Reactive Probabilistic Programming

Guillaume Baudart
MIT-IBM Watson AI Lab,
IBM Research
USA

Louis Mandel
MIT-IBM Watson AI Lab,
IBM Research
USA

Eric Atkinson
MIT
USA

Benjamin Sherman
MIT
USA

Marc Pouzet
École Normale Supérieure,
PSL Research University
France

Michael Carbin
MIT
USA

Abstract

Synchronous modeling is at the heart of programming languages like Lustre, Esterel, or SCADE used routinely for implementing safety critical control software, e.g., fly-by-wire and engine control in planes. However, to date these languages have had limited modern support for modeling uncertainty — probabilistic aspects of the software’s environment or behavior — even though modeling uncertainty is a primary activity when designing a control system.

In this paper we present ProbZelus the first *synchronous probabilistic programming language*. ProbZelus conservatively provides the facilities of a synchronous language to write control software, with probabilistic constructs to model uncertainties and perform *inference-in-the-loop*.

We present the design and implementation of the language. We propose a measure-theoretic semantics of probabilistic stream functions and a simple type discipline to separate deterministic and probabilistic expressions. We demonstrate a semantics-preserving compilation into a first-order functional language that lends itself to a simple presentation of inference algorithms for streaming models. We also redesign the delayed sampling inference algorithm to provide efficient *streaming* inference. Together with an evaluation on several reactive applications, our results demonstrate that ProbZelus enables the design of reactive probabilistic applications and efficient, bounded memory inference.

CCS Concepts: • **Theory of computation** → **Streaming models**; • **Software and its engineering** → **Data flow languages**.

Keywords: Probabilistic programming, Reactive programming, Streaming inference, Semantics, Compilation

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PLDI '20, June 15–20, 2020, London, UK

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7613-6/20/06.

<https://doi.org/10.1145/3385412.3386009>

ACM Reference Format:

Guillaume Baudart, Louis Mandel, Eric Atkinson, Benjamin Sherman, Marc Pouzet, and Michael Carbin. 2020. Reactive Probabilistic Programming. In *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '20)*, June 15–20, 2020, London, UK. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3385412.3386009>

1 Introduction

Synchronous languages [2] were introduced thirty years ago for designing and implementing real-time control software. They are founded on the synchronous abstraction [4] where a system is modeled ideally, as if communications and computations were instantaneous and paced on a global clock. This abstraction is simple but powerful: input, output and local signals are streams that advance synchronously and a system is a stream function. It is at the heart of the data-flow languages Lustre [20] and SCADE [13]; it is also the underlying model behind the discrete-time subset of Simulink.

The data-flow programming style is very well adapted to the direct expression of the classic control blocks of control engineering (e.g., relays, filters, PID controllers, control logic), and a discrete time model of the environment, with the feedback between the two. For example, consider a backward Euler integration method defined by the following stream equations and its corresponding implementation in Zelus [7], a language reminiscent of Lustre:

$$x_0 = x_{00} \quad x_n = x_{n-1} + x'_n \times h \quad \forall n \in \mathbb{N}, n > 0$$

```
let node integr (x0, x') = x where
  rec x = x0 -> (pre x + x' * h)
```

The *node* `integr` is a function from input streams `x0` and `x'` to output stream `x`. The *initialization* operator `->` returns its left-hand side value at the first time step and its right-hand side expression on every time step thereafter. The *unit-delay* operator `pre` returns the value of its expression at the previous time step. The following table presents a sample *timeline* showing the sequences of values taken by the streams defined in the program (where `h` is set to 0.1).

xo	0	0	0	0	0	0	0	...
x'	1	2	1	0	-1	-1	1	...
x' * h	0.1	0.2	0.1	0	-0.1	-0.1	0.1	...
pre x	$\perp$	0	0.2	0.3	0.3	0.2	0.1	...
x	0	0.2	0.3	0.3	0.2	0.1	0.2	...

The node `integr` can be used to define other stream functions, e.g., a PID controller, which can be called in control structures like hierarchical automata, e.g., to express a system that switches between automatic and manual control. Compared to a general purpose functional language (or an embedded DSL), the expressiveness of a synchronous language is purposely constrained to modularly ensure safety properties that are critical for the targeted applications: determinacy, deadlock freedom (reactivity), the generation of statically scheduled code that runs in bounded time and space.

However, to date, these languages have had limited support for modeling uncertainty (e.g., a noisy sensor or channel, a variable delay), to simulate the interaction of a software controller and a partially unknown environment, or to infer parameters from noisy observations. Moreover, uncertainty is a first-order design concern for a controller that operates under the assumption of a probabilistic model of their environment (e.g., object tracking). Using a probabilistic environment model and data gathered from observing the environment, a controller can infer a distribution of likely environments given the observations. Existing approaches consist in hand-coding stochastic controllers that have a known solution (e.g., Kalman filters) which can be tedious and error-prone, or to simply perform off-line statistical testing on the generated code of a controller. Alternatively, in recent years, *probabilistic programming* has developed as an approach to endow general purpose languages with the ability to automate inference.

Probabilistic Programming. Probabilistic programming languages are used to describe probabilistic models and automatically infer distributions of latent (i.e., unobserved) parameters from observations.

A popular approach [6, 18, 28, 36–38] consists in extending a general-purpose programming language with three constructs: (1) `x = sample(d)` introduces a random variable `x` of distribution `d`, (2) `observe(d, y)` conditions on the fact that a sample from distribution `d` produced the observation `y`, and (3) `infer m obs` computes the distribution of the output values of a program or *model* `m` w.r.t. the observation of the input data `obs`.

Probabilistic programming languages offer a variety of automatic inference techniques ranging from exact inference [17] to approximate inference [29] and include hybrid approaches that combine exact and approximate techniques when part of the program is analytically tractable [27]. However, a standing challenge for these programming languages is that none of them meet the design goals of synchronous

reactive languages by being immediately amenable to techniques to ensure that for example a program with an indefinite execution time runs in bounded memory.

Inference in the Loop. In this paper we extend Zelus¹ to provide a synchronous probabilistic programming language, ProbZelus. ProbZelus enables one to combine deterministic reactive data-flow programs, such as `integr` (above), with probabilistic programming constructs to produce reactive probabilistic programs.

Compared to other probabilistic languages (e.g. WebPPL, Church, Stan) where inference is executed on terminating pure functions, our probabilistic models are stateful stream processors. Inference on probabilistic models runs in parallel with deterministic processes that interact with the environment. The distributions computed by `infer` at each step can thus be used by deterministic processes to compute new inputs for the next inference step. We term this capability *inference-in-the-loop*.

Streaming Inference. ProbZelus provides multiple inference algorithms, most notably the *delayed sampling* inference algorithm [27]. This hybrid strategy combines the approximate inference technique of particle filtering [19] with exact inference when it is possible to symbolically determine the exact distribution for some or all of the latent variables of the program [16].

However, the memory consumption of delayed sampling strictly increases with the number of random variables which is not practical for reactive applications that operate on infinite streams. We propose a novel *streaming* implementation of delayed sampling that can operate over infinite streams in constant memory for a large class of models. ProbZelus therefore provides an expressive language for reactive probabilistic programming with appropriate memory consumption properties.

Contributions. We present the following contributions:

Design, Semantics, Compilation. We present ProbZelus, the first synchronous probabilistic programming language, combining language constructs for streams (reactivity) with those for probabilistic programming thus enabling *inference-in-the-loop*. We give a measure-based co-iterative semantics for ProbZelus that forms the basis of a compiler and demonstrate a semantics-preserving compilation strategy to a first-order functional language μF .

Streaming Inference. We define the semantics of multiple inference algorithms on μF including particle filtering and delayed sampling. We then present a novel *streaming delayed sampling* implementation which enables partial exact inference over infinite streams in bounded memory for a large class of models.

¹Language distribution and manual available at <http://zelus.di.ens.fr>.

```

let proba kalman (xo, u, acc, gps) = x where
  rec mu = xo -> (a *@ pre x) +@ (b *@ u)
  and x = sample (mv_gaussian (mu, noise))
  and () = observe (gaussian (vec_get (x, 2), 1.0), acc)
  and () = present gps(pos) ->
    observe (gaussian (vec_get (x, 0), 0.01), pos)
    else ()

let node robot (xo, uo, acc, gps) = u where
  rec x_dist = infer 1000 kalman (xo, u, acc, gps)
  and u = uo -> lqr a b (mean (pre x_dist))

```

Figure 1. Robot controller with inference-in-the-loop. $*@$ and $+@$ are matrix operations, $\text{vec_get } x \ i$ is the i th projection.

Evaluation. We evaluate the performance of ProbZelus on a set of benchmarks that illustrate multiple aspects of the language. We demonstrate that streaming delayed sampling drastically reduces the number of particles required to achieve better accuracy compared to a particle filter.

The result is ProbZelus, a synchronous probabilistic language that enables us to write, in the very same source, a deterministic model for the control software and a probabilistic model with complex interactions between the two. On one hand, a deterministic model of a controller can rely on predictions computed by a probabilistic model. On the other hand, a probabilistic model can be programmed in an expressive reactive language. ProbZelus is open source (<https://github.com/IBM/probzelus>). An extended version with appendices of the paper is also available [1].

