



The Lattice-Theoretic Essence of Property Directed Reachability Analysis

Mayuko Kori^{1,2}(✉) , Natsuki Urabe² , Shin-ya Katsumata² ,
Kohei Suenaga³ , and Ichiro Hasuo^{1,2}

¹ The Graduate University for Advanced Studies (SOKENDAI), Hayama, Japan

² National Institute of Informatics, Tokyo, Japan

{mkori, urabenatsuki, s-katsumata, hasuo}@nii.ac.jp

³ Kyoto University, Kyoto, Japan

ksuenaga@fos.kuis.kyoto-u.ac.jp



Abstract. We present *LT-PDR*, a lattice-theoretic generalization of Bradley’s property directed reachability analysis (PDR) algorithm. LT-PDR identifies the essence of PDR to be an ingenious combination of verification and refutation attempts based on the Knaster–Tarski and Kleene theorems. We introduce four concrete instances of LT-PDR, derive their implementation from a generic Haskell implementation of LT-PDR, and experimentally evaluate them. We also present a categorical structural theory that derives these instances.

Keywords: Property directed reachability analysis · Model checking · Lattice theory · Fixed point theory · Category theory

1 Introduction

Property directed reachability (PDR) (also called *IC3*) introduced in [9, 13] is a model checking algorithm for proving/disproving safety problems. It has been successfully applied to software and hardware model checking, and later it has been extended in several directions, including *fbPDR* [25, 26] that uses both forward and backward predicate transformers and *PrIC3* [6] for the quantitative safety problem for probabilistic systems. See [14] for a concise overview.

The original PDR assumes that systems are given by binary predicates representing transition relations. The PDR algorithm maintains data structures called *frames* and *proof obligations*—these are collections of predicates over states—and updates them. While this logic-based description immediately yields automated tools using SAT/SMT solvers, it limits target systems to qualitative and nondeterministic ones. This limitation was first overcome by *PrIC3* [6] whose target is probabilistic systems. This suggests room for further generalization of PDR.

The authors are supported by ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603). MK is a JSPS DC fellow and supported by JSPS KAKENHI Grant (No. 22J21742). KS is supported by JST CREST Grant (No. JPMJCR2012) and JSPS KAKENHI Grant (No. 19H04084).

© The Author(s) 2022

S. Shoham and Y. Vizel (Eds.): CAV 2022, LNCS 13371, pp. 235–256, 2022.

https://doi.org/10.1007/978-3-031-13185-1_12

In this paper, we propose the first lattice theory-based generalization of the PDR algorithm; we call it *LT-PDR*. This makes the PDR algorithm apply to a wider class of safety problems, including qualitative and quantitative. We also derive a new concrete extension of PDR, namely one for Markov reward models.

We implemented the general algorithm LT-PDR in Haskell, in a way that maintains the theoretical abstraction and clarity. Deriving concrete instances for various types of systems is easy (for Kripke structures, probabilistic systems, etc.). We conducted an experimental evaluation, which shows that these easily-obtained instances have at least reasonable performance.

Preview of the Theoretical Contribution. We generalize the PDR algorithm so that it operates over an arbitrary complete lattice L . This generalization recasts the PDR algorithm to solve a general problem $\mu F \leq^? \alpha$ of over-approximating the least fixed point of an ω -continuous function $F: L \rightarrow L$ by a safety property α . This lattice-theoretic generalization signifies the relationship between the PDR algorithm and the theory of fixed points. This also allows us to incorporate quantitative predicates suited for probabilistic verification.

More specifically, we reconstruct the original PDR algorithm as a combination of two constituent parts. They are called *positive LT-PDR* and *negative LT-PDR*. Positive LT-PDR comes from a witness-based proof method by the *Knaster–Tarski fixed point theorem*, and aims to *verify* $\mu F \leq^? \alpha$. In contrast, negative LT-PDR comes from the *Kleene fixed point theorem* and aims to *refute* $\mu F \leq^? \alpha$. The two algorithms build up witnesses in an iterative and nondeterministic manner, where nondeterminism accommodates guesses and heuristics. We identify the essence of PDR to be an ingenious combination of these two algorithms, in which intermediate results on one side (positive or negative) give informed guesses on the other side. This is how we formulate LT-PDR in Sect. 3.3.

We discuss several instances of our general theory of PDR. We discuss three concrete settings: Kripke structures (where we obtain two instances of LT-PDR), Markov decision processes (MDPs), and Markov reward models. The two in the first setting essentially subsume many existing PDR algorithms, such as the original PDR [9, 13] and Reverse PDR [25, 26], and the one for MDPs resembles PrIC3 [6]. The last one (Markov reward models) is a new algorithm that fully exploits the generality of our framework.

In fact, there is another dimension of theoretical generalization: the derivation of the above concrete instances follows a *structural theory of state-based dynamics and predicate transformers*. We formulate the structural theory in the language of *category theory* [3, 23]—using especially *coalgebras* [17] and *fibrations* [18]—following works such as [8, 15, 21, 28]. The structural theory tells us which safety problems arise under what conditions; it can therefore suggest that certain safety problems are unlikely to be formulatable, too. The structural theory is important because it builds a mathematical order in the PDR literature, in which theoretical developments tend to be closely tied to implementation and thus theoretical essences are often not very explicit. For example, the theory is useful in classifying a plethora of PDR-like algorithms for Kripke structures (the original, Reverse PDR, fbPDR, etc.). See Sect. 5.1.

We present the above structural theory in Sect. 4 and briefly discuss its use in the derivation of concrete instances in Sect. 5. We note, however, that this categorical theory is not needed for reading and using the other parts of the paper.

There are other works on generalization of PDR [16, 24], but our identification of the interplay of Knaster–Tarski and Kleene is new. They do not accommodate probabilistic verification, either. See [22, Appendix A] for further discussions.

Preliminaries. Let (L, \leq) be a poset. $(L, \leq)^{\text{op}}$ denotes the opposite poset (L, \geq) . Note that if (L, \leq) is a complete lattice then so is $(L, \leq)^{\text{op}}$. An ω -chain (resp. ω^{op} -chain) in L is an \mathbb{N} -indexed family of increasing (resp. decreasing) elements in L . A monotone function $F : L \rightarrow L$ is ω -continuous (resp. ω^{op} -continuous) if F preserves existing suprema of ω -chains (resp. infima of ω^{op} -chains).

2 Fixed-points in Complete Lattices

Let (L, \leq) be a complete lattice and $F : L \rightarrow L$ be a monotone function. When we analyze fixed points of F , pre/postfixed points play important roles.

Definition 2.1. A prefixed point of F is an element $x \in L$ satisfying $Fx \leq x$. A postfix point of F is an element $x \in L$ satisfying $x \leq Fx$. We write $\mathbf{Pre}(F)$ and $\mathbf{Post}(F)$ for the set of prefixed points and postfix points of F , respectively.

The following results are central in fixed point theory. They allow us to under/over-approximate the least/greatest fixed points.

Theorem 2.2. A monotone endofunction F on a complete lattice (L, \leq) has the least fixed point μF and the greatest fixed point νF . Moreover,

1. (Knaster–Tarski [30]) The set of fixed points forms a complete lattice. Furthermore, $\mu F = \bigwedge \{x \in L \mid Fx \leq x\}$ and $\nu F = \bigvee \{x \in L \mid x \leq Fx\}$.
2. (Kleene, see e.g. [5]) If F is ω -continuous, $\mu F = \bigvee_{n \in \mathbb{N}} F^n \perp$. Dually, if F is ω^{op} -continuous, $\nu F = \bigwedge_{n \in \mathbb{N}} F^n \top$. \square

Theorem 2.2.2 is known to hold for arbitrary ω -cpols (complete lattices are their special case). A generalization of Theorem 2.2.2 is the Cousot–Cousot characterization [11], where F is assumed to be monotone (but not necessarily ω -continuous) and we have $\mu F = F^\kappa \perp$ for a sufficiently large, possibly transfinite, ordinal κ . In this paper, for the algorithmic study of PDR, we assume the ω -continuity of F . Note that ω -continuous F on a complete lattice is necessarily monotone.

We call the ω -chain $\perp \leq F \perp \leq \dots$ the initial chain of F and the ω^{op} -chain $\top \geq F \top \geq \dots$ the final chain of F . These appear in Theorem 2.2.2.

Theorem 2.2.1 and 2.2.2 yield the following witness notions for proving and disproving $\mu F \leq \alpha$, respectively.

Corollary 2.3. Let (L, \leq) be a complete lattice and $F : L \rightarrow L$ be ω -continuous.

1. (KT) $\mu F \leq \alpha$ if and only if there is $x \in L$ such that $Fx \leq x \leq \alpha$.
2. (Kleene) $\mu F \not\leq \alpha$ if and only if there is $n \in \mathbb{N}$ and $x \in L$ such that $x \leq F^n \perp$ and $x \not\leq \alpha$. \square

By Corollary 2.3.1, proving $\mu F \leq \alpha$ is reduced to searching for $x \in L$ such that $Fx \leq x \leq \alpha$. We call such x a *KT (positive) witness*. In contrast, by Corollary 2.3.2, disproving $\mu F \leq \alpha$ is reduced to searching for $n \in \mathbb{N}$ and $x \in L$ such that $x \leq F^n \perp$ and $x \not\leq \alpha$. We call such x a *Kleene (negative) witness*.

