x^f)^\dagger = t^\dagger] \vdash_\diamond (T^f)^\diamond \rightarrow \top$. Then, $v = v'' = v'$ and (by definition of T^\diamond) $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$, which completes this case.

Next assume then that $[(x^f)^\dagger = t^\dagger] \vdash_\diamond (T^f)^\diamond \rightarrow \perp$. In this case, $v = v''s$ and $v' = v''s'$.

Now

$$[\rho, w] \vdash S \rightarrow s \text{ and } [\rho, \varepsilon] \vdash S \rightarrow s'.$$

Then by induction,

$$(\rho^\dagger \vdash_\diamond S^\diamond \rightarrow \top \wedge s = s') \vee (\rho^\dagger \vdash_\diamond S^\diamond \rightarrow \perp \wedge s = s'w).$$

If $\rho^\dagger \vdash_\diamond S^\diamond \rightarrow \top$, then $v = v''s = v''s' = v'$ and $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$ (by definition of T^\diamond).

If, on the other hand, $\rho^\dagger \vdash_\diamond S^\diamond \rightarrow \perp$, then $v = v''s = v''s'w = v'w$ and $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \perp$ (by definition of T^\diamond). This finishes the case that $T \equiv \mathbf{f}(T', S)$.

- If $T \equiv \text{cons}(T', S)$, then there are strings s and s' and a character $c \in 2$ such that $[\rho, w] \vdash S \rightarrow s$, $[\rho, \varepsilon] \vdash S \rightarrow s'$, $v = cs$, and $v' = cs'$. By induction, either $\rho^\dagger \vdash_\diamond S^\diamond \rightarrow \top$ and $s = s'$ or $\rho^\dagger \vdash_\diamond S^\diamond \rightarrow \perp$ and $s = s'w$. In the first case, $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$ and $v = cs = cs' = v'$. In the second case, $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \perp$ and $v = cs = cs'w = v'w$. ■

Lemma (Lemma 4). *For every RW program p of type $R \rightarrow W$, RW term T from p of type W , strings x, v , and w , natural number $n > |v|$, and x -environment ρ binding w to w :*

$$\rho \vdash_p T \rightarrow v \implies x, n, \rho^\dagger \vdash_{p^\ell} T^\ell \rightarrow v^\ell - \delta,$$

where $\delta = 0$ if $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$ or $\delta = |w|$ if $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \perp$. Moreover, the derivation of the right-hand side is collision-free.

Proof Let $\rho \vdash T \rightarrow v$ abbreviate $\rho \vdash_p T \rightarrow v$ and $\rho \vdash_\ell T \rightarrow v$ abbreviate $x, n, \rho \vdash_{p^\ell} T \rightarrow v$. The proof proceeds by induction on the derivation of $\rho \vdash T \rightarrow v$ and breaks up into cases depending on the form of T . To verify that the derivation is collision-free, we note that any time we add counting modules we do not introduce a new collision. Note that since p^ℓ contains p^\diamond , \vdash_ℓ extends \vdash_\diamond .

- If $T \equiv w$, then $v = w$ and $v^\ell = |w|$. Moreover, $\rho^\dagger \vdash_\ell T^\diamond \rightarrow \perp$, so $\delta = |w|$ and $v^\ell - \delta = 0$. But in this case, $T^\ell \equiv 0$.

