



Delft University of Technology

Document Version

Final published version

Licence

CC BY-NC-ND

Citation (APA)

Solanki, P. (2026). *Safe operation region estimation and certification*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:a34beb10-58f0-459e-b718-55f0a5bae6a1>

Important note

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

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

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

Takedown policy

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

This work is downloaded from Delft University of Technology.

Safe operation region estimation and certification

Safe operation region estimation and certification

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus, prof. dr. ir. H. Bijl
Chair of the Board for Doctorates
to be defended publicly on
18, May 2026 at 15:00

by

Prashant SOLANKI

This dissertation has been approved by the (co)promoters.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Dr. ir. C.C. de Visser,	Delft University of Technology, <i>promotor</i>
Dr. A. Jamshidnejad,	Delft University of Technology, <i>copromotor</i>
<i>Independent members:</i>	
Prof. dr. ir. J.W. van Wingerden,	Delft University of Technology, Netherlands
Prof. dr. P. Mohajerin Esfahani,	Delft University of Technology, Netherlands / University of Toronto, Canada
Prof. dr. A. Abate,	Oxford University, U.K.
Prof. dr. B. Jayawardhana,	University of Groningen, Netherlands
<i>Reserve member:</i>	
Prof. dr. G.C.H.E. de Croon,	Delft University of Technology, Netherlands



Keywords: Reachability analysis, Optimal control, Safe envelope

Copyright © 2026 by P. Solanki

ISBN 978-94-6518-322-0

An electronic copy of this dissertation is available at
<https://repository.tudelft.nl/>.

To my parents, wife and brother who supported me until now and forever beyond.

CONTENTS

Summary	xi
1 Introduction	1
1.1 The Need for Safe and Efficient Autonomous Systems	2
1.1.1 The Growing Role of Automation	2
1.1.2 Consequences of Unsafe Automation	2
1.2 Challenges in Ensuring Safe Automation	5
1.3 Existing approaches and their limitations	9
1.4 Research objectives	18
1.5 Scope of the Research	20
1.6 Methodology	20
1.7 Guide to the reader	22
2 Background	33
2.1 Hamilton-Jacobi Reachability: Foundations and Applications	34
2.1.1 Optimal Control and the Hamilton-Jacobi-Bellman Equation .	34
2.1.2 Robust Optimal Control and Hamilton-Jacobi-Isaacs Equations	35
2.1.3 Hamilton-Jacobi Reachability as a Robust Optimal Control Problem	36
2.1.4 Reachability and the HJI PDE	37
2.1.5 Variations of Reachability	37
2.1.6 Formal Definitions of Reachable Sets and HJI Formulation . .	40
2.1.7 Example: Double Integrator and Backward Reachable Tube (BRT)	42
2.2 Background on Neural Networks	44
2.2.1 Feedforward Neural Networks	44
2.2.2 Universal Approximation Theorem	45
2.2.3 Training Neural Networks	45
2.2.4 Relevance to Control and Robotics	46
2.3 Mathematical Foundations of Satisfiability Modulo Theories (SMT) .	46
2.3.1 Preliminaries and Formal Definitions	46
2.3.2 Common Background Theories	47
2.3.3 The DPLL(T) Framework	47
2.3.4 Model-Theoretic Semantics	48
2.3.5 Applications in Formal Methods and Control	48
2.4 Measure Theory	49
2.4.1 Main Concepts in Measure Theory	49
2.5 Probability Theory: Fundamentals	55

