


First Steps Towards Probabilistic Iris: Harmonizing Independence, Conditioning, and Dynamic Heap Allocation

Iris Workshop 2026

Janine Lohse, Tim Rohde, Jimmy Xin, Niklas Mück, Iona Kuhn
Deepak Garg, Derek Dreyer, Emanuele D'Oswaldo

June 9, 2026



Probabilistic Program Logics

Distributional Probabilistic Logics (DPLs)

Native **assertions over distribution**
of program output/state

PSL [Barthe et al. 2020]

Lilac [Li et al. 2023]

Bluebell [Bao et al. 2025]

pcOL [Zilberstein et al. 2026]

$\{True\}$ $flip()$ $\{V. V \sim Ber_{1/2}\}$

Specialized Probabilistic Logics (SPLs)

Traditional assertions over output/state
plus **probabilistic resources**

Eris [Aguirre et al. 2024]

Tachis [Haselwarter et al. 2024]

Clutch [Gregersen et al. 2024]

Coneris [Li et al. 2025]

$\{\frac{1}{2}(0.5)\}$ $flip()$ $\{v.v = 1\}$

Probabilistic Program Logics

Distributional Probabilistic Logics (DPLs)

Native **assertions over distribution**
of program output/state

Can support **wide range of properties**
including (conditional) independence

Specialized Probabilistic Logics (SPLs)

Traditional assertions over output/state
plus **probabilistic resources**

focus on **specific property** of interest
e.g. expected values, error bounds

often superior language feature support
+ fancy Iris reasoning techniques
+ strong modularity

DPLs have potential for being a **general-purpose foundation** for probabilistic logics,
but so far are **lagging behind** SPLs in e.g. supported language features

Independence as Separation [Barthe et al. 2020]

Example: one time pad encryption: M message, K key:

$$M \sim \text{Ber}_p * K \sim \text{Ber}_{1/2}$$

means that M and K are **probabilistically independent**

Separating Conjunction

$\phi * \psi :=$ can separate probability space into independent parts satisfying ϕ and ψ

One time pad encryption correctness:

$$M \sim \text{Ber}_p * K \sim \text{Ber}_{1/2} \vdash M \sim \text{Ber}_p * (K \text{ xor } M) \sim \text{Ber}_{1/2}$$

The Conditioning Modality [Li et al. 2023]

Conditioning (discrete setting): probabilistic case distinction $C_{\mu}v. \Phi(v)$

$$\begin{aligned} M \sim \text{Ber}_p * K \sim \text{Ber}_{1/2} &\vdash C_{\text{Ber}_p}v. \begin{cases} [M = 0] * [K \text{ xor } M = K] * K \sim \text{Ber}_{1/2} & \text{if } v = 0 \\ [M = 1] * [K \text{ xor } M = \neg K] * K \sim \text{Ber}_{1/2} & \text{else} \end{cases} \\ &\vdash C_{\text{Ber}_p}v. \begin{cases} [M = 0] * (K \text{ xor } M) \sim \text{Ber}_{1/2} & \text{if } v = 0 \\ [M = 1] * (K \text{ xor } M) \sim \text{Ber}_{1/2} & \text{else} \end{cases} \\ &\vdash (K \text{ xor } M) \sim \text{Ber}_{1/2} * C_{\text{Ber}_p}v. \begin{cases} [M = 0] & \text{if } v = 0 \\ [M = 1] & \text{else} \end{cases} \\ &\vdash (K \text{ xor } M) \sim \text{Ber}_{1/2} * M \sim \text{Ber}_p \end{aligned}$$

The Gap

All existing logics with support for independence reasoning
do not support dynamic memory allocation
...let alone more sophisticated Iris-style ghost state

The Challenge of Dynamic Allocation

Probabilistic Logics

Separation = independence

$$M \sim Ber_p * K \sim Ber_{1/2}$$

+?

