Dissertation ETH Zurich No. 21725

Stochastic Motion Planning for Diffusions & Fault Detection and Isolation for Large Scale Nonlinear Systems

A dissertation submitted to
ETH Zurich

for the degree of
Doctor of Sciences

presented by

Peyman Mohajerin Esfahani

M.Sc., Sharif University of Technology, Tehran
born 21.03.1982 in Tehran
citizen of Iran

accepted on the recommendation of

Prof. Dr. John Lygeros, examiner
Prof. Dr. Halil Mete Soner, co-examiner
Prof. Dr. Karl Henrik Johansson, co-examiner

January, 2014

ISBN 978-3-906031-58-3

To my parents, Farahnaz and Hasan
... and to my wife, Farzaneh

Acknowledgment

When I started to write this part, I was going back and forth through countless great memories during all these years in my PhD at the Automatic Control Lab. This was indeed a unique opportunity to see many experts and leaders of different areas, and being privileged to work closely with some of them.

First and foremost, I would like to express my deepest gratitude to my advisor and mentor, Professor John Lygeros, for giving me such a wonderful opportunity and continuously supporting me with his vast knowledge and patience. This thesis only reflects part of what I have learned from him. His enthusiasm in research, sharp thinking, open-mindedness to new territories, and vision are always inspiring to me. John, it was a great pleasure and an honor to work with you; thank you for everything. Had I another chance to redo my PhD, I would definitely come back to you, and hopefully you would accept me again!

It was an incredible luck for me to meet Professor Mete Soner at ETH. I would like to sincerely thank Mete not only for serving in my PhD exam committee, but also for several wonderful courses that I have taken with him during my studies. I greatly appreciate his kind willingness and precious time to share his immense knowledge and experience with me, starting from the beginning when I was just a newbie to finally as my co-examiner.

I am especially grateful to Professor Karl Johansson who kindly accepted to be in my exam committee and provided me detailed and valuable comments on this thesis. Kalle, thank you also for your kind hospitality during my short visit to KTH. I owe a special thanks to Professor Göran Andersson. Our joint internal meetings with him during the VIKING project were always fun, yet instructive occasions. With Göran I had a fantastic chance to learn more about power systems. I was always amazed by his intuition in power systems, in particular when he could guess our numerical results obtained from sophisticated reachability PDEs!

Thanks to Professor Manfred Morari and John, the Automatic Control Lab provided a priceless collaborative environment to work with brightest minds and genuinely friendly people. Let me start with Professor Debasish Chatterjee, to whom I am always indebted. His vast mathematical knowledge was an encouraging motivation to me, and I learned a lot from him during my first three years of PhD when he was a post-doc in our lab. Debasish, thank you so much for guiding me through these years, and of course, I will never forget that scary hand jumping out from the window late one evening at the coffee room!

My special thanks go to Professor Federico Ramponi, who was not only a wonderful office-mate, but also a great teacher. His diverse and deep understanding of different topics from topology and algebra in math to programming and literature were influential to me. Sir Federico, I owe you a large debt for showing me a new perspective to research, and your sentence will always stick in my mind: “never be ashamed of going back to basics”. I would like also to thank Peter Hokayem for his ever kind attention and helpful discussions on scientific as well as social topics during my PhD. Peter, many thanks also for your nice hospitality at ABB; I really appreciate it.

At this point, I would like to express my deep respect to my senior colleagues Andreas Milias and Stefan Richter. Andreas and Stefan, your genuinely deep knowledge and humility were always inspiring to me. You may not know of it, but you both were really influential to my view on research. I would like to thank my fantastic officemate Robin Vujanic with whom hardly do I remember a day without a fight in that office! Without Robin and Stefan I would probably never dare to step in optimization territories. Robin, I owe you also a debt of gratitude for being an ear for my nonsensical daily complaints. My special thanks also go to our chef Robert Nguyen not only for having great lunchtimes and hints on cooking (and, of course, the delicious chess cake at my defense apero!), but also for his unprecedented unbiased opinion on any topic; that was indeed instructive for me.

In these years there have been many people that I have worked with and would like to offer my sincere thanks: Maria Vrakopoulou and Kostas Margellos for the collaboration in the VIKING project, Majid Zamani and Farhad Farokhi my ex-colleagues from Sharif and current collaborators, Maryam Kamgarpour, Angelos Georghiou with special thanks for listening to my last minute PhD defense rehearsal, Frauke Oldewurtel, Georg Schildbach for sharing great times discussing both technical and social topics, Tony Wood not only for being a creative collaborator but also as a good friend and fellow traveler, and Tobias Sutter with whom I have been enjoying a lot to work and learn (also many thanks for translating the thesis abstract in German).

At a personal level, I want to thank my terrific Iranian friends who made my Zurich life memorable: Samira, Mahmoud, Hanieh, Farhad, Elham, and Maysam, you are simply amazing people not only for fun times, but also for always being patient and helpful to share difficult moments. Special thanks go to Mahmoud for the most relieving sentence right after the defense! I also would like to thank Arya, Kyumars, Nemat, Tofigh, Araz, Sahar, Pegah, and many others that I may have forgotten to name here.

Finally, my honest and deepest thanks go to my parents, Farahnaz and Hasan, who have been supporting me through all these steps with their unconditional love and sacrifice. At this time, I want to take this opportunity to sincerely thank my brother, Payam (barat!), whose presence beside our family has always brought me invaluable confidence to continue my career far away from them. Last, but most importantly, I want to express my gratitude to my beloved wife, Farzaneh, who stood by me and patiently embraced all my tough times during these years. It would be impossible to express my feeling in words, so let me keep it simple: I love you.

Peyman Mohajerin Esfahani
Zurich, March 2014

Abstract

The main theme of this thesis is twofold. First, we study a class of specifications, mainly the reachability type questions, in the context of controlled diffusion processes. The second part of the thesis is centered around the fault detection and isolation (FDI) problem for large scale nonlinear dynamical systems.

Reachability is a fundamental concept in the study of dynamical systems, and in view of applications of this concept ranging from engineering, manufacturing, biology, and economics, to name but a few, has been studied extensively in the control theory literature. One particular problem that has turned out to be of fundamental importance in engineering is the so-called "*reach-avoid*" problem. In the deterministic setting this problem consists of determining the set of initial conditions for which one can find at least one control strategy to steer the system to a target set while avoiding certain obstacles. The focus of the first part in this thesis is on the stochastic counterpart of this problem with an extension to more sophisticated maneuvers which we call the "*motion planning*" problem. From the technical standpoint, this part can be viewed as a theoretical bridge between the desired class of specifications and the existing numerical tools (e.g., partial differential equation (PDE) solvers) that can be used for verification and control synthesis purposes.

The second part of the thesis focuses on the FDI problem for large scale nonlinear systems. FDI comprises two stages: residual generation and decision making; the former is the subject addressed here. The thesis presents a novel perspective along with a scalable methodology to design an FDI filter for high dimensional nonlinear systems. Previous approaches on FDI problems are either confined to linear systems, or they are only applicable to low dimensional dynamics with specific structures. In contrast, we propose an optimization-based approach to robustify a linear residual generator given some statistical information about the disturbance signatures, shifting attention from the system dynamics to the disturbance inputs. The proposed scheme provides an alarm threshold whose performance is quantified in a probabilistic fashion.

From the technical standpoint, the proposed FDI methodology is effectively a relaxation from a robust formulation to probabilistic constraints. In this light, the alarm threshold obtained via the optimization program has a probabilistic performance index. Intuitively speaking, one would expect to improve the false alarm rate by increasing the filter threshold. The goal of the last part of the thesis is to quantify this connection rigorously. Namely, in a more general setting including a class of non-convex problems, we establish a theoretical bridge between the optimum values of a robust program and its randomized counterpart. The theoretical results of this part are finally deployed to diagnose and mitigate a cyber-physical attack introduced by the interaction between IT infrastructure and power system.

Zusammenfassung

Diese Dissertation besteht aus zwei Teilen. Zuerst wird eine Klasse von Spezifikationen für kontrollierte Diffusionsprozesse untersucht, vor allem mit Bezug auf Fragen der Erreichbarkeit. Der zweite Teil der Arbeit behandelt das Problem der Fehlererkennung und -isolierung (FDI) in grossen nichtlinearen dynamischen Systemen.

Erreichbarkeit ist ein fundamentales Konzept in der Untersuchung von dynamischen Systemen und wird als solches ausführlich in der Regelungstechnik behandelt. Anwendungen finden sich in verschiedensten Disziplinen wie den Ingenieurwissenschaften, der Systembiologie, der Produktionstechnik und der Ökonomie. Im Bereich der Ingenieurwissenschaften ist insbesondere das sogenannte "*Erreichbarkeit-Vermeidungsproblem*" von fundamentaler Wichtigkeit. Im deterministischen Fall besteht das Problem darin, die Menge der Anfangsbedingungen zu bestimmen, für welche mindestens eine Regelstrategie existiert, die das System in eine gegebene Zielmenge führt, wobei bestimmte Hindernisse zu vermeiden sind. Der Fokus des ersten Teils dieser Arbeit liegt auf dem stochastischen Pendant zu diesem Problem, ergänzt mit komplizierteren Zielvorgaben, das als "*Trajektorienplanungsproblem*" bezeichnet wird. Von einem regelungstechnischen Standpunkt her kann dieser erste Teil als theoretische Verbindung von der gewünschten Spezifikation des Reglers und der vorhandenen numerischen Software (z.B. Lösungsmethoden für partiellen Differentialgleichungen) betrachtet werden, die zur Synthese und Verifikation verwendet werden kann.

Der zweite Teil dieser Dissertation befasst sich mit dem FDI Problem für grosse nichtlineare dynamische Systeme. FDI besteht aus zwei Schritten: die Bestimmung der Regelabweichung sowie die Entscheidungsfindung, wobei der Fokus hier auf dem ersten Schritt liegt. Diese Arbeit präsentiert eine neue Sichtweise zusammen mit einer Methodik für den Entwurf eines FDI Filters für hochdimensionale nichtlineare Systeme. Bisherige Methoden für FDI Probleme beschränken sich entweder auf lineare Systeme oder sind nur für niedrigdimensionale Systeme mit spezifischer Struktur anwendbar. Im Gegensatz dazu wird ein optimierungsbasierter Ansatz für den Entwurf eines robusten Verfahrens zur Bestimmung der Regelabweichung vorgestellt, basierend auf statistischer Information über die Störsignale. Dieser Ansatz verschiebt die Sichtweise weg von der Systemdynamik und hin zu der Störgrösse. Er führt auf ein einfaches Schema, das einen Alarm liefert sobald ein bestimmter Schwellenwert überschritten wird, wobei die Güte dieses Schwellenwertes probabilistisch quantifiziert werden kann.

Von einem mathematischen Standpunkt aus gesehen stellt die vorgeschlagene FDI Methodik eine Relaxation von einer robusten zu einer probabilistischen Formulierung dar. Aus diesem Blickwinkel wird klar, weshalb der Schwellenwert, der durch die Lösung des Optimierungsproblems ermittelt wird, eine probabilistische Güte besitzt. Intuitiv betrachtet wird die Anzahl von Fehlalarmen bei einer Erhöhung dieses Schwellenwertes ansteigen. Das Ziel des letzten Teils

dieser Dissertation ist es, diesen Zusammenhang genauer zu untersuchen. Zu diesem Zweck wird gezeigt, wie in einem sehr allgemeinen Rahmen eine Schranke zwischen dem optimalen Zielfunktionswert eines robusten Optimierungsproblems und dem einer Näherungslösung mittels Stichproben hergeleitet werden kann. Schliesslich wird betrachtet, wie diese theoretischen Ergebnisse verwendet werden können, um cyberphysische Angriffe auf die Schnittstelle zwischen IT-Infrastruktur und Stromversorgungssystem zu erkennen und auszuschalten.

Contents

CHAPTER 1

Introduction

This thesis mainly addresses two problems: motion planning of controlled diffusions, and the problem of fault detection and isolation (FDI) in the context of high dimensional nonlinear systems. The former is a generalized concept of the so-called reachability problem which has received much attention through the study of safety problems in the dynamics and control literature. The latter, the FDI problem, is one of the fundamental subjects in the design of highly reliable control systems.

The first part of the thesis studies the motion planning problem in continuous time and space setting. The two fields of robotics and control have contributed much to motion planning. In the robotics community, research on motion planning typically focuses on the computational issues along with considerations of basic kinematic limitations, while in the control community the emphasis is mainly on the dynamic behavior and specific aspects of trajectory performance that usually involve high order differential constraints. The approach of this thesis on the motion planning is aligned with the latter point of view. This viewpoint has been investigated in the literature in various combinations of deterministic or stochastic dynamics, discrete or continuous time and discrete or continuous space. To date however, there is no treatment in the literature that would allow one to deal with continuous state, continuous time stochastic systems. This thesis fills this gap by investigating motion planning problems for controlled diffusions governed by controlled stochastic differential equations. We tackle the problem from an optimal control perspective based on the dynamic programming argument, which leads to a PDE characterization of the desired set of initial conditions. More detail regarding this part and its contribution is provided in Section 1.1.1

The FDI problem, the central topic in the second part, is one of the main subject in the design of reliable control systems. The FDI task involves generating a diagnosis signal to detect the occurrence of a specific fault. This is typically accomplished by designing a filter with all available signals as inputs (e.g., control signals and given measurements) and a scalar output that implements a non-zero mapping from the fault to the residual while decoupling unknown disturbances. The concept of residual plays a central role for the FDI problem which has been extensively studied in the last two decades. In the literature, the existing approaches are either confined to linear systems or they are only applicable to low dimensional systems with specific structures. Here we will present a novel perspective along with a scalable methodology to robustify a linear residual generator for high dimensional nonlinear systems; the details regarding the proposed approach is discussed in Section 1.1.2.

1.1 Outline and Contributions

Here we outline the organization and contributions of the thesis:

1.1.1 Part I: Stochastic Motion Planning for Diffusions

In Part I the basic object of our study is an $\mathbb{R}^d$-valued controlled random process $(X_s^{t,x;\boldsymbol{u}})_{s\geq t}$, initialized at (t,x) under the control policy $\boldsymbol{u} \in \mathcal{U}_t$, where $\mathcal{U}_t$ is the set of admissible control policies at time t.

A. Stochastic Reach-Avoid Problem

In Chapter 2, we consider a class of stochastic reachability problems with state constraints. The main objective is to characterize the set of initial conditions $x \in \mathbb{R}^d$ for which there exists an admissible control strategy $\boldsymbol{u}$ such that with probability more than a given value $p > 0$ the state trajectory hits a target set A before visiting obstacle B. Previous approaches to solving these problems in continuous time and space context are either studied in the deterministic setting [Aub91] or address almost-sure stochastic notions [AD90]. In contrast, we propose a new methodology to tackle probabilistic specifications that are less stringent than almost sure requirements. More precisely, based on different arriving time requirements, we aim to characterize the following set of initial conditions:

Definition (Reach-Avoid)**.** *Consider a fixed initial time $t \in [0,T]$. Given sets $A, B \subset \mathbb{R}^d$, we define the following* ***reach-avoid*** *initial sets:*

$$\begin{aligned}
\mathrm{RA}(t,p;A,B) &:= \Big\{x \in \mathbb{R}^d \;\big|\; \exists \boldsymbol{u} \in \mathcal{U}_t \;\; : \\
&\qquad\qquad \mathbb{P}\Big(\exists s \in [t,T],\;\; X_s^{t,x;\boldsymbol{u}} \in A \;\; \textit{and} \;\forall r \in [t,s]\;\; X_r^{t,x;\boldsymbol{u}} \notin B\Big) > p\Big\}. \\
\widetilde{\mathrm{RA}}(t,p;A,B) &:= \Big\{x \in \mathbb{R}^d \;\big|\; \exists \boldsymbol{u} \in \mathcal{U}_t \;\; : \\
&\qquad\qquad \mathbb{P}\Big(X_T^{t,x;\boldsymbol{u}} \in A \;\; \textit{and} \;\forall r \in [t,T]\;\; X_r^{t,x;\boldsymbol{u}} \notin B\Big) > p\Big\}.
\end{aligned}$$

In a direct approach, based on the theory of stochastic target problems, the authors of [BET10] recently extended the almost sure requirement of [ST02a, ST02b] to the controlled probability of success; see also the recent book [Tou13]. Here, following the same problem but in an indirect approach, we first establish a link from the above sets of initial conditions to a class of stochastic optimal control problems. In this light, we characterize the desired sets based on the tools from PDEs. Due to the discontinuities of the value functions involved, the PDE is understood in the generalized notion of the so-called discontinuous viscosity solutions. Furthermore, we provide theoretical justifications so that the reach-avoid problem is amenable to numerical solutions by means of off-the-shelf PDE solvers.

B. Stochastic Motion Planning

Chapter 3 generalizes the reach-avoid problem discussed in Chapter 2 to a motion planning specification. Motion planning of dynamical systems can be viewed as a scheme for executing excursions of the state of the system to certain given sets in a specific order according to a specified time schedule. Formally speaking, we consider the following set of initial conditions:

Definition (Motion-Planning)**.** *Consider a fixed initial time $t \in [0,T]$. Given a sequence of set pairs $(W_i, G_i)_{i=1}^n$ and horizon times $(T_i)_{i=1}^n \subset [t,T]$, we define the following* ***motion-planning*** *initial sets:*

$$\begin{aligned}
&\mathrm{MP}\big(t,p;\ (W_i,G_i)_{i=1}^n, T\big) \coloneqq \\
&\qquad \Big\{x \in \mathbb{R}^d \ \Big|\ \exists \boldsymbol{u} \in \mathcal{U}_t \ :\ \mathbb{P}\Big\{\exists (s_i)_{i=1}^n \subset [t,T] \mid X_{s_i}^{t,x;\boldsymbol{u}} \in G_i \ \textit{and}\ X_r^{t,x;\boldsymbol{u}} \in W_i \setminus G_i, \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall r \in [s_{i-1}, s_i[,\ \forall i \le n\Big\} > p\Big\}, \\
&\widetilde{\mathrm{MP}}\big(t,p;\ (W_i,G_i)_{i=1}^n, (T_i)_{i=1}^n\big) \coloneqq \\
&\qquad \Big\{x \in \mathbb{R}^d \ \Big|\ \exists \boldsymbol{u} \in \mathcal{U} \ :\ \mathbb{P}\Big\{X_{T_i}^{t,x;\boldsymbol{u}} \in G_i \ \textit{and}\ X_r^{t,x;\boldsymbol{u}} \in W_i, \\
&\qquad\qquad\qquad\qquad\qquad\qquad \forall r \in [T_{i-1}, T_i], \quad \forall i \le n\Big\} > p\Big\}.
\end{aligned}$$

Despite extensive studies on motion planning objectives in the deterministic setting [Sus91, CS98a, MS90] as well as stochastic but discrete time or space [CCL11, SL10, BHKH05], the continuous time and space settings seem to be investigated much less. In fact, our formal definition of the motion planning above is, to the best of our knowledge, new in the literature. Through a dynamic programming argument, similar to the preceding chapter, we propose a PDE characterization for the desired initial condition sets. The proposed approach leads to a sequence of PDEs, for which the first one has a known boundary condition, while the boundary conditions of the subsequent ones are determined by the solutions to the preceding steps.

During the PhD, we also approached the motion planning objective from different perspective that allows us to extend the class of specifications to more sophisticated maneuvers comprising long (possibly infinite) sequences of actions (e.g., linear temporal logic [CES86]). For this purpose the proposed PDE approach may not be an efficient scheme for two reasons: first, one is required to numerically solve a certain PDE for each excursion in a recursive fashion; second, not every specification can be translated into a finite-time reachability framework. Motivated by that, we develop an approach based on the so-called *symbolic models* which constructs a two-way bridge from a continuous (infinite) system to a discrete (finite) approximation such that a controller designed for the approximation can be refined to a controller for the original systems with an a priori ε-precision. This work is, however, not covered in this thesis, and we refer the interested readers to [ZMEM+13b, ZMEAL13] for further details.

1.1.2 Part II: Fault Detection for Large Scale Nonlinear Systems

The study in Part II is mainly motivated by the FDI problem with the prospect to devise a scalable methodology applicable to relatively high dimensional nonlinear systems.

A. An Optimization-Based Approach with Probabilistic Performance Index

In Chapter 4 we develop a novel approach to FDI which strikes a balance between analytical and computational tractability, and is applicable to relatively large dimensional nonlinear dynamics. For this purpose, we propose a design perspective that basically shifts the emphasis from the system dynamics to the family of disturbances that the system may encounter. Consider a general dynamical system as in Figure 1.1 with its inputs categorized into (i) unknown inputs d, (ii) fault signal f, and (iii) known inputs u.

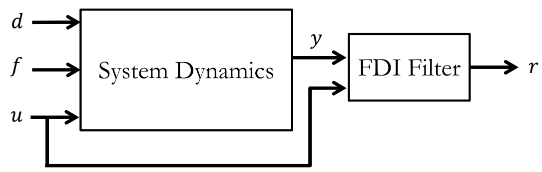

Figure 1.1: General configuration of the FDI filter

The FDI task is to design a filter fed by known signals (u and y) whose output, which is known as the residual and denoted by r, differentiates whether the measurements are a consequence of some accepted input disturbances d, or due to the fault signal f. Formally speaking, the residual may be viewed as the function $r(d, f)$, and the FDI design is ideally translated as the mapping requirements

$$d \mapsto r(d, 0) \equiv 0, \tag{1.1a}$$
$$f \mapsto r(d, f) \neq 0, \qquad \forall d \tag{1.1b}$$

where condition (1.1a) ensures that the residual of the filter, r, is not excited when the system is perturbed by normal disturbances d, while condition (1.1b) guarantees the filter sensitivity to the fault f in the presence of any disturbance d. In practice, however, it may be difficult, or even impossible, to satisfy condition (1.1a) exactly. An attempt to circumvent this issue is to consider the worst case scenario in the robust formulation

$$\text{RP}: \begin{cases} \min\limits_{\gamma} & \gamma \\ \text{s.t.} & \big\|r(d,0)\big\| \le \gamma, \qquad \forall d \in \mathcal{D} \\ & f \mapsto r(d,f) \neq 0, \quad \forall d \in \mathcal{D} \end{cases},$$

where $\mathcal{D}$ is set of normal disturbances, and the above minimization is running over a given class of FDI filters. Note that the parameter γ can be cast as the alarm threshold of the designed filter. In this work, assuming that some statistical information of the disturbance d is available, we relax the robust perspective by introducing probabilistic constraints in the following fashions:

$$\text{AP}:\begin{cases}\min\limits_{\gamma} & \gamma\\ \text{s.t.} & \mathbb{E}\big[J\big(\|r\|\big)\big]\le\gamma\\ & f\mapsto r(d,f)\neq 0,\quad \forall d\in\mathcal{D}\end{cases}\qquad \text{CP}:\begin{cases}\min\limits_{\gamma} & \gamma\\ \text{s.t.} & \mathbb{P}\big(\|r\|\le\gamma\big)\ge 1-\varepsilon\\ & f\mapsto r(d,f)\neq 0,\quad \forall d\in\mathcal{D}\end{cases},$$

where $\mathbb{P}$ is the probability measure on a prescribed probability space and $\mathbb{E}[\,\cdot\,]$ is meant with respect to $\mathbb{P}$. The disturbance d is viewed as a random variable taking values in $\mathcal{D}$ and $r := r(d,0)$. The function J in AP and $\varepsilon \in (0,1)$ in CP are both design parameters. In the sequel, invoking randomized techniques, we propose a scalable optimization-based scheme along with probabilistic certificates to tackle the above relaxed formulations. Finally, the performance of the proposed methodology is illustrated in an application to an emerging problem of cyber security in power transmission systems which led to an EU patent sponsored by ETH Zurich [MEVAL].

B. Performance Bound for Random Programs

The above formulation proposes the chance constrained perspective CP to relax the robust formulation RP of the original FDI problem. In this light, the alarm threshold γ obtained through the proposed optimization program may be violated with probability at most ε. In Chapter 5 our goal is to quantify how the false alarm rate would be improved by increasing the threshold γ. To this end, and in more general setting, we establish a theoretical bridge from the optimal values of the two optimization programs

$$\text{RCP}:\begin{cases}\min\limits_{x} & c^\intercal x\\ \text{s.t.} & f(x,d)\le 0,\quad \forall d\in\mathcal{D}\\ & x\in\mathbb{X}\end{cases},\qquad \text{CCP}_\varepsilon:\begin{cases}\min\limits_{x} & c^\intercal x\\ \text{s.t.} & \mathbb{P}[f(x,d)\le 0]\ge 1-\varepsilon\\ & x\in\mathbb{X}\end{cases},$$

where RCP and CCP_ε stand, respectively, for robust convex program and chance constrained program (cf. the formulation RP and CP in previous part), to a random counterpart so called scenario convex program

$$\text{SCP}:\begin{cases}\min\limits_{x} & c^\intercal x\\ \text{s.t.} & f(x,d_i)\le 0,\quad \forall i\in\{1,\cdots,N\}\\ & x\in\mathbb{X}\end{cases},$$

where $(d_i)_{i=1}^N$ are N independent and identically distributed (i.i.d.) samples drawn according to the probability measure $\mathbb{P}$ supported on $\mathcal{D}$, and $\mathbb{X}$ is a compact convex subset of $\mathbb{R}^n$. Along this way, we also extend the results to a certain class of non-convex problems that allows, for example, binary decision variables and non-convex set $\mathbb{X}$.

1.2 Publications

The work presented in this thesis mainly relies on previously published or submitted articles. The thesis only contains a subset of the research performed throughout my PhD studies and several projects are not featured here. The corresponding articles are listed below according to the related chapters along with relevant application-oriented projects.

1.2.1 Part I

- Chapter 2
  - **P. Mohajerin Esfahani**, D. Chatterjee, and J. Lygeros, "On Stochastic Reach-Avoid Problem and Set Characterization for Diffusions", submitted, Oct 2013, **arXiv: 1202.4375** [MECL13].
  - **P. Mohajerin Esfahani**, D. Chatterjee, and J. Lygeros, "On a Problem of Stochastic Reach-Avoid Set Characterization", in *IEEE Conference on Decision and Control (CDC)*, Orlando, Florida, USA, Dec 2011 [MECL11].
- Chapter 3
  - **P. Mohajerin Esfahani**, D. Chatterjee, and J. Lygeros, "Motion Planning for Continuous Time Stochastic Processes via Optimal Control", submitted to *IEEE Transaction of Automatic Control (TAC)*, Nov 2013, **arXiv: 1211.1138** [MECL12].
  - **P. Mohajerin Esfahani**, A. Milias-Argeitis, D. Chatterjee, "Analysis of Controlled Biological Switches via Stochastic Motion Planning", in *European Control Conference (ECC)*, Zurich, Switzerland, Jul 2013 [MEMAC13].
- Relevant Applications to Part I
  - T. Wood, **P. Mohajerin Esfahani**, and J. Lygeros, "Hybrid Modelling and Reachability on Autonomous RC-Cars", in *IFAC Conference on Analysis and Design of Hybrid Systems (ADHS)*, Eindhoven, Netherland, Jun 2012 [WMEL12].

1.2.2 Part II

- Chapter 4
  - **P. Mohajerin Esfahani** and John Lygeros, "A Tractable Fault Detection and Isolation Approach for Nonlinear Systems with Probabilistic Performance", submitted to *IEEE Transaction of Automatic Control (TAC)*, Feb 2013, Preprint [MEL13].
  - M. Vrakopoulou, **P. Mohajerin Esfahani**, K. Margellos, J. Lygeros, G. Andersson, "Cyber-Attacks in the Automatic Generation Control", submitted to *Cyber Physical Systems Approach to Smart Electric Power Grid*, Understanding Complex Systems, Khaitan, McCalley, and Liu Editors, Springer-Verlag Inc., 2014 [VMEM+].
  - **P. Mohajerin Esfahani**, M. Vrakopoulou, G. Andersson, J. Lygeros, "A Tractable Nonlinear Fault Detection and Isolation Technique with Application to the Cyber-Physical Security of Power Systems", in *IEEE Conference on Decision and Control (CDC)*, Maui, Hawaii, USA, Dec 2012 [MEVAL12].
  - **P. Mohajerin Esfahani**, M. Vrakopoulou, G. Andersson, J. Lygeros, "Intrusion Detection in Electric Power Networks", *Patent applied for PCT-EP-13002162*, filed on 22 July 2013 (Awarded for the best top 20 patents at ETH) [MEVAL].

- Chapter 5
  - **P. Mohajerin Esfahani**, T. Sutter, and J. Lygeros, "Performance Bounds for the Scenario Approach and an Extension to a Class of Non-convex Programs", to appear as a full paper in *IEEE Transaction of Automatic Control (TAC)*, 2014, **arXiv: 1307.0345** [MESL13].
- Relevant Applications to Part II
  - E.E. Tiniou, **P. Mohajerin Esfahani**, J. Lygeros, "Fault detection with discrete-time measurements: An application for the cyber security of power networks", in *IEEE Conference on Decision and Control (CDC)*, Florence, Italy, Dec 2013 [ETMEL13].
  - B. Svetozarevic, **P. Mohajerin Esfahani**, M. Kamgarpour, J. Lygeros, "A Robust Fault Detection and Isolation Filter for a Horizontal Axis Variable Speed Wind Turbine", in *American Control Conference (ACC)*, Washington, USA, Jun 2013 [SMEKL13].

1.2.3 Other Publications

- Journal Publications
  - M. Zamani, **P. Mohajerin Esfahani**, R. Majumdar, A. Abate, and J. Lygeros, "Symbolic Models for Stochastic Control Systems", to appear as a full paper in *IEEE Transaction of Automatic Control (TAC)*, 2014, **arXiv: 1302.3868** [ZMEM+13b].
  - R. Vujanic, **P. Mohajerin Esfahani**, P. Gulart, S. Marièthoz, and M. Morari, "Vanishing Duality Gap in Large Scale Mixed-Integer Optimization: a Solution Method with Power System Applications", submitted to Mathematical Programming, Nov 2013, Preprint [VMEG+].
- Conference Publications
  - T. Sutter, **P. Mohajerin Esfahani**, D. Sutter, J. Lygeros, "Capacity Approximation of Memoryless Channels with Countable Output Alphabets", in *IEEE International Symposium on Information Theory (ISIT)*, Honolulu, Hawaii, USA, Jul 2014 [SMESL14b].
  - D. Sutter, **P. Mohajerin Esfahani**, T. Sutter, J. Lygeros, "Efficient Approximation of Discrete Memoryless Channel Capacities", in *IEEE International Symposium on Information Theory (ISIT)*, Honolulu, Hawaii, USA, Jul 2014 [SMESL14a].
  - F. Heer, **P. Mohajerin Esfahani**, M. Kamgarpour, J. Lygeros, "Model based power optimisation of wind farms", in European Control Conference (ECC), Nov 2013 [HMEKL14].
  - M. Zamani, **P. Mohajerin Esfahani**, R. Majumdar, A. Abate, and J. Lygeros, "Bisimilar finite abstractions of stochastic control systems", in *IEEE Conference on Decision and Control (CDC)*, Florence, Italy, Dec 2013 [ZMEM+13a].

- M. Zamani, **P. Mohajerin Esfahani**, A. Abate, and J. Lygeros, "Symbolic models for stochastic control systems without stability assumptions", in *European Control Conference (ECC)*, Zurich, Switzerland, Jul 2013 [ZMEAL13].
- F. Oldewurtel, D. Sturzenegger, **P. Mohajerin Esfahani**, G. Andersson, M. Morari, J. Lygeros, "Adaptively Constrained Stochastic Model Predictive Control for Closed-Loop Constraint Satisfaction", in *American Control Conference (ACC)*, Washington, DC, USA, Jun 2013 [OSME+13].
- G. Andersson, **P. Mohajerin Esfahani**, M. Vrakopoulou, K. Margellos, J. Lygeros, A. Teixeira, G. Dan, H. Sandberg, K.H. Johansson, "Cyber-security of SCADA systems", in *Innovative Smart Grid Technologies (ISGT), IEEE PES*, Jan 2012 [AEV+12].
- **P. Mohajerin Esfahani**, M. Vrakopoulou, K. Margellos, J. Lygeros, G. Andersson, "A Robust Policy for Automatic Generation Control Cyber Attack in Two Area Power Network", in *IEEE Conference on Decision and Control (CDC)*, Atlanta, Georgia, USA, Dec 2010 [MEVM+11].
- **P. Mohajerin Esfahani**, M. Vrakopoulou, K. Margellos, J. Lygeros, G. Andersson, "Cyber Attack in a Two-Area Power System: Impact Identification using Reachability", in *American Control Conference (ACC)*, Baltimore, USA, Jun 2010 [MEVM+10].

Part I

Stochastic Motion Planning for Diffusions

CHAPTER 2

Stochastic Reach-Avoid Problem

In this chapter we develop a framework for formulating a class of stochastic reachability problems with state constraints as a stochastic optimal control problem. Previous approaches to solving these problems are either confined to the deterministic setting or address almost-sure stochastic notions. In contrast, we propose a new methodology to tackle probabilistic specifications that are less stringent than almost sure requirements. To this end, we first establish a connection between two stochastic reach-avoid problems and three classes of different stochastic optimal control problems involved with discontinuous payoff functions. Subsequently, we focus on solutions to one of the classes of stochastic optimal control problems—the exit-time problem, which solves both the reach-avoid problems mentioned above. We then derive a weak version of a dynamic programming principle (DPP) for the corresponding value function; in this direction our contribution compared to the existing literature is to allow for discontinuous payoff functions. Moreover, based on our DPP, we give an alternative characterization of the value function as a solution to a partial differential equation in the sense of discontinuous viscosity solutions, along with boundary conditions both in Dirichlet and viscosity senses. Theoretical justifications are discussed so as to employ off-the-shelf PDE solvers for numerical computations. Finally, we validate the performance of the proposed framework on the stochastic Zermelo navigation problem.

2.1 Introduction

Reachability is a fundamental concept in the study of dynamical systems, and in view of applications of this concept ranging from engineering, manufacturing, biology, and economics, to name but a few, has been studied extensively in the control theory literature. One particular problem that has turned out to be of fundamental importance in engineering is the so-called "reach-avoid" problem. In the deterministic setting this problem consists of determining the set of initial conditions for which one can find at least one control strategy to steer the system to a target set while avoiding certain obstacles. The set representing the solution to this problem is known as capture basin [Aub91]. This problem finds applications in air traffic management [LTS00] and security of power networks [MEVM+10]. A direct approach to compute the capture basin is formulated in the language of viability theory in [Car96, CQSP02]. Related problems involving pursuit-evasion games are solved in, e.g., [ALQ+02, GLQ06] employing tools from non-smooth analysis, for which computational tools are provided by [CQSP02].

An alternative and indirect approach to reachability involves using level set methods defined by value functions that characterize appropriate optimal control problems. Employing dynamic programming techniques for reachability and viability problems in the absence of state-constraints, one can in turn characterize these value functions by solutions to the standard Hamilton-Jacobi-Bellman (HJB) equations corresponding to these optimal control problems [Lyg04]. Numerical algorithms based on level set methods were developed by [OS88, Set99] and have been coded in efficient computational tools by [MT02, Mit05]. Extending the scope of this technique, the authors of [FG99, BFZ10, ML11] treat the case of time-independent state constraints and characterize the capture basin by means of a control problem whose value function is continuous.

In the stochastic setting, different probabilistic analogs of reachability problems have been studied extensively. In almost-sure setting, stochastic viability and controlled invariance are treated in [AD90, Aub91, APF00, BJ02]; see also the references therein. Methods involving stochastic contingent sets [AP98, APF00], viscosity solutions of second-order partial differential equations [BPQR98, BG99, BJ02], and derivatives of the distance function [DF01] were developed in this context. In [DF04] the authors developed an equivalence for the invariance problem between a stochastic differential equation and a certain deterministic control system. Toward similar objective, the authors in [ST02a] introduced a new class of problems, the so-called stochastic target problem, and characterized the solution via a dynamic programming approach. The differential properties of the reachable set were also studied based on the geometrical partial differential equation which is the analogue of the HJB equation [ST02b].

Although almost sure versions of reachability specifications are interesting in their own right, they may be too strict a concept in some applications. For example, in the safety assessment context, a common specification involves bounding the probability that undesirable events take place. In this regard, in the context of stochastic target problem, the authors of [BET10] recently extended the framework in [ST02a] to allow for unbounded control set so as to address less stringent than the almost sure requirement; see also the recent book [Tou13]. In this chapter, following the same problem but in an indirect approach, we develop a new framework for solving the following stochastic *reach-avoid* problem:

> RA: *Given an initial state $x \in \mathbb{R}^n$, a horizon $T > 0$, a number $p \in [0,1]$, and two disjoint sets $A, B \subset \mathbb{R}^n$, determine whether there exists a policy such that the controlled process reaches A prior to entering B within the interval $[0,T]$ with probability at least p.*

Observe that this is a significantly different problem compared to its almost-sure counterpart referred to above. It is of course immediate that the solution to the above problem is trivial if the initial state is either in B (in which case it is almost surely impossible) or in A (in which case there is nothing to do). However, for generic initial conditions in $\mathbb{R}^n \setminus (A \cup B)$, due to the inherent probabilistic nature of the dynamics, the problem of selecting a policy and determining the probability with which the controlled process reaches the set A prior to hitting B is non-trivial. In addition, we address the following slightly different reach-avoid problem compared to RA above, that requires the process be in the set A at time T:

$\widetilde{\mathrm{RA}}$: *Given an initial state $x \in \mathbb{R}^n$, a horizon $T > 0$, a number $p \in [0,1]$, and two disjoint sets $A, B \subset \mathbb{R}^n$, determine whether there exists a policy such that with probability at least p the controlled process resides in A at time T while avoiding B on the interval $[0,T]$.*

Throughout the chapter, we consider diffusion processes as the solution to a stochastic differential equation (SDE), and establish a connection from the Reach-Avoid problems to different classes of stochastic optimal control problems involving discontinuous payoff functions. One of the stochastic optimal control problems, which in fact addresses both the Reach-Avoid problems alluded above, is known as the *exit-time* problem [FS06, p. 6]. In this light, for the rest of the work we shall focus on the value function corresponding to the exit-time problem, and under some assumptions provide a dynamic programming principle (DPP) for it. The DPP is introduced in a weak sense in the spirit of [BT11], but in the context of exit-time framework; see also the recent work [BN12] with an extension addressing expectation constraints. This weak formulation avoids delicate restrictions related to a measurable selection and allows us to deal with discontinuous payoff functions, which to the best of our knowledge is new compared to the existing literature on exit-time problems. Based on the proposed DPP, we characterize the value function as the (discontinuous) viscosity solution of a partial differential equation (PDE) along with boundary conditions both in viscosity and Dirichlet (pointwise) senses. In this direction, we subsequently provide theoretical justifications so that the Reach-Avoid problem is amenable to numerical solutions by means of off-the-shelf PDE solvers.

Organization of the chapter: In Section 2.2 we formally introduce the stochastic reach-avoid problem RA above. In Section 2.3 we characterize the set of initial conditions that solve the problem RA above in terms of level sets of three different value functions. An identical connection is also established for a solution to the related reach-avoid problem $\widetilde{\mathrm{RA}}$ above. Focusing on the class of exit-time problems, in Section 2.4 we establish a dynamic programming principle (DPP), and characterize it as the solution of a PDE along with some boundary conditions. Section 2.5 presents results connecting those in Sections 2.3 and Section 2.4, and provides a solution to the stochastic reach-avoid problem in an "ε-conservative" sense. One may observe that this ε-precision can be made arbitrarily small. To illustrate the performance of our technique, the theoretical results developed in preceding sections are applied to solve the stochastic Zermelo navigation problem in Section 2.6. We summarize the chapter in 2.7, and for better readability some of the technical proofs are moved to Appendix 2.8.

Notation

Here is a partial notation list which will be also explained in more details later in the chapter:

- $\wedge$ (resp. $\vee$): minimum (resp. maximum) operator;
- $\overline{A}$ (resp. A°): closure (resp. interior) of the set A;
- $\mathbb{B}_r(x)$: open Euclidian ball centered at x and radius r;
- $\mathbb{C}_r(t,x)$: a cylinder with height and radius r, see (2.14);

- $\mathcal{U}_\tau$: set of $\mathbb{F}\tau$-progressively measurable maps into $\mathbb{U}$;
- $\mathcal{T}_{[\tau_1,\tau_2]}$: the collection of all $\mathbb{F}_{\tau_1}$-stopping times τ satisfying $\tau_1 \le \tau \le \tau_2$ $\mathbb{P}$-a.s.
- $(X_s^{t,x;\boldsymbol{u}})_{s\ge 0}$: stochastic process under the control policy $\boldsymbol{u}$ and assumption $X_s^{t,x;\boldsymbol{u}} := x$ for all $s \le t$;
- τ_A: first entry time to A, see Definition 2.3.1;
- V^* (resp. V_*): upper semicontinuous (resp. lower semicontinuous) envelope of function V;
- $\mathrm{USC}(\mathbb{S})$ $\big($resp. $\mathrm{LSC}(\mathbb{S})\big)$: collection of all upper semicontinuous (resp. lower semicontinuous) functions from $\mathbb{S}$ to $\mathbb{R}$;
- $\mathcal{L}^u$: Dynkin operator, see Definition 2.4.9.

2.2 Problem Statement

Consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ whose filtration $\mathbb{F} = (\mathcal{F}_s)_{s\ge 0}$ is generated by an n-dimensional Brownian motion $(W_s)_{s\ge 0}$ adapted to $\mathbb{F}$. Let the natural filtration of the Brownian motion $(W_s)_{s\ge 0}$ be enlarged by its right-continuous completion; — the usual conditions of completeness and right continuity, where $(W_s)_{s\ge 0}$ is a Brownian motion with respect to $\mathbb{F}$ [KS91, p. 48]. For every $t \ge 0$, we introduce an auxiliary subfiltration $\mathbb{F}_t := (\mathcal{F}_{t,s})_{s\ge 0}$, where $\mathcal{F}_{t,s}$ is the $\mathbb{P}$-completion of $\sigma\big(W_{r\vee t} - W_t, r \in [0,s]\big)$. Note that for $s \le t$, $\mathcal{F}_{t,s}$ is the trivial $\sigma-$algebra, and any $\mathcal{F}_{t,s}$-random variable is independent of $\mathcal{F}_t$. By definitions, it is obvious that $\mathcal{F}_{t,s} \subseteq \mathcal{F}_s$ with equality in case of $t = 0$.

Let $\mathbb{U} \subset \mathbb{R}^m$ be a control set, and $\mathcal{U}_t$ denote the set of $\mathbb{F}_t$-progressively measurable maps into $\mathbb{U}$.[1] We employ the shorthand $\mathcal{U}$ instead of $\mathcal{U}_0$ for the set of all $\mathbb{F}$-progressively measurable policies. We also denote by $\mathcal{T}$ the collection of all $\mathbb{F}$-stopping times. For $\tau_1, \tau_2 \in \mathcal{T}$ with $\tau_1 \le \tau_2$ $\mathbb{P}$-a.s. the subset $\mathcal{T}_{[\tau_1,\tau_2]}$ is the collection of all $\mathbb{F}_{\tau_1}$-stopping times τ such that $\tau_1 \le \tau \le \tau_2$ $\mathbb{P}$-a.s. Note that all $\mathbb{F}_\tau$-stopping times and $\mathbb{F}_\tau$-progressively measurable processes are independent of $\mathcal{F}_\tau$.

The basic object of our study concerns the $\mathbb{R}^n$-valued stochastic differential equation (SDE)

$$\mathrm{d}X_s = f(X_s, u_s)\,\mathrm{d}s + \sigma(X_s, u_s)\,\mathrm{d}W_s, \qquad X_0 = x, \quad s \ge 0, \tag{2.1}$$

where $f : \mathbb{R}^n \times \mathbb{U} \longrightarrow \mathbb{R}^n$ and $\sigma : \mathbb{R}^n \times \mathbb{U} \longrightarrow \mathbb{R}^{n\times d}$ are measurable maps, $(W_s)_{s\ge 0}$ is the above standard d-dimensional Brownian motion, and $\boldsymbol{u} := (u_s)_{s\ge 0} \in \mathcal{U}$.[2]

Assumption 2.2.1. *We stipulate that*

[1] Recall [KS91, p. 4] that a $\mathbb{U}$-valued process $(y_s)_{s\ge 0}$ is $\mathbb{F}_t$-progressively measurable if for each $T > 0$ the function $\Omega \times [0,T] \ni (\omega, s) \mapsto y(\omega, s) \in \mathbb{U}$ is measurable, where $\Omega \times [0,T]$ is equipped with $\mathcal{F}_{t,T} \otimes \mathfrak{B}[0,T]$, $\mathbb{U}$ is equipped with $\mathfrak{B}\mathbb{U}$, and $\mathfrak{B}S$ denotes the Borel σ-algebra on a topological space S.

[2] We slightly abuse notation and earlier used σ as a sigma algebra as well. However, it will be always clear from the context to which σ we refer.

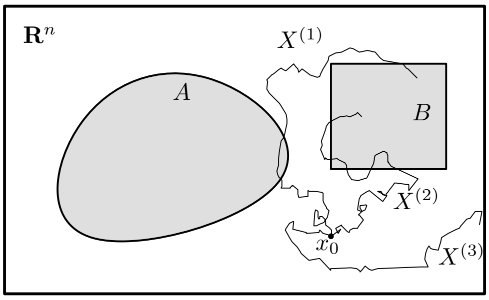

Figure 2.1: The trajectory $X^{(1)}$ hits A prior to B within time $[0, T]$, while $X^{(2)}$ and $X^{(3)}$ do not; all three start from initial state x_0.

a. $\mathbb{U} \subset \mathbb{R}^m$ *is compact;*

b. f *and* σ *are continuous and Lipschitz in their first argument uniformly with respect to the second.*

It is known [Bor05] that under Assumption 2.2.1 there exists a unique strong solution to the SDE (2.1). By definition of the filtration $\mathbb{F}$, we see that the control functions $\boldsymbol{u} \in \mathcal{U}$ satisfy the *non-anticipativity condition* [Bor05]—to wit, the increment $W_t - W_s$ is independent of the past history $\{W_r, \boldsymbol{u}_r \mid r \leq s\}$ of the Brownian motion and the control for every $s \in [0, t[$. (In other words, $\boldsymbol{u}$ does not anticipate the future increments of W). We let $(X_s^{t,x;\boldsymbol{u}})_{s \geq t}$ denote the unique strong solution of (2.1) starting from time t at the state x under the control policy $\boldsymbol{u}$. For future notational simplicity, we slightly modify the definition of $X_s^{t,x;\boldsymbol{u}}$, and extend it to the whole interval $[0, T]$ where $X_s^{t,x;\boldsymbol{u}} := x$ for all s in $[0, t]$. Measurability on $\mathbb{R}^n$ will always refer to Borel-measurability. In the sequel the complement of a set $S \subset \mathbb{R}^n$ is denoted by S^c.

Given an initial time t and the sets $A, B \subset \mathbb{R}^n$, we are interested in the set of initial conditions $x \in \mathbb{R}^n$ where there exists an admissible control strategy $\boldsymbol{u} \in \mathcal{U}$ such that with probability more than p the state trajectory $X_s^{t,x;\boldsymbol{u}}$ hits the set A before set B within the time horizon T. A pictorial representation of our problems is in Figure 2.1. Our main objective in this chapter is to propose a framework in order to characterize this set of initial condition, which is formally introduced as follows:

Definition 2.2.2 (Reach-Avoid within the interval $[0, T]$)**.**

$$\mathrm{RA}(t, p; A, B) := \Big\{ x \in \mathbb{R}^n \;\Big|\; \exists \boldsymbol{u} \in \mathcal{U} \;\; : \\ \mathbb{P}\Big(\exists s \in [t, T],\; X_s^{t,x;\boldsymbol{u}} \in A \;\; \textit{and} \;\; \forall r \in [t, s] \;\; X_r^{t,x;\boldsymbol{u}} \notin B \Big) > p \Big\}.$$

2.3 Connection to Optimal Control Problems

In this section we establish a connection between the stochastic reach-avoid problem RA and three different classes of stochastic optimal control problems. The following definition is one of the key elements in our framework.

Definition 2.3.1 (First entry time). *Given a control $\boldsymbol{u}$, the process $(X_s^{t,x;\boldsymbol{u}})_{s\geq t}$, and a measurable set $A \subset \mathbb{R}^n$, we introduce*[3] *the first entry time to A:*

$$\tau_A(t,x) = \inf\big\{s \geq t \;\big|\; X_s^{t,x;\boldsymbol{u}} \in A\big\}. \tag{2.2}$$

Remark 2.3.2. *Thanks to [EK86, Theorem 1.6, Chapter 2], $\tau_A(t,x)$ is an $\mathbb{F}_t$-stopping time. Moreover, due to the $\mathbb{P}$-a.s. continuity of the solution process, it can be easily deduced that given $\boldsymbol{u} \in \mathcal{U}$:*

$$\tau_{A\cup B} = \tau_A \wedge \tau_B, \tag{2.3a}$$

$$X_s^{t,x;\boldsymbol{u}} \in A \Longrightarrow \tau_A \leq s, \tag{2.3b}$$

$$A \text{ is closed} \Longrightarrow X_{\tau_A}^{t,x;\boldsymbol{u}} \in A. \tag{2.3c}$$

One can think of several different ways of characterizing probabilistic reach avoid sets, see e.g. [CCL11] and the references therein dealing with discrete-time problems. Motivated by these works, we consider value functions involving expectation of indicator functions of certain sets. Three alternative characterizations are considered and we show all three are equivalent. Consider the value functions $V_i : [0,T]\times\mathbb{R}^n \to [0,1]$ for $i = 1,2,3$, defined as follows:

$$V_1(t,x) := \sup_{\boldsymbol{u}\in\mathcal{U}} \mathbb{E}\Big[\mathbb{1}_A(X_{\bar\tau}^{t,x;\boldsymbol{u}})\Big] \qquad \text{where} \quad \bar\tau := \tau_{A\cup B}\wedge T, \tag{2.4a}$$

$$V_2(t,x) := \sup_{\boldsymbol{u}\in\mathcal{U}} \mathbb{E}\bigg[\sup_{s\in[t,T]}\Big\{\mathbb{1}_A(X_s^{t,x;\boldsymbol{u}}) \wedge \inf_{r\in[t,s]} \mathbb{1}_{B^c}(X_r^{t,x;\boldsymbol{u}})\Big\}\bigg], \tag{2.4b}$$

$$V_3(t,x) := \sup_{\boldsymbol{u}\in\mathcal{U}}\ \sup_{\tau\in\mathcal{T}_{[t,T]}}\ \inf_{\sigma\in\mathcal{T}_{[t,\tau]}} \mathbb{E}\Big[\mathbb{1}_A(X_\tau^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_\sigma^{t,x;\boldsymbol{u}})\Big]. \tag{2.4c}$$

Here $\tau_{A\cup B}$ is the hitting time introduced in Definition 2.3.1, and depends on the initial condition (t,x). For notational simplicity, we drop the initial condition in this section.

In the value function (2.4a) the process $X_\cdot^{t,x;\boldsymbol{u}}$ is controlled until the stopping time $\bar\tau$, by which instant the process either exits from the set $A\cup B$ or the terminal time T is reached. A sample $\omega\in\Omega$ is a "successful" path if the stopped process $X_{\bar\tau(\omega)}^{t,x;\boldsymbol{u}}(\omega)$ resides in A. This requirement is captured via the payoff function $\mathbb{1}_A(\cdot)$. In the value function (2.4b) there is no stopping time, and one may observe that the entire process $X_\cdot^{t,x;\boldsymbol{u}}$ is considered. Here the requirement of reaching the target set A before the avoid set B is taken into account by the supremum and infimum operations and payoff functions $\mathbb{1}_A$ and $\mathbb{1}_{B^c}$. In a fashion similar to (2.4a), the value function in (2.4c) involves some stopping time strategies. The stopping strategies are not fixed and the stochastic optimal control problem can be viewed as a game between two players with different authorities. Namely, the first player has both control $\boldsymbol{u}\in\mathcal{U}$ and stopping $\tau\in\mathcal{T}_{[t,T]}$ strategies whereas the second player has only a stopping strategy $\sigma\in\mathcal{T}_{[t,\tau]}$, which is dominated by the first player's stopping time τ; each player contributes through different maps to the payoff function.

The first result of this section, Proposition 2.3.4, asserts that $\mathbb{E}\big[\mathbb{1}_A(X_{\bar\tau}^{t,x;\boldsymbol{u}})\big] = \mathbb{P}\big(\tau_A < \tau_B,\ \tau_A \leq T\big)$. Since τ_A and τ_B are $\mathbb{F}$-stopping times, it then indicates the mapping $(t,x)\mapsto$

[3] By convention, $\inf\emptyset = \infty$.

$\mathbb{E}\big[\mathbb{1}_A(X_{\tilde{\tau}}^{t,x;\boldsymbol{u}})\big]$ is indeed well-defined. Furthermore, in Proposition 2.3.5 we shall establish equality of the three functions V_1, V_2, V_3 that will prove the other value functions are also well-defined.

Assumption 2.3.3. *We assume that the sets A and B are disjoint and closed.*

Proposition 2.3.4. *Consider the system (2.1), and let $A, B \subset R^n$ be given. Under Assumptions 2.2.1 and 2.3.3 we have*

$$\mathrm{RA}(t,p;A,B) = \{x \in \mathbb{R}^n \mid V_1(t,x) > p\},$$

where the set RA *is the set defined in Definition 2.2.2 and V_1 is the value function defined in (2.4a).*

Proof. See Appendix 2.8.1 □

Proposition 2.3.5. *Consider the system (2.1), and let $A, B \subset R^n$ be given. Under Assumptions 2.2.1 and 2.3.3 we have*

$$V_1 = V_2 = V_3 \qquad \textit{on} \quad [0,T] \times \mathbb{R}^n,$$

where the value functions V_1, V_2, V_3 are as defined in (2.4).

Proof. See Appendix 2.8.1. □

We now introduce the reach-avoid problem $\widetilde{\mathrm{RA}}$ mentioned in Section 2.1. Let us recall that the reach-avoid problem in Definition 2.2.2 poses a reach objective while avoiding barriers within the interval $[t,T]$. A similar problem may be formulated as being in the target set at time T while avoiding barriers over the period $[t,T]$. Namely, we define the set $\widetilde{\mathrm{RA}}(t,p;A,B)$ as the set of all initial conditions for which there exists an admissible control strategy $\boldsymbol{u} \in \mathcal{U}$ such that with probability more than p, $X_T^{t,x;\boldsymbol{u}}$ belongs to A and the process avoids the set B over the interval $[t,T]$.

Definition 2.3.6 (Reach-Avoid at the terminal time T)**.**

$$\widetilde{\mathrm{RA}}(t,p;A,B) := \Big\{x \in \mathbb{R}^n \;\Big|\; \exists \boldsymbol{u} \in \mathcal{U} \;: \\ \mathbb{P}\Big(X_T^{t,x;\boldsymbol{u}} \in A \textit{ and } \forall r \in [t,T] \; X_r^{t,x;\boldsymbol{u}} \notin B\Big) > p\Big\}.$$

One can establish a connection between the new reach-avoid problem in Definition 2.3.6 and different classes of stochastic optimal control problems along lines similar to Propositions 2.3.4 and 2.3.5. To this end, let us define the value functions $\widetilde{V}_i : [0,T] \times \mathbb{R}^n \to [0,1]$ for $i = 1,2,3$, as follows:

$$\widetilde{V}_1(t,x) := \sup_{\boldsymbol{u}\in\mathcal{U}} \mathbb{E}\Big[\mathbb{1}_A(X_{\tilde{\tau}}^{t,x;\boldsymbol{u}})\Big] \qquad \text{where} \quad \tilde{\tau} := \tau_B \wedge T, \tag{2.5a}$$

$$\widetilde{V}_2(t,x) := \sup_{\boldsymbol{u}\in\mathcal{U}} \mathbb{E}\Big[\mathbb{1}_A(X_T^{t,x;\boldsymbol{u}}) \wedge \inf_{r\in[t,T]} \mathbb{1}_{B^c}(X_r^{t,x;\boldsymbol{u}})\Big], \tag{2.5b}$$

$$\widetilde{V}_3(t,x) := \sup_{\boldsymbol{u}\in\mathcal{U}} \inf_{\sigma\in\mathcal{T}_{[t,T]}} \mathbb{E}\Big[\mathbb{1}_A(X_T^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_\sigma^{t,x;\boldsymbol{u}})\Big]. \tag{2.5c}$$

In our subsequent work, the measurability of the functions V_i and $\widetilde{V}_i$ will turn out to be irrelevant; see Remark 2.4.8 for details. We state the following proposition concerning assertions identical to those of Propositions 2.3.4 and 2.3.5 for the reach-avoid problem of Definition 2.3.6.

