

## Schedule Agnostic Semantics for Reactive Probabilistic Programming

GUILLAUME BAUDART, Université Paris Cité, Inria, CNRS, IRIF, France

LOUIS MANDEL, IBM Research, USA

CHRISTINE TASSON, ISAE-Supaero, Université de Toulouse, France

Synchronous languages are now a standard industry tool for critical embedded systems. Designers write high-level specifications by composing streams of values using block diagrams. These languages have been recently extended with Bayesian reasoning to program state-space models which compute a stream of distributions given a stream of observations. Yet, the semantics of probabilistic models is only defined for scheduled equations – a significant limitation compared to dataflow synchronous languages and block diagrams.

In this paper we propose a new operational semantics and a new denotational semantics for a probabilistic synchronous language that are both schedule agnostic. The key idea is to externalize the source of randomness and interpret a probabilistic expression as a stream of functions mapping random elements to a value and positive score. The operational semantics interprets expressions as state machines where mutually recursive equations are evaluated using a fixpoint operator. The denotational semantics directly manipulates streams and is thus a better fit to reason about program equivalence. We use the denotational semantics to prove the correctness of a program transformation required to run an optimized inference algorithm for state-space models with constant parameters.

CCS Concepts: • Software and its engineering → Semantics; Data flow languages; • Mathematics of computing → Sequential Monte Carlo methods; • Theory of computation → Probabilistic computation.

Additional Key Words and Phrases: Probabilistic programming languages, Dataflow synchronous languages, Probabilistic stream semantics, Assumed Parameter Filter

### ACM Reference Format:

Guillaume Baudart, Louis Mandel, and Christine Tasson. 2025. Schedule Agnostic Semantics for Reactive Probabilistic Programming. 1, 1 (May 2025), 44 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnn>

## 1 INTRODUCTION

Synchronous programming languages [7] were introduced for the design of critical embedded systems. In dataflow languages such as Lustre [39], system designers write high-level specifications by composing infinite streams of values, called *flows*. Flows progress *synchronously*, paced on a global logical clock. Specialized compilers generate efficient and correct-by-construction embedded code with strong guarantees on execution time and memory consumption. This approach was inspired by block diagrams, a popular notation to describe control systems [40]. Built on these ideas, Scade is now a standard tool in automotive and avionic industries to program safety critical embedded software [20]. The synchronous model of computation is also central for the discrete-time subset of Matlab/Simulink [43].

---

Authors' Contact Information: Guillaume Baudart, Université Paris Cité, Inria, CNRS, IRIF, France; Louis Mandel, IBM Research, USA; Christine Tasson, ISAE-Supaero, Université de Toulouse, France.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Probabilistic languages extend general purpose programming languages with probabilistic constructs for Bayesian reasoning [8, 25, 36, 37]. Following a Bayesian approach, a program describes a probability distribution, the *posterior* distribution, using initial beliefs on random variables, the *prior* distributions, that are conditioned on observations.

At the intersection of these two lines of research, ProbZelus [5] is a probabilistic extension of the synchronous dataflow language Zelus [16]. ProbZelus combines, in a single source program, deterministic controllers and probabilistic models that can interact with each other to perform *inference-in-the-loop*. A classic example is the *Simultaneous Localization and Mapping problem* (SLAM) [49] where an autonomous agent tries to infer both its position and a map of its environment to adapt its trajectory.

The probabilistic model of the SLAM involves two kinds of parameters. The position is a *state parameter* represented by a stream of random variables. At each instant, a new position must be estimated from the previous position and the observations. The map is a *constant parameter* represented by a random variable whose value is progressively refined from the *prior* distribution with each new observation. This type of problem mixing constant parameters and state parameters are instances of *State-Space Models* (SSM) [19]. Any ProbZelus program can be expressed as a SSM.

*Probabilistic semantics and scheduling.* In the original ProbZelus semantics expressions are interpreted as state machines [5]. Following [52], a probabilistic expression computes a stream of *measures*. The semantics of an expression with a set of local declarations integrates the semantics of the main expression over all possible values of the local variables. Unfortunately, this semantics yields nested integrals that are only well defined if the declarations are *scheduled*, i.e., ordered according to data dependencies.

This is a significant limitation compared to synchronous dataflow languages where sets of mutually recursive equations are not ordered: a key requirement for commercial synchronous data-flow languages where programs are written using a block diagram graphical interface. Scheduling should not depend on the placement of the blocks, this motivates their definition as mutually recursive equations. Besides, the compiler implements a series of source-to-source transformations which often introduces new variables in arbitrary order. Scheduling local declarations is one of the very last compilation passes [16]. The semantics of ProbZelus is thus far from what is exposed to the programmer and prevents reasoning about most program transformations and compilation passes.

In this paper, we show how to extend the schedule agnostic semantics of dataflow synchronous languages [11, 18] for probabilistic programming. We first define a new operational semantics where sets of equations in arbitrary order can be interpreted with a fixpoint operator. The key idea of our approach is to interpret probabilistic expressions as a stream of functions mapping random elements in  $[0, 1]$  to a value and positive score. Like its deterministic counterpart [18], this operational semantics is *executable* and has the same structure as the compiled code. A drawback of this semantics is that, at each step, probabilistic state machines compute measures. Proofs of program equivalence must relate measures of states through successive integrations by exhibiting a bisimulation [50]. We introduce a denotational semantics which abstracts away the state machines and directly manipulates streams which simplifies reasoning about program equivalence. This semantics extends the semantics used in the Vélus project to prove an end-to-end compiler for the synchronous language Lustre [11–13].

*Filtering and constant parameters.* As a case study we prove the correctness of a program transformation that is required to run an optimized inference algorithm. To estimate state parameters, Sequential Monte Carlo (SMC) techniques rely on random simulations and filtering to approximate the posterior distribution, i.e., voluntarily dropping information to re-center the inference on the most significant estimations. Unfortunately, this information loss negatively impacts constant parameters estimations.

Manuscript submitted to ACM

```

1  proba motion(f) = x where
2    rec init x = xo and init v = vo
3    and v = sample(gaussian(last v +. f *. h, sv))
4    and x = sample(gaussian(last x +. last v *. h, sx))

6  proba tracker(y_obs) = x where
7    rec x = motion(f)
8    and f = current_force(x)
9    and () = observe(gaussian(g(x), sy), y_obs)

11 node main(y_obs) = u where
12   rec x_dist = infer (tracker (y_obs))
13   and u = controller(x_dist)

```

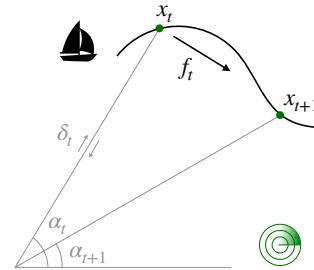


Fig. 1. Tracking a drifting boat with a marine radar in ProbZelus.  $x_t$  is the position of the boat in Cartesian coordinates at time  $t$ . The speed of the boat  $v_t$  is computed from the current force  $f_t$  which depends on the position of the boat. The observation  $g(x_t) = (\alpha_t, \delta_t)$  comprises the radar angle  $\alpha_t = \text{atan}(x_t[1]/x_t[0])$ , and the echo delay  $\delta_t = 2\|x_t\|/c$  where  $c$  is the speed of light.

Inspired by the Assumed Parameter Filter algorithm [35], we can split the inference into two steps: 1) estimate state parameters, and 2) update constant parameters. This technique requires a program transformation to explicitly separate constant from state parameters. A specialized static analysis identifies the constant parameters and their prior distributions. A compilation pass then transforms these parameters into additional inputs of the model. We use our denotational semantics to prove the correctness of this transformation, i.e., the transformation preserves the ideal semantics of the program.

*Contributions.* In this paper, we present the following main contributions:

- We introduce in Section 4 a new operational semantics and show that sets of mutually recursive probabilistic equations can be interpreted using a fixpoint operator.
- We introduce in Section 5 a new denotational semantics that abstracts away the state machines and directly manipulates streams which simplifies reasoning about program equivalence.
- We prove in Section 6 that these two semantics coincide, i.e., they both define the same streams, and that the operational semantics is equivalent to the original ProbZelus semantics for the subset of scheduled programs.
- We define in Section 7 a program transformation required to run an optimized inference algorithm for state-space models with constant parameters. We use our denotational semantics to prove the correctness of the transformation.

## 2 EXAMPLE

To motivate our approach, consider the ProbZelus program of Figure 1 adapted from [19][Section 2.4.1]. ProbZelus is a synchronous dataflow language à la Lustre [39] where programs execute in lock-step on a global logical clock. In this example, a discrete-time probabilistic model is used to track a drifting boat given noisy observations from a marine radar. At each instant, a rotating antenna sweeps a beam of microwaves and detects the boat when the beam is reflected back to the antenna. The radar then estimates the position of the boat from noisy measurements of its angle and of the echo delay. We assume that the strength of the current only depends on the position of the boat.

Manuscript submitted to ACM

*Probabilistic model.* The keyword `proba` indicates the definition of a probabilistic stream function. The `motion` model (Line 1) takes as input the stream `f` of the current force and returns the stream of estimated positions using two noisy integrators to compute the speed `v` and position `x` of the boat. Line 3 uses `sample` to specify that `v` is Gaussian distributed around the (discrete) integral of the current force `f` where `h` is the integration step, and `last v` refers to the previous speed of the boat initialized Line 2 with the `init` keyword. Similarly, Line 4 specifies that `x` is Gaussian distributed around the integral of `v`.

The `tracker` model uses the motion model to estimate the position of the boat (Line 7) and computes the current force from the position using a deterministic function `current_force` (Line 8). The `observe` operator (Line 9) conditions the model assuming that the observations `y_obs` are Gaussian distributed around `g(x)`, where `g` estimates the observable quantities (e.g., angle and echo delay) from the estimated position. The integration step `h`, the initial values `x0` and `v0`, and the noise parameters `sx`, `sv` and `sy` are global constants.

*Inference-in-the-loop.* The tracker model can be used to program a deterministic controller. The keyword `node` indicates the definition of a deterministic stream function that takes as input the stream of radar measurements and returns a stream of commands. Line 12, the explicit `infer` operator computes the stream of distributions `x_dist` defined by the probabilistic model `tracker` conditioned on the observations `y_obs`. Line 13 is an example of *inference-in-the-loop* where the distribution `x_dist` of estimated positions is used to compute the command `u` (e.g., using the mean or the variance of the distribution) that propels the boat which then influences the next observation.<sup>1</sup>

*Mutually recursive equations.* The original ProbZelus semantics focuses on a kernel language where local declarations are all *scheduled*, i.e., ordered according to data dependencies: an equation  $x_i = e_i$  must appear after  $x_j = e_j$  if  $e_j$  uses  $x_i$  without a `last` operator [5, Section 3.1]. But imposing a valid schedule is a significant limitation compared to synchronous dataflow languages which manipulate mutually recursive equations in arbitrary order. In particular, some programs can only be scheduled after several compilation passes which introduce additional equations to relax the scheduling constraints. For instance, in the `tracker` model, `x` depends on `f` (Line 7), but `f` also depends on `x` (Line 8). To break this dependency cycle, the compiler must inline the definition of `motion` to compute `x`, then `f`, then `v`. Such programs cannot be interpreted with the original semantics. In this paper, we propose a new operational semantics, and a new denotational semantics that are both defined on programs before compilation and overcome this limitation.

## 2.1 Operational semantics

We first propose in Section 4 a new operational semantics for ProbZelus inspired by existing probabilistic semantics [10, 38, 45, 48] where the source of randomness is externalized. Following the original ProbZelus co-iterative semantics [5] expressions are interpreted as state machines characterized by an initial state and a transition function. The interpretation of expressions depends on the context: probabilistic (e.g., inside the models `motion` or `tracker`) or deterministic (e.g., inside the controller `main`). Given the current state, the transition function of deterministic expressions returns the next state and a value. The transition function of probabilistic expressions takes one additional argument – a random element for every random variable – and returns one additional output – a positive score, or weight, which measures the quality of the output w.r.t. the model.

On the `motion` model of Figure 1, the state of the transition function memorizes  $(p_v, p_x)$  the previous values of `v` and `x` (required to interpret the `last` operator).<sup>2</sup> The additional argument  $[r_v, r_x]$  contains one random element for

<sup>1</sup>A more complex motion model could also use the value of the command.

<sup>2</sup>After simplifications, omitting empty states for stateless expressions, see Section 4.

each `sample` operator, i.e., an element  $r$  of the interval  $[0, 1]$  that can be mapped to a sample<sup>3</sup> of a distribution  $d$  using inverse transform sampling [32]. The model computes the samples  $v$  and  $x$  associated to the random elements, then returns the new state  $(v, x)$  (i.e., the next previous values for  $v$  and  $x$ ), the result  $x$ , and a weight 1 (there is no `observe`).

$$(\text{motion})_Y^{\text{step}}(f, (p_v, p_x), [r_v, r_x]) = \begin{aligned} & \text{let } \mu_v = \mathcal{N}(p_v + f \times h, s_v) \text{ in let } v = \text{icdf}_{\mu_v}(r_v) \text{ in} \\ & \text{let } \mu_x = \mathcal{N}(p_x + p_v \times h, s_x) \text{ in let } x = \text{icdf}_{\mu_x}(r_x) \text{ in} \\ & (v, x), x, 1 \end{aligned} \quad (1)$$

*Mutually recursive equations.* Following the original co-iterative semantics for dataflow synchronous languages [18] we interpret mutually recursive equations with a fixpoint operator in a flat complete partial order (CPO) where variables are either undefined or set to a value.

For instance, the semantics of the local declarations in `tracker` is the fixpoint of the following function  $F$  starting from the least element  $[x \leftarrow \perp, f \leftarrow \perp]$ . For simplicity, we assume that the semantics of `current_force` is a known stateless deterministic function  $c$  and the semantics of the observation function  $g$  is a stateless deterministic function  $g$ .

$$F(\rho) = \begin{aligned} & \text{let } (v, x), x, 1 = (\text{motion})_Y^{\text{step}}(\rho(f), (p_v, p_x), [r_v, r_x]) \text{ in} \\ & \quad \text{let } f = c(\rho(x)) \text{ in} \\ & \quad \text{let } w_y = \text{pdf}_{\mathcal{N}(g(\rho(x)), s_y)}(y_{\text{obs}}) \text{ in} \\ & \quad [x \leftarrow x, f \leftarrow f] \end{aligned} \quad \begin{aligned} \rho_0 &= [x \leftarrow \perp, f \leftarrow \perp] \\ \rho_1 &= [x \leftarrow \text{icdf}_{\mu_x}(r_x), f \leftarrow \perp] \\ \rho_2 &= [x \leftarrow \text{icdf}_{\mu_x}(r_x), f \leftarrow c(\text{icdf}_{\mu_x}(r_x))] \\ \rho_3 &= [x \leftarrow \text{icdf}_{\mu_x}(r_x), f \leftarrow c(\text{icdf}_{\mu_x}(r_x))] \end{aligned}$$

The fixpoint converges after 3 iterations. Since  $x$  only depends on the state  $(p_v, p_x)$  in Equation (1), the first iteration returns a value for  $x$  even if  $f$  is  $\perp$ . The second iteration can then use this value to compute  $f$ . Using the resulting environment, the semantics of `tracker` computes the next state, the result value, and the weight capturing the likelihood of the observation w.r.t. the estimation (the density of the distribution  $\mathcal{N}(g(x), s_y)$  at  $y_{\text{obs}}$ ).

$$(\text{tracker}(y_{\text{obs}}))_Y^{\text{step}}((p_v, p_x), [r_v, r_x]) = \begin{aligned} & \text{let } \rho = [x \leftarrow \text{icdf}_{\mu_x}(r_x), f \leftarrow c(\text{icdf}_{\mu_x}(r_x))] \text{ in} \\ & \quad \text{let } (v, x), x, 1 = (\text{motion})_Y^{\text{step}}(\rho(f), (p_v, p_x), [r_v, r_x]) \text{ in} \\ & \quad \text{let } f = c(\rho(x)) \text{ in} \\ & \quad \text{let } w_y = \text{pdf}_{\mathcal{N}(g(\rho(x)), s_y)}(y_{\text{obs}}) \text{ in} \\ & \quad (v, x), x, w_y \\ & = \text{let } x = \text{icdf}_{\mathcal{N}(p_x + p_v \times h, s_x)}(r_x) \text{ in} \\ & \quad \text{let } v = \text{icdf}_{\mathcal{N}(p_v + c(x) \times h, s_v)}(r_v) \text{ in} \\ & \quad (v, x), x, \text{pdf}_{\mathcal{N}(g(x), s_y)}(y_{\text{obs}}) \end{aligned} \quad (2)$$

The semantics of a sets of equations is defined via a fixpoint operator even when the solution is obvious like in the `motion` model. If equations are scheduled, the fixpoint yields the same environment as a sequence of local declarations.

*Unnormalized measure.* We obtain an unnormalized measure over pairs (next state, value) by integrating over all possible random elements, i.e., for a model  $e$ , if  $(e)_Y^{\text{step}}(m, r) = m', v, w$  we define  $\psi(m) = \int_{[0,1]^p} w \times \delta_{(m', v)} dr$  where  $p$  is the number of random variables in  $e$  and  $\delta_{(m', v)}$  is the Dirac distribution on the pair of state and value  $(m', v)$ . On

<sup>3</sup> $\text{icdf}_d$  is the generalized inverse cumulative distribution function of  $d$

the example of Figure 1, we get:

$$\begin{aligned}
 \psi(p_v, p_x) &= \int_{[0,1]} \int_{[0,1]} \text{let } (v, x), x, w = (\text{tracker}(y_{\text{obs}}))_Y^{\text{step}}((p_v, p_x), [r_v, r_x]) \text{ in } w \times \delta_{(v,x),x} dr_x dr_v \\
 &= \int_{[0,1]} \text{let } x = \text{icdf}_{\mathcal{N}(p_x + p_v \times h, s_x)}(r_x) \text{ in} \\
 &\quad \int_{[0,1]} \text{let } v = \text{icdf}_{\mathcal{N}(p_v + c(x) \times h, s_v)}(r_v) \text{ in} \\
 &\quad \quad \text{pdf}_{\mathcal{N}(g(x), s_y)}(y_{\text{obs}}) \times \delta_{(v,x),x} dr_x dr_v \\
 &= \int \mathcal{N}(p_x + p_v \times h, s_x)(dx) \int \mathcal{N}(p_v + c(x) \times h, s_v)(dv) \text{pdf}_{\mathcal{N}(g(x), s_y)}(y_{\text{obs}}) \times \delta_{(v,x),x}
 \end{aligned}$$

*Inference.* For a probabilistic model  $e$ ,  $\text{infer}(e)$  is a deterministic expression whose state is a measure over possible states for  $e$ . The transition function 1) integrates the unnormalized measure defined by the model over all possible states, 2) normalizes the measure, and 3) splits the result into a distribution of next states and a distribution of values. The runtime iterates this process from a distribution of initial states to compute a stream of distributions. We prove in Section 5.3 that this semantics coincide with the original ProbZelus semantics for the subset of scheduled programs, i.e., they both define the same streams.

*Program equivalence.* Since deterministic expressions are interpreted as state machines, to prove program equivalence we exhibit a bisimulation [50], i.e., a relation between the states of the two state machines. Two deterministic expressions are equivalent if there exists a relation such that 1) the initial states are in relation, and 2) given two states in relation the transition function produces new states in relation and the same output.

Two probabilistic expressions are equivalent if they describe the same stream of measures of output values obtained by integrating, at each step, the transition functions with respect to the measure over all possible states computed at the previous step. We show in Section 4.4 that we can prove the equivalence of probabilistic expressions with a bisimulation between states, and a *coupling* between the random elements.

## 2.2 Denotational semantics

The denotational semantics directly manipulates streams of values. This approach has been successfully used in the Vélus project<sup>4</sup> to prove an end-to-end compiler for the dataflow synchronous language Lustre [11–13, 15]. The semantics of a node is defined as a function between input streams and output streams. In Vélus, most of the compilation passes are proven correct using this denotational semantics. The translation to state machines is one of the very last passes and focuses on a normalized, scheduled subset of the language.

We extend this semantics to probabilistic models in Section 5. Given a context  $H$  mapping variable names to streams of values, and an array  $R$  of stream of random elements, the semantics of an expression returns a stream of pairs (value, weight):  $H, R \vdash e \Downarrow (v_0, w_0) \cdot (v_1, w_1) \cdot (v_2, w_2) \cdot \dots$  (the operator  $\cdot$  represents the concatenation of stream elements).

<sup>4</sup><https://velus.inria.fr>

We can interpret `tracker` of Figure 1 in a context  $H = [y_{\text{obs}} \leftarrow y_{\text{obs}}]$  that contains the stream of observations  $y_{\text{obs}}$  and in the random context given by the array  $R = [R_v, R_x]$  of two streams of independent random elements in  $[0, 1]$ .

$$\begin{aligned}
 [y_{\text{obs}} \leftarrow y_{\text{obs}}], [R_v, R_x] \vdash \text{tracker}(y_{\text{obs}}) &\Downarrow (x_0, w_0) \cdot (x_1, w_1) \cdot (x_2, w_2) \cdot \dots \\
 \text{where } \mu_x &= \mathcal{N}(x_{\text{init}} + v_{\text{init}} \times h, s_x) \cdot \mathcal{N}(x_0 + v_0 \times h, s_x) \cdot \mathcal{N}(x_1 + v_1 \times h, s_x) \cdot \dots \\
 x &= \text{icdf}_{\mu_{x0}}(R_{x0}) \cdot \text{icdf}_{\mu_{x1}}(R_{x1}) \cdot \text{icdf}_{\mu_{x2}}(R_{x2}) \cdot \dots \\
 \mu_v &= \mathcal{N}(v_{\text{init}} + c(x_0) \times h, s_x) \cdot \mathcal{N}(v_0 + c(x_1) \times h, s_x) \cdot \mathcal{N}(v_1 + c(x_2) \times h, s_x) \cdot \dots \\
 v &= \text{icdf}_{\mu_{v0}}(R_{v0}) \cdot \text{icdf}_{\mu_{v1}}(R_{v1}) \cdot \text{icdf}_{\mu_{v2}}(R_{v2}) \cdot \dots \\
 \mu_y &= \mathcal{N}(g(x_0), s_x) \cdot \mathcal{N}(g(x_1), s_x) \cdot \mathcal{N}(g(x_2), s_x) \cdot \dots \\
 w_y &= \text{pdf}_{\mu_{y0}}(y_{\text{obs}0}) \cdot \text{pdf}_{\mu_{y1}}(y_{\text{obs}1}) \cdot \text{pdf}_{\mu_{y2}}(y_{\text{obs}2}) \cdot \dots
 \end{aligned}$$

The semantics now directly manipulates streams. At each step, the result, a pair (value, weight), is similar to the expression in Equation (2), but states are abstracted away.

