Probabilistic Programming Inference via Intensional Semantics

Simon Castellan[1] and Hugo Paquet[2]

[1] Imperial College London, United Kingdom
`simon.castellan@phis.me`
[2] University of Cambridge, United Kingdom
`hugo.paquet@cl.cam.ac.uk`

Abstract. We define a new denotational semantics for a first-order probabilistic programming language in terms of *probabilistic event structures*. This semantics is *intensional*, meaning that the interpretation of a program contains information about its behaviour throughout execution, rather than a simple distribution on return values. In particular, occurrences of sampling and conditioning are recorded as explicit events, partially ordered according to the data dependencies between the corresponding statements in the program.

This interpretation is *adequate*: we show that the usual measure-theoretic semantics of a program can be recovered from its event structure representation. Moreover it can be leveraged for MCMC inference: we prove correct a version of single-site Metropolis-Hastings with *incremental recomputation*, in which the proposal kernel takes into account the semantic information in order to avoid performing some of the redundant sampling.

Keywords: Probabilistic programming · Denotational Semantics · Event structures · Bayesian inference.

1 Introduction

Probabilistic programming languages [7] were put forward as promising tools for practitioners of Bayesian statistics. By extending traditional programming languages with primitives for sampling and conditioning, they allow the user to express a wide class of statistical models, and provide a simple interface for encoding inference problems. Although the subject of active research, it is still notoriously difficult to design inference methods for probabilistic programs which perform well for the full class of expressible models.

One popular inference technique, proposed by Wingate et al. [21], involves adapting well-known *Monte-Carlo Markov chain* methods from statistics to probabilistic programs, by manipulating *program traces*. One such method is the Metropolis-Hastings algorithm, which relies on a key *proposal* step: given a program trace x (a sequence $x_1, \ldots, x_n$ of random choices with their likelihood), a proposal for the *next* trace sample is generated by choosing $i \in \{1, \ldots, n\}$

uniformly, resampling x_i, and then continuing to execute the program, only performing additional sampling for those random choices not appearing in x. The variables already present in x are not resampled: only their likelihood is updated according to the new value of x_i. Likewise, some conditioning statements must be re-evaluated in case the corresponding weight is affected by the change to x_i.

Observe that there is some redundancy in this process, since the updating process above will only affect variables and observations when their density directly depends on the value of x_i. This may significantly affect performance: to solve an inference problem one must usually perform a large number of proposal steps. To overcome this problem, some recent implementations, notably [11, 25], make use of *incremental recomputation*, whereby some of the redundancy can be avoided via a form of static analysis. However, as pointed out by Kiselyov [12], establishing the correctness of such implementations is tricky.

Here we address this by introducing a theoretical framework in which to reason about data dependencies in probabilistic programs. Specifially, our first contribution is to define a *denotational semantics* for a first-order probabilistic language, in terms of graph-like structures called *event structures* [22]. In event structures, computational events are partially ordered according to the dependencies between them; additionally they can be equipped with quantitative information to represent probabilistic processes [15, 23]. This semantics is *intensional*, unlike most existing semantics for probabilistic programs, in which the interpretation of a program resembles a probability distribution on output values. We relate our approach to a measure-theoretic semantics [17] through an *adequacy* result.

Our second contribution is the design of a Metropolis-Hastings algorithm which exploits the event structure representation of the program at hand. Some of the redundancy in the proposal step of the algorithm is avoided by taking into account the extra dependency information given by the semantics. We provide a proof of correctness for this algorithm, and argue that an implementation is realistically achievable: we show in particular that all graph structures involved and the associated quantitative information admit a finite, concrete representation.

Outline of the paper. In § 2 we give a short introduction to probabilistic programming. We define our main language of study and its measure-theoretic semantics. In § 3.1, we introduce MCMC methods and the Metropolis-Hastings algorithm in the context of probabilistic programming. We then motivate the need for intensional semantics in order to capture data dependency. In § 4 we define our interpretation of programs and prove adequacy. In § 5 we define an updated version of the algorithm, and prove its correctness. We conclude in § 6.

The proofs of the statements are detailed in the appendix.

2   Probabilistic programming

In this section we motivate the need for capturing data dependency in probabilistic programs. Let us start with a brief introduction to probabilistic programming – a more comprehensive account can be found in [7].

2.1 Conditioning and posterior distribution

Let us introduce the problem of inference in probabilistic programming from the point of view of programming language theory.

We consider a first-order programming language enriched with a real number type $\mathbb{R}$ and a primitive `sample` for drawing random values from a given family of standard probability distributions. The language is idealised — but it is assumed that an implementation of the language comprises built-in sampling procedures for those standard distributions. Thus, repeatedly running the program `sample Uniform` $(0,1)$ returns a sequence of values approaching the true uniform distribution on $[0,1]$.

Via other constructs in the language, standard distributions can be combined, as shown in the following example program of type $\mathbb{R}$:

```
let x = sample Uniform(0, 1) in
let y = sample Gaussian(x, 2) in
x + y
```

Here the output will follow a probability distribution built out of the usual uniform and Gaussian distributions. Many probabilistic programming languages will offer more general programming constructs: conditionals, recursion, higher-order functions, data types, *etc.*, enabling a wide range of distributions to be expressed in this way. Such a program is sometimes called a *generative model*.

Conditioning. The process of conditioning involves rescaling the distribution associated with a generative model, so as to reflect some bias. Going back to the example above, say we have made some external measurement indicating that $y = 0$, but we would like to account for possible noise in the measurement using another Gaussian. To express this we modify the program as follows:

```
let x = sample Uniform (0, 1) in
let y = sample Gaussian (x, 2) in
observe y (Gaussian (0, 0.01));
x + y;
```

The purpose of the `observe` statement is to increase the occurrence of executions in which y is close to 0; the original distribution, known as the **prior**, must be updated accordingly. The probabilistic weight of each execution is multiplied by an appropriate *score*, namely the **likelihood** of the current value of y in the Gaussian distribution with parameters $(0, 0.01)$. (This is known as a *soft constraint*. Conditioning via *hard constraints*, *i.e.* only giving a nonzero score to executions where y is exactly 0, is not practically feasible.)

The language studied here does not have an `observe` construct, but instead an explicit `score` primitive; this appears already in [18, 17]. So the third line in the program above would instead be `score(pdf-Gaussian (0, 0.01) (y))` where `pdf-Gaussian (0, 0.01)` is the *density* function of the Gaussian distribution. The resulting distribution is not necessarily normalised. We obtain the

posterior distribution by computing the normalising constant, following Bayes' rule:

$$\text{posterior} \propto \text{likelihood} \times \text{prior}.$$

This process is known as Bayesian inference and has ubiquitous applications. The difficulty lies in computing the normalising constant, which is usually obtained as an integral. Below we discuss *approximate* methods for sampling from the posterior distribution; they do not rely on this normalising step.

Measure theory. Because this work makes heavy use of probability theory, we start with a brief account of measure theory. A standard textbook for this is [1]. Recall that a **measurable space** is a set X equipped with a σ**-algebra** Σ_X: a set of subsets of X containing $\emptyset$ and closed under complements and countable unions. Elements of Σ_X are called **measurable sets**. A **measure** on X is a function $\mu : \Sigma_X \to [0, \infty]$, such that $\mu(\emptyset) = 0$ and, for any countable family $\{U_i\}_{i \in I}$ of measurable sets, $\mu(\bigcup_{i \in I} U_i) = \sum_{i \in I} \mu(U_i)$.

An important example is that of the set $\mathbb{R}$ of real numbers, whose σ-algebra $\Sigma_{\mathbb{R}}$ is generated by the intervals $[a, b)$, for $a, b \in \mathbb{R}$ (in other words, it is the smallest σ-algebra containing those intervals). The **Lebesgue measure** on $(\mathbb{R}, \Sigma_{\mathbb{R}})$ is the (unique) measure λ assigning $b - a$ to every interval $[a, b)$ (with $a \leq b$).

Given measurable spaces (X, Σ_X) and (Y, Σ_Y), a function $f : X \to Y$ is **measurable** if for every $U \in \Sigma_Y$, $f^{-1}U \in \Sigma_X$. A measurable function $f : X \to [0, \infty]$ can be *integrated*: given $U \in \Sigma_X$ the **integral** $\int_U f\, d\lambda$ is a well-defined element of $[0, \infty]$; indeed the map $\mu : U \mapsto \int_U f \mathrm{d}\lambda$ is a measure on X, and f is said to be a **density** for μ. The precise definition of the integral is standard but slightly more involved; we omit it.

We identify the following important classes of measures: a measure μ on (X, Σ_X) is a **probability measure** if $\mu(X) = 1$. It is **finite** if $\mu(X) < \infty$, and it is **s-finite** if $\mu = \sum_{i \in I} \mu_i$, a pointwise, countable sum of finite measures.

We recall the usual product and coproduct constructions for measurable spaces and measures. If $\{X_i\}_{i \in I}$ is a countable family of measurable spaces, their **product** $\prod_{i \in I} X_i$ and **coproduct** $\coprod_{i \in I} X_i = \bigcup_{i \in I} \{i\} \times X_i$ as sets can be turned into measurable spaces, where:

- $\Sigma_{\prod_{i \in I} X_i}$ is generated by $\{\prod_{i \in I} U_i \mid U_i \in \Sigma_{X_i} \text{ for all } i\}$, and
- $\Sigma_{\coprod_{i \in I} X_i}$ is generated by $\{\{i\} \times U_i \mid i \in I \text{ and } U_i \in \Sigma_{X_i}\}$.

The measurable spaces in this paper all belong to a well-behaved subclass: call (X, Σ_X) a **standard Borel space** if it either countable and discrete (*i.e.* all $U \subseteq X$ are in Σ_X), or measurably isomorphic to $(\mathbb{R}, \Sigma_{\mathbb{R}})$. Note that standard Borel spaces are closed under countable products and coproducts, and that in a standard Borel space all singletons are measurable.

2.2   A first-order probabilistic programming language

We consider a first-order, call-by-value language $\mathcal{L}$ with types

$$A, B ::= 1 \mid \mathbb{R} \mid \coprod_{i \in I} A_i \mid \prod_{i \in I} A_i$$

where I ranges over nonempty countable sets. The types denote measurable spaces in a natural way: $[\![1]\!]$ is the singleton space, and $[\![\mathbb{R}]\!] = (\mathbb{R}, \Sigma_{\mathbb{R}})$. Products and coproducts are interpreted via the corresponding measure-theoretic constructions: $[\![\prod_{i \in I} A_i]\!] = \prod_{i \in I} [\![A_i]\!]$ and $[\![\coprod_{i \in I} A_i]\!] = \coprod_{i \in I} [\![A_i]\!] = \bigcup_{i \in I} \{i\} \times [\![A_i]\!]$. Moreover, each measurable space $[\![A]\!]$ has a canonical measure $\mu_{[\![A]\!]} : \Sigma_{[\![A]\!]} \to \mathbb{R}$, induced from the Lebesgue measure on $\mathbb{R}$ and the Dirac measure on $[\![1]\!]$ via standard product and coproduct measure constructions.

The terms of $\mathcal{L}$ are given by the following grammar:

$$M, N ::= () \mid M; N \mid f \mid \texttt{let } a \texttt{ = } M \texttt{ in } N \mid x$$
$$\mid (M_i)_{i \in I} \mid \texttt{case } M \texttt{ of } \{(i, x) \Rightarrow N_i\}_{i \in I}$$
$$\mid \texttt{sample } d \texttt{ } (M) \mid \texttt{score } M$$

and we use standard syntactic sugar to manipulate integers and booleans: $\mathbb{B} = 1 + 1$, $\mathbb{N} = \sum_{i \in \omega} 1$, and constants are given by the appropriate injections. Conditionals and sequencing can be expressed in the usual way: $\texttt{if } M \texttt{ then } N_1 \texttt{ else } N_2 = \texttt{case } M \texttt{ of } \{(i, \_) \Rightarrow N_i\}_{i \in \{1,2\}}$, and $M; N = \texttt{let } a \texttt{ = } M \texttt{ in } N$, where a does not occur in N. In the grammar above:

- f ranges over measurable functions $[\![A]\!] \to [\![B]\!]$, where A and B are types;
- d ranges over a family of *parametric distributions* over the reals, *i.e.* measurable functions $\mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$, for some $n \in \mathbb{N}$, such that for every $\mathbf{r} \in \mathbb{R}^n$, $\int d(\mathbf{r}, -) = 1$. For the purposes of this paper we ignore all issues related to invalid parameters, arising from *e.g.* a call to $\texttt{gaussian}$ with standard deviation $\sigma = 0$. (An implementation could, say, choose to behave according to an alternative distribution in this case.)

The typing rules are as follows:

$$\frac{\Gamma \vdash M : A \qquad \Gamma, a : A \vdash N : B}{\Gamma \vdash \texttt{let } a \texttt{ = } M \texttt{ in } N : B} \qquad \frac{\Gamma \vdash M : \mathbb{R}^n \qquad d : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}}{\Gamma \vdash \texttt{sample } d \texttt{ } (M) : \mathbb{R}}$$

$$\frac{\Gamma \vdash M : \mathbb{R}}{\Gamma \vdash \texttt{score } M : 1} \qquad \frac{}{\Gamma, a : A \vdash a : A} \qquad \frac{}{\Gamma \vdash () : \mathbf{1}}$$

$$\frac{\Gamma \vdash M : \sum_{i \in I} A_i \qquad \Gamma, x : A_i \vdash N_i : C}{\Gamma \vdash \texttt{case } M \texttt{ of } \{(i, x) \Rightarrow N_i\}_{i \in I} : C} \qquad \frac{\Gamma \vdash M_i : A_i}{\Gamma \vdash (M_i)_{i \in I} : \prod_{i \in I} A_i}$$

$$\frac{f : [\![A]\!] \to [\![B]\!] \text{ measurable} \qquad \Gamma \vdash M : A}{\Gamma \vdash f \text{ } M : B}$$

Among the measurable functions f, we point out the following of interest:

- The usual product projections $\pi_i : [\![\prod_{i \in I} A_i]\!] \to [\![A_i]\!]$ and coproduct injections $\iota_i : [\![A_i]\!] \to [\![\coprod_{i \in I} A_i]\!]$;
- The operators $+, \times : \mathbb{R}^2 \to \mathbb{R}$,
- The tests, eg. $\geq 0 : [\![\mathbb{R}]\!] \to [\![\mathbb{B}]\!]$,
- The constant functions $1 \to A$ of the form $() \mapsto a$ for some $a \in [\![A]\!]$.

Examples for d include $\texttt{uniform} : \mathbb{R}^2 \times \mathbb{R} \to \mathbb{R}$, $\texttt{gaussian} : \mathbb{R}^2 \times \mathbb{R} \to \mathbb{R}$, ...

2.3   Measure-theoretic semantics of programs

We now define a semantics of probabilistic programs using the measure-theoretic concept of *kernel*, which we define shortly. The content of this section is not new: using kernels as semantics for probabilistic was originally proposed in [13], while the (more recent) treatment of conditioning (`score`) via *s-finite* kernels is due to Staton [17]. Intuitively, kernels provide a semantics of open terms $\Gamma \vdash M : A$ as measures on $[\![A]\!]$ varying according to the values of variables in Γ.

