

Probabilistic Termination by Monadic Affine Sized Typing^{*}

Ugo Dal Lago¹ and Charles Grellois²

¹ University of Bologna & INRIA Sophia Antipolis

`ugo.dallago@unibo.it`

² INRIA Sophia Antipolis

`charles.grellois@inria.fr`

Abstract. We introduce a system of monadic affine sized types, which substantially generalise usual sized types, and allows this way to capture probabilistic higher-order programs which terminate almost surely. Going beyond plain, strong normalisation without losing soundness turns out to be a hard task, which cannot be accomplished without a richer, quantitative notion of types, but also without imposing some affinity constraints. The proposed type system is powerful enough to type classic examples of probabilistically terminating programs such as random walks. The way typable programs are proved to be almost surely terminating is based on reducibility, but requires a substantial adaptation of the technique.

1 Introduction

Probabilistic models are more and more pervasive in computer science [1–3]. Moreover, the concept of algorithm, originally assuming determinism, has been relaxed so as to allow probabilistic evolution since the very early days of theoretical computer science [4]. All this has given impetus to research on probabilistic programming languages, which however have been studied at a large scale only in the last twenty years, following advances in randomized computation [5], cryptographic protocol verification [6, 7], and machine learning [8]. Probabilistic programs can be seen as ordinary programs in which specific instructions are provided to make the program evolve probabilistically rather than deterministically. The typical example are instructions for sampling from a given distribution toolset, or for performing probabilistic choice.

One of the most crucial properties a program should satisfy is *termination*: the execution process should be guaranteed to end. In (non)deterministic computation, this is easy to formalize, since any possible computation path is only considered qualitatively, and termination is a boolean predicate on programs: any non-deterministic program either terminates – in must or may sense – or it

^{*} This work is partially supported by the ANR projects 12IS02001 PACE and 14CE250005 ELICA.

does not. In probabilistic programs, on the other hand, any terminating computation path is attributed a probability, and thus termination becomes a *quantitative* property. It is therefore natural to consider a program terminating when its terminating paths form a set of measure one or, equivalently, when it terminates with maximal probability. This is dubbed “almost sure termination” (AST for short) in the literature [9], and many techniques for automatically and semi-automatically checking programs for AST have been introduced in the last years [10–13]. All of them, however, focus on imperative programs; while probabilistic functional programming languages are nowadays among the most successful ones in the realm of probabilistic programming [8]. It is not clear at all whether the existing techniques for imperative languages could be easily applied to functional ones, especially when higher-order functions are involved.

In this paper, we introduce a system of monadic affine sized types for a simple probabilistic λ -calculus with recursion and show that it guarantees the AST property for all typable programs. The type system, described in Section 4, can be seen as a non-trivial variation on Hughes et al.’s sized types [14], whose main novelties are the following:

- Types are generalised so as to be *monadic*, this way encapsulating the kind of information we need to type non-trivial examples. This information, in particular, is taken advantage of when typing recursive programs.
- Typing rules are *affine*: higher-order variables cannot be freely duplicated. This is quite similar to what happens when characterising polynomial time functions by restricting higher-order languages akin to the λ -calculus [15]. Without affinity, the type system is bound to be unsound for AST.

The necessity of both these variations is discussed in Section 2 below. The main result of this paper is that typability in monadic affine sized types entails AST, a property which is proved using an adaptation of the Girard-Tait reducibility technique [16]. This adaptation is technically involved, as it needs substantial modifications allowing to deal with possibly infinite and probabilistic computations. In particular, every reducibility set must be parametrized by a quantitative parameter p guaranteeing that terms belonging to this set terminate with probability at least p . The idea of parametrizing such sets already appears in work by the first author and Hofmann [17], in which a notion of realizability parametrized by resource monoids is considered. These realizability models are however studied in relation to linear logic and to the complexity of normalisation, and do not fit as such to our setting, even if they provided some crucial inspiration. In our approach, the fact that recursively-defined terms are AST comes from a continuity argument on this parameter: we can prove, by unfolding such terms, that they terminate with probability p for every $p < 1$, and continuity then allows to take the limit and deduce that they are AST. This soundness result is technically speaking the main contribution of this paper, and is described in Section 5.

An extended version with more details and proofs is available [18].

1.1 Related Works

Sized types have been originally introduced by Hughes, Pareto, and Sabry [14] in the context of reactive programming. A series of papers by Barthe and colleagues [19–21] presents sized types in a way similar to the one we will adopt here, although still for a deterministic functional language. Contrary to the other works on sized types, their type system is proved to admit a decidable type inference, see the unpublished tutorial [20]. Abel developed independently of Barthe and colleagues a similar type system featuring size informations [22]. These three lines of work allow polymorphism, arbitrary inductive data constructors, and ordinal sizes, so that data such as infinite trees can be manipulated. These three features will be absent of our system, in order to focus the challenge on the treatment of probabilistic recursive programs. Another interesting approach is the one of Xi’s Dependent ML [23], in which a system of lightweight dependent types allows a more liberal treatment of the notion of size, over which arithmetic or conditional operations may in particular be applied. Termination is ensured by checking during typing that a given metrics decreases during recursive calls. This system is well-adapted for practical termination checking and can be extended with mutual recursion, inductive types and polymorphism, but does not feature ordinal sizes. See [22] for a detailed comparison of the previously cited systems. Some works along these lines are able to deal with coinductive data, as well as inductive ones [14, 19, 22]. They are related to Amadio and Coupet-Grimal’s work on guarded types ensuring productivity of infinite structures such as streams [24]. None of these works deal with probabilistic computation, and in particular with almost sure termination.

There has been a lot of interest, recently, about probabilistic termination as a verification problem in the context of imperative programming [10–13]. All of them deal, invariably, with some form of while-style language without higher-order functions. A possible approach is to reduce AST for probabilistic programs to termination of non-deterministic programs [10]. Another one is to extend the concept of ranking function to the probabilistic case. Bournez and Garnier obtained in this way the notion of Lyapunov ranking function [25], but such functions capture a notion more restrictive than AST: *positive* almost sure termination, meaning that the program is AST and terminates in expected finite time. To capture AST, the notion of ranking supermartingale [26] has been used. Note that the use of ranking supermartingales allows to deal with programs which are both probabilistic and non-deterministic [11, 13] and even to reason about programs with real-valued variables [12].

Some recent work by Cappai, the first author, and Parisen Toldin [27, 28] introduce type systems ensuring that all typable programs can be evaluated in probabilistic polynomial time. This is too restrictive for our purposes. On the one hand, we aim at termination, and restricting to polynomial time algorithms would be an overkill. On the other hand, the above-mentioned type systems guarantee that the length of *all* probabilistic branches are uniformly bounded (by the same polynomial). In our setting, this would restrict the focus to terms in which infinite computations are forbidden, while we simply want the set of

such computations to have probability 0. In fact, the results we present in this paper can be seen as a first step towards a type system characterizing average polynomial time, in the style of implicit computational complexity [29].

2 Why is Monadic Affine Typing Necessary?

In this section, we justify the design choices that guided us in the development of our type system. As we will see, the nature of AST requires a significant and non-trivial extension of the system of sized types originally introduced to ensure termination in the deterministic case [14].

Sized Types for Deterministic Programs. The simply-typed λ -calculus endowed with a typed recursion operator `letrec` and appropriate constructs for the natural numbers, sometimes called PCF, is already Turing-complete, so that there is no hope to prove it strongly normalizing. Sized types [14] refine the simple type system by enriching base types with annotations, so as to ensure the termination of any recursive definition. Let us explain the idea of sizes in the simple, yet informative case in which the base type is `Nat`. Sizes are defined by the grammar

$$\mathfrak{s} ::= \mathfrak{i} \mid \infty \mid \widehat{\mathfrak{s}}$$

where \mathfrak{i} is a size variable and $\widehat{\mathfrak{s}}$ is the successor of the size \mathfrak{s} — with $\widehat{\infty} = \infty$. These sizes permit to consider decorations $\text{Nat}^{\mathfrak{s}}$ of the base type `Nat`, whose elements are natural numbers of size at most \mathfrak{s} . The type system ensures that the only constant value of type $\text{Nat}^{\widehat{\mathfrak{i}}}$ is 0, that the only constant values of type $\text{Nat}^{\widehat{\mathfrak{i}}}$ are 0 or $\underline{1} = S\ 0$, and so on. The type Nat^{∞} is the one of all natural numbers, and is therefore often denoted as `Nat`.

The crucial rule of the sized type system, which we present here following Barthe et al. [19], allows one to type recursive definitions as follows:

$$\frac{\Gamma, f : \text{Nat}^{\mathfrak{i}} \rightarrow \sigma \vdash M : \text{Nat}^{\widehat{\mathfrak{i}}} \rightarrow \sigma[\widehat{\mathfrak{i}}/\mathfrak{i}] \quad \mathfrak{i} \text{ pos } \sigma}{\Gamma \vdash \text{letrec } f = M : \text{Nat}^{\mathfrak{s}} \rightarrow \sigma[\mathfrak{s}/\mathfrak{i}]} \quad (1)$$

This typing rule ensures that, to recursively define the function $f = M$, the term M taking an input of size $\widehat{\mathfrak{i}}$ calls f on inputs of *strictly lesser* size \mathfrak{i} . This is for instance the case when typing the program

$$M_{DBL} = \text{letrec } f = \lambda x. \text{case } x \text{ of } \{ S \rightarrow \lambda y. S\ S\ (f\ y) \mid 0 \rightarrow 0 \}$$

computing recursively the double of an input integer, as the hypothesis of the fixpoint rule in a typing derivation of M_{DBL} is

$$f : \text{Nat}^{\mathfrak{i}} \rightarrow \text{Nat} \vdash \lambda x. \text{case } x \text{ of } \{ S \rightarrow \lambda y. S\ S\ (f\ y) \mid 0 \rightarrow 0 \} : \text{Nat}^{\widehat{\mathfrak{i}}} \rightarrow \text{Nat}$$

The fact that f is called on an input y of strictly lesser size \mathfrak{i} is ensured by the rule typing the `case` construction:

$$\frac{\Gamma \vdash x : \text{Nat}^{\widehat{\mathfrak{i}}} \quad \Gamma \vdash \lambda y. S\ S\ (f\ y) : \text{Nat}^{\mathfrak{i}} \rightarrow \text{Nat} \quad \Gamma \vdash 0 : \text{Nat}}{\Gamma \vdash \text{case } x \text{ of } \{ S \rightarrow \lambda y. S\ S\ (f\ y) \mid 0 \rightarrow 0 \} : \text{Nat}}$$

where $\Gamma = f : \text{Nat}^{\hat{i}} \rightarrow \text{Nat}$, $x : \text{Nat}^{\hat{i}}$. The soundness of sized types for strong normalization allows to conclude that M_{DBL} is indeed SN.

