



HAL
open science

Coupling proofs are probabilistic product programs

Gilles Barthe, Benjamin Grégoire, Justin Hsu, Pierre-Yves Strub

► **To cite this version:**

Gilles Barthe, Benjamin Grégoire, Justin Hsu, Pierre-Yves Strub. Coupling proofs are probabilistic product programs. POPL 2017 - Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, Jan 2017, Paris, France. pp.161-174, 10.1145/3009837.3009896 . hal-01649028

HAL Id: hal-01649028

<https://hal.science/hal-01649028>

Submitted on 27 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Coupling proofs are probabilistic product programs

Gilles Barthe¹ Benjamin Grégoire² Justin Hsu³ Pierre-Yves Strub¹
IMDEA Software¹ Inria² University of Pennsylvania³

Abstract

Couplings are a powerful mathematical tool for reasoning about pairs of probabilistic processes. Recent developments in formal verification identify a close connection between couplings and pRHL, a relational program logic motivated by applications to provable security, enabling formal construction of couplings from the probability theory literature. However, existing work using pRHL merely shows existence of a coupling and does not give a way to prove quantitative properties about the coupling, which are needed to reason about mixing and convergence of probabilistic processes. Furthermore, pRHL is inherently incomplete, and is not able to capture some advanced forms of couplings such as shift couplings. We address both problems as follows.

First, we define an extension of pRHL, called \times pRHL, which explicitly constructs the coupling in a pRHL derivation in the form of a probabilistic product program that simulates two correlated runs of the original program. Existing verification tools for probabilistic programs can then be directly applied to the probabilistic product to prove quantitative properties of the coupling. Second, we equip \times pRHL with a new rule for **while** loops, where reasoning can freely mix synchronized and unsynchronized loop iterations. Our proof rule can capture examples of *shift couplings*, and the logic is relatively complete for deterministic programs.

We show soundness of \times pRHL and use it to analyze two classes of examples. First, we verify rapid mixing using different tools from coupling: standard coupling, shift coupling, and *path coupling*, a compositional principle for combining local couplings into a global coupling. Second, we verify (approximate) equivalence between a source and an optimized program for several instances of loop optimizations from the literature.

1. Introduction

The *coupling method* [41, 42, 54, 55] is an elegant mathematical tool for analyzing the relationship between probabilistic processes. Informally, couplings correlate outputs of two probabilistic processes by specifying how corresponding sampling statements are correlated; reasoning about the correlated processes can then imply interesting properties of the original processes.

A classical application of couplings is showing that two probabilistic processes converge in distribution. Consider, for instance, a symmetric simple random walk over \mathbb{Z} : starting from some initial position p , the process repeatedly samples a value s uniformly in

$\{-1, 1\}$ and updates its position to $p + s$. As the process evolves, the distribution on position spreads out from its initial position, and converges to a limit distribution that is the same for all initial positions. This property can be proved by constructing a coupling where the probability that the two coupled walks end in the same position approaches 1 as we run process.

Beyond merely showing convergence, typically we are interested in *how quickly* the processes converge. The rate of convergence, sometimes measured by the *mixing time* of a process, is highly important [41]. A common use of probabilistic processes is to efficiently sample from a distribution that approximates a complicated distribution (e.g., Metropolis et al. [43]). For instance, the Glauber dynamics [33] approximately samples from the uniform distribution on proper colorings of a graph—a hard distribution to compute—by maintaining a coloring and randomly re-coloring a single vertex at a time. More generally, the Glauber dynamics is an example of *Markov Chain Monte-Carlo* (MCMC), a family of techniques that underlie many computational simulations in science, and machine learning algorithms for performing probabilistic inference (see Brooks et al. [22] for a survey). The mixing time determines how many steps we need to run the process before we are approximately sampling from the target distribution.

Recently, Barthe et al. [12] noticed a close connection between couplings and probabilistic relational Hoare logic (pRHL), originally designed for reasoning about the computational security of cryptographic constructions [5]. Namely, every valid derivation in pRHL implies the existence of a coupling of the output distributions of the two programs. Using this observation, Barthe et al. [12] verify pRHL judgments that imply properties of random walks, the balls-in-bins process, and the birth-death process. Barthe et al. [15] extend the connection to approximate couplings in approximate probabilistic relational Hoare logic (apRHL), introduced for reasoning about differentially private computations [7], and exploit this connection to prove differential privacy of several examples that had previously been beyond the reach of formal verification.

While pRHL is a useful tool for constructing couplings, it has two limitations. First, it cannot directly reason about the two coupled processes. This poses a problem for proving mixing and convergence properties. For instance, Barthe et al. [12] prove that if a certain property P holds on the coupled samples, then two random walks meet under the coupling. By a theorem about random walks, this means that the (total variation) distance between the two distributions is at most the probability of P under the coupling. However, we do not know what this probability is or how it grows as we run the random walk for more iterations, since pRHL cannot reason about the coupled process. We run into similar difficulties if we try to prove convergence using *path coupling* [23], a general construction that shows fast convergence by upper-bounding the expected distance when we make a transition from two coupled states. While we can express the transition function in code, pRHL cannot reason about expectations.

Second, pRHL cannot express some natural classes of couplings because the rule for **while** loops requires that both loops execute in lockstep. For instance, *shift couplings* [1] allow the two processes to meet at a random time shift, e.g., the first process at time t_1 could track the second process at some other time $t_2 = t_1 + \delta$. This incompleteness of pRHL is also a limitation for more standard applications of pRHL to program analysis, like validation of loop transformations, and more generally for any optimization that alters the control flow of the programs.

We address these problems as follows. First, we deepen the connection between pRHL and couplings with an observation that is reminiscent of proofs-as-programs: not only do pRHL judgements correspond to couplings, but pRHL *proofs* encode a *probabilistic product program* that constructs the distribution witnessing the coupling. This program is similar to existing product program constructions [6, 57] in that it simulates two program executions with a single program. However, the probabilistic product also coordinates the samplings in the two executions, as specified by the coupling encoded in the original pRHL derivation. Second, we propose a general rule for **while** loops; our rule subsumes several existing rules and is sufficiently expressive to capture several examples of shift coupling. The resulting system, which we call *product pRHL* or \times pRHL for short, has several benefits.

First, we obtain a simple, algorithmic procedure to construct the probabilistic product given a pRHL derivation. The product directly simulates the coupled processes, so we can prove quantitative properties about probabilities or expected values of this coupling using existing probabilistic verification techniques. Moreover, intermediate assertions in a pRHL derivation to be soundly transported to the probabilistic product in \times pRHL derivations. For instance, we may prove synchronized control flow in pRHL, and directly assume this fact in the probabilistic product. Many facts are often practically easier to manipulate in pRHL, since pRHL works purely on non-probabilistic assertions.

Second, we obtain a powerful logic that can reason about many examples from the coupling literature, especially shift couplings, and from the translation validation literature, especially loop optimizations. On the foundational side, we prove that our logic is relatively complete for deterministic programs.

Summary of contributions We make the following contributions.

- We present a proof-relevant program logic \times pRHL that extracts a probabilistic product program from a valid derivation, and prove (using the Coq proof assistant) that the logic is sound.
- We propose new rules for loops and random sampling in \times pRHL. We also prove relative completeness for deterministic programs.
- We demonstrate several applications of \times pRHL to showing convergence of probabilistic processes; several of these examples use path coupling, a compositional tool for constructing couplings which bounds convergence in terms of expected properties of the coupling, and shift coupling, a generalization of coupling where the two processes are allowed to meet at different times. We also show how to validate some common loop transformations.

2. Preliminaries

We begin by giving a bird’s eye view on probabilistic couplings, which take output distributions from two probabilistic processes and place them in the same probabilistic space.

In the following, we will work with sub-distributions over discrete (finite or countable) sets.

Definition 1. A discrete sub-distribution over a set A is completely defined by its *mass function* $\mu : A \rightarrow [0, 1]$ such that $\sum_{a \in A} \mu(a)$ is

defined and bound by 1. The quantity $\sum_{a \in A} \mu(a) \in [0, 1]$ is called the *weight* of μ and denoted by $|\mu|$. The support $\text{supp}(\mu)$ of μ is defined as $\{x \in A \mid \mu(x) \neq 0\}$ and is discrete (i.e., countable) by construction. A *distribution* is a sub-distribution with weight 1. The probability of an event P w.r.t. μ , written $\Pr_\mu[P]$ (or $\Pr[P]$ when μ is clear from the context), is defined as $\sum_{\{x \in A \mid P(x)\}} \mu(x)$.

One can equip distributions with a monadic structure using the Dirac distributions $\mathbb{1}_x$ for the unit¹ and *distribution expectation* $\mathbb{E}_{x \sim \mu}[M(x)]$ for the bind, where

$$\mathbb{E}_{x \sim \mu}[M(x)] : x \mapsto \sum_a \mu(a) \cdot M(a)(x).$$

When working with sub-distributions over tuples, the probabilistic versions of the usual projections on tuples are called *marginals*. For distributions over pairs, we define the *first* and *second marginals* $\pi_1(\mu)$ and $\pi_2(\mu)$ of a distribution μ over $A \times B$ by $\pi_1(\mu)(a) = \sum_{b \in B} \mu(a, b)$ and $\pi_2(\mu)(b) = \sum_{a \in A} \mu(a, b)$. For a distribution μ over memories and a set of variables $X \subseteq \mathbf{Vars}$, we define the *X-marginal distribution* of μ as

$$\pi_X(\mu)(\bar{x}) = \sum_{\bar{v} \in \mathbf{Vars} \setminus X} \mu(\bar{x}, \bar{v}).$$

We can also construct a sub-distribution over tuples from two distributions: the *product* sub-distribution of sub-distributions μ_1 and μ_2 over A and B respectively is defined by

$$(\mu_1 \times \mu_2)(a, b) = \mu_1(a) \cdot \mu_2(b).$$

We are now ready to formally define coupling.

Definition 2. Two sub-distributions μ_1, μ_2 resp. over A and B are said to be *coupled* with witness $\mu \in \mathbf{Distr}(A \times B)$, written $\mu \blacktriangleleft \langle \mu_1 \& \mu_2 \rangle$, iff $\pi_1(\mu) = \mu_1$ and $\pi_2(\mu) = \mu_2$.

For $\Psi \subseteq A \times B$, we write $\mu \blacktriangleleft_\Psi \langle \mu_1 \& \mu_2 \rangle$ if $\mu \blacktriangleleft \langle \mu_1 \& \mu_2 \rangle$ and moreover $\text{supp}(\mu) \subseteq \Psi$. We will often abuse notation and call μ a *coupling* of μ_1 and μ_2 .

As an example, suppose that $A = B$ and $\mu_1 = \mu_2$ are the uniform distributions. Then, any bijection $f : A \rightarrow A$ gives a coupling of μ_1 and μ_2 ; we call the resulting coupling $\mathcal{D}_f(A)$ so that $\mathcal{D}_f(A) \blacktriangleleft \langle \mu_1 \& \mu_2 \rangle$. The coupling $\mathcal{D}_f(A)$ assigns positive probability only to pairs $(v, f v)$ with $v \in A$.

For another example, suppose again that $A = B$ and $\mu_1 = \mu_2$, but the distributions are not necessarily uniform. Then, the identity map $\text{id} : A \rightarrow A$ always gives a coupling of μ_1 and μ_2 , correlating samples from both distributions to be the same. We will write $\mathcal{D}^=(\mu_1)$ for the resulting coupling, so that $\mathcal{D}^=(\mu_1) \blacktriangleleft \langle \mu_1 \& \mu_2 \rangle$. If μ_1 is the uniform distribution we will sometime write $\mathcal{D}^=(A)$ instead of $\mathcal{D}^=(\mu_1)$. Note that $\mathcal{D}^=(\mu_1)$ assigns positive probability only to pairs (v, v) with $\mu_1(v), \mu_2(v) \neq 0$.

To reason about convergence of probabilistic processes, we will use the total variation distance on distributions (also known as statistical distance).

Definition 3. Let μ_1 and μ_2 be sub-distributions over a countable set A . The *total variation (TV) distance* between μ_1 and μ_2 is defined by:

$$\text{TV}(\mu_1, \mu_2) = \frac{1}{2} \sum_{a \in A} |\mu_1(a) - \mu_2(a)|.$$

To bound this distance, it is enough to find a coupling and bound the probability that the two coupled variables differ; this is the fundamental idea underlying the coupling method.

