

First Steps Towards Probabilistic Iris: A Separation Logic with Independence, Conditioning, and Dynamic Heap Allocation

Janine Lohse
MPI-SWS
Germany
jlohse@mpi-sws.org

Tim Rohde
MPI-SWS and
Saarland University
Germany
timrohde@mpi-sws.org

Jimmy Xin
MPI-SWS
Germany
jimmyxin@mpi-sws.org

Niklas Mück
MPI-SWS
Germany
mueck@mpi-sws.org

Iona Kuhn
CISPA Helmholtz Center
for Information Security
Germany
iona.kuhn@cispa.de

Derek Dreyer
MPI-SWS
Germany
dreyer@mpi-sws.org

Deepak Garg
MPI-SWS
Germany
dg@mpi-sws.org

Emanuele D’Osualdo
University of Konstanz
Germany
emanuele.dosualdo@uni-
konstanz.de

Abstract

There has recently been exciting progress in the realm of *probabilistic separation logics*. An important subclass of these—including PSL, Lilac, BLUEBELL, and PCOL—are *general-purpose probabilistic logics* (or GPLs, for short), meaning that they provide primitive Hoare-style assertions about probability distributions on the program state, along with powerful modularity principles like *independence* and *conditioning*. However, none of these logics support reasoning about dynamically allocated memory (*i.e.*, pointers into a heap), let alone the more sophisticated resource algebra-based ghost state of modern separation logics like Iris. We argue that this is due to a fundamental obstacle: since the shape of memory (and identity of memory locations) may differ under different random outcomes, it is unclear how pointer ownership can be harmonized with probabilistic independence and conditioning. Furthermore, none of the existing GPLs that support both independence and conditioning have, to our knowledge, been mechanized in a proof assistant.

In this paper, we take a substantial first step towards a marriage of GPLs and modern separation logics like Iris, in the form of **Amaryllis**. Amaryllis is the first GPL to support independence and conditional reasoning while also handling dynamic memory allocation. To overcome the aforementioned obstacle, we propose a new *indexed valuation*-style model of probabilistic assertions, whereby ownership of a standard Iris-style resource (*e.g.*, heaps) can be promoted to ownership at the level of distributions in a generic fashion. We then show how to adapt the central Iris notions of *frame-preserving update*, *authoritative resource algebras*, and the *weakest precondition modality* to be sound for probabilistic reasoning and validate dynamic allocation. Finally, we have mechanized all our results in the Rocq proof assistant and developed an Iris-based proof mode for conducting proofs within Amaryllis.

1 Introduction

The past few years have seen a flurry of work on program logics for reasoning modularly about probabilistic programs. Roughly speaking, these probabilistic logics can be divided into two types, which

may be understood as “property-specific” vs. “general-purpose”. *Property-specific probabilistic logics* (or PPLs) [1, 22, 13, 11, 12, 18] encode various specific probabilistic properties of interest (*e.g.*, error bounds, expected time complexity, etc.) and then derive rules for reasoning about them within the framework of standard Hoare or separation logics [26, 25] (*e.g.*, Iris [17, 16]). *General-purpose probabilistic logics* (or GPLs) [7, 9, 20, 5, 32] instead provide native support for program predicates describing *distributions* over program states, as well as logical connectives describing *independence* and *conditioning* of distributions, neither of which are directly expressible in the property-specific logics.

Thus far, the PPLs have exhibited the advantage that, since they are built from (in some sense) more standard parts (*e.g.*, defining some appropriate form of ghost state and applying the standard Iris methodology), they are easier to construct and soundly apply in the context of more fully-featured programming languages (*e.g.*, with higher-order state). In contrast, the GPLs have the potential advantage of being more general and expressive, supporting a wider range of modular probabilistic reasoning principles in one logic.

However—notably—none of the existing GPLs support reasoning about *dynamically allocated memory* (*i.e.*, pointers into a heap), let alone the more sophisticated resource algebra-based ghost state of modern separation logics like Iris. This may come as a surprise, given that dynamically allocated pointers were the *raison d’être* of the original version of separation logic [25]. But there is good reason for it: multiple, fundamental technical challenges must be overcome in order to harmonize GPLs’ support for independence and conditioning—implemented using a non-standard model of separating conjunction—with the *resource*-based (*e.g.*, heap) model of separation in traditional separation logics.

As a long-term goal, we aim to explore a synthesis of PPLs and GPLs in the form of Probabilistic Iris: an adaptation of the Iris framework that would support rich programming language features and modern separation logic features *alongside* primitive assertions about probability distributions. In this paper, we take a substantial first step towards that goal with **Amaryllis**, the first GPL to support reasoning about Iris-style resources, including dynamically allocated memory (heaps). Though we use heaps as a motivation, Amaryllis’s approach is not actually specific to heap resources at

all: the logic is parametric in a (non-step-indexed) Iris-style resource algebra describing stateful, non-probabilistic resources, and shows how to generically lift such a resource algebra to a model of probabilistic assertions supporting independence, conditioning, and modular reasoning about state. This generic lifting requires a fundamental rethink of various aspects of existing GPL models. To build added confidence in the resulting constructions, we have fully mechanized Amaryllis in the proof assistant (Rocq/Iris).

Contributions. We make the following contributions:

- (1) We observe that existing GPL models—which interpret assertions as predicates on distributions over program states—run into a major road block when we incorporate dynamic memory allocation because they do not account for the possibility that the shape of memory (and identity of memory locations) may differ under different random outcomes.
- (2) We resolve this problem by proposing an *indexed-valuation*-style model for GPLs in which probabilistic resources comprise two parts: (1) a probability distribution on *outcomes* (i.e., results of coin flips), and (2) a mapping from outcomes to separation logic resources (e.g., heaps). In this model, separating conjunction is defined *per random outcome*, rather than requiring disjointness of heaps across all outcomes as prior models do. Consequently, it allows different random outcomes to result in differently shaped heaps, enabling a clean accounting of dynamic heap allocation.
- (3) On top of this probabilistic resource model, we build a program logic, Amaryllis, which supports key reasoning principles of prior GPLs—notably, independence and conditional reasoning—while also supporting dynamic allocation.
- (4) Amaryllis is formalized in the Iris framework. Towards that end, we show for the first time how to adapt central concepts in Iris—such as “frame-preserving update” and “the authoritative resource algebra”, which are used in turn to derive a weakest-precondition modality—so that they remain sound under probabilistic (and conditional) reasoning.
- (5) We demonstrate the effectiveness of Amaryllis on a range of illustrative examples.
- (6) We have mechanized all results of this paper in the Rocq proof assistant, and also developed an Iris-based proof mode for conveniently conducting proofs within Amaryllis. As advised by the program chairs, we have made the (anonymized) proof repository available online [3].

Non-goals, limitations, and future work. In order to focus attention on the core problem of handling dynamic heap allocation and Iris-style resources, we place several restrictions on the probabilistic setting and the types of programs we consider. Specifically, we restrict attention to terminating programs (as Lilac [20] and BLUEBELL [5] do) and discrete probability distributions (as BLUEBELL [5] and pCOL [32] do). We also do not attempt to incorporate Iris’s mechanisms for step-indexing [4, 2], higher-order ghost state [15], and concurrency/invariants [17]. Lastly, for simplicity, we only consider distributions with finite support. In future work, we expect that it should be possible to generalize Amaryllis to handle countably infinite support and almost-sure termination of unbounded loops, but that some of Amaryllis’s proof rules for conditional reasoning will not apply to continuous distributions.

Despite present restrictions, the mechanization of Amaryllis required a significant effort, including approximately 50K lines of Rocq proof.

2 Key Ideas

In this section we give an overview of the main contributions of Amaryllis.

2.1 State-of-the-Art Reasoning in GPLs

Amaryllis’s aim is to provide a modular and expressive program logic for stateful probabilistic languages, able to prove properties of the state distribution induced by such programs. Example of properties are: bounds on the probability of certain events, or on expected values, the exact distribution of a certain variable, the (conditional) probabilistic independence of two variables.

The main judgments of general-purpose probabilistic logics (GPLs) are Hoare triples of the form $\{P\} e \{V. Q(V)\}$. Here, e is a probabilistic expression, the semantics of which can be seen as a function taking a distribution over state and returning a joint distribution of state and return value. The *precondition* P is an assertion constraining the input distribution, while the *postcondition* Q constrains the output distribution. A *random T variable* (for some type T) is a function¹ from the outcomes of a distribution to T . The postcondition is parametric on V , which is a random value variable representing the return value. As an example, $\{X \sim \text{Ber}_{1/2}\} e \{V. V \sim \text{Ber}_{1/2}\}$ states that the program e returns a value distributed as a fair coin if X is initially distributed as a fair coin.

Amaryllis builds on previous work, starting from the seminal PSL [9] and further developed in Lilac [20], BLUEBELL [5], and pCOL [32]. The success of these logics lies in modular reasoning, which is enabled by two crucial ideas: (1) *probabilistic independence* as separation, and (2) *conditioning* as a modality.

Independence as separation. PSL introduced the idea that separating conjunction $P * Q$ can model probabilistic independence in a logic of assertions over distributions. For example, an assertion $X \sim \text{Ber}_{1/2} * Y \sim \text{Ber}_{1/2}$ does not just state that both X and Y are distributed as fair coins, but also that they are *independent*: knowing the outcome of one does not reveal any information about the outcome of the other. More precisely, the probability of seeing both v in X and w in Y is the product of the two constituent probabilities.

Internalizing independence as separation introduces a very useful principle of *locality*, expressed as the Frame rule of Separation Logic:

$$\frac{\text{FRAME} \quad \{P\} e \{V. Q(V)\}}{\{P * R\} e \{V. Q(V) * R\}}$$

Suppose, for instance, that we have $\{\text{True}\} e \{X. X \sim \text{Ber}_{1/2}\}$; then by **FRAME** we can prove $\{Y \sim \text{Ber}_{1/2}\} e \{X. X \sim \text{Ber}_{1/2} * Y \sim \text{Ber}_{1/2}\}$ without having to verify e again from the larger precondition. Since generating/consuming independent variables is commonplace in probabilistic programs, framing turns out to be extremely valuable.

¹Random variables are usually required to be measurable in some ambient probability space. To avoid confusion we will explicitly state, when relevant, if we assume them to be measurable, and in which space they are.

Conditioning as a modality. However, framing alone is not enough to achieve modularity in the probabilistic setting. Imagine we are provided with the triple

$$\forall b \in \mathbb{B}. \{[\ell \mapsto b]\} f(\ell) \{[\ell \mapsto \neg b]\} \quad (1)$$

where $[\cdot]$ states that the predicate inside is true with probability 1. Here the contents of location ℓ are assumed to be the deterministic value b . Intuitively, we should be able to prove that if the contents of ℓ are distributed as a fair coin, written $\ell \rightsquigarrow \text{Ber}_{\frac{1}{2}}$, then

$$\{\ell \rightsquigarrow \text{Ber}_{\frac{1}{2}}\} f(\ell) \{\ell \rightsquigarrow \text{Ber}_{\frac{1}{2}}\}. \quad (2)$$

holds. But the frame rule alone is not sufficient to derive (2) from (1)—indeed, this derivation is not possible in the original PSL. Intuitively, we need to merge the two instances of (1) (for the two choices of b) into a single one by blending them with half probability each. To represent such “blending” at the level of assertions, one can define an operator $P \oplus_q Q$ which, roughly speaking, should hold on a distribution if it can be seen as the sum weighted by $q \in [0, 1]$ and $1 - q$ of two distributions which satisfy P and Q respectively. With this addition we can formulate a second principle of locality—*outcome locality*—in the form of the following rule:

$$\frac{\oplus\text{-LIFT} \quad \frac{\{P_1\} e \{Q_1\} \quad \{P_2\} e \{Q_2\}}{\{P_1 \oplus_q P_2\} e \{Q_1 \oplus_q Q_2\}}}{\{P_1 \oplus_q P_2\} e \{Q_1 \oplus_q Q_2\}}$$

In our example, one can establish that $\ell \rightsquigarrow \text{Ber}_{\frac{1}{2}} \equiv [\ell \mapsto 0] \oplus_{\frac{1}{2}} [\ell \mapsto 1]$ and (2) would then follow from $\oplus\text{-LIFT}$.

This kind of reasoning is an instance of a more general pattern: *reasoning under conditioning*. Often in probabilistic proofs it is necessary to perform case analysis on some random variable X (like the value at ℓ in the example). This formally corresponds to proving something about a distribution μ , by considering the family of *conditional distributions* $\mu|_{X=v}$ (i.e., the restriction of μ to the case where X is the deterministic value v), one for each possible value v .

The Lilac logic [20] proposed to represent conditioning using a *conditioning modality*, which can be used as a generalization of \oplus_q . In Amaryllis we directly adopt this idea, building on later variants of the modality defined for discrete distributions in [5, 32]. Intuitively, Amaryllis’s conditioning modality $C_{X \leftarrow \mu} P(x)$ holds if there is a random variable X distributed as μ such that for every $v \in \text{supp}(\mu)$, $P(v)$ holds on the distribution conditioned on $X = v$; for example, $[\ell \mapsto 0] \oplus_q [\ell \mapsto 1]$ can be expressed as $C_{b \leftarrow \text{Ber}_q} [\ell \mapsto b]$. The modality $C_{X \leftarrow \mu} P(x)$ leaves implicit (i.e., existentially quantified) which variable X is being conditioned; to make X apparent, it is sufficient to write $P(x)$ such that $\forall x. P(x) \vdash [X = x]$.

The growing body of work on probabilistic separation logics provides convincing evidence that these principles are key for making reasoning scale. Yet, existing logics either do not support mutable state (Lilac) or only support a finite variable store (PSL, BLUEBELL, pcOL). This is not by accident: as we now explain, dynamic heap allocation poses serious challenges for the techniques developed so far.

2.2 The Challenge of Dynamic Allocation

The core of the issue with previous models and dynamic allocation is the difficulty of harmonizing in a single logic *independence* and *disjointness* of state.

Lilac uses separation exclusively for *independence*, with the consequence of only supporting immutable state. This is because frames can record too much information, which needs to be preserved for **FRAME** to hold. For example $X \sim \text{Ber}_{\frac{1}{2}} * [X \in \{0, 1\}]$ is satisfiable, so owning $X \sim \text{Ber}_{\frac{1}{2}}$ alone is not enough to be able to mutate X .

