

Verifying Probabilistic Systems: New Algorithms and Complexity Results

Hongfei Fu

The publications of the Department of Computer Science of *RWTH Aachen University* are in general accessible through the World Wide Web.

<http://aib.informatik.rwth-aachen.de/>

Verifying Probabilistic Systems: New Algorithms and Complexity Results

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften
der RWTH Aachen University zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Hongfei Fu, Master of Engineering

aus

Shanghai, Volksrepublik China

Berichter: Prof. Dr.Ir. Joost-Pieter Katoen
Prof. Dr. Antonín Kučera

Tag der mündlichen Prüfung: 21. November, 2014

Diese Dissertation ist auf den Internetseiten der Hochschulbibliothek online verfügbar.

Hongfei Fu
Lehrstuhl für Informatik 2
hongfeifu@cs.rwth-aachen.de

Aachener Informatik Bericht AIB-2014-16

Herausgeber: Fachgruppe Informatik
RWTH Aachen University
Ahornstr. 55
52074 Aachen
GERMANY

ISSN 0935-3232

Die Verifikation Probabilistischer Systeme: Neue Algorithmen und Komplexitätsergebnisse

Hongfei Fu

Abstract

The content of the dissertation falls in the area of *formal verification of probabilistic systems*. It comprises four parts listed below:

1. the decision problem of (probabilistic) simulation preorder between probabilistic pushdown automata (pPDAs) and finite probabilistic automata (fPAs);
2. the decision problem of a bisimilarity metric on finite probabilistic automata (fPAs);
3. the approximation problem of acceptance probability of deterministic-timed-automata (DTA) objectives on continuous-time Markov chains (CTMCs);
4. the approximation problem of cost-bounded reachability probability on continuous-time Markov decision processes (CTMDPs).

The first two parts are concerned with *equivalence checking on probabilistic automata*, where probabilistic automata (PAs) are an analogue of *discrete-time Markov decision processes* that involves both non-determinism and discrete-time stochastic transitions. The last two parts are concerned with *numerical algorithms on Markov jump processes*. In Part 1 and Part 2, we mainly focus on complexity issues; as for Part 3 and Part 4, we mainly focus on numerical approximation algorithms.

In Part 1, we prove that the decision problem of (probabilistic) simulation preorder between pPDAs and fPAs is in EXPTIME. A pPDA is a pushdown automaton extended with probabilistic transitions, and generally it induces an infinite-state PA. The simulation preorder is a preorder that characterizes whether one probabilistic process (modelled as a PA) can mimic the other; technically speaking, it is the one-sided version of (probabilistic) bisimulation, which instead characterizes whether two probabilistic processes are behaviourally equivalent. We demonstrate the EXPTIME-membership of the decision problem through a tableaux system and a partition-refinement algorithm. Combined with the EXPTIME-hardness result by Kučera and Mayr (2010), we are able to show that the decision problem is EXPTIME-complete. The complexity result coincides with the one by Kučera and Mayr (2010) on non-probabilistic pushdown automata. Moreover, we obtain a fixed-parameter-tractable result on this

problem, which again coincides with the counterpart by Kučera and Mayr (2010) on non-probabilistic pushdown automata.

In Part 2, we prove that the decision problem of a bisimilarity metric on fPAs lies in $\text{NP} \cap \text{coNP}$ (and even in $\text{UP} \cap \text{coUP}$). The bisimilarity metric considered here is an undiscounted one defined by van Breugel and Worrell (2005), and is a quantitative extension of (probabilistic) bisimulation in the sense that it measures the distance between states, for which zero distance indicates (probabilistic) bisimilarity. It has a game logical characterization when the underlying model is changed to *stochastic game structures* (instead of fPAs) (cf. de Alfaro *et al.* (2008)). Our result significantly improves the previous complexity results by van Breugel *et al.* (2008) and Chatterjee *et al.* (2010) for the undiscounted case.

In Part 3, we develop a numerical approximation algorithm for acceptance probability by DTA on CTMCs, while correcting errors in the previous work by Chen *et al.* (2011) with new proofs. DTAs are a deterministic subclass of timed automata (by Alur and Dill 1994) which can encode a large class of linear real-time properties. In detail, we present an algorithm which within a given error bound, approximates the probability mass of the set of CTMC-trajectories that satisfy the linear property specified by a multi-clock DTA. As far as we know, this is the first approximation algorithm for acceptance probability by DTA on CTMCs. Previous results such as the ones by Amparore and Donatelli (2010) and Barbot *et al.* (2011) only consider cases where the DTA has only one clock.

In Part 4, we study maximal cost-bounded reachability probability on CTMDPs. In detail, we prove the existence of optimal cost-positional schedulers, where the optimality is considered under all measurable schedulers for CTMDPs. And we develop a numerical approximation algorithm that approximates (within a given error bound) the maximal probability to reach a certain set of target states within a multidimensional cost-bound vector. The time complexity of the algorithm is polynomial in the size of the CTMDP, the unary representation of the cost-bound vector and the reciprocal of the given error bound, and exponential in the dimension of the cost-bound vector. Due to its time complexity, the approximation algorithm is effective for a wide range of applications where the dimension of the cost-bound vector is low. Our results extend the time-bounded case studied by Neuhäüßer and Zhang (2010). Meanwhile, we also point out a proof error in the work on time-bounded case by Neuhäüßer and Zhang (2010) and correct it with new proofs.

Zusammenfassung

Der Inhalt dieser Dissertation fällt in das Gebiet *formaler Verifikation* von probabilistischen Systemen. Die vier Bestandteile sind:

1. Das Entscheidungsproblem von (probabilistischer) Simulationsquasiordnung zwischen probabilistischen Kellerautomaten (pPDAs) und endlichen probabilistischen Automaten (fPAs);
2. Das Entscheidungsproblem einer Bisimulationsmetric auf endlichen probabilistischen Automaten (fPAs);
3. Das Approximationsproblem einer Akzeptanzwahrscheinlichkeit von deterministischen Zeitautomaten auf zeitkontinuierlichen Markow-Ketten (CTMCs);
4. Das Approximationsproblem kostenbeschränkter Erreichbarkeitsprobleme auf zeitkontinuierlichen Markow-Entscheidungsprozessen (CTMDPs);

Die ersten zwei Teile behandeln die *Äquivalenzüberprüfung von probabilistischen Automaten*, wobei probabilistische Automaten (PAs) analog zu *zeitdiskreten Markow-Prozessen* sind, welche sowohl Nichtdeterminismus als auch zeitdiskrete stochastische Transitionen besitzen. Die beiden letzten Teile behandeln *numerische Algorithmen auf Markow-Sprungprozessen*. In Teil 1 und 2 legen wir den Fokus auf Komplexitätsprobleme; in Teil 3 und 4 behandeln wir hauptsächlich numerische Approximationsalgorithmen.

In Teil 1 zeigen wir, dass das Entscheidungsproblem (probabilistischer) Simulationsquasiordnung zwischen pPDAs und fPAs in EXPTIME liegt. Ein pPDA ist ein Kellerautomat welcher um probabilistische Transitionen erweitert ist. Dies induziert im Allgemeinen einen PA mit unendlichem Zustandsraum. Die Simulationsquasiordnung ist eine Quasiordnung, welche charakterisiert, ob ein probabilistischer Prozess (modelliert als PA) einen anderen nachahmen kann; genau genommen ist dies die einseitige Version (probabilistischer) Bisimulation, welche charakterisiert ob zwei probabilistische Prozesse verhaltensäquivalent sind. Wir illustrieren die EXPTIME-Zugehörigkeit des Entscheidungsproblems durch ein Tableausystem und einen Partitionsverfeinerungsalgorithmus. Zusammen mit dem Resultat über die EXPTIME-Schwere von Kučera and Mayr (2010) können wir zeigen, dass

das Entscheidungsproblem EXPTIME-vollständig ist. Das Komplexitätsresultat fällt zusammen mit einem solchen von Kučera and Mayr (2010) für nichtprobabilistische Kellerautomaten.

In Teil 2 zeigen wir, dass das Entscheidungsproblem einer Bisimilaritätsmetrik auf fPAs in $NP \cap \text{coNP}$ liegt (und sogar in $UP \cap \text{coUP}$). Die Bisimulationsmetrik, welche hier betrachtet wird, ist eine undiskontierte, definiert von van Breugel and Worrell (2005). Sie ist eine quantitative Erweiterung (probabilistischer) Bisimulation derart, dass sie die Distanz zwischen Zuständen misst, für welche die Distanz Null Bisimilarität anzeigt. Weiterhin hat sie eine spiellogische Charakterisierung, wenn man das zugrundeliegende Modell zu *stochastischen Spielstrukturen* abändert (anstatt von fPAs) (cf. de Alfaro *et al.* (2008)). Unser Resultat verbessert entscheidend das vorherige Komplexitätsresultat von van Breugel *et al.* (2008) und Chatterjee *et al.* (2010) für den undiskontierten Fall.

In Teil 3 entwickeln wir einen numerischen Approximationsalgorithmus für Akzeptanzwahrscheinlichkeiten von DTAs auf CTMCs, wobei zusätzlich Fehler in einer früheren Arbeit von Chen *et al.* (2011) durch neue Beweise korrigiert werden. DTAs sind eine deterministische Teilklasse von Zeitautomaten (von Alur und Dill 1994), welche eine große Klasse linearer Echtzeiteigenschaften kodieren kann. Genauer gesagt, präsentieren wir einen Algorithmus welcher innerhalb einer gegebenen Fehlerschranke die Wahrscheinlichkeitsmasse einer Menge von CTMC-Trajektorien approximieren kann. Diese erfüllen die lineare Eigenschaft spezifiziert von einem DTA mit mehreren Uhren. Soweit wir wissen, ist dies der erste Approximationsalgorithmus für Akzeptanzwahrscheinlichkeiten von DTAs auf CTMCs. Vorherige Resultate wie beispielsweise von Amparore und Donatelli (2010) und Barbot *et al.* (2011) behandeln nur Fälle wo der DTA nur eine Uhr hat.

In Teil 4 untersuchen wir maximale kostenbeschränkte Erreichbarkeitswahrscheinlichkeiten von CTMDPs. Wir zeigen die Existenz optimaler kostenpositionaler Strategien. Dabei wird Optimalität aller messbaren Strategien für CTMDPs betrachtet. Zusätzlich entwickeln wir einen numerischen Approximationsalgorithmus welcher (innerhalb einer gegebenen Fehlerschranke) die maximale Wahrscheinlichkeit approximiert, eine bestimmte Menge von Zielzuständen innerhalb eines mehrdimensionalen Kostenschrankenvektors zu erreichen. Die Zeitkomplexität des Algorithmus ist polynomiell in der Größe des CTMDP, der unären Darstellung der Kostenschrankenvektors und des Kehrwertes der Fehlerschranke. Sie ist exponentiell in der Dimension des Kostenschrankenvektors. Aufgrund der Zeitkomplexität ist der Approximationsalgorithmus effektiv für ein breites Spektrum von Anwendungen nutzbar, wo die Dimension des Kostenvektors klein ist. Unsere Resultate erweitern den zeitbeschränkten Fall untersucht von Neuhäüßer and Zhang (2010). Weiterhin zeigen wir zudem einen Beweisfehler in dieser Arbeit auf und korrigieren ihn mit neuen Beweisen.

Acknowledgements

I thank my supervisor, Prof. Joost-Pieter Katoen, for his doctoral guidance and for the freedom endowed by him. I thank Nils Jansen for the German translation of the abstract of this dissertation. I also thank my parents for their support on my doctoral study.

Contents

Abstract	v
Zusammenfassung	vii
Acknowledgements	ix
1 Introduction	1
1.1 Background: Formal Methods	1
1.2 Outline of the Dissertation	4
1.3 Origins of the Chapters and Credits	4
1.4 Basic Notations	5
2 Lattice Theory	7
3 Measure Theory	9
3.1 Measure Space	9
3.2 Lebesgue Integral	10
3.3 Product σ -Algebra	13
3.4 Dynkin's π - λ Theorem	13
4 Probabilistic Automata	15
4.1 Probabilistic Automata	16
4.2 Bisimulation and Simulation on PAs	17
5 Simulation Preorder between pPDAs and fPAs	19
5.1 Probabilistic Pushdown Automata	20
5.2 Extended Stack Symbols	22
5.3 Tableaux Proof System	25
5.4 EXPTIME-Hardness	33
5.5 Conclusion	35
6 Bisimilarity Metric on Probabilistic Automata	37
6.1 Bisimilarity Metric on PAs	38
6.2 Approximate Bisimilarity Metrics	44

6.3	Self-Closed Sets	45
6.4	The Membership of $UP \cap coUP$	51
6.5	Conclusion	52
7	Continuous-Time Markov Decision Processes	53
7.1	The Model	54
7.2	Paths and Histories	55
7.3	Measurable Spaces on Paths and Histories	57
7.4	Schedulers and Their Probability Spaces	58
7.5	A General Integral Characterization	61
7.6	Conclusion	62
7.7	Proofs	63
8	Acceptance Probability of CTMC-Paths by DTA	73
8.1	Continuous-Time Markov Chains	75
8.2	Deterministic Timed Automata	76
8.3	Measurability and The Integral Equations	78
8.4	Mathematical Technicalities	82
8.4.1	Equivalence Relations on Clock Valuations	82
8.4.2	Product Region Graph	84
8.4.3	Lipschitz Continuity	86
8.5	A Differential Characterization	91
8.6	Approximation Algorithm	93
8.6.1	Approximation Schemes	94
8.6.2	Error-Bound Analysis	99
8.7	Conclusion	107
9	Cost-Bounded Reachability on CTMDPs	109
9.1	Cost-Bounded Reachability Probability	110
9.2	Optimal Measurable Schedulers	120
9.3	Differential Characterizations	125
9.4	Approximation Algorithm	130
9.5	Conclusion	137
10	Conclusion	139

Chapter 1

Introduction

1.1 Background: Formal Methods

This dissertation falls in the research area of *formal methods*. In general, formal methods are mathematics-based techniques for modelling, verification and synthesis of systems. The term ‘modelling’ refers to describing systems through mathematical formalisms, the term ‘verification’ means to check automatically whether a system satisfies a desired property, and the task of synthesis is to generate automatically a system that satisfies a prescribed property. Traditional systems targeted by formal methods are computer systems such as software systems (e.g., programs, operating systems, protocols...) and hardware systems (e.g. CPU, routers, circuits...) (cf. the textbook [7]). Recently, complex systems such as cyber-physical systems (cf. eg., [54]) and biological systems (cf. eg., [10]) are also targeted by formal methods. As systems become more and more complex in recent years, it is more and more difficult to judge the functionality, reliability or performance of a system. Formal methods are then incorporated to model, verify or synthesize systems with complex behaviours (e.g., concurrency, non-deterministic and stochastic features, etc.).

Two important concepts in formal methods are *model* and *specification*. The concept of models describes in a mathematical sense how a system evolves when time progresses. In other words, models are rigorous descriptions for system evolutions. In general, a model of a system is composed of a set of states and a set of transitions which defines how one state can transit to another; then the behaviour of a system is described as a series of transitions along the time axis. The concept of specifications describes rigorously the desired behaviour of a system. For example, a specification can be a property that a system should obey or optimize. In the following, we describe different types of models and specifications.

Models

Due to different interpretations of time propagation, a model can either be *discrete-time* or *continuous-time*. In a discrete-time model (e.g., labelled transition systems [7, 56], discrete-time Markov chains [36, 68]), time progress is discretized into steps and transitions occur only at those steps; in a continuous-time model (e.g., hybrid systems [44], continuous-time Markov chains [36, 68]), transitions can occur at any time point on the dense time axis. In general, discrete-time models are suitable for discrete-phase systems (e.g., systems relying on a digital clock), whereas continuous-time models are adequate for systems that interact with the real world, where an event can happen at any time point.

Due to different interpretations of non-determinism, a model can also be either *probabilistic* or *non-probabilistic*. A model is probabilistic if it allows stochastic interpretation of non-determinism, while it is non-probabilistic if no stochastic interpretation is allowed. Generally, an instance of a probabilistic model need not to resolve all non-determinism as stochastic transitions, i.e., it can have both stochastic and non-deterministic features. A probabilistic model that resolves all non-determinism (as stochastic transitions) is typically called *fully probabilistic*.

Specifications

A specification can be either a property described by a rigorous linguistic sentence or an instance of a model. When a specification is a rigorous linguistic sentence, it usually specifies the logical property that the system should satisfy or optimize; such specification can be in most cases encoded by a temporal logical formula. Prominent temporal logics for specifications are CTL [26] (Computation Tree Logic), LTL [62] (Linear-Time Temporal Logic), CTL* [26] (the combination of CTL and LTL) and the most expressive logic of μ -calculus [50], together with their probabilistic or continuous-time extensions [45, 67, 11, 4]. Among them some are branching-time logics (e.g., CTL) which focuses on state-based properties, some are linear-time logics (e.g., LTL) which focuses on trajectory-based (or path-based) properties, and others are a combination of the two (e.g., CTL* and μ -calculus). The research area to check whether an instance of a model satisfies a formula of a temporal logic is known as *model checking* [7].

When a specification itself is an instance of a model, it usually describes exactly the desired behaviour of the system. In general, the original instance of the model (for the system) is compared with the specification to check whether two instances are equivalent under some semantical setting. Typical semantical equivalences are bisimulation equivalence [56, 61, 73], simulation preorder [56, 73] and their probabilistic and continuous-time extensions [49, 67, 53, 8]. The research area to check whether two instances of a model are

equivalent is known as *equivalence checking*.

The Main Results of the Dissertation

This dissertation mainly focuses on the *formal verification* of *probabilistic systems*. We consider both discrete-time and continuous-time probabilistic systems. For discrete-time probabilistic systems, we study the computational complexity of two equivalence-checking problems on probabilistic automata (PAs) [67] (which is an analogue of discrete-time Markov decision processes). In detail, we present the following two results.

- The decision problem of the (probabilistic) simulation preorder between a probabilistic pushdown automata (pPDA) and a finite probabilistic automaton (fPA) is in EXPTIME, and is EXPTIME-complete when the hardness result by Kučera and Mayr [52] is imported. The complexity result coincides with the one by Kučera and Mayr [52] on non-probabilistic pushdown automata. Moreover, we obtain a fixed-parameter-tractable result on this problem, which again coincides with the counterpart by Kučera and Mayr [52] on non-probabilistic pushdown automata.
- The bisimilarity metric defined by van Breugel and Worrell [72] on fPA is decidable in $\text{NP} \cap \text{coNP}$ (and even $\text{UP} \cap \text{coUP}$) for the undiscounted case. This result significantly improves the previous one by van Breugel *et al.* [71] (cf. also [22, 25]).

For continuous-time probabilistic systems, we study the model-checking problem on continuous-time Markov chains (CTMCs) and continuous-time Markov decision processes (CTMDPs). In detail, we present the following two results.

- We develop a numerical approximation algorithm for acceptance probability by DTA on CTMCs, while correcting errors in the previous work by Chen *et al.* [24] with new proofs. DTAs are a deterministic subclass of timed automata (by Alur and Dill [1]) which can encode a large class of linear real-time properties. In detail, we present an algorithm which within a given error bound, approximates the probability mass of the set of CTMC-trajectories that satisfy the linear property specified by a multi-clock DTA. The worst-case complexity of the approximation algorithm for CTMC-DTA is double exponential in the input size. As far as we know, this is the first approximation algorithm for acceptance probability by DTA on CTMCs. Previous results such as the ones by Amparore and Donatelli [31] and Barbot *et al.* [9] only consider cases where the DTA has only one clock.
- We study maximal cost-bounded reachability probability on CTMDPs. In detail, we prove the existence of optimal cost-positional schedulers,

where the optimality is considered under all measurable schedulers for CTMDPs. And we develop a numerical approximation algorithm that approximates (within a given error bound) the maximal probability to reach a certain set of target states within a multidimensional cost-bound vector. The time complexity of the algorithm is polynomial in the size of the CTMDP, the unary representation of the cost-bound vector and the reciprocal of the given error bound, and exponential in the dimension of the cost-bound vector. Due to its time complexity, the approximation algorithm is effective for a wide range of applications where the dimension of the cost-bound vector is low. Our results extend the time-bounded case studied by Neuhäüßer and Zhang [59]. Meanwhile, we also point out a proof error in the work on time-bounded case by Neuhäüßer and Zhang [59] and correct it with new proofs.

1.2 Outline of the Dissertation

Chapter 2 and Chapter 3 introduces mathematical preliminaries needed for this dissertation. Chapter 2 briefly introduce lattice theory and Knaster-Tarski's Fixed-Point Theorem. Chapter 3 briefly goes through measure theory and abstract Lebesgue integral.

Chapter 4 and Chapter 7 introduces the formal mathematical models concerned in this dissertation. Chapter 4 introduces PAs [67] and behavioural equivalences on PAs. Chapter 7 introduces the notion of CTMDPs and the notions of related measurable spaces and measure spaces, following the definitions in [74, 58].

Chapters 5, 6, 8 and 9 present the main contributions of this dissertation. Chapter 5 illustrates the complexity to decide (probabilistic) simulation preorder between pPDAs and fPAs. Chapter 6 demonstrates the membership of $UP \cap coUP$ for the bisimilarity metric [72] on PAs. Chapter 8 deals with the approximation algorithm for acceptance probability of CTMC-paths (CTMC-trajectories) by DTAs. Chapter 9 handles the approximation algorithm for maximal cost-bounded reachability probability on CTMDPs.

Finally, Chapter 10 concludes the dissertation.

1.3 Origins of the Chapters and Credits

Chapter 5 is based on the proceeding paper [42], whose detailed information is “Hongfei Fu, Joost-Pieter Katoen: Deciding Probabilistic Simulation between Probabilistic Pushdown Automata and Finite-State Systems. FSTTCS 2011: 445-456”.

Chapter 6 is based on the technical report [37], for which an extension is published as a proceeding paper [39] with detailed information “Hongfei Fu: Computing Game Metrics on Markov Decision Processes. ICALP (2) 2012: 227-238”.

Chapter 8 is based on the proceeding paper [40] whose detail is “Hongfei Fu: Approximating acceptance probabilities of CTMC-paths on multi-clock deterministic timed automata. HSCC 2013: 323-332”.

Chapter 9 is based on the proceeding paper [41] with detail “Hongfei Fu: Maximal Cost-Bounded Reachability Probability on Continuous-Time Markov Decision Processes. FoSSaCS 2014: 73-87”.

Besides, there are several papers published during my PhD study but are not included in this dissertation. They are listed as follows:

- Chaodong He, Yuxi Fu, Hongfei Fu: Decidability of Behavioral Equivalences in Process Calculi with Name Scoping. FSEN 2011: 284-298 [47];
- Hongfei Fu: Model Checking EGF on Basic Parallel Processes. ATVA 2011: 120-134 [38].

1.4 Basic Notations

In the whole dissertation, we use the following convention for notations. We denote by \mathbb{N} the set of natural numbers excluding zero, and by \mathbb{N}_0 the set of non-negative integers. We denote by \mathbb{R} the set of real numbers, by $\mathbb{R}_{\geq 0}$ the set of non-negative real numbers and by $\mathbb{R}_{> 0}$ the set of positive real numbers. Given $\mathbf{c} \in \mathbb{R}^k$ ($k \in \mathbb{N}$), we denote by \mathbf{c}_i ($1 \leq i \leq k$) the i -th coordinate of \mathbf{c} . We extend $\{\leq, \geq\}$ to real vectors and functions in a pointwise fashion: for two real vectors \mathbf{c}, \mathbf{d} , $\mathbf{c} \leq \mathbf{d}$ (resp. $\mathbf{c} \geq \mathbf{d}$) iff $\mathbf{c}_i \leq \mathbf{d}_i$ (resp. $\mathbf{c}_i \geq \mathbf{d}_i$) for all i ; for two real-valued functions g, h , $g \leq h$ (resp. $g \geq h$) iff $g(y) \leq h(y)$ (resp. $g(y) \geq h(y)$) for all y . Given a set Y , we let $\mathbf{1}_Y$ be the indicator function of Y , i.e. $\mathbf{1}_Y(y) = 1$ if $y \in Y$ and $\mathbf{1}_Y(y) = 0$ for $y \in X - Y$, where $X \supseteq Y$ is an implicitly known set.

Chapter 2

Lattice Theory

In this chapter, we briefly introduce lattice theory and the Knaster-Tarski's Fixed-Point Theorem. For a detailed introduction, we refer to the textbook [13] (cf. also [29]). The content of this chapter will be fundamental for Chapter 6 and Chapter 9.

Below we fix a non-empty set X and a partial order $\preceq \subseteq X \times X$ on X .

Definition 2.1. *Let $Y \subseteq X$ be a subset of X and $y \in X$ be an element of X . The following definitions are standard in lattice theory:*

- y is an upper bound of Y if $x \preceq y$ for all $x \in Y$, and dually y is a lower bound of Y if $y \preceq x$ for all $x \in Y$;
- y is the largest element of Y if y is an upper bound of Y and $y \in Y$, and dually y is the least element of Y if y is a lower bound of Y and $y \in Y$;
- y is the supremum of Y if y is the least element of the set of upper bounds of Y , i.e., the least element of the set

$$\{x \in X \mid z \preceq x \text{ for all } z \in Y\} ;$$

- y is the infimum of Y if y is the largest element of the set of lower bounds of Y , i.e., the largest element of the set

$$\{x \in X \mid x \preceq z \text{ for all } z \in Y\} .$$

We denote by $\bigsqcup Y$ (resp. $\bigsqcap Y$) the supremum (resp. infimum) of a subset $Y \subseteq X$. Generally, $\bigsqcup Y$ and $\bigsqcap Y$ may not always exist. The pair (X, \preceq) such that $\bigsqcup Y$ and $\bigsqcap Y$ always exists is called a *complete lattice*.

Definition 2.2. *The pair (X, \preceq) is a complete lattice iff for all $Y \subseteq X$, both $\bigsqcup Y$ and $\bigsqcap Y$ exists. If (X, \preceq) is a complete lattice, then the top element \top of X is defined as $\bigsqcup X$ and the bottom element \perp of X is defined as $\bigsqcap X$.*

Note that if (X, \preceq) is a complete lattice, then \top and \perp are resp. the largest and the least element of X .

In this dissertation, we are interested in fixed-points of monotone functions on a complete lattice, which is also a central part of lattice theory.

Definition 2.3. *Let $h : X \rightarrow X$ be a function and $x \in X$.*

- x is a fixed-point (resp. pre-fixed-point, post-fixed-point) of h if $h(x) = x$ (resp. $h(x) \preceq x$, $x \preceq h(x)$).
- h is monotone if $h(x) \preceq h(y)$ for all $x, y \in X$ such that $x \preceq y$.

The central theorem concerning fixed-points of monotone functions is Knaster-Tarski's Fixed-Point Theorem, which is illustrated as follows.

Theorem 2.1 (Knaster-Tarski's Fixed-Point Theorem). *Assume that (X, \preceq) is a complete lattice. Then every monotone function $h : X \rightarrow X$ has a fixed-point. The least fixed-point $\text{lfp}(h)$ of h (the least element of the set of fixed-points of h) exists and satisfies that*

$$\text{lfp}(h) = \bigcap \{x \in X \mid h(x) \preceq x\} .$$

The greatest fixed-point $\text{gfp}(h)$ of h (the largest element of the set of fixed-points of h) exists and satisfies that

$$\text{gfp}(h) = \bigcup \{x \in X \mid x \preceq h(x)\} .$$

Proof. We first prove the case for $\text{lfp}(h)$. Let $x^* := \bigcap \{x \in X \mid h(x) \preceq x\}$. By definition, $x^* \preceq x$ whenever $x \in X$ and $h(x) \preceq x$. Since h is monotone, $h(x^*) \preceq h(x) \preceq x$ whenever $x \in X$ and $h(x) \preceq x$. Thus, $h(x^*)$ is a lower bound of the set $\{x \in X \mid h(x) \preceq x\}$, which implies that $h(x^*) \preceq x^*$. Hence, $x^* \in \{x \in X \mid h(x) \preceq x\}$. By $h(x^*) \preceq x^*$ and the monotonicity of h , we have $h(h(x^*)) \preceq h(x^*)$, which implies $h(x^*) \in \{x \in X \mid h(x) \preceq x\}$. By definition, $x^* \preceq h(x^*)$. Thus, $x^* = h(x^*)$ and x^* is the least fixed-point of h .

Now we consider the case for $\text{gfp}(h)$. The proof is by dual to the one for $\text{lfp}(h)$. Let $x^* := \bigcup \{x \in X \mid x \preceq h(x)\}$. By definition, $x \preceq x^*$ whenever $x \in X$ and $x \preceq h(x)$. Since h is monotone, $x \preceq h(x) \preceq h(x^*)$ whenever $x \in X$ and $x \preceq h(x)$. Thus, $h(x^*)$ is an upper bound of the set $\{x \in X \mid x \preceq h(x)\}$, which implies that $x^* \preceq h(x^*)$. Hence, $x^* \in \{x \in X \mid x \preceq h(x)\}$. By $x^* \preceq h(x^*)$ and the monotonicity of h , we have $h(x^*) \preceq h(h(x^*))$, which implies $h(x^*) \in \{x \in X \mid x \preceq h(x)\}$. By definition, $h(x^*) \preceq x^*$. Thus, $x^* = h(x^*)$ and x^* is the greatest fixed-point of h . \square

Intuitively, Knaster-Tarski's Fixed-Point Theorem gives a infimum (resp. supremum) characterization for the least (resp. greatest) fixed-point of h .

Chapter 3

Measure Theory

In this chapter, we briefly introduce basic concepts of measure theory. For the details of measure theory, we refer to the textbooks [12, 32]. The content of this chapter will be fundamental in Chapter 7.

In this chapter, we extend the set \mathbb{R} of real numbers with two objects $-\infty$ (which indicates negative infinity) and $+\infty$ (which indicates positive infinity), with the following routine (cf. [32, Page 85]):

- for all $x \in \mathbb{R}$, $-\infty < x < +\infty$;
- for all $x \in \mathbb{R}$, $x + (-\infty) = -\infty$ and $x + (+\infty) = +\infty$;
- for all x which is either a positive real number or $+\infty$, $x \cdot (-\infty) = -\infty$ and $x \cdot (+\infty) = +\infty$;
- for all x which is either a negative real number or $-\infty$, $x \cdot (-\infty) = +\infty$ and $x \cdot (+\infty) = -\infty$;
- $0 \cdot (+\infty) = 0$ and $0 \cdot (-\infty) = 0$.

In some cases, we abbreviate $+\infty$ as ∞ .

3.1 Measure Space

The notion of measure space deals with a non-negative real measure on subsets of a certain set which mimics certain notion of “size” or ”length” on sets. A first notion encountered to introduce measure space is the notion of σ -algebra.

Definition 3.1. *Let Ω be a set. A set $\mathcal{S} \subseteq 2^\Omega$ is a σ -algebra on Ω iff the following conditions hold:*

1. $\emptyset \in \mathcal{S}$;
2. for all $X \in \mathcal{S}$, $\Omega \setminus X \in \mathcal{S}$;
3. for all infinite sequences $\{X_n\}_{n \in \mathbb{N}}$ such that $X_n \in \mathcal{S}$ for all $n \in \mathbb{N}$, $\bigcup_{n \in \mathbb{N}} X_n \in \mathcal{S}$.

If Ω is a set and \mathcal{S} is a σ -algebra on Ω , then (Ω, \mathcal{S}) is called a measurable space on Ω .

Given any sets Ω and $\mathcal{C} \subseteq 2^\Omega$, there is a (unique) smallest σ -algebra that contains \mathcal{C} . This smallest σ -algebra, denoted by $\sigma(\mathcal{C})$ (with Ω implicitly known), is given as follows:

$$\sigma(\mathcal{C}) := \bigcap \{ \mathcal{S} \mid \mathcal{S} \text{ is a } \sigma\text{-algebra on } \Omega \text{ and } \mathcal{C} \subseteq \mathcal{S} \} .$$

It is straightforward to verify by definition that $\sigma(\mathcal{C})$ is a σ -algebra on Ω , and $\sigma(\mathcal{C})$ is the smallest σ -algebra that contains \mathcal{C} . Often, the σ -algebra $\sigma(\mathcal{C})$ is said to be *generated* by \mathcal{C} .

To define the notion of measure on a measurable space, we import the auxiliary notion of *countably-additive function*.

Definition 3.2. Let Ω be a set and $\mathcal{C} \subseteq 2^\Omega$. A function $\mu : \mathcal{C} \rightarrow [0, +\infty]$ is countably-additive on \mathcal{C} if μ satisfies that for all sequences $\{X_n\}_{n \in \mathbb{N}}$ in¹ \mathcal{C} , if (i) $\bigcup_{n \in \mathbb{N}} X_n \in \mathcal{C}$ and (ii) $X_n \cap X_m = \emptyset$ whenever $n, m \in \mathbb{N}$ and $n \neq m$, then $\sum_{n \in \mathbb{N}} \mu(X_n) = \mu(\bigcup_{n \in \mathbb{N}} X_n)$.

Then the notion of measure is given in the following definition.

Definition 3.3. Let (Ω, \mathcal{S}) be a measurable space on Ω . A function $\mu : \mathcal{S} \rightarrow [0, +\infty]$ is called a measure for (Ω, \mathcal{S}) iff μ is countably additive on \mathcal{S} . If $\mu : \mathcal{S} \rightarrow [0, +\infty]$ is a measure for (Ω, \mathcal{S}) , then $(\Omega, \mathcal{S}, \mu)$ is called a measure space. If $(\Omega, \mathcal{S}, \mu)$ is a measure space and $\mu(\Omega) = 1$, then $(\Omega, \mathcal{S}, \mu)$ is also called a probability space and μ is called a probability measure.

If $(\Omega, \mathcal{S}, \mu)$ is a probability space, then normally elements $X \in \mathcal{S}$ are called *events* and $\mu(X)$ is referred as the probability that event X happens (i.e., a randomly chosen $x \in \Omega$ falls in X).

Below we illustrate the Borel σ -fields $\mathcal{B}(\mathbb{R}), \mathcal{B}(\mathbb{R}_{\geq 0})$ and the Borel measure on $\mathcal{B}(\mathbb{R})$.

Definition 3.4. The σ -algebra $\mathcal{B}(\mathbb{R})$ (resp. $\mathcal{B}(\mathbb{R}_{\geq 0})$) is the σ -algebra generated by the set of all open intervals $(a, b) \subseteq \mathbb{R}$ (resp. $(a, b) \subseteq \mathbb{R}_{\geq 0}$). The Borel measure μ_{brl} for the measurable space $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ is the unique measure such that $\mu_{\text{brl}}((a, b)) = b - a$ for all non-empty open intervals $(a, b) \subseteq \mathbb{R}$.

3.2 Lebesgue Integral

In this section, we briefly introduce (abstract) Lebesgue integral. Below we fix a measure space $(\Omega, \mathcal{S}, \mu)$. The following definition illustrates the notion of measurable functions.

¹“in” here means $X_n \in \mathcal{C}$ for all $n \in \mathbb{N}$.

Definition 3.5. Let (Ω', \mathcal{S}') be a measurable space. A function $h : \Omega \rightarrow \Omega'$ is measurable (w.r.t (Ω', \mathcal{S}')) if for all $X \in \mathcal{S}'$, $h^{-1}(X) \in \mathcal{S}$.

The following definition illustrates the notion of simple functions, which is a central notion to define abstract Lebesgue integral.

Definition 3.6. A function $h : \Omega \rightarrow [0, +\infty)$ is a (non-negative) simple function on Ω if there exists $n \in \mathbb{N}$, a finite sequence $\{X_i\}_{1 \leq i \leq n}$ of sets in \mathcal{S} and a finite sequence $\{d_i\}_{1 \leq i \leq n}$ of real numbers such that $h = \sum_{i=1}^n d_i \cdot \mathbf{1}_{X_i}$.

Intuitively, simple functions are piecewise constant functions. By definition, it is not hard to verify that simple functions are measurable functions w.r.t $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. Here, we restrict simple functions to non-negative functions; this restriction is not essential and does not affect the definition of Lebesgue integral.

The following definition introduces Lebesgue integral on simple functions.

Definition 3.7. Let $h : \Omega \rightarrow [0, +\infty)$ be a (non-negative) simple function on Ω such that $h = \sum_{i=1}^n d_i \cdot \mathbf{1}_{X_i}$ (cf. Definition 3.6 for notations). The Lebesgue integral of h w.r.t μ , denoted by $\int h \, d\mu$, is defined by:

$$\int h \, d\mu := \sum_{i=1}^n d_i \cdot \mu(X_i) .$$

The Lebesgue integral is well-defined on simple functions, regardless of the representation of the simple function; see [32, Proposition 4.1.4] for more details. The following definition introduces Lebesgue integral on general non-negative functions.

Definition 3.8. Let $h : \Omega \rightarrow [0, +\infty]$ be a function. The (Lebesgue) integral of h w.r.t μ , denoted by $\int h \, d\mu$, is defined as follows:

$$\int h \, d\mu := \sup \left\{ \int g \, d\mu \mid g \text{ is a simple function on } \Omega \text{ and } g \leq h \right\} ,$$

where for arbitrary functions $g_1, g_2 : \Omega \rightarrow [0, +\infty]$, $g_1 \leq g_2$ means that $g_1(x) \leq g_2(x)$ for all $x \in \Omega$.

Note that by definition, for non-negative function $g_1, g_2 : \Omega \rightarrow [0, +\infty]$, $\int g_1 \, d\mu \leq \int g_2 \, d\mu$ whenever $g_1 \leq g_2$. In this dissertation, we may also write “ $\int h \, d\mu$ ” as “ $\int h(x) \mu(dx)$ ”, stressing the role of the variable x ; we will abbreviate “ $\mu_{\text{Borel}}(dx)$ ” in an integral as “ dx ”.

Remark 3.1. Note that a standard mathematical definition of Lebesgue integral applies only to measurable functions, while Definition 3.8 works also for non-measurable functions. This fact will be used in Chapter 9 to define a monotone operator on a lattice of general functions.

²cf. Definition 3.2 for the meaning of “in”.

For non-negative measurable functions, the next proposition is useful. In the following, we write $x_n \uparrow x$ if the sequence $\{x_n\}_{n \in \mathbb{N}}$ of real numbers converges to x when $n \rightarrow +\infty$ and $x_n \leq x_{n+1}$ for all $n \in \mathbb{N}$ (note that x may be $+\infty$).

Proposition 3.1. *Let $h : \Omega \rightarrow [0, +\infty]$ be a function. If h is measurable w.r.t $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, then there exists a sequence $\{h_n : \Omega \rightarrow [0, +\infty)\}$ of simple functions on Ω such that $h_n(x) \uparrow h(x)$ for all $x \in \Omega$; moreover, for all such sequences $\{h_n : \Omega \rightarrow [0, +\infty)\}$, $\int h_n \, d\mu \uparrow \int h \, d\mu$.*

Proof. See [32, Proposition 4.1.5]. □

The following definition illustrates Lebesgue integral on all functions.

Definition 3.9. *Let $h : \Omega \rightarrow [-\infty, +\infty]$ be a function and define two functions h^+ and h^- by: $h^+ := \max\{h, 0\}$ and $h^- := -\min\{h, 0\}$. The integral $\int h \, d\mu$ is called defined if at most one of the two integrals, namely $\int h^+ \, d\mu$ and $\int h^- \, d\mu$, is equal to $+\infty$. When $\int h \, d\mu$ is defined, it is given by:*

$$\int h \, d\mu := \int h^+ \, d\mu - \int h^- \, d\mu .$$

Moreover, h is called integratable if h is measurable w.r.t $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ and $\int |h| \, d\mu < +\infty$; the set of all integratable functions is denoted by $\mathcal{L}^1(\Omega, \mathcal{S}, \mu)$.

For a function $h : \Omega \rightarrow [-\infty, +\infty]$ and a set $X \subseteq \Omega$, we denote the integral $\int_X h \, d\mu$ to be $\int h \cdot \mathbf{1}_X \, d\mu$ if $\int h \cdot \mathbf{1}_X \, d\mu$ is defined.

Below we state some mathematical facts in measure theory. The following proposition states some basic properties of measurable functions.

Proposition 3.2. *If $c \in \mathbb{R}$ and $h_1, h_2 : \Omega \rightarrow \mathbb{R}$ are two functions measurable w.r.t $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, then both $h_1 + h_2$ and $c \cdot h_1$ is a measurable function w.r.t $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$.*

If $\{h_n : \Omega \rightarrow \mathbb{R}\}_{n \in \mathbb{N}}$ is a sequence of measurable functions w.r.t $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, $h : \Omega \rightarrow \mathbb{R}$ is a function and $\lim_{n \rightarrow \infty} h_n(x) = h(x)$ for all $x \in \Omega$, then h is a measurable function w.r.t $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$.

Proof. See [32, Chapter 4]. □

The following theorem states that Lebesgue integral has good linear properties.

Theorem 3.1. *For all $g, h \in \mathcal{L}^1(\Omega, \mathcal{S}, \mu)$ and $c \in \mathbb{R}$,*

$$\int (g + h) \, d\mu = \left(\int g \, d\mu \right) + \left(\int h \, d\mu \right) \text{ and } \int (c \cdot g) \, d\mu = c \cdot \left(\int g \, d\mu \right) .$$

Proof. See [32, Theorem 4.1.10]. □

Below we introduce Monotone Convergence Theorem, which is an important convergence theorem in measure theory. We would like to mention that there is another important convergence theorem, called Dominated Convergence Theorem, which is not included in this dissertation.

Theorem 3.2 (Monotone Convergence Theorem). *Let $\{h_n\}_{n \in \mathbb{N}}$ be a sequence of measurable functions from Ω to $[-\infty, +\infty]$, and h be a function from Ω to $[-\infty, +\infty]$. If $h_n(x) \uparrow h(x)$ for all $x \in \Omega$ and $\int h_1 d\mu > -\infty$, then $\int h_n d\mu \uparrow \int h d\mu$.*

Proof. See [32, Theorem 4.3.2]. □

3.3 Product σ -Algebra

In this section, we introduce the notion of product measure. Intuitively, a product measure is a measure for certain ‘‘Cartesian product’’ of two measure spaces, while the calculation of the product measure mimics the calculation of the area of rectangles on a two-dimensional plane.

Below we fix two measure spaces $(\Omega_1, \mathcal{S}_1, \mu_1)$ and $(\Omega_2, \mathcal{S}_2, \mu_2)$, where both μ_1 and μ_2 is σ -finite. The definition of the notion of σ -finiteness is given as follows.

Definition 3.10. *Let Ω be a set and $\mathcal{C} \subseteq 2^\Omega$. A function $\mu : \mathcal{C} \rightarrow [0, \infty)$ is said to be σ -finite if there exists $\{X_n\}_{n \in \mathbb{N}}$ such that (i) $X_n \in \mathcal{C}$ for all $n \in \mathbb{N}$ and (ii) $\bigcup_{n \in \mathbb{N}} X_n = \Omega$ and (iii) $\mu(X_n) < +\infty$ for all $n \in \mathbb{N}$.*

Now let $\mathcal{C} := \{X \times Y \mid X \in \mathcal{S}_1, Y \in \mathcal{S}_2\}$. The following definition illustrates the notion of product σ -algebra.

Definition 3.11. *The product σ -algebra $\mathcal{S}_1 \otimes \mathcal{S}_2$ on $\Omega_1 \times \Omega_2$ is defined to be the σ -algebra generated by \mathcal{C} .*

3.4 Dynkin’s π - λ Theorem

In this section, we briefly introduce Dynkin’s π - λ Theorem. For a detailed introduction, we refer to [12]. Below we fix a set Ω . Informally, Dynkin’s π - λ Theorem deals with the relationship between π -systems and λ -systems.

The following definition illustrates the notion of π -system.

Definition 3.12 (π -System). *A set $\mathcal{C} \subseteq 2^\Omega$ is a π -system on Ω iff it is closed under finite intersection, i.e., $X \cap Y \in \mathcal{C}$ whenever $X \in \mathcal{C}$ and $Y \in \mathcal{C}$.*

The notion of λ -system is a weaker notion of σ -algebra, as follows.

Definition 3.13 (λ -System). *A set $\mathcal{C} \subseteq 2^\Omega$ is a λ -system on Ω iff the following conditions hold:*

- $\Omega \in \mathcal{C}$;
- for all $X \in \mathcal{C}$, $\Omega \setminus X \in \mathcal{C}$;
- for all sequences $\{X_n\}_{n \in \mathbb{N}}$ in³ \mathcal{C} such that $X_n \cap X_m = \emptyset$ whenever $n \neq m$, $\bigcup_{n \in \mathbb{N}} X_n \in \mathcal{C}$.

In this thesis, we will use an equivalent version of Definition 3.13⁴, as follows.

Definition 3.14. A set $\mathcal{C} \subseteq 2^\Omega$ is a λ -system on Ω iff the following conditions hold:

- $\Omega \in \mathcal{C}$;
- for all $X, Y \in \mathcal{C}$ such that $X \subseteq Y$, $Y \setminus X \in \mathcal{C}$;
- for all sequences $\{X_n\}_{n \in \mathbb{N}}$ in \mathcal{C} such that $X_n \subseteq X_m$ whenever $n \leq m$, $\bigcup_{n \in \mathbb{N}} X_n \in \mathcal{C}$.

The following proposition shows that the two definitions are equivalent.

Proposition 3.3. Definition 3.13 and Definition 3.14 are equivalent.

Proof. Assume that \mathcal{C} is a λ -system w.r.t Definition 3.13. We prove that \mathcal{C} satisfies the conditions in Definition 3.14. The analysis is as follows: given any $X, Y \in \mathcal{C}$ with $X \subseteq Y$, $Y \setminus X \in \mathcal{C}$ because $Y \setminus X = \Omega \setminus (X \cup (\Omega \setminus Y))$; given any sequences $\{X_n\}_{n \in \mathbb{N}}$ in \mathcal{C} such that $X_n \subseteq X_m$ whenever $n \leq m$, $\bigcup_{n \in \mathbb{N}} X_n \in \mathcal{C}$ because $\bigcup_{n \in \mathbb{N}} X_n = X_1 \cup (\bigcup_{n \in \mathbb{N}} X_{n+1} \setminus X_n)$.

Assume now that \mathcal{C} is a λ -system w.r.t Definition 3.14. We first prove that for all $X, Y \in \mathcal{C}$, if $X \cap Y = \emptyset$ then $X \cup Y \in \mathcal{C}$. This follows directly from the fact that $X \cup Y = \Omega \setminus ((\Omega \setminus X) \setminus Y)$. Then we prove that \mathcal{C} satisfies the conditions in Definition 3.13: for all $X \in \mathcal{C}$, $\Omega \setminus X \in \mathcal{C}$ since $\Omega \in \mathcal{C}$ and $X \subseteq \Omega$; for all sequences $\{X_n\}_{n \in \mathbb{N}}$ in \mathcal{C} such that $X_n \cap X_m = \emptyset$ whenever $n \neq m$, $\bigcup_{n \in \mathbb{N}} X_n \in \mathcal{C}$ since $\bigcup_{n=1}^k X_n \in \mathcal{C}$ for all $k \in \mathbb{N}$. \square

The following theorem, entitled Dynkin's π - λ Theorem, is an important theorem in measure theory.

Theorem 3.3 (Dynkin's π - λ Theorem). Let \mathcal{E} be a π -system on Ω and \mathcal{F} be a λ -system on Ω . If $\mathcal{E} \subseteq \mathcal{F}$, then $\sigma(\mathcal{E}) \subseteq \mathcal{F}$.

Proof. See [12, Theorem 3.2]. \square

³cf. Definition 3.2 for the meaning of "in".

⁴This is an excerpt from the website http://en.wikipedia.org/wiki/Dynkin_system.

Chapter 4

Probabilistic Automata

In this chapter, we consider the notion of probabilistic automata developed by Segala and Lynch [67]. Probabilistic automata (PAs) are an analogue of Markov decision processes [64] which, like Markov decision processes, involves both stochastic and non-deterministic features. The difference between PAs and Markov decision processes is that in a PA, a sole action can lead to different probability distributions, while an action uniquely determines a probability distribution in a Markov decision process. In other words, PAs focus more on reactive features than Markov decision processes do. Probabilistic automata can also be viewed as an orthogonal extension of labelled transition systems [7] with probabilities.

In [67], the semantics of probabilistic automata is defined through two approaches: logical characterization and behavioural equivalences. Logical characterization allows one to reason about probabilistic automata via logical formulae, while behavioural equivalences enables one to check whether two probabilistic automata are equivalent. Behavioural equivalences and logical characterization are closely related as two behaviourally-equivalent probabilistic automata are also logically-equivalent under certain logical characterization [67].

The dissertation mainly focuses on behavioural equivalences of probabilistic automata. A behavioural equivalence can be viewed as an equivalence relation or a preorder on the set of states of the underlying probabilistic automaton, which characterize whether two states are behaviourally indistinguishable or one state can mimic the behaviour of another. Typical behavioural equivalences introduced in [67] are (probabilistic) bisimulation equivalence and (probabilistic) simulation preorder (cf. also [49, 53]). In [67], Segala and Lynch proved that (probabilistic) bisimulation equivalence preserves logical properties encoded by probabilistic computation tree logic (PCTL), and the equivalence relation induced by (probabilistic) simulation preorder preserves a safety fragment of PCTL (see also [7]).

The chapter is organized as follows. In Section 4.1, we briefly introduce

the notion of probabilistic automata. In Section 4.2, we briefly introduce two fundamental notions of behavioural equivalences on probabilistic automata, namely (probabilistic) bisimulation equivalence and (probabilistic) simulation preorder.

4.1 Probabilistic Automata

To introduce the notion of probabilistic automata, we first introduce the notion of (discrete) probability distributions.

Definition 4.1. *Let X be a finite or countable set. A (discrete) probability distribution on X is a function $\mu : X \rightarrow [0, 1]$ such that $\sum_{x \in X} \mu(x) = 1$; μ is Dirac at $x \in X$ iff $\mu(x) = 1$. The Dirac distribution $\mathcal{D}[x] : X \rightarrow [0, 1]$ for an $x \in X$ is defined by: $\mathcal{D}x = 1$ and $\mathcal{D}[x](y) = 0$ for all $y \neq x$.*

A probability distribution μ on X is finite if the support of μ , denoted by $\lfloor \mu \rfloor$ and defined by $\lfloor \mu \rfloor := \{x \in X \mid \mu(x) > 0\}$, is finite. The set of (discrete) probability distributions (resp. finite probability distributions) on X is denoted by $\text{Dist}(X)$ (resp. by $\text{Dist}_f(X)$).

Remark 4.1. *Note that each probability distribution $\mu \in \text{Dist}(X)$ corresponds to the probability space (cf. Chapter 3) $(X, 2^X, \bar{\mu})$ where $\bar{\mu}(Y) = \sum_{x \in Y} \mu(x)$ for all $Y \subseteq X$.*

The notion of probabilistic automata [67] is given as follows.

Definition 4.2. [67] *A probabilistic automaton (PA) is a tuple $(S, \text{Act}, \rightarrow)$ where*

- S is a non-empty, finite or countable set of states;
- Act is a non-empty set of actions;
- $\rightarrow \subseteq S \times \text{Act} \times \text{Dist}(S)$ is a set of transitions.

A PA $(S, \text{Act}, \rightarrow)$ is locally finite if for all $s \in S$ and $a \in \text{Act}$, it holds that $\{\mu \mid (s, a, \mu) \in \rightarrow\} \subseteq \text{Dist}_f(S)$; it is finite if it is locally finite, and both S, Act and \rightarrow is finite.

Technically speaking, the class of probabilistic automata here refers to the class of *simple* probabilistic automata in [67].

Let $(S, \text{Act}, \rightarrow)$ be a PA. For each $s \in S$, we define

- $\text{Act}(s) := \{a \in \text{Act} \mid \exists \mu. (s, a, \mu) \in \rightarrow\}$ to be the set of actions *enabled* at s , and
- $\text{Succ}(s) := \{s' \in S \mid \exists (s, a, \mu) \in \rightarrow. \mu(s') > 0\}$ to be the set of states which can be reached from s within one step with non-zero probability.

Intuitively, the set \rightarrow specifies all possible stochastic transitions from states to states. A transition (s, a, μ) specifies that when the current state is s and the action to be taken is a , the state at the next step is chosen w.r.t the probability distribution μ .

The following definition extends the notions of transitions to *combined transitions*. In [67], *combined transitions* are introduced to model stochastic strategies by, e.g., an adversary.

Definition 4.3. *Let (S, A, \rightarrow) be a PA. The set of combined transitions, $\rightarrow_{\subseteq} \subseteq S \times Act \times Dist(S)$, is defined as follows: $(s, a, \mu) \in \rightarrow_{\subseteq}$ iff there exists a finite or infinite sequence $\{(\mu_n, d_n)\}_{n \in I}$ ($I \subseteq \mathbb{N}$), where $\mu_n \in Dist(S)$ and $d_n \in \mathbb{R}_{\geq 0}$ for all $n \in I$, such that*

- $(s, a, \mu_n) \in \rightarrow$ for all $n \in I$, and
- $\sum_{n \in I} d_n = 1$, and
- $\mu(s) = \sum_{n \in I} d_n \cdot \mu_n(s)$ for all $s \in S$.

By definition, $\rightarrow_{\subseteq} \subseteq \rightarrow$. We write “ $s \xrightarrow[\text{nc}]{a} \mu$ ” for “ $(s, a, \mu) \in \rightarrow$ ” where “nc” stands for “standard” or “non-combined”. And we write “ $s \xrightarrow[\text{c}]{a} \mu$ ” for “ $(s, a, \mu) \in \rightarrow_{\subseteq}$ ”. We will use “ $s \xrightarrow[\text{op}]{a} \mu$ ” to refer to either “ $s \xrightarrow[\text{nc}]{a} \mu$ ” or “ $s \xrightarrow[\text{c}]{a} \mu$ ”, depending on whether $\text{op} = \text{‘nc’}$ or $\text{op} = \text{‘c’}$.

4.2 Bisimulation and Simulation on PAs

In this section, we fix a PA (S, Act, \rightarrow) . To define (probabilistic) bisimulation and (probabilistic) simulation preorder, we first introduce a notion of lifting operation which lifts a binary relation on states to a binary relation on probability distributions.

Definition 4.4. [49, 53] *Let $\mathcal{R} \subseteq S \times S$ be a binary relation on S . The lifting relation $\overline{\mathcal{R}} \subseteq Dist(S) \times Dist(S)$ of \mathcal{R} is defined as follows: $(\mu, \nu) \in \overline{\mathcal{R}}$ iff there exists a weight function $w : S \times S \rightarrow [0, 1]$ such that*

- $\sum_{v \in S} w(u, v) = \mu(u)$ for all $u \in S$, and
- $\sum_{u \in S} w(u, v) = \nu(v)$ for all $v \in S$, and
- for all $u, v \in S$, $(u, v) \in \mathcal{R}$ whenever $w(u, v) > 0$.

For the sake of convenience, we will simply write “ $\mu \mathcal{R} \nu$ ” instead of “ $\mu \overline{\mathcal{R}} \nu$ ”.

Now we define the notion of (probabilistic) bisimulation, following the definition in [67].

Definition 4.5. *An equivalence relation \mathcal{R} on S is an op-bisimulation iff for all $(s, s') \in \mathcal{R}$, the following conditions hold:*

- for all $s \xrightarrow[\text{nc}]{a} \mu$, there exists $s' \xrightarrow[\text{op}]{a} \mu'$ such that $\mu \mathcal{R} \mu'$;
- for all $s' \xrightarrow[\text{nc}]{a} \mu'$, there exists $s \xrightarrow[\text{op}]{a} \mu$ such that $\mu \mathcal{R} \mu'$.

Two states $s, s' \in S$ are op-bisimilar if there exists an op-bisimulation \mathcal{R} such that $(s, s') \in \mathcal{R}$. The op-bisimilarity, denoted by \sim_{op} , is defined as the set of all pairs $(s, s') \in S \times S$ such that s and s' are op-bisimilar.

Here, \sim_{nc} refers to *strong bisimulation* [67] and \sim_c refers to *strong probabilistic bisimulation* [67]. The following definition illustrates the notion of (probabilistic) simulation preorder, which is defined as a one-sided version of (probabilistic) bisimulation.

Definition 4.6. A binary relation \mathcal{R} on S is an op-simulation iff for all $(s, s') \in \mathcal{R}$, the following conditions hold:

- $\text{Act}(s) = \text{Act}(s')$;
- for all $s \xrightarrow[\text{nc}]{a} \mu$, there exists $s' \xrightarrow[\text{op}]{a} \mu'$ such that $\mu \mathcal{R} \mu'$.

The op-simulation preorder, denoted by \sqsubseteq_{op} , is defined by:

$$\sqsubseteq_{\text{op}} := \bigcup \{ \mathcal{R} \subseteq S \times S \mid \mathcal{R} \text{ is an op-simulation} \} .$$

The op-simulation equivalence is defined as $\sqsubseteq_{\text{op}} \cap \sqsubseteq_{\text{op}}^{-1}$, where $\sqsubseteq_{\text{op}}^{-1}$ is the reverse relation of \sqsubseteq_{op} .

\sqsubseteq_{nc} corresponds to strong simulation and \sqsubseteq_c corresponds to strong probabilistic simulation in [67]. It can be easily verified that \sqsubseteq_{nc} (resp. \sqsubseteq_c) itself is an nc-simulation (resp. c-simulation).

Up till now, the notions of bisimulation and simulation preorder are defined on a single probabilistic automaton. They are extended to the counterparts between two probabilistic automata by taking the disjoint union of the two PAs, i.e., by making a PA which precisely comprises all the states, actions and transitions of the original two PAs.

Chapter 5

Simulation Preorder between pPDAs and fPAs

In this chapter, we consider the decision problem of (probabilistic) simulation preorder between a probabilistic pushdown automata (pPDA) and a finite probabilistic automata (fPA). Here, the notions of (finite) probabilistic automata (PAs) and (probabilistic) simulation preorder \sqsubseteq_{op} are defined in Chapter 4. We recall that (i) the notion of probabilistic automata is an analogue of Markov decision processes which models discrete-time stochastic non-deterministic transitions between states, and (ii) the notion of (probabilistic) simulation preorder captures the phenomenon whether one state can mimic another.

The motivation of this problem is to treat pPDAs as implementation and fPAs as specification. On one hand, we treat pPDAs as implementation since they are a natural extension of pushdown automata and can model probabilistic procedural programs. They typically induces an infinite-state probabilistic automaton and are equally expressive as recursive Markov chains [34] or recursive Markov decision processes [33]. On the other hand, we treat fPAs as specification since in many applications, a specification can be described by a finite-state diagram. We use (probabilistic) simulation preorder (or (probabilistic) simulation equivalence) as the comparison semantics between a pPDA and a fPA. It is shown by Segala and Lynch [67] that (probabilistic) simulation equivalence preserves a safety fragment of PCTL* formulae on probabilistic automata.

Verification of (probabilistic) pushdown automata has been studied mainly on the aspect of model checking. The major results in the model checking of pPDAs are on the analysis of *termination probabilities* [51, 34]. As to equivalence checking, Tomáš Brázdil *et al.* proved that probabilistic bisimilarity between pPDAs and fPAs can be decided in EXPTIME [19]. Besides, we would like to mention several novel results on original non-probabilistic pushdown automata [52, 69, 70, 48, 43].

In this chapter, we prove that the decision problem of probabilistic simulation preorder between a pPDA and a fPA is decidable in EXPTIME, and in PTIME if both the number of control states of the pPDA and the size of the finite-state system are fixed. All these results extend the counterparts in the non-probabilistic setting by Kučera and Mayr [52]. We use a different technique from the one adopted in [52]. The technique we use is a tableaux system for probabilistic simulation preorder between pPDA and fPA, which is based on the notion of “extended stack symbols” and a tableaux system both developed by Coling Stirling [69, 70] to prove the decidability of bisimulation equivalence on non-probabilistic pushdown automata. The reason to adopt a new technique is that the original result [52] goes through a reduction to the model-checking problem of μ -calculus on pushdown processes, for which a probabilistic extension is dubious [18]. By applying a EXPTIME-hardness result by Kučera and Mayr [52], we prove that the decision problem is EXPTIME-complete.

The result of this chapter, together with previous results concerning bisimulation equivalence and simulation preorder between pushdown automata and finite-state systems, can be illustrated by Table 5.1.

	non-probabilistic	probabilistic
bisim. equiv.	PSPACE-c. [48, 52]	in EXPTIME [19]
sim. pre.	EXPTIME-c. [52]	EXPTIME-c. (this chapter)

Table 5.1: Related Complexity Results

The chapter is organized as follows. In Section 5.1, we introduce the notion of probabilistic pushdown automata, following the definitions from [51, 19]. In Section 5.2, we introduce the notion of extended stack symbols. In Section 5.3, we present the tableaux system and show its soundness and completeness. In Section 5.4, we briefly describe how one can apply the hardness result by Kučera and Mayr [52] to our case. Finally, Section 5.5 concludes the chapter.

In the whole chapter, we denote by ϵ the empty word.

5.1 Probabilistic Pushdown Automata

Definition 5.1. [51, 19] A probabilistic pushdown automaton (pPDA) is a quadruple (Q, Γ, L, Δ) , where:

- Q is a finite non-empty set of control states;
- Γ is a finite non-empty set of stack symbols;
- L is a finite non-empty set of labels;
- $\mapsto_{\subseteq} (Q \times \Gamma) \times L \times \text{Dist}_f(Q \times \Gamma^*)$ is a finite set of transition rules.

In the whole chapter, we will use p, q to range over control states Q , A, B, C to range over stack symbols Γ , α, β, γ to range over Γ^* , μ, ν to range over $Q \times \Gamma^*$ and a, b, c to range over labels L . Instead of “ $(pA, a, \mu) \in \rightarrow$ ”, we write “ $pA \xrightarrow{a} \mu$ ”. The following definition illustrates how a pPDA generates a probabilistic automaton.

Definition 5.2. *Let $P = (Q, \Gamma, L, \rightarrow)$ be a pPDA. P induces a PA (S, Act, Ω) as follows:*

$$S := Q \times \Gamma^*; Act := L; \rightarrow := \{(pA\gamma, a, \mu_\gamma) \mid pA \xrightarrow{a} \mu, \gamma \in \Gamma^*\} .$$

The probability distribution $\mu_\gamma \in Dist(Q \times \Gamma^*)$ with $\mu \in Dist(Q \times \Gamma^*)$ and $\gamma \in \Gamma^*$ is defined by:

$$\mu_\gamma(p\alpha) = \begin{cases} \mu(p\beta) & \text{if } \alpha = \beta\gamma \text{ for some (unique) } \beta \in \Gamma^* \\ 0 & \text{otherwise} \end{cases} ,$$

for all $p\alpha \in Q \times \Gamma^*$. Elements of $Q \times \Gamma^*$ are also called configurations.

Note that every pPDA induces a locally-finite PA. Intuitively, a pPDA generates a PA by just expanding suffixes to transition rules. Below we state a simple property for such expansion of transition rules.

Lemma 5.1. *Let $(Q, \Gamma, L, \rightarrow)$ be a pPDA. Let $p\beta\gamma \in Q \times \Gamma^*$ with $\beta \in \Gamma^+$, $a \in L$ and $\mu \in Dist(Q \times \Gamma^*)$. Then $p\beta\gamma \xrightarrow[\text{op}]{a} \mu$ iff there exists $\mu' \in Dist(Q \times \Gamma^*)$ such that $p\beta \xrightarrow[\text{op}]{a} \mu'$ and $\mu = \mu'_\gamma$.*

Proof. We first consider the case when $\text{op} = \text{'nc'}$. Let $\beta = A\beta'$. By definition,

$$pA\beta'\gamma \xrightarrow[\text{nc}]{a} \mu \text{ iff } \exists \mu'. (pA \xrightarrow{a} \mu' \wedge \mu = \mu'_{\beta'\gamma}) .$$

Note that $(\nu_\beta)_\gamma = \nu_{\beta\gamma}$ for all $\nu \in Dist(Q \times \Gamma^*)$. Thus, one obtains

$$\exists \mu'. (pA \xrightarrow{a} \mu' \wedge \mu = \mu'_{\beta'\gamma}) \text{ iff } \exists \mu''. (pA\beta' \xrightarrow[\text{nc}]{a} \mu'' \wedge \mu = \mu''_\gamma)$$

(by taking $\mu'' = \mu'_{\beta'}$).