Notation 2.4. We shall use lowercase (Roman and Greek) letters for elements of L (such as $\alpha, x \in L$), and uppercase letters for (finite or infinite) sequences of L (such as $X \in L^*$ or L^ω). The i -th (or $(i - j)$ -th when subscripts are started from j) element of a sequence X is designated by a subscript: $X_i \in L$.

3 Lattice-Theoretic Reconstruction of PDR

Towards the LT-PDR algorithm, we first introduce two simpler algorithms, called positive LT-PDR (Sect. 3.1) and negative LT-PDR (Sect. 3.2). The target problem of the LT-PDR algorithm is the following:

Definition 3.1 (the LFP-OA problem $\mu F \leq^? \alpha$). Let L be a complete lattice, $F : L \rightarrow L$ be ω -continuous, and $\alpha \in L$. The lfp over-approximation (LFP-OA) problem asks if $\mu F \leq \alpha$ holds; the problem is denoted by $\mu F \leq^? \alpha$.

Example 3.2. Consider a transition system, where S be the set of states, $\iota \subseteq S$ be the set of initial states, $\delta : S \rightarrow \mathcal{P}S$ be the transition relation, and $\alpha \subseteq S$ be the set of safe states. Then letting $L := \mathcal{P}S$ and $F := \iota \cup \bigcup_{s \in (-)} \delta(s)$, the lfp over-approximation problem $\mu F \leq^? \alpha$ is the problem whether all reachable states are safe. It is equal to the problem studied by the conventional IC3/PDR [9, 13].

Positive LT-PDR iteratively builds a KT witness in a bottom-up manner that positively answers the LFP-OA problem, while negative LT-PDR iteratively builds a Kleene witness for the same LFP-OA problem. We shall present these two algorithms as clear reflections of two proof principles (Corollary 2.3), each of which comes from the fundamental Knaster–Tarski and Kleene theorems.

The two algorithms build up witnesses in an iterative and nondeterministic manner. The nondeterminism is there for accommodating guesses and heuristics. We identify the essence of PDR to be an ingenious combination of these two algorithms, in which intermediate results on one side (positive or negative) give informed guesses on the other side. This way, each of the positive and negative algorithms provides heuristics in resolving the nondeterminism in the execution of the other. This is how we formulate the LT-PDR algorithm in Sect. 3.3.

The dual of LFP-OA problem is called the *gfp-under-approximation problem* (GFP-UA): the GFP-UA problem for a complete lattice L , an ω^{op} -continuous function $F : L \rightarrow L$ and $\alpha \in L$ is whether the inequality $\alpha \leq \nu F$ holds or not, and is denoted by $\alpha \leq^? \nu F$. It is evident that the GFP-UA problem for (L, F, α) is equivalent to the LFP-OA problem for $(L^{\text{op}}, F, \alpha)$. This suggests the dual algorithm called LT-OpPDR for GFP-UA problem. See Remark 3.24 later.

3.1 Positive LT-PDR: Sequential Positive Witnesses

We introduce the notion of KT^ω witness—a KT witness (Corollary 2.3) constructed in a sequential manner. Positive LT-PDR searches for a KT^ω witness by growing its finitary approximations (called KT sequences).

Let L be a complete lattice. We regard each element $x \in L$ as an abstract presentation of a predicate on states. The inequality $x \leq y$ means that the predicate x is stronger than the predicate y . We introduce the complete lattice $[n, L]$ of increasing chains of length $n \in \mathbb{N}$, whose elements are $(X_0 \leq \dots \leq X_{n-1})$ in L equipped with the element-wise order. We similarly introduce the complete lattice $[\omega, L]$ of ω -chains in L . We lift $F : L \rightarrow L$ to $F^\# : [n, L] \rightarrow [n, L]$ and $F_n^\# : [n, L] \rightarrow [n, L]$ (for $n \geq 2$) as follows. Note that the entries are shifted.

$$\begin{aligned} F^\#(X_0 \leq X_1 \leq \dots) &:= (\perp \leq FX_0 \leq FX_1 \leq \dots) \\ F_n^\#(X_0 \leq \dots \leq X_{n-1}) &:= (\perp \leq FX_0 \leq \dots \leq FX_{n-2}) \end{aligned} \quad (1)$$

Definition 3.3 (KT^ω witness). Let L, F, α be as in Definition 3.1. Define $\Delta\alpha := (\alpha \leq \alpha \leq \dots)$. A KT^ω witness is $X \in [\omega, L]$ such that $F^\#X \leq X \leq \Delta\alpha$.

Theorem 3.4. Let L, F, α be as in Definition 3.1. There exists a KT witness (Corollary 2.3) if and only if there exists a KT^ω witness. \square

Concretely, a KT witness x yields a KT^ω witness $x \leq x \leq \dots$; a KT^ω witness X yields a KT witness $\bigvee_{n \in \omega} X_n$. A full proof (via Galois connections) is in [22].

The initial chain $\perp \leq F\perp \leq \dots$ is always a KT^ω witness for $\mu F \leq \alpha$. There are other KT^ω witnesses whose growth is accelerated by some heuristic guesses—an extreme example is $x \leq x \leq \dots$ with a KT witness x . KT^ω witnesses embrace the spectrum of such different sequential witnesses for $\mu F \leq \alpha$, those which mix routine constructions (i.e. application of F) and heuristic guesses.

Definition 3.5 (KT sequence). Let L, F, α be as in Definition 3.1. A KT sequence for $\mu F \leq^? \alpha$ is a finite chain $(X_0 \leq \dots \leq X_{n-1})$, for $n \geq 2$, satisfying

1. $X_{n-2} \leq \alpha$; and
2. X is a prefixed point of $F_n^\#$, that is, $FX_i \leq X_{i+1}$ for each $i \in [0, n-2]$.

A KT sequence $(X_0 \leq \dots \leq X_{n-1})$ is conclusive if $X_{j+1} \leq X_j$ for some j .

KT sequences are finite by definition. Note that the upper bound α is imposed on all X_i but X_{n-1} . This freedom in the choice of X_{n-1} offers room for heuristics, one that is exploited in the combination with negative LT-PDR (Sect. 3.3).

We take KT sequences as finite approximations of KT^ω witnesses. This view shall be justified by the partial order (\preceq) between KT sequences defined below.

Definition 3.6 (order \preceq between KT sequences). We define a partial order relation \preceq on KT sequences as follows: $(X_0, \dots, X_{n-1}) \preceq (X'_0, \dots, X'_{m-1})$ if $n \leq m$ and $X_j \geq X'_j$ for each $0 \leq j \leq n-1$.

The order $X_j \geq X'_j$ represents that X'_j is a stronger predicate (on states) than X_j . Therefore $X \preceq X'$ expresses that X' is a longer and stronger/more determined chain than X . We obtain KT^ω witnesses as their ω -suprema.

Theorem 3.7. *Let L, F, α be as in Definition 3.1. The set of KT sequences, augmented with the set of KT^ω witnesses $\{X \in [\omega, L] \mid F^\#X \leq X \leq \Delta\alpha\}$ and ordered by the natural extension of \preceq , is an ω -cpo. In this ω -cpo, each KT^ω witness X is represented as the suprema of an ω -chain of KT sequences, namely $X = \bigvee_{n \geq 2} X|_n$ where $X|_n \in [n, L]$ is the length n prefix of X . \square*

Proposition 3.8. *Let L, F, α be as in Definition 3.1. There exists a KT^ω witness if and only if there exists a conclusive KT sequence.*

Proof. (\Rightarrow): If there exists a KT^ω witness, $\mu F \leq \alpha$ holds by Corollary 2.3 and Theorem 3.4. Therefore, the “informed guess” ($\mu F \leq \mu F$) gives a conclusive KT sequence. (\Leftarrow): When X is a conclusive KT sequence with $X_j = X_{j+1}$, $X_0 \leq \dots \leq X_j = X_{j+1} = \dots$ is a KT^ω witness. \square

The proposition above yields the following partial algorithm that aims to answer positively to the LFP-OA problem. It searches for a conclusive KT sequence.

Definition 3.9 (positive LT-PDR). *Let L, F, α be as in Definition 3.1. Positive LT-PDR is the algorithm shown in Algorithm 1, which says ‘True’ to the LFP-OA problem $\mu F \leq^? \alpha$ if successful.*

The rules are designed by the following principles.

Valid is applied when the current X is conclusive.

Unfold extends X with \top . In fact, we can use any element x satisfying $X_{n-1} \leq x$ and $FX_{n-1} \leq x$ in place of \top (by the application of **Induction** with x). The condition $X_{n-1} \leq \alpha$ is checked to ensure that the extended X satisfies the condition in Definition 3.5.1.

Induction strengthens X , replacing the j -th element with its meet with x . The first condition $X_k \not\leq x$ ensures that this rule indeed strengthens X , and the second condition $F(X_{k-1} \wedge x) \leq x$ ensures that the strengthened X satisfies the condition in Definition 3.5.2, that is, $F_n^\#X \leq X$ (see the proof in [22]).

Theorem 3.10. *Let L, F, α be as in Definition 3.1. Then positive LT-PDR is sound, i.e. if it outputs ‘True’ then $\mu F \leq \alpha$ holds.*

*Moreover, assume $\mu F \leq \alpha$ is true. Then positive LT-PDR is weakly terminating (meaning that suitable choices of x when applying **Induction** make the algorithm terminate). \square*

The last “optimistic termination” is realized by the informed guess μF as x in **Induction**. To guarantee the termination of LT-PDR, it suffices to assume that the complete lattice L is well-founded (no infinite decreasing chain exists in L) and there is no strictly increasing ω -chain under α in L , although we cannot hope for this assumption in every instance (Sect. 5.2, 5.3).

