



HAL
open science

Density-Based Semantics for Reactive Probabilistic Programming

Guillaume Baudart, Louis Mandel, Christine Tasson

► **To cite this version:**

Guillaume Baudart, Louis Mandel, Christine Tasson. Density-Based Semantics for Reactive Probabilistic Programming. 2024. hal-04488216

HAL Id: hal-04488216

<https://hal.science/hal-04488216v1>

Preprint submitted on 4 Mar 2024

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

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

Density-Based Semantics for Reactive Probabilistic Programming

GUILLAUME BAUDART, ENS – PSL University – CNRS – Inria, France

LOUIS MANDEL, IBM Research, USA

CHRISTINE TASSON, ISAE Supaero, 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 extended with Bayesian reasoning to program state-space models which compute a stream of distributions given a stream of observations. However, the semantics of probabilistic models is only defined for scheduled equations – a significant limitation compared to dataflow synchronous languages and block diagrams which do not require any ordering.

In this paper we propose two schedule agnostic semantics for a probabilistic synchronous language. The key idea is to interpret probabilistic expressions as a stream of un-normalized density functions which maps random variable values to a result and positive score. The co-iterative semantics interprets programs as state machines and equations are computed using a fixpoint operator. The relational semantics directly manipulates streams and is thus a better fit to reason about program equivalence. We use the relational semantics to prove the correctness of a program transformation required to run an optimized inference algorithm for state-space models with constant parameters.

1 INTRODUCTION

Synchronous programming languages [Benveniste et al. 2003] were introduced for the design of critical embedded systems. In dataflow languages such as Lustre [Halbwachs et al. 1991], system designers write high-level specifications by composing infinite streams of values, called *flows*. Flows progress *synchronously*, paced on a global logical clock. The expressiveness of synchronous languages is deliberately restricted. Specialized compilers can thus 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. Built on these ideas, Scade [Colaço et al. 2017] is now a standard tool in automotive and avionic industries.

Probabilistic languages [Bingham et al. 2019; Cusumano-Towner et al. 2019; Ge et al. 2018; Goodman and Stuhlmüller 2014] extend general purpose programming languages with probabilistic constructs for Bayesian inference. 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 between these two lines of research, ProbZelus [Baudart et al. 2020] is a probabilistic extension of the synchronous dataflow language Zelus [Bourke and Pouzet 2013]. 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) [Montemerlo et al. 2002] 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) [Chopin and Papaspiliopoulos 2020]. Any ProbZelus program can be expressed as a SSM.

Probabilistic semantics and scheduling. ProbZelus semantics [Baudart et al. 2020] is defined in a co-iterative framework where expressions are interpreted as state machines. Following [Kozen 1981; Staton 2017], 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 can be written in any order. 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 [Bourke and Pouzet 2013]. The semantics of ProbZelus is thus far from what is exposed to the programmer, and more importantly, 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 [Bourke et al. 2017a; Caspi and Pouzet 1998] for probabilistic programming. The key idea is to interpret a probabilistic expression as a stream of un-normalized density functions which map random variable values to a result and a positive score.

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

- We introduce in Section 4 a new density-based co-iterative semantics and show that sets of mutually recursive equations in arbitrary order can be interpreted using a fixpoint operator. We prove that this semantics is equivalent to the original ProbZelus semantics.
- We introduce in Section 5 an alternative relational semantics which abstracts away the state machines and directly manipulates streams which simplifies reasoning about program equivalence. We prove that this semantics is equivalent to the co-iterative semantics.
- We define in Section 6 a program transformation required to run an optimized inference algorithm for state-space models with constant parameters. We use the relational semantics to prove the correctness of the transformation.

2 EXAMPLE

To motivate our approach, consider the ProbZelus model of Figure 1 adapted from [Chopin and Papaspiliopoulos 2020][Section 2.4.1]. The goal is to estimate at each instant the position of a moving boat given noisy observations from a marine radar. 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 from noisy measurements of its angle and the echo delay.

The keyword `proba` indicates the definition of a probabilistic stream function. Line 1, the model tracker takes as input a stream of observations `y_obs` and returns a stream of positions `x`. The model uses a transition function `f` to estimate the current position (e.g., using a linear motion model), and a projection function `g` to compute the observable quantities from the state (e.g., angle and echo delay). Line 3 uses the `sample` operator to specify that `x` is Gaussian distributed around `f(last x)`. `last x` refers to the previous position of the boat initialized Line 2 with the `init` keyword. Line 5 uses the `observe` operators to condition the model assuming that the observations `y_obs` are Gaussian distributed around `y = g(x)`. The initial position `x_init` and the noise parameters `sx` and `sy` are global constants.

2.1 Kernel-based co-iterative semantics

ProbZelus original semantics [Baudart et al. 2020] is a co-iterative semantics where expressions are interpreted as state machines characterized by an initial state and a transition function. Given the current state, the transition function of deterministic expression returns the next state and a value.

```

1 proba tracker(y_obs) = x where
2   rec init x = x_init
3   and x = sample(gaussian(f(last x), s_x))
4   and y = g(x)
5   and () = observe(gaussian(y, s_y), y_obs)

7 node main(y_obs) = msg where
8   rec x_dist = infer (tracker (y_obs))
9   and msg = controller(x_dist)

```




Fig. 1. Tracking a moving boat with a marine radar in ProbZelus. x_t is the position of the boat in Cartesian coordinates. We assume a linear motion model $f(x_{t-1}) = x_{t-1} + \theta$, and $g(x_t) = (\alpha_t, \delta_t)$ returns 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 speed of light.

Following [Staton 2017], the transition function of a probabilistic expression returns a *measure* over all possible pairs (next state, value).

For instance, if f and g are deterministic and stateless, the transition function of `tracker` is the following (omitting empty states for stateless expressions) where \mathcal{N} is the normal distribution, δ the Dirac delta distribution, and pdf_d the probability density function of d .

$$\begin{aligned}
\llbracket \text{tracker}(y_obs) \rrbracket_y^{\text{step}}(p_x) &= \int \mathcal{N}(f(p_x), s_x)(dx) \int \delta_{g(x)}(dy) \text{pdf}_{\mathcal{N}(y, s_x)}(y_{obs}) * \delta_{x,x} \\
&= \int \mathcal{N}(f(p_x), s_x)(dx) \text{pdf}_{\mathcal{N}(g(x), s_x)}(y_{obs}) * \delta_{x,x}
\end{aligned} \tag{1}$$

Given the current state p_x (the previous value of x), the transition function integrates over all possible values for x , then all possible values for y , weights each execution by the likelihood of the observation (i.e., the value of the density function on y_obs), and returns a measure over the new state x (that will be used as the previous position in the next step), and the results x .

Inference. In ProbZelus, the explicit `infer` operator computes the stream of distributions described by a model. The state of `infer` is a measure over current states. The transition function integrates the semantics of the model over all possible states, normalizes the resulting measure to obtain a distribution, and splits it 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.

In ProbZelus, inference can be composed with deterministic control to perform *inference-in-the-loop*. For instance, Line 9, the distribution `x_dist` is exploited by a controller, e.g., to send a message to guide the boat.

Scheduling. Local declarations such as x and y yield nested integrals in Equation (1) that are only well-defined if these definitions are ordered according to data dependencies. The ProbZelus semantics defined in [Baudart et al. 2020] thus focuses on a kernel language where local declarations are all *scheduled*. But imposing a valid schedule is a significant limitation compared to synchronous dataflow languages which manipulate mutually recursive equations in arbitrary order.

2.2 Density-based co-iterative semantics

In this paper we first propose a new density-based co-iterative semantics for ProbZelus inspired by the semantics of the popular probabilistic language Stan [Gorinova et al. 2019] where a program defines an un-normalized density function over the random variables. Instead of manipulating

measures via integration, probabilistic expressions are now similar to deterministic expressions, but the transition function now takes one additional argument — a random seed for all random variables — and returns one additional output — a positive score, or weight, which measures the quality of the output w.r.t. the model.

On the example of Figure 1 we have:

$$\begin{aligned} (\text{tracker}(y_obs))_Y^{\text{step}}(p_x, r) = & \text{let } \mu_x = \mathcal{N}(f(p_x), s_x) \text{ in} \\ & \text{let } x = \text{icdf}_{\mu_x}(r) \text{ in} \\ & \text{let } \mu_y = \mathcal{N}(g(x), s_y) \text{ in} \\ & x, x, \text{pdf}_{\mu_y}(y_{obs}) \end{aligned} \quad (2)$$

The additional argument r corresponds to the random seed for the `sample` operator, i.e., an element of $[0, 1]$ that is mapped to a sample of a distribution d using inverse transform sampling [Devroye 2006]. The model computes the sample x associated to the random seed (icdf_d is the inverse of the cumulative distribution function of d), and returns the new state x , the result x , and a weight which captures the likelihood of the observation (the density of the distribution $\mathcal{N}(g(x), s_y)$ at y_{obs}).

Inference. At each step, the `infer` operator first computes the un-normalized measure which associates each pair (state, result) to its weight, i.e., for a model e , if $(e)_Y(m, r) = m', v, w$, `infer`(e) computes the measure $\int_{[0,1]^p} w * \delta_{m',v} dr$ where p is the number of random variables in e . On the example of Figure 1, we can check that this measure corresponds to the original semantics of Equation (1).

$$\begin{aligned} & \int_{[0,1]} \text{let } x, x, w = (\text{tracker}(y_obs))_Y^{\text{step}}(p_x, r) \text{ in } w * \delta_{x,x} dr \\ &= \int_{[0,1]} \text{let } x = \text{icdf}_{\mathcal{N}(f(p_x), s_x)}(r) \text{ in } \text{pdf}_{\mathcal{N}(g(x), s_y)}(y_{obs}) * \delta_{x,x} dr \\ &= \int \mathcal{N}(f(p_x), s_x)(dx) \text{pdf}_{\mathcal{N}(g(x), s_y)}(y_{obs}) * \delta_{x,x} \end{aligned}$$

The semantics of `infer` is then similar to its interpretation in the original kernel-based semantics, i.e., 1) integrate over all possible state, 2) normalize the measure, 3) split the result into a distribution of next states and a distribution of values. We prove in Section 4.2 that this semantics is equivalent to the kernel-based semantics, i.e., the `infer` operator yields the same stream of distributions.

Mutually recursive equations. The original co-iterative semantics for dataflow synchronous languages [Caspi and Pouzet 1998] interprets mutually recursive equations in arbitrary order with a fixpoint operator in a flat complete partial order (CPO) where variables are either undefined or set to a value. In the density-based semantics, the transition functions of probabilistic equations are similar to their deterministic counterparts with additional inputs/outputs. Compared to the kernel-based semantics, there are no longer nested integrals and a fixpoint operator can be defined to interpret sets of probabilistic equations.