¹ $\mathbb{1}_x$ is the distribution where x has probability 1 and all other elements have probability 0.

Theorem 4 (Fundamental theorem of couplings (e.g., [42])). *Let μ_1 and μ_2 be distributions over a countable set. Then for any coupling μ of μ_1, μ_2 , we have*

$$\text{TV}(\mu_1, \mu_2) \leq \Pr_{(x, x') \sim \mu}[x \neq x'].$$

3. Product programs

3.1 Language

We will work with a core, probabilistic imperative language with a command for random sampling from primitive distributions. The set of commands is defined as follows:

$C ::=$	skip	noop
	abort	abort
	$\mathcal{X} \leftarrow \mathcal{E}$	deterministic assignment
	$\mathcal{X} \stackrel{\mathcal{D}}{\leftarrow} \mathcal{D}(\mathcal{E})$	probabilistic assignment
	$C; C$	sequencing
	if \mathcal{E} then C else C	conditional
	while \mathcal{E} do C	while loop

Here, \mathcal{X} is a set of *variables*, \mathcal{E} is a set of (deterministic) *expressions*, and \mathcal{D} is a set of *distribution expressions*. Variables and expressions are typed, ranging over booleans, integers, lists, etc. The expression grammar is entirely standard, and we omit it.

We will use several shorthands for commands: **if** e **then** c for **if** e **then** c **else skip**, **abort** for the looping command **while** \top **do skip**, and $c^{e,k}$ for the k -fold composition of c restricted to e . Formally,

$$c^{e,k} \triangleq \iota = 0; \text{while } (\iota < k) \wedge e \text{ do } c; \iota++$$

where ι is a fresh variable. We allow k to be an arbitrary expression; when k is a constant,

$$\vdash c^{e,k} \equiv \overbrace{\text{if } e \text{ then } c; \dots; \text{if } e \text{ then } c}^{k \text{ times}}$$

for the simple notion of program equivalence introduced in Figure 2.

The denotational semantics of programs is adapted from the seminal work of Kozen [36]. We first interpret every ground type T as a set $\llbracket T \rrbracket$; other constructors C are interpreted as functions $\llbracket C \rrbracket$ that respect their arities. The semantics of expressions and distribution expressions is then parameterized by a state m , and is defined in the usual way. Last, commands are interpreted as a function from memories to sub-distributions over memories, where memories are finite maps from variables to values. Formally, we let State_X denote the set of memories over the finite set of variables X . Moreover, we use \uplus to denote the standard disjoint union on finite maps, so $\uplus : \text{State}_{X_1} \times \text{State}_{X_2} \rightarrow \text{State}_{X_1 \cup X_2}$ for disjoint finite sets X_1 and X_2 . The interpretation of c , written $\llbracket c \rrbracket$, is a function from State_X to $\text{Distr}(\text{State}_X)$, where $\text{Vars}(c) = X$. The definition of $\llbracket \cdot \rrbracket$ is given in Figure 1.

Last, for any predicate over memories Φ , we say that a command is Φ -*lossless* iff for any memory m s.t. $\Phi(m)$, the weight of $\llbracket c \rrbracket_m$ is equal to 1. A command is *lossless* if it is \top -lossless, i.e. if $|\llbracket c \rrbracket_m| = 1$ for every memory m .

3.2 Proof system

Our proof system manipulates judgments of the form:

$$\left\{ \Phi \right\} \begin{array}{c} c_1 \\ c_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c$$

where c_1, c_2 and c are statements over disjoint variables, and Φ and Ψ are assertions over the variables of c_1 and c_2 . We will not fix a specific syntax for assertions, but one natural choice is first-order formulas over the program variables of both programs.

Informally, the judgment is valid if c is a probabilistic product program for c_1 and c_2 under the pre-condition Φ , i.e. for every initial memory $m = m_1 \uplus m_2$ satisfying the pre-condition Φ , the sub-distribution $\llbracket c \rrbracket_m$ is a coupling for $\llbracket c_1 \rrbracket_{m_1}$ and $\llbracket c_2 \rrbracket_{m_2}$, and moreover $\text{supp}(\llbracket c \rrbracket_m)$ only contains states that satisfy Ψ .

Definition 5 (Valid judgment).

- Two commands c_1 and c_2 are *separable* iff their variables are disjoint, i.e. $\text{Vars}(c_1) \cap \text{Vars}(c_2) = \emptyset$.
- A judgment $\left\{ \Phi \right\} \begin{array}{c} c_1 \\ c_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c$ is *valid* iff the commands c_1 and c_2 are separable and for every memory $m = m_1 \uplus m_2$ such that $m \models \Phi$, we have

$$\llbracket c \rrbracket_m \triangleleft_{\Psi} \langle (\llbracket c_1 \rrbracket_{m_1}) \& (\llbracket c_2 \rrbracket_{m_2}) \rangle.$$

For comparison, judgments in pRHL are of the form

$$\left\{ \Phi \right\} \begin{array}{c} c_1 \\ c_2 \end{array} \left\{ \Psi \right\}$$

and assert that for every initial memory $m = m_1 \uplus m_2$ such that $m \models \Phi$, there exists a coupling μ such that

$$\mu \triangleleft_{\Psi} \langle (\llbracket c_1 \rrbracket_{m_1}) \& (\llbracket c_2 \rrbracket_{m_2}) \rangle.$$

In contrast, our notion of judgment is *proof-relevant*, since the derivation of the judgment guides the product construction. We briefly comment on several rules in our system, presented in Fig. 3. The [CONSEQ] rule is similar to the rule of consequence, and reflects that validity is preserved by weakening the post-condition and strengthening the pre-condition; none of the programs is modified in this case.

The [FALSE] rule corresponds to the fact that under a false precondition nothing needs to be proved.

The [CASE] rule allows proving a judgment by case analysis; specifically, the validity of a judgment with pre-condition Φ can be established from the validity of two judgments, one where the pre-condition is strengthened with e and the other where the pre-condition is strengthened with $\neg e$. The product program is of the form **if** e **then** c **else** c' , where c and c' correspond to the product programs built from the proofs of the first and second premises respectively.

The [STRUCT] rule allows replacing programs by provably equivalent programs. The rules for proving program equivalence are given in Figure 2, and manipulate judgments of the form $\Phi \vdash c \equiv c'$, where Φ is a relational assertion. We only introduce equations that are needed for recovering derived rules, striving to keep the notion of structural equivalence as simple as possible.

The [ASSG] rule corresponds to the pRHL rule for assignments; in this case, the product program is simply the sequential composition of the two assignments.

The [RAND] rule informally takes a coupling between the two distributions used for sampling in the left and right program, and produces a product program that draws the pair of correlated samples from the coupling. Since our language supports sampling from distribution expressions and not only distributions, the rule asks for the existence of a coupling for each interpretation of the distribution expressions under a valuation satisfying the pre-condition Φ of the judgment. Furthermore, note that the rule requires that every element in the support of the coupling validates the post-condition. This is similar to pRHL, and is a natural requirement given the notion of valid judgment. However, the rule is strictly more general than the corresponding pRHL rule, which is restricted to the case where Ψ is the graph of some bijection.

The [WHILE] rule for **while** loops constructs a product program that interleaves synchronous and asynchronous executions of the loop bodies. The first premise establishes that k_1 and k_2 are strictly

$$\begin{aligned}
\llbracket \text{skip} \rrbracket_m &= \mathbb{1}_m & \llbracket \text{abort} \rrbracket_m &= 0 \\
\llbracket x \leftarrow e \rrbracket_m &= \mathbb{1}_{m[x:=\llbracket e \rrbracket_m]} & \llbracket x \stackrel{\$}{\leftarrow} g \rrbracket_m &= \mathbb{E}_{v \sim \llbracket g \rrbracket_m} [\mathbb{1}_{m[x:=v]}] \\
\llbracket c_1; c_2 \rrbracket_m &= \mathbb{E}_{\xi \sim \llbracket c_1 \rrbracket_m} [\llbracket c_2 \rrbracket_m] & \llbracket \text{if } e \text{ then } s_1 \text{ else } s_2 \rrbracket_m &= \text{if } \llbracket e \rrbracket_m \text{ then } \llbracket c_1 \rrbracket_m \text{ else } \llbracket c_2 \rrbracket_m \\
\llbracket \text{while } b \text{ do } c \rrbracket_m &= \lim_{n \rightarrow \infty} \llbracket c_{b,n} \rrbracket_m & \text{where } c_{b,n} &\triangleq \overbrace{\text{if } b \text{ then } c; \dots; \text{if } b \text{ then } c}^{k \text{ times}}; \text{if } b \text{ then abort}
\end{aligned}$$

Figure 1. Interpretation of commands

positive when the loop invariant Ψ holds. Then, we specify an expression e —which may mention variables from both sides—that holds true exactly when at least one of the guards is true.

Next, the notation $\oplus\{p_0, p_1, p_2\}$ indicates that exactly one of the tests p_0 , p_1 , and p_2 must hold. These predicates guide the construction of the product program. If p_0 holds, then both guards should be equal and we can execute the two sides k_1 and k_2 iterations respectively, preserving the loop invariant Ψ . If p_1 holds and the right loop has not terminated yet, then the left loop also has not terminated yet and we may execute the left loop one iteration. If p_2 holds and the left loop has not terminated yet, then the right loop also has not terminated yet and we may execute the right loop one iteration.

The last pair of premises deal with termination. Note that some condition on termination is needed for soundness of the logic: if the left loop terminates with probability 1 while the right loop terminates with probability 0 (i.e., never), it is impossible to construct a valid product program since there is no distribution on pairs that has first marginal with weight 1 and second marginal with weight 0. So, we require that the first and second loops are lossless assuming p_1 and p_2 respectively. This ensures that with probability 1, there are only finitely many steps where we execute the left or right loop separately. Note that the product program may still fail to terminate with positive probability since there may be infinite sequences of iterations where p_0 holds so *both* loops advance, but both programs will yield sub-distributions with the *same* weight so we can still find a coupling.

With these premises, the construction of the product program is straightforward. At the top level, the product continues as long as e is true, i.e., while at least one of the two loops can make progress. Each iteration, it performs a case analysis on the three predicates. If p_0 holds, then we execute the product from running the two loops k_1 and k_2 iterations respectively. If p_1 holds, then we execute the product from running the left loop one iteration; if p_2 then we execute the product from running the right loop one iteration.

Derived rules

Presentations of relational Hoare logic often include *one-sided rules*, which are based on the analysis of a single program (rather than two). By reasoning about program equivalence, we can derive all one-sided rules and the two-sided rules of conditionals and loops from pRHL, presented in Fig. 4, within our system.²

Proposition 6. *All the rules in Figure 4 (and their symmetric version) are derived rules.*

We briefly comment on some of the derived rules. The [ASSG-L] rule is the one-sided rule for assignment. It can be derived from its two-sided counterpart, the [ASSG] rule, using the [STRUCT] rule.

²In fact, the one-sided and two-sided rules are inter-derivable for all constructions except random assignments and loops.

$$\begin{aligned}
&\frac{}{\Phi \vdash c \equiv c} & \frac{}{\Phi \vdash c_1 \equiv c_2}{\Phi \vdash c_2 \equiv c_1} & \frac{}{\Phi \vdash x \stackrel{\$}{\leftarrow} \mathbb{1}_x \equiv \text{skip}} \\
&\frac{}{\Phi \vdash (x_1, x_2) \stackrel{\$}{\leftarrow} \mu_1 \times \mu_2 \equiv x_1 \stackrel{\$}{\leftarrow} \mu_1; x_2 \stackrel{\$}{\leftarrow} \mu_2} \\
&\frac{\Phi \implies x = e}{\Phi \vdash x \leftarrow e \equiv \text{skip}} & \frac{}{\Phi \vdash c; \text{skip} \equiv c} & \frac{}{\Phi \vdash \text{skip}; c \equiv c} \\
&\frac{\Phi \vdash c_1 \equiv c'_1}{\Phi \vdash c_1; c_2 \equiv c'_1; c_2} & \frac{\top \vdash c_2 \equiv c'_2}{\Phi \vdash c_1; c_2 \equiv c_1; c'_2} \\
&\frac{\Phi \implies e}{\Phi \vdash \text{if } e \text{ then } c \text{ else } c' \equiv c} & \frac{\Phi \implies \neg e}{\Phi \vdash \text{if } e \text{ then } c \text{ else } c' \equiv c'} \\
&\frac{\Phi \wedge e \vdash c_1 \equiv c_2 \quad \Phi \wedge \neg e \vdash c'_1 \equiv c'_2}{\Phi \vdash \text{if } e \text{ then } c_1 \text{ else } c'_1 \equiv \text{if } e \text{ then } c_2 \text{ else } c'_2} \\
&\frac{e \vdash c \equiv c'}{\Phi \vdash \text{while } e \text{ do } c \equiv \text{while } e \text{ do } c'} \\
&\frac{}{\Phi \vdash \text{while } e \text{ do } c \equiv \text{if } e \text{ then } (c; \text{while } e \text{ do } c)}
\end{aligned}$$

Figure 2. Equivalence rules

The [RAND-L] rule is the one-sided rule for random sampling. It can be derived from its two-sided counterpart, the [RAND] rule, using the [STRUCT] rule and the fact that the product distribution $\mu_1 \times \mu_2$ is a valid coupling of proper distributions μ_1 and μ_2 .