2 Example

In this section, we demonstrate how ProbZelus provides probabilistic modeling, inference-in-the-loop, and bounded-memory inference for a robot navigation system. The results of the inference are continuously used by a controller to correct the robot trajectory.

2.1 Inference in the Loop.

We consider a robot equipped with an accelerometer and a GPS. We assume that the motion of the robot can be described as: $x_{t+1} = Ax_t + Bu_t$ where x_t denotes the state of the robot (position, velocity, and acceleration) at a given time step t , and u_t denotes the command sent to the robot. A and B are constant matrices. In addition, the robot receives at each step noisy *observations* from the accelerometer a_t , and sporadically an estimation of the position from a GPS p_t .

Figure 1 presents a controller, robot, that given an initial state x_0 , an initial command u_0 , and inputs from the accelerometer acc and the GPS gps computes a stream u of commands that drives the robot to a given target. The body of robot is the parallel composition of (1) the inference of a probabilistic process kalman that estimates x_dist the stream of distributions over the robot's state, and (2) a deterministic

Figure 2. Kalman filter for the robot example. Variables are either latent (white, e.g., state x) or observed (gray, e.g., acceleration a). The position p is only sporadically observed.

process that computes the stream u of commands. It is written as two mutually recursive equations that define x_dist from u and u from the previous value of x_dist .

The command u is set to the initial command u_0 at the first time step, and is then computed by a *Linear-Quadratic Regulator* (LQR) [34] — a stable and optimal controller for such dynamic systems — given the estimation of the state at the previous step. Because LQR controllers depend only on mean posterior state, the example in Figure 1 uses the mean function to compute the mean of x_dist before invoking the LQR controller.

Inference. The stream x_dist of distributions of state is inferred from the model defined by the probabilistic node kalman given the initial state x_0 , the command u , and the observations acc and gps . The keyword `proba` indicates a probabilistic model.

In this example, the model is a *Kalman Filter* illustrated in Figure 2. A Kalman filter is a time-dependent probabilistic model used to describe inference problems such as tracking, in which a tracker estimates the true position of an object given noisy, sensed observations. The robot's state x_t is a *latent* random variable in that the tracker is not able to directly observe it. Each arrow connecting two random variables denotes a dependence of the variable at the head of the arrow on the variable at the tail. In this case, the observations at each time step depends on the current state, and the robot's state at a given time step depends only on its states at the previous time step.

Sampling. Inside the `kalman` node, the `sample` operator samples a value from a probability distribution. In this case, the expression samples the current state x from a multivariate Gaussian with mean obtained by applying the motion model to the previous state and the command. This code models the trajectory of a robot where at each time step, the state is Gaussian-distributed around an estimation computed from the motion model.

Observations. The expression `observe` conditions the execution on observed data. Its first parameter denotes a distribution that models the observation and its second parameter denotes the observed value itself. In this case, at each step, the first `observe` statement models a Gaussian-distributed observation of the current acceleration $\text{vec_get } x \ 2$ given by acc . The input gps is a *signal* that is only emitted when the

Figure 3. Particle filter (PF) and streaming delayed sampling (SDS) performances for the robot example of Figure 1. Accuracy is measured using the loss function of the LQR. Speed corresponds to the execution time of 500 steps.

GPS computes a new position. When a value `pos` is emitted on `gps`, the `present` construct executes its left branch, further conditioning the model by adding a Gaussian-distributed observation of the current position `vec_get x 0` given by `pos`.

2.2 Streaming Inference

A classic operational interpretation of a probabilistic model is an *importance sampler* that generates random samples from the model together with an importance weight measuring the quality of the sample. In this model, each execution of a `sample` operator samples a value from the operator’s corresponding distribution. Each execution of an `observe` evaluates the likelihood of the provided observation and multiplies the current importance weight by this value. Then, each execution step of `infer` yields a distribution represented as a set of pairs (output, weight) or *particles*. The particles can be re-sampled at each step to build a *particle filter* [15].

The integer parameter to `infer` determines how many particles to use: the more particles the user specifies, the more accurate the estimate of the distribution becomes. The PF points in Figure 3 present this improvement in accuracy as a function of increasing the number of particles for the robot example. However, as the latency results presents, the more particles the user specifies, the more computation is required for each step because each particle requires a full, independent execution of each time step of the model.

Streaming Delayed Sampling. Delayed sampling [27] can reduce the number of particles required to achieve a given desired quality of inference. Specifically, delayed sampling exploits the opportunity to symbolically reason about the relationships between random variables to compute closed-form distributions whenever possible. To capture relationships between random variables, delayed sampling maintains a graph: a *Bayesian network* that can be used to compute closed-form distributions involving subsets of random variables. For instance, this inference scheme is able to compute the exact posterior distribution for our robot example. The SDS dots in Figure 3 show that the accuracy is independent of the number of particles since each particle computes the exact solution.

Figure 4 illustrates the evolution of the delayed sampling graph as it proceeds through the first four time steps of the robot example (for simplicity we assume that there is no GPS activation in these four steps). A notable challenge with the traditional delayed sampling algorithm is that the graph grows linearly in the number of samples. This property is not tractable in our reactive context because we would like to deploy our programs under the model that they run indefinitely, thus requiring that they execute with bounded resources. To address this problem, we propose a novel *streaming delayed sampling* (SDS) implementation of the delayed sampling algorithm. Specifically, in Figure 4 the node denoting the marginal posterior for `x` at step 1 can be eliminated from the graph at step 3 because the distributions for `pre x` and `x` have fully incorporated its effect on their values and, moreover, the program no longer maintains a reference to the node.

While the standard delayed sampling algorithm will keep this node alive through the edge pointers it maintains, SDS builds a *pointer-minimal* graph representation with a minimal number of edges that 1) ensure that the graph has sufficient connectivity to support operations in the traditional delayed sampling algorithm and 2) only maintain the reachability of nodes that can effect the distribution of future nodes in the graph. The result is that the memory consumption of SDS is constant across the number of steps while the memory consumption of the original delayed sampling implementation DS increases linearly in the number of steps (Figure 5).

3 Language: Syntax, Typing, Semantics

ProbZelus is a reactive probabilistic language with inference-in-the-loop which enables interaction between probabilistic models and deterministic processes. This capability introduces two design requirements. First, a probabilistic model must be able to receive inputs from an evolving environment. Second, instead of awaiting the final result of the inference, deterministic processes running in parallel need access to intermediate results. The resulting inference-in-the-loop enables feedback loops between inferred distributions from probabilistic models and deterministic processes, which our

Figure 4. Evolution of the delayed sampling graph for the model of Figure 1. Each node denotes either a value (dark gray) or a distribution (light gray). Plain arrows represent dependencies in the underlying Bayesian network. The dotted arrow represents the pointers in the original data-structure implementing the graph. Labels indicate the program variables.

Figure 5. Delayed sampling (DS) and streaming delayed sampling (SDS) memory consumption in thousands of live words in the heap per steps for the robot example.

design controls by enforcing a separation between the semantics of probabilistic and deterministic execution.

In this section, we formalize the syntax of ProbZelus, introduce a type system that imposes a clear separation between deterministic and probabilistic expressions, and define the semantics of the language in a co-iteration framework where the semantics of probabilistic processes is adapted from Staton’s measure-theoretic semantics for probabilistic programs [35]. The co-iterative semantics forms the basis of a compiler that is described in Section 4.

3.1 Syntax

We focus on the following kernel of ProbZelus. The missing constructs (e.g., `pre` and `->`) can be compiled into this kernel via a source-to-source transformation.

```

d ::= let node f x = e | let proba f x = e | d d
e ::= c | x | (e, e) | op(e) | f(e) | last x
    | e where rec E
    | present e -> e else e | reset e every e
    | sample(e) | observe(e, e) | infer(e)
E ::= x = e | init x = c | E and E

```

A program is a sequence of *declarations* d of stream functions (`node`) and probabilistic stream functions (`proba`). An *expression* e is either a constant (c), a variable (x), a pair, an external operator application (op), a function application ($f(e)$), a delay (`last x`) that returns a value (x) from the previous step, or a set of locally recursive equations (e `where rec` E).

A set of *equations* E is either an equation $x = e$ that define x with the stream e , the initialization of a variable with a constant `init x = c`, or parallel composition of sets of equations.

Operators (op) include boolean and arithmetic operators. In addition, ProbZelus offers a library of dedicated operators to analyze distributions, such as mean and variance, that can be used in any context (probabilistic or deterministic), e.g., on the result of the inference.

The control structure `present e -> e1 else e2` is an *activation condition* that executes the expression e_1 only when the value of e is true and executes e_2 otherwise. It differs from `if e then e1 else e2`, where both e_1 and e_2 are computed at each step (making their internal states evolve) and the returned value is chosen based on the value of e .² In the following example, o_1 and o_2 are different streams:

```
let node cpt () = o where rec o = 0 -> pre o + 1
```

```

let node present_vs_if (b) = (o1, o2) where
  rec o1 = present (b) -> cpt () else 0
  and o2 = if b then cpt () else 0

```

b	true	true	false	true	false	false	true	...
o1	0	1	0	2	0	0	3	...
o2	0	1	0	3	0	0	6	...