Consider a variant of the example of Figure 1 where we swap Line 3 and Line 4. For a state p_x and a random seed r , the semantics of the local declarations in `tracker` is the fixpoint of the following function F starting from the least element $[x \leftarrow \perp, y \leftarrow \perp]$.

$$\begin{aligned} F(\rho) = & \text{let } y = g(\rho(x)) \text{ in} \\ & \text{let } \mu_x = \mathcal{N}(f(p_x), s_x) \text{ in} \\ & \text{let } x = \text{icdf}_{\mu_x}(r) \text{ in} \\ & [x \leftarrow x, y \leftarrow y] \end{aligned} \quad \begin{aligned} \rho_0 = & [x \leftarrow \perp, y \leftarrow \perp] \\ \rho_1 = & [x \leftarrow \text{icdf}_{\mu_x}(r), y \leftarrow \perp] \\ \rho_2 = & [x \leftarrow \text{icdf}_{\mu_x}(r), y \leftarrow g(\text{icdf}_{\mu_x}(r))] \\ \rho_3 = & [x \leftarrow \text{icdf}_{\mu_x}(r), y \leftarrow g(\text{icdf}_{\mu_x}(r))] \end{aligned}$$

The fixpoint converges after 3 iterations. Using the resulting environment, the semantics of `tracker` then computes the next state, the result, and the weight which, after simplification, yields the same results as Equation (2).

$$\begin{aligned} ((\text{tracker}(y_obs))_Y^{\text{step}}(p_x, r) = & \text{let } \rho = [x \leftarrow \text{icdf}_{\mu_x}(r), y \leftarrow g(\text{icdf}_{\mu_x}(r))] \text{ in} \\ & \text{let } x = \rho(x) \text{ in} \\ & \text{let } \mu_y = \mathcal{N}(\rho(y), s_y) \text{ in} \\ & \rho(x), \rho(x), \text{pdf}_{\mu_y}(y_{obs}) \end{aligned} \quad (3)$$

Program equivalence. Since deterministic expressions are interpreted as state machines, to prove program equivalence one must exhibit a bisimulation [Park 1981], 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. The proof is done by unfolding the definition of the transition function.

Two probabilistic expressions are equivalent if they describe the same stream of measures. Unfortunately, compared to deterministic expressions, the type of probabilistic expressions is asymmetric: given a state, the transition function describes a measure over pairs (next state, value). The corresponding stream of measures is obtained by integrating at each step the transition function over all possible states computed at the previous step. The bisimulation must thus relate measures of states through successive integrations, which is significantly harder to define and to check than in the deterministic case.

2.3 Density-based relational semantics

An alternative is to directly manipulate streams of values. This is the approach used in the *Vélus* project¹ to prove an end-to-end compiler for the dataflow synchronous language *Lustre* [Bourke et al. 2017a, 2020, 2021]. The semantics of a stream function is defined as a relation between input streams and output streams. In *Vélus*, most of the compilation passes are proven correct using this *relational semantics*. The translation to state machines is one of the very last passes and focuses on a normalized, scheduled subset of the language.

In this paper, we extend this relational semantics to probabilistic streams. The key idea is to lift the density-based semantics to streams. Given an environment H mapping variable names to streams of values, and an array R of random streams, the semantics of an expression returns a stream of pairs (value, weight): $H, R \vdash e \Downarrow (v, w)$.

In the example of Figure 1, R is a single stream of independent random elements $R_0 \cdot R_1 \cdot R_2 \cdot \dots$ in $[0, 1]$, we can interpret `tracker` in an environment that contains the observations y_obs :

$$\begin{aligned} [y_obs \leftarrow y_{obs}], R \vdash \text{tracker}(y_obs) \Downarrow & (x_0, w_0) \cdot (x_1, w_1) \cdot (x_2, w_2) \cdot \dots \\ \text{where } \mu_x = & \mathcal{N}(f(x_{\text{init}}), s_x) \cdot \mathcal{N}(f(x_0), s_x) \cdot \mathcal{N}(f(x_1), s_x) \cdot \dots \\ x = & \text{icdf}_{\mu_{x_0}}(R_0) \cdot \text{icdf}_{\mu_{x_1}}(R_1) \cdot \text{icdf}_{\mu_{x_2}}(R_2) \cdot \dots \\ y = & g(x_0) \cdot g(x_1) \cdot g(x_2) \cdot \dots \\ \mu_y = & \mathcal{N}(y_0, s_x) \cdot \mathcal{N}(y_1, s_x) \cdot \mathcal{N}(y_2, s_x) \cdot \dots \\ w = & \text{pdf}_{\mu_{y_0}}(y_{obs0}) \cdot \text{pdf}_{\mu_{y_1}}(y_{obs1}) \cdot \text{pdf}_{\mu_{y_2}}(y_{obs2}) \cdot \dots \end{aligned}$$

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

¹<https://velus.inria.fr>

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

Fig. 2. ProbZelus Syntax.

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 which associates each pair (state, result) to its total weight, and 3) normalizes it to obtain a distribution of values. The key difference with the density-based co-iterative semantics is that the integral is now over the infinite domain of streams. We prove in Section 5.3 that this semantics is equivalent to the co-iterative density-based semantics, i.e., the `infer` operator yields the same stream of distributions.

Mutually recursive equations. Given the random streams R , the semantics of a set of probabilistic equations $H, R \vdash E : W$ checks that an environment mapping variable names to stream of values H is compatible with all the equations in E , and that the combined weight of all sub-expressions is the stream W . Since variables in an environment 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. Of course, compared to the co-iterative semantics, the relational semantics is not executable since equations are only checked a posteriori for a given environment.

Program equivalence. In the relational semantics, states are abstracted away and 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 obtained by integrating at each step the result of the relational semantics over all possible random streams. Since, random streams are uniformly distributed, if we can map the random streams of one expression to the random streams of the other, program equivalence can be reduced to the comparison of the streams of pairs (value, weight) computed by each expression.

3 BACKGROUND

In this section we briefly summarize the key elements of the co-iterative semantics of ProbZelus. Importantly, this semantics is only defined if all equations are ordered according to data dependencies. We then recall the original co-iterative semantics of synchronous dataflow languages where sets of mutually recursive equations in arbitrary order are interpreted using a fixpoint operator.

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 $\text{op}(e)$, 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` $c \rightarrow e_1$ `else` e_2 , or a reset construct `reset` e_1 `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 and E_2 . In a set of equations, every initialized variable must be defined by another equation.

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`($pdf_d(x)$) which conditions the model on the assumption that x was sampled from d .

3.2 Co-iterative semantics

The semantics of ProbZelus presented in [Baudart et al. 2020] extends the co-iterative semantics of dataflow synchronous languages [Caspi and Pouzet 1998; Colaco et al. 2023]. A type system statically identifies deterministic and probabilistic expressions [Baudart et al. 2020, Section 3.2] which have different interpretations.

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

$$(\llbracket e \rrbracket_{\gamma_0}^{\text{init}} = m_0) \xrightarrow{\llbracket e \rrbracket_{\gamma_1}^{\text{step}}} \underset{v_1}{m_1} \xrightarrow{\llbracket e \rrbracket_{\gamma_2}^{\text{step}}} \underset{v_2}{m_2} \xrightarrow{\llbracket e \rrbracket_{\gamma_3}^{\text{step}}} \underset{v_3}{m_3} \rightarrow \dots$$

Following [Staton 2017], the semantics of a probabilistic expression is a state machine which computes a stream of *kernels*. Given the current state, the transition function $\llbracket e \rrbracket_{\gamma}^{\text{step}}$ of type $S \rightarrow \Sigma_{S \times V} \rightarrow [0, \infty)$ returns a measure over pairs (next state, value),² i.e., a function mapping measurable sets of pairs (next state, value) to a positive score.

Figure 3 shows a simplified excerpt of the semantics of probabilistic expressions from [Baudart et al. 2020]. In a probabilistic context, a deterministic expression is interpreted as the Dirac delta measure on the pair (state, value) returned by the deterministic semantics.³ `sample` evaluates its argument into a new state m' and a distribution of values μ , 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 simplify the semantics, the type system ensures that the arguments of the probabilistic operators are always deterministic expressions. To illustrate local declarations, Figure 3 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.

Inference. So far, probabilistic expressions describe a stream of un-normalized measures over pairs (state, value). At each step, the `infer` operator normalizes the measure to obtain a distribution (\top denotes the entire space), that is then split into a distribution of next states, and a distribution of values. The corresponding stream of distributions is obtained by iteratively integrating the transition function along the distribution of states.

$$(\llbracket e \rrbracket_{\gamma_0}^{\text{init}} = \sigma_0) \xrightarrow{\int \sigma_0(dm) \llbracket e \rrbracket_{\gamma_1}^{\text{step}}(m)} \underset{\mu_1}{\sigma_1} \xrightarrow{\int \sigma_1(dm) \llbracket e \rrbracket_{\gamma_2}^{\text{step}}(m)} \underset{\mu_2}{\sigma_2} \xrightarrow{\int \sigma_2(dm) \llbracket e \rrbracket_{\gamma_3}^{\text{step}}(m)} \underset{\mu_3}{\sigma_3} \rightarrow \dots$$

² Σ_A denotes the Borel σ -algebra over values of type A .

³ $\delta_x(U) = 1$ if $x \in U$ and 0 otherwise.

$$\begin{aligned}
\llbracket e \rrbracket_Y^{\text{init}} &= \llbracket e \rrbracket_Y^{\text{init}} \\
\llbracket e \rrbracket_Y^{\text{step}}(m) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } \delta_{m', v} \quad \text{if } e \text{ is deterministic} \\
\llbracket \text{sample}(e) \rrbracket_Y^{\text{init}} &= \llbracket e \rrbracket_Y^{\text{init}} \\
\llbracket \text{sample}(e) \rrbracket_Y^{\text{step}}(m) &= \text{let } m', \mu = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } \int \mu(dv) \delta_{m', v} \\
\llbracket \text{factor}(e) \rrbracket_Y^{\text{init}} &= \llbracket e \rrbracket_Y^{\text{init}} \\
\llbracket \text{factor}(e) \rrbracket_Y^{\text{step}}(m) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } v * \delta_{m', ()} \\
\left\{ \begin{array}{l} e \text{ where } \text{rec init } x = c \\ \quad \text{and } x = e_x \\ \quad \text{and } y = e_y \end{array} \right\}_Y^{\text{init}} &= c, \left(\llbracket e \rrbracket_Y^{\text{init}}, \llbracket e_x \rrbracket_Y^{\text{init}}, \llbracket e_y \rrbracket_Y^{\text{init}} \right) \\
\left\{ \begin{array}{l} e \text{ where } \text{rec init } x = c \\ \quad \text{and } x = e_x \\ \quad \text{and } y = e_y \end{array} \right\}_Y^{\text{step}}(p_x, (m, m_x, m_y)) &= \int \llbracket e_x \rrbracket_{Y+[x.\text{last} \leftarrow p_x]}^{\text{step}}(m_x)(dm'_x, dv_x) \\
&\quad \int \llbracket e_y \rrbracket_{Y+[x.\text{last} \leftarrow p_x, x \leftarrow v_x]}^{\text{step}}(m_y)(dm'_y, dv_y) \\
&\quad \int \llbracket e \rrbracket_{Y+[x.\text{last} \leftarrow p_x, x \leftarrow v_x, y \leftarrow v_y]}^{\text{step}}(m)(dm', dv) \\
&\quad \delta_{(v_x, (m', m'_x, m'_y)), v} \\
\llbracket \text{infer}(e) \rrbracket_Y^{\text{init}} &= \llbracket e \rrbracket_Y^{\text{init}} \\
\llbracket \text{infer}(e) \rrbracket_Y^{\text{step}}(\sigma) &= \text{let } v = \int \sigma(dm) \llbracket e \rrbracket_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}$$

Fig. 3. A simplified excerpt of the original ProbZelus co-iterative probabilistic semantics [Baudart et al. 2020].