- If $T \equiv \text{nil}$, then $v = \varepsilon$. Moreover, $\rho^\dagger \vdash_\ell T^\diamond \rightarrow \top$, so $\delta = \varepsilon$ and $v^\ell - \delta = 0$. But in this case, $T^\ell \equiv 0$ as well.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, assume $\rho \vdash T_0 \rightarrow \top$ (the case \perp is similar). By Lemma 2, $\rho^\dagger \vdash_\ell T_0^\dagger \rightarrow 1$, so by definition of T^\diamond , $\rho^\dagger \vdash_\ell T^\diamond \rightarrow \top \iff \rho^\dagger \vdash_\ell T_1^\dagger \rightarrow \top$. Therefore, the δ associated with T_1 is identical to the δ associated with T , and we can simply say δ unambiguously.
 Since $\rho \vdash T_0 \rightarrow 1$, $\rho \vdash T_1 \rightarrow v$, so (by induction) $\rho^\dagger \vdash_\ell T_1^\ell \rightarrow v^\ell - \delta$. Since $\rho^\dagger \vdash_\ell T_0^\dagger \rightarrow \top$, $\rho^\dagger \vdash_\ell T^\ell \rightarrow v^\ell - \delta$ by definition of T^ℓ .
- If $T \equiv \mathbf{f}(T')$, let $\mathbf{f}(x^\mathbf{f}, \mathbf{w}) = T^\mathbf{f}$ be the definition of \mathbf{f} in p . There exists a value t such that $\rho \vdash T' \rightarrow t$ and $[x^\mathbf{f} = t] \vdash T^\mathbf{f} \rightarrow v$. By induction, $[(x^\mathbf{f})^\dagger = t^\dagger] \vdash_\ell (T^\mathbf{f})^\ell \rightarrow v^\ell$. Hence, $\rho^\dagger \vdash_\ell \mathbf{f}^\ell((T')^\dagger) \rightarrow v^\ell$, and $\rho^\dagger \vdash_\ell T^\ell \rightarrow v^\ell$. (In this case, $\rho^\dagger \vdash_\ell T^\diamond \rightarrow \top$.)
- If $T \equiv \mathbf{f}(T', S)$, let $\mathbf{f}(x^\mathbf{f}, \mathbf{w}) = T^\mathbf{f}$ be the definition of \mathbf{f} in p . There exist values t and s such that $\rho \vdash T' \rightarrow t$, $\rho \vdash S \rightarrow s$, and $[x^\mathbf{f} = t, \mathbf{w} = s] \vdash T^\mathbf{f} \rightarrow v$.
 Let us first suppose that $\rho^\dagger \vdash_\diamond \mathbf{f}^\diamond((T')^\dagger) \rightarrow \top$. By Lemma 2, $\rho^\dagger \vdash_\ell (T')^\dagger \rightarrow t^\dagger$. Since the recursive definition of \mathbf{f}^\diamond is $\mathbf{f}^\diamond((x^\mathbf{f})^\dagger) = (T^\mathbf{f})^\diamond$, we have that $[(x^\mathbf{f})^\dagger = t^\dagger] \vdash_\ell (T^\mathbf{f})^\diamond \rightarrow \top$. By induction, $[(x^\mathbf{f})^\dagger = t^\dagger] \vdash_\ell (T^\mathbf{f})^\ell \rightarrow v^\ell$; hence, $\rho^\dagger \vdash_\ell \mathbf{f}^\ell((T')^\dagger) \rightarrow v^\ell$. By definition of T^ℓ , $\rho^\dagger \vdash_\ell T^\ell \rightarrow v^\ell$, which is what we wanted to show.
 Now suppose that $\rho^\dagger \vdash_\diamond \mathbf{f}^\diamond((T')^\dagger) \rightarrow \perp$. By definition of T^\diamond , $\rho^\dagger \vdash_\ell T^\diamond \rightarrow b$ iff $\rho^\dagger \vdash_\ell S^\diamond \rightarrow b$ for both boolean values b . Therefore, since \mathbf{w} is bound to the same value (w) in S as in T , the value of δ can be used unambiguously in both the inductive hypothesis at S and the conclusion at T .
 Moreover, $[(x^\mathbf{f})^\dagger = t^\dagger] \vdash_\ell (T^\mathbf{f})^\diamond \rightarrow \perp$. Thus, by the inductive hypothesis applied to $[x^\mathbf{f} = t, \mathbf{w} = s] \vdash T^\mathbf{f} \rightarrow v$, $[(x^\mathbf{f})^\dagger = t^\dagger] \vdash_\ell (T^\mathbf{f})^\ell \rightarrow v^\ell - |s|$, and therefore, $\rho^\dagger \vdash_\ell \mathbf{f}^\ell((T')^\dagger) \rightarrow v^\ell - |s|$. By the inductive hypothesis at $\rho \vdash S \rightarrow s$, $\rho^\dagger \vdash_\ell S^\ell \rightarrow s^\ell - \delta$. As $s^\ell = |s|$, $\rho^\dagger \vdash_\ell \mathbf{f}^\ell(T') + S^\ell \rightarrow v^\ell - \delta$. (Note that this uses the assumption that $n > |v|$, which also verifies that this addition does not introduce a new collision.)
 Finally, since $\rho^\dagger \vdash_\diamond \mathbf{f}^\diamond((T')^\dagger) \rightarrow \perp$ and by definition of T^ℓ , $\rho^\dagger \vdash_\ell T^\ell \rightarrow v^\ell - \delta$, which is what we wanted to show.
- If $T \equiv \text{cons}(T', S)$, then there exists a character $c \in 2$ and a string s such that $\rho \vdash T' \rightarrow b$, $\rho \vdash S \rightarrow s$, and $v = cs$. As in the previous case, the values of δ at S and T are identical. By induction, $\rho^\dagger \vdash_\ell S^\ell \rightarrow |s| - \delta$, so $\rho^\dagger \vdash_\ell T^\ell \rightarrow |s| + 1 - \delta$, which is $|v| - \delta$. (We again use the assumption that $n > |v|$ here, which again verifies that this addition does not introduce a new collision.) ■

Lemma (Lemma 5). For every RW program p of type $R \rightarrow W$, RW term T from p of type W , strings x, w and v , natural numbers $n > |v|$ and $1 \leq c \leq |v| - \delta$, and x -environment ρ binding w to w ,

$$\rho \vdash_p T \rightarrow v \implies x, n, \rho^\dagger, [c = c] \vdash_{p^b} T^b \rightarrow v_{c+\delta},$$

where $\delta = 0$ if $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \top$ or $\delta = |w|$ if $\rho^\dagger \vdash_\diamond T^\diamond \rightarrow \perp$.²⁰ Moreover, the computation on the right-hand side is without collisions.

²⁰ Recall that $v_{c+\delta}$ refers to a particular bit of v , where $v = v_{|v|} \dots v_2 v_1$.

Proof Let $\rho \vdash T \rightarrow v$ abbreviate $\rho \vdash_p T \rightarrow v$ and $[\rho^\dagger, c] \vdash_b T^b \rightarrow a$ abbreviate $x, n, \rho^\dagger, [c = c] \vdash_{p^b} T^b \rightarrow a$. The proof is by induction on the size of the derivation of $\rho \vdash T \rightarrow v$ and breaks up into cases depending on the form of T .

- If $T \equiv w$ or $T \equiv \text{nil}$, then $|v| - \delta = 0$, so the universal quantification over c is empty and the conclusion is vacuously true.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, then suppose $\rho \vdash T_0 \rightarrow \top$ (the case $\rho \vdash T_0 \rightarrow \perp$ is similar). Then $\rho \vdash T_1 \rightarrow v$. By Lemma 2, $\rho^\dagger \vdash_\ell T_0^\dagger \rightarrow \top$, and hence $\rho^\dagger \vdash_\ell T^\circ \rightarrow b$ if and only if $\rho^\dagger \vdash_\ell T_1^\circ \rightarrow b$.