In logics that support mutation, one needs to also embed in the model information about disjointness of locations. In all the logics with mutation (except for BLUEBELL which we address later), the independence and disjointness are conflated in the definition of separating conjunction: a distribution of states μ satisfies $P * Q$ just if there are two *disjoint* sets of locations L_1, L_2 such that for some μ_1 distribution of memories of L_1 satisfying P and some μ_2 distribution of memories of L_2 satisfying Q we have $\mu = \mu_1 \otimes \mu_2$ (here \otimes is the product of distributions).

This definition does not work for dynamic allocation. To illustrate why that is, we assume a deterministic allocator, to show that the core issue arises equally in this simpler setting. In the technical development, Amaryllis only assumes a non probabilistic but nondeterministic allocator, obtaining an even stronger result.

What makes dealing with allocation hard is the necessary dependency of the allocator on the current heap (as it needs, at the very least, to know which locations are available). Since the size and shape of the heap may depend on the previous random choices, the specific locations picked by the allocator will also depend on previous random choices.

To see the issue consider the following program:

$$dfl \triangleq \text{let } _ = (\text{if flip}(\frac{1}{2}) \text{ then ref } 0) \text{ in} \\ (\text{ref flip}(\frac{1}{2}), \text{ref flip}(\frac{1}{2})) \quad (3)$$

Depending on an initial coin flip, the heap may grow by one location before the two returned locations are allocated. Ideally, we would like to prove a postcondition that asserts that we end up with two disjoint heap locations storing the results of two independent coins:

$$\{\text{True}\} dfl \{(X, Y). X \rightsquigarrow \text{Ber}_{\frac{1}{2}} * Y \rightsquigarrow \text{Ber}_{\frac{1}{2}}\} \quad (4)$$

Unfortunately, this triple does not hold (semantically) in the model sketched above. The reason is that it is impossible to partition the locations into L_1, L_2 so that X only gets locations from L_1 and Y from L_2 , across all random outcomes. For instance, starting from an empty heap, if the first flip is 0 the two returned locations would be $(0 \times 0, 0 \times 1)$, but would be $(0 \times 1, 0 \times 2)$ if the first flip was 1. Therefore, although in any particular execution the two locations would be different, when considering the whole distribution at once, the possible locations associated with X and the ones with Y overlap in 0×1 .

BLUEBELL proposed a fractional-permission-based way of decoupling ownership of locations and independence; their solution, however, is specific to static variable stores, and does not obviously generalize to heaps with dynamic allocation. Moreover, the BLUEBELL logic only supports the $\oplus\text{-LIFT}$ rule with heavy side conditions which disable the use of frame around it, hindering modularity. (See Section 6 for details.)

2.3 A Model of Per-Outcome Disjointness

We solve the problem from the previous section by defining disjointness of heap resources *per random outcome*, rather than insisting on disjointness of heaps across all outcomes.

To realize this idea, we propose a new model of probabilistic separation logic which is parametric on a standard resource algebra. Specifically, given a standard resource algebra M (e.g., heaps) formalizing the desired notion of memory ownership (e.g., disjoint union of heaps), we show how to build a resource algebra representing *distributions* of elements of M so that can make assertions about distributions of random variables, probabilistic conditional independence, and ownership of memory.

Separation on the standard resources M is represented as a partial commutative binary operator \cdot . For the case of heaps, the resources are finite maps h from locations to values, and $h_1 \cdot h_2$ is defined as disjoint union $h_1 \uplus h_2$ if the two maps have disjoint domains, otherwise the composition is undefined. The crucial question is: what is the correct way of lifting the operation \cdot to an operation $\hat{\cdot}$ on *distributions* of resources $\mathbb{D}(M)$? A natural choice is to make $\mu_1 \hat{\cdot} \mu_2$ only defined if the composition $r_1 \cdot r_2$ is defined for every $r_1 \in \text{supp}(\mu_1)$ and $r_2 \in \text{supp}(\mu_2)$ (i.e., for every pair of resources that have non-zero probability). This choice gives rise to the model of PSL and pCOL.

Given our diagnosis of the problem with allocation, the issue with the above $\hat{\cdot}$ is that it considers pairs of resources that do not live in the same branch of the random choices made by the program. The domain $\mathbb{D}(M)$ does not however record the choices explicitly, making it impossible to fix the situation.

In Amaryllis, we therefore move to a different representation of distributions of resources, which is a variant of so-called *indexed valuations* [28], constructed as follows. We assume an arbitrary countable set of *random choice identifiers* $\iota \in \text{Rid}$. We then fix the outcome space to $\Omega = \text{Rid} \rightarrow \mathbb{B}$ of maps from random choice identifiers to booleans:² these abstractly represent the coin tosses performed by the program so far. At a high level, we define a probabilistic resource to be a pair (\mathcal{P}, R) where \mathcal{P} is a probability space over outcomes Ω , and $R: \Omega \rightarrow M$ is a *random resource variable* (which is *not* required to be measurable in \mathcal{P}). Thanks to the fact that we have an abstract representation of the outcomes of random choices as $\rho \in \Omega$, all our resources are now indexed in R by the specific random branch they belong to. This allows us to define $(\mathcal{P}_1, R_1) \hat{\cdot} (\mathcal{P}_2, R_2)$ as the pair $(\mathcal{P}_1 \otimes \mathcal{P}_2, (\lambda \rho. R_1(\rho) \cdot R_2(\rho)))$, that is: the probability spaces are composed using the independent product \otimes of Lilac, which encodes probabilistic independence, while the stateful resources (e.g., heaps) are composed *outcome-wise*.

In the Amaryllis model, we can give the desired meaning to the triple in (4). The assertion $X \rightsquigarrow \text{Ber}_{\frac{1}{2}}$ is given a meaningful semantics by letting $X: \Omega \rightarrow \text{Loc}$ be a *random location variable*. Given some $\mu \in \mathbb{D}(\text{Val})$, the assertion $X \rightsquigarrow \mu$ holds on (\mathcal{P}, R) if the function $\lambda \rho. R(\rho)(X(\rho))$ is measurable in \mathcal{P} and is distributed as μ in \mathcal{P} . In particular, $X(\rho)$ needs to be a location owned by $R(\rho)$, for every ρ not deemed impossible in \mathcal{P} . Then the triple in (4) expresses exactly that (i) outcome-wise, the two functions X and Y are different locations; and (ii) they point at values distributed as independent fair coins.

²Although Ω is uncountable, the probability spaces we consider only contain information about a finite set $R \subseteq \text{Rid}$. We make this precise in Section 3.

2.4 Probabilistic Specifications in Amaryllis

Now that we have an appropriate model of distributions of heaps, we want to build a program logic. Consider a possible axiom for allocation:

$$\text{HOARE-ALLOC-V} \quad \{\text{True}\} \text{ref } v \{L. L \mapsto v\}$$

The rule uses Amaryllis's "points-to" assertion, which is defined in general as follows. Given $L: \Omega \rightarrow \text{Loc}$ and $V: \Omega \rightarrow \text{Val}$, the assertion $L \mapsto V$ holds on a resource (\mathcal{P}, R) just if $L(\rho) \in \text{dom}(R(\rho))$ and $R(\rho)(L(\rho)) = V(\rho)$ for all $\rho \in \text{supp}(\mathcal{P})$. That is, the assertion $L \mapsto V$ implies per-outcome ownership of the *heap resource* $L(\rho) \mapsto V(\rho)$, but crucially it *does not* imply ownership of the *distribution* of either L or V . This is because (i) in the model, R is not required to be measurable in \mathcal{P} , and (ii) the assertion only constrains R , not \mathcal{P} .

Notice in particular how if $L \mapsto V$ were to own the distribution of L , it would be disastrous. As explained in Section 2.2, L is potentially dependent on (i.e., correlated with) every past random choice. Hence, if $L \mapsto V$ asserted something about the distribution of L , the postcondition of $\text{ref}(v)$ would not be independent of the program context, and the specification of $\text{ref}(v)$ would not be frame-preserving. By confining $L \mapsto V$ to describing heap and not distribution ownership, we ensure that the specification enjoys the frame rule. The assertion $L \rightsquigarrow \mu$ which we used in Section 2.1 can now be defined as shorthand for $\exists V. L \mapsto V * V \sim \mu$.

The **HOARE-ALLOC-V** triple is valid in Amaryllis, but how could we reuse it in a context where the initial value stored at the allocated location is not a deterministic value v but rather a *random* value variable V ? To allow for this, Amaryllis makes a natural generalization: our triples are in fact not over single expressions $e \in \text{Expr}$, but over random expression variables $E: \Omega \rightarrow \text{Expr}$. This allows us to directly prove:

$$\text{HOARE-ALLOC} \quad \{\text{True}\} \text{ref } V \{L. L \mapsto V\}$$

where $\text{ref } V$ stands for $(\lambda \rho. \text{ref } V(\rho))$. **HOARE-ALLOC** subsumes **HOARE-ALLOC-V** by choosing $V = \lambda_. v$, and precisely describes the effect of the command without requiring ownership of the distribution of V ; in the case where we do initially own $V \sim \mu$, by **FRAME** we can also get $\{V \sim \mu\} \text{ref } V \{L. L \mapsto V * V \sim \mu\}$.

Using our triples, we are able to provide a very general version of Iris's "bind" rule—a functional version of the Hoare Logic rule of sequential composition:

$$\text{HOARE-BIND} \quad \frac{\{P\} E \{W. P'(W)\} \quad \forall W. \{P'(W)\} K[W] \{V. Q(V)\}}{\{P\} K[E] \{V. Q(V)\}}$$

Here $K[\cdot]$ is an evaluation context, adapted to work on random expression variables. Using **HOARE-BIND** we can introduce random variables for intermediate results; for example, to handle $\text{ref}(\text{flip}(q))$ we can first bind on the flip and then apply the rules for allocation as sketched before:

$$\frac{\{\text{true}\} \text{flip}(q) \{V. V \sim \text{Ber}_q\} \quad \{V \sim \text{Ber}_q\} \text{ref } V \{L. L \mapsto V * V \sim \text{Ber}_q\}}{\{\text{true}\} \text{ref}(\text{flip}(q)) \{L. L \mapsto V * V \sim \text{Ber}_q\}}$$

2.5 Probabilistic Frame-Preserving Updates

We have already discussed the importance of the \oplus -LIFT rule for supporting outcome-local reasoning. Changing the model as we do in Amaryllis, however, presents a serious challenge: it is unclear how and why the \oplus -LIFT rule should hold. In fact, pcOL supports \oplus -LIFT by relying on the simplifying effect of conflating disjointness and independence; BLUEBELL only supports the rule by disabling further framing via a heavy side-condition.

In Amaryllis, \oplus -LIFT is in fact a special case of the C-LIFT rule:

$$\text{C-LIFT} \quad \frac{\forall x \in \text{supp}(\mu). \{P(x)\} E \{V. Q(x)(V)\}}{\{C_{x \leftarrow \mu}. P(x)\} E \{V. C_{x \leftarrow \mu}. Q(x)(V)\}}$$

We show that C-LIFT is valid in Amaryllis, by strengthening the notion of frame-preservation embedded in the meaning of triples to a new notion we call *probabilistic frame preservation*. We achieve the result in a modular fashion, following the Iris methodology, which prescribes triples and their rules to be derived from the combination of more primitive modalities and their rules: triples are described in terms of a *weakest precondition* (WP) modality $\text{wp } e \{ \Phi \}$ which, in turn, is built around the *update* modality $\text{update } P$. Below, we explain these components, and how we modified them in Amaryllis.

Intuitively, Iris's WP modality $\text{wp } e \{ \Phi \}$ holds on a resource a if every execution of e from a results in a return value v and resource b which satisfies the postcondition $\Phi(v)$. In addition, $\text{wp } e \{ \Phi \}$ insists that every possible frame of a is preserved by e . Using this modality, a triple $\{P\} e \{ \Phi \}$ can be encoded as $P \vdash \text{wp } e \{ \Phi \}$.

A crucial aspect of WP is that it transforms the current resource into a new resource that satisfies the postcondition, while preserving all frames (thus validating **FRAME**). This property is represented by the valid entailment $P * \text{wp } e \{ \Phi \} \vdash \text{wp } e \{ v. P * \Phi(v) \}$. In Iris, WP is itself a derived formula, which uses the so-called *update modality* $\text{update } P$ which, roughly speaking, holds on a resource a if P holds on some resource b such that every frame of a is also a valid frame of b (i.e., we can “update” a to b so that P holds without “breaking” any frame of a). The frame property of WP is thus a reflection of the frame property of updates $P * \text{update } P \vdash \text{update } (P * Q)$.

In Amaryllis, we similarly derive the definition of WP from an update modality. In order for triples to validate C-LIFT , the WP needs to satisfy the C-WP-SWAP rule:

$$\text{C-WP-SWAP} \quad C_{x \leftarrow \mu} \text{wp } E \{V. Q(x)(V)\} \vdash \text{wp } E \{V. C_{x \leftarrow \mu} Q(x)(V)\}$$

The rule says that WP and conditioning commute; consider for example the case where $\mu = \text{Ber}_q$. The rule states that, if E produces a distribution satisfying $Q(b)$ (for $b \in \{0, 1\}$) from the current distribution conditioned on $X = b$ for some (implicit) X , then E , run from the unconditional distribution, produces a distribution that satisfies $Q(0) \oplus_q Q(1)$.

In the same way the frame property of WP follows from the one on update, the proof of C-WP-SWAP boils down to proving a similar rule for update:

$$\text{C-UPD-SWAP} \quad C_{x \leftarrow \mu} \text{update } P(x) \vdash \text{update } C_{x \leftarrow \mu} P(x)$$

Unfortunately, if we only require the update to be frame-preserving, as is done in standard Iris, C-UPD-SWAP would not hold. The reason is that there are frame-preserving updates which do not commute with conditioning. In particular, conditioning on some random variable X (in some \mathcal{P}) is only possible if X is measurable in \mathcal{P} . Consequently, all updates performed under branches of this conditioning must at least preserve the measurability of X or else the conditioning ceases to be meaningful, and yet some frame-preserving updates do not!