Classic Separation Logic

Separation = heap disjointness

$$l_1 \mapsto v_1 * l_2 \mapsto v_2$$

Approach by PSL and pcOL: Separation means independence **and** heap disjointness:

Separating Conjunction in PSL/pcOL

$\phi * \psi \triangleq$ can separate probability space into two independent parts
over disjoint parts of the storage that satisfy ϕ and ψ

\Rightarrow in PSL/pcOL: $\phi * \psi$ requires heap disjointness **across all random outcomes**

The Challenge of Dynamic Allocation

in PSL/pcOL: $\phi * \psi$ requires heap disjointness across all random outcomes

Definition does not work with dynamic allocation!

Example

```
{ $\top$ } (if flip() then ref 0); (ref flip(), ref flip()) {( $L_1, L_2$ ).  $L_1 \rightsquigarrow Ber_{1/2} * L_2 \rightsquigarrow Ber_{1/2}$ }
```

Assume det. allocator that allocates $\ell_0, \ell_1, \ell_2, \dots$ in order

\Rightarrow First flip false $\rightarrow (\ell_0, \ell_1)$ returned, first flip true $\rightarrow (\ell_1, \ell_2)$ returned

$\Rightarrow \ell_1$ can belong to either L_1 or L_2 : no global disjointness!

\Rightarrow Triple is not valid with the PSL/pcOL model of separation!

Key Contributions

Amaryllis

First logic to support independence + **dynamic heap allocation**

- General construction for lifting ownership of Iris resource to probabilistic **ownership per random outcome**
- Instance: Novel model that requires heap disjointness only per random outcome instead over all of them
- Fully mechanized in Rocq, including Iris proof mode

Limitations: Only finite distributions \rightarrow only terminating programs.
No step indexing, invariants, concurrency

Amaryllis' Model of Per-Outcome Disjointness

Key Idea

Require heap disjointness only **per random outcome** instead over all of them

Amaryllis' Model explicitly tracks random choices (indexed valuations)

RID: set of *random choice identifiers*

Model

$$\mathbb{P}(\text{RID} \rightarrow \mathbb{B}) \times RV_{\text{Heap}} \quad \text{where } RV_T \triangleq (\text{RID} \rightarrow \mathbb{B}) \rightarrow T$$

$(\phi * \psi)(\mathcal{P}, R) \triangleq$ can separate \mathcal{P} into independent $\mathcal{P}_1, \mathcal{P}_2$
and R into two pointwise disjoint R_1, R_2
s.t. $\phi(\mathcal{P}_1, R_1)$ and $\psi(\mathcal{P}_2, R_2)$

Logic Assertions

Model

$$\mathbb{P}(\text{RID} \rightarrow \mathbb{B}) \times RV_{\text{Heap}} \quad \text{where } RV_T \triangleq (\text{RID} \rightarrow \mathbb{B}) \rightarrow T$$

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

$$(L \mapsto V)(\mathcal{P}, R) \triangleq \forall \rho \in \text{supp}(\mathcal{P}). (L(\rho) \mapsto V(\rho))(R(\rho))$$

$$L \rightsquigarrow \mu \triangleq \exists V. V \sim \mu \wedge L \mapsto V$$

Amaryllis Logic

Idea: Use random expression variable in the middle of triple!

	Classic Separation Logic	Amaryllis
Points-to	$l \mapsto v$ with $l : Loc, v : Val$	$L \mapsto V$ with $L : RV_{Loc}, V : RV_{Val}$
Allocation Rule	$\{\top\} \mathbf{ref} \ v \ \{l. \ l \mapsto v\}$	$\{\top\} \ \rho. \ \mathbf{ref} \ V(\rho) \ \{L. \ L \mapsto V\}$
Load	$\{l \mapsto v\} \ !l \ \{w. \ v = w\}$	$\{L \mapsto V\} \ \rho. \ !L(\rho) \ \{W. \ V \approx W\}$
Store	$\{l \mapsto v\} \ l \leftarrow w \ \{l \mapsto w\}$	$\{L \mapsto V\} \ \rho. \ L(\rho) \leftarrow W(\rho) \ \{L \mapsto W\}$
Sample	-	$\{\top\} \ \mathbf{flip}(\rho) \ \{V. \ V \sim Ber_\rho\}$

