Master's Thesis

Probabilistic Programming: Applications of Martingales beyond Reachability

Author:
Thomas Haas

Advisor:
Prof. Dr. rer. nat. Roland Meyer
Institute of Theoretical Computer Science

Institute of Theoretical Computer Science
Technische Universität Braunschweig
Germany
October 16, 2019

Abstract

The inclusion of probabilistic aspects into systems, models and algorithms has become more and more common in recent days. Stochastic learning is used in the training of neural networks, randomized incremental constructions are found in the area of computational geometry and Markov decision processes are a frequently used model in control theory. These stochastic aspects are also found in probabilistic programming, an extension of classical programming. Along with programming comes the natural arising problem of verification of programs, that is, the problem of deciding certain properties about programs. Stochastic systems are much harder to verify and deciding qualitative properties such as safety and liveness requires more care. Unlike their deterministic counterpart, stochastic systems can also be analyzed in a quantitative manner by finding or approximating the probability that the system has certain properties. This thesis aims to give an overview over so called martingale-based methods which can be used to verify qualitative properties such as almost sure termination, but also to compute quantitative bounds like reachability probabilities. The models under consideration are probabilistic programs with nondeterminism. We look at the theory of martingale-based verification from two viewpoints, the probabilistic viewpoint via martingale theory and the order-theoretic viewpoint via fixed-point theory. We show the strong similarity between both viewpoints by redeveloping various martingale-based methods in both frameworks. These include martingales for deciding almost sure reachability, but also martingales for bounding reachability and recurrence probabilities. Lastly, we explain the idea of template-based synthesis to automatically find various martingales for program verification by using optimization techniques. Linear and polynomial templates are considered and experiments are done for the former. The results show that while linear template synthesis is a sound technique, it tends to give trivial or very bad approximations of probability bounds in many cases. In contrast, for establishing the quantitative property of positive almost sure termination, even linear templates can be useful.

Declaration of Authorship

I hereby declare that I am the sole author of this master thesis and that I have not used any sources other than those listed in the bibliography and identified as references. I further declare that I have not submitted this thesis at any other institution in order to obtain a degree.

Date

Signature

Contents

Listings

List of Tables

1 Introduction

The analysis and study of systems and their runtime behavior is an important research topic in the field of computer science. There are a variety of different systems and models to describe them, but the most well-known and powerful ones include Turing machines and standard computer programs written in a programming language such as C, Java or Haskell. These systems are sequential, that is, they start from some starting state and then evolve through repeated transitions to successor states. The state space may be described as a set S and the evolution as a sequence of states S^ω, the runs of the system. An often asked question is whether the runs of a given program exhibit a certain property such as guaranteed termination, safety or liveliness. This decision of program properties is known as verification. Termination, however, is known to be an undecidable property since computer programs are known to be theoretically as powerful as Turing machines for which the so called Halting problem is undecidable. So instead of looking at arbitrary, unrestricted systems, more practical approaches have been adopted. One is to restrict the system, such as the memory size of a program or the allowed actions of the system. While these restrictions can make many problems solvable, real life software may or may not adhere to any of the restricted models.

Another approach to deal with these kind of problems is to forgo the need to find an exact solution but to settle with incomplete but sound methods such as system behavior approximations. A common technique used to prove termination in term-rewriting systems is by finding a well-ordering on the terms which decreases as rewriting rules are applied. If such an ordering exists, then only finitely many rewrites can be performed before the system eventually can not perform any more rewrites[1, 2]. For deterministic programs, termination can be shown by finding so called ranking functions[3]. A ranking function can be seen as a kind of potential which decreases as the system evolves and guarantees that the system terminates as soon the potential goes below a certain threshold. In this sense, it induces a well-ordering on the state space of the system. A ranking function f roughly satisfies two properties:

1. $f(s) \geq f(s') + 1$ for all transitions $s \to s'$,
2. $f(s) \leq b \implies s$ is terminating state (where b is some threshold value)

Intuitively, the function assigns to each state the distance towards the terminating state and since it decreases over time (over transitions), it witnesses the fact that the system gets closer to termination as time progresses. A similar idea is used in systems modeled by ordinary differential equations (ODEs) where stability can be shown using so called Lyapunov functions [4, 5]. In either case, a system may be stable or may be terminating but a ranking function to witness it might not exist or might not be computable. Hence, finding a ranking function is generally a sound but incomplete method.

Another type of system where this approach is adopted are stochastic systems such as probabilistic rewriting systems[6, 7] and probabilistic programs. A probabilistic program,

in contrast to usual deterministic programs, can also model stochastic behavior by e.g. sampling from distributions or branching in a probabilistic manner. Thus, probabilistic programs exhibit strictly more behavior than their deterministic counterpart, making problems such as sure termination even harder to handle. Concretely, the notion of termination is replaced by the notion of almost sure termination, asking whether the system terminates with probability 1.

Ranking functions as described before are insufficient to witness almost sure termination. The problem is that many stochastic systems, even when terminating almost surely, occasionally move away from the terminating state caused by their stochastic behavior. So instead of a function which witnesses the systems guaranteed evolution towards termination along all transitions, a function which witnesses an *expected* evolution towards termination is more appropriate. For this purpose, so called probabilistic ranking functions (also sometimes called Lyapunov-style ranking functions) are used[8, 9, 10, 8, 11, 12]. Intuitively, they model the expected distance (expected number of steps/transitions) to the terminating state instead of the real distance.

1. $f(s) \geq \mathbb{E}_{s'}(f(s') \mid s \to s') + 1$ for all states s,
2. $f(s) \leq b \implies s$ is terminating state (where b is some threshold value)

These probabilistic ranking functions turn out to have supermartingale behavior, which is a special kind of behavior of stochastic processes whose expected value is non-increasing over time. This allows the analysis via martingale theory and makes so called martingale concentration inequalities like AzumaHoeffding's inequality applicable, giving certain bounds on the systems stochastic evolution.

Instead of looking at the termination problem in a qualitative, boolean way, the notion of termination probabilities becomes also available. Obviously, determining the exact probabilities encompasses the Halting problem for deterministic systems, making it generally undecidable. However, the problem of computing probabilities as real values allows a natural relaxation, namely, approximation instead of exact computation. Probabilistic ranking functions are not directly applicable in this quantitative setting, but it turns out that the underlying notion of supermartingales (and also submartingales) is able to be extended to witness probability bounds. Difference-bounded repulsing supermartingales are a variation of supermartingale processes that are capable of overapproximating the termination or reachability probability, that is, they give quantitative bounds instead of qualitative assertions[12, 13]. These make use Azuma-Hoeffding's inequality to derive their bounds.

The goal of this thesis is to give a solid understanding on how martingale-based methods can be used to tackle problems such as deciding (almost sure) termination of probabilistic programs, or finding bounds on the termination probability. Furthermore, they can also be used in a variety of other problems such as reachability, bounded reachability and tail reachability. We consider various different notions of martingales for different purposes such as

- Additive ranking supermartingales for almost sure reachability[12, 13]
- Higher order ranking supermartingales for bounding tail probabilities[14]

- Nonnegative repulsing supermartingales for overapproximation of reachability probabilities [11, 13]
- Nonnegative repulsing δ-supermartingales for overapproximation of bounded reachability probabilities[11]
- γ-Scaled submartingles for underapproximation of reachability probabilities[11, 15]
- A combination of super- and submartingale for recurrence probabilities

We develop the necessary theory from two different viewpoints, namely, once from the viewpoint of probabilistic martingale theory and once from the viewpoint of fixed point theory. From the former one, we derive soundness and completeness results by using the well-known Optional Stopping Theorem for martingales. We remark that we do not use any martingale concentration inequalities at all and solely rely on the Optional Stopping Theorem. From the fixed point theoretic point of view we derive the same results in a purely fixed point theoretic manner and establish the connection between martingales and fixed points. The extensive use of fixed point theory in this probabilistic martingale setting has recently been made by Takisaka et al.[11] and Kura et al.[14]. We build upon their work and extend it by providing the link between the commonly used pure martingale theoretic approaches and their new fixed point theoretic approaches. Our notation is in the remainder of this work is heavily influenced by Takisaka's and Kura's work.

Our setting is probabilistic programs with nondeterminism; an extension of probabilistic programs with nondeterministic assignments and branchings. In the presence of nondeterminism, the goal is to analyze the program under angelic and demonic behavior, i.e. to give probability bounds for the worst case and best case behaviors.

Martingale methods are applicable to infinite state systems such as those induced by probabilistic programs and in some cases even when non-deterministic behavior is present. Remarkably, martingales are effective in the sense that they can be computed - under some restrictions - automatically using template-based optimization[12, 13, 11, 14]. We give a simple description on how linear and polynomial templates can be used to compute martingales automatically via linear programming (LP) or semi-definite programming (SDP) and apply the linear template method experimentally.

Remark. For ease of presentation, the word martingale is often used to include the notions of supermartingales and submartingales although they are technically no martingales in the formal sense. In fact, none of the presented methods in this thesis use a strict martingale but always one of its relaxations.

The thesis is structured as follows:

1. Chapter 2 revises the basic theories of probability, measurability and martingales in Section 2.1. Section 2.2 revises order theory with a focus on lattices and fixed points.
2. Chapter 3 presents the main theory of this thesis showing how martingale-based methods can be used in program verification. Section 3.1 introduces probabilistic programs (with nondeterminism) syntactically and gives their semantics in terms of probabilistic control flow graphs (pCFGs) and Markov decision processes (MDPs).

Section 3.2 introduces different concrete super- and submartingales which have been used to verify programs. Each type of martingale has its own dedicated subsection.

3. Chapter 4 explains how different martingales can be automatically synthesized using template based optimization. Section 4.1 shows that martingale constraints for linear templates can be reduced to a linear programming problem. Section 4.2 shows that polynomial template synthesis also works by using semidefinite programming instead. Lastly, Section 4.3 demonstrates linear template-based synthesis by applying it to various programs.

2 Background

This chapter explains the basic concepts, notions and definitions used throughout this thesis. In particular, it revises the basic ideas of probabilistic theory and order theory with an explicit focus on their specialized theories of martingales, measurability and fixed points. Both of these theories are fundamental, and hence, crucial to the understanding of the remainder of this thesis. If the reader is familiar with these topics he may directly skip to Chapter 3 and come back to this chapter when necessary.

2.1 Probability Theory

This section revises the basics of probabilistic theory needed to understand the main content of this paper. A special focus lies on measure theory and martingale theory which both are presented in their own subsections.

Definition 2.1.1 (Probability space)**.** A probability space is a triple $(\Omega, \mathcal{F}, \mu)$ consisting of

1. a set Ω called the *sample space*,
2. a family $\mathcal{F} \subseteq \mathcal{P}(\Omega)$ of subsets of Ω forming a σ-Algebra,
3. a probability measure $\mu : \mathcal{F} \to [0, 1]$.

The elements of $\mathcal{F}$ are called *measurable sets*, *measurable events* or simply just *events*. A set E which is not contained in $\mathcal{F}$ is called *non-measurable*. A σ-Algebra, in short, is a family of subsets which contains the whole space Ω and is closed under countable unions, countable intersections and complements. The probability measure μ simply assigns to each event the probability that that event occurs. It does so in a reasonable manner, given by the following axioms:

- $\mu(\Omega) = 1$,
- $\mu(\bigcup_{i \in \mathbb{N}} E_i) = \sum_{i \in \mathbb{N}} \mu(E_i)$ for countable collections of disjoint events E_i.

The second property of μ is also called σ-additivity or countable additivity. It is easy to see that μ is monotone, i.e. $\mu(A) \leq \mu(B)$ for $A \subseteq B$. The reason why one restricts to a subset $\mathcal{F} \subseteq \mathcal{P}(\Omega)$ of measurable events is discussed in Subsection 2.1.1 in more detail. Broadly speaking, it is by no means easy to assign probabilities to arbitrary subsets of Ω because of rather complicated, pathological subsets (e.g. the Vitali sets in $\mathbb{R}$). The definition of σ-Algebra and measurability is also discussed there.

Definition 2.1.2 (Random variable)**.** A random variable $X : \Omega \to Y$ is a measurable function[1] between the probability space $(\Omega, \mathcal{F}, \mu)$ and the measurable space (Y, Σ).

[1] Measurability is defined in the following subsection

In most cases the target measurable space is $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ or $(\overline{\mathbb{R}}, \mathcal{B}(\overline{\mathbb{R}}))$ where $\mathcal{B}$ denotes the Borel σ-Algebra and $\overline{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$ the extended reals, then X is called a (extended) real-valued random variable. The probability measure μ on $(\Omega, \mathcal{F})$ defines a push-forward measure $\mu_X : \Sigma \to [0,1]$ on (Y, Σ) by assigning to each $S \in \Sigma$ the probability

$$\mu_X(S) \coloneqq \Pr(X \in S) = \mu(\{\omega \in \Omega \mid X(\omega) \in S\}) = \mu(X^{-1}(S)).$$

For real-valued random variables the concept of integration is definable.

Definition 2.1.3 (Integration)**.** Let $X : \Omega \to [0, \infty]$ be a nonnegative, real-valued random variable. Then the integral of X with respect to a measure μ can be defined as

$$\begin{aligned}\int_\Omega X \mathrm{d}\mu &= \int_\Omega X(\omega)\mu(\mathrm{d}\omega) \\ &\coloneqq \sup\{\sum_i \inf\{X(\omega) \mid \omega \in A_i\}\mu(A_i) \mid \{A_i\} \text{ is a finite partition of } \Omega\}.\end{aligned}$$

X is said to be *integrable* if $\int_\Omega X d\mu < \infty$ holds.

This definition may be extended to real-valued random variables $X : \Omega \to \overline{\mathbb{R}}$ by splitting X into its positive and negative parts and integrating them separately:

$$\int_\Omega X \mathrm{d}\mu \coloneqq \int_\Omega X^+ \mathrm{d}\mu - \int_\Omega X^- \mathrm{d}\mu,$$

where $X^+ \coloneqq \max\{0, X\}$ and $X^- \coloneqq \max\{0, -X\}$. This integral is only defined if either of the right-hand side integrals takes a finite value. X is then said to be integrable if $|X| \coloneqq X^+ + X^-$ is integrable, that is, both integrals are finite.

Using the push-forward measure μ_X of a random variable $X : \Omega \to \overline{\mathbb{R}}$ one may alternatively integrate over $\overline{\mathbb{R}}$ instead of Ω via the identity

$$\int_\Omega X \mathrm{d}\mu = \int_{\overline{\mathbb{R}}} x \mathrm{d}\mu_X = \int_{\overline{\mathbb{R}}} x \mu_X(\mathrm{d}x).$$

An important theorem related to integration is the so called *Monotone Convergence Theorem* which shows the exchangability of taking the monotone limit and integration.

Theorem 2.1.1 (Monotone Convergence Theorem)**.** *[16, Theorem 2.146] Let $\{X_n\}_{n\in\mathbb{N}}$ be a monotonically increasing sequence of nonnegative random variables that converges point-wise to a random variable X. Then the limit operation and the integration operation commute:*

$$\lim_{n\to\infty} \int X_n \mathrm{d}\mu = \int \lim_{n\to\infty} X_n \mathrm{d}\mu = \int X \mathrm{d}\mu.$$

Definition 2.1.4 (Moments)**.** The k-th moment of a real-valued random variable X is given by

$$\mathbb{E}(X^k) \coloneqq \int_\Omega X^k \mathrm{d}\mu.$$

The first moment $\mathbb{E}(X)$ is also called the expectation of X. The second moment of the random variable $(X - \mathbb{E}(X))$ is also known as the variance of X.

In Section 3.1 the notion of probabilistic programs over real-valued variables is introduced. The objects of interest are the possible runs of the system, that is, the sample space is the set of all sequences of program configurations. A typical quantity to analyze is the expected termination time. The termination time may be given by a random variable T which assigns to each run the time it takes to enter a terminating state or ∞ if it does not terminate, then the problem of deciding positive[2] almost sure termination is reduced to computing $\mathbb{E}(T)$.

2.1.1 Measurability

Many results about measurability in this section apply to general measurable spaces and are not unique to probability spaces. Accordingly, most results about measurability are stated with respect to arbitrary measurable spaces.

The goal of measure theory is to measure objects (subsets) and assign them some kind of real value indicating mass, size or probability. Consider for example the 2d-plane $\mathbb{R}^2$. A reasonable measure in this setting is the measure of area. For many geometric objects like rectangles, disks or arbitrary polygons, elementary formulas are known which give the area of such objects. An immediately arising question is how one measures more complex objects or even arbitrary subsets of for example $\mathbb{R}^2$. Measure theory tries to give an answer to this question. It turns out that assigning lengths, areas or volumes to arbitrary subsets of R^n is not easy at all. Indeed, the well known Vitali sets discovered by Giuseppe Vitali give an example of subsets of $\mathbb{R}$ which cannot be assigned a geometric length in a consistent manner. Hence, the Vitali sets are said to be *non-measurable*. Another example is the Banach-Tarski paradox which shows that a ball can be decomposed into multiple parts and then reassembled into two balls of the same size as the original, effectively doubling the original ball. This decomposition is non-trivial and needs the Axiom of Choice (so do the Vitali sets). The Banach-Tarski paradox shows that the parts obtained in the decomposition are in a sense so complexly or oddly shaped that the concept of geometric size fails to apply. In both cases, it was assumed that a reasonable geometric measure shall be invariant under translation and rotation. This shows that under these simple geometric conditions we are forced to somehow restrict the objects to be measured. These kind of problems also arise when trying to assign a probability to events instead of a geometric size to geometric object. To circumvent such problems, instead of trying to measure arbitrary subsets, one restricts to a family of so called *measurable* subsets which are given in the form of a σ-Algebra.

Definition 2.1.5 (σ-Algebra)**.** A σ-Algebra Σ of a set X is a family of subsets which

- contains the whole set $X \in \Sigma$,
- is closed under countable unions: $\bigcup_{i \in \mathbb{N}} A_i \in \Sigma$ for $A_i \in \Sigma$,
- and is closed under complementation: $A \in \Sigma \implies (X \setminus A) \in \Sigma$.

[2] Positive a.s. termination means finite termination time, while (standard) a.s. termination refers to the probability of termination to be 1. The former implies the latter, but not vice versa.

The empty set $\emptyset = X \setminus X \in \Sigma$ is also contained in Σ and by De Morgan's laws, Σ is also closed under countable intersections.

Intuitively, one expects that if some finite set of objects are measurable then surely their union, intersection and complement can be measured as well. By applying limit procedures, this concept of measurability can then be extended to countable unions and intersections as well. This intuition leads to the idea to actually construct a σ-Algebra from a set of simple measurable objects. Assume that there is a family of subsets for which it is easy to define some kind of measure, e.g. consider the real line $\mathbb{R}$ and its open intervals (a, b). It is straightforward to assign a size to intervals (a, b), namely the length $b - a$. Starting from these open intervals, one can construct a σ-Algebra by finding the closure with respect to countable unions and complements. This specific σ-Algebra which is generated from the open intervals is the so called Borel σ-Algebra. It will be formally introduced later in this section. Although defining this *generated* σ-Algebra is straightforward, extending the measure on the simple intervals to the whole σ-Algebra is not. Some more assumptions on the measure and the generating set are needed to make it work as expected.

Borel σ-Algebras are the most important σ-Algebras in measure theory, and the most commonly analyzed σ-Algebras are either the Borel σ-Algebras themselves or some larger σ-Algebra containing it. Before further going into this specific σ-Algebra, we first formalize the previously mentioned idea of constructing σ-Algebras.

Definition 2.1.6 (Generated σ-Algebra)**.** Let $\mathcal{A}$ be a family of subsets of a ground set X. Define $\sigma(\mathcal{A})$ to be the smallest σ-Algebra containing $\mathcal{A}$. It is formally given by

$$\sigma(\mathcal{A}) = \bigcap_{\mathcal{A} \subseteq \Sigma} \Sigma$$

where the intersection ranges over all possible σ-Algebras Σ containing $\mathcal{A}$. This intersection is well defined, because the family of σ-Algebras is closed under arbitrary intersections and it is nonempty because it contains a largest element $\mathcal{P}(X)$ which always contains $\mathcal{A}$.

Remark. A pre-measure defined on a family of subsets can in generally not be extended to a measure on its generated σ-Algebra. The generating set has to satisfy some extra properties in order for this to work. Carathéodory's extension theorem[17] states that if the generating set forms a ring, then the pre-measure on this ring can be extended to a measure on the generated σ-Algebra. If the pre-measure is σ-additive, then this extension is unique.

Now we are ready to define the object of interest, namely *measurable* spaces.

Definition 2.1.7 (Measurable space)**.** A *measurable space* is a tuple (X, Σ) where

- X is a set,
- and Σ is a σ-Algebra over X.

The elements of Σ are said to be *measurable*.

As the name suggests, a measurable space can be *measured* by additionally defining a function μ, the so called *measure*, which assigns a nonnegative value to each measurable set

$S \in \Sigma$. Depending on the context, this value can then be interpreted as a size, mass or even probability. Generally, there are many ways to choose μ but once a choice has been fixed, one has a *measure space*.

Definition 2.1.8 (Measure space)**.** A *measure space* is a tuple (X, Σ, μ) where

- the tuple (X, Σ) is a measurable space,
- the *measure* $\mu : \Sigma \to [0, \infty]$ is a σ-additive function satisfying $\mu(\emptyset) = 0$. σ-additivity of μ means that for any countable collection $E_n \in \Sigma$ of disjoint sets, the equality $\mu(\bigcup_n E_n) = \sum_n \mu(E_n)$ is satisfied.

μ is called *finite* if $\mu(X) < \infty$. It is *σ-finite* if X is the countable union of measurable sets with each having finite measure.

If we denote the Borel σ-Algebra of $\mathbb{R}$ by $\mathcal{B}(\mathbb{R})$, then $(\mathbb{R}, \mathcal{B}(\mathbb{R}), \mu)$ is the typical measure space associated with $\mathbb{R}$, where μ is the unique measure that assigns to an open interval (a, b) its length $b - a$. This measure is σ-finite since $\mathbb{R}$ is the countable union of the intervals $I_n = (-n, n)$ $(n \in \mathbb{N})$, each of which has a finite measure $\mu(I_n) = 2n < \infty$. Another more simple example is a discrete set S with the counting measure $\mu = \#$ on the power set $\mathcal{P}(S)$ of S. The counting measure simply assigns to each set the number of elements in that set or ∞ if the set is infinite. Again, $\#$ is σ-finite if S is countable infinite and even finite if S is finite.

A probability space as defined before is nothing but a special kind of measure space, namely one, where the measure is a probability measure and assigns mass 1 to the whole space. In particular, probability spaces are finite measure spaces and every finite measure space (except the trivial constant 0 measure space) is equivalent to a probability space by simply normalizing the measure, that is, by setting $\mu'(A) = \mu(A)/\mu(X)$. Even σ-finite measure spaces can be transformed to probability spaces by making use of a countable partition into finite measure sets. For example, consider throwing a six-sided dice. A possible measure space is $(X, \mathcal{P}(X), \#)$ where $X = \{1, 2, 3, 4, 5, 6\}$ and $\#$ is the counting measure. Since $\#(X) = 6$, this space is no probability space but it is possible to construct the following probability space $(X, \mathcal{P}(X), \mu)$ where $\mu(A) = \#(A)/\#(X) = \#(A)/6$. In this space every possible outcome has equal probability of $1/6$. Note that the original finite measure space and the derived probability space differ only in the associated measure. It is quite common to have different measures on the same ground set and σ-Algebra, especially when dealing with probability spaces. This gives reason to the notion of measurable spaces (X, Σ) where no specific measure is fixed. Also note that the question whether a particular set A is measurable is independent of the actual measure and solely depends on the σ-Algebra. Hence, it is enough to consider measurable spaces instead of measure spaces to decide measurability in many cases. The focus of this subsection lies on the aspect of measurability therefore it is mostly concerned with measurable spaces.

The definition of Borel sets and Borel σ-Algebras heavily rely on the idea of openness and closedness of sets. While both of these concepts are familiar in the setting of euclidean spaces $\mathbb{R}^n$, for more complex or even more simple spaces like discrete ones, these concepts might be obscure. First, let us define openness in a more general setting.

Definition 2.1.9 (Topological space)**.** A tuple (X, τ) is called a *topological space* if X is a set and τ is a family of subsets of X satisfying:

- $\emptyset \in \tau$ and $X \in \tau$,
- τ is closed under arbitrary unions,
- τ is closed under finite intersections.

τ is then called an *open topology* on X and its elements are called *open sets*.

An alternative definition requires closure under finite unions and arbitrary intersections, in which case it is referred to as *closed topology* and its elements are *closed sets*. Both definitions are in a sense equivalent and both induce the same Borel structure. Because of this, every mentioning of topologies in the following sections implicitly refers to open topologies as defined above unless otherwise specified. Analoguous to the euclidean case, a set is said to be *closed* iff its complement is open, i.e. element of the topology. The ground set X is always open and closed (clopen), so is the empty set. Generally, for every open topology one can construct a closed topology by complementing each set in the open topology.

The standard example of a topological space (with open topology) is again the real number line $\mathbb{R}$ together with its open sets as the topology τ. In the reals, a subset $A \subseteq \mathbb{R}$ is considered open if and only if for every point $a \in A$ there exists an interval $I(a, r) = \{x \in \mathbb{R} \mid |x - a| < r\}$ centered at a with a radius $r > 0$ which is fully contained in A. These intervals themselves are open and are also called *open neighborhoods* of a. It is easy to check that these usual open sets of $\mathbb{R}$ indeed satisfy the necessary properties of an open topology. Interestingly, it is enough to start from the open intervals of $\mathbb{R}$ and then find the smallest topology which contains all open intervals to get the open topology τ of $\mathbb{R}$. Here it makes sense to speak of the *smallest* topology because the family of topologies is closed under arbitrary intersections, hence one can intersect all topologies which contain the open intervals and get a smallest topology w.r.t. set inclusion. In this case we say that the open intervals *generate* τ. In general euclidean spaces $\mathbb{R}^n$, the standard topology is generated by the open balls $B(a, r) = \{x \in \mathbb{R}^n \mid \mathrm{d}(x, a) < r\}$ where d is the usual euclidean distance. This topology is naturally defined on $\mathbb{R}^n$ by its metric d. Indeed, for any metric space, i.e. a space equipped with a distance function, there is always a naturally corresponding topology generated via the distance function.

A discrete set S can be endowed with its power set $\mathcal{P}(S)$ to form a topological space. In this case, every subset of S is considered open, in particular, all singletons $\{s\} \subseteq S$ are open. This topology is called the *discrete* topology. In contrast there also exists the *indiscrete* or *trivial* topology which only contains the whole set and the empty set.

Definition 2.1.10 (Continuous function)**.** Let (X, τ) and (Y, τ') be two topological spaces. A function $f : X \to Y$ is called *continuous* if for every open set $B \in \tau'$ in Y the preimage $f^{-1}(B)$ is open in X, that is, $f^{-1}(B) \in \tau$, or equivalently if $f^{-1}(\tau') \subseteq \tau$ holds.

Remark. For the real number line $\mathbb{R}$ with its usual topology, a function f on $\mathbb{R}$ is topologically continuous iff it is continuous in the usual sense. In other words, the above definition is a generalization of the known notion of continuity of real functions.

Topological spaces themselves are a big research topic in mathematics, however, we are mainly interested in σ-Algebras and measurability. Having this in mind, we are now ready to formally define the notion of Borel sets and the Borel σ-Algebra which are fundamental to measure theory.

Definition 2.1.11 (Borel σ-Algebra)**.** Let (X, τ) be a topological space. Let $\mathcal{B}(X, \tau)$ be the smallest σ-Algebra which contains the topology τ. This σ-Algebra is called the Borel σ-Algebra of X and if the topology is clear from context, we simply write $\mathcal{B}(X)$ instead. Its elements are called *Borel sets*. Formally,

$$\mathcal{B}(X, \tau) = \sigma(\tau)$$

is simply the σ-Algebra generated by τ.

Sometimes we say that a Borel σ-Algebra is generated by some family of open sets $\mathcal{A}$, like the open balls in $\mathbb{R}^n$. This means that $\mathcal{A}$ generates a topology $\tau = \tau(\mathcal{A})$ and from that topology we generate $\mathcal{B}(X) = \sigma(\tau)$ which is generally different from the σ-Algebra $\sigma(\mathcal{A})$ generated by $\mathcal{A}$ directly. The reason is that topologies are closed under arbitrary unions while σ-Algebras are only closed under countable unions. In cases where uncountable unions of a generating set $\mathcal{A}$ can be different from countable unions, we have $\sigma(A) \neq \sigma(\tau(\mathcal{A}))$. Consider for example all singleton sets of $\mathbb{R}$ as a generating set. Every subset of $\mathbb{R}$ can be expressed as an uncountable union of all its elements, hence the generated topology is $\mathcal{P}(\mathbb{R})$. On the other hand, the generated σ-Algebra has only countable unions and hence only consists of all countable subsets of $\mathbb{R}$ and its complements (co-countable sets). However, in many practical cases such as euclidean spaces with the open topology $\sigma(\mathcal{A})$ and $\sigma(\tau(\mathcal{A}))$ do indeed coincide.

Since σ-Algebras are closed under complements, starting from a closed topology (e.g. the one generated by the closed intervals on $\mathbb{R}$) yields the same Borel σ-Algebra, making differentiation between closed and open topologies unnecessary in this setting. Generally, the Borel σ-Algebra contains all open and closed sets, meaning all closed and open sets are measurable with respect to that σ-Algebra.

Definition 2.1.12 (Measurable function)**.** Let (X, Σ) and (Y, Σ') be two measurable spaces. A function $f : X \to Y$ is said to be *measurable* if the preimage $f^{-1}(B)$ of every measurable set $B \subseteq Y$ in Y is measurable in X. Formally, $f^{-1}(B) \in \Sigma$ holds for all $B \in \Sigma'$, or equivalently $f^{-1}(\Sigma') \subseteq \Sigma$.

Measurable functions, just like continuous functions, are closed under composition. Oftentimes the topology of a given set X is implied or clear by context. In those cases, the associated Borel σ-Algebra $\mathcal{B}(X)$ is implied as well and is not specifically mentioned. We then call a function $f : X \to Y$ Borel measurable, if f is a measurable function from $(X, \mathcal{B}(X))$ to $(Y, \mathcal{B}(Y))$. The set of all Borel measurable functions from X to Y is denoted by $\mathcal{B}(X, Y)$.

By the definition of measurability of a function f one needs to check that all preimages $f^{-1}(A)$ of measurable $A \subseteq Y$ are also measurable in X. As those sets A can generally be of complicated shape, this task can be difficult. However, it actually suffices to consider only the preimages of a generating set of the σ-Algebra in the codomain. The following lemma makes this precise.

Lemma 2.1.2. *Let (X, Σ) and (Y, Σ') be two measurable spaces and let $\Sigma' = \sigma(\mathcal{A})$ be generated by some family of sets $\mathcal{A}$. A function $f : X \to Y$ is measurable iff $f^{-1}(A)$ is measurable in X for all A in $\mathcal{A}$, or equivalently, iff $f^{-1}(\mathcal{A}) \subseteq \Sigma$.*

Proof. The implication from left to right is immediate. First observe that measurability of f is equivalent to the inclusion $f^{-1}(\Sigma') \subseteq \Sigma$ and that the statement we need to show is equivalent to $f^{-1}(\mathcal{A}) \subseteq \Sigma$. Since $\mathcal{A} \subseteq \Sigma'$ we see that $f^{-1}(\mathcal{A}) \subseteq f^{-1}(\Sigma') \subseteq \Sigma$ holds.

For the other direction, note that inverse images preserve arbitrary unions, intersections and complements. Because of this, we see that the equality $f^{-1}(\sigma(\mathcal{A})) = \sigma(f^{-1}(\mathcal{A}))$ holds. Now it follows that $f^{-1}(\Sigma') = f^{-1}(\sigma(\mathcal{A})) = \sigma(f^{-1}(\mathcal{A})) \subseteq \sigma(\Sigma) = \Sigma$, where the inclusion holds because of the premise and the last equality holds because Σ is already a σ-Algebra and thus the smallest one containing itself. □

Corollary 2.1.2.1. *Let $f : X \to Y$ be a continuous function between two topological spaces, then f is Borel measurable.*

Proof. The Borel σ-Algebra $\mathcal{B}(Y)$ of Y is generated by the open sets of Y's topology. From the previous lemma it suffices to show that the f-preimages of these open sets are measurable in X. Since f is continuous, the preimages of open sets in Y are open in X. By definition of $\mathcal{B}(X)$ all open sets are measurable. □

Remark. The Borel σ-Algebra on X is actually the smallest σ-Algebra which makes all continuous functions $f : X \to X$ measurable. Notice that the identity function $\mathrm{id} : X \to X$ is continuous. Therefore all open subsets $A = \mathrm{id}^{-1}(A)$ must be measurable but $\mathcal{B}(X)$ is precisely the smallest σ-Algebra that makes all open sets measurable.

It is often desirable to have larger σ-Algebras to make more functions measurable and hence allow better analysis of the measure spaces involved. But as mentioned in the introductory example, choosing too large σ-Algebras such as the power set $\mathcal{P}$ may introduce problems in finding actual measures on the measurable space. Therefore, one approaches the problem from the other side and starts with a smaller family of sets which is easily measured, such as the open balls, and then constructs a σ-Algebra accordingly. This approach has given us the Borel σ-Algebra so far. However, it is possible to extend this σ-Algebra even further in natural ways.

One such way is the so called *completion* of measure spaces. The idea is to add *null sets*, i.e. sets which have measure 0, to the family of measurable sets. Unlike before, this construction explicitly relies on the given measure.

Definition 2.1.13 (Complete measure space)**.** A measure space (X, Σ, μ) is said to be *complete* if every subset $N \subseteq A$ of a null set $A \in \Sigma$, i.e. a set satisfying $\mu(A) = 0$, is measurable:

$$N \subseteq A \in \Sigma \text{ and } \mu(A) = 0 \implies N \in \Sigma.$$

Given an incomplete measure space (X, Σ, μ) one can construct its *completion* (X, Σ_μ, μ^*) in the following way:

1. Find $\mathcal{N} = \{N \subseteq X \mid \exists A \in \Sigma : N \subseteq A \wedge \mu(A) = 0\}$,
2. define $\Sigma_\mu = \sigma(\Sigma \cup \mathcal{N})$,

3. let $\mu^*(A) = \inf\{\mu(B) \mid A \subseteq B \in \Sigma\}$.

μ^* is called the *outer measure* and in case μ is σ-finite, this μ^* is the unique extension of μ to Σ_μ[17].

Going back to the example of $\mathbb{R}$ and its Borel σ-Algebra generated by the open intervals, the standard measure can be defined as the unique measure that satisfies $\mu((a,b)) = b - a$. This measure space, however, is not complete. It can be shown that the Cantor set C has measure zero so all of its subsets should be included in a complete measure space. A counting argument shows that $\mathcal{B}(\mathbb{R})$ has cardinality of the continuum and so does the Cantor set C. Therefore the power set $\mathcal{P}(C)$ is strictly larger than the continuum. That means there must exist sets in $\mathcal{P}(C)$ which are not contained in $\mathcal{B}(\mathbb{R})$, however, all those sets are subsets of a null set and hence $\mathcal{B}(\mathbb{R})$ is not complete with respect to μ. Interestingly, this also shows that the completion $\mathcal{B}(\mathbb{R})_\mu$ is strictly larger than $\mathcal{B}(X)$ in cardinality The associated measure μ^* can be identified as the Lebesgue measure known from integration theory.

The obvious disadvantage of completion is the explicit dependency on the measure μ. In Section 3.2 we are interested in analyzing runtime behavior of stochastic systems, in particular, of probabilistic programs. In that case we define a measure on the set of runs (i.e. sequences of program configurations) but this measure will depend on the initial state of the system. Because we also consider nondeterministic behavior, the probability measure will also be dependent on the way in which nondeterminism gets resolved. Through this dependencies, it is not possible to fix any completion of the associated σ-Algebra a priori. Rather, we want to measure runs independent of the concrete completion, i.e. we want to have sets which are measurable with respect to every possible completion. Obviously, the σ-Algebra of the original, uncompleted measure space satisfies this condition but there actually exists a larger σ-Algebra satisfying it.

Definition 2.1.14 (Universal completion)**.** Let (X, Σ) be a measurable space. The *universal completion* Σ^* is defined by

$$\Sigma^* = \bigcap_\mu \Sigma_\mu$$

where μ ranges over all σ-finite measures on Σ. A set $A \in \Sigma^*$ is said to be *universally measurable.*

If $\Sigma = \mathcal{B}(X)$ is a Borel σ-Algebra, then we write $\mathcal{U}(X)$ as the universal completion of $\mathcal{B}(X)$, i.e $\mathcal{U}(X) = \mathcal{B}(X)^*$. It is not immediately clear that the universal completion is larger than the Borel σ-Algebra, but it can be shown that it indeed is in many cases. Specifically, for the real numbers $\mathbb{R}$ the Borel σ-Algebra $\mathcal{B}(\mathbb{R})$ is strictly contained in its universal completion $\mathcal{U}(\mathbb{R})$[18, Appendix B.3]. Similarly to Borel measurable functions, we say a function $f : X \to Y$ is *universally measurable* if f is measurable from $(X, \mathcal{U}(X))$ to $(Y, \mathcal{U}(Y))$. The set of all universally functions from X to Y is denoted by $\mathcal{U}(X, Y)$. The following lemma gives a characterization of universal measurability.

Lemma 2.1.3. *[18, Proposition 7.44, Corollary 7.44.1] A function $f : X \to Y$ is universally measurable iff f is measurable from $(X, \mathcal{U}(X))$ to $(Y, \mathcal{B}(X))$.*

To determine universal measurability, it suffices to look at preimages of Borel sets.

Corollary 2.1.3.1. *Every Borel measurable function $f : X \to Y$ is universally measurable.*

For the analysis of systems, one needs to analyze sequences of states or configurations, i.e. if S denotes the state space then one is interested in the set of finite runs S^* or infinite runs S^ω. It is possible to construct appropriate measurable spaces for either of those cases if a measurable space for S is known. This is done by defining a product operation on σ-Algebras such that one can construct a σ-Algebra for e.g. $X \times Y$ from ones on X and Y.

Definition 2.1.15 (Product measurable space)**.** Let I be an index set and (X_i, Σ_i) a family of measurable spaces indexed by $i \in I$. The product space (X, Σ) is defined such that

- $$X \coloneqq \prod_{i \in I} X_i$$
- Σ is the smallest σ-Algebra such that all the projections $\pi_i : X \to X_i$ are measurable.