Now we consider the case when $\text{op} = \text{'c'}$. By definition, $p\beta\gamma \xrightarrow[\text{c}]{a} \mu$ iff (a) there exists a finite or infinite sequence $\{(\mu_n, d_n)\}_n$ such that $p\beta\gamma \xrightarrow[\text{nc}]{a} \mu_n$ for all n , $\sum_n d_n = 1$, and $\mu(s) = \sum_n d_n \cdot \mu_n(s)$ for all $s \in Q \times \Gamma^*$. By the proof for the case $\text{op} = \text{'nc'}$, we have (a) holds iff (b) there exists a finite or infinite sequence $\{(\mu'_n, d_n)\}_n$ such that $p\beta \xrightarrow[\text{nc}]{a} \mu'_n$ for all n , $\sum_n d_n = 1$, and $\mu(s) = \sum_n d_n \cdot (\mu'_n)_\gamma(s)$ for all $s \in Q \times \Gamma^*$. We can further obtain that (b) holds iff (c) there exists a finite or infinite sequence $\{(\mu'_n, d_n)\}_n$ such that $p\beta \xrightarrow[\text{nc}]{a} \mu'_n$ for all n , $\sum_n d_n = 1$, and $\mu = \mu'_\gamma$, where $\mu'(s) := \sum_n d_n \cdot \mu'_n(s)$ for all $s \in Q \times \Gamma^*$. Then (c) holds iff $\exists \mu'. (p\beta \xrightarrow[\text{c}]{a} \mu' \wedge \mu = \mu'_\gamma)$, from which the result follows. \square

This chapter studies the complexity of the following decision problems:

- **INPUT:** a configuration $p\alpha$ of a pPDA and a state s of a fPA;
- **OUTPUT:** whether $p\alpha \sqsubseteq_{\text{op}} s$ and whether $s \sqsubseteq_{\text{op}} p\alpha$ between the fPA and the PA induced by the pPDA, for $\text{op} \in \{‘nc’, ‘c’\}$.

We prove that both of these problems are EXPTIME-complete, and are in PTIME if both the number of control states of the pPDA and the number of states of the fPA is fixed.

5.2 Extended Stack Symbols

In this section, we extend the notion of “extended stack symbols” by Colin Stirling [70, 69], which is originally used to establish a tableaux proof system that is sound and complete for the bisimulation equivalence on pushdown automata; we follow Colin Stirling’s method to establish extended stack symbols for simulation preorder between (PA induced by) pPDA and fPA. Later in Section 5.3, we present a tableaux proof system for simulation preorder between pPDA and fPA and demonstrate our complexity results.

Below we fix a pPDA $(Q, \Gamma, L, \rightsquigarrow)$ and a fPA (S, A, \rightarrow) . For any two sets X, Y , we define $X \odot Y := (X \times Y) \cup (Y \times X)$. The following definition illustrates the notion of extended stack symbols.

Definition 5.3. *An extended stack symbol U is a function $U : Q \rightarrow 2^S$. The set E is defined as the set of all extended stack symbols.*

W.l.o.g, we assume that $E \cap \Gamma = \emptyset$. Intuitively, an extended stack symbol represents some finite sequence $\gamma \in \Gamma^*$ of (original) stack symbols in the following sense: $U(q)$ tries to capture the set of all states $u \in S$ that either $q\gamma \sqsubseteq_{\text{op}} u$ or $u \sqsubseteq_{\text{op}} q\gamma$ (depending on the context). Thus, U acts like a representative for the “fragment” γ under the context of simulation preorder. Note that E (the set of extended stack symbols) is finite because both Q and S is finite. We now extend the pPDA $(Q, \Gamma, L, \rightsquigarrow)$ with E .

Definition 5.4. *The extended pPDA is defined as $(Q, \Gamma \cup E, L, \rightsquigarrow)$. The set of extended configurations (resp. basically-extended configurations), denoted by \mathcal{E} (resp. by \mathcal{E}_b), is defined by $\mathcal{E} := Q \times (\Gamma^* \cdot (E + \epsilon))$ (resp. $\mathcal{E}_b := Q \times E$).*

Instead of considering $Q \times (\Gamma + E)^*$ as the extended configurations, we only consider elements from \mathcal{E} . This is because we only need elements in \mathcal{E} to complete the tableaux proof system to be established in Section 5.3. Moreover, the elements from \mathcal{E}_b will serve as certain terminal leaves in the tableaux proof system. Note that $\mathcal{E} \setminus \mathcal{E}_b (= Q \times (\epsilon + \Gamma^+ \cdot (E + \epsilon)))$ is the set of all configurations where the extended stack symbol (if it occurs) is preceded by a non-empty sequence of (original) stack symbols. Also note

that we do not modify \succrightarrow in the extension of the pPDA; thus an extended configuration pU with $U \in E$ has no outgoing transitions in the PA induced by the extended pPDA.

In the remaining part of the chapter, if otherwise stated, we will refer to the extended pPDA whenever the notion of pPDA is encountered (e.g., when the PA induced by the pPDA is of concern). We use U, V to range over extended stack symbols, while using α, β, γ to range over $\Gamma^* \cdot (E + \epsilon)$. We also override \rightarrow to be the disjoint union of the transitions of both the fPA and the PA induced by the pPDA.

The following definition extends simulation preorder to extended configurations. Since we focus only on the simulation preorder between the extended pPDA and the fPA, we only consider binary relations between them.

Definition 5.5. *Let*

$$\mathcal{R}_b := \{(qU, s) \in \mathcal{E}_b \times S \mid s \in U(q)\} \cup \{(s, qU) \in S \times \mathcal{E}_b \mid s \in U(q)\} .$$

A binary relation $\mathcal{R} \subseteq \mathcal{E} \odot S$ is an extended op-simulation iff for all $(s, s') \in \mathcal{R}$:

- if $(s, s') \in \mathcal{E}_b \odot S$ then $(s, s') \in \mathcal{R}_b$;
- if $(s, s') \in (\mathcal{E} \setminus \mathcal{E}_b) \odot S$ then (i) $Act(s) = Act(s')$ and (ii) whenever $s \xrightarrow[\text{nc}]{a} \mu$ there exists $s' \xrightarrow[\text{op}]{a} \mu'$ such that $\mu \mathcal{R} \mu'$.

The extended op-similarity, denoted by $\sqsubseteq_{e, \text{op}}$, is the union of all extended op-simulations.

By definition, $\sqsubseteq_{e, \text{op}}$ extends \sqsubseteq_{op} by adding pairs in \mathcal{R}_b . It can be easily verified that $\sqsubseteq_{e, \text{op}}$ itself is an extended op-simulation. Intuitively, \mathcal{R}_b contains all pairs of the form (qU, s) or (s, qU) such that “ $qU \sqsubseteq s$ ” or “ $s \sqsubseteq qU$ ”. The following fact shows that $\sqsubseteq_{e, \text{op}}$ is a legitimate extension of \sqsubseteq_{op} .

Proposition 5.1. $\sqsubseteq_{e, \text{op}} \cap ((Q \times \Gamma^*) \odot S) = \sqsubseteq_{\text{op}} \cap ((Q \times \Gamma^*) \odot S)$, where \sqsubseteq_{op} refers to the simulation preorder on the disjoint union of the fPA and the PA induced by the original pPDA.

Proof. The result follows from the facts that $\sqsubseteq_{\text{op}} \cap ((Q \times \Gamma^*) \odot S)$ is an extended op-simulation, and $\sqsubseteq_{e, \text{op}} \cap ((Q \times \Gamma^*) \odot S)$ is an op-simulation on the disjoint union of the fPA and the original pPDA. \square

Below we define stepwise approximants of $\sqsubseteq_{e, \text{op}}$. The purpose to introduce such notion is to prove the soundness of the tableaux proof system for $\sqsubseteq_{e, \text{op}}$.

Definition 5.6. The family $\{\sqsubseteq_{e, \text{op}}^n\}_{n \in \mathbb{N}_0}$ is inductively defined as follows:

- $\sqsubseteq_{e, \text{op}}^0 := \{(s, s') \in (\mathcal{E} \setminus \mathcal{E}_b) \odot S \mid Act(s) = Act(s')\} \cup \mathcal{R}_b$;

- $\sqsubseteq_{e,op}^{n+1}$ is defined as the following set:

$$\begin{aligned} \sqsubseteq_{e,op}^{n+1} := & \mathcal{R}_b \cup \left\{ (s, s') \in (\mathcal{E} \setminus \mathcal{E}_b) \odot S \mid \text{Act}(s) = \text{Act}(s'), \text{ and} \right. \\ & \left. \text{for all } s \xrightarrow[\text{nc}]{a} \mu, \text{ there exists } s' \xrightarrow[\text{op}]{a} \mu' \text{ such that } \mu \sqsubseteq_{e,op}^n \mu' \right\} . \end{aligned}$$

By definition and an induction on $n \in \mathbb{N}$, we can easily obtain the following fact.

Lemma 5.2. *For all $n \in \mathbb{N}$, $\sqsubseteq_{e,op}^{n+1} \subseteq \sqsubseteq_{e,op}^n$ and $\sqsubseteq_{e,op} \subseteq \sqsubseteq_{e,op}^n$.*

Proof. We proceed by induction on $n \in \mathbb{N}$. The base step $n = 0$ is straightforward. Below we consider the inductive step. Assume $\sqsubseteq_{e,op}^{n+1} \subseteq \sqsubseteq_{e,op}^n$ and $\sqsubseteq_{e,op} \subseteq \sqsubseteq_{e,op}^n$, we prove that the arguments hold for $n + 1$.

Let $(s, s') \in \sqsubseteq_{e,op}^{n+2}$. By definition, either $(s, s') \in \mathcal{R}_b$, or $\text{Act}(s) = \text{Act}(s')$ and for all $s \xrightarrow[\text{nc}]{a} \mu$, there exists $s' \xrightarrow[\text{op}]{a} \mu'$ such that $\mu \sqsubseteq_{e,op}^{n+1} \mu'$. If $(s, s') \in \mathcal{R}_b$, then $(s, s') \in \sqsubseteq_{e,op}^{n+1}$; otherwise, by $\sqsubseteq_{e,op}^{n+1} \subseteq \sqsubseteq_{e,op}^n$ we also have $(s, s') \in \sqsubseteq_{e,op}^{n+1}$ from definition. Since (s, s') is arbitrarily chosen, we obtain $\sqsubseteq_{e,op}^{n+2} \subseteq \sqsubseteq_{e,op}^{n+1}$.

Now let $(s, s') \in \sqsubseteq_{e,op}$. By definition, either $(s, s') \in \mathcal{R}_b$, or $\text{Act}(s) = \text{Act}(s')$ and for all $s \xrightarrow[\text{nc}]{a} \mu$, there exists $s' \xrightarrow[\text{op}]{a} \mu'$ such that $\mu \sqsubseteq_{e,op} \mu'$. If $(s, s') \in \mathcal{R}_b$, then $(s, s') \in \sqsubseteq_{e,op}^{n+1}$; otherwise, by $\sqsubseteq_{e,op} \subseteq \sqsubseteq_{e,op}^n$ we also have $(s, s') \in \sqsubseteq_{e,op}^{n+1}$ from definition. Thus $\sqsubseteq_{e,op} \subseteq \sqsubseteq_{e,op}^{n+1}$ by the arbitrary choice of (s, s') .

From the previous two paragraphs, the inductive step is completed. \square

Intuitively, $s \sqsubseteq_{e,op}^n s'$ holds iff s' can mimic the behaviour of s up to n steps. It is thus seemingly true that $\bigcap_{n \in \mathbb{N}} \sqsubseteq_{e,op}^n$ equals $\sqsubseteq_{e,op}$. Below we prove that they are indeed equal.

Proposition 5.2. *For all $(s, s') \in \mathcal{E} \odot S$, $s \sqsubseteq_{e,op} s'$ iff $s \sqsubseteq_{e,op}^n s'$ for all $n \in \mathbb{N}$.*

Proof. Define $\sqsubseteq_{e,op}^\omega := \bigcap_{n \in \mathbb{N}} \sqsubseteq_{e,op}^n$. We prove that $\sqsubseteq_{e,op}^\omega = \sqsubseteq_{e,op}$. One direction $\sqsubseteq_{e,op} \subseteq \sqsubseteq_{e,op}^\omega$ follows directly from Lemma 5.2. As to the other direction ($\sqsubseteq_{e,op}^\omega \subseteq \sqsubseteq_{e,op}$), we prove that $\sqsubseteq_{e,op}^\omega$ is an extended op-simulation.

Fix $(s, s') \in \sqsubseteq_{e,op}^\omega$ and $a \in L \cup \text{Act}$. If $(s, s') \in \mathcal{E}_b \odot S$, then by definition $(s, s') \in \mathcal{R}_b$. Below we assume that $(s, s') \in (\mathcal{E} \setminus \mathcal{E}_b) \odot S$. Clearly $\text{Act}(s) = \text{Act}(s')$. Define

$$\mathcal{R} := \left\{ (u, v) \in \mathcal{E} \odot S \mid \exists \mu, \nu. \left(s \xrightarrow[\text{nc}]{a} \mu \wedge s' \xrightarrow[\text{nc}]{a} \nu \wedge (u, v) \in [\mu] \times [\nu] \right) \right\} .$$

Then \mathcal{R} is finite. Consider any $(u, v) \in \mathcal{R}$. If $(u, v) \notin \sqsubseteq_{e,op}^\omega$, there is a minimal $N(u, v) \in \mathbb{N}$ such that $(u, v) \notin \sqsubseteq_{e,op}^{N(u,v)}$. Define

$$N := \max\{N(u, v) \mid (u, v) \in \mathcal{R} \setminus \sqsubseteq_{e,op}^\omega\}$$

where $\max \emptyset := 0$. By Lemma 5.2, we have $\mathcal{R} \cap \sqsubseteq_{e, \text{op}}^N = \mathcal{R} \cap \sqsubseteq_{e, \text{op}}^\omega$. Since $s \sqsubseteq_{e, \text{op}}^{N+1} s'$, for all $s \xrightarrow[\text{nc}]{a} \mu$, there exists $s' \xrightarrow[\text{op}]{a} \mu'$ such that $\mu \sqsubseteq_{e, \text{op}}^N \mu'$. Then $\mu(\sqsubseteq_{e, \text{op}}^N \cap \mathcal{R})\mu'$. Thus $\mu \sqsubseteq_{e, \text{op}}^\omega \mu'$ by $\mathcal{R} \cap \sqsubseteq_{e, \text{op}}^N = \mathcal{R} \cap \sqsubseteq_{e, \text{op}}^\omega$. \square

5.3 Tableaux Proof System

In this section, we present a tableaux proof system for the simulation preorder between pPDA and fPA, based on which we prove the EXPTIME-membership of op-simulation preorder. Below we fix a pPDA $(Q, \Gamma, L, \succrightarrow)$ and a fPA $(S, \text{Act}, \rightarrow)$. We extend $(Q, \Gamma, L, \succrightarrow)$ with extended stack symbols as described in Section 5.2. As before, we override \rightarrow to be the set of transitions of both the fPA and the PA induced by the pPDA.

By Proposition 5.1, the decision problem for op-simulation preorder can be reformulated as follows: given a pair $(s, s') \in (Q \times \Gamma) \odot S$, decide if $s \sqsubseteq_{e, \text{op}} s'$. Note that we only consider elements in $Q \times \Gamma$ (instead of $Q \times \Gamma^*$). This is because for any $pX\alpha \in Q \times \Gamma^+$, we can always (i) add a fresh control state p_{init} and a new stack symbol X_{init} , and (ii) augment \succrightarrow with the set

$$\{p_{\text{init}}X_{\text{init}} \xrightarrow{a} \mu_\alpha \mid pX \xrightarrow{a} \mu\}$$

of extra transition rules so that $p_{\text{init}}X_{\text{init}}$ mimics the behaviour of $pX\alpha$. For the sake of simplicity, we abbreviate $\sqsubseteq_{e, \text{op}}$ (resp. $\sqsubseteq_{e, \text{op}}^n$) as \sqsubseteq_{op} (resp. $\sqsubseteq_{\text{op}}^n$).

We follow Stirling's tableaux proof system [70, 69] to develop the tableaux system in our case. A tableaux is a goal-directed proof system that consists of a set of goals **Goals** and a set **RULE** of rules which describes how a goal can be expanded into sub-goals. Graphically, a rule can be viewed as a proof step:

$$\frac{\text{goal}}{\text{goal}_1, \dots, \text{goal}_n},$$

where **goal** is what currently to be “proved” is and $\text{goal}_1, \dots, \text{goal}_n$ are the subgoals to which **goal** is reduced. Each rule is backward sound: in the rule depicted above, if all goal_i ($1 \leq i \leq n$) are true then so is **goal**. An application of a rule is to make all the sub-goals children of **goal** (in a tree). Then a tableaux is a tree built from a specified goal (the root of the tree) and repeated application of rules. The leaves of a tableaux are divided into *terminal* and *nonterminal* leaves. Terminal leaves are further divided into *successful* and *unsuccessful* leaves. A tableaux is *successful* iff it is finite and all its leaves are successful.

Below we formulate our tableaux proof system. Firstly, we introduce goals in our tableaux proof system.

Definition 5.7. Define $\text{Goals} := \mathcal{E} \odot S$ to be the set of goals. A goal $(s, s') \in \text{Goals}$ is successful if either one of the following three conditions hold:

- $(s, s') = (pU, s')$ such that $s' \in U(p)$;
- $(s, s') = (s, pU)$ such that $s \in U(p)$;
- $(s, s') \in (\mathcal{E} \setminus \mathcal{E}_b) \odot S$ and $\text{Act}(s) = \text{Act}(s') = \emptyset$.

A goal $(s, s') \in \text{Goals}$ is unsuccessful if either one of the following three conditions hold:

- $(s, s') = (pU, s')$ such that $s' \notin U(p)$;
- $(s, s') = (s, pU)$ such that $s \notin U(p)$;
- $(s, s') \in (\mathcal{E} \setminus \mathcal{E}_b) \odot S$ and $\text{Act}(s) \neq \text{Act}(s')$.

In this chapter, a goal $(s, s') \in \text{Goals}$ is rewritten as “ $s \sqsubseteq s'$ ”. Intuitively, the goal $s \sqsubseteq s'$ corresponds to a guess that the claim $s \sqsubseteq_{\text{op}} s'$ is correct.

Then, we introduce rules of our tableaux proof system. Formally, a rule

$$\frac{\text{goal}}{\text{goal}_1, \dots, \text{goal}_n}$$

can be viewed as a pair $(\text{goal}, \{\text{goal}_1, \dots, \text{goal}_n\})$, which is an element of $\text{Goals} \times 2^{\text{Goals}}$. There are two kinds of rules: UNF (unfolding) and RED (reduction). Intuitively, a rule of type UNF expands a goal $s \sqsubseteq s'$ one-step further (cf. Lemma 5.4) and a rule of RED reduces a goal $pA\alpha \sqsubseteq u$ (resp. $u \sqsubseteq pA\alpha$) to $pAU \sqsubseteq u$ (resp. $u \sqsubseteq pAU$) together with all information at α encoded in U . For the sake of clarity, we use subscript ‘a’ to indicate the case $s \sqsubseteq s' \in \mathcal{E} \times S$ and subscript ‘b’ to indicate the case for $s \sqsubseteq s' \in S \times \mathcal{E}$.

Definition 5.8. The set $\text{UNF}_{\text{op}} \subseteq \text{Goals} \times 2^{\text{Goals}}$ of unfolding rules is defined as follows: $(s \sqsubseteq s', \mathcal{R}) \in \text{UNF}_{\text{op}}$ iff

- $s \sqsubseteq s' \in (Q \times (\Gamma \cdot (E + \epsilon))) \odot S$ and $\text{Act}(s) = \text{Act}(s') \neq \emptyset$, and
- $\mathcal{R} \subseteq \text{Succ}(s) \times \text{Succ}(s')$ and for all $s \xrightarrow[\text{nc}]{a} \mu$, there exists $s' \xrightarrow[\text{op}]{a} \nu$ such that $\mu \mathcal{R} \nu$.

The set $\text{RED}_{\text{op}} \subseteq \text{Goals} \times 2^{\text{Goals}}$ of reduction rules is defined as $\text{RED}_{\text{op}}^a \cup \text{RED}_{\text{op}}^b$, for which

- $(s \sqsubseteq s', \mathcal{R}) \in \text{RED}_{\text{op}}^a$ iff $s \sqsubseteq s' = pA\alpha \sqsubseteq u$ with $u \in S$ such that
 - $\text{Act}(pA) = \text{Act}(u) \neq \emptyset$ and $\alpha \in \Gamma^+ \cdot (E + \epsilon)$, and
 - $\mathcal{R} = \{pAU \sqsubseteq u\} \cup \{q\alpha \sqsubseteq v \mid q \in Q, v \in U(q)\}$ for some $U \in E$.
- $(s \sqsubseteq s', \mathcal{R}) \in \text{RED}_{\text{op}}^b$ iff $s \sqsubseteq s' = u \sqsubseteq pA\alpha$ with $u \in S$ such that
 - $\text{Act}(u) = \text{Act}(pA) \neq \emptyset$ and $\alpha \in \Gamma^+ \cdot (E + \epsilon)$, and
 - $\mathcal{R} = \{u \sqsubseteq pAU\} \cup \{v \sqsubseteq q\alpha \mid q \in Q, v \in U(q)\}$ for some $U \in E$.

The set RULE_{op} of rules is defined by: $\text{RULE}_{\text{op}} := \text{UNF}_{\text{op}} \cup \text{RED}_{\text{op}}$.

With goals and rules defined, we present our tableaux proof system. Firstly, we introduce the notion of vertex-labelled rooted tree.

Definition 5.9. A vertex-labelled rooted tree is a pair $(\mathcal{T}, \mathcal{L})$ for which $\mathcal{T} = (V(\mathcal{T}), E(\mathcal{T}))$ is a rooted directed tree (with vertex set $V(\mathcal{T})$ and edge set $E(\mathcal{T})$) and \mathcal{L} is a function which assigns to each vertex of \mathcal{T} a goal.

Let $(\mathcal{T}, \mathcal{L})$ be a vertex-labelled rooted tree. A leaf z of \mathcal{T} is successful if either $\mathcal{L}(z)$ is successful, or there is $z' \in V(\mathcal{T})$ such that (i) $z' \neq z$, (ii) z' lies on the path from the root of \mathcal{T} to z and (iii) $\mathcal{L}(z') = \mathcal{L}(z)$. A leaf z of \mathcal{T} is unsuccessful if $\mathcal{L}(z)$ is unsuccessful. A leaf z of \mathcal{T} is terminal if either z is successful or unsuccessful; otherwise it is non-terminal.

Then the notion of tableaux tree (which is the core notion of our tableaux proof system) is defined as a subclass of vertex-labelled rooted trees.

Definition 5.10. An op-tableaux tree is a vertex-labelled rooted tree inductively defined as follows:

- a single vertex labelled with a goal is an op-tableaux tree (in this base case the sole vertex is the root of the tree);
- if
 1. $(\mathcal{T}, \mathcal{L})$ is a op-tableaux tree, and
 2. $z \in V(\mathcal{T})$ is non-terminal in \mathcal{T} and is either a leaf of \mathcal{T} or the sole vertex of \mathcal{T} (which in this case is also the root of \mathcal{T}), and
 3. $(\mathcal{L}(z), \mathcal{R})$ is a rule in RULE_{op} ,

then $((V(\mathcal{T}) \cup V', E(\mathcal{T}) \cup \{(z, z') \mid z' \in V'\}), \mathcal{L} \cup \mathcal{L}')$ is also an op-tableaux tree, where V' is a fresh new set of vertices with $|V'| = |\mathcal{R}|$ and \mathcal{L}' is a bijection from V' to \mathcal{R} .

An op-tableaux tree is successful if either it consists of a sole successful vertex (which is the root) or all its leaves are successful.

Intuitively, a tableaux tree is constructed inductively from an initial goal and (finitely) repeated application of rules. In the following, we show the soundness and completeness of our tableaux proof system, i.e., $s \sqsubseteq_{\text{op}} s'$ iff there exists an op-tableaux tree rooted at $s \sqsubseteq s'$, for all goals $s \sqsubseteq s'$.

We first show that the tableaux trees defined in Definition 5.10 have the following finiteness property. In the following, we define a *suffix* predicate $\text{suff}(\cdot, \cdot)$ by: $\text{suff}(\beta, \alpha)$ holds iff $\alpha = \gamma\beta$ for some γ .

Lemma 5.3. Let Suff be the following set:

$$\left\{ \beta \in \Gamma^* \mid \exists \alpha \in \Gamma^*. \left(\left(\exists q \in Q \exists p A \xrightarrow{a} \mu. (\mu(q\alpha) > 0) \right) \wedge \text{suff}(\beta, \alpha) \right) \right\} .$$

Define

- $G_{\text{suff}}^a := \{p\beta\alpha \sqsubseteq u \mid \beta \in \Gamma \cup \text{Suff}, \alpha \in E \cup \{\epsilon\}, p \in Q, u \in S\}$ and
- $G_{\text{suff}}^b := \{u \sqsubseteq p\beta\alpha \mid \beta \in \Gamma \cup \text{Suff}, \alpha \in E \cup \{\epsilon\}, p \in Q, u \in S\}$.

If $s \sqsubseteq s' \in \text{Gl}_{\text{suff}}^{\text{a}}$, then $\mathcal{L}(V(\mathcal{T})) \subseteq \text{Gl}_{\text{suff}}^{\text{a}}$ for all op-tableaux trees $(\mathcal{T}, \mathcal{L})$ with root label $s \sqsubseteq s'$. Analogously, if $s \sqsubseteq s' \in \text{Gl}_{\text{suff}}^{\text{b}}$, then $\mathcal{L}(V(\mathcal{T})) \subseteq \text{Gl}_{\text{suff}}^{\text{b}}$ for all op-tableaux trees $(\mathcal{T}, \mathcal{L})$ with root label $s \sqsubseteq s'$.

Proof. The proof is a simple induction on the construction of tableaux trees illustrated in Definition 5.10. The base step where $(\mathcal{T}, \mathcal{L})$ is a single root is straightforward. It is also clear that $\text{Gl}_{\text{suff}}^{\text{a}}$ and $\text{Gl}_{\text{suff}}^{\text{b}}$ are closed under rule application of UNF_{op} and RED_{op} (cf. Definition 5.8). \square

Then we show the soundness of our tableaux proof system, i.e., if there exists a successful tableaux tree rooted at $s \sqsubseteq s'$ then $s \sqsubseteq_{\text{op}} s'$. We first show that rules of UNF_{op} and RED_{op} are backward sound, i.e., if all the subgoals are correct then the goal "to be proved" is correct.

Lemma 5.4. *Let $(s \sqsubseteq s', \mathcal{R}) \in \text{UNF}_{\text{op}}$. If $r \sqsubseteq_{\text{op}}^n r'$ for all $r \sqsubseteq r' \in \mathcal{R}$, then $s \sqsubseteq_{\text{op}}^{n+1} s'$.*

Proof. The result follows directly from Definition 5.8 and Definition 5.6. \square

Lemma 5.5. *Let $(pA\alpha \sqsubseteq u, \{pAU \sqsubseteq u\} \cup \text{Gl}_{\alpha, U}^{\text{a}}) \in \text{RED}_{\text{op}}^{\text{a}}$ where*

$$\text{Gl}_{\alpha, U}^{\text{a}} := \{q\alpha \sqsubseteq v \mid q \in Q, v \in U(q)\} .$$

For all $n \in \mathbb{N}_0$, if $pAU \sqsubseteq_{\text{op}}^{n+1} u$ and $q\alpha \sqsubseteq_{\text{op}}^n v$ for all $q\alpha \sqsubseteq v \in \text{Gl}_{\alpha, U}^{\text{a}}$, then $pA\alpha \sqsubseteq_{\text{op}}^{n+1} u$.

Proof. We prove by induction on n that for all $p\gamma\alpha \sqsubseteq u \in \text{Goals}$ with $\gamma \in \Gamma^+$, it holds that for all $U \in E$, if $p\gamma U \sqsubseteq_{\text{op}}^{n+1} u$ and $q\alpha \sqsubseteq_{\text{op}}^n v$ for all $q\alpha \sqsubseteq v \in \text{Gl}_{\alpha, U}^{\text{a}}$, then $p\gamma\alpha \sqsubseteq_{\text{op}}^{n+1} u$.

Base Step: $n = 0$. Assume that $p\gamma U \sqsubseteq_{\text{op}}^1 u$ and $q\alpha \sqsubseteq_{\text{op}}^0 v$ for all $q\alpha \sqsubseteq v \in \text{Gl}_{\alpha, U}^{\text{a}}$. Since $p\gamma U \sqsubseteq_{\text{op}}^1 u$, for all $p\gamma \xrightarrow{\text{nc}} \mu$, there is $u \xrightarrow{\text{nc}} \nu$ such that $\mu U \sqsubseteq_{\text{op}}^0 \nu$ (cf. Lemma 5.1). Let $w : [\mu_U] \times [\nu] \rightarrow [0, 1]$ be a weight function for $\mu_U \sqsubseteq_{\text{op}}^0 \nu$. (We can restrict the domain to $[\mu_U] \times [\nu]$ since all other values are zero.) We define a weight function $w' : [\mu_\alpha] \times [\nu] \rightarrow [0, 1]$ by: $w'(q\beta\alpha, v) = w(q\beta U, v)$ for all $q\beta \in [\mu]$ (note that $\beta \in \Gamma^*$) and $v \in [\nu]$. We prove that w' is a weight function for $\mu_\alpha \sqsubseteq_{\text{op}}^0 \nu$. The first two conditions in Definition 4.4 are straightforward to verify. For the third condition, suppose $w'(q\beta\alpha, v) > 0$ with $q\beta \in [\mu]$. Then $w(q\beta U, v) > 0$ and hence $q\beta U \sqsubseteq_{\text{op}}^0 v$. If $\beta \neq \epsilon$, then by Definition 5.6 $\text{Act}(q\beta) = \text{Act}(v)$ and we have $q\beta\alpha \sqsubseteq_{\text{op}}^0 v$; if $\beta = \epsilon$, then $v \in U(q)$ and $q\alpha \sqsubseteq v \in \text{Gl}_{\alpha, U}^{\text{a}}$, which further implies $q\alpha \sqsubseteq_{\text{op}}^0 v$. In either case $q\beta\alpha \sqsubseteq_{\text{op}}^0 v$. So w' is a weight function for the statement $\mu_\alpha \sqsubseteq_{\text{op}}^0 \nu$. Also from $p\gamma U \sqsubseteq_{\text{op}}^1 u$ we have $\text{Act}(p\gamma\alpha) = \text{Act}(u)$. Thus $p\gamma\alpha \sqsubseteq_{\text{op}}^1 u$.

Inductive Step: Assume that $p\gamma U \sqsubseteq_{\text{op}}^{n+2} u$ and $q\alpha \sqsubseteq_{\text{op}}^{n+1} v$ for all $q\alpha \sqsubseteq v \in \text{Gl}_{\alpha, U}^{\text{a}}$. We prove that $p\gamma\alpha \sqsubseteq_{\text{op}}^{n+2} u$. Since $p\gamma U \sqsubseteq_{\text{op}}^{n+2} u$, for any $p\gamma \xrightarrow{\text{nc}} \mu$, there exists $u \xrightarrow{\text{nc}} \nu$ such that $\mu U \sqsubseteq_{\text{op}}^{n+1} \nu$. Consider any $(q\beta U, v) \in \sqsubseteq_{\text{op}}^{n+1}$

with $\beta \in \Gamma^*$ and $v \in S$: if $\beta = \epsilon$ then $q\beta\alpha \sqsubseteq_{\text{op}}^{n+1} v$ since $q\alpha \sqsubseteq v \in \text{Gl}_{\alpha,U}^a$; if $\beta \in \Gamma^+$ then we have $q\beta\alpha \sqsubseteq_{\text{op}}^{n+1} v$ by induction hypothesis; in any case, we have $q\beta\alpha \sqsubseteq_{\text{op}}^{n+1} v$. Thus by the same construction of weight function in the base step, we have $\mu_\alpha \sqsubseteq_{\text{op}}^{n+1} \nu$. Also we have $\text{Act}(p\gamma\alpha) = \text{Act}(u)$. Thus $p\gamma\alpha \sqsubseteq_{\text{op}}^{n+2} u$. \square

We can prove a symmetrical case for RED_{op}^b as follows.

Lemma 5.6. *Let $(u \sqsubseteq pA\alpha, \{u \sqsubseteq pAU\} \cup \text{Gl}_{\alpha,U}^b) \in \text{RED}_{\text{op}}^b$ where*

$$\text{Gl}_{\alpha,U}^b := \{v \sqsubseteq q\alpha \mid q \in Q, v \in U(q)\} .$$

For all $n \in \mathbb{N}_0$, if $u \sqsubseteq_{\text{op}}^{n+1} pAU$ and $v \sqsubseteq_{\text{op}}^n q\alpha$ for all $v \sqsubseteq q\alpha \in \text{Gl}_{\alpha,U}^b$, then $u \sqsubseteq_{\text{op}}^{n+1} pA\alpha$.

Proof. The proof can be carried out in a completely symmetric fashion from the one of Lemma 5.5. \square

Based on Lemma 5.4 through Lemma 5.6, we prove the soundness of our tableaux system as follows.

Proposition 5.3. *For all goals $s \sqsubseteq s'$, if there exists a successful op-tableaux tree rooted at a vertex labelled with $s \sqsubseteq s'$, then $s \sqsubseteq_{\text{op}} s'$.*

Proof. We only prove the case when $s \sqsubseteq s' \in \mathcal{E} \times S$, the other case is completely symmetrical. Let $p_0\alpha_0 \sqsubseteq u_0 = s \sqsubseteq s'$. Suppose $p_0\alpha_0 \not\sqsubseteq_{\text{op}} u_0$ and $(\mathcal{T}, \mathcal{L})$ is a successful op-tableaux tree rooted at $p_0\alpha_0 \sqsubseteq u_0$. Let the root of $(\mathcal{T}, \mathcal{L})$ be z_0 . By Proposition 5.2, there exists $n_0 \in \mathbb{N}_0$ such that $p_0\alpha_0 \sqsubseteq_{\text{op}}^{n_0} u_0$ but $p_0\alpha_0 \not\sqsubseteq_{\text{op}}^{n_0+1} u_0$. Note that we have $(p_0\alpha_0, u_0) \in \sqsubseteq_{\text{op}}^0$ or otherwise the goal $p_0\alpha_0 \sqsubseteq u_0$ would be unsuccessful. By the backward soundness of UNF_{op} and RED_{op} (cf. Lemma 5.4 and Lemma 5.5), we can obtain the following statements (\dagger):

- if the rule applied to $p_0\alpha_0 \sqsubseteq u_0$ (in the inductive construction of $(\mathcal{T}, \mathcal{L})$) belongs to UNF_{op} , then there exists a child z_1 of z_0 labelled with $p'\alpha' \sqsubseteq u'$ such that $p'\alpha' \not\sqsubseteq_{\text{op}}^{n_0} u'$;
- if the rule applied to $p_0\alpha_0 \sqsubseteq u_0$ lies in RED_{op}^a , then there exists a child z_1 of z_0 whose label is either $p_0AU \sqsubseteq u_0$ with $p_0AU \not\sqsubseteq_{\text{op}}^{n_0+1} u_0$, or $q\alpha \sqsubseteq v$ with $q\alpha \not\sqsubseteq_{\text{op}}^{n_0} v$ for some $q\alpha \sqsubseteq v \in \text{Gl}_{\alpha,U}^a$, where $A \in \Gamma$, $U \in E$ and $\alpha \in \Gamma^+ \cdot (E + \epsilon)$ are specified by the rule (cf. Definition 5.8).

In either case, there is a child z_1 of z_0 labelled with $p_1\alpha_1 \sqsubseteq u_1$ such that $p_1\alpha_1 \not\sqsubseteq_{\text{op}}^{n_0+1} u_1$. Choose such z_1 according to (\dagger) arbitrarily. Let $n_1 \in \mathbb{N}_0$ be such that $p_1\alpha_1 \not\sqsubseteq_{\text{op}}^{n_1+1} u_1$ and $p_1\alpha_1 \sqsubseteq_{\text{op}}^{n_1} u_1$. Then $n_1 \leq n_0$. By performing (\dagger) on z_1 and $p_1\alpha_1 \sqsubseteq u_1$ and so forth, we can recursively construct a finite sequence $\{(z_i, p_i\alpha_i \sqsubseteq u_i, n_i)\}_{0 \leq i \leq k}$ of length $k+1$ ($k \geq 1$) such that $p_i\alpha_i \not\sqsubseteq_{\text{op}}^{n_i+1} u_i$

and $p_i \alpha_i \sqsubseteq_{\text{op}}^{n_i} u_i$, $\{n_i\}$ is decreasing, and the last vertex z_k is a successful leaf. Since $p_k \alpha_k \not\sqsubseteq_{\text{op}}^{n_k+1} u_k$, the goal $p_k \alpha_k \sqsubseteq u_k$ cannot be successful. So the only possibility is that there is $j < k$ such that $p_k \alpha_k \sqsubseteq u_k = p_j \alpha_j \sqsubseteq u_j$. Consider the rule application from $(z_{k-1}, p_{k-1} \alpha_{k-1} \sqsubseteq u_{k-1})$ to $(z_k, p_k \alpha_k \sqsubseteq u_k)$. By (†):

- if the rule lies in UNF_{op} , then $n_k < n_{k-1} \leq n_j$;
- if the rule lies in RED_{op}^a and $p_k \alpha_k \sqsubseteq u_k \in \text{Gl}_{\alpha, U}^a$ where $\alpha \in \Gamma^+ \cdot (E + \epsilon)$ and $U \in E$ are determined by the rule (cf. Definition 5.8), then $n_k < n_{k-1} \leq n_j$.
- If the rule lies in RED_{op}^a and $p_k \alpha_k \sqsubseteq u_k = p_k A U \sqsubseteq u_k$ where $A \in \Gamma$ and $U \in E$ are determined by the rule (cf. Definition 5.8), then $p_{k-1} \alpha_{k-1} \sqsubseteq u_{k-1} \neq p_k A U \sqsubseteq u_k$ and so $j < k - 1$. By $p_k \alpha_k \sqsubseteq u_k = p_j \alpha_j \sqsubseteq u_j$, the rule application from z_j to z_{j+1} belongs to UNF_{op} , which implies $n_{j+1} < n_j$. Hence $n_k < n_j$.

In either case, we have $n_k < n_j$. But then we have $p_k \alpha_k \not\sqsubseteq_{\text{op}}^{n_k+1} u_k$ and $p_k \alpha_k \sqsubseteq_{\text{op}}^{n_j} u_k$. Contradiction. \square

Finally, we prove the completeness of our tableaux proof system, i.e., if $s \sqsubseteq_{\text{op}} s'$ then there exists a successful op-tableaux tree with root label $s \sqsubseteq s'$. We first prove a useful lemma below.

Lemma 5.7. *Let $pA\alpha \sqsubseteq_{\text{op}} u$ and U be an extended stack symbol such that $U(q) := \{v \in S \mid q\alpha \sqsubseteq_{\text{op}} v\}$ for all $q \in Q$. Then $pAU \sqsubseteq_{\text{op}} u$.*

Proof. We prove that the binary relation

$$\mathcal{R} := \{(q\beta U, v) \mid q \in Q, \beta \in \Gamma^*, v \in S, q\beta\alpha \sqsubseteq_{\text{op}} v\}$$

is an extended op-simulation. Consider any $(q\beta U, v) \in \mathcal{R}$. If $\beta = \epsilon$, then $q\alpha \sqsubseteq_{\text{op}} v$ and $v \in U(q)$; it follows that $(qU, v) \in \mathcal{R}_b$. On the other hand, assume that $\beta \in \Gamma^+$. Then $\text{Act}(q\beta U) = \text{Act}(q\beta\alpha) = \text{Act}(v)$. Furthermore, for all $q\beta \xrightarrow[\text{nc}]{a} \mu$, by $q\beta\alpha \sqsubseteq_{\text{op}} v$ there is $v \xrightarrow[\text{op}]{a} \nu$ such that $\mu_\alpha \sqsubseteq_{\text{op}} \nu$. We prove that $\mu_U \mathcal{R} \nu$. By $\mu_\alpha \sqsubseteq_{\text{op}} \nu$, there exists a weight function $w : [\mu_\alpha] \times [\nu] \rightarrow [0, 1]$ for the statement $\mu_\alpha \sqsubseteq_{\text{op}} \nu$. (We can restrict the domain of the weight function to $[\mu_\alpha] \times [\nu]$ since values at other places are zero.) We construct a weight function $w' : [\mu_U] \times [\nu] \rightarrow [0, 1]$ by: $w'(q'\gamma U, v') = w(q'\gamma\alpha, v')$ for all $q'\gamma \in [\mu]$ and $v' \in [\nu]$ (note that $\gamma \in \Gamma^*$). Then we show that w' is a weight function for μ_U and ν . The first two conditions in Definition 4.4 are straightforward to verify. For the third condition, consider any $q'\gamma \in [\mu]$ and $v' \in [\nu]$: assume that $w'(q'\gamma U, v') > 0$; then $w(q'\gamma\alpha, v') > 0$ and $q'\gamma\alpha \sqsubseteq_{\text{op}} v'$, which implies that $(q'\gamma U, v') \in \mathcal{R}$ by definition. Thus \mathcal{R} is an extended op-simulation. \square

By a symmetrical proof, we can obtain the following lemma.

Lemma 5.8. *Let $u \sqsubseteq_{\text{op}} pA\alpha$ and U be an extended stack symbol such that $U(q) := \{v \in S \mid v \sqsubseteq_{\text{op}} q\alpha\}$ for all $q \in Q$. Then $u \sqsubseteq_{\text{op}} pAU$.*

The completeness of our tableaux proof system is as follows.

Proposition 5.4. *Let $s \sqsubseteq s' \in \text{GI}_{\text{suff}}^{\text{a}} \cup \text{GI}_{\text{suff}}^{\text{b}}$. If $s \sqsubseteq_{\text{op}} s'$ then there is a successful op-tableaux tree with root label $s \sqsubseteq s'$.*

Proof. We only prove the case for $s \sqsubseteq s' \in \text{GI}_{\text{suff}}^{\text{a}}$, the other case is completely symmetrical. Below we inductively construct a sequence $\{(\mathcal{T}_n, \mathcal{L}_n)\}_{1 \leq n \leq k}$ of op-tableaux trees, each with root label $s \sqsubseteq s'$, such that $(\mathcal{T}_k, \mathcal{L}_k)$ is a successful op-tableaux tree.

Initially, $(\mathcal{T}_1, \mathcal{L}_1)$ is the tableaux tree which contains only a root labelled with $s \sqsubseteq s'$. Then assume that $(\mathcal{T}_n, \mathcal{L}_n)$ is constructed. If all leaves of $(\mathcal{T}_n, \mathcal{L}_n)$ are terminal, then the construction is ended. Otherwise, we choose an arbitrary non-terminal leaf z of $(\mathcal{T}_n, \mathcal{L}_n)$ and construct $(\mathcal{T}_{n+1}, \mathcal{L}_{n+1})$ as follows:

1. if $\mathcal{L}_n(z) = pA\alpha \sqsubseteq u$ with $\alpha \in E \cup \{\epsilon\}$, then we apply the UNF_{op} rule

$$(pA\alpha \sqsubseteq u, \{(q\beta, v) \in \text{Succ}(pA\alpha) \times \text{Succ}(u) \mid q\beta \sqsubseteq_{\text{op}} v\})$$

to the vertex z (to form $(\mathcal{T}_{n+1}, \mathcal{L}_{n+1})$);

2. if $\mathcal{L}_n(z) = pA\alpha \sqsubseteq u$ with $\alpha \in \Gamma^+ \cdot (E + \epsilon)$, then we apply the RED_{op} rule

$$(pA\alpha \sqsubseteq u, \{pAU \sqsubseteq u\} \cup \text{GI}_{\alpha, U}^{\text{a}})$$

where U is defined by: $U(q) := \{v \in S \mid q\alpha \sqsubseteq_{\text{op}} v\}$ for all $q \in Q$. Note that from Lemma 5.7, we have $pAU \sqsubseteq_{\text{op}} u$.

In the inductive construction above, one easily sees that we only "append" goals $s \sqsubseteq s'$ such that $s \sqsubseteq_{\text{op}} s'$. Thus if the inductive construction ends in a finite number of step, then the last tableaux tree $(\mathcal{T}_k, \mathcal{L}_k)$ should be successful, i.e., it won't contain any unsuccessful leaves. Below we show that the inductive construction above ends in a finite number of steps. Suppose that we obtain an infinite sequence $\{(\mathcal{T}_n, \mathcal{L}_n)\}_{n \in \mathbb{N}}$ of tableaux trees with $\mathcal{T}_n = (V(\mathcal{T}_n), E(\mathcal{T}_n))$ from the inductive construction above. Let $\mathcal{T} := (\bigcup_{n \in \mathbb{N}} V(\mathcal{T}_n), \bigcup_{n \in \mathbb{N}} E(\mathcal{T}_n))$ and $\mathcal{L} := \bigcup_{n \in \mathbb{N}} \mathcal{L}_n$. Then $(\mathcal{T}, \mathcal{L})$ is an infinite vertex-labelled tree. By the way of rule application, \mathcal{T} is also finitely branching. Thus by König's Lemma, there is an infinite path in \mathcal{T} . Then by Lemma 5.3, on such infinite path there must be vertices z, z' such that $z \neq z'$ and $\mathcal{L}(z) = \mathcal{L}(z')$. It follows that either z or z' is a successful leaf of the tableaux tree \mathcal{T}_{n^*} , where n^* is the smallest number such that z, z' are leaves of \mathcal{T}_{n^*} . Then the inductive construction should stop expanding at z (or z' , depending on which one is a leaf), which leads to a contradiction. \square

Below we give an example.

$$\begin{array}{c}
\frac{pZ \sqsubseteq s_2}{\text{UNF}} \\
\frac{\frac{pAZ \sqsubseteq s_1}{\text{RED}}}{\text{UNF}} \\
pZ \sqsubseteq s_2, \frac{pAU \sqsubseteq s_1}{\text{to subtree I, } pU \sqsubseteq s_2} \text{UNF} \\
\frac{\text{subtree I : } pAAU \sqsubseteq s_1}{\text{RED}} \\
\frac{pAV \sqsubseteq s_1}{\text{UNF, } pAU \sqsubseteq s_1} \\
pV \sqsubseteq s_1, \frac{pAAV \sqsubseteq s_1}{pAV \sqsubseteq s_1} \text{RED}
\end{array}$$

Figure 5.1: A Tableaux Tree

Example 5.1. Let the pPDA (Q, L, Γ, \mapsto) and the fPA (S, Act, \rightarrow) be such that

- $Q = \{p\}$, $L = \{a, b\}$, $\Gamma = \{A, Z\}$;
- $\mapsto = \{pA \xrightarrow{a} \{pAA \mapsto 0.5, p \mapsto 0.5\}, pZ \xrightarrow{b} \{pAZ \mapsto 1\}\}$;
- $S = \{s_1, s_2\}$, $Act = \{a, b\}$;
- $\rightarrow = \{s_1 \xrightarrow{a} \{s_1 \mapsto 1\}, s_1 \xrightarrow{a} \{s_1 \mapsto 0.5, s_2 \mapsto 0.5\}, s_2 \xrightarrow{b} \{s_1 \mapsto 1\}\}$.

Intuitively, the pPDA models a counter with random increase and decrease operation. It can be verified that $pZ \sqsubseteq_{\text{nc}} s_2$ since the relation

$$\{(pZ, s_2)\} \cup \{(p\alpha Z, s_1) \mid \alpha \in A^+\}$$

is an nc-simulation. Below we use our tableaux proof system to “prove” that $pZ \sqsubseteq_{\text{nc}} s_2$. A successful nc-tableaux tree is depicted in Fig. 5.1, where $U := \{p \mapsto \{s_2\}\}$ and $V := \{p \mapsto \{s_1\}\}$.

We have ended the demonstration of the soundness and completeness of our tableaux proof system. Based on the tableaux proof system, we develop an algorithm that decides \sqsubseteq_{op} . The algorithm will use a partition-refinement technique to achieve the EXPTIME-upperbound, which is the main result of this chapter. Below we denote by M the integrated size of the pPDA and the fPA, where the numerical values (probability values) are represented in binary.

Theorem 5.1. The problem whether $s \sqsubseteq_{\text{op}} s'$ for a given $(s, s') \in (Q \times \Gamma) \odot S$ can be decided in $\mathcal{O}(h(M) \cdot 8^{|S| \cdot |Q|})$ time where h is a polynomial function. Thus, if $|Q|$ and $|S|$ are fixed, then the problem can be decided in PTIME.

Proof. We assume that $s \sqsubseteq s' \in (Q \times \Gamma) \times S$ and $\text{op} = \text{c}$, the other cases are similar. We present a partition-refinement algorithm to decide whether

$s \sqsubseteq_c s'$. Formally, we construct a finite decreasing sequence of sets of goals $\{X_n\}_{1 \leq n \leq k}$ where the last element X_k contains all the correct goals in $\text{Gl}_{\text{suff}}^a$.

The construction is as follows. Initially, $X_1 = \text{Gl}_{\text{suff}}^a$. Then $X_{n+1} \subseteq \text{Gl}_{\text{suff}}^a$ is constructed from X_n as follows: $s \sqsubseteq s' \in X_{n+1}$ iff (i) $s \sqsubseteq s' \in X_n$ and (ii) either $s \sqsubseteq s'$ is successful or there exists $Y \subseteq X_n$ such that $(s \sqsubseteq s', Y) \in \text{RULE}_c$. Note that $|\text{Gl}_{\text{suff}}^a| = \mathcal{O}(M^4 \cdot 2^{|S| \cdot |Q|})$.

The computation from X_n to X_{n+1} can be done in $\mathcal{O}(h'(M) \cdot 4^{|S| \cdot |Q|})$ time by the following procedure, where h' is a polynomial function. Let $s \sqsubseteq s' \in X_n$. We check whether $s \sqsubseteq s' \in X_{n+1}$ as follows:

- if $s \sqsubseteq s' = pA\alpha \sqsubseteq u$ with $\alpha \in \Gamma^+ \cdot (E + \epsilon)$, we check whether $\{pAU \sqsubseteq u\} \cup \text{Gl}_{\alpha, U}^a \subseteq X_n$ for some $U \in E$;
- if $s \sqsubseteq s' = pA\alpha \sqsubseteq u$ with $\alpha \in E \cup \{\epsilon\}$, we check whether for all $pA\alpha \xrightarrow[\text{nc}]{a} \mu$, there exists $u \xrightarrow[\text{nc}]{a} \nu$ such that $\mu X_n \nu$ (treat X_n as a binary relation such that $s_1 \sqsubseteq s_2 \in X_n$ iff $(s_1, s_2) \in X_n$); this can be checked by examining whether the following linear inequality system (with variables $\{x_\nu\}_{u \xrightarrow[\text{nc}]{a} \nu}$ and $\{y_{(s'', v)}\}_{(s'', v) \in [\mu] \times S}$) has a solution:

- $\sum_{u \xrightarrow[\text{nc}]{a} \nu} x_\nu = 1$;
- $x_\nu \geq 0$ for all $u \xrightarrow[\text{nc}]{a} \nu$;
- $\sum_{v \in S} y_{(s'', v)} = \mu(s'')$ for all $s'' \in [\mu]$;
- $\sum_{s'' \in [\mu]} y_{(s'', v)} = \sum_{u \xrightarrow[\text{nc}]{a} \nu} x_\nu \cdot \nu(v)$ for all $v \in S$;
- $y_{(s'', v)} \geq 0$ for all $(s'', v) \in [\mu] \times S$;
- $y_{(s'', v)} = 0$ whenever $(s'', v) \notin X_n$.

This can be solved in polynomial time in M (cf. [66]).

Since $X_{n+1} \subseteq X_n$, there exists $k \leq |\text{Gl}_{\text{suff}}^a|$ such that $X_{k+1} = X_k$. We show that for all $s \sqsubseteq s' \in \text{Gl}_{\text{suff}}^a$, $s \sqsubseteq_c s'$ iff $s \sqsubseteq s' \in X_k$. On one hand, assume that $s \sqsubseteq_c s'$. Let $(\mathcal{T}, \mathcal{L})$ be the tableaux tree constructed in the proof of Proposition 5.4 for the goal $s \sqsubseteq s'$. An easy induction on n shows that $\mathcal{L}(V(\mathcal{T})) \subseteq X_n$ for all $n \in \mathbb{N}$. Thus $s \sqsubseteq s' \in X_k$. On the other hand, assume that $s \sqsubseteq s' \in X_k$. Since for all goals $s_1 \sqsubseteq s_2 \in X_k$ which are not successful, there is $(s_1 \sqsubseteq s_2, Y) \in \text{RULE}_c$ such that $Y \subseteq X_k$. Thus we can iteratively apply rules to the root $s \sqsubseteq s'$ (and non-terminal leaves), which results in a successful tableaux tree similar to the construction in the proof of Proposition 5.4. Thus, $s \sqsubseteq_c s'$ by Proposition 5.3. Then the result follows from the fact that $k \leq |\text{Gl}_{\text{suff}}^a|$. \square

5.4 EXPTIME-Hardness

In this section, we show that deciding \sqsubseteq_{op} is EXPTIME-hard, whenever $\text{op} = \text{nc}$ or $\text{op} = \text{c}$. We prove this by providing a rather straightforward

reduction from the non-probabilistic EXTPIME-hardness result obtained in [52]. Our main efforts lie in the treatment of the additional “ $Act(s) = Act(s')$ ” condition in Definition 4.6 which is not involved in the definition of non-probabilistic simulation preorder. Firstly, we define a variation of \sqsubseteq_{op} .

Definition 5.11. *Let $\mathcal{M} = (S, Act, \rightarrow)$ be a PA. Define \preceq_{op} to be the union of all binary relations $\mathcal{R} \subseteq S \times S$ such that for all $(s, s') \in \mathcal{R}$, whenever $s \xrightarrow[\text{nc}]{a} \mu$ there is $s' \xrightarrow[\text{op}]{a} \mu'$ with $\mu \mathcal{R} \mu'$.*

In other words, \preceq_{op} is defined in a similar way of \sqsubseteq_{op} , however without the “ $Act(s) = Act(s')$ ” requirement. Then we embed non-probabilistic transition systems into PAs.

Definition 5.12. *A PA (S, Act, \rightarrow) is Dirac if μ is dirac for all $(s, a, \mu) \in \rightarrow$. A pPDA $(Q, \Gamma, L, \rightsquigarrow)$ is Dirac if μ is Dirac for all $(pA, a, \mu) \in \rightsquigarrow$.*

Note that a Dirac pPDA induces a Dirac PA. Dirac PAs correspond to transition systems without probability. It is not hard to verify that \preceq_{n} (over Dirac PAs) coincides with the (non-probabilistic) simulation preorder (cf. [7]) over non-probabilistic transition systems. From [52], deciding \preceq_{n} is EXPTIME-complete between Dirac pPDAs and Dirac fPAs in both direction. Below we reduce \preceq_{op} to \sqsubseteq_{op} under Dirac PAs. The following proposition allows us to focus solely on the case $\text{op} = \text{nc}$.

Proposition 5.5. *Let $\mathcal{M} = (S, Act, \rightarrow)$ be a PA. If \mathcal{M} is Dirac, then $\preceq_{\text{nc}} = \preceq_{\text{c}}$ and $\sqsubseteq_{\text{nc}} = \sqsubseteq_{\text{c}}$.*

Proof. It is clear that $\preceq_{\text{nc}} \subseteq \preceq_{\text{c}}$ and $\sqsubseteq_{\text{nc}} \subseteq \sqsubseteq_{\text{c}}$. Below we prove the reverse direction. We only prove the case $\sqsubseteq_{\text{c}} \subseteq \sqsubseteq_{\text{nc}}$, since the proof for the other is similar. Let $s \sqsubseteq_{\text{c}} s'$ and $s \xrightarrow[\text{nc}]{a} \mu$. From definition, there exists $s' \xrightarrow[\text{c}]{a} \mu'$ such that $\mu \sqsubseteq_{\text{c}} \mu'$. Since μ, μ' are Dirac, there exists $s'' \in [\mu']$ such that $s' \xrightarrow[\text{nc}]{a} \mathcal{D}[s'']$ and $\mu \sqsubseteq_{\text{c}} \mathcal{D}[s'']$. It follows from the arbitrary choice of s, s' and $s \xrightarrow[\text{nc}]{a} \mu$ that \sqsubseteq_{c} is an nc-simulation, which implies $\sqsubseteq_{\text{c}} \subseteq \sqsubseteq_{\text{nc}}$. \square

Now we reduce \preceq_{nc} between a Dirac pPDA $P = (Q, \Gamma, L, \rightsquigarrow)$ and a Dirac fPA $\mathcal{M} = (S, Act, \rightarrow)$, to \sqsubseteq_{nc} between a Dirac pPDA $(Q', \Gamma', L, \rightsquigarrow')$ and a Dirac fPA (S', Act', \rightarrow') . The reduction is as follows:

1. $Q' = Q \cup \{p_{\perp}\}$ and $S' = S \cup \{s_{\perp}\}$ where $p_{\perp} \notin Q$ and $s_{\perp} \notin S$ are fresh elements;
2. $\Gamma' = \Gamma \cup \{A_{\perp}\}$ where $A_{\perp} \notin \Gamma$ is a fresh (bottom) stack symbol;
3. $Act' = Act$;
4. $\rightsquigarrow' = \rightsquigarrow \cup \{(pA, a, \mathcal{D}[p_{\perp}]) \mid p \in Q, A \in \Gamma', a \in L \cup Act\}$;
5. $\rightarrow' = \rightarrow \cup \{(s, a, \mathcal{D}[s_{\perp}]) \mid s \in S, a \in L \cup Act\}$.

The basic idea is that we try to "amend" P and \mathcal{M} so that pairs in \preceq_{nc} will have same action sets. It is not hard to prove that for all $p\alpha \in Q \times \Gamma^*$ and $s \in S$, $p\alpha \preceq_{\text{nc}} s$ (resp. $s \preceq_{\text{nc}} p\alpha$) iff $p\alpha A_{\perp} \sqsubseteq_{\text{nc}} s$ (resp. $s \sqsubseteq_{\text{nc}} p\alpha A_{\perp}$). Thus deciding \sqsubseteq_{nc} between Dirac pPDA's and Dirac finite fPA's is EXPTIME-hard. Then we have the following theorem.

Theorem 5.2. *Deciding \sqsubseteq_{nc} and \sqsubseteq_{c} between probabilistic pushdown automata and finite probabilistic automata in both directions (of the simulation preorder) is EXPTIME-complete.*

5.5 Conclusion

In this chapter, we showed that (probabilistic) simulation preorder between a probabilistic pushdown automaton $(Q, \Gamma, L, \rightrightarrows)$ and a finite probabilistic automaton (S, Act, \rightarrow) is EXPTIME-complete. This result holds for both directions. Furthermore, if $|Q|$ and $|S|$ are fixed, then the problem is decidable in polynomial time. These results extend their non-probabilistic counterparts in [52], and are obtained by extending Colin Stirling's method [69, 70] which is originally used to demonstrate the decidability of bisimilarity on non-probabilistic pushdown automata.

Chapter 6

Bisimilarity Metric on Probabilistic Automata

In this chapter, we consider quantitative variations of (probabilistic) bisimulation equivalence (cf. Chapter 4). The motivation is that the notion of (probabilistic) bisimulation equivalence is sensitive to exact probability values in that a slight change to a probability value may cause two equivalent states inequivalent [72]. To this end, several works [72, 30, 28, 60] developed a notion of bisimilarity metric, which characterizes the distance between two states; states with smaller distance are meant to have similar behaviours. Bisimilarity metric can be viewed as a quantified extension of (probabilistic) bisimulation equivalence in that (i) the distance between two states is zero iff they are equivalent in the sense of probabilistic bisimulation equivalence, and (ii) there are various quantitative logic characterizations for it [72, 30, 28]. Here, we will focus on the bisimilarity metric defined by van Breugel and Worrell [72].

The contribution of this chapter is as follows. We show that the threshold problem of the bisimilarity metric [72], which is to decide whether the pseudometric between two states of a PA is under certain value, is in $\text{NP} \cap \text{coNP}$ and $\text{UP} \cap \text{coUP}$. This complexity result significantly improves the previous PSPACE upperbound by van Breugel *et al* [71] (cf. also [22, 25]). We obtain this result through a core notion called “self-closed sets” to be introduced in Section 6.3. In general, we show a way to check whether an arbitrary pseudometric equals the bisimilarity metric. Then the membership of $\text{UP} \cap \text{coUP}$ follows from the polynomial-size representability of the bisimilarity metric to be proved in Section 6.4.

The chapter is organized as follows. Section 6.1 introduces the notion of bisimilarity metric by van Breugel and Worrell [72]. Section 6.2 defines a notion of approximants of the bisimilarity metric, in order to show that we can focus on the more manageable notion of premetrics instead of pseudometrics. Section 6.3 introduces the notion of self-closed sets and its relation

with the bisimilarity metric. Section 6.4 shows that the bisimilarity metric is of polynomial size, which serves as the last step to the membership of $\text{NP} \cap \text{coNP}$ and $\text{UP} \cap \text{coUP}$.

6.1 Bisimilarity Metric on PAs

In this section, we introduce a quantified notion of probabilistic bisimulation defined by van Breugel and Worrell [72]. This notion will be defined as a pseudometric over states of a finite PA (fPA) for which pairs of states with distance zero are probabilistic bisimilar.

Firstly, we introduce the notions of premetrics and pseudometrics, which assign to each pair of elements a non-negative value representing the distance between the two elements.

Definition 6.1. *Let X be a non-empty set. A function $d : X \times X \rightarrow [0, 1]$ is a (1-bounded) premetric on X if $d(x, x) = 0$ for all $x \in X$. A premetric d is further a (1-bounded) pseudometric on X if for all $x, y, z \in X$, $d(x, y) = d(y, x)$ (symmetry) and $d(x, z) \leq d(x, y) + d(y, z)$ (triangle inequality). We denote the set of premetrics (resp. pseudometrics) on X by $\mathbf{M}_{\text{pr}}(X)$ (resp. $\mathbf{M}_{\text{ps}}(X)$).*

Remark 6.1. *In Definition 6.1, we use 1 to signal the largest distance between pairs of elements.*

In this chapter, we use ‘op’ to indicate either ‘pr’ or ‘ps’. In this way we can use “ $\mathbf{M}_{\text{op}}(X)$ ” to make arguments shared by both $\mathbf{M}_{\text{pr}}(X)$ and $\mathbf{M}_{\text{ps}}(X)$. Below we define a partial order on $\mathbf{M}_{\text{op}}(X)$ in a pointwise fashion.

Definition 6.2. *Let X be a non-empty set. The binary relation $\leq_{X, \text{op}}$ on $\mathbf{M}_{\text{op}}(X)$ is defined as follows: given any $d_1, d_2 \in \mathbf{M}_{\text{op}}(X)$, $d_1 \leq_{X, \text{op}} d_2$ iff $d_1(x, y) \leq d_2(x, y)$ for all $x, y \in X$.*

Clearly, $\leq_{X, \text{op}}$ is a partial order on $\mathbf{M}_{\text{op}}(X)$. The following proposition shows that $(\mathbf{M}_{\text{op}}(X), \leq_{X, \text{op}})$ is a complete lattice. Later on, the notion of bisimilarity metric will be defined as the least fixed-point of certain monotone function under the complete lattice of all pseudometrics on the state space of a PA.