Formally, a **kernel** from (X, Σ_X) to (Y, Σ_Y) is a function $k : X \times \Sigma_Y \to [0, \infty]$ such that for each $x \in X$, $k(x, -)$ is a measure, and for each $U \in \Sigma_Y$, $k(-, U)$ is measurable. (Here the σ-algebra $\Sigma_{[0,\infty]}$ is the restriction of that of $\mathbb{R} + \{\infty\}$.) We say k is **finite** (resp. **probabilistic**) if each $k(x, -)$ is a finite (resp. probability) measure, and it **s-finite** if it is a countable pointwise sum $\sum_{i \in I} k_i$ of finite kernels. We write $k : X \rightsquigarrow Y$ when k is an s-finite kernel from X to Y.

A term $\Gamma \vdash M : A$ will denote an s-finite kernel $[\![M]\!] : [\![\Gamma]\!] \rightsquigarrow [\![A]\!]$, where the context $\Gamma = x_1 : A_1, \ldots, x_n : A_n$ denotes the product of its components: $[\![\Gamma]\!] = [\![A_1]\!] \times \cdots \times [\![A_n]\!]$.

Notice that any measurable function $f : X \to Y$ can be seen as a *deterministic* kernel $f^\dagger : X \rightsquigarrow Y$. Given two s-finite kernels $k : A \rightsquigarrow B$ and $l : A \times B \rightsquigarrow C$, we define their composition $l \circ k : A \rightsquigarrow C$:

$$(l \circ k)(a, X) = \int_{b \in B} l((a, b), C) \times k(a, \mathrm{d}b).$$

Staton [17] proved that $l \circ k$ is a s-finite kernel.

The interpretation of terms is defined by induction:

- $[\![()]\!]$ is the lifting of $[\![\Gamma]\!] \to 1 : x \mapsto ()$.
- $[\![\texttt{let } a \texttt{ = } M \texttt{ in } N]\!]$ is $[\![N]\!] \circ [\![M]\!]$
- $[\![f\, M]\!] = f^\dagger \circ [\![M]\!]$
- $[\![a]\!](x, X) = \delta_x(X)$, the Dirac distribution $\delta_x(X) = 1$ if $x \in X$ and zero otherwise.
- $[\![\texttt{sample } d \texttt{ } (M)]\!] = \texttt{sam} \circ [\![M]\!]$ where $\texttt{sam}_d : \mathbb{R}^n \rightsquigarrow \mathbb{R}$ is given by $\texttt{sam}_d(\mathbf{r}, X) = \int_{x \in X} d(\mathbf{r}, x)\mathrm{d}x$.
- $[\![\texttt{score } M]\!] = \texttt{sco} \circ [\![M]\!]$ where $\texttt{sco} : [\![R]\!] \to [\![1]\!]$ is $\texttt{sco}(r, X) = r \cdot \delta_{()}(X)$.
- $[\![(M_i)_{i \in I}]\!](\gamma, \prod_{i \in I} X_i) = \prod_{i \in I} [\![M_i]\!](\gamma, X_i)$: this is well-defined since the $\prod X_i$ generate the measurable sets of the product space.
- $[\![\texttt{case } M \texttt{ of } \{(i, x) \Rightarrow N_i\}_{i \in I}]\!] = \texttt{coprod} \circ [\![M]\!]$ where $\texttt{coprod} : \Gamma \times [\![\coprod_{i \in I} A_i]\!] \rightsquigarrow [\![B]\!]$ maps $(\gamma, \{i\} \times X)$ to $[\![N_i]\!](\gamma, X)$.

We observe that when M is a program making no use of conditioning (*i.e.* a generative model), the kernel $[\![M]\!]$ is probabilistic:

Lemma 1. *For $\Gamma \vdash M : A$ without scores, $[\![M]\!](\gamma, [\![A]\!]) = 1$ for each $\gamma \in [\![\Gamma]\!]$.*

2.4   Exact inference

Note that a kernel $1 \rightsquigarrow [\![A]\!]$ is the same as a measure on $[\![A]\!]$. Given a closed program $\vdash M : A$, the measure $[\![M]\!]$ is a combination of the prior (occurrences of

sample) and the likelihood (score). Because score can be called on arbitrary arguments, it may be the case that the measure of the total space (that is, the coefficient $[\![M]\!]([\![A]\!])$, often called the *model evidence*) is 0 or ∞.

Whenever this is *not* the case, $[\![M]\!]$ can be normalised to a probability measure, the posterior distribution. For every $U \in \Sigma_{[\![A]\!]}$,

$$\mathrm{norm}[\![M]\!](U) = \frac{[\![M]\!](U)}{[\![M]\!]([\![A]\!])}.$$

However, in many cases, this computation is intractable. Thus the goal of *approximate inference* is to approach $\mathrm{norm}[\![M]\!]$, the *true posterior*, using a well-chosen sequence of samples.

3 Approximate inference via intensional semantics

3.1 An introduction to approximate inference

In this section we describe the Metropolis-Hastings (MH) algorithm for approximate inference in the context of probabilistic programming. Metropolis-Hastings is a generic algorithm to sample from a probability distribution D on a measurable state space $\mathbb{X}$, of which we know the density $d : \mathbb{X} \to \mathbb{R}$ up to some normalising constant.

MH is part of a family of inference algorithms called *Monte-Carlo Markov chain*, in which the posterior distribution is approximated by a series of samples generated using a Markov chain.

Formally, the MH algorithm defines a Markov chain M on the state space $\mathbb{X}$, that is a probabilistic kernel $M : \mathbb{X} \rightsquigarrow \mathbb{X}$. The correctness of the MH algorithm is expressed in terms of convergence. It says that for almost all $x \in \mathbb{X}$, the distribution $M^n(x, \cdot)$ converges to D as n goes to infinity, where M^n is the n-iteration of M: $M \circ \ldots \circ M$. Intuitively, this means that iterated sampling from M gets closer to D with the number of iterations.

The MH algorithm is itself parametrised by a Markov chain, referred to as the **proposal kernel** $P : \mathbb{X} \rightsquigarrow \mathbb{X}$: for each sampled value $x \in \mathbb{X}$, a proposed value for the next sample is drawn according to $P(x, \cdot)$. Note that correctness only holds under certain assumptions on P.

The MH algorithm assumes that we know how to sample from P, and that its density is known, ie. there is a function $p : \mathbb{X}^2 \to \mathbb{R}$ such that $p(x, \cdot)$ is the density of the distribution $P(x, \cdot)$,

The MH algorithm. On an input state x, the MH algorithm samples from $P(x, \cdot)$ and gets a new sample x'. It then compares the likelihood of x and x' by computing an acceptance ratio $\alpha(x, x')$ which says whether the return state is x' or x. In pseudo-code, for an input state $x \in \mathbb{X}$:

1. Sample a new state x' from the distribution $P(x, \cdot)$

2. Compute the acceptance ratio of x' with respect to x:

$$\alpha(x, x') = \min\left(1, \frac{d(x') \times p(x, x')}{d(x) \times p(x', x)}\right)$$

3. With probability $\alpha(x, x')$, return the new sample x', otherwise return the input state x.

The formula for $\alpha(x, x')$ is known as the Hastings acceptance ratio and is key to the correctness of the algorithm.

Very little is assumed of P, which makes the algorithm very flexible; but of course the convergence rate may vary depending on the choice of P. We give a more formal description of MH in § 5.2.

Single-site MH and incremental recomputation. To apply this algorithm to probabilistic programming, we need a proposal kernel. Given a program M, the execution traces of M form a measurable set $\mathbb{X}_M$. In this setting the proposal is given by a kernel $\mathbb{X}_M \rightsquigarrow \mathbb{X}_M$.

A widely adopted choice of proposal is the *single-site proposal kernel* which, given a trace $x \in \mathbb{X}_M$, generates a new trace x' as follows:

1. Select uniformly one of the random choices s encountered in x.
2. Sample a new value for this instruction.
3. Re-execute the program M from that point onwards and with this new value for s, only ever resampling a variable when the corresponding instruction did not already appear in x.

Observe that there is some redundancy in this process: in the final step, the entire program has to be explored even though only a subset of the random choices will be re-evaluated. Some implementations of Trace MH for probabilistic programming make use of *incremental recomputation*.

We propose in this paper to statically compile a program M to an *event structure* G_M which makes explicit the probabilistic dependences between events, thus avoiding unnecessary sampling.

3.2   Capturing probabilistic dependencies using event structures

Consider the program depicted in Figure 1 in which we are interested in learning the parameters μ and σ of a Gaussian distribution from which we have observed two data points, say v_1 and v_2. For $i = 1, 2$ the function $f_i : \mathbb{R} \to \mathbb{R}$ expresses a soft constraint; it can be understood as indicating how much the sampled value of $\texttt{x}i$ matches the observed value v_i.

A *trace* of this program will be of the form

$$\mathsf{Sam}\,\mu \cdot \mathsf{Sam}\,\sigma \cdot \mathsf{Sam}\,x_1 \cdot \mathsf{Sam}\,x_2 \cdot \mathsf{Sco}\,(f_1\,x_1) \cdot \mathsf{Sco}\,(f_2\,x_2) \cdot \mathsf{Rtn}\,(\mu, \sigma),$$

for some μ, σ, x_1, and $x_2 \in \mathbb{R}$ corresponding to sampled values for variables $\texttt{mu}$, $\texttt{sigma}$, $\texttt{x1}$ and $\texttt{x2}$.

```
let mu = sample uniform (150, 200) in
let sigma = sample uniform (1, 50) in
let x1 = sample gaussian (mu, sigma) in
let x2 = sample gaussian (mu, sigma) in
score (f₁ x1); score (f₂ x2);
(mu, sigma)
```
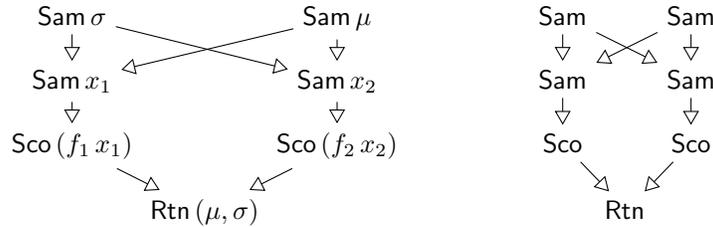
Fig. 1. A simple probabilistic program

A proposal step following the single-site kernel may choose to resample μ; then it must run through the entire trace, checking for potential dependencies to μ, though in this case none of the other variables need to be resampled.

So we argue that viewing a program as tree of traces is not most appropriate in this context: we propose instead to compile a program into a partially ordered structure reflecting the probabilistic dependencies.

With our approach, the example above would yield the partial order displayed below on the right-hand side. The nodes on the first line corresponds to the sample for μ and σ, and those on the second line to x_1 and x_2. This provides an accurate account of the probabilistic dependencies: whenever $e \leq e'$ (where $\leq$ is the reflexive, transitive closure of $\twoheadrightarrow$), it is the case that e' depends on e.

According to this representation of the program, a trace is no longer a linear order, but instead another partial order, similar to the previous one only annotated with a specific value for each variable. This is displayed below, on the left-hand side; note that the order $\leq$ is drawn top to bottom. There is an obvious erasure map from the trace (left) to the graph (right); this will be important later on.

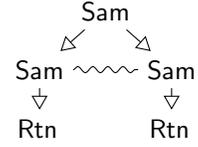*Conflict and control flow.* We have seen that a partial order can be used to faithfully represent the data dependency in the program; it is however not sufficient to accurately describe the control flow. In particular, computational events may live in different *branches* of a conditional statement, as in the following example:

```
let x = sample uniform (0, 5) in
if x ≥ 2 then sample gaussian (3, 1)
else sample uniform (2, 4)
```

The last two samples are independent, but also *incompatible*: in any given trace only one of them will occur. An example of a trace for this program is $\mathsf{Sam}\,1 \cdot \mathsf{Sam}\,3 \cdot \mathsf{Rtn}\,3$.

We represent this information by enriching the partial order with a conflict relation, indicating when two actions are in different branches of a conditional statement. The resulting structure is depicted on the right. Combining partial order and conflict in this way can be conveniently formalised using **event structures** [22]:

Definition 1. *An **event structure** is a tuple $(E, \leq, \#)$ where $(E, \leq)$ is a partially ordered set and $\#$ is an irreflexive, binary relation on E such that*

- *for every $e \in E$, the set $[e] = \{e' \in E \mid e' \leq e\}$ is finite, and*
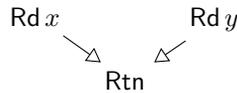- *if $e\#e'$ and $e' \leq e''$, then $e\#e''$.*

From the partial order $\leq$, we extract **immediate causality** $\rightarrowtail$: $e \rightarrowtail e'$ when $e < e'$ with no events in between; and from the conflict relation, we extract **minimal conflict** $\rightsquigarrow$: $e \rightsquigarrow e'$ when $e\#e'$ and there are no other conflicts in $[e] \cup [e']$. In pictures we draw $\rightarrowtail$ and $\rightsquigarrow$ rather than $\leq$ and $\#$.

A subset $x \subseteq E$ is a **configuration** of E if it is down-closed (if $e' \leq e \in x$ then $e' \in x$) and conflict-free (if $e, e' \in x$ then $\neg(e\#e')$). So in this framework, configurations correspond to exactly to *partial* executions traces of E.

The configuration $[e]$ is the **causal history** of e; we also write $[e)$ for $[e] \setminus \{e\}$. We write $\mathscr{C}(E)$ for the set of all finite configurations of E, a partial order under inclusion. A configuration x is **maximal** if it is maximal in $\mathscr{C}(E)$: for every $x' \in \mathscr{C}(E)$, if $x \subseteq x'$ then $x = x'$. We use the notation $x \stackrel{e}{-\!\!\subset} x'$ when $x' = x \cup \{e\}$, and in that case we say x' **covers** x.

An event structure is **confusion-free** if minimal conflict is transitive, and if any two events e, e' in minimal conflict satisfy $[e) = [e')$.

Compositionality. In order to give semantics to the language in a compositional manner, we must consider arbitrary *open* programs, *i.e.* with free parameters. Therefore we also represent each call to a parameter a as a *read* event, marked $\mathsf{Rd}\,a$. For instance the program $x + y$ with two real parameters will become the event structure

Note that the read actions on x and y are independent in the program (no order is specified), and the event structure respects this independence.

Our dependency graphs are event structures where each event carries information about the syntactic operation it comes from, a **label**, which depends on the typing context of the program:

$$\mathscr{L}^{\mathtt{static}}_{\Gamma \vdash B} ::= \mathsf{Rd}\,a \mid \mathsf{Rtn} \mid \mathsf{Sam} \mid \mathsf{Sco}\,,$$

where a ranges over variables $a : A$ in Γ.

Definition 2. *A **dependency graph** over $\Gamma \vdash B$ is an event structure G along with a labelling map $\mathsf{lbl} : G \to \mathscr{L}_{\Gamma \vdash B}^{static}$ where any two events $s, s' \in G$ labelled Rtn are in conflict, and all maximal configurations of G are of the form $[r]$ for $r \in G$ a return event.*

The condition on return events ensures that in any configuration of G there is at most one return event. Events of G are called static events.

