

ULRR

A probabilistic theory of designs based on distributions

Item Type	Meetings and Proceedings
Authors	Bresciani, Riccardo;Butterfield, Andrew
Citation	4th Unifying Theories of Programming Symposium;
Download date	2026-06-06 17:33:17
Item License	https://creativecommons.org/licenses/by-nc-sa/1.0/
Link to Item	https://hdl.handle.net/10344/2588

A probabilistic theory of designs based on distributions

Riccardo Bresciani and Andrew Butterfield*

Foundations and Methods Group,
Trinity College Dublin,
Dublin, Ireland
{bresciar,butrfield}@scss.tcd.ie

Abstract. We present a theory of designs based on functions from the state space to real numbers, which we term *distributions*. This theory uses predicates, in the style of UTP, based on homogeneous relations between distributions, and is richer than the standard UTP theory of designs as it allows us to reason about probabilistic programs; the healthiness conditions H1–H4 of the standard theory are implicitly accounted for in the distributional theory we present. In addition we propose a Galois connection linkage between our distribution-based model of probabilistic designs, and the standard UTP model of (non-probabilistic) designs.

1 Introduction

The Unifying Theories of Programming (UTP) aims at a semantic framework where programs and specifications can be modelled as alphabetised relational predicates, capturing the semantic models normally used for their formal description [HJ98,DS06,But10,Qin10]: the advantage of this common framework is that of enabling formal reasoning on the integration of the different languages.

UTP relies on untyped predicate calculus: programs are expressed by means of logical predicates (programs are predicates! [Heh84,Hoa85]), and different theories can be given a UTP semantics, and here we focus on the theory of designs: this theory allows us to reason about the total correctness of programs from the perspective of what preconditions must be met in order to reach some given postconditions.

A challenging research question is how to add probability to the picture, and in particular how to integrate it in a framework where non-determinism is present: we use a framework based on distributions over the state space, where the predicates used involve a homogeneous relation between before- and after-distributions. This allows us to define a probabilistic theory of designs, which can handle programs where both probabilistic and non-deterministic choice co-exist.

* The present work has emanated from research supported by Science Foundation Ireland grant 08/RFP/CMS1277 and, in part, by Science Foundation Ireland grant 03/CE2/I303.1 to Lero – the Irish Software Engineering Research Centre.

This paper is structured as follows: we describe the background to UTP, with particular focus on the standard theory of designs in that framework (§2); introduce a probabilistic framework based on distributions over the state space (§3); present a probabilistic theory of designs in this new framework (§4) and discuss its application to a well-known example (§5); and conclude (§6).

2 Background

2.1 UTP

UTP uses second-order predicates to represent programs: they are used to express relations among a set of *observable variables* which constitute their alphabet.

Observable variables usually occur as both undecorated and decorated with a dash $'$: the former refer to states before the program starts (*before-states*), whereas the latter refer to the final states reached after the program has run (*after-states*).

For example, a program using two variables x and y might be characterised by having the set $\{x, x', y, y'\}$ as an alphabet, and the meaning of the assignment $x := y + 3$ would be described by the predicate

$$x' = y + 3 \wedge y' = y.$$

In effect UTP uses predicate calculus in a disciplined way to build up a relational calculus for reasoning about programs.

In addition to observations of the values of program variables, often we need to introduce observations of other aspects of program execution via so-called auxiliary variables. For example the theory of reactive programs uses three auxiliary variables — namely *ok*, *wait*, *tr*, *ref* — to keep track of information concerning the current program run, such as termination, reach of a stable state, refusals, . . .

A key notion in UTP is that of *healthiness conditions*: they are usually characterised as monotonic idempotent predicate transformers whose fixpoints characterise sensible (healthy) predicates. In other words they outlaw all arbitrary predicate calculus statements that describe predicates with no sense — an example is $\neg ok \Rightarrow ok'$, which describes a “program” that must terminate when not started.

This notion is closely related to that of *refinement*, defined as the universal closure¹ of reverse implication:

$$S \sqsubseteq P \triangleq [P \Rightarrow S]$$

Healthy predicates form a lattice under the ordering induced by the refinement relation.

¹ Square brackets denote universal closure, *i.e.* $[P]$ asserts that P is true for all values of its free variables.

The refinement calculus enables the derivation of an implementation P from a specification S : such derivation can be proven correct if P is a valid refinement of S .

Most UTP theories developed so far deal only with non-deterministic choice, nevertheless the introduction of a probabilistic choice operator is beneficial to many application requiring a quantitative approach, for example to evaluate reliability of programs.

Nevertheless some lines of research are moving along this direction. In [HS06] the authors present an approach to unification of probabilistic choice with standard constructs. They provide an example of how the laws of pGCL could be captured in UTP as predicates about program equivalence and refinement. However only an axiomatic semantics was presented, and the laws were justified via a Galois connection to an expectation-based semantic model.

The approach presented in [CS09] is that of decomposing non-deterministic choice into a combination of pure probabilistic choice and a unary operator that accounted for its non-deterministic behaviour. It is worth underlining a comment of theirs, on how still unsatisfactory theories are with respect to the issue of having probabilistic and demonic choice to coexist.

The UTP model described in [He10], which is used to give a UTP-style semantics to a probabilistic BPEL-like language, relates an initial state to a final probability distribution over states, rather than relating before-variables to corresponding after-variables of the same type.

We have previously presented an encoding of the semantics of the probabilistic guarded command language (pGCL) in the UTP framework [BB11, BB12]. This encoding captures pGCL programs as predicate-transformers, on predicates over probability before- and after-distributions.

In §3 we will present the underlying distributional framework, which we subsequently use in order to obtain a probabilistic theory of designs.

2.2 The standard theory of designs

Now that we have given a general overview of the UTP framework, we are going to focus on the theory of designs and present its UTP semantics.

The theory of designs patches the relational theory, in the sense that predicates from the relational theory fail to satisfy the following equality:

$$true; \mathcal{P} = true$$

In fact according to the relational theory $true$ is a left identity of the sequential composition operator:

$$\begin{aligned} true; \mathcal{P} &\equiv \exists \underline{v}_m \bullet true\{\underline{v}_m/\underline{v}'\} \wedge \mathcal{P}\{\underline{v}_m/\underline{v}\} \\ &\equiv \exists \underline{v}_m \bullet true \wedge \mathcal{P}\{\underline{v}_m/\underline{v}\} \\ &\equiv \exists \underline{v}_m \bullet \mathcal{P}\{\underline{v}_m/\underline{v}\} \end{aligned}$$

Which reduces to $true$ if $\underline{v} \in fv(\mathcal{P})$, or to \mathcal{P} otherwise.

This has disastrous consequences, as this enables us to show that a program can recover from a never-ending loop:

$$true * skip \equiv \mu X \bullet X \equiv \perp \equiv true$$

... which is surprising, to say the least.