Proposition 6.1 ([60]). *Let X be a non-empty set. Then $(\mathbf{M}_{\text{op}}(X), \leq_{X, \text{op}})$ is a complete lattice.*

Proof. The top element \top_{op} is determined by: $\top_{\text{op}}(x, x) = 0$ and $\top_{\text{op}}(x, y) = 1$ if $x \neq y$, for all $x, y \in X$. The bottom element \perp_{op} is defined by: $\perp_{\text{op}}(x, y) = 0$ for all $x, y \in X$. Given non-empty $Y \subseteq \mathbf{M}_{\text{op}}(X)$, the least upper-bound $\bigsqcup Y$ is given by: $(\bigsqcup Y)(x, y) := \sup\{d(x, y) \mid d \in Y\}$ for arbitrary $x, y \in X$. Note that $\bigsqcup Y \in \mathbf{M}_{\text{op}}(X)$ when $\text{op} = \text{ps}$, i.e.,

$$\sup\{d(x, z) \mid d \in Y\} \leq \sup\{d(x, y) \mid d \in Y\} + \sup\{d(y, z) \mid d \in Y\}$$

for all $x, y, z \in X$; this can be obtained by taking supremum over $d \in Y$ at the both sides of the inequality $d(x, z) \leq d(x, y) + d(y, z)$. The greatest lower bound $\prod Y$ (for non-empty $Y \subseteq M_{\text{op}}(X)$) is given by: $\prod Y = \bigsqcup \{d \in M_{\text{op}}(X) \mid \forall d' \in Y. d \leq_{X, \text{op}} d'\}$. \square

Now we define the bisimilarity metric on fPAs. Below we fix an fPA $\mathcal{M} = (S, \text{Act}, \rightarrow)$. We focus on the complete lattice $(M_{\text{op}}(S), \leq_{S, \text{op}})$. We will omit ‘ S ’ in “ $(M_{\text{op}}(S), \leq_{S, \text{op}})$ ” if the underlying context is clear.

To define the bisimilarity metric, we first lift an element $d \in M_{\text{op}}$ to an element $\bar{d} \in M_{\text{op}}(\text{Dist}(S))$.

Definition 6.3. *Let $d \in M_{\text{op}}$. The element $\bar{d} \in M_{\text{op}}(\text{Dist}(S))$ is defined as follows. Given any $\mu, \nu \in \text{Dist}(S)$, $\bar{d}(\mu, \nu)$ is defined as the optimal value of the following linear program $\text{LP}[d](\mu, \nu)$.*

- $\min \sum_{u, v \in S} d(u, v) \cdot z_{u, v}$ subject to:
- $\sum_{v \in S} z_{u, v} = \mu(u)$ for all $u \in S$;
 - $\sum_{u \in S} z_{u, v} = \nu(v)$ for all $v \in S$;
 - $z_{u, v} \geq 0$ for all $u, v \in S$.

We denote by $\text{FeS}[d](\mu, \nu)$ and resp. $\text{OpS}[d](\mu, \nu)$ the set of feasible solutions and resp. optimum solutions of the linear program $\text{LP}[d](\mu, \nu)$.

Intuitively, $\bar{d}(\mu, \nu)$ equals the minimum cost of the following transshipment problem. The sources (origins) are elements of S and the destinations are elements of an identical copy of S . The amount (probability mass) available at source u equals $\mu(u)$ and the amount needed at destination v equals $\nu(v)$. Probability mass can be moved from any source to any destination. The cost to move mass z from u to v equals $d(u, v) \cdot z$, where $d(u, v)$ is the “unit” transshipment cost.

The feasible region of the linear program specified in Definition 6.3 is not empty, as one can set $z_{u, v} = \mu(u) \cdot \nu(v)$ for $u, v \in S$, which satisfies the linear constraints. For $d \in M_{\text{ps}}$, it is also clear from symmetry that $\bar{d}(\mu, \nu) = \bar{d}(\nu, \mu)$ for $\mu, \nu \in \text{Dist}(S)$. An issue in Definition 6.3 is that \bar{d} may not satisfy the triangle inequality. The following proposition tackles this problem.

Proposition 6.2. *For all $d \in M_{\text{ps}}$, $\bar{d} \in M_{\text{ps}}(\text{Dist}(S))$.*

Proof. The proof follows the lines of [29, Proposition 3.5.6]. Let $d \in M_{\text{ps}}$ and $\mu, \mu', \nu \in \text{Dist}(S)$. It is clear that $\bar{d}(\mu, \mu) = 0$ and $\bar{d}(\mu, \nu) = \bar{d}(\nu, \mu)$ (from symmetry). We prove that $\bar{d}(\mu, \nu) \leq \bar{d}(\mu, \mu') + \bar{d}(\mu', \nu)$. Let $\{z_{u, v}^1\}_{u, v \in S}$ (resp. $\{z_{u, v}^2\}_{u, v \in S}$) be an optimum solution of $\text{LP}[d](\mu, \mu')$ (resp. $\text{LP}[d](\mu', \nu)$). Define $z_{u, v} := \sum_{s \in [\mu']} \frac{z_{u, s}^1 \cdot z_{s, v}^2}{\mu'(s)}$ for $u, v \in S$. We show that $\{z_{u, v}\}_{u, v \in S}$ is a

feasible solution of $\text{LP}[d](\mu, \nu)$. For $u \in S$,

$$\begin{aligned}
& \sum_{v \in S} z_{u,v} \\
&= \sum_{v \in S} \sum_{s \in [\mu']} \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \\
&= \sum_{s \in [\mu']} \sum_{v \in S} \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \\
&= \sum_{s \in S} z_{u,s}^1 = \mu(u) .
\end{aligned}$$

Similarly, for $v \in S$,

$$\begin{aligned}
& \sum_{u \in S} z_{u,v} \\
&= \sum_{u \in S} \sum_{s \in [\mu']} \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \\
&= \sum_{s \in [\mu']} \sum_{u \in S} \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \\
&= \sum_{s \in S} z_{s,v}^2 = \nu(v) .
\end{aligned}$$

Thus, $\{z_{u,v}\}_{u,v \in S}$ is a feasible solution of $\text{LP}[d](\mu, \nu)$. It follows that

$$\begin{aligned}
\bar{d}(\mu, \nu) &\leq \sum_{u,v \in S} d(u, v) \cdot z_{u,v} \\
&= \sum_{u,v \in S} \sum_{s \in [\mu']} d(u, v) \cdot \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \\
&\leq \sum_{u,v \in S} \sum_{s \in [\mu']} (d(u, s) + d(s, v)) \cdot \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \\
&= \left(\sum_{u,v \in S} \sum_{s \in [\mu']} d(u, s) \cdot \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \right) + \left(\sum_{u,v \in S} \sum_{s \in [\mu']} d(s, v) \cdot \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \right) \\
&= \left(\sum_{u \in S} \sum_{s \in [\mu']} \sum_{v \in S} d(u, s) \cdot \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \right) + \left(\sum_{v \in S} \sum_{s \in [\mu']} \sum_{u \in S} d(s, v) \cdot \frac{z_{u,s}^1 \cdot z_{s,v}^2}{\mu'(s)} \right) \\
&= \left(\sum_{u \in S} \sum_{s \in [\mu']} d(u, s) \cdot z_{u,s}^1 \right) + \left(\sum_{v \in S} \sum_{s \in [\mu']} d(s, v) \cdot z_{s,v}^2 \right) \\
&= \bar{d}(\mu, \mu') + \bar{d}(\mu', \nu)
\end{aligned}$$

which implies the result. \square

Proposition 6.2 makes sure that Definition 6.3 is well-defined.

Remark 6.2. *The original version of Definition 6.3 goes through Kantorovich metric. In detail, for $d \in \mathbf{M}_{\text{ps}}$, $\bar{d}(\mu, \nu)$ is equivalently defined as the optimal value of the following linear program*

- $$\max \sum_{s \in S} (\mu(s) - \nu(s)) \cdot x_s \text{ subject to:}$$
- $|x_u - x_v| \leq d(u, v)$ for all $u, v \in S$;
 - $x_s \in [0, 1]$ for all $s \in S$.

It can be proved through dual linear programming that the two definitions coincide (when $\text{op} = \text{ps}$) (cf. [72, 60]).

Now we define the bisimilarity metric as the least fixed-point of certain metric transformer, as follows.

Definition 6.4. *The metric transformer $\mathcal{T}_{\mathcal{M}, \text{op}} : \mathbf{M}_{\text{op}} \rightarrow \mathbf{M}_{\text{op}}$ is defined as a monotone function for the complete lattice $(\mathbf{M}_{\text{op}}, \leq_{\text{op}})$, as follows:*

- $\mathcal{T}_{\mathcal{M}, \text{op}}(d)(u, v) := 1$ if $\text{Act}(u) \neq \text{Act}(v)$;
- $\mathcal{T}_{\mathcal{M}, \text{op}}(d)(u, v) := 0$ if $\text{Act}(u) = \text{Act}(v) = \emptyset$;
- otherwise,

$$\mathcal{T}_{\mathcal{M}, \text{op}}(d)(u, v) := \max_{a \in \text{Act}(u)} \left[\max \left\{ \max_{u \xrightarrow[a]{nc} \mu} \min_{v \xrightarrow[a]{nc} \nu} \bar{d}(\mu, \nu), \max_{v \xrightarrow[a]{nc} \nu} \min_{u \xrightarrow[a]{nc} \mu} \bar{d}(\mu, \nu) \right\} \right];$$

for all $d \in \mathbf{M}_{\text{op}}$ and $u, v \in S$. The bisimilarity metric $\mathfrak{d}_{\mathcal{M}, \text{op}}$ is defined as the least fixed-point of $\mathcal{T}_{\mathcal{M}, \text{op}}$.

Intuitively, $\mathcal{T}_{\mathcal{M}, \text{op}}(d)(u, v)$ measures the distance between u and v in terms of the distance caused by their next-step transitions. The distance caused by next-step transitions is derived in a style similar to probabilistic bisimulation. The well-defined-ness of Definition 6.4 is given by the following proposition.

Proposition 6.3. *$\mathcal{T}_{\mathcal{M}, \text{op}}$ is a monotone function for the complete lattice $(\mathbf{M}_{\text{op}}, \leq_{\text{op}})$.*

Proof. The monotonicity of $\mathcal{T}_{\mathcal{M}, \text{op}}$ is straightforward from Definition 6.4 and Definition 6.3. The nontrivial case is that $\mathcal{T}_{\mathcal{M}, \text{ps}}(d) \in \mathbf{M}_{\text{ps}}$ for all $d \in \mathbf{M}_{\text{ps}}$. Let $d \in \mathbf{M}_{\text{ps}}$. The symmetry of $\mathcal{T}_{\mathcal{M}, \text{ps}}(d)$ follows directly from the symmetry of \bar{d} . Below we prove the triangle inequality of $\mathcal{T}_{\mathcal{M}, \text{ps}}(d)$.

Let $u, u', v \in S$. If either $Act(u) \neq Act(u')$ or $Act(u') \neq Act(v)$ or $Act(u) = Act(u') = Act(v) = \emptyset$, then it is clear that

$$\mathcal{T}_{\mathcal{M},ps}(d)(u, v) \leq \mathcal{T}_{\mathcal{M},ps}(d)(u, u') + \mathcal{T}_{\mathcal{M},ps}(d)(u', v) .$$

Otherwise, $Act(u) = Act(u') = Act(v) \neq \emptyset$. By the definition of $\mathcal{T}_{\mathcal{M},ps}$, there exists $a \in Act$ such that either there exists $u \xrightarrow{a}_{nc} \mu$ such that $\mathcal{T}_{\mathcal{M},ps}(d)(u, v) \leq \bar{d}(\mu, \nu)$ for all $v \xrightarrow{a}_{nc} \nu$, or dually there exists $v \xrightarrow{a}_{nc} \nu$ such that $\mathcal{T}_{\mathcal{M},ps}(d)(u, v) \leq \bar{d}(\mu, \nu)$ for all $u \xrightarrow{a}_{nc} \mu$. Without loss of generality, we assume the former case, i.e., there exists $u \xrightarrow{a}_{nc} \mu$ such that $\mathcal{T}_{\mathcal{M},ps}(d)(u, v) \leq \bar{d}(\mu, \nu)$ for all $v \xrightarrow{a}_{nc} \nu$. From the definition of $\mathcal{T}_{\mathcal{M},ps}$, there exists $u' \xrightarrow{a}_{nc} \mu'$ such that $\bar{d}(\mu, \mu') \leq \mathcal{T}_{\mathcal{M},ps}(d)(u, u')$. Again, there exists $v \xrightarrow{a}_{nc} \nu$ such that $\bar{d}(\mu', \nu) \leq \mathcal{T}_{\mathcal{M},ps}(d)(u', v)$. By Proposition 6.2, $\bar{d}(\mu, \nu) \leq \bar{d}(\mu, \mu') + \bar{d}(\mu', \nu)$. Thus $\mathcal{T}_{\mathcal{M},ps}(d)(u, v) \leq \mathcal{T}_{\mathcal{M},ps}(d)(u, u') + \mathcal{T}_{\mathcal{M},ps}(d)(u', v)$. \square

The bisimilarity metric $\mathfrak{d}_{\mathcal{M},op}$ measures the distance between two states of S . Below we show that pairs of states with distance zero are exactly probabilistic bisimilar states.

Proposition 6.4. *For all $u, v \in S$, $\mathfrak{d}_{\mathcal{M},ps}(u, v) = 0$ iff $u \sim_{nc} v$.*

Proof. Assume that $\mathfrak{d}_{\mathcal{M},ps}(u, v) = 0$. Define

$$\mathcal{R} := \{(u', v') \in S \times S \mid \mathfrak{d}_{\mathcal{M},ps}(u', v') = 0\} .$$

We prove that \mathcal{R} is an nc-bisimulation. Let $(u', v') \in \mathcal{R}$ and $u' \xrightarrow{a}_{nc} \mu'$. By $\mathfrak{d}_{\mathcal{M},ps}(u', v') = 0$, there exists $v' \xrightarrow{a}_{nc} \nu'$ such that $\bar{\mathfrak{d}}_{\mathcal{M},ps}(\mu', \nu') = 0$. Let $\{z_{u'', v''}\}_{u'', v'' \in S}$ be an optimum solution of $LP[\mathfrak{d}_{\mathcal{M},ps}](\mu', \nu')$. By

$$\sum_{u'', v'' \in S} \mathfrak{d}_{\mathcal{M},ps}(u'', v'') \cdot z_{u'', v''} = 0,$$

$z_{u'', v''} = 0$ whenever $\mathfrak{d}_{\mathcal{M},ps}(u'', v'') > 0$. Then w , which is defined such that $w(u'', v'') = z_{u'', v''}$ for all $u'', v'' \in S$, is a weight function for $\mu' \mathcal{R} \nu'$. Symmetrically, we can obtain a similar reasoning for the other direction $v' \xrightarrow{a}_{nc} \nu'$ of the bisimulation conditions. It follows that $u \sim_{nc} v$.

Assume now that $u \sim_{nc} v$. Define d by: $d(u', v') = 0$ if $u' \sim_{nc} v'$, and $d(u', v') = 1$ otherwise, for all $u', v' \in S$. Consider $\mathcal{T}_{\mathcal{M},ps}(d)$. Fix arbitrarily $u', v' \in S$. We clarify two cases.

1. $u' \not\sim_{nc} v'$. Then $\mathcal{T}_{\mathcal{M},ps}(d)(u', v') \leq d(u', v') = 1$.




Figure 6.1: An fPA

2. $u' \sim_{\text{nc}} v'$. Then for all $u' \xrightarrow[\text{nc}]{a} \mu'$, there exists $v' \xrightarrow[\text{nc}]{a} \nu'$ such that $\mu' \sim_{\text{nc}} \nu'$; also for all $v' \xrightarrow[\text{nc}]{a} \nu'$, there exists $u' \xrightarrow[\text{nc}]{a} \mu'$ such that $\mu' \sim_{\text{nc}} \nu'$.

Note that $\mu'' \sim_{\text{nc}} \nu''$ iff $\bar{d}(\mu'', \nu'') = 0$, since a weight function w that witnesses $\mu'' \sim_{\text{nc}} \nu''$ is essentially an optimum solution $\{z_{u'', v''}\}_{u'', v'' \in S}$ of the linear program $\text{LP}[d](\mu'', \nu'')$ (with $w(u'', v'') = z_{u'', v''}$ for all $u'', v'' \in S$) and vice versa. Thus by definition, $\mathcal{T}_{\mathcal{M}, \text{ps}}(d)(u', v') = d(u', v') = 0$.

In either case, we have $\mathcal{T}_{\mathcal{M}, \text{ps}}(d)(u', v') \leq d(u', v')$. Hence, $\mathcal{T}_{\mathcal{M}, \text{ps}}(d) \leq_{\mathcal{M}, \text{ps}} d$ by the arbitrary choice of u', v' . Then by Theorem 2.1, we obtain $\mathfrak{d}_{\mathcal{M}, \text{ps}} \leq d$. It follows that $\mathfrak{d}_{\mathcal{M}, \text{ps}}(u, v) = 0$. \square

Example 6.1. Consider the fPA $\mathcal{M} = (S, \text{Act}, \rightarrow)$ with $S = \{s_0, s_1, s_2, s_3\}$, $\text{Act} = \{a, b, c\}$ (a, b, c are three different actions) and

$$\begin{aligned} \rightarrow = & \{(s_0, b, \{s_0 \mapsto 1\}), (s_1, c, \{s_1 \mapsto 1\})\} \cup \\ & (s_2, a, \{s_0 \mapsto \frac{1}{2}, s_1 \mapsto \frac{1}{2}\}), (s_3, a, \{s_0 \mapsto \frac{1}{2} + \epsilon, s_1 \mapsto \frac{1}{2} - \epsilon\}) \} . \end{aligned}$$

The fPA is depicted in Fig. 6.1. By definition, $\mathfrak{d}_{\mathcal{M}, \text{ps}}(s_2, s_3) = \epsilon$.

When extending the bisimilarity metric to stochastic games, one can also have an equivalent logical characterization for the bisimilarity metric. We refer to [28] for details.

In this chapter, we study the complexity of the following decision problem **BISIMMETRIC**:

- **INPUT:** two states s, s' of an fPA \mathcal{M} and a rational number $c \in [0, 1]$;
- **OUTPUT:** whether $\mathfrak{d}_{\mathcal{M}, \text{ps}}(s, s') \leq c$.

6.2 Approximate Bisimilarity Metrics

In this section, we define approximants of the bisimilarity metric. The main purpose to introduce such notion is to show that the bisimilarity metric can be defined on the lattice of premetrics instead of that of pseudometrics. This fact will be used in Theorem 6.2.

Below we fix an fPA $\mathcal{M} = (S, Act, \rightarrow)$. To ease the notation, we will omit the subscript ‘ \mathcal{M} ’ whenever possible. The following definition illustrates the approximants $\{\mathfrak{d}_{\text{op}}^n\}_{n \in \mathbb{N}_0}$ of \mathfrak{d}_{op} .

Definition 6.5. *The family $\{\mathfrak{d}_{\text{op}}^n\}_{n \in \mathbb{N}_0}$ of approximants of \mathfrak{d}_{op} is inductively defined as follows.*

- $\mathfrak{d}_{\text{op}}^0 \in \mathbf{M}_{\text{op}}$ is given by: $\mathfrak{d}_{\text{op}}^0 := \perp_{\text{op}}$ (i.e., $\mathfrak{d}_{\text{op}}^0(u, v) = 0$ for all $u, v \in S$).
- $\mathfrak{d}_{\text{op}}^{n+1} \in \mathbf{M}_{\text{op}}$ is given by: $\mathfrak{d}_{\text{op}}^{n+1} = \mathcal{T}_{\text{op}}(\mathfrak{d}_{\text{op}}^n)$.

By Proposition 6.3, $\{\mathfrak{d}_{\text{op}}^n\}_{n \in \mathbb{N}_0}$ is well-defined. Note that although the definition of $\{\mathfrak{d}_{\text{op}}^n\}_{n \in \mathbb{N}_0}$ depends on whether $\text{op} = \text{pr}$ or $\text{op} = \text{ps}$, $\{\mathfrak{d}_{\text{pr}}^n\}_{n \in \mathbb{N}_0} = \{\mathfrak{d}_{\text{ps}}^n\}_{n \in \mathbb{N}_0}$ since $\perp_{\text{pr}} = \perp_{\text{ps}}$ and \mathcal{T}_{op} does not essentially differs apart between $\{\mathcal{T}_{\text{pr}}, \mathcal{T}_{\text{ps}}\}$ (cf. Definition 6.4).

Since \mathcal{T}_{op} is a monotone function, the infinite sequence $\{\mathfrak{d}_{\text{op}}^n\}_{n \in \mathbb{N}_0}$ is increasing w.r.t \leq_{op} since $\mathfrak{d}_{\text{op}}^0 \leq_{\text{op}} \mathfrak{d}_{\text{op}}^1$. Proposition 6.5 will show that \mathfrak{d}_{op} is the limit of this sequence. Before proving Proposition 6.5, we first prove an auxiliary lemma.

Lemma 6.1. *Let $\mu, \nu \in \text{Dist}(S)$. Let $d \in \mathbf{M}_{\text{pr}}$ and $\{d_n\}_{n \in \mathbb{N}_0}$ be a sequence with $d_n \in \mathbf{M}_{\text{pr}}$ for all n . If $\lim_{n \rightarrow \infty} d_n = d$, then $\lim_{n \rightarrow \infty} \bar{d}_n(\mu, \nu) = \bar{d}(\mu, \nu)$.*

Proof. Let X be the (finite) set of vertices of the linear program $\text{LP}[d](\mu, \nu)$ (or $\text{LP}[d_n](\mu, \nu)$), which is independent of d and $\{d_n\}_{n \in \mathbb{N}_0}$. Then for all $d' \in \mathbf{M}_{\text{pr}}$,

$$\bar{d}'(\mu, \nu) = \min_{z \in X} \sum_{u, v \in S} d'(u, v) \cdot z_{u, v} ,$$

which is a continuous function w.r.t d' (viewed as a vector on $S \times S$). It follows that $\lim_{n \rightarrow \infty} \bar{d}_n(\mu, \nu) = \bar{d}(\mu, \nu)$ if $\lim_{n \rightarrow \infty} d_n = d$. \square

Proposition 6.5. *For all $u, v \in S$, $\mathfrak{d}_{\text{op}}(u, v) = \lim_{n \rightarrow \infty} \mathfrak{d}_{\text{op}}^n(u, v)$.*

Proof. Let $\mathfrak{d}_{\text{op}}^\infty$ be given by: $\mathfrak{d}_{\text{op}}^\infty(u, v) := \lim_{n \rightarrow \infty} \mathfrak{d}_{\text{op}}^n(u, v)$ for all $u, v \in S$. We prove that $\mathfrak{d}_{\text{op}}^\infty = \mathfrak{d}_{\text{op}}$ through the following two directions.

$\mathfrak{d}_{\text{op}}^\infty \leq_{\text{op}} \mathfrak{d}_{\text{op}}$. By $\mathfrak{d}_{\text{op}}^0 \leq_{\text{op}} \mathfrak{d}_{\text{op}}$ and the monotonicity of \mathcal{T}_{op} , we can prove by induction on n that $\mathfrak{d}_{\text{op}}^n \leq_{\text{op}} \mathfrak{d}_{\text{op}}$ for all $n \in \mathbb{N}_0$. Thus $\mathfrak{d}_{\text{op}}^\infty \leq_{\text{op}} \mathfrak{d}_{\text{op}}$.

$\mathfrak{d}_{\text{op}} \leq_{\text{op}} \mathfrak{d}_{\text{op}}^\infty$. We prove that $\mathcal{T}_{\text{op}}(\mathfrak{d}_{\text{op}}^\infty) \leq_{\text{op}} \mathfrak{d}_{\text{op}}^\infty$. Fix an arbitrary pair $(u, v) \in S \times S$. The situation is clear when $Act(u) = Act(v) = \emptyset$ (in this case $\mathcal{T}_{\text{op}}(\mathfrak{d}_{\text{op}}^\infty)(u, v) = 0$) or $Act(u) \neq Act(v)$ (in this case $\mathfrak{d}_{\text{op}}^\infty(u, v) =$

$\mathcal{T}_{\text{op}}(\mathfrak{d}_{\text{op}}^1)(u, v) = 1$). Below we assume that $\text{Act}(u) = \text{Act}(v) \neq \emptyset$. Consider any $u \xrightarrow[\text{nc}]{a} \mu$. By definition, for all $n \in \mathbb{N}$, there exists $v \xrightarrow[\text{nc}]{a} \nu_n$ such that $\overline{\mathfrak{d}_{\text{op}}^n}(\mu, \nu_n) \leq \mathfrak{d}_{\text{op}}^{n+1}(u, v) \leq \mathfrak{d}_{\text{op}}^\infty(u, v)$. Since the set \rightarrow is finite, there exists $v \xrightarrow[\text{nc}]{a} \nu$ such that $\overline{\mathfrak{d}_{\text{op}}^n}(\mu, \nu) \leq \mathfrak{d}_{\text{op}}^{n+1}(u, v) \leq \mathfrak{d}_{\text{op}}^\infty(u, v)$, for infinitely many $n \in \mathbb{N}$. Then by the fact that $\{\mathfrak{d}_{\text{op}}^n\}_{n \in \mathbb{N}_0}$ is increasing w.r.t \leq_{op} , $\overline{\mathfrak{d}_{\text{op}}^n}(\mu, \nu) \leq \mathfrak{d}_{\text{op}}^\infty(u, v)$ for all $n \in \mathbb{N}_0$. Thus $\overline{\mathfrak{d}_{\text{op}}^\infty}(\mu, \nu) \leq \mathfrak{d}_{\text{op}}^\infty(u, v)$ by Lemma 6.1. Similar arguments can be applied to an arbitrary $v \xrightarrow[\text{nc}]{a} \nu$. Thus $\mathcal{T}_{\text{op}}(\mathfrak{d}_{\text{op}}^\infty) \leq_{\text{op}} \mathfrak{d}_{\text{op}}^\infty$, which implies $\mathfrak{d}_{\text{op}} \leq_{\text{op}} \mathfrak{d}_{\text{op}}^\infty$ from Tarski's Fixed-Point Theorem (Theorem 2.1). \square

Following Proposition 6.5 and the fact that $\mathfrak{d}_{\text{pr}}^0 = \mathfrak{d}_{\text{ps}}^0$, we obtain the main result of this section.

Corollary 6.1. $\mathfrak{d}_{\text{pr}} = \mathfrak{d}_{\text{ps}} = \prod \{d \in \mathbf{M}_{\text{pr}} \mid \mathcal{T}_{\text{pr}}(d) \leq_{\text{pr}} d\}$.

This corollary allows us to focus on the more manageable complete lattice $(\mathbf{M}_{\text{pr}}, \leq_{\text{pr}})$.

6.3 Self-Closed Sets

In this section, we introduce the notion of “self-closed” sets which is the key to prove the membership of $\text{NP} \cap \text{coNP}$ and $\text{UP} \cap \text{coUP}$ for our problem. Below we fix a PA $\mathcal{M} = (S, \text{Act}, \rightarrow)$. We will omit the subscript ‘ \mathcal{M} ’ whenever possible. By Corollary 6.1, we can solely focus on the complete lattice $(\mathbf{M}_{\text{pr}}, \leq_{\text{pr}})$.

The following definition illustrates the notion of “self-closed” sets.

Definition 6.6. Let $d \in \mathbf{M}_{\text{pr}}$ with $d = \mathcal{T}_{\text{pr}}(d)$. A subset $X \subseteq S \times S$ is self-closed w.r.t d iff for all $(u, v) \in X$, the following three conditions hold:

1. $d(u, v) > 0$ (which implies that $u \neq v$) and $\text{Act}(u) = \text{Act}(v)$;
2. for all $u \xrightarrow[\text{nc}]{a} \mu$ such that $d(u, v) = \min\{\overline{d}(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\}$, there exists $v \xrightarrow[\text{nc}]{a} \nu'$ and $z = \{z_{u,v}\}_{u,v \in S} \in \text{OpS}[d](\mu, \nu')$ such that $d(u, v) = \overline{d}(\mu, \nu')$ and $\lfloor z \rfloor \subseteq X$.
3. for all $v \xrightarrow[\text{nc}]{a} \nu$ such that $d(u, v) = \min\{\overline{d}(\mu, \nu) \mid u \xrightarrow[\text{nc}]{a} \mu\}$, there exists $u \xrightarrow[\text{nc}]{a} \mu'$ and $z = \{z_{u,v}\}_{u,v \in S} \in \text{OpS}[d](\mu', \nu)$ such that $d(u, v) = \overline{d}(\mu', \nu)$ and $\lfloor z \rfloor \subseteq X$.

The set $\lfloor z \rfloor$ is defined by: $\lfloor z \rfloor := \{(u, v) \in S \times S \mid z_{u,v} > 0\}$.

Intuitively, a self-closed set X w.r.t d is a set such that all values $\{d(u, v) \mid (u, v) \in X\}$ can be reached on X itself. Below we show that non-empty self-closed sets characterize exactly the least fixed-point \mathfrak{d}_{pr} of the metric transformer \mathcal{T}_{pr} .

Theorem 6.1. *Let $d \in M_{\text{pr}}$ such that $d = \mathcal{T}_{\text{pr}}(d)$. If $d \neq \mathfrak{d}_{\text{pr}}$, then there exists a non-empty self-closed set $X \subseteq S \times S$ w.r.t d .*

Proof. Assume $d \neq \mathfrak{d}_{\text{pr}}$. We construct a non-empty self-closed set X as described below. Define $\delta(u, v) := d(u, v) - \mathfrak{d}_{\text{pr}}(u, v)$ for all $u, v \in S$. It is clear that $\delta(u, v) \geq 0$ for all $u, v \in S$, and there exists $(u, v) \in S \times S$ such that $\delta(u, v) > 0$ (by $d \neq \mathfrak{d}_{\text{pr}}$). Define X to be the following set:

$$X := \{(u, v) \in S \times S \mid \delta(u, v) = \max\{\delta(u', v') \mid (u', v') \in S \times S\}\}$$

We prove that X is a non-empty self-closed set. The non-emptiness of X is obvious. We further prove that any $(u, v) \in X$ satisfies the three conditions specified in Definition 6.6. Fix an arbitrary $(u, v) \in X$.

1. It is clear that $d(u, v) > 0$ since $\delta(u, v) > 0$. We prove that $\text{Act}(u) = \text{Act}(v)$. Suppose $\text{Act}(u) \neq \text{Act}(v)$. Then by definition, $d(u, v) = \mathfrak{d}_{\text{pr}}(u, v) = 1$ which implies $\delta(u, v) = 0$. Contradiction. So (u, v) satisfies the first condition in Definition 6.6.
2. Assume that $u \xrightarrow[\text{nc}]{a} \mu$ satisfy $d(u, v) = \min\{\bar{d}(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\}$, and $v \xrightarrow[\text{nc}]{a} \nu'$ be arbitrarily chosen which satisfies

$$\bar{\mathfrak{d}}_{\text{pr}}(\mu, \nu') = \min\{\bar{\mathfrak{d}}_{\text{pr}}(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\} .$$

Choose an arbitrary optimum solution $z^* = \{z_{u,v}^*\}_{u,v \in S}$ which lies in $\text{OpS}[\bar{\mathfrak{d}}_{\text{pr}}](\mu, \nu')$. We prove that $d(u, v) = \bar{d}(\mu, \nu')$, $z^* \in \text{OpS}[d](\mu, \nu')$ and $[z^*] \subseteq X$ (cf. Definition 6.6). By the definition of X , we have $\delta(u', v') \leq \delta(u, v)$ for all $(u', v') \in S \times S$. Thus for all $z \in \text{FeS}[d](\mu, \nu') (= \text{FeS}[\bar{\mathfrak{d}}_{\text{pr}}](\mu, \nu'))$, we have

$$\sum_{u', v' \in S} \mathfrak{d}_{\text{pr}}(u', v') \cdot z_{u', v'} \geq \sum_{u', v' \in S} (d(u', v') - \delta(u, v)) \cdot z_{u', v'} . \quad (6.1)$$

Since $\sum_{u', v' \in S} z_{u', v'} = 1$, we can further simplify Inequality (6.1) as follows:

$$\sum_{u', v' \in S} \mathfrak{d}_{\text{pr}}(u', v') \cdot z_{u', v'} \geq \left(\sum_{u', v' \in S} d(u', v') \cdot z_{u', v'} \right) - \delta(u, v) \quad (6.2)$$

By taking the infimum at the both sides of Inequality 6.2, we obtain that $\bar{\mathfrak{d}}_{\text{pr}}(\mu, \nu') \geq \bar{d}(\mu, \nu') - \delta(u, v)$. Further from definition, we obtain $\mathfrak{d}_{\text{pr}}(u, v) \geq \bar{\mathfrak{d}}_{\text{pr}}(\mu, \nu')$. Thus, we have:

$$\mathfrak{d}_{\text{pr}}(u, v) \geq \bar{\mathfrak{d}}_{\text{pr}}(\mu, \nu') \geq \bar{d}(\mu, \nu') - \delta(u, v) \geq d(u, v) - \delta(u, v) = \mathfrak{d}_{\text{pr}}(u, v) .$$

This implies that $\mathfrak{d}_{\text{pr}}(u, v) = \overline{\mathfrak{d}_{\text{pr}}}(\mu, \nu')$ and $d(u, v) = \overline{d}(\mu, \nu')$. Then we can form another inequality series as follows:

$$\begin{aligned}
& \sum_{u', v' \in S} d(u', v') \cdot z_{u', v'}^* \\
\geq & d(u, v) && \text{(by } d(u, v) = \overline{d}(\mu, \nu')\text{)} \\
= & \mathfrak{d}_{\text{pr}}(u, v) + \delta(u, v) \\
= & \left(\sum_{u', v' \in S} \mathfrak{d}_{\text{pr}}(u', v') \cdot z_{u', v'}^* \right) + \delta(u, v) && \text{(by } \mathfrak{d}_{\text{pr}}(u, v) = \overline{\mathfrak{d}_{\text{pr}}}(\mu, \nu')\text{)} \\
= & \sum_{u', v' \in S} (\mathfrak{d}_{\text{pr}}(u', v') + \delta(u, v)) \cdot z_{u', v'}^* \\
\geq & \sum_{u', v' \in S} (\mathfrak{d}_{\text{pr}}(u', v') + \delta(u', v')) \cdot z_{u', v'}^* \\
= & \sum_{u', v' \in S} d(u', v') \cdot z_{u', v'}^* .
\end{aligned}$$

Thus it must be the case that $z^* \in \text{OpS}[d](\mu, \nu')$ and $\delta(u', v') = \delta(u, v)$ whenever $z_{u', v'}^* > 0$. Then we have $\lfloor z^* \rfloor \subseteq X$. So (u, v) satisfies the second condition in Definition 6.6.

3. Symmetrically, we can prove that (u, v) satisfies the third condition in Definition 6.6.

Hence in conclusion, X is a self-closed set w.r.t d . □

Theorem 6.2. *Let $d \in \mathbf{M}_{\text{pr}}$ such that $d = \mathcal{T}_{\text{pr}}(d)$. If there exists a non-empty self-closed set $X \subseteq S \times S$ with respect to d , then $d \neq \mathfrak{d}_{\text{pr}}$.*

Proof. Assume $X \subseteq S \times S$ be a non-empty self-closed set w.r.t d . We construct a premetric $d' \neq d$ such that $\mathcal{T}_{\text{pr}}(d') \leq_{\text{pr}} d'$ and $d' \leq_{\text{pr}} d$. For all pairs $(u \xrightarrow[\text{nc}]{a} \mu, v)$ and $(u, v \xrightarrow[\text{nc}]{a} \nu)$ with $(u, v) \in X$, we define the following difference values: (note that $\text{Act}(u) = \text{Act}(v)$ by Definition 6.6)

- $\delta[u \xrightarrow[\text{nc}]{a} \mu, v] := d(u, v) - \min\{\overline{d}(\mu, \nu') \mid v \xrightarrow[\text{nc}]{a} \nu'\}$;
- $\delta[u, v \xrightarrow[\text{nc}]{a} \nu] := d(u, v) - \min\{\overline{d}(\mu', \nu) \mid u \xrightarrow[\text{nc}]{a} \mu'\}$.

All the values above are non-negative since $d = \mathcal{T}_{\text{pr}}(d)$. Furthermore, we define the following two difference values: (where $\min \emptyset := 0$)

- $\delta_1 := \min\{\delta[u \xrightarrow[\text{nc}]{a} \mu, v] \mid (u, v) \in X, u \xrightarrow[\text{nc}]{a} \mu \text{ and } \delta[u \xrightarrow[\text{nc}]{a} \mu, v] > 0\}$;
- $\delta_2 := \min\{\delta[u, v \xrightarrow[\text{nc}]{a} \nu] \mid (u, v) \in X, v \xrightarrow[\text{nc}]{a} \nu \text{ and } \delta[u, v \xrightarrow[\text{nc}]{a} \nu] > 0\}$.

Finally, we define δ as follows:

$$\delta := \begin{cases} \min\{\delta_1, \delta_2, \min\{d(u, v) \mid (u, v) \in X\}\} & \text{if } \delta_1 \neq 0 \text{ and } \delta_2 \neq 0 \\ \min\{\delta_2, \min\{d(u, v) \mid (u, v) \in X\}\} & \text{if } \delta_1 = 0 \text{ and } \delta_2 \neq 0 \\ \min\{\delta_1, \min\{d(u, v) \mid (u, v) \in X\}\} & \text{if } \delta_1 \neq 0 \text{ and } \delta_2 = 0 \\ \min\{d(u, v) \mid (u, v) \in X\} & \text{if } \delta_1 = 0 \text{ and } \delta_2 = 0 \end{cases} .$$

Note that $\delta > 0$. Then we construct $d' \in M_{\text{pr}}$ by:

$$d'(u, v) := \begin{cases} d(u, v) - \frac{1}{2}\delta & \text{if } (u, v) \in X \\ d(u, v) & \text{if } (u, v) \notin X \end{cases}$$

for all $u, v \in S$. It is clear that $d' \neq d$ since X is non-empty. We prove that $\mathcal{T}_{\text{pr}}(d') \leq_{\text{pr}} d'$. Fix an arbitrary $(u, v) \in S \times S$. Assume that $(u, v) \notin X$. From $d' \leq_{\text{pr}} d$, we have $\mathcal{T}_{\text{pr}}(d') \leq_{\text{pr}} \mathcal{T}_{\text{pr}}(d)$, which implies

$$d'(u, v) = d(u, v) = \mathcal{T}_{\text{pr}}(d)(u, v) \geq \mathcal{T}_{\text{pr}}(d')(u, v) .$$

Thus $d'(u, v) \geq \mathcal{T}_{\text{pr}}(d')(u, v)$. Assume now that $(u, v) \in X$. For each $u \xrightarrow{\text{nc}} \mu$, we clarify two cases below:

Case 1: $\delta[u \xrightarrow{\text{nc}} \mu, v] > 0$. Then $\delta_1 > 0$. By definition, we have:

$$d'(u, v) \geq d(u, v) - \frac{1}{2} \cdot \delta > d(u, v) - \delta[u \xrightarrow{\text{nc}} \mu, v] = \min\{\bar{d}(\mu, \nu') \mid v \xrightarrow{\text{nc}} \nu'\} .$$

From $d' \leq_{\text{pr}} d$, we have

$$\min\{\bar{d}(\mu, \nu') \mid v \xrightarrow{\text{nc}} \nu'\} \geq \min\{\bar{d}'(\mu, \nu') \mid v \xrightarrow{\text{nc}} \nu'\} .$$

Thus $d'(u, v) \geq \min\{\bar{d}'(\mu, \nu') \mid v \xrightarrow{\text{nc}} \nu'\}$.

Case 2: $\delta[u \xrightarrow{\text{nc}} \mu, v] = 0$. Since X is self-closed, there exists $v \xrightarrow{\text{nc}} \nu'$ and $z \in \text{OpS}[d](\mu, \nu')$ such that $d(u, v) = \bar{d}(\mu, \nu')$ and $[z] \subseteq X$. From $[z] \subseteq X$, we obtain

$$\begin{aligned} & \sum_{u', v' \in S} d'(u', v') \cdot z_{u', v'} \\ &= \left(\sum_{u', v' \in S} d(u', v') \cdot z_{u', v'} \right) - \frac{1}{2} \cdot \delta \\ &= \bar{d}(\mu, \nu') - \frac{1}{2} \cdot \delta \\ &= d'(u, v) . \end{aligned}$$

Then we have

$$\min\{\bar{d}'(\mu, \nu'') \mid v \xrightarrow{\text{nc}} \nu''\} \leq \bar{d}'(\mu, \nu') \leq \sum_{u', v' \in S} d'(u', v') \cdot z_{u', v'} = d'(u, v) .$$

It follows that $\min\{\bar{d}'(\mu, \nu'') \mid v \xrightarrow{\text{nc}} \nu''\} \leq d'(u, v)$.

Thus $d'(u, v) \geq \min\{\bar{d}'(\mu, \nu') \mid v \xrightarrow{\text{nc}} \nu'\}$ for all $u \xrightarrow{\text{nc}} \mu$. Similarly, we can prove that $d'(u, v) \geq \min\{\bar{d}'(\mu', \nu) \mid u \xrightarrow{\text{nc}} \mu'\}$ for all $v \xrightarrow{\text{nc}} \nu$. Thus $d'(u, v) \geq \mathcal{T}_{\text{pr}}(d')(u, v)$.

Hence, $\mathcal{T}_{\text{pr}}(d') \leq_{\text{pr}} d'$. By $d' \leq_{\text{pr}} d$ and $d' \neq d$, we have

$$d \neq \bigcap \{d'' \in \mathcal{M}_{\text{pr}} \mid \mathcal{T}_{\text{pr}}(d'') \leq_{\text{pr}} d''\} .$$

It follows that $d \neq \mathfrak{d}_{\text{pr}}$. \square

Note that d' in the proof of Theorem 6.2 may not be a pseudometric. This is why Corollary 6.1 is needed.

Thus for all $d \in \mathcal{M}_{\text{pr}}$ with $d = \mathcal{T}_{\text{pr}}(d)$, $d \neq \mathfrak{d}_{\text{pr}}$ iff there exists a non-empty self-closed set w.r.t d . This characterization means that to check whether $d \neq \mathfrak{d}_{\text{pr}}$ or not, we can equivalently check whether there exists a non-empty self-closed set w.r.t d or not. The intuition here is that for any self-closed sets X, Y , $X \cup Y$ is still a self-closed set; thus there is a largest self-closed set. This gives rise to a partition-refinement algorithm that computes the largest self-closed set.

Theorem 6.3. *Define*

$$FP := \{d \in \mathcal{M}_{\text{pr}} \mid d = \mathcal{T}_{\text{pr}}(d) \text{ and all coordinates of } d \text{ are rational.}\}$$

to be the set of rational fixed-points of \mathcal{T}_{pr} . The problem whether a given $d \in FP$ equals \mathfrak{d}_{pr} is decidable in polynomial time in the size of d and \mathcal{M} .

Proof. From Theorem 6.2 and Theorem 6.1, we can check whether $d = \mathfrak{d}_{\text{pr}}$ or not by checking whether there exists a non-empty self-closed set w.r.t the given $d \in FP$. For all self-closed sets X, Y w.r.t d , $X \cup Y$ is still self-closed w.r.t d . So there exists a largest self-closed set w.r.t d , which we denote by Z . Then there exists a non-empty self-closed set w.r.t d iff Z is non-empty. Below we develop a partition-refinement algorithm to compute Z .

Firstly, we define a refining function $ref : \mathcal{X} \rightarrow \mathcal{X}$, where the set \mathcal{X} is given as follows:

$$\mathcal{X} := \{X \subseteq S \times S \mid Act(u) = Act(v) \text{ and } d(u, v) > 0 \text{ for all } (u, v) \in X\} .$$

Note that \mathcal{X} is non-empty since $\emptyset \in \mathcal{X}$. Given $X \in \mathcal{X}$, we define $\delta_X := \min\{d(u, v) \mid (u, v) \in X\}$ (where $\min \emptyset := 0$) and the premetric $d_X \in \mathcal{M}_{\text{pr}}$ as follows:

$$d_X(u, v) = \begin{cases} d(u, v) - \delta_X & \text{if } (u, v) \in X \\ d(u, v) & \text{if } (u, v) \notin X \end{cases} .$$

Then the set $ref(X) \in \mathcal{X}$ for $X \in \mathcal{X}$ is defined as follows: $(u, v) \in ref(X)$ iff $(u, v) \in X$ and furthermore (u, v) satisfies the following two conditions:

1. for all $u \xrightarrow[\text{nc}]{a} \mu$, if $d(u, v) = \min\{\bar{d}(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\}$ then

$$d_X(u, v) \geq \min\{\bar{d}_X(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\} ;$$

2. for all $v \xrightarrow[\text{nc}]{a} \nu$, if $d(u, v) = \min\{\bar{d}(\mu, \nu) \mid u \xrightarrow[\text{nc}]{a} \mu\}$ then

$$d_X(u, v) \geq \min\{\bar{d}_X(\mu, \nu) \mid u \xrightarrow[\text{nc}]{a} \mu\} .$$

Note that $\text{ref}(X) \subseteq X$ for all $X \in \mathcal{X}$.

Now we construct a sequence $\{Z_i\}_{i \in \mathbb{N}_0}$ as follows:

- $Z_0 := \{(u, v) \in S \times S \mid \text{Act}(u) = \text{Act}(v) \text{ and } d(u, v) > 0\}$;
- $Z_{n+1} := \text{ref}(Z_n)$.

By $Z_{n+1} \subseteq Z_n$, there exists $k \leq |Z_0|$ such that $Z_{k+1} = \text{ref}(Z_k) = Z_k$. We show that $Z_k = Z$.

“ $Z \subseteq Z_k$ ”: We prove by induction that $Z \subseteq Z_n$ for all $n \in \mathbb{N}_0$. The base step $Z \subseteq Z_0$ is clear from the definition. For the inductive step, assume that $Z \subseteq Z_n$. We show that $Z \subseteq Z_{n+1}(= \text{ref}(Z_n))$. Fix an arbitrary $(u, v) \in Z$. Consider any $u \xrightarrow[\text{nc}]{a} \mu$ such that $d(u, v) = \min\{\bar{d}(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\}$. Since Z is self-closed, there exists $v \xrightarrow[\text{nc}]{a} \nu'$ and $z \in \text{OpS}[d](\mu, \nu')$ such that $d(u, v) = \bar{d}(\mu, \nu')$ and $[z] \subseteq Z$. Since $Z \subseteq Z_n$, we have $d_{Z_n}(u, v) = d(u, v) - \delta_{Z_n}$ and $d_{Z_n}(u', v') = d(u', v') - \delta_{Z_n}$ for all $(u', v') \in [z]$. Thus we obtain

$$\begin{aligned} & \sum_{u', v' \in S} d_{Z_n}(u', v') \cdot z_{u', v'} \\ &= \sum_{u', v' \in S} (d(u', v') - \delta_{Z_n}) \cdot z_{u', v'} \\ &= d(u, v) - \delta_{Z_n} \\ &= d_{Z_n}(u, v) . \end{aligned}$$

Hence $d_{Z_n}(u, v) \geq \bar{d}_{Z_n}(\mu, \nu) \geq \min\{\bar{d}_{Z_n}(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\}$. By a similar reasoning, we can prove that for all $v \xrightarrow[\text{nc}]{a} \nu$, if $d(u, v) = \min\{\bar{d}(\mu, \nu) \mid u \xrightarrow[\text{nc}]{a} \mu\}$ then $d_{Z_n}(u, v) \geq \min\{\bar{d}_{Z_n}(\mu, \nu) \mid u \xrightarrow[\text{nc}]{a} \mu\}$. So $(u, v) \in Z_{n+1}$. Thus $Z \subseteq Z_{n+1}$.

“ $Z_k \subseteq Z$ ”: We prove that Z_k is a self-closed set w.r.t d , i.e., Z_k satisfies the three conditions specified in Definition 6.6. Without loss of generality, we can assume that $Z_k \neq \emptyset$. The first condition in Definition 6.6 is directly satisfied since $Z_k \subseteq Z_0$. As for the second condition, consider any $(u, v) \in Z_k$ and $u \xrightarrow[\text{nc}]{a} \mu$ which satisfies $d(u, v) = \min\{\bar{d}(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\}$. By $Z_k = \text{ref}(Z_k)$, $d_{Z_k}(u, v) \geq \min\{\bar{d}_{Z_k}(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\}$. Choose $v \xrightarrow[\text{nc}]{a} \nu'$ such that $\bar{d}_{Z_k}(\mu, \nu') = \min\{\bar{d}_{Z_k}(\mu, \nu) \mid v \xrightarrow[\text{nc}]{a} \nu\}$ and an arbitrary $z \in \text{OpS}[d_{Z_k}](\mu, \nu')$.

Since $d_{Z_k}(u, v) = d(u, v) - \delta_{Z_k}$, we have

$$\begin{aligned} d(u, v) &\geq \overline{d_{Z_k}}(\mu, \nu') + \delta_{Z_k} \\ &= \sum_{u', v' \in S} (d_{Z_k}(u', v') + \delta_{Z_k}) \cdot z_{u', v'} \\ &\geq \sum_{u', v' \in S} d(u', v') \cdot z_{u', v'} \\ &\geq \overline{d}(\mu, \nu') \\ &\geq d(u, v) . \end{aligned}$$

Thus $d(u, v) = \overline{d}(\mu, \nu')$, $z \in \text{OpS}[d](\mu, \nu')$ and $d_{Z_k}(u', v') = d(u', v') - \delta_{Z_k}$ for all $(u', v') \in [z]$. This implies that $[z] \subseteq Z_k$ since $\delta_{Z_k} > 0$. It follows that the second condition of Definition 6.6 is satisfied. The reasoning for the third condition can be carried out in the same way as for the second one.

Thus to compute Z , we need only to apply *ref* to Z_0 at most $|Z_0|$ times. Note that the computation of *ref* can be carried out in polynomial time since the optimal value of a linear program can be computed in polynomial time [66]. Hence Z is polynomial-time computable. It follows directly that whether a given $d \in FP$ equals \mathfrak{d}_{pr} is decidable in polynomial time. \square

6.4 The Membership of $\text{UP} \cap \text{coUP}$

In this section, we finish the proof for the membership of $\text{NP} \cap \text{coNP}$ and $\text{UP} \cap \text{coUP}$. Below we fix an fPA $\mathcal{M} = (S, \text{Act}, \rightarrow)$.

By Theorem 6.3, we can decide if a given element d in FP equals \mathfrak{d}_{pr} in polynomial time in the size of d and \mathcal{M} . This indicates a polynomial time verifier as follows: firstly guess a d in FP and then check whether d equals \mathfrak{d}_{pr} in polynomial time. A missing part in the argument for the verifier is to show that the size of d is polynomial in the size of \mathcal{M} . The following proposition tackles this point.

Proposition 6.6. \mathfrak{d}_{pr} is a rational vector and the size of \mathfrak{d}_{pr} is polynomial in the size of \mathcal{M} .

Proof. For each $\mu, \nu \in \text{Dist}(S)$ and $u, v \in S$, we define $\nu[\mu, v] \in \text{Dist}(S)$ and $\mu[\nu, v] \in \text{Dist}(S)$ as follows:

- $\nu[\mu, v] := \underset{v \xrightarrow[\text{nc}]{a} \nu'}{\text{argmin}} \overline{\mathfrak{d}_{\text{pr}}}(\mu, \nu') ;$
- $\mu[\nu, u] := \underset{u \xrightarrow[\text{nc}]{a} \mu'}{\text{argmin}} \overline{\mathfrak{d}_{\text{pr}}}(\mu', \nu) ;$

$\nu[\mu, v]$ and $\mu[\nu, u]$ are chosen to be an arbitrarily optimal one when ties upon the argmin occur. We further define the vectors $w[\mu, v] : S \times S \rightarrow [0, 1]$ and $w[\nu, u] : S \times S \rightarrow [0, 1]$ to be one of the optimum vertices of the linear programs $\text{LP}[\mathfrak{d}_{\text{pr}}](\mu, \nu[\mu, v])$ and $\text{LP}[\mathfrak{d}_{\text{pr}}](\mu[\nu, u], \nu)$, respectively. Again, choices are made to be an arbitrarily optimal one when ties occur. By

the fundamental property of linear programming (cf. [66]), the sizes of $w[\mu, v]$ and $w[\nu, u]$ is polynomial in the size of \mathcal{M} . We prove that $\{\mathfrak{d}_{\text{pr}}(u, v)\}_{u, v \in S}$ is the unique optimum solution of the following linear program on vector z :

- min $\sum_{u, v \in S} z_{u, v}$, subject to:
- $z_{u, u} = 0$ for all $u \in S$
 - $z_{u, v} \in [0, 1]$ for all $u, v \in S$;
 - $z_{u, v} = 1$ if $\text{Act}(u) \neq \text{Act}(v)$;
 - $z_{u, v} = 0$ if $u \neq v$ and $\text{Act}(u) = \text{Act}(v) = \emptyset$;
 - $z_{u, v} \geq \sum_{u', v' \in S} z_{u', v'} \cdot w[\mu, v](u', v')$ for all $(u, v) \in S \times S$ and $u \xrightarrow[\text{nc}]{a} \mu$ such that $u \neq v$ and $\text{Act}(u) = \text{Act}(v) \neq \emptyset$;
 - $z_{u, v} \geq \sum_{u', v' \in S} z_{u', v'} \cdot w[\nu, u](u', v')$ for all $(u, v) \in S \times S$ and $v \xrightarrow[\text{nc}]{a} \nu$ such that $u \neq v$ and $\text{Act}(u) = \text{Act}(v) \neq \emptyset$.

Clearly, the feasible region of the linear program above is non-empty since $\{\mathfrak{d}_{\text{pr}}(u, v)\}_{u, v \in S}$ (with interpretation $z_{u, v} = \mathfrak{d}_{\text{pr}}(u, v)$) is a feasible solution. Also, for all feasible solutions z , z (viewed as a premetric such that $z(u, v) = z_{u, v}$) is a pre-fixed-point of \mathcal{T}_{pr} by Definition 6.4. It follows that $\{\mathfrak{d}_{\text{pr}}(u, v)\}_{u, v \in S}$ is the unique optimum solution of the linear program above since $\mathfrak{d}_{\text{pr}} \leq_{\text{pr}} z$ for all feasible solutions z (Theorem 2.1). Thus \mathfrak{d}_{pr} is of size polynomial in the size of \mathcal{M} from [66]. \square

Following directly from Proposition 6.6, we obtain the main result of the chapter.

Theorem 6.4. *The problem BISIMMETRIC lies in $\text{UP} \cap \text{coUP}$.*

Proof. We describe a polynomial-time verifier simultaneously for the problem and the complement of the problem: the verifier simply guesses a $d \in FP$ using a polynomial number of bits and checks whether $d = \mathfrak{d}_{\text{pr}}$ or not through the algorithm described in Theorem 6.3; if the checking passes, then the verifier compares $d(u, v)$ with ϵ and output the answer which depends on whether we are focusing on the problem or the complement of the problem. It follows that the problem lies in $\text{UP} \cap \text{coUP}$. \square

6.5 Conclusion

In this chapter, we proved that the problem whether the bisimilarity metric [72] between two given states of a finite probabilistic automata is under a given threshold lies in $\text{UP} \cap \text{coUP}$, which significantly improves previous PSPACE upperbound [71]. We prove this by establishing the notion of “self-closed” sets, and then exploring the relationship between the bisimilarity metric and the notion of self-closed sets.

Chapter 7

Continuous-Time Markov Decision Processes

The class of *continuous-time Markov decision processes* (CTMDPs) [64, 63] is a stochastic model that incorporates both features from continuous-time Markov chains (CTMCs) [36] and discrete-time Markov decision processes (MDPs) [64]. Generally, continuous-time Markov processes are systems with a countable set of states, whose timed transitions between states are governed by negative-exponential delays and non-deterministic choices. A CTMDP extends a CTMC in the sense that it allows non-deterministic choices, and it extends an MDP in the sense that it incorporates negative exponential time-delays. Due to its modelling capability of real-time probabilistic behaviour and non-determinism, CTMDPs are widely used in dependability analysis and performance evaluation [6, 3].

In a CTMDP, non-determinism is resolved by *schedulers* [74, 58]. Informally, a scheduler resolves the non-deterministic choices depending on the finite trajectory of the CTMDP accumulated so far and possibly the sojourn time of the current state. A scheduler is assumed to be *measurable* so that it induces a well-defined probability space over the infinite trajectories of the underlying CTMDP. Measurable schedulers are further divided into *early schedulers* and *late schedulers* [58, 74]. A scheduler that makes the choice solely by the trajectory accumulated so far is called an *early scheduler*, while a scheduler that utilizes both the trajectory and the sojourn time (at the current state) is called a *late scheduler*. With schedulers, one can reason about quantitative information such as the maximal/minimal probability/expectation of certain property.

In this chapter, we introduce continuous-time Markov decision processes in a way that combines both early and late schedulers. The chapter is organized as follows. Section 7.1 introduce the basic definition of CTMDPs. Section 7.2 introduces the notion of paths and histories. Section 7.3 introduces measurable spaces on paths and histories. Section 7.4 introduces the

notion of schedulers and probabilities measures under schedulers. Section 7.5 demonstrates a general integral characterization. Section 7.6 concludes this chapter. Section 7.7 collects all the proofs for this chapter.

Given a positive real number $\lambda > 0$, let f_λ be the probability density function of the negative exponential distribution with rate λ , i.e.,

$$f_\lambda(t) := \begin{cases} \lambda \cdot e^{-\lambda \cdot t} & t \geq 0 \\ 0 & t < 0 \end{cases} .$$

7.1 The Model

Definition 7.1. A continuous-time Markov Decision Process (CTMDP) \mathcal{M} is a tuple $(S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$ where

- S is a finite non-empty set of states which is the disjoint union of S_{er} and S_{la} ;
- S_{er} (resp. S_{la}) is a finite set of early-schedulable states (resp. late-schedulable states);
- Act is a finite non-empty set of actions;
- $\mathbf{E}_{\text{er}} : S_{\text{er}} \times \text{Act} \rightarrow \mathbb{R}_{\geq 0}$ (resp. $\mathbf{E}_{\text{la}} : S_{\text{la}} \rightarrow \mathbb{R}_{> 0}$) is the early total exit-rate function (resp. late total exit-rate function);
- $\mathbf{P} : S \times \text{Act} \times S \rightarrow [0, 1]$ is the discrete probability matrix such that for all $s \in S$ and $a \in \text{Act}$, $\sum_{s' \in S} \mathbf{P}(s, a, s') \in \{0, 1\}$.

An action $a \in \text{Act}$ is enabled at a state $s \in S$ if (i) $\sum_{s' \in S} \mathbf{P}(s, a, s') = 1$, and (ii) either $s \in S_{\text{er}}$ and $\mathbf{E}_{\text{er}}(s, a) > 0$, or $s \in S_{\text{la}}$; the set of enabled actions at a state $s \in S$ is denoted by $\text{En}(s)$.

Let $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$ be a CTMDP. Intuitively, $\mathbf{E}_{\text{er}}(s, a)$ is the total exit-rate of an early state s when an action $a \in \text{En}(s)$ is taken, while $\mathbf{E}_{\text{la}}(s)$ is the total exit-rate of a late-schedulable state s , regardless of which $a \in \text{En}(s)$ is taken.

By definition, we only consider finite-state CTMDPs in this dissertation. The actions reflect the non-deterministic feature of CTMDPs: each action represents a possible choice for the evolution of the CTMDP. When the action set is a singleton (i.e., it has only one element) and all states are late-schedulable, then the CTMDP is a *continuous-time Markov chain* (cf. Definition 8.1).

Often, a CTMDP is accompanied with an initial probability distribution which specifies the initial stochastic environment (for the CTMDP).

Definition 7.2. Let $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$ be a CTMDP. An initial distribution (for \mathcal{M}) is a function $\alpha : S \rightarrow [0, 1]$ which satisfies that $\sum_{s \in S} \alpha(s) = 1$.

In this chapter, we will use s, s' (resp. a, b) to range over states (resp. actions) of a CTMDP. In the whole thesis, we will assume that any CTMDP \mathcal{M} encountered will have the following *non-deadlock* property: for all states s of \mathcal{M} , $\text{En}(s)$ is non-empty. In general, states that do not conform to this property can be adjusted by leading them to a fresh new state s_\perp (a “deadlock” state which can either be early-schedulable or late-schedulable), whose exit-rate is an arbitrary positive real number and whose sole enabled action a_\perp leads to a Dirac distribution on s_\perp (i.e., $\mathbf{P}(s_\perp, a_\perp, s_\perp) = 1$).

Intuitively, the evolution of a CTMDP $(S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$ under a *scheduler* D (to be introduced in Section 7.4) and an initial distribution α is as follows.

- 1 At the beginning, an initial state s is chosen (as the current state) w.r.t the initial distribution α ;
- 2.1 On one hand, if $s \in S_{\text{er}}$, then the next state s' is determined as follows: (i) firstly an action $a \in \text{En}(s)$ is chosen by D , and then (ii) a sojourn-time (i.e., a non-negative real number) is triggered at s which is governed by the negative exponential distribution with rate $\mathbf{E}_{\text{er}}(s, a)$, and finally (iii) s' is chosen w.r.t the probability distribution $\mathbf{P}(s, a, \cdot)$.
- 2.2 On the other hand, if $s \in S_{\text{la}}$, then the next state s' is determined as follows: (i) firstly a sojourn-time is trigger at s which is governed by the negative exponential distribution with rate $\mathbf{E}_{\text{la}}(s)$, and then (ii) an action $a \in \text{En}(s)$ is chosen by D , and finally (iii) s' is chosen w.r.t the probability distribution $\mathbf{P}(s, a, \cdot)$.
- 3 The next current state is changed to s' . Then back to Step 2.1/Step 2.2 .

The evolution continues infinitely, and the resultant is an infinite path (or trajectory) (cf. Definition 7.3). The difference between early-schedulable and late-schedulable states is that the action to each early-schedulable state can only be chosen when the state is entered, while the action of each late-schedulable state can be chosen when the sojourn-time at the current state is over. As a result, a scheduler has more “choice” on a late-schedulable state than on an early-schedulable state.

Example 7.1. *A CTMDP with early-schedulable state s_0 , late-schedulable state s_1 is depicted in Fig. 7.1. a_1, a_2 are enabled actions at s_0 and b is the sole enabled action at s_1 . The numbers on the outgoing edges from each action specifies the probability distribution under the action.*

7.2 Paths and Histories

In this section, we introduce the notion of paths (or trajectories) and histories on a CTMDP. Intuitively, paths are resultants of an (infinite) evolution




Figure 7.1: A CTMDP

of a CTMDP, whereas histories are finite prefixes of paths. Below we fix a CTMDP $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$.

Definition 7.3 (Paths and Histories). *A (possibly infinite) path (or trajectory) π is an infinite sequence*

$$\pi = \left\langle s_0 \xrightarrow{a_0, t_0} s_1 \xrightarrow{a_1, t_1} s_2 \dots \right\rangle$$

such that $s_i \in S$, $t_i \in \mathbb{R}_{\geq 0}$ and $a_i \in \text{Act}$ for all $i \geq 0$; we denote s_i, t_i and a_i by $\pi[i]$, $\pi\langle i \rangle$ and $\pi(i)$, respectively. A (finite) history ξ is a finite sequence

$$\xi = \left\langle s_0 \xrightarrow{a_0, t_0} s_1 \xrightarrow{a_1, t_1} s_2 \dots s_m \right\rangle \quad (m \geq 0)$$

such that $s_i \in S$, $t_i \in \mathbb{R}_{\geq 0}$ and $a_i \in \text{Act}$ for all $0 \leq i \leq m-1$, and $s_m \in S$; we denote s_i, t_i, a_i and m by $\xi[i]$, $\xi\langle i \rangle$, $\xi(i)$ and $|\xi|$, respectively; moreover, we define $\xi \downarrow := \xi[|\xi|]$ to be the last state of the history ξ .

Below we introduce more notations on paths and histories. We denote the set of paths and histories (of \mathcal{M}) by $\text{Paths}(\mathcal{M})$ and $\text{Hists}(\mathcal{M})$, respectively. We define $\text{Hists}^n(\mathcal{M}) := \{\xi \in \text{Hists}(\mathcal{M}) \mid |\xi| = n\}$ to be the set of all histories of length n ($n \geq 0$). For each $n \in \mathbb{N}_0$ and $\pi \in \text{Paths}(\mathcal{M})$, we define the history $\pi[0..n]$ to be the finite prefix of π up to n ; Formally,

$$\pi[0..n] := \left\langle \pi[0] \xrightarrow{\pi(0), \pi\langle 0 \rangle} \dots \pi[n] \right\rangle.$$

Given $\pi \in \text{Paths}(\mathcal{M})$ and $(s, a, t) \in S \times \text{Act} \times \mathbb{R}_{\geq 0}$, we denote by $s \xrightarrow{a, t} \pi$ the path obtained by “putting” the prefix “ $s \xrightarrow{a, t}$ ” before π ; Formally,

$$s \xrightarrow{a, t} \pi := \left\langle s \xrightarrow{a, t} \pi[0] \xrightarrow{\pi(0), \pi\langle 0 \rangle} \pi[1] \xrightarrow{\pi(1), \pi\langle 1 \rangle} \dots \right\rangle.$$

Analogously, we define $s \xrightarrow{a,t} \xi$ (for $\xi \in \text{Hists}(\mathcal{M})$) to be the history obtained by “putting” “ $s \xrightarrow{a,t}$ ” before the history ξ .