Consider for example updating some \mathcal{P}_0 to \mathcal{P}_1 where X is measurable in \mathcal{P}_0 but not in \mathcal{P}_1 . “Forgetful” updates like this are always frame-preserving, because if every event that is measurable locally is independent of every event that the frame can measure, then if we make fewer events measurable locally, the frame would remain independent of the local resource. The update, however, *does not preserve conditioning*: conditioning on X in some \mathcal{P} is only possible if X is measurable in \mathcal{P} , and so all updates performed in branches of this conditioning must at least preserve the measurability of X or else the conditioning ceases to be meaningful. To see why, observe the case where we start from some \mathcal{P} where X is measurable, and distributed as a fair coin. If we condition on X we would have $\mathcal{P} = \frac{1}{2} \cdot \mathcal{P}|_{X=0} + \frac{1}{2} \cdot \mathcal{P}|_{X=1}$. But if in the first branch of the conditioning we update $\mathcal{P}|_{X=0}$ to \mathcal{P}'_0 by making X not measurable, then X would no longer be measurable in $\mathcal{P}' = \frac{1}{2} \cdot \mathcal{P}'_0 + \frac{1}{2} \cdot \mathcal{P}|_{X=1}$, causing the conditioning to become senseless.

This is already a problem without considering dynamic allocation: for instance, in pcOL triples only insist on frame preservation, so the logic was forced to impose side conditions that exclude updates like the one above in crucial rules for conditioning.

In Amaryllis, we propose a new definition of the update modality, which imposes stronger constraints on the allowed updates such that both the frame property of updates and C-UPD-SWAP are validated. The idea is to require updates to also preserve the validity of any possible conditioning they might be part of. We call such updates *probabilistic frame-preserving* (or PFP). For this new notion, we show that:

- (1) Any standard frame-preserving update on elements of M (the underlying standard resources) can be lifted to a PFP update on the random resource variable component of our model. As special cases, we get that mutation and dynamic allocation of heap cells are PFP.
- (2) Taking a new sample from a distribution is a PFP update to the probability space component of the model.

2.6 Adapting Iris's “Authoritative” Resource Algebra to the Probabilistic Setting

To fully derive WP within the logic on top of update, standard Iris makes essential use of a construction called the “authoritative resource algebra” $\text{Auth}(M)$. In essence, $\text{Auth}(M)$ consists of two kinds of elements, built on top of the elements of an underlying resource algebra M (e.g., heaps): (1) the *authority* $\bullet g$, which gives exclusive ownership of a resource $g \in M$ representing the “global” view of the state, and (2) the *fragment* $\circ a$, which gives ownership of a resource $a \in M$ representing local ownership of a piece of the state. Fragments compose: $\circ a_1 \cdot \circ a_2 = \circ(a_1 \cdot_M a_2)$. Exclusivity of the authority is represented by the fact that $\bullet g_1 \cdot \bullet g_2$ is undefined,

and thus in any valid separation of resources there is at most one authority. Crucially, $\bullet g \cdot \circ a$ is only valid if a is a “sub-resource” of g (a concept specified in the definition of M).

In Amaryllis we build our model with an authoritative construction on top of the resource algebra we sketched in Section 2.3. Intuitively, fragments $\circ(\mathcal{P}, R)$ support the meaning of the user-facing assertions like $L \mapsto V$ which indicate local resources. The authority³ $\bullet(\mathcal{P}, R)$ is used in the encoding of WP to constrain the local fragments to be sub-resources of the global resource (\mathcal{P}, R) .

The model we defined using the authoritative construction validates all the desired rules, except for one (first proposed in [5]) that represents a key property of conditioning:

$$\text{C-FRAME-BB} \\ P * C_{v \leftarrow \mu} Q(v) \vdash C_{v \leftarrow \mu} (P * Q(v))$$

The rule is involved in virtually every conditional proof. For example, to condition on X in $Y \sim \mu' * X \sim \mu$, one first applies $X \sim \mu \vdash C_{x \leftarrow \mu} [X = x]$ to obtain $Y \sim \mu' * C_{x \leftarrow \mu} [X = x]$, and then **C-FRAME-BB** to push the independent resource $Y \sim \mu'$ under the conditioning modality, obtaining $C_{x \leftarrow \mu} ([X = x] * Y \sim \mu')$. This makes Y available to each branch of the conditioning when we subsequently apply the rule **C-LIFT**.

In Amaryllis there is a counterexample to the validity of **C-FRAME-BB**: suppose we start from $\text{Own}(\bullet(\mathcal{P}, R)) * X \sim \text{Ber}_{1/2}$; then we know that the global distribution \mathcal{P} will see X having a fair coin distribution. If we condition on X , obtaining $\text{Own}(\bullet(\mathcal{P}, R)) * C_{x \leftarrow \text{Ber}_{1/2}} [X = x]$, an application of **C-FRAME-BB** would bring us to $C_{x \leftarrow \text{Ber}_{1/2}} \text{Own}(\bullet(\mathcal{P}, R)) * [X = x]$, which is inconsistent, since in every branch of the conditioning we would be simultaneously claiming that X is deterministic (by $[X = x]$) and distributed as a fair coin (by the information in \mathcal{P}).

In Amaryllis, we resolve this issue by reinterpreting the meaning of the authority: it no longer represents the global resource in absolute terms, but only relative to the current conditional distribution. Specifically, we define the authoritative algebra so that conditioning validates two new variations of **C-FRAME-BB**:

$$\text{C-FRAME} \\ \boxtimes P * C_{v \leftarrow \mu} Q(v) \vdash C_{v \leftarrow \mu} (\boxtimes P * Q(v))$$

$$\text{C-AUTH} \\ \frac{\forall v. Q(v) \vdash [X = v]}{\text{Own}(\bullet g) * C_{v \leftarrow \mu} Q(v) \vdash C_{v \leftarrow \mu} (\text{Own}(\bullet(g|_{X=v})) * Q(v))}$$

The rule **C-FRAME** uses a new modality $\boxtimes P$ to guard the resource to be brought inside a conditioning: it is the identity on fragments ($\text{Own}(\circ a) \dashv\vdash \boxtimes \text{Own}(\circ a)$), but cannot wrap the authority ($\boxtimes \text{Own}(\bullet g) \vdash \text{False}$). The rule **C-AUTH** allows for the sound transfer of global information across a conditioning, by performing the relevant conditioning on the global state. The premise of the rule serves to identify the random variable X being conditioned (which is needed in order to condition the global resource g on X). The two rules together cover all the user-facing uses of **C-FRAME-BB** (which

³In the technical development an authoritative construction is applied both to the probabilistic resources and the underlying standard resources. The points we make in this section apply equally to the variant of Section 3.

would only involve fragments), as well as the new manipulations on the authority required when proving properties of WP.

Equipped with this machinery, Amaryllis supports the encoding of a rather standard-looking definition of WP within the base probabilistic separation logic, which can be manipulated abstractly using our new rules of conditioning and update.

3 The Amaryllis Model

This section describes the model of Amaryllis. We start by recapping prior concepts (Preliminaries) and then describe the specifics of our new model.

3.1 Preliminaries

A set of subsets of Ω is a σ -algebra if it contains Ω and is closed under countable unions and complements. The set of all σ -algebras over Ω is denoted $\mathbb{A}(\Omega)$. Given $\mathcal{F} \in \mathbb{A}(\Omega)$, an element of \mathcal{F} is called *event* and Ω is called *outcome space*. A distribution $\mu \in \mathbb{D}(\mathcal{F})$ over some $\mathcal{F} \in \mathbb{A}(\Omega)$ is an additive function $\mu : \mathcal{F} \rightarrow [0, 1]$ with $\mu(\Omega) = 1$. The *support* of $\mu \in \mathbb{D}(\mathcal{F})$ is defined as $\text{supp } \mu = \Omega \setminus \{\rho \in \Omega \mid \exists X. \rho \in X \wedge \mu(X) = 0\}$. The *Dirac distribution* at $v \in \Omega$ is defined as $\delta_v(v) = 1$ and $\delta_v(w) = 0$ for all $w \neq v$.

A *probability space* \mathcal{P} is a tuple $\mathcal{P} = (\Omega, \mathcal{F}, \mu)$ where $\mathcal{F} \in \mathbb{A}(\Omega)$ and μ is a distribution over \mathcal{F} . When Ω is clear from the context, we often omit it. The probability space $\mathcal{P} = (\mathcal{F}, \mu)$ is *complete* if for all events $X \in \mathcal{F}$ with $\mu(X) = 0$ and all $\rho \in X$, it holds that $\{\rho\} \in \mathcal{F}$. Given $\mathcal{P}_1 = (\mathcal{F}_1, \mu_1), \mathcal{P}_2 = (\mathcal{F}_2, \mu_2)$ with $\mathcal{F}_1 \in \mathbb{A}(\Omega_1)$ and $\mathcal{F}_2 \in \mathbb{A}(\Omega_2)$, their product $\mathcal{P}_1 \otimes \mathcal{P}_2$ is the probability space over $\Omega_1 \times \Omega_2$ with σ -algebra $\sigma(\{X_1 \times X_2 \mid X_1 \in \mathcal{F}_1, X_2 \in \mathcal{F}_2\})$, and distribution $(\mu_1 \otimes \mu_2)(X_1 \times X_2) = \mu_1(X_1) \cdot \mu_2(X_2)$.

The *independent product* [20] $\mathcal{P}_1 \circledast \mathcal{P}_2$ of two probability spaces $\mathcal{P}_1 = (\mathcal{F}_1, \mu_1)$ and $\mathcal{P}_2 = (\mathcal{F}_2, \mu_2)$ over the same outcome space is defined as $(\sigma(\mathcal{F}_1 \cup \mathcal{F}_2), \mu)$, where $\mu(X_1 \cap X_2) = \mu_1(X_1) \cdot \mu_2(X_2)$ for any $X_1 \in \mathcal{F}_1$ and $X_2 \in \mathcal{F}_2$, and $\sigma(\mathcal{F}_1 \cup \mathcal{F}_2)$ is the smallest σ -algebra containing all elements of $\mathcal{F}_1 \cup \mathcal{F}_2$. The independent product is unique if it exists [20]. For a probability space $\mathcal{P} = (\mathcal{F}, \mu)$ and a set $X \in \mathcal{F}$, the *conditioning* $\mathcal{P}|_X$ is defined as (\mathcal{F}, μ') , where $\mu'(Y) = \mu(Y \cap X) / \mu(X)$ for all $Y \in \mathcal{F}$.

If $\mathcal{F} \in \mathbb{A}(\Omega_1)$ and $F : \Omega_1 \rightarrow \Omega_2$ then F is *measurable* if $\forall X \subseteq \Omega_2. F^{-1}(X) \in \mathcal{F}$. Such a measurable function is also called a *random variable*. A random variable F is *distributed* as distribution μ if $\forall i \in \text{supp}(\mu). \mathcal{P}(F^{-1}(i)) = \mu(i)$.

For with $\mathcal{F}_1 \in \mathbb{A}(\Omega)$, a restriction of $\mu \in \mathbb{D}(\mathcal{F}_1)$ to $\mathcal{F}_2 \in \mathbb{A}(\Omega)$ ($\mathcal{F}_2 \subseteq \mathcal{F}_1$) is the distribution $\mu|_{\mathcal{F}_2}$ with $\mu|_{\mathcal{F}_2}(X) = \mu(X)$ for all events in \mathcal{F}_2 . The *extension order* relates $(\mathcal{F}_1, \mu_1) \sqsubseteq (\mathcal{F}_2, \mu_2)$ iff $\mathcal{F}_1 \subseteq \mathcal{F}_2$ and $\mu_1 = \mu_2|_{\mathcal{F}_1}$.

An *ordered resource algebra* (ORA) is a tuple $(M, \preceq, \mathcal{V}, \cdot, _, \varepsilon)$ where M is a set called the *carrier type*, the *resource order* \preceq is a pre-order over M , $\mathcal{V} : M \rightarrow \text{Prop}$ is the *validity predicate*, the *resource composition* \cdot is a partial commutative and associative binary operation on M , the *core* $_$ is a unary operation on M and $\varepsilon \in M$ is the *unit*. The unit should satisfy $\mathcal{V}(\varepsilon)$ and $\varepsilon \cdot a = a$ for any $a \in M$. Further, for all $a, b, c \in M$, $\mathcal{V}(a \cdot b) \Rightarrow \mathcal{V}(a), a \preceq b \Rightarrow \mathcal{V}(b) \Rightarrow \mathcal{V}(a)$ and $a \preceq b \Rightarrow a \cdot c \preceq b \cdot c$ must hold.

3.2 Our Resource Algebra

The core of our model is an ORA defined on probability spaces over the outcome space $\Omega = \text{Rid} \rightarrow \mathbb{B}$. The probability spaces \mathcal{P} we work with can be factorized as $\mathcal{P} = \mathcal{P}_R \otimes \mathbb{1}_{(\text{Rid} \setminus R) \rightarrow \mathbb{B}}$ for a finite $R \subseteq \text{Rid}$, where $\mathbb{1}_\Omega \in \mathbb{P}(\Omega)$ is the σ -algebra $\{\Omega, \emptyset\}$ equipped with the trivial distribution. In our Rocq implementation, we use a concrete representation of such spaces that keeps track of R explicitly. For simplicity, we omit the details of this tracking here and present our definitions directly on probability spaces over $\text{Rid} \rightarrow \mathbb{B}$. The details of our representation can be found in Appendix B.

3.2.1 Our ORA on Probability Spaces. Our ORA's definition is shown below. Crucially, the ORA's resource composition is the independent product of probability spaces, and its persistent core is the *sure* subspace (probability 0 or 1) of the probability space.

$$\begin{aligned} \text{PSp} &\triangleq \{\mathcal{P} \mid \mathcal{P} \text{ is a complete PSp. over } \text{Rid} \rightarrow \mathbb{B}\} \cup \{\perp\} \\ \mathcal{V}(a) &\triangleq a \neq \perp \\ \mathcal{P}_1 \cdot \mathcal{P}_2 &\triangleq \mathcal{P}_1 \otimes \mathcal{P}_2 \\ |(\mathcal{F}, \mu)| &\triangleq (\{X \in \mathcal{F} \mid \mu(X) = 0 \vee \mu(X) = 1\}, \mu) \\ \varepsilon &\triangleq \mathbb{1}_{\text{Rid} \rightarrow \mathbb{B}} \\ \mathcal{P} \preccurlyeq \mathcal{P}' &\triangleq \mathcal{P} \subseteq \mathcal{P}' \quad (\text{in particular, } \perp \not\preccurlyeq \mathcal{P} \text{ if } \mathcal{P} \neq \perp) \end{aligned}$$

3.2.2 Random Resource Variables and Authoritative Component. Next, we present the ORA PSpAuth_M , a slightly simplified version of our full model. The actual full model also contains a *namespace* component that tracks a set of globally free randomness IDs: this is used to ease the proof that allocation of new RIDs is a probabilistic frame-preserving update, but omitted here for simplicity. The details are in Appendix B.