The `reset e1 every e` construct re-initializes the values of the `init` equations and the corresponding `last` expressions in e_1 each time e is true.

The language is extended with the classic probabilistic expressions: `sample` to draw from a distribution, `observe` to condition on observations, and `infer` to compute the distribution described by a model.

Scheduling. In the expression e `where rec` E , E is a set of mutually recursive equations. In practice, a scheduler re-orders the equations according to their dependencies. Initializations `init xj = cj` are grouped at the beginning, and an equation $x_j = e_j$ must be scheduled after the equation $x_i = e_i$ if the expression e_j uses x_i outside a `last` construct. A program satisfying this partial order is said to be *scheduled*. The compiler can also introduce additional equations to relax the scheduling constraints and rejects programs that cannot

²The `if` construct can thus be considered as an external operator.

be statically scheduled [5]. After scheduling, the expression e where `rec` E has the following form.

e where `rec` `init` $x_1 = c_1 \dots$ and `init` $x_k = c_k$
 and $y_1 = e_1 \dots$ and $y_n = e_n$

For simplicity, we also assume that every initialized variable is defined in a subsequent equation, i.e., $\{x_i\}_{1..k} \cap \{y_j\}_{1..n} = \{x_i\}_{1..k}$. If it is not the case, in this kernel we can always add additional equations of the form $x_i = \text{last } x_i$.

Kernel. All ProbZelus programs can be encoded in this kernel language. For instance, the program `integr` of Section 1 can be rewritten in the kernel as follows:

```
let node integr (xo, x') = x where
  rec init first = true and init x = 0.
  and first = false
  and x = if last first then xo else last x + (x' * h)
```

A stream `first` is defined such that `last first` is only true at the first step. The `->` operator is then compiled into an `if` statement. The `pre` operator is compiled into a `last` operator. The initialization value is arbitrary, the compiler's initialization analysis guarantees that this value is never used [14]. Similarly, other constructs like hierarchical automata can be re-written using `present` and `reset` [12].

3.2 Typing: Deterministic vs. Probabilistic

Deterministic and probabilistic expressions have distinct interpretations. A dedicated type system discriminates between the two kinds of expressions, assigning one of two *kinds* to each expression: D for deterministic, or P for probabilistic. The typing judgment $G \vdash^k e : T$ states that in the environment G which maps variable names to their type, the expression e has kind k and type T . Function types $T \rightarrow^k T'$ are extended with the kind k of the body and we introduce a new datatype $T \text{ dist}$ for the probability distribution over values of type T .

The expressions `sample`, and `observe` are probabilistic but their arguments must be deterministic. Any deterministic expression can be lifted to a probabilistic expression using a sub-typing rule. The transition from probabilistic to deterministic is realized via `infer` (the complete set of typing rules is presented in Figure 12 of Appendix A.1).

$$\frac{G \vdash^D e : T \text{ dist}}{G \vdash^P \text{sample}(e) : T} \quad \frac{G \vdash^D e_1 : T \text{ dist}^* \quad G \vdash^D e_2 : T}{G \vdash^P \text{observe}(e_1, e_2) : \text{unit}}$$

$$\frac{G \vdash^D e : T}{G \vdash^P e : T} \quad \frac{G \vdash^P e : T}{G \vdash^D \text{infer}(e) : T \text{ dist}} \quad \frac{G \vdash^D e : T \text{ dist}^*}{G \vdash^D e : T \text{ dist}}$$

The type $T \text{ dist}$ represents distributions over values of type T . Distributions can be sampled (`sample` statement) and analyzed with external operators such as mean and variance.

$$\llbracket x \rrbracket_Y^i = ()$$

$$\llbracket x \rrbracket_Y^s = \lambda s. (\gamma(x), s)$$

$$\llbracket \text{present } e \rightarrow e_1 \text{ else } e_2 \rrbracket_Y^i = (\llbracket e \rrbracket_Y^i, \llbracket e_1 \rrbracket_Y^i, \llbracket e_2 \rrbracket_Y^i)$$

$$\llbracket \text{present } e \rightarrow e_1 \text{ else } e_2 \rrbracket_Y^s =$$

$$\lambda (s, s_1, s_2). \text{let } v, s' = \llbracket e \rrbracket_Y^s(s) \text{ in}$$

$$\text{if } v \text{ then let } v_1, s'_1 = \llbracket e_1 \rrbracket_Y^s(s_1) \text{ in } (v_1, (s', s'_1, s_2))$$

$$\text{else let } v_2, s'_2 = \llbracket e_2 \rrbracket_Y^s(s_2) \text{ in } (v_2, (s', s_1, s'_2))$$

$$\left\| e \text{ where rec init } x_1 = c_1 \dots \text{ and init } x_k = c_k \right\|_Y^i =$$

$$\left((c_1, \dots, c_k), (\llbracket e_1 \rrbracket_Y^i, \dots, \llbracket e_n \rrbracket_Y^i), \llbracket e \rrbracket_Y^i \right)$$

$$\left\| e \text{ where rec init } x_1 = c_1 \dots \text{ and init } x_k = c_k \right\|_Y^s =$$

$$\left((c_1, \dots, c_k), (\llbracket e_1 \rrbracket_Y^s, \dots, \llbracket e_n \rrbracket_Y^s), \llbracket e \rrbracket_Y^s \right)$$

$$\lambda ((m_1, \dots, m_k), (s_1, \dots, s_n), s).$$

$$\text{let } \gamma_1 = \gamma[m_1/x_1_last] \text{ in } \dots \text{let } \gamma_k = \gamma_{k-1}[m_k/x_k_last] \text{ in}$$

$$\text{let } v_1, s'_1 = \llbracket e_1 \rrbracket_Y^s(s_1) \text{ in let } \gamma'_1 = \gamma_k[v_1/y_1] \text{ in } \dots$$

$$\text{let } v_n, s'_n = \llbracket e_n \rrbracket_Y^s(s_n) \text{ in let } \gamma'_n = \gamma'_{n-1}[v_n/y_n] \text{ in}$$

$$\text{let } v, s' = \llbracket e \rrbracket_Y^s(s) \text{ in}$$

$$v, ((\gamma'_n[x_1], \dots, \gamma'_n[x_k]), (s'_1, \dots, s'_n), s')$$

Figure 6. Semantics of deterministic expressions.

The type $T \text{ dist}^*$ is a subtype of $T \text{ dist}$ that represents distributions known to have a density, i.e., discrete distributions (w.r.t. the counting measure) and a subset of continuous distributions (w.r.t. the Lebesgue measure). In ProbZelus, to simplify the semantics of the language, the `observe` statement requires a value of type $T \text{ dist}^*$. In Appendix A, we extend the language with the `factor` statement for arbitrary conditioning.

3.3 Co-Iterative Semantics

We now give the semantics of ProbZelus in a co-iteration framework [11]. In this framework, a *deterministic stream* of type T is defined by an initial state of type S and a transition function of type $S \rightarrow T \times S$.

$$\text{CoStream}(T, S) = S \times (S \rightarrow T \times S)$$

Repeatedly executing the transition function from the initial state yields a stream of values of type T .

The semantics of a *deterministic expression* ($\text{kindOf}(e) = D$) is defined using two auxiliary functions. If γ is an environment mapping variable names to values, $\llbracket e \rrbracket_Y^i$ denotes the initial state, and $\llbracket e \rrbracket_Y^s$ denotes the transition function:

$$\llbracket e \rrbracket_Y : \text{CoStream}(T, S) = \llbracket e \rrbracket_Y^i, \llbracket e \rrbracket_Y^s$$

The deterministic semantics of ProbZelus presented in Figure 6 is an extension of [11] with the control structures `present` and `reset` (see also Figure 13 of Appendix A.2).

The transition function of a variable always returns the corresponding value stored in the environment γ .

The **present** $e \rightarrow e_1$ **else** e_2 construct introduced in Section 2 returns the value of e_1 when e is true and the value of e_2 otherwise. The state (s, s_1, s_2) stores the state of the three sub-expressions. The transition function lazily executes the expression e_1 or e_2 depending on the value of e and returns the updated state.

The state of a set of scheduled locally recursive definitions e **where** **rec** E comprises three parts: the value of the local variables at the previous step which can be accessed via the **last** operator, the state of the defining expressions, and the state of expression e . The initialization stores the initial values introduced by **init** and the initial states of all sub-expressions. The transition function incrementally builds the local environment defined by E . First the environment is populated with a set of fresh variables x_{i_last} initialized with the values stored in the state that can then be accessed via the **last** operator. Then the environment is extended with the definition of all the variables y_i by executing all the defining expressions (where $\{x_i\}_{1..k} \cap \{y_j\}_{1..n} = \{x_i\}_{1..k}$). Finally, the expression e is executed in the final environment. The updated state contains the value of the initialized variables defined in E that will be used to start the next step, and the updated state of the sub-expressions.