Lemma 3.11. *Let L, F, α be as in Definition 3.1. If $\mu F \leq \alpha$, then for any KT sequence X , at least one of the three rules in Algorithm 1 is enabled.*

*Moreover, for any KT sequence X , let X' be obtained by applying either **Unfold** or **Induction**. Then $X \preceq X'$ and $X \neq X'$. \square*

Input : An instance $(\mu F \leq^? \alpha)$ of the LFP-OA problem in L
Output : ‘True’ with a conclusive KT sequence
Data: a KT sequence $X = (X_0 \leq \dots \leq X_{n-1})$
Initially: $X := (\perp \leq F \perp)$
repeat (do one of the following)
 Valid If $X_{j+1} \leq X_j$ for some $j < n - 1$, return ‘True’ with the conclusive KT sequence X .
 Unfold If $X_{n-1} \leq \alpha$, let $X := (X_0 \leq \dots \leq X_{n-1} \leq \top)$. \top
 Induction If some $k \geq 2$ and $x \in L$ satisfy $X_k \not\leq x$ and $F(X_{k-1} \wedge x) \leq x$, let $X := X[X_j := X_j \wedge x]_{2 \leq j \leq k}$.
until any return value is obtained;

Algorithm 1: positive LT-PDR

Input : An instance $(\mu F \leq^? \alpha)$ of the LFP-OA problem in L
Output : ‘False’ with a conclusive Kleene sequence
Data: a Kleene sequence $C = (C_0, \dots, C_{n-1})$
Initially: $C := ()$
repeat (do one of the following)
 Candidate Choose $x \in L$ such that $x \not\leq \alpha$, and let $C := (x)$.
 Model If $C_0 = \perp$, return ‘False’ with the conclusive Kleene sequence C .
 Decide If there exists x such that $C_0 \leq Fx$, then let $C := (x, C_0, \dots, C_{n-1})$.
until any return value is obtained;

Algorithm 2: negative LT-PDR

Input : An instance $(\mu F \leq^? \alpha)$ of the LFP-OA problem in L
Output : ‘True’ with a conclusive KT sequence, or ‘False’ with a conclusive Kleene sequence
Data: $(X; C)$ where X is a KT sequence $(X_0 \leq \dots \leq X_{n-1})$, and C is a Kleene sequence $(C_i, C_{i+1}, \dots, C_{n-1})$ (C is empty if $n = i$).
Initially: $(X; C) := (\perp \leq F \perp; ())$
repeat (do one of the following)
 Valid If $X_{j+1} \leq X_j$ for some $j < n - 1$, return ‘True’ with the conclusive KT sequence X .
 Unfold If $X_{n-1} \leq \alpha$, let $(X; C) := (X_0 \leq \dots \leq X_{n-1} \leq \top; ())$.
 Induction If some $k \geq 2$ and $x \in L$ satisfy $X_k \not\leq x$ and $F(X_{k-1} \wedge x) \leq x$, let $(X; C) := (X[X_j := X_j \wedge x]_{2 \leq j \leq k}; C)$.
 Candidate If $C = ()$ and $X_{n-1} \not\leq \alpha$, choose $x \in L$ such that $x \leq X_{n-1}$ and $x \not\leq \alpha$, and let $(X; C) := (X; (x))$.
 Model If C_1 is defined, return ‘False’ with the conclusive Kleene sequence $(\perp, C_1, \dots, C_{n-1})$.
 Decide If $C_i \leq FX_{i-1}$, choose $x \in L$ satisfying $x \leq X_{i-1}$ and $C_i \leq Fx$, and let $(X; C) := (X; (x, C_i, \dots, C_{n-1}))$.
 Conflict If $C_i \not\leq FX_{i-1}$, choose $x \in L$ satisfying $C_i \not\leq x$ and $F(X_{i-1} \wedge x) \leq x$, and let $(X; C) := (X[X_j := X_j \wedge x]_{2 \leq j \leq i}; (C_{i+1}, \dots, C_{n-1}))$.
until any return value is obtained;

Algorithm 3: LT-PDR

Theorem 3.12. *Let L, F, α be as in Definition 3.1. Assume that \leq in L is well-founded and $\mu F \leq \alpha$. Then, any non-terminating run of positive LT-PDR converges to a KT^ω witness (meaning that it gives a KT^ω witness in ω -steps). Moreover, if there is no strictly increasing ω -chain bounded by α in L , then positive LT-PDR is strongly terminating. \square*

3.2 Negative PDR: Sequential Negative Witnesses

We next introduce *Kleene sequences* as a lattice-theoretic counterpart of *proof obligations* in the standard PDR. Kleene sequences represent a chain of sufficient conditions to conclude that certain unsafe states are reachable.

Definition 3.13 (Kleene sequence). *Let L, F, α be as in Definition 3.1. A Kleene sequence for the LFP-OA problem $\mu F \leq^? \alpha$ is a finite sequence (C_0, \dots, C_{n-1}) , for $n \geq 0$ (C is empty if $n = 0$), satisfying*

1. $C_j \leq FC_{j-1}$ for each $1 \leq j \leq n-1$;
2. $C_{n-1} \not\leq \alpha$.

A Kleene sequence (C_0, \dots, C_{n-1}) is conclusive if $C_0 = \perp$. We may use i ($0 \leq i \leq n$) instead of 0 as the starting index of the Kleene sequence C .

When we have a Kleene sequence $C = (C_0, \dots, C_{n-1})$, the chain of implications $(C_j \leq F^j \perp) \implies (C_{j+1} \leq F^{j+1} \perp)$ hold for $0 \leq j < n-1$. Therefore when C is conclusive, C_{n-1} is a Kleene witness (Corollary 2.3.2).

Proposition 3.14. *Let L, F, α be as in Definition 3.1. There exists a Kleene (negative) witness if and only if there exists a conclusive Kleene sequence.*

Proof. (\implies): If there exists a Kleene witness x such that $x \leq F^n \perp$ and $x \not\leq \alpha$, $(\perp, F\perp, \dots, F^n \perp)$ is a conclusive Kleene sequence. (\impliedby): Assume there exists a conclusive Kleene sequence C . Then C_{n-1} satisfies $C_{n-1} \leq F^{n-1} \perp$ and $C_{n-1} \not\leq \alpha$ because of $C_{n-1} \leq FC_{n-2} \leq \dots \leq F^{n-1} C_0 = F^{n-1} \perp$ and Definition 3.13.2. \square

This proposition suggests the following algorithm to negatively answer to the LFP-OA problem. It searches for a conclusive Kleene sequence. The algorithm updates a Kleene sequence until its first component becomes \perp .

Definition 3.15 (negative LT-PDR). *Let L, F, α be as in Definition 3.1. Negative LT-PDR is the algorithm shown in Algorithm 2, which says ‘False’ to the LFP-OA problem $\mu F \leq^? \alpha$ if successful.*

The rules are designed by the following principles.

Candidate initializes C with only one element x . The element x has to be chosen such that $x \not\leq \alpha$ to ensure Definition 3.13.2.

Model is applied when the current Kleene sequence C is conclusive.

Decide prepends x to C . The condition $C_0 \leq Fx$ ensures Definition 3.13.1.

Theorem 3.16. *Let L, F, α be as in Definition 3.1.*

1. *Negative LT-PDR is sound, i.e. if it outputs ‘False’ then $\mu F \not\leq \alpha$.*
2. *Assume $\mu F \leq \alpha$ is true. Then negative LT-PDR is weakly terminating (meaning that suitable choices of x when applying rules **Candidate** and **Decide** make the algorithm terminate). \square*

3.3 LT-PDR: Integrating Positive and Negative

We have introduced two simple PDR algorithms, called positive LT-PDR (Sect. 3.1) and negative LT-PDR (Sect. 3.2). They are so simple that they have potential inefficiencies. Specifically, in positive LT-PDR, it is unclear that how we choose $x \in L$ in **Induction**, while in negative LT-PDR, it may easily diverge because the rules **Candidate** and **Decide** may choose $x \in L$ that would not lead to a conclusive Kleene sequence. We resolve these inefficiencies by combining positive LT-PDR and negative LT-PDR. The combined PDR algorithm is called LT-PDR, and it is a lattice-theoretic generalization of conventional PDR.

Note that negative LT-PDR is only weakly terminating. Even worse, it is easy to make it diverge—after a choice of x in **Candidate** or **Decide** such that $x \not\leq \mu F$, no continued execution of the algorithm can lead to a conclusive Kleene sequence. For deciding $\mu F \leq^? \alpha$ efficiently, therefore, it is crucial to detect such useless Kleene sequences.

The core fact that underlies the efficiency of PDR is the following proposition, which says that a KT sequence (in positive LT-PDR) can quickly tell that a Kleene sequence (in negative LT-PDR) is useless. This fact is crucially used for many rules in LT-PDR (Definition 3.20).

Proposition 3.17. *Let $C = (C_i, \dots, C_{n-1})$ be a Kleene sequence ($2 \leq n, 0 < i \leq n-1$) and $X = (X_0 \leq \dots \leq X_{n-1})$ be a KT sequence. Then*

1. $C_i \not\leq X_i$ implies that C cannot be extended to a conclusive one, that is, there does not exist C_0, \dots, C_{i-1} such that (C_0, \dots, C_{n-1}) is conclusive.
2. $C_i \not\leq FX_{i-1}$ implies that C cannot be extended to a conclusive one.
3. There is no conclusive Kleene sequence with length $n-1$. □

The proof relies on the following lemmas.