Given an Iris resource algebra M , an element of PSpAuth_M is a triple of a probability space, an M -valued random variable (called a *random resource variable*) and an authoritative component. For the random resource variable component, all resource algebra connectives of M are lifted pointwise. The authoritative component behaves like Iris' authoritative RAs: It is an upper bound on the probability space (in the extension order), and it is *exclusive* – two different authorities cannot be composed. Following Iris terminology, the probability space component is also called a *fragment* of the authoritative part.

For clarity, we highlight expressions related to the random resource variable component and the authoritative component in **orange** and **blue**, respectively.

$$\begin{aligned} \text{PSpAuth}_M &\triangleq \text{PSp} \times \text{RV}_M \times \text{Option}(\text{PSp}) \\ (\mathcal{P}, R, a) \cdot (\mathcal{P}', R', a') &\triangleq (\mathcal{P} \cdot \mathcal{P}', \lambda\rho. R(\rho) \cdot R'(\rho), a \circ a') \\ \mathcal{V}(\mathcal{P}, R, a) &\triangleq \mathcal{V}(\mathcal{P}) \wedge \forall\rho. \mathcal{V}(R(\rho)) \wedge \\ &\quad (a = \mathcal{P}_a \Rightarrow \mathcal{V}(\mathcal{P}_a) \wedge \mathcal{P} \preccurlyeq \mathcal{P}_a) \\ |(\mathcal{P}, R, a)| &\triangleq (|\mathcal{P}|, \lambda\rho. |R(\rho)|, \perp) \\ \varepsilon &\triangleq (\varepsilon, \lambda. \varepsilon, \perp) \\ (\mathcal{P}, R, a) \preccurlyeq (\mathcal{P}', R', a') &\triangleq \mathcal{P} \preccurlyeq \mathcal{P}' \wedge \forall\rho. R(\rho) \preccurlyeq R'(\rho) \wedge \\ &\quad (a \neq \perp \Rightarrow a = a') \end{aligned}$$

where $a \circ \perp = a = \perp \circ a$, $a \circ b = \perp$, and $\text{RV}_M := (\text{Rid} \rightarrow \mathbb{B}) \rightarrow M$.

For projecting out the three components of the resource algebra, we define $\pi_{\text{frag}}(\mathcal{P}, R, a) = \mathcal{P}$, $\pi_{\text{RV}}(\mathcal{P}, R, a) = R$, $\pi_{\text{auth}}(\mathcal{P}, R, a) = a$.

3.3 Weighted Sum and Conditioning

To define the conditioning modality, we find it useful to first define a *weighted sum* operation $wsum : \forall A. \mathbb{D}(A) \rightarrow (A \rightarrow \text{PSpAuth}_M) \rightarrow \text{PSpAuth}_M$ that combines the resources of exclusive probabilistic branches into a single resource. We first define the weighted sum on probability spaces, and then extend it to the full ORA.

3.3.1 Weighed Sum on Probability Spaces. We define the weighted sum on probability spaces as follows.

$$wsum_{\text{PSp}}(\mu, \kappa) \triangleq \begin{cases} \text{bind}\left(\mu, \lambda v. \kappa(v) \upharpoonright_{v \in \text{supp } \mu} \pi_{\mathcal{F}}(\kappa(v))\right) & \text{if } \text{has_rv}(\mu, \kappa) \\ \perp & \text{otherwise} \end{cases}$$

where

$$\text{has_rv}(\mu, \kappa) := \forall v, w \in \text{supp } \mu. \text{supp}(\kappa(v)) \cap \text{supp}(\kappa(w)) = \emptyset$$

The weighted sum is defined as the standard dependent product of μ and κ (the **bind** operation of the Giry monad) but only if the supports of κ 's branches are disjoint, as codified in the condition $\text{has_rv}(\mu, \kappa)$. This is equivalent to requiring that μ be the distribution of *some* random variable, and that κ be an exclusive union of branches defined by the possible values of that random variable. If this is the case, we define the σ -algebra of the weighted sum to be the set of events that are measurable in all branches.

3.3.2 Weighted Sum on our ORA. We extend the definition of weighted sum to our ORA. To take the weighted sum of random resource variables, consider a specific outcome ρ . Since the supports of κ 's branches must be pairwise disjoint (see above), there is at most one branch v such that ρ has nonzero probability in that branch. If such a v exists, define $R(\rho) = \pi_{\text{RV}}(\kappa(v)(\rho))$. If ρ is not possible in any branch, we define the weighted sum to be ε at ρ .⁴

Formally, define $\text{get_rv}(\mu, \kappa)(\rho)$ as the unique $v \in \text{supp}(\mu)$ such that $\rho \in \text{supp } \pi_{\text{frag}}(\kappa(v))$ if such a v exists. Then,

$$\begin{aligned} wsum(\mu, \kappa) &\triangleq (\mathcal{P}, R, a) \\ \mathcal{P} &\triangleq wsum_{\text{PSp}}(\mu, \lambda v. \pi_{\text{frag}}(\kappa v)) \\ R &\triangleq \lambda\rho. \text{if } \rho \in \text{supp}(\mathcal{P}) \text{ then } (\pi_{\text{RV}}(\kappa(\text{get_rv}(\mu, \kappa)))) \rho \text{ else } \varepsilon \end{aligned}$$

Note that the random resource variable of any branch is almost surely equal to R conditional on being in that branch.

We require that the authoritative part either exist in all branches of κ or none of them. If it exists in all of them, we define the resulting authority as the $wsum_{\text{PSp}}$ of the individual authoritative parts.

$$\begin{aligned} a &\triangleq \perp && \text{if } \forall v. \pi_{\text{auth}}(\kappa(v)) = \perp \\ a &\triangleq wsum_{\text{PSp}}(\mu, \kappa') && \text{if } \forall v. \pi_{\text{auth}}(\kappa(v)) = \kappa'(v) \end{aligned}$$

3.3.3 Conditioning. Next, we define conditioning.

$$(\mathcal{C}_{v \leftarrow \mu} \Phi(v))(x) \triangleq \exists \kappa. wsum(\mu, \kappa) \preccurlyeq x \wedge \forall v \in \text{supp } \mu. \Phi(v)(\kappa(v))$$

This definition is equivalent to the following axiomatic characterization of conditioning (here, $\perp|_E \triangleq \perp$ for all E):

$$\begin{aligned} (\mathcal{C}_{v \leftarrow \mu} \Phi(v))(\mathcal{P}, R, a) &\Leftrightarrow \exists X. X \text{ is distributed as } \mu \text{ in } \mathcal{P} \wedge \\ &\quad \forall v \in \text{supp } \mu. \Phi(v)(\mathcal{P}|_{X=v}, a|_{X=v}, R) \end{aligned}$$

⁴Our logic ignores probability 0 branches, so it might seem that ε could be replaced by an arbitrary element of our ORA. However, for technical reasons, we must ensure that those branches are valid. ε is always valid.

3.4 Probabilistic Frame-Preserving Updates

Next, we define PFP updates, which generalize Iris updates. Like Iris updates, PFP updates are upwards closed with respect to the resource order and preserve frames. Additionally, we require PFP updates to *preserve conditioning*, i.e., the validity of weighted sums. Formally, a PFP update from resources x to x' preserves conditioning if

$$\forall \mu, \kappa, v. \mathcal{V}(wsum(\mu, \kappa[v \mapsto x])) \Rightarrow \mathcal{V}(wsum(\mu, \kappa[v \mapsto x']))$$

Next, we consider upwards-closedness of updates. Our weighted sum is not monotone: $(\forall v. \kappa(v) \preceq \kappa'(v)) \Rightarrow wsum(\mu, \kappa) \preceq wsum(\mu, \kappa')$, since κ might not be induced by a random variable even if κ' is, whence $wsum(\mu, \kappa)$ would be $\frac{1}{2}$. As a result, upwards-closedness of updates needs to be encoded explicitly. To do that, we define an operator \circ that characterizes the resource order: $x \preceq y \Leftrightarrow \exists z. x \circ y \equiv z$. The concrete definition of \circ is omitted here and can be found in Appendix B and in our Rocq development [3].

We can now define *probabilistic contexts*, K_M , comprising the operators \circ , $wsum$ and the resource composition (\cdot) , and define context application, $K_M[[x]]$.

$$K_M \triangleq \cdot \mid \text{frame}(cx, K_M) \mid \text{cond}(\mu, \kappa, v, K_M) \mid \text{ext}(cx, K_M)$$

$$\cdot[[x]] \triangleq x$$

$$\text{frame}(cx, K_M)[[x]] \triangleq cx \cdot K_M[[x]]$$

$$\text{cond}(\mu, \kappa, v, K_M)[[x]] \triangleq wsum(\mu, \kappa[v \mapsto K_M[[x]])$$

$$\text{ext}(cx, K_M)[[x]] \triangleq cx \circ K_M[[x]]$$

Equipped with these definitions, we define *probabilistic frame-preserving (PFP) updates*, as those resource updates that preserve validity in all contexts:

$$\models P \triangleq \lambda a. \forall K_M. \mathcal{V}(K_M[[a]]) \Rightarrow \exists a'. \mathcal{V}(K_M[[a']]) \wedge P(a')$$

We have shown that allocating a randomness ID is a probabilistic-frame preserving update, and that frame-preserving updates on the random resource variable (including non-deterministic updates) can be lifted to PFP updates:

THEOREM 3.1 (LIFTING FRAME-PRESERVING UPDATES).

$$(\forall \rho \text{ cx}. \mathcal{V}(R(\rho) \cdot \text{cx}) \Rightarrow \exists x'. \mathcal{V}(x' \cdot \text{cx}) \wedge \Phi(\rho)(x')) \Rightarrow$$

$$\forall K_M. \mathcal{V}(K_M[[\varepsilon, R, \perp]]) \Rightarrow \exists R'. \mathcal{V}(K_M[[\varepsilon, R', \perp]]) \wedge$$

$$\forall \rho. \Phi(\rho)(R'(\rho))$$

This theorem says that if a random resource variable R can be updated in the Iris frame-preserving sense to a resource that satisfies the predicate Φ in all probabilistic branches, then the probabilistic resource (ε, R, \perp) can be PFP updated (in the sense defined above) to satisfy Φ in all branches. Intuitively, the theorem holds in our model because all operations on random resource variables are defined per branch (per outcome). A weighted sum has no effect on non-zero probability branches of the random resource variable, and ignores zero probability branches. Thus, an update in a branch of the random resource variable, whether probability 0 or not, will always preserve the validity of arbitrary weighted sums. See our Rocq development [3] for the formal proof.

4 The Amaryllis Logic

4.1 Propositions and Connectives

Our logic's propositions pProp are predicates over PSPAuth_M that are upwards-closed with respect to the resource order and, additionally, proper with respect to almost-sure equality of the random resource variable. To define them formally, we first define a slight variation of our resource order, the *almost-sure resource order* that ignores probability 0 branches:

$$(\mathcal{P}, R, a) \preceq_{as} (\mathcal{P}', R, a') \triangleq \mathcal{P} \preceq \mathcal{P}' \wedge$$

$$(a \neq \perp \Rightarrow a \equiv a') \wedge \forall \rho \in \text{supp}(\mathcal{P}) \Rightarrow R(\rho) \preceq R'(\rho)$$

$$\text{pProp} := \{\Phi : \text{PSPAuth}_M \rightarrow \text{Prop} \mid \forall x, x'. x \preceq_{as} x' \wedge \Phi(x') \Rightarrow \Phi(x)\}$$

The conditioning $(C_{v \leftarrow \mu} \Phi(v))(x)$ and update $(\models P)$ modalities were already defined in Section 3. $\text{Own}(x)$, \boxtimes , $X \sim \mu$ and sure propositions $(\lceil \Phi \rceil)$ can be defined as follows:

$$\text{Own}(x')(x) \triangleq x' \preceq_{as} x$$

$$(\boxtimes \Phi)(\mathcal{P}, R, a) \triangleq \Phi(\mathcal{P}, R, \varepsilon)$$

$$(X \sim \mu)(\mathcal{P}, R, a) \triangleq X \text{ is distributed as } \mu \text{ in } \mathcal{P}$$

$$\lceil \Phi \rceil \triangleq \Phi \sim \delta_{\text{true}}$$

The remaining primitive connectives $*$, \ast , \forall , \exists , \wedge , \vee , \Rightarrow , pure propositions and the Iris's persistency modality \square for duplicable assertions can all be defined in the standard way. As examples, we show the definitions of $*$, \ast and \square .

$$(P * Q)(x) \triangleq \exists x_1, x_2. x_1 \cdot x_2 \preceq x \wedge P(x_1) \wedge Q(x_2)$$

$$(P \ast Q)(x) \triangleq \forall x'. \mathcal{V}(x \cdot x') \wedge P(x') \Rightarrow Q(x \cdot x')$$

$$(\square P)(x) \triangleq P(|x|)$$

It can be shown that all connectives preserve the almost-sure upwards-closedness of pProp .

In the following, we instantiate M with $\text{Auth}(gmap(\text{Loc} \rightarrow \text{Val}))$, the resource algebra that is used in standard Iris to model (deterministic) points-to assertions. Our pointsto-assertion can then be defined as follows.

$$L \mapsto V \triangleq \text{Own}(\varepsilon, (\lambda \rho. \circ[L(\rho) \mapsto V(\rho)], \perp))$$

Selected primitive proof rules are shown in Figure 1. The rules are straightforward to read, thanks to the setup of our model.