Probabilistic extension. The semantics of a *probabilistic expression* ($\text{kindOf}(e) = P$) follows the same scheme, but the transition function returns a *measure* over the set of possible pairs (result, state):

$$\text{CoPStream}(T, S) = S \times (S \rightarrow (\Sigma_{T \times S} \rightarrow [0, \infty]))$$

A measure μ associates a positive number to each measurable set $U \in \Sigma_{T \times S}$ where $\Sigma_{T \times S}$ denotes the Σ -algebra of $T \times S$, i.e., the set of measurable sets over pairs (result, state). We use the following notation for the semantics of a probabilistic expression e :

$$\llbracket e \rrbracket_\gamma : \text{CoPStream}(T, S) = \llbracket e \rrbracket_\gamma^i, \llbracket e \rrbracket_\gamma^s$$

The semantics of probabilistic expressions is presented in Figure 7 (the complete semantics is in Figure 14 of Appendix A.2). This measure-based semantics is adapted from [35] to explicitly handle the state of the transition functions.

First, any deterministic expression can be lifted as a probabilistic expression. The transition function returns the Dirac delta measure ($\delta_x(U) = 1$ if $x \in U$, 0 otherwise) on the pair returned by the deterministic transition function applied on the current state: $\llbracket e \rrbracket_\gamma^s(s) : T \times S$.

The probabilistic operator **sample**(e) evaluates e which returns a distribution $\mu : T \text{ dist}$ and a new state $s' : S$, and returns a measure over the pair (result, state) where the state is fixed to the value s' . **observe**(e_1, e_2) evaluates e_1 and e_2 into a distribution with density $\mu : T \text{ dist}^*$ and a value $v : T$, and weights execution paths using the likelihood of v w.r.t. μ (μ_{pdf} denotes the density function of the distribution μ).

$$\begin{aligned} \llbracket e \rrbracket_\gamma^i &= \llbracket e \rrbracket_\gamma^i & \text{if } \text{kindOf}(e) = D \\ \llbracket e \rrbracket_\gamma^s &= \lambda s. \lambda U. \delta_{\llbracket e \rrbracket_\gamma^s(s)}(U) & \text{if } \text{kindOf}(e) = D \\ \llbracket \text{sample}(e) \rrbracket_\gamma^i &= \llbracket e \rrbracket_\gamma^i \\ \llbracket \text{sample}(e) \rrbracket_\gamma^s &= \lambda s. \lambda U. \text{let } \mu, s' = \llbracket e \rrbracket_\gamma^s(s) \text{ in } \int_T \mu(dv) \delta_{v, s'}(U) \\ \llbracket \text{observe}(e_1, e_2) \rrbracket_\gamma^i &= (\llbracket e_1 \rrbracket_\gamma^i, \llbracket e_2 \rrbracket_\gamma^i) \\ \llbracket \text{observe}(e_1, e_2) \rrbracket_\gamma^s &= \lambda(s_1, s_2). \lambda U. \\ &\quad \text{let } \mu, s'_1 = \llbracket e_1 \rrbracket_\gamma^s(s_1) \text{ in} \\ &\quad \text{let } v, s'_2 = \llbracket e_2 \rrbracket_\gamma^s(s_2) \text{ in } \mu_{\text{pdf}}(v) * \delta_{(), (s'_1, s'_2)}(U) \\ \llbracket e \text{ where rec init } x_1 = c_1 \dots \text{ and init } x_k = c_k \dots \text{ and } y_1 = e_1 \dots \text{ and } y_n = e_n \rrbracket_\gamma^s &= \\ \lambda((m_1, \dots, m_k), (s_1, \dots, s_n), s). \lambda U. & \\ \text{let } \gamma_1 = \gamma[m_1/x_{1_last}] \text{ in } \dots \text{ let } \gamma_k = \gamma_{k-1}[m_k/x_{k_last}] \text{ in} & \\ \text{let } \mu_1 = \llbracket e_1 \rrbracket_{\gamma_1}^s(s_1) \text{ in} & \\ \int \mu_1(dv_1, ds'_1) \text{let } \gamma'_1 = \gamma_k[v_1/y_1] \text{ in } \dots & \\ \text{let } \mu_n = \llbracket e_n \rrbracket_{\gamma'_{n-1}}^s(s_n) \text{ in} & \\ \int \mu_n(dv_n, ds'_n) \text{let } \gamma'_n = \gamma'_{n-1}[v_n/y_n] \text{ in} & \\ \text{let } \mu = \llbracket e \rrbracket_{\gamma'_n}^s(s) \text{ in} & \\ \int \mu(dv, ds') \delta_{v, ((\gamma'_n[x_1], \dots, \gamma'_n[x_k]), (s'_1, \dots, s'_n), s')}(U) & \end{aligned}$$

Figure 7. Semantics of probabilistic expressions.

The state of a set of locally recursive definitions is the same as in Figure 6 and contains the previous value of the initialized variables and the states of the sub-expressions. The transition starts by adding the variables x_{i_last} to the environment. We note $\int \mu(dv, ds) f(v, s)$ the integral of f w.r.t. the measure μ where variables v and s are the integration variables. The integration measure appears on the right of the integral to maintain the expression order of the source code and we allow local definitions (e.g., $\text{let } x = v \text{ in } \dots$) inside the integral to simplify the presentation. Local definitions are interpreted by successively integrating the measure on pairs (value, state) returned by the defining expressions. In other words, we integrate over all possible executions. Integrals need to be nested to capture the eventual dependencies in the successive expressions. The returned value is a measure on pairs (value, state) where the state captures the value of the initialized variables and the state of the sub-expressions.

Inference in the loop. The **infer** operator is the boundary between the deterministic and the probabilistic expressions. Given a probabilistic model defined by an expression e , at each step the inference computes a distribution of results and a distribution of possible next states. Expression e can contain free variables thus capturing inputs from deterministic processes. The distribution of results can be used by deterministic processes to produce new inputs for the next inference step.

$$\begin{aligned} \llbracket \text{infer}(e) \rrbracket_Y^i &= \lambda U. \delta_{\llbracket e \rrbracket_Y^i}(U) \\ \llbracket \text{infer}(e) \rrbracket_Y^s &= \lambda \sigma. \text{let } \mu = \lambda U. \int_S \sigma(ds) \llbracket e \rrbracket_Y^s(s)(U) \text{ in} \\ &\quad \text{let } \nu = \lambda U. \mu(U) / \mu(\top) \text{ in} \\ &\quad (\pi_{1*}(\nu), \pi_{2*}(\nu)) \end{aligned}$$

The state of $\text{infer}(e)$ is a distribution over the possible states for e . The initial state is the Dirac delta measure on the initial state of e . The transition function integrates the measure defined by e over all the possible states and normalize the result μ to produce a distribution $\nu : T \times S \text{ dist}$ ($\top$ denotes the entire space). This distribution is then split into a pair of marginal distributions using the pushforward of μ across the projections π_1 and π_2 .

Remark. In ProbZelus deterministic processes that interact with the environment cannot rollback on actions based on past estimations (e.g., the command of the robot controller). However, the inferred distribution of a random variable captured in the state may evolve at each time step based on subsequent observations (e.g., the initial position of the robot). These two properties follow from the separation of the distribution of results and the distribution of states in the semantics of infer .

Alternatively, we could define a fix-point semantics of streams based on a Scott order, as simple lazy streams in Haskell, where the value of a stream can depend on future computations (as illustrated Appendix A.3). However, this approach is not practical in a reactive context where processes interact with the environment [10, 22].

4 Compilation

Following the semantics described in Section 3.3 each expression is compiled into a transition function that can be written in a simple functional first-order language extended with probabilistic operators we call μF . Importantly, the compilation process is the same for deterministic and probabilistic expressions. We can then give a classic interpretation to deterministic terms, and a measure-based semantics to probabilistic terms following [35]. We then show that the semantics of the compiled code coincides with the co-iterative semantics described in Section 3.3.

4.1 A First-Order Functional Probabilistic Language

The syntax of μF is the following:

```

d ::= let f = e | d d
e ::= c | x | (e, e) | op(e) | e(e)
    | if e then e else e | let p = e in e | fun p -> e
    | sample(e) | observe(e, e) | infer((fun x -> e), e)
p ::= x | (p, p)

```

A program is a set of definitions. An expression is either a constant, a variable, a pair, an operator, a function call, a conditional, a local definition, an anonymous function, or one

of the probabilistic operators `sample`, `observe`, or `infer`. The `infer` operator is tailored for ProbZelus and always takes two arguments: a transition function of the form `fun x -> e`, and a distribution of states. This operator computes the distribution of results and the distribution of possible next steps. A type system similar to the one of ProbZelus is used to distinguish deterministic from probabilistic expressions (see Appendix B.1).

