

A Separation Logic for Concurrent Randomized Programs

JOSEPH TASSAROTTI, Carnegie Mellon University, USA

ROBERT HARPER, Carnegie Mellon University, USA

We present Polaris, a concurrent separation logic with support for probabilistic reasoning. As part of our logic, we extend the idea of *coupling*, which underlies recent work on probabilistic relational logics, to the setting of programs with both probabilistic and non-deterministic choice. To demonstrate Polaris, we verify a variant of a randomized concurrent counter algorithm and a two-level concurrent skip list. All of our results have been mechanized in Coq.

CCS Concepts: • **Theory of computation** → **Separation logic; Program verification**;

Additional Key Words and Phrases: separation logic, concurrency, probability

ACM Reference Format:

Joseph Tassarotti and Robert Harper. 2019. A Separation Logic for Concurrent Randomized Programs. *Proc. ACM Program. Lang.* 3, POPL, Article 64 (January 2019), 31 pages. <https://doi.org/10.1145/3290377>

1 INTRODUCTION

Many concurrent algorithms use randomization to reduce contention and coordination between threads. Roughly speaking, these algorithms are designed so that if each thread makes a local random choice, then on average the aggregate behavior of the whole system will have some good property.

For example, probabilistic skip lists [Pugh 1990] are known to work well in the concurrent setting [Fraser 2004; Herlihy et al. 2006], because threads can independently insert nodes into the skip list without much synchronization. In contrast, traditional balanced tree structures are difficult to implement in a scalable way because re-balancing operations may require locking access to large parts of the tree.

However, concurrent randomized algorithms are difficult to write and reason about. The use of just concurrency or randomness alone makes it hard to establish the correctness of an algorithm. For that reason, a number of program logics for reasoning about concurrent [Dinsdale-Young et al. 2013, 2010; Fu et al. 2010; Jones 1983; Jung et al. 2015; Nanevski et al. 2014; O’Hearn 2007; Vafeiadis and Parkinson 2007] or randomized [Barthe et al. 2016, 2012; Kaminski et al. 2016; Morgan et al. 1996; Ramshaw 1979] programs have been developed.

But, to our knowledge, the only prior program logic designed for reasoning about programs that are both concurrent *and* randomized is the recent probabilistic rely-guarantee calculus developed by McIver et al. [2016], which extends Jones’s original rely-guarantee logic [Jones 1983] with probabilistic constructs. However, this logic lacks many of the features of modern concurrency logics. For example, starting with the work of Vafeiadis and Parkinson [2007] and Feng et al. [2007], many recent concurrency logics combine rely-guarantee style reasoning with some form of

Authors’ addresses: Joseph Tassarotti, Computer Science Department, Carnegie Mellon University, USA, jtassaro@andrew.cmu.edu; Robert Harper, Computer Science Department, Carnegie Mellon University, USA, rwh@cs.cmu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/1-ART64

<https://doi.org/10.1145/3290377>

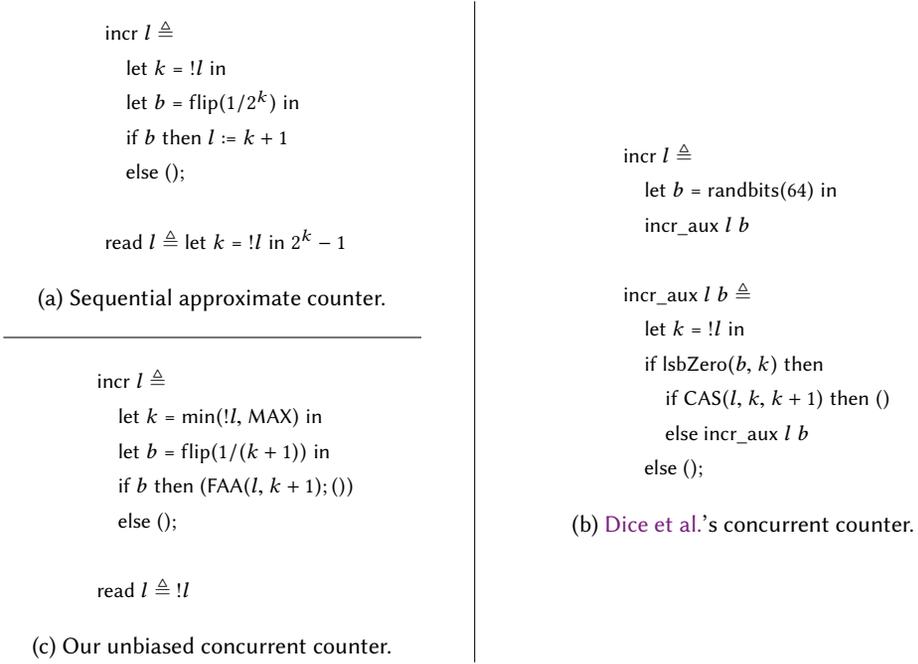


Fig. 1. Approximate counting algorithms.

separation logic, which is useful for modular, local reasoning about fine-grained concurrent data structures.

In this paper we describe Polaris, a logic which extends Iris [Jung et al. 2015], a state of the art concurrency logic, with support for probabilistic reasoning. By extending Iris, we ensure that our logic has the features needed to reason modularly about sophisticated concurrent algorithms. The key to our approach is several recent developments in probabilistic logic, concurrency logics, and denotational semantics. However, before we give an overview of how we build on this related work, let us describe a concurrent randomized algorithm, which will be our running example throughout this paper.

1.1 Example: Concurrent Approximate Counters

In many concurrent systems, threads need to keep counts of events. For example, in OS kernels, these counts can track performance statistics or reference counts. Somewhat surprisingly, Boyd-Wickizer et al. [2010] have shown that maintaining such counts was a serious scalability bottleneck in a prior version of the Linux kernel. In many cases, however, there is no need for these counts to be *exactly* right: an estimate is good enough. Taking advantage of this, Dice et al. [2013] created a scalable concurrent counter by adapting Morris's [Morris 1978] *approximate counting* algorithm.

In order to understand Dice et al.'s concurrent version it is helpful to understand Morris's original work. Morris's motivation was to be able to count up to n using fewer than $O(\log_2(n))$ bits. His idea was that, rather than storing the exact count n , one could instead store something like $\log_2(n)$ rounded to the nearest integer. This would require only $O(\log_2 \log_2(n))$ bits, at the cost of the error introduced by rounding.

Of course, if we round the stored count, then when we need to increment the counter, we do not know how to update the rounded value correctly. Instead, Morris developed a randomized

increment routine: if the counter currently stores the value k , then with probability $\frac{1}{2^k}$ we update the stored value to $k + 1$ and otherwise leave it unchanged. The code for this increment function is shown in [Figure 1a](#), written in an ML-like pseudo-code, where `flip(p)` is a command returning True with probability p and False otherwise. The read function loads the current value k of the counter and returns $2^k - 1$. Let C_n be the random variable giving the value stored in the counter after n calls of the increment function. One can show that $\mathbb{E}[2^{C_n} - 1] = n$. Thus, the expected value of the result returned by read is equal to the true number of increments, and so the counter is said to be an *unbiased estimator*. [Flajolet \[1985\]](#) gave a very detailed analysis of this algorithm by showing that it can be modeled by a simple Markov chain, and proved that it indeed only requires $O(\log_2 \log_2(n))$ bits with high probability.

Although this space-saving property is interesting, the aspect of the algorithm that makes it useful for concurrent counting is that, as the stored count gets larger, the probability that an increment needs to write to memory to update the count gets smaller. A simplified version of the concurrent algorithm proposed by [Dice et al. \[2013\]](#) is shown in [Figure 1b](#). The increment procedure starts by generating a large number of uniformly random bits. It then calls a recursive helper function `incr_aux` with the random bit vector b as an argument. This helper function reads the current value k stored in the counter and checks whether the first k bits of the bitvector are all 0 (performed in the code by the `lsbZero(k, b)` function). If not, the increment is over. If they are all zero, then because this occurs with probability $\frac{1}{2^k}$, the thread tries to atomically update the value stored in the counter from k to $k + 1$ using a compare-and-swap (CAS) operation. If this operation succeeds, it means that no other thread has intervened and modified the counter, and so the increment is finished. If the swap fails, some other thread has modified the counter, so the `incr_aux` function is recursively called to try again. The read procedure is the same as for the sequential algorithm.

As the count gets larger, the probability that a thread will perform a CAS operation during the increment gets smaller, which is useful because these operations are slow. [Dice et al. \[2013\]](#) show that this algorithm works quite well in practice, but do not give a formal argument for its correctness. Therefore, one might ask whether it is still guaranteed to give an unbiased estimate of the count. In fact, the answer is no: the scheduler can bias the count by ordering the compare-and-swap operations in a particular way. Imagine multiple threads are attempting to concurrently perform an increment, and the scheduler lets them each generate their random bits and then pauses them. Suppose the current value in the counter is k . Some of the threads may have drawn values that would cause them to not do an increment, because there is a 1 within their first k bits. Others may have drawn a number where far more than the first k bits will be 0: these threads would have performed an increment even if the value in the counter were larger than k . The scheduler can exploit this fact to maximize the value of the counter by running each thread one after the other in order of how many 0 bits they have at the beginning of their random number.¹

In [Figure 1c](#) we present a new concurrent version that is statistically unbiased, yet retains the same good properties of low contention.² Our increment function reads the current value in the counter, then takes the minimum of that value and a parameter `MAX`. If the minimum value is k , then with probability $\frac{1}{k+1}$, it uses a fetch-and-add operation (FAA) to atomically add $k + 1$ to the counter, otherwise it returns. In our version, the read function just returns the value in the counter. Like CAS, these FAA operations are expensive, so the reason the algorithm scales is that the probability that a FAA happens decreases as the counter value grows. The parameter `MAX`

¹Of course, a real scheduler is unlikely to behave in such an adversarial way. However, if the timing of operations can depend on random values, effects like this can arise even with non-adversarial schedulers, as discussed in [§5](#).

²However, it requires $O(\log_2(n))$ bits to store the count. Nevertheless, it uses less space than other alternatives for decreasing contention (e.g., having each thread maintain its own local counter).

caps how small this probability gets, somewhat like generating only 64 random bits does in the beginning of Figure 1b.

How does one show that this algorithm is unbiased, as we have claimed? Informally, it is because in expectation, each increment adds 1 to the count, so the total expected value is equal to the number of increments. Moreover, because addition is associative and commutative, it does not matter if other threads modify the counter in between when a flip happens and the corresponding FAA occurs. However, it is challenging to make this argument formal. We might try to model the value of the counter as a family of Markov chains,³ as Flajolet did for the sequential algorithm. But this is unwieldy because the relevant state of the chain is not just the current value stored in l , but also the local state of each thread in the middle of an increment operation. Moreover, even if one could model the algorithm in this way, it is hard to justify the connection between the concrete implementation and this mathematical representation.

As we will see, the program logic we have developed makes this algorithm easy to verify.

1.2 Background from Recent Work

Our program logic is based on three strands of recent work:

pRHL: Probabilistic Relational Reasoning. In many program logics for reasoning about probabilistic programs, assertions in the logic either explicitly make statements about probabilities, or are interpreted as being true with some probability (e.g., [Barthe et al. 2016; Kaminski et al. 2016; Morgan et al. 1996; Ramshaw 1979], among others). Although effective, this non-standard semantics of assertions is hard to reconcile with the semantics of concurrent separation logic, in which assertions are understood as claims of ownership of resources.

However, Barthe et al. have shown that reasoning about probabilistic programs can often be done *without* explicitly reasoning about probabilities in the assertions of the logic. Using their pRHL logic [Barthe et al. 2017a, 2012, 2017c], one establishes a refinement between two randomized programs using proof rules that encode a special type of simulation relation. The only time that explicit probabilities arise is a special rule for points in the simulation when *both* programs take a randomized step. The soundness theorem for their logic says that derivations in the logic imply the existence of a *coupling*, a construct from probability theory that is often used to relate two probability distributions. (We describe couplings in detail in §2.5).

Iris: A “Layered” Concurrency Logic. Modern concurrency logics are rather complex, making it hard to adapt them to incorporate probabilistic reasoning. Recent work, however, has sought to unify and simplify these logics [Dinsdale-Young et al. 2013; Jung et al. 2015; Nanevski et al. 2014].

In particular, Iris [Jung et al. 2016, 2015; Krebbers et al. 2017], a recent higher order concurrent separation logic, is composed of two layers: a “base logic” and a derived “program logic” which is encoded on top of the base logic. Crucially, most of the difficult semantic constructions are developed in the base logic. We are able to encode probabilistic relational reasoning *à la* pRHL in Iris by only modifying the second layer. We therefore get all of the results developed in the base logic “for free”, and retain the expressive features of Iris.

Indexed Valuations: A Monadic Encoding. Polaris, like pRHL, is designed for relational reasoning: it establishes a refinement between two programs. That means, if we want to prove a property about some program e , we first come up with some simpler *specification* program e' , use the logic to establish a refinement connecting properties of e' to those of e , and then reason about e' .

Therefore, we need to complement our logic with a way to express the simpler program e' and suitable tools for reasoning about it. We write these specification programs using the monad of

³Or rather, a *Markov decision process*, which accounts for the non-determinism of the ordering of operations.

indexed valuations developed by [Varacca and Winskel \[2006\]](#), which makes it possible to combine operations for both probabilistic and non-deterministic choice effects. This monad has a clean equational theory that makes it possible to reason about probabilistic properties of programs expressed using it.

Integrating this work together, however, is not straightforward. The challenge is that in pRHL, Hoare “quadruples” (which relate a pair of programs, in contrast to a traditional Hoare triple), are defined quite differently from Hoare triples in Iris. In the former, the quadruple is defined to hold just when an appropriate coupling between two programs exists. In contrast, in Iris, the Hoare triple is defined as a step-by-step property showing that various invariants are preserved. To bring pRHL-style probabilistic relational reasoning to Iris, we had to translate a condition about whole-program couplings into a step-by-step property.

1.3 Our Contributions

We make several contributions:

- We develop results for reasoning about computations expressed in the monadic encoding of [Varacca and Winskel](#). Although prior work had used this monad or similar ones combining both effects for denotational semantics [[Goubault-Larrecq 2007, 2015](#); [Mislove 2006](#); [Tix et al. 2009](#); [Varacca and Winskel 2006](#)] and to reason about small program equivalences [[Gibbons and Hinze 2011](#)], we found it necessary to develop new ways to reason about this monad. In particular, we develop an (in)equational theory of orderings between computations, and rewriting rules for bounding expected values. In addition, we adapt the notion of couplings to this setting. Finally, [Varacca and Winskel](#) focused on finite probabilistic and non-deterministic choice, but their constructions generalize to countable probabilistic choice and unbounded non-deterministic choice, which we use.
- We extend Iris with probabilistic relational reasoning in the style of pRHL, which lets us establish refinements between concurrent programs and these monadic representations.
- Using Polaris, we prove that the approximate counter algorithm introduced in §1.1 is unbiased.
- We also verify a fine-grained concurrent two-level skip list, and bound the expected number of comparisons performed when searching for a key.