*Mutually recursive equations.* Given the random streams  $R$ , the semantics of a set of probabilistic mutually recursive equations  $H, R \vdash E : W$  checks that a context  $H$  mapping variable names to stream of values is compatible with all the equations in  $E$ , and that the combined weight of all sub-expressions is the stream  $W$ . Since variables in a context are not ordered, there is nothing special to do to interpret mutually recursive equations. By construction the order of equations does not matter which greatly simplifies reasoning about compilation passes that introduce new equations in arbitrary order. However, compared to the state machines of the operational semantics, the denotational semantics is not executable since equations are only checked a posteriori for a given context.

*Inference.* The semantics of `infer` now operates on a stream of pairs (value, weight):  $(v_0, w_0) \cdot (v_1, w_1) \cdot (v_2, w_2) \cdot \dots$ . The `infer` operator 1) associates to each value  $v_k$  the total weight of its prefix using a cumulative product  $\overline{w_k} = \prod_{i=0}^k w_i$ , 2) computes the un-normalized measure by integrating over all possible values of the stream of random elements, and 3) normalizes this measure to obtain a distribution of values. The key difference with the operational semantics is that the integral is now over the infinite domain of streams. We prove in Section 5.2 that this semantics is equivalent to the operational semantics, i.e., the `infer` operator yields the same stream of distributions.

*Program equivalence.* In the denotational semantics, states are abstracted away. Two deterministic expressions are equivalent if they define the same streams of output values. A probabilistic expression computes a stream of pairs (value, weight) where each element only depends on the random streams. Two probabilistic expressions are equivalent if they describe the same stream of measures over output values, obtained by integrating at each step the values and the weights over all possible random streams. We show in Section 5.4 that we can prove the equivalence of two probabilistic expressions with a coupling over the random streams.

### 2.3 Comparison

The following table summarizes the comparison between the original semantics and the two new semantics. The original ProbZelus semantics is similar to the new operational semantics but probabilistic state machines directly computes a stream of unnormalized measures. Compared to the original semantics, the new semantics are not limited to sets of scheduled equations. In the example of Figure 1 switching the order of equations `x` and `f` in `tracker` does not change the fixpoint. The new semantics are also modular and can interpret function call without inlining the body of

Manuscript submitted to ACM

```

 $d ::= \text{let } x = e \mid \text{node } f x = e \mid \text{proba } f x = e \mid d \ d$ 
 $e ::= c \mid x \mid (e, e) \mid op(e) \mid \text{last } x \mid f(e) \mid e \text{ where rec } E$ 
 $\mid \text{present } e \rightarrow e \text{ else } e \mid \text{reset } e \text{ every } e$ 
 $\mid \text{sample}(e) \mid \text{factor}(e) \mid \text{infer}(e)$ 
 $E ::= x = e \mid \text{init } x = e \mid E \text{ and } E$ 

```

Fig. 2. ProbZelus Syntax.

the callee as illustrated by the example of Figure 1. These semantics can even interpret programs that are not statically schedulable, e.g., programs where variable dependencies can change dynamically [22].

	original	operational	denotational
Executable	✓	✓	✗
Schedule agnostic	✗	✓	✓
Modularity	✗	✓	✓
Mutual recursion	✗	✓	✓
Randomness	internal	external	external
State	explicit	explicit	implicit

### 3 LANGUAGE

In this section we recall the syntax and typing rules of ProbZelus [5]. Advanced constructs, e.g., hierarchical automata, can be reduced to this language. The type system statically identifies deterministic and probabilistic expressions which have different interpretations in the operational semantics of Section 4 and the denotational semantics of Section 5. Recursion, loops, and nested inference are not allowed in the language [1].

#### 3.1 Syntax

The syntax of ProbZelus is presented in Figure 2. A program is a series of declarations  $d$ . A declaration can be a global variable `let`, a deterministic stream function `node`, or a probabilistic model `proba`. Each declaration has a unique name. An expression can be a constant  $c$ , a variable  $x$ , a pair, an operator application  $op(e)$ , the previous value of a variable `last`  $x$ , a function call  $f(e)$ , a local declaration  $e$  `where rec`  $E$  where  $E$  is a set of mutually recursive equations, a lazy conditional `present`  $e \rightarrow e_1 \text{ else } e_2$ , or a reset construct `reset`  $e_1 \text{ every } e_2$ . An equation is either a simple definition  $x = e$ , an initialization `init`  $x = e$  (the delay operator `last`  $x$  can only be used on initialized variables), or a set of equations  $E_1 \text{ and } E_2$ . In a set of equations, every initialized variable must be defined by another equation.

**Example.** The following node implements a fix-step integrator. The table presents the first few steps of a possible timeline with  $h = 0.1$ .

<code>node integr(x) = i where</code>	$x \mid 0 \quad 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad \dots$
<code>rec init i = 0. and i = (last i) +. x *. h</code>	$i \mid 0.0 \quad 0.1 \quad 0.3 \quad 0.6 \quad 1.0 \quad 1.5 \quad \dots$

We add the classic probabilistic constructs to the set of expressions: `sample(d)` creates a random variable with distribution  $d$ , `factor(s)` increments the log-density of the model, and `infer(m)` computes the posterior distribution of a model  $m$ . If  $d$  is a distribution with a density function, we use the syntactic shortcut `observe(d, x)` for `factor(pdfd(x))` which conditions the model on the assumption that  $x$  was sampled from the distribution  $d$  [52].

**Example.** The following model estimates the bias of a coin from a stream of coin tosses. Initially the model makes no assumption, i.e., the bias  $p$  is uniformly distributed over  $[0, 1]$  (0.5 corresponds to a perfect coin, 1 is a coin that always returns head). At each time step, we condition the model assuming that the observation  $o$  is sampled from a Bernoulli distribution with the same bias  $p$ . The deterministic node `main` then computes the mean of the posterior distributions of the bias at each time step.

```
proba coin (o) = p where
  rec init p = sample(uniform(0., 1.))
  and p = last p and () = observe(bernoulli(p), o)

  node main(o) = m where
    rec p_dist = infer(coin(o)) and m = mean(p_dist)      o | 1   0   0   0   0   ...
                                                               m | 0.68 0.51 0.34 0.29 0.25  ...

```

### 3.2 Typing

Figure 3 – adapted from [5, Figure 12] – presents the ProbZelus type system where each expression is associated to one of two *kinds*: D for deterministic, or P for probabilistic. This type system guarantees that probabilistic constructs `sample` and `factor` are only used inside `infer`.

The typing judgment  $G \vdash^k e : T$  asserts that, in a typing environment  $G$  mapping variables names to their types, the expression  $e$  has type  $T$  and kind  $k$ . Function type are annotated with the kind  $k$  of their body. The sub-typing rule can be used to lift a deterministic expression to a probabilistic expression. Deterministic expressions, e.g.,  $x + 1$ , can thus be used inside a probabilistic model (except `infer(e)` to prevent nested inference).

The type  $T \text{ dist}$  represents distributions over values of type  $T$  that can be sampled (`sample` operator). To simplify the presentation of the semantics, we assume that the arguments of the probabilistic operator `sample` and `factor` are deterministic. This is not a limitation since the compiler can always introduce intermediate variables to capture a probabilistic argument, e.g.,  $() = \text{factor}(\text{sample}(d))$  can be written  $x = \text{sample}(d)$  and  $() = \text{factor}(x)$ . The `infer` operator returns a deterministic value (a distribution) that can be used in a classic deterministic program (e.g., via standard functions like `mean` or `var`).

To type local declaration, equations are associated to the kind of their defining expression. The typing judgement  $G \vdash^k E : G'$  asserts that given an initial environment  $G$ , the set of equations  $E$  produces a new environment  $G'$  with kind  $k$ . We require that all equations have the same kind, which may necessitate the use of the sub-typing rule for certain expressions.

## 4 OPERATIONAL SEMANTICS

In this section, we first recall the original deterministic semantics. We then define a new operational semantics for probabilistic expressions. In this semantics we interpret sets of mutually recursive equations with a fixpoint operator. Finally, we define the equivalence of probabilistic programs using probabilistic bisimulation and coupling.

Manuscript submitted to ACM

$$\begin{array}{c}
\frac{\text{typeOf}(c) = t}{G \vdash^D c : t} \quad \frac{G(x) = t}{G \vdash^D x : t} \quad \frac{G(f) = t_1 \rightarrow^k t_2 \quad G \vdash^D e : t_1}{G \vdash^k f(e) : t_2} \\
\frac{G \vdash^D e : T \quad e \neq \text{infer}(e')}{G \vdash^P e : T} \quad \frac{G \vdash^D e : T \text{ dist}}{G \vdash^P \text{sample}(e) : T} \quad \frac{G \vdash^D e : \text{float}}{G \vdash^P \text{factor}(e) : \text{unit}} \quad \frac{G \vdash^P e : T}{G \vdash^D \text{infer}(e) : T \text{ dist}} \\
\frac{G \vdash^k e : t}{G \vdash^k x = e : [x \leftarrow t]} \quad \frac{G + G_1 + G_2 \vdash^k E_1 : G_1 \quad G + G_1 + G_2 \vdash^k E_2 : G_2}{G \vdash^k E_1 \text{ and } E_2 : G_1 + G_2} \quad \frac{G \vdash^k E : G' \quad G + G' \vdash^k e : t}{G \vdash^k e \text{ where rec } E : t}
\end{array}$$

Fig. 3. Typing rules (adapted from [5, Figure 12])

#### 4.1 Background: deterministic operational semantics

Following the co-iterative semantics of dataflow synchronous languages [18, 22], the deterministic semantics interprets expressions as state machines. The main advantage of the co-iterative semantics is that this interpretation is executable. Recent works demonstrated an interpreter from this semantics that can be used to test and validate a compiler against a reference semantics [22].

*State machine.* A deterministic expression  $e$  of type  $V$  is interpreted as a state machine characterized by an initial state  $\llbracket e \rrbracket^{\text{init}} : S$  and a transition function  $\llbracket e \rrbracket^{\text{step}} : \Gamma \rightarrow S \rightarrow S \times V$  which given the environment mapping variable names to values and the current state returns the next state and a value of type  $V$ . A stream of values is obtained by iteratively applying the transition function from the initial state.

$$(\llbracket e \rrbracket^{\text{init}} = m_0) \xrightarrow{\llbracket e \rrbracket^{\text{step}}_{Y_0}} m_1 \xrightarrow{\llbracket e \rrbracket^{\text{step}}_{Y_1}} m_2 \xrightarrow{\llbracket e \rrbracket^{\text{step}}_{Y_2}} m_3 \rightarrow \dots$$

$v_1$                      $v_2$                      $v_3$

Figure 4 shows a few semantics rules for deterministic expressions adapted from [22]. The state of an expression is defined inductively. The structure of the state matches the syntax of the expression and remains the same at each timestep. The semantics thus captures an aspect of the language that is essential for embedded systems: the memory required at runtime can be statically allocated before the initial step. The state of a variable is empty, and the transition function retrieves the value of the variable in the environment  $\gamma$  at each step. The state of a pair comprises the state of each sub-expression, and at each step the transition function executes the transition function of each sub-expression to compute the pair of next states and the pair of value. The initialization of a function call  $f(e)$  combines the initial state of the argument  $e$  with the initial state of  $e_f$  the body of  $f$ . The transition function evaluates the argument  $e$  into a value  $v_e$ , uses this value to evaluate the transition function of  $e_f$  in a fresh environment where the parameter  $x$  is bound to  $v_e$  which returns a result  $v$ , and returns the combined next states and the result.

Figure 5 shows the semantics rules for deterministic equations adapted from [22]. The state of an equation is also defined inductively. The equation  $x = e$  defines a single variable. The transition function evaluates the defining expression  $e$  into a pair (next state, value), and returns the next state and an environment where  $x$  is bound to the value. The `init`  $x = e$  equation manages the special variable `last`  $x$  which refers to the value of  $x$  at the previous time step. Compared to the original ProbZelus semantics, we do not require initial values to be constants. The state contains the previous value of  $x$  initialized with an undefined value of the correct type `nil`, and the initial state  $m_0$  of the expression  $e$ . There are two cases for the transition function. At the first time step (or after a `reset`), the state contains `nil` and the

$\llbracket e : V \rrbracket^{\text{init}}$	$: S$
$\llbracket e : V \rrbracket^{\text{step}}$	$: \Gamma \rightarrow S \times V$
$\llbracket c \rrbracket^{\text{init}}$	$= ()$
$\llbracket c \rrbracket^{\text{step}}_{\gamma} (())$	$= (), c$
$\llbracket x \rrbracket^{\text{init}}$	$= ()$
$\llbracket x \rrbracket^{\text{step}}_{\gamma} (())$	$= (), \gamma(x)$
$\llbracket \text{last } x \rrbracket^{\text{init}}$	$= ()$
$\llbracket \text{last } x \rrbracket^{\text{step}}_{\gamma} (())$	$= (), \gamma(x.\text{last})$
$\llbracket (e_1, e_2) \rrbracket^{\text{init}}$	$= \llbracket e_1 \rrbracket^{\text{init}}, \llbracket e_2 \rrbracket^{\text{init}}$
$\llbracket (e_1, e_2) \rrbracket^{\text{step}}_{\gamma} (m_1, m_2)$	$= \text{let } m'_1, v_1 = \llbracket e_1 \rrbracket^{\text{step}}_{\gamma} (m_1) \text{ in}$ $\text{let } m'_2, v_2 = \llbracket e_2 \rrbracket^{\text{step}}_{\gamma} (m_2) \text{ in}$ $(m'_1, m'_2), (v_1, v_2)$
$\llbracket f(e) \rrbracket^{\text{init}}$	$= \llbracket e \rrbracket^{\text{init}}, \llbracket f \rrbracket^{\text{init}}$ where $\text{node } f x = e_f$
$\llbracket f(e) \rrbracket^{\text{step}}_{\gamma} (m_e, m_f)$	$= \text{let } m'_e, v_e = \llbracket e \rrbracket^{\text{step}}_{\gamma} (m_e) \text{ in}$ $\text{let } m'_f, v = \llbracket f \rrbracket^{\text{step}}_{[x \leftarrow v_e]} (m_f) \text{ in}$ $(m'_e, m'_f), v$
$\llbracket e \text{ where rec } E \rrbracket^{\text{init}}$	$= \llbracket e \rrbracket^{\text{init}}, \llbracket E \rrbracket^{\text{init}}$
$\llbracket e \text{ where rec } E \rrbracket^{\text{step}}_{\gamma} (m, M)$	$= \text{let } F(\rho) = \left( \text{let } M', \rho' = \llbracket E \rrbracket^{\text{step}}_{\gamma+\rho} (M) \text{ in } \rho' \right) \text{ in}$ $\text{let } \rho = \text{fix}(F) \text{ in}$ $\text{let } M', \rho = \llbracket E \rrbracket^{\text{step}}_{\gamma+\rho} (M) \text{ in}$ $\text{let } m', v = \llbracket e \rrbracket^{\text{step}}_{\gamma+\rho} (m) \text{ in}$ $(m', M'), v$

Fig. 4. Operational semantics of deterministic expressions (adapted from [22]).

transition function evaluates  $e$  using  $m_0$  to compute the initial value  $i$  and returns a new state containing the current value of  $x$  and an environment where  $x.\text{last}$  is bound to  $i$ . In any other case, the previous value  $v$  of  $x$  stored in the state is defined. The transition function returns a new state containing the current value of  $x$  and an environment where  $x.\text{last}$  is bound to  $v$ . The initial state of the composition of two sets of equations  $E_1$  and  $E_2$  is the union of the states of  $E_1$  and  $E_2$ . The transition function evaluates  $E_1$  and  $E_2$  on their respective parts of the state but on the same environment  $\gamma$ . This function returns the updated state and the environment containing the variables defined in both sets of equations.

*Mutual recursions.* The original co-iterative semantics [18] and recent works [22] interpret mutually recursive equations using a fixpoint operator. Values  $v \in V$  are interpreted in a flat domain  $V_{\perp} = V + \{\perp\}$  with  $\perp$  as the minimal element and the flat order  $\leq: \forall v \in V. \perp \leq v$ .  $(V_{\perp}, \perp, \leq)$  is a complete partial order (CPO). This flat CPO

$$\begin{array}{ll}
\llbracket E \rrbracket^{\text{init}} & : S \\
\llbracket E \rrbracket^{\text{step}} & : \Gamma \rightarrow S \times \Gamma \\
\\
\llbracket x = e \rrbracket^{\text{init}} & = \llbracket e \rrbracket^{\text{init}} \\
\llbracket x = e \rrbracket^{\text{step}}_{\gamma}(m) & = \text{let } m', v = \llbracket e \rrbracket^{\text{step}}_{\gamma}(m) \text{ in } m', [x \leftarrow v] \\
\\
\llbracket \text{init } x = e \rrbracket^{\text{init}} & = \text{nil}, \llbracket e \rrbracket^{\text{init}} \\
\llbracket \text{init } x = e \rrbracket^{\text{step}}_{\gamma}(nil, m_0) & = \text{let } m', i = \llbracket e \rrbracket^{\text{step}}_{\gamma}(m_0) \text{ in } (\gamma(x), m_0), [x.\text{last} \leftarrow i] \\
\llbracket \text{init } x = e \rrbracket^{\text{step}}_{\gamma}(v, m_0) & = (\gamma(x), m_0), [x.\text{last} \leftarrow v] \\
\\
\llbracket E_1 \text{ and } E_2 \rrbracket^{\text{init}} & = \llbracket E_1 \rrbracket^{\text{init}}, \llbracket E_2 \rrbracket^{\text{init}} \\
\llbracket E_1 \text{ and } E_2 \rrbracket^{\text{step}}_{\gamma}(M_1, M_2) & = \text{let } M'_1, \gamma_1 = \llbracket E_1 \rrbracket^{\text{step}}_{\gamma}(M_1) \text{ in} \\
& \quad \text{let } M'_2, \gamma_2 = \llbracket E_2 \rrbracket^{\text{step}}_{\gamma}(M_2) \text{ in} \\
& \quad (M'_1, M'_2), \gamma_1 + \gamma_2
\end{array}$$

Fig. 5. Operational semantics of deterministic equations (adapted from [22]).

is lifted to environments defining the same set of variables:  $\forall \rho_1, \rho_2$  such that  $\text{dom}(\rho_1) = \text{dom}(\rho_2) = X$ ,  $\rho_1 \leq \rho_2$  iff  $\forall x \in X$ ,  $\rho_1(x) \leq \rho_2(x)$  and the least element is  $\perp = [x \leftarrow \perp]_{x \in X}$ .

To interpret an expression with a set of local declarations  $e$  **where rec**  $E$ , the transition function first computes the environment defined by  $E$  with a fixpoint operator. Given a state  $M$ , the function  $F(\rho) = \text{let } M', \rho' = \llbracket E \rrbracket^{\text{step}}_{\gamma+\rho}(M) \text{ in } \rho'$  is continuous and has a minimal fixpoint  $\rho = \text{fix}(F) = \lim_{n \rightarrow \infty} (F^n(\perp))$ . After convergence, the transition function evaluates  $\llbracket E \rrbracket^{\text{step}}_{\gamma+\rho}(M)$  once more to compute the next state  $M'$  (leaving  $\rho$  unchanged by definition of the fixpoint) and finally evaluates the main expression  $e$  in the environment  $\gamma + \rho$ .

*Causality.* For some programs, the fixpoint operator never converges. For instance,  $x$  **where rec**  $x = x$  returns  $\perp$  at every timestep. In the compiler, a static analysis ensures that dependency cycles between variables are always broken by a delay (**last** operator) [24]. This condition ensures that a valid schedule exists (after several compilation passes, e.g., inlining). If such a valid schedule exists, by monotonicity, each fixpoint iteration computes the value of at least one variable, the fixpoint is thus reached after a finite number of iterations [22].

*Definition 4.1.* A program is *causal* if all dependency cycles between variables are broken by a delay (**last** operator).

Compared to the original ProbZelus semantics, Definition 4.1 does not require that a valid schedule exists for all sets of mutually recursive equations. For instance, the program of Figure 1 is causal because the cycle  $x < f < v < f$  is broken by the delay **last**  $v$  in the definition of  $x$  (Section 2).

**Example.** Consider the **integr** node defined in Section 3.1. For an input expression  $e$  with initial state  $m_e$ , the initial state of the expression  $\text{integr}(e)$  is:

$$m = \frac{m_e, (((), ((\text{nil}, ())), (((), (((), ()))))))}{\begin{array}{ccccccc} \text{e} & \text{i} & \text{0} & \text{x} & \text{h} \\ \hline & & \text{last } \text{i} & & & & \end{array}}$$

$$\begin{array}{lll}
 |e : V| & : & \mathbb{N} \\
 |e| & = & 0 \quad \text{if } e \text{ is deterministic} \\
 |\text{sample}(e)| & = & 1 \\
 |\text{factor}(e)| & = & 0 \\
 |e \text{ where rec } E| & = & |e| + |E|
 \end{array}
 \quad
 \begin{array}{lll}
 |E| & : & \mathbb{N} \\
 |x = e| & = & |e| \\
 |\text{init } x = e| & = & |e| \\
 |E_1 \text{ and } E_2| & = & |E_1| + |E_2|
 \end{array}$$

Fig. 6. Dimension of probabilistic expressions and equations (full version in Figure 16 of the appendix).

During execution, the state at timestep  $t$  comprises the state of the input expression  $m_e$  and the previous value of  $i$ . The following table shows the first steps of the timeline.