2.5.1	Probability Space and Random Variables	55
2.5.2	Distribution Functions and Probability Densities	56
2.5.3	Expectation and Variance	56
2.5.4	Important Inequalities	57
2.5.5	Convergence of Random Variables	57
2.5.6	Law of Large Numbers (LLN)	58
2.5.7	Central Limit Theorem (CLT)	58
2.5.8	Conditional Probability and Independence	58
2.5.9	Markov Chains and Markov Processes	59
2.6	Conclusion	59
3	A novel algorithmic approach to obtaining maneuverable viable sets	63
3.1	Introduction	64
3.1.1	Related work	64
3.1.2	Contributions	64
3.1.3	Organisation	65
3.2	Problem Setup	65
3.3	The maneuver set	67
3.4	Hamilton Jacobi Bellman reachability analysis	72
3.5	Numerical example: Dubins car	74
3.6	Results and discussion	74
3.7	Conclusions	76
4	A novel approach to the safe set and its estimation for systems with bounded disturbances	83
4.1	Introduction	84
4.1.1	Motivation	84
4.1.2	Related work	84
4.1.3	Contributions	86
4.1.4	Structure of the paper	86
4.2	Critical Analysis of the Standard Safe Set Definition	87
4.2.1	Set Definitions and Safe Set Formulation	87
4.3	Problem Formulation and Proposed Approach	90
4.3.1	Absence of Disturbances	90
4.3.2	Bounded Disturbances	102
4.4	Numerical Examples	105
4.4.1	System Dynamics: Double Integrator	105
4.4.2	System Dynamics: Dubins Car	106
4.4.3	System Dynamics: RCAM Slow Flight Envelope Model	106
4.5	Results	107
4.6	Conclusions	111
5	Robust avoid backward reachable tube	117
5.1	Introduction	118
5.2	Robust Dynamic Programming for Control Systems with Uncertainty	119
5.3	Relationship with the avoid backward reachable tube	145

5.4	Numerical Example	149
5.5	Conclusions	150
5.6	Appendix	151
5.7	Proof of equation 5.19	151
6	Certified Approximate Reachability (CARE)	157
6.1	Introduction	158
6.2	Problem Setup	159
6.3	Deep Learning Approach to HJ Reachability	160
6.4	Formal Synthesis of Neural Value Functions	161
6.5	ϵ -Accurate Value Functions	162
6.5.1	Bounding the True Value Function	163
6.5.2	Relating Bounds to the Neural Value Function	163
6.6	Case Studies	164
6.6.1	Training of value function	165
6.6.2	Formal Verification	166
6.7	Conclusions	168
6.8	Appendix: chapter 6	169
6.8.1	Bellman Optimality	169
6.8.2	Monotonicity of the value function	173
6.8.3	Existence and Uniqueness	175
6.8.4	Proof of Theorem 10	182
7	The Stochastic avoid problem with computation	191
7.1	Introduction	192
7.1.1	Contribution and organisation	193
7.2	Setting of the problem statement	194
7.3	Avoid set properties and computation	195
7.3.1	Probable forced backward reachable set computation	198
7.4	Convergence in time	206
7.5	Example	210
7.6	Conclusions	211
7.7	Appendix	212
8	Safe Probable Backward Reachable Sets: Theory and Algorithms	219
8.1	Introduction	220
8.2	Setting and problem statement	223
8.3	An algorithm to compute a Safe Probable Backward Reachable Set	225
8.4	An algorithm to compute the largest possible Safe Probable Backward Reachable Set	230
8.5	Discrete space implementation	236
8.5.1	Discrete space implementation of Algorithm 8.1	237
8.5.2	Discrete space implementation of Algorithm 8.2	237
8.6	Continuous space implementation	238
8.6.1	Continuous space implementation of Algorithm 8.1	240
8.6.2	Continuous space implementation of Algorithm 8.2	244