4.2 Compilation to μF

The compilation function C generates a function that closely follows the transition function defined by the co-iterative semantics presented in Section 3.3. Each expression is compiled into a function of type $S \rightarrow T \times S$ which given a state returns a value and an updated state (see Appendix C for the complete definition). The compilation of `present` is thus:

```

C(present e -> e1 else e2) = fun (s, s1, s2) ->
  let v, zs' = C(e)(s) in
  if v then let v1, s1' = C(e1)(s1) in (v1, (s', s1', s2))
  else let v2, s2' = C(e2)(s2) in (v2, (s', s1, s2'))

```

The probabilistic operators `sample`, and `observe` are treated as external operators. The compilation generates code that simply calls the μF version of these operators. The compilation of `infer` passes the distribution over states to the μF version of `infer`. The inference is thus aware of the distribution over states at the previous step.

```

C(infer(e)) = fun sigma ->
  let mu, sigma' = infer(C(e), sigma)
  in (mu, sigma')

```

The compilation of a node declaration generates two definitions: the transition function f_step and the initial state f_init . The transition function is the result of compiling the body of the node with an additional argument to capture the input. The initial step is generated by the allocation function $\mathcal{A}$ which follows the definition of the initial state in the semantics of Section 3.3 (see the Appendix C for the complete definition).

```

C(let node f x = e) =
  let f_step = fun (s, x) -> C(e)(s)
  let f_init = A(e)

```

4.3 Semantics Equivalence

We showed how to compile ProbZelus to μF a simple functional language with no loops, no recursion, and no higher-order functions, extended with the probabilistic operators. This language is similar to the kernel presented in [35] for which a measure-based probabilistic semantics is defined (see also Figure 16 of Appendix B.2).

We can now prove that the semantics of the generated code corresponds to the semantics of the source language described in Section 3.3.

Theorem. For all ProbZelus expression e , for all state s and environment γ :

- if $\text{kindOf}(e) = D$ then $\llbracket e \rrbracket_{\gamma}^s(s) = \llbracket C(e) \rrbracket_{\gamma}(s)$, and,
- if $\text{kindOf}(e) = P$ then $\llbracket e \rrbracket_{\gamma}^s(s) = \llbracket C(e) \rrbracket_{\gamma}(s)$.

Proof. The proof is done by induction on the structure of e . As an example consider the expression `sample`(e). If this expression is well-typed and since typing is preserved by compilation (see Lemma C.1) $\text{kindOf}(C(e)) = D$. Using the induction hypothesis on $\llbracket C(e) \rrbracket_{\gamma}(s) = \llbracket e \rrbracket_{\gamma}(s)$ we have:

$$\begin{aligned}
 & \llbracket C(\text{sample}(e)) \rrbracket_{\gamma}(s) \\
 &= \left\llbracket \text{fun } s \rightarrow \text{let } \mu, s' = C(e)(s) \text{ in} \right. \\
 & \quad \left. \text{let } v = \text{sample}(\mu) \text{ in } (v, s') \right\rrbracket_{\gamma}(s) \\
 &= \lambda U. \int \delta_{\llbracket C(e) \rrbracket_{\gamma}(s)}(d\mu, ds') \\
 & \quad \llbracket \text{let } v = \text{sample}(\mu) \text{ in } (v, s') \rrbracket_{\gamma[\mu/\mu, s'/s']}(U) \\
 &= \lambda U. \text{let } (\mu, s) = \llbracket C(e) \rrbracket_{\gamma}(s) \text{ in} \\
 & \quad \llbracket \text{let } v = \text{sample}(\mu) \text{ in } (v, s') \rrbracket_{\gamma[\mu/\mu, s'/s']}(U) \\
 &= \lambda U. \text{let } (\mu, s') = \llbracket e \rrbracket_{\gamma}^s(s) \text{ in } \int \mu(dv) \delta_{v, s'}(U) \\
 &= \llbracket \text{sample}(e) \rrbracket_{\gamma}^s(s)
 \end{aligned}$$

□

Remark. The probabilistic semantics of μF is commutative (see [35, Theorem 4]). We can thus show that the semantics of a ProbZelus program does not depend on the schedule used by the compiler to order the equations of local definitions.

5 Inference

The measure-based semantics of `infer` presented in Section 3.3 and Section 4.3 includes often intractable integrals. An additional challenge is to design inference techniques that can operate in bounded memory to be practical in a reactive context where the inference is a non-terminating process.

In this section, we show how to adapt particle filtering [15] to explicitly handle the state of the transition functions. We then present a novel implementation of *delayed sampling*, a recently proposed semi-symbolic inference, which enables partial exact inference over infinite streams in bounded memory for a large class of models including state-space models like the robot example of Figure 1 in Section 2.

5.1 Particle Filtering

In conventional probabilistic programming, the operational interpretation of a model is an *importance sampler* that randomly generates a sample of the model together with an importance weight measuring the quality of the sample.

Following the conventions of Section 3.3 we write $\llbracket e \rrbracket_{\gamma}$ for the semantics of a deterministic expression, and $\llbracket e \rrbracket_{\gamma, w}$ for the semantics of a probabilistic expression. The additional argument w captures the weight. The probabilistic operator `sample` draws a sample from a distribution without changing the score. `observe` increments the score by the likelihood of the observation. A deterministic expression can be lifted in a probabilistic context: the corresponding sample is the return value of the expression and the score is unchanged. The `let` construct illustrates that the score is accumulated along the execution path (the complete semantics is presented in Figure 19 of Appendix D).

$$\begin{aligned}
 \llbracket \text{sample}(e) \rrbracket_{\gamma, w} &= (\text{draw}(\llbracket e \rrbracket_{\gamma}), w) \\
 \llbracket \text{observe}(e_1, e_2) \rrbracket_{\gamma, w} &= \text{let } \mu = \llbracket e_1 \rrbracket_{\gamma} \text{ in } ((\cdot), w * \mu_{\text{pdf}}(\llbracket e_2 \rrbracket_{\gamma})) \\
 \llbracket e \rrbracket_{\gamma, w} &= (\llbracket e \rrbracket_{\gamma}, w) \text{ if } \text{kindOf}(e) = D \\
 \llbracket \text{let } p = e_1 \text{ in } e_2 \rrbracket_{\gamma, w} &= \text{let } v_1, w_1 = \llbracket e_1 \rrbracket_{\gamma, w} \text{ in } \llbracket e_2 \rrbracket_{\gamma[v_1/p], w_1}
 \end{aligned}$$

Such a sampler is the basis of a *particle filter* or a *bootstrap filter* [15]. `infer` independently launches N particles. At each step, each particle samples the distribution of states σ obtained at the previous step and executes the sampler to compute a pair (result, state) along with its weight. The resulting pairs are then normalized according to their weights to form a categorical distribution μ over pairs of values and states (we write $\bar{w}_i = w_i / \sum_{i=1}^N w_i$ for the normalized weights). This distribution is then split into the distribution of returned values and the distribution of next states.

$$\begin{aligned}
 \llbracket \text{infer}(\text{fun } s \rightarrow e, \sigma) \rrbracket_{\gamma} &= \\
 & \text{let } \mu = \lambda U. \sum_{i=1}^N \text{let } s_i = \text{draw}(\llbracket \sigma \rrbracket_{\gamma}) \text{ in} \\
 & \quad \text{let } (v_i, s'_i), w_i = \llbracket \text{fun } s \rightarrow e \rrbracket_{\gamma, 1}(s_i) \text{ in} \\
 & \quad \bar{w}_i * \delta_{v_i, s'_i}(U) \\
 & \text{in } (\pi_{1*}(\mu), \pi_{2*}(\mu))
 \end{aligned}$$

Remark. The resampling step requires the ability to clone particles in the middle of the execution. A classic technique is to compile the model in *continuation passing style* (CPS) [33] and use the probabilistic constructs `sample` and `observe` as checkpoints for resampling. In our context, the compilation presented in Section 4 externalizes the state of the transition functions. Duplicating the state effectively clones a particle during its execution. The code does not need to be compiled in CPS form and we avoid the alignment problem [24].

5.2 Delayed Sampling

The basis of our streaming inference algorithm is delayed sampling, which we first review to explain its conceptual approach. Delayed sampling is an inference technique combining partial exact inference with approximate particle filtering to reduce estimation errors [23, 27].

```

 $\llbracket op(e) \rrbracket_{Y,g,w} =$ 
   $let (e', g_e, w_e) = \llbracket e \rrbracket_{Y,g,w} in (app(op, e'), g_e, w_e)$ 

 $\llbracket if\ e\ then\ e_1\ else\ e_2 \rrbracket_{Y,g,w} =$ 
   $let\ e', g_e, w_e = \llbracket e \rrbracket_{Y,g,w}\ in$ 
   $let\ v, g_v = value(e', g_e)\ in$ 
   $if\ v\ then\ \llbracket e_1 \rrbracket_{Y,g_v,w_e}\ else\ \llbracket e_2 \rrbracket_{Y,g_v,w_e}$ 

 $\llbracket sample(e) \rrbracket_{Y,g,w} =$ 
   $let\ \mu, g_e, w' = \llbracket e \rrbracket_{Y,g,w}\ in$ 
   $let\ X, g' = assume(\mu, g_e)\ in\ (X, g', w')$ 

 $\llbracket observe(e_1, e_2) \rrbracket_{Y,g,w} =$ 
   $let\ \mu, g_1, w_1 = \llbracket e_1 \rrbracket_{Y,g,w}\ in\ let\ X, g_x = assume(\mu, g_1)\ in$ 
   $let\ e'_2, g_2, w_2 = \llbracket e_2 \rrbracket_{Y,g_x,w_1}\ in\ let\ v, g_v = value(e'_2, g_2)\ in$ 
   $let\ g' = observe(X, v, g_v)\ in\ ((), g', w_2 * \mu_{pdf}(v))$ 

```

Figure 8. Delayed sampling sampler. Expressions return a pair (symbolic expression, weight).