For countable products, the second condition that all the projections are measurable can be stated in a different manner. The σ-Algebra Σ can equivalently be expressed as the σ-Algebra generated by the rectangles

$$\{A_1 \times A_2 \times \cdots \times A_i \times \cdots \mid A_n \in \Sigma_n\}.$$

To see this, let $A = A_1 \times A_2 \times \ldots$ be a rectangle. Since A_n is measurable in X_n and the projection π_n is a measurable function, the preimage $\pi_n^{-1}(A_n) = X_1 \times \cdots \times X_{n-1} \times A_n \times X_{n+1} \times \ldots$ is measurable in the product space X. Now $A = \bigcap_{n \in \mathbb{N}} \pi_n^{-1}(A_n)$ is the countable intersection of measurable sets and hence measurable itself.

Conversely, if all rectangles are measurable then in particular also the rectangle $X_1 \times \cdots \times X_{n-1} \times A_n \times X_{n+1} \times \cdots = \pi_n^{-1}(A_n)$, hence π_n is a measurable function for all $n \in \mathbb{N}$.

For uncountable products, a rectangle can no longer be expressed as a countable intersection so the first part of the proof fails. In fact, the product σ-Algebra is then smaller than the σ-Algebra generated by the rectangles.

Yet another way to construct the product σ-Algebra for countable products is by considering cylinder sets of the form $C_n(A_1, A_2, \ldots, A_n) = A_1 \times A_2 \times \cdots \times A_n \times \prod_{i>n} X_i$ where all the A_i are measurable in X_i. These cylinder sets generate the *cylindrical* σ-Algebra which coincides with the product σ-Algebra for countable products. This is easy to see, since every cylinder is a rectangle and every rectangle can be expressed as a countable intersection of cylinder sets.

Cylinder sets are a natural way to represent runtime behavior of infinite state systems. Consider again a system with some countable state space S which runs indefinitely. A run or execution of the system then corresponds to an element in S^ω. A cylinder set $C_n = A_1 \times A_2 \times \cdots \times A_n \times \prod_{i>n} S$ distinguishes runs solely on the finite prefixes of length n and thus containment of a run can be verified after finitely many steps. Many interesting sets such as the set of terminating runs can be expressed in terms of cylinder sets as well. A particularly useful property however, is the fact that a (cylinder) measure on the cylinder sets can be extended to a measure on the infinite runs S^ω. For probabilistic systems, defining a probability on the cylinders then naturally gives rise to a probability

measure on the set of infinite runs (e.g. via Kolmogorov's Extension Theorem[17, Theorem 7.7.1]). We make use of this in the main part of this thesis where we define a measure exactly on the cylinder sets and extend it to all infinite runs.

2.1.2 Martingale Theory

A martingale in probabilistic theory is a special kind of stochastic process, that is, a sequence of random variables for which the expected value at some time $t+1$ equals the current value at time t given all previously observed values. In other words, if one stops the process at some point of time t and observes a current value of m_t then the value m_{t+1} in the next step is expected to stay at the current value on average. For example, consider a simple coin tossing game where the player bets one dollar on head every round. Assuming the player can play as long as he wants and is only forced to stop when he goes broke, this game constitutes a martingale. If at any point in time the player has a positive amount m_t of money, then after playing one more round his expected earning is $\mathbb{E}(m_{t+1}) = 0.5(m_t+1) + 0.5(m_t-1) = m_t$. One can derive the result that the expected earnings over multiple rounds does not change either. In particular, if the player starts with an initial wealth m_0, then the expected earnings after t more rounds stays at value m_0. The player neither gains nor loses money on average no matter how long he decides to play the game.

Now consider the same game but with a different betting strategy. Instead of constantly betting a single dollar, the player doubles his bet every time he loses until he wins once. The player's idea is to make up all his losses in a single win. If he starts with an initial bet m_0 then after n losses in a row he will have lost $\Sigma_{i=1}^{n} m_0 2^{i-1} = m_0(2^n-1)$ dollars, assuming he does not go broke in between rounds. If he wins on the $(n+1)$-th round he gains $m_0 \cdot 2^n$ dollars back giving him a total earning of $m_0 \cdot 2^n - m_0(2^n-1) = m_0$ dollars. Since the probability to permanently lose goes towards 0, the player hopes to eventually win once and make a profit. This type of betting system, where the player tries to make up for all the losses by constantly increasing his bet, is also known as a *martingale* betting system. It is easy to see that this betting strategy is still a martingale since in every round the player either wins or loses the bet with probability $1/2$ each, which averages out to a winning of zero dollars per round. In contrast to naive intuition, Martingale theory then predicts that for any initial bet m_0 and total wealth w, the player will not make any winnings on average despite his clever betting strategy. Only if the player has unbounded wealth, time and bets, his betting strategy can be considered a winning strategy.

Definition 2.1.16 (Simple martingale)**.** A *martingale* is a stochastic random process $\{X_t\}_{t\in\mathbb{N}}$ such that the conditional expectations exist and satisfy

$$\mathbb{E}(X_{t+1} \mid X_t = x_t, X_{t-1} = x_{t-1}, \ldots, X_0 = x_0) = x_t$$

for all $t \geq 0$. The X_t are real-valued random variables. The existence of the conditional expectation will be discussed later.

The equality above formalizes the previously given idea of a game where the expected winnings do not change after playing one more round. A more general definition of martingales is given later in this section; in particular the meaning of conditional expectation is explained then. Related to martingales are the two concepts of super- and submartingales whose definitions are obtained by relaxing the equality on the conditional expectation.

Definition 2.1.17 (Super- and submartingales)**.** A random process $\{X_i\}_{t\in\mathbb{N}}$ is called a *supermartingale* if it satisfies

$$\mathbb{E}(X_{t+1} \mid X_t = x_t, X_{t-1} = x_{t-1}, \ldots, X_0 = x_0) \leq x_t,$$

and a *submartingale* if it satisfies

$$\mathbb{E}(X_{t+1} \mid X_t = x_t, X_{t-1} = x_{t-1}, \ldots, X_0 = x_0) \geq x_t.$$

A martingale can then alternatively be defined to simultaneously be a supermartingale as well as a submartingale. While martingales describe, in a sense, a stable process, supermartingales describe a process whose value decreases over time. Consider for example a roulette game where the player constantly bets on the same color red and gets twice his bet on winning or loses it otherwise. If half of the roulette was red, this gambling process could be considered a martingale. However, casino roulettes also have a green field making strictly less than half of the roulette red. In expectation the player loses money every round while the casino makes money. The player process is thus a supermartingale and the casino process can be considered a submartingale.

The examples of martingales given so far were processes where a player repeatedly gambles until he either decides to stop playing or loses all his money. These points of time where the process is or gets stopped are called *stopping times*. Stopping times are crucial in the analysis of martingales as seen in the *Optional Stopping Theorem*. Informally, a stopping time is usually given by a stopping rule such as "*play until you run out of money*" or "*play* 100 *rounds or until you doubled your money*". However, these rules cannot be arbitrary and need to satisfy an essential property. The decision on whether to stop at a point t in time must be made solely from past and current information. For example, the rule "*stop as soon as you reach the maximum winning you will ever earn*" is not a stopping rule because it inquires information about the future of the process. To formally define these restrictions, the notions of *adapted* processes and *filtrations* are needed.

Remark. Some authors require a stopping rule to be finite almost surely, meaning that it will be satisfied after finite time with probability 1. We do not impose this restriction here.

Definition 2.1.18 (Filtered probability space)**.** Let $(\Omega, \mathcal{F}, \mu)$ be a probability space. A *filtration* is an increasing sequence $\mathbb{F} = \{\mathcal{F}\}_{t\in\mathbb{N}}$ of sub-σ-Algebras of $\mathcal{F}$. The tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mu)$ is then called a *filtered probability space*.

The idea of a filtration is to restrict the accessible information for random processes at any given time in such a way that the available information increases over time. If a random process $\{X_t\}_{t\in\mathbb{N}}$ at time t only depends on the information available in $\mathcal{F}_t$, that is, its value X_t can be determined given the information in $\mathcal{F}_t$, then the process is said to be *adapted* to the filtration $\{\mathcal{F}_t\}_{t\in\mathbb{N}}$.

Definition 2.1.19 (Adapted process)**.** Let $(\Omega, \mathcal{F}, \mathbb{F}, \mu)$ be a filtered probability space where $\mathbb{F} = \{\mathcal{F}\}_{t\in\mathbb{N}}$ is a filtration. A stochastic process $\{X_t\}_{t\in\mathbb{N}}$ is *adapted* to $\mathbb{F}$ iff:

$$\forall t \in \mathbb{N} : X_t \text{ is } \mathcal{F}_t\text{-measurable.}$$

The measurability condition expresses the fact that the value of X_t is determined only by information available at time t. Consider a repeated coin tossing process where X_t takes values in $\{H, T\}$ depending on the outcome of the t-th coin toss. Assume the coin is tossed a total of n times, then the associated sample space is $\Omega = \{H, T\}^n$ and the σ-Algebra is its power set $\mathcal{P}(\Omega)$. The probability measure does not matter for the example but we can assume a uniform distribution, meaning that every outcome is equally likely. Now we want to model the fact that at some time t we have only information about previous results. Consider the filtration defined by $\mathcal{F}_0 = \{\emptyset, \Omega\}$, $\mathcal{F}_t = \sigma(\{C_1 \times C_2 \times \cdots \times C_t \times \{H, T\}^{n-t} \mid C_i \in \{H, T\}\})$. The σ-Algebra $\mathcal{F}_t$ is generated by cylinder sets of bounded length t. This filtration captures the idea that at time t we only have information about the first t coin tosses. To see this, take some event $E \in \mathcal{F}_t$. E can be expressed in terms of the generating sets. For simplicity assume E is one of the generating sets, then it has the form $E = C_1C_2C_3 \dots C_t\{H, T\}^{n-t}$. E contains exact information about the first t outcomes but only has the trivial information that the last $n - t$ tosses are either head or tail. Also note that the sequence of σ-Algebras is indeed increasing because a generating set $C_1C_2 \dots C_{t-1}\{H, T\}^{n-(t-1)} \in \mathcal{F}_{t-1}$ of $\mathcal{F}_{t-1}$ can be expressed as the union $C_1C_2 \dots C_{t-1}H\{H, T\}^{n-t} \cup C_1C_2 \dots C_{t-1}T\{H, T\}^{n-t}$ of events in $\mathcal{F}_t$. Now consider the process $\{X_t\}_{t=1,\dots,n}$ mentioned above. To show that the process is indeed adapted to the filtration, we need to show measurability of the random variable X_t with respect to $\mathcal{F}_t$, i.e. we need to show that the preimages $X_t^{-1}(H)$ and $X_t^{-1}(T)$ are contained in $\mathcal{F}_t$. We have $X_t^{-1}(H) = \{C_1C_2 \dots C_n \in \{H, T\}^n \mid C_t = H\} = \bigcup\{C_1C_2 \dots C_{t-1}H\{H, T\}^{n-t} \mid C_i \in \{H, T\}\} \in \mathcal{F}_t$. Similarly, we have $X_t^{-1}(T) = \{C_1C_2 \dots C_n \in \{H, T\}^n \mid C_t = T\} = \bigcup\{C_1C_2 \dots C_{t-1}T\{H, T\}^{n-t} \mid C_i \in \{H, T\}\} \in \mathcal{F}_t$. The second part also follows from the facts that T is the complement of H, σ-algebras are closed under complementation and preimages preserve complements. Now we are ready to define stopping times formally.

Definition 2.1.20 (Stopping time)**.** Let $(\Omega, \mathcal{F}, \mathbb{F}, \mu)$ be a filtered probability space where $\mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathbb{N}}$ is a filtration. A *stopping time* $T : \Omega \to \mathbb{N} \cup \{\infty\}$ is a random variable such that the events $\{T = t\}$ are $\mathcal{F}_t$-measurable:

$$\forall t \in \mathbb{N} \cup \{\infty\} : \{T = t\} \in \mathcal{F}_t$$

where $\mathcal{F}_\infty \coloneqq \sigma(\bigcup_{t\in\mathbb{N}} \mathcal{F}_t)$. Alternatively, T is a stopping time if the process $\{X_t\}_{t\in\mathbb{N}}$ defined by

$$X_t = \begin{cases} 1, \text{ if } T \leq t \\ 0, \text{ otherwise} \end{cases}$$

is adapted to $\mathbb{F}$.

The notation $\{T = t\}$ here is a shorthand for the set $\{\omega \in \Omega \mid T(\omega) = t\}$. Again, the measurability constraint makes sure that the stopping rule only uses current and past information to decide whether to stop. Note that because $\mathcal{F}_t$ is a σ-Algebra, measurability of $\{T = t\}$ is equivalent to the measurability of $\{T \leq t\}$ or $\{T \geq t\}$.

Notable examples of stopping times include the termination time of a process or the first hitting time of some target state in a state-based system. Clearly, these can be related to termination of programs and reachability in programs respectively, making stopping times a valuable tool in analyzing probabilistic programs.

Stopping times have been introduced as some sort of rule to decide when to stop a given process like a gambling game, however, no formal definition on what stopping a process actually means is given thus far. We define it now.

Definition 2.1.21 (Stopped process)**.** Let $\{X_t\}_{t\in\mathbb{N}}$ be a random process adapted to a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\in\mathbb{N}}, \mu)$ and let T be a stopping time. The *stopped process* $\{X_t^T\}_{t\in\mathbb{N}}$ is defined by

$$X_t^T = X_{\min\{T,t\}}.$$

Some important properties of stopping times and stopped processes include

1. If $\{X_t\}_{t\in\mathbb{N}}$ is a martingale, then so is the stopped process $\{X_t^T\}_{t\in\mathbb{N}}$. Analogous results hold for super- and submartingales.

2. The minimum $T \wedge S$ of two stopping times T and S is again a stopping time. So are the maximum $S \vee T$ and the sum $S + T$. The difference $S - T$, however, is no stopping time.

3. The constant stopping time $k : \omega \mapsto k$ is a stopping time for all natural numbers $k \in \mathbb{N}$.

Theorem 2.1.4 (Optional stopping theorem[19])**.** *Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\in\mathbb{N}}, \mu)$ be a filtered probability space and let $\{X_t\}_{t\in\mathbb{N}}$ be an adapted martingale. If T is a stopping time and one of the three following conditions holds:*

(a) T is almost surely bounded, i.e. $T \leq c$ holds with probability 1 for some $c \in \mathbb{R}$.

(b) The expectation of T is finite ($\mathbb{E}(T) < \infty$) and the conditional expected change of the martingale is almost surely bounded, i.e. $\mathbb{E}(|X_{t+1} - X_t| \mid \mathcal{F}_t) \leq c$ holds almost surely for a constant c.[3]

(c) The stopped process X_t^T is almost surely bounded for all t, i.e. $|X_t^T| = |X_{\min\{t,T\}}| \leq c$ for some constant c holds almost surely.

Then the equality $\mathbb{E}(X_T) = \mathbb{E}(X_0)$ holds. If $\{X_t\}_{t\in\mathbb{N}}$ is a supermartingale instead then $\mathbb{E}(X_T) \leq \mathbb{E}(X_0)$ holds, and if it is a submartingale then $\mathbb{E}(X_T) \geq \mathbb{E}(X_0)$ holds instead.

The Optional Stopping Theorem states that, under any of the specified conditions, no matter the gambling strategy (stopping time), the gambler neither makes nor loses money on average in a martingale game. Let us reconsider the example of a gambler who plays a roulette game and repeatedly bets on red but doubles his bet every time he loses the bet. We already noticed that this game is a martingale if no green field is on the roulette, so for an appropriate stopping time, the above theorem can be applied. Let us look at every one of the three conditions in this context. The first condition states that the number of rounds the gambler can play before he stops (or is forced to stop) is bounded. Examples of such bounds are (1) the bounded life time of the gambler or (2) the bounded opening times of the casino. The second condition states that the gambler will eventually stop playing

[3] this notion of conditional expectation is introduced in the following.

and the casino has a limit on the allowed bets. A limit on bets would disallow the strategy of repeatedly doubling bets. The third condition states that the amount of money he can win or lose in total is bounded because for example either he goes broke and has to stop playing or the casino goes broke and can not pay him out anymore. In any of those three cases, the Optional Stopping Theorem says that the gambler will neither make nor lose any money in expectation. And if a green field is included on the roulette, the gambler will lose money while the casino will make money over time.

For the main part of thesis, we actually need a more general definition of martingales using the so called *conditional expectation*. We already made use of it in part (b) of the Optional Stopping Theorem.

Definition 2.1.22 (Conditional expectation)**.** Let $(\Omega, \mathcal{F}, \mu)$ be a probability space and $\mathcal{H} \subseteq \mathcal{F}$ a sub-σ-Algebra. The conditional expectation $\mathbb{E}(X \mid \mathcal{H})$ of a random variable X with respect to $\mathcal{H}$ is any $\mathcal{H}$-measurable function that satisfies almost surely

$$\int_H \mathbb{E}(X \mid \mathcal{H}) \mathrm{d}\mu = \int_H X \mathrm{d}\mu$$

for all $H \in \mathcal{H}$.

The definition is quite hard to understand, but it is worth noting that the conditional expectation itself is a random variable and not a single value. Intuitively, the conditional expectation $\mathbb{E}(X \mid \mathcal{H})$ gives the expected value of X given that we know all outcomings of events in $\mathcal{H}$. It can be thought of as having some certain fixed information about X and the expectation is only taken with respect to the remaining uncertain information. We give a slightly different and more easily to understand view on it later.

Usually, the conditional expectation assumes integrability of X so that the defining integrals are well-defined and the existence can be proven by Radon-Nikodym's Theorem. But this requirement is not necessary. It is for example possible to extend the above definition to arbitrary positive random variables by considering X as the limit of the sequence $(X \wedge n)$. Every $(X \wedge n)$ is integrable and hence its conditional expectation $\mathbb{E}(X \wedge n \mid \mathcal{H})$ exists. Using the Monotone Convergence Theorem, we get $\mathbb{E}(X \mid \mathcal{H})$ as the limit of the increasing sequence $\mathbb{E}(X \wedge n \mid \mathcal{H})$. This naturally extends to any lower bounded random variable X with lower bound $k \in \mathbb{R}$ by considering $(X + k)$ which is a nonnegative random variable. We do not go further into details but the conditional expectations we consider in the main section of this thesis are always well defined. Some important properties of conditional expectation include:

- $\mathbb{E}(X \mid \{\emptyset, \Omega\}) = \mathbb{E}(X)$. If no information (= the trivial sub-σ-Algebra) is given, then it coincides with the unconditional expectation.
- $\mathbb{E}(X \mid \mathcal{H}) = X$, if X is $\mathcal{H}$-measurable. Intuitively, the information in $\mathcal{H}$ is enough to determine the value of X completely; there is no uncertainty.
- (Linearity) $\mathbb{E}(\alpha X + \beta Y \mid \mathcal{H}) = \alpha \mathbb{E}(X \mid \mathcal{H}) + \beta \mathbb{E}(Y \mid \mathcal{H})$.
- (Monotonicity) For $X \leq Y$ we have $\mathbb{E}(X \mid \mathcal{H}) \leq \mathbb{E}(Y \mid \mathcal{H})$.

- (Law of total expectation) $\mathbb{E}(\mathbb{E}(X \mid \mathcal{H})) = \mathbb{E}(X)$.

The above equalities and inequalities are taken to hold almost surely. Most properties of the regular expectation are also satisfied by the conditional expectation. From the first property, the conditional expectation is easily seen to be a generalization of the regular expectation. The Monotone Convergence Theorem 2.1.1 does also hold for the conditional expectation.

Remark. The meaning of "almost sure equality" between two functions means that these functions are not necessarily equal everywhere, but the set of points where they differ is a null-set, i.e. a set of measure 0. Such functions, while not identical, are indistinguishable when integrated. Since conditional expectation is solely defined by its properties when integrated, there can exist a multitude of almost surely equal random variables that are candidates for the conditional expectation.

Definition 2.1.23 (General martingale with respect to filtration)**.** Let $\{Y_t\}_{t\in\mathbb{N}}$ be a stochastic process adapted to a filtration $\{\mathcal{F}_t\}_{t\in\mathbb{N}}$. $\{Y_t\}_{t\in\mathbb{N}}$ is a martingale if

$$\mathbb{E}(Y_{t+1} \mid \mathcal{F}_t) = Y_t$$

holds almost surely for all $t \in \mathbb{N}$. Super- and submartingales are defined similarly.

Definition 2.1.24 (General martingale with respect to another random process)**.** Let $\{Y_t\}_{t\in\mathbb{N}}$ and $\{X_t\}_{t\in\mathbb{N}}$ be stochastic processes. Y_t is a martingale with respect to X_t if for every $t \in \mathbb{N}$ the equality

$$\mathbb{E}(Y_{t+1} \mid X_t, X_{t-1}, \ldots, X_0) \coloneqq \mathbb{E}(Y_{t+1} \mid \mathcal{F}_t) = Y_t$$

holds almost surely, where $\mathcal{F}_t \coloneqq \sigma(X_t, X_{t-1}, \ldots, X_0)$ is the smallest σ-Algebra that makes $X_0, \ldots, X_t$ measurable. This naturally extends to super- and submartingales as before.

The conditional expectation with respect to another random process has an easy interpretation. $\mathbb{E}(Y_{t+1} \mid X_t, X_{t-1}, \ldots, X_0)$ is the expected value of Y_{t+1} given that we know the values of X_0 up to X_t. The expectation in this case can be thought of as a function that takes the values $x_0, \ldots, x_t$ as input and produces the value $\mathbb{E}(Y_{t+1} \mid X_t = x_t, X_{t-1} = x_{t-1} \ldots, X_0 = x_0)$, which is the expected value of Y at time $(t+1)$, given that we know the history of the process X till now. In Section 3.2 we make use of this particular definition and use the natural random process C_t which returns the system state visited at time t.

2.2 Order Theory

This section revises the fundamentals of order theory with a special focus on complete lattices and fixed points of their monotonous endofunctions. Usual problems in program verification include reachability (is a certain program state reachable?) and safety (does the program always stay in a safe region?). Both of these typical problems admit formulations via fixed points of certain operators. For example, consider a system which has some state space S. We are interested in the reachability region of that system starting from some initial state s_0. Generally, not all possible system states will be reachable so we are looking for a true subset $R \subset S$. Assume we have a single-step-operation $\mathbb{X}$ which computes from a set $A \subseteq S$ all the states $\mathbb{X}(A)$ which are reachable in a single step from A. We can then express the reachability region as a least fixed point of the following monotone operator $A \mapsto A \cup \mathbb{X}(A) \cup \{s_0\}$. In a similar manner, safety properties and invariants of a system or program can be described as greatest fixed points of appropriate operators.

The main topic of this thesis is martingale-based methods for program verification, however, these methods can also be derived from a purely order theoretic point of view. This gives two different viewpoints on the same methods and hence allows reasoning via two different frameworks. While most results can be derived either way, in many cases one of the viewpoints tends to be easier to reason with. The key theorems here are the well-known Knaster-Tarski theorem and the Cousot-Cousot theorem (also often known as Kleene fixed-point theorem). They crucially rely on the concepts of complete partial orders or complete lattices.

Definition 2.2.1 (Lattice)**.** A lattice is a tuple $(L, \leq)$ where

- L is a set,
- $\leq$ is a partial order on L, that is, a reflexive, transitive and antisymmetric relation on L,
- every finite subset $S \subseteq L$ has a supremum $\bigvee S$ (join) and an infimum $\bigwedge S$ (meet) in L.

If join and meet exist for arbitrary subsets of L, then the lattice is called *complete*. Complete lattices always have a least element $\bot = \bigwedge L$ (bottom) and a largest element $\top = \bigvee \emptyset$ (top).

An example of a lattice is the interval $[0, \infty)$ ordered by the usual ordering on the reals. The join is given by the maximum function, while the meet is given by the minimum function. The lattice has a bottom element $\bot = 0$ but no top element, hence it is not complete. If one extends the interval to $[0, \infty]$ by including positive infinity, the interval becomes a complete lattice where join and meet are given by supremum and infimum instead.

Definition 2.2.2 (Monotone function)**.** A function $f : (L, \leq) \to (L', \sqsubseteq)$ is called *monotone* if it preserves order:

$$a \leq b \implies f(a) \sqsubseteq f(b).$$

Monotonically increasing functions on $\mathbb{R}$ are then simply monotone functions on the lattice $(\mathbb{R}, \leq)$. Functions into a lattice $(L, \leq)$ can be ordered by extending the ordering $\leq$ point-wise: $f \sqsubseteq g \iff \forall x \in L : f(x) \leq g(x)$. Interestingly, the set of functions with point-wise ordering $\sqsubseteq$ then again forms a lattice. This function lattice is complete if the codomain lattice is complete.

Lemma 2.2.1. *The set of all functions $[X \to L]$ into a (complete) lattice $(L, \leq)$ is a (complete) lattice ordered by the partial order $f \sqsubseteq g \iff \forall x \in X : f(x) \leq g(x)$.*

Proof. It is straightforward to see that the ordering on the function lattice is indeed a partial order. The join and meet of functions is given by point-wise join and meet on their values. If the codomain is complete then arbitrary point-wise joins and meets exist on the function lattice, hence it is also complete. □

In Section 3.2 we are mainly interested in fixed points in a function lattice, that is, fixed points of monotone mappings on that function lattice, so the stated lemma proves very useful. A powerful theorem which characterizes certain fixed points of monotone endofunctions is due to Knaster and Tarski.

Theorem 2.2.2 (Knaster-Tarski[20])**.** *Let $f : L \to L$ be a monotone endofunction on a complete lattice, then f has a least and a greatest fixed point denoted by μf and νf respectively. Furthermore,*

1. *$\mu f = \inf\{x \in L \mid f(x) \leq x\} = \bigwedge_{f(x) \leq x} x$ is the least prefixed point, and*
2. *$\nu f = \sup\{x \in L \mid x \leq f(x)\} = \bigvee_{x \leq f(x)} x$ is the greatest postfixed point.*

Knaster-Tarski's theorem does not only guarantee the existence of least and greatest fixed points but also gives a way to compute them. Moreover, it provides the following reasoning principle which is fundamental to this thesis.

Corollary 2.2.2.1. *Let $f : L \to L$ be a monotone endofunction on a complete lattice, then*

1. *$f(x) \leq x \implies \mu f \leq x$, and*
2. *$x \leq f(x) \implies x \leq \nu f$.*

Proof. This follows immediately from Knaster-Tarski's theorem. μf is the least pre-fixpoint of f, hence every pre-fixpoint is greater than μf. Similarly, νf is the greatest post-fixpoint and thus every post-fixpoint is less than νf. □

Section 3.1 demonstrates how pre-fixpoints and post-fixpoints are related to super and submartingales. This connection allows the study of martingales in a fixed point theoretic manner. Finding the fixed points exactly is a difficult task so the reasoning principles derived from Knaster-Tarki's theorem give a useful method to at least approximate the desired fixed points. One should note that this reasoning only gives approximations in a single direction and it can not be used to underapproximate least fixed points or overapproximate greatest fixed points.

Another powerful theorem is by Cousot-Cousot. It characterizes the least and greatest fixed points as limits of possibly transfinite iterations. Cousot-Cousot's iteration procedure is essentially equivalent to the well known Kleene fixed point iteration, however, it generalizes Kleene's version by considering transfinite sequences, i.e. ones which are indexed by ordinals instead of usual sequences indexed by natural numbers.

Theorem 2.2.3 (Cousot-Cousot[21]). *Let $f : L \to L$ be a monotone endofunction on a complete lattice $(L, \leq)$. The supremum $\bigvee f^\alpha(\bot)$ of the increasing transfinite sequence $f^\alpha(\bot)$ defined by*

1. $f^0(\bot) = \bot$,
2. $f^{\alpha+1}(x) = f(f^\alpha(x))$,
3. $f^\alpha(x) = \bigvee_{\beta<\alpha} f^\beta(x)$ *if α is a limit ordinal*

is the least fixed point μf of f. Dually, starting from $\top$ and replacing the supremum in 3. by an infimum gives a decreasing sequence $f^\alpha(\top)$ whose infimum is the greatest fixed point νf.

While Cousot-Cousot's theorem provides a practical way to actually construct or at least approximate fixed points, it is mainly used as a reasoning tool in this thesis. The Markov decision processes induced by probabilistic programs in Section 3.1 generally have infinite state space which makes iteration methods impractical. Indeed, for the purpose of automated synthesis of martingales, Knaster-Tarski's theorem gives an effectively usable characterization of fixed points which is demonstrated in Chapter 4. Interesting is the asymmetry between Cousot-Cousot's and Knaster-Tarski's results. The iteration starting from $\bot$ increases to the least fixed point μf, hence underapproximating it. On the other hand, a pre-fixpoint of f gives an overapproximation of μf according to Knaster-Tarski.

Remark. Both theorems actually give stronger results than presented. However, the weaker versions presented here are sufficient for the purpose of this thesis.

3 Martingale-based Program Verification

This chapter is the main part of this thesis and deals with martingale-based methods in program verification for probabilistic programs. Although most methods introduced in this chapter are widely known in this field, we give a unified view via both probabilistic martingales as well as fixed point theoretic martingales. The close connection between both approaches is made apparent. Section 3.1 explains the notion of a probabilistic program with nondeterminism and defines its syntax and semantics formally. Section 3.2 develops the theory of martingales for probabilistic programs and addresses issues such as measurability. Concrete martingale-based methods are then explained in individual subsections.

3.1 Probabilistic programming

Probabilistic programs can model stochastic choices such as probabilistic branchings (e.g. a coin toss) or probabilistic assignments (e.g sampling from a distribution) in addition to what a standard deterministic program can do. In recent years incorporating stochastics into programming has become ever so popular. Applications include stochastic optimal control problems[18, 22], stochastic algorithms[23] and probabilistic machine learning[24]. In the first case, the probabilistic aspect allows to model disturbances on the system such as inaccuracies in the control or influences from outside. In the second and third case, incorporating probabilistic choices into algorithms can give better (expected) runtimes or other nice properties on the expense of being more hard to analyze and having less guarantees.

3.1.1 Programming Languages APP and PPP

This subsection introduces the two probabilistic programming languages APP and PPP. APP stands for *affine probabilistic programs* and PPP for *polynomial probabilistic programs*. Their difference lies in the type of assignment they are allowed to use. In APP every assignment and boolean expression is affine linear, i.e. has the form $x := \alpha x + \beta y + \cdots + c$ and $\alpha x + \beta y + \cdots + c \geq d$. In particular, polynomial expressions of degree at least 2 such as squares are not allowed. PPP in contrast allows arbitrary polynomials as expressions but cannot handle functions like e^x or trigonometric functions. While the theory of martingale-based verification can handle arbitrary measurable functions, the template-based synthesis introduced in Chapter 4 can only handle affine and polynomial functions.

Definition 3.1.1 (Affine and polynomial probabilistic programs)**.** An *affine linear program* (APP) is given by the following grammar.

⟨*stmt*⟩ ::= ⟨*assgn*⟩
| skip
| ⟨*stmt*⟩;⟨*stmt*⟩
| if ⟨*ndbexpr*⟩ then ⟨*stmt*⟩ else ⟨*stmt*⟩ fi
| while ⟨*bexpr*⟩ do ⟨*stmt*⟩ od

$\langle assgn \rangle$ ::= $\langle pvar \rangle$
| $\langle expr \rangle$
| $\langle pvar \rangle$:= $\langle dist \rangle$
| $\langle pvar \rangle$:= ndet($\langle dom \rangle$)

$\langle expr \rangle$ ::= $\langle constant \rangle$
| $\langle pvar \rangle$
| $\langle constant \rangle \cdot \langle pvar \rangle$
| $\langle expr \rangle$ + $\langle expr \rangle$
| $\langle expr \rangle$ - $\langle expr \rangle$

$\langle dom \rangle$ ::= Real
| Real[$\langle const \rangle$, $\langle const \rangle$]
| $\langle dom \rangle$ or $\langle dom \rangle$

$\langle bexpr \rangle$ ::= $\langle conjexpr \rangle$
| $\langle conjexpr \rangle$ or $\langle bexpr \rangle$

$\langle conjexpr \rangle$ ::= $\langle literal \rangle$
| $\langle literal \rangle$ and $\langle conjexpr \rangle$

$\langle literal \rangle$ ::= $\langle expr \rangle \leq \langle expr \rangle$
| $\langle expr \rangle \geq \langle expr \rangle$
| $\neg \langle literal \rangle$

$\langle ndbexpr \rangle$ ::= $\star$
| prob(p)
| $\langle bexpr \rangle$

$\langle pvar \rangle$ ::= $v \in V$

$\langle dist \rangle$ ::= $d \in \mathcal{D}(R)$

$\langle const \rangle$ ::= $c \in \mathbb{R}$

Here the set V is a finite set of variables and $\mathcal{D}(\mathbb{R})$ is the set of distributions over $\mathbb{R}$. The grammar for PPP is essentially the same but the multiplication rule is extended by multiplication of expressions

$\langle expr \rangle$ ::= $\langle expr \rangle \cdot \langle expr \rangle$.

The intended semantics for all purely deterministic program fragments is as usual. The non-deterministic branching simply takes one of the two branches non-deterministically (resolved by a scheduler, which is defined later). The probabilistic branching with parameter $p \in [0, 1]$ takes the first branch with probability p and the other one with probability $(1 - p)$. Probabilistic assignments sample from a given distribution. Non-deterministic assignments pick a value non-deterministically from the reals or from a given set which is the union of real intervals.

3.1.2 Semantics of probabilistic programs

The semantics of probabilistic programs are given by probabilistic control flow graphs (pCFGs). A pCFG is nothing but a Markov Decision Process (MDP) but defined in a way to more closely resemble the operational semantics of probabilistic programs.

Definition 3.1.2 (pCFG[11, 14])**.** A probabilistic control flow graph (pCFG) $\Gamma = (L, V, \mapsto, \mathrm{Up}, \Pr, G)$ consists of

- a finite set of *locations* $L = L_A + L_N + L_P + L_D$ partitioned into *assignment*, *nondeterministic*, *probabilistic* and *deterministic* locations,
- a finite set of *program variables* $V = \{x_1, \ldots, x_{|V|}\}$,
- a *total transition relation* $\mapsto \subseteq L \times L$ on L. Every location $l \in L$ has a non-empty set of successors $\mathrm{succ}(l) := \{l' \in L \mid l \mapsto l'\}$ with the additional condition that the successor is unique for assignment locations,
- an *update function* $\mathrm{Up} : L_A \to V \times \mathcal{U}$ where $\mathcal{U} = \mathcal{B}(\mathbb{R}^V, \mathbb{R}) \cup \mathcal{D}(\mathbb{R}) \cup \mathcal{B}(\mathbb{R})$. $\mathcal{U}$ is partitioned into three parts representing *deterministic*, *probabilistic* and *nondeterministic* assignments, respectively, and $L_A = L_{AD} + L_{AP} + L_{AN}$ can be partitioned accordingly.
- a function $\Pr : L_P \to \mathcal{D}(L)$ assigning to every probabilistic location $l \in L_P$ a distribution $\Pr(l)$ such that $\mathrm{supp}(\Pr(l)) := \{l' \in L \mid \Pr(l)(l') > 0\} \subseteq \mathrm{succ}(l)$. We sometimes write $\Pr_l$ instead of $\Pr(l)$.
- a *guard function* $G : L_D \times L \to \mathcal{B}(\mathbb{R}^V)$ assigning to each deterministic location $l \in L_D$ and target location $l' \in L$ a guard $G(l, l')$ such that $\{G(l, l')\}_{l' \in \mathrm{succ}(l)}$ is a partition of $\mathbb{R}^V$ for each $l \in L_D$. We write $\mathbf{x} \models G(l, l')$ instead of $\mathbf{x} \in G(l, l')$.

Before going into the concrete semantics, we first give a short explanation of the above items. The locations can be differentiated into four types where L_A describes assignments and L_N, L_P and L_D describe the three different branching types, i.e. nondeterministic, probabilistic and deterministic branching. Assignment locations have a unique successor and assign a value to a single program variable $x_j \in V$. They can be further divided into three types corresponding to deterministic assignment L_{AD} (given by a function in $\mathcal{B}(\mathbb{R}^V, \mathbb{R})$), probabilistic assignment L_{AP} (given by a distribution in $\mathcal{D}(\mathbb{R})$) and nondeterministic assignment L_{AN} (given by a set in $\mathcal{B}(\mathbb{R})$). Probabilistic branching is given by the function Pr and the additional condition enforces that only successors $l' \in \mathrm{succ}(l)$ of a location l can have positive probability. The guard function G assigns a subset $G(l, l') \in \mathcal{B}(R^V)$ of R^V to each possible transition $l \mapsto l'$. The transition is enabled if the current valuation $\mathbf{x}$ of the program variables is contained in $G(l, l')$. The condition on G makes sure that there is always exactly one transition enabled.

A configuration of a pCFG Γ is a tuple $c = (l, \mathbf{x}) \in L \times \mathbb{R}^V$ consisting of a program location l and a valuation of the program variables $\mathbf{x}$. A *run* of Γ is an infinite sequence $c_0 c_1 \cdots \in (L \times \mathbb{R}^V)^\omega$ of configurations such that for each pair of successive configurations $c_i = (l, \mathbf{x})$ and $c_{i+1} = (l', \mathbf{x}')$ we have

- if $l \in L_D$ then $\mathbf{x} = \mathbf{x}'$ and l' is the unique location satisfying $\mathbf{x} \models G(l, l')$,
- if $l \in L_P$ then $\mathbf{x} = \mathbf{x}'$ and $\Pr_l(l') > 0$,
- if $l \in L_N$ then $\mathbf{x} = \mathbf{x}'$ and $l \mapsto l'$,
- if $l \in L_A$ and $\mathrm{Up}(l) = (x_j, u)$ then l' is the unique successor $l \mapsto l'$ and
  - if $u = f \in \mathcal{B}(\mathbb{R}^V, \mathbb{R})$ then $\mathbf{x}' = \mathbf{x}(x_j \leftarrow f(\mathbf{x}))$,
  - if $u = d \in \mathcal{D}(\mathbb{R})$ then $\mathbf{x}' = \mathbf{x}(x_j \leftarrow a)$ where $a \in \mathrm{supp}(d)$,
  - if $u = A \in \mathcal{B}(\mathbb{R})$ then $\mathbf{x}' = \mathbf{x}(x_j \leftarrow a)$ where $a \in A$.

