



Contextual Refinement of Higher-Order Concurrent Probabilistic Programs

KWING HEI LI, Aarhus University, Denmark

ALEJANDRO AGUIRRE, Aarhus University, Denmark

JOSEPH TASSAROTTI*, New York University, USA

LARS BIRKEDAL, Aarhus University, Denmark

We present Foxtrot, the first higher-order separation logic for proving contextual refinement of higher-order concurrent probabilistic programs with higher-order local state. From a high level, Foxtrot inherits various concurrency reasoning principles from standard concurrent separation logic, e.g. invariants and ghost resources, and supports advanced probabilistic reasoning principles for reasoning about complex probability distributions induced by concurrent threads, e.g. tape presampling and induction by error amplification. The integration of these strong reasoning principles is highly *non-trivial* due to the combination of probability and concurrency in the language and the complexity of the Foxtrot model; the soundness of the logic relies on a version of the axiom of choice within the Iris logic, which is not used in earlier work on Iris-based logics. We demonstrate the expressiveness of Foxtrot on a wide range of examples, including the adversarial von Neumann coin and the `randombytes_uniform` function of the Sodium cryptography software library.

All results have been mechanized in the Rocq proof assistant and the Iris separation logic framework.

CCS Concepts: • **Theory of computation** → **Separation logic**; **Logic and verification**; **Probabilistic computation**; **Concurrency**; **Program verification**; • **Mathematics of computing** → **Probabilistic algorithms**.

Additional Key Words and Phrases: contextual equivalence, error credits, logical relations

ACM Reference Format:

Kwing Hei Li, Alejandro Aguirre, Joseph Tassarotti, and Lars Birkedal. 2026. Contextual Refinement of Higher-Order Concurrent Probabilistic Programs. *Proc. ACM Program. Lang.* 10, PLDI, Article 187 (June 2026), 25 pages. <https://doi.org/10.1145/3808265>

1 Introduction

Randomization and concurrency are, independently, two important and widely-used language features in the world of computing. The former is used to implement probabilistic data structures, randomized algorithms, and machine learning algorithms; the latter is utilized to increase program throughput and to model distributed systems and multi-core architecture. Together, the combination of these two features can be found in concurrent randomized algorithms [5, 32] and in cryptography protocols [15]. However, formal methods for reasoning about concurrent probabilistic programs are limited and not very expressive, in part because concurrency introduces nondeterminism, which in semantic models does not compose straightforwardly with probabilistic choice [43].

*Also affiliated with Amazon Web Services. This paper does not reflect the views of Amazon Web Services.

Authors' Contact Information: [Kwing Hei Li](mailto:kwling@cs.au.dk), Aarhus University, Aarhus, Denmark, kwling@cs.au.dk; [Alejandro Aguirre](mailto:alejandro@cs.au.dk), Aarhus University, Aarhus, Denmark, alejandro@cs.au.dk; [Joseph Tassarotti](mailto:jt4767@nyu.edu), New York University, New York, USA, jt4767@nyu.edu; [Lars Birkedal](mailto:birkedal@cs.au.dk), Aarhus University, Aarhus, Denmark, birkedal@cs.au.dk.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

ACM 2475-1421/2026/6-ART187

<https://doi.org/10.1145/3808265>

1.1 Motivating Example and Prior Work

Consider the following `flmpl` function, where `rand N` samples an integer uniformly from $\{0, \dots, N\}$:

$$\text{flmpl} \triangleq \lambda _ . \text{let } (x, y) = (\text{rand } 7 \parallel \parallel \text{rand } 31) \text{ in } x \cdot 32 + y$$

Though `flmpl` appears simple, it already exhibits the combination of both probability and concurrency. When we call `flmpl`, it spawns two concurrent threads, which generate a random 3-bit and 5-bit number, respectively. Afterwards, it returns the binary concatenation of the two numbers generated. Note that since the random samples of both threads are independent, this function behaves the same as one that directly samples a random 8-bit number. Hence to reason about its correctness, one might want to show that its behavior is the same as the following `fSpec` function.

$$\text{fSpec} \triangleq \lambda _ . \text{rand } 255$$

To formally argue that the two programs behave the same, we might want to prove that `flmpl` is *contextually equivalent* to `fSpec`. Intuitively, an expression e_1 is contextually equivalent to another expression e_2 at a type τ if no well-typed contexts C of ground type can distinguish them, denoted by $\emptyset \vDash e_1 \simeq_{\text{ctx}} e_2 : \tau$. In other words, the behavior of a client program remains unchanged if we replace any occurrence of the sub-program e_1 with e_2 . Contextual equivalence is defined as the symmetric interior of contextual refinement, denoted by $\emptyset \vDash e_1 \lesssim_{\text{ctx}} e_2 : \tau$, which is roughly defined to mean that, for any context C the observable behaviour of $C[e_1]$ is *included* in the observable behaviour of $C[e_2]$. Proving contextual equivalence or refinement is generally a difficult problem as the statement quantifies over *all* possible contexts, including those that might use concurrency and probability as well. Although there are other techniques for proving contextual equivalence or refinement of concurrent programs [17] or probabilistic programs [7, 45], respectively, to the best of our knowledge, there is *no* prior work on verifying contextual equivalence of higher-order programs utilizing *both* probability and concurrency. (See Table 1 for an overview of a selection of prior work for proving contextual refinement of less expressive languages and §6 for a more detailed summary. Table 1 also compares these techniques on whether they are implemented as a modular program logic, support reasoning about programs with local state, and are mechanized in a proof assistant.)

Table 1. Systems for proving contextual refinement (Logic = implemented as a program logic, ND = only nondeterminism).

	Probability	Concurrency	Logic	Local State	Mechanized
Prob. λ -calculi [9, 11, 34]	●	○	○	○	○
Wand et al. [45]	●	○	○	○	○
Bizjak and Birkedal [7]	●	○	○	●	○
ReLoC [17]	○	●	●	●	●
Aguirre and Birkedal [3]	●	○ (ND)	○	○	○
Clutch [18]	●	○	●	●	●
Approxis [19]	●	○	●	●	●
Foxtrot (this paper)	●	●	●	●	●

1.2 Contributions and Key Ideas

In this paper, we present *Foxtrot*, a higher-order probabilistic concurrent separation logic, which we use to build the *first* logical relations model for proving contextual refinement and equivalence of concurrent probabilistic programs. At a high level, the *Foxtrot* logic supports reasoning about contextual refinement of higher-order probabilistic concurrent imperative programs written in

ConcRandML, an ML-style programming language with higher-order dynamically-allocated local state, discrete probabilistic sampling, and dynamically allocated concurrently running threads.

Since ConcRandML is probabilistic, the observable behaviour of a program refers to the *probability* of observing some behaviour, such as termination. Hence, by contextual refinement $\emptyset \models e_1 \lesssim_{\text{ctx}} e_2 : \tau$, we mean that the (limit of the) probability of termination of $C[e_1]$ is below the (limit of the) probability of termination of $C[e_2]$. The precise definition can be found in §2.4; here it suffices to know that the operational semantics is defined using (probabilistic stateful) schedulers, which at every step of the execution decide which thread to execute next, and that the termination probability is obtained by taking the supremum over all possible schedulers (following the notion of *may-termination* in previous work [3, 6]).

The central new notion in Foxtrot to reason about contextual refinement is a *logical concurrent probabilistic refinement judgment*. The refinement judgment $\Delta \models e \lesssim e' : \tau$ intuitively means that e refines e' at the type τ . The context Δ assigns a relational interpretation to all type variables in τ , but can be ignored for now. The logical refinement provides a sound method to prove contextual refinement, i.e. $\emptyset \models e_1 \lesssim e_2 : \tau$ implies $\emptyset \models e_1 \lesssim_{\text{ctx}} e_2 : \tau$. Following the logical approach to relational refinement [41], the refinement judgment is defined in the Foxtrot logic, which is in turn defined on top of the Iris base logic [21]. The Foxtrot program logic is phrased using Hoare triples $\{P\} e \{Q\}$, which are defined (using ghost state) with an eye for relational refinement, much as in earlier work [18, 19]. Rules in Foxtrot are often expressed in terms of Hoare triples, so to prove a refinement judgment, we often unfold its definition to the Hoare triple layer. We later see how the refinement judgment is defined with Hoare triples in §3.1.

Roughly speaking, Foxtrot inherits proof rules for concurrency (including *invariants* and *ghost resources*) from earlier relational logics for concurrency [17], and *coupling-based proof rules* for probabilistic reasoning from earlier relational logics for probability [18, 19]. In fact, these basic rules are sufficient to verify a wide range of concurrent probabilistic programs, especially those where the concurrency and probability features do not interact with each other in complex ways. However, some programs feature complex behaviors arising from the complicated interaction between probability and concurrency, e.g. the `flmpl` function from §1 where one has to reason about the joint probability distribution produced by sampling in two separate threads. To address these issues, Foxtrot includes more advanced probabilistic reasoning principles, including asynchronous couplings using so-called *presampling tapes* (first introduced in the Clutch logic [18]) and *fragmented couplings* for reasoning about rejection samplers [19]. Moreover, inspired by Eris [4] and Approxis [19], Foxtrot also supports induction by error amplification through the use of *error credit* resources. We explain these reasoning principles in §3, and we demonstrate their expressiveness in the concurrent setting by using them to reason about a range of challenging program refinements in §4, e.g. an adversarial variant of the von Neumann coin program and the implementation of the `randombytes_uniform` function in the open-source cryptography software library Sodium. These examples (even including the smaller examples in §3) are, to the best of our knowledge, beyond the scope of previous techniques, as they require complex reasoning principles, e.g. reasoning about higher-order functions, local state, and execution of unknown code.

In light of this rough overview, the reader may wonder if Foxtrot is a trivial combination of earlier logics and/or what, if anything, is challenging in its development? The key technical novelty and challenges lie in the *definition* of the program logic, specifically in the weakest precondition predicate used in the definition of Hoare triples, and the construction of the model of this logic, which ultimately allows us to prove the soundness theorem above. One of the key challenges is that intuitively, to show that e_1 refines e_2 , for every possible scheduler for the left-hand side program, we have to exhibit a scheduler for the right-hand side program e_2 , so that e_1 can be coupled with e_2 , and this is done by “piecing together” schedulers for subexpressions (to validate compositional program

logic rules). This is significantly different from earlier work on relational logics for concurrency without probabilities [17, 42], where one only has to construct a single *trace* for the right-hand side. Now, with probabilistic features, the challenge is that we have to compose *families* of schedulers (because of the probabilistic branching we get a family of schedulers). To do so, we end up having to use a version of the *axiom of choice* in the Iris logic which, to the best of our knowledge, has not been used in earlier work on Iris-based program logics. Indeed, Foxtrot’s model is significantly more sophisticated than that of its predecessors, such as Clutch [18] and Approxis [19]. Luckily, users of the logic need not understand the model; it suffices to know that we have used the model to prove adequacy of the logic (Theorem 3.1), which intuitively states that proving a Hoare triple in the Foxtrot logic implies that the termination probability of a program is upper-bounded by that of another, and which is used to prove the soundness theorem above (Theorem 3.3).

We believe it is a *strong feature* of Foxtrot that, with the new definition of weakest precondition and model, we are in fact able to prove soundness of the many strong reasoning principles mentioned above. We note that this is not automatic (all the proof rules of Foxtrot have been proved sound from scratch) and not trivial; in fact, as we later emphasize in §3, one of the proof rules that one might expect to hold given earlier work [18, 19], namely “presampling on the right”, is *not* sound in the concurrent setting of Foxtrot.

Contributions. In summary, we provide:

- The *first* higher-order separation logic, Foxtrot, for reasoning about contextual refinement of concurrent probabilistic programs.
- A rich set of coupling logic rules for advanced probabilistic reasoning, proven sound using the novel model of Foxtrot.
- A collection of case studies showcasing the expressiveness of Foxtrot, which are beyond the scope of all previous techniques.
- All results in the paper are mechanized in the Rocq proof assistant [38], using the Iris separation logic framework [20, 21, 23, 24] and Coquelicot real analysis library [8].

Outline. We begin by presenting the preliminaries for the rest of the paper (§2). We then detail the features of Foxtrot and demonstrate the use of these features through small examples (§3). Next, we show how Foxtrot can scale to verify larger and more involved case studies (§4). Afterwards, we explain how the semantic model of Foxtrot is constructed and proven sound (§5). We lastly conclude with a discussion of related and future work (§6, §7).

2 Preliminaries

In this section, we first recall useful facts from probability theory in §2.1. We then describe the syntax and semantics of our programming language ConcRandML in §2.2 and §2.3. Lastly in §2.4, we present the definition of contextual refinement we use in the concurrent probabilistic setting.