8.7	Example	244
8.8	Conclusions	246
9	Conclusions and Future Directions	253
9.1	Overview	254
9.1.1	Restatement of Research Objectives	255
9.2	Conclusions	255
9.2.1	Detailed Summary of Contributions	255
9.3	Limitations	260
9.3.1	Critical Reflections: Limitations of Maneuverable viable Sets and Safe Sets Chapter 3 and 4	261
9.3.2	Critical Reflections: Limitations of Robust Avoid Backward Reachable Tube Chapter 5	262
9.3.3	Critical Reflections: Limitations of Certified Approximate Reachability (CARE) Chapter 6	263
9.3.4	Critical Reflections: Limitations of the Stochastic Avoid Problem with Computation Chapter 7	264
9.3.5	Critical Reflections: Limitations of Safe Probable Backward Reachability Chapter 8	265
9.4	Future Work	266
9.4.1	Future Work: Maneuverable viable Sets and Safe Sets Chapter 3 and 4	267
9.4.2	Future Work: Robust Avoid Backward Reachable Tube Chapter 5	268
9.4.3	Future Work: Certified Approximate Reachability (CARE) Chapter 6	269
9.4.4	Future Work: Stochastic Avoidance and Safe Probable Backward Reachability Chapter 7 and 8	270
9.5	Research Impact	271
9.5.1	Autonomous Transportation Systems	271
9.5.2	Robotics and Human–Robot Interaction	271
9.5.3	Aerospace and Mission-Critical Systems	271
9.5.4	Healthcare and Biomedical Applications	271
9.5.5	Impact on the Foundations of Control Theory	272
9.5.6	Societal Trust in Autonomous Technologies	272
9.6	Final Remarks	272
	Acknowledgements	275
	Curriculum Vitæ	277
	List of Publications	279

SUMMARY

This dissertation develops a unified framework for defining, computing, and certifying safe operating regions for autonomous dynamical systems operating under deterministic, robust, and stochastic dynamics. As these systems are increasingly deployed in safety-critical environments—ranging from autonomous vehicles and air traffic management to robotics and automated decision-making—ensuring that their behaviour remains within safe bounds is of vital importance. Classical reachability methods offer strong theoretical guarantees, but they often rely on idealised models, struggle with high-dimensional systems, and cannot formally certify learned approximations. This work addresses these limitations through new safety definitions, robust generalisations of reachability theory, and a learning-verification pipeline for scalable certification.

The first contribution focuses on the notion of safe sets themselves. [Chapter 3](#) introduces the concept of the maneuverable control-invariant set—a novel refinement of traditional invariant sets. Unlike classical safe sets, which guarantee that the system can remain within a given region, maneuverable sets also ensure the system’s ability to move between any two states within the set in a given time horizon. Although such maneuverability and invariance cannot always be achieved simultaneously, this definition adds a valuable dimension to safety characterisation. The set is computed using a combination of forward and backward reachable sets derived from Hamilton-Jacobi (HJ) reachability analysis. Building on this, [Chapter 4](#) presents an improved definition of a safe set that does guarantee both viability and full internal maneuverability simultaneously. It is obtained through a different combination of forward and backward reachable sets. While the first two chapters establish refined definitions of safe sets under idealised dynamics, real-world systems rarely operate with perfect models. To address this, the next contribution turns to the challenge of ensuring safety under model uncertainty and disturbances.

[Chapter 5](#) addresses robustness under model uncertainty. It extends the HJ reachability framework by formulating a three-player differential game that includes unmodelled dynamics as an additional adversary alongside external disturbances. The resulting Robust Avoid Backward Reachable Tube (RABRT) generalises previous formulations by ensuring safety in the presence of both disturbances and uncertain system dynamics. The associated value function is shown to be the viscosity solution to a well-defined Hamilton-Jacobi-Isaacs (HJI) partial differential equation (PDE), and the method integrates with existing level-set tools, preserving compatibility with standard computational frameworks. While this formulation provides robustness to unmodelled dynamics, its computational demands grow quickly with system dimensionality. As classical level-set methods struggle in high-dimensional settings, recent work has turned to neural network-based approximations. The next

chapter explores how such approximations can be made verifiable and trustworthy through formal certification.

To overcome the curse of dimensionality inherent in classical reachability methods, [Chapter 6](#) introduces the Certified Approximate Reachability (CARE) framework. Instead of solving PDEs directly on grids, CARE leverages neural networks to approximate the value function and then uses formal verification tools—specifically Satisfiability Modulo Theory (SMT) solvers and Counterexample-Guided Inductive Synthesis (CEGIS)—to certify the learned approximation within a bounded error. This framework bridges the gap between deep learning and formal safety verification, offering a scalable method for high-dimensional systems while preserving correctness guarantees. While CARE provides a scalable and certifiable solution for deterministic systems, many real-world applications involve stochastic dynamics and probabilistic safety requirements. To address this, the final part of the thesis extends the reachability framework to systems governed by uncertainty and noise.