Proposition 2.3.7. *Consider the system (2.1), and let $A, B \subset R^n$ be given. If the set B is closed, then under Assumption 2.2.1 we have $\widetilde{\mathrm{RA}}(t,p;A,B) = \{x \in \mathbb{R}^n \mid \widetilde{V}_1(t,x) > p\}$, where the set $\widetilde{\mathrm{RA}}$ is the set defined in Definition 2.3.6. Moreover, we have $\widetilde{V}_1 = \widetilde{V}_2 = \widetilde{V}_3$ on $[0,T]\times\mathbb{R}^n$ where the value functions $\widetilde{V}_1, \widetilde{V}_2, \widetilde{V}_3$ are as defined in (2.5).*

Proof. The proof follows effectively the same arguments as in the proofs of Propositions 2.3.4 and 2.3.5. □

2.4 Alternative Characterization of the Exit-Time Problem

The stochastic control problems introduced in (2.4a) and (2.5a) are well-known as the exit-time problem [FS06, p. 6]. Note that in light of Propositions 2.3.4 and 2.3.7, both problems in Definitions 2.2.2 and 2.3.6 can alternatively be characterized in the framework of exit-time problems, see (2.4a) and (2.5a), respectively. Motivated by this, in this section we present an alternative characterization of the exit-time problem based on solutions to certain partial differential equations. To this end, we generalize the value functions to

$$V(t,x) := \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\ell\big(X^{t,x;\boldsymbol{u}}_{\bar\tau(t,x)}\big)\Big], \qquad \bar\tau(t,x) := \tau_O(t,x)\wedge T, \tag{2.6}$$

with

$$\ell : \mathbb{R}^n \to \mathbb{R} \tag{2.7}$$

a given bounded measurable function, and O a measurable set. Note that τ_O is the stopping time defined in Definition 2.3.1 that in case of value function (2.4a) can be considered as $O = A \cup B$.

Hereafter we shall restrict our control processes to $\mathcal{U}_t$, the set $\mathcal{U}_t$ denotes the collection of all $\mathbb{F}_t$-progressively measurable processes $\boldsymbol{u} \in \mathcal{U}$. We will show that the function V in (2.6) is well-defined, Fact 2.4.2. In view of independence of the increments of Brownian motion, the restriction of control processes to $\mathcal{U}_t$ is not restrictive, and one can show that the value function in (2.6) remains the same if $\mathcal{U}_t$ is replaced by $\mathcal{U}$; see, for instance, [Kry09, Theorem 3.1.7, p. 132] and [BT11, Remark 5.2].

Our objective is to characterize the value function (2.6) as a (discontinuous) viscosity solution of a suitable Hamilton-Jacobi-Bellman equation. We introduce the set $\mathbb{S} := [0,T]\times\mathbb{R}^n$ and define the lower and upper semicontinuous envelopes of function $V : \mathbb{S} \to \mathbb{R}$:

$$V_*(t,x) := \liminf_{(t',x')\to(t,x)} V(t',x') \qquad V^*(t,x) := \limsup_{(t',x')\to(t,x)} V(t',x')$$

and also denote by $\mathrm{USC}(\mathbb{S})$ and $\mathrm{LSC}(\mathbb{S})$ the collection of all upper-semicontinuous and lower-semicontinuous functions from $\mathbb{S}$ to $\mathbb{R}$ respectively. Note that, by definition, $V_* \in \mathrm{LSC}(\mathbb{S})$ and $V^* \in \mathrm{USC}(\mathbb{S})$.

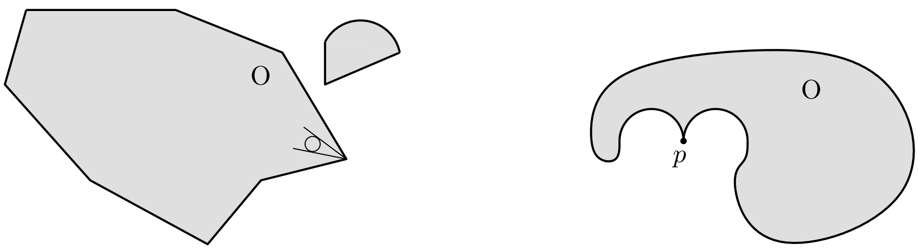

(a) Interior cone condition holds at every point of the boundary.

(b) Interior cone condition fails at the point p—the only possible interior cone at p is a line.

Figure 2.2: Interior cone condition of the boundary.

2.4.1 Assumptions and Preliminaries

Assumption 2.4.1. *In addition to Assumption 2.2.1, we stipulate the following:*

a. *(Non-degeneracy) The controlled processes are uniformly non-degenerate, i.e., there exists $\delta > 0$ such that for all $x \in \mathbb{R}^n$ and $u \in \mathbb{U}$, $\sigma(x,u)\sigma^\intercal(x,u) > \delta I$ where $\sigma(x,u)$ is the diffusion term in SDE (2.1).*

b. *(Interior Cone Condition) There are positive constants h, r, and an $\mathbb{R}^n$-value bounded map $\eta : \overline{O} \to \mathbb{R}^n$ satisfying*

$$\mathbb{B}_{rt}\big(x + \eta(x)t\big) \subset O \qquad \textit{for all } x \in \overline{O} \textit{ and } t \in (0, h]$$

where $\mathbb{B}_r(x)$ denotes an open ball centered at x and radius r, and $\overline{O}$ stands for the closure of the set O.

c. *(Lower Semicontinuity) The function ℓ defined in (2.7) is lower semicontinuous.*

Note that if the set A in Section 2.3 is open, then $\ell(\,\cdot\,) = \mathbb{1}_A(\,\cdot\,)$ satisfies Assumption 2.4.1.c. The interior cone condition in Assumption 2.4.1.b. concerns shapes of the set O; figure 2.2 illustrates two typical scenarios.

Fact 2.4.2 (Measurability)**.** *Consider the system (2.1), and suppose that Assumption 2.2.1 holds. Fix $(t, x, \boldsymbol{u}) \in \mathbb{S} \times \mathcal{U}$ and take an $\mathbb{F}$-stopping time $\theta : \Omega \to [0, T]$. For every measurable function $f : \mathbb{R}^n \to \mathbb{R}$, the function*

$$\Omega \ni \omega \mapsto g(\omega) := f\big(X^{t,x;\boldsymbol{u}}_{\theta(\omega)}(\omega)\big) \in \mathbb{R}$$

is $\mathcal{F}$-measurable (Recall that $\big(X^{t,x;\boldsymbol{u}}_s\big)_{s \geq t}$ is the unique strong solution of (2.1)).

Let us define the function $J : \mathbb{S} \times \mathcal{U} \to \mathbb{R}$:

$$J\big(t, x, \boldsymbol{u}\big) := \mathbb{E}\Big[\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big)\Big], \qquad \bar{\tau}(t,x) := \tau_O(t,x) \wedge T. \tag{2.8}$$

In the following proposition, we establish continuity of $\bar{\tau}(t,x)$ and lower semicontinuity of $J(t, x, \boldsymbol{u})$ with respect to (t, x).

Proposition 2.4.3. *Consider the system (2.1), and suppose that Assumptions 2.2.1 and 2.4.1 hold. Then for any strategy* $\boldsymbol{u} \in \mathcal{U}$ *and* $(t_0, x_0) \in \mathbb{S}$*,* $\mathbb{P}$*-a.s. the function* $(t, x) \mapsto \bar{\tau}(t, x)$ *is continuous at* (t_0, x_0)*. Moreover, the function* $(t, x) \mapsto J\big(t, x, \boldsymbol{u}\big)$ *defined in* (5.19) *is uniformly bounded and lower semicontinuous:*

$$J\big(t, x, \boldsymbol{u}\big) \leq \liminf_{(t',x')\to(t,x)} J\big(t', x', \boldsymbol{u}\big).$$

Proof. See Appendix 2.8.2. ☐

Remark 2.4.4 (Measurability)**.** *As a consequence of Fact 2.4.2 and Proposition 2.4.3, one can observe that for fixed* $(t, x, \boldsymbol{u}) \in \mathbb{S} \times \mathbb{U}$ *the function*

$$\Omega \ni \omega \mapsto J\big(\theta(\omega), X^{t,x;\boldsymbol{u}}_{\theta(\omega)}(\omega), \boldsymbol{u}\big) \in \mathbb{R}$$

is $\mathcal{F}$*-measurable.*

Fact 2.4.5 (Stability under Concatenation)**.** *For every* $\boldsymbol{u}$ *and* $\boldsymbol{v}$ *in* $\mathcal{U}_t$*, and* $\theta \in \mathcal{T}_{[t,T]}$

$$\mathbb{1}_{[t,\theta]}\boldsymbol{u} + \mathbb{1}_{]\theta,T]}\boldsymbol{v} \in \mathcal{U}_t.$$

Due to the definition of admissibility, the control process $\boldsymbol{u} := (u_s)_{s\geq 0} \in \mathcal{U}_t$ at time $s \geq 0$ can be viewed as a measurable mapping $(W_{r\vee t} - W_t)_{[0,s]} \mapsto u_s \in \mathbb{U}$, where $(W_s)_{s\geq 0}$ is the d-dimensional Brownian motion in (2.1); see [KS91, Def. 1.11, p. 4] for the details. Given $\theta \in \mathcal{T}_{[t,T]}$ and $u \in \mathcal{U}_t$, for each $\omega \in \Omega$ and the Brownian path up to the stopping time θ, i.e. $(W_r)_{r\in[0,\theta(\omega)]}$, we define the *random* policy $\boldsymbol{u}_\theta \in \mathcal{U}_{\theta(\omega)}$ as

$$\big(W_{\cdot\,\vee\theta(\omega)} - W_{\theta(\omega)}\big) \mapsto \boldsymbol{u}\big(W_{\cdot\,\wedge\theta(\omega)} + W_{\cdot\,\vee\theta(\omega)} - W_{\theta(\omega)}\big) =: \boldsymbol{u}_\theta. \tag{2.9}$$

Notice that $W_{\cdot} \equiv W_{\cdot\wedge\theta(\omega)} + W_{\cdot\vee\theta(\omega)} - W_{\theta(\omega)}$. Thus, the randomness of $\boldsymbol{u}_\theta$ is referred to the term $W_{\cdot\wedge\theta(\omega)}$.

Lemma 2.4.6 (Strong Markov Property)**.** *Consider the system (2.1) satisfying Assumptions 2.2.1. Then, for a stopping time* $\theta \in \mathcal{T}_{[t,T]}$ *and an admissible control* $\boldsymbol{u} \in \mathcal{U}_t$*, we have*

$$\mathbb{E}\Big[\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) \;\Big|\; \mathcal{F}_\theta\Big] = \mathbb{1}_{\{\bar{\tau}(t,x)<\theta\}}\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) + \mathbb{1}_{\{\bar{\tau}(t,x)\geq\theta\}}J\big(\theta, X^{t,x;\boldsymbol{u}}_\theta, \boldsymbol{u}_\theta\big) \qquad \mathbb{P}\text{-a.s.},$$

where $\boldsymbol{u}_\theta$ *is the random policy in the sense of* (2.9)*.*

Proof. See Appendix 2.8.2. ☐

2.4.2 Dynamic Programming Principle

The following Theorem provides a dynamic programming principle (DPP) for the exit time problem introduced in (2.6).

Theorem 2.4.7 (Dynamic Programming Principle)**.** *Consider the system* (2.1)*, and suppose that Assumptions 2.2.1 and 2.4.1 hold. Then for every* $(t,x) \in \mathbb{S}$ *and family of stopping times* $\{\theta^{\boldsymbol{u}}, \boldsymbol{u} \in \mathcal{U}_t\} \subset \mathcal{T}_{[t,T]}$*, we have*

$$V(t,x) \leq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)\leq\theta^{\boldsymbol{u}}\}}\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) + \mathbb{1}_{\{\bar{\tau}(t,x)>\theta^{\boldsymbol{u}}\}}V^*\big(\theta^{\boldsymbol{u}}, X^{t,x;\boldsymbol{u}}_{\theta^{\boldsymbol{u}}}\big)\Big], \tag{2.10}$$

and

$$V(t,x) \geq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)\leq\theta^{\boldsymbol{u}}\}}\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) + \mathbb{1}_{\{\bar{\tau}(t,x)>\theta^{\boldsymbol{u}}\}}V_*\big(\theta^{\boldsymbol{u}}, X^{t,x;\boldsymbol{u}}_{\theta^{\boldsymbol{u}}}\big)\Big], \tag{2.11}$$

where V *is the value function defined in* (2.6)*.*

Proof. The proof is based on techniques developed in [BT11]. We first assemble an appropriate covering for the set $\mathbb{S}$, and use this covering to construct a control strategy which satisfies the required conditions within ε precision, $\varepsilon > 0$ being pre-assigned and arbitrary. For notational simplicity, in the following we set $\theta := \theta^{\boldsymbol{u}}$.

Proof of (2.10)**:** In view of Strong Markov Property, Lemma 2.4.6, and the tower property of conditional expectation [Kal97, Theorem 5.1], for any $(t,x) \in \mathbb{S}$ we have

$$\begin{aligned}
\mathbb{E}\Big[\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big)\Big] &= \mathbb{E}\Big[\mathbb{E}\Big[\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) \,\Big|\, \mathcal{F}_\theta\Big]\;\Big] \\
&= \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)\leq\theta\}}\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) + \mathbb{1}_{\{\bar{\tau}(t,x)>\theta\}}J\big(\theta, X^{t,x;\boldsymbol{u}}_{\theta}, \boldsymbol{u}_\theta\big)\Big] \\
&\leq \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)\leq\theta\}}\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) + \mathbb{1}_{\{\bar{\tau}(t,x)>\theta\}}V^*\big(\theta, X^{t,x;\boldsymbol{u}}_{\theta}\big)\Big],
\end{aligned}$$

where $\boldsymbol{u}_\theta$ is the random control as introduced in (2.9). Note that the last inequality follows from the fact that $\boldsymbol{\theta} \in \mathcal{U}_{\theta(\omega)}$ for each $\omega \in \Omega$. Now taking supremum over all admissible controls $\boldsymbol{u} \in \mathcal{U}_t$ leads to the desired dynamic programming inequality (2.10).

Proof of (2.11)**:** Suppose $\phi : \mathbb{S} \to \mathbb{R}$ is uniformly bounded such that

$$\phi \in \mathrm{USC}(\mathbb{S}) \quad \text{and} \quad \phi \leq V_* \qquad \text{on} \quad \mathbb{S}. \tag{2.12}$$

According to (2.12) and Fact 2.4.3, given $\varepsilon > 0$, for all $(t_0, x_0) \in \mathbb{S}$ and $\boldsymbol{u} \in \mathcal{U}_{t_0}$ there exists $r_\varepsilon > 0$ such that

$$\begin{aligned}
\phi(t,x) - \varepsilon \leq \phi(t_0,x_0) \leq V_*(t_0,x_0), &\quad \forall (t,x) \in \mathbb{C}_{r_\varepsilon}(t_0,x_0) \cap \mathbb{S}, \\
J\big(t_0,x_0,\boldsymbol{u}\big) \leq J\big(t,x,\boldsymbol{u}\big) + \varepsilon, &\quad \forall (t,x) \in \mathbb{C}_{r_\varepsilon}(t_0,x_0) \cap \mathbb{S},
\end{aligned} \tag{2.13}$$

where $\mathbb{C}_r(t,x)$ is a cylinder defined as:

$$\mathbb{C}_r(t,x) := \{(s,y) \in \mathbb{R} \times \mathbb{R}^n \mid s \in]t-r, t] \;,\; \|x-y\| < r\}. \tag{2.14}$$

Moreover, by definition of (5.19) and (2.6), given $\varepsilon > 0$ and $(t_0,x_0) \in \mathbb{S}$ there exists $\boldsymbol{u}^{t_0,x_0}_\varepsilon \in \mathcal{U}_{t_0}$ such that

$$V_*(t_0,x_0) \leq V(t_0,x_0) \leq J\big(t_0,x_0,\boldsymbol{u}^{t_0,x_0}_\varepsilon\big) + \varepsilon.$$

By the above inequality and (2.13), one can conclude that given $\varepsilon > 0$, for all $(t_0, x_0) \in \mathbb{S}$ there exist $\boldsymbol{u}_\varepsilon^{t_0,x_0} \in \mathcal{U}_{t_0}$ and $r_\varepsilon := r_\varepsilon(t_0, x_0) > 0$ such that

$$\phi(t,x) - 3\varepsilon \leq J\big(t, x, \boldsymbol{u}_\varepsilon^{t_0,x_0}\big) \quad \forall (t,x) \in \mathbb{C}_{r_\varepsilon}(t_0, x_0) \cap \mathbb{S}. \tag{2.15}$$

Therefore, given $\varepsilon > 0$, the family of cylinders $\{\mathbb{C}_{r_\varepsilon}(t,x) \ : \ (t,x) \in \mathbb{S}, \quad r_\varepsilon(t_0,x_0) > 0\}$ forms an open covering of $[0,T[\times\mathbb{R}^n$. By the *Lindelöf* covering Theorem [Dug66, Theorem 6.3 Chapter VIII], there exists a countable sequence $(t_i, x_i, r_i)_{i\in\mathbb{N}}$ of elements of $\mathbb{S} \times \mathbb{R}^+$ such that

$$[0,T[\times\mathbb{R}^n \subset \bigcup_{i\in\mathbb{N}} \mathbb{C}_{r_i}(t_i, x_i).$$

Note that the implication of (2.10) simply holds for $(t,x) \in \{T\} \times \mathbb{R}^n$. Let us construct a sequence $(\mathbb{C}^i)_{i\in\mathbb{N}_0}$ as

$$\mathbb{C}^0 := \{T\} \times \mathbb{R}^n, \qquad \mathbb{C}^i := \mathbb{C}_{r_i}(t_i, x_i) \setminus \bigcup_{j\leq i-1} \mathbb{C}^j.$$

By definition $\mathbb{C}^i$ are pairwise disjoint and $\mathbb{S} \subset \bigcup_{i\in\mathbb{N}_0} \mathbb{C}^i$. Furthermore, $\mathbb{P}$- a.s., $(\theta, X_\theta^{t,x;\boldsymbol{u}}) \in \bigcup_{i\in\mathbb{N}_0} \mathbb{C}^i$, and for all $i \in \mathbb{N}_0$ there exists $\boldsymbol{u}_\varepsilon^{t_i,x_i} \in \mathcal{U}_{t_i}$ such that

$$\phi(t,x) - 3\varepsilon \leq J\big(t, x, \boldsymbol{u}_\varepsilon^{t_i,x_i}\big), \qquad \forall (t,x) \in \mathbb{C}^i \cap \mathbb{S}. \tag{2.16}$$

To prove (2.11), let us fix $\boldsymbol{u} \in \mathcal{U}_t$ and $\theta \in \mathcal{T}_{[t,T]}$. Given $\varepsilon > 0$ we define

$$\boldsymbol{v}_\varepsilon := \mathbb{1}_{[t,\theta]}\boldsymbol{u} + \mathbb{1}_{]\theta,T]} \sum_{i\in\mathbb{N}_0} \mathbb{1}_{\mathbb{C}^i}(\theta, X_\theta^{t,x;\boldsymbol{u}})\boldsymbol{u}_\varepsilon^{t_i,x_i}. \tag{2.17}$$

Notice that by Fact 2.4.5, the set $\mathcal{U}_t$ is closed under countable concatenation operations, and consequently $\boldsymbol{v}_\varepsilon \in \mathcal{U}_t$. In view of Lemma 2.4.6 and (2.16), it can be deduced that, $\mathbb{P}$-a.e. on Ω under $\boldsymbol{v}_\varepsilon$ in (2.17),

$$\begin{aligned}
&\mathbb{E}\Big[\ell\big(X_{\bar\tau(t,x)}^{t,x;\boldsymbol{v}_\varepsilon}\big) \mid \mathcal{F}_\theta\Big] \\
&\quad = \mathbb{1}_{\{\bar\tau(t,x)\leq\theta\}}\ell\big(X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}\big) + \mathbb{1}_{\{\bar\tau(t,x)>\theta\}} J\Big(\theta, X_\theta^{t,x;\boldsymbol{u}}, \sum_{i\in\mathbb{N}_0} \mathbb{1}_{\mathbb{C}^i}(\theta, X_\theta^{t,x;\boldsymbol{u}})\boldsymbol{u}_\varepsilon^{t_i,x_i}\Big) \\
&\quad = \mathbb{1}_{\{\bar\tau(t,x)\leq\theta\}}\ell\big(X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}\big) + \mathbb{1}_{\{\bar\tau(t,x)>\theta\}} \sum_{i\in\mathbb{N}_0} J\big(\theta, X_\theta^{t,x;\boldsymbol{u}}, \boldsymbol{u}_\varepsilon^{t_i,x_i}\big)\mathbb{1}_{\mathbb{C}^i}\big(\theta, X_\theta^{t,x;\boldsymbol{u}}\big) \\
&\quad \geq \mathbb{1}_{\{\bar\tau(t,x)\leq\theta\}}\ell\big(X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}\big) + \mathbb{1}_{\{\bar\tau(t,x)>\theta\}} \sum_{i\in\mathbb{N}_0} \Big(\phi\big(\theta, X_\theta^{t,x;\boldsymbol{u}}\big) - 3\varepsilon\Big)\mathbb{1}_{\mathbb{C}^i}\big(\theta, X_\theta^{t,x;\boldsymbol{u}}\big) \\
&\quad = \mathbb{1}_{\{\bar\tau(t,x)\leq\theta\}}\ell\big(X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}\big) + \mathbb{1}_{\{\bar\tau(t,x)>\theta\}}\Big(\phi\big(\theta, X_\theta^{t,x;\boldsymbol{u}}\big) - 3\varepsilon\Big).
\end{aligned}$$

By the definition of V and the tower property of conditional expectations,

$$\begin{aligned}
V(t,x) \geq J(t,x,\boldsymbol{v}_\varepsilon) &= \mathbb{E}\Big[\mathbb{E}\Big[\ell\big(X_{\bar\tau(t,x)}^{t,x;\boldsymbol{v}_\varepsilon}\big) \mid \mathcal{F}_\theta\Big]\Big] \\
&\geq \mathbb{E}\Big[\mathbb{1}_{\{\bar\tau(t,x)\leq\theta\}}\ell\big(X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}\big) + \mathbb{1}_{\{\bar\tau(t,x)>\theta\}}\phi\big(\theta, X_\theta^{t,x;\boldsymbol{u}}\big)\Big] - 3\varepsilon\mathbb{E}\Big[\mathbb{1}_{\{\bar\tau(t,x)>\theta\}}\Big].
\end{aligned}$$

The arbitrariness of $\boldsymbol{u} \in \mathcal{U}_t$ and $\varepsilon > 0$ imply that

$$V(t,x) \geq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)\leq\theta\}}\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) + \phi\big(\theta, X^{t,x;\boldsymbol{u}}_{\theta}\big)\Big].$$

It suffices to find a sequence of continuous functions $(\phi_i)_{i\in\mathbb{N}}$ such that $\Phi_i \leq V_*$ on $\mathbb{S}$ and converges pointwise to V_*. The existence of such a sequence is guaranteed by [Ren99, Lemma 3.5]. Thus, by Fatou's lemma,

$$\begin{aligned} V(t,x) &\geq \liminf_{i\to\infty} \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)<\theta\}}\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) + \mathbb{1}_{\{\bar{\tau}(t,x)\geq\theta\}}\phi_i\big(\theta, X^{t,x;\boldsymbol{u}}_{\theta}\big)\Big] \\ &\geq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)<\theta\}}\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) + \mathbb{1}_{\{\bar{\tau}(t,x)\geq\theta\}}\liminf_{i\to\infty}\phi_i\big(\theta, X^{t,x;\boldsymbol{u}}_{\theta}\big)\Big] \\ &= \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)<\theta\}}\ell\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}(t,x)}\big) + \mathbb{1}_{\{\bar{\tau}(t,x)\geq\theta\}}V_*\big(\theta, X^{t,x;\boldsymbol{u}}_{\theta}\big)\Big]. \end{aligned}$$

□

Remark 2.4.8. *The dynamic programming principles in* (2.10) *and* (2.11) *are introduced in a weaker sense than the standard DPP for stochastic optimal control problems [FS06]. To wit, note that one does not have to verify the measurability of the value function V defined in (2.6) to apply our DPP.*

2.4.3 Dynamic Programming Equation

Our objective in this subsection is to demonstrate how the DPP derived in Subsection 2.4.2 characterizes the value function V as a (discontinuous) viscosity solution to an appropriate HJB equation; for the general theory of viscosity solutions we refer to [CIL92] and [FS06]. To complete the PDE characterization and provide numerical solutions for this PDE, one also needs appropriate boundary conditions which will be the objective of the next subsection.

Definition 2.4.9 (Dynkin Operator)**.** *Given $u \in \mathbb{U}$, we denote by $\mathcal{L}^u$ the Dynkin operator (also known as the infinitesimal generator) associated to the controlled diffusion* (2.1) *as*

$$\mathcal{L}^u\Phi(t,x) := \partial_t\Phi(t,x) + f(x,u).\partial_x\Phi(t,x) + \frac{1}{2}\mathrm{Tr}[\sigma(x,u)\sigma^\top(x,u)\partial_x^2\Phi(t,x)],$$

where Φ is a real-valued function smooth on the interior of $\mathbb{S}$, with $\partial_t\Phi$ and $\partial_x\Phi$ denoting the partial derivatives with respect to t and x respectively, and $\partial_x^2\Phi$ denoting the Hessian matrix with respect to x.

We refer to [Kal97, Theorem 17.23] for more details on the above differential operator.

Theorem 2.4.10 (Dynamic Programming Equation)**.** *Consider the system* (2.1)*, and suppose that Assumptions 2.2.1 and 2.4.1 hold. Then,*

- *the lower semicontinuous envelope of V introduced in* (2.6) *is a viscosity supersolution of*

$$-\sup_{u\in\mathbb{U}}\mathcal{L}^u V_*(t,x) \geq 0 \qquad \text{on} \quad [0,T[\times\overline{O}^c,$$

∘ *the upper semicontinuous envelope of V is a viscosity subsolution of*

$$-\sup_{u\in\mathbb{U}} \mathcal{L}^u V^*(t,x) \le 0 \qquad \textit{on} \quad [0,T[\times\overline{O}^c,$$

Proof. We first prove the supersolution part:

Supersolution: For the sake of contradiction, assume that there exists $(t_0,x_0) \in [0,T[\times\overline{O}^c$ and a smooth function $\phi : \mathbb{S} \to \mathbb{R}$ satisfying

$$\min_{(t,x)\in\mathbb{S}} \big(V_* - \phi\big)(t,x) = \big(V_* - \phi\big)(t_0,x_0) = 0$$

such that for some $\delta > 0$

$$-\sup_{u\in\mathbb{U}} \mathcal{L}^u \phi(t_0,x_0) < -2\delta$$

Notice that, without loss of generality, one can assume that (t_0,x_0) is the strict minimizer of $V_* - \phi$ [FS06, Lemma II 6.1, p. 87]. Since ϕ is smooth, the map $(t,x) \mapsto \mathcal{L}^u\phi(t,x)$ is continuous. Therefore, there exist $u \in \mathbb{U}$ and $r > 0$ such that $\mathbb{B}_r(t_0,x_0) \subset [0,T) \times \overline{O}^c$ and

$$-\mathcal{L}^u\phi(t,x) < -\delta \quad \forall (t,x) \in \mathbb{B}_r(t_0,x_0). \tag{2.18}$$

Let us define the stopping time $\theta(t,x) \in \mathcal{T}_{[t,T]}$

$$\theta(t,x) = \inf\{s \ge t \ : \ (s, X_s^{t,x;u}) \notin \mathbb{B}_r(t_0,x_0)\}, \tag{2.19}$$

where $(t,x) \in \mathbb{B}_r(t_0,x_0)$. Note that by continuity of solutions to (2.1), $t < \theta(t,x) < T$ $\mathbb{P}$- a.s. for all $(t,x) \in \mathbb{B}_r(t_0,x_0)$. Moreover, selecting $r > 0$ sufficiently small so that $\theta(t,x) < \tau_O$, we have

$$\theta(t,x) < \tau_O \wedge T = \bar{\tau}(t,x) \quad \mathbb{P}\text{- a.s.} \quad \forall (t,x) \in \mathbb{B}_r(t_0,x_0) \tag{2.20}$$

Applying Itô's formula and using (2.18), we see that for all $(t,x) \in \mathbb{B}_r(t_0,x_0)$,

$$\begin{aligned}
\phi(t,x) &= \mathbb{E}\bigg[\phi\big(\theta(t,x), X_{\theta(t,x)}^{t,x;u}\big) + \int_t^{\theta(t,x)} -\mathcal{L}^u\phi\big(s, X_s^{t,x;u}\big)ds\bigg] \\
&\le \mathbb{E}\Big[\phi\big(\theta(t,x), X_{\theta(t,x)}^{t,x;u}\big)\Big] - \delta(\mathbb{E}[\theta(t,x)] - t) < \mathbb{E}\Big[\phi\big(\theta(t,x), X_{\theta(t,x)}^{t,x;u}\big)\Big].
\end{aligned}$$

Now it suffices to take a sequence $(t_n, x_n, V(t_n,x_n))_{n\in\mathbb{N}}$ converging to $(t_0, x_0, V_*(t_0,x_0))$ to see that

$$\phi(t_n,x_n) \to \phi(t_0,x_0) = V_*(t_0,x_0).$$

Therefore, for sufficiently large n we have

$$V(t_n,x_n) < \mathbb{E}\Big[\phi\big(\theta(t_n,x_n), X_{\theta(t_n,x_n)}^{t_n,x_n;u}\big)\Big] < \mathbb{E}\Big[V_*\big(\theta(t_n,x_n), X_{\theta(t_n,x_n)}^{t_n,x_n;u}\big)\Big],$$

which, in accordance with (2.20), can be expressed as

$$V(t_n,x_n) < \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t_n,x_n) < \theta(t_n,x_n)\}}\ell\big(X_{\bar{\tau}(t_n,x_n)}^{t_n,x_n;u}\big) + \mathbb{1}_{\{\bar{\tau}(t_n,x_n) \ge \theta(t_n,x_n)\}}V_*\big(\theta, X_{\theta(t_n,x_n)}^{t_n,x_n;u}\big)\Big].$$

This contradicts the DPP in (2.11).

Subsolution: The subsolution property is proved in a fashion similar to the supersolution part but with slightly more care. For the sake of contradiction, assume that there exists $(t_0, x_0) \in [0, T[\times \overline{O}^c$ and a smooth function $\phi : \mathbb{S} \to \mathbb{R}$ satisfying

$$\max_{(t,x)\in\mathbb{S}} \big(V^* - \phi\big)(t, x) = \big(V^* - \phi\big)(t_0, x_0) = 0$$

such that for some $\delta > 0$

$$-\sup_{u\in\mathbb{U}} \mathcal{L}^u \phi(t_0, x_0) > 2\delta.$$

By continuity of the mapping $(t, x, u) \mapsto \mathcal{L}^u \phi(t, x)$ and compactness of the control set $\mathbb{U}$, Assumption 2.2.1.a, there exists $r > 0$ such that for all $u \in \mathbb{U}$

$$-\mathcal{L}^u \phi(t, x) > \delta, \quad \forall (t, x) \in \mathbb{B}_r(t_0, x_0), \tag{2.21}$$

where $\mathbb{B}_r(t_0, x_0) \subset [0, T) \times \overline{O}^c$. Note as in the preceding part, (t_0, x_0) can be considered as the strict maximizer of $V^* - \phi$ that consequently implies that there exists $\gamma > 0$ such that

$$\big(V^* - \phi\big)(t, x) < -\gamma, \qquad \forall (t, x) \in \partial\mathbb{B}_r(t_0, x_0). \tag{2.22}$$

where $\partial\mathbb{B}_r(t_0, x_0)$ stands for the boundary of the ball $\mathbb{B}_r(t_0, x_0)$. Let $\theta(t, x) \in \mathcal{T}_{[t,T]}$ be the stopping time defined in (2.19); notice that θ may, of course, depend on the policy $\boldsymbol{u}$. Applying Itô's formula and using (2.21), one can observe that given $\boldsymbol{u} \in \mathcal{U}_t$,

$$\begin{aligned}
\phi(t, x) &= \mathbb{E}\left[\phi\big(\theta(t, x), X^{t,x;\boldsymbol{u}}_{\theta(t,x)}\big) + \int_t^{\theta(t,x)} -\mathcal{L}^{u_s} \phi\big(s, X^{t,x;\boldsymbol{u}}_s\big) ds\right] \\
&\geq \mathbb{E}\Big[\phi\big(\theta(t, x), X^{t,x;\boldsymbol{u}}_{\theta(t,x)}\big)\Big] + \delta(\mathbb{E}[\theta(t, x)] - t) > \mathbb{E}\Big[\phi\big(\theta(t, x), X^{t,x;\boldsymbol{u}}_{\theta(t,x)}\big)\Big].
\end{aligned}$$

Now it suffices to take a sequence $(t_n, x_n, V(t_n, x_n))_{n\in\mathbb{N}}$ converging to $(t_0, x_0, V^*(t_0, x_0))$ to see that

$$\phi(t_n, x_n) \to \phi(t_0, x_0) = V^*(t_0, x_0).$$

As argued in the supersolution part above, for sufficiently large n, for given $\boldsymbol{u} \in \mathcal{U}_t$,

$$V(t_n, x_n) > \mathbb{E}\Big[\phi\big(\theta(t_n, x_n), X^{t_n,x_n;\boldsymbol{u}}_{\theta(t_n,x_n)}\big)\Big] > \mathbb{E}\Big[V^*\big(\theta(t_n, x_n), X^{t_n,x_n;\boldsymbol{u}}_{\theta(t_n,x_n)}\big)\Big] + \gamma,$$

where the last inequality is deduced from the fact that $\big(\theta(t_n, x_n), X^{t_n,x_n;\boldsymbol{u}}_{\theta(t_n,x_n)}\big) \in \partial\mathbb{B}_r(t_0, x_0)$ together with (2.22). Thus, in view of (2.20), we arrive at

$$V(t_n, x_n) > \mathbb{E}\Big[\mathbb{1}_{\{\bar{\tau}(t,x)<\theta(t_n,x_n)\}} \ell\big(X^{t_n,x_n;\boldsymbol{u}}_{\bar{\tau}}\big) + \mathbb{1}_{\{\bar{\tau}(t,x)\geq\theta(t_n,x_n)\}} V^*\big(\theta, X^{t_n,x_n;\boldsymbol{u}}_{\theta(t_n,x_n)}\big)\Big] + \gamma.$$

This contradicts the DPP in (2.10) as γ is chosen uniformly with respect to $\boldsymbol{u} \in \mathcal{U}_t$. □

2.4.4 Boundary Conditions

Before proceeding with the main result of this subsection on boundary conditions, we need a preparatory lemma that provides a stronger assertion than Proposition 2.4.3.

Lemma 2.4.11. *Suppose that the conditions of Proposition 2.4.3 hold. Given a sequence of control policies* $(\boldsymbol{u_n})_{n\in\mathbb{N}} \subset \mathcal{U}$ *and initial conditions* $(t_n, x_n) \to (t, x)$*, we have*

$$\lim_{n\to\infty} \left\| X^{t,x;\boldsymbol{u}_n}_{\bar\tau(t,x)} - X^{t_n,x_n;\boldsymbol{u}_n}_{\bar\tau(t_n,x_n)} \right\| = 0, \qquad \mathbb{P}\text{-a.s.},$$

where the stopping time $\bar\tau$ *is introduced in* (2.6)*.*

Proof. See Appendix 2.8.2. □

The following Proposition provides boundary conditions for the value function V both in viscosity and Dirichlet (pointwise) senses:

Proposition 2.4.12 (Boundary Conditions)**.** *Suppose that the condition of Theorem 2.4.10 holds. Then the value function* V *introduced in* (2.6) *satisfies the following boundary value conditions:*

$$\textit{Dirichlet:} \qquad V(t,x) = \ell(x) \qquad \textit{on} \quad [0,T]\times\overline{O}\bigcup\{T\}\times\mathbb{R}^n \tag{2.23a}$$

$$\textit{Viscosity:} \qquad \begin{cases} \limsup\limits_{\substack{(\overline{O})^c\ni x'\to x\\ t'\uparrow t}} V(t',x') \le \ell^*(x) & \textit{on} \quad [0,T]\times\partial O\bigcup\{T\}\times\mathbb{R}^n \\ \liminf\limits_{\substack{(\overline{O})^c\ni x'\to x\\ t'\uparrow t}} V(t',x') \ge \ell(x) & \textit{on} \quad [0,T]\times\partial O\bigcup\{T\}\times\mathbb{R}^n \end{cases} \tag{2.23b}$$

Proof. In light of [RB98, Corollary 3.2, p. 65], Assumptions 2.4.1.a. and 2.4.1.b. ensure that

$$\bar\tau(t,x) = t, \quad \forall (t,x)\in[0,T]\times\overline{O}\cup\{T\}\times\mathbb{R}^n \qquad \mathbb{P}\text{-a.s.}$$

which readily implies the pointwise boundary condition (2.23a). To prove the discontinuous viscosity boundary condition (2.23b), we only show the first assertion; the second one follows from similar arguments. Let $(t,x)\in[0,T]\times\partial O\bigcup\{T\}\times\mathbb{R}^n$ and $(t_n,x_n)\to(t,x)$, where $t_n < T$ and $x\in(\overline{O})^c$. In the definition of V in (2.6), one can choose a sequence of policies that is increasing and attains the supremum value. This sequence, of course, depends on the initial condition. Thus, let us denote it via two indices $(\boldsymbol{u}_{n,j})_{j\in\mathbb{N}}$ as a sequence of policies corresponding to the initial condition (t_n,x_n) corresponding to the value $V(t_n,x_n)$. In this light, there exists a subsequence of $(\boldsymbol{u}_{n_j})_{j\in\mathbb{N}}$ such that

$$\begin{aligned} V^*(t,x) &= \lim_{n\to\infty} V(t_n,x_n) = \lim_{n\to\infty}\lim_{j\to\infty}\mathbb{E}\Big[\ell\big(X^{t_n,x_n;\boldsymbol{u}_{n,j}}_{\bar\tau(t_n,x_n)}\big)\Big] \\ &\le \lim_{j\to\infty}\mathbb{E}\Big[\ell\big(X^{t_j,x_j;\boldsymbol{u}_{n_j}}_{\bar\tau(t_j,x_j)}\big)\Big] \le \mathbb{E}\Big[\lim_{j\to\infty}\ell\big(X^{t_j,x_j;\boldsymbol{u}_{n_j}}_{\bar\tau(t_j,x_j)}\big)\Big] \le \ell^*(x) \end{aligned} \tag{2.24}$$

where the second and third inequality in (2.24) follow, respectively, from Fatou's lemma and the almost sure uniform continuity assertion in Lemma 2.4.11. Let us recall that $\bar\tau(t,x) = t$ and consequently $X^{t,x;\boldsymbol{u}_{n_j}}_{\bar\tau(t,x)} = x$. □

Proposition 2.4.12 provides boundary condition for V in both Dirichlet (pointwise) and viscosity senses. The Dirichlet boundary condition (2.23a) is the one usually employed to numerically compute the solution via PDE solvers, whereas the viscosity boundary condition (2.23b) is required for theoretical support of the numerical schemes and comparison results.

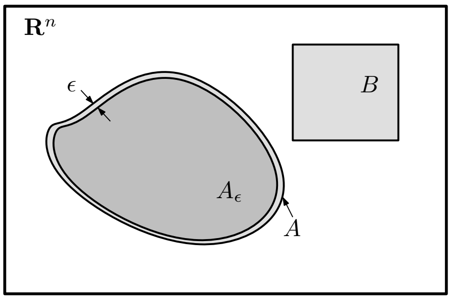

Figure 2.3: Construction of the sets A_ε from A as described in Section 2.5.

2.5 From the Reach-Avoid Problem to the PDE Characterization

In this section we draw a connection between the reach-avoid problem of Section 2.2 and the stochastic optimal control problems detailed in Section 2.3. To this end, note that on the one hand, an assumption on the sets A and B in the reach-avoid problem (Definition 2.2.2) within the time interval $[0, T]$ is that they are closed. On the other hand, our solution to the stochastic optimal control problem (defined in Section 2.2 and solved in Section 2.4) relies on lower semicontinuity of the payoff function ℓ in (2.6), see Assumption 2.4.1.c.

To achieve a reconciliation between the two sets of hypotheses, given sets A and B satisfying Assumption 2.3.3, we construct a smaller measurable set $A_\varepsilon \subset A^\circ$ such that $A_\varepsilon := \{x \in A \mid \mathrm{dist}(x, A^c) \geq \varepsilon\}$ [4] and A_ε satisfies Assumption 2.4.1.b. Note that this is always possible if $O := A \cup B$ satisfies Assumption 2.4.1.b.—indeed, simply take $\varepsilon < h/2$ to see this, where h is as defined in Assumption 2.4.1.b. Figure 2.3 depicts this case. To be precise, we define

$$V_\varepsilon(t, x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\Big[\ell_\varepsilon\big(X_{\tau_\varepsilon}^{t,x;\boldsymbol{u}}\big)\Big], \qquad \tau_\varepsilon := \tau_{A_\varepsilon \cup B} \wedge T, \tag{2.25}$$

where the function $\ell_\varepsilon : \mathbb{R}^n \to \mathbb{R}$ is defined as

$$\ell_\varepsilon(x) := \left(1 - \frac{\mathrm{dist}(x, A_\varepsilon)}{\varepsilon}\right) \vee 0.$$

The following Theorem asserts that the above technique affords an ε-conservative but precise way of characterizing the solution to the reach-avoid problem defined in Definition 2.2.2 in the framework of Section 2.4.

Theorem 2.5.1. *Consider the system* (2.1)*, and suppose that Assumptions 2.2.1, 2.3.3, 2.4.1.a. and 2.4.1.b. hold. Then, for all* $(t, x) \in [t, T[\times \mathbb{R}^n$ *and* $\varepsilon_1 \geq \varepsilon_2 > 0$*, we have* $V_{\varepsilon_2}(t, x) \geq V_{\varepsilon_1}(t, x)$*, and* $V(t, x) = \lim_{\varepsilon \downarrow 0} V_\varepsilon(t, x)$ *where the functions* V *and* V_ε *are defined as* (2.4a) *and* (2.25)*, respectively.*

Proof. See Appendix 2.8.3. □

[4] $\mathrm{dist}(x, A) := \inf_{y \in A} \|x - y\|$, where $\|\cdot\|$ stands for the Euclidean norm.

Remark 2.5.2. *Observe that for the problem of reachability at the time T, (as defined in Definition 2.3.6,) the above procedure is unnecessary if the set A is open; see the required conditions for Proposition 2.3.7.*

The following corollary addresses continuity of the value function V_ε in (2.25). It not only simplifies the PDE characterization developed in Subsection 2.4.3 from discontinuous to continuous regime, but also provides a theoretical justification for existing tools to numerically solve the corresponding PDE.

Corollary 2.5.3. *Consider the system in (2.1), and suppose that Assumptions 2.2.1 and 2.4.1 hold. Then, for any $\varepsilon > 0$ the value function $V_\varepsilon : \mathbb{S} \to [0,1]$ defined as in (2.25) is continuous. Furthermore, if $(A_\varepsilon \cup B)^c$ is bounded*[5] *then V_ε is the unique viscosity solution of*

$$\begin{cases} -\sup_{u\in\mathbb{U}} \mathcal{L}^u V_\varepsilon(t,x) = 0 & \text{in} \quad [0,T[\times(A_\varepsilon \cup B)^c \\ V_\varepsilon(t,x) = \ell_\varepsilon(x) & \text{on} \quad [0,T]\times(A_\varepsilon\cup B)\bigcup\{T\}\times\mathbb{R}^n \end{cases} \tag{2.26}$$

Proof. The continuity of the value function V_ε defined as in (2.25) readily follows from Lipschitz continuity of the payoff function ℓ_ε and uniform continuity of the stopped solution process in Lemma 2.4.11.[6] The PDE characterization of V_ε in (2.26) is the straightforward consequence of its continuity and Theorem 2.4.10 with boundary condition in Proposition 2.4.12. The uniqueness follows from the weak comparison principle, [FS06, Theorem 8.1 Chap. VII, p. 274], that in fact requires $(A_\varepsilon \cup B)^c$ being bounded. □

The following Remark summarizes the preceding results and pave the analytical ground on so that the Reach-Avoid problem is amenable to numerical solutions by means of off-the-shelf PDE solvers.

Remark 2.5.4. *Theorem 2.5.1 implies that the conservative approximation V_ε can be arbitrarily precise, i.e., $V(t,x) = \lim_{\varepsilon\downarrow 0} V_\varepsilon(t,x)$. Corollary 2.5.3 implies that V_ε is continuous, i.e., the PDE characterization in Theorem 2.4.10 can be simplified to the continuous version. Continuous viscosity solution can be numerically solved by invoking existing toolboxes, e.g. [Mit05]. The precision of numerical solutions can also be made arbitrarily accurate at the cost of computational time and storage. In other words, let V_ε^δ be the numerical solution of V_ε obtained through a numerical routine, and let δ be the discretization parameter (grid size) as required by [Mit05]. Then, since the continuous PDE characterization meets the hypothesis required for the toolbox [Mit05], we have $V_\varepsilon = \lim_{\delta\downarrow 0} V_\varepsilon^\delta$, and consequently we have $V(t,x) = \lim_{\varepsilon\downarrow 0}\lim_{\delta\downarrow 0} V_\varepsilon^\delta(t,x)$.*

2.6 Numerical Example: Zermelo Navigation Problem

To illustrate the theoretical results of the preceding sections, we apply the proposed reach-avoid formulation to the *Zermelo* navigation problem with constraints and stochastic uncertainties.

[5] One may replace this condition by imposing the drift and diffusion terms to be bounded.

[6] This continuity result can, alternatively, be deduced via the comparison result of the viscosity characterization of Theorem 2.4.10 together with boundary conditions (2.23b) [CIL92].

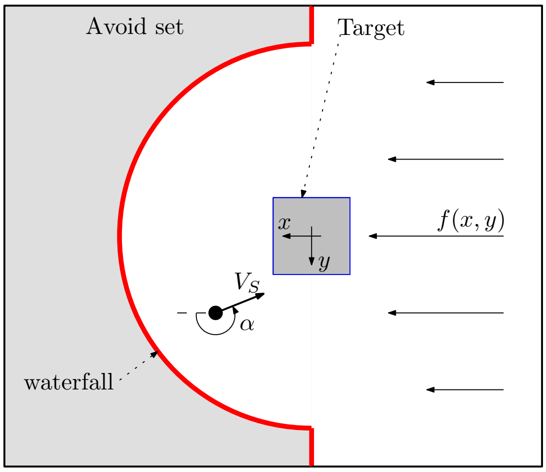

Figure 2.4: Zermelo navigation problem : a swimmer in the river

In control theory, the Zermelo navigation problem consists of a swimmer who aims to reach an island (Target) in the middle of a river while avoiding the waterfall, with the river current leading towards the waterfall. The situation is depicted in Figure 2.4. We say that the swimmer "succeeds" if he reaches the target before going over the waterfall, the latter forming a part of his Avoid set.

2.6.1 Mathematical Modeling

The dynamics of the river current are nonlinear; we let $f(x, y)$ denote the river current at position (x, y) [CQSP97]. We assume that the current flows with constant direction towards the waterfall, with the magnitude of f decreasing in distance from the middle of the river:

$$f(x, y) := \begin{bmatrix} 1 - ay^2 \\ 0 \end{bmatrix}.$$

To describe the uncertainty of the river current, we consider the diffusion term

$$\sigma(x, y) := \begin{bmatrix} \sigma_x & 0 \\ 0 & \sigma_y \end{bmatrix}.$$

We assume that the swimmer moves with constant velocity V_S, and we assume that he can change his direction α instantaneously. The complete dynamics of the swimmer in the river is given by

$$\begin{bmatrix} \mathrm{d}x_s \\ \mathrm{d}y_s \end{bmatrix} = \begin{bmatrix} 1 - ay^2 + V_S \cos(\alpha) \\ V_S \sin(\alpha) \end{bmatrix} \mathrm{d}s + \begin{bmatrix} \sigma_x & 0 \\ 0 & \sigma_y \end{bmatrix} \mathrm{d}W_s, \tag{2.27}$$

where W_s is a two-dimensional Brownian motion, and $\alpha \in [\pi, \pi]$ is the direction of the swimmer with respect to the x axis and plays the role of the controller for the swimmer.

2.6.2 Reach-Avoid Formulation

Obviously, the probability of the swimmer's "success" starting from some initial position in the navigation region depends on starting point (x, y). As shown in Section 2.3, this probability can be characterized as the level set of a value function, and by Theorem 2.4.10 this value function is the discontinuous viscosity solution of a certain differential equation on the navigation region with particular lateral and terminal boundary conditions. The differential operator $\mathcal{L}$ in Theorem 2.4.10 can be analytically calculated in this case as follows:

$$\begin{aligned}\sup_{u\in\mathbb{U}} \mathcal{L}^u\Phi(t,x,y) = \sup_{\alpha\in[-\pi,\pi]} \Big(&\partial_t\Phi(t,x,y) + \big(1 - ay^2 + V_S\cos(\alpha)\big)\partial_x\Phi(t,x,y) \\ &+ V_S\sin(\alpha)\partial_y\Phi(t,x,y) + \frac{1}{2}\sigma_x^2\partial_x^2\Phi(t,x,y) + \frac{1}{2}\sigma_y^2\partial_y^2\Phi(t,x,y)\Big).\end{aligned}$$

It can be shown that the controller value maximizing the above Dynkin operator is

$$\begin{aligned}\alpha^*(t,x,y) &:= \arg\max_{\alpha\in[-\pi,\pi]} \Big(\cos(\alpha)\partial_x\Phi(t,x,y) + \sin(\alpha)\partial_y\Phi(t,x,y)\Big) \\ &= \arctan\Big(\frac{\partial_y\Phi}{\partial_x\Phi}\Big)(t,x,y).\end{aligned}$$

Therefore, the differential operator can be simplified to

$$\begin{aligned}\sup_{u\in\mathbb{U}} \mathcal{L}^u\Phi(t,x,y) = \partial_t\Phi(t,x,y) &+ (1 - ay^2)\partial_x\Phi(t,x,y) + \frac{1}{2}\sigma_x^2\partial_x^2\Phi(t,x,y) \\ &+ \frac{1}{2}\sigma_y^2\partial_y^2\Phi(t,x,y) + V_S\|\nabla\Phi(t,x,y)\|,\end{aligned}$$

where $\nabla\Phi(t,x,y) := \big[\partial_x\Phi(t,x,y) \quad \partial_y\Phi(t,x,y)\big]$.

2.6.3 Simulation Results

For the following numerical simulations we fix the diffusion coefficients $\sigma_x = 0.5$ and $\sigma_y = 0.2$. We investigate three different scenarios: first, we assume that the river current is uniform, i.e., $a = 0\ \mathrm{m}^{-1}\mathrm{s}^{-1}$ in (2.27). Moreover, we consider the case that the swimmer velocity is less than the current flow, e.g., $V_S = 0.6\ \mathrm{ms}^{-1}$. Based on the above calculations, Figure 2.5(a) depicts the value function which is the numerical solution of the differential operator equation in Theorem 2.4.10 with the corresponding terminal and lateral conditions. As expected, since the swimmer's speed is less than the river current, if he starts from the beyond the target he has less chance of reach the island. This scenario is also captured by the value function shown in Figure 2.5(a).

Second, we assume that the river current is non-uniform and decreases with respect to the distance from the middle of the river. This means that the swimmer, even in the case that his speed is less than the current, has a non-zero probability of success if he initially swims to the sides of the river partially against its direction, followed by swimming in the direction of the current to reach the target. This scenario is depicted in Figure 2.5(b), where a non-uniform river current $a = 0.04\ \mathrm{m}^{-1}\mathrm{s}^{-1}$ in (2.27) is considered.

Third, we consider the case that the swimmer can swim faster than the river current. In this case we expect the swimmer to succeed with some probability even if he starts from beyond

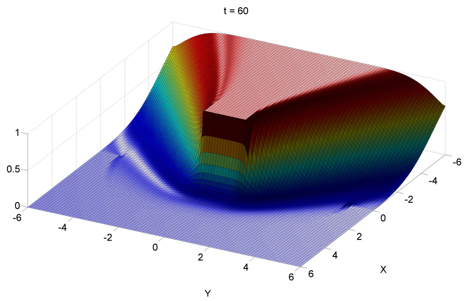

(a) The first scenario: the swimmer's speed is slower than the river current, the current being assumed uniform.

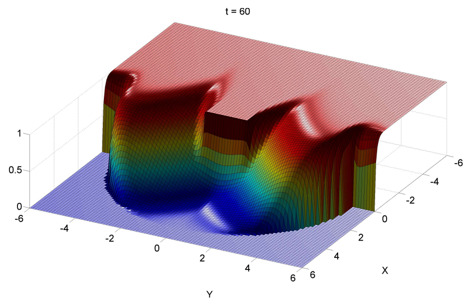

(b) The second scenario: the swimmer's speed is slower than the maximum river current.

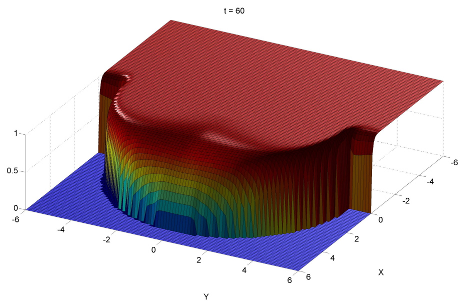

(c) The third scenario: the swimmer can swim faster than the maximum river current.

Figure 2.5: The value functions for the different scenarios

the target. This scenario is captured in Figure 2.5(c), where the reachable set (of course in probabilistic fashion) covers the entire navigation region of the river except the region near the waterfall.

In the following we show the level sets of the aforementioned value functions for $p = 0.9$. To wit, as defined in Section 2.3 (and in particular in Proposition 2.3.4), these level sets, roughly speaking, correspond to the reachable sets with probability $p = 90\%$ in certain time horizons while the swimmer is avoiding the waterfall. By definition, as can be seen in Figure 2.6, these sets are nested with respect to the time horizon.

All simulations were obtained using the Level Set Method Toolbox [Mit05] (version 1.1), with a grid 101×101 in the region of simulation.

2.7 Summary and Outlook

In this chapter we presented a new method to address a class of stochastic reachability problems with state constraints. The proposed framework provides a set characterization of the stochastic reach-avoid set based on discontinuous viscosity solutions of a second order PDE. In contrast to earlier approaches, this methodology is not restricted to almost-sure notions, and one can compute the set of initial conditions that can satisfy the reach-avoid specification with any

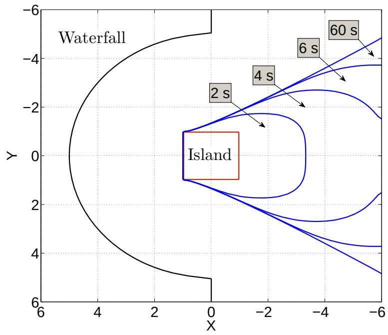

(a) The first scenario: the swimmer's speed is slower than the river current, the current being assumed uniform.

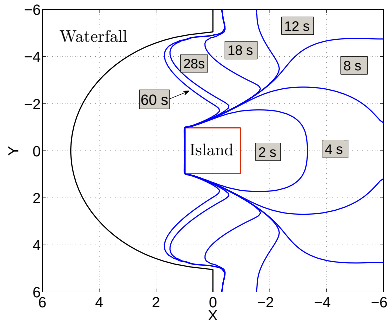

(b) The second scenario: the swimmer's speed is slower than the maximum river current.

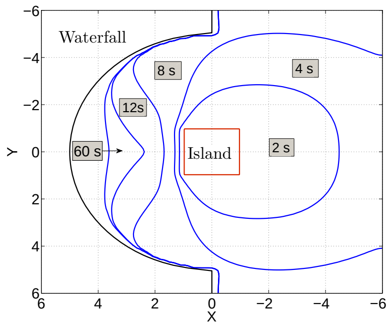

(c) The third scenario: the swimmer can swim faster than the maximum river current.