A Naïve Generalization to Probabilistic Terms. The aim of this paper is to obtain a probabilistic, *quantitative* counterpart to this soundness result for sized types. Note that unlike the result for sized types, which was focusing on *all* reduction strategies of terms, we only consider a *call-by-value* calculus³. Terms can now contain a probabilistic choice operator \oplus_p , such that $M \oplus_p N$ reduces to the term M with probability $p \in \mathbb{R}_{[0,1]}$, and to N with probability $1 - p$. The language and its operational semantics will be defined more extensively in Section 3. Suppose for the moment that we type the choice operator in a naïve way:

$$\text{Choice} \quad \frac{\Gamma \vdash M : \sigma \quad \Gamma \vdash N : \sigma}{\Gamma \vdash M \oplus_p N : \sigma}$$

On the one hand, the original system of sized types features subtyping, which allows some flexibility to “unify” the types of M and N to σ . On the other hand, it is easy to realise that *all* probabilistic branches would have to be terminating, without any hope of capturing interesting AST programs: nothing has been done to capture the *quantitative* nature of probabilistic termination. An instance of a term which is not strongly normalizing but which is almost-surely terminating — meaning that it normalizes with probability 1 — is

$$M_{BIAS} = \left(\text{letrec } f = \lambda x. \text{case } x \text{ of } \left\{ S \rightarrow \lambda y. f(y) \oplus_{\frac{2}{3}} (f(SSy)) \mid 0 \rightarrow 0 \right\} \right) \bar{n} \quad (2)$$

simulating a *biased random walk* which, on $x = m+1$, goes to m with probability $\frac{2}{3}$ and to $m+2$ with probability $\frac{1}{3}$. The naïve generalization of the sized type system only allows us to type the body of the recursive definition as follows:

$$f : \text{Nat}^{\hat{i}} \rightarrow \text{Nat}^{\infty} \vdash \lambda y. f(y) \oplus_{\frac{2}{3}} (f(SSy)) : \text{Nat}^{\hat{i}} \rightarrow \text{Nat}^{\infty} \quad (3)$$

and thus does not allow us to deduce any relevant information on the *quantitative* termination of this term: nothing tells us that the recursive call $f(SSy)$ is performed with a relatively low probability.

A Monadic Type System. Along the evaluation of M_{BIAS} , there is *indeed* a quantity which decreases during each recursive call to the function f : the *average* size of the input on which the call is performed. Indeed, on an input of size \hat{i} , f calls itself on an input of smaller size \hat{i} with probability $\frac{2}{3}$, and on an input of greater size \hat{i} with probability only $\frac{1}{3}$. To capture such a relevant *quantitative* information on the recursive calls of f , and with the aim to capture almost sure termination, we introduce a *monadic* type system, in which *distributions of types* can be used to type in a finer way the functions to be used recursively. Contexts $\Gamma | \Theta$ will be generated by a context Γ attributing sized types to any number

³ Please notice that choosing a reduction strategy is crucial in a probabilistic setting, otherwise one risks getting nasty forms of non-confluence [30].

of variables, while Θ will attribute a *distribution* of sized types to at most *one* variable — typically the one we want to use to recursively define a function. In such a context, terms will be typed by a distribution type, formed by combining the Dirac distributions of types introduced in the Axiom rules using the following rule for probabilistic choice:

$$\text{Choice} \quad \frac{\Gamma | \Theta \vdash M : \mu \quad \Gamma | \Psi \vdash N : \nu \quad \langle \mu \rangle = \langle \nu \rangle}{\Gamma | \Theta \oplus_p \Psi \vdash M \oplus_p N : \mu \oplus_p \nu}$$

The guard condition $\langle \mu \rangle = \langle \nu \rangle$ ensures that μ and ν are distributions of types decorating of the *same* simple type. Without this condition, there is no hope to aim for a decidable type inference algorithm.

The Fixpoint Rule. Using these monadic types, instead of the insufficiently informative typing (3), we can derive the sequent

$$f : \left\{ \left(\text{Nat}^i \rightarrow \text{Nat}^\infty \right)^{\frac{2}{3}}, \left(\text{Nat}^{\widehat{i}} \rightarrow \text{Nat}^\infty \right)^{\frac{1}{3}} \right\} \vdash \lambda y. f(y) \oplus_{\frac{2}{3}} (f(SSy)) : \text{Nat}^{\widehat{i}} \rightarrow \text{Nat}^\infty \quad (4)$$

in which the type of f contains finer information on the sizes of arguments over which it is called recursively, and with which probability. This information enables us to perform a first switch from a qualitative to a quantitative notion of termination: we will adapt the hypothesis

$$\Gamma, f : \text{Nat}^i \rightarrow \sigma \vdash M : \text{Nat}^{\widehat{i}} \rightarrow \sigma[\widehat{i}/i] \quad (5)$$

of the original fix rule (1) of sized types, expressing that f is called on an argument of size one less than the one on which M is called, to a condition meaning that there is probability 1 to call f on arguments of a lesser size *after enough iterations of recursive calls*. We therefore define a random walk associated to the distribution type μ of f , the *sized walk* associated to μ , and which is as follows for the typing (4):

- the random walk starts on 1, corresponding to the size \widehat{i} ,
- on an integer $n + 1$, the random walk jumps to n with probability $\frac{2}{3}$ and to $n + 2$ with probability $\frac{1}{3}$,
- 0 is stationary: on it, the random walk loops.

This random walk – as all sized walks will be – is an instance of *one-counter Markov decision problem* [31], so that it is decidable in polynomial time whether the walk reaches 0 with probability 1. We will therefore replace the hypothesis (5) of the `letrec` rule by the quantitative counterpart we just sketched, obtaining

$$\text{letrec} \quad \frac{\left\{ (\text{Nat}^{s_j} \rightarrow \nu[s_j/i])^{p_j} \mid j \in \mathcal{J} \right\} \text{ induces an AST sized walk}}{\Gamma | f : \left\{ (\text{Nat}^{s_j} \rightarrow \nu[s_j/i])^{p_j} \mid j \in \mathcal{J} \right\} \vdash V : \text{Nat}^{\widehat{i}} \rightarrow \nu[\widehat{i}/i]} \quad \Gamma, \Delta | \Theta \vdash \text{letrec } f = V : \text{Nat}^{\tau} \rightarrow \nu[\tau/i]$$

where we omit two additional technical conditions to be found in Section 4 and which justify the weakening on contexts incorporated to this rule. The resulting

type system allows to type a varieties of examples, among which the following program computing the geometric distribution over the natural numbers:

$$M_{EXP} = \left(\text{letrec } f = \lambda x.x \oplus_{\frac{1}{2}} S (f x) \right) 0 \quad (6)$$

and for which the decreasing quantity is the size of the set of probabilistic branches of the term making recursive calls to f . Another example is the unbiased random walk

$$M_{UNB} = \left(\text{letrec } f = \lambda x.\text{case } x \text{ of } \left\{ S \rightarrow \lambda y.f(y) \oplus_{\frac{1}{2}} (f(SSy)) \mid 0 \rightarrow 0 \right\} \right) n \quad (7)$$

for which there is no clear notion of decreasing measure during recursive calls, but which yet terminates almost surely, as witnessed by the sized walk associated to an appropriate derivation in the sized type system. We therefore claim that the use of this external guard condition on associated sized walks, allowing us to give a general condition of termination, is satisfying as it both captures an interesting class of examples, and is computable in polynomial time.

In Section 5, we prove that this shift from a qualitative to a quantitative hypothesis in the type system results in a shift from the soundness for strong normalization of the original sized type system to a soundness for its quantitative counterpart: *almost-sure termination*.

Why Affinity? To ensure the soundness of the `letrec` rule, we need one more structural restriction on the type system. For the sized walk argument to be adequate, we must ensure that the recursive calls of f are indeed precisely modelled by the sized walk, and this is not the case when considering for instance the following term:

$$M_{NAFF} = \left(\text{letrec } f = \lambda x.\text{case } x \text{ of } \left\{ S \rightarrow \lambda y.f(y) \oplus_{\frac{2}{3}} (f(SSy); f(SSy)) \mid 0 \rightarrow 0 \right\} \right) n \quad (8)$$

where the sequential composition $;$ is defined in this call-by-value calculus as $M; N = (\lambda x.\lambda y.0) M N$. Note that M_{NAFF} calls recursively f *twice* in the right branch of its probabilistic choice, and is not therefore modelled appropriately by the sized walk associated to its type. In fact, we would need a generalized notion of random walk to model the recursive calls of this process; it would be a random walk on *stacks* of integers. In the case where $n = 1$, the recursive calls to f can indeed be represented by a tree of stacks as depicted in Figure 1, where leftmost edges have probability $\frac{2}{3}$ and rightmost ones $\frac{1}{3}$. The root indicates that




Fig. 1. A Tree of Recursive Calls.

the first call on f was on the integer 1. From it, there is either a call of f on 0 which terminates, or *two* calls on 2 which are put into a stack of calls, and so on. We could prove that, without the *affine* restriction we are about to formulate, the term M_{NAFF} is typable with monadic sized types and the fixpoint rule we just designed. However, this term is not almost-surely terminating. Notice, indeed, that the *sum* of the integers appearing in a stack labelling a node of the tree in Figure 1 decreases by 1 when the left edge of probability $\frac{2}{3}$ is taken, and increases by *at least* 3 when the right edge of probability $\frac{1}{3}$ is taken. It follows that the expected increase of the sum of the elements of the stack during one step is at least $-1 \times \frac{2}{3} + 3 \times \frac{1}{3} = \frac{1}{3} > 0$. This implies that the probability that f is called on an input of size 0 after enough iterations is strictly less than 1, so that the term M_{NAFF} cannot be almost surely terminating.

Such general random processes have stacks as states and are rather complex to analyse. To the best of the authors' knowledge, they do not seem to have been considered in the literature. We also believe that the complexity of determining whether 0 can be reached almost surely in such a process, if decidable, would be very high. This leads us to the design of an *affine* type system, in which the management of contexts ensures that a given probabilistic branch of a term may only use at most once a given higher-order symbol. We do not however formulate restrictions on variables of simple type Nat , as affinity is only used on the *letrec* rule and thus on higher-order symbols. Remark that this is in the spirit of certain systems from implicit computational complexity [15, 32].

Another restriction imposed by this reduction of almost-sure termination checking for higher-order programs to almost-sure termination checking for one-counter Markov decision processes is the fact that we do not allow a general form of nested recursion. This restriction is encoded in the system by allowing *at most* one variable to have a distribution of types in the context. It follows that programs making use of mutual recursion can not be typed in this system.

3 A Simple Probabilistic Functional Programming Language

We consider the language λ_{\oplus} , which is an extension of the λ -calculus with recursion, constructors for the natural numbers, and a choice operator. In this section, we introduce this language and its operational semantics, and use them to define the crucial notion of *almost-sure termination*.

Terms and Values. Given a set of variables \mathcal{X} , terms and values of the language λ_{\oplus} are defined by mutual induction as follows:

$$\text{Terms:} \quad M, N, \dots ::= V \mid V W \mid \text{let } x = M \text{ in } N \mid M \oplus_p N \\ \mid \text{case } V \text{ of } \{S \rightarrow W \mid 0 \rightarrow Z\}$$

$$\text{Values:} \quad V, W, Z, \dots ::= x \mid 0 \mid S V \mid \lambda x.M \mid \text{letrec } f = V$$

where $x, f \in \mathcal{X}$, $p \in]0, 1[$. When $p = \frac{1}{2}$, we often write \oplus as a shorthand for $\oplus_{\frac{1}{2}}$. The set of terms is denoted Λ_{\oplus} and the set of values is denoted Λ_{\oplus}^V .

Terms of the calculus are assumed to be in A-normal form [33]. This allows to formulate crucial definitions in a simpler way, concentrating in the Let construct the study of the probabilistic behaviour of terms. We claim that all traditional constructions can be encoded in this formalism. For instance, the usual application $M N$ of two terms can be harmlessly recovered via the encoding let $x = M$ in (let $y = N$ in $x y$). In the sequel, we write $c \vec{V}$ when a value may be either 0 or of the shape $S V$.