The theory of designs uses an additional auxiliary variable $o\mathcal{K}$ (along with its dashed version $o\mathcal{K}'$) to record start (and termination) of a program.

A design (specification) is made of a precondition Pre that has to be met when the program starts, and if so the program establishes $Post$ upon termination, which is guaranteed:

$$o\mathcal{K} \wedge Pre \Rightarrow o\mathcal{K}' \wedge Post$$

for which we use the following shorthand:

$$Pre \vdash Post$$

The semantics of the assignment $x := y + 3$ in this theory is the following:

$$true \vdash x' = y + 3 \wedge y' = y$$

(if started, it will terminate, and the final value of x will equal the initial value of y plus three, with y unchanged).

The behaviour of $true$ with respect to sequential composition is the desirable one, as now we have:

$$\begin{aligned} true; (Pre \vdash Post) &\equiv true; o\mathcal{K} \wedge Pre \Rightarrow o\mathcal{K}' \wedge Post \\ &\equiv \exists o\mathcal{K}_m, \underline{v}_m \bullet true\{o\mathcal{K}_m/o\mathcal{K}'\}\{\underline{v}_m/\underline{v}'\} \wedge (o\mathcal{K}_m \wedge Pre\{\underline{v}_m/\underline{v}\} \Rightarrow o\mathcal{K}' \wedge Post) \\ &\equiv \exists o\mathcal{K}_m, \underline{v}_m \bullet true \wedge (o\mathcal{K}_m \wedge Pre\{\underline{v}_m/\underline{v}\} \Rightarrow o\mathcal{K}' \wedge Post) \\ &\equiv \exists o\mathcal{K}_m, \underline{v}_m \bullet o\mathcal{K}_m \wedge Pre\{\underline{v}_m/\underline{v}\} \Rightarrow o\mathcal{K}' \wedge Post \\ &\equiv true \end{aligned}$$

and therefore $true$ is a left zero for sequential composition.

Designs form a lattice, whose bottom and top elements are respectively:

$$abort \triangleq false \vdash false \equiv false \vdash true$$

and

$$miracle \triangleq true \vdash false \equiv \neg o\mathcal{K}$$

It should be noted that $miracle$ is a (infeasible) program that cannot be started.

Valid designs are predicates R which comply with four healthiness conditions [HJ98]. The first one (*unpredictability*, H1) excludes from observation all programs that have not started, and therefore restricts valid relations to those such that:

$$R = (o\mathcal{K} \Rightarrow R)$$

All H1-healthy predicates satisfy the left zero and left unit laws:

$$true; R = true \quad \text{and} \quad skip; R = R$$

The second one (*possible termination*, H2) states that a valid relation cannot require nontermination:

$$R\{false/o\kappa'\} \Rightarrow R\{true/o\kappa'\}$$

The third one (*dischargeable assumptions*, H3) states that preconditions cannot use dashed variables. All H3-healthy predicates satisfy the right unit law:

$$R; skip = R$$

The fourth one (*feasibility* or *excluded miracle*, H4) requires the existence of final values for the dashed variables that satisfy the relation:

$$\exists o\kappa', \underline{v}' \bullet R = true$$

H4 excludes *miracle* from the valid designs, and this implies that all H4-healthy predicates satisfy the right zero law:

$$R; true = true$$

This condition cannot be expressed as an idempotent healthiness transformer, and does not preserve the predicate lattice structure. It serves solely to identify and/or eliminate predicates that characterise infeasible behaviour.

Through our distributional framework (§4) we obtain a richer theory where corresponding healthiness conditions hold (§4.1), even without the introduction of the auxiliary variables $o\kappa, o\kappa'$. Moreover the use of distributions enables us to evaluate the probability both of termination and of meeting a set of arbitrary postconditions as a function of the initial distribution (which determines the probability of meeting any required precondition).

3 The distributional framework

In [BB11] we have presented a UTP framework to deal with demonic probabilistic programs.

This framework relies on the concept of *distributions* over the state space: a generic distribution is a real-valued function $\chi : \mathcal{S} \rightarrow \mathbb{R}$ that assigns a *weight* x_i (a real number) to each state σ_i in the state space \mathcal{S} . The mathematics we employ is valid provided the probabilities constitute what is known as a measure space [Hal50]. If the state is not finite, then the limitation we face is that property predicates (pre- and post-expectations, for example) can only talk about probabilities associated with sets of observations, rather than single ones — in effect χ has to be interpreted as a probability density function.

A state σ bundles all the information regarding program variables into a single observation, in a style shared with many presentations of *Circus*-like languages: program variable values are modelled with a single state observation $\sigma : \mathcal{V} \rightarrow \mathcal{W}$, which is treated as a finite map from variables (\mathcal{V}) to values (\mathcal{W}). This choice simplifies the treatment of alphabets to a considerable degree.

The weight of a distribution is defined as:

$$\|\chi\| \triangleq \sum_{\sigma \in \text{dom } \chi} \chi(\sigma)$$

This operation can be lifted to a set \mathcal{X} of distributions:

$$\|\mathcal{X}\| \triangleq \{\|\chi\| \mid \chi \in \mathcal{X}\}$$

Among all generic distributions, the following two sub-classes play important roles in our framework:

- a *weighting distribution* π has the property that for every state σ we have $0 \leq \pi(\sigma) \leq 1$ — we define two particular weighting distributions, ϵ and ι , as the ones mapping every state to 0 and 1 respectively. There is no limit for the distribution weight;
- a *probability distribution* δ is a weighting distribution with the additional property that $\|\delta\| \leq 1$.

We will use the term *sub-distribution* to refer to a probability distribution where $\|\delta\| < 1$ and the term *full distribution* to refer to a probability distribution where $\|\delta\| = 1$.

Generally speaking, it is possible to operate on distributions by lifting pointwise operators such as addition, multiplication and multiplication by a scalar. Analogously we can lift pointwise all traditional relations and functions on real numbers².

In the case of pointwise multiplication, it is interesting to see it as a way of “re-weighting” a distribution. We have a particular interest in the case when one of the operands is a weighting distribution π , as we will use this operation to give semantics to choice constructs. We opt for a postfix notation to write this operation, as this is an effective way of marking when pointwise multiplication happens in the operational flow: for example if we multiply the probability distribution δ by the weighting distribution π , we write this as $\delta\langle\pi\rangle$. We use notation ϵ and ι to denote the everywhere zero and unit distributions, respectively:

$$\epsilon(\sigma) = 0 \wedge \iota(\sigma) = 1, \quad \text{for all } \sigma$$

Given a condition (predicate on state) c , we can define the weighting distribution that maps every state where c evaluates to *true* to 1, and every other state to 0: as the value of each state can be seen as the boolean value of c in that state

² Distributions form a vector space, which we have explored elsewhere. We omit discussion of this aspect of our theory for clarity and brevity.

multiplied by 1, we overload the above notation and note this distribution as $\iota\langle c \rangle$. In general whenever we have the multiplication of a distribution by $\iota\langle c \rangle$, we can use the postfix operator $\langle c \rangle$ for short, instead of using $\langle \iota\langle c \rangle \rangle$. It is worth pointing out that if we multiply a probability distribution δ by $\iota\langle c \rangle$, we obtain a distribution whose weight $\|\delta\langle c \rangle\|$ is exactly the probability of being in a state satisfying c .