Example Verification in Amaryllis

Example

$\{\top\}$ (if flip() then ref 0); (ref flip(), ref flip()) $\{(L_1, L_2). L_1 \rightsquigarrow Ber_{1/2} * L_2 \rightsquigarrow Ber_{1/2}\}$

Proof idea:

Condition on outcome of first flip and reason **reason case-by-case** using C-lift!

$$\begin{array}{c} \text{C-LIFT} \\ \forall v. \{\Phi(v)\} E \{\Psi(v)\} \\ \hline \{C_\mu v. \Phi(v)\} E \{C_\mu v. \Psi(v)\} \end{array}$$

First, show:

$$\{\top\} (\text{ref flip}(), \text{ref flip}()) \{\exists L_1, L_2. V \approx (L_1, L_2) * L_1 \rightsquigarrow \text{Ber}_{1/2} * L_2 \rightsquigarrow \text{Ber}_{1/2}\}$$

Then:

$$\begin{array}{ll} \{\top\} (\text{if flip}() \text{ then ref } 0); \dots & \{V. \exists L_1, L_2. V \approx (L_1, L_2) * \dots\} \dashv \\ \{V_f \sim \text{Ber}_{1/2}\} (\text{if } V_f \text{ then ref } 0); \dots & \{V. \exists L_1, L_2. V \approx (L_1, L_2) * \dots\} \dashv \\ \{C_{\text{Ber}_{1/2}} v. V_f \approx v\} (\text{if } V_f \text{ then ref } 0); \dots & \{V. C_{\text{Ber}_{1/2}} v. \exists L_1, L_2. V \approx (L_1, L_2) * \dots\} \dashv \\ \{V_f \approx v\} (\text{if } V_f \text{ then ref } 0); \dots & \{\exists L_1, L_2. V \approx (L_1, L_2) * \dots\} \dashv \\ \{\top\} (\text{if } v \text{ then ref } 0); \dots & \{\exists L_1, L_2. V \approx (L_1, L_2) * \dots\} \end{array}$$

Do case analysis on v :

$$\begin{array}{ll} \{\top\} (\text{if true then ref } 0); \dots & \{\exists L_1, L_2. V \approx (L_1, L_2) * \dots\} \dashv \\ \{\top\} \text{ref } 0; (\text{ref flip}(), \text{ref flip}()) & \{\exists L_1, L_2. V \approx (L_1, L_2) * \dots\} \dashv \\ \{L \mapsto 0\} (\text{ref flip}(), \text{ref flip}()) & \{\exists L_1, L_2. V \approx (L_1, L_2) * \dots\} \end{array}$$

Example Verification in Amaryllis

Example

$\{\top\}$ (if flip() then ref 0); (ref flip(), ref flip()) $\{(L_1, L_2). L_1 \rightsquigarrow Ber_{1/2} * L_2 \rightsquigarrow Ber_{1/2}\}$

Key Proof Step

In Amaryllis Model, the postcondition is **convex**:

$$C_{Ber_{1/2}} \text{ v. } \exists L_1, L_2. V \approx (L_1, L_2) * \dots \vdash \exists L_1, L_2. V \approx (L_1, L_2) * \dots$$

Convexity in general not trivial, also challenging on paper

To ease this: Amaryllis' proof mode automates away std cases like this via typeclasses

Summary

- Amaryllis combines independence, conditioning, and dynamic allocation in one logic
- Key idea to combine separation=independence and separation=heap disjointness:
Require only **per-outcome disjointness** of the heap
- In Amaryllis' Hoare triples $\{\Phi\} E \{V. \Psi(V)\}$, E and V are **random variables**
- Amaryllis' model is parametric in resource algebra, not specific to heaps

What else is in the paper?

- Generalize Iris' update modality to DPL setting
 - ▶ Novel notion: *probabilistic frame-preserving update* that preserves conditionings
 - ▶ General concept for achieving modular reasoning in DPLs
- First to apply Iris methodology to a DPL:
 - ▶ Re-interpret *authoritative resource algebra* for DPL setting
 - ▶ Define *base logic* (including conditioning modality) then derive *weakest precondition modality* and *program logic* on top