Term Distributions. The introduction of a probabilistic choice operator in the syntax leads to a *probabilistic* reduction relation. It is therefore meaningful to consider the (operational) semantics of a term as a *distribution* of values modelling the outcome of *all* the finite probabilistic reduction paths of the term. For instance, the term M_{EXP} defined in (6) evaluates to the term distribution assigning probability $\frac{1}{2^{n+1}}$ to the value \underline{n} . Let us define this notion more formally:

Definition 1 (Distribution). A distribution on X is a function $\mathcal{D} : X \rightarrow [0, 1]$ satisfying the constraint $\sum \mathcal{D} = \sum_{x \in X} \mathcal{D}(x) \leq 1$, where $\sum \mathcal{D}$ is called the sum of the distribution \mathcal{D} . We say that \mathcal{D} is proper precisely when $\sum \mathcal{D} = 1$. We denote by \mathcal{P} the set of all distributions, would they be proper or not. We define the support $\mathsf{S}(\mathcal{D})$ of a distribution \mathcal{D} as: $\mathsf{S}(\mathcal{D}) = \{x \in X \mid \mathcal{D}(x) > 0\}$. When $\mathsf{S}(\mathcal{D})$ consists only of closed terms, we say that \mathcal{D} is a closed distribution. When it is finite, we say that \mathcal{D} is a finite distribution. We call Dirac a proper distribution \mathcal{D} such that $\mathsf{S}(\mathcal{D})$ is a singleton. We denote by 0 the null distribution, mapping every term to the probability 0.

When $X = \Lambda_{\oplus}$, we say that \mathcal{D} is a *term distribution*. In the sequel, we will use a more practical notion of *representation* of distributions, which enumerates the terms with their probabilities as a family of assignments. For technical reasons, notably related to the subject reduction property, we will also need *pseudo-representations*, which are essentially multiset-like decompositions of the representation of a distribution.

Definition 2 (Representations and Pseudo-Representations). Let $\mathcal{D} \in \mathcal{P}$ be of support $\{x_i \mid i \in \mathcal{I}\}$, where $x_i = x_j$ implies $i = j$ for every $i, j \in \mathcal{I}$. The representation of \mathcal{D} is the set $\mathcal{R} = \{x_i^{\mathcal{D}(x_i)} \mid i \in \mathcal{I}\}$ where $x_i^{\mathcal{D}(x_i)}$ is just an intuitive way to write the pair $(x_i, \mathcal{D}(x_i))$. A pseudo-representation of \mathcal{D} is any multiset $[y_j^{p_j} \mid j \in \mathcal{J}]$ such that

$$\forall j \in \mathcal{J}, y_j \in \mathsf{S}(\mathcal{D}) \quad \forall i \in \mathcal{I}, \mathcal{D}(x_i) = \sum_{y_j = x_i} p_j.$$

By abuse of notation, we will simply write $\mathcal{D} = [y_j^{p_j} \mid j \in \mathcal{J}]$ to mean that \mathcal{D} admits $[y_j^{p_j} \mid j \in \mathcal{J}]$ as pseudo-representation. Any distribution has a canonical pseudo-representation obtained by simply replacing the set-theoretic notation with the multiset-theoretic one.

Distributions support operations like affine combinations and sums – the latter being only a partial operation. We extend these operations to (pseudo)-representations, in a natural way. Distributions, endowed with the pointwise partial-order \preceq , form an ω -CPO, but not a lattice, since the join of two distributions is not guaranteed to exist.

Definition 3 (Value Decomposition of a Term Distribution). *Let \mathcal{D} be a term distribution. We write its value decomposition as $\mathcal{D} \stackrel{VD}{=} \mathcal{D}_{|V} + \mathcal{D}_{|T}$, where $\mathcal{D}_{|V}$ is the maximal subdistribution of \mathcal{D} whose support consists of values, and $\mathcal{D}_{|T} = \mathcal{D} - \mathcal{D}_{|V}$ is the subdistribution of “non-values” contained in \mathcal{D} .*

Operational Semantics. The semantics of a term will be the value distribution to which it reduces via the probabilistic reduction relation, iterated up to the limit. As a first step, we define the call-by-value reduction relation $\rightarrow_v \subseteq \mathcal{P} \times \mathcal{P}$ on Figure 2. Note that we write Dirac distributions simply as terms on the left side of \rightarrow_v , to improve readability. As usual, we denote by \rightarrow_v^n the n -th iterate of the relation \rightarrow_v , with \rightarrow_v^0 being the identity relation. We then define the relation \rightleftharpoons_v^n as follows. Let $\mathcal{D} \rightarrow_v^n \mathcal{E} \stackrel{VD}{=} \mathcal{E}_{|V} + \mathcal{E}_{|T}$. Then $\mathcal{D} \rightleftharpoons_v^n \mathcal{E}_{|V}$. Note that, for every $n \in \mathbb{N}$ and $\mathcal{D} \in \mathcal{P}$, there is a unique distribution \mathcal{E} such that $\mathcal{D} \rightarrow_v^n \mathcal{E}$. Moreover, $\mathcal{E}_{|V}$ is the only distribution such that $\mathcal{D} \rightleftharpoons_v^n \mathcal{E}_{|V}$.

Lemma 1. *Let $n, m \in \mathbb{N}$ with $n < m$. Let \mathcal{D}_n (resp \mathcal{D}_m) be the distribution such that $M \rightleftharpoons_v^n \mathcal{D}_n$ (resp $M \rightleftharpoons_v^m \mathcal{D}_m$). Then $\mathcal{D}_n \preceq \mathcal{D}_m$.*

Definition 4 (Semantics of a Term, of a Distribution). *The semantics of a distribution \mathcal{D} is the distribution $\llbracket \mathcal{D} \rrbracket = \sup_{n \in \mathbb{N}} (\{\mathcal{D}_n \mid \mathcal{D} \rightleftharpoons_v^n \mathcal{D}_n\})$. This supremum exists thanks to Lemma 1, combined with the fact that (\mathcal{P}, \preceq) is an ω -CPO. We define the semantics of a term M as $\llbracket M \rrbracket = \llbracket \{M^1\} \rrbracket$.*

We now have all the ingredients required to define the central concept of this paper, the one of almost-surely terminating term:

Definition 5 (Almost-Sure Termination). *We say that a term M is almost-surely terminating precisely when $\sum \llbracket M \rrbracket = 1$.*

4 Monadic Affine Sized Typing

Following the discussion from Section 2, we introduce in this section a non-trivial lifting of sized types to our probabilistic setting. As a first step, we design an *affine* simple type system for λ_{\oplus} . This means that no higher-order variable may be used more than once in the same probabilistic branch. However, variables of base type \mathbf{Nat} may be used freely. In spite of this restriction, the resulting system allows to type terms corresponding to any probabilistic Turing machine. In Section 4.2, we introduce a more sophisticated type system, which will be *monadic* and affine, and which will be sound for almost-sure termination as we prove in Section 5.

$$\begin{array}{c}
\frac{}{\text{let } x = V \text{ in } M \rightarrow_v \{(M[V/x])^1\}} \quad \frac{}{(\lambda x.M) V \rightarrow_v \{(M[V/x])^1\}} \\
\frac{}{M \oplus_p N \rightarrow_v \{M^p, N^{1-p}\}} \\
\frac{M \rightarrow_v \{L_i^{p_i} \mid i \in \mathcal{I}\}}{\text{let } x = M \text{ in } N \rightarrow_v \{(\text{let } x = L_i \text{ in } N)^{p_i} \mid i \in \mathcal{I}\}} \\
\frac{}{\text{case } S \text{ } V \text{ of } \{S \rightarrow W \mid 0 \rightarrow Z\} \rightarrow_v \{(W \text{ } V)^1\}} \\
\frac{}{\text{case } 0 \text{ of } \{S \rightarrow W \mid 0 \rightarrow Z\} \rightarrow_v \{(Z)^1\}} \\
\frac{}{(\text{letrec } f = V) (c \vec{W}) \rightarrow_v \left\{ \left(V[(\text{letrec } f = V)/f] (c \vec{W}) \right)^1 \right\}} \\
\frac{\mathcal{D} \stackrel{VD}{=} \{M_j^{p_j} \mid j \in \mathcal{J}\} + \mathcal{D}|_V \quad \forall j \in \mathcal{J}, M_j \rightarrow_v \mathcal{E}_j}{\mathcal{D} \rightarrow_v \left(\sum_{j \in \mathcal{J}} p_j \cdot \mathcal{E}_j \right) + \mathcal{D}|_V}
\end{array}$$

Fig. 2. Call-by-value reduction relation \rightarrow_v on distributions.

4.1 Affine Simple Types for λ_{\oplus}

The terms of the language λ_{\oplus} can be typed using a variant of the simple types of the λ -calculus, extended to type **letrec** and \oplus_p , but also restricted to an *affine* management of contexts. Recall that the constraint of affinity ensures that a given higher-order symbol is used at most *once* in a probabilistic branch. We define simple types over the base type **Nat** in the usual way: $\kappa, \kappa', \dots ::= \mathbf{Nat} \mid \kappa \rightarrow \kappa'$ where, by convention, the arrow associates to the right. Contexts Γ, Δ, \dots are sequences of simply-typed variables $x :: \kappa$. We write sequents as $\Gamma \vdash M :: \kappa$ to distinguish these sequents from the ones using distribution types appearing later in this section. Before giving the rules of the type system, we need to define two policies for contracting contexts: an affine and a general one.

Context Contraction. Contexts can be combined in two different ways. On the one hand, one can form the *non-affine contraction* $\Gamma \cup \Delta$ of two contexts, for which Γ and Δ are allowed to share some variables, but these variables must be attributed the *same type* in both contexts. On the other hand, one can form the *affine contraction* $\Gamma \uplus \Delta$, in which variables in common between Γ and Δ must be attributed the type **Nat**.

Var	$\frac{}{\Gamma, x :: \kappa \vdash x :: \kappa}$	$\frac{\Gamma \vdash V :: \text{Nat}}{\Gamma \vdash \mathbf{S} V :: \text{Nat}}$	$\frac{}{\Gamma \vdash 0 :: \text{Nat}}$
λ	$\frac{\Gamma, x :: \kappa \vdash M :: \kappa'}{\Gamma \vdash \lambda x.M :: \kappa \rightarrow \kappa'}$	$\frac{\Gamma \vdash V :: \kappa \rightarrow \kappa' \quad \Delta \vdash W :: \kappa}{\Gamma \uplus \Delta \vdash V W :: \kappa'}$	App
Choice	$\frac{\Gamma \vdash M :: \kappa \quad \Delta \vdash N :: \kappa}{\Gamma \cup \Delta \vdash M \oplus_p N :: \kappa}$		
Let	$\frac{\Gamma \vdash M :: \kappa \quad \Delta, x :: \kappa \vdash N :: \kappa'}{\Gamma \uplus \Delta \vdash \text{let } x = M \text{ in } N :: \kappa'}$		
Case	$\frac{\Gamma \vdash V :: \text{Nat} \quad \Delta \vdash W :: \text{Nat} \rightarrow \kappa \quad \Delta \vdash Z :: \kappa}{\Gamma \uplus \Delta \vdash \text{case } V \text{ of } \{ \mathbf{S} \rightarrow W \mid 0 \rightarrow Z \} :: \kappa}$		
letrec	$\frac{\Gamma, f :: \text{Nat} \rightarrow \kappa \vdash V :: \text{Nat} \rightarrow \kappa \quad \forall x \in \Gamma, x :: \text{Nat}}{\Gamma \vdash \text{letrec } f = V :: \text{Nat} \rightarrow \kappa}$		

Fig. 3. Affine simple types for λ_{\oplus} .