4.2 Language and Semantics

We consider the language pHeapLang , a variant of Iris's HeapLang language, extended with real numbers r and a sampling primitive $\text{flip}(e)$ that samples a boolean value from the Bernoulli distribution with bias e .

$$\begin{aligned} \text{Val } \exists v ::= z \in \mathbb{Z} \mid r \in \mathbb{R} \mid b \in \mathbb{B} \mid \ell \in \text{Loc} \mid \text{rec } f x = e \mid (v_1, v_2) \\ \text{Expr } \exists e ::= v \mid x \mid e_1 e_2 \mid \text{ref } e \mid \text{get } e \mid \text{set } e e \mid e_1 + e_2 \mid e_1 \leq e_2 \mid \dots \\ \mid (e_1, e_2) \mid \text{fst } e \mid \text{snd } e \mid \text{if } e \text{ then } e \text{ else } e \mid \text{flip}(e) \end{aligned}$$

We use notation $\lambda x. e \triangleq \text{rec } _x = e$, and $\text{let } x = e_1 \text{ in } e_2 \triangleq (\lambda x. e_2)e_1$.

pHeapLang comes equipped with a small step semantics:

$$\rightsquigarrow : (\text{Expr} \times \text{State}) \times (\text{Expr} \times \text{State} \times \text{Expr} \times \text{State} \times \mathbb{R})$$

929	DISTR-SUBST $X \sim \mu * [X = Y] \vdash Y \sim \mu$	DISTR-DIRAC $\vdash v \sim \delta(v)$	OWN-SUBST $\text{Own}(\varepsilon, X) * [X = Y] \vdash \text{Own}(\varepsilon, Y)$	SURE-PERSISTENT $[X = Y] \vdash \square [X = Y]$	987
930					988
931					989
932	C-SKOLEM $C_{v \leftarrow \mu} \exists a. \Phi(a, v) \dashv\vdash \exists f. C_{v \leftarrow \mu} \Phi(f(v), v)$	C-BIND $C_{v \leftarrow \mu} X \sim \kappa(v) \vdash X \sim \text{bind}(\mu, \kappa)$	C-FUSE $C_{v \leftarrow \mu} C_{w \leftarrow \kappa(v)} \Phi(v, w) \dashv\vdash C_{(v, w) \leftarrow \mu \leftarrow \kappa} \Phi(v, w)$		990
933					991
934					992
935					993
936	C-RV $C_{v \leftarrow \mu} \Phi(v) \vdash \exists X. C_{v \leftarrow \mu} [X = v] * \Phi(v)$	C-INTRO $X \sim \mu \dashv\vdash C_{v \leftarrow \mu} [X = v]$	C-DIRAC $C_{v \leftarrow \delta(i)} \Phi(v) \dashv\vdash \Phi(i)$	C-CONS $\frac{\forall v \in \text{supp}(\mu). \Phi_1(v) \vdash \Phi_2(v)}{C_{v \leftarrow \mu} \Phi_1(v) \vdash C_{v \leftarrow \mu} \Phi_2(v)}$	994
937					995
938					996
939					997
940	POINTSTO-STR-CONVEX $\text{StrConvex}(L \mapsto V)$	DISTR-STR-CONVEX $\text{StrConvex}(X \sim \mu)$	DISTR-FRAMEABLE $X \sim \mu \vdash \boxtimes X \sim \mu$	PURE-FRAMEABLE $\Phi \text{ pure} \Rightarrow \Phi \vdash \boxtimes \Phi$	998
941				SEP-FRAMABLE $\boxtimes P * \boxtimes Q \dashv\vdash \boxtimes (P * Q)$	999
942		$\mu \prec \kappa \triangleq \lambda(v, w). \mu(v)\kappa(v, w)$		$\text{StrConvex}(\Psi) \triangleq C_{v \leftarrow \mu} \Phi(v) * \Psi \vdash \Psi * C_{v \leftarrow \mu} \Phi(v)$	1000
943					1001

Figure 1: Selected proof rules of Amaryllis

where $e, \sigma \rightsquigarrow e_1, \sigma_1, e_2, \sigma_2, p$ can be read as: e, σ steps to e_1, σ_1 with probability p , and to e_2, σ_2 with probability $1 - p$. For instance, we have that $\text{flip}(p), \sigma \rightsquigarrow \text{true}, \sigma, \text{false}, \sigma, p$. All standard non-probabilistic steps $e, \sigma \rightsquigarrow_{np} e', \sigma'$ are lifted to probabilistic steps $e, \sigma \rightsquigarrow e', \sigma', e_2, \sigma_2, 1$ for all e_2, σ_2 . For example, $\text{ref } v, \sigma \rightsquigarrow (), \sigma[l \mapsto v], e_2, \sigma_2, 1$ for all $l \notin \text{dom } \sigma, e_2, \sigma_2$. We use this representation for non-probabilistic steps to avoid having to represent them differently from probabilistic steps.

We lift this semantics to a collecting big-step semantics on indexed valuations:

$$\Downarrow: (\text{PSpD} \times \text{RV}_{\text{Expr}} \times \text{RV}_{\text{State}}) \times (\text{PSpD} \times \text{RV}_{\text{Expr}} \times \text{RV}_{\text{State}})$$

Intuitively, for each step \rightsquigarrow that is executed in some branch, a new randomness ID is allocated and the random variables are updated accordingly. For example, consider the program `let x = flip(1/2) in ref (flip(1/2) + x)`, executed in an initially empty state and probability space. The first flip will create an RID ι_1 and leave us with the random expression variable $\lambda\rho$. `ref (flip(1/2) + \rho(\iota_1))`. The second flip creates two new randomness IDs, ι_2 and ι_3 , one for each of the two outcomes of the first flip, leaving us with the random expression variable $\lambda\rho$. if $\rho(\iota_1)$ then $\rho(\iota_2)$ else $\rho(\iota_3) + 1$. The last step, which is the allocation of a reference, will result in the random expression variable λ . $()$ and the random *state* variable $\lambda\rho$. if $\rho(\iota_1)$ then $\rho(\iota_2)$ else $\rho(\iota_3) + 1$. Note that this is not the only possible reduction as the RIDs and the RV values for probability 0 branches are chosen nondeterministically.

The formal definition of our language's operational semantics is in Appendix C.

4.3 The Weakest Precondition Modality

We define the *weakest precondition modality* $\text{wp} : \text{RV}_{\text{Expr}} \rightarrow (\text{RV}_{\text{Val}} \rightarrow \text{pProp})$ using the primitive connectives.

$$\begin{aligned} \text{wp } E \{v. \Phi(v)\} &\triangleq \forall \mathcal{P}, \Sigma. E, \Sigma \text{ measurable in } \mathcal{P} * S(\mathcal{P}, \Sigma) * \\ &\forall V, \mathcal{P}', \Sigma'. \langle \mathcal{P}, E, \Sigma \rangle \Downarrow \langle \mathcal{P}', V, \Sigma' \rangle * \models S(\mathcal{P}', \Sigma') * \Phi(V) \end{aligned}$$

where the *state interpretation* S is defined as

$$S(\mathcal{P}, \Sigma) := \text{Own}(|\mathcal{P}|, (\lambda\rho. \bullet\Sigma(\rho)), \mathcal{P})$$

and the postcondition Φ is required to be proper with respect to almost-sure equality: $\lceil V = V' \rceil * \Phi(V) \vdash \Phi(V')$.

In the weakest precondition, we first quantify over all possible current 'probabilistic states': A pair (\mathcal{P}, Σ) is a *possible probabilistic state* if (1) both E and Σ are measurable in \mathcal{P} , and (2) \mathcal{P} and Σ are both consistent with the 'local information', i.e. is is not contradictory to assume that \mathcal{P} and Σ represent the ground truth. Statement (2) is equivalent to assuming the state interpretation.

Next, we quantify over the all random value variables and probabilistic states that the current random expression variable might evaluate to. For each option, we require that we can update our local information such that (1) our local information is again consistent with the new probabilistic state \mathcal{P}', Σ' , and (2) the postcondition holds on the output random value variable.

Hoare triples are defined on top of the weakest precondition modality: $\{P\} E \{Q\} \triangleq P \vdash \text{wp } E \{Q\}$.

Proof rules for the weakest precondition can be found in Figure 2. The rules should be intuitive, and those for non-probabilistic constructs are analogous to standard proof rules in Iris, except that they now apply to random expression variables rather than deterministic expressions.

The following adequacy theorem, our main result, states that our WP definition is sound.

THEOREM 4.1 (ADEQUACY). *If $\langle \varepsilon, (\lambda_{-}. e), (\lambda_{-}. \sigma) \rangle \Downarrow \langle \mathcal{P}, V, \Sigma \rangle$ and $\vdash \text{wp } (\lambda_{-}. e) \{V. V \sim \mu\}$, then V is distributed as μ in \mathcal{P} .*

5 Illustrative Examples

We present three illustrative examples of verification in Amaryllis.

5.1 Deriving Uniform Sampling from Bernoulli Sampling

Our first example illustrates how our rule **C-WP-SWAP** enables reasoning about conditional program expressions modularly. We verify the following function `u`, which implements a uniform sampler over the integers $\{0, 1, \dots, n-1\}$ using our Bernoulli sampler, `flip`.

`rec u n \triangleq if n = 1 then 0 else if flip(1/n) then (n-1) else u(n-1)`

<p>1045 WP-ALLOC</p> <p>1046 $\text{wp } \text{ref } V \{V'. \exists L. \lceil V' = L \rceil * L \mapsto V\}$</p> <p>1047</p> <p>1048</p> <p>1049 WP-LOAD</p> <p>1050 $L \mapsto V \vdash \text{wp } \text{get } L \{V'. L \mapsto V * \lceil V = V' \rceil\}$</p> <p>1051</p> <p>1052 WP-VALUE</p> <p>1053 $\models \Phi(V) \vdash \text{wp } V \{\Phi\}$</p> <p>1054</p>	<p>1052 WP-MONO</p> <p>1053 $\text{wp } E \{\Phi\} * (\forall V. \Phi(V) * \models \Psi(V)) * \text{wp } E \{\Psi\}$</p> <p>1054</p>	<p>1045 WP-SAMPLE</p> <p>1046 $0 < p < 1 \vdash \text{wp } \text{flip}(p) \{V. V \sim \text{Ber}_p\}$</p> <p>1047</p> <p>1048</p> <p>1049 WP-PURE</p> <p>1050 $\frac{E \rightsquigarrow_{\text{pure}} E'}{\text{wp } E' \{\Phi\} \vdash \text{wp } E \{\Phi\}}$</p> <p>1051</p> <p>1052 WP-BIND</p> <p>1053 $\text{wp } E \{V. \text{wp } K[V] \{\Phi\}\} \vdash \text{wp } K[E] \{\Phi\}$</p> <p>1054</p>	<p>1045 WP-STORE</p> <p>1046 $L \mapsto V \vdash \text{wp } \text{set } L V' \{L \mapsto V'\}$</p> <p>1047</p> <p>1048</p> <p>1049 WP-AS-EQ</p> <p>1050 $\lceil E_1 = E_2 \rceil * \text{wp } E_1 \{\Phi\} * \text{wp } E_2 \{\Phi\}$</p> <p>1051</p>
---	---	---	---

Figure 2: Selected proof rules about weakest preconditions

We show $\{n \in \mathbb{N}_{>0}\} u_n \{V. V \sim \text{Unif}_n\}$ by induction on n . The base case, $n = 1$, follows directly from **WP-PURE**, **WP-VALUE** and **DISTR-DIRAC**. For the induction step, we first apply **WP-BIND** and **WP-SAMPLE** to obtain the proof goal:

$$V \sim \text{Ber}_{1/n} \vdash \text{wp } \text{if } V \text{ then } (n-1) \text{ else } u(n-1) \{V'. V' \sim \text{Unif}_n\}$$

Next, we would like to condition on V but to do that, we have to transform the left and right sides of the entailment to forms that have the conditioning connective $C_{v \leftarrow \text{Ber}_{1/n}}$ at the top-level so that we can apply **C-CONS**. On the left side, we apply **C-INTRO** to obtain $C_{v \leftarrow \text{Ber}_{1/n}} \lceil V = v \rceil$. On the right side, we first apply **WP-MONO** and **C-BIND** to reduce the postcondition $V'. V' \sim \text{Unif}_n$ to $C_{v \leftarrow \text{Ber}_{1/n}} V' \sim (\text{if } v \text{ then } \delta_{n-1} \text{ else } \text{Unif}_{n-1})$ and obtaining a revised goal:

$$C_{v \leftarrow \text{Ber}_{1/n}} \lceil V = v \rceil \vdash \text{wp } \text{if } V \text{ then } (n-1) \text{ else } u(n-1) \\ \left\{ C_{v \leftarrow \text{Ber}_{1/n}} V' \sim (\text{if } v \text{ then } \delta_n \text{ else } \text{Unif}_{n-1}) \right\}$$

This brings us to the key step in the proof: We apply **C-WP-SWAP** to rewrite the right to a form with $C_{v \leftarrow \text{Ber}_{1/n}}$ at the top level.

$$C_{v \leftarrow \text{Ber}_{1/n}} \lceil V = v \rceil \vdash C_{v \leftarrow \text{Ber}_{1/n}} \text{wp } \text{if } V \text{ then } (n-1) \text{ else } u(n-1) \\ \left\{ V' \sim (\text{if } v \text{ then } \delta_n \text{ else } \text{Unif}_{n-1}) \right\}$$

Next, we apply **C-CONS** and use **WP-AS-EQ** to rewrite with $\lceil V = v \rceil$. A case analysis on v leaves us with two subgoals corresponding to the **then** and the **else** branches: $\text{wp } (n-1) \{V'. V' \sim \delta_{n-1}\}$ and $\text{wp } u(n-1) \{V'. V' \sim \text{Unif}_{n-1}\}$. These subgoals follow from **C-DIRAC** and the induction hypothesis, respectively. \square

5.2 Markov Blankets

Our next example is a Markov blanket: We sample three values x_1, x_2, x_3 with x_2, x_3 being dependent on x_1, x_2 respectively, and show that x_1 and x_3 are independent *conditional on* x_2 (x_2 blankets or covers x_1 's influence on x_3 ; hence, the term Markov blanket). The example, shown below, further illustrates reasoning with conditioning, and is an adaptation of a similar example included in the BLUEBELL paper [5]. The triple we obtain is stronger than BLUEBELL's in that it fully supports further framing around chain, enabling reuse.

$$\text{chain}(f, g, h) \triangleq \text{let } x_1 = f() \text{ in let } x_2 = g(x_1) \text{ in} \\ \text{let } x_3 = h(x_2) \text{ in } (x_1, x_2, x_3)$$