[Chapters 7 and 8](#) extend the thesis into the stochastic domain. [Chapter 7](#) introduces the Probabilistic Forced Backward Reachable Set, which identifies states from which the system enters an undesirable set with high probability, regardless of the control policy. The chapter develops convergent algorithms for computing this set in both discrete and continuous settings and provides error bounds under discretisation. [Chapter 8](#) then presents the Safe Probable Backward Reachable Set (SPBRS), which guarantees not only that the system will reach a target with a given probability but also that it will remain inside a certified safe set throughout the trajectory. This is especially important in stochastic systems where intermediate safety violations can occur despite a safe terminal state. Two algorithms are developed—one computationally efficient and the other optimal—and their convergence and tractability are formally analysed.

Each chapter is supported by theoretical proofs, algorithmic implementation, and numerical validation on benchmark systems such as the Dubins car and double integrator. This dissertation offers a cohesive and extensible framework for analysing and certifying safety in both deterministic and stochastic systems. By introducing new definitions of safety, developing scalable algorithms, and integrating formal verification with deep learning, the work makes contributions that are both foundational and practical. It addresses longstanding limitations in the scalability, conservatism, and verifiability of reachability-based safety analysis. The results have direct implications for real-world autonomous systems, including autonomous vehicles, robotics, air traffic management, and safety-critical embedded systems. More broadly, the methods contribute to the growing effort to make machine learning and autonomy formally trustworthy, especially in domains where safety cannot be left to statistical performance alone. Looking ahead, this work lays the groundwork for future research in adaptive safety verification, stochastic control synthesis, and trustworthy decision-making in uncertain environments—pushing the boundaries of what is possible at the intersection of formal methods, control theory, and autonomy.

1

INTRODUCTION

The motivation for this research stems from a fundamental requirement: ensuring that autonomous systems operate safely in real-world environments. Traditionally, safety assurance meant verifying every possible trajectory a system could take to prevent failures. However, as autonomous systems become more complex, this task has grown increasingly challenging.

Modern autonomous systems, such as quadcopters and self-driving cars, must operate in dynamic and unstructured environments. These environments often include unpredictable factors, such as other agents (e.g., humans, other vehicles), environmental disturbances, and changing conditions. As a result, many real-world implementations rely on heuristic safety measures, which may not be formally verified. This lack of rigorous safety guarantees has, in some cases, led to severe consequences, including accidents and system failures.

This dissertation addresses these issues by leveraging the concepts of backward and forward reachable tubes to develop a unified framework for defining, computing, and certifying safe operating regions for autonomous dynamic systems operating under deterministic, robust, and stochastic dynamics. It provides a more effective method for finding these regions and offers theoretical guarantees that the identified safe region is indeed safe, ensuring that no accidents will occur as long as the system remains within it. This chapter offers an introduction to the research by detailing the motivation, surveying the existing literature, setting out the research objective and questions, and describing its scope and methodology.

1.1. THE NEED FOR SAFE AND EFFICIENT AUTONOMOUS SYSTEMS

1.1.1. The Growing Role of Automation

Automation has become an essential part of modern technological advancement, revolutionizing industries such as transportation [1–3], aerospace [4–6], manufacturing [7], and robotics [8]. As systems grow more complex and expectations for performance increase, automation is no longer just an enhancement—it is a necessity. Several key factors drive the increasing need for autonomous and automated systems, the first of which being efficiency [9, 10]. Automated systems can perform tasks faster, more accurately, and with greater consistency than human operators. In industries such as logistics, autonomous transportation, and industrial manufacturing, automation allows systems to operate continuously with minimal downtime, maximizing output and reducing operational costs [11]. For example, in autonomous vehicles, automation optimizes fuel consumption and traffic flow, improving overall transportation efficiency [12, 13]. Secondly, many industries involve inherently dangerous tasks where automation can enhance safety. Robots and autonomous systems are increasingly used in high-risk environments, reducing human exposure [14]. For example, autonomous drones assess structural damage after earthquakes, eliminating the need for rescue teams to enter unstable buildings, while autonomous mining equipment operates underground, reducing accidents from cave-ins or gas leaks.