Figure 2.6: The level sets of the value functions for the different scenarios

non-zero probability by means of off-the-shelf PDE solvers.

A natural extension of the reach-avoid problem is the motion planning objective as an excursion through more than one target set in a specific order, while avoiding certain obstacles. This extension will be the topic of the next chapter. Moreover, in light of Proposition 2.3.5 (resp. Proposition 2.3.7), we know that the stochastic optimal control problems in (2.4) (resp. (2.5)) have a close connection to the reach-avoid problem. This chapter only studied the exit-time formulation. As a step further, however, it would be interesting to study the other formulations, in particular the interpretation of the reach avoid problem as a dynamic game between two players with different authorities, e.g., a differential game between a stopper and controller.

2.8 Appendix

This Appendix collects the proofs omitted throughout the chapter.

2.8.1 Proofs of Section 2.3

Proof of Proposition 2.3.4. In view of Assumption 2.3.3, the implication (2.3b), and the definition of reach-avoid set in 2.2.2, we can express the set $\mathrm{RA}(t,p;A,B)$ as

$$\mathrm{RA}(t,p;A,B) = \Big\{x \in \mathbb{R}^n \;\big|\; \exists \boldsymbol{u} \in \mathcal{U} \;:\; \mathbb{P}\big(\tau_A < \tau_B \;\text{ and }\; \tau_A \le T\big) > p\Big\}. \tag{2.28}$$

Also, by Assumption 2.3.3, the properties (2.3a) and (2.3c), and the definition of stopping time $\bar{\tau}$ in (2.4a), given $\boldsymbol{u} \in \mathcal{U}$ we have

$$X_{\bar{\tau}}^{t,x;\boldsymbol{u}} \in A \Longrightarrow \tau_A \le \bar{\tau} \text{ and } \bar{\tau} \neq \tau_B \Longrightarrow T \ge \bar{\tau} = \tau_A < \tau_B,$$

which means the sample path $X_{\cdot}^{t,x;\boldsymbol{u}}$ hits the set A before B at the time $\bar{\tau} \le T$. Moreover,

$$X_{\bar{\tau}}^{t,x;\boldsymbol{u}} \notin A \Longrightarrow \bar{\tau} \neq \tau_A \Longrightarrow \bar{\tau} = (\tau_B \wedge T) < \tau_A,$$

and this means that the sample path does not succeed in reaching A while avoiding set B within time T. Therefore, the event $\{\tau_A < \tau_B \;\text{ and }\; \tau_A \le T\}$ is equivalent to $\{X_{\bar{\tau}}^{t,x;\boldsymbol{u}} \in A\}$, and

$$\mathbb{P}\big(\tau_A < \tau_B \;\text{ and }\; \tau_A \le T\big) = \mathbb{E}\big[\mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}})\big].$$

This, in view of (2.28) and arbitrariness of control strategy $\boldsymbol{u} \in \mathcal{U}$ leads to the assertion. □

Proof of Proposition 2.3.5. We first establish the equality of $V_1 = V_2$. To this end, let us fix $\boldsymbol{u} \in \mathcal{U}$ and (t,x) in $[0,T] \times \mathbb{R}^n$. Observe that it suffices to show that pointwise on Ω,

$$\mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}}) = \sup_{s\in[t,T]} \{\mathbb{1}_A(X_s^{t,x;\boldsymbol{u}}) \wedge \inf_{r\in[t,s]} \mathbb{1}_{B^c}(X_r^{t,x;\boldsymbol{u}})\}.$$

According to the Assumption 2.3.3 and Remark 2.3.2, one can simply see that

$$\begin{aligned}
&\sup_{s\in[t,T]} \{\mathbb{1}_A(X_s^{t,x;\boldsymbol{u}}) \wedge \inf_{r\in[t,s]} \mathbb{1}_{B^c}(X_r^{t,x;\boldsymbol{u}})\} = 1 \\
&\qquad \Longleftrightarrow \exists s \in [t,T]\; X_s^{t,x;\boldsymbol{u}} \in A \text{ and } \forall r \in [t,s]\; X_r^{t,x;\boldsymbol{u}} \in B^c \\
&\qquad \Longleftrightarrow \exists s \in [t,T]\; \tau_A \le s \le T \text{ and } \tau_B > s \\
&\qquad \Longleftrightarrow X_{\tau_A}^{t,x;\boldsymbol{u}} = X_{\tau_A\wedge\tau_B\wedge T}^{t,x;\boldsymbol{u}} = X_{\tau_{A\cup B}\wedge T}^{t,x;\boldsymbol{u}} \in A \\
&\qquad \Longleftrightarrow \mathbb{1}_A\big(X_{\bar{\tau}}^{t,x;\boldsymbol{u}}\big) = 1
\end{aligned}$$

and since the functions take values in $\{0,1\}$, we have $V_1(t,x) = V_2(t,x)$.

As a first step towards proving $V_1 = V_3$, we start with establishing $V_3 \ge V_1$. It is straightforward from the definition that

$$\sup_{\tau\in\mathcal{T}_{[t,T]}} \inf_{\sigma\in\mathcal{T}_{[t,\tau]}} \mathbb{E}\Big[\mathbb{1}_A(X_\tau^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_\sigma^{t,x;\boldsymbol{u}})\Big] \ge \inf_{\sigma\in\mathcal{T}_{[t,\bar{\tau}]}} \mathbb{E}\Big[\mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_\sigma^{t,x;\boldsymbol{u}})\Big], \tag{2.29}$$

where $\bar{\tau}$ is the stopping time defined in (2.4a). For all stopping times $\sigma \in \mathcal{T}_{[t,\bar{\tau}]}$, in view of (2.3b) we have

$$\begin{aligned}
\mathbb{1}_{B^c}(X_\sigma^{t,x;\boldsymbol{u}}) = 0 &\Longrightarrow X_\sigma^{t,x;\boldsymbol{u}} \in B \Longrightarrow \tau_B \le \sigma \le \bar{\tau} = \tau_A \wedge \tau_B \wedge T \\
&\Longrightarrow \tau_B = \sigma = \bar{\tau} < \tau_A \Longrightarrow X_{\bar{\tau}}^{t,x;\boldsymbol{u}} \notin A \\
&\Longrightarrow \mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}}) = 0
\end{aligned}$$

This implies that for all $\sigma \in \mathcal{T}_{[t,\bar{\tau}]}$,

$$\mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_{\sigma}^{t,x;\boldsymbol{u}}) = \mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}}) \qquad \mathbb{P}\text{-a.s.}$$

which, in connection with (2.29) leads to

$$\sup_{\tau \in \mathcal{T}_{[t,T]}} \inf_{\sigma \in \mathcal{T}_{[t,\tau]}} \mathbb{E}\Big[\mathbb{1}_A(X_{\tau}^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_{\sigma}^{t,x;\boldsymbol{u}})\Big] \geq \mathbb{E}\Big[\mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}})\Big].$$

By arbitrariness of the control strategy $\boldsymbol{u} \in \mathcal{U}$, we get $V_3 \geq V_1$. It remains to show $V_2 \leq V_1$. Given $\boldsymbol{u} \in \mathcal{U}$ and $\tau \in \mathcal{T}_{[t,T]}$, let us choose $\bar{\sigma} := \tau \wedge \tau_B$. Note that since $t \leq \bar{\sigma} \leq \tau$ then $\bar{\sigma} \in \mathcal{T}_{[t,\tau]}$. Hence,

$$\inf_{\sigma \in \mathcal{T}_{[t,\tau]}} \mathbb{E}\Big[\mathbb{1}_A(X_{\tau}^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_{\sigma}^{t,x;\boldsymbol{u}})\Big] \leq \mathbb{E}\Big[\mathbb{1}_A(X_{\tau}^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_{\bar{\sigma}}^{t,x;\boldsymbol{u}})\Big]. \tag{2.30}$$

Note that by an argument similar to the proof of Proposition 2.3.4, for all $\tau \in \mathcal{T}_{[t,T]}$:

$$\begin{aligned}
\mathbb{1}_A(X_{\tau}^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_{\bar{\sigma}}^{t,x;\boldsymbol{u}}) = 1 &\Longrightarrow X_{\tau}^{t,x;\boldsymbol{u}} \in A \text{ and } X_{\bar{\sigma}}^{t,x;\boldsymbol{u}} \notin B \\
&\Longrightarrow \tau_A \leq \tau \leq T \text{ and } \bar{\sigma} \neq \tau_B \\
&\Longrightarrow \tau_A \leq \tau \leq T \text{ and } \tau_A \leq \bar{\sigma} = \tau < \tau_B \\
&\Longrightarrow \bar{\tau} = \tau_A \wedge \tau_B \wedge T = \tau_A \Longrightarrow \mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}}) = 1.
\end{aligned}$$

It follows that for all $\tau \in \mathcal{T}_{[t,\tau]}$,

$$\mathbb{1}_A(X_{\tau}^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_{\bar{\sigma}}^{t,x;\boldsymbol{u}}) \leq \mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}}) \qquad \mathbb{P}\text{-a.s.}$$

which in connection with (2.30) leads to

$$\sup_{\tau \in \mathcal{T}_{[t,T]}} \inf_{\sigma \in \mathcal{T}_{[t,\tau]}} \mathbb{E}\Big[\mathbb{1}_A(X_{\tau}^{t,x;\boldsymbol{u}}) \wedge \mathbb{1}_{B^c}(X_{\sigma}^{t,x;\boldsymbol{u}})\Big] \leq \mathbb{E}\Big[\mathbb{1}_A(X_{\bar{\tau}}^{t,x;\boldsymbol{u}})\Big].$$

By arbitrariness of the control strategy $\boldsymbol{u} \in \mathcal{U}$ we arrive at $V_3 \leq V_1$. □

2.8.2 Proofs of Section 2.4

Proof of Proposition 2.4.3. We first prove continuity of $\bar{\tau}(t,x)$ with respect to (t,x). Let us take a sequence $(t_n, x_n) \to (t_0, x_0)$, and let $\big(X_r^{t_n,x_n;\boldsymbol{u}}\big)_{r \geq t_n}$ be the solution of (2.1) for a given policy $\boldsymbol{u} \in \mathcal{U}$. Let us recall that by definition we assume that $X_s^{t,x;\boldsymbol{u}} := x$ for all $s \in [0,t]$. Here we assume that $t_n \leq t$, but one can effectively follow the same technique for $t_n > t$. Notice that it is straightforward to observe that by the definition of stochastic integral in (2.1) we have

$$X_r^{t_n,x_n;\boldsymbol{u}} = X_t^{t_n,x_n;\boldsymbol{u}} + \int_t^r f\big(X_s^{t_n,x_n;\boldsymbol{u}}, u_s\big)\mathrm{d}s + \int_t^r \sigma\big(X_s^{t_n,x_n;\boldsymbol{u}}, u_s\big)\mathrm{d}W_s \qquad \mathbb{P}\text{-a.s.}$$

Therefore, by virtue of [Kry09, Theorem 2.5.9, p. 83], for all $q \geq 1$ we have

$$\begin{aligned}
\mathbb{E}\Big[\sup_{r \in [t,T]} \big\|X_r^{t,x;\boldsymbol{u}} - X_r^{t_n,x_n;\boldsymbol{u}}\big\|^{2q}\Big] &\leq C_1(q,T,K)\mathbb{E}\Big[\big\|x - X_t^{t_n,x_n;\boldsymbol{u}}\big\|^{2q}\Big] \\
&\leq 2^{2q-1} C_1(q,T,K)\mathbb{E}\Big[\|x - x_n\|^{2q} + \big\|x_n - X_t^{t_n,x_n;\boldsymbol{u}}\big\|^{2q}\Big],
\end{aligned}$$

where in light of [Kry09, Corollary 2.5.12, p. 86], it leads to

$$\mathbb{E}\Big[\sup_{r\in[t,T]}\big\|X_r^{t,x;\boldsymbol{u}}-X_r^{t_n,x_n;\boldsymbol{u}}\big\|^{2q}\Big]\leq C_2(q,T,K,\|x\|)\big(\|x-x_n\|^{2q}+|t-t_n|^q\big). \tag{2.31}$$

In the above relations K is the Lipschitz constant of f and σ mentioned in Assumption 2.2.1; C_1 and C_2 are constant depending on the indicated parameters. Hence, in view of Kolmogorov's Continuity Criterion [Pro05, Corollary 1 Chap. IV, p. 220], one may consider a version of the stochastic process $X_\cdot^{t,x;\boldsymbol{u}}$ which is continuous in (t,x) in the topology of uniform convergence on compacts. This yields to the fact that $\mathbb{P}$-a.s, for any $\varepsilon>0$, for all sufficiently large n,

$$X_r^{t_n,x_n;\boldsymbol{u}}\in\mathbb{B}_\varepsilon\big(X_r^{t_0,x_0;\boldsymbol{u}}\big),\qquad\forall r\in[t_n,T], \tag{2.32}$$

where $\mathbb{B}_\varepsilon(y)$ denotes the ball centered at y and radius ε. Based on the Assumptions 2.4.1.a. and 2.4.1.b., it is a well-known property of non-degenerate processes that the set of sample paths that hit the boundary of O and do not enter the set is negligible [RB98, Corollary 3.2, p. 65]. Hence, by the definition of $\bar\tau$ and (2.3b), one can conclude that

$$\forall\delta>0,\ \exists\varepsilon>0,\qquad\bigcup_{s\in[t_0,\bar\tau(t_0,x_0)-\delta]}\mathbb{B}_\varepsilon\big(X_s^{t_0,x_0;\boldsymbol{u}}\big)\cap\overline{O}=\emptyset\qquad\mathbb{P}\text{-a.s.}$$

This together with (2.32) indicates that $\mathbb{P}$-a.s. for all sufficiently large n,

$$X_r^{t_n,x_n;\boldsymbol{u}}\notin\overline{O},\qquad\forall r\in[t_n,\bar\tau(t_0,x_0)[\ ,$$

which in conjunction with $\mathbb{P}$-a.s. continuity of sample paths immediately leads to

$$\liminf_{(t_n,x_n)\to(t,x)}\bar\tau(t_n,x_n)\geq\bar\tau(t_0,x_0)\qquad\mathbb{P}\text{-a.s.} \tag{2.33}$$

On the other hand by the definition of $\bar\tau$ and Assumptions 2.4.1.a. and 2.4.1.b., again in view of [RB98, Corollary 3.2, p. 65],

$$\forall\delta>0,\quad\exists s\in[\tau_O(t_0,x_0),\tau_O(t_0,x_0)+\delta[,\quad X_s^{t_0,x_0;\boldsymbol{u}}\in O^\circ\quad\mathbb{P}\text{-a.s.},$$

where τ_O is the first entry time to O, and O° denotes the interior of the set O. Hence, in light of (2.32), $\mathbb{P}$-a.s. there exists $\varepsilon>0$, possibly depending on δ, such that for all sufficiently large n we have $X_s^{t_n,x_n;\boldsymbol{u}}\in\mathbb{B}_\varepsilon(X_s^{t_0,x_0;\boldsymbol{u}})\subset O$. According to the definition of $\tau_O(t_n,x_n)$ and (2.3b), this implies $\tau_O(t_n,x_n)\leq s<\tau_O(t_0,x_0)+\delta$. From arbitrariness of δ and the definition of $\bar\tau$ in (5.19), it leads to

$$\limsup_{(t_n,x_n)\to(t,x)}\bar\tau(t_n,x_n)\leq\bar\tau(t_0,x_0)\qquad\mathbb{P}\text{-a.s.},$$

where in conjunction with (2.33), $\mathbb{P}$-a.s. continuity of the map $(t,x)\mapsto\bar\tau(t,x)$ at (t_0,x_0) follows.

It remains to show lower semicontinuity of J. Note that J is bounded since ℓ is. In accordance with the $\mathbb{P}$-a.s. continuity of $X_r^{t,x;\boldsymbol{u}}$ and $\bar\tau(t,x)$ with respect to (t,x), and Fatou's lemma, we have

$$\begin{aligned}\liminf_{n\to\infty}J\big(t_n,x_n,\boldsymbol{u}\big)&=\liminf_{n\to\infty}\mathbb{E}\Big[\ell\big(X_{\bar\tau(t_n,x_n)}^{t_n,x_n;\boldsymbol{u}}\big)\Big]\\&=\liminf_{n\to\infty}\mathbb{E}\Big[\ell\big(X_{\bar\tau(t_n,x_n)}^{t_n,x_n;\boldsymbol{u}}-X_{\bar\tau(t_n,x_n)}^{t,x;\boldsymbol{u}}+X_{\bar\tau(t_n,x_n)}^{t,x;\boldsymbol{u}}-X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}+X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}\big)\Big]\\&=\liminf_{n\to\infty}\mathbb{E}\Big[\ell\big(\varepsilon_n+X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}\big)\Big]\geq\mathbb{E}\Big[\liminf_{n\to\infty}\ell\big(\varepsilon_n+X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}\big)\Big]\\&\geq\mathbb{E}\Big[\ell\big(X_{\bar\tau(t,x)}^{t,x;\boldsymbol{u}}\big)\Big]=J(t,x,\boldsymbol{u}),\end{aligned} \tag{2.34}$$

where inequality in (2.34) follows from Fatou's Lemma, and $\varepsilon_n \to 0$ $\mathbb{P}$-a.s. as n tends to ∞. Note that by definition $X^{t,x;\boldsymbol{u}}_{\bar\tau(t_n,x_n)} = x$ on the set $\{\bar\tau(t_n,x_n) < t\}$. $\square$

Proof of Lemma 2.4.6. By Definition 2.3.1, one has

$$\mathbb{1}_{\{\bar\tau(t,x)\geq\theta\}}\bar\tau(t,x) = \mathbb{1}_{\{\bar\tau(t,x)\geq\theta\}}\big(\bar\tau(\theta, X^{t,x;\boldsymbol{u}}_\theta) + \theta - t\big) \qquad \mathbb{P}\text{-a.s.}$$

One can now follow effectively the same computations as in the proof of [BT11, Proposition 5.1] to conclude the assertion. $\square$

Proof of Lemma 2.4.11. Let us consider a version of $X^{t,x;\boldsymbol{u}}_\cdot$ which is almost surely continuous in (t,x) uniformly respect to the policy $\boldsymbol{u}$; this is always possible since the constant C_2 in (2.31) does not depend on $\boldsymbol{u}$. That is, $\boldsymbol{u}$ may only affect a negligible subset of Ω; we refer to [Pro05, Theorem 72 Chap. IV, p. 218] for further details on this issue. Hence, all the relations in the proof of Proposition 2.4.3, in particular (2.32), hold if we permit the control policy $\boldsymbol{u}$ to depend on n in an arbitrary way. Therefore, the assertions of Proposition 2.4.3 holds uniformly with respect to $(\boldsymbol{u}_n)_{n\in\mathbb{N}} \subset \mathcal{U}$. That is, for all $(t,x)\in\mathbb{S}$, $(t_n,x_n)\to(t,x)$, and $(\boldsymbol{u}_n)_{n\in\mathbb{N}}$, with probability one we have

$$\lim_{n\to\infty}\sup_{s\in[0,T]}\big\|X^{t_n,x_n;\boldsymbol{u}_n}_s - X^{t,x;\boldsymbol{u}_n}_s\big\| = 0, \qquad \lim_{n\to\infty}\big|\bar\tau(t_n,x_n) - \bar\tau(t,x)\big| = 0 \tag{2.35}$$

where $\bar\tau$ is as defined in (2.6) while the solution process is driven by control policies $\boldsymbol{u}_n$. Moreover, according to [Kry09, Corollary 2.5.10, p. 85]

$$\mathbb{E}\Big[\big\|X^{t,x;\boldsymbol{u}}_r - X^{t,x;\boldsymbol{u}}_s\big\|^{2q}\Big] \leq C_3\big(q,T,K,\|x\|\big)\big|r-s\big|^q, \qquad \forall r,s\in[t,T] \quad \forall q\geq 1,$$

following the arguments in the proof of Proposition 2.4.3 in conjunction with above inequality, one can also deduce that the mapping $s\mapsto X^{t,x;\boldsymbol{u}}_s$ is $\mathbb{P}$-a.s. continuous uniformly with respect to $\boldsymbol{u}$. Hence, one can infer that for all $(t,x)\in\mathbb{S}$, with probability one we have

$$\begin{aligned}\lim_{n\to\infty}\big\|X^{t_n,x_n;\boldsymbol{u}_n}_{\bar\tau(t_n,x_n)} - X^{t,x;\boldsymbol{u}_n}_{\bar\tau(t,x)}\big\| &\leq \lim_{n\to\infty}\big\|X^{t_n,x_n;\boldsymbol{u}_n}_{\bar\tau(t_n,x_n)} - X^{t,x;\boldsymbol{u}_n}_{\bar\tau(t_n,x_n)}\big\| \\ &\quad + \lim_{n\to\infty}\big\|X^{t,x;\boldsymbol{u}_n}_{\bar\tau(t_n,x_n)} - X^{t,x;\boldsymbol{u}_n}_{\bar\tau(t,x)}\big\| = 0.\end{aligned}$$

Notice that the first limit term above tends to zero as the version of the solution process $X^{t,x;\boldsymbol{u}_n}_\cdot$ on the compact set $[0,T]$ is continuous in the initial condition (t,x) uniformly with respect to n. The second term is the consequence of limits in (2.35) and continuity of the mapping $s\mapsto X^{t,x;\boldsymbol{u}_n}_s$ uniformly in $n\in\mathbb{N}$. $\square$

2.8.3 Proofs of Section 2.5

Proof of Theorem 2.5.1. By definition, the family of the sets $(A_\varepsilon)_{\varepsilon>0}$ is nested and increasing as $\varepsilon\downarrow 0$. Therefore, in view of (2.3a), τ_ε is nonincreasing as $\varepsilon\downarrow 0$ pathwise on Ω. Moreover it is obvious to see that the family of functions ℓ_ε is increasing with respect to ε. Hence, given an initial condition $(t,x) \in \mathbb{S}$, an admissible control $\boldsymbol{u} \in \mathcal{U}_t$, and $\varepsilon_1 \geq \varepsilon_2 > 0$, pathwise on Ω we have

$$\begin{aligned}\ell_{\varepsilon_2}\big(X^{t,x;\boldsymbol{u}}_{\tau_{\varepsilon_2}}\big) < 1 &\Longrightarrow \tau_{\varepsilon_2} = \tau_B \wedge T < \tau_{A_{\varepsilon_2}} < \tau_{A_{\varepsilon_1}} \\ &\Longrightarrow \tau_{\varepsilon_1} = \tau_B \wedge T = \tau_{\varepsilon_2} \Longrightarrow \ell_{\varepsilon_2}\big(X^{t,x;\boldsymbol{u}}_{\tau_{\varepsilon_2}}\big) \geq \ell_{\varepsilon_1}\big(X^{t,x;\boldsymbol{u}}_{\tau_{\varepsilon_1}}\big),\end{aligned}$$

which immediately leads to $V_{\varepsilon_2}(t,x) \geq V_{\varepsilon_1}(t,x)$. Now let $(\varepsilon_i)_{i\in\mathbb{N}}$ be a decreasing sequence of positive numbers that converges to zero, and for the simplicity of notation let $A_n := A_{\varepsilon_n}$, $\tau_n := \tau_{\varepsilon_n}$, and $\ell_n := \ell_{\varepsilon_n}$. According to the definitions (2.4a) and (2.25), we have

$$\begin{aligned}
V(t,x) - \lim_{n\to\infty} V_{\varepsilon_n}(t,x) &= \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_A\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}}\big)\Big] - \lim_{n\to\infty} \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\big[\ell_n\big(X^{t,x;\boldsymbol{u}}_{\tau_n}\big)\big] \\
&= \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_A\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}}\big)\Big] - \sup_{n\in\mathbb{N}} \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\big[\ell_n\big(X^{t,x;\boldsymbol{u}}_{\tau_n}\big)\big] && (2.36a) \\
&\leq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \Big(\mathbb{E}\Big[\mathbb{1}_A\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}}\big)\Big] - \sup_{n\in\mathbb{N}} \mathbb{E}\big[\ell_n\big(X^{t,x;\boldsymbol{u}}_{\tau_n}\big)\big]\Big) \\
&\leq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \inf_{n\in\mathbb{N}} \mathbb{E}\Big[\mathbb{1}_A\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}}\big) - \mathbb{1}_{A_n}\big(X^{t,x;\boldsymbol{u}}_{\tau_n}\big)\Big] \\
&= \sup_{\boldsymbol{u}\in\mathcal{U}_t} \inf_{n\in\mathbb{N}} \mathbb{P}\Big(\{\tau_{A_n} > \tau_B \wedge T\} \cap \{\tau_A \leq T\} \cap \{\tau_A < \tau_B\}\Big) && (2.36b) \\
&= \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{P}\Big(\bigcap_{n\in\mathbb{N}} \{\tau_{A_n} > \tau_B \wedge T\} \cap \{\tau_A \leq T\} \cap \{\tau_A < \tau_B\}\Big) && (2.36c) \\
&\leq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{P}\Big(\{\tau_{A^\circ} \geq \tau_B \wedge T\} \cap \{\tau_A \leq T\} \cap \{\tau_A < \tau_B\}\Big) && (2.36d) \\
&\leq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{P}\big(\{\tau_{A^\circ} > \tau_A\} \cup \{\tau_A = T\}\big) = 0 && (2.36e)
\end{aligned}$$

Note that the equality in (2.36a) is due to the fact that the sequence of the value functions $\big(V_{\varepsilon_n}\big)_{n\in\mathbb{N}}$ is increasing pointwise. One can infer the equality (2.36b) when $\mathbb{1}_A\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}}\big) = 1$ and $\mathbb{1}_{A_n}\big(X^{t,x;\boldsymbol{u}}_{\tau_n}\big) = 0$ as $\mathbb{1}_A\big(X^{t,x;\boldsymbol{u}}_{\bar{\tau}}\big) \geq \mathbb{1}_{A_n}\big(X^{t,x;\boldsymbol{u}}_{\tau_n}\big)$ pathwise on Ω. Moreover, since the sequence of the stopping times $(\tau_n)_{n\in\mathbb{N}}$ is decreasing $\mathbb{P}$-a.s., the family of sets $\big(\{\tau_{A_n} > \tau_A\}\big)_{n\in\mathbb{N}}$ is also decreasing; consequently, the equality (2.36c) follows. In order to show (2.36d), it is not hard to inspect that

$$\begin{aligned}
\omega \in \bigcap_{n\in\mathbb{N}} \{\tau_{A_n} > \tau_B \wedge T\} &\Longrightarrow \forall n \in \mathbb{N}, \quad \tau_{A_n}(\omega) > \tau_B(\omega) \wedge T \\
&\Longrightarrow \forall n \in \mathbb{N}, \quad \forall s \leq \tau_B(\omega) \wedge T, \quad X^{t,x;\boldsymbol{u}}_s(\omega) \notin A_n \\
&\Longrightarrow \forall s \leq \tau_B(\omega) \wedge T, \quad X^{t,x;\boldsymbol{u}}_s(\omega) \notin \bigcup_{n\in\mathbb{N}} A_n = A^\circ \\
&\Longrightarrow \omega \in \{\tau_{A^\circ} \geq \tau_B \wedge T\}.
\end{aligned}$$

Based on non-degeneracy and the interior cone condition in Assumptions 2.4.1.a. and 2.4.1.b. respectively, by virtue of [RB98, Corollary 3.2, p. 65], we see that the set $\{\tau_{A^\circ} > \tau_A\}$ is negligible. Moreover, the interior cone condition implies that the Lebesgue measure of ∂A, boundary of A, is zero. In view of non-degeneracy and Girsanov's Theorem [KS91, Theorem 5.1, p. 191], $X^{t,x;\boldsymbol{u}}_r$ has a probability density $d(r,y)$ for $r \in]t,T]$; see [FS06, Section IV.4] and references therein. Hence, the aforesaid property of ∂A results in $\mathbb{P}\{\tau_A = T\} \le \mathbb{P}\big\{X_T^{t,x;\boldsymbol{u}} \in \partial A\big\} = \int_{\partial A} d(T,y)\mathrm{d}y = 0$, and the second equality of (2.36e) follows. It is straightforward to see $V \ge V_{\varepsilon_n}$ pointwise on $\mathbb{S}$ for all $n \in \mathbb{N}$. The assertion now follows at once. $\square$

CHAPTER 3

Stochastic Motion Planning

In this chapter, we extend the reach-avoid problem discussed in the preceding chapter to different classes of stochastic motion planning problems which involve a controlled process, with possibly discontinuous sample paths, visiting certain subsets of the state-space while avoiding others in a sequential fashion. For this purpose, we first introduce two basic notions of motion planning, and then establish a connection to a class of stochastic optimal control problems concerned with sequential stopping times. A weak dynamic programming principle (DPP) is then proposed, which characterizes the set of initial states that admit a policy enabling the process to execute the desired maneuver with probability no less than some pre-specified value. The proposed DPP comprises auxiliary value functions defined in terms of discontinuous payoff functions. A concrete instance of the use of this novel DPP in the case of diffusion processes is also presented. In this case, we establish that the aforementioned set of initial states can be characterized as the level set of a discontinuous viscosity solution to a sequence of partial differential equations, for which the first one has a known boundary condition, while the boundary conditions of the subsequent ones are determined by the solutions to the preceding steps. Finally, the generality and flexibility of the theoretical results are illustrated on an example involving a biological switch.

3.1 Introduction

Motion planning of dynamical systems can be viewed as a scheme for executing excursions of the state of the system to certain given sets in a specific order according to a specified time schedule. The two fields of robotics and control have contributed much to motion planning. In the robotics community, research on motion planning typically focuses on the computational issues along with considerations of basic kinematic limitations; see examples of navigation of unmanned air vehicles [MB00, BMGA02], and recent surveys on motion planning algorithms [GKM10] and dynamic vehicle routing [BFP$^+$11]. In the control community motion planning emphasizes the dynamic behavior and specific aspects of trajectory performance that usually involve high order differential constraints. This chapter deals with motion planning from the latter point of view.

In the control literature, motion planning problems have been studied extensively in a deterministic setting from the differential geometric [Sus91, CS98a, MS90] and dynamic programming [SM86] perspectives. However, motion planning in the stochastic setting has received relatively little attention. In fact, it was not until recently when the basic motion planning problem involving one target and one obstacle set—the so-called reach-avoid problem, has been investigated in the context of finite probability spaces for a class of continuous-time Markov decision processes [BHKH05], and in the discrete-time stochastic hybrid systems context [CCL11, SL10].

In the continuous time and space settings, one may tackle the dynamic programming formulation of the reach-avoid problem from two perspectives: a direct technique based on the theory of stochastic target problems, and an indirect approach via an exit-time stochastic optimal control formulation. For the former, we refer the reader to [ST02b, BET10]; see also the recent book [Tou13] for details. In Chapter 2 we focused on the latter perspective for reachability of controlled diffusion processes. Here we continue in the same spirit by going beyond the reach-avoid problem to more complex motion planning problems for a larger class of stochastic processes with possibly discontinuous sample paths.

The contributions of this work are outlined as follows:

(i) we formalize the stochastic motion planning problem for continuous time, continuous space stochastic processes (Section 3.2);

(ii) we establish a connection between different motion planning maneuvers and a class of stochastic optimal control problems (Section 3.3);

(iii) we propose a weak dynamic programming principle (DPP) under mild assumptions on the admissible policies and the stochastic process (Section 3.4);

(iv) we derive a partial differential equation (PDE) characterization of the desired set of initial conditions in the context of controlled diffusions processes based on the proposed DPP (Section 3.5).

Concerning item (i), we start with the formal definition of a motion planning objective comprising of two fundamental reachability maneuvers. To the best of our knowledge, this is new in the literature. We address the following natural question: for which initial states do there exist admissible policies such that the controlled stochastic processes satisfy the motion planning objective with a probability greater than a given value p? Under item (ii), we then characterize this set of initial states by establishing a connection between the motion planning specifications and a class of stochastic optimal control problems involving discontinuous payoff functions and a sequence of successive stopping times.

Concerning item (iii), we should highlight that due to the discontinuity of the payoff functions, the classical results on stochastic optimal control problems and its connection to Hamilton-Jacobi-Bellman PDE are not applicable here. In the spirit of [BT11], under some mild assumptions, we propose a weak DPP involving auxiliary value functions. As opposed to the classical DPP results, this formulation does not need to verify the measurability of the value functions. The only non-trivial assumption required for the proposed DPP is the continuity of the sequence of stopping times with respect to the initial states.

Finally, concerning item (iv), we focus on a class of controlled diffusion processes in which the required assumptions of the proposed DPP are investigated. Indeed, it turns out that the standard uniform non-degeneracy and exterior cone conditions of the involved sets suffice to fulfill the DPP requirements. Subsequently, we demonstrate how the DPP leads to a new framework for characterizing the desired set of initial conditions based on tools from PDEs. Due to the discontinuities of the value functions involved, all the PDEs are understood in the generalized notion of the so-called discontinuous viscosity solutions. In this context, we show how the value functions can be solved by a means of a sequence of PDEs, in which the preceding PDE provides the boundary condition of the following one.

On the computational side, it is well-known that PDE techniques suffer from the curse of dimensionality. In the literature a class of suboptimal control methods referred to as Approximate Dynamic Programming (ADP) have been developed for dealing with this difficulty; for a sampling of recent works see [dFVR03, DFVR04, CRVRL06] for a linear programming approach, [KT03, VL10] for actor-critic algorithms, and [Ber05] for a comprehensive survey on the entire area. Besides the ADP literature, very recent progress on numerical methods based on tensor train decompositions holds the potential of substantially ameliorating this curse of dimensionality; see two representative articles [KO10, KS11] and the references therein. In this light, taken in its entirety, the results in this study can be viewed as a theoretical bridge between the motion planning objective formalized in (i) and sophisticated numerical methods that can be used to address real problem instances. Here we demonstrate the practical use of this bridge by addressing a stochastic motion planning problem for a biological switch.

As indicated above, the organization of the chapter follows the steps (i)-(iv). In Section 3.2 we formally introduce the stochastic motion planning problems. In Section 3.3 we establish a connection between the motion planning objectives and a class of stochastic optimal control problems, for which a weak DPP is proposed in Section 3.4. A concrete instance of the use of the novel DPP in the case of controlled diffusion processes is presented in Section V, leading to characterization of the motion planning objective with the help of a sequence of PDE's in an iterative fashion. To validate the performance of the proposed methodology, in Section 3.6 the theoretical results are applied to a biological two-gene network. For better readability, the technical proofs along with required preliminaries are provided in Appendix 3.8.

Notation

Here is a partial notation list which will be also explained in more details later in the chapter:

- $a \wedge b := \min\{a, b\}$ and $a \vee b := \max\{a, b\}$ for $a, b \in \mathbb{R}$;
- A^c (resp. A°): complement (resp. interior) of the set A;
- $\overline{A}$ (resp. ∂A): closure (resp. boundary) of the set A;
- $\mathbb{B}_r(x)$: open Euclidean ball centered at x and radius r;
- $\mathfrak{B}(\mathbb{A})$: Borel σ-algebra on a topological space $\mathbb{A}$;
- $\mathcal{U}_t$: set of admissible policies at time t;
- $(X_s^{t,x;\boldsymbol{u}})_{s\geq 0}$: stochastic process under the control policy $\boldsymbol{u}$ and convention $X_s^{t,x;\boldsymbol{u}} := x$ for all $s \leq t$;

- $(W_i \rightsquigarrow G_i)_{\leq T_i}$ (resp. $W_i \xrightarrow{T_i} G_i$) : motion planning event of reaching G_i sometime before time T_i (resp. at time T_i) while staying in W_i, see Definition 3.2.1;

- $\left(\Theta_i^{A_{k:n}}\right)_{i=k}^n$: sequential exit-times from the sets $(A_i)_{i=k}^n$ in order, see Definition 3.3.1;

- V^* (resp. V_*): upper (resp. lower) semicontinuous envelope of the function V;

- $\mathcal{L}^u$: Dynkin operator, see Definition 2.4.9.

3.2 General Setting and Problem Description

Consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ whose filtration $\mathbb{F} := (\mathcal{F}_s)_{s\geq 0}$ is generated by an $\mathbb{R}^{d_z}$-valued process $(z_{:=})_{:=\geq 0}(z_s)_{s\geq 0}$ with independent increments. Let this natural filtration be enlarged by its right-continuous completion, i.e., it satisfies the usual conditions of completeness and right continuity [KS91, p. 48]. Consider also an auxiliary subfiltration $\mathbb{F}_t := (\mathcal{F}_{t,s})_{s\geq 0}$, where $\mathcal{F}_{t,s}$ is the $\mathbb{P}$-completion of $\sigma\big(z_{r\vee t} - z_t, r \in [0, s]\big)$. It is obvious to observe that any $\mathcal{F}_{t,s}$-random variable is independent of $\mathcal{F}_t$, $\mathcal{F}_{t,s} \subseteq \mathcal{F}_s$ with equality in case of $t = 0$, and for $s \leq t$, $\mathcal{F}_{t,s}$ is the trivial $\sigma-$algebra.

The object of our study is an $\mathbb{R}^d$-valued controlled random process $\big(X_s^{t,x;\boldsymbol{u}}\big)_{s\geq t}$, initialized at (t, x) under the control policy $\boldsymbol{u} \in \mathcal{U}_t$, where $\mathcal{U}_t$ is the set of admissible policies at time t. Since the precise class of admissible policies does not play a role until Section IV we defer the formal definition of these until then. Let $T > 0$ be a fixed time horizon, and let $\mathbb{S} := [0, T] \times \mathbb{R}^d$. We assume that for every $(t, x) \in \mathbb{S}$ and $\boldsymbol{u} \in \mathcal{U}_t$, the process $\big(X_s^{t,x;\boldsymbol{u}}\big)_{s\geq t}$ is $\mathbb{F}_t$-adapted process whose sample paths are right continuous with left limits. We denote by $\mathcal{T}$ the collection of all $\mathbb{F}$-stopping times; for $\tau_1, \tau_2 \in \mathcal{T}$ with $\tau_1 \leq \tau_2$ $\mathbb{P}$-a.s. we let the subset $\mathcal{T}_{[\tau_1,\tau_2]}$ denote the collection of all $\mathbb{F}_{\tau_1}$-stopping times τ such that $\tau_1 \leq \tau \leq \tau_2$ $\mathbb{P}$-a.s. Measurability on $\mathbb{R}^d$ will always refer to Borel-measurability, and $\mathfrak{B}(\mathbb{A})$ stands for the Borel σ-algebra on a topological space $\mathbb{A}$. Throughout this chapter all the (in)equalities between random variables are understood in almost sure sense.

Given sets $(W_i, G_i) \in \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d)$ for $i \in \{1, \cdots, n\}$, we are interested in a set of initial conditions $(t, x) \in \mathbb{S}$ such that there exists an admissible strategy $\boldsymbol{u} \in \mathcal{U}_t$ steering the process $X_\cdot^{t,x;\boldsymbol{u}}$ through $(W_i)_{i=1}^n$ ("way point" sets) while visiting $(G_i)_{i=1}^n$ ("goal" sets) in a pre-assigned order. One may pose this objective from different perspectives based on different time scheduling for the excursions between the sets. We formally introduce some of these notions which will be addressed throughout this chapter.

Definition 3.2.1 (Motion Planning Events). *Consider a fixed initial condition $(t, x) \in \mathbb{S}$ and admissible policy $\boldsymbol{u} \in \mathcal{U}_t$. Given a sequence of pairs $(W_i, G_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d)$ and horizon*

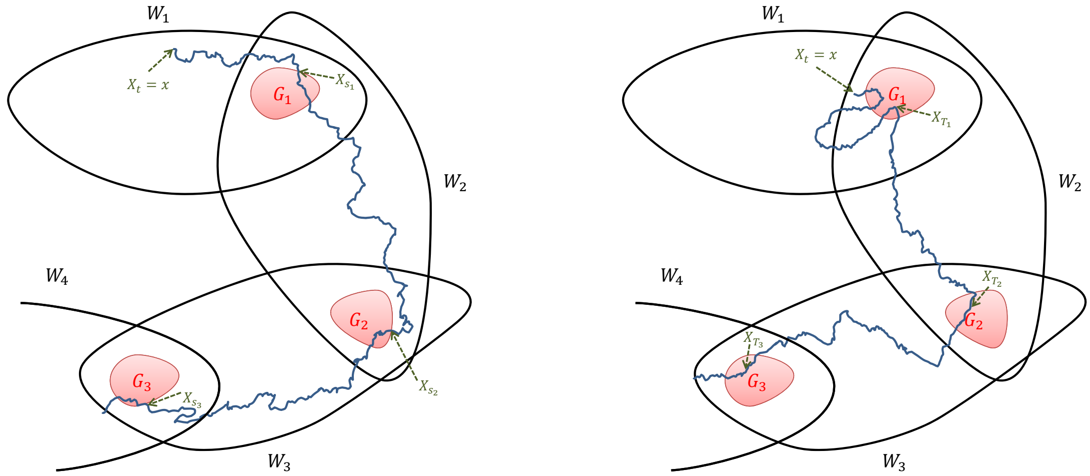

(a) A sample path satisfying the first three phases of the specification in the sense of (3.1a)

(b) A sample path satisfying the first three phases of the specification in the sense of (3.1b)

Figure 3.1: Sample paths of the process $X_\cdot^{t,x;\boldsymbol{u}}$ for a fix policy $\boldsymbol{u} \in \mathcal{U}_t$

times $(T_i)_{i=1}^n \subset [t,T]$*, we introduce the following* ***motion planning events:***

$$\Big\{X_\cdot^{t,x;\boldsymbol{u}} \models \big[(W_1 \rightsquigarrow G_1) \circ \cdots \circ (W_n \rightsquigarrow G_n)\big]_{\le T}\Big\} \coloneqq \tag{3.1a}$$
$$\Big\{\exists (s_i)_{i=1}^n \subset [t,T] \;\Big|\; X_{s_i}^{t,x;\boldsymbol{u}} \in G_i \text{ and } X_r^{t,x;\boldsymbol{u}} \in W_i \setminus G_i, \;\forall r \in [s_{i-1}, s_i[, \quad \forall i \le n\Big\},$$
$$\Big\{X_\cdot^{t,x;\boldsymbol{u}} \models (W_1 \xrightarrow{T_1} G_1) \circ \cdots \circ (W_n \xrightarrow{T_n} G_n)\Big\} \coloneqq \tag{3.1b}$$
$$\Big\{X_{T_i}^{t,x;\boldsymbol{u}} \in G_i \text{ and } X_r^{t,x;\boldsymbol{u}} \in W_i, \quad \forall r \in [T_{i-1}, T_i], \quad \forall i \le n\Big\},$$

where $s_0 = T_0 \coloneqq t$*.*

The set in (3.1a), roughly speaking, contains the events in the underlying probability space that the trajectory $X_\cdot^{t,x;\boldsymbol{u}}$, initialized at $(t,x) \in \mathbb{S}$ and controlled via $\boldsymbol{u} \in \mathcal{U}_t$, succeeds in visiting $(G_i)_{i=1}^n$ in a certain order, while the entire duration between the two visits to G_{i-1} and G_i is spent in W_i, all within the time horizon T. In other words, the journey from G_{i-1} to the next destination G_i must belong to W_i for all i. Figure 3.1(a) depicts a sample path that successfully contributes to the first three phases of the excursion in the sense of (3.1a). In the case of (3.1b), the set of paths is usually more restricted in comparison to (3.1a). Indeed, not only is the trajectory confined to W_i on the way between G_{i-1} and G_i, but also there is a time schedule $(T_i)_{i=1}^n$ that a priori forces the process to be at the goal sets G_i at the specific times $(T_i)_{i=1}^n$. Figure 3.1(b) demonstrates one sample path in which the first three phases of the excursion are successfully fulfilled.

Note that once a trajectory belonging to the set in (3.1a) visits G_i for the first time, it is required to remain in W_{i+1} until the next goal G_{i+1} is reached, whereas a trajectory belonging to the set in definition (3.1b) may visit the destination G_i several times, while staying in W_i until the intermediate time schedule T_i. The only requirement, in contrast to (3.1a), is to confine the trajectory to be at G_i at the time T_i. As an illustration, one can easily inspect that the sample path in Figure 3.1(b) indeed violates the requirements of the definition (3.1a) as it leaves W_2 after it visits G_1 for the first time. In other words, the definition (3.1a) changes the admissible way set W_i to W_{i+1} *immediately after* the trajectory visits G_i, while the definition (3.1b) only changes the admissible way set *only after* the intermediate time T_i irrespective of whether the trajectory visits G_i prior to T_i.

For simplicity we may impose the following assumptions:

Assumption 3.2.2. *We stipulate that*

a. The sets $(G_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d)$ *are closed.*

b. The sets $(W_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d)$ *are open.*

Concerning Assumption 3.2.2.a., if G_i is not closed, then it is not difficult to see that there could be some continuous transitions through the boundary of G_i that are not admissible in view of the definition (3.1a) since the trajectory must reside in $W_i \setminus G_i$ for the whole interval $[s_{i-1}, s_i[$ and just hit G_i at the time s_i. Notice that this is not the case for the definition (3.1b) since the trajectory only visits the sets G_i at the specific times T_i while any continuous transition and maneuver inside G_i are allowed. Assumption 3.2.2.b. is rather technical and required for the analysis employed in the subsequent sections.

The events introduced in Definition 3.2.1 depend, of course, on the control policy $\boldsymbol{u} \in \mathcal{U}$ and initial condition $(t, x) \in \mathbb{S}$. The main objective of this chapter is to determine the set of initial conditions $x \in \mathbb{R}^d$ such that there exists an admissible policy $\boldsymbol{u}$ where the probability of the motion planning events is higher than a certain threshold. Let us formally introduce these sets as follows:

Definition 3.2.3 (Motion Planning Initial Condition Set)**.** *Consider a fixed initial time* $t \in [0, T]$*. Given a sequence of set pairs* $(W_i, G_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d)$ *and horizon times* $(T_i)_{i=1}^n \subset [t, T]$*, we define the following* ***motion planning initial condition sets****:*

$$\mathrm{MP}\big(t, p;\ (W_i, G_i)_{i=1}^n, T\big) := \tag{3.2a}$$
$$\Big\{x \in \mathbb{R}^d \ \big|\ \exists \boldsymbol{u} \in \mathcal{U}_t \ :\ \mathbb{P}\big\{X_{\cdot}^{t,x;\boldsymbol{u}} \models \big[(W_1 \rightsquigarrow G_1) \circ \cdots \circ (W_n \rightsquigarrow G_n)\big]_{\leq T}\big\} > p\Big\},$$
$$\widetilde{\mathrm{MP}}\big(t, p;\ (W_i, G_i)_{i=1}^n, (T_i)_{i=1}^n\big) := \tag{3.2b}$$
$$\Big\{x \in \mathbb{R}^d \ \big|\ \exists \boldsymbol{u} \in \mathcal{U}_t \ :\ \mathbb{P}\big\{X_{\cdot}^{t,x;\boldsymbol{u}} \models (W_1 \xrightarrow{T_1} G_1) \circ \cdots \circ (W_n \xrightarrow{T_n} G_n)\big\} > p\Big\}.$$

Remark 3.2.4 (Stochastic Reach-Avoid Problem)**.** *The motion planning scenarios for only two sets* (W_1, G_1) *basically reduce to the basic reach-avoid maneuver studied in Chapter 2 by setting the reach set to* G_1 *and the avoid set to* $\mathbb{R}^d \setminus W_1$*. See also [GLQ06, ML11] for the corresponding deterministic and [SL10] for the corresponding discrete time stochastic reach-avoid problems.*

Remark 3.2.5 (Mixed Motion Planning Events)**.** *One may also consider an event that consists of a mixture of the events in (4.28), e.g.,* $(W_1 \rightsquigarrow G_1)_{\leq T_1} \circ (W_2 \xrightarrow{T_2} G_2)$*. Following essentially* *the same analytical techniques as the ones proposed in the subsequent sections, one can also address these mixed motion planning objectives. We shall provide an example of this nature in Section 3.6.*

Remark 3.2.6 (Time-varying Goal and Way Point Sets)**.** *In Definition 3.2.3 the motion planning objective is introduced in terms of stationary (time-independent) goal and way point sets. However, note that one can always augment the state space with time, and introduce a new stochastic process $Y_t := [X_t^\intercal, t]^\intercal$. Therefore, a motion planning concerning moving sets for X_t can be viewed as a motion planning with stationary sets for the process Y_t.*

3.3 Connection to Optimal Control Problems

In this section we establish a connection between stochastic motion planning initial condition sets MP and $\widetilde{\mathrm{MP}}$ of Definition 3.2.3 and a class of stochastic optimal control problems involving stopping times. First, given a sequence of sets we introduce a sequence of random times that corresponds to the times that the process $X_\cdot^{t,x;\boldsymbol{u}}$ exits each set in the sequence for the first time.

Definition 3.3.1 (Sequential Exit-Time)**.** *Given an initial condition $(t,x) \in \mathbb{S}$ and a sequence of measurable sets $(A_i)_{i=k}^n \subset \mathfrak{B}(\mathbb{R}^d)$, the sequence of random times $\big(\Theta_i^{A_{k:n}}\big)_{i=k}^n$ defined*[1] *by*

$$\Theta_i^{A_{k:n}}(t,x) := \inf\big\{r \geq \Theta_{i-1}^{A_{k:n}}(t,x) \ : \ X_r^{t,x;\boldsymbol{u}} \notin A_i\big\}, \qquad \Theta_{k-1}^{A_{k:n}}(t,x) := t,$$

is called the ***sequential exit-time*** *through the set A_k to A_n.*

Note that the sequential exit-time $\Theta_i^{A_{k:n}}$ depends on the control policy $\boldsymbol{u}$ in addition to the initial condition (t,x), but here and later in the sequel we shall suppress this dependence. For notational simplicity, we may also drop (t,x) in the subsequent sections.

In Figure 3.2 a sample path of the process $X_\cdot^{t,x;\boldsymbol{u}}$ along with the sequential exit-times $(\Theta_i^{A_{k:3}})_{i=k}^n$ is depicted for different $k \in \{1,2,3\}$. Note that since the initial condition x does not belongs to A_3, the first exit-time of the set A_3 is indeed the start time t, i.e., $\Theta_3^{A_{3:3}} = t$. Let us highlight the difference between stopping times $\Theta_2^{A_{1:3}}$ and $\Theta_2^{A_{2:3}}$. The former is the first exit-time of the set A_2 after the time that the process leaves A_1, whereas the latter is the first exit-time of the set A_2 from the very beginning. In Section 3.5 we shall see that these differences will lead to different definitions of value functions in order to derive a dynamic programming argument.

The following lemma shows that the sequential stopping times are indeed well defined.

Lemma 3.3.2 (Measurability)**.** *Consider a sequence of $(A_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d)$ and $(t,x) \in \mathbb{S}$. The sequential exit-time $\Theta_i^{A_{1:n}}(t,x)$ is an $\mathbb{F}_t$-stopping time for all $i \in \{1,\cdots,n\}$, i.e., $\big\{\Theta_i^{A_{1:n}}(t,x) \leq s\big\} \in \mathcal{F}_{t,s}$ for all $s \geq 0$.*

Proof. See Appendix 3.8.1. □

[1]By convention, $\inf \emptyset = \infty$.

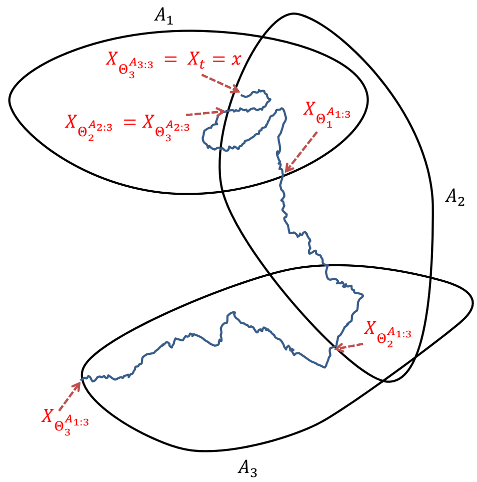

Figure 3.2: Sequential exit-times of a sample path through the sets $(A_i)_{i=k}^3$ for different values of k

Given $(W_i, G_i, T_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d) \times [t, T]$, we introduce two value functions $V, \widetilde{V} : \mathbb{S} \to [0, 1]$ defined by

$$V(t,x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\left[\prod_{i=1}^n \mathbb{1}_{G_i}\big(X_{\eta_i}^{t,x;\boldsymbol{u}}\big)\right], \qquad \eta_i := \Theta_i^{B_{1:n}} \wedge T, \quad B_i := W_i \setminus G_i, \tag{3.3a}$$

$$\widetilde{V}(t,x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\left[\prod_{i=1}^n \mathbb{1}_{G_i \cap W_i}\big(X_{\widetilde{\eta}_i}^{t,x;\boldsymbol{u}}\big)\right], \qquad \widetilde{\eta}_i := \Theta_i^{W_{1:n}} \wedge T_i, \tag{3.3b}$$

where $\Theta_i^{W_{1:n}}, \Theta_i^{B_{1:n}}$ are the sequential exit-times in the sense of Definition 3.3.1. Figure 3.3(a) and 3.3(b) illustrate the sequential exit-times corresponding to the sets B_i and W_i, respectively. The main result of this section, Theorem 3.3.3 below, establishes a connection from the sets $\mathrm{MP}, \widetilde{\mathrm{MP}}$ and superlevel sets of the value functions V and $\widetilde{V}$.

Theorem 3.3.3. *Fix a probability level $p \in [0,1]$, a sequence of set pairs $(W_i, G_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d)$, an initial time $t \in [0, T]$, and intermediate times $(T_i)_{i=1}^n \subset [t, T]$. Then,*

$$\mathrm{MP}\big(t, p; (W_i, G_i)_{i=0}^n, T\big) = \big\{x \in \mathbb{R}^d \;\big|\; V(t,x) > p\big\}. \tag{3.4}$$

Moreover, suppose Assumption 3.2.2.b. holds. Then,

$$\widetilde{\mathrm{MP}}\big(t, p; (W_i, G_i)_{i=0}^n, (T_i)_{i=1}^n\big) = \big\{x \in \mathbb{R}^d \;\big|\; \widetilde{V}(t,x) > p\big\}, \tag{3.5}$$

where the value functions V and $\widetilde{V}$ are as defined in (3.3).

Proof. See Appendix 3.8.1. ∎

Intuitively speaking, observe that the value functions (3.3) consist of a sequence of indicator functions, where the reward is 1 when the corresponding phase (i.e., reaching G_i while staying in W_i) of motion planning is fulfilled, while the reward is 0 if it fails. Let us also highlight that the difference between the time schedule between the two motion planning problems in

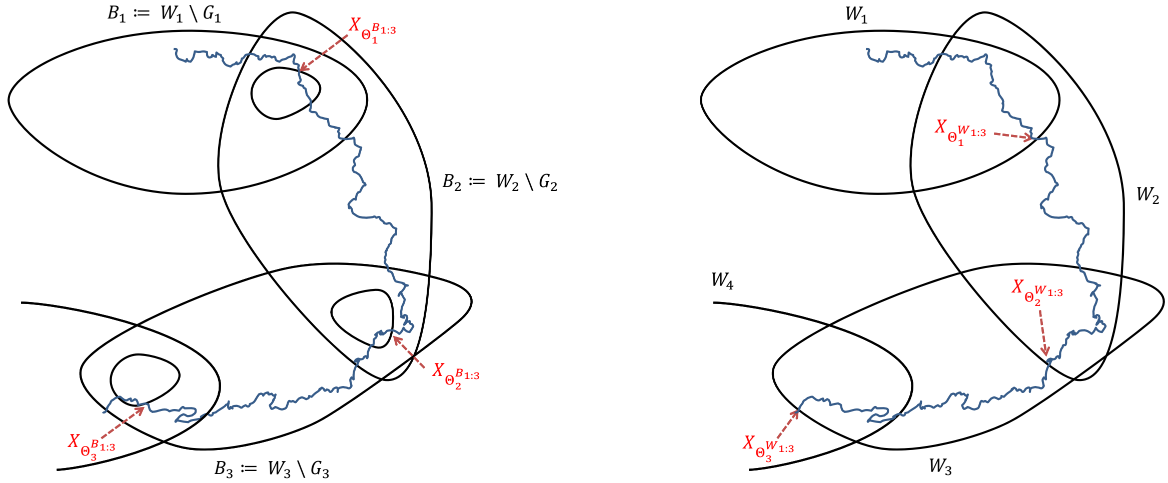

(a) Sequential exit-times related to motion planning event (3.1a)

(b) Sequential exit-times related to motion planning event (3.1b)

Figure 3.3: Sequential exit-times corresponding to different motion planning events as introduced in (4.28)

(3.3) is captured via the stopping times η_i and $\widetilde{\eta}_i$: the former refers to the first time to leave W_i or hit G_i before T, and the latter only considers the exit time from W_i prior to T_i. Hence, the product of the indicators evaluates to 1 if and only if the entire journey comprising n phases is successfully accomplished. In this light, taking expectations yields the probability of the desired event, and the supremum over admissible policies leads to the assertion that there exists a policy for which the desired properties hold.