All of the results in this paper, including the soundness of Polaris and the examples, have been mechanized in Coq.

We start by describing [Varacca and Winskel](#)’s [[Varacca and Winskel 2006](#)] monad for probabilistic and non-deterministic choice, and our results for reasoning about computations expressed in it (§2). We then describe Polaris (§3). In §4, we give a detailed explanation of how we use Polaris to verify the approximate counter example. Then in §5 we give an overview of the skip list example. Finally, we discuss additional related work in §6.

2 MONADIC REPRESENTATION

A common approach to reasoning about effectful programs is to model effects using a suitable monad, M . Using this monad, one represents an effectful program that returns a value of type T as a term of type $M(T)$. Next, one usually proves a series of equational rules for simplifying terms of type $M(T)$, and other lemmas for reasoning about such terms. This approach has been used for reasoning about a number of effects, including: state [[Nanevski et al. 2008](#); [Swierstra 2009](#)], non-termination [[Chlipala 2013](#)], non-determinism [[Gibbons and Hinze 2011](#)], and probabilistic choice [[Audebaud and Paulin-Mohring 2009](#); [Gibbons and Hinze 2011](#); [Petcher and Morrisett 2015](#); [van der Weegen and McKinna 2008](#)].

2.1 Monads for Non-determinism or Probability

Let us recall common monadic encodings for non-deterministic and probabilistic choice (separately). For non-determinism, we can define $M_N(T)$ as the type consisting of predicates $A : T \rightarrow \text{Prop}$ for which there exists at least some $t : T$ such that $A(t)$ holds. We think of these predicates as non-empty sets of terms of type T , where each element of the set represents one of the different non-deterministic outcomes. We say two terms A and B of type $M_N(T)$ are equivalent, written $A \equiv B$, if their sets of elements are the same: for all x , $x \in A \leftrightarrow x \in B$. In addition to the standard monadic operations (bind and return), we can represent non-deterministic choice between two computations A and B as the union $A \cup B$ of the two sets, defined by:

$$A \cup B \equiv \lambda t. A(t) \vee B(t)$$

This operation satisfies a number of natural rules:

$$A \cup B \equiv B \cup A \qquad A \cup (B \cup C) \equiv (A \cup B) \cup C \qquad A \cup A \equiv A$$

These, along with the usual monad laws, can be used to prove that one non-deterministic computation is equivalent to another.

We can represent a (discrete) probabilistic computation of type T as a function $f : T \rightarrow [0, 1]$, mapping values of type T to the probabilities that they occur. The *support* of f , written $\text{supp}(f)$ is the set of t such that $f(t) > 0$. Naturally, we want the sum of all the probabilities for the values in $\text{supp}(f)$ to be equal to 1. In order to make sense of such an infinite sum, we require the support to be countable. We define $M_P(T)$ to be the type of all functions $f : T \rightarrow [0, 1]$ such that $\text{supp}(f)$ is countable and

$$\sum_{x \in \text{supp}(f)} f(x) = 1$$

Given $A, B : M_P(T)$, we say $A \equiv B$ if for all x , $A(x) = B(x)$. We can define an operation which selects between a computation A with probability p and another computation B of the same type with probability $(1 - p)$:

$$A \oplus_p B \triangleq \lambda x. p \cdot A(x) + (1 - p) \cdot B(x)$$

This operation satisfies equational rules such as:

$$A \oplus_p B \equiv B \oplus_{1-p} A \qquad A \oplus_p A \equiv A$$

2.2 Combining Effects

In order to reason about programs that use both probability and non-determinism, we would like some way to *combine* the monads we have just described. We might try to represent computations of type T combining both effects as terms of type $M_N(M_P(T))$, *i.e.*, non-empty sets of probability distributions.

But how do we define the monad operations for this combination? One way to derive the monad operations for the combination is to specify a *distributive law* [Beck 1969]. However, Varacca and Winskel [2006] have given a proof (based on an idea they attribute to Plotkin) that no distributive law exists between the monads⁴ we have described above. For our purposes, it is not necessary to understand the impossibility proof. What is important is that, based on their impossibility arguments, Varacca and Winskel observe that the following equational law, which holds in the probabilistic choice monad, is problematic if we want to have a distributive law:

$$A \oplus_p A \equiv A$$

⁴More precisely, they consider the case where M_N is the monad of *finite* non-empty sets of terms of type T , and M_P consists of *finite* distributions, instead of countable ones. However, the impossibility of a distributive law in the finitary case precludes one for the non-finitary versions we have defined.

At first this equivalence seems entirely natural: if in either case we choose A , then the probabilistic choice was irrelevant. However, when we later add in the effect of non-determinism, removing this law becomes more justifiable, since it allows us to account for the fact that subsequent non-determinism in the computation can be *resolved differently* on the basis of this seemingly irrelevant probabilistic choice.

Using this observation, [Varacca and Winskel](#) describe an alternative way of representing probabilistic choice, which they call the *indexed valuation monad*, in which this equivalence does not hold, and they then describe a distributive law between non-empty lists and these indexed valuations to obtain a monad combining both effects.

An indexed valuation of type T is a tuple $(I, \text{ind}, \text{val})$, where I is a countable set whose elements are called *indices*; ind is a function of type $I \rightarrow T$, and val is a function of type $I \rightarrow \mathbb{R}^{\geq 0}$ such that ⁵

$$\sum_{i \in I} \text{val}(i) = 1$$

Informally, we can think of the indices as a set of “codes” or identifiers, the val function gives the probability of a particular index occurring, and ind maps these codes to elements of type T . Importantly, the ind function is not required to be injective, so that different codes can lead to the same observable result. We write $M_1(T)$ for the type of indexed valuations of type T . The *indicial support* of the valuation,⁶ notated $\text{isupp}(\text{val})$, is the set of indices i for which $\text{val}(i) > 0$. We say $(I_1, \text{ind}_1, \text{val}_1) \equiv (I_2, \text{ind}_2, \text{val}_2)$ if there exists a bijection $h : \text{isupp}(\text{val}_1) \rightarrow \text{isupp}(\text{val}_2)$ such that for all $i \in \text{isupp}(\text{val}_1)$, $\text{val}_1(i) = \text{val}_2(h(i))$ and $\text{ind}_1(i) = \text{ind}_2(h(i))$. That is, the bijection can only “relabel” indices in a way that preserves their probabilities and what they decode to. There is a map H which takes indexed valuations of type T to elements of $M_p(T)$:

$$H(I, \text{ind}, \text{val}) = \lambda x. \sum_{i \in \text{ind}^{-1}(\{x\})} \text{val}(i)$$

It is clear that if $\mathbb{I}_1 \equiv \mathbb{I}_2$, then $H(\mathbb{I}_1) \equiv H(\mathbb{I}_2)$. However, the converse is not true because \mathbb{I}_1 and \mathbb{I}_2 could have indicial supports with different cardinalities.

The probabilistic choice between two indexed valuations is defined by:

$$(I_1, \text{ind}_1, \text{val}_1) \oplus_p (I_2, \text{ind}_2, \text{val}_2) \triangleq (I_1 + I_2, \text{ind}', \text{val}')$$

where:

$$\text{ind}'(i) = \begin{cases} \text{ind}_1(i') & \text{if } i = \text{inl}(i') \\ \text{ind}_2(i') & \text{if } i = \text{inr}(i') \end{cases}$$

and

$$\text{val}'(i) = \begin{cases} p \cdot \text{val}_1(i') & \text{if } i = \text{inl}(i') \\ (1-p) \cdot \text{val}_2(i') & \text{if } i = \text{inr}(i') \end{cases}$$

One can show that for all indexed valuations \mathbb{I}_1 and \mathbb{I}_2 and $0 \leq p \leq 1$, we have $\mathbb{I}_1 \oplus_p \mathbb{I}_2 \equiv \mathbb{I}_2 \oplus_{1-p} \mathbb{I}_1$.

However, unlike the original probabilistic choice monad we described before, $\mathbb{I} \oplus_p \mathbb{I} \not\equiv \mathbb{I}$, unless $p = 0$ or $p = 1$. The reason is that, when p is neither 0 nor 1, the indicial support of $\mathbb{I} \oplus_p \mathbb{I}$ will have

⁵In fact, [Varacca and Winskel](#) first define a more general structure in which the sums of $\text{val}(i)$ do not have to equal 1, and the indices need not be countable. After working out some of the theory of these more general objects, they restrict to the subcategory where the indices are finite sets and the probabilities sum to 1. We will not restrict to finite sets of indices, since by letting them be countable we can model sampling from arbitrary discrete distributions.

⁶[Varacca and Winskel](#) call this simply the “support” of the indexed valuation, however we prefer to use that term for something different, defined below.

$$\begin{aligned}
\mathcal{I}_1 \oplus_p \mathcal{I}_2 &\equiv \mathcal{I}_2 \oplus_{1-p} \mathcal{I}_1 & \mathcal{I}_1 \oplus_1 \mathcal{I}_2 &\equiv \mathcal{I}_1 & \mathcal{I} \cup \mathcal{I} &\equiv \mathcal{I} & \mathcal{I}_1 \cup \mathcal{I}_2 &\equiv \mathcal{I}_2 \cup \mathcal{I}_1 \\
\mathcal{I}_1 \cup (\mathcal{I}_2 \cup \mathcal{I}_3) &\equiv (\mathcal{I}_1 \cup \mathcal{I}_2) \cup \mathcal{I}_3 & \mathcal{I}_1 \oplus_p (\mathcal{I}_2 \cup \mathcal{I}_3) &\equiv (\mathcal{I}_1 \oplus_p \mathcal{I}_2) \cup (\mathcal{I}_1 \oplus_p \mathcal{I}_3) \\
x \leftarrow \mathcal{I}_1 \cup \mathcal{I}_2; F(x) &\equiv (x \leftarrow \mathcal{I}_1; F(x)) \cup (x \leftarrow \mathcal{I}_2; F(x)) \\
x \leftarrow \mathcal{I}_1 \oplus_p \mathcal{I}_2; F(x) &\equiv (x \leftarrow \mathcal{I}_1; F(x)) \oplus_p (x \leftarrow \mathcal{I}_2; F(x))
\end{aligned}$$

Fig. 2. Equational laws for $M_N \circ M_1$ monad.

a larger cardinality than the indicial support of \mathbb{I} , so there can be no bijection between them. Recall that we do *not* want this equivalence to hold while trying to define the combined monad, because it is the one that [Varacca and Winskel \[2006\]](#) identified as problematic. With this obstruction removed, it is possible to define the monad operations on $M_N \circ M_1$, and we write M_{N1} for this composition. Given \mathcal{I}_1 and \mathcal{I}_2 of type $M_{N1}(T)$, the probabilistic choice operation $\mathcal{I}_1 \oplus_p \mathcal{I}_2$ is defined by taking the pairwise probabilistic choice of each indexed valuation in the respective sets:

$$\mathcal{I}_1 \oplus_p \mathcal{I}_2 \equiv \{\mathbb{I}_1 \oplus_p \mathbb{I}_2 \mid \mathbb{I}_1 \in \mathcal{I}_1, \mathbb{I}_2 \in \mathcal{I}_2\}$$

while the non-deterministic choice is once again the union of the sets. We say $\mathcal{I}_1 \equiv \mathcal{I}_2$ if for each $\mathbb{I}_1 \in \mathcal{I}_1$, there exists some $\mathbb{I}_2 \in \mathcal{I}_2$ such that $\mathbb{I}_1 \equiv \mathbb{I}_2$, and vice versa. The full definition of the bind operation is somewhat involved, as are the proofs of the monad laws, so we refer to [Varacca and Winskel \[2006\]](#). What is important is the equational properties that hold, of which a selection are shown in [Figure 2](#) (the standard monad laws are omitted).

2.3 Example: Modeling Approximate Counters

In [Figure 3](#) we show how to model the approximate counter code from [Figure 1c](#) using this monad. The `approxNcr` computation first non-deterministically selects a number k up to `MAX` – this models the process of taking the minimum of the value in `l` and `MAX` in the code. The non-determinism accounts for the fact that the value that will be read depends on what other threads do. The monadic encoding then makes a probabilistic choice, returning $k + 1$ with probability $\frac{1}{k+1}$ and 0 otherwise, which represents the probabilistic choice that the code will make about whether to do the `fetch-and-add`.

Finally, the process of repeatedly incrementing the counter n times is modeled by `approxN`. The first argument n tracks the number of pending increments to perform, while the second argument l accumulates the sum of the values returned by the calls to `approxNcr`. Note that this model *does not* try to represent multiple threads in the middle of an increment each waiting to add its value to the shared count – rather, it is *as if* the actual calls to `incr` all happened atomically in sequential order, with the effects of concurrency captured by the non-determinism in the `approxNcr` computation.

Of course, we need to show that this model accurately captures the behavior of the code from [Figure 1c](#) – this is what the program logic we describe in [§3](#) will do. First, however, we need to describe the new results we have developed to reason further about the monadic encoding itself.

2.4 Reasoning about Quantitative Properties

With what we have described so far, we can express computations with randomness and non-determinism and derive equivalences between them, but we do not yet have a way to talk about the standard concerns of probability theory (e.g., expected values, variances, tail bounds).

$$\begin{array}{l|l}
\text{approxIncr} \triangleq & \text{approxN } 0 \text{ } l \triangleq \text{ret } l \\
k \leftarrow \text{ret } 0 \cup \dots & \text{approxN } (n+1) \text{ } l \triangleq \\
\cup \text{ret MAX;} & k \leftarrow \text{approxIncr;} \\
\text{ret } (k+1) \oplus_{\frac{1}{k+1}} \text{ret } 0 & \text{approxN } n \text{ } (l+k)
\end{array}$$

Fig. 3. Monadic encoding of approximate counter algorithm from Figure 1c.

$$\begin{array}{l}
\mathbb{E}_f^{\min}[\text{ret } v] = f(v) \qquad \frac{k \geq 0}{\mathbb{E}_{(\lambda x.k \cdot f(x)+c)}^{\min}[I] = k \cdot \mathbb{E}_f^{\min}[I] + c} \\
\mathbb{E}_f^{\min}[I_1 \oplus_p I_2] = p \cdot \mathbb{E}_f^{\min}[I_1] + (1-p) \cdot \mathbb{E}_f^{\min}[I_2] \qquad \mathbb{E}_{gf}^{\min}[I] = \mathbb{E}_g^{\min}[x \leftarrow I; \text{ret } f(x)] \\
\frac{\forall x. k_1 \leq \mathbb{E}_f^{\min}[F(x)] \leq k_2}{k_1 \leq \mathbb{E}_f^{\min}[x \leftarrow I_1; F(x)] \leq k_2}
\end{array}$$

Fig. 4. Selection of rules for calculating extrema of expected values (analogous rules for $\mathbb{E}_f^{\max}[-]$ omitted).