3.1 Assignment

A challenge we have faced has been describing how assignment, which is very much oriented towards individual variables, is given a semantics in terms of a distribution that involves complete entanglement of those variables. In effect an assignment statement $x := e$ involves a partial entanglement of variable x with the variables mentioned in e . In general as we build up larger programs using single assignment as the basic component we observe an increasing degree of entanglement, which can often be captured as an appropriate simultaneous assignment, so we shall work at this level here.

Given a simultaneous assignment $\underline{v} := \underline{e}$, where underlining indicates that we have lists of variables and expressions of the same length, we denote its effect on an initial probability distribution δ by $\delta\{\underline{e}/\underline{v}\}$. The postfix operator $\{\underline{e}/\underline{v}\}$ reflects the modifications introduced by the assignment — the intuition behind this, roughly speaking, is that all states σ where the expression \underline{e} evaluates to the same value $\underline{w} = \text{eval}_\sigma(\underline{e})$ are replaced by a single state $\sigma' = (\underline{v} \mapsto \underline{w})$ that maps to a probability that is the sum of the probabilities of the states it replaces.

$$(\delta\{\underline{e}/\underline{v}\})(\sigma') \triangleq (\sum \delta(\sigma) \mid \sigma' = \sigma \dagger \{\underline{v} \mapsto \text{eval}_\sigma(\underline{e})\})$$

Here we treat the state as a map, where \dagger denotes map override; this operator essentially implements the concept of “push-forward” used in measure theory, and is therefore a linear operator. An example is given in Figure 1.

Assignment preserves the overall weight of a probability distribution if \underline{e} can be evaluated in every state, and if not the assignment returns a sub-distribution, where the “missing” weight accounts for the assignment failing on some states (this failure prevents a program from proceeding and causes non-termination).

3.2 Programming constructs

The semantic definitions of various programming constructs are based on a homogeneous relation between distributions and are listed in Figure 2; we will now proceed to discuss each one.

The failing program *abort* is represented by the predicate $\|\delta'\| \leq \|\delta\|$, which captures the fact that it is maximally unpredictable, given that it cannot increase distribution weight. Such an increase would describe a program whose probability of termination was higher than that of it starting, and is infeasible.

The miraculous program *miracle* is defined as $(\delta = \epsilon) \wedge (\delta' = \epsilon)$: this is a difference in comparison with the standard UTP theory, where it is simply *false*.




Fig. 1. The assignment $x := y$ from an initial uniform distribution on the state space $\mathcal{S} = \{0, 1\} \times \{0, 1\}$.

$$\begin{aligned}
\mathit{abort} &\triangleq \|\delta'\| \leq \|\delta\| \\
\mathit{miracle} &\triangleq (\delta = \epsilon) \wedge (\delta' = \epsilon) \\
\mathit{skip} &\triangleq \delta' = \delta \\
\mathit{v} := \mathit{e} &\triangleq \delta' = \delta \{\mathit{e}/\mathit{v}\} \\
A; B &\triangleq \exists \delta_m \bullet A(\delta, \delta_m) \wedge B(\delta_m, \delta') \\
\mathit{choice}(A, B, \mathcal{X}) &\triangleq \exists \pi, \delta_A, \delta_B \bullet \pi \in \mathcal{X} \wedge A(\delta\{\pi\}, \delta_A) \wedge B(\delta\{\bar{\pi}\}, \delta_B) \wedge \delta' = \delta_A + \delta_B \\
c * A &\triangleq \mu X \bullet \mathit{choice}((A; X), \mathit{skip}, \{\iota\{c\}\})
\end{aligned}$$

Fig. 2. UTP semantics for different programming constructs.

This definition coincides with the standard one for most pairs of before- and after-distributions, with the exception of (ϵ, ϵ) : this makes sure that *miracle* is a unit for nondeterministic choice.

Program *skip* makes no changes and immediately terminates.

Assignment remaps the distribution as has already been discussed in §3.1.

Sequential composition is characterised by the existence of a “mid-point” distribution that is the outcome of the first program, and is then fed into the second. It should be noted at this juncture that we are quantifying over function quantities, such as δ or π — this makes our logic at least second-order, even if the state spaces are finite (the range $[0, 1]$ is not).

The choice operator is probably the strangest-looking element of the list: it takes a weighting distribution π , uses it with its complementary distribution $\bar{\pi} = \iota - \pi$) to weigh the distributions resulting from the left- and right-hand side respectively, and existentially quantifies it over the set of distributions $\mathcal{X} \subseteq \mathcal{D}_w$. We have termed this operator as the *generic choice* as it can emulate the behaviour of all standard choices (and more):

– for $\mathcal{X} = \{\iota\langle c \rangle\}$ we have conditional choice:

$$\begin{aligned} A \triangleleft c \triangleright B &= \mathbf{choice}(A, B, \{\iota\langle c \rangle\}) \\ &= \exists \delta_A, \delta_B \bullet A(\delta\langle c \rangle, \delta_A) \wedge B(\delta\langle -c \rangle, \delta_B) \wedge \delta' = \delta_A + \delta_B \end{aligned}$$

– for $\mathcal{X} = \{p \cdot \iota\}$ we have probabilistic choice:

$$\begin{aligned} A \oplus_p B &= \mathbf{choice}(A, B, \{p \cdot \iota\}) \\ &= \exists \delta_A, \delta_B \bullet A(p \cdot \delta, \delta_A) \wedge B((1-p) \cdot \delta, \delta_B) \wedge \delta' = \delta_A + \delta_B \end{aligned}$$

– for $\mathcal{X} = \mathcal{D}_w$ we have non-deterministic choice:

$$\begin{aligned} A \sqcap B &= \mathbf{choice}(A, B, \mathcal{D}_w) \\ &= \exists \pi, \delta_A, \delta_B \bullet A(\delta\langle \pi \rangle, \delta_A) \wedge B(\delta\langle \bar{\pi} \rangle, \delta_B) \wedge \delta' = \delta_A + \delta_B \end{aligned}$$

The usual notations for conditional, probabilistic and non-deterministic choice will be used as syntactic sugar in the remainder of this document. Program *abort* is a zero for non-deterministic choice, as entering $\|\delta'\| \leq \|\delta\|$ for B in the definition, results in

$$\exists \pi, \delta_A, \delta_B \bullet A(\delta\langle \pi \rangle, \delta_A) \wedge \|\delta_B\| \leq \|\delta\langle \bar{\pi} \rangle\| \wedge \delta' = \delta_A + \delta_B$$

which, after the one-point rule with $\delta_B = \delta' - \delta_A$ reduces to

$$\exists \pi, \delta_A \bullet A(\delta\langle \pi \rangle, \delta_A) \wedge \|\delta' - \delta_A\| \leq \|\delta\langle \bar{\pi} \rangle\|$$

We can take $\pi = \epsilon$ as a witness, which forces $\delta_A = \epsilon$ (by healthiness condition *Dist1*, see Section 3.3) and we obtain

$$(A(\epsilon, \epsilon) \wedge \|\delta' - \epsilon\| \leq \|\delta\|) \vee \exists \pi, \delta_A \bullet \dots$$

A consequence of *Dist1* is that $A(\epsilon, \epsilon)$ is always true for healthy A so this reduces to $\|\delta'\| \leq \|\delta\|$ in a disjunction with a predicate that it subsumes, and hence equivalent to *abort*.