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

3.3 Equations and fixpoints

In the interpretation of local declarations in Figure 3, the nested integrals are only well defined if equations are ordered according to data dependencies. The original semantics in [Baudart et al. 2020] thus focuses on a kernel language where local declarations are all *scheduled*: initializations are grouped at the beginning and an equation $y = e_y$ must appear after $x = e_x$ if x appears in e_y outside a *last*. In the compiler, a specialized type system, the *causality analysis* statically checks that a program is *causal*, i.e., that all local declarations can be scheduled [Cuoq and Pouzet 2001]. The kernel-based semantics is commutative, i.e., yields the same results for any valid schedule [Staton 2017], but imposing a scheduled order on equations is a significant limitation compared to block diagrams or synchronous dataflow languages which manipulate set of equations in arbitrary order.

The original co-iterative semantics [Caspi and Pouzet 1998] and more recent works [Colaco et al. 2023] interpret mutually recursive equations using a fixpoint operator. Values $v \in V$ are interpreted in a flat domain $V^\perp = V + \{\perp\}$ with minimal element \perp 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 is lifted to environments defining the

$$\begin{aligned}
\llbracket x = e \rrbracket_Y^{\text{step}}(m) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', [x \leftarrow v] \\
\llbracket \text{init } x = c \rrbracket_Y^{\text{step}}(p_x) &= \gamma(x), [x.\text{last} \leftarrow p_x] \\
\llbracket E_1 \text{ and } E_2 \rrbracket_Y^{\text{step}}(M_1, M_2) &= \text{let } M'_1, \gamma_1 = \llbracket E_1 \rrbracket_Y^{\text{step}}(M_1) \text{ in} \\
&\quad \text{let } M'_2, \gamma_2 = \llbracket E_2 \rrbracket_Y^{\text{step}}(M_2) \text{ in} \\
&\quad (M'_1, M'_2), \gamma_1 + \gamma_2 \\
\llbracket e \text{ where rec } E \rrbracket_Y^{\text{step}}(m, M) &= \text{let } F(\rho) = \left(\text{let } M', \rho' = \llbracket E \rrbracket_{\gamma+\rho}^{\text{step}}(M) \text{ in } \rho' \right) \text{ in} \\
&\quad \text{let } \rho = \text{fix}(F) \text{ in} \\
&\quad \text{let } M', \rho = \llbracket E \rrbracket_{\gamma+\rho}^{\text{step}}(M) \text{ in} \\
&\quad \text{let } m', v = \llbracket e \rrbracket_{\gamma+\rho}^{\text{step}}(m) \text{ in} \\
&\quad (m', M'), v
\end{aligned}$$

Fig. 4. Co-iterative semantics of deterministic equations with a fixpoint operator (based on [Pouzet 2021]).

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 $\perp = [x \leftarrow \perp]_{x \in X}$.

Figure 4 shows the semantics rules for deterministic equations adapted from [Colaco et al. 2023]. The initial state of an equation is the initial state of its defining expression. Given a state, the transition function returns a new state and an environment. 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 $F(\rho) = \text{let } M', \rho' = \llbracket e \rrbracket_{\gamma+\rho}^{\text{step}}(M) \text{ in } \rho'$ is continuous and has a minimal fixpoint $\rho = \text{fix}(F) = \lim_{n \rightarrow \infty} (F^n(\perp))$. After convergence, the transition function evaluates $\llbracket E \rrbracket_{\gamma+\rho}^{\text{step}}(M)$ once more to compute the next state *M'* (leaving ρ unchanged by definition of the fixpoint) and finally evaluates the main expression *e* in the environment $\gamma + \rho$.

If the program is *causal* a valid schedule exists for *E*, and by monotony, each fixpoint iteration computes the value of at least one variable and the fixpoint is reached after a finite number of iterations [Colaco et al. 2023].

4 DENSITY-BASED CO-ITERATIVE SEMANTICS

In this section we detail the new density-based co-iterative semantics for probabilistic expressions. We show that, in this semantics we can now interpret sets of mutually recursive equations with a fixpoint operator as in the original co-iterative semantics. We then prove that the density-based semantics is equivalent to the kernel-based semantics, i.e., describes the same stream of distributions.

4.1 Probabilistic co-iterative semantics with fixpoint

The key idea of the density-based 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 seeds containing 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.

Expressions. More formally, the initialization function of a probabilistic expression *e*, $(e)_Y^{\text{init}} : S \times \mathbb{N}$ returns the initial state and the number of random variables in *e*. Loops and recursive calls are not allowed in the language of Figure 2. The number of calls to *sample* can thus be statically computed.

$$\begin{aligned}
(e)_Y^{\text{init}} &= \llbracket e \rrbracket_Y^{\text{init}}, 0 \\
(e)_Y^{\text{step}}(m, []) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', v, 1 \quad \text{if } e \text{ is deterministic} \\
(\text{sample}(e))_Y^{\text{init}} &= \text{let } m = \llbracket e \rrbracket_Y^{\text{init}} \text{ in } m, 1 \\
(\text{sample}(e))_Y^{\text{step}}(m, [r]) &= \text{let } m', \mu = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', \text{icdf}_\mu(r), 1 \\
(\text{factor}(e))_Y^{\text{init}} &= \text{let } m = \llbracket e \rrbracket_Y^{\text{init}} \text{ in } m, 0 \\
(\text{factor}(e))_Y^{\text{step}}(m, []) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', (), v \\
(f(e))_Y^{\text{init}} &= \text{let } m_f, p_f = \gamma(f.\text{init}) \text{ in} \\
&\quad \text{let } m_e, p_e = (e)_Y^{\text{init}} \text{ in} \\
&\quad (m_f, m_e), p_f + p_e \\
(f(e))_Y^{\text{step}}((m_f, m_e), [r_f : r_e]) &= \text{let } m'_e, v_e, w_e = (e)_Y^{\text{step}}(m_e, r_e) \text{ in} \\
&\quad \text{let } m'_f, v, w_f = \gamma(f.\text{step})(v_e, m_f, r_f) \text{ in} \\
&\quad (m'_f, m'_e), v, w_e * w_f
\end{aligned}$$

Fig. 5. Density-based co-iterative semantics for ProbZelus expressions (full version in Figure 12 of the appendix).

Given the current state and a value for all random seeds (an array of p values in $[0, 1]$ where p is the number of random variables computed by initialization function) the transition function $(e)_Y^{\text{step}} : S \times [0, 1]^p \rightarrow S \times V \times [0, \infty)$ returns a triple (next state, result, weight).

An excerpt of the density-based co-iterative semantics is presented in Figure 5. If e is deterministic, there is no random variable and no conditioning. The transition function takes an empty array of random seeds, evaluates the expression, and returns the next state, the value, and a weight of 1. `sample` defines one random variable. The transition function takes an array containing one random seed, evaluates the argument into a distribution, converts the random seed into a sample of the distribution, and returns the next state, the sample, and a weight of 1. `factor` updates the weight. The transition function evaluates its arguments into a real value v , and returns the next state, an empty value $()$, and the score v . The initialization of a function call $f(e)$ evaluates the initialization functions of f and e , combines the initial states and sums the numbers of random variables. The transition function takes an array containing the random seeds for e and f ,⁴ evaluates the argument e into a value v_e and a weight w_e , uses the value to evaluate the transition function of f which returns a result v and a weight w_f , and returns the combined next states, the result, and the total weight $w_e * w_f$.

Mutually recursive equations. The semantics of probabilistic equations is presented in Figure 6. As for probabilistic expressions, the initialization function returns the initial states, and the number of random variables. Given a state and an array of random seeds, the transition function returns a tuple (next state, environment, weight). The equation $x = e$ defines a single variable. The transition function evaluates the defining expression e into a tuple (next state, value, weight), and returns the next state, an environment where x is bound to the value v , and the weight. 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

⁴We note $[r_1 : r_2]$ the concatenation of two arrays.

$$\begin{aligned}
(x = e)_Y^{\text{init}} &= (e)_Y^{\text{init}} \\
(x = e)_Y^{\text{step}}(m, r) &= \text{let } m', v, w = (e)_Y^{\text{step}}(m, r) \text{ in } m', [x \leftarrow v], w \\
(\text{init } x = e)_Y^{\text{init}} &= \text{let } m_0, p = (e)_Y^{\text{init}} \text{ in } (\text{nil}, m_0), p \\
(\text{init } x = e)_Y^{\text{step}}((\text{nil}, m_0), r) &= \text{let } m', i, w = (e)_Y^{\text{step}}(m_0, r) \text{ in } (\gamma(x), m_0), [x.\text{last} \leftarrow i], w \\
(\text{init } x = e)_Y^{\text{step}}((v, m_0), r) &= (\gamma(x), m_0), [x.\text{last} \leftarrow v], 1 \\
(E_1 \text{ and } E_2)_Y^{\text{init}} &= \text{let } M_1, p_1 = (E_1)_Y^{\text{init}} \text{ in} \\
&\quad \text{let } M_2, p_2 = (E_2)_Y^{\text{init}} \text{ in} \\
&\quad (M_1, M_2), p_1 + p_2 \\
(E_1 \text{ and } E_2)_Y^{\text{step}}((M_1, M_2), [r_1 : r_2]) &= \text{let } M'_1, \rho_1, w_1 = (E_1)_Y^{\text{step}}(M_1, r_1) \text{ in} \\
&\quad \text{let } M'_2, \rho_2, w_2 = (E_2)_Y^{\text{step}}(M_2, r_2) \text{ in} \\
&\quad (M'_1, M'_2), \rho_1 + \rho_2, w_1 * w_2 \\
(e \text{ where rec } E)_Y^{\text{init}} &= \text{let } m, p_e = (e)_Y^{\text{init}} \text{ in} \\
&\quad \text{let } M, p_E = (E)_Y^{\text{init}} \text{ in} \\
&\quad (m, M), p_e + p_E \\
(e \text{ where rec } E)_Y^{\text{step}}((m, M), [r_e : r_E]) &= \text{let } F(\rho) = \left(\text{let } M', \rho, w = (E)_{\gamma+\rho}(M, r_E) \text{ in } \rho \right) \text{ in} \\
&\quad \text{let } \rho = \text{fix}(F) \text{ in} \\
&\quad \text{let } M', \rho, W = (E)_{\gamma+\rho}^{\text{step}}(M, r_E) \text{ in} \\
&\quad \text{let } m', v, w = (e)_{\gamma+\rho}^{\text{step}}(m, r_e) \text{ in} \\
&\quad (m', M'), v, w * W
\end{aligned}$$

Fig. 6. Density-based co-iterative semantics for ProbZelus equations (full version in Figure 13 of the appendix).