We have one-sided and two-sided rules for conditionals. The [COND-L] rule is the one-sided version; it can be derived from the [CASE] and [STRUCT] rules. The [COND-S] rule is the two-sided version; the rule assumes that the two guards of the conditional statements are synchronized, so that one must only need to reason about the cases where both statements enter the true branch, and the case where both statements enter the false branch. It can be derived from the [CASE], [STRUCT], and [FALSE] rules.

We also have one-sided and two-sided rules for loops. The [WHILE-L] rule corresponds to the pRHL one-sided rule for **while** loops; it can be derived from the [WHILE] rule by setting $p_1 = \top$, $p_0 = p_2 = \perp$, and $k_1 = k_2 = 1$, $e = e_1$, $e_2 = \perp$, $c_2 = \text{skip}$ and using the [STRUCT] and [FALSE] rules. The rule [WHILE-S] corresponds to the pRHL two-sided rule for **while** loops. This rule assumes that the two loops are synchronized, i.e., the guards of the two loops are equal assuming the loop invariant. This rule can be derived from the general [WHILE] rule by setting $p_0 = \top$,

$$\begin{array}{c}
\text{CONSEQ} \frac{\left\{ \Phi \right\} \begin{array}{l} c_1 \\ c_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c \quad \Phi' \Longrightarrow \Phi \quad \Psi \Longrightarrow \Psi'}{\left\{ \Phi' \right\} \begin{array}{l} c_1 \\ c_2 \end{array} \left\{ \Psi' \right\} \rightsquigarrow c} \quad \text{FALSE} \frac{}{\left\{ \perp \right\} \begin{array}{l} c_1 \\ c_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow \text{skip}} \\
\text{STRUCT} \frac{\left\{ \Phi \right\} \begin{array}{l} c_1 \\ c_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c \quad \Phi \vdash c_1 \equiv c'_1 \quad \Phi \vdash c_2 \equiv c'_2 \quad \Phi \vdash c \equiv c'}{\left\{ \Phi \right\} \begin{array}{l} c'_1 \\ c'_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c'} \\
\text{CASE} \frac{\left\{ \Phi \wedge e \right\} \begin{array}{l} c_1 \\ c_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c \quad \left\{ \Phi \wedge \neg e \right\} \begin{array}{l} c_1 \\ c_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c'}{\left\{ \Phi \right\} \begin{array}{l} c_1 \\ c_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow \text{if } e \text{ then } c \text{ else } c'} \\
\text{ASSG} \frac{}{\left\{ \Psi[e_1/x_1, e_2/x_2] \right\} \begin{array}{l} x_1 \leftarrow e_1 \\ x_2 \leftarrow e_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow x_1 \leftarrow e_1; x_2 \leftarrow e_2} \\
\text{RAND} \frac{\forall m. m \models \Phi \Longrightarrow \llbracket \mu \rrbracket_m \triangleleft \{(v_1, v_2) \mid m[x_1 := v_1, x_2 := v_2] \models \Psi\} \langle \llbracket \mu_1 \rrbracket_m \ \& \ \llbracket \mu_2 \rrbracket_m \rangle}{\left\{ \Phi \right\} \begin{array}{l} x_1 \xleftarrow{\S} \mu_1 \\ x_2 \xleftarrow{\S} \mu_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow (x_1, x_2) \xleftarrow{\S} \mu} \\
\text{SEQ} \frac{\left\{ \Phi \right\} \begin{array}{l} c_1 \\ c_2 \end{array} \left\{ \Xi \right\} \rightsquigarrow c \quad \left\{ \Xi \right\} \begin{array}{l} c'_1 \\ c'_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c'}{\left\{ \Phi \right\} \begin{array}{l} c_1; c'_1 \\ c_2; c'_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c; c'} \\
\text{WHILE} \frac{\Psi \Longrightarrow k_1 > 0 \wedge k_2 > 0 \quad \Psi \Longrightarrow (e_1 \vee e_2) = e \quad \Psi \wedge e \Longrightarrow \oplus\{p_0, p_1, p_2\} \\ \Psi \wedge p_0 \wedge e \Longrightarrow e_1 = e_2 \quad \Psi \wedge p_1 \wedge e \Longrightarrow e_1 \quad \Psi \wedge p_2 \wedge e \Longrightarrow e_2 \\ \text{while } e_1 \wedge p_1 \text{ do } c_1 \quad \Psi\text{-lossless} \quad \text{while } e_2 \wedge p_2 \text{ do } c_2 \quad \Psi\text{-lossless}}{\left\{ \Psi \wedge p_0 \right\} \begin{array}{l} c_1^{e_1, k_1} \\ c_2^{e_2, k_2} \end{array} \left\{ \Psi \right\} \rightsquigarrow c_0 \quad \left\{ \Psi \wedge e_1 \wedge p_1 \right\} \begin{array}{l} c_1 \\ \text{skip} \end{array} \left\{ \Psi \right\} \rightsquigarrow c_1 \quad \left\{ \Psi \wedge e_2 \wedge p_2 \right\} \begin{array}{l} \text{skip} \\ c_2 \end{array} \left\{ \Psi \right\} \rightsquigarrow c_2} \\ \left\{ \Psi \right\} \text{while } e_1 \text{ do } c_1 \\ \text{while } e_2 \text{ do } c_2 \left\{ \Psi \wedge \neg e_1 \wedge \neg e_2 \right\} \rightsquigarrow \text{while } e \text{ do if } p_0 \text{ then } c_0 \text{ else if } p_1 \text{ then } c_1 \text{ else } c_2
\end{array}$$

Figure 3. Proof rules

$p_1 = p_2 = \perp$, and $k_1 = k_2 = 1$, and $e = e_1$, and using the [FALSE] rule.

3.3 Soundness and (relative) completeness

We have formally verified the soundness theorem below in the Coq proof assistant, with Ssreflect.

Theorem 7 (Soundness). *Every derivable judgment is valid.*

Proof sketch. By induction on the proof derivation. The most interesting case is the loop rule, which we briefly sketch here. We will use the same notations from Fig. 3. Proving that the product program satisfies the post-condition is immediate: by construction, the condition e does not hold after the execution of the loop. Moreover, the premises ensure that the loop body of the product preserves the invariant Ψ . So, we just need to prove that the semantics of the projections of the product program correspond to respective semantics of the original programs. We detail the first projection, i.e. that for any initial memory m s.t. $m \models \Psi$, we have

$$\llbracket w_1 \rrbracket_{\pi_1 m} = \pi_1(\llbracket w \rrbracket_m) \quad (1)$$

where $w \triangleq \text{while } e \text{ do } c$ and $w_i \triangleq \text{while } e_i \text{ do } c_i$; the other projection is similar. By definition, we have for any e and c

$$\llbracket w \rrbracket_m = \lim_{n \rightarrow \infty} \llbracket c_{e,n} \rrbracket_m$$

Since $\llbracket c_{e,n} \rrbracket_m$ is monotonic in n , proving Eq. (1) amounts to showing

$$\llbracket c_{1e_1,n} \rrbracket_{\pi_1 m} \leq \pi_1(\llbracket w \rrbracket_m) \quad (2)$$

$$\pi_1(\llbracket c_{e,n} \rrbracket_m) \leq \llbracket w_1 \rrbracket_{\pi_1 m} \quad (3)$$

for all n . We proceed by induction on n . Let $m_i \triangleq \pi_i(m)$ for $i \in \{1, 2\}$. The base case is immediate and we consider the $n + 1$ case. We will prove (2) by case analysis on $\llbracket e_1 \rrbracket_{m_1}$ and $\llbracket e \rrbracket_m$. If $\llbracket e_1 \rrbracket_{m_1}$ does not hold, then $\llbracket c_{1e_1,n} \rrbracket_{m_1} = \mathbb{1}_{m_1}$. Now if

$\llbracket e_2 \rrbracket_{m_2}$ is false, then $\llbracket e \rrbracket_m$ is also false and (2) holds. Otherwise if $\llbracket e_2 \rrbracket_{m_2}$ is true, $\llbracket p_0 \rrbracket_m$ and $\llbracket p_1 \rrbracket_m$ cannot hold, since that would imply that $\llbracket e_1 \rrbracket_{m_1}$ holds. So in this case, we can conclude (2) since $\llbracket w \rrbracket_m = \mathbb{1}_{m_1} \times \llbracket w_2 \rrbracket_{m_2}$ and w_2 is Ψ -lossless.

Otherwise, $\llbracket e_1 \rrbracket_{m_1}$ holds. We proceed by case analysis on $\llbracket p_0 \rrbracket_m$, $\llbracket p_1 \rrbracket_m$ and $\llbracket p_1 \rrbracket_m$, as they are mutually exclusive by the premise.

$$\begin{array}{c}
\text{ASSG-L} \frac{}{\left\{ \Psi[e_1/x_1] \right\} \quad x_1 \leftarrow e_1 \quad \text{skip} \quad \left\{ \Psi \right\} \rightsquigarrow x_1 \leftarrow e_1} \\
\\
\text{RAND-L} \frac{\mu_1 \text{ lossless}}{\left\{ \forall v_1 \in \text{supp}(\mu_1), \Psi[v_1/x_1] \right\} \quad x_1 \stackrel{\$}{\leftarrow} \mu_1 \quad \text{skip} \quad \left\{ \Psi \right\} \rightsquigarrow x_1 \stackrel{\$}{\leftarrow} \mu_1} \\
\\
\text{COND-L} \frac{\left\{ \Phi \wedge e_1 \right\} \quad c_1 \quad \left\{ \Psi \right\} \rightsquigarrow c \quad \left\{ \Phi \wedge \neg e_1 \right\} \quad c'_1 \quad \left\{ \Psi \right\} \rightsquigarrow c'}{\left\{ \Phi \right\} \quad \text{if } e_1 \text{ then } c_1 \text{ else } c'_1 \quad c_2 \quad \left\{ \Psi \right\} \rightsquigarrow \text{if } e_1 \text{ then } c \text{ else } c'} \\
\\
\text{COND-S} \frac{\Phi \implies e_1 = e_2 \quad \left\{ \Phi \wedge e_1 \right\} \quad c_1 \quad \left\{ \Psi \right\} \rightsquigarrow c \quad \left\{ \Phi \wedge \neg e_1 \right\} \quad c'_1 \quad \left\{ \Psi \right\} \rightsquigarrow c'}{\left\{ \Phi \right\} \quad \text{if } e_1 \text{ then } c_1 \text{ else } c'_1 \quad c_2 \quad \left\{ \Psi \right\} \rightsquigarrow \text{if } e_1 \text{ then } c \text{ else } c'} \\
\\
\text{WHILE-L} \frac{\left\{ \Psi \wedge e_1 \right\} \quad c_1 \quad \left\{ \Psi \right\} \rightsquigarrow c_1 \quad \text{while } e_1 \text{ do } c_1 \quad \Psi\text{-lossless}}{\left\{ \Psi \right\} \quad \text{while } e_1 \text{ do } c_1 \quad \text{skip} \quad \left\{ \Psi \wedge \neg e_1 \right\} \rightsquigarrow \text{while } e_1 \text{ do } c_1} \\
\\
\text{WHILE-S} \frac{\Psi \implies e_1 = e_2 \quad \left\{ \Psi \wedge e_1 \right\} \quad c_1 \quad \left\{ \Psi \right\} \rightsquigarrow c}{\left\{ \Psi \right\} \quad \text{while } e_1 \text{ do } c_1 \quad \left\{ \Psi \wedge \neg e_1 \right\} \rightsquigarrow \text{while } e_1 \text{ do } c}
\end{array}$$

Figure 4. Derived rules

- If $\llbracket p_0 \rrbracket_m$ holds, $k_1 > 0$ implies

$$\llbracket c_{1e_1, n+1} \rrbracket_{m_1} \leq \llbracket c_1; c_{1e_1, n} \rrbracket_{m_1} \leq \llbracket c_1^{e_1, k_1}; c_{1e_1, n} \rrbracket_{m_1}$$

which allows us to conclude using the induction hypothesis, since $\llbracket w \rrbracket_m = \llbracket c_0; w \rrbracket_m$ in this case and $\llbracket c_1^{e_1, k_1} \rrbracket_{m_1} = \pi_1(\llbracket c_0 \rrbracket_m)$

- If $\llbracket p_1 \rrbracket_m$ holds, we have

$$\begin{aligned} \llbracket c_{1e_1, n+1} \rrbracket_{m_1} &= \llbracket c_1; c_{1e_1, n} \rrbracket_{m_1} \\ \llbracket w \rrbracket_m &= \llbracket c_1; w \rrbracket_m \end{aligned}$$

which allows to conclude using the induction hypothesis.