	$m_e$	$m_{e0}$	$m_{e1}$	$m_{e2}$	...
$e$		0.0	1.0	2.0	...
$m_{\text{integr}}$	$m_{e0}, (((), (0.0, ()), (((), (((), ())))$	$m_{e1}, (((), ((0.1, ()), (((), (((), ())))$	$m_{e2}, (((), ((0.3, ()), (((), (((), ())))$	...	...
$i$	0.0		0.1	0.3	...

## 4.2 Probabilistic operational semantics

The key idea of the probabilistic operational semantics is to externalize all sources of randomness. Compared to the deterministic case, the transition function of a probabilistic expression takes one additional argument: an array of random elements with one random element for each random variable introduced by `sample`. To capture the effect of the `factor` operator, the transition function also returns a weight which measures the quality of the result w.r.t. the model.

*Model dimension.* The dimension of a probabilistic model is defined as the number of random variables in the model, i.e., the number of calls to the `sample` operator. Since loops and recursive calls are not allowed in the language of Figure 2, the dimension of a model  $e$  can be statically computed with the function  $|e|$  defined in Figure 6.

*Probabilistic state machine.* The initialization function of a probabilistic expression  $e$  of type  $V$ ,  $(e)^{\text{init}} : S$  returns the initial state. Given an environment, the current state, and a value for all random elements (an array of  $|e|$  values in  $[0, 1]$ ) the transition function  $(e)^{\text{step}} : \Gamma \rightarrow S \times [0, 1] \text{ array} \rightarrow S \times V \times [0, \infty)$  returns a triple (next state, value, weight).

An excerpt of the probabilistic operational semantics is presented in Figure 7. If  $e$  is deterministic, there is no random variable and no conditioning. The transition function takes an empty array of random elements, evaluates the expression, and returns the next state, the value, and a weight of 1. The transition function of `sample`( $e$ ) takes an array containing one random element, evaluates the argument  $e$  into a distribution  $\mu$ , converts the random element into a sample of the distribution using inverse transform sampling  $icdf_\mu(r)$ , and returns the next state, the sample, and a weight of 1. The transition function of `factor`( $e$ ) evaluates its arguments  $e$  into a real value  $v$ , and returns the next state, an empty value  $()$ , and the weight  $v$ . To simplify the semantics, the type system of Section 3.2 ensures that the arguments of the probabilistic operators are always deterministic expressions. These arguments can thus be interpreted with the deterministic semantics of Section 4.1. The transition function of a function call  $f(e)$  is similar to the deterministic case with an extra inputs for the random elements<sup>5</sup> and an extra output for the weight.

<sup>5</sup>We note  $[r_1 : r_2]$  the concatenation of two arrays.

$$\begin{aligned}
\langle e : V \rangle^{\text{init}} &: S \\
\langle e : V \rangle^{\text{step}} &: \Gamma \rightarrow S \times [0, 1] \text{ array} \rightarrow S \times V \times [0, \infty) \\
\langle e \rangle^{\text{init}} &= \llbracket e \rrbracket^{\text{init}} \quad \text{if } e \text{ is deterministic} \\
\langle e \rangle_Y^{\text{step}}(m, []) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', v, 1 \\
\langle \text{sample}(e) \rangle^{\text{init}} &= \llbracket e \rrbracket^{\text{init}} \\
\langle \text{sample}(e) \rangle_Y^{\text{step}}(m, [r]) &= \text{let } m', \mu = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', \text{icdf}_\mu(r), 1 \\
\langle \text{factor}(e) \rangle^{\text{init}} &= \llbracket e \rrbracket^{\text{init}} \\
\langle \text{factor}(e) \rangle_Y^{\text{step}}(m, []) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', (), v \\
\langle f(e) \rangle^{\text{init}} &= \langle e \rangle^{\text{init}}, \langle e_f \rangle^{\text{init}} \quad \text{where } \text{proba } f \ x = e_f \\
\langle f(e) \rangle_Y^{\text{step}}((m_e, m_f), [r_e : r_f]) &= \text{let } m'_e, v_e, w_e = \langle e \rangle_Y^{\text{step}}(m_e, r_e) \text{ in} \\
&\quad \text{let } m'_f, v, w_f = \langle e_f \rangle_{[x \leftarrow v_e]}^{\text{step}}(m_f, r_f) \text{ in} \\
&\quad (m'_f, m'_e), v, w_e \times w_f \\
\langle e \text{ where rec } E \rangle^{\text{init}} &= \langle e \rangle^{\text{init}}, \langle E \rangle^{\text{init}} \\
\langle e \text{ where rec } E \rangle_Y^{\text{step}}((m, M), [r_e : r_E]) &= \text{let } F(\rho) = \left( \text{let } M', \rho, w = \langle E \rangle_{Y+\rho}(M, r_E) \text{ in } \rho \right) \text{ in} \\
&\quad \text{let } \rho = \text{fix}(F) \text{ in} \\
&\quad \text{let } M', \rho, W = \langle E \rangle_{Y+\rho}^{\text{step}}(M, r_E) \text{ in} \\
&\quad \text{let } m', v, w = \langle e \rangle_{Y+\rho}^{\text{step}}(m, r_e) \text{ in} \\
&\quad (m', M'), v, w \times W
\end{aligned}$$

Fig. 7. Probabilistic operational semantics of expressions (full version in Figure 17 of the appendix).

The semantics of probabilistic equations is presented in Figure 8. Given an environment, a state, and an array of random elements, the transition function  $\langle E \rangle^{\text{step}} : \Gamma \rightarrow S \times [0, 1] \text{ array} \rightarrow S \times \Gamma \times [0, \infty)$  returns a tuple (next state, environment, weight). The only difference with the deterministic case is that the transition functions now take the random elements for the sub-expressions as arguments and accumulates the weights.

**Example.** Consider the coin node defined in Section 3.1. The number of random variables is 1 because there is only one parameter:  $p$ . For an input expression  $e$  with initial state  $m_e$ , the initial state of the expression  $\text{coin}(e)$  is:

$$\begin{array}{c}
\text{coin.init} \\
\hline
\text{init } p = \text{last } p \\
\text{init } p = \text{sample}(\text{U}(0,1)) \quad | \quad () = \text{observe}(\text{B}(p), 0) \\
m = m_e, ((), ((\text{nil}, (((), ())), (), (((), ()))) \\
\text{e} \quad p \quad \quad \quad 0 \quad 1 \quad | \quad p \quad o \\
\text{last } p
\end{array}$$

### 4.3 Inference

*Notations.* In the following,  $\text{Meas}(U)$  is the set of measures on a measurable space  $U$ ,  $\delta_x$  is the Dirac delta measure ( $\delta_x(U) = 1$  if  $x \in U$ , 0 otherwise), we note  $\int_U \mu(dx)$   $g(x)$  the integral of a measurable function  $g$  w.r.t. the measure  $\mu$ .

Manuscript submitted to ACM

$$\begin{aligned}
(E)^{\text{init}} &: S \\
(E)^{\text{step}} &: \Gamma \rightarrow S \times [0, 1] \text{ array} \rightarrow S \times \Gamma \times [0, \infty) \\
\\
(x = e)^{\text{init}} &= (e)^{\text{init}} \\
(x = e)_Y^{\text{step}}(m, r) &= \text{let } m', v, w = (e)_Y^{\text{step}}(m, r) \text{ in } m', [x \leftarrow v], w \\
\\
(\text{init } x = e)^{\text{init}} &= \text{nil}, (e)^{\text{init}} \\
(\text{init } x = e)_Y^{\text{step}}((\text{nil}, m_0), r) &= \text{let } m', i, w = (e)_Y^{\text{step}}(m_0, r) \text{ in } (\gamma(x), m_0), [x.\text{last} \leftarrow i], w \\
(\text{init } x = e)_Y^{\text{step}}((v, m_0), r) &= (\gamma(x), m_0), [x.\text{last} \leftarrow v], 1 \\
\\
(E_1 \text{ and } E_2)^{\text{init}} &= (E_1)^{\text{init}}, (E_2)^{\text{init}} \\
(E_1 \text{ and } E_2)_Y^{\text{step}}((M_1, M_2), [r_1 : r_2]) &= \text{let } M'_1, \rho_1, w_1 = (E_1)_Y^{\text{step}}(M_1, r_1) \text{ in} \\
&\quad \text{let } M'_2, \rho_2, w_2 = (E_2)_Y^{\text{step}}(M_2, r_2) \text{ in} \\
&\quad (M'_1, M'_2), \rho_1 + \rho_2, w_1 \times w_2
\end{aligned}$$

Fig. 8. Probabilistic operational semantics of equations.

The Lebesgues measure over  $[0, 1]^p$  is denoted  $\lambda^p$  and by abuse of notation, we note  $\int_{[0,1]^p} g(x)dx$  the integral of  $g$  w.r.t to the measure  $\lambda^p$ . If  $f : U \rightarrow V$  is a measurable function and  $\mu \in \text{Meas}(V)$ , then the pushforward of  $\mu$  along  $f$  is the measure  $f_*\mu \in \text{Meas}(U)$  defined as  $f_*\mu = \mu \circ f^{-1}$ . With these notations, the change of variable formula is:  $\int g(y)f_*\mu(dy) = \int g(f(x))\mu(dx)$ , for any measurable function  $g : V \rightarrow W$ .

Given a probabilistic model  $e$  of type  $V$  with a state of type  $S$ , the deterministic expression  $\text{infer}(e)$  computes a stream of distributions. At each time step, given a measure over possible states, the transition function returns a measure over possible next states  $\sigma$ , and a measure over return values  $\mu$ .

$$\begin{aligned}
\llbracket \text{infer}(e) \rrbracket^{\text{init}} &: \text{Meas}(S) \\
\llbracket \text{infer}(e) \rrbracket^{\text{step}} &: \Gamma \rightarrow \text{Meas}(S) \rightarrow \text{Meas}(S) \times \text{Meas}(V)
\end{aligned}$$

The initial state of  $\text{infer}(e)$  is a Dirac delta measure over the initial state of the model  $e$ . The transition function first computes a function  $\psi$  mapping a state to the un-normalized measure which associates each pair (next state, value) to its weight. The  $\text{infer}$  operator then 1) integrates this function along all possible values of the state, 2) normalizes it ( $\top$  denotes the entire space), and 3) splits the result into a distribution of next states  $\sigma$  and a distribution of values  $\mu$ .

$$\begin{aligned}
\llbracket \text{infer}(e) \rrbracket^{\text{init}} &= \delta_{(e)^{\text{init}}} \\
\llbracket \text{infer}(e) \rrbracket_Y^{\text{step}}(\sigma) &= \text{let } \psi(m) = \int_{[0,1]^{|e|}} \text{let } m', v, w = (e)_Y^{\text{step}}(m, r) \text{ in } w \times \delta_{(m', v)} dr \text{ in} \\
&\quad \text{let } v = \int \sigma(dm) \psi(m) \text{ in} \\
&\quad \text{let } \bar{v} = v/v(\top) \text{ in} \\
&\quad \pi_{1*}(\bar{v}), \pi_{2*}(\bar{v})
\end{aligned} \tag{3}$$

If the model is ill-defined, the normalization constant can be 0 or  $\infty$ , which triggers an exception and stops the execution. It is the programmer's responsibility to avoid such error cases when defining the model.

The stream of distributions corresponding to  $\text{infer}(e)$  is obtained by iteratively integrating the transition function along the distribution of states.

$$((e)^{\text{init}} = \delta_{m_0}) \xrightarrow{\int \sigma_0(dm)\psi_0(m)} \sigma_1 \xrightarrow{\int \sigma_1(dm)\psi_1(m)} \sigma_2 \xrightarrow{\int \sigma_2(dm)\psi_2(m)} \sigma_3 \rightarrow \dots$$

$$\mu_1 \qquad \qquad \qquad \mu_2 \qquad \qquad \qquad \mu_3$$

#### 4.4 Program equivalence

Two expressions are equivalent if they compute the same stream of output values. The semantics is defined with an initial state and a transition function. To prove equivalence of two state machines one can exhibit a bisimulation [50] that relates the states and ensure the equality of output values.

*Definition 4.2.* Deterministic expressions are equivalent if there is a bisimulation, i.e., a relation on states  $\mathcal{P} \subseteq S \times S$  such that:

- $(\llbracket e_1 \rrbracket^{\text{init}}, \llbracket e_2 \rrbracket^{\text{init}}) \in \mathcal{P}$
- $\forall \gamma, \text{ if } (m_1, m_2) \in \mathcal{P}, \text{ then } (m'_1, m'_2) \in \mathcal{P} \text{ and } v_1 = v_2 \text{ where } m'_1, v_1 = \llbracket e_1 \rrbracket^{\text{step}}_\gamma(m_1) \text{ and } m'_2, v_2 = \llbracket e_2 \rrbracket^{\text{step}}_\gamma(m_2).$

A bisimulation on states can be lifted to a bisimulation on measures of states to characterize equivalence between probabilistic processes [9, 47].

*Definition 4.3.* Let  $\mathcal{P} \subseteq S \times S$  be a relation on states. The lifted relation  $\mathcal{P}^+$  is the relation on measures of states defined as  $(\sigma_1, \sigma_2) \in \mathcal{P}^+$ , when for any measurable subset  $\mathcal{P}' \subseteq \mathcal{P}$ ,  $\sigma_1(\pi_1 \mathcal{P}') = \sigma_2(\pi_2 \mathcal{P}')$ .

Two probabilistic expressions are equivalent if starting from bisimilar states, integrating at each steps the pairs (value, weight) computed by the probabilistic operational semantics, the output measures are equal and the measures on states are bisimilar.

*Definition 4.4.* Two probabilistic expressions  $e_1$  and  $e_2$  are equivalent if there is a bisimulation  $\mathcal{P} \subseteq S \times S$  such that:

- $((e_1)^{\text{init}}, (e_2)^{\text{init}}) \in \mathcal{P}$
- $\forall \gamma, \text{ if } (m_1, m_2) \in \mathcal{P}, \text{ then } (\sigma'_1, \sigma'_2) \in \mathcal{P}^+ \text{ and } \mu_1 = \mu_2, \text{ where}$

$$\begin{aligned} \sigma'_1, \mu_1 &= \text{let } \psi_1(m_1) = \int_{[0,1]^{|e_1|}} \text{let } m'_1, v_1, w_1 = \llbracket e_1 \rrbracket^{\text{step}}_\gamma(m_1, r_1) \text{ in } w_1 \times \delta_{(m'_1, v_1)} dr_1 \text{ in} \\ &\qquad \pi_{1*}(\psi_1), \pi_{2*}(\psi_1) \\ \sigma'_2, \mu_2 &= \text{let } \psi_2(m_2) = \int_{[0,1]^{|e_2|}} \text{let } m'_2, v_2, w_2 = \llbracket e_2 \rrbracket^{\text{step}}_\gamma(m_2, r_2) \text{ in } w_2 \times \delta_{(m'_2, v_2)} dr_2 \text{ in} \\ &\qquad \pi_{1*}(\psi_2), \pi_{2*}(\psi_2) \end{aligned}$$

Using this definition, we can show that two probabilistic expressions are equivalent if they describe the same stream of measures, i.e., if at each step, integrating the transition function along the distributions of state in both expressions yields the same measures.

**PROPOSITION 4.5.** *Let  $e_1$  and  $e_2$  be two equivalent expressions with probabilistic bisimulation  $\mathcal{P}$ . If  $(\sigma_1, \sigma_2) \in \mathcal{P}^+$ , then with the notations of Definition 4.4,  $(\sigma_1'', \sigma_2'') \in \mathcal{P}^+$  and  $\mu_1' = \mu_2'$  where*

$$\begin{aligned}\sigma_1'', \mu_1' &= \int \sigma_1(dm_1) \sigma_1'(m_1), \int \sigma_1(dm_1) \mu_1(m_1) \\ \sigma_2'', \mu_2' &= \int \sigma_2(dm_2) \sigma_2'(m_2), \int \sigma_2(dm_2) \mu_2(m_2)\end{aligned}$$

**PROOF.** Let  $\sigma$  and  $\sigma'$  be the measures on  $\mathcal{P}$  such that for any measurable  $\mathcal{P}' \subseteq \mathcal{P}$ ,  $\sigma(\mathcal{P}') = \sigma_1(\pi_1 \mathcal{P}') = \sigma_2(\pi_2 \mathcal{P}')$  and  $\sigma'(m)(\mathcal{P}') = \sigma_1'(m_1)(\pi_1 \mathcal{P}') = \sigma_2'(m_2)(\pi_2 \mathcal{P}')$ , where  $m = (m_1, m_2)$ . By the change of variable formula, for  $i \in \{1, 2\}$ ,

$$\sigma_i''(\pi_i \mathcal{P}') = \int \pi_{i*}(\sigma)(dm_i) \sigma_i'(m_i)(\pi_i \mathcal{P}') = \int \sigma(dm) \sigma_i'(\pi_i(m))(\pi_i \mathcal{P}') = \int \sigma(dm) \sigma'(m)(\mathcal{P}')$$

Similarly, let  $\mu(m) = \mu_1(m_1) = \mu_2(m_2)$ . Then, for  $i \in \{1, 2\}$ ,

$$\mu_i' = \int \pi_{i*}(\sigma)(dm_i) \mu_i(m_i) = \int \sigma(dm) \mu_i(\pi_i(m)) = \int \sigma(dm) \mu(m)$$

□

For a given context, each triplet (state, value, weight) is a function of the previous state and of the random elements. If we can relate the states of  $e_1$  with the states of  $e_2$  and the random elements of  $e_1$  with the random elements of  $e_2$ , then program equivalence can be reduced to the comparison of the stream of pairs (value, weight). To relate two distributions, we use a *coupling*: a joint measure over the product space such that each marginal corresponds to one of the measure. We adapt the presentation of coupling for discrete probabilistic systems of [3] and follow the co-algebraic formulation of [29] to get a characterization of probabilistic equivalence by coupling for our operational semantics.

**PROPOSITION 4.6 (OPERATIONAL COUPLING EQUIVALENCE).** *Probabilistic expressions  $e_1$  and  $e_2$  are equivalent if there is a pair of measurable relations  $\mathcal{P} \subseteq S \times S$  and  $\mathcal{R} \subseteq [0, 1]^{|e_1|} \times [0, 1]^{|e_2|}$  such that:*

- $\pi_1 \mathcal{R} = [0, 1]^{|e_1|}$  and  $\pi_2 \mathcal{R} = [0, 1]^{|e_2|}$
- There is a coupling  $\Lambda$  of the Lebesgue measures over  $[0, 1]^{|e_1|}$ , and  $[0, 1]^{|e_2|}$ :  
 $\Lambda(\mathcal{R}) = 1$ ,  $\pi_{1*} \Lambda = \lambda^{|e_1|}$ , and  $\pi_{2*} \Lambda = \lambda^{|e_2|}$
- $((e_1))^{init}, ((e_2))^{init} \in \mathcal{P}$
- $\forall \gamma, \text{if } (m_1, m_2) \in \mathcal{P} \text{ and } (r_1, r_2) \in \mathcal{R}, \text{ then } (m'_1, m'_2) \in \mathcal{P}, v_1 = v_2 \text{ and } w_1 = w_2, \text{ where}$

$$\begin{aligned}m'_1, v_1, w_1 &= (e_1)_\gamma^{step}(m_1, r_1) \\ m'_2, v_2, w_2 &= (e_2)_\gamma^{step}(m_2, r_2)\end{aligned}$$

**PROOF.** The proof is adapted from [29, Lemma 5.5]. Assume  $(m_1, m_2) \in \mathcal{P}$ . Let us define a measure on  $(S \times V) \times (S \times V)$  as  $C = \int_{\mathcal{R}} w \times \delta_{(m'_1, v), (m'_2, v)} \Lambda(dr_1, dr_2)$ , where  $w = w_1 = w_2$  and  $v = v_1 = v_2$ . It is a coupling of  $\psi_1(m_1)$  and  $\psi_2(m_2)$ . Indeed, for any  $i \in \{1, 2\}$ ,

$$\pi_{i*} C = \int_{\mathcal{R}} w \times \delta_{(m'_i, v)} \Lambda(dr_1, dr_2) = \int_{[0, 1]^{|e_i|}} w_i \times \delta_{(m'_i, v_i)} \lambda^{|e_i|}(dr_i) = \psi_i(m_i)$$

We prove that  $\mu_1 = \mu_2$  by applying the change of variable formula along  $\pi_i : [0, 1]^{|e_1|} \times [0, 1]^{|e_2|} \rightarrow [0, 1]^{|e_i|}$ ,

$$\mu_i = \int_{[0, 1]^{|e_i|}} w_i \times \delta_{v_i} \lambda^{|e_i|}(dr_i) = \int_{\pi_i \mathcal{R}} w \times \delta_v \pi_{i*} \Lambda(dr_i) = \int_{\mathcal{R}} w \times \delta_v \Lambda(dr_1, dr_2).$$

Let us prove that  $(\sigma'_1, \sigma'_2) \in \mathcal{P}^+$ . For any  $\mathcal{P}' \subseteq \mathcal{P}$ , let us set  $Q = \{(m'_1, v), (m'_2, v) \mid (m'_1, m'_2) \in \mathcal{P}'\}$ , we prove that  $\sigma'_1(\pi_1 \mathcal{P}') = \sigma'_2(\pi_2 \mathcal{P}')$  by applying the definition of the pushforward measure along  $\pi_i \pi_1$ , noticing that  $\pi_i \mathcal{P}' = \pi_1 \pi_i Q$ ,

$$\sigma'_i(\pi_i \mathcal{P}') = \pi_{1*} \psi_i(m_i)(\pi_i \mathcal{P}') = \pi_{1*} \pi_{i*} C(\pi_1 \pi_i Q) = \pi_{i*} C(\pi_i Q) = C(Q).$$

□

Finding a relation and a coupling is in general difficult. A useful simple case is when the two programs involve the same random variables in different orders, e.g., a program and its compiled version after a source-to-source transformation. Then, the measurable relation on random elements is a permutation, and two expressions are equivalent if they compute the same stream of pairs (value, weight).