As commonly seen in UTP, disjunction of two programs is a kind of choice (usually non-deterministic, in other theories), which here can be defined using generic choice:

$$A \vee B = \mathbf{choice}(A, B, \{\epsilon, \iota\}).$$

Disjunction is the usual semantics for non-deterministic choice, but here we see that non-deterministic choice has a richer behaviour as it exhibits more variability. Nevertheless with the appropriate definition of refinement we can introduce a concept of equivalence (*i.e.* two programs mutually refine each other), that restores the equivalence between disjunction and non-deterministic choice. [BB11]

Using the customary notation for conditional choice enlightens the definition of while-loops, which can be rewritten in a more familiar fashion as:

$$c * A \hat{=} \mu X \bullet (A; X) \triangleleft c \triangleright \mathit{skip}$$

They are characterized as fixpoints of the appropriate functional, with respect to the ordering defined by the refinement relation, details of which can be found in [MM04, BB11] and are beyond the scope of this paper.

These are the most significant elements and constructs that characterise our framework: this has been a presentation from a fairly high level, and it should have provided the reader with a working knowledge of the framework; a formal and rigorous definition of the elements presented so far is beyond the scope of this paper and can be found in [BB11], along with some soundness proofs.

3.3 Healthiness conditions

Before moving further on, we are going to list quickly the healthiness conditions that characterise this framework.

The first one (*feasibility*, **Dist1**) assures that for any program $\mathcal{P}(\delta, \delta')$ the probability of termination cannot be greater than that of having started:

$$\|\delta'\| \leq \|\delta\|$$

Another healthiness condition (*monotonicity*, **Dist2**), states that, for any deterministic program \mathcal{P} , increasing δ implies that the resulting δ' increases as well:

$$\mathcal{P}(\delta_1, \delta'_1) \wedge \mathcal{P}(\delta_2, \delta'_2) \wedge \delta_2 > \delta_1 \Rightarrow \delta'_2 \geq \delta'_1$$

A third healthiness condition is that multiplication by a (not too large and non-negative³) constant distributes through commands (*scaling*, **Dist3**):

$$\forall a \in \mathbb{R}^+ \wedge \|a \cdot \delta\| \leq 1 \bullet \mathcal{P}(\delta, \delta') \Leftrightarrow \mathcal{P}(a \cdot \delta, a \cdot \delta')$$

Finally the purely random non-deterministic model adopted in the distributional framework yields a fourth healthiness condition **Dist4** (*convexity*):

$$(\mathcal{P}_1 \sqcap \mathcal{P}_2)(\delta, \delta') \Rightarrow \delta' \geq \min(\mathcal{P}_1(\delta) \cup \mathcal{P}_2(\delta))$$

Here $\mathcal{P}_1(\delta)$ denotes the set of all δ' that satisfy $\mathcal{P}_1(\delta, \delta')$.

This poses restrictions on the space of possible program images, which is strictly a subset of $\wp\mathcal{D}$: this is analogous to the set $\mathbb{H}S$ from [MM04].

4 A probabilistic theory of designs

A distinguishing characteristic of designs is the use of the auxiliary variables $\mathcal{o}\mathcal{K}$ and $\mathcal{o}\mathcal{K}'$. They are not sufficient in a probabilistic setting, as we need to be able to express quantitative information about the program also in terms of it having started or finished. We argue that this information is embedded in the distributions used to express programming constructs.

³ Mathematically the relation holds also if this is not met, but in that case $a \cdot \delta$ is not a probability distribution.

In fact the variable δ records implicitly if the program has started, as for each state σ it gives a precise probability that the program is in that initial state.

If δ is a full distribution (*i.e.* $\|\delta\| = 1$), then the program has started with probability 1: in some sense we can translate the statement $o\mathcal{K} = \text{true}$ with the statement $\|\delta\| = 1$. Conversely a program for which $\delta = \epsilon$ has not started. Obviously there are all situations in between, where the fact of δ being a sub-distribution accounts for the program having started with probability $\|\delta\| < 1$.

Similarly if δ' is a full distribution, then the program terminates with probability 1: coherently we can translate the statement $o\mathcal{K}' = \text{true}$ with the statement $\|\delta'\| = 1$. In general the weight of δ' is the probability of termination: if the program reaches an after-distribution whose weight is strictly less than 1, then termination is not guaranteed (and in particular if $\delta' = \epsilon$ it is certain that it will not terminate).

4.1 From standard designs to probabilistic designs

Given a standard design $\mathcal{P}re \vdash \mathcal{P}ost$ we can easily derive the corresponding probabilistic design by using the observation above:

$$\begin{aligned} \mathcal{P}re \vdash \mathcal{P}ost &\equiv o\mathcal{K} \wedge \mathcal{P}re \Rightarrow o\mathcal{K}' \wedge \mathcal{P}ost \\ &\equiv \|\delta\| = 1 \wedge \mathcal{P}re \Rightarrow \|\delta'\| = 1 \wedge \mathcal{P}ost \\ &\equiv \|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1 \end{aligned}$$

This expression tells us that we have a valid design if whenever the before-distribution δ is a full distribution which is null everywhere $\mathcal{P}re$ is not satisfied (and therefore $\delta = \delta\langle \mathcal{P}re \rangle$), then the resulting after-distribution δ' is a full distribution which is null everywhere $\mathcal{P}ost$ is not satisfied (and therefore $\delta' = \delta'\langle \mathcal{P}ost \rangle$).

We can easily redefine assignment, in the same style as it has been redefined to make it a valid construct according to the theory of designs:

$$\begin{aligned} \underline{v} := \underline{e} \hat{=} \text{true} \vdash \delta' &= \delta\{\underline{e}/\underline{v}\} \\ &\equiv o\mathcal{K} \wedge \text{true} \Rightarrow o\mathcal{K}' \wedge \delta\{\underline{e}/\underline{v}\} \\ &\equiv \|\delta\| = 1 \Rightarrow \|\delta'\| = 1 \wedge \delta' = \delta\{\underline{e}/\underline{v}\} \end{aligned}$$

This states that an assignment is a valid design only if the expression e is defined everywhere in the state space: in fact undefinedness of e causes $\delta\{\underline{e}/\underline{v}\}$ to be a sub-distribution and therefore $\underline{v} := e$ reduces to *false*.