Given an indexed valuation $\mathbb{I} = (I, \text{ind}, \text{val})$ of type T and a function $f : T \rightarrow \mathbb{R}$, we can define the *expected value* of f on \mathbb{I} as:

$$\mathbb{E}_f[\mathbb{I}] \triangleq \sum_{i \in I} f(\text{ind}(i)) \cdot \text{val}(i)$$

(this coincides with the usual notion of expected value of a random variable if we interpret the indexed valuation as a distribution using the map H defined above). Since I is a countable set, the above series may not necessarily converge. We say that the expected value of f on \mathbb{I} exists if the above series converges absolutely. Throughout this paper, when we mention expected values in rules and derivations, we will implicitly assume side conditions stating that all the relevant expected values exist.

If $\mathbb{I}_1 \equiv \mathbb{I}_2$, then for all f , $\mathbb{E}_f[\mathbb{I}_1] = \mathbb{E}_f[\mathbb{I}_2]$. Moreover, given a value t of type T , if we define $[t]$ to be the *indicator function* that returns 1 if its input is equal to t and 0 otherwise, then $\mathbb{E}_{[t]}[\mathbb{I}]$ is equal to the probability that \mathbb{I} yields the value t , so we can encode probabilities as expected values.

Since an \mathcal{I} of type $M_{\text{NI}}(T)$ is just a set of indexed valuations, we can apply $\mathbb{E}_f[-]$ to each $\mathbb{I} \in \mathcal{I}$ to get the set of expected values that can arise depending on how non-deterministic choices are resolved. Generally speaking, we will be interested in bounding the smallest or largest possible value that these expected values can take. We can define the *minimal* and *maximal* expected value of f on \mathcal{I} as:

$$\mathbb{E}_f^{\min}[\mathcal{I}] \triangleq \inf_{\mathbb{I} \in \mathcal{I}} \mathbb{E}_f[\mathbb{I}] \qquad \mathbb{E}_f^{\max}[\mathcal{I}] \triangleq \sup_{\mathbb{I} \in \mathcal{I}} \mathbb{E}_f[\mathbb{I}]$$

We say that these extrema exist if for all $\mathbb{I} \in \mathcal{I}$, $\mathbb{E}_f[\mathbb{I}]$ exists. Since \mathcal{I} may be an infinite set, $\mathbb{E}_f^{\min}[\mathcal{I}]$ and $\mathbb{E}_f^{\max}[\mathcal{I}]$ can be $-\infty$ and $+\infty$ respectively. Rules for calculating these values are given in Figure 4. As before, we implicitly assume that all of the stated extrema exist and are finite.

To help reason about these extrema, we introduce a partial order on terms of type $M_{\text{NI}}(T)$: We say $\mathcal{I}_1 \subseteq \mathcal{I}_2$ if for each $\mathbb{I}_1 \in \mathcal{I}_1$, there exists some $\mathbb{I}_2 \in \mathcal{I}_2$ such that $\mathbb{I}_1 \equiv \mathbb{I}_2$. If $\mathcal{I}_1 \subseteq \mathcal{I}_2$ then

$\mathbb{E}_f^{\max}[\mathcal{I}_1] \leq \mathbb{E}_f^{\max}[\mathcal{I}_2]$ and $\mathbb{E}_f^{\min}[\mathcal{I}_2] \leq \mathbb{E}_f^{\min}[\mathcal{I}_1]$. Thus, we can bound \mathcal{I}_1 's extrema by first finding some \mathcal{I}_2 such that $\mathcal{I}_1 \subseteq \mathcal{I}_2$, and then bounding the latter's extrema.

One way to ensure that expected values exist is to show that the functions we are computing expected values of are suitably bounded. We first define the *support* of \mathbb{I} as the set of all values that occur with non-zero probability:

$$\text{supp}(\mathbb{I}) \triangleq \{v \mid \exists i \in \mathbb{I}. \text{ind}(i) = v \wedge \text{val}(i) > 0\}$$

The support of a set of indexed valuations \mathcal{I} is then the union of their supports:

$$\text{supp}(\mathcal{I}) \triangleq \bigcup_{\mathbb{I} \in \mathcal{I}} \text{supp}(\mathbb{I})$$

We say that f is bounded on the support of \mathcal{I} if there exists some c such that $|f(v)| \leq c$ for all $v \in \text{supp}(\mathcal{I})$. If this holds, then $\mathbb{E}_f^{\min}[\mathcal{I}]$ and $\mathbb{E}_f^{\max}[\mathcal{I}]$ exist and are finite.

Using the above rules, we can show that $\mathbb{E}_{\text{id}}^{\min}[\text{approxN } n \ 0] = \mathbb{E}_{\text{id}}^{\max}[\text{approxN } n \ 0] = n$, which implies that no matter how the non-determinism is resolved in our model of the counter, the expected value of the result will be the number of increments. Let us just consider the case for the minimum, since the maximum is the same. The proof proceeds by induction on n , after first strengthening the induction hypothesis to the claim that $\mathbb{E}_{\text{id}}^{\min}[\text{approxN } n \ l] = n + l$. The key step of the proof is to show that $\mathbb{E}_{\text{id}}^{\min}[\text{approxInc}] = 1$, *i.e.*, each increment contributes 1 to the expected value. From the last rule in Figure 4, it suffices to show that whatever value of k is non-deterministically selected, the resulting expected value will be 1. We have that for all k :

$$\begin{aligned} & \mathbb{E}_{\text{id}}^{\min} \left[\text{ret } (k+1) \oplus_{\frac{1}{k+1}} \text{ret } 0 \right] \\ &= \left(\frac{1}{k+1} \right) \cdot (k+1) + \left(1 - \frac{1}{k+1} \right) \cdot 0 \\ &= 1 \end{aligned}$$

Let us summarize the discussion so far. Because the non-determinism monad M_{N} could not be combined with the standard probabilistic choice monad M_{P} , we replaced the latter with the monad of indexed valuations, M_{I} . The distinction between M_{I} and M_{P} is that indexed valuations carry additional data (the indices) and have a finer notion of equivalence. Because of this finer equivalence, the problematic equational law does not hold for M_{I} , so that the monad operations can be defined on M_{NI} .

We then re-developed the notions of expected values and probabilities for M_{I} and defined corresponding extrema of expected values for M_{NI} . Since these definitions respect the equivalence relations on M_{I} and M_{NI} , we could bound the extrema of some \mathcal{I} by first finding \mathcal{I}' such that $\mathcal{I} \equiv \mathcal{I}'$, and then bounding the extrema of \mathcal{I}' . More generally, we also had the \subseteq relation, so that similar bounds could be obtained solely by showing $\mathcal{I} \subseteq \mathcal{I}'$.

However, the relations \equiv and \subseteq above are finer than necessary if our goal is to use them to translate bounds on extrema of \mathcal{I}' to bounds on \mathcal{I} , and similarly so for translating bounds on expected values of one indexed valuation to another. For example, if f is a bounded function, then $\mathbb{E}_f[\mathbb{I}] = \mathbb{E}_f[\mathbb{I} \oplus_{\text{P}} \mathbb{I}]$, yet we know that $\mathbb{I} \not\equiv \mathbb{I} \oplus_{\text{P}} \mathbb{I}$.

Because we will be interested in using relational reasoning in order to bound expected values, it is useful to define the coarsest relations that suffice for this purpose, and derive some properties about them. We define $\mathbb{I} \equiv_{\text{P}} \mathbb{I}'$ to hold if for all bounded⁷ functions f , $\mathbb{E}_f[\mathbb{I}] = \mathbb{E}_f[\mathbb{I}']$. Because

⁷The reason for quantifying over bounded functions is to ensure that the two expected values exist.

$$\begin{array}{c}
\mathbb{I} \equiv_p \mathbb{I} \\
\frac{\mathbb{I}_1 \equiv \mathbb{I}'_1 \quad \mathbb{I}_2 \equiv \mathbb{I}'_2 \quad \mathbb{I}_1 \equiv_p \mathbb{I}_2}{\mathbb{I}'_1 \equiv_p \mathbb{I}'_2} \quad \frac{\mathbb{I}_1 \equiv_p \mathbb{I}_2 \quad \mathbb{I}_2 \equiv_p \mathbb{I}_3}{\mathbb{I}_1 \equiv_p \mathbb{I}_3} \\
\frac{\mathbb{I}_1 \equiv_p \mathbb{I}_2 \quad \forall x. F_1(x) \equiv_p F_2(x)}{x \leftarrow \mathbb{I}_1; F_1(x) \equiv_p x \leftarrow \mathbb{I}_2; F_2(x)} \quad \frac{\mathbb{I}_1 \equiv_p \mathbb{I}_2 \quad \mathbb{I}'_1 \equiv_p \mathbb{I}'_2}{\mathbb{I}_1 \oplus_p \mathbb{I}'_1 \equiv_p \mathbb{I}_2 \oplus_p \mathbb{I}'_2} \quad (x \leftarrow \mathbb{I}_1; \mathbb{I}_2) \equiv_p \mathbb{I}_2
\end{array}$$

Fig. 5. Rules for the \equiv_p relation on indexed valuations.

$$\begin{array}{c}
\mathcal{I} \subseteq_p \mathcal{I} \\
\frac{\mathcal{I}_1 \subseteq_p \mathcal{I}_2 \quad \mathcal{I}_2 \subseteq_p \mathcal{I}_3}{\mathcal{I}_1 \subseteq_p \mathcal{I}_3} \quad \frac{\mathcal{I}'_1 \subseteq \mathcal{I}_1 \quad \mathcal{I}_2 \subseteq \mathcal{I}'_2 \quad \mathcal{I}_1 \subseteq_p \mathcal{I}_2}{\mathcal{I}'_1 \subseteq_p \mathcal{I}'_2} \\
\frac{\mathcal{I}_1 \subseteq_p \mathcal{I}_2 \quad \forall x. F_1(x) \subseteq_p F_2(x)}{x \leftarrow \mathcal{I}_1; F_1(x) \subseteq_p x \leftarrow \mathcal{I}_2; F_2(x)} \quad \frac{\mathcal{I}_1 \subseteq_p \mathcal{I}_2 \quad \mathcal{I}'_1 \subseteq_p \mathcal{I}'_2}{\mathcal{I}_1 \oplus_p \mathcal{I}'_1 \subseteq_p \mathcal{I}_2 \oplus_p \mathcal{I}'_2} \quad (x \leftarrow \mathcal{I}_1; \mathcal{I}_2) \subseteq_p \mathcal{I}_2
\end{array}$$

Fig. 6. Rules for \subseteq_p relation.

indicator functions are bounded, observe that if $\mathbb{I} \equiv_p \mathbb{I}'$, then $\mathbb{E}_{[t]}[\mathbb{I}] = \mathbb{E}_{[t]}[\mathbb{I}']$ for all t . In other words, this notion of equivalence is the same as saying that the probability distributions $H(\mathbb{I})$ and $H(\mathbb{I}')$ are equal. Rules for this relation are shown in Figure 5. Crucially, it is a congruence with respect to all the monad operations.

Analogously, we define $\mathcal{I} \subseteq_p \mathcal{I}'$ to hold if for all bounded functions f , $\mathbb{E}_f^{\max}[\mathcal{I}] \leq \mathbb{E}_f^{\max}[\mathcal{I}']$. Thus, if this relation holds, one can bound the maxima of \mathcal{I} by bounding the maxima of \mathcal{I}' . Since the negation of f is bounded if and only if f is, this also implies that $\mathbb{E}_f^{\min}[\mathcal{I}'] \leq \mathbb{E}_f^{\min}[\mathcal{I}]$ for all bounded f . Some rules for this relation are shown in Figure 6.

The following lemma shows that the definition of \subseteq_p generalizes to a larger class of functions than just the bounded ones:

LEMMA 2.1. *If $\mathcal{I} \subseteq_p \mathcal{I}'$ and f is bounded on the support of \mathcal{I}' , then f is bounded on the support of \mathcal{I} and $\mathbb{E}_f^{\max}[\mathcal{I}] \leq \mathbb{E}_f^{\max}[\mathcal{I}']$.*

2.5 Non-deterministic Couplings

Recent work by Barthe et al. [2015, 2017a,c] has shown that the notion of a probabilistic coupling [Lindvall 2002] is fundamental for relational reasoning in probabilistic program logics. Given two distributions $A : M_p(T_A)$ and $B : M_p(T_B)$, a coupling between A and B is a distribution $C : M_p(T_A \times T_B)$ such that:

- (1) $\forall x : T_A. A(x) = \sum_y C(x, y)$
- (2) $\forall y : T_B. B(y) = \sum_x C(x, y)$

That is, C is a joint distribution whose marginals equal A and B . These two conditions are equivalent to requiring that:

- (1) $A \equiv ((x, y) \leftarrow C; \text{ret } x)$
- (2) $B \equiv ((x, y) \leftarrow C; \text{ret } y)$

Given a predicate $P : A \times B \rightarrow \text{Prop}$, we say that C is a P -coupling, if, in addition to the above, we have:

$$\forall x, y. C(x, y) > 0 \rightarrow P(x, y)$$

$\frac{\text{RET} \quad P(a, b)}{\text{ret } a \sim \text{ret } b : P}$	$\frac{\text{EQUIV} \quad \mathbb{I} \equiv \mathbb{I}' \quad \mathcal{I} \subseteq \mathcal{I}' \quad \mathbb{I} \sim \mathcal{I} : P}{\mathbb{I}' \sim \mathcal{I}' : P}$	$\frac{\text{CONSEQ} \quad \mathbb{I} \sim \mathcal{I} : P \quad \forall x, y. P(x, y) \rightarrow P'(x, y)}{\mathbb{I} \sim \mathcal{I} : P'}$
$\frac{\text{BIND} \quad \mathbb{I} \sim \mathcal{I} : P \quad \forall x, y. P(x, y) \rightarrow F(x) \sim F'(y) : Q}{(x \leftarrow \mathbb{I}; F(x)) \sim (y \leftarrow \mathcal{I}; F'(y)) : Q}$	$\frac{\text{P-CHOICE} \quad \mathbb{I} \sim \mathcal{I} : P \quad \mathbb{I}' \sim \mathcal{I}' : P}{\mathbb{I} \oplus_p \mathbb{I}' \sim \mathcal{I} \oplus_p \mathcal{I}' : P}$	$\text{TRIVIAL} \quad \mathbb{I} \sim \mathcal{I} : \text{True}$

Fig. 7. Rules for constructing non-deterministic couplings.

i.e., all pairs (x, y) in the support of the distribution C satisfy P . The existence of a P -coupling can tell us important things about the two distributions. For example, if $P(x, y) = (x = y)$, then the existence of a P -coupling tells us the two distributions are equivalent. Moreover, there are rules for systematically constructing couplings between distributions. We will explain some of these rules once we have described how to adapt couplings to the monad M_{NI} .