- If $\llbracket p_2 \rrbracket_m$ holds, we have

$$\llbracket c_{1e_1, n+1} \rrbracket_{\pi_1 m} = \mathbb{E}_{m'_1 \sim \pi_1} \llbracket \text{while } e_2 \wedge p_2 \text{ do } c_2; w \rrbracket_m \llbracket c_{1e_1, n+1} \rrbracket_{m'_1}.$$

Note that after the executing **while** $e_2 \wedge p_2$ **do** c_2 —which is lossless— Ψ and e_1 still hold while e_2 does not. By the premises of the rule, p_2 must be false so p_0 or p_1 holds on every memory m'_1 in the support. By the first two cases with initial memory m'_1 , we have:

$$\llbracket c_{1e_1, n+1} \rrbracket_{m'_1} \leq \pi_1 \llbracket w \rrbracket_{m'_1}.$$

We can then conclude:

$$\begin{aligned} &\llbracket c_{1e_1, n+1} \rrbracket_{\pi_1 m} \\ &\leq \mathbb{E}_{m'_1 \sim \pi_1} \llbracket \text{while } e_2 \wedge p_2 \text{ do } c_2; w \rrbracket_m [\pi_1 \llbracket w \rrbracket_{m'_1}] \quad (\text{linearity}) \\ &= \pi_1 (\mathbb{E}_{m'_1 \sim \llbracket \text{while } e_2 \wedge p_2 \text{ do } c_2; w \rrbracket_m} \llbracket w \rrbracket_{m'_1}) \quad (\text{separable}) \\ &= \pi_1 \llbracket w \rrbracket_m. \quad (\text{definition}) \end{aligned}$$

Where the second to last step is because **while** $e_2 \wedge p_2$ **do** c_2 does not modify $\pi_1 m$ since c_1, c_2 are separable.

The proof of Eq. (3) follows the same pattern. \square

Although it is not a primary objective of our work, we briefly comment on completeness of the logic. First, the coupling method is not complete for proving rapid mixing of Markov chains.³ Second, it is not clear that our proof system is complete with respect to hoisting random assignments out of loops.

However, we note that the *deterministic* fragment of our logic achieves completeness for programs that satisfy a sufficiently strong termination property; the key is that the new rule for while subsumes self-composition for **while** loops, provided they terminate on all initial memories satisfying the invariant. More precisely, we can prove the following completeness theorem.

Theorem 8. *Let c_1 and c_2 be separable deterministic programs. If $\{\Phi\} \quad c_1; c_2 \quad \{\Psi\}$ is derivable using the rules from Figure 9, then $\left\{ \Phi \right\} \quad c_1 \quad \left\{ \Psi \right\} \rightsquigarrow _$ is derivable. Therefore, $\times\text{pRHL}$ is relatively complete for deterministic programs.*

Proof sketch. It suffices to prove that if $\{\Phi\} \quad c_1 \quad \{\Psi\}$ is derivable using Hoare logic, then

$$\left\{ \Phi \right\} \quad c_1 \quad \text{skip} \quad \left\{ \Psi \right\} \rightsquigarrow _$$

is derivable. The proof proceeds by induction on the derivation. \square

³Kumar and Ramesh [39] show that the class of *causal couplings*—which contains all couplings in our logic—are unable to prove rapid mixing for some rapidly-mixing Markov chains [29, 34]

3.4 Convergence from couplings

The fundamental theorem of couplings (Theorem 4) gives a powerful method to prove convergence of random processes. First, we recast it in terms of \times pRHL.

Proposition 9. *Let c_1 and c_2 be separable programs and assume that the following judgment is valid:*

$$\left\{ \Phi \right\} \begin{array}{c} c_1 \\ c_2 \end{array} \left\{ \Psi \implies x_1 = x_2 \right\} \rightsquigarrow c$$

Then for every memory m such that $m \models \Phi$, we have

$$\text{TV}(\mu_1, \mu_2) \leq \Pr_{x \leftarrow [c]_m} [\neg \Psi]$$

where μ_1, μ_2 are the distributions obtained by sampling x_1 from $[[c_1]]_m$ and x_2 from $[[c_2]]_m$ respectively.

This result is a direct consequence of the soundness of the logic, and allows proving convergence in two parts. First, we use \times pRHL to establish a valid post-condition of the form $\Psi \implies x_1 = x_2$. Second, we prove that for every memory m satisfying some pre-condition, the product program c built from the derivation satisfies:

$$\Pr_{x \leftarrow [c]_m} [\neg \Psi] \leq \beta$$

There are multiple approaches for proving properties of this form—reasoning directly about the semantics of programs; existing formalisms for bounding probabilities and reasoning about expectations (e.g., [37, 45]); program logics for probabilistic programs (e.g., [16, 28, 48]). We will check the property on pen and paper; mechanizing the proofs is left for future work.

4. Application: convergence of Markov chains

We now turn to our first group of examples: proving convergence of probabilistic processes. Suppose we have a probabilistic process on a set Ω of possible states. At each time step, the process selects the next state probabilistically. Consider two runs of the same probabilistic process started from two different states in Ω . We would like to know how many steps we need to run before the two distributions on states converges to a common distribution. We consider several classic examples.

Notation. Throughout this section, we consider two copies of the same program. To ensure that the two copies are separable, we tag all the variables of the first copy with 1, and all the variables of the second copy with 2.

4.1 Simple, symmetric random walk

```

x ← s; v ← 0; l ← [];
while v < T do
  r  $\stackrel{\$}{\leftarrow}$  {−1, 1};
  x ← x + r;
  l ← x :: l;
  v ← v + 1;
return x

```

Figure 5. Random walk

track of the list of sampled values, and will be used to state assertions when we construct the coupling.

Now, consider the random walk started from starting positions s_1, s_2 . We want to show that the TV-distance between the two distributions on positions decreases as we run for more steps T ; roughly, the random walk forgets its initial position. It is not hard to see that if $s_1 - s_2$ is an odd integer, then we will not have

Our first example is a simple random walk on the integers. Let the state space Ω be \mathbb{Z} . At each step, the process chooses uniformly to move left (decreasing the position by 1) or right (increasing the position by 1). The program `rwalk` in Figure 5 implements the process, executed for T steps. The variable l is a *ghost variable*. While it does not influence the process, it keeps

track of the list of sampled values, and will be used to state assertions

when we construct the coupling.

Now, consider the random walk started from starting positions s_1, s_2 . We want to show that the TV-distance between the two distributions on positions decreases as we run for more steps T ; roughly, the random walk forgets its initial position. It is not hard to see that if $s_1 - s_2$ is an odd integer, then we will not have

convergence: at any timestep t , the support of one distribution will be on even integers while the support of the other distribution will be on odd integers, so the TV-distance remains 1.

When $s_1 - s_2 = 2k$ is even, we can construct a coupling to show convergence. Barthe et al. [12] used pRHL to couple these random walks by *mirroring*; informally, the coupled walks make mirror moves until they meet, when they make identical moves to stay equal. Specifically, they show that

$$\left\{ \Phi \right\} \begin{array}{c} \text{rwalk}_1 \\ \text{rwalk}_2 \end{array} \left\{ k \in \text{psum}(\text{rev}(l_1)) \implies x_1 = x_2 \right\}$$

where $\Phi \triangleq s_1 + s_2 = 2k$, $\text{rev}(l)$ reverses the list, and $\text{psum}(l)$ is the list of partial sums of l (sums over its initial segments). We can lift the judgment to \times pRHL

$$\left\{ \Phi \right\} \begin{array}{c} \text{rwalk}_1 \\ \text{rwalk}_2 \end{array} \left\{ k \in \text{psum}(\text{rev}(l_1)) \implies x_1 = x_2 \right\} \rightsquigarrow \text{rwalk}_0$$

where `rwalk0` is the following product program:

```

x1 ← s1; x2 ← s2; v1 ← 0; v2 ← 0; l1 ← []; l2 ← [];
while v1 < T do
  if x1 = x2 then
    (r1, r2)  $\stackrel{\$}{\leftarrow}$  D=({−1, 1});
  else
    (r1, r2)  $\stackrel{\$}{\leftarrow}$  Dopp({−1, 1});
    x1 ← x1 + r1; x2 ← x2 + r2;
    l1 ← x1 :: l1; l2 ← x2 :: l2;
    v1 ← v1 + 1; v2 ← v2 + 1;
return (x1, x2)

```

where $\text{opp } x \triangleq -x$.

We briefly sketch the derivation of the product. We start by an application of the [WHILE-S] rule with the invariant

$$\begin{aligned} \Psi &\triangleq k \in \text{psum}(\text{rev}(l_1)) \implies x_1 = x_2 \\ &\wedge k \notin \text{psum}(\text{rev}(l_1)) \implies x_1 - x_2 = 2k - \text{sum}(l_1), \end{aligned}$$

where $\text{sum}(l)$ is the sum of the list. Then, we apply a [SEQ] rule (consuming the first random sampling on each side) with intermediate assertion

$$\Xi \triangleq \Psi \wedge (x_1 = x_2 \implies r_1 = r_2) \wedge (x_1 \neq x_2 \implies r_1 = -r_2)$$

The sub-proof obligation on tails is straightforward, the interesting one is for the random sampling. We start by using the [CASE] rule with $e \triangleq x_1 = x_2$, which introduce the conditional in the product program. If the equality holds, then the two random values are synchronized using the $\mathcal{D}^=(\{-1, 1\})$; if not, they are mirrored using $\mathcal{D}^{\text{opp}}(\{-1, 1\})$. In our mirror coupling, $k \in \text{psum}(l_1)$ implies that both the have already met, and continue to have the same position.

However, the derivation by itself does not tell us *how far* the two distributions are, as a function of T . To get this information, we will use the probabilistic product construction and the following classical result from the theory of random walks.

Theorem 10 (e.g., [3]). *Let X_0, X_1, \dots be a symmetric random walk on the integers with initial position $X_0 = 0$. Then, for any position $k \in \mathbb{Z}$, the probability that the walk does not reach k within t steps is at most*

$$\Pr[X_0, \dots, X_t \neq k] \leq \frac{ke\sqrt{2}}{\pi\sqrt{t}}.$$

Now, we can analyze how quickly the two walks mix.

Theorem 11. *If we perform a simple random walk for T steps from two positions that are $2k$ apart, then the resulting distributions μ_1 and μ_2 on final positions satisfy $\text{TV}(\mu_1, \mu_2) \leq \frac{ke\sqrt{2}}{\pi\sqrt{T}}$. Formally, for every two memories m_1 and m_2 such that $m_1(x) - m_2(x) = 2k$,*

we have

$$\text{TV}(\llbracket \text{rwalk} \rrbracket_{m_1}, \llbracket \text{rwalk} \rrbracket_{m_2}) \leq \frac{ke\sqrt{2}}{\pi\sqrt{T}},$$

Proof. Conceptually, we can think of the difference $x_1 - x_2$ as a random walk which increases by 2 with half probability and decreases by 2 with half probability. By applying Theorem 10 to this random walk, we find that in the product program c_0 ,

$$\Pr[k \notin \text{psum}(l_1)] \leq \frac{ke\sqrt{2}}{\pi\sqrt{T}}.$$

Then, we can conclude by Proposition 9. \square

4.2 The Dynkin process

Our second example models a process called the *Dynkin process*. There is a sequence of N concealed cards, each with a number drawn uniformly at random from $\{1, \dots, 10\}$. A player starts at some position in $\{1, \dots, 10\}$. Repeatedly, the player looks at the number at their current position, and moves forward that many steps. For instance, if the player reveals 2 at their current location, then she moves forward two spaces. The player stops when she passes the last card of the sequence. We want to show fast convergence of this process if we start from any two initial positions. In code, the Dynkin process is captured by the program `dynkin` defined in Figure 6. Here, l stores the history of positions of the player; this ghost variable will be useful both for writing assertions about the coupling, and for assertions in the product program. Just like for random walks, we can consider the mixing rate of this process, starting from two positions s_1, s_2 .

```

x ← s; l ← [x];
while x < N do
  r  $\stackrel{\$}{\leftarrow}$  [1, 10];
  x ← x + r;
  l ← x :: l;
return x

```

We will couple two runs of the Dynkin process from two starting positions asynchronously: we will move whichever process is behind, holding the other process temporarily fixed. If both processes are at the same position, then they move together.