Lemma 3.18. *Any KT sequence $(X_0 \leq \dots \leq X_{n-1})$ over-approximates the initial sequence: $F^i \perp \leq X_i$ holds for any i such that $0 \leq i \leq n-1$. □*

Lemma 3.19. *Let $C = (C_i, \dots, C_{n-1})$ be a Kleene sequence ($0 < i \leq n-1$) and $(X_0 \leq \dots \leq X_{n-1})$ be a KT sequence. The following satisfy $1 \Leftrightarrow 2 \Rightarrow 3$.*

1. The Kleene sequence C can be extended to a conclusive one.
2. $C_i \leq F^i \perp$.
3. $C_i \leq F^j X_{i-j}$ for each j with $0 \leq j \leq i$. □

Using the above lattice-theoretic properties, we combine positive and negative LT-PDRs into the following *LT-PDR* algorithm. It is also a lattice-theoretic generalization of the original PDR algorithm. The combination exploits the mutual relationship between KT sequences and Kleene sequences, exhibited as Proposition 3.17, for narrowing down choices in positive and negative LT-PDRs.

Definition 3.20 (LT-PDR). *Given a complete lattice L , an ω -continuous function $F : L \rightarrow L$, and an element $\alpha \in L$, LT-PDR is the algorithm shown in Algorithm 3 for the LFP-OA problem $\mu F \leq^? \alpha$.*

The rules are designed by the following principles.

(**Valid**, **Unfold**, and **Induction**): These rules are almost the same as in positive LT-PDR. In **Unfold**, we reset the Kleene sequence because of Proposition 3.17.3. Occurrences of **Unfold** punctuate an execution of the algorithm: between two occurrences of **Unfold**, a main goal (towards a negative conclusion) is to construct a conclusive Kleene sequence with the same length as the X .

(**Candidate**, **Model**, and **Decide**): These rules have many similarities to those in negative LT-PDR. Differences are as follows: the **Candidate** and **Decide** rules impose $x \leq X_i$ on the new element x in $(x, C_{i+1}, \dots, C_{n-1})$ because Proposition 3.17.1 tells us that other choices are useless. In **Model**, we only need to check whether C_1 is defined instead of $C_0 = \perp$. Indeed, since C_1 is added in **Candidate** or **Decide**, $C_1 \leq X_1 = F\perp$ always holds. Therefore, $2 \Rightarrow 1$ in Lemma 3.19 shows that $(\perp, C_1, \dots, C_{n-1})$ is conclusive.

(**Conflict**): This new rule emerges from the combination of positive and negative LT-PDRs. This rule is applied when $C_i \not\leq FX_{i-1}$, which confirms that the current C cannot be extended to a conclusive one (Proposition 3.17.2). Therefore, we eliminate C_i from C and strengthen X so that we cannot choose C_i again, that is, so that $C_i \not\leq (X_i \wedge x)$. Let us explain how X is strengthened. The element x has to be chosen so that $C_i \not\leq x$ and $F(X_{i-1} \wedge x) \leq x$. The former dis-inequality ensures the strengthened X satisfies $C_i \not\leq (X_i \wedge x)$, and the latter inequality implies $F(X_{i-1} \wedge x) \leq x$. One can see that **Conflict** is **Induction** with additional condition $C_i \not\leq x$, which enhances so that the search space for x is narrowed down using the Kleene sequence C .

Canonical choices of $x \in L$ in **Candidate**, **Decide**, and **Conflict** are $x := X_{n-1}$, $x := X_{i-1}$, and $x := FX_{i-1}$, respectively. However, there can be cleverer choices; e.g. $x := S \setminus (C_i \setminus FX_{i-1})$ in **Conflict** when $L = \mathcal{PS}$.

Lemma 3.21. *Each rule of LT-PDR, when applied to a pair of a KT and a Kleene sequence, yields a pair of a KT and a Kleene sequence.* \square

Theorem 3.22 (correctness). *LT-PDR is sound, i.e. if it outputs ‘True’ then $\mu F \leq \alpha$ holds, and if it outputs ‘False’ then $\mu F \not\leq \alpha$ holds.* \square

Many existing PDR algorithms ensure termination if the state space is finite. A general principle behind is stated below. Note that it rarely applies to infinitary or quantitative settings, where we would need some abstraction for termination.

Proposition 3.23 (termination). *LT-PDR terminates regardless of the order of the rule-applications if the following conditions are satisfied.*

1. **Valid** and **Model** rules are immediately applied if applicable.
2. (L, \leq) is well-founded.
3. Either of the following is satisfied: a) $\mu F \leq \alpha$ and (L, \leq) has no strictly increasing ω -chain bounded by α , or b) $\mu F \not\leq \alpha$. \square

Cond 1 is natural: it just requires LT-PDR to immediately conclude ‘True’ or ‘False’ if it can. Cond. 2–3 are always satisfied when L is finite.

Table 1. Categorical modeling of state-based dynamics and predicate transformers

a transition system as a <i>coalgebra</i> [17] in the base category \mathbb{B} of sets and functions	
objects X, Y, \dots in \mathbb{B}	sets (in our examples where $\mathbb{B} = \mathbf{Set}$)
an arrow $f: X \rightarrow Y$ in \mathbb{B}	a function (in our examples where $\mathbb{B} = \mathbf{Set}$)
a functor $G: \mathbb{B} \rightarrow \mathbb{B}$	a transition type $\left(\begin{array}{l} G = \mathcal{P} \text{ for Kripke structures (§5.1),} \\ G = (\mathcal{D}(-) + 1)^{\text{Act}} \text{ for MDPs (§5.2), etc.} \end{array} \right)$
a <i>coalgebra</i> $\delta: S \rightarrow GS$ in \mathbb{B} [17]	a transition system (Kripke structure, MDP, etc.)
a <i>fibration</i> $p: \mathbb{E} \rightarrow \mathbb{B}$ [18] that equips sets in \mathbb{B} with <i>predicates</i>	
the fiber category \mathbb{E}_S over S in \mathbb{B}	the lattice of predicates over a set S
the <i>pullback</i> functor $l^*: \mathbb{E}_Y \rightarrow \mathbb{E}_X$ for $l: X \rightarrow Y$ in \mathbb{B}	substitution $P(y) \mapsto P(l(x))$ in predicates $P \in \mathbb{E}_Y$ over Y
a <i>lifting</i> $\dot{G}: \mathbb{E} \rightarrow \mathbb{E}$ of G along p	logical interpretation of the transition type G (specifies e.g. the may vs. must modalities)
the <i>predicate transformer</i> , whose fixed points are of our interest	
the composite $\delta^* \dot{G}: \mathbb{E}_S \rightarrow \mathbb{E}_S$	the predicate transformer associated with the transition system δ

Theorem 3.22 and Proposition 3.23 still hold if **Induction** rule is dropped. However, the rule can accelerate the convergence of KT sequences and improve efficiency.

Remark 3.24 (LT-OpPDR). The GFP-UA problem $\alpha \leq^? \nu F$ is the dual of LFP-OA, obtained by opposing the order \leq in L . We can also dualize the LT-PDR algorithm (Algorithm 3), obtaining what we call the *LT-OpPDR* algorithm for GFP-UA. Moreover, we can express LT-OpPDR as LT-PDR if a suitable *involution* $\neg: L \rightarrow L^{\text{op}}$ is present. See [22, Appendix B] for further details; see also Proposition 4.3.

4 Structural Theory of PDR by Category Theory

Before we discuss concrete instances of LT-PDR in Sect. 5, we develop a structural theory of transition systems and predicate transformers as a basis of LT-PDR. The theory is formulated in the language of *category theory* [3, 17, 18, 23]. We use category theory because 1) categorical modeling of relevant notions is well established in the community (see e.g. [2, 8, 17, 18, 27]), and 2) it gives us the right level of abstraction that accommodates a variety of instances. In particular, qualitative and quantitative settings are described in a uniform manner.

Our structural theory (Sect. 4) serves as a backend, not a frontend. That is,

- the theory in Sect. 4 is important in that it explains how the instances in Sect. 5 arise and why others do not, but
- the instances in Sect. 5 are described in non-categorical terms, so readers who skipped Sect. 4 will have no difficulties following Sect. 5 and using those instances.

4.1 Categorical Modeling of Dynamics and Predicate Transformers

Our interests are in instances of the LFP-OA problem $\mu F \leq^? \alpha$ (Definition 3.1) that appear in *model checking*. In this context, 1) the underlying lattice L is that of *predicates* over a state space, and 2) the function $F: L \rightarrow L$ arises from the dynamic/transition structure, specifically as a *predicate transformer*. The categorical notions in Table 1 model these ideas (state-based dynamics, predicate transformers). This modeling is well-established in the community.

Our introduction of Table 1 here is minimal, due to the limited space. See [22, Appendix C] and the references therein for more details.

A *category* consists of *objects* and *arrows* between them. In Table 1, categories occur twice: 1) a *base category* \mathbb{B} where objects are typically sets and arrows are typically functions; and 2) *fiber categories* \mathbb{E}_S , defined for each object S of \mathbb{B} , that are identified with the lattices of *predicates*. Specifically, objects P, Q, \dots of \mathbb{E}_S are predicates over S , and an arrow $P \rightarrow Q$ represents logical implication. A general fact behind the last is that every preorder is a category—see e.g. [3].