First, using the monadic formulation of the coupling conditions, it is straightforward to define an analogous idea for M_1 : Given $\mathbb{I}_1 : M_1(T_1)$ and $\mathbb{I}_2 : M_1(T_2)$, a coupling between \mathbb{I}_1 and \mathbb{I}_2 is an $\mathbb{I} : M_1(T_1 \times T_2)$ such that:

- (1) $\mathbb{I}_1 \equiv_p ((x, y) \leftarrow \mathbb{I}; \text{ret } x)$
- (2) $\mathbb{I}_2 \equiv_p ((x, y) \leftarrow \mathbb{I}; \text{ret } y)$

and $\mathbb{I} = (I, \text{ind}, \text{val})$ is a P -coupling if for all i such that $\text{val}(i) > 0$, $P(\text{ind}(i))$ holds. As before, if P is the equality predicate, then the existence of a P -coupling between \mathbb{I}_1 and \mathbb{I}_2 implies $\mathbb{I}_1 \equiv_p \mathbb{I}_2$.

We can lift this to a relation between a single indexed valuation \mathbb{I} and a set of indexed valuations \mathcal{I} . We say⁸ there is a *non-deterministic P -coupling* between \mathbb{I} and \mathcal{I} if there exists some \mathbb{I}' such that $\{\mathbb{I}'\} \subseteq_p \mathcal{I}$ and a P -coupling between \mathbb{I} and \mathbb{I}' . We write $\mathbb{I} \sim \mathcal{I} : P$ to denote the existence of such a coupling.

Rules for constructing these couplings are shown in [Figure 7](#). If we interpret the P in $\mathbb{I} \sim \mathcal{I} : P$ as a kind of “post-condition” for the execution of the computations \mathbb{I} and \mathcal{I} , then these coupling rules have the structure of a Hoare-like relational logic [[Benton 2004](#)], as in the work of [Barthe et al. \[2015\]](#): *e.g.*, the rule **BIND** is analogous to the usual sequencing rule in Hoare logic.

The rule **P-CHOICE** lets us couple probabilistic choices $\mathbb{I} \oplus_p \mathbb{I}'$ and $\mathcal{I} \oplus_p \mathcal{I}'$ with post-condition P by coupling \mathbb{I} to \mathcal{I} and \mathbb{I}' to \mathcal{I}' . This is somewhat surprising: we get to reason about these two probabilistic choices as if they both chose the left alternative or both chose the right alternative, rather than considering the full set of four combinations. This counter-intuitive rule is quite useful, as demonstrated in the many examples given in the work of [Barthe et al.](#) We will see an example of its use in [§4](#).

The following theorem lets us use the existence of a non-deterministic coupling to bound expected values:

THEOREM 2.2. *Let g be bounded on $\text{supp}(\mathcal{I})$ and let $P(x, y) = (f(x) = g(y))$. If $\mathbb{I} \sim \mathcal{I} : P$, then $\mathbb{E}_f[\mathbb{I}]$ exists and*

$$\mathbb{E}_g^{\min}[\mathcal{I}] \leq \mathbb{E}_f[\mathbb{I}] \leq \mathbb{E}_g^{\max}[\mathcal{I}]$$

⁸[Barthe et al. \[2015\]](#) use “non-deterministic coupling” to refer to a particular kind of coupling which is unrelated to adversarial non-deterministic choice.

Syntax:

<i>Val</i>	$v ::= \lambda x. e_1 \mid (v_1, v_2) \mid () \mid n \mid b \mid \dots$
<i>Expr</i>	$e ::= x \mid v \mid e_1 e_2 \mid \text{fork}\{e\} \mid \text{flip}(e_1, e_2) \mid \dots$
<i>Eval Ctx</i>	$K ::= [] \mid K e \mid V K \mid \text{flip}(K, e) \mid \text{flip}(v, K) \mid \dots$
<i>State</i>	$\sigma \in \mathbb{N} \rightarrow \text{Val}$
<i>Config</i>	$\rho \in \{l : \text{List Expr} \mid l \neq \emptyset\} \times \text{State}$
<i>Trace</i>	$T \in \{l : \text{List Config} \mid l \neq \emptyset\}$
<i>Scheduler</i>	$\varphi \in \text{Trace} \rightarrow \mathbb{N}$

Per-Thread Reduction: $e; \sigma \xrightarrow{p} e'; \sigma'$

FLIP-TRUE	FLIP-FALSE	
$0 \leq \frac{n_1}{n_2} \leq 1$	$0 \leq \frac{n_1}{n_2} \leq 1$	
<hr/>	<hr/>	
$\text{flip}(n_1, n_2); \sigma \xrightarrow{\frac{n_1}{n_2}} \text{True}; \sigma$	$\text{flip}(n_1, n_2); \sigma \xrightarrow{1 - \frac{n_1}{n_2}} \text{False}; \sigma$	(Standard rules omitted.)

Concurrent Semantics: $\rho \xrightarrow{i} \rho'$

$$\frac{e_i; \sigma \xrightarrow{p} e'_i; \sigma'}{[\dots, K[e_i], \dots]; \sigma \xrightarrow{i} [\dots, K[e'_i], \dots]; \sigma'}$$

$$[e_1, \dots, e_{i-1}, K[\text{fork}\{e_f\}], \dots]; \sigma \xrightarrow{i} [e_1, \dots, e_{i-1}, K[()], \dots, e_f]; \sigma$$

Trace Semantics: $T \xrightarrow[\varphi]{p} T'$

$$\frac{\varphi(T, \rho) = i \quad \rho \xrightarrow{i} \rho'}{T, \rho \xrightarrow[\varphi]{p} T, \rho, \rho'} \quad \frac{\varphi(T, \rho) = i \quad \neg(\exists \rho', p. \rho \xrightarrow{i} \rho')}{T, \rho \xrightarrow[\varphi]{1} T, \rho, \rho}$$

Fig. 8. Syntax and semantics of concurrent language.

3 PROGRAM LOGIC

We now describe Polaris, the program logic we have developed for proving that a program is modeled by the monadic specifications from the previous section.

3.1 Program Semantics

Polaris is parameterized by a generic probabilistic concurrent language. However, for concreteness, we instantiate it with the ML-like language used in the examples from §1. Figure 8 gives the syntax and semantics of this language. We omit the standard rules for things like tuples, recursive functions, and references. The per-thread reduction relation $e; \sigma \xrightarrow{p} e'; \sigma'$ is annotated with a probability p that the transition takes place. We say e is atomic, written $\text{atomic}(e)$, if e reduces to a value in a single step.

The $\text{flip}(n_1, n_2)$ command takes two integers as arguments and simulates a biased coin flip: it transitions to True with probability $\frac{n_1}{n_2}$ and False with probability $1 - \frac{n_1}{n_2}$. (In the introduction we somewhat informally wrote $\text{flip}(n_1/n_2)$ as if the language had rational numbers as a primitive). There is a side condition to ensure that $\frac{n_1}{n_2}$ actually corresponds to a valid probability. Other than this command, the per-thread transition system for this language is deterministic. The generic framework for our logic allows us to extend this language with other probabilistic commands, so long as they only sample from discrete probability distributions.

This per-thread reduction relation is then lifted to a concurrent transition system. A *configuration* ρ is a pair consisting of a list of expressions (representing a pool of threads) and a state σ . We say $\rho \xrightarrow[p]{i} \rho'$ when the i^{th} thread of ρ transitions with probability p leading to a new configuration ρ' . The $\text{fork}\{e_f\}$ command adds a new thread e_f to the pool.

A *scheduler* decides which thread will get to step at each point in an execution. We model a scheduler as a function φ of type $\text{Trace} \rightarrow \mathbb{N}$, where a trace is a non-empty list of configurations representing a partial execution. The scheduler is permitted to inspect the entire history and complete state of the program when deciding which thread gets to go next. Of course, a real implementation of a scheduler does not actually do this, but by conservatively considering this strong class of *adversarial* schedulers, results we prove will also hold for realistic schedulers. We write $T \xrightarrow[\varphi]{p} T'$ to indicate that the thread selected by $\varphi(T)$ steps with probability p to a new configuration which is appended to T to obtain T' . We write $\text{curr}(T)$ for the last configuration in a trace. If the scheduler returns a thread number which cannot take a step or which does not exist, the system takes a “stutter” step and the current configuration is repeated again at the end of the trace. We say T reduces to T' in n steps under φ if:

$$T \xrightarrow[\varphi]{p_1} \dots \xrightarrow[\varphi]{p_n} T'$$

for some p_1, \dots, p_n where each $p_i > 0$. A configuration ρ has terminated if the first thread in the pool is a value. We say that T is terminating in at most n steps under φ , if for all T' which T reduces to under φ for $n' \geq n$ steps, $\text{curr}(T')$ has terminated.

We now want to interpret this reduction relation as defining a distribution on program executions. However, in general, this would require measure theoretic probability to handle properly: even though our language only features sampling from countable discrete distributions, the set of all executions of a program is uncountable if the program does not necessarily terminate.⁹ However, if we restrict consideration to programs that terminate in a bounded number of steps, we can avoid these issues. Since most concurrency logics only handle partial correctness specifications anyway, this does not lead to much further loss of generality.

With this restriction in place, we can interpret program executions as indexed valuations (as we explained in §2, indexed valuations can be interpreted as probability distributions, and vice versa). Given a scheduler φ and a trace T , we first convert the trace step relation to an indexed valuation. Since the set of traces T' which T can step to under φ is countable, we can take the set of indices I to be any set in bijection with this set of traces. Take ind to be this bijection, and set $\text{val}(i)$ equal to the probability p such that $T \xrightarrow[\varphi]{p} \text{ind}(i)$. We refer to the resulting indexed valuation $(I, \text{ind}, \text{val})$ as

⁹A more denotational alternative, based on an approach due to [Kozen \[1981\]](#), is to interpret programs as monotone maps on sub-distributions of states. Then recursive commands are interpreted as least fixed points. However, since the original soundness proof of Iris is given in terms of a language with an operational semantics, we found it easier to use the semantics we describe in this section.

$$\begin{array}{c}
\text{ML-ALLOC} \\
\{\text{True}\} \text{ ref } v \{x. x \hookrightarrow v\} \\
\\
\text{ML-LOAD} \\
\{l \hookrightarrow v\} !l \{x. x = v \wedge l \hookrightarrow v\} \\
\\
\text{ML-STORE} \\
\{l \hookrightarrow v\} l := w \{l \hookrightarrow w\} \\
\\
\text{ML-FAA} \\
\{l \hookrightarrow n\} \text{ FAA}(l, k) \{x. x = n \wedge l \hookrightarrow n + k\} \\
\\
\text{HT-FRAME} \\
\frac{\{P\} e \{v. Q\}}{\{P * R\} e \{v. Q * R\}} \\
\\
\text{HT-CSQ} \\
\frac{P \Rightarrow P' \quad \{P'\} e \{v. Q'\} \quad \forall v. Q' \Rightarrow Q}{\{P\} e \{v. Q\}} \\
\\
\text{ML-FORK} \\
\frac{P \Rightarrow Q_0 * Q_1 \quad \{Q_0\} e \{\text{True}\} \quad \{Q_1\} e' \{R\}}{\{P\} \text{ fork}\{e; e'\} \{R\}}
\end{array}$$

Fig. 9. Selection of rules from Iris.

$\text{tstep}_\varphi^n(T)$. For each n , we define the indexed valuation $\text{resStep}_\varphi^n(T)$ recursively by:

$$\begin{aligned}
\text{resStep}_\varphi^0(T) &\triangleq \text{ match curr}(T) \text{ with } ([e_1, \dots], \sigma) \Rightarrow \text{ret } e_1 \text{ end} \\
\text{resStep}_\varphi^{n+1}(T) &\triangleq T' \leftarrow \text{tstep}_\varphi(T); \text{resStep}_\varphi^n(T')
\end{aligned}$$

This corresponds to stepping the trace n times and returning the first thread from the final configuration of the resulting trace. We regard the “return value” of a concurrent program to be the value that the first thread evaluates to, so in the event that the program terminates in n steps under the scheduler, $\text{resStep}_\varphi^n(T)$ gives this return value.

3.2 Background on Iris

As we have mentioned, Polaris is an extension of Iris, a recent concurrency logic with many expressive features. For reasons of space, we cannot explain all of Iris. We refer the reader to the Iris papers and manual [Jung et al. 2016, 2015; Krebbers et al. 2017; Team 2017] for a full account. Instead, we will just describe some essential aspects needed to understand our extensions and examples.

Figure 9 shows the basic concurrent separation logic rules of Iris. (Treat the \Rightarrow connective as just a kind of implication for now, we explain its use below.) These rules are used to establish triples of the form:

$$\{P\} e \{x. Q\}$$

which imply that if e is executed in a state that initially satisfies P , and it terminates with value v , then the terminating state will satisfy $[v/x]Q$. Furthermore, at no point will e go wrong and reach a stuck state, and neither will any of the threads forked by e during its execution. This is a partial correctness property: the post-condition only holds under executions where e terminates.

The fundamental idea of separation logic [Reynolds 2002] is the separating conjunction $P * Q$, which says that the program heap can be split into two disjoint pieces satisfying P and Q respectively. Thus, assertions are interpreted as claims of *ownership of resources*, where a “resource” is just a fragment of the heap. In particular, the $l \hookrightarrow v$ assertion claims ownership of a part of the heap containing the location l , and moreover says that l maps to the value v . Ownership of this resource licenses a thread to access and modify this location (see `ML-LOAD` and `ML-STORE` in Figure 9).