3.4 Dynamic Programming Principle

The objective of this section is to derive a DPP for the value functions V and $\widetilde{V}$ introduced in (3.3). The DPP provides a bridge between the theoretical characterization of the solution to our motion planning problem through value functions (Section 3.3) and explicit characterizations of these value functions using, for example, PDEs (Section 3.5), which can then be used to solve the original problem numerically.

Let $(T_i)_{i=1}^n \subset [0, T]$ be a sequence of times, $(A_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d)$ be a sequence of open sets, and $\ell_i : \mathbb{R}^r \to \mathbb{R}$ for $i = 1, \cdots, n$ be a sequence of measurable and bounded payoff functions. We define the sequence of value functions $V_k : [0, T] \times \mathbb{R}^d \to \mathbb{R}^d$ for each $k \in \{1, \ldots, n\}$ as

$$V_k(t,x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\Big[\prod_{i=k}^{n} \ell_i\big(X_{\tau_i^k}^{t,x;\boldsymbol{u}}\big)\Big], \qquad \tau_i^k(t,x) := \Theta_i^{A_{k:n}}(t,x) \wedge T_i, \quad i \in \{k, \cdots, n\}, \tag{3.6}$$

where the stopping times $(\Theta_i^{W_{k:n}})_{i=k}^n$ are sequential exit-times in the sense of Definition 3.3.1. Recall that the sequential exit-times of V_k correspond to an excursion through the sets $(A_i)_{i=k}^n$ irrespective of the first $(k-1)$ sets. It is straightforward to observe that the value function

V (resp. $\widetilde{V}$) in (3.3) is a particular case of the value function V_1 defined as in (3.6) when $A_i := W_i \setminus G_i$ (resp. $A_i := W_i$), $\ell_i := \mathbb{1}_{G_i}$ (resp. $\ell_i := \mathbb{1}_{G_i \cap W_i}$), and $T_i := T$.

Given a metric space $\mathbb{A}$ and function $f : \mathbb{A} \longrightarrow \mathbb{R}$, the lower and upper semicontinuous envelopes of f are defined, respectively, as

$$f_*(x) := \liminf_{x' \to x} f(x'), \qquad f^*(x) := \limsup_{x' \to x} f(x').$$

We denote by $\mathrm{USC}(\mathbb{A})$ and $\mathrm{LSC}(\mathbb{A})$ the collection of all upper-semicontinuous and lower-semicontinuous functions from $\mathbb{A}$ to $\mathbb{R}$, respectively. To state the main result of this section, Theorem 3.4.3 below, some technical definitions and assumptions concerning the stochastic processes $X_\cdot^{t,x;\boldsymbol{u}}$, admissible strategies $\mathcal{U}_t$, and the payoff functions ℓ_i, are needed:

Assumption 3.4.1. *For all $(t,x) \in \mathbb{S}$, $\theta \in \mathcal{T}_{[t,T]}$, and $\boldsymbol{u}, \boldsymbol{v} \in \mathcal{U}_t$, we stipulate the following assumptions*

a. ***Admissible control policies:***
$\mathcal{U}_t$ is the set of $\mathbb{F}_t$-progressively measurable processes with values in a given control set. That is, the value of $\boldsymbol{u} := (u_s)_{s\in[0,T]}$ at time s can be viewed as a measurable mapping $(z_{r\vee t} - z_t)_{r\in[0,s]} \mapsto u_s$ for all $s \in [0,T]$, see [KS91, Def. 1.11, p. 4] for the details.

b. ***Stochastic process:***

   i. Causality: *If $\mathbb{1}_{[t,\theta]}\boldsymbol{u} = \mathbb{1}_{[t,\theta]}\boldsymbol{v}$, then we have $\mathbb{1}_{[t,\theta]}X_\cdot^{t,x;\boldsymbol{u}} = \mathbb{1}_{[t,\theta]}X_\cdot^{t,x;\boldsymbol{v}}$.*

   ii. Strong Markov property: *For each $\omega \in \Omega$ and the sample path $(z_r)_{r\in[0,\theta(\omega)]}$ up to the stopping time θ, let the random policy $\boldsymbol{u}_\theta \in \mathcal{U}_{\theta(\omega)}$ be the mapping $(z_{\cdot\vee\theta(\omega)} - z_{\theta(\omega)}) \mapsto \boldsymbol{u}(z_{\cdot\wedge\theta(\omega)} + z_{\cdot\vee\theta(\omega)} - z_{\theta(\omega)}) =: \boldsymbol{u}_\theta$.[2] Then,*

   $$\mathbb{E}\Big[\ell\big(X_{\theta+s}^{t,x;\boldsymbol{u}}\big) \;\Big|\; \mathcal{F}_\theta\Big] = \mathbb{E}\Big[\ell\big(X_{\theta+s}^{\bar{t},\bar{x};\bar{\boldsymbol{u}}}\big) \;\Big|\; \bar{t} = \theta, \bar{x} = X_\theta^{t,x;\boldsymbol{u}}, \bar{\boldsymbol{u}} = \boldsymbol{u}_\theta\Big], \qquad \mathbb{P}\text{-a.s.}$$

   for all bounded measurable functions $\ell : \mathbb{R}^d \to \mathbb{R}$ and $s \geq 0$.

   iii. Continuity of the exit-times: *Given initial condition $(t_0, x_0) \in \mathbb{S}$, for all $k \in \{1, \cdots, n\}$ and $i \in \{k, \cdots, n\}$ the stochastic mapping $(t,x) \mapsto X_{\tau_i^k(t,x)}^{t,x;\boldsymbol{u}}$ is $\mathbb{P}$-a.s. continuous at (t_0, x_0) where the stopping times τ_i^k are defined as in (3.6).*

c. ***Payoff functions:***
$(\ell_i)_{i=1}^n$ are lower semicontinuous, i.e., $\ell_i \in \mathrm{LSC}(\mathbb{R}^d)$ for all $i \leq n$.

Remark 3.4.2. *Some remarks on the above assumptions are in order:*

- *Assumption 3.4.1.a. implies that the admissible policies $\boldsymbol{u} \in \mathcal{U}_t$ take action at time t independent of future information arriving at $s > t$. This is known as the* non-anticipative *strategy and is a standard assumption in stochastic optimal control [Bor05].*

[2] Notice that $z_\cdot \equiv z_{\cdot\wedge\theta(\omega)} + z_{\cdot\vee\theta(\omega)} - z_{\theta(\omega)}$. Thus, the randomness of $\boldsymbol{u}_\theta$ is referred to the term $z_{\cdot\wedge\theta(\omega)}$.

- *Assumption 3.4.1.b. imposes three constraints on the process $X_\cdot^{t,x;\boldsymbol{u}}$ defined on the prescribed probability space: i) causality of the solution processes for a given admissible policy ii) strong Markov property iii) continuity of exit-time. The causality property is always satisfied in practical applications; uniqueness of the solution process $X_\cdot^{t,x;\boldsymbol{u}}$ under any admissible control process $\boldsymbol{u}$ guarantees it. The class of Strong Markov processes is fairly large; for instance, it contains the solution of SDEs under some mild assumptions on the drift and diffusion terms [Kry09, Thm. 2.9.4]. The almost sure continuity of the exit-time with respect to the initial condition of the process is the only potentially restrictive part of the assumptions. Note that this condition does not always hold even for deterministic processes with continuous trajectories. One may need to impose conditions on the process and possibly the sets involved in motion planning in order to satisfy continuity of the mapping $(t,x) \mapsto X^{t,x;\boldsymbol{u}}_{\tau_i^k(t,x)}$ at the given initial condition with probability one. We shall elaborate on this issue and its ramifications for a class of diffusion processes in Section 3.5.*

- *Assumption 3.4.1.c. imposes a fairly standard assumption on the payoff functions. In case ℓ_i is the indicator function of a given set, for example in (3.3), this assumption requires the set to be open. This issue will be addressed in more details in Subsection 3.5.3, in particular how it can be reconciled with Assumption 3.2.2.a..*

Let function $J_k : \mathbb{S} \times \mathcal{U}_0 \to \mathbb{R}$ be

$$J_k(t,x;\boldsymbol{u}) \coloneqq \mathbb{E}\Big[\prod_{i=k}^{n} \ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big)\Big],$$

where $\big(\tau_i^k\big)_{i=k}^n$ are as defined in (3.6).

The following Theorem, the main result of this section, establishes a dynamic programming argument for the value function V_k in terms of the "successor" value functions $(V_j)_{j=k+1}^n$, all defined as in (3.6). For ease of notation, we shall introduce deterministic times $\tau_{k-1}^k, \tau_{n+1}^k$, and a trivial constant value function V_{n+1}.

Theorem 3.4.3 (Dynamic Programming Principle)**.** *Consider the value functions $(V_j)_{j=1}^n$ and the sequential stopping times $(\tau_j^k)_{j=k}^n$ introduced in (3.6) where $k \in \{1, \cdots, n\}$. Under Assumptions 3.4.1, for all $(t,x) \in \mathbb{S}$ and family of stopping times $\{\theta^{\boldsymbol{u}}, \boldsymbol{u} \in \mathcal{U}_t\} \subset \mathcal{T}_{[t,T]}$, we have*

$$V_k(t,x) \le \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\sum_{j=k}^{n+1} \mathbb{1}_{\{\tau_{j-1}^k \le \theta^{\boldsymbol{u}} < \tau_j^k\}} V_j^*\big(\theta^{\boldsymbol{u}}, X^{t,x;\boldsymbol{u}}_{\theta^{\boldsymbol{u}}}\big) \prod_{i=k}^{j-1} \ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big)\Big], \tag{3.7a}$$

$$V_k(t,x) \ge \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\sum_{j=k}^{n+1} \mathbb{1}_{\{\tau_{j-1}^k \le \theta^{\boldsymbol{u}} < \tau_j^k\}} V_{j*}\big(\theta^{\boldsymbol{u}}, X^{t,x;\boldsymbol{u}}_{\theta^{\boldsymbol{u}}}\big) \prod_{i=k}^{j-1} \ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big)\Big], \tag{3.7b}$$

where V_j^* and V_{j*} are, respectively, the upper and the lower semicontinuous envelope of V_j, $\tau_{k-1}^k \coloneqq t$, $V_{n+1} \equiv 1$, and τ_{n+1}^k is any constant time strictly greater than T, say $\tau_{n+1}^k \coloneqq T+1$.

Proof. See Appendix 3.8.2. □

In our context the DPP proposed in Theorem 3.4.3 allows us to characterize the value function (3.6) through a sequence of value functions $(V_j)_{j=k+1}^n$. That is, (3.7a) and (3.7b) impose mutual constraints on a value function and subsequent value functions in the sequence. Moreover, the last function in the sequence is fixed to a constant by construction. Therefore, assuming that an algorithm to sequentially solve for these mutual constraints can be established, one could in principle use it to compute all value functions in the sequence and solve the original motion planning problem. In Section 3.5 we show how, for a class of controlled diffusion processes, the constraints imposed by (3.7) reduce to PDEs that the value functions need to satisfy. This enables the use of numerical PDE solution algorithms for this purpose.

Remark 3.4.4 (Measurability)**.** *Theorem 3.4.3 introduces DPP's in a weaker sense than the standard DPP in stochastic optimal control problems [FS06]. Namely, one does not need to verify the measurability of the value functions V_k in* (3.3) *so as to apply the DPP's. Notice that in general this measurability issue is non-trivial due to the supremum operation running over possibly uncountably many policies.*

3.5 The Case of Controlled Diffusions

In this section we come to the last step in our construction. We demonstrate how the DPP derived in Section 3.4, in the context of controlled diffusion processes, gives rise to a series of PDE's. Each PDE is understood in the discontinuous viscosity sense with boundary conditions in both Dirichlet (pointwise) and viscosity senses. This paves the way for using PDE numerical solvers to numerically approximate the solution of our original motion planning problem for specific examples. We demonstrate an instance of such an example in Section 3.6.

We first introduce formally the standard probability space setup for SDEs, then proceed with some preliminaries to ensure that the requirements of the proposed DPP, Assumptions 3.4.1, hold. The section consists of subsections concerning PDE derivation and boundary conditions along with further discussions on how to deploy existing PDE solvers to numerically compute our PDE characterization.

Let Ω be $\mathcal{C}\big([0,T],\mathbb{R}^{z_d}\big)$, the set of continuous functions from $[0,T]$ into $\mathbb{R}^{z_d}$, and let $(z_t)_{t\geq 0}$ be the canonical process, i.e., $z_t(\omega) := \omega_t$. We consider $\mathbb{P}$ as the Wiener measure on the filtered probability space $(\Omega,\mathcal{F},\mathbb{F})$, where $\mathbb{F}$ is the smallest right continuous filtration on Ω to which the process $(z_t)_{t\geq 0}$ is adapted. Let us recall that $\mathbb{F}_t := (\mathcal{F}_{t,s})_{s\geq 0}$ is the auxiliary subfiltration defined as $\mathcal{F}_{t,s} := \sigma\big(z_{r\vee t} - z_t, r\in[0,s]\big)$. Let $\mathbb{U}\subset\mathbb{R}^{d_u}$ be a control set, and $\mathcal{U}_t$ denote the set of all $\mathbb{F}_t$- progressively measurable mappings into $\mathbb{U}$. For every $\boldsymbol{u} = (u_t)_{t\geq 0}$ we consider the $\mathbb{R}^d$-valued SDE[3]

$$\mathrm{d}X_s = f(X_s,u_s)\,\mathrm{d}s + \sigma(X_s,u_s)\,\mathrm{d}W_s, \qquad X_t = x, \quad s\geq t, \tag{3.8}$$

where $f:\mathbb{R}^d\times\mathbb{U}\to\mathbb{R}^d$ and $\sigma:\mathbb{R}^d\times\mathbb{U}\to\mathbb{R}^{d\times d_z}$ are measurable functions, and $W_s := z_s$ is the canonical process.

[3]We slightly abuse notation and earlier used σ for the sigma algebra as well. However, it will be always clear from the context to which σ we refer.

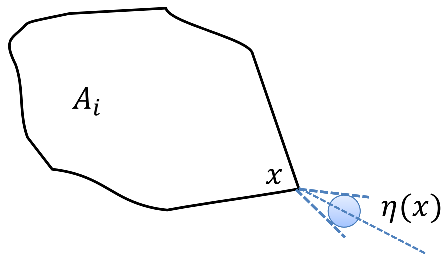

(a) Exterior cone condition holds at every point of the boundary.

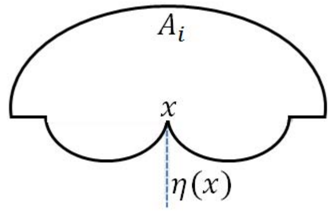

(b) exterior cone condition fails at the point x—the only possible exterior cone at x is a line.

Figure 3.4: Exterior cone condition of the boundary

Assumption 3.5.1. *We stipulate that*

a. $\mathbb{U} \subset \mathbb{R}^m$ *is compact;*

b. f *and* σ *are continuous and Lipschitz in first argument uniformly with respect to the second;*

c. *The diffusion term* σ *of the SDE* (3.8) *is uniformly non-degenerate, i.e., there exists* $\delta > 0$ *such that for all* $x \in \mathbb{R}^d$ *and* $u \in \mathbb{U}$, $\|\sigma(x,u)\sigma^\top(x,u)\| > \delta$.

It is well-known that under Assumptions 3.5.1.a. and 3.5.1.b. there exits a unique strong solution to the SDE (3.8) [Bor05]; let us denote it by $\left(X_s^{t,x;\boldsymbol{u}}\right)_{s\geq t}$. For future notational simplicity, we slightly modify the definition of $X_s^{t,x;\boldsymbol{u}}$, and extend it to the whole interval $[0,T]$ where $X_s^{t,x;\boldsymbol{u}} := x$ for all s in $[0,t]$.

In addition to Assumptions 3.5.1 on the SDE (3.8), we impose the following assumption on the motion planning sets that allows us to guarantee the continuity of sequential exit-times, as required for the DPP obtained in the preceding section.

Assumption 3.5.2 (Exterior Cone Condition)**.** *The open sets* $(A_i)_{i=1}^n$ *satisfy the following condition: for every* $i \in \{1, \cdots, n\}$, *there are positive constants* h, r *an* $\mathbb{R}^d$*-value bounded map* $\eta : A_i^c \to \mathbb{R}^d$ *such that*

$$\mathbb{B}_{rt}(x + \eta(x)t) \subset A_i^c \qquad \text{for all } x \in A_i^c \text{ and } t \in (0,h]$$

where $\mathbb{B}_r(x)$ *denotes an open ball centered at* x *and radius* r *and* A_i^c *stands for the complement of the set* A_i.

Remark 3.5.3 (Smooth Boundary)**.** *If the set* A_i *is bounded and its boundary* ∂A_i *is smooth, then Assumption 3.5.2 holds. Furthermore, boundaries with corners may also satisfy Assumption 3.5.2; Figure 3.4 depicts two different examples.*

3.5.1 A Sequential PDEs

In the context of SDEs, we show how the abstract DPP of Theorem 3.4.3 results in a sequence of PDEs, to be interpreted in the sense of discontinuous viscosity solutions; for the general theory of viscosity solutions we refer to [CIL92] and [FS06]. For numerical solutions to these PDEs, one also needs appropriate boundary conditions which is addressed in the next subsection.

To apply the proposed DPP, one has to make sure that Assumptions 3.4.1 are satisfied. As pointed out in Remark 3.4.2, the only nontrivial assumption in the context of SDEs is Assumption 3.4.1.b.iii. The following proposition addresses this issue, and allows us to employ the DPP of Theorem 3.4.3 for the main result of this subsection.

Proposition 3.5.4. *Consider the SDE (3.8) where Assumptions 3.5.1 hold. Suppose the open sets $(A_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d)$ satisfy the exterior cone condition in Assumption 3.5.2. Let $\big(\Theta_i^{A_{1:n}}\big)_{i=1}^n$ be the respective sequential exit-times as defined in Definition 3.3.1. Given intermediate times $(T_i)_{i=1}^n$ and control policy $\boldsymbol{u} \in \mathcal{U}_t$, for any $i \in \{1 \cdots, n\}$, initial condition $(t,x) \in \mathbb{S}$, and sequence of initial conditions $(t_m, x_m) \to (t,x)$, we have*

$$\lim_{m\to\infty} \tau_i(t_m, x_m) = \tau_i(t,x) \quad \mathbb{P}\text{-a.s.}, \qquad \tau_i(t,x) := \Theta_i^{A_{1:n}}(t,x) \wedge T_i.$$

As a consequence, the stochastic mapping $(t,x) \mapsto X^{t,x;\boldsymbol{u}}_{\tau_i(t,x)}$ is continuous with probability one, i.e., $\lim\limits_{m\to\infty} X^{t_m,x_m;\boldsymbol{u}}_{\tau_i(t_m,x_m)} = X^{t,x;\boldsymbol{u}}_{\tau_i(t,x)}$ $\mathbb{P}$-a.s. for all i.

Proof. See Appendix 3.8.3. □

Let us recall again the Dynkin Operator associated with the SDE (3.8), as also introduced in the Chapter 2.

Definition 3.5.5 (Dynkin Operator)**.** *Given $u \in \mathbb{U}$, we denote by $\mathcal{L}^u$ the Dynkin operator (also known as the infinitesimal generator) associated to the SDE (3.8) as*

$$\mathcal{L}^u\Phi(t,x) := \partial_t\Phi(t,x) + f(x,u).\partial_x\Phi(t,x) + \frac{1}{2}Tr[\sigma(x,u)\sigma^\top(x,u)\partial_x^2\Phi(t,x)],$$

where Φ is a real-valued function smooth on the interior of $\mathbb{S}$, with $\partial_t\Phi$ and $\partial_x\Phi$ denoting the partial derivatives with respect to t and x, respectively, and $\partial_x^2\Phi$ denoting the Hessian matrix with respect to x.

Theorem 3.5.6 is the main result of this subsection, which provides a characterization of the value functions V_k in terms of Dynkin operator in Definition 3.5.5 in the interior of the set of interest, i.e., $[0, T_k[\times A_k$. We refer to [Kal97, Thm. 17.23] for details on the above differential operator.

Theorem 3.5.6 (Dynamic Programming Equation)**.** *Consider the system (3.8), and suppose that Assumptions 3.5.1 hold. Let the value functions $V_k : \mathbb{S} \to \mathbb{R}^d$ be as defined in (3.6), where the sets $(A_i)_{i=1}^n$ satisfy Assumption 3.5.2, and the payoff functions $(\ell_i)_{i=1}^n$ are all lower semicontinuous. Then,*

∘ V_{k*} *is a viscosity supersolution of*

$$-\sup_{u\in\mathbb{U}} \mathcal{L}^u V_{k*}(t,x) \geq 0 \qquad \text{on} \quad [0,T_k[\times A_k;$$

∘ V_k^* *is a viscosity subsolution of*

$$-\sup_{u\in\mathbb{U}} \mathcal{L}^u V_k^*(t,x) \leq 0 \qquad \text{on} \quad [0,T_k[\times A_k.$$

Proof. The proof follows the same technique as in the proof of Theorem 2.4.10, which is briefly sketched in Appendix 3.8.3. □

3.5.2 Boundary Conditions

To numerically solve the PDE of Theorem 3.5.6, one needs boundary conditions on the complement of the set where the PDE is defined. This requirement is addressed in the following proposition.

Proposition 3.5.7 (Boundary Conditions)**.** *Suppose that the hypotheses of Theorem 3.5.6 hold. Then the value functions V_k introduced in (3.6) satisfy the following boundary value conditions:*

$$\textit{Dirichlet:} \qquad V_k(t,x) = V_{k+1}(t,x)\ell_k(x) \qquad \text{on} \quad [0,T_k]\times A_k^c \bigcup \{T_k\}\times\mathbb{R}^d \tag{3.9a}$$

$$\textit{Viscosity:} \qquad \begin{cases} \limsup\limits_{\substack{A_k\ni x'\to x\\ t'\uparrow t}} V_k(t',x') \leq V_{k+1}^*(t,x)\ell_k^*(x) & \text{on} \quad [0,T_k]\times\partial A_k \bigcup \{T_k\}\times\overline{A}_k \\ \liminf\limits_{\substack{A_k\ni x'\to x\\ t'\uparrow t}} V_k(t',x') \geq V_{k+1\,*}(t,x)\ell_k(x) & \text{on} \quad [0,T_k]\times\partial A_k \bigcup \{T_k\}\times\overline{A}_k \end{cases} \tag{3.9b}$$

Proof. See Appendix 3.8.3. □

Proposition 3.5.7 provides boundary condition for V_k in both Dirichlet (pointwise) and viscosity senses. The Dirichlet boundary condition (3.9a) is the one usually employed to numerically compute the solution via PDE solvers, whereas the viscosity boundary condition (3.9b) is required for theoretical support of the numerical schemes and comparison results.

Remark 3.5.8. *In the SDE setting, one can, without loss of generality, extend the class of admissible policies in the definition of V_k to $\mathcal{U}_0$, i.e., $V_k(t,x) = \sup_{\boldsymbol{u}\in\mathcal{U}_0} J_k(t,x;\boldsymbol{u})$; for a rigorous technology to prove this assertion see [BT11, Remark 5.2]. Thus, V_k is lower semicontinuous as it is a supremum over a fixed family of lower semicontinuous functions, see Lemma 3.8.2 in Appendix 3.8.3. In this light, one may argue that in the viscosity boundary condition (3.9b), the second assertion is subsumed by the Dirichlet boundary condition (3.9a).*

3.5.3 Discussion on Numerical Issues

For the class of controlled diffusion processes (3.8), Subsection 3.5.1 developed a PDE characterization of the value function V_k within the set $[0,T_k[\times A_k$ along with boundary conditions in terms of the successor value function V_{k+1} provided in Subsection 3.5.2. Since $V_{n+1} \equiv 1$, one can infer that Theorem 3.5.6 and Proposition 3.5.7 provide a series of PDE where the last one has known boundary condition, while the boundary conditions of earlier in the sequence are determined by the solution of subsequent PDE, i.e., V_{k+1} provides boundary conditions for the PDE corresponding to the value function V_k. Let us highlight once again that the basic motion planning maneuver involving only two sets is effectively the same as the first step of this series of PDEs and was studied in Chapter 2.

Before proceeding with numerical solutions, we need to properly address two technical concerns:

(i) On the one hand, for the definition (3.2a) we need to assume that the goal set G_i is closed so as to allow continuous transition into G_i; see Assumption 3.2.2.a. and the following discussion. On the other hand, in order to invoke the DPP argument of Section 3.4 and its consequent PDE in Subsection 3.5.1, we need to impose that the payoff functions $(\ell_i)_{i=1}^n$ are all lower semicontinuous; see Assumption 3.4.1.c. In the case of the value function V in (3.3a), this constraint results in $(G_i)_{i=1}^n$ all being open, which in general contradicts Assumption 3.2.2.a..

(ii) Most of the existing PDE solvers provide theoretical guarantees for *continuous* viscosity solutions, e.g., [Mit05]. Theorem 3.5.6, on the other hand, characterizes the solution to the motion planning problem in terms of *discontinuous* viscosity solutions. Therefore, it is a natural question whether we could employ any of available numerical methods to approximate the solution of our desired value function.

Let us initially highlight the following points: Concerning (i) the contradiction is not applicable for the motion planning initial set (3.2b) since the goal set G_i can be simply chosen to be open without confining the continuous transitions. Concerning (ii), we would like to stress that this discontinuous formulation is inevitable since the value functions defined in (3.3) are in general discontinuous, and any PDE approach has to rely on discontinuous versions.

To address the above concerns, we propose an ε-conservative but precise way of characterizing the motion planning initial set. Given $(W_i, G_i) \in \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d)$, let us construct a smaller goal set $G_i^\varepsilon \subset G_i$ such that $G_i^\varepsilon := \{x \in G_i |\, \mathrm{dist}(x, G_i^c) \geq \varepsilon\}$.[4] For sufficiently small $\varepsilon > 0$ one may observe that $W_i \setminus G_i^\varepsilon$ satisfies Assumption 3.5.2. Note that this is always possible if $W_i \setminus G_i$ satisfies Assumption 3.5.2 since one can simply take $\varepsilon < h/2$, where h is as defined in Assumption 3.5.2. Figure 3.5 depicts this situation.

Formally we define the payoff function $\ell_i^\varepsilon : \mathbb{R}^d \to \mathbb{R}$ as follows:

$$\ell_i^\varepsilon(x) := \left(1 - \frac{\mathrm{dist}(x, G_i^\varepsilon)}{\varepsilon}\right) \vee 0.$$

Replacing the goal sets G_i^ε and payoff functions ℓ_i^ε in (3.3a), we arrive at the value function

$$V^\varepsilon(t,x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\left[\prod_{i=1}^{n} \ell_i^\varepsilon\big(X_{\eta_i^\varepsilon}^{t,x;\boldsymbol{u}}\big)\right], \qquad \eta_i^\varepsilon := \Theta_i^{B_{1:n}^\varepsilon} \wedge T, \qquad B_i^\varepsilon := W_i \setminus G_i^\varepsilon.$$

[4] $\mathrm{dist}(x, A) := \inf_{y \in A} \|x - y\|$, where $\|\cdot\|$ stands for the Euclidean norm.

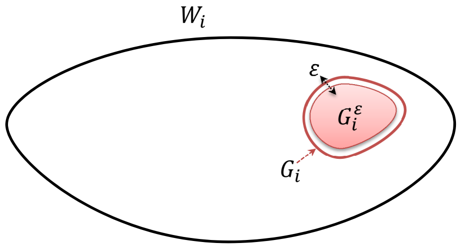

Figure 3.5: Construction of the sets G_i^ε from G_i as described in Subsection 3.5.3

It is straightforward to inspect that $V^\varepsilon \le V$ since $G_i^\varepsilon \subset G_i$. Moreover, with a similar technique as in Theorem 2.5.1 (see the proof in Section 2.8.3), one may show that $V(t,x) = \lim_{\varepsilon\downarrow 0} V^\varepsilon(t,x)$ on the set $(t,x) \in [t,T[\times\mathbb{R}^d$, which indicates that the approximation scheme can be arbitrarily precise. Note that the approximated payoff functions ℓ_i^ε are, by construction, Lipschitz continuous that in light of uniform continuity of the process, Lemma 3.8.4 in Appendix 3.8.3, leads to the continuity of the value function V^ε.[5] Hence, the discontinuous PDE characterization of Subsection 3.5.1 can be approximated arbitrarily closely in the continuous regime.

Let us recall that having reduced the motion planning problems to PDEs, numerical methods and computational algorithms exist to approximate its solution [Mit05]. In Section 3.6 we demonstrate how to use such methods to address practically relevant problems. In practice, such methods are effective for systems of relatively small dimension due to the curse of dimensionally. To alleviate this difficulty and extend the method to large problems, we can leverage on ADP [dFVR03, DFVR04, CRVRL06] or other advances in numerical mathematics, such as tensor trains [KO10, KS11]. The link between motion planning and the PDEs through DPP is precisely what allows us to capitalize on any such developments in the numerics.

3.6 Application to a Biological Switch Example

When modeling uncertainty in biochemical reactions, one often resorts to countable Markov chain models [Wil06] which describe the evolution of molecular numbers. Due to the Markov property of chemical reactions, one can track the time evolution of the probability distribution for molecular populations as a family of ordinary differential equations called the *chemical master equation* (CME) [AGA09, ESKPG05], also known as the forward Kolmogorov equation.

Though close to the physical reality, the CME is particularly difficult to work with analytically. One therefore typically employs different approximate solution methods, for example the Finite State Projection method [Kha] or the moment closure method [SH07]. Such approximation method resorts to approximating discrete molecule numbers by a continuum and capturing the stochasticity in their evolution through a stochastic differential equation. This stochastic continuous-time approximation is called the *chemical Langevin equation* or the *diffusion approximation*, see for example [Kha] and the reference therein. The Langevin approximation can be inaccurate for chemical species with low copy numbers; it may even assign a negative number

[5] This continuity result can, alternatively, be deduced via the comparison result of the viscosity characterization of Theorem 3.5.6 together with boundary conditions (3.9b) [CIL92].

to some molecular species. To circumvent this issue we assume here that the species of interest come in sufficiently high copy numbers to make the Langevin approximation reasonable.

Multistable biological systems are often encountered in nature [BSS01]. In this section we consider the following chemical Langevin formulation of a bistable two gene network:

$$\begin{cases} \mathrm{d}X_t = \big(f(Y_t, \boldsymbol{u}_x) - \mu_x X_t\big)\mathrm{d}t + \sqrt{f(Y_t, \boldsymbol{u}_x)}\mathrm{d}W_t^1 + \sqrt{\mu_x X_t}\mathrm{d}W_t^2, & X_0 = x_0 \\ \mathrm{d}Y_t = \big(g(X_t, \boldsymbol{u}_y) - \mu_y Y_t\big)\mathrm{d}t + \sqrt{g(X_t, \boldsymbol{u}_y)}\mathrm{d}W_t^3 + \sqrt{\mu_y Y_t}\mathrm{d}W_t^4, & Y_0 = y_0 \end{cases} \tag{3.10}$$

where X_t and Y_t are the concentration of the two repressor proteins with the respective degradation rates μ_x and μ_y; $(W_t^i)_{t\geq 0}$ are independent standard Brownian motion processes. Functions f and g are repression functions that describe the impact of each protein on the other's rate of synthesis controlled via some external inputs $\boldsymbol{u}_x$ and $\boldsymbol{u}_y$.

In the absence of exogenous control signals, the authors of [Che00] study sufficient conditions on the drifts f and g under which the system dynamic (3.10) without the diffusion term has two (or more) stable equilibria. In this case, system (3.10) can be viewed as a biological switch network. The theoretical results of [Che00] are also experimentally investigated in [GCC00] for a genetic toggle switch in *Escherichia coli*.

Here we consider the biological switch dynamics where the production rates of proteins are influenced by external control signals; experimental constructs that can be used to provide such inputs have recently been reported in the literature [MSSO+11]. The level of repression is described by a *Hill* function, which models cooperativity of binding as follows:

$$f(y, u) := \frac{\theta_1^{n_1} k_1}{y^{n_1} + \theta_1^{n_1}} u, \qquad g(x, u) := \frac{\theta_2^{n_2} k_2}{x^{n_2} + \theta_2^{n_2}} u,$$

where θ_i are the threshold of the production rate with respective exponents n_i, and k_i are the production scaling factors. The parameter u represents the role of external signals that affect the production rates, for which the control sets are $\mathbb{U}_x := [\underline{u}_x, \overline{u}_x]$ and $\mathbb{U}_y := [\underline{u}_y, \overline{u}_y]$. In this example we consider system (3.10) with the following parameters: $\theta_i = 40$, $\mu_i = 0.04$, $k_i = 4$ for both $i \in \{1, 2\}$, and exponents $n_1 = 4$, $n_2 = 6$. Figure 3.6(a) depicts the drift nullclines and the equilibria of the system. The equilibria z_a and z_c are stable, while z_b is the unstable one. We should remark that the "stable equilibrium" of SDE (3.10) is understood in the absence of the diffusion term as the noise may very well push the states from one stable equilibrium to another.

We first aim to steer the number of proteins toward a target set around the unstable equilibrium within a certain time horizon, say T_1, by synthesizing appropriate input signals $\boldsymbol{u}_x$ and $\boldsymbol{u}_y$. During this task we opt to avoid the region of attraction of the stable equilibria as well as low numbers for each protein; the latter justifies our Langevin model being well-posed in the region of interest. These target and avoid sets are denoted, respectively, by the closed sets B and A in Figure 3.6(b). In the second phase of the task, once the trajectory visits the target set B, it is required to keep the molecular populations within a slightly larger margin around the unstable equilibrium for some time, say T_2; Figure 3.6(b) depicts this maintenance margin by the open set C. In the context of reachability, the second phase is known as *viability* [Aub91, AP98].

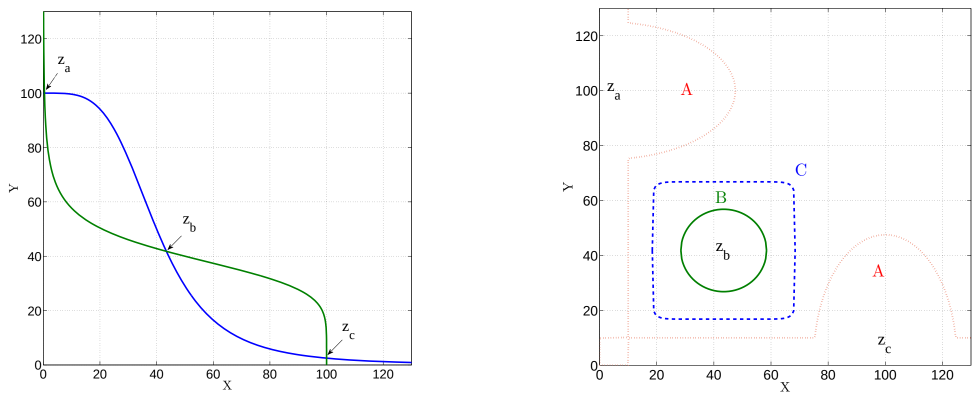

(a) Nullclines and equilibria of the drift of the SDE (3.10); x_a and x_b are stable, and x_c is unstable.

(b) The set A is an avoidance region contained in the region of attraction of the stable equilibria x_a and x_c, B is the target set around the unstable equilibrium x_b, and C is the maintenance margin.

Figure 3.6: State space of the biological switch (3.10) with desired motion planning sets.

In view of motion planning events introduced in Definition 3.2.1, the first phase of the path can be expressed as $(A^c \rightsquigarrow B)_{\le T_1}$, and the second phase as $(C \xrightarrow{T_2} C)$; see (4.28) for the detailed definitions of these symbols. By defining the joint process $Z^{t,z;\boldsymbol{u}}_{\cdot} := \big[X^{t,x;\boldsymbol{u}}_{\cdot}, Y^{t,y;\boldsymbol{u}}_{\cdot}\big]$, with the initial condition $z := [x,y]$ and controller $\boldsymbol{u} := [\boldsymbol{u}_x, \boldsymbol{u}_y]$, the desired excursion is a combination of the events studied in the preceding sections and, with a slight abuse of notation, can be expressed by

$$\Big\{Z^{t,z;\boldsymbol{u}}_{\cdot} \models (A^c \rightsquigarrow B)_{\le T_1} \circ (C \xrightarrow{T_2} C)\Big\}.$$

Though the desired path mixes the two events of Definition 3.2.1, one can still invoke the framework of Section 3.3 and introduce the following value functions:

$$V_1(t,z) := \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_B\big(Z^{t,z;\boldsymbol{u}}_{\tau_1^1}\big)\mathbb{1}_C\big(Z^{t,z;\boldsymbol{u}}_{\tau_2^1}\big)\Big], \tag{3.11a}$$

$$V_2(t,z) := \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_C\big(Z^{t,z;\boldsymbol{u}}_{\tau_2^2}\big)\Big], \tag{3.11b}$$

where τ_1^1 and τ_2^2 are defined in (3.6) with sets $A_1 := (A \cup B)^c$ and $A_2 := C$. We define the stopping time $\tau_2^1 := \Theta_2^{A_{1:2}} \wedge (\tau_1^1 + T_2)$ to address the concatenation of heterogeneous events in our specification..

The solution of the motion planning objective is the value function V_1 in (3.11a), which in view of Theorem 3.5.6 is characterized by the Dynkin PDE operator in the interior of $[0,T_1[\times(A\cup B)^c$. However, we first need to compute V_2 in (3.11b) to provide boundary conditions for V_1 according to

$$V_1(t,z) = \mathbb{1}_B(z)V_2(t,z), \qquad (t,z) \in [0,T_1] \times (A \cup B) \bigcup \{T_1\} \times \mathbb{R}^2. \tag{3.12}$$

It is straightforward to observe that the boundary condition for the value function V_2 is

$$V_2(t,z) = \mathbb{1}_C(z), \qquad (t,z) \in [0, T_1 + T_2] \times C^c \bigcup \{T_1 + T_2\} \times \mathbb{R}^2.$$

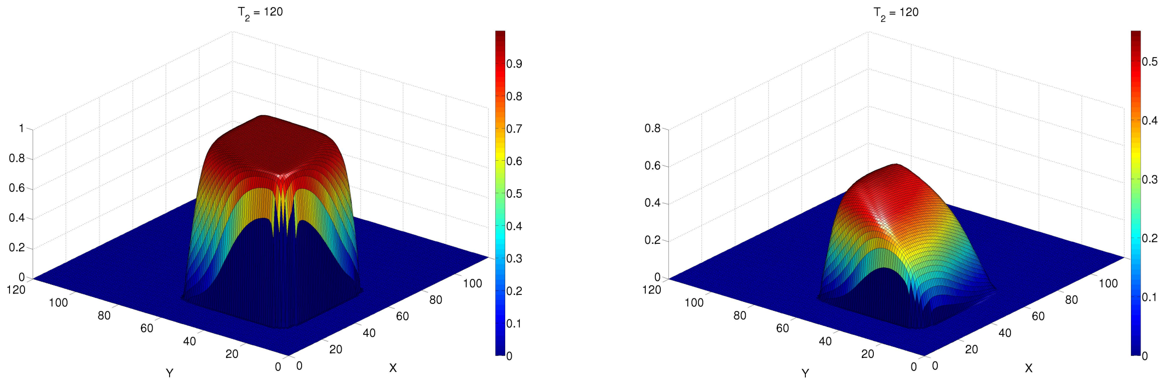

(a) V_2 in case of full controllability over both production rates.

(b) V_2 in case only the production rate of protein x is controllable.

Figure 3.7: The value function V_2 as defined in (3.11b) corresponding to probability of staying in C for 120 time units.

Therefore, we need to solve the PDE of V_2 backward from the time $T_1 + T_2$ to T_1 together with the above boundary condition. Then, the value function V_1 can be computed via solving the same PDE from T_1 to 0 with boundary condition (3.12). The Dynkin operator $\mathcal{L}^u$ reduces to

$$
\begin{aligned}
&\sup_{u\in\mathbb{U}} \mathcal{L}^u \phi(t,x,y)\\
&\quad= \max_{u\in\mathbb{U}} \Big[\partial_t\phi + \partial_x\phi\big(f(y,u_x) - \mu_x x\big) + \partial_y\phi\big(g(x,u_y) - \mu_y y\big)\\
&\qquad\qquad + \frac{1}{2}\partial_x^2\phi\big(f(y,u_x) + \mu_x x\big) + \frac{1}{2}\partial_y^2\phi\big(g(x,u_y) + \mu_y y\big)\Big]\\
&\quad= \partial_t\phi - \big(\partial_x\phi - \frac{1}{2}\partial_x^2\phi\big)\mu_x x - \big(\partial_y\phi - \frac{1}{2}\partial_y^2\phi\big)\mu_y y\\
&\qquad\qquad + \max_{u_x\in[\underline{u}_x,\overline{u}_x]} \big[f(y,u_x)\big(\partial_x\phi + \frac{1}{2}\partial_x^2\phi\big)\big] + \max_{u_y\in[\underline{u}_y,\overline{u}_y]} \big[g(x,u_y)\big(\partial_y\phi + \frac{1}{2}\partial_y^2\phi\big)\big].
\end{aligned}
$$

Thanks to the linearity of the drift term in $\boldsymbol{u}$, an optimal policy can be expressed in terms of derivatives of the value functions V_1 and V_2 as

$$
\boldsymbol{u}_x^*(t,x,y) = \begin{cases} \overline{u}_x(t,x,y) & \text{if } \ \partial_x V_i(t,x,y) + \frac{1}{2}\partial_x^2 V_i(t,x,y) \ge 0,\\ \underline{u}_x(t,x,y) & \text{if } \ \partial_x V_i(t,x,y) + \frac{1}{2}\partial_x^2 V_i(t,x,y) < 0,\end{cases}
$$

$$
\boldsymbol{u}_y^*(t,x,y) = \begin{cases} \overline{u}_y(t,x,y) & \text{if } \ \partial_y V_i(t,x,y) + \frac{1}{2}\partial_y^2 V_i(t,x,y) \ge 0,\\ \underline{u}_y(t,x,y) & \text{if } \ \partial_y V_i(t,x,y) + \frac{1}{2}\partial_y^2 V_i(t,x,y) < 0,\end{cases}
$$

where $i \in \{1,2\}$ corresponds to the phase of the motion.

For this system we investigate two scenarios: one where full control over both production rates is possible and one where only the production rate of protein x can be controlled. Accordingly, in the first scenario we set $\underline{u}_x = \underline{u}_y = 0$ and $\overline{u}_x = \overline{u}_y = 2$ while in the second we set $\underline{u}_x = 0$, $\overline{u_x} = 2$ and $\overline{u}_y = \underline{u}_y = 1$. Figure 3.7 depicts the probability distribution of staying in set C within the time horizon $T_2 = 120$ time units [6] in terms of the initial conditions $(x,y) \in \mathbb{R}^2$.

[6]Notice that the half-life of each protein is assumed to be 17.32 time units

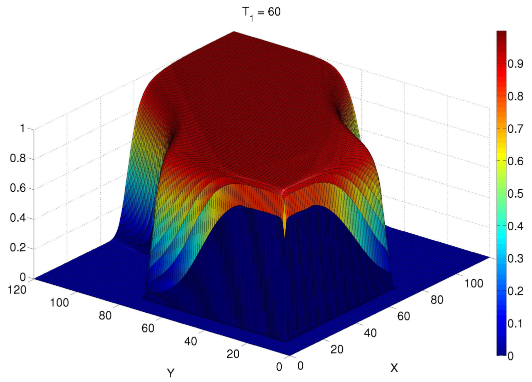

(a) V_1 in case of full controllability over the production rates.

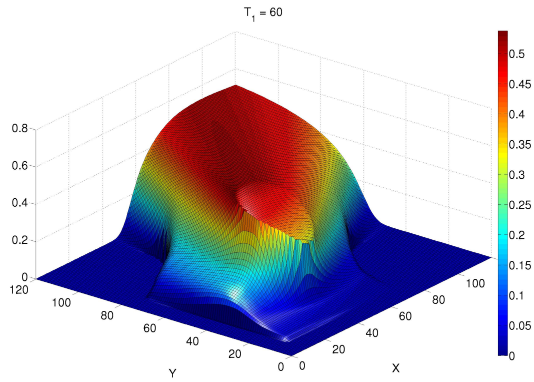

(b) V_1 in case only the production rate of protein x is controllable.

Figure 3.8: The value function V_1 as defined in (3.11a) corresponding to probability of staying in C for 120 time units, once it reaches B while avoiding A within 60 time units.

V_2 is zero outside set C, as the process has obviously left C if it starts outside it. Figures 3.7(a) and 3.7(b) demonstrate the first and second scenarios, respectively. Note that in the second case the probability of success dramatically decreases in comparison to the first. This result indicates the importance of full controllability of the production rates for the achievement of the desired control objective.

Figure 3.8 depicts the probability of successively reaching set B within the time horizon $T_1 = 60$ time units and staying in set C for $T_2 = 120$ time units thereafter. Since the objective is to avoid A, the value function V_1 takes zero value on A. Figures 3.8(a) and 3.8(b) demonstrate the first and second control scenarios, respectively. It is easy to observe the non-smooth behavior of the value function V_1 on the boundary of set B in Figure 3.8(b). This is indeed a consequence of the boundary condition (3.12). All simulations in this subsection were obtained using the Level Set Method Toolbox [Mit05] (version 1.1), with a grid 121×121 in the region of interest.

3.7 Summary and Outlook

We introduced different notions of stochastic motion planning problems. Based on a class of stochastic optimal control problems, we characterized the set of initial conditions from which there exists an admissible policy to execute the desired maneuver with probability no less than some pre-specified value. We then established a weak DPP in terms of auxiliary value functions. Subsequently, we focused on a case of diffusions as the solution of a controlled SDE, and investigated the required conditions to apply the proposed DPP. It turned out that invoking the DPP one can solve a series of PDEs in a recursive fashion to numerically approximate the desired initial set as well as the admissible policy for the motion planning specifications. Finally, the performance of the proposed stochastic motion planning notions was illustrated for a biological switch network.

For future work, as Theorem 3.4.3 holds for the broad class of stochastic processes whose sample paths are right continuous with left limits, one may study the required conditions of the proposed DPP (Assumptions 3.4.1) for a larger class of stochastic processes, e.g., controlled Markov jump-diffusions. Furthermore, motivated by the fact that full state measurements may not be available in practice, an interesting question is to address the motion planning objective with imperfect information, i.e., an admissible control policy would be only allowed to utilize the information of the process $Y_s := h(X_s)$ where $h : \mathbb{R}^d \to \mathbb{R}^{d_y}$ is a given measurable mapping.

3.8 Appendix

This Appendix collects the missing proofs throughout the chapter.

3.8.1 Proofs of Section 3.3

Proof of Lemma 3.3.2. Let τ_A be the first exit-time from the set A_i:

$$\tau_{A_i}(t,x) := \inf\{s \geq 0 \ : \ X^{t,x;\boldsymbol{u}}_{t+s} \notin A_i\}. \tag{3.13}$$

We know that τ_A is an $\mathbb{F}_t$-stopping time [EK86, Thm. 1.6, Chapter 2]. Let $\omega(\,\cdot\,) \mapsto \vartheta_s\big(\omega(\,\cdot\,)\big) := \omega(s + \,\cdot\,)$ be the time-shift operator. From the definition it follows that for all $i \geq 0$

$$\Theta^{A_{1:n}}_{i+1} = \Theta^{A_{1:n}}_{i} + \tau_{A_i} \circ \vartheta_{\Theta^{A_{1:n}}_i}.$$

Now the assertion follows directly in light of the measurability of the mapping ϑ and right continuity of the filtration $\mathbb{F}_t$; see [EK86, Prop. 1.4, Chapter 2] for more details in this regard. □

Before proceeding with the proof of Theorem 3.3.3, we start with a fact which is an immediate consequence of right continuity of the process $X^{t,x;\boldsymbol{u}}_{\cdot}$:

Fact 3.8.1. *Fix a control policy $\boldsymbol{u} \in \mathcal{U}_t$ and an initial condition $(t,x) \in \mathbb{S}$. Let $(A_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d)$ be a sequence of open sets. Then, for all $i \in \{1,\cdots,n\}$*

$$X^{t,x;\boldsymbol{u}}_{\Theta^{A_{1:n}}_i} \notin A_i, \qquad \text{on} \quad \big\{\Theta^{A_{1:n}}_i < \infty\big\},$$

where $\big(\Theta^{A_{1:n}}_i\big)_{i=1}^n$ are the sequential exit-times in the sense of Definition 3.3.1.

Proof of Theorem 3.3.3. We first show (3.4). Observe that it suffices to prove that

$$\Big\{X^{t,x;\boldsymbol{u}}_{\cdot} \models \big[(W_1 \rightsquigarrow G_1) \circ \cdots \circ (W_n \rightsquigarrow G_n)\big]_{\leq T}\Big\} = \bigcap_{i=1}^{n} \Big\{X^{t,x;\boldsymbol{u}}_{\eta_i} \in G_i\Big\} \tag{3.14}$$

for all initial conditions (t,x) and policies $\boldsymbol{u}$, where the stopping time η_i is as defined in (3.3a). Let ω belong to the left-hand side of (3.14). In view of the definition (3.1a), there exists a set of instants $(s_i)_{i=1}^n \subset [t,T]$ such that for all i, $X^{t,x;\boldsymbol{u}}_{s_i}(\omega) \in G_i$ while $X^{t,x;\boldsymbol{u}}_{r}(\omega) \in W_i \setminus G_i =: B_i$ for all $r \in [s_{i-1}, s_i[$, where we set $s_0 = t$. It then follows by an induction argument that

$\eta_i(\omega) = \Theta_i^{B_{1:n}} = s_i$, which immediately leads to $X^{t,x;\boldsymbol{u}}_{\eta_i(\omega)}(\omega) \in G_i$ for all $i \le n$. This proves the relation " $\subset$ " between the left- and right-hand sides of (3.14). Now suppose that ω belongs to the right-hand side of (3.14). Then, we have $X^{t,x;\boldsymbol{u}}_{\eta_i(\omega)}(\omega) \in G_i$ for all $i \le n$. In view of the definition of stopping times η_i in (3.3a), it follows that $X^{t,x;\boldsymbol{u}}_r(\omega) \in B_i := W_i \setminus G_i$ for all $r \in [\eta_{i-1}(\omega), \eta_i(\omega)[$. Introducing the time sequence $s_i := \eta_i(\omega)$ implies the relation " $\supset$ " between the left- and right-hand sides of (3.14). Together with preceding argument, this implies (3.14).

To prove (3.5) we only need to show that

$$\Big\{X^{t,x;\boldsymbol{u}}_{\cdot} \models (W_1 \xrightarrow{T_1} G_1) \circ \cdots \circ (W_n \xrightarrow{T_n} G_n)\Big\} = \bigcap_{i=1}^{n} \Big\{X^{t,x;\boldsymbol{u}}_{\widetilde{\eta}_i} \in G_i \cap W_i\Big\} \tag{3.15}$$

for all initial conditions (t,x) and policies $\boldsymbol{u}$, where the stopping time $\widetilde{\eta}_i$ is introduced in (3.3b). To this end, let us fix $(t,x) \in \mathbb{S}$ and $\boldsymbol{u} \in \mathcal{U}_t$, and assume that ω belongs to the left-hand side of (3.15). By definition (3.1b), for all $i \le n$ we have $X^{t,x;\boldsymbol{u}}_{T_i}(\omega) \in G_i$ and $X^{t,x;\boldsymbol{u}}_r(\omega) \in W_i$ for all $r \in [T_{i-1}, T_i]$. By a straightforward induction, we see that $\widetilde{\eta}_i(\omega) = T_i$, and consequently $X^{t,x;\boldsymbol{u}}_{\widetilde{\eta}_i(\omega)}(\omega) \in G_i \cap W_i$ for all $i \le n$. This establishes the relation " $\subset$ " between the left- and right-hand sides of (3.15). Now suppose ω belongs to the right-hand side of (3.15). Then, for all $i \le n$ we have $X^{t,x;\boldsymbol{u}}_{\widetilde{\eta}_i(\omega)}(\omega) \in G_i \cap W_i$. By virtue of Fact 3.8.1 and an induction argument once again, it is guaranteed that $\widetilde{\eta}_i(\omega) = T_i$, and consequently it follows that $X^{t,x;\boldsymbol{u}}_{T_i}(\omega) \in G_i$ and $X^{t,x;\boldsymbol{u}}_r(\omega) \in W_i$ for all $r \in [T_{i-1}, T_i]$. This establishes the relation " $\supset$ " in (3.15), and the assertion follows. □

3.8.2 Proofs of Section 3.4

Before proceeding with the proof of Theorem 3.4.3, we need a preparatory lemma.

Lemma 3.8.2. *Under Assumptions 3.4.1.b.iii. and 3.4.1.c., the function* $\mathbb{S} \ni (t,x) \mapsto J_k(t,x;\boldsymbol{u}) \in \mathbb{R}$ *is lower semicontinuous for all* $k \in \{1, \cdots, n\}$ *and control policy* $\boldsymbol{u} \in \mathcal{U}_0$.

Proof. Fix $k \in \{1,\ldots,n\}$. It is obvious that the function J_k is uniformly bounded since ℓ_k are. Therefore,

$$\begin{aligned}\liminf_{(s,y)\to(t,x)} J_k\big(s,y;\boldsymbol{u}\big) &= \liminf_{(s,y)\to(t,x)} \mathbb{E}\Big[\prod_{i=k}^{n} \ell_i\big(X^{s,y;\boldsymbol{u}}_{\tau_i^k(s,y)}\big)\Big] \ge \mathbb{E}\Big[\liminf_{(s,y)\to(t,x)} \prod_{i=k}^{n} \ell_i\big(X^{s,y;\boldsymbol{u}}_{\tau_i^k(s,y)}\big)\Big] \\ &\ge \mathbb{E}\Big[\prod_{i=k}^{n} \liminf_{(s,y)\to(t,x)} \ell_i\big(X^{s,y;\boldsymbol{u}}_{\tau_i^k(s,y)}\big)\Big] \ge \mathbb{E}\Big[\prod_{i=k}^{n} \ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k(s,y)}\big)\Big] = J_k(t,x;\boldsymbol{u}),\end{aligned}$$

where the inequality in the first line follows from the Fatou's lemma, and the second inequality in the second line is a direct consequence of Assumptions 3.4.1.b.iii. and 3.4.1.c. □

Proof of Theorem 3.4.3. The proof extends the main result of Theorem 2.4.7 on the so-called reach-avoid maneuver. Let $u \in \mathcal{U}_t$, $\theta := \theta^{\boldsymbol{u}} \in \mathcal{T}_{[t,T]}$, and $\boldsymbol{u}_\theta$ be the random policy as introduced in Assumption 3.4.1.b.ii. Then we have

$$
\begin{aligned}
\mathbb{E}\Big[\prod_{i=k}^{n}\ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big)\;\Big|\;\mathcal{F}_\theta\Big] &= \sum_{j=k}^{n+1}\mathbb{1}_{\{\tau_{j-1}^k\le\theta<\tau_j^k\}}\mathbb{E}\Big[\prod_{i=j}^{n}\ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big)\;\Big|\;\mathcal{F}_\theta\Big]\prod_{i=k}^{j-1}\ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big)\\
&= \sum_{j=k}^{n+1}\mathbb{1}_{\{\tau_{j-1}^k\le\theta<\tau_j^k\}}J_j\big(\theta, X^{t,x;\boldsymbol{u}}_\theta;\boldsymbol{u}_\theta\big)\prod_{i=k}^{j-1}\ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big) \qquad (3.16a)\\
&\le \sum_{j=k}^{n+1}\mathbb{1}_{\{\tau_{j-1}^k\le\theta<\tau_j^k\}}V_j\big(\theta, X^{t,x;\boldsymbol{u}}_\theta\big)\prod_{i=k}^{j-1}\ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big) \qquad (3.16b)
\end{aligned}
$$

where (3.16a) follows from Assumption 3.4.1.b.ii. and right continuity of the process, and (3.16b) is due to the fact that $\boldsymbol{u}_\theta\in\mathcal{U}_{\theta(\omega)}$ for each realization $\omega\in\Omega$. In light of the tower property of conditional expectation [Kal97, Thm. 5.1], arbitrariness of $\boldsymbol{u}\in\mathcal{U}_t$, and obvious inequality $V_j\le V_j^*$, we arrive at (3.7a).