Let us apply the induction hypothesis to $\rho \vdash T_1 \rightarrow v$. The value of δ at the induction hypothesis is identical to the value of δ for $\rho \vdash T \rightarrow v$, so by induction $[\rho^\dagger, c] \vdash_b T_1^b \rightarrow v_{c+\delta}$, and by definition of T^b , $[\rho^\dagger, c] \vdash_b T^b \rightarrow v_{c+\delta}$.

- If $T \equiv f(T')$, let $f(x^f) = T^f$ be the definition of f in p . Then, there exists a value t such that $\rho \vdash T' \rightarrow t$ and $[x^f = t] \vdash T^f \rightarrow v$. In this case, $\rho^\dagger \vdash T^\circ \rightarrow \top$, so $\delta = 0$. By Lemma 2, $\rho^\dagger \vdash_b (T')^\dagger \rightarrow t^\dagger$.

Since $[x^f = t] \vdash T^f \rightarrow v$, $[x^f = t, w = \varepsilon] \vdash T^f \rightarrow v$. By induction on the latter, $[(x^f)^\dagger = t^\dagger, c = c] \vdash_b (T^f)^b \rightarrow v_c$. (Notice that the denotation of $(T^f)^\circ$ is irrelevant if $|w| = 0$.) Hence, $[\rho^\dagger, c] \vdash_b f((T')^\dagger, c) \rightarrow v_c$, i.e., $[\rho^\dagger, c] \vdash_b T^b \rightarrow v_c$.

- If $T \equiv f(T', S)$, let $f(x, w) = T^f$ be the definition of f in p . Then, there are values (t, s) such that $\rho \vdash T' \rightarrow s$, $\rho \vdash S \rightarrow s$, and $[x^f = t, w = s] \vdash T^f \rightarrow v$.

First let us suppose that $\rho^\dagger \vdash f^\circ((T')^\dagger) \rightarrow \top$. In this case, $\rho^\dagger \vdash_b T^\circ \rightarrow \top$, so $\delta = 0$. Moreover, $[(x^f)^\dagger = t^\dagger] \vdash_b (T^f)^\circ \rightarrow \top$, so by induction applied to T^f , $[(x^f)^\dagger = t^\dagger, c = c] \vdash_b (T^f)^b \rightarrow v_c$, and hence, $[\rho^\dagger, c] \vdash_b f^b((T')^\dagger, c) \rightarrow v_c$. But by definition of T^b , $[\rho^\dagger, c] \vdash_b T^b \rightarrow v_c$.

Next let us suppose that $\rho^\dagger \vdash f^\circ((T')^\dagger) \rightarrow \perp$. In this case, $[(x^f)^\dagger] \vdash_b (T^f)^\circ \rightarrow \perp$, and the denotation of T° is equivalent to that of S° . Since it is the same environment ρ at T and S as well, the value of $\delta = |\rho(w)|$ is identical at both S and T . Therefore, by induction applied to $\rho \vdash S \rightarrow s$, $[\rho^\dagger, c] \vdash_b S^b \rightarrow s_{c+\delta}$ for any $1 \leq c \leq |s| - \delta$. By Lemma 3, since $[x^f = t, w = s] \vdash T^f \rightarrow v$, s appears as a suffix of v , and hence for $1 \leq c \leq s - |\delta|$, $s_{c+\delta} = v_{c+\delta}$. Hence for any $1 \leq c \leq |s| - \delta$, $[\rho^\dagger, c] \vdash_b S^b \rightarrow v_{c+\delta}$.

By induction applied to $[x^f = t, w = s] \vdash T^f \rightarrow v$, $[(x^f)^\dagger, c = c] \vdash_b (T^f)^b \rightarrow v_{c+|s|}$, and hence $[\rho^\dagger, c] \vdash_b f^b((T')^\dagger, c) \rightarrow v_{c+|s|}$, for any $1 \leq c \leq |v| - |s|$. Said another way, for any $|s| - \delta + 1 \leq c \leq |v| - \delta$, $[\rho^\dagger, c - |s| + \delta] \vdash_b f^b((T')^\dagger, c) \rightarrow v_{c+\delta}$. By Lemma 4, $\rho^\dagger \vdash_b S^\ell \rightarrow |s| - \delta$. Hence for any $|s| - \delta + 1 \leq c \leq |v| - \delta$, $[\rho^\dagger, c] \vdash_b f^b((T')^\dagger, c - S^\ell) \rightarrow v_{c+\delta}$.

Since $\rho^\dagger \vdash_b S^\ell \rightarrow |s| - \delta$, $[\rho^\dagger, c] \vdash_b c \leq S^\ell \rightarrow \top$ just in case $c \leq |s| - \delta$, otherwise $[\rho^\dagger, c] \vdash_b c \leq S^\ell \rightarrow \perp$. We have now shown that for all $1 \leq c \leq |v| - \delta$,

$$[\rho^\dagger, c] \vdash_b \text{if } c \leq S^\ell \text{ then } S^b \text{ else } f^b((T')^\dagger, c - S^\ell) \rightarrow v_{c+\delta}.$$

Since we are in the case $\rho^\dagger \vdash f^\circ((T')^\dagger) \rightarrow \perp$, we have now shown that $[\rho^\dagger, c] \vdash_b T^b \rightarrow v_{c+\delta}$, which is what we wanted to show.