In addition to the importance weight, each particle exploits *conjugacy* relationships between pairs of random variables to maintain a graph: a *Bayesian network* representing closed-form distributions involving subsets of random variables. Observations are incorporated by analytically *conditioning* the network. Particles are thus only required to draw sample when forced to, i.e., when exact computation is not possible, or when a concrete value is required.

To perform analytic computations, delayed sampling manipulates symbolic terms where random variables are referenced in the graph. The semantics of an expression $\llbracket e \rrbracket_{Y,g,w}$ takes an additional argument g for the graph and returns a symbolic term, an updated weight, and an updated graph. Given a graph, a symbolic term can be evaluated into a concrete value by sampling the random variables that appear in the term. The graph can be accessed and modified using the three following functions defined in [27].

$v, g' = value(e, g)$ evaluate a symbolic term and return a concrete value.

$X, g' = assume(\mu, g)$ add a random variable $X \sim \mu$ to the graph and return the variable.

$g' = observe(X, v, g)$ condition the graph by observing the value v for the variable X .

Compared to the particle filter, any expression, probabilistic or deterministic, can contribute to a symbolic term. The evaluation function $\llbracket e \rrbracket_{Y,g,w}$ partially presented in Figure 8 must thus be defined on the entire language and not only on probabilistic constructs. For instance, the application of an operator $op(e)$ returns a symbolic term $app(op, e')$ that represents the application of op on the evaluation of e . Some terms are partially evaluated when symbolic computation is not possible. For instance, in the general case, to compute the importance weight of `if  $e$  then  $e_1$  else  $e_2$` , each particle must compute a concrete value for the condition e .

The probabilistic `sample( $e$ )` adds a new random variable to the graph without drawing a sample. `observe( $e_1, e_2$ )` adds a new random variable $X \sim \mu$ where μ is defined by e_1 , then computes a concrete value v for e_2 and conditions the graph by observing the value v for X . As for the particle filter, the score is incremented by the likelihood of the observation.

Symbolic Computations. The functions *value*, *assume*, and *observe* used in Figure 8 rely on the following mutually recursive lower level operations (Y is the parent of X) [27]:

$X, g' = initialize(\mu, Y, g)$ add a new node X with a distribution $p_{X|Y} = \mu$ as a child of Y in g .

$g' = marginalize(X, g)$ compute and store p_X in g' from p_Y and $p_{X|Y}$ where p_Y and $p_{X|Y}$ are in g .

$g' = realize(X, v, g)$ assign in g' a concrete value to a random variable X .

$g' = condition(Y, g)$ compute $p_{Y|X}$ given $p_X, p_{X|Y}$, and a concrete value $X = v$ where v is in g .

In the class of Bayesian networks maintained by the delayed sampler, marginalization w.r.t. a parent node, and conditioning a parent on the value of a child are tractable operations.

To reflect these operations, nodes are characterized by a state (see Figures 4 and 9). *Initialized* nodes $\bigcirc$ are random variables with a conditional distribution $p_{X|Y}$ where the parent Y has no concrete value yet. *Marginalized* nodes $\bigcirc$ are random variables with a marginal distribution p_X that incorporate the distributions of the ancestors. *Realized* nodes $\bullet$ are random variables that have been assigned a concrete value via sampling or observation.

The evaluation function *value(e, g)* forces the realization by sampling of all the random variables referenced in e to produce a concrete value. Similarly, the function *observe(X, v, g)* realizes a variable X with a given observation v . The realization of a random variable comprises three steps: (1) compute the distribution $p(x)$ by recursively marginalizing the parents from a root node, (2) sample a value, or use the observation, and (3) use the concrete value to update the children and condition the parent which removes the dependencies.

The function *assume(μ, g)* adds a new node to the graph and is defined case by case on the shape of the symbolic term μ . If there is a conjugacy relationship between μ and a random variable Y present in the graph, e.g., $\mu = Bernoulli(Y)$ with $Y \sim Beta(\alpha, \beta)$, a new initialized node $X \sim \mu$ is added as a child of Y . Otherwise, since symbolic computation is not possible, dependencies are broken by realizing the random variables that appear in μ , e.g., $\mu' = Bernoulli(value(Y, g))$, and $X \sim \mu'$ is added as a new root node.

Inference. The inference scheme is similar to the particle filter. At each step, the inference draws N states from σ to execute the transition function. For each particle, execution starts with the graph computed at the previous step and returns a pair of symbolic terms (result, state), the particle

Figure 9. One step of the robot example of Figure 1 with SDS. Plain arrows represent dependencies and dotted arrows represent pointers at runtime. The `sample` statement adds the initialized nodes `x` (b). The `observe` statement adds the initialized node `a` (c), triggers the marginalizations of `x` (d) and `a` (e), and assigns to `a` its observed value (f). When the state is updated, the value `x` becomes `pre x`. The previous values are not referenced anymore and can be removed (g).

weight, and the updated graph. The function $distribution(e, g)$ returns the distribution corresponding to the expression e without altering the graph. Results are then aggregated in a mixture distribution (concrete values are lifted to Dirac distribution) where the distribution d_i operates on the value component of U and we use the pair (symbolic term, graph) computed by the transition function for the distribution of state. This distribution is then split into the distribution of returned values and the distribution of next states.

$$\begin{aligned} \llbracket \text{infer}(\text{fun } s \rightarrow e, \sigma) \rrbracket_Y = \\ \text{let } \mu = \lambda U. \sum_{i=1}^N \text{let } s_i, g_i = \text{draw}(\llbracket \sigma \rrbracket_Y) \text{ in} \\ \text{let } (e_i, s'_i), w_i, g'_i = \llbracket \text{fun } s \rightarrow e \rrbracket_{Y, 1, g_i}(s_i) \text{ in} \\ \text{let } d_i = \text{distribution}(e_i, g'_i) \text{ in} \\ \frac{w_i}{\sum w_i} * d_i(\pi_1(U)) * \delta_{s'_i, g'_i}(\pi_2(U)) \\ \text{in } (\pi_{1*}(\mu), \pi_{2*}(\mu)) \end{aligned}$$

5.3 Streaming Delayed Sampling

As illustrated in Section 2.2, a notable challenge with the traditional delayed sampling algorithm is that graph grows linearly in the number of samples. In the original formulation of delayed sampling [27], graph edges are only removed when a node is realized. All nodes that have been neither sampled nor observed are thus kept in the graph even if they are no longer referenced by the program. In a reactive programming context, such an implementation can consume unbounded memory.

Bounded Delayed Sampling. A simple mitigation is to limit the scope of symbolic computations to one time step and discard the graph at the end of each time step. We call this inference technique *bounded delayed sampling* (BDS).

BDS performs symbolic computations during the execution of a time step, and whenever possible, delays the sampling until the end of the instant. Like the particle filter, BDS guarantees a bounded-memory execution. For each particle, the size of the graph is bounded by the number of variables introduced during a time step, which by construction, is bounded for any valid ProbZelus program.

Streaming Delayed Sampling. Compared to the original delayed sampling algorithm, BDS loses the ability to perform symbolic computations using variables defined at different time steps. This can result in a significant loss of precision

for models with inter-steps dependencies such as the robot example of Figure 1. To adapt delayed sampling to streaming settings while keeping its maximum accuracy, we designed a delayed sampler that is *pointer-minimal* where nodes that are no longer referenced by the program can be eventually removed. We call this inference *streaming delayed sampling* (SDS). SDS enables partial exact inference in bounded memory for a large class of models.

In the original implementation of delayed sampling, graph nodes need to access their parents and children. Marginalization requires access to the parent to incorporate the ancestor distribution. Realization requires access to both the parent and the children of a node to update their respective distributions with the concrete value assigned to the node.

In the pointer minimal implementation, initialized nodes only keep a pointer to their parent to follow the ancestor chain during marginalization and marginalized nodes only keep a pointer to a marginalized child (Delayed Sampling imposes that a node always has at most one marginalized child). Compared to the original implementation, marginalized nodes only keep track of one child, and marginalization turns backward pointers to the parent node into forward pointers to the marginalized child. Note that this implementation prevents updating the children when the parent is realized, and prevents conditioning a parent when a child is realized. Instead, when marginalizing a node, the sampler first checks if the parent is realized to apply the update. Symmetrically, to realize a node, the sampler first checks if the children are realized and, if necessary, conditions the distribution before assigning the concrete value.