We use dependency graphs as a causal representation of programs, reflecting the dependency between different parts of the program. In what follows we enrich this representation with runtime information in order to keep track of the dataflow of the program (in §3.3), and the associated distributions (in §3.4).

3.3   Runtime values and dataflow graphs

We have seen how data dependency can be captured by representing a program P as a dependency graph G_P. But observe that this graph does not give any runtime information about the data in P; every event $s \in G_P$ only carries a label $\mathsf{lbl}(s)$ indicating the class of action it belongs to. (For an event labelled $\mathsf{Rd}\, a$, G does not specify the value at a; whereas at runtime this will be filled by an element of $\llbracket A \rrbracket$ where A is the type of a.)

To each label, we can associate a measurable space of possible runtime values:

$$\mathscr{Q}(\mathsf{Rd}\, b) = \llbracket \Gamma(b) \rrbracket \qquad \mathscr{Q}(\mathsf{Rtn}) = \llbracket A \rrbracket \qquad \mathscr{Q}(\mathsf{Sam}) = (\mathbb{R}, \Sigma_{\mathbb{R}}) \qquad \mathscr{Q}(\mathsf{Sco}) = (\mathbb{R}, \Sigma_{\mathbb{R}}).$$

Then, in a particular execution, an event $s \in G_P$ has a value in $\mathscr{Q}(\mathsf{lbl}(s))$, and can be instead labelled by the following expanded set:

$$\mathscr{L}_{\Gamma \vdash B}^{\mathbf{run}} ::= \mathsf{Rd}\, a\, v \mid \mathsf{Rtn}\, v \mid \mathsf{Sam}\, r \mid \mathsf{Sco}\, r$$

where r ranges over real numbers; in $\mathsf{Rd}\, a\, v$, $a : A \in \Gamma$ and $v \in \llbracket A \rrbracket$; and in $\mathsf{Rtn}\, v$, v ranges over elements of $\llbracket B \rrbracket$. Notice that there is an obvious forgetful map $\alpha : \mathscr{L}_{\Gamma \vdash A}^{\mathbf{run}} \to \mathscr{L}_{\Gamma \vdash A}^{\mathtt{static}}$, discarding the runtime value. This runtime value can be extracted from a label in $\mathscr{L}_{\Gamma \vdash B}^{\mathbf{run}}$ as follows:

$$\mathbf{q}(\mathsf{Rd}\, b\, v) = v \qquad \mathbf{q}(\mathsf{Rtn}\, v) = v \qquad \mathbf{q}(\mathsf{Sam}\, r) = r \qquad \mathbf{q}(\mathsf{Sco}\, r) = r.$$

In particular, we have $\mathbf{q}(\ell) \in \mathscr{Q}(\alpha(\ell))$.

Such runtime events organise themselves in an event structure E_P, labelled over $\mathscr{L}_{\Gamma \vdash B}^{\mathbf{run}}$, the **runtime graph** of P. Runtime graphs are in general uncountable, and so difficult to represent pictorially. It can be done in some

$$\mathsf{Rd}\, a\, \mathsf{tt} \sim \mathsf{Rd}\, a\, \mathsf{ff}$$
$$\mathord{\Downarrow} \qquad\qquad \mathord{\Downarrow}$$
$$\mathsf{Rtn}\, 2 \qquad\quad \mathsf{Rtn}\, 3$$

simple, finite cases: the graph for `if` a `then 2 else 3` is depicted on the right. Recall that in dependency graphs conflict was used to represent conditional branches; here instead conflict is used to keep disjoint the possible outcomes of the same static event. (Necessarily, this static event must be a sample or a read, since other actions (return, score) are deterministic.)

Intuitively one can project runtime events to static events by erasing the runtime information; this suggests the existence of a function $\pi_P : E_P \to G_P$. This function will turn out to satisfy the axioms of a *rigid map of event structures*:

Definition 3. *Given event structures $(E, \leq_E, \#_E)$ and $(G, \leq_G, \#_G)$ a function $\pi : E \to G$ is a **rigid map** if*

- *it preserves configurations: for every $x \in \mathscr{C}(E)$, $\pi x \in \mathscr{C}(G)$*
- *it is* locally injective*: for every $x \in \mathscr{C}(E)$ and $e, e' \in x$, if $\pi(e) = \pi(e')$ then $e = e'$.*
- *it preserves dependency: if $e \leq_E e'$ then $\pi(e) \leq_G \pi(e')$.*

In general π is not injective, since many runtime events may correspond to the same static event – in that case however the axioms will require them to be in conflict. The last condition in the definition ensures that all causal dependencies come from G.

Given $x \in \mathscr{C}(G_P)$ we define the possible runtime values for x as the set $\mathscr{Q}(x)$ of functions mapping $s \in x$ to a runtime value in $\mathscr{Q}(\mathsf{lbl}(s))$; in other words $\mathscr{Q}(x) = \prod_{s \in x} \mathscr{Q}(\mathsf{lbl}(s))$. A configuration x' of E_P can be viewed as a trace over $\pi_P x'$; hence $\pi_P^{-1}\{x\} := \{x' \in \mathscr{C}(E_P) \mid \pi_P x' = x\}$ is the set of traces of P over x. We can now define dataflow graphs:

Definition 4. *A **dataflow graph** on $\Gamma \vdash B$ is a triple $\mathbf{S} = (E_S, G_S, \pi_S : E_S \to G_S)$ with G_S a dependency graph and E_S a runtime graph, such that:*

- *π_S is a rigid map and $\mathsf{lbl} \circ \pi_S = \alpha \circ \mathsf{lbl} : E_S \to \mathscr{L}_{\Gamma \vdash B}^{static}$*
- *for each $x \in \mathscr{C}(G_S)$, the following function is injective*

$$q_x : \pi_S^{-1}\{x\} \;\to\; \mathscr{Q}(x)$$
$$x' \;\mapsto\; (s \mapsto \mathbf{q}(\mathsf{lbl}(s)))$$

- *if $e, e' \in E_S$ with $e \smallsmile e'$ then $\pi e = \pi e'$, and moreover e and e' are either both sample or both read events.*

As mentioned above, maximal configurations of E_P correspond to total traces of P, and will be the states of the Markov chain in § 5. By the second axiom, they can be seen as pairs $(x \in \mathscr{C}(G_S), q \in \mathscr{Q}(x))$. Because of the third axiom, E_S is always confusion-free.

Measurable fibres. Rigid maps are convenient in this context because, they allow for reasoning about program traces by organising them as *fibres*. The key property we rely on is the following:

Lemma 2. *If $\pi : E \to G$ is a rigid map of event structures, then the induced map $\pi : \mathscr{C}(E) \to \mathscr{C}(G)$ is a* discrete fibration*: that is, for every $y \in \mathscr{C}(E)$, if $x \subseteq \pi y$ for some $x \in \mathscr{C}(G)$, then there is a unique $y' \in \mathscr{C}(E)$ such that $y' \subseteq y$ and $\pi y' = x$.*

This enables an essential feature of our approach: given a configuration x of the dataflow graph G, the fibre $\pi^{-1}\{x\}$ over it contains all the (possibly partial) program traces over x, *i.e.* those whose path through the program corresponds to that of x. Additionally the lemma implies that every pair of configurations

$x, x' \in \mathscr{C}(G)$ such that $x \subseteq x'$ induces a **restriction map** $r_{x,x'} : \pi^{-1}\{x'\} \to \pi^{-1}\{x\}$, whose action on a program trace over x' is to return its *prefix* over x.

Although there is no measure-theoretic structure in the definition of dataflow graphs, we can recover it: for every $x \in \mathscr{C}(G_S)$, the fibre $\pi_S^{-1}\{x\}$ can be equipped with the σ-algebra induced from $\Sigma_{\mathscr{Q}(x)}$ via q_x; it is generated by sets $q_x^{-1}U$ for $U \in \Sigma_{\mathscr{Q}(x)}$.

It is easy to check that this makes the restriction map $r_{x,x'} : \pi_S^{-1}\{x'\} \to \pi_S^{-1}\{x\}$ measurable for each pair x, x' of configurations with $x \subseteq x'$. (Note that this makes **S** a *measurable event structure* in the sense of [15].) Moreover, the map $q_{x,s} : \pi_S^{-1}\{x\} \to \mathscr{Q}(\mathsf{lbl}(s))$ for $s \in x \in \mathscr{C}(G_S)$, mapping $x' \in \pi_S^{-1}\{x\}$ to $\mathbf{q}(\mathsf{lbl}(s'))$ for s' the unique antecedent by π_S of s in x', is also measurable.

We will also make use of the following result:

Lemma 3. *Consider a dataflow* **S** *and* $x, y, z \in \mathscr{C}(G_S)$ *with* $x \subseteq y$, $x \subseteq z$, *and* $y \cup z \in \mathscr{C}(G_S)$. *If* $y \cap z = x$, *then the space* $\pi_S^{-1}\{y \cup z\}$ *is isomorphic to the set*

$$\{(u_y, u_z) \in \pi_S^{-1}\{y\} \times \pi_S^{-1}\{z\} \mid r_{x,y}(u_y) = r_{x,z}(u_z)\},$$

with σ-*algebra generated by sets of the form* $\{(u_y, u_z) \in X_y \times X_z \mid X_y \in \Sigma_{\pi_S^{-1}\{y\}}, X_z \in \Sigma_{\pi_S^{-1}\{z\}}$ *and* $r_{x,y}(u_y) = r_{x,z}(u_z)\}$.

(For the reader with knowledge of category theory, this says exactly that the diagram

$$
\begin{array}{ccc}
\pi_S^{-1}\{y \cup z\} & \xrightarrow{\ r_{y,y\cup z}\ } & \pi_S^{-1}\{y\} \\
{\scriptstyle r_{z,y\cup z}}\downarrow & & \downarrow{\scriptstyle r_{x,y}} \\
\pi_S^{-1}\{z\} & \xrightarrow[\ r_{x,z}\ ]{} & \pi_S^{-1}\{x\}
\end{array}
$$

is a pullback in the category of measurable spaces.)

3.4 Quantitative dataflow graphs

We can finally introduce the last bit of information we need about programs in order to perform inference: the probabilistic information. So far, in a dataflow graph, we know when the program is sampling, but not from which distribution. This is resolved by adding for each sample event s in the dependency graph a kernel $k_s : \pi^{-1}\{[s)\} \rightsquigarrow \pi^{-1}\{[s]\}$. Given a trace x over $[s)$, k_s specifies a probability distribution according to which x will be extended to a trace over $[s]$. This distribution must of course have support contained in the set $r_{[s),[s]}^{-1}\{x\}$ of traces over $[s]$ of which x is a prefix; this is the meaning of the technical condition in the definition below.

Definition 5. *A **quantitative dataflow graph** is a tuple* $\mathbf{S} = (E_S, G_S, \pi : E_S \to G_S, (k_s^S))$ *where for each sample event* $s \in G_S$, k_s^S *is a kernel* $\pi^{-1}\{[s)\} \rightsquigarrow \pi^{-1}\{[s]\}$ *satisfying for all* $x \in \pi^{-1}\{[s)\}$,

$$k_s^S(x, \pi^{-1}\{[s]\} \setminus r_{[s),[s]}^{-1}\{x\}) = 0.$$

This axiom stipulates that any extension $x' \in \pi_S^{-1}\{[s]\}$ of $x \in \pi_S^{-1}\{[s]\}$ drawn by k_s must contain x; in effect k_s only samples the runtime value for s.

From graphs to kernels. We show how to collapse a quantitative dataflow graph $\mathbf{S}$ on $\Gamma \vdash B$ to a kernel $[\![\Gamma]\!] \rightsquigarrow [\![B]\!]$. First, we extend the kernel family on sampling events $(k_s^S : \pi^{-1}\{[s]\} \rightsquigarrow \pi^{-1}\{[s]\})$ to a family $(k_s^{S[\gamma]} : \pi^{-1}\{[s]\} \rightsquigarrow \pi^{-1}\{[s]\})$ defined on *all* events $s \in S$, parametrised by the value of the environment $\gamma \in [\![\Gamma]\!]$. To define $k_s^{S[\gamma]}(x, \cdot)$ it is enough to specify its value on the generating set for $\Sigma_{\pi^{-1}\{[s]\}}$. As we have seen this contains elements of the form $q_{[s]}^{-1}(U)$ with $U \in \Sigma_{\mathscr{Q}([s])}$. We distinguish the following cases corresponding to the nature of s:

- If s is a sample event, $k_s^{S[\gamma]} = k_s^S$
- If s is a read on $a : A$, any $x \in \pi^{-1}[s]$ has runtime information $q_{[s]}(x)$ in $\mathscr{Q}([s]))$ which can be extended to $\mathscr{Q}([s])$ by mapping s to $\gamma(a)$:

$$k_s^{S[\gamma]}(x, q_{[s]}^{-1}U) = \delta_{q_{[s]}(x)[s:=\gamma(a)]}(U)$$

- If s is a return or a score event: any $x \in \pi^{-1}\{[s]\}$ has at most one extension to $o(x) \in \pi^{-1}\{[s]\}$ (because return and score events cannot be involved in a minimal conflict): $k_s^{S[\gamma]}(x, q_{[s]}^{-1}(U)) = \delta_{q_{[s]}(o(x))}(U)$. If $o(x)$ does not exist, we let $k_s^{S[\gamma]}(x, X) = 0$.

We can now define a kernel $k_{x,s}^{S[\gamma]} : \pi^{-1}\{x\} \rightsquigarrow \pi^{-1}\{x'\}$ for every atomic extension $x \stackrel{s}{-\!\!\subset} x'$ in G_S, ie. when $x' \setminus x = \{s\}$, as follows:

$$k_{x,s}^{S[\gamma]}(y, U) = k_s(r_{[s],x}(y), \{w \in \pi_S^{-1}\{[s]\} \mid (y, w) \in U\}).$$

The second argument to k_s above is always measurable, by a standard measure-theoretic argument based on Lemma 3, as $x \cap [s] = [s]$.

From this definition we derive:

Lemma 4. *If $x \stackrel{s_1}{-\!\!\subset} x_1$ and $x \stackrel{s_2}{-\!\!\subset} x_2$ are concurrent extensions of x (i.e. s_1 and s_2 are not in conflict), then $k_{x_1,s_2}^{S[\gamma]} \circ k_{x,s_1}^{S[\gamma]} = k_{x_2,s_1}^{S[\gamma]} \circ k_{x,s_2}^{S[\gamma]}$.*

Given a configuration $x \in \mathscr{C}(G_S)$ and a covering chain $\emptyset \stackrel{s_1}{-\!\!\subset} x_1 \ldots \stackrel{s_n}{-\!\!\subset} x_n = x$, we can finally define a measure on $\pi^{-1}\{x\}$:

$$\mu_x^{S[\gamma]} = k_{x_{n-1},s_n}^{S[\gamma]} \circ \ldots \circ k_{\emptyset,s_1}^{S[\gamma]}(*, \cdot),$$

where $*$ is the only trace over $\emptyset$. The particular covering chain used does not matter by the previous lemma. Using this, we can define the kernel of a quantitative dataflow graph $\mathbf{S}$ as follows:

$$\mathsf{kernel}(\mathbf{S})(\gamma, X) = \sum_{r \in G_S, \mathsf{lbl}(r) = \mathsf{Rtn}} \mu_{[r]}^{S[\gamma]}(q_{[r],r}^{-1}(X)),$$

where the measurable map $q_{[r],r} : \pi^{-1}\{r\} \to [\![B]\!]$ looks up the runtime value of r in an element of the fibre over $[r]$ (defined in § 3.3).