2.1 Probability Theory

A *subdistribution* (henceforth simply *distribution*) on a countable set A is a function $\mu : A \rightarrow [0, 1]$ s.t. $\sum_{a \in A} \mu(a) \leq 1$. We write the collection of distributions on A as $\mathcal{D}(A)$. Distributions admit monadic operations, with ($\text{ret} : A \rightarrow \mathcal{D}(A)$) defined as $\text{ret}(a)(a') \triangleq \text{if } a = a' \text{ then } 1 \text{ else } 0$ and ($\text{bind} : (A \rightarrow \mathcal{D}(B)) \times \mathcal{D}(A) \rightarrow \mathcal{D}(B)$) defined as $\text{bind}(f, \mu)(b) \triangleq \sum_{a \in A} \mu(a) \cdot f(a)(b)$. We often use the notation $\mu \gg f$ for $\text{bind}(f, \mu)$.

Some examples of distributions include the *null distribution* $\mathbf{0} : \mathcal{D}(A)$, given by $\lambda x.0$, and the uniform distribution $\mathcal{U}N : \mathcal{D}(\mathbb{N})$, which returns $1/(N+1)$ for every $n \in \{0, \dots, N\}$ and 0 otherwise.

The *expected value* of $X : A \rightarrow [0, 1]$ w.r.t. μ is defined as $\mathbb{E}_\mu[X] \triangleq \sum_{a \in A} \mu(a) \cdot X(a)$. The *mass* of μ is defined as $|\mu| \triangleq \mathbb{E}_\mu[\lambda x.1]$.

Foxtrot uses *approximate* probabilistic coupling [19, 35] to relate pairs of probabilistic programs.

DEFINITION 2.1 (APPROXIMATE COUPLING). *Let $\mu_1 \in \mathcal{D}(A)$ and $\mu_2 \in \mathcal{D}(B)$. Given $\varepsilon \in [0, 1]$ and a relation $R \subseteq A \times B$, we say that there exists an (ε, R) -coupling of μ_1 and μ_2 if for all $[0, 1]$ -valued random variables $X : A \rightarrow [0, 1]$ and $Y : B \rightarrow [0, 1]$, s.t. $(a, b) \in R$ implies $X(a) \leq Y(b)$, we have the inequality $\mathbb{E}_{\mu_1}[X] \leq \mathbb{E}_{\mu_2}[Y] + \varepsilon$. We write $\mu_1 \lesssim_\varepsilon \mu_2 : R$ if such an (ε, R) -coupling exists.*

2.2 The ConcRandML Language

Our examples are written in the ConcRandML language, a higher-order ML-style programming language with higher-order local state, extended with discrete probabilistic sampling and fork-based unstructured concurrency. The syntax of the language is defined by the grammar below:

$$\begin{aligned}
v, w \in \text{Val} &::= z \in \mathbb{Z} \mid b \in \mathbb{B} \mid () \mid \ell \in \text{Loc} \mid \text{rec } f \ x = e \mid (v, w) \mid \text{inl } v \mid \text{inr } v \mid \\
e \in \text{Expr} &::= v \mid x \mid e_1 \ e_2 \mid e_1 + e_2 \mid e_1 - e_2 \mid \dots \mid \text{if } e \ \text{then } e_1 \ \text{else } e_2 \mid (e_1, e_2) \mid \text{fst } e \mid \dots \\
&\quad \text{ref } e \mid !e \mid e_1 \leftarrow e_2 \mid \text{rand } e \mid \text{fork } e \mid \text{faa } e_1 \ e_2 \mid \dots \\
\text{Heap} &::= \text{Loc} \xrightarrow{\text{fin}} \text{Val} & \text{Tapes} &::= \text{Label} \xrightarrow{\text{fin}} \text{Tape} \\
\sigma \in \text{State} &::= \text{Heap} \times \text{Tapes} & \rho \in \text{Cfg} &::= \text{List}(\text{Expr}) \times \text{State} \\
\tau \in \text{Type} &::= \alpha \mid \text{unit} \mid \text{bool} \mid \text{nat} \mid \text{int} \mid \tau \times \tau \mid \tau + \tau \mid \tau \rightarrow \tau \mid \forall \alpha. \tau \mid \exists \alpha. \tau \mid \mu \alpha. \tau \mid \text{ref } \tau
\end{aligned}$$

The syntax is mostly standard, e.g. `ref e`, `!e`, and `e1 ← e2` respectively allocate, read from, and write onto a reference. The expression `rand N` samples from \mathcal{UN} , and `fork e` spawns a new thread with `e`. ConcRandML provides various *atomic* operations for synchronizing threads, e.g. `fetch-and-add faa e1 e2` adds the integer `e2` to the value `v` stored at location `e1` and returns `v`.

The state σ is given by a pair of maps. The first map models the heap as a finite map from locations to values. The second map is used for keeping track of *presampling tapes*, the discussion of which is postponed until §3.3. A program configuration $\rho \in \text{List}(\text{Expr}) \times \text{State}$ is given by a pair containing the list of currently executing threads and the state. We say a configuration is *final* if the first expression in the thread list is a value.

To define contextual refinement of ConcRandML programs (§2.4), we use its typing system, which follows standard ML-style rules. Typing is expressed as $\Theta \mid \Gamma \vdash e : \tau$ where Γ is a context assigning types to program variables, and Θ contains type variables that may occur in Γ and τ .

2.3 Operational Semantics

The operational semantics of ConcRandML programs is defined in 3 steps: (1) a step function for stepping a single expression, (2) a stepping function for a configuration w.r.t. a scheduler, and (3) full program execution w.r.t. a scheduler¹.

Expressions. We consider a mostly standard call-by-value semantics. The step: $(\text{Expr}, \text{State}) \rightarrow \mathcal{D}(\text{Expr}, \text{State}, \text{List}(\text{Expr}))$ function takes an expression and a state, and returns a distribution over a new expression, state, and a (possibly empty) list of newly spawned threads. We display some reductions below; in particular `rand N` steps to an integer between 0 and `N` uniformly at random.

$$\begin{aligned}
\text{step}(\text{if true then } e_1 \ \text{else } e_2, \sigma) &= \text{ret}(e_1, \sigma, []) \\
\text{step}(\text{fork } e, \sigma) &= \text{ret}(() , \sigma, [e]) \\
\text{step}(\text{rand } N, \sigma) &= \lambda (n, \sigma, []). \frac{1}{N+1} \text{ if } n \in \{0, \dots, N\} \text{ and } 0 \text{ otherwise}
\end{aligned}$$

¹ConcRandML is the same language studied in Coneris [26] and we refer readers to that paper for more details.

Thread Pools and Schedulers. The tpStep function takes a configuration $\rho = (\vec{e}, \sigma)$ and an index j and returns a distribution over new configurations after stepping the j -th thread:

$$\text{tpStep}(\vec{e}, \sigma)(j) \triangleq \begin{cases} \mathbf{0} & \text{if } (\vec{e}, \sigma) \text{ is final,} \\ \text{ret}(\vec{e}, \sigma) & \text{if } e_j \in \text{Val or } j \geq |\vec{e}|, \\ \text{step}(e_j, \sigma) \gg \lambda (e'_j, \sigma', \vec{e}'). \text{ ret}(\vec{e}[j \mapsto e'_j] \# \vec{e}', \sigma') & \text{otherwise.} \end{cases}$$

If ρ is final, we return immediately. If e_j is a value or does not exist, we take a stutter step. Otherwise, we step the j -th thread and add the spawned threads to the thread pool.

The choice of which thread to step is determined by a *scheduler*. A (probabilistic, stateful) scheduler is defined by a transition function $\zeta : (\text{SchedSt} \times \text{Cfg}) \rightarrow \mathcal{D}(\text{SchedSt} \times \mathbb{N})$, which takes in an internal state $\Xi \in \text{SchedSt}$ and a configuration ρ , and returns a distribution on its new internal state and the index of the thread to step next. Intuitively, the scheduler is used to resolve all non-determinism in a program (which is here limited to which thread is to be executed at each step), and its internal state can be used to fully store the history of the program execution and all the previous non-deterministic results. Here, we consider probabilistic schedulers as it is a fairly standard model in literature (see [47]) and it better reflects some real-world schedulers that make use of randomness in an essential way, e.g. work-stealing schedulers.

Given a configuration ρ , a scheduler ζ , and a scheduler state Ξ , we can now define the single scheduler-step reduction function $\text{schStep}_\zeta(\Xi, \rho) \in \mathcal{D}(\text{SchedSt} \times \text{Cfg})$ as follows:

$$\text{schStep}_\zeta(\Xi, \rho) \triangleq \zeta(\Xi, \rho) \gg \lambda (\Xi', j). \text{tpStep}(\rho, j) \gg \lambda \rho'. \text{ret}(\Xi', \rho')$$

Informally, $\text{schStep}_\zeta(\Xi, \rho)$ first evaluates $\zeta(\Xi, \rho)$ to get a new state Ξ' and index j , steps the j -th thread to obtain the new configuration ρ' , and returns the new scheduler state and configuration.

Full Program Execution. We now define n -step program execution w.r.t. a scheduler ζ with the recursive function $\text{exec}_{\zeta, n} : (\text{SchedSt} \times \text{Cfg}) \rightarrow \mathcal{D}(\text{Val})$ defined below. Intuitively, $\text{exec}_{\zeta, n}(\Xi, \rho)(v)$ represents the probability of returning a value v in the first thread after at most n steps of ρ w.r.t. the scheduler ζ with starting scheduler state Ξ . The full program execution is then defined as the limit of $\text{exec}_{\zeta, n}$, which exists by monotonicity and continuity, i.e. $\text{exec}_\zeta(\Xi, \rho) \triangleq \lim_{n \rightarrow \infty} \text{exec}_{\zeta, n}(\Xi, \rho)$.

$$\text{exec}_{\zeta, n}(\Xi, \rho) \triangleq \begin{cases} \text{ret } v & \text{if } \rho = (v \cdot \vec{e}, \sigma) \text{ for some } v \in \text{Val}, \\ \mathbf{0} & \text{if } n = 0 \text{ and } \rho \text{ is not final,} \\ \text{schStep}_\zeta(\Xi, \rho) \gg \text{exec}_{\zeta, n-1} & \text{otherwise.} \end{cases}$$

The probability that a starting configuration ρ terminates under scheduler ζ with starting scheduler state Ξ is hence defined as $\text{exec}_\zeta^\downarrow(\Xi, \rho) \triangleq |\text{exec}_\zeta(\Xi, \rho)|$.

The upper termination probability of an arbitrary configuration ρ is taken as the supremum of the termination probability ranging over all possible schedulers and initial scheduler states, i.e. $\text{exec}^{\sqcup\downarrow}(\rho) \triangleq \sup_{\zeta, \Xi} \text{exec}_\zeta^\downarrow(\Xi, \rho)$. We use the upper termination probability $\text{exec}^{\sqcup\downarrow}$ as the observation predicate for contextual refinement. In other words, we consider *may-termination* [3, 6] as the main property of interest for the refinement². Defining observational equivalence in terms of (non-)termination is common for non-probabilistic sequential languages, where it is provably equivalent to alternate definitions stated in terms of returning the same Boolean value [33].

²We leave the problem of *must-termination* refinement of concurrent probabilistic programs for future work; we believe an Iris-based logic would appear to require transfinite step-indexing [3, 36].

We here also provide an auxiliary definition of partial program execution $\text{pexec}_{\zeta, n} : (\text{SchedSt} \times \text{Cfg}) \rightarrow \mathcal{D}(\text{SchedSt} \times \text{Cfg})$, used in the model of Foxtrot.

$$\text{pexec}_{\zeta, n}(\Xi, \rho) = \begin{cases} \text{ret}(\Xi, \rho) & \text{if } \rho \text{ is final or } n = 0, \\ \text{schStep}_{\zeta}(\Xi, \rho) \gg \text{pexec}_{\zeta, n-1} & \text{otherwise.} \end{cases}$$

We can view pexec as a relaxation of exec which keeps probability mass on configurations that are not final, whereas the latter only considers final configurations.

2.4 Contextual Refinement

Since we are in a typed setting, we consider only typed contexts. A program context is well-typed, written $C : (\Theta \mid \Gamma \vdash \tau) \Rightarrow (\Theta' \mid \Gamma' \vdash \tau')$, if for any term e where $\Theta \mid \Gamma \vdash e : \tau$ we have $\Theta' \mid \Gamma' \vdash C[e] : \tau'$. We say expression e_1 *contextually refines* expression e_2 if for all well-typed program contexts C resulting in a closed program, the upper termination probability of $C[e_1]$ is bounded by the upper termination probability of $C[e_2]$:

$$\begin{aligned} \Theta \mid \Gamma \vDash e_1 \lesssim_{\text{ctx}} e_2 : \tau \triangleq \forall \tau', (C : (\Theta \mid \Gamma \vdash \tau) \Rightarrow (\emptyset \mid \emptyset \vdash \tau')), \sigma. \\ \text{exec}^{\sqcup\downarrow}([C[e_1]], \sigma) \leq \text{exec}^{\sqcup\downarrow}([C[e_2]], \sigma) \end{aligned}$$

Note that the statement itself is in the meta-logic (e.g., Rocq) and makes no mention of Foxtrot or Iris. Contextual equivalence $\Theta \mid \Gamma \vDash e_1 \simeq_{\text{ctx}} e_2 : \tau$ is defined as a refinement in both directions, i.e. $(\Theta \mid \Gamma \vDash e_1 \lesssim_{\text{ctx}} e_2 : \tau) \wedge (\Theta \mid \Gamma \vDash e_2 \lesssim_{\text{ctx}} e_1 : \tau)$. We simply write $e_1 \lesssim_{\text{ctx}} e_2$ or $e_1 \simeq_{\text{ctx}} e_2$ if the type τ can be inferred and that Γ and Θ are empty.