**Example.** If  $x$  and  $y$  are not free variables in  $e_1$  and  $e_2$ , and each expression has one random variable ( $|e_1| = |e_2| = 1$ ):

$$\text{sample}(e_1) + \text{sample}(e_2) \sim x + y \text{ where } \text{rec } x = \text{sample}(e_2) \text{ and } y = \text{sample}(e_1)$$

We define the following:

$$\begin{aligned} m'_1, v_1, 1 &= (e_1)_Y^{\text{step}}(m_1, r_1) & v_y &= \text{icdf}_{v_1}(r_1) \\ m'_2, v_2, 1 &= (e_2)_Y^{\text{step}}(m_2, r_2) & v_x &= \text{icdf}_{v_2}(r_2) \end{aligned}$$

The left hand side term  $e_\ell = \text{sample}(e_1) + \text{sample}(e_2)$  is interpreted by the state machine:

$$\begin{aligned} (e_\ell)^{\text{init}} &= ((e_1)^{\text{init}}, (e_2)^{\text{init}}) \\ (e_\ell)_Y^{\text{step}}((m_1, m_2), [r_1 : r_2]) &= (m'_1, m'_2), v_y + v_x, 1 \end{aligned}$$

The right hand side term  $e_r = x + y$  where  $\text{rec } x = \text{sample}(e_2)$  and  $y = \text{sample}(e_1)$  is interpreted by the state machine:

$$\begin{aligned} (e_r)^{\text{init}} &= ((((), ((), ())), ((e_2)^{\text{init}}, (e_1)^{\text{init}}))) \\ (e_r)_Y^{\text{step}}((((), ((), ())), (m_2, m_1)), [r_2 : r_1]) &= ((((), ((), ())), (m'_2, m'_1)), v_x + v_y, 1) \end{aligned}$$

With  $\mathcal{P} = \{((m_1, m_2), (((), ((), ())), (m_2, m_1))) \mid (m_1, m_2) \in S \times S\}$

and  $\mathcal{R} = \{([r_1 : r_2], [r_2 : r_1]) \mid r_1 \in [0, 1], r_2 \in [0, 1]\}$  we have:

- $\pi_1 \mathcal{R} = [0, 1]^2$  and  $\pi_2 \mathcal{R} = [0, 1]^2$
- the Lebesgues measure on  $[0, 1]^4$ ,  $\Lambda = \lambda^4$  satisfies  $\pi_{1*} \Lambda = \lambda^2$  and  $\pi_{2*} \Lambda = \lambda^2$
- $\mathcal{P}$  relates initial states  $((e_1)^{\text{init}}, (e_2)^{\text{init}})$  and  $(((), ((), ())), ((e_2)^{\text{init}}, (e_1)^{\text{init}}))$
- if  $\mathcal{P}$  relates the current states  $(m_1, m_2)$  and  $(((), ((), ())), (m_2, m_1))$  and the random elements are permuted, then  $\mathcal{P}$  relates the next states  $(m'_1, m'_2)$  and  $(((), ((), ())), (m'_2, m'_1))$ , and the two state machines yield the same pairs of values  $v_y + v_x = v_x + v_y$  and weights  $1 = 1$ .

We can thus apply Proposition 4.6 to conclude that the two probabilistic expressions are equivalent.

## 5 DENOTATIONAL SEMANTICS

An alternative to the operational semantics where expressions are interpreted as state machines is to define a denotational semantics where expressions directly return streams of values [21]. This formalism has been successfully used in the Vélus project to prove an end-to-end dataflow synchronous compiler within the Rocq proof assistant [11–13, 15].

In this section, we first present a denotational semantics for the deterministic expressions of our language. We then define a probabilistic denotational semantics for probabilistic expressions and prove that this semantics is equivalent to the probabilistic operational semantics of Section 4.2, i.e., the `infer` operator yields the same stream of distributions.

Manuscript submitted to ACM

$$\begin{array}{c}
 \frac{}{H \vdash c \downarrow c} \quad \frac{}{H \vdash x \downarrow H(x)} \quad \frac{H \vdash e_1 \downarrow s_1 \quad H \vdash e_2 \downarrow s_2}{H \vdash (e_1, e_2) \downarrow (s_1, s_2)} \quad \frac{H \vdash e \downarrow s}{H \vdash op(e) \downarrow op(s)} \quad \frac{H(x.\text{last}) = s}{H \vdash \text{last } x \downarrow s} \\
 \\ 
 \frac{H \vdash e \downarrow s_e \quad \text{node } f \ x = e_f}{H \vdash f(e) \downarrow s} \quad \frac{[x \leftarrow s_e] \vdash e_f \downarrow s}{H \vdash e \downarrow s} \quad \frac{H + H_E \vdash E \quad H + H_E \vdash e \downarrow s}{H \vdash e \text{ where rec } E \downarrow s} \\
 \\ 
 \frac{}{H \vdash e \downarrow H(x)} \quad \frac{H \vdash e \downarrow i \cdot s \quad H(x.\text{last}) = i \cdot H(x)}{H \vdash \text{init } x = e} \quad \frac{H \vdash E_1 \quad H \vdash E_2}{H \vdash E_1 \text{ and } E_2}
 \end{array}$$

Fig. 9. Denotational semantics for deterministic expressions and equations (full version in Figure 20 of the appendix).

### 5.1 Deterministic denotational semantics

*Notations.* In the following,  $V^\omega$  is the type of infinite streams of values of type  $V$ . The infix operator  $(\cdot) : V \rightarrow V^\omega \rightarrow V^\omega$  is the stream constructor (e.g.,  $1 \cdot 2 \cdot 3 \cdot \dots$ ). Constants are lifted to constant streams (e.g.,  $1 = 1 \cdot 1 \cdot 1 \cdot \dots$ ) and when the context is clear we write  $f(s) = f(s_0) \cdot f(s_1) \cdot \dots$  for *map*  $f$   $s$ , and  $(s, t) = (s_1, t_1) \cdot (s_2, t_2) \cdot \dots$  to cast a pair of streams into a stream of pairs.

In the denotational semantics, deterministic expressions compute streams of values. In a context  $H$  which maps variable names to stream of values, the semantics of a deterministic expression  $e$  returns a stream  $s: H \vdash e \downarrow s$ . The semantics of a set of equations  $E$  checks that the context  $H$  is compatible with all the equations:  $H \vdash E$ . The semantics of a set of equations thus defines a *relation* between the streams stored in the context. Compared to the operational semantics, the denotational semantics is not executable since the context must be guessed a priori and validated against the equations of the program.

Figure 9 presents the denotational semantics for deterministic expressions and equations. A constant is interpreted as a constant stream, and a variable returns the corresponding stream in the context. The semantics of a pair evaluates each component independently and packs the results into a stream of pairs. The application of an operator evaluates its argument into a stream of values and maps the operator on the result. `last`  $x$  fetches a special variable  $x.\text{last}$  in the context. A function call first evaluates its argument, and then evaluates the body of the function in a context where the parameter is bound to the argument value.

To interpret the expression  $e \text{ where rec } E$ , equations  $E$  are evaluated in a new context  $H_E$  that is also used to evaluate the main expression  $e$ . The semantics of a simple equation checks that a variable is associated to the stream computed by its defining expression. The initialization operator `init`  $x = e$  prepends an initial value  $i$  to the stream associated to  $x$  and checks that the special variable  $x.\text{last}$  is bound to this delayed version of  $x$ . In the denotational semantics, contexts are un-ordered maps and scheduling equations does not change the semantics.

*Causality.* For some programs, the denotational semantics is not uniquely defined. For instance, the semantics of the expression  $x \text{ where rec } x = x$  can be any stream of value (the equation is verified for any context containing  $x$ ). Such programs are rejected by the causality analysis of the compiler (Definition 4.1).

### 5.2 Probabilistic denotational semantics

The key idea of the probabilistic denotational semantics is the same as in Section 4.2: instead of manipulating streams of measures, probabilistic expressions compute streams of pairs (value, score) using external streams of random elements, and integration is deferred to the `infer` operator.

Manuscript submitted to ACM

$$\begin{array}{c}
\frac{H \vdash e \downarrow s}{H, [] \vdash e \Downarrow (s, 1)} \quad \frac{H \vdash e \downarrow s_\mu}{H, [R] \vdash \text{sample}(e) \Downarrow (icdf_{s_\mu}(R), 1)} \quad \frac{H \vdash e \downarrow w}{H, [] \vdash \text{factor}(e) \Downarrow (((), w))} \\
\\
\frac{H, R_e \vdash e \downarrow (s_e, w_e) \quad \text{proba } f \ x = e_f \quad [x \leftarrow s_e], R_f \vdash e_f \Downarrow (s, w)}{H, [R_e : R_f] \vdash f(e) \Downarrow (s, w \times w_e)} \\
\\
\frac{H + H_E, R_E \vdash E : w_E \quad H + H_E, R_E \vdash e \Downarrow (s, w)}{H, [R_E : R_E] \vdash e \text{ where rec } E \Downarrow (s, w \times w_E)} \quad \frac{H, R \vdash e \Downarrow (H(x), w)}{H, R \vdash x = e : w} \\
\\
\frac{H, R \vdash e \Downarrow (i \cdot s, w_i \cdot w) \quad H(x.\text{last}) = i \cdot H(x)}{H, R \vdash \text{init } x = e : w_i \cdot 1} \quad \frac{H, R_1 \vdash E_1 : w_1 \quad H, R_2 \vdash E_2 : w_2}{H, [R_1 : R_2] \vdash E_1 \text{ and } E_2 : w_1 \times w_2} \\
\\
\frac{[H, R \vdash e \Downarrow (s, w) \quad \bar{w} = \prod w]_{R \in ([0,1]^\omega)^{|e|}}}{H \vdash \text{infer}(e) \downarrow \text{integ}_{|e|} \bar{w} s}
\end{array}$$

Fig. 10. Denotational semantics for probabilistic expressions and equations (full version in Figure 21 of the appendix).

Figure 10 presents the semantics of probabilistic expressions and equations. In a context  $H$  which maps variable names to stream of values, the semantics of a probabilistic expression  $e$  takes an array of random streams  $R$  and returns a stream of pairs (value, weight):  $H, R \vdash e \Downarrow (s, w)$ . The semantics of a set of equations  $E$  takes an array containing the random streams of all sub-expressions, checks that the context  $H$  is compatible with all the equations, and returns the total weight  $w_E$  of all sub-expressions:  $H, R \vdash E : w_E$ .

The semantics of deterministic expressions (e.g., constant or variable) returns the expected stream of values associated to a constant weight of 1. The semantics of `sample` takes an array containing one random stream  $R$ , evaluates its argument into a stream of distributions  $s_\mu$ , and uses  $R$  to compute a stream of samples associated to the weight 1:  $(icdf_{s_\mu_0}(R_0), 1) \cdot (icdf_{s_\mu_1}(R_1), 1) \cdot \dots$ . The semantics of `factor` evaluates its argument into a stream of values  $w$  which is used as the weight associated to a stream of empty values:  $(((), w_0) \cdot (((), w_1) \cdot \dots$ . The semantics of a function call is similar to the deterministic case, but the random streams are split between the argument and the function body, and the total weight captures the weight of the argument and the weight of the function body. Similarly, for a set of local definitions the random streams are split between sub-expressions and the weight is the total weight of all sub-expressions.

### 5.3 Inference

As in the operational semantics, the `infer` operator is defined by integrating at each step the semantics of the model over all possible values of the streams of random elements. The semantics of a probabilistic model returns a pair of stream functions (value, weight) which both depend on the random streams. Given the random streams, at each time step, the semantics of `infer` first computes the total weight of the prefix to capture all the conditioning since the beginning of the execution:  $\bar{w} = \prod w = w_0 \cdot (w_0 \times w_1) \cdot (w_0 \times w_1 \times w_2) \cdot \dots$ . Then the function `integ` 1) turns the current value  $v_n$  and the total weight  $\bar{w}_n$  into an un-normalized measure by integrating over all possible values of the random streams, and 2) normalizes the result to obtain a stream of distributions of values. If  $\lambda_\omega^{|e|}$  is the uniform measure over the cube of random streams  $([0, 1]^\omega)^{|e|}$ , then:

$$\text{integ}_{|e|} (\bar{w}_n \cdot \bar{w}s) (v \cdot vs) = \left( \text{let } \mu = \int \bar{w}_n \times \delta_v \lambda_\omega^{|e|} (dR) \text{ in } \mu / \mu(\top) \right) \cdot (\text{integ } \bar{w}s vs) \quad (4)$$

*Cube of random streams.* The uniform measure over the cube of random streams is defined as follows. Let  $[0, 1]^\omega$  be the countable product of the measurable spaces on the interval  $[0, 1]$  endowed with the Lebesgue  $\sigma$ -algebra, i.e., the coarsest  $\sigma$ -algebra such that projections are measurable. We define  $\lambda_\omega$  as the uniform distribution on the *continuous cube* defined by a Kolmogorov extension such that for any  $k \in \mathbb{N}$ , the pushforward measure of  $\lambda_\omega$  along the projection  $\pi_{\leq k} : [0, 1]^\omega \rightarrow [0, 1]^k$  on the first  $k$  coordinates is the Lebesgue measure on  $[0, 1]^k$ :  $\lambda_{\leq k} = \pi_{\leq k*}(\lambda_\omega)$ . For any measurable function  $g : [0, 1]^k \rightarrow V$  we have the following change of variable formula:

$$\int g(\pi_{\leq k}(R)) \lambda_\omega(dR) = \int g(R_{\leq k}) \lambda_{\leq k}(dR_{\leq k})$$

Integrating a function which only depends on the  $k$  first coordinates of  $R$  can thus be reduced to integrating over these coordinates. We can then define the uniform measure on the cube of random streams  $\lambda_\omega^{|e|}$  as the  $|e|$ -ary product measure of  $\lambda_\omega$ , and lift the change of variable formula.

#### 5.4 Program equivalence

Compared to the operational semantics where proving the equivalence between two state machines requires a bisimulation, in the denotational semantics, to prove the equivalence between two programs one need only to check that they define the same streams.

*Definition 5.1.* Deterministic expressions  $e_1$  and  $e_2$  are equivalent if for any context  $H$ ,  $H \vdash e_1 \downarrow s$  and  $H \vdash e_2 \downarrow s$ .

Equivalence between probabilistic expressions is more complex. Two probabilistic expressions are equivalent if they describe the same stream of measures obtained by integrating at each step the streams of pairs (value, weight) computed by the denotational semantics.

*Definition 5.2.* Probabilistic expressions  $e_1$  and  $e_2$  are equivalent if for all contexts  $H$ , and  $\forall k > 0$ ,

$$\int \overline{w_1}_k \times \delta_{s_1 k} \lambda_\omega^{|e_1|}(dR_1) = \int \overline{w_2}_k \times \delta_{s_2 k} \lambda_\omega^{|e_2|}(dR_2)$$

where  $H, R_1 \vdash e_1 \Downarrow (s_1, w_1)$  and  $H, R_2 \vdash e_2 \Downarrow (s_2, w_2)$ .

In the denotational semantics, for a given context, each pair (value, weight) is a function of the random streams. Since, random streams are uniformly distributed, if we can relate the random streams of  $e_1$  with the random streams of  $e_2$  with a coupling on uniform distributions, program equivalence can be reduced to the comparison of the streams of pairs (value, weight) computed by each expression.

**PROPOSITION 5.3 (DENOTATIONAL COUPLING EQUIVALENCE).** *Probabilistic expressions  $e_1$  and  $e_2$  are equivalent if there is  $\mathcal{R} \subseteq ([0, 1]^\omega)^{|e_1|} \times ([0, 1]^\omega)^{|e_2|}$  measurable such that,*

- $\pi_1 \mathcal{R} = [0, 1]^{|e_1|}$  and  $\pi_2 \mathcal{R} = [0, 1]^{|e_2|}$
- there is a coupling  $\Lambda_\omega$  on  $([0, 1]^\omega)^{|e_1|} \times ([0, 1]^\omega)^{|e_2|}$ , such that  $\Lambda_\omega(\mathcal{R}) = 1$ ,  $\pi_{1*} \Lambda_\omega = \lambda_\omega^{|e_1|}$  and  $\pi_{2*} \Lambda_\omega = \lambda_\omega^{|e_2|}$
- for all contexts  $H$  and arrays of random streams  $(R_1, R_2) \in \mathcal{R}$ ,  $s_1 = s_2$  and  $w_1 = w_2$ , where  $H, R_1 \vdash e_1 \Downarrow (s_1, w_1)$  and  $H, R_2 \vdash e_2 \Downarrow (s_2, w_2)$ .

**PROOF.** With  $w_k = w_{1k} = w_{2k}$  and  $s_k = s_{1k} = s_{2k}$ , for all  $k > 0$ ,

$$\int \overline{w_1}_k \times \delta_{s_1 k} \lambda_\omega^{|e_1|}(dR_1) = \int \overline{w}_k \times \delta_{s_k} \Lambda_\omega(dR) = \int \overline{w_2}_k \times \delta_{s_2 k} \lambda_\omega^{|e_2|}(dR_2)$$

□

Manuscript submitted to ACM

Finding such a coupling is in general difficult. As in Section 4.4, a useful simple case is when the two programs involve the same random variables in different orders. Then, the coupling permutes the random streams.

The denotational semantics of an expression is described by a derivation tree where each relation is a consequence of smaller relations on all the sub-expressions, up to atomic expressions. Two expressions compute the same streams if from the derivation tree of the first, one can build a derivation tree for the second and vice-versa.

**Example.** If  $x$  and  $y$  are not free variables in  $e_1$  and  $e_2$ :

$$\text{sample}(e_1) + \text{sample}(e_2) \sim x + y \text{ where } \text{rec } x = \text{sample}(e_2) \text{ and } y = \text{sample}(e_1)$$

Let  $R_i$  be the array of random streams associated to the expressions  $\text{sample}(e_i)$ . For all contexts  $H$ , if  $H \vdash e_i \downarrow \mu_i$ , then we define  $s_i = \text{icdf}_{\mu_i}(R_i)$ . Then, the derivation tree for the lhs expression is:

$$\frac{H, R_1 \vdash \text{sample}(e_1) \Downarrow (s_1, 1) \quad H, R_2 \vdash \text{sample}(e_2) \Downarrow (s_2, 1)}{H, [R_1 : R_2] \vdash \text{sample}(e_1) + \text{sample}(e_2) \Downarrow (s_1 + s_2, 1)}$$

With  $H_E = [x \leftarrow s_2, y \leftarrow s_1]$ , the derivation tree for the rhs expression is:

$$\frac{\begin{array}{c} H + H_E, R_2 \vdash \text{sample}(e_2) \Downarrow (s_2, 1) \quad H + H_E, R_1 \vdash \text{sample}(e_1) \Downarrow (s_1, 1) \\ \hline H + H_E, R_2 \vdash x = \text{sample}(e_2) : 1 \quad H + H_E, R_1 \vdash y = \text{sample}(e_1) : 1 \end{array}}{\begin{array}{c} H + H_E, [] \vdash x + y \Downarrow (s_2 + s_1, 1) \quad H + H_E, [R_2 : R_1] \vdash x = \text{sample}(e_2) \text{ and } y = \text{sample}(e_1) : 1 \\ \hline H, [R_2 : R_1] \vdash x + y \text{ where } \text{rec } x = \text{sample}(e_2) \text{ and } y = \text{sample}(e_1) \Downarrow (s_2 + s_1, 1) \end{array}}$$

Since both programs compute the same stream of pairs (value, weight) and the relation permuting random streams  $\mathcal{R} = \{([R_1 : R_2], [R_2 : R_1]) \mid R_1 \in [0, 1]^\omega, R_2 \in [0, 1]^\omega\}$  is the support of the coupling  $\Lambda_\omega = \lambda_\omega^4$  of  $\lambda_\omega^2$  and  $\lambda_\omega^2$ , the two programs are equivalent.

## 6 SEMANTICS COMPARISON

In this section we show that the operational and the denotational semantics coincide for causal programs, i.e., they both define the same streams. We then prove that these new semantics coincide with the original kernel semantics for the restricted set of schedulable programs that can be interpreted by the original semantics.

### 6.1 Adequacy between the operational and denotational semantics

For a probabilistic expression  $e$ , we first relate the operational semantics of Section 4.2 and the denotational semantics of Section 5.2. If  $H$  is a context mapping variables names to streams of values,  $H_k$  is the context where streams are projected on their  $k$ -th coordinate and  $H_{\leq k}$  is the context where streams are truncated at  $k$ . We define similarly  $R_{\leq k}$ , and  $R_k$  for an array of random streams  $R$ . The following proposition states that if a program is causal (Definition 4.1), the probabilistic operational semantics and the probabilistic denotational semantics coincide.

**PROPOSITION 6.1.** *For a correct model  $e$  (well-typed, causal, and without runtime error), the probabilistic denotational semantics is well defined and matches the trace of the probabilistic operational semantics, i.e., for all contexts  $H, R, H, R \vdash e \Downarrow (s, w)$  is well defined and  $\forall k \geq 0$ ,  $(m_{k+1}, s_{k+1}, w_{k+1}) = (e)_{H_{k+1}}^{\text{step}}(m_k, R_{k+1})$  with  $m_0 = (e)^{\text{init}}$ . By construction, for all  $k > 0$  the value  $s_k$  and the weight  $w_k$  only depend on  $H_{\leq k}$  and  $R_{\leq k}$ .*

Manuscript submitted to ACM

**PROOF SKETCH.** Probabilistic constructs aside, the kernel language of Figure 2 is a subset of the Vélus language defined in [11]. We compile the model to this language. Probabilistic nodes become deterministic nodes with additional inputs (the random streams) and one additional output (the score). The `sample` operator is also compiled to a function call with one additional input for the random stream, and the `factor` operator simply updates the score. Following the semantics of Figure 10, the distribution of random streams in sub-expressions is performed by the compilation function. We can then apply the correctness theorem of [14, Theorem 5] : *if a program is correct and without runtime error, the denotational semantics is well defined (there is a valid environment for all equations sets) and matches the trace of the compiled state machine.*

The execution of the compiled state machine is deterministic and corresponds to the operational semantics of the normalized scheduled program which does not require any fixpoint computation. Since, the normalization and scheduling passes preserve the operational semantics [22], the execution of the state machine also corresponds to the operational semantics of the original unscheduled program.  $\square$

We can now state the adequacy theorem, i.e., the `infer` operator yields the same stream of distributions in the operational semantics and in the denotational semantics.