Lemma 5. *$\mathsf{kernel}(\mathbf{S})$ is an s-finite kernel $[\![\Gamma]\!] \rightsquigarrow [\![B]\!]$.*

4   Programs as labelled event structures

We now detail our interpretation of programs as quantitative dataflow graphs. Our interpretation is given by induction, similarly to the measure-theoretic interpretation given in § 2.3, in which composition of kernels plays a central role. In § 4.1, we discuss how to compose quantitative dataflow graphs, and in § 4.2, we define our interpretation.

4.1   Composition of probablistic event structures

Consider two quantitative dataflow graphs, S on $\Gamma \vdash A$, and T on $\Gamma, a : A \vdash B$ where a does not occur in Γ. In what follows we show how they can be composed to form a quantitative dataflow graph $T \odot^a S$ on $\Gamma \vdash B$.

Unlike in the kernel model of § 2.3, we will need two notions of composition. The first one is akin to the usual sequential composition: actions in T must wait on S to return before they can proceed. The second is closer to parallel composition: actions on T which do not depend on a read of the variable a can be executed in parallel with S. The latter composition is used to interpret the let construct. In `let` a `=` M `in` N, we want all the probabilistic actions or reads on other variables which do not depend on the value of a to be in parallel with M. However, in a program such as `case` M `of` $\{(i, x) \Rightarrow N_i\}_{i \in I}$ we do not want any actions of N_i to start before the selected branch is known, *i.e.* before the return value of M is known.

By way of illustration, consider the following simple example, in which we only consider runtime graphs, ignoring the rest of the structure for now. Suppose S and T are given by

$$S = \begin{array}{cc} \mathsf{Rd}\,b\,\mathsf{tt} \sim \mathsf{Rd}\,b\,\mathsf{ff} \\ \Downarrow \qquad \Downarrow \\ \mathsf{Rtn}\,\mathsf{ff} \quad \mathsf{Rtn}\,\mathsf{tt} \end{array} \qquad T = \begin{array}{c} \mathsf{Sam}\,r \quad \mathsf{Rd}\,a\,\mathsf{tt} \sim\!\!\sim \mathsf{Rd}\,a\,\mathsf{ff} \\ \mathsf{Rtn}\,((),\mathsf{tt}) \quad \mathsf{Rtn}\,((),\mathsf{ff}) \end{array}$$

The graph S can be seen to correspond to the program `if` b `then` `ff` `else` `tt` and T to the pairing $(\mathtt{sample}\ d\ (0), a)$ for any d. Here S is a runtime graph on $b : \mathbb{B} \vdash \mathbb{B}$ and T on $a : \mathbb{B}, b : \mathbb{B} \vdash \mathbb{B}$.

Both notions of compositions are displayed in the diagram below. The sequential composition (left) corresponds to

$$\mathtt{if}\ b\ \mathtt{then}\ (\mathtt{sample}\ d\ (0), \mathsf{ff})\ \mathtt{else}\ (\mathtt{sample}\ d\ (0), \mathsf{tt})$$

and the parallel composition to $(\mathtt{sample}\ d\ (0), \mathtt{if}\ b\ \mathtt{then}\ \mathsf{ff}\ \mathtt{else}\ \mathsf{tt})$:

$$T \odot^a_{\mathrm{seq}} S = \left(\begin{array}{cc} \mathsf{Rd}\,b\,\mathsf{tt} \sim \mathsf{Rd}\,b\,\mathsf{ff} \\ \Downarrow \qquad \Downarrow \\ \mathsf{Sam}\,r \quad \mathsf{Sam}\,r \\ \Downarrow \qquad \Downarrow \\ \mathsf{Rtn}\,\mathsf{ff} \quad \mathsf{Rtn}\,\mathsf{tt} \end{array} \right) \qquad T \odot^a_{\mathrm{par}} S = \left(\begin{array}{c} \mathsf{Sam}\,r \quad \mathsf{Rd}\,b\,\mathsf{tt} \sim \mathsf{Rd}\,b\,\mathsf{ff} \\ \mathsf{Rtn}\,\mathsf{ff} \quad \mathsf{Rtn}\,\mathsf{tt} \end{array} \right)$$

Composition of runtime and dependency graphs. Let us now define both composition operators at the level of the event structures. Through the bijection $\mathscr{L}^{\texttt{static}}_{\Gamma \vdash B} \simeq \mathscr{L}^{\texttt{run}}_{\Gamma' \vdash 1}$ where $\Gamma'(a) = 1$ for all $a \in \mathrm{dom}(\Gamma)$, we will see dependency graphs and runtime graphs as the same kind of objects, event structures labelled over $\mathscr{L}^{\texttt{run}}_{\Gamma \vdash A}$.

The two compositions $S \odot^a_{\mathrm{par}} T$ and $S \odot^a_{\mathrm{seq}} T$ are two instances of the same construction, parametrised by a set of labels $D \subseteq \mathscr{L}^{\texttt{run}}_{\Gamma, a:A \vdash B}$. Informally, D specifies which events of T are to depend on the return value of S in the resulting composition graph. It is natural to assume in particular that D contains all reads on a, and all return events.

Sequential and parallel composition are instances of this construction where D is set to one of the following:

$$D^{\Gamma, a:A \vdash B}_{\mathrm{seq}} = \mathscr{L}^{\texttt{run}}_{\Gamma, a:A \vdash B} \qquad\qquad D^{\Gamma, a:A \vdash B}_{\mathrm{par}} = \{\mathsf{Rd}\, a\, v, \mathsf{Rtn}\, v \in \mathscr{L}^{\texttt{run}}_{\Gamma, a:A \vdash B}\}.$$

We proceed to describe the construction for an abstract D. Let T be an event structure labelled by $\mathscr{L}^{\texttt{run}}_{\Gamma, a:A \vdash B}$ and S labelled by $\mathscr{L}^{\texttt{run}}_{\Gamma \vdash A}$. A configuration $x \in \mathscr{C}(S)$ is a **justification** of $y \in \mathscr{C}(T)$ when

1. if $\mathsf{lbl}(y)$ intersects D, then x contains a return event
2. for all $t \in y$ with label $\mathsf{Rd}\, a\, v$, there exists an event $s \in x$ labelled $\mathsf{Rtn}\, v$.

In particular if $\mathsf{lbl}(y)$ does not intersect D, then any configuration of S is a justification of y. A **minimal justification** of y is a justification that admits no proper subset which is also a justification of y. We now define the event structure $S \cdot_D T$ as follows:

- *Events*: $S \cup \{(x, t) \mid x \in \mathscr{C}(S), t \in T, x \text{ minimal justification for } [t]\}$;
- *Causality*: $\leq_S \cup \{(x, t), (x', t') \mid x \subseteq x' \wedge t \leq t'\} \cup \{s, (x, t) \mid s \in x\}$;
- *Conflict*: the symmetric closure of

$$\begin{aligned} &\#_S \cup \{(x, t), (x', t') \mid x \cup x' \notin \mathscr{C}(T) \vee t \#_B t'\} \\ &\cup \{s, (x, t) \mid \{s\} \cup x \notin \mathscr{C}(S)\}. \end{aligned}$$

Lemma 6. *$S \cdot_D T$ is an event structure, and the following is an order-isomorphism:*

$$\langle \cdot, \cdot \rangle : \{(x, y) \in \mathscr{C}(S) \times \mathscr{C}(T) \mid x \text{ is a justification of } y\} \cong \mathscr{C}(S \cdot_D T).$$

This event structure is not quite what we want, since it still contains return events from S and reads on a from T. To remove them, we use the following general construction. Given a Σ-labelled event structure E and $V \subseteq E$ a set of visible events, its **projection** $E \downarrow V$ has events V and causality, conflict and labelling inherited from E. Thus the composition of S and T is:

$$S \odot^a_D T := S \cdot_D T \downarrow (\{s \in S \mid s \text{ not a return}\} \cup \{(x, t) \mid t \text{ not a read on } a\}).$$

As a result $S \odot^a_D T$ is labelled over $\mathscr{L}^{\texttt{run}}_{\Gamma \vdash B}$ as needed.

Dataflow information. We now explain how this construction lifts to dataflow graphs. Consider dataflow graphs $S = (E_S, G_S, \pi_S : E_S \to G_S)$ on $\Gamma \vdash A$ and $T = (E_T, G_T, \pi_T : E_T \to E_T)$ on $\Gamma, a : A \vdash B$. Given $D \subseteq \mathscr{L}^{\text{static}}_{\Gamma, a : A \vdash B}$ we define

$$E_{S \cdot_D T} = E_S \cdot_{\alpha^{-1} D} E_T \qquad G_{S \cdot_D T} = G_S \cdot_D G_T$$
$$E_{S \odot^a_D T} = E_S \odot^a_{\alpha^{-1} D} E_T \qquad G_{S \odot^a_D T} = G_S \odot^a_D G_T$$

Lemma 7. *The maps π_S and π_T extend to rigid maps*

$$\pi_{S \cdot_D T} : E_{S \cdot_{\alpha^{-1} D} T} \to G_{S \cdot_D T}$$
$$\pi_{S \odot^a_D T} : E_{S \odot^a_{\alpha^{-1} D} T} \to G_{S \odot^a_D T}$$

Moreover, if $\langle x, y \rangle \in \mathscr{C}(E_{S \cdot_D T})$, $\langle \pi_S\, x, \pi_T\, y \rangle$ is a well-defined configuration of $G_{S \cdot_D T}$. As a result, for $\langle x, y \rangle \in \mathscr{C}(G_{S \cdot_D T})$, we have a injection $\varphi_{x,y} : \pi^{-1}\{\langle x, y \rangle\} \to \pi^{-1}\{x\} \times \pi^{-1}\{y\}$ making the following diagram commute:

$$
\begin{array}{ccc}
\pi^{-1}\{\langle x, y \rangle\} & \xrightarrow{\ \varphi_{x,y}\ } & \pi^{-1}\{x\} \times \pi^{-1}\{y\} \\
q_{\langle x,y \rangle} \downarrow & & \downarrow q_x \times q_y \\
\mathscr{Q}(\langle x, y \rangle) & \xrightarrow[\ \cong\ ]{} & \mathscr{Q}(x) \times \mathscr{Q}(y)
\end{array}
$$

In particular, $\varphi_{x,y}$ is measurable and induces the σ-algebra on $\pi^{-1}\{\langle x, y \rangle\}$. We write φ_x for the map $\varphi_{x,\emptyset}$, an isomorphism.

Adding probability. At this point we have defined all the components of dataflow graphs $S \odot^a_D T$ and $S \cdot_D T$. We proceed to make them quantitative.

Observe first that each sampling event of $G_{S \cdot_D T}$ (or equivalently of $G_{S \odot^a_D T}$ – sampling events are never hidden) corresponds either to a sampling event of G_S, or to an event (x, t) where t is a sampling event of G_T. We consider both cases to define a family of kernels $(k_s^{S \cdot_D T})$ between the fibres of $S \cdot_D T$. This will in turn induce a family $(k_s^{S \odot^a_D T})$ on $S \odot^a_D T$.

- If s is a sample event of G_S, we use the isomorphisms $\varphi_{[s)}$ and $\varphi_{[s]}$ of Lemma 7 to define:
$$k_s^{S \odot^a_D T}(v, X) = k_s^S(\varphi^{-1}_{[s)}\, v, \varphi^{-1}_{[s]} X).$$

- If s corresponds to (x, t) for t a sample event of G_T, then for every $X_x \in \Sigma_{\pi_S^{-1}\{x\}}$ and $X_t \in \Sigma_{\pi_T^{-1}\{[t)\}}$ we define

$$k_{(x,t)}^{S \odot^a_D T}(\langle x', y' \rangle, \varphi^{-1}_{x,[t]}(X_x \times X_t)) = \delta_{x'}(X_x) \times k_t^T(y', X_t).$$

By Lemma 7, the sets $\varphi^{-1}_{x,[t]}(X_x \times X_t)$ form a basis for $\Sigma_{\pi^{-1}\{\langle x,[t) \rangle\}}$, so that this definition determines the entire kernel.

So we have defined a kernel $k_s^{S \cdot_D T}$ for each sample event s of $G_{S \cdot_D T}$. We move to the composition $(S \odot_D^a T)$. Recall that the *causal history* of a configuration $z \in \mathscr{C}(G_{S \odot_D^a T})$ is the set $[z]$, a configuration of $G_{S \cdot_D T}$. We see that hiding does not affect the fibre structure:

Lemma 8. *For any $z \in \mathscr{C}(G_{S \odot_D^a T})$, there is a measurable isomorphism $\psi_z :$ $\pi_{S \odot_D^a T}^{-1}\{z\} \cong \pi_{S \cdot_D T}^{-1}\{[z]\}$.*

Using this result and the fact that $G_{S \odot_D^a T} \subseteq G_{S \cdot_D T}$, we may define for each s:

$$k_s^{S \odot_D^a T}(v, X) = k_s^{S \cdot_D T}(\psi_{[s]}(v), \psi_{[s]}X).$$

We conclude:

Lemma 9. $S \odot_D^a T := (G_{S \odot_D^a T}, E_{S \odot_D^a T}, \pi_{S \odot_D^a T}, (k_s^{S \odot_D^a T}))$ *is a quantitative dataflow graph on $\Gamma \vdash B$.*

Multicomposition. By chaining this composition, we can compose on several variables at once. Given quantitative dataflow graphs S_i on $\Gamma \vdash A_i$ and T on $\Gamma, a_1 : A_1, \ldots, a_n : A_n \vdash A$ we define

$$(S_i) \odot_{\mathrm{par}}^{(a_i)} T := S_1 \odot_{\mathrm{par}}^{a_1} (\ldots \odot_{\mathrm{par}}^{a_n} T)$$
$$(S_i) \odot_{\mathrm{seq}}^{(a_i)} T := S_1 \odot_{\mathrm{seq}}^{a_1} (\ldots \odot_{\mathrm{seq}}^{a_n} T)$$

4.2   Interpretation of programs

We now describe how to interpret programs of our language using quantitative dataflow graphs. To do so we follow the same pattern as for the measure-theoretical interpretation given in § 2.3.

Interpretation of functions. Given a measurable function $f : [\![A]\!] \to [\![B]\!]$, we define the quantitative dataflow graph

$$S_f^a = \left(\sum_{v \in [\![A]\!]} \begin{array}{c} \mathsf{Rd}\, a\, v \\ \curlyvee \\ \mathsf{Rtn}\, (f\, v) \end{array} \quad \to \quad \begin{array}{c} \mathsf{Rd}\, a \\ \curlyvee \\ \mathsf{Rtn} \end{array} \right).$$

We then define $[\![f\, M]\!]_{\mathcal{G}}$ as $[\![M]\!]_{\mathcal{G}} \odot_{\mathrm{par}}^a S_f^a$ where a is chosen so as not to occur free in M.