Transition Systems as Coalgebras. State-based transition systems are modeled as *coalgebras* in the base category \mathbb{B} [17]. We use a *functor* $G: \mathbb{B} \rightarrow \mathbb{B}$ to represent a transition type. A *G-coalgebra* is an arrow $\delta: S \rightarrow GS$, where S is a state space and δ describes the dynamics. For example, a Kripke structure can be identified with a pair (S, δ) of a set S and a function $\delta: S \rightarrow \mathcal{P}S$, where $\mathcal{P}S$ denotes the powerset. The powerset construction \mathcal{P} is known to be a functor $\mathcal{P}: \mathbf{Set} \rightarrow \mathbf{Set}$; therefore Kripke structures are \mathcal{P} -coalgebras. For other choices of G , *G-coalgebras* become different types of transition systems, such as MDPs (Sect. 5.2) and Markov Reward Models (Sect. 5.3).

Predicates Form a Fibration. Fibrations are powerful categorical constructs that can model various indexed entities; see e.g. [18] for its general theory. Our use of them is for organizing the lattices \mathbb{E}_S of *predicates* over a set S , indexed by the choice of S . For example, $\mathbb{E}_S = 2^S$ —the lattice of subsets of S —for modeling qualitative predicates. For quantitative reasoning (e.g. for MDPs), we use $\mathbb{E}_S = [0, 1]^S$, where $[0, 1]$ is the unit interval. This way, qualitative and quantitative reasonings are mathematically unified in the language of fibrations.

A *fibration* is a functor $p: \mathbb{E} \rightarrow \mathbb{B}$ with suitable properties; it can be thought of as a collection $(\mathbb{E}_S)_{S \in \mathbb{B}}$ of *fiber categories* \mathbb{E}_S —indexed by objects S of \mathbb{B} —suitably organized as a single category \mathbb{E} . Notable in this organization is that we obtain the *pullback* functor $l^*: \mathbb{E}_Y \rightarrow \mathbb{E}_X$ for each arrow $l: X \rightarrow Y$ in \mathbb{B} . In our examples, l^* is a *substitution* along l in predicates— l^* is the monotone map that carries a predicate $P(y)$ over Y to the predicate $P(l(x))$ over X .

In this paper, we restrict to a subclass of fibrations (called **CLat_∧-fibrations**) in which every fiber category \mathbb{E}_S is a complete lattice, and each pullback functor preserves all meets. We therefore write $P \leq Q$ for arrows in \mathbb{E}_S ; this represents logical implication, as announced above. Notice that each f^* has a left adjoint (lower adjoint in terms of Galois connection), which exists by Freyd’s adjoint functor theorem. The left adjoint is denoted by f_* .

We also consider a *lifting* $\dot{G}: \mathbb{E} \rightarrow \mathbb{E}$ of G along p ; it is a functor \dot{G} such that $p\dot{G} = Gp$. See the diagram on the right. It specifies the *logical interpretation* of the transition type G . For example, for $G = \mathcal{P}$ (the powerset functor) from the above, two choices of \dot{G} are for the *may* and *must* modalities. See e.g. [2, 15, 20, 21].

$$\begin{array}{ccc} \mathbb{E} & \xrightarrow{\dot{G}} & \mathbb{E} \\ p \downarrow & & p \downarrow \\ \mathbb{B} & \xrightarrow{G} & \mathbb{B} \end{array}$$

Categorical Predicate Transformer. The above constructs allow us to model predicate transformers— F in our examples of the LFP-OA problem $\mu F \leq^? \alpha$ —in categorical terms. A *predicate transformer* along a coalgebra $\delta: S \rightarrow GS$ with respect to the lifting \dot{G} is simply the composite $\mathbb{E}_S \xrightarrow{\dot{G}} \mathbb{E}_{GS} \xrightarrow{\delta^*} \mathbb{E}_S$, where the first \dot{G} is the restriction of $\dot{G}: \mathbb{E} \rightarrow \mathbb{E}$ to \mathbb{E}_S . Intuitively, 1) given a *postcondition* P in \mathbb{E}_S , 2) it is first interpreted as the predicate $\dot{G}P$ over GS , and then 3) it is pulled back along the dynamics δ to yield a *precondition* $\delta^*\dot{G}P$. Such (backward) predicate transformers are fundamental in a variety of model checking problems.

4.2 Structural Theory of PDR from Transition Systems

We formulate a few general *safety* problems. We show how they are amenable to the LT-PDR (Definition 3.20) and LT-OpPDR (Remark 3.24) algorithms.

Definition 4.1 (backward safety problem, BSP). Let p be a \mathbf{CLat}_\wedge -fibration, $\delta: S \rightarrow GS$ be a coalgebra in \mathbb{B} , and $\dot{G}: \mathbb{E} \rightarrow \mathbb{E}$ be a lifting of G along p such that $\dot{G}_X: \mathbb{E}_X \rightarrow \mathbb{E}_{GX}$ is ω^{op} -continuous for each $X \in \mathbb{B}$. The backward safety problem for $(\iota \in \mathbb{E}_S, \delta, \alpha \in \mathbb{E}_S)$ in (p, G, \dot{G}) is the GFP-UA problem for $(\mathbb{E}_S, \alpha \wedge \delta^*\dot{G}\iota)$, that is,

$$\iota \leq^? \nu x. \alpha \wedge \delta^*\dot{G}x. \quad (2)$$

Here, ι represents the initial states and α represents the safe states. The predicate transformer $x \mapsto \alpha \wedge \delta^*\dot{G}x$ in (2) is the standard one for modeling safety—currently safe (α), and the next time x ($\delta^*\dot{G}x$). Its gfp is the safety property; (2) asks if all initial states (ι) satisfy the safety property. Since the backward safety problem is a GFP-UA problem, we can solve it by LT-OpPDR (Remark 3.24).

Additional assumptions allow us to reduce the backward safety problem to LFP-OA problems, which are solvable by LT-PDR, as shown on the right.

$$\begin{array}{ccc} \text{BSP} & \xrightarrow{\text{as-is}} & \text{GFP-UA} \xrightarrow{\text{LT-OpPDR}} \text{True/False} \\ & \searrow \text{suitable adjoints} & \\ & \text{involution } \neg & \searrow \\ & & \text{LFP-OA} \xrightarrow{\text{LT-PDR}} \text{True/False} \end{array}$$

The first case requires the existence of the *left adjoint* to the predicate transformer $\delta^*\dot{G}_S: \mathbb{E}_S \rightarrow \mathbb{E}_S$. Then we can translate BSP to the following LFP-OA problem. It directly asks whether all reachable states are safe.

Proposition 4.2 (forward safety problem, FSP). In the setting of Definition 4.1, assume that each $\dot{G}_X: \mathbb{E}_X \rightarrow \mathbb{E}_{GX}$ preserves all meets. Then by

letting $\dot{H}_S : \mathbb{E}_{GS} \rightarrow \mathbb{E}_S$ be the left adjoint of \dot{G}_S , the BSP (2) is equivalent to the LFP-OA problem for $(\mathbb{E}_S, \iota \vee \dot{H}_S \delta_*, \alpha)$:

$$\mu x. \iota \vee \dot{H}_S \delta_* x \leq^? \alpha. \quad (3)$$

This problem is called the forward safety problem for (ι, δ, α) in (p, G, \dot{G}) . \square

The second case assumes that the complete lattice \mathbb{E}_S of predicates admits an involution operator $\neg : \mathbb{E}_S \rightarrow \mathbb{E}_S^{\text{op}}$ (cf. [22, Appendix B]).

Proposition 4.3 (inverse backward safety problem, IBSP). *In the setting of Definition 4.1, assume further that there is a monotone function $\neg : \mathbb{E}_S \rightarrow \mathbb{E}_S^{\text{op}}$ satisfying $\neg \circ \neg = \text{id}$. Then the backward safety problem (2) is equivalent to the LFP-OA problem for $(\mathbb{E}_S, (\neg\alpha) \vee (\neg \circ \delta^* \dot{G} \circ \neg), \neg\iota)$, that is,*

$$\mu x. (\neg\alpha) \vee (\neg \circ \delta^* \dot{G} \circ \neg x) \leq^? \neg\iota. \quad (4)$$

We call (4) the inverse backward safety problem for (ι, δ, α) in (p, G, \dot{G}) . Here $(\neg\alpha) \vee (\neg \circ \delta^* \dot{G} \circ \neg(-))$ is the inverse backward predicate transformer. \square

When both additional assumptions are fulfilled (in Proposition 4.2 and 4.3), we obtain two LT-PDR algorithms to solve BSP. One can even simultaneously run these two algorithms—this is done in fbPDR [25, 26]. See also Sect. 5.1.

5 Known and New PDR Algorithms as Instances

We present several concrete instances of our LT-PDR algorithms. The one for Markov reward models is new (Sect. 5.3). We also sketch how those instances can be systematically derived by the theory in Sect. 4; details are in [22, Appendix D].

5.1 LT-PDRs for Kripke Structures: $\text{PDR}^{\text{F-Kr}}$ and $\text{PDR}^{\text{IB-Kr}}$

In most of the PDR literature, the target system is a Kripke structure that arises from a program's operational semantics. A *Kripke structure* consists of a set S of states and a transition relation $\delta \subseteq S \times S$ (here we ignore initial states and atomic propositions). The basic problem formulation is as follows.