**THEOREM 6.2 (ADEQUACY).** *For a correct model  $e$  (well-typed, causal, and without runtime errors), for all context  $H$  such that  $H \vdash \text{infer}(e) \downarrow \mu$ , the operational execution trace yields the same stream of distributions, that is,  $\sigma_0 = \llbracket \text{infer}(e) \rrbracket^{init}$  and  $\forall k \geq 0, \sigma_{k+1}, \mu_{k+1} = \llbracket \text{infer}(e) \rrbracket_{H_{k+1}}^{step}(\sigma_k)$ .*

**PROOF.** We show  $\forall k > 0$ :

$$\begin{aligned} \sigma_k &\propto \int_{([0,1]^k)^{|e|}} \overline{w_k} \times \delta_{m_k} \lambda_{\leq k}^{|e|}(dR_{\leq k}) \\ \mu_k &\propto \int_{([0,1]^k)^{|e|}} \overline{w_k} \times \delta_{s_k} \lambda_{\leq k}^{|e|}(dR_{\leq k}) \end{aligned}$$

By Proposition 6.1, for all  $k > 0$ , the values  $s_k$ , and the weights  $w_k$  computed by the operational semantics corresponds to the stream  $(s, w)$  of the denotational semantics. From the definition of  $\psi$  in Equation (3) in Section 4.3 and Fubini's theorem we have:<sup>6</sup>

$$\begin{aligned} v_{k+1} &= \int \sigma_k(dm) \psi(m) \\ &\propto \int_{([0,1]^k)^{|e|}} \overline{w_k} \times \psi(m_k) \lambda_{\leq k}^{|e|}(dR_{\leq k}) \\ &\propto \int_{([0,1]^k)^{|e|}} \int_{[0,1]^{|e|}} \overline{w_k} \times w_{k+1} \times \delta_{m_{k+1}, v_{k+1}} \lambda^{|e|}(dR_{k+1}) \lambda_{\leq k}^{|e|}(dR_{\leq k}) \\ &\propto \int_{([0,1]^{k+1})^{|e|}} \overline{w_{k+1}} \times \delta_{m_{k+1}, v_{k+1}} \lambda_{\leq k+1}^{|e|}(dR_{\leq k+1}) \end{aligned}$$

The normalization and marginalization by  $\pi_{1*}$  and  $\pi_{2*}$  concludes the inductive case. Then using the change of variable formula on the cube of random streams we get:

$$\mu_k \propto \int \overline{w_k} \times \delta_{s_k} \lambda_{\omega}^{|e|}(dR)$$

which corresponds to Equation (4) in Section 5.3 and concludes the proof.  $\square$

<sup>6</sup>The symbol  $\propto$  means proportional to.

$$\begin{aligned}
\llbracket e : V \rrbracket^{\text{init}} &: S \\
\llbracket e : V \rrbracket^{\text{step}} &: \Gamma \rightarrow S \rightarrow \text{Meas}(S \times V) \\
\llbracket e \rrbracket^{\text{init}} &= \llbracket e \rrbracket^{\text{init}} \\
\llbracket e \rrbracket^{\text{step}}_\gamma(m) &= \text{let } m', v = \llbracket e \rrbracket^{\text{step}}_\gamma(m) \text{ in } \delta_{m', v} \text{ if } e \text{ is deterministic} \\
\llbracket \text{sample}(e) \rrbracket^{\text{init}} &= \llbracket e \rrbracket^{\text{init}} \\
\llbracket \text{sample}(e) \rrbracket^{\text{step}}_\gamma(m) &= \text{let } m', \mu = \llbracket e \rrbracket^{\text{step}}_\gamma(m) \text{ in } \int \mu(dv) \delta_{m', v} \\
\llbracket \text{factor}(e) \rrbracket^{\text{init}} &= \llbracket e \rrbracket^{\text{init}} \\
\llbracket \text{factor}(e) \rrbracket^{\text{step}}_\gamma(m) &= \text{let } m', v = \llbracket e \rrbracket^{\text{step}}_\gamma(m) \text{ in } v \times \delta_{m', ()} \\
\left\{ \begin{array}{l} e \text{ where rec init } x = c \\ \text{and } x = e_x \\ \text{and } y = e_y \end{array} \right\}^{\text{init}} &= c, (\llbracket e \rrbracket^{\text{init}}, \llbracket e_x \rrbracket^{\text{init}}, \llbracket e_y \rrbracket^{\text{init}}) \\
\left\{ \begin{array}{l} e \text{ where rec init } x = c \\ \text{and } x = e_x \\ \text{and } y = e_y \end{array} \right\}^{\text{step}}_\gamma(p_x, (m, m_x, m_y)) &= \int \llbracket e_x \rrbracket^{\text{step}}_{\gamma + [x.\text{last} \leftarrow p_x]}(m_x)(dm'_x, dv_x) \\
&\quad \int \llbracket e_y \rrbracket^{\text{step}}_{\gamma + [x.\text{last} \leftarrow p_x, x \leftarrow v_x]}(m_y)(dm'_y, dv_y) \\
&\quad \int \llbracket e \rrbracket^{\text{step}}_{\gamma + [x.\text{last} \leftarrow p_x, x \leftarrow v_x, y \leftarrow v_y]}(m)(dm', dv) \\
&\quad \delta_{(v_x, (m', m'_x, m'_y)), v}
\end{aligned}$$

Fig. 11. A simplified excerpt of the original ProbZelus probabilistic kernel operational semantics [5].

## 6.2 Correspondance between the original and the operational semantics

In this section we show that for scheduled programs where local declarations are all ordered according to data dependencies, the original kernel semantics coincide with the operational semantics, and thus with the denotational semantics.

*Kernel semantics.* The original ProbZelus semantics [5, Section 3.1] is a co-iterative semantics similar to the operational semantics presented in Section 4 where expressions are interpreted as state machines. The key difference is that probabilistic state machines directly computes a stream of unnormalized measures. Given a current state, the transition function  $\llbracket e \rrbracket^{\text{step}}$  of type  $\Gamma \rightarrow S \rightarrow \text{Meas}(S \times V)$  returns a measure over pairs (next state, value).

Figure 11 shows a simplified excerpt of the semantics of probabilistic expressions from [5]. In a probabilistic context, a deterministic expression is interpreted as the Dirac delta measure on the pair (state, value) returned by the deterministic semantics. `sample` evaluates its argument into a new state  $m'$  and a distribution of values  $\mu$ , and returns a measure over pairs (new state, value). `factor` evaluates its argument into a new state  $m'$  and a real value  $v$ , and returns a Dirac delta measure on the pair  $(m', ())$  weighted by  $v$ . To illustrate local declarations, Figure 11 shows the semantics of a simple expression with two local variables  $x$  and  $y$ . The state captures the previous value of the initialized variable  $x$ , and the state of all sub-expressions. The transition function starts in a context where the previous value of  $x$  is bound to a special variable  $x.\text{last}$ , and integrates over all possible executions of the sub-expressions to compute the main expression.

Manuscript submitted to ACM

Since the original ProbZelus semantics focuses on a kernel language where local declarations are all *scheduled* the nested integrals are well-defined.

The `infer` operator is similar to its counterpart in Section 4.2. Given a probabilistic model  $e$  with a state of type  $S$  and a return value of type  $V$ , the state of `infer` is a measure over possible states of the model, and at each step, the transition function returns a measure over possible next states, and a measure over return values.

$$\begin{aligned} \llbracket \text{infer}(e) \rrbracket^{\text{K init}} &: \text{Meas}(S) \\ \llbracket \text{infer}(e) \rrbracket^{\text{K step}} &: \Gamma \rightarrow \text{Meas}(S) \rightarrow \text{Meas}(S) \times \text{Meas}(V) \end{aligned}$$

Given the unnormalized measure defined by a model  $e$ , `infer`( $e$ ) performs the same operations as in the operational semantics (Equation (3) in Section 4.2): 1) integrate over all possible states, 2) normalize the measure, 3) split the result into a distribution of next states and a distribution of values. We then obtain a stream of distributions by iteratively integrating the transition function along the distribution of states.

$$\begin{aligned} \llbracket \text{infer}(e) \rrbracket^{\text{K init}} &= \delta_{\llbracket e \rrbracket^{\text{init}}} \\ \llbracket \text{infer}(e) \rrbracket^{\text{K step}}_Y(\sigma) &= \text{let } v = \int \sigma(dm) \llbracket e \rrbracket^{\text{step}}_Y(m) \text{ in} \\ &\quad \text{let } \bar{v} = v/v(\top) \text{ in} \\ &\quad (\pi_{1*}(\bar{v}), \pi_{2*}(\bar{v})) \end{aligned}$$

*Correspondance.* The definition of `infer` is very similar to its operational semantics counterpart where the function  $\psi(m)$  in Equation (3) plays the role of the semantics of the model. We now show that these two notions coincide.

**PROPOSITION 6.3.** *For all probabilistic expression  $e$  where all equations are scheduled, the kernel semantics corresponds to the operational semantics by integrating over random elements, i.e., for all environment  $\gamma$  and state  $m$ :*

$$\llbracket e \rrbracket^{\text{step}}_Y(m) = \left( \int_{[0,1]^{|e|}} \text{let } m', v, w = \llbracket e \rrbracket^{\text{step}}_Y(m, r) \text{ in } w \times \delta_{(m', v)} dr \right)$$

**PROOF.** The kernel semantics is only defined for a scheduled language. We first prove by induction on the structure of  $e$  that the kernel semantics coincides with the operational semantics where local declarations are ordered and can be computed incrementally without a fixpoint operator. For a set of scheduled equations, the ordered sequence of local declarations yields the same environment as the fixpoint operator (see Proposition A.1). Since all sets of equations are scheduled, the operational semantics (with a fixpoint operator) also coincides with the kernel semantics.

The case `sample`( $\mu$ ) is a simple variable substitution  $x = \text{icdf}_\mu(r)$ . Indeed, any real continuous distribution  $\mu$  is the pushforward by  $\text{icdf}_\mu$  of the uniform distribution over  $[0, 1]$  denoted  $\lambda$ :

$$\int_{[0,1]} \delta_{\text{icdf}_\mu(r)} dr = \int \delta_{\text{icdf}_\mu(r)} \lambda(dr) = \int \delta_x \text{icdf}_{\mu*}(\lambda)(dx) = \int \delta_x \mu(dx) = \mu$$

This property generalizes to discrete distributions, multivariate distributions, and any distributions over Polish spaces. By analogy, we use the notation  $\text{icdf}_\mu$  in all cases.

The case  $E_1$  and  $E_2$  is a consequence of Fubini's theorem.

$$\begin{aligned}
& \int \llbracket E_1 \rrbracket_{\gamma}^{\text{step}} (M_1) (dM'_1, d\rho_1) \int \llbracket E_2 \rrbracket_{\gamma+\rho_1}^{\text{step}} (M_2) (dM'_2, d\rho_2) \delta_{(M'_1, M'_2), \rho_1 + \rho_2} \\
&= \int \left( \int_{[0,1]^{|e_1|}} \text{let } M'_1, \rho_1, w_1 = (E_1)_{\gamma}^{\text{step}} (M_1, r_1) \text{ in } w_1 \times \delta_{M_1, \rho_1} dr_1 \right) (dM'_1, d\rho_1) \\
& \quad \int \left( \int_{[0,1]^{|e_2|}} \text{let } M'_2, \rho_2, w_2 = (E_2)_{\gamma+\rho_1}^{\text{step}} (M_2, r_2) \text{ in } w_2 \times \delta_{M_2, \rho_2} dr_2 \right) (dM'_2, d\rho_2) \\
& \quad \delta_{(M'_1, M'_2), \rho_1 + \rho_2} \\
&= \int_{[0,1]^{|e_1|}} \int_{[0,1]^{|e_2|}} \int \left( \text{let } M'_1, \rho_1, w_1 = (E_1)_{\gamma}^{\text{step}} (M_1, r_1) \text{ in } w_1 \times \delta_{M_1, \rho_1} \right) (dM'_1, d\rho_1) \\
& \quad \int \left( \text{let } M'_2, \rho_2, w_2 = (E_2)_{\gamma+\rho_1}^{\text{step}} (M_2, r_2) \text{ in } w_2 \times \delta_{M_2, \rho_2} \right) (dM'_2, d\rho_2) dr_1 dr_2 \\
& \quad \delta_{(M'_1, M'_2), \rho_1 + \rho_2} \\
&= \int_{[0,1]^{|e_1|+|e_2|}} \text{let } M'_1, \rho_1, w_1 = (E_1)_{\gamma}^{\text{step}} (M_1, r_1) \text{ in} \\
& \quad \text{let } M'_2, \rho_2, w_2 = (E_2)_{\gamma+\rho_1}^{\text{step}} (M_2, r_2) \text{ in} \\
& \quad w_1 \times w_2 \times \delta_{(M'_1, M'_2), \rho_1 + \rho_2} dr_1 dr_2
\end{aligned}$$

Other cases are similar.  $\square$

We can now state the main correctness theorem, i.e., the `infer` operator yields the same stream of distributions in the operational semantics and in the kernel based semantics.

**THEOREM 6.4 (OPERATIONAL SEMANTICS CORRECTNESS).** *For every probabilistic model  $e$  where all equations sets are scheduled, for all distribution of states  $\sigma$ , and for all environment  $\gamma$ :*

$$\begin{aligned}
\llbracket \text{infer}(e) \rrbracket^{\text{init}} &= \llbracket \text{infer}(e) \rrbracket^{\text{K init}} \\
\llbracket \text{infer}(e) \rrbracket_{\gamma}^{\text{step}} (\sigma) &= \llbracket \text{infer}(e) \rrbracket_{\gamma}^{\text{K step}} (\sigma)
\end{aligned}$$

**PROOF.** By construction, in the operational semantics, the first element of the initial state of `infer` is a Dirac delta measure on the initial state of the model which corresponds to the initial state of `infer` in the kernel semantics.

By Proposition 6.3 the un-normalized measure defined by the operational semantics matches the measure computed by the kernel semantics. Given this measure, the rest of the transition function of `infer` is the same in both cases.  $\square$

## 7 APPLICATION: ASSUMED PARAMETER FILTER

In ProbZelus, state-space models can involve two kinds of random variables. *State parameters* are represented by a stream of random variables which evolve over time depending on the previous values and the observations. *Constant parameters* are represented by a random variable whose value is progressively refined from the *prior* distribution with each new observation.

**Example.** Consider the radar example of Section 2 where the current force is an unknown constant  $\theta$ . We want to estimate the moving position (state parameter), and the current force (constant parameter). To estimate  $\theta$ , the `current_force` is now a probabilistic model defined as follows (where the noise parameter `st` is a global constant):

```

proba current_force(x) = theta where
  rec init theta = sample(gaussian(zeros, st))
  and theta = last theta

```

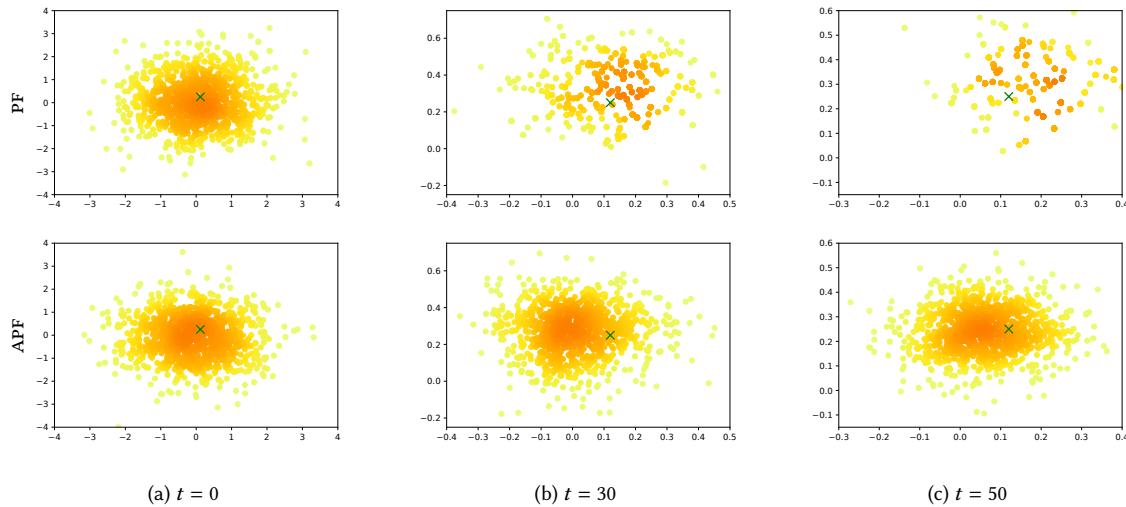


Fig. 12. Estimates of the theta parameter of the radar example over time with a particle filter (PF at the top) and assumed parameter filter (APF at the bottom). The true drift speed is indicated by a green cross. The color gradient represents the dot density. The scale shrinks over time. Results may differ in between runs.

**Filtering.** To estimate state parameters, Sequential Monte Carlo (SMC) inference algorithms rely on filtering techniques [19, 30]. Filtering is an approximate method which consists of deliberately losing information on the current approximation to refocus future estimations on the most significant information. These methods are particularly well suited to the reactive context where a system in interaction with its environment never stops and must execute with bounded resources. All ProbZelus inference methods are SMC algorithms [2, 5, 6]. Unfortunately, this loss of information is harmful for the estimation of constant parameters which do not change over time.

The most basic SMC algorithm, the *particle filter*, approximates the posterior distribution by launching a set of independent executions, called *particles*. At each step, each particle returns a value associated to a score which measures the quality of the value w.r.t. the model. To recenter the inference on the most significant particles, the inference runtime periodically resamples the set of particles according to their weights. The most significant particles are then duplicated while the least interesting ones are dropped.

Unfortunately, constant parameters are only sampled at the beginning of the execution of each particle. After each resampling step, the duplicated particles share the same value for the `theta`. The quantity of information used to estimate constant parameters thus strictly decreases over time until eventually, there is only one possible value left. The upper part of the Figure 12 graphically illustrates this phenomenon.

**Assumed Parameter Filter.** To mitigate this issue, the *Assumed Parameter Filter* (APF) maintains a symbolic distribution for the constant parameters and splits the inference into two steps: 1) estimate the state parameters distributions, and 2) update the constant parameters distributions [35]. The lower part of Figure 12 illustrates the results of APF on the estimation of the drift speed for our radar example.

The APF algorithm assumes that constant parameters are an input of the model, and their prior distributions an input of a new inference operator `APF.infer` (the APF algorithm is described in Appendix B.1). In this section, we present a program transformation which generates models that are exploitable by the APF algorithm. First, a static

Manuscript submitted to ACM

$$\begin{array}{c}
 \begin{array}{ccccc}
 \text{CEXP-C} & \text{CEXP-VAR} & \text{CEXP-PAIR} & \text{CEXP-OP} & \text{CEXP-WHERE} \\
 \frac{}{C \vdash^c c} & \frac{x \in C}{C \vdash^c x} & \frac{C \vdash^c e_1 \quad C \vdash^c e_2}{C \vdash^c (e_1, e_2)} & \frac{C \vdash^c e}{C \vdash^c op(e)} & \frac{C + \text{dom}(E) \vdash^c e \quad C \vdash^c E : \text{dom}(E)}{C \vdash^c e \text{ where rec } E}
 \end{array} \\
 \text{CEQ-INIT} \qquad \qquad \qquad \text{CEQ-LAST} \qquad \qquad \qquad \text{CEQ-CEXP} \qquad \qquad \qquad \text{CEQ-EXP} \\
 \frac{}{C \vdash^c \text{init } x = e : \emptyset} \qquad \qquad \frac{}{C \vdash^c x = \text{last } x : \{x\}} \qquad \qquad \frac{C \vdash^c e}{C \vdash^c x = e : \{x\}} \qquad \qquad \frac{C \not\vdash^c e}{C \vdash^c x = e : \emptyset} \\
 \text{CEQ-AND} \qquad \qquad \qquad \qquad \qquad \frac{C + C_2 \vdash^c E_1 : C_1 \quad C + C_1 \vdash^c E_2 : C_2}{C \vdash^c E_1 \text{ and } E_2 : C_1 + C_2}
 \end{array}$$

Fig. 13. Constant expressions and equations.

analysis identifies the constant parameters and their prior distributions. Then a compilation pass transforms these parameters into additional inputs of the model. We use the denotational semantics of Section 5 to prove the correctness of this transformation, i.e., the transformation preserves the ideal semantics of the program.

**Example.** The radar model (see Figure 1) is compiled into the following model for APF. The constant parameter  $f$  is an argument of the model and the corresponding prior distribution is defined outside the model.

```

let current_force.prior = gaussian(zeros, st)
proba current_force.model(theta, x) = theta

let tracker_prior = current_force.prior
proba tracker.model(theta, y_obs) = x where
  rec x = motion(f)
  and f = current_force.model(theta, x)
  and () = observe(gaussian(g(x), sy), y_obs)

node main(y_obs) = msg where
  rec x_dist = APF.infer(tracker.model, tracker.prior, y_obs)
  and msg = controller(x_dist)

```

### 7.1 Static Analysis

The goal of the static analysis is to identify the constant parameters of each probabilistic node, i.e., initialized random variables ( $\text{init } x = \text{sample}(e)$ ) that are also constant ( $x = \text{last } x$ ). For a program  $\text{prog}$  the judgement  $[], \emptyset \vdash \text{prog} : \Phi, C$  incrementally analyzes all the declarations to build the environment  $\Phi$  which associates to each probabilistic node a type  $\phi$  mapping constant parameters to their prior distributions. The environment  $C$  contains the global constant variables that can be used to define the prior distributions. An excerpt of the type system is given in Figures 13 and 14.