Probablistic actions. In order to interpret scoring and sampling primitives, we need the following two quantitative dataflow graphs:

$$\mathtt{score} = \left(\sum_{r \in \mathbb{R}} \begin{array}{c} \mathsf{Rd}\, a\, r \\ \curlyvee \\ \mathsf{Sco}\, r \\ \curlyvee \\ \mathsf{Rtn}\, () \end{array} \quad \to \quad \begin{array}{c} \mathsf{Rd}\, a \\ \curlyvee \\ \mathsf{Sco} \\ \curlyvee \\ \mathsf{Rtn} \end{array} \right) \qquad \mathtt{sample}_d = \left(\sum_{\mathbf{r} \in \mathbb{R}^n} \begin{array}{c} \mathsf{Rd}\, a\, \mathbf{r} \\ \curlyvee \\ \mathsf{Sam}\, s \\ \curlyvee \\ \mathsf{Rtn}\, () \end{array} \quad \to \quad \begin{array}{c} \mathsf{Rd}\, a \\ \curlyvee \\ \mathsf{Sam} \\ \curlyvee \\ \mathsf{Rtn} \end{array} \;, k_{\mathsf{Sam}} \right)$$

and we define k_{Sam} by integrating the density function d; here we identify $\mathscr{Q}(\{\mathsf{Rd}\,a\,,\mathsf{Sam}\})$ and $\pi^{-1}\{\{\mathsf{Rd}\,a\,,\mathsf{Sam}\}\}$:

$$k_{\mathsf{Sam}}(\{\mathsf{Rd}\,a\,\mathbf{r}\}, U) = \int_{q \in U, q(\mathsf{Rd}\,a)=\mathbf{r}} d(\mathbf{r}, q(\mathsf{Sam}))\mathrm{d}\lambda.$$

We can now interpret scoring and sampling constructs:

$$[\![\mathtt{score}\ M]\!]_{\mathcal{G}} = [\![M]\!]_{\mathcal{G}} \odot_{\mathrm{par}}^{a} \mathtt{score} \qquad [\![\mathtt{sample}\ d\ (M)]\!]_{\mathcal{G}} = [\![M]\!]_{\mathcal{G}} \odot_{\mathrm{par}}^{a} \mathtt{sample}_{d}.$$

Interpretation of tuples and variables. Given a family $(a_i)_{i \in I}$, we define the dataflow graph $\mathtt{tuple}_{(a_i:A_i)}$ on $a_1 : A_1, \ldots, a_n : A_n \vdash A_1 \times \ldots \times A_n$ as follows. Its set of events is the disjoint union

$$\bigcup_{i \in I, v \in [\![A_i]\!]} \mathsf{Rd}\,a_i\,v + \bigcup_{\mathbf{v} \in [\![A_1 \times \ldots \times A_n]\!]} \mathsf{Rtn}\,\mathbf{v}$$

where the conflict is induced by $\mathsf{Rd}\,a_i\,v \smallfrown \mathsf{Rd}\,a_i\,v'$ for $v \neq v'$; and causality contains all the pairs $\mathsf{Rd}\,a_i\,v \rightarrowtail \mathsf{Rtn}\,(v_1, \ldots, v_n)$ where $v_i = v$. Then we form a quantitative dataflow graph $\mathtt{Tuple}_{(a_i:A_i)}$, whose dependency graph is $\mathtt{tuple}_{(a_i:1)}$ (up to the bijection $\mathscr{L}_{\Gamma \vdash A}^{\mathtt{run}} \simeq \mathscr{L}_{\Gamma' \vdash 1}^{\mathtt{static}}$ where $\Gamma'(a) = 1$ for $a \in \mathrm{dom}(\Gamma)$); and the runtime graph is $\mathtt{tuple}_{(a_i:A_i)}$, along with the obvious rigid map between them.

We then define the semantics of $(M_1, \ldots, M_n)$:

$$[\![(M_1, \ldots, M_n)]\!]_{\mathcal{G}} = ([\![M_i]\!]_{\mathcal{G}})_i \odot_{\mathrm{par}}^{(a_i)} \mathtt{Tuple}_{\mathtt{a_i:A_i}},$$

where the a_i are chosen free in all of the M_j. This construction is also useful to interpret variables:

$$[\![a]\!]_{\mathcal{G}} = \mathtt{Tuple}_{a:A} \qquad \text{where } \Gamma \vdash a : A.$$

Interpretation of pattern matching. Consider now a term of the form $\mathtt{case}\ M\ \mathtt{of}$ $\{(i, a) \Rightarrow N_i\}_{i \in i}$. By induction, we have that $[\![N_i]\!]_{\mathcal{G}}$ is a quantitative dataflow graph on $\Gamma, a : A_i \vdash B$. Let us write $[\![N_i]\!]_{\mathcal{G}}^{*}$ for the quantitative dataflow graph on $\Gamma, a : (\sum_{i \in I} A_i) \vdash B$ obtained by relabelling events of the form $\mathsf{Rd}\,a\,v$ to $\mathsf{Rd}\,a\,(i, v)$, and sequentially precomposing with $\mathtt{Tuple}_{a:\sum_{i \in I} A_i}$. This ensures that minimal events in $[\![N_i]\!]_{\mathcal{G}}^{*}$ are reads on a. We then build the quantitative dataflow graph $\sum_{i \in I}[\![N_i]\!]_{\mathcal{G}}^{*}$ on $\Gamma, a : \sum_{i \in I} A_i \vdash B$. This can be composed with $[\![M]\!]_{\mathcal{G}}$:

$$[\![\mathtt{case}\ M\ \mathtt{of}\ \{(i, a) \Rightarrow N_i\}_{i \in I}]\!]_{\mathcal{G}} = [\![M]\!]_{\mathcal{G}} \odot_{\mathrm{seq}}^{a} \left(\sum_{i \in I}[\![N_i]\!]_{\mathcal{G}}^{*}\right).$$

It is crucial here that one uses *sequential* composition: none of the branches must be evaluated until the outcome of M is known.

Adequacy of composition. We now prove that our interpretation is adequate with respect to the measure-theoretic semantics described in § 2.3. Given any subset $D \subseteq \mathscr{L}^{\mathtt{static}}_{\Gamma, a : A \vdash B}$ containing returns and reads on a, we show that the composition $S \odot^a_D T$ does implement the composition of kernels:

Theorem 1. *For S a quantitative dataflow graph on $\Gamma \vdash A$ and T on $\Gamma, a : A \vdash B$, we have*

$$\mathsf{kernel}(S \odot^a_D T) = \mathsf{kernel}(T) \circ \mathsf{kernel}(S) : [\![\Gamma]\!] \to [\![B]\!].$$

From this result, we can deduce that the semantics in terms of quantitative dataflow graphs is adequate with respect to the measure-theoretic semantics:

Theorem 2. *For every term $\Gamma \vdash M : A$, $\mathsf{kernel}([\![M]\!]_{\mathcal{G}}) = [\![M]\!]$.*

5 An inference algorithm

In this section, we exploit the intensional semantics defined above and define a Metropolis-Hastings inference algorithm. We start, in § 5.1, by giving a concrete presentation of those quantitative dataflow graphs arising as the interpretation of probabilistic programs; we argue this makes them well-suited for manipulation by an algorithm. Then, in § 5.2, we give a more formal introduction to the Metropolis-Hastings sampling methods than that given in § 3.1. Finally, in § 5.3, we build the proposal kernel on which our implementation relies, and conclude.

5.1 A concrete presentation of probabilistic dataflow graphs

Quantitative dataflow graphs as presented in the previous sections are not easy to handle inside of an algorithm: among other things, the runtime graph has an uncountable set of events. In this section we show that some dataflow graphs, in particular those needed for modelling programs, admit a finite representation.

Recovering fibres. Consider a dataflow graph $\mathbf{S} = (E_S, G_S, \pi_S)$ on $\Gamma \vdash B$. It follows from Lemma 3 that the fibre structure of $\mathbf{S}$ is completely determined by the spaces $\pi_S^{-1}\{[s]\}$, for $s \in G_S$, so we focus on trying to give a simplified representation for those spaces.

First, let us notice that if s is a return or score event, given $x \in \pi^{-1}\{x\}$, the value $q_x(s)$ is determined by $q|_{[s)}$. In other words the map $\pi^{-1}\{[s]\} \to \mathscr{Q}([s))$ is an injection. This is due to the fact that minimal conflict in E_S cannot involve return or score events. As a result, E_S induces a partial function $o^S_s : \mathscr{Q}([s)) \rightharpoonup \mathscr{Q}(\mathsf{lbl}(s))$, called the **outcome function**. It is defined as follows:

$$o^S_s(q) = \begin{cases} q_{[s]}(x')(s) & \text{if there exists } x' \in \pi^{-1}\{x'\},\, q_{[s]}(x')|_{[s)} = q, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Note that x' must be unique by the remark above since its projection to $\mathscr{Q}([s))$ is determined by q. The function o^S is partial, because it might be the case that the event s occurs conditionally on the runtime value on $[s)$.

In fact this structure is all we need in order to describe a dataflow graph:

Lemma 10. *Given G_S a dependency graph on $\Gamma \vdash B$, and partial functions $(o_s) : \mathscr{Q}([s)) \rightharpoonup \mathscr{Q}(\mathsf{lbl}(s))$ for score and return events of S. There exists a dataflow graph $(E_S, G_S, \pi_S : E_S \to G_S)$ whose outcome functions coincide with the o_s. Moreover, there is an order-isomorphism*

$$\mathscr{C}(E_S) \cong \{(x, q) \mid x \in \mathscr{C}(G_S), q \in \mathscr{Q}(x), \forall s \in x, o_s(q|_{[s)}) = q(s)\}.$$

Adding probabilities. To add probabilities, we simply equip each sample event s of G_S with a density function $d_s : \mathscr{Q}([s)) \times \mathbb{R} \rightharpoonup \mathbb{R}$.

Definition 6. *A **concrete quantitative dataflow graph** is a tuple $(G_S, (o_s : \mathscr{Q}([s)) \rightharpoonup \mathscr{Q}(\mathsf{lbl}(s))), (d_s : \mathscr{Q}([s)) \times \mathbb{R} \rightharpoonup \mathbb{R})_{s \in sample(G_S)})$ where $d_s(x, \cdot)$ is normalised.*

Lemma 11. *Any concrete quantitative dataflow graph $\mathcal{S}$ unfolds to a quantitative dataflow graph* `unfold` $\mathcal{S}$.

We see now that the quantitative dataflow graphs arising as the interpretation of a program must be the unfolding of a concrete quantitative dataflow graph:

Lemma 12. *For any concrete quantitative dataflow graphs $\mathcal{S}$ on $\Gamma \vdash A$ and $\mathcal{T}$ on $\Gamma, a : A \vdash B$,* `unfold` $\mathcal{S} \odot_D^a T$`unfold` $\mathcal{T}$ *is the unfolding of a concrete quantitative dataflow graph. It follows that for any program $\Gamma \vdash M : B$, $[\![M]\!]_{\mathcal{G}}$ is the unfolding of a concrete quantitative dataflow graph.*

5.2 Metropolis-Hastings

Recall that the Metropolis-Hastings algorithm is used to sample from a density function $d : \mathbb{A} \to \mathbb{R}$ which may not be normalised. Here $\mathbb{A}$ is a measurable *state space*, equipped with a measure λ. The algorithm works by building a Markov chain whose stationary distribution is D, the probability distribution obtained from d after normalisation:

$$\forall X \in \Sigma_\mathbb{A}, D(X) = \frac{\int_{x \in X} d(x)}{\int_{x \in \mathbb{A}} d(x)}.$$

Our presentation and reasoning in the rest of this section are inspired by the work of Borgström et al. [2].

Preliminaries on Markov chains. A Markov chain on a measurable state space $\mathbb{A}$ is a probability kernel $k : \mathbb{A} \rightsquigarrow \mathbb{A}$, viewed as a transition function: given a state $x \in \mathbb{A}$, the distribution $k(x, \cdot)$ is the distribution from which a next sample state will be drawn. Usually, each $k(x, \cdot)$ comes with a procedure for sampling: we will treat this as a probabilistic program $M(x)$ whose output is the next state. Given an initial state $x \in \mathbb{A}$ and a natural number $n \in \mathbb{N}$, we have a distribution

$k^n(x, \cdot)$ on $\mathbb{A}$ obtained by iterating k n times. We say that the Markov chain k has **limit** the distribution μ on $\mathbb{A}$ when

$$\lim_{n \to \infty} ||k^n(x, \cdot) - \mu|| = 0 \qquad \text{where } ||\mu_1 - \mu_2|| = \sup_{A \in \Sigma_{\mathbb{A}}} \mu_1(A) - \mu_2(A).$$

For the purposes of this paper, we call a Markov chain $k : \mathbb{A} \to \mathbb{A}$ **computable** when there exists a type A such that $[\![A]\!] = \mathbb{A}$ (up to iso) and an expression *without scores* $x : A \vdash K : A$ such that $[\![K]\!] = k$. (Recall that programs without conditioning denote probabilistic kernels, and are easily sampled from, since all standard distributions in the language are assumed to come with a built-in sampler.)

We will use terms of our language to describe computable Markov chains language, taking mild liberties with syntax. We assume in particular that programs may call each other as subroutines (this can be done via substitutions), and that manipulating finite structures is computable and thus representable in the language.

The Metropolis-Hastings algorithm. Recall that we wish to sample from a distribution with un-normalised density $d : \mathbb{A} \to \mathbb{R}$; d is assumed to be computable. The Markov chain defined by the Metropolis-Hastings algorithm has two parameters: a computable Markov chain $x : A \vdash P : A$, the *proposal kernel*, and a measurable, computable function $p : \mathbb{A}^2 \to \mathbb{R}$ representing the kernel $[\![P]\!]$, *i.e.*

$$[\![P]\!](x, X') = \int_{x' \in X'} p(x, x') \, d\lambda(x').$$

The Markov-chain $\mathtt{MH}(P, p, d)$ is defined as

$$\mathtt{MH}(P, p, d)(x) := \texttt{let } x' = P(x) \texttt{ in}$$
$$\texttt{let } \alpha = \min\left(1, \frac{d(x') \times p(x, x')}{d(x) \times p(x', x)}\right) \texttt{ in}$$
$$\texttt{let } u = \texttt{sample uniform } (0, 1) \texttt{ in}$$
$$\texttt{if } u < \alpha \texttt{ then } x' \texttt{ else } x$$

In words, the Markov chain works as follows: given a start state x, it generates a proposal for the next state x' using P. It then computes an *acceptance ratio* α, which is the probability with which the new sample will be *accepted*: the return state will then either be the original x or x', accordingly.

Assuming P and p satisfy a number of conditions, the algorithm is correct:

Theorem 3. *Assume that P and p satisfies the following properties:*

1. **Strong irreducibility:** *There exists $n \in \mathbb{N}$ such that for all $x \in \mathbb{A}$ and $X \in \Sigma_{\mathbb{A}}$ such that $D(X) \neq \emptyset$ and $d(x) > 0$, there exists $n \in \mathbb{N}$ such that $[\![P]\!]^n(x, X) > 0$.*
2. $[\![P]\!](x, X') = \int_{x' \in X'} p(x, x')$.
3. *If $d(x) > 0$ and $p(x, y) > 0$ then $d(y) > 0$.*

4. If $d(x) > 0$ and $d(y) > 0$, then $p(x, y) > 0$ iff $p(y, x) > 0$.

Then, the limit of $\mathtt{MH}(P, p, d)$ for any initial state $x \in \mathbb{A}$ with $d(x) > 0$ is equal to D, the distribution obtained after normalising d.