To prove (3.7b), consider uniformly bounded upper semicontinuous functions $(\phi_j)_{j=k}^n\subset \mathrm{USC}(\mathbb{S})$ such that $\phi_j\le V_{j*}$ on $\mathbb{S}$. Mimicking the ideas in the proof of Theorem 2.4.7 and due to Lemma 3.8.2, one can construct an admissible control policy $\boldsymbol{u}_j^\varepsilon$ for any $\varepsilon>0$ and $j\in\{k,\cdots,n\}$ such that

$$
\phi_j(t,x)-3\varepsilon\le J_j(t,x;\boldsymbol{u}_j^\varepsilon)\qquad\forall (t,x)\in\mathbb{S}. \tag{3.17}
$$

Let us fix $\boldsymbol{u}\in\mathcal{U}_t$ and $\varepsilon>0$, and define

$$
\boldsymbol{v}^\varepsilon := \mathbb{1}_{[t,\theta]}\boldsymbol{u} + \mathbb{1}_{]\theta,T]}\sum_{j=k}^{n}\mathbb{1}_{\{\tau_{j-1}^k\le\theta<\tau_j^k\}}\boldsymbol{u}_j^\varepsilon, \tag{3.18}
$$

where $\boldsymbol{u}_j^\varepsilon$ satisfies (3.17). Notice that Assumption 3.4.1.a. ensures $\boldsymbol{v}^\varepsilon\in\mathcal{U}_t$. By virtue of the tower property, Assumptions 3.4.1.b.i. and 3.4.1.b.ii., and the assertions in (3.17) and (3.18), it follows

$$
\begin{aligned}
V_k(t,x)\ge J_k(t,x;\boldsymbol{v}^\varepsilon) &= \mathbb{E}\Big[\mathbb{E}\Big[\prod_{i=k}^{n}\ell_i\big(X^{t,x;\boldsymbol{v}^\varepsilon}_{\tau_i^k}\big)\;\Big|\;\mathcal{F}_\theta\Big]\Big]\\
&= \mathbb{E}\Big[\sum_{j=k}^{n+1}\mathbb{1}_{\{\tau_{j-1}^k\le\theta<\tau_j^k\}}J_j\big(\theta, X^{t,x;\boldsymbol{u}}_\theta;\boldsymbol{u}_j^\varepsilon\big)\prod_{i=k}^{j-1}\ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big)\Big]\\
&= \mathbb{E}\Big[\sum_{j=k}^{n+1}\mathbb{1}_{\{\tau_{j-1}^k\le\theta<\tau_j^k\}}\Big(\phi_j\big(\theta, X^{t,x;\boldsymbol{u}}_\theta\big)-3\varepsilon\Big)\prod_{i=k}^{j-1}\ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big)\Big].
\end{aligned}
$$

Now, consider a sequence of increasing continuous functions $(\phi_j^m)_{m\in\mathbb{N}}$ that converges point-wise to V_{j*}. The existence of such sequence is ensured by Lemma 3.8.2, see [Ren99, Lemma 3.5]. By boundedness of $(\ell_j)_{i=1}^n$ and the dominated convergence Theorem, we get

$$
V_k(t,x)\ge\mathbb{E}\Big[\sum_{j=k}^{n+1}\mathbb{1}_{\{\tau_{j-1}^k\le\theta<\tau_j^k\}}\Big(V_{j*}\big(\theta, X^{t,x;\boldsymbol{u}}_\theta\big)-3\varepsilon\Big)\prod_{i=k}^{j-1}\ell_i\big(X^{t,x;\boldsymbol{u}}_{\tau_i^k}\big)\Big]
$$

Since $\boldsymbol{u}\in\mathcal{U}_t$ and $\varepsilon>0$ are arbitrary, this leads to (3.7b). $\square$

3.8.3 Proofs of Section 3.5

Proof of Proposition 3.5.4. The key step in the proof relies on the two Assumptions 3.5.1.c. and 3.5.2. There is a classical result on non-degenerate diffusion processes indicating that if the process starts from the tip of a cone, then it enters the cone with probability one [RB98, Corollary 3.2, p. 65]. This hints at the possibility that the aforementioned Assumptions together with almost sure continuity of the strong solution of the SDE (3.8) result in the continuity of sequential exit-times $\Theta_i^{A_{1:n}}$ and consequently τ_i. In the following we shall formally work around this idea.

Let us assume that $t_m \le t$ for notational simplicity, but one can effectively follow similar arguments for $t_m > t$. By the definition of the SDE (3.8),

$$X_r^{t_m,x_m;\boldsymbol{u}} = X_t^{t_m,x_m;\boldsymbol{u}} + \int_t^r f\big(X_s^{t_m,x_m;\boldsymbol{u}}, u_s\big)\mathrm{d}s + \int_t^r \sigma\big(X_s^{t_m,x_m;\boldsymbol{u}}, u_s\big)\mathrm{d}W_s \qquad \mathbb{P}\text{-a.s.}$$

By virtue of [Kry09, Thm. 2.5.9, p. 83], for all $q \ge 1$ we have

$$\begin{aligned}\mathbb{E}\Big[\sup_{r\in[t,T]} \big\|X_r^{t,x;\boldsymbol{u}} - X_r^{t_m,x_m;\boldsymbol{u}}\big\|^{2q}\Big] &\le C_1(q,T,K)\mathbb{E}\Big[\big\|x - X_t^{t_m,x_m;\boldsymbol{u}}\big\|^{2q}\Big] \\ &\le 2^{2q-1}C_1(q,T,K)\mathbb{E}\Big[\|x-x_m\|^{2q} + \big\|x_m - X_t^{t_m,x_m;\boldsymbol{u}}\big\|^{2q}\Big],\end{aligned}$$

whence, in light of [Kry09, Corollary 2.5.12, p. 86], we get

$$\mathbb{E}\Big[\sup_{r\in[t,T]} \big\|X_r^{t,x;\boldsymbol{u}} - X_r^{t_n,x_n;\boldsymbol{u}}\big\|^{2q}\Big] \le C_2(q,T,K,\|x\|)\big(\|x-x_n\|^{2q} + |t-t_n|^q\big). \tag{3.19}$$

In the above inequalities, K is the Lipschitz constant of f and σ mentioned in Assumption 3.5.1.b.; C_1 and C_2 are constant depending on the indicated parameters. Hence, in view of Kolmogorov's continuity criterion [Pro05, Corollary 1 Chap. IV, p. 220], one may consider a version of the stochastic process $X_\cdot^{t,x;\boldsymbol{u}}$ which is continuous in (t,x) in the topology of uniform convergence on compacts. This leads to the fact that $\mathbb{P}$-a.s, for any $\varepsilon > 0$, for all sufficiently large m,

$$X_r^{t_m,x_m;\boldsymbol{u}} \in \mathbb{B}_\varepsilon\big(X_r^{t_0,x_0;\boldsymbol{u}}\big), \qquad \forall r \in [t_m, T], \tag{3.20}$$

where $\mathbb{B}_\varepsilon(y)$ denotes the ball centered at y and radius ε. For simplicity, let us define the shorthand $\tau_i^m := \tau_i(t_m, x_m)$.[7] By the definition of τ_i and Definition 3.3.1, since the set A_i is open, we conclude that

$$\exists \varepsilon > 0, \qquad \bigcup_{s\in[\tau_{i-1}^0,\tau_i^0[} \mathbb{B}_\varepsilon(X_s^{t_0,x_0;\boldsymbol{u}}) \cap A_i^c = \emptyset \qquad \mathbb{P}\text{-a.s.} \tag{3.21}$$

By definition $\tau_0^0 := \tau_0(t_0, x_0) = t_0$. As an induction hypothesis, let us assume τ_{i-1}^0 is $\mathbb{P}$-a.s. continuous, and we proceed with the induction step. One can deduce that (3.21) together with (3.20) implies that $\mathbb{P}$-a.s. for all sufficiently large m,

$$X_r^{t_m,x_m;\boldsymbol{u}} \in A_i, \qquad \forall r \in [t_m, \tau_i^0[.$$

[7] This notation is only employed in this proof.

In conjunction with $\mathbb{P}$-a.s. continuity of sample paths, this immediately leads to

$$\liminf_{m\to\infty} \tau_i^m := \liminf_{m\to\infty} \tau_i(t_m, x_m) \geq \tau_i(t_0, x_0) \qquad \mathbb{P}\text{-a.s.} \tag{3.22}$$

On the other hand, as mentioned earlier, the Assumptions 3.5.1.c. and 3.5.2 imply that the set of sample paths that hit the boundary of A_i and do not enter the set is negligible [RB98, Corollary 3.2, p. 65]. Hence

$$\forall \delta > 0, \quad \exists s \in \big[\Theta_i^{A_{1:n}}(t_0, x_0), \Theta_i^{A_{1:n}}(t_0, x_0) + \delta\big[, \quad X_s^{t_0,x_0;\boldsymbol{u}} \in A_i \quad \mathbb{P}\text{-a.s.}$$

Hence, in light of (3.20), $\mathbb{P}$-a.s. there exists $\varepsilon > 0$, possibly depending on δ, such that for all sufficiently large m we have

$$X_s^{t_m,x_m;\boldsymbol{u}} \in \mathbb{B}_\varepsilon(X_s^{t_0,x_0;\boldsymbol{u}}) \subset A_i^c$$

Recalling the induction hypothesis, we note that in accordance with the definition of sequential stopping times $\Theta_i^{A_{1:n}}$, one can infer that $\Theta_i^{A_{1:n}}(t_m, x_m) \leq s < \Theta_i^{A_{1:n}}(t_0, x_0) + \delta$. From arbitrariness of δ and the definition of τ_i, this leads to

$$\limsup_{m\to\infty} \tau_i(t_m, x_m) := \limsup_{m\to\infty} \Big(\Theta_i^{A_{1:n}}(t_m, x_m) \wedge T_i\Big) \leq \tau_i(t_0, x_0) \qquad \mathbb{P}\text{-a.s.},$$

where in conjunction with (3.22), $\mathbb{P}$-a.s. continuity of the map $(t, x) \mapsto \tau_i(t, x)$ at (t_0, x_0) for any $i \in \{1, \cdots, n\}$ follows. The assertion follows by induction.

The continuity of the mapping $(t, x) \mapsto X_{\tau_i(t,x)}^{t,x;\boldsymbol{u}}$ follows immediately from the almost sure continuity of the stopping time $\tau_i(t, x)$ in conjunction with the almost sure continuity of the version of the stochastic process $X_\cdot^{t,x;\boldsymbol{u}}$ in (t, x); for the latter let us note again that Kolmogorov's continuity criterion guarantees the existence of such a version in light of (3.19). $\square$

Proof of Theorem 3.5.6. Here we briefly sketch the proof in words, and refer the reader to the proof of Theorem 2.4.10 for detailed arguments concerning the same technology to prove the assertion of the Theorem. Note that any $\mathbb{F}_t$-progressively measurable policy $\boldsymbol{u} \in \mathcal{U}_t$ satisfies Assumptions 3.4.1.a.. It is a classical result [Øks03, Chap. 7] that the strong solution $X_\cdot^{t,x;\boldsymbol{u}}$ satisfies Assumptions 3.4.1.b.i. and 3.4.1.b.ii. Furthermore, Proposition 3.5.4 together with almost sure path-continuity of the strong solution guarantees Assumption 3.4.1.b.iii. Hence, having met all the required assumptions of Theorem 3.4.3, one can employ the DPP (3.7). Namely, to establish the assertion concerning the *supersolution*, for the sake of contradiction, one can assume that there exists $(t_0, x_0) \in [0, T_k[\times A_k$, and a smooth function ϕ dominated by the value function V_{k*} where $(V_{k*} - \phi)(t_0, x_0) = 0$, such that for some $\delta > 0$, $-\sup_{u\in\mathbb{U}} \mathcal{L}^u\phi(t_0, x_0) < -2\delta$. Since ϕ is smooth, the map $(t, x) \mapsto \mathcal{L}^u\phi(t, x)$ is continuous. Therefore, there exist $u \in \mathbb{U}$ and $r > 0$ such that $\mathbb{B}_r(t_0, x_0) \subset [0, T_k[\times A_k$ and $-\mathcal{L}^u\phi(t, x) < -\delta$ for all (t, x) in $\mathbb{B}_r(t_0, x_0)$. Let us define the stopping time $\theta(t, x)$ as the first exit time of trajectory $X_\cdot^{t,x;u}$ from the ball $\mathbb{B}_r(t_0, x_0)$. Note that by continuity of solutions to (3.8), $t < \theta(t, x)$ $\mathbb{P}$- a.s. for all $(t, x) \in \mathbb{B}_r(t_0, x_0)$. Therefore, selecting $r > 0$ sufficiently small so that $\theta < \tau_k$, and applying Itô's formula, we see that for all $(t, x) \in \mathbb{B}_r(t_0, x_0)$, $\phi(t, x) < \mathbb{E}\big[\phi\big(\theta(t, x), X_{\theta(t,x)}^{t,x;u}\big)\big]$. Now it suffices to take a sequence $(t_m, x_m, V_k(t_m, x_m))_{m\in\mathbb{N}}$ converging to $(t_0, x_0, V_{k*}(t_0, x_0))$. For sufficiently large m we have $V(t_m, x_m) < \mathbb{E}\big[V_{k*}\big(\theta(t_m, x_m), X_{\theta(t_m,x_m)}^{t_m,x_m;u}\big)\big]$ which, in view of the fact that $\theta(t_m, x_m) < \tau_k \wedge T_k$, contradicts the DPP in (3.7a). The *subsolution* property is proved effectively in a similar fashion. $\square$

To provide boundary conditions, in particular in the viscosity sense (3.9b), we need some preliminaries as follows:

Fact 3.8.3. *Consider a control policy $\boldsymbol{u} \in \mathcal{U}_t$ and initial condition $(t, x) \in \mathbb{S}$. Given a sequence of $(A_i)_{i=k}^n \subset \mathfrak{B}(\mathbb{R}^d)$ and stopping time $\theta \in \mathcal{T}_{[t,T]}$, for all $k \in \{1, \cdots, n\}$ and $j \geq i \geq k$ we have*

$$\Theta_j^{A_{k:n}}(t, x) = \Theta_j^{A_{i:n}}\big(\theta, X_\theta^{t,x;\boldsymbol{u}}\big) \qquad on \quad \big\{\Theta_{i-1}^{A_{k:n}}(t, x) \leq \theta < \Theta_i^{A_{k:n}}(t, x)\big\}$$

Lemma 3.8.4. *Suppose that the conditions of Proposition 3.5.4 hold. Given a sequence of control policies $(\boldsymbol{u_m})_{m\in\mathbb{N}} \subset \mathcal{U}$ and initial conditions $(t_m, x_m) \to (t, x)$, we have*

$$\lim_{m\to\infty} \Big\| X_{\tau_i(t,x)}^{t,x;\boldsymbol{u}_m} - X_{\tau_i(t_m,x_m)}^{t_m,x_m;\boldsymbol{u}_m} \Big\| = 0 \quad \mathbb{P}\text{-}a.s., \qquad \tau_i(t, x) \coloneqq \Theta_i^{A_{1:n}}(t, x) \wedge T_i.$$

Note that Lemma 3.8.4 is indeed a stronger statement than Proposition 3.5.4 as the desired continuity is required uniformly with respective to the control policy. Let us highlight that the stopping times $\tau_i(t, x)$ and $\tau_i(t_m, x_m)$ are both effected by control policies $\boldsymbol{u}_m$. But nonetheless, the mapping $(t, x) \mapsto X_{\tau_i}^{t,x;\boldsymbol{u_m}}$ is almost surely continuous irrespective of the policies $(\boldsymbol{u}_m)_{m\in\mathbb{N}}$. For the proof we refer to an identical technique used in the proof of Lemma 2.4.11 in Section 2.8.2.

Proof of Proposition 3.5.7. The boundary condition in (3.9a) is an immediate consequence of the definition of the sequential exit-times introduced in Definition 3.3.1. Namely, for any initial state $x \in A_k^c$ we have $\Theta_k^{A_{k:n}}(t, x) = t$, and in light of Fact 3.8.3 for all $i \in \{k, \cdots, n\}$

$$\Theta_i^{A_{k:n}}(t, x) = \Theta_i^{A_{k+1:n}}(t, x), \qquad \forall (t, x) \ \in [0, T_k] \times A_k^c \bigcup \{T_k\} \times \mathbb{R}^d.$$

Since $\tau_k^k = t$ for the above initial conditions, then $X_{\tau_k^k}^{t,x;\boldsymbol{u}} = x$ which yields to (3.9a).

For the boundary conditions (3.9b), we show the first assertion; the second follows similarly. Let $\big(t_m, x_m\big) \to \big(t, x\big)$ where $t_m < T_k$ and $x_m \in A_k$. Invoking the DPP in Theorem 3.4.3 and introducing $\theta \coloneqq \tau_{k+1}^k$ in (3.7a), we reach

$$V_k(t_m, x_m) \leq \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[V_{k+1}^*\big(\tau_k^k, X_{\tau_k^k}^{t_m,x_m;\boldsymbol{u}}\big)\ell_k\big(X_{\tau_k^k}^{t_m,x_m;\boldsymbol{u}}\big)\Big].$$

Note that one can replace a sequence of policies in the above inequalities to attain the supremum running over all policies. This sequence, of course, depends on the initial condition (t_m, x_m). Hence, let us denote it via two indices $(\boldsymbol{u}_{m,j})_{j\in\mathbb{N}}$. One can deduce that there exists a subsequence of $(\boldsymbol{u}_{m_j})_{j\in\mathbb{N}}$ such that

$$\begin{aligned}
\lim_{m\to\infty} V_k(t_m, x_m) &\leq \lim_{m\to\infty} \lim_{j\to\infty} \mathbb{E}\Big[V_{k+1}^*\big(\tau_k^k, X_{\tau_k^k}^{t_m,x_m;\boldsymbol{u}_{m,j}}\big)\ell_k\big(X_{\tau_k^k}^{t_m,x_m;\boldsymbol{u}_{m,j}}\big)\Big] \\
&\leq \lim_{j\to\infty} \mathbb{E}\Big[V_{k+1}^*\big(\tau_k^k, X_{\tau_k^k}^{t_j,x_j;\boldsymbol{u}_{m_j}}\big)\ell_k\big(X_{\tau_k^k}^{t_j,x_j;\boldsymbol{u}_{m_j}}\big)\Big] \\
&\leq \mathbb{E}\Big[\lim_{j\to\infty} V_{k+1}^*\big(\tau_k^k, X_{\tau_k^k}^{t_j,x_j;\boldsymbol{u}_{m_j}}\big)\ell_k^*\big(X_{\tau_k^k}^{t_j,x_j;\boldsymbol{u}_{m_j}}\big)\Big] \qquad (3.23) \\
&= V_{k+1}^*(t, x)\ell_k^*(x) \qquad (3.24)
\end{aligned}$$

where (3.23) and (3.24) follow, respectively, from Fatou's lemma and the uniform continuity assertion in Lemma 3.8.4. Let us recall that by Lemma 3.8.4 we know $\tau_k^k(t_j, x_j) \to \tau_k^k(t, x) = t$ as $j \to \infty$ uniformly with respect to the policies $(\boldsymbol{u}_{m_j})_{j\in\mathbb{N}}$. Similar analysis would follow for the second part of (3.9b) by using the other side of DPP in (3.7b). $\square$

Part II

Fault Detection for Large Scale Nonlinear Systems

CHAPTER 4

A Tractable Approach with Probabilistic Performance Index

The second part of this thesis presents a novel perspective along with a scalable methodology to design a fault detection and isolation (FDI) filter for high dimensional nonlinear systems. Previous approaches on FDI problems are either confined to linear systems or they are only applicable to low dimensional dynamics with specific structures. In contrast, shifting attention from the system dynamics to the disturbance inputs, we propose a relaxed design perspective to train a linear residual generator given some statistical information about the disturbance patterns. That is, we propose an optimization-based approach to robustify the filter with respect to finitely many signatures of the nonlinearity. We then invoke existing results in randomized optimization to provide theoretical guarantees for the performance of the proposed filer. Finally, motivated by a cyber-physical attack emanating from the vulnerabilities introduced by the interaction between IT infrastructure and power system, we deploy the developed theoretical results to detect such an intrusion before the functionality of the power system is disrupted.

4.1 Introduction

The task of FDI in control systems involves generating a diagnostic signal sensitive to the occurrence of specific faults. This task is typically accomplished by designing a filter with all available information as inputs (e.g., control signals and given measurements) and a scalar output that implements a non-zero mapping from the fault to the diagnostic signal, which is known as the residual, while decoupling unknown disturbances. The concept of residual plays a central role for the FDI problem which has been extensively studied in the last two decades.

In the context of linear systems, Beard and Jones [Bea71, Jon73] pioneered an observer-based approach whose intrinsic limitation was later improved by Massoumnia et al. [MVW89]. Following the same principles but from a game theoretic perspective, Speyer and coauthors thoroughly investigated the approach in the presence of noisy measurements [CS98b]. Nyberg and Frisk extended the class of systems to linear differential-algebraic equation (DAE) apparently subsuming all the previous linear classes [NF06]. This extension greatly enhanced the applicability of FDI methods since the DAE models appear in a wide range of applications, including electrical systems, robotic manipulators, and mechanical systems.

For nonlinear systems, a natural approach is to linearize the model at an operating point, treat the nonlinear higher order terms as disturbances, and decouple their contributions from the residual by employing robust techniques [SF91, HP96]. This strategy only works well if either the system remains close to the chosen operating point, or the exact decoupling is possible. The former approach is often limited, since in the presence of unknown inputs the system may have a wide dynamic operating range, which in case linearization leads to a large mismatch between linear model and nonlinear behavior. The latter approach was explored in detail by De Persis and Isidori, who in [PI01] proposed a differential geometric approach to extend the unobservibility subspaces of [Mas86, Section IV], and by Chen and Patton, who in [CP82, Section 9.2] dealt with a particular class of bilinear systems. These methods are, however, practically limited by the need to verify the required conditions on the system dynamics and transfer them into a standard form, which essentially involve solving partial differential equations, restricting the application of the method to relatively low dimensional systems.

Motivated by this shortcoming, in this chapter we develop a novel approach to FDI which strikes a balance between analytical and computational tractability, and is applicable to high dimensional nonlinear dynamics. For this purpose, we propose a design perspective that basically shifts the emphasis from the system dynamics to the family of disturbances that the system may encounter. We assume that some statistical information of the disturbance patterns is available. Following [NF06] we restrict the FDI filters to a class of linear operators that fully decouple the contribution of the linear part of the dynamics. Thanks to the linearity of the resulting filter, we then trace the contribution of the nonlinear term to the residual, and propose an optimization-based methodology to robustify the filter to the nonlinearity signatures of the dynamics by exploiting the statistical properties on the disturbance signals. The optimization formulation is effectively convex and hence tractable for high dimensional dynamics. Some preliminary results in this direction were reported in [MEVAL12], while an application of our approach in the presence of measurement noise was successfully tested for wind turbines in [SMEKL13].

The performance of the proposed methodology is illustrated in an application to an emerging problem in cyber security of power networks. In modern power systems, the cyber-physical interaction of IT infrastructure (SCADA systems) with physical power systems renders the system vulnerable not only to operational errors but also to malicious external intrusions. As an example of this type of cyber-physical interaction we consider here the Automatic Generation Control (AGC) system, which is one of the few control loops in power networks that are closed over the SCADA system without human operator intervention. In earlier work [MEVM+10, MEVM+11] we have shown that, having gained access to the AGC signal, an attacker can provoke frequency deviations and power oscillations by applying sophisticated attack signals. The resulting disruption can be serious enough to trigger generator out-of-step protection relays, leading to load shedding and generator tripping. Our earlier work, however, also indicated that an early detection of the intrusion may allow one to disconnect the AGC and limit the damage by relying solely on the so-called primary frequency controllers. In this chapter we show how to mitigate this cyber-physical security concerns by using the proposed FDI scheme to develop a protection layer which quickly detects the abnormal signals generated by the attacker. This approach to enhancing the cyber-security of power transmission systems led to an EU patent sponsored by ETH Zurich [MEVAL].

The chapter is organized as follows. In Section 4.2 a formal description of the FDI problem as well as the outline of the proposed methodology is presented. A general class of nonlinear models is described in Section 4.3. Then, reviewing residual generation for the linear models, we develop an optimization-based framework for nonlinear systems in Section 4.4. Theoretical guarantees are also provided in the context of randomized algorithms. We apply the developed methodology to the AGC case study in Section 4.5, and finally conclude with some remarks and directions for future work in Section 4.6. For better readability, the technical proofs of Sections 4.4.2 and 4.4.3 are moved to the appendices.

Notation

Here is a partial notation list which will be also explained in more details later in the chapter:

- The symbols $\mathbb{N}$ and $\mathbb{R}_+$ denote the set of natural and nonnegative real numbers, respectively.
- Let $A \in \mathbb{R}^{n\times m}$ be an $n \times m$ matrix with real values. Then, the transpose of the matrix A is denoted by $A^\intercal \in \mathbb{R}^{m\times n}$.
- Let $v := [v_1, \cdots, v_n]^\intercal$ be a vector in $\mathbb{R}^n$. Then $\|v\|_2 := \sqrt{\sum_{i=1}^n v_i^2}$ is the Euclidean vector norm. The infinity norm of v is also denoted by $\|v\|_\infty := \max_{i\le n} |v_i|$.
- Given $A \in \mathbb{R}^{n\times m}$, $\|A\|_2 := \overline{\sigma}(A)$, where $\overline{\sigma}$ is the maximum singular value of the matrix.
- $\mathcal{W}^n$ denotes the set of piece-wise continuous (p.w.c.) functions taking values in $\mathbb{R}^n$.
- $\mathcal{W}^n_T$ is the restriction of $\mathcal{W}^n$ into the time interval $[0, T]$, which is endowed with the $\mathcal{L}_2$-inner product, i.e., $\langle e_1, e_2\rangle := \int_0^T e_1^\intercal(t)e_2(t)\mathrm{d}t$ with associated $\mathcal{L}_2$-norm $\|e\|_{\mathcal{L}_2} := \sqrt{\langle e, e\rangle}$.
- The linear operator $p : \mathcal{W}^n \to \mathcal{W}^n$ is the distributional derivative operator. In particular, if $e : \mathbb{R}_+ \to \mathbb{R}^n$ is a smooth mapping then $p[e(t)] := \frac{\mathrm{d}}{\mathrm{d}t}e(t)$.
- Let G be a linear matrix transfer function. Then $\|G\|_{\mathcal{H}_\infty} := \sup_{\omega\in\mathbb{R}} \overline{\sigma}\big(G(j\omega)\big)$, where $\overline{\sigma}$ is the maximum singular value of the matrix $G(j\omega)$.
- Given a probability space $(\Omega, \mathbb{F}, \mathbb{P})$, we denote by $\mathbb{P}^n$ the n-fold product probability measure on $(\Omega^n, \mathbb{F}^n)$.

4.2 Problem Statement and Outline of the Proposed Approach

In this section, we provide the formal description of the FDI problem as well as our new design perspective. We will also outline our methodology to tackle the proposed perspective.

4.2.1 Formal Description

The objective of the FDI design is to use all information to generate a diagnostic signal to alert the operators to the occurrence of a specific fault. Consider a general dynamical system as in Figure 4.1 with its inputs categorized into (i) unknown inputs d, (ii) fault signal f, and (iii) known inputs u. The unknown input d represents unknown disturbances that the dynamical system encounters during normal operation. The known input u contains all known signals injected to the system which together with the measurements y are available for FDI tasks. Finally, the input f is a fault (or an intrusion) which cannot be directly measured and represents the signal to be detected.

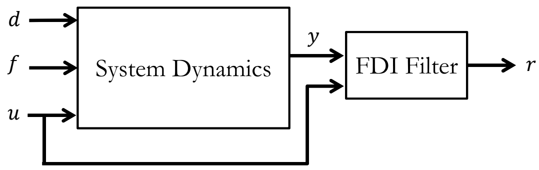

Figure 4.1: General configuration of the FDI filter

The FDI task is to design a filter whose input are the known signals (u and y) and whose output (known as the residual and denoted by r) differentiates whether the measurements are a consequence of some normal disturbance input d, or due to the fault signal f. Formally speaking, the residual can be viewed as a function $r(d, f)$, and the FDI design is ideally translated as the mapping requirements

$$d \mapsto r(d, 0) \equiv 0, \tag{4.1a}$$

$$f \mapsto r(d, f) \neq 0, \quad \forall d \tag{4.1b}$$

where condition (4.1a) ensures that the residual of the filter, r, is not excited when the system is perturbed by normal disturbances d, while condition (4.1b) guarantees the filter sensitivity to the fault f in the presence of any disturbance d.

The state of the art in FDI concentrates on the system dynamics, and imposes restrictions to provide theoretical guarantees for the required mapping conditions (4.1). For example, the authors in [NF06] restrict the system to linear dynamics, whereas [HKEY99, PI01] treat nonlinear systems but impose necessary conditions in terms of a certain distribution connected to their dynamics. In an attempt to relax the perfect decoupling condition, one may consider the worst case scenario of the mapping (4.1) in a robust formulation as follows:

$$\text{RP}: \begin{cases} \min\limits_{\gamma} & \gamma \\ \text{s.t.} & \big\|r(d,0)\big\| \leq \gamma, \qquad \forall d \in \mathcal{D} \\ & f \mapsto r(d,f) \neq 0, \quad \forall d \in \mathcal{D}, \end{cases} \tag{4.2}$$

where $\mathcal{D}$ is set of normal disturbances, γ is the alarm threshold of the designed filter, and the minimization is running over a given class of FDI filters. In view of formulation (4.2), an alarm is only raised whenever the residual exceeds γ, i.e., the filter avoids any false alarm. This, however, comes at the cost of missed detections of the faults whose residual is not bigger than the threshold γ. In the literature, the robust perspective RP has also been studied in order for a trade-off between disturbance rejection and fault sensitivity for a certain class of dynamics, e.g., see [CP82, Section 9.2] for bilinear dynamics and [FF12] for multivariate polynomial systems.

4.2.2 New Design Perspective

Here we shift our attention from the system dynamics to the class of unknown inputs $\mathcal{D}$. We assume that the disturbance signal d comes from a prescribed probability space and relax the robust formulation RP by introducing probabilistic constraints instead. In this view, the performance of the FDI filter is characterized in a probabilistic fashion.

Assume that the signal d is modeled as a random variable on the prescribed probability space $(\Omega, \mathcal{F}, \mathbb{P})$, which takes values in a metric space endowed with the corresponding Borel sigma-algebra. Assume further that the class of FDI filters ensures the measurability of the mapping $d \mapsto r$ where r also belongs to a metric space. In light of this probabilistic framework, one may quantify the filter performance from different perspectives; in the following we propose two of them:

$$\text{AP}: \begin{cases} \min\limits_{\gamma} & \gamma \\ \text{s.t.} & \mathbb{E}\big[J\big(\|r(d,0)\|\big)\big] \le \gamma \\ & f \mapsto r(d,f) \neq 0, \quad \forall d \in \mathcal{D}, \end{cases} \qquad \text{CP}: \begin{cases} \min\limits_{\gamma} & \gamma \\ \text{s.t.} & \mathbb{P}\big(\|r(d,0)\| \le \gamma\big) \ge 1-\varepsilon \\ & f \mapsto r(d,f) \neq 0, \quad \forall d \in \mathcal{D}, \end{cases} \tag{4.3}$$

where $\mathbb{E}[\,\cdot\,]$ in AP is meant with respect to the probability measure $\mathbb{P}$, and $\|\cdot\|$ is the corresponding norm in the r space. The function $J : \mathbb{R}_+ \to \mathbb{R}_+$ in AP and $\varepsilon \in (0,1)$ in CP are design parameters. To control the filter residual generated by d, the payoff function J is required to be in class $\mathcal{K}_\infty$, i.e., J is strictly increasing and $J(0) = 0$ [Kha92, Def 4.2, p. 144]. Minimizations in the above optimization problems are over a class of FDI filters which is chosen a priori, as detailed in subsequent sections.

Two formulations provide different probabilistic interpretations of fault detection. The program AP stands for "*Average Performance*" and takes all possible disturbances into account, but in accordance with their occurrence probability in an averaging sense. The program CP stands for "*Chance Performance*" and ignores an ε-fraction of the disturbance patterns and only aims to optimize the performance over the rest of the disturbance space. Note that in the CP perspective, the parameter ε is an additional design parameter to be chosen a priori.

Let us highlight that the proposed perspectives rely on the probability distribution $\mathbb{P}$, which requires prior information about possible disturbance patterns. That is, unlike the existing literature, the proposed design prioritizes between disturbance patterns in terms of their occurrence likelihood. From a practical point of view this requirement may be natural; in Section 4.5 we will describe an application of this nature.

4.2.3 Outline of the Proposed Methodology

We employ randomized algorithms to tackle the formulations in (4.3). We generate n independent and identically distributed (i.i.d.) scenarios $(d_i)_{i=1}^n$ from the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and consider the following optimization problems as random counterparts of those in (4.3):

$$\widetilde{\mathrm{AP}}: \begin{cases} \min\limits_{\gamma} & \gamma \\ \text{s.t.} & \frac{1}{n}\sum_{i=1}^{n} J\big(\|r(d_i,0)\|\big) \le \gamma \\ & f \mapsto r(d,f) \neq 0, \quad \forall d \in \mathcal{D} \end{cases} \qquad \widetilde{\mathrm{CP}}: \begin{cases} \min\limits_{\gamma} & \gamma \\ \text{s.t.} & \max\limits_{i\le n} \|r(d_i,0)\| \le \gamma \\ & f \mapsto r(d,f) \neq 0, \quad \forall d \in \mathcal{D}, \end{cases} \tag{4.4}$$

Notice that the optimization problems $\widetilde{\mathrm{AP}}$ and $\widetilde{\mathrm{CP}}$ are naturally stochastic as they depend on the generated scenarios $(d_i)_{i=1}^n$, which is indeed a random variable defined on n-fold product probability space $(\Omega^n, \mathcal{F}^n, \mathbb{P}^n)$. Therefore, their solutions are also random variables. In this work, we first restrict the FDI filters to a class of linear operators in which the random programs (4.4) are effectively convex, and hence tractable. In this step, the FDI filter is essentially robustified to n signatures of the dynamic nonlinearity. Subsequently, invoking existing results on randomized optimization (e.g., [CG08, Cal10, Han12]) we will provide probabilistic guarantees on the relation of programs (4.3) and their probabilistic counterparts in (4.4), whose precision is characterized in terms of the number of scenarios n. We should highlight that the design parameter ε of CP in (4.3) does not explicitly appear in the random counterpart $\widetilde{\mathrm{CP}}$ in (4.4). However, as we will clarify in 4.4.3, the parameter ε contributes to the probabilistic guarantees of the design.

4.3 Model Description and Basic Definitions

In this section we introduce a class of nonlinear models along with some basic definitions, which will be considered as the system dynamics in Figure 4.1 throughout the chapter. Consider the nonlinear differential-algebraic equation (DAE) model

$$E(x) + H(p)x + L(p)z + F(p)f = 0, \tag{4.5}$$

where the signals x, z, f are assumed to be piece-wise continuous (p.w.c.) functions from $\mathbb{R}_+$ into $\mathbb{R}^{n_x}, \mathbb{R}^{n_z}, \mathbb{R}^{n_f}$, respectively; we denote the spaces of such signals by $\mathcal{W}^{n_x}, \mathcal{W}^{n_z}, \mathcal{W}^{n_f}$, respectively. Let n_r be the number of rows in (4.5), and $E : \mathbb{R}^{n_x} \to \mathbb{R}^{n_r}$ be a Lipschitz continuous mapping. The operator p is the distributional derivative operator [Ada75, Section I], and H, L, F are polynomial matrices in the operator p with n_r rows and n_x, n_z, n_f columns, respectively. In the setup of Figure 4.1, the signal x represents all unknowns signals, e.g., internal states of the system dynamics and unknown disturbances d. The signal z contains all known signals, i.e., it is an augmented signal including control input u and available measurements y. The signal f stands for faults or intrusion which is the target of detection. We refer to [Shc07] and the references therein for general theory of nonlinear DAE systems and the regularity of their solutions.

One may extend the space of functions x, z, f to Sobolev spaces, but this is outside the scope of our study. On the other hand, if these spaces are restricted to the (resp. right) smooth functions, then the operator p can be understood as the classical (resp. right) differentiation operator. Throughout this chapter we will focus on continuous-time models, but one can obtain similar results for discrete-time models by changing the operator p to the time-shift operator. We will think of the matrices $H(p)$, $L(p)$ and $F(p)$ above either as linear operators on the function spaces (in which case p will be interpreted as a generalized derivative operator as explained above) or as algebraic objects (in which case p will be interpreted as simply a complex variable). The reader is asked to excuse this slight abuse of the notation, but the interpretation should be clear from the context.

Let us first show the generality of the DAE framework of (4.5) by the following example. Consider the classical nonlinear ordinary differential equation

$$\begin{cases} G\dot{X}(t) & = E_X\big(X(t), d(t)\big) + AX(t) + B_u u(t) + B_d d(t) + B_f f(t) \\ Y(t) & = E_Y\big(X(t), d(t)\big) + CX(t) + D_u u(t) + D_d d(t) + D_f f(t) \end{cases} \tag{4.6}$$

where $u(\cdot)$ is the input signal, $d(\cdot)$ the unknown disturbance, $Y(\cdot)$ the measured output, $X(\cdot)$ the internal variables, and $f(\cdot)$ a faults (or an attack) signal to be detected. Parameters $G, A, B_u, B_d, B_f, D_u, D_d, D_f$ are constant matrices and functions E_X, E_Y are Lipschitz continuous mappings with appropriate dimensions. One can easily fit the model (4.6) into the DAE framework of (4.5) by defining

$$x := \begin{bmatrix} X \\ d \end{bmatrix}, \qquad z := \begin{bmatrix} Y \\ u \end{bmatrix},$$

$$E(x) := \begin{bmatrix} E_X(x) \\ E_Y(x) \end{bmatrix}, \quad H(p) := \begin{bmatrix} -pG + A & B_d \\ C & D_d \end{bmatrix}, \quad L(p) := \begin{bmatrix} 0 & B_u \\ -I & D_u \end{bmatrix}, \quad F(p) := \begin{bmatrix} B_f \\ D_f \end{bmatrix}.$$

Following [NF06], with a slight extension to a nonlinear dynamics, let us formally characterize all possible observations of the model (4.5) in the absence of the fault signal f:

$$\mathcal{M} := \big\{ z \in \mathcal{W}^{n_z} \,\big|\, \exists x \in \mathcal{W}^{n_x} : \quad E(x) + H(p)x + L(p)z = 0 \big\}; \tag{4.7}$$

This set is known as the *behavior* of the system [PW98].

Definition 4.3.1 (Residual Generator)**.** *A proper linear time invariant filter $r := R(p)z$ is a residual generator for* (4.5) *if for all $z \in \mathcal{M}$, it holds that $\lim_{t\to\infty} r(t) = 0$.*

Note that by Definition 4.3.1 the class of residual generators in this study is restricted to a class of *linear* transfer functions where $R(p)$ is a matrix of proper rational functions of p.

Definition 4.3.2 (Fault Sensitivity)**.** *The residual generator introduced in Definition 4.3.1 is sensitive to fault f_i if the transfer function from f_i to r is nonzero, where f_i is the i^{th} elements of the signal f.*

One can inspect that Definition 4.3.1 and Definition 4.3.2 essentially encode the basic mapping requirements (4.1a) and (4.1b), respectively.

4.4 Fault Detection and Isolation Filters

The main objective of this section is to establish a scalable framework geared towards the design perspectives AP and CP as explained in Section 4.2. To this end, we first review a polynomial characterization of the residual generators and its linear program formulation counterpart for linear systems (i.e., the case where $E(x) \equiv 0$). We then extend the approach to the nonlinear model (4.5) to account for the contribution of $E(\cdot)$ to the residual, and subsequently provide probabilistic performance guarantees for the resulting filter.

4.4.1 Residual Generators for Linear Systems

In this subsection we assume $E(x) \equiv 0$, i.e., we restrict our attention to the class of linear DAEs. One can observe that the behavior set $\mathcal{M}$ can alternatively be defined as

$$\mathcal{M} = \big\{z \in \mathcal{W}^{n_z} \big| \ N_H(p)L(p)z = 0\big\},$$

where the collection of the rows of $N_H(p)$ forms an irreducible polynomial basis for the left null-space of the matrix $H(p)$ [PW98, Section 2.5.2]. This representation allows one to describe the residual generators in terms of polynomial matrix equations. That is, by picking a linear combination of the rows of $N_H(p)$ and considering an arbitrary polynomial $a(p)$ of sufficiently high order with roots with negative real parts, we arrive at a residual generator in the sense of Definition 4.3.1 with transfer operator

$$R(p) = a^{-1}(p)\gamma(p)N_H(p)L(p) := a^{-1}(p)N(p)L(p). \tag{4.8}$$

The above filter can easily be realized by an explicit state-space description with input z and output r. Multiplying the left hand-side of (4.5) by $a^{-1}(p)N(p)$ leads to

$$r = -a^{-1}(p)N(p)F(p)f.$$

Thus, a sensitive residual generator, in the sense of Definition 4.3.1 and Definition 4.3.2, is characterized by the polynomial matrix equations

$$N(p)H(p) = 0, \tag{4.9a}$$

$$N(p)F(p) \neq 0, \tag{4.9b}$$

where (4.9a) implements condition (4.1a) above (cf. Definition 4.3.1) while (4.9b) implements condition (4.1b) (cf. Definition 4.3.2). Both row polynomial vector $N(p)$ and denominator polynomial $a(p)$ can be viewed as design parameters. Throughout this study we, however, fix $a(p)$ and aim to find an optimal $N(p)$ with respect to a certain objective criterion related to the filter performance.

In case there are more than one faults ($n_f > 1$), it might be of interest to isolate the impact of one fault in the residual from the others. The following remark implies that the isolation problem is effectively a detection problem.

Remark 4.4.1 (Fault Isolation)**.** *Consider model (4.5) and suppose $n_f > 1$. In order to detect only one of the fault signals, say f_1, and isolate it from the other faults, $f_i, i \in \{2, \cdots, n_f\}$, one may consider the detection problem for the same model but in new representation*

$$E(x) + [H(p) \ \ \widetilde{F}(p)] \begin{bmatrix} x \\ \tilde{f} \end{bmatrix} + L(p)z + F_1(p)f = 0,$$

where $F_1(p)$ is the first column of $F(p)$, and $\widetilde{F}(p) := [F_2(p), \cdots, F_{n_f}(p)]$, and $\tilde{f} := [f_2, \cdots, f_{n_f}]$, see [FKA09, Thm. 2] for more details on fault isolation.

Next, we show how to transform the possibly complex matrix polynomial equations (4.9) into a linear programming framework.

Lemma 4.4.2 (Linear Programming Characterization)**.** *Let $N(p)$ be a feasible polynomial matrix of degree d_N for the inequalities (4.9), where*

$$H(p) := \sum_{i=0}^{d_H} H_i p^i, \quad F(p) := \sum_{i=0}^{d_F} F_i p^i, \quad N(p) := \sum_{i=0}^{d_N} N_i p^i,$$

and $H_i \in \mathbb{R}^{n_r \times n_x}$, $F_i \in \mathbb{R}^{n_r \times n_f}$, and $N_i \in \mathbb{R}^{1 \times n_r}$ are constant matrices. Then, the polynomial matrix inequalities (4.9) are equivalent, up to a scalar, to

$$\bar{N}\bar{H} = 0, \tag{4.10a}$$

$$\left\|\bar{N}\bar{F}\right\|_\infty \geq 1, \tag{4.10b}$$

where $\|\cdot\|_\infty$ is the infinity vector norm, and

$$\bar{N} := \begin{bmatrix} N_0 & N_1 & \cdots & N_{d_N} \end{bmatrix}$$

$$\bar{H} := \begin{bmatrix} H_0 & H_1 & \cdots & H_{d_H} & 0 & \cdots & 0 \\ 0 & H_0 & H_1 & \cdots & H_{d_H} & 0 & \vdots \\ \vdots & & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & H_0 & H_1 & \cdots & H_{d_H} \end{bmatrix},$$

$$\bar{F} := \begin{bmatrix} F_0 & F_1 & \cdots & F_{d_F} & 0 & \cdots & 0 \\ 0 & F_0 & F_1 & \cdots & F_{d_F} & 0 & \vdots \\ \vdots & & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & F_0 & F_1 & \cdots & F_{d_F} \end{bmatrix}.$$

Proof. It is easy to observe that

$$\begin{aligned} N(p)H(p) &= \bar{N}\bar{H}[I \; pI \; \cdots \; p^i I]^\intercal, \qquad & i := d_N + d_H, \\ N(p)F(p) &= \bar{N}\bar{F}[I \; pI \; \cdots \; p^j I]^\intercal, \qquad & j := d_N + d_F. \end{aligned}$$

Moreover, in light of the linear structure of equations (4.9), one can scale the inequality (4.9b) and arrive at the assertion of the lemma. □

Remark 4.4.3. *Strictly speaking, the formulation in Lemma 4.4.2 is not a linear program, due to the non-convex constraint (4.10b). It is, however, easy to see that if $\bar{N}$ is a solution to (4.10), then so is $-\bar{N}$. Hence, the inequality (4.10b) can be understood as a family of m linear programs where $m = n_f(d_F + d_N + 1)$ is the number of columns of $\bar{F}$, and n_f is the dimension* *of signal f in (4.5). Each of these linear programs focuses on a component of the vector $\bar{N}\bar{F}$, replacing the inequality (4.10b) with*

$$\bar{N}\bar{F}v \geq 1, \qquad v := [0, \cdots, 1, \cdots, 0]^\intercal.$$

Fact 4.4.4. *There exists a solution $N(p)$ to (4.9) if and only if Rank $[H(p)\ F(p)] >$ Rank $H(p)$.*

Fact 4.4.4 provides necessary and sufficient conditions for the feasibility of the linear program formulation in Lemma 4.4.2; proof is omitted as it is an easy adaptation of the one in [FKA09, Corollary 3].

4.4.2 Extension to Nonlinear Systems

In the presence of nonlinear terms $E(x) \neq 0$, it is straightforward to observe that the residual of filter (4.8) consists of two terms:

$$r := R(p)z = -\underbrace{a^{-1}(p)N(p)F(p)f}_{(i)} - \underbrace{a^{-1}(p)N(p)E(x)}_{(ii)}. \tag{4.11}$$

Term (i) is the desired contribution of the fault f and is in common with the linear setup. Term (ii) is due to the nonlinear term $E(\cdot)$ in (4.5). Our aim here will be to reduce the impact of $E(x)$ while increasing the sensitivity to the fault f. To achieve this objective, we develop two approaches to control each of the two terms separately; in both cases we assume that the degree of the filter (i.e., d_N in Lemma 4.4.2) and the denominator (i.e., $a(p)$ in (4.11)) are fixed, and the aim is to design the numerator coefficients (i.e., $N(p)$ in (4.11)).

Approach (I): Fault Sensitivity

We first focus on fault sensitivity and neglect the contribution of the nonlinear term. To this end, we assume that the system operates close to an equilibrium point $x_e \in \mathbb{R}^{n_x}$. Even though in case of a fault the system may eventually deviate substantially from its nominal operating point, if the FDI filter succeeds in identifying the fault early, this assumption may still be valid over the time horizon relevant for fault detection. We assume, without loss of generality, that

$$\lim_{x \to x_e} \frac{\|E(x)\|_2}{\|x - x_e\|_2} = 0,$$

where $\|\cdot\|_2$ stands for the Euclidean norm of a vector. If this is not the case, the linear part of $E(\cdot)$ can be extracted and included in the linear part of the system.

To increase the sensitivity of the linear filter to the fault f, we revisit the linear programming formulation (4.10) and seek a feasible numerator $N(p)$ such that the coefficients of the transfer function $N(p)F(p)$ attain maximum values within the admissible range. This gives rise to the following optimization problem:

$$\begin{cases} \max\limits_{\bar{N}} & \|\bar{N}\bar{F}\|_\infty \\ \text{s.t.} & \bar{N}\bar{H} = 0 \\ & \|\bar{N}\|_\infty \leq 1 \end{cases} \tag{4.12}$$

where the objective function targets the contribution of the signal f to the residual r. Let us recall that $\bar{N}\bar{F}$ is the vector containing all numerator coefficients of the transfer function $f \mapsto r$. The second constraint in (4.12) is added to ensure that the solutions remain bounded; note that thanks to the linearity of the filter this constraint does not influence the performance. Though strictly speaking (4.12) is not a linear program, in view of Remark 4.4.3 it is easy to transform it to a family of m different linear programs, where m is the number of columns of $\bar{F}$.

The main problem with this approach is the lack of prior information when the linearization technique is "precise enough". That is, how well the filter designed by (4.12) will work depends on the magnitude of the second term in (4.11), which is due to the nonlinearities $E(x)$ and is ignored in (4.12). If the term generated by $E(x)$ is large enough for nominal excursions of x from x_e, the filter may lead to false alarms, whereas if we set our thresholds high to tolerate the disturbance generated by $E(x)$ in nominal conditions, the filter may lead to missed detections. A direct way toward controlling this trade-off involving the nonlinear term will be the focus of the second approach.

Approach (II): Robustify to Nonlinearity Signatures

This approach is the main step toward the theoretical contribution of the chapter, and provides the principle ingredients to tackle the proposed perspectives AP and CP introduced in (4.3). The focus is on term (ii) of the residual (4.11), in relation to the mapping (4.1a). The idea is to robustify the filter against certain signatures of the nonlinearity during nominal operation. In the following we restrict the class of filters to the feasible solutions of polynomial matrix equations (4.9), characterized in Lemma 4.4.2.

Let us denote the space of all p.w.c. functions from the interval $[0, T]$ to $\mathbb{R}^n$ by $\mathcal{W}_T^n$. We equip this space with the $\mathcal{L}_2$-inner product and the corresponding norm

$$\|e\|_{\mathcal{L}_2} := \sqrt{\langle e, e\rangle}, \qquad \langle e, g\rangle := \int_0^T e^\intercal(t) g(t) \mathrm{d}t, \quad e, g \in \mathcal{W}_T^n.$$

Consider an unknown signal $x \in \mathcal{W}_T^{n_x}$. In the context of the ODEs (4.6) that means we excite the system with the disturbance $d(\cdot)$ for the time horizon T. We then stack $d(\cdot)$ together with the internal state $X(\cdot)$ to introduce $x := \left[\begin{smallmatrix} X \\ d \end{smallmatrix}\right]$. We define the signals $e_x \in \mathcal{W}_T^{n_r}$ and $r_x \in \mathcal{W}_T^1$ as follows:

$$e_x(t) := E\big(x(t)\big), \qquad r_x(t) := -a^{-1}(p) N(p)[e_x](t), \qquad \forall t \in [0, T]. \tag{4.13}$$

The signal e_x is the "*nonlinearity signature*" in the presence of the unknown signal x, and the signal r_x is the contribution of the nonlinear term to the residual of the linear filter. Our goal now is to minimize $\|r_x\|_{\mathcal{L}_2}$ in an optimization framework in which the coefficients of polynomial $N(p)$ are the decision variables and the denominator $a(p)$ is a fixed stable polynomial with the degree at least the same as $N(p)$.

Lemma 4.4.5. *Let $N(p)$ be a polynomial row vector of dimension n_r and degree d_N, and $a(p)$ be a stable scalar polynomial with the degree at least d_N. For any $x \in \mathcal{W}_T^{n_x}$ there exists*

$\psi_x \in \mathcal{W}_T^{n_r(d_N+1)}$ such that

$$r_x(t) = \bar{N}\psi_x(t), \qquad \forall t \in [0,T] \tag{4.14a}$$

$$\|\psi_x\|_{\mathcal{L}_2} \le C\|e_x\|_{\mathcal{L}_2}, \qquad C := \sqrt{n_r(d_N+1)}\|a^{-1}\|_{\mathcal{H}_\infty}, \tag{4.14b}$$

where $\bar{N}$ is the vector collecting all the coefficients of the numerator $N(p)$ as introduced in Lemma 4.4.2, and the signals e_x and r_x are defined as in (4.13).

Proof. See Appendix 4.7.1. □

Given $x \in \mathcal{W}_T^{n_x}$ and the corresponding function ψ_x as defined in Lemma 4.4.5, we have

$$\|r_x\|_{\mathcal{L}_2}^2 = \bar{N}Q_x\bar{N}^\intercal, \qquad Q_x := \int_0^T \psi_x(t)\psi_x^\intercal(t)\mathrm{d}t. \tag{4.15}$$

We call Q_x the "*signature matrix*" of the nonlinearity signature $t \mapsto e_x(t)$ resulting from the unknown signal x. Given x and the corresponding signature matrix Q_x, the $\mathcal{L}_2$-norm of r_x in (4.13) can be minimized by considering an objective which is a quadratic function of the filter coefficients $\bar{N}$ subject to the linear constraints in (4.10):

$$\begin{cases} \min\limits_{\bar{N}} & \bar{N}Q_x\bar{N}^\intercal \\ \text{s.t.} & \bar{N}\bar{H} \quad = 0 \\ & \left\|\bar{N}\bar{F}\right\|_\infty \ge 1 \end{cases} \tag{4.16}$$

The program (4.16) is not a true quadratic program due to the second constraint. Following Remark 4.4.3, however, one can show that the optimization program (4.16) can be viewed as a family of m quadratic programs where $m = n_f(d_F + d_N + 1)$.

In the rest of the subsection, we establish an algorithmic approach to approximate the matrix Q_x for a given $x \in \mathcal{W}_T^{n_x}$, with an arbitrary high precision. We first introduce a finite dimensional subspace of $\mathcal{W}_T^1$ denoted by

$$\mathcal{B} := span\{b_0, b_1, \cdots, b_k\}, \tag{4.17}$$

where the collection of $b_i : [0,T] \to \mathbb{R}$ is a basis for $\mathcal{B}$. Let $\mathcal{B}^{n_r} := \bigotimes_{i=1}^{n_r}\mathcal{B}$ be the n_r Cartesian product of the set $\mathcal{B}$, and $\mathbb{T}_\mathcal{B} : \mathcal{W}_T^{n_r} \to \mathcal{B}^{n_r}$ be the $\mathcal{L}_2$-orthogonal projection operator onto $\mathcal{B}^{n_r}$, i.e.,

$$\mathbb{T}_\mathcal{B}(e_x) = \sum_{i=0}^k \beta_i^\star b_i, \qquad \beta^\star := \arg\min_\beta \left\|e_x - \sum_{i=0}^k \beta_i b_i\right\|_{\mathcal{L}_2} \tag{4.18}$$

Let us remark that if the basis of $\mathcal{B}$ is orthonormal (i.e., $\langle b_i, b_j\rangle = 0$ for $i \neq j$), then $\beta_i^\star = \int_0^T b_i(t)e_x(t)\mathrm{d}t$; we refer to [Lue69, Sec. 3.6] for more details on the projection operator.