The support of a distribution d is defined as $\mathrm{supp}(d) := \{s \in \mathbb{R} \mid \forall N \subseteq \mathbb{R} : N \text{ is open } \wedge\ s \in N \implies d(N) > 0\}$. It is the set of all values s such that any open neighborhood containing s has positive measure.

The set of all runs is denoted by $\mathrm{Run}(\Gamma)$. A finite prefix $\pi = c_0 c_1 \dots c_k \in (L \times \mathbb{R}^V)^k$ of a run is called a *path*. The set of all paths is denoted by $\Pi = \mathrm{Paths}(\Gamma)$. We use a subscript to denote the paths which end in a specific type of configuration, e.g. $\Pi_{AN} \subseteq \Pi \times (L_{AN} \times \mathbb{R}^V)$ is the set of all path whose last location is a nondeterministic assignment.

The goal is to analyze the runtime behavior of the system, i.e. computing probabilities that runs exhibit a certain behavior such as being terminating, being recurrent or always being safe. In order to do so, one needs to associate a probability space to a pCFG Γ. The sample space Ω will simply be the set of all infinite sequences of configurations $(L \times \mathbb{R}^V)^\omega$. Restricting to only the runs is not necessary because we can define a probability measure ν on the sequences such that $\nu(\mathrm{Run}(\Gamma)) = 1$ and no set of non-valid runs has positive probability. Once a σ-Algebra Σ over $L \times \mathbb{R}^V$ is fixed, such as $\mathcal{B}(L \times \mathbb{R}^V)$, a natural candidate for the σ-Algebra over $(L \times \mathbb{R}^V)^\omega$ is the product σ-Algebra over infinitely many copies of Σ. And indeed, the product σ-Algebra will be used, however, for ease of analysis we actually enlarge it. The reason here is twofold. First, larger σ-Algebras make more functions measurable, thus potentially allowing for more properties to be analyzed. Secondly, Borel sets are missing some nice properties, e.g. the projection of a product Borel measurable set may fail to be Borel measurable[18, Appendix 3.B]. For these reasons, it is common to go over to universally measurable sets instead. So the σ-Algebra in use will be $\mathcal{U}((L \times \mathbb{R}^V)^\omega)$ which is obtained by first taking the countable product of Borel σ-Algebras $\mathcal{B}(L \times \mathbb{R}^V)$ and then forming its universal completion.

In order to define an appropriate probability measure, one first needs to resolve the nondeterminism of the system. For this purpose, the notion of a *scheduler* (also known as policy or strategy) is defined.

Definition 3.1.3 (Scheduler[11, 14])**.** Let $\Gamma = (L, V, \mapsto, \mathrm{Up}, \Pr, G)$ be a pCFG. A scheduler $\sigma = (\sigma_t, \sigma_a)$ is a pair of functions $\sigma_t : \Pi_N \to \mathcal{D}(L)$ and $\sigma_a : \Pi_{AN} \to \mathcal{D}(\mathbb{R})$ such that for any path $\pi = c_0 c_1 \dots c_k$ with $c_k = (l, \mathbf{x})$ we have

- if $l \in L_{AN}$ and $\mathrm{Up}(l) = (x_j, A)$ then $\mathrm{supp}(\sigma_a(\pi)) \subseteq A$,
- if $l \in L_N$ then $\mathrm{supp}(\sigma_t(\pi)) \subseteq \mathrm{succ}(l)$.

- Both σ_a and σ_t are universally measurable, i.e. they are universally measurable stochastic kernels on $\mathbb{R}$ and L given Π_{AN} and Π_N, respectively.

The set of all schedulers associated with Γ is denoted by Sch_Γ.

The component σ_a resolves nondeterministic assignments by defining a probability distribution on A, essentially replacing it by probabilistic assignments. Similarly, σ_t replaces nondeterministic branching by probabilistic branching. Schedulers are in general *history-dependent*, meaning they can resolve nondeterminism depending on the past behavior of the system. For any given history $\pi \in \Pi$ and scheduler σ, the scheduler σ_π with past history π is defined such that $\sigma_\pi(\pi') = \sigma(\pi\pi')$ holds.

We enlarge the domain of the schedulers to arbitrary finite, non-empty sequences $(L \times \mathbb{R}^V)^+$ such that for any non-valid path $\pi = c_0 \dots c_k$ we have $\sigma(\pi) = \delta_{c_k}$. That is, the scheduler will simply loop inside the current configuration.

The measurability condition on the scheduler is important. Consider a system with a single real valued variable x which takes values in the unit interval $[0,1]$ uniformly distributed. If no restriction on the scheduler is assumed, then the scheduler may resolve a nondeterministic binary branching depending on x such that branch (1) is taken if x is in a Vitali set A (or any other non-measurable set) and else branch (2). The probability that the system takes branch (1) is then exactly the probability that x is contained in A which is given by its Lebesgue-measure $\mu(A)$ (since x is uniformly chosen). This is not well-defined since A is not measurable. To prevent this, the measurability condition is needed.

Once a scheduler $\sigma \in \mathrm{Sch}_\Gamma$ and an initial configuration $c_0 \in L \times \mathbb{R}^V$ has been fixed, the system becomes a fully probabilistic system and its behavior can be intuitively described as follows. The system starts in c_0 and performs deterministic and probabilistic transitions as seen above. Whenever it reaches a nondeterministic location after a finite, non-empty sequence $\pi = c_0 c_1 \dots c_k$ of configurations, it invokes the scheduler σ via $\sigma(\pi)$ to determine the next step. For every such finite path π, one then can define a probability distribution $\mu^\sigma_\pi \in \mathcal{D}(L \times \mathbb{R}^V)$ over the configuration space, describing the probabilities to reach the successor configurations in the next step.

Definition 3.1.4 (Stochastic kernel $\mu^\sigma_\_$)**.** Let $\sigma \in \mathrm{Sch}_\Gamma$ be a scheduler for a pCFG Γ. For each non-empty, finite sequence $\pi = c_0 c_1 \dots c_{k-1}(l, \mathbf{x}) \in (L \times \mathbb{R}^V)^+$, define the probability measure $\mu^\sigma_\pi \in \mathcal{D}(L \times \mathbb{R}^V)$ such that

- if $l \in L_D$, then $\mu^\sigma_\pi = \delta_{(l', \mathbf{x})}$ where l' is the unique successor satisfying $(l, x) \models G(l, l')$,
- if $l \in L_P$, then $\mu^\sigma_\pi = \sum_{l \mapsto l'} \Pr_l(l') \delta_{(l', \mathbf{x})}$,
- if $l \in L_N$, then $\mu^\sigma_\pi = \sum_{l \mapsto l'} \sigma_t(\pi)(l') \delta_{(l', \mathbf{x})}$
- if $l \in L_A$, $l' = \mathrm{succ}(l)$ and $\mathrm{Up}(l, \mathbf{x}) = (x_j, u)$ then
  - if $l \in L_{AD}$ and $u = f \in \mathcal{B}(\mathbb{R}^V, \mathbb{R})$, then $\mu^\sigma_\pi = \delta_{(l', \mathbf{x}(x_j \leftarrow f(\mathbf{x})))}$,
  - if $l \in L_{AP}$ and $u = d \in \mathcal{D}(\mathbb{R})$, then μ^σ_π is the unique measure that satisfies

$$\forall A \in \mathcal{B}(\mathbb{R}) : \mu^\sigma_\pi(\{l'\} \times (\{\mathbf{x}_1\} \times \dots \times \{\mathbf{x}_{j-1}\} \times A \times \dots \times \{\mathbf{x}_{|V|}\})) = d(A).$$

- if $l \in L_{AN}$, then μ_π^σ is the unique measure that satisfies

$$\forall A \in \mathcal{B}(\mathbb{R}) : \mu_\pi^\sigma(\{l'\} \times (\{\mathbf{x}_1\} \times \cdots \times \{\mathbf{x}_{j-1}\} \times A \times \cdots \times \{\mathbf{x}_{|V|}\})) = \sigma_a(\pi)(A).$$

μ_-^σ is a universally measurable stochastic kernel on $L \times \mathbb{R}^V$ given $(L \times \mathbb{R}^V)^+$, that is, for every sequence $\pi \in (L \times \mathbb{R}^V)^+$ the function μ_π^σ is a probability measure on $\mathcal{U}(L \times \mathbb{R}^V)$ and the mapping $\mu_-^\sigma(U) : (L \times \mathbb{R}^V)^+ \to [0,1]$, $\pi \mapsto \mu_\pi^\sigma(U)$ is universally measurable for every fixed, universally measurable set U.

μ_π^σ is defined for arbitrary sequences of $\pi \in (L \times \mathbb{R}^V)^+$ which need not be actual paths of the pCFG Γ. For any non-empty finite sequence $\pi = \pi_1\pi_2$, the equality $\mu_{\pi_1\pi_2}^\sigma = \mu_{\pi_2}^{\sigma_{\pi_1}}$ holds. From this stochastic kernel and a fixed starting configuration $c_0 \in L \times \mathbb{R}^V$, one can construct a probability measure on $(L \times \mathbb{R}^V)^n$ as follows.

Lemma 3.1.1. *[18, Proposition 7.45] Let Γ be a pCFG, $\sigma \in \mathrm{Sch}_\Gamma$ be a scheduler and $c_0 \in L \times \mathbb{R}^V$ be a starting configuration, then there exists for all $n \in \mathbb{N} \setminus \{0\}$ a unique probability measure $\nu_n^{c_0,\sigma}$ on $\mathcal{B}((L \times \mathbb{R}^V)^n)$ such that*

$$\nu_n^{c_0,\sigma}(C_1 \times C_2 \times \cdots \times C_n) = \int_{C_1} \mu_{c_0}^\sigma(\mathrm{d}c_1) \int_{C_2} \mu_{c_0c_1}^\sigma(\mathrm{d}c_2) \cdots \int_{C_n} \mu_{c_0c_1\ldots c_{n-1}}^\sigma(\mathrm{d}c_n).$$

For $n = 0$ we can define $\nu_0^{c_0,\sigma} = \delta_{c_0}$. These probability measures extend uniquely to a probability measure $\nu^{c_0,\sigma}$ on $\mathcal{B}((L \times \mathbb{R}^V)^\omega)$ such that all its marginals on $(L \times \mathbb{R}^V)^n$ coincide with $\nu^{c_0,\sigma}n$.

Proof. We refer to Proposition 7.45 in Bertsekas' Stochastic Optimal Control book[18]. Note that the measures $\nu_n^{c_0,\sigma}$ induce a measure on the cylinder sets induced by the n-fold product of Borel sets C_i. This family of measures can then be extended uniquely by Kolmogorov's extension theorem to the measure $\nu^{c_0,\sigma}$. □

The probability measure $\nu^{c_0,\sigma}$ from the above lemma may be completed to be a measure on $\mathcal{U}((L \times \mathbb{R}^V)^\omega)$. From now on we work on different probability spaces in the form of $((L \times \mathbb{R}^V)^\omega, \mathcal{U}((L \times \mathbb{R}^V)^\omega), \nu^{c_0,\sigma})$ where the measure depends on a scheduler σ and a starting configuration c_0. Note that despite working on a multitude of different probability spaces, the underlying measurable space is always the same regardless of the chosen scheduler and starting configuration.

3.2 Application of Martingales

With the notion of probabilistic control flow graphs and their semantics at hand, we can now develop the theory of martingales for program verification. Concrete martingale-based methods are each introduced in their own subsection, but they all rely on the same principles and definitions we introduce in the following.

In the introduction we gave a short descriptive idea of so called ranking functions and a fixed point theoretic characterization of the reachability region. The former relied on the ability to compute the expected outcome of a ranking function after performing a transition in the system. The latter can be constructed by starting with a reachable set (e.g. the starting configuration) and iteratively enlarging it by the set of immediate successors until it reaches a fixed point. In either case, there is a necessity to make computations about the (expected) next step of the system. We formalize this in the form of so called *nexttime* operators $\mathbb{X}, \overline{\mathbb{X}}$ and $\underline{\mathbb{X}}$. For any given function η on the configuration space, these operators compute the expected value of η after performing a transition in the system. The nexttime operator $\mathbb{X}$ performs this transition with respect to a given scheduler, while the other two nexttime operators try to maximize or minimize the expected value over all possible available actions. Other authors call these operators *pre-expectation* (e.g. [13, 12]) but we use the naming convention by Takisaka et al.[11] which more clearly captures the idea of performing a single transition step in the transition system underlying the control flow graph.

Definition 3.2.1 (the "nexttime" operators $\mathbb{X}, \overline{\mathbb{X}}, \underline{\mathbb{X}}$, [14, 11])**.** Let Γ be a pCFG, and $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to \mathbb{K}$ be universally measurable for every $\sigma \in \mathrm{Sch}_\Gamma$ where $\mathbb{K}$ is a closed and convex proper subset of $\mathbb{R} \cup \{\pm\infty\}$. We define the *nexttime* operator $\mathbb{X}\eta$ as follows.

$$(\mathbb{X}\eta)(c, \sigma) = \int_{L \times \mathbb{R}^V} \eta(c', \sigma_c) \mu_c^\sigma(\mathrm{d}c')$$

Let $c = (l, \mathbf{x})$ then evaluating the above expression over all possible location types gives the following case-by-case definition.

- For $l \in L_D$, $(\mathbb{X}\eta)(c, \sigma) = \eta((l', \mathbf{x}), \sigma_c)$ where l' is the unique location s.t. $\mathbf{x} \models G(l, l')$.
- For $l \in L_P$, $(\mathbb{X}\eta)(c, \sigma) = \sum_{l \mapsto l'} \Pr_l(l') \eta((l', \mathbf{x}), \sigma_c)$.
- For $l \in L_N$, $(\mathbb{X}\eta)(c, \sigma) = \sum_{l \mapsto l'} \sigma_T(c)(l') \eta((l', \mathbf{x}), \sigma_c)$.
- For $l \in L_A$, let $\mathrm{Up}(l) = (x_j, u)$.
  - $(\mathbb{X}\eta)(c, \sigma) = \eta((\mathrm{succ}(l), \mathbf{x}(x_j \leftarrow f(\mathbf{x})), \sigma_c)$ if $u = f \in \mathcal{B}(L \times \mathbb{R}^V, \mathbb{K})$ is a measurable function.
  - $(\mathbb{X}\eta)(c, \sigma) = \int_{x \in \mathrm{supp}(s)} \eta((\mathrm{succ}(l), \mathbf{x}(x_j \leftarrow x)), \sigma_c) \mathrm{d}s$ if $u = s \in \mathcal{D}(L \times \mathbb{R}^V)$ is a distribution.
  - $(\mathbb{X}\eta)(c, \sigma) = \int_{x \in \mathrm{supp}(s)} \eta((\mathrm{succ}(l), \mathbf{x}(x_j \leftarrow x)), \sigma_c) \mathrm{d}s$ if $u = A \in \mathcal{B}(L \times \mathbb{R}^V)$ is a measurable set and $s = \sigma(c) \in \mathcal{D}(A)$ a distribution.

Let $\eta : L \times \mathbb{R}^V \to \mathbb{K}$ be universally measurable. We define the *upper nexttime* operator $\overline{\mathbb{X}}$ as follows.

$$(\overline{\mathbb{X}}\eta)(c) = \sup_{\sigma \in \mathrm{Sch}_\Gamma} \int_{L \times \mathbb{R}^V} \eta(c')\mu_c^\sigma(\mathrm{d}c')$$

Similarly to $\mathbb{X}$, the following case-by-case definition can be made, where $c = (l, \mathbf{x})$.

- For $l \in L_D$, $(\overline{\mathbb{X}}\eta)(c) = \eta(l', \mathbf{x})$ where l' is the unique location s.t. $\mathbf{x} \models G(l, l')$.
- For $l \in L_P$, $(\overline{\mathbb{X}}\eta)(c) = \sum_{l \mapsto l'} \Pr_l(l')\eta(l', \mathbf{x})$.
- For $l \in L_N$, $(\overline{\mathbb{X}}\eta)(c) = \max_{l \mapsto l'} \eta(l', \mathbf{x})$.
- For $l \in L_A$, let $\mathrm{Up}(l) = (x_j, u)$.
  - $(\overline{\mathbb{X}}\eta)(c) = \eta(\mathrm{succ}(l), \mathbf{x}(x_j \leftarrow f(\mathbf{x}))$ if $u = f \in \mathcal{B}(L \times \mathbb{R}^V, \mathbb{K})$ is a measurable function.
  - $(\overline{\mathbb{X}}\eta)(c) = \int_{x \in \mathrm{supp}(s)} \eta(\mathrm{succ}(l), \mathbf{x}(x_j \leftarrow x))\mathrm{d}s$ if $u = s \in \mathcal{D}(L \times \mathbb{R}^V)$ is a distribution.
  - $(\overline{\mathbb{X}}\eta)(c) = \sup_{x \in A} \eta(\mathrm{succ}(l), \mathbf{x}(x_j \leftarrow x))$ if $u = A \in \mathcal{B}(L \times \mathbb{R}^V)$ is a measurable set.

The *lower nexttime* operator $\underline{\mathbb{X}}\eta : L \times \mathbb{R}^V \to \mathbb{K}$ is defined as above, but replacing max with min and sup with inf in the respective lines.

These upper and lower nexttime operators are endofunctions on $\mathcal{U}(L \times \mathbb{R}^V, \mathbb{K})$ and the standard nexttime operator is an endofunction on $\mathcal{U}((L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma, \mathbb{K})$. They can be restricted to a pure invariant $I \in \mathcal{B}(L \times \mathbb{R}^V)$, i.e. to a transition closed subset of the configuration space. We can see $\overline{\mathbb{X}}$, $\underline{\mathbb{X}}$ and $\mathbb{X}$ as monotone functions on the complete lattices $(\mathcal{U}(I, \mathbb{K}), \sqsubseteq)$ and $(\mathcal{U}(I \times \mathrm{Sch}_\Gamma, \mathbb{K}), \sqsubseteq)$ where $\sqsubseteq$ is the partial ordering obtained by comparing functions pointwise with respect to the standard ordering $\leq$ on $\mathbb{K} \subset \mathbb{R} \cup \{\pm\infty\}$. These lattices are complete because the codomain $\mathbb{K}$ is closed and hence complete with respect to the standard ordering. Monotonicity is easily seen, since integrals are monotone operators. The restriction of the codomain makes sure that the integrals are well-defined, that is, they never take the indeterminate form $\infty - \infty$ because $\mathbb{K}$ is either lower bounded or upper bounded.

Remark. For scheduler dependent functions $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to L$ into a complete lattice L, we sometimes write the parameter σ in the subscript, i.e. we write $\eta_\sigma(c)$ instead of $\eta(c, \sigma)$. Furthermore, we implicitly define the *upper* and *lower* versions of the function by taking the supremum and the infimum over all schedulers. These are then denoted by $\overline{\eta}(c) \coloneqq \sup_\sigma \eta_\sigma(c)$ and $\underline{\eta}(c) \coloneqq \inf_\sigma \eta_\sigma(c)$, respectively.

If η_σ is a function that assigns to each configuration $c \in L \times \mathbb{R}^V$ of the system a real value $\eta_\sigma(c) \in \mathbb{R}$, then the nexttime operators describe the expected value of η after performing a single transition in the system, e.g. $\mathbb{X}\eta(c, \sigma) = \mathbb{E}(\eta(c', \sigma_c) \mid c \mapsto c', \sigma)$ is the expected value of $\eta(c', \sigma_c)$ where c' is the successor of c that gets (potentially randomly) chosen under the scheduler σ. To properly define the expectation, a scheduler and a starting configuration

c_0 has to be given (because the underlying measure ν is dependent on them). We use the notation $\mathbb{E}_{c_0,\sigma}$ to make the dependencies clear but may drop one or both subscripts if they are clear from context. Furthermore, we define the upper expectation $\overline{\mathbb{E}}_{c_0} \coloneqq \sup_\sigma \mathbb{E}_{c_0,\sigma}$ and the lower expectation $\underline{\mathbb{E}}_{c_0} \coloneqq \inf_\sigma \mathbb{E}_{c_0,\sigma}$. The upper and lower nexttime operators then satisfy $\overline{\mathbb{X}}\eta(c_0) = \overline{\mathbb{E}}_{c_0}(\eta(c_1) \mid c_0 \mapsto c_1)$ and $\underline{\mathbb{X}}\eta(c_0) = \underline{\mathbb{E}}_{c_0}(\eta(c_1) \mid c_0 \mapsto c_1)$, respectively. These various types of expectation give rise to various martingale notions.

Definition 3.2.2 (σ-martingales)**.** A stochastic process $\{X\}_{i\in\mathbb{N}}$ where $X_i : (L \times \mathbb{R}^V)^\omega \to \mathbb{K}$ for some closed, convex proper subset $\mathbb{K} \subset \mathbb{R} \cup \{\pm\infty\}$ is called a σ-martingale at $c_0 \in L \times \mathbb{R}^V$ if it satisfies

$$\mathbb{E}_{c_0,\sigma}(X_{i+1} \mid C_i, C_{i-1}, \ldots, C_0) = X_i$$

for a fixed scheduler σ, where $C_i : (L \times \mathbb{R}^V)^\omega \to L \times \mathbb{R}^V$, $c_0c_1 \ldots c_i \cdots \mapsto c_i$ is the canonical process. The subscript c_0 from the expectation is usually dropped since $C_0 = c_0$ is given in the conditions. Similarly, it is an upper martingale if

$$\overline{\mathbb{E}}(X_{i+1} \mid C_i, C_{i-1}, \ldots, C_0) = X_i$$

and a lower martingale if

$$\underline{\mathbb{E}}(X_{i+1} \mid C_i, C_{i-1}, \ldots, C_0) = X_i$$

holds. Super and supmartingales are defined accordingly. The restriction to a closed, convex and proper subset of the extended reals makes sure that the conditional expectation is well-defined.

A reason why we deal with universally measurable functions instead of Borel measurable ones comes from the supremizing and infimizing upper and lower nexttime operators. While the regular nexttime operator preserves Borel measurability if we only consider Borel schedulers, the extremizing ones do not. A well known fact about Borel measurable sets is that they are not closed under projection, i.e. projecting a product measurable set $D \in \mathcal{B}(\mathbb{R})\mathcal{B}(\mathbb{R})$ onto its first component may yield a set that fails to be Borel measurable. However, it is known that such a projection is universally measurable (actually analytic, and hence universally measurable[18, Corollary 7.42.1]). Extremization and projections are closely related, in fact, if $f : X \times A \to \overline{\mathbb{R}}$ is a function, then for any $c \in \overline{\mathbb{R}}$ we have

$$\{x \in X \mid \inf_{a\in A} f(x,a) < c\} = \mathrm{proj}_X\{(x,a) \in X \times A \mid f(x,a) < c\}.$$

If the function f is Borel measurable, then so is the set $\{(x,a) \in X \times A \mid f(x,a) < c\}$ defined by f, but its projection is generally non-Borel. The left-hand side set defined by the infimization thus may fail to be Borel, making the infimization itself a non-Borel measurable function. However, it is universally measurable. This shows the reason why we work on universally measurable functions instead of Borel measurable ones.

Remark. An analytic set is defined to be the continuous image of a Borel set of a Polish space (a separable, completely metrizable topological space). All spaces we consider in this thesis are Polish spaces. It may seem like it is sufficient to work on analytic sets instead of universally measurable ones, but the analytic sets have two problems. First, they do

not form a σ-Algebra by themselves since they are not closed under complementation. Secondly, if one computes the smallest σ-Algebra containing them, and then defines the appropriate notion of *analytically measurable* function, it turns out that these functions are not closed under composition. The universal σ-Algebra, however, contains all analytic sets and its measurable functions are closed under composition, making it a reasonable choice for our analysis. There does actually exist a smaller σ-Algebra called the *limit σ-Algebra*[18, Appendix] that does satisfy these two requirements as well but it is not widely used.

An example of such a function is the reachability probability $\mathbb{P}^{\text{reach}}_{C,\sigma}(c)$ which gives the probability to reach a region $C \in \mathcal{B}(L \times \mathbb{R}^V)$ when starting from configuration $c \in L \times \mathbb{R}^V$ under scheduler σ. Using LTL, this can be stated as $\mathbb{P}^{\text{reach}}_{C,\sigma}(c) = \Pr_\sigma(c \models \Diamond C)$. The nexttime operator $\mathbb{X}$ then corresponds to the LTL next operator $\bigcirc$, e.g. $\mathbb{X}\mathbb{P}^{\text{reach}}_{C,\sigma}(c) = \Pr_\sigma(c \models \bigcirc\Diamond C)$. While $\mathbb{P}^{\text{reach}}_{C,\sigma}$ is Borel measurable if σ is a Borel measurable scheduler, the infimization $\underline{\mathbb{P}}^{\text{reach}}_{C}$ is generally not Borel measurable. It can be shown to be lower semianalytic and hence universally measurable[18].

Lemma 3.2.1 (Continuity of $\overline{\mathbb{X}}$ and $\underline{\mathbb{X}}$)**.** *The operators $\overline{\mathbb{X}}$ and $\underline{\mathbb{X}}$ are ω-continuous and ω^{op}-continuous, respectively. This means that $\overline{\mathbb{X}}$ preserves suprema of ascending ω-chains, while $\underline{\mathbb{X}}$ preserves infima of descending ω-chains.*

Proof. We prove the ω-continuity of $\overline{\mathbb{X}}$. Let $\eta_0 \sqsubseteq \eta_1 \sqsubseteq \dots$ be an ascending chain and let $\eta = \sup_{i\in\mathbb{N}} \eta_i$ be its supremum. Then,

$$\begin{aligned}
\sup_{i\in\mathbb{N}} \overline{\mathbb{X}}(\eta_i)(c) &= \sup_{i\in\mathbb{N}} \sup_{\sigma} \int_{L\times\mathbb{R}^V} \eta_i(c')\mu^\sigma_c(\mathrm{d}c') \\
&= \sup_{\sigma} \sup_{i\in\mathbb{N}} \int_{L\times\mathbb{R}^V} \eta_i(c')\mu^\sigma_c(\mathrm{d}c') \\
&= \sup_{\sigma} \int_{L\times\mathbb{R}^V} \sup_{i\in\mathbb{N}} \eta_i(c')\mu^\sigma_c(\mathrm{d}c') \\
&= \sup_{\sigma} \int_{L\times\mathbb{R}^V} \eta(c')\mu^\sigma_c(\mathrm{d}c') = \overline{\mathbb{X}}(\eta)(c)
\end{aligned}$$

In the second line we exchange the order of suprema. In the second to last line, we use the Monotone Convergence Theorem. The proof for $\underline{\mathbb{X}}$ and descending chains is similar, but we exchange two infima instead. □

The nexttime operator $\mathbb{X}$ is both ω-continuous as well as ω^{op}continuous which follows directly from the Monotone Convergence Theorem.

A naturally arising question is whether the equalities $\overline{\mathbb{X}}\overline{\eta} = \overline{\mathbb{X}\eta}$ and $\underline{\mathbb{X}}\underline{\eta} = \underline{\mathbb{X}\eta}$ hold. This is generally not the case, but the following inequalities $\overline{\mathbb{X}}\overline{\eta} \sqsupseteq \overline{\mathbb{X}\eta}$ and $\underline{\mathbb{X}}\underline{\eta} \sqsubseteq \underline{\mathbb{X}\eta}$ are easily seen to be true. For equality to hold, the existence of so called ϵ-optimal schedulers is needed.

Definition 3.2.3 (ϵ-optimal schedulers)**.** Let Γ be a pCFG and let $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to \overline{\mathbb{R}}$ be a scheduler-dependent function. A scheduler σ' is called *upper ϵ-optimal* if for all $c \in (L \times \mathbb{R}^V)$ it holds that $\eta(c,\sigma') \geq \overline{\eta}(c) - \epsilon$ when $\overline{\eta}(c) < +\infty$ and $\eta(c,\sigma') > \frac{1}{\epsilon}$ when

$\overline{\eta}(c) = +\infty$. Similarly, a scheduler σ' is called *lower ϵ-optimal* if $\eta(c,\sigma') \leq \underline{\eta}(c) + \epsilon$ when $\underline{\eta}(c) > -\infty$ and $\eta(c,\sigma') < -\frac{1}{\epsilon}$ when $\underline{\eta}(c) = -\infty$.

By definition of $\overline{\eta}$ there always exists a scheduler σ^c for every $c \in L \times \mathbb{R}^V$ such that σ^c is ϵ-optimal at c. In contrast to this, an ϵ-optimal scheduler is ϵ-optimal for every input $c \in L \times \mathbb{R}^V$ simultaneously. Such schedulers may not exist in general.

Lemma 3.2.2. *Let $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to \mathbb{K}$ be a scheduler dependent function where $\mathbb{K}$ is a closed and convex proper subset of $\overline{\mathbb{R}}$. If there exist upper ϵ-optimal schedulers for all $\epsilon > 0$, then the equality $\overline{\mathbb{X}}\overline{\eta} = \overline{\mathbb{X}\eta}$ holds. Analogously, if lower ϵ-schedulers exists, then $\underline{\mathbb{X}}\underline{\eta} = \underline{\mathbb{X}\eta}$ holds.*

Proof. We prove the case of upper ϵ-optimal schedulers. The other case is analogous. The direction $\overline{\mathbb{X}\eta} \sqsubseteq \overline{\mathbb{X}}\overline{\eta}$ is easily seen to be true. For the other direction, let $(\rho^n)_{n\in\mathbb{N}}$ be a sequence of $\frac{1}{n}$-optimal schedulers such that $\eta(\cdot,\rho^k) \sqsubseteq \eta(\cdot,\rho^n)$ for $k \leq n$. For any scheduler σ, define schedulers σ^n such that $\sigma^n(c) = \sigma(c)$, and $\sigma^n(c\pi) = \rho^n(\pi)$. We have $\mu_c^\sigma = \mu_c^{\sigma^n}$ (see definition 3.1.4 of μ_π^σ) and $\sigma_c^n = \rho^n$.

$$
\begin{aligned}
\overline{\mathbb{X}}\overline{\eta}(c) &= \sup_\sigma \int_{L\times\mathbb{R}^V} \sup_\rho \eta(c',\rho)\mu_c^\sigma(\mathrm{d}c') \\
&= \sup_\sigma \int_{L\times\mathbb{R}^V} \lim_{n\in\mathbb{N}} \eta(c',\rho^n)\mu_c^\sigma(\mathrm{d}c') \\
&= \sup_\sigma \lim_{n\in\mathbb{N}} \int_{L\times\mathbb{R}^V} \eta(c',\rho^n)\mu_c^\sigma(\mathrm{d}c') \\
&= \sup_\sigma \lim_{n\in\mathbb{N}} \int_{L\times\mathbb{R}^V} \eta(c',\sigma_c^n)\mu_c^{\sigma^n}(\mathrm{d}c') \\
&\leq \sup_\sigma \int_{L\times\mathbb{R}^V} \eta(c',\sigma_c)\mu_c^\sigma(\mathrm{d}c') \\
&= \sup_\sigma \mathbb{X}\eta(c,\sigma) = \overline{\mathbb{X}\eta}(c)
\end{aligned}
$$

From the second to third line, we use the Monotone Convergence Theorem. We then use the properties of σ^n in line four. □

In the above proof, the existence of ϵ-optimal schedulers is important. If this is not given, then in line two one has maximizing scheduler sequences of the form $\rho^{n,c'}$ which specifically depend on the successor c' of c. Despite that, it is still possible to define the combined scheduler σ^n as above but with a slight modification such that $\sigma^n(cc'\pi) = \rho^{n,c'}(c'\pi)$. This scheduler would simultaneously maximize the integrand for all c' and the integral itself to yield the desired result. However, the schedulers σ^n defined like this are a combination of potentially uncountably many schedulers (since there can be uncountably many successors c') and may fail to be universally measurable.

To any measurable function $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to \mathbb{K}$ and scheduler σ we can naturally associate a stochastic process as follows.

Definition 3.2.4 (Stochastic process Y_i)**.** Let Γ be a pCFG, $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to \mathbb{K}$ be a universally measurable function and $\sigma \in \mathrm{Sch}_\Gamma$ be a scheduler. The stochastic process $\{Y_i\}_{i\in\mathbb{N}}$ of random variables $Y_i : (L \times \mathbb{R}^V)^\omega \to \mathbb{K}$ is defined by

$$Y_i(c_0c_1c_2\dots) = \eta(c_i, \sigma_{c_0c_1\dots c_{i-1}}).$$

Similarly, the upper and lower versions are defined as

$$\overline{Y}_i(c_0c_1c_2\dots) = \overline{\eta}(c_i),$$

and

$$\underline{Y}_i(c_0c_1c_2\dots) = \underline{\eta}(c_i).$$

We analyze the process $\{Y_i\}_{i\in\mathbb{N}}$ with respect to the martingale conditions and get the following result.

Lemma 3.2.3. *The stochastic process $\{Y_i\}_{i\in\mathbb{N}}$ is a σ-martingale (with respect to the canonical process $C_i : c_0c_1\cdots \mapsto c_i$), iff the equality $\eta(c,\sigma_\pi) = \mathbb{X}\eta(c,\sigma_\pi)$ holds for all paths $\pi c = c_0\dots c$. It is a σ-supermartingale iff $\eta(c,\sigma_\pi) \sqsupseteq \mathbb{X}\eta(c,\sigma_\pi)$, and a σ-submartingale iff $\eta(c,\sigma_\pi) \sqsubseteq \mathbb{X}\eta(c,\sigma_\pi)$ holds. Similarly, $\{\overline{Y}_i\}_{i\in\mathbb{N}}$ is an upper martingale iff the respective equalities and inequalities hold with $\overline{\eta}$ and $\overline{\mathbb{X}}$ replacing η and $\mathbb{X}$. An analogous results holds for $\{\underline{Y}_i\}_{i\in\mathbb{N}}$ as well.*

Proof. We prove the case when $\eta(c,\sigma_\pi) = \mathbb{X}\eta(c,\sigma_\pi)$ holds.

$$\begin{aligned}
\mathbb{E}_\sigma(Y_{i+1} \mid C_i, C_{i-1}, \dots, C_0) &= \mathbb{E}_\sigma(\eta(C_{i+1}, \sigma_{C_0C_1\dots C_i}) \mid C_i, C_{i-1}, \dots, C_0) \\
&= \int_{L\times\mathbb{R}^V} \eta(C_{i+1}, \sigma_{C_0C_1\dots C_i}) \mu^\sigma_{C_0C_1\dots C_i}(\mathrm{d}C_{i+1}) \\
&= \int_{L\times\mathbb{R}^V} \eta(C_{i+1}, \sigma_{C_0C_1\dots C_i}) \mu^{\sigma_{C_0C_1\dots C_{i-1}}}_{C_i}(\mathrm{d}C_{i+1}) \\
&= \mathbb{X}\eta(C_i, \sigma_{C_0C_1\dots C_{i-1}}) = \eta(C_i, \sigma_{C_0C_1\dots C_{i-1}}) = Y_i
\end{aligned}$$

For the super- and submartingale cases the second to last equality simply becomes the respective inequality. The proof for the upper and lower case are almost identical. We prove the case of $\overline{\eta}(c) = \overline{\mathbb{X}}\overline{\eta}(c)$.

$$\begin{aligned}
\overline{\mathbb{E}}(\overline{Y}_{i+1} \mid C_i, C_{i-1}, \dots, C_0) &= \overline{\mathbb{E}}(\overline{\eta}(C_{i+1}) \mid C_i, C_{i-1}, \dots, C_0) \\
&= \sup_\sigma \int_{L\times\mathbb{R}^V} \overline{\eta}(C_{i+1}) \mu^\sigma_{C_0C_1\dots C_i}(\mathrm{d}C_{i+1}) \\
&= \sup_\sigma \int_{L\times\mathbb{R}^V} \overline{\eta}(C_{i+1}) \mu^{\sigma_{C_0C_1\dots C_{i-1}}}_{C_i}(\mathrm{d}C_{i+1}) \\
&= \sup_\sigma \int_{L\times\mathbb{R}^V} \overline{\eta}(C_{i+1}) \mu^\sigma_{C_i}(\mathrm{d}C_{i+1}) \\
&= \overline{\mathbb{X}}\overline{\eta}(C_i) = \overline{\eta}(C_i) = \overline{Y}_i
\end{aligned}$$

□

From now on we simply call η a (upper/lower) martingale or a super/submartingale if it satisfies the conditions in the above Lemma 3.2.3. That is, if the induced process $\{Y_i\}_{i\in\mathbb{N}}$ has the respective martingale properties. This notion extends to scheduler-independent functions $\eta : L \times \mathbb{R}^V \to \mathbb{K}$ by simply defining $\eta'(c,\sigma) = \eta(c)$ for all schedulers σ and then considering η' instead. By abuse of notation we simply write $\eta(c,\sigma) = \eta(c)$ for such functions and do not consider η' explicitly.

3.2.1 Additive Ranking Supermartingales

Ranking supermartingales are probably the most well-known supermartingales used in the analysis of probabilistic systems. A ranking supermartingale η assigns real values to each configuration of the system such that two properties hold;

- It is nonnegative on all (reachable) configurations.
- Its expected value after a single transition decreases by at least some fixed value $\epsilon > 0$ outside of some target region C.

If such a function on the configuration space exists, then it witnesses the fact that the system will eventually reach the region C after finitely many steps, i.e. the system reaches C positive almost surely. The intuitive reason being, that from any starting configuration c_0, the value of η decreases in expectation by at least ϵ per step and after $\eta(c_0)/\epsilon$ many steps, the expectation reaches zero. When it reaches zero, the function can no longer decrease because of the nonnegativity condition, hence it must have entered the target region C. Not only does it witness positive a.s. reachability but it also gives upper bounds on the expected number of steps, namely $\eta(c_0)/\epsilon \geq$ "expected reaching time of C".

Instead of having an additive decrease per step, one can also have a multiplicative decrease instead. The positivity condition then gets replaced by the condition $\eta(c) \geq 1$ outside of the target region. This gives a so called *multiplicative* ranking supermartingale. The idea remains the same and, in fact, by using the exponential function one can transform both types of martingales into each other. So our focus here will be on the additive version.