5.3   Our proposal kernel

Consider a closed program $\vdash M : A$ in which every measurable function is a computable one. Then, its interpretation as a concrete quantitative dataflow graph is computable, and we write $\mathcal{S}$ for the quantitative dataflow graph whose unfolding is $[\![M]\!]_{\mathcal{G}}$. Moreover, because M is closed, its measure-theoretic semantics gives a measure $[\![M]\!]$ on $[\![A]\!]$. Assume that $\mathrm{norm}([\![M]\!])$ is well-defined: it is a probability distribution on $[\![A]\!]$. We describe how a Metropolis-Hastings algorithm may be used to sample from it, by reducing this problem to that of sampling from configurations of E_S according to the following density:

$$d_S(x, q) := \left(\prod_{s \in \mathrm{sample}(x)} d_s(q(s)) \right) \left(\prod_{s \in \mathrm{score}(x)} q(s) \right).$$

Lemma 10 induces a natural measure on $\mathscr{C}(E_S)$. We have:

Lemma 13. *For all $X \in \Sigma_{\mathscr{C}(E_S)}$, $\mu^{\mathcal{S}}(X) = \displaystyle\int_{y \in X} d_S(y)\mathrm{d}y$.*

Note that $d_S(x, q)$ is easy to compute, but it is not normalised. Computing the normalising factor is in general intractable, but the Metropolis-Hastings algorithm does not require the density to be normalised.

Let us write $\mu^{\mathcal{S}}_{\mathrm{norm}}(X) = \frac{\mu^{\mathcal{S}}(X)}{\mu^{\mathcal{S}}(\mathscr{C}(E_S))}$ for the normalised distribution. By adequacy, we have for all $X \in \Sigma_{[\![A]\!]}$:

$$\mathrm{norm}[\![M]\!](X) = \mu^{\mathcal{S}}_{\mathrm{norm}}(\mathtt{result}^{-1}(X)).$$

where $\mathtt{result} : \max \mathscr{C}(E_S) \rightharpoonup [\![A]\!]$ maps a maximal configuration of E_S to its return value, if any. This says that sampling from $\mathrm{norm}[\![M]\!]$ amounts to sampling from $\mu^{\mathcal{S}}_{\mathrm{norm}}$ and only keeping the return value.

Accordingly, we focus on designing a Metropolis-Hastings algorithm for sampling values in $\mathscr{C}(E_S)$ following the (unnormalised) density d_S. We start by defining a proposal kernel for this algorithm.

To avoid overburdening the notation, we will no longer distinguish between a type and its denotation. Since G_S is finite, it can be represented by a type, and so can $\mathscr{C}(G_S)$. Moreover, $\mathscr{C}(E_S)$ is a subset of $\sum_{x \in \mathscr{C}(G_S)} \mathscr{Q}(x)$ which is also representable as the type of pairs $(x \in \mathscr{C}(G_S), q \in \mathscr{Q}(x))$. Operations on G_S and related objects are all computable and measurable so we can directly use them in the syntax. In particular, we will make use of the function $\mathtt{ext} : \mathscr{C}(E_S) \to G_S + 1$ which for each configuration $(x, q) \in \mathscr{C}(E_S)$ returns $(1, \mathfrak{s})$ if there exists $x \overset{s}{-\!\!\!-\!\!\subset}$ with $o_s(q|_{[s)})$ defined, and $(2, *)$ if (x, q) is maximal.

Informally, for $(x, q) \in \mathscr{C}(E_S)$, the algorithm is:

- Pick a sample event $s \in x$, randomly over the set of sample events of x.
- Construct $x_0 := x \setminus \{s' \in x \mid s' \geq s\} \cup \{s\} \in \mathscr{C}(G_S)$.
- Return a maximal extension (x', q') of $(x_0, q|_{x_0})$ by only resampling the sample events of x' which are not in x.

The last step follows the single-site MH principle: sample events in $x \cap x'$ have already been evaluated in x, and are not updated. However, events which are in $x' \setminus x$ belong to conditional branches not explored in x; they must be sampled.

We start by formalising the last step of the algorithm. We give a probabilistic program `complete` which has three parameters: the original configuration (x, q), the current modification (x_0, q_0) and returns a possible maximal extension:

$$\texttt{complete}(x, q, x_0, q_0) = \texttt{case ext}(x_0, q_0) \texttt{ of}$$
$$(2, ()) \Rightarrow (x_0, q_0)$$
$$(1, s) \Rightarrow$$

```
    if s is a return or a score event then
        complete(x, v, x_0 ∪ {s}, q_0[s := o_s(q_0)])
    else if s ∈ x
            complete(x, q, x_0 ∪ {s}, q_0[s := q(s)])
    else
            complete(x, q, x_0 ∪ {s}, q_0[s := sample d (q_0)])
```

The program starts by trying to extend (x_0, q_0) by calling `ext`. If (x_0, q_0) is already maximal, we directly return it. Otherwise, we get an event s. To extend the quantitative information, there are three cases:

- if s is not a sample event, ie. since S is closed it must be a return or a score event, we use the function o_s.
- if s is a sample event occurring in x, we use the value in q
- if s is a sample event not occurring in x, we sample a value for it.

This program is recursive, but because G_S is finite, there is a static bound on the number of recursive calls; thus this program can be unfolded to a program expressible in our language. We can now define the proposal kernel:

$$P_S(x, q) =$$

```
    let s = sample uniformly over sample events in x in
    let r = sample d_s (q|[s)) in
    let x_0 = x \ {s' ≥ s | s' ∈ x} in
    complete(x, q, x_0, q[s := r])
```

We now need to compute the density for P_S to be able to apply Metropolis-Hastings. Given $(x, q), (x', q') \in \mathscr{C}(E_S)$, we define:

$$p_S((x, q), (x', q')) = \sum_{s \in \texttt{sample}(x)} \left(\frac{q_s(v'|_{[s)})}{|\texttt{sample}(x)|} \times \prod_{s' \in \texttt{sample}(x' \setminus x)} q_{s'}(v|_{[s')}) \right).$$

Theorem 4. *The Markov chain P_S and density p satisfy the hypothesis of Theorem 3, as a result for any $(x, q) \in \mathscr{C}(E_S)$ the distribution $[\![\mathtt{MH}(d_S, P_S, p_S)^n]\!]((x, q), \cdot)$ tends to μ_{norm}^P as n goes to infinity.*

One can thus sample from norm($[\![M]\!]$) using the algorithm above, keeping only the return value of the obtained configuration.

Let us re-state the key advantage of our approach: having access to the data dependency information, `complete` requires fewer steps in general, because at each proposal step only a portion of the graph needs exploring.

6   Conclusion

Related work. There are numerous approaches to the semantics of programs with random choice. Among those concerned with statistical applications of probabilistic programming are Staton et al. [18, 17], Ehrhard et al. [6], and Dahlqvist et al. [5]. A game semantics model was announced in [14].

The work of Scibior et al. [16] was influential in suggesting a denotational approach for proving correctness of inference, in the framework of quasi-Borel spaces [8]. It is not clear however how one could reason about data dependencies in this framework, because of the absence of explicit causal information.

Hur et al. [10] gives a proof of correctness for Trace MCMC using new forms of operational semantics for probabilistic programs. This method is extended to higher-order programs with *soft constraints* in Borgström et al. [2]. However, these approaches do not consider incremental recomputation.

To the best of our knowledge, this is the first work adressing formal correctness of incremental recomputation in MCMC. However, methods exist which take advantage of data dependency information to improve the performance of each proposal step in "naive" Trace MCMC. We mention in particular the work on *slicing* by Hur et al. [9]; other approaches include [4], [24]. In the present work we claim no immediate improvement in performance over these techniques, but only a mathematical framework for reasoning about the structures involved.

It is worth remarking that our event structure representation is reminiscent of *graphical model* representation made explicit in some languages. Indeed, for a first-order language such as the one of this paper, Bayesian networks can directly be used as a semantics, see [20]. We claim that the alternative view offered by event structures will allow for an easier extension to higher-order programs, using ideas from game semantics.

Perspectives. This is the start of an investigation into intensional semantics for probabilistic programs. Note that the framework of event structures is very flexible and the semantics presented here is by no means the only possible one. Additionally, though the present work only treats the case of a first-order language, we believe that building on recent advances in probabilistic concurrent game semantics [3, 15] (from which the present work draws much inspiration), we can extend the techniques of this paper to arbitrary higher-order probabilistic programs with recursion.

References

1. Patrick Billingsley. *Probability and measure.* John Wiley & Sons, 2008.
2. Johannes Borgström, Ugo Dal Lago, Andrew D Gordon, and Marcin Szymczak. A lambda-calculus foundation for universal probabilistic programming. In *ACM SIGPLAN Notices*, volume 51, pages 33–46. ACM, 2016.
3. Simon Castellan, Pierre Clairambault, Hugo Paquet, and Glynn Winskel. The concurrent game semantics of probabilistic PCF. In *Logic in Computer Science (LICS), 2018 33rd Annual ACM/IEEE Symposium on, ACM/IEEE*, 2018.
4. Yutian Chen, Vikash Mansinghka, and Zoubin Ghahramani. Sublinear approximate inference for probabilistic programs. *stat*, 1050:6, 2014.
5. Fredrik Dahlqvist, Vincent Danos, Ilias Garnier, and Alexandra Silva. Borel kernels and their approximation, categorically. *arXiv preprint arXiv:1803.02651*, 2018.
6. Thomas Ehrhard, Michele Pagani, and Christine Tasson. Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. volume 2, pages 59:1–59:28, 2018.
7. Andrew D Gordon, Thomas A Henzinger, Aditya V Nori, and Sriram K Rajamani. Probabilistic programming. In *Proceedings of the on Future of Software Engineering*, pages 167–181. ACM, 2014.
8. Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. A convenient category for higher-order probability theory. In *LICS'17, Reykjavik.*, pages 1–12, 2017.
9. Chung-Kil Hur, Aditya V Nori, Sriram K Rajamani, and Selva Samuel. Slicing probabilistic programs. In *ACM SIGPLAN Notices*, volume 49, pages 133–144. ACM, 2014.
10. Chung-Kil Hur, Aditya V Nori, Sriram K Rajamani, and Selva Samuel. A provably correct sampler for probabilistic programs. In *LIPIcs-Leibniz International Proceedings in Informatics*, volume 45. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
11. Oleg Kiselyov. Probabilistic programming language and its incremental evaluation. In *Asian Symposium on Programming Languages and Systems*, pages 357–376. Springer, 2016.
12. Oleg Kiselyov. Problems of the lightweight implementation of probabilistic programming. In *Proceedings of Workshop on Probabilistic Programming Semantics*, 2016.
13. Dexter Kozen. Semantics of probabilistic programs. *J. Comput. Syst. Sci.*, 22(3):328–350, 1981.
14. Luke Ong and Matthijs Vákár. S-finite kernels and game semantics for probabilistic programming. In *POPL'18 Workshop on Probabilistic Programming Semantics (PPS)*, 2018.
15. Hugo Paquet and Glynn Winskel. Continuous probability distributions in concurrent games. *Electronic Notes in Theoretical Computer Science*, 341:321–344, 2018.

16. Adam Ścibior, Ohad Kammar, Matthijs Vákár, Sam Staton, Hongseok Yang, Yufei Cai, Klaus Ostermann, Sean K Moss, Chris Heunen, and Zoubin Ghahramani. Denotational validation of higher-order bayesian inference. *Proceedings of the ACM on Programming Languages*, 2(POPL):60, 2017.

17. Sam Staton. Commutative semantics for probabilistic programming. In *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, pages 855–879, 2017.

18. Sam Staton, Hongseok Yang, Frank D. Wood, Chris Heunen, and Ohad Kammar. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Proceedings of LICS '16, New York, NY, USA, July 5-8, 2016*, pages 525–534, 2016.

19. Luke Tierney. Markov chains for exploring posterior distributions. *Annals of Statistics*, 22:1701–1762, 1994.

20. Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, and Frank Wood. An introduction to probabilistic programming. *arXiv preprint arXiv:1809.10756*, 2018.

21. David Wingate, Andreas Stuhlmüller, and Noah Goodman. Lightweight implementations of probabilistic programming languages via transformational compilation. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 770–778, 2011.

22. Glynn Winskel. Event structures. In *Advances in Petri Nets*, pages 325–392, 1986.

23. Glynn Winskel. Distributed probabilistic and quantum strategies. *Electr. Notes Theor. Comput. Sci.*, 298:403–425, 2013.

24. Yi Wu, Lei Li, Stuart Russell, and Rastislav Bodik. Swift: Compiled inference for probabilistic programming languages. *arXiv preprint arXiv:1606.09242*, 2016.

25. Lingfeng Yang, Patrick Hanrahan, and Noah Goodman. Generating efficient MCMC kernels from probabilistic programs. In *Artificial Intelligence and Statistics*, pages 1068–1076, 2014.

A   Proofs of § 2

Lemma 1. *For $\Gamma \vdash M : A$ without scores, $[\![M]\!](\gamma, [\![A]\!]) = 1$ for each $\gamma \in [\![\Gamma]\!]$.*

Proof. By induction on M: probabilistic kernels are stable under composition.

B   Proofs of § 3.1

Lemma 3. *Consider a dataflow $\mathbf{S}$ and $x, y, z \in \mathscr{C}(G_S)$ with $x \subseteq y$, $x \subseteq z$, and $y \cup z \in \mathscr{C}(G_S)$. If $y \cap z = x$, then the space $\pi_S^{-1}\{y \cup z\}$ is isomorphic to the set*

$$\{(u_y, u_z) \in \pi_S^{-1}\{y\} \times \pi_S^{-1}\{z\} \mid r_{x,y}(u_y) = r_{x,z}(u_z)\},$$

with σ-algebra generated by sets of the form $\{(u_y, u_z) \in X_y \times X_z \mid X_y \in \Sigma_{\pi_S^{-1}\{y\}}, X_z \in \Sigma_{\pi_S^{-1}\{z\}} \text{ and } r_{x,y}(u_y) = r_{x,z}(u_z)\}$.
  (For the reader with knowledge of category theory, this says exactly that the diagram

$$\begin{array}{ccc}
\pi_S^{-1}\{y \cup z\} & \xrightarrow{\ r_{y,y\cup z}\ } & \pi_S^{-1}\{y\} \\
\downarrow{\scriptstyle r_{z,y\cup z}} & & \downarrow{\scriptstyle r_{x,y}} \\
\pi_S^{-1}\{z\} & \xrightarrow[\ r_{x,z}\ ]{} & \pi_S^{-1}\{x\}
\end{array}$$

is a pullback in the category of measurable spaces.)

Proof. We show that the diagram is a pullback. Consider $\varphi : X \to \pi^{-1}\{y\}$ and $\psi : X \to \pi^{-1}\{z\}$. We define $\langle \varphi, \psi \rangle : X \to \pi^{-1}\{y \cup z\}$, mapping α to $\varphi\,\alpha \cup \psi\,\alpha$. This is well-defined, first $\varphi\,\alpha \cup \psi\,\alpha$ is downclosed. Furthermore, is there is a minimal conflict $s \smallsmile t$ with $s \in \varphi\,\alpha$ and $t \in \psi\,\alpha$, then we know that $\pi\,s = \pi\,t$, hence since $y \cap z = x$, we have $s, t \in \pi^{-1}x$, and that means that $s \in r_{x,y}(\varphi\,\alpha) = r_{x,z}(\psi\,\alpha) \ni t$. By local injectivity of π we deduce that $s = t$, absurd.