Like other recent concurrency logics [Appel 2014; Dinsdale-Young et al. 2013; Nanevski et al. 2014], Iris lets users of the logic extend the notion of resource beyond just heap fragments. These user-defined resources let us model the complex protocols that govern how threads access shared

$$\begin{array}{c}
\text{COUNTGEQ} \\
\boxed{\bullet n}^Y * \boxed{\circ(q, n')}^Y \Rightarrow n \geq n' \\
\\
\text{COUNTPERM} \\
\boxed{\circ(q, n)}^Y * \boxed{\circ(q', n')}^Y \Rightarrow q + q' \leq 1 \\
\\
\text{COUNTALLO} \\
\text{True} \Rightarrow \exists \gamma. \boxed{\bullet n}^Y * \boxed{\circ(1, n)}^Y \\
\\
\text{COUNTUPD} \\
\boxed{\bullet(n+k)}^Y * \boxed{\circ(q, n)}^Y \Rightarrow \boxed{\bullet(n'+k)}^Y * \boxed{\circ(q, n')}^Y \\
\\
\text{COUNTSEP} \\
\boxed{\circ(q, n)}^Y * \boxed{\circ(q', n')}^Y \Leftrightarrow \boxed{\circ(q+q', n+n')}^Y
\end{array}$$

Fig. 10. Counter resource rules.

$$\begin{array}{c}
\text{INV-ALLOC} \\
P \Rightarrow \exists \iota. \boxed{P}^\iota \\
\\
\text{INV-DUP} \\
\boxed{P}^\iota \Rightarrow \boxed{P}^\iota * \boxed{P}^\iota \\
\\
\text{INV-OPEN} \\
\frac{\text{(additional side conditions omitted)} \quad \{P * R\} e \{v. P * Q\} \quad \text{atomic}(e)}{\boxed{\boxed{P}^\iota * R} e \{v. Q\}}
\end{array}$$

Fig. 11. Invariant rules.

state in a concurrent system. Instead of describing the machinery that makes this work, let us give an example of one of these user-defined resources: the abstract “counter” resource. There are two types of these counter resources, represented by the following assertions:

$$\boxed{\bullet n}^Y \quad \text{and} \quad \boxed{\circ(q, n')}^Y$$

where n and n' are natural numbers, $0 < q \leq 1$ is a rational number, and γ is an abstract name assigned to a particular counter. The $\bullet n$ resource represents a shared counter that contains the value n . If we think of such a counter as being composed of n “units”, then the resource $\circ(q, n')$ represents a “stake” or ownership of n' of the units in the global counter. The parameter q is a fractional permission [Boyland 2003] that lets us track how many threads have such a stake; when $q = 1$, this represents full ownership, so no other threads have a stake.¹⁰

Rules for using these assertions are given in Figure 10. The rules **COUNTGEQ** and **COUNTSEQ** let us conclude that the global counter value must be at least as big as any stake’s value; and when a stake’s q value is 1, we furthermore know that the counter and the stake value are the same. The rule **COUNTSEP** lets us join (or conversely, split) two stakes by summing their permissions and their count values, subject to the (implicit) constraint that q , q' , and $q + q'$ all lie in the interval $(0, 1]$. The **COUNTALLO** rule lets us create a new counter with some existentially quantified name; the \Rightarrow connective here is a kind of implication in Iris which lets one modify or create a resource. Finally, **COUNTUPD** lets us modify a counter: if we own the global value and a stake, we can update the value and the stake, so long as we preserve the part of the counter value owned by other stakes (represented by k in the rule).

Of course, we need some way to connect these “abstract” resources to the actual state of the program. The mechanism for doing this is an *invariant*. An invariant is an assertion that is dynamically established at some point in the program, and then is guaranteed to hold thereafter. We write \boxed{P}^ι for the assertion which says that the invariant P has been established with the abstract name ι . If we have ownership of resources satisfying P , we can use the rule **INV-ALLOC** from Figure 11 to

¹⁰Note that the q is *not* the fraction of the global counter value represented by the stake’s value.

establish P as an invariant; we lose the resources and get back \boxed{P}^t with some fresh name t . If we know the invariant \boxed{P}^t holds and we are trying to prove some Hoare triple about an expression e , we can use **INV-OPEN** to “open” the invariant. This lets us add P to the pre-condition of the triple we are trying to prove, but we need to re-establish and give up P in the post-condition in order to “close” the invariant. Moreover, to use this rule, e must be atomic: Since e will reduce to a value in a single-step, this ensures there is no intermediate step in which the invariant did not hold. (In the statement of **INV-OPEN** we have omitted certain side conditions that are used to ensure that the same invariant is not opened multiple times simultaneously.)

For example, we can create the invariant $\boxed{\exists n. l \leftrightarrow n * \bullet \overline{n}}^t$ to ensure the physical heap location l will always store the value represented by the counter resource. Now a thread that owns a stake $\overline{(g, n')^y}$ can read from l or modify it using a compare-and-swap by opening the invariant and updating the global counter resource suitably with **COUNTUPD**. This resource and invariant pattern were used to verify a (non-approximate) concurrent counter in the Iris Coq development. As we shall see, we are able to use this same resource to verify the approximate counter with our extensions.

3.3 Probabilistic Rules

We now describe how we extend Iris with probabilistic relational reasoning to obtain Polaris. Our goal is to be able to prove that there exists a suitable coupling between the indexed valuation of a program and some set \mathcal{I} of indexed valuations. The existence of an appropriate coupling will let us use **Theorem 2.2**, so that we can bound the expected values of our program by bounding the extrema of \mathcal{I} .

To motivate the rules of our extension, let us first give some more background on how this kind of relational reasoning works in the pRHL logic of **Barthe et al.** The idea there, following **Benton’s** Relational Hoare Logic [**Benton 2004**], is to replace Hoare triples with Hoare *quadruples*, which replace pre and post-conditions with pre and post-relations about *pairs* of programs. Translated to our setting, we would have judgments of the form:

$$\{P\} e - \mathcal{I} \{x, y. Q\}$$

where e is the program we are trying to relate to \mathcal{I} , and in the post-relation Q , we would substitute the return value of e in for x and the return value of \mathcal{I} for y . Then, we would adapt the standard Hoare rules to consider the pairs of steps of e and \mathcal{I} . Although the work of **Barthe et al.** shows that this approach can be useful for reasoning about non-concurrent probabilistic programs, there is an issue with applying it in the concurrent setting: what do we do when e forks a new thread? We would then need to relate \mathcal{I} to multiple program expressions, and it’s not clear how to adapt **ML-FORK** to this quadruple style.

Instead, we adapt an idea originally developed by **Turon et al. [2013]** for non-probabilistic concurrent relational reasoning: rather than including the specification program \mathcal{I} as part of the Hoare judgment, we add a new assertion $\text{Prob}(\mathcal{I})$ to the logic. That is, the specification computation becomes just another kind of “resource” that can be transferred between threads. We then add the probabilistic rules shown in **Figure 12**. The rule **HT-COUPLE** lets us establish a triple about a $\text{flip}(n_1, n_2)$ command. The precondition requires us to own a monadic computation of the form $x \leftarrow \mathcal{I} ; F(x)$, and we must exhibit a coupling between a random choice between True and False, (weighted by $\frac{n_1}{n_2}$) and the monadic specification \mathcal{I} . The post condition says that we get back the monadic resource, but updated so that it is now of the form $\text{Prob}(F(v'))$ for some v' . In addition, v (the outcome of the $\text{flip}(n_1, n_2)$ command) and this v' are related by R , the postcondition of the

$$\begin{array}{c}
\text{HT-COUPLE} \\
\frac{0 \leq n_1/n_2 \leq 1 \quad (\text{ret True} \oplus_{\frac{n_1}{n_2}} \text{ret False}) \sim \mathcal{I} : R}{\{\text{Prob}(x \leftarrow \mathcal{I}; F(x))\} \text{flip}(n_1, n_2)} \\
\{v. \exists v'. \text{Prob}(F(v')) \wedge R(v, v')\} \\
\\
\text{HT-NONCOUPLE} \qquad \qquad \qquad \text{PROBLE} \\
\frac{0 \leq n_1/n_2 \leq 1}{\{\text{True}\} \text{flip}(n_1, n_2) \{v. (v = \text{True} \vee v = \text{False})\}} \qquad \frac{\mathcal{I}' \subseteq_p \mathcal{I}}{\text{Prob}(\mathcal{I}) \Rightarrow \text{Prob}(\mathcal{I}')}
\end{array}$$

Fig. 12. Probabilistic rules.

coupling we exhibited. This rule gives us a way to relate the execution of a concrete expression e to an execution of \mathcal{I} .

Rule **HT-NONCOUPLE** lets us handle a case where the concrete program executes a probabilistic choice yet we do not want to relate this to a step in the probabilistic specification. In this case, in the post-condition we merely know that the return value was True or False.

Finally, **PROBLE** lets us replace our \mathcal{I} resource with any \mathcal{I}' such that $\mathcal{I}' \subseteq_p \mathcal{I}$. We use this to manipulate the \mathcal{I} into a form that matches the precondition required by something like **HT-COUPLE**. Since $\mathcal{I}' \equiv \mathcal{I}$ implies $\mathcal{I}' \subseteq_p \mathcal{I}$, it follows that if $\mathcal{I}' \equiv \mathcal{I}$ then $\text{Prob}(\mathcal{I}) \Leftrightarrow \text{Prob}(\mathcal{I}')$.

Because $\text{Prob}(\mathcal{I})$ is just an assertion like any other, we can control access to it between threads by storing it in an invariant. This idea of representing a specification computation as a resource assertion has been used in other separation logics based on Iris [Krogh-Jespersen et al. 2017; Tassarotti et al. 2017].

3.4 Soundness

The following soundness theorem for the logic will guarantee that if we prove an appropriate triple involving $\text{Prob}(\mathcal{I})$, the expected value of the concrete program will lie in the range of the extrema of \mathcal{I} :

THEOREM 3.1. *Let $\mathcal{I} : M_{NI}(T)$ for some type T , and let $f : \text{Val} \rightarrow \mathbb{R}$, $g : T \rightarrow \mathbb{R}$, and assume that g is bounded on the support of \mathcal{I} . Suppose*

$$\{\text{Prob}(\mathcal{I})\} e \{v. \exists v'. \text{Prob}(\text{ret } v') \wedge f(v) = g(v')\}$$

holds. Let φ be a scheduler such that $([e], \sigma)$ terminates in at most n steps under φ . Then $\mathbb{E}_f[\text{resStep}_\varphi^n([e], \sigma)]$ exists and

$$\mathbb{E}_g^{\min}[\mathcal{I}] \leq \mathbb{E}_f[\text{resStep}_\varphi^n([e], \sigma)] \leq \mathbb{E}_g^{\max}[\mathcal{I}]$$

This result only holds for schedulers under which the program is guaranteed to terminate in some number of steps; this is not that surprising, since the original Iris is a partial correctness logic.

To prove this soundness theorem and validate the rules we have given, we first change the definition of the Hoare triple in Iris so that if the probabilistic resource is of the form $\text{Prob}(x \leftarrow \mathcal{I}; F(x))$ and the expression e takes a step, we must exhibit a coupling between e 's transition (interpreted as an indexed valuation) and \mathcal{I} . Then in the soundness proof, as e takes successive steps, we combine these couplings together using **BIND** from Figure 7; if e terminates and the post-condition matches the form stated in Theorem 3.1, then we will have constructed a complete

$$\begin{aligned}
& \text{ACOUNTERNEW} \\
& \{\text{Prob}(\text{approxN } n \ 0)\} \text{ ref } 0 \{l. \exists \gamma_l, \gamma_p, \gamma_c. \text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(l, 1, n)\} \\
& \text{ACOUNTERSEP} \\
& \text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(l, q + q', n + n') \Leftrightarrow \text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(l, q, n) * \text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(l, q', n') \\
& \text{ACOUNTERINCR} \\
& \{\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(l, q, n + 1)\} \text{ incr } l \{\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(l, q, n)\} \\
& \text{ACOUNTERREAD} \\
& \{\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(l, 1, 0)\} \text{ read } l \{v. \exists v'. \text{Prob}(\text{ret } v') \wedge v = v'\}
\end{aligned}$$

Fig. 13. Specification for approximate counters.

$$\text{countTrue } c \ lb \triangleq \text{foldLeft } (\lambda _ b. \text{if } b \text{ then } (\text{incr } c) \text{ else } ()) \ lb \ ()$$

$$\begin{aligned}
& \{\text{Prob}(\text{approxN } (|lb_1|_t + |lb_2|_t) \ 0)\} \\
& \text{let } c = \text{ref } 0 \text{ in} \\
& \{\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(c, 1, |lb_1|_t + |lb_2|_t)\} \\
& \{\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(c, 1/2, |lb_1|_t)\} \quad \parallel \quad \{\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(c, 1/2, |lb_2|_t)\} \\
& \text{countTrue } c \ lb_1 \quad \parallel \quad \text{countTrue } c \ lb_2 \\
& \{\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(c, 1/2, 0)\} \quad \parallel \quad \{\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(c, 1/2, 0)\} \\
& \{\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(c, 1, 0)\} \\
& \text{read } c \\
& \{v. \exists v'. \text{Prob}(\text{ret } v') \wedge v = v'\}
\end{aligned}$$

Fig. 14. Example client using approximate counters.

coupling between $\text{resStep}_\varphi^n([e], \sigma)$ and the monadic specification. Moreover, this will be an R -coupling with $R(x, y) \triangleq f(x) = g(y)$. Hence, we can apply [Theorem 2.2](#) to conclude the claim about the expected values.

4 EXAMPLE 1: APPROXIMATE COUNTERS

In this section we prove triples that relate the approximate counter algorithm from [Figure 1c](#) to the monadic computation approxN from [Figure 3](#).