Figure 6. Dynkin process

Formally, we prove the following \times pRHL judgment:

$$\left\{ \begin{array}{l} \top \\ \text{dynkin}_1 \\ \text{dynkin}_2 \end{array} \right\} \left\{ (l_1 \cap l_2 \neq \emptyset) \implies x_1 = x_2 \right\} \rightsquigarrow \text{dynkin}_0$$

where `dynkin0` is the following program:

```

x1 ← s1; x2 ← s2;
l1 ← [x1]; l2 ← [x2];
while x1 < N ∨ x2 < N do
  if x1 = x2 then
    (r1, r2)  $\stackrel{\$}{\leftarrow}$   $\mathcal{D}^=$ ([1, 10]);
    x1 ← x1 + r1; x2 ← x2 + r2;
    l1 ← x1 :: l1; l2 ← x2 :: l2;
  else
    if x1 < x2 then
      r1  $\stackrel{\$}{\leftarrow}$  [1, 10];
      x1 ← x1 + r1;
      l1 ← x1 :: l1;
    else
      r2  $\stackrel{\$}{\leftarrow}$  [1, 10];
      x2 ← x2 + r2;
      l2 ← x2 :: l2;
return (x1, x2)

```

To couple the two programs, we use the rule [WHILE] with $k_1 = k_2 = 1$. To control which process will advance, we define:

$$p_0 \triangleq (x_1 = x_2) \quad p_1 \triangleq (x_1 < x_2) \quad p_2 \triangleq (x_1 > x_2).$$

The lossless conditions are satisfied since the distance between x_1 and x_2 strictly decreases at each iteration of the loops. For the samplings, when p_1 or p_2 hold we use the one-sided rules for random sampling ([RAND-L] and corresponding [RAND-R]); when p_0 holds, we use the identity coupling.

Theorem 12. *Let m_1 and m_2 two memories such that $m_1(s), m_2(s) \in [1, 10]$, and suppose $N > 10$. Then:*

$$\text{TV}(\llbracket \text{dynkin} \rrbracket_{m_1}, \llbracket \text{dynkin} \rrbracket_{m_2}) \leq (9/10)^{N/5-2}.$$

The two distributions converge exponentially fast as N grows.

Proof. In the product program `dynkin0`, we want to bound the probability that l_1 and l_2 are disjoint; i.e., the probability that the two processes never meet. We proceed in two steps. First, it is not hard to show that $|x_1 - x_2| < 10$ is an invariant. Thus, at each iteration, there is a $1/10$ chance that the lagging process hits the leading process. Second, each process moves at most 10 positions each iteration and we finish when both processes reach the end, so there are at least $2(N/10 - 1) = N/5 - 2$ chances to hit. Therefore, in the product we can show

$$\Pr[(l_1 \cap l_2) = \emptyset] \leq (9/10)^{N/5-2}. \quad (4)$$

We can then conclude by Proposition 9. \square

To highlight the quantitative information verified by our approach, we note that the corresponding theorem for random walks (Theorem 11) shows that the total variation distance between two random walks decreases as $O(1/\sqrt{T})$. In contrast, Theorem 12 shows that total variation distance between two Dynkin processes converges as $O(0.9^N)$, giving much faster mixing (exponentially fast instead of polynomially fast).

Remark. The Dynkin process is inspired by the following two-player game. Both players pick a starting position. There is one sequence of random cards that is shared by both players, and players look at the card at their current position and move forward that many spaces. The random cards are shared, so a player samples the card only if the other player has not yet visited the position. If the other player has already landed on the position, the later player looks at the revealed card and moves forward.

While the Dynkin process samples every card that the player lands on, it is not hard to see that the product program exactly models the two-player game. More specifically, the product interleaves the players so that at each turn, the player that is lagging behind makes the next move. By this scheduling, as long as one player is lagging behind, the players have not landed on the same position and so each player lands on unseen cards and draws random samples to reveal. Once the players meet, the product program makes the same move for both players.

For this game, Eq. (4) bounds the probability that the players do not land on the same final position from any two initial positions. This result is the basis of the magic trick called *Dynkin's card trick* or the *Kruskal count*. If one player is the magician and the other player is a spectator, if the spectator starts at a secret position and runs the process mentally, the magician can guess the correct final position with high probability by starting at any position and counting along.

4.3 Path coupling

So far, we have seen how to prove convergence of probabilistic processes by constructing a coupling, reasoning about the probability of the processes not meeting, and applying the coupling theorem. While this reasoning is quite powerful, for more complicated processes it may be difficult to directly construct a coupling that shows fast mixing. For example, it can be difficult to find and reason

about a coupling on the distributions from two states s, s' if there are many transitions apart in the Markov chain.

To address this problem, Bubley and Dyer [23] proposed the *path coupling* technique, which allows us to consider just pairs of *adjacent* states, that is states where s can transition in one step to s' . Roughly speaking, if we can give a good coupling on the distributions from two adjacent states for every pair of adjacent states, then path coupling shows that the state distributions started from two arbitrary states converge quickly.

As the name suggests, path coupling considers paths of states in a probabilistic process. For this to make sense, we need to equip the state space with additional structure. For the basic setup, let Ω be a finite set of states and suppose that we have a metric $d : \Omega \times \Omega \rightarrow \mathbb{N}$. We require that d is a *path metric*: if $d(s, s') > 1$, then there exists $s'' \neq s, s'$ such that $d(s, s') = d(s, s'') + d(s'', s')$. Two states are said to be *adjacent* if $d(s, s') = 1$. We will assume that the *diameter* of the state space, i.e. the distance between any two states, is finite: $\Delta < \infty$. The Markov chain is then defined by iterating a transition function $f : \Omega \rightarrow \text{Distr}(\Omega)$.

The main idea behind path couplings is that if we can couple the distributions from any two adjacent states, then there exists a coupling for the distributions from two states at an arbitrary distance, constructed by piecing together the couplings along the path. Furthermore, if the expected distance between states contracts under the coupling on adjacent states, i.e., the resulting expected distance is strictly less than 1, then the same holds for the coupling on two states at any distance. More formally, we have the following main theorem of path coupling.

Lemma 13 (Bubley and Dyer [23]). *Consider a Markov chain with transition function f over a set Ω with diameter at most Δ . Suppose that for any two states s and s' such that $d(s, s') = 1$, there exists a coupling μ of $f(s), f(s')$ such that*

$$\mathbb{E}_{(r,r') \sim \mu}[d(r, r')] \leq \beta.$$

Then, starting from any two states s and s' and running t steps of the chain, there is a coupling μ of $f^t(s), f^t(s')$ such that

$$\text{TV}(f^t(s), f^t(s')) \leq \Pr_{(r,r') \sim \mu}[r \neq r'] \leq \beta^t \Delta.$$

This lemma applies for all β , but is most interesting for $\beta < 1$ when it implies that the Markov chain mixes quickly. With the main theorem in hand, we will show how to verify the conditions for path coupling on two examples from Bubley and Dyer [23].

4.3.1 Graph coloring: the Glauber dynamics.

Our first example is a Markov chain to provide approximately uniform samples from the set of colorings of a finite graph; it was first analyzed by Jerrum [33], our analysis follows Bubley and Dyer [23]. Recall that a finite *graph* G is defined by a finite set of vertices V , and a symmetric relation E relating pairs vertices that are connected by an edge; we will let $\mathcal{N}_G(v) \subseteq V$ denote the set of neighbors of v , i.e. the set vertices that have an edge to v in G . Let C be the set of *colors*; throughout, we assume that C is finite. A *coloring* w of G is a map from V to C . A coloring is *valid* (sometimes called *proper*) if all neighboring vertices have different colors: for all $v' \in \mathcal{N}_G(v)$ we have $w(v) \neq w(v')$. We write $\mathcal{V}_G(w)$ if w is a valid coloring.

The following program $\text{glauber}(T)$ models T steps of the *Glauber dynamics* in statistical physics:

```

 $\iota \leftarrow 0;$ 
while  $\iota < T$  do;
   $v \stackrel{\$}{\leftarrow} V; c \stackrel{\$}{\leftarrow} C;$ 
  if  $\mathcal{V}_G(w[v \mapsto c])$  then  $w \leftarrow w[v \mapsto c]$ 
   $\iota \leftarrow \iota + 1$ 

```

Informally, the algorithm starts from a valid coloring w and iteratively modifies it by sampling uniformly a vertex v and a color c , recoloring the v with c if this continues to be a valid coloring.

We want to measure the statistical distance between two executions of the process starting from two initial colorings w_1 and w_2 . There are two natural approaches. The first option is to couple the two copies of *glauber* directly, analyze the product program and apply Proposition 9. The problem is that when the two colorings are far apart, it is hard to reason about how the processes might meet under a coupling; Jerrum [33] adopted this strategy, but the resulting proof is dense and complex.

The second, far simpler option is to apply path coupling. Here, we build a product for just one iteration of the loop, and it suffices to consider cases where the two initial states are adjacent. This drastically simplifies the coupling and analysis of the product program, so we adopt this approach here. For the sake of clarity, we adapt the transition function so that its output is stored in a fresh variable w' , and call the resulting program glauber^\dagger :

```

 $v \stackrel{\$}{\leftarrow} V; c \stackrel{\$}{\leftarrow} C;$ 
if  $\mathcal{V}_G(w[v \mapsto c])$  then  $w' \leftarrow w[v \mapsto c]$  else  $w' \leftarrow w$ 

```

Note that $\text{glauber}^\dagger; w \leftarrow w'; \iota \leftarrow \iota + 1$ is semantically equivalent to the loop body of *glauber*.

To apply the path coupling theorem Lemma 13, we need to define a path metric on Ω and construct a coupling for the process started from two adjacent states. For the path metric, we define the distance $d(w_1, w_2)$ to be the Hamming distance: the number of vertices where w_1 and w_2 provide different colors. We say two states are *adjacent* if $d(w_1, w_2) = 1$; these states differ in the color of exactly one vertex. In order to apply path coupling, we need to find a coupling of the transition function on adjacent states such that the expected distance shrinks. We first build the coupling using $\times \text{pRHL}$. Specifically, we prove

$$\left\{ d(w_1, w_2) = 1 \right\} \frac{\text{glauber}_1^\dagger}{\text{glauber}_2^\dagger} \left\{ d(w'_1, w'_2) \leq 2 \right\} \rightsquigarrow c_0 \quad (5)$$

where c_0 is the following program:

```

 $v_1, v_2 \stackrel{\$}{\leftarrow} \mathcal{D}^\neq(V);$ 
if  $v_1 \in \mathcal{N}_G(v_0)$  then  $c_1, c_2 \stackrel{\$}{\leftarrow} \mathcal{D}_{\pi^{ab}}(C)$  else  $c_1, c_2 \stackrel{\$}{\leftarrow} \mathcal{D}^\neq(C)$ 
if  $\mathcal{V}_G(w_1[v_1 \mapsto c_1])$  then  $w'_1 \leftarrow w_1[v_1 \mapsto c_1]$  else  $w'_1 \leftarrow w_1$ 
if  $\mathcal{V}_G(w_2[v_2 \mapsto c_2])$  then  $w'_2 \leftarrow w_2[v_2 \mapsto c_2]$  else  $w'_2 \leftarrow w_2$ 

```

We briefly sketch how to prove the judgment. Note that the two states must agree at all vertices, except at a single vertex v_0 . Let $w_1(v_0) = a$ and $w_2(v_0) = b$. First, we couple the vertex sampling with the rule [RAND] using the identity coupling, ensuring $v_1 = v_2$. Then, we use the rule [CASE] to perform a case analysis on the sampled vertex, call it v . If v is a neighbor of the differing vertex v_0 , we use the rule [RAND] and the transposition bijection $\pi^{ab} : C \rightarrow C$ defined by the clauses:

$$\pi^{ab}(a) = b \quad \pi^{ab}(b) = a \quad \pi^{ab}(x) = x \quad \text{otherwise}$$

to ensure that $c_1 = \pi^{ab}(c_2)$. Otherwise, we use the rule [RAND] and the identity coupling to ensure $c_1 = c_2$. By applying the one-sided rules for conditionals to the left and the right sides ([COND-L] and [COND-R]), we can conclude the derivation.