Definition 3.2.5 (Additive ranking supermartingle)**.** Let Γ be a pCFG, $c_0 \in L \times \mathbb{R}^V$ a starting configuration, $I = I_{c_0} \in \mathcal{B}(L \times \mathbb{R}^V)$ an invariant given c_0 (i.e. all runs starting from c_0 reach only configurations in I) and $C \in \mathcal{B}(I)$ be a target region. Let $\sigma \in \mathrm{Sch}_\Gamma$ be a scheduler, then a universally measurable function $\eta : I \times \mathrm{Sch}_\Gamma \to [0,\infty]$ is called an *additive ranking σ-supermartingale* (σ-ARnkSupM) at c_0 for C supported by I if it satisfies

$$\eta(c,\sigma_\pi) \geq 1 + \mathbb{X}\eta(c,\sigma_\pi) \text{ for all } c \in I \setminus C \text{ and all paths } \pi c = c_0 \dots c.$$

Similarly, a universally measurable function $\eta : I \to [0,\infty]$ is called an *upper additive ranking supermartingale* (UARnkSupM) if

$$\eta(c) \geq 1 + \overline{\mathbb{X}}\eta(c) \text{ for all c } \in I \setminus C \text{ and all paths } \pi c = c_0 \dots c$$

holds. A *lower additive ranking supermartingale* (LARnkSupM) is defined analogously by replacing the upper nexttime operator by the lower nexttime operator.

There are variations of this definition. For example, instead of forcing a decrease of at least 1 per transition, one can instead assume a decrease of some $\epsilon > 0$. By rescaling, both definitions are equivalent. Another commonly used formulation is using the drift operator $\Delta\eta = \eta - \mathbb{X}\eta$ and then formulate the ranking supermartingale condition to be a negative drift, e.g. $\Delta\eta(c, \sigma_\pi) \leq -\epsilon$ ([6] uses this kind of formulation).

Despite its name, the associated stochastic process Y_i from Lemma 3.2.3 is no supermartingale. The problem is that an ARnkSupM η may increase in value (and in expectation) inside the target region since no conditions are enforced there. However, outside the target region the process is decreasing in expectation, hence, we can define an appropriate stopping time T^C such that the stopped process is indeed a supermartingale.

Lemma 3.2.4. *Let η be a σ-ARnkSupM for some scheduler σ and target region C. Let $T = T^C : (L \times \mathbb{R}^V)^\omega \to [0, \infty]$ be the first hitting time of C, i.e. $T(c_0 c_1 \dots) = \inf\{i \in \mathbb{N} \mid c_i \in C\}$. The first hitting time is indeed a stopping time according to Definition 2.1.20. Then the stopped process $\{Y_i^T\}_{i \in \mathbb{N}}$ is a σ-supermartingale with respect to the canonical process C_i.*

Proof. The stopped process is defined by $Y_i^T = Y_{\min\{i,T\}}$. We need to show that $\mathbb{E}_\sigma(Y_{i+1}^T \mid C_i, \dots, C_0) = \mathbb{X}\eta(C_i, \sigma_{C_0 C_1 \dots C_{i-1}}) \leq Y_i^T$ holds. We do a case distinction. Assume $i < T$, then $C_i \notin C$ and so

$$\begin{aligned}
\mathbb{E}_\sigma(Y_{i+1}^T \mid C_i, \dots, C_0) &= \mathbb{E}_\sigma(Y_{i+1} \mid C_i, \dots, C_0) \\
&= \mathbb{X}\eta(C_i, \sigma_{C_0 C_1 \dots C_{i-1}}) \\
&= 1 + \mathbb{X}\eta(C_i, \sigma_{C_0 C_1 \dots C_{i-1}}) - 1 \\
&\leq \eta(C_i, \sigma_{C_0 C_1 \dots C_{i-1}}) - 1 = Y_i^T - 1 \leq Y_i^T.
\end{aligned}$$

For the case $i \geq T$, we have

$$\begin{aligned}
\mathbb{E}_\sigma(Y_{i+1}^T \mid C_i, \dots, C_0) &= \mathbb{E}_\sigma(Y_T \mid C_i, \dots, C_T, \dots, C_0) \\
&= Y_T = Y_{\min\{i,T\}} = Y_i^T.
\end{aligned}$$

We use the fact that Y_T is measurable with respect to the condition, since its value is determined by the values of C_0 up to C_T.

Note that the case distinction is only valid, because the information whether $i < T$ or $i \geq T$ holds is contained in the given information $C_0, \dots, C_i$. Formally, this can be expressed as the events $\{i < T\}$ and $\{i \geq T\}$ being measurable with respect to the σ-Algebra induced by the canonical process $C_0, \dots, C_i$. □

The idea of this supermartingale is as follows. The value of η decreases by one in expectation each step until the target region C is reached. When reaching the target region after expected $\mathbb{E}_{c_0,\sigma}(T)$ many steps, then the function η will have decreased by $\mathbb{E}_{c_0,\sigma}(T)$ points in expectation. Since η is nonnegative, we conclude that the starting value $\eta(c_0, \sigma)$ must have been greater than $\mathbb{E}_{c_0,\sigma}(T)$. So $\eta(c_0, \sigma)$ gives an upper bound on the reaching time $\mathbb{E}_{c_0,\sigma}(T)$, and if $\eta(c_0, \sigma)$ is finite then so is the expected reaching time and hence the target region is reached within finite time with probability 1. We prove this idea now.

Applying the Optional Stopping Theorem (Theorem 2.1.4) directly to Y_i^T (by using the constant stopping time i) does not give any useful result. The problem is that the information on how much Y_i^T decreases as i gets larger is not used at all. To make use of that property, we define the following stochastic process $Z_i = Y_i^T + \min\{i, T\}$ and show that this is still a σ-supermartingale. Consider again the case $i < T$, then

$$\begin{aligned}
\mathbb{E}_\sigma(Z_{i+1} \mid C_i, \ldots, C_0) &= \mathbb{E}_\sigma(Y_{i+1} + (i+1) \mid C_i, \ldots, C_0) \\
&= \mathbb{E}_\sigma(Y_{i+1} \mid C_i, \ldots, C_0) + (i+1) \\
&= \mathbb{X}\eta(C_i, \sigma_{C_0 C_1 \ldots C_{i-1}}) + (i+1) \\
&= 1 + \mathbb{X}\eta(C_i, \sigma_{C_0 C_1 \ldots C_{i-1}}) + i \\
&\leq \eta(C_i, \sigma_{C_0 C_1 \ldots C_{i-1}}) + i \\
&= Y_i^T + i = Z_i
\end{aligned}$$

and for $i \geq T$ we have

$$\begin{aligned}
\mathbb{E}_\sigma(Z_{i+1}^T \mid C_i, \ldots, C_0) &= \mathbb{E}_\sigma(Y_T + T \mid C_i, \ldots, C_T, \ldots, C_0) \\
&= Y_T + T = Y_{\min\{i,T\}} + \min\{i, T\} = Z_i.
\end{aligned}$$

Using the σ-supermartingale Z_i, the main theorem of σ-ARnkSupMs can be proven.

Theorem 3.2.5. *Let η be a σ-ARnkSupM for a starting configuration c_0, a scheduler σ and a target region C, and let $T = T^C$ be the first hitting time of C. Then $\eta(c_0, \sigma) \geq \mathbb{E}_{c_0,\sigma}(T)$.*

Proof. Consider the previously described stochastic process Z_i which is a σ-supermartingale. Because Z_i is a supermartingale, the inequality $\mathbb{E}_{c_0,\sigma}(Z_0) \geq \mathbb{E}_{c_0,\sigma}(Z_i)$ holds for all $i \in \mathbb{N}$. This gives

$$\begin{aligned}
\eta(c_0, \sigma) = Y_0 = Z_0 = \mathbb{E}_{c_0,\sigma}(Z_0) \geq \mathbb{E}_{c_0,\sigma}(Z_i) &= \mathbb{E}_{c_0,\sigma}(Y_i^T + \min\{i, T\}) \\
&= \mathbb{E}_{c_0,\sigma}(Y_i^T) + \mathbb{E}_{c_0,\sigma}(\min\{i, T\}) \geq \mathbb{E}_{c_0,\sigma}(\min\{i, T\}) = \mathbb{E}_{c_0,\sigma}(T \wedge i).
\end{aligned}$$

The sequence $(T \wedge i)$ of stopping times a.s converges to T as $i \to \infty$. Because $(T \wedge i) \leq (T \wedge (i+1))$ trivially holds for all i, the sequence of stopping times is monotonically converging. By the Monotone Convergence Theorem, we then immediately get

$$\mathbb{E}_{c_0,\sigma}(T) = \mathbb{E}_{c_0,\sigma}(\lim_{i\to\infty}(T \wedge i)) = \lim_{i\to\infty} \mathbb{E}_{c_0,\sigma}(T \wedge i) \leq Y_0 = \eta(c_0, \sigma).$$

□

We can generalize this result to not only hold at the starting configuration c_0 but along all paths from c_0.

Corollary 3.2.5.1. *Let η be a σ-ARnkSupM at c_0, then for any path $\pi c = c_0 c_1 \ldots c$ from c_0, we have $\eta(c, \sigma_\pi) \geq \mathbb{E}_{c,\sigma_\pi}(T)$.*

Proof. Note that if η is a σ-ARnkSupM at c_0, then it is also a σ_π-ARnkSupM at c. To see this, let $\pi' c' = c \ldots c'$ be any path from c to some $c' \notin C$, then we have to show that $\eta(c', (\sigma_\pi)_{\pi'}) \geq 1 + \mathbb{X}\eta(c', (\sigma_\pi)_{\pi'})$.

$$\eta(c', (\sigma_\pi)_{\pi'}) = \eta(c', \sigma_{\pi\pi'}) \geq 1 + \mathbb{X}\eta(c', \sigma_{\pi\pi'}) = 1 + \mathbb{X}\eta(c', (\sigma_\pi)_{\pi'})$$

holds because $\pi\pi'$ is a path from c_0 to c' and η is a σ-ARnkSupM at c_0. Now we can apply the previous theorem and get the desired result. □

Similar theorems hold for UARnkSupMs and under restrictions also for LARnkSupMs. While the prove of the former is straightforward, the latter does need a extra condition to be satisfied.

Theorem 3.2.6. *Let η be an UARnkSupM for a starting configuration c_0 and a target region C. Let $T = T^C$ be the first hitting time of C, then $\eta(c_0) \geq \overline{\mathbb{E}}_{c_0}(T)$.*

Proof. From the definition of a UARnkSupM, we have for all paths $\pi c = c_0 \dots c$ where $c \in I \setminus C$ for some invariant I, that $\eta(c) \geq 1 + \overline{\mathbb{X}}\eta(c) \geq 1 + \mathbb{X}\eta(c, \sigma)$ holds for all schedulers $\sigma \in \mathrm{Sch}_\Gamma$. We abuse the notation and see η as a function which takes a scheduler as parameter but simply ignores it, i.e. $\eta(c, \sigma) = \eta(c)$. By this inequality, η is a σ-ARnkSupM for any σ. Now we apply Theorem 3.2.5 and get

$$\eta(c_0) = \eta(c_0, \sigma) \geq \mathbb{E}_{c_0,\sigma}(T).$$

This holds for all schedulers σ, so we take the supremum on both sides and conclude $\eta(c_0) \geq \overline{\mathbb{E}}_{c_0}(T)$. □

We can extend this theorem again to all paths from c_0 as we did before. This can be done for all coming martingale-based methods, so we generally just prove the special case for the starting configuration.

Theorem 3.2.7. *Let η be an LARnkSupM for a starting configuration c_0 and a target region C such that $\mathbb{X}\eta$ admits lower ϵ-optimal schedulers. Let $T = T^C$ be the first hitting time of C, then $\eta(c_0) \geq \underline{\mathbb{E}}_{c_0}(T)$.*

Proof. In contrast to the upper version, there does not necessarily exist any scheduler σ such that η is a σ-ARnkSupM. However, consider a lower ϵ-optimal scheduler σ (for $\mathbb{X}\eta$) such that $\eta(c, \sigma_\pi) \geq \underline{\mathbb{X}}\eta(c) + 1 \geq \mathbb{X}\eta(c, \sigma_\pi) + 1 - \epsilon$ holds for all paths $\pi = c_0 \dots c$ where $c \notin C$. We again abuse notation and define $\eta(c, \sigma) = \eta(c)$. For $\epsilon \in (0, 1)$ we can divide both sides by $(1 - \epsilon)$ and obtain

$$\frac{\eta}{1-\epsilon}(c, \sigma_\pi) \geq \mathbb{X}\left(\frac{\eta}{1-\epsilon}\right)(c, \sigma_\pi) + 1$$

Setting $\eta' = \frac{\eta}{1-\epsilon}$ we see that η' is a σ-ARnkSupM. By Theorem 3.2.5 we have

$$\eta'(c_0) = \eta'(c_0, \sigma) \geq \mathbb{E}_{c_0,\sigma}(T)$$

which implies

$$\eta(c_0) \geq (1 - \epsilon)\mathbb{E}_{c_0,\sigma}(T) \geq (1 - \epsilon)\underline{\mathbb{E}}_{c_0}(T).$$

This holds for all $\epsilon \in (0, 1)$ and in particular for $\epsilon \to 0$ we conclude $\eta(c_0) \geq \underline{\mathbb{E}}_{c_0}(T)$. □

The extra condition assumed in the above theorem is restrictive but satisfied in many cases. If η is lower semianalytic than such schedulers do exist[18, Proposition 7.50]. In particular, any Borel measurable and hence also any continuous function is lower semianalytic and admits ϵ-optimal schedulers.

While additive ranking supermartingales do give sound upper bounds on the reaching time, it is not clear how tight the bounds are. We wish to address this in the following and show that there exist supermartingales such that the inequality is actually satisfied with equality for all three presented versions of ARnkSupMs.

Lemma 3.2.8. *The function $\mathbb{E}_C^{\text{steps}}$ defined as $\mathbb{E}_C^{\text{steps}}(c,\sigma) = \mathbb{E}_{c,\sigma}(T^C)$ is a σ-ARnkSupM for any scheduler σ, any starting configuration c_0 and any invariant I.*

Proof. Let c_0 and σ be given. Let $C \in \mathcal{B}(L \times \mathbb{R}^V)$ be a target region. We need to show that for any path $\pi c = c_0 c_1 \dots c$ with $c \notin C$ we have $\mathbb{E}_C^{\text{steps}}(c,\sigma_\pi) \geq \mathbb{X}\mathbb{E}_C^{\text{steps}}(c,\sigma_\pi) + 1$. Let $T = T^C$ be the first hitting time of C.

$$\begin{aligned}
\mathbb{E}_C^{\text{steps}}(c,\sigma_\pi) &= \mathbb{E}_{c,\sigma_\pi}(T) = \mathbb{E}_{c,\sigma_\pi}(T \mid C_0 = c \notin C) \\
&= \int \mathbb{E}_{c',\sigma_{\pi c}}(T+1)\mu_c^{\sigma_\pi}(\mathrm{d}c') \\
&= \int \mathbb{E}_{c',\sigma_{\pi c}}(T)\mu_c^{\sigma_\pi}(\mathrm{d}c') + 1 \\
&= \int \mathbb{E}_C^{\text{steps}}(c',\sigma_{\pi c})\mu_c^{\sigma_\pi}(\mathrm{d}c') + 1 \\
&= \mathbb{X}\mathbb{E}_C^{\text{steps}}(c,\sigma_\pi) + 1
\end{aligned}$$

The important equality is from the first line to the second line. Here we use the fact, that $T(c_0 c_1 c_2 \dots) = T(c_1 c_2 \dots) + 1$ holds whenever $c_0 \notin C$. This gives $\mathbb{E}_{c_0,\sigma}(T) = \int \mathbb{E}_{c_1,\sigma_{c_0}}(T+1)\mu_{c_0}^{\sigma}(\mathrm{d}c_1)$. □

Corollary 3.2.8.1. *σ-additive ranking supermartingales are sound and complete in witnessing positive almost sure reachability, that is, if $\mathbb{E}_{c_0,\sigma}(T^C) < \infty$, then there is exists a σ-ARnkSupM for c_0 witnessing it.*

Proof. For completeness we just take $\mathbb{E}_C^{\text{steps}}$ as the σ-additive ranking supermartingale. For soundness we use Theorem 3.2.5. □

Lemma 3.2.9. *The function $\mathbb{E}_C^{\text{steps}}$ admits lower and upper ϵ-optimal schedulers. Furthermore, they can be chosen to be history-independent.*

Proof. Instead of giving a proof, we refer to Lemma 3.2 by Takisaka et al.[11]. However, we give a short explanation. A pCFG can be seen as a general Markov Decision Process, which is a well studied type of system in control theory. A commonly analyzed problem is the expected reward problem, where to every state and action of the system a reward is associated and the task is to calculate what average reward is accumulated over a run of the system (oftentimes discounted or time bounded). It can be shown that if the rewards are positive and upper bounded, then ϵ-optimal schedulers exist that minimize or maximize the expected reward. The expected reaching time can be transformed into such a model by associating a constant reward of 1 to every transition that does not lead the target region, and by modifying the system to loop inside the target region once reached. Now the expected reaching time coincides with the average accumulated reward and hence ϵ-optimal schedulers exist. We use this argument later to argue about the existence of ϵ-optimal schedulers for other functions as well. A concrete construction from pCFG to the stochastic optimal control model is given in [11, Lemma 3.2], and the existence of ϵ-optimal schedulers is proven in [18, Proposition 7.50 and Proposition 9.19]. □

Corollary 3.2.9.1. *UARnkSupM and LARnkSupM are sound and complete in witnessing upper and lower positive almost sure reachability. With upper and lower positive almost sure reachability we refer to $\overline{\mathbb{E}}(T^C)$ and $\underline{\mathbb{E}}(T^C)$ being finite.*

Proof. By Lemma 3.2.2 and 3.2.9 we have $\overline{\mathbb{X}\mathbb{E}_C^{\text{steps}}} = \overline{\mathbb{X}}\,\overline{\mathbb{E}}_C^{\text{steps}}$ and taking the supremum over all schedulers in the equation $\mathbb{E}_C^{\text{steps}}(c,\sigma_\pi) = \mathbb{X}\mathbb{E}_C^{\text{steps}}(c,\sigma_\pi) + 1\,(c \notin C \text{ and } \pi c = c_0 c_1 \dots c)$ gives

$$\overline{\mathbb{E}}_C^{\text{steps}}(c) = \overline{\mathbb{X}}\,\overline{\mathbb{E}}_C^{\text{steps}}(c) + 1 \geq \overline{\mathbb{X}}\,\overline{\mathbb{E}}_C^{\text{steps}}(c) + 1.$$

$\overline{\mathbb{E}}_C^{\text{steps}}$ is a UARnkSupM satisfying $\overline{\mathbb{E}}_C^{\text{steps}}(c_0) = \overline{\mathbb{E}}_{c_0}(T)$. By the same argument it can be shown that $\underline{\mathbb{E}}_C^{\text{steps}}$ is a LARnkSupM satisfying $\underline{\mathbb{E}}_C^{\text{steps}}(c_0) = \underline{\mathbb{E}}_{c_0}(T)$. We need to show that $\mathbb{X}\underline{\mathbb{E}}_C^{\text{steps}}$ admits lower ϵ-optimal schedulers to apply Theorem 3.2.7. The proof for this is quite technical and requires more background in measure theory, analytic sets and selection theorems. We refer to the excellent book by Bertsekas and Shreve on stochastic optimal control[18]. According to them $\underline{\mathbb{E}}_C^{\text{steps}}$ is lower semianalytic as the infimization of a lower semianalytic function $\mathbb{E}_C^{\text{steps}}$. Applying the nexttime operator to a lower semianalytic function yields again a lower semianalytic function [18, Corollary 7.48.1]. Lower semianalytic functions admit ϵ-optimal selections [18, Proposition 7.50]. □

Upper positive almost sure reachability is different from positive almost sure reachability for all schedulers σ, that is, it does not witness $\forall\sigma : \mathbb{E}_\sigma(T^C) < \infty$ but rather $\forall\sigma : \mathbb{E}_\sigma(T^C) < k$ for some constant k. On the other hand, lower positive almost sure reachability does witness $\exists\sigma : \mathbb{E}_\sigma(T^C) < \infty$, i.e. the existence of a scheduler which reaches the target almost surely in finite time.

The following program does terminate in finite time for any scheduler σ, but this time can be chosen to be arbitrarily large and as such, $\overline{\mathbb{E}}(T^C)$ is infinite. Termination under all schedulers is not witnessed by any UARnkSupM.

Listing 3.1: Example: Finite but unbounded termination times

```
n := ndet(N);
while n > 0 do
        n := n - 1;
od
```

The Corollaries 3.2.8.1 and 3.2.9.1 show that $\mathbb{E}_C^{\text{steps}}$, $\overline{\mathbb{E}}_C^{\text{steps}}$ and $\underline{\mathbb{E}}_C^{\text{steps}}$ are the least functions satisfying their respective inequalities. Indeed, they satisfy their respective ARnkSupM conditions with equality. Any σ/upper/lower additive ranking supermartingale is an overapproximation of $\mathbb{E}_C^{\text{steps}}/\overline{\mathbb{E}}_C^{\text{steps}}/\underline{\mathbb{E}}_C^{\text{steps}}$. These properties are reminiscent of least fixed points and their overapproximations via pre-fixpoints. In light of this observation, we briefly redevelop the obtained results in a fixed point theoretic framework in accordance to the approach by Takisaka et al.[11].

Theorem 3.2.10. *$\mathbb{E}_C^{\text{steps}}$ defined by $\mathbb{E}_C^{\text{steps}}(c,\sigma) = \mathbb{E}_{c,\sigma}(T^C)$ is the least fixed point of the monotone operator*

$$\Phi_C(\eta) = \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\mathbb{X}\eta.$$

Any pre-fixpoint $\eta \sqsupseteq \Phi_C(\eta)$ of Φ_C overapproximates $\mathbb{E}_C^{\text{steps}}$. Here we consider the complete lattice of functions $[(L \times \mathbb{R}^V) \times \text{Sch}_\Gamma \to [0,\infty]]$.

Proof. First we note that $\mathbb{E}_C^{\text{steps}}$ is indeed a fixed point which was proven in Lemma 3.2.8. We define $\mathbb{E}_C^{\text{steps}\leq n} := \mathbb{E}(T \wedge n)$ where $T = T^C$ is the first hitting time of C. It is easy to see that the sequence $(\mathbb{E}_C^{\text{steps}\leq n})_{n\in\mathbb{N}}$ converges monotonically to $\mathbb{E}_C^{\text{steps}} = \mathbb{E}(T)$ as we have shown previously using the Monotone Convergence Theorem. We now show that this sequence is actually the Cousot-Cousot sequence for the least fixed point, that is, $\mathbb{E}_C^{\text{steps}\leq n} = \Phi_C^n(\bot)$.

We proceed by induction. For $n = 0$, both sides equal the constant zero function for every scheduler σ. We look at the case $n+1$ and assume that $\mathbb{E}_C^{\text{steps}\leq n} = \Phi_C^n(\bot)$ holds for all schedulers σ. First note that both sides are equal to 0 on C. We consider $c \notin C$. Notice that $(T \wedge (n+1))(c_0c_1\dots) = (T \wedge n)(c_1c_2\dots) + 1$ for all $n \in \mathbb{N}$ and $c_0 \notin C$. Then we have

$$\begin{aligned}
\mathbb{E}_C^{\text{steps}\leq n+1}(c,\sigma) &= \mathbb{E}_{c,\sigma}(T \wedge (n+1)) = \mathbb{E}_{c,\sigma}(T \wedge (n+1) \mid C_0 = c \notin C) \\
&= \int \mathbb{E}_{c',\sigma_c}((T \wedge n) + 1)\mu_c^\sigma(\mathrm{d}c') \\
&= 1 + \int \mathbb{E}_{c',\sigma_c}(T \wedge n)\mu_c^\sigma(\mathrm{d}c') \\
&= 1 + \int \mathbb{E}_C^{\text{steps}\leq n}(c',\sigma_c)\mu_c^\sigma(\mathrm{d}c') = 1 + \mathbb{X}\mathbb{E}_C^{\text{steps}\leq n}(c,\sigma) \\
&= \Phi_C(\mathbb{E}_C^{\text{steps}\leq n})(c,\sigma) = \Phi_C(\Phi_C^n\bot)(c,\sigma) = (\Phi_C^{n+1}\bot)(c,\sigma)
\end{aligned}$$

This proves that $\mathbb{E}_C^{\text{steps}} = \lim_{n\to\infty} \mathbb{E}_C^{\text{steps}\leq n}$ is the least fixed point of the operator Φ_C. By Knaster-Tarski we have that it is also the least pre-fixpoint and therefore every pre-fixpoint of Φ_C overapproximates $\mathbb{E}_C^{\text{steps}}$. □

Remark. The arguments of the function $\Phi_C(\eta)$ consist of both configuration and scheduler so the indicator function that disregards the scheduler is actually $\mathbb{1}_{\overline{C}\times\text{Sch}_\Gamma}$, but for notational simplicity we just denote it by $\mathbb{1}_{\overline{C}}$.

Notice that a pre-fixpoint of Φ_C is a function η satisfying the condition $\eta \geq 1 + \mathbb{X}\eta$ on $\overline{C}$, i.e. it satisfies the conditions of a σ-ARnkSupM. The tight connection between pre-fixpoints and supermartingales becomes apparent.

Remarkably, neither martingale theory nor the fact that T is a stopping time is used. Soundness and completeness follow directly from Knaster-Tarski and Cousot-Cousot. It is possible to define $\mathbb{E}_C^{\text{steps}}$ and $\mathbb{E}_C^{\text{steps}\leq n}$ without directly but rather implicitly relying on the first hitting time T as shown by Takisaka et al. in [11]. Their approach reduces the use of probability theory to a minimum. However, they only applied this reasoning to the extremizations $\overline{\mathbb{E}}_C^{\text{steps}}$ and $\underline{\mathbb{E}}_C^{\text{steps}}$. Furthermore, ranking supermartingales defined this way are different in the sense that they satisfy the inequality for every scheduler everywhere (possible restricted to an invariant) and not just along valid paths from some starting configuration c_0.

The fixed point characterization of Theorem 3.2.10 is not practically feasible because one has to compute the nexttime operator with respect to every possible scheduler in each

iteration. For a more practical approach, it is possible to fix a single history-independent scheduler σ and characterize $\mathbb{E}^{\text{steps}}_{C,\sigma}$ as the least fixed point of an adapted operator $\Phi_{C,\sigma}$ instead.

Lemma 3.2.11. *Let σ be a history-independent scheduler, then $\mathbb{E}^{\text{steps}}_{C,\sigma}$ defined by $\mathbb{E}^{\text{steps}}_{C,\sigma}(c) = \mathbb{E}_{c,\sigma}(T^C)$ is the least fixed point of the monotone operator*

$$\Phi_{C,\sigma}(\eta) = \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\mathbb{X}\eta(\cdot,\sigma).$$

where $\eta : (L \times \mathbb{R}^V) \to [0,\infty]$ is a scheduler-independent function. We set $\eta(c,\sigma) = \eta(c)$ for all schedulers σ to make it compatible with the definition of $\mathbb{X}$.

Proof. Follows by the same arguments as in the proof of Theorem 3.2.10 and using the fact that $\sigma = \sigma_\pi$ holds for all valid paths π. □

It is possible to generalize this to an arbitrary fixed scheduler. We do not further follow this approach in detail, but give a brief description instead. For a fixed scheduler σ, it is sufficient to store its history π for the fixed point computation. One considers a history-dependent function $\eta : (L \times \mathbb{R}^V) \times (L \times \mathbb{R}^V)^* \to [0,\infty]$ and defines the operator $\Phi_{C,\sigma}(\eta)(c,\pi) = 1 + \int \eta(c',\pi c)\mu_c^{\sigma_\pi}(\mathrm{d}c')$ for $c \notin C$, and 0 elsewhere. The mapping $(c,\pi) \mapsto \mathbb{E}^{\text{steps}}_C(c,\sigma_\pi)$ is then the least fixed point of the operator.

The proof of 3.2.10 can be straightforwardly adapted to UARnkSupM and the operator $\overline{\Phi}_C(\eta) = \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\overline{\mathbb{X}}\eta$. One essentially takes all equations presented in the proof and adds suprema over all schedulers in every line. The desired result is obtained by the fact that ϵ-optimal schedulers exist, and that $\overline{\Phi}_C$ is ω-continuous and hence preserves the chain supremum of the Cousot-Cousot sequence.

Unfortunately, the proof does not transfer to the case of LARnkSupM. While the first arguments are straightforwardly adapted to the lower case, the problem is in the convergence of the sequence $\underline{\mathbb{E}}^{\text{steps}\le n}$. This sequence is increasing but the operator $\underline{\Phi}_C$ is ω^{op}-continuous and not ω-continuous. The increasing sequence $\underline{\mathbb{E}}^{\text{steps}\le n}_C$ does not necessarily converge after ω steps, and transfinite induction steps are needed. However, under the assumption that any fixed point η of $\underline{\Phi}_C(\eta) = \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\underline{\mathbb{X}}\eta$ admits ϵ-optimal schedulers, it can be shown that $\underline{\mathbb{E}}^{\text{steps}}_C$ is indeed the least fixed point. Takisaka et al.[11] construct an ϵ-optimal scheduler for an arbitrary fixed point to prove this statement but we do not see any justification as to why their constructed scheduler is measurable in their proof. While this statement may hold even in the absence of such schedulers, we have no such proof at hand and as such consider the additional assumption as described above necessary.

Theorem 3.2.12. *$\overline{\mathbb{E}}^{\text{steps}}_C$ is the least fixed point of the monotone operator*

$$\overline{\Phi}_C(\eta) = \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\overline{\mathbb{X}}\eta.$$

Any pre-fixpoint $\eta \sqsupseteq \overline{\Phi}_C(\eta)$ of $\overline{\Phi}_C$ overapproximates $\overline{\mathbb{E}}^{\text{steps}}_C$.

Proof. Instead of adapting the proof of Theorem 3.2.10, we directly use its result in a similar fashion as we did in the proof of Theorem 3.2.6. Consider an arbitrary fixed point of η of

$\overline{\Phi}_C$. By abuse of notation we define $\eta(c,\sigma) = \eta(c)$. Now we have $\eta = \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\overline{\mathbb{X}}\eta \geq \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\mathbb{X}\eta$. η is thus a pre-fixpoint of the operator Φ_C and as such it is larger than $\mathbb{E}_C^{\text{steps}}$ (from Theorem 3.2.10). We have $\eta(c) = \eta(c,\sigma) \geq \mathbb{E}_C^{\text{steps}}(c,\sigma)$. Since this holds for all schedulers we conclude $\eta \geq \overline{\mathbb{E}}_C^{\text{steps}}$. As shown in the proof of Corollary 3.2.9.1 $\overline{\mathbb{E}}_C^{\text{steps}}$ is a fixed point and by the previous observation it must be the least one.

The operator $\overline{\Phi}_C$ is ω-continuous which follows from the fact that $\overline{\mathbb{X}}$ is ω-continuous (see Lemma 3.2.1). Using this fact, we see that the least fixed point is reached after at most ω steps in the Cousot-Cousot sequence.

$$\overline{\Phi}_C(\lim_{n\to\infty}\overline{\Phi}_C^n\bot) = \lim_{n\to\infty}\overline{\Phi}_C^n(\overline{\Phi}_C\bot) = \lim_{n\to\infty}\overline{\Phi}_C^{n+1}(\bot) = \lim_{n\to\infty}\overline{\Phi}_C^n(\bot)$$

And again, every pre-fixpoint is an overapproximation of $\overline{\mathbb{E}}_C^{\text{steps}}$ by Knaster-Tarski. □

For the lower case we consider an extra condition to guarantee that the function $\underline{\mathbb{E}}_C^{\text{steps}}$ is the least fixed point. Namely, we show that this function is the least fixed point among all functions η for which $\mathbb{X}\eta$ admits ϵ-optimal schedulers. As previously mentioned, this can be guaranteed for a wide class of functions, namely all lower semianalytic functions and this includes all Borel measurable functions.

Theorem 3.2.13. *$\underline{\mathbb{E}}_C^{\text{steps}}$ is the least fixed point of the monotone operator*

$$\underline{\Phi}_C(\eta) = \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\underline{\mathbb{X}}\eta.$$

under the assumption that for any fixed point η, the function $\mathbb{X}\eta$ admits lower ϵ-optimal schedulers. Any pre-fixpoint $\eta \sqsupseteq \underline{\Phi}_C(\eta)$ of $\underline{\Phi}_C$ overapproximates $\underline{\mathbb{E}}_C^{\text{steps}}$.

Proof. Let η be any fixed point of $\underline{\Phi}_C$. By assumption there exists for every $\epsilon > 0$ a scheduler σ^ϵ such that $\underline{\mathbb{X}}\eta + \epsilon \geq \mathbb{X}\eta(\cdot,\sigma^\epsilon)$. This scheduler can be chosen to be history-independent by defining σ as $\sigma(\pi c) = \sigma^\epsilon(c)$. Now we get

$$\begin{aligned}\eta = \underline{\Phi}_C(\eta) &= \mathbb{1}_{\overline{C}} + \mathbb{1}_{\underline{C}}\underline{\mathbb{X}}\eta \\ &\geq \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}(\mathbb{X}\eta(\cdot,\sigma) - \epsilon\mathbb{1}) = (1-\epsilon)\mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\mathbb{X}\eta(\cdot,\sigma).\end{aligned}$$

Dividing both sides by $(1-\epsilon)$ for $\epsilon \in (0,1)$ gives

$$\frac{\eta}{1-\epsilon} \geq \mathbb{1}_{\overline{C}} + \mathbb{1}_{\overline{C}}\mathbb{X}(\frac{\eta}{1-\epsilon})(\cdot,\sigma)$$

so the function $\eta' = \frac{\eta}{1-\epsilon}$ is a pre-fixpoint of $\Phi_{C,\sigma}$ and by Lemma 3.2.11 it overapproximates $\mathbb{E}_{C,\sigma}^{\text{steps}} \geq \underline{\mathbb{E}}_C^{\text{steps}}$. From that it follows that $\eta \geq (1-\epsilon)\underline{\mathbb{E}}_C^{\text{steps}}$ for all $\epsilon \in (0,1)$. Letting $\epsilon \to 0$ gives the desired result $\eta \geq \underline{\mathbb{E}}_C^{\text{steps}}$. $\underline{\mathbb{E}}_C^{\text{steps}}$ itself is a fixed point of $\underline{\Phi}_C$ and hence it is the least fixed point. By Knaster-Tarski any pre-fixed point is larger than the least fixed point and therefore overapproximates $\underline{\mathbb{E}}_C^{\text{steps}}$. □

Both presented proofs regarding UARnkSupM and LARnkSupM are essentially identical to the previously shown proofs using martingale theory. The difference is solely in the invoked theorem for σ-ARnkSupM which was once proven via martingale theory (Theorem 3.2.5) and once via fixed point theory (Theorem 3.2.10). This concludes our discussion about expected reaching times in probabilistic programs. Effective ways to find such supermartingales via template-based synthesis methods are presented in Chapter 4.

3.2.2 Ranking Supermartingales for Higher Order Moments

The ranking supermartingales in Subsection 3.2.1 are used to overapproximate the expected reaching time and, if it is finite, conclude that the program almost surely reaches the target region eventually. Instead of overapproximating the expectation (i.e. the first moment), one can also use supermartingales to get approximations on higher moments. Higher moments can be used to bound tail probabilities, i.e. the probability that the program takes at least a certain amount of steps before reaching the target. This is established by the generalized Markov's inequality.

Proposition 3.2.14 (Markov's inequality)**.** *Let T be a non-negative, real-valued random variable and ϕ a non-negative, monotonically increasing function, then for every $d \geq 0$ with $\phi(d) > 0$ it holds that*

$$\Pr(T \geq d) \leq \frac{\mathbb{E}(\phi(T))}{\phi(d)}$$

Corollary 3.2.14.1. *Under the assumptions of Proposition 3.2.14, for any $d > 0$ the inequality*

$$\Pr(T \geq d) \leq \frac{\mathbb{E}(T^n)}{d^n}.$$

holds.

Proof. We set $\phi(x) = x^n$ for any $n \in \mathbb{N}$ and apply Proposition 3.2.14. □

The goal is to find upper bounds on the higher moments $\mathbb{E}(T^n)$ and use them to bound the tail probabilities of T. Since the inequality gives a bound on the same tail probability for any $n \in N$, it can be useful to compute bounds on multiple higher order moments for various values of n and then choose the one which gives the tightest bound.

The ranking supermartingales in Subsection 3.2.1, in a sense, model the change of expectation as time elapses in the system via the nexttime operators. We made use of the fact that $T(c_0c_1\dots) = T(c_1c_2\dots) + 1$ is true for all c_0 which are not already in some target region C. From this we concluded the relation $\mathbb{E}_{c_0,\sigma}(T \mid c_0 \notin C) = \int \mathbb{E}_{c_1,\sigma_{c_0}}(T+1)\mu^{\sigma}_{c_0}(\mathrm{d}c_1) = 1 + \int \mathbb{E}_{c_1,\sigma_{c_0}}(T)\mu^{\sigma}_{c_0}(\mathrm{d}c_1)$. Identifying $\mathbb{E}_{c,\sigma}(T)$ with the expression $\eta(c,\sigma)$ then gives rise to the fixed point characterization $\eta(c,\sigma) = 1 + \mathbb{X}\eta(c,\sigma)$. However, higher moments of T do not change linearly with time. Indeed, for the second moment we have

$$\begin{aligned}\mathbb{E}_{c_0,\sigma}(T^2 \mid c_0 \notin C) &= \int \mathbb{E}_{c_1,\sigma_{c_0}}((T+1)^2)\mu^{\sigma}_{c_0}(\mathrm{d}c_1) \\ &= 1 + 2\int \mathbb{E}_{c_1,\sigma_{c_0}}(T)\mu^{\sigma}_{c_0}(\mathrm{d}c_1) + \int \mathbb{E}_{c_1,\sigma_{c_0}}(T^2)\mu^{\sigma}_{c_0}(\mathrm{d}c_1).\end{aligned}$$

So to compute the second moment after a transition, knowledge about the first moment is needed. Generally, to compute the n-th moment this way, knowledge about all lower moments up to the n-th is needed. This simultaneous computation can be expressed by defining the n-th order ranking supermartingale to be an n-dimensional vector containing an entry for each moment. In the case of $n = 2$, we would consider $\eta = (\eta_1, \eta_2)$ and the (simplified) fixed point equations $\eta_1 = 1 + \mathbb{X}\eta$ and $\eta_2 = 1 + 2\mathbb{X}\eta_1 + \mathbb{X}\eta_2$. The constant 1 term may also be considered as the 0-th moment, i.e. $\mathbb{E}(T^0)$.