3 Logic

Now we introduce the rules of the Foxtrot logic. We first present a small set of its propositions:

$$\begin{aligned} P, Q \in iProp ::= \text{True} \mid \text{False} \mid P \wedge Q \mid P \vee Q \mid P \Rightarrow Q \mid \forall x. P \mid \exists x. P \mid \\ P * Q \mid P \multimap Q \mid \triangleright P \mid \boxed{P}^I \mid \overline{g}^Y \mid \varepsilon_1 \Rightarrow_{\varepsilon_2} P \mid j \Rightarrow e \mid \ell \mapsto v \mid \ell \mapsto_s v \mid \\ \{P\} e \{Q\}_{\varepsilon} \mid \overset{\approx}{\rightsquigarrow} P \mid \kappa \hookrightarrow (N, \vec{n}) \mid \kappa \hookrightarrow_s (N, \vec{n}) \mid \sharp(\varepsilon) \mid \dots \end{aligned}$$

Foxtrot is built on top of the Iris base logic [23] and it inherits many of the basic propositions found in standard Iris separation logics, such as the *later* modality \triangleright , *invariants* \boxed{I}^I , *ghost resources* \overline{g}^Y , and *fancy updates* $\varepsilon_1 \Rightarrow_{\varepsilon_2} P$. The specification program on the right-hand side is represented by the specification thread connective $j \Rightarrow e$, which intuitively states that the j -th right-hand side thread is currently executing expression e . We use the points-to connectives $\ell \mapsto v$ and $\ell \mapsto_s v$ to represent exclusive ownership of the location ℓ storing value v in the left-hand side program and right-hand side program, respectively.

Central to Foxtrot's logic is the Hoare triple proposition. These Hoare triples are annotated with a mask ε for tracking opened invariants (we omit it if it is the top mask \top). We give meaning to the Foxtrot Hoare triple via the adequacy theorem [Theorem 3.1](#):

THEOREM 3.1 (ADEQUACY). *If we can prove $\{0 \Rightarrow e'\} e \{v. \exists v'. 0 \Rightarrow v'\}$ in Foxtrot, then for all states σ, σ' , we have $\text{exec}^{\sqcup\downarrow}([e], \sigma) \leq \text{exec}^{\sqcup\downarrow}([e'], \sigma')$.*

Roughly speaking, the theorem states that if we are able to prove the Hoare triple for the expression e assuming that the 0-th thread on the right stores e' initially, and we can match up the execution traces of e and e' such that whenever we have a trace where e terminates, the corresponding trace for e' also terminates, then the upper termination probability of e is upper-bounded by that of e' .

Foxtrot also supports other propositions for advanced probabilistic reasoning such as the *probabilistic update modality* $\bowtie P$, *presampling-tape* connectives $\kappa \hookrightarrow (N, \vec{n})$ and $\kappa \hookrightarrow_s (N, \vec{n})$, as well as *error credits* $f(\varepsilon)$. We describe these in more detail later in this section.

In §3.1, we first present the logical relation that gives rise to the refinement judgment introduced in §1.2, as well as the key results of the refinement judgment. We then present an overview of the rules of Foxtrot, from basic program-logic rules in §3.2 to more advanced rules for more complicated couplings in §3.3. We emphasize that though various rules and features are inspired by prior work, as mentioned in §1.2, proving all the rules are sound in our novel Foxtrot model is highly *non-trivial* and is *not automatic*, which we explain more in detail in §5.

3.1 Logical Relations

In §1.2, we introduced the logical concurrent probabilistic refinement judgment for proving contextual refinement of ConcRandML programs. In Foxtrot, this judgment $\Delta \vDash e \lesssim e' : \tau$ is defined as $\forall K, j. \{j \Rightarrow K[e']\} e \{v. \exists v'. j \Rightarrow K[v'] * \llbracket \tau \rrbracket_\Delta(v, v')\}$, where K quantifies over all evaluation contexts. Here $\llbracket \tau \rrbracket_\Delta$ represents the semantic interpretation of type τ defined inductively on the type, which intuitively describes when two values of a certain type are considered related. The way we define the logical relations and encode it within the logic is standard and similar to that in previous work [41]. We then extend the refinement judgment to open terms in the standard way, i.e. $\Delta \mid \Gamma \vDash e \lesssim e' : \tau \triangleq \forall \gamma, \gamma'. \llbracket \Gamma \rrbracket_\Delta(\gamma, \gamma') \multimap \Delta \vDash e\gamma \lesssim e'\gamma' : \tau$, and prove the fundamental lemma of the binary logical relations:

LEMMA 3.2 (FUNDAMENTAL LEMMA OF BINARY LOGICAL RELATIONS). *If $\Theta \mid \Gamma \vdash e : \tau$ then for all Δ assigning a relational interpretation to every type variable in Θ , we have $\Delta \mid \Gamma \vDash e \lesssim e : \tau$.*

We then show that the refinement relation is sound with respect to contextual refinement:

THEOREM 3.3 (SOUNDNESS). *Let Θ be a type variable context, and assume that, for all Δ assigning a relational interpretation to all type variables in Θ , we can derive $\Delta \mid \Gamma \vDash e_1 \lesssim e_2 : \tau$. Then $\Theta \mid \Gamma \vDash e_1 \lesssim_{ctx} e_2 : \tau$.*

Here, note that the type system of ConcRandML is considerably rich, supporting complex features such as polymorphism and higher-order references. Therefore, later in the paper, even though the statement of our examples might not directly involve these features, the contexts that interact with them might. Hence, to prove contextual refinement, our logic Foxtrot has to be, and is, expressive enough to support higher-order propositions and define our logical relations (which utilizes second-order impredicative quantification).

In addition to the binary refinement relations, we also additionally construct a *unary* logical relations model for proving type safety of ConcRandML programs where $\Delta \vDash e : \tau \triangleq \{\text{True}\} e \{v. \llbracket \tau \rrbracket_\Delta(v)\}$. This is then similarly extended to open terms: $\Delta \mid \Gamma \vDash e : \tau \triangleq \forall \gamma. \llbracket \Gamma \rrbracket_\Delta(\gamma) \multimap \Delta \vDash e\gamma : \tau$ and we prove its expected fundamental lemma, which is used later in the adversarial von Neumann coin example in §4.1 to reason about unknown adversaries.

LEMMA 3.4 (FUNDAMENTAL LEMMA OF UNARY LOGICAL RELATIONS). *If $\Theta \mid \Gamma \vdash e : \tau$ then for all Δ assigning a unary interpretation to all type variables in Θ , we have $\Delta \mid \Gamma \vDash e : \tau$.*

3.2 Basic Rules

In this subsection, we cover structural rules, rules for symbolically executing the left or right-hand side programs, and the standard rule for coupling programs.

Standard Iris Rules. Foxtrot inherits many standard rules in similar Iris separation logics, such as **HT-BIND** for reasoning about sequences of instructions and the frame rule **HT-FRAME** for

framing resources around a Hoare triple. We symbolically execute the left-hand side program with familiar-looking rules such as that of **HT-LOAD** and **HT-FORK**. As mentioned in §1.2, Foxtrot is a concurrent separation logic and it supports invariants (with rules such as **HT-INV-ALLOC** and **HT-INV-OPEN**) and ghost resources (standard rules omitted here for brevity).

$$\begin{array}{c}
\text{HT-BIND} \\
\frac{\{P\} e \{v.Q\} \quad \forall v. \{Q\} K[v] \{R\}}{\{P\} K[e] \{R\}} \\
\\
\text{HT-FORK} \\
\frac{\{\text{True}\} e \{\text{True}\}}{\{\text{True}\} \text{fork } e \{w.w = ()\}} \\
\\
\text{HT-FRAME} \\
\frac{\{P\} e \{Q\}}{\{P * R\} e \{Q * R\}} \\
\\
\text{HT-INV-ALLOC} \\
\frac{\{[P]^t * Q\} e \{R\}}{\{P * Q\} e \{R\}} \\
\\
\text{HT-LOAD} \\
\frac{\{l \mapsto v\} !l \{w.w = v * l \mapsto v\}}{\{[I]^t * P\} e \{Q\}_{\mathcal{E}\cup\{t\}}} \\
\\
\text{HT-INV-OPEN} \\
\frac{e \text{ atomic} \quad \{\triangleright I * P\} e \{\triangleright I * Q\}_{\mathcal{E}}}{\{[I]^t * P\} e \{Q\}_{\mathcal{E}\cup\{t\}}}
\end{array}$$

We implement the parallel composition operator with the **fork** primitive, and with ghost resources, we prove the familiar-looking rule **HT-PAR-COMP** for parallel composition:

$$\frac{\{P_1\} e_1 \{v_1. Q_1 v_1\} \quad \{P_2\} e_2 \{v_2. Q_2 v_2\}}{\{P_1 * P_2\} e_1 \parallel e_2 \{(v_1, v_2). Q_1 v_1 * Q_2 v_2\}} \text{HT-PAR-COMP}$$

Probabilistic Update Modality. We previously saw that we can symbolically execute the left-hand side program with the Hoare triple directly. Naturally, one might ask, how can we do so for the right? The answer is that we do so with the *probabilistic update modality* \rightsquigarrow , a construct first introduced in Coneri [26] for reasoning about randomized logical atomicity. Here, we extend the modality for manipulating the right-hand side program; we also later see how the probabilistic update modality is used in more complicated rules, such as those for presampling onto tapes and for generating error credits.

The probabilistic update modality supports monadic-like rules, i.e. **PUPD-RET** and **PUPD-BIND**. Most importantly, the probabilistic update modality can be eliminated in the precondition of a Hoare triple (**HT-PUPD-ELIM**). As a consequence, rules for symbolically executing the right-hand side program can be succinctly expressed with the probabilistic update modality, which is different in presentation from earlier relational logics [17, 41]. For example, **PUPD-LOAD** and **PUPD-FORK** symbolically steps a load and **fork** expression on the right-hand side program, respectively.

$$\begin{array}{c}
\frac{P}{\rightsquigarrow P} \text{PUPD-RET} \quad \frac{\rightsquigarrow P \quad P * \rightsquigarrow Q}{\rightsquigarrow Q} \text{PUPD-BIND} \quad \frac{\{P * Q\} e \{R\}}{\{(\rightsquigarrow P) * Q\} e \{R\}} \text{HT-PUPD-ELIM} \\
\\
\frac{l \mapsto_s v \quad j \Rightarrow K[!l]}{\rightsquigarrow (l \mapsto_s v * j \Rightarrow K[v])} \text{PUPD-LOAD} \quad \frac{j \Rightarrow K[\text{fork } e]}{\rightsquigarrow (j \Rightarrow K[()]) * \exists j'. j' \Rightarrow e} \text{PUPD-FORK}
\end{array}$$

In addition, Foxtrot provides **PUPD-PAR-COMP** for parallel composition on the specification side:

$$\frac{j \Rightarrow K[e_1 \parallel e_2]}{\rightsquigarrow \exists j_1, j_2, K_1, K_2. \left(j_1 \Rightarrow K_1[e_1] * j_2 \Rightarrow K_2[e_2] * \forall v_1, v_2. j_1 \Rightarrow K_1[v_1] * j_2 \Rightarrow K_2[v_2] * \rightsquigarrow j \Rightarrow K[(v_1, v_2)] \right)} \text{PUPD-PAR-COMP}$$

Couplings. Foxtrot makes use of (approximate) *probabilistic couplings* [29, 39] to relate the randomness of two programs. We present the following simple yet useful coupling rule **HT-COUPLE**, which allows us to reason as if the two sampled values on both programs are related by a bijective function f under the same sampling space $\{0, \dots, N\}$.

$$\frac{f \text{ bijection} \quad \forall n \leq N. \{j \Rightarrow K[f(n)]\} n \{\Phi\}}{\{j \Rightarrow K[\text{rand } N]\} \text{rand } N \{\Phi\}} \text{HT-COUPLE}$$

Together with the other logical facilities for reasoning about concurrency, e.g. invariants, the coupling rule **HT-COUPLE** is sufficient for reasoning about various concurrent probabilistic programs. As an example, consider the left-hand side entropy-mixer program in [Figure 1](#), originally studied in Probabilistic Concurrent Outcome Logic [47]. Here, we are given two sources of entropy, x_1 and x_2 , the former being unreliable as it can be adversarially controlled by the scheduler, while x_2 is of high quality. Nonetheless, we are still able to “combine” these two sources and produce a high-quality output that behaves like a `rand 1`. We can verify this claim by showing contextual equivalence between the entropy-mixer program and a `rand 1` program with Foxtrot. Even though this program uses both concurrency (parallel composition) and random sampling (`rand`), the bulk of the proof simply boils down to applying the soundness theorem ([Theorem 3.3](#)), defining an invariant to share the reference y between the concurrent threads, and applying **HT-COUPLE** with the right bijective function f .

$$\begin{array}{l} \text{let } (y, r) = (\text{ref } 0, \text{ref } 0) \text{ in} \\ \left(\begin{array}{l} \text{let } x_1 = !y \text{ in} \\ \text{let } x_2 = \text{rand } 1 \text{ in} \\ r \leftarrow ((x_1 + x_2) \bmod 2) \end{array} \parallel \parallel y \leftarrow 1 \right); \quad \approx_{\text{ctx}} \quad \text{rand } 1 \\ !r \end{array}$$

Fig. 1. Entropy mixer example.