We can redefine *skip* in a similar way:

$$\begin{aligned} \text{skip} \hat{=} \text{true} \vdash \delta' &= \delta \\ &\equiv o\mathcal{K} \wedge \text{true} \Rightarrow o\mathcal{K}' \wedge \delta \\ &\equiv \|\delta\| = 1 \Rightarrow \|\delta'\| = 1 \wedge \delta' = \delta \\ &\equiv \|\delta\| = 1 \Rightarrow \delta' = \delta \end{aligned}$$

This new version of *skip* states that the after-distribution is the same as the before-distribution (and therefore it does not alter the weight, so this can be left implicit), but as any other design it reduces to *true* if δ is not a full distribution.

The bottom of the lattice is *abort*, which is again *true* as in the standard theory:

$$\begin{aligned}
\mathit{abort} &\triangleq \mathit{false} \vdash \mathit{false} \\
&\equiv \mathit{ok} \wedge \mathit{false} \Rightarrow \mathit{ok}' \wedge \mathit{false} \\
&\equiv \mathit{false} \Rightarrow \mathit{false} \\
&\equiv \mathit{true} \\
&\equiv \mathit{false} \Rightarrow \mathit{true} \\
&\equiv \mathit{ok} \wedge \mathit{false} \Rightarrow \mathit{ok}' \wedge \mathit{true} \\
&\equiv \mathit{false} \vdash \mathit{true}
\end{aligned}$$

The standard definition of the construct *chaos* is

$$\begin{aligned}
\mathit{chaos} &\triangleq \mathit{true} \vdash \mathit{true} \\
&\equiv \mathit{ok} \wedge \mathit{true} \Rightarrow \mathit{ok}' \wedge \mathit{true} \\
&\equiv \mathit{ok} \Rightarrow \mathit{ok}' \\
&\equiv \|\delta\| = 1 \Rightarrow \|\delta'\| = 1
\end{aligned}$$

This is a program that guarantees termination, but in an unspecified state. It is equivalent to:

$$\mathit{chaos} \equiv \mathit{true} \vdash \mathit{abort}_R,$$

where the subscript *R* indicates that we are talking of the relational version of *abort*, from Figure 2.

The top of the lattice is *miracle*:

$$\begin{aligned}
\mathit{miracle} &\triangleq \mathit{true} \vdash \mathit{false} \\
&\equiv \mathit{ok} \wedge \mathit{true} \Rightarrow \mathit{ok}' \wedge \mathit{false} \\
&\equiv \mathit{ok} \Rightarrow \mathit{false} \\
&\equiv \neg \mathit{ok} \\
&\equiv \neg(\|\delta\| = 1) \\
&\equiv \|\delta\| < 1
\end{aligned}$$

This is equivalent to

$$\mathit{miracle} \equiv \mathit{true} \vdash \mathit{miracle}_R.$$

Healthiness conditions These new definitions relying on the distributional framework satisfy the healthiness conditions H1–H4 as well (§2.2).

We can in fact prove that the following laws hold:

– left unit law:

$$\begin{aligned}
\mathit{skip}; \mathcal{P}re \vdash \mathcal{P}ost &\equiv (\|\delta\| = 1 \Rightarrow \delta' = \delta); (\|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1) \\
&\equiv \exists \delta_m \bullet (\|\delta\| = 1 \Rightarrow \delta_m = \delta) \wedge (\|\delta_m\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1) \\
&\equiv \|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1 \\
&\equiv \mathcal{P}re \vdash \mathcal{P}ost
\end{aligned}$$

– right unit law:

$$\begin{aligned}
\mathcal{P}re \vdash \mathcal{P}ost; \mathit{skip} &\equiv (\|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1); (\|\delta\| = 1 \Rightarrow \delta' = \delta) \\
&\equiv \exists \delta_m \bullet (\|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta_m\langle \mathcal{P}ost \rangle\| = 1) \wedge (\|\delta_m\| = 1 \Rightarrow \delta' = \delta_m) \\
&\equiv \|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1 \\
&\equiv \mathcal{P}re \vdash \mathcal{P}ost
\end{aligned}$$

– left zero law:

$$\begin{aligned}
\mathit{true}; \mathcal{P}re \vdash \mathcal{P}ost &\equiv \mathit{true}; (\|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1) \\
&\equiv \exists \delta_m \bullet \mathit{true} \wedge (\|\delta_m\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1) \\
&\equiv \exists \delta_m \bullet \|\delta_m\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1 \\
&\equiv \mathit{true}
\end{aligned}$$

– right zero law:

$$\begin{aligned}
\mathcal{P}re \vdash \mathcal{P}ost; \mathit{true} &\equiv (\|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta'\langle \mathcal{P}ost \rangle\| = 1); \mathit{true} \\
&\equiv \exists \delta_m \bullet (\|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta_m\langle \mathcal{P}ost \rangle\| = 1) \wedge \mathit{true} \\
&\equiv \exists \delta_m \bullet \|\delta\langle \mathcal{P}re \rangle\| = 1 \Rightarrow \|\delta_m\langle \mathcal{P}ost \rangle\| = 1 \\
&\equiv \mathit{true}
\end{aligned}$$

4.2 Recasting total correctness

The reason that led to the standard theory of designs was that programs fail to satisfy the left zero law in the relational theory.

In the distributional framework programming constructs do satisfy this law, as for any programming construct \mathcal{P} other than *miracle* it is never the case that $\delta \notin \mathit{fv}(\mathcal{P})$.

For this reason we have:

$$\begin{aligned}
\mathit{true}; \mathcal{P}(\delta, \delta') &\equiv \exists \delta_m \bullet \mathit{true} \wedge \mathcal{P}(\delta_m, \delta') \\
&\equiv \exists \delta_m \bullet \mathcal{P}(\delta_m, \delta') \\
&\equiv \mathit{true}
\end{aligned}$$

Similarly the right zero law is satisfied as well, along with the left and right unit laws: healthiness conditions equivalent to **H1–H4** hold here as well.

Following this observation it appears that restricting the reasoning to programs with guaranteed termination is somehow limiting, as guaranteed termination is not an actual real-world feature of programs: programs must be reasonably reliable, but failure is always a possibility.

The reason for this may be inherent to the fact that programs are run on hardware which is susceptible of failure, as well as being imputable to the way a program is designed (for example the implementation of a probabilistic algorithm where termination is probabilistic as well).

We can fully exploit the potential of the distributional framework towards modelling these situations by removing the constraints on the weights of the before- and after-distributions — so we use the programming constructs in Figure 2 exactly with the semantics presented there.