The Affine Type System. The affine simple type system is then defined in Figure 3. All the rules are quite standard. Higher-order variables can occur at most once in any probabilistic branch because all binary typing rules – except probabilistic choice – treat contexts affinely. We set $\Lambda_{\oplus}^V(\Gamma, \kappa) = \{V \in \Lambda_{\oplus}^V \mid \Gamma \vdash V :: \kappa\}$ and $\Lambda_{\oplus}(\Gamma, \kappa) = \{M \in \Lambda_{\oplus} \mid \Gamma \vdash M :: \kappa\}$. We simply write $\Lambda_{\oplus}^V(\kappa) = \Lambda_{\oplus}^V(\emptyset, \kappa)$ and $\Lambda_{\oplus}(\kappa) = \Lambda_{\oplus}(\emptyset, \kappa)$ when the terms or values are closed. These closed, typable terms enjoy subject reduction and the progress property.

4.2 Monadic Affine Sized Types

This section is devoted to giving the basic definitions and results about monadic affine sized types (MASTs, for short), which can be seen as decorations of the affine simple types with some *size information*.

Sized Types. We consider a set \mathcal{S} of *size variables*, denoted i, j, \dots and define *sizes* (called *stages* in [19]) as:

$$\mathfrak{s}, \mathfrak{t} ::= i \mid \infty \mid \widehat{\mathfrak{s}}$$

where $\widehat{\cdot}$ denotes the *successor* operation. We denote the iterations of $\widehat{\cdot}$ as follows: $\widehat{\widehat{\mathfrak{s}}}$ is denoted $\widetilde{\mathfrak{s}}^2$, $\widehat{\widetilde{\mathfrak{s}}}$ is denoted $\widetilde{\mathfrak{s}}^3$, and so on. By definition, at most one variable $i \in \mathcal{S}$ appears in a given size \mathfrak{s} . We call it its *spine variable*, denoted as $\text{spine}(\mathfrak{s})$. We write $\text{spine}(\mathfrak{s}) = \emptyset$ when there is no variable in \mathfrak{s} . An order \preceq on sizes can

be defined as follows:

$$\frac{}{s \preccurlyeq s} \quad \frac{s \preccurlyeq \tau \quad \tau \preccurlyeq t}{s \preccurlyeq t} \quad \frac{}{s \preccurlyeq \widehat{s}} \quad \frac{}{s \preccurlyeq \infty}$$

Notice that these rules imply notably that $\widehat{\infty}$ is equivalent to ∞ , i.e., $\widehat{\infty} \preccurlyeq \infty$ and $\infty \preccurlyeq \widehat{\infty}$. We consider sizes modulo this equivalence. We can now define sized types and distribution types by mutual induction, calling distributions of (sized) types the distributions over the set of sized types:

Definition 6 (Sized Types, Distribution Types). *Sized types and distribution types are defined by mutual induction, contextually with the function $\langle \cdot \rangle$ which maps any sized or distribution type to its underlying affine type.*

$$\begin{array}{ll} \text{Sized types:} & \sigma, \tau ::= \sigma \rightarrow \mu \mid \mathbf{Nat}^s \\ \text{Distribution types:} & \mu, \nu ::= \{ \sigma_i^{p_i} \mid i \in \mathcal{I} \}, \\ \text{Underlying map:} & \langle \sigma \rightarrow \mu \rangle = \langle \sigma \rangle \rightarrow \langle \mu \rangle \\ & \langle \mathbf{Nat}^s \rangle = \mathbf{Nat} \\ & \langle \{ \sigma_i^{p_i} \mid i \in \mathcal{I} \} \rangle = \langle \sigma_j \rangle \end{array}$$

For distribution types we require additionally that $\sum_{i \in \mathcal{I}} p_i \leq 1$, that \mathcal{I} is a finite non-empty set, and that $\langle \sigma_i \rangle = \langle \sigma_j \rangle$ for every $i, j \in \mathcal{I}$. In the last equation, j is any element of \mathcal{I} .

The definition of sized types is *monadic* in that a higher-order sized type is of the shape $\sigma \rightarrow \mu$ where σ is again a sized type, and μ is a *distribution* of sized types.

Contexts and Operations on Them. Contexts are sequences of variables together with a sized type, and at most one distinguished variable with a distribution type:

Definition 7 (Contexts). *Contexts are of the shape $\Gamma \mid \Theta$, with*

$$\begin{array}{ll} \text{Sized contexts:} & \Gamma, \Delta, \dots ::= \emptyset \mid x : \sigma, \Gamma \quad (x \notin \text{dom}(\Gamma)) \\ \text{Distribution contexts:} & \Theta, \Psi, \dots ::= \emptyset \mid x : \mu \end{array}$$

As usual, we define the domain $\text{dom}(\Gamma)$ of a sized context Γ by induction: $\text{dom}(\emptyset) = \emptyset$ and $\text{dom}(x : \sigma, \Gamma) = \{x\} \uplus \text{dom}(\Gamma)$. We proceed similarly for the domain $\text{dom}(\Theta)$ of a distribution context Θ . When a sized context $\Gamma = x_1 : \sigma_1, \dots, x_n : \sigma_n$ ($n \geq 1$) is such that there is a simple type κ with $\forall i \in \{1, \dots, n\}, \langle \sigma_i \rangle = \kappa$, we say that Γ is *uniform of simple type κ* . We write this as $\langle \Gamma \rangle = \kappa$.

We write Γ, Δ for the disjoint union of these sized contexts: it is defined whenever $\text{dom}(\Gamma) \cap \text{dom}(\Delta) = \emptyset$. We proceed similarly for Θ, Ψ , but note that due to the restriction on the cardinality of such contexts, there is the additional requirement that $\Theta = \emptyset$ or $\Psi = \emptyset$.

We finally define contexts as pairs $\Gamma \mid \Theta$ of a sized context and of a distribution context, with the constraint that $\text{dom}(\Gamma) \cap \text{dom}(\Theta) = \emptyset$.

Definition 8 (Probabilistic Sum). Let μ and ν be two distribution types. We define their probabilistic sum $\mu \oplus_p \nu$ as the distribution type $p \cdot \mu + (1 - p) \cdot \nu$. We extend this operation to a partial operation on distribution contexts:

- For two distribution types μ and ν such that $\langle \mu \rangle = \langle \nu \rangle$, we define $(x : \mu) \oplus_p (x : \nu) = x : \mu \oplus_p \nu$,
- $(x : \mu) \oplus_p \emptyset = x : p \cdot \mu$,
- $\emptyset \oplus_p (x : \mu) = x : (1 - p) \cdot \mu$,
- In any other case, the operation is undefined.

Definition 9 (Weighted Sum of Distribution Contexts). Let $(\Theta_i)_{i \in \mathcal{I}}$ be a non-empty family of distribution contexts and $(p_i)_{i \in \mathcal{I}}$ be a family of reals of $[0, 1]$. We define the weighted sum $\sum_{i \in \mathcal{I}} p_i \cdot \Theta_i$ as the distribution context $x : \sum_{i \in \mathcal{I}} p_i \cdot \mu_i$ when the following conditions are met:

1. $\exists x, \forall i \in \mathcal{I}, \Theta_i = x : \mu_i$,
2. $\forall (i, j) \in \mathcal{I}^2, \langle \Theta_i \rangle = \langle \Theta_j \rangle$,
3. and $\sum_{i \in \mathcal{I}} p_i \leq 1$,

In any other case, the operation is undefined.

We define the substitution $[\tau/i]$ of a size variable in a size or in a sized or distribution type in the expected way; see the long version [18] for details. A subtyping relation allows to lift the order \preceq on sizes to monadic sized types:

Definition 10 (Subtyping). We define the subtyping relation \sqsubseteq on sized types and distribution types as follows:

$$\frac{}{\sigma \sqsubseteq \sigma} \quad \frac{\mathfrak{s} \preceq \mathfrak{r}}{\text{Nat}^{\mathfrak{s}} \sqsubseteq \text{Nat}^{\mathfrak{r}}} \quad \frac{\tau \sqsubseteq \sigma \quad \mu \sqsubseteq \nu}{\sigma \rightarrow \mu \sqsubseteq \tau \rightarrow \nu}$$

$$\frac{\exists f : \mathcal{I} \rightarrow \mathcal{J}, (\forall i \in \mathcal{I}, \sigma_i \sqsubseteq \tau_{f(i)}) \text{ and } (\forall j \in \mathcal{J}, \sum_{i \in f^{-1}(j)} p_i \leq p'_j)}{\{\sigma_i^{p_i} \mid i \in \mathcal{I}\} \sqsubseteq \{\tau_j^{p'_j} \mid j \in \mathcal{J}\}}$$

Sized Walks and Distribution Types. As we explained in Section 2, the rule typing `letrec` in the monadic, affine type system relies on an external decision procedure, computable in polynomial time. This procedure ensures that the *sized walk* — a particular instance of *one-counter Markov decision process* (OC-MDP, see [31]), but which does not make use of non-determinism — associated to the type of the recursive function of interest indeed ensures almost sure termination. Let us now define the sized walk associated to a distribution type μ . For the precise connection with OC-MDPs, see the long version [18].

Definition 11 (Sized Walk). Let $\mathcal{I} \subseteq_{\text{fin}} \mathbb{N}$ be a finite set of integers. Let $\{p_i\}_{i \in \mathcal{I}}$ be such that $\sum_{i \in \mathcal{I}} p_i \leq 1$. These parameters define a Markov chain whose set of states is \mathbb{N} and whose transition relation is defined as follows:

- the state $0 \in \mathbb{N}$ is stationary (i.e. one goes from 0 to 0 with probability 1),
- from the state $s + 1 \in \mathbb{N}$ one moves:

- to the state $s + i$ with probability p_i , for every $i \in \mathcal{I}$;
- to 0 with probability $1 - (\sum_{i \in \mathcal{I}} p_i)$.

We call this Markov chain the sized walk on \mathbb{N} associated to $(\mathcal{I}, (p_i)_{i \in \mathcal{I}})$. A sized walk is almost surely terminating when it reaches 0 with probability 1 from any initial state.

Notably, checking whether a sized walk is terminating is relatively easy:

Proposition 1 (Decidability of AST for Sized Walks). *It is decidable in polynomial time whether a sized walk is AST.*

Proof. By encoding sized walks into OC-MDPs, which enjoy this property [31]. See the long version [18].

Definition 12 (From Types to Sized Walks). *Consider a distribution type $\mu = \{(\text{Nat}^{s_j} \rightarrow \nu_j)^{p_j} \mid j \in \mathcal{J}\}$ such that $\forall j \in \mathcal{J}$, $\text{spine}(s_j) = \mathbf{i}$. Then μ induces a sized walk, defined as follows. First, by definition, s_j must be of the shape $\widehat{\mathbf{i}}^{k_j}$ with $k_j \geq 0$ for every $j \in \mathcal{J}$. We set $\mathcal{I} = \{k_j \mid j \in \mathcal{J}\}$ and $q_{k_j} = p_j$ for every $j \in \mathcal{J}$. The sized walk induced by the distribution type μ is then the sized walk associated to $(\mathcal{I}, (q_i)_{i \in \mathcal{I}})$.*

Example 1. Let $\mu = \left\{ \left(\text{Nat}^{\mathbf{i}} \rightarrow \text{Nat}^\infty \right)^{\frac{1}{2}}, \left(\text{Nat}^{\widehat{\mathbf{i}}^2} \rightarrow \text{Nat}^\infty \right)^{\frac{1}{3}} \right\}$. Then the induced sized walk is the one associated to $(\{0, 2\}, (p_0 = \frac{1}{2}, p_2 = \frac{1}{3}))$. In other words, it is the random walk on \mathbb{N} which is stationary on 0, and which on non-null integers $i + 1$ moves to i with probability $\frac{1}{2}$, to $i + 2$ with probability $\frac{1}{3}$, and jumps to 0 with probability $\frac{1}{6}$. Note that the type μ , and therefore the associated sized walk, models a recursive function which calls itself on a size lesser by one unit with probability $\frac{1}{2}$, on a size greater by one unit with probability $\frac{1}{3}$, and which does not call itself with probability $\frac{1}{6}$.