Assumption 4.4.6. *We stipulate that*

(i) *The basis functions b_i of subspace $\mathcal{B}$ are smooth and $\mathcal{B}$ is closed under the differentiation operator p, i.e., for any $b \in \mathcal{B}$ we have $p[b] = \frac{\mathrm{d}}{\mathrm{d}t}b \in \mathcal{B}$.*

(ii) *The basis vectors in (4.17) are selected from an $\mathcal{L}_2$-complete basis for $\mathcal{W}_T^1$, i.e., for any $e \in \mathcal{W}_T^{n_r}$, the projection error $\|e - \mathbb{T}_\mathcal{B}(e)\|_{\mathcal{L}_2}$ can be made arbitrarily small by increasing the dimension, k, of subspace $\mathcal{B}$.*

The requirements of Assumptions 4.4.6 can be fulfilled for subspaces generated by, for example, the polynomial or Fourier basis. Thanks to Assumption 4.4.6(i), the linear operator p can be viewed as a matrix operator. That is, there exists a square matrix D with dimension $k+1$ such that

$$p[B(t)] = \frac{\mathrm{d}}{\mathrm{d}t}B(t) = DB(t), \qquad B(t) := [b_0(t), \cdots, b_k(t)]^\intercal. \tag{4.19}$$

In Section 4.5.2 we will provide an example of such matrix operator for the Fourier basis. By virtue of matrix representations of (4.19) we have

$$N(p)\mathbb{T}_\mathcal{B}(e_x) = \sum_{i=0}^{d_N} N_i p^i \beta^\star B = \sum_{i=0}^{d_N} N_i \beta^\star D^i B = \bar{N}\bar{D}B, \qquad \bar{D} := \begin{bmatrix} \beta^\star \\ \beta^\star D \\ \vdots \\ \beta^\star D^{d_N} \end{bmatrix}, \tag{4.20}$$

where the vector $\beta^\star := [\beta_0^\star, \cdots, \beta_k^\star]$ is introduced in (4.18). Let us define the positive semidefinite matrix $G := [G_{ij}]$ of dimension $k+1$ by

$$G_{ij} := \left\langle a^{-1}(p)[b_i], a^{-1}(p)[b_j] \right\rangle \tag{4.21}$$

Therefore, with the aid of equations (4.20) and (4.21) we arrive at

$$\left\|a^{-1}(p)N(p)\mathbb{T}_\mathcal{B}(e)\right\|_{\mathcal{L}_2}^2 = \bar{N}Q_\mathcal{B}\bar{N}^\intercal, \qquad Q_\mathcal{B} := \bar{D}G\bar{D}^\intercal, \tag{4.22}$$

where $\bar{D}$ and G are defined in (4.20) and (4.21), respectively. Note that the matrices G and D are built by the data of the subspace $\mathcal{B}$ and denominator $a(p)$, whereas the nonlinearity signature only influences the coefficient $\beta^\star$. The above discussion is summarized in Algorithm 1 with an emphasis on models described by the ODE (4.6). The precision of the algorithm output is quantified in Theorem 4.4.7.

Theorem 4.4.7. *Consider an unknown signal $x : [0,T] \to \mathbb{R}^{n_x}$ in $\mathcal{W}_T^{n_x}$ and the corresponding nonlinearity signature e_x and signature matrix Q_x as defined in (4.13) and (4.15), respectively. Let $(b_i)_{i\in\mathbb{N}} \subset \mathcal{W}_T^1$ be a family of basis functions $(b_i)_{i\in\mathbb{N}} \subset \mathcal{W}_T^1$ satisfying Assumptions 4.4.6, and let $\mathcal{B}$ be the finite dimensional subspace in (4.17). If $\|e_x - \mathbb{T}_\mathcal{B}(e_x)\|_{\mathcal{L}_2} < \delta$, where $\mathbb{T}_\mathcal{B}$ is the projection operator onto $\mathcal{B}^{n_r}$, then*

$$\left\|Q_x - Q_\mathcal{B}\right\|_2 < \bar{C}\delta, \qquad \bar{C} := \left(1 + 2\|e_x\|_{\mathcal{L}_2}\right)C\|a^{-1}\|_{\mathcal{H}_\infty}, \tag{4.23}$$

where $Q_\mathcal{B}$ is obtained by (4.22) (the output of Algorithm 1), and C is the same constant as in (4.14b).

Proof. See Appendix 4.7.1. □

Algorithm 1 Computing the signature matrix Q_x in (4.15)

(i) **Initialization of the Filter Paramters:**
1. Set the filter denominator $a(p)$, the numerator degree d_N, and horizon T
2. Set the basis $\{b_i\}_{i=1}^{k} \subset \mathcal{W}_T^1$ satisfying Assumptions 4.4.6
3. Compute the differentiation matrix D (see (4.19))
4. Compute the matrix G (see (4.21))

(ii) **Identification of the Nonlinearity Signature:**
1. Input the disturbance pattern $d(\cdot)$ for time horizon T
2. Set $f \equiv 0$ in (4.6) and run the system by $d(\cdot)$ to obtain the internal state $X(\cdot)$
3. Set the unknown signal $x(t) := [X^\intercal(t),\ d^\intercal(t)]^\intercal$
4. Set the nonlinearity signature $e_x(t) := \big[E_X^\intercal\big(x(t)\big),\ E_Y^\intercal\big(x(t)\big)\big]^\intercal$

(iii) **Computation of the Signature Matrix**
1. Compute $\beta^\star$ from (4.18) (in case of orthonormal basis $\beta_i^\star = \int_0^T b_i(t)e_x(t)\mathrm{d}t$)
2. Compute $\bar{D}$ from (4.20)
3. Ouput $Q_{\mathcal{B}} := \bar{D}G\bar{D}^\intercal$ (see (4.22))

We close this subsection with a following remark, a natural extension to robustify the filter to multiple nonlinearity signatures.

Remark 4.4.8 (Multi Signatures Training)**.** *In order to robustify the FDI filter to more than one unknown signal, say $\{x_i(\cdot)\}_{i=1}^n$, one may introduce an objective function as an average cost $\bar{N}\big(\frac{1}{n}\sum_{i=1}^n Q_{x_i}\big)\bar{N}^\intercal$ or the worst case viewpoint $\max_{i\le n}\bar{N}Q_{x_i}\bar{N}^\intercal$, where Q_{x_i} is the signature matrix corresponding to x_i as defined in* (4.15)*.*

4.4.3 Proposed Methodology and Probabilistic Performance

The preceding subsection proposed two optimization-based approaches to enhance the FDI filter design from linear to nonlinear system dynamics. Approach (I) targets the fault sensitivity while neglecting the nonlinear term of the system dynamics, and Approach (II) offers a QP framework to robustify the residual with respect to signatures of the dynamic nonlinearities. Here our aim is to achieve a reconciliation between these two approaches. We subsequently provide theoretical results from the proposed solutions to the original design perspectives (4.3).

Let $(d_i)_{i=1}^n \subset \mathcal{D}$ be i.i.d. disturbance patterns generated from the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. For each d_i, let x_i be the corresponding unknown signal with the associated signature matrix Q_{x_i} as defined in (4.15). In regard to the average perspective AP, we propose the two-stage (random) optimization program

$$\widetilde{\mathrm{AP}}_1: \begin{cases} \min\limits_{\gamma,\bar{N}} & \gamma \\ \text{s.t.} & \bar{N}\bar{H} & = 0 \\ & \left\|\bar{N}\bar{F}\right\|_\infty & \geq 1 \\ & \frac{1}{n}\sum\limits_{i=1}^{n} J\Big(\sqrt{\bar{N}Q_{x_i}\bar{N}^\intercal}\Big) & \leq \gamma \end{cases} \tag{4.24a}$$

$$\widetilde{\mathrm{AP}}_2: \begin{cases} \max\limits_{\bar{N}} & \left\|\bar{N}\bar{F}\right\|_\infty \\ \text{s.t.} & \bar{N}\bar{H} & = 0 \\ & \left\|\bar{N}\right\|_\infty & \leq 1 \\ & \frac{1}{n}\sum\limits_{i=1}^{n} J\Big(\|\bar{N}_1^\star\|_\infty\sqrt{\bar{N}Q_{x_i}\bar{N}^\intercal}\Big) & \leq \gamma_1^\star \end{cases} \tag{4.24b}$$

where $J : \mathbb{R}_+ \to \mathbb{R}_+$ is an increasing and convex payoff function, and in the second stage (4.24b) $\bar{N}_1^\star$ and $\gamma_1^\star$ are the optimizers of the first stage (4.24a), i.e., the programs (4.24) need to be solved sequentially. Let us recall that the filter coefficients can always be normalized with no performance deterioration. Hence, it is straightforward to observe that the main goal of the second stage is only to improve the coefficients of $\bar{N}\bar{F}$ (concerning the fault sensitivity) while the optimality of the first stage (concerning the robustification to nonlinearity signatures) is guaranteed. Similarly, we also propose the following two-stage program for the perspective CP:

$$\widetilde{\mathrm{CP}}_1: \begin{cases} \min\limits_{\gamma,\bar{N}} & \gamma \\ \text{s.t.} & \bar{N}\bar{H} & = 0 \\ & \left\|\bar{N}\bar{F}\right\|_\infty & \geq 1 \\ & \max\limits_{i\leq n} \bar{N}Q_{x_i}\bar{N}^\intercal & \leq \gamma \end{cases} \tag{4.25a}$$

$$\widetilde{\mathrm{CP}}_2: \begin{cases} \max\limits_{\bar{N}} & \left\|\bar{N}\bar{F}\right\|_\infty \\ \text{s.t.} & \bar{N}\bar{H} & = 0 \\ & \left\|\bar{N}\right\|_\infty & \leq 1 \\ & \|\bar{N}_1^\star\|_\infty^2\big(\max\limits_{i\leq n} \bar{N}Q_{x_i}\bar{N}^\intercal\big) & \leq \gamma_1^\star \end{cases} \tag{4.25b}$$

On the computational side, in view of Remark 4.4.3, all the programs in (4.24) and (4.25) are effectively convex, and hence tractable. In the rest of the subsection we establish a probabilistic bridge between the solutions to the program (4.24) (resp. (4.25)) and the original perspective AP (resp. CP) in (4.3) when the class of filters is confined to the linear residuals characterized in Lemma 4.4.2. For this purpose, we need a technical measurability assumption which is always expected to hold in practice.

Assumption 4.4.9 (Measurability)**.** *We assume that the mapping $\mathcal{D} \ni d \mapsto x \in \mathcal{W}_T^{n_x}$ is measurable.*[1] *In particular, x can be viewed as a random variable on the same probability space as d.*

[1]The function spaces are endowed with the $\mathcal{L}_2$-topology and the respective Borel sigma-algebra.

Assumption 4.4.9 is referred to the behavior of the system dynamics as a mapping from the disturbance d to the internal states. In the context of ODEs (4.6), it is well-known that under mild assumptions (e.g., Lipschitz continuity of E_X) the mapping $d \mapsto X$ is indeed continuous [Kha92, Chap. 5], which readily ensures Assumption 4.4.9.

Probabilistic performance of $\widetilde{\mathrm{AP}}$:

Here we study the asymptotic behavior of the empirical average of $\mathbb{E}\big[J(\|r\|)\big]$ uniformly in the filter coefficients $\bar{N}$, which allows us to link the solutions of programs (4.24) to AP. Let $\mathcal{N} := \big\{\bar{N} \in \mathbb{R}^{nr(d_N+1)} : \|\bar{N}\|_\infty \leq 1\big\}$ and consider the payoff function of AP in (4.3) as the mapping $\phi : \mathcal{N} \times \mathcal{W}_T^{n_x} \to \mathbb{R}_+$:

$$\phi(\bar{N}, x) := J\big(\|r_x\|_{\mathcal{L}_2}\big) = J\big(\|\bar{N}\psi_x\|_{\mathcal{L}_2}\big), \tag{4.26}$$

where the second equality follows from Lemma 4.4.5.

Theorem 4.4.10 (Average Performance)**.** *Suppose Assumption 4.4.9 holds and the random variable x is almost surely bounded*[2]*. Then, the mapping $\bar{N} \mapsto \phi(\bar{N}, x)$ is a random function. Moreover, if $(x_i)_{i=1}^n \subset \mathcal{W}_T^{n_x}$ are i.i.d. random variables and e_n is the uniform empirical average error*

$$e_n := \sup_{\bar{N} \in \mathcal{N}} \Big\{ \frac{1}{n} \sum_{i=1}^{n} \phi(\bar{N}, x_i) - \mathbb{E}\big[\phi(\bar{N}, x)\big] \Big\}, \tag{4.27}$$

then,

(i) the Strong Law of Large Numbers (SLLN) holds, i.e., $\lim_{n\to\infty} e_n = 0$ almost surely.

(ii) the Uniform Central Limit Theorem (UCLT) holds, i.e., $\sqrt{n}e_n$ converges in law to a Gaussian variable with distribution $N(0, \sigma)$ for some $\sigma \geq 0$.

Proof. See Appendix 4.7.2 along with required preliminaries. □

The following Corollary is an immediate consequence of the UCLT in Theorem 4.4.10 (ii).

Corollary 4.4.11. *Let assumptions of Theorem 4.4.10 hold, and e_n be the empirical average error (4.27). For all $\varepsilon > 0$ and $k < \frac{1}{2}$, we have*

$$\lim_{n\to\infty} \mathbb{P}^n\big(n^k e_n \geq \varepsilon\big) = 0,$$

where $\mathbb{P}^n$ denotes the n-fold product probability measure on Ω^n.

[2]This assumption may be relaxed in terms of the moments of x, tough this will not be pursued further here.

Probabilistic performance of $\widetilde{\text{CP}}$:

The formulation CP in (4.3) is known as chance constrained program which has received increasing attention due to recent developments toward tractable approaches, in particular via the scenario counterpart (cf. $\widetilde{\text{CP}}$ in (4.4)) [CC05, CC06, CG08, Cal10]. The crucial requirement to invoke these results is the convexity of the optimization program to the decision variables. Due to the non-convexity arising from the constraint $\|\bar{N}\bar{F}\|_\infty \geq 1$, these studies are not directly applicable to our problem. Here our aim is to exploit the specific structure of this non-convexity and adapt the aforesaid results accordingly.

Let $(\bar{N}_n^\star, \gamma_n^\star)$ be the optimizer obtained through the two-stage programs (4.25) where $\bar{N}_n^\star$ is the filter coefficients and $\gamma_n^\star$ represents the filter threshold; n is referred to the number of disturbance patterns. Given the filter $\bar{N}_n^\star$, let us denote the corresponding filter residual due to the signal x by $r_x[\bar{N}_n^*]$; this is a slight modification of our notation r_x in (4.13) to specify the filter coefficients. To quantify the filter performance, one may ask for the probability that a new unknown signal x violates the threshold $\gamma_n^\star$ when the FDI filter is set to $\bar{N}_n^\star$ (i.e., $\|r_x[\bar{N}_n^*]\|_{\mathcal{L}_2}^2 > \gamma_n^*$). In the FDI literature this violation is known as false alarm, and from the CP standpoint its occurrence probability is allowed at most to the ε level. In this view the performance of the filter can be quantified by the event

$$\mathcal{E}(\bar{N}_n^\star, \gamma_n^*) := \Big\{\mathbb{P}\Big(\|r_x[\bar{N}_n^*]\|_{\mathcal{L}_2}^2 > \gamma_n^*\Big) > \varepsilon\Big\}. \tag{4.28}$$

The event (4.28) accounts for the feasibility of the $\widetilde{\text{CP}}$ solution from the original perspective CP. Note that the measure $\mathbb{P}$ in (4.28) is referred to x whereas the stochasticity of the event stems from the random solutions $(\bar{N}_n^\star, \gamma_n^\star)$.[3] We proceed with the main result of this part in regard to the likelihood of the event (4.28).

Theorem 4.4.12 (Chance Performance)**.** *Suppose Assumption 4.4.9 holds and $(x_i)_{i=1}^n$ are i.i.d. random variables on $(\Omega, \mathcal{F}, \mathbb{P})$. Let $\bar{N}_n^\star \in \mathbb{R}^{n_r(d_N+1)}$ and $\gamma_n^* \in \mathbb{R}_+$ be the solutions of $\widetilde{\text{CP}}$, and measurable in $\mathcal{F}^n$. Then, the set (4.28) is $\mathcal{F}^n$-measurable, and for every $\beta \in (0,1)$ and any n such that*

$$n \geq \frac{2}{\varepsilon}\Big(\ln\frac{n_f(d_F+d_N+1)}{\beta} + n_r(d_N+1)+1\Big),$$

where d_N is the degree of the filter and n_f, n_r, d_F are the system size parameters of (4.5), then we have

$$\mathbb{P}^n\Big(\mathcal{E}(\bar{N}_n^\star, \gamma_n^*)\Big) < \beta.$$

Proof. See Appendix 4.7.2. □

4.5 Cyber-Physical Security of Power Systems: AGC Case Study

In this section, we illustrate the performance of our theoretical results to detect a cyber intrusion in a two-area power system. Motivated by our earlier studies [MEVM+10, MEVM+11], we

[3]The measure $\mathbb{P}$ is, with slight abuse of notation, the induced measure via the mapping addressed in Assumption 4.4.9

consider the IEEE 118-bus power network equipped with primary and secondary frequency control, whose essential objective is to regulate frequency and power exchange between the controlled areas. While the primary frequency control is implemented locally, the secondary loop, referred also as AGC (Automatic Generation Control), is closed over the SCADA system without human operator intervention. As investigated in [MEVM+10], the aforesaid interaction with IT infrastructure may give rise to cyber intrusion that causes unacceptable frequency deviations with potential consequences toward load shedding or generation tripping. If the intrusion is, however, detected on time, one may prevent further damage by disconnecting the AGC. Thus, the FDI scheme may offer a protection layer to address this security concern. To achieve this goal, invoking the proposed FDI methodology, we construct an FDI filter to utilize the available measurements to diagnose the AGC intrusion sufficiently fast, despite the presence of unknown load deviations.

4.5.1 Mathematical Model Description

In this section a multi-machine power system, based only on frequency dynamics, is described [Andb]. The system is arbitrarily divided into two control areas. The generators are equipped with primary frequency control and each area is under AGC which adjusts the generating setpoints of specific generators so as to regulate frequency and maintain the power exchange between the two areas to its scheduled value.

A. System description

We consider a system comprising of n buses and g number of generators. Let $G = \{i\}_1^g$ denote the set of generator indices and $A_1 = \{i \in G \mid i \text{ in Area } 1\}$, $A_2 = \{i \in G \mid i \text{ in Area } 2\}$ the sets of generators that belong to Area 1 and Area 2, respectively. Let also $L_{tie}^k = \{(i,j)|i,j \text{ edges of a tie line from area } k \text{ to the other areas}\}$ where a tie line is a line connecting the two independently controlled areas and let also $K = \{1, 2\}$ be the set of the indices of the control areas in the system.

Using the classical generator model every synchronous machine is modeled as constant voltage source behind its transient reactance. The dynamic states of the system are the rotor angle δ_i (rad), the rotor electrical frequency f_i (Hz) and the mechanical power (output of the turbine) P_{mi} (MW) for each generator $i \in G$. We also have one more state that represents the output of the AGC ΔP_{agc_k} for each control area $k \in K$.

We denote by $E_G \in \mathbb{C}^g$ a vector consisting of the generator internal node voltages $E_{Gi} = |E_{Gi}^0|\angle\delta_i$ for $i \in G$. The phase angle of the generator voltage node is assumed to coincide with the rotor angle δ_i and $|E_{Gi}^0|$ is a constant. The voltages of the rest of the nodes are included in $V_N \in \mathbb{C}^n$, whose entries are $V_{Ni} = |V_{Ni}|\angle\theta_i$ for $i = 1, \ldots, n$. To remove the algebraic constraints that appear due to the Kirchhoff's first law for each node, we retain the internal nodes (behind the transient reactance) of the generators and eliminate the rest of the nodes. This could be achieved only under the assumption of constant impedance loads since in that way they can be included in the network admittance matrix. The node voltages can then be linearly connected to the internal node voltages and hence to the dynamic state δ_i. Moreover, this results in a reduced admittance matrix that corresponds only to the internal nodes of the generators. The power flows, which are a function of the node voltages, can be now expressed directly by the dynamic states of the system. The resulting model of the two area power system is described by the following set of equations.

$$\begin{aligned}
\dot{\delta}_i &= 2\pi(f_i - f_0),\\
\dot{f}_i &= \frac{f_0}{2H_i S_{B_i}}(P_{m_i} - P_{e_i}(\delta) - \frac{1}{D_i}(f_i - f_0) - \Delta P_{load_i}),\\
\dot{P}_{m,a_k} &= \frac{1}{T_{ch,a_k}}(P^0_{m,a_k} + v_{a_k}\Delta P^{sat}_{p,a_k} + w_{a_k}\Delta P^{sat}_{agc,k} - P_{m,a_k}),\\
\Delta\dot{P}_{agc,k} &= \sum_{j\in A_k} c_{kj}(f_j - f_0) + \sum_{j\in A_k} b_{kj}(P_{m_j} - P_{e_j}(\delta) - \Delta P_{load_j})\\
&\qquad - \frac{1}{T_{N_k}} g_k(\delta, f) - C_{p_k} h_k(\delta, f) - \frac{K_k}{T_{N_k}}(\Delta P_{agc,k} - \Delta P^{sat}_{agc,k}).
\end{aligned}$$

where $i \in G$, $a_k \in A_k$ for $k \in K$. Supperscript *sat* on the AGC output signal $\Delta P_{agc,k}$ and on the primary frequency control signal $\Delta P_{p,a_k}$ highlights the saturation to which the signals are subjected. The primary frequency control is given by $\Delta P_{p,i} = -(f_i - f_0)/S_i$. Based on the reduced admittance matrix, the generator electric power output is given by

$$P_{ei} = \sum_{j=1}^{g} E_{G_i} E_{G_j}(G^{red}_{ij}\cos(\delta_i - \delta_j) + B^{red}_{ij}\sin(\delta_i - \delta_j)).$$

Moreover, $g_k = \sum_{(i,j)\in L^k_{tie}}(P_{ij} - P_{T^0_{12}})$ and $h_k = dg_k/dt$, where the power flow P_{ij}, based on the initial admittance matrix of the system, is given by

$$P_{ij} = |V_{N_i}||V_{N_j}|(G_{ij}\cos(\theta_i - \theta_j) + B_{ij}\sin(\theta_i - \theta_j))$$

All undefined variables are constants, and details on the derivation of the models can be found in [MEVAL12]. The AGC attack is modeled as an additive signal to the AGC signal. For instance, if the attack signal is imposed in Area 1, the mechanical power dynamics of Area 1 will be modified as

$$\dot{P}_{m,a_1} = \frac{1}{T_{ch,a_1}}(P^0_{m,a_1} + v_{a_1}\Delta P^{sat}_{p,a_1} + w_{a_1}\big(\Delta P^{sat}_{agc_1} + f(t)\big) - P_{m,a_1}),$$

The described model above can be compactly written as

$$\begin{cases}\dot{X}(t) = h(X(t)) + B_d d(t) + B_f f(t)\\ Y(t) = CX(t)\end{cases}, \tag{4.29}$$

where $X := \big[\{\delta_i\}_{1:g}, \{f_i\}_{1:g}, \{P_{m,i}\}_{1:g}, \{\Delta P_{agc_i}\}_{1:2}\big]^\intercal \in \mathbb{R}^{3g+2}$ denotes the internal states vector comprising rotor angles δ_i, generators frequencies f_i, generated mechanical powers $P_{m,i}$, and the AGC control signal ΔP_{agc_i} for each area. The external input $d := \big[\{\Delta P_{load_i}\}_{1:g}\big]^\intercal$ represents the unknown load disturbances which will be discussed in the next subsection. The other external input, f, represents the intrusion signal injected to the AGC of the first area, which is the target to detect. We assume that the measurements of all the frequencies and generated mechanical power are available, i.e., $Y = \big[\{f_i\}_{1:g}, \{P_{m,i}\}_{1:g}\big]^\intercal$. The nonlinear function $h(\cdot)$ and the constant matrices B_d, B_f and C can be easily obtained by the mapping between the analytical model and (4.29). To transfer the ODE dynamic expression (4.29) into the DAE (4.5) it suffices to introduce

$$x := \begin{bmatrix} X - X_e \\ d \end{bmatrix}, \qquad z := Y - CX_e$$

$$E(x) := \begin{bmatrix} h(X) - A(X - X_e) \\ 0 \end{bmatrix}, \quad H(p) := \begin{bmatrix} -pI + A & B_d \\ C & 0 \end{bmatrix}, \quad L(p) := \begin{bmatrix} 0 \\ -I \end{bmatrix}, \quad F(p) := \begin{bmatrix} B_f \\ 0 \end{bmatrix},$$

where X_e is the equilibrium of (4.29), i.e., $h(X_e) = 0$, and $A := \frac{\partial h}{\partial X}\big|_{X=X_e}$.

B. Load Deviations and Disturbances

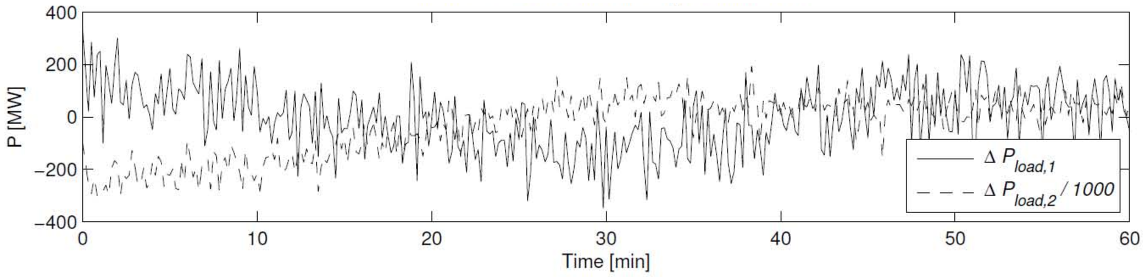

Figure 4.2: Stochastic load fluctuation and prediction error [Anda, p. 59]

In normal operation of power networks there are different sources with different time scales that give rise to power imbalances, e.g., load fluctuation, load forecast errors, and trading on electricity market. The high frequency fluctuation often refers to largely uncorrelated stochastic noises on a second or minute time scale, whereas forecast errors usually stems from the mismatch of predicted and actual consumption on a 15-minute time scale. Figure 4.2 demonstrates two samples of stochastic load fluctuation and forecast error which may appear at two different nodes of the network [Anda, p. 59]. The trading on the electricity market has also impact on the AGC in an hourly framework due to schedule changes at the full hours.

Despite the inherent stochasticity of these disturbances, one may exploit these information to model the behavior of the load deviations. For example, we can describe the disturbances of the first two categories in Figure 4.2 via a family of sinusoidal functions concentrated on different frequency regions, i.e., the high frequency modes correspond to the slower modes concern the prediction mismatch. In regard to the electricity market, the hourly abrupt changes in power imbalances can be captured by a step function with a random power. In this light, the space of load deviations (i.e., the disturbance patterns $\mathcal{D}$ in our FDI setting) is described by

$$\Delta P_{load}(t) := \alpha_0 + \sum_{i=1}^{\eta} \alpha_i \sin(\omega_i t + \phi_i), \qquad t \in [0, T] \tag{4.30}$$

where the parameters $(\alpha_i)_{i=0}^{\eta}$, $(\omega_i)_{i=1}^{\eta}$ $(\phi_i)_{i=1}^{\eta}$, and η are random variables whose distributions induce the probability measure on $\mathcal{D}$.

4.5.2 Diagnosis Filter Design

To design the FDI filter, we set the degree of the filter $d_N = 7$, the denominator $a(p) = (p+2)^{d_N}$, and the finite time horizon $T = 10\,\text{sec}$. Note that the degree of the filter is significantly less than the dimension of the system (4.29), which is 59. This is a general advantage of the residual generator approach in comparison to the observer-based approach where the filter order is effectively the same as the system dynamics. To compute the signature matrix Q_x, we resort to the finite dimensional approximation $Q_{\mathcal{B}}$ in Theorem 4.4.7. Inspired by the class of disturbances in (4.30), we first choose the Fourier basis

$$b_i(t) := \begin{cases} \cos(\frac{i}{2}\omega t) & i : \text{even} \\ \sin(\frac{i+1}{2}\omega t) & i : \text{odd} \end{cases}, \qquad \omega := \frac{2\pi}{T}, \quad i \in \{0, 1, \cdots, k\} \tag{4.31}$$

where the number of the basis is chosen $k = 80$. We should emphasize that there is no restriction on the basis selection as long as Assumptions 4.4.6 are fulfilled; we refer to [MEVAL12, Sec. V.B] for another example with a polynomial basis. Given the basis (4.31), it is easy to see that the differentiation matrix D introduced in (4.19) is

$$D = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \omega & \cdots & 0 & 0 \\ 0 & -\omega & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & & 0 & k\omega \\ 0 & 0 & 0 & \cdots & -k\omega & 0 \end{bmatrix}.$$

We can also compute offline (independent of x) the matrix G in (4.21) with the help of the basis (4.31) and the denominator $a(p)$. To proceed with Q_x of a sample ΔP_{load} we need to run the system dynamic (4.29) with the input $d(\cdot) := \Delta P_{load}$ and compute $x(t) := [X(t)^\intercal, \Delta P_{load}(t)]^\intercal$ where X is the internal states of the system. Given the signal x, we then project the nonlinearity signature $t \mapsto e_x(t) =: E\big(x(t)\big)$ onto the subspace $\mathcal{B}$ (i.e., $\mathbb{T}_{\mathcal{B}}(e_x)$), and finally obtain Q_x from (4.22). In the following simulations, we deploy the YALMIP toolbox [Lof04] to solve the corresponding optimization problems.

4.5.3 Simulation Results

A. Test system

To illustrate the FDI methodology we employed the IEEE 118-bus system. The data of the model are retrieved from a snapshot available at [ref]. It includes 19 generators, 177 lines, 99 load buses and 7 transmission level transformers. Since there were no dynamic data available, typical values provided by [AF02] were used for the simulations. The network was arbitrarily divided into two control areas whose nonlinear frequency model was developed in the preceding

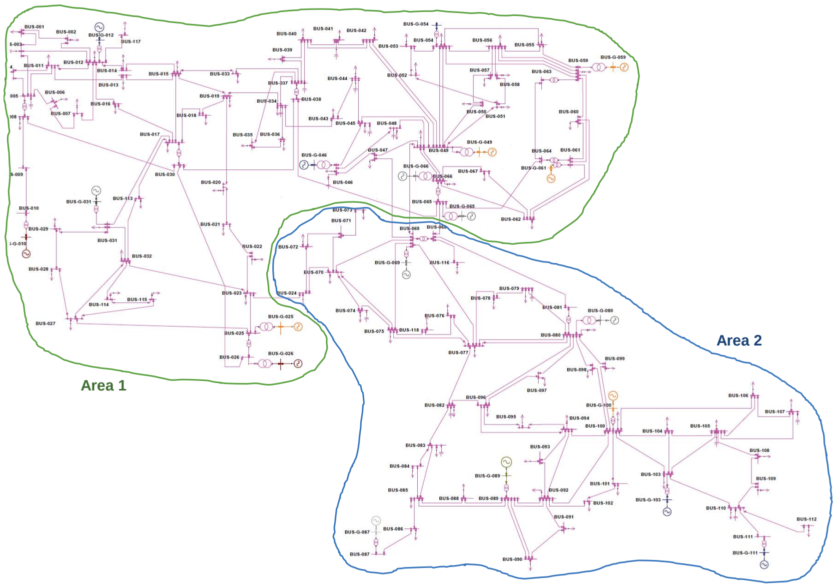

Figure 4.3: IEEE 118-bus system divided into two control areas

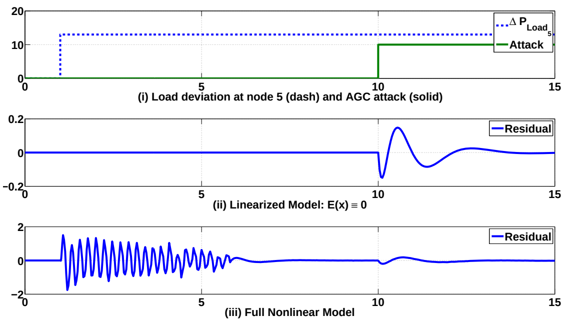

(a) Performance of the filter neglecting the nonlinear term

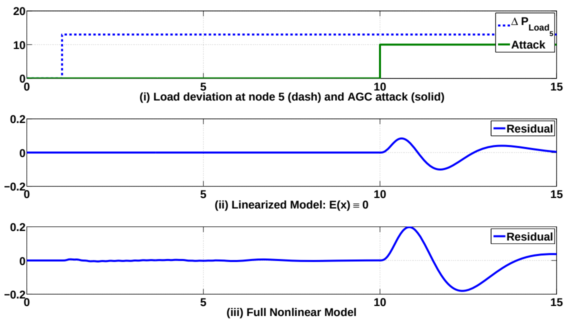

(b) Performance of the filter trained for the step signatures

Figure 4.4: Performance of the FDI filters with step inputs

subsections. Figure 4.3 depicts a single-line diagram of the network and the boundaries of the two controlled areas where the first and second area contain, respectively, 12 and 7 generators.

B. Numerical results

In the first simulation we consider the scenario that an attacker manipulates the AGC signal of the first area at $T_{ack} = 10\,\text{sec}$. We model this intrusion as a step signal equal to 14 MW injected into the AGC in Area 1. To challenge the filter, we also assume that a step load deviation occurs at $T_{load} = 1\,\text{sec}$ at node 5. In the following we present the results of two filters: Figure 4.4(a) shows the filter based on formulation (4.12) in Approach (I), which basically neglects the nonlinear term; Figure 4.4(b) shows the proposed filter in (4.24) based on AP perspective where the payoff function is $J(\alpha) := \alpha^2$.

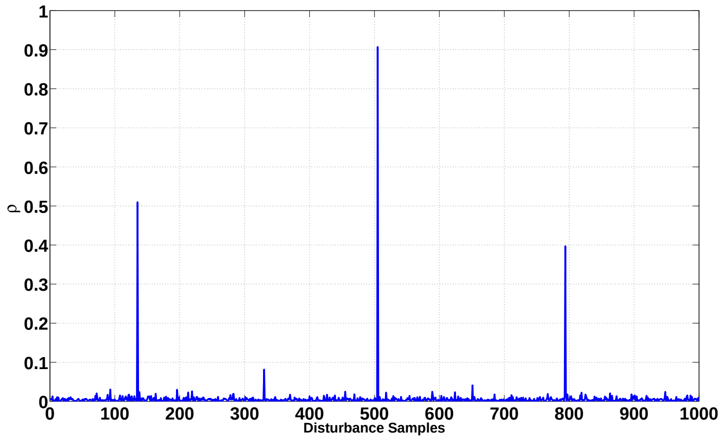

Figure 4.5: The indicator ρ defined in (4.32)

We validate the filters performance with two sets of measurements: first the measurements obtained from the linearized dynamic (i.e. $E(x) \equiv 0$); second the measurements obtained from the full nonlinear model (4.29). As shown in Fig. 4.4(a)(ii) and Fig. 4.4(b)(ii), both filters work perfectly well with linear dynamics measurements. One can even inspect that the first filter seems more sensitive. However, Fig. 4.4(a)(iii) and Fig. 4.4(b)(iii) demonstrate that in the nonlinear setting the first filter fails whereas the robustified filter works effectively similar to the linear setting.

In the second simulation, to evaluate the filter performance in more realistic setup, we robustify the filter to random disturbance patterns, and then verify it with new generated samples. To measure the performance in the presence of the attack, we introduce the following indicator:

$$\rho := \frac{\max\limits_{t \le T_{ack}} \|r(t)\|_\infty}{\max\limits_{t \le T} \|r(t)\|_\infty}, \tag{4.32}$$

where r is the residual (4.11), and T_{ack} is when the attack starts. Observe that $\rho \in [0, 1]$, and the lower ρ the better performance for the filter, e.g., in Fig. 4.4(a)(iii) $\rho = 1$, and in Fig. 4.4(b)(iii) $\rho \approx 0$.

In the training phase, we randomly generate five sinusoidal load deviations as described in (4.30), and excite the dynamics for $T = 10$ *sec* in the presence of each of the load deviations individually. Hence, in total we have $n = 19 \times 5 = 95$ disturbance signatures. Then, we compute the filter coefficients by virtue of $\widetilde{\mathrm{AP}}$ in (4.24) and these 95 samples. In the operation phase, we generate two new disturbance patterns with the same distribution as in the training phase and run the system in the presence of both load deviations simultaneously at two random nodes for the horizon $T = 120\,\mathrm{sec}$. Meanwhile, we inject an attack signal at $T_{ack} = 110\,\mathrm{sec}$ in the AGC, and compute the indicator ρ in (4.32) for each of the scenarios. Figure 4.5 demonstrates the result of this simulation for 1000 experiments.

4.6 Summary and Outlook

In this chapter, we proposed a novel perspective toward the FDI filter design, which is tackled via an optimization-based methodology along with probabilistic performance guarantees. Thanks to the convex formulation, the methodology is applicable to high dimensional nonlinear systems in which some statistical information of exogenous disturbances are available. Motivated by our earlier works, we deployed the proposed technique to design a diagnosis filter to detect the AGC malfunction in two-area power network. The simulation results validated the filter performance, particularly when the disturbance patterns are different from training to the operation phase.

The central focus of the work in this chapter is to robustify the filter to certain signatures of dynamic nonlinearities in the presence of given disturbance patterns. As a next step, motivated by applications that the disruptive attack may follow certain patterns, a natural question is whether the filter can be trained to these attack patterns. From the technical standpoint, this problem in principle may be different from the robustification process since the former may involve maximization of the residual norm as opposed to the minimization for the robustification discussed in this chapter. Therefore, this problem offers a challenge to reconcile the disturbance rejection and the fault sensitivity objectives.

The methodology studied in this chapter is applicable to both discrete and continuous-time dynamics and measurements. In reality, however, we often have different time-setting in different parts, i.e., we only have discrete-time measurements while the system dynamics follows a continuous-time behavior. We believe this setup introduces new challenges to the field. We recently reported heuristic attempts toward this objective in [ETMEL13], though there is still a need to address this problem in a rigorous and systematic framework.

4.7 Appendix

4.7.1 Proofs of Section 4.4.2

Let us start with a preliminary required for the main proof of this section.

Lemma 4.7.1. *Let $N(p) := \sum_{i=0}^{d_N} N_i p^i$ be an $\mathbb{R}^{n_r}$ row polynomial vector with degree d_N, and $a(p)$ be a stable polynomial with the degree at least d_N. Let $\bar{N} := [N_0 \ N_1 \ \cdots \ N_{d_N}]$ be the collection of the coefficients of $N(p)$. Then,*

$$\left\|a^{-1}N\right\|_{\mathcal{H}_\infty} \le \widetilde{C}\|\bar{N}\|_\infty, \qquad \widetilde{C} := \sqrt{n_r(d_N+1)}\ \|a^{-1}\|_{\mathcal{H}_\infty}.$$

Proof. Let $b(p) := \sum_{i=0}^{d_N} b_i p^i$ be a polynomial scaler function. By $\mathcal{H}_\infty$-norm definition we have

$$\left\|a^{-1}b\right\|^2_{\mathcal{H}_\infty} = \sup_{\omega\in(-\infty,\infty)} \left|\frac{b(j\omega)}{a(j\omega)}\right|^2 \le \sup_{\omega\in[0,\infty)} \frac{\sum_{i=0}^{d_N}|b_i|^2\omega^{2i}}{|a(j\omega)|^2}. \tag{4.33}$$

Let $\bar{b} := \begin{bmatrix} b_0 & b_1 & \cdots & b_{d_N} \end{bmatrix}$. It is then straightforward to inspect that

$$\sum_{i=0}^{d_N} |b_i|^2 \omega^{2i} \le \begin{cases} (d_N+1)\|\bar{b}\|_\infty^2 & \text{if} \quad \omega \in [0,1] \\ (d_N+1)\|\bar{b}\|_\infty^2 \omega^{2d_N} & \text{if} \quad \omega \in (1,\infty) \end{cases} \tag{4.34}$$

Therefore, (4.33) together with (4.34) yields to

$$\left\|a^{-1}b\right\|_{\mathcal{H}_\infty}^2 \le (d_N+1)\left\|a^{-1}\right\|_{\mathcal{H}_\infty}^2 \|\bar{b}\|_\infty^2.$$

Now, taking the dimension of the vector $N(p)$ into consideration, we conclude the desired assertion. □

Proof of Lemma 4.4.5. Let $\ell \ge d_N$ be the degree of the scalar polynomial $a(p)$. Then, taking advantage of the state-space representation of the matrix transfer function $a^{-1}(p)N(p)$, in particular the observable canonical form [ZD97, Sec. 3.5], we have

$$r_x(t) = \int_0^t C\,\mathrm{e}^{-A(t-\tau)}\,Be_x(\tau)\mathrm{d}\tau + De_x(t),$$

where $C \in \mathbb{R}^{1\times\ell}$ is a constant vector, $A \in \mathbb{R}^{\ell\times\ell}$ is the state matrix depending only on $a(p)$, and $B \in \mathbb{R}^{\ell\times n_r}$ and $D \in \mathbb{R}^{1\times n_r}$ are matrices that depend linearly on all the coefficients of the numerator $\bar{N} \in \mathbb{R}^{n_r(d_N+1)}$. Therefore, it can be readily deduced that (4.14a) holds for some function $\psi_x \in \mathcal{W}_T^{n_r(d_N+1)}$. In regard to (4.14a) and the definition (4.13), we have

$$\|\bar{N}\psi_x\|_{\mathcal{L}_2} = \|r_x\|_{\mathcal{L}_2} = \left\|a^{-1}(p)N(p)e_x\right\|_{\mathcal{L}_2} \le \left\|a^{-1}N\right\|_{\mathcal{H}_\infty}\|e_x\|_{\mathcal{L}_2} \le \widetilde{C}\|\bar{N}\|_\infty\|e_x\|_{\mathcal{L}_2}, \tag{4.35}$$

where the first inequality follows from the classical result that the $\mathcal{L}_2$-gain of a matrix transfer function is the $\mathcal{H}_\infty$-norm of the matrix [ZD97, Thm. 4.3, p. 51], and the second inequality follows from Lemma 4.7.1. Since (4.35) holds for every $\bar{N} \in \mathbb{R}^{n_r(d_N+1)}$, then

$$\|\psi_x\|_{\mathcal{L}_2} \le \sqrt{n_r(d_N+1)}\;\widetilde{C}\|e_x\|_{\mathcal{L}_2},$$

which implies (4.14b). □

Proof of Theorem 4.4.7. Observe that by virtue of the triangle inequality and linearity of the projection mapping we have

$$\Big|\|r_x\|_{\mathcal{L}_2} - \left\|a^{-1}(p)N(p)\mathbb{T}_{\mathcal{B}}(e_x)\right\|_{\mathcal{L}_2}\Big| \le \left\|a^{-1}(p)N(p)\big(e_x - \mathbb{T}_{\mathcal{B}}(e_x)\big)\right\|_{\mathcal{L}_2} \le \widetilde{C}\|\bar{N}\|_\infty\delta,$$

where the second inequality follows in the same spirit as (4.35) and $\|e_x - \mathbb{T}_{\mathcal{B}}(e_x)\|_{\mathcal{L}_2} \le \delta$. Note that by definitions of Q_x and $Q_{\mathcal{B}}$ in (4.15) and (4.22), respectively, we have

$$\begin{aligned} \left|\bar{N}(Q_x - Q_{\mathcal{B}})\bar{N}^\intercal\right| &= \Big|\|r_x\|_{\mathcal{L}_2}^2 - \left\|a^{-1}(p)N(p)\mathbb{T}_{\mathcal{B}}(e_x)\right\|_{\mathcal{L}_2}^2\Big| \le \widetilde{C}\|\bar{N}\|_\infty\delta\big(\widetilde{C}\|\bar{N}\|_\infty\delta + 2\|r_x\|_{\mathcal{L}_2}\big) \\ &\le \widetilde{C}^2\|\bar{N}\|_\infty^2\delta\big(\delta + 2\|e_x\|_{\mathcal{L}_2}\big) \le C\|a^{-1}\|_{\mathcal{H}_\infty}\|\bar{N}\|_2^2\delta\big(1 + 2\|e_x\|_{\mathcal{L}_2}\big) \end{aligned}$$

where the inequality of the first line stems from the simple inequality $|\alpha^2 - \beta^2| \le |\alpha - \beta|\big(2|\alpha| + |\alpha - \beta|\big)$, and C is the constant as in (4.14b). □

4.7.2 Proofs of Section 4.4.3

To prove Theorem 4.4.10 we need a preparatory result addressing the continuity of the mapping ϕ in (4.26).

Lemma 4.7.2. *Consider the function ϕ as defined in (4.26). Then, there exists a constant $L > 0$ such that for any $\bar{N}_1, \bar{N}_2 \in \mathcal{N}$ and $x_1, x_2 \in \mathcal{W}_T^{n_x}$ where $\|x_i\|_{\mathcal{L}_2} \le M$, we have*

$$\big|\phi(\bar{N}_1, x_1) - \phi(\bar{N}_2, x_2)\big| \le L\big(\big\|\bar{N}_1 - \bar{N}_2\big\|_\infty + \|x_1 - x_2\|_{\mathcal{L}_2}\big).$$

Proof. Let L_E be the Lipschitz continuity constant of the mapping $E : \mathbb{R}^{n_x} \to \mathbb{R}^{n_r}$ in (4.13). We modify the notation of r_x in (4.13) with a new argument as $r_x[\bar{N}]$, in which $\bar{N}$ represents the filter coefficients. Then, with the aid of (4.35), we have

$$\sup_{\|x\|_{\mathcal{L}_2} \le M} \sup_{\bar{N} \in \mathcal{N}} \|r_x[\bar{N}]\|_{\mathcal{L}_2} \le \sup_{\|x\|_{\mathcal{L}_2} \le M} \sup_{\bar{N} \in \mathcal{N}} \widetilde{C} L_E \|\bar{N}\|_\infty \|x\|_{\mathcal{L}_2} \le \widetilde{M}, \qquad \widetilde{M} \coloneqq \widetilde{C} L_E M,$$

where the constant $\widetilde{C}$ is introduced in Lemma 4.7.1. As the payoff function J is convex, it is then Lipschitz continuous over the compact set $[0, \widetilde{M}]$ [Ber09, Prop. 5.4.2, p. 185]; we denote this Lipschitz constant by L_J. Then for any $\bar{N}_i \in \mathcal{N}$ and $\|x_i\|_{\mathcal{L}_2} \le M$, $i \in \{1, 2\}$, we have,

$$\begin{aligned}
\big|\phi(\bar{N}_1, x_1) - \phi(\bar{N}_2, x_2)\big| &\le L_J \big| \big\|r_{x_1}[N_1]\big\|_{\mathcal{L}_2} - \big\|r_{x_2}[N_2]\big\|_{\mathcal{L}_2} \big| \\
&\le L_J \Big(\big\|r_{x_1}[N_1] - r_{x_1}[N_2]\big\|_{\mathcal{L}_2} + \big\|r_{x_1}[N_2] - r_{x_2}[N_2]\big\|_{\mathcal{L}_2} \Big) \\
&\le L_J \Big(\widetilde{C} \|e_{x_1}\|_{\mathcal{L}_2} \|N_1 - N_2\|_\infty + \widetilde{C} \|e_{x_1} - e_{x_2}\|_{\mathcal{L}_2} \|N_2\|_\infty \Big) \qquad (4.36) \\
&\le L_J \widetilde{C} L_E \big(M \|N_1 - N_2\|_\infty + \|x_1 - x_2\|_{\mathcal{L}_2} \big).
\end{aligned}$$

where (4.36) follows from (4.35) and the fact that the mapping $(\bar{N}, e_x) \mapsto r_x[\bar{N}]$ is bilinear. □

Proof of Theorem 4.4.10. By virtue of Lemma 4.7.2, one can infer that for every $\bar{N} \in \mathcal{N}$ the mapping $x \mapsto \phi(\bar{N}, x)$ is continuous, and hence measurable. Therefore, $\phi(\bar{N}, x)$ can be viewed as a random variable for each $\bar{N} \in \mathcal{N}$, which yields to the first assertion, see [Bil99, Chap. 2, p. 84] for more details.

By uniform (almost sure) boundedness and again Lemma 4.7.2, the mapping $\bar{N} \mapsto \phi(\bar{N}, x)$ is uniformly Lipschitz continuous (except on a negligible set), and consequently first moment continuous in the sense of [Han12, Def. 2.5]. We then reach (i) by invoking [Han12, Thm. 2.1].

For assertion (ii), note that the compact set $\mathcal{N}$ is finite dimensional, and thus admits a logarithmic ε-capacity in the sense of [Dud99, Sec. 1.2, p. 11]. Therefore, the condition [Dud99, (6.3.4), p. 209] is satisfied. Since the other requirements of [Dud99, Thm. 6.3.3, p. 208] are readily fulfilled by the uniform boundedness assumption and Lemma 4.7.2, we arrive at the desired UCLT assertion in (ii). □

Proof of Theorem 4.4.12. The measurability of $\mathcal{E}$ is a straightforward consequence of the measurability of $[\bar{N}_n^\star, \gamma_n^\star]$ and Fubini's Theorem [Bil95, Thm. 18.3, p. 234]. For notational simplicity, we introduce the following notation. Let $\ell \coloneqq n_r(d_N + 1) + 1$ and define the function

$f : \mathbb{R}^\ell \times \mathcal{W}_T^{n_x} \to \mathbb{R}$

$$f(\theta, x) := \bar{N} Q_x \bar{N}^\intercal - \gamma, \qquad \theta := [\bar{N}, \gamma]^\intercal \in \mathbb{R}^\ell,$$

where Q_x is the nonlinearity signature matrix of x as defined in (4.15), and θ is the augmented vector collecting all the decision variables. Consider the convex sets $\Theta_j \subset \mathbb{R}^\ell$

$$\Theta_j := \Big\{ \theta = [\bar{N}, \gamma]^\intercal \ \big|\ \bar{N}\bar{H} = 0,\ \bar{N}\bar{F}v_j \geq 1 \Big\}, \qquad v_j := \big[0, \cdots, \overset{\downarrow\, j^{\text{th}}}{1}, \cdots, 0\big]^\intercal,$$

where the size of v_j is $m := n_f(d_F + d_N + 1)$. We then express the program CP in (4.3) and its random counterpart $\widetilde{\mathrm{CP}}_1$ in (4.25a) as follows:

$$\mathrm{CP} : \begin{cases} \min\limits_{\theta \in \bigcup_{j=1}^m \Theta_j} & c^\intercal \theta \\ \text{s.t.} & \mathbb{P}\big(f(\theta, x) \leq 0\big) \geq 1 - \varepsilon \end{cases} \qquad \widetilde{\mathrm{CP}}_1 : \begin{cases} \min\limits_{\theta \in \bigcup_{j=1}^m \Theta_j} & c^\intercal \theta \\ \text{s.t.} & \max\limits_{i \leq n} f(\theta, x_i) \leq 0, \end{cases}$$

where c is the constant vector with 0 elements except the last which is 1. It is straightforward to observe that the optimal threshold $\gamma_n^\star$ of the two-stage program $\widetilde{\mathrm{CP}}$ in (4.25) is the same as the optimal threshold obtained in the first stage $\widetilde{\mathrm{CP}}_1$. Thus, it suffices to show the desired assertion considering only the first stage. Let $\theta_n^\star := [\bar{N}_n^\star, \gamma_n^\star]$ denote the optimizer of $\widetilde{\mathrm{CP}}_1$. Now, consider m sub-programs denoted by $\mathrm{CP}\,(j)$ and $\widetilde{\mathrm{CP}}\,(j)$ for $j \in \{1, \cdots, m\}$:

$$\mathrm{CP}\,(j) : \begin{cases} \min\limits_{\theta \in \Theta_j} & c^\intercal \theta \\ \text{s.t.} & \mathbb{P}\big(f(\theta, x) \leq 0\big) \geq 1 - \varepsilon \end{cases} \qquad \widetilde{\mathrm{CP}}_1 : \begin{cases} \min\limits_{\theta \in \Theta_j} & c^\intercal \theta \\ \text{s.t.} & \max\limits_{i \leq n} f(\theta, x_i) \leq 0, \end{cases}$$

Let us denote the optimal solution of $\widetilde{\mathrm{CP}}_1\,(j)$ by $\theta_{n,j}^*$. Note that for all j, the set Θ_j is convex and the corresponding random program $\widetilde{\mathrm{CP}}_1\,(j)$ is feasible if $\Theta_j \neq \emptyset$. Therefore, we can readily employ the existing results of the random convex problems. Namely, by [CG08, Thm. 1] we have

$$\mathbb{P}^n\big(\mathcal{E}(\theta_{n,j}^*)\big) < \sum_{i=0}^{\ell-1} \binom{n}{i} \varepsilon^i (1-\varepsilon)^{n-i}, \qquad \forall j \in \{1, \cdots, m\}$$

where $\mathcal{E}$ is introduced in (4.28). Furthermore, it is not hard to inspect that $\theta_n^* \in \big(\theta_{n,j}^*\big)_{j=1}^m$. Thus, $\mathcal{E}(\theta_n^*) \subseteq \bigcup_{j=1}^m \mathcal{E}(\theta_{n,j}^*)$ which yields

$$\mathbb{P}^n\big(\mathcal{E}(\theta_n^*)\big) \leq \mathbb{P}^n\Big(\bigcup_{j=1}^m \mathcal{E}(\theta_{n,j}^*)\Big) \leq \sum_{j=1}^m \mathbb{P}^n\big(\mathcal{E}(\theta_{n,j}^*)\big) < m \sum_{i=0}^{\ell-1} \binom{n}{i} \varepsilon^i (1-\varepsilon)^{n-i}.$$

Now, considering β as an upper bound, the desired assertion can be obtained by similar calculation as in [Cal09] to make the above inequality explicit for n in terms of ε and β. □

CHAPTER 5

Performance Bound for Random Programs

In Chapter 4, we proposed the chance constrained perspective to relax the robust formulation of the original FDI problem. Employing randomized algorithms, we provided a (random) solution to the relaxed problem whose performance is guaranteed with high probability. Along this way an alarm threshold is introduced which may be violated with probability at most an $\varepsilon \in (0, 1)$ (a design parameter). This violation probability is allowed from the relaxed perspective and known in literature as the false alarm rate. In this chapter we aim to study the behavior of the false alarm rate as the threshold level obtained via the random programs changes.

In more general setting, we consider the Scenario Convex Program (SCP) for two classes of optimization problems that are not tractable in general: Robust Convex Programs (RCPs) and Chance-Constrained Programs (CCPs). We establish a probabilistic bridge from the optimal value of SCP to the optimal values of RCP and CCP in which the uncertainty takes values in a general, possibly infinite dimensional, metric space. We then extend our results to a certain class of non-convex problems that includes, for example, binary decision variables. In the process, we also settle a measurability issue for a general class of scenario programs, which to date has been addressed by an assumption. Finally, we demonstrate the applicability of our results on a benchmark problem and a problem in fault detection and isolation.

5.1 Introduction

Optimization problems under uncertainty have considerable applications in disciplines ranging from mathematical finance to control engineering. For example most control systems involve some level of uncertainty; the aim of a robust control design is to provide a guaranteed level of performance for all admissible values of the uncertain parameters. In the convex case, two well-known approaches for dealing with such uncertain programs are robust convex programs (RCPs) and chance-constrained programs (CCPs). RCPs consider constraint satisfaction for all, possibly infinitely many, realizations of the uncertainty. While it is known that certain classes of RCPs can be solved as effectively as their non-robust counterparts [BS06] in other cases RCPs can be intractable [BtN98, BtN99, GOL98, BtNR01]. For example, the class of parametric linear matrix inequalities, which occur in many control problems, is NP-hard [BGFB94, Gah96]. CCPs, on the other hand, allow constraint violation with a low probability. The resulting optimization problem, however, is in general non-convex [Pré95, SDR09].

Computationally tractable approximations to the aforesaid optimization problems can be obtained through the scenario convex programs (SCPs) in which only finitely many uncertainty samples are considered. A natural question in this case is how many samples would be "enough" to provide a good solution. To answer this question, one may view the problem from two perspectives: feasibility and objective performance. The literature mainly focuses on the first perspective. In this direction, the authors in [CC05, CC06] initialized a feasibility theory for CCP refined subsequently in [CG08, Cal10]. They established an explicit probabilistic lower bound for the sample size to guarantee the feasibility of the SCP solutions from a chance-constrained perspective. By contrast, the issue of performance bounds for both RCP and CCP via SCP has not been settled up to now. [CG11] provides a novel perspective in this direction that leads to optimal performance bounds for CCPs. However, it involves the problem of optimal constraint removal, which in general is computationally intractable.

The first contribution of this chapter is to address the SCP performance issue from the objective viewpoint. The key element of our analysis relies on the concept of the worst-case violation inspired by the recent work [KT12]. The authors of [KT12] derived an upper bound of the worst-case violation for the SCPs where the uncertainty takes values in a finite dimensional Euclidean space. This result leads to a performance bound for a particular class of RCPs where the uncertainly appears in the objective function, e.g., min-max optimization problems. Motivated by different applications such as control problems with saturation constraints [CGP09], fault detection and isolation in dynamical systems [MEL13], and approximate dynamic programming [DPR13], in this chapter we first extend this result to infinite dimensional uncertainty spaces. In the sequel, we establish a theoretical bridge from the optimal values of SCP to the optimal values of both RCP and CCP. Along this direction, under mild assumptions on the constraint function (measurability with respect to the uncertainty and lower semicontinuity with respect to the decision variables), we shall also rigorously settle a measurability issue of the SCP optimizer, which to date has been addressed in the literature by an assumption, e.g. [CC06, CG08]. Our second contribution is to extend these results to a class of non-convex programs that, in particular, allows for binary decision variables. In the context of mixed integer programs, the recent work [CLF12] investigates the feasibility perspective of CCPs, which leads to a bound of the required number of scenarios with exponential growth rate in the number of integer variables, whereas our proposed bound scales linearly.