Intuitively, a path π reflects a whole evolution of a CTMDP, where $\pi[i]$ is the current state at the i -th stage, $\pi(i)$ is the action chosen at $\pi[i]$ (by a scheduler) and $\pi\langle i \rangle$ is the sojourn-time on $\pi[i]$. On the other hand, a history ξ is a finite prefix of a path which reflects the evolution up to $|\xi|$ stages.

Below we extend sets of histories to sets of paths in a cylindrical fashion.

Definition 7.4. *Let $n \in \mathbb{N}_0$ and $\Xi \subseteq \text{Hists}^n(\mathcal{M})$. The cylinder extension of Ξ , denoted by $\text{Cyl}(\Xi)$, is defined as follows:*

$$\text{Cyl}(\Xi) := \{\pi \in \text{Paths}(\mathcal{M}) \mid \pi[0..n] \in \Xi\} .$$

Intuitively, the *cylinder extension* of $\Xi \subseteq \text{Hists}^n(\mathcal{M})$ is the set of all paths whose prefixes up to n steps are in Ξ .

7.3 Measurable Spaces on Paths and Histories

In this section, we define the measurable spaces for paths and histories on CTMDPs by combining the definitions from [74, 58]. Below we fix a CTMDP $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$.

Firstly, we introduce the notion of combined actions and its measurable space.

Definition 7.5 (Combined Actions). *A combined action is a tuple (a, t, s) where $a \in \text{Act}$, $t \in \mathbb{R}_{\geq 0}$ and $s \in S$. The measurable space $(\Gamma_{\mathcal{M}}, \mathcal{U}_{\mathcal{M}})$ over combined actions is defined as follows:*

- $\Gamma_{\mathcal{M}} := \text{Act} \times \mathbb{R}_{\geq 0} \times S$ is the set of combined actions;
- $\mathcal{U}_{\mathcal{M}} := 2^{\text{Act}} \otimes \mathcal{B}(\mathbb{R}_{\geq 0}) \otimes 2^S$ is the product σ -algebra (cf. Definition 3.11).

Then we introduce the notion of templates, which will be used to define the measurable spaces.

Definition 7.6 (Templates). *A template θ is a finite sequence*

$$\theta = \langle s, U_1, \dots, U_m \rangle (m \geq 0)$$

such that $s \in S$ and $U_i \in \mathcal{U}_{\mathcal{M}}$ for $1 \leq i \leq m$; The length of θ , denoted by $|\theta|$, is defined to be m . The set of histories $\text{Hists}(\theta)$ spanned by a template θ is defined by:

$$\text{Hists}(\langle s, U_1, \dots, U_m \rangle) := \left\{ \xi \in \text{Hists}^m(\mathcal{M}) \mid \xi[0] = s \text{ and } (\xi(i), \xi\langle i \rangle, \xi[i+1]) \in U_{i+1} \text{ for all } 0 \leq i < m \right\} .$$

Now we introduce the measurable spaces on paths and histories, as in the following definition.

Definition 7.7 (Measurable Spaces). *The measurable space $(\Omega_{\mathcal{M}}^n, \mathcal{S}_{\mathcal{M}}^n)$ over $Hists^n(\mathcal{M})$ ($n \in \mathbb{N}_0$) is defined as follows: $\Omega_{\mathcal{M}}^n = Hists^n(\mathcal{M})$ and $\mathcal{S}_{\mathcal{M}}^n$ is generated by the family*

$$\{\text{Hists}(\theta) \mid \theta \text{ is a template and } |\theta| = n\}$$

of subsets of $Hists^n(\mathcal{M})$.

The measurable space $(\Omega_{\mathcal{M}}, \mathcal{S}_{\mathcal{M}})$ over $Paths(\mathcal{M})$ is defined as follows: $\Omega_{\mathcal{M}} = Paths(\mathcal{M})$ and $\mathcal{S}_{\mathcal{M}}$ is the smallest σ -algebra generated by the family

$$\{\text{Cyl}(\Xi) \mid \Xi \in \mathcal{S}_{\mathcal{M}}^n \text{ for some } n \geq 0\} \quad (\S)$$

of subsets of $Paths(\mathcal{M})$.

Remark 7.1. *An alternative way to define the measurable space on paths can be done by changing (\S) to the following set:*

$$\mathcal{C} := \{\text{Cyl}(\text{Hists}(\theta)) \mid \theta \text{ is a template}\} .$$

This can be seen as follows. Let \mathcal{S}' be the σ -algebra on paths generated by \mathcal{C} . Clearly, $\mathcal{S}' \subseteq \mathcal{S}_{\mathcal{M}}$. For each $n \in \mathbb{N}_0$, define $\mathcal{S}'_n := \{\Xi \subseteq \Omega_{\mathcal{M}}^n \mid \text{Cyl}(\Xi) \in \mathcal{S}'\}$. One can verify that \mathcal{S}'_n is a σ -algebra on $\Omega_{\mathcal{M}}^n$ by the following facts:

1. $\Omega_{\mathcal{M}}^n \in \mathcal{S}'_n$;
2. If $\Xi \in \mathcal{S}'_n$ then $\Omega_{\mathcal{M}}^n - \Xi \in \mathcal{S}'_n$;
3. If $\Xi_1, \Xi_2, \dots \in \mathcal{S}'_n$ then $\bigcup_{m \geq 0} \Xi_m \in \mathcal{S}'_n$.

The second and third fact follows from $\text{Cyl}(\Omega_{\mathcal{M}}^n - \Xi) = \Omega_{\mathcal{M}} - \text{Cyl}(\Xi)$ and $\text{Cyl}(\bigcup_{m \geq 0} \Xi_m) = \bigcup_{m \geq 0} \text{Cyl}(\Xi_m)$, respectively. Then one obtains $\mathcal{S}_{\mathcal{M}}^n \subseteq \mathcal{S}'_n$ for all $n \geq 0$ since $\{\text{Hists}(\theta) \mid \theta \text{ is a template and } |\theta| = n\} \subseteq \mathcal{S}'_n$. This implies that $\text{Cyl}(\Xi) \in \mathcal{S}'$ for all $n \geq 0$ and $\Xi \in \mathcal{S}_{\mathcal{M}}^n$. It follows that $\mathcal{S}_{\mathcal{M}} \subseteq \mathcal{S}'$.

7.4 Schedulers and Their Probability Spaces

In this section, we introduce the notion of schedulers. Informally, a scheduler resolves the actions to be chosen (i.e., non-determinism) at each (current) state of a CTMDP, so that a unique probability space on the set of trajectories of the CTMDP can be established. As in the case of the previous section, we combine the notions stemming from [74, 58].

Below we fix a CTMDP $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, Act, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$. We distinguish histories through their last states by defining the following sets of histories:

- $Hists_{\text{er}}(\mathcal{M}) := \{\xi \in Hists(\mathcal{M}) \mid \xi \downarrow \in S_{\text{er}}\}$;

- $Hists_{\text{la}}(\mathcal{M}) := \{\xi \in Hists(\mathcal{M}) \mid \xi \downarrow \in S_{\text{la}}\}$.

It is easy to see that $Hists_{\text{er}}(\mathcal{M}) \cup Hists_{\text{la}}(\mathcal{M}) = Hists(\mathcal{M})$.

Definition 7.8 (Schedulers). *A scheduler D is a function*

$$D : (Hists_{\text{er}}(\mathcal{M}) \cup (Hists_{\text{la}}(\mathcal{M}) \times \mathbb{R}_{\geq 0})) \times Act \rightarrow [0, 1]$$

such that for all $\xi \in Hists(\mathcal{M})$ and $t \in \mathbb{R}_{\geq 0}$, the following conditions hold:

- either $\xi \in Hists_{\text{er}}(\mathcal{M})$ and $\sum_{a \in Act} D(\xi, a) = 1$, or $\xi \in Hists_{\text{la}}(\mathcal{M})$ and $\sum_{a \in Act} D(\xi, t, a) = 1$;
- for all $a \in Act$, if either (i) $\xi \in Hists_{\text{er}}(\mathcal{M})$ and $D(\xi, a) > 0$ or (ii) $\xi \in Hists_{\text{la}}(\mathcal{M})$ and $D(\xi, t, a) > 0$, then $a \in \text{En}(\xi \downarrow)$.

D is called measurable if for all $a \in Act$ and $n \geq 0$, the following conditions hold:

- the function $D(\cdot, a)$ is measurable w.r.t $(\Omega_{\mathcal{M}}^n, \mathcal{S}_{\mathcal{M}}^n)$, provided that the domain of $D(\cdot, a)$ is restricted to $Hists_{\text{er}}(\mathcal{M}) \cap Hists^n(\mathcal{M})$;
- the function $D(\cdot, \cdot, a)$ is measurable w.r.t $(\Omega_{\mathcal{M}}^n \times \mathbb{R}_{\geq 0}, \mathcal{S}_{\mathcal{M}}^n \otimes \mathcal{B}(\mathbb{R}_{\geq 0}))$, provided that the domain of $D(\cdot, \cdot, a)$ is restricted to $(Hists_{\text{la}}(\mathcal{M}) \cap Hists^n(\mathcal{M})) \times \mathbb{R}_{\geq 0}$.

From the definition, we can see that the difference between early- and late-schedulable states. In early-schedulable states, a scheduler chooses a probability distribution on actions immediately on entering a new current state; in late-schedulable states, a scheduler has the option to choose such a probability distribution after the sojourn-time at the new current state is over (i.e., the state is about to be left). Intuitively, a scheduler generally has more “free space” at late-schedulable states than at early-schedulable states. This allows a scheduler to better optimize certain property, e.g., maximal time-bounded reachability probability (cf. [59]). The measurability condition (in the definition) will be needed to define a probability measure for the measurable space $(\Omega_{\mathcal{M}}, \mathcal{S}_{\mathcal{M}})$.

Each measurable scheduler directly induces a probability measure on combined actions, when applied to a specific history. This probability measure serves as a basis for the definition of the probability measure on trajectories of the CTMDP.

Definition 7.9. *Let $\xi \in Hists(\mathcal{M})$ be a history and D a measurable scheduler. The probability measure $\mu_{\mathcal{M}}^D(\xi, \cdot)$ for the measurable space $(\Gamma_{\mathcal{M}}, \mathcal{U}_{\mathcal{M}})$ is defined as follows: if $\xi \in Hists_{\text{er}}(\mathcal{M})$, then*

$$\mu_{\mathcal{M}}^D(\xi, U) := \sum_{a \in \text{En}(\xi \downarrow)} D(\xi, a) \cdot \int_{\mathbb{R}_{\geq 0}} f_{\mathbf{E}_{\text{er}}(\xi \downarrow, a)}(t) \cdot \left[\sum_{s \in S} \mathbf{1}_U(a, t, s) \cdot \mathbf{P}(\xi \downarrow, a, s) \right] dt$$

for each $U \in \mathcal{U}_{\mathcal{M}}$; if $\xi \in \text{Hists}_{\text{la}}(\mathcal{M})$, then

$$\mu_{\mathcal{M}}^D(\xi, U) := \int_{\mathbb{R}_{\geq 0}} f_{\mathbf{E}_{\text{la}}(\xi \downarrow)}(t) \cdot \left\{ \sum_{a \in \text{En}(\xi \downarrow)} D(\xi, t, a) \cdot \left[\sum_{s \in S} \mathbf{1}_U(a, t, s) \cdot \mathbf{P}(\xi \downarrow, a, s) \right] \right\} dt$$

for each $U \in \mathcal{U}_{\mathcal{M}}$.

It can be shown that all integrand functions in Definition 7.9 are measurable (cf. [57]).

Based on Definition 7.9, we define the probability spaces on histories and paths (trajectories). Firstly, we define the probability space on histories. To this end, we introduce the notion of concatenation as follows.

Definition 7.10. *Let $\xi \in \text{Hists}(\mathcal{M})$ be a history and $(a, t, s) \in \Gamma_{\mathcal{M}}$ be a combined action. Define $\xi \circ (a, t, s) \in \text{Hists}(\mathcal{M})$ to be the history obtained by concatenating (a, t, s) to $\xi \downarrow$ (i.e. $\xi \circ (a, t, s) = \xi[0] \dots \xi \downarrow \xrightarrow{a, t} s$).*

Then the probability space on histories of fixed length is given as follows.

Definition 7.11. *Suppose D is a measurable scheduler and α is an initial distribution. The sequence $\{\text{Pr}_{\mathcal{M}, D, \alpha}^n : \mathcal{S}_{\mathcal{M}}^n \rightarrow [0, 1]\}_{n \geq 0}$ of probability measures is inductively as follows:*

$$\begin{aligned} \text{Pr}_{\mathcal{M}, D, \alpha}^0(\Xi) &:= \sum_{s \in \Xi} \alpha(s) ; \\ \text{Pr}_{\mathcal{M}, D, \alpha}^{n+1}(\Xi) &:= \int_{\Omega_{\mathcal{M}}^n} \left[\int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi}(\xi \circ \gamma) \mu_{\mathcal{M}}^D(\xi, d\gamma) \right] \text{Pr}_{\mathcal{M}, D, \alpha}^n(d\xi) . \end{aligned}$$

Again, it can be shown that all integrand functions in Definition 7.11 are measurable (cf. [57]). Then, the probability space on paths (trajectories) is given as follows.

Definition 7.12 (Probability Space on Paths). *Let D be a measurable scheduler and α be an initial distribution. The probability space*

$$(\Omega_{\mathcal{M}}, \mathcal{S}_{\mathcal{M}}, \text{Pr}_{\mathcal{M}, D, \alpha})$$

is defined as follows:

- $\Omega_{\mathcal{M}}$ and $\mathcal{S}_{\mathcal{M}}$ is defined as in Definition 7.7;
- $\text{Pr}_{\mathcal{M}, D, \alpha}$ is the unique probability measure such that

$$\text{Pr}_{\mathcal{M}, D, \alpha}(\text{Cyl}(\Xi)) = \text{Pr}_{\mathcal{M}, D, \alpha}^n(\Xi)$$

for all $n \geq 0$ and $\Xi \in \mathcal{S}_{\mathcal{M}}^n$.

We refer for the detailed construction of $(\Omega_{\mathcal{M}}, \mathcal{S}_{\mathcal{M}}, \text{Pr}_{\mathcal{M}, D, \alpha})$ to [58, 74, 57].

We end this section with a fundamental property asserting that the role of initial distribution can be decomposed into Dirac distributions on individual states.

Proposition 7.1. *For each measurable scheduler D and initial distribution α , $\text{Pr}_{\mathcal{M}, D, \alpha}(\Pi) = \sum_{s \in S} \alpha(s) \cdot \text{Pr}_{\mathcal{M}, D, \mathcal{D}[s]}(\Pi)$ for all $\Pi \in \mathcal{S}_{\mathcal{M}}$.*

Recall that $\mathcal{D}[s]$ is the Dirac distribution at s (cf. Definition 4.1). This proposition allows one to focus only on Dirac distributions when one wants to compute/approximate probability mass of certain measurable sets of trajectories.

7.5 A General Integral Characterization

In this section, we derive a general integral characterization for the probability measure on paths (trajectories). Below we fix a CTMDP $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$. For the sake of simplicity, we will omit all ‘ \mathcal{M} ’s which appear in the subscripts (of e.g., the notation ‘Pr’). The proofs for this section are highly measure-theoretic, and are put in Section 7.7.

Firstly, we define shifting functions on histories and paths which shifts each path/history by one transition step.

Definition 7.13. *Given $\Pi \in \mathcal{S}_{\mathcal{M}}$ and $(s, a) \in S \times \text{Act}$, we define the function $P_{\Pi}^{s, a} : \mathbb{R}_{\geq 0} \rightarrow 2^{\text{Paths}(\mathcal{M})}$ by:*

$$P_{\Pi}^{s, a}(t) := \{\pi \in \text{Paths}(\mathcal{M}) \mid s \xrightarrow{a, t} \pi \in \Pi\} .$$

Analogously, given $\Xi \in \mathcal{S}_{\mathcal{M}}^n$ with $n \geq 1$ and $(s, a) \in S \times \text{Act}$, we define $H_{\Xi}^{s, a} : \mathbb{R}_{\geq 0} \rightarrow 2^{\text{Hists}^{n-1}(\mathcal{M})}$ by:

$$H_{\Xi}^{s, a}(t) := \{\xi \in \text{Hists}^{n-1}(\mathcal{M}) \mid s \xrightarrow{a, t} \xi \in \Xi\} .$$

We also define the shifted version of a measurable scheduler as follows.

Definition 7.14. *Let $s \in S$, $a \in \text{Act}$ and $t \in \mathbb{R}_{\geq 0}$. For each measurable scheduler D , the scheduler $D[s \xrightarrow{a, t}]$ is defined as follows:*

- $D[s \xrightarrow{a, t}](\xi, \cdot) := D(s \xrightarrow{a, t} \xi, \cdot)$ for all $\xi \in \text{Hists}_{\text{er}}(\mathcal{M})$;
- $D[s \xrightarrow{a, t}](\xi, \tau, \cdot) := D(s \xrightarrow{a, t} \xi, \tau, \cdot)$ for all $(\xi, \tau) \in \text{Hists}_{\text{la}}(\mathcal{M}) \times \mathbb{R}_{\geq 0}$.

The following lemma states that each shifted set of paths/histories is measurable w.r.t corresponding measurable space.

Lemma 7.1. $P_{\Pi}^{s,a}(t) \in \mathcal{S}_{\mathcal{M}}$ for all $\Pi \in \mathcal{S}_{\mathcal{M}}$, $(s, a) \in S \times \text{Act}$ and $t \in \mathbb{R}_{\geq 0}$. Analogously, $H_{\Xi}^{s,a}(t) \in \mathcal{S}_{\mathcal{M}}^{n-1}$ for all $n \geq 1$, $\Xi \in \mathcal{S}_{\mathcal{M}}^n$, $(s, a) \in S \times \text{Act}$ and $t \in \mathbb{R}_{\geq 0}$.

Moreover, each shifted scheduler is measurable, as is illustrated by the following lemma.

Lemma 7.2. Let $s \in S$, $a \in \text{Act}$ and $t \in \mathbb{R}_{\geq 0}$. For each measurable scheduler D , $D[s \xrightarrow{a,t}]$ is a measurable scheduler.

Based on Lemma 7.1 and Lemma 7.2, we define a shift probability function as follows.

Definition 7.15. Let $s \in S$, $a \in \text{Act}$, $\Pi \in \mathcal{S}_{\mathcal{M}}$ and D a measurable scheduler. Define $p_{\Pi,D}^{s,a} : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ by:

$$p_{\Pi,D}^{s,a}(t) := \Pr_{D[s \xrightarrow{a,t}], \mathbf{P}(s,a,\cdot)}(P_{\Pi}^{s,a}(t))$$

for all $t \geq 0$.

The following proposition states that the shift probability function is measurable w.r.t $(\mathbb{R}_{\geq 0}, \mathcal{B}(\mathbb{R}_{\geq 0}))$.

Proposition 7.2. $p_{\Pi,D}^{s,a}$ is a measurable function w.r.t $(\mathbb{R}_{\geq 0}, \mathcal{B}(\mathbb{R}_{\geq 0}))$ given any $\Pi \in \mathcal{S}_{\mathcal{M}}$, $s \in S$, $a \in \text{Act}$ and measurable scheduler D .

Below we present the integral characterization for measurable schedulers, which is the main result of this section.

Theorem 7.1. Let D be a measurable scheduler. For each $\Pi \in \mathcal{S}_{\mathcal{M}}$ and $s \in S_{\text{er}}$, we have

$$\Pr_{D, \mathcal{D}[s]}(\Pi) = \sum_{a \in \text{En}(s)} D(s, a) \cdot \int_0^{\infty} f_{\mathbf{E}_{\text{er}}(s,a)}(t) \cdot p_{\Pi,D}^{s,a}(t) dt .$$

For each $\Pi \in \mathcal{S}_{\mathcal{M}}$ and $s \in S_{\text{la}}$, we have

$$\Pr_{D, \mathcal{D}[s]}(\Pi) = \int_0^{\infty} f_{\mathbf{E}_{\text{la}}(s)}(t) \cdot \left[\sum_{a \in \text{En}(s)} D(s, t, a) \cdot p_{\Pi,D}^{s,a}(t) \right] dt .$$

7.6 Conclusion

In this chapter, we introduced continuous-time Markov decision processes (CTMDPs) [64, 63] and several related notions from [74, 58]. These notions are namely paths, histories, schedulers and probability spaces on paths and histories. We illustrated them in a way that combines the setting of early [74] and late [58] schedulers. Moreover, as a contribution, we proved a general integral characterization for the probability space on paths of a CTMDP.

7.7 Proofs

In the following, we fix a CTMDP $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$.

Proposition 7.1. For each measurable scheduler D and each initial distribution α , $\Pr_{\mathcal{M}, D, \alpha}(\Pi) = \sum_{s \in S} \alpha(s) \cdot \Pr_{\mathcal{M}, D, \mathcal{D}[s]}(\Pi)$ for all $\Pi \in \mathcal{S}_{\mathcal{M}}$.

Proof. Define $\Pr'(\Pi) := \sum_{s \in S} \alpha(s) \cdot \Pr_{\mathcal{M}, D, \mathcal{D}[s]}(\Pi)$ for $\Pi \in \mathcal{S}_{\mathcal{M}}$. We prove that \Pr' coincides with $\Pr_{\mathcal{M}, D, \alpha}$. It is clear that \Pr' is a probability measure since each $\Pr_{\mathcal{M}, D, \mathcal{D}[s]}$ is a probability measure. So it suffices to prove that \Pr' coincides with $\Pr_{\mathcal{M}, D, \alpha}$ on $\bigcup_{n \geq 0} \{\text{Cyl}(\Xi) \mid \Xi \in \mathcal{S}_{\mathcal{M}}^n\}$. To this end, we proceed by induction on n .

Base Step: $n = 0$ and $\Xi \in \mathcal{S}_{\mathcal{M}}^n$. By definition, we have

$$\Pr_{\mathcal{M}, D, \alpha}(\text{Cyl}(\Xi)) = \Pr_{\mathcal{M}, D, \alpha}^0(\Xi) = \sum_{s \in S} \alpha(s) \cdot \mathbf{1}_{\Xi}(s) = \Pr'(\text{Cyl}(\Xi)) .$$

Inductive Step: Assume $\Xi \in \mathcal{S}_{\mathcal{M}}^{n+1}$. Define

$$g(\xi) := \int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi}(\xi \circ \gamma) \mu_{\mathcal{M}}^D(\xi, d\gamma) .$$

Let $\{g_m\}_{m \geq 0}$ be a sequence of simple functions that converges increasingly to g (cf. Proposition 3.1), which are denoted by $g_m = \sum_{i=1}^{l_m} d_m^i \cdot \mathbf{1}_{\Xi_i}$, where $l_m \geq 1$, $d_m^i \geq 0$ and $\Xi_i \in \mathcal{S}_{\mathcal{M}}^n$ for all $1 \leq i \leq l_m$. Then

$$\Pr_{\mathcal{M}, D, \alpha}^{n+1}(\Xi) = \int_{\Omega_{\mathcal{M}}^n} g(\xi) \Pr_{\mathcal{M}, D, \alpha}^n(d\xi) = \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{\mathcal{M}, D, \alpha}^n(\Xi_i) .$$

By induction hypothesis,

$$\Pr_{\mathcal{M}, D, \alpha}(\text{Cyl}(\Xi_i)) = \Pr_{\mathcal{M}, D, \alpha}^n(\Xi_i) = \sum_{s \in S} \alpha(s) \cdot \Pr_{\mathcal{M}, D, \mathcal{D}[s]}(\text{Cyl}(\Xi_i)) .$$

Thus we have

$$\begin{aligned}
& \Pr_{\mathcal{M},D,\alpha}(\text{Cyl}(\Xi)) \\
&= \Pr_{\mathcal{M},D,\alpha}^{n+1}(\Xi) \\
&= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \sum_{s \in S} \alpha(s) \cdot \Pr_{\mathcal{M},D,\mathcal{D}[s]}(\text{Cyl}(\Xi_i)) \\
&= \lim_{m \rightarrow \infty} \sum_{s \in S} \alpha(s) \cdot \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{\mathcal{M},D,\mathcal{D}[s]}(\text{Cyl}(\Xi_i)) \\
&= \sum_{s \in S} \alpha(s) \cdot \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{\mathcal{M},D,\mathcal{D}[s]}^n(\Xi_i) \\
&= \sum_{s \in S} \alpha(s) \cdot \int_{\Omega_{\mathcal{M}}^n} g(\xi) \Pr_{\mathcal{M},D,\mathcal{D}[s]}^n(d\xi) \\
&= \sum_{s \in S} \alpha(s) \cdot \Pr_{\mathcal{M},D,\mathcal{D}[s]}^{n+1}(\Xi) \\
&= \sum_{s \in S} \alpha(s) \cdot \Pr_{\mathcal{M},D,\mathcal{D}[s]}(\text{Cyl}(\Xi)) \quad ,
\end{aligned}$$

which justifies the induction hypothesis. \square

Lemma 7.1. $P_{\Pi}^{s,a}(t) \in \mathcal{S}_{\mathcal{M}}$ for all $\Pi \in \mathcal{S}_{\mathcal{M}}$, $(s, a) \in S \times \text{Act}$ and $t \in \mathbb{R}_{\geq 0}$. Analogously, $H_{\Xi}^{s,a}(t) \in \mathcal{S}_{\mathcal{M}}^{n-1}$ for all $n \geq 1$, $\Xi \in \mathcal{S}_{\mathcal{M}}^n$, $(s, a) \in S \times \text{Act}$ and $t \in \mathbb{R}_{\geq 0}$.

Proof. We first prove the case for paths. Fix some $s \in S$, $a \in \text{Act}$ and $t \geq 0$. Define the set \mathcal{S}' by: $\mathcal{S}' := \{\Pi \in \mathcal{S}_{\mathcal{M}} \mid P_{\Pi}^{s,a}(t) \in \mathcal{S}_{\mathcal{M}}\}$. We prove that $\mathcal{S}' = \mathcal{S}_{\mathcal{M}}$. By Remark 7.1, it suffices to prove that

1. $\{\text{Cyl}(\text{Hists}(\theta)) \mid \theta \text{ is a template.}\} \subseteq \mathcal{S}'$, and
2. \mathcal{S}' is a σ -algebra.

The first point follows directly from the definition of templates. To see the second point, one can verify that (i) $P_{\Omega_{\mathcal{M}}}^{s,a}(t) = \Omega_{\mathcal{M}}$, (ii) $P_{\Pi^c}^{s,a}(t) = [P_{\Pi}^{s,a}(t)]^c$ and (iii) $P_{\bigcup_{n \geq 0} \Pi_n}^{s,a}(t) = \bigcup_{n \geq 0} P_{\Pi_n}^{s,a}(t)$.

The proof for histories can be similarly obtained by proving the following fact: $\{\Xi \in \mathcal{S}_{\mathcal{M}}^n \mid H_{\Xi}^{s,a}(t) \in \mathcal{S}_{\mathcal{M}}^{n-1}\} = \mathcal{S}_{\mathcal{M}}^n$ for $n \geq 1$. \square

Lemma 7.2. Let $s \in S$, $a \in \text{Act}$ and $t \in \mathbb{R}_{\geq 0}$. For each measurable scheduler D , $D[s \xrightarrow{a,t}]$ is a measurable scheduler.

Proof. Fix some arbitrary $b \in \text{Act}$, $\epsilon \in [0, 1]$ and $n \geq 0$. Let D be a measurable scheduler. We have

$$\{\xi \in \text{Hists}^n(\mathcal{M}) \cap \text{Hists}_{\text{er}}(\mathcal{M}) \mid D[s \xrightarrow{a,t}](\xi, b) \leq \epsilon\} = H_{\Xi}^{s,a}(t) \quad ,$$

where $\Xi := \{\xi \in Hists^{n+1}(\mathcal{M}) \cap Hists_{er}(\mathcal{M}) \mid D(\xi, b) \leq \epsilon\}$. By Lemma 7.1, $H_{\Xi}^{s,a}(t)$ is measurable. As for histories in $Hists_{1a}(\mathcal{M})$, we define

$$\text{shift}(X) := \{(\xi, \tau) \in (Hists^n(\mathcal{M}) \cap Hists_{1a}(\mathcal{M})) \times \mathbb{R}_{\geq 0} \mid (s \xrightarrow{a,t} \xi, \tau) \in X\}$$

for each $X \in \mathcal{S}_{\mathcal{M}}^{n+1} \otimes \mathcal{B}(\mathbb{R}_{\geq 0})$. Let

$$\mathcal{X}' := \{X \in \mathcal{S}_{\mathcal{M}}^{n+1} \otimes \mathcal{B}(\mathbb{R}_{\geq 0}) \mid \text{shift}(X) \in \mathcal{S}_{\mathcal{M}}^n \otimes \mathcal{B}(\mathbb{R}_{\geq 0})\}.$$

Then one can prove $\mathcal{X}' = \mathcal{S}^{n+1} \otimes \mathcal{B}(\mathbb{R}_{\geq 0})$ in a similar way to the proof of Lemma 7.1. In detail, one can proceed by proving that

1. $\{\text{Hists}(\theta) \times I \mid |\theta| = n + 1 \text{ and } I \text{ is an interval of } \mathbb{R}_{\geq 0}\} \subseteq \mathcal{X}'$, and
2. \mathcal{X}' is a σ -algebra.

Then, from

$$\begin{aligned} & \{(\xi, \tau) \in Hists^n(\mathcal{M}) \times \mathbb{R}_{\geq 0} \mid D[s \xrightarrow{a,t}](\xi, \tau, b) \leq \epsilon\} = \\ & \text{shift}(\{(\xi, \tau) \in Hists^{n+1}(\mathcal{M}) \times \mathbb{R}_{\geq 0} \mid D(\xi, \tau, b) \leq \epsilon\}), \end{aligned}$$

we obtain that $\{(\xi, \tau) \in Hists^n(\mathcal{M}) \times \mathbb{R}_{\geq 0} \mid D[s \xrightarrow{a,t}](\xi, \tau, b) \leq \epsilon\}$ is measurable. \square

Proposition 7.2. $p_{\Pi, D}^{s,a}$ is a measurable function w.r.t $(\mathbb{R}_{\geq 0}, \mathcal{B}(\mathbb{R}_{\geq 0}))$ given any $\Pi \in \mathcal{S}_{\mathcal{M}}$, $s \in S$, $a \in Act$ and measurable scheduler D .

Proof. Fix some $s \in S$, $a \in Act$ and measurable scheduler D . Define the set

$$\mathcal{S}' := \{\Pi \in \mathcal{S}_{\mathcal{M}} \mid p_{\Pi, D}^{s,a} \text{ is a measurable function.}\}.$$

We show that (cf. Definition 3.14 and Definition 3.12)

1. \mathcal{S}' is a λ -system, and
2. the π -system $\{\text{Cyl}(\Xi) \mid \Xi \in \mathcal{S}_{\mathcal{M}}^n \text{ for some } n \geq 0\} \subseteq \mathcal{S}'$.

We first prove the second point. The proof proceeds through an induction on n such that $\text{Cyl}(\Xi) \in \mathcal{S}'$ for all $\Xi \in \mathcal{S}_{\mathcal{M}}^n$. The base step $n = 0$ is straightforward since $p_{\Pi, D}^{s,a}$ is constantly 0 or 1, depending on whether $s \in \Xi$ or not. The base step $n = 1$ follows from the definition of templates (Definition 7.6) and the finiteness of S and Act , which implies that $p_{\Pi, D}^{s,a}$ is a simple function (cf. Definition 3.6). Below we consider the inductive step for $n + 1$ with $n \geq 1$. Let $\Xi \in \mathcal{S}_{\mathcal{M}}^{n+1}$ with $n \geq 1$. For a measurable set of histories Ξ' and a measurable scheduler D' , we denote the measurable function $g_{\Xi', D'}$ by:

$$g_{\Xi', D'}(\xi) := \int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi'}(\xi \circ \gamma) \mu_{\mathcal{M}}^{D'}(\xi, d\gamma).$$

Let $\{g_m\}_{m \in \mathbb{N}}$ be a sequence of simple functions that converges increasingly and pointwisely to $g_{\Xi, D}$ (cf. Proposition 3.1). Denote $g_m = \sum_{i=1}^{l_m} d_m^i \cdot \mathbf{1}_{\Xi_m^i}$. By definition,

- $\mathbf{1}_{H_{\Xi}^{s,a}(t)}(\xi \circ \gamma) = \mathbf{1}_{\Xi}(s \xrightarrow{a,t} \xi \circ \gamma)$ for all $t \geq 0$, $\xi \in \Omega_{\mathcal{M}}^{n-1}$ and $\gamma \in \Gamma_{\mathcal{M}}$;
- $\mu_{\mathcal{M}}^{D[s \xrightarrow{a,t} \cdot]}(\xi, \cdot) = \mu_{\mathcal{M}}^D(s \xrightarrow{a,t} \xi, \cdot)$ for all $t \geq 0$ and $\xi \in \Omega_{\mathcal{M}}^{n-1}$.

Then, we have $g_{H_{\Xi}^{s,a}(t), D[s \xrightarrow{a,t} \cdot]}(\xi) = g_{\Xi, D}(s \xrightarrow{a,t} \xi)$ for all $t \geq 0$ and $\xi \in \Omega_{\mathcal{M}}^{n-1}$; it follows that

$$\lim_{m \rightarrow \infty} g_m(s \xrightarrow{a,t} \xi) = g_{\Xi, D}(s \xrightarrow{a,t} \xi) = g_{H_{\Xi}^{s,a}(t), D[s \xrightarrow{a,t} \cdot]}(\xi) .$$

Note that by definition, $g_m(s \xrightarrow{a,t} \xi) = \sum_{i=1}^{l_m} d_m^i \cdot \mathbf{1}_{H_{\Xi_m^i}^{s,a}(t)}(\xi)$ for all $t \geq 0$ and $\xi \in \Omega_{\mathcal{M}}^{n-1}$. Thus, by Proposition 3.1, we obtain

$$\begin{aligned} & p_{\text{Cyl}(\Xi), D}^{s,a}(t) \\ &= \Pr_{D[s \xrightarrow{a,t} \cdot], \mathbf{P}(s, a, \cdot)} \left(P_{\text{Cyl}(\Xi)}^{s,a}(t) \right) \\ &= \Pr_{D[s \xrightarrow{a,t} \cdot], \mathbf{P}(s, a, \cdot)}^n \left(H_{\Xi}^{s,a}(t) \right) \\ &= \int_{\Omega_{\mathcal{M}}^{n-1}} g_{H_{\Xi}^{s,a}(t), D[s \xrightarrow{a,t} \cdot]}(\xi) \Pr_{D[s \xrightarrow{a,t} \cdot], \mathbf{P}(s, a, \cdot)}^{n-1}(\mathrm{d}\xi) \\ &= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D[s \xrightarrow{a,t} \cdot], \mathbf{P}(s, a, \cdot)}^{n-1} \left(H_{\Xi_m^i}^{s,a}(t) \right) \\ &= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D[s \xrightarrow{a,t} \cdot], \mathbf{P}(s, a, \cdot)} \left(P_{\text{Cyl}(\Xi_m^i)}^{s,a}(t) \right) \\ &= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot p_{\text{Cyl}(\Xi_m^i), D}^{s,a}(t) \end{aligned}$$

for all $t \geq 0$. It follows from Proposition 3.2 and the induction hypothesis that $p_{\text{Cyl}(\Xi), D}^{s,a}$ is a measurable function.

Now the first point follows from the following facts:

- $p_{\Omega_{\mathcal{M}}, D}^{s,a}$ is measurable since $p_{\Omega_{\mathcal{M}}, D}^{s,a}(t) = 1$ for all $t \geq 0$;
- Whenever $\Pi_1, \Pi_2 \in \mathcal{S}'$ and $\Pi_1 \subseteq \Pi_2$, we have $\Pi_2 \setminus \Pi_1 \in \mathcal{S}'$ since $p_{\Pi_2 \setminus \Pi_1, D}^{s,a} = p_{\Pi_2, D}^{s,a} - p_{\Pi_1, D}^{s,a}$;
- For any sequence $\{\Pi_n\}_{n \geq 0}$ such that $\Pi_n \in \mathcal{S}'$ and $\Pi_n \subseteq \Pi_{n+1}$ for all $n \geq 0$, we have $\bigcup_{n \geq 0} \Pi_n \in \mathcal{S}'$ since $p_{\bigcup_{n \geq 0} \Pi_n, D}^{s,a} = \lim_{n \rightarrow \infty} p_{\Pi_n, D}^{s,a}$.

By applying Dynkin's π - λ Theorem (Theorem 3.3), we obtain that $\mathcal{S}_{\mathcal{M}} \subseteq \mathcal{S}'$, which implies the result. \square

Theorem 7.1. Let D be a measurable scheduler. For each $\Pi \in \mathcal{S}_{\mathcal{M}}$ and $s \in S_{\text{er}}$, we have

$$\Pr_{D, D[s]}(\Pi) = \sum_{a \in \text{En}(s)} D(s, a) \cdot \int_0^{\infty} f_{\mathbf{E}_{\text{er}}(s, a)}(t) \cdot p_{\Pi, D}^{s, a}(t) \mathrm{d}t .$$

For each $\Pi \in \mathcal{S}_{\mathcal{M}}$ and $s \in S_{\text{la}}$, we have

$$\Pr_{D, \mathcal{D}[s]}(\Pi) = \int_0^\infty f_{\mathbf{E}_{\text{la}}(s)}(t) \cdot \left[\sum_{a \in \text{En}(s)} D(s, t, a) \cdot p_{\Pi, D}^{s, a}(t) \right] dt .$$

Proof. Let $s \in S$. Define $\Pr' : \mathcal{S}_{\mathcal{M}} \rightarrow [0, 1]$ by:

$$\Pr'(\Pi) := \sum_{a \in \text{En}(s)} D(s, a) \cdot \int_0^\infty f_{\mathbf{E}_{\text{er}}(s, a)}(t) \cdot p_{\Pi, D}^{s, a}(t) dt$$

if $s \in S_{\text{er}}$, and

$$\Pr'(\Pi) := \int_0^\infty f_{\mathbf{E}_{\text{la}}(s)}(t) \cdot \left[\sum_{a \in \text{En}(s)} D(s, t, a) \cdot p_{\Pi, D}^{s, a}(t) \right] dt$$

if $s \in S_{\text{la}}$. We prove that \Pr' coincides with $\Pr_{D, \mathcal{D}[s]}$. It is straightforward from Proposition 7.2, Theorem 3.1 and Theorem 3.2 that \Pr' is a probability measure. Below we prove that \Pr' and $\Pr_{D, \mathcal{D}[s]}$ coincide on $\bigcup_{n \geq 0} \{\text{Cyl}(\Xi) \mid \Xi \in \mathcal{S}_{\mathcal{M}}^n\}$. The proof proceeds through induction on n .

Base Step: $n \in \{0, 1\}$ and $\Xi \in \mathcal{S}_{\mathcal{M}}^n$. If $n = 0$, we have

$$\Pr_{D, \mathcal{D}[s]}(\text{Cyl}(\Xi)) = \Pr_{D, \mathcal{D}[s]}^0(\Xi) = \mathbf{1}_{\Xi}(s) = \Pr'(\text{Cyl}(\Xi)) .$$

Otherwise, $n = 1$. If $s \in S_{\text{er}}$, then we have

$$\begin{aligned} & \Pr_{D, \mathcal{D}[s]}(\text{Cyl}(\Xi)) \\ &= \Pr_{D, \mathcal{D}[s]}^1(\Xi) \\ &= \int_{\Omega_{\mathcal{M}}^0} \left[\int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi}(\xi \circ \gamma) \mu_{\mathcal{M}}^D(\xi, d\gamma) \right] \Pr_{D, \mathcal{D}[s]}^0(d\xi) \\ &= \int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi}(s \circ \gamma) \mu_{\mathcal{M}}^D(s, d\gamma) \end{aligned}$$

Let $U := \{\gamma \in \Gamma_{\mathcal{M}} \mid s \circ \gamma \in \Xi\}$. Then $U \in \mathcal{U}_{\mathcal{M}}$. Thus,

$$\begin{aligned} & \int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi}(s \circ \gamma) \mu_{\mathcal{M}}^D(s, d\gamma) \\ &= \mu_{\mathcal{M}}^D(s, U) \\ &= \sum_{a \in \text{En}(s)} D(s, a) \cdot \int_{\mathbb{R}_{\geq 0}} f_{\mathbf{E}_{\text{er}}(s, a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{1}_U(a, t, s') \cdot \mathbf{P}(s, a, s') \right] dt \\ &= \sum_{a \in \text{En}(s)} D(s, a) \cdot \int_{\mathbb{R}_{\geq 0}} f_{\mathbf{E}_{\text{er}}(s, a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{1}_{\Xi}(s \xrightarrow{a, t} s') \cdot \mathbf{P}(s, a, s') \right] dt \\ &= \sum_{a \in \text{En}(s)} D(s, a) \cdot \int_{\mathbb{R}_{\geq 0}} f_{\mathbf{E}_{\text{er}}(s, a)}(t) \cdot p_{\text{Cyl}(\Xi), D}^{s, a}(t) dt \end{aligned}$$

where the last equality follows from Proposition 7.1. If $s \in S_{\text{la}}$, then we have

$$\begin{aligned}
& \Pr_{D, \mathcal{D}[s]}(\text{Cyl}(\Xi)) \\
&= \Pr_{D, \mathcal{D}[s]}^1(\Xi) \\
&= \int_{\Omega_{\mathcal{M}}^0} \left[\int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi}(\xi \circ \gamma) \mu_{\mathcal{M}}^D(\xi, d\gamma) \right] \Pr_{D, \mathcal{D}[s]}^0(d\xi) \\
&= \int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi}(s \circ \gamma) \mu_{\mathcal{M}}^D(s, d\gamma) .
\end{aligned}$$

Let $U := \{\gamma \in \Gamma_{\mathcal{M}} \mid s \circ \gamma \in \Xi\}$. Then $U \in \mathcal{U}_{\mathcal{M}}$. Thus,

$$\begin{aligned}
& \int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi}(s \circ \gamma) \mu_{\mathcal{M}}^D(s, d\gamma) \\
&= \mu_{\mathcal{M}}^D(s, U) \\
&= \int_{\mathbb{R}_{\geq 0}} f_{\mathbf{E}_{\text{la}}(s)}(t) \cdot \left\{ \sum_{a \in \text{En}(s)} D(s, t, a) \cdot \left[\sum_{s' \in S} \mathbf{1}_U(a, t, s') \cdot \mathbf{P}(s, a, s') \right] \right\} dt \\
&= \int_{\mathbb{R}_{\geq 0}} f_{\mathbf{E}_{\text{la}}(s)}(t) \cdot \left\{ \sum_{a \in \text{En}(s)} D(s, t, a) \cdot \left[\sum_{s' \in S} \mathbf{1}_{\Xi}(s \xrightarrow{a, t} s') \cdot \mathbf{P}(s, a, s') \right] \right\} dt \\
&= \int_0^{\infty} f_{\mathbf{E}_{\text{la}}(s)}(t) \cdot \left[\sum_{a \in \text{En}(s)} D(s, t, a) \cdot p_{\text{Cyl}(\Xi), D}^{s, a}(t) \right] dt .
\end{aligned}$$

where the last step follows from Proposition 7.1.

Inductive Step: Assume $\Xi \in \mathcal{S}_{\mathcal{M}}^{n+1}$ with $n \geq 1$. Denote

$$g_{\Xi', D'}(\xi) := \int_{\Gamma_{\mathcal{M}}} \mathbf{1}_{\Xi'}(\xi \circ \gamma) \mu_{\mathcal{M}}^{D'}(\xi, d\gamma) .$$

Let $\{g_m : \Omega_{\mathcal{M}}^n \rightarrow [0, 1]\}_{m \geq 0}$ be a sequence of simple functions that converges

increasingly to $g_{\Xi, D}$. Denote $g_m = \sum_{i=1}^{l_m} d_m^i \cdot \mathbf{1}_{\Xi_m^i}$. If $s \in S_{\text{er}}$, then

$$\begin{aligned}
& \Pr_{D, \mathcal{D}[s]} (\text{Cyl}(\Xi)) \\
&= \Pr_{D, \mathcal{D}[s]}^{n+1} (\Xi) \\
&= \int_{\Omega_{\mathcal{M}}^n} g_{\Xi, D}(\xi) \Pr_{D, \mathcal{D}[s]}^n(d\xi) \\
&= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D, \mathcal{D}[s]}^n(\Xi_m^i) \\
&= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \left[\sum_{a \in \text{En}(s)} D(s, a) \cdot \int_0^\infty f_{\mathbf{E}_{\text{er}}(s, a)}(t) \cdot p_{\text{Cyl}(\Xi_m^i), D}^{s, a}(t) dt \right] \\
&= \lim_{m \rightarrow \infty} \sum_{a \in \text{En}(s)} D(s, a) \cdot \left[\int_0^\infty f_{\mathbf{E}_{\text{er}}(s, a)}(t) \cdot \sum_{i=1}^{l_m} d_m^i \cdot p_{\text{Cyl}(\Xi_m^i), D}^{s, a}(t) dt \right] \\
&= \sum_{a \in \text{En}(s)} D(s, a) \cdot \left[\int_0^\infty f_{\mathbf{E}_{\text{er}}(s, a)}(t) \cdot \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot p_{\text{Cyl}(\Xi_m^i), D}^{s, a}(t) dt \right],
\end{aligned}$$

where the fourth equality follows from the induction hypothesis and the last equality follows from Theorem 3.2. Note that

$$\begin{aligned}
& \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot p_{\text{Cyl}(\Xi_m^i), D}^{s, a}(t) \\
&= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D[s \xrightarrow{a, t}], \mathbf{P}(s, a, \cdot)} \left(P_{\text{Cyl}(\Xi_m^i)}^{s, a}(t) \right) \\
&= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D[s \xrightarrow{a, t}], \mathbf{P}(s, a, \cdot)}^{n-1} \left(H_{\Xi_m^i}^{s, a}(t) \right).
\end{aligned}$$

Denote $g'_m(\xi) := \sum_{i=1}^{l_m} d_m^i \cdot \mathbf{1}_{H_{\Xi_m^i}^{s, a}(t)}(\xi)$. Then $g'_m(\xi) = g_m(s \xrightarrow{a, t} \xi)$. It follows that $\lim_{m \rightarrow \infty} g'_m(\xi) = g_{\Xi, D}(s \xrightarrow{a, t} \xi)$. By definition,

1. $\mathbf{1}_{\Xi} \left((s \xrightarrow{a, t} \xi) \circ \gamma \right) = \mathbf{1}_{H_{\Xi}^{s, a}(t)}(\xi \circ \gamma)$ for all combined actions γ ;
2. $\mu_{\mathcal{M}}^D(s \xrightarrow{a, t} \xi, U) = \mu_{\mathcal{M}}^{D[s \xrightarrow{a, t}]}(\xi, U)$ for all $U \in \mathcal{U}_{\mathcal{M}}$.

Thus $g_{\Xi, D}(s \xrightarrow{a, t} \xi) = g_{H_{\Xi}^{s, a}(t), D[s \xrightarrow{a, t} \cdot]}(\xi) = \lim_{m \rightarrow \infty} g'_m(\xi)$. It follows that

$$\begin{aligned} & \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D[s \xrightarrow{a, t} \cdot], \mathbf{P}(s, a, \cdot)}^{n-1} \left(H_{\Xi_m^i}^{s, a}(t) \right) \\ &= \int_{\Omega_{\mathcal{M}}^{n-1}} g_{H_{\Xi}^{s, a}(t), D[s \xrightarrow{a, t} \cdot]}(\xi) \Pr_{D[s \xrightarrow{a, t} \cdot], \mathbf{P}(s, a, \cdot)}^{n-1}(\mathrm{d}\xi) \\ &= \Pr_{D[s \xrightarrow{a, t} \cdot], \mathbf{P}(s, a, \cdot)}^n \left(H_{\Xi}^{s, a}(t) \right) . \end{aligned}$$

Then, we have

$$\begin{aligned} & \Pr_{D, \mathcal{D}[s]}(\mathrm{Cyl}(\Xi)) \\ &= \sum_{a \in \mathrm{En}(s)} D(s, a) \cdot \left[\int_0^\infty f_{\mathbf{E}_{\mathrm{er}}(s, a)}(t) \cdot \Pr_{D[s \xrightarrow{a, t} \cdot], \mathbf{P}(s, a, \cdot)}^n \left(H_{\Xi}^{s, a}(t) \right) \mathrm{d}t \right] \\ &= \sum_{a \in \mathrm{En}(s)} D(s, a) \cdot \left[\int_0^\infty f_{\mathbf{E}_{\mathrm{er}}(s, a)}(t) \cdot p_{\mathrm{Cyl}(\Xi), D}^{s, a}(t) \mathrm{d}t \right] . \end{aligned}$$

Now assume that $s \in \mathbb{S}_{\mathrm{la}}$. We have

$$\begin{aligned} & \Pr_{D, \mathcal{D}[s]}(\mathrm{Cyl}(\Xi)) \\ &= \Pr_{D, \mathcal{D}[s]}^{n+1}(\Xi) \\ &= \int_{\Omega_{\mathcal{M}}^n} g_{\Xi, D}(\xi) \Pr_{D, \mathcal{D}[s]}^n(\mathrm{d}\xi) \\ &= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D, \mathcal{D}[s]}^n(\Xi_m^i) \\ &= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \left\{ \int_0^\infty f_{\mathbf{E}_{\mathrm{la}}(s)}(t) \cdot \left[\sum_{a \in \mathrm{En}(s)} D(s, t, a) \cdot p_{\mathrm{Cyl}(\Xi_m^i), D}^{s, a}(t) \right] \mathrm{d}t \right\} \\ &= \lim_{m \rightarrow \infty} \int_0^\infty f_{\mathbf{E}_{\mathrm{la}}(s)}(t) \cdot \left\{ \sum_{a \in \mathrm{En}(s)} D(s, t, a) \cdot \left[\sum_{i=1}^{l_m} d_m^i \cdot p_{\mathrm{Cyl}(\Xi_m^i), D}^{s, a}(t) \right] \right\} \mathrm{d}t \\ &= \int_0^\infty f_{\mathbf{E}_{\mathrm{la}}(s)}(t) \cdot \left\{ \sum_{a \in \mathrm{En}(s)} D(s, t, a) \cdot \left[\lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot p_{\mathrm{Cyl}(\Xi_m^i), D}^{s, a}(t) \right] \right\} \mathrm{d}t \end{aligned}$$

where the fourth equality is from the induction hypothesis and the last

equality follows from Theorem 3.2. Note that

$$\begin{aligned}
& \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot p_{\text{Cyl}(\Xi_m^i), D}^{s,a}(t) \\
&= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D[s \xrightarrow{a,t}], \mathbf{P}(s,a,\cdot)} \left(P_{\text{Cyl}(\Xi_m^i)}^{s,a}(t) \right) \\
&= \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D[s \xrightarrow{a,t}], \mathbf{P}(s,a,\cdot)}^{n-1} \left(H_{\Xi_m^i}^{s,a}(t) \right)
\end{aligned}$$

Denote $g'_m(\xi) := \sum_{i=1}^{l_m} d_m^i \cdot \mathbf{1}_{H_{\Xi_m^i}^{s,a}(t)}(\xi)$. Then $g'_m(\xi) = g_m(s \xrightarrow{a,t} \xi)$. It follows that $\lim_{m \rightarrow \infty} g'_m(\xi) = g_{\Xi, D}(s \xrightarrow{a,t} \xi)$. By definition,

1. $\mathbf{1}_{\Xi} \left((s \xrightarrow{a,t} \xi) \circ \gamma \right) = \mathbf{1}_{H_{\Xi}^{s,a}(t)}(\xi \circ \gamma)$ for all combined action γ ;
2. $\mu_{\mathcal{M}}^D(s \xrightarrow{a,t} \xi, U) = \mu_{\mathcal{M}}^{D[s \xrightarrow{a,t}]}(\xi, U)$ for all $U \in \mathcal{U}_{\mathcal{M}}$.

Thus $g_{\Xi, D}(s \xrightarrow{a,t} \xi) = g_{H_{\Xi}^{s,a}(t), D[s \xrightarrow{a,t}]}(\xi) = \lim_{m \rightarrow \infty} g'_m(\xi)$. It follows that

$$\begin{aligned}
& \lim_{m \rightarrow \infty} \sum_{i=1}^{l_m} d_m^i \cdot \Pr_{D[s \xrightarrow{a,t}], \mathbf{P}(s,a,\cdot)}^{n-1} \left(H_{\Xi_m^i}^{s,a}(t) \right) \\
&= \int_{\Omega_{\mathcal{M}}^{n-1}} g_{H_{\Xi}^{s,a}(t), D[s \xrightarrow{a,t}]}(\xi) \Pr_{D[s \xrightarrow{a,t}], \mathbf{P}(s,a,\cdot)}^{n-1}(\mathrm{d}\xi) \\
&= \Pr_{D[s \xrightarrow{a,t}], \mathbf{P}(s,a,\cdot)}^n \left(H_{\Xi}^{s,a}(t) \right) .
\end{aligned}$$

Then we have

$$\begin{aligned}
& \Pr_{D, \mathcal{D}[s]}(\text{Cyl}(\Xi)) \\
&= \int_0^\infty f_{\mathbf{E}_{\text{Ia}}(s)}(t) \cdot \left[\sum_{a \in \mathbf{E}_{\text{N}}(s)} D(s, t, a) \cdot \Pr_{D[s \xrightarrow{a,t}], \mathbf{P}(s,a,\cdot)}^n \left(H_{\Xi}^{s,a}(t) \right) \right] \mathrm{d}t \\
&= \int_0^\infty f_{\mathbf{E}_{\text{Ia}}(s)}(t) \cdot \left[\sum_{a \in \mathbf{E}_{\text{N}}(s)} D(s, t, a) \cdot p_{\text{Cyl}(\Xi), D}^{s,a}(t) \right] \mathrm{d}t .
\end{aligned}$$

It follows that the inductive step is completed. \square



Chapter 8

Acceptance Probability of CTMC-Paths by DTA

A continuous-time Markov chain (CTMC) is a continuous-time Markov decision process (cf. Chapter 7) such that all states of the CTMDP are late-schedulable and have exactly one enabled action (i.e., exactly one non-deterministic choice). (It does not matter whether states of a CTMC are early-schedulable or late schedulable, since each state has exactly one enabled action.) Under a CTMC, there is a unique measurable scheduler which selects upon each history the only enabled action of the last state of the history. Intuitively, the class of CTMCs is a “deterministic” subclass of the one of CTMDPs, where no actual non-determinism is present. Applications of CTMCs include Markovian queueing networks [14, 68], stochastic Petri nets, system biology and so forth.

Formal verification of CTMCs has received much attention in recent years [6]. Many applicable results have been obtained on time-bounded reachability [4, 35], CSL model checking [4, 75], and so forth. In this chapter, we focus on verifying CTMCs against timed-automata specification. In particular we consider approximating the probabilities of sets of CTMC-paths accepted by a deterministic timed automata (DTA) [1, 17]. In general, DTA represents a wide class of linear real-time specifications. For example, we can describe time-bounded reachability probability “to reach target set $G \subseteq S$ within time bound T while avoiding unsafe states $U \subseteq S$ ” ($G \cap U = \emptyset$) by the single-clock DTA \mathcal{A}_1 (Fig. 8.1), and the property “to reach target set $G \subseteq S$ within time bound T_1 while successively remaining in unsafe states $U \subseteq S$ for at most T_2 time” ($G \cap U = \emptyset$) by the two-clock DTA \mathcal{A}_2 (Fig. 8.2), both with initial configuration $(q_0, \vec{0})$. (We omit redundant locations that cannot reach the accepting state.)

The problem to verify CTMCs against DTA-specifications is first considered by Donatelli *et al.* [31] where they enriched CSL [4] with an acceptance condition of one-clock DTA, yielding the logic CSL^{TA} . In their paper, they

Figure 8.1: DTA \mathcal{A}_1 Figure 8.2: DTA \mathcal{A}_2

proved that CSL^{TA} is at least as expressive as CSL and asCSL [4, 2], and is strictly more expressive than CSL. Moreover, they presented a model-checking algorithm for CSL^{TA} using Markov regenerative processes [31]. Chen *et al.* [24] systematically studied the DTA-acceptance condition on CTMC-paths. More specifically, they proved that the set of CTMC-path accepted by a multi-clock DTA is measurable and proposed a system of integral equations which characterizes the acceptance probabilities. Moreover, they demonstrated that the product of a CTMC and a DTA is a piecewise deterministic Markov process [27], a stochastic dynamic system which integrates both discrete control and continuous evolution. Afterwards, Barbot *et al.* [9] put the approximation of DTA-acceptance probabilities of CTMC-paths into practice, especially the algorithm on one-clock DTA, which is first devised by Donatelli *et al.* [31] and then re-written by Chen *et al.* [24]. Later on, Chen *et al.* [23] proposed approximation algorithms for time-bounded verification of several linear real-time specifications, where the restricted time-bounded case, in which the time guard $x < T$ with a fresh clock x and a time bound T is enforced on each edge that leads to some final state of the DTA, is covered. Very recently, Mikeev *et al.* [55] applies the notion of DTA-acceptance condition on CTMC-paths to system biology. It is worth noting that Brázdil *et al.* also studied DTA-specifications in [17]. However they focused on semi-Markov processes as the underlying continuous-time stochastic model and limit frequencies of locations (in the DTA) as the performance measures, rather than path-acceptance conditions.

The contributions of this chapter are as follows. We start by providing a rigorous proof for the measurability of the set of CTMC-paths accepted by

a DTA, correcting the proof provided by Chen *et al.* [24]. We confirm the correctness of the integral equation system characterizing acceptance probabilities provided by Chen *et al.* [24] through a direct application of Theorem 7.1, and derive a differential characterization. Based on the differential characterization, we present an approximation algorithm to approximate acceptance probabilities, and provide tight error bounds for the approximation algorithm. Whereas other works [9, 55, 31] focus on single-clock DTA, our approximation algorithm is applicable to any (multi-clock) DTA. To our knowledge, this is the first such approximation algorithm with explicit error bounds. Barbot *et al.* [9] suggested an approximation scheme, but did not provide error bounds.

The chapter is organized as follows. In Section 8.1, we introduce the notion of continuous-time Markov chains. In Section 8.2, we introduce the notion of deterministic timed automata. In Section 8.3, we prove the measurability of accepted paths, and present the proof for the integral equations [24] that characterize the acceptance probability. In Section 8.4, we develop several mathematical technicalities useful to our main result. In Section 8.5, we propose a differential characterization for the family of acceptance probability functions. Base on these results, we establish and solve our approximation scheme in Section 8.6. Finally, Section 8.7 concludes the chapter.

8.1 Continuous-Time Markov Chains

In this section, we present basic definitions for continuous-time Markov chains [36, 68]. These definitions are just simplified definitions from the ones in Chapter 7, which omit the single enabled actions at each state and the different between early- and late-schedulable states.

Definition 8.1. A continuous-time Markov chain (CTMC) is a triple $(S, \mathbf{E}, \mathbf{P})$, where

- S is a finite non-empty set of states,
- $\mathbf{E} : S \rightarrow \mathbb{R}_{>0}$ is a total exit-rate function, and
- $\mathbf{P} : S \times S \rightarrow [0, 1]$ is a transition matrix such that $\sum_{u \in S} \mathbf{P}(s, u) = 1$ for all $s \in S$.

A path of a CTMC $(S, \mathbf{E}, \mathbf{P})$ is an infinite sequence $\langle s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} \dots \rangle$ such that $s_n \in S$ and $t_n \in \mathbb{R}_{\geq 0}$ for all $n \in \mathbb{N}_0$. The set of paths of a CTMC \mathcal{M} is denoted by $\text{Paths}(\mathcal{M})$.

The evolution of a CTMC and the notions of paths and histories are the same as the one of a CTMDP (cf. Chapter 7). The definition of the probability space $\text{Pr}_{\mathcal{M}, D, \alpha}$, with a CTMC \mathcal{M} , the unique measurable scheduler D for \mathcal{M} which always chooses the Dirac distribution at sole enabled

action, and an initial distribution α for \mathcal{M} , is defined in the same way as in Chapter 7. Since the scheduler is unique in the context of CTMCs, we will abbreviate ‘ $\text{Pr}_{\mathcal{M},D,\alpha}$ ’ as ‘ $\text{Pr}_{\mathcal{M},\alpha}$ ’, and omit the subscript \mathcal{M} whenever possible.

In this chapter, we consider CTMCs with labelled states. Informally, a labelling function assigns to each state of a CTMC a label which specifies the atomic property that holds at the state.

Definition 8.2. *Let $\mathcal{M} = (S, \mathbf{E}, \mathbf{P})$ be a CTMC. A labelling function \mathcal{L} is a function $S \rightarrow L$, where L is a set of labels.*

With labels, one can view any path (trajectory) $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} \dots$ of a CTMC as a *timed word* $t_0\mathcal{L}(s_0)t_1\mathcal{L}(s_1)\dots$; this allows one to specify desired linear properties on a CTMC via acceptance conditions of a timed automaton (cf. Section 8.2).

8.2 Deterministic Timed Automata

The class of *timed automata* [1] is an extension of that of finite automata with *clocks*. It can be used either as a language acceptor for words integrated with time information (i.e., *timed words*), or as a formal description for timed transition systems (just like finite automata for finite reactive systems). The class of *deterministic timed automata* is a subclass of timed automata, which is an analogue to deterministic finite automata.

The key concept in timed automata is the concept of *clocks*. Generally, clocks are abstract variables holding non-negative real values interpreted as time-elapse quantities.

Definition 8.3. *Let \mathcal{X} be a finite set of clocks. A (clock) valuation on \mathcal{X} is a function $\eta : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$. We denote by $\text{Val}(\mathcal{X})$ the set of valuations on \mathcal{X} .*

Clocks play a role in the notion of timed automata through the notion of *guards* (or *clock constraints*). Intuitively, a guard is a logical formula interpreted over clock valuations.

Definition 8.4. *A guard g (or clock constraint) over a finite set of clocks \mathcal{X} is a logical formula generated by the following grammar:*

$$g ::= x \bowtie c \mid g \wedge g ,$$

where $x \in \mathcal{X}$, $\bowtie \in \{<, \leq, >, \geq\}$, $c \in \mathbb{N}_0$ and ‘ \wedge ’ represents the logical ‘and’ operator. We denote the set of guards over \mathcal{X} by $\Phi(\mathcal{X})$. For each $\eta \in \text{Val}(\mathcal{X})$ and $g \in \Phi(\mathcal{X})$, the satisfaction relation $\eta \models g$ is inductively defined by: $\eta \models x \bowtie c$ iff $\eta(x) \bowtie c$, and $\eta \models g_1 \wedge g_2$ iff $\eta \models g_1$ and $\eta \models g_2$.

Given $X \subseteq \mathcal{X}$, $\eta \in \text{Val}(\mathcal{X})$ and $t \in \mathbb{R}_{\geq 0}$, the valuations $\eta[X := 0]$, $\eta + t$, and $\eta - t$ are defined as follows:

1. $\eta[X := 0](x) := 0$ for all $x \in X$, and $\eta[X := 0](x) := \eta(x)$ for all $x \in \mathcal{X} \setminus X$;
2. $(\eta + t)(x) := \eta(x) + t$ for all $x \in \mathcal{X}$;
3. $(\eta - t)(x) := \eta(x) - t$ for all $x \in \mathcal{X}$, provided that $\eta(x) \geq t$ for all $x \in \mathcal{X}$.

Intuitively, $\eta[X := 0]$ is obtained by resetting all clocks of X to zero on η , and $\eta + t$ resp. $\eta - t$ is obtained by delaying resp. backtracking t time units from η .

In this chapter, we will also view a clock valuation as a real vector indexed by the elements of \mathcal{X} . We may also refer $g \in \Phi(\mathcal{X})$ to the set of valuations that satisfy g : this may happen in the phrases such as “ $g_1 \cap g_2$ ”, etc.