We formalize the specification “ x_1 and x_3 are independent conditional on x_2 ” as follows:

$$\{\text{True}\} \lambda. f() \{V. V \sim \mu_1\} \Rightarrow \\ (\forall x. \{\text{True}\} \lambda. g(x) \{V. V \sim \kappa_2(x)\}) \Rightarrow \\ (\forall y. \{\text{True}\} \lambda. h(y) \{V. V \sim \kappa_3(y)\}) \Rightarrow \\ \{\text{True}\} \text{chain}(f, g, h) \{V. \exists V_1, V_2, V_3. \lceil V = (V_1, V_2, V_3) \rceil * \\ \exists \mu. C_{v \leftarrow \mu}(\lceil V_2 = v \rceil * \exists \mu'. V_1 \sim \mu' * V_3 \sim \kappa_3(v))\}$$

The first step of the proof is replacing the postcondition of chain with a stronger, easier to use postcondition by proving the following entailment. The proof of this entailment is nontrivial, but it is the same as in Bluebell's original example, so we defer it to the appendix.

$$C_{v_1 \leftarrow \mu_1} \left(\lceil V_1 = v_1 \rceil * C_{v_2 \leftarrow \kappa_2(v_1)}(\lceil V_2 = v_2 \rceil * V_3 \sim \kappa_3(v_3)) \right) \\ \vdash \exists \mu. C_{v \leftarrow \mu}(\lceil V_2 = v \rceil * \exists \mu'. V_1 \sim \mu' * V_3 \sim \kappa_3(v))$$

Using **WP-MONO**, our proof goal becomes

$$\{\text{True}\} \text{chain}(f, g, h) \{V. \exists V_1, V_2, V_3. \lceil V = (V_1, V_2, V_3) \rceil * \\ C_{v_1 \leftarrow \mu_1} \left(\lceil V_1 = v_1 \rceil * C_{v_2 \leftarrow \kappa_2(v_1)}(\lceil V_2 = v_2 \rceil * V_3 \sim \kappa_3(v_3)) \right)\}$$

To prove this, we first apply f 's Hoare triple using **WP-BIND** to obtain a random variable $V_1 \sim \mu_1$ that is stored in x_1 . Next, to use g 's Hoare triple, we must condition on V_1 after applying **C-WP-SWAP** to the goal. However, we cannot immediately apply **C-WP-SWAP** since the conditioning $C_{v_1 \leftarrow \mu_1}(\dots)$ in the goal is under an existential quantifier. We commute the conditioning and existential quantification use a *derived* rule.

$$\text{C-EXISTS-RV} \\ \frac{C_{v \leftarrow \mu} \exists V. \Phi(v)(V) \quad \forall v, V_1, V_2. \lceil V_1 = V_2 \rceil * \Phi(v)(V_1) \vdash \Phi(v)(V_2)}{\exists V. C_{v \leftarrow \mu} \Phi(v)(V)}$$

C-EXISTS-RV is derived as follows. First, apply **C-SKOLEM** to the first premise to obtain a function f satisfying $C_{v \leftarrow \mu} \Phi(v)(f(v))$. Next, use **C-RV** to obtain a random variable $X \sim \mu$ that the conditioning is actually on. Define $V := \lambda \rho. f(X(\rho))\rho$ and apply **C-CONS** to get the goal $\lceil X = v \rceil * \Phi(v)(f(v)) \vdash \Phi(v)(\lambda \rho. f(X(\rho)))$, which follows from the rule's second premise.

Using **C-EXISTS-RV**, we can continue the proof of the example: after moving the existentials below the conditioning, we can apply **C-WP-SWAP** to condition on V_1 , and **WP-BIND** to use the Hoare triple

for g . To complete the proof, we repeat this pattern once more for h . \square

This example cannot be verified easily in existing GPLs for different reasons. BLUEBELL, due to its C-WP-SWAP rule which must be used to verify the three function calls, requires ownership of every variable in the precondition of the triple, thus prohibiting framing of resources around the final chain triple, hindering its reuse. All other GPLs conflate probabilistic independence and disjointness and, therefore, they cannot express assertions like chain's postcondition, which mentions all three variables V_1, V_2, V_3 on both sides of a separating conjunction.

5.3 Sorted Linked Lists with Random Elements

Our next example is a standard pointer-based sorted linked list, with three operations: `init` (new empty list), `ins` (insert a value) and `del` (remove a value and return it) operations. The list is an interesting example for us because it leverages dynamic allocation, which prior GPLs do not support. Additionally, *sorted* lists interact with randomness in challenging ways. If a sampled value is inserted into a sorted list, the position at which it ends up depends on what samples drawn for other elements of the list. Hence, the *structure* of the list is itself randomized. Despite this entanglement between data and list structure, we are able to prove a strong specification: If each inserted value's distribution is independent of the list's state at the time of insertion, then any value removed from the list is independent of the remaining list's state at the time of removal.

We start with an implementation of the sorted list's three operations in our language, which should be unsurprising.

```

init  $\triangleq$   $\lambda x$ . ref ()
ins  $\triangleq$   $\lambda x \ell$ . (rec loop  $\ell' = \text{let } v = \text{get } \ell' \text{ in}$ 
  if  $v = () \vee x \leq \text{fst } v$  then ref ( $x, \ell'$ ) else
  set  $\ell' (\text{fst } v, \text{loop}(\text{snd } v)); ; \ell'$ )
del  $\triangleq$   $\lambda x \ell$ . (rec loop  $\ell' = \text{let } v = \text{get } \ell' \text{ in}$ 
  if  $v = ()$  then  $\ell'$  else
  if  $x = \text{fst } v$  then  $\text{snd } v$  else
  set  $\ell' (\text{fst } v, \text{loop}(\text{snd } v)); ; \ell'$ )

```

To verify our implementation, we will define a representation predicate $\text{psll} : RV_{Loc} \times 2^{RV_Z \times \mathbb{D}(\mathbb{Z})} \rightarrow \text{pProp}$ for randomized sorted linked lists. Informally, $\text{psll}(L, ds)$ holds iff ds is a set of the form $\{(V_1, \mu_1), \dots, (V_n, \mu_n)\}$, L is a random list containing the random elements V_1, \dots, V_n in sorted order, and each V_i is distributed as μ_i . (Technically, L, V_i s are random variables.) Equipped with this definition, we prove the following specification for the list operations in Amaryllis:

$$\{\text{True}\} \text{init } () \{\exists L. [V = L] * \text{psll}(L, \emptyset)\}$$

$$\{\text{psll}(L, ds) * X \sim \mu\} \text{ins } X L$$

$$\{V. \exists L'. [V = L'] * \text{psll}(L', (X, \mu) \cup ds)\}$$

$$\{\text{psll}(L, (X, \mu) \cup ds)\} \text{del } X L$$

$$\{V. \exists L'. [V = L'] * \text{psll}(L', ds) * X \sim \mu\}$$

The precondition of `ins` requires that the inserted element X 's distribution be independent of $\text{psll}(L, ds)$. Dually, the postcondition of

`del` states that the removed element's distribution is independent of the rest of the list.

We define psll in three steps. First, we define the standard list representation predicate $\text{ll} : RV_{Loc} \rightarrow \text{list } \mathbb{Z} \rightarrow \text{pProp}$. $\text{ll}(L, zs)$ holds if the *random location* L holds the *deterministic elements* of zs in order.

$$\text{ll}(L, zs) \triangleq \begin{cases} L \mapsto \lambda. () & \text{for } zs = [] \\ \exists L'. L \mapsto (\lambda \rho. (z, L'(\rho))) * \text{ll}(L', zs) & \text{for } zs = z :: zs \end{cases}$$

Next, we define $\text{sll}(L, zs) := \text{ll}(L, zs) * \text{sorted}(zs)$ to add that the lists be sorted. Using $\text{sll}(L, zs)$ we prove standard non-probabilistic specifications for the list operations, e.g.,

$$\{\text{sll}(L, zs_1 ++ z :: zs_2)\} \text{del } z L \{V. \exists L'. [V = L'] * \text{sll}(L', zs_1 ++ zs_2)\}$$

Finally, we use Amaryllis's conditioning operator \mathbf{C} to lift sll to the predicate $\text{psll}(L, ds)$ that relates a random list L to a set ds of random elements and their distributions. The distributions must be independent.

$$\text{psll}(L, ds) \triangleq \mathbf{C}_{(v_1, \dots, v_n) \leftarrow \pi_2(ds_1) \otimes \dots \otimes \pi_2(ds_n)} \exists zs. (\text{sll}(L, zs) * \left[\pi_1(ds_1) = v_1 \wedge \dots \wedge \pi_1(ds_n) = v_n \right] * zs \text{ is a permutation of } v_1, \dots, v_n)$$

The proof of the probabilistic specification of `init` is straightforward. In verifying `del`, a key step is the transformation of the postcondition into a conditioning on $\mu \otimes \pi_2(ds_1) \otimes \dots \otimes \pi_2(ds_n)$ using the strong convexity of the sll predicate. After that, an application of C-WP-SWAP and C-CONS leads to a proof goal with deterministic values, which is established using the non-probabilistic specification of `del`. The proof strategy for `ins` is similar. For more details, refer to our Rocq development [3]. \square

To further test modularity of Amaryllis, we verified a client for our list. The client iterates from 0 to i and, at each step, randomly decides whether to insert an element sampled from Unif_i or to not insert anything:

```

rec rand_list i = if i ≤ 0 then init() else
  if flip(½) then ins(u(i), rand_list(i - 1)) else rand_list(i - 1)

```

We have shown that, conditional on the outcomes of the `flip(½)` operations, we obtain a probabilistic sorted linked list with a length corresponding to the number of successful flips:

$$\{\text{True}\} \text{rand_list } n \{V. \exists L. [V = L] * \mathbf{C}_{bs \leftarrow \mu_{n \text{ flips}}} (\exists xs. \text{psll}(L, xs) * \text{len}(xs) = \# \text{ones}(bs))\}$$

This postcondition also entails that the elements of the list are conditionally independent. The proof can be found in our Rocq development [3]. \square

6 Related Work

We divide probabilistic logics in two main groups: the *general-purpose probabilistic logics* (GPLs) and the *property-specific probabilistic logics* (PPLs). GPLs provide native support for assertions describing properties of the whole *distribution* over program states (e.g., probabilistic (conditional) independence, which is not directly expressible in PPLs). PPLs instead provide ways to reason about

specific probabilistic properties of interest (e.g., error bounds, expected time complexity, etc.) within the framework of standard Hoare/separation logics (e.g., Iris, WP calculi).

General-purpose probabilistic logics. Ellora [7] advocate for assertions on probability distributions in the context of Hoare-style logic. PSL [9] pioneered the idea of internalizing probabilistic independence via separating conjunction for a discrete probabilistic imperative language on a fixed set of variables. The separation of PSL insists on both disjointness of variables (i.e., locations) and probabilistic independence *at the same time*, which (as explained in Section 2.2) is problematic when considering dynamic memory allocation. Reasoning under conditioning was then studied in DIBI [6], which internalized conditioning using a non-commutative connective in addition to PSL-style separation.

Lilac [20] introduced a measure-theoretic model of separation that internalizes *purely* probabilistic independence of *continuous* random variables. Lilac's programs are first-order functional programs with bounded loops, immutable state, and the ability to sample from continuous distributions. On top of the new model of separation (refined in follow-up work [21]), Lilac proposed to express reasoning under conditioning by using a new *conditioning* modality based on disintegrations.

Later, BLUEBELL [5] proposed a variant of Lilac's conditioning modality (in the context of *discrete* distributions) which was shown to satisfy very modular entailment rules, and could encode and mix both unary and relational concepts within a unified framework. The BLUEBELL logic was the first to show how mutable state could be integrated with Lilac's notion of separation without conflating disjointness and independence, but its solution is specific to finite variable stores and does not generalize to dynamic allocation. Furthermore, BLUEBELL illustrated the advantages of **C-WP-SWAP**, but their model only supports the rule with a heavy side-condition that disables the use of frame around it, thus inhibiting modularity.

BASL [14] looks at the problem of supporting the Bayesian updating statement **observe** in a Lilac-like logic. Like Lilac, BASL also lacks a general **C-WP-SWAP** rule, and remedies this by providing some specialized rules for conditional sampling. Understanding the conditions under which **C-WP-SWAP** can be validated in a model for continuous distributions remains an open problem.

Outcome Separation Logic [31] is a probabilistic logic for a language with dynamic allocation, but where separation is used exclusively for heap disjointness, not probabilistic independence.

pcOL [32] is the first logic to integrate invariants and concurrency into a GPL. pcOL supports almost-surely terminating programs with finite variable stores, sampling from discrete distributions, and adopts a Lilac/BLUEBELL-style conditioning modality, that also enjoys a rather general version of the **C-WP-SWAP** rule (their **SPLIT1** and **SPLIT2** rules). The proof of soundness, however, relies on the simplifying effect of adopting a PSL-style separation conflating disjointness and independence, and thus is fundamentally incompatible with dynamic allocation. Moreover, our approach based on PFP updates allows us to unify **SPLIT1** and **SPLIT2** and remove some of their side-conditions, widening their applicability.

Concerning mechanization of GPLs, there is some recent work by Yan et al. [29] that mechanizes PSL in Isabelle/HOL and applies it to security verification for oblivious algorithms. However, to our

knowledge, Amaryllis is the first GPL supporting both independence and conditioning that has been mechanized in a proof assistant.

Property-specific probabilistic logics. The property-specific logics for probabilistic programs stick to traditional assertions over state, but include resources that can be used to prove certain specific properties of the induced distributions (e.g., expectations, bounds on probability of events). One subclass of this category are *quantitative logics* where assertions evaluate to a quantity, like probability or expected value, instead of a binary truth value [10, 18, 19]. This idea nicely enables modular reasoning and leads to automation, but can't be applied to other probabilistic properties like independence.

Most closely related to our work are PPLs based on Iris [16], where assertions are standard separation logic assertions on state, augmented with ghost state representing some quantity related to the distribution. Eris [1] aims at proving bounds on the probability that the postcondition will be violated. A very nice feature of Eris is that the bounds themselves are treated as ghost resources, called "error credits". Tachis [13] uses the same technique to prove bounds on expected runtime. Coneris [22] extends the approach to handle concurrency.