Unicity of $\langle \varphi, \psi \rangle$ is easy to see. Finally, to show that $\langle \varphi, \psi \rangle$ is measurable, we just use that the simila diagram with $\mathscr{Q}(\cdot)$ instead of $\pi^{-1}\{\cdot\}$ is also a pullback.

Lemma 5. *kernel*(**S**) *is an s-finite kernel* $\llbracket \Gamma \rrbracket \rightsquigarrow \llbracket B \rrbracket$.

Proof. We first notice that $k_s^{S[\gamma]}$ is a s-finite kernel $\pi^{-1}\{[s)\} \rightsquigarrow \pi^{-1}\{[s]\}$. From there, $k_{x,s}$ is a kernel as the lifting of a kernel through a pullback square. This clearly preseves s-finiteness. Finally $\mu_x^{S[\gamma]}$ is s-finite as a composition of s-finite kernels. So we deduce that for every $\gamma \in \llbracket \Gamma \rrbracket$, kernel$((\,)\mathbf{S})(\gamma, \cdot)$ is a s-finite measure. Now, the map $\gamma \mapsto k_s^{S[\gamma]}$ is trivially seen to be measurable. This implies that kernel$((\,)\mathbf{S})$ is indeed a kernel, hence it is a s-finite kernel.

C   Proofs of § 4

C.1   Composition of probablistic event structures

Lemma 6. $S \cdot_D T$ *is an event structure, and the following is an order-isomorphism:*

$$\langle \cdot, \cdot \rangle : \{(x,y) \in \mathscr{C}(S) \times \mathscr{C}(T) \mid x \text{ is a justification of } y\} \cong \mathscr{C}(S \cdot_D T).$$

Proof. We check that $S \cdot_D T$ is an event structure. The relation $\leq$ is clearly reflexive and antisymmetric. We check transitivity and conflict inheritance.

Transivity The two interesting cases are:
  - When $s \leq (x,t) \leq (x',t')$: then $s \in x$ and $x \subseteq x'$ hence $s \in x'$ and $s \leq (x',t')$.
  - When $s \leq s' \leq (x,t)$: then $s \in x$ since x is downclosed and $s \leq (x',t)$.

Conflict inheritance There are two interesting cases:
  - When $s\#s' \leq (x,t)$: since $s' \in x$, we cannot have $x \cup \{s\} \in \mathscr{C}(S)$ hence $s\#(x,t)$

    – When $s\#(x,t) \leq (x',t')$: since $x \cup \{s\}$ is not a configuration, $x' \cup \{s\}$ cannot be a configuration either and $s\#(x',t')$.

We now check the isomorphism. If $z \in \mathscr{C}(S \cdot_D T)$, then write $x = \bigcup_{s \in x \cap S} s$. It is easy to check that $x \in \mathscr{C}(S)$. Similarly we write $y = \bigcup_{(x,t) \in x} t$, which is similarly a configuration of T. We check that x justifies y.

– If there exists $t \in y$ such that $\mathsf{lbl}(t) \in D$, then we have $(x_0, t) \in z$ hence there exists a return event s in x_0. Since $s \leq (x_0, t)$ by downclosure we have $s \in x$ as desired.
– Same reasoning as before.

Moreover the map $z \mapsto (x, y)$ preserves and reflect the order. Conversely, if we are given a pair (x, y), for each $t \in y$ there exists a unique minimal justification x_t (either x or $\emptyset$) of t in x and we can consider $z = \bigcup x \cup \bigcup_{t \in y} (x_t, t)$ for the inverse map.

Lemma 7. *The maps π_S and π_T extend to rigid maps*

$$\pi_{S \cdot_D T} : E_{S \cdot_{\alpha^{-1}D} T} \to G_{S \cdot_D T}$$

$$\pi_{S \odot_D^a T} : E_{S \odot_{\alpha^{-1}D}^a T} \to G_{S \odot_D^a T}$$

Moreover, if $\langle x, y \rangle \in \mathscr{C}(E_{S \cdot_D T})$, $\langle \pi_S\, x, \pi_T\, y \rangle$ is a well-defined configuration of $G_{S \cdot_D T}$. As a result, for $\langle x, y \rangle \in \mathscr{C}(G_{S \cdot_D T})$, we have a injection $\varphi_{x,y} : \pi^{-1}\{\langle x, y \rangle\} \to \pi^{-1}\{x\} \times \pi^{-1}\{y\}$ making the following diagram commute:

$$
\begin{array}{ccc}
\pi^{-1}\{\langle x, y \rangle\} & \xrightarrow{\ \varphi_{x,y}\ } & \pi^{-1}\{x\} \times \pi^{-1}\{y\} \\
{\scriptstyle q_{\langle x,y \rangle}} \downarrow & & \downarrow {\scriptstyle q_x \times q_y} \\
\mathscr{Q}(\langle x, y \rangle) & \xrightarrow[\ \cong\ ]{} & \mathscr{Q}(x) \times \mathscr{Q}(y)
\end{array}
$$

In particular, $\varphi_{x,y}$ is measurable and induces the σ-algebra on $\pi^{-1}\{\langle x, y \rangle\}$. We write φ_x for the map $\varphi_{x,\emptyset}$, an isomorphism.

Proof. We focus the proof on $S \odot_D^a T$, the reasoning is similar for $S \cdot_D T$.

    We first show that if $\langle x, y \rangle$ correponds to a configuration of $E_S \odot_D^a E_T$, then $\langle \pi_S\, x, \pi_T\, y \rangle$ corresponds to a configuration of $G_S \odot_D^a G_T$:

– Consider an event $\pi_T(t) \in \pi_T\, y$ whose label is D. By construction the label of t is in $\alpha^{-1} D$ hence there exists a return event in $s \in x$, and $\pi_S\, s$ is a return event of $\pi_S\, x$ as desired. If furthermore the label of $\pi_T\, t$ is $\mathsf{Rd}\, v$, then the label of t must be $\mathsf{Rtn}\, v'$ with $\alpha\, v' = v$ and we can conclude.
– If there exists a return event in $\pi_S\, x$, then there must be a return event in x as well and y must contain an event t labelled in $\alpha^{-1} D$. Hence $\pi_T\, t$ is labelled in D as desired.
– Same reasoning as before.

Let us define $\pi_{S \odot_D^a T}$ as follows:

$$\pi_{S \odot_D^a T} : E_{S \odot_D^a T} \to G_{S \odot_D^a T}$$
$$s \in E_S \mapsto \pi_S\, s$$
$$(x, t) \mapsto (\pi_S x, \pi_T\, t)$$

This is well-defined: indeed if $(x, [t])$ is a composition pair, then $\pi_T\,[t] = [\pi_T\, t]$ by rigidity and $(\pi_S\, x, [\pi_T t])$ is also a composition pair. Furthermore if $x' \subseteq \pi_S\, x$ is a smaller justification, then $\pi_S^{-1}\, x'$ (well-defined by rigidity again) would be a smaller justification of $[t]$. Rigidity is a routine check.

Lemma 8. *For any $z \in \mathscr{C}(G_{S \odot_D^a T})$, there is a measurable isomorphism ψ_z : $\pi_{S \odot_D^a T}^{-1}\{z\} \cong \pi_{S \cdot_D T}^{-1}\{[z]\}$.*

Lemma 9. *$S \odot_D^a T := (G_{S \odot_D^a T}, E_{S \odot_D^a T}, \pi_{S \odot_D^a T}, (k_s^{S \odot_D^a T}))$ is a quantitative dataflow graph on $\Gamma \vdash B$.*

Lemma 4. *If $x \overset{s_1}{-\!\!\subset} x_1$ and $x \overset{s_2}{-\!\!\subset} x_2$ are concurrent extensions of x (i.e. s_1 and s_2 are not in conflict), then $k_{x_1,s_2}^{S[\gamma]} \circ k_{x,s_1}^{S[\gamma]} = k_{x_2,s_1}^{S[\gamma]} \circ k_{x,s_2}^{S[\gamma]}$.*

Proof. We know that the measurable space $\pi^{-1}(x_1 \cup x_2)$ is the pullback along the two inclusions $x \subseteq x_1$ and $x \subseteq x_2$. This implies that given $v \in \pi^{-1}(x)$, the measurable space obtained by the preimage of $r_{x,x_1 \cup x_2}^{-1}(v)$ is generated by the rectangles $\langle X_1, X_2 \rangle$ where $X_i \in \pi^{-1}\{x_i\} \cap r_{x,x_i}^{-1}(v)$.

Hence given $v \in \pi^{-1}\{x\}$, we have that for any rectangle:

$$k_{x_1,s_2} \circ k_{x,s_1}(v, \langle X_1, X_2 \rangle) = \int_{v' \in \pi^{-1}\{x_1\}} k_{x_1,s_2}(v', X) k_{x,s_1}(v, dv')$$
$$= \int_{v' \in r_{x_1, x_1 \cup x_2}(X)} k_{x,s_2}(r_{x,x_1}\, v', X_2) k_{x,s_1}(v, dv')$$
$$= k_{x,s_2}(v, X_2) k(v, X_1)$$

Hence by symmetry we have $k_{x_1,s_2} \circ k_{x,s_1}(v, \langle X_1, X_2 \rangle) = k_{x_2,s_1} \circ k_{x,s_2}(v, \langle X_1, X_2 \rangle)$ which implies the desired result.

C.2   Adequacy of composition

We now prove that the kernel of a composition is the composition of a kernel. For that, we need a few intermediate result. In this section, we assume S a quantitative dataflow graph on $\Gamma \vdash A$ and T on $\Gamma, a : A \vdash B$ as well as a suitable subset D.

Lemma 14. *We have a bijection $\max(S \odot_D^a T) \simeq \max(S) \times \max(T)$.*

Proof. A maximal element of $S \odot_D^a T$ is of the form (x, t) where x is maximal in S, ie. of the form $[r]$ and t maximal in T, ie. a return event. The map $(x, t) \mapsto (r, t)$ is clearly injective. Given $(r_S, r_T) \in \max(S) \times \max(T)$, $([r_S], r_T)$ is the desired maximal event of $S \odot_D^a T$.

Lemma 15. *If $([r_S], r_T)$ is a maximal in $S \odot_D^a T$, then we have for all $\gamma \in [\![\Gamma]\!]$ and $X \in \Sigma_{[\![B]\!]}$:*

$$\mu_{\langle [r_S], r_T \rangle}^{(S \odot_D^a T)[\gamma]}(o_{([r_S], r_T)}^{-1}(X)) = \int_{a \in [\![A]\!]} \mu_{[r_T]}^{T[\gamma, a]}(o_{r_T}^{-1}(X)) \mu_{[r_S]}^{S[\gamma]}(o_{r_S}^{-1}(\mathrm{d}a)).$$

Proof. We use the fact that the distribution induced by a configuration is compositional: since we have the inclusions: $\emptyset \subseteq \langle [r_S], \emptyset \rangle \subseteq \langle [r_S], [r_T] \rangle$, we have:

$$\mu_{\langle [r_S], r_T \rangle}^{(S \odot_D^a T)[\gamma]}(o_{([r_S], r_T)}^{-1}(X)) = \left(k_{\langle [r_S], \emptyset \rangle, \langle [r_S], [r_T] \rangle}^{(S \odot_D^a T)[\gamma]} \circ k_{\emptyset, \langle [r_S], \emptyset \rangle}^{(S \odot_D^a T)[\gamma]} \right) (o_{([r_S], r_T)}^{-1}(X))$$

$$= \int_{v \in \pi^{-1} \langle [r_S], \emptyset \rangle} k_{\langle [r_S], \emptyset \rangle, \langle [r_S], [r_T] \rangle}^{(S \odot_D^a T)[\gamma]}(v, o_{([r_S], r_T)}^{-1}(X)) \times \mu_{\langle [r_S], \emptyset \rangle}^{(S \odot_D^a T)[\gamma]}(\mathrm{d}v)$$

We are almost there. First, we have $\pi^{-1}\langle [r_S], \emptyset \rangle \cong \langle \pi^{-1}[r_S], \emptyset \rangle \cong \pi^{-1}[r_S]$, we will leave this isomorphism silent for the sake of clarity. Then, remark that, up to this isomorphism we have

$$\mu_{\langle [r_S], \emptyset \rangle}^{(S \odot_D^a T)[\gamma]} = \mu_{[r_S]}^{S[\gamma]}.$$

and also, that

$$k_{\langle [r_S], \emptyset \rangle, \langle [r_S], [r_T] \rangle}^{(S \odot_D^a T)[\gamma]}(v, o_{([r_S], r_T)}^{-1}(X)) = k_{[r_T]}^{T[\gamma, o_r(v)]}(o_{r_T}^{-1}(X)).$$

By substitution, we are one change of variable away:

$$\mu_{\langle [r_S], r_T \rangle}^{(S \odot_D^a T)[\gamma]}(o_{([r_S], r_T)}^{-1}(X)) = \int_{v \in \pi^{-1}[r_S]} k_{[r_T]}^{T[\gamma, o_r(v)]}(o_{r_T}^{-1}(X)) \times \mu_{[r_S]}^{S[\gamma]}(\mathrm{d}v)$$

$$= \int_{a \in o_r^{-1}(\pi^{-1}[r_S])} k_{[r_T]}^{T[\gamma, a]}(o_{r_T}^{-1}(X)) \times \mu_{[r_S]}^{S[\gamma]}(o_r^{-1}(\mathrm{d}a))$$

$$= \int_{a \in [\![A]\!]} k_{[r_T]}^{T[\gamma, a]}(o_{r_T}^{-1}(X)) \times \mu_{[r_S]}^{S[\gamma]}(o_r^{-1}(\mathrm{d}a))$$

Theorem 1. *For S a quantitative dataflow graph on $\Gamma \vdash A$ and T on $\Gamma, a : A \vdash B$, we have*

$$\mathsf{kernel}(S \odot_D^a T) = \mathsf{kernel}(T) \circ \mathsf{kernel}(S) : [\![\Gamma]\!] \to [\![B]\!].$$

Proof. By Lemma 14 and Lemma 15, we have:

$$\mathsf{kernel}(S \odot_D^a T)(\gamma, X) = \sum_{r \in \max(S \odot_D^a T)} \mu_{[r]}^{(S \odot_D^a T)[\gamma]}(o_r^{-1}(X))$$

$$= \sum_{r_S \in \max S, r_T \in \max T} \mu_{\langle [r_S], [r_T] \rangle}^{(S \odot_D^a T)[\gamma]}(o_{([r_S], r_T)}^{-1}(X))$$

$$= \sum_{r_S \in \max S, r_T \in \max T} \int_{a \in [\![A]\!]} \mu_{[r_T]}^{T[\gamma, a]}(o_{r_T}^{-X}(X)) \times \mu_{[r_S]}^{S[\gamma]}(o_{r_S}^{-1}(\mathrm{d}a))$$

$$= \int_{a \in [\![A]\!]} \left(\sum_{r_T \in T} \mu_{[r_T]}^{T[\gamma, a]}(o_{r_T}^{-X}(X)) \right) \left(\sum_{r_S \in \max S} \mu_{[r_S]}^{S[\gamma]}(o_{r_S}^{-1}(\mathrm{d}a)) \right)$$

$$= \int_{a \in [\![A]\!]} \mathsf{kernel}(T)((\gamma, a), X) \times \mathsf{kernel}(S)(\gamma, \mathrm{d}a)$$

$$= (\mathsf{kernel}(T) \circ \mathsf{kernel}(S))(\gamma, X)$$

Theorem 2. *For every term $\Gamma \vdash M : A$, $\mathsf{kernel}([\![M]\!]_\mathcal{G}) = [\![M]\!]$.*

Proof. For most cases, the proof directly relies on Theorem 1, and simple calculation, such that $\mathsf{kernel}(S_f) = \bar{f}$ where $\bar{f}$ is the lift of $f : [\![A]\!] \to [\![B]\!]$ to a kernel $[\![A]\!] \rightsquigarrow [\![B]\!]$. The only interesting case is that of pairing.