We summarize this subsection with the following slogan:

Slogan 1: Invariants, ghost resources, and **HT-COUPLE are (almost) all you need.**

3.3 Advanced Rules

Previously, we saw the standard coupling rule **HT-COUPLE** for coupling a `rand N` on the left with one on the right under a bijective function. However, unsurprisingly, there are times where it is not the case that there is conveniently a single `rand` of the same range on both programs. In those cases, Foxtrot uses other advanced logical facilities and rules to establish the couplings. In the following examples, we also move away from reasoning about programs of base types (e.g. `nat`), to reasoning about programs of function types (e.g. $A \rightarrow B$), which requires nested Hoare triples.

Combining Samplings in Separate Threads. Foxtrot supports rules for reasoning about programs that perform random samplings across different threads. Consider the batch sampling program `batchImpl` in [Figure 2](#), which is a generalization of `flmpl` from [§1](#). It computes `rand N` and `rand M` concurrently and returns the sum of the result with the first value multiplied by $(M + 1)$. In other words, it aggregates two parallelized random samples to produce another one. We can show that this function returns a number that is uniformly distributed by showing its contextual equivalence to the function `batchSpec` that directly samples from `rand((N + 1) · (M + 1) - 1)`.

$$\begin{array}{l} \text{batchImpl} \triangleq \\ \lambda _ . \text{let } (x, y) = (\text{rand } N \parallel \parallel \text{rand } M) \text{ in } x \cdot (M + 1) + y \quad \approx_{\text{ctx}} \quad \text{batchSpec} \triangleq \\ \lambda _ . \text{rand}((N + 1) \cdot (M + 1) - 1) \end{array}$$

Fig. 2. Batch-sampling example.

Let us consider the left-to-right refinement first. To reason about random samplings occurring in different threads in `batchImpl`, Foxtrot uses presampling tapes (first introduced in Clutch [18]) to “compress” multiple samplings and reason about them as if they are one.

$$\begin{aligned}
v \in \text{Val} ::= \dots \mid \kappa \in \text{Label} & & \sigma \in \text{State} \triangleq (\text{Loc} \xrightarrow{\text{fin}} \text{Val}) \times (\text{Label} \xrightarrow{\text{fin}} \text{Tape}) \\
e \in \text{Expr} ::= \dots \mid \text{tape } e \mid \text{rand } e_1 e_2 & & t \in \text{Tape} \triangleq \{(N, \vec{n}) \mid N \in \mathbb{N} \wedge \vec{n} \in \mathbb{N}_{\leq N}^*\}
\end{aligned}$$

As mentioned previously in §2.2, the state of our language contains a finite map from tape labels to presampling tapes. We also introduce two new constructs in ConcRandML, `tape e` and `rand $e_1 e_2$` . The expression `tape N` allocates an empty tape and returns a new tape label (1). In addition to the regular unlabelled random sampling we have shown before, we can perform a *labelled* random sampling by supplying a tape label to the `rand` operation. If the tape label points to an empty tape, then the labelled random sampling instruction acts like a normal unlabelled one (2), *i.e.* it samples a number uniformly from the range. However, if it is not empty, we pop and return the first number in the tape *deterministically* (3). Perhaps surprisingly, there are *no* constructs in ConcRandML for presampling elements onto tapes, and so, tapes remain empty throughout the execution of the program (and hence, (3) is never applied normally). Instead, the tape presampling action is done only within the logic during the proof.

$$\text{step}(\text{tape } N, \sigma) = \text{ret}(\kappa, \sigma[\kappa := (N, \epsilon)], []) \quad (\text{where } \kappa \text{ is fresh w.r.t. } \sigma) \quad (1)$$

$$\text{step}(\text{rand } \kappa N, \sigma) = \lambda (n, \sigma, []) . \text{if } 0 \leq n \leq N \text{ then } \frac{1}{N+1} \text{ else } 0 \quad (\text{where } \sigma[\kappa] = (N, \epsilon)) \quad (2)$$

$$\text{step}(\text{rand } \kappa N, \sigma) = \text{ret}(n, \sigma[\kappa := \vec{n}], []) \quad (\text{where } \sigma[\kappa] = (N, n \cdot \vec{n})) \quad (3)$$

Just like we have the heap points-to connectives on both sides of the refinement, we have the tapes connectives for both sides as well: $\kappa \hookrightarrow (N, \vec{n})$ and $\kappa \hookrightarrow_s (N, \vec{n})$. The rules **HT-ALLOC-TAPE** and **PUPD-ALLOC-TAPE** allocates a new empty tape, and **HT-RAND-TAPE** returns the first element of a tape deterministically for a labelled `rand` on the left-hand side program. There is however *no* counterpart rule of **HT-RAND-TAPE** for the right-hand side program, as presampling onto tapes on the specification side is unsound³.

$$\begin{array}{c}
\text{HT-ALLOC-TAPE} \\
\hline
\{\text{True}\} \text{tape } N \{\kappa. \kappa \hookrightarrow (N, \epsilon)\} \\
\text{HT-RAND-TAPE} \\
\hline
\{\kappa \hookrightarrow (N, n \cdot \vec{n})\} \text{rand } \kappa N \{x. x = n * \kappa \hookrightarrow (N, \vec{n})\}
\end{array}
\qquad
\begin{array}{c}
\text{PUPD-ALLOC-TAPE} \\
j \Vdash K[\text{tape } N] \\
\hline
\approx\approx\approx \exists \kappa. j \Vdash K[\kappa] * \kappa \hookrightarrow_s (N, \epsilon)
\end{array}$$

We emphasize that presampling tapes are only used for verification purposes and they do not change the behavior of programs (recall that a labelled `rand` with an empty tape behaves the same as an unlabelled one). To relate a program annotated with tapes and one without, we provide the two rules **HT-COUPLE-RAND-LBL** and **PUPD-COUPLE-LBL-RAND** to couple a pair of labelled and unlabelled `rand` operations. The rule **HT-COUPLE-RAND-LBL** couples an unlabelled `rand` on the left and a labelled one on the right with the empty tape resource on the specification side provided. On the other hand, the rule **PUPD-COUPLE-LBL-RAND** couples a presampling action of a tape resource on the implementation side and a `rand` construct on the right; after which, applying **HT-RAND-TAPE** later reads out the presampled value in the tape resource.

$$\begin{array}{c}
\text{HT-COUPLE-RAND-LBL} \\
f \text{ bijection} \\
\forall n \leq N. \{j \Vdash K[f(n)] * \kappa \hookrightarrow_s (N, \epsilon)\} n \{\Phi\} \\
\hline
\{j \Vdash K[\text{rand } \kappa N] * \kappa \hookrightarrow_s (N, \epsilon)\} \text{rand } N \{\Phi\}
\end{array}
\qquad
\begin{array}{c}
\text{PUPD-COUPLE-LBL-RAND} \\
f \text{ bijection } j \Vdash K[\text{rand } N] \quad \kappa \hookrightarrow (N, \vec{n}) \\
\hline
\approx\approx\approx \exists n \leq N. \kappa \hookrightarrow (N, \vec{n} \cdot [n]) * j \Vdash K[f(n)]
\end{array}$$

³Even though presampling onto tapes on the specification side is unsound, the (empty) specification-side tape connective is still needed for **HT-COUPLE-RAND-LBL**.

Returning to our example, we prove the refinement via the following intermediate program `batchImpl'` annotated with tapes.

$$\begin{aligned} \text{batchImpl}' &\triangleq \lambda _ . \text{let } (\kappa, \kappa') = (\text{tape } N, \text{tape } M) \text{ in} \\ &\quad \text{let } (x, y) = (\text{rand } \kappa \ N \ ||| \ \text{rand } \kappa' \ M) \text{ in} \\ &\quad x \cdot (M + 1) + y \end{aligned}$$

We split the refinement into two parts by transitivity of contextual refinement.

In the first part, we show that `batchImpl` refines the intermediate program `batchImpl'`. This is relatively straightforward as the structure of the two programs are similar; After allocating the specification tapes with `PUPD-ALLOC-TAPE`, we apply `HT-COUPLE-RAND-LBL` to couple the unlabelled `rand`s on the left with labelled ones on the right.

Next, we show that the intermediate program `batchImpl'` refines `batchSpec`. We first allocate the tapes in `batchImpl'` with `HT-ALLOC-TAPE`. We can take advantage of the two tape resources and perform a coupling, where we couple two presampling operations on the left tapes with an unlabelled `rand` on the right with the following rule `PUPD-COUPLE-TWO-RANDS-RAND`, choosing $f(n, m) \triangleq n \cdot (M + 1) + m$. Subsequently, each thread in `batchImpl'` deterministically reads from their corresponding tape resource with `HT-RAND-TAPE`.

$$\frac{\text{PUPD-COUPLE-TWO-RANDS-RAND} \quad j \Rightarrow K[\text{rand}((N + 1) \cdot (M + 1) - 1)] \quad \kappa \hookrightarrow (N, \vec{n}) \quad \kappa' \hookrightarrow (M, \vec{m}) \quad f \text{ bijection}}{\rightsquigarrow \exists n \leq N, m \leq M. j \Rightarrow K[f(n, m)] * \kappa \hookrightarrow (N, \vec{n} \cdot [n]) * \kappa' \hookrightarrow (M, \vec{m} \cdot [m])}$$

At a high level, the key step of this proof is that we take advantage of presampling tapes to perform multiple (possibly future) random samplings in advance, and to couple this operation with a single sampling on the right.

Now that we proved the left-to-right refinement, one might ask, can we adopt a similar proof structure with tape presampling for the right-to-left refinement? Unfortunately the answer is *no*! It turns out that, unlike in previous logics supporting tapes [18, 19], it is *unsound* to perform any presampling on the right-hand side program in Foxtrot, and we provide a proof of this in the extended version of this paper [27]. Intuitively, presampling is an operation that enables us to predict future random outcomes from `rand`. Allowing arbitrary presampling operations on the right commutes probabilistic and non-deterministic choices in a program, and as a result, the scheduler constructed by the model is one that can make scheduling decisions based on future probabilistic outcomes, which is not valid according to the operational semantics of `ConcRandML`. In previous logics that support presampling on the right, e.g. Clutch [18] and Approxis [19], this paradox does not occur since the underlying language is sequential, and the operational semantics of programs does not involve any scheduler to resolve non-deterministic choices (e.g. thread scheduling).

That said, we can still show that `batchSpec` refines `batchImpl` without presampling on the right-hand side program. To see why, consider symbolically stepping through `batchImpl` now residing on the right-hand side of our refinement. By applying `PUPD-PAR-COMP`, we then create two thread resources via the parallel composition operator. We now have access to both `rand` operations as specification resources at the same time, and we can couple the single `rand` operation in `batchSpec` with the two `rand` operations residing in two *separate* threads on the right with the following rule `HT-COUPLE-RAND-TWO-RANDS`.

$$\frac{\text{HT-COUPLE-RAND-TWO-RANDS} \quad f \text{ bijection} \quad \forall n \leq N, m \leq M. \{j \Rightarrow K[n] * j' \Rightarrow K'[m]\} f(n, m) \{\Phi\}}{\{j \Rightarrow K[\text{rand } N] * j' \Rightarrow K'[\text{rand } M]\} \text{rand}((N + 1) \cdot (M + 1) - 1) \{\Phi\}}$$

We note that **PUPD-COUPLE-TWO-RANDS-RAND** and **HT-COUPLE-RAND-TWO-RANDS** can be generalized in a straightforward manner to allow couplings between multiple random samplings on one side with one on the other. This enables us to show contextual refinement between parallel programs that utilize more than two threads.

Let us take a step back and summarize the key ideas of this example. The main challenge of this batch sampling example is that of relating two **rand** operations with one. Because there is always one left-hand side expression in a Hoare triple, in order for us to reason about multiple **rand** operations in different threads on the left-hand side program, we utilize presampling tapes and presampling coupling rules to perform the coupling in advance before the actual samplings on the left occur. When the multiple **rand** operations occur in different threads on the right, we can directly make use of the multiple thread resources produced by the concurrent nature of the program to perform the coupling. We summarize these ideas with the following slogan.

Slogan 2: Tape the left; parallelize the right.

Rejection Sampling. We now turn our attention to another difficult-to-verify pattern in probabilistic computation: rejection sampling schemes. Rejection samplers are a class of Las Vegas algorithms; they generate random samples from a target distribution by repeatedly sampling from another distribution and only accepting a sample if it satisfies certain conditions. Though this rejection sampling technique is ubiquitous, it is also challenging to reason about, particularly due to its use of unbounded looping in the implementation.

