

Bisimilar stochastic systems

Tkachev, Ilya

DOI

[10.4233/uuid:d3d6a8ba-45ff-406d-a81a-087e8bc8dbf0](https://doi.org/10.4233/uuid:d3d6a8ba-45ff-406d-a81a-087e8bc8dbf0)

Publication date

2019

Document Version

Final published version

Citation (APA)

Tkachev, I. (2019). *Bisimilar stochastic systems*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:d3d6a8ba-45ff-406d-a81a-087e8bc8dbf0>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Bisimilar Stochastic Systems

Ilya Tkachev

BISIMILAR STOCHASTIC SYSTEMS

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus Prof. dr. ir. T.H.J.J. van der
Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op
donderdag 27 juni 2019 om 12:00 uur

door

Ilya TKACHEV

Master of Science in Financial Mathematics,
Halmstad University, Sweden
geboren te Volgograd, Russia

Dit proefschrift is goedgekeurd door de promotoren:

Prof. dr. ir. B. De Schutter

Prof. A. Abate

Dr. P. Mohajerin Esfahani

Samenstelling promotiecommissie bestaat uit:

Rector Magnificus,	voorzitter
Prof. dr. ir. B. De Schutter,	Technische Universiteit Delft, promotor
Prof. dr. A. Abate,	University of Oxford, promotor
Dr. P. Mohajerin Esfahani,	Technische Universiteit Delft, copromotor

onafhankelijke leden:

This dissertation has been completed in fulfillment of the requirements of the Dutch Institute of Systems and Control (DISC) for graduate study.

The work presented in this thesis has been supported by the project *Modeling, Verification and Control of Complex Systems (MoVeS)*, funded by the European Union seventh framework FP7/2007-2013, under grant agreement n°257005.

ISBN: 978-94-6384-048-4

Copyright © 2019 by Ilya Tkachev.

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission from the copyright owner.

Printed by CPI-Wöhrmann Print Service, The Netherlands.

*To my father, for he is the one who deserves
this the most*

Acknowledgments

Though unexpected and ever unintended, the timeline of this thesis happened to be quite long. It took around a thousand days of intensive research activities to come up with the content for it, and another four years of (alas) not intensive enough trials to put all this content together. For this reason, here I mention my gratitude both to those who supported or influenced me during the university days, and those who was around after – without either of them this work would not have been possible to finish.

I would like to start with thanking signor Alessandro Abate, my PhD supervisor, for supporting my wish to try myself in research, and giving me so much freedom in doing what I deemed important. I greatly appreciate the amount of time and attention he dedicated to our joint work, especially since every month – unlikely less often – his brain was attacked with new ideas that came from my roaming mind. His attentiveness spanned far beyond the research, and even besides his great support in the department affairs, I do warmly recall a few rare, but hence even more impactful, talks we had regarding life, relationships and people during our travels. I am certain that with a different supervisor my university time likely would have been less exciting.

One of the first things that Alessandro has shared with me during one of our first talks is that he wanted to create an environment where his students learn how to be independent researchers. It is unlikely this task could have been carried in a better way. Given that with time direction of my research required more and more expertise in the fields we have not been familiar with, it is a big luck that I have discovered MathStackexchange and MathOverflow on early stages of my PhD track, as its users have answered many questions I had, and gave some of interesting ideas to dig into. I have enjoyed their community and support, and would like to specifically mention such people as Andre Nicolas, Asaf Karagila, Theo Buehler, Nate Eldredge, George Lowther, J.M., my kernel buddy Michael Greinecker and my TU Delft colleague Jonas Teuwen. I am also particularly grateful to Rob Johnson not only for his comments and chat jokes, but also for showing me L.A. for my first time. With the development of my thesis I could no longer get enough serious input on the forums, and I appreciate a lot the enthusiasm of another TU Delft professor Frank Redig in co-authoring a paper with me. This casted me to be braver in contacting other senior scientist around the world, and I was surprised how easy it was to get replies from them. In particular, I would like to mention

invaluable comments regarding gambling models (that made it to be the framework for this thesis) by Bill Sudderth, and even a closer touch regarding Markov Decision Processes with Eugene Feinberg, who was kind enough to support my visit to NYSE in Stony Brook and write a paper together.

Before my thesis departed from control theory to more abstract areas of mathematics, and since this transition never happened in full, I got a lot of insights from the systems and control community. I am definitely indebted to bright young and older minds, with some of whom I even had a pleasure of collaborative projects. Joost-Pieter Katoen for disseminating his expertise in compute science and helping to write an award-winning paper together with his student Alexanru Mereacre. Maria Prandini for hosting me in Milano and shedding light on alternative approaches to stochastic bisimulation. John Lygeros for helping me to learn a better balance between what is practical and what is just elegant. In fact, all the people from the MoVeS project, as it was a very exciting joint effort from which I have gained a lot of experience, and not only in science. Manuel Mazo Jr. for his energy, honesty and non-stochastic control theory discussions. Paulo Tabuada for sharing his ideas, interest in my research, a very nice interview and a delicious dinner in one of the best (in my humble opinion) Portugese restaurants, and Matthias Runnger for his respectfully formal and precise attitude to math, and for great days in Souther California. Ufuk Topcu for his critical view, and another great interview, and another great dinner. Claire Tomlin for inviting me to for an extended visit towards the end of my university track, hence giving me the opportunity to fall in love with the vibes of UC Berkeley, the Bay Area and California overall. Max Balandat for my finest student collaboration and for bringing me closer to game theory – perhaps, the greatest field in math, and Roel Dobbe whose enthusiasm in pursuing his own PhD and company in my free time at Berkeley enlightened the trip even more. Marco Diciolla for his fire and focus on the goal, and valuable ideas on what to do after PhD. Antoine Girard for sharing the first-hand memories on how the approximate bisimulation methods started developing in control theory, and Agung Julius for elucidating their first applications to stochastic systems. Aaron Ames for his strong opinions, interesting thoughts and personal shows of how should bipedal robots walk. Calin Belta and all his group for making a very nice company during almost every conference I have been to, always having something compelling to chat about, and among them Ebru Aydin Gol for her in-depth analysis and Yushan Chen for showing me what a startupper life in Silicon Valley looks like. Gregory Batt for introducing me to the field of system biology and making it possible to visit Paris for so many times. Marielle Stoëlinga for talks on the differences between stochasticity and non-determinism. The whole department I have studied at, as people in Delft Center for Systems and Control (DCSC) always seemed to knew what is the best balance between the work and the fun. In particular, my MSc studens Lakshmeesh, Yifan and Pepijn for their dedication, asking me so many diverse questions, and leaving my brain no other option but to stretch further and further. Gabriel Lopez and Ivo Grondman for their insights regarding robotics and great soccer games. The team I have worked in together with Sadegh, Dieky, Katerina and Majid for their feedback. My special DCSC thanks would definitely go to Pawel Stano and Marco Forgione, for they were not only my colleagues, but are also friends I am lucky to have, even

though we only rarely have a chance to meet in person.

I am grateful to Optiver for after leaving TU Delft I happened to be in an environment where people were at least as smart as in academia, and in terms of energy, honesty and dedication likely are setting even a higher benchmark. This experience allowed me to look from a bit different angle at my results, and the approach to research. I am happy I had a chance to work with Tejaswi Navilarekallu, an extremely talented mathematician, unexpectedly persistent in his efforts to make me enjoy brain teasers, which even though deemed futile by me in the beginning succeeded in the end. I am thankful to Prasad Chebolu for his subtle and elegant sense of humor and being a great research partner, showing depths in feeling the science I have rarely encountered in others. Finally, Matthijs Snel was a colleague and is a great friend who supported me in some of my hardest times, and also showed be a an example of someone finishing thesis so long after finishing the university, and was kind enough to provide his help with translating the summary of this manuscript into Dutch. After going to industry and putting this thesis on hold, I have been always reminded by some people that it is worth finishing it. Some of them already had a PhD degree and were praising my abilities so high, and so sincerely, I felt ashamed I have not done it yet. Among them I want to mention my father, a professor of math whose level of pure and undissolved interest in this subject likely I will never be able to reach, even though it is perhaps mostly due to his genes I managed to come up with the contents of what follows below. A parent, whose wise advices I could have been listening more carefully to. Another person in that group would be Dmitry Shcherbakov, who was extremely persistent in finishing his own thesis, no matter how hard the obstacles were. A friend who is an example of optimism and energy, who never hesitated to tell me what is right or wrong in his opinion, but only after first supporting me in the moments of need. And of course my deep gratitude goes to Evgeniya, who even though have not started pursuing the degree, leaves me no doubt that she would have accomplished it easily if she ever wished for it. She was close wherever I went, and supported my views regarding this world however unconventional they might have seemed. It was her example that helped me to learn how to set priorities, find the goals, fight for them no matter what, and reach them.

I am also grateful to Mikhail Nechaev, who supervised my Master Thesis, and by virtue of that helped me rediscover how much do I love doing research. Likely it is for this reason that I decided to postpone trying to enter the gates of the financial industry by another four years, and just focused entirely on mathematics, given an infinitesimal amount of thoughts to how practical is my research.

Finally, I am forever indebted to DCSC for their trust in me being able to finish this thesis. For the support given by Hans Hellendoorn and Bart de Schutter on the one hand, and an in-depth effort from Peyman Mohajerin Esfahani to help me finish this very manuscript.

I was and am happy that all those people took part in my life, even though apparently none of them had taught me to write shorter acknowledgements. For though unexpected and ever unintended, writing this part took long enough time and big enough effort to shed some light why it was not such an easy task to finish this thesis.

Ilya Tkachev
Moscow, 20 May 2019

Contents

Acknowledgments	vii
1 Introduction	1
1.1 Motivation	1
1.2 Literature review	4
1.3 Research goals and original contributions	6
1.4 Outline of the thesis	9
2 Models and problems	11
2.1 Transition systems	11
2.2 Problem specifications	15
2.3 Stochastic systems	20
2.4 SS behaviors	26
2.5 Comments on models and problem formulation	30
3 Finite-horizon case	33
3.1 Approximate simulation of TSs	33
3.2 Precise simulation of SSs	35
3.3 Approximate simulation of SSs	44
3.4 Theory for MDPs	52
3.5 Comments on approximate stochastic bisimulation	58
4 Infinite-horizon case	65
4.1 Motivation	65
4.2 Reachability problem	65
4.2.1 Characterization	65
4.2.2 Computation	74
4.2.3 Comments on the reachability problem	81

4.3	Repeated reachability	83
4.3.1	Characterization	83
4.3.2	Computation	84
4.3.3	Characterization through harmonic functions	86
4.3.4	Comments on the repeated reachability problem	97
5	Conclusions and future research	99
5.1	Conclusions	99
5.2	Future research	101
A	General concepts	103
A.1	Notation and conventions	103
A.2	Pseudometrics	108
A.3	Approximate relations	109
A.4	Modelling terminology	112
A.5	Weak inclusion of families of probability measures	117
A.6	Strong and weak inclusions	119
B	Categorical concepts	121
B.1	Categories	121
B.2	Category of Borel relations	124
C	Other topics	129
C.1	Important fragments of LTL	129
C.2	Stochastic equations	131
C.3	Auxiliary results	131
	Bibliography	139
	List of Symbols and Notation	151
	List of Abbreviations	153
	Summary	155
	Samenvatting	157
	Curriculum Vitae	159

Introduction

This work concerns optimal control of complex stochastic systems by virtue of approximating them with simpler ones. We start with outlining the problem, its relevance and the previous results on this topic. It follows by a summary of the authors' contributions, and is concluded by a scheme of the thesis.

1.1 Motivation

The focus of this thesis is on decision making in presence of uncertainty. A common approach is to come up with a mathematical model for the system under consideration, e.g. a robot, biological cell, plant, and a performance criterion to be optimized over this system. Such situations are present in a most of practical applications, and depending on the kind of uncertainty, different models are used.

For example, whenever the uncertainty can be assumed to have only marginal effect, the system can be described using deterministic dynamics. In continuous time the classical framework for that is provided by ordinary differential equations (ODEs) [109], whereas transition systems (TSs) are often used in discrete-time, e.g. event-driven, systems.

If uncertainty cannot be easily neglected, but is known to be bounded, the classical control methods of control of deterministic systems can be extended to robust control techniques [150], which look into performance under worst case scenario . As in deterministic case, TSs are used to model the discrete-time case, whereas an example of a continuous-time model is given by differential inclusions (DIs) [14]. Such framework is successfully used e.g. in safety-critical applications such as air traffic control.

Finally, when uncertainty in the system can be statistically quantified, a natural way to model it is by means of probability theory [50]. There are at least two

reasons to use this approach over robust control: first, often uncertainty cannot be easily bounded and second, in many cases probabilistic bounds are sufficient, whereas worst case scenario analysis may provide very conservative results. Continuous-time systems are modeled with controlled diffusions [86], and Markov Decision Processes (MDPs) [111] are used for the discrete case. This approach is used in cases where random behavior of a system is crucial, and one can allow summing up several events to get to the average results, e.g. in gambling, trading, insurance etc.

In terms defined above, this thesis focuses on discrete-time probabilistic case. As mentioned, MDPs is a common way to model such systems, and it will be the benchmark model in this work. There are several reasons for that: a lot of results were already developed for MDPs, and its semantics is arguably natural to model the class of systems we are interested in. There are other models available in the literature – similar either in syntax or semantics – and some of the results relevant to our study are only available in those frameworks. Below we introduce these models in more details whenever we need it, and restate those results for MDPs.

So far we focus on MDPs, so let us look at their semantics. At some step the system resides in state x belonging to the state space X . The agent can pick any action u from a set $U(x)$ of actions enabled in x , and the new state x' is randomly generated according to the transition probability, which depends both on previous step and the choice of action: $x' \sim p(x, u)$. Such choice can be done based on the full history of previous states and actions. A set of these choices is referred to as the strategy of the agent: a sequential decision rule that chooses actions based on available historical information. The agent may prefer one action to another according to the performance criterion he seeks to optimize. For example, an additive criterion (AC) rewards the agent on any transition with some value $r(x, u, x')$, and agent focuses on maximizing the expected cumulative reward: this classical problem was treated in many studies [20, 70], and solution relies on dynamic programming (DP) introduced in [18]. As an alternative, an agent may be interested in maximizing probability of reaching some set of goal states $G \subseteq X$, which was studied in a setting of gambling [48, 95], and more recently revisited by compute science community [145]. There it was also proposed to use modal logics such as LTL to characterize more complex performance criteria, for example maximizing probability of reaching a goal state G via a checkpoint set A , while avoiding an unsafe set B . Using the performance criterion to be optimized, the agent can evaluate a given strategy. The strategy with the best value is called optimal; to solve MDP problem is to find such strategy and/or its value.

When solving an MDP, one rarely expects to find an analytical solution. In case of continuous state space such solutions are available e.g. for linear Gaussian models with additive performance criterion, where optimal value and policy can be found by solving a system of matrix equations [77]. A large body of research was devoted to optimal control of finite MDPs: those with finite state space X , and finite set of actions for agent to choose from. For this models most of studied performance criteria allow for computable algorithmic solutions, which enabled development of several dedicated software packages such as PRISM [73] and MRMC [80]. In particular, extending results from non-probabilistic finite models, these al-

gorithms employ reducing control for general LTL specifications to basic problem such as probabilistic reachability by enlarging the state space and using automata expressions of LTL formulae [17].

Due to the fact that for continuous spaces MDPs only allow for precise solutions in a limited class of models and performance criteria, one resorts to numerical methods to find approximate solutions. Ideally such methods provide bounds on the approximation error and devise a way to refine an approximation scheme to achieve any a priori given accuracy. These methods can be split among the following three paradigms, which we list in order of increasing generality of problems they can tackle:

- The first in the list is a classical approach comprising direct solution of a particular problem. For example, [6] provided approximation scheme for a special case of MDP without control structure, also known as Markov Chain (MC), with performance criterion given by finite-horizon probabilistic invariance. Namely, one seeks to compute a probability that a state of a system stays in a safe set A for a given finite number of transitions n . Similarly, [57] developed an approximation scheme for controlled MDPs and additive cost criterion.
- An issue with the direct approach is that it is hard to extend it beyond some fundamental but basic problems such as reachability, invariance, or additive cost, for which there are well-known DP procedures. To cope with this issue, one could use results from automata theory mentioned above to reduce complex specification, such as ones given by LTL formulae, to basic problems. Authors of [7] successfully applied this approach to bounded time horizon properties over MCs.
- A drawback of a second approach is the need to modify the underlying model of a system, which in most cases enlarges its state space. Since approximate errors are often sensitive to the size of the latter, it is preferable to find methods whose accuracy does not depend of the performance criterion, or at least those that provide guarantees that hold uniformly for a large class of criteria. Such goal was achieved for non-probabilistic systems with the use of precise and approximate bisimulations [128], which was missing for probabilistic systems at the moment of the start of the title project.

This work provides novel results that were not yet achieved when following the paradigms above. Starting from the most general one, we suggest criterion-agnostic methods (approximate bisimulations) for controlled MDPs. Assumptions used there cannot guarantee bounded errors for infinite-horizon performance criteria, so to cover this case we resort to the second solution paradigm, extending results of [7] to controlled MDPs. Finally, we provide ways to solve infinite-horizon probabilistic reachability problem, which in particular is needed to complete the previous step. Our contribution is detailed below in Section 1.3, but before that we provide a literature review of the methods that were available before.

1.2 Literature review

As many other numerical methods, approximation methods for general state-space MDPs were studied already several decades ago, i.e. relatively early given that the DP procedure itself was developed in mid 50s [18]. In most cases such methods only provided rate of convergence, rather than formal bounds on the error, see e.g. [88] for the summary of some first results, and [66] for an example of some recent developments.

Move to more precise guarantees came from the following two directions: approximation work for finite MDPs and similar problems for general non-stochastic TSs. With focus on the former, [90] seems to be the first work that tried extending (precise) bisimulation methods from TSs to an MDP-like model. The latter was given the name Probabilistic Transition System (PTS), and syntactically was very close to a discrete MDP, however its semantics was not as rich. Those ideas were further developed in [40] for Labelled Markov Processes (LMP), again a model with a syntax very close to PTS and MDP, but without a clearly defined semantics – even though it was suggested the one may be inherited from PTSs. In contrast, [114] studied more conventional and rich semantics of MDPs under the name of Probabilistic Automata (PA), essentially MDPs with general state space, but only discrete transition distributions. Focus again was on precise match between behaviors of the systems. It is interesting though, that it was in fact the area of probabilistic systems where the idea of approximate bisimulation first popped up, likely due to genuine quantitateness in them, i.e. each property could have been satisfied only with a certain probability. That casted results to be continuous rather than boolean as in TSs, hence motivating looking for solutions that can obtain those results with a given precision. Perhaps, the first work to introduce this idea was [42], which replaced the standard Boolean logic used to characterize precise bisimulation with a functions family, for which extreme levels sets were exactly the states satisfying formulae of that logic. That was pretty much the state where the development of the methods relevant to this thesis was put on hold, e.g. [34] and [29] were refining the theory of earlier results and looking at them from different angles. It is worth mentioning that a similar line of work started by [37], picked up in [139] and extended in [143] tried to characterize bisimulations in PTSs from the point of view of category theory, employing ideas from [92] and [65] on probabilistic functors.

In the meantime, a relevant research was going for general TSs with focus on optimal control policy synthesis. For example, [107] focuses on precise relations between such systems, but already in [64] ideas on probabilistic approximations receive development in the domain of TSs, where they – for the first time – are applied to a natural linear-time semantics, thus allowing to solve many interesting problems inspired by practical applications. This approach found an immediate response back in the probabilistic community. For example, even though [4] did not try introducing a new conceptual framework like the works above, it has provided formal error bounds and a useful abstraction technique for a rather general class of MDPs, even though with finite actions space only. Those errors were derived with focus on an important task of optimal reachability. It is worth paying

attention to the latter fact: in contrast to early works on probabilistic approximations that focused on the methods, perhaps hoping that they will find a later use (with perhaps an exception of [57] that was also problem-oriented), the latter tried to find a concrete solution to a concrete problem that was already known to be relevant. Inspired by those results, [5] tried formalizing the developed approximation approach, however results were quite conservative and only applied to marginal distributions of uncontrolled MDPs. As a next step, [6] merged those ideas and applied them to derive approximation error bounds for probabilistic safety, again over uncontrolled MDPs. In parallel, there was also an attempt to extend a theory of approximate bisimulations [64] MDPs: [76] focused on continuous time case showed that under the assumption of the existence of a Lyapunov-like approximate bisimulation function (ABF) that bounds divergence between two MDPs, one can use a probability computed over the first (simpler) MDP to bound that over the second (more complex) one. The paper provides sufficient condition for existence and explicit examples of such ABFs for linear uncontrolled diffusions: multi-dimensional version of geometric Brownian motion [106]. It was followed by [1] that further related ABFs to stochastic stability [59], and [8] which relaxed the bounds to allow for Monte-Carlo randomized methods, hence only providing guarantees that would hold with a certain probability for continuous-time uncontrolled MDPs. Similar relaxation was applied in [79] to classical DP problems over general discrete-time MDPs. Coming back to the formal guarantees, [119] laid down the way to refined approximation techniques by using adaptive gridding of the state space, and since most of the works on precise bounds for MDPs mentioned in this paragraph were focused on a particular problem – stochastic reachability – [7] extended these results to a bigger class of properties by means of the automata theory [17].

A more elaborate survey of that progress one can find in [2], and more details are provided at the end of each section below. For now it is important for us to summarize what was left to do in this field, which will serve as a reasoning for the list of the results of this thesis that follow in the next section. The first group of work (e.g. [42]) has derived interesting bounds in a unified framework of bisimulation for PTSs, but it was not clear whether those bounds can be used for any popular optimal control problems. Another group (a benchmark would be [7]) derived bounds for a problem whose relevance is leaving no doubts, but the class of properties covered coincides with bounded linear temporal logic (BLTL) [17], and the derived bounds were growing linearly both in time horizon of the problem as well as in the size of the automata needed to express this property. The latter part is especially important, as for bounded properties the automata tend to have large sizes, hence one possibly gets bounds with higher polynomial degree (rather than just linear) of dependence on the time horizon. For this reason, those works lacked both the solutions to the infinite time-horizon case, and the unified framework which could have allowed for bounds that are independent on the automata size. The third group (starting from [76]) tried to tackle those problems and due to the use of stability assumptions, the bounds provided in their paper hold over infinite time horizon for a large number of properties, however this assumption does rarely hold in practice, and in the end the problem is only reduced to a smaller dimension. Hence, even in case of success, one still has to solve the

problem over a continuous-time MDP. This open questions in previous research paved the way to the goals of this thesis.

1.3 Research goals and original contributions

The previous section describes the state of art in the field on the moment when this thesis was started, hence the author was facing the following two challenges.

1. In case no stability assumptions are made and the focus is on the bounded time horizon, is it necessary for bounds to depend on the shape of the property? Namely, is it possible for the error bound to only depend on the time horizon of the property and yet is less conservative than the one in [7]?
2. What are reasonable assumptions that one needs to provide good bounds that hold on the infinite time horizon, at least with focus on a particular problem of stochastic reachability?

Those were the main two topics the research of the author evolved over during the work on this thesis. In the beginning just the second topic (infinite-horizon approximations) was chosen as the core one, however on the way some results were derived for the former topic as well. What follows is the list of the author's original contributions, separated in those two groups.

Infinite-horizon stochastic reachability over MDPs.

- [129] seeks to refine the results of [76] for the case of uncontrolled discrete-time MDP. It establishes relations between the concepts of the latter paper and classical stochastic theory methods, such as optimal stopping problem and minimal dominating supermartingales. For example, it shows that reachability probability function is the best (in a certain sense) Lyapunov function, and a ABF is its conservative approximate. This paper is the first to notice the importance of absorbing sets for the solution of the infinite-horizon reachability. It introduces a concept of the largest absorbing subset (LAS) of a set, together with the explicit procedure to construct it. The latter uses finite-horizon stochastic reachability value functions, and this interdependence leads to the equivalence between empty LAS and trivial solution of infinite horizon reachability, hence only leaving non-empty LAS case to be tackled. For the latter situation, the paper first provides a solution to infinite-horizon stochastic reach-avoid problem, and further uses it to decompose infinite-horizon reachability into two simpler problems. Finally, it does find a suitable Lyapunov-like condition needed to wrap up the solution.
- [130] Builds upon the results of the previous paper and extends its results to a bigger class of infinite-horizon problem, including the stochastic shortest path/mean exit time problem. It relates Bellman equations to be solved to find value functions to Dirichlet problems over partial differential equations

[63]: in both cases one deals with a fixpoint of a linear operator that satisfies some boundary conditions, just in case of discrete-time MDPs those operators are bounded. As a result, the decomposition technique of [129] can be considered as a regularization of a Dirichlet problem, and its derivation can be obtained from the monotonicity of the solutions of the latter problem with respect to the boundary conditions. This provides an analytic derivation of the method, that complements the probabilistic approach used in [129].

As it was established in the previous paper, the main challenge in finding value function for infinite-horizon stochastic properties was reduced to finding the LAS of a given set. Unfortunately, this problem is undecidable much like the halting problem: if LAS is empty then the procedure will stop at some point, but until that moment one cannot say whether LAS is empty or not. For this reason, the current paper has provided an example of a class of uncontrolled MDPs, over which the LAS problem is decidable.

- [134] Is the journal version of the previous two papers. It enriches and formalizes the ideas of those papers into complete results and provides comprehensive proofs. In addition, it applies the developed infinite-horizon theory to solve the problems that appear as specifications of probabilistic computational tree logic (PCTL).
- [131] extends [129] in a different direction, by going beyond unbounded horizon stochastic reachability to a genuinely infinite-horizon property: stochastic repeated reachability. It establishes asymptotic infinite-horizon behavior of uncontrolled MDPs as LASs are natural candidates for stochastic attractors. It further introduces a natural form of stochastic stability with respect to a given set, and a relevant notion of stochastically attractive sets, and shows that an absorbing set is stochastically attractive if and only if it admits for a Lyapunov-like function that first appeared in the decomposition technique [129], thus providing necessary and sufficient conditions for the solution of infinite-horizon properties over uncontrolled MDPs, and fully characterizing connection between stochastic stability and Lyapunov functions, similarly to celebrated results for non-stochastic systems [118].
- [137] achieves several goals at once.

First, it expresses stochastic reachability as a classical additive cost performance criterion (PC) over a controlled MDP. This fact allows using the rich theory of [20] to fully characterize solution of the controlled stochastic reachability problem, greatly simplifying and enhancing previous results of [9] and [124] on that topic.

Second, it provides bounds for bounded-horizon controlled stochastic reachability. This bounds still have linear growth, but relax the assumption on the finiteness of the action space.

Third, it generalizes the concepts of the LAS to a controlled MDP, and provides a solution to the controlled stochastic reach-avoid problem in case of empty LAS, thus extending the result of [129] to a more general class of systems.

Finally, it uses ideas of [7] to motivate importance of stochastic reachability as a key problem for verifying more complex properties, exploiting the automata theory to cover all properties expressible as formulae of linear temporal logic (LTL).

[132] Focuses on the infinite-horizon part of the challenge tackled by the previous paper, and provides a solution for the infinite stochastic reach-avoid problem in case of non-empty LAS. It also shows the application of this theory to the problem of ruin probability dealt with in the field of actuarial mathematics [13].

- [138] is the journal version of the previous two papers. Besides providing comprehensive proofs, it even further completes characterization of the stochastic reachability. Moreover, its main focus is set on the infinite-horizon properties, extending now the results of [134] to obtain equivalence conditions between the triviality of stochastic controlled reachability solution and emptiness of relevant versions of LASs. This is further used to obtain develop a decomposition technique and a Lyapunov-like method to solve infinite-horizon reachability by means of stochastic reach-avoid problem solved in a previous paper.

While deriving most of the results on the unbounded-horizon problems, it was presumed that one has a suitable procedure to solve their bounded-horizon counterparts. This assumption was realistic as [4] already provided a numerical scheme which yielded approximate solution with any given precision, and these results were further refined by [119] and the followup research. These bounds work well for the bounded-horizon reachability problem, however when applied to more complex problems by means of the automata method of [7], the bound grows with the size of the automaton used to express the problem, and may become of a limited usability. Hence, the second branch of the author's research was focused on studying whether this is a feature of the problem, or just the bounds are not tight enough. Given the idea of approaching verification of a class of properties, rather than a single property only, this research also happened to be a quest for a natural definition of approximate bisimulation for stochastic systems.

- [133] shows that the approximation error does only depend on the time horizon of a specification, and as a result any bound derived for bounded-horizon reachability can be applied to any BLTL formula. Moreover, the generality of the result makes it useful for quantitative model-checking, providing bounds also for verification of properties with bounded reward.
- [10] the bounds in the previous paper are general and less conservative than those available before, but they still have linear growth with respect to the time horizon of the property. The current paper showed that these bounds can be improved to be of bounded growth in time, slowly approaching the value of 1 (the maximal possible difference between two probabilities), as the time horizon of the property goes to infinity. This paper also shows that these bounds are tight, and there are MDPs with probability differences

that follow those bounds exactly, so that they can only be improved in some specific cases, but not in general.

- [135]. The previous two papers provided all the essential results needed for approximate model-checking, but it is still some fragments of the theory rather than a solid framework. The current paper makes the first attempt on formalizing previous results to fit them in a proper framework, similar to the one of approximate bisimulations for non-stochastic systems. In particular, it puts forward a definition of behaviors of MDPs, and suggests which properties should a natural behavioral metric on the space of MDPs satisfy. As an example, it provides two such metrics: the total variation one use in the previous two papers and pretty much all previous research on numerical analysis of MDPs, and the Wasserstein one – a less conservative generalization of metric induced by ABFs of [76]. Finally, this paper provides Monte-Carlo based bounds on those metrics, in contrast to more precise but computationally more expensive bounds of [133, 10].
- This thesis. The previous paper introduces some important ideas, but is yet far away from delivering a unified framework of approximate bisimulations of MDPs that would fully and elegantly exploit the right bounds of [10]. Also it was never shown explicitly how these bounds can be applied to the case of controlled MDPs. To the best of the author’s knowledge, this thesis provides such methods together with the whole framework for the first time. It studies approximate stochastic bisimulations also from the categorical perspective, similar to attempts performed by [37].

1.4 Outline of the thesis

As mentioned in the previous section, this thesis contains original material on how to put forward the framework of approximate stochastic bisimulations for MDPs. We start with defining the scope of problems we deal with and concepts we apply in Chapter 2. The bisimulation framework is then developed in Chapter 3. As the bounds provided by the latter cannot be extended to unbounded-horizon properties in general, Chapter 4 applies a group of more focused method to solve that case. As some results appear to be rather technical, or just branching out of the main content, in order not to break the flow the are put in Appendices A, B and C.

Models and problems

This chapter introduces problems and models we deal with throughout the thesis. It starts from a simpler, and likely more familiar, setting of non-stochastic transition systems to introduce the main concepts, and then applies those ideas for the stochastic case.

2.1 Transition systems

To motivate and elucidate concepts and methods presented below for stochastic systems (SSs), we plan to use the rich theory of transition systems (TSs) as a benchmark. For this reason, we start with a short recapitulation of the latter model, its syntax and semantics. Our exposition of TSs closely follows [128], with minor modifications made to better connect with the theory of SSs below.

Definition 2.1 *A transition system (TS) is a tuple $\mathcal{T} = (X, T, Y, L)$ where X and Y are arbitrary sets, $T \subseteq X \times X$ is an l.t.r. and $L : X \rightarrow Y$. We say that X is the state space, T is the transition relation, Y is the output space and L is the output map of \mathcal{T} . The TS \mathcal{T} is said to be finite if the set T is finite, otherwise the TS \mathcal{T} is called infinite. The set of all TSs with the output space Y is denoted by TS_Y .*

Unless the contrary is specified, further in this section we always assume that the TS $\mathcal{T} = (X, T, Y, L)$ is given and fixed. Definition 2.1 provides the syntax of the TS model only, whereas to define its semantics we need to describe the dynamics of its models. If the current state of \mathcal{T} is x_n , we observe the output value of $y_n = L(x_n)$ and then the new state is chosen among admissible successor states $x_{n+1} \in T|_{x_n}$. To formalize this procedure we use the notion of a (control) strategy. It does not often appear explicitly in the literature on TSs, but it is crucial below for SSs, and to better emphasize similarity between these two models we introduce strategies for TSs.

Definition 2.2 A strategy for the TS \mathcal{T} is a sequential decision rule $\sigma = (\sigma_n)_{n \in \mathbb{N}}$, where the map $\sigma_n : X^n \rightarrow X$ is such that $\sigma_n(x_0, \dots, x_n) \in T|_{x_n}$ for each $x_i \in X$, $i \in [0; n]$ and $n \in \mathbb{N}$. The set of all such strategies we denote by Σ^T .

Notice that only the transition structure of the \mathcal{T} given by a pair (X, T) appears in the definition of strategies, whereas the output structure (Y, L) has no effect. In fact, T uniquely determines the transition structure as it is an l.t.r. and hence X can be defined as a left projection of T . In particular, T alone uniquely defines the set of strategies: due to this reason we use the notation Σ^T instead of more cumbersome $\Sigma^{X,T}$.

Based on the definition of strategies for TSs we can now formalize their semantics. Each combination of an initial state $x \in X$ and strategy $\sigma \in \Sigma^T$ generate a unique *internal run* of \mathcal{T} denoted by $V_x^\sigma \in X^\mathbb{N}$ and defined as follows: $v = V_x^\sigma$ iff $v_0 = x$ and $v_{n+1} = \sigma_n(v_0, \dots, v_n)$. Sets of all internal runs and those that start at $x \in X$ are denoted by $V(\Sigma^T)$ and $V(\Sigma^T, x)$ respectively, where clearly $V(\Sigma^T) := \bigcup_{x \in X} V(\Sigma^T, x)$. Note that given an internal run $v \in V(\Sigma^T)$ we can uniquely deduce both the initial state $x \in X$ and the strategy $\sigma \in \Sigma^T$ that generate v . A non-stochastic nature of TS makes the proof of this fact trivial, however interestingly a similar result also holds for the SSs under some technical conditions as we see below.

Since only outputs are observed externally, to each internal run $v \in V(\Sigma^T)$ we assign a corresponding *external run* $w = L(v) \in Y^\mathbb{N}$. We further denote by $V_L(\Sigma^T) := L(V(\Sigma^T))$ and $V_L(\Sigma^T, x) := L(V(\Sigma^T, x))$ sets of all external runs, and those initiated at x , respectively. We refer to $V_L(\Sigma^T)$ as the *behavior*¹ of the TS \mathcal{T} .

Given the set of outputs Y , by a linear temporal (LT) property² we mean an arbitrary subset $H \subseteq Y^\mathbb{N}$. For example, if $S \subseteq Y$ is a set of “safe” outputs and $G \subseteq Y$ is that of “goal outputs”, then $H = S^\mathbb{N}$ corresponds to the safety property, whereas $H = Y^\mathbb{N} \setminus (G^c)^\mathbb{N}$ corresponds to the reachability property. A lot of other interesting and more intricate properties can be expressed e.g. by means of the automata theory [17, Chapter 4], or modal logics [17, Chapter 5] using expressions similar to those of natural languages.

We say that \mathcal{T} satisfies LT H if there exists a behavior $w \in V_L(\Sigma^T)$ such that $w \in H$: in such case we write $\mathcal{T} \models H$, and we write $\mathcal{T} \not\models H$ otherwise. As a result, we have defined the satisfaction and refutation relations $\models, \not\models \subseteq \text{TS}_Y \times Y^\mathbb{N}$. Clearly,

$$\mathcal{T} \models H \iff V_L(\Sigma^T) \cap H \neq \emptyset. \quad (2.1)$$

Note that due to (2.1) the behavior of \mathcal{T} uniquely determines which LT properties does \mathcal{T} satisfy, regardless of what is the transition structure of \mathcal{T} . Namely, if two TSs induce the same behaviors, they satisfy exactly the same LT properties, even if their transition structures happen to be different. This fact is crucial for us later, when we want to describe behaviors of a complex system using a much simpler one.

¹ The terminology here may differ, for example in [64] such set is called the language of \mathcal{T} . Our use of the term “behavior” is inspired by [128].

² For details on LT properties for TSs see e.g. [17, Chapter 3] and in particular [17, Definition 3.11].

Note also that although we have defined the satisfaction relation \models using an existential quantifier, the ability to verify whether $\mathcal{T} \models H$ for any given H gives us a way to check whether all the behavior of \mathcal{T} belong to H as well. Indeed, the answer to such question is affirmative if $\mathcal{T} \models H^c$ and negative if $\mathcal{T} \not\models H^c$. The main reason \models is defined above using the existential quantifier is that we are not only focused on the verification problem (whether $\mathcal{T} \models H$), but also on the synthesis problem: given that $\mathcal{T} \models H$, how to find x and σ such that $V_x^\sigma \in H$. Computer science community developed efficient automatic algorithms and dedicated software (model checkers) for both problems over finite TSs: see [17] for the details and references.

Despite the fact that the definition of TS is rather simple, it can be used e.g. to describe continuous, control or hybrid systems – see [128] for the details. However, due to the fact that such systems are genuinely infinite, the resulting TS is unavoidably infinite as well. Due to the reason that aforementioned verification and synthesis algorithms were developed for the finite TSs, [64] proposed to approximate an infinite system of interest with a finite one whose behavior is somewhat related to that of the original infinite system. The new system is often referred to as an *abstraction*, as it often provides a model for the real-world phenomenon described by the original system on a higher level of abstraction. The original system is referred to as the *concrete* one.

With focus on LT problems exclusively, *behavioral inclusion* plays a predominant role. Given a new TS $\bar{\mathcal{T}} = (\bar{X}, \bar{T}, Y, \bar{L})$ over the same output space as \mathcal{T} , we say that $\bar{\mathcal{T}}$ behaviorally includes \mathcal{T} if $V_L(\Sigma^{\mathcal{T}}) \subseteq V_{\bar{L}}(\Sigma^{\bar{\mathcal{T}}})$; in such case we write $\mathcal{T} \leq \bar{\mathcal{T}}$. Why would that be interesting to study behavior inclusions between TSs? Suppose that $\bar{\mathcal{T}}$ is a finite system which is an abstraction of an infinite system \mathcal{T} . If $\mathcal{T} \leq \bar{\mathcal{T}}$ and $\bar{\mathcal{T}} \not\models H$ for some LT property H , then we can immediately conclude that $\mathcal{T} \not\models H$. Similarly, if $\mathcal{T} \models H'$ then by definition there exist \bar{x} and $\bar{\sigma}$ such that $\bar{V}_{\bar{x}}^{\bar{\sigma}} \in H'$. Thus, behavioral inclusion of \mathcal{T} in $\bar{\mathcal{T}}$ helps solving verification (synthesis) problems on \mathcal{T} ($\bar{\mathcal{T}}$) exploiting the results obtained for $\bar{\mathcal{T}}$ (\mathcal{T}). Due to this reason, if we are interested in a finite abstraction useful both for verification and synthesis, we shall talk of *behavioral equivalence*. We say that \mathcal{T} and $\bar{\mathcal{T}}$ are behaviorally equivalent if $\mathcal{T} \leq \bar{\mathcal{T}}$ and $\bar{\mathcal{T}} \leq \mathcal{T}$; in such case we write $\mathcal{T} \approx \bar{\mathcal{T}}$.

Remark 2.3 *The fact that behavioral equivalence allows solving synthesis problems shall emphasize the nature of the version of the TS model considered in this thesis. Often the TS is defined as a quintuple (X, U, T, Y, L) where U is an input space and $T \subseteq X \times U \times X$, see e.g. [128, Section 1.1]. In such case each transition from x requires the input u to be chosen, and the successor state must then belong to $T|_{(x,u)}$. The choice of u is called the external nondeterminism of \mathcal{T} , whereas the choice of $x' \in T|_{(x,u)}$ is referred to as the internal nondeterminism of \mathcal{T} . The TS with no internal nondeterminism is called deterministic [128, Section 1.1]. Behavioral equivalence helps solving synthesis problems only for deterministic systems. In our setting of Definition 2.1 the input space U is omitted, however the resulting nondeterminism we treat as an external one although it looks much more similar to the internal one – same approach has been taken e.g. in [127]. To connect our setting with that of [128] one just needs to take $U = X$ and extend T from $X \times X$ to $X \times U \times X$ in the obvious way.*

Although constructing a finite abstraction $\tilde{\mathcal{T}}$ that behaviorally includes a concrete infinite TS \mathcal{T} is not a hard task, it is rarely the case that $\mathcal{T} \approx \tilde{\mathcal{T}}$. In fact, if there exists a finite $\tilde{\mathcal{T}}$ such that the latter equivalence holds, then \mathcal{T} must have a very special, somewhat piecewise constant, transition structure. Obviously, this happens only for a very limited class of continuous, control or hybrid systems. To cope with this issue, [64] proposed the use of approximate relations (rather than exact ones defined above) and corresponding pseudometrics: see the Appendix A.3 for the theoretical background on exact and approximate relations, and their connections with pseudometrics (see Appendix A.2) – we use this terminology below. Let us briefly recapitulate how these ideas apply in our setting.

The main idea is to focus on the class of problems to be solved over TSs. Even though there may not be a finite TS that is behaviorally equivalent to a given infinite one \mathcal{T} , often it is still useful to have a finite TS $\tilde{\mathcal{T}}$ such that if $\tilde{\mathcal{T}} \models H$ then $\mathcal{T} \models H'$ where H and H' are close in some sense. To formalize this idea, [64] suggested to look at output spaces that are endowed with metrics. Assume that the output space Y is endowed with some pseudometric d_Y , and endow $Y^{\mathbb{N}}$ with the pseudometric d_Y^∞ given by³

$$d_Y^\infty(y_0, y_1, \dots, y'_0, y'_1, \dots) := \bigvee_{n=0}^{\infty} d_Y(y_n, y'_n).$$

One says that $\tilde{\mathcal{T}}$ *behaviorally ε -includes* \mathcal{T} whenever $\mathsf{V}_L(\Sigma^T) \subseteq \left[\mathsf{V}_{\tilde{L}}(\Sigma^{\tilde{T}}) \right]^\varepsilon$; in such case we write $\mathcal{T} \leq_\varepsilon \tilde{\mathcal{T}}$. Clearly, $\leq_\varepsilon := (\leq_\varepsilon)_{\varepsilon \in \mathbb{R}_+}$ is an ε -preorder on TS_Y . Its symmetrization we denote by $\approx_\varepsilon := (\approx_\varepsilon)_{\varepsilon \in \mathbb{R}_+}$. Note that \leq_0 and \approx_0 are exactly the (exact) behavioral inclusion and equivalence relations defined above. See Appendix A.3 on notation for and properties of approximate relations.

Due to monotonicity of \leq , for each $\varepsilon > 0$ the relation \leq_ε is weaker than the exact behavioral inclusion \leq_0 , however it is still useful for the purposes of verification and synthesis. For example, if $\mathcal{T} \leq_\varepsilon \tilde{\mathcal{T}}$ and $\tilde{\mathcal{T}} \not\models H$ for some LT property $H \subseteq Y^{\mathbb{N}}$, then $\mathcal{T} \not\models H'$ for any H' satisfying $[H']^\varepsilon \subseteq H$. Conversely, if $\mathcal{T} \models H'$ then there exist \bar{x} and $\bar{\sigma}$ such that $\tilde{V}_{\bar{x}} \models [H']^\varepsilon$. As a result, approximate behavioral inclusion implies verification of strengthened properties, and existence of a strategy that generates behavior satisfying a relaxed property.

To provide a stronger bridge between models for TSs above and that for SSs below, let us remark on the definition of approximate behavioral inclusions. So far we have considered LT properties that are subsets of the space of output trajectories, which is also the conventional approach [17]. However, the following generalization is useful conceptually and in some applications. Imagine that every time a system has a transition a certain cost is reward is received, and one tries to maximize this reward. Formally, there is a reward function $r : X \rightarrow \mathbb{R}$ and the value to be optimized is e.g. $\sum_{n \in \mathbb{N}} \beta^n r(y_n)$, where $\beta \in (0, 1)$ is some discount factor. Notice that the latter expression is a function defined on the output trajectories, so let us say that an LT property is now any function $h : Y^{\mathbb{N}} \rightarrow \mathbb{R}$, and instead of the

³ Note that the topology of $(Y^{\mathbb{N}}, d_Y^\infty)$ often differs from the product topology induced on $Y^{\mathbb{N}}$ by (Y, d_Y) , however this fact does not cause any issues in our considerations.

satisfaction relation \models let us define the following functional:

$$\mathcal{T}(h) := \sup_{w \in \mathcal{V}_L(\Sigma^T)} h(w). \quad (2.2)$$

To see that $\mathcal{T}(\cdot)$ is indeed an extension of \models , consider the case where $h = 1_H$ is an indicator function of some $H \in Y^{\mathbb{N}}$. Notice that $\mathcal{T}(1_H) = 1$ iff $\mathcal{T} \models H$ and $\mathcal{T}(1_H) = 0$ iff $\mathcal{T} \not\models H$, which makes the satisfaction relation \models a special case of $\mathcal{T}(\cdot)$. Note also that it is not crucial whether \sup or \inf is used in (2.2). Indeed, if one changes h to $-h$ then $\mathcal{T}(\cdot)$ can be used to compute minimal values as well, rather than only maximal ones.

This new richer class of LT specifications relates to what is called quantitative model-checking of TSs [17]. There instead of saying whether a particular TS satisfies a given property or not, one needs to find the best (e.g. maximal) value the TS can achieve w.r.t. this property. As we have just shown, that covers the case of usual boolean (qualitative) satisfaction relation thanks to the use of indicator functions. In fact, the functional \mathcal{T} fully describes the behaviors of the TS \mathcal{T} , which is exactly why we denote them with the same symbol. It should come as no surprise that \mathcal{T} can be used to define behavior inclusions. In fact, this is exactly the most natural way to define them for SSs below. Proposition A.1 provides an equivalent quantitative characterization of \leq . For any $\varepsilon \in \mathbb{R}_+$ it holds that $\mathcal{T} \leq_{\varepsilon} \tilde{\mathcal{T}}$ iff

$$\mathcal{T}(f) \leq \tilde{\mathcal{T}}(f) + \varepsilon \quad \forall f \in \text{Lip}_1(Y^{\mathbb{N}}, d_Y^{\infty}). \quad (2.3)$$

The latter fact implies that if one is interested in quantitative properties of the TSs that are Lipschitz continuous functions of the output trajectory, such as a discounted accumulated cost for a Lipschitz continuous reward function, then ε -relations \approx and \simeq provide 2ε -wide estimates together with approximately optimal policies. In fact, a version of (2.3) we use below as *the* definition of behavioral inclusion for SSs, for details see the discussion in Appendix A.6.

2.2 Problem specifications

The framework of TSs and SSs is often used in optimization. In particular, one of the most prominent questions to answer is the following: what is the maximal achievable value of a given performance measure, and how can a strategy that achieves such a value be derived? Clearly, the answer crucially depends on the chosen criterion: this choice is quite broad in the literature, so let us discuss some important cases.⁴

We do not consider multi-objective optimization where the performance criterion has a partial order on its co-domain (see e.g. [26]), and instead focus on numerical criteria, namely measures taking values on \mathbb{R} . As we mentioned above, one of the most general approaches for TSs would be to define a numerical performance criterion as a function on its output trajectories, that is $h : Y^{\mathbb{N}} \rightarrow \mathbb{R}$. Unfortunately,

⁴ A comprehensive survey on different performance criteria, as well as on the general development of the theory of MDP, is given in [12, Section 3].

the generality of such approach does not allow for specific results related to the computability of the optimal solutions. Due to this reason, more specific performance criteria have attracted a significant interest, in particular the *discounted cost* (DC) and the *average cost* (AC) criteria [12]. Consider some $r : Y \rightarrow \mathbb{R}$ and define

$$\text{DC} := \sum_{k=0}^n \beta^k r(y_k),$$

$$\text{AC} := \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^n r(y_k)$$

where $\beta \in (0, 1]$ is the *discounting factor* and $n \in \bar{\mathbb{N}}$ is the *time horizon*. Clearly, both provide examples of an output-trajectory dependent function h . The case $\beta = 1$ is often referred to as the *total cost* (TC) criterion or, alternatively, the *additive cost*. These problems are extensively studied in the literature: see e.g. [20, 70] for the DC, and [12] for the AC.

Reward-based numerical performance criteria are mostly popular in the SS literature, whereas the recent research on TSs more often that not focused on specifications where $h = 1_H$ for some desired set of output trajectories $H \subseteq Y^{\mathbb{N}}$. Modal logic provides both a handy way to describe such sets, similar to expressions of natural languages, and methods to solve verification and synthesis problems for them in a unified fashion [17]. To give examples of such expressions, consider the following tasks: “the output must always be a ” or “if a was observed, then the output value of b must appear infinitely many times”. The first task corresponds to a *safety* problem, whereas the second task is more complicated, even in its qualitative description. For this purpose we introduce a modal logic called Linear Temporal Logic (LTL), which is useful in the following two aspects. First of all, it provides “a very intuitive but mathematically precise notation” [17, Section 5.1] to operate a large class of complex and interesting events. Secondly, LTL allows reducing the optimization problems for any of such events to one of the following two fundamental problems: *reachability*, requiring visiting a specified target set at least once; or *constrained repeated reachability*, requiring visiting a target set infinitely often and visiting an unsafe set only finitely often.

LTL is introduced using its *grammar*, namely the set of rules determining the construction of LTL formulae. The meaning of each formula (that is, the event corresponding to the formula) is formalized by the LTL *semantics*. It is common to refer to the output space Y as an *alphabet*, since its elements $y \in Y$ appear as *letters* in the formulae. Sequences of letters are called (infinite) *words* $w \in Y^{\mathbb{N}}$ – the output trajectories, and sets of words are *languages* $H \subseteq Y^{\mathbb{N}}$. Languages corresponding to LTL formulae fall within an important class of so-called ω -regular languages [17]. The definition of ω -regular languages is lengthy and is omitted here, for us it is only important below that such languages are measurable subsets of $Y^{\mathbb{N}}$, and hence will be valid specifications for SSs below. The definition of ω -regular languages is available in [17, Section 4.3.1]

The grammar of LTL over the alphabet Y is given by the following set of rules

$$\Phi ::= y \in Y \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid X\Phi \mid \Phi_1 U_\infty \Phi_2. \quad (2.4)$$

The definition (2.4) shall be understood as follows: if Φ_1 and Φ_2 are LTL formulae, so are the expressions $\Phi_1 \wedge \Phi_2$, $\Phi_1 U_\infty \Phi_2$, $\neg\Phi_1$ etc. Here \wedge is the standard logical *conjunction* and \neg is the logical *negation*, which allows us defining *disjunction* as $\Phi_1 \vee \Phi_2 := \neg(\neg\Phi_1 \wedge \neg\Phi_2)$. Furthermore, X and U_∞ are the *neXt* and *unbounded Until* temporal modalities whose meaning is clarified below.

The semantics of LTL formulae is defined using the notion of *accepted language*, that is $\mathcal{L}(\Phi) \subseteq Y^\mathbb{N}$ is the collection of all infinite words over Y that are accepted by the formula Φ . Firstly, we define the shift operator on infinite words $\theta : Y^\mathbb{N} \rightarrow Y^\mathbb{N}$ by

$$\theta(w_0, w_1, w_2, \dots) = (w_1, w_2, \dots).$$

The semantics of LTL formulae is defined recursively as:

$$\begin{aligned} w \in \mathcal{L}(y) &\iff w_0 = y \\ w \in \mathcal{L}(\neg\Phi) &\iff w \notin \mathcal{L}(\Phi) \\ w \in \mathcal{L}(\Phi_1 \wedge \Phi_2) &\iff w \in \mathcal{L}(\Phi_1) \cap \mathcal{L}(\Phi_2) \\ w \in \mathcal{L}(X\Phi) &\iff \theta(w) \in \mathcal{L}(\Phi), \end{aligned}$$

and in addition the semantics of the U_∞ modality is as follows:

$$\begin{aligned} w \in \mathcal{L}(\Phi_1 U_\infty \Phi_2) &\iff \theta^i(w) \in \mathcal{L}(\Phi_2) \text{ for some } i \in \mathbb{N} \text{ and} \\ &\theta^j(w) \in \mathcal{L}(\Phi_1) \text{ for all } 0 \leq j < i. \end{aligned} \quad (2.5)$$

It is useful to consider formulae describing bounded time horizon properties. We first introduce powers of X inductively as $X^0\Phi := \Phi$ and $X^n\Phi := X(X^{n-1}\Phi)$ for $n \geq 1$. Using the latter notation, it is now possible for any $n \in \mathbb{N}$ to define the formula

$$\Phi_1 U_n \Phi_2 := \bigvee_{i=0}^n \left(\bigwedge_{j=0}^{i-1} X^j \Phi_1 \wedge X^i \Phi_2 \right), \quad (2.6)$$

whose semantics is a finite-horizon equivalent of (2.5), that is

$$\begin{aligned} w \in \mathcal{L}(\Phi_1 U_n \Phi_2) &\iff \theta^i(w) \in \mathcal{L}(\Phi_2) \text{ for some } 0 \leq i \leq n \text{ and} \\ &\theta^j(w) \in \mathcal{L}(\Phi_1) \text{ for all } 0 \leq j < i. \end{aligned}$$

Note that U_∞ could be also expressed using (2.6), but the countably infinite number of operations of conjunction needed are not explicitly allowed in the syntax of LTL. We further denote $\text{true} := \bigvee_{y \in Y} y$, and introduce new temporal modalities: *eventually*, $\Diamond_n \Phi := \text{true} U_n \Phi$, and *always*, $\Box_n \Phi := \neg \Diamond_n \neg \Phi$, for all $n \in \mathbb{N}$. We further simplify the notation as $U := U_\infty$, $\Diamond := \Diamond_\infty$ and $\Box := \Box_\infty$.

Let us provide some examples of how LTL formulae can be used to describe events of interest. The event $\{y_k = a, k \geq 0\}$ can be expressed as $\Box a$, $\{\exists k \leq n : y_k = b\}$

as $\Diamond_n b$, $\{y_k = b \text{ infinitely often}\}$ as $\Box \Diamond b$, $\{\exists k : y_j = a, j \geq k\}$ as $\Diamond \Box a$, and finally the event $\{\exists k \leq n : y_k = b \text{ and } y_j = a, j < k\}$ can be expressed as $a \cup_n b$. As an additional example, the specification we mentioned above “if a was observed, then the output value of b must appear infinitely many times” can be expressed with the following formula:

$$\Box(\neg a) \vee (\Diamond a \wedge \Box \Diamond b)$$

We show now how automata theory can be used to reduce rich but complex LTL specifications to some basic ones. One may think of an automaton as a TS whose transitions are labeled with inputs over a finite alphabet⁵, and which is in addition endowed with a simple *acceptance condition* [17, Chapter 4]. An input word is accepted by an automaton if its corresponding *run* of the automaton satisfies its acceptance condition. Unlike TSs where we care about the outputs, in automata inputs play this role. Let us introduce these concepts formally.

Definition 2.4 *Given an alphabet Y , a deterministic ω -automaton over Y is a tuple $\mathcal{D} = (Q, q^s, Y, \mathbf{t}, A)$ where Q is a finite set, $q^s \in Q$, $\mathbf{t} : Q \times Y \rightarrow Q$ is some map and $A \subseteq Y^\mathbb{N}$. We say that Q is the state space, q^s is the initial condition, Y is the alphabet, \mathbf{t} is the transition map and A is the acceptance condition of \mathcal{D} .*

Any word $w \in Y^\mathbb{N}$ induces a run $z \in Q^\mathbb{N}$ of \mathcal{D} which is defined as follows: $z_0 = q^s$ and $z_{k+1} = \mathbf{t}(z_k, w_k)$ for any $k \in \mathbb{N}$. We can then introduce a map $\mathbf{T} : Y^\mathbb{N} \rightarrow Q^\mathbb{N}$ that assigns to any input word the corresponding run.

The acceptance condition of an automaton indicates which runs are accepted by the automaton ($z \in A$) and which are not ($z \notin A$). Similarly, we say that a word is accepted by a deterministic automaton if the corresponding run is accepted. There are several versions of acceptance conditions for automata in the literature. In the context of this work the following three are the most important:

- (DRA) for a *deterministic Rabin automaton* the A is defined using $(F'_i, F''_i)_{i \in I}$, where I is some finite index set and $F'_i, F''_i \subseteq Q$ for each $i \in I$. A DRA accepts a run z if there exists $i \in I$ such that z visits F'_i an infinite number of times and F''_i only a finite number of times.
- (DBA) a *deterministic Büchi automaton* is a special case of a DRA with I being a singleton and $F'' = \emptyset$, that is precisely runs that visit $F' \subseteq Q$ infinitely often are accepted.
- (DFA) a *deterministic finite automaton* can be seen as a special case of a DBA⁶ with all final states having self-loops ($\mathbf{t}(q, y) = q$ for any $q \in F'$, $y \in Y$), that is it

⁵ Here we only consider deterministic automata – those for which the current input and state uniquely determine the next state.

⁶ While it is canonical to introduce a DFA on finite words [17, Definition 4.9], we introduce it here on infinite words for the sake of consistency: in that way we do not have to consider both spaces of finite and infinite words over the alphabet Y , and can just focus on the latter. It should be clear that our definition is also consistent with the canonical one in [17, Definition 4.9]: an infinite word $w \in Y^\mathbb{N}$ is accepted by our DFA if and only if there exists a finite prefix $w' \in \Sigma^*$ that is accepted by their DFA.

accepts precisely those runs that eventually visit F'^7 .

For an automaton \mathcal{D} we define its *accepted language* as the set of all infinite words that are accepted by \mathcal{D} ; we further denote this language by $\mathcal{L}(\mathcal{D}) := T^{-1}(A)$.

Accepted languages of DRA are exactly ω -regular languages [17, Theorem 10.55], so in particular for any LTL formula Φ there exists a DRA \mathcal{D}^Φ such that $\mathcal{L}(\Phi) = \mathcal{L}(\mathcal{D}^\Phi)$. Furthermore, DBA (DFA) are strictly less expressive than DRA (DBA) – for details see [17, Chapter 4]. We consider all three kinds of automata, rather than focusing on the most expressive DRA, due to the following reason. For any automaton \mathcal{D} , verification of specification given by its accepted language can be reduced to a basic specification encoded by the accepting condition of such automaton. Unfortunately, solving this for DRA is rather difficult over SSs and we only provide partial results for the DBA case (Section 4.3), whereas the acceptance condition of the DFA allows for a much more complete characterization (Section 4.2).

Before we proceed, let us provide some examples of automata. The DBA for the formula task is given in Figure 2.1(a): here if we do not label the transition (as the loop at q^1) it means that the transition happens for any label. The final state is q^0 as indicated by a double circle. As we have mentioned above, the analysis of the DBA acceptance condition is more complicated than that of the DFA one, hence even if the original LTL formula does not allow for the DFA expression, it is worth checking whether its negation does allow for one. For example, the DFA for the negation of the first task is given in Figure 2.1(b).




Figure 2.1: Automata representation of the first task of the case study

To explain how to reduce verification of a specification encoded as an automaton to verification of its accepting conditions, we need to introduce a composition between a TS and an automaton.

Definition 2.5 Given a TS $\mathcal{T} = (X, T, L, Y)$ and an automaton $\mathcal{D} = (Q, q^s, Y, t, A)$, their composition is a TS $\hat{\mathcal{T}} = \mathcal{T} \otimes \mathcal{D} = (\hat{X}, \hat{T}, \hat{L}, \hat{Q})$, where $\hat{X} := X \times Q$, $\hat{L} : (x, q) \mapsto q$ and

$$(x, q)\hat{T}(x', q') \iff xTx' \text{ and } q' = t(q, L(x)).$$

⁷ An important version of the DFA has an n -horizon acceptance condition [137, Section 2.4], which requires the run to visit F in at most n steps. This is useful when one needs to express formulae in bounded LTL (BLTL) – a fragment of LTL (for details see Section C.1).

To elucidate the idea behind this notion of composition, note that \hat{T} can be understood as follows: if the composed system is in the state (x_k, q_k) , the next state must satisfy the following recursion formula

$$\begin{cases} x_{k+1} & \in T|_{x_k}, \\ q_{k+1} & = \mathfrak{t}(q_k, L(x_k)). \end{cases}$$

This dynamics should be understood as follows: the x -coordinate of the new state evolves according to the law T of the original TS \mathcal{T} , and its output $L(x)$ is used as an input to the transition system, which produces the q -coordinate. It should be clear that if \mathcal{D} is e.g. a DFA with the set of final states being F' , then \mathcal{T} satisfies $\mathcal{L}(\mathcal{D})$ iff the output of $\mathcal{T} \otimes \mathcal{D}$ visits F' at least once: for details see e.g. [17]. We also provide more formal exposition in the next section when talking about SSs.

2.3 Stochastic systems

There are many models of discrete-time probabilistic systems available in the literature, most of which have a dynamic feature similar to the TSs: once one knows the current state $x \in X$, one has all the necessary information about the successor state. In case of TSs this information contains a subset of a state space $T|_x$, that is the set of all possible successor states. In contrast, in the stochastic case one gets a distribution (or a collection thereof) over the successor states.

For probabilistic systems such feature is called the Markovian property by the name of a mathematician A.A. Markov who studied stochastic processes whose conditional distributions only depended on the current state, rather than the whole past: such processes are now called Markov processes⁸. The difference between a probabilistic model and a stochastic process is that to each model there may correspond different stochastic processes, or non at all. For example, a discrete-time Markov Chain (MC) [105] is a probabilistic model that assigns to any initial distribution on its state space the corresponding (unique) stochastic process: a distribution over the state trajectories. In turn, a probabilistic model called Markov Decision Process (MDP) model [111] is a generalization of the discrete-time MC one, which in addition to fixed initial distribution requires specifying a sequential decision rule to determine the stochastic process⁹. Such mapping from a probabilistic model to stochastic processes is actually the LT semantics of probabilistic models. Not all probabilistic models are endowed with a natural LT semantics, e.g. probabilistic transition systems (PTSs) [75], also known as labelled Markov processes (LMPs) [25]. However, to the best of our knowledge, all the models where such semantics is explicitly defined, have it defined in a very same way – namely there is no argument about what exactly is the LT semantics of a probabilistic model

⁸ As a historical remark, interestingly one of the applications Markov considered for these processes was to compute conditional probabilities of vowels and consonants in one of the most famous Russian poems, see [99] for the English translation of the original paper.

⁹ As another historical remark, it is commonly assumed that Bellman introduced the MDP together with the technique to solve optimization problem called dynamic programming (DP), see e.g. [19].

(in case it allows for one), which makes all these models somewhat equivalent to MDPs¹⁰. Such model frameworks e.g. comprise probabilistic automata (PAs)¹¹ [114] and gambling models (GMs) [48]. The differences in frameworks are mostly technical, and the interpretation is very similar as we have mentioned above.

Due to the fact that LT semantics of probabilistic models is pretty much unique, one can work out all necessary results for just one of those models, and then seamlessly extrapolate the results to the rest of them. We choose to work in a setting inspired by the GMs: similarly to TSs and a model from [89, Definition 1] and in contrast to MDPs, it does not have a distinguished action space which simplifies its analysis. Yet, this model is as expressible as MDP [136], which we emphasize by showing how to extrapolate the obtained results to the MDP framework in Section 3.4.