4.1 Triples and Example Client

The Hoare triples we have proved about this data structure are given in [Figure 13](#). The specification uses a predicate $\text{ACounter}_{\gamma_l, \gamma_p, \gamma_c}(l, q, n)$, which can be treated by a user as an abstract predicate representing the permission to perform n increments to the counter at l . The parameter q is a fractional permission that we use to track how many threads can access the counter. (Ignore the names γ_l , γ_p , and γ_c – we will describe how they are used when we give the definition of ACounter later). The triple [ACOUNTERNEW](#) says that we can create a new counter by allocating a reference cell

$$\begin{aligned}
\text{LocInv}_{\gamma_1}(l) &\triangleq \exists n. l \hookrightarrow n * \overset{\circ}{\bullet} \overset{\circ}{n}^{\gamma_1} \\
\text{ProbInv}_{\gamma_p, \gamma_c} &\triangleq \exists n_1, n_2. \overset{\circ}{\bullet} \overset{\circ}{n_1}^{\gamma_p} * \overset{\circ}{\bullet} \overset{\circ}{n_2}^{\gamma_c} * (\text{Prob}(\text{approxN } n_1 \ n_2) \vee \overset{\circ}{\circ} \overset{\circ}{(1, n_1)}^{\gamma_p}) \\
\text{ACounter}_{\gamma_1, \gamma_p, \gamma_c}(l, q, n) &\triangleq \exists l_1, l_2, n'. \boxed{\text{LocInv}_{\gamma_1}(l)}^{l_1} * \boxed{\text{ProbInv}_{\gamma_p, \gamma_c}}^{l_2} * \overset{\circ}{\circ} \overset{\circ}{(q, n')}^{\gamma_1} * \overset{\circ}{\circ} \overset{\circ}{(q, n)}^{\gamma_p} \\
&\quad * \overset{\circ}{\circ} \overset{\circ}{(q, n')}^{\gamma_c}
\end{aligned}$$

Fig. 15. Invariants and definitions for proof.

containing 0. It takes the monadic specification $\text{Prob}(\text{approxN } n \ 0)$ as a precondition, and returns the full **ACounter** permission for n increments. The rule **ACOUNTERSEP** lets us split or join this **ACounter** permission into pieces. If we have permission to perform at least one increment, we can use **ACOUNTERINCR**, which gives us back **ACounter** with permission to do one fewer increment. Finally, if we have **ACounter** with the full fractional permission 1, and there are 0 pending increments, we can use **ACOUNTERREAD**. In the post condition we get back $\text{Prob}(\text{ret } v)$, where v is the value that the call to read returns.

At first this specification seems weak, but this is exactly what we need for **Theorem 3.1**. To see how we can use these triples to reason about a client program that uses the approximate counter, consider the example client in **Figure 14**. We start with a helper function `countTrue`, which takes an approximate counter c and a list of booleans lb , and counts the number of times `True` occurs in lb using the counter. The client begins by creating a new counter c . It then runs two threads in parallel that run `countTrue` on two lists lb_1 and lb_2 , using the shared counter c – we denote this parallel composition using \parallel . The parent blocks until both threads finish and then reads from the counter.¹¹

Refer to this client code as e . If we write $|lb|_t$ for the logical function giving the number of times `True` occurs in lb , then we would like to show that in expectation, e returns $|lb_1|_t + |lb_2|_t$. The derivation in **Figure 13** shows that the triple

$$\{\text{Prob}(\text{approxN } (|lb_1|_t + |lb_2|_t) \ 0)\} e \{v. \exists v'. \text{Prob}(\text{ret } v') \wedge v = v'\}$$

holds. Moreover, it is not hard to show that for each k , there is an upper bound on the value returned by `approxN` k , so by **Theorem 3.1** we have:

$$\mathbb{E}_{\text{id}}^{\min}[\text{approxN } (|lb_1|_t + |lb_2|_t)] \leq \mathbb{E}_{\text{id}}[\text{resStep}_{\varphi}^n([e], \sigma)]$$

(and similarly for \mathbb{E}^{\max}) for suitable φ and n . And, we have shown that $\mathbb{E}_{\text{id}}^{\min}[\text{approxN } (|lb_1|_t + |lb_2|_t)] = |lb_1|_t + |lb_2|_t$ in **§2.4**, so we are done.

4.2 Proofs of Triples

The definition of **ACounter** and the invariants used in the proof are given in **Figure 15**. The proof uses three counter resources to track: (1) the number of increments left to perform in the monadic specification, (2) the accumulated count in the monadic specification, and (3) the actual count currently stored in the concrete program. We use two invariants to connect the counter resources to these intended interpretations. First, we have $\text{LocInv}_{\gamma_1}(l)$ which says that the counter resource

¹¹Of course, here the threads may as well maintain their own exact counters and combine them at the end. But in a real application such as [Tristan et al. 2015], there are tens of millions of counters and hundreds of threads, so having each thread maintain its own set of counters would be expensive.

named γ_1 stores some value n and the physical location l points to that same value n . Then, assertion $\text{Problnv}_{\gamma_p, \gamma_c}$ says that there are two counter resources containing some n_1 and n_2 , and the invariant either contains (a) the monadic specification resource $\text{Prob}(\text{approxN } n_1 \ n_2)$ (i.e., there are n_1 further increments to perform, and the monadic counter has accumulated a value of n_2) or (b) it contains the complete stake for one of the counter resources. Then ACounter says that these two invariants have been set up with some names, and we own a stake in the γ_p permission corresponding to the number of increments this permission allows. Further, for some n' there is a stake in the γ_1 and γ_c counters both equal to n' , which represents the total amount that this permission has been used to add to the counter.

We will only describe the proofs of ACOUNTERINCR and ACOUNTERREAD , since ACOUNTERNEW is straight-forward.

Proof of ACOUNTERINCR . Eliminating the existentials in the definition of ACounter , we get that the appropriate invariants have been set up and there is some n' -stake in γ_1 and γ_c , along with the $n + 1$ stake in γ_p . The first step of $\text{incr } l$ reads the value of l ; to perform this read the thread needs to own $l \hookrightarrow v$ for some v . To get this resource, it opens the $\text{LocInV}_{\gamma_1}(l)$ invariant; after completing the read, the $l \hookrightarrow v$ resource is returned to close the invariant. The code then takes the minimum of the value read and MAX , and binds this value to k .

It then performs $\text{flip}(1, k + 1)$. We want to use HT-COUPLE to couple this flip with the monadic code. To do so, we first open the invariant $\text{Problnv}_{\gamma_p, \gamma_c}$. We know this will contain $\{\bullet n'_1\}^{\gamma_p}$ and $\{\bullet n'_2\}^{\gamma_c}$ for some n'_1 and n'_2 , and either $\text{Prob}(\text{approxN } n'_1 \ n'_2)$ or a full stake $\{\circ(1, n'_1)\}^{\gamma_p}$. However, the latter is impossible because the $\text{ACounter}_{\gamma_1, \gamma_p, \gamma_c}(l, q, n + 1)$ resource entails ownership of $\{\circ(q, n + 1)\}^{\gamma_p}$, but $q + 1 > 1$, contradicting COUNTPERM . So, we obtain $\text{Prob}(\text{approxN } n'_1 \ n'_2)$. Now, by COUNTGEQ we know that $n'_1 \geq n + 1$, hence we can unfold $\text{approxN } n'_1 \ n'_2$ to obtain $\text{Prob}(k \leftarrow \text{approxIncr}; \text{approxN } (n'_1 - 1) \ n'_2)$.

We can now use HT-COUPLE so long as we can exhibit a coupling between the concrete program's coin flip and approxIncr . First, since $0 \leq k \leq \text{MAX}$, we can show that:

$$\begin{aligned} & (\text{ret } k + 1) \oplus_{\frac{1}{k+1}} (\text{ret } 0) \\ & \quad \subseteq (x \leftarrow \text{ret } 0 \cup \dots \cup \text{ret } \text{MAX}; (\text{ret } x + 1 \oplus_{\frac{1}{x+1}} \text{ret } 0)) \\ & \quad \equiv \text{approxIncr} \end{aligned}$$

hence by EQUIV , it suffices to exhibit a coupling between $(\text{ret } \text{True} \oplus_{\frac{1}{k+1}} \text{ret } \text{False})$ and $(\text{ret } k + 1 \oplus_{\frac{1}{k+1}} \text{ret } 0)$. Take $R(x, y)$ to be $(x = \text{True} \wedge y = k + 1) \vee (x = \text{False} \wedge y = 0)$, then we can use P-CHOICE and RET to prove the existence of an R -coupling.

Applying HT-COUPLE with this coupling, we then have $\text{Prob}(\text{approxN } (n'_1 - 1) \ (n'_2 + v'))$ where v' and the return value v of the $\text{flip}(1, k)$ are related by R . We use COUNTUPD to update the thread's stake in γ_p resource to n , and the global value to $n'_1 - 1$ (to record that a simulated increment has been performed), similarly, we update the thread's stake in the γ_c counter to $n' + v'$ and the global value to $n'_2 + v'$ (to record the new total) and then close the $\text{Problnv}_{\gamma_p, \gamma_c}$ invariant.

The code then cases on the value v returned by the flip. If it is false, then v' is 0, the code returns, and the post condition holds. If v is true, then $v' = k + 1$, the amount that the code adds using a fetch-and-add. We therefore open the $\text{LocInV}_{\gamma_1}(l)$ invariant again to get access to l , perform the increment and update the γ_1 counter and stake using COUNTUPD to record the fact that we are adding $k + 1$.

Proof of ACOUNTERREAD . The precondition $\text{ACounter}_{\gamma_1, \gamma_p, \gamma_c}(l, 1, 0)$ represents the full stake in each counter, and the 0 argument means there are no pending increments to perform. Thus, when we

open the $\text{LocInV}_{\gamma_l}(l)$ and $\text{ProbInV}_{\gamma_p, \gamma_c}$ invariants we know that for some n' , $l \hookrightarrow n'$ and we have $\text{Prob}(\text{ret } n')$. So, we can read from l , knowing the returned value will be n' . After reading, we must close the invariant. This time we will keep the $\text{Prob}(\text{ret } n')$ resource so that we can put it in the post condition, instead we give up $[\text{O}(\overline{1, 0})]_{\gamma_p}^{\gamma_c}$ to satisfy the disjunction in $\text{ProbInV}_{\gamma_p, \gamma_c}$.

4.3 Variations

In our mechanized proofs, we have verified two additional variations on this approximate counter example. For the first variation, we consider a version of `incr` which directly uses the current value it reads from the counter, rather than taking the minimum of this value and `MAX`.

For the second variation, we address a limitation of the specification we have described so far. Notice that to use the rules in [Figure 13](#) and obtain a suitable triple to use with [Theorem 3.1](#), the total number of calls to `incr` must be a deterministic function of the program: we have to pick some n when we initialize the counter using `ACOUNTERNEW`. In the case of our example client, we chose n to be the number of times that `true` occurred in the two lists. But what if the number of calls to increment is itself probabilistic or non-deterministic? In this case we still would like to know that the expected value returned by the approximate counter is equal to the expected number of times the counter was incremented. However, if the number of times the counter is incremented is completely arbitrary, this expected value may not exist! To guarantee that the expected value will exist, our specification imposes an upper bound on the total number of increments that can be performed, and then allows us to establish a coupling with the following monadic computation:

$$\begin{aligned} \text{approxN}' \ 0 \ t \ l &\triangleq \text{ret } (t, l) \\ \text{approxN}' \ (n + 1) \ t \ l &\triangleq (\text{ret } (t, l)) \cup (k \leftarrow \text{approxIncr}; \text{approxN}' \ n \ (t + 1) \ (l + k)) \end{aligned}$$

The first argument of `approxN'` gives an upper bound on the remaining number of increments that can be performed, the second argument t tracks the total number of increments that have been done so far, and l again tracks the current value in the counter. In contrast to the original `approxN`, before performing each increment, there is a non-deterministic choice to simply return (t, l) . Let f be the function $\lambda(x, y). x - y$. We prove that

$$\mathbb{E}_f^{\min}[\text{approxN}' \ n \ 0 \ 0] = \mathbb{E}_f^{\max}[\text{approxN}' \ n \ 0 \ 0] = 0$$

i.e., the expected value of the difference between the total number of increments and the value in the counter is 0. We have proved more flexible versions of the rules in [Figure 13](#) that use this `approxN'` instead.

5 EXAMPLE 2: CONCURRENT SKIP LIST

For our next example, we verify properties of a probabilistic concurrent skip list. The code and proofs for this example are more complex, so for space reasons we will give a high-level description of the algorithm and the triples we have proved about it.

5.1 Implementation

A skip list [[Pugh 1990](#)] is a data structure that can be used to implement a dynamic set interface for ordered data. The implementation we consider will only allow integer keys to be stored in the set. The skip list consists of several sorted linked lists, where the nodes in each list contain a key. We visualize each list as running horizontally from left to right, with the different lists stacked vertically above one another (see [Figure 16](#)). For simplicity we only allow 2 lists in our implementation – this still exposes most of the main concurrency issues involved. The set of keys contained in the top list is a subset of the keys contained in the bottom list, and the node containing a key k in the top list includes a pointer to the corresponding node for k in the list below it. At the beginning and ends of

each list, there are sentinel nodes containing the minimum and maximum representable integer (which we write as $-\infty$ and $+\infty$ in Figure 16).

Non-concurrent Implementation. We first describe how operations on this data structure are implemented in the sequential case. To check whether a key k is contained in the set, we first search for the key in the top list starting at the left sentinel. If we find a node containing it, we return true. If not, we stop at the largest key $j < k$ in the list, and then follow the pointer in j 's node to the copy of j in the bottom list. We then resume searching for k starting at node j in the bottom list. If k is found in the bottom list we return true, otherwise the key is not in the set so we return false. To insert k , we first find the nodes N_t and N_b with the largest keys $\leq k$ in the top and bottom list, respectively. If we find that k is already in either list, we stop and return. Otherwise we execute `flip(1/2)`. If it returns true, we insert new nodes for key k into both the top and bottom lists, after N_t and N_b . Otherwise, if it returns false, we only insert a node in the bottom list after N_b . We call N_t and N_b the “predecessor nodes”, because they become the predecessors of k if it is inserted into each list.

If we insert n keys into the set, then in expectation $n/2$ of them will appear in the top list. Then when searching for a key, we will be able to more quickly descend down the top list, and either find the key there, or if not, only have to examine a few additional nodes in the bottom list. Of course, it is possible (though unlikely) that none or all of the nodes are inserted into the top list, in which case we are effectively searching in a regular sorted linked list. Later on, we will show how to use our program logic to derive a bound on the expected number of comparisons needed to find a key in the list. We will not handle deletion in our implementation, because if an adversarial client can observe the state of the list, it can repeatedly delete and re-insert any key that happens to end up in the top list, forcing the top list to be empty.

Adding Concurrency. There are several ways to add support for concurrent operations to a skip list. We will consider a simplified implementation inspired by that of Herlihy et al. [2006]. We add a lock to each node in the lists. Checking for whether a key is in the set is the same as in the non-concurrent case, and no locks need to be acquired.

To insert a key k , we again search for the predecessor nodes N_t and N_b . When we identify one of these nodes, we acquire its lock and then check that the node after it has not changed in the time between when we examined its successor and when the lock was acquired. If it has, that means another thread may have inserted a new node with a larger key less than k , so we release the lock and search for the predecessor again. Otherwise, so long as we hold the locks, we are guaranteed that N_t and N_b will remain the proper predecessors for key k . Having acquired both locks, we proceed as in the sequential case by generating a random bit, and on the basis of that bit we insert new nodes for k into either both lists or just the bottom list. We then release the locks and return.