Consider the rejection sampler program on the left-hand side of the refinement in **Figure 3**. Given $M < N$, the rejection sampler program is a function that first samples from $\{0, \dots, N\}$. We then return the value if it is upper bounded by M , otherwise we repeat the process. It should not be surprising that this function returns an integer from $\{0, \dots, M\}$ uniformly, and hence one would want to show that it is contextually equivalent to one that calls **rand** M . We emphasize that this result is *novel* and is stronger than those found in previous work [19], as we show that this contextual equivalence holds even in the concurrent setting.

$$\begin{array}{l} \text{rec } f _ = \text{let } x = \text{rand } N \text{ in} \\ \text{if } x \leq M \text{ then } x \text{ else } f _ () \end{array} \quad \simeq_{\text{ctx}} \quad \lambda _ . \text{rand } M$$

Fig. 3. Rejection sampler example.

Here, we sadly cannot use the trick presented in **Slogan 2**; instead of trying to relate two **rand**s with one, we are trying to relate a pair of **rand**s in a somewhat unnatural manner. If the rejection sampler samples a value $\leq M$, we want that value to be coupled with **rand** M . Otherwise if we sampled a value $> M$, we do not want to step **rand** M at all. In other words, we only want to perform a coupling with the last sampled value in the rejection sampler that gets accepted with **rand** M .

Foxtrot solves this issue by supporting *fragmented couplings*, a form of couplings first introduced in Approxis [19]⁴. To see fragmented couplings in action, let us first consider the left-to-right direction of the refinement. After stepping through the program where we reach **rand** N on the left and **rand** M on the right, we apply **HT-COUPLE-FRAGMENTED** to couple the **rand** expressions in a

⁴The presentation of the fragmented coupling rules in Foxtrot is however slightly different from those in Approxis. This is because the Approxis fragmented coupling rules are expressed with presampling on the right-hand side program, which is unsound in Foxtrot as shown in the extended version of this paper [27].

fragmented way, taking f as the identity function.

$$\frac{\text{HT-COUPLE-FRAGMENTED} \quad M \leq N \quad f \text{ injective} \quad \forall n \leq N. \left\{ \begin{array}{l} (\exists m \leq M. f(m) = n * j \Rightarrow K[m]) \vee \\ \neg(\exists m \leq M. f(m) = n) * j \Rightarrow K[\text{rand } M] \end{array} \right\} n \{\Phi\}}{\{j \Rightarrow K[\text{rand } M]\} \text{ rand } N \{\Phi\}}$$

Intuitively, **HT-COUPLE-FRAGMENTED** performs a case split on the value of $\text{rand } N$. If it is accepted, then we also take a step for $\text{rand } M$ on the right, and their values are related by some injective function f (the identity function for this example). Otherwise, we do not take a step on the right. From here, we continue stepping through the proof after a case split. If we sample an accepted value, we can establish the postcondition directly. In the other case, we step the left-hand side program and return to the start of the loop, at which point, we can apply the hypothesis of Löb induction, a standard rule for reasoning about recursive functions in Iris logics [21].

Again, one might ask, can we do the same for the other right-to-left direction of the refinement? Alas, reality is often disappointing and the answer is *no*. This is due to two reasons. The first reason is more technical; we can perform a fragmentation step on the right-hand side program with **HT-COUPLE-FRAGMENTED** where we can choose to step the program or not, but we cannot do something similar for the left program due to how the Foxtrot model is defined. Secondly, we cannot apply Löb induction for this direction; in the case that the right-hand side program samples a rejected value, we can only proceed by stepping the right-hand side program back to the start of the recursive function with the left-hand side program kept unchanged. We hence cannot eliminate the later modality in the hypothesis since we do not take any steps on the left program.

To overcome these two limitations, we utilize presampling tapes which we saw previously and *error credits*, a separation logic resource first introduced in Eris [4]. Error credits are written with the $\$ (\epsilon)$ assertion, where ϵ is a non-negative real. Intuitively, ownership of $\$ (\epsilon)$ error credits means we are attempting to prove a goal allowing some error of at most ϵ probability. In Foxtrot, error credits are mainly used for a proof technique called *induction by error amplification*, specifically in examples where it is easier to prove that the two programs are equivalent by showing that the distance between them is bounded by errors ϵ for all $\epsilon > 0$, like this rejection sampler.

We now review the basic rules of error credits. The rule **ERR-SPLIT** states that we can split error credits by addition. We can derive False with $\$ (1)$ error credits, which holds because any proposition, including False, holds with probability at least 0 (**ERR-1**). We can generate some arbitrarily-small positive error credits from thin air with **PUPD-ERR**. Finally, the error amplification rule **IND-ERR-AMP** allows us to prove any proposition if we are able to amplify some positive error credit by an amplification factor $k > 1$.

$$\begin{array}{ccc} \text{ERR-SPLIT} & \text{ERR-1} & \text{PUPD-ERR} \\ \frac{\$ (\epsilon_1 + \epsilon_2)}{\$ (\epsilon_1) * \$ (\epsilon_2)} & \frac{\$ (1)}{\text{False}} & \frac{}{\rightsquigarrow (\exists \epsilon. 0 < \epsilon * \$ (\epsilon))} \end{array} \quad \begin{array}{c} \text{IND-ERR-AMP} \\ \frac{\epsilon > 0 \quad k > 1 \quad \$ (\epsilon)}{(\$ (k \cdot \epsilon) \multimap P) \multimap \$ (\epsilon) \multimap P} \\ P \end{array}$$

The right-to-left direction of the refinement is established via an intermediate program, similar to the right-to-left refinement of the batch sampling example. In particular our intermediate program is one that matches the ideal $\text{rand } M$ program but is modified to use tapes, i.e. $\lambda _ . \text{rand} (\text{tape } M) M$. It is easy to show that the original $\text{rand } M$ program refines this intermediate program with **HT-COUPLE-RAND-LBL**. Now, for the refinement between the intermediate program and the rejection sampler program, we first apply **PUPD-ERR** to generate some arbitrarily-small error credit. Then we apply **IND-ERR-AMP**, choosing $k \triangleq \frac{N+1}{N-M} > 1$. Subsequently we apply **PUPD-COUPLE-FRAGMENTED**, a rule similar to **HT-COUPLE-FRAGMENTED** except that it couples a tape presampling operation on the

left and a `rand N` on the right, and it amplifies some error credit by the amplification factor k in the case that we sample a rejected value on the right (we take f to be the identity function).

$$\frac{\text{PUPD-COUPLE-FRAGMENTED} \quad M < N \quad j \Rightarrow K[\text{rand } N] \quad \kappa \hookrightarrow (M, \vec{n}) \quad \not\prec(\varepsilon) \quad f \text{ injective}}{\approx\approx \exists n \leq N. j \Rightarrow K[n] * \left(\begin{array}{l} (\exists m \leq M. f(m) = n * \kappa \hookrightarrow (M, \vec{n} \cdot [n])) \vee \\ \neg(\exists m \leq M. f(m) = n) * \kappa \hookrightarrow (M, \vec{n}) * \not\prec\left(\frac{N+1}{N-M} \cdot \varepsilon\right) \end{array} \right)}$$

The proof then proceeds in a straightforward way after a case split on the result of `rand N`. If it is accepted, we directly read out the sampled value in the tape on the left-hand side with `HT-RAND-TAPE`. Otherwise, we step the right-hand side program back to the start of the recursive function, and apply the induction hypothesis of `IND-ERR-AMP`, where we returned to where we started while successfully amplifying the error by k .

To summarize, for programs exhibiting a rejection sampling pattern, we use fragmented couplings to couple the rejection sampling with the simplified ideal one and apply induction by error amplification to show approximate equivalence of the programs for any positive error ε . This brings us to our third slogan below that captures the key ideas of proving rejection samplers in Foxtrot.

Slogan 3: Fragmented couplings and Error Amplification for Rejections, a.k.a. FEAR.

4 Case Studies

In this section, we present more complex case studies that demonstrate the use of the advanced features of Foxtrot featured in §3 and highlight many subtleties in reasoning about concurrent probabilistic programs. Results in this section are out of scope of previous techniques since they involve a combination of concurrency, probability, and higher-order functions.

4.1 Adversarial von Neumann Coin

John von Neumann [44] showed that a biased coin can be used to simulate the result of flipping a fair coin. This is done by repeatedly tossing the coin twice until the results of a pair of tosses are different, at which point we return the result of the first toss. Zilberstein et al. [47] showed that this result is also true if the bias is stored in shared memory and can be modified as long as it satisfies some invariant, and only fetched right before a pair of tosses. Informally, this means that the program still behaves like a fair coin even when composed with *any* other terminating threads. We formally show that this claim holds by showing contextual equivalence between the following von Neumann coin program and a simple program that samples from `rand 1` (see Figure 4).

$$\lambda \mathcal{A}. \text{let } l = \text{ref } 0 \text{ in fork}(\mathcal{A} \ l); \quad \left(\begin{array}{l} \text{rec } f \ _ = \text{let } b = \min (!l) \ N \text{ in} \\ \text{let } x = \text{rand}(N+1) \leq b \text{ in} \\ \text{let } y = \text{rand}(N+1) \leq b \text{ in} \\ \text{if } x = y \text{ then } f \ () \text{ else } x \end{array} \right) \approx_{\text{ctx}} \lambda \mathcal{A} _ . \text{flip}$$

Fig. 4. Adversarial von Neumann coin example.

Both programs in Figure 4 first take in an adversary argument \mathcal{A} of type `ref nat → unit` and return a closure. Specifically for the adversarial coin program, it creates a reference l containing 0 and forks a thread that calls \mathcal{A} with it. This process allows \mathcal{A} to modify the reference to any natural number in any way, which might involve forking more threads or modifying the reference according to some probabilistic result. The closure returned by the adversarial coin program is a

recursive function that sets b to be the minimum of N and the value read from l . This represents setting the bias of the coin to be $(b + 1)/(N + 1)$, and we flip the coin twice, and return the first result if they differ. If they are the same, we call the closure again.

Most of this proof follows similarly to that of the rejection sampler example seen before (**Slogan 3** in §3.3) albeit a more complicated fragmented coupling rule. In this example, we also have to reason about unknown adversarial code \mathcal{A} , which we do so with the *unary* logical relations introduced in §3.1. When proving the left-to-right refinement, we apply **Lemma 3.4** to symbolically execute $\mathcal{A} l$ in the forked thread on the left. Note that we *cannot* apply the regular binary logical relations to step $\mathcal{A} l$ since we cannot couple the adversarial code with anything on the right. The right-to-left direction is even simpler; after forking $\mathcal{A} l$, the angelic nondeterminism of the right-hand side program enables us to directly discard the newly-forked thread resource and focus only on the expression in the main thread.

4.2 Sodium

Sodium [12] is a secure cross-platform cross-language open-source cryptography software library. We show the correctness of the `randombytes_uniform` function in the library for generating uniform samples (see **Figure 5**). Our result is particularly strong because it proves that the function behaves as intended even in contexts that might involve *concurrency*, thus showing that the function is *thread-safe* even though its implementation uses a rejection sampling technique and that the random sampling is not atomic (though it appears logically atomic to an outside observer).

The Sodium function first takes in an argument N . We modify our programs slightly by adding in an additional check that N is smaller than $\text{MAX} \triangleq 2^{32}$, which is not needed in the original C implementation since the argument is an unsigned integer of 32 bits. We then perform a simple check on whether N is smaller than 2 and we simply return 0 if that is the case. Otherwise, we take min to be $\text{MAX} \bmod N$ and we repeatedly sample a random number of 32 bits until it is larger or equal to min ⁵, and we eventually return the remainder of the number divided by N . We show that this somewhat convoluted implementation of a rejection sampler is contextually equivalent to a function that samples from $\{0, \dots, N - 1\}$ directly.

$$\begin{array}{l} \lambda N. \text{ if } \text{MAX} \leq N \text{ then } 0 \\ \quad \text{else if } N < 2 \text{ then } 0 \\ \quad \text{else let } \text{min} = \text{MAX} \bmod N \text{ in} \\ \quad \text{let } r = \text{ref } 0 \text{ in} \\ \quad \left(\begin{array}{l} \text{rec } f _ = r \leftarrow \text{rand}(\text{MAX} - 1); \\ \text{if } !r < \text{min} \text{ then } f () \text{ else } (!r \bmod N) \end{array} \right) () \end{array} \quad \simeq_{\text{ctx}} \quad \begin{array}{l} \lambda N. \text{ if } (\text{MAX} \leq N \mid\mid N = 0) \\ \quad \text{then } 0 \\ \quad \text{else } \text{rand}(N - 1) \end{array}$$

Fig. 5. Sodium `randombytes_uniform` implementation.