Next, we must reason about the expected value of the distance between w'_1 and w'_2 after executing the product program.

Lemma 14. *Let $n = |V|$ and $k = |C|$, and suppose that the graph G has degree bounded by D . That is, for any $v \in V$, there are at most D vertices v' such that $E(v, v')$. If $k \geq D$, then there is a coupling μ the distributions after running glauber^\dagger on adjacent states such that*

$$\mathbb{E}_{(w_1, w_2) \sim \mu}[d(w'_1, w'_2)] \leq 1 - 1/n + 2D/kn.$$

Proof. Let w_1 and w_2 be adjacent states. We must bound the expected distance between the states w'_1, w'_2 in the product program. Let $d' = d(w'_1, w'_2)$, we have:

$$\begin{aligned} \mathbb{E}_{(w_1, w_2) \sim \mu}[d'] &= 0 \cdot \Pr[d' = 0] + 1 \cdot \Pr[d' = 1] + 2 \cdot \Pr[d' = 2] \\ &= 1 - \Pr[d' = 0] + \Pr[d' = 2] \\ &\leq 1 - \Pr[v_1 = v_0 \wedge \mathcal{V}_G(w')] + \Pr[v_1 \in \mathcal{N}_G(v_0) \wedge c_1 = b] \\ &\leq 1 - \frac{1}{n} \left(1 - \frac{D}{k}\right) + \frac{D}{nk} = 1 - \frac{1}{n} + \frac{2D}{nk}. \end{aligned}$$

where $w' = w_1[v_0 \mapsto c_1]$. The first equality holds because the distance between the two resulting coloring will be at most 2 by judgment Eq. (5). The second equality holds since $1 = \Pr[d' = 0] + \Pr[d' = 1] + \Pr[d' = 2]$. The second to last step follows since each vertex has at most D neighbors, so there are at least $k - D$ valid colors at any vertex. \square

Applying the path coupling lemma (Lemma 13), noting that the diameter is n since there are n vertices, proves that the Glauber dynamics mixes quickly if there are sufficiently many colors k .

Theorem 15. *Consider the Glauber dynamics on k colors with a graph G with n vertices and degree at most D , and suppose $k \geq 2D + 1$. Then, for some constant $\beta < 1$,*

$$\text{TV}(\llbracket \text{glauber}(T) \rrbracket_{m_1}, \llbracket \text{glauber}(T) \rrbracket_{m_2}) \leq \beta^T n$$

for any two initial memories m_1 and m_2 containing valid colorings.

This theorem recovers the result by Bubley and Dyer [23]; this is the key step to showing that running Glauber dynamics for a small number of steps and taking a sample is almost equivalent to drawing a uniformly random sample from all proper colorings of the graph.

4.3.2 Independent sets: the conserved hard-core model

Our second example is from graph theory and statistical physics, modeling the evolution of a physical system in the *conserved hard-core lattice gas* (CHLG) model [23]. Suppose that we have a set P of particles, and we have a graph G . A *placement* is a map $w : P \rightarrow V$, placing each particle at a vertex of the graph. We wish to place the particles so that each vertex has at most one particle, and no two particles are located at adjacent vertices; we call such a placement *safe* and denote it by $\mathcal{S}_G(w)$. For a specific graph, there could be multiple safe placements.

If we want to sample a uniformly random safe placement, we can use a simple Markov chain. We take the state space $\Omega = P \rightarrow V$ to be the set of placements. Again, we take G and P to be finite. We start using a safe initial placement. Each step, we sample a particle p from P and a vertex v from V uniformly at random and try to move p to v . If $w[p \mapsto v]$ is safe, then we make this the new placement; otherwise, we keep the same placement. We can model T steps of this dynamics with the following program $\text{chlg}(T)$:

```

 $\iota \leftarrow 0;$ 
while  $\iota < T$  do;
   $p \xleftarrow{\$} P; v \xleftarrow{\$} V;$ 
  if  $\mathcal{S}_G(w[p \mapsto v])$  then  $w \leftarrow w[p \mapsto v]$ 
   $\iota \leftarrow \iota + 1$ 

```

As in the previous example, we adapt the loop body to form chlg^\dagger :

```

 $p \xleftarrow{\$} P; v \xleftarrow{\$} V;$ 
if  $\mathcal{S}_G(w[p \mapsto v])$  then  $w' \leftarrow w[p \mapsto v]$  else  $w' \leftarrow w$ 

```

Like the graph coloring sampler, we take the path metric on placements to be Hamming distance and try to find a coupling on the distributions from adjacent initial placements.

Lemma 16. *Let $s = |P|$ and $n = |V|$, and suppose that the graph G has degree bounded by D . Starting from any two adjacent safe placements w_1 and w_2 , there is a coupling μ on the distributions after one step such that*

$$\mathbb{E}_{(w_1, w_2) \sim \mu}[d(w'_1, w'_2)] \leq \left(1 - \frac{1}{s}\right) \left(\frac{3(D+1)}{n}\right).$$

Proof. Let $\text{chlg}_1^\dagger, \text{chlg}_2^\dagger$ be two copies of the transition function, with variables tagged. Consider two adjacent placements w_1 and w_2 . We will sketch how to couple the transitions.

We use rule [RAND] twice to couple the particle and vertex samplings with the identity coupling, ensuring $p_1 = p_2$ and $v_1 = v_2$. Then, we can apply the one-sided rules for conditionals to the left and the right sides ([COND-L] and [COND-R]) to conclude the following judgment:

$$\left\{ d(w_1, w_2) = 1 \right\} \begin{array}{l} \text{chlg}_1^\dagger \\ \text{chlg}_2^\dagger \end{array} \left\{ d(w_1, w_2) \leq 2 \right\} \rightsquigarrow c_0,$$

where c_0 is the following product program:

```

 $(p_1, p_2) \xleftarrow{\$} \mathcal{D}^=(P); (v_1, v_2) \xleftarrow{\$} \mathcal{D}^=(V);$ 
if  $\mathcal{S}_G(w_1[p_1 \mapsto v_1])$  then  $w'_1 \leftarrow w_1[p_1 \mapsto v_1]$  else  $w'_1 \leftarrow w_1$ 
if  $\mathcal{S}_G(w_2[p_2 \mapsto v_2])$  then  $w'_2 \leftarrow w_2[p_2 \mapsto v_2]$  else  $w'_2 \leftarrow w_2$ 

```

Then, we can bound the expected distance between w'_1, w'_2 in the product program. Let $d' = d(w'_1, w'_2)$, we have:

$$\begin{aligned} \mathbb{E}_{(w_1, w_2) \sim \mu}[d'] &= 1 - \Pr[d' = 0] + \Pr[d' = 2] \\ &= 1 - \Pr[p = p_0 \wedge \mathcal{S}_G(w_1[p \mapsto v])] \\ &\quad + \Pr[p \neq p_0 \wedge \neg(\mathcal{S}_G(w_1[p \mapsto v]) \Leftrightarrow \mathcal{S}_G(w_2[p \mapsto v]))] \\ &\leq 1 - \Pr[p = p_0 \wedge \mathcal{S}_G(w_1[p \mapsto v])] \\ &\quad + \Pr[p \neq p_0 \wedge \neg(\mathcal{S}_G(w_1[p \mapsto v]) \wedge \mathcal{S}_G(w_2[p \mapsto v]))] \end{aligned}$$

We can bound the two probability terms. For the first term, we know that the probability of selecting $p = p_0$ is $1/s$, and the probability that p is safe at v if it avoids all other points (at most $s-1$) and all the neighbors of the other points (at most $(s-1)D$); this probability is the same for both placements w_1 and w_2 , since the two placements are identical on points besides p_0 .

For the second term, we know that the probability of selecting $p \neq p_0$ is $1 - 1/s$, and p is not safe at v in placement w_1 or in w_2 if we select the position a, b , or one of their neighbors. Putting everything together, we can conclude:

$$\begin{aligned} \mathbb{E}_{(w_1, w_2) \sim \mu}[d(w'_1, w'_2)] &\leq 1 - \frac{1}{s} \left(1 - \frac{(s-1)(D+1)}{n}\right) \\ &\quad + (s-1) \left(\frac{|\mathcal{N}_G(a)| + 1 + |\mathcal{N}_G(b)| + 1}{sn}\right) \\ &\leq 1 - \frac{1}{s} \left(1 - \frac{(s-1)(D+1)}{n}\right) + (s-1) \left(\frac{2(D+1)}{sn}\right) \\ &= \left(1 - \frac{1}{s}\right) \left(\frac{3(D+1)}{n}\right). \quad \square \end{aligned}$$

Applying the path coupling lemma (Lemma 13) shows that if we iterate the transition function on two initial placements, the resulting distributions on placements converge quickly.

Theorem 17. *Consider the conserved lattice gas model with $s = |P|$ particles on a graph G with $n = |V|$ vertices and degree at most D , where $s \leq n/3(D+1) + 1$. Then, for a constant $\beta < 1$,*

$$\text{TV}(\llbracket \text{chlg} \rrbracket_{m_1}, \llbracket \text{chlg} \rrbracket_{m_2}) \leq \beta^T s$$

for any two initial memories m_1, m_2 containing safe placements.

Remark. This theorem is slightly weaker than the corresponding result by Bubley and Dyer [23], who prove rapid mixing under the weaker condition $s \leq n/2(D+1)+1$. Roughly, they use the maximal coupling on the two transition distributions, giving a tighter analysis and better bound. It is also possible use the maximal coupling in \times pRHL, but the corresponding specification of the coupling distribution would be proved as part of soundness of the logic, rather than as a property of a probabilistic program.

5. Application: loop optimizations

Program equivalence is one of the original motivation for relational program logics [17]. In this section, we demonstrate the effectiveness of our logic using several examples of exact and approximate program equivalence. Our first example is a loop transformation which originates from the literature on parallelizing compilers but also has applications in computer-aided cryptography. Our second example is drawn from the recent literature on approximate computing, and is an instance of loop perforation.

5.1 Loop strip-mining

Loop *strip-mining* (or *loop sectioning*) is a transformation that turns a loop into a nested loop. While the transformation originates from literature on parallelizing compilers and is primarily used to take advantage of vectorized instructions, it is also useful for formally proving the computational security of certain cryptographic constructions. The following example, Fig. 7, is inspired from a proof of indifferenciability of the SHA3 hash function [24]⁴ Using the rule for **while** loops, we can prove the following:

$$\left\{ x_1 = x_2 \right\} \begin{array}{c} c_1 \\ c_2 \end{array} \left\{ x_1 = x_2 \right\} \rightsquigarrow c$$

The crux of the proof is applying the [WHILE] rule with $k_1 = 1$ and $k_2 = M$, and $e = v_1 < N$, and $p_0 = \top$, and $p_1 = p_2 = \perp$ and an invariant Ψ which strengthens the assertion $x_1 = x_2$ mainly by adding $l_2 = v_1 * M$. Side conditions using p_1 and p_2 are trivial to prove (using the FALSE rule) since they have \perp in hypothesis. It remains to check the premise for p_0 , but we now have two synchronized loops; we can use the [STRUCT] rule to remove the conditional on p_0 which is always true in this case.