As technology advances, modern systems are becoming increasingly complex [15, 16]. Manual management is impractical, as human operators cannot process vast real-time data. For example, modern aircraft use automated flight control systems to integrate sensor data, stabilize flight, and optimize performance [17]. Without automation, such systems would be underutilized, error-prone, or inefficient. Finally, a significant challenge in traditional manually operated systems is the underutilization of system capabilities due to human limitations. Automation enables these systems to function at their maximum potential while maintaining safety and reliability [9, 18].

1.1.2. Consequences of Unsafe Automation

Automation has the potential to enhance efficiency, reduce human risk, and optimize system performance, but when implemented without proper safety guarantees, it can lead to unintended system behaviours, severe failures, financial and reputational damage, legal and regulatory consequences, and even injury or loss of life. Several major failures in automated systems highlight the risks associated with inadequate safety verification and incorrect assumptions about operational safety, for example the Boeing 737 MAX crashes from 2018 and 2019 [19]. The Manoeuvring Characteristics Augmentation System (MCAS) was introduced to adjust the aircraft's pitch in certain flight conditions, but it was designed based on incomplete safety verification and a lack of redundancy in sensor input. A proper rigorous safety analysis should have identified the risks



Figure 1.1: Industrial robotics in automotive assembly



Figure 1.2: Automation of logistics using small warehouse robots for reducing human workload



Figure 1.3: Autonomous multirotor inspecting a wind-turbine blade—shifting hazardous, high-altitude maintenance to unmanned systems while requiring robust perception, gust-resilient control, and fail-safe behaviours.

associated with using a single sensor without cross-checking redundancy, but the verification process was inadequate. This failure resulted in two fatal crashes, the loss of 346 lives, a global grounding of the fleet, and billions of dollars in financial and reputational damage [19]. Automated robots in manufacturing have caused severe workplace injuries due to incorrect safety zoning and misclassification of human presence [20]. Dynamic locomotion in legged robots, which has applications in industrial collaboration, is facing major barriers to commercialization due to the lack of standardized safety testing. Unlike traditional industrial automation, where regulatory safety frameworks exist, legged robots operate in unstructured environments, requiring a new category of safety validation [21]. Self-driving cars have been involved in multiple accidents due to sensor misinterpretations, software errors, and failure to handle rare edge cases [22]. These failures highlight the danger of overestimating system capability and the lack of rigorous safety validation for real-world conditions.

Automation has the potential to enhance efficiency, reduce human risk, and optimize system performance, but when implemented without proper safety guarantees, it can lead to unintended system behaviours, severe failures, financial and reputational damage, legal and regulatory consequences, and even injury or loss of life. Several major failures in automated systems highlight the risks associated with inadequate safety verification and incorrect assumptions about