Figure 9 shows the evolution of the graph during one step for the robot example of Figure 1. At the end of the step, the value of `pre x` is updated. The previous value is not referenced anymore by the program and the node can be removed from the graph. In the original implementation, backward pointers between marginalized nodes prevent the collection (see Figure 4).

Limitations. With SDS, models like the robot example that only maintain bounded chains of dependencies between variables are guaranteed to be executed in bounded memory. The class of models that can be executed in bounded memory

with our pointer-minimal implementation thus already comprises state-space models like Kalman filters, and models for learning unknown constant parameters from a series of observations (e.g., computing the bias of a coin from a succession of flips) where variables introduced at each step are immediately realized.

However, unbounded chains can still be formed if the program keeps a reference to a constant variable that is never realized. In the following example, at each step, a new variable `x` is added as a child of `pre x` and then marginalized for the observation. But `p1` keeps a reference to the initial variable `i` which is never realized and thus forms an unbounded chain between `i` and `x`.

```
let proba p1 (xo, obs) = (i, o) where
  rec init i = sample (gaussian (xo, 1.))
  and x = sample (gaussian (i -> pre x, 1.))
  and () = observe (gaussian (x, 1.), obs)
```

In addition, in ProbZelus, at each step the inference returns a snapshot of the current distribution without forcing the realization of any node in the graph. Compared to the original delayed sampling implementation, initialized nodes can be inspected without being realized. It is thus possible to form unbounded chains of initialized nodes which cannot be pruned even when nodes are no longer referenced in the program due to the backward pointers to the parent in initialized nodes. In the following example, at each step, the variable `x` is added as a child of `pre x`, but without observation these variables remain initialized for ever.

```
let proba p2 (xo) = x where
  rec x = sample (gaussian (xo -> pre x, 1.))
```

To mitigate these issues, we can force the realization of trailing nodes at each step as in *bounded delayed sampling* or use a sliding window. Alternatively, the `value (eval)` function is available to the programmer and can be used to implement any strategy to force the evaluation of the nodes. For instance, the previous example can be adapted to execute in bounded memory:

```
let proba p2' (xo) = x where
  rec x = sample (gaussian (xo -> pre x, 1.))
  and _ = eval (xo -> pre x)
```

6 Evaluation

We next evaluate the performance of ProbZelus on a set of benchmarks that illustrate multiple aspects of the language: inferring fixed parameters from observations, online trajectory estimation, inference-in-the-loop. For these examples, we compare the accuracy and the latency cost of the three inference techniques: PF, BDS, and SDS. Appendix E details our implementation as an extension of the Zelus compiler.

Table 1. Benchmarks with: inference of *fixed* parameters from observations, estimation of a *moving* state (state-space model), and *inference-in-the-loop* (IITL).

	Fixed	Moving	IITL	Metric
Beta-Bernoulli	✓			MSE
Gaussian-Gaussian	✓			MSE
Kalman-1D		✓		MSE
Outlier	✓	✓		MSE
Robot		✓	✓	LQR
SLAM	✓	✓	✓	MSE
MTT		✓		MOTA*

Benchmarks. The models used in the experiments are summarized in Table 1 (a detailed description along with the code of the benchmarks is given in Appendix F.1). Two models infer fixed parameters from a series of observations. Beta-Bernoulli estimates the parameter of a Bernoulli distribution from a series of binary observations (e.g., the bias of a coin). Gaussian-Gaussian estimates the mean and variance of a Gaussian distribution from a series of observations. The accuracy metric is the *Mean Squared Error* (MSE) of the inferred parameters compared to their exact values.

Two models infer the state of a moving agent from noisy observations. Kalman-1D is a one-dimensional Kalman filter that models an agent that estimates its trajectory from noisy observations. Outlier adapted from [26] models the same situation as Kalman-1D, but the sensor occasionally produces invalid readings. This models infer both the trajectory of the agent, and the bias of the sensor. The accuracy metric is the *Mean Squared Error* (MSE) of the inferred trajectory compared to the exact positions.

Two models use inference-in-the-loop (IITL). Robot is the robot example of Section 2 and the accuracy metric is the LQR loss. SLAM (*Simultaneous Localization and Mapping*) adapted from [16] models an agent that estimates its position and a map of its environment. In this simplified version, the agent moves in a one-dimensional grid where each cell is either black or white. The robot’s wheels may slip causing the robot to unknowingly stay in place (noisy motion), and the sensor is not perfect and may accidentally report the wrong color (noisy observations). At each step the robot uses the inferred position to decide its next move. The accuracy metric is the MSE of both the position and the map.

MTT (*Multi-Target Tracker*) adapted from [28] is a model where there are a variable number of targets with linear-Gaussian motion models with a state space of 2D position and velocity, producing linear-Gaussian measurements of the position at each time step. Targets randomly appear according to a Poisson process and each disappear with fixed probability at each step. Measurements do not identify which target they came from, and “clutter” measurements that come not from targets but from some underlying distribution add

Figure 10. Execution time comparison when 90% of 1000 runs reach an accuracy similar to the baseline (median accuracy of SDS with 1000 particles) after 500 steps. The number of particles required to reach this accuracy is shown on top of the bars. The error bars show the 10th and 90th percentiles.

to observations, complicating inference of which measurements are associated to which targets. The accuracy metric is expected $MOTA^* = (1/MOTA) - 1$ where $MOTA \in [0, 1]$ is the *Multiple Object Tracking Accuracy* [3].

Experimental Setup. All the experiments were run on a server with 32 CPUs (2.60 GHz) and 128 GB memory. We ran all the benchmarks for 500 steps. In all cases, the inference runs in bounded memory (see Appendix F.3).

For each algorithm, we evaluated how much time it required to achieve 90% of runs close to a loss target (out of 1000 runs total):

$$|\log(P_{90\%}(loss)) - \log(loss_{target})| < 0.5.$$

For each benchmark, the baseline is the median loss of SDS at 1000 particles as $loss_{target}$ for that benchmark. We measured the number of particles required to achieve this loss, and then measured the total execution time at this particle count for 500 steps (in Appendix F.2 we also evaluate loss and step latency across a fixed range of particle counts).

Results. Figure 10 shows the results. The height of each bar is the median total execution time, and the error bars are 90% and 10% quantiles, aggregated over 1000 runs. Each bar is labeled with the minimum number of particles required to achieve the accuracy threshold, accurate to 1.5 significant digits (100, 150, 200, 250, ...). We observe that SDS is able to compute an exact solution for Beta-Bernoulli, Kalman-1D, and Robot. In all these examples 1 particle is already enough to reach the target accuracy. Overall, the results show that the number of particles required to reach the desired accuracy with PF implies a significant slowdown compared to SDS. Moreover, the SLAM and MTT benchmarks show that, in some cases, PF is not an option: the target accuracy was not reached with 15,000 and 2,500 particles, respectively, at which point PF was already 10 times slower than SDS and we stopped the experiments.

As expected, BDS performance numbers are between those of PF and SDS. At worst, when there is no possible intra-step

symbolic computations (e.g., Beta-Bernoulli), BDS behaves like a particle filter and requires as many particles as PF. At best, BDS performs as well as SDS (e.g., Outlier).

Additionally, Figure 10 also shows that for a given number of particles, the overhead induced by managing the delayed sampling graph is significant. Compared to BDS and SDS, depending on the benchmark, it is possible to use 2 to 4 times as many particles for PF with the same execution time. However, this is not enough to match the gain in accuracy.

Alternative Baselines. The results presented in Figure 10 do not quantify the speedup of SDS on the SLAM and MTT benchmarks because the other inference algorithms time out. To evaluate speedups on these two benchmarks, we used PF as an alternative baseline instead of SDS. Figure 11 presents the execution time of PF, BDS, and SDS to reach a loss close to the median of PF with 2000 and 4000 particles.

We observe that SDS requires a much smaller number of particles to reach similar accuracy which translates into speedups ranging from 10^1 (MTT-2000) to 10^4 (SLAM-4000). BDS requires either a similar or smaller numbers of particles. But the overhead introduced by the graph manipulations mostly translates in slowdowns compared to PF.

7 Related Work

Probabilistic Programming. Over the last few years there has been a growing interest on probabilistic programming languages. Some languages like BUGS [25], Stan [9], or Augur [21] offer optimized inference technique for a constrained subset of models. Other languages like WebPPL [18], Edward [37], Pyro [6], or Birch [28] focus on expressivity allowing the specification of arbitrary complex models. Compared to these languages, ProbZelus can be used to program *reactive models* that typically do not terminate, and inference can be run in parallel with deterministic components that interact with an environment.

Figure 11. Execution time comparison with two different baselines: median accuracy of PF with 2000 and 4000 particles, respectively.