*Constants (CEQ-\* and CEXP-\* rules, Figure 13):* The auxiliary judgement  $C \vdash^c E : C'$  (CEQ-\* rules) identifies constant streams  $C'$  in the set of equations  $E$  given the constant variables  $C$ . A stream  $x$  is constant if it is always equal to its previous value ( $x = \text{last } x$ ) or if it is defined by a constant expression. The auxiliary judgement  $C \vdash^c e$  (CEXP-\* rules)

Manuscript submitted to ACM

$$\begin{array}{c}
 \text{DECL-CST} \quad \text{DECL-PROBA} \\
 \frac{C \vdash^c e}{\Phi, C \vdash \text{let } x = e : \Phi, C + \{x\}} \quad \frac{C \vdash e : \phi}{\Phi, C \vdash \text{proba } f x = e : \Phi + [f \leftarrow \phi], C}
 \end{array}$$

$$\begin{array}{c}
 \text{EXP-RESET} \quad \text{EXP-APP} \quad \text{EXP-WHERE} \\
 \frac{C \vdash e_2 : \phi}{C \vdash \text{reset } e_1 \text{ every } e_2 : \phi} \quad \frac{C \vdash f_\theta(e) : [\theta \leftarrow f.\text{prior}]}{C \vdash e \text{ where rec } E : \phi_e + \phi_E} \\
 \frac{C \vdash e : \phi_e \quad C \vdash^c E : D \quad C, D \vdash E : \phi_E}{C \vdash e \text{ where rec } E : \phi_e + \phi_E}
 \end{array}$$

$$\begin{array}{c}
 \text{EXP-SAMPLE} \quad \text{EQ-ONE} \\
 \frac{x \in D \quad C \vdash^c e}{C, D \vdash \text{init } x = \text{sample}(e) : [x \leftarrow e]} \quad \frac{C \vdash e : \phi}{C, D \vdash x = e : \phi}
 \end{array}$$

$$\begin{array}{c}
 \text{EQ-AND} \\
 \frac{C, D \vdash E_1 : \phi_1 \quad C, D \vdash E_2 : \phi_2}{C, D \vdash E_1 \text{ and } E_2 : \phi_1 + \phi_2}
 \end{array}$$

Fig. 14. Extract constant parameters and associated prior distributions (full type system in Appendix B.2).

checks that an expression defines a constant stream. An expression with a set of local declarations is constant if all the equations define constant streams.

*Declarations* (DECL-\* rules, Figure 14): A global declaration `let x = e` typed in the environment  $\Phi, C$  adds the name  $x$  to the global constant set  $C$  if the expression  $e$  is constant. A probabilistic node `proba f x = e` is associated to the map  $\phi$  computed by the judgement  $C \vdash e : \phi$ .

*Expressions* (EXP-\* rules, Figure 14): The auxiliary judgement  $C \vdash e : \phi$  collects the constant parameters of the sub-expressions. To simplify the analysis, we associate a unique instance name  $\theta$  to each function call and we assume that all variables and instances names are unique, e.g.,  $f(1) + f(2)$  becomes  $f_{\theta_1}(1) + f_{\theta_2}(2)$ . The rule for  $f_\theta(e)$  associates to  $\theta$  the prior distribution of the constant parameters of the body of  $f$ :  $f.\text{prior}$  which is defined as a global variable by the compilation pass. The rule for `reset e1 every e2` focuses only on the condition  $e_2$  because  $e_1$  can be re-initialized and thus is not constant.

*Equations* (EQ-\* rules, Figure 14): The typing of  $e$  `where rec`  $E$  identifies the set  $D$  of constant variables in  $E$ , then types the equations with the judgement  $C, D \vdash E : \phi$  where  $C$  is the set of constant free variables in  $E$ . If a variable  $x$  is introduced by the equation `init x = sample(e)` and is constant ( $x \in D$ ), then  $x$  is a constant parameter and the result type maps  $x$  to the distribution  $e$ .

**Example.** On our example, the variable `theta` is identified as a constant parameter of the node `current_force` and is propagated through the node `tracker` that calls `current_force`. The final environment is:

$$\Phi = [\text{current\_force} \leftarrow [\text{theta} \leftarrow \text{gaussian}(\dots)], \text{tracker} \leftarrow [\theta \leftarrow \text{current\_force}.\text{prior}]]$$

## 7.2 Compilation

To run the APP algorithm, constant parameters must become additional inputs of the model. The inference runtime can thus execute the model multiple times with different values of the constant parameters to update their distributions. The compilation function is defined in Figure 15 by induction on the syntax and relies on the result of the static analysis. The compilation function  $C$  is thus parameterized by the typing environment  $\Phi$  for declarations and the type  $\phi$  for expressions.

Manuscript submitted to ACM

$$\begin{aligned}
C_\Phi(\text{proba } f \ x = e) &= \begin{cases} \text{proba } f \ x = C_\Phi(e) & \text{if } \Phi(f) = [] \\ \text{let } f.\text{prior} = \text{im}(\phi) & \text{with } \phi = \Phi(f) \\ \text{proba } f.\text{model}(\text{dom}(\phi), x) = C_\Phi(e) \end{cases} \\
C_\phi(e \text{ where rec } E) &= C_\phi(e) \text{ where rec } C_\phi(E) \\
C_\phi(\text{init } x = e) &= \begin{cases} \emptyset & \text{if } x \in \text{dom}(\phi) \\ \text{init } x = C_\phi(e) & \text{otherwise} \end{cases} \\
C_\phi(x = e) &= \begin{cases} \emptyset & \text{if } x \in \text{dom}(\phi) \\ x = C_\phi(e) & \text{otherwise} \end{cases} \\
C_\phi(f_\theta(e)) &= \begin{cases} f(C_\phi(e)) & \text{if } f \text{ is deterministic or } \phi = [] \\ f.\text{model}(\theta, C_\phi(e)) & \text{if } \theta \in \text{dom}(\phi) \\ f.\text{model}(\theta, C_\phi(e)) \text{ where} \\ \text{rec init } \theta = \text{sample}(f.\text{prior}) \\ \text{and } \theta = \text{last } \theta & \text{otherwise} \end{cases} \\
C_\phi(\text{infer}(f(e))) &= \text{APF.infer}(f.\text{model}, f.\text{prior}, C_\phi(e))
\end{aligned}$$

Fig. 15. APF compilation (full definition in Appendix B.3).

A model `proba f x = e` such that  $\Phi(f) = \phi$  (i.e.,  $C \vdash e : \phi$ ) is compiled into two statements: 1) the prior distribution of the constant parameters in  $e$ , `let f.prior = im(ϕ)`, and 2) a new model that takes the constant parameters  $\text{dom}(\phi)$  as additional arguments, `proba f.model(dom(ϕ), x) = C_ϕ(e)`.

The compilation function of an expression  $C_\phi$  removes the definitions of the constant parameters  $x \in \text{dom}(\phi)$ . The `where/rec` case effectively removes the constant parameters by keeping only the equations defining variables  $x \notin \text{dom}(\phi)$ . The main difficulty is to handle constant parameters introduced by a function call.

- If the node is deterministic, there is no constant parameter.
- If the constant parameters of the callee are also constant parameters of the caller, we have  $\theta \in \text{dom}(\phi)$  and we just replace the call to  $f_\theta$  with a call to  $f.\text{model}$  using the instance name for the constant parameters.
- Otherwise, the constant parameters of the callee are not constant for the caller because the instance  $f_\theta$  is used inside a `reset/every` or a `present/else`. In this case, we redefine these parameters locally by sampling their prior distribution  $f.\text{prior}$ .

Finally, the call to `infer(f(e))` is replaced by a call to `APF.infer(f.model, f.prior, e)`.

### 7.3 Correctness

We use the denotational semantics defined in Section 5.2 to prove the correctness of the APF compilation pass. First, we prove that any probabilistic expression is equivalent to its compiled version computed in an environment which already contains the definition of the constant parameters. The main theorem which relates `infer(f(e))` and `APF.infer(f.model, f.prior, e)` then corresponds to the case  $f_\theta(e)$  in Figure 15.

*Definition 7.1.* For a model  $f$  that is compiled into  $f.\text{prior}$  and  $f.\text{model}$ , the ideal semantics of `APF.infer` externalizes the definition of the constant parameters: `APF.infer(f.model, f.prior, e)` denotes the expression:

`infer(f.model(θ, e) where rec init θ = sample(f.prior))`

We first prove the two following correctness lemmas for expressions and equations in parallel.

LEMMA 7.2. *For all probabilistic functions  $f$  such that  $\Phi(f) = \phi$ , for all expressions  $e$  in the body of  $f$  such that  $C \vdash e : \phi_e$ , for all contexts  $H, R$ , there is a permutation relating  $R$  and  $[R' : R^p]$  such that*

$$H, R \vdash e \Downarrow (s, w) \iff H + H_f, R' \vdash C_\phi(e) \Downarrow (s, w).$$

where  $\vec{p} = \text{dom}(\phi_e)$  is the list of constant parameters in  $e$ ,  $\vec{\mu} = \text{im}(\phi_e)$  are the corresponding prior distributions and  $H_f$  is the context that already contains the definitions of all the constant parameters in  $f$  including  $\vec{p}$ , i.e.,  $H_f(\vec{p}) = \text{icdf}_{\vec{\mu}}(R_0^p)$ .

LEMMA 7.3. *For all equations  $E$  in the body of  $f$  such that  $C, D \vdash E : \phi_E$ , for all contexts  $H, R$ , there is a permutation relating  $R$  and  $[R' : R^p]$  such that*

$$H, R \vdash E : W \iff H + H_f, R' \vdash C_\phi(E) : W.$$

PROOF. The proof is by induction on the expression and the size of the context, i.e., the number of declarations before the expression. For each induction case, we apply Proposition 5.3 by giving a permutation relating  $R$  and  $R' + R^p$  inducing the coupling, and show that a semantics derivation for  $e$  in a context  $H, R$  is equivalent to a semantics derivation for  $C_\phi(e)$  in the context  $H + H_f, R'$ . We focus on the most interesting cases, i.e., expressions altered by the compilation function.

*Case `init`  $x = \text{sample}(d)$  and  $x = \text{last}$   $x$  where  $x \in \text{dom}(\phi)$ .* We use the permutation relating  $[R_x] [] : [R_x]$  and with  $v_x = \text{icdf}_\mu(R_{x0})$

$$\frac{\overline{H \vdash d \downarrow \mu \cdot s_\mu}}{H + [x \leftarrow v_x], [R_x] \vdash \text{init } x = \text{sample}(d) : 1 \quad H + [x \leftarrow v_x], [] \vdash x = \text{last } x : 1} \quad H + [x \leftarrow v_x], [R_x] \vdash \text{init } x = \text{sample}(d) \text{ and } x = \text{last } x : 1$$

On the other hand, because  $x$  is a constant parameter,  $C_\phi(d) = \emptyset$  and  $H_f(x) = v_x$ .

$$\frac{H_f(x) = v_x \quad \overline{H + H_f, [] \vdash \emptyset : 1}}{H + H_f, [] \vdash C_\phi(\text{init } x = \text{sample}(d) \text{ and } x = \text{last } x) : 1}$$

This results can then be generalized to arbitrary sets of equations where the two equations are not necessarily grouped together at the cost of an additional permutation.

*Case  $g_\theta(e)$ .* By induction we have two permutations relating  $R_e$  with  $[R'_e : R_e^p]$  and relating  $R_g$  with  $[R'_g : R_g^p]$ . With `proba`  $g x = e_g$  and  $C \vdash e_g : \phi_g$ , we can apply the induction hypothesis on  $e_g$  because there is no possible recursive call. The callee context for  $e_g$  is thus strictly smaller than the caller context. We also have  $H_g = [\vec{p}_g \leftarrow \vec{v}_p]$  with  $\vec{p}_g = \text{dom}(\phi_g)$ ,  $\vec{\mu}_g = \text{im}(\phi_g)$ , and  $\vec{v}_p = \text{icdf}_{\vec{\mu}_g}(R_{g0}^p)$ .

$$\frac{\overline{H + H_f, R'_e \vdash C_\phi(e) \Downarrow (s_e, w_e)}}{H, R_e \vdash e \Downarrow (s_e, w_e)} \quad \frac{[x \leftarrow s_e] + [\vec{p}_g \leftarrow \vec{v}_p], R'_g \vdash C_{\phi_g}(e_g) \Downarrow (s, w)}{[x \leftarrow s_e], R_g \vdash e_g \Downarrow (s, w)}$$

$$\frac{}{H, [R_e : R_g] \vdash g_\theta(e) \Downarrow (s, w_e \times w)}$$

On the other end, by construction we have  $\text{proba } g.\text{model}(\vec{p}_g, x) = C_{\phi_g}(e_g)$  and there are two cases. If  $\theta \in \text{dom}(\phi)$ , then the constant parameters are already in the context and  $H_f(\theta) = \vec{v}_p$ . Using the permutation relating  $[R_e : R_g]$  with  $[R'_e : R'_g] : [R_e^P : R_g^P]$  we have:

$$\frac{\begin{array}{c} H + H_f, [] \vdash \theta \Downarrow (\vec{v}_p, 1) \quad H + H_f, R'_e \vdash C_\phi(e) \Downarrow (s_e, w_e) \\ \hline H + H_f, R'_e \vdash (\theta, C_\phi(e)) \Downarrow ((\vec{v}_p, s_e), w_e) \end{array}}{H + H_f, [R'_e, R'_g] \vdash g.\text{model}(\theta, C_\phi(e)) \Downarrow (s, w_e \times w)}$$

$$\frac{H + H_f, [R'_e, R'_g] \vdash g.\text{model}(\theta, C_\phi(e)) \Downarrow (s, w_e \times w)}{H + H_f, [R'_e, R'_g] \vdash C_\phi(g_\theta(e)) \Downarrow (s, w_e \times w)}$$

Finally if  $\theta \notin \text{dom}(\phi)$ , the constant parameters are not in the context and the compilation adds a defining equation for  $\theta$ . Using the permutation relating  $[R_e : R_g]$  with  $[R'_e : R'_g] : [R_e^P : R_g^P]$  we have:

$$\frac{\begin{array}{c} g.\text{prior} = \vec{\mu}_g \\ \dots \\ H + H_f + [\theta \leftarrow \vec{v}_p], [R'_e, R'_g] \vdash g.\text{model}(\theta, C_\phi(e)) \Downarrow (s, w) \end{array}}{H + H_f + [\theta \leftarrow \vec{v}_p], R_g^P \vdash \text{sample}(g.\text{prior}) \Downarrow (\vec{v}_p, 1)}$$

$$\frac{H + H_f + [\theta \leftarrow \vec{v}_p], R_g^P \vdash \text{sample}(g.\text{prior}) \Downarrow (\vec{v}_p, 1)}{H + H_f + [\theta \leftarrow \vec{v}_p], R_g^P \vdash \text{init } \theta = \text{sample}(g.\text{prior}) : 1}$$

$$\frac{H + H_f, [R'_e, R'_g, R_g^P] \vdash g.\text{model}(\theta, C_\phi(e)) \text{ where rec init } \theta = \text{sample}(g.\text{prior}) \Downarrow (s, w)}{H + H_f, [R'_e, R'_g, R_g^P] \vdash C_\phi(g_\theta(e)) \Downarrow (s, w)}$$

□

We can now state and prove the correctness of the APF compilation pass.

**THEOREM 7.4 (APF COMPILATION).** *For all probabilistic nodes  $f$ , for all contexts  $H$ ,*

$$H \vdash \text{infer}(f(e)) \downarrow d \iff H \vdash \text{APF.infer}(f.\text{model}, f.\text{prior}, e) \downarrow d$$

**PROOF.** From Definition 7.1, we need to show that for all random streams  $R$ :

$$H, R \vdash f_\theta(e) \Downarrow (s, w) \iff H + [\theta \leftarrow \vec{v}_p], R' \vdash f.\text{model}(\theta, e) \Downarrow (s, w)$$

with  $f.\text{prior} = \vec{\mu}$  and  $\vec{v}_p = \text{icdf}_{\vec{\mu}}(R_0^P)$ . This corresponds to the case  $f_\theta(e)$  with  $\theta \in \text{dom}(\phi)$ . □

## 8 RELATED WORK

*Probabilistic Semantics.* In the seminal work [45], two semantics are introduced for a probabilistic imperative language. The first one is already sampling semantics that first picks an infinite stack of random numbers and then executes the program deterministically. In the second semantics, programs are interpreted as distribution transformer: input distributions are transformed to an output sub-probability distribution. This second semantics is defined using measurable functions and kernels.

Probabilistic Coherent Spaces is a generalization of this idea to higher-order types but for discrete probability [27]. This setting has been extended to continuous distributions with models based on positive cones [26, 34], a variation on Banach spaces with positive scalars [51]. To interpret the sampling operation, cones have to be equipped with a measurability structure such that measures and integration can be defined for any types [33].

Quasi-Borel spaces (QBS) define an alternative semantics of higher-order probabilistic programs with conditioning [41, 53] based on measurable spaces and kernels. A probabilistic expression is interpreted as a quasi-Borel measure, i.e., an equivalence class of pairs  $[\alpha, \mu]$  where  $\mu$  is a measure over  $\mathbb{R}$ , and  $\alpha$  is a measurable function from  $\mathbb{R}$  to values. Intuitively, the corresponding distribution is obtained as the pushforward of  $\mu$  along  $\alpha$ . Recent work extends this

Manuscript submitted to ACM

formalism to capture lazy data structures and streams in a functional probabilistic language [28]. Our new semantics rely on a similar representation: probabilistic expressions are interpreted by pushingforward a uniform measure over  $[0, 1]^{|e|}$  along a measurable function.

Our semantics are also closely related to the sampling semantics of [10] for a higher-order lambda calculus where a sequence of random draws is a parameter of the evaluation rules. The main difference is that we focus on a stream based reactive language. On the one hand, we show how to lift the sampling semantics to the reactive setting. On the other hand, we show how to adapt the reactive semantics (operational and denotational) to the probabilistic setting. In addition, there is no recursion and no loop in our language. We thus don't need to define the semantics using step-indexing.

Other works [23, 55, 56] define a similar semantics using  $[0, 1]^\omega$  as the *entropy* which describes an idealized random numbers generator. The stream head returns a random element, and a stream can be split (e.g., filtering odd/even indices) to generate two new entropies. This technique is used to interpret programs where random variables can be dynamically created (e.g., in recursive calls). An originality of our language is that we can statically compute the number of streams of random variables in a program, i.e., the calls to `sample`. The random streams exactly correspond to these sample sites and there is no need for a splittable entropy source. This is reminiscent of the Stan language [17] where all random variables must be declared in the `parameters` block, except that in our case, parameters are streams of random variables.

*Mutually recursive equations.* In the original ProbZelus semantics, the interpretation of local declarations yields nested integrals that are only well defined if equations are ordered according to data dependencies (Figure 11 in Section 6.2). Our new semantics relaxes this assumptions using a fixpoint operator.

Existing works define a fixpoint operator to interpret loops and recursions in probabilistic lambda-calculi [34, 44, 54]. The least element is the null measure, i.e., for all measurable set  $U$ ,  $\perp(U) = 0$ , and the partial order is  $\mu_1 \leq \mu_2$  iff  $\forall U. \mu_1(U) \leq \mu_2(U)$ . Unfortunately, using this CPO the semantics of a set of equations is always the null measure. An alternative is to lift the deterministic operational semantics of Section 4 to measures, i.e., on measure over  $V_\perp = V + \{\perp\}$  with the least element  $\perp = \delta_{[x \leftarrow \perp]_{x \in X}}$ . However, using this CPO the semantics of `sample`( $e$ ) is re-computed at each fixpoint iteration which breaks the correlation between mutually recursive equations. Therefore, we cannot directly define a schedule agnostic kernel semantics.

**Example.** Consider the following example:  $x = y$  and  $y = \text{sample}(\mu)$ . In the kernel semantics, local declarations are interpreted as integration. The semantics is thus the fixpoint of the following function  $F$  where the argument  $\rho$  is a measure over environments.

$$F(\rho) = \int \rho(dr) \int \delta_{r(y)}(dx) \int \mu(dy) \delta_{[x \leftarrow x, y \leftarrow y]}$$

If the least element  $\rho_0$  is the null measure, then  $\rho_1 = F(\rho_0)$  is also the null measure. If  $\rho_0 = \delta_{[x \leftarrow \perp, y \leftarrow \perp]}$  we get :

$$\begin{aligned} \rho_0 &= \delta_{[x \leftarrow \perp, y \leftarrow \perp]} \\ \rho_1 &= \int \mu(dy_1) \delta_{[x \leftarrow \perp, y \leftarrow y_1]} \\ \rho_2 &= \int \mu(dy_1) \int \mu(dy_2) \delta_{[x \leftarrow y_1, y \leftarrow y_2]} \\ \rho_3 &= \int \mu(dy_1) \int \mu(dy_2) \int \mu(dy_3) \delta_{[x \leftarrow y_2, y \leftarrow y_3]} \end{aligned}$$

The fixpoint is reached after 3 iterations, but the correlation between  $x$  and  $y$  is lost. The semantics differs from the scheduled version  $y = \text{sample}(\mu) \text{ and } x = y$  which yields the expected measure:  $\int \mu(dy) \delta_{[x \leftarrow y, y \leftarrow y]}$ .

Positive cones and  $\omega$ QBS are endowed with a structure of CPO resulting in adequate denotational semantics for probabilistic programming with sampling from continuous distributions, recursive types and term recursion [34, 54]. Although our language does not support term recursion, fixpoint operators in positive cones or  $\omega$ QBS might be adapted to give a probabilistic semantics of mutually recursive equations related to our operational semantics. We leave for future investigations these connections.

*Program Equivalence.* Probabilistic bisimulation has been introduced for testing equivalence of (discrete) probabilistic systems [47] and generalized to study Labelled Markov Process (continuous systems) [31]. Probabilistic coupling is a classic proof technique for probabilistic bisimulation, as shown in [29] using a categorical framework and a co-algebraic presentation of probabilistic systems. Our presentation of probabilistic bisimulation and coupling for state machines follows this line of research.

Probabilistic coupling has been extensively used for discrete systems. In [3], an algorithm, based on couplings, has been introduced to effectively construct probabilistic bisimulation for discrete systems. Coupling has been applied to the formal verification of distribution equivalence, convergence, and stochastic domination [4, 42]. Probabilistic coupling has also been used to define a probabilistic bisimulation for a (higher-order) lambda-calculus with continuous probabilistic choice [46].

A logical relation is proposed in [23, 55, 56] to reason about contextual equivalence for probabilistic programs. In particular [55] uses this framework to prove that reordering declarations preserves the semantics with a permutation of the entropy stream. This reasoning is reminiscent of the proof of Lemma 7.2 (APF Correctness). We thus show that we can apply similar techniques with our denotational semantics which manipulates infinite streams and mutually recursive equations. Extending these works to define an equational theory to reason about reactive program equivalence beyond permutations of the random variables streams is a promising future work.