- If $T \equiv \text{cons}(T', S)$, then T° and S° are identical, so we can use δ unambiguously at both T and S as in the previous case. Let $\rho \vdash S \rightarrow s$. By induction, for any $1 \leq c \leq |s| - \delta$, $[\rho^\dagger, c] \vdash_b S^b \rightarrow v_{c+\delta}$. If $c = |v| - \delta$, then $v_{c+\delta}$ is the first character of v , so $\rho \vdash T' \rightarrow v_{c+\delta}$. By Lemma 2, $\rho^\dagger \vdash T' \rightarrow v_{c+\delta}$ ($v_{c+\delta} = v_{c+\delta}^\dagger$ as it is a boolean value).

Since $|v| - \delta = (|s| - \delta) + 1$ and $[\rho^\dagger, c] \vdash_b c \leq S^\ell \rightarrow \top$ if and only if $c \leq |s| - \delta$, otherwise $[\rho^\dagger, c] \vdash_b c \leq S^\ell \rightarrow \perp$, then for any $1 \leq c \leq |v| - \delta$,

$$[\rho^\dagger, c] \vdash_b \text{ if } c \leq S^\ell \text{ then } S^b \text{ else } (T')^\dagger \rightarrow v_{c+\delta};$$

hence $[\rho^\dagger, c] \vdash_b T^b \rightarrow v_{c+\delta}$, which is what we wanted to show.

That there are no collisions stems from the fact that the transformation $T \mapsto T^b$ introduces no new occurrences of counting module addition that were not already contained in the transformation $T \mapsto T^\ell$. ■

Lemma (Lemma 6). *For any cons-free program with global input p , values x u , and v , variable x , and x -term T ,*

$$x, [x = u] \vdash T \rightarrow v \implies [x_R = u \circ x] \vdash T_R^x \rightarrow v.$$

Proof As usual, the proof proceeds by induction on the size of the derivation of the hypothesis and breaks up into cases according to the form of T . Let us abbreviate T_R^x by T_R .

- If $T \equiv x$, then $v = u$, and $[x_R = u \circ x] \vdash x_R[0, n - 1] \rightarrow u$.
- If $T \equiv \text{in}$, then $v = x$, and $[x_R = u \circ x] \vdash x_R[n] \rightarrow x$.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, suppose that $x, [x = u] \vdash T_0 \rightarrow \top$. (The case $x, [x = u] \vdash T_0 \rightarrow \perp$ is similar.) Then, $x, [x = u] \vdash T_1 \rightarrow v$, so by induction, $[x_R = u \circ x] \vdash (T_0)_R \rightarrow \top$ and $[x_R = u \circ x] \vdash (T_1)_R \rightarrow v$. Therefore, $[x_R = u \circ x] \vdash T_R \rightarrow v$.
- Suppose that $T \equiv \varphi(S)$, where $\varphi \in \{\text{hd}, \text{tl}, \text{null}\}$. Then, there exists some v' such that $x, [x = u] \vdash S \rightarrow v'$ and $\varphi(v') = v$. By induction, $[x_R = u \circ x] \vdash S_R \rightarrow v'$. Therefore, $[x_R = u \circ x] \vdash T_R \rightarrow v$.
- Suppose that $T \equiv T_0 \oplus \dots \oplus T_{n-1}$. Then, there exist v_0, \dots, v_{n-1} such that $x, [x = u] \vdash T_i \rightarrow v_i$ for each $i < n$ and $v = v_0 \circ \dots \circ v_{n-1}$. By induction, for each $i < n$, $[x_R = u \circ x] \vdash T_i \rightarrow v_i$, so $[x_R = u \circ x] \vdash T \rightarrow v$.
- If $T \equiv S[i, j]$, there is some v' such that $x, [x = u] \vdash S \rightarrow v'$. Let (v_0, \dots, v_{n-1}) be a decomposition of v' into a tuple of atomic values; then, $v = (v_i, \dots, v_j)$. By induction, $[x_R = u \circ x] \vdash S_R \rightarrow (v_0, \dots, v_{n-1})$, so $[x_R = u \circ x] \vdash S_R[i, j] \rightarrow (v_i, \dots, v_j)$, i.e., $[x_R = u \circ x] \vdash T_R \rightarrow v$.
- Finally, suppose $T \equiv \mathbf{f}(S)$. Let $\mathbf{f}(x^\mathbf{f}) = T^\mathbf{f}$ be the definition of \mathbf{f} in p . Then, there is some v' such that $x, [x = u] \vdash S \rightarrow v'$ and $x, [x^\mathbf{f} = v'] \vdash T^\mathbf{f} \rightarrow v$. By induction, $[x_R = u \circ x] \vdash S_R \rightarrow v'$ and $[x_R^\mathbf{f} = v' \circ x] \vdash (T^\mathbf{f})_R^\mathbf{f} \rightarrow v$. Since $[x_R = u \circ x] \vdash S_R \rightarrow v'$, $[x_R = u \circ x] \vdash S_R \oplus x_R[n] \rightarrow v' \circ x$.
But $\mathbf{f}_R(x_R^\mathbf{f}) = (T^\mathbf{f})_R^\mathbf{f}$ is the recursive definition of \mathbf{f}_R in p_R . Hence, $[x_R = u \circ x] \vdash \mathbf{f}_R(S_R \oplus x_R[n]) \rightarrow v$, which is what we wanted to prove. ■

Lemma (Lemma 7). *For every bit-length program p , bit-length term T from p , string x , x -environment ρ , and value v ,*

$$x, (|x| + 1)^k, \rho \vdash_p T \rightarrow v \implies x, \rho^\ddagger \vdash_{p^\ddagger} T^\ddagger \rightarrow v^\ddagger.$$

Proof The proof is by induction on the length of x , $(|x| + 1)^k, \rho \vdash_p T \rightarrow v$ and then breaks up into cases depending on the form of T . Abbreviate $x, (|x| + 1)^k, \rho \vdash_p T \rightarrow v$ by $\rho \vdash T \rightarrow v$ and $x, \rho^\dagger \vdash_{p^\dagger} T^\dagger \rightarrow v^\dagger$ by $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.