Typing Rules. Judgements are of the shape $\Gamma \mid \Theta \vdash M : \mu$. When a distribution $\mu = \{\sigma^1\}$ is Dirac, we simply write it σ . The type system is defined in Figure 4. As earlier, we define sets of typable terms, and set $\Lambda_\oplus^{s,V}(\Gamma \mid \Theta, \sigma) = \{V \mid \Gamma \mid \Theta \vdash V : \sigma\}$, and $\Lambda_\oplus^s(\Gamma \mid \Theta, \mu) = \{M \mid \Gamma \mid \Theta \vdash M : \mu\}$. We abbreviate $\Lambda_\oplus^{s,V}(\emptyset \mid \emptyset, \sigma)$ as $\Lambda_\oplus^{s,V}(\sigma)$ and $\Lambda_\oplus^s(\emptyset \mid \emptyset, \mu)$ as $\Lambda_\oplus^s(\mu)$.

This sized type system is a refinement of the affine simple type system for λ_\oplus : if $x_1 : \sigma_1, \dots, x_n : \sigma_n \mid f : \mu \vdash M : \nu$, then it is easily checked that $x_1 :: \langle \sigma_1 \rangle, \dots, x_n :: \langle \sigma_n \rangle, f :: \langle \mu \rangle \vdash M :: \langle \nu \rangle$.

Lemma 2 (Properties of Distribution Types).

- $\Gamma \mid \Theta \vdash V : \mu \implies \mu$ is Dirac.
- $\Gamma \mid \Theta \vdash M : \mu \implies \mu$ is proper.

Var	$\frac{}{\Gamma, x : \sigma \Theta \vdash x : \sigma}$	$\frac{}{\Gamma x : \sigma \vdash x : \sigma}$	Var'
Succ	$\frac{\Gamma \Theta \vdash V : \text{Nat}^s}{\Gamma \Theta \vdash \mathbf{S} V : \widehat{\text{Nat}}^s}$	$\frac{}{\Gamma \Theta \vdash 0 : \text{Nat}^s}$	Zero
λ	$\frac{\Gamma, x : \sigma \Theta \vdash M : \mu}{\Gamma \Theta \vdash \lambda x.M : \sigma \rightarrow \mu}$	$\frac{\Gamma \Theta \vdash M : \mu \quad \mu \sqsubseteq \nu}{\Gamma \Theta \vdash M : \nu}$	Sub
App	$\frac{\Gamma, \Delta \Theta \vdash V : \sigma \rightarrow \mu \quad \Gamma, \Xi \Psi \vdash W : \sigma \quad \langle \Gamma \rangle = \text{Nat}}{\Gamma, \Delta, \Xi \Theta, \Psi \vdash V W : \mu}$		
Choice	$\frac{\Gamma \Theta \vdash M : \mu \quad \Gamma \Psi \vdash N : \nu \quad \langle \mu \rangle = \langle \nu \rangle}{\Gamma \Theta \oplus_p \Psi \vdash M \oplus_p N : \mu \oplus_p \nu}$		
Let	$\frac{\Gamma, \Delta \Theta \vdash M : \{ \sigma_i^{p_i} \mid i \in \mathcal{I} \} \quad \langle \Gamma \rangle = \text{Nat} \quad \Gamma, \Xi, x : \sigma_i \Psi_i \vdash N : \mu_i \quad (\forall i \in \mathcal{I})}{\Gamma, \Delta, \Xi \Theta, (\sum_{i \in \mathcal{I}} p_i \cdot \Psi_i) \vdash \text{let } x = M \text{ in } N : \sum_{i \in \mathcal{I}} p_i \cdot \mu_i}$		
Case	$\frac{\Gamma \emptyset \vdash V : \widehat{\text{Nat}}^s \quad \Delta \Theta \vdash W : \text{Nat}^s \rightarrow \mu \quad \Delta \Theta \vdash Z : \mu}{\Gamma, \Delta \Theta \vdash \text{case } V \text{ of } \{ \mathbf{S} \rightarrow W \mid 0 \rightarrow Z \} : \mu}$		
letrec	$\frac{\langle \Gamma \rangle = \text{Nat} \quad i \notin \Gamma \text{ and } i \text{ positive in } \nu \text{ and } \forall j \in \mathcal{J}, \text{ spine}(s_j) = i \quad \{ (\text{Nat}^{s_j} \rightarrow \nu[s_j/i])^{p_j} \mid j \in \mathcal{J} \} \text{ induces an AST sized walk}}{\Gamma f : \{ (\text{Nat}^{s_j} \rightarrow \nu[s_j/i])^{p_j} \mid j \in \mathcal{J} \} \vdash V : \widehat{\text{Nat}}^i \rightarrow \nu[\widehat{i}/i]} \quad \Gamma, \Delta \Theta \vdash \text{letrec } f = V : \text{Nat}^r \rightarrow \nu[\tau/i]}$		

Fig. 4. Affine distribution types for λ_{\oplus} .

Subject Reduction for Monadic Affine Sized Types. The type system enjoys a form of subject reduction adapted to the probabilistic case and more specifically to the fact that terms reduce to *distributions* of terms. Let us sketch the idea of this adapted subject reduction property on an example. Remark that the type system allows us to derive the sequent

$$\emptyset | \emptyset \vdash 0 \oplus 0 : \left\{ \left(\text{Nat}^{\widehat{s}} \right)^{\frac{1}{2}}, \left(\text{Nat}^{\widehat{r}} \right)^{\frac{1}{2}} \right\} \quad (9)$$

where this distribution type is formed by typing a copy of 0 with $\text{Nat}^{\widehat{s}}$ and the other with $\text{Nat}^{\widehat{r}}$. Then, the term $0 \oplus 0$ reduces to $\left\{ 0^{\frac{1}{2}} \right\} + \left\{ 0^{\frac{1}{2}} \right\} = \left\{ 0^1 \right\} = \llbracket 0 \oplus 0 \rrbracket$: the operational semantics collapses the two copies of 0 appearing during

the reduction. However, in the spirit of the usual subject reduction for deterministic languages, we would like to type the two copies of 0 appearing during the reduction with different types. We therefore use the notion of *pseudo-representation*: $\left[0^{\frac{1}{2}}, 0^{\frac{1}{2}}\right]$ is a pseudo-representation of $\llbracket 0 \oplus 0 \rrbracket$, and we attribute the type $\text{Nat}^{\widehat{s}}$ to the first element of this pseudo-representation and the type $\text{Nat}^{\widehat{t}}$ to the other, obtaining the following *closed distribution of typed terms*:

$$\left\{ \left(0 : \text{Nat}^{\widehat{s}}\right)^{\frac{1}{2}}, \left(0 : \text{Nat}^{\widehat{t}}\right)^{\frac{1}{2}} \right\} \quad (10)$$

We can then compute the *average* type of (10), which we call the *expectation type* of this closed distribution of typed terms:

$$\frac{1}{2} \cdot \left\{ \left(\text{Nat}^{\widehat{s}}\right)^1 \right\} + \frac{1}{2} \cdot \left\{ \left(\text{Nat}^{\widehat{t}}\right)^1 \right\} = \left\{ \left(\text{Nat}^{\widehat{s}}\right)^{\frac{1}{2}}, \left(\text{Nat}^{\widehat{t}}\right)^{\frac{1}{2}} \right\}$$

Remark that it coincides with the type of the initial term (9). This will be our result of subject reduction: when a closed term M of distribution type μ reduces to a distribution \mathcal{D} of terms, we can type all the terms appearing in a pseudo-representation of \mathcal{D} to obtain a closed distribution of typed terms whose expectation type is μ . Let us now introduce the definitions necessary to the formal statement of the subject reduction property.

Definition 13 (Distributions of Distribution Types, of Typed Terms).

- A distribution of distribution types is a distribution \mathcal{D} over the set of distribution types, and such that $\mu, \nu \in \mathcal{S}(\mathcal{D}) \Rightarrow \langle \mu \rangle = \langle \nu \rangle$.
- A distribution of typed terms, or typed distribution, is a distribution of typing sequents which are derivable in the monadic, affine sized type system. The representation of such a distribution has thus the following form: $\left\{ (\Gamma_i \mid \Theta_i \vdash M_i : \mu_i)^{p_i} \mid i \in \mathcal{I} \right\}$. In the sequel, we restrict to the uniform case in which all the terms appearing in the sequents are typed with distribution types of the same fixed underlying type.
- A distribution of closed typed terms, or closed typed distribution, is a typed distribution in which all contexts are $\emptyset \mid \emptyset$. In this case, we simply write the representation of the distribution as $\left\{ (M_i : \mu_i)^{p_i} \mid i \in \mathcal{I} \right\}$, or even as $(M_i : \mu_i)^{p_i}$ when the indexing is clear from context. We write pseudo-representations in a similar way.

Definition 14 (Expectation Types). Let $(M_i : \mu_i)^{p_i}$ be a closed typed distribution. We define its expectation type as the distribution type $\mathbb{E}((M_i : \mu_i)^{p_i}) = \sum_{i \in \mathcal{I}} p_i \mu_i$.

We can now state the main lemma of subject reduction:

Lemma 3 (Subject Reduction, Fundamental Lemma). Let $M \in A_{\oplus}^s(\mu)$ and \mathcal{D} be the unique closed term distribution such that $M \rightarrow_v \mathcal{D}$. Then there exists a closed typed distribution $\left\{ (L_j : \nu_j)^{p_j} \mid j \in \mathcal{J} \right\}$ such that

- $\mathbb{E}((L_j : \nu_j)^{p_j}) = \mu$,
- $[(L_j)^{p_j} \mid j \in \mathcal{J}]$ is a pseudo-representation of \mathcal{D} .

Note that the condition on expectations implies that $\bigcup_{j \in \mathcal{J}} \mathbb{S}(\nu_j) = \mathbb{S}(\mu)$.

The proof of this result, and its generalization to the iterated reduction of closed typed distributions, appear in the long version. They allow us to deduce the following property on the operational semantics of λ_{\oplus} -terms:

Theorem 1 (Subject Reduction). *Let $M \in \Lambda_{\oplus}^{\mathbb{S}}(\mu)$. Then there exists a closed typed distribution $\{(W_j : \sigma_j)^{p_j} \mid j \in \mathcal{J}\}$ such that*

- $\mathbb{E}((W_j : \sigma_j)^{p_j}) \preceq \mu$,
- and that $[(W_j)^{p_j} \mid j \in \mathcal{J}]$ is a pseudo-representation of $\llbracket M \rrbracket$.

Note that $\mathbb{E}((W_j : \sigma_j)^{p_j}) \preceq \mu$ since the semantics of a term may not be a proper distribution at this stage. In fact, it will follow from the soundness theorem of Section 5 that the typability of M implies that $\sum \llbracket M \rrbracket = 1$ and thus that the previous statement is an equality.