Definition 2.6 *A stochastic system (SS) is a tuple $\mathcal{S} = (X, \Gamma, Y, L)$ where X and Y are Borel spaces, $\Gamma \in \mathcal{A}(X \times \mathcal{P}(X))$ is an l.t.r. and $L \in \mathcal{B}(X, Y)$. We say that X is the state space, Γ is the stochastic relation, Y is the output space and L is the output map of \mathcal{S} . The SS \mathcal{S} is said to be finite if the set Γ is finite, otherwise the SS \mathcal{S} is called infinite. The set of all SSs with the output space Y is denoted by SS_Y . If $\Gamma|_x$ contains exactly one element for each $x \in X$, we denote it by $\Gamma(x)$, and say that the SS \mathcal{S} is autonomous.*

Let us give some comments on Definition 2.6. First of all, a pair (X, Γ) with X a Borel space and $\Gamma \in \mathcal{A}(X \times \mathcal{P}(X))$ an l.t.r. is a GM as per [94]. Although in classical GMs there is no specified output structure (Y, L) , a GM can be understood as an SS with $Y = X$ and $L = \text{id}_X$. The PTS model is syntactically similar to SS, and in the PTS literature it is often assumed that Γ relates states to discrete distributions only, rather than arbitrary elements of $\mathcal{P}(X)$. This allows avoiding issues related to measurability, so X and Y can be arbitrary sets, Γ an arbitrary l.t.r. and L an arbitrary map, similar to a TS. The fact that we work with general, not necessarily discrete, probability measures should explain why we require X and Y to be Borel spaces and L to be a Borel map. We comment on the measurability of Γ later. Unless the contrary is specified, further in this section we always assume that the SS $\mathcal{S} := (X, \Gamma, Y, L)$ is given and fixed.

As in case of TSs, Definition 2.6 only describes the syntax of the SS, so let us which dynamics does SS define. The evolution of a sample path of SS \mathcal{S} is similar to that of TSs with the only difference that first the distribution over states is chosen, and then the successor state is drawn from that distribution. That is, if the current state of the SS is x_n , the output we also observe the output value of $y_n = L(x_n)$, choose the distribution of the successor state among admissible probability measures $\gamma_n \in \Gamma|_{x_n}$, and the new state x_{n+1} is drawn randomly according to γ_n . To formalize this procedure again we need to define how exactly his choice is made,

¹⁰ MDPs themselves also appear in the literature under different names such as stochastic optimal control models [20], Markov Control Models [70], controlled discrete-time Stochastic Hybrid Systems [4] and controlled discrete-time Markov processes [137].

¹¹ The PA model is slightly different from the MDP one in that it allows for two types of non-determinism: both internal and external [17, 128], so it can be interpreted as a stochastic game [89] with an agent playing for the external non-determinism, and an adversary playing for the internal one.

namely introduce the strategies. In contrast to the case of TSs, strategies is a necessary component of the definition of semantics of SSs¹².

Definition 2.7 *A strategy for the SS \mathcal{S} is a sequential decision rule $\sigma = (\sigma_n)_{n \in \mathbb{N}}$, where the map $\sigma_n \in \mathcal{U}(X|X^n)$ is such that $\sigma_n(x_0, \dots, x_n) \in \Gamma|_{x_n}$ for each $x_i \in X$, $i \in [0; n]$ and $n \in \mathbb{N}$. The set of all such strategies we denote by Σ^Γ .*

We can now comment on why Γ is required to be an analytic set in the definition of SSs. To work with general distributions over successor states, not only the discrete ones, we needed measurability assumptions on the strategies. Although it is easier to work with Borel strategies, it may happen that Γ does not contain a graph of a Borel map even if Γ is a Borel space itself [23], hence it may happen that there would not be any strategies at all. To guarantee the existence of such a Borel map we would need rather restrictive assumptions on Γ . At the same time, if Γ is at least analytic, let alone Borel, it is guaranteed to contain a graph of a universally measurable map by Proposition C.17. Due to this reason, in Definition 2.6 we require stochastic relations of SSs to be analytic; this is also conventional in the literature on stochastic control both for MDPs and GMs [20, 97]. Note that if \mathcal{S} is autonomous, then if Γ is analytic, it is Borel and hence is a graph of some Borel map. Like in case of TSs only the GM component of the SS \mathcal{S} , which is given by a pair (X, Γ) , plays a role in the definition of strategies and the output structure (Y, L) has no effect¹³.

To understand how to define LT semantics for SSs, let us think what are the natural questions that we ask when dealing with probabilistic systems. For example, in case of TS one would wonder whether it has a sample trajectory that reaches a goal set at some point, or what is the reward one accumulates while moving along such trajectory. Since an SS would have several sample trajectories with different likelihoods, there instead one asks what is the *probability* of a sample trajectory passing by a goal set, or what is the *expected* reward accumulated over it. For this reason, naturally the LT semantics of an SS is not characterized by single trajectories (runs), but instead by probability measures defined on the space of those trajectories. These measures are also known as stochastic processes. Such stochastic process is defined as follows: given any initial distribution $\alpha \in \mathcal{P}(X)$ and a strategy $\sigma \in \Sigma^\Gamma$ there exists a unique probability measure¹⁴ $p \in \mathcal{P}(X^\mathbb{N})$ whose first marginal is α and whose transition probabilities are given by $(\sigma_n)_{n \in \mathbb{N}}$, i.e.

$$p|_0 = \alpha, \quad \frac{dp|_{n+1}}{dp|_n} = \sigma_n \text{ (} p|_n \text{-a.s.)} \quad \forall n \in \mathbb{N}. \quad (2.7)$$

The measure satisfying (2.7) we denote by P_α^σ and call a *strategic measure*. The sets of all strategic measures and those that start at $\alpha \in \mathcal{P}(X)$ are denoted by $S(\Sigma^\Gamma)$ and $S(\Sigma^\Gamma, \alpha)$ respectively. Clearly, here $S(\Sigma^\Gamma) := \bigcup_{\alpha \in \mathcal{P}(X)} S(\Sigma^\Gamma, \alpha)$.

¹² For probabilistic models, strategies are also known as control policies [20] or schedulers [17, Chapter 10].

¹³ Note also that Γ uniquely determines the transition structure as it is an l.t.r., and hence X can be defined as a left projection of Γ . In particular, Γ alone uniquely defines the set of strategies: due to this reason we use the notation Σ^Γ instead of cumbersome $\Sigma^{X, \Gamma}$.

¹⁴ See the Appendix A.1, the part on kernels.

Recall that for TSs we are able to deduce the initial state and the strategy uniquely given a run they generate. A similar result applies to the case of SSs, although it requires more technical machinery such as the existence of conditional distributions for a given product measure¹⁵. Given any $p \in S(\Sigma^\Gamma)$ we can obtain the corresponding α and σ by applying (2.7); note however that although α is unique in such case, each σ_n is only determined ($p|_n$ -a.s.)-uniquely.

To each strategic measure $P_\alpha^\sigma \in S(\Sigma^\Gamma) \subseteq \mathcal{P}(X^\mathbb{N})$ (that corresponds to an internal run in TSs) we assign a corresponding *observation measure* $Q_\alpha^\sigma := L_*P_\alpha^\sigma \in \mathcal{P}(Y^\mathbb{N})$. We further denote by $S_L(\Sigma^\Gamma) := L_*(S(\Sigma^\Gamma))$ and $S_L(\Sigma^\Gamma, \alpha) := L_*(S(\Sigma^\Gamma, \alpha))$ sets of all observation measures, and those initiated at α respectively. To define the behaviors of the SS \mathcal{S} , and their relations to outputs, let us first discuss LT properties of SSs. Connecting to the case of TSs, consider some set of desired output trajectories $H \subseteq Y^\mathbb{N}$, defined e.g. using an LTL formula, and think of the evolution of the SS for a fixed initial distribution α and strategy σ . The output trajectory hence is a random element with the distribution P_α^σ . A realization of the output trajectory may belong to H or not, depending on the result of each probabilistic draw, so in each non-trivial case neither H nor its negation H^c is surely satisfied by the output trajectory. Perhaps in some applications it would be of interest whether there does exist a single element of H which is “possible”, meaning it has positive probability. Such statement would make some sense for SSs where only discrete distributions are allowed (e.g. finite SSs): in that case some authors even define the strategic measures not on all realizations of state trajectories (space $X^\mathbb{N}$), but only on those where each transition $x_n \rightarrow x_{n+1}$ happens with a positive probability [17, Chapter 10]. This would make a little difference from the non-determinism of the TSs, will not use any specifics of the probabilistic case and this is not applicable to general SSs, since often each measure in $\Gamma|_{x_n}$ gives zero mass to any single candidate for the successor state x_{n+1} , as also each strategic measure assigns zero probability to any separate output trajectory. Due to this reason, we define strategic and observation measures on spaces of all realizations of state and output trajectories respectively: this is also a common practice in literature on probability and stochastic processes. Thus it one cannot distinguish which realizations of output trajectories are “possible” and which are not. Instead we can only talk about the probability that the realization of an output trajectory belongs to a given set H , the now conventional way to define LT properties of SSs in the computer science literature [17, Section 10.1], and the only way considered in the literature on probability theory, see e.g. [94]. See also [133, 135] for a similar discussion and Appendix A.4 for more details.

Based on the ideas of the previous paragraph, one can define an LT specification over SS_Y as some set $H \in \mathcal{U}(Y^\mathbb{N})$: the measurability of H is needed to make sure that $q(H)$ is well-defined for each $q \in S_L(\Sigma^\Gamma)$. We have already seen in Section 2.2 that automata and LTL formulae can be used to define some interesting subsets of $Y^\mathbb{N}$ in a handy way. In fact, not any language (subset of $Y^\mathbb{N}$) is a valid specification for a SS, as it may happen not to be measurable, and hence one could not say what is the probability of an output trajectory being in this set: recall that $Y^\mathbb{N}$

¹⁵ See the Appendix A.1, the part on kernels. For Borel spaces (regular) conditional distributions do always exist [52], but for general measurable space that may not be the case.

is uncountable, and any uncountable Borel space admits for a set which is not universally measurable [82]. Fortunately, languages accepted by all automata we consider here and LTL formulae are ω -regular [147], and hence are Borel subsets of $Y^{\mathbb{N}}$ [145, Proposition 2.3]. On the other hand, not any measurable language is ω -regular: clearly any singleton $\{w\}$ generated by a word $w \in Y^{\mathbb{N}}$ is measurable, but the language $\{w\}$ may not be ω -regular if w is not a periodic word.

Similarly to TSs, for SSs we do not only focus on specifications expressible as subsets of the output trajectories, and instead we can define a more general quantitative LT property as a map $h \in \text{b}\mathcal{U}(Y^{\mathbb{N}})$ and talk about its expected value $q[h]$. It is a more general as one can always take $h = 1_H$ to express a set as a function. As for TSs, we also define a maximization functional associated with an SS

$$\mathcal{S}_\alpha(h) := \sup_{q \in \mathcal{S}_L(\Sigma^\Gamma, \alpha)} q[h], \quad (2.8)$$

where the arguments are $\alpha \in \mathcal{P}(X)$ and $h \in \text{b}\mathcal{U}(Y^{\mathbb{N}})$. Recall that as a first step for TSs we defined a satisfaction relation, and only after that we introduced a maximization functional as a more general version of the former. In contrast, even if we only focus on properties expressed as sets of output trajectories for SSs, we cannot have a single natural satisfaction relation which would allow us to say whether a particular SS satisfies a given property or not. When dealing with SSs such satisfiability will genuinely be a continuous quantity, rather than a boolean outcome. For this reason it seems reasonable to care about those quantities, which are exactly the values of the functional in (2.8), namely the maximal probabilities and/or expectations. Note that one can use \mathcal{S}_α to express the minimal values as well by changing h to $-h$. Furthermore, if one still prefers to deal with relations, this functional can be further used to introduce them as $\models_r \subseteq \text{SS}_Y \times \text{b}\mathcal{U}(Y^{\mathbb{N}})$ in the following way:

$$\mathcal{S} \models_r h \iff \sup_{\alpha \in \mathcal{P}(X)} \mathcal{S}_\alpha(h) \geq r$$

where a parameter r is any real number for \mathbb{R} . For example, that way we can ask questions of the kind “does there exist a strategy which makes a trajectory of \mathcal{S} belong to a set H with probability of at least r ?” This is similar to what is done in probabilistic computational tree logic (PCTL) [17, Section 10.2], however we do not focus much on this approach here. Instead we just ask what exactly is the maximal probability that a trajectory of the SS \mathcal{S} belongs to H , and the functional \mathcal{S} precisely gives us the answer to this question. Like in TSs, to emphasize that this functional is everything we need to know about the behavior of its underlying SS we denote them with the same symbol.

Unlike the non-stochastic case, the class of problems expressible as a function h on output trajectories is not the largest one that can be defined for SSs. Arguably one of the most general approaches to the definition of numerical performance criteria over SSs has been considered in [54]. There, a criterion is simply any function on strategic measures $F : \mathcal{P}(Y^{\mathbb{N}}) \rightarrow \mathbb{R}$. Here we only focus on a special case of it, namely the *expected utility* [83, 84, 85], where h is precisely is a utility function. One can see that this class only provides linear functions F of strategic measures, of a shape $F_h(q) = q[h]$. An example of a criterion not expressible in this form is a

probabilistic version of AC

$$\text{AC} := \limsup_{n \rightarrow \infty} \mathbb{Q}_\alpha^\sigma \left[\frac{1}{n} \sum_{k=0}^n r(y_k) \right]$$

For the expected utility criteria there are general results on characterization of the optimal values and strategies, and more specific results were obtained when the focus was on a specific criterion. For example, the probabilistic version of DC

$$\text{DC} := \mathbb{Q}_\alpha^\sigma \left[\sum_{n \in \mathbb{N}} \beta^n r(y_n) \right]$$

was studied in [20, 70], whereas [95] contains results on criteria considered in gambling literature. We are particularly interested in properties expressible as automata, solving which for SSs can also be reduced to reachability and repeated reachability like for TSs. This technique has been employed to study finite SSs [32], leading to analytical solutions for that setup, however there results seem to crucially depend upon the finiteness assumption, which we relax here. As above, for the reduction technique we introduce a composition between an SS and an automaton.

Definition 2.8 *Given an SS $\mathcal{S} = (X, \Gamma, L, Y)$ and an automaton $\mathcal{D} = (Q, q^s, Y, \mathbf{t}, \mathbf{A})$, their composition is an SS $\hat{\mathcal{S}} = \mathcal{S} \otimes \mathcal{D} = (\hat{X}, \hat{\Gamma}, \hat{L}, Q)$, where $\hat{X} := X \times Q$, $\hat{L} : (x, q) \mapsto q$ and*

$$(x, q) \hat{\Gamma}(\mu, \nu) \iff x\Gamma\mu \text{ and } \nu = \delta(\mathbf{t}(q, L(x))).$$

Similarly to TSs, the dynamics of the composed SS just follows that of the original SS, the outputs of which steer the state of the automaton: if the composed system is in the state (x_k, q_k) , the next state must satisfy the following recursion formula

$$\begin{cases} \mu_{k+1} & \in \Gamma|_{x_k}, \\ x_{k+1} & \sim \mu_{k+1}, \\ q_{k+1} & = \mathbf{t}(q_k, L(x_k)). \end{cases}$$

One of the first differences in working with TSs versus SSs we face now: in the former case the composition system was obviously a TS, whereas in the latter we need to show it satisfies the desired measurability condition. For example, \hat{X} is a Borel space as a product of two Borel spaces, \hat{L} is a projection and hence Borel-measurable, but the fact that $\hat{\Gamma}$ is an analytic subset of $\hat{X} \times \mathcal{P}(\hat{X})$ is less straightforward. Notice that

$$\begin{aligned} \hat{\Gamma} &= \{(x, q, \mu, \nu) : (x, \mu) \in \Gamma \text{ and } \nu = \delta(\mathbf{t}(q, L(x)))\} \\ &= \left(\text{proj}_{X \times \mathcal{P}(X)}^{-1}(\Gamma) \right) \cap (\mathcal{P}(X) \times \text{Gr}(\delta \circ \mathbf{t} \circ (\text{id}_Q \sqcup L))) \end{aligned}$$

and hence is analytic as an intersection of two analytic sets.

To formalize the composition technique, we first need to establish a strategy equivalence between optimal utilities over \mathcal{S} and $\hat{\mathcal{S}}$. More precisely, we relate classes Σ^Γ and $\Sigma^{\hat{\Gamma}}$ as follows. The former class can be treated as a subclass of the latter, where strategies do not depend on q -coordinates of the history, so we let $\mathfrak{I} : \Sigma^\Gamma \rightarrow \Sigma^{\hat{\Gamma}}$ denote the corresponding embedding map. For the reverse, we introduce a projection map $\mathfrak{P} : \Sigma^{\hat{\Gamma}} \rightarrow \Sigma^\Gamma$ by the formula

$$(\mathfrak{P}\hat{\sigma})_n(x_0, x_1, \dots, x_n) := \hat{\sigma}_n(x_0, q_0, x_1, q_1, \dots, x_n),$$

where $q_0 = q^s$ and $q_{k+1} = \mathfrak{t}(q_k, L(x_k))$, for all $0 \leq k < n$.

Lemma 2.9 *For any $\alpha \in \mathcal{P}(X)$, and any strategies $\sigma \in \Sigma^\Gamma$ and $\hat{\sigma} \in \Sigma^{\hat{\Gamma}}$, it holds that*

$$\mathcal{Q}_\alpha^\sigma(\mathcal{L}(\mathcal{D})) = \hat{\mathcal{Q}}_{\alpha \otimes \delta_{q^s}}^{\mathfrak{I}\pi} (Y^\mathbb{N} \times A), \quad \hat{\mathcal{Q}}_{\alpha \otimes \delta_{q^s}}^{\hat{\sigma}} (Y^\mathbb{N} \times A) = \mathcal{Q}_\alpha^{\mathfrak{P}\sigma}(\mathcal{L}(\mathcal{D})).$$

Proof: Let us introduce a map $\chi : X^\mathbb{N} \rightarrow \hat{X}^\mathbb{N}$ as $\chi := \text{id}_{X^\mathbb{N}} \sqcup (\mathfrak{T} \circ L)$, so that given a path $v \in X^\mathbb{B}$ this map returns a path $\hat{v} = \chi(v) \in \hat{X}^\mathbb{N}$ which has the same x -coordinates, and the q -coordinates of which are obtained using the automaton transition map. As a result, for any $\alpha \in \mathcal{P}(X)$ and any $\sigma \in \Sigma^\Gamma$ it holds that

$$\mathcal{P}_\alpha^\sigma(L(v) \in \mathcal{L}(\mathcal{D})) = \mathcal{P}_\alpha^\sigma((\mathfrak{T} \circ L)(v) \in D) = (\chi_* \mathcal{P}_\alpha^\pi) (\hat{L}(\hat{v}) \in Y^\mathbb{N} \sqcup A).$$

Applying definitions of maps \mathfrak{I} and \mathfrak{P} immediately yields the desired result. \square

Given that Lemma 2.9 matches probability strategy per strategy, obviously it applies also to maximization and minimization, which gives us the desired reduction method. For a wide class of such LT specifications there were developed efficient algorithms over the finite SSs, and the dedicated model-checking software was built [80, 73]. This motivates us to look for finite models that approximate infinite SSs, which in turn leads to formalization of behaviors for them.

2.4 SS behaviors

Recall that in case of TSs a realization of an output trajectory alone determines the satisfaction of any LT property, e.g. given $w \in Y^\mathbb{N}$ we can say whether it belongs to $H \subseteq Y^\mathbb{N}$ or not, or which value a utility function h achieves on w . In contrast, when dealing with SSs if we are given a particular realization of the output trajectory it tells us nothing about which LT properties are satisfied: we need to know its strategic measure instead. Even if we know in addition the realization of the state trajectory and which distributions were chosen at each time step, it will only determine a limited subset of LT properties – precisely those expressible in the form $h(q|_0, q|_1, q|_2, \dots)$ such as the discounted cost criterion. Such approach was proposed in [148], however note that this information would not be sufficient even to compute the value of $q(y_0 \in A, y_1 \in B)$, not to mention LTL formulae

which contain some operations besides the “Next” modality and the negation¹⁶. As a result, not the particular choice of marginal distributions $\{q|_0, q|_1, q|_2, \dots\}$ determines LT properties of the SS, but rather the way they have been chosen conditional on the previous history: namely, the strategy σ itself.

Based on the ideas above it is natural to say that the *behavior* of the SS \mathcal{S} is the family of its observation measures. In most of the applications not all the behaviors of an SS are of interest, but only extreme ones – those where the maximum in the definition of \mathcal{S}_α is achieved; of course, such maximum may depend on a particular LT property. Due to this reason, instead of defining the approximate behavioral inclusion for SS as a (strong) set inclusion of families of observation measures, we exploit the concept of the weak inclusion introduced in the Appendix A.5. The former would be a direct analogue of the method we used for TSs, however the latter also has its similarities – see Appendix A.6 for a discussion.

The case of TSs strongly motivates to use approximate relations, so we skip the step of precise behavioral inclusion for SSs. Given an SS \mathcal{S} and a new SS over the same output space $\bar{\mathcal{S}} := (\bar{X}, \bar{\Gamma}, Y, \bar{L})$ we say that the SS $\bar{\mathcal{S}}$ *behaviorally ε -includes* \mathcal{S} whenever for any $\alpha \in \mathcal{P}(X)$ there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ such that $S_L(\Sigma^\Gamma, \alpha) \sqsubseteq^\varepsilon S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})$, that is

$$\mathcal{S}_\alpha(h) \leq \bar{\mathcal{S}}_{\bar{\alpha}}(h) + \varepsilon, \quad \forall h \in \text{b}\mathcal{U}_1(Y^\mathbb{N}).$$

In such case we write $\mathcal{S} \leq_\varepsilon \bar{\mathcal{S}}$. From the properties of \sqsubseteq it follows that \leq is an ε -preorder. Note that $\mathcal{S} \leq_\varepsilon \bar{\mathcal{S}}$ implies that

$$\inf_{q \in S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})} \bar{q}[h] - \varepsilon \leq \inf_{q \in S_L(\Sigma^\Gamma, \alpha)} q[h] \leq \sup_{q \in S_L(\Sigma^\Gamma, \alpha)} q[h] \leq \sup_{q \in S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})} \bar{q}[h] + \varepsilon \quad (2.9)$$

for any LT property $f \in \text{b}\mathcal{U}_1(Y^\mathbb{N})$, hence \leq provides bounds for the verification problem. However, due to the way it is defined it may not help solving the synthesis problem: we show below the approximate bisimulation relation between SS (which is stronger and implies approximate behavioral equivalence) is useful for the both problems. Also, the leftmost and rightmost values in (2.9) are often rather conservative bounds for $\sup_{q \in S_L(\Sigma^\Gamma, \alpha)} q[f]$ and $\inf_{q \in S_L(\Sigma^\Gamma, \alpha)} q[f]$ respectively, as the quality of such bounds crucially depends on the gap between the latter two values, which cannot be controlled by ε . This gap is obviously zero for autonomous SSs, since there is only one element in the set of observation measures, however in general it may be rather large. Unfortunately, even if $\mathcal{S} \leq_\varepsilon \bar{\mathcal{S}}$ and $\bar{\mathcal{S}} \leq_\varepsilon \mathcal{S}$, it may still happen that for some $\alpha \in \mathcal{P}(X)$ there is no $\bar{\alpha} \in \mathcal{P}(\bar{X})$ such that \mathcal{S}_α can be bounded in terms of $\bar{\mathcal{S}}_{\bar{\alpha}}$ and ε . Due to this reason, we do not define approximate behavioral equivalence for SS as a symmetrization of an approximate behavioral inclusion, and instead propose the following definition.

Definition 2.10 *SSs \mathcal{S} and $\bar{\mathcal{S}}$ are behaviorally ε -equivalent if for any initial distribution $\alpha \in \mathcal{P}(X)$ ($\bar{\alpha} \in \mathcal{P}(\bar{X})$) there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ ($\alpha \in \mathcal{P}(X)$) such that $S_L(\Sigma^\Gamma, \alpha) \equiv_\varepsilon S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})$*

¹⁶ It shall be clear now that although probabilistic models of [148] are endowed with LTL specifications, the LT semantics of these models is different from ours or the ones in [17, Chapter 10]. For example, it is neither clear how to interpret the fact that a probabilistic model of [148] satisfies a given LTL specification, nor how to use guarantees provided there.

$S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})$; in such case we write $\mathcal{S} \approx_{\varepsilon} \bar{\mathcal{S}}$.

If $\mathcal{S} \approx_{\varepsilon} \bar{\mathcal{S}}$ we can obtain a solution of a verification problem over the former system by solving it on the latter since the following inequality applies:

$$\begin{aligned} |\mathcal{S}_{\alpha}(h) - \bar{\mathcal{S}}_{\alpha}(h)| &= \left| \sup_{q \in S_L(\Sigma^{\Gamma}, \alpha)} q[h] - \sup_{\bar{q} \in S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})} \bar{q}[h] \right| \\ &= \left| \inf_{q \in S_L(\Sigma^{\Gamma}, \alpha)} q[h] - \inf_{\bar{q} \in S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})} \bar{q}[h] \right| \leq \varepsilon \end{aligned} \quad \forall h \in \text{b}\mathcal{U}_1(Y^{\mathbb{N}}). \quad (2.10)$$

Note that in (2.10), in contrast to (2.9), the upper and lower bounds for the solution of the maximization/minimization problem over \mathcal{S} are given in terms of the solution of the same type problem over $\bar{\mathcal{S}}$. As a result, when providing such bounds the gap between solutions for the maximization and minimization problem of $\bar{\mathcal{S}}$ has no effect on bounds, which hence converge to each other when ε goes to 0.

Let us emphasize the difference in the notation $\sqsubseteq (\equiv)$ and $\leq (\approx)$: the former is a relation on the families of measures, whereas the latter is a relation on SSs. Even though \leq and \approx are defined using \sqsubseteq and \equiv , another part of the definition puts important conditions concerning initial distributions. See the Appendix A.3 for the notation used here for relations. There is an interesting connection between the ways the approximate behavioral inclusion is defined for TSs and SSs, see the discussion in the Appendix A.6.

The following classification of strategies is often used in the literature. Although we do not focus on any particular class here, later we are interested in preserving the strategy class when refining a strategy from the abstraction SS to the concrete one. All strategies are by default referred to as *history-dependent*.

Definition 2.11 A strategy $\sigma = (\sigma_n)_{n \in \mathbb{N}} \in \Sigma^{\Gamma}$ is called

- Markov if σ_n depends only on x_n for all $n \in \mathbb{N}$, subclass of such strategies is denoted by Σ_M^{Γ} ;
- stationary if it is Markov and $\sigma_n = \sigma_0$ for all $n \in \mathbb{N}$;
- (ε, α) -optimal for $h \in \text{b}\mathcal{U}(Y^{\mathbb{N}})$ if $Q_{\alpha}^{\sigma}[h] \geq \mathcal{S}_{\alpha}(h) - \varepsilon$;
- uniformly ε -optimal for $h \in \text{b}\mathcal{U}(Y^{\mathbb{N}})$ if it is (ε, α) -optimal for each $\alpha \in \mathcal{P}(X)$.

Let us provide an example of two simple finite autonomous SSs which intuitively are behaviorally equivalent. After we introduce bisimulation for SSs, we prove that these two systems are bisimilar and hence indeed behaviorally equivalent.

Example 2.12 Let the output space be a two-letter alphabet $Y = \{A, B\}$, and consider an SS \mathcal{S} on Figure 2.2(a). Here $X := \{a_1, a_2, b\}$, $L(\{a_1, a_2\}) := A$ and $L(b) := B$, and Γ is given by $\Gamma(a_1) := (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, $\Gamma(a_2) := (\frac{2}{3}, 0, \frac{1}{3})$ and $\Gamma(b) := (0, 0, 1)$. The SS $\bar{\mathcal{S}}$ is

depicted on Figure 2.2(b) with $\bar{X} := \{\bar{a}, \bar{b}\}$, $\bar{L}(\bar{a}) := A$, $\bar{L}(\bar{b}) := B$ and $\bar{\Gamma}(a) = (\frac{2}{3}, \frac{1}{3})$, $\bar{\Gamma}(b) = (1, 0)$.

Although states $a_1, a_2 \in X$ have different successor distributions, they share the same output A , and those distributions give the same probability to b . Thus, we may expect that \mathcal{S} can be “lumped” as $\bar{\mathcal{S}}$ so that they have same behaviors. This is perhaps one of the simplest example of behaviorally equivalent SSs, and in particular we would like our notion of bisimulation to hold true for \mathcal{S} and $\bar{\mathcal{S}}$. We formalize this idea in Example 3.9.




Figure 2.2: Example of two finite autonomous SSs.

The main focus of this chapter is on providing sufficient conditions for approximate behavioral inclusion of SSs, which we build off the notion of approximate simulation for TSs [64], and they ways to construct an approximately similar finite SS for a given possibly infinite one.

Before we proceed, let us just comment on the way we have defined behaviors and their inclusions above. If we departed from the case of TSs, how would we define behaviors for SSs? Recall that a TS $\bar{\mathcal{T}}$ behaviorally includes \mathcal{T} if any output trajectory of the latter system can be matched by that of the former one. This is equivalent to say that any (qualitative) LT property satisfied by \mathcal{T} is also satisfied by $\bar{\mathcal{T}}$. Since every trajectory of SSs occurs only with a certain probability, it might have been an interesting idea to say that $\bar{\mathcal{T}}$ has a richer set of behaviors if it gives higher probability to any trajectory relatively to \mathcal{T} . This is however impossible unless two systems give exactly same probabilities to every trajectory: indeed, if we look at sets of trajectories instead and $H \subseteq Y^{\mathbb{N}}$ appears in \mathcal{T} and $\bar{\mathcal{T}}$ with probabilities p and $\bar{p} > p$ respectively, then necessary H^c appears with probabilities $(1 - p)$ and $(1 - \bar{p}) < (1 - p)$. The latter fact is a consequence of the full probability being exactly 1, so to cope with this issue one may also propose to use substochastic measures: those whose full mass is possibly less than 1. This way was chosen by [41], but it is by no means natural and does not lead to any comprehensive set of results that would help to solve problems that we are interested in. The trick is to understand that not a single trajectory is an element of behavior of an SSs, but a distribution over the trajectories is. For example, even though autonomous SSs (i.e. Markov Chains) may seem similar to TSs in that in both cases one has

a choice over the successor states – non-deterministic in the former setting and probabilistic in the latter – this similarity is misleading, and it makes little sense to say that one MC behaviorally includes another one, as each MC only has one behavior.

2.5 Comments on models and problem formulation

The exposition of the model in this work is rather standard and is similar to that in [70, Section 2.2]. However the present model is more general: for example we allow for a feasibility set K that is analytic, and for universally measurable policies. It can be shown that whenever the initial distribution $\alpha \in \mathcal{P}(X)$ is fixed, for a large class of performance criteria including all expected bounded utility cases it is sufficient to consider only analytically measurable deterministic policies depending exclusively on state coordinates of the history [24]. Moreover, one can sufficiently deal with Borel measurable policies, provided they do exist. However, if one is interested in finding a policy that is optimal or ε -optimal for any initial distribution, it is more convenient to deal with the class of universally measurable policies: the latter is rich enough to assure the existence of policies for many interesting problems – see e.g. the discussion in [20, Section 1.2]. This class also possesses some nice closure properties in contrast to the class of analytically measurable policies: e.g. the composition of two universally measurable functions is again universally measurable, but the composition of analytically measurable functions may not be analytically measurable. Such closure properties are important to ensure the appropriate measurability of the performance criterion with respect to the initial state. More details on this topic can be found in [117].

It is worth mentioning that there is an alternative approach to sequential decision making in a stochastic environment, which is known as *gambling* [48]. The difference with the MDP is mainly conceptual: if the current state is x , instead of first making a choice of a control action u and drawing a new state according to the distribution $\mathfrak{T}(x, u)$, in gambling the agent is allowed to choose the distribution of the new state directly, from the set of available *gambles* $\Gamma|_x$ ¹⁷. The set $\Gamma \subseteq X \times \mathcal{P}(X)$ is called the *gambling house*. On the methodological level, the difference between the MDP and gambling is that the latter extensively uses stopping time-like methods to derive most of the results, whereas the former is more focused on techniques based on DP. Finally, the difference between MDP and gambling models is also technical. First of all, initially the research on gambling theory has been done in the framework of finitely-additive probability measures [48]. Later, gambling models have also been considered in the σ -additive framework, which made it possible to compare them with MDP: for example, [24] showed the equivalence between some classes of MDP and gambling models – this result also holds for

¹⁷ Note that in MDP the choice of the distribution of the successor state is “labelled” by actions, whereas in gambling models such choice is unlabelled. One may think of this being similar to internal and external non-determinism in probabilistic automata [114], however there is no semantic difference between MDP and gambling models, and in both cases non-determinism can be considered both as an internal one or as an external one.

the MDP model we consider in this thesis. Further gambling models have been used more recently, e.g. in [94] and [96].

Research on gambling has broadly looked into the optimization of probabilities of given events. For example, [94] has obtained results for safety properties (that are clearly also applicable to the reachability), and [94, 96] has characterized the repeated reachability property. Due to this reason, although we do not use the gambling framework explicitly, sometimes we recall the results obtained there. For example, using the MDP framework for reachability properties seems more beneficial, however we mostly use the results of gambling for the repeated reachability. Another important point is that [95, Chapter 6] proposes an idea to optimize the probabilities of events, which is an alternative to the one we convey here. More precisely, it is shown that in the case of a countable state space the functional \mathfrak{M}^* possesses some useful properties of capacities [38]. In particular, [95, Theorem (1.2), Chapter 6] claims that for any state $x \in X$ and any event $A \in \mathcal{B}(H)$ it holds that

$$\mathcal{S}_x(A) = \inf \{ \mathcal{S}_x(B) : B \text{ is open and } B \supseteq A \}. \quad (2.11)$$

Furthermore, \mathfrak{M}^* for open events can be obtained by means of stopping times – see [95, Chapter 6] for more details. This result may be extendable to the more general case we deal with, where X is uncountable and one is interested only in events that can be described using some finite alphabet Σ . Unfortunately (2.11) does not provide a direct and explicit way to compute quantities of interest, or to derive optimal policies, so we do not pursue such direction here, preferring instead more explicit methods based on LTL formulae and automata theory.

The problem of optimizing the probability of a given event (or a property) is a problem that often appears in computer science, see e.g. a wide range of examples described in [17, Section 10.6]. Using LTL and automata theory for finite state-space MDP has a long history, part of which can be consulted in [17, Section 10.8]. However, extensions to the general state-space case have only appeared recently: [7] has provided an extension to the uncontrolled case (where trivially $U = \{u\}$ is a singleton), whereas [78] and [137] worked out the controlled case¹⁸.

¹⁸ The difference between the approaches in these two works is that [78] has allowed for Markov policies only, but clearly the policies over the composed system may depend on the state of the transition system: the map \mathfrak{P} can map Markov policies to history-dependent ones. To cope with this issue, extended Markov policies have been proposed in [78], namely policies that can depend also on an additional historical variable – the state of the transition system, which is a deterministic function of the MDP state history.

Finite-horizon case

This chapter formally develops ideas of precise and stochastic bisimulation for stochastic systems and introduces main results for the finite-horizon properties.

3.1 Approximate simulation of TSs

In most of the cases solving verification or synthesis problems over the concrete TS is easier than checking behavioral inclusion or equivalence over two *given* TSs directly, which questions usefulness of the latter relations. That task of providing a finite system that behaviorally includes or is equivalent to a given one is even more complicated. To cope with this issue, a different way to compare TSs was developed on the level of transition relations, rather than on the level of behaviors. These new relations tend to be more conservative than behavioral ones, but are easier to check and serve as sufficient conditions for the latter to hold.

Definition 3.1 *The TS $\tilde{\mathcal{T}}$ simulates \mathcal{T} if there exists an l.t.r. $R \subseteq X \times \bar{X}$ such that*

1. *for any $(x, \bar{x}) \in R$ it holds that $L(x) = \bar{L}(\bar{x})$,*
2. *for any $(x, \bar{x}) \in R$ and $x' \in T|_x$ there exists $\bar{x}' \in \bar{T}|\bar{x}$ such that $(x', \bar{x}') \in R$.*

In such case we say that R is an abstraction relation from \mathcal{T} to $\tilde{\mathcal{T}}$ and write $\mathcal{T} \preceq \tilde{\mathcal{T}}$ ¹.

¹ Often the definition of TS requires specification of the set of initial conditions $X_0 \subseteq X$, and in the definition of simulation R is required to be such that $R|_x$ contains some element of \bar{X}_0 for each $x \in X_0$, so that R is not necessary an l.t.r. At the same time [128, Definition 4.7] requires existence of simulating abstract states also for $x \notin X_0$. In our case it is always assumed that $X_0 = X$ which does not have any significant effect on the results.

As we said, verifying simulation for two TSs or constructing a TS that simulates the given one is often much easier than solving analogous problems over the behavioral inclusion [64]. At the same time, simulation is a sufficient condition for behavioral inclusion: this is easy to prove using induction, see e.g. [128, Proposition 4.9]. We provide a proof here just to compare it later with a much more technically intricate proof of the analogous result for SSs.

Theorem 3.2 *If the TS $\bar{\mathcal{T}}$ simulates \mathcal{T} , then for any $x \in X$ there exists $\bar{x} \in \bar{X}$ such that $V_L(\sigma^T, x) \subseteq V_{\bar{L}}(\Sigma^{\bar{T}}, \bar{x})$. In particular, $\mathcal{T} \leq \bar{\mathcal{T}}$.*

Proof: Consider an arbitrary $x \in X$ and pick any $\bar{x} \in R|_x$. Let us show that for any $v \in V(\sigma^T, x)$ there exists $\bar{\sigma} \in \Sigma^{\bar{T}}$ such that $\bar{v} := \bar{V}(\bar{x}, \bar{\sigma})$ satisfies

$$(v_0, \dots, v_n) R^{n+1} (\bar{v}_0, \dots, \bar{v}_n) \quad \forall n \in \mathbb{N}. \quad (3.1)$$

For $n = 0$ this result is trivial, so suppose it holds for some $n \in \mathbb{N}$. In particular (3.1) implies that $v_n R \bar{v}_n$, so there exists $\bar{v}_{n+1} \in \bar{T}|_{\bar{v}_n}$ such that $v_{n+1} R \bar{v}_{n+1}$. The desired map $\bar{\sigma}_n$ is given by $\bar{\sigma}_n(\bar{v}_0, \dots, \bar{v}_n) := \bar{v}_{n+1}$ and in arbitrary way for all other arguments. To finish the proof just note that (3.1) implies that $\bar{L}(\bar{v}) = L(v)$, and since $v \in V(\sigma^T, x)$ is arbitrary, we obtain that $V_L(\sigma^T, x) \subseteq V_{\bar{L}}(\Sigma^{\bar{T}}, \bar{x})$. \square

Hence, if we would like to use $\bar{\mathcal{T}}$ to solve verification and synthesis problems over \mathcal{T} it is sufficient to have that $\mathcal{T} \preceq \bar{\mathcal{T}}$ and $\bar{\mathcal{T}} \preceq \mathcal{T}$. If the two latter conditions are satisfied we say that \mathcal{T} and $\bar{\mathcal{T}}$ are *simulation equivalent* and denote it by $\mathcal{T} \simeq \bar{\mathcal{T}}$. Clearly both the behavioral inclusion \leq and the simulation relation \preceq are preorders on TS_Y according to the terminology of the Appendix A.3. As a result, behavioral equivalence \approx and simulation equivalence \simeq are indeed equivalence relations, as their names suggest.

The approximation version of simulation is given as follows:

Definition 3.3 *The TS $\bar{\mathcal{T}}$ ε -simulates \mathcal{T} if there exists an l.t.r. $R \subseteq X \times \bar{X}$ such that*

1. *for any $(x, \bar{x}) \in R$ it holds that $d_Y(L(x), \bar{L}(\bar{x})) \leq \varepsilon$,*
2. *for any $(x, \bar{x}) \in R$ and $x' \in T|_x$ there exists $\bar{x}' \in \bar{T}|\bar{x}$ such that $(x', \bar{x}') \in R$.*

In such case we say that R is an ε -abstraction relation from \mathcal{T} to $\bar{\mathcal{T}}$ and write $\mathcal{T} \preceq_\varepsilon \bar{\mathcal{T}}$.

Again, $\preceq := (\preceq_\varepsilon)_{\varepsilon \in \mathbb{R}_+}$ is an ε -preorder on TS_Y ; its symmetrization we denote by $\simeq := (\simeq_\varepsilon)_{\varepsilon \in \mathbb{R}_+}$. From Proposition A.2 it follows that \approx and \simeq define pseudometrics on TS_Y : they were first considered in [64]. Relations \preceq_0 and \simeq_0 are precisely (exact) simulation and simulation equivalence defined above. As in the case of their exact counterparts, for each $\varepsilon \in \mathbb{R}_+$ it holds that $\mathcal{T} \preceq_\varepsilon \bar{\mathcal{T}}$ implies that $\mathcal{T} \leq_\varepsilon \bar{\mathcal{T}}$, which can also be easily proved by induction similar to [128, Proposition 9.4].

After the introduction of ε -relation \preceq , the main challenge in the area of infinite TSs with metric output spaces become the construction of a finite TS $\bar{\mathcal{T}}$ for a given

infinite \mathcal{T} such that $\mathcal{T} \simeq_\varepsilon \bar{\mathcal{T}}$ for ε small enough. One of the leading methods is to come up with relevant stability assumption on the continuous, control or hybrid system underlying \mathcal{T} and use these assumptions to partition the state space of such system in order to assure the desired result. An example of such assumption is provided by approximate simulation functions that were extensively used in [64].

3.2 Precise simulation of SSs

The notion of simulation for SSs provided in this section serves as a sufficient condition for the (exact) behavioral inclusion. The following lemma provides simpler sufficient conditions which are used in proofs below.

Lemma 3.4 *For every $\alpha \in \mathcal{P}(X)$ it holds that*

- i. $S(\Sigma^\Gamma, \alpha) \subseteq S(\Sigma^{\Gamma'}, \alpha)$ for any $\Gamma' \supseteq \Gamma$ and
- ii. $S(\Sigma^\Gamma, \alpha) \equiv_0 S(\Sigma^{\text{sco } \Gamma}, \alpha)$.

Proof: The proof of the statement [i] is trivial, so let us focus on the second statement. It follows from [97, Theorem 2.1] that $S(\Sigma^{\text{sco } \Gamma}, \alpha) = \text{sco } S(\Sigma^\Gamma, \alpha)$, and hence one is only left to apply Proposition A.7.[ii]. \square

To prove Lemma 3.4.[ii] we use a very important result concerning the randomized choice. Indeed, choosing measures from $\text{sco } \Gamma$ is equivalent to choosing measures from Γ at random, according to some distribution [97]. Note that [97, Theorem 2.1] used in the proof of Lemma 3.4.[ii] states that the family of strategic measures for a convex set Γ is convex itself, which in particular implies that any strategic measure generated in $\text{sco } \Gamma$ (using a randomized choice over Γ) is representable as a convex combination of those generated in Γ , and that nothing is gained by applying the randomized choice. Similar results were obtained for different kind of stochastic models as well – see [55] and references therein.

We are now ready to introduce simulation for SSs.

Definition 3.5 *The SS $\bar{\mathcal{S}}$ simulates \mathcal{S} if there exists an l.t.r. $R \in \mathcal{A}(X \times \bar{X})$ such that*

- 1. *for any $(x, \bar{x}) \in R$ it holds that $L(x) = \bar{L}(\bar{x})$,*
- 2. *for any $(x, \bar{x}) \in R$ and $\gamma \in \Gamma|_x$ there exists $\bar{\gamma} \in \bar{\Gamma}|_{\bar{x}}$ such that $\gamma R_* \bar{\gamma}$.*

In such case we say that R is an abstraction relation from \mathcal{S} to $\bar{\mathcal{S}}$ and write $\mathcal{S} \preceq \bar{\mathcal{S}}$. If in addition \mathcal{S} simulates $\bar{\mathcal{S}}$ via R^{-1} , we say that \mathcal{S} and $\bar{\mathcal{S}}$ are bisimilar and write $\mathcal{S} \sim \bar{\mathcal{S}}$.

Remark 3.6 Note that by Lemma B.9 the requirement for R to be an l.t.r. in Definition 3.5 is equivalent to a seemingly stronger condition: for any $\alpha \in \mathcal{P}(X)$ there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ satisfying $\alpha R_* \bar{\alpha}$.

Note also that if $\mathcal{S} \preceq \bar{\mathcal{S}}$ via R and $\bar{\mathcal{S}}'$ is an SS over the state space \bar{X} which agree with $\bar{\mathcal{S}}$ on $\text{proj}_{\bar{X}}(R)$, then $\mathcal{S} \preceq \bar{\mathcal{S}}'$ as well.

The next result shows that similarity is sufficient for the behavioral inclusion.

Theorem 3.7 If the SS $\bar{\mathcal{S}}$ simulates \mathcal{S} , then $\mathcal{S} \leq_0 \bar{\mathcal{S}}$.

Proof: Let $\alpha \in \mathcal{P}(X)$ be arbitrary. Remark 3.6 implies that there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ such that $\alpha R_* \bar{\alpha}$. We consider an arbitrary $p \in \mathcal{S}(\Sigma^\Gamma, \alpha)$ and claim that there exists $\bar{p} \in \mathcal{S}(\Sigma^{\text{sco}\bar{\Gamma}}, \bar{\alpha})$ such that $p \downharpoonright_n R_*^{n+1} \bar{p} \downharpoonright_n$ for all $n \in \mathbb{N}$. This is the main part of the proof and the most technical one; after establishing this result we show how it implies the statement of the theorem. Let $\sigma \in \Sigma^\Gamma$ be such that $p = P_\alpha^\sigma$; we construct $\bar{\sigma}$ by induction and then define $\bar{p} := \bar{P}_{\bar{\alpha}}^{\bar{\sigma}}$. The induction hypothesis is satisfied for $n = 0$ since $p \downharpoonright_0 = \alpha$ and $\bar{p} \downharpoonright_0 = \bar{\alpha}$. Suppose the induction hypothesis is satisfied for some $n \in \mathbb{N}$. For any $(\omega, \bar{\omega}) := (x_0, \dots, x_n, \bar{x}_0, \dots, \bar{x}_n)$ such that $\omega R^{n+1} \bar{\omega}$ there exists $\bar{\gamma}(\omega, \bar{\omega}) \in \Gamma|_{\bar{x}_n}$ such that $\sigma_n(\omega) R_* \bar{\gamma}(\omega, \bar{\omega})$. Let $G(\omega, \bar{\omega}) \in \mathfrak{C}(\sigma_n(\omega), \bar{\gamma}(\omega, \bar{\omega}))$ by any coupling satisfying $G(R|_{\omega, \bar{\omega}}) = 1$. In Lemma C.21 put

$$\begin{array}{lllll} \Omega := X^{n+1} & \Xi := X & \Phi := R^{n+1} & \mu := p \downharpoonright_n & \kappa := \sigma_n \\ \bar{\Omega} := \bar{X}^{n+1} & \bar{\Xi} := \bar{X} & \Psi := R & \bar{\mu} := \bar{p} \downharpoonright_n & \bar{\Upsilon} := \bar{X}^n \times \bar{\Gamma}. \end{array}$$

We obtain that there exists $\bar{\sigma}_{n+1}$ such that $\bar{\sigma}_{n+1}(\bar{\omega}) \in \text{sco } \bar{\Gamma}|_{x_n}$ for all $\bar{\omega} \in \bar{X}^{n+1}$ and such that $p \downharpoonright_{n+1} R_*^{n+2} \bar{p} \downharpoonright_{n+1}$ as desired.

To show the main result, let us fix $n \in \mathbb{N}$ and pick up $\mathbb{P} \in \mathfrak{C}(p \downharpoonright_n, \bar{p} \downharpoonright_n)$ such that $\mathbb{P}(R^{n+1}) = 1$. Since $\mathcal{S} \preceq \bar{\mathcal{S}}$, $(\omega, \bar{\omega}) \in R^{n+1}$ implies that $L(\omega) = \bar{L}(\bar{\omega})$. The latter statement can be equivalently written as $(L \times \bar{L})^{-1}(\Delta_{Y^{n+1}}) \supseteq R^{n+1}$, hence

$$(L \times \bar{L})_* \mathbb{P}(\Delta_{Y^{n+1}}) = \mathbb{P}((L \times \bar{L})^{-1}(\Delta_{Y^{n+1}})) \geq \mathbb{P}(R^{n+1}) = 1. \quad (3.2)$$

Since $(L \times \bar{L})_* \mathbb{P} \in \mathfrak{C}(L_*(p \downharpoonright_n), \bar{L}_*(\bar{p} \downharpoonright_n))$ by Lemma C.18, we obtain $L_*(p \downharpoonright_n) = \bar{L}_*(\bar{p} \downharpoonright_n)$. Due to the fact that $n \in \mathbb{N}$ is arbitrary, $L_* p = \bar{L}_* \bar{p}$. As a result, $S_L(\Sigma^\Gamma) \subseteq S_{\bar{L}}(\Sigma^{\text{sco}\bar{\Gamma}})$, which together with Lemma 3.4 yields the statement of the theorem. \square

Corollary 3.8 If SSs \mathcal{S} and $\bar{\mathcal{S}}$ are bisimilar, then $\mathcal{S} \approx_0 \bar{\mathcal{S}}$.

Proof: Let R be the corresponding abstraction relation from \mathcal{S} to $\bar{\mathcal{S}}$. Since $\mathcal{S} \preceq \bar{\mathcal{S}}$, it follows from the proof of Theorem 3.7 that any pair $(\alpha, \bar{\alpha}) \in R_*$ satisfies $S_L(\Sigma^\Gamma, \alpha) \sqsubseteq_0 S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})$. Similarly, since $\bar{\mathcal{S}} \preceq \mathcal{S}$ we obtain that any pair $(\bar{\alpha}, \alpha) \in R_*^{-1}$ satisfies $S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha}) \sqsubseteq_0 S_L(\Sigma^\Gamma, \alpha)$. As $(R^{-1})_* = (R_*)^{-1}$ (cf. Appendix B.2), we obtain that $(\alpha, \bar{\alpha}) \in R_*$ iff $(\bar{\alpha}, \alpha) \in R_*^{-1}$, so any such pair satisfies $S_L(\Sigma^\Gamma, \alpha) \sqsubseteq_0 S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})$. Since both R and R^{-1} are l.t.r., by Remark 3.6 there always exists $\bar{\alpha}(\alpha)$

for a given α ($\bar{\alpha}$) such that $\alpha R_* \bar{\alpha}$, hence all the conditions of behavioral inclusion are satisfied. \square

Let us show that systems in Example 2.12 are bisimilar, which would in turn imply their behavioral equivalence by Theorem 3.7.

Example 3.9 Consider a map $f : X \rightarrow \bar{X}$ given by $f(\{a_1, a_2\}) := \bar{a}$ and $f(b) := \bar{b}$. Let us show that \mathcal{S} and $\bar{\mathcal{S}}$ are bisimilar via $\text{Gr}(f)$. The graph of a map is always an l.t.r., and since f is surjective, so is $\text{Gr}(f)^{-1}$. Recall that $(x, \bar{x}) \in \text{Gr}(f)$ iff $\bar{x} = f(x)$, so the first conditions of bisimilarity reads as $\bar{L} \circ f = L$, which is satisfied in our case.

The second condition for $\mathcal{S} \preceq \bar{\mathcal{S}}$ via $\text{Gr}(f)$ reads as $\Gamma(x) \text{Gr}(f)_* \bar{\Gamma}(f(x))$, or equivalently $\bar{\Gamma}(f(x)) = f_* \Gamma(x)$, for each $x \in X$. It is easy to check that indeed $\bar{\Gamma}(f(x))$ is a push-forward of $\Gamma(x)$ along f , however to elucidate the connection between f_* and $\text{Gr}(f)_*$, and in particular to elucidate Example B.2, let us provide explicit coupling measures $G(x, \bar{x}) \in \mathcal{P}(X \times \bar{X})$ of $\Gamma(x)$ and $\bar{\Gamma}(\bar{x})$ that are supported on $\text{Gr}(f)$ for every $x \in X$ and $\bar{x} \in \bar{X}$. These couplings are depicted on Figure 3.1.

For the converse direction, $\bar{\mathcal{S}} \preceq \mathcal{S}$ via $\text{Gr}(f)^{-1}$, note that again the second condition is $\Gamma(x) \text{Gr}(f)_* \bar{\Gamma}(f(x))$ for each $x \in X$, hence it is obviously satisfied, and the same coupling measures $G(x, \bar{x})$ can be used.




Figure 3.1: Coupling measures for Example 3.9; $\text{Gr}(f)$ is depicted in dark.

Before establishing some properties of \preceq and \sim , let us introduce their relaxed versions which are also sufficient for the behavioral inclusion and equivalence respectively. For the SS \mathcal{S} let us denote its *convex hull* as another SS given by $\text{sco } \mathcal{S} := (X, \text{sco } \Gamma, Y, L)$. Here sco is applied to Γ section-wise, that is $\text{sco } \Gamma|_x := \text{sco}(\Gamma|_x)$. It follows from [98, Lemma 2.3] that $\text{sco } \Gamma$ is analytic whenever Γ is, so $\text{sco } \mathcal{S}$ is a well-defined SS.

Definition 3.10 The SS $\bar{\mathcal{S}}$ probabilistically simulates \mathcal{S} via R whenever $\mathcal{S} \preceq \text{sco } \bar{\mathcal{S}}$ via R ; in such case we write $\mathcal{S} \trianglelefteq \bar{\mathcal{S}}$ and say that R is a probabilistic abstraction relation from \mathcal{S} to $\bar{\mathcal{S}}$. If in addition $\bar{\mathcal{S}}$ probabilistically simulates \mathcal{S} via R^{-1} , we say that \mathcal{S} and $\bar{\mathcal{S}}$ are probabilistically bisimilar and write $\mathcal{S} = \bar{\mathcal{S}}$.

Note that $\mathcal{S} \preceq \bar{\mathcal{S}}$ ($\mathcal{S} \sim \bar{\mathcal{S}}$) implies that $\mathcal{S} \trianglelefteq \bar{\mathcal{S}}$ ($\mathcal{S} = \bar{\mathcal{S}}$), but not vice-versa.

Corollary 3.11 If $\mathcal{S} \trianglelefteq \bar{\mathcal{S}}$ ($\mathcal{S} = \bar{\mathcal{S}}$), then $\mathcal{S} \leq_0 \bar{\mathcal{S}}$ ($\mathcal{S} \approx_0 \bar{\mathcal{S}}$).

Proof: If $\mathcal{S} \sqsubseteq \bar{\mathcal{S}}$, Theorem 3.7 implies that for any $\alpha \in \mathcal{P}(X)$ there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ such that $S_L(\Sigma^\Gamma, \alpha) \sqsubseteq_0 S_{\bar{L}}(\Sigma^{\text{sco}\bar{\Gamma}}, \bar{\alpha})$. Since $S_{\bar{L}}(\Sigma^{\text{sco}\bar{\Gamma}}, \bar{\alpha}) \equiv_0 S_{\bar{L}}(\Sigma^{\bar{\Gamma}}, \bar{\alpha})$ by Lemma 3.4, we obtain that $\mathcal{S} \leq_0 \bar{\mathcal{S}}$. The proof that probabilistic bisimilarity of \mathcal{S} and $\bar{\mathcal{S}}$ implies $\mathcal{S} \approx_0 \bar{\mathcal{S}}$ is similar to that of Corollary 3.8. \square

Let us now formulate some properties for simulation and its probabilistic version.

Theorem 3.12 *The following statements hold:*

- i. if $\bar{\mathcal{S}} = (X, \bar{\Gamma}, Y, L)$ and $\Gamma \subseteq \bar{\Gamma}$ then $\mathcal{S} \preceq \bar{\mathcal{S}}$ and $\mathcal{S} \sqsubseteq \bar{\mathcal{S}}$ via Δ_X ;
- ii. if $\mathcal{S} \sqsubseteq \bar{\mathcal{S}}$ via R then $\text{sco } \mathcal{S} \preceq \text{sco } \bar{\mathcal{S}}$ via R ;
- iii. if $\mathcal{S} \preceq \bar{\mathcal{S}}$ ($\mathcal{S} \sqsubseteq \bar{\mathcal{S}}$) via R and $\bar{\mathcal{S}} \preceq \hat{\mathcal{S}}$ ($\bar{\mathcal{S}} \sqsubseteq \hat{\mathcal{S}}$) via \bar{R} , then $\mathcal{S} \preceq \hat{\mathcal{S}}$ ($\mathcal{S} \sqsubseteq \hat{\mathcal{S}}$) via $\bar{R} \circ R$;
- iv. \preceq and \sqsubseteq are preorders on SS_Y ;
- v. \sim and $=$ are equivalences on SS_Y ;

Proof: The proof is as follows:

- i. The proof is trivial.
- ii. R satisfies the first condition of similarity since only Γ is transformed into $\text{sco } \Gamma$, and other components of SSs are left unchanged. To prove that the second condition of similarity holds, consider $(x, \bar{x}) \in R$ and apply Lemma B.10 to $\Gamma|_x$ and $\text{sco } \bar{\Gamma}|_{\bar{x}} \supseteq R_*|_{(\Gamma|_x)}$ recalling that $(\text{sco})^2 = \text{sco}$ over analytic sets.
- iii. Let $\hat{\mathcal{S}} := (\hat{X}, \hat{\Gamma}, Y, \hat{L})$. We first consider the case of \preceq . For any $(x, \hat{x}) \in \bar{R} \circ R$ there exists $\bar{x} \in \bar{X}$ such that $xR\bar{x}$ and $\bar{x}\bar{R}\hat{x}$ by definition of composition of relations (A.4). The first condition of similarity now follows immediately, and for the second condition consider any $\gamma \in \Gamma|_x$. Since $\mathcal{S} \preceq \bar{\mathcal{S}}$ and $xR\bar{x}$, there exists $\bar{\gamma} \in \bar{\Gamma}|_{\bar{x}}$ such that $\gamma R_* \bar{\gamma}$. Similarly, since $\bar{\mathcal{S}} \preceq \hat{\mathcal{S}}$ and $\bar{x}\bar{R}\hat{x}$, there exists a measure $\hat{\gamma} \in \hat{\Gamma}|_{\hat{x}}$ satisfying $\bar{\gamma} \bar{R}_* \hat{\gamma}$. By Lemma B.6 we obtain that $\gamma(\bar{R} \circ R)_* \hat{\gamma}$.
 With focus on \sqsubseteq : if $\bar{\mathcal{S}} \sqsubseteq \hat{\mathcal{S}}$ via \bar{R} then $\text{sco } \bar{\mathcal{S}} \preceq \text{sco } \hat{\mathcal{S}}$ via \bar{R} by [ii]. Since $\mathcal{S} \preceq \text{sco } \bar{\mathcal{S}}$ via R , using the just shown property of \preceq , we obtain that $\mathcal{S} \preceq \text{sco } \hat{\mathcal{S}}$, or equivalently $\mathcal{S} \sqsubseteq \hat{\mathcal{S}}$, via $\bar{R} \circ R$ as desired.
- iv. \preceq and \sqsubseteq are reflexive by [i] and transitive by [iii].
- v. \sim and $=$ are reflexive by [i], symmetric by definition, and transitive by [iii].

\square

Results above suggest that simulations, bisimulations and their probabilistic counterparts are useful in solving verification problems, similar to the case of TSs. The

bisimulation relation is also useful for the synthesis problem: we discuss this procedure in the general setting of approximate bisimulations below in Section 3.3.

Let us provide some guidelines on how the theory above is applied in practice, and in particular how to check (bi-)similarity or to construct a (bi-)similar SS for a given one. Here we assume that \mathcal{S} represents a concrete system and $\bar{\mathcal{S}}$ is its abstraction. Perhaps, the most important case to consider concerns abstraction relation $R \in \mathcal{A}(X \times \bar{X})$ being a graph of some map $f \in \mathcal{B}(X, \bar{X})$. Namely, we say that f is an *abstraction map* from \mathcal{S} to $\bar{\mathcal{S}}$ if $\text{Gr}(f)$ is an abstraction relation from \mathcal{S} to $\bar{\mathcal{S}}$. Note that $\text{Gr}(f)$ is always an l.t.r. Furthermore the two conditions of similarity translate into two simpler conditions for f to be abstraction map from \mathcal{S} to $\bar{\mathcal{S}}$,

1. $\bar{L} \circ f = L$, that is $\bar{L}(f(x)) = L(x)$ for each $x \in X$;
2. $(f \parallel f_*)(\Gamma) \subseteq \bar{\Gamma}$, that is $f_*\gamma \in \bar{\Gamma}|_{f(x)}$ for each $x \in X$ and $\gamma \in \Gamma|_x$.

In fact, one can propose a version of Definition 3.5 based purely on abstraction maps rather than general abstraction relations. Such approach allows avoiding introduction of lifting of relations from states to measures described in the Appendix B.2, thus significantly simplifying the proofs and yet yielding similarity in a number cases where it intuitively should hold. However, although the two SSs in Example 2.12 are bisimilar, there does not exist an abstraction map from $\bar{\mathcal{S}}$ back to \mathcal{S} as the following example shows.

Example 3.13 Suppose that there exists an abstraction map from $\bar{\mathcal{S}}$ to \mathcal{S} in Example 2.12, The first condition of similarity implies that such map \bar{f} satisfies $\bar{f}(\bar{b}) = b$ and $\bar{f}(\bar{a}) \in \{a_1, a_2\}$, so there are only two choices for \bar{f} . For $i \in \{1, 2\}$ let \bar{f}_i denote the choices satisfying $\bar{f}_i(\bar{a}) = a_i$. We obtain that

$$(\bar{f}_i)_*\bar{\Gamma}(\bar{b}) = \{0, 0, 1\} = \Gamma(b) = \Gamma(\bar{f}_i(\bar{b}))$$

as desired for either choice, however

$$\begin{aligned} (\bar{f}_1)_*\bar{\Gamma}(\bar{a}) &= \left\{ \frac{2}{3}, 0, \frac{1}{3} \right\} \neq \Gamma(a_1) = \Gamma(\bar{f}_1(\bar{a})), \\ (\bar{f}_2)_*\bar{\Gamma}(\bar{a}) &= \left\{ 0, \frac{2}{3}, \frac{1}{3} \right\} \neq \Gamma(a_2) = \Gamma(\bar{f}_2(\bar{a})) \end{aligned}$$

so neither choice of \bar{f}_i satisfies the second condition of similarity.

Note that the second condition of similarity for abstraction maps suggests that there is the “best” choice of $\bar{\Gamma}$ that ensures similarity. More precisely:

Definition 3.14 An abstraction pair for \mathcal{S} is an arbitrary Borel space \bar{X} together with an arbitrary surjective map $f \in \mathcal{B}(X, \bar{X})$ consistent with L . The minimal abstraction of \mathcal{S} by an abstraction pair (\bar{X}, f) is an SS $\mathcal{S}_{(\bar{X}, f)} := (\bar{X}, \bar{\Gamma}, Y, \bar{L})$, where $\bar{L} \in \mathcal{B}(\bar{X}, Y)$ is a unique map satisfying $\bar{L} \circ f = L$, and $\bar{\Gamma} := (f \parallel f_*)(\Gamma)$.