- If $T \equiv \text{true}$ or false , then $v = v^\dagger = \top$ or \perp , and $T^\dagger \equiv \text{true}$ or false , respectively.
- If $T \equiv x$, then $v = \rho(x)$, so $v^\dagger = \rho^\dagger(x^\dagger)$, $T^\dagger \equiv x^\dagger$, and $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- If $T \equiv 0$ then $v = 0$ (as a counting module), $v^\dagger = (0, \dots, 0, 0)$ (as an element of R^k), and $T^\dagger \equiv \text{zero}$. By correctness of zero , $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- The case $T \equiv 1$ is similar to the previous one, replacing $(0, \dots, 0, 0)$ by $(0, \dots, 0, 1)$ and zero by one .
- If $T \equiv T_0 + T_1$, then $T^\dagger \equiv \text{add}(T_0^\dagger, T_1^\dagger)$. By induction, there exist u_0 and u_1 such that $\rho \vdash T_i \rightarrow u_i$ for $i < 2$ and $v = \max\{u_0 + u_1, (|x| + 1)^k - 1\}$. By induction, $\rho^\dagger \vdash T_i^\dagger \rightarrow u_i^\dagger$ for $i < 2$. By correctness of add , $[u_0 = u_0^\dagger, u = u_1^\dagger] \vdash \text{add}(u_0, u_1) \rightarrow v^\dagger$. Hence, $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- The cases $T \equiv T_0 - T_1$ and $T \equiv T_0 \leq T_1$ are similar to the previous one, by the correctness of minus and less respectively.
- If $T \equiv \text{min}$, then $v = 0$ (as an index), $v^\dagger = \varepsilon$, and $T^\dagger = \text{zero}_I$. By correctness of zero_I , $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- If $T \equiv \text{max}$, then $v = |x|$, $v^\dagger = x$, $T^\dagger \equiv \text{in}$, so $x, \rho^\dagger \vdash \text{in} \rightarrow x$ implies $x, \rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- If $T \equiv P(S)$, then there exists some u such that $\rho \vdash S \rightarrow u$ and $u - 1 = v$. In this case, v^\dagger is the tail of u^\dagger ; by induction, $\rho^\dagger \vdash S^\dagger \rightarrow u^\dagger$, and thus, $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- If $T \equiv \text{null}(S)$, then there exists some u such that $\rho \vdash S \rightarrow u$ and $v = \text{true} \iff u = 0$. In this case $v^\dagger = v$ and $u^\dagger = \varepsilon \iff u = 0$; by induction, $\rho^\dagger \vdash S^\dagger \rightarrow u^\dagger$, and thus $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- If $T \equiv \text{bit}(S)$, then there exists some u such that $\rho \vdash S \rightarrow u$ and $x_u = v$. In this case, $v^\dagger = v$ and u^\dagger is a suffix of x with head $x_u = v = v^\dagger$; by induction, $\rho^\dagger \vdash S^\dagger \rightarrow u^\dagger$, and thus, $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- If $T \equiv T_0 \oplus \dots \oplus T_{n-1}$, then there exist v_0, \dots, v_{n-1} such that $\rho \vdash T_i \rightarrow v_i$ for each $i < n$; by induction, $\rho^\dagger \vdash T_i^\dagger \rightarrow v_i^\dagger$ and $\rho^\dagger \vdash T^\dagger \rightarrow v_0^\dagger \circ \dots \circ v_{n-1}^\dagger$. But this is exactly v^\dagger .
- If $T \equiv S[i, j]$, then there exists a u such that $\rho \vdash S \rightarrow u$. Let (u_0, \dots, u_{m-1}) be a decomposition of u into atomic values; then $v = u_i \circ \dots \circ u_j$. Now $\rho^\dagger \vdash S^\dagger \rightarrow u^\dagger$ by induction, and $u^\dagger = u_0^\dagger \circ \dots \circ u_{m-1}^\dagger$. Suppose that (t_0, \dots, t_{n-1}) is a decomposition of u^\dagger into atomic values. By definition of s_α , where α is the type of S , $t_i \circ \dots \circ t_j = u_i^\dagger \circ \dots \circ u_j^\dagger = v^\dagger$. Hence, $\rho^\dagger \vdash S^\dagger[i, j] \rightarrow v^\dagger$, which is what we wanted to show.
- If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, then suppose that $\rho \vdash T_0 \rightarrow \top$. (The case $\rho \vdash T_0 \rightarrow \perp$ is similar.) Then, $\rho \vdash T_1 \rightarrow v$. By induction, $\rho^\dagger \vdash T_0^\dagger \rightarrow \top$ and $\rho^\dagger \vdash T_1^\dagger \rightarrow v^\dagger$; hence, $\rho^\dagger \vdash T^\dagger \rightarrow v^\dagger$.
- If $T \equiv f(S)$, let $f(x^f) = T^f$ be the recursive definition of f in p . Then, there exists a u such that $\rho \vdash S \rightarrow u$ and $[x^f = u] \vdash T^f \rightarrow v$. By induction $\rho^\dagger \vdash S^\dagger \rightarrow u^\dagger$ and $[(x^f)^\dagger = u^\dagger] \vdash (T^f)^\dagger \rightarrow v^\dagger$. Since $f^\dagger((x^f)^\dagger) = (T^f)^\dagger$ is the definition of f^\dagger in p^\dagger , $\rho^\dagger \vdash f^\dagger(S^\dagger) \rightarrow v^\dagger$. ■