operational safety, for example the Boeing 737 MAX crashes from 2018 and 2019 [19]. The Manoeuvring Characteristics Augmentation System (MCAS) was introduced to adjust the aircraft's pitch in certain flight conditions, but it was designed based on incomplete safety verification and a lack of redundancy in sensor input. A proper rigorous safety analysis should have identified the risks associated with using a single sensor without cross-checking redundancy, but the verification process was inadequate. This failure resulted in two fatal crashes, the loss of 346 lives, a global grounding of the fleet, and billions of dollars in financial and reputational damage [19]. Automated robots in manufacturing have caused severe workplace injuries due to incorrect safety zoning and misclassification of human presence [20]. Dynamic locomotion in legged robots, which has applications in industrial collaboration, is facing major barriers to commercialization due to the lack of standardized safety testing. Unlike traditional industrial automation, where regulatory safety frameworks exist, legged robots operate in unstructured environments, requiring a new category of safety validation [21]. Self-driving cars have been involved in multiple accidents due to sensor misinterpretations, software errors, and failure to handle rare edge cases [22]. These failures highlight the danger of overestimating system capability and the lack of rigorous safety validation for real-world conditions. Notable cases include the Uber ATG test vehicle striking a pedestrian in Tempe, Arizona (2018) [23], the Tesla Autopilot crash near Williston, Florida (2016) [24], and the California suspension and subsequent federal actions following a Cruise robotaxi's post-collision behaviour in San Francisco (2023) [25, 26]. In rail transit, the Washington Metro Red Line collision (2009) revealed hazards in automatic train control and track-circuit occupancy detection [27]. In commercial aviation, automation-mode confusion and logic limits played roles in Asiana Airlines 214 at SFO (2013) [28] and Qantas Flight 72's uncommanded pitch-down events due to ADIRU data spikes and flight-control computer logic (2008) [29]. And historically, the first recorded robot-related fatality in a factory (1979) underscores the need for robust safety guarantee for automation.[30, 31].

1.2. CHALLENGES IN ENSURING SAFE AUTOMATION

Ensuring the safety of autonomous systems is a complex and multi-faceted challenge, stemming from inherent uncertainties, high-dimensional state spaces, and unpredictable environments in which these systems operate. Unlike manually controlled systems, autonomous agents must continuously evaluate their safety in real time, often while operating under adversarial disturbances and dynamically changing conditions.

For discrete-time dynamical systems, safety, liveness, and fairness properties can be formally verified by analysing transition systems or by checking whether the system satisfies specific properties using finite automata representations [32, 33]. While these techniques offer strong theoretical guarantees, they do not scale well to continuous or hybrid systems, where the state space is infinite, and



Figure 1.4: A crashed drone



Figure 1.5: Boeing 737 MAX nose with the angle-of-attack (AoA) vane circled; MCAS relied on a single AoA input, a key factor in the 2018–2019 accidents [19].

system dynamics are often non-linear and coupled. In such settings, formal safety verification techniques like reachability analysis and set-based methods become essential for providing rigorous safety guarantees [34–36].

One of the primary challenges in this context is the ‘curse of dimensionality’: As the number of state variables increases, the computational cost of verification grows exponentially [34, 35]. This problem, also known as state-space explosion, significantly limits the feasibility of traditional grid-based methods, especially for real-time or high-dimensional applications. Formal verification techniques such as Hamilton-Jacobi (HJ) reachability rely on solving partial differential equations (PDEs) over a discretized grid of the state space, which incurs exponential computational complexity with respect to the number of dimensions



Figure 1.6: Cabin overhead panel damage on Qantas Flight 72 after uncommanded pitch-downs driven by erroneous ADIRU data and flight-control logic (in-flight upset, 7 Oct 2008, 154 km W of Learmonth, WA) [29].

[34–37]. While there have been efforts to address this—for instance, using the Hopf-Lax formula combined with convex optimization for linear systems [38, 39], or applying scalable methods for affine or polynomial systems using polytopic or ellipsoidal set representations [40–45]—the general case of nonlinear, high-dimensional systems remains computationally prohibitive.

Another key difficulty lies in our incomplete knowledge of a system’s true safe operating region. These regions are often defined using empirical testing, simplified mathematical models, or conservative rule-based heuristics, which may not accurately reflect real-world behaviour. This can lead either to overestimation, where the system operates beyond its actual safety limits and becomes vulnerable to failures, or to over-conservatism, where it operates far below its capacity, reducing efficiency [46, 47].

Compounding the issue is the dependence of most verification techniques on precise system models, typically expressed through differential equations or discrete-time transitions [34, 35, 37, 48]. However, in practice, these models are often incomplete, only partially observed, or uncertain. As a result, even formally verified safety guarantees may become invalid or misleading when deployed on the actual system.