The following results are mainly due to Kura et al.[14]. Unlike the one dimensional case of ARnkSupMs, a purely martingale theoretic approach to the following results seems more complicated than the fixed point theoretic approach. We use the advantage of having two frameworks to work with and, in this case, only use the fixed point theoretic one.

Definition 3.2.6 (Higher order additive ranking supermartingale[14])**.** Let $\sigma \in \mathrm{Sch}_\Gamma$ be a scheduler, $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region and c_0 be a starting configuration. A universally measurable vector-valued function $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to [0,\infty]^n$ is called an *n-order additive ranking σ-supermartingale* ((n,σ)-ARnkSupM) at c_0 if it satisfies for all paths $\pi c = c_0 c_1 \dots c$ with $c \notin C$ and $k \in \{1,\dots,n\}$:

$$\eta_k(c,\sigma_\pi) \geq \sum_{i=1}^{k} \binom{k}{i} \mathbb{X}\eta_i(c,\sigma_\pi) + 1.$$

In particular, for $k = 1$ it reduces to the definition of an ordinary additive ranking σ-supermartingale. As before, for a universally measurable function $\eta : L \times \mathbb{R}^V \to [0,\infty]^n$ we define *n-order upper additive ranking supermartingales* (n-UARnkSupM) to satisfy

$$\eta_k(c) \geq \sum_{i=1}^{k} \binom{k}{i} \overline{\mathbb{X}}\eta_i(c) + 1,$$

under the previous conditions. The lower version is defined analogously.

The domain may be restricted to an invariant I without any problems, but to reduce notational overload we consider the trivial invariant $I = L \times \mathbb{R}^V$ only from now on.

The k-th entry of η is associated to the k-th moment of the first hitting time T^C. As a single elapsing time unit has different effects on the different moments, we introduce an *elapse* function which represents this effect.

Definition 3.2.7 (Elapse function El^n[14])**.** The function $\mathrm{El}^n : [0,\infty]^n \to [0,\infty]^n$ is defined as

$$\mathrm{El}^n_k(x_1,x_2,\dots,x_n) = \sum_{i=1}^{k} \binom{k}{i} x_i + 1.$$

Using the elapse function and defining $\mathbb{X}\eta = (\mathbb{X}\eta_1, \mathbb{X}\eta_2, \dots, \mathbb{X}\eta_n)$, the supermartingale condition can be restated as $\eta(c,\sigma_\pi) \geq \mathbb{X}(\mathrm{El}^n \circ \eta)(c,\sigma_\pi)$, where the inequality sign $\geq$ is to be interpreted componentwise. Note by linearity of the $\mathbb{X}$ operator, we can also exchange $\mathbb{X}$ and El^n, that is, $\mathbb{X}(\mathrm{El}^n \circ \eta) = \mathrm{El}^n \circ \mathbb{X}\eta$

We are ready to state the main theorem of this section.

Theorem 3.2.15. *Let $C \in \mathcal{B}(L \times \mathbb{R}^V)$ be a target region, σ a scheduler and $n \in \mathbb{N} \setminus \{0\}$ a positive natural number. Let η be an additive ranking σ-supermartingale of order n for C at c_0. Then for all $k \in \{1,\dots,n\}$ we have*

$$\forall c \in (L \times \mathbb{R}^V) : \forall \pi c = c_0 c_1 \dots c : \eta_k(c,\sigma_\pi) \geq \mathbb{E}_{c,\sigma_\pi}(T^k)$$

where $T = T^C$ is the first hitting time of C.

Proof. By monotone convergence it is easy to see that $\mathbb{E}_{c,\sigma}((T \wedge m)^k)$ converges to $\mathbb{E}_{c,\sigma}(T^k)$ as $m \to \infty$. Let η be some (n,σ)-ARnkSupM at c_0. We show by induction that for all $m \in \mathbb{N}$ and $k \in \{1,\dots,n\}$ we have $\eta_k(c,\sigma_\pi) \geq \mathbb{E}_{c,\sigma_\pi}((T \wedge m)^k)$. For $m = 0$ it holds trivially since the right-hand side equals 0 everywhere. Assume it holds for all values $< m+1$. The right-hand side $\mathbb{E}_{c,\sigma_\pi}((T \wedge m)^k)$ is constantly 0 on C so we only have to consider $c \notin C$. For any $k \leq n$ and $c_0 \notin C$, we have

$$\begin{aligned}(T \wedge (m+1))^k(c_0c_1\dots) &= ((T \wedge m)+1)^k(c_1c_2\dots)\\ &= (1 + \sum_{i=1}^{k} \binom{k}{i} (T \wedge m)^i)(c_1c_2\dots)\\ &= \mathrm{El}^n_k \circ ((T \wedge m), (T \wedge m)^2, \dots, (T \wedge m)^k, \dots, (T \wedge m)^n)(c_1c_2\dots)\end{aligned}$$

In the following we use the notation $\mathbb{E}(T)(c,\sigma) = \mathbb{E}_{c,\sigma}(T)$ to make it compatible with the nexttime operator $\mathbb{X}$. Taking the expectation $\mathbb{E}_{c,\sigma_\pi}$ over the above equality for $c \notin C$ then gives

$$\begin{aligned}&\mathbb{E}((T \wedge (m+1))^k)(c,\sigma_\pi) = \mathbb{X}\mathbb{E}(\mathrm{El}^n_k \circ ((T \wedge m), (T \wedge m)^2, \dots, (T \wedge m)^k, \dots, (T \wedge m)^n))(c,\sigma_\pi)\\ &= \mathrm{El}^n_k \circ (\mathbb{X}\mathbb{E}(T \wedge m)(c,\sigma_\pi), \mathbb{X}\mathbb{E}((T \wedge m)^2)(c,\sigma_\pi), \dots, \mathbb{X}\mathbb{E}((T \wedge m)^n)(c,\sigma_\pi))\\ &\leq \mathrm{El}^n_k \circ (\mathbb{X}\eta_1(c,\sigma_\pi), \mathbb{X}\eta_2(c,\sigma_\pi), \dots, \mathbb{X}\eta_n(c,\sigma_\pi))\\ &= \mathrm{El}^n_k \circ (\mathbb{X}\eta)(c,\sigma_\pi) \leq \eta_k(c,\sigma_\pi).\end{aligned}$$

From the first to second line we used the linearity of the expectation and the nexttime operator. From the second to third we used the induction hypothesis, and in the last line we used the defining inequality of (n,σ)-ARnkSupM.

The above inequality holds for all $m \in \mathbb{N}$, in particular for $m \to \infty$ we conclude $\mathbb{E}(T^k)(c,\sigma_\pi) \leq \eta_k(c,\sigma_\pi)$. □

Corollary 3.2.15.1. *$(\mathbb{E}(T), \mathbb{E}(T^2), \dots, \mathbb{E}(T^n))$ is a (n,σ)-ARnkSupM for all $n \in \mathbb{N}$, scheduler σ and starting configurations c_0. Furthermore, it is the least (n,σ)-ARnkSupM.*

Proof. From the proof of Theorem 3.2.15 we have the equality (for $c \notin C$)

$$\begin{aligned}&\mathbb{E}((T \wedge (m+1))^k)(c,\sigma_\pi)\\ &= \mathrm{El}^n_k \circ (\mathbb{X}\mathbb{E}(T \wedge m)(c,\sigma_\pi), \mathbb{X}\mathbb{E}((T \wedge m)^2)(c,\sigma_\pi), \dots, \mathbb{X}\mathbb{E}((T \wedge m)^n)(c,\sigma_\pi)).\end{aligned}$$

By monotone convergence and ω-continuity of $\mathbb{X}$, we get for $m \to \infty$

$$\mathbb{E}(T^k)(c,\sigma_\pi) = \mathrm{El}^n_k \circ (\mathbb{X}\mathbb{E}(T)(c,\sigma_\pi), \mathbb{X}\mathbb{E}(T^2)(c,\sigma_\pi), \dots, \mathbb{X}\mathbb{E}(T^n)(c,\sigma_\pi))$$

showing that $(\mathbb{E}(T), \mathbb{E}(T^2), \dots, \mathbb{E}(T^n))$ satisfies the conditions of a (n,σ)-ARnkSupM. From Theorem 3.2.15 it must be the least one. □

Similar to the one dimensional case of σ-ARnkSupM, higher order ARnkSupMs may be characterized as pre-fixpoints of an appropriate operator. Indeed, the proof of Theorem 3.2.15 implicitly makes use of the Cousot-Cousot sequence given by the stopped processes $(T \wedge m)$. We shortly formalize this observation in the following theorem.

Theorem 3.2.16. *Let $C \in \mathcal{B}(L \times \mathbb{R}^V)$ be some target region and let σ be some scheduler. For any n, the tuple $\vec{\mathbb{E}} = (\mathbb{E}(T), \mathbb{E}(T^2), \ldots, \mathbb{E}(T^n))$ is the least fixed point of the monotone operator*

$$\Phi_{n,C}(\eta) = \mathbb{1}_{\overline{C}}(\mathrm{El}^n \circ \mathbb{X}\eta).$$

Any pre-fixpoint is an overapproximation of the first n moments of $T = T^C$.

Proof. The second claim follows by Knaster-Tarski if we can show that the tuple of moments is indeed a least fixed point. Denote by $\vec{\mathbb{E}}^m$ the tuple $(\mathbb{E}(T \wedge m), \mathbb{E}((T \wedge m)^2), \ldots, \mathbb{E}((T \wedge m)^n))$ of bounded moments. From the calculations of the proof of Theorem 3.2.15, it is easily seen that $\vec{\mathbb{E}}^{m+1} = \Phi_{n,C}\vec{\mathbb{E}}^m$ holds. From $\bot = \vec{\mathbb{E}}^0$ we immediately get that the Cousot-Cousot sequence $\Phi^m_{n,C}(\bot) = \vec{\mathbb{E}}^m$ converges to $\vec{\mathbb{E}}$ as m tends towards infinity. □

Next we want to analyze the upper and lower versions, that is, the supremization and infimization over all schedulers. Unlike the simple ARnkSupM case, the upper version for higher order ARnkSupM turns out to be sound but not complete, and the lower version does not even achieve soundness. The reason for this discrepancy is easy to see. Maximizing a higher order martingale entails maximizing multiple objectives, namely, the various moments $\mathbb{E}(T^k)$. There might not exist a single sequence of schedulers which uniformly maximizes all moments simultaneously. In other words, for every moment there is generally a different sequence of schedulers maximizing that specific moment. The higher order UARnkSupMs, however, overapproximate all moments simultaneously. By the dependence of a component on all the components of lower index, once a lower index component is strictly overapproximating its corresponding moment, all higher index components will necessarily be strict overapproximations of their respective moments. For the higher order LARnkSupMs this same reasoning shows that higher index components might actually underapproximate their respective moment, making the method unsound.

Theorem 3.2.17. *Let $C \in \mathcal{B}(L \times \mathbb{R}^V)$ be a target region and $n \in \mathbb{N} \setminus \{0\}$ a positive natural number. Let η be an upper additive ranking supermartingale of order n for C at c_0. Then for all $k \in \{1, \ldots, n\}$ we have*

$$\forall c \in (L \times \mathbb{R}^V) : \forall \pi c = c_0 c_1 \ldots c : \eta_k(c) \geq \overline{\mathbb{E}}_c(T^k)$$

where $T = T^C$ is the first hitting time of C.

Proof. We abuse notation and define $\eta(c, \sigma) = \eta(c)$ for all schedulers σ. For $c \in C$ the inequality holds trivially. Otherwise we have for all paths $\pi c = c_0 c_1 \ldots c$ with $c \notin C$:

$$\eta(c, \sigma_\pi) \geq \mathrm{El}^n \circ \overline{\mathbb{X}}\eta(c) \geq \mathrm{El}^n \circ \mathbb{X}\eta(c, \sigma_\pi).$$

So η is a (n, σ)-ARnkSupM for any scheduler σ, and hence we have $\eta_k \geq \mathbb{E}_\sigma(T^k)$ by Theorem 3.2.15. Taking the supremum over all schedulers then gives the desired result. □

Remark. It is again possible to view a higher order URnkSupM as a pre-fixpoint of a suitable operator

$$\overline{\Phi}_{n,C}(\eta) = \mathbb{1}_{\overline{C}}(\mathrm{El}^n \circ \overline{\mathbb{X}}\eta).$$

The least fixed point of this operator, however, is in general strictly larger than $(\overline{\mathbb{E}}(T), \ldots, \overline{\mathbb{E}}(T^n))$. This fixed point does not admit a nice and concise description in terms of naturally arising expectations of suitable random variables.

To show that the method is incomplete, consider the following example 3.2. The program has to choose one of two branches in the beginning. Under branch (a), the program takes exactly 3 steps to reach the assert (which marks the target region), so $\mathbb{E}_a(T) = 3$ and $\mathbb{E}_a(T^2) = 9$. Under branch (b) there is a probability of p that the program takes only 2 steps, and a probability $(1-p)$ that it needs 4 steps. The moments associated to that action are $\mathbb{E}_b(T) = 2p + 4(1-p) = 4 - 2p$ and $\mathbb{E}_b(T^2) = 2^2p + 4^2(1-p) = 16 - 12p$. As previously discussed, we consider a situation where one action maximizes the first moment while another maximizes the second moment. It turns out that for $p = \frac{13}{24}$ we have $\mathbb{E}_a(T) = 3 > 3 - \frac{1}{13} = \mathbb{E}_b(T)$ but $\mathbb{E}_a(T^2) = 9 < \frac{19}{2} = \mathbb{E}_b(T^2)$. Action (a) maximizes the first moment while action (b) maximizes the second. Now consider a URnkSupM of order 2 for this program. By Theorem 3.2.17, we have $\eta_1 \geq \overline{\mathbb{E}}(T)$ and $\eta_2 \geq \overline{\mathbb{E}}(T^2)$. From the definition of η we also have $\eta_2 \geq 1 + 2\overline{\mathbb{X}}\eta_1 + \overline{\mathbb{X}}\eta_2 \geq 1 + 2\overline{\mathbb{X}}\overline{\mathbb{E}}(T) + \overline{\mathbb{X}}\overline{\mathbb{E}}(T^2)$. From the starting configuration we can calculate $\overline{\mathbb{X}}\overline{\mathbb{E}}(T)(c_0) = 2$ (by taking branch (a)) and $\overline{\mathbb{X}}\overline{\mathbb{E}}(T^2)(c_0) = \frac{14}{3}$ (by taking branch (b)). It follows that $\eta_2(c_0) \geq 1 + 2 \cdot 2 + \frac{14}{3} = \frac{29}{3} > \frac{19}{2} = \mathbb{E}_b(T^2)(c_0) = \overline{\mathbb{E}}(T^2)(c_0)$. η_2 is strictly larger than $\overline{\mathbb{E}}(T^2)$ at c_0.

Listing 3.2: Example: Incompleteness of higher order UARnkSupM

```
if * then
        skip; // Branch (a)
        skip;
else
        if prob(1 - p) then // Branch (b)
                skip;
                skip;
        fi
fi
assert true;
```

Taking this exact example also shows that LRnkSupM of order 2 are unsound. If we take $\eta = (\underline{\mathbb{E}}(T), \underline{\mathbb{E}}(T^2))$, then it is easy to see that $\eta_2 \geq 1 + 2\underline{\mathbb{X}}\eta_1 + \underline{\mathbb{X}}\eta_2 = \mathrm{El}_2^2(\underline{\mathbb{X}}\eta)$ holds. By monotonicity of El^2 and $\underline{\mathbb{X}}$, we see that $\mathrm{El}^2(\underline{\mathbb{X}}\eta)$ is also a LRnkSupM of order 2, but calculating it shows that $\mathrm{El}_2^2(\underline{\mathbb{X}}\eta)(c_0) = 9 - \frac{2}{13} < 9 = \underline{\mathbb{E}}(T^2)$, contradicting the soundness. In particular, the tuple $\eta = (\underline{\mathbb{E}}(T), \underline{\mathbb{E}}(T^2))$ becomes smaller through the operator $\eta \mapsto \mathrm{El}^2(\underline{\mathbb{X}}\eta)$ and is thus just a pre-fixpoint of the operator but no fixed point. There exists pre-fixpoints, such as the least fixed point, which are smaller than $(\underline{\mathbb{E}}(T), \underline{\mathbb{E}}(T^2))$ making the method unsound.

3.2.3 Nonnegative Repulsing Supermartingales

In the previous two subsections we introduced supermartingales which give overapproximations of the expected reaching time or higher moments of it. The expected reaching time can be used to give an qualitative statement about almost sure reachability in the program. In this subsection we introduce a notion of supermartingale which gives quantitative information about reachability, concretely, it overapproximates the reachability probability. This can be used to refute almost sure reachability in programs.

Definition 3.2.8 (Nonnegative repulsing supermartingale (NNRepSupM))**.** Let Γ be a pCFG, $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region, c_0 a starting configuration and $\sigma \in \mathrm{Sch}_\Gamma$ a scheduler. A universally measurable function $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to [0,1]$ is called a *nonnegative repulsing σ-supermartingale* (σ-NNRepSupM) at c_0 if it satisfies

- $\eta(c, \sigma_\pi) = 1$ for all $c \in C$ and all paths $\pi c = c_0 c_1 \dots c$,
- $\eta(c, \sigma_\pi) \geq \mathbb{X}\eta(c, \sigma_\pi)$ for all paths $\pi c = c_0 c_1 \dots c$ with $c \notin C$.

A universally measurable function $\eta : L \times \mathbb{R}^V \to [0,1]$ is called an *upper nonnegative repulsing supermartingale* (UNNRepSupM) if it satisfies

- $\eta(c) = 1$ for all $c \in C$ and all paths $\pi c = c_0 c_1 \dots c$,
- $\eta(c) \geq \overline{\mathbb{X}}\eta(c)$ for all paths $\pi c = c_0 c_1 \dots c$ with $c \notin C$.

A lower nonnegative repulsing supermartingale (LNNRepSupM) is defined by replacing $\overline{\mathbb{X}}$ with $\underline{\mathbb{X}}$ in the above definition. The domains of these functions may be restricted to an invariant I containing c_0.

We show that NNRepSupMs give a complete and sound overapproximation of the reachability probability $\mathbb{P}^{\mathrm{reach}}_C(c_0, \sigma) := \Pr_{c_0,\sigma}(T^C < \infty)$ by employing martingale theory and fixed point theory in a similar fashion as we did in Subsection 3.2.1. First we show that the process $Y_i = \eta(C_i, \sigma_{C_0 \dots C_{i-1}})$ induced by a σ-NNRepSupM is indeed a supermartingale.

Remark. The probability measure $\Pr_{c,\sigma}$ on the set of all runs is simply $\nu^{c,\sigma}$ from Lemma 3.1.1. The term $T^C < \infty$ expresses the event $\{c_0 c_1 \dots \in (L \times \mathbb{R}^V)^\omega \mid T^C(c_0 c_1 \dots) < \infty\}$.

Lemma 3.2.18. *Let η be a σ-NNRepSupM for C at c_0 for some scheduler σ, starting configuration c_0 and target region C. Then the induced process Y_i given by $Y_i(c_0 c_1 \dots) = \eta(c_i, \sigma_{c_0 c_1 \dots c_{i-1}})$ is a supermartingale.*

Proof. We need to show that $\mathbb{E}_\sigma(Y_{i+1} \mid C_i, \dots, C_0) \leq Y_i$ holds. We distinguish two cases. Either $C_i \in C$, then $\mathbb{E}_\sigma(Y_{i+1} \mid C_i, \dots, C_0) \leq 1 = Y_i$ holds trivially. Otherwise we have for $C_i \notin C$:

$$\begin{aligned}
\mathbb{E}_\sigma(Y_{i+1} \mid C_i, \dots, C_0) &= \int \eta(C_{i+1}, \sigma_{C_0 \dots C_i}) \mu^\sigma_{C_0 \dots C_i}(\mathrm{d}C_{i+1}) \\
&= \int \eta(C_{i+1}, \sigma_{C_0 \dots C_i}) \mu_{C_i}^{\sigma_{C_0 \dots C_{i-1}}}(\mathrm{d}C_{i+1}) \\
&= \mathbb{X}\eta(C_i, \sigma_{C_0 \dots C_{i-1}}) \leq \eta(C_i, \sigma_{C_0 \dots C_{i-1}}) = Y_i.
\end{aligned}$$

□

Since the process Y_i is bounded by the value 1, the Optional Stopping Theorem may be applied directly (see condition (c) in Theorem 2.1.4) giving us the following result.

Theorem 3.2.19. *Let η be a σ-NNRepSupM for C at c_0 for some scheduler σ, starting configuration c_0 and target region C. Then we have $\eta(c_0, \sigma) \geq \mathbb{P}^{\text{reach}}_C(c_0, \sigma)$.*

Proof. Consider the associated stochastic process Y_i as before. It is bounded by the constant 1 function, so the Optional Stopping Theorem can be applied directly with respect to the first reaching time $T = T^C$. We get

$$\begin{aligned}
\eta(c_0, \sigma) = Y_0 &= \mathbb{E}_{c_0,\sigma}(Y_0) \geq \mathbb{E}_{c_0,\sigma}(Y_T) \\
&= \mathbb{E}_{c_0,\sigma}(Y_T \mid T < \infty)\Pr_{c_0,\sigma}(T < \infty) + \mathbb{E}_{c_0,\sigma}(Y_T \mid T = \infty)\Pr_{c_0,\sigma}(T = \infty) \\
&= 1 \cdot \Pr_{c_0,\sigma}(T < \infty) + \mathbb{E}_{c_0,\sigma}(Y_T \mid T = \infty)\Pr_{c_0,\sigma}(T = \infty) \\
&\geq \Pr_{c_0,\sigma}(T < \infty) = \mathbb{P}^{\text{reach}}_C(c_0, \sigma)
\end{aligned}$$

□

The name repulsing supermartingale is actually derived from a somewhat different definition. One can define a repulsing supermartingale to be a supermartingale that can take negative values and is guaranteed to be positive inside the target region[13]. If the supermartingale is also difference bounded (i.e. each step changes its expected value in a bounded manner) then it is possible to derive a probability bound on never reaching the target region if the starting value is negative. The idea here is that the supermartingale has the tendency to decrease over time and hence it tends to stay negative if it started negative. With a bound on the deviation per step and the use of concentration inequalities (such as Azuma-Hoeffding's inequality) the probability that it deviates into the positive region can be bounded. Because of this tendency to move away (being repulsed) from the positive region and hence the target region, the name *repulsing* supermartingale is used.

While this definition is usable, it does not admit a direct adaption to a fixed point theoretic setting. We opt for a definition that closely relates to the reachability probability, and for this it necessarily needs to have the same codomain, in particular, it should never be negative.

Theorem 3.2.20. *For any scheduler σ, starting configuration c_0 and target region C, the function $\mathbb{P}^{\text{reach}}_C$ is a σ-NNRepSupM at c_0.*

Proof. Let $\pi c = c_0 c_1 \dots c$ be any path and $T = T^C$ the first reaching time of C. If $c \in C$, then $\mathbb{P}^{\text{reach}}(c, \sigma_\pi) = 1$. For $c \notin C$, note that $T(cc' \dots) < \infty$ if and only if $T(c' \dots) < \infty$. From this we can conclude for $c \notin C$

$$\begin{aligned}
\mathbb{P}^{\text{reach}}(c, \sigma_\pi) &= \Pr_{c,\sigma_\pi}(T < \infty) = \Pr_{c,\sigma_\pi}(T < \infty \mid c \notin C) \\
&= \int \Pr_{c',\sigma_{\pi c}}(T < \infty)\mu_c^{\sigma_\pi}(\mathrm{d}c') \\
&= \int \mathbb{P}^{\text{reach}}_C(c', \sigma_{\pi c})\mu_c^{\sigma_\pi}(\mathrm{d}c') = \mathbb{X}\mathbb{P}^{\text{reach}}_C(c, \sigma_\pi)
\end{aligned}$$

□

These two theorems establish soundness and completeness of nonnegative repulsing supermartingales. The proof of the latter actually shows that the value of $\mathbb{P}_C^{\text{reach}}$ at configurations outside of C stays fixed under the nexttime operator $\mathbb{X}$. This can be used to establish the following fixed point characterization of $\mathbb{P}_C^{\text{reach}}$.

Theorem 3.2.21. *Let Γ be a pCFG and $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region. $\mathbb{P}_C^{\text{reach}}$ is the least fixed point of the monotone operator*

$$\Psi_C(\eta) = \mathbb{1}_C + \mathbb{1}_{\overline{C}} \mathbb{X}\eta.$$

Proof. From the completeness proof in Theorem 3.2.20, we easily see that $\mathbb{P}_C^{\text{reach}}$ is indeed a fixed point of Ψ_C. We need to show that it is the least fixed point. Towards this, consider the bounded reachability probability $\mathbb{P}_C^{\text{reach}\leq n}(c,\sigma) = \Pr_{c,\sigma}(T \leq n)$. The sequence of events $\{T < n\}$ is monotonically increasing, hence we have $\lim_{n\to\infty} \Pr_{c,\sigma}(T \leq n) = \Pr_{c,\sigma}(T < \infty) = \mathbb{P}_C^{\text{reach}}(c,\sigma)$. Now we show that $\Psi_C(\mathbb{P}_C^{\text{reach}\leq n}) = \mathbb{P}_C^{\text{reach}\leq n+1}$ holds and conclude from Cousot-Cousot's Theorem that $\mathbb{P}_C^{\text{reach}} = \Psi_C^{\omega}(\bot)$ is the least fixed point. For $c \in C$, we have $\Psi_C(\mathbb{P}_C^{\text{reach}\leq n})(c,\sigma) = 1 = \mathbb{P}_C^{\text{reach}\leq n+1}(c,\sigma)$. Consider $c \notin C$, then $T(cc'\dots) \leq n+1$ if and only if $T(c'\dots) \leq n$.

$$\begin{aligned}
\mathbb{P}_C^{\text{reach}\leq n+1}(c,\sigma) &= \Pr_{c,\sigma}(T \leq n+1 \mid c \notin C) \\
&= \int \Pr_{c',\sigma_c}(T \leq n)\mu_c^{\sigma}(\mathrm{d}c') \\
&= \int \mathbb{P}_C^{\text{reach}\leq n}(c',\sigma_c)\mu_c^{\sigma}(\mathrm{d}c') \\
&= \mathbb{X}\mathbb{P}_C^{\text{reach}\leq n}(c,\sigma) = \Psi_C(\mathbb{P}_C^{\text{reach}\leq n})(c,\sigma)
\end{aligned}$$

Note that $\mathbb{P}_C^{\text{reach}\leq 0} = \mathbb{1}_C = \Psi_C(\bot) \neq \bot$ (unless C is empty). The n-th element of the Cousot-Cousot sequence actually corresponds to $\mathbb{P}_C^{\text{reach}\leq n-1}$ and not to $\mathbb{P}_C^{\text{reach}\leq n}$. The latter would be the case, if we used a strict inequality to define $\mathbb{P}_C^{\text{reach}\leq n}$. □

It is well-known that reachability in deterministic structures like graphs and transition systems can be characterized as the least fixed point of an operator of the form $S \mapsto S \cup R(S)$ where S is some state/node set and $R(S)$ is the set of states/nodes immediately reachable within one step from S. So it is to no surprise that reachability in probabilistic systems admits an analogous characterization of similar form. It can be interpreted as the direct extension from boolean reasoning to quantitative reasoning.

Now we want to specialize the previous result to the case of a fixed scheduler, as done before for ARnkSupM.

Lemma 3.2.22. *Let σ be a history-independent scheduler, then $\mathbb{P}_{C,\sigma}^{\text{reach}} = \mathbb{P}_C^{\text{reach}}(\cdot,\sigma)$ is the least fixed point of the monotone operator*

$$\Psi_{C,\sigma}(\eta) = \mathbb{1}_C + \mathbb{1}_{\overline{C}} \mathbb{X}\eta(\cdot,\sigma).$$

where $\eta : (L \times \mathbb{R}^V) \to [0,1]$ is a scheduler-independent function.

Proof. Follows by the same arguments as in the proof of Theorem 3.2.21 and applying the equality $\sigma = \sigma_\pi$ for all valid paths π. □

Now we again consider the upper and lower versions of NNRepSupM. Similar to the case of $\mathbb{E}_C^{\text{steps}}$, $\mathbb{P}_C^{\text{reach}}$ also admits lower and upper ϵ-optimal schedulers which we shortly state in a lemma.

Lemma 3.2.23. *The functions $\mathbb{P}_C^{\text{reach}}$ and $\mathbb{P}_C^{\text{reach}\leq n}$ admit lower and upper ϵ-optimal schedulers. These schedulers can be chosen to be history-independent.*

Proof. The argument is analogous to the argument in Lemma 3.2.9 and a proof is found in [11, Lemma 3.2]. A reward function is chosen which gives a reward of 1 when entering the target region C, and 0 otherwise. The expected reward then coincides with the probability to reach the target region. □

The results regarding the extremizations of NNRepSupMs are in essence identical to the results proven in Subsection 3.2.1 about ARnkSupMs. It is again possible to derive them via either martingale theory or fixed point theory. We limit ourselves to the fixed point theoretic approach here but stress the fact that either theory may be used in a similar fashion as done for ARnkSupMs. Handling the lower version via martingale theory is more difficult, while the upper version follows the same proof principle as shown before.

Theorem 3.2.24. *Let Γ be a pCFG and $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region. $\overline{\mathbb{P}}_C^{\text{reach}}$ is the least fixed point of the monotone operator*

$$\overline{\Psi}_C(\eta) = \mathbb{1}_C + \mathbb{1}_{\overline{C}}\overline{\mathbb{X}}\eta.$$

Proof. Let η be any fixed point of $\overline{\Psi}_C$, then it satisfies $\eta = \mathbb{1}_C + \mathbb{1}_{\overline{C}}\overline{\mathbb{X}}\eta \geq \mathbb{1}_C + \mathbb{1}_{\overline{C}}\mathbb{X}\eta = \Psi_C(\eta)$ for all schedulers σ. We write $\eta(c,\sigma) = \eta(c)$. η is thus a pre-fixpoint of Ψ_C and by Knaster-Tarski it is larger than its least fixed point $\mathbb{P}_C^{\text{reach}}$ (see Theorem 3.2.21). Taking the supremum over all schedulers gives $\eta \geq \overline{\mathbb{P}}_C^{\text{reach}}$. We need to show that $\overline{\mathbb{P}}_C^{\text{reach}}$ itself is a fixed point to complete the proof. From the facts that $\mathbb{P}_C^{\text{reach}} = \mathbb{1}_C + \mathbb{1}_{\overline{C}}\mathbb{X}\mathbb{P}_C^{\text{reach}}$ (Theorem 3.2.21) holds and $\mathbb{P}_C^{\text{reach}}$ admits upper ϵ-optimal schedulers (Lemma 3.2.23), we can apply Lemma 3.2.2 and conclude $\overline{\mathbb{P}}_C^{\text{reach}} = \mathbb{1}_C + \mathbb{1}_{\overline{C}}\overline{\mathbb{X}}\,\overline{\mathbb{P}}_C^{\text{reach}} = \overline{\Psi}_C(\overline{\mathbb{P}}_C^{\text{reach}})$. □

To prove the analogous result for lower NNRepSupMs, we again make an additional assumption on the existence of ϵ-optimal schedulers as we did in Theorem 3.2.13.

Theorem 3.2.25. *Let Γ be a pCFG and $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region. $\underline{\mathbb{P}}_C^{\text{reach}}$ is a fixed point of the monotone operator*

$$\underline{\Psi}_C(\eta) = \mathbb{1}_C + \mathbb{1}_{\overline{C}}\underline{\mathbb{X}}\eta.$$

Moreover, if for every fixed point η of $\underline{\Psi}_C$ the function $\mathbb{X}\eta$ admits lower ϵ-optimal schedulers, then $\underline{\mathbb{P}}_C^{\text{reach}}$ is the least fixed point.

Proof. We first show that $\underline{\mathbb{P}}_C^{\text{reach}}$ is indeed a fixed point. As in the upper case, we use the fact that $\underline{\mathbb{P}}_C^{\text{reach}}$ admits lower ϵ-optimal schedulers and apply Lemma 3.2.2 to the equality $\mathbb{P}_C^{\text{reach}} = \mathbb{1}_C + \mathbb{1}_{\overline{C}}\mathbb{X}\mathbb{P}_C^{\text{reach}}$ from Theorem 3.2.21 to get the desired result.

Now let η be an arbitrary fixed point of $\underline{\Psi}_C$. For any $\epsilon > 0$, let σ^ϵ be a lower ϵ-optimal scheduler for $\mathbb{X}\eta$. Construct a universally measurable scheduler σ satisfying $\sigma(\pi) = \sigma^{\epsilon 2^{-|\pi|}}(\pi)$

such that the scheduler σ picks an $\epsilon 2^{-(n+1)}$-optimal action for $\mathbb{X}\eta$ at the n-th step. In particular, the scheduler σ_π is lower $\epsilon 2^{-(|\pi|+1)}$-optimal for $\mathbb{X}\eta$. We show by induction that $\mathbb{P}_C^{\text{reach}\leq n}(\cdot, \sigma_\pi) \leq \eta + \sum_{i=|\pi|+1}^{n+|\pi|} \epsilon 2^{-i}$ holds for all paths π and $n \in \mathbb{N}$. For $n = 0$ we have $\mathbb{P}_C^{\text{reach}\leq 0} = \mathbb{1}_C \leq \mathbb{1}_C + \mathbb{1}_{\overline{C}}\mathbb{X}\eta = \eta$ for all paths π. Now we consider the case $n+1$. For $c \in C$ the inequality holds true. For $c \notin C$ we have:

$$\begin{aligned}
\mathbb{P}_C^{\text{reach}\leq n+1}(c, \sigma_\pi) &= \mathbb{X}\mathbb{P}_C^{\text{reach}\leq n}(c, \sigma_\pi) \\
&= \int \mathbb{P}_C^{\text{reach}\leq n}(c', \sigma_{\pi c})\mu_c^{\sigma_\pi}(\mathrm{d}c') \\
&\leq \int \eta(c')\mu_c^{\sigma_\pi}(\mathrm{d}c') + \sum_{i=|\pi c|+1}^{n+|\pi c|} \epsilon 2^{-i} \\
&= \mathbb{X}\eta(c, \sigma_\pi) + \sum_{i=|\pi|+2}^{(n+1)+|\pi|} \epsilon 2^{-i} \\
&\leq \underline{\mathbb{X}}\eta(c) + \epsilon 2^{-(|\pi|+1)} + \sum_{i=|\pi|+2}^{(n+1)+|\pi|} \epsilon 2^{-i} \\
&= \eta(c) + \sum_{i=|\pi|+1}^{(n+1)+|\pi|} \epsilon 2^{-i}
\end{aligned}$$

From the second to third line we use the induction hypothesis and the linearity of the integral. From the fourth to fifth line we use the fact that σ_π is $\epsilon 2^{-(|\pi|+1)}$-optimal. For the empty path and $n \to \infty$ we get

$$\underline{\mathbb{P}}_C^{\text{reach}}(c) \leq \mathbb{P}_C^{\text{reach}}(c, \sigma) = \lim_{n\to\infty} \mathbb{P}_C^{\text{reach}\leq n}(c, \sigma) \leq \eta(c) + \sum_{i=1}^{\infty} \epsilon 2^{-i} = \eta(c) + \epsilon.$$

Since $\epsilon > 0$ was chosen arbitrarily, we conclude $\underline{\mathbb{P}}^{\text{reach}}(c) \leq \eta(c)$. □

The proof actually shows that $\underline{\mathbb{P}}_C^{\text{reach}}$ is the least fixed point η of $\underline{\Psi}_C$ such that $\mathbb{X}\eta$ admits lower ϵ-optimal schedulers. However, we do not have any example of a fixed point which does not satisfy this property. In particular, we do not have an example where there is a fixed point smaller than $\underline{\mathbb{P}}_C^{\text{reach}}$ that does not admit ϵ-optimal schedulers.

The operators $\overline{\Psi}_C$ and $\underline{\Psi}_C$ make use of Bellman's principle of optimality for dynamic programming which is widely applied in the context of optimal control (e.g. [25, 18, 26]). It states that an optimal policy or strategy for some starting state must remain optimal for the resulting state after the first transition. These operators may be understood as maximizing or minimizing the expected rewards of the system as it progresses. Here a reward of 1 is obtained whenever the system enters the target region C. If we assume the system to stop after entering once (or potentially loop indefinitely in some dummy state afterwards), then the reward can be obtained only once and the expected reward coincides with the probability to reach C. A similar viewpoint can be taken on ARnkSupM, where a reward of 1 is accumulated whenever the system performs a step outside of the target region. The

expected reward is then equivalent to the expected reaching time, assuming the system does not reenter the target region afterwards. This exact reduction to the expected reward model is used to prove the existence of ϵ-optimal schedulers in [18].

Before we conclude this subsection, we shortly introduce another use of UNNRepSupM to *underapproximate* reachability probability proposed by Chatterjee et al.[13]. Chatterjee et al. made use of the concept of probabilistic invariants and coupled them with qualitative reasoning via ranking supermartingales. The idea is to find a predicate which is strong enough to infer almost sure reachability but potentially too restrictive to be a pure invariant. Instead, there is a certain probability that the predicate actually fails to be an invariant. A case where such reasoning can be useful, is when dealing with unbounded distributions such as normal distributions. A pure invariant has to consider all possible outcomes of sampling such a distribution which in generally gives a trivial invariant of the form $-\infty < x < \infty$. On the other hand, the probability that a sample from for example $N(0,1)$ (normal distribution with mean 0 and variance 1) exceeds the finite value 10 is around the order of 10^{-24}. For qualitative reasoning an invariant such as $x < 10$ may be strong enough to show almost sure reachability whereas $x < \infty$ may fail to do so. In the example 3.3 the probabilistic invariant $-50 < x < 50$ is sufficient to guarantee reachability under every scheduler but the pure invariant fails to do so for any scheduler.