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 transition function evaluates e using m_0 to compute the initial value i and the corresponding weight, and returns a new state containing the current value of x , an environment where $x.\text{last}$ is bound to i , and the weight. 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 , an environment where $x.\text{last}$ is bound to v , and a weight of 1.

Compared to the original kernel-based semantics described in Section 3.2 which combines measures via integration, the density-based semantics only manipulates deterministic values for which the flat CPO on environments described in Section 3.3 is well defined. The initialization function of an expression with a set of local declarations $e \text{ where rec } E$ combines the initial states of e and E and returns the total number of random variables. The transition function takes an array containing the random seeds for e and E , computes the environment ρ defined by E with a fixpoint operator, evaluates $(E)_{\gamma+\rho}(M, r)$ once more to compute the next state M' and the weight W , evaluates the main expression e in the environment $\gamma + \rho$ which returns the result v and a weight w , and returns the combined next states, the result, and the total weight $w * W$. The only difference with the deterministic case is that the transition functions of e and E now take the random seeds as arguments and return the weights.

Scheduling. To compare the density-based semantics with the kernel-based 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 density-based semantics except for the two following rules:

$$\begin{aligned} (E_1 \text{ and } E_2)_{\gamma}^{\text{s step}} ((M_1, M_2), [r_1 : r_2]) &= \text{let } M'_1, \rho_1, w_1 = (E_1)_{\gamma}^{\text{s step}} (M_1, r_1) \text{ in} \\ &\quad \text{let } M'_2, \rho_2, w_2 = (E_2)_{\gamma+\rho_1}^{\text{s step}} (M_2, r_2) \text{ in} \\ &\quad (M'_1, M'_2), \rho_1 + \rho_2, w_1 * w_2 \\ (e \text{ where rec } E)_{\gamma}^{\text{s step}} ((m, M), [r_e : r_E]) &= \text{let } M', \rho, W = (E)_{\gamma}^{\text{s step}} (M, r_E) \text{ in} \\ &\quad \text{let } m', v, w = (e)_{\gamma+\rho}^{\text{s step}} (m, r) \text{ in} \\ &\quad (m', M'_1, M'_2), v, w * W \end{aligned}$$

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.

PROPOSITION 4.1. *For an expression where all equations are scheduled, the scheduled density-based semantics is equal to the density-based semantics with a fixpoint.*

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

LEMMA 4.1. *For all scheduled equations set E , the scheduled semantics yields the same environment as the fixpoint operator, i.e., for an environment γ , a state M and an array of random seeds r :*

$$\text{fix}(\lambda \rho. \text{let } M', \rho', w = (E)_{\gamma+\rho}^{\text{step}} (M, r) \text{ in } \rho') = \text{let } M', \rho, w = (E)_{\gamma}^{\text{s step}} (M, r) \text{ in } \rho$$

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

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

□

4.2 Inference

The `infer` operator first turns the result of the density-based semantics into an un-normalized measure, and then performs the same operation as in the kernel-based semantics: 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.

$$\begin{aligned}
\llbracket \text{infer}(e) \rrbracket_Y^{\text{Dinit}} &= \text{let } m, p = (e)_Y^{\text{init}} \text{ in } \delta_m, p \\
\llbracket \text{infer}(e) \rrbracket_Y^{\text{Dstep}}(\sigma, p) &= \text{let } \psi(m) = \int_{[0,1]^p} \text{let } m', v, w = (e)_Y^{\text{step}}(m, r) \text{ in } w * \delta_{(m', v)} \text{ dr in} \\
&\quad \text{let } \nu = \int \sigma(dm) \psi(m) \text{ in} \\
&\quad \text{let } \bar{\nu} = \nu / \nu(\top) \text{ in} \\
&\quad (\pi_{1*}(\bar{\nu}), p), \pi_{2*}(\bar{\nu})
\end{aligned} \tag{4}$$

The state of the `infer` operator contains the number of random variables in the model p (which remains constant) and a distribution of possible states. The initial distribution of states is a Dirac delta measure over the initial state of the model. The transition function first computes a function ψ mapping a state to the un-normalized measure which associates each pair (next state, value) to its weight. The `infer` operator then integrates this function along all possible values of the state, normalizes it, and splits it into a pair of distributions.

Correctness. The previous definition is very similar to its kernel-based semantics counterpart where the function $\psi(m)$ in Equation (4) plays the role of the semantics of the model. We now show that these two notions coincide.

PROPOSITION 4.2. *For all probabilistic expression e with p random variables where all equations are scheduled, the density-based semantics is the density of the measure computed by the kernel semantics, i.e., for all environment γ and state m :*

$$\int_{[0,1]^p} \text{let } m', v, w = (e)_Y^{\text{step}}(m) \text{ in } w * \delta_{m', v} \text{ dr} = \llbracket e \rrbracket_Y^{\text{step}}(m)$$

PROOF. The kernel-based semantics is only defined for a scheduled language. We first prove by induction on the structure of e that the scheduled density-based semantics coincide with the kernel-based semantics. We can then conclude with Proposition 4.1.

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

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

This property generalizes to discrete distributions, multivariate distributions, and any distributions over Polish spaces. By analogy, we use the notation icdf_μ in all cases.

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

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

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 density-based semantics and in the kernel based semantics.

THEOREM 4.3 (CO-ITERATIVE SEMANTICS CORRECTNESS). *For all probabilistic model e where all equations sets are scheduled, for all environment γ , and for all distribution of state σ :*

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

PROOF. By construction, in the density-based 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-based semantics.

By Proposition 4.2 the un-normalized measure defined by the density-based semantics matches the measure computed by the kernel-based semantics. Given this measure, the rest of the transition function of `infer` is the same in both cases. \square

5 DENSITY-BASED RELATIONAL SEMANTICS

An alternative to the operational view of the co-iterative semantics where expressions are interpreted as state machines is to define a relational semantics where expressions directly return streams of values [Colaço and Pouzet 2003]. This formalism has been used in the Vélus project to prove an end-to-end dataflow synchronous compiler within the Coq proof assistant [Bourke et al. 2017a, 2020, 2021].

In this section, we first present a relational semantics for the deterministic expressions of our language. We then define a relational density-based semantics for probabilistic expressions and prove that this semantics is equivalent to the co-iterative density-based semantics, i.e., the `infer` operator yields the same stream of distributions.

Notations. In the following, V^{ω} 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.

5.1 Deterministic relational semantics

In the relational semantics, deterministic expressions compute streams of values. In a context H which maps variables names to stream of values, the semantics of a deterministic expression

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

Fig. 7. Deterministic relational semantics (full version in Figure 14 of the appendix).

e returns a stream s : $G, H \vdash e \downarrow s$. The additional context G stores global declarations (global constants and function definitions). The semantics of a set of equations E checks that the context H is compatible with all the equations: $G, H \vdash E$. The semantics of a set of equations thus defines a *relation* between the streams stored in the context. Compared to the co-iterative semantics, the relational semantics is not executable since the context must be guessed a priori and validated against the equations.

Figure 7 presents the relational 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.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 an expression with a set of local declarations `e 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.last` is bound to this delayed version of x . In the relational semantics, context are un-ordered maps and scheduling equations does not change the semantics.

5.2 Probabilistic relational semantics

The key idea of the probabilistic relational semantics is similar to the density-based co-iterative semantics: instead of manipulating streams of measures, probabilistic expressions compute streams of pairs (value, score) using external streams of random seeds, and integration is deferred to the `infer` operator.

Figure 8 presents the density-based relational semantics for probabilistic expressions and equations. In a context H which maps variables names to values, the semantics of a probabilistic expression e takes an array of random streams R and returns a stream of pairs (value, weight): $G, H, R \vdash e \Downarrow (s, w)$. The semantics of a set of equations E takes an array containing the random streams of all sub-expressions, checks that the contexts H is compatible with all the equations, and returns the total weight W of all sub-expressions: $G, H, R \vdash E : W$.

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

Fig. 8. Probabilistic relational semantics (full version in Figure 15 of the appendix).

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_μ , and uses the random stream to compute a stream of samples associated to the constant weight 1: $(\mathit{icdf}_{s_{\mu_0}}(R_0), 1) \cdot (\mathit{icdf}_{s_{\mu_1}}(R_1), 1) \cdot \dots$. The semantics of **factor** evaluates its arguments 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 an expression with a set of local definitions the random streams are split between sub-expressions and the weight is the total weight of all sub-expressions.

By construction, for any probabilistic expression e , the size of the array of random streams is the number of random variables defined in e , i.e., the number of **sample**. This information can be computed statically (as in the initialization functions of the co-iterative semantics in Section 4.1), and in the following $\mathit{RV}(e)$ returns the number of random variables in e .

5.3 Inference

As in the density-based co-iterative semantics, the **infer** operator is defined by integrating at each step an un-normalized density function over all possible values of the streams of random seeds. 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: $\bar{w}(R) = \Pi w(R) = w_0(R) \cdot (w_0(R) * w_1(R)) \cdot (w_0(R) * w_1(R) * w_2(R)) \cdot \dots$. Then the function $\mathit{integ} \ 1$) turns v_n and \bar{w}_n into an un-normalized measures by integrating over all possible values of the random streams, and 2) normalizes the result to obtain a stream of distributions of values. If $p = \mathit{RV}(e)$ is the number of random variables in the model and λ_ω^p is the uniform measure

over the cube of random streams $([0, 1]^\omega)^p$:

$$\text{integ}_p(\overline{w} \cdot \overline{ws})(v \cdot vs) = \left(\text{let } \mu = \int \overline{w}(R) \delta_{v(R)} \lambda_\omega^p(dR) \text{ in } \mu / \mu(\top) \right) \cdot (\text{integ } ws \ vs) \quad (5)$$

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 σ -algebra, i.e., the coarsest σ -algebra such that projections are measurable. We define λ_ω 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 λ_ω 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 λ_ω^p as the p -ary product measure of λ_ω , and lift the change of variable formula.