5.2 Loop perforation

Loop perforation [44, 51] is a program transformation that delivers good trade-offs between performance and accuracy, and is practical in many applications, including image and audio processing, simulations and machine-learning. Informally, loop perforation transforms a loop that performs n iterations of its body into a loop that performs $m < n$ iterations of its body, followed by a simple post-processing statement. Figure 8 shows an example of loop perforation inspired from a financial analysis application, called swaptions. In this example, every other loop iteration is skipped, and the post-processing statement simply multiplies by 2 the value s computed by the optimized loop. As for the previous example, we can prove the following judgment:

$$\left\{ \top \right\} \begin{array}{c} c_1 \\ c_2 \end{array} \left\{ s_1 = s_2 \right\} \rightsquigarrow c$$

The product program can be built using the [WHILE] rule. We use $e = v_1 < 2 * n$, $p_0 = \top$, $p_1 = p_2 = \perp$ and $k_1 = 2$ and $k_2 = 1$, and the invariant is $v_1 = 2 * v_2$. The invariant allows to show, using the

⁴for simplicity, the programs use an operator f which takes randomness as an argument (note that the value r is sampled immediately before the assignment using f), although in the proof of the SHA3 hash function f is a procedure call whose body performs random samplings.

<p>Nested loop:</p> <pre> v1 ← 0; while v1 < N do j1 ← 0; while j1 < M do l1 ← v1 * M + j; r1 ← μ; x1 ← f(l1, x1, r1); j1 ← j1 + 1; v1 ← v1 + 1; l1 ← N * M; </pre>	<p>Single loop:</p> <pre> l2 ← 0; while l2 < N * M do v2 ← l2 / M; j2 ← l2 % M; r2 ← μ; x2 ← f(l2, x2, r2); l2 ← l2 + 1; v2 ← N; j2 ← M; </pre>
<p>Product program:</p> <pre> v1 ← 0; l2 ← 0; while v1 < N do j1 ← 0; while j1 < M do l1 ← v1 * M + j1; v2 ← l2 / M; j2 ← l2 % M; (r1, r2) ← D(μ); x1 ← f(l1, x1, r1); x2 ← f(l2, x2, r2); j1 ← j1 + 1; l2 ← l2 + 1; v1 ← v1 + 1; l1 ← N * M; v2 ← N; j2 ← M; </pre>	

Figure 7. Loop strip-mining

<p>Original program:</p> <pre> s1 ← 0; for (v1 ← 0, v1 < 2 * n, v1 ← v1 + 1) do x1 ← μ; s1 ← s1 + x1; </pre>	<p>Perforated program:</p> <pre> s2 ← 0; for (v2 ← 0, v2 < n, v2 ← v2 + 1) do x2 ← μ; s2 ← s2 + x2; s2 ← 2 * s2; </pre>
<p>Product program:</p> <pre> s1 ← 0; s2 ← 0; v1 ← 0; v2 ← 0; while v1 < 2 * n do x1, x2 ← D(μ); s1 ← s1 + x1; s2 ← s2 + x2; v1 ← v1 + 1; x1 ← μ; s1 ← s1 + x1; v1 ← v1 + 1; v2 ← v2 + 1; s2 ← 2 * s2; </pre>	

Figure 8. Loop perforation

[STRUCT] rule, that the loop on k_1 (denoted by $c_1^{e_1, k_1}$ in [WHILE]) perform exactly 2 iterations. Using the product, one can also analyze the (probabilistic) accuracy rate of the transformed program, using concentration bounds to achieve a more precise bound.

Finally, in some applications the number of iterations performed by the perforated loop is probabilistic; for instance, the program

```

k ← factors(n); s ← 0;
for (i ← 0, i < n, i ← i + k) do
  x ← μ; s ← s + x;
s ← k * s;

```

selects uniformly at random a factor k of the original number n of iterations, and performs n/k iterations. It is possible to relate the original and the perforated loop, using the [WHILE] rule as before.

5.3 Other optimizations and program transformations

Barthe et al. [6] define an inductive method for building valid product programs, and use their method for validating instances of loop optimizations. Their method combines a rule for each program construction and a rule akin to our [STRUCT] rule. Despite this similarity, the two methods are fundamentally different: their treatment of **while** loops is restricted to synchronized executions. As a consequence, their structural rule is based on a more advanced refinement relation between programs. Nevertheless, we can reproduce all their examples in our formalism, taking advantage of our more powerful rule for loops.

6. Related work

Relational logics can be seen as a proof-theoretical counterpart of semantics-based relational methods such as logical relations. Under this view, our logic bears strong similarities with proof-relevant logical relations [18]. As for proof-relevant logical relations, we expect that manipulating explicit witnesses rather than existentially quantified can help developing meta-theoretical studies of our logic.

Much of the recent work on product programs and relational logics has been motivated by applications to security and compiler correctness. For instance, Barthe et al. [4] explore self-composition for a variety of programming languages and show that it induces a sound and complete reduction of an information flow policy to a safety property. Independently, Darvas et al. [27] consider self-composition and deductive verification based on dynamic logic, also for verifying information flow policies. Later, Terauchi and Aiken [53] introduce the class of 2-safety properties and show a reduction from 2-safety to safety of the self-composed program. Their reduction is more efficient than self-composition as it selectively applies self-composition or a synchronous product construction akin to cross-products (described below). Further improvements appear in Kovács et al. [35], Müller et al. [46]. In a related work, Beringer and Hofmann [20] observe that one can encode 2-safety properties in standard Hoare logic, provided that assertions are sufficiently expressive to model the denotational semantics of programs. Beringer [19] further refines this approach, by introducing the notion of relational decomposition.

Zaks and Pnueli [57] define a cross-product construction, which is well-suited for reasoning about programs with identical control-flow. Barthe et al. [6] generalize the notion of cross-product by proposing a more general notion of product program which subsumes self-composition and cross-products, and show how it enables validation of common loop optimizations. Specifically, they define an inductive relation for proving that c is a valid product for c_1 and c_2 ; informally, their rules closely follow those of our system (for deterministic constructs), except for the general rule for **while** loops; instead, they use a rule that is closer to the pRHL rule, and a rule akin to our [STRUCT] rule, with a much stronger relationship between programs in order to compensate for the lack of generality of their rule. Barthe et al. [14] carry a more precise study of the relative expressiveness of product program constructions and relational program logics. In a different thread of work, Barthe et al. [8] generalize the notion of product program so that it supports verification of refinement properties (modeled by universal quantification over runs of the first program and existential quantification over runs of the second program), as well as the 2-safety properties (modeled by universal quantification over runs of the first and the second programs). These constructions are focused on non-probabilistic programs. Motivated by applications to differential privacy, Barthe et al. [11] define a specialized product construction from probabilistic programs to deterministic programs, so that the original program is differentially private, provided its deterministic product program satisfies some Hoare specification. To the best of our knowledge,

this is the sole product construction that goes from a probabilistic language to a deterministic one.

Benton [17] and Yang [56] were among the first to consider relational program logics that support direct reasoning about two programs. Benton [17] introduces Relational Hoare Logic, proves correctness of several program transformations, and soundly embeds a type system for information flow security into his logic. Yang [56] defines Relational Separation Logic and proves the equivalence between Depth-First Search and the Schorr-Waite algorithm. Barthe et al. [5] develop probabilistic Relational Hoare Logic, and use it for proving computational security of cryptographic constructions. In a follow-up work, Barthe et al. [7] develop an approximate variant of probabilistic Relational Hoare Logic, and verify differential privacy of several algorithms. More recently, Sousa and Dillig [52] propose Cartesian Hoare Logic, an extension of relational Hoare logic to an arbitrary finite number of executions.

Several authors have considered relational logics for higher-order programs. Nanevski et al. [47] develop a relational logic to reason about information flow properties of a higher-order language with mutable state. Ghani et al. [31] introduce a relational type theory, and a supporting categorical model, for reasoning about parametricity. Barthe et al. [10] propose a relational extension of a subset of the F^* ; in a follow-up work, Barthe et al. [13] combine a relational refinement type system with a graded monad which they use for modeling differentially private computations.

There are several works that develop more specialized program logics for analyzing relational properties of programs. For instance, Amtoft et al. [2] introduce independence assertions and a supporting program logic for proving information flow security. In a similar way, Chaudhuri et al. [26] propose a logical approach for proving continuity properties of programs, and Carbin et al. [25] develop a logical approach for reasoning about the reliability of approximate computation.

Further afield, there has been a significant amount of work on semantical methods for probabilistic programs and processes initiated by Kozen [36], see e.g. [21, 38, 40, 49] for some recent developments. In conjunction with these semantics, research in deductive verification methods for non-relational properties of probabilistic programs is an active area of research; examples include [28, 37, 45, 48] to cite only a few systems.

7. Conclusion and future directions

We have introduced \times pRHL, a new program logic that deepens the connection between probabilistic couplings and relational verification of probabilistic programs in two different ways. First, \times pRHL broadens the class of couplings supported by relational verification. Second, \times pRHL derivations explicitly build a probabilistic product program, which can be used to analyze mixing times. We have shown the flexibility of our approach on several examples.

There is ample room for future work. On the theoretical side, it would be interesting to extend \times pRHL to handle continuous distributions as was recently done by Sato [50] for apRHL. Also, we believe that we are just scratching the surface of probabilistic product programs; there should be many further applications, notably in relationship with path couplings, in domains such as Brownian motion [42], molecular evolution [30], and anonymity [32]. On the more practical side, it would be natural to integrate \times pRHL in EasyCrypt [9], a proof-assistant used for reasoning about computational security of cryptographic constructions, and currently the main implementation of pRHL. We expect that several proofs of cryptographic constructions can be simplified using the new loop rule, and we are also planning to use the loop rule in an ongoing formalization of indifferenciability of the SHA3 standard for hash functions.