Similarity conditions on abstraction maps imply that $\mathcal{S} \preceq \mathcal{S}_{(\bar{X}, f)}$ and that $\bar{L} = \bar{L}'$, $\bar{\Gamma} \subseteq \bar{\Gamma}'$ whenever $\mathcal{S} \preceq (\bar{X}, \bar{\Gamma}', Y, \bar{L}')$ via f , which justifies the name “minimal abstraction”. Let us comment on Definition 3.14, and show that assuming f to be surjective and consistent with L is necessary. If f is not consistent with L , then it is never an abstraction map since it violates the first condition of similarity. If f is consistent, then $\bar{L}(\bar{x}) = L(x)$ for any $x \in f^{-1}(\bar{x})$ and $\bar{x} \in \bar{X}$, which gives us an explicit expression for \bar{L} . Furthermore, $\bar{\Gamma} = (f \parallel f_*)(\Gamma)$ is analytic, and it is an l.t.r. since f is surjective. A more convenient explicit expression for sections of $\bar{\Gamma}$ is given by $\bar{\Gamma}|_{\bar{x}} = \bigcup_{x \in f^{-1}(\bar{x})} f_*(\Gamma|_x)$ for each $\bar{x} \in \bar{X}$.

One may try extending Definition 3.14 to non-surjective maps by saying that $\bar{\Gamma}$ coincides with $(f \parallel f_*)(\Gamma)$ only on $f(X)$, \bar{L} is uniquely determined on $f(X)$ and that $\bar{\Gamma}$ and \bar{L} are defined in an arbitrary way over $f(X)^c$. In such case the minimal abstraction would be unique unless f is not surjective, but this is not the main issue. Unfortunately, there may not exist an analytic $\bar{\Gamma}$ which coincides with $(f \parallel f_*)(\Gamma)$ on $f(X)$. For example, consider the case when $(f \parallel f_*)(\Gamma)$ is Borel and $f(X)$ is analytic but not Borel, then $A := \bar{\Gamma} \setminus (f \parallel f_*)(\Gamma)$ is analytic, hence so is $\text{proj}_{\bar{X}}(A)$. On the other hand, the latter set equals $\bar{X} \setminus f(X)$ which is strictly co-analytic; this leads to a contradiction.

The minimal abstraction is related to the abstraction by partitioning. Indeed, given an abstraction pair (\bar{X}, f) we can define a partition $X_f := (X_{\bar{x}})_{\bar{x} \in \bar{X}}$ with $X_{\bar{x}} := f^{-1}(\bar{x})$ for each $\bar{x} \in \bar{X}$. As a result, we can treat the abstract state space \bar{X} as an index set for a partition of the concrete state space X . The first condition of similarity for f implies that such partition is done “within labels”, that is L is constant on partition cells $X_{\bar{x}}$. Recall that f defines the equivalence relation $\mathcal{E}_f \in \mathcal{B}(X^2)$ such that $x \mathcal{E}_f x'$ iff $f(x) = f(x')$; then X_f is exactly a collection of equivalence classes for \mathcal{E}_f . One may also expect a converse result to hold true, and try constructing an abstraction pair given some equivalence relation $\mathcal{E} \in \mathcal{B}(X^2)$ satisfying $\mathcal{E} \subseteq L^{-1}(\Delta_Y)$, however this task cannot always be accomplished with our framework. More precisely, since f is a surjective map, $\bar{X} = X/\mathcal{E}$ is the quotient space, which may even fail to be analytic, let alone Borel [82]. Even if do not require f to be a surjective map hence relaxing the condition on \bar{X} to $\bar{X} \supseteq X/\mathcal{E}$, then we can take a Borel \bar{X} iff \mathcal{E} is a smooth equivalence [82]. A sufficient condition for smoothness is the existence of a Borel selector for \mathcal{E} , which in addition implies that X/\mathcal{E} is a Borel space so that a projection map $f : X \rightarrow X/\mathcal{E}$ can be chosen to be surjective. Sufficient conditions for existence of a Borel selector can be found in [146]. To summarize, unless the partition has countably many cells, it is not a trivial task to construct an abstraction pair and hence the minimal abstraction even in theory. Due to this reason, here we always assume that such a pair is given – in each practical case it is rather easy to find it as one rarely deals with Borel spaces whose structure is too complicated.

Let us now discuss bisimilarity in terms of abstraction maps. As above, suppose that \mathcal{S} represents the concrete system, and consider any abstract state space \bar{X} which has the same cardinality as X does. Recall that X and \bar{X} are Borel isomorphic, so let $f \in \mathcal{B}(X, \bar{X})$ be any Borel isomorphism. By checking the similarity conditions for abstraction maps we obtain that $\mathcal{S} \preceq \mathcal{S}_{(\bar{X}, f)}$ via f , and that

$\mathcal{S}_{(\bar{X},f)} \preceq \mathcal{S}$ via f^{-1} . Since f is bijective, $\text{Gr}(f^{-1}) = \text{Gr}(f)^{-1}$, hence $\mathcal{S} \sim \mathcal{S}_{(\bar{X},f)}$ via f . In particular, for any SS over an uncountable state space we can construct a bisimilar SS over $\bar{X} = [0, 1]^2$. Although one can even come up with a constructive version of f , in most of the cases it does not preserve the structure of interest such as continuity, hence not much useful in practice. As a result, the question whether an SS admits a bisimilar abstraction over the state space of the same cardinality, even if it of lesser dimension, has an affirmative answer, so one shall instead think of bisimilar abstractions with nice structure rather than over a “smaller” state space.

Due to the argument above, we study bisimilarity conditions a *given* abstraction map f has to satisfy, rather than looking for such maps for a given abstract state space \bar{X} . Suppose that $\mathcal{S} \preceq \bar{\mathcal{S}}$ via f , then $\bar{\mathcal{S}} \preceq \mathcal{S}$ via $\text{Gr}(f)^{-1}$ iff f is surjective and for each $\bar{x} \in \bar{X}$, $\bar{\gamma} \in \bar{\Gamma}_{\bar{x}}$ and $x \in f^{-1}(\bar{x})$ there exists $\gamma \in \Gamma_x$ such that $\bar{\gamma} = f_*\gamma$. As a result, if $\mathcal{S} \sim \bar{\mathcal{S}}$ via f , then necessarily $\bar{\mathcal{S}} = \mathcal{S}_{(\bar{X},f)}$ which emphasizes importance of the concept of the minimal abstraction and leads to the following result.

Theorem 3.15 *If SSs \mathcal{S} and $\bar{\mathcal{S}}$ are bisimilar via f , then $\bar{\mathcal{S}} = \mathcal{S}_{(\bar{X},f)}$. In addition, $\mathcal{S} \sim \mathcal{S}_{(\bar{X},f)}$ via f for an abstraction pair (\bar{X}, f) iff $f_*(\Gamma_x) = f_*(\Gamma_{x'})$ for each $(x, x') \in \mathcal{E}_f$.*

Proof: The proof follows from the discussion above the statement of the theorem. \square

Note that the condition of Theorem 3.15 reminds of those used as the definition of bisimulation for PTSs in [91]. It is likely that PTSs can be expressed as SSs and that in such case our version of bisimilarity generalizes the one used for PTSs. However, since PTSs were not given an explicit semantics, we cannot formally claim that such an expression is valid so we omit it here. Nevertheless, it is worth discussing the ideas used in the aforementioned literature, focusing on our setting of SSs. As we have observed above, each SS admits a bisimilar one over the state space of the same cardinality, so it is interesting to find an SS from the latter class with simpler structure. One way to do this is to consider an equivalence relation $\mathcal{E} \in \mathcal{B}(X^2)$ which satisfies the following property: if $x\mathcal{E}x'$ then for any $\gamma \in \Gamma_x$ there exists $\gamma' \in \Gamma_{x'}$ such that $\gamma(A) = \gamma'(A)$ for any \mathcal{E} -closed set $A \in \mathcal{B}(X)$. Clearly, if X/\mathcal{E} is a Borel space and f is a projection map, then $\mathcal{S} \sim \mathcal{S}_{(X/\mathcal{E},f)}$ by Theorem 3.15. Hence, it is of interest whether there does exist “the largest” \mathcal{E} of such kind: that would lead to “the smallest” bisimilar representation of \mathcal{S} . Unfortunately, as we have discussed above, verifying that X/\mathcal{E} is a Borel space is hard even in theory. Even if some logical characterization of \mathcal{E} would be available (such as in [43]), then it still may happen that X/\mathcal{E} is analytic³. Even though analytic spaces possess nice closure properties, our analysis of the interplay between simulation (bisimulation) and behavioral inclusion (equivalence) required some features of Borel spaces that analytic ones may fail to have, so introducing bisimulation from the relation on states does not seem to be worth dealing with

² Although perhaps an interesting fact, this shall come as no surprise since a similar result holds for the TSs, where also each bijection is an isomorphism as state spaces of TSs are arbitrary sets.

³ See e.g. the discussion in [40, Section 1] and also [40, Lemma 4.9].

SSs over analytic spaces. A possible alternative is to work with relations that are compact sets: the latter are closed under continuous images much as analytic sets are. In such case even the theory of relation lifting from the Appendix B.2 is likely to be simpler yield richer results. This may be an interesting direction to pursue, however the compactness assumption seems to be restrictive – in particular, compactness of Γ is a kind of continuity assumption on the dynamics of \mathcal{S} – so such investigation goes beyond the scope this work. Let us emphasize that the arguments above further clarify why do we prefer working with bisimulation defined via relations between states of two systems, rather than via an equivalence relation on the states of one (e.g. joint) system.

Let us also briefly comment on probabilistic (bi-)similarity via probabilistic abstraction maps. Since $\mathcal{S} \leq \bar{\mathcal{S}}$ via f iff $\mathcal{S} \preceq \text{sco } \bar{\mathcal{S}}$ via f , for f to be a probabilistic abstraction map from \mathcal{S} to $\bar{\mathcal{S}}$ the first condition is the same: $\bar{L} \circ f = L$ and the second takes form $(f \parallel f_*)(\Gamma) \subseteq \text{sco } \bar{\Gamma}$, or equivalently $\text{sco}(f \parallel f_*)(\Gamma) \subseteq \text{sco } \bar{\Gamma}$. There may not be the smallest $\bar{\Gamma}$ satisfying the latter condition unless some regularity assumptions are made⁴, however to study probabilistic bisimulation any $\bar{\Gamma}$ for which the equality holds suffice our purposes.

Definition 3.16 *The SS $\bar{\mathcal{S}}$ is a minimal probabilistic simulation of \mathcal{S} by an abstraction pair (\bar{X}, f) if $\bar{L} \in \mathcal{B}(\bar{X}, Y)$ is such that $\bar{L} \circ f = L$, and $\text{sco } \bar{\Gamma} = \text{sco}(f \parallel f_*)(\Gamma)$. In case $\bar{\Gamma} = \text{sco}(f \parallel f_*)(\Gamma)$, we say that $\bar{\mathcal{S}}$ is the largest minimal probabilistic simulation of \mathcal{S} by (\bar{X}, f) , and write $\bar{\mathcal{S}} = \mathcal{S}_{p(\bar{X}, f)}$.*

Note that any minimal probabilistic simulation $\bar{\mathcal{S}}$ probabilistically simulates \mathcal{S} by definition. Also at least one minimal probabilistic simulation does always exist for \mathcal{S} , e.g. $\bar{\mathcal{S}} = \mathcal{S}_{p(\bar{X}, f)}$. Furthermore, if $\bar{\mathcal{S}}' = (\bar{X}, \bar{\Gamma}', Y, \bar{L})$ is a minimal probabilistic simulation for \mathcal{S} via (\bar{X}, f) then necessarily $\bar{\Gamma}' \subseteq \bar{\Gamma}$ which justifies the name “largest minimal probabilistic simulation”. Recall that by Corollary B.11 the push-forward f_* and sco commute, so $\bar{\Gamma} = (f \parallel f_*)(\text{sco } \Gamma)$, however in case \bar{X} is finite it is likely to be simpler to characterize and compute a finite-dimensional convex set $\text{sco}(f \parallel f_*)(\Gamma)$, rather than $\text{sco } \Gamma$.

The next theorem provides necessary and sufficient conditions for probabilistic abstraction map to ensure probabilistic bisimilarity in terms of minimality.

Theorem 3.17 *If \mathcal{S} and $\bar{\mathcal{S}}$ are probabilistically bisimilar via f , then $\bar{\mathcal{S}}$ is a minimal probabilistic simulation of \mathcal{S} . In such case $\mathcal{S} = \bar{\mathcal{S}}'$ via f for each minimal probabilistic simulation $\bar{\mathcal{S}}'$. Furthermore, $\mathcal{S} = \mathcal{S}_{p(\bar{X}, f)}$ via f iff $\text{sco } f_*(\Gamma|_x) = \text{sco } f_*(\Gamma|_{x'})$ for each $(x, x') \in \mathcal{E}_f$.*

Proof: Recall that \mathcal{S} and $\bar{\mathcal{S}}$ are probabilistically bisimilar via f iff $\bar{\Gamma}|_{\bar{x}} \subseteq f_*(\text{sco } \Gamma|_x)$ for each $\bar{x} \in \bar{X}$ and $x \in f^{-1}(\bar{x})$, or equivalently $\text{sco } \bar{\Gamma}|_{\bar{x}} \subseteq \text{sco } f_*(\Gamma|_x)$, from which

⁴ If $\text{sco}(f \parallel f_*)(\Gamma)$ is a compact convex set, then the Choquet-Bishop-de Leeuw theorem [21] implies that $\text{sco}(f \parallel f_*)(\Gamma) = \text{sco } E$ where E is a set of extreme points of $\text{sco}(f \parallel f_*)(\Gamma)$. In particular, $E \subseteq \bar{\Gamma}$ whenever $\text{sco}(f \parallel f_*)(\Gamma) \subseteq \text{sco } \bar{\Gamma}$, so in such case E can be taken to be the *smallest* candidate for $\bar{\Gamma}$.

the rest of the proof follows. For example, to show the first part just note that since $\mathcal{S} \trianglelefteq \bar{\mathcal{S}}$ via f , $\text{sco } \bar{\Gamma}|_{\bar{x}} \supseteq \text{sco } f_*(\Gamma|_x)$ which leads to $\text{sco } \bar{\Gamma} = \text{sco}(f \parallel f_*)(\Gamma)$. \square

Similar to the case of exact bisimilarity, the last part of Theorem 3.17 can be used to define probabilistic bisimilarity through equivalence relations on states. It comes as no surprise that in general case one still faces the same problems related to quotients of Borel spaces, so we do not pursue this approach to probabilistic bisimilarity either.

We conclude the discussion on abstraction maps with an example of constructing a sound finite abstraction for a given SS \mathcal{S} . Consider an abstraction pair (\bar{X}, f) where $\bar{X} = [0; n]$ and $n \in \mathbb{N}$. For f to be consistent with L it is necessary that the cardinality of $L(X)$ is less or equal to $n + 1$, so we assume that \mathcal{S} satisfies this restriction. Note that although $\mathcal{S}_{(\bar{X}, f)}$ has a finite state space, it may not be a finite SS: in most of the cases if \mathcal{S} is infinite then so is $\bar{\Gamma} = (f \parallel f_*)(\Gamma)$. To cope with this issue, as a next step we look for a finite SS on \bar{X} which probabilistically simulates $\mathcal{S}_{(\bar{X}, f)}$. Recall that $\mathcal{S}_{(\bar{X}, f)} = \text{sco } \mathcal{S}_{(\bar{X}, f)}$, so if we find finite $\bar{\Gamma}'$ satisfying $\text{sco } \bar{\Gamma}' = \text{sco } \bar{\Gamma}$, then $\bar{\mathcal{S}}' := (\bar{X}, \bar{\Gamma}', Y, \bar{L})$ would be the desired abstraction since $\mathcal{S} \trianglelefteq \bar{\mathcal{S}}'$. One candidate for $\bar{\Gamma}'|_i$ is a set of extreme points of $\text{sco } \bar{\Gamma}|_i$ due to [21], however we may know little about the structure of $\text{sco } \bar{\Gamma}|_i$ so such a set may be hard to find. Needless to say, it may appear to be infinite as well. Due to this reason, instead of requiring $\bar{\mathcal{S}}'$ to be a minimal probabilistic simulation of $\mathcal{S}_{(\bar{X}, f)}$, we look for a satisfaction of a relaxed condition $\text{sco } \bar{\Gamma}' \supseteq \text{sco } \bar{\Gamma}$, where e.g. $\text{sco } \bar{\Gamma}'|_i$ is a convex polytope which covers $\text{sco } \bar{\Gamma}|_i$, so that $\bar{\Gamma}'|_i$ is a finite set of its vertices.

The construction of such polytopes in fact can be done without introducing $\bar{\Gamma}$ explicitly, and by using a little information about Γ . Define

$$l(i, j) := \inf_{\substack{\gamma \in \Gamma|_x \\ x \in f^{-1}(i)}} f_* \gamma(\{j\}), \quad u(i, j) := \sup_{\substack{\gamma \in \Gamma|_x \\ x \in f^{-1}(i)}} f_* \gamma(\{j\}) \quad i, j \in \bar{X},$$

and let $\bar{\Gamma}'|_i$ be the set of the vertices of the polytope $\theta_n \cap \prod_{j=0}^n [l(i, j), u(i, j)]$, where θ_n is the n -dimensional probability simplex as in (A.1). Such procedure have been proposed for finite autonomous SSs [53, 81] using *abstract Markov Chains (AMCs)*. Some details of the construction of $\bar{\Gamma}'$ for a finite autonomous \mathcal{S} can be found in aforementioned works, and we do not further elaborate on this for the general case as such abstraction is just sound and only guarantees behavioral inclusion. In fact, although defining $\bar{\Gamma}'$ via l and u is practically feasible, verifying properties over $\bar{\mathcal{S}}'$ is likely to lead to rather conservative result due to the conservative nature of l and u . Clearly, although $\mathcal{S} \trianglelefteq \bar{\mathcal{S}}'$, by no means we shall expect that $\mathcal{S} = \bar{\mathcal{S}}'$. Our main interest concerns approximate notions introduce below, which provide stronger guarantees. Still, the construction of $\bar{\mathcal{S}}'$ may come handy in some cases, so we elucidate it on Figure 3.2.




Figure 3.2: The scheme of the construction of a sound abstraction, and relations that hold between systems constructed on each step.

3.3 Approximate simulation of SSs

Above we have used the transitivity of \preccurlyeq to prove the transitivity of \sqsubseteq , taking advantage of the way the latter relation has been defined through sco and \preccurlyeq . Bearing this in mind, we define approximate simulation of SSs in a similar fashion, just instead of convexification we consider an inflation of an SS, which is defined as follows.

Definition 3.18 Given $\varepsilon \in \mathbb{R}_+$, the ε -inflation of \mathcal{S} is given by $\mathcal{S}^\varepsilon := (X, \Gamma^\varepsilon, Y, L)$ where Γ^ε is defined section-wise by $\Gamma^\varepsilon|_x := (\Gamma|_x)^\varepsilon$ for all $x \in X$.

Notice that Γ^ε is an image of the following set

$$\{(x, p, q) : x\Gamma p \text{ and } \|p - q\| \leq \varepsilon\} = (\Gamma \times \mathcal{P}(X)) \cap (X \times \Delta_{d_{\text{TV}}}^\varepsilon)$$

under proj_{02} . Since $d_{\text{TV}} \in \mathcal{B}(\mathcal{P}(X) \times \mathcal{P}(X))$ by Corollary C.14, $\Gamma^\varepsilon \in \mathcal{A}(X \times \mathcal{P}(X))$ and thus \mathcal{S}^ε is a well-defined SS. Let us show that inflation of SSs preserves similarity as much as sco does by Theorem 3.12.[ii].

Lemma 3.19 If $\mathcal{S} \preccurlyeq \bar{\mathcal{S}}$ via R , then $\mathcal{S}^\varepsilon \preccurlyeq \bar{\mathcal{S}}^\varepsilon$ via R for all $\varepsilon \in \mathbb{R}_+$.

Proof: Consider $(x, \bar{x}) \in R$ and $\gamma \in \Gamma^\varepsilon|_x$. There exist $\gamma' \in \Gamma|_x$ and $\bar{\gamma}' \in \bar{\Gamma}|_{\bar{x}}$ such that $\|\gamma - \gamma'\| \leq \varepsilon$ and $\gamma' R_* \bar{\gamma}'$, so by Lemma B.12 there exists $\bar{\gamma} \in \bar{\Gamma}^\varepsilon|_{\bar{x}}$ satisfying $\gamma R_* \bar{\gamma}$. \square

Before defining approximate simulation of SSs, let us first relate behaviors of \mathcal{S} and \mathcal{S}^ε . Unfortunately, it is often the case that the behavioral distance between the original SS and its inflation is maximal, and $d_\approx(\mathcal{S}, \mathcal{S}^\varepsilon) = 2$ regardless of how small $\varepsilon > 0$ is. This is due to the reason \approx compares systems on the infinite time horizon. The following example supports this statement.

Example 3.20 Consider a simple autonomous SS with only two states: $X = \{a, b\}$, $Y = X$ and $L = \text{id}_X$. Suppose that $\Gamma(a), \Gamma(b) \in (0, 1)^2$. It follows that the Markov Chain corresponding to this SS is ergodic; let $\tilde{\alpha} \in \mathcal{P}(X)$ denote its unique invariant distribution [100]. For any function $f : Y^2 \rightarrow \mathbb{R}$ define

$$A_f := \left\{ (y_0, y_1, \dots) : \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(y_i, y_{i+1}) = \int_{Y^2} f(y, y') \Gamma(dy'|y) \tilde{\alpha}(dy) \right\}.$$

It holds that $\mathbf{Q}_x(A_h) = 1$ for any initial distribution $\alpha \in \mathcal{P}(X)$. For \mathcal{S}^ε let σ be any stationary strategy satisfying $\sigma(x) \neq \Gamma(x)$ for some $x \in X$. The corresponding Markov Chain is also ergodic for ε small enough, so let $\tilde{\alpha}^\varepsilon$ be its invariant distribution. Define

$$A_f^\varepsilon := \left\{ (y_0, y_1, \dots) : \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(y_i, y_{i+1}) = \int_{Y^2} f(y, y') \sigma(dy'|y) \tilde{\alpha}^\varepsilon(dy) \right\}.$$

It holds that $\mathbf{Q}_x^\sigma(A_h^\varepsilon) = 1$ for any initial distribution $\alpha \in \mathcal{P}(X)$, so choosing f such that

$$\int_{Y^2} f(y, y') \Gamma(dy'|y) \tilde{\alpha}(dy) \neq \int_{Y^2} f(y, y') \tilde{\sigma}(dy'|y) \tilde{\alpha}^\varepsilon(dy)$$

we obtain that $\|\mathbf{Q}_\alpha - \mathbf{Q}_{\alpha'}^\varepsilon\| = 2$ for any $\alpha, \alpha' \in \mathcal{P}(X)$, hence $d_\approx(\mathcal{S}, \mathcal{S}^\varepsilon) = 2$.

Remark 3.21 Note that the set A_f over which the maximum is achieved for $d_{\text{TV}}(\mathbf{Q}_\alpha, \mathbf{Q}_{\alpha'}^\varepsilon)$ in Example 3.20 is unlikely to be ω -regular⁵. Thus one may hope that the distance between measures only over regular or ω -regular events may be less conservative. Unfortunately, this is not the case. By Lemma C.13 the variation distance between \mathcal{S}_α and $\mathcal{S}_{\alpha'}^\varepsilon$ over any algebra that generates $\mathcal{B}(Y^\mathbb{N})$ equals $d_{\text{TV}}(\mathcal{S}_\alpha, \mathcal{S}_{\alpha'}^\varepsilon)$, so in particular the distance over acceptance languages of BLTL, which form an algebra, coincides with the total variation. The same applies to the collection of regular and ω -regular subsets of $Y^\mathbb{N}$.

To obtain non-trivial and meaningful results we thus restrict our attention to the bounded time horizon case. For each $n \in \mathbb{N}$ define $S_L^n(\Sigma^\Gamma) := (S_L(\Sigma^\Gamma))|_n$ and $S_L^n(\Sigma^\Gamma, \alpha) := (S_L(\Sigma^\Gamma, \alpha))|_n$ to be families of bounded-horizon marginals of all observation measures and those that start at $\alpha \in \mathcal{P}(X)$ respectively. We say that $\bar{\mathcal{S}}$ behaviorally (ε, n) -includes \mathcal{S} whenever $S_L^n(\Sigma^\Gamma) \sqsubseteq_\varepsilon S_L^n(\Sigma^\Gamma)$; in such case we write $\mathcal{S} \leq_n^\varepsilon \bar{\mathcal{S}}$.

⁵ See e.g. the discussion in [31], which also provide a different version of a maximizing set.

Lemma 3.22 *For each $\alpha \in \mathcal{P}(X)$, $n \in \mathbb{N}$ and $\varepsilon \in \mathbb{R}_+$ it holds that $S_L^n(\Sigma^\Gamma, \alpha) \equiv_{\varepsilon \otimes n} S_L^n(\Sigma^{\Gamma^\varepsilon})$, and in particular $\mathcal{S} \approx_{\varepsilon \otimes n, n} \mathcal{S}^\varepsilon$.*

Proof: Fix an arbitrary initial distribution $\alpha \in \mathcal{P}(X)$. Clearly, $S_L^n(\Sigma^\Gamma, \alpha) \subseteq S_L^n(\Sigma^{\Gamma^\varepsilon})$ so we only have to show that $S_L^n(\Sigma^{\Gamma^\varepsilon}) \leq_{\varepsilon_n} S_L^n(\Sigma^\Gamma, \alpha)$ for all $n \in \mathbb{N}$. Consider some $\sigma \in \Sigma^{\Gamma^\varepsilon}$ and denote $p := P_\alpha^\sigma$. Let us construct $\bar{\sigma} \in \Sigma^\Gamma$ such that $\bar{p} := P_\alpha^{\bar{\sigma}}$ satisfies $\|p|_n - \bar{p}|_n\| \leq \varepsilon_n$ for each $n \in \mathbb{N}$: that would imply the desired result. Let us prove the latter inequality by induction; since it trivially holds for $n = 0$, assume that it is true for some $n \in \mathbb{N}$. Denote $q := \frac{1}{2}(p|_n + \bar{p}|_n)$; there exists a kernel $\sigma'_n \in \mathcal{B}(X|X^{n+1})$ and a set $E_1 \in \mathcal{B}(X^{n+1})$ such that $\sigma_n(\omega) = \sigma'_n(\omega)$ for all $\omega \in E_1$ and such that $q(E_1) = 1$. By definition of Γ^ε and σ'_n , for any $\omega = (x_0, \dots, x_n) \in E_1$ there exists $\gamma(\omega) \in \Gamma|_{x_n}$ satisfying the inequality $\|\sigma'_n(\omega) - \gamma(\omega)\| \leq \varepsilon$. Let us show that such choice of γ can be done in a universally measurable way. Consider the set $J = J_1 \cap J_2 \subseteq X^{n+1} \times \mathcal{P}(X)$ we have to choose over. Here $J_1 = X^n \times \Gamma$ is analytic, and J_2 is an image of the Borel set

$$\{(\omega, \mu, \nu) : \mu = \sigma'_n(\omega) \text{ and } \|\mu - \nu\| \leq \varepsilon\} = (\text{Gr}(\sigma'_n) \times \mathcal{P}(X)) \cap (X^{n+1} \times \Delta_{d_{\text{TV}}}^\varepsilon)$$

under the projection $\text{proj}_{X^{n+1} \times \mathcal{P}(X)} : (\omega, \mu, \nu) \mapsto (\omega, \nu)$, hence analytic as well. As a result, Proposition C.17 applies and we can assume that $\gamma \in \mathcal{U}(X|X^{n+1})$.

Consider a kernel $\bar{\sigma}'_n \in \mathcal{B}(X|X^{n+1})$ and a set $E_2 \in \mathcal{B}(X^{n+1})$ such that $\bar{\sigma}'_n(\omega) = \gamma(\omega)$ for all $\omega \in E_2$ and such that $q(E_2) = 1$. Define $E = E_1 \cap E_2$ and put $\bar{\sigma}'_n(\omega) := \sigma'_n(\omega)$ for all $\omega \notin E$. Given that $\|p|_n - \bar{p}|_n\| \leq \varepsilon_n$, from [10, Lemma 2] it follows that

$$\|(p|_n \otimes \sigma'_n) - (\bar{p}|_n \otimes \bar{\sigma}'_n)\| \leq \varepsilon_{n+1}.$$

Since $q(E) = 1$ we obtain that $p|_n(E) = \bar{p}|_n(E) = 1$, so in particular $p|_n \otimes \sigma'_n = p|_{n+1}$. Finally, as $X^n \times \Gamma$ is an analytic l.t.r. it contains a graph of some universally measurable map. Let $\bar{\sigma}_n$ coincide with that map on E^c and with $\bar{\sigma}'_n$ on E , then $\bar{\sigma}_n \in \Sigma^\Gamma$ and satisfies $\|p|_{n+1} - (\bar{p}|_n \otimes \bar{\sigma}_n)\| \leq \varepsilon_{n+1}$ as desired. \square

We are now ready to define approximate similarity for SS and to show that it is a sufficient condition for approximate behavioral inclusion on the bounded time horizon.

Definition 3.23 *The SS $\bar{\mathcal{S}}$ ε -simulates \mathcal{S} if there exists an l.t.r. $R \in \mathcal{A}(X \times \bar{X})$ such that*

1. *for any $x \in X$ there exists $\bar{x} \in \bar{X}$ satisfying $(x, \bar{x}) \in R$,*
2. *for any $(x, \bar{x}) \in R$ it holds that $L(x) = \bar{L}(\bar{x})$,*
3. *for any $(x, \bar{x}) \in R$ and $\gamma \in \Gamma|_x$ there exists $\bar{\gamma} \in \bar{\Gamma}^\varepsilon|_{\bar{x}}$ such that $\gamma R_* \bar{\gamma}$.*

In such case we say that R is an ε -abstraction relation from \mathcal{S} to $\bar{\mathcal{S}}$ and write $\mathcal{S} \preceq_\varepsilon \bar{\mathcal{S}}$. If in addition \mathcal{S} ε -simulates $\bar{\mathcal{S}}$ via R^{-1} , we say that \mathcal{S} and $\bar{\mathcal{S}}$ are ε -bisimilar and write $\mathcal{S} \simeq_\varepsilon \bar{\mathcal{S}}$.

Notice that $\mathcal{S} \preceq_\varepsilon \bar{\mathcal{S}}$ iff $\mathcal{S} \preccurlyeq \bar{\mathcal{S}}^\varepsilon$; this is useful in proofs of the following results.

Theorem 3.24 *If the SS $\bar{\mathcal{S}}$ ε -simulates \mathcal{S} , then for any $\alpha \in \mathcal{P}(X)$ there exists $\bar{\alpha} \in \bar{\mathcal{P}}(X)$ such that $S_L^n(\Sigma^\Gamma, \alpha) \sqsubseteq_{\varepsilon, n} S_L^n(\Sigma^{\bar{\Gamma}}, \bar{\alpha})$ for all $n \in \mathbb{N}$ and in particular, $\mathcal{S} \leq_{\varepsilon, n} \bar{\mathcal{S}}$ where*

$$\varepsilon_n := 2(1 - (1 - \varepsilon)^n).$$

Proof: Fix an arbitrary initial distribution $\alpha \in \mathcal{P}(X)$. Since $\mathcal{S} \preccurlyeq \bar{\mathcal{S}}^\varepsilon$, by Theorem 3.7 there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ such that $S_L(\Sigma^\Gamma, \alpha) \sqsubseteq_0 S_L(\Sigma^{\bar{\Gamma}^\varepsilon}, \bar{\alpha})$, so in particular it follows that $S_L^n(\Sigma^\Gamma, \alpha) \sqsubseteq_0 S_L^n(\Sigma^{\bar{\Gamma}^\varepsilon}, \bar{\alpha})$ for all $n \in \mathbb{N}$. To finish the proof of the theorem one is only left of us the triangularity of \sqsubseteq and the result of Lemma 3.22. \square

Let us elaborate on the properties of \preceq and \simeq .

Theorem 3.25 *The following statements hold:*

- i. \preceq_0 is \preccurlyeq and \simeq_0 is \sim ;
- ii. if $\mathcal{S} \preceq_\varepsilon \bar{\mathcal{S}}$ via R and $\bar{\mathcal{S}} \preceq_\varepsilon \hat{\mathcal{S}}$ via \bar{R} , then $\mathcal{S} \preceq_{\varepsilon+\bar{\varepsilon}} \hat{\mathcal{S}}$ via $\bar{R} \circ R$;
- iii. \preceq (\simeq) is an ε -preorder (ε -equivalence) on SS_Y .

Proof: The proof is as follows:

- i. This fact immediately follows from the definition.
- ii. Lemma 3.19 implies that $\bar{\mathcal{S}}^\varepsilon \preccurlyeq (\hat{\mathcal{Q}}^\varepsilon)^\varepsilon$ via \bar{R} , and since $(\hat{\Gamma}^\varepsilon)^\varepsilon \subseteq \hat{\Gamma}^{\varepsilon+\bar{\varepsilon}}$ we obtain that $\bar{\mathcal{S}}^\varepsilon \preccurlyeq \hat{\mathcal{Q}}^{\varepsilon+\bar{\varepsilon}}$ by Theorem 3.12.[i]. The result follows from Theorem 3.12.[iii].
- iii. This result follows from [ii].

\square

Remark 3.26 *One can also define probabilistic ε -simulation by saying that $\bar{\mathcal{S}}$ probabilistically ε -simulates \mathcal{S} if $\mathcal{Q} \trianglelefteq \bar{\mathcal{S}}^\varepsilon$. Unfortunately, such a relation is not triangular since sco and ε -inflations do not commute. More precisely, $\text{sco}(\Gamma^\varepsilon) \subset (\text{sco } \Gamma)^\varepsilon$ and there exist cases when such inclusion is strict. Perhaps, for a metric different from TV such operations do commute, however this goes beyond the scope of the thesis.*

The examples above show how to use ε -bisimulation to obtain a solution of the verification problem for the concrete SS by solving the same problem over its ε -bisimilar abstraction. Let us now discuss the synthesis problem. Suppose that over \mathcal{S} we are given an initial distribution $\alpha \in \mathcal{P}(X)$ and an LT specification $h \in \mathcal{H}_1^n(Y)$ to maximize. Since there may not be an optimal strategy, let us also fix some precision level ϵ . We would like to find a strategy $\sigma \in \Sigma^\Gamma$ such that $\mathcal{Q}_\alpha^\sigma[h] \geq$

$\mathcal{S}_\alpha(h) - \epsilon$. How this can be done using the relations introduced above. Let us show that even in case of an exact bisimulation it is not always an easy task.

If $\mathcal{S} \sim \bar{\mathcal{S}}$ via some relation R , then there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ such that $\mathcal{S}_\alpha(h) = \bar{\mathcal{S}}_{\bar{\alpha}}(h)$. Let $\bar{\sigma} \in \Sigma^\Gamma$ be any strategy satisfying $\bar{Q}_{\bar{\alpha}}^{\bar{\sigma}}[h] \geq \bar{\mathcal{S}}_{\bar{\alpha}}(h) - \epsilon/2$. Since in particular $\bar{\mathcal{S}} \preceq \mathcal{S}$, it follows from the proof of Theorem 3.7 that there exist $\sigma' \in \Sigma^{\text{sco}\Gamma}$ such that $Q_{\alpha}^{\sigma'} = \bar{Q}_{\bar{\alpha}}^{\bar{\sigma}}$. Since $S(\Sigma^{\text{sco}\Gamma}, \alpha) \equiv_0 S(\Sigma^\Gamma, \alpha)$ by Lemma 3.4, there exists a strategy $\sigma \in \Sigma^\Gamma$ such that $Q_{\alpha}^{\sigma}[h] \geq Q_{\alpha}^{\sigma'}[h] + \epsilon/2$, and combining all the inequalities: $Q_{\alpha}^{\sigma}[h] \geq \mathcal{S}_\alpha(h) - \epsilon$. Note that to construct σ from $\alpha, \bar{\alpha}$ and $\bar{\sigma}$ we performed the following two steps:

1. constructed $\sigma' \in \Sigma^{\text{sco}\Gamma}$ satisfying $Q_{\alpha}^{\sigma'} = \bar{Q}_{\bar{\alpha}}^{\bar{\sigma}}$ as in the proof of Theorem 3.7;
2. constructed $\sigma \in \Sigma^\Gamma$ satisfying $Q_{\alpha}^{\sigma}[h] \geq Q_{\alpha}^{\sigma'}[h] - \epsilon/2$.

The procedure in the first step is “constructive” in the sense that in the proof of Theorem 3.7 one inductively applies Lemma C.21 which provides an analytic expression for σ'_n . However, such construction is rather complicated and is unlikely to be applicable in practice, at least since as a part of its procedure it requires computing conditional distributions of abstract trajectories given concrete trajectories. Moreover, even if the abstract strategy $\bar{\sigma}$ is Markov, stationary or uniformly ϵ -optimal, the corresponding σ' may not satisfy any of these properties. The use of aforementioned conditional distributions may turn Markov or stationary strategy into a history-dependent one, and dependence on α may compromise uniform ϵ -optimality. The second step is even harder.

Perhaps, one can avoid decomposing the problem into two steps as above, and given $\bar{\sigma} \in \Sigma^{\bar{\Gamma}}$ directly construct $\sigma \in \Sigma^\Gamma$ satisfying $Q_{\alpha}^{\sigma}[h] \geq \bar{Q}_{\bar{\alpha}}^{\bar{\sigma}}[h] - \epsilon/2$ as such strategy does always exist. At the same time, we are not aware of any such direct method which would significantly differ from the procedure above, and it is unlikely that such a method would be of practical importance. Due to this reason, we focus again on abstraction maps and provide a neat construction of σ for a given $\bar{\sigma}$.

As a motivational example, suppose that $\mathcal{S} \preceq \bar{\mathcal{S}}$ via $\text{Gr}(f)^{-1}$ for some f . Given an abstract strategy $\bar{\sigma} \in \Sigma^{\bar{\Gamma}}$ we can “imitate” it over the concrete SS as follows. If the current state history is (x_0, \dots, x_n) , the corresponding abstract history is $f(x_0, \dots, x_n)$ and hence the distribution $\bar{\gamma} = \bar{\sigma}_n(f(x_0, \dots, x_n))$ is chosen over the abstraction. To match it, over the concrete system we choose any distribution $\gamma \in \Gamma|_{x_n}$ satisfying $\bar{\gamma} = f_*\gamma$: it always exists since $\bar{\mathcal{S}} \preceq \mathcal{S}$. Intuitively, such choice shall generate same observation measure over the concrete system as $\bar{\sigma}$ generates over the abstraction. Before we prove this, let us show that such choice can be made dependent only on x_n and $\bar{\gamma}$ in a formal and consistent way. We do this via a concept of a refinement map.

Definition 3.27 *If the SS \mathcal{S} ϵ -simulates $\bar{\mathcal{S}}$ via $\text{Gr}(f)^{-1}$ for some $f \in \mathcal{B}(X, \bar{X})$, we say that $\mathbf{r}_f : X \times \mathcal{P}(\bar{X}) \rightarrow \mathcal{P}(X)$ is an ϵ -refinement map for f if for each $x \in X$ and $\bar{\gamma} \in \bar{\Gamma}|_{f(x)}$ it holds that $\mathbf{r}_f(x, \bar{\gamma}) \in \Gamma|_x$ and there exists γ' such that $d_{\text{TV}}(\mathbf{r}_f(x, \bar{\gamma}), \gamma') \leq \epsilon$ and $f_*\gamma' = \bar{\gamma}$.*

The existence of an ε -refinement map follows from the definition of approximate simulation, however not any such map serves out purposes well. Since it is ought to be used in the construction of strategies, we want it to be at least universally measurable.

Lemma 3.28 *If the SS \mathcal{S} ε -simulates $\bar{\mathcal{S}}$ via $\text{Gr}(f)^{-1}$ for some $f \in \mathcal{B}(X, \bar{X})$, then there exists a universally measurable ε -refinement map.*

Proof: Conditions an ε -refinement map has to satisfy lead to the following choice set:

$$A = \{(x, \bar{\gamma}, \gamma', \gamma) : (x, \gamma) \in \Gamma, d_{\text{TV}}(\gamma, \gamma') \leq \varepsilon, \bar{\gamma} = f_* \gamma'\}$$

which is analytic as an intersection of analytic sets: Γ is analytic, d_{TV} and f_* are Borel maps. Define $A' := \text{proj}_{124}(A)$, so that \mathfrak{r}_f is an ε -measurable map iff $\mathfrak{r}_f(x, \bar{\gamma}) \in A'|_{(x, \bar{\gamma})}$ for each $x \in X$ and $\bar{\gamma} \in \bar{\Gamma}|_{f(x)}$. Since A' is a projection of an analytic map, it is analytic as well, so that Proposition C.17 applies. \square

Given an ε -refinement map $\mathfrak{r}_f \in \mathcal{U}(\mathcal{P}(X)|X \times \mathcal{P}(\bar{X}))$, we define its action on strategies $\mathfrak{R} : \Sigma^\Gamma \rightarrow \Sigma^\Gamma$ for each $x_0, \dots, x_n \in X$, $n \in \mathbb{N}$ and $\bar{\sigma} \in \Sigma^{\bar{\Gamma}}$ as follows:

$$(\mathfrak{R}_f \bar{\sigma})_n(x_0, \dots, x_n) := \mathfrak{r}_f(x_n, \bar{\sigma}_n(f(x_0, \dots, x_n))).$$

The following theorem characterizes the main property of the refinement maps.

Theorem 3.29 *Suppose that the SS \mathcal{S} ε -simulates $\bar{\mathcal{S}}$ via $\text{Gr}(f)^{-1}$ for some $f \in \mathcal{B}(X, \bar{X})$, and consider any ε -refinement map $\mathfrak{r}_f \in \mathcal{U}(\mathcal{P}(X)|X \times \mathcal{P}(\bar{X}))$. For each $\bar{\sigma} \in \Sigma^{\bar{\Gamma}}$, $\bar{\alpha} \in \mathcal{P}(\bar{X})$ and any $\alpha \in \mathcal{P}(X)$ such that $f_* \alpha = \bar{\alpha}$, the ε -refined strategy $\sigma = \mathfrak{R}_f \bar{\sigma} \in \Sigma^\Gamma$ satisfies*

$$\|(f_* P_\alpha^\sigma)_n - (\bar{P}_{\bar{\alpha}}^{\bar{\sigma}})_n\| \leq \varepsilon^{\otimes n}, \quad \forall n \in \mathbb{N}. \quad (3.3)$$

Proof: The idea of the proof is the following: since $\bar{\mathcal{S}} \preceq \mathcal{S}^\varepsilon$, we can construct a strategy $\sigma' \in \Sigma^{\Gamma^\varepsilon}$ such that $f_* P_\alpha^{\sigma'} = \bar{P}_{\bar{\alpha}}^{\bar{\sigma}}$. Although Lemma 3.22 further implies the existence of $\sigma'' \in \Sigma^\Gamma$ such that $\|(P_\alpha^{\sigma'})_n - (P_\alpha^{\sigma''})_n\| \leq \varepsilon^{\otimes n}$ for all $n \in \mathbb{N}$, the question is whether σ can be taken as an instance of such σ'' : in this case (3.3) would follow immediately.

To formalize the idea above, denote $p_n := P_\alpha^\sigma|_n$ and $\bar{p}_n := \bar{P}_{\bar{\alpha}}^{\bar{\sigma}}|_n$ for each $n \in \mathbb{N}$. Let us show that there exists a measure $p' \in \mathcal{P}(X^\mathbb{N})$ such that $p'_n := p'|_n$ satisfies $p'_n = f_* \bar{p}_n$ and $\|p_n - p'_n\| \leq \varepsilon^{\otimes n}$. The proof is done by induction: for $n = 0$ one just takes $\mathcal{P}'_0 = \alpha$. Suppose we have constructed p'_n with the desired properties for some $n \in \mathbb{N}$. There exists a set $E \in \mathcal{B}(X^{n+1})$ and a kernel $\sigma'_n \in \mathcal{B}(X|X^{n+1})$ such that $p'_n(E) = 1$ and such that $\sigma_n(\omega) = \sigma'_n(\omega)$ for all $\omega \in E$. From Lemma C.21 and the definition of an ε -refinement map it follows that there exists $\kappa \in \mathcal{U}(X|X^{n+1})$ such that $f_*(p'_n \otimes \kappa) = \bar{p}_{n+1}$ and such that $\|\kappa(\omega) - \sigma_n(\omega)\| \leq \varepsilon$ for all $\omega \in E$. Hence $p'_{n+1} := p'_n \otimes \kappa$ satisfies $f_* p'_{n+1} = \bar{p}_{n+1}$, and also by Lemma C.19 we obtain that $\|p'_{n+1} - p_{n+1}\| \leq \varepsilon^{\otimes(n+1)}$ as desired. \square

Corollary 3.30 *Suppose that the SSs \mathcal{S} and $\bar{\mathcal{S}}$ are ε -bisimilar via $\text{Gr}(f)$ for some $f \in \mathcal{B}(X, \bar{X})$, and consider some $h \in \mathcal{H}_1^n(Y)$ and $\alpha \in \mathcal{P}(X)$. If $\bar{\sigma} \in \Sigma^{\bar{\Gamma}}$ is $(\bar{\varepsilon}, f_*\alpha)$ -optimal for h , then $\mathfrak{R}_f \bar{\sigma}$ is $(\bar{\varepsilon} + 2\varepsilon^{\otimes n}, \alpha)$ -optimal for h . In particular, if $\bar{\sigma}$ is uniformly $\bar{\varepsilon}$ -optimal for h , then $\mathfrak{R}_f \bar{\sigma}$ is uniformly $(\bar{\varepsilon} + 2\varepsilon^{\otimes n}, \alpha)$ -optimal for h .*

Proof: Denote $\sigma := \mathfrak{R}_f \bar{\sigma}$. Since \mathcal{S} ε -simulates $\bar{\mathcal{S}}$ via $\text{Gr}(f)^{-1}$, Theorem 3.29 implies that

$$(\mathcal{Q}_\alpha^\sigma)|_n[h] \geq (\bar{\mathcal{Q}}_{f_*\alpha}^{\bar{\sigma}})|_n[h] - \varepsilon^{\otimes n} \geq \bar{\mathcal{S}}_{f_*\alpha}(h) - \bar{\varepsilon} - \varepsilon^{\otimes n} \geq \mathcal{S}_\alpha(h) - \bar{\varepsilon} - 2\varepsilon^{\otimes n},$$

so that σ is $(\bar{\varepsilon} + 2\varepsilon^{\otimes n}, \alpha)$ -optimal for h as desired. \square

In Section 3.2 we have seen that exact and probabilistic (bi-)simulations via abstraction maps, as opposed to general abstraction relations, possess some nice properties. Theorem 3.29 further shows that abstraction maps are also useful for the solution of the synthesis problem, whereas even exact bisimilarity via general abstraction relation only allows for a refinement procedure which is difficult to implement in practice. For this reason, let us discuss approximate (bi-)simulation in terms of abstraction maps.

One of the properties probabilistic (bi-)simulations via abstraction maps have is minimality. This property allows for necessary and sufficiency conditions for exact and probabilistic bisimilarity given only the abstraction pair (\bar{X}, f) , without $\bar{\Gamma}$ specified. Provided such conditions are satisfied, one can take any minimal $\bar{\Gamma}$, which is also unique for exact bisimulation. Clearly, the ε -inflation is neither idempotent, unlike sco , nor does it have an inverse operation. For example, f is an ε -abstraction map from \mathcal{S} to $\bar{\mathcal{S}}$ iff $(f \parallel f_*)(\Gamma) \subseteq \bar{\Gamma}^\varepsilon$, but it is by no means clear how to define minimality here in such a way that if \mathcal{S} and $\bar{\mathcal{S}}$ are ε -bisimilar via f , then $\bar{\mathcal{S}}$ is necessary minimal. Due to the reasons, for purposes of approximate bisimilarity we introduce the following object.

Definition 3.31 *An abstraction triple for \mathcal{S} is an abstraction pair (\bar{X}, f) together with an l.t.r. $D \in \mathcal{A}(X \times \mathcal{P}(X))$. The abstraction triple (\bar{X}, f, D) is said to be regular whenever $D|_x = D|_{x'}$ for each $(x, x') \in \mathcal{E}_f$. The candidate approximate abstraction of \mathcal{S} by the abstraction triple (\bar{X}, f, D) is an SS $\mathcal{S}_{(\bar{X}, f, D)} = (\bar{X}, \bar{\Gamma}, Y, \bar{L})$ where $\bar{L} \in \mathcal{B}(\bar{X}, Y)$ is a unique map satisfying $\bar{L} \circ f = L$, and $\bar{\Gamma} := (f \parallel f_*)(D)$.*

Note that the “discretization” set D in Definition 3.31 is not required to be a subset of the concrete stochastic relation Γ . The simplest case when it can be useful is given by autonomous SSs: often $\Gamma(x) \neq \Gamma(x')$ whenever $x \neq x'$, however these distributions may be close enough for us to want x to represent x' in an abstraction. Note also that the minimal abstraction is a special case of the candidate approximate abstraction given by the choice $D = \Gamma$. Notice that $\mathcal{S} \simeq \mathcal{S}_{(\bar{X}, f, D)}$ via f iff

$$(f \parallel f_*)(\Gamma) \subseteq ((f \parallel f_*)(D))^\varepsilon, \quad (f \parallel f_*)(D)|_{\bar{x}} \subseteq \bigcap_{x \in f^{-1}(\bar{x})} f_*(\Gamma^\varepsilon|_x) \quad \forall \bar{x} \in \bar{X}. \quad (3.4)$$

This formulation allows us providing sufficient conditions for ε -bisimilarity.

Theorem 3.32 *For any regular abstraction triple (\bar{X}, f, D) it holds that*

$$d_{\simeq}(\mathcal{S}, \mathcal{S}_{(\bar{X}, f, D)}) \leq \sup_{x \in X} d_H(D|_x, \Gamma|_x).$$

Proof: It is equivalent to show that if $d_H(D|_x, \Gamma|_x) \leq \varepsilon$ for each $x \in X$, then $\mathcal{S} \simeq^\varepsilon \mathcal{S}_{(\bar{X}, f, D)}$ via f . We have $f_*(D|_x) \subseteq (f_*(\Gamma|_x))^\varepsilon$ and $f_*(\Gamma|_x) \subseteq (f_*(D|_x))^\varepsilon$ for each $x \in X$, and in particular $f_*(\Gamma|_x) \subseteq (f_*(D|_x))^\varepsilon$, so $\mathcal{S} \preceq_\varepsilon \mathcal{S}_{(\bar{X}, f, D)}$ via f . To show the converse direction, note that $(f \parallel f_*)(D)|_{\bar{x}} = f_*(D|_x) \subseteq f_*(\Gamma^\varepsilon|_x)$ for every $\bar{x} \in \bar{X}$ and $x \in f^{-1}(\bar{x})$, hence (3.4) holds and $\mathcal{S} \simeq^\varepsilon \mathcal{S}_{(\bar{X}, f, D)}$ via f . \square

To give a more specific example, consider some $m, n \in \mathbb{N}$ and a pointed partition $(X_i, x_i)_{i \in [0; n]}$ of X consistent with L , together with a collection of distributions $(\gamma_{ij})_{i \in [0; n], j \in [0; m]}$ such that $\gamma_{ij} \in \Gamma|_{x_i}$ for all $i \in [0; n]$ and $j \in [0; m]$.

Corollary 3.33 *Consider an arbitrary metric d_X on X , and suppose that for all $i \in [0; n]$: $\Gamma \in \text{Lip}_{\beta_i}((X_i, d_X), (\mathcal{P}(X), d_{\text{TV}}))$ and $\Gamma|_{x_i} \subseteq \left(\bigcup_{j \in [0; m]} \gamma_{ij}\right)^{\varepsilon_i}$. Then*

$$d_{\simeq}(\mathcal{S}, \mathcal{S}_{(\bar{X}, f, D)}) \leq \max_{i \in [0; n]} (\varepsilon_i + \beta_i \cdot \text{diam}_{d_X}(X_i))$$

where $\bar{X} = [0; n]$, f is the natural projection and $D = \bigcup_{i \in [0; n]} X_i \times (\bigcup_{j \in [0; m]} \gamma_{ij})$.

Proof: The result follows directly from Theorem 3.32. \square

Let us also show that introduced notions of (bi-)simulation are compositional.

Definition 3.34 *The product of two SSs $\mathcal{S}_1 = (X_1, \Gamma_1, Y_1, L_1)$ and $\mathcal{S}_2 = (X_2, \Gamma_2, Y_2, L_2)$ is a SS $\mathcal{S}_1 \times \mathcal{S}_2$ given by $(X_1 \times X_2, \Gamma_1 \times \Gamma_2, Y_1 \times Y_2, L_1 \times L_2)$.*

Theorem 3.35 *Denote $R := R_1 \times R_2$. It holds that*

- i. if $\mathcal{S}_1 \preceq \bar{\mathcal{S}}_1$ via R_1 and $\mathcal{S}_2 \preceq \bar{\mathcal{S}}_2$ via R_2 , then $(\mathcal{S}_1 \times \mathcal{S}_2) \preceq (\bar{\mathcal{S}}_1 \times \bar{\mathcal{S}}_2)$ via R ;
- ii. if $\mathcal{S}_1 \trianglelefteq \bar{\mathcal{S}}_1$ via R_1 and $\mathcal{S}_2 \trianglelefteq \bar{\mathcal{S}}_2$ via R_2 , then $(\mathcal{S}_1 \times \mathcal{S}_2) \trianglelefteq (\bar{\mathcal{S}}_1 \times \bar{\mathcal{S}}_2)$ via R ;
- iii. if $\mathcal{S}_1 \preceq_{\varepsilon_1} \bar{\mathcal{S}}_1$ via R_1 and $\mathcal{S}_2 \preceq_{\varepsilon_2} \bar{\mathcal{S}}_2$ via R_2 , then $(\mathcal{S}_1 \times \mathcal{S}_2) \preceq_{\varepsilon_1 \otimes \varepsilon_2} (\bar{\mathcal{S}}_1 \times \bar{\mathcal{S}}_2)$ via R .

Proof: Recall that $(x_1, x_2, \bar{x}_1, \bar{x}_2) \in R$ iff $(x_1, \bar{x}_1) \in R_1$ and $(x_2, \bar{x}_2) \in R_2$. Thus, R is an l.t.r. whenever R_1 and R_2 are, and $L_1(x_1) = \bar{L}_1(\bar{x}_1)$, $L_2(x_2) = \bar{L}_2(\bar{x}_2)$ for all $(x_1, x_2, \bar{x}_1, \bar{x}_2) \in R$. Due to this reason, we only need to check the second condition of similarity in each of the cases. Fix $(x_1, x_2, \bar{x}_1, \bar{x}_2) \in R$ and denote $x := (x_1, x_2)$, $\bar{x} := (\bar{x}_1, \bar{x}_2)$, $\Gamma = \Gamma_1 \times \Gamma_2$ and $\bar{\Gamma} = \bar{\Gamma}_1 \times \bar{\Gamma}_2$. Consider any $\gamma \in \Gamma|_x$, that is $\gamma = (\gamma_1, \gamma_2)$ where $\gamma_i \in \Gamma_i|_{x_i}$ for $i = 1, 2$.

- i. There exist $\bar{\gamma}_i \in \bar{\Gamma}_i|_{\bar{x}_i}$ such that $\gamma_i(R_i)_*\bar{\gamma}_i$ for $i = 1, 2$, and hence $\gamma R_*\bar{\gamma}$ where $\bar{\gamma} = (\bar{\gamma}_1, \bar{\gamma}_2)$.
- ii. There exist $\bar{\gamma}_i \in \text{sco } \bar{\Gamma}_i|_{\bar{x}_i}$ such that $\gamma_i(R_i)_*\bar{\gamma}_i$ for $i = 1, 2$. Further, $(\bar{\gamma}_1, \bar{\gamma}_2) \in \text{sco}(\bar{\Gamma}_1|_{\bar{x}_1} \times \bar{\Gamma}_2|_{\bar{x}_2})$, and hence $\gamma R_*\bar{\gamma}$ where $\bar{\gamma} = (\bar{\gamma}_1, \bar{\gamma}_2)$.
- iii. There exist $\bar{\gamma}_i \in (\bar{\Gamma}_i|_{\bar{x}_i})^{\varepsilon_i}$ such that $\gamma_i(R_i)_*\bar{\gamma}_i$ for $i = 1, 2$. By Lemma C.19, $(\bar{\gamma}_1, \bar{\gamma}_2) \in (\bar{\Gamma}_1|_{\bar{x}_1} \times \bar{\Gamma}_2|_{\bar{x}_2})^{\varepsilon_1 \otimes \varepsilon_2}$.

□

Corollary 3.36 $d_{\simeq}(\mathcal{S}_1 \times \mathcal{S}_2, \bar{\mathcal{S}}_1 \times \bar{\mathcal{S}}_2) \leq d_{\simeq}(\mathcal{S}_1, \bar{\mathcal{S}}_1) \otimes d_{\simeq}(\mathcal{S}_2, \bar{\mathcal{S}}_2)$.

Proof: The proof follows immediately from Theorem 3.35. □

3.4 Theory for MDPs

Above we have developed a theory of exact and approximate (bi)simulations for SSs. As we have mentioned in the beginning of Section 2.3, the results we obtained are likely to be applicable also in other modelling frameworks that have similar syntax and semantics. To support this statement, here we provide an interpretation of the above theory over the MDPs, perhaps the most popular stochastic control model. We also hope that the MDP framework elucidates technical exposition of the theory for SSs. A particular version of the MDP model we consider here is inspired by [20] and have been previously studied in [137, 138].

Definition 3.37 A Markov Decision Process (MDP) is a tuple $\mathcal{M} = (X, U, K, \tau, Y, L)$ where X, U and Y are Borel spaces, $K \in \mathcal{A}(X \times U)$ is an l.t.r., $\tau \in \mathcal{B}(X|X \times U)$ and $L \in \mathcal{B}(X, Y)$. We say that X is the state space, U is the action space, K is the feasibility set, τ is a transition kernel, Y is the output space and L is the output map of \mathcal{M} . The MDP \mathcal{M} is said to be finite if the set K is finite, otherwise the MDP \mathcal{M} is called infinite. The set of all MDPs with the output space Y is denoted by MDP_Y . If $K|_x$ contains exactly one element for each $x \in X$, we denote it by $K(x)$, and $\tau(x, K(x))$ by $\tau(x)$, and say that the MDP \mathcal{M} is autonomous.

Unless the contrary is specified, we assume the MDP $\mathcal{M} = (X, U, K, \tau, Y, L)$ to be given and fixed. The evolution of \mathcal{M} shall be understood as follows: given a current state x_n we observe its output $y_n = L(x_n)$ and choose some feasible action $u_n \in K|_{x_n}$, the successor state is distributed according to $\tau(x_n, u_n)$. This is very similar to the case of SSs, the only difference is that instead of choosing the distribution of the successor state γ_n directly, we first choose the action u_n and then obtain γ_n as $\tau(x_n, u_n)$. Thus, one can think of the set U as an indexation or parametrization of the successor distributions. This is particularly important for infinite system: for them Γ is rarely given directly in practice, as it would require

specifying sets of probability measures over uncountable Borel spaces. The latter task is often done through parametrization in one way or another, so one can think of MDPs as practically useful representations of SSs. For finite models there is no such problem: there are only finitely many measures to index; in fact, finite MDPs are sometimes defined as SSs, see e.g. [89, Definition 1].

The fact that the choice of the successor distribution is done directly in SSs and in a parameterized way in MDPs constitute to the main difference between them⁶. Following the previous paragraph, it is thus natural to define for a each MDP \mathcal{M} the corresponding SS as $\mathfrak{S}(\mathcal{M}) = (X, \Gamma, Y, L)$ where $\Gamma|_x := \bigcup_{u \in K|_x} \tau(x, u)$ for each $x \in X$. This SS is well-defined since

$$\Gamma := \text{proj}_{X \times \mathcal{P}(X)}(\text{Gr}(\tau) \cap (K \times \mathcal{P}(X)))$$

is an analytic set whenever K is. Note that the map $\mathfrak{S} : \text{MDP}_Y \rightarrow \text{SS}_Y$ is not injective, that is different MDPs can generate the same SSs. Yet, it is a surjective map since each SS $\mathcal{S} = (X, \Gamma, Y, L)$ can be also expressed as an MDP $\mathfrak{M}(\mathcal{S}) = (X, \mathcal{P}(X), \Gamma, \text{proj}_{\mathcal{P}(X)}, Y, L)$. The map $\mathfrak{M} : \text{SS}_Y \rightarrow \text{MDP}_Y$ is injective, but not surjective since its range does not contain MDPs with $U \neq \mathcal{P}(X)$; it further holds that $\mathfrak{S} \circ \mathfrak{M} = \text{id}_{\text{SS}_Y}$. We call \mathfrak{S} the embedding map; below it is used to pull back (bi)simulation relations on MDP_Y from SS_Y .

Similarly to TSs and SSs above, the LT semantics of MDPs is introduced by means of the sequential decision rules.

Definition 3.38 *A policy for the MDP \mathcal{M} is a sequential decision rule $\pi = (\pi_n)_{n \in \mathbb{N}}$, where the map $\pi_n \in \mathcal{U}(X_n, U)$ is such that $\pi_n(x_0, \dots, x_n) \in K|_{x_n}$ for each $x_i \in X$, $i \in [0; n]$ and $n \in \mathbb{N}$. The set of all such policies we denote by Π^K .*

We use the term “policy” for MDPs to distinguish them from strategies used for SSs: we comment more on this slightly later. By Proposition C.17 there does always exist at least one policy since K is an analytic set. Moreover, if \mathcal{M} is an autonomous MDP, then K is a graph of a Borel map⁷. Given any initial distribution $\alpha \in \mathcal{P}(X)$ and a policy $\pi \in \Pi^K$ there exists a unique probability measure $p \in \mathcal{P}(X^{\mathbb{N}})$ satisfying

$$p|_0 = \alpha, \quad \frac{d(p|_{n+1})}{d(p|_n)} = \tau(\text{proj}_n, \pi_n) \quad (p|_n \text{-a.s.}) \quad \forall n \in \mathbb{N}. \quad (3.5)$$

We denote this measure by P_α^π and also call it a strategic measure. The sets of all strategic measures and those that start at $\alpha \in \mathcal{P}(X)$ are denoted by $S(\Pi^K)$ and $S(\Pi^K, \alpha)$ respectively. Clearly, here $S(\Pi^K) := \bigcup_{\alpha \in \mathcal{P}(X)} S(\Pi^K, \alpha)$. To each strategic measure $P_\alpha^\pi \in S(\Pi^K)$ there corresponds an *observation measure* $Q_\alpha^\pi := L_* P_\alpha^\pi$. We further denote by $S_L(\Pi^K) := L_*(S(\Pi^K))$ and $S_L(\Pi^K, \alpha) := L_*(S(\Pi^K, \alpha))$ sets of

⁶ In fact, MDP and GM are equally expressible [136].