The layout of this chapter is as follows: In Section 5.2 we formally introduce the optimization problems that will be addressed. Our results on probabilistic objective performance for both RCPs and CCPs based on SCPs are reported in Section 5.3. In Section 5.4 we extend our results to a class of non-convex programs, including mixed-integer programs with binary variables. To illustrate the proposed methodology, in Section 5.5 the theoretical results are applied to two examples: a benchmark problem whose solution can be computed explicitly, and a fault detection and isolation study with an application to the security of power networks. We conclude in Section 5.6 with a summary of our work and comment on possible subjects of further research. For better readability, some of the technical proofs and details are given in the appendices.

Notation

Let $\mathbb{R}_+$ denote the non-negative real numbers. Given a metric space $\mathcal{D}$, its Borel σ-algebra is denoted by $\mathfrak{B}(\mathcal{D})$. Throughout this chapter, measurability is always referred to Borel measurability. Given a probability space $(\mathcal{D}, \mathfrak{B}(\mathcal{D}), \mathbb{P})$, we denote the N-Cartesian product set of $\mathcal{D}$ by $\mathcal{D}^N$ and the respective product measure by $\mathbb{P}^N$. An open ball in $\mathcal{D}$ with radius r and center v is denoted by $\mathbb{B}_r(v) := \{d \in \mathcal{D} : \|d - v\| < r\}$. The symbol $\models$ refers to the feasibility satisfaction, i.e., $x \models$ RCP implies that x is a feasible solution for the program RCP. Similarly, $x \not\models$ RCP implies that x is not a feasible solution for the optimization problem RCP.

5.2 Problem Statement

Let $\mathbb{X} \subset \mathbb{R}^n$ be a compact convex set and $c \in \mathbb{R}^n$ a constant vector. Let $(\mathcal{D}, \mathfrak{B}(\mathcal{D}), \mathbb{P})$ be a probability space where $\mathcal{D}$ is a metric space with the respective Borel σ-algebra $\mathfrak{B}(\mathcal{D})$. Consider the measurable function $f : \mathbb{X} \times \mathcal{D} \to \mathbb{R}$, which is convex in the first argument for each $d \in \mathcal{D}$, and bounded in the second argument for each $x \in \mathbb{X}$. We then consider the following optimization problems:

$$\text{RCP}: \begin{cases} \min\limits_{x} & c^\intercal x \\ \text{s.t.} & f(x,d) \leq 0, \quad \forall d \in \mathcal{D} \\ & x \in \mathbb{X} \end{cases}, \qquad \text{CCP}_\varepsilon: \begin{cases} \min\limits_{x} & c^\intercal x \\ \text{s.t.} & \mathbb{P}[f(x,d) \leq 0] \geq 1 - \varepsilon \\ & x \in \mathbb{X} \end{cases}, \tag{5.1}$$

where $\varepsilon \in [0,1]$ is the constraint violation level for the chance-constrained program. We denote the optimal value of the program RCP (resp. CCP_ε) by $J^\star_{\text{RCP}}$ (resp. $J^\star_{\text{CCP}_\varepsilon}$). Suppose $(d_i)_{i=1}^N$ are N independent and identically distributed (i.i.d.) samples drawn according to the probability measure $\mathbb{P}$. The centerpiece of this study is the scenario program

$$\text{SCP}: \begin{cases} \min\limits_{x} & c^\intercal x \\ \text{s.t.} & f(x,d_i) \leq 0, \quad \forall i \in \{1, \cdots, N\} \\ & x \in \mathbb{X} \end{cases}, \tag{5.2}$$

where the optimal solution and optimal value of SCP are denoted, respectively, by $x^\star_N$ and $J^\star_N$. Notice that SCP is naturally random as it depends on the random samples $(d_i)_{i=1}^N$.

We assume throughout our subsequent analysis that the following measurability assumption holds, though we shall show in Subsection 5.3.3 how one may rigorously address this issue without any assumption for a large class of optimization programs (not necessarily convex).

Assumption 5.2.1. *The* SCP *optimizer generates a Borel measurable mapping from* $(\mathcal{D}^N, \mathfrak{B}(\mathcal{D}^N))$ *to* $(\mathbb{X}, \mathfrak{B}(\mathbb{X}))$ *that associates each* $(d_i)_{i=1}^N$ *with a unique* $x^\star_N$*.*

The optimization program SCP in (5.2) is convex and hence tractable even for cases where the problems (5.1) are NP-hard. Motivated by this, a natural question is whether there exist theoretical links from SCP to RCP and CCP_ε. As mentioned in the introduction, this question can be addressed from two different perspectives: feasibility and objective performance. From the feasibility perspective, we recall the explicit bound of [CG08] which measures the finite sample behavior of SCP:

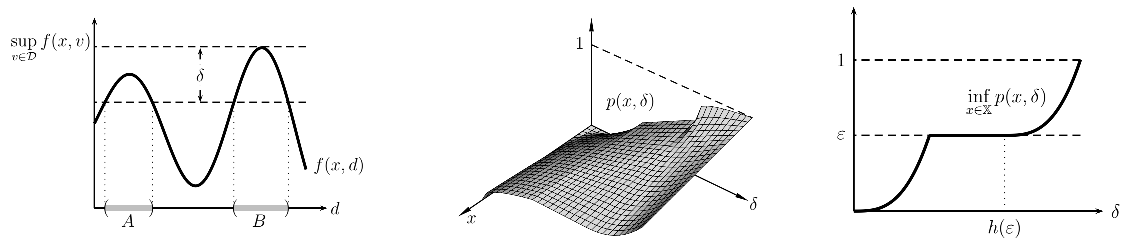

(a) $p(x,\delta) = \mathbb{P}[A \cup B]$ (b) Tail probability of the worst-case violation (c) Uniform level set bound

Figure 5.1: Pictorial representation of Definition 5.3.1

Theorem 5.2.2 (CCP$_\varepsilon$ Feasibility)**.** *Let $\beta \in [0,1]$ and $N \geq N(\varepsilon,\beta)$ where*

$$N(\varepsilon,\beta) := \min\Big\{ N \in \mathbb{N} \ \Big| \ \sum_{i=0}^{n-1} \binom{N}{i} \varepsilon^i (1-\varepsilon)^{N-i} \leq \beta \Big\}. \tag{5.3}$$

Then, the optimizer of SCP *is a feasible solution of* CCP$_\varepsilon$ *with probability at least* $1-\beta$*.*

With "$\models$" notation, the assertion of Theorem 5.2.2 is alternatively stated by $\mathbb{P}^N\big[x^\star_N \models \text{CCP}_\varepsilon\big] \geq 1-\beta$, where $\mathbb{P}^N$ stands for the N-fold product probability measure.[1]

To the best of our knowledge, there is no clear connection between the feasibility of RCP and the solution of SCP. Furthermore, in Subsection 5.3.2 we provide an example to challenge the possibility of such a connection. The focus of our study is on the second perspective to seek a (probabilistic) bound for the optimal values $J^\star_{\text{RCP}}$ and $J^\star_{\text{CCP}_\varepsilon}$ in terms of $J^\star_N$.

5.3 Probabilistic Objective Performance

5.3.1 Confidence interval for the objective functions

The following definition inspired by the recent work [KT12] is the key object for our analysis.

Definition 5.3.1. *The* tail probability of the worst-case violation *is the function* $p : \mathbb{X} \times \mathbb{R}_+ \to [0,1]$ *defined as*

$$p(x,\delta) := \mathbb{P}\Big[\sup_{v\in\mathcal{D}} f(x,v) - \delta < f(x,d)\Big].$$

We call $h : [0,1] \to \mathbb{R}_+$ *a* uniform level-set bound *(ULB) of* p *if for all* $\varepsilon \in [0,1]$

$$h(\varepsilon) \geq \sup\Big\{\delta \in \mathbb{R}_+ \ \Big| \ \inf_{x\in\mathbb{X}} p(x,\delta) \leq \varepsilon\Big\}.$$

[1] Note that $\mathbb{P}$ is the probability measure on $\mathfrak{B}(\mathcal{D})$; for simplicity we slightly abuse the notation, and will be doing so hereinafter. Strictly speaking, one has to define a new probability measure, say $\mathbb{Q}$, which is the induced measure on $\mathfrak{B}(\mathbb{X})$ via the mapping introduced in Assumption 5.2.1.

A pictorial representation of Definition 5.3.1 is given in Figure 5.1. Note that from a statistical perspective the ULB function may be alternatively viewed as an upper bound for the *quantile* function of the $\mathbb{R}$-valued random variable $d \mapsto \big(\sup_{v\in\mathcal{D}} f(x,v) - f(x,d)\big)$ uniformly in decision variable $x \in \mathbb{X}$ (cf. [Sha03, Section 5.2]). Proposition 5.3.8 at the end of this subsection provides sufficient conditions under which a candidate ULB can be constructed. If the uncertainty set $\mathcal{D}$ is a specific compact subset of a Euclidean space, namely a norm-constrained or more generally a star-shaped set, the authors in [KT12] provide a constructive approach to obtain an admissible ULB.

Consider the relaxed version of the program RCP for $\gamma > 0$:

$$\mathrm{RCP}_\gamma : \begin{cases} \min\limits_{x} & c^\intercal x \\ \text{s.t.} & f(x,d) \le \gamma, \quad \forall d \in \mathcal{D} \\ & x \in \mathbb{X} \end{cases}, \tag{5.4}$$

with the optimal value $J^\star_{\mathrm{RCP}_\gamma}$.

Lemma 5.3.2. *Let $h : [0,1] \to \mathbb{R}_+$ be a ULB. Then,*

$$x \models \mathrm{CCP}_\varepsilon \quad \Longrightarrow \quad x \models \mathrm{RCP}_{h(\varepsilon)}$$

that is, the feasible set of the program CCP_ε with constraint violation level ε is a subset of the feasible set of the relaxed program RCP_γ with $\gamma := h(\varepsilon)$.

Proof. See Appendix 5.7.1. □

Assumption 5.3.3 (Slater Point)**.** *There exists an $x_0 \in \mathbb{X}$ such that $\sup\limits_{d\in\mathcal{D}} f(x_0,d) < 0$.*

Under Assumption 5.3.3, we define the constant

$$L_{\mathbf{SP}} := \frac{\min_{x\in\mathbb{X}} c^\intercal x - c^\intercal x_0}{\sup_{d\in\mathcal{D}} f(x_0,d)}. \tag{5.5}$$

The following lemma is a classical result in perturbation theory of convex programs, which is a significant ingredient for the first result of this chapter.

Lemma 5.3.4. *Consider the relaxed program RCP_γ and its optimal value $J^\star_{\mathrm{RCP}_\gamma}$ as introduced in (5.4). Under Assumption 5.3.3, the mapping $\mathbb{R}_+ \ni \gamma \mapsto J^\star_{\mathrm{RCP}_\gamma} \in \mathbb{R}$ is Lipschitz continuous with constant bounded by $L_{\mathbf{SP}}$ in (5.5), i.e., for all $\gamma_2 \ge \gamma_1 \ge 0$ we have*

$$0 \le J^\star_{\mathrm{RCP}_{\gamma_1}} - J^\star_{\mathrm{RCP}_{\gamma_2}} \le L_{\mathbf{SP}}(\gamma_2 - \gamma_1).$$

Proof. See Appendix 5.7.1. □

Assumption 5.3.3 requires the existence of a strictly feasible solution x_0 which, in general, may not exist. However, in applications where a "risk-free" decision is available such an assumption is not really restrictive; the portfolio selection problem is an example of this kind [PRC12]. In addition, for the class of min-max problems, as a particular case of the program RCP, it is not difficult to see that Assumption 5.3.3 holds; see the following remark for more details and Section 5.5.2 for an application to the problem of fault detection and isolation.

Remark 5.3.5 ($L_{\mathbf{SP}}$ for Min-Max Problems)**.** *In min-max problems, one may inspect that there always exists a Slater point (in the sense of Assumption 5.3.3) with the corresponding constant $L_{\mathbf{SP}}$ arbitrarily close to 1. In fact, it is straightforward to observe that for min-max problems $J^\star_{\mathrm{RCP}_\gamma} = J^\star_{\mathrm{RCP}} - \gamma$, which readily implies that the Lipschitz constant of Lemma 5.3.4 is 1.*

The following results are the main contributions of the first part of the chapter.

Theorem 5.3.6 (RCP Confidence Interval)**.** *Consider the programs* RCP *and* SCP *in* (5.1) *and* (5.2) *with the associated optimal values $J^\star_{\mathrm{RCP}}$ and $J^\star_N$, respectively. Suppose Assumption 5.3.3 holds and $L_{\mathbf{SP}}$ is the constant in* (5.5)*. Given a ULB h and ε, β in $[0,1]$, for all $N \geq N(\varepsilon,\beta)$ as defined in* (5.3)*, we have*

$$\mathbb{P}^N\Big[J^\star_{\mathrm{RCP}} - J^\star_N \in \big[0, I(\varepsilon)\big]\Big] \geq 1-\beta, \tag{5.6}$$

where

$$I(\varepsilon) \coloneqq \min\Big\{L_{\mathbf{SP}} h(\varepsilon),\ \max_{x\in\mathbb{X}} c^\intercal x - \min_{x\in\mathbb{X}} c^\intercal x\Big\}. \tag{5.7}$$

Proof. Due to the definition of the optimization problems RCP and SCP, the second term of the confidence interval (5.7) is a trivial bound. It then suffices to establish the bound for the first term of (5.7). By Theorem 5.2.2, we know $\mathbb{P}^N\big[x^\star_N \models \mathrm{CCP}_\varepsilon\big] \geq 1-\beta$ that in view of Lemma 5.3.2 implies

$$\mathbb{P}^N\big[x^\star_N \models \mathrm{RCP}_{h(\varepsilon)}\big] \geq 1-\beta \implies \mathbb{P}^N\big[J^\star_{\mathrm{RCP}_{h(\varepsilon)}} \leq J^\star_N\big] \geq 1-\beta,$$

where h is the ULB, and $J^\star_{\mathrm{RCP}_{h(\varepsilon)}}$ is the optimal value of the relaxed robust program (5.4) with $\gamma \coloneqq h(\varepsilon)$. By virtue of Lemma 5.3.4, we have $J^\star_{\mathrm{RCP}} \leq J^\star_{\mathrm{RCP}_{h(\varepsilon)}} + L_{\mathbf{SP}} h(\varepsilon)$, that in conjunction with the above implication leads to

$$\mathbb{P}^N\Big[J^\star_{\mathrm{RCP}} \leq J^\star_N + L_{\mathbf{SP}} h(\varepsilon)\Big] \geq 1-\beta.$$

Since the program SCP is just a restricted version of RCP, it is trivial that $J^\star_N \leq J^\star_{\mathrm{RCP}}$ pointwise on Ω^N, which concludes (5.6). □

In accordance with the optimization problem CCP_ε, the following theorem provides similar performance assessment but in both a priori and a posteriori fashions.

Theorem 5.3.7 (CCP_ε Confidence Interval)**.** *Consider the programs* CCP_ε *and* SCP *in* (5.1) *and* (5.2) *with the associated optimal values $J^\star_{\mathrm{CCP}_\varepsilon}$ and $J^\star_N$, respectively. Suppose Assumption 5.3.3 holds and $L_{\mathbf{SP}}$ is the constant in* (5.5)*. Let h be a ULB and $\lambda^\star_N$ the dual optimizer of* SCP*. Given β in $[0,1]$, for all $N \geq N(\varepsilon,\beta)$ defined in* (5.3)*, we have*

$$\textit{A Priori Assessment:}\qquad \mathbb{P}^N\Big[J^\star_{\mathrm{CCP}_\varepsilon} - J^\star_N \in \big[-I(\varepsilon), 0\big]\Big] \geq 1-\beta, \tag{5.8a}$$

$$\textit{A Posteriori Assessment:}\qquad \mathbb{P}^N\Big[J^\star_{\mathrm{CCP}_\varepsilon} - J^\star_N \in \big[-I_N(\varepsilon), 0\big]\Big] \geq 1-\beta, \tag{5.8b}$$

where the a priori interval $I(\varepsilon)$ is defined as in (5.7)*, and the a posteriori interval is*

$$I_N(\varepsilon) \coloneqq \min\Big\{\big\|\lambda^\star_N\big\|_1 h(\varepsilon),\ \max_{x\in\mathbb{X}} c^\intercal x - \min_{x\in\mathbb{X}} c^\intercal x\Big\}. \tag{5.9}$$

Proof. Similar to the proof of Theorem 5.3.6, we only need to show the first term of the confidence interval (5.9). In light of Theorem 5.2.2 and Lemma 5.3.2, we know that

$$\mathbb{P}^N\big[J^\star_{\mathrm{RCP}_{h(\varepsilon)}} \le J^\star_{\mathrm{CCP}_\varepsilon} \le J^\star_N\big] \ge 1-\beta. \tag{5.10}$$

In the same spirit as the previous proof, Lemma 5.3.4 ensures $J^\star_N \le J^\star_{\mathrm{RCP}} \le J^\star_{\mathrm{RCP}_{h(\varepsilon)}} + L_{\mathbf{SP}}h(\varepsilon)$ everywhere on Ω^N, which together with (5.10) arrives at (5.8a).

To show (5.8b), let us consider the scenario counterpart of the relaxed program RCP_γ in (5.4) with $\gamma := h(\varepsilon)$. We denote the optimal value of this scenario program by $J^\star_{N,h(\varepsilon)}$. Thus, we have $J^\star_{N,h(\varepsilon)} \le J^\star_{\mathrm{RCP}_{h(\varepsilon)}}$ with probability 1. Notice that Assumption 5.3.3 also holds for the scenario program SCP, and consequently Lemma 5.3.4 is applicable to SCP as well. In fact, following the proof of Lemma 5.3.4 [BV04, p. 250], one can infer that the Lipschitz constant of the perturbation function can be over approximated by the ℓ_1-norm of a dual optimizer of the optimization program. Therefore, applying Lemma 5.3.4 to SCP yields to $J^\star_N - \|\lambda^\star_N\|_1 h(\varepsilon) \le J^\star_{N,h(\varepsilon)} \le J^\star_{\mathrm{RCP}_{h(\varepsilon)}}$ pointwise on Ω^N. Substituting into (5.10) leads to (5.8b). □

The parameter ε in Theorem 5.3.6 is a design choice which can be tuned to shrink the confidence interval $[0, I(\varepsilon)]$. On the contrary, in Theorem 5.3.7 the parameter ε is part of the problem data associated with the program CCP_ε. That is, in Theorem 5.3.7 $I(\varepsilon)$ is indeed fixed and the number of scenarios N in SCP only improves the confidence level β. In a same spirit but along a different approach, [Cal10, Theorem 6.1] bounds $J^\star_N$ by the optimal solutions of two chance-constrained programs associated with different constraint violation levels, say $\bar{\varepsilon} < \varepsilon$. This value gap between the chance-constrained program and its scenario counterpart (either as explicitly derived in Theorem 5.3.7 or implicitly by two chance-constrained programs in [Cal10, Theorem 6.1]) represents an inherent difference. To arbitrarily reduce the gap for CCP_ε, one may resort to *optimally* discarding a fraction of scenarios, which is in general computationally intractable; see for example [CG11, Theorem 6.1] and [Cal10, Theorem 6.2].

By virtue of Theorem 5.3.6, the gap between $J^\star_{\mathrm{RCP}}$ and $J^\star_N$ is effectively quantified by a ULB $h(\varepsilon)$ as introduced in Definition 5.3.1. To control the behavior of $h(\varepsilon)$ as $\varepsilon \to 0$, one may require more structure on the measure $\mathbb{P}$ defined on $\big(\mathcal{D}, \mathfrak{B}(\mathcal{D})\big)$. Proposition 5.3.8 addresses this issue by introducing sufficient conditions concerning the measure of open balls in $\mathfrak{B}(\mathcal{D})$ and the continuity of the constraint mapping in the uncertainty argument.

Proposition 5.3.8. *Assume that the mapping $\mathcal{D} \ni d \mapsto f(x,d) \in \mathbb{R}$ is Lipschitz continuous with constant L_d uniformly in $x \in \mathbb{X}$. Suppose there exists a strictly increasing function $g : \mathbb{R}_+ \to [0,1]$ such that*

$$\mathbb{P}\big[\mathbb{B}_r(d)\big] \ge g(r), \qquad \forall d \in \mathcal{D},$$

where $\mathbb{B}_r(d) \subset \mathcal{D}$ is an open ball centered at d with radius r. Then, $h(\varepsilon) := L_d\, g^{-1}(\varepsilon)$ is a ULB in the sense of Definition 5.3.1, where g^{-1} is the inverse function of g.

Proof. See Appendix 5.7.1. □

Proposition 5.3.8 generalizes the corresponding result of [KT12, Lemma 3.1] by allowing the uncertainty space $\mathcal{D}$ to be possibly an infinite dimensional space. Note that the required assumptions in Proposition 5.3.8 implicitly require $\mathcal{D}$ to be bounded, though in practice this may not be really restrictive.

Remark 5.3.9. *Two remarks regarding the function g in Proposition 5.3.8 are in order:*

(i) ***Explicit expression:*** *Under the hypotheses of Proposition 5.3.8, Theorem 5.3.6 can be expressed in more explicit form. Let ε and β be in $[0,1]$, L_d be the Lipschitz constant of the constraint function f in d, $L_{\mathbf{SP}}$ be the constant (5.5), and $N(\cdot,\cdot)$ be as defined in (5.3). Then, for any $N \geq N\big(g(\frac{\varepsilon}{L_{\mathbf{SP}}L_d}),\beta\big)$ we have*

$$\mathbb{P}^N\Big[J^\star_{\mathrm{RCP}} - J^\star_N \in \big[0,\varepsilon\big]\Big] \geq 1-\beta.$$

(ii) ***Curse of dimensionality:*** *For an n_d-dimensional uncertainty set $\mathcal{D}$, the number of disjoint balls in $\mathcal{D}$ with radius r grows proportional to r^{-n_d} as r decreases. Thus, the assumptions of Proposition 5.3.8 imply that $g(r)$ is of the order of r^{n_d}. Therefore, for the desired precision ε, as detailed in the preceding remark, the required number of samples N grows exponentially as ε^{-n_d}. This appears to be an inherent feature when one seeks to bound the optimal value via scenario programs; see [LVLM08, LVLM10] for similar observations.*

5.3.2 Feasibility of RCP via SCP

In this subsection we provide an example to show the inherent difficulty of the feasibility connection from SCP to the original problem RCP. Consider the following RCP with its SCP counterpart in which both decision and uncertainty space are compact subsets of $\mathbb{R}$:

$$\begin{cases} \min\limits_x & -x \\ \text{s.t.} & x-d\leq 0, \quad \forall d\in\mathcal{D}:=[0,1] \\ & x\in\mathbb{X}:=[-1,1] \end{cases} \qquad\qquad \begin{cases} \min\limits_x & -x \\ \text{s.t.} & x-d_i\leq 0, \quad \forall i\in\{1,\cdots,N\} \\ & x\in\mathbb{X}:=[-1,1] \end{cases}.$$

It is not difficult to see that the feasible set of the robust program is $[-1,0]$ with the optimizer $x^\star=0$, whereas the optimizer of its scenario program is $x^\star_N=\min_{i\leq N} d_i$. If the probability measure $\mathbb{P}$ does not have atoms (point measure), we have $\mathbb{P}^N\big[\min_{i\leq N} d_i>0\big]=1$. Thus, one can deduce that

$$\mathbb{P}^N\big[x^\star_N\models\mathrm{RCP}\big]=0, \qquad \forall\mathbb{P}\in\mathcal{P}, \quad \forall N\in\mathbb{N},$$

where $\mathcal{P}$ is the family of all nonatomic measures on $\big(\mathcal{D},\mathfrak{B}(\mathcal{D})\big)$. More generally, if the set $\arg\max_{d\in\mathcal{D}} f(x,d)$ has measure zero for any $x\models\mathrm{RCP}$ (e.g., when f is convex in d and the boundary of $\mathcal{D}$ has zero measure), then the program SCP will almost surely return infeasible solutions to the program RCP, as the worst-case scenarios are almost surely neglected.

5.3.3 Measurability of the SCP optimizer

The objective of this subsection is to address the standing Assumption 5.2.1. The measurability of the optimizer $x^\star_N$ for the scenario program SCP is a rather involved technical issue. In fact, to the best of our knowledge, in the literature this issue is always resolved by introducing an assumption. Let us highlight that the measurability of optimal values and the set of optimizers as well as the existence of a measurable selection are classical results in this context, see for instance [RW10, Theorem 14.37, p. 664]. However, there is no a priori guarantee that the obtained optimizer of the program SCP can be viewed as a measurable mapping from $\mathcal{D}^N$ to $\mathbb{X}$. Toward this issue, we propose a "two-stage" optimization program, in the *lexicographic* sense in the context of multi-objective optimization problems [MA04], in which the measurability of this mapping is ensured for a large class of programs (not necessarily convex).

For the rest of this section we assume that $\mathbb{X} \subset \mathbb{R}^n$ is closed and the mapping $x \mapsto f(x,d)$ is lower semicontinuous. Consider the scenario program SCP as defined in (5.2) with the corresponding optimal value $J_N^\star$; SCP is assumed to be feasible with probability one. Given the same uncertainty samples $(d_i)_{i=1}^N$ as in SCP, we introduce the second program

$$\begin{cases} \min\limits_{x} & \phi(x) \\ \text{s.t.} & f(x,d_i) \le 0, \quad \forall i \in \{1,\cdots,N\} \\ & c^\intercal x \le J_N^\star \\ & x \in \mathbb{X} \end{cases}, \tag{5.11}$$

where $\phi : \mathbb{R}^n \to \mathbb{R}$ is a strictly convex function. Let us denote the optimizer of the above program by $\widetilde{x}_N^\star$. It is straightforward to observe that $\widetilde{x}_N^\star$ is indeed an optimizer of the program SCP.

Proposition 5.3.10 (Measurability of the Optimizer)**.** *Consider the sequential two-stage programs* SCP *and* (5.11)*, with the optimizer* $\widetilde{x}_N^\star$ *for the latter program. Then, the mapping* $\mathcal{D}^N \ni (d_i)_{i=1}^N \mapsto \widetilde{x}_N^\star \in \mathbb{X}$ *is a singleton and measurable.*

Proof. See Appendix 5.7.1 along with some preparatory lemmas. □

The above two-stage program may be viewed as a tie-break rule [CC05] or a regularization procedure [Cal10, Section 2.1], which was proposed to resolve the uniqueness property of the SCP optimizer. Proposition 5.3.10 indeed asserts that the same trick ensures the measurability of the optimizer as well.

Remark 5.3.11 (Measurability of the Feasible Set)**.** *The measurability of the feasibility event* $\widetilde{x}_N^\star \models \mathrm{CCP}_\varepsilon$ *(equivalently the measurability of the mapping* $x \mapsto \mathbb{P}[f(x,d) \le 0]$*) is a straightforward consequence of Proposition 5.3.10 and Fubini's Theorem [Bil95, Thm. 18.3, p. 234].*

5.4 Extension to a Class of Non-Convex Programs

This section extends the results developed in Section 5.3.1 to a class of non-convex problems. Consider a family of programs introduced in (5.1) in which the program data are indexed by k, i.e., $(\mathbb{X}_k, f_k, \varepsilon_k)_{k=1}^m$. We assume that each tuple $(\mathbb{X}_k, f_k, \varepsilon_k)$ satisfies the required conditions in Section 5.2 (i.e., $\mathbb{X}_k$ is a compact convex set and the mapping $x \mapsto f_k(x,d)$ is convex for every $d \in \mathcal{D}$), and the corresponding programs are denoted by $\mathrm{RCP}^{(\mathrm{k})}$ and $\mathrm{CCP}^{(\mathrm{k})}_{\varepsilon_{\mathrm{k}}}$ as defined in (5.1). Consider the following (non-convex) optimization problems:

$$\mathrm{RP}: \begin{cases} \min\limits_{x} & c^{\intercal}x \\ \text{s.t.} & x \models \bigcup\limits_{k=1}^{m} \mathrm{RCP}^{(\mathrm{k})} \end{cases} \qquad \mathrm{CP}: \begin{cases} \min\limits_{x} & c^{\intercal}x \\ \text{s.t.} & x \models \bigcup\limits_{k=1}^{m} \mathrm{CCP}^{(\mathrm{k})}_{\varepsilon_{\mathrm{k}}} \end{cases}, \tag{5.12}$$

where $x \models \bigcup_{k=1}^{m} \mathrm{RCP}^{(\mathrm{k})}$ $\big($resp. $x \models \bigcup_{k=1}^{m} \mathrm{CCP}^{(\mathrm{k})}_{\varepsilon_{\mathrm{k}}}\big)$ indicates that there exists $k \in \{1, \cdots, m\}$ such that $x \models \mathrm{RCP}^{(\mathrm{k})}$ $\big($resp. $x \models \mathrm{CCP}^{(\mathrm{k})}_{\varepsilon_{\mathrm{k}}}\big)$. In other words, the programs (5.12) seek an optimal solution which is feasible for at least one of the subprograms indexed by k, while the uncertainty space $\mathcal{D}$ as well as the associated measure $\mathbb{P}$ is shared between all the subprograms. Similarly, given i.i.d. samples $(d_i)_{i=1}^{N} \subset \mathcal{D}$ with respect to the probability measure $\mathbb{P}$, consider the scenario (non-convex) program

$$\mathrm{SP}: \begin{cases} \min\limits_{x} & c^{\intercal}x \\ \text{s.t.} & x \models \bigcup\limits_{k=1}^{m} \mathrm{SCP}^{(\mathrm{k})} \end{cases}. \tag{5.13}$$

Each subprogram $\mathrm{SCP}^{(\mathrm{k})}$ is defined according to the scenario convex program (5.2) associated with the program data $(\mathbb{X}_k, f_k)$ while the uncertainty samples $(d_i)_{i=1}^{N}$ are the same for all $k \in \{1, \cdots, m\}$. Before proceeding with the main result of this section, let us point out that the programs (5.12) contain, for example, a class of mixed integer programs. Let $f : \mathbb{R}^n \times \{0,1\}^{\ell} \times \mathcal{D} \to \mathbb{R}$ be the constraint function in (5.1). It is straightforward to see that a chance-constrained mixed integer program can be formulated as

$$\begin{cases} \min\limits_{x,y} & c^{\intercal}x \\ \text{s.t.} & \mathbb{P}[f(x,y,d) \le 0] \ge 1-\varepsilon \\ & x \in \mathbb{X}, \quad y \in \{0,1\}^{\ell} \end{cases} \iff \begin{cases} \min\limits_{x} & c^{\intercal}x \\ \text{s.t.} & \max\limits_{k \in \{1,\cdots,2^{\ell}\}} \mathbb{P}[f_k(x,d) \le 0] \ge 1-\varepsilon \\ & x \in \mathbb{X} \end{cases},$$

where $f_k(x,d) := f(x, y_k, d)$ for each selection of the binary variables $y_k \in \{0,1\}^{\ell}$. Then, by setting $m := 2^{\ell}$, $\mathbb{X}_k := \mathbb{X}$, $\varepsilon_k := \varepsilon$, the right-hand side of the above relation is readily in the framework of (5.12). A similar argument also holds for the robust mixed integer problems counterparts.

As a first step, we extend the feasibility result of Theorem 5.2.2 to the non-convex setting in (5.12).

Theorem 5.4.1 (CP Feasibility)**.** *Let* $\vec{\varepsilon} := (\varepsilon_1, \cdots, \varepsilon_m)$*,* $\beta \in (0,1]$*, and* $N \ge \widetilde{N}(\vec{\varepsilon}, \beta)$ *where*

$$\widetilde{N}(\vec{\varepsilon}, \beta) := \min\Big\{ N \in \mathbb{N} \;\Big|\; \sum_{k=1}^{m}\sum_{i=0}^{n-1} \binom{N}{i} \varepsilon_k^i (1-\varepsilon_k)^{N-i} \le \beta \Big\}. \tag{5.14}$$

Then, the optimizer of SP *is a feasible solution of* CP *with probability at least* $1-\beta$*.*

Proof. Let $x^{\star}_{N,k}$ be the optimizer of the subprogram $\mathrm{SCP}^{(\mathrm{k})}$. By virtue of Theorem 5.2.2, one can infer that

$$\mathbb{P}^N\Big[x^{\star}_{N,k} \not\models \mathrm{CCP}^{(\mathrm{k})}_{\varepsilon_{\mathrm{k}}}\Big] < \sum_{i=0}^{n-1} \binom{N}{i} \varepsilon_k^i (1-\varepsilon_k)^{N-i}.$$

On the other hand, it is straightforward to observe that the optimizer of the program SP, denoted by $x^\star_N$, belongs to the set $(x^\star_{N,k})_{k=1}^m$. Therefore,

$$\mathbb{P}^N\Big[x^\star_N \not\models \mathrm{CP}\Big] \le \mathbb{P}^N\Big[\exists k \in \{1,\cdots,m\} \mid x^\star_{N,k} \not\models \mathrm{CCP}^{(\mathrm{k})}_{\varepsilon_\mathrm{k}}\Big] \le \sum_{k=1}^{m} \mathbb{P}^N\Big[x^\star_{N,k} \not\models \mathrm{CCP}^{(\mathrm{k})}_{\varepsilon_\mathrm{k}}\Big]$$
$$< \sum_{k=1}^{m}\sum_{i=0}^{n-1}\binom{N}{i}\varepsilon_k^i(1-\varepsilon_k)^{N-i},$$

leading to the desired assertion. □

Remark 5.4.2 (Growth rate)**.** *Notice that the number of subprograms, m, contributes to the confidence level β in a linear fashion. As an illustration, suppose $\varepsilon_k \coloneqq \varepsilon$. In this case, one can easily verify that the confidence level of the non-convex program* SP *can be set equal to $\frac{\beta}{m}$, where β is the confidence level of each of the subprograms* $\mathrm{SCP}^{(\mathrm{k})}$*. From a computational perspective, one can follow the same calculation as in [Cal09], and deduce that the contribution of m to the number of the required samples $\widetilde{N}$ appears in a logarithm. Thus, in our example of mixed integer programming above, the required number of samples grows linearly in the number of binary variables, which for most of applications could be considered a reasonable growth rate.*

The literature on computational schemes on non-convex problems is mainly based on statistical learning methods. A recent example of this nature is [ATC09], which considers a class of problems involving Boolean expressions of polynomial functions. Given the degree and number of polynomial functions (α and k, respectively), the explicit sample bounds of [ATC09] scale with $\varepsilon^{-1}\log(\alpha k\varepsilon^{-1})$ as opposed to our result in (5.14) which grows proportional to $\varepsilon^{-1}\log(m)$. We now proceed to extend the main results of Subsection 5.3.1, i.e., Theorems 5.3.6 and 5.3.7, to the non-convex settings (5.12) and (5.13) at once.

Theorem 5.4.3 (RP & CP Confidence Intervals)**.** *Consider the programs* RP*,* CP*, and* SP *in (5.12) and (5.13) with the corresponding optimal values $J^\star_{\mathrm{RP}}$, $J^\star_{\mathrm{CP}}$, and $J^\star_N$. Given $k \in \{1,\cdots,m\}$ and the program data $(\mathbb{X}_k, f_k)$, let Assumption 5.3.3 hold and $I^{(k)}$ and $I^{(k)}_N$ be the a priori and a posteriori confidence intervals of the k^{th} subprogram as defined in (5.7) and (5.9). Then, given $\beta \in [0,1]$ and $\vec{\varepsilon} \coloneqq (\varepsilon_1,\cdots,\varepsilon_m) \in [0,1]^m$, for all $N \ge \widetilde{N}(\vec{\varepsilon},\beta)$ as defined in (5.14) we have*

$$\textit{A Priori Assessment:}\quad \begin{cases} \mathbb{P}^N\Big[J^\star_{\mathrm{RP}} - J^\star_N \in \big[0,\ \max_{k\le m} I^{(k)}(\varepsilon)\big]\Big] & \ge 1-\beta, \\ \mathbb{P}^N\Big[J^\star_{\mathrm{CP}} - J^\star_N \in \big[-\max_{k\le m} I^{(k)}(\varepsilon),\ 0\big]\Big] & \ge 1-\beta, \end{cases}$$

$$\textit{A Posteriori Assessment:}\quad \mathbb{P}^N\Big[J^\star_{\mathrm{CP}} - J^\star_N \in \big[-\max_{k\le m} I^{(k)}_N(\varepsilon),\ 0\big]\Big] \ge 1-\beta.$$

Sketch of the proof. The proof effectively follows the same lines as in the proofs of Theorems 5.3.6 and 5.3.7. To adapt the required preliminaries, let us recall again that the optimizer of the programs (5.12) is one of the optimizers of the respective subprograms. The same assertion holds for the random program (5.13) as well. Moreover, since each subprogram of (5.13) fulfills the assumptions of Subsection 5.3.1, Lemmas 5.3.2 and 5.3.4 also hold for each subprogram with the corresponding data $(\mathbb{X}_k, f_k)$. Therefore, in light of Theorem 5.4.1, it only suffices to consider the worst-case possibility among all the subprograms. □

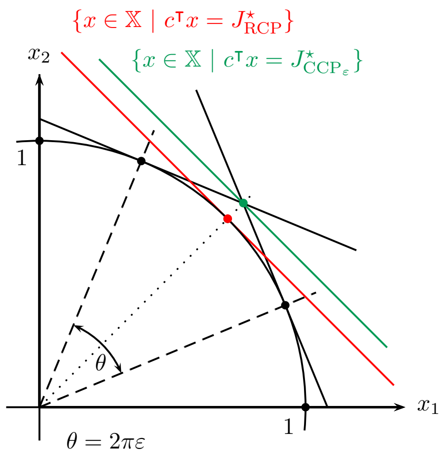

Figure 5.2: Analytical solutions of Example 1

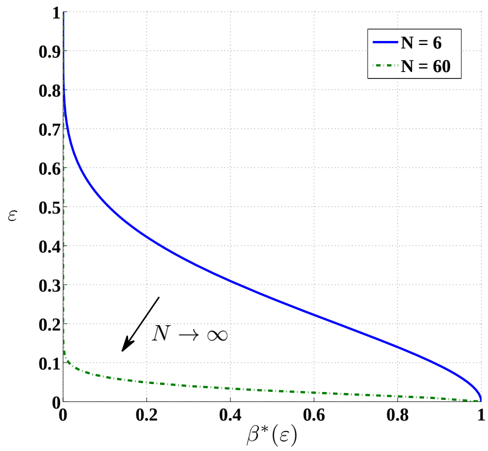

Figure 5.3: Behavior of the confidence level $\beta^*(\varepsilon)$ in terms of scenarios numbers N

5.5 Numerical Examples

This section presents two examples to illustrate the theoretical results developed in the preceding sections and their performance. We first apply the results to a simple example whose analytical solution is available.

5.5.1 Example 1: Quadratic Constraints via Infinite Hyperplanes

Let $x = [x_1, x_2]^\intercal$ be the decision variables selected in the compact set $\mathbb{X} := [0,1]^2 \subset \mathbb{R}^2$, the linear objective function defined by $c := [-1,-1]^\intercal$, and the constraint function $f(x,d) := x_1\cos(d) + x_2\sin(d) - 1$ where the uncertainty d comes from the set $\mathcal{D} := [0, 2\pi]$. Consider the optimization problems introduced in (5.1) where $\mathbb{P}$ is the uniform probability measure on $\mathcal{D}$. It is not difficult to infer that the infinitely many hyperplane constraints can be replaced by a simple quadratic constraint. That is, for any $\gamma \geq 0$

$$\max_{d\in[0,2\pi]} x_1\cos(d) + x_2\sin(d) - 1 \leq \gamma \qquad \Longleftrightarrow \qquad x_1^2 + x_2^2 \leq (\gamma+1)^2.$$

In the light of the above observation, we have the analytical solutions

$$J^\star_{\mathrm{RCP}_\gamma} = \max\Big\{-\sqrt{2}(\gamma+1), -2\Big\}, \qquad J^\star_{\mathrm{CCP}_\varepsilon} = \max\Big\{\frac{-\sqrt{2}}{\cos(\pi\varepsilon)}, -2\Big\}, \tag{5.15}$$

where $J^\star_{\mathrm{RCP}_\gamma}$ and $J^\star_{\mathrm{CCP}_\varepsilon}$ are the optimal values of the optimization problems RCP_γ and CCP_ε as defined in (5.4) and (5.1), respectively. The pictorial representation of the solutions is in Figure 5.2.

Let us fix the number of scenarios N for SCP in (5.2) with the optimal value $J^\star_N$. Given N and $\varepsilon \in [0,1]$, the confidence level $\beta \in [0,1]$ associated with our theoretical results is

$$\beta^*(\varepsilon) := \sum_{i=0}^{n-1}\binom{N}{i}\varepsilon^i(1-\varepsilon)^{N-i} = (1-\varepsilon)^N + N\varepsilon(1-\varepsilon)^{N-1},$$

where $n = 2$ in this example. Figure 5.3 depicts the behavior of $\beta^*(\varepsilon)$ for different values of N. Note that $x_0 = [0,0]^\intercal$ is a Slater point in the sense of Assumption 5.3.3 with the corresponding constant $L_{\mathbf{SP}} := \frac{-2-0}{-1} = 2$ (cf. (5.5)). Moreover, it is easy to see that the mapping $d \mapsto f(x,d)$ has the Lipschitz constant $L_d = \sqrt{2}$ over the compact set $\mathbb{X} = [0,1]^2$. Thanks to Proposition 5.3.8 (and periodicity of the constraint function over the interval $[0,2\pi]$), it is straightforward to introduce $g(r) = \frac{r}{\pi}$, and consequently obtain the ULB candidate $h(\varepsilon) := \sqrt{2}\pi\varepsilon$. Then, the confidence interval defined in (5.7) is $I(\varepsilon) := \max\{2\sqrt{2}\pi\varepsilon, 2\}$. As shown in Theorem 5.3.6 (resp. Theorem 5.3.7) we know that $J^\star_{\mathrm{RCP}} - J^\star_N \in [0, I(\varepsilon)]$ $\big($resp. $J^\star_{\mathrm{CCP}_\varepsilon} - J^\star_N \in [-I(\varepsilon), 0]\big)$ with probability at least $1 - \beta^*(\varepsilon)$ for any $\varepsilon \in [0,1]$. To validate this result, we solve the program SCP for M different experiments. For each experiment $k \in \{1, \cdots, M\}$, we draw N scenarios $\big(d_i(k)\big)_{i=1}^N \subset [0,2\pi]$ with respect to the uniform probability distribution $\mathbb{P}$ and solve the program SCP. Let $J^\star_N(k)$ be the optimal value of the k^{th} experiment. Given $\beta \in [0,1]$, the empirical confidence interval of the program RCP can be represented by the minimal $\widetilde{I}(\beta)$ so that the interval $[0, \widetilde{I}(\beta)]$ contains $J^\star_N(m) - J^\star_{\mathrm{RCP}}$ for at least m experiments where $\frac{m}{M} \geq 1 - \beta$, i.e.,

$$\widetilde{I}(\beta) := \min\Big\{\widetilde{I} \in \mathbb{R}_+ \;\big|\; \exists A \subset \{1, \cdots, M\} \;: \\ |A| \geq (1-\beta)M \quad \text{and} \quad J^\star_{\mathrm{RCP}} - J^\star_N(k) \in [0, \widetilde{I}] \;\forall k \in A\Big\}.$$

Regarding the program CCP_ε, notice that the empirical confidence interval depends on both parameters ε and β since the analytical optimal values $J^\star_{\mathrm{CCP}_\varepsilon}$ depends on ε as well. Hence, we define

$$\widetilde{I}_\varepsilon(\beta) := \min\Big\{\widetilde{I} \in \mathbb{R}_+ \;\big|\; \exists A \subset \{1, \cdots, M\} \;: \\ |A| \geq (1-\beta)M \quad \text{and} \quad J^\star_{\mathrm{CCP}_\varepsilon} - J^\star_N(k) \in [-\widetilde{I}, 0] \;\forall k \in A\Big\}.$$

The sets $\widetilde{I}(\beta)$ and $\widetilde{I}_\varepsilon(\beta)$ are in close relation with sample quantiles in the sense of [Sha03, Section 5.3.1]. In the following simulations the number of experiments is set to $M = 2000$. Figures 5.4(a) and 5.4(b) depict our theoretical performance bound $I(\varepsilon)$ for $N = 6$ and $N = 60$ in comparison with the empirical bounds $\widetilde{I}\big(\beta^*(\varepsilon)\big)$ and $\widetilde{I}_\varepsilon\big(\beta^*(\varepsilon)\big)$ where $\beta^*(\varepsilon)$ is the confidence level in Figure 5.3. As our theoretical results suggest, the confidence interval $[0, I(\varepsilon)]$ (resp. $[-I(\varepsilon), 0]$) contains the empirical interval $\big[0, \widetilde{I}\big(\beta^*(\varepsilon)\big)\big]$ (resp. $\big[-\widetilde{I}_\varepsilon\big(\beta^*(\varepsilon)\big), 0\big]$). Moreover, to demonstrate the a posteriori confidence interval in Theorem 5.3.7, we choose one of the experiments and depict the corresponding confidence interval $I_N(\varepsilon)$ versus $\beta^*(\varepsilon)$ as well. Note that in both cases of Figure 5.4 the a posteriori confidence interval proposes a tighter bound than the a priori confidence interval. With this observation, we conjecture that in general the dual optimizer of SCP may happen to be a better approximation in comparison with the constant $L_{\mathbf{SP}}$ introduced in (5.5).

5.5.2 Example 2: Fault Detection and Isolation

In the second example we illustrate the theoretical results developed in this chapter to the problem of fault detection and isolation (FDI) discussed in Chapter 4. Let us recall that the

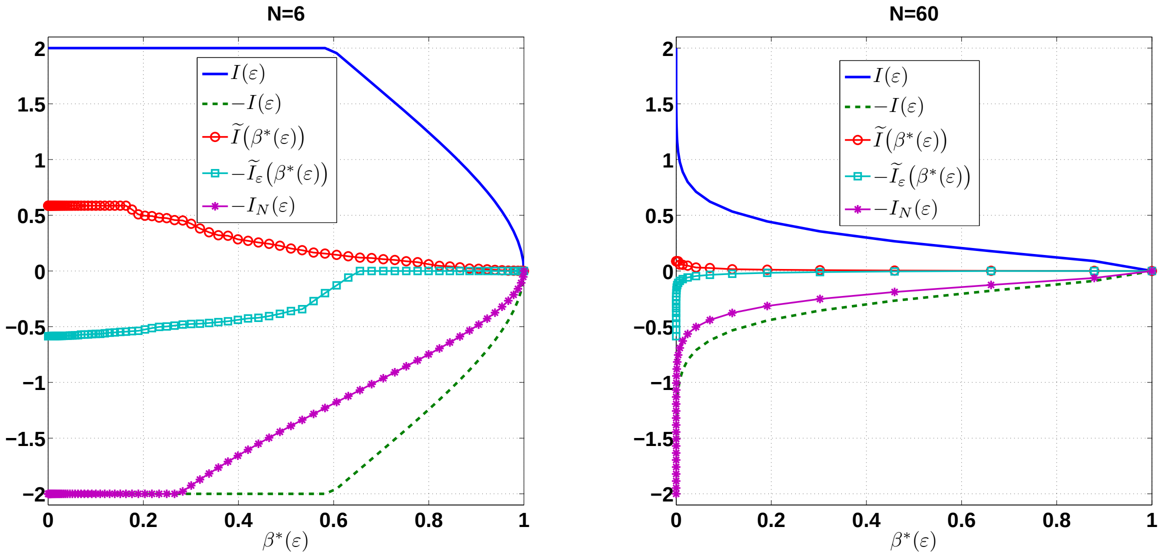

(a) Simulations for $N = 6$

(b) Simulations for $N = 60$

Figure 5.4: Numerical results for Example 1

FDI problem is essentially designing a filter fed by available signals as inputs (e.g., control signals and given measurements) whose output implements a non-zero mapping from the fault to the residual while decoupling unknown disturbances (cf. the mapping requirements (4.1) and the robust version RP in (4.2)).

As a particular subclass of DAEs, consider the nonlinear differential equation

$$\begin{cases} \dot{X}(t) = E\big(X(t)\big) + AX(t) + B_d d(t) + B_f f(t) \\ Y(t) = CX(t) \end{cases}, \tag{5.16}$$

where the matrices A, B_d, B_f, C and the function $E(\,\cdot\,)$ describe the linear and nonlinear dynamics of the model, respectively (see (4.6) and the following discussion). Restricting the class of FDI filters to linear operator, we obtain a residual consisting of two terms: $r = G[x](f) + r[x](d)$ where $G[x]$ is a linear time invariant transfer function expressing the mapping from the fault $f(\,\cdot\,)$ to the residual, and $r[x](d)$ is the contribution of the unknown disturbance $d(\,\cdot\,)$, and $x \in \mathbb{R}^n$ denotes the coefficients of the FDI filter to be designed (cf. (4.11)[2]). In this light, to minimize the impact of nonlinearities and disturbances on the residual, an optimal FDI filter can be obtained by the min-max program

$$\begin{cases} \min\limits_{x,\gamma} & \gamma \\ \text{s.t.} & x^\intercal Q_d x \le \gamma, \quad \forall d \in \mathcal{D} \\ & Hx = 0 \\ & \big\|Fx\big\|_\infty \ge 1 \end{cases}, \tag{5.17}$$

where the quadratic term $x^\intercal Q_d x$ represents the $\mathcal{L}_2$-norm of $r[x](d)$ over a given receding horizon, $\mathcal{D}$ is the space of possible disturbance patterns, and the last (non-convex) constraint is

[2] Note that the decision variable x here corresponds to the filter coefficients $\bar{N}$ in (4.11).

concerned with the norm of $G[x]$ as an operator (cf. (4.16)). The matrices H and F are determined by the linear terms of the system dynamics (5.16), and the positive semidefinite matrix Q_d reflects the nonlinearity signature of the system dynamics in the presence of a disturbance pattern d; it depends on d and the nonlinear term $E(\cdot)$ of (5.16) (cf. the signature matrix (4.15)). We refer to Approach (II) in Section 4.4.2 for details of the derivation of the above program.

For numerical case study, we consider an application of the above FDI design to detect a cyber intrusion in a two-area power network discussed in Section 4.5. The setup in this example is a simplified version of the test system in Section 4.5.3 where each power area contains one generator ($g \coloneqq 2$). Thus, the state in (5.16) is described by

$$X \coloneqq \left[\Delta\phi, \{\Delta f_i\}_{1:2}, \{\Delta P_{m_i}\}_{1:2}, \{\Delta P_{agc_i}\}_{1:2}\right]^{\intercal},$$

where $\Delta\phi$ is the voltage angle difference between the ends of the tie line, Δf_i the generator frequency, ΔP_{m_i} the generated mechanical power, and ΔP_{agc_i} the automatic generation control (AGC) signal in each area.[3] The system dynamics is modeled in the framework of (5.16); the details are provided in Appendix 5.7.2. The disturbance signal $d(\cdot)$ represents a load deviation that may occur in the first area. The signal f models the intrusion signal in the AGC of the first area, and the measurement signals are the frequencies and output power of the turbines, i.e., $Y = \left[\{\Delta f_i\}_{1:2}, \{\Delta P_{m_i}\}_{1:2}\right]^{\intercal}$.

For a given horizon $T > 0$, we consider the class of disturbance patterns

$$\mathcal{D} \coloneqq \left\{ d : [0,T] \to \mathbb{R} \;\middle|\; \exists \alpha \in [0,1],\ d(t) \coloneqq \sum_{k=0}^{p} a_k(\alpha)\cos\left(\frac{2\pi}{T}kt\right) \right\},$$

where $a_k(\alpha)$ are the constant coefficients parametrized by α (cf. (4.30)). The choice of $\mathcal{D}$ allows one to exploit available spectrum information of the disturbance signals. In this example, motivated by the emphasis on both low and high frequency regions, we assume $a_k(\alpha) \coloneqq 5\big(\alpha 0.5^k + (1-\alpha)0.5^{|10-k|}\big)$, $p = 30$, and $T = 4$ sec. For scenario generation, we consider a uniform probability distribution for the parameter $\alpha \in [0,1]$, which in fact induces the probability measure $\mathbb{P}$ on $\mathcal{D}$. Let $d_0 \in \mathcal{D}$ be a disturbance signature with the corresponding parameter α_0. It is straightforward to observe that

$$\begin{aligned}
\mathbb{P}\big[\|d - d_0\|_{\mathcal{L}_2} < r\big] &= \mathbb{P}\left[\frac{T}{2}\sum_{k=0}^{p} \big|a_k(\alpha) - a_k(\alpha_0)\big|^2 < r^2\right] \\
&= \mathbb{P}\left[|\alpha - \alpha_0| < \frac{\sqrt{2}\, r}{5\sqrt{T\sum_{k=0}^{p}\big(0.5^k - 0.5^{|10-k|}\big)^2}}\right] \\
&= \mathbb{P}\big[|\alpha - \alpha_0| < 0.142r\big] \geq 0.142r =: g(r),
\end{aligned}$$

where the function g, denoted in view of Proposition 5.3.8, is an invertible lower bound for the measure of open balls in $\mathcal{D}$. For the particular set of parameters in this example and specific operating region of interest, one can show that the mapping $d \mapsto Q_d$ is Lipschitz continuous

[3] The symbol Δ stands for the deviation from the nominal value.

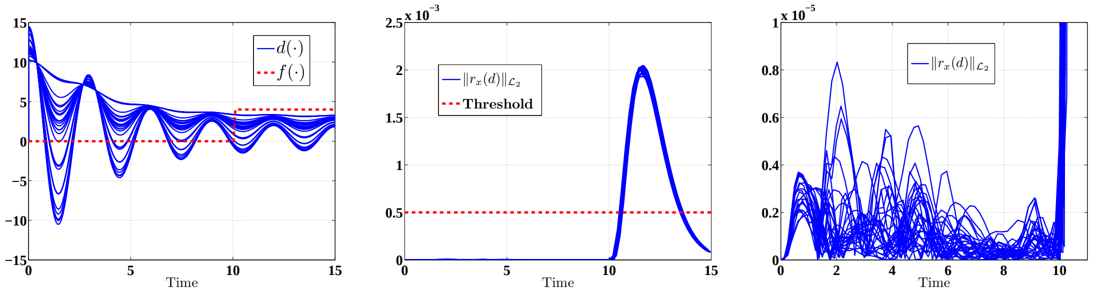

(a) Scenarios of the disturbance signatures (solid), and intrusion signal (dash)

(b) Energy of the filter residual (solid), and the threshold level (dash)

(c) Residual response before the intrusion starts

Figure 5.5: Numerical results for Example 2

with the constant $L_d = 0.02$; see Appendix 5.7.2 for more details. By virtue of Proposition 5.3.8 and normalizing[4] the optimizer of the SCP counterpart of the program (5.17), we can introduce the ULB candidate

$$h(\varepsilon) \coloneqq L_d\, g^{-1}(\varepsilon) = 0.14\varepsilon.$$

Notice that the Infinite norm constraint in (5.17) is in fact a non-convex constraint. However, one may view it as the union of a finite number of constraint sets, see Remark 4.4.3. Therefore, the optimization problem (5.17) is already in the framework of RP as introduced in (5.12) where m is the number of rows in matrix F. It is remarkable that $m-1$ equals the degree of the FDI filter chosen a priori. Thanks to the min-max structure of the robust program (5.17), the Lipschitz constant of Lemma 5.3.4 for each subprogram of (5.17) is $L_{\mathbf{SP}} = 1$, see Remark 5.3.5.