Practical autonomous systems deployed in real-world settings must operate under unpredictable external disturbances and may even face adversarial agents attempting to disrupt their behaviour [48, 49]. These factors, combined with




Figure 1.7: Accident scene of Asiana Airlines Flight 214 at SFO: a low-energy approach with automation-mode mismanagement (autothrottle HOLD) led to impact short of runway 28L and a seawall strike [28].

continuous-time, nonlinear dynamics and uncertain environments, make scalable and robust safety certification an extremely demanding problem [50, 51].

Finally, when addressing stochastic systems, the challenge goes beyond ensuring safety in a purely deterministic sense. In safety-critical domains such as robotics, biomedical engineering, and aerospace, it is essential to guarantee that hard constraints are satisfied with high probability, rather than absolute certainty [52–57]. A fundamental question in this setting is: From which initial conditions can the system remain within a prescribed, time-varying set with at least a specified probability, despite having limited control authority?

In summary, while a rich body of theoretical work has laid the groundwork for verifying safety in autonomous systems, bridging the gap between formal guarantees and real-world deployment remains an open and pressing challenge. The limitations imposed by high-dimensional dynamics, model uncertainty, environmental unpredictability, computational intractability, and the need for probabilistic assurances underscore the need for scalable, adaptive, and model-aware verification methods. Addressing these challenges is essential to enabling



Figure 1.8: Wreckage of the Tesla Model S (Williston, Florida). With Autopilot engaged, the system failed to detect a crossing tractor-trailer; insufficient driver oversight resulted in a fatal underride [24].

the safe, efficient, and reliable operation of autonomous systems in complex, real-world environments. The following section reviews the current state of the art in safety verification and highlights key methodologies that have been developed to tackle these problems.

1.3. EXISTING APPROACHES AND THEIR LIMITATIONS

Over the years, a variety of approaches have been proposed to compute the reachable sets, given a range of initial conditions, control inputs, and external disturbances. These approaches differ in their underlying assumptions, scalability, treatment of non-linearities, and the types of systems they can accommodate. Broadly, existing work in reachability analysis can be grouped into seven methodological categories: Hamilton–Jacobi-based approaches, set-based over-approximation techniques [58], simulation-guided or data-driven analysis, SMT (Satisfiability Modulo Theories) and logic-based verification, probabilistic and stochastic reachability, learning-based approximations, and game-theoretic or adversarial reachability (Hamilton–Jacobi-based approaches can be classified as a special case of game-theoretic reachability).

Table 1.1 summarises the key features, strengths, and limitations of each category, along with representative toolboxes while figure 1.15 provides a decision tree for selecting a reachability analysis method.



Figure 1.9: Stock market crash

Hamilton–Jacobi (HJ) based methods offer rigorous and expressive frameworks for computing reachable sets by solving partial differential equations such as the Hamilton–Jacobi–Bellman (HJB) or Hamilton–Jacobi–Isaacs (HJI) equations [48, 59, 60]. These approaches are particularly well-suited to non-linear systems and are capable of encoding both safety and optimality objectives. Tools such as the Level Set Toolbox [61] and HelperOC [62], both implemented in MATLAB, provide powerful abstractions for computing different types of reachable sets under various control scenarios. BEACLS [63] extends these capabilities with CUDA-based acceleration for improved performance. Despite their theoretical robustness, a key drawback of HJ-based methods lies in their limited scalability: Due to the curse of dimensionality, they are typically restricted to systems with low dimensionality (often no more than five to six states) [36, 37].

Set-based reachability methods, on the other hand, take a geometric perspective by representing reachable sets using mathematical objects such as zonotopes, polytopes, or Taylor models [58]. These methods can accommodate non-linear dynamics to an extent, often by leveraging linearisation techniques or higher-order Taylor expansions. For instance, CORA [64] and JuliaReach [65] offer efficient and modular reachability computations based on zonotopic and support function representations. Flow* [66] employs Taylor model arithmetic to deliver accurate approximations for systems with polynomial dynamics, while SpaceEx [43] excels in scalable analysis of hybrid automata using polyhedral over-approximations. Although set-based methods are computationally more



Figure 1.10: Washington Metro Red Line collision near Fort Totten. An automatic train-control/track-circuit occupancy failure commanded a following train into an occupied block, causing a rear-end collision [27].

scalable than grid-based PDE solvers, they tend to produce conservative estimates for highly non-linear systems and can still face challenges as system complexity increases.