Consider a term of the form `case` M `of` $\{(i, x) \Rightarrow N_i\}_{i \in I}$. First, it is easy to see that

$$\mathsf{kernel}([\![N_i]\!]^*)(\gamma, X) = \begin{cases} \mathsf{kernel}([\![N_i]\!])(\gamma[a := a'], X) & \gamma(a) = (i, a') \\ 0 & \text{otherwise} \end{cases}$$

and

$$\mathsf{kernel}(\sum [\![N_i]\!]^*(\gamma, X)) = \sum_{i \in I} \mathsf{kernel}([\![N_i]\!]^*)(\gamma, X).$$

From this equality and Theorem 1, the desired result follows.

D   Proofs of § 5

D.1   Concrete quantitative dataflow graphs

Lemma 10. *Given G_S a dependency graph on $\Gamma \vdash B$, and partial functions $(o_s) : \mathcal{Q}([s)) \rightharpoonup \mathcal{Q}(\mathsf{lbl}(s))$ for score and return events of S. There exists a dataflow graph $(E_S, G_S, \pi_S : E_S \to G_S)$ whose outcome functions coincide with the o_s. Moreover, there is an order-isomorphism*

$$\mathscr{C}(E_S) \cong \{(x, q) \mid x \in \mathscr{C}(G_S), q \in \mathcal{Q}(x), \forall s \in x, o_s(q|_{[s)}) = q(s)\}.$$

Proof. Say that $q \in \mathcal{Q}(x)$ for $x \in \mathcal{C}(G_S)$ is reachable when for all $s \in x$, $o_s(q|_{[s)}) = q(s)$. The events of E_S are pairs (s, q) of an event $s \in S$ and $q \in \mathcal{Q}([s])$ reachable. Causality is pairwise conflict and $(s, q)\#(s', q')$ when $s\#s'$ or $q(s_0) \neq q'(s_0)$ for $s_0 \in [s] \cap [s']$. The labelling is the obvious thing. It is straightforward to check that this is an event structure and that $(s, q) \mapsto s$ is a rigid map $E_S \to G_S$ such that (E_S, G_S, π_S) is a dataflow graph with outcome function o_s. The characterisation of the configurations follows directly from the configuration.

Lemma 11. *Any concrete quantitative dataflow graph $\mathcal{S}$ unfolds to a quantitative dataflow graph* **unfold** *$\mathcal{S}$.*

Proof. Straightforward, for every sample event s let

$$k_s(x', q_x^{-1}U) = \int_{q \in U} \delta_{q_x(x') \subseteq q} \times d_s(q_x(x'), q(s)).$$

Lemma 12. *For any concrete quantitative dataflow graphs $\mathcal{S}$ on $\Gamma \vdash A$ and $\mathcal{T}$ on $\Gamma, a : A \vdash B$,* **unfold** *$\mathcal{S} \odot_D^a T$* **unfold** *$\mathcal{T}$ is the unfolding of a concrete quantitative dataflow graph. It follows that for any program $\Gamma \vdash M : B$, $[\![M]\!]_{\mathcal{G}}$ is the unfolding of a concrete quantitative dataflow graph.*

Proof. Given (G_s, o^S, k^S) and (G_T, o^T, k^T), we define:

- if $s \in G_s$ is a score or return event:

$$o_s^{S \odot_D^a T}(q) = o_s^S(\varphi\, q) \qquad \text{where } \varphi_s : \pi_{S \odot_D^a T}^{-1}\langle[s), \emptyset\rangle \cong \pi_S^{-1}[s)$$

- If $s \in G_S$ is a sample event, we define

$$d_s^{S \odot_D^a T}(q, r) = d_s^S(\varphi\, q, r)$$

- If $(x, t) \in G_{S \odot_D^a T}$ is a score or return event, we let:

$$o_{(x,t)}^{S \odot_D^a T}(q) = o_t^T(\psi\, q) \qquad \text{where } \psi_t : \pi_{S \odot_D^a T}^{-1}\langle x, [t)\rangle \to \pi_T^{-1}[t)$$

- If $(x, t) \in G_{S \odot_D^a T}$ is a sample event:

$$d_{(x,t)}^{S \odot_D^a T}(q, r) = d_t^T(\psi_t\, q, r)$$

It is straightforward to see that $(G_S \odot_D^a G_T, o^{S \odot_D^a T}, d^{S \odot_D^a T})$ is a concrete probabilistic graph whose unfolding is **unfold** $S \odot_D^a$ **unfold** T.

D.2   Metropolis-hastings

Lemma 1. *For $\Gamma \vdash M : A$ without scores, $[\![M]\!](\gamma, [\![A]\!]) = 1$ for each $\gamma \in [\![\Gamma]\!]$.*

Correction of the Metropolis-Hastings algorithm We now focus on the proof of Theorem 3. Remember that we assume a program $x : \mathbb{X} \vdash Q : \mathbb{X}$ and a measurable computable function $q : \mathbb{X}^2 \to \mathbb{R}$ such that

$$\llbracket Q \rrbracket(x, Y) = \int_{y \in Y} q(x, y) \mathrm{d}y$$

We assume moreover a computable density function $d : \mathbb{X} \to \mathbb{R}$ and write D the distribution defined by d:

$$D(X) = \int_{x \in X} d(x) \mathrm{d}x.$$

First, let us define abstractly the transition kernel of the Metropolis-Hastings Markov chain. Consider $M : \mathbb{X} \rightsquigarrow \mathbb{X}$ a markov chain with density $m : \mathbb{X}^2 \to \mathbb{R}$. Such a pair (M, m) is (D, d)-valid when it satisfies the assumptions (3) and (4) of Theorem 3 namely:

1. if $d(x) > 0$ and $m(x, y) > 0$ then $m(y) > 0$
2. if $d(x)d(y) > 0$ then $m(x, y) > 0$ iff $m(y, x) > 0$.

We define the metropolis-hastings kernel as follows

$$\alpha_m(x, x') = \min(1, \frac{d(y) \cdot m(y, x)}{d(x) \cdot m(x, y)})$$

$$\mathbb{MH}_{M,m}(x, Y) = \int_{y \in Y} \alpha_m(x, y) \cdot M(x, \mathrm{d}y)$$

$$+ [x \in X] \cdot \int_{y \in Y} (1 - \alpha_m(x, y))M(s, \mathrm{d}y).$$

Let us first compute the transition kernel of the Markov chain $\mathtt{MH}(Q, q, d)$, and show it coincides with the usual definition:

Lemma 16. *We have* $\llbracket \mathtt{MH}(Q, q, d) \rrbracket = \mathbb{MH}_{\llbracket Q \rrbracket, q}$.

Proof. A simple unfolding of the definition of the measure-theoretic semantics.

As a result, we can apply the results of Tierney [19].

Lemma 17 (From [19], Theorem 1 and Corollary 2). *If there exists $n \in \mathbb{N}$ such that $\mathbb{MH}^n_{M,m}$ is D-strongly irreducible then $\mathbb{MH}^k_{M,m}(x, \cdot)$ converges towards the distribution D for any $x \in \mathbb{X}$ with $d(x) > 0$, ie.*

$$\lim_{k \to \infty} ||\mathbb{MH}^k_{M,m}(x, \cdot) - D|| = 0.$$

To prove Theorem 3, we need to show that if M^n is D-strongly irreducible, then so is $\mathbb{MH}^n_{M,m}$. Let us define the partial Metropolis-Hastings kernel $\mathbb{W}_{M,m} : \mathbb{X} \rightsquigarrow \mathbb{X}$

$$\mathbb{W}_{M,m}(x, Y) = \int_{y \in Y} \alpha_m(x, y) \times M(x, \mathrm{d}y) = \int_{y \in Y} \alpha_m(x, y) \times m(x, y) \mathrm{d}y.$$

It is straightforward to see that:

– If $W_{M,m}(x, Y) > 0$ implies $\mathbb{MH}_{M,m}(x, Y) > 0$: as a result $W_{\llbracket Q \rrbracket}$ D-strongly irreductible implies that $\mathbb{MH}_{M,m}$ is.
– By (D, d)-validity, it is easy to see that $m(x, y) > 0$ implies $\alpha_m(x, y) > 0$, hence $M(x, Y) > 0$ implies $W_{M,m}(x, Y) > 0$.
– $W_{M,m}$ has density $w_{M,m}(x, y) = \alpha_m(x, y) \times m(x, y)$.

If we have two markov chains $M_1, M_2 : \mathbb{X} \rightsquigarrow \mathbb{X}$ remember that their composition has density $m_1 \otimes m_2(x, z) = \int_{y \in \mathbb{X}} m_1(x, y) m_2(y, z) \mathrm{d}y$.

Lemma 18. *Consider $((M_k, m_k))$ a family of (D, d)-valid markov chains equipped with density. If the composition $W_n \circ \dots \circ W_1$ is strongly irreducible, then so is $\mathbb{W}_{M_n,m_n} \circ \dots \circ \mathbb{W}_{M_1,m_1}$.*

Proof. We show the case for $n = 2$ but the reasoning is the same for arbitrary n.

Consider $x \in \mathbb{X}$ with nonzero density and $Z \in \Sigma_{\mathbb{X}}$ with $D(Z) > 0$. We have

$$(\mathbb{W}_{M_2,m_2} \circ \mathbb{W}_{M_1,m_1})(x, Z) = \int_{z \in Z} \int_{y \in \mathbb{Y}} \alpha_{m_1}(x, y) \alpha_{m_2}(y, z) m_1(x, y) m_2(y, z) \mathrm{d}z\mathrm{d}y$$

$$= \int_{(z,y) \in C} \alpha_{m_1}(x, y) \alpha_{m_2}(y, z) m_1(x, y) m_2(y, z) \mathrm{d}z\mathrm{d}y$$

where C is defined as the set of $(z, y) \in Z \times \mathbb{X}$ such that $m_1(x, y) m_2(y, z)$. Since $(M_2 \circ M_1)(x, Z) = \int_{(z,y) \in C} m_1(x, y) m_2(y, z) \mathrm{d}z\mathrm{d}y > 0$ we know that this set has nonzero mesure (on the product measure of D with itself), and that $m_1(x, y)$ and $m_2(y, z)$ is nonzero for $(z, y) \in C$. By d-validity this implies that $\alpha_{m_1}(x, y)$ and $\alpha_{m_2}(y, z)$ are nonzero as well hence the whole integral is nonzero, as desired.

Theorem 3. *Assume that P and p satisfies the following properties:*

1. **Strong irreducibility**: *There exists $n \in \mathbb{N}$ such that for all $x \in \mathbb{A}$ and $X \in \Sigma_{\mathbb{A}}$ such that $D(X) \neq \emptyset$ and $d(x) > 0$, there exists $n \in \mathbb{N}$ such that $\llbracket P \rrbracket^n(x, X) > 0$.*
2. $\llbracket P \rrbracket(x, X') = \int_{x' \in X'} p(x, x')$.
3. *If $d(x) > 0$ and $p(x, y) > 0$ then $d(y) > 0$.*
4. *If $d(x) > 0$ and $d(y) > 0$, then $p(x, y) > 0$ iff $p(y, x) > 0$.*

Then, the limit of $\mathtt{MH}(P, p, d)$ for any initial state $x \in \mathbb{A}$ with $d(x) > 0$ is equal to D, the distribution obtained after normalising d.

Proof. We know that $\llbracket Q \rrbracket^n$ is strongly irreducible and $(\llbracket Q \rrbracket, q)$ is (D, d)-valid. As a result $\mathtt{MH}(Q, q, d)^n$ is strongly irreducible by Lemma 18, which entails the result by Lemma 17.

Lemma 19. *The kernel $\llbracket P_S \rrbracket$ has density p_S.*

Proof. Let us first show that $[\![\texttt{complete}]\!]$ has density:

$$q_{\texttt{complete}}(x, v, x_0, v_0, y, v') = \begin{cases} \prod_{s \in \texttt{sample}\, y \setminus (x_0 \cup x)} d_s(v'(s)) & x_0 \subseteq y \wedge v \subseteq v' \\ 0 & \text{otherwise} \end{cases}$$

The 0 clause is easy to see. Otherwise if $x_0 \subseteq y$ and $v \subseteq v'$, we prove the result by induction on the size of $y \setminus x_0$. If there are no sample in y outside $x_0 \cup x$, then it is clear that the density of (y, v') is 1 (no probabilistic choice). Otherwise, to go from (x_0, v_0) to (y, v'), it means we have sampled a $s \in y$ with exactly the same value as in v', and then continued, so we have:

$$q_{\texttt{complete}}(x, v, x_0, v_0, y, v') = d_s(v'(s)) \times q_{\texttt{complete}}(x, v, x_0 \cup \{s\}, v_0[s := v'(s)]).$$

which entails the desired equation.

From the formula for $\texttt{complete}$, it is easy to establish the formula for q_S.

Theorem 4. *The Markov chain P_S and density p satisfy the hypothesis of Theorem 3, as a result for any $(x, q) \in \mathscr{C}(E_S)$ the distribution $[\![\texttt{MH}(d_S, P_S, p_S)^n]\!]((x, q), \cdot)$ tends to μ_{norm}^P as n goes to infinity.*

Proof. **Strong-irreducibility.** Take n to be the size of G_S. Let $(x, v) \in E$ and $Y \in \Sigma_E$ with $d_S(Y) > 0$. Consider Y_0 the subset of Y of elements with nonzero. We must also have $d(Y_0) > 0$. For each $(y, v') \in Y_0$, we show that $q_S^k((x, v), (y', v')) > 0$ for k the number of elements of y which are not in x or such that $v(s) \neq v'(s)$. Since $k \leq n$ this implies the result

- If $x = y$, then clearly, we can take any sample event s with value r in x, resample it with the same value, and that shows that $q(x, y)$ has density $d_s(r) > 0$.
- Consider a minimal event $s \in y \setminus x$. We can resamble it to the value $v'(s)$ with density $d_s(v'(s))$ and then extend it to a maximal configuration (y', v'). By construction, (y', v') is closer to (v, y) and we can conclude by induction.

The fact that q_S is the density of Q_S is a consequence of Lemma 19. Property (3) and (4) are easy to establish.