Definition 5.1 (backward safety problem (BSP) for Kripke structures). *The BSP for a Kripke structure (S, δ) , a set $\iota \in 2^S$ of initial states, and a set $\alpha \in 2^S$ of safe states, is the GFP-UA problem $\iota \leq^? \nu x. \alpha \wedge F'x$, where $F' : 2^S \rightarrow 2^S$ is defined by $F'(A) := \{s \mid \forall s'. ((s, s') \in \delta \Rightarrow s' \in A)\}$.*

It is clear that the GFP in Definition 5.1 represents the set of states from which all reachable states are in α . Therefore the BSP is the usual safety problem.

The above BSP is easily seen to be equivalent to the following problems.

Proposition 5.2 (forward safety problem (FSP) for Kripke structures). *The BSP in Definition 5.1 is equivalent to the LFP-OA problem $\mu x. \iota \vee F''x \leq^? \alpha$, where $F'' : 2^S \rightarrow 2^S$ is defined by $F''(A) := \bigcup_{s \in A} \{s' \mid (s, s') \in \delta\}$. \square*

Proposition 5.3 (inverse backward safety problem (IBSP) for Kripke structures). *The BSP in Definition 5.1 is equivalent to the LFP-OA problem $\mu x. \neg \alpha \vee \neg F'(\neg x) \leq^? \neg \iota$, where $\neg: 2^S \rightarrow 2^S$ is the complement function $A \mapsto S \setminus A$.* \square

Instances of LT-PDR. The FSP and IBSP (Propositions 5.2–5.3), being LFP-OA, are amenable to the LT-PDR algorithm (Definition 3.20). Thus we obtain two instances of LT-PDR; we call them $\mathbf{PDR}^{F\text{-Kr}}$ and $\mathbf{PDR}^{IB\text{-Kr}}$. $\mathbf{PDR}^{IB\text{-Kr}}$ is a step-by-step dual to the application of LT-OpPDR to the BSP (Definition 5.1)—see Remark 3.24.

We compare these two instances of LT-PDR with algorithms in the literature. If we impose $|C_i| = 1$ on each element C_i of Kleene sequences, the $\mathbf{PDR}^{F\text{-Kr}}$ instance of LT-PDR coincides with the conventional IC3/PDR [9, 13]. In contrast, $\mathbf{PDR}^{IB\text{-Kr}}$ coincides with *Reverse PDR* in [25, 26]. The parallel execution of $\mathbf{PDR}^{F\text{-Kr}}$ and $\mathbf{PDR}^{IB\text{-Kr}}$ roughly corresponds to fbPDR [25, 26].

Structural Derivation. The equivalent problems (Propositions 5.2–5.3) are derived systematically from the categorical theory in Sect. 4.2. Indeed, using a lifting $\hat{\mathcal{P}}: 2^S \rightarrow 2^{2^S}$ such that $A \mapsto \{A' \mid A' \subseteq A\}$ (the *must modality* \square), F' in Definition 5.1 coincides with $\delta^* \hat{\mathcal{P}}$ in (2). The above $\hat{\mathcal{P}}$ preserves meets (cf. the modal axiom $\square(\varphi \wedge \psi) \cong \square\varphi \wedge \square\psi$, see e.g. [7]); thus Proposition 4.2 derives the FSP. Finally, \neg in Proposition 5.3 allows the use of Proposition 4.3. More details are in [22, Appendix D].

5.2 LT-PDR for MDPs: $\mathbf{PDR}^{IB\text{-MDP}}$

The only known PDR-like algorithm for *quantitative* verification is *PrIC3* [6] for Markov decision processes s(MDPs). Here we instantiate LT-PDR for MDPs and compare it with PrIC3.

An *MDP* consists of a set S of states, a set Act of actions and a transition function δ mapping $s \in S$ and $a \in \text{Act}$ to either $*$ (“the action a is unavailable at s ”) or a probability distribution $\delta(s)(a)$ over S .

Definition 5.4 (IBSP for MDPs). *The inverse backward safety problem (IBSP) for an MDP (S, δ) , an initial state $s_\iota \in S$, a real number $\lambda \in [0, 1]$, and a set $\alpha \subseteq S$ of safe states, is the LFP-OA problem $\mu x. F'(x) \leq^? d_{\iota, \lambda}$. Here $d_{\iota, \lambda}: S \rightarrow [0, 1]$ is the predicate such that $d_{\iota, \lambda}(s_\iota) = \lambda$ and $d_{\iota, \lambda}(s) = 1$ otherwise. $F': [0, 1]^S \rightarrow [0, 1]^S$ is defined by $F'(d)(s) = 1$ if $s \notin \alpha$, and $F'(d)(s) = \max\{\sum_{s' \in S} d(s') \cdot \delta(s)(a)(s') \mid a \in \text{Act}, \delta(s)(a) \neq *\}$ if $s \in \alpha$.*

The function F' in Definition 5.4 is a *Bellman operator* for MDPs—it takes the average of d over $\delta(s)(a)$ and takes the maximum over a . Therefore the lfp in Definition 5.4 is the maximum reachability probability to $S \setminus \alpha$; the problem asks if it is $\leq \lambda$. In other words, it asks whether the *safety* probability—of staying in α henceforth, under any choices of actions—is $\geq 1 - \lambda$. This problem is the same as in [6].

Instance of PDR. The IBSP (Definition 5.4) is LFP-OA and thus amenable to LT-PDR. We call this instance $\mathbf{PDR}^{\text{IB-MDP}}$; See [22, Appendix E] for details.

$\mathbf{PDR}^{\text{IB-MDP}}$ shares many essences with PrIC3 [6]. It uses the operator F' in Definition 5.4, which coincides with the one in [6, Def. 2]. PrIC3 maintains frames; they coincide with KT sequences in $\mathbf{PDR}^{\text{IB-MDP}}$.

Our Kleene sequences correspond to *obligations* in PrIC3, modulo the following difference. Kleene sequences aim at a negative witness (Sect. 3.2), but they happen to help the positive proof efforts too (Sect. 3.3); obligations in PrIC3 are solely for accelerating the positive proof efforts. Thus, if PrIC3 cannot solve these efforts, we need to check whether obligations yield a negative witness.

Structural Derivation. One can derive the IBSP (Definition 5.4) from the categorical theory in Sect. 4.2. Specifically, we first formulate the $BSP \neg d_\lambda \leq^? \nu x. d_\alpha \wedge \delta^* \dot{G}x$, where \dot{G} is a suitable lifting (of G for MDPs, Table 1) that combines average and minimum, $\neg: [0, 1]^S \rightarrow [0, 1]^S$ is defined by $(\neg d)(s) := 1 - d(s)$, and d_α is such that $d_\alpha(s) = 1$ if $s \in \alpha$ and $d_\alpha(s) = 0$ otherwise. Using $\neg: [0, 1]^S \rightarrow [0, 1]^S$ in the above as an involution, we apply Proposition 4.3 and obtain the IBSP (Definition 5.4).

Another benefit of the categorical theory is that it can tell us a forward instance of LT-PDR (much like $\mathbf{PDR}^{\text{F-Kr}}$ in Sect. 5.1) is unlikely for MDPs. Indeed, we showed in Proposition 4.2 that \dot{G}' 's preservation of meets is essential (existence of a left adjoint is equivalent to meet preservation). We can easily show that our \dot{G} for MDPs does not preserve meets. See [22, Appendix G].

5.3 LT-PDR for Markov Reward Models: $\mathbf{PDR}^{\text{MRM}}$

We present a PDR-like algorithm for *Markov reward models (MRMs)*, which seems to be new, as an instance of LT-PDR. An MRM consists of a set S of states and a transition function δ that maps $s \in S$ (the current state) and $c \in \mathbb{N}$ (the reward) to a function $\delta(s)(c) : S \rightarrow [0, 1]$; the last represents the probability distribution of next states.

We solve the following problem. We use $[0, \infty]$ -valued predicates—representing accumulated rewards—where $[0, \infty]$ is the set of extended nonnegative reals.

Definition 5.5 (SP for MRMs). *The safety problem (SP) for an MRM (S, δ) , an initial state $s_i \in S$, $\lambda \in [0, \infty]$, and a set $\alpha \subseteq S$ of safe states is $\mu x. F'(x) \leq^? d_{i,\lambda}$. Here $d_{i,\lambda} : S \rightarrow [0, \infty]$ maps s_i to λ and others to ∞ , and $F' : [0, \infty]^S \rightarrow [0, \infty]^S$ is defined by $F'(d)(s) = 0$ if $s \notin \alpha$, and $F'(d)(s) = \sum_{s' \in S, c \in \mathbb{N}} (c + d(s')) \cdot \delta(s)(c)(s')$ if $s \in \alpha$.*

The function F' accumulates expected reward in α . Thus the problem asks if the expected accumulated reward, starting from s_i and until leaving α , is $\leq \lambda$.

Instance of PDR. The SP (Definition 5.5) is LFP-OA thus amenable to LT-PDR. We call this instance $\mathbf{PDR}^{\text{MRM}}$. It seems new. See [22, Appendix F] for details.