PPLs have been also shown to work well for proving relational properties, e.g., equivalence of the distribution induced by two programs. The pRHL [8] logic pioneered the use of probabilistic couplings to reduce relational properties to proving relational triples with standard assertions on state. Polaris [27] and Clutch [11], for example, prove contextual refinement in Iris, using probabilistic couplings. Approxix [12] provides a way to prove *approximate* refinements. BLUEBELL showed how couplings can be derived on top of conditioning in a relational GPL, which could offer an approach for a future relational extension of Amaryllis.

As a result of using standard assertions over state, the Iris-based PPLs can reuse the Iris infrastructure as is (including step-indexing and impredicative invariants), and support very rich source languages including higher-order features and sometimes concurrency. None of them internalize probabilistic independence, however, and they are specialized to proving specific properties of the program semantics. As a future direction, it would be interesting to encode these abstractions (e.g., error credits) on top of the Amaryllis base logic, for example by deriving variations of our WP that encode the abstraction-specific invariants.

Nondeterminism and Outcome Logic. Integrating nondeterminism and random choices in a compositional way is a longstanding issue, given that general solutions have been proven impossible [33]. A well-known solution, proposed by Varacca and Winskel [28], shows that by representing distributions with the more detailed *indexed valuations*, which record the random choices that lead to an outcome through indexes, probabilistic choice can be combined with nondeterminism compositionally. As explained in Section 2.2, the use of an indexed valuations-style model in Amaryllis is not motivated by nondeterminism (of the allocator), but by our need to lift the underlying notion of separation of state to separation of distributions. As a byproduct of using indexed valuations, we can also combine random choices with nondeterminism in the allocator, and prove triples that are valid independently of the specific (non-probabilistic) implementation of allocation one picks.

Vancouver, BC, Canada, June 24-27, 2019. IEEE, 1–13. DOI: [10.1109/LICS.2019.8785707](https://doi.org/10.1109/LICS.2019.8785707).

A Proof Rules

We give a list of proof rules in this section. The proof rules for conditioning can be found in figure 3, for C-frameable in figure 4 and for WP in figure 5.

A.1 Example Proofs

Here we write down the details of the derivation that we used for the Markov blanket example in Section 5.2.

We define $\mu_0 \triangleq \mu_1 \prec \kappa_2$. Following Bayes' theorem, we define $\mu_2 \triangleq \lambda v_2. \sum_{v_1} \mu_0(v_1, v_2)$ and $\kappa_1 v_2 \triangleq \lambda v_1. \frac{\mu_0(v_1, v_2)}{\mu_2(v_2)}$ such that we can invert the conditional probability: $\mu_1 \prec \kappa_2(v_1, v_2) = \mu_2 \prec \kappa_1(v_2, v_1)$. We use this equality together with C-TRANSF and derive the following:

$$\begin{aligned}
& C_{v_1 \leftarrow \mu_1} \llbracket V_1 = v_1 \rrbracket * C_{v_2 \leftarrow \kappa_2(v_1)} \llbracket V_2 = v_2 \rrbracket * V_3 \sim \kappa_3(v_3) \\
\vdash & C_{v_1 \leftarrow \mu_1} C_{v_2 \leftarrow \kappa_2(v_1)} \llbracket V_1 = v_1 \rrbracket * \llbracket V_2 = v_2 \rrbracket * V_3 \sim \kappa_3(v_3) && \text{C-FRAME} \\
\vdash & C_{(v_1, v_2) \leftarrow \mu_1 \prec \kappa_2} \llbracket V_1 = v_1 \rrbracket * \llbracket V_2 = v_2 \rrbracket * V_3 \sim \kappa_3(v_3) && \text{C-FUSE} \\
\vdash & C_{(v_2, v_1) \leftarrow \mu_1 \prec \kappa_2} \llbracket V_1 = v_1 \rrbracket * \llbracket V_2 = v_2 \rrbracket * V_3 \sim \kappa_3(v_3) && \text{C-TRANSF} \\
\vdash & C_{v_2 \leftarrow \mu} C_{v_1 \leftarrow \kappa(v_2)} \llbracket V_1 = v_1 \rrbracket * \llbracket V_2 = v_2 \rrbracket * V_3 \sim \kappa_3(v_3) && \text{C-FUSE} \\
\vdash & C_{v_2 \leftarrow \mu} \llbracket V_2 = v_2 \rrbracket * \left(C_{v_1 \leftarrow \kappa(v_2)} \llbracket V_1 = v_1 \rrbracket \right) * V_3 \sim \kappa_3(v_3) && \text{DISTR-STR-COMPLEX} \\
\vdash & C_{v_2 \leftarrow \mu} \llbracket V_2 = v_2 \rrbracket * V_1 \sim \kappa(v_2) * V_3 \sim \kappa_3(v_3) && \text{C-INTRO} \\
\vdash & \exists \mu \kappa. C_{v \leftarrow \mu} \llbracket V_2 = v \rrbracket * V_1 \sim \kappa(v) * V_2 \sim \kappa_3(v) \\
\vdash & \exists \mu. C_{v \leftarrow \mu} \llbracket V_2 = v \rrbracket * \exists \mu'. V_1 \sim \mu' * V_2 \sim \kappa_3(v) && \text{C-CONS}
\end{aligned}$$

B Full Model

In this section, we present the full model by presenting five layers of resource algebras that build on top of each other:

B.1 Bluebell's PSp Resource Algebra

A unary and finite version of the Bluebell's PSp $_{\Omega}$ resource algebra forms the basis for our probabilistic RA [5]. The resources are probability spaces, and the resource composition is the *Independent Product*.

In case the independent product does not exist, the probability spaces compose to an invalid element \perp .

B.2 Dynamic Domain Extensions with PSpD

In this section, we present a finite representation of probability spaces over $Rid \rightarrow \mathbb{B}$. We use Bluebell's PSp $_{\Omega}$ but do not choose a fixed Ω , but instead dynamically extend Ω on the fly as we encounter new randomness IDs. All randomness IDs that are not mentioned in the current probability space are considered to have an unknown distribution. That way, elements of the resource algebra can still represent the infinite probability space that we are interested in, while actually being finite.

The resource algebra PSpD in this layer explicitly tracks which RIDs are mentioned in the probability space. Its type is therefore $Rid \times PSp$. The *domain extension operation* Ext extends an element $(I, \mathcal{F}, \mu) \in PSpD$ with a set of new RIDs I' , without adding any information about the distribution of those RIDs. This operation

can be seen as going from a resource algebra element where the RIDs in I' are implicitly unknown to an equivalent resource algebra element where they are explicitly mentioned but still unknown.

$$\text{Ext}((I, \mathcal{F}, \mu), I') :=$$

$$\text{Compl}(\{\{m_1 \cup m_2 \mid m_1 \in E \wedge m_2 : I' \rightarrow \mathbb{B}\} \mid E \in \mathcal{F}\}, \mu_{\text{ext}})$$

$$\mu_{\text{ext}}(E) := \mu(\{m : I \rightarrow \mathbb{B} \mid \exists m' \in E, m \preceq m'\})$$

The function Compl calculates the completion of the probability space. We can use the domain extension operation to define an equivalence relation on this resource algebra:

$$(I_1, \mathcal{F}_1, \mu_1) \equiv (I_2, \mathcal{F}_2, \mu_2) :=$$

$$\text{Ext}((I_1, \mathcal{F}_1, \mu_1), I_2 \setminus I_1) = \text{Ext}((I_2, \mathcal{F}_2, \mu_2), I_1 \setminus I_2)$$

The resource composition is defined as follows:

$$(I_1, \mathcal{F}_1, \mu_1) \cdot (I_2, \mathcal{F}_2, \mu_2) :=$$

$$(I_1 \cup I_2, \text{Ext}((I_1, \mathcal{F}_1, \mu_1), I_2 \setminus I_1) \cdot_{\text{PSp}} \text{Ext}((I_2, \mathcal{F}_2, \mu_2), I_1 \setminus I_2))$$

The persistent core of a resource algebra element is the *sure* part of the probability space: it consists of all probability 0 and 1 events.

$$|(I, \mathcal{F}, \mu)| = (I, \{E \in \mathcal{F} \mid \mu(E) = 0 \vee \mu(E) = 1\}, \mu)$$

Weighted Sum. Next, we'll define the weighted sum on this layer. We would like the weighted sum to represent *conditioning*, thus we require that the decomposition given by the branches of the weighted sum is induced by a random variable (the random variable we are conditioning on). This is equivalent to requiring the support of the branches to be pairwise disjoint.

$$\text{has_rv}(\mu, \kappa) :=$$

$$\forall \rho. \forall i j \in \text{supp } \mu. \text{poss}(\rho, \kappa(i)) \rightarrow \text{poss}(\rho, \kappa(j)) \rightarrow i = j$$

where $\text{poss}(\rho, (\mathcal{F}, \mu)) := \exists m \in \text{supp}(\mathcal{F}, \mu). \forall i \in \text{dom}(m). m[i] = \rho(i)$.

In the weighted sum on this layer, an event is measurable if it is measurable in all branches. Its probability is the sum of the probabilities in the branches.

$$\sigma_intersec(\mu, \kappa) := \bigcap \pi_{\sigma}(\text{Ext}(\kappa(i), \bigcup_{i \in \text{supp}(\mu)} I(\kappa(i))))$$

$$\text{wsum}(\mu, \kappa) := \text{bind}(\mu, \lambda i. \text{Ext}(\kappa(i), \bigcup_{i \in \text{supp}(\mu)} I(\kappa(i)))|_{\sigma_intersec(\mu, \kappa)})$$

$$\text{if } \text{has_rv}(\mu, \kappa) \wedge \forall i \in \text{supp}(\mu). \mathcal{V}(\kappa(i))$$

$$\pi_{\mu}(\text{bind}(\mu, \kappa))(X) = \sum_{v \in \text{supp}(\mu)} \mu(v) \cdot \pi_{\mu}(\kappa(v))(X)$$

Operation \circ . We define the operation \circ on this layer using the weighted sum:

$$(I_1, \mathcal{F}_1, \mu_1) \circ (I_2, \mathcal{F}_2, \mu_2) =$$

$$\text{wsum}(\mu_1|_{\text{atoms}(\mathcal{F}_1 \cap \mathcal{F}_2)}, \lambda A. (F_1, \mu_1)|_A \cdot (F_2, \mu_2)|_A)$$

$$\text{if } \forall E \in F_1 \cap F_2. \mu_1(E) = \mu_2(E)$$

1625	$\vdash X \sim \mu \multimap \lceil X = Y \rceil \multimap Y \sim \mu$	$\vdash (\lambda. v) \sim \delta(v)$	$\text{Own}(\varepsilon, X) * \lceil X = Y \rceil \vdash \text{Own}(\varepsilon, Y)$	$\lceil X = Y \rceil \vdash \square \lceil X = Y \rceil$	1683
1626					1684
1627					1685
1628					1686
1629					1687
1630	$\text{StrConvex}(\Psi) \triangleq C_{v \leftarrow \mu} \Phi(v) * \Psi \vdash \Psi * C_{v \leftarrow \mu} \Phi(v)$		$C_{x \leftarrow \mu} \text{true} \vdash C_{x \leftarrow \mu} P(x)$	$\frac{\forall v \in \text{supp}(\mu). \Phi_1(v) \vdash \Phi_2(v)}{C_{v \leftarrow \mu} \Phi_1(v) \vdash C_{v \leftarrow \mu} \Phi_2(v)}$	1688
1631					1689
1632					1690
1633	$C_{v \leftarrow \mu} C_{w \leftarrow \kappa(v)} \Phi(v, w) \vdash C_{(v, w) \leftarrow \mu \times \kappa} \Phi(v, w)$	$C_{v \leftarrow \mu} \forall x. \Phi(x, v) \vdash \forall x. C_{v \leftarrow \mu} \Phi(x, v)$		$X \sim \mu \vdash C_{v \leftarrow \mu} \lceil X = v \rceil$	1691
1634					1692
1635					1693
1636					1694
1637	$C_{v \leftarrow \mu} \Phi(v) \vdash \exists X. C_{v \leftarrow \mu} \lceil X = v \rceil * \Phi(v)$	$\frac{f \text{ injective} \quad \forall v \in \text{supp}(\mu_1). \mu_1(v) = \mu_2(f(v))}{C_{v \leftarrow \mu_2} \Phi(v) \vdash C_{v \leftarrow \mu_1} \Phi(f(v))}$			1695
1638					1696
1639					1697
1640					1698
1641	$C_{v \leftarrow \mu} \exists a. \Phi(a, v) \vdash \exists f. C_{v \leftarrow \mu} \Phi(f(v), v)$	$C_{v \leftarrow \mu} X \sim \kappa(v) \vdash X \sim \text{bind}(\mu, \kappa)$		$\frac{\Phi(v) \vdash \perp \quad v \in \text{supp}(\mu)}{C_{v \leftarrow \mu} \Phi(v) \vdash \perp}$	1699
1642					1700
1643					1701
1644					1702
1645	$C_{v \leftarrow \delta(i)} \Phi(v) \vdash \Phi(i)$	$\frac{\forall v. Q(v) \vdash \lceil X = v \rceil}{\text{Own}(\bullet, g) * C_{v \leftarrow \mu} Q(v) \vdash C_{v \leftarrow \mu} (\text{Own}(\bullet, g _{X=v})) * Q(v)}$	$\text{supp } \mu \subseteq E * C_{v \leftarrow \mu} \Phi(v) \vdash C_{v \leftarrow \mu} v \in E * \Phi(v)$		1703
1646					1704
1647					1705
1648					1706
1649					1707
1650	$\frac{\Psi \text{ pure}}{\text{StrConvex}(\Psi)}$	$\text{StrConvex}(L \mapsto V)$		$\text{StrConvex}(X \sim \mu)$	1708
1651					1709
1652					1710
1653					1711