⁷ Due to this fact, autonomous MDPs are exactly the models known in the literature under the names general Markov Chains [112] or discrete-time Markov processes [134].

all observation measures, and those initiated at α respectively. We treat observation measures as behaviors of MDPs and introduce the following functional

$$\mathcal{M}_\alpha(h) := \sup_{q \in \mathcal{S}_L(\Pi^K, \alpha)} q[h], \quad \alpha \in \mathcal{P}(X), h \in \text{bU}(Y^\mathbb{N}).$$

For each $n \in \mathbb{N}$ we further define $\mathcal{S}_L^n(\Pi^K) := (\mathcal{S}_L(\Pi^K)) \downarrow_n$ and $\mathcal{S}_L^n(\Pi^K, \alpha) := (\mathcal{S}_L(\Pi^K, \alpha)) \downarrow_n$ to be families of bounded-horizon marginals of all observation measures and those that start at $\alpha \in \mathcal{P}(X)$ respectively.

Equipped with the notion of behaviors for MDPs, we can now define approximate behavioral inclusion and equivalence directly, as we have done for SSs in Section 2.4. We can also go further and introduce various (bi)simulation relations over MDPs, and then try to show that they imply their behavioral counterparts. Of course this does not make much sense as we have planned to use the results obtained above, rather than to redevelop them over the new modelling framework. For this purpose, below we define approximate behavioral inclusion, equivalence and exact, probabilistic or approximate (bi)simulation relations on MDP_Y by pulling them back from SS_Y along \mathfrak{T} ; in particular, all the properties of these relations are preserved (see Appendix A.4).

Let us emphasize *why* is that fine to pull relations back from SSs to MDPs, and why it would not work for TSs. For example, we can define another embedding map $\mathfrak{T} : \text{MDP}'_Y \rightarrow \text{TS}_Y$ by $\mathfrak{T}(\mathcal{M}) := (X, T, Y, L)$, where the transition relation is given by

$$T|_x = \{x' \in X : \tau(\{x'\}|x, u) > 0\} \quad x \in X,$$

where MDP'_Y is a collection of all finite MDPs⁸. As a result, the transition from x to x' is allowed in $\mathfrak{T}(\mathcal{M})$ iff it can happen with some non-zero probability in \mathcal{M} , that is $\mathfrak{T}(\mathcal{M})$ contains all “possible” transitions of \mathcal{M} . One can say that the map \mathfrak{T} “forgets” the probabilistic structure of MDPs⁹. Using the map \mathfrak{T} we can pull back e.g. exact behavioral inclusion and equivalence or exact (bi-)simulation from TS_Y to MDP'_Y and obtain well-defined relations over the latter modelling framework. Moreover, defined in such way (bi)simulation and behavioral inclusion (equivalence) would be consistent for finite MDPs, that is the former would be a stronger relation than the latter¹⁰. However, such behavioral equivalence would not be natural for MDPs and would not be strong enough in the following sense. Consider another MDP $\bar{\mathcal{M}} := (\bar{X}, \bar{U}, \bar{K}, \bar{\tau}, Y, \bar{L})$. In case \mathcal{M} and $\bar{\mathcal{M}}$ are finite we could say that they are behaviorally equivalent iff $\mathfrak{T}(\mathcal{M}) \approx_0 \mathfrak{T}(\bar{\mathcal{M}})$, however it may happen that for some $\alpha \in \mathcal{P}(X)$ there does not exist $\bar{\alpha} \in \mathcal{P}(\bar{X})$ such that $\mathcal{M}_\alpha = \bar{\mathcal{M}}_{\bar{\alpha}}$ as the next example shows.

Example 3.39 Suppose that \mathcal{M} and $\bar{\mathcal{M}}$ are finite autonomous MDPs given as follows: $X = \bar{X} := \{a, b, c\}$ with $U := \{0\}$, $K := X \times U$, $Y := \{A, B, C\}$ and $L(a) := A$, $L(b) := B$, $L(c) := C$. Suppose further that their transition kernels are given by $\tau(a) =$

⁸ We cannot define \mathfrak{T} as above for all MDPs since it may happen that T is not an l.t.r. if \mathcal{M} is infinite.

⁹ This statement can be made more formal by treating \mathfrak{T} as a functor between two categories, e.g. if one treats MDPs and TSs as coalgebras (see Appendix B), so that \mathfrak{T} would be a forgetful functor.

¹⁰ Such version of bisimulation for PTSs was introduced in [91].

$\{0, \frac{1}{3}, \frac{2}{3}\}$, $\bar{\tau}(a) = \{0, \frac{1}{2}, \frac{1}{2}\}$ and $\tau(b) = \bar{\tau}(b) = \{0, 1, 0\}$, $\tau(c) = \bar{\tau}(c) = \{0, 0, 1\}$. Note that $\mathfrak{T}(\mathcal{M}) = \mathfrak{T}(\bar{\mathcal{M}})$: all the systems are depicted on Figure 3.3. Consider h being an indicator of the event $\{y_0 = A, y_1 = C\}$ and let $\alpha = \delta(a)$, then $\mathcal{M}_\alpha(h) = \frac{2}{3}$, but $\bar{\mathcal{M}}_{\bar{\alpha}}(h) \leq \frac{1}{2}$ for all $\bar{\alpha} \in \mathcal{P}(X)$.




Figure 3.3: Example of two finite autonomous SSs.

To summarize, we could always pull back relations defined over one modelling framework to another as long as we have an embedding map from the latter to the former, however such induced relations may fail to be useful or natural. This issue can be caused e.g. by a “forgetting” structure of the embedding map used, as it happens for $\mathfrak{T} : \text{MDP}'_Y \rightarrow \text{TS}_Y$. Although the map $\mathfrak{S} : \text{MDP}_Y \rightarrow \text{SS}_Y$ has forgetting features as well since it eliminates action spaces and transition kernels from the model syntax, it still serves well our purposes due to the following fact:

$$\mathfrak{S}(\mathcal{M})_\alpha = \mathcal{M}_\alpha \quad \forall \alpha \in \mathcal{P}(X) \quad (3.6)$$

for any $\mathcal{M} \in \text{MDP}_Y$. For models with quantitative semantics, (3.6) can be thought of a sufficient condition on an embedding map to pull back behavioral relations to the desired ones. As a trivial example, \mathfrak{T} does not satisfy (3.6) even if one considers only deterministic initial conditions. The identity (3.6) says that essentially MDPs and SSs (or GMs) are equally expressive models: this have been first shown for some versions of these modelling frameworks in [24], and [136] extended this result to our current setting in particular. As a result, due to (3.6) we can pull back the relations defined for SSs to MDPs in such a way that they provide us desired behavioral properties. At the same time, since one of the goals of this section is to provide clarifying examples to the above theory, we give explicit definitions.

Definition 3.40 *The MDP $\bar{\mathcal{M}}$ behaviorally ε -includes \mathcal{M} if for any $\alpha \in \mathcal{P}(X)$ there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ such that $S_L(\Pi^K, \alpha) \sqsubseteq_\varepsilon S_{\bar{L}}(\Pi^K, \bar{\alpha})$, that is*

$$\mathcal{M}_\alpha(h) \leq \bar{\mathcal{M}}_{\bar{\alpha}}(h) + \varepsilon \quad \forall h \in \text{bU}_1(Y^{\mathbb{N}}).$$

In such case we write $\mathcal{M} \leq_\varepsilon \bar{\mathcal{M}}$. MDPs \mathcal{M} and $\bar{\mathcal{M}}$ are behaviorally ε -equivalent if for any initial distribution $\alpha \in \mathcal{P}(X)$ ($\bar{\alpha} \in \mathcal{P}(\bar{X})$) there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ ($\alpha \in \mathcal{P}(X)$) such that $S_L(\Pi^K, \alpha) \equiv_\varepsilon S_L(\Pi^{\bar{K}}, \bar{\alpha})$; in such case we write $\mathcal{M} \approx_\varepsilon \bar{\mathcal{M}}$.

The MDP $\bar{\mathcal{M}}$ behaviorally (ε, n) -includes \mathcal{M} whenever $S_L^n(\Pi^K) \sqsubseteq_\varepsilon S_L^n(\Pi^{\bar{K}})$; in such case we write $\mathcal{S} \leq_n^\varepsilon \bar{\mathcal{S}}$. MDPs \mathcal{M} and $\bar{\mathcal{M}}$ are behaviorally (n, ε) -equivalent if for any initial distribution $\alpha \in \mathcal{P}(X)$ ($\bar{\alpha} \in \mathcal{P}(\bar{X})$) there exists $\bar{\alpha} \in \mathcal{P}(\bar{X})$ ($\alpha \in \mathcal{P}(X)$) such that $S_L^n(\Pi^K, \alpha) \equiv_\varepsilon S_L^n(\Pi^{\bar{K}}, \bar{\alpha})$; in such case we write $\mathcal{M} \approx_n^\varepsilon \bar{\mathcal{M}}$.

As we have mentioned, by (3.6) approximate behavioral inclusion and equivalence for MDPs per Definition 3.40 are exactly pullbacks of their SS counterparts. Hence, $\mathcal{M} \leq_\varepsilon \bar{\mathcal{M}}$ ($\mathcal{M} \approx_\varepsilon \bar{\mathcal{M}}$) iff $\mathfrak{S}(\mathcal{M}) \leq_\varepsilon \mathfrak{S}(\bar{\mathcal{M}})$ ($\mathfrak{S}(\mathcal{M}) \approx_\varepsilon \mathfrak{S}(\bar{\mathcal{M}})$). Let us now proceed to the definition of (bi)simulations for MDPs.

Definition 3.41 Consider MDPs \mathcal{M} and $\bar{\mathcal{M}}$ and suppose there exists an l.t.r. $R \in \mathcal{A}(X \times \bar{X})$ such that $L(x) = \bar{L}(\bar{x})$ for all $(x, \bar{x}) \in R$.

- We say that $\bar{\mathcal{M}}$ simulates \mathcal{M} via R , and write $\mathcal{M} \preceq \bar{\mathcal{M}}$, if for any $(x, \bar{x}) \in R$ and $u \in K|_x$ there exists $\bar{u} \in \bar{K}|\bar{x}$ such that $\tau(x, u)R_*\bar{\tau}(\bar{x}, \bar{u})$. If in addition $\bar{\mathcal{M}} \preceq \mathcal{M}$ via R^{-1} , we say that \mathcal{M} and $\bar{\mathcal{M}}$ are bisimilar. In such case we write $\mathcal{M} \sim \bar{\mathcal{M}}$.
- We say that $\bar{\mathcal{M}}$ probabilistically simulates \mathcal{M} via R , and write $\mathcal{M} \trianglelefteq \bar{\mathcal{M}}$, if for any $(x, \bar{x}) \in R$ and $u \in K|_x$ there exists $\bar{v} \in \mathcal{P}(U)$ such that $\nu(K|_{\bar{x}}) = 1$ and such that $\tau(x, u)R_* \int \bar{\tau}(\bar{x}, \bar{u})\bar{\nu}(d\bar{u})$. If in addition $\bar{\mathcal{M}} \trianglelefteq \mathcal{M}$ via R^{-1} , we say that \mathcal{M} and $\bar{\mathcal{M}}$ are probabilistically bisimilar. In such case we write $\mathcal{M} = \bar{\mathcal{M}}$.
- We say that $\bar{\mathcal{M}}$ ε -simulates \mathcal{M} via R , and write $\mathcal{M} \preceq_\varepsilon \bar{\mathcal{M}}$, if for any $(x, \bar{x}) \in R$ and $u \in K|_x$ there exists $\bar{u} \in \bar{K}|\bar{x}$ and $\bar{\gamma} \in \mathcal{P}(\bar{X})$ such that $d_{TV}(\bar{\tau}(\bar{x}, \bar{u}), \bar{\gamma}) \leq \varepsilon$ and such that $\tau(x, u)R_*\bar{\gamma}$. If in addition $\bar{\mathcal{M}} \preceq_\varepsilon \mathcal{M}$ via R^{-1} , we say that \mathcal{M} and $\bar{\mathcal{M}}$ are ε -bisimilar. In such case we write $\mathcal{M} \simeq_\varepsilon \bar{\mathcal{M}}$.

Again, exact (probabilistic, approximate) simulation and bisimulation relations for MDPs are exactly the pullbacks of the corresponding relations over SSs. Since pullback of relations preserves properties of reflexivity, symmetry, transitivity, monotonicity and triangularity, we obtain the following result.

Theorem 3.42 Over MDP_Y it holds that

- i. \preceq and \trianglelefteq (\sim and $=$) imply \leq_0 (\approx_0);
- ii. \preceq and \trianglelefteq (\sim and $=$) are preorders (equivalences);
- iii. for each $\varepsilon \in \mathbb{R}_+$ and $n \in \mathbb{N}$, \preceq_ε (\simeq_ε) imply $\leq_{\varepsilon \otimes n, n}$ ($\approx_{\varepsilon \otimes n, n}$);
- iv. \preceq (\simeq) is an ε -preorder (ε -equivalence);

Recall that whenever two SSs \mathcal{S} and $\bar{\mathcal{S}}$ are ε -bisimilar, one can use $\bar{\mathcal{S}}$ to solve synthesis problems over \mathcal{S} by means of refinement maps (cf. Definition 3.27). A similar result applies to MDPs, however we cannot use refinement maps directly in this setting since any refinement map returns a successor distribution rather than an action. Although the former is sufficient to define a strategy over an SS, the latter is needed to construct a policy over an MDP. For this reason, we introduce a related concept of an interpolation map.

Definition 3.43 *If the MDP \mathcal{M} ε -simulates $\bar{\mathcal{M}}$ via $\text{Gr}(f)^{-1}$ for some $f \in \mathcal{B}(X, \bar{X})$, we say that $\mathbf{i}_f : X \times \bar{U} \rightarrow U$ is an ε -interpolation map for f if $\tau \circ (\text{proj}_X \times \mathbf{i}_f)$ is an ε -refinement map for f . That is, for each $x \in X$ and $\bar{u} \in \bar{K}|_{f(x)}$ it holds that $\mathbf{i}_f(x, \bar{u}) \in K|x$ and there exists γ' such that $d_{\text{TV}}(\tau(x, \mathbf{i}_f(x, \bar{u})), \gamma') \leq \varepsilon$ and $f_*\gamma' = \bar{\tau}(f(x), \bar{u})$.*

The next result shows that there always exists a measurable interpolation map.

Proposition 3.44 *If the MDP \mathcal{M} ε -simulates $\bar{\mathcal{M}}$ via $\text{Gr}(f)^{-1}$ for some $f \in \mathcal{B}(X, \bar{X})$, then there exists a universally measurable ε -interpolation map.*

Proof: Just notice that an ε -interpolation map has to choose over

$$A = \{(x, u, \gamma, \gamma', \bar{u}, \bar{\gamma}) : (x, u) \in K, \bar{\gamma} = \tau(f(x), \bar{u}), \bar{\gamma} = f_*\gamma', \gamma = \tau(x, \bar{u}), d_{\text{TV}}(\gamma, \gamma') \leq \varepsilon\}$$

which is an analytic set. \square

Finally, let us provide a procedure to build a finite ε -bisimilar abstraction for a given concrete system. With focus on the MDP \mathcal{M} consider some $m, n \in \mathbb{N}$ and a pointed partition $(X_i, x_i)_{i \in [0; n]}$ of X consistent with L . For each i pick control actions $(u_{ij})_{i \in [0; n], j \in [0; m]}$ such that $u_{ij} \in \Gamma|_{x_i}$ for all $i \in [0; n]$ and $j \in [0; m]$. The abstraction MDP is $\bar{\mathcal{M}} = ([0; n], [0; m], \bar{K}, \bar{\tau}, Y, \bar{L})$ where $\bar{K} := [0; n] \times [0; m]$, $\bar{\tau}(\{i'\}|i, j) := \tau(X_{i'}|x_i, u_{ij})$ and $\bar{L}(i) := L(x_i)$ for all $i, i' \in [0; n]$ and $j \in [0; m]$.

Proposition 3.45 *Consider arbitrary metrics d_X and d_U on X and U respectively, and let $d_{X \times U}$ be the product metric. Suppose that for all $i \in [0; n]$: $K \in \text{Lip}_{\beta_i}((X_i, d_X), (U, d_U))$, $K|x_i \subseteq \left(\bigcup_{j \in [0; m]} u_{ij}\right)^{\varepsilon_i}$ and $\tau \in \text{Lip}_{\rho_i}((X_i \times U, d_{X \times U}), (\mathcal{P}(X), d_{\text{TV}}))$. Then*

$$d_{\simeq}(\mathcal{M}, \bar{\mathcal{M}}) \leq \max_{i \in [0; n]} \rho_i \cdot (\varepsilon_i + (1 \vee \beta_i) \cdot \text{diam}_{d_X}(X_i)).$$

Proof: The result follows directly from Corollary 3.33. \square

Lipschitz continuity of the transition kernel τ in Proposition 3.45 can be established based on the following sufficient condition.

Proposition 3.46 *Suppose that τ admits an integral representation $\tau(x, u) = t(x, u, \tilde{x})\mu(\tilde{x})$ where $t \in \mathcal{B}(X \times U \times X)$ and μ is some positive σ -finite Borel measure on X . If*

$$|t(x, u, \tilde{x}) - t(x', u', \tilde{x})| \leq r_i(\tilde{x})d_{X \times U}((x, u), (x', u'))$$

for all $x, x' \in X_i, u \in U, \tilde{x} \in X$ and some $r_i \in \mathcal{B}(X)$ and metric $d_{X \times U}$ on $X \times U$, then $\tau \in \text{Lip}_{\rho_i}((X_i \times U, d_{X \times U}), (\mathcal{P}(X), d_{\text{TV}}))$ where $\rho_i = \int_X r_i(\tilde{x}) \mu(d\tilde{x})$.

Proof: For any $f \in \text{b}\mathcal{U}_1(X)$ we obtain that

$$\begin{aligned} |\tau(x, u)[f] - \tau(x', u')[f]| &= \left| \int_X f(\tilde{x}) (t(x, u, \tilde{x}) - t(x', u', \tilde{x})) \mu(d\tilde{x}) \right| \\ &\leq \int_X |f(\tilde{x})| \cdot |t(x, u, \tilde{x}) - t(x', u', \tilde{x})| \mu(d\tilde{x}) \\ &\leq \int_X r_i(\tilde{x}) d_{X \times U}((x, u), (x', u')) \mu(d\tilde{x}) \\ &\leq \rho_i \cdot d_{X \times U}((x, u), (x', u')), \end{aligned}$$

and since f is arbitrary, $d_{\text{TV}}(\tau(x, u), \tau(x', u')) \leq \rho_i \cdot d_{X \times U}((x, u), (x', u'))$. \square

3.5 Comments on approximate stochastic bisimulation

Some historical remarks regarding (bi-)simulation for TSs can be found in [17, Section 7.10], which in particular states that these notions are originated from [101]. Similarly, [17, Section 10.8] provides some bibliographical notes regarding (bi-)simulation for probabilistic systems. We discuss the latter topic in some detail here as well.

There are two approaches used when defining (bi-)simulations for different models. The first one is to define (bi-)simulation directly as relation between systems: we refer to this approach as *system-based*. This approach is the one used in our thesis, see e.g. Definition 3.5. Alternatively, one can define (bi-)simulation as a relation between states of a single system. In such case, one then says that systems A and B are (bi-)similar if their initial states are (bi-)similar as states over a “joint” system $A \cup B$, which is defined in some intuitive way. We call such approach here *state-based*, see e.g. [43] for clarifying examples. Within the state-based approach it is often required that (bi-)simulation is a preorder (equivalence) on states, which casts the corresponding relation between systems being a preorder (equivalence) as well.

Clearly, the system-based approach to the definition of (bi-)simulation is more general: given a state-based (bi-)simulation one can always define a corresponding relation between systems, by relating their initial state over a joint system. However, it may happen that some system-based version of (bi-)simulation does not allow for a suitable state-based analogue. For example, exact and probabilistic bisimulations for SSs from Definitions 3.5 and 3.16 have natural state-based versions discussed after Theorems 3.15 and 3.17, however if two SSs are bisimilar via some analytic relation R , there may not exist a corresponding Borel relation \mathcal{E} over a joint system. Yet, for some syntaxes of systems such as TSs and finite SSs the two approaches are equivalent.

The paper [90] was among the first works on equivalence relations for probabilistic systems, and in particular proposed a state-based version of bisimulation for PTSs through partitioning. That is, two states of a PTS were said to be bisimilar whenever their associated transition probabilities coincided over each equivalence class. Such procedure is hard to generalize for non-symmetric relations, so a follow-up work [75] introduced state-based simulation for PTSs by lifting relations from states to measures. It also showed that a version of bisimulation defined through such lifting is equivalent to the one defined through partitioning in [90]: this result is an analogue of Theorem 3.15 here. Since PTSs were not given explicit semantics, one could not say that (bi-)simulation implies behavioral inclusion (equivalence) as no notion of behaviors was introduced. Due to this reason, [90] and [75] only showed that proposed versions of (bi-)simulation preserve a certain limited class of properties, such as testing.

This theory was enriched by the work [114] which introduced a model of probabilistic automaton (PA). The PA model is similar to the MDP one and resolves non-determinism by means of policies (called *schedulers* in [114]), and is thus endowed with the conventional quantitative semantics¹¹. The paper [114] defined system-based (bi-)simulation for PAs via a familiar method of relation lifting¹². Furthermore, it also proposed a notion of probabilistic (bi-)simulation and showed that both exact and probabilistic versions preserve PCTL properties. A somewhat stronger result, even though only for finite autonomous SSSs, was obtained in [16] which showed that bisimulation preserves PCTL* properties, and thus implies behavioral equivalence.

Let us give some details regarding probabilistic (bi-)simulation from Definition 3.10. Here we generalize the concept introduced in [114] to general SSSs via the operation of *sco*, whereas in the original work probabilistic (bi-)simulation was defined using “mixed” transitions as much as we have had to do this for MDPs in Definition 3.41 using an auxiliary measure ν . Note that a name “probabilistic bisimulation” also appears in [90], but it has a different meaning. In fact, “probabilistic bisimulation” of [90] is just what we call here a bisimulation for probabilistic systems, whereas a “bisimulation” there was defined as a coarser relation that obtained by interpreting PTSs as TSs – see the Appendix A.4 for the details.

It was well-understood that a complex system rarely admits a bisimilar one which is much simpler. For this reason, to solve verification problems over concrete systems [35] proposed the following abstraction procedure for discrete SSSs. One starts with an abstraction which simulates the concrete system. Such abstraction is assumed to be obtained by partitioning the state space of the concrete system; by refining such partition, one improves the quality of abstraction in a certain sense. This procedure was developed to provide bounds on maximal and minimal reachability probabilities for discrete MDPs. A similar method for finite autonomous SSSs was proposed in [53] which used the concept of abstract Markov

¹¹ In fact, the semantics of PAs is richer than that of MDPs as it allows for two kinds of non-determinism, of which one is resolved as adversarial, and hence is close to stochastic zero-one games [89].

¹² The paper is written in a technically cumbersome language as it introduces a new probability space for each probability measure it considers. To follow its ideas in an easy way, one shall abstract away all but a probability measure whenever a probability space is mentioned in [114].

Chains (AMCs): an AMC is a representation of SSs where one is only given minimal and maximal probabilities of going from one state to another. The latter work proposed an abstraction procedure similar to the one on Figure 3.2, and used the fact that Γ set of an SS obtained from an AMC consists of convex polytopes, and hence one can deal just with a finite set of their extreme points. This procedure was further elaborated on in [81] which also extended it to (finite) continuous-time Markov Chains. The abstraction methods above, when applied to autonomous systems, could perform well, but over controlled system the quality of bounds on the reachability probabilities they provide would crucially depend on the gap between the values of minimal and maximal reachability probability over the original system, which clearly cannot be controlled by refining an abstraction: see e.g. the discussion following (2.9). To cope with this issue, [89] developed a theory of abstracting finite SSs with stochastic games where a non-determinism obtained by an abstraction was treated as an adversary. Thus, both minimal and maximal reachability probabilities over the concrete system could be bounded by the corresponding pairs of probabilities over a abstract stochastic game, and the quality of such abstraction would not depend on the aforementioned gap.

All the works mentioned above dealt with discrete probabilistic models, finite in most of the cases. However, due to inherently quantitative semantics, even such models motivated the development of approximate relations and metrics between probabilistic systems. Similarly to exact notion of (bi-)simulation, its approximate counterpart could be defined either as an ε -relation on the level of systems (system-based approach), or first as an ε -relation between states of a single system (model-based approach). Recall further than any approximate equivalence defines a pseudometric and vice-versa. A major part of the relevant literature focused on state-based approximate bisimulations (rather than simulations) introduced as a pseudometric. Less often it is introduced as state-based ε -relation and even more rarely approximate simulation is defined.

It is likely that [61] was the first work to introduce notion of ε -bisimulation over discrete probabilistic systems: the latter was defined as a state-based ε -relation. This has further led to a series of works on pseudometrics for discrete and general PTSs such as [42], [142] and [58] to name a few. The former work used ideas based on a logical characterization of state-based exact bisimulation for PTSs, in this and some following papers called *Labelled Markov Processes (LMPs)*. Namely, [25] has shown that two states of a PTS are bisimilar iff they satisfy the same formulas of a certain logic \mathcal{L} . Such logic was extended in [42] to a family of functions \mathcal{F} , and the value of the \mathcal{F} -distance between two states was defined as a difference between functions evaluations over these states. In particular, \mathcal{F} -distance equals zero iff the two states are bisimilar. Again, due to the reason that PTS were not given explicit semantics, one could not claim that such pseudometric implies some kind of approximate behavioral inclusion. The work in [142] followed a coalgebraic approach to PTSs introduced in [37], see the Appendix B for some details on coalgebras. The authors of [142] proposed a functor in the category of pseudometric spaces based on the Kantorovich metric [62], and have shown that such functor admits the final coalgebra (with a corresponding unique pseudometric), so that one could define a distance between two states of a PTS as a value of pseudometric

between their (unique) images in the final co-algebra. Interestingly, although defined in completely different manners, pseudometrics of [42] and [142] appeared to be the same up to a scaling factor [139, Theorem 2], so below we refer to it just as the PTS pseudometric.

Some part of the follow-up work was devoted to computing this pseudometric over finite PTSs [140], however the reason for this computational challenge is not quite clear. Although [42] motivated the use pseudometric for PTSs as a tool to use abstraction for computation of interesting properties over concrete systems, as we have above mentioned, it is still unknown which exactly valuable properties such pseudometric preserves for a simple reason that PTSs are not given explicit semantics. For example, since the syntax of PTSs and MDPs is the same, one can endow PTSs with the semantics of MDPs, however it is by no means clear whether the pseudometric of [42] and [142] is capable of providing non-trivial bounds for such basic quantitative properties of MDPs as maximal reachability probabilities. Indeed, although by [16] exact bisimulation implies behavioral inclusion over finite autonomous PTSs endowed with the semantics of Markov Chains, and two states are bisimilar iff their \mathcal{F} -distance is zero, there is no guarantee that non-zero \mathcal{F} -distance bounds difference in quantitative properties, and even if it does, how to derive such bounds¹³. Unfortunately, none of the papers on the PTS pseudometric contains an illustrative case study at least with finite model of some interesting process, which could help understanding usefulness of this pseudometric. Hence, from the approximation point of view, the value of such pseudometric is rather questionable, and it is not clear what is the reason for its computation if not for practical purposes. One attempt to cope with this issue was performed in [58] where a similar pseudometric was used to bound DC criterion over MDPs with finite action spaces. A state-based version of approximate simulation for finite PTSs was proposed in [44], which studied its connection with some of the logics that are richer than \mathcal{L} . A corresponding notion of state-based approximate bisimulation was exploited in [46] to provide guarantees for the satisfaction of PCTL properties over finite autonomous SSs.

The research on approximate relations for probabilistic systems inspired studied on this topic over TSs. The work in [64] introduced approximate behavioral inclusion (equivalence), together with approximate (bi-)simulation for TSs. Since showing that the former is implied by the latter is considerably easier over TSs than over SSs, the goal of [64] and the follow-up works was to provide sufficient conditions to assure that two given systems are ε -bisimilar, or that a given concrete system admits a finite ε -bisimilar abstraction. Since ε -bisimulation over TSs implies behavioral ε -inclusion over the infinite time horizon, such conditions were provided via the notion of Lyapunov-like approximate bisimulation functions.

The methods developed over TSs in turn have led to two novel approaches to approximate bisimulation over probabilistic systems. The first was developed in [76] where a stochastic version of an approximate bisimulation function was used as

¹³ In fact, there can be other families of functions different from \mathcal{F} that induce different distances which are yet zero exactly over bisimilar states. As a result, unless such bounds are derived there is no particular reason for the choice of \mathcal{F} done in [42] to be the natural one.

an upper-bound on the probability of ε -divergence, that is the probability that outputs of the concrete system and the abstraction become ε -far from each other on the infinite time horizon. For linear continuous-time autonomous stochastic processes [76] provided ways to fine quadratic approximate stochastic bisimulation functions in terms of matrix inequalities. A discrete-time version of this result was obtained in [2]¹⁴. Furthermore, [1] proposed sufficient conditions for non-linear stochastic processes and [8] suggested a Monte-Carlo procedure to compute the probability of ε -divergence over a finite time horizon when no approximate stochastic bisimulation function is available in a closed form. These line of work can be seen as a system-based approach to approximate bisimulation for probabilistic systems. Although the probability of ε -divergence provides upper-bounds on the reachability probabilities of the concrete systems in terms of those computed over the abstraction, such bounds are often conservative. The reason is that these works assumed a given coupling of a concrete system and an abstraction, whereas such coupling can be chosen freely to minimize the distance between systems – see e.g. the discussion in [135]. The second approach was pursued in [148]: there stochastic systems were given a semantics of transition systems on with states being probability measures over the original state space, see the Appendix A.4 for the details. With such interpretation one can use all the machinery of [64] to develop approximate bisimulation theory over probabilistic systems via approximate bisimulation functions. However, as we have mentioned in Section 2.4, exactly due to this interpretation, the semantics of systems considered in [148] differs from the conventional semantics of probabilistic systems, and it is not clear which properties do their result provide guarantees on.

Our work proposes using the total variation distance for measures, and as we have shown this choice allows approximating solutions of verification and synthesis problems on the one hand, and yet admits finite approximate abstraction of many non-trivial infinite systems on the other. Clearly, the proposed approach has some drawbacks, the main one concerning the validity of the approximation theory only for finite time horizons. Although one may think of posing some stability assumptions on SSs to extend the finite-time horizon bounds of Theorem 3.32 to the infinite horizon, we have not found any interesting sufficient conditions. Perhaps, one such condition is to assume that there exists a state which is always reached by the trajectory of the SS, and from which there are no transitions to other states, however such assumption is too conservative. For this reason, so far we can only provide specification-dependent approximate solution method for the infinite-time horizon verification and synthesis problems. Due to the aforementioned reasons, it also shall be clear that any metric between SSs that upper-bounds the total-variation distance between their infinite-time horizon strategic measures is likely to give the trivial value 1 in most of the cases unless on restricts attention to discrete SSs exclusively. A similar approach was later independently pursued in [67], where the authors work directly with MDPs rather than SSs, and notion of precise (bi)similarity introduced there is close to the one in Definition 3.41 here, whereas their notion of approximate (bi)similarity allows for small differences both in available distributions, and in observations (the set Y

¹⁴ This work also contains a survey on approximate bisimulation for probabilistic systems.

is supposed to be endowed with some metric there). Unfortunately, even though claimed so by [67, Theorem 5], the introduced relations there fail to be transitive: state relations are supposed to be Borel-measurable there¹⁵, and according to the proof of that theorem, a composition of relations is used to show the transitivity. As we have mentioned above, the resulting relation itself may fail to be Borel-measurable, which is exactly one of the reasons why in this thesis relations are assumed to be analytic sets. The authors also follow framework of [114], operating with probability spaces, rather than simply focusing on measures.

One way to overcome the issue mentioned in the previous paragraph is to consider a metric on distribution different from the total variation distance. It seems that the Kantorovich metric between strategic measures could be a good choice, see e.g. examples in [135]. Unfortunately, it is unclear how to bound such metric just based on the distances between stochastic relations Γ ¹⁶. For example, although [58] and [142] use Kantorovich metric to define approximate bisimulation for MDPs and PTSs, in their case this metric is only used over state spaces and is not propagated to measures over trajectories. There is a large choice of metrics for probabilities [62], and some of them may be as useful as the total variation distance to approximate behaviors, yet being easier to work with. Perhaps, some also behave nicely w.r.t. the operation sco and may even allow defining a combined version of probabilistic and approximate simulation which is triangular (cf. Remark 3.26).

¹⁵ The measurability requirements are somewhat imprecise there, e.g. in [67, Definition 6] the relation is assumed to be a Borel set, whereas in [67, Definitions 7, 9, 10] explicit measurability requirements are replaced by an unclear statement “under the same conditions as above”. In case this means that relations are just any subsets of the product space, those definitions are not formal as they use lifting of relations, which was only defined for Borel sets in [67, Definitions 5, 8].

¹⁶ The authors has recently discovered that imposing strong stability assumptions on autonomous SSs, such as Lipschitz continuity of transition kernels in Kantorovich metric, and uniform stability of them w.r.t. some state is enough to bound infinite trajectories in this metric. However, these results seem to require too conservative assumptions, and estimates made there are not tight: for example, the maximum of the infinite sequence of distances is replace by a sum thereof.

Infinite-horizon case

This chapter tackles the case of infinite-horizon properties, as the framework of approximate stochastic bisimulation in general does not allow for non-trivial results off the finite time bound.

4.1 Motivation

As we have seen in Chapter 3, in general we cannot use ε -bisimulation to control error over the infinite horizon specifications. To cope with this issue, we use the automata reduction techniques and focus on problems, whose specification can be expressed as some DFA or DBA. Even though that does cover all ω -regular expressions, or even the full LTL, it is still enough to describe a lot of practical situations.

Recall from Lemma 2.9 that due to the reasons above, for our purposes it is sufficient to learn how to solve reachability (for DFA) and repeated reachability (for DBA) problems for any given SS, which is exactly the goal of this section. We start the former problem, and provide a rather detailed approach to its solution. Unfortunately, results for the latter problem are relatively scarce.

4.2 Reachability problem

4.2.1 Characterization

It is more convenient to consider a slightly more general setup, called the *constrained reachability* problem [17, Section 10.1.1].¹ To satisfy the constrained reachability property, the path of an SS does not only have to reach a given goal set, but also to stay within some safe set before hitting the goal one. In terms of the LTL

¹ The constrained reachability problem is also known as the *reach-avoid* problem [124].

grammar, we are going to deal with the property SU_nG , where S is a safe set and G is a goal set. The (unconstrained) reachability problem corresponds to the special case $\Diamond_n G = \text{true} \cup_n G$.

More precisely, consider an SS $\mathcal{S} = (X, \Gamma, Y, L)$, where $Y = \{G, S, D\}$ are the labels corresponding to goal, safe and dangerous states. For example, $L^{-1}(S) \in \mathcal{B}(X)$ is a set of safe states. For any initial distribution α , any strategy $\sigma \in \Sigma^\Gamma$, and any time horizon $n \in \mathbb{N}$ we are thus interested in the value of $Q_\alpha^\sigma(SU_nG)$. It is more convenient to focus on the initial distribution supported on single points and thus consider a function $Q_{\delta(\cdot)}^\sigma(SU_nG) : X \rightarrow [0, 1]$, extending the results to arbitrary initial distributions at a later stage. To make the notation easier, we write Q_x in place of more cumbersome but formal $Q_{\delta(x)}$. Clearly, $Q_{(\cdot)}^\sigma(SU_nG) \in \text{bU}(X)$ for any $\sigma \in \Sigma^\Gamma$ and $n \in \mathbb{N}$. Moreover, the sequence $(Q_x^\sigma(SU_nG))_{n \in \mathbb{N}}$ is non-decreasing in n and furthermore for any fixed $x \in X$

$$Q_x^\sigma(SUG) = \lim_{n \rightarrow \infty} Q_x^\sigma(SU_nG). \quad (4.1)$$

Obviously, the unconstrained reachability is a special instance of the constrained reachability in case the safe set is the whole state space, i.e. $D = \emptyset$. On the other hand, the constrained reachability can be also obtained from the unconstrained one by changing the dynamics of the SS on the set D [137, Section 3.1].

Note that on a part of the state space the value function is already known:

$$Q_x^\sigma(SU_nG) = \begin{cases} 1, & \text{if } x \in G, \\ 0, & \text{if } x \in D \end{cases} \quad (4.2)$$

and, as a result, the constrained reachability problem needs to be solved only for states in S . Recall from Section 2.2 that DFAs are not closed under negations, and hence if we consider both minimization and maximization reachability problems it would allow us to cover a bigger class of specifications (cf. Appendix C.1). As in Section 2.3, we use notation $\mathcal{S}_x(SU_nG)$ for the maximal probability; and now we also adopt the notation $\mathcal{S}_x^\nabla(SU_nG)$ for the minimal one.

In optimal control, the principles of dynamic programming (DP) allow decomposing the general optimization problem into smaller and simpler subproblems [19]. In the literature there have been several results developing DP characterizations of the constrained reachability problem. One of the main differences in these studies has been the choice of the structural representation of the value function $Q_{(\cdot)}^\sigma(SU_nG)$. For example, the work in [9] has considered the max cost representation for the unconstrained reachability, as

$$Q_x^\sigma(XU_nG) = Q_x^\sigma \left[\max_{k \leq n} 1_G(y_k) \right], \quad (4.3)$$

and using the dual safety problem, an alternative multiplicative cost representation

$$Q_x^\sigma(SU_nG) = 1 - Q_x^\sigma \left[\prod_{k=0}^n 1_{G^c}(y_k) \right]. \quad (4.4)$$

These results have been extended in [124], which has dealt with the general constrained reachability problem in the form of a sum-multiplicative cost

$$Q_x^\sigma(SU_n G) = Q_x^\sigma \left[\sum_{k=0}^n \left(\prod_{j=0}^{k-1} 1_{S \setminus G}(y_j) \right) 1_G(y_k) \right]. \quad (4.5)$$

Later, [30] suggested a cost formulation using the notion of a first hitting time as

$$Q_x^\sigma(SU_n G) = Q_x^\sigma \left[\sum_{k=0}^{n \wedge \tau_G \wedge \tau_D} 1_G(\mathbf{x}_k) \right], \quad (4.6)$$

where $\tau_A := \inf\{k \geq 0 : \mathbf{x}_k \in A\}$ is the first hitting time of the set $A \in \mathcal{B}(X)$. Unfortunately, none of the above cost functions was well-studied in the literature. Our idea instead is to express the reachability problem through the additive discounted cost (DC) criterion: recall from Section 2.2 that such case allows for a rich theory of DP in a rather general setting. For example, the aforementioned studies in [9], [124] and [30] have recovered only a subset of these results for the reachability problem, sometimes requiring restrictive assumptions on the structure of the model. In contrast, here we show that the reachability problem has an equivalent DC formulation, which allows us proving all results available for this general performance criterion.

In general it may not be possible to characterize the constrained reachability problem as a DC criterion over the original SS \mathcal{S} . The key idea is to consider an auxiliary SS $\hat{\mathcal{S}}$, constructed from the original one by adding a new binary variable that represents whether the path of \mathcal{S} has left the safe set S or not. To our knowledge, the first time such construction has been explicitly used in [137].² For the sake of consistency, here we introduce a new SS, using the notion of the composition between SSs and automata from Definition 2.8




Figure 4.1: Transition system for the TC formulation of constrained reachability

Let us consider an automaton $\mathcal{D} = (Q, q^s, Y, t, A)$ as in Figure 4.1 with a state space $Q = \{q^s, q^f\}$, an input alphabet $Y = (D, G, S)$, and transition function given by

$$t(q^s, S) = q^s, \quad t(q^s, \{D, G\}) = q^f, \quad t(q^f, Y) = q^f.$$

Let $\hat{\mathcal{S}} := \mathcal{S} \otimes \mathcal{D} = (\hat{X}, \hat{\Gamma}, \hat{L}, Q)$ denote the composed SS. Of special interest are functions $\hat{f} : \hat{X} \rightarrow \mathbb{R}$ that are zero off q^s , namely $\hat{f}(\cdot, q^f) \equiv 0$: they can be always represented in the form

$$\hat{f}(x, q) = 1\{q = q^f\} \cdot f(x) \quad (4.7)$$

² In [45] a similar construction was used to formulate reachability problem as a final cost problem.

for some $f : X \rightarrow \mathbb{R}$. Let $c : \hat{X} \rightarrow \{0, 1\}$ be a cost function $c(x, q) := 1\{q = q^s\} \cdot 1_G(x)$ that is zero off q^s , and define the corresponding DC utility for any $n \in \bar{\mathbb{N}}$ as follows:

$$\hat{J}_n := \sum_{k=0}^n c(x_k, q_k), \quad (4.8)$$

where x and q are the components of the state process of the composed system. The corresponding maximization and minimization problems are given by

$$\hat{S}_{(x,q)}(\hat{J}_n) = \sup_{\hat{\sigma} \in \Sigma^{\hat{\Gamma}}} \hat{P}_{(x,q)}^{\sigma} [\hat{J}_n], \quad \hat{S}_{(x,q)}^{\nabla}(\hat{J}_n) = \inf_{\hat{\sigma} \in \Sigma^{\hat{\Gamma}}} \hat{P}_{(x,q)}^{\sigma} [\hat{J}_n]. \quad (4.9)$$

In order to show the equivalence between the optimal constrained reachability problem over the SS \mathcal{S} and the formulation in (4.9) over the SS $\hat{\mathcal{S}}$, we apply the technique from Lemma 2.9. Let us again denote by $\mathcal{J} : \Sigma^{\Gamma} \rightarrow \Sigma^{\hat{\Gamma}}$ the embedding map used there, and consider a slightly different the projection map $\hat{\mathfrak{P}} : \Sigma^{\hat{\Gamma}} \rightarrow \Sigma^{\Gamma}$ given by

$$(\hat{\mathfrak{P}}\hat{\sigma})_n(x_0, \dots, x_{n-1}, x_n) := \hat{\sigma}_n(x_0, q^s, \dots, x_{n-1}, q^s, x_n, q^s). \quad (4.10)$$

Note that $\hat{\mathfrak{P}}$ is not the projection map \mathfrak{P} used in Lemma 2.9, in particular, below we use the fact that $\hat{\mathfrak{P}}$ preserves Markovianity of strategies, in contrast to \mathfrak{P} . The following equivalence holds true:

Lemma 4.1 *For any $n \in \bar{\mathbb{N}}$, $\sigma \in \Sigma^{\Gamma}$ and $\hat{\sigma} \in \Sigma^{\hat{\Gamma}}$, it holds that*

$$\hat{Q}_{(x,q^s)}^{\hat{\sigma}}(\hat{J}_n) = Q_x^{\hat{\sigma}}(SU_n G), \quad Q_x^{\sigma}(SU_n G) = \hat{Q}_{(x,q^s)}^{\hat{\sigma}}(\hat{J}_n). \quad (4.11)$$

Proof: We prove this theorem by induction. First of all, both equalities in (4.11) clearly hold true for $n = 0$ as in this case all functions are simply $1_G(x)$. Furthermore, with focus on the first equality, we have that

$$\hat{Q}_{(x,q^s)}^{\hat{\sigma}}(\hat{J}_{n+1}) - \hat{Q}_{(x,q^s)}^{\hat{\sigma}}(\hat{J}_n) = \hat{P}_{(x,q^s)}^{\hat{\sigma}}[c(\mathbf{x}_{n+1}, \mathbf{q}_{n+1})].$$

As $c(\mathbf{x}_{n+1}, \mathbf{q}_{n+1})$ is a Bernoulli random variable supported on $\{0, 1\}$, we obtain that

$$\begin{aligned} \hat{P}_{(x,q^s)}^{\hat{\sigma}}[c(\mathbf{x}_{n+1}, \mathbf{q}_{n+1})] &= \hat{P}_{(x,q^s)}^{\hat{\sigma}}(c(\mathbf{x}_{n+1}, \mathbf{q}_{n+1}) = 1) \\ &= \hat{P}_{(x,q^s)}^{\hat{\sigma}}(\{\mathbf{x}_k \in S, k \leq n\}, \{\mathbf{x}_{n+1} \in G\}, \{\mathbf{q}_k = q^s, k \leq n+1\}). \end{aligned}$$

On the other hand,

$$Q_x^{\hat{\sigma}}(SU_{n+1} G) - Q_x^{\hat{\sigma}}(SU_n G) = P_x^{\hat{\mathfrak{P}}\hat{\sigma}}(\{\mathbf{x}_k \in S, k \leq n\}, \{\mathbf{x}_{n+1} \in G\}).$$

The fact that these probabilities are equal follows immediately from their integral expressions in from the definition of the projection map $\hat{\mathfrak{P}}$. By induction we obtain the first part in (4.11) for $n < \infty$, and the case $n = \infty$ follows by taking the limit.

Finally, the proof of the second part of (4.11) is obtained the same way, mutatis mutandis. \square

Lemma 4.1 leads to several important results that allow us to develop a DP framework for constrained reachability. First of all, it clearly implies that both optimization problems are equivalent in the following sense:

Theorem 4.2 *For all $n \in \bar{\mathbb{N}}$ and $x \in X$ we have $\hat{\mathcal{S}}_{(x,q^f)}(\hat{J}_n) = \hat{\mathcal{S}}_{(x,q^f)}^\nabla(\hat{J}_n) = 0$ and*

$$\mathcal{S}_x(SU_n G) = \hat{\mathcal{S}}_{(x,q^s)}(\hat{J}_n), \quad \mathcal{S}_x^\nabla(SU_n G) = \hat{\mathcal{S}}_{(x,q^s)}^\nabla(\hat{J}_n). \quad (4.12)$$

Proof: To prove the first part, one has to notice that if $q_0 = q^f$, then $q_n = q^f$ for all $n \in \mathbb{N}$, hence $\hat{Q}_{(x,q^f)}^{\hat{\sigma}}(\hat{J}_n) = 0$ for all $n \in \mathbb{N}$, $x \in X$, and $\hat{\sigma} \in \Sigma^{\hat{\Gamma}}$. Furthermore, (4.12) is an immediate consequence of Lemma 4.1 and Lemma C.9 in the Appendix. \square

As we have mentioned above, Theorem 4.2 allows us to extrapolate the rich theory developed for the DC criterion to the case of the constrained reachability problem. However, most of the results for DC are developed for the minimization case [20, 70], considering either positive or negative costs c . As such, we can directly derive the results for the minimization problem since $\mathcal{S}_x^\nabla(SU_n G) = \hat{\mathcal{S}}_{(x,q^s)}^\nabla(\hat{J}_n)$, however for the maximal constrained reachability we need to interpret

$$\mathcal{S}_x(SU_n G) = -\hat{\mathcal{S}}_{(x,q^s)}^\nabla(-\hat{J}_n),$$

thus characterizing both optimization problems as a minimization of some DC. Note that for the minimization of the constrained reachability we use a positive cost c , thus falling into the setting of the positive DP [22] corresponding to [20, Assumption (P), Chapter 9]. On the other hand, for the maximization of the constrained reachability a negative cost $-c$ is used, hence leading to the case of the negative DP [121] corresponding to [20, Assumption (N), Chapter 9]. This difference is not always important and only matters in the case $n = \infty$. In particular, we show below that it affects the convergence of bounded-horizon functions to the unbounded-horizon ones, as well as the existence of optimal policies.

Let us proceed with the application of Lemma 4.1 and Theorem 4.2 to the characterization of the optimal constrained reachability problems. The next results shows that it is sufficient to deal with Markov policies.

Proposition 4.3 *For any $n \in \bar{\mathbb{N}}$ and any policy $\sigma \in \Sigma^\Gamma$, there exists a Markov strategy $\sigma' \in \Sigma_M^\Gamma$ such that $Q_{(\cdot)}^\sigma(SU_n G) = Q_{(\cdot)}^{\sigma'}(SU_n G)$, and as a consequence*

$$\mathcal{S}_x(SU_n G) = \sup_{\sigma \in \Sigma_M^\Gamma} Q_x^\sigma(SU_n G), \quad \mathcal{S}_x^\nabla(SU_n G) = \inf_{\sigma \in \Sigma_M^\Gamma} Q_x^\sigma(SU_n G). \quad (4.13)$$

Proof: Fix any state $x \in X$ and any strategy $\sigma \in \Sigma$. It follows from Lemma 4.1 that $Q_x^\sigma(SU_n G) = \hat{Q}_{(x,q^s)}^{\hat{\sigma}}(\hat{J}_n)$. On the other hand, [20, Proposition 8.1] assures the

existence of a Markov strategy $\hat{\sigma}' \in \Sigma_M^{\hat{\Gamma}}$ satisfying $\hat{Q}_{(x,q^s)}^{\hat{\sigma}}(\hat{J}_n) = \hat{Q}_{(x,q^s)}^{\hat{\sigma}'}(\hat{J}_n)$. From the definition of the projection map \mathfrak{P} it follows that $\sigma' := \mathfrak{P}\hat{\sigma}' \in \Sigma_M^{\Gamma}$ and as a result

$$Q_x^{\sigma'}(SU_n G) = \hat{Q}_{(x,q^s)}^{\hat{\sigma}'}(\hat{J}_n) = \hat{Q}_{(x,q^s)}^{\hat{\sigma}}(\hat{J}_n) = Q_x^{\sigma}(SU_n G),$$

as desired. In order to obtain (4.13) we only have to apply Lemma C.9. \square

The results above, obtained for deterministic initial conditions, can be extended to the case of general initial distributions: we show that a value function over an initial distribution $\alpha \in \mathcal{P}(X)$ can be obtained by integrating value functions over deterministic initial conditions. Although this result is obvious in case the strategy is fixed, it is not trivial to show this for optimal value functions. We show a proof for the case of the minimization problem on the unbounded time horizon, however similar results can be obtained for the unbounded-time maximization case, as well as for both bounded-horizon problems.

Proposition 4.4 *For any distribution $\alpha \in \mathcal{P}(X)$ it holds that*

$$S_{\alpha}^{\nabla}(SUG) = \int_X S_x^{\nabla}(SUG) \alpha(dx). \quad (4.14)$$

Proof: From [20, Propositions 9.2, 9.3, 9.5] it follows that

$$\hat{S}_{\hat{\alpha}}^{\nabla}(\hat{J}_{\infty}) = \int_{\hat{X}} \hat{S}_{(x,q)}^{\nabla}(\hat{J}_{\infty}) \hat{\alpha}(dx \times dq)$$

for any distribution $\hat{\alpha} \in \mathcal{P}(\hat{X})$. As a result, for any $\alpha \in \mathcal{P}(X)$ it holds that

$$\begin{aligned} S_{\alpha}^{\nabla}(SUG) &= \inf_{\sigma \in \Sigma} \int_{\hat{X}} \hat{Q}_{(x,q)}^{\sigma}(\hat{J}_{\infty})(\alpha \otimes \delta_{q^s})(dx \times dq) \\ &\geq \int_{\hat{X}} \hat{S}_{(x,q)}^{\nabla}(\hat{J}_{\infty})(\alpha \otimes \delta_{q^s})(dx \times dq) \\ &= \int_{\hat{X}} \hat{S}_{(x,q^s)}^{\nabla}(\hat{J}_{\infty}) \alpha(dx) = \int_X S_x^{\nabla}(SUG) \alpha(dx). \end{aligned}$$

The converse inequality we get as follows:

$$\begin{aligned} \int_X S_x^{\nabla}(SUG) \alpha(dx) &= \int_{\hat{X}} \hat{S}_{(x,q^s)}^{\nabla}(\hat{J}_{\infty}) \alpha(dx) \\ &= \int_{\hat{X}} \hat{S}_{(x,q)}^{\nabla}(\hat{J}_{\infty})(\alpha \otimes \delta_{q^s})(dx \times dq) \\ &= \inf_{\hat{\sigma} \in \Sigma^{\hat{\Gamma}}} \int_{\hat{X}} \hat{Q}_{(x,q)}^{\hat{\sigma}}(\hat{J}_{\infty})(\alpha \otimes \delta_{q^s})(dx \times dq) \\ &= \inf_{\hat{\sigma} \in \Sigma^{\hat{\Gamma}}} \int_X Q_x^{\hat{\sigma}}(SUG) \alpha(dx) \geq S_{\alpha}^{\nabla}(SUG). \end{aligned}$$

Since both inequalities hold true, we obtain the desired result. \square

Although in general one cannot switch the order of the minimization (or maximization) and of the integral, Proposition 4.4 shows this can be done in the case of (4.14). Thus, it is sufficient to deal with deterministic initial distributions: the value function for the general one can be obtained by integrating with respect to the initial distribution of interest.

Before proceeding, we need to introduce some concepts useful in solving the reachability problem. The reader should take heart from now on, as despite our best effort, notation is going to become even more cumbersome than before. Alas, to the best of our knowledge, better and simpler notation was not yet introduced to describe the things that follow.

In Chapter 3 we have discussed how to approximate quantitative behaviors of one SS with those of another one, however we have never mentioned how does one actually compute the latter values, as it was not important back then. Now we have to face it directly, and it should come as no surprise that since we are using DP, the methods are going to be iterative. In particular, we are considering effect of one transition in an SS on some function defined over its state space. For this purpose, we introduce the following two operators:

We are ready to formulate one of the most relevant outcomes of Theorem 4.2: a DP procedure for the constrained reachability problem over a general class of policies. For this purpose we introduce the following DP operators:

$$\begin{aligned} r^\Delta f(x) &= 1_G(x) + 1_S(x) \cdot \mathfrak{T}^\Delta f(x), & f \in \mathbf{bA}^\Delta(X), \\ r^\nabla f(x) &= 1_G(x) + 1_S(x) \cdot \mathfrak{T}^\nabla f(x), & f \in \mathbf{bA}^\nabla(X). \end{aligned}$$

From the properties of operators \mathfrak{T}^Δ and \mathfrak{T}^∇ , it follows that r^Δ maps $\mathbf{bA}^\Delta(X)$ into itself and r^∇ maps $\mathbf{bA}^\nabla(X)$ into itself.

Theorem 4.5 *It holds that $\mathcal{S}(\cdot; \mathbf{SU}_0 G) = \mathcal{S}^\nabla(\cdot; \mathbf{SU}_0 G) = 1_G(\cdot)$, and for any $n \in \bar{\mathbb{N}}$*

$$\mathcal{S}(\cdot; \mathbf{SU}_{n+1} G) = r^\Delta [\mathcal{S}(\cdot; \mathbf{SU}_n G)], \quad \mathcal{S}^\nabla(\cdot; \mathbf{SU}_{n+1} G) = r^\nabla [\mathcal{S}^\nabla(\cdot; \mathbf{SU}_n G)].$$

Moreover, $\mathcal{S}(\cdot; \mathbf{SUG})$ and $\mathcal{S}^\nabla(\cdot; \mathbf{SUG})$ are the least non-negative fixpoints of the corresponding operators, that is if there exists a non-negative function $f \in \mathbf{bA}^\Delta(X)$ (or $f \in \mathbf{bA}^\nabla(X)$) that satisfies the inequality $f \geq r^\Delta[f]$ (or $f \geq r^\nabla[f]$), then it holds that $f(\cdot) \geq \mathcal{S}(\cdot; \mathbf{SUG})$ (or $f(\cdot) \geq \mathcal{S}^\nabla(\cdot; \mathbf{SUG})$).

Proof: We provide an explicit proof for the minimization problem, and appeal to duality for the maximization case.

First of all, the fact that $\mathcal{S}^\nabla(\mathbf{SU}_0 G) = 1_G(\cdot)$ follows immediately from the definition of the constrained reachability. Furthermore, for any $n \in \bar{\mathbb{N}}$ by Theorem 4.2 we have that $\mathcal{S}_x^\nabla(\mathbf{SU}_n G) = \hat{\mathcal{S}}_{(x, q^s)}^\nabla(\cdot; \hat{J}_n)$. The DP recursion for the TC is given in [20,

Proposition 8.2, Proposition 9.8], and applied here yields the following:

$$\begin{aligned} \mathcal{S}_x^\nabla(\text{SU}_{n+1}G) &= \hat{\mathcal{S}}_{(x,q^s)}^\nabla(\hat{J}_{n+1}) = \inf_{\hat{\gamma} \in \hat{\Gamma}|_{(x,q^s)}} \left(c(x, q^s) + \hat{\gamma} \mathcal{S}_{(x,q^s)}^\nabla(\hat{J}_n) \right) \\ &= \inf_{\gamma \in \Gamma|_x} \left(1_G(x) + 1_S(x) \gamma \mathcal{S}_{(x,q^s)}^\nabla(\hat{J}_n) \right) \\ &= 1_G(x) + 1_S(x) (\Gamma|_x)^\nabla \mathcal{S}^\nabla(\text{SU}_n G) = r^\nabla [\mathcal{S}_x^\nabla(\text{SU}_n G)]. \end{aligned}$$

This results in both the DP recursion ($n < \infty$) and in the fixpoint equation (for $n = \infty$).

Consider now a non-negative function $f \in \text{bA}^\nabla(X)$ satisfying $f \geq r^\nabla[f]$, and define a new function $\hat{f} : \hat{X} \rightarrow [0, \infty)$ by the formula $\hat{f}(x, q) := 1\{q = q^s\} \cdot f(x)$. Clearly, the function \hat{f} is zero off q^s , so that we obtain:

$$\inf_{\hat{\gamma} \in \hat{\Gamma}|_{(x,q^s)}} \left(c(x, q) + \hat{\gamma} \hat{f}(x, q) \right) = 1\{q = q^s\} \cdot r^\nabla f(x) \leq 1\{q = q^s\} \cdot f(x) = \hat{f}(x, q).$$

As a result, [20, Proposition 9.10 (P)] implies that $\hat{\mathcal{S}}^\nabla(J_\infty) \leq \hat{f}(\cdot)$ and thus

$$\mathcal{S}_x^\nabla(\text{SUG}) = \hat{\mathcal{S}}_{(x,q^s)}^\nabla(\hat{J}_\infty) \leq \hat{f}(x, q^s) = f(x),$$

so $\mathcal{S}^\nabla(\text{SUG})$ is the least fixpoint in the class of non-negative bA^∇ functions. \square

In view of Theorem 4.5 we can compute the value of the bounded horizon optimal constrained reachability problems backward-recursively, starting from the indicator function 1_G . The computation of the fixpoint problem is more intricate and is addressed below in Section 4.2.2. Due to this reason, it is worth discussing the relation between the solution of the constrained reachability problem on the bounded time horizon, and that on the unbounded time horizon. In particular, an interesting question is whether the latter can be in general obtained as the limit of the former, as the time index n goes to infinity. This is one of the anticipated cases where the difference between the maximization and minimization problems becomes important: the answer is positive in the first case and is negative in the second.

Proposition 4.6 *For every state $x \in X$ it holds that*

$$\mathcal{S}_x(\text{SUG}) = \lim_{n \rightarrow \infty} \mathcal{S}_x(\text{SU}_n G). \quad (4.15)$$

Furthermore, for any $x \in X$ there exists a limit

$$f^\nabla := x \mapsto \lim_{n \rightarrow \infty} \mathcal{S}_x^\nabla(\text{SU}_n G) \leq \mathcal{S}^\nabla(\text{SUG}). \quad (4.16)$$

Moreover, $\mathcal{S}^\nabla(\text{SUG}) = f^\nabla$ if and only if f^∇ is a fixpoint of the DP operator r^∇ .

Proof: We start with the maximization case, which corresponds to Assumption (N) of [20, Chapter 9] since $\mathcal{S}_x(\text{SU}_n G) = -\hat{\mathcal{S}}_{(x,q^s)}^\nabla(-\hat{J}_n)$ for any $x \in X$. It fol-

lows from [20, Section 9.5] that the sequence $\left(\hat{\mathcal{S}}_x^\nabla(q; -\hat{J}_n)\right)_{n \in \mathbb{N}}$ has a limit for any $x \in X$ and $q \in Q$. Furthermore, [20, Proposition 9.14] implies that this limit is $\hat{\mathcal{S}}_x^\nabla(q; -\hat{J}_\infty)$, which leads to (4.15).

For the minimization case we satisfy Assumption (P) of [20, Chapter 9]. The discussion in [20, Section 9.5] implies the existence of the point-wise limit for the sequence $(\hat{\mathcal{S}}_x^\nabla(q; \hat{J}_n))_{n \in \mathbb{N}}$: we denote this limit by \hat{f}^∇ . Furthermore, it follows from [20, Proposition 9.16] that $\hat{f}^\nabla(\cdot) \leq \hat{\mathcal{S}}^\nabla(\cdot; \hat{J}_\infty)$, and that the equality holds if and only if \hat{f}^∇ is a fixpoint of the corresponding DP operator, i.e.

$$\hat{f}^\nabla(x, q) = c(x, q) + \hat{\mathcal{T}}^\nabla \hat{f}^\nabla(x, q). \quad (4.17)$$

For the constrained reachability case, we now obviously have the existence of the limit

$$f^\nabla(x) := \lim_{n \rightarrow \infty} \mathcal{S}_x^\nabla(SU_n G) = 1\{q = q^s\} \hat{f}^\nabla(x, q).$$

Clearly, $f^\nabla(\cdot) \geq \mathcal{S}^\nabla(\cdot; SUG)$; if f^∇ is a fixpoint of r^∇ , then \hat{f}^∇ satisfies (4.17), thus $\hat{f}^\nabla(\cdot) = \hat{\mathcal{M}}(\cdot; \hat{J}_\infty)$ and hence $f^\nabla(\cdot) = \mathcal{S}^\nabla(\cdot; SUG)$. Conversely, if $f^\nabla(\cdot) = \mathcal{S}^\nabla(\cdot; SUG)$ then by Theorem 4.5 it has to be a fixpoint of the DP operator r^∇ . \square

The following example shows that the inequality in (4.16) can be strict.³

Example 4.7 Let $X = \mathbb{N}$ and $\Gamma|_0 = \{\delta(n)\}_{n \in \mathbb{N}}$ and $\Gamma|_x = \delta(x - 1)$ for $x \neq 0$, so that the dynamics is deterministic. Let $G := \{1\}$ be the goal set, and let the safe set be its complement $S := X \setminus G$. Let us focus on the case when $\mathbf{x}_0 = 0$. If we would like to minimize the probability of reaching G over some finite horizon $n \in \mathbb{N}$, one of the optimal strategies is to choose $\sigma_0(0) = \delta(n + 2)$. Then $\mathbf{x}_1 = n + 1$, $\mathbf{x}_2 = n$ and $\mathbf{x}_n = 2$, so that G is not reached. As a result, for any finite $n \in \mathbb{N}$ we have that $\mathcal{S}_0^\nabla(SU_n G) = 0$. However, regardless of the first choice $\sigma_0(0) = \delta(n)$, the set G is reached by the path of the process in at most $n < \infty$ steps. Thus,

$$1 = \mathcal{S}_0^\nabla(SUG) \neq \lim_{n \rightarrow \infty} \mathcal{S}_0^\nabla(SU_n G) = 0.$$

So far we have developed DP over the value functions for the constrained reachability problem. The main tool we have used is a TC reformulation of the original performance criterion, which makes it possible to apply the rich theory that has been developed for the DC problem. Following similar lines as in the proofs of the theorems above, one can reformulate for the constrained reachability problem almost any result developed for the DC criterion. While in this thesis we do not have a focus on the existence of optimal strategies, one can easily tailor a number of results from [20], as we overview next. Recall that Assumption (P) in [20, Chapter 9] corresponds to the minimization problem, whereas Assumption (N) corresponds to the maximization one.

(P) [20, Proposition 9.12] and its corollary provide necessary and sufficient conditions for the optimality of stationary policies, together with results to com-

³ The example is obtained by modifying [20, Example 1].

pute such policies. Moreover, [20, Propositions 9.17, 9.18] and their corollaries provide various sufficient conditions for the existence of optimal stationary policies, for their Borel measurability, and for the equality in (4.16).