Reactive Languages with Uncertainty. Lutin is a language for describing non-deterministic reactive systems for testing and simulation [32], but while Lutin supports weighted sampling to describe constrained random scenarios, it does not support inference. ProPL [30] is a language to describe probabilistic models for process that evolve over a period of time. This language also extends a probabilistic language with a notion of processes that can be composed in parallel, but compared to ProbZelus, ProPL focuses on a constrained class of models that can be interpreted as *Dynamic Bayesian Networks* (DBN), and relies on standard DBN inference techniques. In the same vein, CTPPL [31] is a language to describe continuous-time processes where the amount of time taken by a sub-process can be specified by a probabilistic model. These models cannot be expressed in ProbZelus which relies on the synchronous model of computation. It would be interesting to investigate how to extend ProbZelus to continuous-time models based on Zelus’ support for ordinary differential equations (ODE) [7].

Inference. Researchers have proposed streaming inference algorithms, including variational [8], or sampling-based [16, 19] approaches. Popular languages like Stan, Edward, or Pyro, offer support to stream data through the model during inference to handle large datasets. However, compared to ProbZelus, the model must be defined a priori and does not evolve during the inference.

The Anglican and Birch probabilistic programming languages support delayed sampling [27]. These languages do not support streaming inference or reactive programming. Again, their interfaces only support inference on a complete probabilistic model.

8 Conclusion

Modeling uncertainty is a primary element of control systems for tasks that operate under the assumption of a probabilistic model of their environment (e.g., object tracking).

While synchronous languages have developed as a prominent way to develop control applications, to date there has been limited work in these languages on programming language support for modeling uncertainty.

In this paper we present ProbZelus, the first synchronous probabilistic programming language that lifts emerging abstractions for probabilistic programming into the reactive setting thus enabling *inference-in-the-loop*. Moreover, our streaming delayed sampling algorithm provides efficient semi-symbolic inference while still satisfying a key requirement of control applications in that they must execute with bounded resources.

Our results demonstrate that ProbZelus enables us to write, in the very same source, a deterministic model for the control software and a probabilistic model for its behavior and environment with complex interactions between the two.

Acknowledgments

This work was supported in part by the MIT-IBM Watson AI Lab and the Office of Naval Research (ONR N00014-17-1-2699).

References

- [1] Guillaume Baudart, Louis Mandel, Eric Atkinson, Benjamin Sherman, Marc Pouzet, and Michael Carbin. 2020. Reactive Probabilistic Programming. *CoRR* abs/1908.07563 (2020).
- [2] 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.
- [3] Keni Bernardin and Rainer Stiefelhausen. 2008. Evaluating Multiple Object Tracking Performance: The CLEAR MOT Metrics. *EURASIP J. Image and Video Processing* 2008 (2008).
- [4] Gérard Berry. 1989. Real Time Programming: Special Purpose or General Purpose Languages. In *IFIP Congress*. North-Holland/IFIP, 11–17.
- [5] Dariusz Biernacki, Jean-Louis Colaço, Grégoire Hamon, and Marc Pouzet. 2008. Clock-directed modular code generation for synchronous data-flow languages. In *LCTES*. ACM, 121–130.
- [6] 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.
- [7] Timothy Bourke and Marc Pouzet. 2013. Zelus: a synchronous language with ODEs. In *HSCC*. ACM, 113–118.
- [8] Tamara Broderick, Nicholas Boyd, Andre Wibisono, Ashia C. Wilson, and Michael I. Jordan. 2013. Streaming Variational Bayes. In *NIPS*. 1727–1735.
- [9] Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. Stan: A probabilistic programming language. *J. Statistical Software* 76, 1 (2017), 1–37.
- [10] Paul Caspi. 1992. Clocks in Dataflow Languages. *Theor. Comput. Sci.* 94, 1 (1992), 125–140.
- [11] 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, 1–21.
- [12] Jean-Louis Colaço, Grégoire Hamon, and Marc Pouzet. 2006. Mixing signals and modes in synchronous data-flow systems. In *EMSOFT*. ACM, 73–82.

- [13] Jean-Louis Colaço, Bruno Pagano, and Marc Pouzet. 2017. SCADE 6: A formal language for embedded critical software development (invited paper). In *TASE*. IEEE Computer Society, 1–11.
- [14] Jean-Louis Colaço and Marc Pouzet. 2004. Type-based initialization analysis of a synchronous dataflow language. *Int. J. Softw. Tools Technol. Transf.* 6, 3 (2004), 245–255.
- [15] 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.
- [16] Arnaud Doucet, Nando de Freitas, Kevin P. Murphy, and Stuart J. Russell. 2000. Rao-Blackwellised Particle Filtering for Dynamic Bayesian Networks. In *UAI*. Morgan Kaufmann, 176–183.
- [17] Timon Gehr, Sasa Misailovic, and Martin T. Vechev. 2016. PSI: Exact Symbolic Inference for Probabilistic Programs. In *CAV (1) (Lecture Notes in Computer Science)*, Vol. 9779. Springer, 62–83.
- [18] Noah D. Goodman and Andreas Stuhlmüller. 2014. The Design and Implementation of Probabilistic Programming Languages. <http://dippl.org>. Accessed April 2020.
- [19] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. 1993. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F - Radar and Signal Processing* 140, 2, 107–113.
- [20] N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud. 1991. The Synchronous Dataflow Programming Language LUSTRE. *Proc. IEEE* 79, 9 (September 1991), 1305–1320.
- [21] Daniel Huang, Jean-Baptiste Tristan, and Greg Morrisett. 2017. Compiling Markov chain Monte Carlo algorithms for probabilistic modeling. In *PLDI*. ACM, 111–125.
- [22] Gilles Kahn. 1974. The Semantics of a Simple Language for Parallel Programming. In *IFIP Congress*. North-Holland, 471–475.
- [23] Daniel Lundén. 2017. *Delayed sampling in the probabilistic programming language Anglican*. Master's thesis. KTH Royal Institute of Technology.
- [24] Daniel Lundén, David Broman, Fredrik Ronquist, and Lawrence M. Murray. 2018. Automatic Alignment of Sequential Monte Carlo Inference in Higher-Order Probabilistic Programs. *CoRR* abs/1812.07439 (2018).
- [25] David Lunn, David Spiegelhalter, Andrew Thomas, and Nicky Best. 2009. The BUGS project: Evolution, critique and future directions. *Statistics in medicine* 28, 25 (2009), 3049–3067.
- [26] Thomas P. Minka. 2001. Expectation Propagation for approximate Bayesian inference. In *UAI*. Morgan Kaufmann, 362–369.
- [27] Lawrence M. Murray, Daniel Lundén, Jan Kudlicka, David Broman, and Thomas B. Schön. 2018. Delayed Sampling and Automatic Rao-Blackwellization of Probabilistic Programs. In *AISTATS (Proceedings of Machine Learning Research)*, Vol. 84. PMLR, 1037–1046.
- [28] Lawrence M. Murray and Thomas B. Schön. 2018. Automated learning with a probabilistic programming language: Birch. *Annual Reviews in Control* 46 (2018), 29–43.
- [29] Praveen Narayanan, Jacques Carette, Wren Romano, Chung-chieh Shan, and Robert Zinkov. 2016. Probabilistic Inference by Program Transformation in Hakaru (System Description). In *FLOPS (Lecture Notes in Computer Science)*, Vol. 9613. Springer, 62–79.
- [30] Avi Pfeffer. 2005. Functional Specification of Probabilistic Process Models. In *AAAI*. AAAI Press / The MIT Press, 663–669.
- [31] Avi Pfeffer. 2009. CTPPL: A Continuous Time Probabilistic Programming Language. In *IJCAI*. 1943–1950.
- [32] Pascal Raymond, Yvan Roux, and Erwan Jahier. 2008. Lutin: A Language for Specifying and Executing Reactive Scenarios. *EURASIP Journal of Embedded Systems* 2008 (2008).
- [33] Daniel Ritchie, Andreas Stuhlmüller, and Noah D. Goodman. 2016. C3: Lightweight Incrementalized MCMC for Probabilistic Programs using Continuations and Callsite Caching. In *AISTATS (JMLR Workshop and Conference Proceedings)*, Vol. 51. JMLR.org, 28–37.
- [34] Eduardo D Sontag. 2013. *Mathematical control theory: deterministic finite dimensional systems*. Vol. 6. Springer Science & Business Media.
- [35] Sam Staton. 2017. Commutative Semantics for Probabilistic Programming. In *ESOP (Lecture Notes in Computer Science)*, Vol. 10201. Springer, 855–879.
- [36] David Tolpin, Jan-Willem van de Meent, Hongseok Yang, and Frank D. Wood. 2016. Design and Implementation of Probabilistic Programming Language Anglican. In *IFL*. ACM, 6:1–6:12.
- [37] Dustin Tran, Matthew D. Hoffman, Rif A. Saurous, Eugene Brevdo, Kevin Murphy, and David M. Blei. 2017. Deep Probabilistic Programming. In *ICLR (Poster)*. OpenReview.net.
- [38] Yi Wu, Lei Li, Stuart J. Russell, and Rastislav Bodík. 2016. Swift: Compiled Inference for Probabilistic Programming Languages. In *IJCAI*. IJCAI/AAAI Press, 3637–3645.