Lemma (Lemma 8). *For any program p , environment ρ , string x , term T , and value v ,*

$$g(x), \lambda_f(|g(x)|), \rho \vdash_p T \rightarrow v \implies x, \mu(|x|), \rho_g \vdash_{p(g)} T_g \rightarrow v_g,$$

without collisions, if the derivation of the left-hand side has no collisions.

Proof For legibility, we will abbreviate $g(x), \lambda_f(|g(x)|), \rho \vdash_p$ by $\rho \vdash$ and $x, \mu(|x|), \rho_g \vdash_{p(g)}$ by $\rho_g \vdash_g$. The proof proceeds by induction on the size of the derivation of $\rho \vdash T \rightarrow v$ and is broken into cases depending on the form of T .

1. If $T \equiv x$, a single variable, then $T_g \equiv x_g$, and $v = \rho(x)$, so $v_g = \rho(x_g)$. Hence, $\rho_g \vdash_g T_g \rightarrow v_g$.
2. If $T \equiv f(T')$, let $\mathbf{f}(x) = T^{\mathbf{f}}$ be the recursive definition of \mathbf{f} in p . Then, there exists a w such that $\rho \vdash T' \rightarrow w$ and $[x = w] \vdash T^{\mathbf{f}} \rightarrow v$. By induction, $\rho_g \vdash_g T'_g \rightarrow w_g$ and $[x_g = w_g] \vdash T_g^{\mathbf{f}} \rightarrow v_g$.
But $\mathbf{f}_g(x_g) = T_g^{\mathbf{f}}$ is the recursive definition of \mathbf{f}_g in p_g . Hence, $\rho_g \vdash_g \mathbf{f}_g(T'_g) \rightarrow v_g$.
But $T_g \equiv \mathbf{f}_g(T'_g)$, so $\rho_g \vdash_g T_g \rightarrow v_g$.
3. If $T \equiv \text{true}, \text{false}, 0$, or 1 , then the conclusion is immediate.
4. If $T \equiv \varphi(T_0, T_1)$, for $\varphi \in \{+, -, \leq\}$, then there exists v_0, v_1 , such that $\rho \vdash T_i \rightarrow v_i$ for $i < 2$ and $\varphi(v_0, v_1) = v$ in $C(\lambda_f(|g(x)|))$. Since, by assumption, there are no collisions, $\varphi(v_0, v_1) = v$ in $C(\mu(|x|))$. Note that $v_i = (v_i)_g$ and $v = v_g$, when we identify $C(\lambda_f(|g(x)|))$ as a subset of $C(\mu(|x|))$.
By induction, $\rho_g \vdash_g (T_i)_g \rightarrow (v_i)_g$. Since $T_g \equiv \varphi((T_0)_g, (T_1)_g)$, $\rho_g \vdash T_g \rightarrow v_g$. Since $C(\mu(|x|))$ contains $C(\lambda_f(|g(x)|))$, no new collisions are introduced.
5. If $T \equiv \text{min}$, then v is the index 0 , so $T_g \equiv 0$, and v_g is $0 \in C(\mu(|x|))$. Hence, $\rho_g \vdash_g T_g \rightarrow v_g$.
6. If $T \equiv \text{max}$, then v is the index $|g(x)|$, so $v_g = |g(x)| \in C(\mu(|x|))$.
Since $x, \lambda_g(|x|) \vdash_{q_g} h_q \rightarrow |g(x)|$ with no collisions, $x, \mu(|x|) \vdash_{q_g} h_q \rightarrow |g(x)|$. Since q_g is a fragment of p_g , $x, \mu(|x|) \vdash_{p_g} h_q \rightarrow |g(x)|$, i.e., $\vdash_g h_q \rightarrow |g(x)|$. Hence, $\vdash_g h_q \rightarrow v_g$, but $T_g \equiv h_q$, so $\vdash_g T_g \rightarrow v_g$.
7. If $T \equiv \varphi(T')$ for $\varphi \in \{\text{P}, \text{null}\}$, then there exists an index w of $g(x)$ such that $\rho \vdash T' \rightarrow w$ and $v = \varphi(w)$. By induction, there exists some $w_g \in C(\mu(|x|))$ such that $v_g = w_g - 1$ (if $\varphi = \text{P}$) or $v_g \iff w_g = 0$ (if $\varphi = \text{null}$), and $\rho_g \vdash_g T'_g \rightarrow w_g$. In the former case, $\rho_g \vdash_g T'_g - 1 \rightarrow v_g$ and in the latter, $\rho_g \vdash_g (T'_g) \rightarrow v_g$; in either case, $\rho_g \vdash T_g \rightarrow v_g$.
8. If $T \equiv \text{bit}(T')$, then there is some w such that $\rho \vdash T' \rightarrow w$ and $(g(x))_{w-1} = v$. Notice that $v_g = v$ and w_g is natural number equivalent to w in the counting module $C(\mu(|x|))$. If we identify $C(\lambda_g(|x|))$ as a subtype of $C(\mu(|x|))$, w_g becomes a member of $C(\lambda_g(|x|))$ too. Then,

$$x, \lambda_g(|x|), [c = w_g] \vdash_{p_g} h_p(c) \rightarrow v,$$

moreover with no collisions. Hence,

$$x, \mu(|x|), [c = w_g] \vdash_{p_g} h_p(c) \rightarrow v,$$

and therefore $[c = w_g] \vdash_g h_p(c) \rightarrow v$.