The main bulk of the proof can be split into two stages and it follows very similarly to the arguments we have seen in the batch-sampling and rejection sampler example in §3.3. In the first stage, we show that the implementation of the Sodium function is contextually equivalent to one that samples an integer n from `rand` ($\text{MAX} - (\text{MAX} \bmod N) - 1$) and returns $n \bmod N$. This is done by applying the fragmented coupling rules captured in **Slogan 3** for reasoning about rejection sampling. Then, in the second stage, we show that this intermediate program p_1 is contextually equivalent to the idealized function that samples from `rand`($N - 1$) directly. This equivalence is

⁵In the actual implementation, this sampling of a random number of 32 bits is done by reading from the `/dev/random` file, which we model directly with a `rand`.

established via another intermediate program p_2 that concurrently samples from $\text{rand}(N - 1)$ and $\text{rand}(\text{MAX}/N - 1)$ with two threads, but only returns the result of the first thread. For the equivalence between p_1 and p_2 , we apply the advanced coupling rules captured in **Slogan 2** to couple the results of two threads with one. Finally, we show that p_2 is contextually equivalent to the idealized program directly with the regular coupling rule **HT-COUPLE**.

4.3 Other Case Studies

For reasons of space, other case studies demonstrating the flexibility of our approach can be found in the extended version of this paper [27]. For example, similar to a result in Aguirre and Birkedal [3], we show that there exist programs which are contextually equivalent to one another even though there are no optimal schedulers that witness the equivalence. We also present a concurrent one-time pad example that is beyond the scope of all previous techniques, where presampling tapes are utilized to linearize concurrent random samplings. Finally, we study the equations of the algebraic theory induced by the contextual equivalence relation in Foxtrot.

5 Semantic Model and Soundness

In this section, we describe the semantic model of Foxtrot, defined in the Iris base logic [21], and the key ideas of its soundness proof.

5.1 Full-Information Schedulers

As described in §2.3, since ConcRandML is concurrent, we resolve the non-determinism of scheduling the threads through a (stateful probabilistic) scheduler. However, the definition of a scheduler in §2.3 is not easy to work with directly when constructing the model of Foxtrot. Here, we introduce an alternative scheduler called a *full-information scheduler* that is only used internally for defining the logic, which we simply write as FIsch(es). We write FIsch to denote the type of FIsches.

Basic Definition of FIsches. From a high-level perspective, FIsches are different from normal schedulers in three aspects. Firstly, the internal state of FIsches are fixed to be $\text{FISchSt} \triangleq \text{List}(\text{Cfg}' \times \mathbb{N})$, where $\text{Cfg}' \triangleq \text{List}(\text{Expr}) \times \text{Heap}$. We implicitly coerce configurations ρ of type Cfg to that of type Cfg' by removing the tape map (Tapes)⁶. Given two FIsch states $\Xi, \Xi' \in \text{FISchSt}$, we write $\Xi \leq \Xi'$ if Ξ is a prefix of Ξ' .

Secondly, the type of FIsches is different from a normal scheduler. Specifically, FIsches are defined by a transition function Υ of the type $(\text{FISchSt} \times \text{Cfg}) \rightarrow \text{option}(\mathcal{D}(\mathbb{N}))$.

Lastly, all FIsches Υ satisfy a *consistency* condition, where for all FIsch states Ξ and configurations ρ , if $\Upsilon(\Xi, \rho) = \text{None}$, then for all other configurations ρ' we have $\Upsilon(\Xi, \rho') = \text{None}$.

Semantics of FIsches. We modify the operational semantics of ConcRandML with respect to FIsches, such that the thread-stepping function FltpStep of FIsches is more “eager”, i.e. it attempts to step the corresponding thread regardless of whether the configuration is final or not:

$$\text{FltpStep}(\vec{e}, \sigma)(j) \triangleq \begin{cases} \text{ret}(\vec{e}, \sigma) & \text{if } e_j \in \text{Val} \text{ or } j \geq |\vec{e}|, \\ \text{step}(e_j, \sigma) \gg \lambda (e'_j, \sigma', \vec{e}'). \text{ret}(\vec{e}[j \mapsto e'_j] \# \vec{e}', \sigma') & \text{otherwise.} \end{cases}$$

Given a configuration ρ , a FIsch Υ , and a FIsch state Ξ , we can now define the single FIsch-step reduction function $\text{FIschStep}_\Upsilon(\Xi, \rho) \in \mathcal{D}(\text{SchedSt} \times \text{Cfg})$ as follows:

$$\text{FIschStep}_\Upsilon(\Xi, \rho) \triangleq \begin{cases} \mu \gg \lambda j. \text{FltpStep}(\rho, j) \gg \lambda \rho'. \text{ret}(\Xi \cdot [(\rho, j)], \rho') & \Upsilon(\Xi, \rho) = \text{Some } \mu \\ \text{ret}(\Xi, \rho) & \text{otherwise} \end{cases}$$

⁶We remove the tape map to ensure FIsches cannot schedule threads based on the content of the tapes.

Here, in addition to stepping the configuration according to the FIsch Υ , FIschStep also updates the FIsch's internal state by extending it with the configuration ρ and the thread j it attempts to step. Intuitively, this captures the idea that FIsches are keeping track of all the information it has access to since the beginning of the execution, hence the term "full-information".

The n -step program execution of a FIsch Υ returns both its internal state and the configuration if the FIsch reaches a None state:

$$\text{Flexec}_{\Upsilon,n}(\Xi, \rho) = \begin{cases} \text{ret}(\Xi, \rho) & \Upsilon(\Xi, \rho) = \text{None} \\ \mathbf{0} & n = 0 \\ \text{FIschStep}_{\Upsilon}(\Xi, \rho) \gg \text{Flexec}_{\Upsilon,n-1} & \text{otherwise.} \end{cases}$$

In contrast to the n -step program execution of a normal scheduler, $\text{Flexec}_{\Upsilon,n}$ returns a distribution on the FIsch state and configuration, instead of a distribution of values. Intuitively, $\text{Flexec}_{\Upsilon,n}$ simulates the n steps of the execution of a FIsch Υ until it returns a None. The full program execution of a FIsch is taken as the limit of its n -step program execution, i.e. $\text{Flexec}_{\Upsilon}(\Xi, \rho) \triangleq \lim_{n \rightarrow \infty} \text{Flexec}_{\Upsilon,n}(\Xi, \rho)$.

Properties of FIsches. As mentioned, FIsches are only used internally to define the model of Foxtrot. We relate FIsches and our original definition of schedulers via the following lemma:

LEMMA 5.1. *Given a FIsch Υ , there exists a (regular) scheduler ζ such that for all FIsch states Ξ and values v , we have $(\text{Flexec}_{\Upsilon}(\Xi, \rho) \gg g)(v) \leq \text{exec}_{\zeta}(\Xi, \rho)(v)$ where*

$$g \triangleq \lambda(\Xi', \rho'). \begin{cases} \text{ret } v & v \in \text{Val} \wedge \rho'.1 = v \cdot \vec{e} \\ \mathbf{0} & \text{otherwise} \end{cases}$$

We now consider various ways to construct FIsches, which are used in proving the soundness of the model of Foxtrot and its proof rules.

Firstly, the FIschLift function takes a FIsch Υ and FIsch state Ξ as input, and returns a FIsch that behaves the same as Υ but as if all states are additionally prepended with Ξ in front.

LEMMA 5.2. *There exists a function FIschLift such that given FIsch Υ and FIsch state Ξ , FIschLift(Ξ, Υ) is a FIsch where for all scheduler state Ξ' and configurations ρ , we have $\text{Flexec}_{\text{FIschLift}(\Xi, \Upsilon)}(\Xi \cdot \Xi', \rho) = \Upsilon(\Xi', \rho) \gg \lambda(\Xi'', \rho'). (\Xi \cdot \Xi'', \rho')$.*

Using FIschLift, we can define the FIschApp function, allowing us to chain a FIsch with a family of FIsches. Intuitively, FIschApp(Υ, f) is a FIsch that first acts like Υ until we reach a state Ξ_2 for which Υ always returns None, at which point, we continue to behave as the FIsch $f(\Xi_2)$. Later in the proof of the adequacy theorem (Theorem 3.1), this FIschApp function is our main weapon in piecing various FIsches together (see §5.3).

LEMMA 5.3. *There exists a function FIschApp such that given FIsch Υ and a function f of type FIschSt \rightarrow FIsch, FIschApp(Υ, f) is a FIsch where for all ρ, x , we have $(\text{Flexec}_{\Upsilon}(\epsilon, \rho) \gg g)(x) \leq \text{Flexec}_{\text{FIschApp}(\Upsilon, f)}(\epsilon, \rho)(x)$ where*

$$g(\Xi_1, \rho_1) \triangleq \begin{cases} \text{Flexec}_{\text{FIschLift}(\Xi_2, f(\Xi_2))}(\Xi_1, \rho_1) & \text{exists } \Xi_2 \leq \Xi_1 \text{ minimal s.t. } \forall \rho'. \Upsilon(\Xi_2, \rho') = \text{None} \\ \mathbf{0} & \text{otherwise} \end{cases}$$

5.2 Semantic Model

Weakest Precondition. We follow the standard trick of expressing Hoare triples in terms of a weakest precondition [21], i.e. $\{P\} e \{Q\} \triangleq P \multimap^* \text{wp } e \{Q\}$. The definition of the weakest precondition (shown below) follows similar structure as previous logics [19, 26] where it is defined as a guarded fixed point (all recursive occurrences of the weakest precondition appear under the

later modality \triangleright). We emphasize that the novelty of the Foxtrot weakest precondition lies in the significant differences in how the spec-coupling (sstep) and program-coupling (pstep) preconditions are defined, which we detail later.

$$\begin{aligned} \text{wp } e_1 \{ \Phi \} \triangleq & \forall \sigma_1, \rho_1, \varepsilon_1. S \sigma_1 \rho_1 \varepsilon_1 \text{ --* } \top \Vdash_{\emptyset} \text{sstep } \sigma_1 \rho_1 \varepsilon_1 \{ \sigma_2, \rho_2, \varepsilon_2. \\ & (e_1 \in \text{Val} * \emptyset \Vdash_{\top} S \sigma_2 \rho_2 \varepsilon_2 * \Phi e_1) \vee \\ & (e_1 \notin \text{Val} * \text{pstep } (e_1, \sigma_2) \rho_2 \varepsilon_2 \{ e_2, \sigma_3, l, \rho_3, \varepsilon_3. \\ & \triangleright \text{sstep } \sigma_3 \rho_3 \varepsilon_3 \{ \sigma_4, \rho_4, \varepsilon_4. \emptyset \Vdash_{\top} S \sigma_4 \rho_4 \varepsilon_4 * \text{wp } e_2 \{ \Phi \} * *_{e' \in l} \text{wp } e' \{ \text{True} \} \} \} \} \} \end{aligned}$$

To prove a weakest precondition, we initially assume the ownership of the *state interpretation* $S \sigma_1 \rho_1 \varepsilon_1$, which interprets the physical state of the implementation program, the specification program, and the approximate error as resources in the standard way [19, 21, 26]. It hence gives meaning to various resources e.g. the points-to connective $\ell \mapsto v$, the specification thread connective $j \Rightarrow e$, and the error credits $\mathcal{f}(\varepsilon)$. We then get access to the resources in all invariants through an update modality $\top \Vdash_{\emptyset}$ and we need to prove a *spec-coupling* precondition $\text{sstep } \sigma_1 \rho_1 \varepsilon_1 \{ \dots \}$. We explain its definition later in Figure 6; for now it suffices to interpret it as a precondition for allowing the right-hand side program to progress.

We subsequently perform a case split on whether e_1 is a value. If it is, we do a view shift $\emptyset \Vdash_{\top}$ where we re-establish all invariants, return the updated state interpretation and prove that e_1 satisfies the postcondition Φ . Otherwise, we prove a *program-coupling* precondition $\text{pstep } (e_1, \sigma_2) \rho_2 \varepsilon_2 \{ \dots \}$. As before, we explain the details of this precondition later, but for now one can think of it as a precondition similar to that of sstep , but it allows us to also take an actual step on the configuration (e_1, σ_2) to obtain resulting expression e_2 , state σ_3 , a list of forked expressions l , and resulting error budget ε_3 . We then prove another sstep precondition (which can be ignored here and is only needed to establish **HT-INV-OPEN**). We finally re-establish all invariants through the view shift $\emptyset \Vdash_{\top}$ and return the newly updated state interpretation, prove that $\text{wp } e_2 \{ \Phi \}$ holds, and show that all forked expressions are safe to execute, i.e. $\text{wp } e' \{ \text{True} \}$ for every expression e' in the list l .

Coupling Preconditions. We now turn our attention to the spec-coupling (sstep) and program-coupling (pstep) preconditions. At a high level, these preconditions are defined with FISches that satisfy coupling properties, which then become the building blocks to be “pieced together” in the proof of the adequacy theorem (Theorem 3.1).

We first present two auxiliary definitions. Firstly, we say a relation $R \subseteq A \times (B \times C)$ is state injective, written as $\text{stateinj}(R)$, if it is functional in B , i.e. for all a, a', b, c, c' , if $R a (b, c)$ and $R a' (b, c')$, then $a = a'$ and $c = c'$.

Secondly, we define the notion of a scheduler erasable state update, which intuitively describes actions on the state (mostly on the tape heap) that do not alter the operational semantics of the expressions; examples of scheduler erasable state updates include tape presamplings actions.