Structural Derivation. The function F' in Definition 5.5 can be expressed categorically as $F'(x) = d_\alpha \wedge \delta^* \dot{G}(x)$, where $d_\alpha: S \rightarrow [0, \infty]$ carries $s \in \alpha$ to ∞ and $s \notin \alpha$ to 0, and \dot{G} is a suitable lifting that accumulates expected reward. However, the SP (Definition 5.5) is *not* an instance of the three general safety problems in Sect. 4.2. Consequently, we expect that other instances of LT-PDR than $\mathbf{PDR}^{\text{MRM}}$ (such as $\mathbf{PDR}^{\text{F-Kr}}$ and $\mathbf{PDR}^{\text{IB-Kr}}$ in Sect. 5.1) are hard for MRMs.

6 Implementation and Evaluation

Implementation. **LTPDR** We implemented LT-PDR in Haskell. Exploiting Haskell’s language features, it is succinct (~ 50 lines) and almost a literal translation of Algorithm 3 to Haskell. Its main part is presented in [22, Appendix K]. In particular, using suitable type classes, the code is as abstract and generic as Algorithm 3.

Specifically, our implementation is a Haskell module named **LTPDR**. It has two interfaces, namely the type class **CLat** τ (the lattice of predicates) and the type **Heuristics** τ (the definitions of **Candidate**, **Decide**, and **Conflict**). The main function for LT-PDR is $\text{ltPDR} :: \mathbf{CLat} \tau \Rightarrow \mathbf{Heuristics} \tau \rightarrow (\tau \rightarrow \tau) \rightarrow \tau \rightarrow \mathbf{IO} (\mathbf{PDRAnswer} \tau)$, where the second argument is for a monotone function F of type $\tau \rightarrow \tau$ and the last is for the safety predicate α .

Obtaining concrete instances is easy by fixing τ and **Heuristics** τ . A simple implementation of $\mathbf{PDR}^{\text{F-Kr}}$ takes 15 lines; a more serious SAT-based one for $\mathbf{PDR}^{\text{F-Kr}}$ takes ~ 130 lines; $\mathbf{PDR}^{\text{IB-MDP}}$ and $\mathbf{PDR}^{\text{MRM}}$ take ~ 80 lines each.

Heuristics. We briefly discuss the heuristics, i.e. how to choose $x \in L$ in **Candidate**, **Decide**, and **Conflict**, used in our experiments. The heuristics of $\mathbf{PDR}^{\text{F-Kr}}$ is based on the conventional PDR [9]. The heuristics of $\mathbf{PDR}^{\text{IB-MDP}}$ is based on the idea of representing the smallest possible x greater than some real number $v \in [0, 1]$ (e.g. x taken in **Candidate**) as $x = v + \epsilon$, where ϵ is a symbolic variable. This implies that **Unfold** (or **Valid**, **Model**) is always applied in finite steps, which further guarantees finite-step termination for invalid cases and ω -step termination for valid cases (see [22, Appendix H] for more detail). The heuristics of $\mathbf{PDR}^{\text{MRM}}$ is similar to that of $\mathbf{PDR}^{\text{IB-MDP}}$.

Experiment Setting. We experimentally assessed the performance of instances of LTPDR. The settings are as follows: 1.2 GHz Quad-Core Intel Core i7 with 10 GB memory using Docker, for $\mathbf{PDR}^{\text{IB-MDP}}$; Apple M1 Chip with 16 GB memory for the other. The different setting is because we needed Docker to run PrIC3 [6].

Experiments with $\mathbf{PDR}^{\text{MRM}}$. Table 2a shows the results. We observe that $\mathbf{PDR}^{\text{MRM}}$ answered correctly, and that the execution time is reasonable. Further performance analysis (e.g. comparison with [19]) and improvement is future work; the point here, nevertheless, is the fact that we obtained a reasonable MRM model checker by adding ~ 80 lines to the generic solver LTPDR.

Experiments with $\text{PDR}^{\text{IB-MDP}}$. Table 2c shows the results. Both PrIC3 and our $\text{PDR}^{\text{IB-MDP}}$ solve a linear programming (LP) problem in **Decide**. PrIC3 uses Z3 for this; $\text{PDR}^{\text{IB-MDP}}$ uses GLPK. PrIC3 represents an MDP symbolically, while $\text{PDR}^{\text{IB-MDP}}$ do so concretely. Symbolic representation in $\text{PDR}^{\text{IB-MDP}}$ is possible—it is future work. PrIC3 can use four different *interpolation generalization* methods, leading to different performance (Table 2c).

We observe that $\text{PDR}^{\text{IB-MDP}}$ outperforms PrIC3 for some benchmarks with smaller state spaces. We believe that the failure of $\text{PDR}^{\text{IB-MDP}}$ in many instances can be attributed to our current choice of a generalization method (it is the closest to the linear one for PrIC3). Table 2c suggests that use of *polynomial* or *hybrid* can enhance the performance.

Experiments with $\text{PDR}^{\text{F-Kr}}$. Table 2b shows the results. The benchmarks are mostly from the HWMCC’15 competition [1], except for `latch0.smv`¹ and `counter.smv` (our own).

IC3ref vastly outperforms $\text{PDR}^{\text{F-Kr}}$ in many instances. This is hardly a surprise—IC3ref was developed towards superior performance, while $\text{PDR}^{\text{F-Kr}}$ ’s emphasis is on its theoretical simplicity and genericity. We nevertheless see that $\text{PDR}^{\text{F-Kr}}$ solves some benchmarks of substantial size, such as `power2bit8.smv`. This demonstrates the practical potential of LT-PDR, especially in view of the following improvement opportunities (we will pursue them as future work): 1) use of well-developed SAT solvers (we currently use `toysolver`² for its good interface but we could use Z3); 2) allowing $|C_i| > 1$, a technique discussed in Sect. 5.1 and implemented in IC3ref but not in $\text{PDR}^{\text{F-Kr}}$; and 3) other small improvements, e.g. in our CNF-based handling of propositional formulas.

Ablation Study. To assess the value of the key concept of PDR (namely the *positive-negative interplay* between the Knaster–Tarski and Kleene theorems (Sect. 3.3)), we compared $\text{PDR}^{\text{F-Kr}}$ with the instances of positive and negative LT-PDR (Sects. 3.1–3.2) for Kripke structures.

Table 2d shows the results. Note that the value of the positive-negative interplay is already theoretically established; see e.g. Proposition 3.17 (the interplay detects executions that lead to nowhere). This value was also experimentally witnessed: see `power2bit8.smv` and `simpleTrans.smv`, where the one-sided methods made wrong choices and timed out. One-sided methods can be efficient when they get lucky (e.g. in `counter.smv`). LT-PDR may be slower because of the overhead of running two sides, but that is a trade-off for the increased chance of termination.

Discussion. We observe that all of the studied instances exhibited at least reasonable performance. We note again that detailed performance analysis and improvement is out of our current scope. Being able to derive these model checkers, with such a small effort as ~ 100 lines of Haskell code each, demonstrates the value of our abstract theory and its generic Haskell implementation LTPDR.

¹ <https://github.com/arminbiere/aiger>.

² <https://github.com/msakai/toysolver>.

Table 2. experimental results for our $\mathbf{PDR}^{\mathbf{F-Kr}}$, $\mathbf{PDR}^{\mathbf{IB-MDP}}$, and $\mathbf{PDR}^{\mathbf{MRM}}$

(a) Results with $\mathbf{PDR}^{\mathbf{MRM}}$. The MRM is from [4, Example 10.72], whose ground truth expected reward is $\frac{4}{3}$. The benchmarks ask if the expected reward (not known to the solver) is ≤ 1.5 or ≤ 1.3 .

Benchmark	Result	Time
DIEByCOIN $\leq^{1.5}$	True	6.01 ms
DIEByCOIN $\leq^{1.3}$	False	43.1 μ s

(b) Results with $\mathbf{PDR}^{\mathbf{F-Kr}}$ in comparison with IC3ref, a reference implementation of [9] (<https://github.com/arbrad/IC3ref>). Both solvers answered correctly. Timeout (TO) is 600 sec.

Benchmark	S	Result	$\mathbf{PDR}^{\mathbf{F-Kr}}$	IC3ref
latch0.smv	2^3	True	317 μ s	270 μ s
counter.smv	2^5	False	1.620 s	3.27 ms
power2bit8.smv	2^{15}	True	1.516 s	4.13 ms
ndista128.smv	2^{17}	True	TO	73.1 ms
shift1add256.smv	2^{21}	True	TO	174 ms

(c) Results with $\mathbf{PDR}^{\mathbf{IB-MDP}}$ (an excerpt of [22, Table 3]). Comparison is against PrIC3 [6] with four different interpolation generalization methods (none, linear, polynomial, hybrid). The benchmarks are from [6]. |S| is the number of states of the benchmark MDP. “GT pr.” is for the *ground truth probability*, that is the reachability probability $Pr^{max}(s_t \models \diamond(S \setminus \alpha))$ computed outside the solvers under experiments. The solvers were asked whether the GT pr. (which they do not know) is $\leq \lambda$ or not; they all answered correctly. The last five columns show the average execution time in seconds. – is for “did not finish,” for out of memory or timeout (600 sec.)