What effect does concurrency have on the number of nodes that must be examined to find a key? None, so long as there are no concurrent insertions happening while searching. The reason is that in the implementation we have just described, the random choice is made *after* acquiring the locks for insertion. Thus, at the point the random choice is made, the ordering of operations by threads cannot affect where the node will be inserted.

However, to illustrate the subtleties involved, consider the following variant. If we generate the random bit *before* acquiring locks for the predecessors, and the resulting bit says we will only insert the node in the bottom list, then we only need to acquire the lock for the bottom predecessor. Now the distribution can be affected by the scheduler. To see why, imagine two threads are trying to concurrently insert key k into the list. Suppose that the outcome of the first thread's random bit generation indicates that it will insert the node only into the bottom list, but the second thread will try to insert into both lists. Then the scheduler can influence the distribution by pausing the

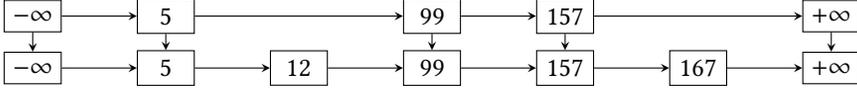


Fig. 16. Diagram for 2-Level Concurrent Skip List

second thread and letting the first thread finish; when the second thread is eventually allowed to run, it will find that k is already in the list and so it will return without doing anything. Although we have once more used adversarial language to describe the scheduler, in this case such behavior could arise without having to imagine any malice. Because the first thread only has to acquire a single lock, it is plausible that it might tend to finish before the second thread. In the following, we will just analyze the original version that we described.

5.2 Monadic Model

We follow the same pattern as in our verification of the approximate counter example: we first define a monadic model of the data structure, bound appropriate expected values of the monadic computation, and then describe triples that can be used to prove the existence of a coupling between programs using the skip list and the monadic model.

Our monadic model is the following:

$$\begin{aligned}
 \text{skiplist nil } tl \ bl &\triangleq \text{ret } (\text{sort}(tl), \text{sort}(bl)) \\
 \text{skiplist } (k :: l) \ tl \ bl &\triangleq k' \leftarrow \bigcup_{i \in k :: l} \text{ret } i; \\
 &tl' \leftarrow (\text{ret } tl) \oplus_{1/2} (\text{ret } k' :: tl); \\
 &\text{skiplist } (\text{remove } k' (k :: l)) \ tl' \ (k' :: bl)
 \end{aligned}$$

The computation $\text{skiplist } l \ tl \ bl$ simulates adding keys from the list l to a skip list, where the arguments tl and bl represent the keys in the top and bottom lists of the skip list, respectively. If the first argument l is empty, it sorts tl and bl and returns the result. If l is non-empty, it first non-deterministically selects a key k' from l . Then, with probability $1/2$ it adds this key to tl . It then removes any copies of k' from l , and recurses to process the remaining elements with the updated top and bottom lists. (There is no point in keeping the arguments tl and bl sorted throughout the recursive calls in this monadic formulation.)

There are many quantitative properties of the skip list that one might want to analyze (*e.g.*, the total number of nodes in both lists, the probability that a large fraction of nodes lie only in the bottom list, *etc.*). As we alluded to above, we will bound the expected number of inequality comparisons required to test for membership of a key in the skiplist. We define a function $\text{skipcost}(tl, bl, k)$ which gives the number of comparisons needed to check if k is in the skip list when the elements in the top and bottom lists are tl and bl , respectively:

$$\begin{aligned}
 \text{topcost}(tl, k) &= 1 + |\{i \in tl \mid \text{INTMIN} < i < k\}| \\
 \text{maxbelow}(tl, k) &= \max(\{i \in tl \mid i < k\} \cup \{\text{INTMIN}\}) \\
 \text{botcost}(tl, bl, k) &= 1 + |\{i \in bl \mid \text{maxbelow}(tl, k) < i < k\}| \\
 \text{skipcost}(tl, bl, k) &= \begin{cases} \text{topcost}(tl, k) & \text{if } k \in tl \\ \text{topcost}(tl, k) + \text{botcost}(tl, bl, k) & \text{if } k \notin tl \end{cases}
 \end{aligned}$$

If the key k is in the top list, then the number of comparisons is 1 plus the number of elements in the list less than k ($\text{topcost}(tl, k)$). If k is not in the top list, then we must first still perform

the same number of comparisons while searching through the top list. Then we search in the bottom list starting from the largest key less than k that was in tl ($\text{maxbelow}(tl, k)$). The total number of comparisons in the second list is the number of keys between $\text{maxbelow}(tl, k)$ and k ($\text{botcost}(tl, bl, k)$).

We then bound $\mathbb{E}_{\text{skipcost}(-, -, k)}^{\max}[\text{skiplist } l \text{ nil nil}]$, to obtain an upper bound on the expected value of searching for a key k . Assuming l has no duplicates, the key k and all keys in l lie between INTMIN and INTMAX , and there are n keys less than k in l , we show that:

$$\mathbb{E}_{\text{skipcost}(-, -, k)}^{\max}[\text{skiplist } l \text{ nil nil}] \leq 1 + \frac{n}{2} + 2 \left(1 - \frac{1}{2^{n+1}}\right)$$

This means that on average we have to do about half the number of comparisons that would be required to search for the key in a regular sorted linked list.

5.3 Triples

Figure 17 shows the triples we have proved about the skip list. The specification defines an assertion $\text{SkipP}_{\Gamma}(q, v, S, S_t, S_b)$, which represents permission to access a skip list whose top left sentinel is v . The argument Γ is just a set of resource names (like the γ 's in the counter example), q is a fractional permission, S is the finite set of keys which may be added to the list, and S_t and S_b are a subset of the keys currently in the top and bottom lists. Additional keys from S may be in either S_t or S_b , but the owner of this permission assertion knows that they contain *at least* these sets.

The expression `newSkipList` creates a new skip list. The precondition for the triple in `SKIPNEW` requires us to own the monadic computation¹² $\text{Prob}(\text{skiplist } S \text{ nil nil})$. The post condition gives the full permission ($q = 1$) to access the skip list, with empty top and bottom lists. To use this rule, all the keys in S must be between INTMIN and INTMAX . Notice here that the set of keys S which will be added to the skiplist must be deterministic, so that it can be decided in this precondition (much like our original specification for the approximate counters required the total number of increments to be deterministic). This restriction is important: if the keys to be added are non-deterministically selected, and a client can observe the state of the skip list, it can insert a special sequence of keys in such a way so as to force a large number of comparisons to find a particular target key.

We use `addSkipList` to insert a key k into the skip list. The post condition in `SKIPADD` indicates that we now know that the added key k is in the bottom list. On the other hand, the client does not know whether the key was added to the top list or not, so in the permission for the post condition, the contents of the top list are given by some existentially quantified S'_t .

The function `memSkipList` checks whether a key is in the skip list. It returns a pair (b, z) , where b is a boolean indicating whether the key was in the set or not, and z is the number of key comparisons performed. The triple `SKIPMEM` says that if we have the full permission for the skip list, then the boolean b indeed reflects whether the key is in the set or not, and z is in fact equal to the cost function $\text{skipcost}(S_t, S_b, k)$ we defined above. (One can also prove triples for when we have less than the full permission of the list (*i.e.*, $q < 1$), but we have not done so.)

The rule `SKIPSEP` lets us split and join together the `SkipP` permission so that separate threads can use the skip list. Finally, `SKIPSHIFT` lets us do a view shift to convert a full `SkipP` permission in which we have added all the keys in S to the skip list back into a `Prob` permission in which the monadic computation is finished. This lets us prove triples of the form required by [Theorem 3.1](#).

¹² Here, S is a set, whereas the arguments to `skiplist` are lists. However, it is easy to show that if l' , tl' , and bl' are permutations of the lists l , tl , and bl , respectively, then `skiplist l tl bl` \equiv `skiplist l' tl' bl'`, so it makes no difference if we treat the first argument instead as an unordered set.

$$\begin{array}{c}
\text{SKIPNEW} \\
\frac{\forall k \in S. \text{INTMIN} < k < \text{INTMAX}}{\{\text{Prob}(\text{skiplist } S \text{ nil nil})\} \text{newSkiplist}} \\
\{v. \exists \Gamma. \text{SkipP}_{\Gamma}(1, v, S, \emptyset, \emptyset)\}
\end{array}
\qquad
\begin{array}{c}
\text{SKIPADD} \\
\frac{k \in S}{\{\text{SkipP}_{\Gamma}(q, v, S, S_t, S_b)\} \text{addSkiplist } v \ k} \\
\{\exists S'_t. \text{SkipP}_{\Gamma}(q, v, S, S'_t, S_b \cup \{k\})\}
\end{array}$$

$$\begin{array}{c}
\text{SKIPMEM} \\
\frac{\text{INTMIN} < k < \text{INTMAX}}{\{\text{SkipP}_{\Gamma}(1, v, S, S_t, S_b)\} \text{memSkiplist } v \ k} \\
\left\{ \begin{array}{l} \text{SkipP}_{\Gamma}(1, v, S, S_t, S_b) * (b = \text{True} \Rightarrow k \in S_b) \\ (b, z). * (b = \text{False} \Rightarrow k \notin S_t \cup S_b) * (z = \text{skipcost}(S_t, S_b, k)) \end{array} \right\}
\end{array}$$

$$\begin{array}{c}
\text{SKIPSEP} \\
\text{SkipP}_{\Gamma}(q + q', v, S, S_t \cup S'_t, S_b \cup S'_b) \\
\Leftrightarrow \text{SkipP}_{\Gamma}(q, v, S, S_t, S_b) * \text{SkipP}_{\Gamma}(q', v, S, S'_t, S'_b)
\end{array}
\qquad
\begin{array}{c}
\text{SKIPSHIFT} \\
\text{SkipP}_{\Gamma}(1, v, S, S_t, S) \Rightarrow \text{Prob}(\text{ret}(\text{sort}(S_t), \text{sort}(S)))
\end{array}$$

Fig. 17. Specification for skip list.

In our mechanized proofs we have used this specification to verify a simple client in which two threads concurrently add lists of integers to a skip list set, and then after they both finish, one looks up a key using `memSkiplist` and returns the number of comparisons performed.

6 RELATED WORK

Our logic used ideas from the related work described in §1. We now describe further related work.

Probabilistic Logics. [McIver et al. \[2016\]](#) present a probabilistic form of rely-guarantee logic [[Jones 1983](#)]. Like the original rely-guarantee logic, this logic does not permit local reasoning: one must check stability against rely-guarantee conditions that refer to the global state of the program. More recent concurrency logics have combined rely-guarantee reasoning with the local reasoning features of concurrent separation logic [[Feng 2009](#); [Vafeiadis and Parkinson 2007](#)]. Iris, and our extensions, incorporate these ideas, which is what enables us to give specifications like those in [Figure 17](#) – clients can use the skip list without having to reason about interference involving the underlying state of the list. Since many uses of randomization in the concurrent setting are in implementations of *data structures*, it is important to provide these more abstract specifications.

Recently, [Batz et al. \[2018\]](#) have developed a version of (non-concurrent) separation logic for reasoning about sequential probabilistic programs with dynamic memory allocation. They verify an example of a program which probabilistically appends nodes to a list (so that the length of the list is geometrically distributed), and a tree deletion procedure which only probabilistically deletes nodes. Instead of using relational reasoning, assertions in their logic denote probabilities/expected values, and rules are given for computing and bounding these probabilities.

Several program logics for probabilistic reasoning (e.g., [[Morgan et al. 1996](#)]) are designed to reason about languages that have primitives for both probabilistic choice and (demonic) non-deterministic choice. However, in that work, non-determinism was not used for modeling concurrency. Instead, it was used to model “underspecified” programs with several implementations of a component, of which one is selected non-deterministically, as in GCL [[Dijkstra 1975](#)].

[Barthe et al. \[2015\]](#) were the first to connect the idea of coupling to the kind of probabilistic relational reasoning done in pRHL, an earlier logic by [Barthe et al. \[2012\]](#). Since then, different results from the theory of coupling and variants of couplings have been used to extend pRHL [[Barthe et al. 2017a,b,c](#); [Hsu 2017](#)].

Iris uses a step-indexed semantic model to support impredicative features of the logic, but we do not make special use of step-indexing in our extensions. However, Aguirre et al. [2018] have shown that a kind of step-indexed model can be used to reason about more general kinds of couplings (so-called “shift couplings”). Previously, Bizjak and Birkedal [2015] developed a step-indexed logical relation for a higher-order language with random choice.

Denotational Semantics. A number of denotational models combining probabilistic and non-deterministic choice have been developed [Goubault-Larrecq 2015; Jones 1990; Mislove 2000; Tix et al. 2009; Varacca 2002; Varacca and Winskel 2006]. Our soundness theorem considers a scheduler which *deterministically* selects which thread to run next. Varacca and Winskel [2006] showed that their monadic encoding, which we have used in our work, gives an adequate semantic model for an imperative language with this kind of deterministic scheduler. An alternative is to permit the scheduler to also make random choices when selecting which thread to run. Varacca and Winskel show that in this case, an alternative monad developed by Mislove [2000] and Tix et al. [2009] gives an adequate model. It would be interesting to use this latter monad in our program logic to reason about behavior under probabilistic schedulers.

Linearizability and Quiescent Consistency. For non-randomized concurrent data structures, one important correctness criterion is *linearizability* [Herlihy and Wing 1990], which ensures that we can treat the execution of operations on the data structure as if they happened in atomic steps. Many program logics have been developed for establishing linearizability itself or related notions of atomicity [da Rocha Pinto et al. 2014; Frumin et al. 2018; Turon et al. 2013; Vafeiadis 2007]. However, Golab et al. [2011] show that when clients of data structures can make randomized choices, linearizable implementations can indeed be distinguished from versions that are truly atomic. They propose an alternative condition called *strong linearizability* which resolves this issue, but they do not consider randomized implementations of data structures, only randomized clients.

The bounds we have proved only apply to “stable” states of the data structure in which there are no on-going modifications. In the analysis of non-probabilistic concurrent data structures, there is a condition weaker than linearizability known as *quiescent consistency*, whose definition considers groups of operations separated by gaps of time in which no modifications occur. Sergey et al. [2016] have developed ways of reasoning about quiescent consistent data structures in a separation logic, and it might be possible to adapt their approach to the probabilistic setting.

7 CONCLUSION

We have developed a concurrent program logic that can be used for probabilistic relational reasoning, and have used it to verify two realistic examples of randomized concurrent algorithms. Moreover, we have mechanized all the results described here in Coq by modifying the prior Coq formalization of Iris. The development is available at <https://github.com/jtassarotti/polaris>.