Listing 3.3: Example: Probabilistic invariants

```
x := norm(0, 1);
if (x < 0) then
        x := -x;
fi
z := 0;
n := ndet(Int[1, 10000]);
while (n > 0) do
        z := z + x;
        x := x/2;
        n := n + 1;
od
assert z <= 100;
```

To find such probabilistic invariants, one can consider a candidate invariant P and over-approximate the probability to reach $\overline{P}$ via UNNRepSupMs. This gives an upper bound on the probability that the system fails to satisfy P under any scheduler, or alternatively, a lower bound on the probability that P is an invariant. This can then be used to lower bound the reachability probability to some target region C via ranking supermartingales for C supported by P, since whenever P is in invariant, the system is guaranteed to reach the target C.

3.2.4 Nonnegative Repulsing δ-Supermartingales

A slight variation of the previously introduced nonnegative repulsing supermartingales (NNRepSupM) in Subsection 3.2.3 can be used to reason about bounded reachability probabilities, that is, to reason about $\mathrm{P}(T^C \leq n)$ instead of $\Pr(T^C < \infty)$. This variation, in a sense, gives the counterpart to reasoning about tail probabilities via higher order ranking supermartingales (Subsection 3.2.2). The idea is to relax the conditions on NNRepSupM which have to strictly decrease in expectation, and allow them to also increase in expectation but in a bounded manner. The bound is given by a value $\delta \geq 0$.

Definition 3.2.9 (δ-NNRepSupM)**.** Let Γ be a pCFG, $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region, c_0 a starting configuration and $\sigma \in \mathrm{Sch}_\Gamma$ a scheduler. Let $\delta \geq 0$ be given. A universally measurable function $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to [0,1]$ is called a *nonnegative δ-repulsing σ-supermartingale* ((δ,σ)-NNRepSupM) at c_0 if it satisfies

- $\eta(c,\sigma_\pi) = 1$ for all $c \in C$ and all paths $\pi c = c_0 c_1 \dots c$,
- $\eta(c,\sigma_\pi) \geq \mathbb{X}\eta(c,\sigma_\pi) - \delta$ for all paths $\pi c = c_0 c_1 \dots c$ with $c \notin C$.

A universally measurable function $\eta : L \times \mathbb{R}^V \to [0,1]$ is called an *upper nonnegative δ-repulsing supermartingale* (δ-UNNRepSupM) if it satisfies

- $\eta(c) = 1$ for all $c \in C$ and all paths $\pi c = c_0 c_1 \dots c$,
- $\eta(c) \geq \overline{\mathbb{X}}\eta(c) - \delta$ for all paths $\pi c = c_0 c_1 \dots c$ with $c \notin C$.

A lower nonnegative repulsing δ-supermartingale (δ-LNNRepSupM) is defined by replacing $\overline{\mathbb{X}}$ with $\underline{\mathbb{X}}$ in the above definition. As usual, the domains of these functions may be restricted to an invariant I containing c_0.

We want to show that $\eta(c,\sigma_\pi) + n\delta \geq \mathbb{P}_C^{\mathrm{reach}\leq n}(c,\sigma_\pi)$ holds for paths $\pi = c_0 c_1 \dots c$. The extra summand $n\delta$ here gives an uncertainty introduced through the δ-relaxation. Smaller values of δ yield better bounds, and in particular for $\delta = 0$ we get back the inequality of a regular NNRepSupM. The purpose of the δ-relaxation is to increase the set of functions satisfying the supermartingale condition. This may give smaller solutions for η which can give tighter bounds for small time horizons. Increasing the set of solutions can be particularly useful when automatic template synthesis methods are used.

Lemma 3.2.26. *Let η be a (δ,σ)-NNRepSupM at c_0 for some scheduler σ. Let $Z_i = Y_i - \delta \cdot i$ define a stochastic process where Y_i is the stochastic process induced by η. The stopped process Z_i^T is a supermartingale with respect to the canonical process C_i.*

Proof. We have to show that $\mathbb{E}_\sigma(Z_{i+1}^T \mid C_i,\dots,C_0) \leq Z_i^T$ holds. We do a case distinction on the event $\{T < i+1\}$. If $T < i+1$, then

$$\begin{aligned}\mathbb{E}_\sigma(Z_{i+1}^T \mid C_i,\dots,C_0) &= \mathbb{E}_\sigma(Z_T \mid C_i,\dots,C_T,\dots,C_0)\\ &= Z_T = Z_i^T\end{aligned}$$

If $T \geq i+1$, then

$$\begin{aligned}
\mathbb{E}_\sigma(Z_{i+1}^T \mid C_i, \ldots, C_0) &= \mathbb{E}_\sigma(Z_{i+1} \mid C_i, \ldots, C_0) \\
&= \mathbb{E}_\sigma(Y_{i+1} \mid C_i, \ldots, C_0) - \delta \cdot (i+1) \\
&= \mathbb{X}\eta(C_i, \sigma_{C_0 \ldots C_{i-1}}) - \delta \cdot (i+1) \\
&\leq \eta(C_i, \sigma_{C_0 \ldots C_{i-1}}) + \delta - \delta \cdot (i+1) \\
&= \eta(C_i, \sigma_{C_0 \ldots C_{i-1}}) - \delta \cdot i = Z_i = Z_i^T
\end{aligned}$$

□

Note that the process Z_i defined as above is not bounded and the Optional Stopping Theorem may not be applied directly to the stopping time T. However, stopping the process at any fixed time $n \in \mathbb{N}$, we get the desired result.

Theorem 3.2.27. *Let η be a (δ, σ)-NNRepSupM at c_0, then $\eta(c_0, \sigma) + \delta \cdot n \geq \mathbb{P}^{\text{reach} \leq n}(c_0, \sigma)$.*

Proof. Let Z_i^T be the supermartingale from Lemma 3.2.26. For any fixed $n \in \mathbb{N}$, we can apply the Optional Stopping Theorem to the process Z_i^T stopped at n and get

$$\begin{aligned}
\eta(c_0, \sigma) &= Z_0 = \mathbb{E}_{c_0,\sigma}(Z_0^T) \geq \mathbb{E}_{c_0,\sigma}(Z_n^T) \\
&= \mathbb{E}_{c_0,\sigma}(Z_T \mid T \leq n)\text{Pr}_{c_0,\sigma}(T \leq n) + \mathbb{E}_{c_0,\sigma}(Z_n \mid T > n)\text{Pr}_{c_0,\sigma}(T > n) \\
&= \mathbb{E}_{c_0,\sigma}(Y_T - \delta \cdot T \mid T \leq n)\text{Pr}_{c_0,\sigma}(T \leq n) + \mathbb{E}_{c_0,\sigma}(Y_n - \delta \cdot n \mid T > n)\text{Pr}_{c_0,\sigma}(T > n) \\
&= (1 - \delta \cdot \mathbb{E}_{c_0,\sigma}(T \mid T \leq n))\text{Pr}_{c_0,\sigma}(T \leq n) + (\mathbb{E}_{c_0,\sigma}(Y_n \mid T > n) - \delta \cdot n)\text{Pr}_{c_0,\sigma}(T > n) \\
&\geq (1 - \delta \cdot n)\text{Pr}_{c_0,\sigma}(T \leq n) + (0 - \delta \cdot n)\text{Pr}_{c_0,\sigma}(T > n) \\
&= \text{Pr}_{c_0,\sigma}(T \leq n) - \delta \cdot n(\text{Pr}_{c_0,\sigma}(T \leq n) + \text{Pr}_{c_0,\sigma}(T > n)) \\
&= \text{Pr}_{c_0,\sigma}(T \leq n) - \delta \cdot n = \mathbb{P}_C^{\text{reach} \leq n}(c_0, \sigma) - \delta \cdot n
\end{aligned}$$

□

We now consider the upper version. The result and proof are similar to Theorem 3.2.5 for ARnkSupMs and to Theorem 3.2.19 for NNRepSupMs.

Theorem 3.2.28. *Let η be a δ-UNNRepSupM at c_0, then $\eta(c_0) + \delta \cdot n \geq \overline{\mathbb{P}}_C^{\text{reach} \leq n}(c_0)$.*

Proof. We set $\eta(c, \sigma) = \eta(c)$, then for every path $\pi c = c_0 c_1 \ldots c$ we have $\eta(c, \sigma_\pi) = \eta(c) \geq \overline{\mathbb{X}}\eta(c) - \delta \geq \mathbb{X}\eta(c, \sigma_\pi) - \delta$. Hence, η is a (δ, σ)-NNRepSupM for any scheduler σ and thus by Theorem 3.2.27 it satisfies $\eta(c_0) + \delta \cdot n \geq \mathbb{P}_C^{\text{reach} \leq n}(c_0, \sigma)$. Taking the supremum over all schedulers gives the desired result. □

As before, a fixed point theoretic approach can derive similar results. We do not further develop the fixed point theoretic approach here but instead refer to Takisaka et al.[11] who give a characterization for the upper and lower version only. The characterization as well as the proofs are a simple and straightforward adaption of the corresponding proofs for NNRepSupM in Subsection 3.2.3.

3.2.5 γ-scaled Submartingales

Reachability is usually seen as a least fixed point of a suitable operator and by employing Knaster-Tarski one naturally can find overapproximations via pre-fixpoints. In the previous sections this gave rise to many supermartingale methods which are able to give sound overapproximations on either reaching times or reachability probabilities. On the other hand, Cousot-Cousot gives a way to underapproximate least fixed points by constructing an increasing, infinite sequence that converges to the least fixed point. Such a sequence is calculated iteratively and is especially hard to compute on infinite state spaces. Furthermore, within finite time it can only give sound but incomplete underapproximations on most systems. For these reasons, we would like to use Knaster-Tarski instead to obtain sound underapproximations. Towards this, we need to reframe reachability as a greatest fixed point such that any post-fixpoint will be a sound underapproximation. Unlike before, this way we obtain so called submartingales instead of supermartingales. Reachability does not naturally admit a greatest fixed point characterization, instead, we consider discounted reachability where longer paths contribute less to the reachability than shorter ones. This can be achieved by adding a discount factor γ to the natural fixed point characterization of the reachability probability. It turns out that by adding such a factor, all previously existing fixed points will get contracted to a single fixed point which is simultaneously the least and the greatest fixed point, allowing the usage of Knaster-Tarski to obtain overapproximations as well as underapproximations. However, through the addition of the discount factor, this method necessarily becomes incomplete.

Definition 3.2.10 (γ-scaled Submartingale (γ-SclSubM)[11, 15])**.** Let Γ be a pCFG, $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region, c_0 a starting configuration and $\sigma \in \mathrm{Sch}_\Gamma$ a scheduler. Let $\gamma \in (0,1)$ be a scaling factor. A universally measurable function $\eta : (L \times \mathbb{R}^V) \times \mathrm{Sch}_\Gamma \to [0,1]$ is called a *γ-scaled σ-submartingale* (γ-Scl-σ-SubM) at c_0 if it satisfies

- $\eta(c, \sigma_\pi) = 1$ for all $c \in C$ and all paths $\pi c = c_0 c_1 \dots c$,
- $\eta(c, \sigma_\pi) \leq \gamma \mathbb{X} \eta(c, \sigma_\pi)$ for all paths $\pi c = c_0 c_1 \dots c$ with $c \notin C$.

A universally measurable function $\eta : L \times \mathbb{R}^V \to [0,1]$ is called an *upper γ-scaled submartingale* (U-γ-SclSubM) if it satisfies

- $\eta(c) = 1$ for all $c \in C$ and all paths $\pi c = c_0 c_1 \dots c$,
- $\eta(c) \leq \gamma \overline{\mathbb{X}} \eta(c)$ for all paths $\pi c = c_0 c_1 \dots c$ with $c \notin C$.

A lower γ-scaled submartingale (L-γ-SclSubM) is defined by replacing $\overline{\mathbb{X}}$ with $\underline{\mathbb{X}}$ in the above definition. The domains of these functions may be restricted to an invariant I containing c_0.

The above definition is very similar to that of NNRepSupM. There are just two differences, namely, the scaling factor γ and the direction of the inequality. It is easy to see that such γ scaled submartingale induces a submartingale if stopped at T^C, simply by the fact that $\eta \leq \gamma \mathbb{X} \eta \leq \mathbb{X} \eta$ holds and Lemma 3.2.3. However, we look at the modified stochastic process $Z_i = \gamma^i Y_i$ instead.

Lemma 3.2.29. *Let η be a γ-scaled σ-submartingale at c_0 for C for some scaling factor $\gamma \in (0,1)$, scheduler σ, starting configuration c_0 and target region C. Let $\{Y_i\}_{i\in\mathbb{N}}$ be the induced process according to Definition 3.2.4. The stochastic process $Z_i = \gamma^i Y_i$ stopped at $T = T^C$ is a submartingale.*

Proof. We need to show that $\mathbb{E}_\sigma(Z^T_{i+1} \mid C_0, \ldots, C_i) \geq Z^T_i$ holds. We make a case distinction based on the event $\{T \leq i\}$. For $T \leq i$ we have

$$\mathbb{E}_\sigma(Z^T_{i+1} \mid C_0, \ldots, C_i) = \mathbb{E}_\sigma(Z_T \mid C_0, \ldots, C_T \ldots, C_i) = Z_T = Z^T_i$$

For $T > i$ we have

$$\begin{aligned}\mathbb{E}_\sigma(Z^T_{i+1} \mid C_0, \ldots, C_i) &= \mathbb{E}_\sigma(Z_{i+1} \mid C_0, \ldots, C_i) \\ &= \mathbb{E}_\sigma(\gamma^{i+1} Y_{i+1} \mid C_0, \ldots, C_i) = \gamma^{i+1}\mathbb{E}_\sigma(Y_{i+1} \mid C_0, \ldots, C_i) \\ &= \gamma^i \gamma \mathbb{X}\eta(C_i, \sigma_{C_0,\ldots,C_{i-1}}) \geq \gamma^i \eta(C_i, \sigma_{C_0,\ldots,C_{i-1}}) = Z^T_i\end{aligned}$$

□

Now we want to apply the Optional Stopping Theorem. The process Z^T_i is obviously bounded by $0 \leq Z^T_i \leq 1$, so we may use the Optional Stopping Theorem by condition (c). This gives the following main theorem.

Theorem 3.2.30. *Let η be a γ-scaled σ-submartingale as in Lemma 3.2.29. Then $\eta(c_0, \sigma) \leq \mathbb{P}^{\text{reach}}_C(c_0, \sigma)$.*

Proof. We take the process Z^T_i from Lemma 3.2.29 and apply the Optional Stopping Theorem (by condition (c)). This gives

$$\begin{aligned}\eta(c_0, \sigma) &= \mathbb{E}_{c_0,\sigma}(Z^T_0) \leq \mathbb{E}_{c_0,\sigma}(Z_T) = \mathbb{E}_{c_0,\sigma}(\gamma^T Y_T) \\ &= \mathbb{E}_{c_0,\sigma}(\gamma^T Y_T \mid T < \infty)\Pr_{c_0,\sigma}(T < \infty) + \mathbb{E}_{c_0,\sigma}(\gamma^T Y_T \mid T = \infty)\Pr_{c_0,\sigma}(T = \infty) \\ &= \mathbb{E}_{c_0,\sigma}(\gamma^T \mid T < \infty)\Pr_{c_0,\sigma}(T < \infty) + 0 \leq \Pr_{c_0,\sigma}(T < \infty)\end{aligned}$$

We note that $\gamma^T \leq 1$ holds and for $T = \infty$ we get $\gamma^T Y_T = 0$ since $\gamma < 1$. □

The above theorem shows two crucial points about γ-scaled submartingales. First, it is essential that $\gamma < 1$ holds. For $\gamma = 1$, we are not able to conclude $\mathbb{E}_\sigma(\gamma^T Y_T \mid T = \infty) = 0$ and hence the proof fails. Secondly, since γ needs to be strictly less than 1, we have $\mathbb{E}_\sigma(\gamma^T \mid 0 < T < \infty) < 1$, which in generally leads to a gap between η and $\mathbb{P}^{\text{reach}}_C$ (unless $T = 0$ holds). The higher the expected reaching time, the larger the approximation gap becomes. The method is sound but not complete and in particular the function $\Pr(T < \infty)$ is no γ-scaled submartingale.

If put into the framework of the general MDP setting, then γ-SclSubMs correspond to the expected *discounted* reward model and may be reframed as such. This discount or scaling factor actually lets all fixed points of the corresponding Bellman optimality equation coincide. We will later give a fixed point characterization based on this Bellman equation

and see that the corresponding operator is a contraction on the space of universally measurable functions. By Banach's fixed point theorem, we see that this operator has a unique fixed point as we already pointed out in the introduction.

We want to get analogous results for both versions as done for other martingale-based methods. For supermartingales the upper versions tended to be easier to prove while the lower versions made use of extra assumptions. For submartingales it is the opposite; the lower version is the easier one to prove. This discrepancy will appear again in Chapter 4 about automatic synthesis, where only lower submartingales and upper supermartingales admit straightforward synthesis procedures. Let us consider the lower version first.

Theorem 3.2.31. *Let η be a lower γ-scaled submartingale at c_0 for some target region C and scaling factor $\gamma \in (0,1)$. Then $\eta(c_0) \leq \mathbb{P}^{\text{reach}}_{C}(c_0)$ holds.*

Proof. For any arbitrary scheduler σ and all paths $\pi c = c_0 c_1 \dots c$ we have $\eta(c, \sigma_\pi) = \eta(c) \leq \gamma \underline{\mathbb{X}} \eta(c) \leq \gamma \mathbb{X} \eta(c, \sigma_\pi)$. Hence, η is a γ-scaled σ-submartingale for any scheduler σ, and therefore $\eta(c_0) \leq \mathbb{P}^{\text{reach}}_{C}(c_0, \sigma)$ holds for all schedulers. Taking the infimum on both sides proves the statement. □

We will handle the upper case in the fixed point theoretic framework. First, let us define the fixed point characterization of γ-Scl-σ-SubM. The proof of Theorem 3.2.30 actually gives a stronger result than stated, namely, it gives a concrete bound on the largest possible γ-SclSubM. The mapping $(c, \sigma) \mapsto \mathbb{E}_{c,\sigma}(\gamma^T \mid T < \infty)\Pr_{c,\sigma}(T < \infty) = \mathbb{E}_{c,\sigma}(\gamma^T)$ is an upper bound. We now show that it is indeed the tightest bound by characterizing it as the unique fixed point of the corresponding fixed point theoretic operator.

Theorem 3.2.32. *Let Γ be a pCFG, $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region and $\gamma \in (0,1)$ a scaling factor. $\mathbb{P}^{\text{reach}}_{C,\gamma}(c, \sigma) := \mathbb{E}_{c,\sigma}(\gamma^T)$ is the unique fixed point of the monotone operator*

$$\Psi_{C,\gamma}(\eta) = \mathbb{1}_C + \gamma \mathbb{1}_{\overline{C}} \mathbb{X} \eta.$$

Proof. We proceed in two steps. First we show that $\mathbb{P}^{\text{reach}}_{C,\gamma}$ is indeed a fixed point and then we show that $\Psi_{C,\gamma}$ is a contraction. From Banach's fixed-point theorem we conclude that it must be the unique fixed point. First note that for all $c \in C$, we have $\mathbb{P}^{\text{reach}}_{C,\gamma}(c, \sigma) = \mathbb{E}_{c,\sigma}(\gamma^T) = \mathbb{E}_{c,\sigma}(\gamma^0) = 1 = \Psi_{C,\gamma}(\mathbb{P}^{\text{reach}}_{C,\gamma})(c, \sigma)$. For $c \notin C$, note that $T(cc' \dots) = T(c' \dots) + 1$ holds. This gives

$$\begin{aligned}
\mathbb{P}^{\text{reach}}_{C,\gamma}(c, \sigma) &= \mathbb{E}_{c,\sigma}(\gamma^T) = \mathbb{E}_{c,\sigma}(\gamma^T \mid c \notin C) \\
&= \int \mathbb{E}_{c', \sigma_c}(\gamma^{T+1}) \mu^\sigma_c(\mathrm{d}c') \\
&= \gamma \cdot \int \mathbb{E}_{c', \sigma_c}(\gamma^T) \mu^\sigma_c(\mathrm{d}c') \\
&= \gamma \cdot \int \mathbb{P}^{\text{reach}}_{C,\gamma}(c', \sigma_c) \mu^\sigma_c(\mathrm{d}c') \\
&= \gamma \mathbb{X} \mathbb{P}^{\text{reach}}_{C,\gamma}(c, \sigma) = \Psi_{C,\gamma}(\mathbb{P}^{\text{reach}}_{C,\gamma})(c, \sigma).
\end{aligned}$$

Next we show that the operator is a contraction on the Banach space of universally measurable functions into $[0,1]$ with the supremum norm $\|\eta\| = \sup_{c,\sigma} \eta(c, \sigma)$. It is easy to see

that this space is indeed a complete metric space, because the functions codomain $[0,1]$ is complete and the limit of measurable functions is again measurable. For any two such functions η_1 and η_2, we have:

$$\begin{aligned}\|\Psi_{C,\gamma}(\eta_1) - \Psi_{C,\gamma}(\eta_2)\| &= \|\gamma \mathbb{1}_{\overline{C}} \mathbb{X}\eta_1 - \gamma \mathbb{1}_{\overline{C}} \mathbb{X}\eta_2\| \\ &= \gamma \|\mathbb{1}_{\overline{C}} \mathbb{X}(\eta_1 - \eta_2)\| \leq \gamma \|\mathbb{X}(\eta_1 - \eta_2)\| \\ &\leq \gamma \|\eta_1 - \eta_2\|.\end{aligned}$$

Since $\gamma < 1$, $\Psi_{C,\gamma}$ is a contraction. We made use of the fact that $\mathbb{X}$ is a linear operator and the underlying integral is taken over a space with total measure 1. □

This proof does not directly apply to $\underline{\mathbb{X}}$ and $\overline{\mathbb{X}}$ because these operators are not linear but rather super- and subadditive. However, slight modifications to the proof give us analogous results. Before we go into the upper and lower version, we first establish the existence of ϵ-optimal schedulers.

Lemma 3.2.33. *The function $\mathbb{P}^{\text{reach}}_{C,\gamma}$ admits lower and upper ϵ-optimal schedulers.*

Proof. As mentioned previously, this function can be expressed in the discounted average reward model for general MDPs with positive, bounded rewards and as such admits ϵ-optimal schedulers[18, 11]. A full proof is given in [11, Lemma 3.2]. □

Theorem 3.2.34. *Let Γ be a pCFG, $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region and $\gamma \in (0,1)$ a scaling factor. $\underline{\mathbb{P}}^{\text{reach}}_{C,\gamma}(c) = \underline{\mathbb{E}}_c(\gamma^T)$ is the unique fixed point of the monotone operator*

$$\underline{\Psi}_{C,\gamma}(\eta) = \mathbb{1}_C + \gamma \mathbb{1}_{\overline{C}} \underline{\mathbb{X}}\eta.$$

Proof. From Theorem 3.2.32 we have $\Psi_{C,\gamma}(\mathbb{P}^{\text{reach}}_{C,\gamma}) = \mathbb{1}_C + \gamma \mathbb{1}_{\overline{C}} \mathbb{X}\mathbb{P}^{\text{reach}}_{C,\gamma} = \mathbb{P}^{\text{reach}}_{C,\gamma}$. From the existence of ϵ-optimal schedulers and Lemma 3.2.2 we can take the infimum over all schedulers in the previous equation and get $\underline{\Psi}_{C,\gamma}(\underline{\mathbb{P}}^{\text{reach}}_{C,\gamma}) = \mathbb{1}_C + \gamma \mathbb{1}_{\overline{C}} \underline{\mathbb{X}}\, \underline{\mathbb{P}}^{\text{reach}}_{C,\gamma} = \underline{\mathbb{P}}^{\text{reach}}_{C,\gamma}$, so we have a fixed point. Next we show that the operator is contractive on the space of universally measurable functions into $[0,1]$. Let η_1 and η_2 be two such functions. We define the sets $C_P = \{c \in L \times \mathbb{R}^V \mid \underline{\mathbb{X}}\eta_1(c) - \underline{\mathbb{X}}\eta_2(c) > 0\}$ and $C_N = \{c \in L \times \mathbb{R}^V \mid \underline{\mathbb{X}}\eta_1(c) - \underline{\mathbb{X}}\eta_2(c) \leq 0\}$. These two sets partition the configuration space. Consider the following inequality:

$$\|\underline{\Psi}_{C,\gamma}(\eta_1) - \underline{\Psi}_{C,\gamma}(\eta_2)\| = \|\gamma \mathbb{1}_{\overline{C}} \underline{\mathbb{X}}\eta_1 - \gamma \mathbb{1}_{\overline{C}} \underline{\mathbb{X}}\eta_2\| \leq \gamma \|\underline{\mathbb{X}}\eta_1 - \underline{\mathbb{X}}\eta_2\| \tag{3.1}$$

We can do a case distinction with respect to the sets C_P and C_N because we have

$$\gamma \|\underline{\mathbb{X}}\eta_1 - \underline{\mathbb{X}}\eta_2\| = \gamma \max(\|\mathbb{1}_{C_P}(\underline{\mathbb{X}}\eta_1 - \underline{\mathbb{X}}\eta_2)\|, \|\mathbb{1}_{C_N}(\underline{\mathbb{X}}\eta_1 - \underline{\mathbb{X}}\eta_2)\|), \tag{3.2}$$

since the maximum value of the difference is either taken in C_P or in C_N (or both). First we consider the second term, i.e. the case of C_N. The difference $\underline{\mathbb{X}}\eta_1 - \underline{\mathbb{X}}\eta_2$ is nonpositive, hence the norm is largest when the difference takes the smallest, negative value. Thus we have

$$\|\mathbb{1}_{C_N}(\underline{\mathbb{X}}\eta_1 - \underline{\mathbb{X}}\eta_2)\| \leq \|\mathbb{1}_{C_N} \underline{\mathbb{X}}(\eta_1 - \eta_2)\| \leq \|\eta_1 - \eta_2\|.$$

For the case of C_P, note that $c \in C_P$ implies $\underline{\mathbb{X}}\eta_2(c) - \underline{\mathbb{X}}\eta_1(c) \leq 0$. Using the same argument as above, we can conclude

$$\|\mathbb{1}_{C_P}(\underline{\mathbb{X}}\eta_1 - \underline{\mathbb{X}}\eta_2)\| = \|\mathbb{1}_{C_P}(\underline{\mathbb{X}}\eta_2 - \underline{\mathbb{X}}\eta_1)\| \leq \|\mathbb{1}_{C_P}\underline{\mathbb{X}}(\eta_2 - \eta_1)\| \leq \|\eta_2 - \eta_1\| = \|\eta_1 - \eta_2\|.$$

Applying these results to Equations 3.1 and 3.2, we get

$$\|\underline{\Psi}_{C,\gamma}(\eta_1) - \underline{\Psi}_{C,\gamma}(\eta_2)\| \leq \gamma\|\eta_1 - \eta_2\|,$$

hence $\underline{\Psi}_{C,\gamma}$ is a contraction. □

Because $\underline{\Psi}_{C,\gamma}$ is contractive, we also know that its unique fixed point may be computed via the Cousot-Cousot iteration within ω many steps. For LNNRepSupM and LARnkSupM the convergence after ω steps was not guaranteed. Using this convergence property, we could actually remove the usage of Lemma 3.2.33 and instead define the fixed point as the limit of the Cousot-Cousot sequence.

Theorem 3.2.35. *Let Γ be a pCFG, $C \in \mathcal{B}(L \times \mathbb{R}^V)$ a target region and $\gamma \in (0,1)$ a scaling factor. $\overline{\mathbb{P}}^{\text{reach}}_{C,\gamma}(c) = \overline{\mathbb{E}}_c(\gamma^T)$ is the unique fixed point of the monotone operator*

$$\overline{\Psi}_{C,\gamma}(\eta) = \mathbb{1}_C + \gamma\mathbb{1}_{\overline{C}}\overline{\mathbb{X}}\eta.$$

Proof. The proof follows the same arguments as used in the proof of the lower version in Theorem 3.2.34. The only argument that changes slightly is the use of the partition into C_P and C_N. On the set C_P we have $\|\overline{\mathbb{X}}\eta_1 - \overline{\mathbb{X}}\eta_2\| \leq \|\overline{\mathbb{X}}(\eta_1 - \eta_2)\|$ and on the set C_N we can simply consider the difference $\|\overline{\mathbb{X}}\eta_2 - \overline{\mathbb{X}}\eta_1\|$ and argue as before. □

As mentioned in the introduction, we can simply apply Knaster-Tarski to any of the three variations of γ-scaled submartingales and get that any postfix point is a underapproximation of the unique fixed point. Note that $\overline{\mathbb{P}}^{\text{reach}}_{C,\gamma} \leq \overline{\mathbb{P}}^{\text{reach}}_{C}$ and $\underline{\mathbb{P}}^{\text{reach}}_{C,\gamma} \leq \underline{\mathbb{P}}^{\text{reach}}_{C}$ hold, so every underapproximation via upper or lower γ-scaled submartingales is sound. We want to stress the strong correspondence between submartingales and post-fixpoints, supermartingales and pre-fixpoints and lastly between martingales and fixed points. This correspondence is the key component making both the fixed point theoretic and the martingale theoretic approach so similar in appliance and results.

3.2.6 Martingales for Recurrence

In this section, we introduce a method to characterize recurrence probabilities via fixed point theoretic martingales. A run of a system is called recurrent with respect to a region C, if the run visits C infinitely often. Recurrence is also known as liveliness property, which states that no matter what happens, the system will eventually come back to a safe state. In LTL this property can be specified via the formula $\Box\Diamond C$. Recurrence is known to be a combination of a least fixed point ($\Diamond C$) and a greatest fixed point ($\Box(\Diamond C)$) and we show that in the demonic case, these two fixed points can be computed/approximated sequentially via a combination of super- and submartingales. In the angelic case, this sequential computation is not directly possible.

Definition 3.2.11 (Recurrence probability)**.** Let $C \in \mathcal{B}(L \times \mathbb{R}^V)$ be a measurable target region. A run $c_0 c_1 \dots$ is said to be C-recurrent, if for all $n \in N$, there is a $i \geq n$ such that $c_i \in C$ holds. We denote the set of all C-recurrent runs as Π_C^{rec}. For any scheduler σ and starting configuration c, we define $\mathbb{P}_C^{\text{rec}}(c, \sigma) := \Pr_{c,\sigma}(\Pi_C^{\text{rec}})$ to be the probability that a run starting from c under scheduler σ is C-recurrent. Naturally, we define the upper and lower versions by supremizing and infimizing over all schedulers respectively.

Note that the above definition is well-defined only if the set Π_C^{rec} is measurable. We now give a characterization of this set as a limit of reachability sets, which also shows its Borel measurability. We also remind the reader that $\Pr_{c,\sigma}$ is nothing but the path measure $\nu^{c,\sigma}$ from Lemma 3.1.1.

Lemma 3.2.36. *For any $C \in \mathcal{B}(L \times \mathbb{R}^V)$, the set Π_C^{rec} is equivalent to the set $\bigcap_n \Pi_C^{\text{reach}\geq n}$, where $\Pi_C^{\text{reach}\geq n} := \{c_0 c_1 \cdots \in \Pi \mid \exists i \geq n : c_i \in C\}$ are delayed reachability sets. Each of these sets is Borel measurable.*

Proof. The following relation is apparent: $\Pi_C^{\text{rec}} \subseteq \Pi_C^{\text{reach}\geq n+1} \subseteq \Pi_C^{\text{reach}\geq n}$. Every run that enters C after at least $n+1$ steps, also enters it after at least n steps. A C-recurrent run visits C infinitely often and hence by definition is contained in any of the delayed reachability sets. We have $\Pi_C^{\text{rec}} \subseteq \bigcap_n \Pi_C^{\text{reach}\geq n}$. For the other direction, consider a run in $c_0 c_1 \cdots \in \bigcap_n \Pi_C^{\text{reach}\geq n}$. For all $n \in \mathbb{N}$, the run is contained in $\Pi_C^{\text{reach}\geq n}$, so there exists an $i \geq n$ such that $c_i \in C$ holds. But this is exactly a C-recurrent run.

To see that these sets are measurable, we first denote by $\text{Rect}_n(C)$ the measurable rectangle $(L \times \mathbb{R}^V)^n \times C \times (L \times \mathbb{R}^V)^\omega$. Then we have $\Pi_C^{\text{reach}\geq n} = \bigcup_{i \geq n} \text{Rect}_i(C)$ is the countable union of measurable sets and thus measurable itself. The set Π_C^{rec} is then also measurable as the countable intersection of measurable sets. □

If we denote the probability $\Pr_{c,\sigma}(\Pi_C^{\text{reach}\geq n})$ by $\mathbb{P}_C^{\text{reach}\geq n}(c, \sigma)$, then we want to show that $\mathbb{P}_C^{\text{reach}\geq n} = \mathbb{X}^n \mathbb{P}_C^{\text{reach}\geq 0} = \mathbb{X}^n \mathbb{P}_C^{\text{reach}}$ converges to $\mathbb{P}_C^{\text{rec}}$ as n goes to infinity. This leads to a sequential fixed point characterization by first computing $\mathbb{P}_C^{\text{reach}}$ as a least fixed point (see NNRepSupM) and then seeing $\mathbb{P}_C^{\text{rec}}$ as the greatest fixed point given by the decreasing Cousot-Cousot sequence $(\mathbb{X}^n \mathbb{P}_C^{\text{reach}})_{n \in \mathbb{N}}$.

Lemma 3.2.37. *For all $n \in \mathbb{N}$ we have $\mathbb{P}_C^{\text{reach}\geq n+1} = \mathbb{X}\mathbb{P}_C^{\text{reach}\geq n}$.*

Proof. A path $c_0c_1c_2\dots$ is contained in $\Pi_C^{\text{reach}\geq n+1}$ if and only if the suffix $c_1c_2\dots$ is contained in $\Pi_C^{\text{reach}\geq n}$. From this we get

$$\mathbb{P}_C^{\text{reach}\geq n+1}(c,\sigma) = \int \mathbb{P}_C^{\text{reach}\geq n}(c',\sigma_c)\mu_c^\sigma(\mathrm{d}c') = \mathbb{X}\mathbb{P}_C^{\text{reach}\geq n}(c,\sigma)$$

□

In terms of LTL, the nexttime operator $\mathbb{X}$ can be identified with the LTL next operator $\bigcirc$. We generally have $\mathbb{X}\Pr_{c,\sigma}(\pi \models \varphi) = \Pr_{c,\sigma}(\pi \models \bigcirc\varphi)$ for all LTL formulas φ. The above lemma corresponds to the equivalence of the LTL formulas $\Diamond^{\geq n+1}$ and $\bigcirc\Diamond^{\geq n}$. Indeed, in terms of LTL, the operator $\Diamond^{\geq n}$ is shorthand for $\bigcirc^n\Diamond$ which makes the previous claim trivial. Now we need show that this sequence $(\mathbb{X}^n\mathbb{P}^{\text{reach}})_{n\in\mathbb{N}}$ indeed converges to $\mathbb{P}_C^{\text{rec}}$, but this is just an application of Lemmas 3.2.36 and 3.2.37.

Lemma 3.2.38. *The sequence $(\mathbb{X}^n\mathbb{P}_C^{\text{reach}})_{n\in\mathbb{N}}$ decreases monotonically and converges to $\mathbb{P}_C^{\text{rec}}$.*

Proof. First of all note that $\Pr_{c,\sigma}$ is a probability measure for all σ and c and therefore monotone. The sequence $\Pi_C^{\text{reach}\geq n}$ is monotonically decreasing and hence the associated sequence of probabilities $\Pr_{c,\sigma}(\Pi_C^{\text{reach}\geq n}) = \mathbb{P}_C^{\text{reach}\geq n}(c,\sigma) = \mathbb{X}^n\mathbb{P}_C^{\text{reach}}(c,\sigma)$ is monotonically decreasing. For the convergence we use Lemma 3.2.36 and get

$$\begin{aligned}
\mathbb{P}_C^{\text{rec}}(c,\sigma) &= \Pr_{c,\sigma}(\bigcap_n \Pi_C^{\text{reach}\geq n}) \\
&= \inf_n \Pr_{c,\sigma}(\Pi_C^{\text{reach}\geq n}) \\
&= \lim_{n\to\infty} \mathbb{P}_C^{\text{reach}\geq n}(c,\sigma) \\
&= \lim_{n\to\infty} \mathbb{X}^n\mathbb{P}_C^{\text{reach}}(c,\sigma) = \mathbb{X}^\omega\mathbb{P}_C^{\text{reach}}(c,\sigma).
\end{aligned}$$

We made use of the fact that the sequence is decreasing and hence infimum and limit coincide. □

Now we are ready to state the main theorem of this section.

Theorem 3.2.39. *The tuple $(\eta_1,\eta_2) = (\mathbb{P}_C^{\text{reach}},\mathbb{P}_C^{\text{rec}})$ is the solution to the fixed point equation system given by*

$$\begin{aligned}
\eta_1 &=_\mu \mathbb{1}_C + \mathbb{1}_{\overline{C}}\mathbb{X}\eta_1 \\
\eta_2 &=_\nu \min\{\eta_1, \mathbb{X}\eta_2\}
\end{aligned}$$

Proof. The solution to the first fixed point equation is simply the least NNRepSupM as shown in Theorem 3.2.21. This is exactly $\mathbb{P}_C^{\text{reach}}$. The solution to the second equation is the greatest fixed point of $\eta_2 \mapsto \mathbb{X}\eta_2$ that is less than $\eta_1 = \mathbb{P}_C^{\text{reach}}$. We can consider η_1 as the top element $\top$ and compute the greatest fixed point under η_1 by considering the Cousot-Cousot sequence $\mathbb{X}^n(\top)$. For $\top = \eta_1 = \mathbb{P}_C^{\text{reach}}$ this converges to $\mathbb{P}_C^{\text{rec}}$ by Lemma 3.2.38. □

In order to find sound approximations for $\mathbb{P}_C^{\text{rec}}$, one has to either underapproximate or overapproximate both fixed points in Theorem 3.2.39. Knaster-Tarski is insufficient for this task since it gives approximations in different directions for both types of fixed points. However we can use γ-scaled submartingales to get an underapproximation $\eta_1 \leq \mathbb{P}_C^{\text{reach}}$ and then use Knaster-Tarski to get an underapproximation of the second fixed point. This method is generally not complete because γ-scaled submartingales are not complete.

Corollary 3.2.39.1. *Let $C \in \mathcal{B}(L \times \mathbb{R}^V)$ be a target region, $\gamma \in (0,1)$ a scaling factor and let (η_1, η_2) be functions such that*

1. $\eta_1 \leq \mathbb{1}_C + \gamma \mathbb{1}_{\overline{C}} \mathbb{X} \eta_1$,
2. $\eta_2 \leq \mathbb{X}\eta_2$ *and*
3. $\eta_2 \leq \eta_1$,

then $(\eta_1, \eta_2) \leq (\mathbb{P}_C^{\text{reach}}, \mathbb{P}_C^{\text{rec}})$ holds.