The role of preconditions and postconditions is that of restricting the range of acceptable before- and after-distributions (and therefore act as restrictions to be applied to δ and δ' respectively) — this allows us to express desirable characteristics of a program in great detail, for example:

- $\mathcal{P} \wedge \|\delta'\| = 1$ requires \mathcal{P} to guarantee termination;
- $\mathcal{P} \wedge \|\delta'\| > 0.95$ requires \mathcal{P} to terminate with at least 95% probability;
- $\mathcal{P} \wedge \|\delta'\langle \mathcal{Post} \rangle\| > 0.95$ requires \mathcal{P} to terminate with at least 95% probability in a state satisfying \mathcal{Post} ;
- $\mathcal{Pre} \Rightarrow \mathcal{P} \wedge \|\delta'\langle \mathcal{Post} \rangle\| > 0.95$ requires \mathcal{P} to terminate with at least 95% probability in a state satisfying \mathcal{Post} whenever it starts in a state satisfying \mathcal{Pre} ;
- $\|\delta\langle \mathcal{Pre} \rangle\| > 0.98 \Rightarrow \mathcal{P} \wedge \|\delta'\langle \mathcal{Post} \rangle\| > 0.95$ requires \mathcal{P} to terminate with at least 95% probability in a state satisfying \mathcal{Post} whenever the probability of \mathcal{Pre} being satisfied at the beginning is at least 0.98;
- ...

All healthiness conditions deriving from the distributional framework (**Dist1–Dist4**) obviously hold here as well; with a small modification we can recast the notion of total correctness by restricting **Dist1** to a variant **Dist1-TC** (which implies **Dist1**), stating that:

$$\|\delta\| = \|\delta'\|$$

This requires a program to terminate with the same probability p with which it has started:

$$\|\delta\| = p \wedge \mathcal{Pre} \Rightarrow \|\delta'\| = p \wedge \mathcal{Post}$$

4.3 Link with the standard model

Standard designs have observations $\mathit{o}\mathcal{K}, \mathit{o}\mathcal{K}', \sigma, \sigma'$.

$$\begin{aligned} \mathit{o}\mathcal{K}, \mathit{o}\mathcal{K}' &: \mathbb{B} \\ \sigma, \sigma' &: \mathcal{S} \end{aligned}$$

A standard design is a predicate $\mathcal{P}_S(\sigma, \sigma', \mathcal{OK}, \mathcal{OK}')$ that states that a program started (if \mathcal{OK} is true) in the state σ ends (if \mathcal{OK}' is true) in the state σ' .

Probabilistic designs have observations δ, δ'

$$\delta, \delta' : \mathcal{S} \rightarrow [0, 1]$$

A probabilistic design is a predicate $\mathcal{P}_D(\delta, \delta')$ stating that a before-distribution δ will be transformed into the after-distribution δ' .

Informally we require the two approaches to yield the same results when we are dealing with point distributions, *i.e.* when the probability of being in a given state is 1.

In order to formalise the link between these two worlds, we define the linking predicate L as:

$$\begin{aligned} L((\delta, \delta'), (\sigma, \sigma', \mathcal{OK}, \mathcal{OK}')) &\triangleq \mathcal{OK} \Leftrightarrow (\|\delta'\| = 1) \wedge \mathcal{OK}' \Leftrightarrow (\|\delta'\| = 1) \\ &\wedge \delta = \eta_\sigma \wedge \delta' = \eta_{\sigma'} \end{aligned}$$

Here the notation η_σ denotes the point distribution returning 1 for state σ , and 0 elsewhere.

This linking predicate allows us to introduce the following *Galois connections*; first we define the weakest probabilistic design corresponding to a standard design \mathcal{P}_S :

$$\forall \sigma, \sigma', \mathcal{OK}, \mathcal{OK}' \bullet L((\delta, \delta'), (\sigma, \sigma', \mathcal{OK}, \mathcal{OK}')) \Rightarrow \mathcal{P}_S(\sigma, \sigma', \mathcal{OK}, \mathcal{OK}')$$

Analogously, the strongest standard design corresponding to a probabilistic design \mathcal{P}_D is:

$$\exists \delta, \delta' \bullet L((\delta, \delta'), (\sigma, \sigma', \mathcal{OK}, \mathcal{OK}')) \wedge \mathcal{P}_D(\delta, \delta')$$

It is easy to see that all programming constructs from the probabilistic theory that have homologue ones in the standard theory are linked to them, with the restriction of operating only on point distributions, otherwise they reduce to *abort*.

Weakening the link This linking predicate is a bit too strong, as it maps many interesting program constructs to the aborting program: an example is that of generic choice, which has no homologue in the standard theory. Ideally a better option would be to relax some constraints and to map generic choice to non-deterministic choice rather than to *abort*.

In other words we are aiming at a link that loses all probabilistic information about the possible after-states and flattens it to a mere list of them.

This is not straightforward, as the linking predicate L in some sense verifies consistency of δ with respect to σ, \mathcal{OK} and of δ' with respect to σ', \mathcal{OK}' : when the support⁴ of the distribution has more than one element, the relation between δ and a state from its domain is too weak to be useful.

⁴ We remind the reader that the support of a function is the set of points where the function is not zero-valued: $\text{supp}(\delta) \triangleq \text{dom}(\delta) \setminus \ker(\delta)$.

The situation is similar to that of a 3D-space, where dots are characterised by their x, y, z coordinates: a transformation creates a space with coordinates x', y', z' , whose relation with the undashed coordinates cannot in general be captured by a relation that mentions only one undashed and one dashed coordinate.

So far we have seen standard designs as relations:

$$\mathcal{P}_S : \mathcal{S} \times \mathbb{B} \rightarrow \mathcal{S} \times \mathbb{B}$$

but in order to build a more useful link we turn to this other interpretation:

$$\mathcal{P}_{\wp S} : \mathcal{S} \times \mathbb{B} \rightarrow \wp \mathcal{S} \times \mathbb{B}$$

which maps a state to what we may term its program image $\mathcal{P}(\sigma)$ (as it is a similar concept to that of program image introduced in §3), which contains all of the possible after-states reachable from a given before-state:

$$\mathcal{P}(\sigma) = \{\sigma' \mid \mathcal{P}_S(\sigma, \sigma')\}$$

All deterministic standard constructs map a state to a singleton set, whereas non-deterministic choice maps it to larger sets.