By induction, $\rho_g \vdash T'_g \rightarrow w_g$. Hence, $\rho_g \vdash h_p(T'_g) \rightarrow v$; as $T_g \equiv h_p(T'_g)$ and $v = v_g$, this is what we wanted to show.

9. If $T \equiv (T_0, \dots, T_{n-1})$, then $v = (v_0, \dots, v_{n-1})$, where $\rho \vdash T_i \rightarrow v_i$. By induction, $\rho_g \vdash (T_i)_g \rightarrow (v_i)_g$, and $\rho_g \vdash T_g \rightarrow v_g$, as $T_g \equiv ((T_0)_g, \dots, (T_{n-1})_g)$ and $v_g = ((v_0)_g, \dots, (v_{n-1})_g)$.
10. If $T \equiv T'[i]$, then there exists some (v_0, \dots, v_{n-1}) such that $v = v_i$ and $\rho \vdash T' \rightarrow (v_0, \dots, v_{n-1})$. By induction, $\rho_g \vdash_g T'_g \rightarrow ((v_0)_g, \dots, (v_{n-1})_g)$, and therefore, $\rho_g \vdash_g T_g \rightarrow (v_i)_g$.
11. If $T \equiv \text{if } T_0 \text{ then } T_1 \text{ else } T_2$, then there exists some $b \in 2$ such that $\rho \vdash T_0 \rightarrow b$ and $\rho \vdash T_b \rightarrow v$. By induction, $\rho_g \vdash_g (T_0)_g \rightarrow b$ (as $b = b_g$) and $\rho_g \vdash (T_b)_g \rightarrow v_g$. Therefore, since $T_g \equiv \text{if } (T_0)_g \text{ then } (T_1)_g \text{ else } (T_2)_g$, $\rho_g \vdash T_g \rightarrow v_g$. ■

Appendix B: Bit-length programs and Turing machines

In this section, we sketch a proof of Theorem 2. We split the recursive/polynomial time correspondence and the tail-recursive/logarithmic space correspondence into two different parts, as they require different constructions. We try to give just enough detail so that these are perspicuous and unburdened by excess notation.

Theorem. *For any function $f : 2^* \rightarrow 1^*$, the following are equivalent:*

- *f is computable in polynomial time.*
- *There is a polynomially bounded $\lambda : \omega \rightarrow \omega$ and a bit-length program p such that $\llbracket p \rrbracket_\lambda(x) = f(x)$, without collisions, for any $x \in 2^*$.*

For any function $f : 2^ \times 1^* \rightarrow 2$ and polynomially bounded function $\pi : \omega \rightarrow \omega$, the following are equivalent:*

- *There is a polynomial-time computable function $g : 2^* \times 1^* \rightarrow 2$ such that $g(x, y) = f(x, y)$ for any string $x \in 2^*$ and $y < \pi(|x|)$.*
- *There is a polynomially bounded function $\lambda : \omega \rightarrow \omega$ and a bit-length program p such that $\llbracket p \rrbracket_\lambda(x, y) = f(x, y)$, without collisions, for any $x \in 2^*$ and $y < \pi(|x|)$.*

Proof Let f be a function of type $2^* \times (1^*)^k \rightarrow \alpha$ for some $k \in \omega$ and $\alpha \in \{1^*, 2\}$. (This generalizes both cases above.) Suppose that f is computable in polynomial time, and let \mathcal{M} be a Turing machine witnessing as much. We may assume that \mathcal{M} has read-only input tapes, a one-way write-only output tape if the output alphabet is 1^* , and any finite number of work tapes, each with tape alphabet Σ containing a blank character. All tapes of \mathcal{M} extend infinitely to the right; the heads of \mathcal{M} are initialized to the left; the work and output tapes are initialized to all blank.

We now describe a bit-length program simulating \mathcal{M} . The global input variable is identified with the single input of type 2^* . The tuple \vec{x} of counting module variables is identified with the remaining unary inputs; these keep getting passed around, unchanged, in each recursive call. For each tape of \mathcal{M} we define a recursive functions $\text{position} : C \rightarrow C$ and $\text{character} : C \times C \rightarrow \Sigma$. The function $\text{position}(t, \vec{x})$ describes the position of the head after t steps, and $\text{character}(t, p, \vec{x})$ describes the character on that tape after t steps at position p . In addition, there is a single recursive function $\text{state}(t, \vec{x})$ describing the state

of \mathcal{M} after t steps. The codomain of `state` is a finite type, which we can encode by a sufficiently wide tuple of booleans.

The *surface configuration* of \mathcal{M} at a given time t , consists of the current state (`state`(t, \vec{x})) and the characters read by the heads (`character`($t, \text{position}(t), \vec{x}$) for each tape). The surface configuration tells us what state \mathcal{M} transitions to, how the heads move, and what they write in each step. (Exactly how this happens is specified by the transition function of \mathcal{M} .) More precisely, for $t > 0$, `state`(t, \vec{x}) is given by the previous surface configuration, `position`(t, \vec{x}) = `position`($t - 1, \vec{x}$) + δ , where $\delta = 0, 1$, or -1 depending on the previous surface configuration, and `character`(t, p, \vec{x}) is equal to `character`($t - 1, p, \vec{x}$) unless `position`($t - 1$) = p , in which case it is given by the previous surface configuration.