- (N) [20, Proposition 9.13] gives necessary and sufficient conditions for the optimality of stationary policies. However, it does not give a way to construct a policy (such as the one available for (P)). This is almost the only result on the optimality of policies under Assumption (N).

4.2.2 Computation

The DC formulation of the constrained reachability problem not only leads to results for the characterization of its solution, but also connects to computational methods developed for this criterion. For example, see [49] and references therein. Alternatively, one could use the theory of Chapter 3 to solve bounded-horizon reachability with precise bounds on the error. In the present context we are interested in extending these results to the unbounded time horizon case.

Let us recall the classical theory for the DC performance criterion. If its discounting factor satisfies $\gamma < 1$, one falls into the setting of discounted problems for which the corresponding DP operator is contractive on some function space. Such a property has nice consequences: the unbounded-horizon value function is the unique fixpoint of this operator, and it can also be efficiently approximated by means of the bounded-horizon value functions, as it follows from the contraction mapping theorem.⁴ This approach is clearly interesting to us because of the computational techniques developed for the bounded time horizon case. Unfortunately, the DC formulation of the constrained reachability problem (4.8) has a discounting factor $\gamma = 1$, so the contractivity of the DP operators r^Δ and r^∇ cannot be established using classical techniques. Due to this reason, we come up with new sufficient conditions for the DP operators associated to the constrained reachability problem to be contractive: the approach is based on Lemma C.12, which is inspired by [69, Proposition A.2].

The DP operators for the constrained reachability problem are rarely contractive over the whole state space X , so it is worth restricting attention to the safe set S exclusively. This also emphasizes the leading role of the set S in the solution of the problem (in contrast to the goal set G), as we discussed before: we have already mentioned that the solution of the constrained reachability problem is trivial outside of the safe set (4.2), so we can work with the restriction of value functions to the set S . Let us define

$$\Gamma_S^\Delta f(x) := \sup_{\gamma \in \Gamma|_x} \gamma(1_S f), \quad \Gamma_S^\nabla f(x) := \inf_{\gamma \in \Gamma|_x} \gamma(1_S f).$$

Note that the operators Γ_S^Δ and Γ_S^∇ map spaces $\mathcal{BA}^\Delta(S)$ and $\mathcal{BA}^\nabla(S)$ into themselves respectively. Moreover, for $f \in \mathcal{BA}^\nabla(X)$ it holds that $f|_S \in \mathcal{BA}^\nabla(S)$, which

⁴ The contraction mapping theorem is alternatively known as Banach's Fixed Point Theorem [69, Proposition A.1].

follows immediately from the definition of lower-semianalytic functions and Borel measurability of S . Clearly, the same applies to the restrictions of functions in $\text{bA}^\Delta(X)$. In particular, if we define

$$w_n^\Delta(x) := \mathcal{S}_x^\Delta(SU_n G)|_S, \quad w_n^\nabla(x) := \mathcal{S}_x^\nabla(SU_n G)|_S$$

for any $x \in X$ and $n \in \bar{\mathbb{N}}$, then $w_n^\Delta \in \text{bA}^\Delta(S)$ and $w_n^\nabla \in \text{bA}^\nabla(S)$. Thus, we can rewrite the DP over the safe set S as follows:

$$w_{n+1}^\Delta = \mathbf{r}_S^\Delta[w_n^\Delta], \quad w_{n+1}^\nabla = \mathbf{r}_S^\nabla[w_n^\nabla]$$

for any $n \in \bar{\mathbb{N}}$, where $w_0^\Delta = w_0^\nabla = 0$, and the truncated DP operators are given by

$$\begin{aligned} \mathbf{r}_S^\Delta f(x) &:= \sup_{\gamma \in \Gamma|_x} \gamma(1_G + 1_S f), & f \in \text{bA}^\Delta(S), \\ \mathbf{r}_S^\nabla f(x) &:= \inf_{\gamma \in \Gamma|_x} \gamma(1_G + 1_S f), & f \in \text{bA}^\nabla(S). \end{aligned}$$

Clearly, these operators map their domains into themselves, so that they can be applied recursively.

In order to formulate the main result on the contractivity of the DP operators, we are only left to introduce a very important special case of constrained reachability: safety [9]. This can be characterized by the LTL formula $\Box_n S$ and thus

$$\mathbf{Q}_x^\sigma(\Box_n S) = 1 - \mathbf{Q}_x^\sigma(SU_n S^c)$$

for all $x \in X$ and any $n \in \bar{\mathbb{N}}$. We are interested in the restriction of the safety problem to the safe set S itself, the main focus being the characterization of contractivity. Let

$$v_n^\Delta(x) := \mathcal{S}_x(\Box_n S)|_S, \quad v_n^\nabla(x) := \mathcal{S}_x^\nabla(\Box_n S)|_S.$$

The DP for safety over S is hence given by

$$v_{n+1}^\Delta = \Gamma_S^\Delta v_n^\Delta, \quad v_{n+1}^\nabla = \Gamma_S^\nabla v_n^\nabla, \quad n \in \bar{\mathbb{N}}.$$

with $v_0^\Delta = v_0^\nabla = 1$. Clearly, we have that $0 \leq v_n^\Delta \leq 1$ for all $n \in \bar{\mathbb{N}}$. Let us define

$$\beta_n(S) := \sup_{x \in S} v_n^\Delta(x) = \sup_{x \in X} \mathcal{S}_x^\Delta(\Box_n S) \in [0, 1],$$

$$m(S) := \inf\{n \in \bar{\mathbb{N}} : \beta_n(S) < 1\} \in \bar{\mathbb{N}},$$

and note that both β_n and m are monotonic functions of S with respect to set inclusion. We are now ready to provide sufficient conditions for contractivity.

Theorem 4.8 *If $m := m(S) < \infty$, then operators $(\mathbf{r}_S^\Delta)^m$ and $(\mathbf{r}_S^\nabla)^m$ are contractions with modulus $\beta_m(S)$ on the spaces $\text{bA}^\Delta(S)$ and $\text{bA}^\nabla(S)$ respectively. In particular, each of them has a unique fixpoint, for any $n \in \bar{\mathbb{N}}$ the following inequalities hold true:*

$$|w_\infty^\Delta(x) - w_{mn}^\Delta(x)| \leq \beta^n, \quad |w_\infty^\nabla(x) - w_{mn}^\nabla(x)| \leq \beta^n \quad (4.18)$$

for all $x \in S$, operators $(\Gamma_S^\Delta)^m$ and $(\Gamma_S^\nabla)^m$ are contractions as well and $v^\Delta = v_\infty^\nabla = 0$.

Proof: We are going to apply Lemma C.12 in order to establish the contractivity property. Let us consider the case of r_S^Δ first, so in Lemma C.12 we put $\mathcal{F} = \text{b}\mathcal{A}^\Delta(S)$. The condition (1) of the lemma is obviously satisfied for r_S^Δ and hence for $(r_S^\Delta)^n$ regardless of $n \in \mathbb{N}$. Furthermore, for any two functions $f, g \in \text{b}\mathcal{A}^\Delta(S)$ we have that

$$\begin{aligned} r_S^\Delta(f(x) + g(x)) &= \sup_{\gamma \in \Gamma|_x} \gamma(1_G + 1_S f + 1_S g) \\ &\leq \sup_{\gamma \in \Gamma|_x} \gamma(1_G + 1_S f) + \sup_{\gamma \in \Gamma|_x} \gamma(1_S g) \\ &= r_S^\Delta f(x) + \Gamma_S^\Delta g(x). \end{aligned}$$

As a result, for any $f \in \text{b}\mathcal{A}^\Delta(S)$ and any $c \geq 0$ it holds that

$$r_S^\Delta(f + c) \leq r_S^\Delta f + c \cdot v_1^\Delta,$$

and further by induction for any $n \in \mathbb{N}$

$$(r_S^\Delta)^n(f + c) \leq (r_S^\Delta)^n f + c \cdot v_n^\Delta.$$

In particular, for the case $n = m$ we obtain the following:

$$(r_S^\Delta)^m(f + c) \leq (r_S^\Delta)^m f + c \cdot v_m^\Delta \leq (r_S^\Delta)^n f + c \cdot \beta.$$

Hence, $(r_S^\Delta)^m$ satisfies all the assumptions of Lemma C.12 and thus is a contraction on $\text{b}\mathcal{A}^\Delta(S)$. The contractivity of $(r_S^\nabla)^m$ can be shown by a similar argument, with the only difference being the inequality

$$r_S^\nabla(f + g) \leq r_S^\nabla f + \Gamma_S^\Delta g,$$

rather than the one with $\Gamma_S^\nabla g$, and with conditions on contractivity that are state in terms of functions v_n^Δ rather than v_n^∇ .

After the contractivity of the operators is established, the uniqueness of the solutions of fixpoint equations and the bounds in (4.18) follow immediately from the contraction mapping theorem [69, Proposition A.1]. Finally, the statement for operators Γ_S^Δ and Γ_S^∇ follows directly if one considers the special case $G = \emptyset$. \square

Theorem 4.8 shows that in the case of contractive operators the unbounded-horizon value function can be approximated by bounded-horizon ones with any precision level. However, there are some questions left: what are the cases in which the contractivity conditions are violated, and what would be a solution for such cases? Let us first address the former question. For example, whenever the conditions of Theorem 4.8 are met, the equality holds in (4.16). As a result, Example 4.7 does not admit contractive operators since the equality does not hold there. Some of other important examples can be given using the notion of absorbing set.

Definition 4.9 The set $A \in \mathcal{B}(X)$ is called *strongly absorbing* if $\gamma(A) = 1$ for all $x \in A$ and $\gamma \in \Gamma|_x$. The set $A \in \mathcal{B}(X)$ is called *weakly absorbing* if there exists a selector $\kappa \in \mathcal{U}(X|X)$ for Γ such that $\kappa(A|x) = 1$ for all $x \in A$. We say that the set $A \in \mathcal{B}(X)$ is *simple* if it does not have non-empty weakly absorbing subsets.

Note that a strongly absorbing set is always also weakly absorbing, one can e.g. take $\kappa = \sigma_0$ regardless of $\sigma \in \Sigma^\Gamma$, the converse is obviously not true in general, which the use of the adjectives “weak” and “strong” in Definition 4.9. Furthermore, in case of autonomous SSSs, the notion of weak and strong sets coincide with that of an absorbing set [100]. Intuitively, a strongly absorbing set remains absorbing under any possible control action, whereas for a weakly absorbing set there has to exist a strategy that makes this set absorbing. Similarly to [138] one could show that for a weakly absorbing set it is sufficient to find a selector over $\text{sco } \Gamma$, rather than just over Γ itself, which corresponds to randomized control policies over MDPs, however this is not of our concern here when dealing with SSSs.

As promised, absorbing sets can be used to provide examples when the contractivity of truncated DP operators is violated, and in particular when the fixpoint equations do not have unique solutions. Note that in the case of the unconstrained reachability $G = S^c$, it holds that the operators r_S^Δ and r_S^∇ always admit the trivial fixpoint 1. However, if S is not simple (that is, if it admits absorbing subsets), then the optimal value functions are different than 1. For example, if a trajectory starts in an absorbing subset of S then it never reaches the goal set. More precisely:

Proposition 4.10 Let $A \in \mathcal{B}(X)$ be a subset of S . If A is

- i. *strongly absorbing* then $\mathcal{S}_x^\Delta(SUS^c) = 0$ for each $x \in A$. In particular, $w_\infty^\Delta(x) = 0$ for each $x \in A$, and $(r_S^\Delta)^n$ is not a contraction for any $n \in \mathbb{N}$.
- ii. *weakly absorbing* then $\mathcal{S}_x^\nabla(SUS^c) = 0$ for each $x \in A$. In particular, $w_\infty^\nabla(x) = 0$ for each $x \in A$, and $(r_S^\nabla)^n$ is not a contraction for any $n \in \mathbb{N}$.

Proof: With focus on case i. fix a point $x \in A$. Then $P_x^\sigma(\mathbf{x}_n \in A) = 1$ for all $n \in \mathbb{N}$ regardless of a strategy $\sigma \in \Sigma^\Gamma$. As a result, $P_x^\sigma(\mathbf{x}_n \in S^c) = 0$ for each $n \in \mathbb{N}$, so

$$Q_x^\sigma(SUS^c) \leq \sum_{n=0}^{\infty} P_x^\sigma(\mathbf{x}_n \in S^c) = 0$$

for each strategy $\sigma \in \Sigma^\Gamma$. Thus, we obtain that $\mathcal{S}_x^\Delta(SUS^c) = 0$ for each $x \in A$. Clearly, it follows immediately that $w_\infty^\Delta(x) = 0$ for each $x \in A$. Suppose now that $(r_S^\Delta)^n$ is contractive for some n . In such a case the solution of the fixpoint equation would be unique and hence it would imply that $w_\infty^\Delta = 1$, which is clearly not the case.

Let now A be a weakly absorbing set and consider a stationary strategy $\sigma' \in \Sigma_S^\Gamma$ with $\sigma'_0 = \kappa$ where κ is as per Definition 4.9. It follows that $P_x^{\sigma'}(\mathbf{x}_n \in A) = 1$ and hence $P_x^{\sigma'}(\mathbf{x}_n \in S^c) = 0$ for all $x \in A$, so

$$\mathcal{S}^\nabla(\cdot; SUS^c) \leq Q_x^{\sigma'}(SUS^c) \leq \sum_{n=0}^{\infty} P_x^{\sigma'}(\mathbf{x}_n \in S^c) = 0.$$

As for r_S^Δ , one can now show that $(r_S^\nabla)^n$ is not a contraction for any $n \in \mathbb{N}$. \square

In general the presence of absorbing sets is not the only reason that may violate contractivity. For example, it is easy to see that the set S in Example 4.7 does not have weakly absorbing subsets, and still the contractivity does not hold. However, the following assumption allows to characterize precisely the relationship between absorbing sets and contractivity.

Assumption 4.11 *The SS S is continuous and the set S is compact.*

We are going to show that, under Assumption 4.11, the case $m(S) < \infty$ precisely coincides with the case when S does not admit weakly absorbing sets. In order to prove this fact some preparation is required: let us define $S_n := \{v_n^\Delta = 1\}$ for each $n \in \mathbb{N}$. Note that for any $x \in S$ and $\sigma \in \Sigma$ the sequence $(v_n^\Delta(x))_{n \in \mathbb{N}}$ is non-increasing in n , hence so is $(S_n)_{n \in \mathbb{N}}$, i.e. $S_{n+1} \subseteq S_n$ for each $n \in \mathbb{N}$. Let us denote by $S_\infty := \bigcap_{n=0}^\infty S_n$ the limit of this sequence. We are now ready to formulate the main result connecting absorbing sets to the contractivity.

Theorem 4.12 *Under Assumption 4.11 the set S_∞ is weakly absorbing and coincides with $\{v^\Delta = 1\}$. Moreover, the following statements are equivalent:*

- i. *it holds that $m(S) < \infty$ (contractivity);*
- ii. *the operator Γ_S^Δ has a unique fixpoint (uniqueness);*
- iii. *it holds that $v^\Delta = 0$ (triviality);*
- iv. *it holds that $S_\infty = \emptyset$ (simplicity).*

Proof: We start with characterizing the set S_∞ . First of all, let us show that S_n satisfies

$$S_{n+1} = \{x \in S : \exists \gamma' \in \Gamma|_x \text{ s.t. } \gamma'(S_n) = 1\} \quad (4.19)$$

for any $n \in \mathbb{N}$. Indeed, if such γ' exists for a given $x \in S$, then

$$v_{n+1}^\Delta(x) = \sup_{\gamma \in \Gamma|_x} \gamma v_n^\Delta \geq \gamma' v_n^\Delta \geq \gamma(1_{S_n} v_n^\Delta) = 1,$$

and hence $x \in S_{n+1}$. Conversely, assume that the latter inclusion holds, then

$$v_{n+1}^\Delta(x) = \sup_{\gamma \in \Gamma|_x} \gamma(1_{S_n} v_n^\Delta) = 1. \quad (4.20)$$

Let us show that the maximum is achieved in (4.20). Notice that $v_0^\Delta = 1_S \in \text{b}\mathcal{C}^\Delta(X)$ since S is closed by being a compact subset of a metrizable space. If $v_n^\Delta \in \text{b}\mathcal{C}^\Delta(X)$ for some $n \in \mathbb{N}$, then $\sup_{\gamma \in \Gamma|_x} \gamma v_n^\Delta \in \text{b}\mathcal{C}^\Delta(X)$ by [20, Proposition 7.33] and as a result, $v_n^\Delta \in \text{b}\mathcal{C}^\Delta(X)$ by Lemma C.11. Since the supremum in (4.20) is taken over a $\text{b}\mathcal{C}^\Delta(X)$ function over a compact set, there exists $\gamma'' \in \Gamma|_x$ where this maximum is achieved. Hence $\gamma''(1_{S_n} v_n^\Delta) = 1$ which implies $\gamma''(S_n) = 1$ since $v_n^\Delta \leq 1$, so (4.19) is proved.

The case $S_\infty = \emptyset$ is trivial. Indeed, the empty set is weakly absorbing by definition, and if $v^\Delta(x) = 1$ for some x , then $x \in S_n$ for each $n \in \mathbb{N}$ and hence $x \in S_\infty$. Let us hence focus the case when S_∞ contains at least one element, some x . Define a sequence of sets $P_n = \{\gamma' \in \Gamma|_x : \gamma'(S_n) = 1\}$, then P_n is non-empty for each n , otherwise from (4.19) we would get that $x \notin S_{n+1}$, which contradicts the fact that $x \in S_\infty$. Furthermore, since S_n are non-increasing, so are P_n . Finally, since each S_n is a level set of a $\text{bc}^\Delta(X)$ function, it is closed, and compact as a subset of a compact set S . Hence, P_n is also compact for each $n \in \mathbb{N}$, and as a result $P_\infty := \bigcap_{n \in \mathbb{N}} P_n$ is non-empty. Let γ' be some element of P_∞ , then $\gamma'(S_n) = 1$ for each $n \in \mathbb{N}$ and hence by continuity of probability $\gamma'(S_\infty) = 1$. As a result, for each $x \in S_\infty$ there exists γ' such that the latter equality holds. Since S_∞ is a intersection of a non-increasing sequence of compact sets, it is compact, and hence by [20, Proposition 7.33] there exists a selector $\kappa \in \mathcal{B}(X|X)$ from Γ such that $\kappa(S_\infty|x) = 1$ for each $x \in S_\infty$, which casts the latter set to be weakly absorbing. Trivially, $v^\Delta = 1$ on S_∞ , and conversely if $v^\Delta(x) = 1$ for some x , then $x \in S_n$ for each $n \in \mathbb{N}$ and hence $x \in S_\infty$.

Let us now prove the equivalence result. The fact that $i. \implies ii.$ has been proven in Theorem 4.8. Also, $\Gamma_S^\Delta f = f$ always has a solution $f = 0$, so the uniqueness of a fixpoint of Γ_S^Δ implies $v^\Delta = 0$ and thus $ii. \implies iii.$ If $v^\Delta = 0$, then by just proven characterization of S_∞ it is empty, so $iii. \implies iv.$ Finally, if $m(S) = \infty$ then $\sup_{x \in X} v_n^\Delta(x) = 1$ for all $n \in \mathbb{N}$. Since each of the functions is in $\text{bc}^\Delta(X)$, the maximum over a compact set S is attained, so that $m(S) = \infty$ implies $S_n \neq \emptyset$ for all $n \in \mathbb{N}$. As we have seen, that casts $S_\infty \neq \emptyset$ and hence $iv. \implies i.$ \square

Note that the above results also implies that under Assumption 4.11 the set S_∞ is not only weakly absorbing itself, it also does contain any other weakly absorbing subset of S . We have obtained a rather precise characterization of the contractivity condition $m(S) < \infty$ in terms of presence or absence of weakly absorbing subsets of the safe set. In particular, if both Assumption 4.11 and the condition $S_\infty = \emptyset$ are satisfied, then regardless of the set G we are able to approximate $S_x^\Delta(\text{SUG})$ and $S_x^\nabla(\text{SUG})$ by their bounded horizon counterparts. Moreover, Theorem 4.12 also justifies the following intuitive statement: if one wants to keep the path of the process inside a set with some non-zero probability, there has to be an “attractor” within such set, which in our case appears to be the largest weak absorbing subset of S , that is S_∞ : If such attractor is absent, no matter what control policy is chosen, the path will leave the desired set almost surely – we discuss this in greater detail for autonomous SSS below in Section 4.3. The “if and only if” nature of Theorem 4.12 also implies that for the maximal safety problem such condition is necessary. However, it still may be the case that $S_\infty \neq \emptyset$ but r_S^∇ is a contraction, even though r^Δ is not. Although such cases are interesting to study, this goes beyond the scope of this thesis: we are now interested in techniques that allow us reducing the unbounded horizon problem to the bounded horizon one in the situation where $S_\infty \neq \emptyset$. These results are particularly powerful under the following assumption.

Assumption 4.13 *Stationary policies are sufficient for the solution of the constrained*

reachability problem on the unbounded time horizon, that is for any $x \in X$:

$$\mathcal{S}_x^\Delta(\text{SUG}) = \sup_{\sigma \in \Sigma_S^\Gamma} Q_x^\sigma(\text{SUG}), \quad \mathcal{S}_x^\nabla(\text{SUG}) = \inf_{\sigma \in \Sigma_S^\Gamma} Q_x^\sigma(\text{SUG}).$$

Before we provide the main result, the following technical lemma is needed.

Lemma 4.14 *Let $C \in \mathcal{B}(X)$ be any subset of S . Under Assumption 4.13, for each $x \in X$*

$$|\mathcal{S}_x(\text{SUG}) - \mathcal{S}_x((S \setminus C)\text{UG})| \leq \chi^*(C) := \sup_{\sigma \in \Sigma_S^\Gamma} \sup_{x' \in C} Q_{x'}^\sigma(\text{SUG}),$$

$$|\mathcal{S}_x^\nabla(\text{SUG}) - \mathcal{S}_x^\nabla((S \setminus C)\text{UG})| \leq \chi_*(C) := \inf_{\sigma \in \Sigma_S^\Gamma} \sup_{x' \in C} Q_{x'}^\sigma(\text{SUG}).$$

Proof: Let us fix an arbitrary policy $\sigma \in \Sigma_S^\Gamma$ and an arbitrary state $x \in X$. Denote $\chi^\sigma(C) := \sup_{x' \in C} Q_{x'}^\sigma(\text{SUG})$. Clearly, $S \setminus C \subseteq S$ means that $Q_x^\sigma(\text{SUG}) \geq Q_x^\sigma((S \setminus C)\text{UG})$. On the other hand

$$Q^\sigma(\text{SUG}) - Q^\sigma((S \setminus C)\text{UG}) \leq \chi^\sigma(C)$$

which together with Lemma C.10 immediately yields the desired result. \square

Let us discuss how Lemma 4.14 can be useful. Suppose that Assumption 4.11 holds true and that for the original problem we have that $S_\infty \neq \emptyset$, so that $m(S) = \infty$, and hence we cannot apply Theorem 4.8 to compute the optimal value functions. If we find a set $C \supseteq S_\infty$ such that $m(S \setminus C) < \infty$, then we can solve the unconstrained problem with truncated safe set $S \setminus C$. Also, since C contains S_∞ we can expect that $\chi^*(C)$ and $\chi_*(C)$ are close enough to zero, which would make the bounds in Lemma 4.14 useful. To further elaborate this idea we need the notion of a *locally excessive* function.

Definition 4.15 *Let $\kappa \in \mathcal{U}(X|X)$ be a selector from Γ . A non-negative function $g \in \text{b}\mathcal{B}(X)$ is called locally κ -excessive, if for any $x \in \{g \leq 1\}$ it holds that $\kappa g(x) \leq g(x)$. If in addition for some $A \in \mathcal{B}(X)$ we have that $A_\infty \subseteq \{g = 0\}$, $\{g \leq 1\} \subseteq A$ and $\{g < \varepsilon\}$ is an open set for all $\varepsilon > 0$, we say that g is locally κ -excessive on A .*

A non-negative function $g \in \text{b}\mathcal{B}(X)$ is called locally uniformly excessive if for any $x \in \{g \leq 1\}$ and $\gamma \in \Gamma_x$ it holds that $\gamma g \leq g(x)$. If in addition for some $A \in \mathcal{B}(X)$ we have that $A_\infty \subseteq \{g = 0\}$, $\{g \leq 1\} \subseteq A$ and $\{g < \varepsilon\}$ is an open set for all $\varepsilon > 0$, we say that g is locally uniformly excessive on A .

Theorem 4.16 *Let Assumptions 4.11 and 4.13 hold true. Suppose that g^Δ is locally uniformly excessive on S , and that g^∇ is locally σ'_0 -excessive for some $\sigma' \in \Sigma_S^\Gamma$. For any $\varepsilon \in (0, 1]$ the following inequalities are valid:*

$$\chi^*(\{g^\Delta < \varepsilon\}) \leq \varepsilon \quad \chi_*(\{g^\Delta < \varepsilon\}) \leq \varepsilon$$

and that sets $S \setminus \{g^\Delta < \varepsilon\}$, $S \setminus \{g^\nabla < \varepsilon\}$ are simple.

Proof: We start with the case of the maximization. For any policy $\sigma \in \Sigma_S^\Gamma$ we have that

$$Q_x^\sigma(XU\{g^\Delta > 1\}) \leq g^\Delta(x)$$

whenever $x \in \{g^\Delta \leq 1\}$, as it follows from [134, Lemma 3]. Furthermore, since $\{g \leq 1\} \subseteq S$ and $G \subseteq S^c$, it holds that $G \subseteq \{g^\Delta > 1\}$. As a result,

$$Q^\sigma(SUG) \leq Q^\sigma(XU\{g^\Delta > 1\}).$$

Combining both inequalities, we obtain that

$$\sup_{x' \in \{g^\Delta < \varepsilon\}} Q_{x'}^\sigma(SUG) \leq \varepsilon,$$

and thus after maximizing over all stationary policies we obtain that $\chi^\Delta(\{g^\Delta < \varepsilon\}) \leq \varepsilon$.

For the case of the minimization we similarly have

$$\sup_{x' \in \{g^\nabla < \varepsilon\}} Q_{x'}^{\sigma'}(SUG) \leq \varepsilon,$$

and since $\chi^\nabla(C) \leq \sup_{x' \in C} Q_{x'}^{\sigma'}(SUG)$ for any set $C \in \mathcal{B}(X)$, we immediately obtain that $\chi^\nabla(\{g^\Delta < \varepsilon\}) \leq \varepsilon$ for all $\varepsilon \leq 1$, as desired.

Finally, simplicity of $S \setminus \{g^\Delta < \varepsilon\}$ and $S \setminus \{g^\nabla < \varepsilon\}$ follows from the fact that they are compact. Indeed, S is compact whereas $\{g^\Delta < \varepsilon\}$ and $\{g^\nabla < \varepsilon\}$ are open. Moreover, the simplicity follows from the definition of functions locally excessive on S which implies that $S_\infty \subseteq \{g^\Delta < \varepsilon\}$ and $S_\infty \subseteq \{g^\nabla < \varepsilon\}$. \square

4.2.3 Comments on the reachability problem

Let us mention how the DP formulation has been developed for the (un)constrained reachability problem in the SS setting. To our knowledge, the first work with this goal has been [9], which has considered a class of models called controlled discrete-time Stochastic Hybrid Systems (cdt-SHS), namely a class of SS with a state space comprised of a collection of Borel subsets of \mathbb{R}^n . It has treated the unconstrained reachability property $\diamond_n G = \text{true}U_n G$ and the dual safety one $\square_n S = \neg \diamond_n S^c$, and has proposed their characterization using a maximal cost (4.3) for the first problem, and a multiplicative cost (4.4) for the second. Within this formulations, the DP recursion has been derived for the bounded time horizon $n < \infty$, while restricting the attention to Markov policies. [124] has addressed a more general⁵ constrained reachability problem $SU_n G$ within a simi-

⁵ Although the constrained reachability includes the unconstrained one as a special case, the latter can be used to solve the former if one just slightly modifies dynamic by making the set of unsafe sets $D = X \setminus (S \cup G)$ absorbing. Indeed, in such case $\diamond_n G$ is equivalent to $SU_n G$ since G is never reached by a trajectory that has visited D at least once [137, Proposition 1]. In particular, one immediately obtains [124, Theorem 8] by applying [9, Theorem 1] over a modified model. Similarly, rendering the set D absorbing allows one to recast a related terminal hitting-time reach-avoid problem [124, Section 4] as a special case of a terminal cost problem [70, Section 3].

lar setting: cdt-SHS models, Markov policies, and bounded time horizons: a new sum-multiplicative cost (4.5) has been proposed, leading to the DP scheme in [124, Theorem 8]. In contrast to these studies, here we have proposed a TC formulation, which has allowed dealing with non-Markovian policies, and to show that Markov policies are sufficient. In particular, one obtains [9, Theorems 1, 2] and [124, Theorem 8] as special cases of Theorem 4.5. At the same time, the TC formulation has also led to simpler proofs, which mostly rely on known results for the TC performance criterion [20, Chapters 8,9].

The case of the unbounded time horizon problem has received some attention already in [124, Section 3.3] and [3, Section V]. There it was suggested to use the convergence of the bounded-horizon values to the unbounded-horizon one, which led to considering the fixpoint equations. Although we have shown in Theorem 4.5 that fixpoint equations are indeed valid, they can not be obtained using limiting arguments as the latter may fail as shown in Example 4.7. An alternative approach via a hitting time formulation (4.6) has been proposed in [30], and the fixpoint equation for the maximal constrained reachability has been obtained in [30, Theorem 2.10 (i)]. However, one of the assumptions of this theorem required the first hitting time of the complement of the safe set τ_{S^c} to be almost surely finite for any Markov policy. As a result, in the case of the unconstrained reachability this theorem assumes that the value is fixed and constant. Finally, [78, Theorem 2] has shown the convergence of the *maximal* bounded-horizon unconstrained reachability to the unbounded-horizon one, and has showed that the latter satisfies the fixpoint equation. In contrast to the aforementioned contributions, Theorem 4.5 does not pose any limitations and establishes fixpoint equations for both the maximization and the minimization problems in generality, without for example requiring any continuity assumptions that are often imposed otherwise (cf. [78, Assumption 1] or [30, Assumption 2.9]). In addition, Proposition 4.6 provides a complete characterization of the convergence of bounded-horizon problems to the unbounded-horizon ones, and is further supported by Example 4.7. On a parallel note, related results for the reach-avoid problem – but for the continuous-time MDPs were obtained in [103] and [102].

The approximation of the unbounded-horizon reachability problem with bounded-horizon counterparts is an extension to the controlled case of the result in [134]. This extension requires no additional assumptions and (weak) continuity of the kernel \mathfrak{T} is sufficient to establish important results such as Theorems 4.12 and 4.16. At the same time, in the proofs we have extensively used continuity assumption, and so the equivalence in Theorem 4.12 may fail to hold without such assumptions – see e.g. [134, Appendix]. In particular, we acknowledge that [137, Proposition 2] is not correct: although uniqueness of fixpoint indeed yields trivial constant solutions for the maximal and minimal unconstrained reachability in the general case, without continuity assumptions it may happen that the solution is trivial but yet there are multiple fixpoints. In emphasizing the role of absorbing sets, it is crucial to use the connection between $m(S)$ and the contractivity of powers of the operators r^Δ and r^∇ in Theorem 4.8. In particular, as a special case we obtain [78, Proposition 1], which has obtained conditions for the contractivity in the special case $m(S) = 1$, but that would not be enough to obtain stronger results on the

connection with absorbing sets. The characterization of the absorbing sets, as well as finding an appropriate μ -excessive function, is an interesting and important problem. For example, there seems to be a connection between weakly absorbing sets (such as S_∞) and maximal controlled invariant sets in non-stochastic systems [113]. Another related concept is that of the maximal end component (MEC) [17, Section 10.6], which is used to solve both the reachability and the repeated reachability problems in the case of finite-state SS. Such techniques are extremely powerful and allow for the full solution of those problems, but unfortunately the discrete structure of the finite state and control spaces is crucial, and most of the nice properties MEC has are lost in the more general case of uncountable state spaces.

An alternative approach to the computation of the unbounded-horizon maximal reachability is in [78, Proposition 3], where it is proposed to recast the original fix-point equation as a linear constrained optimization over the infinite-dimensional space $\text{bU}(X)$, and to apply numerical methods for its solution. However, the uniqueness of the solution of this problem has not been addressed yet. Other possible alternatives are the theory of Poisson's equations [71, Chapter 7] and the theory of transient SS [71, Section 9.6], both of which should be applied over the truncated operator Γ_S^Δ . Another interesting way to approach this problem is to impose the ψ -irreducibility on the model and to tailor the results in [56, Chapter 10] developed for the AC performance criterion. All those extensions, however, are out of the scope of the present contribution.

4.3 Repeated reachability

4.3.1 Characterization

It follows from Lemma 2.9 that model-checking a SS against any property expressed as a DRA can be reduced to solving the Rabin-like conditions $\Box\Diamond F' \wedge (\neg\Box\Diamond F'')$ over the composition of the SS with the underlying transition system of the DRA. This result applies in particular to all ω -regular languages and LTL formulae. Unfortunately, we cannot provide a theory that is as comprehensive as for the reachability case (namely, for DFA or safe LTL specifications), as it has been presented in Section 4.2, and only focus on some partial results. In particular, we focus only on the case of the Büchi acceptance condition $\Box\Diamond F$, which is also easier to characterize by means of its dual property $\Diamond\Box S$, known as *persistence*. As mentioned in Section 2.5, we show how results developed in the setting of gambling theory apply to the SS case discussed here. This problem was studied in [94, 96]. In accordance with those works we call a function $f \in \text{bU}(X)$ *excessive*⁶ if $\Gamma^\Delta f \leq f$, *deficient* if $(-f)$ is excessive, and *invariant* if its both deficient and excessive. Clearly, invariant functions are precisely the fixpoints of the operator Γ^Δ .

⁶ Note that excessive functions are similar to locally uniformly excessive ones as per Definition 4.15.

The next result provides a characterization of the maximal persistence probability $\mathcal{S}^\Delta(\Diamond\Box S)$ and emphasizes its connection with the maximal safety probability $\mathcal{S}^\Delta(\Box S)$.

Theorem 4.17 *For any set $S \in \mathcal{B}(X)$ it holds that $\mathcal{S}^\Delta(\Diamond\Box S) \in \mathbf{bA}^\Delta(X)$. It is also an invariant function, and for any excessive function $f \in \mathbf{bA}^\Delta(X)$ satisfying the inequality $f \geq \Gamma^\Delta \mathcal{S}^\Delta(\Box S)$ it holds that $f \geq \mathcal{S}^\Delta(\Diamond\Box S)$. Moreover, the following DP holds true:*

$$\mathcal{S}_x(\Diamond\Box S) = \lim_{n \rightarrow \infty} (\Gamma^\Delta)^n \mathcal{S}_x(\Box S), \quad (4.21)$$

where the limit is non-increasing point-wise, for all $x \in X$.

Proof: The result follows immediately from [94, Theorem 1.2] and [96, Theorem 4.5, Corollary 5.5]. \square

Note that Theorem 4.17 connects the maximal safety probability $\mathcal{S}^\Delta(\cdot; \Box S)$ and the maximal persistence probability $\mathcal{S}^\Delta(\Diamond\Box S)$. As a result, we can use results on the former function obtained in Section 4.2 to derive properties of the latter one.

Proposition 4.18 $\mathcal{S}^\Delta(\Box S) = 0$ if and only if $\mathcal{S}^\Delta(\Diamond\Box S) = 0$

Proof: Note that (4.21) immediately implies that $\mathcal{S}^\Delta(\Diamond\Box S) = 0$ is sufficient to claim that $\mathcal{S}^\Delta(\Box S) = 0$. On the other hand, since $\mathcal{S}^\Delta(\Diamond\Box S) \geq \mathcal{S}^\Delta(\Box S)$, thanks to Theorem 4.17 we obtain the converse implication. \square

4.3.2 Computation

Although the recursions in (4.21) already suggest a possible computational procedure for computing the value of the maximal probability of persistence $\mathcal{S}^\Delta(\Diamond\Box S)$, the scheme requires an infinite number of iterations that are initialized at the maximal safety probability $\mathcal{S}^\Delta(\Box S)$, which in turn has to be computed in advance. For the latter quantity we have already discussed non-trivial issues in Section 4.2, so the result of Theorem 4.17 is not in general practically applicable. Instead, we propose tailoring the technique developed in Theorem 4.16 to the problem at hand. Since the latter result is obtained for stationary policies, we can in fact focus on the autonomous case and provide a rather comprehensive characterization. Unfortunately, extending the results below to the general case is a rather challenging task.

Recall that in autonomous systems Γ is a graph of some Borel map, which we also denote by Γ , i.e. in each $x \in X$ we can only choose a single distribution $\Gamma(x) \in \mathcal{P}(X)$. As stated above, absorbing sets are crucial when dealing with infinite-horizon problems over SSs, and they also hint upon certain convergence properties of sample paths of SSs. We start with discussing the latter concept.

Definition 4.19 For a sequence $(x_n)_{n \in \mathbb{N}}$ of elements of X , we say that $x_n \rightarrow A$ with $n \rightarrow \infty$ whenever $\lim_n \rho(x_n, A) = 0$. If for $f \in \mathcal{B}(X)$ it holds that $f|_A = c \equiv \text{const}$, then the notation $\lim_{x \rightarrow A} f(x) = c$ denotes that $f(x_n) \rightarrow c$ for any $x_n \rightarrow A$. Equivalently, for any $\varepsilon > 0$ there is $\theta > 0$ such that $|f(x) - c| < \varepsilon$ for all $x \in A^\theta$.

The definition of an attractive set $A \subseteq X$ for a classical deterministic dynamical system [109] requires the set A to satisfy the following conditions: 1) to be closed, 2) to be invariant under the flow of the system and 3) to have a neighborhood with the property that the system, starting within this neighborhood, never leaves it and converges to the set A . The set of all points in X which are initial points for trajectories converging to A is called the domain of attraction of A . For SSs such requirements may be in general too conservative, which inspired us to introduce the following definition.

Definition 4.20 We call a set $A \subseteq X$ stochastically attractive for \mathcal{S} if

$$\lim_{x \rightarrow A} P_x(\mathbf{x}_m \rightarrow A) = 1. \quad (4.22)$$

The equation (4.22) means that selecting an initial condition x closer to A makes the probability of the event that “the process X converges to A ” closer to 1 (below we show that this probability is in fact well-defined). Thus, Definition 4.20 is a modified version of the condition 3) above. Also, it captures the closure property: if (4.22) holds for the set A then it holds as well for its closure \bar{A} . However, it does not imply that within some neighborhood of set A the probability to stay always within this neighborhood is close to 1 (safety). Moreover, the notion of being invariant required in deterministic systems and here interpreted as in Definition 4.9 is too restrictive for the discrete-time stochastic framework (we enlighten it below in Example 4.34) — this is the reason why such a requirement is not included in Definition 4.20.

The definition of stochastic attractivity leads to studying probabilities related to events, as in (4.22). In contrast to deterministic dynamical systems, where the study is often restricted to the attractivity of equilibria and closed orbits, in this work we embrace a more general approach which first selects potentially attractive subsets of X and thereafter verifies their attractivity. A special focus is given to compact subsets of X , thus excluding possible divergent behaviors of the process X . Due to this reason we introduce appropriate value functions that depend both on the point $x \in X$ and on the set $A \in \mathcal{B}$. In line with the notation of Section 4.2, we let $v(x; A) = \mathcal{S}_x(\Box A)$ and $w(x; A, B) = \mathcal{S}_x(A \cup B)$ for each $A, B \in \mathcal{B}(X)$.

We also introduce a value function for the convergence event: $\{\mathbf{x}_n \rightarrow A\}$. Clearly, unlike the properties above, it does not have a finite horizon version. We start by defining another, closely related event. By direct application of the definition of limit, we have that $x_n \rightarrow A$ if and only if for any $m \in \mathbb{N}$ there is $N \in \mathbb{N}$, such that $x_n \in A^{1/m}$, for all $n \geq N$. We call the property that, for some $B \subseteq X$ there is an N such that $x_n \in B$ for all $n \geq N$, “ x_n is eventually always in B ,” and denote it as $x_n \in B$. Now, if $x_n \in B$ then obviously $x_n \rightarrow B$ and $x_n \rightarrow B$ if for any $m \in \mathbb{N}$, $x_n \in B^{1/m}$.

For a set $A \in \mathcal{B}(X)$ the event $\{\mathbf{x}_n \in A\}$ is measurable. Indeed, it is first the complement of the event “ \mathbf{x}_n visits A^c infinitely often,” which is used to characterize properties such as transience and recurrence, and is proven to be measurable [100]. Secondly, this event is invariant, i.e. it is independent on any finite prefix of the sequence \mathbf{x}_n . For more detailed discussion, see [100, Section 15] and [112, Section 3.3]. As a result, it is legitimate to introduce the value function defined as $h(x; A) := P_x(\mathbf{x}_n \in A)$. Let us now introduce a value function for the convergence probability: $c(x; A) = P_x(\mathbf{x}_n \rightarrow A)$.

$$\{\mathbf{x}_n \rightarrow A\} = \bigcap_{m=0}^{\infty} \{\mathbf{x}_n \in A^{1/m}\},$$

hence the probability in (4.22) is well defined. Furthermore:

$$c(x; A) = \lim_m h\left(x; A^{1/m}\right), \quad (4.23)$$

for any $x \in X$ and $A \in \mathcal{B}(X)$. As a result, (4.22) is equivalent to

$$\lim_{x \rightarrow A} c(x; A) = 1. \quad (4.24)$$

Functions c and h play a prominent role below.

4.3.3 Characterization through harmonic functions

We provide a formal way to derive h through the safety value function v . First of all, let us define a safety operator i_A on $\mathcal{B}_1(X)$ by $i_A f(x) = 1_A(x) \Gamma f(x)$, so from Section 4.2

$$v(x; A) = i_A v(x; A). \quad (4.25)$$

We say that function $f \in \mathcal{B}(X)$ is superharmonic if $f \geq \Gamma f$, subharmonic if $f \leq \Gamma f$, and harmonic if $f = \Gamma f$. Clearly, the function v is subharmonic: if $x \in A$ then

$$v(x; A) = i_A v(x; A) = \Gamma v(x; A),$$

whereas if $x \in A^c$, then $v(x; A) = 0 \leq \Gamma v(x; A)$.

Theorem 4.21 *For all $x \in X$ and $A \in \mathcal{B}(X)$ it holds that*

$$h(x; A) = \lim_n \Gamma^n v(x; A), \quad (4.26)$$

so that $h = \Gamma h$ and $c = \Gamma c$. Moreover, if f is a harmonic function and $f \geq v$ then $f \geq h$, i.e. h is the smallest harmonic majorant of the function b . In particular,

$$\inf_{x \in X} |h(x; A) - v(x; A)| = 0. \quad (4.27)$$

Proof: First of all

$$\{\mathbf{x}_n \in A\} = \bigcup_{n=0}^{\infty} \left\{ \prod_{k=n}^{\infty} 1_A(\mathbf{x}_k) = 1 \right\},$$

so

$$\begin{aligned} h(x; A) &= P_x(\mathbf{x}_n \in A) = P_x \left(\bigcup_{n=0}^{\infty} \left\{ \prod_{k=n}^{\infty} 1_A(\mathbf{x}_k) = 1 \right\} \right) \\ &= \lim_n P_x \left(\prod_{k=n}^{\infty} 1_A(\mathbf{x}_k) = 1 \right) \\ &= \lim_n \Gamma^n \left(P_x \left(\prod_{k=0}^{\infty} 1_A(\mathbf{x}_k) = 1 \right) \right) = \lim_n \Gamma^n v(x; A), \end{aligned}$$

which proves the first part. Now notice that for all $x \in X$ and $A \in \mathcal{B}(X)$

$$\begin{aligned} \Gamma h(x; A) &= \int_X \lim_n \Gamma^n v(y; A) \Gamma(dy|x) \\ &= \lim_n \int_X \Gamma^n v(y; A) \Gamma(dy|x) \\ &= \lim_n \Gamma^{n+1} v(x; A) = h(x; A) \end{aligned}$$

where we used dominated convergence theorem [50] to interchange the limit and the integral operators.

Let now f be a harmonic function such that $f \geq v$. The operator Γ is clearly monotone, so $f = \Gamma^n f \geq \Gamma^n v$ for all $n \in \mathbb{N}$. The limit $n \rightarrow \infty$ yields that $f \geq h$, so h is the least harmonic majorant of v . Suppose that

$$\varepsilon = \inf_{x \in X} |h(x; A) - v(x; A)| > 0$$

then $h(x; A) \geq v(x; A) + \varepsilon$ and hence $f(x) := h(x; A) - \varepsilon \geq v(x; A)$ being in addition a harmonic function. But it contradicts with the fact that $f \geq h$, so (4.27) holds.

Finally, with regards to the function c we have

$$\begin{aligned} \Gamma c(x; A) &= \Gamma \lim_{m \rightarrow \infty} h(x; A^{1/m}) \\ &= \lim_{m \rightarrow \infty} \Gamma h(x; A^{1/m}) = c(x; A) \end{aligned}$$

since functions h are harmonic. □

Let us discuss the result above. First, the convergence $\Gamma^n v \rightarrow h$ is monotonically non-decreasing: $\Gamma u \geq u$, because v is subharmonic, furthermore $\Gamma^{n+1} v \geq \Gamma^n v$, because Γ is a monotone operator. On the other hand, the convergence may not be uniform, which makes it difficult to find the bounds on the quantity $\|\Gamma^n v - h\|$.

Second, the equation $h = \Gamma h$ is called the *Laplace equation* and admits infinitely many solutions, e.g. any constant function. Finally, although h is the least harmonic majorant of v and (4.27) holds, there are cases when $h(x; A) > v(x; A)$ for all $x \in X$.

These remarks emphasize possible difficulties related to finding the function h . Moreover, recall from above that the problem of finding the safety function v (and thus $\Gamma^n v$) is difficult by itself. On the other hand, although the function h in general cannot be computed with any accuracy, we show that there are cases when this problem has a solution. Also, here we provide techniques that allow one to find the value of h directly, without calculating v in advance.

With focus on the function c , its evaluation is even more difficult: from a computational perspective view we have that

$$c(x; A) = \lim_i \lim_j \lim_k (\Gamma^j i_{A^{1/i}}^k) 1_{A^{1/i}}(x), \quad (4.28)$$

for all $x \in X, A \in \mathcal{B}(X)$. In order to tackle this problem, we show how to eliminate the limit with respect to i , which leads to the calculation of function h .

Although it is in general hard to find an analytical expression for $h(x; A)$, given $x \in X$ and an $A \in \mathcal{B}(X)$, in some cases it is possible. We characterize such instances using absorbing sets: in autonomous SSs there is no difference between weakly and strongly absorbing sets. They are analogues of equilibrium points, closed orbits and more generally, of invariant manifolds for classical deterministic dynamical systems [109]. The next lemma further highlights this similarity.

Definition 4.22 *An autonomous SS \mathcal{S} is called continuous if Γf is continuous for any continuous function f .*

Lemma 4.23 *If \mathcal{S} is continuous and A is absorbing then \overline{A} is absorbing.*

Proof: Note that $\Gamma(A|x) = \Gamma 1_A(x) = 1$ for all $x \in A$, so it is sufficient to prove that $\Gamma(A|x)$ is continuous on \overline{A} . Define

$$g(x) = \min \{ \rho(x, \overline{A}), 1 \}$$

so $g \in \mathcal{C}(X)$. Put $f_n(x) = (1 - g(x))^n$ then $f_n(x) \in [0, 1]$ for all $x \in X$ and $n \in \mathbb{N}$ and $f_n \in \mathcal{C}(X)$. Moreover, $f_n(x) = 1$ for $x \in \overline{A}$ and $f_n \downarrow 1_{\overline{A}}$ pointwise.

Now, since $f_n(x) \geq 1_A(x)$ then $\Gamma f_n(x) \geq \Gamma 1_A(x) = 1$ for all $x \in A$. Since \mathcal{S} is continuous, a function $\Gamma f_n \in \mathcal{C}(X)$ and hence $\Gamma f_n(x) \geq 1$ for all $x \in \overline{A}$. By monotone convergence theorem we obtain that $\Gamma f_n(x) \downarrow \Gamma 1_{\overline{A}}(x) = \Gamma(\overline{A}|x)$, so $\Gamma(\overline{A}|x) = 1$ for all $x \in \overline{A}$, which proves the statement of the lemma. \square

We denote with $\text{l.a.s.}(A)$ the largest absorbing subset of a given set $A \in \mathcal{B}(X)$. In contrast to the general case, for autonomous SSs it is always well-defined and is provided by $A_\infty = \{v(\cdot; A) = 1\}$. We further introduce the following notation: for any $f \in \text{b}\mathcal{B}(X)$ we define $m_f^\Delta = \{f = \sup_{x \in X} f(x)\}$ and $m_f^\nabla = \{f = \inf_{x \in X} f(x)\}$.

Lemma 4.24 *For a superharmonic f , the set m_f^∇ is absorbing.*

Proof: If m_f is empty, it is absorbing by the definition. Hence, we assume that there is at least one $x \in m_f^\nabla$. We have

$$0 \leq f(x) - \Gamma f(x) = \int_X \left(\inf_{x'' \in X} f(x'') - f(x') \right) \Gamma(dx'|x) \leq 0$$

where the left inequality holds since f is superharmonic and the right one holds because $\inf_{x'' \in X} f(x'') \leq f(x)$ for all $x \in X$. As a result, $\Gamma(m_f^\nabla|x) = 1$ as needed. \square

As a corollary, we have that for a subharmonic function f the set m_f^Δ is absorbing, and furthermore that if f is harmonic both m_f^∇ and m_f^Δ are absorbing. Let us show how these results are employed to find function h .

Theorem 4.25 *For any set $A \in \mathcal{B}$ the function $v(\cdot; A)$ is harmonic iff $m_v^\nabla = \{v(\cdot; A) = 0\}$ is absorbing. In that case $h(\cdot; A) = v(x; A)$.*

Proof: If v is harmonic then m_v^∇ is absorbing by Lemma 4.24. On the other hand, let m_v^∇ be absorbing. For each $x \in A$ we have $v(x; A) = \Gamma v(x; A)$. Now let $x \in A^c$, then $x \in m_v^\nabla$ and so

$$\Gamma v(x; A) = \int_X v(x'; A) \Gamma(dx'|x) = \int_{m_v^\nabla} v(x'; A) \Gamma(dx'|x) = 0$$

so $0 = v(x; A) = \Gamma v(x; A)$ and hence v is harmonic, so v is the least harmonic majorant of itself and hence $h = v$. \square

Theorem 4.25 implies that if the set m_v^∇ is absorbing then the problem of finding h is reduced to that of finding v . Although the analytical expression for v is in general hard to obtain and thus the verification of absorbance of m_v^∇ is not a trivial problem, there exist cases with an analytical solution.

We show now which measurable subsets of X are essentially not attractive. First, since equation (4.25) is linear, it always admits a trivial zero solution, though this may happen when $v(\cdot; A)$ is not a constant zero function – a simple example being $A = X$. From Theorem 4.12 we have that $\text{l.a.s.}(A) = \emptyset$ iff $v(\cdot; A) = 0$ under the assumption that \mathcal{S} is continuous A is compact. We can now formulate criteria for $h(x; A)$ and $c(x; A)$ to be constantly zero based on the simplicity of the set A .

Theorem 4.26 *For a set $A \in \mathcal{B}(X)$ it holds that:*

- i. $h(\cdot; A) = 0$ iff $v(\cdot; A) = 0$;
- ii. in particular, if \mathcal{S} is continuous and A is compact, then $h(\cdot; A) = 0$ iff A is simple;

- iii. if \mathcal{S} is continuous, X is locally compact and A is compact then $c(\cdot; A) = 0$ iff A is simple.

Proof: The proof is as follows:

- i. If $v(\cdot; A) = 0$ then it is harmonic and hence $h = v = 0$. On the other hand, since $h(\cdot; A) \geq v(\cdot; A) \geq 0$, the fact that $h = 0$ implies $v = 0$.
- ii. If \mathcal{S} is continuous and A is compact, $v(\cdot; A) = 0$ iff if A is simple by Theorem 4.12. Hence in that case *ii.* easily follows from *i.*
- iii. Clearly, $\{x_n \in A\} \subseteq \{x_n \rightarrow A\}$, thus $h(\cdot; A) \leq c(\cdot; A)$. If $c(x; A) = 0$ then $h(\cdot; A) = 0$ and by *ii.* the set A is simple.

Conversely, assume that A is a compact simple set. Let us denote $h_m(x) = h(x; A^{1/m})$ so $c(x; A) = \lim_m h_m(x)$ for all $x \in X$, and the idea is to show that $h_m = 0$ for m big enough, provided that A is simple. First, we show that there is $M > 0$ such that $\overline{A^{1/M}}$ is a compact set. Since X is locally compact, each $x \in A$ has a compact neighborhood, so let $\varepsilon : A \rightarrow \mathbb{R}$ be such that $\{x\}^{\varepsilon(x)}$ is contained in some compact set. We have that $\{\{x\}^{\varepsilon(x)} : x \in A\}$ is an open cover of A , thus there is an open subcover $\{\{x_k\}^{\varepsilon_k}\}_{k \leq n}$ with $\varepsilon_k := \varepsilon(x_k)$. Now, $\{x_k\}^{\varepsilon_k}$ is compact for each $k \leq n$, as closed subsets of compact sets. Hence the set

$$C = \bigcup_{k \leq n} \overline{\{x_k\}^{\varepsilon_k}}$$

is compact and $\overline{A^{\varepsilon'}} \subseteq C$ where $\varepsilon' = \min_{k \leq n} \varepsilon_k$. We only need now to pick up $M > \frac{1}{\varepsilon'}$, then $\overline{A^{1/M}}$ is a compact as a closed subset of a compact C . For $m \geq M$ we denote $B_m = \overline{A^{1/m}}$ and $B'_m = \text{l.a.s.}(B_m)$. Then we have $A^{1/(m+1)} \subseteq B_{m+1} \subseteq A^{1/m}$ and so $h_m(x) \leq h(x; B_m)$ for all $x \in X$.

Second, let us show that there is $M' \geq M$ such that B_m is simple for all $m \geq M'$. Suppose contrary: namely that $B'_m \neq \emptyset$ for all $m \geq M$. B'_m are compact sets so $B' := \bigcap_{m=M}^{\infty} B'_m \subseteq A$ is not empty. For any $x \in B'$ it holds that $x \in B'_m$ for all $m \geq M$, thus $\Gamma(B'_m|x) = 1$. We have

$$\Gamma(B'|x) = \Gamma\left(\bigcap_{m=M}^{\infty} B'_m \middle| x\right) = \lim_m \Gamma(B'_m|x) = 1$$

which means that the set B' is a non-empty absorbing subset of A and contradicts with the simplicity of A .

Finally, since for some $M' \geq M$ sets $(B_m)_{m \geq M'}$ are compact and simple, $h(x; B_m) = 0$ by *ii.* On the other hand, $0 = h(x; B_m) \geq h_m(x)$ and so $c(x; A) = \lim_m h_m(x) = 0$.

□

Corollary 4.27 *It follows that:*

- i. *If $A \in \mathcal{B}(X)$ and for some $\theta > 0$ the set A^θ is trivial, then A is not stochastically attractive.*
- ii. *If X is locally compact, S is continuous and A is compact and simple then A is not stochastically attractive.*

Proof: The proof is as follows:

- i. For all $m \geq 1/\theta$ we have $h(x; A^{1/m}) = 0$ and hence $c(\cdot; A) = 0$. As a result, (4.24) does not hold for A .
- Follows directly from the statement *iii.* of Theorem 4.26.

□

Corollary 4.27 provides conditions for sets not to be stochastically contractive. Although the problem of verification of simplicity or triviality of a given set A does not have a general (respectively analytical or computational) solution, there exist sufficient conditions. The first (analytical) conditions require super- or subharmonic functions to be constants, implying that $v(\cdot; A) = 0$ for all $A \neq X$. The second (computational) conditions require that A_n is empty for some $n \in \mathbb{N}$ [129, Theorem 2].

So far we have discussed sets that do not satisfy the given definition of stochastic attractivity. Next, the attention is shifted over a class of sets that satisfies it. We start with the following useful result. Notice that $\lim_n v(\mathbf{x}_n; A)$ exists P_x -a.s. for all $x \in X$, since v is a bounded subharmonic function, hence the existence of the limit is insured by the martingale convergence theorem [50].

Lemma 4.28 *For any $x \in X$ and $A \in \mathcal{B}(X)$ we have*

$$P_x \left(\lim_n v(\mathbf{x}_n; A) = \liminf_n 1_A(\mathbf{x}_n) \right) = 1, \quad P_x \left(\lim_n h(\mathbf{x}_n; A) = \liminf_n 1_A(\mathbf{x}_n) \right) = 1.$$

In particular,

$$h(x; A) = P_x \left(\lim_n v(\mathbf{x}_n; A) = 1 \right) = P_x \left(\lim_n h(\mathbf{x}_n; A) = 1 \right).$$

Proof: We define an operator q on $b\mathcal{B}(X)$ as $qf(x) = \max\{f(x), \Gamma f(x)\}$. From [116, Lemma 6, Lemma 8, p. 43] it follows that if $f \in b\mathcal{B}(X)$ and $g(x) = \lim_{n \rightarrow \infty} q^n g(x)$ then

$$P_x \left(\lim_n g(\mathbf{x}_n) = \limsup_n f(\mathbf{x}_n) \right) = 1. \quad (4.29)$$

We take $f = 1_{A^c}$ gives us the first equality of the lemma. Furthermore, $h(x; A) = \lim_n \Gamma^n v(x; A) = \lim_n q^n v(x; A)$ since v is subharmonic, $\Gamma^n v$ are subharmonic for all

$n \in \mathbb{N}$. Hence, $q\Gamma^n v = \Gamma^{n+1}v$ for all $n \in \mathbb{N}$ and we obtain the second equality of the lemma from (4.29). For any x , \mathbb{P}_x -a.s. we have that $\lim_n v(\mathbf{x}_n; A) = 1$ and $\lim_n h(\mathbf{x}_n; A) = 1$ iff $\liminf_n 1_A(\mathbf{x}_n) = 1$ and the latter statement mean $\{\mathbf{x}_n \in A\}$. Recalling of the definition of h completes the proof. \square

Lemma 4.29 *If \mathcal{S} is continuous and A is a compact set, then $\lim_n v(x_n; A) \rightarrow 1$ implies $x_n \rightarrow \text{l.a.s.}(A)$, for any sequence $(x_n)_{n \in \mathbb{N}}$ of elements of X .*

Proof: Suppose, that $\lim_n v(x_n; A) = 1$ and $x_n \not\rightarrow A$, i.e.

$$\limsup_n \rho(x_n, \text{l.a.s.}(A)) > 0.$$

Since $v(x; A) = 0$ for all $x \in A^c$, we have $x_n \in A$. Compactness of A implies that there is a convergent subsequence $x'_n \rightarrow x' \in A$ such that $\lim_n \rho(x'_n, \text{l.a.s.}(A)) > 0$ but $\lim_n v(x'_n; A) = 1$. Since $v \in \text{b}\mathcal{C}^\Delta(X)$ we get that $v(x'; A) = 1$ and hence $x' \in \text{l.a.s.}(A)$, but $\rho(x', \text{l.a.s.}(A)) > 0$, which leads us to a contradiction. \square

Lemmas 4.28 and 4.29 show that under assumptions of continuity of \mathcal{S} and compactness of A , $\mathbf{x}_n \rightarrow \text{l.a.s.}(A)$ is a necessary condition for $\mathbf{x}_n \in A$, \mathbb{P}_x -a.s. for all $x \in X$. On the other hand, this is not a sufficient condition in general. Due to this reason we introduce the concept of *stable absorbing set*.

Definition 4.30 *An absorbing set $A \in \mathcal{B}$ is called stable if there exists a compact neighborhood U_A of A such that $A = \text{l.a.s.}(U_A)$ and*

$$\lim_{x \rightarrow A} v(x; U_A) = 1. \quad (4.30)$$

Remark 4.31 *The stability property of absorbing sets can be related to the Lyapunov stability for classical deterministic dynamical systems [109]. Indeed, (4.30) means that if A is a stable absorbing subset, then for any $\varepsilon > 0$ there is a neighborhood of A starting from which the process never leaves such neighborhood with a probability at least $1 - \varepsilon$.*

The compactness of set U_A plays a role in the following result.

Theorem 4.32 *If \mathcal{S} is continuous and A is a stable absorbing set, then A is stochastically attractive and there exists $M \in \mathbb{N}$ such that*

$$h\left(x; \overline{A^{1/m}}\right) = h(x; U_A) = c(x; A)$$

for all $x \in X$ and $m \geq M$.

Proof: Let $A \in \mathcal{B}$ be a stable absorbing set and U_A be as in Definition 4.30. Since U_A is compact and $A = \text{l.a.s.}(U_A)$, A is compact too. Moreover, since U_A is a compact neighborhood of A , each $x \in A$ has a compact neighborhood and similarly

to the proof of Theorem 4.26 we pick up M such that $\overline{A^{1/M}}$ is a compact set. Thus for all $m \geq M$ the set $\overline{A^{1/m}}$ is compact.

Let us consider an arbitrary $m \geq M$. By Lemma 4.29 we obtain that $v(x_n; \overline{A^{1/m}}) \rightarrow 1$ implies $x_n \rightarrow A$. Let us show that the converse statement is also true. Since X is a metric space, it is equivalent to show that for any $\varepsilon > 0$ there is $\theta(\varepsilon) > 0$ such that $\rho(x, A) < \theta(\varepsilon)$ implies

$$v(x; \overline{A^{1/m}}) \geq 1 - \varepsilon. \quad (4.31)$$

We fix $\varepsilon > 0$ and denote $f(x) := 1 - v(x; U_A)$. Since f is a superharmonic function with a range in $[0, 1]$, the process $(f(\mathbf{x}_n))_{n \in \mathbb{N}}$ is a non-negative P_x -supermartingale for all $x \in X$ [116]. Hence, the Doob's inequality [50] holds:

$$P_x \left(\sup_{n \in \mathbb{N}} f(\mathbf{x}_n) > r \right) \leq \frac{1}{r} f(x) \quad (4.32)$$

for all $x \in X$ and $r > 0$. In the level sets notation, the inequality (4.32) takes the form $v(x; \{f \leq r\}) \geq 1 - \frac{1}{r} f(x)$.

Let us show that the stability of A implies an existence of $r > 0$ such that $\{f \leq r\} \subseteq \overline{A^{1/m}}$. Indeed, if it was true, we would be able to pick up a sequence $x_k \notin \overline{A^{1/m}}$ such that $f(x_k) \leq 1/k$. Clearly,

$$\lim_k v(x_k; U_A) = 1 - \lim_k f(x_k) = 1$$

but $\rho(x_k, A) \geq \frac{1}{m}$ which contradicts with Lemma 4.29.

It follows from the existence of r that

$$v(x; \overline{A^{1/m}}) \geq v(x; \{f \leq r\}) \geq 1 - \frac{1}{r} f(x).$$

Leveraging the stability of A again, we obtain that there exists $\theta > 0$ such that $f(x) \leq r\varepsilon$ for all $x \in A^\theta$, and hence for all such x the inequality (4.31) holds.