5 Typability Implies Termination: Reducibility Strikes Again

This section is technically the most advanced one of the paper, and proves that the typing discipline we have introduced indeed enforces almost sure termination. As already mentioned, the technique we will employ is a substantial generalisation of Girard-Tait’s reducibility. In particular, reducibility must be made quantitative, in that terms can be said to be reducible *with a certain probability*. This means that reducibility sets will be defined as sets parametrised by a real number p , called the *degree of reducibility* of the set. As Lemma 4 will emphasize, this degree of reducibility ensures that terms contained in a reducibility set parametrised by p terminate with probability at least p . These “intermediate” degrees of reducibility are required to handle the fixpoint construction, and show that recursively-defined terms that are typable are indeed AST — that is, that they belong to the appropriate reducibility set, parametrised by 1.

The first preliminary notion we need is that of a size environment:

Definition 15 (Size Environment). *A size environment is any function ρ from \mathcal{S} to $\mathbb{N} \cup \{\infty\}$. Given a size environment ρ and a size expression \mathfrak{s} , there is a naturally defined element of $\mathbb{N} \cup \{\infty\}$, which we indicate as $\llbracket \mathfrak{s} \rrbracket_{\rho}$:*

- $\llbracket \widehat{\mathfrak{i}}^n \rrbracket_{\rho} = \rho(\mathfrak{i}) + n$,
- $\llbracket \infty \rrbracket_{\rho} = \infty$.

In other words, the purpose of size environments is to give a semantic meaning to size expressions. Our reducibility sets will be parametrised not only on a probability, but also on a size environment.

Definition 16 (Reducibility Sets).

- For values of simple type Nat , we define the reducibility sets

$$\text{VRed}_{\text{Nat}^s, \rho}^p = \{S^n 0 \mid p > 0 \implies n < \llbracket s \rrbracket_\rho\}.$$

- Values of higher-order type are in a reducibility set when their applications to appropriate values are reducible terms, with an adequate degree of reducibility:

$$\text{VRed}_{\sigma \rightarrow \mu, \rho}^p = \left\{ V \in \Lambda_{\oplus}^V(\langle \sigma \rightarrow \mu \rangle) \mid \forall q \in (0, 1], \quad \forall W \in \text{VRed}_{\sigma, \rho}^q, \right. \\ \left. V W \in \text{TRed}_{\mu, \rho}^{pq} \right\}$$

- Distributions of values are reducible with degree p when they consist of values which are themselves globally reducible “enough”. Formally, $\text{DRed}_{\mu, \rho}^p$ is the set of finite distributions of values – in the sense that they have a finite support – admitting a pseudo-representation $\mathcal{D} = [(V_i)^{p_i} \mid i \in \mathcal{I}]$ such that, setting $\mu = \{(\sigma_j)^{p'_j} \mid j \in \mathcal{J}\}$, there exists a family $(p_{ij})_{i \in \mathcal{I}, j \in \mathcal{J}} \in [0, 1]^{|\mathcal{I}| \times |\mathcal{J}|}$ of probabilities and a family $(q_{ij})_{i \in \mathcal{I}, j \in \mathcal{J}} \in [0, 1]^{|\mathcal{I}| \times |\mathcal{J}|}$ of degrees of reducibility, satisfying:

1. $\forall i \in \mathcal{I}, \forall j \in \mathcal{J}, V_i \in \text{VRed}_{\sigma_j, \rho}^{q_{ij}}$,
2. $\forall i \in \mathcal{I}, \sum_{j \in \mathcal{J}} p_{ij} = p_i$,
3. $\forall j \in \mathcal{J}, \sum_{i \in \mathcal{I}} p_{ij} = \mu(\sigma_j)$,
4. $p \leq \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} q_{ij} p_{ij}$.

Note that (2) and (3) imply that $\sum \mathcal{D} = \sum \mu$. We say that $[(V_i)^{p_i} \mid i \in \mathcal{I}]$ witnesses that $\mathcal{D} \in \text{DRed}_{\mu, \rho}^p$.

- A term is reducible with degree p when its finite approximations compute distributions of values of degree of reducibility arbitrarily close to p :

$$\text{TRed}_{\mu, \rho}^p = \left\{ M \in \Lambda_{\oplus}(\langle \mu \rangle) \mid \forall 0 \leq r < p, \quad \exists \nu_r \preceq \mu, \quad \exists n_r \in \mathbb{N}, \right. \\ \left. M \Rightarrow_{\nu}^{n_r} \mathcal{D}_r \text{ and } \mathcal{D}_r \in \text{DRed}_{\nu_r, \rho}^r \right\}$$

Note that here, unlike to the case of DRed , the fact that $M \in \Lambda_{\oplus}(\langle \mu \rangle)$ implies that μ is proper.

The first thing to observe about reducibility sets as given in Definition 16 is that they only deal with closed terms, and not with arbitrary terms. As such, we cannot rely *directly* on them when proving AST for typable terms, at least if we want to prove it by induction on the structure of type derivations. We will therefore define in the sequel an extension of these sets to open terms, which will be based on these sets of closed terms, and therefore enjoy similar properties. The following lemma, relatively easy to prove, is crucial for the understanding of the reducibility sets, for that it shows that the degree of reducibility of a term gives information on the sum of its operational semantics:

Lemma 4 (Reducibility and Termination).

- Let $\mathcal{D} \in \text{DRed}_{\mu, \rho}^p$. Then $\sum \mathcal{D} \geq p$.
- Let $M \in \text{TRed}_{\mu, \rho}^p$. Then $\sum \llbracket M \rrbracket \geq p$.

It follows from this lemma that terms with degree of reducibility 1 are AST:

Corollary 1 (Reducibility and AST). *Let $M \in \text{TRed}_{\mu, \rho}^1$. Then M is AST.*

Fundamental Properties. Before embarking in the proof that typability implies reducibility, it is convenient to prove some fundamental properties of reducibility sets, which inform us about how these sets are structured, and which will be crucial in the sequel. First of all, if the degree of reducibility p is 0, then no assumption is made on the probability of termination of terms, distributions or values. It follows that the three kinds of reducibility sets collapse to the set of all affinely simply typable terms, distributions or values:

Lemma 5 (Candidates of Null Reducibility).

- If $V \in \Lambda_{\oplus}^V(\kappa)$, then $V \in \text{VRed}_{\sigma,\rho}^0$ for every σ such that $\langle \sigma \rangle = \kappa$ and every size environment ρ .
- Let $\mathcal{D} = \{(V_i)^{p_i} \mid i \in \mathcal{I}\}$ be a finite distribution of values. If $\forall i \in \mathcal{I}, V_i \in \Lambda_{\oplus}^V(\kappa)$, then $\mathcal{D} \in \text{DRed}_{\mu,\rho}^0$ for every μ such that $\langle \mu \rangle = \kappa$ and $\sum \mu = \sum \mathcal{D}$ and every ρ .
- If $M \in \Lambda_{\oplus}(\kappa)$, then $M \in \text{TRed}_{\mu,\rho}^0$ for μ such that $\langle \mu \rangle = \kappa$ and every ρ .

As p gives us a lower bound on the sum of the semantics of terms, it is easily guessed that a term having degree of reducibility p must also have degree of reducibility $q < p$. The following lemma makes this statement precise:

Lemma 6 (Downward Closure). *Let σ be a sized type, μ be a distribution type and ρ be a size environment. Let $0 \leq q < p \leq 1$. Then:*

- For any value V , $V \in \text{VRed}_{\sigma,\rho}^p \implies V \in \text{VRed}_{\sigma,\rho}^q$,
- For any finite distribution of values \mathcal{D} , $\mathcal{D} \in \text{DRed}_{\mu,\rho}^p \implies \mathcal{D} \in \text{DRed}_{\mu,\rho}^q$,
- For any term M , $M \in \text{TRed}_{\mu,\rho}^p \implies M \in \text{TRed}_{\mu,\rho}^q$.

To analyse the letrec construction, we will prove that, for every $\varepsilon \in (0, 1]$, performing enough unfoldings of the fixpoint allows to prove that the recursively-defined term is in a reducibility set parametrised by $1 - \varepsilon$. We will be able to conclude on the AST nature of recursive constructions using the following continuity lemma, proved using the theory of linear programming [18]:

Lemma 7 (Continuity). *Let σ be a sized type, μ be a distribution type and ρ be a size environment. Let $p \in (0, 1]$. Then:*

- $\text{VRed}_{\sigma,\rho}^p = \bigcap_{0 < q < p} \text{VRed}_{\sigma,\rho}^q$,
- $\text{DRed}_{\mu,\rho}^p = \bigcap_{0 < q < p} \text{DRed}_{\mu,\rho}^q$,
- $\text{TRed}_{\mu,\rho}^p = \bigcap_{0 < q < p} \text{TRed}_{\mu,\rho}^q$.

The last fundamental property about reducibility sets which will be crucial to treat the recursive case is the following, stating that the sizes appearing in a sized type may be recovered in the reducibility set by using an appropriate semantics of the size variables, and conversely:

Lemma 8 (Size Commutation). *Let i be a size variable, \mathfrak{s} be a size such that $\mathfrak{s} = \infty$ or that $\text{spine}(\mathfrak{s}) \neq i$ and ρ be a size environment. Then:*

- $\text{VRed}_{\sigma[\mathfrak{s}/i],\rho}^p = \text{VRed}_{\sigma,\rho[i \mapsto \llbracket \mathfrak{s} \rrbracket_\rho]}^p$,
- $\text{DRed}_{\mu[\mathfrak{s}/i],\rho}^p = \text{DRed}_{\mu,\rho[i \mapsto \llbracket \mathfrak{s} \rrbracket_\rho]}^p$,
- $\text{TRed}_{\mu[\mathfrak{s}/i],\rho}^p = \text{TRed}_{\mu,\rho[i \mapsto \llbracket \mathfrak{s} \rrbracket_\rho]}^p$.

Unfoldings. The most difficult step in proving all typable terms to be reducible is, unexpectedly, proving that terms involving *recursion* are reducible whenever their respective *unfoldings* are. This very natural concept expresses simply that any term in the form $\text{letrec } f = W$ is assumed to compute the fixpoint of the function defined by W .

Definition 17 (*n*-Unfolding). *Suppose that $V = (\text{letrec } f = W)$ is closed, then the n -unfolding of V is:*

- V if $n = 0$;
- $W[Z/f]$ if $n = m + 1$ and Z is the m -unfolding of V .

We write the set of unfoldings of V as $\text{Unfold}(V)$. Note that if V admits a simple type, then all its unfoldings have this same simple type as well. In the sequel, we implicitly consider that V is simply typed.

Any unfolding of $V = (\text{letrec } f = W)$ should behave like V itself: all unfoldings of V should be equivalent. This, however, cannot be proved using simply the operational semantics. It requires some work, and techniques akin to logical relations, to prove (see [18]) this behavioural equivalence between a recursive definition and its unfoldings.