Correctness. For a probabilistic expression e , we first relate the relational semantics of Section 4.1 and the co-iterative semantics of Section 5.1. If H is an environment mapping variables names to streams of values, H_k is the environment where streams are projected on their k -th coordinate and $H_{\leq k}$ is the environment where streams are truncated at k . We define similarly $R_{\leq k}$, and R_k for an array of random streams R .

PROPOSITION 5.1. *For a causal probabilistic model e , if $G, H, R \vdash e \Downarrow (s, w)$ there is a co-iterative execution trace $m_0 = (e)_G^{\text{init}}$ and $\forall k > 0$, $(m_{k+1}, v_{k+1}, w'_{k+1}) = (e)_{H_{k+1}}^{\text{step}}(m_k, R_{k+1})$ such that $\forall k > 0$, m_k, v_k, w'_k only depend on $H_{\leq k}$ and $R_{\leq k}$, and $s_k(H) = v_k(H_{\leq k}, R_{\leq k})$ and $w_k(H) = w'_k(H_{\leq k}, R_{\leq k})$.*

PROOF SKETCH. This proposition states that if a program is causal, i.e., if all equations can be scheduled, the co-iterative semantics and the relational semantics coincide.

First, we can show that if a relational semantics exists, there exists a state machine whose execution matches the relational semantics. A similar proof is at the heart of the Vélus compiler [Bourke et al. 2017a]. While adapting this proof to the ProbZelus language with explicit scores is far from easy, it does not offer any new insights. Second, the semantics of the compiled state-machines is deterministic and corresponds to the co-iterative semantics of the normalized scheduled program which does not require any fixpoint computation. Finally, the co-iterative semantics can be used to prove that source-to-source transformations preserve the semantics, in particular the normalization and scheduling passes [Colaco et al. 2023]. The semantics of the state machine thus corresponds to the co-iterative semantics of the original program.

The property also states that at each instant, the output of a causal model only depends on past inputs and states which can be proved by induction on the structure of the program. \square

As in Section 4.2, we can now state the main correctness theorem, i.e., the `infer` operator yields the same stream of distributions in the co-iterative semantics and in the relational semantics.

THEOREM 5.2 (RELATIONAL SEMANTICS CORRECTNESS). *For a causal probabilistic model e , and for all environments G, H , if $G, H \vdash \text{infer}(e) \Downarrow \mu$ then the co-iterative execution trace yields the same stream of distributions, i.e., $\sigma_0, p = \llbracket \text{infer}(e) \rrbracket_G^{\text{init}}$ and $\forall k > 0$, $(\sigma_k, p), \mu_k = \llbracket \text{infer}(e) \rrbracket_{H_{k+1}}^{\text{step}}(\sigma_k, p)$.*

PROOF. If p is the number of random variables in the model, we show $\forall k > 0$:

$$\begin{aligned}\sigma_{k+1}(H_{k+1}, \sigma_k) &\propto \int_{([0,1]^k)^p} \overline{w}_k(H_{\leq k}, R_{\leq k}) * \delta_{m_k(H_{\leq k}, R_{\leq k})} \lambda_{\leq k}^p(dR_{\leq k}) \\ \mu_{k+1}(H_{k+1}, \sigma_k) &\propto \int_{([0,1]^k)^p} \overline{w}_k(H_{\leq k}, R_{\leq k}) * \delta_{v_k(H_{\leq k}, R_{\leq k})} \lambda_{\leq k}^p(dR_{\leq k})\end{aligned}$$

By the induction hypothesis, $\forall k, H, R, s(H, R)_k = v_k(H_{\leq k}, R_{\leq k})$ and $w(H, R)_k = w'_k(H_{\leq k}, R_{\leq k})$. From the definition of ψ in Equation (4) and Fubini's theorem we have:

$$\begin{aligned}v_{k+1} &= \int \sigma_k(dm) \psi(m) \\ &\propto \int_{([0,1]^k)^p} \overline{w}_k(H_{\leq k}, R_{\leq k}) * \psi(m_k(H_{\leq k}, R_{\leq k})) \lambda_{\leq k}^p(dR_{\leq k}) \\ &\propto \int_{([0,1]^k)^p} \int_{[0,1]^p} \overline{w}_k(H_{\leq k}, R_{\leq k}) * w_{k+1}(H_{k+1}, R_{k+1}) \\ &\quad * \delta_{m_{k+1}(H_{\leq k+1}, R_{\leq k+1}), v_{k+1}(H_{\leq k+1}, R_{\leq k+1})} \lambda^p(dR_{k+1}) \lambda_{\leq k}^p(dR_{\leq k}) \\ &\propto \int_{([0,1]^{k+1})^p} \overline{w}_{k+1}(H_{\leq k+1}, R_{\leq k+1}) * \delta_{m_{k+1}(H_{\leq k+1}, R_{\leq k+1}), v_{k+1}(H_{\leq k+1}, R_{\leq k+1})} \lambda_{\leq k+1}^p(dR_{\leq k+1})\end{aligned}$$

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

$$\mu_k \propto \int \overline{w}_k(R) \delta_{s_k(R)} \lambda_\omega^p(dR)$$

which corresponds to Equation (5) and concludes the proof. \square

5.4 Program equivalence

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

Definition 5.3. Deterministic expressions e_1 and e_2 are equivalent if for all contexts G, H :

$$G, H \vdash e_1 \Downarrow s \quad \text{and} \quad G, H \vdash e_2 \Downarrow s$$

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 density-based relational semantics.

Definition 5.4. Probabilistic expressions e_1 and e_2 with $\text{RV}(e_1) = p_1$ and $\text{RV}(e_2) = p_2$ are equivalent if for any context G, H :

$$\forall k > 0. \int \overline{w}_{1k}(R_1) * \delta_{s_{1k}(R_1)} d\lambda_\omega^{p_1}(R_1) = \int \overline{w}_{2k}(R_2) * \delta_{s_{2k}(R_2)} d\lambda_\omega^{p_2}(R_2)$$

where for all random streams $R_1, R_2: G, H, R_1 \vdash e_1 \Downarrow (s_1, w_1)$ and $G, H, R_2 \vdash e_2 \Downarrow (s_2, w_2)$.

In the relational 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 map the random streams of expression e_1 to the random streams of expression e_2 while preserving uniform distributions, program equivalence can be reduced to the comparison of the streams of pairs (value, weight) computed by each expression.

PROPOSITION 5.5 (PROBABILISTIC EQUIVALENCE). *Probabilistic expressions e_1 and e_2 with $\text{RV}(e_1) = p_1$ and $\text{RV}(e_2) = p_2$ are equivalent if there is a measurable function $f : ([0, 1]^\omega)^{p_1} \rightarrow ([0, 1]^\omega)^{p_2}$ such that, for all contexts G, H :*

- if $f(R_1) = R_2$ and $G, H, R_i \vdash e_i \Downarrow (s_i, w_i)$, then $s_1(R_1) = s_2(R_2)$ and $w_1(R_1) = w_2(R_2)$, and
- $\lambda_\omega^{p_2}$ is the pushforward of $\lambda_\omega^{p_1}$ along f , i.e., $\lambda_\omega^{p_2} = f_*(\lambda_\omega^{p_1})$.

PROOF. The proof is a direct application of the change of variable formula. For all contexts G, H , and $k > 0$:

$$\begin{aligned} \int \overline{w}_{2k}(R_2) * \delta_{s_{2k}(R_2)} \lambda_\omega^{p_2}(dR_2) &= \int \overline{w}_{2k}(f(R_1)) * \delta_{s_{2k}(f(R_2))} f_*(\lambda_\omega^{p_1})(dR_2) \\ &= \int \overline{w}_{1k}(R_1) * \delta_{s_{1k}(R_1)} \lambda_\omega^{p_1}(dR_1) \end{aligned}$$

□

Finding such a mapping 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. In this case, the measurable function is a permutation of the random streams, and two expressions are equivalent if they compute the same stream of pairs (value, weight).

The relational 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 expressions 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 random streams associated to the expressions $\text{sample}(e_i)$. For all context G, H , if $G, H \vdash e_i \Downarrow \mu_i$ we define $s_i = \text{icdf}_{\mu_i}(R_i)$. Then, the derivation tree for the lhs expression is:

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

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

$$\frac{\frac{G, H + H_E, R_2 \vdash \text{sample}(e_2) \Downarrow (s_2, 1) \quad G, H + H_E, R_1 \vdash \text{sample}(e_1) \Downarrow (s_1, 1)}{G, H + H_E, R_2 \vdash x = \text{sample}(e_2) : w_2 \quad G, H + H_E, R_1 \vdash y = \text{sample}(e_1) : w_1}}{G, H + H_E, [] \vdash x + y \Downarrow (s_2 + s_1, 1) \quad G, H + H_E, [R_2 : R_1] \vdash x = \text{sample}(e_2) \text{ and } y = \text{sample}(e_1) : w_1 * w_2}}{G, H, [R_2 : R_1] \vdash x + y \text{ where } \text{rec } x = \text{sample}(e_2) \text{ and } y = \text{sample}(e_1) \Downarrow (s_2 + s_1, 1)}$$

Since both programs compute the same stream of pairs (value, weight), and the permutation $f([R_1 : R_2]) = [R_2 : R_1]$ preserves the uniform distribution, the two programs are equivalent.

6 APPLICATION: ASSUMED PARAMETERS FILTERING

ProbZelus probabilistic models are state-space models that 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.




Fig. 9. 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.

Example. Consider the example of Section 2 where the boat is drifting at a constant speed θ . We want to estimate both the moving position (state parameter), and the drift speed (constant parameter). The motion model f is now defined as follows (where the noise parameter st is a global constant):

```
let proba f(pre_x) = pre_x + theta where
  rec init theta = sample(gaussian(zeros, st))
  and theta = last theta
```

Filtering. To estimate state parameters, Sequential Monte Carlo (SMC) inference algorithms rely on filtering techniques [Chopin and Papaspiliopoulos 2020; Del Moral et al. 2006]. 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 [Atkinson et al. 2022; Baudart et al. 2020, 2022]. 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 θ . 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 9 graphically illustrates this phenomenon.

Assumed Parameter Filter. To mitigate this issue, the *Assumed Parameter Filter* (APF) [Erol et al. 2017] split the inference into two steps: 1) estimate the state parameters, and 2) update the constant parameters. The lower part of Figure 9 illustrates the results of APF on the estimation of the drift speed for our radar example.

The APF algorithm assumes that constant parameters are well identified and that the model is written in a form which makes it possible to update the constant parameters given the state parameters. The constant parameters must be 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 analysis identifies the constant parameters and their prior distributions. Then a compilation pass transforms these parameters into additional input of the model. We use the relational semantics of Section 5 to prove the correctness of this transformation, i.e., the transformation preserves the ideal semantics of the program.