The initial position of each head is 0. The initial characters all work and output tapes are blank, and the initial characters of the input tapes are given by the \vec{x} . Once \mathcal{M} reaches a halting state, the state, positions, and characters of the program remain fixed. Therefore, to name the “final” state, tape positions, etc., we need a function `seed` : C which is big enough to bound the largest possible running time on some given input.

Notice that we can always “copy” an index value to a counting module value in a bit-length program. In particular, we can copy the maximum index `max` into the counting module value n , the length of the input. Using the counting module primitives, we can define multiplication on counting modules. Then, all we need to do is construct $(n + 1)^k$ for a large enough constant k . All of this is done without collisions in the context of a sufficiently large (but still polynomially bounded) counting module.²¹

The final state is named by `state`(`seed`, \vec{x}), which is the output of \mathcal{M} if f is boolean-valued. If \mathcal{M} has output of type 1^* , then the output is `position`(`seed`, \vec{x}), the final position of the output tape. Since the output tape is one-way and write-only, this position is precisely the output string. Finally, this program contains no collisions, since it contains no counting module additions except those in the construction of `seed`.

Conversely, suppose that there is a polynomially bounded $\lambda : \omega \rightarrow \omega$ and a bit-length program p such that $\llbracket p \rrbracket_\lambda(x) = f(x)$, without collisions, for any $x \in 2^*$. For a fixed input x , there are at most polynomially many-in- $|x|$ possible environments ρ obtained by binding variables of p to indices of x and counting modules bounded by $\lambda(|x|)$. We make a table of all pairs (ρ, T) of environments and possible terms and then, using dynamic programming, evaluate each term on each environment. One of these is the output of the program, and the whole process takes polynomial time. ■

Theorem. *For any function $f : 2^* \rightarrow 1^*$, the following are equivalent:*

- f is computable in logarithmic space.
- There is a polynomially bounded $\lambda : \omega \rightarrow \omega$ and a tail-recursive bit-length program p such that $\llbracket p \rrbracket_\lambda(x) = f(x)$, without collisions, for any $x \in 2^*$.

For any function $f : 2^ \times 1^* \rightarrow 2$ and polynomially bounded function $\pi : \omega \rightarrow \omega$, the following are equivalent:*

²¹ In earlier papers (Jones, 2001; Kop & Simonsen, 2017), `seed` was originally a counting module *primitive* naming the largest element of the counting module. We did not follow that approach here, because its inclusion would break Lemma 1, the independence of collision-free computation from the bound of the counting module.

- There is a logarithmic-space computable function $g: 2^* \times 1^* \rightarrow 2$ such that $g(x, y) = f(x, y)$ for any string $x \in 2^*$ and $y < \pi(|x|)$.
- There is a polynomially bounded function $\lambda: \omega \rightarrow \omega$ and a tail-recursive bit-length program p such that $\llbracket p \rrbracket_\lambda(x, y) = f(x, y)$, without collisions, for any $x \in 2^*$ and $y < \pi(|x|)$.

Proof As above, let f be a function of type $2^* \times (1^*)^k \rightarrow \alpha$ for some $k \in \omega$ and $\alpha \in \{1^*, 2\}$, and let \mathcal{M} be a Turing machine computing f in logarithmic space. In case $\alpha = 1^*$, \mathcal{M} has a write-only one-way output tape; note that the logarithmic space bound does not apply to this tape.

For a fixed finite alphabet Δ , call a *short string* a Δ -string of length $O(\log n)$, for any given n . We can encode short strings by polynomially bounded-in- n counting modules, roughly, by identifying Δ with some finite initial segment of ω , and identifying bounded-width Δ -strings with base- Δ numerals.

Let the *configuration* of \mathcal{M} at a given time consist of the current state, the positions of each head, and the entire contents of each tape. (Contrast this with the surface configuration above.) The point is, in a logarithmic-space Turing machine on input of length n , the entire configuration of n can be encoded by a short string in n and thus in a polynomially bounded counting module. Moreover, we can compute the next configuration from the previous one by a tail-recursive program using the counting module primitives.

Ultimately, the program p looks something like this. As above, the single binary string input is identified with the (invisible) global input variable, and the string \vec{x} is identified with any remaining variables of type 1^* .

$$\begin{aligned} f_0(\vec{x}) &= f_1(\text{init}(\vec{x})) \\ f_1(c) &= \text{if halt}(c) \text{ then out}(c) \text{ else } f_1(\text{next}(c)), \end{aligned}$$

where $\text{halt}(c)$ detects whether the configuration c is halting, $\text{out}(c)$ correctly returns the output of a halting configuration, $\text{init}(\vec{x})$ is the initial configuration, and $\text{next}(c)$ is the configuration following c . This program is tail-recursive and correctly computes f .

Conversely, suppose that there is a polynomially bounded $\lambda: \omega \rightarrow \omega$ and a tail-recursive bit-length program p such that $\llbracket p \rrbracket_\lambda(x) = f(x)$, without collisions, for any $x \in 2^*$. It is a well-known fact that every tail-recursive program can be rewritten as a while program with a single while loop and no additional data structures (e.g., a stack). We can directly simulate a program with a single while loop on a Turing machine by, e.g., storing each program variable on a tape of the machine. But it takes only logarithmic space to encode indices and counting modules which are bounded polynomially in the length of the input. ■