In this example, the dimension of the decision variable x is $n = 55$, the number of rows in F is $m = 5$, and the confidence level is set to $\beta = 0.01$. Therefore, to achieve the confidence interval $I(\varepsilon) = h(\varepsilon) = 5 \times 10^{-4}$, we need to set $\varepsilon = 3.57 \times 10^{-3}$ which, due to Theorem 5.4.1, requires to generate N disturbance signatures $d \in \mathcal{D}$ so that

$$N \geq \min\left\{ N \in \mathbb{N} \;\middle|\; \sum_{i=0}^{n-1} \binom{N}{i} \varepsilon^i (1-\varepsilon)^{N-i} \leq \frac{\beta}{m} \right\} = 22618.$$

Figures 5.5 demonstrate the numerical results of Example 2 over the course of 15 seconds. In Figure 5.5(a), 30 different realizations of disturbance inputs as well as an intrusion signal starting from $t = 10$ are shown in solid and dash curves, respectively. Figure 5.5(b) depicts the energy of the filter residual for the last $T = 4$ seconds (solid), and the threshold level associated with confidence $\beta = 0.01$ (dash). Notice that the proposed threshold is $\gamma^\star + 0.0005$, where $\gamma^\star$ is the optimal solution of the random counterpart of the program (5.17) with N $=$ 22618 scenarios. Figure 5.5(c) presents the filter response which is the same figure as 5.5(b) but zoomed in on the period prior to the intrusion.

[4]Due to the linearity of the filter operator, one can always normalize the filter coefficients with no performance deterioration; see Section 4.4.3.

5.6 Summary and Outlook

In this chapter we presented probabilistic performance bounds for both RCP and CCP$_\varepsilon$ via SCP. The proposed bounds are based on considering the tail probability of the worst-case constraint violation of the SCP solution as introduced in [KT12] together with some classical results from perturbation theory of convex optimization. In contrast to earlier approaches, this methodology is, to the best of our knowledge, the first confidence bounds for the objective performance of RCPs and CCPs based on scenario programs. Subsequently, we extended our results to a certain class of non-convex programs allowing for binary decision variables.

For future work, in light of Theorems 5.3.6 and 5.3.7, we aim to study the derivation of ULBs as introduced in Definition 5.3.1. Meaningful ULBs may depend highly on the individual structure of the optimization problems, in particular the uncertainty set and the constraint functions. For certain classes of problems, [KT12] provides a constructive approach to obtain ULBs. Another potential direction is the estimation of constant $L_{\mathbf{SP}}$ in Theorems 5.3.6 and 5.3.7, see Remark 5.3.5. This problem may be closely related to the estimation of the dual optimizers of RCPs.

5.7 Appendix

5.7.1 Proofs

This section collects the technical proofs skipped throughout the chapter.

Proof of Lemma 5.3.2. Let h is a ULB as introduced in Definition 5.3.1, $x_0 \models \mathrm{CCP}_\varepsilon$, and $f^*(x_0) \coloneqq \sup_{v\in\mathcal{D}} f(x_0, v)$. By definition of CCP$_\varepsilon$ and p, the tail probability of the worst-case violation, we have

$$p\big(x_0, f^*(x_0)\big) \le \varepsilon \implies \inf_{x\in\mathbb{X}} p\big(x, f^*(x_0)\big) \le \varepsilon \implies f^*(x_0) \le h(\varepsilon) \implies x_0 \models \mathrm{RCP}_{h(\varepsilon)} \qquad \square$$

Proof of Proposition 5.3.8. Given $x \in \mathbb{X}$, let $(v_i)_{i\in\mathbb{N}}$ be a sequence in $\mathcal{D}$ so that

$$\limsup_{i\in\mathbb{N}} f(x, v_i) = \sup_{v\in\mathcal{D}} f(x, v).$$

Thus, in light of Definition 5.3.1 we have

$$\begin{aligned} p(x,\delta) &= \mathbb{P}\Big[\sup_{v\in\mathcal{D}} f(x,v) - f(x,d) < \delta\Big] = \mathbb{P}\Big[\limsup_{i\in\mathbb{N}} f(x,v_i) - f(x,d) < \delta\Big] \\ &\ge \mathbb{P}\Big[\limsup_{i\in\mathbb{N}} L_d\|v_i - d\| < \delta\Big] \ge \limsup_{i\in\mathbb{N}} \mathbb{P}\Big[\|v_i - d\| < \frac{\delta}{L_d}\Big] \\ &= \limsup_{i\in\mathbb{N}} \mathbb{P}\Big[\mathbb{B}_{\frac{\delta}{L_d}}(v_i)\Big] \ge g\big(\frac{\delta}{L_d}\big), \end{aligned} \tag{5.18}$$

where the first inequality in (5.18) follows from the Lipschitz continuity of f with respect to d, and the second inequality in (5.18) is due to Fatou's lemma [Rud87, p. 23]. Hence, in view of the ULB definition and the above analysis, we arrive at

$$\sup\Big\{\delta \in \mathbb{R}_+ \;\big|\; \inf_{x\in\mathbb{X}} p(x,\delta) \le \varepsilon\Big\} \le \sup\Big\{\delta \in \mathbb{R}_+ \mid g\big(\frac{\delta}{L_d}\big) \le \varepsilon\Big\} = L_d g^{-1}(\varepsilon). \qquad \square$$

Proof of Lemma 5.3.4. It is well-known that under the strong duality condition the mapping $\gamma \mapsto \mathrm{RCP}_\gamma$, the so-called perturbation function, is Lipschitz continuous with the constant $\|\lambda^\star\|_1$ where $\lambda^\star$ is a dual optimizer of the program RCP; see [BV04, p. 250] for the proof and [Roc97, Section 28] for more details in this direction. Now Lemma 5.3.4 follows from [NO08, Lemma 1], which essentially implies $\|\lambda^\star\|_1 \le L_{\mathbf{SP}}$ where $L_{\mathbf{SP}}$ is the constant (5.5) corresponding to any Slater point in the sense of Assumption 5.3.3. $\square$

To prove Proposition 5.3.10, we need some preliminaries.

Lemma 5.7.1. *Let $\mathcal{C}$ be the set of all lower semicontinuous functions from $\mathbb{X} \subset \mathbb{R}^n$ to $\mathbb{R}$. Consider the mapping $J : \mathcal{C} \to \mathbb{R}$ defined by the optimization program*

$$\begin{cases} J(g) := & \min\limits_{x} & c^\intercal x \\ & \text{s.t.} & g(x) \le 0 \\ & & x \in \mathbb{X} \end{cases}. \tag{5.19}$$

Then, the function J is measurable where the space of $\mathcal{C}$ is endowed with the infinite norm and the respective Borel σ-algebra.[5]

Proof. The proof is an application of [RW10, Theorem 14.37, p. 664]. Let us define the set-valued mapping $S : \mathcal{C} \rightrightarrows \mathbb{X} \times \mathbb{R}$ as follows:

$$S(g) := \big\{(x,\alpha) \in \mathbb{X} \times \mathbb{R} \;\big|\; \{g(x) \le 0\} \;\&\; \{c^\intercal x \le \alpha\}\big\}.$$

We first show that S is a normal integrand in the sense of [RW10, Definition 14.27, p. 661]. Since g is lower semicontinuous, then S is clearly closed-valued. We then only need to show that S is measurable according to [RW10, Definition 14.1, p. 643]. Let $O \subset \mathbb{X} \times \mathbb{R}$ be an open set, $(x_0, \alpha_0) \in O$ and $g_0 \in S^{-1}(x_0, \alpha_0)$. Observe that for sufficiently small $\varepsilon > 0$ we have $\mathbb{B}_\varepsilon(g_0) \subset S^{-1}(O)$ where $\mathbb{B}_\varepsilon(g_0) := \{g \in \mathcal{C} \mid \sup_{x\in\mathbb{X}} \|g(x) - g_0(x)\| \le \varepsilon\}$, that implies that $S^{-1}(O)$ is open, and in particular measurable. Thereby, S is measurable and hence a normal integral. Now the desired measurability readily follows from [RW10, Theorem 14.37, p. 664]. $\square$

Lemma 5.7.2. *Let $\phi : \mathbb{R}^n \to \mathbb{R}$ be a strictly convex function, and $\widetilde{J} : \mathcal{C} \to \mathbb{R}$ defined as follows:*

$$\begin{cases} \widetilde{J}(g) := & \min\limits_{x} & \phi(x) \\ & \text{s.t.} & g(x) \le 0 \\ & & c^\intercal x \le J(g) \\ & & x \in \mathbb{X} \end{cases}, \tag{5.20}$$

[5] Under assumptions of Section 5.2, one can show a stronger assertion that the mapping $g \mapsto J(g)$ is indeed lower semicontinuous; see for instance [BGK+83, Theorem 4.3.2, p. 67]. Thanks to a personal communication with Diethard Klatte, it turns out that the statement can be even extended to continuity if Assumption 5.3.3 also holds.

where $J(g)$ is the function introduced in (5.19). Let $\widetilde{x}^\star(g)$ denote the set of optimizers of the program (5.20). Then, the mapping $\mathcal{C} \ni g \mapsto \widetilde{x}^\star \in \mathbb{R}^d$ is a measurable singleton.

Proof. Let us define the set-valued mapping $S : \mathcal{C} \rightrightarrows \mathbb{X} \times \mathbb{R}$

$$S(g) \coloneqq \big\{(x, \alpha) \in \mathbb{X} \times \mathbb{R} \;\big|\; \{g(x) \leq 0\} \;\&\; \{c^\intercal x - J(g) \leq 0\} \;\&\; \{\phi(x) \leq \alpha\}\big\}.$$

By virtue of the measurability of the mapping $g \mapsto J(g)$ in Lemma 5.7.1 and along the same line of its proof, we know that S is a normal integral. Now, by [RW10, Theorem 14.37, p. 664] the existence of a measurable selection for the optimizer $\widetilde{x}^\star(g)$ as a function of $g \in \mathcal{C}$ is guaranteed. On the other hand, since $\phi : \mathbb{R}^n \to \mathbb{R}$ is strictly convex, the minimizer of the program (5.20) is unique. Therefore, $\widetilde{x}^\star(g)$ is a singleton and the desired measurability property follows at once. □

We now have all the required results to prove Proposition 5.3.10:

Proof of Proposition 5.3.10. Let $g : \mathcal{D}^N \to \mathcal{C}$ defined as

$$g(d_1, \cdots, d_N) \coloneqq \max_{i \in \{1, \cdots, N\}} f(x, d_i). \tag{5.21}$$

The measurability of the mapping (5.21) is ensured by the measurability assumption of the mapping $d \mapsto f(x, d)$ for each x. It is straightforward to observe that the optimizer of the program (5.11) can be viewed as the composition $\widetilde{x}^\star_N = \widetilde{x}^\star \circ g(d_1, \cdots, d_N)$ where $\widetilde{x}^\star$ is the optimizer of the program (5.20) and g is defined as in (5.21). Hence, the desired implication follows directly from the measurability of the mapping (5.21) and Lemma 5.7.2. □

5.7.2 Details of Example 2

This appendix provides details of Example 2 in Subsection 5.5.2.

A. Mathematical model description

The two-area power network is described by the set of nonlinear ordinary differential equations

$$\begin{aligned}
\Delta\dot{\phi} &= 2\pi(\Delta f_1 - \Delta f_2), \\
\Delta\dot{f}_i &= \frac{f_0}{2H_i S_{B_i}}\Big(-\frac{1}{D_i}\Delta f_i - P_T \sin\Delta\phi + \Delta P_{m_i} - \Delta P_{load_i}\Big), \\
\Delta\dot{P}_{m_i} &= \frac{1}{T_{ch_i}}\Big(-\frac{1}{S_i}\Delta f_i - \Delta P_{m_i} + \Delta P_{agc_i}\Big), \\
\Delta\dot{P}_{agc_i} &= \Big(\frac{1}{D_i}\frac{C_i f_0}{2S_i H_i S_{B_i}} - \frac{1}{S_i}\frac{1}{T_{N_i}}\Big)\Delta f_i \\
&\quad - \frac{C_i f_0}{2S_i H_i S_{B_i}}(\Delta P_{m_i} - \Delta P_{load_i}) - \frac{C_i f_0}{2S_i H_i S_{B_i}}\Delta P_{agc_i} \\
&\quad - \Big(\frac{1}{T_{N_i}} - \frac{C_i f_0}{2S_i H_i S_{B_i}}\Big)P_T \sin\Delta\phi - 2\pi C_i P_T(\Delta f_1 - \Delta f_2)\cos\Delta\phi,
\end{aligned}$$

where $i \in \{1,2\}$ is the index of each area, $X := [\Delta\phi, \{\Delta f_i\}_{1:2}, \{\Delta P_{m_i}\}_{1:2}, \{\Delta P_{agc_i}\}_{1:2}]^\intercal \in \mathbb{R}^7$ is the state vector, and the constant parameters in this example are chosen the same for both areas as $T_{ch_i} = 5\,\text{sec}, S_{B_i} = 1.8$ GW, $f_0 = 50$ Hz, $H_i = 6.5\,\text{sec}$, $D_i = 428.6$ Hz/GW, $S_i = 1.389$ Hz/GW, $C_i = 0.1$, $T_{N_i} = 30$, $P_T = 0.15$ GW. We refer to [MEVM+10] for physical interpretation of these parameters and more details on the model equations. In the example, we assume that $\Delta P_{load_1} = d$ where $d \in \mathcal{D}$ is the disturbance signal and $\Delta P_{load_2} \equiv 0$.

B. Lipschitz constant of the mapping $d \mapsto Q_d$

This mapping can be viewed in two steps: $d \mapsto E(X)$ and $E(X) \mapsto Q_d$ where X is the solution process in the presence of the disturbance input d, and E the nonlinear term of the ODE (5.16). The key step is to approximate the Lipschitz constant of the first mapping $d \mapsto E(X)$. The classical result of the continuity of the ODEs solution, obtained by Lipschitz continuity of the vector field and Gronwall's inequality, turns out to be too conservative in this case. We then invoke a Lyapunov-like approach to address this issue more efficiently. Let us define the shorthand $h(X,d) := E(X) + AX + B_d d$. Suppose there exist a function $V : \mathbb{R}^7 \times \mathbb{R}^7 \to \mathbb{R}_+$ and positive constants κ, ρ so that for every $X, \widetilde{X} \in \mathbb{R}^7$ and $d, \widetilde{d} \in \mathbb{R}$

$$\big\|E(X) - E(\widetilde{X})\big\|^2 \le V(X, \widetilde{X}) \tag{5.22a}$$

$$\partial_X V(X,\widetilde{X}) h(X,d) + \partial_{\widetilde{X}} V(X,\widetilde{X}) h(\widetilde{X},\widetilde{d}) \le -\kappa V(X,\widetilde{X}) + \rho\big|d - \widetilde{d}\big|. \tag{5.22b}$$

Using standard Gronwall's inequality, one can show that under conditions (5.22) we have

$$\begin{aligned}\big\|E(X) - E(\widetilde{X})\big\|^2_{\mathcal{L}_2} &\le \int_0^T \big\|E\big(X(t)\big) - E\big(\widetilde{X}(t)\big)\big\|^2 \mathrm{d}t \le \int_0^T V\big(X(t), \widetilde{X}(t)\big)\mathrm{d}t \\ &\le \rho \int_0^T \mathrm{e}^{-\kappa t} \int_0^t \big|d(s) - \widetilde{d}(s)\big| \mathrm{d}s\mathrm{d}t \le \frac{2\rho}{3} T\sqrt{T} \|d - \widetilde{d}\|_{\mathcal{L}_2}.\end{aligned}$$

In [ZMEM+13b, Theorem 3.3], a similar technique is discussed in more detail to establish a connection between the Lyapunov function and continuity of the solution trajectories. In order to find a Lyapunov function in the above sense, we limit our search domain to the quadratic functions, i.e., $V(X,\widetilde{X}) = (X - \widetilde{X})^\intercal Q (X - \widetilde{X})$ for some positive semidefinite matrix Q. It is not difficult to deduce that the nonlinear term E effectively depends only on the state $\Delta\phi$. Hence, to fulfill the requirement (5.22a) it suffices to guarantee $Q \succeq vv^\intercal$ where $v = [0,0,1,0,0,0,0]^\intercal$. Setting $\kappa = 0.01$, we then solve the set of linear matrix inequalities (LMIs)

$$\begin{cases} \min\limits_{\sigma, Q} & \sigma \\ \text{s.t.} & QA^\intercal + AQ \preceq -\kappa Q \\ & vv^\intercal \preceq Q \preceq \sigma I \end{cases},$$

which provides a local Lyapunov function in the sense of (5.22). Note that one can always extract the linear part of E and add it to the matrix A. Now, by numerical inspection, it turns out that for the specific system parameters of this example, V obtained from the above LMIs is a Lyapunov function in the domain of $\Delta f_i \in [-0.1, 0.1]$ Hz, $\Delta\phi \in [-10°, 10°]$, $\Delta p_{m_i} \in [-10, 10]$ MW, $\Delta p_{agc_i} \in [-15, 15]$ MW. Therefore, the parameter ρ in (5.22b) can be numerically approximated via the optimal σ in the LMIs together with matrix B_d and the region of interest described above. Besides, since the FDI filter is a stable linear time invariant transfer function with normalized coefficients, the Lipschitz constant of the second mapping $E(X) \mapsto Q_d$ can be explicitly computed based on the filter denominator which is fixed prior to the design procedure; see [MEL13, Lemma 4.5].

Conclusions

CHAPTER 6

Conclusions and Future Directions

In this thesis we studied two problems: first, the motion planning of controlled diffusions, and second the problem of fault detection for large scale nonlinear systems. The motion planning problem is an extension to the reachability problem which is well known in the context of safety problems in the dynamics and control literature, and the fault detection problem is a crucial concept in the design of reliable control systems. These problems were addressed separately in Part I and II, respectively.

6.1 Part I: Stochastic Motion Planning for Diffusions

6.1.1 Chapter 2: Stochastic Reach-Avoid Problem

As a first step toward the main objective of this part, in Chapter 2 we studied a class of stochastic reach-avoid problems with state constraints in the context of SDEs. We proposed a framework to characterize the set of initial conditions based on discontinuous viscosity solutions of a second order PDE. In contrast to earlier approaches, this methodology is not restricted to almost-sure notions and one can compute the desired set with any non-zero probability by means of off-the-shelf PDE solvers.

6.1.2 Chapter 3: Stochastic Motion Planning

In Chapter 3, continuing our studies of Chapter 2, we extended the class of reachability maneuvers as well as the stochastic process dynamics to different notions of stochastic motion planning problems which involve a controlled process, with possibly discontinuous sample paths, visiting certain subsets of the state-space while avoiding others in a sequential fashion. We established a weak DPP comprising auxiliary value functions defined in terms of discontinuous payoff functions. Subsequently, we focused on a case of diffusions as the solution of a controlled SDE, and investigated the required conditions to apply the proposed DPP. It turned out that the proposed DPP leads to a sequence of PDEs, for which the first one has a known boundary condition, while the boundary conditions of the subsequent ones are determined by the solutions to the preceding steps. Finally, the performance of the proposed stochastic motion planning notions was illustrated for a biological switch network.

For potential directions to pursue research toward the motion planning problem, one may consider the following questions:

(i) Note that full-state measurement is one of the standing assumptions in the problems discussed in Part I. Motivated by the fact that in some applications only partial and possibly imperfect measurements may be available, a natural question is to address the motion planning objective under measurement constraints, i.e., an admissible control policy would be only allowed to utilize the information of the process $Y_s := h(X_s)$ where $h : \mathbb{R}^d \to \mathbb{R}^{d_y}$ is a given measurable mapping.

(ii) In light of Proposition 2.3.5 (resp. Proposition 2.3.7), we know that all the classes of stochastic optimal control problems in (2.4) (resp. (2.5)) have a close connection to the reach-avoid problem. Chapter 2 only studied the exit-time formulation which was also exploited in Chapter 3 to address the motion planning objectives. As an alternative step, however, it would be interesting to study the other formulations that possible may shed light on other features of the problem. In particular that the other value functions basically reflect a dynamic game between two players with different authorities, e.g., a stopper versus a controller in (2.5c).

(iii) Theorem 3.4.3 holds for the broad class of stochastic processes whose sample paths are right continuous with left limits. Therefore, as a step toward the generalization of the process dynamics, a potential research would be to investigate the required conditions of the proposed DPP (Assumptions 3.4.1) for a larger class of stochastic processes, e.g., controlled Markov jump-diffusions.

6.2 Part II: Fault Detection for Large Scale Nonlinear Systems

6.2.1 Chapter 4: A Tractable Approach with Probabilistic Performance Index

In Chapter 4 we proposed a novel perspective toward the FDI filter design along with a tractable optimization-based methodology. Previous approaches on FDI problems are either confined to linear systems or they are only applicable to low dimensional dynamics with specific structures. In contrast, thanks to the convex formulation, the methodology is applicable to large scale nonlinear systems in which some statistical information of exogenous disturbances are available. From a technical viewpoint, the crucial step in the proposed approach is based on robustification of the FDI filter to finitely many signatures of the dynamics nonlinearity. Motivated by our earlier works, we deployed the proposed technique to design a diagnosis filter to detect the AGC malfunction in two-area power network. The simulation results validated the filter performance, particularly that the filter was encountered the disturbance patterns different than the training ones. That is, the test disturbances only shared the same statistical properties with the training disturbances.

For further research on the FDI problem discussed in Chapter 4, one may look into the following directions:

(i) The central focus of the work in Chapter 4 is to robustify the filter to certain signatures of dynamic nonlinearities in the presence of given disturbance patterns. As a next step, motivated by applications that disruptive attacks may follow certain patterns, a natural question is whether the filter can be trained for these attack patterns. From the technical standpoint, this problem in principle may be different from the robustification phase since the former may involve maximization of the residual norm as opposed to the minimization for the robustification discussed in this chapter. Therefore, this problem may require a reconciliation between the disturbance rejection and the fault sensitivity objectives.

(ii) The methodology studied in this chapter is applicable to both discrete and continuous-time dynamics and measurements. In reality, however, we often have different time-setting in different parts, i.e., we only have discrete-time measurements while the system dynamics follows a continuous-time behavior. We believe this setup introduces new challenges to the problem. We recently reported heuristic attempts toward this objective in [ETMEL13], though there is still a need to address the problem in a rigorous and systematic framework.

(iii) Another standing assumption throughout Chapter 4 is the accessibility of perfect measurements, i.e., the system dynamics in Figure 4.1 is deterministic and the output signal y is noiseless. Another question to answer is how the model stochasticity as well as measurements noise can be incorporated into the FDI design. We reported some preliminaries in [SMEKL13] to address noisy measurements but effectively for linear systems where the dynamic nonlinearities can be treated as auxiliary disturbances.

6.2.2 Chapter 5: Performance Bound for Random Programs

One of the motivation of the study in Chapter 5 is to quantify the behavior of the false alarm rate of the proposed FDI in Chapter 4 with respect to the threshold level obtained via the corresponding random programs. Chapter 5 presented probabilistic performance bounds for both robust and chance constrained programs via the so-called scenario program (SCP). This result is, to the best of our knowledge, the first confidence bounds for the objective performance of RCPs and CCPs based on scenario programs. Subsequently, we extended our results to a certain class of non-convex programs which, in particular, allows for binary decision variables with linear growth rate for the required number of samples.

For potential research directions on the content of this chapter, one may think of the following possibilities:

(i) In light of Theorem 5.3.6, one may look more closely into the derivation of ULBs as introduced in Definition 5.3.1. Meaningful ULBs may depend highly on the individual structure of the optimization problems, in particular the uncertainty set and the constraint functions. In fact, the probability measure $\mathbb{P}$ can also be viewed as a decision variable to propose the most "*informative*" scenario program to tackle the robust formulation RCP.

(ii) Another potential direction is the estimation of the constant $L_{\mathbf{SP}}$ in Theorems 5.3.6 and 5.3.7. This problem may be closely related to the estimation of the dual optimizers of RCPs, and enables better theoretical bounds particularly for a posteriori assessments (cf. Theorem 5.3.7).

(iii) The theoretical results developed in Chapter 5 essentially establishes a theoretical bridge between a semi-infinite program (RCP) and a finite counterpart (SCP). Motivating by different applications involving an infinite optimization programs (e.g., approximate dynamic programming) a potential extension would be to complete this bridge up to a class of full-infinite programs. This link, however, may need to resort to an asymptotic as well as probabilistic performance index.

Bibliography

[AD90] J.P. Aubin and G. Da Prato, *Stochastic viability and invariance*, Annali della Scuola Normale Superiore di Pisa. Classe di Scienze. Serie IV **17** (1990), no. 4, 595–613.

[Ada75] Robert A. Adams, *Sobolev spaces*, Academic Press [A subsidiary of Harcourt Brace Jovanovich, Publishers], New York-London, 1975, Pure and Applied Mathematics, Vol. 65.

[AEV$^+$12] Göran Andersson, Peyman Mohajerin Esfahani, Maria Vrakopoulou, Kostas Margellos, John Lygeros, André Teixeira, György Dán, Henrik Sandberg, and Karl H Johansson, *Cyber-security of scada systems*, Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES, IEEE, 2012, pp. 1–2.

[AF02] P. M. Anderson and A. A. Fouad, *Power System Control and Stability*, IEEE Computer Society Press, 2002.

[AGA09] Steven Andrews, Tuan Ginh, and Adam Arkin, *Stochastic models of biological processes*, Meyers RA, ed. Encyclopedia of Complexity and System Science (2009), 87308749.

[ALQ$^+$02] J.P. Aubin, J. Lygeros, M. Quincampoix, S.S. Sastry, and N. Seube, *Impulse differential inclusions: a viability approach to hybrid systems*, Institute of Electrical and Electronics Engineers. Transactions on Automatic Control **47** (2002), no. 1, 2–20.

[Anda] Göran Andersson, *Dynamics and control of electric power systems*, Power System Laboratory, ETH Zurich.

[Andb] ———, *Power system analysis*, Power System Laboratory, ETH Zurich.

[AP98] J.P. Aubin and G. Da Prato, *The viability theorem for stochastic differential inclusions*, Stochastic Analysis and Applications **16** (1998), no. 1, 1–15.

[APF00] J.P Aubin, G. Da Prato, and H. Frankowska, *Stochastic invariance for differential inclusions*, Set-Valued Analysis. An International Journal Devoted to the Theory of Multifunctions and its Applications **8** (2000), no. 1-2, 181–201.

[ATC09] Teodoro Alamo, Roberto Tempo, and Eduardo F. Camacho, *Randomized strategies for probabilistic solutions of uncertain feasibility and optimization problems*, IEEE Trans. Automat. Control **54** (2009), no. 11, 2545–2559. MR 2571919 (2010i:90137)

[Aub91] J.P. Aubin, *Viability Theory*, Systems & Control: Foundations & Applications, Birkhäuser Boston Inc., Boston, MA, 1991.

[Bea71] R. V. Beard, *Failure accommodation in linear systems through self-reorganization*, Ph.D. thesis, Massachusetts Inst. Technol., Cambridge, MA, 1971.

[Ber05] D. Bertsekas, *Dynamic programming and suboptimal control: A survey from {ADP} to mpc*, European Journal of Control **11** (2005), no. 45, 310 – 334.

[Ber09] Dimitri P. Bertsekas, *Convex optimization theory*, Athena Scientific, Nashua, NH, 2009. MR 2830150 (2012f:90001)

[BET10] Bruno Bouchard, Romuald Elie, and Nizar Touzi, *Stochastic target problems with controlled loss*, SIAM Journal on Control and Optimization **48** (2009/10), no. 5, 3123–3150. MR 2599913 (2011e:49039)

[BFP+11] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S.L. Smith, *Dynamic vehicle routing for robotic systems*, Proceedings of the IEEE **99** (2011), no. 9, 1482–1504.

[BFZ10] Olivier Bokanowski, Nicolas Forcadel, and Hasnaa Zidani, *Reachability and minimal times for state constrained nonlinear problems without any controllability assumption*, SIAM J. Control and Optimization **48(7)** (2010), 4292–4316.

[BG99] M. Bardi and P. Goatin, *Invariant sets for controlled degenerate diffusions: a viscosity solutions approach*, Stochastic analysis, control, optimization and applications, Systems Control Found. Appl., Birkhäuser Boston, Boston, MA, 1999, pp. 191–208.

[BGFB94] Stephen Boyd, Laurent El Ghaoul, Eric Feron, and Venkataramanan Balakrishnan, *Linear matrix inequalities in system and control theory*, Society for Industrial and Applied Mathematics: SIAM studies in applied mathematics, Society for Industrial and Applied Mathematics, 1994.

[BGK+83] Bernd Bank, Jürgen Guddat, Diethard Klatte, Bernd Kummer, and Klaus Tammer, *Nonlinear parametric optimization*, Birkhäuser Verlag, Basel, 1983.

[BHKH05] Christel Baier, Holger Hermanns, Joost-Pieter Katoen, and Boudewijn R. Haverkort, *Efficient computation of time-bounded reachability probabilities in uniform continuous-time markov decision processes*, Theor. Comput. Sci. **345** (2005), no. 1, 2–26.

[Bil95] Patrick Billingsley, *Probability and measure*, third ed., Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons Inc., New York, 1995, A Wiley-Interscience Publication. MR 1324786 (95k:60001)

[Bil99] ______, *Convergence of probability measures*, second ed., Wiley Series in Probability and Statistics: Probability and Statistics, John Wiley & Sons Inc., New York, 1999, A Wiley-Interscience Publication. MR 1700749 (2000e:60008)

[BJ02] M. Bardi and R. Jensen, *A geometric characterization of viable sets for controlled degenerate diffusions*, Set-Valued Analysis **10** (2002), no. 2-3, 129–141, Calculus of variations, nonsmooth analysis and related topics.

[BMGA02] R.W. Beard, T.W. McLain, M.A. Goodrich, and E.P. Anderson, *Coordinated target assignment and intercept for unmanned air vehicles*, Robotics and Automation, IEEE Transactions on **18** (2002), no. 6, 911–922.

[BN12] Bruno Bouchard and Marcel Nutz, *Weak dynamic programming for generalized state constraints*, SIAM Journal on Control and Optimization **50** (2012), no. 6, 3344–3373. MR 3024163

[Bor05] V. S. Borkar, *Controlled diffusion processes*, Probability Surveys **2** (2005), 213–244 (electronic).

[BPQR98] R. Buckdahn, Sh. Peng, M. Quincampoix, and C. Rainere, *Existence of stochastic control under state constraints*, Comptes Rendus de l'Académie des Sciences. Série I. Mathématique **327** (1998), no. 1, 17–22.

[BS06] Dimitris Bertsimas and Melvyn Sim, *Tractable approximations to robust conic optimization problems*, Mathematical Programming **107** (2006), no. 1-2, 5–36.

[BSS01] A. Becskei, B. Seraphin, and L. Serrano, *Positive feedback in eukaryotic gene networks: cell differentiation by graded to binary response conversion*, EMBO J **20** (2001), no. 10, 2528–2535.

[BT11] B. Bouchard and N. Touzi, *Weak dynamic programming principle for viscosity solutions*, SIAM Journal on Control and Optimization **49** (2011), no. 3, 948–962.

[BtN98] Aharon Ben-tal and Arkadi Nemirovski, *Robust convex optimization*, Mathematics of Operations Research **23** (1998), 769–805.

[BtN99] ______, *Robust solutions of uncertain linear programs*, Operations Research Letters **25** (1999), 1–13.

[BtNR01] Aharon Ben-tal, Arkadi Nemirovski, and C. Roos, *Robust solutions of uncertain quadratic and conic-quadratic problems*, Solutions of Uncertain Linear Programs: Math. Program, Kluwer, 2001, pp. 351–376.

[BV04] Stephen Boyd and Lieven Vandenberghe, *Convex optimization*, Cambridge University Press, New York, NY, USA, 2004.

[Cal09] Giuseppe C. Calafiore, *A note on the expected probability of constraint violation in sampled convex programs*, 18th IEEE International Conference on Control Applications Part of 2009 IEEE Multi-conference on Systems and Control, july 2009, pp. 1788 – 1791.

[Cal10] Giuseppe Carlo Calafiore, *Random convex programs*, SIAM J. Optim. **20** (2010), no. 6, 3427–3464. MR 2763511 (2012c:90094)

[Car96] P. Cardaliaguet, *A differential game with two players and one target*, SIAM Journal on Control and Optimization **34** (1996), no. 4, 1441–1460.

[CC05] Giuseppe Calafiore and M. C. Campi, *Uncertain convex programs: randomized solutions and confidence levels*, Math. Program. **102** (2005), no. 1, Ser. A, 25–46. MR 2115479 (2005k:90104)

[CC06] Giuseppe C. Calafiore and Marco C. Campi, *The scenario approach to robust control design*, IEEE Trans. Automat. Control **51** (2006), no. 5, 742–753. MR 2232597 (2007a:93075)

[CCL11] Debasish Chatterjee, Eugenio Cinquemani, and John Lygeros, *Maximizing the probability of attaining a target prior to extinction*, Nonlinear Analysis: Hybrid Systems (2011), http://dx.doi.org/10.1016/j.nahs.2010.12.003.

[CES86] E. M. Clarke, E. A. Emerson, and A. P. Sistla, *Automatic verification of finite-state concurrent systems using temporal logic specifications*, ACM Transactions on Programming Languages and Systems (TOPLAS) **8** (1986), no. 2, 244–263.

[CG08] M. C. Campi and S. Garatti, *The exact feasibility of randomized solutions of uncertain convex programs*, SIAM J. Optim. **19** (2008), no. 3, 1211–1230. MR 2460739 (2009j:90081)

[CG11] ———, *A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality*, J. Optim. Theory Appl. **148** (2011), no. 2, 257–280. MR 2780563 (2012b:90118)

[CGP09] Marco C. Campi, Simone Garatti, and Maria Prandini, *The scenario approach for systems and control design*, Annual Reviews in Control (2009), no. 2, 149–157.

[Che00] J. Cherry, *How to make a Biological Switch*, Journal of Theoretical Biology **203** (2000), no. 2, 117–133.

[CIL92] M. G. Crandall, H. Ishii, and P. L. Lions, *User's guide to viscosity solutions of second order partial differential equations*, American Mathematical Society **27** (1992), 1–67.

[CLF12] Giuseppe C. Calafiore, D. Lyons, and L. Fagiano, *On mixed-integer random convex programs*, Decision and Control (CDC), 2012 IEEE 51st Annual Conference on, 2012, pp. 3508–3513.

[CP82] J. Chen and R. Patton, *Robust model based faults diagnosis for dynamic systems*, Dordrecht: Kluwer Academic Publishers, New York, 1982.

[CQSP97] P. Cardaliaguet, M. Quincampoix, and P. Saint-Pierre, *Optimal times for constrained nonlinear control problems without local controllability*, Applied Mathematics and Optimization **36** (1997), no. 1, 21–42.

[CQSP02] ———, *Differential Games with State-Constraints*, ISDG2002, Vol. I, II (St. Petersburg), St. Petersburg State Univ. Inst. Chem., St. Petersburg, 2002, pp. 179–182.

[CRVRL06] Randy Cogill, Michael Rotkowitz, Benjamin Van Roy, and Sanjay Lall, *An approximate dynamic programming approach to decentralized control of stochastic systems*, Control of Uncertain Systems: Modelling, Approximation, and Design, Lecture Notes in Control and Information Science, Springer, 2006, pp. 243–256.

[CS98a] Y. Chitour and H. J. Sussmann, *Line-integral estimates and motion planning using the continuation method*, Essays on Mathematical Robotics, proceedings of the February 1993 workshop on Robotics held at the Institute for Mathematics and its Applications (IMA) at the University of Minnesota, Springer-Verlag, New York, 1998, pp. 91–126.

[CS98b] Walter H. Chung and Jason L. Speyer, *A game-theoretic fault detection filter*, IEEE Trans. Automat. Control **43** (1998), no. 2, 143–161.

[DF01] G. Da Prato and H. Frankowska, *Stochastic viability for compact sets in terms of the distance function*, Dynamic Systems and Applications **10** (2001), no. 2, 177–184.

[DF04] ———, *Invariance of stochastic control systems with deterministic arguments*, Journal of Differential Equations **200** (2004), no. 1, 18–52.

[dFVR03] D. P. de Farias and B. Van Roy, *The linear programming approach to approximate dynamic programming*, Operations Research **51** (2003), no. 6, 850–865.

[DFVR04] Daniela Pucci De Farias and Benjamin Van Roy, *On constraint sampling in the linear programming approach to approximate dynamic programming*, Mathematics of operations research **29** (2004), no. 3, 462–478.

[DPR13] François Dufour and Tomás Prieto-Rumeau, *Finite linear programming approximations of constrained discounted Markov decision processes*, SIAM J. Control Optim. **51** (2013), no. 2, 1298–1324. MR 3036990

[Dud99] R. M. Dudley, *Uniform central limit theorems*, Cambridge Studies in Advanced Mathematics, vol. 63, Cambridge University Press, Cambridge, 1999. MR 1720712 (2000k:60050)

[Dug66] J. Dugundji, *Topolgy*, Boston: Allyn and Bacon, US, 1966.

[EK86] S.N. Ethier and T.G. Kurtz, *Markov processes: Characterization and convergence*, Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons, Ltd., New York, 1986.

[ESKPG05] Hana El Samad, Mustafa Khammash, Linda Petzold, and Dan Gillespie, *Stochastic modelling of gene regulatory networks*, Internat. J. Robust Nonlinear Control **15** (2005), no. 15, 691–711.

[ETMEL13] Erasmia Evangelia Tiniou, Peyman Mohajerin Esfahani, and John Lygeros, *Fault detection with discrete-time measurements: an application for the cyber security of power networks*, 52th IEEE Conference Decision and Control, Dec 2013.

[FF12] Giuseppe Franzè and Domenico Famularo, *A robust fault detection filter for polynomial nonlinear systems via sum-of-squares decompositions*, Systems & Control Letters **61** (2012), no. 8, 839–848.

[FG99] I. J. Fialho and T. Georgiou, *Worst case analysis of nonlinear systems*, IEEE transactions on Automatic Control **44(6)** (1999), 4292–4316.

[FKA09] Erik Frisk, Mattias Krysander, and Jan Aslund, *Sensor placement for fault isolation in linear differential-algebraic systems*, Automatica **45** (2009), no. 6, 364–371.

[FS06] W.H. Fleming and H.M. Soner, *Controlled Markov Processes and Viscosity Solution*, 3 ed., Springer-Verlag, 2006.

[Gah96] Pascal Gahinet, *Explicit controller formulas for lmi-based h-infinity synthesis*, Automatica **32** (1996), no. 7, 1007–1014.

[GCC00] T. S. Gardner, C. R. Cantor, and J. J. Collins, *Construction of a genetic toggle switch in Escherichia coli*, Nature **403** (2000), no. 6767, 339–42.

[GKM10] C. Goerzen, Z. Kong, and B. Mettler, *A survey of motion planning algorithms from the perspective of autonomous uav guidance*, Journal of Intelligent and Robotic Systems **57** (2010), no. 1-4, 65–100 (English).

[GLQ06] Y. Gao, J. Lygeros, and M. Quincampoix, *The Reachability Problem for Uncertain Hybrid Systems Revisited: a Viability Theory Perspective*, Hybrid systems: computation and control, Lecture Notes in Comput. Sci., vol. 3927, Springer, Berlin, 2006, pp. 242–256.

[GOL98] Laurent El Ghaoui, Francois Oustry, and Herv Lebret, *Robust solutions to uncertain semidefinite programs*, SIAM Journal on Optimization **9** (1998), no. 1, 33–52.

[Han12] Lars Peter Hansen, *Proofs for large sample properties of generalized method of moments estimators*, J. Econometrics **170** (2012), no. 2, 325–330. MR 2970318

[HKEY99] H. Hammouri, M. Kinnaert, and E.H. El Yaagoubi, *Observer-based approach to fault detection and isolation for nonlinear systems*, Automatic Control, IEEE Transactions on **44** (1999), no. 10, 1879 –1884.

[HMEKL14] Flavio Heer, Peyman Mohajerin Esfahani, Maryam Kamgarpour, and John Lygeros, *Model based power optimisation of wind farms*, European Control Conference (ECC), Jun 2014.

[HP96] M. Hou and R.J. Patton, *An lmi approach to H_-/H_∞ fault detection observers*, Control '96, UKACC International Conference on (Conf. Publ. No. 427), vol. 1, sept. 1996, pp. 305 – 310 vol.1.

[Jon73] H. L. Jones, *Failure detection in linear systems*, Ph.D. thesis, Massachusetts Inst. Technol., Cambridge, MA, 1973.

[Kal97] O. Kallenberg, *Foundations of Modern Probability*, Probability and its Applications (New York), Springer-Verlag, New York, 1997.

[Kha] Mustafa Khammash, *Modeling and analysis of stochastic biochemical networks*, Control Theory and System Biology.

[Kha92] Hassan K. Khalil, *Nonlinear systems*, Macmillan Publishing Company, New York, 1992. MR 1201326 (93k:34001)

[KO10] B. N. Khoromskij and I. Oseledets, *Quantics-TT collocation approximation of parameter-dependent and stochastic elliptic PDEs*, Computational Methods in Applied Mathematics **10** (2010), no. 4, 376–394. MR 2770302 (2012c:65020)

[Kry09] N.V. Krylov, *Controlled Diffusion Processes*, Stochastic Modelling and Applied Probability, vol. 14, Springer-Verlag, Berlin Heidelberg, 2009, Reprint of the 1980 Edition.

[KS91] I. Karatzas and S.E. Shreve, *Brownian Motion and Stochastic Calculus*, 2 ed., Graduate Texts in Mathematics, vol. 113, Springer-Verlag, New York, 1991.

[KS11] Boris N. Khoromskij and Christoph Schwab, *Tensor-structured Galerkin approximation of parametric and stochastic elliptic PDEs*, SIAM Journal on Scientific Computing **33** (2011), no. 1, 364–385. MR 2783199 (2012e:65273)

[KT03] Vijay R. Konda and John N. Tsitsiklis, *On actor-critic algorithms*, SIAM Journal on Control and Optimization **42** (2003), no. 4, 1143–1166 (electronic). MR 2044789 (2004m:93151)

[KT12] Takafumi Kanamori and Akiko Takeda, *Worst-case violation of sampled convex programs for optimization with uncertainty*, Journal of Optimization Theory and Applications **152** (2012), no. 1, 171–197. MR 2872517

[Lof04] J. Lofberg, *Yalmip : a toolbox for modeling and optimization in matlab*, Computer Aided Control Systems Design, 2004 IEEE International Symposium on, sept. 2004, pp. 284 –289.

[LTS00] J. Lygeros, C. Tomlin, and S.S. Sastry, *A game theorretic approach to controller design for hybrid systems*, Proceedings of IEEE **88** (2000), no. 7, 949–969.

[Lue69] David G. Luenberger, *Optimization by vector space methods*, John Wiley & Sons Inc., New York, 1969.

[LVLM08] Andrea Lecchini-Visintini, John Lygeros, and Jan M. Maciejowski, *Approximate domain optimization for deterministic and expected value criteria*, Tech. report, March 2008, [Online]. Available: http://control.ee.ethz.ch/index.cgi?page=publications;action=details;id=3048.

[LVLM10] ______, *Stochastic optimization on continuous domains with finite-time guarantees by Markov chain Monte Carlo methods*, IEEE Trans. Automat. Control **55** (2010), no. 12, 2858–2863. MR 2767154 (2012b:90121)

[Lyg04] J. Lygeros, *On reachability and minimum cost optimal control*, Automatica. A Journal of IFAC, the International Federation of Automatic Control **40** (2004), no. 6, 917–927 (2005).

[MA04] R Timothy Marler and Jasbir S Arora, *Survey of multi-objective optimization methods for engineering*, Structural and multidisciplinary optimization **26** (2004), no. 6, 369–395.

[Mas86] Mohammad-Ali Massoumnia, *A geometric approach to the synthesis of failure detection filters*, IEEE Trans. Automat. Control **31** (1986), no. 9, 839–846.

[MB00] Timothy W. Mclain and Randal W. Beard, *Trajectory planning for coordinated rendezvous of unmanned air vehicles*, Proc. GNC'2000, 2000, pp. 1247–1254.

[MECL11] Peyman Mohajerin Esfahani, Debasish Chatterjee, and John Lygeros, *On a problem of stochastic reach-avoid set characterization for diffusions*, IEEE Conference on Decision and Control, December 2011, pp. 7069 –7074.

[MECL12] ______, *Motion planning via optimal control for stochastic processes*, submitted to IEEE Transactions on Automatic Control (2012), [Online]. Available: http://arxiv.org/abs/1211.1138.

[MECL13] ______, *On a stochastic reach-avoid problem and set characterization*, submitted for publication (2013), [Online]. Available: http://arxiv.org/abs/1202.4375.

[MEL13] Peyman Mohajerin Esfahani and John Lygeros, *A tractable fault detection and isolation approach for nonlinear systems with probabilistic performance*, submitted to IEEE Transactions on Automatic Control (2013), [Online]. Available: http://control.ee.ethz.ch/index.cgi?page=publications&action=details&id=4344.

[MEMAC13] Peyman Mohajerin Esfahani, Andreas Milias-Argeitis, and Debasish Chatterjee, *Analysis of controlled biological switches via stochastic motion planning*, European Control Conference, July 2013.

[MESL13] Peyman Mohajerin Esfahani, Tobias Sutter, and John Lygeros, *Performance bounds for the scenario approach and an extension to a class of non-convex programs*, to appear in as a full paper in IEEE Transactions on Automatic Control (2013), [Online]. Available: http://arxiv.org/abs/1307.0345.

[MEVAL] Peyman Mohajerin Esfahani, Maria Vrakopoulou, Goran Andersson, and John Lygeros, *Intrusion detection in electric power networks*, Patent applied for EP-12005375, filed 24 July 2012.

[MEVAL12] ______, *A tractable nonlinear fault detection and isolation technique with application to the cyber-physical security of power systems*, 51th IEEE Conference Decision and Control, 2012, [Online]. Full version: http://control.ee.ethz.ch/index.cgi?page=publications;action=details;id=4196.

[MEVM+10] Peyman Mohajerin Esfahani, Maria Vrakopoulou, Kostas Margellos, John Lygeros, and Goran Andersson, *Cyber attack in a two-area power system: Impact identification using reachability*, American Control Conference, 2010, pp. 962 – 967.

[MEVM+11] ______, *A robust policy for automatic generation control cyber attack in two area power network*, 49th IEEE Conference Decision and Control, 2011, pp. 5973 – 5978.

[Mit05] I. Mitchell, *A toolbox of hamilton-jacobi solvers for analysis of nondeterministic continuous and hybrid systems*, Hybrid systems: computation and control (M. Morari and L. Thiele, eds.), Lecture Notes in Comput. Sci., no. 3414, Springer-Verlag, 2005, pp. 480–494.

[ML11] Kostas Margellos and John Lygeros, *Hamilton-Jacobi formulation for reach-avoid differential games*, IEEE Trans. Automat. Control **56** (2011), no. 8, 1849–1861.

[MS90] R.M. Murray and S.S. Sastry, *Steering nonholonomic systems using sinusoids*, Decision and Control, 1990., Proceedings of the 29th IEEE Conference on, 1990, pp. 2097–2101 vol.4.

[MSSO+11] A. Milias, S. Summers, J. Stewart-Ornstein, I. Zuleta, D. Pincus, H. El-Samad, M. Khammash, and J. Lygeros, *In silico feedback for in vivo regulation of a gene expression circuit*, Nature Biotechnology **29** (2011), no. 12, 11141116.

[MT02] I. Mitchell and C. J. Tomlin, *Level set methods for computation in hybrid systems*, Hybrid systems: computation and control, Lecture Notes in Comput. Sci., vol. 1790, Springer-Verlag, New York, 2002, pp. 310–323.

[MVW89] Mohammad-Ali Massoumnia, G. C. Verghese, and A. S. Willsky, *Failure detection and identification*, IEEE Transaction on Automatic Control **34** (1989), no. 3, 316–321.

[NF06] Mattias Nyberg and Erik Frisk, *Residual generation for fault diagnosis of system described by linear differential-algebraic equations*, IEEE Transaction on Automatic Control **51** (2006), no. 12, 1995–2000.

[NO08] Angelia Nedić and Asuman Ozdaglar, *Approximate primal solutions and rate analysis for dual subgradient methods*, SIAM J. Optim. **19** (2008), no. 4, 1757–1780. MR 2486049 (2010f:90123)

[Øks03] Bernt Øksendal, *Stochastic differential equations*, sixth ed., Universitext, Springer-Verlag, Berlin, 2003, An introduction with applications. MR 2001996 (2004e:60102)

[OS88] S. Osher and J.A. Sethian, *Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations*, Journal of Computational Physics **79** (1988), no. 1, 12–49.

[OSME$^+$13] Frauke Oldewurtel, David Sturzenegger, Peyman Mohajerin Esfahani, Goran Andersson, Manfred Morari, and John Lygeros, *Adaptively constrained stochastic model predictive control for closed-loop constraint satisfaction*, American Control Conference (ACC), 2013, IEEE, 2013, pp. 4674–4681.

[PI01] Claudio De Persis and Alberto Isidori, *A geometric approach to nonlinear fault detection and isolation*, IEEE Trans. Automat. Control **46** (2001), no. 6, 853–865.

[PRC12] B. K. Pagnoncelli, D. Reich, and M. C. Campi, *Risk-return trade-off with the scenario approach in practice: a case study in portfolio selection*, J. Optim. Theory Appl. **155** (2012), no. 2, 707–722. MR 2994071

[Pré95] András Prékopa, *Stochastic programming*, Mathematics and Its Applications, Springer, 1995.

[Pro05] Philip E. Protter, *Stochastic integration and differential equations*, Stochastic Modelling and Applied Probability, vol. 21, Springer-Verlag, Berlin, 2005, Second edition. Version 2.1, Corrected third printing.

[PW98] Jan Willem Polderman and Jan C. Willems, *Introduction to mathematical systems theory*, Texts in Applied Mathematics, vol. 26, Springer-Verlag, New York, 1998, A behavioral approach.

[RB98] Richard Bass, *Diffusions and elliptic operators*, Probability and its Applications (New York), Springer-Verlag, New York, 1998.

[ref] *Power systems test case archive, college of engineering,university of washington*, URL: http://www.ee.washington.edu/research/pstca/.

[Ren99] P. J. Reny, *On the existence of pure and mixed strategy nash equilibria in discontinuous games*, Econometrica **67** (1999), 1029–1056.

[Roc97] R. Tyrrell Rockafellar, *Convex analysis*, Princeton Landmarks in Mathematics and Physics Series, PRINCETON University Press, 1997.

[Rud87] Walter Rudin, *Real and complex analysis*, Mathematics series, McGraw-Hill, 1987.

[RW10] R. Tyrrell Rockafellar and Roger J.B. Wets, *Variational analysis*, Grundlehren der mathematischen Wissenschaften, Springer, 2010.

[SDR09] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński, *Lectures on stochastic programming: Modeling and theory*, MPS-SIAM series on optimization, Society for Industrial and Applied Mathematics, 2009.

[Set99] J.A. Sethian, *Level set methods and fast marching methods*, second ed., Cambridge Monographs on Applied and Computational Mathematics, vol. 3, Cambridge University Press, Cambridge, 1999, Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science.

[SF91] R. Seliger and P.M. Frank, *Fault diagnosis by disturbance-decoupled nonlinear observers*, Proceedings of the 30th IEEE Conference on Decision and Control, 1991, pp. 2248–2253.

[SH07] Abhyudai Singh and JooPedro Hespanha, *A derivative matching approach to moment closure for the stochastic logistic model*, Bulletin of Mathematical Biology **69** (2007), no. 6, 1909–1925 (English).

[Sha03] Jun Shao, *Mathematical statistics*, second ed., Springer Texts in Statistics, Springer-Verlag, New York, 2003. MR 2002723 (2004g:62002)

[Shc07] A. A. Shcheglova, *Nonlinear differential-algebraic systems*, Sibirsk. Mat. Zh. **48** (2007), no. 4, 931–948. MR 2355385 (2009c:34002)

[SL10] Sean Summers and John Lygeros, *Verification of discrete time stochastic hybrid systems: a stochastic reach-avoid decision problem*, Automatica J. IFAC **46** (2010), no. 12, 1951–1961. MR 2878218

[SM86] Kang G. Shin and Neil D. McKay, *A dynamic programming approach to trajectory planning of robotic manipulators*, Automatic Control, IEEE Transactions on **31** (1986), no. 6, 491–500.

[SMEKL13] Bratislav Svetozarevic, Peyman Mohajerin Esfahani, Maryam Kamgarpour, and John Lygeros, *A robust fault detection and isolation filter for a horizontal axis variable speed wind turbine*, American Control Conference, Jun 2013.

[SMESL14a] David Sutter, Peyman Mohajerin Esfahani, Tobias Sutter, and John Lygeros, *Efficient approximation of discrete memoryless channel capacities*, IEEE International Symposium on Information Theory (ISIT), Jul 2014, [Online]. Available: http://www.phys.ethz.ch/~suttedav/123/123.pdf.

[SMESL14b] Tobias Sutter, Peyman Mohajerin Esfahani, David Sutter, and John Lygeros, *Capacity approximation of memoryless channels with countable output alphabets*, IEEE International Symposium on Information Theory (ISIT), Jul 2014, [Online]. Available: http://www.phys.ethz.ch/~suttedav/789/789.pdf.

[ST02a] H Mete Soner and Nizar Touzi, *Stochastic target problems, dynamic programming, and viscosity solutions*, SIAM Journal on Control and Optimization **41** (2002), no. 2, 404–424.

[ST02b] H.M. Soner and N. Touzi, *Dynamic programming for stochastic target problems and geometric flows*, Journal of the European Mathematical Society (JEMS) **4** (2002), no. 3, 201–236.

[Sus91] H. J. Sussmann, *Two new methods for motion planning for controllable systems without drift*, Proceedings of the First European Control Conference (ECC91) (Grenoble, France), 1991, pp. 1501–1506.

[Tou13] Nizar Touzi, *Optimal stochastic control, stochastic target problems, and backward SDE*, Fields Institute Monographs, vol. 29, Springer, New York, 2013, With Chapter 13 by Angès Tourin. MR 2976505

[VL10] Kyriakos G. Vamvoudakis and Frank L. Lewis, *Online actor-critic algorithm to solve the continuous-time infinite horizon optimal control problem*, Automatica. A Journal of IFAC, the International Federation of Automatic Control **46** (2010), no. 5, 878–888. MR 2877161 (2012m:93183)

[VMEG+] Robin Vujanic, Peyman Mohajerin Esfahani, Paul Goulart, Sèbastian Marièthoz, and Manfred Morari, *Vanishing duality gap in large scale mixed-integer optimization: a solution method with power system applications*, submitted to Journal of Mathematical Programming.

[VMEM+] Maria Vrakopoulou, Peyman Mohajerin Esfahani, Kostas Margellos, John Lygeros, and Goran Andersson, *Cyber-attacks in the automatic generation control*, Cyber Physical Systems Approach to Smart Electric Power Grid.

[Wil06] Darren James Wilkinson, *Stochastic modelling for systems biology*, Chapman & Hall/CRC, Boca Raton, FL, 2006. MR 2222876 (2006k:92040)

[WMEL12] Tony Wood, Peyman Mohajerin Esfahani, and John Lygeros, *Hybrid Modelling and Reachability on Autonomous RC-Cars*, IFAC Conference on Analysis and Design of Hybrid Systems (ADHS) (Eindhoven, Netherlands), June 2012, [Videos]. https://sites.google.com/site/orcaracer/videos.

[ZD97] Kemin Zhou and John C. Doyle, *Essentials of robust control*, Prentice Hall, September 1997.

[ZMEAL13] Majid Zamani, Peyman Mohajerin Esfahani, Alessandro Abate, and John Lygeros, *Symbolic models for stochastic control systems without stability assumptions*, European Control Conference, July 2013.

[ZMEM+13a] Majid Zamani, Peyman Mohajerin Esfahani, Rupak Majumdar, Alessandro Abate, and John Lygeros, *Bisimilar finite abstractions of stochastic control systems*, 52th IEEE Conference Decision and Control, December 2013.

[ZMEM+13b] ______, *Symbolic models for stochastic control systems*, to appear as a full paper in IEEE Transactions on Automatic Control (2013), [Online]. Available: http://arxiv.org/abs/1302.3868.

Curriculum Vitae

Peyman Mohajerin Esfahani
born on March 21th, 1982 in Tehran, Iran

Jun 2008 – Jan 2014	**ETH Zürich**, Switzerland Doctorate at the Automatic Control Laboratory Department of Information Technology and Electrical Engineering
Sep 2005 – Dec 2007	**Sharif University of Technology**, Tehran, Iran Master of Science, major in Control Engineering Department of Electrical Engineering
Sep 2000 – Sep 2005	**Sharif University of Technology**, Tehran, Iran Bachelor of Science, major in Control Engineering Department of Electrical Engineering
Sep 1997 – Sep 2000	**Sadr High School**, Tehran, Iran