The notion of deterministic time automata is a subclass of timed automata [1], which is illustrated as follows.

Definition 8.5. [1, 17] A deterministic timed automaton (DTA) is a tuple $(Q, \Sigma, \mathcal{X}, \Delta, F)$, where:

- Q is a finite set of locations;
- $F \subseteq Q$ is a set of final locations;
- Σ is a finite alphabet of signatures;
- \mathcal{X} is a finite set of clocks;
- $\Delta \subseteq Q \times \Sigma \times \Phi(\mathcal{X}) \times 2^{\mathcal{X}} \times Q$ is a finite set of rules such that
 1. Δ is deterministic: for all $(q_1, a_1, g_1, X_1, q'_1), (q_2, a_2, g_2, X_2, q'_2) \in \Delta$, if $(q_1, a_1) = (q_2, a_2)$ and $g_1 \cap g_2 \neq \emptyset$ then $(g_1, X_1, q'_1) = (g_2, X_2, q'_2)$.
 2. Δ is total: for all $(q, a) \in Q \times \Sigma$ and $\eta \in \text{Val}(\mathcal{X})$, there exists $(q, a, g, X, q') \in \Delta$ such that $\eta \models g$.

In the following, we fix a DTA $\mathcal{A} = (Q, \Sigma, \mathcal{X}, \Delta, F)$. The next definition illustrates the notion of configurations which is closely related to runs of DTAs on timed words. To ease the notation, given any triple $(q, \eta, a) \in Q \times \text{Val}(\mathcal{X}) \times \Sigma$, we define $(\mathbf{g}_{q,a}^\eta, \mathbf{X}_{q,a}^\eta, \mathbf{q}_{q,a}^\eta) \in \Phi(\mathcal{X}) \times 2^{\mathcal{X}} \times Q$ to be the unique triple such that the rule $(q, a, \mathbf{g}_{q,a}^\eta, \mathbf{X}_{q,a}^\eta, \mathbf{q}_{q,a}^\eta) \in \Delta$ satisfies $\eta \models \mathbf{g}_{q,a}^\eta$; the triple is well-defined and unique since Δ is deterministic and total.

Definition 8.6. [1] A configuration of \mathcal{A} is a pair (q, η) , where $q \in Q$ and $\eta \in \text{Val}(\mathcal{X})$. A timed signature is a pair (t, a) where $t \in \mathbb{R}_{\geq 0}$ and $a \in \Sigma$. The one-step transition function

$$\kappa : (Q \times \text{Val}(\mathcal{X})) \times (\mathbb{R}_{\geq 0} \times \Sigma) \rightarrow Q \times \text{Val}(\mathcal{X})$$

is defined by: $\kappa((q, \eta), (t, a)) = (\mathbf{q}_{q,a}^{\eta+t}, (\eta + t)[\mathbf{X}_{q,a}^{\eta+t} := 0])$.

Intuitively, the configuration $\kappa((q, \eta), (t, a))$ is obtained as follows: firstly, we delay t time-units at (q, η) to obtain $(q, \eta + t)$; then we find the unique rule $(q, a, g, X, q') \in \Delta$ such that $\eta + t \models g$; finally, we obtain $\kappa((q, \eta), (t, a))$ by changing the location to q' and resetting $\eta + t$ with X . For the sake of simplicity, we may represent “ $\kappa((q, \eta), (t, a)) = (q', \eta')$ ” by the more intuitive phrase “ $(q, \eta) \xrightarrow{(t, a)} (q', \eta')$ ”.

Then the notion of timed words and runs on timed words is demonstrated in the following definition.

Definition 8.7. [1] *A timed word is an infinite sequence of timed signatures. The run of \mathcal{A} on a timed word $w = \{(t_n, a_n)\}_{n \in \mathbb{N}_0}$ with initial configuration (q, η) , denoted by $\mathcal{A}_{q, \eta}(w)$, is the unique infinite sequence $\{(q_n, \eta_n)(t_n, a_n)\}_{n \in \mathbb{N}_0}$ which satisfies that $(q_0, \eta_0) = (q, \eta)$ and $(q_{n+1}, \eta_{n+1}) = \kappa((q_n, \eta_n), (t_n, a_n))$ for $n \geq 0$.*

A timed word w is accepted by \mathcal{A} with initial configuration (q, η) iff $\mathcal{A}_{q, \eta}(w) = \{(q_n, \eta_n)(t_n, a_n)\}_{n \in \mathbb{N}_0}$ satisfies that $q_n \in F$ for some $n \geq 0$. Moreover, w is accepted by \mathcal{A} with initial configuration (q, η) within k steps ($k \geq 0$) iff $\mathcal{A}_{q, \eta}(w) = \{(q_n, \eta_n)(t_n, a_n)\}_{n \in \mathbb{N}_0}$ satisfies that $q_n \in F$ for some $0 \leq n \leq k$.

8.3 Measurability and The Integral Equations

In this section, we first formally define the notion of DTA-acceptance on CTMC-trajectories, following the definition from [24], which is the central theme of this chapter. Then we provide a rigorous proof for the measurability of the set of CTMC-trajectories accepted by a DTA, and the system of integral equations that characterizes the acceptance probability.

Below we fix a CTMC $\mathcal{M} = (S, \mathbf{E}, \mathbf{P})$, a DTA $\mathcal{A} = (Q, \Sigma, \mathcal{X}, \Delta, F)$ and a labelling function $\mathcal{L} : S \rightarrow \Sigma$. Firstly, we formally define the notion of DTA-acceptance on CTMC-trajectories.

Definition 8.8. [24] *The set of \mathcal{M} -paths accepted by \mathcal{A} w.r.t the triple $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$, denoted by $\text{Paths}^{\mathcal{M}, \mathcal{A}}(s, q, \eta)$, is defined by:*

$$\begin{aligned} \text{Paths}^{\mathcal{M}, \mathcal{A}}(s, q, \eta) &:= \{\pi \in \text{Paths}(\mathcal{M}) \mid \\ &\pi[0] = s \text{ and } \mathcal{L}(\pi) \text{ is accepted by } \mathcal{A} \text{ with initial configuration } (q, \eta)\} , \end{aligned}$$

where $\mathcal{L}(\pi)$ is the timed word given by:

$$\mathcal{L}(\pi) := \langle (\pi\langle 0 \rangle, \mathcal{L}(\pi[0])) (\pi\langle 1 \rangle, \mathcal{L}(\pi[1])) \dots (\pi\langle n \rangle, \mathcal{L}(\pi[n])) \rangle .$$

Moreover, the set of \mathcal{M} -paths accepted by \mathcal{A} w.r.t $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$ within k -steps ($k \geq 0$), denoted by $\text{Paths}_k^{\mathcal{M}, \mathcal{A}}(s, q, \eta)$, is defined as the set of paths $\pi \in \text{Paths}(\mathcal{M})$ such that $\pi[0] = s$ and $\mathcal{L}(\pi)$ is accepted by \mathcal{A} with initial configuration (q, η) within k -steps.

In this chapter, we omit the superscript “ \mathcal{M}, \mathcal{A} ” if the underlying context is clear.

Below we prove that $\text{Paths}(s, q, \eta)$ is measurable w.r.t the measurable space $(\Omega_{\mathcal{M}}, \mathcal{S}_{\mathcal{M}})$ (cf. Definition 7.7) for all $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$, where the measurable space results from omitting the deterministic choices of actions in the definition for CTMDPs. Given measurability, the integral-equation system that characterizes the acceptance probability [24] follows directly from Theorem 7.1.

Remark 8.1. *We point out the flaw in the measurability proof by Chen et al. [24]. The error appears on Page 11 under the label “(1b)” which handles the equality guards in timed transitions. In (1b), for an timed transition e emitted from q with guard $x = K$, four DTAs $\mathcal{A}_e, \bar{\mathcal{A}}_e, \mathcal{A}_e^>, \mathcal{A}_e^<$ are defined w.r.t the original DTA \mathcal{A} . Then it is argued that*

$$\text{Paths}^c(\mathcal{A}_e) = \text{Paths}^c(\bar{\mathcal{A}}_e) \setminus (\text{Paths}^c(\mathcal{A}_e^>) \cup \text{Paths}^c(\mathcal{A}_e^<)) .$$

This is incorrect. The left part $\text{Paths}^c(\mathcal{A}_e)$ excludes all timed paths which involve both the guard $x > K$ and the guard $x < K$ (from q). However the right part does not. So the left and right part are not equal. \square

By definition, $\bigcup_{k \geq 0} \text{Paths}_k(s, q, \eta) = \text{Paths}(s, q, \eta)$. Thus in order to prove the measurability of $\text{Paths}(s, q, \eta)$, it suffices to prove that each set $\text{Paths}_k(s, q, \eta)$ is measurable under $(\Omega_{\mathcal{M}}, \mathcal{S}_{\mathcal{M}})$. The following proposition presents this point, whose proof uses a decomposition of $\text{Path}_k(s, q, \eta)$ into subsets of paths. Below given a word β of length k ($k \geq 1$) on some alphabet, we denote by β_i the i -th symbol of β (i.e., $\beta = \beta_1 \dots \beta_k$).

Proposition 8.1. *For all $k \in \mathbb{N}_0$ and $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$, the set $\text{Paths}_k(s, q, \eta)$ is measurable w.r.t. $(\Omega_{\mathcal{M}}, \mathcal{S}_{\mathcal{M}})$.*

Proof. Fix some $k \in \mathbb{N}_0$ and $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$. If $k = 0$, then the measurability follows from the fact that either $\text{Paths}_0(s, q, \eta) = \text{Paths}(\mathcal{M})$ ($q \in F$) or $\text{Paths}_0(s, q, \eta) = \emptyset$ ($q \notin F$). From now on, we assume that $k \geq 1$.

Given a word $\beta \in S^k$ and a word $\gamma \in \Delta^k$, we define the set $\text{Paths}_{\beta, \gamma}(s, q, \eta)$ to be the set of all $\pi \in \text{Paths}(\mathcal{M})$ such that $\pi[0] = s$, $\pi[i] = \beta_{i+1}$ for all $0 \leq i \leq k-1$, and the run $\mathcal{A}_{q, \eta}(\mathcal{L}(\pi)) = \{(q_n, \eta_n)(t_n, a_n)\}_{n \geq 0}$ satisfies that $(q_i, a_i, \mathbf{g}_{q_i, a_i}^{\eta_i + t_i}, \mathbf{X}_{q_i, a_i}^{\eta_i + t_i}, \mathbf{q}_{q_i, a_i}^{\eta_i + t_i}) = \gamma_{i+1}$ for all $0 \leq i \leq k-1$. Intuitively, the set $\text{Paths}_{\beta, \gamma}(s, q, \eta)$ consists of all paths whose runs sequentially follow the state sequence β and the rule sequence γ . By definition,

$$\bigcup_{(\beta, \gamma)} \text{Paths}_{\beta, \gamma}(s, q, \eta) = \text{Paths}_k(s, q, \eta) ,$$

where (β, γ) ranges over the set

$$S^k \times \left\{ \gamma' \in \Delta^k \mid \exists 1 \leq i \leq k. \exists q \in F. (q \text{ is a component of the tuple } \gamma_i) \right\} .$$

Thus, it suffices to show that each $\text{Paths}_{\beta,\gamma}(s, q, \eta)$ is measurable.

Fix some $\beta \in S^k$ and $\gamma \in \Delta^k$. If either $\beta_1 \neq s$, or the first component of the tuple γ_1 is not q , or β and γ mismatch (i.e., the action in some γ_i does not match $\mathcal{L}(\beta_i)$), or the last component of γ_i is not equal to the first component of γ_{i+1} for some i , then $\text{Paths}_{\beta,\gamma}(s, q, \eta) = \emptyset$, which is clearly measurable. From now on, we assume that $\beta \in S^k$ and $\gamma \in \Delta^k$ survive the failure conditions in this paragraph.

Consider an arbitrary path

$$\pi = s_0 \xrightarrow{t_0} \dots s_{k-1} \xrightarrow{t_{k-1}} \dots \in \text{Paths}_{\beta,\gamma}(s, q, \eta) .$$

(Note that $\text{Paths}_{\beta,\gamma}(s, q, \eta)$ may still be empty.) Let the run $\mathcal{A}_{q,\eta}(\mathcal{L}(\pi))$ be $\{(q_n, \eta_n)(t_n, a_n)\}_{n \geq 0}$. Since the rule sequence γ is fixed, we can represent each η_i ($0 \leq i \leq k-1$) as a vector-valued *linear* function h_i on t_0, \dots, t_{k-1} . (In fact, h_i is a function on t_0, \dots, t_{i-1} .) Thus, $\pi \in \text{Paths}_{\beta,\gamma}(s, q, \eta)$ implies that (i) $\pi[i] = \beta_{i+1}$ for all $0 \leq i \leq k-1$ and (ii) $h_i(t_0, \dots, t_{k-1}) + t_i \models g_{i+1}$ for all $0 \leq i \leq k-1$, where g_j is the rule component of the tuple γ_j (for $1 \leq j \leq k$). Conversely, one easily verifies that for all $\pi \in \text{Paths}(\mathcal{M})$, if (i) $\pi[i] = \beta_{i+1}$ for all $0 \leq i \leq k-1$ and (ii) $h_i(t_0, \dots, t_{k-1}) + t_i \models g_{i+1}$ for all $0 \leq i \leq k-1$, then $\pi \in \text{Paths}_{\beta,\gamma}(s, q, \eta)$. Thus, there exists some set $\mathcal{J} \subseteq \mathbb{R}^k$ definable through a system of a finite number of linear inequalities, dependent only on γ and η , such that $\pi \in \text{Paths}_{\beta,\gamma}(s, q, \eta)$ iff $\pi[i] = \beta_{i+1}$ and $(\pi(0), \dots, \pi(k-1)) \in \mathcal{J}$. We first assume that \mathcal{J} can be defined through a system of finitely many linear inequalities where all comparison operators are \leq or \geq . Then \mathcal{J} is a closed subset of \mathbb{R}^k , which implies that

$$\text{Paths}_{\beta,\gamma}(s, q, \eta) = \bigcap_{n \in \mathbb{N}_0} \bigcup \{ \text{Cyl}(\text{Hists}(\theta)) \mid \theta \in \mathcal{C}_n \} ,$$

where \mathcal{C}_n is the set of all templates θ (cf. Definition 7.6) such that there exists non-negative integers n_0, \dots, n_{k-1} such that

- $\left(\prod_{i=0}^{k-1} \left[\frac{n_i}{n}, \frac{n_i+1}{n} \right] \right) \cap \mathcal{J} \neq \emptyset$ and
-

$$\theta = \langle \beta_1, \left[\frac{n_0}{n}, \frac{n_0+1}{n} \right] \times \{\beta_2\}, \dots, \left[\frac{n_{k-2}}{n}, \frac{n_{k-2}+1}{n} \right] \times \{\beta_k\}, \left[\frac{n_{k-1}}{n}, \frac{n_{k-1}+1}{n} \right] \times S \rangle .$$

Intuitively, the closed-ness of \mathcal{J} allows us to “cover” \mathcal{J} with arbitrarily small grids, where each grid (with β) can be interpreted directly as a template (cf. Definition 7.6). It follows that $\text{Paths}_{\beta,\gamma}(s, q, \eta)$ is measurable. Now we assume that \mathcal{J} is defined through a system of finitely many linear inequalities which involves the comparison operator ‘ $<$ ’ or ‘ $>$ ’. This case can be reduced to the previous case (where only ‘ \leq ’ and ‘ \geq ’ are present) as follows: firstly,

we modify each inequality $\mathbf{a}^T \cdot \mathbf{t} < c$ in \mathcal{J} to $\mathbf{a}^T \cdot \mathbf{t} \leq c - \frac{1}{n}$ and each inequality $\mathbf{a}^T \cdot \mathbf{t} > c$ in \mathcal{J} to $\mathbf{a}^T \cdot \mathbf{t} \geq c + \frac{1}{n}$, to obtain a new system of linear inequalities which defines a closed set $\mathcal{J}_n \subseteq \mathbb{R}^k$. From previous analysis, the set Π_n of all paths π such that (i) $\pi[i] = \beta_{i+1}$ for all $0 \leq i \leq k-1$ and (ii) $(\pi\langle 0 \rangle, \dots, \pi\langle k-1 \rangle) \in \mathcal{J}_n$ is measurable, for all $n \in \mathbb{N}$. By definition and the representation of \mathcal{J} through a finite system of linear inequalities, $\mathcal{J} = \bigcup_{n \in \mathbb{N}} \mathcal{J}_n$. Thus by $\text{Paths}_{\beta, \gamma}(s, q, \eta) = \bigcup_{n \in \mathbb{N}} \Pi_n$, $\text{Paths}_{\beta, \gamma}(s, q, \eta)$ is measurable. \square

Based on Proposition 8.1 and Theorem 7.1, we directly obtain the integral characterization as follows. Below we define the following functions:

- for each $n \in \mathbb{N}_0$, we define the function

$$\text{prob}_n : S \times Q \times \text{Val}(\mathcal{X}) \rightarrow [0, 1]$$

by: $\text{prob}_n(s, q, \eta) = \Pr_{\mathcal{D}[s]}(\text{Paths}_n(s, q, \eta))$ (cf. Definition 8.8);

- we define the function

$$\text{prob} : S \times Q \times \text{Val}(\mathcal{X}) \rightarrow [0, 1]$$

by: $\text{prob}(s, q, \eta) = \Pr_{\mathcal{D}[s]}(\text{Paths}(s, q, \eta))$.

Moreover, we abbreviate $(\mathbf{g}_{q, \mathcal{L}(s)}^\eta, \mathbf{X}_{q, \mathcal{L}(s)}^\eta, \mathbf{q}_{q, \mathcal{L}(s)}^\eta)$ as $(\mathbf{g}_{q, s}^\eta, \mathbf{X}_{q, s}^\eta, \mathbf{q}_{q, s}^\eta)$.

Theorem 8.1. *The family $\{\text{prob}_n\}_{n \in \mathbb{N}_0}$ satisfies the following properties: (i) $\text{prob}_0(s, q, \eta) = \mathbf{1}_F(q)$; (ii) for all $n \in \mathbb{N}_0$, if $q \in F$ then $\text{prob}_{n+1}(s, q, \eta) = 1$, otherwise*

$$\text{prob}_{n+1}(s, q, \eta) = \int_0^\infty \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}_n(u, \mathbf{q}_{q, s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q, s}^{\eta+t} := 0]) \right] \right\} dt .$$

The function prob satisfies the following system of integral equations: If $q \in F$ then $\text{prob}(s, q, \eta) = 1$, otherwise

$$\text{prob}(s, q, \eta) = \int_0^\infty \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}(u, \mathbf{q}_{q, s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q, s}^{\eta+t} := 0]) \right] \right\} dt .$$

Proof. We first consider the case for the family $\{\text{prob}_n\}_{n \in \mathbb{N}_0}$. Let $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$. By definition, one can verify that $\text{prob}_0(s, q, \eta) = \mathbf{1}_F(q)$ and $\text{prob}_n(s, q, \eta) = 1$ with $q \in F$. Given $n \in \mathbb{N}_0$ and $q \in Q \setminus F$, we obtain directly from Theorem 7.1 that

$$\text{prob}_{n+1}(s, q, \eta) = \int_0^\infty \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}_n(u, \mathbf{q}_{q, s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q, s}^{\eta+t} := 0]) \right] \right\} dt .$$

Then the case for prob follows from Monotone Convergence Theorem (cf. Theorem 3.2). \square

In this chapter, we study the approximation of the function prob.

8.4 Mathematical Technicalities

In this section, we prepare several mathematical technicalities to derive the differential characterization for the function prob (cf. Section 8.5). In detail, we review several equivalence relations on clock valuations [1] and the product region graph between CTMC and DTA [24], and prove that the function prob is Lipschitz continuous.

Below we fix a CTMC $\mathcal{M} = (S, \mathbf{E}, \mathbf{P})$, a DTA $\mathcal{A} = (Q, \Sigma, \mathcal{X}, \Delta, F)$ and a labelling function $\mathcal{L} : S \rightarrow \Sigma$. For a clock $x \in \mathcal{X}$, we denote by T_x the largest integer c that appears in some guard $x \bowtie c$ of \mathcal{A} , by T_{\max} the integer $\max_{x \in \mathcal{X}} T_x$, and by \mathbf{E}_{\max} the value $\max\{\mathbf{E}(s) \mid s \in S\}$.

8.4.1 Equivalence Relations on Clock Valuations

Definition 8.9. [1] *Two valuations $\eta, \eta' \in \text{Val}(\mathcal{X})$ are guard-equivalent, denoted by $\eta \equiv_{\text{gd}} \eta'$, if they satisfy the following conditions:*

1. for all $x \in \mathcal{X}$, $\eta(x) > T_x$ iff $\eta'(x) > T_x$;
2. for all $x \in \mathcal{X}$, if $\eta(x) \leq T_x$ and $\eta'(x) \leq T_x$, then (i) $\text{int}(\eta(x)) = \text{int}(\eta'(x))$ and (ii) $\text{frac}(\eta(x)) > 0$ iff $\text{frac}(\eta'(x)) > 0$.

where $\text{int}(\cdot)$ and $\text{frac}(\cdot)$ is the integral and fractional part of a real number, respectively. Moreover, η and η' are equivalent, denoted by $\eta \equiv \eta'$, if (i) $\eta \equiv_{\text{gd}} \eta'$ and (ii) for all $x, y \in \mathcal{X}$, if $\eta(x), \eta'(x) \leq T_x$ and $\eta(y), \eta'(y) \leq T_y$, then $\text{frac}(\eta(x)) \bowtie \text{frac}(\eta(y))$ iff $\text{frac}(\eta'(x)) \bowtie \text{frac}(\eta'(y))$ for all $\bowtie \in \{<, =, >\}$. We will call equivalence classes of \equiv regions. Given a region $[\eta]_{\equiv}$, we say that $[\eta]_{\equiv}$ is marginal if for some clock $x \in \mathcal{X}$, $\eta(x) \leq T_x$ and $\text{frac}(\eta(x)) = 0$.

In other words, equivalence classes of \equiv_{gd} are captured by (i) a boolean vector over \mathcal{X} which indicates whether $\eta(x) > T_x$ or not, (ii) an integer vector which indicates the integral parts of clocks in $\{x \in \mathcal{X} \mid \eta(x) \leq T_x\}$, and (iii) a boolean vector which indicates whether $\eta(x)$ is an integer when $\eta(x) \leq T_x$. Equivalence classes of \equiv is further captured by a linear order on the set $\{x \in \mathcal{X} \mid \eta(x) \leq T_x\}$ w.r.t the ordering on the values $\{\text{frac}(\eta(x))\}_{\eta(x) \leq T_x}$. Below we state some basic properties of \equiv_{gd} and \equiv .

Proposition 8.2. [1] *The following properties on \equiv_{gd} and \equiv hold:*

1. both \equiv_{g} and \equiv is an equivalence relation over clock valuations, and has finite index;

2. if $\eta \equiv_{\text{gd}} \eta'$, then η and η' satisfy the same set of guards that appear in the rules of \mathcal{A} ;
3. if $\eta \equiv \eta'$, then
 - for all $t > 0$, there exists $t' > 0$ such that $\eta + t \equiv \eta' + t'$, and
 - for all $t' > 0$, there exists $t > 0$ such that $\eta + t \equiv \eta' + t'$;
4. if $\eta \equiv \eta'$, then $\eta[X := 0] \equiv \eta'[X := 0]$ for all subsets $X \subseteq \mathcal{X}$;
5. for all $\eta \in \text{Val}(\mathcal{X})$ and $X \subseteq \mathcal{X}$, $\{\eta'[X := 0] \mid \eta' \in [\eta]_{\equiv}\}$ is a region.

Besides these two equivalence notions, we define another finer equivalence notion as follows.

Definition 8.10. *Two valuations $\eta, \eta' \in \text{Val}(\mathcal{X})$ are bound-equivalent, denoted by $\eta \equiv_{\text{bd}} \eta'$, if for all $x \in \mathcal{X}$, either $\eta(x) > T_x$ and $\eta'(x) > T_x$, or $\eta(x) = \eta'(x)$.*

Intuitively, $\eta \equiv_{\text{bd}} \eta'$ means that the behaviours of η and η' are almost equal, as they either have the same value on a clock, or their values on a clock x are both over the relevance threshold T_x . It is straightforward to verify that \equiv_{bd} is an equivalence relation.

The following lemma specifies the relation between \equiv_{bd} and prob (cf. Barbot *et al.* [9]). Below we present an alternative proof.

Proposition 8.3. *Let $s \in S$, $q \in Q$ and $\eta, \eta' \in \text{Val}(\mathcal{X})$. If $\eta \equiv_{\text{bd}} \eta'$, then $\text{prob}(s, q, \eta) = \text{prob}(s, q, \eta')$.*

Proof. We prove that $\text{Paths}(s, q, \eta) = \text{Paths}(s, q, \eta')$. Assume that $\pi \in \text{Paths}(s, q, \eta)$. Then the run $\mathcal{A}_{q, \eta}(\mathcal{L}(\pi)) = \{(q_n, \eta_n)(\pi\langle n \rangle, \mathcal{L}(\pi[n]))\}_{n \geq 0}$ satisfies that $q_n \in F$ for some $n \geq 0$. Denote

$$\mathcal{A}_{q, \eta'}(\mathcal{L}(\pi)) = \{(q'_n, \eta'_n)(\pi\langle n \rangle, \mathcal{L}(\pi[n]))\}_{n \geq 0} .$$

We prove inductively on n that $q_n = q'_n$ and $\eta_n \equiv_{\text{bd}} \eta'_n$ for all $n \geq 0$. This would imply $\pi \in \text{Paths}(s, q, \eta')$. The inductive proof can be carried out by the fact that $\eta_n \equiv_{\text{bd}} \eta'_n$ implies

- $\eta_n + \pi\langle n \rangle \equiv_{\text{bd}} \eta'_n + \pi\langle n \rangle$ and
- $(\eta_n + \pi\langle n \rangle)[X := 0] \equiv_{\text{bd}} (\eta'_n + \pi\langle n \rangle)[X := 0]$ for all $X \subseteq \mathcal{X}$.

Thus $\text{Paths}(s, q, \eta) \subseteq \text{Paths}(s, q, \eta')$ due to the arbitrary choice of π . The other direction of inclusion can be proved symmetrically. \square

In the following, we further introduce a useful proposition.

Proposition 8.4. *For all $\eta \in \text{Val}(\mathcal{X})$, there exists a real number $t_{\text{rht}} > 0$ such that $\eta + t \equiv \eta + t'$ for all $t, t' \in (0, t_{\text{rht}})$. For all $\eta \in \text{Val}(\mathcal{X})$, if $\eta(x) > 0$ for all $x \in \mathcal{X}$, then there exists a real number $t_{\text{lft}} > 0$ such that $\eta - t \equiv \eta - t'$ for all $t, t' \in (0, t_{\text{lft}})$.*

Proof. Define $Z := \{\eta(x) - T_x \mid x \in \mathcal{X} \text{ and } \eta(x) > T_x\}$. If $\eta(x) > T_x$ for all clocks x , then we can choose t_{rht} to be any positive real number and $t_{\text{lft}} = \min Z$. Below we assume that there exists $x \in \mathcal{X}$ such that $\eta(x) \leq T_x$.

Define $Z' := \{\text{frac}(\eta(x)) \mid x \in \mathcal{X} \text{ and } \eta(x) \leq T_x\}$. Let c_1, c_2 be the maximum and the minimum value of Z' , respectively. Note that $0 \leq c_2 \leq c_1 < 1$. Then we choose t_{rhs} to be $1 - c_1$. The choice of t_{lft} subjects to the two cases below.

1. $c_2 > 0$. Then we can choose t_{lft} to be $\min(\{c_2\} \cup Z)$.
2. $c_2 = 0$. If $Z' = \{c_2\}$ then we can choose $t_{\text{lft}} = \min(\{1\} \cup Z)$. Otherwise, let $c' > c_2$ be the second minimum value in Z' . Then we can choose $t_{\text{lft}} = \min(\{c'\} \cup Z)$.

It is straightforward to verify that $t_{\text{rht}}, t_{\text{lft}}$ satisfy the desired property. \square

Let $\eta \in \text{Val}(\mathcal{X})$. We denote η^+ to be a representative in $\{\eta + t \mid t \in (0, t_{\text{rht}})\}$, and η^- to be a representative in $\{\eta - t \mid t \in (0, t_{\text{lft}})\}$, where $t_{\text{rht}}, t_{\text{lft}}$ are specified in Proposition 8.4. The choice among the representatives will be irrelevant due to the fact that all of them are equivalent under \equiv . Note that if a region $[\eta]_{\equiv}$ is not marginal, then $[\eta]_{\equiv} = [\eta^+]_{\equiv} = [\eta^-]_{\equiv}$; this is because one can always find a real number δ small enough such that integral parts of all values $(\eta \pm \delta)(x)$ ($x \in \mathcal{X}$) do not change, and the corresponding fractional parts remain positive.

8.4.2 Product Region Graph

In this part, we define a qualitative variation of the product region graph proposed by Chen *et al.* [24], mainly to derive a qualitative property of the function prob.

Definition 8.11. *The product region graph $\mathcal{G}^{\mathcal{M}, \mathcal{A}} = (V^{\mathcal{M}, \mathcal{A}}, E^{\mathcal{M}, \mathcal{A}})$ of \mathcal{M} and \mathcal{A} is a digraph defined as follows:*

- $V^{\mathcal{M}, \mathcal{A}} = S \times Q \times (\text{Val}(\mathcal{X}) / \equiv)$;
- $((s, q, [\eta]_{\equiv}), (s', q', [\eta']_{\equiv})) \in E^{\mathcal{M}, \mathcal{A}}$ iff (i) $\mathbf{P}(s, s') > 0$ and (ii) there exists $t \in \mathbb{R}_{>0}$ such that $[\eta + t]_{\equiv}$ is not a marginal region and $(q', \eta') = \kappa((q, \eta), (t, \mathcal{L}(s)))$.

A vertex $(s, q, [\eta]_{\equiv}) \in V^{\mathcal{M}, \mathcal{A}}$ is called final if $q \in F$.

By Proposition 8.2, Definition 8.11 is well-defined. We omit the superscript ‘ \mathcal{M}, \mathcal{A} ’ in “ $\mathcal{G}^{\mathcal{M}, \mathcal{A}} = (V^{\mathcal{M}, \mathcal{A}}, E^{\mathcal{M}, \mathcal{A}})$ ” if the underlying context is clear.

The following lemma states the relationship between prob and the product region graph. Below we define

$$Z_\eta := \{0, 1\} \cup \{\text{frac}(\eta(x)) \mid x \in \mathcal{X} \text{ and } \eta(x) \leq T_x\}$$

for each $\eta \in \text{Val}(\mathcal{X})$. Intuitively, Z_η captures the relevant fractional values of η .

Proposition 8.5. *For all $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$, $\text{prob}(s, q, \eta) > 0$ iff $(s, q, [\eta]_{\equiv})$ can reach some final vertex in \mathcal{G} .*

Proof. “only if”: It is clear that $\text{prob}(s, q, \eta) > 0$ iff $\text{prob}_n(s, q, \eta) > 0$ for some $n \in \mathbb{N}_0$. We prove by induction on n that for all $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$, if $\text{prob}_n(s, q, \eta) > 0$ then $(s, q, [\eta]_{\equiv})$ can reach some final vertex in \mathcal{G} . The base step $n = 0$ is straightforward from the definition. Assume that $\text{prob}_{n+1}(s, q, \eta) > 0$ with $q \notin F$. (If $q \in F$, then $(s, q, [\eta]_{\equiv})$ itself is a final vertex.) By Theorem 8.1, we can obtain that

$$\int_0^\infty \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}_n(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) \right] \right\} dt > 0 . \quad (8.1)$$

Consider the regions traversed by $\eta + t$ when t goes from 0 to ∞ . Denote $Z_\eta = \{w_0, \dots, w_m\}$ (cf. the line before the proposition) such that $m \geq 1$ and $w_i > w_{i+1}$ for all $0 \leq i < m$. Note that $w_0 = 1$ and $w_m = 0$. We divide $[0, \infty)$ into open intervals $(0, 1), (1, 2), (T_{\max} - 1, T_{\max}), (T_{\max}, \infty)$. For each integer $k < T_{\max}$, we further divide the interval $(k, k+1)$ into the following open sub-intervals, excluding a finite number of isolating points:

$$(k+1 - w_0, k+1 - w_1), \dots, (k+1 - w_{m-1}, k+1 - w_m) .$$

Then we define the cluster \mathcal{I} of intervals by:

$$\mathcal{I} := \{(k+1 - w_i, k+1 - w_{i+1}) \mid 0 \leq k < T_{\max}, 0 \leq i < m\} \cup \{(T_{\max}, \infty)\} .$$

One can verify by definition that for all $I \in \mathcal{I}$ and $t', t \in I$, $\eta + t \equiv \eta + t'$. In other words, $[\eta + t]_{\equiv}$ does not change when t is restricted to one of the intervals in \mathcal{I} . By Inequality (8.1), there exists $I \in \mathcal{I}$ such that

$$\int_I \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}_n(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) \right] \right\} dt > 0 .$$

This means that there exists $u^* \in S$ and $t^* \in I$ such that

$$\mathbf{P}(s, u^*) \cdot \text{prob}_n(u^*, \mathbf{q}_{q,s}^{\eta+t^*}, (\eta+t^*)[\mathbf{X}_{q,s}^{\eta+t^*} := 0]) > 0 .$$

Since I is nonempty, $[\eta + t^*]_{\equiv}$ is not a marginal region. Thus there exists an edge from $(s, q, [\eta]_{\equiv})$ to $(u^*, \mathbf{q}_{q,s}^{\eta+t^*}, [(\eta+t^*)[\mathbf{X}_{q,s}^{\eta+t^*} := 0]])_{\equiv}$ in \mathcal{G} . By the induction hypothesis, the vertex $(u^*, \mathbf{q}_{q,s}^{\eta+t^*}, [(\eta+t^*)[\mathbf{X}_{q,s}^{\eta+t^*} := 0]])_{\equiv}$ can reach some final vertex \mathcal{G} . Thus $(s, q, [\eta]_{\equiv})$ can reach some final vertex in \mathcal{G} .

“if”: Assume $(s, q, [\eta]_{\equiv})$ can reach some final vertex in \mathcal{G} . Let the reachability route be

$$(s, q, [\eta]_{\equiv}) = (s_n, q_n, [\eta_n]_{\equiv}) \rightarrow (s_{n-1}, q_{n-1}, [\eta_{n-1}]_{\equiv}) \cdots \rightarrow (s_0, q_0, [\eta_0]_{\equiv})$$

with $q_0 \in F$. We prove inductively on $m \leq n$ that $\text{prob}_m(s_m, q_m, \eta') > 0$ for all $\eta' \in [\eta_m]_{\equiv}$. The case $m = 0$ is straightforward. Assume that $\text{prob}_m(s_m, q_m, \eta') > 0$ for all $\eta' \in [\eta_m]_{\equiv}$. Let $\eta'' \in [\eta_{m+1}]_{\equiv}$ be an arbitrary clock valuation. By $(s_{m+1}, q_{m+1}, [\eta_{m+1}]_{\equiv}) \rightarrow (s_m, q_m, [\eta_m]_{\equiv})$, $\mathbf{P}(s_{m+1}, s_m) > 0$ and there exists $\eta'_{m+1} \in [\eta_{m+1}]_{\equiv}$, $\eta'_m \in [\eta_m]_{\equiv}$ and $t \in \mathbb{R}_{>0}$ such that $[\eta_{m+1} + t]_{\equiv}$ is not marginal and $(q_m, \eta'_m) = \kappa((q_{m+1}, \eta'_{m+1}), (t, \mathcal{L}(s_{m+1})))$. By $\eta'' \equiv \eta'_{m+1}$, there exists $t' \in \mathbb{R}_{>0}$ such that $\eta'' + t' \equiv \eta'_{m+1} + t$, which further implies that

$$(\eta'' + t')[\mathbf{X}_{q_{m+1}, s_{m+1}}^{\eta'' + t'} := 0] \equiv (\eta'_{m+1} + t)[\mathbf{X}_{q_{m+1}, s_{m+1}}^{\eta'_{m+1} + t} := 0] \quad (= \eta'_m) .$$

By the fact that $[\eta'' + t']_{\equiv}$ is not marginal, there exists an interval $I \subseteq \mathbb{R}_{\geq 0}$ with positive length such that for all $\tau \in I$, $\eta'' + \tau \equiv \eta'' + t'$ and

$$(\eta'' + \tau)[\mathbf{X}_{q_{m+1}, s_{m+1}}^{\eta'' + \tau} := 0] \equiv (\eta'' + t')[\mathbf{X}_{q_{m+1}, s_{m+1}}^{\eta'' + t'} := 0] \equiv \eta'_m .$$

Thus by induction hypothesis,

$$\text{prob}_m \left(s_m, q_m, (\eta'' + \tau)[\mathbf{X}_{q_{m+1}, s_{m+1}}^{\eta'' + \tau} := 0] \right) > 0$$

for all $\tau \in I$. Below we prove by contradiction that

$$\int_I \left\{ f_{\mathbf{E}(s)}(\tau) \cdot \left[\sum_{u \in S} \mathbf{P}(s_{m+1}, u) \cdot \text{prob}_m \left(u, q_m, (\eta'' + \tau)[\mathbf{X}_{q_{m+1}, s_{m+1}}^{\eta'' + \tau} := 0] \right) \right] \right\} d\tau > 0 . \quad (\dagger)$$

This would imply $\text{prob}_{m+1}(s_{m+1}, q_{m+1}, \eta'') > 0$ by Theorem 8.1. Suppose that (\dagger) does not hold. By Proposition 7.2, the integrand function $h^* : I \rightarrow \mathbb{R}$ at the left hand side of (\dagger) is measurable. Thus, there exists a sequence $\{h_k\}_{k \in \mathbb{N}}$ of simple functions such that $h_k \uparrow h^*$ and $\int_I h_k(\tau) d\tau \uparrow \int_I h^*(\tau) d\tau$ (cf. Proposition 3.1). Since (\dagger) does not hold (i.e., $\int_I h^*(\tau) d\tau = 0$), for all $k \in \mathbb{N}$, the set $C_k := \{\tau \in I \mid h_k(\tau) > 0\}$ has Borel measure zero. Thus $\bigcup_{k \in \mathbb{N}} C_k$ also has Borel measure zero. It follows that h^* is zero on a non-empty subset of I ; contradiction to the induction hypothesis. \square

8.4.3 Lipschitz Continuity

In this part, we prove that the function prob is Lipschitz continuous. More specifically, we prove that all functions that satisfy a boundness condition related to \equiv_{bd} and the system of integral equations specified in Theorem 8.1 are Lipschitz continuous. The Lipschitz continuity will be fundamental to our differential characterization and the error bound of our approximation result.

Theorem 8.2. *Let $h : S \times Q \times \text{Val}(\mathcal{X}) \rightarrow [0, 1]$ be a function which satisfies the following conditions for all $s \in S$, $q \in Q$ and $\eta, \eta' \in \text{Val}(\mathcal{X})$:*

- if $\eta \equiv_{\text{bd}} \eta'$ then $h(s, q, \eta) = h(s, q, \eta')$;
- if $q \in F$ then $h(s, q, \eta) = 1$, otherwise $h(s, q, \eta)$ is equal to the integral

$$\int_0^\infty \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot h(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) \right] \right\} dt .$$

For all $s \in S$, $q \in Q$ and $\eta, \eta' \in \text{Val}(\mathcal{X})$, if $\|\eta - \eta'\|_\infty < 1$ then

$$|h(s, q, \eta) - h(s, q, \eta')| \leq M_1 \cdot \|\eta - \eta'\|_\infty ,$$

where $M_1 := |\mathcal{X}| \cdot \mathbf{E}_{\max} T_{\max} \cdot e^{\mathbf{E}_{\max} T_{\max}}$.

Proof. If $q \in F$, then the result follows directly from $h(s, q, \eta) = h(s, q, \eta') = 1$. From now on, we assume that $q \notin F$. To prove the theorem, it suffices to prove that

$$|h(s, q, \eta) - h(s, q, \eta')| \leq \mathbf{E}_{\max} T_{\max} \cdot e^{\mathbf{E}_{\max} T_{\max}} \cdot \|\eta - \eta'\|_\infty ,$$

provided that $\|\eta - \eta'\|_\infty < 1$ and η, η' differ exactly on one clock, i.e. $|\{x \in \mathcal{X} \mid \eta(x) \neq \eta'(x)\}| = 1$. To this end we define $\delta(\epsilon)$ for each $\epsilon \in (0, 1)$ as follows:

$$\delta(\epsilon) := \sup \{ |h(s, q, \eta) - h(s, q, \eta')| \mid s \in S, q \in Q, \eta, \eta' \in \text{Val}(\mathcal{X}), \\ \|\eta - \eta'\|_\infty \leq \epsilon \text{ and } \eta, \eta' \text{ differ only on one clock.} \}$$

Note that for all $\eta, \eta' \in \text{Val}(\mathcal{X})$ and $X \subseteq \mathcal{X}$:

- if η and η' differ at most on one clock, then so are $\eta[X := 0]$ and $\eta'[X := 0]$;
- $\|\eta[X := 0] - \eta'[X := 0]\|_\infty \leq \|\eta - \eta'\|_\infty$.

Let $\epsilon \in (0, 1)$. Let $s \in S$, $q \in Q \setminus F$ and $\eta, \eta' \in \text{Val}(\mathcal{X})$ which satisfies $\|\eta - \eta'\|_\infty \leq \epsilon < 1$ and differ exactly on the clock x , i.e., $\eta(x) \neq \eta'(x)$ and $\eta(y) = \eta'(y)$ for all $y \neq x$. W.l.o.g, we can assume that $\eta(x) < \eta'(x)$. We clarify two cases below.

Case 1: $\text{int}(\eta(x)) = \text{int}(\eta'(x))$. Then by $\eta(x) < \eta'(x)$, $\text{frac}(\eta(x)) < \text{frac}(\eta'(x))$. Consider the ‘‘trajectories’’ of $\eta + t$ and $\eta' + t$ when t goes from 0 to ∞ . We divide $[0, \infty)$ into open integer intervals $(0, 1), (1, 2), \dots, (T_{\max} - 1, T_{\max})$ and (T_{\max}, ∞) . For each $n < T_{\max}$, we further divide the interval $(n, n + 1)$ into the following open sub-intervals:

$$(n, n + 1 - \text{frac}(\eta'(x))), (n + 1 - \text{frac}(\eta'(x)), n + 1 - \text{frac}(\eta(x))), \\ (n + 1 - \text{frac}(\eta(x)), n + 1) .$$

One can observe that for $t \in (n, n+1 - \text{frac}(\eta'(x))) \cup (n+1 - \text{frac}(\eta(x)), n+1)$, we have $\eta + t \equiv_{\text{gd}} \eta' + t$, which implies that $\eta + t$ and $\eta' + t$ satisfies the same set of guards in \mathcal{A} . However for $t \in (n+1 - \text{frac}(\eta'(x)), n+1 - \text{frac}(\eta(x)))$, it may be the case that $\eta + t \not\equiv_{\text{gd}} \eta' + t$ due to their difference on the clock x . Thus the total length for t within $(n, n+1)$ such that $\eta + t \not\equiv_{\text{gd}} \eta' + t$ is no greater than $|\eta(x) - \eta'(x)|$. Then we have (\dagger):

$$\begin{aligned} \delta_n &:= \left| \int_n^{n+1} f_{\mathbf{E}(s)}(t) \cdot \left\{ \sum_{u \in S} \left[\mathbf{P}(s, u) \cdot h(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) \right. \right. \right. \\ &\quad \left. \left. \left. - \mathbf{P}(s, u) \cdot h(u, \mathbf{q}_{q,s}^{\eta'+t}, (\eta'+t)[\mathbf{X}_{q,s}^{\eta'+t} := 0]) \right] \right\} dt \right| \\ &\leq \int_n^{n+1} \{f_{\mathbf{E}(s)}(t) \cdot \delta(\epsilon)\} dt + \mathbf{E}(s) \cdot e^{-\mathbf{E}(s) \cdot n} \cdot |\eta(x) - \eta'(x)| \\ &\leq \delta(\epsilon) \cdot \int_n^{n+1} \{f_{\mathbf{E}(s)}(t)\} dt + \mathbf{E}(s) \cdot e^{-\mathbf{E}(s) \cdot n} \cdot \epsilon \end{aligned}$$

Note that for all $t \in (T_{\max}, \infty)$ and $X \subseteq \mathcal{X}$,

$$(\eta + t)[X := 0] \equiv_{\text{bd}} (\eta' + t)[X := 0] ;$$

this implies (from Proposition 8.3)

$$h(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) = h(u, \mathbf{q}_{q,s}^{\eta'+t}, (\eta'+t)[\mathbf{X}_{q,s}^{\eta'+t} := 0]) .$$

Therefore we have (\ddagger):

$$\begin{aligned} &|h(s, q, \eta) - h(s, q, \eta')| \\ &\leq \sum_{n=0}^{T_{\max}-1} \delta_n \\ &\leq \delta(\epsilon) \cdot \int_0^{T_{\max}} \{f_{\mathbf{E}(s)}(t)\} dt + \mathbf{E}(s) \cdot \epsilon \cdot \sum_{n=0}^{T_{\max}-1} e^{-\mathbf{E}(s) \cdot n} \\ &\leq \delta(\epsilon) \cdot (1 - e^{-\mathbf{E}(s) \cdot T_{\max}}) + \epsilon \cdot \mathbf{E}(s) \cdot T_{\max} \\ &\leq \delta(\epsilon) \cdot (1 - e^{-\mathbf{E}_{\max} \cdot T_{\max}}) + \epsilon \cdot \mathbf{E}_{\max} \cdot T_{\max} \end{aligned}$$

Case 2: $\text{int}(\eta(x)) < \text{int}(\eta'(x))$. By $|\eta(x) - \eta'(x)| < 1$, we have $\text{int}(\eta(x)) + 1 = \text{int}(\eta'(x))$ and $\text{frac}(\eta'(x)) < \text{frac}(\eta(x))$. Similarly, we divide the interval $[0, \infty)$ into integer intervals $(0, 1), (1, 2), \dots, (T_{\max} - 1, T_{\max}), (T_{\max}, \infty)$. And in each interval $(n, n+1)$, we divide the interval into the following open sub-intervals:

$$\begin{aligned} &(n, n+1 - \text{frac}(\eta(x))), (n+1 - \text{frac}(\eta(x)), n+1 - \text{frac}(\eta'(x))), \\ &\quad (n+1 - \text{frac}(\eta'(x)), n+1) . \end{aligned}$$

If $t \in (n + 1 - \text{frac}(\eta(x)), n + 1 - \text{frac}(\eta'(x)))$, then $\eta + t \equiv_{\text{gd}} \eta' + t$. And if t lies in either $(n, n + 1 - \text{frac}(\eta(x)))$ or $(n + 1 - \text{frac}(\eta'(x)), n + 1)$, then it may be the case that $\eta + t \not\equiv_{\text{gd}} \eta' + t$. Thus the total length within $(n, n + 1)$ such that $\eta + t \not\equiv_{\text{gd}} \eta' + t$ is still smaller than $|\eta(x) - \eta'(x)|$. Therefore we can apply the analysis (†) and (‡), to obtain that

$$|h(s, q, \eta) - h(s, q, \eta')| \leq \delta(\epsilon) \cdot (1 - e^{-\mathbf{E}_{\max} T_{\max}}) + \epsilon \cdot \mathbf{E}_{\max} \cdot T_{\max}$$

Thus by the definition of $\delta(\epsilon)$, we obtain

$$\delta(\epsilon) \leq \delta(\epsilon) \cdot (1 - e^{-\mathbf{E}_{\max} T_{\max}}) + \epsilon \cdot \mathbf{E}_{\max} \cdot T_{\max}$$

which implies $\delta(\epsilon) \leq \epsilon \cdot e^{\mathbf{E}_{\max} T_{\max}} \cdot \mathbf{E}_{\max} \cdot T_{\max}$. By letting $\epsilon = \|\eta - \eta'\|_{\infty}$, we obtain the desired result. \square

The Lipschitz continuity of the function prob follows directly from Theorem 8.2.

Corollary 8.1. *For all $s \in S$, $q \in Q$ and $\eta, \eta' \in \text{Val}(\mathcal{X})$, if $\|\eta - \eta'\|_{\infty} < 1$ then*

$$|\text{prob}(s, q, \eta) - \text{prob}(s, q, \eta')| \leq M_1 \cdot \|\eta - \eta'\|_{\infty}$$

where M_1 is defined as in Theorem 8.2.

Proof. Directly from Theorem 8.1 and Theorem 8.2. \square

By Corollary 8.1, we can further prove that prob is the unique solution of a revised system of integral equations from the one specified in Theorem 8.2.

Theorem 8.3. *The function prob is the unique solution of the following system of integral equations on $h : S \times Q \times \text{Val}(\mathcal{X}) \rightarrow [0, 1]$:*

1. for all $s \in S$, $q \in Q$ and $\eta, \eta' \in \text{Val}(\mathcal{X})$, if $\eta \equiv_{\text{bd}} \eta'$ then $h(s, q, \eta) = h(s, q, \eta')$;
2. for all $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$, (i) if $q \in F$ then $h(s, q, \eta) = 1$, and (ii) if $(s, q, [\eta]_{\equiv})$ cannot reach a final vertex in \mathcal{G} then $h(s, q, \eta) = 0$;
3. for all $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$, if $(s, q, [\eta]_{\equiv})$ can reach a final vertex in \mathcal{G} and $q \notin F$, then $h(s, q, \eta)$ equals

$$\int_0^{\infty} \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot h(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta + t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) \right] \right\} dt .$$

Proof. By Theorem 8.1, Proposition 8.3 and Proposition 8.5, prob satisfies the integral-equation system. Below we prove that the integral-equation system has only one solution.

We first prove that if $h : S \times Q \times \text{Val}(\mathcal{X}) \rightarrow [0, 1]$ satisfies the integral-equation system, then h satisfies the prerequisite of Theorem 8.2. Let h be such a function which satisfies the integral-equation system. We only

need to consider the case for $h(s, q, \eta)$ where $(s, q, [\eta]_{\equiv})$ cannot reach a final vertex in \mathcal{G} . Assume (s, q, η) such that $(s, q, [\eta]_{\equiv})$ cannot reach a final vertex in \mathcal{G} . Note that $h(s, q, \eta) = 0$. From the proof of Proposition 8.5, we can construct a cluster \mathcal{I} of disjoint open intervals which satisfies the following conditions: (i) $\bigcup \mathcal{I} \subseteq \mathbb{R}_{\geq 0}$; (ii) $\mathbb{R}_{\geq 0} \setminus \bigcup \mathcal{I}$ is a finite set; (iii) for all $I \in \mathcal{I}$ and $t, t' \in I$, $\eta + t \equiv \eta + t'$. Choose any $t \in \bigcup \mathcal{I}$ and $u \in S$ such that $\mathbf{P}(s, u) > 0$. Then $(u, \mathbf{q}_{q,s}^{\eta+t}, [(\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]]_{\equiv})$ cannot reach some final vertex in \mathcal{G} since $[\eta+t]_{\equiv}$ is not marginal. Thus $h(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) = 0$ by Proposition 8.5. It follows that

$$\int_0^{\infty} \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot h(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) \right] \right\} dt = 0 ,$$

which shows that h satisfies the prerequisite of Theorem 8.2.

Then suppose that $h_1, h_2 : S \times Q \times \text{Val}(\mathcal{X}) \rightarrow [0, 1]$ are two distinct solutions of the integral-equation system. Define $h := |h_1 - h_2|$. By Theorem 8.2, h is Lipschitz continuous on $\text{Val}(\mathcal{X})$. Furthermore, by the fact that $\eta \equiv_{\text{bd}} \eta'$ implies $h(s, q, \eta) = h(s, q, \eta')$, the image of h can be obtained on $S \times Q \times \prod_{x \in \mathcal{X}} [0, T_x]$. Thus the maximum value

$$M := \sup\{h(s, q, \eta) \mid (s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})\}$$

can be reached. Since $h_1 \neq h_2$, $M > 0$. Denote

$$\mathcal{C} := \{(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X}) \mid h(s, q, \eta) = M\} .$$

We first prove that (†): for all $(s, q, \eta) \in \mathcal{C}$ and all edge $(s, q, [\eta]_{\equiv}) \rightarrow (s', q', [\eta']_{\equiv})$ in \mathcal{G} , there exists $\eta'' \in [\eta']_{\equiv}$ such that $(s', q', \eta'') \in \mathcal{C}$.

Consider an arbitrary $(s, q, \eta) \in \mathcal{C}$. Since $M > 0$, $(s, q, [\eta]_{\equiv})$ can reach a final vertex in \mathcal{G} and $q \notin F$. From the proof of Proposition 8.5, we can divide $[0, \infty)$ into a cluster \mathcal{I} of disjoint open intervals, disregarding only a finite number of isolating points, such that $[\eta + \tau]_{\equiv}$ ($\tau \in I$) does not change for all $I \in \mathcal{I}$. Thus $h(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0])$ is piecewise continuous on $t \in \mathbb{R}_{\geq 0}$, for all $u \in S$. Note that

$$\begin{aligned} & h(s, q, \eta) \\ & \leq \int_0^{\infty} \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot h(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) \right] \right\} dt \\ & \leq \int_0^{\infty} \left\{ f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot h(s, q, \eta) \right] \right\} dt \\ & = h(s, q, \eta) \end{aligned}$$

By the piecewise continuity, $h(u, \mathbf{q}_{q,s}^{\eta+t}, (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) = M$ whenever $t \in \bigcup \mathcal{I}$ and $\mathbf{P}(s, u) > 0$. Note that $[0, \infty) \setminus (\bigcup \mathcal{I})$ is finite. Thus for all edge

$(s, q, [\eta]_{\equiv}) \rightarrow (s', q', [\eta']_{\equiv})$ in \mathcal{G} , there exists $t \in \bigcup \mathcal{I}$ such that $\mathbf{q}_{q,s}^{\eta+t} = q'$ and $(\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0] \in [\eta']_{\equiv}$. It follows from $(s', q', (\eta+t)[\mathbf{X}_{q,s}^{\eta+t} := 0]) \in \mathcal{C}$ that (\dagger) holds.

Let $(s, q, \eta) \in \mathcal{C}$. Then there exists a path

$$(s, q, [\eta]_{\equiv}) = (s_0, q_0, [\eta_0]_{\equiv}) \rightarrow (s_1, q_1, [\eta_1]_{\equiv}) \cdots (s_n, q_n, [\eta_n]_{\equiv})$$

in \mathcal{G} with $q_n \in F$. However, from (\dagger) , one can prove through induction that there exists $\eta' \in [\eta_n]_{\equiv}$ such that $(s_n, q_n, \eta') \in \mathcal{C}$, which implies $q_n \notin F$. Contradiction. Thus $M = 0$ and the solution to the integral-equation system is unique. \square

8.5 A Differential Characterization

In this section, we present a differential characterization for the function prob (cf. Section 8.3). Below we fix a CTMC $\mathcal{M} = (S, \mathbf{E}, \mathbf{P})$, a DTA $\mathcal{A} = (Q, \Sigma, \mathcal{X}, \Delta, F)$ and a labelling function $\mathcal{L} : S \rightarrow \Sigma$.

The following definition introduces our notion of derivative.

Definition 8.12. *Given a function $h : S \times Q \times \text{Val}(\mathcal{X}) \rightarrow [0, 1]$, we denote by $\nabla_{\mathbf{1}}^+ h$ (resp. $\nabla_{\mathbf{1}}^- h$) the right directional derivative (resp. the left directional derivative) of h along the direction $\mathbf{1}$ (the vector whose coordinates are all one) if the derivative exists. Formally, we define:*

- $\nabla_{\mathbf{1}}^+ h(s, q, \eta) := \lim_{t \rightarrow 0^+} \frac{1}{t} \cdot (h(s, q, \eta + t) - h(s, q, \eta))$, if the limit exists;
- $\nabla_{\mathbf{1}}^- h(s, q, \eta) := \lim_{t \rightarrow 0^+} \frac{1}{t} \cdot (h(s, q, \eta) - h(s, q, \eta - t))$, if $\eta(x) > 0$ for all $x \in \mathcal{X}$ and the limit exists;

for each $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$.

Below we calculate these directional derivatives for the function prob.

Theorem 8.4. *For all $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$ with $q \notin F$, $\nabla_{\mathbf{1}}^+ \text{prob}(s, q, \eta)$ exists, and $\nabla_{\mathbf{1}}^- \text{prob}(s, q, \eta)$ exists given that $\eta(x) > 0$ for all $x \in \mathcal{X}$. Furthermore,*

$$\nabla_{\mathbf{1}}^+ \text{prob}(s, q, \eta) = \mathbf{E}(s) \cdot \text{prob}(s, q, \eta) - \mathbf{E}(s) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}(u, \mathbf{q}_{q,s}^{\eta^+}, \eta[\mathbf{X}_{q,s}^{\eta^+} := 0]) \right]$$

and

$$\nabla_{\mathbf{1}}^- \text{prob}(s, q, \eta) = \mathbf{E}(s) \cdot \text{prob}(s, q, \eta) - \mathbf{E}(s) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}(u, \mathbf{q}_{q,s}^{\eta^-}, \eta[\mathbf{X}_{q,s}^{\eta^-} := 0]) \right]$$

whenever $\nabla_{\mathbf{1}}^- \text{prob}(s, q, \eta)$ exists.

Proof. Let $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$ with $q \notin F$. To ease the notation, we temporarily denote by $h[s, q, \eta]$ the function

$$\tau \mapsto f_{\mathbf{E}(s)}(\tau) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}(u, \mathbf{q}_{q,s}^{\eta+\tau}, (\eta + \tau)[\mathbf{X}_{q,s}^{\eta+\tau} := 0]) \right] .$$

We first prove the case for $\nabla_{\mathbf{1}}^+ \text{prob}(s, q, \eta)$. By Theorem 8.1,

$$\text{prob}(s, q, \eta) = \int_0^\infty h[s, q, \eta](\tau) \, d\tau$$

and

$$\text{prob}(s, q, \eta + t) = \int_0^\infty h[s, q, \eta + t](\tau) \, d\tau$$

for all $t \geq 0$. Note that $h[s, q, \eta], h[s, q, \eta + t]$ is Riemann integratable since it is piecewise continuous on τ (cf. the cluster \mathcal{I} constructed in the proof of Proposition 8.5). Thus, $\int_0^\infty h[s, q, \eta](\tau) \, d\tau$ and $\int_0^\infty h[s, q, \eta + t](\tau) \, d\tau$ can be deemed as Riemann integral. By the variable substitution $\tau' = t + \tau$, we have

$$\text{prob}(s, q, \eta + t) = e^{\mathbf{E}(s) \cdot t} \cdot \int_t^\infty h[s, q, \eta](\tau) \, d\tau$$

for all $t \geq 0$. Then we have

$$\begin{aligned} & \text{prob}(s, q, \eta + t) - \text{prob}(s, q, \eta) \\ &= e^{\mathbf{E}(s) \cdot t} \cdot \int_t^\infty h[s, q, \eta](\tau) \, d\tau - \int_0^\infty h[s, q, \eta](\tau) \, d\tau \\ &= (e^{\mathbf{E}(s) \cdot t} - 1) \cdot \int_t^\infty h[s, q, \eta](\tau) \, d\tau - \int_0^t h[s, q, \eta](\tau) \, d\tau \end{aligned}$$

By Proposition 8.4, there exists $t_{\text{rhs}} > 0$ such that $\mathbf{q}_{q,s}^{\eta+\tau}$ and $\mathbf{X}_{q,s}^{\eta+\tau}$ does not change for $\tau \in (0, t_{\text{rhs}})$. Thus $h[s, q, \eta]$ is continuous on τ when $t \in (0, t_{\text{rhs}})$. Moreover, the point $\tau = 0$ can be continuously redefined on $h[s, q, \eta]$. Thus by L'Hôpital's Rule, we obtain

$$\begin{aligned} & \nabla_{\mathbf{1}}^+ \text{prob}(s, q, \eta) = \\ & \mathbf{E}(s) \cdot \text{prob}(s, q, \eta) - \mathbf{E}(s) \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}(u, \mathbf{q}_{q,s}^{\eta+}, \eta[\mathbf{X}_{q,s}^{\eta+} := 0]) \right] . \end{aligned}$$

Then we handle the case for $\nabla_{\mathbf{1}}^- \text{prob}(s, q, \eta)$ given that $\eta(x) > 0$ for all

$x \in \mathcal{X}$. For $t \in (0, \min\{\eta(x) \mid x \in \mathcal{X}\})$, we have

$$\begin{aligned}
& \text{prob}(s, q, \eta) - \text{prob}(s, q, \eta - t) \\
&= \text{prob}(s, q, (\eta - t) + t) - \text{prob}(s, q, \eta - t) \\
&= (e^{\mathbf{E}(s) \cdot t} - 1) \cdot \int_t^\infty h[s, q, \eta - t](\tau) \, d\tau - \int_0^t h[s, q, \eta - t](\tau) \, d\tau \\
&= (1 - e^{-\mathbf{E}(s) \cdot t}) \cdot \int_0^\infty h[s, q, \eta](\tau) \, d\tau - e^{-\mathbf{E}(s) \cdot t} \cdot \mathbf{E}(s) \cdot \\
&\quad \int_0^t \left\{ e^{\mathbf{E}(s) \cdot \tau} \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}(u, \mathbf{q}_{q,s}^{\eta-\tau}, (\eta - \tau)[\mathbf{X}_{q,s}^{\eta-\tau} := 0]) \right] \right\} \, d\tau
\end{aligned}$$

where the last step is obtained by performing the variable substitutions $\tau' = \tau - t$ in the first integral and $\tau' = t - \tau$ in the second integral. By Proposition 8.4, there exists $t_{\text{fft}} > 0$ such that $\mathbf{q}_{q,s}^{\eta-t}$ and $\mathbf{X}_{q,s}^{\eta-t}$ does not change for $t \in (0, t_{\text{fft}})$. Thus the integrand function in the integral

$$\int_0^t \left\{ e^{\mathbf{E}(s) \cdot \tau} \cdot \left[\sum_{u \in S} \mathbf{P}(s, u) \cdot \text{prob}(u, \mathbf{q}_{q,s}^{\eta-\tau}, (\eta - \tau)[\mathbf{X}_{q,s}^{\eta-\tau} := 0]) \right] \right\} \, d\tau$$

is continuous on τ when $t \in (0, t_{\text{fft}})$; furthermore, the point $\tau = 0$ can be continuously redefined on the integrand function. Thus we can also apply L'Hôpital's Rule and obtain the desired result. \square

Remark 8.2. Note that if $[\eta]_{\equiv}$ is not marginal, then $[\eta^+]_{\equiv} = [\eta^-]_{\equiv} = [\eta]_{\equiv}$. This tells us that $\nabla_{\mathbf{1}} \text{prob}(s, q, \eta)$ exists when $[\eta]_{\equiv}$ is not marginal, i.e., $\nabla_{\mathbf{1}}^+ \text{prob}(s, q, \eta) = \nabla_{\mathbf{1}}^- \text{prob}(s, q, \eta)$. \square

Theorem 8.4 will serve as a basis for our approximation algorithm.

8.6 Approximation Algorithm

In this section, we present an algorithm that approximates the function prob through finite approximation schemes. We establish our approximation scheme based on Theorem 8.4. Then we prove that our approximation scheme converges to prob with a derived error bound.

Below we fix a CTMC $\mathcal{M} = (S, \mathbf{E}, \mathbf{P})$, a DTA $\mathcal{A} = (Q, \Sigma, \mathcal{X}, \Delta, F)$ and a labelling function $\mathcal{L} : S \rightarrow \Sigma$. For computational purpose, we assume that all numerical values in \mathcal{M} are rational.

Given clock valuation η and $t \geq 0$, we define $\eta \oplus t \in \text{Val}(\mathcal{X})$ by:

$$(\eta \oplus t)(x) := \min\{T_x, \eta(x) + t\}$$

for all $x \in \mathcal{X}$. Note that $\eta \oplus 0 = \eta$ iff $\eta(x) \leq T_x$ for all clocks x . Intuitively, ' \oplus ' is a variant operator of '+' which takes into account the relevance threshold for clock valuations.