The interpretation of the predicate $\mathcal{P}_{\wp S}(\sigma, \alpha', \mathcal{OK}, \mathcal{OK}')$ is therefore that \mathcal{P} has started (if \mathcal{OK} is true) in the state σ and has ended (if \mathcal{OK}' is true) in a state $\sigma' \in \alpha'$:

$$\mathcal{P}_{\wp S}(\sigma, \alpha', \mathcal{OK}, \mathcal{OK}') \equiv \bigvee_{\sigma' \in \alpha'} \mathcal{P}_S(\sigma, \sigma', \mathcal{OK}, \mathcal{OK}')$$

With this in mind we can define the following linking predicate:

$$\begin{aligned} L_{\wp}((\delta, \delta'), (\sigma, \alpha', \mathcal{OK}, \mathcal{OK}')) \triangleq \mathcal{OK} \Leftrightarrow (\|\delta'\| = 1) \wedge \mathcal{OK}' \Leftrightarrow (\|\delta'\| = 1) \\ \wedge \delta = \eta_{\sigma} \wedge \text{supp}(\delta') = \alpha' \end{aligned}$$

We can state the variants of the Galois connections above as:

$$\begin{aligned} \forall \sigma, \alpha', \mathcal{OK}, \mathcal{OK}' \bullet L_{\wp}((\delta, \delta'), (\sigma, \alpha', \mathcal{OK}, \mathcal{OK}')) \Rightarrow \mathcal{P}_{\wp S}(\sigma, \alpha', \mathcal{OK}, \mathcal{OK}') \\ \exists \delta, \delta' \bullet L_{\wp}((\delta, \delta'), (\sigma, \alpha', \mathcal{OK}, \mathcal{OK}')) \wedge \mathcal{P}_D(\delta, \delta') \end{aligned}$$

5 An example: interaction of probabilistic and non-deterministic choice

This brief classical example is meant to show the interaction and the difference between probabilistic and non-deterministic choice: we will use this to show the effect of projecting the probabilistic design on the space of standard designs.

Let us take these two simple programs:

$$\begin{aligned} A \triangleq x := 0 \sqcap x := 1 \ ; \ y := 0 \ \frac{1}{2} \oplus \ y := 1 \\ B \triangleq x := 0 \ \frac{1}{2} \oplus x := 1 \ ; \ y := 0 \sqcap y := 1 \end{aligned}$$

$$\begin{aligned}
A &\hat{=} x := 0 \sqcap x := 1 ; y := 0 \frac{1}{2} \oplus y := 1 \\
&\equiv \exists \pi \bullet \delta' = \delta \langle \pi \rangle \{0/x\} + \delta \langle \bar{\pi} \rangle \{1/x\} ; \delta' = 1/2 \cdot \delta \{0/y\} + 1/2 \cdot \delta \{1/y\} \\
&\equiv \exists \pi, \delta_m \bullet \delta_m = \delta \langle \pi \rangle \{0/x\} + \delta \langle \bar{\pi} \rangle \{1/x\} \wedge \delta' = 1/2 \cdot \delta_m \{0/y\} + 1/2 \cdot \delta_m \{1/y\} \\
&\equiv \exists \pi \bullet \delta' = 1/2 \cdot (\delta \langle \pi \rangle \{0/x\} + \delta \langle \bar{\pi} \rangle \{1/x\}) \{0/y\} + \\
&\quad + 1/2 \cdot (\delta \langle \pi \rangle \{0/x\} + \delta \langle \pi \rangle \{1/x\}) \{1/y\} \\
&\equiv \exists \pi \bullet \delta' = \delta'_A(\pi) \wedge \delta'_A(\pi) \hat{=} 1/2 \cdot \delta \langle \pi \rangle \{0/x\} \{0/y\} + 1/2 \cdot \delta \langle \bar{\pi} \rangle \{1/x\} \{0/y\} + \\
&\quad + 1/2 \cdot \delta \langle \pi \rangle \{0/x\} \{1/y\} + 1/2 \cdot \delta \langle \bar{\pi} \rangle \{1/x\} \{1/y\} \\
\\
B &\hat{=} x := 0 \frac{1}{2} \oplus x := 1 ; y := 0 \sqcap y := 1 \\
&\equiv \delta' = 1/2 \cdot \delta \{0/x\} + 1/2 \cdot \delta \{1/x\} ; \exists \pi \bullet \delta' = \delta \langle \pi \rangle \{0/y\} + \delta \langle \bar{\pi} \rangle \{1/y\} \\
&\equiv \exists \pi, \delta_m \bullet \delta_m = 1/2 \cdot \delta \{0/x\} + 1/2 \cdot \delta \{1/x\} \wedge \delta' = \delta_m \langle \pi \rangle \{0/y\} + \delta_m \langle \bar{\pi} \rangle \{1/y\} \\
&\equiv \exists \pi \bullet \delta' = (1/2 \cdot \delta \{0/x\} + 1/2 \cdot \delta \{1/x\}) \langle \pi \rangle \{0/y\} + \\
&\quad + (1/2 \cdot \delta \{0/x\} + 1/2 \cdot \delta \{1/x\}) \langle \bar{\pi} \rangle \{1/y\} \\
&\equiv \exists \pi \bullet \delta' = \delta'_B(\pi) \wedge \delta'_B(\pi) \hat{=} 1/2 \cdot \delta \{0/x\} \langle \pi \rangle \{0/y\} + 1/2 \cdot \delta \{1/x\} \langle \pi \rangle \{0/y\} + \\
&\quad + 1/2 \cdot \delta \{0/x\} \langle \bar{\pi} \rangle \{1/y\} + 1/2 \cdot \delta \{1/x\} \langle \bar{\pi} \rangle \{1/y\}
\end{aligned}$$

Fig. 3. Programs A and B .

In Figure 3 we have worked out parametric expression for the final distribution for each program — they are parametric in the weighting distribution π which accounts for the non-deterministic choice performed in both programs:

$$\begin{aligned}
\delta'_A(\pi) &\hat{=} 1/2 \cdot (\delta \langle \pi \rangle \{0/x\} \{0/y\} + \delta \langle \bar{\pi} \rangle \{1/x\} \{0/y\} + \delta \langle \pi \rangle \{0/x\} \{1/y\} + \delta \langle \bar{\pi} \rangle \{1/x\} \{1/y\}) \\
\delta'_B(\pi) &\hat{=} 1/2 \cdot (\delta \{0/x\} \langle \pi \rangle \{0/y\} + \delta \{1/x\} \langle \pi \rangle \{0/y\} + \delta \{0/x\} \langle \bar{\pi} \rangle \{1/y\} + \delta \{1/x\} \langle \bar{\pi} \rangle \{1/y\})
\end{aligned}$$

The two after-distributions are very similar, but with one crucial difference: the position of $\langle \pi \rangle$, which clearly marks when the non-deterministic choice was made; this is reflected in the different after-distributions reached by each program:

- $\forall \pi \bullet \|\delta'_A(\pi) \langle x = y \rangle\| = 1/2$, *i.e.* regardless of the non-deterministic choice and of the initial distribution program A terminates in a state satisfying the condition $x = y$ with probability $1/2$, whereas in program B we cannot remove the dependence on π so if we turn to a worst-case analysis (in the non-deterministic choice the left-hand side $y := 0$ is picked whenever $x = 1$, *i.e.* $\pi = \iota \langle x = 1 \rangle$) we have that $\|\delta'_B(\iota \langle x = 1 \rangle) \langle x = y \rangle\| = 0$ and therefore the minimum guaranteed probability that $x = y$ is 0;
- viceversa for program B we can show that $\forall \pi \bullet \|\delta'_B(\pi) \langle x = 1 \rangle\| = 1/2$ similarly as above, and so the probability that $x = 1$ after program B is $1/2$, whereas if we take program A we can see that if $\pi = \iota$ (*i.e.* in the non-deterministic choice the left-hand side $x := 0$ is always picked) then we have that $\|\delta'_A(\iota) \langle x = 1 \rangle\| = 0$, so the minimum guaranteed probability that $x = 1$ is 0.