Example. The compiled radar model for APF is the following:

```
let f_prior = gaussian(zeros, id)
proba f_model(theta, pre_x) = pre_x + theta

let tracker_prior = f_prior
proba tracker_model(theta, y_obs) = x where
  rec init x = x_init
  and x = sample(gaussian(f(theta, last x), sx))
  and y = g(x)
  and () = observe(gaussian(y, 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)
```

6.1 Static Analysis

The goal of the static analysis is to identify the constant parameters of each probabilistic node, i.e., initialized random variables (`init x = sample(e)`) that are also constant (`x = last x`). For a program *prog* the judgement $\emptyset, \emptyset \vdash prog : \Phi, C$ builds the environment Φ which associates to each probabilistic node a type ϕ which maps constant parameters to their prior distributions. The environment C contains the global constant variables that can be used to define the prior distributions. The main type system is given in Figure 10.

Constants. The auxiliary judgement $C \vdash^c E : C'$ 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 = last x`) or if it is defined by a constant expression. The auxiliary judgement $C \vdash^c e$ 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. A global declaration `let x = e` typed in the environment Φ, 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 ϕ computed by the judgement $C \vdash e : \phi$.

Expressions. Typing an expression collects the constant parameters of the sub-expressions. To simplify the analysis, we associate a unique instance name θ to each function call and we assume

$$\begin{array}{c}
\frac{C \vdash^c e}{\Phi, C \vdash \text{let } x = e : \Phi, C + \{x\}} \qquad \frac{C \vdash e : \phi}{\Phi, C \vdash \text{proba } f \ x = e : \Phi + \{f \leftarrow \phi\}, C} \\
\\
\frac{C \vdash e_1 : \phi}{C \vdash \text{present } e_1 \rightarrow e_2 \text{ else } e_3 : \phi} \qquad \frac{C \vdash e_2 : \phi}{C \vdash \text{reset } e_1 \text{ every } e_2 : \phi} \qquad \frac{}{C \vdash f_\theta(e) : \{\theta \leftarrow f.\text{prior}\}} \\
\\
\frac{C \vdash e : \phi_e \quad C \vdash^c E : D \quad C, D \vdash E : \phi_E}{C \vdash e \text{ where rec } E : \phi_e + \phi_E} \qquad \frac{x \in D \quad C \vdash^c e}{C, D \vdash \text{init } x = \text{sample}(e) : \{x \leftarrow e\}} \\
\\
\frac{C \vdash e : \phi}{C, D \vdash x = e : \phi} \\
\\
\frac{}{C \vdash^c x = \text{last } x : \{x\}} \qquad \frac{C \vdash^c e}{C \vdash^c x = e : \{x\}} \qquad \frac{x \in C}{C \vdash^c x} \qquad \frac{C + \text{dom}(E) \vdash^c e \quad C \vdash^c E : \text{dom}(E)}{C \vdash^c e \text{ where rec } E}
\end{array}$$

Fig. 10. Extract constant parameters and associated prior distributions (full type system in Appendix B.2).

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 θ 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` e_1 `every` e_2 focuses only on the condition e_2 because e_1 can be re-initialized and thus is not constant. Similarly, the rule for `present` $e_1 \rightarrow e_2$ `else` e_3 focuses only on the condition e_1 .

Equations. 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` $x = \text{sample}(e)$ 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 `f` and is propagated through the node `tracker` that calls `f`. The final environment is:

$$\Phi = \{f \leftarrow \{\text{theta} \leftarrow \text{gaussian}(\dots)\}, \text{tracker} \leftarrow \{\theta \leftarrow f.\text{prior}\}\}$$

6.2 Compilation

To run the APF 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 11 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 Φ for declarations and the type ϕ for expressions.

A model `proba` $f \ x = e$ such that $\Phi(f) = \phi$ (i.e., $C \vdash e : \phi$) is compiled into two statements:

- `let` $f.\text{prior} = \text{im}(\phi)$: the prior distribution of the constant parameters in e , and
- `proba` $f.\text{model}(\text{dom}(\phi), x) = C_\phi(e)$: a new model that takes the constant parameters $\text{dom}(\phi)$ as additional arguments.

The compilation function of an expression C_ϕ 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 $f_\theta(e)$.

- If the node is deterministic, there is no constant parameter.

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

Fig. 11. APF compilation (full definition in Appendix B.3).

- 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_θ 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_θ 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`($f.\text{model}$, $f.\text{prior}$, e).

6.3 Correctness

We use the relational semantics 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`($f.\text{model}$, $f.\text{prior}$, e) then corresponds to the case $f_\theta(e)$.

Definition 6.1. 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.

$$\begin{aligned}
&\text{APF.infer}(f.\text{model}, f.\text{prior}, e) \\
&\quad \equiv \\
&\text{infer}(f.\text{model}(\theta, e) \text{ where rec init } \theta = \text{sample}(f.\text{prior}))
\end{aligned}$$

Notations. In the following, G^+ is the context G for global declarations augmented with the definition of $f.\text{model}$ and $f.\text{prior}$ for all probabilistic functions f . To simplify the presentation we use vectorized notations, e.g., $\text{icdf}_{\vec{\mu}}(R_0)$ is a syntactic shortcut for the list $[\text{icdf}_{\vec{\mu}[0]}(R[0]_0), \text{icdf}_{\vec{\mu}[1]}(R[1]_0), \dots]$.

THEOREM 6.2. *For all probabilistic functions f such that $\Phi(f) = \phi$, for all expressions e in the body of f such that $C \vdash e : \phi_e$, there is a permutation $R \rightarrow [R' : R^P]$ such that*

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

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., $H_f(\vec{p}) = \text{icdf}_{\vec{\mu}}(R_0^P)$.

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 give the mapping between R and $R' + R^p$, and show that a semantics derivation for e in a context G, H, R is equivalent to a semantics derivation for $C_\phi(e)$ in the context $G^+, H + H_f, R'$. We focus on the most interesting cases, i.e., expressions altered by the compilation function.

To simplify the proof we prove the following lemma for equations in parallel with Theorem 6.2.

LEMMA 6.1. *For all equations E in the body of f such that $C, D \vdash E : \phi_E$, there is a permutation $R \rightarrow [R' : R^p]$ such that*

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

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

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

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

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

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 the two permutations $R_e \rightarrow [R'_e : R_e^p]$ and $R_g \rightarrow [R'_g, R_g^p]$. 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 $\vec{v}_p = \text{icdf}_{\vec{\mu}_g}(R_{g0}^p)$.

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

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

$$\frac{\frac{G^+, H + H_f, [] \vdash \theta \Downarrow (\vec{v}_p, 1) \quad G^+, H + H_f, R'_e \vdash C_\phi(e) \Downarrow (s_e, w_e)}{G^+, H + H_f, R'_e \vdash (\theta, C_\phi(e)) \Downarrow ((\vec{v}_p, s_e), w_e)} \quad G^+, [x \leftarrow s_e, \vec{p}_g \leftarrow \vec{v}_p], R'_g \vdash C_{\phi_g}(e_g) \Downarrow (s, w)}{G^+, H + H_f, [R'_e, R'_g] \vdash g.\text{model}(\theta, C_\phi(e)) \Downarrow (s, w_e * w)} \\ \frac{G^+, H + H_f, [R'_e, R'_g] \vdash g.\text{model}(\theta, C_\phi(e)) \Downarrow (s, w_e * w)}{G^+, H + H_f, [R'_e, R'_g] \vdash C_\phi(g_\theta(e)) \Downarrow (s, w_e * w)}$$

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

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

□

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

THEOREM 6.3 (APF COMPILATION). *For all probabilistic nodes f ,*

$$G, H \vdash \text{infer}(f(e)) \Downarrow d \iff G^+, H \vdash \text{APF.infer}(f.\text{model}, f.\text{prior}, e) \Downarrow d$$

PROOF. From Definition 6.1, we need to show that for all random streams R :

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

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

7 RELATED WORK

Probabilistic Semantics. Measurable functions and kernels have been used to define the semantics of first-order probabilistic programs as probability distribution transformers [Kozen 1981]. Probabilistic Coherent Spaces is a generalization of this idea to higher-order types but for discrete probability [Danos and Ehrhard 2011]. This setting has been extended to continuous distributions with models based on positive cones [Dahlqvist and Kozen 2020; Ehrhard et al. 2018], a variation on Banach spaces with positive scalars [Selinger 2004]. 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 [Ehrhard and Geoffroy 2023].

Quasi-Borel spaces (QBS) are another alternative to classic measurable spaces to define the semantics of higher-order probabilistic programs with conditioning [Heunen et al. 2017; Staton et al. 2016]. A probabilistic expression is interpreted as a quasi-Borel measure, i.e., an equivalence class of pairs $[\alpha, \mu]$ where μ is a measure over \mathbb{R} , and α is a measurable function from \mathbb{R} to values. Intuitively, the corresponding distribution is obtained as the pushforward of μ along α . Recent work extends this formalism to capture lazy data structures and streams in a functional probabilistic language [Dash et al. 2023].

Our density-based semantics relies on a similar representation: probabilistic expressions are interpreted by pushingforward a uniform measure over $[0, 1]^n$ along a measurable function. The main difference is that we focus on a domain specific dataflow synchronous language. The set of random variables can be computed statically, and integration is entirely deferred to the `infer` operator. Importantly, we recover the fact that equations sets can be interpreted in any order, a key property for dataflow synchronous languages.

Program Equivalence. Probabilistic bisimulation has been introduced for testing equivalence of discrete probabilistic systems [Larsen and Skou 1989] and generalized to study Labelled Markov Process (continuous systems) [Desharnais et al. 2002]. Following [Lago and Gavazzo 2019] which defines a notion of bisimulation for a (higher-order) probabilistic lambda calculus, we could define a notion of probabilistic bisimulation for the co-iterative semantics. One challenge is to adapt the bisimulation to reactive state machines with explicit state.

Probabilistic coupling, is an alternative classic proof technique for probabilistic programs equivalence [Hsu 2017]. A coupling describes correlated executions by associating pairs of samples. Proposition 5.5 that we use to prove the correctness of the APF compilation pass is an instance of probabilistic coupling where the relation is made explicit through a measurable function.

Static analysis and compilation. The static analysis presented in Section 6.1 is inspired by static analyses designed for dataflow synchronous languages, in particular the initialization analysis [Coço and Pouzet 2004] which guarantees that all streams are well defined at the first time step, and the typing of Zelus’ static arguments which checks that some value can be statically computed at compile time [Bourke et al. 2017b].

The compilation pass presented in Section 6.2 turns constant parameters into additional arguments of the model. An alternative, closer to our density-based semantics, would be to externalize all random variables. This would give a lot of flexibility to the inference runtime to apply different inference strategies to different sets of variables. This approach is also reminiscent of the compilation of Zelus’ hybrid models [Benveniste et al. 2011].