To ease the notation, we extend \oplus , $+$, $[\cdot]_{\equiv}$, $\mathbf{E}(\cdot)$, $\mathbf{P}(\cdot, \cdot)$ and \cdot^+ to triples $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$ in a straightforward fashion, as follows:

- $(s, q, \eta) \oplus t := (s, q, \eta \oplus t)$ and $(s, q, \eta) + t = (s, q, \eta + t)$;
- $[(s, q, \eta)]_{\equiv} := (s, q, [\eta]_{\equiv})$ is a vertex of \mathcal{G} ;
- $\mathbf{E}(s, q, \eta) := \mathbf{E}(s)$, and $\mathbf{P}((s, q, \eta), u) := \mathbf{P}(s, u)$ for each $u \in S$;
- $(s, q, \eta)^+ := (s, q, \eta^+)$.

Moreover, we say that $[(s, q, \eta)]_{\equiv}$ is *marginal* if $[\eta]_{\equiv}$ is marginal; we denote the triple $(u, \mathbf{q}_{q,s}^{\eta^+}, \eta[\mathbf{X}_{q,s}^{\eta^+} := 0]) \in S \times Q \times \text{Val}(\mathcal{X})$ by $(s, q, \eta)_u^+$, for $u \in S$.

Remark 8.3. *By Corollary 8.1 and Proposition 8.3, one obtains easily that $\text{prob}(v + t) = \text{prob}(v \oplus t)$ for all $v \in S \times Q \times \text{Val}(\mathcal{X})$ and $t \geq 0$.*

8.6.1 Approximation Schemes

In this part, we establish our approximation scheme in two steps: firstly, we discretize the hypercube $\prod_{x \in \mathcal{X}} [0, T_x] \subseteq \text{Val}(\mathcal{X})$ into small grids; secondly, we establish our approximation scheme by building constraints between these discrete values through finite difference methods. By the Lipschitz continuity (Corollary 8.1) and Proposition 8.3, we don't need to consider clock valuations outside $\prod_{x \in \mathcal{X}} [0, T_x]$. The discretization is as follows.

Definition 8.13. *Let $m \in \mathbb{N}$. The set of discretized points \mathbf{D}_m is defined as follows:*

$$\mathbf{D}_m := \{ \mathbf{h}[(s, q, \eta)] \mid (s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X}), \\ \text{and for all } x \in \mathcal{X}, \eta(x) \in [0, T_x] \text{ and } \eta(x) \cdot m \text{ is an integer.} \} .$$

Intuitively, \mathbf{D}_m results from discretizing the hypercube $\prod_{x \in \mathcal{X}} [0, T_x] \subseteq \text{Val}(\mathcal{X})$ with discretization step m . Note that the point $\mathbf{h}[v]$ is in fact $v \in S \times Q \times \text{Val}(\mathcal{X})$. To simplify the presentation, sometimes we do not distinguish between the point $\mathbf{h}[v]$ and the element v .

Below we fix a $m \in \mathbb{N}$ and define $\rho := m^{-1}$ (i.e., the discretization step size). Based on Theorem 8.4, we render our basic approximation scheme as follows.

Definition 8.14. *The approximation scheme Υ_m consists of the discrete points in \mathbf{D}_m and a system of linear equations on \mathbf{D}_m . The system of linear equations contains one of the following equations for each $\mathbf{h}[v] \in \mathbf{D}_m$:*

- $\mathbf{h}[v] = 0$ if $[v]_{\equiv}$ (as a vertex of \mathcal{G}) cannot reach a final vertex in \mathcal{G} ;
- $\mathbf{h}[v] = 1$ if $[v]_{\equiv}$ is a final vertex in \mathcal{G} ;

- If $[v]_{\equiv}$ can reach a final vertex in \mathcal{G} and itself is not a final vertex, then

$$\frac{\mathbf{h}[v \oplus \rho] - \mathbf{h}[v]}{\rho} = \mathbf{E}(v) \cdot \mathbf{h}[v] - \mathbf{E}(v) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \mathbf{h}[v_u^+] .$$

Intuitively, $\mathbf{h}[v]$ approximates $\text{prob}(v)$ for $\mathbf{h}[v] \in \mathbf{D}_m$. We relate elements in \mathbf{D}_m by using $\nabla_{\mathbf{1}}^+ \mathbf{h}$ described in Theorem 8.4.

Remark 8.4. We would like to remark that Υ_m does not have initial values from which we can approximate prob incrementally. This phenomenon is an inherent feature from the automata-theoretic definition of the problem (cf. Definition 8.8). In contrast to approximation schemes with initial values, this increases the difficulty to solve the problem and leads to the following two problems: one is whether Υ_m has a solution, or even a unique solution; the other is the error bound $\max\{|h^*[v] - \text{prob}(v)| \mid \mathbf{h}[v] \in \mathbf{D}_m\}$ provided that h^* is the unique solution of Υ_m .

Below we first derive the *error bound* of Υ_m , which is defined as the maximal error of each linear equality when we substitute all $\mathbf{h}[v]$'s by the corresponding $\text{prob}(v)$'s. We would like to note that generally the error bound of an approximation scheme without initial value does not imply any information of the error between the solution to the approximation scheme and the function which the approximation scheme approximates.

Proposition 8.6. For all $\mathbf{h}[v] \in \mathbf{D}_m$, if $[v]_{\equiv}$ is not a final vertex and can reach some final vertex (in \mathcal{G}) then

$$\left| \frac{1}{\rho} \cdot (\text{prob}(v \oplus \rho) - \text{prob}(v)) - \nabla_{\mathbf{1}}^+ \text{prob}(v) \right| < M_2 \cdot \rho ,$$

where $M_2 := 2 \cdot \mathbf{E}_{\max} \cdot M_1$.

Proof. Let $\mathbf{h}[v] \in \mathbf{D}_m$ such that $[v]_{\equiv}$ is not a final vertex and can reach some final vertex in \mathcal{G} . Since $\mathbf{h}[v] \in \mathbf{D}_m$, the function $f[v] : [0, \rho] \rightarrow [0, 1]$, defined by $f[v](t) := \text{prob}(v \oplus t) (= \text{prob}(v + t))$, is continuous on $[0, \rho]$ and is differentiable on $(0, \rho)$ (cf. Theorem 8.4 and Remark 8.2). By Lagrange's Mean-Value Theorem, there exists $\rho' \in (0, \rho)$ such that $\frac{1}{\rho} \cdot (\text{prob}(v \oplus \rho) - \text{prob}(v)) = \left(\frac{d}{dt} f[v]\right)(\rho')$. By Theorem 8.4, we obtain

$$\left(\frac{d}{dt} f[v]\right)(\rho') = \mathbf{E}(v) \cdot \text{prob}(v + \rho') - \mathbf{E}(v) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \text{prob}((v + \rho')_u^+)$$

and

$$\nabla_{\mathbf{1}}^+ \text{prob}(v) = \mathbf{E}(v) \cdot \text{prob}(v) - \mathbf{E}(v) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \text{prob}(v_u^+) .$$

Let $v = (s, q, \eta)$. By the definition of ρ , $[\eta + \rho']_{\equiv} = [\eta^+]_{\equiv}$. Then by Corollary 8.1, we obtain the desired result. \square

To analyze Υ_m , we further define several auxiliary approximation schemes. Below we define subsets $\mathbf{B}_m, \mathbf{B}_m^{\max}$ of \mathbf{D}_m as follows:

$$\begin{aligned} \mathbf{B}_m &= \{h[v] \in \mathbf{D}_m \mid [v]_{\equiv} \text{ is not final and can reach some final vertex in } \mathcal{G}\}; \\ \mathbf{B}_m^{\max} &= \{h[v] \in \mathbf{B}_m \mid v = (s, q, \eta) \text{ and } \eta(x) = T_x \text{ for all } x \in \mathcal{X}\}. \end{aligned}$$

Intuitively, \mathbf{B}_m contains the discrete points to be determined in \mathbf{D}_m , and \mathbf{B}_m^{\max} is the extreme discrete points in \mathbf{B}_m . For each $h[v] \in \mathbf{B}_m$, we define $N_v \in \mathbb{N}_0$ to be the minimum number such that either $h[v \oplus (N_v \cdot \rho)] \in \mathbf{B}_m^{\max}$, or $[v \oplus (N_v \cdot \rho)]_{\equiv}$ cannot reach some final vertex in \mathcal{G} .

Below we transform Υ_m into an equivalent form.

Definition 8.15. *The approximation scheme Υ'_m consists of the discrete points \mathbf{D}_m , and the system of linear equations which contains one of the following linear equalities for each $h[v] \in \mathbf{D}_m$:*

- $h[v] = 0$ if $[v]_{\equiv}$ cannot reach a final vertex in \mathcal{G} ;
- $h[v] = 1$ if $[v]_{\equiv}$ is a final vertex of \mathcal{G} .
- if $h[v] \in \mathbf{B}_m \setminus \mathbf{B}_m^{\max}$, then

$$h[v] = \frac{1}{1 + \rho \cdot \mathbf{E}(v)} \cdot h[v \oplus \rho] + \frac{\rho \cdot \mathbf{E}(v)}{1 + \rho \cdot \mathbf{E}(v)} \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot h[v_u^+] \quad (8.2)$$

- if $h[v] \in \mathbf{B}_m^{\max}$ then $h[v] = \sum_{u \in S} \mathbf{P}(v, u) \cdot h[v_u^+] .$

It is clear that Υ'_m is an equivalent reformulation of Υ_m . Note that the case for $h[v] \in \mathbf{B}_m^{\max}$ in Υ'_m is derived from the fact that $v \oplus \rho = v$. The error bound of the approximation scheme Υ'_m is as follows. To ease the notation, in the following we define

$$d_{\rho, v} := \frac{1}{1 + \rho \cdot \mathbf{E}(v)} .$$

Proposition 8.7. *For all $h[v] \in \mathbf{B}_m^{\max}$,*

$$\text{prob}(v) = \sum_{u \in S} \mathbf{P}(v, u) \cdot \text{prob}(v_u^+) .$$

For all $h[v] \in \mathbf{B}_m \setminus \mathbf{B}_m^{\max}$, the value

$$\left| \text{prob}(v) - \left(d_{\rho, v} \cdot \text{prob}(v \oplus \rho) + (1 - d_{\rho, v}) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \text{prob}(v_u^+) \right) \right|$$

is smaller than $M_2 \cdot \rho^2$.

Proof. The case for $h[v] \in \mathbf{B}_m^{\max}$ is due to the fact that $\nabla_1^+ \text{prob}(v) = 0$ by definition. The case $h[v] \in \mathbf{B}_m \setminus \mathbf{B}_m^{\max}$ can be directly derived from the statement of Proposition 8.6, using the fact that Υ'_m is a direct equivalent reformulation of Υ_m . \square

In the following, we unfold Υ'_m into another equivalent form Υ''_m .

Definition 8.16. *The approximation scheme Υ''_m consists of the discrete points \mathcal{D}_m , and one of the following linear equality for each $\mathbf{h}[v] \in \mathcal{D}_m$:*

- $\mathbf{h}[v] = 0$ if $[v]_{\equiv}$ cannot reach some final vertex in \mathcal{G} , and $\mathbf{h}[v] = 1$ if $[v]_{\equiv}$ is a final vertex in \mathcal{G} ;
- if $\mathbf{h}[v] \in \mathcal{B}_m \setminus \mathcal{B}_m^{\max}$, then

$$\mathbf{h}[v] = \sum_{l=0}^{N_v-1} \left\{ d_{\rho,v}^l \cdot (1 - d_{\rho,v}) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \mathbf{h}[(v \oplus (l \cdot \rho))_u^+] \right\} + d_{\rho,v}^{N_v} \cdot f(v) \quad (8.3)$$

- where $f(v) := 0$ if $[v \oplus (N_v \cdot \rho)]_{\equiv}$ cannot reach some final vertex in \mathcal{G} , and $f(v) := \sum_{u \in S} \mathbf{P}(v, u) \cdot \mathbf{h}[(v \oplus (N_v \cdot \rho))_u^+]$ if $\mathbf{h}[v \oplus (N_v \cdot \rho)] \in \mathcal{B}_m^{\max}$;
- if $\mathbf{h}[v] \in \mathcal{B}_m^{\max}$ then $\mathbf{h}[v] = \sum_{u \in S} \mathbf{P}(v, u) \cdot \mathbf{h}[v_u^+]$.

Intuitively, Υ''_m is obtained by unfolding $\mathbf{h}[v \oplus \rho]$ further in Equation (8.2). In the following, we prove that Υ'_m and Υ''_m are equivalent, i.e., they have the same set of solutions.

To ease the notation, we describe Υ''_m by a matrix equation $\mathbf{v} = \mathbf{A}\mathbf{v} + \mathbf{b}$ where \mathbf{v} is the vector over \mathcal{B}_m to be solved, $\mathbf{b} : \mathcal{B}_m \rightarrow \mathbb{R}$ is a vector and $\mathbf{A} : \mathcal{B}_m \times \mathcal{B}_m \rightarrow \mathbb{R}$ is a matrix. For example, for $\mathbf{h}[v] \in \mathcal{B}_m \setminus \mathcal{B}_m^{\max}$, the row $\mathbf{A}(\mathbf{h}[v], -)$ is specified by the coefficients on $\mathbf{h}[v'] \in \mathcal{B}_m$ in Equation (8.3); the value $\mathbf{b}(\mathbf{h}[v])$ is the sum of the coefficients on $\mathcal{D}_m \setminus \mathcal{B}_m$ in Equation (8.3). The exact permutation among \mathcal{B}_m is irrelevant. Analogously, we describe Υ'_m by a matrix equation $\mathbf{v} = \mathbf{C}\mathbf{v} + \mathbf{d}$.

Proposition 8.8. *Υ'_m and Υ''_m are equivalent, i.e., they have the same set of solutions.*

Proof. We first consider the direction from Υ'_m to Υ''_m . However, it is clear that Υ''_m is obtained directly from Υ'_m , by expanding $\mathbf{h}[v \oplus \rho]$ repeatedly in Equation (8.2) whenever $v \oplus \rho \in \mathcal{B}_m \setminus \mathcal{B}_m^{\max}$.

Then we consider the non-trivial direction from Υ''_m to Υ'_m . Let $\{\mathbf{h}[v] \mid \mathbf{h}[v] \in \mathcal{D}_m\}$ be a solution of Υ''_m . Define the function $\hat{\mathbf{h}}$ by: $\hat{\mathbf{h}} := 1_{\mathcal{B}_m} \cdot \mathbf{h}$. We prove that for all $\mathbf{h}[v] \in \mathcal{B}_m \setminus \mathcal{B}_m^{\max}$ and all $0 \leq n < N_v$,

$$\begin{aligned} \mathbf{h}[v] &= \sum_{l=0}^n \left\{ d_{\rho,v}^l \cdot (1 - d_{\rho,v}) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \hat{\mathbf{h}}[(v \oplus (l \cdot \rho))_u^+] \right\} \\ &\quad + \sum_{l=0}^n \left\{ d_{\rho,v}^l \cdot \mathbf{d}(\mathbf{h}[v \oplus (l \cdot \rho)]) \right\} \\ &\quad + d_{\rho,v}^{n+1} \cdot \mathbf{h}[v \oplus ((n+1) \cdot \rho)] . \end{aligned} \quad (8.4)$$

We prove this by induction on $N_v - n$. The case when $n = N_v - 1$ is directly specified by Υ_m'' . Let $0 \leq n < n + 1 < N_v$ and assume that Equation 8.4 holds when $N_{v'} - n' < N_v - n$. By induction hypothesis, we have

$$\begin{aligned} h[v] &= \sum_{l=0}^{n+1} \left\{ d_{\rho,v}^l \cdot (1 - d_{\rho,v}) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \hat{h} [(v \oplus (l \cdot \rho))_u^+] \right\} \\ &\quad + \sum_{l=0}^{n+1} \left\{ d_{\rho,v}^l \cdot \mathbf{d} (h[v \oplus (l \cdot \rho)]) \right\} \\ &\quad + d_{\rho,v}^{n+2} \cdot h[v \oplus ((n+2) \cdot \rho)] . \end{aligned}$$

Then, we have (†):

$$\begin{aligned} h[v] &= \sum_{l=0}^n \left\{ d_{\rho,v}^l \cdot (1 - d_{\rho,v}) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \hat{h} [(v \oplus l \cdot \rho)_u^+] \right\} \\ &\quad + \sum_{l=0}^n \left\{ d_{\rho,v}^l \cdot \mathbf{d} (h[v \oplus (l \cdot \rho)]) \right\} \\ &\quad + d_{\rho,v}^{n+1} \cdot \left\{ (1 - d_{\rho,v}) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \hat{h} [(v \oplus ((n+1) \cdot \rho))_u^+] \right. \\ &\quad \left. + \mathbf{d} (h[v \oplus ((n+1) \cdot \rho)]) + d_{\rho,v} \cdot h[v \oplus ((n+2) \cdot \rho)] \right\} . \end{aligned}$$

Note that $N_{v \oplus (n+1) \cdot \rho} - 0 = N_v - (n+1) < N_v - n$. Thus by the induction hypothesis,

$$\begin{aligned} &h[v \oplus ((n+1) \cdot \rho)] \\ &= (1 - d_{\rho,v}) \cdot \sum_{u \in S} \left\{ \mathbf{P}(v, u) \cdot \hat{h} [(v \oplus ((n+1) \cdot \rho))_u^+] \right\} \\ &\quad + \mathbf{d} (h[v \oplus ((n+1) \cdot \rho)]) + d_{\rho,v} \cdot h[v \oplus ((n+2) \cdot \rho)] . \end{aligned}$$

Thus, Equation (8.4) holds when we substitute $h[v \oplus ((n+1) \cdot \rho)]$ into (†). By taking $n = 0$ in Equation (8.4), we obtain that $\{h[v] \mid h[v] \in \mathbf{D}_m\}$ is a solution of Υ_m' . \square

Below we derive the error bound of Υ_m'' .

Proposition 8.9. *The error bound of the approximation scheme Υ_m'' is $M_3 \cdot \rho$, where $M_3 := T_{\max} \cdot M_2$.*

Proof. We only need to consider $h[v] \in \mathbf{B}_m \setminus \mathbf{B}_m^{\max}$. By Proposition 8.7,

$$\begin{aligned} &\left| \text{prob}(v) - \left(d_{\rho,v} \cdot \text{prob}(v \oplus \rho) + (1 - d_{\rho,v}) \cdot \sum_{u \in S} \mathbf{P}(v, u) \cdot \text{prob}(v_u^+) \right) \right| \\ &< M_2 \cdot \rho^2 . \end{aligned} \tag{8.5}$$

Expanding $\text{prob}[v \oplus \rho]$ one step further in (8.5) will result in another error of $d_{\rho,v} \cdot M_2 \cdot \rho^2$. By repeated expansion up to N_v ($\leq T_{\max} \cdot \rho^{-1}$) steps, the error bound of Υ''_m is no greater than $M_2 \cdot \rho^2 \cdot \sum_{n=0}^{N_v-1} d_{\rho,v}^n$, which is smaller than $M_2 \cdot T_{\max} \cdot \rho$. \square

8.6.2 Error-Bound Analysis

In this part, we analyze the error between prob and the solution to the approximation scheme Υ_m (or equivalently $\Upsilon'_m, \Upsilon''_m$). We fix some $m \in \mathbb{N}$ and $\rho = m^{-1}$. We define

- $\mathbf{E}_{\min} = \min\{\mathbf{E}(s) \mid s \in S\}$ and
- $p_{\min} = \min\{\mathbf{P}(s, u) \mid s, u \in S \text{ and } \mathbf{P}(s, u) > 0\}$.

We note that $\mathbf{E}_{\min} > 0$ by Definition 8.1.

Recall that we describe Υ''_m by $\mathbf{v} = \mathbf{A}\mathbf{v} + \mathbf{b}$ and Υ'_m by $\mathbf{v} = \mathbf{C}\mathbf{v} + \mathbf{d}$ in the previous part. Below we analyse the equation $\mathbf{v} = \mathbf{A}\mathbf{v} + \mathbf{b}$.

We first reproduce (on CTMCs and DTAs) the notions of δ -separateness and δ -wideness, which is originally discovered by Brazdil *et al.* [17] on semi-Markov processes and DTAs. These notions are used to derive the error bound. The following definition introduces a related technical notion of transition relations.

Definition 8.17. For each $t \in \mathbb{R}_{>0}$, the binary relation \xrightarrow{t} on $S \times Q \times \text{Val}(\mathcal{X})$ is defined by: $((s, q, \eta), (u, q', \eta')) \in \xrightarrow{t}$ iff $\mathbf{P}(s, u) > 0$, $[\eta+t]_{\equiv}$ is not marginal, and $\kappa((q, \eta), (t, \mathcal{L}(s))) = (q', \eta')$.

We write “ $(s, q, \eta) \xrightarrow{t} (u, q', \eta')$ ” instead of “ $((s, q, \eta), (u, q', \eta')) \in \xrightarrow{t}$ ”; we will call an $(s, q, \eta) \xrightarrow{t} (u, q', \eta')$ a *transit*. The notions of separateness and wideness are defined as follows. Below we recall that

$$Z_{\eta} = \{0, 1\} \cup \{\text{frac}(\eta(x)) \mid x \in \mathcal{X} \text{ and } \eta(x) \leq T_x\} .$$

Definition 8.18. Let $\delta \in \mathbb{R}_{>0}$. A clock valuation η is δ -separated if for all $d_1, d_2 \in Z_{\eta}$, either $d_1 = d_2$ or $|d_1 - d_2| \geq \delta$. A transit $(s, q, \eta) \xrightarrow{t} (u, q', \eta')$ is δ -wide if $t \geq \delta$ and for all $\tau \in (t - \delta, t + \delta)$, $\eta + \tau \equiv \eta + t$. Furthermore, a sequence of transits

$$(s_0, q_0, \eta_0) \xrightarrow{t_1} (s_1, q_1, \eta_1) \dots \xrightarrow{t_n} (s_n, q_n, \eta_n) \quad (n \geq 1)$$

from (s_0, q_0, η_0) to (s_n, q_n, η_n) , where $(s_i, q_i, \eta_i) \xrightarrow{t_{i+1}} (s_{i+1}, q_{i+1}, \eta_{i+1})$ for all $0 \leq i < n$, is δ -wide if $(s_i, q_i, \eta_i) \xrightarrow{t_{i+1}} (s_{i+1}, q_{i+1}, \eta_{i+1})$ is δ -wide for all $0 \leq i < n$.

Intuitively, a transit is δ -wide if one can adjust the transit by up to δ time units, while keeping the DTA rule used to obtain this transition (cf. Definition 8.7).

To explain the usage of separateness and wideness, we further import the following technical notions.

Definition 8.19. *A collection \mathcal{I} of disjoint non-empty open intervals in $\mathbb{R}_{\geq 0}$ is an open partition (of $[0, 1]$) if $\bigcup \mathcal{I} \subseteq [0, 1]$ and $[0, 1] \setminus (\bigcup \mathcal{I})$ is a finite set. Given a non-empty open interval $I \subseteq [0, 1]$ with $I = (c_1, c_2)$ and a $t \in \mathbb{R}_{\geq 0}$, we define $I \circ t$ to be the (possibly empty) interval $(\text{frac}(c_1+t), \text{frac}(c_2+t))$.*

The following proposition illustrates the usage of separateness and wideness, which is also the counterpart of the one on semi-Markov processes and DTA [16, 17]. In the following, $|\mathcal{G}|$ denotes the number of vertices of \mathcal{G} .

Proposition 8.10. *Let $\delta \in \mathbb{R}_{> 0}$. For all $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$, if η is δ -separated, $q \notin F$ and $(s, q, [\eta]_{\equiv})$ can reach some final vertex in \mathcal{G} , then there exists a $\frac{\delta}{|\mathcal{G}|}$ -wide and at most $|\mathcal{G}|$ -long sequence of transits from (s, q, η) to some (s', q', η') with $q' \in F$.*

Proof. Let $(s, q, \eta) \in S \times Q \times \text{Val}(\mathcal{X})$ such that $(s, q, [\eta]_{\equiv})$ is not final and can reach some final vertex in \mathcal{G} . Assume that η is δ -separated. Then there exists a path

$$(s, q, [\eta]_{\equiv}) = (s_1, q_1, \mathbf{r}_1) \rightarrow \cdots \rightarrow (s_n, q_n, \mathbf{r}_n)$$

in \mathcal{G} such that $1 < n \leq |\mathcal{G}|$ and $q_n \in F$. Firstly, we inductively construct a sequence of transits

$$(s, q, \eta) = (s_1, q_1, \eta_1) \xrightarrow{t_1} \cdots \xrightarrow{t_{n-1}} (s_n, q_n, \eta_n)$$

such that $\eta_i \in \mathbf{r}_i$ for all $1 \leq i \leq n$, while maintaining the following structures:

- two open partitions $\mathcal{I}'_i, \mathcal{I}_i$ with $Z_{\eta_i} \subseteq [0, 1] \setminus \bigcup \mathcal{I}'_i$, for each $1 \leq i \leq n$;
- a bijection $\varphi_i : \mathcal{I}'_i \rightarrow \mathcal{I}_i$, for each $1 \leq i \leq n$;
- two intervals $(c_1^i, c_2^i) \in \mathcal{I}_i$, $(w_1^i, w_2^i) \in \mathcal{I}'_i$ with $\varphi_i((w_1^i, w_2^i)) = (c_1^i, c_2^i)$, for each $1 \leq i \leq n-1$;
- a $c^i \in (c_1^i, c_2^i)$, for each $1 \leq i \leq n-1$.

Initially, we set $\eta_1 = \eta$ and $\mathcal{I}_1 = \mathcal{I}'_1 = \{(w_j, w_{j+1}) \mid 0 \leq j < k\}$, where the unique finite sequence $\{w_j\}_{0 \leq j \leq k}$ satisfies that $Z_\eta = \{w_0, w_1, \dots, w_k\}$ and $w_j < w_{j+1}$ for all $0 \leq j < k$; (note that $w_0 = 0$ and $w_k = 1$); we let φ_1 be the identity mapping.

Assume that the sequence of transits until (s_i, q_i, η_i) , together with $\mathcal{I}'_i, \mathcal{I}_i$ and φ_i , are constructed. Since $(s_i, q_i, \mathbf{r}_i) \rightarrow (s_{i+1}, q_{i+1}, \mathbf{r}_{i+1})$ in \mathcal{G} , there

exists $t_i > 0$ such that $[\eta_i + t_i]_{\equiv}$ is not marginal, $q_{i+1} = \mathbf{q}_{q_i, s_i}^{\eta_i + t_i}$ and $(\eta_i + t_i)[\mathbf{X}_{q_i, s_i}^{\eta_i + t_i} := 0] \in \mathfrak{r}_{i+1}$. Since $[\eta_i + t_i]_{\equiv}$ is not marginal, we can adjust t_i , while without changing $[\eta_i + t_i]_{\equiv}$, so that we can choose $(w_1^i, w_2^i) \in \mathcal{I}'_i$ with $1 \in (w_1^i + \text{frac}(t_i), w_2^i + \text{frac}(t_i))$. Define $(c_1^i, c_2^i) := \varphi_i((w_1^i, w_2^i))$ and choose $c^i \in (c_1^i, c_2^i)$ arbitrarily (e.g., $c^i := \frac{1}{2} \cdot (c_1^i + c_2^i)$). Then we set $\eta_{i+1} := (\eta_i + t_i)[\mathbf{X}_{q_i, s_i}^{\eta_i + t_i} := 0] \in \mathfrak{r}_{i+1}$, and split $\mathcal{I}_i, \mathcal{I}'_i$ as follows:

$$\begin{aligned} \mathcal{I}_{i+1} &:= (\mathcal{I}_i - \{\varphi_i((w_1^i, w_2^i))\}) \cup \{(c_1^i, c^i), (c^i, c_2^i)\} ; \\ \mathcal{I}'_{i+1} &:= \{(w_1^i + \text{frac}(t_i), 1), (0, \text{frac}(w_2^i + \text{frac}(t_i)))\} \\ &\quad \cup \{I \diamond t_i \mid I \in \mathcal{I}'_i - \{(w_1^i, w_2^i)\}\} . \end{aligned}$$

The mapping $\varphi_{i+1} : \mathcal{I}'_{i+1} \rightarrow \mathcal{I}_{i+1}$ is defined as follows:

$$\begin{aligned} \varphi_{i+1}((w_1^i + \text{frac}(t_i), 1)) &= (c_1^i, c^i) , \varphi_{i+1}((0, \text{frac}(w_2^i + \text{frac}(t_i)))) = (c^i, c_2^i); \\ \varphi_{i+1}(I \diamond t_i) &= \varphi_i(I) \text{ for all } I \in \mathcal{I}'_i - \{(w_1^i, w_2^i)\} . \end{aligned}$$

Intuitively, we record by \mathcal{I}'_i every possible splitting point caused by a transit which may make the wideness of $\eta_i + t_i$ decrease, and we record the un-timed splitting information by \mathcal{I}_i , where the correspondence between them is maintained by φ_i .

Since $n \leq |\mathcal{G}|$, at most $|\mathcal{G}| - 1$ splitting operations occur on intervals during the inductive construction described above. Based on this point, we inductively construct a new $\frac{\delta}{|\mathcal{G}|}$ -wide sequence of transits

$$(s, q, \eta) = (s'_1, q'_1, \eta'_1) \xrightarrow{t'_1} \dots \xrightarrow{t'_{n-1}} (s'_n, q'_n, \eta'_n)$$

such that $s_i = s'_i$, $q_i = q'_i$ and $\eta'_i \equiv \eta_i$ for all $1 \leq i \leq n$, while maintaining an open partition \mathcal{I}''_i and a bijection $\psi_i : \mathcal{I}''_i \rightarrow \mathcal{I}'_i$ for each $1 \leq i \leq n$, which satisfy the following conditions for all $1 \leq i \leq n$:

1. $Z_{\eta'_i} \subseteq [0, 1] \setminus \bigcup \mathcal{I}''_i$;
2. for all $I_1, I_2 \in \mathcal{I}''_i$, $\sup I_1 \leq \inf I_2$ iff $\sup \psi_i(I_1) \leq \inf \psi_i(I_2)$;
3. for all $x \in \mathcal{X}$, $(d'_1, d'_2) \in \mathcal{I}''_i$ and $(d_1, d_2) \in \mathcal{I}'_i$, if $\psi_i((d'_1, d'_2)) = (d_1, d_2)$ then (i) $\eta'_i(x) = d'_1$ iff $\eta_i(x) = d_1$ and (ii) $\eta'_i(x) = d'_2$ iff $\eta_i(x) = d_2$.

Intuitively, the new inductive construction maintains the order on the fractional values in the previous sequence of transits, while adjusting the timed information on each transit to meet the wideness requirement. In the new inductive construction, we define N_I (for each $I \in \bigcup_{i=1}^n \mathcal{I}_i$) to be the number of splittings on the interval I in the previous inductive construction, i.e., $N_I := |\{I' \in \mathcal{I}_n \mid I' \subseteq I\}| - 1$.

Initially, we set $(s'_1, q'_1, \eta'_1) := (s, q, \eta)$, $\mathcal{I}''_1 = \mathcal{I}'_1$ and ψ_1 to be the identity mapping. Assume that the sequence of transits until (s'_i, q'_i, η'_i) is constructed. Let $(d_1^i, d_2^i) \in \mathcal{I}''_i$ be such that $\psi_i((d_1^i, d_2^i)) = (w_1^i, w_2^i)$. We

choose t'_i such that $\text{int}(t'_i) = \text{int}(t_i)$, $1 \in (d_1^i + \text{frac}(t'_i), d_2^i + \text{frac}(t'_i))$ and the length of $(d_1^i + \text{frac}(t'_i), 1)$ (resp. $(0, \text{frac}(d_2^i + \text{frac}(t'_i)))$) is no smaller than $(N_{(c_1^i, c^i)} + 1) \cdot \frac{\delta}{|\mathcal{G}|}$ (resp. $(N_{(c^i, c_2^i)} + 1) \cdot \frac{\delta}{|\mathcal{G}|}$). Then we set

$$(s'_{i+1}, q'_{i+1}, \eta'_{i+1}) = (s_{i+1}, \mathbf{q}_{q'_i, s'_i}^{\eta'_i + t'_i}, (\eta'_i + t'_i)[\mathbf{X}_{q'_i, s'_i}^{\eta'_i + t'_i} := 0])$$

and split \mathcal{I}_i'', ψ_i as follows:

$$\begin{aligned} \mathcal{I}_{i+1}'' &:= \{I \diamond t'_i \mid I \in \mathcal{I}_i'' - \{(d_1^i, d_2^i)\}\} \\ &\quad \cup \{(d_1^i + \text{frac}(t'_i), 1), (0, \text{frac}(d_2^i + \text{frac}(t'_i)))\}; \\ \psi_{i+1}(I' \diamond t'_i) &= I \diamond t_i \text{ whenever } \psi_i(I') = I \text{ and } I' \neq (d_1^i, d_2^i); \\ \psi_{i+1}((d_1^i + \text{frac}(t'_i), 1)) &= (w_1^i + \text{frac}(t_i), 1); \\ \psi_{i+1}((0, \text{frac}(d_2^i + \text{frac}(t'_i)))) &= (0, \text{frac}(w_2^i + \text{frac}(t_i))) . \end{aligned}$$

By the choice of t'_i and t_i , one can prove inductively on i that for all $1 \leq i \leq n$, (i) $(s_i, q_i) = (s'_i, q'_i)$, (ii) $\eta_i \equiv \eta'_i$ and $\eta_i + t_i \equiv \eta'_i + t'_i$, and (iii) for all $I \in \mathcal{I}_i''$, the length of I is no smaller than $(N_{\varphi_i(\psi_i(I))} + 1) \cdot \frac{\delta}{|\mathcal{G}|}$. Thus, the newly-constructed sequence of transits is $\frac{\delta}{|\mathcal{G}|}$ -wide. \square

Based on Proposition 8.10, we study the equation $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta$, where $\zeta : \mathbf{B}_m \rightarrow \mathbb{R}$ is an arbitrary real vector. We define $\zeta_{\max} : \mathbf{B}_m \rightarrow \mathbb{R}$ by: $\zeta_{\max}(\mathbf{h}[v]) = M_3 \cdot \rho$ for all $\mathbf{h}[v] \in \mathbf{B}_m$. Given a vector $\zeta : \mathbf{B}_m \rightarrow \mathbb{R}$, we denote by $|\zeta|$ the vector such that $|\zeta|(\mathbf{h}[v]) = |\zeta(\mathbf{h}[v])|$ for all $\mathbf{h}[v] \in \mathbf{B}_m$. We denote by $\vec{0}$ the vector with all coordinates zero.

Proposition 8.11. *Assume $m > 2|\mathcal{G}|^2$. Let $\zeta : \mathbf{B}_m \rightarrow \mathbb{R}$ be a vector such that $|\zeta| \leq \zeta_{\max}$. Then the matrix series $\sum_{n=0}^{\infty} \mathbf{A}^n \zeta$ converges. Moreover,*

$$\left\| \sum_{n=0}^{\infty} \mathbf{A}^n \zeta \right\|_{\infty} \leq |\mathcal{G}| \cdot \mathbf{c}^{-|\mathcal{G}|} \cdot M_3 \cdot \rho ,$$

where $\mathbf{c} := e^{-\mathbf{E}_{\max} \cdot T_{\max}} \cdot p_{\min} \cdot \frac{\mathbf{E}_{\min}}{2|\mathcal{G}|^2 + \mathbf{E}_{\min}}$.

Proof. Let $\delta := |\mathcal{G}|^{-2}$ and $k := \lfloor \frac{m}{|\mathcal{G}|^2} \rfloor$. We analyse $(\sum_{n=0}^{\infty} \mathbf{A}^n \zeta)(\mathbf{h}[v])$ for each $\mathbf{h}[v] \in \mathbf{B}_m$. Let $\mathbf{h}[v] \in \mathbf{B}_m$ with $v = (s, q, \eta)$.

Firstly, we consider the case when $\zeta = \zeta_{\max}$ and $\mathbf{h}[v] \in \mathbf{B}_m^{\max}$. By definition, η is 1-separated. Then by Proposition 8.10, there exists a shortest $|\mathcal{G}|^{-1}$ -wide sequence of transits

$$(s, q, \eta) = (s_1, q_1, \eta_1) \xrightarrow{t_1} \dots \xrightarrow{t_{n-1}} (s_n, q_n, \eta_n)$$

with $1 < n \leq |\mathcal{G}|$, $q_i \notin F$ for all $1 \leq i \leq n-1$ and $q_n \in F$. By the definition of wideness, $[\eta_i + t_i]_{\equiv}$ is not marginal for all $1 \leq i \leq n-1$. We adjust

the timed information in the sequence up to δ by deviations $\{\delta_i\}_{1 \leq i \leq n-1}$, to obtain a new sequence of transits

$$(s, q, \eta) = (s_1, q'_1, \eta'_1) \xrightarrow{t_1 + \delta_1} \dots \xrightarrow{t_{n-1} + \delta_{n-1}} (s_n, q'_n, \eta'_n) ,$$

where $\delta_i \in [0, \delta)$ for all $1 \leq i \leq n-1$. Given any such deviations $\{\delta_i\}_{1 \leq i \leq n-1}$, we can prove inductively on i that for all $1 \leq i \leq n-1$ (\dagger):

- $q'_i = q_i$, $\eta'_i \equiv_{\text{gd}} \eta_i$;
- $\eta_i(x) + t_i \leq \eta'_i(x) + (t_i + \delta_i) \leq \eta_i(x) + t_i + \sum_{j=1}^i \delta_j < \eta_i(x) + t_i + |\mathcal{G}|^{-1}$ for all clocks x .
- $\eta'_i + (t_i + \delta_i) \equiv_{\text{gd}} \eta_i + t_i$ (which follows from the previous two items and the $|\mathcal{G}|^{-1}$ -wideness of the original sequence of transits).

It follows that one can deviate the original sequence of transits up to δ amount, while maintaining the reachability to some (s_n, q_n, η'_n) with $q_n \in F$.

Then, we inductively define $\{\mathbf{V}_i\}_{1 \leq i \leq n}$ with each $\mathbf{V}_i \subseteq \mathbf{D}_m$ by:

- $\mathbf{V}_1 = \{\mathbf{h}[v]\}$;
- $\mathbf{V}_{i+1} = \{\mathbf{h}[(v' \oplus \tau)_{s_{i+1}}^+ \mid \mathbf{h}[v'] \in \mathbf{V}_i, \tau \in [t_i, t_i + \delta), \mathbf{h}[v' \oplus \tau] \in \mathbf{D}_m]\}$.

We prove that (\ddagger): for all $1 \leq k \leq n$ and $(s', q', \eta') \in \mathbf{V}_k$, $(s', q') = (s_k, q_k)$, and $(s', q', \eta') \in \mathbf{B}_m$ if $k \neq n$. The case when $k = 1$ is straightforward. Below we fix an arbitrary $1 < k \leq n$. Let $(s', q', \eta') \in \mathbf{V}_k$. From the inductive definition of \mathbf{V}_i , there exists a finite sequence $\{(s''_i, q''_i, \eta''_i)\}_{1 \leq i \leq k}$ such that

- $(s''_1, q''_1, \eta''_1) = (s, q, \eta)$ and $(s''_k, q''_k, \eta''_k) = (s', q', \eta')$, and
- for all $1 \leq i < k$,

$$(s''_{i+1}, q''_{i+1}, \eta''_{i+1}) = \left(s''_i, \mathbf{q}_{q''_i, s''_i}^{(\eta''_i \oplus (t_i + \delta_i))^+}, \eta''_i \oplus (t_i + \delta_i) \left[\mathbf{X}_{q''_i, s''_i}^{(\eta''_i \oplus (t_i + \delta_i))^+} \right] \right)$$

for some $\delta_i \in [0, \delta)$.

By the definition of \mathbf{V}_i , $s''_i = s_i$ for all $1 \leq i \leq k$. We extend the sequence $\{\delta_i\}_{1 \leq i < k}$ to a whole collection of deviations $\{\delta_i\}_{1 \leq i < n}$, where the values $\{\delta_i\}_{k \leq i < n}$ are arbitrarily chosen from the interval $[0, \delta)$. By the previous analysis in this proof, we can construct a new sequence of transits

$$(s, q, \eta) = (s_1, q_1, \eta''_1) \xrightarrow{t_1 + \delta_1} \dots \xrightarrow{t_{n-1} + \delta_{n-1}} (s_n, q_n, \eta'''_n) .$$

We prove by induction on i that for all $1 \leq i \leq k$, (i) $q''_i = q_i$ and (ii) for all clocks x , either $\eta''_i(x) = \eta'''_i(x)$, or both $\eta''_i(x) \geq T_x$ and $\eta'''_i(x) \geq T_x$ holds.

The base step $i = 1$ is straightforward. Assume the inductive hypothesis for i . By the definition of transits, $t_i + \delta_i > 0$. It follows that

$$\eta''_i + t_i + \delta_i \equiv_{\text{bd}} \eta'''_i + t_i + \delta_i .$$

This further implies that

$$[\eta_i''' + t_i + \delta_i]_{\equiv} = [(\eta_i'' \oplus (t_i + \delta_i))^+]_{\equiv}$$

since $[\eta_i''' + t_i + \delta_i]_{\equiv}$ is not marginal. Thus, we have $q_{i+1}'' = q_{i+1}$ and for all clocks x , either $\eta_{i+1}''(x) = \eta_{i+1}'''(x)$, or both $\eta_{i+1}''(x) \geq T_x$ and $\eta_{i+1}'''(x) \geq T_x$ holds; this completes the inductive step.

Thus, we have $(s', q') = (s_k, q_k)$, and for all clocks x , either $\eta'(x) = \eta_k'''(x)$ or both $\eta'(x)$ and $\eta_k'''(x)$ is greater than or equal to T_x . Assume that $k \neq n$. By the definition of transits, $(s_k, q_k, [(\eta_k''')^+]_{\equiv})$ is not final and can reach some final vertex in \mathcal{G} . So we also have that $(s', q', [(\eta')^+]_{\equiv})$ is not final and can reach some final vertex in \mathcal{G} , which implies that $(s', q', [\eta']_{\equiv})$ is not final and can reach some final vertex in \mathcal{G} . By the arbitrary choice of (s', q', η') , we obtain that $\mathbf{V}_k \subseteq \mathbf{B}_m$, for $1 \leq k < n$.

Below we prove by induction on $i \geq 1$ that for all $v' \in \mathbf{V}_{n-i}$,

$$|(\mathbf{A}^i \zeta_{\max})(v')| \leq (1 - \mathbf{c}^i) \cdot M_3 \rho .$$

Note that (§): $\mathbf{A} \zeta_{\max} \leq \zeta_{\max}$ and $\mathbf{A} \zeta_1 \leq \mathbf{A} \zeta_2$ for all $\vec{0} \leq \zeta_1 \leq \zeta_2$. It follows that $\mathbf{A}^j \zeta_{\max} \leq \mathbf{A}^i \zeta_{\max}$ for all $0 \leq i \leq j$.

Base Step: $i = 1$. Consider an arbitrary $v' \in \mathbf{V}_{n-1}$. If $N_{v'} \cdot \rho < t_{n-1} + \delta_{n-1}$, then from Υ_m'' and (§), we have

$$\begin{aligned} 1 - \sum_{\mathbf{h}[v''] \in \mathbf{B}_m} \mathbf{A}(\mathbf{h}[v'], \mathbf{h}[v'']) &\geq p_{\min} \cdot \left(\frac{1}{1 + \rho \cdot \mathbf{E}(v')} \right)^{N_{v'}} \\ &\geq p_{\min} \cdot \left(\frac{1}{1 + \rho \cdot \mathbf{E}(v')} \right)^{T_{\max}/\rho} \\ &\geq p_{\min} \cdot e^{-\mathbf{E}_{\max} \cdot T_{\max}} \\ &\geq \mathbf{c} . \end{aligned}$$

Otherwise, there are at least $k := \lfloor \delta/\rho \rfloor$ distinct τ 's from the interval $[t_{n-1}, t_{n-1} + \delta_{n-1})$ such that $\mathbf{h}[v' \oplus \tau] \in \mathbf{D}_m$. Note that $k\rho \geq \frac{1}{2}|\mathcal{G}|^{-2}$. By (§), we have

$$\begin{aligned} &1 - \sum_{\mathbf{h}[v''] \in \mathbf{B}_m} \mathbf{A}(\mathbf{h}[v'], \mathbf{h}[v'']) \\ &\geq \left(\frac{1}{1 + \rho \cdot \mathbf{E}(v')} \right)^{T_{\max}/\rho} \cdot \frac{k\rho \cdot \mathbf{E}(v')}{1 + \rho \cdot \mathbf{E}(v')} \cdot p_{\min} \\ &\geq e^{-\mathbf{E}_{\max} \cdot T_{\max}} \cdot \frac{\frac{1}{2}|\mathcal{G}|^{-2} \cdot \mathbf{E}(v')}{1 + \frac{1}{2}|\mathcal{G}|^{-2} \cdot \mathbf{E}(v')} \cdot p_{\min} \\ &\geq e^{-\mathbf{E}_{\max} \cdot T_{\max}} \cdot p_{\min} \cdot \frac{\mathbf{E}_{\min}}{2|\mathcal{G}|^2 + \mathbf{E}_{\min}} . \end{aligned}$$

Thus, we obtain $|(\mathbf{A} \zeta_{\max})(v')| \leq (1 - \mathbf{c}) \cdot M_3 \rho$.

Inductive Step: Assume that $(\mathbf{A}^i \zeta_{\max})(v'') \leq (1 - \mathbf{c}^i) \cdot M_3 \rho$ for all $v'' \in \mathbf{V}'_{n-i}$. We prove the case for $i + 1$. Fix some $v' \in \mathbf{V}_{n-(i+1)}$. If $N_{v'} \cdot \rho < t_{n-1} + \delta_{n-1}$, then by a similar analysis in the base step, we have

$$(\mathbf{A}^{i+1} \zeta_{\max})(v') \leq \left(1 - p_{\min} \cdot \left(\frac{1}{1 + \rho \cdot \mathbf{E}(v')} \right)^{N_{v'}} \cdot \mathbf{c}^i \right) \cdot M_3 \rho \leq (1 - \mathbf{c}^{i+1}) \cdot M_3 \rho$$

Otherwise, there are at least $k := \lfloor \delta/\rho \rfloor$ distinct τ 's from the interval $[t_{n-(i+1)}, t_{n-(i+1)} + \delta_{n-(i+1)})$ such that $\mathbf{h}[v' \oplus \tau] \in \mathbf{D}_m$. By the induction hypothesis, we have

$$\begin{aligned} & (\mathbf{A}^{i+1} \zeta_{\max})(v') \\ & \leq M_3 \rho \cdot \left\{ 1 - \left(\frac{1}{1 + \rho \cdot \mathbf{E}(v')} \right)^{N_{v'}} \cdot \frac{k\rho \cdot \mathbf{E}(v')}{1 + \rho \cdot \mathbf{E}(v')} \cdot p_{\min} \cdot \mathbf{c}^i \right\} \\ & \leq M_3 \rho \cdot \left\{ 1 - e^{-\mathbf{E}_{\max} \cdot T_{\max}} \cdot p_{\min} \cdot \frac{\mathbf{E}_{\min}}{2|\mathcal{G}|^2 + \mathbf{E}_{\min}} \cdot \mathbf{c}^i \right\} \\ & = M_3 \rho \cdot (1 - \mathbf{c}^{i+1}) . \end{aligned}$$

Then the inductive step is completed.

Then we obtain

$$(\mathbf{A}^{|\mathcal{G}|-1} \zeta_{\max})(v') \leq (\mathbf{A}^i \zeta_{\max})(v') \leq (1 - \mathbf{c}^i) \cdot M_3 \rho \leq (1 - \mathbf{c}^{|\mathcal{G}|-1}) \cdot M_3 \rho$$

for all $1 \leq i \leq n-1$ and $v' \in \mathbf{V}_{n-i}$. Thus, $\mathbf{A}^{|\mathcal{G}|-1} \zeta_{\max}(v) \leq (1 - \mathbf{c}^{|\mathcal{G}|-1}) \cdot M_3 \rho$, for all $v \in \mathbf{B}_m^{\max}$.

Now consider an arbitrary $v \in \mathbf{B}_m$ while $\zeta = \zeta_{\max}$. If either $v \oplus (N_v \rho) \notin \mathbf{B}_m^{\max}$ or $\mathbf{h}[(v \oplus (N_v \rho))_u^+] \notin \mathbf{B}_m$ for some $u \in S$ with $\mathbf{P}(v, u) > 0$, then

$$\begin{aligned} (\mathbf{A}^{|\mathcal{G}|} \zeta_{\max})(v) & \leq (\mathbf{A} \zeta_{\max})(v) \\ & \leq \left(1 - \left(\frac{1}{1 + \rho \mathbf{E}(v)} \right)^{T_{\max}/\rho} \cdot p_{\min} \right) \cdot M_3 \rho \\ & \leq (1 - e^{-\mathbf{E}_{\max} \cdot T_{\max}} \cdot p_{\min}) \cdot M_3 \rho \\ & \leq (1 - \mathbf{c}^{|\mathcal{G}|}) \cdot M_3 \rho \end{aligned}$$

Otherwise, by

$$\sum_{u \in S} \mathbf{P}(v^*, u) \cdot (\mathbf{A}^{|\mathcal{G}|-1} \zeta_{\max})((v^*)_u^+) = (\mathbf{A}^{|\mathcal{G}|} \zeta_{\max})(v^*) \leq (\mathbf{A}^{|\mathcal{G}|-1} \zeta_{\max})(v^*)$$

where $v^* := v \oplus (N_v \rho)$, we have

$$\begin{aligned} (\mathbf{A} \cdot \mathbf{A}^{|\mathcal{G}|-1} \zeta_{\max})(v) & \leq \left(1 - \left(\frac{1}{1 + \rho \cdot \mathbf{E}(v)} \right)^{T_{\max}/\rho} \cdot \mathbf{c}^{|\mathcal{G}|-1} \right) \cdot M_3 \rho \\ & \leq (1 - \mathbf{c}^{|\mathcal{G}|}) \cdot M_3 \rho . \end{aligned}$$

Then $\mathbf{A}^{|\mathcal{G}|}\zeta_{\max} \leq (1 - \mathbf{c}^{|\mathcal{G}|}) \cdot \zeta_{\max}$. It follows that

$$\mathbf{A}^{i|\mathcal{G}|}\zeta_{\max} \leq (1 - \mathbf{c}^{|\mathcal{G}|})^i \zeta_{\max}$$

for all $i \in \mathbb{N}$. Thus by (§), $\sum_{i=0}^{\infty} \mathbf{A}^i \zeta_{\max}$ converges since $\sum_{i=0}^{\infty} \mathbf{A}^i \zeta_{\max}$ is bounded by $|\mathcal{G}| \cdot \mathbf{c}^{-|\mathcal{G}|} \cdot \zeta_{\max}$.

Finally, we consider any ζ such that $|\zeta| \leq \zeta_{\max}$. Since all entries of \mathbf{A}^i are non-negative, $|\mathbf{A}^i \zeta| \leq \mathbf{A}^i \zeta_{\max}$. Thus by Cauchy's criterion, $\sum_{i=0}^{\infty} \mathbf{A}^i \zeta$ converges and $\|\sum_{i=0}^{\infty} \mathbf{A}^i \zeta\|_{\infty} \leq |\mathcal{G}| \cdot \mathbf{c}^{-|\mathcal{G}|} \cdot M_3 \rho$. \square

By Proposition 8.11, the system of linear equation $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta$ has a solution for all $|\zeta| \leq \zeta_{\max}$, when $m > 2|\mathcal{G}|^2$. The following propositions show that the linear equation has a unique solution.

Proposition 8.12. *Assume that $m > 2|\mathcal{G}|^2$. For all solutions \mathbf{v} of $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta$ with $\|\zeta\|_{\infty} < M_3 \rho$, $|\mathbf{v}| \leq \mathbf{v}^*$, where $\mathbf{v}^* := \sum_{i=0}^{\infty} \mathbf{A}^i \zeta_{\max}$.*

Proof. Let \mathbf{v} be an arbitrary solution of $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta$. Define $\mathbf{v}' = \mathbf{v}^* - \mathbf{v}$. By the fact that $\mathbf{v}^* = \mathbf{A}\mathbf{v}^* + \zeta_{\max}$, $\mathbf{v}'(\mathbf{h}[v]) > (\mathbf{A}\mathbf{v}')(\mathbf{h}[v])$ for all $\mathbf{h}[v] \in \mathcal{B}_m$. Suppose that there exists some $\mathbf{h}[v] \in \mathcal{B}_m$ such that $\mathbf{v}'(\mathbf{h}[v]) < 0$. W.l.o.g, we assume that $\mathbf{v}'(\mathbf{h}[v])$ is the least element of $\{\mathbf{v}'(\mathbf{h}[v']) \mid \mathbf{h}[v'] \in \mathcal{B}_m\}$. Denote $c := \sum_{\mathbf{h}[v'] \in \mathcal{B}_m} \mathbf{A}(\mathbf{h}[v], \mathbf{h}[v']) \in [0, 1]$. Then $\mathbf{v}'(\mathbf{h}[v]) > c \cdot \mathbf{v}'(\mathbf{h}[v])$, which implies $c > 1$. Contradiction. Thus $\mathbf{v}' \geq \vec{0}$. Similar arguments holds if we define $\mathbf{v}' = \mathbf{v}^* + \mathbf{v}$. Thus we have $|\mathbf{v}| \leq \mathbf{v}^*$. \square

Proposition 8.13. *The matrix equation $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta$ has a unique solution for all ζ such that $\|\zeta\|_{\infty} < M_3 \rho$. It follows that $\mathbf{I} - \mathbf{A}$ is invertible, where \mathbf{I} is the identity matrix.*

Proof. By Proposition 8.11, the equation $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta$ has a solution. By Proposition 8.12, all solutions of $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta$ are bounded by \mathbf{v}^* . Suppose that the equation has two distinct solutions. Then the homogeneous equation $\mathbf{v} = \mathbf{A}\mathbf{v}$ has a non-trivial solution, which implies that the solutions of $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta$ cannot be bounded. Contradiction. Thus $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta$ has a unique solution and $\mathbf{I} - \mathbf{A}$ is invertible. \square

Now we analyse the approximation scheme Υ'_m (Υ_m). In the following theorem (which is the main result of this chapter), we prove that the equation $\mathbf{v} = \mathbf{C}\mathbf{v} + \mathbf{d}$ has a unique solution (i.e. $\mathbf{I} - \mathbf{C}$ is invertible), and give the error bound between the unique solution and the function prob.

Theorem 8.5. *The matrix equation $\mathbf{v} = \mathbf{C}\mathbf{v} + \mathbf{d}$ (for Υ'_m) has a unique solution $\bar{\mathbf{v}}$. Moreover, $\max_{\mathbf{h}[v] \in \mathcal{B}_m} |\bar{\mathbf{v}}(\mathbf{h}[v]) - \text{prob}(v)| \leq |\mathcal{G}| \cdot \mathbf{c}^{-|\mathcal{G}|} \cdot M_3 \rho$.*

Proof. We first prove that $\mathbf{v} = \mathbf{C}\mathbf{v} + \mathbf{d}$ has a unique solution. Let $\mathbf{v} = \mathbf{C}\mathbf{v} + \zeta$ be a matrix equation such that $\|\zeta\|_{\infty} < M_2 \rho^2$. From the proof of Proposition 8.8, we can equivalently expand this equation into some equation

$\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta'$ with $\|\zeta'\|_\infty < T_{\max}/\rho \cdot M_2\rho^2 = M_3 \cdot \rho$. Since $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta'$ has a unique solution, $\mathbf{v} = \mathbf{C}\mathbf{v} + \zeta$ also has a unique solution. Thus $\mathbf{I} - \mathbf{C}$ is invertible and $\mathbf{v} = \mathbf{C}\mathbf{v} + \mathbf{d}$ has a unique solution.

Now we prove the error bound between $\bar{\mathbf{v}}$ and prob . Define the vector \mathbf{v}' such that $\mathbf{v}'(\mathbf{h}[v]) = \bar{\mathbf{v}}(\mathbf{h}[v]) - \text{prob}(v)$ for all $\mathbf{h}[v] \in \mathbf{B}_m$. By Proposition 8.7, \mathbf{v}' is the unique solution of $\mathbf{v} = \mathbf{C}\mathbf{v} + \zeta$, for some $\|\zeta\|_\infty < M_2\rho^2$. Then \mathbf{v}' is also the unique solution of the equation $\mathbf{v} = \mathbf{A}\mathbf{v} + \zeta'$, for some $\|\zeta'\|_\infty < M_3\rho$. By Proposition 8.11, $\|\mathbf{v}'\|_\infty \leq |\mathcal{G}| \cdot \mathbf{c}^{-|\mathcal{G}|} \cdot M_3\rho$. \square

By Theorem 8.5 and the Lipschitz Continuity (Corollary 8.1), we can approximate the value $\text{prob}(s, q, \eta)$ as follows: given $\epsilon \in (0, 1)$, we choose m sufficiently large and some $\mathbf{h}[v] \in \mathbf{D}_m$ such that $|\text{prob}(v) - \text{prob}(s, q, \eta)| < \frac{1}{2} \cdot \epsilon$ and $M_3|\mathcal{G}|\mathbf{c}^{-|\mathcal{G}|} \cdot \rho < \frac{1}{2} \cdot \epsilon$. Then we solve the approximation scheme Υ'_m to obtain $\bar{\mathbf{v}}(\mathbf{h}[v])$.

8.7 Conclusion

In this chapter, we corrected the errors in the paper [24] by new proofs, namely, the proof for the measurability of the set of CTMC-paths accepted by a DTA and the proof for the integral characterization. And we presented the first algorithm to approximate the acceptance probabilities of CTMC-paths by a multi-clock DTA under finite acceptance condition. Unlike the result by Barbot *et al.* [9], we are able to derive a tight approximation error.

Chapter 9

Cost-Bounded Reachability on CTMDPs

In this chapter, we focus on the problem to compute *max/min resource-bounded reachability probability* on a CTMDP (cf. Chapter 7). Typical resource types considered here are time and cost, where a time bound can be deemed as a special cost bound with unit-cost one. In general, the task is to compute or to approximate the optimal (max/min) reachability probability to certain target states within a given resource bound (e.g., a time bound).

Optimal time-bounded reachability probability over CTMDPs has been widely studied in recent years. Neuhäüßer *et al.* [59] proved that the maximal time-bounded reachability probability function is the least fixed point of a system of integral equations. Rabe and Schewe [65] showed that the max/min time-bounded reachability probability can be attained by a deterministic piecewise-constant time-positional scheduler. Efficient approximation algorithms are also developed by, e.g., Neuhäüßer *et al.* [59], Brázdil *et al.* [15], Hatefi *et al.* [46] and Rabe *et al.* [35].

As to optimal cost-bounded reachability probability, much less is known. To the best of the author's knowledge, the only prominent result is by Baier *et al.* [5], which establishes a certain duality property between time and cost bound, under the setting of time-abstract schedulers. This duality results in an approximation algorithm for the case of one-dimensional cost-bounds. Their result is restrictive in the sense that (i) it assumes that the CTMDP have everywhere positive unit-cost values, (ii) it only takes into account one-dimensional cost-bound aside the time-bound, and (iii) it does not really provide an approximation algorithm when both time- and cost-bounds are present.

Besides resource-bounded reachability probability, we would like to mention another well-investigated objective on CTMDPs with costs (or dually, rewards), which is (discounted) accumulated reward over finite/infinite horizon (cf. [20, 63], just to mention a few).

In this chapter, we consider multi-dimensional maximal cost-bounded reachability probability (*abbr.* MMCRP) over CTMDPs under the setting of both early and late schedulers, for which the unit-cost is constant. We first prove that the MMCRP function is the least fixed-point of a system of integral equations. Then we prove that deterministic cost-positional measurable schedulers suffice to achieve the MMCRP value. Finally, we describe a numerical algorithm which approximates the MMCRP value with an error bound. The approximation algorithm relies on a differential characterization which in turn is derived from the least fixed-point characterization. The complexity of the approximation algorithm is polynomial in the size of the CTMDP and the reciprocal of the error bound, and exponential in the dimension of cost-bound vectors.

Besides, we point out a proof error in the treatment of maximal time-bounded reachability probability on continuous-time Markov decision processes [57, 59]. We fix this error in the more general setting of maximal cost-bounded reachability probability through a new methodology.

The chapter is organized as follows. In Section 9.1, we define the notion of maximal cost-bounded reachability probability and derive the least-fixed-point characterization, while we also point out the proof error in [57, 59]. In Section 9.2, we prove that the maximal cost-bounded reachability probability can be reached by a measurable deterministic cost-positional scheduler. In Section 9.3, we derive a differential characterization which is crucial to our approximation algorithm. In Section 9.4, we present our approximation algorithm. Finally, Section 9.5 concludes the chapter.

We denote by \vec{x} the real vector whose coordinates are all equal to $x \in \mathbb{R}$ (with the implicitly known dimension).

9.1 Cost-Bounded Reachability Probability

In this section, we introduce the notion of maximal cost-bounded reachability probability. Below we fix a CTMDP $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$ (cf. Chapter 7). We define:

- $\mathbf{E}_{\text{max}} := \max(\{\mathbf{E}_{\text{la}}(s) \mid s \in S_{\text{la}}\} \cup \{\mathbf{E}_{\text{er}}(s, a) \mid s \in S_{\text{er}}, a \in \text{En}(s)\})$;
- $\mathbf{P}_{\text{min}} := \min\{\mathbf{P}(s, a, s') \mid s, s' \in S, a \in \text{Act}, \mathbf{P}(s, a, s') > 0\}$.

In the whole chapter, we will use ‘ $\mathbf{E}(s, a)$ ’ to denote $\mathbf{E}_{\text{er}}(s, a)$ when $s \in S_{\text{er}}$ and $a \in \text{Act}$, and ‘ $\mathbf{E}(s)$ ’ to denote $\mathbf{E}_{\text{la}}(s)$ when $s \in S_{\text{la}}$.

Firstly, we introduce the notion of *cost functions* which associates costs (or dually, rewards) to a CTMDP.

Definition 9.1. *Let $k \in \mathbb{N}$. A cost function \mathbf{w} (for \mathcal{M}) is a function $\mathbf{w} : S \times \text{Act} \rightarrow \mathbb{R}_{\geq 0}^k$ such that for all $s \in S_{\text{la}}$ and $a, b \in \text{En}(s)$, $\mathbf{w}(s, a) = \mathbf{w}(s, b)$.*

Intuitively, a cost function assigns to each pair in $S \times Act$ a non-negative real vector \mathbf{c} , where each component \mathbf{c}_j of the vector can be viewed as the unit cost per time w.r.t certain resource represented by the coordinate index j , when certain action is chosen at certain current state. A cost function should observe the restriction that when it is applied to a late-schedulable state, the cost should be independent of the action chosen (by a scheduler) at the state; this restriction is to make cost functions compatible with the notion of late-schedulable states (cf. Chapter 7).

In the following, we fix a cost function $\mathbf{w} : S \times Act \rightarrow \mathbb{R}_{\geq 0}^k$ (for \mathcal{M}), where $k \in \mathbb{N}$ is a fixed natural number. We abbreviate $'(\mathbf{w}(s, a))_i'$ ($s \in S$, $a \in Act$, $1 \leq i \leq k$) as $'\mathbf{w}_i(s, a)'$. For a late-schedulable states $s \in S_{\text{la}}$, we simple use $'\mathbf{w}(s)'$ to denote $'\mathbf{w}(s, a)'$ for an arbitrary $a \in \text{En}(s)$. We define

- $\mathbf{w}_{\min} := \min\{\mathbf{w}_i(s, a) \mid 1 \leq i \leq k, s \in S, a \in \text{En}(s), \mathbf{w}_i(s, a) > 0\}$, and
- $\mathbf{w}_{\max} := \max\{\mathbf{w}_i(s, a) \mid 1 \leq i \leq k, s \in S, a \in \text{En}(s)\}$.

We assume that \mathbf{w}_{\min} is well-defined and $\mathbf{w}_{\min} > 0$ (i.e., \mathbf{w} is not the zero function.): if \mathbf{w} is the zero function, then the cost-bounded reachability on \mathcal{M} will be equivalent to reachability on the discrete-time Markov decision process with state space S , actions Act and the probability transition matrix \mathbf{P} (cf. [7, Section 10.6]).

The main focus of this chapter is on costs of paths and histories. The cost is assigned linearly w.r.t the unit-cost and the time spent in a state. The following definition presents the details.