We are now going to derive the strongest standard design corresponding to A and B using the linking predicate L_\wp :

$$\begin{aligned}
& \exists \delta, \delta' \bullet L_\wp((\delta, \delta'), (\sigma, \alpha', \mathcal{OK}, \mathcal{OK}')) \wedge A(\delta, \delta') \\
\equiv & \quad \text{Definition of } L_\wp \\
& \exists \delta, \delta' \bullet (\mathcal{OK} \Leftrightarrow (\|\delta'\| = 1) \wedge \mathcal{OK}' \Leftrightarrow (\|\delta'\| = 1) \wedge \delta = \eta_\sigma \wedge \text{supp}(\delta') = \alpha') \wedge A(\delta, \delta') \\
\equiv & \quad A \text{ and } B \text{ return after-distributions with the same support} \\
& \exists \delta, \delta' \bullet (\mathcal{OK} \Leftrightarrow (\|\delta'\| = 1) \wedge \mathcal{OK}' \Leftrightarrow (\|\delta'\| = 1) \wedge \delta = \eta_\sigma \wedge \text{supp}(\delta') = \alpha') \wedge B(\delta, \delta')
\end{aligned}$$

The last line of this derivation clearly shows the effect of the link, which flattens out all probabilistic information and as a result the programs A and B are mapped to the same program in the world of standard designs: such a program records that a choice⁵ was made, but there is no discrimination among choices of different kind. Moreover what matters is the set of possible after-states and this is also not affected by altering the order in which the choices are made.

6 Conclusion

We have presented a probabilistic theory of designs, which relies on a UTP-style framework based on distributions over the state space.

We have shown that we are able to embed the standard UTP theory by requiring guaranteed termination from all program constructs, and treating them as the aborting program otherwise.

We have later relaxed this constraint to be able to reason about probabilistic programs: the advantages of this richer approach is that it allows us to express in fine detail the desired behaviour of a program, including its probabilistic aspects.

References

- [BB11] Riccardo Bresciani and Andrew Butterfield. Towards a UTP-style framework to deal with probabilities. Technical Report TCD-CS-2011-09, FMG, Trinity College Dublin, Ireland, August 2011.
- [BB12] Riccardo Bresciani and Andrew Butterfield. A UTP semantics of pGCL as a homogeneous relation. In *iFM 2012*, 2012.
- [But10] Andrew Butterfield, editor. *Unifying Theories of Programming, Second International Symposium, UTP 2008, Dublin, Ireland, September 8-10, 2008, Revised Selected Papers*, volume 5713 of *Lecture Notes in Computer Science*. Springer, 2010.
- [CS09] Yifeng Chen and Jeff W. Sanders. Unifying probability with nondeterminism. In *FM 2009, LNCS 5850*, pages 467–482, 2009.

⁵ Conditional choice is excluded, as it is not really a choice but rather a different evolution of the program which was determined by the current program state.

- [DS06] Steve Dunne and Bill Stoddart, editors. *Unifying Theories of Programming, First International Symposium, UTP 2006, Walworth Castle, County Durham, UK, February 5-7, 2006, Revised Selected Papers*, volume 4010 of *Lecture Notes in Computer Science*. Springer, 2006.
- [Hal50] Paul R. Halmos. *Measure Theory*. University Series in Higher Mathematics. D. Van Nostrand Company, Inc., Princeton, New Jersey, 1950.
- [He10] Jifeng He. A probabilistic bpeL-like language. In Qin [Qin10], pages 74–100.
- [Heh84] Eric C. R. Hehner. Predicative programming part i&ii. *Commun. ACM*, 27(2):134–151, February 1984.
- [HJ98] C. A. R. Hoare and He Jifeng. *Unifying Theories of Programming*. Prentice Hall International Series in Computer Science, 1998.
- [Hoa85] C. A. R. Hoare. Programs are predicates. In *Proceedings of a discussion meeting of the Royal Society of London on Mathematical logic and programming languages*, pages 141–155, Upper Saddle River, NJ, USA, 1985. Prentice-Hall.
- [HS06] Jifeng He and Jeff W. Sanders. Unifying probability. In Dunne and Stoddart [DS06], pages 173–199.
- [MM04] Annabelle McIver and Carroll Morgan. *Abstraction, Refinement And Proof For Probabilistic Systems (Monographs in Computer Science)*. SpringerVerlag, 2004.
- [Qin10] Shengchao Qin, editor. *Unifying Theories of Programming - Third International Symposium, UTP 2010, Shanghai, China, November 15-16, 2010. Proceedings*, volume 6445 of *Lecture Notes in Computer Science*. Springer, 2010.

A Keisli Composition

Assume a semantic model of the form $S \rightarrow \mathbb{F}S$ where \mathbb{F} is a type constructor (functor). The question that naturally arises is how to compose such functions, i.e., given $p : S \rightarrow \mathbb{F}T$ and $q : T \rightarrow \mathbb{F}U$, how do we compose these to get $(p; q) : S \rightarrow \mathbb{F}U$? The standard solution for this is Kleisli lifting and composition which involves two functions with the following signatures:

$$\eta_S : S \rightarrow \mathbb{F}S \qquad _{}^* : (S \rightarrow \mathbb{F}T) \rightarrow (\mathbb{F}S \rightarrow \mathbb{F}T)$$

that obey the following laws:

$$\eta_S^* = id_{\mathbb{F}S} \qquad p^* \circ \eta_S = p \qquad (q^* \circ p)^* = q^* \circ p^*$$

The intuition behind these is best understood in a diagram:

$$\begin{array}{ccccc}
 \mathbb{F}S & \xrightarrow{p^*} & \mathbb{F}T & \xrightarrow{q^*} & \mathbb{F}U \\
 \eta_S \uparrow & \nearrow p & \eta_T \uparrow & \nearrow q & \uparrow \eta_U \\
 S & & T & & U
 \end{array}$$

The Kleisli composition of p and q is given by $q^* \circ p$, where \circ denotes regular function composition.

In this paper $\mathbb{F}S = \mathbb{C}(S \rightarrow [0, 1])$, and we do not use the full lifting (which results in $\mathbb{C}(S \rightarrow [0, 1]) \rightarrow \mathbb{C}(S \rightarrow [0, 1])$), but instead lift partway to get $((S \rightarrow [0, 1]) \rightarrow \mathbb{C}(S \rightarrow [0, 1]))$. This “partway” lifting is one of the stages in giving an explicit definition of the full lifting.