Inference. ProbZelus inference runtime relies on semi-symbolic sequential Monte Carlo samplers [Atkinson et al. 2022; Baudart et al. 2020; Lundén 2017]. The posterior distribution is approximated by a set of independent executions, the particles. Each particle tries to compute a closed form solution using symbolic computation, and only samples concrete values when symbolic computations fails. For some models, symbolic computations can be used to compute the distribution of constant parameters, and ProbZelus default inference algorithms thus outperforms APF. Unfortunately symbolic computations are only possible with conjugate random variables and affine transformations. APF is thus more robust for arbitrary models.

8 CONCLUSION

In this paper we proposed two semantics for a reactive probabilistic programming languages: a density-based co-iterative semantics and a density-based relational 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. The relational semantics directly manipulates streams which can significantly simplify 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 relational 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, Thomas Ehrhard, Patrick Hoscheit, Adrien Guatto, Paul Jeanmaire, Basile Pesin, and Marc Pouzet. This work was supported by the project ReLiSe, Émergence Ville de Paris 2021-2025.

REFERENCES

- Eric Atkinson, Charles Yuan, Guillaume Baudart, Louis Mandel, and M. Carbin. 2022. Semi-Symbolic Inference for Efficient Streaming Probabilistic Programming. In *OOPSLA*.
- Guillaume Baudart, Louis Mandel, Eric Atkinson, Benjamin Sherman, Marc Pouzet, and Michael Carbin. 2020. Reactive Probabilistic Programming. In *PLDI*.
- Guillaume Baudart, Louis Mandel, and Reyyan Tekin. 2022. JAX Based Parallel Inference for Reactive Probabilistic Programming. In *LCTES*.
- Albert Benveniste, Timothy Bourke, Benoit Caillaud, and Marc Pouzet. 2011. Divide and recycle: types and compilation for a hybrid synchronous language. In *LCTES*.
- 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.

- 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.
- Timothy Bourke, L elio Brun, Pierre- Evariste Dagand, Xavier Leroy, Marc Pouzet, and Lionel Rieg. 2017a. A formally verified compiler for Lustre. In *PLDI*.
- Timothy Bourke, L elio 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.
- Timothy Bourke, Francois Carcenac, Jean-Louis Cola o, Bruno Pagano, C edric Pasteur, and Marc Pouzet. 2017b. A Synchronous Look at the Simulink Standard Library. *ACM Trans. Embed. Comput. Syst.* 16, 5s (2017), 176:1–176:24.
- Timothy Bourke, Paul Jeanmaire, Basile Pesin, and Marc Pouzet. 2021. Verified Lustre Normalization with Node Subsampling. *ACM Trans. Embed. Comput. Syst.* 20, 5s (2021), 98:1–98:25.
- Timothy Bourke and Marc Pouzet. 2013. Z elus: a synchronous language with ODEs. In *HSCC*.
- Paul Caspi and Marc Pouzet. 1998. A Co-iterative Characterization of Synchronous Stream Functions. In *CMCS (Electronic Notes in Theoretical Computer Science, Vol. 11)*. Elsevier.
- Nicolas Chopin and Omiros Papaspiliopoulos. 2020. *An introduction to sequential Monte Carlo*. Springer.
- Jean-Louis Cola o and Marc Pouzet. 2004. Type-based Initialization Analysis of a Synchronous Data-flow Language. *International Journal on Software Tools for Technology Transfer* 6, 3 (2004), 245–255.
- Jean-Louis Cola o, Bruno Pagano, and Marc Pouzet. 2017. SCADE 6: A formal language for embedded critical software development. In *TASE*.
- Jean-Louis Cola o and Marc Pouzet. 2003. Clocks as First Class Abstract Types. In *EMSOFT (Lecture Notes in Computer Science, Vol. 2855)*. Springer, 134–155.
- Jean-Louis Colaco, 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*.
- Pascal Cuoq and Marc Pouzet. 2001. Modular Causality in a Synchronous Stream Language. In *ESOP (Lecture Notes in Computer Science, Vol. 2028)*. Springer, 237–251.
- 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*.
- 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.
- Vincent Danos and Thomas Ehrhard. 2011. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Inf. Comput.* 209, 6 (2011), 966–991.
- Swaraj Dash, Younesse Kaddar, Hugo Paquet, and Sam Staton. 2023. Affine Monads and Lazy Structures for Bayesian Programming. *Proc. ACM Program. Lang.* 7, POPL (2023), 1338–1368.
- 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.
- Jos ee Desharnais, Abbas Edalat, and Prakash Panangaden. 2002. Bisimulation for Labelled Markov Processes. *Inf. Comput.* 179, 2 (2002), 163–193.
- Luc Devroye. 2006. Nonuniform random variate generation. *Handbooks in operations research and management science* 13 (2006), 83–121.
- Thomas Ehrhard and Guillaume Geoffroy. 2023. *Integration in cones*. Technical Report. IRIF.
- 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.
- 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*.
- Hong Ge, Kai Xu, and Zoubin Ghahramani. 2018. Turing: A Language for Flexible Probabilistic Inference. In *Proceedings of Machine Learning Research*.
- Noah D. Goodman and Andreas Stuhlm uller. 2014. The Design and Implementation of Probabilistic Programming Languages. <http://dippl.org>
- 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.
- Nicolas Halbhwachs, Paul Caspi, Pascal Raymond, and Daniel Pilaud. 1991. The Synchronous Dataflow Programming Language LUSTRE. *Proc. of the IEEE* 79, 9 (September 1991), 1305–1320.
- Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. 2017. A convenient category for higher-order probability theory. In *LICS*.
- Justin Hsu. 2017. *Probabilistic Couplings for Probabilistic Reasoning*. Ph. D. Dissertation. University of Pennsylvania.
- Dexter Kozen. 1981. Semantics of Probabilistic Programs. *J. Comput. Syst. Sci.* 22, 3 (1981), 328–350.

- Ugo Dal Lago and Francesco Gavazzo. 2019. On Bisimilarity in Lambda Calculi with Continuous Probabilistic Choice. In *MFPS (Electronic Notes in Theoretical Computer Science, Vol. 347)*. Elsevier, 121–141.
- Kim Guldstrand Larsen and Arne Skou. 1989. Bisimulation Through Probabilistic Testing. In *POPL*. ACM Press, 344–352.
- Daniel Lundén. 2017. *Delayed sampling in the probabilistic programming language Anglican*. Master’s thesis. KTH Royal Institute of Technology.
- Michael Montemerlo, Sebastian Thrun, Daphne Koller, and Ben Wegbreit. 2002. FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping Problem. In *AAAI*.
- David Michael Ritchie Park. 1981. Concurrency and Automata on Infinite Sequences. In *Theoretical Computer Science*.
- Marc Pouzet. 2021. An Executable Constructive Semantics for Zelus. Synchron.
- Peter Selinger. 2004. Towards a semantics for higher-order quantum computation.. In *QPL*.
- Sam Staton. 2017. Commutative Semantics for Probabilistic Programming. In *ESOP*.
- 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*.

$$\begin{aligned}
(e)_Y^{\text{init}} &= \llbracket e \rrbracket_Y^{\text{init}}, 0 \\
(e)_Y^{\text{step}}(m, []) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', v, 1 \quad \text{if } e \text{ is deterministic} \\
(\text{sample}(e))_Y^{\text{init}} &= \text{let } m = \llbracket e \rrbracket_Y^{\text{init}} \text{ in } m, 1 \\
(\text{sample}(e))_Y^{\text{step}}(m, [r]) &= \text{let } m', \mu = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', \text{icdf}_\mu(r), 1 \\
(\text{factor}(e))_Y^{\text{init}} &= \text{let } m = \llbracket e \rrbracket_Y^{\text{init}} \text{ in } m, 0 \\
(\text{factor}(e))_Y^{\text{step}}(m, []) &= \text{let } m', v = \llbracket e \rrbracket_Y^{\text{step}}(m) \text{ in } m', (), v \\
(f(e))_Y^{\text{init}} &= \text{let } m_f, p_f = \gamma(f, \text{init}) \text{ in} \\
&\quad \text{let } m_e, p_e = (e)_Y^{\text{init}} \text{ in} \\
&\quad (m_f, m_e), p_f + p_e \\
(f(e))_Y^{\text{step}}((m_f, m_e), [r_f : r_e]) &= \text{let } m'_e, v_e, w_e = (e)_Y^{\text{step}}(m_e, r_e) \text{ in} \\
&\quad \text{let } m'_f, v_f, w_f = \gamma(f, \text{step})(v_e, m_f, r_f) \text{ in} \\
&\quad (m'_f, m'_e), v, w_e * w_f \\
(\text{present } e \rightarrow e_1 \text{ else } e_2)_Y^{\text{init}} &= \text{let } m_1, p_1 = (e_1)_Y^{\text{init}} \text{ in} \\
&\quad \text{let } m_2, p_2 = (e_2)_Y^{\text{init}} \text{ in} \\
&\quad (\llbracket e \rrbracket_Y^{\text{init}}, m_1, m_2), p_1 + p_2 \\
(\text{present } e \rightarrow e_1 \text{ else } e_2)_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) = (e_1)_Y^{\text{step}}(m_1, r_1) \\
&\quad \quad \text{in } (m', m'_1, m_2), v_1, w \\
&\quad \quad \text{else let } (m'_2, v_2, w) = (e_2)_Y^{\text{step}}(m_2, r_2) \\
&\quad \quad \text{in } (m', m_1, m'_2), v_2, w \\
(\text{reset } e_1 \text{ every } e_2)_Y^{\text{init}} &= \text{let } m_1, p = (e_1)_Y^{\text{init}} \text{ in } (m_1, m_1, \llbracket e_2 \rrbracket_Y^{\text{init}}), p \\
(\text{reset } e_1 \text{ every } e_2)_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 = (e_1)_Y^{\text{step}}(\text{if } v_2 \text{ then } m_0 \text{ else } m_1, r) \text{ in} \\
&\quad (m_0, m'_1, m'_2), v_1, w
\end{aligned}$$

Fig. 12. Density-based co-iterative semantics for ProbZelus probabilistic expressions.

A SEMANTICS

A.1 Density-based co-iterative semantics

The density-based co-iterative semantics is presented in Section 4. Figures 12 and 13 presents the full semantics for expressions and equations. The additional rules are for **present** and **reset**.