ACKNOWLEDGMENTS

The authors thank Jean-Baptiste Tristan, Jan Hoffmann, Derek Dreyer, Guy L. Steele, Victor Luchangco, Jeremy Avigad, Jon Sterling, Justin Hsu, and Daniel Gratzer for feedback and discussions related to this work.

This work was supported by a gift from Oracle Labs. This research was conducted with U.S. Government support under and awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of these organizations.

REFERENCES

- Alejandro Aguirre, Gilles Barthe, Lars Birkedal, Ales Bizjak, Marco Gaboardi, and Deepak Garg. 2018. Relational Reasoning for Markov Chains in a Probabilistic Guarded Lambda Calculus. In *ESOP*. 214–241.
- Andrew W. Appel. 2014. *Program Logics - for Certified Compilers*. Cambridge University Press.
- Philippe Audebaud and Christine Paulin-Mohring. 2009. Proofs of randomized algorithms in Coq. *Sci. Comput. Program.* 74, 8 (2009), 568–589.
- Gilles Barthe, Thomas Espitau, Benjamin Grégoire, Justin Hsu, Léo Stefanescu, and Pierre-Yves Strub. 2015. Relational Reasoning via Probabilistic Coupling. In *Logic for Programming, Artificial Intelligence, and Reasoning - 20th International Conference, LPAR-20 2015, Suva, Fiji, November 24-28, 2015, Proceedings*. 387–401.
- Gilles Barthe, Thomas Espitau, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2017a. Proving uniformity and independence by self-composition and coupling. In *LPAR*.
- Gilles Barthe, Thomas Espitau, Justin Hsu, Tetsuya Sato, and Pierre-Yves Strub. 2017b. *-Liftings for Differential Privacy. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*. 102:1–102:12.
- Gilles Barthe, Marco Gaboardi, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2016. A Program Logic for Union Bounds. In *ICALP*. 107:1–107:15.
- Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. 2012. Probabilistic Relational Hoare Logics for Computer-Aided Security Proofs. In *Mathematics of Program Construction - 11th International Conference, MPC 2012, Madrid, Spain, June 25-27, 2012. Proceedings*. 1–6.
- Gilles Barthe, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2017c. Coupling proofs are probabilistic product programs. In *POPL*. 161–174.
- Kevin Batz, Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Thomas Noll. 2018. Quantitative Separation Logic. *CoRR* abs/1802.10467 (2018). arXiv:1802.10467 <http://arxiv.org/abs/1802.10467>
- Jon Beck. 1969. Distributive laws. In *Seminar on Triples and Categorical Homology Theory*, B. Eckmann (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 119–140.
- Nick Benton. 2004. Simple Relational Correctness Proofs for Static Analyses and Program Transformations. In *POPL*.
- Ales Bizjak and Lars Birkedal. 2015. Step-Indexed Logical Relations for Probability. In *Foundations of Software Science and Computation Structures - 18th International Conference, FoSSaCS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*. 279–294.
- Silas Boyd-Wickizer, Austin T. Clements, Yandong Mao, Aleksey Pesterev, M. Frans Kaashoek, Robert Tappan Morris, and Nikolai Zeldovich. 2010. An Analysis of Linux Scalability to Many Cores. In *9th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2010, October 4-6, 2010, Vancouver, BC, Canada, Proceedings*. 1–16.
- John Boyland. 2003. Checking Interference with Fractional Permissions. In *Static Analysis, 10th International Symposium, SAS 2003, San Diego, CA, USA, June 11-13, 2003, Proceedings*. 55–72.
- Adam Chlipala. 2013. *Certified Programming with Dependent Types - A Pragmatic Introduction to the Coq Proof Assistant*. MIT Press. <http://mitpress.mit.edu/books/certified-programming-dependent-types>
- Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. 2014. TaDA: A Logic for Time and Data Abstraction. In *ECOOP*. 207–231.
- Dave Dice, Yossi Lev, and Mark Moir. 2013. Scalable statistics counters. In *SPAA*. 43–52.
- Edsger W. Dijkstra. 1975. Guarded Commands, Nondeterminacy and Formal Derivation of Programs. *Commun. ACM* 18, 8 (1975), 453–457.
- Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew J. Parkinson, and Hongseok Yang. 2013. Views: Compositional reasoning for concurrent programs. In *POPL*.
- T. Dinsdale-Young, M. Dodds, P. Gardner, M. Parkinson, and V. Vafeiadis. 2010. Concurrent abstract predicates. In *ECOOP*. 504–528.
- Xinyu Feng. 2009. Local rely-guarantee reasoning. In *POPL*. 315–327.
- Xinyu Feng, Rodrigo Ferreira, and Zhong Shao. 2007. On the relationship between concurrent separation logic and assume-guarantee reasoning. In *ESOP*. 173–188.
- Philippe Flajolet. 1985. Approximate Counting: A Detailed Analysis. *BIT* 25, 1 (1985), 113–134.
- Keir Fraser. 2004. *Practical lock-freedom*. Ph.D. Dissertation. University of Cambridge.
- Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2018. ReLoC: A Mechanised Relational Logic for Fine-Grained Concurrency. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*. 442–451.
- Ming Fu, Yong Li, Xinyu Feng, Zhong Shao, and Yu Zhang. 2010. Reasoning about optimistic concurrency using a program logic for history. In *CONCUR*. 388–402.
- Jeremy Gibbons and Ralf Hinze. 2011. Just do it: simple monadic equational reasoning. In *ICFP*. 2–14.

- Wojciech M. Golab, Lisa Higham, and Philipp Woelfel. 2011. Linearizable implementations do not suffice for randomized distributed computation. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*. 373–382.
- Jean Goubault-Larrecq. 2007. Continuous Previsions. In *Computer Science Logic, 21st International Workshop, CSL 2007, 16th Annual Conference of the EACSL, Lausanne, Switzerland, September 11-15, 2007, Proceedings*. 542–557.
- Jean Goubault-Larrecq. 2015. Full abstraction for non-deterministic and probabilistic extensions of PCF I: The angelic cases. *J. Log. Algebr. Meth. Program.* 84, 1 (2015), 155–184.
- Maurice Herlihy, Yossi Lev, Victor Luchangco, and Nir Shavit. 2006. A Provably Correct Scalable Concurrent Skip List (Brief Announcement). In *OPODIS*.
- Maurice P. Herlihy and Jeannette M. Wing. 1990. Linearizability: a correctness condition for concurrent objects. *TOPLAS* 12, 3 (1990), 463–492.
- J. Hsu. 2017. Probabilistic Couplings for Probabilistic Reasoning. *ArXiv e-prints* (Oct. 2017). arXiv:cs.LO/1710.09951
- Claire Jones. 1990. *Probabilistic non-determinism*. Ph.D. Dissertation. University of Edinburgh, UK. <http://hdl.handle.net/1842/413>
- C. B. Jones. 1983. Tentative steps toward a development method for interfering programs. *TOPLAS* 5, 4 (1983), 596–619.
- Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-Order Ghost State. In *ICFP*.
- 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 *POPL*. 637–650.
- Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Federico Olmedo. 2016. Weakest Precondition Reasoning for Expected Run-Times of Probabilistic Programs. In *ESOP*. 364–389.
- Dexter Kozen. 1981. Semantics of Probabilistic Programs. *J. Comput. Syst. Sci.* 22, 3 (1981), 328–350.
- Robbert Krebbers, Ralf Jung, Ales Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017. The Essence of Higher-Order Concurrent Separation Logic. In *ESOP*. 696–723.
- Morten Krogh-Jespersen, Kasper Svendsen, and Lars Birkedal. 2017. A Relational Model of Types-and-Effects in Higher-Order Concurrent Separation Logic. In *POPL*. 218–231.
- T. Lindvall. 2002. *Lectures on the Coupling Method*. Dover Publications, Incorporated.
- Annabelle McIver, Tahiry M. Rabehaja, and Georg Struth. 2016. Probabilistic rely-guarantee calculus. *Theor. Comput. Sci.* 655 (2016), 120–134.
- Michael W. Mislove. 2000. Nondeterminism and Probabilistic Choice: Obeying the Laws. In *CONCUR 2000 - Concurrency Theory, 11th International Conference, University Park, PA, USA, August 22-25, 2000, Proceedings*. 350–364.
- Michael W. Mislove. 2006. On Combining Probability and Nondeterminism. *Electr. Notes Theor. Comput. Sci.* 162 (2006), 261–265. <https://doi.org/10.1016/j.entcs.2005.12.113>
- Carroll Morgan, Annabelle McIver, and Karen Seidel. 1996. Probabilistic Predicate Transformers. *ACM Trans. Program. Lang. Syst.* 18, 3 (1996), 325–353.
- Robert Morris. 1978. Counting Large Numbers of Events in Small Registers. *Commun. ACM* 21, 10 (1978), 840–842.
- Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. 2014. Communicating State Transition Systems for Fine-Grained Concurrent Resources. In *ESOP*. 290–310.
- Aleksandar Nanevski, J. Gregory Morrisett, and Lars Birkedal. 2008. Hoare type theory, polymorphism and separation. *J. Funct. Program.* 18, 5-6 (2008), 865–911. <https://doi.org/10.1017/S0956796808006953>
- P.W. O’Hearn. 2007. Resources, concurrency, and local reasoning. *TCS* 375, 1 (2007), 271–307.
- Adam Petcher and Greg Morrisett. 2015. The Foundational Cryptography Framework. In *POST*. 53–72.
- William Pugh. 1990. Skip Lists: A Probabilistic Alternative to Balanced Trees. *Commun. ACM* 33, 6 (1990), 668–676.
- Lyle Harold Ramshaw. 1979. *Formalizing the Analysis of Algorithms*. Ph.D. Dissertation. Stanford University.
- John C. Reynolds. 2002. Separation logic: A logic for shared mutable data structures. In *LICS*.
- Ilya Sergey, Aleksandar Nanevski, Anindya Banerjee, and Germán Andrés Delbianco. 2016. Hoare-style specifications as correctness conditions for non-linearizable concurrent objects. In *OOPSLA*. 92–110.
- Wouter Swierstra. 2009. A Hoare Logic for the State Monad. In *TPHOLs*. 440–451.
- Joseph Tassarotti, Ralf Jung, and Robert Harper. 2017. A Higher-Order Logic for Concurrent Termination-Preserving Refinement. In *ESOP*. 909–936.
- Iris Team. 2017. Iris 3.0 Documentation. <http://plv.mpi-sws.org/iris/appendix-3.0.pdf>
- Regina Tix, Klaus Keimel, and Gordon D. Plotkin. 2009. Semantic Domains for Combining Probability and Non-Determinism. *Electr. Notes Theor. Comput. Sci.* 222 (2009), 3–99.
- Jean-Baptiste Tristan, Joseph Tassarotti, and Guy L. Steele Jr. 2015. Efficient Training of LDA on a GPU by Mean-for-Mode Estimation. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*. 59–68.
- Aaron Turon, Derek Dreyer, and Lars Birkedal. 2013. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *ICFP*. 377–390.

- Viktor Vafeiadis. 2007. *Modular fine-grained concurrency verification*. Ph.D. Dissertation. University of Cambridge.
- V. Vafeiadis and M. Parkinson. 2007. A marriage of rely/guarantee and separation logic. In *CONCUR*. 256–271.
- Elis van der Weegen and James McKinna. 2008. A Machine-Checked Proof of the Average-Case Complexity of Quicksort in Coq. In *TYPES*. 256–271.
- Daniele Varacca. 2002. The Powerdomain of Indexed Valuations. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*. 299.
- Daniele Varacca and Glynn Winskel. 2006. Distributing probability over non-determinism. *Mathematical Structures in Computer Science* 16, 1 (2006), 87–113.

A APPENDIX

In this appendix, we describe in more detail how we modify the definition of Hoare triples in Iris in order to support our probabilistic extensions. Recall from the description in the paper that the main idea is to augment the definition of Hoare triples so that a derivation of a Hoare triple will encode a coupling between each step of the concrete program e and the monadic specification.

From here on, we assume familiarity with the Iris model. Let $ProbState$ be the type $\Sigma_{T:Type} M_N(M_1(T))$. Given two terms (T_1, \mathcal{I}_1) and (T_2, \mathcal{I}_2) of type $ProbState$, we say $(T_1, \mathcal{I}_1) \equiv (T_2, \mathcal{I}_2)$ if $T_1 = T_2$ and $\mathcal{I}_1 \equiv \mathcal{I}_2$. Using this equivalence relation, we impose a discrete OFE structure on $ProbState$. Here, we will often simply omit the type T when writing an element of $ProbState$.

The monad specification resource is handled much in the same way that physical state is in Iris. We represent the monoid specification code using the “authoritative exclusive resource” construction. Define:

$$P\text{Interp}(\mathcal{I}) \triangleq \bullet \text{Ex } \mathcal{I} \qquad \text{Prob}(\mathcal{I}) \triangleq \exists \mathcal{I}'. \mathcal{I} \subseteq_p \mathcal{I}' * \circ \text{ex}(\mathcal{I}')$$

In Iris, Hoare triples are defined in terms of weakest-preconditions, so we actually need to modify the definition of the latter. Recall the following definition of weakest-precondition in Iris 3.0 [Team 2017], which is indexed by a state interpretation function S .

$$\begin{aligned} \text{wp}^S \triangleq & \mu \text{wp}. \lambda \mathcal{E}, e, \varphi. \\ & (\exists v. \text{expr2val}(e) = v \wedge \Vdash_{\mathcal{E}} \varphi(v)) \vee \\ & \left(\text{expr2val}(e) = \text{None} \wedge \forall \sigma. S(\sigma) \text{ -*} \right. \\ & \quad \mathcal{E} \Vdash^{\emptyset} \left(\text{red}(e, \sigma) * \triangleright \forall e', \sigma', T. (e, \sigma \rightarrow e', \sigma', T) \text{ -*} \right. \\ & \quad \quad \left. \left. \left. \emptyset \Vdash^{\mathcal{E}} \left(S(\sigma') * \text{wp}(\mathcal{E}, e', \varphi) * \bigstar_{e'' \in T} \text{wp}(\top, e'', \lambda _ . \text{True}) \right) \right) \right) \end{aligned}$$

The definition is a guarded fixed-point, composed of a disjunction which handles two cases: (1) either the expression e is value, in which case the post-condition φ should hold for that value, or (2) it is not a value, in which case for each possible state σ , given the interpretation of σ , we need to show e is reducible, and then recursively show that for each thing which e could step to, we will be able to update the state interpretation appropriately and recursively prove weakest-precondition for the reduct.

We write $\text{primStep}(e, \sigma)$ for the indexed valuation of type $\text{Option}(\text{Expr} \times \text{State} \times \text{List Expr})$ which returns reducts of $e; \sigma$ or None if $e; \sigma$ is not reducible.