## 9 CONCLUSION

In this paper we proposed two semantics for a reactive probabilistic programming languages: an operational semantics and a denotational semantics. Both semantics are schedule agnostic, i.e., sets of mutually recursive equations can be interpreted in arbitrary order, a key property of synchronous dataflow languages. We defined for both semantics equivalence of programs. The operational semantics manipulates state machines and equivalence reasoning requires the description of bisimulations on states. The denotational semantics directly manipulates streams, which can lighten program equivalence reasoning for probabilistic expressions. We then defined a program transformation required to run an optimized inference algorithm for state-space models with constant parameters and used the denotational semantics to prove the correctness of the transformation.

*Acknowledgements.* The authors are grateful to the following people for many discussions and helpful feedback: Timothy Bourke, Grégoire Bussone, Raphaëlle Crubillé, Thomas Ehrhard, Guillaume Geoffroy, Adrien Guatto, Patrick Hoscheit, Paul Jeanmaire, Dexter Kozen, Basile Pesin, Gordon Plotkin, and Marc Pouzet. This work was supported by the project ReaLiSe, Émergence Ville de Paris 2021-2025.

Manuscript submitted to ACM

## REFERENCES

- [1] Eric Atkinson, Guillaume Baudart, Louis Mandel, Charles Yuan, and Michael Carbin. 2021. Statically bounded-memory delayed sampling for probabilistic streams. *Proc. ACM Program. Lang.* 5, OOPSLA (2021), 1–28.
- [2] Eric Atkinson, Charles Yuan, Guillaume Baudart, Louis Mandel, and Michael Carbin. 2022. Semi-symbolic inference for efficient streaming probabilistic programming. *Proc. ACM Program. Lang.* 6, OOPSLA2 (2022), 1668–1696.
- [3] Christel Baier, Bettina Engelen, and Mila E. Majster-Cederbaum. 2000. Deciding Bisimilarity and Similarity for Probabilistic Processes. *J. Comput. Syst. Sci.* 60, 1 (2000), 187–231.
- [4] Gilles Barthe, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2017. Coupling proofs are probabilistic product programs. In *POPL*.
- [5] Guillaume Baudart, Louis Mandel, Eric Atkinson, Benjamin Sherman, Marc Pouzet, and Michael Carbin. 2020. Reactive Probabilistic Programming. In *PLDI*.
- [6] Guillaume Baudart, Louis Mandel, and Reyyan Tekin. 2022. JAX Based Parallel Inference for Reactive Probabilistic Programming. In *LCTES*.
- [7] Albert Benveniste, Paul Caspi, Stephen A. Edwards, Nicolas Halbwachs, Paul Le Guernic, and Robert de Simone. 2003. The synchronous languages 12 years later. *Proc. IEEE* 91, 1 (2003), 64–83.
- [8] Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul A. Szerlip, Paul Horsfall, and Noah D. Goodman. 2019. Pyro: Deep Universal Probabilistic Programming. *J. Mach. Learn. Res.* 20 (2019), 28:1–28:6.
- [9] Richard Blute, Josée Desharnais, Abbas Edalat, and Prakash Panangaden. 1997. Bisimulation for Labelled Markov Processes. In *LICS*.
- [10] Johannes Borgström, Ugo Dal Lago, Andrew D. Gordon, and Marcin Szymbczak. 2016. A lambda-calculus foundation for universal probabilistic programming. In *ICFP*.
- [11] Timothy Bourke, Lélio Brun, Pierre-Évariste Dagand, Xavier Leroy, Marc Pouzet, and Lionel Rieg. 2017. A formally verified compiler for Lustre. In *PLDI*.
- [12] Timothy Bourke, Lélio Brun, and Marc Pouzet. 2020. Mechanized semantics and verified compilation for a dataflow synchronous language with reset. *Proc. ACM Program. Lang.* 4, POPL (2020), 44:1–44:29.
- [13] Timothy Bourke, Paul Jeannaire, Basile Pesin, and Marc Pouzet. 2021. Verified Lustre Normalization with Node Subsampling. *ACM Trans. Embed. Comput. Syst.* 20, 5s (2021), 98:1–98:25.
- [14] Timothy Bourke, Paul Jeannaire, and Marc Pouzet. 2025. Functional Stream Semantics for a Synchronous Block-Diagram Compiler. In *LICS*.
- [15] Timothy Bourke, Basile Pesin, and Marc Pouzet. 2023. Verified Compilation of Synchronous Dataflow with State Machines. *ACM Trans. Embed. Comput. Syst.* 22, 5s (2023), 137:1–137:26.
- [16] Timothy Bourke and Marc Pouzet. 2013. Zélus: a synchronous language with ODEs. In *HSCC*.
- [17] Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A probabilistic programming language. *J. Statistical Software* 76, 1 (2017), 1–37.
- [18] Paul Caspi and Marc Pouzet. 1998. A Co-iterative Characterization of Synchronous Stream Functions. In *CMCS*.
- [19] Nicolas Chopin and Omiros Papaspiliopoulos. 2020. *An introduction to sequential Monte Carlo*. Springer.
- [20] Jean-Louis Colaço, Bruno Pagano, and Marc Pouzet. 2017. SCADE 6: A formal language for embedded critical software development. In *TASE*.
- [21] Jean-Louis Colaço and Marc Pouzet. 2003. Clocks as First Class Abstract Types. In *EMSOFT*.
- [22] Jean-Louis Colaço, Michael Mendler, Baptiste Pauget, and Marc Pouzet. 2023. A Constructive State-based Semantics and Interpreter for a Synchronous Data-flow Language with State machines. In *EMSOFT*.
- [23] Ryan Culpepper and Andrew Cobb. 2017. Contextual Equivalence for Probabilistic Programs with Continuous Random Variables and Scoring. In *ESOP*.
- [24] Pascal Cuoq and Marc Pouzet. 2001. Modular Causality in a Synchronous Stream Language. In *ESOP*.
- [25] Marco F. Cusumano-Towner, Feras A. Saad, Alexander K. Lew, and Vikash K. Mansinghka. 2019. Gen: a general-purpose probabilistic programming system with programmable inference. In *PLDI*.
- [26] Fredrik Dahlqvist and Dexter Kozen. 2020. Semantics of higher-order probabilistic programs with conditioning. *Proc. ACM Program. Lang.* 4, POPL (2020), 57:1–57:29.
- [27] Vincent Danos and Thomas Ehrhard. 2011. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Inf. Comput.* 209, 6 (2011), 966–991.
- [28] Swaraj Dash, Younousse Kaddar, Hugo Paquet, and Sam Staton. 2023. Affine Monads and Lazy Structures for Bayesian Programming. *Proc. ACM Program. Lang.* 7, POPL (2023), 1338–1368.
- [29] Erik P. de Vink and Jan J. M. M. Rutten. 1997. Bisimulation for Probabilistic Transition Systems: A Coalgebraic Approach. In *ICALP*.
- [30] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. 2006. Sequential Monte Carlo samplers. *J. Royal Statistical Society: Series B (Statistical Methodology)* 68, 3 (2006), 411–436.
- [31] Josée Desharnais, Abbas Edalat, and Prakash Panangaden. 2002. Bisimulation for Labelled Markov Processes. *Inf. Comput.* 179, 2 (2002), 163–193.
- [32] Luc Devroye. 2006. Nonuniform random variate generation. *Handbooks in operations research and management science* 13 (2006), 83–121.
- [33] Thomas Ehrhard and Guillaume Geoffroy. 2023. *Integration in cones*. Technical Report. IRIF.
- [34] Thomas Ehrhard, Michele Pagani, and Christine Tasson. 2018. Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. *Proc. ACM Program. Lang.* 2, POPL (2018), 59:1–59:28.

- [35] Yusuf Bugra Erol, Yi Wu, Lei Li, and Stuart Russell. 2017. A Nearly-Black-Box Online Algorithm for Joint Parameter and State Estimation in Temporal Models. In *AAAI*.
- [36] Hong Ge, Kai Xu, and Zoubin Ghahramani. 2018. Turing: A Language for Flexible Probabilistic Inference. In *Proceedings of Machine Learning Research*.
- [37] Noah D. Goodman and Andreas Stuhlmüller. 2014. The Design and Implementation of Probabilistic Programming Languages. <http://dippl.org>
- [38] Maria I. Gorinova, Andrew D. Gordon, and Charles Sutton. 2019. Probabilistic programming with densities in SlicStan: efficient, flexible, and deterministic. *Proc. ACM Program. Lang.* 3, POPL (2019), 35:1–35:30.
- [39] Nicolas Halbwachs, Paul Caspi, Pascal Raymond, and Daniel Pilaud. 1991. The Synchronous Dataflow Programming Language LUSTRE. *Proc. of the IEEE* 79, 9 (September 1991), 1305–1320.
- [40] David Harel. 1987. Statecharts: A Visual Formalism for Complex Systems. *Sci. Comput. Program.* 8, 3 (1987), 231–274.
- [41] Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. 2017. A convenient category for higher-order probability theory. In *LICS*.
- [42] Justin Hsu. 2017. *Probabilistic Couplings for Probabilistic Reasoning*. Ph.D. Dissertation. University of Pennsylvania.
- [43] The MathWorks Inc. 2022. *Simulation and Model-Based Design (R2024a)*. Natick, Massachusetts, United States. <https://www.mathworks.com/products/simulink.html>
- [44] Claire Jones and Gordon D. Plotkin. 1989. A Probabilistic Powerdomain of Evaluations. In *LICS*.
- [45] Dexter Kozen. 1981. Semantics of Probabilistic Programs. *J. Comput. Syst. Sci.* 22, 3 (1981), 328–350.
- [46] Ugo Dal Lago and Francesco Gavazzo. 2019. On Bisimilarity in Lambda Calculi with Continuous Probabilistic Choice. In *MFPS*.
- [47] Kim Guldstrand Larsen and Arne Skou. 1989. Bisimulation Through Probabilistic Testing. In *POPL*.
- [48] Wonyeol Lee, Hangyeol Yu, Xavier Rival, and Hongseok Yang. 2020. Towards verified stochastic variational inference for probabilistic programs. *Proc. ACM Program. Lang.* 4, POPL (2020), 16:1–16:33.
- [49] Michael Montemerlo, Sebastian Thrun, Daphne Koller, and Ben Wegbreit. 2002. FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping Problem. In *AAAI*.
- [50] David Michael Ritchie Park. 1981. Concurrency and Automata on Infinite Sequences. In *Theoretical Computer Science*.
- [51] Peter Selinger. 2004. Towards a semantics for higher-order quantum computation.. In *QPL*.
- [52] Sam Staton. 2017. Commutative Semantics for Probabilistic Programming. In *ESOP*.
- [53] Sam Staton, Hongseok Yang, Frank D. Wood, Chris Heunen, and Ohad Kammar. 2016. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *LICS*.
- [54] Matthijs Vákár, Ohad Kammar, and Sam Staton. 2019. A domain theory for statistical probabilistic programming. *Proc. ACM Program. Lang.* 3, POPL (2019), 36:1–36:29.
- [55] Mitchell Wand, Ryan Culpepper, Theophilos Giannakopoulos, and Andrew Cobb. 2018. Contextual equivalence for a probabilistic language with continuous random variables and recursion. *Proc. ACM Program. Lang.* 2, ICFP (2018), 87:1–87:30.
- [56] Yizhou Zhang and Nada Amin. 2022. Reasoning about “reasoning about reasoning”: semantics and contextual equivalence for probabilistic programs with nested queries and recursion. *Proc. ACM Program. Lang.* 6, POPL (2022), 1–28.

$$\begin{aligned}
|e| &= 0 & \text{if } e \text{ is deterministic} \\
|\text{sample}(e)| &= 1 \\
|\text{factor}(e)| &= 0 \\
|f(e)| &= |e| + |e_f| & \text{where } \text{proba } f x = e_f \\
|\text{present } e \rightarrow e_1 \text{ else } e_2| &= |e_1| + |e_2| \\
|\text{reset } e_1 \text{ every } e_2| &= |e_1| \\
|e \text{ where rec } E| &= |e| + |E|
\end{aligned}$$
  

$$\begin{aligned}
|x = e| &= |e| \\
|\text{init } x = e| &= |e| \\
|E_1 \text{ and } E_2| &= |E_1| + |E_2|
\end{aligned}$$

Fig. 16. Dimension of probabilistic expressions and equations.

## A SEMANTICS

### A.1 Operational semantics

*Model dimension.* The dimension of a probabilistic model is defined as the number of random variables in the model, i.e., the number of calls to the `sample` operator. Since loops and recursive calls are not allowed in the language of Figure 2, the dimension of a model can be statically computed with the function  $|e : V| : \mathbb{N}$  defined in Figure 16.

The operational semantics is presented in Section 4. Figures 17 and 18 presents the full semantics for expressions and equations. The additional rules are for `present` and `reset`.

The initialization of  $\text{present } e \rightarrow e_1 \text{ else } e_2$  allocates memory for  $e$ ,  $e_1$  and  $e_2$  and counts the number of random variables in  $e_1$  and  $e_2$  ( $e$  is deterministic and does not have any random variable). The step function first executes  $e$  and depending on its value executes  $e_1$  or  $e_2$ . The initialization of `reset`  $e_1 \text{ every } e_2$  duplicates the memory needed to execute  $e_1$ . That way, in the step function, only the second copy is updated by the transition and if  $e_1$  is reset, the execution restarts from the initial memory state.

*Scheduling.* To compare the new denotational semantics with the kernel semantics, it is useful to define an alternative semantics for a scheduled language without a fixpoint operator. This alternative semantics  $(e)^{\text{s step}}$  exactly matches the denotational semantics except for the two following rules:

$$\begin{aligned}
(E_1 \text{ and } E_2)_{\gamma}^{\text{s step}} ((M_1, M_2), [r_1 : r_2]) &= \text{let } M'_1, \rho_1, w_1 = (E_1)_{\gamma}^{\text{s step}} (M_1, r_1) \text{ in} \\
&\quad \text{let } M'_2, \rho_2, w_2 = (E_2)_{\gamma+\rho_1}^{\text{s step}} (M_2, r_2) \text{ in} \\
&\quad (M'_1, M'_2), \rho_1 + \rho_2, w_1 \times w_2
\end{aligned}$$
  

$$\begin{aligned}
(e \text{ where rec } E)_{\gamma}^{\text{s step}} ((m, M), [r_e : r_E]) &= \text{let } M', \rho, W = (E)_{\gamma}^{\text{s step}} (M, r_E) \text{ in} \\
&\quad \text{let } m', v, w = (e)_{\gamma+\rho}^{\text{s step}} (m, r_e) \text{ in} \\
&\quad (m', M'), v, w \times W
\end{aligned}$$

Since all equations are scheduled, the environment produced by a set of equations can be computed incrementally and there is no need for a fixpoint operator to interpret local declarations.

$$\begin{aligned}
\langle e \rangle_Y^{\text{init}} &= \llbracket e \rrbracket^{\text{init}} & \text{if } e \text{ is deterministic} \\
\langle e \rangle_Y^{\text{step}} (m, []) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}} (m) \text{ in } m', v, 1 \\
\langle \text{sample}(e) \rangle_Y^{\text{init}} &= \llbracket e \rrbracket^{\text{init}} \\
\langle \text{sample}(e) \rangle_Y^{\text{step}} (m, [r]) &= \text{let } m', \mu = \llbracket e \rrbracket_Y^{\text{step}} (m) \text{ in } m', \text{icdf}_{\mu}(r), 1 \\
\langle \text{factor}(e) \rangle_Y^{\text{init}} &= \llbracket e \rrbracket^{\text{init}} \\
\langle \text{factor}(e) \rangle_Y^{\text{step}} (m, []) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}} (m) \text{ in } m', (), v \\
\langle f(e) \rangle_Y^{\text{init}} &= \langle e \rangle^{\text{init}}, \langle e_f \rangle^{\text{init}} & \text{where } \text{proba } f x = e_f \\
\langle f(e) \rangle_Y^{\text{step}} ((m_e, m_f), [r_e : r_f]) &= \text{let } m'_e, v_e, w_e = \langle e \rangle_Y^{\text{step}} (m_e, r_e) \text{ in} \\
&\quad \text{let } m'_f, v, w_f = \langle e_f \rangle_{[x \leftarrow v_e]}^{\text{step}} (m_f, r_f) \text{ in} \\
&\quad (m'_f, m'_e), v, w_e \times w_f \\
\langle \text{present } e \rightarrow e_1 \text{ else } e_2 \rangle_Y^{\text{init}} &= \llbracket e \rrbracket^{\text{init}}, \langle e_1 \rangle^{\text{init}}, \langle e_2 \rangle^{\text{init}} \\
\langle \text{present } e \rightarrow e_1 \text{ else } e_2 \rangle_Y^{\text{step}} ((m, m_1, m_2), [r_1 : r_2]) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}} (m) \text{ in} \\
&\quad \text{if } v \text{ then let } (m'_1, v_1, w) = \langle e_1 \rangle_Y^{\text{step}} (m_1, r_1) \\
&\quad \quad \text{in } (m', m'_1, m_2), v_1, w \\
&\quad \text{else let } (m'_2, v_2, w) = \langle e_2 \rangle_Y^{\text{step}} (m_2, r_2) \\
&\quad \quad \text{in } (m', m_1, m'_2), v_2, w \\
\langle \text{reset } e_1 \text{ every } e_2 \rangle_Y^{\text{init}} &= \langle e_1 \rangle^{\text{init}}, \langle e_1 \rangle^{\text{init}}, \llbracket e_2 \rrbracket^{\text{init}} \\
\langle \text{reset } e_1 \text{ every } e_2 \rangle_Y^{\text{step}} ((m_0, m_1, m_2), w, r) &= \text{let } m'_2, v_2 = \llbracket e_2 \rrbracket_Y^{\text{step}} (m_2) \text{ in} \\
&\quad \text{let } m'_1, v_1, w = \langle e_1 \rangle_Y^{\text{step}} (\text{if } v_2 \text{ then } m_0 \text{ else } m_1, r) \text{ in} \\
&\quad \quad (m_0, m'_1, m'_2), v_1, w \\
\langle e \text{ where rec } E \rangle_Y^{\text{init}} &= \langle e \rangle^{\text{init}}, \langle E \rangle^{\text{init}} \\
\langle e \text{ where rec } E \rangle_Y^{\text{step}} ((m, M), [r_e : r_E]) &= \text{let } F(\rho) = \left( \text{let } M', \rho, w = \langle E \rangle_{Y+\rho}^{\text{step}} (M, r_E) \text{ in } \rho \right) \text{ in} \\
&\quad \text{let } \rho = \text{fix}(F) \text{ in} \\
&\quad \text{let } M', \rho, W = \langle E \rangle_{Y+\rho}^{\text{step}} (M, r_E) \text{ in} \\
&\quad \text{let } m', v, w = \langle e \rangle_{Y+\rho}^{\text{step}} (m, r_e) \text{ in} \\
&\quad \quad (m', M'), v, w \times W
\end{aligned}$$

Fig. 17. Operational semantics for probabilistic expressions.

PROPOSITION A.1. *For an expression where all equations are scheduled, the scheduled semantics is equal to the operational semantics with a fixpoint.*

PROOF. This result is a consequence of the following lemma:

LEMMA A.2. *For all scheduled equations set  $E$ , the scheduled semantics yields the same environment as the fixpoint operator, i.e., for an environment  $\gamma$ , a state  $M$  and an array of random elements  $r$ :*

$$\text{fix}(\lambda \rho. \text{let } M', \rho', w = \langle E \rangle_{Y+\rho}^{\text{step}} (M, r) \text{ in } \rho') = \text{let } M', \rho, w = \langle E \rangle_Y^{\text{step}} (M, r) \text{ in } \rho$$

$$\begin{aligned}
(x = e)^{\text{init}} &= (e)^{\text{init}} \\
(x = e)_\gamma^{\text{step}} (m, r) &= \text{let } m', v, w = (e)_\gamma^{\text{step}} (m, r) \text{ in } m', [x \leftarrow v], w \\
(\text{init } x = e)^{\text{init}} &= \text{nil}, (e)^{\text{init}} \\
(\text{init } x = e)_\gamma^{\text{step}} ((\text{nil}, m_0), r) &= \text{let } m', i, w = (e)_\gamma^{\text{step}} (m_0, r) \text{ in } (\gamma(x), m_0), [x.\text{last} \leftarrow i], w \\
(\text{init } x = e)_\gamma^{\text{step}} ((v, m_0), r) &= (\gamma(x), m_0), [x.\text{last} \leftarrow v], 1 \\
(E_1 \text{ and } E_2)^{\text{init}} &= (E_1)^{\text{init}}, (E_2)^{\text{init}} \\
(E_1 \text{ and } E_2)_\gamma^{\text{step}} ((M_1, M_2), [r_1 : r_2]) &= \text{let } M'_1, \rho_1, w_1 = (E_1)_\gamma^{\text{step}} (M_1, r_1) \text{ in} \\
&\quad \text{let } M'_2, \rho_2, w_2 = (E_2)_\gamma^{\text{step}} (M_2, r_2) \text{ in} \\
&\quad (M'_1, M'_2), \rho_1 + \rho_2, w_1 \times w_2
\end{aligned}$$

Fig. 18. Operational semantics for equations.