Proof. η_1 is just a γ-scaled submartingale and by Theorem 3.2.32 underapproximates $\mathbb{P}_C^{\text{reach}}$. η_2 is a post-fixpoint of the mapping $\eta \mapsto \min\{\eta_1, \mathbb{X}\eta\}$ and by monotonicity of the minimum function, it is also a post-fixpoint of $\eta \mapsto \min\{\mathbb{P}_C^{\text{reach}}, \mathbb{X}\eta\}$. Hence by Knaster-Tarski, it underapproximates the greatest fixed point of the later, which is exactly $\mathbb{P}_C^{\text{rec}}$. □

Recurrence can be thought of as a two objective function. The first objective is to reach the target region C, the second one is to stay inside a region that is capable of revisiting C. Because of this, replacing the nexttime operator $\mathbb{X}$ in the fixed point characterization of $\mathbb{P}_C^{\text{rec}}$ by $\overline{\mathbb{X}}$ yields a system whose solution is not $\overline{\mathbb{P}}_C^{\text{rec}}$, but rather an overapproximation of it. The key observation is that $\overline{\mathbb{P}}_C^{\text{rec}}$ is determined by schedulers which simultaneously maximize both objectives, while a modified fixed point characterization allows the usage of different schedulers for each objective. Consider the following trivial example.

Listing 3.4: Example: Upper recurrence incompleteness

```
1. x := 0;
2. while x = 0 do
3.         if * then
4.                 x := 1
           fi
   od
5. assert true;
6. skip;
```

The target region is marked by the assert statement. Obviously, this statement can only be reached once so the system is never recurrent under any scheduler. Now consider the upper reachability probability $\overline{\mathbb{P}}_C^{\text{reach}}$, which assigns to every configuration with label $l \in \{1,2,3,4,5\}$ the value 1, and 0 for the last skip statement. We can simply use a scheduler which executes $x := 1$ in the first loop iteration to break out of the loop and reach the assert. If we now compute a few iterations of $\overline{\mathbb{X}}^n \overline{\mathbb{P}}_C^{\text{reach}}$, then it will converge to a function that assigns value 1 to the labels $\{1,2,3\}$ and 0 elsewhere. The reason is that we can take

a scheduler which never leaves the loop and thus stays inside a region R which satisfies $\overline{\mathbb{P}}_C^{\text{reach}}(\mathrm{R}) = 1$. The problem here is that the objective of staying inside R and the goal of reaching C are mutually exclusive and no single scheduler can achieve both. By considering both objectives separately, we can find two schedulers that satisfy one objective each.

Interestingly, lower recurrence $\underline{\mathbb{P}}_C^{\text{rec}}$ does not suffer from this problem. It is indeed possible to take Theorem 3.2.39 and simply replace $\mathbb{X}$ by $\underline{\mathbb{X}}$ for both fixed points to get a correct characterization of $\underline{\mathbb{P}}_C^{\text{rec}}$. Before we proof this statement, we first need to establish the existence of ϵ-optimal schedulers for $\mathbb{P}^{\text{reach}\geq n}$.

Lemma 3.2.40. *For all $n \in \mathbb{N}$ and target regions $C \in \mathcal{B}(L \times \mathbb{R}^V)$, the function $\underline{\mathbb{P}}_C^{\text{reach}\geq n}$ admits lower and upper ϵ-optimal schedulers*

Proof. Unlike before, we cannot directly reference results from Takisaka et al. but we use their proof technique to show that ϵ-optimal schedulers do exist. We use the construction by Takisaka et al.[11, Lemma 3.2] to construct an infinite horizon stochastic optimal control model from a given pCFG for which we know the existence of ϵ-optimal schedulers[18]. In this model we assume that a configuration in C is reached only once (by modifying it to go to a fixed self-looping configuration after entering C). Furthermore, we extend the configuration space to $(L \times \mathbb{R}^V) \times \{0, ..., N\}$ such that every transition moves from (c, k) to some $(c', \min(k+1, n))$. This is done to differentiate the first n steps of the system. Note that since $\{0, ..., n\}$ is finite, we do not introduce any measurability issues here. Now we can define a single-step cost function $g_n : ((L \times \mathbb{R}^V) \times \{0, ..., n\}) \times (L \cup \mathbb{R}) \to [0, 1]$ as follows.

$$g_n((c,k),\xi) = \begin{cases} 1 & k = n \text{ and } c \in C, \\ 0 & \text{else} \end{cases}$$

The parameter $\xi \in (L \cup \mathbb{R})$ is the action that was chosen at state (c, k). It can either be a nondeterministic transition to a different location $l \in L$ or the real value $x \in \mathbb{R}$ of a nondeterministic assignment. In either case, we ignore the action and give a reward of 1 if we are in the target region C and at least n steps have been made so far. Then this cost function induces an infinite horizon expected reward function

$$\mathbf{J}_n(c_0, \sigma) := \int \sum_{i=0}^{\infty} g_n(c_k, \xi_k) \nu^{c,\sigma}(\mathrm{d}\pi)$$

where $\pi = (c_0, \xi_0)(c_1, \xi_1)(c_2, \xi_2) \dots$ is a path augmented with actions such that action ξ_i is chosen according to $\sigma(c_0 c_1 \dots c_i)$. It is easy to see that $\mathbf{J}_n = \mathbb{P}_C^{\text{reach}\geq n}$ holds because exactly the paths in $\Pi_C^{\text{reach}\geq n}$ get a reward of 1, while all others get a reward of 0. We made some adaptions to $\mathbf{J}_n$ to fit it in our setting. Originally it is formulated via policies instead of schedulers, but they are essentially identical and can be translated to each other[11, Lemma 3.2]. The cost function g_n is the indicator function of the Borel set $C \times \{n\}$ and as such it is bounded, nonnegative and lower semianalytic. There exist ϵ-optimal policies and hence schedulers for such cost functions[18, Proposition 9.19] □

Theorem 3.2.41. *The tuple $(\eta_1, \eta_2) = (\underline{\mathbb{P}}_C^{\text{reach}}, \underline{\mathbb{P}}_C^{\text{rec}})$ is the unique solution to the fixed point*

equation system given by

$$\eta_1 =_\mu \mathbb{1}_C + \mathbb{1}_{\overline{C}} \underline{\mathbb{X}} \eta_1$$
$$\eta_2 =_\nu \min\{\eta_1, \underline{\mathbb{X}} \eta_2\}$$

Proof. The solution to the first fixed point equation is simply the least LNNRepSupM as shown in Theorem 3.2.25. This is exactly $\underline{\mathbb{P}}_C^{\text{reach}}$. For the second part, we first note that $\underline{\mathbb{P}}_C^{\text{reach}\geq n+1} = \underline{\mathbb{X}}\, \underline{\mathbb{P}}_C^{\text{reach}\geq n}$ holds by the existence of ϵ-optimal schedulers and Lemma 3.2.2. The sequence defined by $\underline{\mathbb{P}}_C^{\text{reach}\geq n} = \underline{\mathbb{X}}^n \underline{\mathbb{P}}_C^{\text{reach}}$ is monotonically decreasing hence converges. We show that it converges to $\underline{\mathbb{P}}_C^{\text{rec}}$.

$$\begin{aligned}
\underline{\mathbb{P}}_C^{\text{rec}}(c) &= \inf_\sigma \mathbb{P}_C^{\text{rec}}(c, \sigma) \\
&= \inf_\sigma \inf_n \mathbb{P}_C^{\text{reach}\geq n}(c, \sigma) \\
&= \inf_n \inf_\sigma \mathbb{P}_C^{\text{reach}\geq n}(c, \sigma) \\
&= \inf_n \underline{\mathbb{P}}_C^{\text{reach}\geq n}(c) \\
&= \inf_n \underline{\mathbb{X}}^n \underline{\mathbb{P}}_C^{\text{reach}}(c) = \underline{\mathbb{X}}^\omega \underline{\mathbb{P}}_C^{\text{reach}}(c)
\end{aligned}$$

We made use of the fact that $\mathbb{P}_C^{\text{rec}} = \lim_{n\to\infty} \mathbb{P}_C^{\text{reach}\geq n} = \inf_n \mathbb{P}_C^{\text{reach}\geq n}$ holds from Lemma 3.2.38 and that the order of taking infima can be exchanged. □

This proof is not adaptable to the case of upper recurrence because the exchange of order of sup and inf is generally not possible. In fact, by exchanging sup inf to inf sup, the result gets larger, which explains why a similar fixed point characterization via $\overline{\mathbb{X}}$ fails and leads to a fixed point that is larger than $\overline{\mathbb{P}}_C^{\text{rec}}$ as demonstrated in Example 3.4.

Now we take a look at the problem of deciding almost sure recurrence, i.e. deciding whether $\mathbb{P}_C^{\text{rec}}(c) = 1$ holds. The problem of qualitatively deciding almost sure recurrence in a probabilistic system is surprisingly similar to the problem of deciding positive almost sure reachability. Indeed, ranking functions are capable of this. If we can find a ranking function for a target region C that is finite on an invariant I containing C then the system must be recurrent. Simply put, finiteness of an ARnkSupM on I implies a finite expected reaching time $\mathbb{E}_C^{\text{steps}}(c)$ for every $c \in I$ by Theorem 3.2.5, so every configuration is eventually driven towards C. We then have for all $c \in I$:

1. $\Pr(c \models \Box I) = 1$ and
2. $\Pr(c \models \Diamond C) = 1$.

These conditions imply that the systems always stays in I (since it is an invariant) and whenever it leaves the target region C, it will eventually reenter it. This technique is reminiscent of Foster's Theorem for recurrence in countable state Markov chains. Indeed such reasoning is found in [27] and a generalization to uncountable Markov chains is given by Meyn et al. in [28]. We formalize the above mentioned idea know.

Definition 3.2.12 (Recurrence set $\mathrm{Rec}_{C,\sigma}$)**.** For a target region $C \in \mathcal{B}(L \times \mathbb{R}^V)$ and a *history-independent* scheduler σ, we define the recurrence set $\mathrm{Rec}_{C,\sigma} = \left\{c \in L \times \mathbb{R}^V \mid \mathbb{P}_C^{\mathrm{rec}}(c,\sigma) = 1\right\}$.

We use history-independent schedulers for this definition to make the following characterization more concise. There is no gain in considering history-dependent schedulers.

Theorem 3.2.42. $\mathrm{R} = \mathrm{Rec}_{C,\sigma}$ *satisfies the following two properties.*

1. *for all* $c \in \mathrm{R}$: $\Pr_\sigma(c \models \Box \mathrm{R}) = 1$ *and*
2. *for all* $c \in \mathrm{R}$: $\Pr_\sigma(c \models \Diamond C) = \mathbb{P}_C^{\mathrm{reach}}(c,\sigma) = 1$.

Furthermore, any set satisfying the above properties is a subset of $\mathrm{Rec}_{C,\sigma}$*, making it the largest such set.*

Proof. First we show that $\mathrm{R} = \mathrm{Rec}_{C,\sigma}$ indeed satisfies both properties. Property 2 holds directly because $1 = \mathbb{P}_C^{\mathrm{rec}}(c,\sigma) \leq \mathbb{P}_C^{\mathrm{reach}}(c,\sigma)$ holds. For property 1 assume that there is a configuration $c_0 \in \mathrm{R}$ such that $L \times \mathbb{R}^V \setminus \mathrm{R}$ has positive measure (i.e. $\mu_{c_0}^\sigma(L \times \mathbb{R}^V \setminus \mathrm{R}) > 0$), then we get the following contradiction:

$$\begin{aligned} 1 = \mathbb{P}_C^{\mathrm{rec}}(c_0,\sigma) = \mathbb{X}\mathbb{P}_C^{\mathrm{rec}}(c_0,\sigma) &= \int_{L\times\mathbb{R}^V} \mathbb{P}_C^{\mathrm{rec}}(c_1,\sigma)\mu_{c_0}^\sigma(\mathrm{d}c_1) \\ &= \int_{\mathrm{R}} \mathbf{1}\mu_{c_0}^\sigma(\mathrm{d}c_1) + \int_{L\times\mathbb{R}^V\setminus\mathrm{R}} \mathbb{P}_C^{\mathrm{rec}}(c_1,\sigma)\mu_{c_0}^\sigma(\mathrm{d}c_1) \\ &< \mu_{c_0}^\sigma(\mathrm{R}) + \mu_{c_0}^\sigma(L \times \mathbb{R}^V \setminus \mathrm{R}) = 1. \end{aligned}$$

The set $L \times \mathbb{R}^V \setminus \mathrm{R}$ must have measure zero starting from any configuration c_0 in R. The last inequality is obtained because $\mathbb{P}_C^{\mathrm{rec}}(c_1,\sigma) < 1$ for all $c_1 \in L \times \mathbb{R}^V \setminus \mathrm{R}$. The system stays almost surely in R.

We proceed to show that any set R satisfying 1 and 2 is a subset of $\mathrm{Rec}_{C,\sigma}$ making it the largest such set. Suppose both properties hold and let $c \in \mathrm{R}$ be a configuration. By property 1 we can conclude $\Pr_\sigma(c \models \bigcirc^n \mathrm{R}) = \mathbb{X}^n(\mathbb{1}_\mathrm{R})(c,\sigma) = 1$ for all $n \in \mathbb{N}$, where $\mathbb{1}_\mathrm{R}$ maps input (c,σ) to 1 if $c \in \mathrm{R}$ and else to 0. Note that we need to define $\mathbf{1}_\mathrm{R}$ scheduler-dependent to be compatible with the definition of the $\mathbb{X}$ operator. Property 2 implies $\mathbb{P}_C^{\mathrm{reach}} \cdot \mathbb{1}_\mathrm{R} = \mathbb{1}_\mathrm{R}$. By 3.2.38 we get the following equality

$$\begin{aligned} \mathbb{P}_C^{\mathrm{rec}}(c,\sigma) &= \inf_n \mathbb{P}_C^{\mathrm{reach}\leq n}(c,\sigma) = \inf_n \Pr_\sigma(c \models \Diamond^{\geq n} C) \\ &= \inf_n \Pr_\sigma(c \models \bigcirc^n \Diamond C) = \inf_n \Pr_\sigma(c \models \bigcirc^n(\Diamond C \wedge \mathrm{R})) \\ &= \inf_n \mathbb{X}^n(\mathbb{P}_C^{\mathrm{reach}} \cdot \mathbb{1}_\mathrm{R})(c,\sigma) = \inf_n \mathbb{X}^n(\mathbb{1}_\mathrm{R})(c,\sigma) = 1. \end{aligned}$$

In the fourth inequality we use the fact that $c \models \bigcirc^n R$ is an almost sure event, in the second to last equality we use property 2 and in the last one use property 1. □

Recurrence sets in probabilistic programs have strong similarities to bottom strongly connected components (BSCCs) in finite Markov chains. A BSCC is a strongly connected

component that has no outgoing transitions. In finite Markov chains recurrence can be decided by finding all BSCCs that intersect C and find whether one of them is eventually reached[29]. The bottomness of BSCCs corresponds to property 1 in that both the BSCC and the recurrence set are not left once entered (they both are kinds of invariants). If a BSCC intersects with the target region C then finiteness as well as strongly connectedness guarantee that the region C is reached. This corresponds to property 2 of the recurrence sets. In fact, the union of all BSCCs of a finite Markov chain which intersect some region C form a recurrence set for that region (since a Markov Chain is scheduler-independent, recurrence with respect to a scheduler σ and regular recurrence coincide). As previously mentioned, these properties can be decided by finding an invariant I and a ranking supermartingale on that invariant.

4 Template-based Synthesis

This section explains methods to automatically synthesize martingales which are introduced in Section 3.2. The method of choice is so called template-based synthesis, where a template function is fixed and then its variable parameters are computed such that it is indeed a martingale for the program in question. We consider two types of templates; linear templates which can be solved using linear programming (LP), and polynomial templates which are solved by semi-definite programming (SDP). In both cases, the problem becomes an optimization problem which allows us to compute a solution that not only satisfies the martingale constraints but also minimizes or maximizes the probability bounds for the initial configuration. Template-based synthesis is the state of the art method to construct martingale-like expressions for verification. Linear as well as polynomial templates are widely used[12, 13, 9, 11, 14, 27, 8]. We apply synthesis of linear martingales to short programs and show its results in Section 4.3.

4.1 Linear Templates

In the following we assume a pCFG Γ over a finite variable set $V = \{x_1, x_2, \ldots, x_n\}$ with a finite location set L. We also assume that an invariant I and a target region C are given.

Definition 4.1.1 (Linear expressions)**.** A linear expression is given by the grammar

$$\langle a \rangle ::= r \mid x \mid r \cdot \langle a \rangle \mid \langle a \rangle + \langle a \rangle$$

where $r \in R$ is a real value and $c \in V$ a variable. Its semantics $[\![a]\!] : \mathbb{R}^V \to \mathbb{R}$ are inductively defined by $[\![r]\!](\mathbf{x}) = r$, $[\![x]\!](\mathbf{x}) = \mathbf{x}(x)$, $[\![r \cdot a]\!](\mathbf{x}) = r \cdot [\![a]\!](\mathbf{x})$ and $[\![a + b]\!](\mathbf{x}) = [\![a]\!](\mathbf{x}) + [\![b]\!](\mathbf{x})$. A linear expression map is a mapping η that assigns to each location $l \in L$ a linear expression over V. Its semantics are defined by $[\![\eta]\!](l, \mathbf{x}) = [\![\eta(l)]\!](\mathbf{x})$.

Definition 4.1.2 (Linear constraints, conjunctions and predicates)**.** A linear constraint is an expression of the form $a > 0$ or $a \geq 0$ where a is a linear expression. We define $[\![a \triangleright 0]\!] = \{\mathbf{x} \in \mathbb{R}^V \mid [\![a]\!](\mathbf{x}) \triangleright 0\} \subseteq \mathbb{R}^V$ where $\triangleright \in \{>, \geq\}$. A linear conjunction is an expression of the form $r_1 \wedge r_2 \wedge \cdots \wedge r_m$ where each r_i is a linear constraint. A linear predicate $p_1 \vee p_2 \vee \cdots \vee p_j$ is a disjunction of linear conjunctions. Their semantics are given by $[\![r_1 \wedge \cdots \wedge r_m]\!] = \bigcap_{i=1}^{m} [\![r_i]\!]$ and $[\![p_1 \vee \cdots \vee p_j]\!] = \bigcup_{i=1}^{j} [\![p_i]\!]$. A linear predicate map P assigns to each label $l \in L$ a linear predicate $P(l)$. We define $[\![P]\!](l) = [\![P(l)]\!]$.

Our goal is to automatically synthesize a linear expression map η that satisfies the martingale conditions and is as close as possible to the bounds (over or underapproximation depends on the concrete martingale we want to synthesize). For this, we need some restrictions on the pCFG Γ that ensure *linearity*. We have the following conditions and assumptions

1. For assignments $l \in L_A$ and update function $\mathrm{Up}(l) = (x_j, u)$ we have

a) if $u = f \in \mathcal{B}(\mathbb{R}^V, \mathbb{R})$, then $f = [\![\mathfrak{f}]\!]$ for a linear expression $\mathfrak{f}$,

b) if $u = d \in \mathcal{D}(\mathbb{R})$, then the expectation $\mathbb{E}(d)$ is known,

c) if $u = A \in \mathcal{B}(\mathbb{R})$, then $A = [\![\mathfrak{A}]\!]$ for some linear predicate $\mathfrak{A}$ over $\{x_j\}$,

2. if $l \in L_D$, then each guard $G(l, l') = [\![\mathfrak{G}]\!]$ is given by a linear predicate $\mathfrak{G}$,

3. the invariant $I = [\![\mathfrak{I}]\!]$ is given by a linear predicate map $\mathfrak{I}$,

4. the target region $C = [\![\mathfrak{C}]\!]$ is given by a linear predicate map $\mathfrak{C}$,

5. the pCFG is finitely branching. Every location has only finitely many successors.

For pCFGs that are induced by a affine probabilistic program (APP), most of these conditions are automatically satisfied. We will use a slightly adapted syntax for APP to annotate labels with predicates that specify the invariant and the target region. This syntactical extension is described in the Experiments section.

Now we fix a linear expression map template for the martingale η of the form $\eta_l = c_1^l x_1 + c_2^l x_2 + \cdots + c_n^l x_n + d^l$ for each $l \in L$, where the coefficients c_i^l and d^l are unknown parameters to be optimized for. We denote this set of coefficients by $\mathcal{A}$. We introduce a new variable x_{ndet} that is used to represent nondeterministic choice and extend the set of variables to $V' = V \cup \{x_{\text{ndet}}\}$. We construct 3 sets of formulas for each label l with the following intentions:

- The first set of formulas $\mathfrak{F}_1^l$ describes the invariant conditions (e.g. positivity inside the invariant).
- The second set $\mathfrak{F}_2^l$ describes the conditions inside the target region (such as taking value 1 for NNRepSupM).
- The third set $\mathfrak{F}_3^l$ describes the martingale conditions outside of C (increase or decrease in expectation).

Depending on the martingale under consideration, the formulas have different concrete structure but their general structure is always the same. All the formulas we use have the form $\phi \implies \psi$, where ϕ is a linear conjunction over V' with inequalities $\geq$, and ψ is a formula of the form $c_1 x_1 + c_2 x_2 + \cdots + c_n x_n + c_{\text{ndet}} x_{\text{ndet}} + d \geq 0$ where each c_i and d is a linear expression over the coefficient set $\mathcal{A}$.

Since all expressions are linear, we can equivalently describe each such formula via $A\mathbf{x} + b \geq 0 \implies c^T\mathbf{x} + d \geq 0$. Note that the coefficient vector c contains linear expressions over $\mathcal{A}$. We make use of Farka's Lemma to transcribe the implications into a linear program to optimize. For this, we relax all strict inequalities $>$ to $\geq$ to allow the usage of Farkas' Lemma. This affects completeness of the method but soundness is preserved. LP solvers can not handle strict inequalities so this kind of relaxation is necessary either way.

Proposition 4.1.1 (Farkas' Lemma (affine version))**.** *Let $A\mathbf{x} \leq b$ have at least one solution, then $A\mathbf{x} \leq b \implies c^T\mathbf{x} \leq d$ holds if and only if there exists $y \geq 0$ such that $y^T A = c^T$ and $y^T b \leq d$.*

For a proof we refer to [30, Corollary 7.1h]. The idea behind this lemma is as follows. If $A\mathbf{x} \leq b$ is feasible, then the inequalities of this system can be conically combined to derive new inequalities $y^T A\mathbf{x} \leq y^T b$ (with $y \geq 0$). Then $c^T\mathbf{x} \leq d$ holds if and only if it is possible to derive a (potentially) stronger inequality of the form $c^T\mathbf{x} \leq d'$ where $d' \leq d$, since $c^T\mathbf{x} \leq d'$ implies $c^T\mathbf{x} \leq d$.

By bringing the inequalities $A\mathbf{x} + b \geq 0$ and $c^T\mathbf{x} + d \geq 0$ of our original system into the form $(-A)\mathbf{x} \leq b$ and $(-c^T)x \leq d$ we can apply Farkas' Lemma and obtain: There exists y such that $y^T(-A) = -c$ and $y^T b \leq d$, or equivalently, $y^T A = c$ and $y^T b \leq d$. This application is valid if $A\mathbf{x} \leq b$ has a solution. If it has not, then the implications $A\mathbf{x} \leq b \implies c^T x \leq d$ are trivially satisfied anyway.

Corollary 4.1.1.1. *An implication $A\mathbf{x} + b \geq 0 \implies c^T\mathbf{x} + d \geq 0$ is true if either $A\mathbf{x} + b \geq 0$ has no solutions or there exists a nonnegative vector $y \geq 0$ such that $y^T A = c$ and $y^T b \leq d$ holds.*

This reduces satisfiability of linear implications into a feasibility problem over the variable set $\mathcal{A} \cup \{y_1, \ldots, y_m\}$ where m is the number of rows in A and can be solved using standard linear programming algorithms. Now it remains to concretely construct such formulas for different martingales. We show the constructions for UNNRepSupM and L-γ-SclSubM. UARnkSupM and higher order UARnkSupM can be handled similarly (See [9] and [14], respectively). From now on we assume the following form of the predicates describing the invariant, target region and the guards:

- the invariant is given by $\mathfrak{I}(l) = \bigvee_{i=1}^{N_l^{\mathfrak{I}}} \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{I}}} (a_{i,j}^l \mathbf{x} + b_{i,j}^l \geq 0) = \bigvee_{i=1}^{N_l^{\mathfrak{I}}} \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{I}}} (\alpha_{i,j}^l \geq 0)$, where $a_{i,j}^l$ is a row vector of coefficients, and $\alpha_{i,j}^l$ is just a shorthand for the linear expression $a_{i,j}^l \mathbf{x} + b_{i,j}^l \geq 0$,
- the target region is given by $\mathfrak{C}(l) = \bigvee_{i=1}^{N_l^{\mathfrak{C}}} \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{C}}} (a_{i,j}^l \mathbf{x} + b_{i,j}^l \geq 0) = \bigvee_{i=1}^{N_l^{\mathfrak{C}}} \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{C}}} (\beta_{i,j}^l \geq 0)$,
- the guards are given by $\mathfrak{G}(l, l') = \bigvee_{i=1}^{N_{l,l'}^{\mathfrak{G}}} \bigwedge_{j=1}^{N_{l,l',i}^{\mathfrak{G}}} (a_{i,j}^{l,l'} \mathbf{x} + b_{i,j}^{l,l'} \geq 0) = \bigvee_{i=1}^{N_{l,l'}^{\mathfrak{G}}} \bigwedge_{j=1}^{N_{l,l',i}^{\mathfrak{G}}} (g_{i,j}^{l,l'} \geq 0)$.

We consider the special case of UNNRepSupM and assume its form to be $\eta_l = c_1^l x_1 + c_2^l x_2 + \cdots + c_n^l x_n + d^l$ for each location $l \in L$. The first axiom we need to encode is non-negativity inside the invariant ($\eta \geq 0$ inside I):

$$\mathfrak{F}_1^l = \{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{I}}} (\alpha_{i,j}^l \geq 0) \implies (c_1^l x_1 + c_2^l x_2 + \cdots + c_n^l x_n + d^l \geq 0) \mid 1 \leq i \leq N_l^{\mathfrak{I}} \}$$

If any of the invariant disjunctions are satisfied, then the supermartingale has to be positive.

The second axiom we encode is the property that a UNNRepSupM takes value 1 inside the target region ($\eta(c) = 1, c \in C$). We relax this axiom and allow it to take values greater than 1 without affecting soundness nor completeness. Furthermore, this condition is only

enforced inside the invariant, that is inside the intersection of C and I.

$$\mathfrak{F}_2^l = \{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{j=1}^{N_{l,m}^{\mathfrak{C}}} (\beta_{m,j}^l \geq 0) \implies$$
$$(c_1^l x_1 + c_2^l x_2 + \cdots + c_n^l x_n + (d^l - 1) \geq 0) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq m \leq N_l^{\mathfrak{C}} \}$$

If we are inside the target region and invariant, then the supermartingale has to take value of at least 1.

Lastly, we need to encode the supermartingale axiom itself, i.e. the decrease in expectation per transition ($\eta \geq \overline{\mathbb{X}}\eta$ outside of C). This condition only needs to be satisfied inside the invariant but outside the target region, that is, inside $I \cap \overline{C}$. Note that $(l, \mathbf{x})$ is outside of C if and only if for each $1 \leq m \leq N_l^{\mathfrak{C}}$ there is a $1 \leq j_m \leq N_{l,m}^{\mathfrak{C}}$ such that $\beta_{m,j_m}^l < 0$, or equivalently, $-\beta_{m,j_m}^l > 0$. We relax this inequality to $-\beta_{m,j_m}^l \geq 0$ which affects the completeness but not the soundness of this method. We do a case distinction based on the location type l:

- If $l \in L_N$, then

$$\mathfrak{F}_3^l = \Big\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \implies \big((c_1^l - c_1^{l'}) x_1 + (c_2^l - c_2^{l'}) x_2 + \cdots +$$
$$(c_n^l - c_n^{l'}) x_n + (d^l - d^{l'}) \geq 0 \big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}}, l' \in \mathrm{succ}(l) \Big\}$$

  Note that we iterate over all successors l' of l and over all tuples $\vec{j} = (j_1, \ldots, j_{N_l^{\mathfrak{C}}})$ with $1 \leq j_m \leq N_{l,m}^{\mathfrak{C}}$.

- If $l \in L_P$, then

$$\mathfrak{F}_3^l = \Big\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \implies \big((c_1^l - \sum_{l \mapsto l'} \Pr(l,l') c_1^{l'}) x_1 + \cdots +$$
$$(c_n^l - \sum_{l \mapsto l'} \Pr(l,l') c_n^{l'}) x_n + (d^l - \sum_{l \mapsto l'} \Pr(l,l') d^{l'}) \geq 0 \big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}} \Big\}$$

- If $l \in L_D$, then for each successor $l' \in \mathrm{succ}(l)$ there is a unique guard $G(l,l')$ and

$$\mathfrak{F}_3^l = \Big\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \wedge \bigwedge_{j=1}^{N_{l,l',h}^{\mathfrak{G}}} (g_{h,j}^{l,l'} \geq 0) \implies \big((c_1^l - c_1^{l'}) x_1 + \cdots +$$
$$(c_n^l - c_n^{l'}) x_n + (d^l - d^{l'}) \geq 0 \big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}}, 1 \leq h \leq N_{l,l'}^{\mathfrak{G}}, l' \in \mathrm{succ}(l) \Big\}$$

- If $l \in L_A$, $\mathrm{Up}(l) = (x_v, u)$ and l' is the unique successor of l then

- if $u = f \in \mathcal{B}(\mathbb{R}^V, \mathbb{R})$ and $f = [\![\mathfrak{f}]\!] = [\![r_1 x_1 + r_2 x_2 + \cdots + r_n x_n + s]\!]$ is a linear expression, then

$$\mathfrak{F}_3^l = \Bigg\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \implies$$
$$\big((c_1^l - c_1^{l'} - r_1 c_v^{l'}) x_1 + \cdots + (c_v^l - r_v c_v^{l'}) x_v + \ldots$$
$$+ (c_n^l - c_n^{l'} - r_n c_v^{l'}) x_n + (d^l - d^{l'} - s c_v^{l'}) \geq 0\big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}} \Bigg\}$$

- if $u = d \in \mathcal{D}(\mathbb{R})$ and $\mathbb{E}(d) = s$, then

$$\mathfrak{F}_3^l = \Bigg\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \implies \big((c_1^l - c_1^{l'}) x_1 + \cdots + (c_v^l - s) x_v$$
$$+ \cdots + (c_n^l - c_n^{l'}) x_n + (d^l - d^{l'}) \geq 0\big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}} \Bigg\}$$

- if $u = [\![\mathfrak{A}]\!] = [\![\bigvee_{k=1}^{N_l^{\mathfrak{A}}} \bigwedge_{j=1}^{N_{l,k}^{\mathfrak{A}}} (p_{k,j}^l x_v + q_{k,j}^l \geq 0)]\!] \in \mathcal{B}(\mathbb{R})$ is given by a linear predicate over $\{x_v\}$, then

$$\mathfrak{F}_3^l = \Bigg\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \wedge \bigwedge_{j=1}^{N_{l,k}^{\mathfrak{A}}} (p_{k,j}^l x_{\text{ndet}} + q_{k,j}^l \geq 0) \implies$$
$$\big((c_1^l - c_1^{l'}) x_1 + \cdots + c_v^l x_v + \cdots + (c_n^l - c_n^{l'}) x_n + (-c_v^{l'}) x_{\text{ndet}} + (d^l - d^{l'}) \geq 0\big)$$
$$\mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}}, 1 \leq k \leq N_l^{\mathfrak{A}} \Bigg\}.$$

  For every choice of x_{ndet}, i.e. every possible way to resolve the nondeterminism, the expectation has to decrease.

It is worth noting that for the nondeterministic locations this type of construction only works because we are maximizing the next step expectation, i.e. no matter how the nondeterminism is resolved, the expectation has to decrease. This is expressible by a conjunction of formulas for each possible way to resolve nondeterminism. For LNNRepSupM, the nondeterminism results in disjunctive formulas because there only has to exist a single way to resolve it such that the supermartingale's expectation decreases. This is not solvable by linear programming, but satisfiability of such formulas is still decidable because of the decidability of first-order theories of the real numbers[31].

As previously described, each of these formulas can be expressed as $A\mathbf{x} + b \geq 0 \implies c^T \mathbf{x} + d \geq 0$. Farkas' Lemma then reduces this satisfiability problem to an optimization problem $y^T A = c$ and $y^T b \leq d$ subject to $y \geq 0$ with parameters in $\mathcal{A} \cup \{y_1, \ldots, y_m\}$. Any solution to this problem provides a certificate for the satisfiability of the above formulas in the form of a supermartingale η. Concretely, we have the following theorem:

Theorem 4.1.2 (Synthesis Theorem of UNNRepSupM). *Let* $\mathrm{sol} : \mathcal{A} \cup \{y_1, \ldots, y_m\} \to \mathbb{R}$ *be a solution to the optimization problem induced by the family of formulas* $\mathfrak{F}_i^l$ *where* $i \in \{1,2,3\}$ *and* $l \in L$*, then the function* $\eta(l, \mathbf{x}) = \mathrm{sol}(c_1^l)x_1 + \mathrm{sol}(c_2^l)x_2 + \cdots + \mathrm{sol}(c_n^l)x_n + \mathrm{sol}(d^l)$ *is a UNNRepSupM for* C *supported by the invariant* I*.*

We have $\overline{\mathbb{P}}_C^{\mathrm{reach}}(l, \mathbf{x}) \leq \eta(l, \mathbf{x})$ and in particular for a starting configuration $c_0 = (l_{\mathrm{init}}, \mathbf{x}_{\mathrm{init}})$, we have $\overline{\mathbb{P}}_C^{\mathrm{reach}}(l_{\mathrm{init}}, \mathbf{x}_{\mathrm{init}}) \leq \eta(l_{\mathrm{init}}, \mathbf{x}_{\mathrm{init}})$. We can minimize the linear objective function $[\![\eta_{l_{\mathrm{init}}}]\!](\mathbf{x}_{\mathrm{init}})$ obtained from the template $\eta_l = c_1^l x_1 + c_2^l x_2 + \cdots + c_n^l x_n + d^l$ to find a best UNNRepSupM at c_0.

With a slight adaption to the formula set $\mathfrak{F}_3$ we can also synthesize UNNRep-δ-SupM easily. We just encode the axiom $\eta - \overline{\mathbb{X}}\eta + \delta \geq 0$ instead of $\eta - \overline{\mathbb{X}}\eta$ in a straightforward manner. Because this δ appears linearly in the conditions, we can make it part of the search space and find a δ that gives the best bounds on bounded reachability probability. In order to do so, we fix a time bound $t \in \mathbb{N}$ and minimize the function $[\![\eta_{l_{\mathrm{init}}}]\!](\mathbf{x}_{\mathrm{init}}) + \delta t \geq \overline{\mathbb{P}}_C^{\mathrm{reach}\leq t}(l_{\mathrm{init}}, \mathbf{x}_{\mathrm{init}})$.

Now we do a similar construction for L-γ-SclSubM. Conceptually, we do the same as for UNNRepSupM, that is, we fix a template $\eta_l = c_1^l x_1 + \cdots + c_n^l x_n + d^l$ and encode the submartingale conditions into formulas of the form $\phi \implies \psi$. The first two sets of formulas $\mathfrak{F}_1^l$ and $\mathfrak{F}_2^l$ are identical to the UNNRepSupM case since L-γ-SclSubM need to satisfy the same conditions, namely, they are nonnegative in I and take value 1 inside of the target region C. The only difference is that we need to encode the submartingale conditions $\eta \leq \gamma \underline{\mathbb{X}}\eta$ ($\equiv \gamma \underline{\mathbb{X}}\eta - \eta \geq 0$) which gives a slightly different set of formulas $\mathfrak{F}_3^l$.