The initialization of **present** $e \rightarrow e_1$ **else** e_2 allocates memory for e , e_1 and e_2 and count 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 **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.

$$\begin{aligned}
(x = e)_Y^{\text{init}} &= (e)_Y^{\text{init}} \\
(x = e)_Y^{\text{step}}(m, r) &= \text{let } m', v, w = (e)_Y^{\text{step}}(m, r) \text{ in } m', [x \leftarrow v], w \\
(\text{init } x = e)_Y^{\text{init}} &= \text{let } m_0, p = (e)_Y^{\text{init}} \text{ in } (\text{nil}, m_0), p \\
(\text{init } x = e)_Y^{\text{step}}((\text{nil}, m_0), r) &= \text{let } m', i, w = (e)_Y^{\text{step}}(m_0, r) \text{ in } (\gamma(x), m_0), [x.\text{last} \leftarrow i], w \\
(\text{init } x = e)_Y^{\text{step}}((v, m_0), r) &= (\gamma(x), m_0), [x.\text{last} \leftarrow v], 1 \\
(E_1 \text{ and } E_2)_Y^{\text{init}} &= \text{let } M_1, p_1 = (E_1)_Y^{\text{init}} \text{ in} \\
&\quad \text{let } M_2, p_2 = (E_2)_Y^{\text{init}} \text{ in} \\
&\quad (M_1, M_2), p_1 + p_2 \\
(E_1 \text{ and } E_2)_Y^{\text{step}}((M_1, M_2), [r_1 : r_2]) &= \text{let } M'_1, \rho_1, w_1 = (E_1)_Y^{\text{step}}(M_1, r_1) \text{ in} \\
&\quad \text{let } M'_2, \rho_2, w_2 = (E_2)_Y^{\text{step}}(M_2, r_2) \text{ in} \\
&\quad (M'_1, M'_2), \rho_1 + \rho_2, w_1 * w_2 \\
(e \text{ where rec } E)_Y^{\text{init}} &= \text{let } m, p_e = (e)_Y^{\text{init}} \text{ in} \\
&\quad \text{let } M, p_E = (E)_Y^{\text{init}} \text{ in} \\
&\quad (m, M), p_e + p_E \\
(e \text{ where rec } E)_Y^{\text{step}}((m, M), [r_e : r_E]) &= \text{let } F(\rho) = \left(\text{let } M', \rho, w = (E)_{Y+\rho}(M, r_E) \text{ in } \rho \right) \text{ in} \\
&\quad \text{let } \rho = \text{fix}(F) \text{ in} \\
&\quad \text{let } M', \rho, W = (E)_{Y+\rho}^{\text{step}}(M, r_E) \text{ in} \\
&\quad \text{let } m', v, w = (e)_{Y'}^{\text{step}}(m, r_e) \text{ in} \\
&\quad (m', M'), v, w * W
\end{aligned}$$

Fig. 13. Density-based co-iterative semantics for ProbZelus equations.

A.2 Density-based relational semantics

Stream functions. The density-based relational semantics is presented in Section 5. The definition of this semantics relies on a few stream functions.

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

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

Fig. 14. Deterministic relational semantics.

The function tl drops the first element of a stream (tl^n drops the n first elements). The function $\text{map } f \ s$ applies f to each element of the stream s . $\text{merge } cs \ as \ bs$ merges the streams as and bs according to the condition cs . $as \text{ when } cs$ keeps the values of as only when the condition cs is true. The function $\text{slicer } ss \ cs$ 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. It means that 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.

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 $\text{map } f \ H$, e.g., for all $x \in \text{dom}(H)$, $(tl \ H)(x) = tl \ (H(x))$.

Relational semantics. The full deterministic and probabilistic density-based relational semantics including the rules for present and reset are given in Figures 14 and 15. The semantics of $\text{present } e \rightarrow e_1 \text{ else } e_2$ 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 $\text{reset } e_1 \text{ every } e_2$ is based on the slicer function. $[G, (tl^n H) \vdash e_1 \downarrow s_n]_{n \in \mathbb{N}}$ represents the stream of streams where s_n 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.

$$\begin{array}{c}
\frac{}{G, H, [] \vdash c \Downarrow (c, 1)} \quad \frac{}{G, H, [] \vdash x \Downarrow (H(x), 1)} \quad \frac{x \notin H}{G, H, [] \vdash x \Downarrow (G(x), 1)} \\
\frac{G, H, R_e \vdash e \Downarrow (s_e, w_e) \quad G(f) = \text{proba } f \ x = e_f \quad G, [x \leftarrow s_e], R_f \vdash e_f \Downarrow (s, w)}{G, H, [R_e : R_f] \vdash f(e) \Downarrow (s, w * w_e)} \\
\frac{G, H + H_E, R_E \vdash E : w_E \quad G, H + H_E, R_e \vdash e \Downarrow (s, w)}{G, H, [R_e : R_E] \vdash e \text{ where } \text{rec } E \Downarrow (s, w * w_E)} \\
\frac{G, H, R_e \vdash e \Downarrow s_c \quad G, (H, R_1 \text{ when } s_c) \vdash e_1 \Downarrow s_{w_1} \quad G, (H, R_2 \text{ when not } s_c) \vdash e_2 \Downarrow s_{w_2}}{G, H, [R_1 : R_2] \vdash \text{present } e \rightarrow e_1 \text{ else } e_2 \Downarrow \text{merge } s_c \ s_{w_1} \ s_{w_2}} \\
\frac{[G, (tl^n H, R_1) \vdash e_1 \Downarrow s_{w_n}]_{n \in \mathbb{N}} \quad G, H, R_2 \vdash e_2 \Downarrow s_c \quad (s, w) = \text{slicer } (s_{w_0} \cdot s_{w_0} \cdot s_{w_1} \cdot s_{w_2} \cdot \dots) \ s_c}{G, H, [R_1 : R_2] \vdash \text{reset } e_1 \text{ every } e_2 \Downarrow (s, w)} \\
\frac{G, H \vdash e \Downarrow s_\mu}{G, H, [R] \vdash \text{sample}(e) \Downarrow (\text{icdf}_{s_\mu}(R), 1)} \quad \frac{G, H \vdash e \Downarrow w}{G, H, [] \vdash \text{factor}(e) \Downarrow ((), w)} \quad \frac{G, H, R \vdash e \Downarrow (H(x), w)}{G, H, R \vdash x = e : w} \\
\frac{G, H, R \vdash e \Downarrow (i \cdot s, w_i \cdot w) \quad H(x.\text{last}) = i \cdot H(x)}{G, H, R \vdash \text{init } x = e : w_i \cdot 1} \quad \frac{G, H, R_1 \vdash E_1 : w_1 \quad G, H, R_2 \vdash E_2 : w_2}{G, H, [R_1 : R_2] \vdash E_1 \text{ and } E_2 : w_1 * w_2} \\
\frac{p = \text{RV}(e) \quad [G, H, R \vdash e \Downarrow (s, w) \quad \bar{w} = \Pi w]_{R \in ([0,1]^\omega)^p}}{G, H \vdash \text{infer}(e) \Downarrow \text{integ}_p \bar{w} \ s}
\end{array}$$

Fig. 15. Probabilistic relational semantics.

B ASSUMED PARAMETERS FILTERING

B.1 Algorithm

The inference methods proposed by ProbZelus [Atkinson et al. 2022; Baudart et al. 2020, 2022] 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

Data: probabilistic model model , observation y_t , and previous result μ_{t-1} .

Result: μ_t an approximation of the distribution of p_t .

for each particle $i = 1$ **to** N **do**

$p_{t-1}^i = \text{sample}()$
 $p_t^i, w_t^i = \text{model}(y_t \mid p_{t-1}^i)$

$\mu_t = \mathcal{M}(\{w_t^i, p_t^i\}_{1 \leq i \leq N})$

return μ_t

Algorithm 1: Particle Filter.

Data: probabilistic model model , observation y_t , and previous result μ_{t-1} .

Result: μ_t an approximation of the distributions of state parameter x_t and constant parameter θ .

for each particle $i = 1$ **to** N **do**

$x_{t-1}^i, \Theta_{t-1}^i = \text{sample}()$
 $\theta^i = \text{sample}()$
 $x_t^i, w_t^i = \text{model}(y_t \mid \theta^i, x_{t-1}^i)$
 $\Theta_t^i = \text{Udpate}(\Theta_{t-1}^i, \lambda \theta, \text{model}(y_t \mid \theta, x_{t-1}^i, x_t^i))$

$\mu_t = \mathcal{M}(\{w_t^i, (x_t^i, \Theta_t^i)\}_{1 \leq i \leq N})$

return μ_t

Algorithm 2: Assumed Parameter Filter [Erol et al. 2017].

variables) p_t^i of the model. We begin by sampling a new set of particles 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 μ_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 θ . At each instant, the duplicated particles share the same value for θ . The quantity of information useful for estimating θ 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 an optimization pass which updates 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-1}^i and a distribution of constant parameters Θ_{t-1}^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 θ^i in Θ_{t-1}^i . Knowing the value of the constant parameters θ^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

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

Fig. 16. Extract constant parameters and associated prior distributions.

$$\begin{array}{c}
\frac{}{C \vdash^c c} \quad \frac{x \in C}{C \vdash^c x} \quad \frac{C \vdash^c e_1 \quad C \vdash^c e_2}{C \vdash^c (e_1, e_2)} \quad \frac{C \vdash^c e}{C \vdash^c \text{op}(e)} \quad \frac{C + \text{dom}(E) \vdash^c e \quad C \vdash^c E : \text{dom}(E)}{C \vdash^c e \text{ where rec } E} \\
\frac{}{C \vdash^c \text{init } x = e : \emptyset} \quad \frac{}{C \vdash^c x = \text{last } x : \{x\}} \quad \frac{C \vdash^c e}{C \vdash^c x = e : \{x\}} \quad \frac{C \not\vdash^c e}{C \vdash^c x = e : \emptyset} \\
\frac{C + C_2 \vdash^c E_1 : C_1 \quad C + C_1 \vdash^c E_2 : C_2}{C \vdash^c E_1 \text{ and } E_2 : C_1 + C_2}
\end{array}$$

Fig. 17. Constant expressions and equations.

update Θ_t^i by exploring the other possible values for θ 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 μ_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, Θ_t^i) is associated to the probability $w_t^i / \sum_{i=1}^N w_t^i$.

B.2 Static Analysis

Figure 17 presents the auxiliary type system identifying the constant expressions and extracting the names of the constant variables. An expression is constant if it is a constant value c , a variable referring to a constant stream ($x \in C$), or an expression whose all sub-expressions are constants. The judgement $C \vdash^c E : C'$ extracts all the name of all the streams of the set of equations E that are constant. So if $C \vdash^c E : \text{dom}(E)$, all the streams of E are constants. An equation defines a constant

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

Fig. 18. APF Compilation.

streams if it is defined by a constant expression or by the equation $x = \text{last } x$. The full type system presented in Section 6.1 is given in Figure 16.

B.3 Compilation

The entire compilation function to transform a ProbZelus model into a model compatible with `APF.infer` is given in Figure 18. Most cases simply call the compilation functions on all sub-expressions. The interesting cases are presented in Section 6.2.