As a result, for any $m \geq M$ we obtain that $v(x_n; \overline{A^{1/m}}) \rightarrow 1$ iff $x_n \rightarrow A$. Since

$$h(x; \overline{A^{1/m}}) = P_x \left(v(\mathbf{x}_n; \overline{A^{1/m}}) \rightarrow 1 \right) = P_x(\mathbf{x}_n \rightarrow A)$$

we obtain that

$$h(x; \overline{A^{1/m'}}) = h(x; \overline{A^{1/m''}}).$$

for all $m', m'' \geq M$ and $x \in X$. Furthermore, since

$$h(x; \overline{A^{1/m}}) \geq h(x; A^{1/m}) \geq h(x; \overline{A^{1/(m+1)}})$$

for all $m \geq M$ and $x \in X$, we obtain that

$$c(x; A) = \lim_m h\left(x; \overline{A^{1/m}}\right) = h\left(x; \overline{A^{1/m'}}\right)$$

for all $m' \geq M$ and $x \in X$. Moreover, $v(x_n; U_A) \rightarrow 1$ iff $x_n \rightarrow A$, so $c(x; A) = h(x; U_A)$.

To complete the proof of the theorem we observe that $h\left(x; \overline{A^{1/m'}}\right) \geq v\left(x; \overline{A^{1/m'}}\right)$, and for the latter we proved that it converges to 1 on any sequence which converges to A . As a result, A is stochastically attractive. \square

Let us discuss examples showing that some of the conditions we have provided are sufficient but not necessary in general. Let $X = \{\pm \frac{1}{n}\}_{n \in \mathbb{N}} \cup \{0\}$ be endowed with the Euclidean metric, which makes it a complete separable compact (and locally compact) metric space. We first show that the reverse statement of Theorem 4.32 does not hold.

Example 4.33 (A stochastically attractive absorbing set is not necessary stable) Let $\Gamma(\{0\}|0) = 1, \Gamma(\{-1\}|1) = 1, \Gamma\left(\left\{\frac{1}{n-1}\right\} \middle| \frac{1}{n}\right) = 1$ for all $n \in \mathbb{N} \setminus \{1\}$ and $\Gamma\left(\left\{-\frac{1}{n+1}\right\} \middle| -\frac{1}{n}\right) = 1$ for all $n \in \mathbb{N}$. The corresponding dynamics is clearly deterministic, still it is a continuous SS. This process converges to an absorbing set $A = \{0\}$ starting from any initial condition, so $c(\cdot; A) = 1$ and hence A is stochastically attractive. However, there is no such neighborhood U_A of A such that $A = \text{l.a.s.}(U_A)$.

The second example shows that there may exist a set A such that $c(\cdot; A) = 1$, which in particular means that A is stochastically attractive, but it is not absorbing. This fact relates to the discussion given after Definition 4.20.

Example 4.34 (A stochastically attractive set is not necessary absorbing) In the previous example we only change $\Gamma(\{1\}|0) = 1$. We still have $A = \{0\}$ stochastically attractive since $c(x; A) \equiv 1$ but now A is not absorbing. Note that the update in the dynamics leads to the lack of continuity of the process, thus this assumption cannot be relaxed in Theorem 4.26, statement iii.

Theorem 4.32 shows that the stability of an absorbing set under mild conditions implies its stochastic attractivity. Moreover, it helps eliminating the outermost limit in the computation of the function c as in (4.28). The question is hence in how to show a stability of a given absorbing set. We introduce the following Lyapunov-like functions for that purpose.

Lemma 4.35 *An absorbing set A is stable iff the following stabilizing pair exists:*

- a compact neighborhood U_A of A such that $A = \text{l.a.s.}(U_A)$;
- a function $f \in \text{bB}(X)$ such that $\inf_{x \in X} f(x) = 0, m_f^\nabla = A, \lim_{x \rightarrow A} f(x) = 0$ and there exists $r > 0$ such that $\{f \leq r\} \subseteq U_A$ and $f(x) \geq \Gamma f(x)$ for all $x \in \{f \leq r\}$.

If \mathcal{S} is continuous, A is stable iff there is a stabilizing pair with an l.s.c. function f .

Proof: Suppose that such f and U_A exist. From the local version of Doob's inequality [129, Theorem 6] it follows that

$$v(x; \{f \leq r\}) \geq 1 - \frac{1}{r}f(x)$$

for all $x \in X$. As a result, $\lim_{x \rightarrow A} v(x; U_A) = 1$ since $v(x; U_A) \geq v(x; \{f \leq r\})$.

Now, let A be a stable absorbing set. Then there exists an compact neighborhood U_A of A such that $A = \text{l.a.s.}(U_A)$. Clearly $(U_A, 1 - v(x; U_A))$ is a stabilizing pair. \square

We have provided conditions for an absorbing set A to be stable, and hence stochastically attractive, under the conditions of Theorem 4.32. With regards to its domain of attraction, the set of points for which it holds with probability 1 is clearly given by $\{c(\cdot; A) = 1\}$. Since c is a harmonic function, such set is itself absorbing and hence may coincide either with A or with X . So, the claim that convergence must hold with probability 1 may be too conservative and instead one may consider ε -domains of attraction given by $\{c(\cdot; A) \geq 1 - \varepsilon\}$. To characterize such domains the procedure of computing the function c with explicit bounds on the error is needed. We show that the knowledge of a Lyapunov-like function as in Lemma 4.35 is not only useful to establish the stability of A , but also for such a computational procedure.

Lemma 4.36 *For any $A \in \mathcal{B}(X)$ the following trichotomy holds: either $h(\cdot; A) = 0$ or $h(\cdot; A) = 1$, or $\inf_{x \in X} h(x; A) = 0$ and $\sup_{x \in X} h(x; A) = 1$.*

Proof: Suppose that $\inf_{x \in X} h(x; A) > 0$. Clearly, $P_x \left(\lim_{n \rightarrow \infty} h(\mathbf{x}_n; A) = 0 \right) = 0$, hence by Lemma 4.28 we obtain that $h(\cdot; A) = 1$. Applying the same argument to the case $\sup_{x \in X} h(x; A) < 0$ we obtain that $h(\cdot; A) = 0$. \square

Theorem 4.37 *Assume that \mathcal{S} is continuous, A is a stable absorbing set which admits a stabilizing pair (U_A, f) with $f \in \text{b}\mathcal{C}^\nabla(X)$, and $r > 0$ is as in the statement of Lemma 4.35. Assume also that there exists an open set E such that $\inf_{x \in E} c(x; A) = 0$, $U_A \cap E = \emptyset$, $A = \text{l.a.s.}(E^c)$ and put $D_{r'} = (f_{\leq r'} \cup E)^c$ for all $r' \in \mathbb{R}$. Then*

$$|c(x; A) - w(x; D_{\varepsilon r}, \{f \leq r\})| \leq \max \left(\varepsilon, \sup_{x' \in E} c(x'; A) \right), \quad (4.33)$$

for any $\varepsilon \in (0, 1)$.

Proof: From the proof of Theorem 4.32 we obtain that $c(x; A) = h(x; U_A)$. Moreover, since $\lim_{x \rightarrow A} f(x) = 0$, we obtain that there exists $m \in \mathbb{N}$ such that $A^{1/m} \subseteq \{f \leq r\}$ and since

$$\overline{A^{1/(m+1)}} \subseteq \{f \leq r\} \subseteq U_A,$$

we obtain that $c(x; A) = h(x; \{f \leq r\})$.

Let us denote by $\tau^\varepsilon = \inf \{n \in \mathbb{N} : \mathbf{x}_n \in \{f \leq \varepsilon_r\} \cup E\}$. Then:

$$h(x) = P_x(\mathbf{x}_n \in \{f \leq r\}, \tau^\varepsilon = \infty) + P_x(\mathbf{x}_n \in \{f \leq r\}, \tau^\varepsilon < \infty). \quad (4.34)$$

For the first term of (4.34) we have:

$$P_x(\mathbf{x}_n \in \{f \leq r\}, \tau^\varepsilon = \infty) = P_x(\mathbf{x}_n \in \{f \leq r\} \setminus \{f \leq \varepsilon_r\}, \tau^\varepsilon = \infty) = 0.$$

To prove it we show that $\{f \leq r\} \setminus \{f \leq \varepsilon_r\}$ has function v zero everywhere. Since $\lim_{x \rightarrow A} f(x) = 0$ there exists $m \in \mathbb{N}$ such that $A^{1/m} \subseteq \{f \leq \varepsilon_r\}$, so $U_A \setminus A^{1/m}$ is compact and simple, hence has a zero v function. Hence $\{f \leq r\} \setminus \{f \leq \varepsilon_r\} \subseteq U_A \setminus A^{1/m}$ its v function is zero as well.

For the second term of (4.34) we have

$$P_x\{\mathbf{x}_n \in \{f \leq r\}, \tau^\varepsilon < \infty\} = \int_E h(y) \kappa(dx'|x) + \int_{\{f \leq \varepsilon_r\}} h(x') \kappa(dx'|x)$$

where $\kappa(B|x) = P_x\{\mathbf{x}_{\tau^\varepsilon} \in B, \tau^\varepsilon < \infty\}$. We obtain:

$$\inf_{x' \in \{f \leq \varepsilon_r\}} h(x') w(x; D_{\varepsilon_r}, \{f \leq \varepsilon_r\}) \leq h(x) \leq \sup_{x' \in E} h(x') + w(x; D_{\varepsilon_r}, \{f \leq \varepsilon_r\})$$

since $\inf_{x' \in E} h(x') = 0$. Now, $h(x) \geq v(x; \{f \leq r\}) \geq 1 - \frac{1}{r} f(x)$ and hence for all $x \in \{f \leq \varepsilon_r\}$ we have $h(x) \geq 1 - \varepsilon$, which completes the proof. \square

Let us make some remarks on Theorem 4.37. First, the reach-avoid value function w in (4.33) can be computed with explicit bounds on the error [129]. Combining such bounds with the right-hand side in (4.33), we obtain an approximate value of c with a known precision. Second, weak continuity of X and stability of A ensure that a necessary stabilizing pair exists by Lemma 4.35. Moreover, since $c(x; A) = h(x; U_A)$, it is either a constant function equal to 1, or the set E exists by Lemma 4.36.

Applying the result of Lemma C.10 one could extend Theorem 4.37 to the non-autonomous SSs.

Theorem 4.38 *Let Assumption 4.11 hold true and further assume stationary policies are sufficient, that is for all $x \in X$ assume that*

$$\mathcal{S}(x; \Diamond \Box A) = \sup_{\sigma \in \Sigma_S^\Gamma} Q_x^\sigma(\Diamond \Box A).$$

Suppose that g is a locally σ'_0 -excessive function on A for some stationary policy $\sigma' \in \Sigma_S^\Gamma$. Let $E \in \mathcal{B}(X)$ be any open set such that $\inf_{x \in E} \mathcal{S}_x^\Delta(\Diamond \Box A) = 0$, $E \cap \{g \leq 1\} = \emptyset$, E^c is

a compact set, and $(E^c)_\infty = A_\infty$. Then for all $x \in X$ and $\varepsilon \in (0, 1]$ it holds that

$$|\mathcal{S}(x; \Diamond \Box A) - \mathcal{S}(x; F_\varepsilon \cup B_\varepsilon)| \leq \max \left(\varepsilon, \sup_{x \in E} \mathcal{S}_x^\Delta(\Diamond \Box A) \right), \quad (4.35)$$

where $B_\varepsilon := \{g \leq \varepsilon\}$ and $F_\varepsilon = (G_\varepsilon \cup E)^c$.

4.3.4 Comments on the repeated reachability problem

We have mentioned that the characterization in Theorem 4.17 is taken from the literature on gambling: indeed we have not been able to find similar results obtained for the SS framework. It is interesting to see that the function $\mathcal{S}(x; \Diamond \Box S)$ satisfies a fixpoint equation, similarly to the uncontrolled case [131]. The connection between the solution of this problem and the value of the maximal safety probability $\mathcal{S}(x; \Box S)$ appears to be useful in characterizing simple instances, as we have encountered in Proposition 4.18.

There is range of literature in gambling on utilities with the form $J := \limsup_{n \rightarrow \infty} c(\mathbf{x}_n)$ and $J := \liminf_{n \rightarrow \infty} c(\mathbf{x}_n)$, which turn out to be repeated reachability specifications in the case the cost is an indicator function, namely $c(x) = 1_S(x)$. For the \limsup criterion, conditions on sufficiency of stationary policies have been obtained in [122] and [72], while for the \liminf case in [123]. A number of results valid for these criteria are summarized in [95, Section 4], in particular [95, Theorem 9.1, Chapter 4] provides a procedure to find $\mathcal{S}(x; \Box S)$ using the transfinite induction algorithm over all countable ordinals, rather than a simple recursion like in (4.21). Although this book only focuses on the case when the state space is countable, some of those results seem to allow for extensions to general Borel state spaces – more research is needed towards this goal. Unfortunately however, they do not seem to lead to practical computational procedures. To the best of our knowledge the result of Theorem 4.38 is novel, and is an extension of a version for uncontrolled processes in [131], where the focus was on studying the stability properties of the absorbing sets. Alternatively, it may be worth invoking some results obtained for recurrence [100]: however, such results are only strong when obtained under assumption of ψ -irreducibility of the transition kernel Γ [56, Chapter 10], which are often restrictive and lead to results that are rarely computational. The AC criterion also seems to be related to the \limsup and \liminf criteria in general, and to the repeated reachability property in particular, however much more research is needed to formally clarify the precise relationship. To summarize, on the one hand there are many results in gambling related to the repeated reachability problem, however they do not seem to lead to practically useful computational methods. On the other hand, in the SS setting such criteria have not received much attention, and although some related methods for other criteria [56] may be useful, such relationship is by no means direct or clear. The current contribution only makes an initial step towards numerical procedures for repeated reachability properties over SS, and much more research on the topic is needed.

Conclusions and future research

This chapter summarizes the thesis and shortly discusses its main contributions. We also provide some possible directions of future research.

5.1 Conclusions

In this thesis we have explored approximate solutions of diverse stochastic optimization problems over general Markov Decision Processes (MDP) in the framework of stochastic systems (SS), and equally expressible model, for which it is a bit easier and less cumbersome to derive the desired results. We have studied both optimization criteria popular in the control community, e.g. the discounted additive cost criterion, and even more focus was given to less explored group of specifications expressed as formulae of linear temporal logic (LTL), more common in the computer science community.

In order to make things somewhat more natural and intuitive, we have departed from the problems at hand, motivating each of our concept through the fact it is useful for the solution of the problems we think of and discarding alternatives which would not bring us closer to that solution, while keeping a direction in some sense parallel to a well-developed theory for non-probabilistic transition systems (TS). For example, as a first step we have defined behaviors for SSs in Section 2.4, showed that they do not exactly coincide with those of TSs, but yet are similar to the latter in the sense that concepts of behaviors for both models are helping in solving the linear temporal optimization problems for each of them. Due to slight difference in the problems one should not expect the precise match of behaviors in absolute sense, but only in the relative: how comprehensive is the concept of behavior for this particular model. We believe that contribution of this thesis covers both concrete problems and conceptual advancements that came up while trying to solve the former. Hence, in the list below we go over results of both categories, increasing the degree of abstractness.

- Bisimulations theory for MDPs. In Chapter 3 we presented a novel theory of approximate bisimulation of MDPs strong enough to cover error in optimization of any bounded-horizon linear temporal specification, or any discounted cost criterion. Relations we have introduced exhibit many desired properties such as transitivity, which makes it possible to combine them, and achieve stronger results step by step. We have shown that the bounds provided on approximation techniques are tight in a sense that they cannot be improved in general, without introducing additional assumptions on the closeness of the models. Finally, we have provided concrete algorithmic procedures of constructing finite abstractions for infinite systems: sound abstractions in general case, and approximate abstractions with any precision for MDPs whose dynamics is continuous.
- Analysis of infinite-horizon LTL optimization for MDPs. As bounds provided by approximate bisimulation methods do not always extend well beyond the finite time horizon, for infinite-horizon problems we gave a more specific treatment in Chapter 4, using automata theory to reduce those problems to fundamental ones: reachability and repeated reachability. The former problem we have solved by restating it as an additive cost criterion, which gave us an access to a rich toolbox of optimization results developed for that criterion in classical stochastic control literature. We have also studied more thoroughly cases which went beyond that literature, and discovered strong relations between stability of MDPs through absorbing sets and solutions of reachability problem, which led us to a strong equivalence result: if there are no absorbing sets outside the goal set the latter will be reached with probability 1, and otherwise one needs to “cut-off” those sets with some Lyapunov-like function, and solve the problem on the rest of the space. The results for repeated reachability in the most general setting did not yield so strong results, so we further focused on the case when the control strategy is fixed or the MDP is autonomous, i.e. a Markov Chain (MC). For that we have connected repeated reachability with harmonic functions theory and stochastic stability, which allowed us to achieve an equivalence result between existing of stable absorbing sets and Lyapunov functions for them. This result is practically important, as Lyapunov functions are needed in solutions of both reachability and repeated reachability problems when non-trivial absorbing sets are present.
- More conceptually, we have formalized notions of approximate bisimulations in Appendix A.3, providing general theory for approximate relations, defining their properties that are of use for us, and relating them to a more familiar notion of pseudometrics. This theory is agnostic of the model definition, and applies to TSs, SSs, MDPs and likely many other models.
- In Appendix A.4 we have given a try to formalize the notion of model syntax and semantics the way we use it in the thesis, and perhaps also useful for other works. In particular, we have argued what are the relevant versions of behaviors one should focus on, and that they should be motivated by the semantics of the model, rather than by its syntax. We have provided several example from this thesis and from the literature.

- In Appendix A.6 we have stressed how strongly similar definitions of behavioral inclusions for TSs and SSs in fact are, and that they are like to be particular cases of more general framework, which perhaps can cover more models in one go. We have done this with the help of theory of weak (convex) inclusions developed in A.5
- Finally, in Appendix B.2 we have in detail discussed a known lifting procedures of operations from states to measure, and to the best of our knowledge we are the first to study comprehensively properties of this lifting applied to lifting analytic relations between state to those between measures. In particular, we have shown that it can be considered as an endofunctor in the category where objects are Borel spaces and morphisms are analytic relations between them. This can be considered as an extension of the work by Lawvere and Giri on probabilistic functor, which is a special case of our functor, when the latter is restricted to graphs of Borel maps. Those results were particularly useful in showing transitivity-like properties of precise and approximate relations between SSs.

5.2 Future research

We envision the following three directions of how to extend the results of this thesis.

- i. Finite-horizon theory of approximate bisimulation is rather complete and comprehensive in here, which is supported by the tightness of the bounds. For the infinite horizon case, though, we had focus on fundamental problems separately. It is likely that one still can come up with a notion of approximate bisimulation of SSs that is more helpful for the latter type of tasks. One approach would be to try out bisimulation functions: those are very likely to provide strong bounds on dynamics between different models, but they often exist only in a small subclass of cases when the models exhibit certain global stability properties. Conservativeness of this approach is even stronger in SSs compared to TSs, as in the former case one has to deal with the issue of choosing the best coupling to bound the dynamics – see e.g. [135]. As an alternative, likely replacing the total variation distance with the Kantorovich metric in comparing behaviors of the SSs and hence in the notions of approximate bisimulations would yield results that are less conservative both w.r.t. the current approach, and the one based on approximate bisimulation functions.
- ii. Even though we think that finite horizon theory for bounding optimization results in total variation is now in a good shape, yet from authors personal experience those bounds are still happen to be rather conservative in some cases, especially when it comes to high-dimensional models. Even though precise estimates are hard to improve in such cases, one can allow for less certain errors by using e.g. randomized methods, for example an interesting

procedure was discussed in [79]. Even though the bounds there are less formal than the one we focused on, from a practical point of view they may be sufficient in many tasks, while being easier to achieve even for very complex systems.

- iii. Finally, if we shift the focus away from MDPs and look at things from more general perspective, we believe that results on abstract models and their behaviors in this system can be further extended to a more general framework, which perhaps will make it easier to extrapolate results from one model type to another.

General concepts

A.1 Notation and conventions

This work is motivated by the problems arising in stochastic optimal control. Even over finite models one cannot avoid dealing with uncountable sets, for example a set of all trajectories of the system – see e.g. [95]. For this reason, one needs measure-theoretical probability theory even in such case, let alone when solving control problems over stochastic models with uncountable state or actions sets. Therefore, the structure of such sets is of little importance, and there is no difference whether the states of a system belong to the interval $[0, 1]$ or to some infinite-dimensional manifold. Due to this reason, instead of restricting ourselves to e.g. hybrid spaces [28], we present most of the results in a rich and neat framework of Borel spaces which combines topology and measure theory. We assume the reader to be familiar with basic notions in topology, measure theory and probability; they can also be consulted e.g. in [60, Chapters 0-4 & 10]. Whenever we use more advanced results from these areas, we try our best to put precise references. Usually such results concern Borel spaces and their analytic subsets: most of the relevant material can be found in [20, Chapter 7], with a specific focus on stochastic optimal control, or in dedicated books [108], [82] and [120]¹.

A Borel space is a topological space homeomorphic to a Borel subset of a complete separable metric space. Examples of Borel spaces are the sets of reals \mathbb{R} endowed with the Euclidean topology, its subset of non-negative reals \mathbb{R}_+ with a subspace topology, the set of non-negative integers \mathbb{N} endowed with the discrete topology, or its finite subsets $[0; n] := \{0, 1, \dots, n\}$ for each $n \in \mathbb{N}$ with the discrete topology. A countable product of Borel spaces is a Borel space, and its Borel σ -algebra equals its product σ -algebra [20, Proposition 7.13]

The Borel σ -algebra of a Borel space Ω is denoted by $\mathcal{B}(\Omega)$, and the space of all Borel probability measures on Ω by $\mathcal{P}(\Omega)$. In this thesis we use terms “probability

¹ There was also developed a framework for stochastic optimal control in discrete time over general sets which avoids measurability issues, by dealing with finitely-additive measures [36], as opposed to countably-additive ones usually met in probability theory. The details of this framework can be found in [48]. Although an interesting approach, it comes with its own trade-offs, so we do not follow it here.

measure” and “distribution” interchangeably. We always assume that $\mathcal{P}(\Omega)$ is endowed with the topology of weak convergence, so that it is a Borel space itself [20, Corollary 7.25.1]. It follows that $\mathcal{B}(\mathcal{P}(\Omega))$ coincides with a σ -algebra generated by the evaluation maps $(\epsilon_A)_{A \in \mathcal{B}(\Omega)}$, where $\epsilon_A : \mu \mapsto \mu(A)$.

Given another Borel space Ξ , the set of all Borel measurable maps from Ω to Ξ is denoted by $\mathcal{B}(\Omega, \Xi)$. We use a shorthand $\mathcal{B}(\Omega)$ for $\mathcal{B}(\Omega, \mathbb{R})$. A bijection $\varphi \in \mathcal{B}(\Omega, \Xi)$ is called an isomorphism if $\varphi^{-1} \in \mathcal{B}(\Xi, \Omega)$ ². We say that Ω and Ξ are isomorphic if there exists an isomorphism between them. If $\varphi \in \mathcal{B}(\Omega, \Xi)$ is an injection, then $\varphi(A) \in \mathcal{B}(\Xi)$ for each $A \in \mathcal{B}(\Omega)$, and $\varphi^{-1} \in \mathcal{B}(\varphi(\Omega), \Omega)$ [20, Proposition 15]. Furthermore, a powerful result called the Borel isomorphism theorem claims that any two uncountable Borel spaces are isomorphic, and any countable Borel space is isomorphic to a subset of \mathbb{N} [20, Proposition 16]³. Often to prove a fact concerning a general Borel space, one proves it first for a “convenient” Borel space (e.g. $\mathbb{N}^{\mathbb{N}}$) and then applies the Borel isomorphism theorem. Similarly, if a counterexample is found in for some particular Borel space, it can be translated to a general setting.

In case we relax the assumption that $\varphi \in \mathcal{B}(\Omega, \Xi)$ is injective, then it may happen that $\varphi(A) \notin \mathcal{B}(\Xi)$ for some $A \in \mathcal{B}(\Omega)$ even if φ is continuous [125]. Such maps are called analytic; applying the Borel isomorphism theorem, we can define analytic subsets of Ω as images of $[0, 1]$ under all possible maps $\varphi \in \mathcal{B}([0, 1], \Omega)$, so in particular any Borel set is analytic. We denote the collection of analytic subsets of Ω by $\mathcal{A}(\Omega)$; it admits several other equivalent characterizations [82]. The collection $\mathcal{A}(\Omega)$ is closed under countable unions and intersections, and under Borel images: $\varphi(A) \in \mathcal{A}(\Xi)$ for each $A \in \mathcal{A}(\Omega)$ and $\varphi \in \mathcal{B}(\Omega, \Xi)$. Yet $\mathcal{A}(\Omega)$ is not a σ -algebra as the complement of an analytic set is not necessarily analytic [82]. Sets whose complements are analytic are often referred to as co-analytic sets, and a subset of a Borel space is both analytic and co-analytic iff it is Borel [125]. A countable product of analytic sets is an analytic subset of a product space [82]. The connection between $\mathcal{A}(\Omega)$ and $\mathcal{A}(\mathcal{P}(\Omega))$ is as follows: for each $A \in \mathcal{A}(\Omega)$ it holds that $\{\mu : \mu(A) \geq c\} \in \mathcal{A}(\mathcal{P}(\Omega))$ for every $c \in \mathbb{R}$.

Given a measure $\mu \in \mathcal{P}(\Omega)$, we denote by $\mathcal{B}^\mu(\Omega)$ the p -completion of $\mathcal{B}(\Omega)$. The universal σ -algebra of Ω is given by $\mathcal{U}(\Omega) := \bigcap_{\mu \in \mathcal{P}(\Omega)} \mathcal{B}^\mu(\Omega)$; its elements are called universally measurable sets. The universal σ -algebra of a countable product of Borel spaces contains the product of their universal σ -algebras, but does not equal it in general [20]. Furthermore, every analytic subset of a Borel space is universally measurable. A map $\varphi : \Omega \rightarrow \Xi$ is called universally measurable if $\varphi^{-1}(B) \in \mathcal{U}(\Omega)$ for each $B \in \mathcal{B}(\Xi)$; in such case we write $\varphi \in \mathcal{U}(\Omega, \Xi)$. Despite the asymmetry in their definition, for Borel spaces a composition of universally measurable maps is universally measurable [20], so in particular $\varphi^{-1}(B) \in \mathcal{U}(\Omega)$ even if $B \in \mathcal{U}(\Xi)$. The following characterization is often useful: universally measurable maps are essentially Borel measurable, that is $\varphi \in \mathcal{U}(\Omega, \Xi)$ iff for any $\mu \in \mathcal{P}(\Omega)$ there exists $\varphi' \in \mathcal{B}(\Omega, \Xi)$ such that $\varphi = \varphi'$ (μ -a.s.) [20]. In particular, $A \in \mathcal{U}(\Omega)$ iff for each $\mu \in \mathcal{P}(\Omega)$ there exists $A' \in \mathcal{B}(\Omega)$ such that $\mu(A \Delta A') = 0$. As

² Isomorphisms for Borel spaces play a role of homeomorphisms for topological spaces.

³ This is a formalization of the idea we have mentioned above that from the point of view of Borel spaces, there is no difference between $[0, 1]$ or infinite-dimensional spaces.

a result, any $\mu \in \mathcal{B}(\Omega)$ extends uniquely to $\mathcal{U}(\Omega)$.

For a product of Borel spaces $\prod_{n \in \mathbb{N}} \Omega_n$ and an arbitrary index set $I \subseteq \mathbb{N}$ we denote by $\text{proj}_I : \prod_{n \in \mathbb{N}} \Omega_n \rightarrow \prod_{n \in I} \Omega_n$ the corresponding projection map, i.e. $\text{proj}_I : ((\omega_n)_{n \in \mathbb{N}}) \mapsto ((\omega_n)_{n \in I})$. If $A_n \in \mathcal{B}(\Omega_n)$ for each $n \in I$, then the set $(\text{proj}_I)^{-1}(\prod_{n \in I} A_n)$ is called a measurable rectangle. Given a sequence of maps $(\varphi_n)_{n \in I}$, where $\varphi_n : \Xi \rightarrow \Omega_n$ for each $n \in I$, we define

$$\left(\prod_{n \in I} \varphi_n \right) : \xi \mapsto (\varphi_n(\xi))_{n \in I}$$

to be the product map. We further denote their parallelization⁴ by

$$\left(\prod_{n \in I} \varphi_n \right) : (\xi_n)_{n \in I} \mapsto (\varphi_n(\xi_n))_{n \in I}.$$

In case $I = \{i, j\}$ we simply write $\varphi_i \times \varphi_j$ and $\varphi_i \sqcup \varphi_j$ instead of more cumbersome $\prod_{n \in \{i, j\}} \varphi_n$ and $\coprod_{n \in \{i, j\}} \varphi_n$. Moreover, we slightly abuse notation by using φ for $\prod_{n \in I} \varphi_n$ if $\varphi_n = \varphi$ for all $n \in I$, namely we extend φ to its element-wise application on vectors of values. If $\varphi_n \in \mathcal{B}(\Xi, \Omega_n)$ or $\varphi_n \in \mathcal{U}(\Xi, \Omega_n)$ for each $n \in I$, then $\prod_{n \in I} \varphi_n, \coprod_{n \in I} \varphi_n$ belong to $\mathcal{B}(\Xi, \prod_{n \in I} \Omega_n)$ or $\mathcal{U}(\Xi, \prod_{n \in I} \Omega_n)$ respectively.

Given a measure $\mu \in \mathcal{P}(\Omega)$ and a map $f \in \mathcal{U}(\Omega, \Xi)$, the pushforward of μ along f is a measure $f_*\mu \in \mathcal{P}(\Xi)$ satisfying $f_*\mu(B) = \mu(f^{-1}(B))$ for each $B \in \mathcal{B}(\Xi)$. If $M \in \mathcal{P}(\Omega \times \Xi)$, then $(\text{proj}_0)_*M$ and $(\text{proj}_1)_*M$ are called left and right marginals of M respectively. Given $\mu \in \mathcal{P}(\Omega)$ and $\mu' \in \mathcal{P}(\Xi)$, their coupling is any measure $M \in \mathcal{P}(\Omega \times \Xi)$ whose left and right marginals are μ and μ' respectively. The set of all couplings of μ and μ' is denoted by $\mathfrak{C}(\mu, \mu')$. If $M \in \mathcal{P}(\prod_{n \in \mathbb{N}} \Omega_n)$, then $(\text{proj}_I)_*M$ is called a finite-dimensional marginal of M . Whenever $I = [0; n]$, we use a shorthand $M|_n$ instead of $(\text{proj}_I)_*M$. Recall that if \mathcal{F} generates $\mathcal{B}(\Omega)$ and is closed under finite intersections, then measures $\mu, \mu' \in \mathcal{P}(\Omega)$ are equal iff $\mu(A) = \mu'(A)$ for each $A \in \mathcal{F}$. In particular, the measures $M, M' \in \mathcal{P}(\prod_{n \in \mathbb{N}} \Omega_n)$ are equal iff they agree on measurable rectangles, or equivalently if $M|_n = M'|_n$ for each $n \in \mathbb{N}$.

A map $\kappa : \Omega \rightarrow \mathcal{P}(\Xi)$ is alternatively called a stochastic kernel on Ξ given Ω . Given $\omega \in \Omega$ and $B \in \mathcal{B}(\Xi)$, instead of $\kappa(\omega)(B)$ we use a less cumbersome notation $\kappa(B|\omega)$. A shorthand $\kappa \in \mathcal{B}(\Xi|\Omega)$ is used for $\kappa \in \mathcal{B}(\Omega, \mathcal{P}(\Xi))$, and similarly $\kappa \in \mathcal{U}(\Xi|\Omega)$ means that $\kappa \in \mathcal{U}(\Omega, \mathcal{P}(\Xi))$. Given $\mu \in \mathcal{P}(\Omega)$ and $\kappa \in \mathcal{U}(\Xi|\Omega)$, the corresponding product measure in $\mathcal{P}(\Omega \times \Xi)$ is denoted by $\mu \otimes \kappa$, and is uniquely defined by the following formula:

$$(\mu \otimes \kappa)(A \times B) := \int_A \kappa'(B|\omega) \mu(d\omega), \quad A \in \mathcal{B}(\Omega), B \in \mathcal{B}(\Xi),$$

where $\kappa' \in \mathcal{B}(\Xi|\Omega)$ is any kernel satisfying $\kappa = \kappa'$ (μ -a.s.). Conversely, for $M \in$

⁴ The \prod and \sqcup symbols usually denote co-products in respective categories, so here we may use them in a non-canonical way. At the same time, we do not use co-products of maps in this thesis, so there should no be any confusion.

$\mathcal{P}(\Omega \times \Xi)$ and $\mu \in \mathcal{P}(\Omega)$ the regular conditional probability kernel $\frac{dM}{d\mu}$ has to satisfy $\mu \otimes \frac{dM}{d\mu} = M$. Since Ξ is a Borel space, such kernel always exists [52] and is defined uniquely only (μ -a.s.) since $\mu \otimes \kappa = \mu \otimes \kappa'$ iff $\kappa = \kappa'$ (μ -a.s.) for each $\kappa, \kappa' \in \mathcal{U}(\Xi|\Omega)$. A version of the Ionescu Tulcea theorem in [20, Proposition 7.45] implies that for any sequence of Borel spaces $(\Omega_n)_{n \in \mathbb{N}}$, initial distribution $\mu \in \mathcal{P}(\Omega_0)$ and a sequence of transition kernels $(\kappa_n)_{n \in \mathbb{N}}$, where $\kappa_n \in \mathcal{U}(\Omega_{n+1} | \prod_{k=0}^n \Omega_k)$ for each $n \in \mathbb{N}$, there exists a unique product measure $M \in \mathcal{P}(\prod_{n \in \mathbb{N}} \Omega_n)$ satisfying

$$M|_0 = \mu, \quad \frac{d(M|_{n+1})}{d(M|_n)} = \kappa_n \text{ (M|}_n \text{-a.s.)} \quad \forall n \in \mathbb{N}.$$

One of the simplest examples of kernels is given by a Dirac kernel $\delta \in \mathcal{B}(\Omega|\Omega)$ satisfying

$$\delta(A|\omega) = 1_A(\omega) = \begin{cases} 1, & \text{if } \omega \in A, \\ 0, & \text{if } \omega \notin A, \end{cases} \quad \omega \in \Omega, A \in \mathcal{B}(\Omega).$$

Here 1_A is an indicator function of the set A . Furthermore, for any $\varphi \in \mathcal{U}(\Omega, \Xi)$ we denote $\delta_\varphi := \delta \circ \varphi \in \mathcal{U}(\Xi|\Omega)$.

The graph of a map $\varphi : \Omega \rightarrow \Xi$ is denoted by

$$\text{Gr}(\varphi) := \{(\omega, \varphi(\omega)) : \omega \in \Omega\} \subseteq \Omega \times \Xi.$$

It holds that $\varphi \in \mathcal{B}(\Omega, \Xi)$ iff $\text{Gr}(f) \in \mathcal{B}(\Omega \times \Xi)$ [27]. For any $A \subseteq \Omega \times \Xi$ and any $\omega \in \Omega$, ω - and ξ -sections of A is given by

$$A|_\omega := \{\xi : (\omega, \xi) \in A\}, \quad A|^\xi := \{\omega : (\omega, \xi) \in A\}.$$

For $B \subseteq \Omega$ and $C \subseteq \Xi$ we also write $A|_B := \bigcup_{\omega \in B} A|_\omega$ and $A|^\mathcal{C} := \bigcup_{\xi \in C} A|^\xi$. We say that a map φ is a selector for a set A if $\varphi(\omega) \in A|_\omega$ for each $\omega \in \text{proj}_\Omega(A)$. Similarly, a kernel $\kappa : \Omega \rightarrow \mathcal{P}(\Xi)$ is a randomized selector for A if $\kappa^*(A|_\omega|\omega) = 1$ for all $\omega \in \text{proj}_\Omega(A)$. Note that for every selector φ one can put in a correspondence a randomized selector $\delta(\varphi)$. Obviously, every set admits a selector (let along a randomized one), so the question usually is whether such a selector can be chosen to be measurable. Since δ is a Borel-measurable kernel, $\delta(\varphi)$ preserves measurability of φ , so existence of a measurable selector implies existence of a randomized selector with the same measurability properties. The converse is not necessarily true [82]. Moreover, even if $A \in \mathcal{U}(\Omega \times \Xi)$, there may not exist a universally measurable randomized selector for A [82]. There is a number of results on the existence of measurable selectors, but here we mostly use the fact that any analytic set admits a universally measurable selector [20].

We say that a measure $\mu \in \mathcal{P}(X)$ is discrete if it is a linear combination of Dirac measures, i.e. there exists $(\omega_n)_{n \in \mathbb{N}} \subseteq \Omega$ and $\nu \in \mathcal{P}(\mathbb{N})$ such that $\mu = \sum_{n \in \mathbb{N}} \delta(\omega_n) \nu(\{n\})$. More generally, given a family of measures $P \subset \mathcal{P}(X)$, we say that μ is a barycenter of P if there exists $\nu \in \mathcal{P}(\mathcal{P}(X))$ and $P' \in \mathcal{U}(\mathcal{P}(\Omega))$ such that $P' \subseteq P$ and $\mu = \int_{P'} \mu' \nu(d\mu')$. Note that if $P \in \mathcal{U}(\mathcal{P}(\Omega))$, one can always take $P' = P$. We de-

note the set of all barycentres of P by $\text{sco } P$ and refer to it as the strong convex hull of P ⁵; for $P \in \mathcal{A}(\mathcal{P}(\Omega))$ it holds that $\text{sco } P \in \mathcal{A}(\mathcal{P}(\Omega))$ and that $\text{sco}(\text{sco } P) = \text{sco } P$ [39, 33], so that sco is an idempotent operation on analytic subsets of $\mathcal{P}(\Omega)$. In particular, $P \subseteq \text{sco } P'$ iff $\text{sco } P \subseteq \text{sco } P'$ for each $P, P' \in \mathcal{A}(\mathcal{P}(\Omega))$.

For any $\varphi \in \mathcal{U}(\Omega)$ we define its norm as $\|\varphi\| := \sup_{\omega \in \Omega} |f(\omega)|$, and denote

$$\text{b}\mathcal{U}_c(\Omega) := \{\varphi \in \text{b}\mathcal{U}(\Omega) : \|\varphi\| \leq c\}, \quad c \in \mathbb{R}_+ \cup \{\infty\}.$$

Given $\mu \in \mathcal{P}(\Omega)$ we often write $\mu\varphi$ instead of $\int_{\Omega} \varphi d\mu$. We write that $\varphi \in \mathcal{H}_c^n(\Omega)$ if $\varphi \in \text{b}\mathcal{U}_c(\Omega^{\mathbb{N}})$ and $\mu\varphi = \mu'\varphi$ whenever $\mu|_n = \mu'|_n$ for any $\mu, \mu' \in \mathcal{P}(\Omega^{\mathbb{N}})$, that is φ depends only on first $n + 1$ coordinates. Although the topological space $\mathcal{P}(\Omega)$ is metrizable with the Prokhorov metric [108], we are interested in considering a different metric on $\mathcal{P}(\Omega)$ called the total variation metric

$$d_{\text{TV}}(\mu, \mu') := \sup_{\varphi \in \text{b}\mathcal{U}_1(\Omega)} \left| \int_{\Omega} \varphi d\mu - \int_{\Omega} \varphi d\mu' \right|, \quad \mu, \mu' \in \mathcal{P}(\Omega).$$

This definition of d_{TV} for general measures admits a special formula in our case:

$$d_{\text{TV}}(\mu, \mu') = 2 \cdot \sup_{A \in \mathcal{U}(\Omega)} |\mu(A) - \mu'(A)|, \quad \forall \mu, \mu' \in \mathcal{P}(\Omega).$$

Note also that if $\varphi \in \text{b}\mathcal{U}_{\infty}(\Omega)$ is such that $\varphi(\omega) \in [a, b]$ for each $\omega \in \Omega$, then

$$\left| \int_{\Omega} \varphi d\mu - \int_{\Omega} \varphi d\mu' \right| \leq \frac{1}{2} \cdot (b - a) \cdot d_{\text{TV}}(\mu, \mu').$$

If Ω is a countable space, each probability measure $\mu \in \mathcal{P}(\Omega)$ can be decomposed as $\mu = \sum_{\omega \in \Omega} \mu(\{\omega\})\delta(\omega)$, and hence $d_{\text{TV}}(\mu, \mu') = \sum_{\omega \in \Omega} |\mu(\omega) - \mu'(\omega)|$. Furthermore,

$$\int_P \mu \nu(d\mu) = \int_P \sum_{\omega \in \Omega} \mu(\{\omega\})\delta(\omega)\nu(d\mu) = \sum_{\omega \in \Omega} \left(\int_P \mu(\{\omega\})\nu(d\mu) \right) \delta(\omega)$$

In particular, when Ω is a finite space of cardinality $n + 1$, say $\Omega = [0; n]$, $\mathcal{P}(\Omega)$ is an n -dimensional space which can be identified with a probability simplex

$$\theta_n := \left\{ \theta \in \mathbb{R}_+^n : \sum_{i=0}^n \theta_i = 1 \right\}, \quad (\text{A.1})$$

the homeomorphism obviously being $\theta_i = \mu(\{i\})$ for each $i \in [0; n]$. Note that d_{TV} in such case is just the 1-norm on θ_n , and the metric topology of d_{TV} coincides with the topology of weak convergence $\mathcal{P}(\Omega)$ is endowed with by default, however for general Borel spaces these topologies differ [62].

⁵ Here we use notation and terminology of [97, 98], although perhaps it would be more correct to say that the set of all barycentres of P is its barycentric closure. For more details on this topic, see [39].

For $\varepsilon, \varepsilon' \in \mathbb{R}$ we introduce the following arithmetic operation

$$\varepsilon \otimes \varepsilon' := 2 \left(1 - \left(1 - \frac{\varepsilon}{2} \right) \left(1 - \frac{\varepsilon'}{2} \right) \right).$$

Note that it is a symmetric and associative binary operation on reals. We use the following shorthand: $\varepsilon^{\otimes 1} := \varepsilon$ and $\varepsilon^{\otimes n} := \varepsilon^{\otimes(n-1)} \otimes \varepsilon$, so $\varepsilon^{\otimes n} = 2(1 - (1 - \varepsilon/2)^n)$. We further write $\varepsilon \vee \varepsilon'$ and $\varepsilon \wedge \varepsilon'$ for the maximum and minimum of $\varepsilon, \varepsilon' \in \mathbb{R}$ respectively.

A.2 Pseudometrics

A *pseudometric* on Ω is a function $d : \Omega^2 \rightarrow [0, \infty)$ which

- is symmetric: $d(\omega, \omega') = d(\omega', \omega)$ for all $\omega, \omega' \in \Omega$;
- satisfies the triangular inequality: $d(\omega, \omega'') \leq d(\omega, \omega') + d(\omega', \omega'')$ for all points $\omega, \omega', \omega'' \in \Omega$.

Note that the triangular inequality in particular implies that $d(\omega, \omega) = 0$ for all $\omega \in \Omega$, however it may happen that $d(\omega, \omega') = 0$ for $\omega \neq \omega'$. In case $d(\omega, \omega') = 0$ only for $\omega = \omega'$ we say that d is a *metric* on Ω .

By a *pseudometric space* we mean a pair (Ω, d) where Ω is an arbitrary set and d is some pseudometric on Ω . A pseudometric topology on (Ω, d) is the one generated by the family $(\{\omega' \in \Omega : d(\omega, \omega') < \varepsilon\})_{\omega \in \Omega, \varepsilon \in \mathbb{R}_+}$ of open d -balls. We always assume a pseudometric space to be endowed with its pseudometric topology. A pseudometric space is called bounded if $\|d\| < \infty$. If the original pseudometric space is not bounded, we can introduce a new pseudometric $d' := d/(1 + d)$: it is bounded by 1 and generates the same topology as the original pseudometric does.

If $(\bar{\Omega}, \bar{d})$ is another pseudometric space, for any map $f : \Omega \rightarrow \bar{\Omega}$ we let

$$\mathfrak{c}(f) := \sup_{d(\omega, \omega') > 0} \frac{\bar{d}(f(\omega), f(\omega'))}{d(\omega, \omega')}$$

denote the modulus of continuity of f . If $\beta = \mathfrak{c}(f) < \infty$, we say that f is a β -Lipschitz map. In particular, each such map is continuous. For each $\beta \in \mathbb{R}_+$ the space of all β -Lipschitz maps we denote by

$$\text{Lip}_\beta((\Omega, d), (\bar{\Omega}, \bar{d})) := \{f : \Omega \rightarrow \bar{\Omega} \text{ such that } \mathfrak{c}(f) \leq \beta\}.$$

When $\bar{\Omega} = \mathbb{R}$ and \bar{d} is the Euclidean metric we use a shorthand $\text{Lip}_\beta(\Omega, d)$.

An important example of 1-Lipschitz function is given by a pseudometric itself. For each $A \subseteq \Omega$ define the distance to A by $d(\omega, A) := \inf_{\omega' \in A} d(\omega, \omega')$. From the triangular inequality it follows then that $d(\cdot, A) \in \text{Lip}_1(\Omega, d)$. We also denote

$$A^\varepsilon := \{\omega \in \Omega : d(\omega, A) \leq \varepsilon\}.$$

The following proposition connects set inclusion with Lipschitz functions.

Proposition A.1 *Consider two arbitrary sets $A, B \subseteq \Omega$, then $A \subseteq B^\varepsilon$ iff*

$$\sup_{\omega \in A} f(\omega) \leq \sup_{\omega' \in B} f(\omega') + \varepsilon, \quad \forall f \in \text{Lip}_1(\Omega, d). \quad (\text{A.2})$$

Proof: Let us show first that for all $\omega \in B^\varepsilon$ and $f \in \text{Lip}_1(\Omega, d)$ it holds that

$$f(\omega) \leq \sup_{\omega' \in B} f(\omega') + \varepsilon. \quad (\text{A.3})$$

Consider a sequence $(\omega_n)_{n \in \mathbb{N}} \subseteq B$ such that $d(\omega, \omega_n) \leq \varepsilon + \frac{1}{n}$. Since f is 1-Lipschitz

$$|f(\omega) - f(\omega_n)| \leq \varepsilon + \frac{1}{n} \implies f(\omega_n) \geq f(\omega) - \varepsilon - \frac{1}{n}.$$

As a result, $\sup_{\omega' \in B} f(\omega') \geq \sup_{n \in \mathbb{N}} f(\omega_n) \geq f(\omega) - \varepsilon$ which is equivalent to (A.3). The latter inequality further implies (A.2) in case $A \subseteq B^\varepsilon$.

Conversely, if (A.2) holds, then letting $f = d(\cdot, B)$ yields $\sup_{\omega \in A} d(\omega, B) \leq \varepsilon$. \square

A.3 Approximate relations

Let Ω and $\bar{\Omega}$ be arbitrary sets. A relation on Ω and $\bar{\Omega}$ is just a subset of their product: $\Phi \subseteq \Omega \times \bar{\Omega}$. Sometimes we write $\omega \Phi \bar{\omega}$ to mean $(\omega, \bar{\omega}) \in \Phi$ in order to emphasize that we treat the set Φ as a relation.

One can consider a category Rel where objects are sets and morphisms are relations. The identity relation on any set Ω is given by its diagonal Δ_Ω , and the composition of relations $\Phi \subseteq \Omega \times \bar{\Omega}$ and $\bar{\Phi} \subseteq \bar{\Omega} \times \hat{\Omega}$ is given by

$$\bar{\Phi} \circ \Phi := \{(\omega, \hat{\omega}) : \exists \bar{\omega} \in \bar{\Omega} \text{ such that } (\omega, \bar{\omega}) \in \Phi \text{ and } (\bar{\omega}, \hat{\omega}) \in \bar{\Phi}\}. \quad (\text{A.4})$$

The inverse of a relation $\Phi \subseteq \Omega \times \bar{\Omega}$ is defined as

$$\Phi^{-1} := \{(\bar{\omega}, \omega) : (\omega, \bar{\omega}) \in \Phi\}.$$

It further holds that $(\Phi^{-1})^{-1} = \Phi$ and $(\bar{\Phi} \circ \Phi)^{-1} = \Phi^{-1} \circ \bar{\Phi}^{-1}$.

Whenever $\Phi \subseteq \Omega \times \bar{\Omega}$ and $\Psi \subseteq \Xi \times \bar{\Xi}$ are two relations, we can define their product $\Phi \times \Psi \subseteq \Omega \times \bar{\Omega} \times \Xi \times \bar{\Xi}$ as a usual Cartesian product of sets:

$$\Phi \times \Psi = \{(\omega, \bar{\omega}, \xi, \bar{\xi}) : (\omega, \bar{\omega}) \in \Phi \text{ and } (\xi, \bar{\xi}) \in \Psi\}.$$

We say that $\Phi \subseteq \Omega \times \bar{\Omega}$ is a *left-total relation* (l.t.r.) if $\text{proj}_\Omega(\Phi) = \Omega$. Any l.t.r. can be regarded as graph a multi-valued map $\phi : \Omega \rightarrow 2^{\bar{\Omega}} \setminus \{\emptyset\}$ given by $\phi(\omega) = \Phi|_\omega$. In particular, whenever $\varphi : \Omega \rightarrow \bar{\Omega}$ is a usual single-valued map, $\omega \text{ Gr}(\varphi) \bar{\omega}$ iff $\bar{\omega} = \varphi(\omega)$.

Whenever $\bar{\Omega} = \Omega$, the following classification applies. A relation $\Phi \subseteq \Omega^2$ is said to be

- *reflexive* if $\omega\Phi\omega$ for any $\omega \in \Omega$, or equivalently $\Delta_\Omega \in \Phi$;
- *symmetric* if $\omega\Phi\omega'$ implies $\omega'\Phi\omega$ for any $\omega, \omega' \in \Omega$, or equivalently $\Phi = \Phi^{-1}$;
- *transitive* if $\omega\Phi\omega'$ and $\omega'\Phi\omega''$ together imply $\omega\Phi\omega''$ for any $\omega, \omega', \omega'' \in \Phi$, or equivalently $\Phi \circ \Phi \subseteq \Phi$.

We say that the relation Φ is a *preorder* if it is both reflexive and transitive; a symmetric preorder is called *equivalence*. The symmetrization of a relation $\Phi \subseteq \Omega^2$ is another relation $\tilde{\Phi}$ given by $\tilde{\Phi} := \Phi \cap \Phi^{-1}$, so that $\omega\tilde{\Phi}\omega'$ iff $\omega\Phi\omega'$ and $\omega'\Phi\omega$. If Φ is symmetric that its symmetrization is clearly Φ itself. Furthermore, note that if Φ is reflexive or transitive, then so is its symmetrization. In particular, a symmetrization of any preorder defines an equivalence. An example of a preorder is the linear order on reals $\leq \subseteq \mathbb{R}^2$; its symmetrization is an equivalence relation given by the equality of reals $= \subseteq \mathbb{R}^2$.

In our work relations appears on different levels: those between states of different models are usually defined for Ω and $\bar{\Omega}$ being unequal sets, so that we cannot apply the classification above to them. At the same time, we lift relations from states to models themselves hence inducing relations on the set of models – these can be classified as reflexive, symmetric or transitive. Such “exact” relations between systems are known to be insufficient for the practical purposes [64]. Due to this reason we also study families of relations on Ω indexed by some “approximation” parameter $\varepsilon \in \mathbb{R}_+$. We refer to such objects as *approximate relations* or ε -relations. Note that an ε -relation is not a single relation but rather a family thereof.

Let $\Psi := (\Psi_\varepsilon)_{\varepsilon \in \mathbb{R}_+}$ be an ε -relation on Ω . We say that Ψ is reflexive (symmetric) whenever Ψ_ε is reflexive (symmetric) for all $\varepsilon \in \mathbb{R}_+$. Furthermore, for the ε -relation Ψ it is also important to study the interplay between its components Ψ_ε for different values of ε . We say that an ε -relation Ψ is

- *monotone* if $\Psi_\varepsilon \subseteq \Psi_{\varepsilon'}$ for all $\varepsilon' \geq \varepsilon \in \mathbb{R}_+$;
- *finite* if it is monotone and for each $\omega, \omega' \in \Omega$ there exists $\varepsilon \in \mathbb{R}_+$ such that $\omega\Psi_\varepsilon\omega'$, or equivalently $\bigcup_{\varepsilon \in \mathbb{R}_+} \Psi_\varepsilon = \Omega^2$;
- *triangular* if $\omega\Psi_\varepsilon\omega'$ and $\omega'\Psi_{\varepsilon'}\omega''$ together imply that $\omega\Psi_{\varepsilon+\varepsilon'}\omega''$, or equivalently $\Psi_{\varepsilon'} \circ \Psi_\varepsilon \subseteq \Psi_{\varepsilon+\varepsilon'}$, for all $\varepsilon, \varepsilon' \in \mathbb{R}_+$;
- *continuous at* $\rho \in \mathbb{R}_+$ if Ψ is monotone and $\Psi_\rho = \bigcap_{\varepsilon > \rho} \Psi_\varepsilon$;
- *continuous* if it is continuous at each $\rho \in \mathbb{R}_+$.

Note that $\Psi_\rho \subseteq \bigcap_{\varepsilon > \rho} \Psi_\varepsilon$ for any monotone Ψ and $\rho \in \mathbb{R}_+$. Hence, for the continuity it is sufficient to check only the converse inclusion. Notice also, that if Ψ is triangular then necessary Ψ_0 is transitive. Due to this reason, we say that Ψ is an ε -preorder if it is reflexive, monotone, finite and triangular. A symmetric ε -preorder

we call an ε -equivalence; notice that in such case Ψ_0 is an equivalence relation. For an ε -relation Ψ its symmetrization is defined as an ε -relation $\tilde{\Psi} = (\tilde{\Psi})_{\varepsilon \in \mathbb{R}_+}$ where for each $\varepsilon \in \mathbb{R}_+$ the relation $\tilde{\Psi}_\varepsilon$ is a symmetrization of Ψ_ε . If Ψ is monotone, finite, triangular or continuous at $\rho \in \mathbb{R}_+$ then its symmetrization is monotone, finite, triangular or continuous at ρ respectively. In particular, the symmetrization of any ε -preorder is an ε -equivalence.

Examples of ε -relations can be given using pseudometrics. For any pseudometric d on Ω we can define its ε -diagonals for $\varepsilon \in \mathbb{R}_+$ by

$$\Delta_{\varepsilon,d} := \{(\omega, \omega') : d(\omega, \omega') \leq \varepsilon\} \subseteq \Omega^2.$$

Their connection with the diagonal of Ω is obviously given by $\Delta_\Omega \subseteq \Delta_{0,d}$; also, these sets coincide iff d is a metric. It is easy to check that $\Delta_d := (\Delta_{\varepsilon,d})_{\varepsilon \in \mathbb{R}_+}$ is a continuous ε -equivalence (in particular, $\Delta_{0,d}$ is an equivalence relation), so it is a natural example of an ε -relation having all the properties of ε -relations introduced above.

It follows that a somewhat converse result holds true: any ε -equivalence Ψ induces a corresponding pseudometric d_Ψ such that $\Psi_\varepsilon \subseteq \Delta_{\varepsilon,d_\Psi}$ for all $\varepsilon \in \mathbb{R}_+$, given by

$$d_\Psi(\omega, \omega') := \inf\{\varepsilon \in \mathbb{R}_+ : \omega \Psi_\varepsilon \omega'\}.$$

The function d_Ψ is indeed a pseudometric as the following result shows.

Proposition A.2 *Let Ω be an any set, and $(\Psi_\varepsilon)_{\varepsilon \in \mathbb{R}_+}$ be an ε -equivalence on Ω . Then:*

- i. Ψ_0 is an equivalence relation on Ω ;
- ii. $\omega \Psi_{d_\Psi(\omega, \omega') + \delta} \omega'$ for all $\delta > 0$;
- iii. d_Ψ is a pseudometric on Ω ;
- iv. $\Delta_{\varepsilon,d_\Psi} = \bigcap_{\delta > \varepsilon} \Psi_\delta$, and in particular $\Psi_\varepsilon \subseteq \Delta_{\varepsilon,d_\Psi}$ for all $\varepsilon \in \mathbb{R}_+$;
- v. $\Psi_\rho = \Delta_{\rho,d_\Psi}$ iff Ψ is continuous at ρ ;
- vi. d_Ψ is a metric iff Ψ is continuous at 0 and $\Psi_0 = \Delta_\Omega$.

Proof: The proof is as follows:

- i. To show that Ψ_0 is an equivalence relation, we only need to notice that it is reflexive and symmetric by assumptions, and as we have mentioned above, triangularity of Ψ implies transitivity of Ψ_0 .
- ii. Suppose that $(\omega, \omega') \notin \Psi_{d_\Psi(\omega, \omega') + \delta}$ for some $\delta > 0$, then

$$d_\Psi(\omega, \omega') = \inf\{\varepsilon \in \mathbb{R}_+ : \omega \Psi_\varepsilon \omega'\} \geq d_\Psi(\omega, \omega') + \delta$$

which is obviously a contradiction.

- iii. Clearly, d_Ψ is a symmetric non-negative function, so let us show that d_Ψ satisfies the triangular inequality. By [ii] and triangularity of Ψ , for any points $\omega, \omega', \omega'' \in \Omega$ and any $\delta > 0$ it holds that $\omega \Psi_{d_\Psi(\omega, \omega') + d_\Psi(\omega', \omega'') + 2\delta} \omega''$ so that $d_\Psi(\omega, \omega'') \leq d_\Psi(\omega, \omega') + d_\Psi(\omega', \omega'') + 2\delta$. Since $\delta > 0$ is arbitrary, d_Ψ satisfies the triangular inequality. Finally, since Ψ is finite, $d_\Psi(\omega, \omega') < \infty$ for each $\omega, \omega' \in \Omega$, so d_Ψ satisfies all the conditions in the definition of pseudometrics.
- iv. The first assertion follows directly from the definition of d_Ψ , and the second one from the monotonicity of Ψ .
- v. Follows immediately from [iv].
- vi. Recall that d_Ψ is a metric iff $\Delta_{0, d_\Psi} = \Delta_\Omega$. Since Ψ_0 is an equivalence relation by [i], $\Delta_\Omega \subseteq \Psi_0$. Furthermore, by [iv] $\Psi_0 \subseteq \Delta_{0, d_\Psi}$, hence d_Ψ is a metric iff $\Delta_{0, d_\Psi} = \Psi_0 = \Delta_\Omega$, and the rest follows from [v].

□

Consider an arbitrary map $\varphi : \Omega' \rightarrow \Omega$, and given an ε -relation $\Psi = (\Psi_\varepsilon)_{\varepsilon \in \mathbb{R}_+}$ on Ω define an ε -relation $\Psi' = (\Psi'_\varepsilon)_{\varepsilon \in \mathbb{R}_+}$ on Ω' as follows: $\omega' \Psi'_\varepsilon \bar{\omega}'$ iff $\varphi(\omega') \Psi_\varepsilon \varphi(\bar{\omega}')$. In this case we say that Ψ' is a *pullback* of Ψ along φ . Pullback preserves many properties of ε -relations: Ψ' is reflexive (symmetric, triangular, monotone, finite, continuous at ρ) whenever Ψ is. Note also that as sets, $\Phi_\varepsilon = \varphi^{-1}(\Psi_\varepsilon)$ for all $\varepsilon \in \mathbb{R}_+$.

A.4 Modelling terminology

This thesis describes several models together with corresponding (bi-)simulation relations. Since the latter serve as sufficient conditions for behavioral inclusion (equivalence) over the underlying models, here we introduce some formal terminology for models, behaviors and related concepts we use throughout the whole thesis. This also helps supporting the ideas we have used to introduce behaviors of SSs above, and to connect model relations for different models.

Definition A.3 *A model is a triple $\text{Mod} = (\text{Syn}, \text{Spec}, \text{Sem})$ where Syn and Spec are arbitrary sets, and $\text{Sem} : \text{Syn} \times \text{Spec} \rightarrow \mathbb{R}$. We say that Syn is the syntax, Spec is the specification set and Sem is the semantics of Mod . The semantics Sem is said to be qualitative whenever the range of Sem is finite, and quantitative otherwise.*

To clarify concepts introduced in this section, a running example of the TS model is used. For a better fit we consider only pointed versions of the models below, that is those for which initial conditions are specified. For a fixed output set Y , the syntax of the model TS_Y consists exactly of all tuples $\mathcal{T} = (X, x_0, T, Y, L)$ where X

is an arbitrary set, $x_0 \in X$, $T \subseteq X \times X$ is an l.t.r. and $L : X \rightarrow Y$ (cf. Definition 2.1). The specification set of TS_Y is $2^{Y^{\mathbb{N}}}$ and the semantics is given by

$$\text{Sem}(\mathcal{T}, H) = \begin{cases} 1, & \text{if } \bigvee_L(\Sigma^T, x_0) \cap H \neq \emptyset \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A.5})$$

Clearly, the semantics of the TS_Y is qualitative since the range of Sem is $\{0, 1\}$ in this case. As a result, Sem can be identified with a satisfaction relation $\models \subseteq \text{Syn} \times \text{Spec}$, which is exactly what have been done for TSs above (cf. (2.1)). Notice that (A.5) does not directly define the semantics of TSs for each instance \mathcal{T} of the syntax, but in fact uses the set of output words $\bigvee_L(\Sigma^T, x_0)$ as a proxy for the definition of the semantics. This construction we generalize through the notion of behaviors.

Definition A.4 A behavior of a model Mod is a triple $\text{Beh} = (\mathfrak{B}, \mathfrak{b}_s, \mathfrak{b}_t)$ where \mathfrak{B} is an arbitrary set, and $\mathfrak{b}_s : \text{Syn} \rightarrow \mathfrak{B}$, $\mathfrak{b}_t : \mathfrak{B} \times \text{Spec} \rightarrow \mathbb{R}$. We say that \mathfrak{B} is the domain, \mathfrak{b}_s is the source map and \mathfrak{b}_t is the target map of Beh . The behavior Beh is called regular if

$$\mathfrak{b}_t \circ (\mathfrak{b}_s \sqcup \text{id}_{\text{Spec}}) = \text{Sem}. \quad (\text{A.6})$$

Let us motivate concepts in the definition above. First of all, \mathfrak{B} in Beh is a set of behaviors that we think is useful, a trivial example is a set of paths produced by a system, but for example in case of an SS a better choice is a set of probabilities measures on those paths. In turn, \mathfrak{b}_s connects between the model syntax and a particular subset of behaviors, i.e. it tells what is the set of behaviors for each model. Finally, the aim of \mathfrak{b}_t is similar to Sem , but instead of quantifying specifications on models directly, we now quantify the behaviors those models produce against the very specifications. As a result, (A.6) means that Sem and Beh are consistent.