DEFINITION 5.4. A distribution on states μ is a scheduler erasable state update of $\sigma \in \text{State}$, written as $\text{schErasable}(\mu, \sigma)$, if for all schedulers ζ , scheduler states Ξ , thread pools \vec{e} , and integers n , we have

$$(\mu \gg (\lambda \sigma'. \text{pexec}_{\zeta, n}(\Xi, (\vec{e}, \sigma')))).\text{tp} = (\text{pexec}_{\zeta, n}(\Xi, (\vec{e}, \sigma))).\text{tp}$$

where the function $-\text{.tp}$ projects out the thread pool component from a distribution on configurations.

The spec-coupling precondition sstep is defined inductively by four inference rules shown in Figure 6. If the error budget ε is lower-bounded by 1, the precondition holds trivially as all sub-distributions have mass smaller or equal to 1 (**SPEC-STEP-ERR-1**). If the postcondition holds for the input parameters, the precondition also holds trivially (**SPEC-STEP-RET**). If the precondition holds for

$$\begin{array}{c}
\text{SPEC-STEP-ERR-1} \quad \frac{1 \leq \varepsilon}{\text{sstep } \sigma \rho \varepsilon \{ \Phi \}} \quad \text{SPEC-STEP-RET} \quad \frac{\Phi(\sigma, \rho, \varepsilon)}{\text{sstep } \sigma \rho \varepsilon \{ \Phi \}} \quad \text{SPEC-STEP-CONTINUOUS} \quad \frac{\forall \varepsilon'. \varepsilon < \varepsilon' \multimap \text{sstep } \sigma \rho \varepsilon' \{ \Phi \}}{\text{sstep } \sigma \rho \varepsilon \{ \Phi \}} \\
\text{SPEC-STEP-EXP} \quad \frac{\text{schErasable}(\mu, \sigma_1) \quad \exists r. \forall x. \mathcal{F}_2(x) \leq r \quad (\varepsilon' + \mathbb{E}_{\text{Flexec}_Y(\varepsilon, \rho_1)}[\mathcal{F}]) \leq \varepsilon \quad \mu \lesssim_{\varepsilon'} \text{Flexec}_Y(\varepsilon, \rho_1) : R \quad \text{stateinj}(R) \quad \forall \sigma_2, (\Xi, \rho_2). R \sigma_2 (\Xi, \rho_2) \multimap \text{sstep } \sigma_2 \rho_2 (\mathcal{F}(\Xi, \rho_2)) \{ \Phi \}}{\text{sstep } \sigma_1 \rho_1 \varepsilon \{ \Phi \}}
\end{array}$$

Fig. 6. Inductive Definition of the spec-coupling precondition $\text{sstep } \sigma \rho \varepsilon \{ \Phi \}$.

any error budget ε' larger than ε , then the precondition also holds for ε (**SPEC-STEP-CONTINUOUS**). This rule enables us to prove **PUPD-ERR** and create error credits from thin air.

We now focus on **SPEC-STEP-EXP**, which is the most important inference rule in the sstep precondition. Intuitively, this rule establishes an ε' -approximate coupling between a scheduler erasable state update μ , e.g. a tape presampling action on the left, with some execution steps of the right-hand side specification program under some state injective relation R . Here, the execution of the specification program is captured by the limit execution of a Flsch Y with the empty list ε as the starting state. In addition, the rule allows us to distribute the remaining error according to some function \mathcal{F} as long as it satisfies an expectation-preserving inequality.

The program-coupling precondition pstep is defined by the single inference rule **PROG-STEP-EXP** which is of a similar structure as that of **SPEC-STEP-EXP**. The main difference here is that this rule couples a single execution step on the left-hand side program with some execution steps on the right-hand side program according to some Flsch Y . (The proposition $\text{red}(e, \sigma)$ means that the configuration (e, σ) is reducible for one execution step.)

$$\text{PROG-STEP-EXP} \quad \frac{\exists r. \forall x. \mathcal{F}(x) \leq r \quad (\varepsilon' + \mathbb{E}_{\text{Flexec}_Y(\varepsilon, \rho_1)}[\mathcal{F}]) \leq \varepsilon \quad \text{step}(e_1, \sigma_1) \lesssim_{\varepsilon'} \text{Flexec}_Y(\varepsilon, \rho_1) : R \quad \text{red}(e_1, \sigma_1) \quad \text{stateinj}(R) \quad \forall (e_2, \sigma_2, l), (\Xi, \rho_2). R (e_2, \sigma_2, l) (\Xi, \rho_2) \multimap \Phi(e_2, \sigma_2, l, \rho_2, \mathcal{F}(\Xi, \rho_2))}{\text{pstep}(e_1, \sigma_1) \rho_1 \varepsilon \{ \Phi \}}$$

Probabilistic Update Modality. We follow the approach of Coneris [26] in defining the probabilistic update modality. Just like the fancy update modality $\varepsilon_1 \Rightarrow \varepsilon_2$, the probabilistic update modality is in fact additionally annotated with two sets of invariants, i.e. $\varepsilon_1 \overset{\text{mask}}{\Rightarrow} \varepsilon_2$. Intuitively, $\varepsilon_1 \overset{\text{mask}}{\Rightarrow} \varepsilon_2 P$ is defined by first assuming ownership of some state interpretation, opening all invariants in ε_1 , proving a spec-coupling precondition with input parameters σ_1, ρ_1 , and ε_1 , re-establishing all invariants in the mask ε_2 , and finally giving back the state interpretation and resource P .

$$\varepsilon_1 \overset{\text{mask}}{\Rightarrow} \varepsilon_2 P \triangleq \forall \sigma_1, \rho_1, \varepsilon_1. S \sigma_1 \rho_1 \varepsilon_1 \multimap \varepsilon_1 \Rightarrow_{\emptyset} \text{sstep } \sigma_1 \rho_1 \varepsilon_1 \{ \sigma_2, \rho_2, \varepsilon_2. \emptyset \Rightarrow_{\varepsilon_2} S \sigma_2 \rho_2 \varepsilon_2 * P \}$$

5.3 Soundness

As hinted in §1.2, to compose families of Flsches, we rely on an Iris version of the axiom of choice:

LEMMA 5.5 (IRIS CHOICE). *We have $(\forall a. \exists b. P a b) \vdash \exists f. \forall a. P a (f(a))$ within the Iris logic⁷.*

Perhaps this result might appear suspicious at first glance. Diaconescu [13] showed that the original axiom of choice implies the law of excluded middle, which is known to not hold in the intuitionistic Iris separation logic [23]. The seeming contradiction is avoided by restricting Lemma 5.5 to only apply when the variables a and b are quantified over meta-level Rocq types. Consequently, the types of a and b cannot utilize the step-indexing and resources of the Iris logic, as would be necessary to replay Diaconescu's proof.

⁷Assuming the axiom of choice in the meta-logic.

The adequacy theorem ([Theorem 3.1](#)) is proven via the following intermediate lemma [Lemma 5.6](#)⁸. We then recover [Theorem 3.1](#) by first taking \vec{e} to be the empty list ϵ , applying [Lemma 5.1](#) to eliminate the use of FISches, and finally taking the limit of n .

LEMMA 5.6. *If we can prove*

$$\not\vdash(\epsilon) * \bigstar_{(j,e') \in \text{enum}(\vec{e}_s)} j \models e' \vdash \text{wp } e \{v.\exists v'. 0 \Rightarrow v' * \phi v v'\} * \bigstar_{e' \in \vec{e}} \text{wp } e' \{\text{True}\}$$

in Foxtrot, then for all schedulers ζ , scheduler states Ξ , states σ, σ' , natural numbers n , and error $\epsilon' > 0$, there exists a FIsch Υ such that $\text{exec}_{\zeta,n}(\Xi, (e \cdot \vec{e}, \sigma)) \lesssim_{\epsilon+\epsilon'} \text{Flexec}_{\Upsilon}(\epsilon, (\vec{e}_s, \sigma')) : \psi$ where

$$\psi v (\Xi', \rho') \triangleq \exists v' \vec{e}'. \rho'. 1 = v' \cdot \vec{e}' \wedge \phi v v'$$

Intuitively, [Lemma 5.6](#) states that for any scheduler ζ on the left and any execution steps n , we can find a FIsch Υ for the right such that we can couple the n -step execution of the left-hand side program with the limit execution of the one on the right. In the proof of [Lemma 5.6](#), we start by fixing the arbitrary scheduler for the left-hand side program. Recall that the preconditions pstep and sstep are defined with FISches for the right-hand side program that satisfy various coupling properties. For each step of the program on the left, unfolding the weakest precondition of Foxtrot (specifically the spec-coupling and program-coupling preconditions) gives us various FISches for the right and an associated probabilistic coupling property for each. The key idea of the proof is to “piece together” all the FISches and couplings from pstep and sstep to produce a single FIsch for the right-hand side program, along with a coupling that covers the entire execution.

This piecing operation is done by induction on n ; given a FIsch that couples the first step and a family of FISches that couples n more steps for each possible outcome after the first step, we can apply [Lemma 5.5](#) and the FIschApp function to glue the first FIsch with the family of FISches to form a new FIsch that couples the first $n + 1$ steps of execution. Here [Lemma 5.5](#) is needed because the size of the family of FISches might be countably infinite. Interestingly, instead of performing a *single* “FIsch-piecing” operation in the inductive step, we actually have to glue the FISches a total of 4 times in the inductive step! We provide a more detailed proof sketch of [Lemma 5.6](#) in the extended version of this paper [[27](#)].

6 Related Work

Contextual Equivalences of Higher-Order Programs. There are various prior studies on utilizing binary logical relations to prove contextual equivalences of higher-order programs. For example, there is a long line of work using step-indexed logical relations for establishing refinements of higher-order concurrent languages. In particular, we highlight CaReSL [[42](#)], one of the first logics for fine-grained concurrency, which was then mechanized in Iris [[25, 40](#)]. The mechanized concurrent separation logic ReLoC [[17](#)] extended CaReSL to internalize refinement judgements as first-class logical statements, providing more concise rules for reasoning about invariants and logical atomicity. Similarly, there are also many logical relations specialized for probabilistic languages. Bizjak and Birkedal [[7](#)] defined a step-indexed, biorthogonal logical relation for reasoning about contextual equivalences of a higher-order probabilistic language with discrete random variables and state. This was then extended for more complicated language features such as continuous random sampling [[10, 45](#)] and countable nondeterminism [[3](#)]. The mechanized logic Clutch [[18](#)] first introduced presampling tapes for asynchronous couplings, which we also use to implement advanced coupling rules for reasoning about rejection sampling and multiple concurrent samplings

⁸We quantify over errors $\epsilon' > 0$ since sometimes an optimal scheduler cannot be found for the right-hand side program, but we can find one that approximately couples with the left for any positive error (see the extended version of this paper [[27](#)]).

in Foxtrot. We also took inspiration from Approxis [19], which used error credits to establish contextual equivalences of rejection samplers by means of approximation.

There are other techniques other than logical relations to verify contextual equivalences, e.g. Ehrhard et al. [14] utilized probabilistic coherence spaces to construct a fully abstract denotational model for PCF programs with a discrete random primitive. Much work has also been done in using bisimulation techniques to establish the contextual equivalences of programs written in the probabilistic λ -calculus, whether it be call-by-value [9], call-by-name [11], or call-by-need [34].

Compared to all the work above, our work in proving contextual equivalences of higher-order (stateful) programs is the *first* to be extended to the concurrent and randomized setting, which we believe is a *significant* contribution as reasoning about this mixture is notoriously hard.

Logics for Probability and Concurrency. There has been previous work on program logics for reasoning about properties (other than contextual refinement) of concurrent probabilistic programs.

McIver et al. [31] first introduced the probabilistic rely-guarantee calculus for reasoning about the quantitative correctness of probabilistic concurrent programs without support for dynamically-allocated local state. Fesefeldt et al. [16] developed concurrent quantitative separation logic for reasoning about quantitative properties of concurrent, randomized, heap-manipulating programs. Polaris [37] is similar to our work in that it is also a mechanized relational logic in Iris, but the refinement is constructed between a concurrent probabilistic program and a simpler monadic representation model, which can then further be analyzed to establish properties of the original program, e.g. bounds on expected values. ExpIris [30] is another mechanized program logic in Iris; its Hoare triples are annotated with a parameter called the potential, which is used to prove bounds on expected costs of concurrent probabilistic programs. The recent Probabilistic Concurrent Outcome Logic [47], a variant of Outcome Logic [46], supports compositional reasoning of the distribution of possible outcomes from the execution of concurrent randomized programs.

Most similar to our logic is that of Coneris [26], a separation logic for proving error bounds of concurrent probabilistic programs. Firstly, the underlying language ConcRandML of Foxtrot is identical to that in Coneris. Moreover, our use of error credits is heavily inspired by that of Coneris [26] (and similar logics such as Eris [4] and Approxis [19]), where the error credits were originally used to prove error bounds of (concurrent) probabilistic programs. We also adopt Coneris' probabilistic update modality and extend it into our relational setting for symbolically executing right-hand side programs in addition to presampling onto tapes. Instead of proving contextual refinement between two programs, Coneris is a unary logic and is used to establish error probability bounds of individual programs.