Definition 9.2. *Given a path $\pi \in \text{Paths}(\mathcal{M})$ and a set $G \subseteq S$ of states, we denote by $\mathbf{C}(\pi, G)$ the accumulated cost vector along π until G is reached; formally, if $\pi[m] \in G$ for some $m \geq 0$ then*

$$\mathbf{C}(\pi, G) := \sum_{i=0}^n \pi\langle i \rangle \cdot \mathbf{w}(\pi[i], \pi(i)) \ ,$$

where $n \in \mathbb{N}_0 \cup \{-1\}$ is the smallest integer such that $\pi[n+1] \in G$; otherwise, $\mathbf{C}(\pi, G) := \vec{\infty}$, for which $\vec{\infty}$ is the k -dimensional vector whose all coordinates are ∞ .

Given a history $\xi \in \text{Hists}(\mathcal{M})$, we denote by $\mathbf{C}(\xi)$ the accumulated cost vector of ξ ; formally,

$$\mathbf{C}(\xi) := \sum_{i=0}^{|\xi|-1} \xi\langle i \rangle \cdot \mathbf{w}(\xi[i], \xi(i)) \ .$$

Then, we introduce the notion of maximal cost-bounded reachability probability, which is the main subject of this chapter.

Definition 9.3. Let $G \subseteq S$. For each measurable scheduler D , we define the function $\text{prob}_G^D : S \times \mathbb{R}^k \rightarrow [0, 1]$ by: $\text{prob}_G^D(s, \mathbf{c}) := \Pr_{D, \mathcal{D}[s]}(\Pi_G^{\mathbf{c}})$ where

$$\Pi_G^{\mathbf{c}} := \{\pi \in \text{Paths}(\mathcal{M}) \mid \mathbf{C}(\pi, G) \leq \mathbf{c}\} .$$

Moreover, we define $\text{prob}_G^{\max} : S \times \mathbb{R}^k \rightarrow [0, 1]$ by:

$$\text{prob}_G^{\max}(s, \mathbf{c}) := \sup_{D \in \text{MSched}(\mathcal{M})} \text{prob}_G^D(s, \mathbf{c})$$

for all $s \in S$ and $\mathbf{c} \in \mathbb{R}^k$, where $\text{MSched}(\mathcal{M})$ is the set of all measurable schedulers (for \mathcal{M}).

Remark 9.1. It is not hard to prove that $\Pi_G^{\mathbf{c}}$ is measurable. One can proceed by showing that each $\Pi_{n,G}^{\mathbf{c}}$, which is the set of paths π that can reach G within cost-bound \mathbf{c} and n transition steps, is a closed set on the collection of variables $(\pi\langle 0 \rangle, \dots, \pi\langle n-1 \rangle)$. Then the result follows from the fact that $\Pi_G^{\mathbf{c}} = \bigcup_{n \in \mathbb{N}} \Pi_{n,G}^{\mathbf{c}}$. Thus, all functions in Definition 9.3 are well-defined.

From the definition, we can see that $\Pi_G^{\mathbf{c}}$ is the set of paths which can reach G within the cost-bound vector \mathbf{c} , and $\text{prob}_G^{\max}(s, \mathbf{c})$ is the maximal probability of $\Pi_G^{\mathbf{c}}$ with initial distribution $\mathcal{D}[s]$ (i.e., with fixed initial state s) ranging over all measurable schedulers. It is worth noting that if $\mathbf{c} \not\geq \vec{0}$, then both $\text{prob}_G^D(s, \mathbf{c})$ and $\text{prob}_G^{\max}(s, \mathbf{c})$ is zero.

From Theorem 7.1 and Proposition 7.1, we can directly obtain the following results.

Corollary 9.1. Let D be a measurable scheduler and $G \subseteq S$. The function prob_G^D satisfies the following conditions for all $s \in S$:

1. if $s \in G$ then $\text{prob}_G^D(s, \mathbf{c}) = \mathbf{1}_{\mathbb{R}_{\geq 0}^k}(\mathbf{c})$ for all $\mathbf{c} \in \mathbb{R}^k$;
2. if $s \in S_{\text{er}} - G$ then

$$\begin{aligned} \text{prob}_G^D(s, \mathbf{c}) &= \sum_{a \in \text{En}(s)} D(s, a) \cdot \\ &\int_0^\infty f_{\mathbf{E}(s,a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{D[s \xrightarrow{a,t}]}(s', \mathbf{c} - t \cdot \mathbf{w}(s, a)) \right] dt ; \end{aligned}$$

3. if $s \in S_{\text{la}} - G$ then

$$\begin{aligned} \text{prob}_G^D(s, \mathbf{c}) &= \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \left\{ \sum_{a \in \text{En}(s)} D(s, t, a) \cdot \right. \\ &\left. \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{D[s \xrightarrow{a,t}]}(s', \mathbf{c} - t \cdot \mathbf{w}(s)) \right] \right\} dt . \end{aligned}$$

Intuitively, Corollary 9.1 expands the function prob_G^D to an integral representation.

The following theorem mainly presents the fixed-point characterization for prob_G^{\max} , while it also states that prob_G^{\max} is Lipschitz continuous.

Theorem 9.1. *Let $G \subseteq S$. The function $\text{prob}_G^{\max}(\cdot, \cdot)$ is the least fixed-point (w.r.t \leq) of the higher-order operator*

$$\mathcal{T}_G : \left[S \times \mathbb{R}^k \rightarrow [0, 1] \right] \rightarrow \left[S \times \mathbb{R}^k \rightarrow [0, 1] \right]$$

defined by:

- $\mathcal{T}_G(h)(s, \mathbf{c}) := \mathbf{1}_{\mathbb{R}_{\geq 0}^k}(\mathbf{c})$ for all $s \in G$ and $\mathbf{c} \in \mathbb{R}^k$;
- for all $s \in S_{\text{er}} - G$ and $\mathbf{c} \in \mathbb{R}^k$,

$$\mathcal{T}_G(h)(s, \mathbf{c}) := \max_{a \in \text{En}(s)} \int_0^\infty f_{\mathbf{E}(s,a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot h(s', \mathbf{c} - t \cdot \mathbf{w}(s, a)) \right] dt ;$$

- for all $s \in S_{\text{la}} - G$ and $\mathbf{c} \in \mathbb{R}^k$,

$$\mathcal{T}_G(h)(s, \mathbf{c}) := \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot h(s', \mathbf{c} - t \cdot \mathbf{w}(s)) \right] dt ;$$

for each $h : S \times \mathbb{R}^k \rightarrow [0, 1]$ (cf. Definition 3.8). Moreover,

$$\left| \text{prob}_G^{\max}(s, \mathbf{c}) - \text{prob}_G^{\max}(s, \mathbf{c}') \right| \leq \frac{\mathbf{E}_{\max}}{\mathbf{w}_{\min}} \cdot \|\mathbf{c} - \mathbf{c}'\|_\infty$$

for all $\mathbf{c}, \mathbf{c}' \geq \vec{0}$ and $s \in S$.

Proof. For each $n \in \mathbb{N}_0$, we define the function $\text{prob}_{n,G}^{\max} : S \times \mathbb{R}^k \rightarrow [0, 1]$ by

$$\text{prob}_{n,G}^{\max}(s, \mathbf{c}) := \sup_{D \in \text{MSched}(\mathcal{M})} \Pr_{D, \mathcal{D}[s]} \left(\Pi_{n,G}^{\mathbf{c}} \right) ,$$

where

$$\Pi_{n,G}^{\mathbf{c}} := \{ \pi \in \text{Paths}(\mathcal{M}) \mid \mathbf{C}(\pi, G) \leq \mathbf{c} \text{ and } \pi[m] \in G \text{ for some } 0 \leq m \leq n \} .$$

Intuitively, $\text{prob}_{n,G}^{\max}$ is the maximal cost-bounded reachability probability function within n steps. For each $n \in \mathbb{N}_0$ and $\delta > 0$, define

$$\epsilon(n, \delta) := \sup \left\{ \left| \text{prob}_{n,G}^{\max}(s, \mathbf{c}) - \text{prob}_{n,G}^{\max}(s, \mathbf{c}') \right| \mid s \in S, \mathbf{c}, \mathbf{c}' \geq \vec{0} \text{ and } \|\mathbf{c} - \mathbf{c}'\|_\infty \leq \delta \right\} .$$

Note that $\epsilon(0, \delta) = 0$ for all $\delta > 0$. Firstly, we prove by induction on $n \geq 0$ that the following assertions hold:

- (a) $\text{prob}_{n,G}^{\max}(s, \mathbf{c}) = \mathbf{1}_{\mathbb{R}_{\geq 0}^k}(\mathbf{c})$ if $s \in G$;
- (b) given any $\delta > 0$, $\epsilon(n, \delta) \leq \frac{\mathbf{E}_{\max}}{\mathbf{w}_{\min}} \cdot \delta$;
- (c) if $n > 0$ then $\text{prob}_{n+1,G}^{\max}(s, \mathbf{c}) = \mathcal{T}_G(\text{prob}_{n,G}^{\max})(s, \mathbf{c})$ for all $s \in S - G$ and $\mathbf{c} \in \mathbb{R}^k$.

The base step when $n = 0$ is straightforward: it is clear that $\Pr_{D, \mathcal{D}[s]}(\Pi_{0,G}^{\mathbf{c}}) = \mathbf{1}_G(s) \cdot \mathbf{1}_{\mathbb{R}_{\geq 0}^k}(\mathbf{c})$ for all measurable schedulers D ; thus (a), (b) holds and (c) is vacuously true. For the inductive step, let $n = m + 1$ with $m \geq 0$ and assume that (a),(b),(c) hold at m . It is easy to see that (a) holds at n . We prove that (b) and (c) hold for n .

We first prove the inductive step for (c). Let $\mathbf{c} \geq \vec{0}$ and $s \in S - G$. (The situation when $\mathbf{c} \not\geq \vec{0}$ is straightforward.) Assume that $s \in S_{\text{er}} - G$. By Theorem 7.1 and Proposition 7.1, for all measurable schedulers D ,

$$\Pr_{D, \mathcal{D}[s]}(\Pi_{m+1,G}^{\mathbf{c}}) = \sum_{a \in \text{En}(s)} D(s, a) \cdot \int_0^\infty f_{\mathbf{E}(s,a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \Pr_{D[s \xrightarrow{a,t}], \mathcal{D}[s']} \left(\Pi_{m,G}^{\mathbf{c} - t \cdot \mathbf{w}(s,a)} \right) \right] dt .$$

If we modify D to D' by setting $D'(s, \cdot)$ to the Dirac distribution at the action

$$\operatorname{argmax}_{a \in \text{En}(s)} \int_0^\infty f_{\mathbf{E}(s,a)}(t) \cdot \left(\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \Pr_{D[s \xrightarrow{a,t}], \mathcal{D}[s']} \left(\Pi_{m,G}^{\mathbf{c} - t \cdot \mathbf{w}(s,a)} \right) \right) dt$$

and $D'(\xi, \cdot) = D(\xi, \cdot)$, $D'(\xi, \cdot, \cdot) = D(\xi, \cdot, \cdot)$ for $\xi \neq s$, then D' is a measurable scheduler which satisfies that $\Pr_{D, \mathcal{D}[s]}(\Pi_{m+1,G}^{\mathbf{c}}) \leq \Pr_{D', \mathcal{D}[s]}(\Pi_{m+1,G}^{\mathbf{c}})$ (since $D'[s \xrightarrow{a,t}] = D[s \xrightarrow{a,t}]$). Thus,

$$\text{prob}_{m+1,G}^{\max}(s, \mathbf{c}) = \sup_{D \in \text{MSched}(\mathcal{M})} \max_{a \in \text{En}(s)} \int_0^\infty f_{\mathbf{E}(s,a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \Pr_{D[s \xrightarrow{a,t}], \mathcal{D}[s']} \left(\Pi_{m,G}^{\mathbf{c} - t \cdot \mathbf{w}(s,a)} \right) \right] dt .$$

We further prove that $\text{prob}_{m+1,G}^{\max}(s, \mathbf{c})$ equals $\max_{a \in \text{En}(s)} \mathbf{G}_{\text{er}}^s(a, \mathbf{c})$, where

$$\mathbf{G}_{\text{er}}^s(a, \mathbf{c}) := \int_0^\infty f_{\mathbf{E}(s,a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_{m,G}^{\max}(s', \mathbf{c} - t \cdot \mathbf{w}(s,a)) \right] dt .$$

Denote $\text{prob}_{m+1,G}^{\max}(s, \mathbf{c}) = \sup_{D \in \text{MSched}(\mathcal{M})} \max_{a \in \text{En}(s)} \mathbf{F}_{\text{er}}^s(D, a, \mathbf{c})$ with

$$\mathbf{F}_{\text{er}}^s(D, a, \mathbf{c}) := \int_0^\infty f_{\mathbf{E}(s,a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \Pr_{D[s \xrightarrow{a,t}], \mathcal{D}[s']} \left(\Pi_{m,G}^{\mathbf{c} - t \cdot \mathbf{w}(s,a)} \right) \right] dt .$$

It is not difficult to see that $\text{prob}_{m+1,G}^{\max}(s, \mathbf{c}) \leq \max_{a \in \text{En}(s)} \mathbf{G}_{\text{er}}^s(a, \mathbf{c})$. Below we prove the reverse direction. Let

$$a^* := \underset{a \in \text{En}(s)}{\text{argmax}} \mathbf{G}_{\text{er}}^s(a, \mathbf{c}) ,$$

where the action is chosen w.r.t an arbitrarily-fixed linear order when ties occur. We clarify two cases below.

Case 1: $\mathbf{w}(s, a^*) = \vec{0}$. For each $\epsilon > 0$, we can choose the measurable scheduler D^ϵ such that the following conditions hold:

- $D^\epsilon(s, \bullet) = \mathcal{D}[a^*]$;
- $D^\epsilon(s \xrightarrow{a^*, t} \xi, \bullet) = D_{\xi[0]}^{m, \epsilon}(\xi, \bullet)$ for all $t \in \mathbb{R}_{\geq 0}$ and $\xi \in \text{Hists}_{\text{er}}(\mathcal{M})$;
- $D^\epsilon(s \xrightarrow{a^*, t} \xi, \tau, \bullet) = D_{\xi[0]}^{m, \epsilon}(\xi, \tau, \bullet)$ for all $t, \tau \in \mathbb{R}_{\geq 0}$ and $\xi \in \text{Hists}_{\text{la}}(\mathcal{M})$;

the measurable scheduler $D_{\xi[0]}^{m, \epsilon}$ is chosen such that

$$\Pr_{D_{\xi[0]}^{m, \epsilon}, \mathcal{D}[\xi[0]]}(\Pi_{m, G}^{\mathbf{c}}) \geq \text{prob}_{m, G}^{\max}(\xi[0], \mathbf{c}) - \epsilon .$$

The decisions $D(\xi, \bullet)$ or $D(\xi, \tau, \bullet)$ (for all other $\xi \in \text{Hists}(\mathcal{M})$ and $\tau \in \mathbb{R}_{\geq 0}$) are irrelevant and can be set to an arbitrary canonical distribution. It is not hard to verify that D^ϵ satisfies

$$\max_{a \in \text{En}(s)} \mathbf{F}_{\text{er}}^s(D^\epsilon, a, \mathbf{c}) \geq \mathbf{F}_{\text{er}}^s(D^\epsilon, a^*, \mathbf{c}) \geq \mathbf{G}_{\text{er}}^s(a^*, \mathbf{c}) - \epsilon .$$

Thus $\text{prob}_{m+1, G}^{\max}(s, \mathbf{c}) = \mathbf{G}_{\text{er}}^s(a^*, \mathbf{c})$ by the arbitrary choice of ϵ .

Case 2: $\mathbf{w}(s, a^*) \neq \vec{0}$. Then the integrand function of $\mathbf{G}_{\text{er}}^s(a^*, \mathbf{c})$ takes non-zero values only on $[0, T_{s, a^*}^{\mathbf{c}}]$ with

$$T_{s, a^*}^{\mathbf{c}} := \min \left\{ \frac{\mathbf{c}_i}{\mathbf{w}_i(s, a^*)} \mid 1 \leq i \leq k, \mathbf{w}_i(s, a^*) > 0 \right\} .$$

By induction hypothesis (b), the integrand of $\mathbf{G}_{\text{er}}^s(a^*, \mathbf{c})$ is Lipschitz continuous on $[0, T_{s, a^*}^{\mathbf{c}}]$. For each $N \in \mathbb{N}$, we divide the interval $[0, T_{s, a^*}^{\mathbf{c}}]$ into N pieces I_1, \dots, I_N with $[t_{j-1}, t_j] := I_j = [\frac{j-1}{N} \cdot T_{s, a^*}^{\mathbf{c}}, \frac{j}{N} \cdot T_{s, a^*}^{\mathbf{c}}]$ for $1 \leq j \leq N$. Then we construct the measurable scheduler $D^{N, \epsilon}$ as follows:

- $D^{N, \epsilon}(s, \bullet) = \mathcal{D}[a^*]$;
- $D^{N, \epsilon}(s \xrightarrow{a^*, t} \xi, \bullet) = D_{\xi[0], j}^{m, \epsilon}(\xi, \bullet)$ whenever $\xi \in \text{Hists}_{\text{er}}(\mathcal{M})$ and $t \in I_j$;
- $D^{N, \epsilon}(s \xrightarrow{a^*, t} \xi, \tau, \bullet) = D_{\xi[0], j}^{m, \epsilon}(\xi, \tau, \bullet)$ whenever $\xi \in \text{Hists}_{\text{la}}(\mathcal{M})$, $\tau \in \mathbb{R}_{\geq 0}$ and $t \in I_j$;

the measurable scheduler $D_{\xi[0], j}^{m, \epsilon}$ is defined such that

$$\Pr_{D_{\xi[0], j}^{m, \epsilon}, \mathcal{D}[\xi[0]]}(\Pi_{m, G}^{\mathbf{c} - t_j \cdot \mathbf{w}(s, a^*)}) \geq \text{prob}_{m, G}^{\max}(\xi[0], \mathbf{c} - t_j \cdot \mathbf{w}(s, a^*)) - \epsilon .$$

The decisions $D^{N,\epsilon}(\xi, \bullet)$ or $D^{N,\epsilon}(\xi, \bullet, \bullet)$ for all other ξ 's are irrelevant. By induction hypothesis (b) and the monotonicity of $\Pi_{m,G}^{\mathbf{c}-t \cdot \mathbf{w}(s,a^*)}$ on t ,

$$\lim_{N \rightarrow \infty, \epsilon \rightarrow 0^+} F_{\text{er}}^s(D^{N,\epsilon}, a^*, \mathbf{c}) = G_{\text{er}}^s(a^*, \mathbf{c}) .$$

Then, since N, ϵ can be arbitrarily chosen, $\text{prob}_{m+1,G}^{\max}(s, \mathbf{c}) = G_{\text{er}}^s(a^*, \mathbf{c})$.

Now assume that $s \in S_{\text{la}} - G$. By Theorem 7.1 and Proposition 7.1,

$$\begin{aligned} \Pr_{D, \mathcal{D}[s]}(\Pi_{m+1,G}^{\mathbf{c}}) &= \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \sum_{a \in \text{En}(s)} D(s, t, a) \cdot \\ &\quad \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \Pr_{D[s \xrightarrow{a,t}], \mathcal{D}[s']} \left(\Pi_{m,G}^{\mathbf{c}-t \cdot \mathbf{w}(s)} \right) \right] dt \end{aligned}$$

for all measurable scheduler D . Note that for each $a \in \text{En}(s)$, the function

$$t \mapsto \sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \Pr_{D[s \xrightarrow{a,t}], \mathcal{D}[s']} \left(\Pi_{m,G}^{\mathbf{c}-t \cdot \mathbf{w}(s)} \right)$$

is measurable w.r.t $(\mathbb{R}_{\geq 0}, \mathcal{B}(\mathbb{R}_{\geq 0}))$ due to Proposition 7.1 and Proposition 7.2. Thus, if we modify D to D' by setting (i) $D'(s, t, \bullet)$ to

$$D \left[\operatorname{argmax}_{a \in \text{En}(s)} \sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \Pr_{D[s \xrightarrow{a,t}], \mathcal{D}[s']} \left(\Pi_{m,G}^{\mathbf{c}-t \cdot \mathbf{w}(s)} \right) \right]$$

for each $t \in \mathbb{R}_{\geq 0}$, (ii) $D'(\xi, \bullet, \bullet) = D(\xi, \bullet, \bullet)$ for all $\xi \in \text{Hist}_{\text{sla}}(\mathcal{M}) \setminus \{s\}$ and (iii) $D'(\xi, \bullet) = D(\xi, \bullet)$ for all $\xi \in \text{Hist}_{\text{er}}(\mathcal{M})$, then D' is still a measurable scheduler which satisfies $\Pr_{D, \mathcal{D}[s]}(\Pi_{m+1,G}^{\mathbf{c}}) \leq \Pr_{D', \mathcal{D}[s]}(\Pi_{m+1,G}^{\mathbf{c}})$ (since $D'[s \xrightarrow{a,t}] = D[s \xrightarrow{a,t}]$). Then,

$$\begin{aligned} \text{prob}_{m+1,G}^{\max}(s, \mathbf{c}) &= \sup_{D \in \text{MSched}(\mathcal{M})} \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \\ &\quad \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \Pr_{D[s \xrightarrow{a,t}], \mathcal{D}[s']} \left(\Pi_{m,G}^{\mathbf{c}-t \cdot \mathbf{w}(s)} \right) \right] dt . \end{aligned}$$

We prove that $\sup_{D \in \text{MSched}(\mathcal{M})} \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \max_{a \in \text{En}(s)} F_{\text{la}}^s(D, a, t, \mathbf{c}) dt$ with

$$F_{\text{la}}^s(D, a, t, \mathbf{c}) := \sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \Pr_{D[s \xrightarrow{a,t}], \mathcal{D}[s']} \left(\Pi_{m,G}^{\mathbf{c}-t \cdot \mathbf{w}(s)} \right)$$

(which is essentially $\text{prob}_{m+1,G}^{\max}(s, \mathbf{c})$) equals

$$\int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \max_{a \in \text{En}(s)} G_{\text{la}}^s(a, t, \mathbf{c}) dt$$

with

$$\mathbf{G}_{\text{la}}^s(a, t, \mathbf{c}) := \sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_{m, G}^{\max}(s', \mathbf{c} - t \cdot \mathbf{w}(s)) .$$

It is not difficult to see that the former item is no greater than the latter one. Below we prove the reverse direction. Define

$$a^*(t) := \underset{a \in \text{En}(s)}{\text{argmax}} \mathbf{G}_{\text{la}}^s(a, t, \mathbf{c})$$

where the maximum is chosen w.r.t an arbitrarily-fixed linear order when ties occur. We consider two cases.

Case 1: $\mathbf{w}(s) = \vec{0}$. Then $a^*(t)$ is independent of t , which we shall denote by a^* . For each $\epsilon > 0$, we define the measurable scheduler D^ϵ as follows:

- $D^\epsilon(s, t, \cdot) = \mathcal{D}[a^*]$ for all $t \geq 0$;
- $D^\epsilon(s \xrightarrow{a^*, t} \xi, \tau, \cdot) = D_{\xi[0]}^{m, \epsilon}(\xi, \tau, \cdot)$ for all $t, \tau \in \mathbb{R}_{\geq 0}$ and $\xi \in \text{Hist}_{s_{\text{la}}}(\mathcal{M})$;
- $D^\epsilon(s \xrightarrow{a^*, t} \xi, \cdot) = D_{\xi[0]}^{m, \epsilon}(\xi, \cdot)$ for all $t \in \mathbb{R}_{\geq 0}$ and $\xi \in \text{Hist}_{s_{\text{er}}}(\mathcal{M})$;

the measurable scheduler $D_{\xi[0]}^{m, \epsilon}$ is defined such that

$$\Pr_{D_{\xi[0]}^{m, \epsilon}, \mathcal{D}[\xi[0]]}(\Pi_{m, G}^{\mathbf{c}}) \geq \text{prob}_{m, G}^{\max}(\xi[0], \mathbf{c}) - \epsilon .$$

D^ϵ satisfies that

$$\max_{a \in \text{En}(s)} \mathbf{F}_{\text{la}}^s(D^\epsilon, a, t, \mathbf{c}) \geq \mathbf{F}_{\text{la}}^s(D^\epsilon, a^*, t, \mathbf{c}) \geq \mathbf{G}_{\text{la}}^s(a^*, t, \mathbf{c}) - \epsilon$$

for all $t \geq 0$. Thus,

$$\text{prob}_{m+1, G}^{\max}(s, \mathbf{c}) = \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \mathbf{G}_{\text{la}}^s(a^*, t, \mathbf{c}) dt$$

since ϵ can be arbitrarily chosen.

Case 2: $\mathbf{w}(s) \neq \vec{0}$. Then for all $a \in \text{En}(s)$, $\mathbf{G}_{\text{la}}^s(a, t, \mathbf{c})$ takes non-zero value only on $t \in [0, T_s^{\mathbf{c}}]$ with

$$T_s^{\mathbf{c}} := \min \left\{ \frac{\mathbf{c}_i}{\mathbf{w}_i(s)} \mid 1 \leq i \leq k, \mathbf{w}_i(s) > 0 \right\} .$$

From the induction hypothesis (b), $\max_{a \in \text{En}(s)} \mathbf{G}_{\text{la}}^s(a, \cdot, \mathbf{c})$ is Lipschitz continuous on $[0, T_s^{\mathbf{c}}]$. For each $N \in \mathbb{N}$, we divide the interval $[0, T_s^{\mathbf{c}}]$ into N equal pieces I_1, \dots, I_N with $[t_{j-1}, t_j] := I_j = [\frac{j-1}{N} \cdot T_s^{\mathbf{c}}, \frac{j}{N} \cdot T_s^{\mathbf{c}}]$. Then we construct the measurable scheduler $D^{N, \epsilon}$ for each $\epsilon > 0$ as follows:

- $D^{N, \epsilon}(s, t, \cdot) = \mathcal{D}[a^*(t_j)]$ when $t \in I_j$;
- $D^{N, \epsilon}(s \xrightarrow{a^*(t_j), t} \xi, \cdot, \cdot) = D_{\xi[0], j}^{m, \epsilon}(\xi, \cdot, \cdot)$ when $t \in I_j$ and $\xi \in \text{Hist}_{s_{\text{la}}}(\mathcal{M})$;

- $D^{N,\epsilon}(s \xrightarrow{a^*(t_j),t} \xi, \cdot) = D_{\xi[0],j}^{m,\epsilon}(\xi, \cdot)$ when $t \in I_j$ and $\xi \in \text{Hist}_{S_{\text{er}}}(\mathcal{M})$;

the measurable scheduler $D_{\xi[0],j}^{m,\epsilon}$ is chosen such that

$$\Pr_{D_{\xi[0],j}^{m,\epsilon}, \mathcal{D}[\xi[0]]} \left(\Pi_{m,G}^{\mathbf{c}-t_j \cdot \mathbf{w}(s)} \right) \geq \text{prob}_{m,G}^{\max}(\xi[0], \mathbf{c} - t_j \cdot \mathbf{w}(s)) - \epsilon .$$

From the inductive hypothesis (b) and the monotonicity of $\Pi_{m,G}^{\mathbf{c}-t \cdot \mathbf{w}(s)}$ on t , we obtain

$$\begin{aligned} \lim_{\substack{N \rightarrow +\infty \\ \epsilon \rightarrow 0^+}} \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \max_{a \in \text{En}(s)} F_{\text{la}}^s(D^{N,\epsilon}, a, t, \mathbf{c}) dt = \\ \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \max_{a \in \text{En}(s)} G_{\text{la}}^s(a, t, \mathbf{c}) dt \end{aligned}$$

which implies the inductive step for (c).

It remains to show that the inductive step for (b) holds. Let $\mathbf{c}, \mathbf{c}' \geq \vec{0}$ and $s \in S - G$. Denote $\delta := \|\mathbf{c} - \mathbf{c}'\|_\infty$. Assume that $s \in S_{\text{er}}$. Consider an arbitrary $a \in \text{En}(s)$. If $\mathbf{w}(s, a) = \vec{0}$, then clearly $|G_{\text{er}}^s(a, \mathbf{c}) - G_{\text{er}}^s(a, \mathbf{c}')| \leq \epsilon(m, \delta)$. Otherwise,

$$\begin{aligned} & |G_{\text{er}}^s(a, \mathbf{c}) - G_{\text{er}}^s(a, \mathbf{c}')| \\ & \leq \int_0^T f_{\mathbf{E}(s,a)}(t) \cdot \epsilon(m, \delta) dt + f_{\mathbf{E}(s,a)}(T) \cdot |T_{s,a}^{\mathbf{c}} - T_{s,a}^{\mathbf{c}'}| \\ & \leq (1 - e^{-\mathbf{E}(s,a) \cdot T}) \cdot \epsilon(m, \delta) + e^{-\mathbf{E}(s,a) \cdot T} \cdot \frac{\mathbf{E}_{\max}}{\mathbf{w}_{\min}} \cdot \delta \\ & \leq \frac{\mathbf{E}_{\max}}{\mathbf{w}_{\min}} \cdot \delta \end{aligned}$$

where $T := \min\{T_{s,a}^{\mathbf{c}}, T_{s,a}^{\mathbf{c}'}\}$ and the last step is obtained through induction hypothesis (b). It follows that

$$\begin{aligned} & |\text{prob}_{m+1,G}^{\max}(s, \mathbf{c}) - \text{prob}_{m+1,G}^{\max}(s, \mathbf{c}')| \\ & \leq \max_{a \in \text{En}(s)} |G_{\text{er}}^s(a, \mathbf{c}) - G_{\text{er}}^s(a, \mathbf{c}')| \\ & \leq \frac{\mathbf{E}_{\max}}{\mathbf{w}_{\min}} \cdot \delta . \end{aligned}$$

Assume now that $s \in S_{\text{la}}$. Consider any $a \in \text{En}(s)$. If $\mathbf{w}(s) = \vec{0}$, then

$$|G_{\text{la}}^s(a, t, \mathbf{c}) - G_{\text{la}}^s(a, t, \mathbf{c}')| \leq \epsilon(m, \delta)$$

for all $t \geq 0$. It follows that $|\text{prob}_{m+1,G}^{\max}(s, \mathbf{c}) - \text{prob}_{m+1,G}^{\max}(s, \mathbf{c}')| \leq \epsilon(m, \delta)$. Otherwise, by the inductive step (c) and the induction hypothesis (b), we

have

$$\begin{aligned}
& \left| \text{prob}_{m+1,G}^{\max}(s, \mathbf{c}) - \text{prob}_{m+1,G}^{\max}(s, \mathbf{c}') \right| \\
& \leq \int_0^T f_{\mathbf{E}(s)}(t) \cdot \epsilon(m, \delta) dt + f_{\mathbf{E}(s)}(T) \cdot \left| T_s^{\mathbf{c}} - T_s^{\mathbf{c}'} \right| \\
& \leq (1 - e^{-\mathbf{E}(s) \cdot T}) \cdot \epsilon(m, \delta) + e^{-\mathbf{E}(s) \cdot T} \cdot \frac{\mathbf{E}_{\max}}{\mathbf{w}_{\min}} \cdot \delta \\
& \leq \frac{\mathbf{E}_{\max}}{\mathbf{w}_{\min}} \cdot \delta,
\end{aligned}$$

where $T := \min\{T_s^{\mathbf{c}}, T_s^{\mathbf{c}'}\}$. Thus the inductive step for (b) is completed.

Secondly, we prove that $\lim_{n \rightarrow \infty} \text{prob}_{n,G}^{\max} = \text{prob}_G^{\max}$. From definition, $\text{prob}_{n,G}^{\max} \leq \text{prob}_{n+1,G}^{\max}$ and $\text{prob}_{n,G}^{\max} \leq \text{prob}_G^{\max}$ for all $n \geq 0$. Thus the limit function $\lim_{n \rightarrow \infty} \text{prob}_{n,G}^{\max}$ exists and is no greater than prob_G^{\max} . For the reverse direction, let $s \in S$ and $\mathbf{c} \in \mathbb{R}^k$. Fix an arbitrary $\epsilon > 0$. Let D be a measurable scheduler such that $\Pr_{D, \mathcal{D}[s]}(\Pi_G^{\mathbf{c}}) \geq \text{prob}_G^{\max}(s, \mathbf{c}) - \epsilon$. By definition, $\text{prob}_{n,G}^{\max}(s, \mathbf{c}) \geq \Pr_{D, \mathcal{D}[s]}(\Pi_{n,G}^{\mathbf{c}})$. It follows that

$$\lim_{n \rightarrow \infty} \text{prob}_{n,G}^{\max}(s, \mathbf{c}) \geq \lim_{n \rightarrow \infty} \Pr_{D, \mathcal{D}[s]}(\Pi_{n,G}^{\mathbf{c}}) = \Pr_{D, \mathcal{D}[s]}(\Pi_G^{\mathbf{c}}) \geq \text{prob}_G^{\max}(s, \mathbf{c}) - \epsilon.$$

Thus $\lim_{n \rightarrow \infty} \text{prob}_{n,G}^{\max}(s, \mathbf{c}) = \text{prob}_G^{\max}(s, \mathbf{c})$ since ϵ is arbitrarily chosen.

Thirdly, we prove that prob_G^{\max} is the least fixed-point of \mathcal{T}_G . It is clear that $\text{prob}_G^{\max}(s, \mathbf{c}) = \mathbf{1}_{\mathbb{R}_{\geq 0}^k}(\mathbf{c})$ if $s \in G$. By applying Monotone Convergence Theorem (Theorem 3.2) on (c), we obtain $\text{prob}_G^{\max}(s, \mathbf{c}) = \mathcal{T}_G(\text{prob}_G^{\max})(s, \mathbf{c})$ when $s \in S \setminus G$. Thus, prob_G^{\max} is a fixed-point of \mathcal{T}_G . To see that it is the least fixed-point of \mathcal{T}_G , one can proceed by induction on $n \geq 0$ that given any fixed-point h of \mathcal{T}_G , $\text{prob}_{n,G}^{\max} \leq h$ for all $n \geq 0$ by the facts that $\text{prob}_{0,G}^{\max} \leq h$ and $\text{prob}_{n+1,G}^{\max} = \mathcal{T}_G(\text{prob}_{n,G}^{\max})$. It follows that $\text{prob}_G^{\max} \leq h$ for any fixed-point h of \mathcal{T}_G .

Finally, by taking the limit from (b) we can obtain that

$$\left| \text{prob}_G^{\max}(s, \mathbf{c}) - \text{prob}_G^{\max}(s, \mathbf{c}') \right| \leq \frac{\mathbf{E}_{\max}}{\mathbf{w}_{\min}} \cdot \|\mathbf{c} - \mathbf{c}'\|_{\infty}$$

for all $\mathbf{c}, \mathbf{c}' \geq \vec{0}$ and $s \in S$. □

The Lipschitz constant $\frac{\mathbf{E}_{\max}}{\mathbf{w}_{\min}}$ will be crucial to the error bound of our approximation algorithm.

Remark 9.2. We describe the proof error in [57, 59]. The error lies in the proof of [57, Lemma 5.1 on Pages 119] which tries to prove that the time-bounded reachability probability functions are continuous. In detail, the error is at the proof for right-continuity of the functions. Let us take the sentence “This implies ... for some $\xi \leq \frac{\epsilon}{2}$.” from line -3 to line -2 on page

119 as (*). (*) is wrong in general, as one can treat D 's as natural numbers, and define

$$\Pr_n(\text{"reach } G \text{ within } z") := \begin{cases} n \cdot z & \text{if } z \in [0, \frac{1}{n}] \\ 1 & \text{if } z \in (\frac{1}{n}, \infty) \end{cases} .$$

Then $\sup_n \Pr_n(\text{"reach } G \text{ within } z")$ equals 1 for $z > 0$ and 0 for $z = 0$. Thus $\sup_D \Pr_D(\text{"reach } G \text{ within } z")$ on $z \geq 0$ is right-discontinuous at $z = 0$, which does not satisfy (*) (treat D as a natural number). Note that [57, Lemma 5.1] is correct; it is the proof that is flawed. Also note that Lemma 5.1 is important as the least fixed-point characterization [57, Theorem 5.1 on Page 120] and the optimal scheduler [57, Theorem 5.2 on page 124] directly rely on it. We fix this error in the more general setting of cost-bounded reachability probability through a new methodology as illustrated in this section.

9.2 Optimal Measurable Schedulers

In this section, we establish optimal measurable schedulers for maximal cost-bounded reachability probability. We show that there exists a deterministic cost-positional measurable scheduler that achieves the maximal cost-bounded reachability probability. Below we fix a CTMDP $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$. We first introduce the notion of deterministic cost-positional schedulers.

Definition 9.4. *A measurable scheduler D is called deterministic cost-positional iff the following conditions hold:*

- for all $\xi, \xi' \in \text{Hists}_{\text{er}}(\mathcal{M})$, if $\xi \downarrow = \xi' \downarrow$ and $\mathbf{C}(\xi) = \mathbf{C}(\xi')$, then $D(\xi, \cdot) = D(\xi', \cdot)$;
- for all $\xi, \xi' \in \text{Hists}_{\text{la}}(\mathcal{M})$ and $t, t' \in \mathbb{R}_{\geq 0}$, if $\xi \downarrow = \xi' \downarrow$ and $\mathbf{C}(\xi) + t \cdot \mathbf{w}(\xi \downarrow) = \mathbf{C}(\xi') + t' \cdot \mathbf{w}(\xi' \downarrow)$, then $D(\xi, t, \cdot) = D(\xi', t', \cdot)$;
- for all $\xi \in \text{Hists}_{\text{er}}(\mathcal{M})$, $D(\xi, \cdot)$ is Dirac;
- for all $\xi \in \text{Hists}_{\text{la}}(\mathcal{M})$ and $t \in \mathbb{R}_{\geq 0}$, $D(\xi, t, \cdot)$ is Dirac.

Intuitively, a deterministic cost-positional scheduler makes its decision solely on the current state and the cost accumulated so far, and its decision is always Dirac. Below we show that such a scheduler suffices to achieve maximal cost-bounded reachability probability.

Theorem 9.2. *For all $\mathbf{c} \in \mathbb{R}^k$ and $G \subseteq S$, there exists a measurable deterministic cost-positional scheduler $D_{\mathbf{c}}$ such that $\text{prob}_G^{\max}(s, \mathbf{c}) = \Pr_{D_{\mathbf{c}}, D[s]}(\Pi_G^{\mathbf{c}})$ for all $s \in S$.*

Proof. Let $G \subseteq S$. Fix an arbitrary linear order \preceq on Act . Consider some $\mathbf{c} \in \mathbb{R}^k$. We define the function $\mathbf{G}_{\text{er}} : S_{\text{er}} \times Act \times \mathbb{R}^k \rightarrow [0, 1]$ by:

$$\mathbf{G}_{\text{er}}(s, a, \mathbf{x}) := \int_0^\infty f_{\mathbf{E}(s,a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} - \mathbf{x} - t \cdot \mathbf{w}(s, a)) \right] dt .$$

We further define the function $\mathbf{G}_{\text{la}} : S_{\text{la}} \times Act \times \mathbb{R}^k \times \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ by:

$$\mathbf{G}_{\text{la}}(s, a, \mathbf{x}, t) := \sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} - \mathbf{x} - t \cdot \mathbf{w}(s)) .$$

Note that $\text{prob}_G^{\max}(s, \mathbf{c} - \mathbf{x}) = \max_{a \in \text{En}(s)} \mathbf{G}_{\text{er}}(s, a, \mathbf{x})$ if $s \in S_{\text{er}} - G$; and

$$\text{prob}_G^{\max}(s, \mathbf{c} - \mathbf{x}) = \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \max_{a \in \text{En}(s)} \mathbf{G}_{\text{la}}(s, a, \mathbf{x}, t) dt$$

if $s \in S_{\text{la}} - G$. We construct the measurable scheduler $\mathbf{D}_{\mathbf{c}}$ as follows. Consider an arbitrary $\xi \in \text{Hists}(\mathcal{M})$. Define

$$X_{\neq \vec{0}}^\xi := \left\{ s \in S_{\text{er}} \mid \exists a^* \in \text{En}(s) \cdot \left[\mathbf{G}_{\text{er}}(s, a^*, \mathbf{C}(\xi)) = \max_{a \in \text{En}(s)} \mathbf{G}_{\text{er}}(s, a, \mathbf{C}(\xi)) \right. \right. \\ \left. \left. \wedge \mathbf{w}(s, a^*) \neq \vec{0} \right] \right\} \cup \left\{ s \in S_{\text{la}} \mid \mathbf{w}(s) \neq \vec{0} \right\}$$

and $X_{=0}^\xi := \{s \in S \mid \text{prob}_G^{\max}(s, \mathbf{c} - \mathbf{C}(\xi)) = 0\}$. Note that the definitions of $X_{\neq \vec{0}}^\xi$ and $X_{=0}^\xi$ depend only on $\mathbf{C}(\xi)$. The probability distribution $\mathbf{D}_{\mathbf{c}}(\xi, \cdot), \mathbf{D}_{\mathbf{c}}(\xi, \cdot, \cdot)$ is determined by the following procedure.

1. If $\xi \downarrow \in X_{=0}^\xi \cap S_{\text{er}}$, then we set $\mathbf{D}_{\mathbf{c}}(\xi, \cdot) = \mathcal{D}[a_{=0}^{\xi \downarrow}]$ where $a_{=0}^{\xi \downarrow} \in \text{En}(\xi \downarrow)$ is an arbitrarily fixed action.
2. If $\xi \downarrow \in X_{=0}^\xi \cap S_{\text{la}}$ and $t \in \mathbb{R}_{\geq 0}$, then we set $\mathbf{D}_{\mathbf{c}}(\xi, t, \cdot) = \mathcal{D}[a_{=0}^{\xi \downarrow}]$, where the choice of $a_{=0}^{\xi \downarrow}$ is the same as above.
3. If $\xi \downarrow \in S_{\text{er}} \cap (X_{\neq \vec{0}}^\xi \setminus X_{=0}^\xi)$, then we set $\mathbf{D}_{\mathbf{c}}(\xi, \cdot) = \mathcal{D}[a^\xi]$, where $a^\xi \in \text{En}(\xi \downarrow)$ satisfies that

$$\mathbf{G}_{\text{er}}(\xi \downarrow, a^\xi, \mathbf{C}(\xi)) = \max_{a \in \text{En}(\xi \downarrow)} \mathbf{G}_{\text{er}}(\xi \downarrow, a, \mathbf{C}(\xi)) \text{ and } \mathbf{w}(\xi \downarrow, a^\xi) \neq \vec{0} ;$$

if there are multiple such a^ξ 's, we choose the least of them w.r.t \preceq . Note that the choice of a^ξ depends only on $\mathbf{C}(\xi)$ and $\xi \downarrow$.

4. If $\xi \downarrow \in S_{\text{la}} \cap (X_{\neq \vec{0}}^\xi \setminus X_{=0}^\xi)$ and $t \in \mathbb{R}_{\geq 0}$, then we set $\mathbf{D}_{\mathbf{c}}(\xi, t, \cdot) = \mathcal{D}[a^{\xi, t}]$, where $a^{\xi, t} \in \text{En}(\xi \downarrow)$ satisfies

$$\mathbf{G}_{\text{la}}(\xi \downarrow, a^{\xi, t}, \mathbf{C}(\xi), t) = \max_{a \in \text{En}(\xi \downarrow)} \mathbf{G}_{\text{la}}(\xi \downarrow, a, \mathbf{C}(\xi), t) ;$$

if there are multiple such $a^{\xi, t}$'s, we choose the least of them w.r.t \preceq . Note that the choice of $a^{\xi, t}$ depends only on $\mathbf{C}(\xi) + t \cdot \mathbf{w}(\xi \downarrow)$ and $\xi \downarrow$.

5. If $\xi \downarrow \in S_{\text{er}} \setminus (X_{\neq 0}^{\xi} \cup X_{=0}^{\xi})$, then we set $D_{\mathbf{c}}(\xi, \cdot)$ to be the Dirac distribution at an action $a^{\xi} \in \text{En}(\xi \downarrow)$ which satisfies that

$$G_{\text{er}}(\xi \downarrow, a^{\xi}, \mathbf{C}(\xi)) = \max_{a \in \text{En}(\xi \downarrow)} G_{\text{er}}(\xi \downarrow, a, \mathbf{C}(\xi))$$

and there exists $s \in S$ such that $\mathbf{P}(\xi \downarrow, a^{\xi}, s) > 0$ and the distance from s to $X_{\neq 0}^{\xi} \setminus X_{=0}^{\xi}$ is (one-step) smaller than that from $\xi \downarrow$ in the digraph \mathcal{G}^{ξ} (to be defined later in this proof). If there are multiple such a^{ξ} 's, we choose the least of them w.r.t \preceq .

6. If $\xi \downarrow \in S_{\text{la}} \setminus (X_{\neq 0}^{\xi} \cup X_{=0}^{\xi})$ and $t \in \mathbb{R}_{\geq 0}$, then we set $D_{\mathbf{c}}(\xi, t, \cdot)$ to be the Dirac distribution at an action $a^{\xi, t} \in \text{En}(\xi \downarrow)$ which satisfies that

$$G_{\text{la}}(\xi \downarrow, a^{\xi, t}, \mathbf{C}(\xi), t) = \max_{a \in \text{En}(\xi \downarrow)} G_{\text{la}}(\xi \downarrow, a^{\xi, t}, \mathbf{C}(\xi), t)$$

and there exists $s \in S$ such that $\mathbf{P}(\xi \downarrow, a^{\xi, t}, s) > 0$ and the distance from s to $X_{\neq 0}^{\xi} \setminus X_{=0}^{\xi}$ is (one-step) smaller than that from $\xi \downarrow$ in the digraph \mathcal{G}^{ξ} . If there are multiple such $a^{\xi, t}$'s, we choose the least of them w.r.t \preceq . Note that $a^{\xi, t}$ only depends on $\mathbf{C}(\xi) + t \cdot \mathbf{w}(\xi \downarrow)$ and $\xi \downarrow$, and the dwell-time t does not affect the choice of $a^{\xi, t}$ since $\mathbf{w}(\xi \downarrow) = \vec{0}$.

The digraph \mathcal{G}^{ξ} is defined such that its vertex set is S , and its edge set is the set of all $(u, v) \in S \times S$ such that either $u \in S_{\text{er}}$ and

$$\exists b \in \text{En}(u). [\mathbf{P}(u, b, v) > 0 \wedge G_{\text{er}}(u, b, \mathbf{C}(\xi)) = \max_{a \in \text{En}(u)} G_{\text{er}}(u, a, \mathbf{C}(\xi))]$$

holds, or $u \in S_{\text{la}}$ and

$$\exists b \in \text{En}(u). [\mathbf{P}(u, b, v) > 0 \wedge G_{\text{la}}(u, b, \mathbf{C}(\xi), 0) = \max_{a \in \text{En}(u)} G_{\text{la}}(u, a, \mathbf{C}(\xi), 0)]$$

holds. The legitimacy of the third step in the procedure above follows from Proposition 9.1 to be proved later: the set $X' \subseteq S \setminus (X_{\neq 0}^{\xi} \cup X_{=0}^{\xi})$ of states that cannot reach $X_{\neq 0}^{\xi} \setminus X_{=0}^{\xi}$ should be empty, or otherwise one can reduce all values in

$$\{\text{prob}_G^{\max}(s, \mathbf{c} - \mathbf{C}(\xi))\}_{s \in X'}$$

by a small amount, to obtain a pre-fixed-point which is strictly smaller than $\{\text{prob}_G^{\max}(s, \mathbf{c} - \mathbf{C}(\xi))\}_{s \in S}$. By definition, $D_{\mathbf{c}}$ is deterministic cost-positional. Note that there are finitely many triples $(X_{\neq 0}^{\xi}, X_{=0}^{\xi}, \mathcal{G}^{\xi})$ since S is finite. By Theorem 9.1, the function $\mathbf{x} \mapsto G_{\text{er}}(s, a, \mathbf{x})$ (resp. $(\mathbf{x}, t) \mapsto G_{\text{la}}(s, a, \mathbf{x}, t)$) is separately continuous on $\{\mathbf{x} \mid \mathbf{x} \leq \mathbf{c}\}$ (resp. $\{(\mathbf{x}, t) \mid \mathbf{x} + t \cdot \mathbf{w}(s) \leq \mathbf{c}\}$) and its complement set, for all $s \in S$ and $a \in \text{En}(s)$. Thus, the set of all histories ξ with length n such that the triple $(X_{\neq 0}^{\xi}, X_{=0}^{\xi}, \mathcal{G}^{\xi})$ happens to be a specific one

is measurable w.r.t $(\Omega_{\mathcal{M}}^n, \mathcal{S}_{\mathcal{M}}^n)$. It follows that $\mathbf{D}_{\mathbf{c}}$ is a measurable scheduler. Below we prove by contradiction that $\text{prob}_G^{\max}(s, \mathbf{c}) = \Pr_{\mathbf{D}_{\mathbf{c}}, \mathcal{D}[s]}(\Pi_G^{\mathbf{c}})$ for all $s \in S$ and $\mathbf{c} \geq \vec{0}$.

Define the function $h : S \times \mathbb{R}^k \rightarrow [0, 1]$ by $h(s, \mathbf{c}) := \Pr_{\mathbf{D}_{\mathbf{c}}, \mathcal{D}[s]}(\Pi_G^{\mathbf{c}})$. Suppose that $\text{prob}_G^{\max} \neq h$. Let $\bar{h} := \text{prob}_G^{\max} - h$. (Note that $\text{prob}_G^{\max} \geq h$.) Then there exists $s \in S$ and $\mathbf{c} \geq \vec{0}$ such that $\bar{h}(s, \mathbf{c}) > 0$. Define $T := \max_{1 \leq i \leq k} \left\{ \frac{c_i}{\mathbf{w}_{\min}} \right\}$. Let

$$d := \sup \left\{ \bar{h}(s, \mathbf{c}') \mid s \in S, \vec{0} \leq \mathbf{c}' \leq \mathbf{c} \right\}$$

and

$$d' := \sup \left\{ \bar{h}(s, \mathbf{c}') \mid \vec{0} \leq \mathbf{c}' \leq \mathbf{c}, \text{ and either } s \in S_{\text{er}} \text{ and } \mathbf{w}(s, a^*) \neq \vec{0} \right. \\ \left. \text{with } \mathcal{D}[a^*] = \mathbf{D}_{\mathbf{c}'}(s, \cdot) \text{ or } s \in S_{\text{la}} \text{ with } \mathbf{w}(s) \neq \vec{0} \right\} .$$

We first show that $d' < d$ by a (nested) contradiction proof.

Suppose $d' = d$. Choose $\epsilon > 0$ such that $d > e^{\mathbf{E}_{\max} \cdot T} \cdot \epsilon$. Then choose $s \in S$ and $\vec{0} \leq \mathbf{c}' \leq \mathbf{c}$ such that $d - \epsilon < \bar{h}(s, \mathbf{c}') \leq d$, and either $s \in S_{\text{er}}$ with $\mathbf{w}(s, a^*) \neq \vec{0}$ ($\mathcal{D}[a^*] = \mathbf{D}_{\mathbf{c}'}(s, \cdot)$) or $s \in S_{\text{la}}$ with $\mathbf{w}(s) \neq \vec{0}$. On one hand, by Theorem 7.1, if $s \in S_{\text{er}}$ then

$$h(s, \mathbf{c}') = \int_0^\infty f_{\mathbf{E}(s, a^*)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a^*, s') \cdot h(s', \mathbf{c}' - t \cdot \mathbf{w}(s, a^*)) \right] dt;$$

then with Theorem 9.1, we obtain

$$\bar{h}(s, \mathbf{c}') \leq \int_0^T f_{\mathbf{E}(s, a^*)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a^*, s') \cdot \bar{h}(s', \mathbf{c}' - t \cdot \mathbf{w}(s, a^*)) \right] dt .$$

On the other hand, by Theorem 7.1, if $s \in S_{\text{la}}$ then

$$h(s, \mathbf{c}') = \int_0^\infty f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a^*(t), s') \cdot h(s', \mathbf{c}' - t \cdot \mathbf{w}(s)) \right] dt$$

with $\mathcal{D}[a^*(t)] = \mathbf{D}_{\mathbf{c}'}(s, t, \cdot)$. Then with Theorem 9.1, we obtain

$$\bar{h}(s, \mathbf{c}') \leq \int_0^T f_{\mathbf{E}(s)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a^*(t), s') \cdot \bar{h}(s', \mathbf{c}' - t \cdot \mathbf{w}(s)) \right] dt .$$

In either case, we can obtain $d - \epsilon \leq (1 - e^{-\mathbf{E}_{\max} \cdot T}) \cdot d$. This implies $d \leq e^{\mathbf{E}_{\max} \cdot T} \cdot \epsilon$. Contradiction to the choice of ϵ .

Thus $d > d'$. Let $\delta := d - d'$ and $\epsilon := \mathbf{P}_{\min}^{|S|} \cdot \delta$. We inductively construct a finite sequence s_0, s_1, \dots, s_l in S ($1 \leq l \leq |S|$) which satisfies

$$\bar{h}(s_i, \mathbf{c}') > d - \mathbf{P}_{\min}^{-i} \cdot \epsilon \quad (i = 0, \dots, l)$$

for some fixed \mathbf{c}' , as follows. Note that the triple $(X_{\neq \vec{0}}^s, X_{=0}^s, \mathcal{G}^s)$ for $s \in S$ (w.r.t any $\mathbf{c}'' \in \mathbb{R}^k$) remains constant as s varies, since $\mathbf{C}(s) = \vec{0}$.

1. Initially, we set $i = 0$ and choose $s_0 \in S$ and $\vec{0} \leq \mathbf{c}' \leq \mathbf{c}$ such that

$$d - \epsilon < \bar{h}(s_0, \mathbf{c}') \leq d .$$

2. As long as both $i \leq l$ and $0 \leq d - \mathbf{P}_{\min}^{-i} \cdot \epsilon < \bar{h}(s_i, \mathbf{c}') \leq d$ holds, $s_i \in S \setminus (X_{\neq \vec{0}}^{s_i} \cup X_{=0}^{s_i})$ (w.r.t $\mathbf{D}_{\mathbf{c}'}$) since $\mathbf{P}_{\min}^{-i} \cdot \epsilon \leq \delta$. On one hand, assume that $s_i \in S_{\text{er}}$. Then

$$h(s_i, \mathbf{c}') = \sum_{s' \in S} \mathbf{P}(s_i, a^*, s') \cdot h(s', \mathbf{c}') \text{ with } \mathcal{D}[a^*] = \mathbf{D}_{\mathbf{c}'}(s_i)$$

and there exists $u \in S$ such that (i) $\mathbf{P}(s_i, a^*, u) > 0$ and (ii) via u the distance to $X_{\neq \vec{0}}^{s_i} \setminus X_{=0}^{s_i}$ in \mathcal{G}^{s_i} (w.r.t $\mathbf{D}_{\mathbf{c}'}$) is decreased by one. Moreover, from

$$\bar{h}(s_i, \mathbf{c}') \leq \sum_{s' \in S} \mathbf{P}(s_i, a^*, s') \cdot \bar{h}(s', \mathbf{c}')$$

we obtain

$$d - \mathbf{P}_{\min}^{-i} \cdot \epsilon < (1 - \mathbf{P}_{\min}) \cdot d + \mathbf{P}_{\min} \cdot \bar{h}(u, \mathbf{c}') ,$$

which can be further reduced to $\bar{h}(u, \mathbf{c}') > d - \mathbf{P}_{\min}^{-(i+1)} \cdot \epsilon$. On the other hand, assume that $s_i \in S_{\text{la}}$. Since $\mathbf{w}(s_i) = \vec{0}$, $\mathbf{D}_{\mathbf{c}'}(s_i, t, \cdot)$ remains constant when t varies. Let a^* be the action such that $\mathcal{D}[a^*] = \mathbf{D}_{\mathbf{c}'}(s_i, 0, \cdot)$. Then

$$h(s_i, \mathbf{c}') = \sum_{s' \in S} \mathbf{P}(s_i, a^*, s') \cdot h(s', \mathbf{c}')$$

and there exists $u \in S$ such that $\mathbf{P}(s_i, a^*, u) > 0$ and via u the distance to $X_{\neq \vec{0}}^{s_i} \setminus X_{=0}^{s_i}$ in \mathcal{G}^{s_i} (w.r.t $\mathbf{D}_{\mathbf{c}'}$) is decreased by one. Moreover, from

$$\bar{h}(s_i, \mathbf{c}') \leq \sum_{s' \in S} \mathbf{P}(s_i, a^*, s') \cdot \bar{h}(s', \mathbf{c}')$$

we obtain

$$d - \mathbf{P}_{\min}^{-i} \cdot \epsilon < (1 - \mathbf{P}_{\min}) \cdot d + \mathbf{P}_{\min} \cdot \bar{h}(u, \mathbf{c}') ,$$

which is further reduced to $\bar{h}(u, \mathbf{c}') > d - \mathbf{P}_{\min}^{-(i+1)} \cdot \epsilon$.

3. In either case in the previous step, we set $s_{i+1} := u$.
4. If $s_{i+1} \in X_{\neq \vec{0}}^{s_{i+1}} \setminus X_{=0}^{s_{i+1}}$, then the construction is terminated. Otherwise, back to Step 2.

The legitimacy and termination (within $|S|$ steps) of the inductive construction follows directly from the definition of $\mathbf{D}_{\mathbf{c}'}$. By $s_l \in X_{\neq \vec{0}}^{s_l} \setminus X_{=0}^{s_l}$, we obtain

$$d - \delta \geq \bar{h}(s_l, \mathbf{c}') > d - \mathbf{P}_{\min}^{-l} \cdot \epsilon \quad (l \leq |S|) ,$$

which is a contradiction due to $\epsilon = \mathbf{P}_{\min}^{|S|} \cdot \delta$. Thus $\text{prob}_G^{\max} = h$. \square

9.3 Differential Characterizations

In this section, we derive differential characterizations for the functions prob_G^{\max} . These differential characterizations will be fundamental to our approximation algorithms.

Below we fix a CTMDP $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$, a set $G \subseteq S$ and a cost function $\mathbf{w} : S \times \text{Act} \rightarrow \mathbb{R}_{\geq 0}^k$. To introduce the differential characterization, we first extend the function prob_G^{\max} as follows.

Definition 9.5. Let $Z_G := \{(s, a) \in (S_{\text{er}} - G) \times \text{Act} \mid a \in \text{En}(s)\}$. Define $\text{prob}_G^{\max} : Z_G \times \mathbb{R}^k \rightarrow [0, 1]$ by

$$\text{prob}_G^{\max}((s, a), \mathbf{c}) := \int_0^\infty f_{\mathbf{E}(s, a)}(t) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} - t \cdot \mathbf{w}(s, a)) \right] dt$$

for $(s, a) \in Z_G$ and $\mathbf{c} \in \mathbb{R}^k$.

Intuitively, we extend prob_G^{\max} to Z_G . By Theorem 9.1, one easily sees that $\text{prob}_G^{\max}(s, \mathbf{c}) = \max_{a \in \text{En}(s)} \text{prob}_G^{\max}((s, a), \mathbf{c})$ for all $s \in S_{\text{er}} - G$ and $\mathbf{c} \in \mathbb{R}^k$.

The following two definitions introduce a sort of directional derivative which will be crucial in our approximation algorithm.

Definition 9.6. Let $\mathbf{c} \in \mathbb{R}_{\geq 0}^k$. For $z \in Z_G \cup (S_{\text{la}} - G)$, we define

$$\nabla^+ \text{prob}_G^{\max}(z, \mathbf{c}) := \lim_{t \rightarrow 0^+} \frac{\text{prob}_G^{\max}(z, \mathbf{c} + t \cdot \mathbf{w}(z)) - \text{prob}_G^{\max}(z, \mathbf{c})}{t} .$$

If $\mathbf{c}_i > 0$ whenever $\mathbf{w}_i(z) > 0$ ($1 \leq i \leq k$), we further define

$$\nabla^- \text{prob}_G^{\max}(z, \mathbf{c}) := \lim_{t \rightarrow 0^-} \frac{\text{prob}_G^{\max}(z, \mathbf{c} + t \cdot \mathbf{w}(z)) - \text{prob}_G^{\max}(z, \mathbf{c})}{t} ;$$

otherwise, let $\nabla^- \text{prob}_G^{\max}(z, \mathbf{c}) := \nabla^+ \text{prob}_G^{\max}(z, \mathbf{c})$.

Thus $\nabla^+ \text{prob}_G^{\max}(z, \mathbf{c})$ (resp. $\nabla^- \text{prob}_G^{\max}(z, \mathbf{c})$) is the right (resp. left) directional derivative along the vector $\mathbf{w}(z)$. The following theorem gives a characterization for $\nabla^+ \text{prob}_G^{\max}(z, \mathbf{c})$ and $\nabla^- \text{prob}_G^{\max}(z, \mathbf{c})$.

Theorem 9.3. For all $z \in Z_G \cup (S_{\text{la}} \setminus G)$ and $\mathbf{c} \in \mathbb{R}_{\geq 0}^k$, $\nabla^+ \text{prob}_G^{\max}(z, \mathbf{c}) = \nabla^- \text{prob}_G^{\max}(z, \mathbf{c})$. Moreover,

$$\nabla^+ \text{prob}_G^{\max}((s, a), \mathbf{c}) = \sum_{s' \in S} \mathbf{R}(s, a, s') \cdot (\text{prob}_G^{\max}(s', \mathbf{c}) - \text{prob}_G^{\max}((s, a), \mathbf{c}))$$

for all $(s, a) \in Z_G$ and $\mathbf{c} \in \mathbb{R}_{\geq 0}^k$. And

$$\nabla^+ \text{prob}_G^{\max}(s, \mathbf{c}) = \max_{a \in \text{En}(s)} \sum_{s' \in S} \mathbf{R}(s, a, s') \cdot (\text{prob}_G^{\max}(s', \mathbf{c}) - \text{prob}_G^{\max}(s, \mathbf{c})) .$$

for all $s \in S_{\text{la}} \setminus G$ and $\mathbf{c} \in \mathbb{R}_{\geq 0}^k$. The function $\mathbf{R} : S \times \text{Act} \times S \rightarrow \mathbb{R}_{\geq 0}$ is defined as follows:

- $\mathbf{R}(s, a, s') := \mathbf{E}(s, a) \cdot \mathbf{P}(s, a, s')$ when $s \in S_{\text{er}}$;
- $\mathbf{R}(s, a, s') := \mathbf{E}(s) \cdot \mathbf{P}(s, a, s')$ when $s \in S_{\text{la}}$;

Proof. Let $z \in Z_G \cup (S_{\text{la}} \setminus G)$ and $\mathbf{c} \geq \vec{0}$. We first assume that $z \in Z_G$ and $z = (s, a)$. Define the function $h[(s, a), \mathbf{c}']$ to be

$$\tau \mapsto f_{\mathbf{E}(s, a)}(\tau) \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c}' - \tau \cdot \mathbf{w}(s, a)) \right].$$

Denote $\lambda := \mathbf{E}(s, a)$. Consider $\nabla^+ \text{prob}_G^{\max}$. By definition, for all $t > 0$,

$$\begin{aligned} & \text{prob}_G^{\max}((s, a), \mathbf{c} + t \cdot \mathbf{w}(s, a)) \\ &= \int_0^\infty h[(s, a), \mathbf{c} + t \cdot \mathbf{w}(s, a)](\tau) \, d\tau \\ &= \int_0^t h[(s, a), \mathbf{c} + t \cdot \mathbf{w}(s, a)](\tau) \, d\tau + \int_t^\infty h[(s, a), \mathbf{c} + t \cdot \mathbf{w}(s, a)](\tau) \, d\tau \\ &= e^{-\lambda t} \cdot \int_0^t \lambda \cdot e^{\lambda \tau} \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} + \tau \cdot \mathbf{w}(s, a)) \right] \, d\tau \\ &\quad + e^{-\lambda t} \cdot \int_0^\infty h[(s, a), \mathbf{c}](\tau) \, d\tau \\ &= e^{-\lambda t} \cdot \int_0^t \lambda \cdot e^{\lambda \tau} \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} + \tau \cdot \mathbf{w}(s, a)) \right] \, d\tau \\ &\quad + e^{-\lambda t} \cdot \text{prob}_G^{\max}((s, a), \mathbf{c}) \end{aligned}$$

where the third equality is obtained by the variable substitution $\tau' = t - \tau$ for the first integral, and $\tau' = \tau - t$ in the second integral. The legitimacy of the variable substitution follows from the fact that the integrand is piecewise continuous (thus Riemann integratable) (cf. Theorem 9.1). (Hence, the concerned integrals can be deemed as Riemann integrals.) Thus by the continuity of prob_G^{\max} (Theorem 9.1) and an application of L'Hospital's Rule to Definition 9.6, we obtain

$$\begin{aligned} & \nabla^+ \text{prob}_G^{\max}((s, a), \mathbf{c}) = \\ & \lambda \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c}) \right] - \lambda \cdot \text{prob}_G^{\max}((s, a), \mathbf{c}), \end{aligned}$$

which implies the result.