Behaviors can be used as a proxy to define semantics of the model as follows. Suppose we are only given two sets Syn and Spec , which we would like to be the syntax and the specification set of our model respectively, but we have not defined its semantics so far. For any triple $\text{Beh} = (\mathfrak{B}, \mathfrak{b}_s, \mathfrak{b}_t)$ such that $\mathfrak{b}_s : \text{Syn} \rightarrow \mathfrak{B}$ and $\mathfrak{b}_t : \mathfrak{B} \times \text{Spec} \rightarrow \mathbb{R}$, we can define the semantics $\text{Sem}_{\text{Beh}} := \mathfrak{b}_t \circ (\mathfrak{b}_s \sqcup \text{id}_{\text{Spec}})$. Clearly, in such case Beh is a regular behavior for the model $(\text{Syn}, \text{Spec}, \text{Sem}_{\text{Beh}})$. Hence, by changing Beh we can come up with various models over $(\text{Syn}, \text{Spec})$ that differ in their semantics. The behavioral method of the semantics construction has been used for the TSs in Section 2.1, where we have set $\mathfrak{B} = 2^{Y^{\mathbb{N}}}$, $\mathfrak{b}_s(\mathcal{T}) = \bigvee_L(\Sigma^T, x_0)$ and $\mathfrak{b}_t(b, H) = 1\{b \cap H \neq \emptyset\}$. On the other hand, in Section 2.4 for SSs we have chosen $\mathfrak{B} = \mathcal{P}(Y^{\mathbb{N}})$, $\mathfrak{b}_s(\mathcal{S}) = S_L(\Sigma^\Gamma, \alpha)$ and $\mathfrak{b}_t(\mathcal{S}, h) = \sup_{q \in S_L(\Sigma^\Gamma, \alpha)} qh$, where now $\text{Spec} = \text{bU}_1(Y^{\mathbb{N}})$. We could have left \mathfrak{B} to be just the set of paths, but that would only allow us to come up with rather unnatural semantics, which perhaps are not rich enough to cover most of the interesting properties one would think of. Below we further elaborate on this point.

Recall from Section 2.1 that all versions of (bi-)simulation relations for TSs have been defined only in terms of syntax. Due to this reason, we refer to them as *syntactical relations* and to the corresponding pseudometrics as *syntactical pseudometrics*. Similarly, we say *semantical relations* for behavioral inclusion and equiv-

alence (exact or approximate), and *semantical pseudometrics* for the corresponding pseudometrics.

Whenever the bisimulation theory is developed for Mod , it may be of interest to use this theory in another modeling framework, say $\text{Mod}' = (\text{Syn}', \text{Spec}', \text{Sem}')$. In such case, one may want to interpret the new model Mod' in terms of Mod , for which the theory is available. Here this is done via the concept of interpretation maps.

Definition A.5 An interpretation map between two models $\mathcal{I} : \text{Mod}' \rightarrow \text{Mod}$ is a pair $\mathcal{I} = (\mathcal{I}_0, \mathcal{I}_1)$, where $\mathcal{I}_0 : \text{Syn}' \rightarrow \text{Syn}$ and $\mathcal{I}_1 : \text{Spec}' \rightarrow \text{Spec}$. We say that \mathcal{I} is regular whenever $\text{Sem}' = \text{Sem} \circ \mathcal{I}$. If $\text{Syn}' = \text{Syn}$ and $\mathcal{I}_0 = \text{id}_{\text{Syn}'}$, then a regular interpretation map is called a *semantical enlargement* of Mod' .

The idea of a semantical enlargement is that for a syntax of given model Mod' we would like to consider a new larger specification set Spec , leaving the semantics of the old specifications from Spec' untouched. An example of a semantical enlargement over TS_Y is given in the Appendix A.6. There, the qualitative semantics is enlarged to a more general quantitative one as follows: $\text{Spec}' = 2^{Y^{\mathbb{N}}}$, $\text{Spec} = \mathbb{R}^{Y^{\mathbb{N}}}$ and $\mathcal{I}_1(H') = 1_{H'}$ for each $H' \in \text{Spec}'$. Since the quantitative semantics over Spec is

$$\text{Sem}(\mathcal{I}, h) := \sup_{w \in \mathbb{V}_L(\Sigma^T, x_0)} h(w),$$

we obtain a regular interpretation map, hence a semantical enlargement.

As we have mentioned above, regular interpretation maps provide a tool to extend the bisimulation theory from one modelling framework to another. Suppose that Mod is some benchmark model, and let Φ and Ψ be two ε -relations on Syn that we regard as syntactical and semantical ones, respectively. Assume that $\Phi^\varepsilon \subseteq \Psi^\varepsilon$ for all $\varepsilon \in \mathbb{R}_+$, that is $S\Psi^\varepsilon \bar{S}$ whenever $S\Phi^\varepsilon \bar{S}$. For example, if $\text{Mod} = \text{TS}_Y$ then Φ and Ψ could be ε -simulation and ε -behavioral inclusion, respectively. Let Mod' be another model and $\mathcal{I} : \text{Mod}' \rightarrow \text{Mod}$ some interpretation map. Recall from the Appendix A.3 that we can pull back Φ from Syn to Syn' along \mathcal{I}_0 , and define a new ε -relation Φ' on Syn' satisfying $S'\Phi'^\varepsilon \bar{S}'$ iff $\mathcal{I}_0(S')\Phi^\varepsilon \mathcal{I}_0(\bar{S}')$. Similarly, let Ψ' be a pullback of Ψ along \mathcal{I}_0 . Since pullbacks of ε -relations preserve many useful properties, and in particular monotonicity, we obtain that $\Phi'^\varepsilon \subseteq \Psi'^\varepsilon$ for all $\varepsilon \in \mathbb{R}_+$. To highlight the importance of this fact, let us consider two examples of Ψ . Define Ψ_1 to be such that $S\Psi_1^\varepsilon \bar{S}$ iff

$$\text{Sem}(S, H) \leq \text{Sem}(\bar{S}, H) + \varepsilon \quad \forall H \in \text{Spec}.$$

For the second example, let Λ be an ε -relation on Spec and let $S\Psi_2^\varepsilon \bar{S}$ iff

$$\text{Sem}(S, H) \leq \text{Sem}(\bar{S}, \bar{H}) \quad \forall (H, \bar{H}) \in \Lambda^\varepsilon.$$

Instances of Ψ_1 and Ψ_2 are behavioral ε -inclusions for SSs and TSs, respectively. In particular, recall that for the latter case $\text{Spec} = 2^{Y^{\mathbb{N}}}$, so the corresponding example

of Λ would be $H\Lambda^\varepsilon\bar{H}$ iff $H^\varepsilon \subseteq \bar{H}$. In case \mathcal{I} is regular, we obtain that

$$\text{Sem}'(S', H') = \text{Sem}(\mathcal{I}_0(S'), \mathcal{I}_1(H')) \leq \text{Sem}(\mathcal{I}_0(\bar{S}'), \mathcal{I}_1(H')) + \varepsilon = \text{Sem}'(\bar{S}', H') + \varepsilon \quad (\text{A.7})$$

for all $H' \in \text{Spec}'$ whenever $S'\Psi_1^{\varepsilon}\bar{S}'$. Similarly,

$$\text{Sem}'(S', H') = \text{Sem}(\mathcal{I}_0(S'), \mathcal{I}_1(H')) \leq \text{Sem}(\mathcal{I}_0(\bar{S}'), \mathcal{I}_1(\bar{H}')) = \text{Sem}'(\bar{S}', \bar{H}') \quad (\text{A.8})$$

for all $(H', \bar{H}') \in \Lambda'^\varepsilon$ whenever $S'\Psi_2^{\varepsilon}\bar{S}'$, where Λ' is the pullback of Λ along \mathcal{I}_1 . As a result, whenever \mathcal{I} is a regular interpretation map and Φ and Ψ are syntactical and semantical relations such that Φ is stronger than Ψ , then their pullbacks Φ' and Ψ' have the same property. Using this fact, in Section 3.4 we have developed the bisimulation theory for MDPs based on that for SSs.

Interpretation maps do not only provide a way of extending bisimulation theory from one model to another. In addition, they can be used to define behaviors and semantics of a model. Suppose that we are only given $(\text{Syn}', \text{Spec}')$, whereas Sem' is left unspecified. For any map $\mathcal{I} = (\mathcal{I}_0, \mathcal{I}_1)$ satisfying $\mathcal{I}_0 : \text{Syn}' \rightarrow \text{Syn}$ and $\mathcal{I}_1 : \text{Spec}' \rightarrow \text{Spec}$, we can pull back the semantics from Mod to Mod' and define $\text{Sem}'_{\mathcal{I}} := \text{Sem} \circ \mathcal{I}$. Clearly, \mathcal{I} is a regular interpretation map for $(\text{Syn}', \text{Spec}', \text{Sem}'_{\mathcal{I}})$ and Mod . Assume that Beh is a regular behavior for Mod , then $\text{Beh}' = (\mathfrak{B}, \mathfrak{b}'_s, \mathfrak{b}'_t)$ given by

$$\mathfrak{b}'_s := \mathfrak{b}_s \circ \mathcal{I}_0, \quad \mathfrak{b}'_t := \mathfrak{b}_t \circ (\text{id}_{\mathfrak{B}} \sqcup \mathcal{I}_1)$$

is a regular behavior for $(\text{Syn}', \text{Spec}', \text{Sem}'_{\mathcal{I}})$. Even without given Spec' , one can always define $\text{Spec}' := \text{Spec}$ and let $\mathcal{I}_1 := \text{id}_{\text{Spec}}$. By changing \mathcal{I} we obtain various models for $(\text{Syn}', \text{Spec}')$ with different behaviors and semantics. This feature is extremely important when dealing with probabilistic models, in contrast to the case of non-probabilistic ones. Although in the latter framework there may still be several distinct semantics of interest, e.g. linear-time and branching-time semantics of TSs, such distinction is usually transparent and thus rarely casts the meaning of the model being confusing. At the same time, within the probabilistic framework one can introduce two rather different semantics that look very similar, which can cause issues when interpreting solution of verification or synthesis problems obtained over such models.

The situation described in the previous paragraph may happen when Mod' is endowed with a commonly accepted semantics that does not admit regular interpretation maps to a model for which the bisimulation theory is mature. Although in such case it may still be tempting to pull back the relations from Mod to Mod' , one also does pull back the new semantics to Mod' without caring about the old one. To give a concrete example, let us consider the approach to the bisimulation theory for probabilistic systems taken in [148] and the follow-up works. In their setting, Mod is the TS model whereas Mod' is a continuous controlled stochastic system. That is, Syn' comprises all tuples $S' = (n', f', g', U', Y', L')$, where f' and g' are the drift and diffusion terms in

$$dx'_t = f'(x'_t, u'_t)dt + g'(x'_t)dB'_t \quad (\text{A.9})$$

over the state space $\mathbb{R}^{n'}$ and the action space U' , Y' is some Borel (output) space and L' is a continuous (output) map. It would be natural to define $\text{Spec}' = \text{bU}(\mathcal{C}(Y'), \mathbb{R})$, behaviors based on decision strategies and corresponding strategic measures from $\mathcal{P}(\mathcal{C}(Y'))$ [55], and $\text{Sem}'(S', H')$ to be the maximal expectation of H' over all strategies/strategic measure available for S' . It is highly unlikely to find a regular interpretation map from Mod' to \mathcal{T}_Y . Due to this reason, [148] implicitly suggested an interpretation map constructed as follows. For a fixed time step δt , given $(n', f', g', U', Y', L') \in \text{Syn}'$ the corresponding TS (X, T, Y, L) is given by $X = \mathcal{P}(\mathbb{R}^n)$ endowed with a mean-square metric, $Y = X$ and $L = \text{id}_X$, whereas the transition relation T is such that $\tilde{x} \in T|_x$ iff there exists a solution of (A.9) such that x'_0 and x'_δ have distributions x and \tilde{x} , respectively.

In addition, [148] does not comment on how to embed the natural specifications for Mod' as those of \mathcal{T}_Y , and instead puts $\text{Spec}' = \text{Spec}$ by definition. Thus, only the syntax of Mod' is left whereas specifications, behaviors and semantics are all taken from Mod in order to develop the bisimulation theory over continuous-time stochastic systems. As a result, not only solutions of verification and synthesis problems over Mod' using relations pulled back from TSs are not suitable for the conventional semantics of probabilistic systems, but also their meaning is uncertain. For example, the solution of the safety problem would tell that the distribution of x'_t from (A.9) belongs to the safe set $A \subseteq X$ (of probability distributions) for each $t \in \mathbb{R}_+$. It is rather unclear how such a statement shall be interpreted over a real-life process modeled as a probabilistic system. However, in case A is an ε -inflation of “deterministic points” that is Dirac distributions $A' := (\eta(a_1), \dots, \eta(a_n))$, one may misinterpret the safety result above as follows: x'_t belongs to the ε -inflation of A' (in the Euclidian metric) for all $t \in \mathbb{R}_+$ which is a much stronger result⁶. Due to this reason, a person not experienced in probability theory may overestimate the guarantees provided by [148] and the follow-up works.

Another interesting example of an irregular interpretation map is inspired by [90]. Let Mod' be the SS model from Section 2.4 that only consists of discrete SSs, and let Mod be the semantical enlargement of the TS model to the quantitative semantics as in the Appendix A.6, so that $\text{Spec}' = \text{Spec}$. Define \mathcal{J} by $\mathcal{J}_1 = \text{id}_{\text{Spec}}$ and $\mathcal{J}_0(\mathcal{S}) = (X, T, Y, L)$ for $\mathcal{S} = (X, \Gamma, Y, L)$, where

$$T|_x = \{\tilde{x} \in X : \gamma(\tilde{x}) > 0 \text{ for some } \gamma \in \Gamma|_x\} \quad x \in X.$$

Note that we cannot define \mathcal{J} as above for all SSs since T is not an l.t.r. if $\Gamma|_x$ does not contain a discrete distribution for some x . As a result, the transition from x to \tilde{x} is allowed in $\mathcal{J}_0(\mathcal{S})$ iff it can happen with a positive probability in \mathcal{S} , that is $\mathcal{J}_0(\mathcal{S})$ contains all “possible” transitions of \mathcal{M} . One can say that the map \mathcal{J} “forgets” the probabilistic structure of SSs⁷. Recall discussion in the beginning of this section, where we have mentioned that one could have define $\mathcal{B} = 2^{Y^{\mathbb{N}}}$ for SSs as well, but that would have caused problems with the semantics. Pulling back the bisimulation from TSs to SSs along \mathcal{J} , we obtain the exact bisimulation

⁶ A related discussion can be found in [149, Sections 5.1, 5.3].

⁷ This statement can be made more formal by treating \mathcal{J} as a functor between two categories, e.g. if one treats MDPs and TSs as coalgebras (see Appendix B), so that \mathcal{J} would be a forgetful functor.

in the terminology of [90], whereas their probabilistic bisimulation is the exact bisimulation of SSs in our terminology. The issues of bisimulation pulled back to probabilistic systems along such forgetful interpretation maps are discussed in Section 3.4 on the example of MDPs.

Let us further comment on PTSs introduced in [90]. The notions of exact and approximate bisimulations for PTSs were developed purely based on their syntax, which is similar to that of MDPs. One could think that it is easy to provide an interpretation map for such case, or at least its \mathcal{I}_0 component, however since the PTSs are not given concise semantics, it is not clear what to take as Spec and Sem in this case. One can still pull back exact and approximate bisimulation from PTSs to MDPs just based on \mathcal{I}_0 , but such relations are not guaranteed to be of any use for MDPs.

A.5 Weak inclusion of families of probability measures

In this section we introduce two useful ε -relations (cf. the Appendix A.3) between families of probability measures. The first one is the *weak inclusion* which generalizes a usual inclusion of families of measures as sets, and the second one is *weak equivalence* which symmetrizes the weak inclusion. We prove some basic results regarding these concepts, and in particular define the corresponding pseudometric.

Let Ω be an arbitrary Borel space, and let $P, \bar{P} \subseteq \mathcal{P}(\Omega)$ be two families of Borel probability measures on Ω . For $\varepsilon \in \mathbb{R}_+$ we say that P is *weakly ε -included* in \bar{P} , and write $P \sqsubseteq_\varepsilon \bar{P}$, if $P^*f \leq \bar{P}^*f + \varepsilon$ for any $f \in \text{b}\mathcal{B}_1(\Omega)$.

We start with some fundamental properties of the weak inclusion $\sqsubseteq := (\sqsubseteq_\varepsilon)_{\varepsilon \in \mathbb{R}_+}$.

Proposition A.6 *Consider some Borel space Ω and $P, \bar{P} \subseteq \mathcal{P}(\Omega)$. It holds that*

- i. *if $P \subseteq \bar{P}$ then $P \sqsubseteq_0 \bar{P}$;*
- ii. *\sqsubseteq is a reflexive, monotone and triangular ε -relation;*
- iii. *\sqsubseteq is a continuous ε -relation;*
- iv. *if P, \bar{P} are non-empty then $P \sqsubseteq_2 \bar{P}$;*
- v. *if $P \sqsubseteq_\varepsilon \bar{P}$ for some $\varepsilon \in \mathbb{R}_+$, then $P^*f \leq \bar{P}^*f + r\varepsilon$ for any $r \in \mathbb{R}_+$ and $f \in \text{b}\mathcal{B}_r(\Omega)$; in particular $P^*1_B \leq \bar{P}^*1_B + \varepsilon/2$ for any $B \in \mathcal{B}(\Omega)$.*

Proof: The first two statements follow immediately from the definition of \sqsubseteq_ε . For the rest, let $f \in \text{b}\mathcal{B}_1(\Omega)$ be an arbitrary function.

With focus on [ii], if $P \sqsubseteq_\varepsilon \bar{P}$ for all $\varepsilon > \rho$ then $P^*f \leq \bar{P}^*f + \varepsilon$ for all $\varepsilon > \rho$ and hence $P^*f \leq \bar{P}^*f + \rho$. Since f is arbitrary, $P \sqsubseteq \bar{P}$.

Regarding [iv], since P and \bar{P} are non-empty we have $P^*f, \bar{P}^*f \in \mathbb{R}$. Moreover, for any $p \in P$ and $\bar{p} \in \bar{P}$:

$$|pf - \bar{p}f| \leq \|f\| \cdot \|p - \bar{p}\| \leq 2$$

and hence $|P^*f - \bar{P}^*f| \leq 2$ by Lemma C.16, which in particular leads to [iv].

Finally, for [vi] the case $r = 0$ is trivial. Let $r > 0$, then $f \in \text{b}\mathcal{B}_r(\Omega)$ iff $f/r \in \text{b}\mathcal{B}_1(\Omega)$. Since $P \sqsubseteq_\varepsilon \bar{P}$, we obtain $P^*(f/r) \leq \bar{P}^*(f/r) + \varepsilon$ which is equivalent to

$$P^*f \leq \bar{P}^*f + r\varepsilon$$

as $r > 0$. Now, let $B \in \mathcal{B}(X)$ be any set and define $f := 1_B - \frac{1}{2}$. Clearly, $f \in \text{b}\mathcal{B}_{\frac{1}{2}}(X)$ and hence the desired result follows from the shift-invariance of P^* and \bar{P}^* . \square

In accordance to the terminology of the Appendix A.3, \sqsubseteq is a continuous ε -preorder. Furthermore, Proposition A.6.[iv] implies that \sqsubseteq_2 is the trivial relation on the collection of non-empty families of probability measures. Finally, although the weak inclusion is introduced in terms of functions $f \in \text{b}\mathcal{B}_1(\Omega)$, the scaling property in Proposition A.6.[vi] shows that \sqsubseteq can be also used to argue about any class of functions with a known uniform bound.

Let us discuss which operations on families of measures preserve the ε -relation \sqsubseteq .

Proposition A.7 *Consider some $P, \bar{P} \subseteq \mathcal{P}(\Omega)$ and $\varepsilon \in \mathbb{R}_+$, and let Ξ be an arbitrary Borel space. If $\bar{P} \sqsubseteq_\varepsilon P$ then*

- i. $\varphi_*\bar{P} \sqsubseteq_\varepsilon \varphi_*P$ for any $\varphi \in \mathcal{B}(\Omega, \Xi)$;
- ii. $\text{sco } P \sqsubseteq_\varepsilon \bar{P}$.

Proof: The proof is as follows:

- i. Consider any $g \in \text{b}\mathcal{B}_1(\Xi)$ and recall that $(\varphi_*p)g = p[g \circ \varphi]$ for all $p \in \mathcal{P}(\Omega)$. Since $g \circ \varphi \in \text{b}\mathcal{B}_1(\Omega)$ for any $g \in \text{b}\mathcal{B}_1(\Xi)$, we obtain the desired result.
- ii. Let $f \in \text{b}\mathcal{B}_1(\Omega)$ be an arbitrary function and consider some measure $p \in \text{sco } P$, then there exists $\nu \in \mathcal{P}(\mathcal{P}(\Omega))$ such that $pf = \int_P p'f\nu(dp')$. Thus

$$pf \leq \int_P (\bar{P}^*f + \varepsilon)d\nu = \bar{P}^*f + \varepsilon$$

and so $(\text{sco } P)^*f \leq \bar{P}^*f + \varepsilon$. Since f is arbitrary, we obtain that $\text{sco } P \sqsubseteq_\varepsilon \bar{P}$. \square

We denote the symmetrization of \sqsubseteq by \equiv and say that $P, \bar{P} \subseteq \mathcal{P}(\Omega)$ are *weakly ε -equivalent* whenever $P \equiv_\varepsilon \bar{P}$. In particular this means that

$$|P^*f - \bar{P}^*f| \leq \varepsilon \quad \forall f \in \text{b}\mathcal{B}_1(\Omega).$$

It follows from Proposition A.6 that \equiv is a continuous ε -equivalence, and hence induces a pseudometric d_{\equiv} on $\mathcal{P}(\Omega)$. Let us briefly summarize the rest of basic properties of \equiv .

Proposition A.8 *For any Borel space Ω and families $P, \bar{P}, \hat{P} \subseteq \mathcal{P}(\Omega)$:*

- i. *if P, \bar{P} are non-empty then $P \equiv_2 \bar{P}$;*
- ii. *if $P \equiv_{\varepsilon} \bar{P}$ for some $\varepsilon \in \mathbb{R}_+$, then $|P^*f - \bar{P}^*f| \leq r\varepsilon$ for all $r \in \mathbb{R}_+$ and $f \in \mathbf{b}\mathcal{B}_r(\Omega)$; in particular $|P^*1_B - \bar{P}^*1_B| \leq \varepsilon/2$ for any $B \in \mathcal{B}(\Omega)$;*
- iii. *if $P \equiv_{\varepsilon} \bar{P}$ for some $\varepsilon \in \mathbb{R}_+$ and Ξ is some Borel space, then $\varphi_*P \equiv_{\varepsilon} \varphi_*\bar{P}$ for any Borel map $\varphi : \Omega \rightarrow \Xi$;*
- iv. *$\text{sco } P \equiv_0 P$.*

Proof: The proof immediately follows from the corresponding properties of \sqsubseteq . \square

A.6 Strong and weak inclusions

Let us comment on the difference between the weak inclusion and the usual subset inclusion on the example of $\mathcal{P}(\Omega)$ and use it to underline the similarity between behavioral comparisons for TSs and SSs. Endow $\mathcal{P}(\Omega)$ with the total variation distance d_{TV} and recall that we say that P is ε -included in \bar{P} whenever $P \subseteq \bar{P}^{\varepsilon}$. From Proposition A.1 it follows that this is equivalent to

$$\sup_{p \in P} G(p) \leq \sup_{\bar{p} \in \bar{P}} G(\bar{p}) + \varepsilon \quad \forall g \in \mathcal{G}, \quad (\text{A.10})$$

where $\mathcal{G} = \text{Lip}_1(\mathcal{P}(\Omega), d_{\text{TV}})$. In contrast, for the weak inclusion instead of arbitrary $g \in \text{Lip}_1(\mathcal{P}(\Omega), d_{\text{TV}})$ we require (A.10) to hold only for functions of the form $G(p) = pf$ for $f \in \mathbf{b}\mathcal{B}_1(\Omega)$. Clearly, g of such form always belong to $\text{Lip}_1(\mathcal{P}(\Omega), d_{\text{TV}})$, moreover it is a linear function, however not every such linear function admits this representation⁸. Hence, both weak and strong ε -inclusions can be defined as special cases of more general condition (A.10), where \mathcal{G} is some class of functions. As a result, we see that approximate behavioral inclusions for TSs and SSs are just versions of the following condition

$$\sup_{b \in \mathbf{b}_s(\text{Syn})} G(b) \leq \sup_{b \in \mathbf{b}_s(\text{Syn})} G(b) + \varepsilon \quad \forall g \in \mathcal{G}$$

⁸ As an example, consider a functional that maps the measure to its discrete part

$$g : p \mapsto \sup\{p(S) : S \subseteq \Omega \text{ is countable}\}.$$

Linearity can be verified by noticing, that the maximum is always achieved, which also help to show the Lip_1 property. To prove that there does not exist f such that $g(p) = pf$, take $\Omega = [0, 1]$ and consider $p = a\delta(0) + (1-a)\lambda$ for different $a \in [0, 1]$, where λ is the Lebesgue measure. What is $f(0)$?

for some class \mathcal{G} of functions $g : \mathfrak{B} \rightarrow \mathbb{R}$. Here we use the terminology of Appendix A.4, in particular $\mathfrak{b}_s(\text{Syn}) \subseteq \mathfrak{B}$ is a set of all behaviors of some generic model with a syntax Syn .

Categorical concepts

B.1 Categories

Category theory provides a unified approach to study different mathematical structures that yet have similar features. A category \mathbf{C} consists of objects and morphisms (or arrows) relate by the following fundamental maps:

- dom and cod assign to each morphism φ objects $\Omega = \text{dom}(\varphi)$ and $\Xi = \text{cod}(\varphi)$ called the *domain* and *codomain* of φ . We write $\varphi \in \text{Hom}(\Omega, \Xi)$ to mean that $\text{dom}(\varphi) = \Omega$ and $\text{cod}(\varphi) = \Xi$.
- \circ assigns to each pair of morphisms (φ, ψ) satisfying $\text{dom}(\psi) = \text{cod}(\varphi)$ a morphism $\psi \circ \varphi$ called the *composition* of ψ and φ . The composition of morphisms must be associative: $(\chi \circ \psi) \circ \varphi = \chi \circ (\psi \circ \varphi)$.
- id assigns to each object Ω the *identity* morphism $\text{id}_\Omega \in \text{Hom}(\Omega, \Omega)$ satisfying $\varphi \circ \text{id}_\Omega = \text{id}_\Xi \circ \varphi = \varphi$ for any morphism $\varphi \in \text{Hom}(\Omega, \Xi)$.

For each $n \in \mathbb{N}$ let \mathbf{C}_n denote the collection of n -tuples of \mathbf{C} -morphisms $(\varphi_1, \dots, \varphi_n)$ satisfying $\text{dom}(\varphi_{k+1}) = \text{cod}(\varphi_k)$ for $k \in [0; n-1]$, with the convention that \mathbf{C}_0 is the collection of objects of \mathbf{C} .

One of the simplest examples of categories is a category of sets, denoted by \mathbf{Set} ; there objects are sets, morphisms are functions and composition and identity have usual meanings. For example, a morphism $\varphi \in \mathbf{Set}_1$ given by a function $\varphi : \Omega \rightarrow \Xi$ satisfies $\text{dom}(\varphi) = \Omega$ and $\text{cod}(\varphi) = \Xi$. Similarly, the category $\mathbf{Mes}(\mathbf{Top})$ of measurable (topological) spaces has sets endowed with σ -algebras (topologies) as objects, and measurable (continuous) maps as morphisms: recall that composition of measurable (continuous) maps is measurable (continuous) as well, so the property of composition of morphisms holds in this case. The category \mathbf{Rel} provides an example where morphisms are not functions; it has sets as objects and relations as morphisms. The composition of morphisms in \mathbf{Rel} is given by (A.4), and the diagonal of a set is its identity morphism. For more examples of categories and an accessible introduction to category theory see e.g. [1].

One of the key features of categories is that one can define diverse abstract constructions and prove a multitude of useful facts about all categories in one go just in terms of objects and morphisms. For example, a *product* of two objects $\Omega, \bar{\Omega} \in C_0$ is a triple $(\Xi, \psi, \bar{\psi})$ with $\psi \in \text{Hom}(\Xi, \Omega)$ and $\bar{\psi} \in \text{Hom}(\Xi, \bar{\Omega})$ such that for any other triple $(\Xi', \psi', \bar{\psi}')$ of that kind there exists a unique $\chi \in \text{Hom}(\Xi', \Xi)$ satisfying $\psi' = \psi \circ \chi$ and $\bar{\psi}' = \bar{\psi} \circ \chi$. In case of Set we obtain $\Xi = \Omega \times \bar{\Omega}$ is a usual Cartesian product of sets and $\psi = \text{proj}_\Omega, \bar{\psi} = \text{proj}_{\bar{\Omega}}$ are usual projection maps. For more complex structures such as $\text{Mes}(\text{Top})$, the categorical product of two measurable (topological) spaces is a usual Cartesian product endowed with the classical product σ -algebra (topology). At the same time, the product in Rel has Ξ being a disjoint union of sets.

The product construction is one of the basic in category theory, and yet it shows that natural objects (such as products of spaces) can be unambiguously defined solely in terms of morphisms. Hence, instead of showing some properties of a product separately in measurable and topological spaces, one can try establishing such a property purely in categorical terms, so that it automatically holds for all instances of categories with products as special cases. The very same logic applies to other constructions; although arguably the biggest success category theory had is in algebraic topology and related fields, it appeared to be useful to relate existing and introduce new concepts in measure theory, probability theory, dynamical systems and logic, which is the reason we discuss categories in this work. In particular, the categorical notion of *co-algebra* discussed below is useful in representing dynamical systems and studying their behaviors and corresponding equivalence relations. It is in our interest to relate the framework developed in this thesis to the co-algebraic approach.

To introduce co-algebras, we need to discuss a couple of another categorical concepts. First of all, some objects are equivalent from the point of view of category theory, which can be formalized via the notion of *isomorphism*. An morphism $\varphi \in \text{Hom}(\Omega, \Xi)$ is said to be an isomorphism if there exists an *inverse* morphism $\psi \in \text{Hom}(\Xi, \Omega)$ satisfying $\psi \circ \varphi = \text{id}_\Omega$ and $\varphi \circ \psi = \text{id}_\Xi$. Two objects are called isomorphic if there exists an isomorphism between them. For example, in Set an isomorphism is any bijection so that two sets are isomorphic iff they have the same cardinality, whereas in Rel an isomorphism is any l.t.r. whose inverse it l.t.r. as well so that any two sets are isomorphic. In $\text{Mes}(\text{Top})$ an isomorphism is a measurable (continuous) bijection whose inverse is measurable (continuous) as well; in case of Top isomorphisms are exactly homeomorphisms.

We say that $1 \in C_0$ is the *final object* of C if for any $\Omega \in C_0$ it holds that $\text{Hom}(\Omega, 1)$ has cardinality of 1, that is from any object there exists exactly one morphism to a final object; such morphism is often denoted as $!_\Omega$. Clearly, final objects (whenever they exist) are defined up to isomorphism. For example, in Set , Rel , Mes and Top singletons are final objects.

As much as morphisms can be considered abstract versions of maps between sets, *functors* are regarded as maps between categories. A functor $\mathcal{F} : C \rightarrow \bar{C}$ assigns to each object $\Omega \in C_0$ an object $\mathcal{F}(\Omega) \in \bar{C}_0$, and to each morphism $\varphi \in \text{Hom}(\Omega, \Xi) \subseteq C_0$ a morphism $\mathcal{F}(\varphi) \in \text{Hom}(\mathcal{F}(\Omega), \mathcal{F}(\Xi)) \subseteq \bar{C}_1$ such that $\mathcal{F}(\text{id}_\Omega) = \text{id}_{\mathcal{F}(\Omega)}$ and $\mathcal{F}(\psi \circ \varphi) = \mathcal{F}(\psi) \circ \mathcal{F}(\varphi)$. As a result, to define a functor one needs to define its action

both on objects, and on morphisms. In addition, its action on morphisms must preserve identities and composition. For example, the functor $\mathcal{B} : \text{Top} \rightarrow \text{Mes}$ does not change morphisms and assigns to any topological space object in Top_0 a measurable space object in Mes_0 with the same underlying set endowed with the Borel σ -algebra generated by the topology of the original object. A functor $\mathcal{F} : \mathcal{C} \rightarrow \mathcal{C}$ is called an *endofunctor* on \mathcal{C} . For example, if \mathcal{C} is a category where object are Borel spaces and morphisms are measurable maps, then an endofunctor $\mathcal{P} : \text{Mes} \rightarrow \text{Mes}$ sends a Borel space Ω to $\mathcal{P}(\Omega)$: the set of probability measures over Ω endowed with the topology of weak convergence with respect to the topology of Ω , and a map $\varphi : \Omega \rightarrow \Xi$ to a pushforward map $\varphi_* : \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\Xi)$. One of our results below shows that the latter endofunctor can be generalized to relations, which is useful to use when working with MDPs.

Given an endofunctor \mathcal{F} on some category \mathcal{C} , an \mathcal{F} -*coalgebra* is a pair (Ω, φ) where $\varphi \in \text{Hom}(\Omega, \mathcal{F}(\Omega))$. Coalgebras appeared to be extremely useful in describing dynamical systems, where one considers $\mathcal{F}(\Omega)$ to be a dynamical structure on Ω . For example, in Set let Y be some fixed set of observations or label, and define $\mathcal{F}(\Omega) := \Omega \times Y$ with $\mathcal{F}(\psi) := (\psi \sqcup \text{id}_Y)$. In that case φ assigns to any point $\omega \in \Omega$ its successor $\varphi_0(\omega) \in \Omega$ and its observation $\varphi_1(\omega) \in Y$. For each category \mathcal{C} we can define a new category of its \mathcal{F} -coalgebras, latter being objects in this category. Morphisms between coalgebras are exactly the maps between objects of \mathcal{C} that are consistent with \mathcal{F} , i.e. $(\Omega, \varphi) \xrightarrow{\psi} (\bar{\Omega}, \bar{\varphi})$ iff $\bar{\varphi} \circ \psi = \mathcal{F}(\psi) \circ \varphi$. Those morphisms were noticed to correspond to the known notions of bisimulation by model represented by coalgebras, see e.g. [37], [139] and [68]. This results further inspired the following approach: if a model of interest can be encoded through a coalgebra in a suitable category, then one can just use a coalgebraic definition of bisimulation, and hence get that theory for the desired models. However, apparently notions of bisimulation – let alone approximate bisimulations – as being dependent on the syntax of the model, not necessarily agree with its semantics, as the latter is not always unique for the given syntax (nor does it have to be, see discussion in Appendix A.4). As a result, coalgebras provide just *one possible* way to define bisimulations, not necessarily the most natural or useful one.

With focus on the approximate notions, it is worth mentioning that some of the aforementioned works suggested using the final coalgebra in pseudometric categories to introduce pseudometrics between the models: when the latter are encoded as coalgebras, each has exactly one morphism to the final coalgebra, so if it has a pseudometric structure, the distance between coalgebras can be defined as distance between the “images” of them under those morphisms. From pseudometric one can hope to go to bisimulations, or directly to the approximation bounds. This approach is unlikely to be useful in practice, since as a first step one has to show existence of a final coalgebra – see e.g. [104] for such methods. But even if such an object does exist, it is neither clear how to compute the desired distances efficiently (see [141, 143, 144, 140]), nor how to use them even one happens to finish those computations.

Below we present theory for the functor that sends analytic relations between states to those between measures on state spaces, which is extensively used in Chapter 3. Even though we show that this is indeed a well-defined functor in

a certain category, and likely stochastic systems can be described as coalgebras w.r.t. this functor, due to the reason we are not certain about the usefulness of the coalgebraic approach, we do not particularly focus on studying the coalgebraic properties of that functor.

B.2 Category of Borel relations

As we have mentioned above, the ε -relations are important for us on the level of systems. To show that two systems belong to a particular ε -relation, we deal with usual (exact) relations on the level of states. The latter we can lift from states to measures, the idea which seems to be mentioned for the first time over discrete probability spaces in [75] for purposes similar to ours. Independently, such lifting procedure was studied in a greater generality over Polish spaces in [93] as an extension of the notion of stochastic orders (see e.g. [115]). Before comparing our account with the aforementioned works, let us first introduce the lifting procedure.

For any two Borel spaces Ω and $\bar{\Omega}$, and a relation $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$ we define a corresponding relation $\Phi_* \subseteq \mathcal{P}(\Omega) \times \mathcal{P}(\bar{\Omega})$ as follows: $(\mu, \bar{\mu}) \in \Phi_*$ if there exists a coupling $M \in \mathfrak{C}(\mu, \bar{\mu})$ such that $M(\Phi) = 1$. Equivalently, two probability measures μ and $\bar{\mu}$ are related via Φ_* whenever there exist two random variables \mathfrak{w} and $\bar{\mathfrak{w}}$ on the common probability space such that $\mathfrak{w}(\bar{\mathfrak{w}})$ is distributed according to μ ($\bar{\mu}$) and $(\mathfrak{w}, \bar{\mathfrak{w}}) \in \Phi$ (a.s.). Clearly, \mathcal{P} is a monotone operation with respect to set inclusion. Let us consider two important special cases of \mathcal{P} applied to relations.

Example B.1 *It holds that $\delta(\omega)\Phi_*\delta(\bar{\omega})$ iff $\omega\Phi\bar{\omega}$. As a result, Φ_* indeed can be seen as an extension of relations between points, embedded as Dirac distributions, to more general probability measures. In case only one of measures is Dirac, the following useful property holds true: $\delta(\omega)\Phi_*\bar{\mu}$ iff $\bar{\mu}(\Phi|_{\omega}) = 1$, where $\bar{\mu} \in \mathcal{P}(\bar{\Omega})$ is arbitrary.*

Example B.2 *Let μ and $\bar{\mu}$ be arbitrary and let $\Phi := \text{Gr}(\varphi)$ for some $\varphi \in \mathcal{B}(\Omega, \bar{\Omega})$. In this case $\mu\Phi_*\bar{\mu}$ iff $\bar{\mu} = \varphi_*\mu$. To show this define $M := (\text{id}_{\Omega} \times \varphi)_*\mu$. We obtain:*

$$M(\text{Gr}(\varphi)) = \mu((\text{id}_{\Omega} \times \varphi)^{-1}(\text{Gr}(\varphi))) = \mu(\Omega) = 1.$$

Let us now show that $M \in \mathfrak{C}(\mu, \bar{\mu})$:

$$\begin{aligned} (\text{proj}_{\Omega})_*M &= (\text{proj}_{\Omega} \circ (\text{id}_{\Omega} \times \varphi))_*\mu = (\text{id}_{\Omega})_*\mu = \mu, \\ (\text{proj}_{\bar{\Omega}})_*M &= (\text{proj}_{\bar{\Omega}} \circ (\text{id}_{\Omega} \times \varphi))_*\mu = \varphi_*\mu = \bar{\mu}. \end{aligned}$$

Conversely, if $\mu\Phi_\bar{\mu}$ and $M \in \mathfrak{C}(\mu, \bar{\mu})$ is such that $M(\text{Gr}(\varphi)) = 1$, then*

$$\begin{aligned} \bar{\mu}(\bar{B}) &= M(\Omega \times \bar{B}) = M(\text{Gr}(\varphi) \cap (\Omega \times \bar{B})) = M((\omega, \bar{\omega}) : \bar{\omega} = \varphi(\omega) \text{ and } \bar{\omega} \in \bar{B}) \\ &= M(\Omega \times \varphi^{-1}(\bar{B})) = \mu(\varphi^{-1}(\bar{B})) = \varphi_*\mu. \end{aligned}$$

As a result, $\mu\mathcal{P}(\text{Gr}(\varphi))\bar{\mu}$ is equivalent to $\bar{\mu} = \varphi_*\mu$ and hence relations between measures can be seen as a generalization of pushforwards, as much as relations between points can be seen as generalizations of the usual maps.

The coupling M is referred to as a "weight function" in works on discrete probability spaces that followed up on [75], e.g. [81] which among others interprets M as a solution of a maximal flow problem, but does not mention a perhaps more intuitive interpretation we have provided above. Such interpretation is yet provided in [93]; the difference between our approaches is that our exposition is slightly more general, as we consider relations being analytic subsets of Borel spaces rather than closed subsets of Polish spaces. It is likely that focusing on the latter framework would not reduce much the practical problems we can tackle here, but it might require certain auxiliary continuity assumptions on SSs, which we prefer to avoid. At the same time, working with non-closed measurable relations is more technically involved. For example, we cannot focus only on Borel relations: as we show now, their composition may fail to be Borel again, whereas compositions of analytic relations are always analytic.

We can in fact define a category where the objects are Borel spaces and arrows are analytic relations the are composed according to (A.4), we further refer to this category as BoRel. Our main result here concerns the fact that BoRel is indeed a category, and that \mathcal{P} is an endofunctor in this category. To show this, we need some supplementary statements that are of use for us throughout the thesis.

Lemma B.3 *Let $\Omega, \bar{\Omega}, \hat{\Omega}$ be arbitrary Borel spaces. For any two relations $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$ and $\Phi' \in \mathcal{A}(\bar{\Omega} \times \hat{\Omega})$ their composition satisfies $\Phi' \circ \Phi \in \mathcal{A}(\Omega \times \hat{\Omega})$.*

Proof: The composition of relations can be equivalently expressed as follows:

$$\Phi' \circ \Phi = \text{proj}_{\Omega \times \hat{\Omega}} \left(\left(\Phi \times \hat{\Omega} \right) \cap \left(\Omega \times \Phi' \right) \right). \quad (\text{B.1})$$

Note that $\Phi \times \hat{\Omega}$ is analytic as a product of analytic sets, and so is the intersection in parentheses. Hence $\Phi' \circ \Phi$ is analytic as a projection of an analytic set. \square

Theorem B.4 *The category BoRel where objects are Borel spaces, morphisms are analytic relations with identities being diagonals and compositions given by (A.4), is well-defined.*

Proof: To prove the theorem we only need to verify that morphisms satisfy the desired property. The associativity of the composition follows from Lemma B.3, and the diagonal is an identity since it is in Rel since diagonals of Borel spaces are Borel sets. \square

Let us show now that \mathcal{P} is an endofunctor on BoRel. We first show that \mathcal{P} sends morphisms to morphisms, that is it preserves the analyticity of relations.

Lemma B.5 $\Phi_* \in \mathcal{A}(\mathcal{P}(\Omega) \times \mathcal{P}(\bar{\Omega}))$ for any Borel spaces $\Omega, \bar{\Omega}$ and $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$.

Proof: Let $A := (\mathfrak{c}_\Phi)^{-1}(1)$ be set of all measures in $\mathcal{P}(\Omega \times \bar{\Omega})$ concentrated on Φ , so that $A \in \mathcal{A}(\mathcal{P}(\Omega \times \bar{\Omega}))$. Notice further that

$$\Phi_* = ((\text{proj}_\Omega)_* \times (\text{proj}_{\bar{\Omega}})_*)(A).$$

From Lemma C.15 we obtain that $(\text{proj}_\Omega)_* \times (\text{proj}_{\bar{\Omega}})_* \in \mathcal{B}(\mathcal{P}(\Omega \times \bar{\Omega}), \mathcal{P}(\Omega) \times \mathcal{P}(\bar{\Omega}))$ as a product of two Borel maps, so Φ_* is an analytic set being the image of the analytic set A under a Borel map. \square

Note that $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$ iff $\Phi^{-1} \in \mathcal{A}(\bar{\Omega} \times \Omega)$. It follows immediately from the definition of \mathcal{P} that $\mathcal{P}(\Phi^{-1}) = \Phi_*^{-1}$. The proof that \mathcal{P} does not only preserve inverses, but also compositions, requires some work and is presented in the two next lemmas.

Lemma B.6 *Consider some Borel spaces $\Omega, \bar{\Omega}, \hat{\Omega}$, $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$ and $\Phi' \in \mathcal{A}(\bar{\Omega} \times \hat{\Omega})$. For any measures $\mu \in \mathcal{P}(\Omega)$, $\bar{\mu} \in \mathcal{P}(\bar{\Omega})$ and $\hat{\mu} \in \mathcal{P}(\hat{\Omega})$ it holds that*

$$\mu\Phi_*\bar{\mu} \text{ and } \bar{\mu}\Phi'_*\hat{\mu} \implies \mu(\Phi' \circ \Phi)_*\hat{\mu}.$$

Proof: Let $M \in \mathfrak{C}(\mu, \bar{\mu})$ and $M' \in \mathfrak{C}(\bar{\mu}, \hat{\mu})$ be such that $M(\Phi) = 1$ and $M'(\Phi') = 1$. It follows from [126] that there exists a conditional coupling of M and M' , that is a measure $\mathbb{M} \in \mathcal{P}(\Omega \times \bar{\Omega} \times \hat{\Omega})$ satisfying $(\text{proj}_{\Omega \times \bar{\Omega}})_*\mathbb{M} = M$ and $(\text{proj}_{\bar{\Omega} \times \hat{\Omega}})_*\mathbb{M} = M'$, so that

$$(\text{proj}_\Omega)_*\mathbb{M} = \mu, \quad (\text{proj}_{\bar{\Omega}})_*\mathbb{M} = \bar{\mu}, \quad (\text{proj}_{\hat{\Omega}})_*\mathbb{M} = \hat{\mu}.$$

Let us define $M'' := (\text{proj}_{\Omega \times \hat{\Omega}})_*\mathbb{M}$, so clearly $M'' \in \mathfrak{C}(\mu, \hat{\mu})$. Furthermore,

$$\begin{aligned} M''(\Phi' \circ \Phi) &= \mathbb{M}((\omega, \bar{\omega}, \hat{\omega}) : (\omega, \hat{\omega}) \in \Phi' \circ \Phi) \\ &\geq \mathbb{M}((\omega, \bar{\omega}, \hat{\omega}) : (\omega, \bar{\omega}) \in \Phi \text{ and } (\bar{\omega}, \hat{\omega}) \in \Phi') = 1 \end{aligned}$$

since $\mathbb{M}((\omega, \bar{\omega}) \in \Phi) = M(\Phi) = 1$ and $\mathbb{M}((\bar{\omega}, \hat{\omega}) \in \Phi') = M'(\Phi') = 1$. As a result, we obtain that $M'' \in \mathfrak{C}(\mu, \hat{\mu})$ and $M''(\Phi' \circ \Phi) = 1$ so that $\mu(\Phi' \circ \Phi)_*\hat{\mu}$ as desired. \square

Lemma B.7 *Let $\Omega, \bar{\Omega}, \hat{\Omega}$ be arbitrary Borel spaces, $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$ and $\Phi' \in \mathcal{A}(\bar{\Omega} \times \hat{\Omega})$. For any two measures $\mu \in \mathcal{P}(\Omega)$ and $\hat{\mu} \in \mathcal{P}(\hat{\Omega})$ satisfying $\mu(\Phi' \circ \Phi)_*\hat{\mu}$ there exists a third probability measure $\bar{\mu} \in \mathcal{P}(\bar{\Omega})$ such that $\mu\Phi_*\bar{\mu}$ and $\bar{\mu}\Phi'_*\hat{\mu}$.*

Proof: Let $M'' \in \mathfrak{C}(\mu, \hat{\mu})$ be such that $M''(\Phi' \circ \Phi) = 1$. Define a set

$$A := (\Phi \times \hat{\Omega}) \cap (\Omega \times \Phi') \in \mathcal{A}(\Omega \times \bar{\Omega} \times \hat{\Omega})$$

and recall from (B.1) that $\Phi' \circ \Phi = \text{proj}_{\Omega \times \hat{\Omega}}(A)$. In Lemma C.1 put $\nu := M''$ and $\varphi := \text{proj}_{\Omega \times \hat{\Omega}}$; it follows that there exists some measure $\mathbb{P} \in \mathcal{P}(\Omega \times \bar{\Omega} \times \hat{\Omega})$ such that $\mathbb{P}(A) = 1$ and $(\text{proj}_{\Omega \times \hat{\Omega}})_*\mathbb{P} = M''$, so in particular $(\text{proj}_\Omega)_*\mathbb{P} = \mu$ and $(\text{proj}_{\hat{\Omega}})_*\mathbb{P} = \hat{\mu}$. Let us define $\bar{M} := (\text{proj}_{\Omega \times \bar{\Omega}})_*\mathbb{P}$, $M' := (\text{proj}_{\bar{\Omega} \times \hat{\Omega}})_*\mathbb{P}$ and $\bar{\mu} := (\text{proj}_{\bar{\Omega}})_*\mathbb{P}$.

Clearly, we immediately obtain that $M \in \mathfrak{C}(\mu, \bar{\mu})$ and $M' \in \mathfrak{C}(\bar{\mu}, \hat{\mu})$. Furthermore,

$$\begin{aligned} M(\Phi) &= \mathbb{P}(\Phi \times \hat{\Omega}) \geq \mathbb{P}(A) = 1, \\ M'(\Phi) &= \mathbb{P}(\Omega \times \Phi') \geq \mathbb{P}(A) = 1, \end{aligned}$$

so that $\mu\Phi_*\bar{\mu}$ and $\bar{\mu}\Phi'_*\hat{\mu}$ as desired. \square

Theorem B.8 *The endofunctor \mathcal{P} on BoRel is well-defined.*

Proof: We check properties \mathcal{P} has to satisfy in order to be an endofunctor. First of all, for any Borel space Ω we have $\mathcal{P}(\Omega)$ being again a Borel space, hence \mathcal{P} sends objects of BoRel to objects of this category again.

Second, for any morphism in BoRel between objects Ω and $\bar{\Omega}$ given by $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$ it follows from Lemma B.5 that $\Phi_* \in \mathcal{A}(\mathcal{P}(\Omega) \times \mathcal{P}(\bar{\Omega}))$, hence it is again a morphism in BoRel now between objects $\mathcal{P}(\Omega)$ and $\mathcal{P}(\bar{\Omega})$. To show that \mathcal{P} preserves identity notice that $\mu(\Delta_\Omega)_*\mu' \text{ iff } \mu' = (\text{id}_\Omega)_*\mu = \mu$ as in Example B.2 since $\Delta_\Omega = \text{Gr}(\text{id}_\Omega)$. Finally, to show that $(\Phi' \circ \Phi)_* = \Phi'_* \circ \Phi_*$ recall that $(\Phi' \circ \Phi)_* \subseteq \Phi'_* \circ \Phi_*$ by Lemma B.6 whereas the converse inclusion holds by Lemma B.7. \square

Example B.2 shows that \mathcal{P} is an extension of the probabilistic functor developed by Lawvere and Giry [92, 65]. However, it is not clear whether \mathcal{P} defines a monad: at least δ is not a natural transformation anymore in a sense that

$$\text{Gr}(\delta) \circ \Phi = \{(\omega, \delta(\bar{\omega})) : \omega\Phi\bar{\omega}\}$$

in general differs from

$$\Phi_* \circ \text{Gr}(\delta) = \{(\omega, p) : \delta(\omega)\Phi_*p\} = \{(\omega, p) : p(\Phi|_\omega) = 1\}$$

where $\Omega, \bar{\Omega}$ are generic Borel spaces and $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$. In fact it is unlikely that there are other natural candidates for such a transformation which makes the existence of a monad for \mathcal{P} questionable.

The next results describe other useful properties of the lifting of relations.

Lemma B.9 *Consider some Borel spaces $\Omega, \bar{\Omega}$ and $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$, then Φ is an l.t.r. iff so is Φ_* .*

Proof: If Φ_* is an l.t.r., then for any $\omega \in \Omega$ there exists $\bar{\mu} \in \mathcal{P}(\bar{\Omega})$ such that $\delta(\omega)\Phi_*\bar{\mu}$, that is $\bar{\mu}(\Phi|_\omega) = 1$ and hence $\Phi|_\omega \neq \emptyset$. Conversely, if Φ is an l.t.r., Proposition C.17 implies that there exists a map $\varphi \in \mathcal{U}(\Omega, \bar{\Omega})$ such that $\text{Gr}(\varphi) \subseteq \Phi$. For any $\mu \in \mathcal{P}(\Omega)$ we obtain that $\mu\Phi_*\varphi_*\mu$ and hence Φ_* is an l.t.r. \square

Lemma B.10 *Consider some Borel spaces $\Omega, \bar{\Omega}$ and $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$, then $\text{sco}(\Phi_*|_P) = \Phi_*|_{\text{sco } P}$ for each $P \in \mathcal{A}(\mathcal{P}(\Omega))$.*

Proof: Let us denote $\bar{P} := \Phi_*|_P$; note that $\bar{P} = \text{proj}_{\mathcal{P}(\bar{\Omega})}(\Phi_* \cap (P \times \mathcal{P}(\bar{\Omega}))) \in \mathcal{A}(\bar{\Omega})$. We prove the lemma by showing that $\Phi_*|_{\text{sco } P} \subseteq \text{sco}(\bar{P})$ and vice-versa. The former inclusion can be equivalently restated as follows: for any $\mu' \in \text{sco } P$ there exists $\bar{\mu}' \in \text{sco } \bar{P}$ such that $\mu' \Phi_* \bar{\mu}'$. By definition of \bar{P} , there exists a map $M : \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\Omega \times \bar{\Omega})$, which for any $\mu \in P$ satisfies the following three conditions:

1. $M(\Phi|\mu) = 1$;
2. $(\text{proj}_{\bar{\Omega}})_* M(\mu) = \mu$;
3. $(\text{proj}_{\bar{\Omega}})_* M(\mu) \in \bar{P}$;

By Lemma C.20 M can be chosen to be universally measurable. Consider an arbitrary $\mu' \in \text{sco } P$ and let $\nu \in \mathcal{P}(\mathcal{P}(\Omega))$ be any measure satisfying $\mu' = \int_P \mu \nu(d\mu)$. Define $M' := \int_P M(\mu) \nu(d\mu)$. From the definition of M it follows that $\bar{\mu}' := (\text{proj}_{\bar{\Omega}})_* M' \in \text{sco } \bar{P}$, $M' \in \mathfrak{C}(\mu', \bar{\mu}')$ and $M'(\Phi) = 1$.

The converse direction is proven in a similar way. Let $\bar{M} : \mathcal{P}(\bar{\Omega}) \rightarrow \mathcal{P}(\Omega \times \bar{\Omega})$ be such that $M(\bar{\Phi}|\bar{\mu}) = 1$, $(\text{proj}_{\bar{\Omega}})_* \bar{M}(\bar{\mu}) = \bar{\mu}$ and $(\text{proj}_{\Omega})_* \bar{M}(\bar{\mu}) \in P$ for each $\bar{\mu} \in \bar{P}$. As above, we can assume that \bar{M} is universally measurable. Given some $\bar{\mu}' \in \text{sco } \bar{P}$, define $\bar{M}' := \int_{\bar{P}} \bar{M}(\bar{\mu}) \bar{\nu}(d\bar{\mu})$. The latter measure satisfies $\mu' := (\text{proj}_{\Omega})_* \bar{M}' \in \text{sco } P$, $\bar{M}' \in \mathfrak{C}(\mu', \bar{\mu}')$ and $\bar{M}'(\bar{\Phi}) = 1$, which completes the proof. \square

Corollary B.11 Consider some Borel spaces $\Omega, \bar{\Omega}$ and $\varphi \in \mathcal{B}(\Omega, \bar{\Omega})$, then $\text{sco}(\varphi_* P) = \varphi_*(\text{sco } P)$ for each $P \in \mathcal{A}(\mathcal{P}(\Omega))$.

Proof: Apply Lemma B.10 to $\Phi := \text{Gr}(\varphi)$. \square

Lemma B.12 Consider some Borel spaces $\Omega, \bar{\Omega}$ and let $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$. If measures $\mu \in \mathcal{P}(\Omega)$ and $\bar{\mu} \in \mathcal{P}(\bar{\Omega})$ satisfy $\mu \Phi_* \bar{\mu}$, then for any $\mu' \in \mathcal{P}(\Omega)$ there exists $\bar{\mu}' \in \mathcal{P}(\bar{\Omega})$ such that $\|\bar{\mu} - \bar{\mu}'\| \leq \|\mu - \mu'\|$ and $\mu' \Phi_* \bar{\mu}'$.

Proof: Let $M \in \mathfrak{C}(\mu, \bar{\mu})$ satisfy $M(\Phi) = 1$ and denote $m := \frac{dM}{d\mu} \in \mathcal{B}(\bar{\Omega}|\Omega)$. Clearly, $m(\Phi|\omega) = 1$ (μ -a.s.), and since Φ contains a graph of a universally measurable map, there exists $m' \in \mathcal{U}(\bar{\Omega}|\Omega)$ such that $m' = m$ (μ -a.s.) and $m'(\Phi|\omega) = 1$ for all $\omega \in \Omega$. Define $M' := \mu' \otimes m'$ and $\bar{\mu}' := \mu' m'$, then $M' \in \mathfrak{C}(\mu', \bar{\mu}')$ and $M'(\Phi) = 1$. Finally,

$$\begin{aligned} \|\bar{\mu} - \bar{\mu}'\| &= 2 \sup_{\bar{A} \in \bar{\Omega}} \left| \int_{\Omega} m'(A|\omega) \mu(d\omega) - \int_{\Omega} m'(A|\omega) \mu'(d\omega) \right| \\ &= 2 \sup_{\bar{A} \in \bar{\Omega}} \left| \int_{\Omega} \left(m'(A|\omega) - \frac{1}{2} \right) \mu(d\omega) - \int_{\Omega} \left(m'(A|\omega) - \frac{1}{2} \right) \mu'(d\omega) \right| \leq \|\mu - \mu'\| \end{aligned}$$

since $m'(A|\cdot) - \frac{1}{2} \in \text{b}\mathcal{U}_{\frac{1}{2}}(\Omega)$ for each $A \in \mathcal{B}(\bar{\Omega})$. \square

Other topics

C.1 Important fragments of LTL

Although any LTL formula can be expressed as a DRA, such generality is not very useful in practice. Even when dealing with finite cdt-MP \mathcal{D} , expressing a given formula as a DFA $A = (\text{TS}, D)$ (if possible) may reduce the complexity of the automaton comparing to some DRA expressions of the formula, as well as allows applying simpler solution methods, which altogether leads to a smaller state space of the composition $\mathcal{D} \sqcup \text{TS}$ and hence to a lower computational time. In the case when the cdt-MP \mathcal{D} is not finite, in addition the solution methods are much more involved and as Sections 4.2 and 4.3 suggest, solution of a bounded-horizon reachability problem simpler than the one of an unbounded horizon reachability, which in turn is easier than the repeated reachability problem. As a result, e.g. although any LTL formula that encodes some bounded-horizon property can be expressed as a DRA, it is worth analyzing the formula to check whether it allows for an automaton expression with a simpler acceptance condition. In this section we describe how to perform such analysis, and what are the useful fragments of LTL that allow for an expression via an automaton that is simpler than a DRA.

The syntactically safe LTL (sLTL) [87] expresses safety languages. A language $\phi \subseteq Y^{\mathbb{N}}$ is called a safety property iff any word $w \notin \phi$ has a finite “bad” prefix:

$$w \notin \phi \iff \exists n \in \mathbb{N} : \text{proj}_{Y^n}^{-1}(\text{proj}_{Y^n}(w)) \cap \phi = \emptyset.$$

The syntactically co-safe LTL (scLTL) [87] expresses co-safety languages, where a co-safety language ϕ is the one for which any word $w \in \phi$ has a good prefix, that is

$$w \in \phi \iff \exists n \in \mathbb{N} : \text{proj}_{Y^n}^{-1}(\text{proj}_{Y^n}(w)) \subseteq \phi.$$

Clearly ϕ is a safety language if and only if $Y^{\mathbb{N}} \setminus \phi$ is a co-safety one. This comes as no surprise as safety languages are exactly closed subsets of $Y^{\mathbb{N}}$ in the product topology, whereas co-safety languages are open [11]. It follows that any co-safety language can be expressed as a DFA, and hence DFA can be used for negations of

safety languages. Here we only give a grammar of sLTL¹. For this purpose, in the LTL setting let us define a temporal modality *Weak until* W^∞ by

$$\Phi_1 W^\infty \Phi_2 := \Phi_1 \cup \Phi_2 \vee \Box \Phi_1.$$

The grammar of sLTL is given as follows:

$$\Phi ::= \sigma \in \Sigma \mid \neg \sigma \mid \Phi_1 \wedge \Phi_2 \mid \Phi_1 \vee \Phi_2 \mid X\Phi \mid \Phi_1 W^\infty \Phi_2.$$

Note that in sLTL the negation can be only applied on the level of letters, so that \vee could not be expressed through \wedge in general in sLTL in contrast to the LTL setting. Moreover, in general it is not possible to express $\Phi_1 \cup \Phi_2$ using sLTL grammar. An example of an sLTL formula is $\Box^n \sigma$, and that of an csLTL formula are $\Diamond^n \sigma$ and $\sigma_1 \cup^n \sigma_2$ where $n \in \mathbb{N}^\infty$ in all three cases. One immediate way to see whether a given LTL formula belongs to sLTL is to write it in a negation normal form (NNF), where the negation is presented on the level of atomic propositions by means of the following identities: $\neg X\Phi = X(\neg\Phi)$, $\neg(\Phi_1 \cup \Phi_2) = \neg\Phi_1 W^\infty \neg\Phi_2$ etc. However, even a LTL formula corresponding to a safety language may lead to a NNF which does not belong to sLTL, so for more elaborate methods see [87]. Recent examples of applications of sLTL and of csLTL can be found in papers [113] and [15] respectively.

Although sLTL and scLTL are related to the expression of formulae via DFA rather than DRA, they still lead to the unbounded-horizon reachability problem over $\mathcal{D} \sqcup \text{TS}$, even in case when the original formula encodes a bounded-horizon specification. A useful framework to deal with the latter is given by the bounded LTL (BLTL) [133] which expresses bounded languages: a language $\phi \subseteq Y^\mathbb{N}$ is called bounded if there exists $n \in \mathbb{N}$ such that

$$w \in \phi \iff \text{proj}_{Y^n}^{-1}(\text{proj}_{Y^n}(w)) \subseteq \phi.$$

In particular, it appears that bounded languages are exactly those that are both safety and co-safety languages [110, Proposition 3.10, Chapter III], that is they are clopen subsets of $Y^\mathbb{N}$. The grammar of BLTL is given as follows:

$$\Phi ::= \sigma \in \Sigma \mid \neg \Phi \mid \Phi_1 \wedge \Phi_2 \mid X\Phi \tag{C.1}$$

so that it still allows for negations to be applied on all the levels, but \cup is not absent. On the other hand, (2.6) implies that $\Phi_1 \cup^n \Phi_2$ belongs to BLTL for finite $n \in \mathbb{N}$. It is likely that any BLTL formula allows to be expressed as a bounded-horizon version of the DFA [133, Section 3.4] which accepts only those runs that visit the set of final states in at most n steps, where n is specified a priori, in the definition of the automaton. For the applications of BLTL see e.g. [74].

¹ The grammar of scLTL can be easily deduced from the one of sLTL; see also [15, Definition 2.1].

C.2 Stochastic equations

Throughout the thesis sometimes we face the following problem: let Ω and Ξ be two Borel spaces, $\varphi \in \mathcal{B}(\Omega, \Xi)$ and $\nu \in \mathcal{P}(\Xi)$. Does there exist a probability measure $\mu \in \mathcal{P}(\Omega)$ such that $\nu = \varphi_*\mu$?

The expression $\nu = \varphi_*\mu$ with given φ, ν and unknown μ is called a *stochastic equation* [51], and can be considered as an abstract form of Itô stochastic differential equations [106]. Also, a stochastic equation can be seen as an extension of a measure from a smaller sub- σ -algebra to a bigger one: clearly, given φ and ν we can define a measure ν' on a sub- σ -algebra $\varphi^{-1}(\mathcal{B}(\Xi)) \subseteq \mathcal{B}(\Omega)$. The question is whether it is possible to extend ν' to the whole Borel σ -algebra $\mathcal{B}(\Xi)$: such an extension μ would be a desired solution of a stochastic equation $\nu = \varphi_*\mu$.