- If $l \in L_N$, then

$$\mathfrak{F}_3^l = \Big\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \implies \big((\gamma c_1^{l'} - c_1^l)x_1 + (\gamma c_2^{l'} - c_2^l)x_2 + \cdots + (\gamma c_n^{l'} - c_n^l)x_n + (\gamma d^{l'} - d^l) \geq 0\big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}}, l' \in \mathrm{succ}(l) \Big\}$$

- If $l \in L_P$, then

$$\mathfrak{F}_3^l = \Big\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \implies \big((\gamma \sum_{l \mapsto l'} \Pr(l,l')c_1^{l'} - c_1^l)x_1 + \cdots + (\gamma \sum_{l \mapsto l'} \Pr(l,l')c_n^{l'} - c_n^l)x_n + (\gamma \sum_{l \mapsto l'} \Pr(l,l')d^{l'} - d^l) \geq 0\big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}} \Big\}$$

- If $l \in L_D$, then for each successor $l' \in \mathrm{succ}(l)$ there is a unique guard $G(l,l')$ and

$$\mathfrak{F}_3^l = \Big\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \wedge \bigwedge_{j=1}^{N_{l,l',h}^{\mathfrak{G}}} (g_{h,j}^{l,l'} \geq 0) \implies$$
$$\big((\gamma c_1^{l'} - c_1^l)x_1 + \cdots + (\gamma c_n^{l'} - c_n^l)x_n + (\gamma d^{l'} - d_n^l) \geq 0\big)$$
$$\mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}}, 1 \leq h \leq N_{l,l'}^{\mathfrak{G}}, l' \in \mathrm{succ}(l) \Big\}$$

- If $l \in L_A$, $\mathrm{Up}(l) = (x_v, u)$ and l' is the unique successor of l then
  - if $u = f \in \mathcal{B}(\mathbb{R}^V, \mathbb{R})$ and $f = [\![\mathrm{f}]\!] = [\![r_1x_1 + r_2x_2 + \cdots + r_nx_n + s]\!]$ is a linear expression, then

$$\mathfrak{F}_3^l = \Big\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \implies$$
$$\big((\gamma(c_1^{l'} + r_1c_v^{l'}) - c_1^l)x_1 + \cdots + (\gamma r_v c_v^{l'} - c_v^l)x_v + \cdots +$$
$$(\gamma(c_n^{l'} + r_nc_v^{l'}) - c_n^l)x_n + (\gamma(d^{l'} + sc_v^{l'}) - d^l) \geq 0\big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}} \Big\}$$

  - if $u = d \in \mathcal{D}(\mathbb{R})$ and $\mathbb{E}(d) = s$, then

$$\mathfrak{F}_3^l = \Big\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \implies \big((\gamma c_1^{l'} - c_1^l)x_1 + \cdots + (\gamma s - c_v^l)x_v$$
$$+ \cdots + (\gamma c_n^{l'} - c_n^l)x_n + (\gamma d^{l'} - d^l) \geq 0\big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}} \Big\}$$

  - if $u = [\![\mathfrak{A}]\!] = [\![\bigvee_{k=1}^{N_l^{\mathfrak{A}}} \bigwedge_{j=1}^{N_{l,k}^{\mathfrak{A}}} (p_{k,j}^l x_v + q_{k,j}^l \geq 0)]\!] \in \mathcal{B}(\mathbb{R})$ is given by a linear predicate over $\{x_v\}$, then

$$\mathfrak{F}_3^l = \Big\{ \bigwedge_{j=1}^{N_{l,i}^{\mathfrak{J}}} (\alpha_{i,j}^l \geq 0) \wedge \bigwedge_{m=1}^{N_l^{\mathfrak{C}}} (-\beta_{m,j_m}^l \geq 0) \wedge \bigwedge_{j=1}^{N_{l,k}^{\mathfrak{A}}} (p_{k,j}^l x_{\mathrm{ndet}} + q_{k,j}^l \geq 0) \implies$$
$$\big((\gamma c_1^{l'} - c_1^l)x_1 + \cdots + (-c_v^l)x_v + \cdots + (\gamma c_n^{l'} - c_n^l)x_n + (\gamma c_v^{l'})x_{\mathrm{ndet}}$$
$$+ (\gamma d^{l'} - d^l) \geq 0\big) \mid 1 \leq i \leq N_l^{\mathfrak{J}}, 1 \leq j_m \leq N_{l,m}^{\mathfrak{C}}, 1 \leq k \leq N_l^{\mathfrak{A}} \Big\}$$

This set of formulas has two key differences to the previous formulas for UNNRepSupM. Firstly, the differences are swapped around resulting from encoding $(\gamma \underline{\mathbb{X}}\eta - \eta \geq 0)$ instead of $(\eta - \overline{\mathbb{X}}\eta \geq 0)$. Secondly, the factor γ is added. The condition $(\gamma \underline{\mathbb{X}}\eta - \eta \geq 0)$ is satisfied

if and only if $(\gamma \mathbb{X}\eta(\cdot,\sigma) - \eta \geq 0)$ is satisfied for every possible scheduler. This observation allows us to encode it as a conjunctive formula again.

Observe that Farkas' Lemma transforms the satisfiability problem into a feasibility problem of the form $y^T A = (\gamma c + c')^T$ and $y^T b \leq \gamma d + d'$ where $y \geq 0$. For any fixed parameter γ, this is just a linear feasibility problem. One might want to consider γ itself as a parameter to this problem and optimize for that as well. This, however, does not seem to bring any benefits. First of all, the problem becomes a quadratically constrained linear program (QCLP) which is harder to solve. Secondly, since γ needs to be strictly less than 1, we have to introduce a constraint $\gamma < 1$ resulting in a non-closed search space where no optimum may exist. Lastly, larger values of γ are always desirable since if η is a γ-SclSubM, then it is also a γ'-SclSubM for any $\gamma' > \gamma$, so the optimizer will necessarily want to bring γ as close as possible to 1 and as a result will not reach an optimal value.

Theorem 4.1.3 (Synthesis theorem of L-γ-SclSubM). *Let* $\mathrm{sol} : \mathcal{A} \cup \{y_1, \ldots, y_m\} \to \mathbb{R}$ *be a solution to the optimization problem induced by the above family of formulas* $\mathfrak{F}_i^l$ *where* $i \in \{1,2,3\}$ *and* $l \in L$*, then the function* $\eta(l, \mathbf{x}) = \mathrm{sol}(c_1^l)x_1 + \mathrm{sol}(c_2^l)x_2 + \cdots + \mathrm{sol}(c_n^l)x_n + \mathrm{sol}(d^l)$ *is a L-*γ*-SclSubM for C supported by the invariant I.*

Since γ-scaled submartingales give lower bounds on the reachability probability, we naturally want to maximize its value. This is easily done by setting the objective function to $[\![\eta_{l_{\text{init}}}]\!](\mathbf{x}_{\text{init}})$ (same as for UNNRepSupM) and maximizing it.

4.2 Polynomial Templates

In this section we take a look into polynomial templates-based synthesis methods. The synthesis procedure is more difficult compared to simple linear templates but polynomial templates can achieve much better bounds. Consider for example a pendulum-like setting, where the reachability probability is highest at the endpoints of some interval, but small at the center. E.g. assume that a variable x takes values in $[-1,1]$ and the true reachability probability is given by $\mathbb{P}_C^{\text{reach}}(x) = x^2$, then the highest reachability probability of 1 is obtained at $x = -1$ and $x = 1$. The best linear overapproximation of this probability has to take value 1 at the points $x = 1$ and $x = -1$ which necessarily means it has to take value 1 in between as well. No linear supermartingale can be better than the trivial upper bound in this case. A polynomial of degree 2, however, is able to give tight bounds.

The setting is as follows. As in the linear case, we assume a pCFG Γ over a finite set of variables, an invariant I and a target region C are given. Unlike before, we allow all expressions and predicates to be polynomial, that is, they are induced by the grammar

$$\langle a \rangle ::= r \mid x \mid \langle a \rangle \cdot \langle a \rangle \mid \langle a \rangle + \langle a \rangle$$

The template we fix is of the form $\eta_l = \sum c_i^l m_i + d^l$ where m_i is a monomial over the variables V with degree less than or equal to $k \in \mathbb{N}$. In other words, we fix a maximum degree $k \in \mathbb{N}$ and let η^l be a linear combination of all those monomials with degree bounded by k. We denote this set of monomials by $\mathcal{M}_k$ and the set of coefficients by $\mathcal{A}$. The size of η^l is $\Theta(V^{k+1})$. Furthermore, we assume that for any used distribution d in the pCFG, all higher moments $\mathbb{E}(d^m) := \int x^m \mathrm{d}d$ are computable. Actually it is sufficient to be able to

compute all moments up to the $(k \cdot k')$-th moment where k' is the highest degree among all polynomial expressions used in assignments. The reason is that while η has degree at most k, the degree of $\mathbb{X}\eta$ can be $(k \cdot k')$ after an assignment of a k'-th degree polynomial to a variable x_j.

The construction of the formula sets $\mathfrak{F}_1^l$ to $\mathfrak{F}_3^l$ follow the same pattern as in the linear case. The only difference is that all formulas $\phi \implies \psi$ constructed this way are polynomial, i.e., the formula ϕ is a polynomial conjunctive predicate over the variables $V' = V \cup \{x_{\text{ndet}}\}$ (of degree at most k') and the formula ψ is a polynomial inequality (of potentially degree up to $k \cdot k'$) over V' with coefficients that are affine linear over $\mathcal{A}$. Concretely we have,

- $\phi \equiv f_1 \geq 0 \wedge f_2 \geq 0 \wedge \cdots \wedge f_r \geq 0$ where each f_i is a polynomial expression over the variables V', and
- $\psi \equiv \sum_{m \in \mathcal{M}_z} c_m m \geq 0$ where m is a monomial over V' and c_m is a affine linear expression over the coefficient set $\mathcal{A}$.

This is again a satisfiability problem and may be solved using the decidability of first-order theories of the real numbers. Decision procedures for this general problem are not practical, especially considering the size of the formula set. A more practical approach is to use variations of the so called *Positivstellensatz* which characterizes polynomials that are positive on some semialgebraic set.

Definition 4.2.1 (Semialgebraic set)**.** A set $W \subseteq \mathbb{R}^n$ is called semialgebraic if it can be described as a finite union of sets

$$W_k = \{x \in \mathbb{R}^n \mid \forall i \in I_k : P_i(x_1, x_2, \ldots, x_n) \geq 0 \wedge \forall j \in J_k : Q_j(x_1, x_2, \ldots, x_n) = 0\}$$

where I_k and J_k are finite index sets and P_i and Q_j are real polynomials over $\{x_1, \ldots, x_n\}$.

Note that the formula ϕ exactly represents a semialgebraic set W_ϕ. Satisfiability of the implication $\phi \implies \psi$ (from above) is then equivalent to showing that the polynomial defining ψ is nonnegative on W_ϕ. There are multiple different variations of Positivstellensatz that may be used to characterize nonnegative (positive) polynomials on such sets, but we will consider the version by Schmüdgen.

Proposition 4.2.1 (Schmüdgen's Positivstellensatz[32])**.** *If the semialgebraic set $W = \{x \in \mathbb{R}^n \mid \forall f \in F : f(x) \geq 0\}$ defined by a finite set $F = \{f_1, \ldots, f_r\}$ of polynomials is compact, then each polynomial $p \in \mathbb{R}[x_1, \ldots, x_n]$ that is strictly positive on W (i.e. $\forall x \in W : p(x) > 0$) can be written as*

$$p = \sum_{\alpha \in \{0,1\}^r} \sigma_\alpha f_1^{\alpha_1} f_2^{\alpha_2} \ldots f_r^{\alpha_r}$$

where σ_α is is a sum-of-squares (SOS) polynomial over $\{x_1, \ldots, x_n\}$.

A sum-of-squares polynomial σ is a polynomial of the form $\sigma = \sum_{i \in I} h_i^2$ where each h_i is a polynomial. Note that the degree of σ is necessarily even. It is easy to see that sum-of-square polynomials are everywhere positive and the polynomials f_i in Schmüdgen's Positivstellensatz are nonnegative on W. Therefore every polynomial p defined this way

is necessarily nonnegative on W, and the Positivstellensatz guarantees that every strictly positive polynomial can actually be characterized this way under some assumptions.

If we consider an implication $\phi \implies \psi$ again, then it is of the form $f_1(x) \geq 0 \wedge f_2(x) \geq 0 \wedge \cdots \wedge f_r(x) \geq 0 \implies p(x) \geq 0$ where p is a polynomial with coefficients in $\mathcal{A}$. The premise exactly defines the semialgebraic set W and if we change the inequality to a strict one on the righthand side, then we can apply Schmüdgen's Positivstellensatz assuming W is compact. If it is not compact then we might still be able to express p as described but we lose any guarantees. This method is sound but not complete. Assuming now that p is of the form $p = \sum_{\alpha \in \{0,1\}^r} \sigma_\alpha f_1^{\alpha_1} f_2^{\alpha_2} \dots f_r^{\alpha_r}$, we need to synthesize the sum-of-square coefficients σ_α. The following characterization of SOS polynomials is useful.

Proposition 4.2.2 (Characterization of SOS polynomials). *A SOS polynomial $\sigma \in \mathbb{R}[x_1, \dots, x_n]$ of degree at most $2k$ can be expressed as*

$$\sigma = \mathfrak{m}^T A \mathfrak{m}$$

where $\mathfrak{m}$ is a vector of all monomials $\mathcal{M}_k$ (monomials of degree at most k) and $A \in \mathbb{R}^{|\mathcal{M}_k| \times |\mathcal{M}_k|}$ is a positive-semidefinite matrix.

Combining this characterization of SOS polynomials with Schmüdgen's Positivstellensatz, we can express a positive polynomial p on a semialgebraic set W defined by $F = \{f_1, \dots, f_r\}$ as

$$p = \sum_{\alpha \in \{0,1\}^r} (\mathfrak{m}^T A^\alpha \mathfrak{m}) f_1^{\alpha_1} f_2^{\alpha_2} \dots f_r^{\alpha_r}$$

where A^α is a positive-semidefinite matrix. We fix some maximum degree z for the sum-of-square polynomials and let $u = |\mathcal{M}_z|$,

$$A^\alpha = \begin{bmatrix} a^\alpha_{1,1} & \dots & a^\alpha_{1,u} \\ \vdots & \ddots & \vdots \\ a^\alpha_{u,1} & \dots & a^\alpha_{u,u} \end{bmatrix}$$

and add the all entries $a^\alpha_{i,j}$ into a set $\mathcal{A}'$ of unknown parameters. Note that the degree of the SOS polynomials may be different than the degree of the martingale we want to synthesize. We can choose a large enough degree to guarantee that we find a solution if it exists, or use lower degree polynomials to speed up the synthesis process. Either way, we have to fix a maximum degree $z \in \mathbb{N}$. Then for any given implication $f_1(x) \geq 0 \wedge \cdots \wedge f_r(x) \geq 0 \implies p(x) \geq 0$ with p having coefficients in $\mathcal{A}$, we can set $p = \sum_{\alpha \in \{0,1\}^r} (\mathfrak{m}^T A^\alpha \mathfrak{m}) f_1^{\alpha_1} f_2^{\alpha_2} \dots f_r^{\alpha_r}$ and by comparing coefficients of the monomials on each side, we get a set of linear equalities over the variables $\mathcal{A} \cup \mathcal{A}'$ (one equality for each monomial). The only non-linear conditions we have are the semi-definiteness conditions on A^α. This problem is readily solvable by semi-definite programming[33].

There are other versions of the Positivstellensatz which have been successfully applied in practice (see [9]) such as Putinar's Positivstellensatz[34] or Handelman's Theorem[35]. These are less general than Schmüdgen's version but allow a much more succinct description of the polynomial p, which in turn accelerates the synthesis procedure.

4.3 Experiments

For the experiments we synthesize the 3 following types of martingales via linear template synthesis.

1. UNNRepSupM to overapproximate $\overline{\mathbb{P}}_C^{\text{reach}}$,
2. L-γ-SclSubM to underapproximate $\underline{\mathbb{P}}_C^{\text{reach}}$,
3. UARnkSupM to overapproximate $\overline{\mathbb{E}}_C^{\text{steps}}$ and decide a.s. termination.

We have not implemented the polynomial synthesis algorithm but for results on that we refer to

- Polynomial additive ranking supermartingales by Chaterjee et al.[9]
- Polynomial UNNRepSupM and L-γ-SclSubM by Takisaka et al. [11]
- Polynomial higher order ranking supermartingales by Kura et al. [14]

In order to specify the invariant and the target region conveniently, we enhanced the syntax of APP by two new constructs to annotate locations.

```
        {x = 0}                  // Invariant annotation
1.      x := unif(0,1);
        {0 <= x and x <= 1}      // Invariant annotation
        [x <= 0.5]               // Target region annotation
2.      skip;
```

The first location is annotated with the invariant $x = 0$ and the second location is annotated with an invariant as well as a target region. The target region is $C = \{(l, x) \in (L \times \mathbb{R}) \mid l = 2 \wedge x \leq 0.5\}$. In this case, the annotated skip statement is equivalent to the statement "assert $<= 0.5$".

The first examples we look into are variations of the 1d random walk (Listing 4.1) with different starting values and transition probabilities. We start from a value $s = N \in [0, 50]$ and increase that value by 1 with probability P or decrease it by 1 with probability $(1 - P)$ until the value reaches either one of the bounds 0 or 50. We are interested in the probability that it reaches 50, i.e. $\Pr(\Diamond(s >= 50) \mid s_0 = N, P)$.

Parameters (N, P)	$\overline{\mathbb{P}}_C^{\text{reach}}$	UNNRepSupM	L-γ-SclSubM ($\gamma = 0.9999$)
(25, 0.5)	$= 0.5$	≤ 0.5	≥ 0
(35, 0.5)	$= 0.7$	≤ 0.7	≥ 0
(15, 0.5)	$= 0.3$	≤ 0.3	≥ 0
(25, 0.51)	≈ 0.7311	≤ 1	> 0.4896
(10, 0.51)	≈ 0.3813	≤ 1	> 0.1836
(25, 0.49)	≈ 0.2689	≤ 0.5	≥ 0
(40, 0.49)	≈ 0.6187	≤ 0.8	≥ 0
(25, 0.2)	< 0.0001	≤ 0.5	≥ 0
(15, 0.8)	> 0.9999	≤ 1	> 0.2856
(40, 0.35)	≈ 0.002	≤ 0.8	≥ 0

Table 4.1: Probability bounds for 1d-Random Walk

Listing 4.1: 1d Random Walk

```
s := N;
{s >= 0 and s <= 50}
while (s >= 1) do
        {s >= 0 and s <= 50}
        while s >= 50 do
                {s = 50}
                assert true;
        od;
        {s >= 1 and s <= 49}
        if prob(P) then
                {s >= 1 and s <= 49}
                s := s + 1;
        else
                {s >= 1 and s <= 49}
                s := s - 1;
        fi;
od;
assert false;
```

In the first three cases, the true probability function is indeed a linear function so the linear synthesis procedure for UNNRepSupM is able to find it as shown in Table 4.3. Also, the expected number of steps to termination in those cases is infinite. If T denotes the first hitting time, then the expectation $\mathbb{E}(T)$ is infinite for all cases and in particular the best possible γ-scaled submartingale is $\mathbb{E}(\gamma^T)$ which equals 0 if $T = \infty$. Also note that the best possible linear function in s that overapproximates $\mathbb{P}_C^{\text{reach}}$ at l_{init} is $\eta(l_{\text{init}}, s) = s/50$ since it is the least linear function that satisfies $\eta(l_{\text{init}}, 0) \geq 0 = \overline{\mathbb{P}}_C^{\text{reach}}(l_{\text{init}}, 0)$ and $\eta(l_{\text{init}}, 50) \geq 1 = \overline{\mathbb{P}}_C^{\text{reach}}(l_{\text{init}}, 50)$. This function was found whenever possible. The γ-scaled supermartingale is stricly below the optimal linear bound resulting from the γ-scaling. In all cases, only one of the supermartingales is able to give non-trivial bounds.

For the next example we consider the cooling system (Listing 4.2) of a refrigerator which constantly tries to cool down the refrigerator from a starting temperature until it gets close to a desired target temperature. Working against this process is a natural temperature exchange between the refrigerator and the outside. We want to make sure that the refrigerator is guaranteed to reach its target temperature eventually.

For the concrete scenario, we start with a non-deterministic outside temperature given by some closed interval $[T_1, T_2]$ and a starting temperature T_{init} of the refrigerator. In every iteration, the refrigerator first exchanges temperature with the outside in the following manner: with a probability of 0.9 the refrigerator is closed and the exchange is slow; with a smaller probability of 0.1 the refrigerator is open and the exchange is faster. Afterwards, the cooling system cools the refrigerator towards the target temperature T_{targ}. This process continues until the target temperature is reached within an error of some $\epsilon > 0$. The exchange rates are dependent on the temperature gradients between the refrigerator's inside and the outside or the target, respectively. We note that the expected reaching time is measured in transitions, not loop iterations.

Listing 4.2: Refrigerator cooling system

```
// I_global := {0 <= f and f <= oT and T_1 <= oT and oT <= T_2}

oT := ndet(Real[T_1, T_2]);
{T_1 <= oT and oT <= T_2}
f := T_init;
{I_global}
while (f - T_targ >= ε) do
        {I_global and f - T_targ >= ε}
        if prob(0.1) then
                {I_global and f - T_targ >= ε}
                f := f + 0.2*(oT - f);
        else
                {I_global and f - T_targ >= ε}
                f := f + 0.05*(oT - f);
        fi;
        {I_global and f - T_targ >= ε}
        f := f - 0.5*(f - T_targ);
od;
assert true;
```

The results in Table 4.3 show that, while the overapproximation of the upper expected reaching time is orders of magnitude off, the ranking supermartingale method is able to decide positive almost sure termination whenever possible. As a reference, we also used γ-scaled martingales to underapproximate the true reachability probability. For the first two cases, the UARnkSupM is able to show positive a.s. termination for different starting temperatures. In the third case, the outside temperature can be 24 units, and the cooling system is not able to overcome the natural heat exchange in this case and reach close to the

Parameters ($[T_1, T_2]$, T_{init}, T_{targ}, ϵ)	$\overline{\mathbb{E}}_C^{\text{steps}}$	UARnkSupM	L-γ-SclSubM ($\gamma = 0.9999$)
([20, 23], 15, 7, 1)	≈ 33	≤ 4803	> 0.6175
([20, 23], 23, 7, 1)	≈ 37	≤ 7363	> 0.4137
([20, 24], 15, 7, 1)	∞	$\leq \infty$	≥ 0
([20, 23], 15, 7, 0.977)	≈ 53.52	≤ 237627	≥ 0
([20, 23], 15, 7, 0.9765)	∞	$\leq \infty$	≥ 0
([20, 23], 15, 6.5, 1)	∞	$\leq \infty$	≥ 0
([20, 23], 15, 6.7, 1)	≈ 41.16	≤ 21822	≥ 0

Table 4.2: Expected cooling time

target temperature of 7 units. In the fifth case, the cooling system in not able to go below a temperature of 7.9765 units, while case (4) shows that it will eventually go below 7.977. Comparing the parameter sets (4) and (5), we see that the inspected case (4) is very close to not be positive a.s. terminating, yet a ranking supermartingale could still be synthesized, despite its huge inaccuracy. For the last two cases we changed the target temperature, which also accelerates the cooling process. A target temperature of 6.5 is not positive almost sure reachable, while 6.7 is eventually reached. In all cases, the γ-scaled submartingale gave less information about the process than the ranking supermartingale.

The high discrepancy between the true expected reaching time and the ranking supermartingale can be explained when considering the sequence of refrigerator temperatures after each loop iteration. Generally, the temperature falls steeply during the first few iterations but quickly starts to change in small decrements (exponential decay). In case (1) for example, if the temperature of the refrigerator is exactly at the boundary of 8 units, then an iteration brings it down to 7.9875 which equals a decrement of 0.0125 units. Since the ranking supermartingale has to overapproximate this decrease, it can decrease by at most 0.0125 units, and by linearity, it must globally decrease by at most 0.0125 per iteration. From a starting temperature of 15 with a decrease of at most 0.0125 per iteration, it takes at least 560 loop iterations, equaling 2240 transitions, until the temperature goes below 8, that is, the loop terminates. Furthermore, the ranking supermartingale has to be nonnegative across the whole invariant, including the never reached case of a zero temperature inside the refrigerator. And dropping down to 0 instead of 8 under the exact same conditions does indeed take 4800 transitions inside the loop, and 3 outside, totaling the computed 4803. Doing the same calculation for a starting temperature of 23 units, we see that is takes 7363 transitions to reach 0. The synthesized linear ranking supermartingale is thus optimal under the given conditions. Raising the lower bound of the refrigerator's temperature from 0 to some higher value does improve the computed expected reaching times.

5 Related Work

This chapter discusses work related to the topic of this thesis. This serves to inform the reader on similar work for further reading and comparison.

The most similar work to this thesis is done by Takisaka et al.[11] and Kura et al.[14]. Indeed, much of thesis is inspired by their work and our notation follows their's very closely. Takisaka et al. derive martingale-based methods purely via their fixed point theoretic framework based on Knaster-Tarski's theorem as well as Cousot-Cousot's theorem. Their reasoning principles do not rely on probabilistic theory at all and yet derive mostly the same results as was commonly done before using martingale theory. In this thesis we make a direct comparison between their fixed point theoretic approach and commonly used martingale theory. To our knowledge, the results of Kura et al. regarding higher order ranking supermartingales are completely new, in the sense that higher order ranking supermaringales have not been characterized via martingale-theoretic methods. Here the fixed point theoretic approach seems to have an advantage.

Chatterjee et al. use common martingale theory in [9] and [36] to argue about positive almost sure termination via additive ranking supermartingales and show how polynomial-template synthesis can be performed using Positivstellensatz. A similar work was done by Chakarov et al. in [12]. In another work for theirs [13], they combine the idea of repulsing supermartingales and ranking martingales together with a new notion of probabilistic invariants to deduce upper bounds on reachability probabilities. We briefly discussed this idea in Subsection 3.2.3 and explained it advantages. In short, a probabilistic invariant is a predicate which is not guaranteed to be an invariant but may be one only with a certain probability. The idea is that this kind of predicate may be stronger than pure invariants, and possibly strong enough to deduce positive almost sure reachability via ranking supermartingales. A lower bound on this probability can be computed using repulsing supermartingales.

Chakarov et al. give sound rules to decide the qualitative properties of almost sure recurrence and almost sure persistence[27]. Similar to our work on recurrence in Subsection 3.2.6, they give slight variations of ranking supermartingales to witness almost sure recurrence but also almost sure persistence. They use polynomial synthesis methods to find such supermartingale expressions automatically. Unlike our work, they do not consider recurrence probabilities at all.

Chakarov and Chatterjee focus their work directly on polynomial systems from the very beginning and do not develop a more general theory on general probabilistic programs. Furthermore, they generally only deal with a single form of non-determinism

In a different setting, Avanzini et al. consider probabilistic term rewriting[7], an extension

of usual term rewriting. Unlike standard programs, term rewriting naturally incorporates non-determinism in form of choice of the rewriting rule to be applied. Avanzini et al. show that probabilistic ranking functions are applicable in term rewriting to decide almost sure termination. Their work uses polynomial as well as matrix interpretations to automatically derive ranking expressions. Unlike the other works, they employ SMT-based synthesis instead of optimization-based synthesis.

In a different direction, Kaminski et al. [37] use a weakest precondition calculus to compute expected running times of programs and hence decide almost sure termination. They use a runtime transformer which easily computes the expected runtime of loop-free programs but needs a least fixed point computation to handle loops. In order to avoid this fixed point computation they consider using approximations via upper invariants and incremental invariants to bound loop runtimes.

In comparison to the mentioned works, this thesis gives an overview over various martingale-based methods and derives known results in a uniform manner. In particular, we give a unified view on the fixed point theoretic approach as well as the martingale based approach and establish the close resemblance between martingales and fixed points.

6 Conclusion

This thesis shows that the martingale-based approach to the verification of probabilistic programs can be developed from two quite distinctive theoretic foundations. The martingale theoretic approach comes from the probabilistic nature of probabilistic programs and as such heavily relies on probability theory. On the other hand, the fixed point theoretic approach puts the focus on the programs whose semantics are classically characterized by fixed point equations (e.g. semantics of loops). As such, the fixed point theoretic approach is often adopted in verification of many deterministic systems. This work shows that it can also handle probabilistic programs effectively in a uniform manner. The key observation is that the concept of probabilistic martingales, sub- and supermartingales are very closely related to order-theoretic fixed points, post-fixpoints and pre-fixpoints, respectively.

Martingale-based techniques are not restricted to the classical boolean decision problems of deciding certain properties (e.g. termination, safety or liveliness) in typical deterministic verification, but they also naturally extend to quantitative reasoning about probabilities of programs. They theoretically span a variety of problems, ranging from deciding boolean almost sure termination, reachability and recurrence to quantitatively lower and upper bounding the respective probabilities. Martingales are also naturally able to handle nondeterminism; an even further extension of probabilistic programs.

Furthermore, martingales of certain shape can automatically be found by employing well-known optimization procedures such as linear programming and semi-definite programming. While the martingale-based approach generally gives sound and oftentimes even complete ways to show certain properties or bounds on probabilistic programs, the practical application via template-based synthesis seems to lack in many ways. Especially linear repulsing martingales tend to give very bad or even trivial results about rather simple programs.

An obvious problem here is that many programs that are not surely terminating, especially ones containing loops have a termination probability that is exponential in the variables. The example of a random walk clearly shows that linear supermartingales can work very well and give tight bounds in the unbiased case because the true termination probability is indeed a linear function. However, a slight bias already gives a exponential true probability function which is only badly approximated by linear functions. Indeed, even polynomial supermartingales are unable to give tight bounds in this case. On the other hand, even if a linear ranking supermartingale is off by orders of magnitude, as long as they witness finiteness they can witness termination, making them more useful in general. The approach by Chatterjee et al. in [13] by using probabilistic invariants might be the right approach to leverage the robustness of ranking supermartingales against inaccuracies and to use them in conjunction with repulsing supermartingales to witness probabilistic bounds.

7 Future Work

The presented theory and practical methods in this work allow further research and extensions in possible future work. In this chapter we summarize some of these possibilities.

Extending to Other Martingale-based Methods. We only considered a few of the known martingale-based methods in this work and can potentially enlarge this unified framework to include other martingale-based methods, such as supermartingales for persistence[27] and nonnegative repulsing ϵ-supermartingales[13] for reachability probabilities. The former of these two is a supermartingale that strictly drives away from the target region by some fixed value ϵ in each transition. In this sense, it has similarities to ranking supermartingales which drive towards the target region with a lower bounded step size. These repulsing supermartingales may be used to derive upper bounds on reachability probabilities.

Improved Synthesis. Linear template-based synthesis turned out to lack accuracy to give meaningful probabilistic bounds. An obvious extension as mentioned in Chapter 4 is by using polynomial templates in conjunction with semi-definite programming to achieve better bounds. Another possibility is to use piece-wise linear or polynomial templates which can break down the problem of globally overapproximating into locally overapproximating. This could give tighter bounds in many cases. Ideally, we would like find exponential martingales which are much more capable of tracking probabilities in looped programs, but we do not know of possible ways of synthesizing exponential functions automatically.

Stronger Selection Theorems. For many proofs we needed the existence of ϵ-optimal schedulers, which was not guaranteed for general universally measurable functions. This lead to the need of extra existence assumptions. For lower semianalytic functions it is known that ϵ-optimal selection is possible as proven in [18]. The fundamental theorem used to prove this result is Jankov-von Neumann's Selection Theorem for analytic sets. There exist stronger versions of this theorem, some of which deal with (optimal) selection related to universally measurable sets (e.g. in [38]). Maybe these can be used to find a more general theorem that shows the existence of ϵ-optimal schedulers for a bigger class of functions, ideally the class of universally measurable functions.

Bibliography

[1] N. Dershowitz, "Term Rewriting Systems by Terese (Marc Bezem, Jan Willem Klop, and Roel de Vrijer, eds.), Cambridge University Press, Cambridge Tracts in Theoretical Computer Science 55, 2003, hard cover: ISBN 0-521-39115-6, xxii+884 pages," *Theory and Practice of Logic Programming*, vol. 5, pp. 395–399, May 2005.

[2] N. Dershowitz, "Termination of rewriting," *Journal of Symbolic Computation*, vol. 3, pp. 69–115, Feb. 1987.

[3] B. Cook, A. Podelski, and A. Rybalchenko, "Proving program termination," *Communications of the ACM*, vol. 54, p. 88, May 2011.

[4] J. Tsinias, N. Kalouptsidis, and A. Bacciotti, "Lyapunov functions and stability of dynamical polysystems," *Mathematical systems theory*, vol. 19, pp. 333–354, Dec. 1986.

[5] H. Nakamura, Y. Yamashita, and H. Nishitani, "Lyapunov functions for homogeneous differential inclusions," *IFAC Proceedings Volumes*, vol. 37, pp. 733–738, Sept. 2004.

[6] O. Bournez and F. Garnier, "Proving Positive Almost-Sure Termination," in *Term Rewriting and Applications* (J. Giesl, ed.), Lecture Notes in Computer Science, pp. 323–337, Springer Berlin Heidelberg, 2005.

[7] M. Avanzini, U. D. Lago, and A. Yamada, "On Probabilistic Term Rewriting," *arXiv:1802.09774 [cs]*, Feb. 2018. arXiv: 1802.09774.

[8] S. Agrawal, K. Chatterjee, and P. Novotný, "Lexicographic Ranking Supermartingales: An Efficient Approach to Termination of Probabilistic Programs," *arXiv:1709.04037 [cs]*, Sept. 2017. arXiv: 1709.04037.

[9] K. Chatterjee, H. Fu, and A. K. Goharshady, "Termination Analysis of Probabilistic Programs through Positivstellensatz's," *arXiv:1604.07169 [cs]*, Apr. 2016. arXiv: 1604.07169.

[10] O. Bournez and F. Garnier, "Proving Positive Almost-Sure Termination," in *Term Rewriting and Applications* (D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, and J. Giesl, eds.), vol. 3467, pp. 323–337, Berlin, Heidelberg: Springer Berlin Heidelberg, 2005.

[11] T. Takisaka, Y. Oyabu, N. Urabe, and I. Hasuo, "Ranking and Repulsing Supermartingales for Reachability in Probabilistic Programs," *arXiv:1805.10749 [cs]*, vol. 11138, pp. 476–493, 2018. arXiv: 1805.10749.

[12] A. Chakarov and S. Sankaranarayanan, "Probabilistic Program Analysis with Martingales," in *Computer Aided Verification* (N. Sharygina and H. Veith, eds.), Lecture Notes in Computer Science, pp. 511–526, Springer Berlin Heidelberg, 2013.

[13] K. Chatterjee, P. Novotný, and o. ikeli, "Stochastic Invariants for Probabilistic Termination," *arXiv:1611.01063 [cs]*, Nov. 2016. arXiv: 1611.01063.

[14] S. Kura, N. Urabe, and I. Hasuo, "Tail Probabilities for Randomized Program Runtimes via Martingales for Higher Moments," *arXiv:1811.06779 [cs]*, Nov. 2018. arXiv: 1811.06779.

[15] N. Urabe, M. Hara, and I. Hasuo, "Categorical Liveness Checking by Corecursive Algebras," *arXiv:1704.04872 [cs]*, Apr. 2017. arXiv: 1704.04872.

[16] E.-E. Doberkat, "Measures and all that — A Tutorial," *arXiv:1409.2662 [math]*, Sept. 2014. arXiv: 1409.2662.

[17] V. Bogachev, *Measure Theory*. Berlin Heidelberg: Springer-Verlag, 2007.

[18] D. P. Bertsekas and S. E. Shreve, "Stochastic optimal control: The discrete-time case," in *Stochastic Optimal Control: The Discrete-Time Case*, 1978.

[19] J. L. Doob, "What is a martingale?," *The American Mathematical Monthly*, vol. 78, no. 5, pp. 451–463, 1971.

[20] A. Granas and J. Dugundji, *Elementary Fixed Point Theorems*, pp. 9–84. New York, NY: Springer New York, 2003.

[21] R. Cousot and P. Cousot, "Constructive versions of tarski's fixed point theorems.," *Pacific Journal of Mathematics*, vol. 82, no. 1, pp. 43–57, 1979.

[22] H. Shousong and Z. Qixin, "Stochastic optimal control and analysis of stability of networked control systems with long delay," *Automatica*, vol. 39, pp. 1877–1884, Nov. 2003.

[23] C. Graham and D. Talay, *Stochastic Algorithms*, pp. 213–230. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013.

[24] Z. Ghahramani, "Probabilistic machine learning and artificial intelligence," *Nature*, vol. 521, pp. 452–459, May 2015.

[25] S. J. Bradtke and M. O. Duff, "Reinforcement Learning Methods for Continuous-Time Markov Decision Problems," p. 8.

[26] I. C. Dolcetta and H. Ishii, "Approximate solutions of the bellman equation of deterministic control theory," *Applied Mathematics and Optimization*, vol. 11, pp. 161–181, Feb 1984.

[27] A. Chakarov, Y.-L. Voronin, and S. Sankaranarayanan, "Deductive Proofs of Almost Sure Persistence and Recurrence Properties," in *Tools and Algorithms for the Construction and Analysis of Systems* (M. Chechik and J.-F. Raskin, eds.), vol. 9636, pp. 260–279, Berlin, Heidelberg: Springer Berlin Heidelberg, 2016.

[28] S. P. Meyn and R. L. Tweedie, "A Survey of Foster-Lyapunov Techniques for General State Space Markov Processes," p. 13.

[29] C. Baier and J.-P. Katoen, *Principles of model checking*. Cambridge, Mass: The MIT Press, 2008. OCLC: ocn171152628.

[30] A. Schrijver, *Theory of linear and integer programming*. Wiley-Interscience series in discrete mathematics and optimization, Chichester: Wiley, reprinted ed., 2000. OCLC: 247967491.

[31] L. van den Dries, "Alfred tarski's elimination theory for real closed fields," *The Journal of Symbolic Logic*, vol. 53, no. 1, pp. 7–19, 1988.

[32] K. Schmüdgen, "Thek-moment problem for compact semi-algebraic sets," *Mathematische Annalen*, vol. 289, pp. 203–206, Mar 1991.

[33] R. Fruend, "Introduction to Semidefinite Programming," p. 51.

[34] M. Putinar, "Positive polynomials on compact semi-algebraic sets," *Indiana Univ. Math. J.*, vol. 42, pp. 969–984, 1993.

[35] V. De Angelis and S. Tuncel, "Handelman's theorem on polynomials with positive multiples," in *Codes, Systems, and Graphical Models* (B. Marcus and J. Rosenthal, eds.), (New York, NY), pp. 439–445, Springer New York, 2001.

[36] K. Chatterjee, H. Fu, P. Novotny, and R. Hasheminezhad, "Algorithmic Analysis of Qualitative and Quantitative Termination Problems for Affine Probabilistic Programs," *arXiv:1510.08517 [cs]*, Oct. 2015. arXiv: 1510.08517.

[37] B. L. Kaminski, J. Katoen, C. Matheja, and F. Olmedo, "Weakest precondition reasoning for expected run-times of probabilistic programs," *CoRR*, vol. abs/1601.01001, 2016.

[38] M. F. Sainte-Beuve, "On the extension of von Neumann-Aumann's theorem," *Journal of Functional Analysis*, vol. 17, pp. 112–129, Sept. 1974.

Task

With applications in learning, randomized construction, and control theory, probabilistic programming currently receives considerable attention and is on its way to becoming an established part of the computing infrastructure we interact with and rely on every day. It is thus not surprising that there is similarly high interest in verification techniques for probabilistic programs. A modern such verificication technique are martingale methods. A martingale is a stochastic process that is (in a precise sense) fair, and thus suitable for modeling long-running systems in which the components frequently interact. A classic verification technique, dating back to the analysis of deterministic programs, are fixed points. What is interesting about fixed points is that the domain of computation can be chosen freely (as long as very mild assumptions apply). It is thus natural to ask whether the recent martingal-based methods can be cast in terms of fixed points, and perhaps vice versa. Addressing this question is the task of the present Masters thesis.

The Masters thesis should compare martingale-based and fixed-point-based verification techniques for probabilistic programs. The correctness properties should include qualitative as well as quantitative aspects and safety as well as liveness concerns. On the algorithmic side, the comparison should be as broad as possible. Once a relationship between martingales and fixed points has been established, generalizations to submartingales and prefix points resp. supermartingales and postfix points shall be studied.