The proof is by induction on  $E$ . It is sufficient to focus on the case  $E_1 \text{ and } E_2$ . Since equations are scheduled,  $E_1$  does not depend on variables defined in  $E_2$  and we have for an environment  $\gamma$ , a state  $(M_1, M_2)$  and an array of random elements  $[r_1 : r_2]$ :

$$\begin{aligned}
&\text{fix } (\lambda(\rho_1 + \rho_2). \text{let } M'_1, \rho'_1, w_1 = (E_1)_{\gamma+\rho_1+\rho_2}^{\text{step}} (M_1, r_1) \text{ in} \\
&\quad \text{let } M'_2, \rho'_2, w_2 = (E_2)_{\gamma+\rho_1+\rho_2}^{\text{step}} (M_2, r_2) \text{ in } \rho'_1 + \rho'_2) \\
&= \text{fix } (\lambda(\rho_1 + \rho_2). \text{let } M'_1, \rho'_1, w_1 = (E_1)_{\gamma+\rho_1}^{\text{step}} (M_1, r_1) \text{ in} \\
&\quad \text{let } M'_2, \rho'_2, w_2 = (E_2)_{\gamma+\rho_1+\rho_2}^{\text{step}} (M_2, r_2) \text{ in } \rho'_1 + \rho'_2) \\
&= \text{let } \rho''_1 = \text{fix}(\lambda \rho_1. \text{let } M'_1, \rho'_1, w_1 = (E_1)_{\gamma+\rho_1}^{\text{step}} (M_1, r_1) \text{ in } \rho'_1) \text{ in} \\
&\quad \text{let } \rho''_2 = \text{fix}(\lambda \rho_2. \text{let } M'_2, \rho'_2, w_2 = (E_2)_{\gamma+\rho''_1+\rho_2}^{\text{step}} (M_2, r_2) \text{ in } \rho''_1 + \rho''_2) \\
&= \text{let } M'_1, \rho'_1, w_1 = (E_1)_\gamma^{\text{step}} (M_1, r_1) \text{ in} \\
&\quad \text{let } M'_2, \rho'_2, w_2 = (E_2)_{\gamma+\rho'_1}^{\text{step}} (M_2, r_2) \text{ in } \rho'_1 + \rho'_2
\end{aligned}$$

□

## A.2 Denotational semantics

*Stream functions.* The denotational semantics is presented in Section 5. The definition of this semantics relies on a few stream functions presented in Figure 19. The function  $tl$  drops the first element of a stream ( $tl^n$  drops the  $n$  first elements). The function  $map f s$  applies  $f$  to each element of the stream  $s$ .  $merge cs as bs$  merges the streams  $as$  and  $bs$  according to the condition  $cs$ .  $as$  when  $cs$  keeps the values of  $as$  only when the condition  $cs$  is true. The function  $slicer ss cs$  is used to define the semantics of  $\text{reset } e_1 \text{ every } e_2$ . The first argument  $ss$  is a stream of streams where each stream represents  $e_1$  restarted at each time step, and  $cs$  the the reset condition. When the condition is false, the first value of the first stream of  $ss$  is returned and the second stream of  $ss$  is discarded. We progress by one step in  $e_1$  and the stream representing  $e_1$  restarted at the current iteration is not useful since the expression was not reset. When the condition is true, the first stream of  $ss$  which represents the current state of  $e_1$  is discarded and the execution restarts with the first value of the second stream of  $ss$  which represents  $e_1$  restarted at the current step.

Manuscript submitted to ACM

$$\begin{aligned}
\begin{array}{rcl}
\mathit{tl} & : & A^\omega \rightarrow A^\omega \\
\mathit{tl} (a \cdot as) & = & as
\end{array} \\
\\
\begin{array}{rcl}
\mathit{map} & : & (A \rightarrow B) \rightarrow (A^\omega \rightarrow B^\omega) \\
\mathit{map} f (a \cdot as) & = & f(a) \cdot (\mathit{map} f as)
\end{array} \\
\\
\begin{array}{rcl}
\mathit{merge} & : & \mathbb{B}^\omega \rightarrow A^\omega \rightarrow A^\omega \rightarrow A^\omega \\
\mathit{merge} (T \cdot cs) (a \cdot as) bs & = & a \cdot (\mathit{merge} cs as bs) \\
\mathit{merge} (F \cdot cs) as (b \cdot bs) & = & b \cdot (\mathit{merge} cs as bs)
\end{array} \\
\\
\begin{array}{rcl}
\mathit{when} & : & \mathbb{A}^\omega \rightarrow \mathbb{B}^\omega \rightarrow A^\omega \\
(a \cdot as) \mathit{when} (T \cdot cs) & = & a \cdot (as \mathit{when} cs) \\
(a \cdot as) \mathit{when} (F \cdot cs) & = & as \mathit{when} cs
\end{array} \\
\\
\begin{array}{rcl}
\mathit{slicer} & : & (A^\omega)^\omega \rightarrow \mathbb{B}^\omega \rightarrow A^\omega \\
\mathit{slicer} ((a \cdot as) \cdot bs \cdot ss) (F \cdot cs) & = & a \cdot (\mathit{slicer} (as \cdot ss) cs) \\
\mathit{slicer} (as \cdot (b \cdot bs) \cdot ss) (T \cdot cs) & = & b \cdot (\mathit{slicer} (bs \cdot ss) cs)
\end{array}
\end{aligned}$$

Fig. 19. Stream functions for the denotational semantics.

$$\begin{array}{c}
\begin{array}{c}
\begin{array}{c}
H \vdash c \downarrow c \quad H \vdash x \downarrow H(x) \quad \frac{H \vdash e_1 \downarrow s_1 \quad H \vdash e_2 \downarrow s_2}{H \vdash (e_1, e_2) \downarrow (s_1, s_2)} \quad \frac{H \vdash e \downarrow s}{H \vdash op(e) \downarrow op(s)} \quad \frac{H(x.\text{last}) = s}{H \vdash \text{last } x \downarrow s}
\end{array} \\
\\
\frac{H \vdash e \downarrow s_e \quad \text{node } f x = e_f \quad [x \leftarrow s_e] \vdash e_f \downarrow s}{H \vdash f(e) \downarrow s} \quad \frac{H + H_E \vdash E \quad H + H_E \vdash e \downarrow s}{H \vdash e \text{ where rec } E \downarrow s}
\end{array} \\
\\
\frac{H \vdash e \downarrow s_c \quad (H \text{ when } s_c) \vdash e_1 \downarrow s_1 \quad (H \text{ when not } s_c) \vdash e_2 \downarrow s_2}{H \vdash \text{present } e \rightarrow e_1 \text{ else } e_2 \downarrow \mathit{merge} s_c s_1 s_2}
\end{array} \\
\\
\frac{[(\mathit{tl}^n H) \vdash e_1 \downarrow s_n]_{n \in \mathbb{N}} \quad H \vdash e_2 \downarrow s_c}{H \vdash \text{reset } e_1 \text{ every } e_2 \downarrow \mathit{slicer} (s_0 \cdot s_0 \cdot s_1 \cdot s_2 \cdot \dots) s_c} \quad \frac{H \vdash e \downarrow H(x)}{H \vdash x = e} \quad \frac{H \vdash e \downarrow v_i \cdot s_i \quad H(x.\text{last}) = v_i \cdot H(x)}{H \vdash \text{init } x = e}
\end{array} \\
\\
\frac{H \vdash E_1 \quad H \vdash E_2}{H \vdash E_1 \text{ and } E_2}$$

Fig. 20. Deterministic denotational semantics.

*Environment.* An *environment*  $H$  is a map from variable names to streams of values, for any bound variable  $x \in \text{dom}(H), H(x) : A^\omega$ . When the context is clear, we write  $f H$  for  $\mathit{map} f H$ , e.g., for all  $x \in \text{dom}(H)$ ,  $(\mathit{tl} H)(x) = \mathit{tl}(H(x))$ .

*Denotational semantics.* The full deterministic and probabilistic denotational semantics including the rules for *present* and *reset* are given in Figures 20 and 21. The semantics of the control structure *present*  $e \rightarrow e_1 \text{ else } e_2$  uses the *when* function on the environment  $H$  such that the execution of  $e_1$  and  $e_2$  respectively progress only when the condition is true or false. Then the value of these two streams are merged using the *merge* function. The semantics of *reset*  $e_1 \text{ every } e_2$  is based on the *slicer* function.  $[(\mathit{tl}^n H) \vdash e_1 \downarrow s_n]_{n \in \mathbb{N}}$  represents the stream of streams where  $s_n$

Manuscript submitted to ACM

$$\begin{array}{c}
H, [] \vdash c \Downarrow (c, 1) \quad H, [] \vdash x \Downarrow (H(x), 1) \quad \frac{H, R_e \vdash e \Downarrow (s_e, w_e) \quad \text{proba } f x = e_f \quad [x \leftarrow s_e], R_f \vdash e_f \Downarrow (s, w)}{H, [R_e : R_f] \vdash f(e) \Downarrow (s, w \times w_e)} \\
\\
\frac{H + H_E, R_E \vdash E : w_E \quad H + H_E, R_E \vdash e \Downarrow (s, w)}{H, [R_E : R_E] \vdash e \text{ where rec } E \Downarrow (s, w \times w_E)} \\
\\
\frac{H, R_e \vdash e \Downarrow s_c \quad (H, R_1 \text{ when } s_c) \vdash e_1 \Downarrow sw_1 \quad (H, R_2 \text{ when not } s_c) \vdash e_2 \Downarrow sw_2}{H, [R_1 : R_2] \vdash \text{present } e \rightarrow e_1 \text{ else } e_2 \Downarrow \text{merge } s_c sw_1 sw_2} \\
\\
\frac{[(tl^n H, R_1) \vdash e_1 \Downarrow sw_n]_{n \in \mathbb{N}} \quad H, R_2 \vdash e_2 \Downarrow s_c \quad (s, w) = \text{slicer } (sw_0 \cdot sw_0 \cdot sw_1 \cdot sw_2 \cdot \dots) s_c}{H, [R_1 : R_2] \vdash \text{reset } e_1 \text{ every } e_2 \Downarrow (s, w)} \\
\\
\frac{H \vdash e \Downarrow s_\mu}{H, [R] \vdash \text{sample}(e) \Downarrow (icdf_{s_\mu}(R), 1)} \quad \frac{H \vdash e \Downarrow w}{H, [] \vdash \text{factor}(e) \Downarrow (((), w))} \quad \frac{H, R \vdash e \Downarrow (H(x), w)}{H, R \vdash x = e : w} \\
\\
\frac{H, R \vdash e \Downarrow (i \cdot s, w_i \cdot w) \quad H(x.\text{last}) = i \cdot H(x)}{H, R \vdash \text{init } x = e : w_i \cdot 1} \quad \frac{H, R_1 \vdash E_1 : w_1 \quad H, R_2 \vdash E_2 : w_2}{H, [R_1 : R_2] \vdash E_1 \text{ and } E_2 : w_1 \times w_2} \\
\\
\frac{[H, R \vdash e \Downarrow (s, w) \quad \bar{w} = \prod w]_{R \in ([0,1]^\omega)^{|e|}}}{H \vdash \text{infer}(e) \downarrow \text{integ}_{|e|} \bar{w} s}
\end{array}$$

Fig. 21. Probabilistic denotational semantics.

is the stream of values computed by  $e_1$  restarted at time step  $n$ . In the slicer, the stream  $s_0$  is duplicated because  $\text{reset } e_1 \text{ every } e_2$  returns the same value whether or not  $e_2$  is true at the initial step.

## B APPLICATION: ASSUMED PARAMETERS FILTERING

### B.1 Algorithm

The inference methods proposed by ProbZelus [2, 5, 6] belong to the family of SMC algorithms. These methods rely on a set of independent simulations, called *particles*. Each particle returns an output value associated with a score. The score represents the quality of the simulation. A large number of particles makes it possible to approximate the desired distribution.

More concretely, the `sample(d)` construct randomly draws a value from the  $d$  distribution, and the `factor(x)` construct multiplies the current score of the particle by  $x$ . At each instant, the `infer` operator accumulates the values calculated by each particle weighted by their scores to approximate the posterior distribution.

If the model calls on the operator `sample` at each instant, for example to estimate the position of the boat in the radar example (Section 2), the previous method implements a simple random walk for each particle. As time progresses, it becomes increasingly unlikely that one of the random walks will coincide with the stream of observations. The score associated with each particle quickly goes down towards 0.

To solve this issue, sequential Monte Carlo methods (SMC) add a filtering step. Algorithm 1 describes the execution of one instant for a particle filter, the most basic SMC algorithm. At each instant  $t$ , a particle  $1 \leq i \leq N$  corresponds to a possible value of the parameters (*i.e.*, random variables)  $p_t^i$  of the model. We begin by sampling a new set of particles

Manuscript submitted to ACM

**Data:** probabilistic model `model`, observation  $y_t$ , and previous result  $\mu_{t-1}$ .  
**Result:**  $\mu_t$  an approximation of the distribution of  $p_t$ .

```

for each particle  $i = 1$  to  $N$  do
   $p_{t-1}^i = \text{sample}()$ 
   $p_t^i, w_t^i = \text{model}(y_t | p_{t-1}^i)$ 
   $\mu_t = \mathcal{M}(\{w_t^i, p_t^i\}_{1 \leq i \leq N})$ 
return  $\mu_t$ 

```

**Algorithm 1:** Particle Filter.

**Data:** probabilistic model `model`, observation  $y_t$ , and previous result  $\mu_{t-1}$ .  
**Result:**  $\mu_t$  an approximation of the distributions of state parameter  $x_t$  and constant parameter  $\theta$ .

```

for each particle  $i = 1$  to  $N$  do
   $x_{t-1}^i, \Theta_{t-1}^i = \text{sample}()$ 
   $\theta^i = \text{sample}()$ 
   $x_t^i, w_t^i = \text{model}(y_t | \theta^i, x_{t-1}^i)$ 
   $\Theta_t^i = \text{Update}(\Theta_{t-1}^i, \lambda\theta, \text{model}(y_t | \theta, x_{t-1}^i, x_t^i))$ 
   $\mu_t = \mathcal{M}(\{w_t^i, (x_t^i, \Theta_t^i)\}_{1 \leq i \leq N})$ 
return  $\mu_t$ 

```

**Algorithm 2:** Assumed Parameter Filter [35].

in the distribution obtained at the previous step. The most probable particles are thus duplicated and the less probable ones are eliminated. This refocuses the inference around the most significant information while maintaining the same number of particles throughout the execution. Knowing the previous state  $p_{t-1}^i$ , each particle then executes a step of the model to obtain a sample of the parameters  $p_t^i$  associated with a score  $w_t^i$ . At the end of the instant, we construct a distribution  $\mu_t$  where each particle is associated with its score.  $\mathcal{M}(\{w_t^i, p_t^i\}_{1 \leq i \leq N})$  is a multinomial distribution, where the value  $p_t^i$  is associated with the probability  $w_t^i / \sum_{i=1}^N w_t^i$ .

Unfortunately, this approach generates a loss of information for the estimation of constant parameters. On our radar example, at the first instant, each particle draws a random value for the parameter `theta`. At each instant, the duplicated particles share the same value for `theta`. The quantity of information useful for estimating `theta` therefore decreases with each new filtering and, after a certain time, only one possible value remains.

Rather than sampling at the start of execution a set of values for the constant parameters that will impoverish with each filtering, in the APF algorithm, each particle computes a symbolic distribution of constant parameters. At runtime, the inference then alternates between a sampling pass to estimate the state parameters, and the update of the constant parameters. This avoids impoverishment for the estimation of the constant parameters.

More formally, Algorithm 2 describes the execution of one step of APF. At each instant  $t$ , a particle  $1 \leq i \leq N$  corresponds to a possible value of the state parameters  $x_t^i$  and a distribution of constant parameters  $\Theta_t^i$ . As for the particle filter, we begin by sampling a set of particles in the distribution obtained at the previous instant. We then sample a value  $\theta^i$  in  $\Theta_{t-1}^i$ . Knowing the value of the constant parameters  $\theta^i$  and the previous state  $x_{t-1}^i$ , we can execute a step of the model to obtain a sample of the state parameters  $x_t^i$  associated with a score  $w_t^i$ . We can then update  $\Theta_t^i$  by exploring the other possible values for  $\theta$  knowing that the particle has chosen the transition  $x_{t-1}^i \rightarrow x_t^i$ . At the end of the instant, we construct a distribution  $\mu_t$  where each particle is associated with its score.  $\mathcal{M}(\{w_t^i, (x_t^i, \Theta_t^i)\}_{1 \leq i \leq N})$  is a multinomial distribution where the pair of values  $(x_t^i, \Theta_t^i)$  is associated to the probability  $w_t^i / \sum_{i=1}^N w_t^i$ .

Manuscript submitted to ACM

$$\begin{array}{c}
\frac{C \vdash^c e}{\Phi, C \vdash \text{let } x = e : \Phi, C + \{x\}} \quad \frac{\Phi, C \vdash \text{node } f x = e : \Phi + [f \leftarrow \emptyset], C}{\Phi, C \vdash \text{proba } f x = e : \Phi + [f \leftarrow \phi], C} \\
\frac{\Phi, C \vdash d_1 : \Phi_1, C_1 \quad \Phi_1, C_1 \vdash d_2 : \Phi', C'}{\Phi, C \vdash d_1 d_2 : \Phi', C'} \\
\frac{}{C \vdash c : \emptyset} \quad \frac{}{C \vdash x : \emptyset} \quad \frac{C \vdash e_1 : \phi_1 \quad C \vdash e_2 : \phi_2}{C \vdash (e_1, e_2) : \phi_1 + \phi_2} \quad \frac{C \vdash e : \phi}{C \vdash \text{op}(e) : \phi} \quad \frac{C \vdash e : \phi}{C \vdash \text{sample}(e) : \phi} \\
\frac{C \vdash e : \phi}{C \vdash \text{factor}(e) : \phi} \quad \frac{}{C \vdash \text{last } x : \emptyset} \quad \frac{C \vdash e_1 : \phi}{C \vdash \text{present } e_1 \rightarrow e_2 \text{ else } e_3 : \phi} \quad \frac{C \vdash e_2 : \phi}{C \vdash \text{reset } e_1 \text{ every } e_2 : \phi} \\
\frac{}{C \vdash f_\theta(e) : [\theta \leftarrow f.\text{prior}]} \quad \frac{C \vdash e : \phi_e \quad C \vdash^c E : D \quad C, D \vdash E : \phi_E}{C \vdash e \text{ where rec } E : \phi_e + \phi_E} \quad \frac{x \in D \quad C \vdash^c e}{C, D \vdash \text{init } x = \text{sample}(e) : [x \leftarrow e]} \\
\frac{C \vdash e : \phi}{C, D \vdash \text{init } x = e : \phi} \quad \frac{C \vdash e : \phi}{C, D \vdash x = e : \phi} \quad \frac{C, D \vdash E_1 : \phi_1 \quad C, D \vdash E_2 : \phi_2}{C, D \vdash E_1 \text{ and } E_2 : \phi_1 + \phi_2}
\end{array}$$

Fig. 22. Extract constant parameters and associated prior distributions.

$$\begin{array}{c}
\frac{}{C \vdash^c c} \quad \frac{x \in C}{C \vdash^c x} \quad \frac{C \vdash^c e_1 \quad C \vdash^c e_2}{C \vdash^c (e_1, e_2)} \quad \frac{C \vdash^c e}{C \vdash^c \text{op}(e)} \quad \frac{C + \text{dom}(E) \vdash^c e \quad C \vdash^c E : \text{dom}(E)}{C \vdash^c e \text{ where rec } E} \\
\frac{}{C \vdash^c \text{init } x = e : \emptyset} \quad \frac{}{C \vdash^c x = \text{last } x : \{x\}} \quad \frac{C \vdash^c e}{C \vdash^c x = e : \{x\}} \quad \frac{C \vdash^c e}{C \vdash^c x = e : \emptyset} \\
\frac{C + C_2 \vdash^c E_1 : C_1 \quad C + C_1 \vdash^c E_2 : C_2}{C \vdash^c E_1 \text{ and } E_2 : C_1 + C_2}
\end{array}$$

Fig. 23. Constant expressions and equations.

## B.2 Static Analysis

The full type system is given in Figures 22 and 23. The interesting cases are presented in Section 7.1.

## B.3 Compilation

The complete compilation function to transform a ProbZelus model into a model compatible with `APF.infer` is given in Figure 24. Most cases simply call the compilation functions on all sub-expressions. The interesting cases are presented in Section 7.2.

$C_\phi(c)$	$= c$
$C_\phi(x)$	$= x$
$C_\phi((e_1, e_2))$	$= (C_\phi(e_1), C_\phi(e_2))$
$C_\phi(op(e))$	$= op(C_\phi(e))$
$C_\phi(\text{last } x)$	$= \text{last } x$
$C_\phi(\text{present } e_1 \rightarrow e_2 \text{ else } e_3)$	$= \text{present } C_\phi(e_1) \rightarrow C_\phi(e_2) \text{ else } C_\phi(e_3)$
$C_\phi(\text{reset } e_1 \text{ every } e_2)$	$= \text{reset } C_\phi(e_1) \text{ every } C_\phi(e_2)$
$C_\phi(\text{sample}(e))$	$= \text{sample}(C_\phi(e))$
$C_\phi(\text{factor}(e))$	$= \text{factor}(C_\phi(e))$
$C_\phi(e \text{ where rec } E)$	$= C_\phi(e) \text{ where rec } C_\phi(E)$
$C_\phi(\text{init } x = e)$	$= \begin{cases} \emptyset & \text{if } x \in \text{dom}(\phi) \\ \text{init } x = C_\phi(e) & \text{otherwise} \end{cases}$
$C_\phi(x = e)$	$= \begin{cases} \emptyset & \text{if } x \in \text{dom}(\phi) \\ x = C_\phi(e) & \text{otherwise} \end{cases}$
$C_\phi(f_\theta(e))$	$= \begin{cases} f(C_\phi(e)) & \text{if } f \text{ is deterministic or } \phi = [] \\ f.\text{model}(\theta, C_\phi(e)) & \text{if } \theta \in \text{dom}(\phi) \\ f.\text{model}(\theta, C_\phi(e)) \text{ where} \\ \quad \text{rec init } \theta = \text{sample}(f.\text{prior}) \\ \quad \text{and } \theta = \text{last } \theta & \text{otherwise} \end{cases}$
$C_\phi(\text{infer}(f(e)))$	$= \text{APF.infer}(f.\text{model}, f.\text{prior}, C_\phi(e))$
$C_\phi(E_1 \text{ and } E_2)$	$= C_\phi(E_1) \text{ and } C_\phi(E_2)$
$C_\Phi(\text{let } x = e)$	$= \text{let } x = e$
$C_\Phi(\text{node } f x = e)$	$= \text{node } f x = C_{[]} (e)$
$C_\Phi(\text{proba } f x = e)$	$= \begin{cases} \text{proba } f x = C_{[]} (e) & \text{if } \Phi(f) = [] \\ \text{let } f.\text{prior} = im(\phi) & \text{with } \phi = \Phi(f) \\ \text{proba } f.\text{model}(\text{dom}(\phi), x) = C_\phi(e) & \end{cases}$

Fig. 24. APF Compilation.