References

- [1] D. J. Aldous and H. Thorisson. Shift-Coupling. *Stochastic Proc. Appl.*, 44:1–14, 1993.
- [2] T. Amtoft, S. Bandhakavi, and A. Banerjee. A logic for information flow in object-oriented programs. In J. G. Morrisett and S. L. P. Jones, editors, *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, Charleston, South Carolina, pages 91–102, 2006.
- [3] L. Avena, M. Heydenreich, F. den Hollander, E. Verbitskiy, and W. van Zuijlen. Random walks (lecture notes). Technical report, Mathematical Institute, Leiden University.
- [4] G. Barthe, P. D’Argenio, and T. Rezk. Secure Information Flow by Self-Composition. In R. Foccardi, editor, *IEEE Computer Security Foundations Workshop (CSFW)*, Pacific Grove, California, pages 100–114, 2004.
- [5] G. Barthe, B. Grégoire, and S. Zanella-Béguelin. Formal certification of code-based cryptographic proofs. In *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, Savannah, Georgia, pages 90–101, New York, 2009.
- [6] G. Barthe, J. M. Crespo, and C. Kunz. Relational verification using product programs. In M. J. Butler and W. Schulte, editors, *International Symposium on Formal Methods (FM)*, Limerick, Ireland, volume 6664 of *Lecture Notes in Computer Science*, pages 200–214. Springer, 2011.
- [7] G. Barthe, B. Köpf, F. Olmedo, and S. Z. Béguelin. Probabilistic relational reasoning for differential privacy. In J. Field and M. Hicks, editors, *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, Philadelphia, Pennsylvania, pages 97–110, 2012.
- [8] G. Barthe, J. M. Crespo, and C. Kunz. Beyond 2-safety: Asymmetric product programs for relational program verification. In S. N. Artëmov and A. Nerode, editors, *Symposium on the Logical Foundations of Computer Science (LFCS)*, San Diego, California, volume 7734 of *Lecture Notes in Computer Science*, pages 29–43. Springer, 2013.
- [9] G. Barthe, F. Dupressoir, B. Grégoire, C. Kunz, B. Schmidt, and P. Strub. EasyCrypt: A tutorial. In A. Aldini, J. Lopez, and F. Martinelli, editors, *Foundations of Security Analysis and Design VII (FOSAD)*, volume 8604 of *Lecture Notes in Computer Science*, pages 146–166. Springer, 2013. Tutorial Lectures.
- [10] G. Barthe, C. Fournet, B. Grégoire, P. Strub, N. Swamy, and S. Z. Béguelin. Probabilistic relational verification for cryptographic implementations. In S. Jagannathan and P. Sewell, editors, *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, San Diego, California, pages 193–206, 2014.
- [11] G. Barthe, M. Gaboardi, E. J. Gallego Arias, J. Hsu, C. Kunz, and P.-Y. Strub. Proving differential privacy in Hoare logic. In *IEEE Computer Security Foundations Symposium (CSF)*, Vienna, Austria, pages 411–424, 2014.
- [12] G. Barthe, T. Espitau, B. Grégoire, J. Hsu, L. Stefanescu, and P.-Y. Strub. Relational reasoning via probabilistic coupling. In *International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR)*, Suva, Fiji, volume 9450, pages 387–401, 2015.
- [13] G. Barthe, M. Gaboardi, E. J. G. Arias, J. Hsu, A. Roth, and P. Strub. Higher-order approximate relational refinement types for mechanism design and differential privacy. In S. K. Rajamani and D. Walker, editors, *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, Mumbai, India, pages 55–68, 2015.
- [14] G. Barthe, J. M. Crespo, and C. Kunz. Product programs and relational program logics. *Journal of Logical and Algebraic Methods in Programming*, 2016.
- [15] G. Barthe, M. Gaboardi, B. Grégoire, J. Hsu, and P.-Y. Strub. Proving differential privacy via probabilistic couplings. In *IEEE Symposium on Logic in Computer Science (LICS)*, New York, New York, 2016.
- [16] G. Barthe, M. Gaboardi, B. Grégoire, J. Hsu, and P.-Y. Strub. A program logic for union bounds. In *International Colloquium on Automata, Languages and Programming (ICALP)*, Rome, Italy, 2016.
- [17] N. Benton. Simple relational correctness proofs for static analyses and program transformations. In N. D. Jones and X. Leroy, editors, *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, Venice, Italy, pages 14–25, 2004.
- [18] N. Benton, M. Hofmann, and V. Nigam. Proof-relevant logical relations for name generation. In M. Hasegawa, editor, *International Conference on Typed Lambda Calculi and Applications (TLCA)*, Eindhoven, The Netherlands, volume 7941 of *Lecture Notes in Computer Science*, pages 48–60. Springer, 2013.
- [19] L. Beringer. Relational decomposition. In M. C. J. D. van Eekelen, H. Geuvers, J. Schmaltz, and F. Wiedijk, editors, *Interactive Theorem Proving (ITP)*, Nijmegen, The Netherlands, volume 6898 of *Lecture Notes in Computer Science*, pages 39–54. Springer, 2011.
- [20] L. Beringer and M. Hofmann. Secure information flow and program logics. In *IEEE Computer Security Foundations Symposium (CSF)*, Venice, Italy, pages 233–248, 2007.
- [21] A. Bizjak and L. Birkedal. Step-indexed logical relations for probability. In A. M. Pitts, editor, *International Conference on Foundations of Software Science and Computation Structures (FoSSaCS)*, London, England, volume 9034 of *Lecture Notes in Computer Science*, pages 279–294. Springer, 2015.
- [22] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011.
- [23] R. Bubley and M. Dyer. Path coupling: A technique for proving rapid mixing in Markov chains. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, Miami Beach, Florida, pages 223–231, 1997.
- [24] A. Canteaut, T. Fuhr, M. Naya-Plasencia, P. Paillier, J. Reinhard, and M. Videau. A unified indifferenciability proof for permutation- or block cipher-based hash functions. *IACR Cryptology ePrint Archive*, 2012: 363, 2012.
- [25] M. Carbin, D. Kim, S. Misailovic, and M. C. Rinard. Proving acceptability properties of relaxed nondeterministic approximate programs. In J. Vitek, H. Lin, and F. Tip, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Beijing, China, pages 169–180, 2012.
- [26] S. Chaudhuri, S. Gulwani, and R. Lublinerman. Continuity analysis of programs. In M. V. Hermenegildo and J. Palsberg, editors, *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, Madrid, Spain, pages 57–70, 2010.
- [27] A. Darvas, R. Hähnle, and D. Sands. A theorem proving approach to analysis of secure information flow. In D. Hutter and M. Ullmann, editors, *Security in Pervasive Computing*, volume 3450 of *Lecture Notes in Computer Science*, pages 193–209. Springer, 2005. Preliminary version in the informal proceedings of WITS 2003.
- [28] J. den Hartog. *Probabilistic extensions of semantical models*. PhD thesis, Vrije Universiteit Amsterdam, 2002.
- [29] P. Diaconis and D. Stroock. Geometric Bounds for Eigenvalues of Markov Chains. *The Annals of Applied Probability*, 1(1):36–61, Feb. 1991.
- [30] N. M. Dixit, P. Srivastava, and N. K. Vishnoi. A finite population model of molecular evolution: Theory and computation. *Journal of Computational Biology*, 19(10):1176–1202, 2012.
- [31] N. Ghani, F. N. Forsberg, and A. Simpson. Comprehensive parametric polymorphism: Categorical models and type theory. In B. Jacobs and C. Löding, editors, *International Conference on Foundations of Software Science and Computation Structures (FoSSaCS)*, Eindhoven, The Netherlands, volume 9634 of *Lecture Notes in Computer Science*, pages 3–19. Springer, 2016.
- [32] M. Gomulkiewicz, M. Klonowski, and M. Kutylowski. Rapid mixing and security of chaum’s visual electronic voting. In E. Snekkenes and D. Gollmann, editors, *European Symposium on Programming on Research in Computer Security (ESORICS)*, Gjøvic, Norway, volume 2808 of *Lecture Notes in Computer Science*, pages 132–145. Springer, 2008.
- [33] M. Jerrum. A very simple algorithm for estimating the number of k-colorings of a low-degree graph. *Random Struct. Algorithms*, 7(2): 157–166, 1995.
- [34] M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, 1989.

- [35] M. Kovács, H. Seidl, and B. Finkbeiner. Relational abstract interpretation for the verification of 2-hypersafety properties. In A. Sadeghi, V. D. Gligor, and M. Yung, editors, *ACM SIGSAC Conference on Computer and Communications Security (CCS), Berlin, Germany*, pages 211–222, 2013.
- [36] D. Kozen. Semantics of probabilistic programs. In *IEEE Symposium on Foundations of Computer Science (FOCS), San Juan, Puerto Rico*, pages 101–114, 1979.
- [37] D. Kozen. A probabilistic PDL. *Journal of Computer and System Sciences*, 30(2):162–178, 1985.
- [38] D. Kozen. Kolmogorov extension, martingale convergence, and compositionality of processes. In *IEEE Symposium on Logic in Computer Science (LICS), New York, New York*, 2016.
- [39] V. S. A. Kumar and H. Ramesh. Coupling vs. conductance for the Jerrum-Sinclair chain. *Random Struct. Algorithms*, 18(1):1–17, 2001.
- [40] U. D. Lago, D. Sangiorgi, and M. Alberti. On coinductive equivalences for higher-order probabilistic functional programs. In S. Jagannathan and P. Sewell, editors, *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), San Diego, California*, pages 297–308, 2014.
- [41] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009.
- [42] T. Lindvall. *Lectures on the coupling method*. Courier Corporation, 2002.
- [43] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [44] S. Misailovic, D. M. Roy, and M. C. Rinard. Probabilistically accurate program transformations. In E. Yahav, editor, *International Symposium on Static Analysis (SAS), Venice, Italy*, volume 6887 of *Lecture Notes in Computer Science*, pages 316–333. Springer, 2011.
- [45] C. Morgan, A. McIver, and K. Seidel. Probabilistic predicate transformers. *ACM Transactions on Programming Languages and Systems*, 18(3):325–353, 1996.
- [46] C. Müller, M. Kovács, and H. Seidl. An analysis of universal information flow based on self-composition. In C. Fournet, M. W. Hicks, and L. Viganò, editors, *IEEE Computer Security Foundations Symposium (CSF), Venice, Italy*, pages 380–393, 2015.
- [47] A. Nanevski, A. Banerjee, and D. Garg. Verification of information flow and access control policies with dependent types. In *IEEE Symposium on Security and Privacy (S&P), Oakland, California*, pages 165–179, 2011.
- [48] L. H. Ramshaw. *Formalizing the Analysis of Algorithms*. PhD thesis, Computer Science, 1979.
- [49] D. Sangiorgi and V. Vignudelli. Environmental bisimulations for probabilistic higher-order languages. In R. Bodík and R. Majumdar, editors, *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), Saint Petersburg, Florida*, pages 595–607, 2016.
- [50] T. Sato. Approximate relational hoare logic for continuous random samplings. In *Workshop on the Mathematical Foundations of Programming Semantics (MFPS), Pittsburgh, Pennsylvania*, 2016.
- [51] S. Sidiroglou-Douskos, S. Misailovic, H. Hoffmann, and M. C. Rinard. Managing performance vs. accuracy trade-offs with loop perforation. In T. Gyimóthy and A. Zeller, editors, *Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE), Szeged, Hungary*, pages 124–134, 2011.
- [52] M. Sousa and I. Dillig. Cartesian Hoare logic for verifying k-safety properties. In C. Krintz and E. Berger, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Santa Barbara, California*, pages 57–69, 2016.
- [53] T. Terauchi and A. Aiken. Secure information flow as a safety problem. In C. Hankin and I. Siveroni, editors, *International Symposium on Static Analysis (SAS), London, England*, volume 3672 of *Lecture Notes in Computer Science*, pages 352–367. Springer, 2005.
- [54] H. Thorisson. *Coupling, Stationarity, and Regeneration*. Springer, 2000.
- [55] C. Villani. *Optimal transport: old and new*. Springer Science, 2008.
- [56] H. Yang. Relational separation logic. *Theoretical Computer Science*, 375(1-3):308–334, 2007.
- [57] A. Zaks and A. Pnueli. CoVaC: Compiler Validation by Program Analysis of the Cross-Product. In J. Cuéllar, T. S. E. Maibaum, and K. Sere, editors, *International Symposium on Formal Methods (FM), Turku, Finland*, volume 5014 of *Lecture Notes in Computer Science*, pages 35–51. Springer, 2008.

Hoare Logic for Deterministic Programs

We only consider deterministic programs in this section, i.e. programs without random samplings. In that case, the semantic $\llbracket s \rrbracket_m$ of a deterministic command is either \emptyset if the program does not terminate, or of the form $\mathbb{1}_{m'}$ if it does. From that remark, we define the operator $\llbracket \cdot \rrbracket_m^*$ as follow:

$$\llbracket s \rrbracket_m^* \triangleq \begin{cases} m' & \text{if } \llbracket s \rrbracket_m = \mathbb{1}_{m'} \\ \perp & \text{otherwise} \end{cases}$$

A while loop **while** e **do** s is said to *terminate certainly under condition* Φ , written $\Phi \Downarrow$ **while** e **do** s , iff for any m s.t. $\Phi(m)$, we have $\llbracket \text{while } e \text{ do } s \rrbracket_m^* \neq \perp$.

A *Hoare triple* is a triple of the form $\{\Phi\} s \{\Psi\}$ where Φ and Ψ are assertions and s is a deterministic command. A Hoare triple is said to be valid, written $\models \{\Phi\} s \{\Psi\}$, iff for any memory m s.t. $\Phi(m)$, either $\llbracket s \rrbracket_m^* = \perp$ or $\Psi(\llbracket s \rrbracket_m^*)$. Figure 9 gives a sound inference system for Hoare triples: If $\vdash \{\Phi\} s \{\Psi\}$ is derivable from the rules of Figure 9, then $\models \{\Phi\} s \{\Psi\}$. Moreover, this system is complete for programs that only contain certainly terminating loops.

$$\begin{array}{c} \frac{}{\vdash \{\Phi\} \text{skip } \{\Phi\}} \text{[SKIP]} \quad \frac{}{\vdash \{\top\} \text{abort } \{\perp\}} \text{[ABORT]} \\ \\ \frac{}{\vdash \{\Phi[e/x]\} x \leftarrow e \{\Phi\}} \text{[ASSIGN]} \\ \\ \frac{\vdash \{\Phi \wedge e\} c_1 \{\Psi\} \quad \vdash \{\Phi \wedge \neg e\} c_2 \{\Psi\}}{\vdash \{\Phi\} \text{if } e \text{ then } c_1 \text{ else } c_2 \{\Psi\}} \text{[COND]} \\ \\ \frac{\vdash \{\Phi \wedge e\} c \{\Phi\} \quad \Phi \Downarrow \text{while } e \text{ do } c}{\vdash \{\Phi\} \text{while } e \text{ do } c \{\Phi \wedge \neg e\}} \text{[WHILE]} \\ \\ \frac{\vdash \{\Phi\} c_1 \{\Xi\} \quad \vdash \{\Xi\} c_2 \{\Psi\}}{\vdash \{\Phi\} c_1; c_2 \{\Psi\}} \text{[SEQ]} \\ \\ \frac{\vdash \{\Phi'\} c \{\Psi'\} \quad \models \Phi \implies \Phi' \quad \models \Psi' \implies \Psi}{\vdash \{\Phi\} c \{\Psi\}} \text{[CONSEQ]} \end{array}$$

Figure 9. Hoare Logic Rules for Deterministic Programs