Benchmark	S	GT pr.	λ	$\mathbf{PDR}^{\mathbf{IB-MDP}}$	PrIC3			
					none	lin.	pol.	hyb.
Grid	10^2	$1.2E^{-3}$	0.3	0.31	1.31	19.34	–	–
			0.2	0.48	1.75	24.62	–	–
Grid	10^3	$4.4E^{-10}$	0.3	122.29	–	–	–	–
			0.2	136.46	–	–	–	–
BRP	10^3	0.035	0.1	–	–	–	–	–
			0.01	18.52	56.55	594.89	–	722.38
			0.005	1.36	11.68	238.09	–	–
ZeroConf	10^4	0.5	0.9	–	–	–	0.58	0.51
			0.75	–	–	–	0.55	0.46
			0.52	–	–	–	0.48	0.46
			0.45	<0.1	<0.1	<0.1	<0.1	<0.1
Chain	10^3	0.394	0.9	–	72.37	–	0.91	0.70
			0.4	–	80.83	–	0.93	–
			0.35	177.12	115.98	–	–	–
			0.3	88.27	66.89	557.68	–	–
DoubleChain	10^3	0.215	0.9	–	–	–	1.83	1.99
			0.3	–	–	–	1.88	1.96
			0.216	–	–	–	139.76	–
			0.15	7.46	–	–	–	–

(d) Ablation experiments: LT-PDR ($\mathbf{PDR}^{\mathbf{F-Kr}}$) vs. positive and negative LT-PDRs, implemented for the FSP for Kripke structures. The benchmarks are as in Table 2b, except for a new micro benchmark `simpleTrans.smv`. Timeout (TO) is 600 sec.

Benchmark	Result	LT-PDR	positive	negative
latch0.smv	True	317 μ s	1.68 ms	TO
power2bit8.smv	True	1.516 s	TO	TO
counter.smv	False	1.620 s	TO	2.88 μ s
simpleTrans.smv	False	295 μ s	TO	TO

7 Conclusions and Future Work

We have presented a lattice-theoretic generalization of the PDR algorithm called LT-PDR. This involves the decomposition of the PDR algorithm into positive and negative ones, which are tightly connected to the Knaster–Tarski and Kleene fixed point theorems, respectively. We then combined it with the coalgebraic and fibrational theory for modeling transition systems with predicates. We instantiated it with several transition systems, deriving existing PDR algorithms as well as a new one over Markov reward models. We leave instantiating our LT-PDR and categorical safety problems to derive other PDR-like algorithms, such as PDR for hybrid systems [29], for future work.

We will also work on the combination of our work and the theory of *abstract interpretation* [10, 12]. Our current framework axiomatizes what is needed of heuristics, but it does not tell how to realize such heuristics (that differ a lot in different concrete settings). We expect abstract interpretation to provide some general recipes for realizing such heuristics.

References

1. The 8th competitive event for hardware model checkers (HWMCC 2015) (2015). <http://fmv.jku.at/hwmcc15/>
2. Aguirre, A., Katsumata, S.: Weakest preconditions in fibrations. *Electronic Notes in Theoretical Comput. Sci.* **352**, 5–27 (2020). <http://www.sciencedirect.com/science/article/pii/S1571066120300487>, the 36th Mathematical Foundations of Programming Semantics Conference (2020)
3. Awodey, S.: *Category Theory*. Oxford Logic Guides. Oxford Univ. Press, Oxford (2006)
4. Baier, C., Katoen, J.: *Principles of Model Checking*. MIT Press, Cambridge (2008)
5. Baranga, A.: The contraction principle as a particular case of Kleene’s fixed point theorem. *Discret. Math.* **98**(1), 75–79 (1991)
6. Batz, K., Junges, S., Kaminski, B.L., Katoen, J.-P., Matheja, C., Schröder, P.: PrIC3: property directed reachability for MDPs. In: Lahiri, S.K., Wang, C. (eds.) *CAV 2020*. LNCS, vol. 12225, pp. 512–538. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-53291-8_27
7. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. In: *Tracts in Theoretical Computer Science* No. 53 (2001)
8. Bonchi, F., König, B., Petrisan, D.: Up-to techniques for behavioural metrics via fibrations. In: Schewe, S., Zhang, L. (eds.) *29th International Conference on Concurrency Theory, CONCUR 2018*. LIPIcs, vol. 118, pp. 17:1–17:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018). <https://doi.org/10.4230/LIPIcs.CONCUR.2018.17>
9. Bradley, A.R.: SAT-based model checking without unrolling. In: Jhala, R., Schmidt, D. (eds.) *VMCAI 2011*. LNCS, vol. 6538, pp. 70–87. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-18275-4_7
10. Cousot, P., Cousot, R.: Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: Graham, R.M., Harrison, M.A., Sethi, R. (eds.) *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages*, pp. 238–252. ACM (1977). <https://doi.org/10.1145/512950.512973>

11. Cousot, P., Cousot, R.: Constructive versions of Tarski's fixed point theorems. *Pacific J. Math.* **82**(1), 43–57 (1979). <https://doi.org/10.2140/pjm.1979.82.43>
12. Cousot, P., Cousot, R.: Systematic design of program analysis frameworks. In: Aho, A.V., Zilles, S.N., Rosen, B.K. (eds.) *Conference Record of the Sixth Annual ACM Symposium on Principles of Programming Languages*, pp. 269–282. ACM Press (1979). <https://doi.org/10.1145/567752.567778>
13. Eén, N., Mishchenko, A., Brayton, R.K.: Efficient implementation of property directed reachability. In: Bjesse, P., Slobodová, A. (eds.) *International Conference on Formal Methods in Computer-Aided Design, FMCAD 2011*, pp. 125–134. FMCAD Inc. (2011). <http://dl.acm.org/citation.cfm?id=2157675>
14. Gurfinkel, A.: IC3, PDR, and friends (2015). <https://arieg.bitbucket.io/pdf/gurfinkel.ssft15.pdf>
15. Hermida, C., Jacobs, B.: Structural induction and coinduction in a fibrational setting. *Inf. Comput.* **145**(2), 107–152 (1998). <https://doi.org/10.1006/inco.1998.2725>
16. Hoder, K., Bjørner, N.: Generalized property directed reachability. In: Cimatti, A., Sebastiani, R. (eds.) *SAT 2012. LNCS*, vol. 7317, pp. 157–171. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31612-8_13
17. Jacobs, B.: *Introduction to Coalgebra: Towards Mathematics of States and Observation*, Cambridge Tracts in Theoretical Computer Science, vol. 59. Cambridge University Press (2016). <https://doi.org/10.1017/CBO9781316823187>
18. Jacobs, B.P.: *Categorical Logic and Type Theory, Studies in logic and the foundations of mathematics*, vol. 141. North-Holland (2001). <http://www.elsevierdirect.com/product.jsp?isbn=9780444508539>
19. Katoen, J., Khattri, M., Zapreev, I.S.: A Markov reward model checker. In: *Second International Conference on the Quantitative Evaluation of Systems (QEST 2005)*, pp. 243–244. IEEE Computer Society (2005). <https://doi.org/10.1109/QEST.2005.2>
20. Komorida, Y., Katsumata, S., Hu, N., Klin, B., Hasuo, I.: Codensity games for bisimilarity. In: *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24–27, 2019*, pp. 1–13. IEEE (2019). <https://doi.org/10.1109/LICS.2019.8785691>
21. Kori, M., Hasuo, I., Katsumata, S.: Fibrational initial algebra-final coalgebra coincidence over initial algebras: turning verification witnesses upside down. In: Hadad, S., Varacca, D. (eds.) *32nd International Conference on Concurrency Theory, CONCUR 2021. LIPIcs*, vol. 203, pp. 21:1–21:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2021). <https://doi.org/10.4230/LIPIcs.CONCUR.2021.21>
22. Kori, M., Urabe, N., Katsumata, S.Y., Suenaga, K., Hasuo, I.: The lattice-theoretic essence of propertydirected reachability analysis (2022). <https://arxiv.org/abs/2203.14261>, a longer version
23. Mac Lane, S.: *Categories for the Working Mathematician*, 2nd edn. Springer, Heidelberg (1998). <https://doi.org/10.1007/978-1-4612-9839-7>
24. Rinetzky, N., Shoham, S.: Property directed abstract interpretation. In: Jobstmann, B., Leino, K.R.M. (eds.) *VMCAI 2016. LNCS*, vol. 9583, pp. 104–123. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49122-5_5
25. Seufert, T., Scholl, C.: Combining PDR and reverse PDR for hardware model checking. In: Madsen, J., Coskun, A.K. (eds.) *2018 Design, Automation & Test in Europe Conference & Exhibition. DATE 2018*, pp. 49–54. IEEE (2018). <https://doi.org/10.23919/DATE.2018.8341978>

26. Seufert, T., Scholl, C.: fbPDR: In-depth combination of forward and backward analysis in property directed reachability. In: Teich, J., Fummi, F. (eds.) *Design, Automation & Test in Europe Conference & Exhibition, DATE 2019*. pp. 456–461. IEEE (2019). <https://doi.org/10.23919/DATE.2019.8714819>
27. Sokolova, A.: Probabilistic systems coalgebraically: a survey. *Theor. Comput. Sci.* **412**(38), 5095–5110 (2011). <https://doi.org/10.1016/j.tcs.2011.05.008>
28. Sprunger, D., Katsumata, S., Dubut, J., Hasuo, I.: Fibrational bisimulations and quantitative reasoning. In: Cirstea, C. (ed.) *CMCS 2018*. LNCS, vol. 11202, pp. 190–213. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00389-0_11
29. Suenaga, K., Ishizawa, T.: Generalized property-directed reachability for hybrid systems. In: Beyer, D., Zufferey, D. (eds.) *VMCAI 2020*. LNCS, vol. 11990, pp. 293–313. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-39322-9_14
30. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math.* **5**(2), 285–309 (1955)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