Proposition 2 (Reducibility is Stable by Unfolding). *Let $n \in \mathbb{N}$ and $V = (\text{letrec } f = W)$ be a closed value. Suppose that Z is the n -unfolding of V . Then $V \in \text{VRed}_{\text{Nat}^s \rightarrow \mu, \rho}^p$ if and only if $Z \in \text{VRed}_{\text{Nat}^s \rightarrow \mu, \rho}^p$.*

Extension to Open Terms. We are now ready to extend the notion of reducibility set from the realm of *closed* terms to the one of *open* terms. This turns out to be subtle. The guiding intuition is that one would like to define a term M with free variables in \vec{x} to be reducible iff any closure $M[\vec{V}/\vec{x}]$ is itself reducible in the sense of Definition 16. What happens, however, to the underlying degree of reducibility p ? How do we relate the degrees of reducibility of \vec{V} with the one of $M[\vec{V}/\vec{x}]$? The answer is contained in the following definition:

Definition 18 (Reducibility Sets for Open Terms). *Suppose that Γ is a sized context in the form $x_1 : \sigma_1, \dots, x_n : \sigma_n$, and that y is a variable distinct from x_1, \dots, x_n . Then we define the following sets of terms and values:*

$$\begin{aligned} \text{OTRed}_{\mu, \rho}^{\Gamma | \emptyset} &= \left\{ M \mid \forall (q_i)_i \in [0, 1]^n, \forall (V_1, \dots, V_n) \in \prod_{i=1}^n \text{VRed}_{\sigma_i, \rho}^{q_i}, \right. \\ &\quad \left. M[\vec{V}/\vec{x}] \in \text{TRed}_{\mu, \rho}^{\prod_{i=1}^n q_i} \right\} \\ \text{OVRD}_{\mu, \rho}^{\Gamma | \emptyset} &= \left\{ W \mid \forall (q_i)_i \in [0, 1]^n, \forall (V_1, \dots, V_n) \in \prod_{i=1}^n \text{VRed}_{\sigma_i, \rho}^{q_i}, \right. \\ &\quad \left. W[\vec{V}/\vec{x}] \in \text{VRed}_{\mu, \rho}^{\prod_{i=1}^n q_i} \right\} \end{aligned}$$

$$\text{OTRed}_{\mu,\rho}^{\Gamma|y:\{\tau_j^{p_j}\}_{j \in \mathcal{J}}} = \left\{ M \mid \begin{array}{l} \forall (q_i)_i \in [0, 1]^n, \quad \forall \vec{V} \in \prod_{i=1}^n \text{VRed}_{\sigma_i, \rho}^{q_i}, \\ \forall (q'_j)_j \in [0, 1]^{\mathcal{J}}, \quad \forall W \in \bigcap_{j \in \mathcal{J}} \text{VRed}_{\tau_j, \rho}^{q'_j}, \\ M[\vec{V}, W/\vec{x}, y] \in \text{TRed}_{\mu, \rho}^{\alpha} \end{array} \right\}$$

$$\text{OVRed}_{\mu,\rho}^{\Gamma|y:\{\tau_j^{p_j}\}_{j \in \mathcal{J}}} = \left\{ Z \mid \begin{array}{l} \forall (q_i)_i \in [0, 1]^n, \quad \forall \vec{V} \in \prod_{i=1}^n \text{VRed}_{\sigma_i, \rho}^{q_i}, \\ \forall (q'_j)_j \in [0, 1]^{\mathcal{J}}, \quad \forall W \in \bigcap_{j \in \mathcal{J}} \text{VRed}_{\tau_j, \rho}^{q'_j}, \\ Z[\vec{V}, W/\vec{x}, y] \in \text{VRed}_{\mu, \rho}^{\alpha} \end{array} \right\}$$

where $\alpha = (\prod_{i=1}^n q_i) \left((\sum_{j \in \mathcal{J}} p_j q'_j) + 1 - (\sum_{j \in \mathcal{J}} p_j) \right)$ is called the degree of reducibility. Note that these sets extend the ones for closed terms: in particular, $\text{OTRed}_{\mu, \rho}^{\emptyset|\emptyset} = \text{TRed}_{\mu, \rho}^1$.

Lemma 9 (Reducible Values are Reducible Terms). *For every Γ, Θ, σ and $\rho, V \in \text{OVRed}_{\sigma, \rho}^{\Gamma|\Theta}$ if and only if $V \in \text{OTRed}_{\{\sigma^1\}, \rho}^{\Gamma|\Theta}$. An immediate consequence is that $\text{OVRed}_{\sigma, \rho}^{\Gamma|\Theta} \subseteq \text{OTRed}_{\{\sigma^1\}, \rho}^{\Gamma|\Theta}$.*

Reducibility and Sized Walks. To handle the fixpoint rule, we need to relate the notion of sized walk which guards it with the reducibility sets, and in particular with the degrees of reducibility we can attribute to recursively-defined terms.

Definition 19 (Probabilities of Convergence in Finite Time). *Let us consider a sized walk. We define the associated probabilities of convergence in finite time $(Pr_{n,m})_{n \in \mathbb{N}, m \in \mathbb{N}}$ as follows: $\forall n \in \mathbb{N}, \forall m \in \mathbb{N}$, the real number $Pr_{n,m}$ is defined as the probability that, starting from m , the sized walk reaches 0 in at most n steps.*

The point is that, for an AST sized walk, the more we iterate, the closer we get to reaching 0 in finite time n with probability 1.

Lemma 10 (Finite Approximations of AST). *Let $m \in \mathbb{N}$ and $\varepsilon \in (0, 1]$. Consider a sized walk, and its associated probabilities of convergence in finite time $(Pr_{n,m})_{n \in \mathbb{N}, m \in \mathbb{N}}$. If the sized walk is AST, there exists $n \in \mathbb{N}$ such that $Pr_{n,m} \geq 1 - \varepsilon$.*

The following lemma is the crucial result relating sized walks with the reducibility sets. It proves that, when the sized walk is AST, and after substitution of the variables of the context by reducible values in the recursively-defined term, we can prove the degree of reducibility to be any probability $Pr_{n,m}$ of convergence in finite time.

Lemma 11 (Convergence in Finite Time and letrec). *Consider the distribution type $\mu = \{ (\text{Nat}^{s_j} \rightarrow \nu[\mathfrak{s}_j/\mathfrak{i}])^{p_j} \mid j \in \mathcal{J} \}$. Let Γ be the sized context $x_1 : \text{Nat}^{\tau_1}, \dots, x_l : \text{Nat}^{\tau_l}$. Suppose that $\Gamma \mid f : \mu \vdash V : \text{Nat}^{\hat{\mathfrak{i}}} \rightarrow \nu[\hat{\mathfrak{i}}/\mathfrak{i}]$ and that*

μ induces an AST sized walk. Denote $(Pr_{n,m})_{n \in \mathbb{N}, m \in \mathbb{N}}$ its associated probabilities of convergence in finite time. Suppose that $V \in \text{OVRed}_{\text{Nat}^i \rightarrow \nu[\widehat{i}/i], \rho}^{\Gamma | f : \mu}$ for every ρ . Let $\vec{W} \in \prod_{i=1}^l \text{VRed}_{\text{Nat}^r i, \rho}^1$, then for every $(n, m) \in \mathbb{N}^2$, we have that

$$\text{letrec } f = V[\vec{W}/\vec{x}] \in \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[i \rightarrow m]}^{Pr_{n,m}}$$

Proof. We give a sketch of the proof, to be found in the long version [18]. The proof is by recurrence on n . The main case relies on the decomposition $Pr_{n+1, m'+1} = \sum_{j \in \mathcal{J}} p_j Pr_{n, m'+k_j} + 1 - (\sum_{j \in \mathcal{J}} p_j)$. The induction hypothesis allows then to state that for every $j \in \mathcal{J}$ we have $\text{letrec } f = V[\vec{W}/\vec{x}] \in \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[i \rightarrow m'+k_j]}^{Pr_{n, m'+k_j}}$. We use the Size Commutation lemma (Lemma 8) to obtain that $\text{letrec } f = V[\vec{W}/\vec{x}]$ is in an appropriate intersection of reducibility sets, and the hypothesis that $V \in \text{OVRed}_{\text{Nat}^i \rightarrow \nu[\widehat{i}/i], \rho[i \rightarrow m']}$ then implies that $V[\vec{W}, \text{letrec } f = V[\vec{W}/\vec{x}]/\vec{x}, f] \in \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[i \rightarrow m'+1]}^{Pr_{n+1, m'+1}}$ using the Size Commutation lemma once again. As this term is an unfolding of $\text{letrec } f = V[\vec{W}/\vec{x}]$, we conclude using Proposition 2. \square

When $m = \infty$, the previous lemma does not allow to conclude, and an additional argument is required. Indeed, it does not make sense to consider a sized walk beginning from ∞ : the meaning of this size is in fact *any integer*, not the ordinal ω . The following lemma justifies this vision by proving that, if a term is in a reducibility set for any finite interpretation of a size, then it is also in the set where the size is interpreted as ∞ .

Lemma 12 (Reducibility for Infinite Sizes). *Suppose that $i \text{ pos } \nu$ and that W is the value $\text{letrec } f = V$. If $W \in \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[i \rightarrow n]}^p$ for every $n \in \mathbb{N}$, then $W \in \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[i \rightarrow \infty]}^p$.*

All these fundamental lemmas allow us to prove the following proposition, which expresses that all typable terms are reducible and is the key step towards the fact that typability implies AST:

Proposition 3 (Typing Soundness). *If $\Gamma | \Theta \vdash M : \mu$, then $M \in \text{OTRed}_{\mu, \rho}^{\Gamma | \Theta}$ for every ρ . Similarly, if $\Gamma | \Theta \vdash V : \sigma$, then $V \in \text{OVRed}_{\sigma, \rho}^{\Gamma | \Theta}$ for every ρ .*

Proof. We proceed by induction on the derivation of the sequent $\Gamma | \Theta \vdash M : \mu$. When $M = V$ is a value, we know by Lemma 2 that $\mu = \{\sigma^1\}$; and we prove that $V \in \text{OVRed}_{\sigma, \rho}^{\Gamma | \Theta}$ for every ρ . By Lemma 9 we obtain that $V \in \text{OTRed}_{\mu, \rho}^{\Gamma | \Theta}$ for every ρ . We proceed by case analysis on the last rule of the derivation:

- **letrec:** Suppose that $\Gamma, \Delta | \Theta \vdash \text{letrec } f = V : \text{Nat}^r \rightarrow \nu[r/i]$. We treat the case where $\Delta = \Theta = \emptyset$. The general case is easily deduced using the downward-closure of the reducibility sets (Lemma 6). Let $\Gamma = x_1 :$

$\text{Nat}^{\tau_1}, \dots, x_n : \text{Nat}^{\tau_n}$. We need to prove that, for every family $(q_i)_i \in [0, 1]^n$ and every $(W_1, \dots, W_n) \in \prod_{i=1}^n \text{VRed}_{\text{Nat}^{\tau_i}, \rho}^{q_i}$, we have

$$(\text{letrec } f = V) [\vec{W}/\vec{x}] = \left(\text{letrec } f = V[\vec{W}/\vec{x}] \right) \in \text{VRed}_{\text{Nat}^{\tau} \rightarrow \nu[\tau/i], \rho}^{\prod_{i=1}^n q_i}$$

If there exists $i \in \mathcal{I}$ such that $q_i = 0$, the result is immediate as the term is simply-typed and Lemma 5 applies. Else, for every $i \in \mathcal{I}$, we have by definition that $\text{VRed}_{\text{Nat}^{\tau_i}, \rho}^{q_i} = \text{VRed}_{\text{Nat}^{\tau_i}, \rho}^1$. Since the sets VRed are downward-closed (Lemma 6), it is in fact enough to prove that for every $(W_1, \dots, W_n) \in \prod_{i=1}^n \text{VRed}_{\text{Nat}^{\tau_i}, \rho}^1$, we have

$$\text{letrec } f = V[\vec{W}/\vec{x}] \in \text{VRed}_{\text{Nat}^{\tau} \rightarrow \nu[\tau/i], \rho}^1$$