The aforementioned paper [51] provides several useful conditions for the existence of solutions of stochastic equations, however the case we deal with is slightly different. Due to this reason, we provide a proof for a specific situation of our interest, which is based on methods used in [51]: construction of a measurable inverse of φ .

Lemma C.1 *Consider some Borel spaces Ω, Ξ , and $\varphi \in \mathcal{B}(\Omega, \Xi)$. For each $A \in \mathcal{A}(\Omega)$ and probability measure $\nu \in \mathcal{P}(\Xi)$ satisfying $\nu(\varphi(A)) = 1$ there exists a solution of the stochastic equation $\mu \in \mathcal{P}(\Omega)$ such that $\mu(A) = 1$ and such that $\nu = \varphi_*\mu$.*

Proof: Since φ is a Borel map, $\text{Gr}(\varphi) \in \mathcal{B}(\Omega \times \Xi)$ and $\text{Gr}(\varphi)^{-1} \in \mathcal{B}(\Xi \times \Omega)$. Proposition C.17 implies the existence of $\psi \in \mathcal{U}(\Xi, \Omega)$ satisfying $\varphi(\psi(\xi)) = \xi$ for each $\xi \in \varphi(A)$. We extend the map ψ to Ξ by defining $\psi(\xi) = \omega' \in \Omega$, an arbitrary auxiliary point. By [33, Lemma 1.2] there exists $\psi' \in \mathcal{B}(\Xi, \Omega)$ such that $\psi = \psi'$ (ν -a.s.).

Let us define $\mu := \psi'_*\nu$ and show that it satisfies the desired properties. First of all,

$$\mu(A) = \nu(\xi \in \Xi : \psi'(\xi) \in A) = \nu(\xi \in \Xi : \psi(\xi) \in A) \geq \nu(\varphi(A)) = 1.$$

Second, for any $B \in \mathcal{B}(\Xi)$ we obtain that

$$\varphi_*\mu(B) = (\varphi \circ \psi)_*\nu(B) = \nu(\xi : \varphi(\psi(\xi)) \in B) \geq \nu(B \cap \varphi(A)) = \nu(B).$$

which completes the proof. \square

C.3 Auxiliary results

Lemma C.2 *Consider some Borel space Ω , and for any class of sets $\mathcal{F} \subseteq 2^\Omega$ denote*

$$d_{\text{TV}, \mathcal{F}}(\mu, \mu') = 2 \cdot \sup_{F \in \mathcal{F}} |\mu(F) - \mu'(F)|, \quad \mu, \mu' \in \mathcal{P}(\Omega).$$

If \mathcal{F} is an algebra that generates $\mathcal{B}(\Omega)$, then $d_{\text{TV}, \mathcal{F}} = d_{\text{TV}}$.

Proof: If \mathcal{F} is an algebra that generates $\mathcal{B}(\Omega)$ then $\mathcal{F} \subseteq \mathcal{B}(\Omega)$, so clearly $d_{\text{TV}, \mathcal{F}} \leq d_{\text{TV}}$, hence only the converse inequality needs to be proved. It is easy to show that for any $\nu \in \mathcal{P}(\Omega)$, $\varepsilon > 0$ and $A \in \mathcal{B}(\Omega)$ there exists $F \in \mathcal{F}$ such that $\nu(A \Delta F) \leq \varepsilon$. In particular, if $\nu = \frac{1}{2}(\mu + \mu')$ for some $\mu, \mu' \in \mathcal{P}(\Omega)$, then $\mu(A \Delta F) \leq 2\varepsilon$ and hence $|\mu(A) - \mu(F)| \leq 2\varepsilon$, and similarly for μ' . Thus

$$|\mu(A) - \mu'(A)| \leq |\mu(A) - \mu(F)| + |\mu(F) - \mu'(F)| + |\mu'(F) - \mu'(A)| \leq d_{\text{TV}, \mathcal{F}}(\mu, \mu') + 4\varepsilon.$$

Since μ, μ', A and ε are arbitrary, we obtain the desired result. \square

Corollary C.3 $d_{\text{TV}} \in \mathcal{B}(\mathcal{P}(\Omega) \times \mathcal{P}(\Omega))$ for each Borel space Ω .

Proof: Let \mathcal{F} be an algebra that generates $\mathcal{B}(\Omega)$. Since the latter is countably generated, we can assume \mathcal{F} to be countable. Note that $d_{\text{TV}, \mathcal{F}}(\mu, \mu') = 2 \cdot \sup_{F \in \mathcal{F}} |\mathfrak{c}_F(\mu) - \mathfrak{c}_F(\mu')|$ and hence is a Borel function as a countable supremum of Borel functions. Since $d_{\text{TV}} = d_{\text{TV}, \mathcal{F}}$ by Lemma C.13, the desired result follows immediately. \square

Lemma C.4 Let $\Omega, \hat{\Omega}$ be arbitrary Borel spaces and let $f : \Omega \rightarrow \hat{\Omega}$ be a Borel map. Then $f_* : \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\hat{\Omega})$ is also a Borel map.

Proof: For the proof see e.g. [47, Proposition 1.27]. \square

Lemma C.5 Let Ω be an arbitrary set and consider any two functions $f, g : \Omega \rightarrow \mathbb{R}$. If $|f(\omega) - g(\omega)| \leq \varepsilon$ for all $\omega \in \Omega$ then $|\sup_{\omega \in \Omega} f(\omega) - \sup_{\omega \in \Omega} g(\omega)| \leq \varepsilon$.

Proof: For the proof see e.g. [69, Appendix A.3]. \square

Proposition C.6 Consider some Borel spaces Ω, Ξ and $\Phi \in \mathcal{A}(\Omega \times \Xi)$. There exists a map $\varphi \in \mathcal{U}(\Omega, \Xi)$ such that $\varphi(\omega) \in \Phi|_{\omega}$ for all $\omega \in \text{proj}_{\Omega}(\Phi)$.

Proof: For the proof see e.g. [20, Proposition 7.49]. \square

Lemma C.7 Coupling can be pushed.

Lemma C.8 Consider some Borel spaces Ω, Ξ , and let $\mu, \bar{\mu} \in \mathcal{P}(\Omega)$ and $\kappa, \bar{\kappa} \in \mathcal{U}(\Xi|\Omega)$. If there exists a set $E \in \mathcal{B}(\Omega)$ such that $\bar{\mu}(E) = 1$ and $\|\kappa(\omega) - \bar{\kappa}(\omega)\| \leq \varepsilon$ for all $\omega \in E$, then $\|\mu \otimes \kappa - \bar{\mu} \otimes \bar{\kappa}\| \leq \|\mu - \bar{\mu}\| \otimes \varepsilon$.

Proof: Define $\nu := \frac{1}{2}(\mu + \bar{\mu})$, so that $A \in \mathcal{B}(\Omega)$ satisfies $\nu(A) = 1$ iff $\mu(A) = \bar{\mu}(A) = 1$. Then there exists $E' \in \mathcal{B}(\Omega)$ and $\kappa', \bar{\kappa}' \in \mathcal{B}(\Xi|\Omega)$ such that $\nu(E') = 1$ and such that $\kappa(\omega) = \kappa'(\omega)$ and $\bar{\kappa}(\omega) = \bar{\kappa}'(\omega)$ for each $\omega \in E'$. Define $E'' = E \cap E'$, then

$\nu(E'') = 1$ and $\|\kappa'(\omega) - \bar{\kappa}'(\omega)\| \leq \varepsilon$ for all $\omega \in E''$. Define $\kappa'' := 1_{E''}\kappa' + 1_{(E'')^c}\bar{\kappa}'$, so $\|\kappa'(\omega) - \bar{\kappa}'(\omega)\| \leq \varepsilon$ for all $\omega \in \Omega$ and by [10, Lemma 2] it holds that $\|\mu \otimes \kappa'' - \bar{\mu} \otimes \bar{\kappa}'\| \leq \|\mu - \bar{\mu}\| \otimes \varepsilon$. The result follows from the fact that $\mu \otimes \kappa = \mu \otimes \kappa''$ and $\bar{\mu} \otimes \bar{\kappa} = \bar{\mu} \otimes \bar{\kappa}'$. \square

Lemma C.9 *Let Y, Y' be arbitrary sets and let $g : Y \rightarrow \mathbb{R}$ and $g' : Y' \rightarrow \mathbb{R}$ be some functions. Suppose that there exist maps $a : Y \rightarrow Y'$ and $a' : Y' \rightarrow Y$ such that*

$$g(y) = g'(a(y)), \quad g'(y') = g(a'(y')), \quad \forall y \in Y, y' \in Y'$$

Then: $\inf_{y \in Y} g(y) = \inf_{y' \in Y'} g'(y')$ and $\sup_{y \in Y} g(y) = \sup_{y' \in Y'} g'(y')$.

Proof: The following sequences of inequalities

$$\begin{aligned} \inf_{y \in Y} g(y) &= \inf_{y \in Y} g'(a(y)) \geq \inf_{y' \in Y'} g'(y') = \inf_{y' \in Y'} g(a'(y')) \geq \inf_{y \in Y} g(y) \\ \sup_{y \in Y} g(y) &= \sup_{y \in Y} g'(a(y)) \leq \sup_{y' \in Y'} g'(y') = \sup_{y' \in Y'} g(a'(y')) \leq \sup_{y \in Y} g(y) \end{aligned}$$

yield the desired result. \square

The next lemma shows that point-wise bounds also hold for the optimal values.

Lemma C.10 *Let Y be an arbitrary set and consider any two function $f, g : Y \rightarrow \mathbb{R}$. If $|f(y) - g(y)| \leq \varepsilon$ for all $y \in Y$ then $|\sup_{y \in Y} f(y) - \sup_{y \in Y} g(y)| \leq \varepsilon$.*

Proof: The proof is given in [69, Appendix A.3]. \square

Lemma C.11 *If Y is a Borel space, the set S is closed in Y and the function $f \in \text{b}\mathcal{C}^*(X)$ is such that $f \geq 0$, then it holds that $1_S \cdot f \in \text{b}\mathcal{C}^*(X)$.*

Proof: Notice that for any $c \leq 0$ it holds that $\{1_S \cdot f \geq c\} = X$, whereas for $c > 0$ we obtain $\{1_S \cdot f \geq c\} = S \cap \{f \geq c\}$ which is a closed set as well. \square

Lemma C.12 *Let A be any set, and let (\mathcal{F}, ρ) be a metric space where \mathcal{F} is any class of bounded functions $f : A \rightarrow \mathbb{R}$ and $\rho : (f', f'') \mapsto \sup_{a \in A} |f''(a) - f'(a)|$. Consider an arbitrary operator $\mathfrak{G} : \mathcal{F} \rightarrow \mathcal{F}$ that satisfies the following two properties:*

1. *if $f, g \in \mathcal{F}$ such that $f \leq g$, then $\mathfrak{G}f \leq \mathfrak{G}g$,*
2. *there exists $\beta \in [0, 1)$ such that if $f \in \mathcal{F}$ and $c \geq 0$ then $\mathfrak{G}(f + c) \leq \mathfrak{G}f + \beta c$.*

Then \mathfrak{G} is a contraction on \mathcal{F} with a modulus β .

Proof: Let $f, g \in \mathcal{F}$ be arbitrary, then $f \leq g + \rho(f, g)$ and thus

$$\mathfrak{G}f \leq \mathfrak{G}g + \beta\rho(f, g) \implies \mathfrak{G}f - \mathfrak{G}g \leq \beta\rho(f, g).$$

By a symmetric argument, we obtain that

$$\mathfrak{G}f - \mathfrak{G}g \leq \beta\rho(f, g) \implies |\mathfrak{G}f - \mathfrak{G}g| \leq \beta\rho(f, g) \implies \rho(\mathfrak{G}f, \mathfrak{G}g) \leq \beta\rho(f, g),$$

so that \mathfrak{G} is a contraction with a modulus β . \square

Lemma C.13 Consider some Borel space Ω , and for any class of sets $\mathcal{F} \subseteq 2^\Omega$ denote

$$d_{\text{TV}, \mathcal{F}}(\mu, \mu') = 2 \cdot \sup_{F \in \mathcal{F}} |\mu(F) - \mu'(F)|, \quad \mu, \mu' \in \mathcal{P}(\Omega).$$

If \mathcal{F} is an algebra that generates $\mathcal{B}(\Omega)$, then $d_{\text{TV}, \mathcal{F}} = d_{\text{TV}}$.

Proof: If \mathcal{F} is an algebra that generates $\mathcal{B}(\Omega)$ then $\mathcal{F} \subseteq \mathcal{B}(\Omega)$, so clearly $d_{\text{TV}, \mathcal{F}} \leq d_{\text{TV}}$, hence only the converse inequality needs to be proved. It is easy to show that for any $\nu \in \mathcal{P}(\Omega)$, $\varepsilon > 0$ and $A \in \mathcal{B}(\Omega)$ there exists $F \in \mathcal{F}$ such that $\nu(A \triangle F) \leq \varepsilon$. In particular, if $\nu = \frac{1}{2}(\mu + \mu')$ for some $\mu, \mu' \in \mathcal{P}(\Omega)$, then $\mu(A \triangle F) \leq 2\varepsilon$ and hence $|\mu(A) - \mu(F)| \leq 2\varepsilon$, and similarly for μ' . Thus

$$|\mu(A) - \mu'(A)| \leq |\mu(A) - \mu(F)| + |\mu(F) - \mu'(F)| + |\mu'(F) - \mu'(A)| \leq d_{\text{TV}, \mathcal{F}}(\mu, \mu') + 4\varepsilon.$$

Since μ, μ', A and ε are arbitrary, we obtain the desired result. \square

Corollary C.14 $d_{\text{TV}} \in \mathcal{B}(\mathcal{P}(\Omega) \times \mathcal{P}(\Omega))$ for each Borel space Ω .

Proof: Let \mathcal{F} be an algebra that generates $\mathcal{B}(\Omega)$. Since the latter is countably generated, we can assume \mathcal{F} to be countable. Note that $d_{\text{TV}, \mathcal{F}}(\mu, \mu') = 2 \cdot \sup_{F \in \mathcal{F}} |\mathfrak{e}_F(\mu) - \mathfrak{e}_F(\mu')|$ and hence is a Borel function as a countable supremum of Borel functions. Since $d_{\text{TV}} = d_{\text{TV}, \mathcal{F}}$ by Lemma C.13, the desired result follows immediately. \square

Lemma C.15 Let $\Omega, \hat{\Omega}$ be arbitrary Borel spaces and let $f : \Omega \rightarrow \hat{\Omega}$ be a Borel map. Then $f_* : \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\hat{\Omega})$ is also a Borel map.

Proof: For the proof see e.g. [47, Proposition 1.27]. \square

Lemma C.16 Let Ω be an arbitrary set and consider any two functions $f, g : \Omega \rightarrow \mathbb{R}$. If $|f(\omega) - g(\omega)| \leq \varepsilon$ for all $\omega \in \Omega$ then $|\sup_{\omega \in \Omega} f(\omega) - \sup_{\omega \in \Omega} g(\omega)| \leq \varepsilon$.

Proof: For the proof see e.g. [69, Appendix A.3]. \square

Proposition C.17 Consider some Borel spaces Ω, Ξ and $\Phi \in \mathcal{A}(\Omega \times \Xi)$. There exists a map $\varphi \in \mathcal{U}(\Omega, \Xi)$ such that $\varphi(\omega) \in \Phi|_{\omega}$ for all $\omega \in \text{proj}_{\Omega}(\Phi)$.

Proof: For the proof see e.g. [20, Proposition 7.49]. \square

Lemma C.18 Coupling can be pushed.

Lemma C.19 Consider some Borel spaces Ω, Ξ , and let $\mu, \bar{\mu} \in \mathcal{P}(\Omega)$ and $\kappa, \bar{\kappa} \in \mathcal{U}(\Xi|\Omega)$. If there exists a set $E \in \mathcal{B}(\Omega)$ such that $\bar{\mu}(E) = 1$ and $\|\kappa(\omega) - \bar{\kappa}(\omega)\| \leq \varepsilon$ for all $\omega \in E$, then $\|\mu \otimes \kappa - \bar{\mu} \otimes \bar{\kappa}\| \leq \|\mu - \bar{\mu}\| \otimes \varepsilon$.

Proof: Define $\nu := \frac{1}{2}(\mu + \bar{\mu})$, so that $A \in \mathcal{B}(\Omega)$ satisfies $\nu(A) = 1$ iff $\mu(A) = \bar{\mu}(A) = 1$. Then there exists $E' \in \mathcal{B}(\Omega)$ and $\kappa', \bar{\kappa}' \in \mathcal{B}(\Xi|\Omega)$ such that $\nu(E') = 1$ and such that $\kappa(\omega) = \kappa'(\omega)$ and $\bar{\kappa}(\omega) = \bar{\kappa}'(\omega)$ for each $\omega \in E'$. Define $E'' = E \cap E'$, then $\nu(E'') = 1$ and $\|\kappa'(\omega) - \bar{\kappa}'(\omega)\| \leq \varepsilon$ for all $\omega \in E''$. Define $\kappa'' := 1_{E''}\kappa' + 1_{(E'')^c}\bar{\kappa}'$, so $\|\kappa'(\omega) - \bar{\kappa}'(\omega)\| \leq \varepsilon$ for all $\omega \in \Omega$ and by [10, Lemma 2] it holds that $\|\mu \otimes \kappa'' - \bar{\mu} \otimes \bar{\kappa}'\| \leq \|\mu - \bar{\mu}\| \otimes \varepsilon$. The result follows from the fact that $\mu \otimes \kappa = \mu \otimes \kappa''$ and $\bar{\mu} \otimes \bar{\kappa} = \bar{\mu} \otimes \bar{\kappa}'$. \square

Lemma C.20 Consider some Borel spaces $\Omega, \bar{\Omega}$ and $\Xi, \bar{\Xi}$, and let $\Upsilon \in \mathcal{A}(\Omega \times \mathcal{P}(\Xi))$, $\bar{\Upsilon} \in \mathcal{A}(\bar{\Omega} \times \mathcal{P}(\bar{\Xi}))$ be l.t.r. and $\Phi \in \mathcal{A}(\Omega \times \bar{\Omega})$, $\Psi \in \mathcal{A}(\Xi \times \bar{\Xi})$. Suppose that for any $(\omega, \bar{\omega}) \in \Phi$ there exists $G(\omega, \bar{\omega}) \in \mathcal{P}(\Xi \times \bar{\Xi})$ satisfying

1. $G(\Psi|_{\omega}, \bar{\omega}) = 1$;
2. $(\text{proj}_{\Xi})_* G(\omega, \bar{\omega}) \in \Upsilon|_{\omega}$;
3. $(\text{proj}_{\bar{\Xi}})_* G(\omega, \bar{\omega}) \in \bar{\Upsilon}|_{\bar{\omega}}$.

Then there exists a kernel in $\mathcal{U}(\Xi \times \bar{\Xi}|\Omega \times \bar{\Omega})$ satisfying conditions (1) – (3).

Proof: Clearly, G is a particular choice in $\Phi \times \mathcal{P}(\Xi \times \bar{\Xi})$ that satisfies the three conditions, so we need to show that such choice could be done in a universally measurable way. For this purpose, we are going to show that the set $J \subseteq \Phi \times \mathcal{P}(\Xi \times \bar{\Xi})$ to do the choice over is analytic, so that Proposition C.17 applies.

It follows that $J = \bigcap_{i=1}^3 J_i$, where J_i corresponds to the i -th condition. In particular,

$$J_1 = \{(\omega, \bar{\omega}, p) : (\omega, \bar{\omega}) \in \Phi, p(\Psi) = 1\} = \Phi \times (\mathfrak{e}_{\Psi})^{-1}(1) \in \mathcal{A}(\Phi \times \mathcal{P}(\Xi \times \bar{\Xi}))$$

as a product of two analytic sets. Furthermore, for J_2 we have

$$\begin{aligned} J_2 &= \{(\omega, \bar{\omega}, p) : (\omega, \bar{\omega}) \in \Phi, (\omega, (\text{proj}_{\Xi})_* p) \in \Upsilon\} \\ &= (\Phi \times \mathcal{P}(\Xi \times \bar{\Xi})) \cap \left(\bar{\Omega} \times (\text{id}_{\Omega} \times (\text{proj}_{\Xi})_*)^{-1}(\Upsilon) \right). \end{aligned}$$

As proj_{Ξ} is Borel so is $(\text{proj}_{\Xi})_* : \mathcal{P}(\Xi \times \bar{\Xi}) \rightarrow \mathcal{P}(\Xi)$ by Lemma C.15. Since the product of two Borel maps is Borel, by [20, Proposition 7.40] $(\text{id}_{\Omega} \times (\text{proj}_{\Xi})_*)^{-1}(\Upsilon)$ is an analytic set, hence so is J_2 . A similar argument implies that J_3 is an analytic set as well, so that J is analytic as an intersection of analytic sets. \square

Lemma C.21 *Consider some Borel spaces $\Omega, \bar{\Omega}$ and $\Xi, \bar{\Xi}, \Phi \in \mathcal{B}(\Omega \times \bar{\Omega})$ and $\Psi \in \mathcal{B}(\Xi \times \bar{\Xi})$, $\mu \in \mathcal{P}(\Omega)$ and $\bar{\mu} \in \mathcal{P}(\bar{\Omega})$, a stochastic kernel $\kappa \in \mathcal{U}(\Xi|\Omega)$ and an l.t.r. $\Upsilon \in \mathcal{A}(\bar{\Omega} \times \mathcal{P}(\bar{\Xi}))$. Suppose that $\mu\Phi_*\bar{\mu}$ and that for any $(\omega, \bar{\omega}) \in \Phi$ there exists a probability measure $G(\omega, \bar{\omega}) \in \mathcal{P}(\Xi \times \bar{\Xi})$ satisfying*

1. $G(\Psi|\omega, \bar{\omega}) = 1$;
2. $(\text{proj}_{\Omega})_*G(\omega, \bar{\omega}) = \kappa(\omega)$;
3. $(\text{proj}_{\bar{\Omega}})_*G(\omega, \bar{\omega}) \in \tilde{\Upsilon}|_{\bar{\omega}}$.

Then there exists a stochastic kernel $\bar{\kappa} \in \mathcal{U}(\bar{\Xi}|\bar{\Omega})$ that satisfies $(\mu \otimes \kappa)\mathcal{P}(\Phi \times \Psi)(\bar{\mu} \otimes \bar{\kappa})$ and such that $\bar{\kappa}(\bar{\omega}) \in \text{sco } \tilde{\Upsilon}|_{\bar{\omega}}$ for any $\bar{\omega} \in \bar{\Omega}$.

Proof: Since $\mu\Phi_*\bar{\mu}$, there exists $M \in \mathcal{C}(\mu, \bar{\mu})$ such that $M(\Phi) = 1$. Our main goal is to use G to construct a kernel $K \in \mathcal{U}(\Xi \times \bar{\Xi}|\Omega \times \bar{\Omega})$ such that $(M \otimes K)(\Phi \times \Psi) = 1$ and the left marginal of $M \otimes K$ is $\mu \otimes \kappa$. We then define $\bar{\kappa}$ using K and show that it satisfies desired properties.

Consider an arbitrary $\kappa' \in \mathcal{B}(\Xi|\Omega)$ such that $\kappa = \kappa'$ (μ -a.s.). By taking $\Upsilon := \text{Gr}(\kappa)$ we obtain $\Upsilon \in \mathcal{B}(\bar{\Omega} \times \mathcal{P}(\bar{\Xi}))$, so that Lemma C.20 implies the existence of $K \in \mathcal{U}(\Xi \times \bar{\Xi}|\Omega \times \bar{\Omega})$ satisfying conditions (1) and (3) from the statement and such that

$$(\text{proj}_{\Omega})_*K(\omega, \bar{\omega}) = \kappa'(\omega)$$

for all $(\omega, \bar{\omega}) \in \Phi$. Denote $\bar{m} := \frac{dM}{d\bar{\mu}} \in \mathcal{B}(\Omega|\bar{\Omega})$ and define

$$\bar{\kappa}(\bar{\omega}) := \int_{\Omega} (\text{proj}_{\Omega})_*K(\omega, \bar{\omega})\bar{m}(d\omega|\bar{\omega}) \quad \bar{\omega} \in \bar{\Omega}. \quad (\text{C.2})$$

By [20, Proposition 7.46] $\bar{\kappa}(\bar{B}|\cdot) \in \text{b}\mathcal{U}(\bar{\Omega})$ for any $\bar{B} \in \mathcal{B}(\bar{\Omega})$, and thus $\bar{\kappa} \in \mathcal{U}(\bar{\Xi}|\bar{\Omega})$ by [20, Lemma 7.28(b \Rightarrow a)]. Note that

$$1 = M(\Phi) = \int_{\Omega} \bar{m}(\Phi|_{\bar{\omega}}|\bar{\omega})\bar{\mu}(d\bar{\omega})$$

and hence there exists some set $\bar{E} \in \mathcal{B}(\bar{\Omega})$ such that $\bar{\mu}(\bar{E}) = 1$ and $\bar{m}(\Phi|_{\bar{\omega}}|\bar{\omega}) = 1$, so that $\bar{\kappa}(\bar{\omega}) \in \text{sco } \tilde{\Upsilon}|_{\bar{\omega}}$ for all $\bar{\omega} \in \bar{E}$. Let $\bar{v} \in \mathcal{U}(\bar{\Xi}|\bar{\Omega})$ be any kernel satisfying $\text{Gr}(\bar{v}) \subseteq \tilde{\Upsilon}$ and redefine $\bar{\kappa}(\bar{\omega}) := \bar{v}(\bar{\omega})$ for all $\bar{\omega} \in \bar{E}^c$, so that $\bar{\kappa}(\bar{\omega}) \in \text{sco } \tilde{\Upsilon}|_{\bar{\omega}}$ for all $\bar{\omega} \in \bar{\Omega}$.

Now, let us show that $(\mu \otimes \kappa)\mathcal{P}(\Phi \times \Psi)(\bar{\mu} \otimes \bar{\kappa})$; we are going to use $M \otimes K$ as a desired coupling. Let us check properties of $M \otimes K$: first of all

$$(M \otimes K)(\Phi \times \Psi) = \int_{\Phi} K(\Psi|\omega, \bar{\omega})M(d\omega \times d\bar{\omega}) = 1$$

by condition (1). To show that $M \otimes K$ has the desired marginals, it is sufficient to focus only on measurable rectangles. For any $A \in \mathcal{B}(\Omega)$ and $B \in \mathcal{B}(\Xi)$:

$$\begin{aligned} (M \otimes K)(A \times \bar{\Omega} \times B \times \bar{\Xi}) &= \int_{A \times \bar{\Omega}} K(B \times \bar{\Xi} | \omega, \bar{\omega}) M(d\omega \times d\bar{\omega}) \\ &= \int_{A \times \bar{\Omega}} \kappa'(B | \omega) M(d\omega \times d\bar{\omega}) \\ &= \int_A \kappa'(\omega, B) \mu(d\omega) = (\kappa \otimes \mu)(A \times B) \end{aligned}$$

where the second equality is satisfied thanks to condition (2). Finally, for any $\bar{A} \in \mathcal{B}(\bar{\Omega})$ and $\bar{B} \in \mathcal{B}(\bar{\Xi})$ it holds that

$$\begin{aligned} (M \otimes K)(\Omega \times \bar{A} \times \Xi \times \bar{B}) &= \int_{\Omega \times \bar{A}} K(\Xi \times \bar{B} | \omega, \bar{\omega}) M(d\omega \times d\bar{\omega}) \\ &= \int_{\Omega \times \bar{A}} (\text{proj}_{\bar{\Omega}})_* K(\bar{B} | \omega, \bar{\omega}) M(d\omega \times d\bar{\omega}) \\ &= \int_{\bar{A}} \bar{\kappa}(\bar{B} | \bar{\omega}) \bar{\mu}(d\bar{\omega}) = (\bar{\mu} \otimes \bar{\kappa})(\bar{A} \times \bar{B}). \end{aligned}$$

□

Bibliography

- [1] A. Abate. Probabilistic bisimulations of switching and resetting diffusions. In *49th IEEE Conference on Decision and Control (CDC)*, pages 5918–5923, Dec 2010.
- [2] A. Abate. Approximation metrics based on probabilistic bisimulations for general state-space markov processes: a survey. *Electronic Notes in Theoretical Computer Sciences*, 2013. In Print.
- [3] A. Abate, S. Amin, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safe sets computation for discrete time stochastic hybrid systems. In *Proceedings of the 45th IEEE Conference of Decision and Control*, pages 258–263, 2006.
- [4] A. Abate, S. Amin, M. Prandini, J. Lygeros, and S. Sastry. Computational approaches to reachability analysis of stochastic hybrid systems. In *Hybrid Systems: Computation and Control*, pages 4–17. Springer Verlag, 2007.
- [5] A. Abate, A. D’Innocenzo, M. Di Benedetto, and S. Sastry. *Markov Set-Chains as Abstractions of Stochastic Hybrid Systems*, pages 1–15. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [6] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini. Approximate model checking of stochastic hybrid systems. *European Journal of Control*, 16:624–641, 2010.
- [7] A. Abate, J.-P. Katoen, and A. Mereacre. Quantitative automata model checking of autonomous stochastic hybrid systems. In *Proceedings of the 14th international conference on Hybrid Systems: Computation and Control, HSCC ’11*, pages 83–92, New York, NY, USA, 2011. ACM.
- [8] A. Abate and M. Prandini. Approximate abstractions of stochastic systems: A randomized method. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pages 4861–4866, 2011.
- [9] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008.

- [10] A. Abate, F. Redig, and I. Tkachev. On the effect of perturbation of conditional probabilities in total variation. *Statistics & Probability Letters*, 2014. In Press.
- [11] B. Alpern and F. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181 – 185, 1985.
- [12] A. Arapostathis, V. Borkar, E. Fernández-Gaucherand, M. Ghosh, and S. Marcus. Discrete-time controlled Markov processes with average cost criterion: a survey. *SIAM J. Control Optim.*, 31(2):282–344, 1993.
- [13] S. Asmussen. *Ruin probabilities*, volume 2 of *Advanced Series on Statistical Science & Applied Probability*. World Scientific Publishing Co. Inc., River Edge, NJ, 2000.
- [14] J.-P. Aubin and A. Cellina. *Differential inclusions: set-valued maps and viability theory*, volume 264. Springer Science & Business Media, 2012.
- [15] E. Aydin Gol, M. Lazar, and C. Belta. Language-guided controller synthesis for discrete-time linear systems. In *Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control*, HSCC '12, pages 95–104, New York, NY, USA, 2012. ACM.
- [16] A. Aziz, V. Singhal, F. Balarin, R. Brayton, and A. Sangiovanni-Vincentelli. It usually works: The temporal logic of stochastic systems. In Pierre Wolper, editor, *Computer Aided Verification*, volume 939 of *Lecture Notes in Computer Science*, pages 155–165. Springer Berlin Heidelberg, 1995.
- [17] C. Baier and J.-P. Katoen. *Principles of model checking*. The MIT Press, 2008.
- [18] R. Bellman. The theory of dynamic programming. *Bull. Amer. Math. Soc.*, 60:503–515, 1954.
- [19] R. Bellman. A Markovian decision process. *J. Math. Mech.*, 6:679–684, 1957.
- [20] D. Bertsekas and S. Shreve. *Stochastic optimal control: The discrete time case*, volume 139. Academic Press, 1978.
- [21] E. Bishop and K. de Leeuw. The representations of linear functionals by measures on sets of extreme points. *Annales de l’institut Fourier*, 9:305–331, 1959.
- [22] D. Blackwell. Positive dynamic programming. In *Proc. Fifth Berkeley Sympos. Math. Statist. and Probability (Berkeley, Calif., 1965/66)*, Vol. I: Statistics, pages 415–418. Univ. California Press, Berkeley, Calif., 1967.
- [23] D. Blackwell. A Borel set not containing a graph. *The Annals of Mathematical Statistics*, 39(4):1345–1347, 08 1968.
- [24] D. Blackwell. The stochastic processes of Borel gambling and dynamic programming. *Ann. Statist.*, 4(2):370–374, 1976.

- [25] R. Blute, J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for labelled markov processes. In *Proceedings of 12th Annual IEEE Symposium on Logic in Computer Science*, 1997, pages 149–158, 1997.
- [26] V. Borkar. *Topics in controlled Markov chains*, volume 240 of *Pitman Research Notes in Mathematics Series*. Longman Scientific & Technical, Harlow, 1991.
- [27] J. J. Buckley. Graphs of measurable functions. *Proceedings of AMS*, 44:78–80, 1974.
- [28] C. Cassandras and J. Lygeros, editors. *Stochastic hybrid systems*, volume 24. CRC Press, 2007.
- [29] P. Chaput, V. Danos, P. Panangaden, and G. Plotkin. Approximating Markov processes by averaging. In *Proceedings of the 36th Internatilonal Colloquium on Automata, Languages and Programming: Part II, ICALP '09*, pages 127–138, Berlin, Heidelberg, 2009. Springer-Verlag.
- [30] D. Chatterjee, E. Cinquemani, and J. Lygeros. Maximizing the probability of attaining a target prior to extinction. *Nonlinear Analysis: Hybrid Systems*, 5(2):367 – 381, 2011.
- [31] T. Chen and S. Kiefer. On the total variation distance of labelled Markov chains. *CoRR*, abs/1405.2852, 2014.
- [32] C. Courcoubetis and M. Yannakakis. Markov decision processes and regular events. *IEEE Trans. Automat. Control*, 43(10):1399–1418, 1998.
- [33] H. Crauel. *Random probability measures on Polish spaces*, volume 11 of *Stochastics Monographs*. Taylor & Francis, London, 2002.
- [34] V. Danos, J. Desharnais, and P. Panangaden. Labelled Markov processes: Stronger and faster approximations. *Electronic Notes in Theoretical Computer Science*, 87:157–203, November 2004.
- [35] P. D’Argenio, B. Jeannet, H. Jensen, and K. Larsen. Reachability analysis of probabilistic systems by successive refinements. In *Proceedings of the Joint International Workshop on Process Algebra and Probabilistic Methods, Performance Modeling and Verification*, PAPM-PROBMIV ’01, pages 39–56, London, UK, UK, 2001. Springer-Verlag.
- [36] B. De Finetti. *Theory of probability: a critical introductory treatment*, volume 6. John Wiley & Sons, 2017.
- [37] E. de Vink and J. Rutten. Bisimulation for probabilistic transition systems: a coalgebraic approach. *Theoretical Computer Science*, 221(1-2):271–293, June 1999.
- [38] C. Dellacherie. Capacities and analytic sets. In A. Kechris, D. Martin, and Y. Moschovakis, editors, *Cabal Seminar 77-79*, volume 839 of *Lecture Notes in Mathematics*, pages 1–31. Springer Berlin Heidelberg, 1981.

- [39] C. Dellacherie and P.-A. Meyer. *Probabilities and Potential*. C. Elsevier, 1988.
- [40] J. Desharnais, A. Edalat, and P. Panangaden. A logical characterization of bisimulation for labeled Markov processes. In *Logic in Computer Science, 1998. Proceedings. Thirteenth Annual IEEE Symposium on*, pages 478–487, Jun 1998.
- [41] J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for Labelled Markov Processes. *Information and Computation*, 179(2):163 – 193, 2002.
- [42] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for Labeled Markov Systems. In *Proceedings of the 10th International Conference on Concurrency Theory, CONCUR '99*, pages 258–273, London, UK, 1999. Springer-Verlag.
- [43] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labelled Markov processes. *Theoretical Computer Science*, 318(3):323–354, 2004.
- [44] J. Desharnais, F. Laviolette, and M. Tracol. Approximate analysis of probabilistic processes: Logic, simulation and games. In *Fifth International Conference on Quantitative Evaluation of Systems, 2008. QEST '08.*, pages 264–273, Sept 2008.
- [45] J. Ding, A. Abate, and C. Tomlin. Optimal control of partially observable discrete time stochastic hybrid systems for safety specifications. *Proceedings of the 32nd American Control Conference*, pages 6231–6236, 2013.
- [46] A. D’Innocenzo, A. Abate, and J.-P. Katoen. Robust PCTL model checking. In *Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control*, pages 275–286. ACM, 2012.
- [47] E.-E. Doberkat. *Stochastic relations*. Chapman & Hall/CRC Studies in Informatics Series. Chapman & Hall/CRC, Boca Raton, FL, 2007. Foundations for Markov transition systems.
- [48] L. E. Dubins and L. J. Savage. *How to gamble if you must. Inequalities for stochastic processes*. McGraw-Hill Book Co., New York, 1965.
- [49] F. Dufour and T. Prieto-Rumeau. Approximation of Markov decision processes with general state space. *J. Math. Anal. Appl.*, 388(2):1254–1267, 2012.
- [50] R. Durrett. *Probability: Theory and Examples - Third Edition*. Duxbury Press, 2004.
- [51] M. Ershov. Extension of measures and stochastic equations. *Theory of Probability & Its Applications*, 19(3):431–444, 1975.
- [52] A. Faden. The existence of regular conditional probabilities: Necessary and sufficient conditions. *The Annals of Probability*, 13(1):288–298, 02 1985.
- [53] H. Fecher, M. Leucker, and V. Wolf. Don’t know in probabilistic systems. In *Proceedings of the 13th international conference on Model Checking Software, SPIN’06*, pages 71–88, Berlin, Heidelberg, 2006. Springer-Verlag.

- [54] E. Feinberg. Controlled Markov processes with arbitrary numerical criteria. *Theory of Probability & Its Applications*, 27(3):486–503, 1983.
- [55] E. Feinberg. On measurability and representation of strategic measures in Markov decision processes. In *Statistics, probability and game theory*, volume 30 of *IMS Lecture Notes Monogr. Ser.*, pages 29–43. Inst. Math. Statist., Hayward, CA, 1996.
- [56] E. Feinberg and A. Shwartz, editors. *Handbook of Markov decision processes*. International Series in Operations Research & Management Science, 40. Kluwer Academic Publishers, Boston, MA, 2002. Methods and applications.
- [57] N. Ferns, P. Panangaden, and D. Precup. Metrics for markov decision processes with infinite state spaces. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, UAI’05, pages 201–208, Arlington, Virginia, United States, 2005. AUAI Press.
- [58] N. Ferns, P. Panangaden, and D. Precup. Bisimulation metrics for continuous Markov decision processes. *SIAM Journal on Computing*, 40(6):1662–1714, 2011.
- [59] P. Florchinger. Lyapunov-like techniques for stochastic stability. *SIAM Journal on Control and Optimization*, 33(4):1151–1169, 1995.
- [60] G. Folland. *Real analysis*. Pure and Applied Mathematics. John Wiley & Sons Inc., New York, second edition, 1999.
- [61] A. Giacalone, C. Jou, and S. Smolka. Algebraic reasoning for probabilistic concurrent systems. In *Proceedings of IFIP TC2 Working Conference on Programming Concepts and Methods*, pages 443–458. North-Holland, 1990.
- [62] A. Gibbs and F. Su. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, 2002.
- [63] D. Gilbarg and N. Trudinger. *Elliptic partial differential equations of second order*. springer, 2015.
- [64] A. Girard and G. Pappas. Approximation metrics for discrete and continuous systems. *IEEE Transactions on Automatic Control*, 52(5):782–798, 2007.
- [65] M. Giry. A categorical approach to probability theory. In B. Banaschewski, editor, *Categorical Aspects of Topology and Analysis*, volume 915 of *Lecture Notes in Mathematics*, pages 68–85. Springer Berlin Heidelberg, 1982.
- [66] I. Gyöngy and D. Šiška. On finite-difference approximations for normalized bellman equations. *Applied Mathematics and Optimization*, 60(3):297–339, 2009.
- [67] S. Haesaert, S. Soudjani, and A. Abate. Verification of general markov decision processes by approximate similarity relations and policy refinement. *SIAM Journal on Control and Optimization*, 55(4):2333–2367, 2017.

- [68] E. Haghverdi, P. Tabuada, and G. Pappas. Bisimulation relations for dynamical, control, and hybrid systems. *Theoretical Computer Science*, 342(2-3):229–261, 2005.
- [69] O. Hernández-Lerma. *Adaptive Markov control processes*, volume 79 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 1989.
- [70] O. Hernández-Lerma and J. B. Lasserre. *Discrete-time Markov control processes*, volume 30 of *Applications of Mathematics (New York)*. Springer Verlag, New York, 1996.
- [71] O. Hernández-Lerma and J. B. Lasserre. *Further topics on discrete-time Markov control processes*, volume 42 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 1999.
- [72] T. Hill. On the existence of good Markov strategies. *Trans. Amer. Math. Soc.*, 247:157–176, 1979.
- [73] A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker. PRISM: A tool for automatic verification of probabilistic systems. In H. Hermanns and J. Palsberg, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 3920 of *Lecture Notes in Computer Science*, pages 441–444. Springer Verlag, 2006.
- [74] S. Jha, E. Clarke, C. Langmead, A. Legay, A. Platzer, and P. Zuliani. A Bayesian approach to model checking biological systems. In *Computational Methods in Systems Biology*, pages 218–234. Springer, 2009.
- [75] B. Jonsson and K. Larsen. Specification and refinement of probabilistic processes. In *Proceedings of Sixth Annual IEEE Symposium on Logic in Computer Science*, pages 266–277, 1991.
- [76] A. Julius and G. Pappas. Approximations of stochastic hybrid systems. *IEEE Transactions on Automatic Control*, 54(6):1193–1203, 2009.
- [77] R. Kalman et al. Contributions to the theory of optimal control. *Bol. Soc. Mat. Mexicana*, 5(2):102–119, 1960.
- [78] M. Kamgarpour, S. Summers, and J. Lygeros. Control design for specifications on stochastic hybrid systems. In *Proceedings of the 16th international conference on Hybrid systems: computation and control, HSCC '13*, pages 303–312, New York, NY, USA, 2013. ACM.
- [79] N. Kariotoglou, S. Summers, T. Summers, M. Kamgarpour, and J. Lygeros. Approximate dynamic programming for stochastic reachability. In *2013 European Control Conference (ECC)*, pages 584–589, July 2013.
- [80] J.-P. Katoen, M. Khattri, and I. Zapreev. A Markov reward model checker. In *Proceedings of the Second International Conference on the Quantitative Evaluation of Systems, QEST '05*, pages 243–244, Washington, DC, USA, 2005. IEEE Computer Society.

- [81] J.-P. Katoen, D. Klink, M. Leucker, and V. Wolf. Three-valued abstraction for probabilistic systems. *The Journal of Logic and Algebraic Programming*, 81(4):356 – 389, 2012.
- [82] A. Kechris. *Classical descriptive set theory*, volume 156 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.
- [83] D. Kreps. Decision problems with expected utility criteria. I. Upper and lower convergent utility. *Math. Oper. Res.*, 2(1):45–53, 1977.
- [84] D. Kreps. Decision problems with expected utility criteria. II. Stationarity. *Math. Oper. Res.*, 2(3):266–274, 1977.
- [85] D. Kreps. Decision problems with expected utility criteria. III. Upper and lower transience. *SIAM J. Control Optim.*, 16:420–428, 1978.
- [86] N. Krylov. *Controlled diffusion processes*, volume 14. Springer Science & Business Media, 2008.
- [87] O. Kupferman and M. Vardi. Model checking of safety properties. In *Computer aided verification (Trento, 1999)*, volume 1633 of *Lecture Notes in Comput. Sci.*, pages 172–183. Springer, Berlin, 1999.
- [88] H. Kushner. *Approximation and weak convergence methods for random processes, with applications to stochastic systems theory*, volume 6. MIT press, 1984.
- [89] M. Kwiatkowska, G. Norman, and D. Parker. Game-based abstraction for Markov decision processes. In *Proceedings of the 3rd International Conference on Quantitative Evaluation of Systems (QEST’06)*, pages 157–166. IEEE CS Press, 2006.
- [90] K. Larsen and A. Skou. Bisimulation through probabilistic testing (preliminary report). In *Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL ’89*, pages 344–352, New York, NY, USA, 1989. ACM.
- [91] K. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, 1991.
- [92] F. Lawvere. The category of probabilistic mappings. *preprint*, 1962. <http://www.fuw.edu.pl/~kostecki/scans/lawvere1962.pdf>.
- [93] L. Leskelä. Stochastic relations of random variables and processes. *Journal of Theoretical Probability*, 23(2):523–546, 2010.
- [94] A. Maitra, R. Purves, and W. Sudderth. A Borel measurable version of König’s lemma for random paths. *Ann. Probab.*, 19(1):423–451, 1991.
- [95] A. Maitra and W. Sudderth. *Discrete gambling and stochastic games*, volume 32 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 1996.

- [96] A. Maitra and W. Sudderth. The gambler and the stopper. In *Statistics, probability and game theory*, volume 30 of *IMS Lecture Notes Monogr. Ser.*, pages 191–208. Inst. Math. Statist., Hayward, CA, 1996.
- [97] A. Maitra and W. Sudderth. Randomized strategies and terminal distributions. In *Game theory, optimal stopping, probability and statistics*, volume 35 of *IMS Lecture Notes Monogr. Ser.*, pages 39–52. Inst. Math. Statist., Beachwood, OH, 2000.
- [98] A. Maitra and W. Sudderth. Saturations of gambling houses. In J. Azema, M. Ledoux, M. Emery, and M. Yor, editors, *Seminaire de Probabilites XXXIV*, volume 1729 of *Lecture Notes in Mathematics*, pages 218–238. Springer Berlin Heidelberg, 2000.
- [99] A. Markov. An example of statistical investigation of the text Eugene Onegin concerning the connection of samples in chains. *Science in Context*, 19:591–600, 12 2006.
- [100] S. Meyn and R. Tweedie. *Markov chains and stochastic stability*. Communications and Control Engineering Series. Springer-Verlag London Ltd., London, 1993.
- [101] R. Milner. An algebraic definition of simulation between programs. In *Proceedings of the 2nd International Joint Conference on Artificial Intelligence, IJCAI'71*, pages 481–489, San Francisco, CA, USA, 1971. Morgan Kaufmann Publishers Inc.
- [102] P. Mohajerin Esfahani, D. Chatterjee, and J. Lygeros. Motion planning for continuous-time stochastic processes: A dynamic programming approach. *IEEE Transactions on Automatic Control*, 61(8):2155–2170, 2016.
- [103] P. Mohajerin Esfahani, D. Chatterjee, and J. Lygeros. The stochastic reach-avoid problem and set characterization for diffusions. *Automatica*, 70:43–56, 2016.
- [104] L. Moss and I. Viglizzo. Final coalgebras for functors on measurable spaces. *Information and Computation*, 204(4):610 – 636, 2006. Seventh Workshop on Coalgebraic Methods in Computer Science 2004.
- [105] J. Norris. *Markov chains*. Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, 1998.
- [106] B. Øksendal. *Stochastic differential equations: an introduction with applications*. Universitext. Springer-Verlag, Berlin, sixth edition, 2003.
- [107] G. Pappas. Bisimilar linear systems. *Automatica*, 39(12):2035–2047, 2003.
- [108] K. Parthasarathy. *Probability measures on metric spaces*. Probability and Mathematical Statistics, No. 3. Academic Press Inc., New York, 1967.
- [109] L. Perko. *Differential equations and dynamical systems*, volume 7 of *Texts in Applied Mathematics*. Springer Verlag, New York, third edition, 2001.

- [110] D. Perrin and J.-E. Pin. *Infinite words: automata, semigroups, logic and games*, volume 141. Academic Press, 2004.
- [111] M. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, Inc., 1994.
- [112] D. Revuz. *Markov chains*. North-Holland Publishing, Amsterdam, second edition, 1984.
- [113] M. Rungger, M. Mazo, and P. Tabuada. Specification-guided controller synthesis for linear systems and safe linear-time temporal logic. In *Proceedings of the 16th international conference on Hybrid systems: computation and control, HSCC '13*, pages 333–342, New York, NY, USA, 2013. ACM.
- [114] R. Segala and N. Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2(2):250–273, 1995.
- [115] M. Shaked and J. Shanthikumar. *Stochastic orders and their applications*. Academic Press, 1994.
- [116] A. Shiryaev, B. Rozovskii, and G. Grimmett. *Optimal stopping rules*. Springer Verlag Berlin Heidelberg, 2008.
- [117] S. Shreve and D. Bertsekas. Universally measurable policies in dynamic programming. *Math. Oper. Res.*, 4(1):15–30, 1979.
- [118] E. Sontag. *Mathematical control theory*, volume 6 of *Texts in Applied Mathematics*. Springer-Verlag, New York, second edition, 1998. Deterministic finite-dimensional systems.
- [119] S. Soudjani and A. Abate. Adaptive gridding for abstraction and verification of stochastic hybrid systems. In *Proceedings of the 2011 Eighth International Conference on Quantitative Evaluation of SysTems, QEST '11*, pages 59–68, Washington, DC, USA, 2011. IEEE Computer Society.
- [120] S. Srivastava. *A course on Borel sets*, volume 180 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998.
- [121] R. E. Strauch. Negative dynamic programming. *Ann. Math. Statist.*, 37:871–890, 1966.
- [122] W. Sudderth. On the existence of good stationary strategies. *Trans. Amer. Math. Soc.*, 135:399–414, 1969.
- [123] W. Sudderth. Gambling problems with a limit inferior payoff. *Math. Oper. Res.*, 8(2):287–297, 1983.
- [124] S. Summers and J. Lygeros. Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46(12):1951–1961, 2010.
- [125] M. Suslin. Sur une définition des ensembles mesurables B sans nombres transfinis. *CR Acad. Sci. Paris*, 164(2):88–91, 1917.

- [126] J. Swart. A conditional product measure theorem. *Statist. Probab. Lett.*, 28(2):131–135, 1996.
- [127] P. Tabuada. Symbolic control of linear systems based on symbolic subsystems. *IEEE Transactions on Automatic Control*, 51(6):1003–1013, June 2006.
- [128] P. Tabuada. *Verification and control of hybrid systems: A symbolic approach*. Springer Verlag, New York, 2009.
- [129] I. Tkachev and A. Abate. On infinite-horizon probabilistic properties and stochastic bisimulation functions. In *2011 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, pages 526–531, 2011.
- [130] I. Tkachev and A. Abate. Regularization of Bellman equations for infinite-horizon probabilistic properties. In *Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control, HSCC '12*, pages 227–236, New York, NY, USA, 2012. ACM.
- [131] I. Tkachev and A. Abate. Stability and attractivity of absorbing sets for discrete-time markov processes. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 7652–7657, 2012.
- [132] I. Tkachev and A. Abate. A control lyapunov function approach for the computation of the infinite-horizon stochastic reach-avoid problem. In *52nd IEEE Conference on Decision and Control*, pages 3211–3216, Dec 2013.
- [133] I. Tkachev and A. Abate. Formula-free finite abstractions for linear temporal verification of stochastic hybrid systems. In *Proceedings of the 16th international conference on Hybrid Systems: Computation and Control, HSCC '13*, pages 283–292, New York, NY, USA, 2013. ACM.
- [134] I. Tkachev and A. Abate. Characterization and computation of infinite-horizon specifications over Markov processes. *Theoretical Computer Science*, 515:1 – 18, 2014.
- [135] I. Tkachev and A. Abate. On approximation metrics for linear temporal model-checking of stochastic systems. In *HSCC*, pages 193–202, 2014.
- [136] I. Tkachev and E. Feinberg. On equivalence between markov decision processes and gambling models. In *arxiv*, pages 3211–3216, Dec 2018.
- [137] I. Tkachev, A. Mereacre, J.-P. Katoen, and A. Abate. Quantitative automata-based controller synthesis for non-autonomous stochastic hybrid systems. In *Proceedings of the 16th international conference on Hybrid Systems: Computation and Control, HSCC '13*, pages 293–302, New York, NY, USA, 2013. ACM.
- [138] I. Tkachev, A. Mereacre, J.-P. Katoen, and A. Abate. Quantitative model-checking of controlled discrete-time Markov processes. *arXiv preprint*, 2014. arXiv:1407.5449.

- [139] F. van Breugel, S. Shalit, and J. Worrell. Testing labelled Markov processes. In *Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, ICALP '02, pages 537–548, London, UK, 2002. Springer-Verlag.
- [140] F. van Breugel, B. Sharma, and J. Worrell. Approximating a behavioural pseudometric without discount for probabilistic systems. In *Proceedings of the 10th international conference on Foundations of software science and computational structures*, FOSSACS'07, pages 123–137, Berlin, Heidelberg, 2007. Springer-Verlag.
- [141] F. van Breugel and J. Worrell. An algorithm for quantitative verification of probabilistic transition systems. In *Proceedings of the 12th International Conference on Concurrency Theory*, CONCUR '01, pages 336–350, London, UK, 2001. Springer-Verlag.
- [142] F. van Breugel and J. Worrell. Towards quantitative verification of probabilistic transition systems. In *Proceedings of the 28th International Colloquium on Automata, Languages and Programming*, ICALP '01, pages 421–432, London, UK, 2001. Springer-Verlag.
- [143] F. van Breugel and J. Worrell. A behavioural pseudometric for probabilistic transition systems. *Theoretical Computer Science*, 331(1):115–142, February 2005.
- [144] F. van Breugel and J. Worrell. Approximating and computing behavioural distances in probabilistic transition systems. *Theoretical Computer Science*, 360(1):373–385, August 2006.
- [145] M. Vardi. Automatic verification of probabilistic concurrent finite state programs. In *26th Annual Symposium on Foundations of Computer Science*, pages 327–338, 1985.
- [146] D. Wagner. Survey of measurable selection theorems. *SIAM Journal on Control and Optimization*, 15(5):859–903, 1977.
- [147] P. Wolper. Temporal logic can be more expressive. In *Foundations of Computer Science, 1981. SFCS '81. 22nd Annual Symposium on*, pages 340–348, 1981.
- [148] M. Zamani, P. M. Esfahani, A. Abate, and J. Lygeros. Symbolic Models for Stochastic Control Systems without Stability Assumptions. In *European Control Conference*, Zurich, Switzerland, July 2013.
- [149] M. Zamani, I. Tkachev, and A. Abate. Bisimilar symbolic models for stochastic control systems without state-space discretization. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control*, HSCC '14, pages 41–50, New York, NY, USA, 2014. ACM.
- [150] K. Zhou and John C. Doyle. *Essentials of robust control*, volume 104. Prentice hall Upper Saddle River, NJ, 1998.

List of symbols

$\mathbb{N} = \{0, 1, 2, \dots\}$	set of non-negative integer numbers
$[a; b] = \{a, a + 1, \dots, b\}$	finite set of integers from a to b
\mathbb{R}	set of real numbers
\rightarrow	going to
\mapsto	maps to
$\mathcal{B}(\Omega)$	Borel σ -algebra of a Borel space Ω
$\mathcal{U}(\Omega)$	universal σ -algebra of a Borel space Ω
$\mathcal{A}(\Omega)$	collection of analytic subsets of a Borel space Ω
$\mathcal{P}(\Omega)$	Borel space of all probability measure over a Borel space Ω
f_*	a pushforward of a map f from points to measures
R_*	a lifting of a relation R from points to measures
$\kappa(A \omega)$	evaluation of a stochastic kernel κ on a set A given a point x
1_A	an indicator function of a set A
δ	a Dirac stochastic kernel
$R _x$	a left-section of a set R at x
$R _x^x$	a right-section of a set R at x
R^n	an n -times product of R with itself

List of Abbreviations

i.i.d.	independent identically distributed
TS	Transition System
SS	Stochastic System
PA	Probabilistic Automaton
PTS	Probabilistic Transition System
LMP	Labelled Markov Process
MDP	Markov Decision Process
GM	Gambling Model
MC	Markov Chain
AMC	Abstract Markov Chain
DFA	Deterministic Finite Automaton
DBA	Deterministic Büchi Automaton
DRA	Deterministic Rabin Automaton
LT	Linear Temporal
LTL	Linear Temporal Logic
BLTL	Bounded Linear Temporal Logic
PCTL	Probabilistic Computation Tree Logic
DP	Dynamic Programming
DC	Discounted additive cost Criterion
AC	Average cost Criterion
TC	Terminal cost Criterion
ABF	Approximate Bisimulation Function
l.t.r.	left-total relation

Summary

Bisimilar Stochastic Systems

Ilya Tkachev

Stochastic systems have been widely investigated and employed in numerous applications in different areas such as finance, biology and engineering as they allow accounting for imprecisions so often faced in every practical tasks. Often that task would require to find the best action sequence in order to optimize the outcome. When the model is small, one can efficiently employ algorithmic techniques to synthesize such a control policy. Hence, in case of more complex models, instead of solving control tasks there directly, one may want to approximate them with simpler ones and then use those algorithms. This method is called *abstraction* for it abstracts the original “physical” model to an “abstract” one, only needed to ease the computations. Ideally, this abstract model is somewhat similar to the original one, as we want to extrapolate results achieved over the former to the setting of the latter. One way this similarity can be ensured is by means of the (bi)simulation methods, that give sufficient conditions to the closeness of behaviors of the two systems being compared. Such techniques became popular in discrete non-stochastic models, then advanced to continuous ones and started making steps to discrete stochastic systems. Yet, definite results were not achieved for abstractions of continuous stochastic models. There were trials to extend ideas from continuous non-stochastic framework, or discrete stochastic one, but they were mostly fragmentary. This thesis brings those methods together to build a unified framework and shows immediate benefits of doing this.

To define the closeness between the systems we look at their path-wise properties, which cover most of the tasks whose relevance was praised in the literature. That comprises both additive cost-like criteria and formal specifications, e.g. encoded by LTL formulae of the kind “reach the goal set through the safe set while avoiding dangerous states”. We derive guarantees on the approximation error and suggest how to build an abstraction for a given tolerance level. These guarantees work mostly for the finite time horizon properties, hence for the rest we develop task-dependent solution methods, further connecting with the existing literature. Besides those concrete results, we also put some effort in developing the conceptual side of the bisimulation framework for stochastic systems. For example, we show how important it is to choose a definition of behavior here, since bisimilarity is useful as long as it guarantees closeness of behaviors one is interested in.

We hence stress the importance of keeping in mind the final goal while extrapolating abstract solution methods, and show which issues may arise when this goal is forgotten. We also extend the framework we deal with beyond bisimulation of stochastic systems only, providing a formalization of approximate relations and their connections with pseudo-metrics, proving several theorems in probabilistic approximation, whose generality is greater than the scope of this thesis, and also provide a category-theoretical basis for bisimulations of stochastic systems, hence opening one more door from which this problem can be approached.

Samenvatting

Bisimilaar Stochastische Systemen

Ilya Tkachev

Stochastische systemen zijn op grote schaal onderzocht en gebruikt in tal van toepassingen in verschillende gebieden, zoals financiën, biologie en bouwkunde, omdat ze rekening kunnen houden met de onnauwkeurigheden die zo vaak bij praktische taken voor komen. Vaak vereist een taak het vinden van de beste actiereeks om de uitkomst te optimaliseren. Als het model klein is, kan men op efficiënte wijze algoritmische technieken gebruiken om een dergelijke control policy te vinden. Complexere modellen zou men kunnen benaderen met eenvoudigere modellen teneinde dezelfde efficiënte algoritmische technieken te kunnen gebruiken om ze op te lossen. Deze benadering wordt abstractie genoemd omdat men het oorspronkelijke “fysieke” model abstraheert tot een “abstract” model, dat gebruikt wordt om de berekeningen te vergemakkelijken. Idealiter lijkt dit abstracte model enigszins op het originele model, omdat we de resultaten van het één willen extrapoleren naar het ander. Een manier om deze gelijkenis te waarborgen is door middel van (bi) simulatiemethoden, die voldoende voorwaarden geven aan de nabijheid van het gedrag van de twee systemen die worden vergeleken. Dergelijke technieken werden populair in discrete niet-stochastische modellen, vervolgens doorontwikkeld naar continue niet-stochastische systemen, en begonnen met het maken van stappen naar discrete stochastische systemen. Toch zijn geen definitieve resultaten bereikt voor abstracties van continue stochastische modellen. Er is geprobeerd om ideeën uit te breiden vanuit het continue of discrete niet-stochastisch raamwerk, maar deze pogingen waren meestal fragmentarisch. Dit proefschrift brengt die methoden bij elkaar om een uniform raamwerk op te bouwen en laat de directe voordelen zien die hier uit voortkomen.

Om de nabijheid van systemen te bepalen, bekijken we hun padsgewijze eigenschappen; deze bestrijken de meeste relevante taken in de literatuur. Dat omvat zowel additieve kosten-achtige criteria als formele specificaties, bijv. gecodeerd door LTL-formules van het soort “bereik de doel set door de veilige set en vermijdt gevaarlijke toestanden”. We leiden garanties op de benaderingsfout af en suggereren hoe we een abstractie kunnen opbouwen voor een bepaald tolerantieniveau. Deze garanties werken meestal voor de eigenschappen van de eindige tijdshorizon, daarom ontwikkelen we voor de rest taakafhankelijke oplossingsme-

thoden, wat verder aansluit op de bestaande literatuur. Naast die concrete resultaten hebben we ook enige moeite gedaan om de conceptuele kant van het bisimulatiekader voor stochastische systemen te ontwikkelen. We laten bijvoorbeeld zien hoe belangrijk het is om hier een definitie van systeem gedrag te kiezen, omdat bisimilariteit alleen nuttig is als het garanties geeft op dat nabijheid van het systeem gedrag waarin iemand geïnteresseerd is. We benadrukken daarom het belang van het in gedachten houden van het uiteindelijke doel terwijl we abstracte oplossingsmethoden extrapoleren, en laten zien welke problemen zich kunnen voordoen wanneer dit doel uit het oog wordt verloren. Hiernaast breiden we het raamwerk uit met methoden die verder gaan dan alleen bisimulatie van stochastische systemen, en verschaffen een formalisering van benaderde relaties en hun connecties met pseudo-metrics. Verder bewijzen we verschillende stellingen in de probabilistische benadering waarvan de algemeenheid verder gaat dan alleen het toepassingsgebied van dit proefschrift, en die ook een categorietheoretische basis bieden voor bisimulaties van stochastische systemen, wat een nieuwe invalshoek verschaft van waaruit dit probleem kan worden benaderd.

Curriculum Vitae

Ilya Tkachev was born on 15 December 1987 in Volgograd Russia. Towards the end of his bachelor's program in Mathematics in Volgograd State University in he started working as an SAP consultant and in parallel became interested in trading. After graduating summa cum laude in 2008, he focused on his work for a year, and then inspired by his fascination in finance, he went to Halmstad University in Sweden in 2009, where he completed an MSc degree in Financial Mathematics in 2010. At that moment he thought of going to financial industry, but while writing his master thesis on options pricing under the supervision of Professor Mikhail Nechaev, he found excitement in research, and decided to join a group of Professor Alessabdro Abate at Delft Center of Systems and Control, Delft University of Technology, to pursue a career in academia. In his research, he was involved in the European project "MoVeS" on modeling, verification, and control of complex systems. With the time his main focus became on abstract problems arising from comparing behaviors of different stochastic dynamic models, and he tried applying them in diverse fields such as Systems Biology, Finance and control of Power Networks. After leaving the university in 2014 he joined trading and research in a Dutch company Optiver, where he also started educating himself in Machine Learning (ML). After three years there, he moved back to Russia and settled in Moscow, where he currently works on diverse ML projects.