Figure 3: Proof rules for conditioning

1656	$\frac{P \vdash Q}{\boxtimes P \vdash \boxtimes Q}$	$\boxtimes P * C_{v \leftarrow \mu} \Phi(v) \vdash C_{v \leftarrow \mu} (\boxtimes P) * \Phi(v)$	$\boxtimes P \vdash P$	$\boxtimes \boxtimes P \vdash \boxtimes P$	$\boxtimes \forall v. \Phi(v) \vdash \forall v. \boxtimes \Phi(v)$	1716
1657						1717
1658						1718
1659						1719
1660						1720
1661	$\boxtimes \exists v. \Phi(v) \vdash \exists v. \boxtimes \Phi(v)$	$\blacksquare P \vdash \boxtimes P$	$\boxtimes P * \boxtimes Q \vdash \boxtimes (P * Q)$	$\square P \vdash \boxtimes P$	$X \sim \mu \vdash \boxtimes (X \sim \mu)$	1721
1662						1722
1663						1723
1664						1724
1665						1725
1666						1726
1667						1727
1668						1728
1669						1729
1670						1730
1671						1731
1672						1732
1673						1733
1674						1734
1675						1735
1676						1736
1677						1737
1678						1738
1679						1739
1680						1740
1681						1741
1682						1742

Figure 4: Proof rules for C-frameable

B.3 Tracking undeterminate RIDs with PSpU

On this layer, we additionally track in each resource algebra element a set of RIDs that is *undeterminate*, i.e. that the probability space has no information on. Let $\mathcal{N}(a)$ be the largest probability 0 event in the probability space of a .

$$\text{PSpU} := \{(U, a) \mid a \in \text{PSpD} \wedge \text{is_undeterminate}(U, a)\}$$

$$(U, a) \equiv (U', a') := U = U' \wedge a \equiv a'$$

$$(U, a) \cdot (U', a') := (U \cap U', a \cdot a')$$

$$\mathcal{V}((U, a)) := \mathcal{V}(a)$$

$$(U, a) \circ (U', a') := (U \cap U', a \circ a')$$

$$\text{wsum}(\mu, \kappa) := \left(\bigcap_{i \in \text{supp}(\mu)} \pi_U(\kappa(i)), \text{wsum}(\mu, \pi_a(\kappa(i))) \right)$$

$$|(U, a)| := (U, |a|)$$

$$\text{is_undeterminate}(U, a) := \forall E \in a. E \cap \mathcal{N}(a) = \emptyset \rightarrow$$

$$\forall m_1. \forall m_2 m_3 : (U \cap I(a)) \rightarrow \mathbb{B}. m_1 \cup m_2 \in E \rightarrow m_1 \cup m_3 \in E$$

1741	WP-BIND	WP-LOAD	WP-ALLOC	1799
1742	$\text{wp } E \{V. \text{wp } K[V] \{\Phi\}\} \vdash \text{wp } K[E] \{\Phi\}$	$L \mapsto V \vdash \text{wp } \text{get } L \{X. L \mapsto V * [X = V]\}$	$\text{wp } \text{ref } V \{V'. \exists L. [V' = L] * L \mapsto V\}$	1800
1743				1801
1744	WP-SAMPLE	WP-STORE	WP-PURE	1802
1745	$\frac{0 < p < 1}{\text{wp } \text{flip}(p) \{V. V \sim \text{Ber}_p\}}$	$L \mapsto V \vdash \text{wp } \text{set } L V' \{V''. [V'' = ()] * L \mapsto V'\}$	$\frac{E \rightsquigarrow_{\text{pure}} E'}{\text{wp } E' \{\Phi\} \vdash \text{wp } E \{\Phi\}}$	1803
1746				1804
1747				1805
1748				1806
1749	WP-AS-EQ	WP-VALUE	WP-MONO	1807
1750	$[E_1 = E_2] * \text{wp } E_1 \{\Phi\} * \text{wp } E_2 \{\Phi\}$	$\models \Phi(V) \vdash \text{wp } V \{\Phi\}$	$\text{wp } E \{\Phi\} * (\forall V. \Phi(V) * \models \Psi(V)) * \text{wp } E \{\Psi\}$	1808
1751				1809
1752	UPD-WP	WP-UPD	C-WP-SWAP	1810
1753	$\models \text{wp } E \{\Phi\} \vdash \text{wp } E \{\Phi\}$	$\text{wp } E \{V. \models \Phi(V)\} \vdash \text{wp } E \{\Phi\}$	$C_{v \leftarrow \mu} V. \text{wp } E \{\Phi(v, V)\} \vdash \text{wp } E \{V. C_{v \leftarrow \mu} \Phi(v, V)\}$	1811
1754				1812
1755				1813
1756				1814
1757				1815

Figure 5: WP rules

B.4 The Authoritative Probability Space RA

PSpU_Auth

Since we want to use our RA in the state interpretation, we need to make it authoritative. Additionally, we add a *Namespace* to the authoritative part that tracks a set of globally free randomness IDs. This will simplify allocating a new RID since we can ensure that it doesn't yet exist anywhere in the context.

$$\text{PSpU_Auth} := \text{PSpU} \times \text{Option}(\text{Ex}(\text{PSpD} \times N))$$

Composition, \circ and equivalence are pointwise on the underlying pRAs. Validity requires the fragment to be included in the authoritative part, and the RIDs in the namespace to be globally undetermined.

$$\begin{aligned} \mathcal{V}(a, \text{Some}(\text{ex}(b, N))) &:= \mathcal{V}(a) \wedge \mathcal{V}(b) \wedge \pi_{\text{PSpD}}(a) \preceq_{\circ} b \wedge \\ &\quad N \subseteq \pi_U(b) \wedge \text{is_undetermined}(N, b) \\ \mathcal{V}(a, \perp) &:= \mathcal{V}(a) \end{aligned}$$

The weighted sum is also pointwise, given that all branches are valid:

$$\begin{aligned} \text{wsum}(\mu, \kappa) &:= (\text{wsum}(\mu, \pi_{\text{frag}}(\kappa)), \text{wsum}(\mu, \pi_{\text{auth}}(\kappa))) \\ &\quad \text{if } \forall i \in \text{supp}(\mu). \mathcal{V}(\kappa(i)) \end{aligned}$$

B.5 Adding Iris Resources as Random Variables with PSp_RV_M

Next, we add (non-probabilistic) Iris resources as random variables. For example, $\lambda \rho. L(\rho) \mapsto \rho(i)$ says that the location $L : \text{RV}_{\text{loc}}$ (which might be a different location based on the random choices) maps to whatever value the randomness ID i resolves to.

The key challenge here is how to combine multiple Iris RVs in the weighted sum. We can use that we already require conditioning on a random variable.

First, we need a way to extract the random variable that is conditioned on from the branches (up to almost sure equality). Define $\text{get_rv}(\mu, \kappa)(\rho)$ as the unique $i \in \text{supp}(\mu)$ such that $\text{poss}(\kappa(i), \rho)$ if such an i exists. For example, for the weighted sum $\text{wsum}(\mu, \lambda v. [X = v])$, we have that $\text{get_rv}(\mu, \kappa) =_{\text{as}} X$.

This helps us in combining Iris resources across branches, for example:

$$\text{wsum}(\mu, \lambda v. ([X = v], \lambda \rho. L(\rho) \mapsto v)) = (X \sim \mu, \lambda \rho. L(\rho) \mapsto X(\rho))$$

Let M be a unital RA.

$$\begin{aligned} \text{PSp_RV}_M &:= \text{PSpU_Auth} \times (\text{RV}_M) \\ \mathcal{V}(a, X) &:= \mathcal{V}(a) \wedge \forall \rho. \mathcal{V}(X(\rho)) \\ (a, X) \equiv (a', X') &:= (a \equiv a') \wedge \forall \rho. X(\rho) \equiv X'(\rho) \\ (a, X) \cdot (b, Y) &:= (a \cdot b, \lambda \rho. X(\rho) \cdot Y(\rho)) \\ (a, X) \circ (b, Y) &:= (a \circ b, \lambda \rho. X(\rho) \cdot Y(\rho)) \\ \text{wsum}(\mu, \kappa) &:= (\text{wsum}(\mu, \pi_1(\kappa)), \\ &\quad \lambda \rho. \text{if } \text{poss}(\text{wsum}(\mu, \pi_{1, \text{frag}}(\kappa)), \rho) \\ &\quad \text{then } \pi_2(\kappa)(\text{get_rv}(\mu, \pi_{1, \text{frag}}(\kappa))(\rho)) \\ &\quad \text{else } \varepsilon \end{aligned}$$

Note that in the random variable that we get from the weighted sum, ρ s that are probability 0 in all branches are mapped to the unit ε . Those cases don't matter on the level of the logic and they cannot be mapped to a unique branch - we choose ε to ensure that those cases still remain valid.

C Operational Semantics

The language is given by a small step relation \rightsquigarrow . The proposition $e, \sigma \rightsquigarrow e_1, \sigma_1, e_2, \sigma_2, p$ states that expression e under state σ steps to e_1 and σ_1 with probability p and to e_2 and σ_2 with probability $1 - p$.

We lift this small step to a multi step $E_1, \Sigma_1, \mathcal{P}_1, N_1 \rightsquigarrow_{\text{RV}} E_2, \Sigma_2, \mathcal{P}_2, N_2$ which allows us to compose several steps. We are using PSpD as model for probability spaces to simplify the allocation of RIDs. Additionally, we track all globally free RIDs in N_1 and N_2 . Because the expressions and state are lifted to RV_{Expr} resp. RV_{State} but the probability space only contains partial maps we relate partial maps and functions using \prec . We define $m \prec \rho := \forall i \in \text{dom}(m). m[i] = \rho(i)$.

$$\begin{aligned}
& E_1, \Sigma_1, \mathcal{P}_1, N_1 \rightsquigarrow_{RV} E_2, \Sigma_2, \mathcal{P}_2, N_2 := \\
& \exists X, \iota, e_1, \sigma_1, e_2, \sigma_2, p. \\
& X \in \text{atoms}(\mathcal{P}_1) \wedge \\
& \iota \in N_1 \wedge \\
& \mathcal{P}_2 = \text{wsum}(\pi_\mu(\mathcal{P}_1)|_{\text{atoms}(\mathcal{P}_1)}, \\
& \quad \lambda Y. \mathcal{P}_1|_Y \cdot \text{if } X = Y \text{ then } \iota \sim \text{Ber}_p \text{ else } \varepsilon) \wedge \\
& N_2 = N_1 \setminus \{\iota\} \wedge \\
& (\forall m \in X. \\
& \quad E_1(m), \Sigma_1(m) \rightsquigarrow e_1, \sigma_1, e_2, \sigma_2, p) \wedge \\
& (\forall m \in X, \rho_t, \rho_f. \\
& \quad m[\iota \mapsto \text{true}] \prec \rho_t \rightarrow m[\iota \mapsto \text{false}] \prec \rho_f \rightarrow \\
& \quad E_2(\rho_t) = e_1 \wedge \Sigma_2(\rho_t) = \sigma_1 \wedge \\
& \quad (p \neq 1 \rightarrow E_2(\rho_f) = e_2 \wedge \Sigma_2(\rho_f) = \sigma_2)) \wedge \\
& (\forall Y \in \text{atoms}(\mathcal{P}_1), m \in Y, \rho. \\
& \quad Y \neq X \rightarrow m \prec \rho \rightarrow \\
& \quad E_1(\rho) = E_2(\rho) \wedge \Sigma_1(\rho) = \Sigma_2(\rho))
\end{aligned}$$

An atom X of \mathcal{P}_1 and an $\iota \in N_1$ are nondeterministically chosen. The step will happen in branch X of \mathcal{P}_1 and it will allocate RID ι . The step in X goes to $e_1, \sigma_1, e_2, \sigma_2, p$ which means that we split X into two branches: one containing e_1, σ_1 , having probability p and one containing e_2, σ_2 having probability $1 - p$. \mathcal{P}_2 is therefore constructed by leaving every branch of \mathcal{P}_1 unchanged except X which we split into two atoms X_1 and X_2 by composing with $\iota \sim \text{Ber}_p$. The RID ι is undeterminate in each atom of \mathcal{P}_2 except in X_1 and X_2 , where it is true in X_1 and false in X_2 .

E_2 and Σ_2 should be equal to E_1 and Σ_1 for any ρ that is neither in X_1 nor X_2 because there was no step done in this branch. However, for a ρ_t in X_1 they should be valued as e_1 and σ_1 and similarly for X_2 to create two different branches with (e_1, σ_1) and (e_2, σ_2) .

Assuming that E_1 and Σ_1 are measurable in \mathcal{P}_1 and the RIDs in \mathcal{P}_1 are disjoint of N_1 then this definition of $E_1, \Sigma_1, \mathcal{P}_1, N_1 \rightsquigarrow_{RV} E_2, \Sigma_2, \mathcal{P}_2, N_2$ has the following properties:

- E_2 and Σ_2 are measurable in \mathcal{P}_2
- We can rewrite with almost-sure equalities at the positions of $E_1, \Sigma_1, E_2, \Sigma_2$
- The RIDs in \mathcal{P}_2 are disjoint of N_2
- Two consecutive steps in different atoms of \mathcal{P}_1 can be swapped

From this we build our big step semantics:

BIG-STEP-REFL

$$\frac{\forall m \in \text{supp}(\mathcal{P}), \rho. m \prec \rho \rightarrow E(\rho) = V(\rho) \quad [\Sigma_1 = \Sigma_2]_{\mathcal{P}}}{E, \Sigma_1, \mathcal{P}, N \Downarrow V, \Sigma_2, \mathcal{P}, N}$$

BIG-STEP-STEP

$$\frac{E_1, \Sigma_1, \mathcal{P}_1, N_1 \rightsquigarrow_{RV} E_2, \Sigma_2, \mathcal{P}_2, N_2 \quad E_2, \Sigma_2, \mathcal{P}_2, N_2 \Downarrow V, \Sigma_3, \mathcal{P}_3, N_3}{E_1, \Sigma_1, \mathcal{P}_1, N_1 \Downarrow V, \Sigma_3, \mathcal{P}_3, N_3}$$

They are collecting semantics, collecting all computations from each branch of the computation tree into one evaluation. Each branch

of the tree is evaluated until it results into a value but the order in which the branches are evaluated is nondeterministic. The fact that a multi step only changes one branch, let's us easily decompose the computation tree into its subtrees. If we have given $E_1, \Sigma_1, \mathcal{P}_1, N_1 \Downarrow V, \Sigma_2, \mathcal{P}_2, N_2$ and some random variable $(X \sim \mu)_{\mathcal{P}_1}$, then we can find a big step evaluation starting from $E_1, \Sigma_1, \mathcal{P}_1|_{X=v}, N'$ for all $v \in \text{supp}(\mu)$. This fact is crucial for the proof of **C-WP-SWAP**.