Moreover, by size commutation (Lemma 8),

$$\text{VRed}_{\text{Nat}^{\tau} \rightarrow \nu[\tau/i], \rho}^1 = \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[\mathbf{i} \mapsto \llbracket \tau \rrbracket_\rho]}^1$$

Let us therefore prove the stronger fact that, for *every* integer $m \in \mathbb{N} \cup \{\infty\}$,

$$\text{letrec } f = V[\vec{W}/\vec{x}] \in \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[\mathbf{i} \mapsto m]}^1$$

Now, the typing derivation gives us that $\Gamma \mid f : \mu \vdash V : \text{Nat}^{\widehat{\mathbf{i}}} \rightarrow \nu[\widehat{\mathbf{i}}/i]$ and that μ induces an AST sized walk. Denote $(Pr_{n,m})_{n \in \mathbb{N}, m \in \mathbb{N}}$ its associated probabilities of convergence in finite time. By induction hypothesis, $V \in \text{OVRed}_{\text{Nat}^i \rightarrow \nu[\widehat{\mathbf{i}}/i], \rho}^{\Gamma \mid f : \mu}$ for every ρ and we can apply Lemma 11. It follows that, for every $(n, m) \in \mathbb{N}$,

$$\text{letrec } f = V[\vec{W}/\vec{x}] \in \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[\mathbf{i} \mapsto m]}^{Pr_{n,m}}$$

Let $\varepsilon \in (0, 1)$. By Lemma 10, there exists $n \in \mathbb{N}$ such that $Pr_{n,m} \geq 1 - \varepsilon$. Using downward closure (Lemma 6) and quantifying over all the ε , we obtain

$$\text{letrec } f = V[\vec{W}/\vec{x}] \in \bigcap_{0 < \varepsilon < 1} \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[\mathbf{i} \mapsto m]}^{1-\varepsilon}$$

so that, by continuity of VRed (Lemma 7), we obtain

$$\text{letrec } f = V[\vec{W}/\vec{x}] \in \text{VRed}_{\text{Nat}^i \rightarrow \nu, \rho[\mathbf{i} \mapsto m]}^1 \quad (11)$$

for every $m \in \mathbb{N}$, allowing us to conclude. It remains however to treat the case where $m = \infty$. Since $\mathbf{i} \text{ pos } \nu$ and that (11) holds for every $m \in \mathbb{N}$, Lemma 12 applies and we obtain the result.

- **Other cases:** the other cases are treated in the long version [18]. □

This proposition, together with the definition of OTRed , implies the main result of the paper, namely that typability implies almost-sure termination:

Theorem 2. *Suppose that $M \in A_{\oplus}^s(\mu)$. Then M is AST.*

Proof. Suppose that $M \in A_{\oplus}^s(\mu)$, then by Proposition 3 we have $M \in \text{OTRed}_{\mu, \rho}^{\emptyset \mid \emptyset}$ for every ρ . By definition, $\text{OTRed}_{\mu, \rho}^{\emptyset \mid \emptyset} = \text{TRed}_{\mu, \rho}^1$. Corollary 1 then implies that M is AST.

6 Conclusions and Perspectives

We presented a type system for an affine, simply-typed λ -calculus enriched with a probabilistic choice operator, constructors for the natural numbers, and recursion. This affinity constraint implies that a given higher-order variable may occur (freely) at most once in any probabilistic branch of a program. The type system we designed decorates the affine simple types with *size information*, allowing to incorporate in the types relevant information about the recursive behaviour of the functions contained in the program. A guard condition on the typing rule for `letrec`, formulated with reference to an appropriate Markov chain, ensures that typable terms are AST. The proof of soundness of this type system for AST relies on a quantitative extension of the reducibility method, to accommodate sets of candidates to the infinitary and probabilistic nature of the computations we consider.

A first natural question is the one of the decidability of type inference for our system. In the deterministic case, this question was only addressed by Barthe and colleagues in an unpublished tutorial [20], and their solution is technically involved, especially when it comes to dealing with the fixpoint rule. We believe that their approach could be extended to our system of monadic sized types, and hope that it could provide a decidable type inference procedure for it. However, this extension will certainly be challenging, as we need to appropriately infer distribution types associated with AST sized walks in the `letrec` rule.

Another perspective would be to study the general, *non-affine* case. This is challenging, for two reasons. First, the system of size annotations needs to be more expressive in order to distinguish between various occurrences of a same function symbol in a same probabilistic branch. A solution would be to use the combined power of dependent types – which already allowed Xi to formulate an interesting type system for termination in the deterministic case [23] – and of linearity: we could use *linear dependent types* [34] to formulate an extension of the monadic sized type system keeping track of *how many* recursive calls are performed, and of the size of *each* recursive argument. The second challenge would then be to associate, in the typing rule for `letrec`, this information contained in linear dependent types with an appropriate random process. This random process should be kept decidable to guarantee that at least *derivation* checking can be automated, and there will probably be a trade-off between the duplication power we allow in programs and the complexity of deciding AST for the guard in the `letrec` rule.

The extension of our type system to deal with general inductive datatypes is essentially straightforward. Other perspectives would be to enrich the type system so as to be able to treat coinductive data, polymorphic types, or ordinal sizes, three features present in most system of sized types dealing with the traditional deterministic case, but which we chose not to address in this paper to focus on the already complex task of accommodating sized types to a probabilistic and higher-order framework.

References

1. Manning, C.D., Schütze, H.: Foundations of statistical natural language processing. MIT Press (2001)
2. Pearl, J.: Probabilistic reasoning in intelligent systems - networks of plausible inference. Morgan Kaufmann series in representation and reasoning. (1989)
3. Thrun, S.: Robotic mapping: A survey. In: Exploring Artificial Intelligence in the New Millenium, Morgan Kaufmann (2002)
4. de Leeuw, K., Moore, E.F., Shannon, C.E., Shapiro, N.: Computability by probabilistic machines. Automata Studies **34** (1956) 183–212
5. Motwani, R., Raghavan, P.: Randomized Algorithms. Cambridge University Press (1995)
6. Barthe, G., Grégoire, B., Béguelin, S.Z.: Formal certification of code-based cryptographic proofs. In Shao, Z., Pierce, B.C., eds.: POPL 2009, ACM (2009) 90–101
7. Barthe, G., Grégoire, B., Heraud, S., Béguelin, S.Z.: Computer-aided security proofs for the working cryptographer. In Rogaway, P., ed.: CRYPTO 2011. Volume 6841 of LNCS., Springer (2011) 71–90
8. Goodman, N.D., Mansinghka, V.K., Roy, D.M., Bonawitz, K., Tenenbaum, J.B.: Church: a language for generative models. In McAllester, D.A., Myllymäki, P., eds.: UAI 2008, AUAI Press (2008) 220–229
9. Bournez, O., Kirchner, C.: Probabilistic rewrite strategies. applications to ELAN. In Tison, S., ed.: RTA 2002. Volume 2378 of LNCS., Springer (2002) 252–266
10. Esparza, J., Gaiser, A., Kiefer, S.: Proving termination of probabilistic programs using patterns. In Madhusudan, P., Seshia, S.A., eds.: CAV 2012. Volume 7358 of LNCS., Springer (2012) 123–138
11. Fioriti, L.M.F., Hermanns, H.: Probabilistic termination: Soundness, completeness, and compositionality. In Rajamani, S.K., Walker, D., eds.: POPL 2015, ACM (2015) 489–501
12. Chatterjee, K., Fu, H., Novotný, P., Hasheminezhad, R.: Algorithmic analysis of qualitative and quantitative termination problems for affine probabilistic programs. In Bodík, R., Majumdar, R., eds.: POPL 2016, ACM (2016) 327–342
13. Chatterjee, K., Fu, H., Goharshady, A.K.: Termination analysis of probabilistic programs through positivstellensatz’s. In Chaudhuri, S., Farzan, A., eds.: CAV 2016. Volume 9779 of LNCS., Springer (2016) 3–22
14. Hughes, J., Pareto, L., Sabry, A.: Proving the correctness of reactive systems using sized types. In Boehm, H., Jr., G.L.S., eds.: POPL’96, ACM Press (1996) 410–423
15. Hofmann, M.: A mixed modal/linear lambda calculus with applications to Bellantoni-Cook safe recursion. In Nielsen, M., Thomas, W., eds.: CSL ’97. Volume 1414 of LNCS., Springer (1997) 275–294
16. Girard, J.Y., Taylor, P., Lafont, Y.: Proofs and Types. Cambridge University Press, New York, NY, USA (1989)
17. Dal Lago, U., Hofmann, M.: Realizability models and implicit complexity. Theor. Comput. Sci. **412**(20) (2011) 2029–2047
18. Dal Lago, U., Grellois, C.: Probabilistic termination by monadic affine sized typing (long version). <https://arxiv.org/abs/1701.04089> (2016)
19. Barthe, G., Frade, M.J., Giménez, E., Pinto, L., Uustalu, T.: Type-based termination of recursive definitions. MSCS **14**(1) (2004) 97–141
20. Barthe, G., Grégoire, B., Riba, C.: A tutorial on type-based termination. In Bove, A., Barbosa, L.S., Pardo, A., Pinto, J.S., eds.: ALFA Summer School 2008. Volume 5520 of LNCS., Springer (2008) 100–152

21. Barthe, G., Grégoire, B., Riba, C.: Type-based termination with sized products. In Kaminski, M., Martini, S., eds.: CSL 2008. Volume 5213 of LNCS., Springer (2008) 493–507
22. Abel, A.: Termination checking with types. *ITA* **38**(4) (2004) 277–319
23. Xi, H.: Dependent types for program termination verification. *Higher-Order and Symbolic Computation* **15**(1) (2002) 91–131
24. Amadio, R.M., Coupet-Grimal, S.: Analysis of a guard condition in type theory. In Nivat, M., ed.: FoSSaCS’98. Volume 1378 of LNCS., Springer (1998) 48–62
25. Bournez, O., Garnier, F.: Proving positive almost-sure termination. In Giesl, J., ed.: RTA 2005. Volume 3467 of LNCS., Springer (2005) 323–337
26. Chakarov, A., Sankaranarayanan, S.: Probabilistic program analysis with martingales. In Sharygina, N., Veith, H., eds.: CAV 2013. Volume 8044 of LNCS., Springer (2013) 511–526
27. Cappai, A., Dal Lago, U.: On equivalences, metrics, and polynomial time. In Kosowski, A., Walukiewicz, I., eds.: FCT 2015. Volume 9210 of LNCS., Springer (2015) 311–323
28. Dal Lago, U., Parisen Toldin, P.: A higher-order characterization of probabilistic polynomial time. *Inf. Comput.* **241** (2015) 114–141
29. Dal Lago, U.: A short introduction to implicit computational complexity. In: ESLLI 2010. (2011) 89–109
30. Dal Lago, U., Zorzi, M.: Probabilistic operational semantics for the lambda calculus. *RAIRO - Theor. Inf. and Applic.* **46**(3) (2012) 413–450
31. Brázdil, T., Brožek, V., Etessami, K., Kučera, A., Wojtczak, D.: One-counter markov decision processes. 21st ACM-SIAM Symposium on Discrete Algorithms (2010)
32. Dal Lago, U.: The geometry of linear higher-order recursion. In: LICS 2005, IEEE Computer Society (2005) 366–375
33. Sabry, A., Felleisen, M.: Reasoning about programs in continuation-passing style. *Lisp and Symbolic Computation* **6**(3-4) (1993) 289–360
34. Dal Lago, U., Gaboardi, M.: Linear dependent types and relative completeness. In: LICS 2011, IEEE Computer Society (2011) 133–142