7 Conclusion

We presented Foxtrot, the first higher-order separation logic for proving contextual refinements of higher-order concurrent probabilistic programs with higher-order local state. In addition, we demonstrated its strengths on a wide range of examples involving complex local state, random sampling, and concurrent behavior, that were previously out of scope.

In the future, we would like to extend Foxtrot to reason about contextual refinement of programs under schedulers with restricted power for proving security guarantees of distributed cryptographic protocols, e.g. schedulers which can only rely on their internal state to decide which thread to step next. We also aim to implement prophecy variables [1, 2, 22] within Foxtrot to prove more interesting equivalences. Lastly, it would be interesting to consider modifying Foxtrot for proving must-termination contextual refinements [3, 6] of concurrent probabilistic programs.

Data Availability Statement

The Rocq formalization accompanying this work is available online on Zenodo [28] and on Github at github.com/logsem/clutch.

Acknowledgments

The first author would like to thank Amin Timany and Daniel Gratzer for insightful discussions regarding the axiom of choice.

This work was supported in part by the National Science Foundation, grant no. 2338317, a Villum Investigator grant, no. VIL25804 and no. VIL73403, Center for Basic Research in Program Verification (CPV), from the VILLUM Foundation, and the European Union (ERC, CHORDS, 101096090). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

- [1] M. Abadi and L. Lamport. 1988. The existence of refinement mappings. In *[1988] Proceedings. Third Annual Symposium on Logic in Computer Science*. 165–175. doi:10.1109/LICS.1988.5115
- [2] Martin Abadi and Leslie Lamport. 1991. The existence of refinement mappings. *Theoretical Computer Science* 82, 2 (1991), 253–284. doi:10.1016/0304-3975(91)90224-P
- [3] Alejandro Aguirre and Lars Birkedal. 2023. Step-Indexed Logical Relations for Countable Nondeterminism and Probabilistic Choice. *Proceedings of the ACM on Programming Languages* 7 (9 Jan. 2023), 33–60. doi:10.1145/3571195
Publisher Copyright: © 2023 Owner/Author.
- [4] Alejandro Aguirre, Philipp G. Haselwarter, Markus de Medeiros, Kwing Hei Li, Simon Oddershede Gregersen, Joseph Tassarotti, and Lars Birkedal. 2024. Error Credits: Resourceful Reasoning about Error Bounds for Higher-Order Probabilistic Programs. *Proc. ACM Program. Lang.* 8, ICFP, Article 246 (Aug. 2024), 33 pages. doi:10.1145/3674635
- [5] James Aspnes and Eric Ruppert. 2016. Depth of a Random Binary Search Tree with Concurrent Insertions. In *Distributed Computing*, Cyril Gavoille and David Ilcinkas (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 371–384.
- [6] Lars Birkedal, Aleš Bizjak, and Jan Schwinghammer. 2013. Step-Indexed Relational Reasoning for Countable Nondeterminism. *Logical Methods in Computer Science* Volume 9, Issue 4 (Oct. 2013). doi:10.2168/lmcs-9(4:4)2013
- [7] Aleš Bizjak and Lars Birkedal. 2015. Step-Indexed Logical Relations for Probability. In *Foundations of Software Science and Computation Structures*, Andrew Pitts (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 279–294.
- [8] Sylvie Boldo, Catherine Lelay, and Guillaume Melquiond. 2015. Coqelicot: A User-Friendly Library of Real Analysis for Coq. *Mathematics in Computer Science* 9, 1 (March 2015), 41–62. doi:10.1007/s11786-014-0181-1
- [9] Raphaëlle Crubillé and Ugo Dal Lago. 2014. On Probabilistic Applicative Bisimulation and Call-by-Value λ -Calculi. In *Programming Languages and Systems*, Zhong Shao (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 209–228.
- [10] Ryan Culpepper and Andrew Cobb. 2017. Contextual Equivalence for Probabilistic Programs with Continuous Random Variables and Scoring. In *Programming Languages and Systems*, Hongseok Yang (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 368–392.
- [11] Ugo Dal Lago, Davide Sangiorgi, and Michele Alberti. 2014. On coinductive equivalences for higher-order probabilistic functional programs. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Diego, California, USA) (POPL '14). Association for Computing Machinery, New York, NY, USA, 297–308. doi:10.1145/2535838.2535872
- [12] Frank Denis. 2013. The Sodium cryptography library. <https://download.libsodium.org/doc/>
- [13] Radu Diaconescu. 1975. Axiom of Choice and Complementation. *Proc. Amer. Math. Soc.* 51, 1 (1975), 176–178. <http://www.jstor.org/stable/2039868>
- [14] Thomas Ehrhard, Christine Tasson, and Michele Pagani. 2014. Probabilistic coherence spaces are fully abstract for probabilistic PCF. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Diego, California, USA) (POPL '14). Association for Computing Machinery, New York, NY, USA, 309–320. doi:10.1145/2535838.2535865
- [15] David Evans, Vladimir Kolesnikov, and Mike Rosulek. 2018. *A Pragmatic Introduction to Secure Multi-Party Computation*. doi:10.1561/33000000019
- [16] Ira Fesefeldt, Joost-Pieter Katoen, and Thomas Noll. 2022. Towards Concurrent Quantitative Separation Logic. In *33rd International Conference on Concurrency Theory (CONCUR 2022) (Leibniz International Proceedings in Informatics*

- (*LIPICs*, Vol. 243), Bartek Klin, Sławomir Lasota, and Anca Muscholl (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 25:1–25:24. doi:10.4230/LIPICs.CONCUR.2022.25
- [17] Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2021. ReLoC Reloaded: A Mechanized Relational Logic for Fine-Grained Concurrency and Logical Atomicity. *Log. Methods Comput. Sci.* 17, 3 (2021). doi:10.46298/LMCS-17(3:9)2021
- [18] Simon Oddershede Gregersen, Alejandro Aguirre, Philipp G. Haselwarter, Joseph Tassarotti, and Lars Birkedal. 2024. Asynchronous Probabilistic Couplings in Higher-Order Separation Logic. *Proc. ACM Program. Lang.* 8, POPL, Article 26 (Jan. 2024), 32 pages. doi:10.1145/3632868
- [19] Philipp G. Haselwarter, Kwing Hei Li, Alejandro Aguirre, Simon Oddershede Gregersen, Joseph Tassarotti, and Lars Birkedal. 2025. Approximate Relational Reasoning for Higher-Order Probabilistic Programs. *Proc. ACM Program. Lang.* 9, POPL, Article 41 (Jan. 2025), 31 pages. doi:10.1145/3704877
- [20] Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-order ghost state. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*. 256–269. doi:10.1145/2951913.2951943
- [21] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018), e20. doi:10.1017/S0956796818000151
- [22] Ralf Jung, Rodolphe Lepigre, Gaurav Parthasarathy, Marianna Rapoport, Amin Timany, Derek Dreyer, and Bart Jacobs. 2019. The future is ours: prophecy variables in separation logic. *Proc. ACM Program. Lang.* 4, POPL, Article 45 (Dec. 2019), 32 pages. doi:10.1145/3371113
- [23] Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*. 637–650. doi:10.1145/2676726.2676980
- [24] Robbert Krebbers, Ralf Jung, Ales Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017. The Essence of Higher-Order Concurrent Separation Logic. In *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*. 696–723. doi:10.1007/978-3-662-54434-1_26
- [25] Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017. Interactive proofs in higher-order concurrent separation logic. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 205–217. doi:10.1145/3009837.3009855
- [26] Kwing Hei Li, Alejandro Aguirre, Simon Oddershede Gregersen, Philipp G. Haselwarter, Joseph Tassarotti, and Lars Birkedal. 2025. Modular Reasoning about Error Bounds for Concurrent Probabilistic Programs. *Proc. ACM Program. Lang.* 9, ICFP, Article 245 (Aug. 2025), 30 pages. doi:10.1145/3747514
- [27] Kwing Hei Li, Alejandro Aguirre, Joseph Tassarotti, and Lars Birkedal. 2025. Contextual Refinement of Higher-Order Concurrent Probabilistic Programs (Extended Version). arXiv:2511.10135 [cs.LO] <https://arxiv.org/abs/2511.10135>
- [28] Kwing Hei Li, Alejandro Aguirre, Joseph Tassarotti, and Lars Birkedal. 2026. Contextual Refinement of Higher-Order Concurrent Probabilistic Programs. doi:10.5281/zenodo.18924289
- [29] T. Lindvall. 2002. *Lectures on the Coupling Method*. Dover Publications, Incorporated.
- [30] Janine Lohse and Deepak Garg. 2024. An Iris for Expected Cost Analysis. arXiv:2406.00884 [cs.PL]
- [31] Annabelle McIver, Tahiry Rabehaja, and Georg Struth. 2016. Probabilistic rely-guarantee calculus. *Theoretical Computer Science* 655 (2016), 120–134. doi:10.1016/j.tcs.2016.01.016 Quantitative Aspects of Programming Languages and Systems (2013-14).
- [32] Robert Morris. 1978. Counting large numbers of events in small registers. *Commun. ACM* 21, 10 (Oct. 1978), 840–842. doi:10.1145/359619.359627
- [33] Andrew M. Pitts. 2002. Operational Semantics and Program Equivalence. In *Applied Semantics*, Gilles Barthe, Peter Dybjer, Luís Pinto, and João Saraiva (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 378–412.
- [34] David Sabel, Manfred Schmidt-Schauß, and Luca Maio. 2022. Contextual Equivalence in a Probabilistic Call-by-Need Lambda-Calculus. In *Proceedings of the 24th International Symposium on Principles and Practice of Declarative Programming (Tbilisi, Georgia) (PPDP '22)*. Association for Computing Machinery, New York, NY, USA, Article 4, 15 pages. doi:10.1145/3551357.3551374
- [35] Tetsuya Sato. 2016. Approximate Relational Hoare Logic for Continuous Random Samplings. In *The Thirty-second Conference on the Mathematical Foundations of Programming Semantics, MFPS 2016*. doi:10.1016/J.ENTCS.2016.09.043
- [36] Simon Spies, Lennard Gäher, Daniel Grätzer, Joseph Tassarotti, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. 2021. Transfinite Iris: resolving an existential dilemma of step-indexed separation logic. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (Virtual, Canada) (PLDI 2021)*. Association for Computing Machinery, New York, NY, USA, 80–95. doi:10.1145/3453483.3454031

- [37] Joseph Tassarotti and Robert Harper. 2019. A separation logic for concurrent randomized programs. *Proc. ACM Program. Lang.* 3, POPL, Article 64 (Jan. 2019), 30 pages. doi:10.1145/3290377
- [38] The Rocq Development Team. 2024. *The Rocq Prover*. doi:10.5281/zenodo.11551307
- [39] Hermann Thorisson. 2000. *Coupling, stationarity, and regeneration*. Springer-Verlag, New York. xiv+517 pages.
- [40] Amin Timany. 2018. Contributions in Programming Languages Theory: Logical Relations and Type Theory. <https://lirias.kuleuven.be/1838165>
- [41] Amin Timany, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. 2024. A Logical Approach to Type Soundness. *J. ACM* 71, 6, Article 40 (Nov. 2024), 75 pages. doi:10.1145/3676954
- [42] Aaron Turon, Derek Dreyer, and Lars Birkedal. 2013. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming (Boston, Massachusetts, USA) (ICFP '13)*. Association for Computing Machinery, New York, NY, USA, 377–390. doi:10.1145/2500365.2500600
- [43] Daniele Varacca and Glynn Winskel. 2006. Distributing probability over non-determinism. *Mathematical Structures in Comp. Sci.* 16, 1 (Feb. 2006), 87–113. doi:10.1017/S0960129505005074
- [44] John von Neumann. 1961. *John von Neumann Collected Works*. Vol. 5: Design of Computers, Theory of Automata and Numerical Analysis. Pergamon Press, Oxford, England, Chapter Various Techniques Used in Connection with Random Digits, 768–770. Notes by George E. Forsythe and reproduced from *J. Res. Nat. Bus. Stand. Appl. Math. Series*, vol. 3, pp. 36–38 (1955).
- [45] Mitchell Wand, Ryan Culpepper, Theophilos Giannakopoulos, and Andrew Cobb. 2018. Contextual equivalence for a probabilistic language with continuous random variables and recursion. *Proc. ACM Program. Lang.* 2, ICFP, Article 87 (July 2018), 30 pages. doi:10.1145/3236782
- [46] Noam Zilberstein, Derek Dreyer, and Alexandra Silva. 2023. Outcome Logic: A Unifying Foundation for Correctness and Incorrectness Reasoning. *Proc. ACM Program. Lang.* 7, OOPSLA1, Article 93 (April 2023), 29 pages. doi:10.1145/3586045
- [47] Noam Zilberstein, Alexandra Silva, and Joseph Tassarotti. 2026. Probabilistic Concurrent Reasoning in Outcome Logic: Independence, Conditioning, and Invariants. *Proc. ACM Program. Lang.* 10, POPL, Article 9 (Jan. 2026), 30 pages. doi:10.1145/3776651

Received 2025-11-10; accepted 2026-04-03