The proof for $\nabla^+ \text{prob}_G^{\max}((s, a), \mathbf{c}) = \nabla^- \text{prob}_G^{\max}((s, a), \mathbf{c})$ follows a similar argument. By definition and the previous derivation, for adequate $t > 0$, we have

$$\begin{aligned}
& \text{prob}_G^{\max}((s, a), \mathbf{c}) \\
&= \text{prob}_G^{\max}((s, a), (\mathbf{c} - t \cdot \mathbf{w}(s, a)) + t \cdot \mathbf{w}(s, a)) \\
&= e^{-\lambda \cdot t} \cdot \int_0^t \lambda \cdot e^{\lambda \cdot \tau} \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} - (t - \tau) \cdot \mathbf{w}(s, a)) \right] d\tau \\
&\quad + e^{-\lambda \cdot t} \cdot \text{prob}_G^{\max}((s, a), \mathbf{c} - t \cdot \mathbf{w}(s, a)) \\
&= \int_0^t \lambda \cdot e^{-\lambda \cdot \tau} \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} - \tau \cdot \mathbf{w}(s, a)) \right] d\tau \\
&\quad + e^{-\lambda \cdot t} \cdot \text{prob}_G^{\max}((s, a), \mathbf{c} - t \cdot \mathbf{w}(s, a)) .
\end{aligned}$$

where the last equality is obtained through the variable substitution $\tau' = t - \tau$. Thus by continuity and L'Hôpital's rule, we obtain

$$\begin{aligned}
& \nabla^- \text{prob}_G^{\max}((s, a), \mathbf{c}) = \\
& \lambda \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c}) \right] - \lambda \cdot \text{prob}_G^{\max}((s, a), \mathbf{c}) .
\end{aligned}$$

It follows that $\nabla^+ \text{prob}_G^{\max}((s, a), \mathbf{c}) = \nabla^- \text{prob}_G^{\max}((s, a), \mathbf{c})$.

Now we assume that $z = s \in S_{\text{la}} \setminus G$. Denote $\lambda := \mathbf{E}(s)$. Define $h[s, \mathbf{c}]$ to be the function

$$\tau \mapsto f_\lambda(\tau) \cdot \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} - \tau \cdot \mathbf{w}(s)) \right] .$$

We first consider $\nabla^+ \text{prob}_G^{\max}$. By Theorem 9.1, for all $t > 0$, we have

$$\begin{aligned}
& \text{prob}_G^{\max}(s, \mathbf{c} + t \cdot \mathbf{w}(s)) \\
&= \int_0^\infty h[s, \mathbf{c} + t \cdot \mathbf{w}(s)](\tau) \, d\tau \\
&= \int_0^t h[s, \mathbf{c} + t \cdot \mathbf{w}(s)](\tau) \, d\tau + \int_t^\infty h[s, \mathbf{c} + t \cdot \mathbf{w}(s)](\tau) \, d\tau \\
&= e^{-\lambda t} \cdot \int_0^t \lambda \cdot e^{\lambda \tau} \cdot \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} + \tau \cdot \mathbf{w}(s)) \right] \, d\tau \\
&\quad + e^{-\lambda t} \cdot \int_0^\infty h[s, \mathbf{c}](\tau) \, d\tau \\
&= e^{-\lambda t} \cdot \int_0^t \lambda \cdot e^{\lambda \tau} \cdot \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} + \tau \cdot \mathbf{w}(s)) \right] \, d\tau \\
&\quad + e^{-\lambda t} \cdot \text{prob}_G^{\max}(s, \mathbf{c})
\end{aligned}$$

where the third equality is obtained by the variable substitution $\tau' = t - \tau$ for the first integral, and $\tau' = \tau - t$ in the second integral. The legitimacy of the variable substitution follows from the fact that the integrand is piecewise continuous (cf. Theorem 9.1). Thus by applying L'Hospital's rule to Definition 9.6, we obtain that

$$\begin{aligned}
& \nabla^+ \text{prob}_G^{\max}(s, \mathbf{c}) = \\
& \max_{a \in \text{En}(s)} \lambda \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot (\text{prob}_G^{\max}(s', \mathbf{c}) - \text{prob}_G^{\max}(s, \mathbf{c})) \right],
\end{aligned}$$

which implies the result. The proof for $\nabla^- \text{prob}_G^{\max}(s, \mathbf{c})$ follows a similar argument. By Theorem 9.1 and previous derivations, for adequate $t > 0$,

$$\begin{aligned}
& \text{prob}_G^{\max}(s, \mathbf{c}) \\
&= \text{prob}_G^{\max}(s, (\mathbf{c} - t \cdot \mathbf{w}(s)) + t \cdot \mathbf{w}(s)) \\
&= e^{-\lambda t} \cdot \int_0^t \lambda \cdot e^{\lambda \tau} \cdot \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} - (t - \tau) \cdot \mathbf{w}(s)) \right] \, d\tau \\
&\quad + e^{-\lambda t} \cdot \text{prob}_G^{\max}(s, \mathbf{c} - t \cdot \mathbf{w}(s)) \\
&= \int_0^t \lambda \cdot e^{-\lambda \tau} \cdot \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c} - \tau \cdot \mathbf{w}(s)) \right] \, d\tau \\
&\quad + e^{-\lambda t} \cdot \text{prob}_G^{\max}(s, \mathbf{c} - t \cdot \mathbf{w}(s)).
\end{aligned}$$

where the last equality is obtained through the variable substitution $\tau' = t - \tau$. Thus by applying L'Hospital's Rule to Definition 9.6, we obtain

$$\begin{aligned} \nabla^- \text{prob}_G^{\max}(s, \mathbf{c}) = \\ \max_{a \in \text{En}(s)} \lambda \cdot \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot (\text{prob}_G^{\max}(s', \mathbf{c}) - \text{prob}_G^{\max}(s, \mathbf{c})) \right] \end{aligned}$$

which directly shows that $\nabla^+ \text{prob}_G^{\max}(s, \mathbf{c}) = \nabla^- \text{prob}_G^{\max}(s, \mathbf{c})$. \square

Remark 9.3. *The value $\mathbf{R}(s, a, s')$ can be viewed as the exit-rate of s via a to s' (cf. [57]).*

Theorem 9.3 gives a differential characterization for $\text{prob}_G^{\max}(z, \cdot)$ with $z \in Z_G \cup (S_{\text{la}} - G)$. Since $\nabla^+ \text{prob}_G^{\max}(z, \mathbf{c}) = \nabla^- \text{prob}_G^{\max}(z, \mathbf{c})$, we will solely use $\nabla \text{prob}_G^{\max}(z, \mathbf{c})$ to denote both of them.

Theorem 9.3 allows one to approximate $\text{prob}_G^{\max}(z, \mathbf{c} + t \cdot \mathbf{w}(z))$ through $\text{prob}_G^{\max}(z, \mathbf{c})$ and $\nabla \text{prob}_G^{\max}(z, \mathbf{c})$. This suggests an approximation algorithm which approximates $\text{prob}_G^{\max}(z, \mathbf{c})$ from $\{\text{prob}_G^{\max}(z, \mathbf{c}') \mid \mathbf{c}' \preceq \mathbf{c}\}$. An exception is the case when $\mathbf{w}(z) = \vec{0}$. Below we tackle this situation.

Proposition 9.1. *Let*

$$\begin{aligned} Y_G := \{z \in Z_G \cup (S_{\text{la}} \setminus G) \mid \mathbf{w}(z) = \vec{0}\} \\ \cup \{s \in S_{\text{er}} - G \mid (\exists a \in \text{En}(s)) \cdot \mathbf{w}(s, a) = \vec{0}\} . \end{aligned}$$

For all $\mathbf{c} \in \mathbb{R}_{\geq 0}^k$, the function $y \mapsto \text{prob}_G^{\max}(y, \mathbf{c})$ ($y \in Y_G$) is the least fixed-point (w.r.t \leq) of the higher order operator $\mathcal{Y}_{\mathbf{c}, G} : [Y_G \rightarrow [0, 1]] \rightarrow [Y_G \rightarrow [0, 1]]$ defined as follows: for each $(s, a) \in Y_G \cap Z_G$,

$$\begin{aligned} \mathcal{Y}_{\mathbf{c}, G}(h)(s, a) := \\ \sum_{s' \in Y_G \cap S} \mathbf{P}(s, a, s') \cdot h(s') + \sum_{s' \in S \setminus Y_G} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c}) ; \end{aligned}$$

for each $s \in Y_G \cap S_{\text{er}}$,

$$\begin{aligned} \mathcal{Y}_{\mathbf{c}, G}(h)(s) := \max\{\max\{h(s, a) \mid a \in \text{En}(s), (s, a) \in Y_G\}, \\ \max\{\text{prob}_G^{\max}((s, a), \mathbf{c}) \mid a \in \text{En}(s), (s, a) \notin Y_G\}\} ; \end{aligned}$$

for each $s \in Y_G \cap S_{\text{la}}$,

$$\begin{aligned} \mathcal{Y}_{\mathbf{c}, G}(h)(s) := \\ \max_{a \in \text{En}(s)} \left[\sum_{s' \in Y_G \cap S} \mathbf{P}(s, a, s') \cdot h(s') + \sum_{s' \in S \setminus Y_G} \mathbf{P}(s, a, s') \cdot \text{prob}_G^{\max}(s', \mathbf{c}) \right] . \end{aligned}$$

Proof. Let $\mathbf{c} \in \mathbb{R}_{\geq 0}^k$. By Theorem 9.1, $y \mapsto \text{prob}_G^{\max}(y, \mathbf{c})$ is a fixed-point of $\mathcal{Y}_{\mathbf{c}, G}$. Suppose that it is not the least fixed-point of $\mathcal{Y}_{\mathbf{c}, G}$. Let the least fixed-point be h^* . Define

- $\delta := \max\{\text{prob}_G^{\max}(y, \mathbf{c}) - h^*(y) \mid y \in Y_G\}$, and
- $Y' := \{y \in Y_G \mid \text{prob}_G^{\max}(y, \mathbf{c}) - h^*(y) = \delta\}$.

Since $y \mapsto \text{prob}_G^{\max}(y, \mathbf{c}) \neq h^*$, we have $\delta > 0$. Consider an arbitrary $y \in Y'$. By the maximal choice of δ and Y' , one can obtain that

1. for all $(s, a) \in Y' \cap Z_G$, $s' \in Y'$ whenever $s' \in S$ and $\mathbf{P}(s, a, s') > 0$;
2. for all $s \in Y' \cap S_{\text{er}}$, $(s, a) \in Y'$ whenever $a \in \text{En}(s)$ and $\text{prob}_G^{\max}(s, \mathbf{c}) = \text{prob}_G^{\max}((s, a), \mathbf{c})$.
3. for all $s \in Y' \cap S_{\text{la}}$, $s' \in Y'$ whenever there exists an $a \in \text{En}(s)$ such that $\mathbf{P}(s, a, s') > 0$ and $\text{prob}_G^{\max}(s, \mathbf{c}) = \sum_{s'' \in S} \mathbf{P}(s, a, s'') \cdot \text{prob}_G^{\max}(s'', \mathbf{c})$.

Intuitively, one can decrease every coordinate in Y' on prob_G^{\max} by a same amount so that a certain “balance” still holds. Choose $\delta' \in (0, \delta)$ such that

$$\text{prob}_G^{\max}(s, \mathbf{c}) - \delta' \geq \max\{\text{prob}_G^{\max}((s, a), \mathbf{c}) - \delta' \cdot \mathbf{1}_{Y'}(s, a) \mid a \in \text{En}(s)\}$$

for all $s \in Y' \cap S_{\text{er}}$, and

$$\text{prob}_G^{\max}(s, \mathbf{c}) - \delta' \geq \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{P}(s, a, s') \cdot (\text{prob}_G^{\max}(s', \mathbf{c}) - \delta' \cdot \mathbf{1}_{Y'}(s')) \right]$$

for all $s \in Y' \cap S_{\text{la}}$. Define $h : S \times \mathbb{R}^k \rightarrow [0, 1]$ by: $h(s, \mathbf{c}') := \text{prob}_G^{\max}(s, \mathbf{c}') - \delta'$ if $\mathbf{c}' = \mathbf{c}$ and $s \in Y'$, and $h(s, \mathbf{c}') := \text{prob}_G^{\max}(s, \mathbf{c}')$ otherwise. Then h is a pre-fixed-point of \mathcal{T}_G which satisfies that $h \not\preceq \text{prob}_G^{\max}$. Contradiction to Theorem 9.1. \square

9.4 Approximation Algorithm

In this section, we develop approximation algorithms to compute the maximal cost-bounded reachability probability under measurable schedulers. Our numerical algorithms will achieve the following task:

- **Input:** a CTMDP $(S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$, a cost function $\mathbf{w} : S \times \text{Act} \rightarrow \mathbb{R}_{\geq 0}^k$, a set $G \subseteq S$, a state $s \in S$, a vector $\mathbf{c} \in \mathbb{N}_0^k$ and an error bound $\epsilon > 0$;
- **Output:** a value $x \in [0, 1]$ such that $|\text{prob}_G^{\max}(s, \mathbf{c}) - x| \leq \epsilon$.

For computational purposes, we assume that each $\mathbf{w}_i(s, a)$ is an integer; rational numbers (and simultaneously the input cost-bound vector) can be adjusted to integers by multiplying a common multiplier of the denominators, without changing the maximal probability value to be approximated.

In the following, we fix a CTMDP $\mathcal{M} = (S, S_{\text{er}}, S_{\text{la}}, \text{Act}, \mathbf{E}_{\text{er}}, \mathbf{E}_{\text{la}}, \mathbf{P})$ and a cost function $\mathbf{w} : S \times \text{Act} \rightarrow \mathbb{R}_{\geq 0}^k$. And we fix a set $G \subseteq S$.

Based on Theorem 9.3 and Proposition 9.1, we design our approximation scheme as follows. First we introduce our discretization for a given $\mathbf{c} \in \mathbb{N}_0^k$ and a discretization step $\frac{1}{N}$ ($N \in \mathbb{N}$). Note that $\text{prob}_G^{\max}(s, \mathbf{c}) = 1$ whenever $s \in G$ and $\mathbf{c} \geq \vec{0}$. Thus we do not need to incorporate those points into discretization.

Definition 9.7. Let $\mathbf{c} \in \mathbb{N}_0^k$ and $N \in \mathbb{N}$. Define

$$\text{Disc}(\mathbf{c}, N) := \{\mathbf{d} \in \mathbb{R}^k \mid \vec{0} \leq \mathbf{d} \leq \mathbf{c} \text{ and } N \cdot \mathbf{d}_i \in \mathbb{N}_0 \text{ for all } 1 \leq i \leq k\} .$$

The set $\mathbf{D}_N^{\mathbf{c}}$ of discretized points is defined as follows:

$$\mathbf{D}_N^{\mathbf{c}} := ((S - G) \cup Z_G) \times \text{Disc}(\mathbf{c}, N) .$$

Thus $\mathbf{D}_N^{\mathbf{c}}$ is the set of “grid points” that are within the scope of \mathbf{c} and are discretized with discretization step $\frac{1}{N}$. The following definition presents the approximation scheme on $\mathbf{D}_N^{\mathbf{c}}$. Intuitively, the approximation scheme describes how those “points” are related.

Definition 9.8. Define $X_G := ((S - G) \cup Z_G) - Y_G$. The approximation scheme $\Upsilon_{\mathbf{c}, N}^G$ on $\mathbf{D}_N^{\mathbf{c}}$ consists of the following items:

- exactly one rounding argument for each element of $\mathbf{D}_N^{\mathbf{c}}$;
- a system of equations for elements in $X_G \times \text{Disc}(\mathbf{c}, N)$;
- a linear program on Y_G for each $\mathbf{d} \in \text{Disc}(\mathbf{c}, N)$.

Rounding Arguments: For each element $y \in \mathbf{D}_N^{\mathbf{c}}$, the rounding argument for y is as follows:

$$\overline{\text{prob}}_G(y) = \frac{K}{N^2} \text{ if } \text{prob}_G(y) \in \left[\frac{K}{N^2}, \frac{K+1}{N^2} \right) \text{ for some integer } 0 \leq K \leq N^2 .$$

Equations: The system of equations is described as follows. For all points $((s, a), \mathbf{d}) \in \mathbf{D}_N^{\mathbf{c}}$ with $(s, a) \in Z_G$, $\mathbf{w}(s, a) \neq \vec{0}$ and $\mathbf{d} - \frac{1}{N} \cdot \mathbf{w}(s, a) \geq \vec{0}$, there is a linear equation

$$\frac{\text{prob}_G((s, a), \mathbf{d}) - \overline{\text{prob}}_G((s, a), \text{pre}(\mathbf{d}, (s, a)))}{\frac{1}{N}} = \sum_{s' \in S} \mathbf{R}(s, a, s') \cdot (\overline{\text{prob}}_G(s', \text{pre}(\mathbf{d}, (s, a))) - \overline{\text{prob}}_G((s, a), \text{pre}(\mathbf{d}, (s, a)))) \quad (\text{E1})$$

where $\text{pre}(\mathbf{d}, z) := \mathbf{d} - \frac{1}{N} \cdot \mathbf{w}(z)$ for $z \in Z_G \cup (S_{\text{la}} \setminus G)$. For all $((s, a), \mathbf{d}) \in \mathbf{D}_N^{\mathbf{c}}$ with $(s, a) \in Z_G$, $\mathbf{w}(s, a) \neq \vec{0}$ and $\mathbf{d} - \frac{1}{N} \cdot \mathbf{w}(s, a) \not\geq \vec{0}$, there is a linear equation

$$\text{prob}_G((s, a), \mathbf{d}) = 0 . \quad (\text{E2})$$

For all $(s, \mathbf{d}) \in \mathcal{D}_N^c$ such that $s \in S_{\text{er}}$ and $\mathbf{w}(s, a) \neq \vec{0}$ for all $a \in \text{En}(s)$, there is an equation

$$\text{prob}_G(s, \mathbf{d}) = \max_{a \in \text{En}(s)} \text{prob}_G((s, a), \mathbf{d}) . \quad (\text{E3})$$

For all $(s, \mathbf{d}) \in \mathcal{D}_N^c$ with $s \in S_{\text{la}}$, $\mathbf{w}(s) \neq \vec{0}$ and $\mathbf{d} - \frac{1}{N} \cdot \mathbf{w}(s) \geq \vec{0}$, there is a linear equation

$$\begin{aligned} \frac{\text{prob}_G(s, \mathbf{d}) - \overline{\text{prob}}_G(s, \text{pre}(\mathbf{d}, s))}{\frac{1}{N}} = & \quad (\text{E4}) \\ \max_{a \in \text{En}(s)} \sum_{s' \in S} \mathbf{R}(s, a, s') \cdot (\overline{\text{prob}}_G(s', \text{pre}(\mathbf{d}, s)) - \overline{\text{prob}}_G(s, \text{pre}(\mathbf{d}, s))) & \end{aligned}$$

For all $(s, \mathbf{d}) \in \mathcal{D}_N^c$ with $s \in S_{\text{la}}$, $\mathbf{w}(s) \neq \vec{0}$ and $\mathbf{d} - \frac{1}{N} \cdot \mathbf{w}(s) \not\geq \vec{0}$, there is a linear equation

$$\text{prob}_G(s, \mathbf{d}) = 0 . \quad (\text{E5})$$

Linear Programs: For each $\mathbf{d} \in \text{Disc}(\mathbf{c}, N)$, the collection $\{\text{prob}_G(y, \mathbf{d})\}_{y \in Y_G}$ of values is the unique optimum solution of the linear program as follows:

- $\min \sum_{y \in Y_G} \text{prob}_G(y, \mathbf{d})$, subject to:
- $\text{prob}_G((s, a), \mathbf{d}) \geq \sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G(s', \mathbf{d})$ for all $(s, a) \in Y_G \cap Z_G$;
 - $\text{prob}_G((s, a), \mathbf{d}) \leq \text{prob}_G(s, \mathbf{d})$ for all $(s, a) \in Y_G \cap Z_G$ and $s \in Y_G \cap S_{\text{er}}$;
 - $\text{prob}_G(s, \mathbf{d}) \geq \sum_{s' \in S} \mathbf{P}(s, a, s') \cdot \text{prob}_G(s', \mathbf{d})$ for all $s \in Y_G \cap S_{\text{la}}$ and $a \in \text{En}(s)$;
 - $\text{prob}_G(y, \mathbf{d}) \in [0, 1]$ for all $y \in Y_G$;

where the values $\{\text{prob}_G(y, \mathbf{d})\}_{y \in X_G}$ are assumed to be known.

In all of the statements above, both $\text{prob}_G(s, \mathbf{d})$ and $\overline{\text{prob}}_G(s, \mathbf{d})$ represents 1 for $s \in G$.

Generally, $\text{prob}_G(y, \mathbf{d})$ approximates $\text{prob}_G^{\max}(y, \mathbf{d})$ and $\overline{\text{prob}}_G(y, \mathbf{d})$ approximates the same value with a rounding operation. A detailed computational sequence of the approximation scheme is described in Algorithm 1.

In principle, we compute the “higher” grid point $\text{prob}_G(z, \mathbf{d} + \frac{1}{N} \cdot \mathbf{w}(z))$ by $\text{prob}_G(z, \mathbf{d})$ and (E1) (or (E4)), and then update other “higher” points through (E3) and the linear program. The rounding argument is incorporated to avoid precision explosion caused by linear programming. The following proposition shows that Algorithm 1 indeed terminates after a finite number of steps.

Proposition 9.2. *Algorithm 1 terminates after a finite number of steps for all $\mathbf{c} \in \mathbb{N}_0^k$ and $N \in \mathbb{N}$.*

Algorithm 1 The Computation of $\Upsilon_{\mathbf{c},N}^G$

-
- 1: Set all relevant discretized points to zero by (E2) and (E5);
 - 2: Compute all $\text{prob}_G(s, \mathbf{d})$ that can be directly obtained through (E3);
 - 3: Compute all $\text{prob}_G(y, \mathbf{d})$ that can be directly obtained through the linear program;
 - 4: Compute all $\overline{\text{prob}}_G(y, \mathbf{d})$ that can be directly obtained by the rounding argument;
 - 5: Compute all relevant discretized points that can be directly obtained through (E1) and (E4), and back to Step 2. until all grid points in $D_N^{\mathbf{c}}$ are computed;
-

Proof. Let $\mathbf{c} \geq \vec{0}$ and $N \in \mathbb{N}$. The proposition can be proved through a straightforward induction on $\sum_{i=1}^k \mathbf{d}_i$ that both $\text{prob}_G(y, \mathbf{d})$ and $\overline{\text{prob}}_G(y, \mathbf{d})$ can be computed after a finite number of steps for all $(y, \mathbf{d}) \in D_N^{\mathbf{c}}$. The base step where $\text{prob}_G(y, \mathbf{d})$ is computed through (E2) or (E5) is easy. For the inductive step, suppose that for all $(y, \mathbf{d}') \in D_N^{\mathbf{c}}$ with $\sum_{i=1}^k \mathbf{d}'_i < \sum_{i=1}^k \mathbf{d}_i$, both $\text{prob}_G(y, \mathbf{d}')$ and $\overline{\text{prob}}_G(y, \mathbf{d}')$ can be computed after a finite number of steps by Algorithm 1. Then the inductive step can be sequentially justified by (E1) or (E4), (E3), the linear program and the rounding argument. \square

Below we prove that the approximation scheme really approximates prob_G^{\max} . To ease the notation, we shall use $\text{prob}_G(y, \mathbf{d})$ or $\overline{\text{prob}}_G(y, \mathbf{d})$ to denote both the variable at the grid point and the value it holds under the approximation scheme.

Theorem 9.4. *Let $\mathbf{c} \in \mathbb{N}_0^k$ and $N \in \mathbb{N}$ with $N \geq \mathbf{E}_{\max}$. For each $(y, \mathbf{d}) \in D_N^{\mathbf{c}}$,*

$$|\text{prob}_G(y, \mathbf{d}) - \text{prob}_G^{\max}(y, \mathbf{d})| \leq \left(\frac{2 \cdot \mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{N \cdot \mathbf{w}_{\min}} + \frac{1}{N} \right) \cdot \left[\sum_{i=1}^k \mathbf{d}_i \right] + \frac{\mathbf{E}_{\max}}{N}$$

and

$$\begin{aligned} |\overline{\text{prob}}_G(y, \mathbf{d}) - \text{prob}_G^{\max}(y, \mathbf{d})| &\leq \\ &\left(\frac{2 \cdot \mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{N \cdot \mathbf{w}_{\min}} + \frac{1}{N} \right) \cdot \left[\sum_{i=1}^k \mathbf{d}_i \right] + \frac{\mathbf{E}_{\max}}{N} + \frac{1}{N^2} . \end{aligned}$$

Proof. We proceed by induction on the number of computation steps illustrated by Algorithm 1.

Base Step: (y, \mathbf{d}) satisfies that $\mathbf{d} - \frac{1}{N} \cdot \mathbf{w}(y) \not\geq \vec{0}$. Then we know that $\text{prob}_G^{\max}(y, \mathbf{d} - x \cdot \mathbf{w}(y)) = 0$ where $x \in [0, \frac{1}{N}]$ is the largest real number such

that $\mathbf{d} - x \cdot \mathbf{w}(y) \geq \vec{0}$. Then by Theorem 9.3 and Lagrange's Mean-Value Theorem, we obtain that

$$\text{prob}_G^{\max}(y, \mathbf{d}) - \text{prob}_G^{\max}(y, \mathbf{d} - x \cdot \mathbf{w}(y)) = x \cdot \nabla \text{prob}_G^{\max}(y, \mathbf{d} - x' \cdot \mathbf{w}(y))$$

for some $x' \in (0, x)$. It follows that $\text{prob}_G^{\max}(y, \mathbf{d}) \leq \frac{1}{N} \cdot \mathbf{E}(y)$. Thus we have:

$$|\text{prob}_G^{\max}(y, \mathbf{d}) - \text{prob}_G(y, \mathbf{d})| \leq \frac{1}{N} \cdot \mathbf{E}_{\max} .$$

Inductive Step. The inductive step can be classified into the following cases:

Case 1: $(y, \mathbf{d}) = (s, \mathbf{d})$ with $s \in S_{\text{er}}$ and $\text{prob}_G(s, \mathbf{d})$ is computed through (E3). Then the result follows directly from the fact that

$$|\text{prob}_G^{\max}(s, \mathbf{d}) - \text{prob}_G(s, \mathbf{d})| \leq \max_{a \in \text{En}(s)} |\text{prob}_G^{\max}((s, a), \mathbf{d}) - \text{prob}_G((s, a), \mathbf{d})| .$$

Case 2: (y, \mathbf{d}) is computed through the linear program for \mathbf{d} . From Knaster-Tarski's Fixed-Point Theorem (Theorem 2.1), the linear program indeed computes the least fixed-point of $\mathcal{Y}_{\mathbf{d}, G}$. By induction hypothesis, for all $y' \in X_G$,

$$\begin{aligned} |\text{prob}_G(y', \mathbf{d}) - \text{prob}_G^{\max}(y', \mathbf{d})| \leq \\ \left(\frac{2 \cdot \mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{N \cdot \mathbf{w}_{\min}} + \frac{1}{N} \right) \cdot \left[\sum_{i=1}^k \mathbf{d}_i \right] + \frac{\mathbf{E}_{\max}}{N} . \end{aligned}$$

Thus by induction on n , one can prove that

$$\left| \mathcal{Y}_{\mathbf{d}, n}(\vec{0}) - \mathcal{Y}'_{\mathbf{d}, n}(\vec{0}) \right| \leq \left(\frac{2 \cdot \mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{N \cdot \mathbf{w}_{\min}} + \frac{1}{N} \right) \cdot \left[\sum_{i=1}^k \mathbf{d}_i \right] + \frac{\mathbf{E}_{\max}}{N} ,$$

where $\mathcal{Y}'_{\mathbf{d}}$ is the operator obtained by replacing $\{\text{prob}_G^{\max}(y, \mathbf{c})\}_{y \in X_G}$ in the definition of $\mathcal{Y}_{\mathbf{d}, G}$ (cf. Proposition 9.1) with $\{\text{prob}_G(y, \mathbf{c})\}_{y \in X_G}$, and $\mathcal{Y}_{\mathbf{d}, n}$ (resp. $\mathcal{Y}'_{\mathbf{d}, n}$) is the n -th Picard's iteration of $\mathcal{Y}_{\mathbf{d}, G}$ (resp. $\mathcal{Y}'_{\mathbf{d}}$). It follows that

$$|\text{prob}_G(y, \mathbf{d}) - \text{prob}_G^{\max}(y, \mathbf{d})| \leq \left(\frac{2 \cdot \mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{N \cdot \mathbf{w}_{\min}} + \frac{1}{N} \right) \cdot \left[\sum_{i=1}^k \mathbf{d}_i \right] + \frac{\mathbf{E}_{\max}}{N} .$$

Case 3: $(y, \mathbf{d}) = ((s, a), \mathbf{d})$ and $\text{prob}(y, \mathbf{d})$ is computed through (E1). By Lagrange's Mean-Value Theorem and Theorem 9.3, we have

$$\text{prob}_G^{\max}(y, \mathbf{d}) - \text{prob}_G^{\max}(y, \text{pre}(\mathbf{d}, y)) = \frac{1}{N} \cdot \nabla \text{prob}_G^{\max}(y, \mathbf{d} - x \cdot \mathbf{w}(y))$$

for some $x \in (0, \frac{1}{N})$. By Theorem 9.3 and Theorem 9.1, we can obtain that

$$\begin{aligned} \text{prob}_G^{\max}((s, a), \mathbf{d}) &= \text{prob}_G^{\max}((s, a), \text{pre}(\mathbf{d}, (s, a))) + \delta + \\ &\frac{1}{N} \cdot \sum_{s' \in S} \mathbf{R}(s, a, s') \cdot [\text{prob}_G^{\max}(s', \text{pre}(\mathbf{d}, y)) - \text{prob}_G^{\max}(y, \text{pre}(\mathbf{d}, y))] \end{aligned} \quad (*)$$

for some $\delta \in [-\frac{2}{N^2} \cdot \frac{\mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{\mathbf{w}_{\min}}, \frac{2}{N^2} \cdot \frac{\mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{\mathbf{w}_{\min}}]$. Rewriting (*) and (E1), we obtain that

$$\begin{aligned} \text{prob}_G^{\max}(y, \mathbf{d}) &= \frac{1}{N} \cdot \left[\sum_{s' \in S} \mathbf{R}(s, a, s') \cdot \text{prob}_G^{\max}(s', \text{pre}(\mathbf{d}, y)) \right] + \\ &\left(1 - \frac{\mathbf{E}(s, a)}{N} \right) \cdot \text{prob}_G^{\max}(y, \text{pre}(\mathbf{d}, y)) + \delta \end{aligned}$$

and

$$\begin{aligned} \text{prob}_G(y, \mathbf{d}) &= \frac{1}{N} \cdot \left[\sum_{s' \in S} \mathbf{R}(s, a, s') \cdot \overline{\text{prob}}_G(s', \text{pre}(\mathbf{d}, y)) \right] + \\ &\left(1 - \frac{\mathbf{E}(s, a)}{N} \right) \cdot \overline{\text{prob}}_G(y, \text{pre}(\mathbf{d}, y)) . \end{aligned}$$

By induction hypothesis, we have

$$\begin{aligned} |\text{prob}_G^{\max}(y, \mathbf{d}) - \text{prob}_G(y, \mathbf{d})| &\leq \\ &\left(\frac{2 \cdot \mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{N \cdot \mathbf{w}_{\min}} + \frac{1}{N} \right) \cdot \left[\sum_{i=1}^k (\text{pre}(\mathbf{d}, y))_i \right] + |\delta| + \frac{1}{N^2} + \frac{\mathbf{E}_{\max}}{N} , \end{aligned}$$

from which the induction step can be obtained.

Case 4: $(y, \mathbf{d}) = (s, \mathbf{d})$ and $\text{prob}_G(y, \mathbf{d})$ is computed through (E4). This case is very much similar to the previous case. By Lagrange's Mean-Value Theorem and Theorem 9.3, we have

$$\text{prob}_G^{\max}(s, \mathbf{d}) - \text{prob}_G^{\max}(s, \text{pre}(\mathbf{d}, s)) = \frac{1}{N} \cdot \nabla \text{prob}_G^{\max}(s, \mathbf{d} - x \cdot \mathbf{w}(s))$$

for some $x \in (0, \frac{1}{N})$. By Theorem 9.3 and Theorem 9.1, we can obtain that

$$\begin{aligned} \text{prob}_G^{\max}(s, \mathbf{d}) &= \text{prob}_G^{\max}(s, \text{pre}(\mathbf{d}, s)) + \delta + \\ &\frac{1}{N} \cdot \max_{a \in \text{En}(s)} \sum_{s' \in S} \mathbf{R}(s, a, s') \cdot [\text{prob}_G^{\max}(s', \text{pre}(\mathbf{d}, s)) - \text{prob}_G^{\max}(s, \text{pre}(\mathbf{d}, s))] \end{aligned} \quad (**)$$

for some $\delta \in [-\frac{2 \cdot \mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{N^2 \cdot \mathbf{w}_{\min}}, \frac{2 \cdot \mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{N^2 \cdot \mathbf{w}_{\min}}]$. Rewriting (**) and (E4), we obtain

$$\begin{aligned} \text{prob}_G^{\max}(s, \mathbf{d}) &= \frac{1}{N} \cdot \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{R}(s, a, s') \cdot \text{prob}_G^{\max}(s', \text{pre}(\mathbf{d}, s)) \right] + \\ &\left(1 - \frac{\mathbf{E}(s)}{N} \right) \cdot \text{prob}_G^{\max}(s, \text{pre}(\mathbf{d}, s)) + \delta \end{aligned}$$

and

$$\begin{aligned} \text{prob}_G(s, \mathbf{d}) &= \frac{1}{N} \cdot \max_{a \in \text{En}(s)} \left[\sum_{s' \in S} \mathbf{R}(s, a, s') \cdot \overline{\text{prob}}_G(s', \text{pre}(\mathbf{d}, s)) \right] + \\ &\quad \left(1 - \frac{\mathbf{E}(s)}{N} \right) \cdot \overline{\text{prob}}_G(s, \text{pre}(\mathbf{d}, s)) . \end{aligned}$$

By induction hypothesis, we have

$$\begin{aligned} |\text{prob}_G^{\max}(s, \mathbf{d}) - \text{prob}_G(s, \mathbf{d})| &\leq \\ &\left(\frac{2 \cdot \mathbf{E}_{\max}^2 \cdot \mathbf{w}_{\max}}{N \cdot \mathbf{w}_{\min}} + \frac{1}{N} \right) \cdot \left[\sum_{i=1}^k (\text{pre}(\mathbf{d}, s))_i \right] + \frac{\mathbf{E}_{\max}}{N} + |\delta| + \frac{1}{N^2} , \end{aligned}$$

from which the induction step can be obtained.

Case 5: $\overline{\text{prob}}(y, \mathbf{d})$ is computed through rounding. The induction step for this case is straightforward. \square

From Theorem 9.4, we derive our approximation algorithm as follows.

Corollary 9.2. *There exists an algorithm such that given any $\epsilon > 0$, $s \in S$, $G \subseteq S$ and $\mathbf{c} \in \mathbb{N}_0^k$, the algorithm outputs a $d \in [0, 1]$ which satisfies that $|d - \text{prob}_G^{\max}(s, \mathbf{c})| \leq \epsilon$. Moreover, the algorithm runs in*

$$\mathcal{O}\left(\left(\max\left\{\mathbf{E}_{\max}, \frac{M}{\epsilon}\right\}\right)^k \cdot \left(\prod_{i=1}^k \mathbf{c}_i\right) \cdot (|\mathcal{M}| + \log \frac{M}{\epsilon})^8\right)$$

time, where $M := \left(2 \cdot \mathbf{E}_{\max}^2 \cdot \frac{\mathbf{w}_{\max}}{\mathbf{w}_{\min}} + 1\right) \cdot \left[\sum_{i=1}^k \mathbf{c}_i\right] + \mathbf{E}_{\max} + 1$ and $|\mathcal{M}|$ is the size of \mathcal{M} .

Proof. The algorithm is a simple application of Theorem 9.4. If $s \in G$, the algorithm just returns 1; otherwise, the algorithm just calls Algorithm 1 with $N := \lfloor \max\{\mathbf{E}_{\max}, \frac{M}{\epsilon}\} \rfloor + 1$ and set $d := \text{prob}_G(s, \mathbf{c})$. By Theorem 9.4, we can directly obtain that $|d - \text{prob}_G^{\max}(s, \mathbf{c})| \leq M \cdot \frac{1}{N}$. For each $\mathbf{d} \in \text{Disc}(\mathbf{c}, N)$, the total computation of $\{\text{prob}_G(y, \mathbf{d})\}_{y \in X_G \cup Y_G}$ and $\{\overline{\text{prob}}_G(y, \mathbf{d})\}_{y \in X_G \cup Y_G}$ takes $\mathcal{O}\left(\left(|\mathcal{M}| + \log \frac{M}{\epsilon}\right)^8\right)$ time, since the most time consuming part is the linear program which takes $\mathcal{O}\left(\left(|\mathcal{M}| + \log N\right)^8\right)$ time (cf. [66]). Thus the total running time of the algorithm is

$$\mathcal{O}\left(\left(\max\left\{\mathbf{E}_{\max}, \frac{M}{\epsilon}\right\}\right)^k \cdot \left(\prod_{i=1}^k \mathbf{c}_i\right) \cdot (|\mathcal{M}| + \log \frac{M}{\epsilon})^8\right)$$

since the size of $\text{Disc}(\mathbf{c}, N)$ is $\mathcal{O}\left(N^k \cdot \left(\prod_{i=1}^k \mathbf{c}_i\right)\right)$. \square

9.5 Conclusion

In this chapter, we established an integral characterization for (multi dimensional) maximal cost-bounded reachability probability on continuous-time Markov decision processes, the existence of deterministic cost-positional optimal scheduler and an algorithm to approximate the cost-bounded reachability probability with an error bound. The approximation algorithm is based on a differential characterization of cost-bounded reachability probability, which in turn is derived from the integral characterization. The error bound is obtained through the differential characterization and a certain Lipschitz property. Moreover, the approximation algorithm runs in polynomial time in the size of the CTMDP and the reciprocal of the error bound, and exponential in the dimension of the cost-bound vector. Besides, we pointed out a proof error in the treatment of maximal time-bounded reachability probability on continuous-time Markov decision processes [57, 59]. We fixed this error in the more general setting of maximal cost-bounded reachability probability through a new methodology.

Chapter 10

Conclusion

This dissertation focuses on both algorithmic and complexity aspect of formal verification of probabilistic systems. The contributions of this dissertation are as follows.

In Chapter 5, we prove that the decision problem of simulation preorder between probabilistic pushdown automata and finite probabilistic automata is in EXPTIME. We demonstrate the EXPTIME-membership of the decision problem through a tableaux system obtained by a variation of the one by Colin Stirling [70, 69] and a partition-refinement technique. Combined with the EXPTIME-hardness result by Kučera and Mayr [52], we are able to show that the decision problem is EXPTIME-complete.

In Chapter 6, we prove that the threshold problem of the bisimilarity metric on finite probabilistic automata defined by van Breugel and Worrell [72] lies in $UP \cap coUP$, which is a subclass of $NP \cap coNP$. Our result significantly improves the previous complexity result by van Breugel *et al.* [71].

In Chapter 8, we correct errors in the paper [24] with new proofs and develop a numerical approximation algorithm for acceptance probability of CTMC-paths by a general (multi-clock) deterministic timed automata. Unlike the work by Barbot *et al.* [9], we are able to provide a tight error bound for our approximation algorithm.

In Chapter 9, we study maximal cost-bounded reachability probability on continuous-time Markov decision processes. In detail, we provide an integral characterization for the maximal cost-bounded reachability probability function, prove the existence of optimal cost-positional scheduler, and develop a numerical approximation algorithm for maximal cost-bounded reachability probability on CTMDPs. Meanwhile, we also fix a proof error in the work by Neuhäüßer and Zhang [59, 57]. The time complexity of the algorithm is polynomial in the size of the CTMDP, the unary representation of the cost-bound vector and the reciprocal of the given error bound, and exponential in the dimension of the cost-bound vector.

Bibliography

- [1] Rajeev Alur and David L. Dill. A theory of timed automata. *Theor. Comput. Sci.*, 126(2):183–235, 1994.
- [2] Christel Baier, Lucia Cloth, Boudewijn R. Haverkort, Matthias Kuntz, and Markus Siegle. Model checking Markov chains with actions and state labels. *IEEE Trans. Software Eng.*, 33(4):209–224, 2007.
- [3] Christel Baier, Ernst Moritz Hahn, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Model checking for performability. *Mathematical Structures in Computer Science*, 23(4):751–795, 2013.
- [4] Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Trans. Software Eng.*, 29(6):524–541, 2003.
- [5] Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Reachability in continuous-time Markov reward decision processes. In Jörg Flum, Erich Grädel, and Thomas Wilke, editors, *Logic and Automata*, volume 2 of *Texts in Logic and Games*, pages 53–72. Amsterdam University Press, 2008.
- [6] Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Performance evaluation and model checking join forces. *Commun. ACM*, 53(9):76–85, 2010.
- [7] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [8] Christel Baier, Joost-Pieter Katoen, Holger Hermanns, and Verena Wolf. Comparative branching-time semantics for Markov chains. *Inf. Comput.*, 200(2):149–214, 2005.
- [9] Benoît Barbot, Taolue Chen, Tingting Han, Joost-Pieter Katoen, and Alexandru Mereacre. Efficient CTMC model checking of linear real-time objectives. In Parosh Aziz Abdulla and K. Rustan M. Leino, editors, *TACAS*, volume 6605 of *Lecture Notes in Computer Science*, pages 128–142. Springer, 2011.

- [10] Gilles Bernot, Jean-Paul Comet, Adrien Richard, and Janine Guespin. Application of formal methods to biological regulatory networks: extending Thomas asynchronous logical approach with temporal logic. *Journal of Theoretical Biology*, 229(3):339–347, 2004.
- [11] Andrea Bianco and Luca de Alfaro. Model checking of probabilistic and nondeterministic systems. In P. S. Thiagarajan, editor, *FSTTCS*, volume 1026 of *Lecture Notes in Computer Science*, pages 499–513. Springer, 1995.
- [12] Patrick Billingsley. *Probability and Measure*. John Wiley & Sons, New York, NY, USA, 2nd edition, 1986.
- [13] Garrett Birkhoff. *Lattice Theory*, volume 25 of *Colloquium Publications*. American Math. Society, New York, USA, rev. ed. edition, 1948.
- [14] Gunter Bolch, Stefan Greiner, Hermann de Meer, and Kishor S. Trivedi. *Queueing Networks and Markov Chains - Modeling and Performance Evaluation with Computer Science Applications; 2nd Edition*. Wiley, 2006.
- [15] Tomáš Brázdil, Vojtech Forejt, Jan Krcál, Jan Kretínský, and Antonín Kučera. Continuous-time stochastic games with time-bounded reachability. *Inf. Comput.*, 224:46–70, 2013.
- [16] Tomáš Brázdil, Jan Krcál, Jan Kretínský, Antonín Kučera, and Vojtech Reháč. Stochastic real-time games with qualitative timed automata objectives. In Paul Gastin and François Laroussinie, editors, *CONCUR*, volume 6269 of *Lecture Notes in Computer Science*, pages 207–221. Springer, 2010.
- [17] Tomáš Brázdil, Jan Krcál, Jan Kretínský, Antonín Kučera, and Vojtech Reháč. Measuring performance of continuous-time stochastic processes using timed automata. In Marco Caccamo, Emilio Frazzoli, and Radu Grosu, editors, *HSCC*, pages 33–42. ACM, 2011.
- [18] Tomáš Brázdil, Antonín Kucera, and Oldrich Strazovský. On the decidability of temporal properties of probabilistic pushdown automata. In Volker Diekert and Bruno Durand, editors, *STACS*, volume 3404 of *Lecture Notes in Computer Science*, pages 145–157. Springer, 2005.
- [19] Tomáš Brázdil, Antonín Kučera, and Oldrich Strazovský. Deciding probabilistic bisimilarity over infinite-state probabilistic systems. *Acta Inf.*, 45(2):131–154, 2008.
- [20] Peter Buchholz and Ingo Schulz. Numerical analysis of continuous time Markov decision processes over finite horizons. *Computers & OR*, 38(3):651–659, 2011.

- [21] Supratik Chakraborty and Amit Kumar, editors. *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2011, December 12-14, 2011, Mumbai, India*, volume 13 of *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2011.
- [22] Krishnendu Chatterjee, Luca de Alfaro, Rupak Majumdar, and Vishwanath Raman. Algorithms for game metrics (full version). *Logical Methods in Computer Science*, 6(3), 2010.
- [23] Taolue Chen, Marco Diciolla, Marta Z. Kwiatkowska, and Alexandru Mereacre. Time-bounded verification of CTMCs against real-time specifications. In Uli Fahrenberg and Stavros Tripakis, editors, *FORMATS*, volume 6919 of *Lecture Notes in Computer Science*, pages 26–42. Springer, 2011.
- [24] Taolue Chen, Tingting Han, Joost-Pieter Katoen, and Alexandru Mereacre. Model checking of continuous-time Markov chains against timed automata specifications. *Logical Methods in Computer Science*, 7(1), 2011.
- [25] Taolue Chen, Tingting Han, and Jian Lu. On metrics for probabilistic systems: Definitions and algorithms. *Computers & Mathematics with Applications*, 57(6):991–999, 2009.
- [26] Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronization skeletons using branching-time temporal logic. In Dexter Kozen, editor, *Logic of Programs*, volume 131 of *Lecture Notes in Computer Science*, pages 52–71. Springer, 1981.
- [27] M.H.A. Davis. *Markov Models and Optimizations*. Chapman & Hall, New York, NY, USA, 1993.
- [28] Luca de Alfaro, Rupak Majumdar, Vishwanath Raman, and Mariëlle Stoelinga. Game refinement relations and metrics. *Logical Methods in Computer Science*, 4(3), 2008.
- [29] Yuxin Deng. *Lecture Notes on Probabilistic Concurrency*. Lecture Notes, available at <http://basics.sjtu.edu.cn/~yuxin/> ., 2009.
- [30] Josee Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Metrics for labelled Markov processes. *Theor. Comput. Sci.*, 318(3):323–354, 2004.
- [31] Susanna Donatelli, Serge Haddad, and Jeremy Sproston. Model checking timed and stochastic properties with CSL^{TA}. *IEEE Trans. Software Eng.*, 35(2):224–240, 2009.

- [32] R. M. Dudley. *Real Analysis and Probability*. Cambridge University Press, 2002.
- [33] Kousha Etessami and Mihalis Yannakakis. Recursive Markov decision processes and recursive stochastic games. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *ICALP*, volume 3580 of *Lecture Notes in Computer Science*, pages 891–903. Springer, 2005.
- [34] Kousha Etessami and Mihalis Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. *J. ACM*, 56(1), 2009.
- [35] John Fearnley, Markus Rabe, Sven Schewe, and Lijun Zhang. Efficient approximation of optimal control for continuous-time Markov games. In Chakraborty and Kumar [21], pages 399–410.
- [36] William Feller. *An Introduction to Probability Theory and Its Applications*. John Wiley & Sons, New York, NY, USA, 3rd edition, 1968.
- [37] Hongfei Fu. The complexity of deciding a behavioural pseudometric on probabilistic automata. Technical Report AIB-2011-26, RWTH Aachen, December 2011.
- [38] Hongfei Fu. Model checking EGF on basic parallel processes. In Tevfik Bultan and Pao-Ann Hsiung, editors, *ATVA*, volume 6996 of *Lecture Notes in Computer Science*, pages 120–134. Springer, 2011.
- [39] Hongfei Fu. Computing game metrics on Markov decision processes. In Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer, editors, *ICALP (2)*, volume 7392 of *Lecture Notes in Computer Science*, pages 227–238. Springer, 2012.
- [40] Hongfei Fu. Approximating acceptance probabilities of CTMC-paths on multi-clock deterministic timed automata. In Calin Belta and Franjo Ivancic, editors, *HSCC*, pages 323–332. ACM, 2013.
- [41] Hongfei Fu. Maximal cost-bounded reachability probability on continuous-time Markov decision processes. In Anca Muscholl, editor, *FoSSaCS*, volume 8412 of *Lecture Notes in Computer Science*, pages 73–87. Springer, 2014.
- [42] Hongfei Fu and Joost-Pieter Katoen. Deciding probabilistic simulation between probabilistic pushdown automata and finite-state systems. In Chakraborty and Kumar [21], pages 445–456.
- [43] Jan Friso Groote and Hans Hüttel. Undecidable equivalences for basic process algebra. *Inf. Comput.*, 115(2):354–371, 1994.

- [44] Robert L. Grossman, Anil Nerode, Anders P. Ravn, and Hans Rischel, editors. *Hybrid Systems*, volume 736 of *Lecture Notes in Computer Science*. Springer, 1993.
- [45] Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reliability. *Formal Asp. Comput.*, 6(5):512–535, 1994.
- [46] Hassan Hatefi and Holger Hermanns. Improving time bounded reachability computations in interactive Markov chains. In Farhad Arbab and Marjan Sirjani, editors, *FSEN*, volume 8161 of *Lecture Notes in Computer Science*, pages 250–266. Springer, 2013.
- [47] Chaodong He, Yuxi Fu, and Hongfei Fu. Decidability of behavioral equivalences in process calculi with name scoping. In Farhad Arbab and Marjan Sirjani, editors, *FSEN*, volume 7141 of *Lecture Notes in Computer Science*, pages 284–298. Springer, 2011.
- [48] Petr Jancar, Antonín Kucera, and Richard Mayr. Deciding bisimulation-like equivalences with finite-state processes. *Theor. Comput. Sci.*, 258(1-2):409–433, 2001.
- [49] Bengt Jonsson and Kim Guldstrand Larsen. Specification and refinement of probabilistic processes. In *LICS*, pages 266–277. IEEE Computer Society, 1991.
- [50] Dexter Kozen. Results on the propositional mu-calculus. *Theor. Comput. Sci.*, 27:333–354, 1983.
- [51] Antonín Kučera, Javier Esparza, and Richard Mayr. Model checking probabilistic pushdown automata. *Logical Methods in Computer Science*, 2(1), 2006.
- [52] Antonín Kučera and Richard Mayr. On the complexity of checking semantic equivalences between pushdown processes and finite-state processes. *Inf. Comput.*, 208(7):772–796, 2010.
- [53] Kim Guldstrand Larsen and Arne Skou. Bisimulation through probabilistic testing. *Inf. Comput.*, 94(1):1–28, 1991.
- [54] Edward A. Lee and Sanjit A. Seshia. *Introduction to Embedded Systems, A Cyber-Physical Systems Approach*. <http://LeeSeshia.org>, 2011.
- [55] Linar Mikeev, Martin R. Neuhäüßer, David Spieler, and Verena Wolf. On-the-fly verification and optimization of DTA-properties for large Markov chains. *Formal Methods in System Design*, 43(2):313–337, 2013.
- [56] Robin Milner. *Communication and Concurrency*. PHI Series in computer science. Prentice Hall, 1989.

- [57] Martin R. Neuhäüßer. *Model Checking Nondeterministic and Randomly Timed Systems*. PhD thesis, RWTH Aachen, 2010.
- [58] Martin R. Neuhäüßer, Mariëlle Stoelinga, and Joost-Pieter Katoen. Delayed nondeterminism in continuous-time Markov decision processes. In Luca de Alfaro, editor, *FOSSACS*, volume 5504 of *Lecture Notes in Computer Science*, pages 364–379. Springer, 2009.
- [59] Martin R. Neuhäüßer and Lijun Zhang. Time-bounded reachability probabilities in continuous-time Markov decision processes. In *QEST*, pages 209–218. IEEE Computer Society, 2010.
- [60] Prakash Panangaden. *Labelled Markov Processes*. Imperial College Press, 2009.
- [61] David Park. Concurrency and automata on infinite sequences. In *Proceedings of the 5th GI-Conference on Theoretical Computer Science*, pages 167–183, London, UK, 1981. Springer-Verlag.
- [62] Amir Pnueli. The temporal logic of programs. In *FOCS*, pages 46–57. IEEE Computer Society, 1977.
- [63] Tomás Prieto-Rumeau and Onésimo Hernández-Lerma. *Selected Topics on Continuous-Time Controlled Markov Chains and Markov Games*. Imperial College Press, London, UK, 2012.
- [64] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1994.
- [65] Markus N. Rabe and Sven Schewe. Finite optimal control for time-bounded reachability in CTMDPs and continuous-time Markov games. *Acta Inf.*, 48(5-6):291–315, 2011.
- [66] Alexander Schrijver. *Theory of Linear and Integer Programming*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 1999.
- [67] Roberto Segala and Nancy A. Lynch. Probabilistic simulations for probabilistic processes. *Nord. J. Comput.*, 2(2):250–273, 1995.
- [68] William J. Stewart. *Probability, Markov Chains, Queues, and Simulation - The Mathematical Basis of Performance Modeling*. Princeton University Press, 2009.
- [69] Colin Stirling. Decidability of bisimulation equivalence for normed pushdown processes. *Theor. Comput. Sci.*, 195(2):113–131, 1998.

- [70] Colin Stirling. Decidability of bisimulation equivalence for pushdown processes. *Unpublished manuscript, available at <http://homepages.inf.ed.ac.uk/cps/>*, 2000.
- [71] Franck van Breugel, Babita Sharma, and James Worrell. Approximating a behavioural pseudometric without discount for probabilistic systems. *Logical Methods in Computer Science*, 4(2), 2008.
- [72] Franck van Breugel and James Worrell. A behavioural pseudometric for probabilistic transition systems. *Theor. Comput. Sci.*, 331(1):115–142, 2005.
- [73] Rob J. van Glabbeek. The linear time-branching time spectrum (extended abstract). In Jos C. M. Baeten and Jan Willem Klop, editors, *CONCUR*, volume 458 of *Lecture Notes in Computer Science*, pages 278–297. Springer, 1990.
- [74] Nicolás Wolovick and Sven Johr. A characterization of meaningful schedulers for continuous-time Markov decision processes. In Eugene Asarin and Patricia Bouyer, editors, *FORMATS*, volume 4202 of *Lecture Notes in Computer Science*, pages 352–367. Springer, 2006.
- [75] Lijun Zhang, David N. Jansen, Flemming Nielson, and Holger Hermanns. Automata-based CSL model checking. *Logical Methods in Computer Science*, 8(2), 2011.

Aachener Informatik-Berichte

This list contains all technical reports published during the past three years. A complete list of reports dating back to 1987 is available from:

<http://aib.informatik.rwth-aachen.de/>

To obtain copies please consult the above URL or send your request to:

Informatik-Bibliothek, RWTH Aachen, Ahornstr. 55, 52056 Aachen,
Email: biblio@informatik.rwth-aachen.de

- 2011-01 * Fachgruppe Informatik: Jahresbericht 2011
- 2011-02 Marc Brockschmidt, Carsten Otto, Jürgen Giesl: Modular Termination Proofs of Recursive Java Bytecode Programs by Term Rewriting
- 2011-03 Lars Noschinski, Fabian Emmes, Jürgen Giesl: A Dependency Pair Framework for Innermost Complexity Analysis of Term Rewrite Systems
- 2011-04 Christina Jansen, Jonathan Heinen, Joost-Pieter Katoen, Thomas Noll: A Local Greibach Normal Form for Hyperedge Replacement Grammars
- 2011-06 Johannes Lotz, Klaus Leppkes, and Uwe Naumann: dco/c++ - Derivative Code by Overloading in C++
- 2011-07 Shahar Maoz, Jan Oliver Ringert, Bernhard Rumpe: An Operational Semantics for Activity Diagrams using SMV
- 2011-08 Thomas Ströder, Fabian Emmes, Peter Schneider-Kamp, Jürgen Giesl, Carsten Fuhs: A Linear Operational Semantics for Termination and Complexity Analysis of ISO Prolog
- 2011-09 Markus Beckers, Johannes Lotz, Viktor Mosenkis, Uwe Naumann (Editors): Fifth SIAM Workshop on Combinatorial Scientific Computing
- 2011-10 Markus Beckers, Viktor Mosenkis, Michael Maier, Uwe Naumann: Adjoint Subgradient Calculation for McCormick Relaxations
- 2011-11 Nils Jansen, Erika brahám, Jens Katelaan, Ralf Wimmer, Joost-Pieter Katoen, Bernd Becker: Hierarchical Counterexamples for Discrete-Time Markov Chains
- 2011-12 Ingo Felscher, Wolfgang Thomas: On Compositional Failure Detection in Structured Transition Systems
- 2011-13 Michael Förster, Uwe Naumann, Jean Utke: Toward Adjoint OpenMP
- 2011-14 Daniel Neider, Roman Rabinovich, Martin Zimmermann: Solving Muller Games via Safety Games
- 2011-16 Niloofar Safiran, Uwe Naumann: Toward Adjoint OpenFOAM
- 2011-17 Carsten Fuhs: SAT Encodings: From Constraint-Based Termination Analysis to Circuit Synthesis
- 2011-18 Kamal Barakat: Introducing Timers to pi-Calculus
- 2011-19 Marc Brockschmidt, Thomas Ströder, Carsten Otto, Jürgen Giesl: Automated Detection of Non-Termination and NullPointerExceptions for Java Bytecode

- 2011-24 Callum Corbett, Uwe Naumann, Alexander Mitsos: Demonstration of a Branch-and-Bound Algorithm for Global Optimization using McCormick Relaxations
- 2011-25 Callum Corbett, Michael Maier, Markus Beckers, Uwe Naumann, Amin Ghoheity, Alexander Mitsos: Compiler-Generated Subgradient Code for McCormick Relaxations
- 2011-26 Hongfei Fu: The Complexity of Deciding a Behavioural Pseudometric on Probabilistic Automata
- 2012-01 Fachgruppe Informatik: Annual Report 2012
- 2012-02 Thomas Heer: Controlling Development Processes
- 2012-03 Arne Haber, Jan Oliver Ringert, Bernhard Rumpe: MontiArc - Architectural Modeling of Interactive Distributed and Cyber-Physical Systems
- 2012-04 Marcus Gelderie: Strategy Machines and their Complexity
- 2012-05 Thomas Ströder, Fabian Emmes, Jürgen Giesl, Peter Schneider-Kamp, and Carsten Fuhs: Automated Complexity Analysis for Prolog by Term Rewriting
- 2012-06 Marc Brockschmidt, Richard Musiol, Carsten Otto, Jürgen Giesl: Automated Termination Proofs for Java Programs with Cyclic Data
- 2012-07 André Egners, Björn Marschollek, and Ulrike Meyer: Hackers in Your Pocket: A Survey of Smartphone Security Across Platforms
- 2012-08 Hongfei Fu: Computing Game Metrics on Markov Decision Processes
- 2012-09 Dennis Guck, Tingting Han, Joost-Pieter Katoen, and Martin R. Neuhäuser: Quantitative Timed Analysis of Interactive Markov Chains
- 2012-10 Uwe Naumann and Johannes Lotz: Algorithmic Differentiation of Numerical Methods: Tangent-Linear and Adjoint Direct Solvers for Systems of Linear Equations
- 2012-12 Jürgen Giesl, Thomas Ströder, Peter Schneider-Kamp, Fabian Emmes, and Carsten Fuhs: Symbolic Evaluation Graphs and Term Rewriting — A General Methodology for Analyzing Logic Programs
- 2012-15 Uwe Naumann, Johannes Lotz, Klaus Leppkes, and Markus Towara: Algorithmic Differentiation of Numerical Methods: Tangent-Linear and Adjoint Solvers for Systems of Nonlinear Equations
- 2012-16 Georg Neugebauer and Ulrike Meyer: SMC-MuSe: A Framework for Secure Multi-Party Computation on MultiSets
- 2012-17 Viet Yen Nguyen: Trustworthy Spacecraft Design Using Formal Methods
- 2013-01 * Fachgruppe Informatik: Annual Report 2013
- 2013-02 Michael Reke: Modellbasierte Entwicklung automobiler Steuerungssysteme in Klein- und mittelständischen Unternehmen
- 2013-03 Markus Towara and Uwe Naumann: A Discrete Adjoint Model for OpenFOAM
- 2013-04 Max Sagebaum, Nicolas R. Gauger, Uwe Naumann, Johannes Lotz, and Klaus Leppkes: Algorithmic Differentiation of a Complex C++ Code with Underlying Libraries

- 2013-05 Andreas Rausch and Marc Sihling: Software & Systems Engineering Essentials 2013
- 2013-06 Marc Brockschmidt, Byron Cook, and Carsten Fuhs: Better termination proving through cooperation
- 2013-07 André Stollenwerk: Ein modellbasiertes Sicherheitskonzept für die extrakorporale Lungenunterstützung
- 2013-08 Sebastian Junges, Ulrich Loup, Florian Corzilius and Erika brahám: On Gröbner Bases in the Context of Satisfiability-Modulo-Theories Solving over the Real Numbers
- 2013-10 Joost-Pieter Katoen, Thomas Noll, Thomas Santen, Dirk Seifert, and Hao Wu: Performance Analysis of Computing Servers using Stochastic Petri Nets and Markov Automata
- 2013-12 Marc Brockschmidt, Fabian Emmes, Stephan Falke, Carsten Fuhs, and Jürgen Giesl: Alternating Runtime and Size Complexity Analysis of Integer Programs
- 2013-13 Michael Eggert, Roger Häußling, Martin Henze, Lars Hermerschmidt, René Hummen, Daniel Kerpen, Antonio Navarro Pérez, Bernhard Rumpe, Dirk Thißen, and Klaus Wehrle: SensorCloud: Towards the Interdisciplinary Development of a Trustworthy Platform for Globally Interconnected Sensors and Actuators
- 2013-14 Jörg Brauer: Automatic Abstraction for Bit-Vectors using Decision Procedures
- 2013-19 Florian Schmidt, David Orlea, and Klaus Wehrle: Support for error tolerance in the Real-Time Transport Protocol
- 2013-20 Jacob Palczynski: Time-Continuous Behaviour Comparison Based on Abstract Models
- 2014-01 * Fachgruppe Informatik: Annual Report 2014
- 2014-02 Daniel Merschen: Integration und Analyse von Artefakten in der modellbasierten Entwicklung eingebetteter Software
- 2014-03 Uwe Naumann, Klaus Leppkes, and Johannes Lotz: dco/c++ User Guide
- 2014-04 Namit Chaturvedi: Languages of Infinite Traces and Deterministic Asynchronous Automata
- 2014-05 Thomas Ströder, Jürgen Giesl, Marc Brockschmidt, Florian Frohn, Carsten Fuhs, Jera Hensel, and Peter Schneider-Kamp: Automated Termination Analysis for Programs with Pointer Arithmetic
- 2014-06 Esther Horbert, Germán Martín García, Simone Frintrop, and Bastian Leibe: Sequence Level Salient Object Proposals for Generic Object Detection in Video
- 2014-07 Niloofar Safran, Johannes Lotz, and Uwe Naumann: Algorithmic Differentiation of Numerical Methods: Second-Order Tangent and Adjoint Solvers for Systems of Parametrized Nonlinear Equations
- 2014-08 Christina Jansen, Florian Göbe, and Thomas Noll: Generating Inductive Predicates for Symbolic Execution of Pointer-Manipulating Programs

- 2014-09 Thomas Ströder and Terrance Swift (Editors): Proceedings of the International Joint Workshop on Implementation of Constraint and Logic Programming Systems and Logic-based Methods in Programming Environments 2014

* These reports are only available as a printed version.

Please contact biblio@informatik.rwth-aachen.de to obtain copies.