Simulation-guided or data-driven reachability offers a more empirical alternative, especially suitable when analytical models are incomplete or unavailable. These methods rely on system simulations or empirical trajectories to estimate the reachable set by sweeping through a range of initial conditions and input sequences. They are particularly attractive in scenarios involving black-box models, such as legacy code or learned systems [67, 68]. DryVR [68] is a Python-based tool that combines simulation traces with local sensitivity analysis to infer bounded reachable tubes for hybrid systems. Hylaa [69] focuses on linear approximations generated from simulations, while Verisig [70] enables reachability analysis for systems that integrate neural networks with hybrid dynamics. These tools are highly flexible and scalable, but they lack hard formal guarantees, and their reliability hinges on the breadth and resolution of the sampled trajectories.

In contrast, SMT [71] and logic-based reachability tools offer symbolic




Figure 1.11: Uber ATG test vehicle (Volvo XC90) front-end damage and the pedestrian’s bicycle. The ADS misclassified the pedestrian; OEM AEB was disabled for testing and the safety driver did not intervene in time [23].

verification of bounded-time safety properties. These tools translate reachability questions into logical formulas and solve them using specialised decision procedures. Non-linear dynamics and hybrid transitions can be encoded directly, although the complexity of the resulting formulas limits their use to systems with small state dimensions. dReal [72] is a δ -satisfiability solver that supports non-linear ODEs and hybrid models, while HyST [73] acts as a model converter interfacing with SMT solvers and other analysis tools. Unlike set-based or PDE-based reachability tools that aim to construct reachable sets over time, SMT-based approaches do not compute sets directly [74, 75]. Instead, they answer bounded-time satisfiability queries—such as whether a particular unsafe state or region is reachable from a given initial set. The result is a binary output (satisfiable or unsatisfiable), rather than a geometric over-approximation of the system’s evolution. While precise, these methods are not designed for full set computation and generally do not scale well beyond three to five dimensions due to their symbolic complexity [74].

Probabilistic or stochastic reachability extends the classical formulation by incorporating randomness in the form of stochastic disturbances, uncertain parameters, or probabilistic initial conditions. Rather than computing hard reachable sets, these approaches estimate the likelihood of the system entering a particular region within a given time horizon [57, 76–78]. This probabilistic

framing is particularly useful in risk-sensitive applications such as autonomous navigation under sensor noise. Tools such as SReachTools [79] and ProbReach [80] combine sampling-based methods with probabilistic decision procedures, while PRISM [81] offers model checking capabilities for Markov decision processes. These tools provide valuable insight into system behaviour under uncertainty but are computationally intensive and do not provide formal guarantees in the classical sense but in stochastic sense.

The rise of neural networks and learning-based control has inspired a new class of reachability tools that use deep learning to approximate value functions, reachable sets, or safety envelopes. These methods are especially effective for systems with high-dimensional or black-box dynamics, such as neural ODEs or reinforcement learning agents [35, 82, 83]. Tools such as DeepReach [35] employ unsupervised training paradigms to infer the reachability structure of complex systems from data. While these methods are highly scalable and computationally efficient at runtime, they offer only approximate guarantees and require careful validation to ensure robustness, especially in safety-critical applications.

Finally, game-theoretic or adversarial reachability casts the problem as a two-player zero-sum differential game between a controller and a worst-case disturbance [48, 59]. This framework naturally arises in safety-critical settings where robustness against adversarial inputs is essential. Solving the corresponding Hamilton–Jacobi–Isaacs (HJI) equation yields a value function that encodes the optimal safety policy under worst-case scenarios. HelperOC [62] and HJReachability.jl [65] support such formulations, while FaSTrack [39, 60] integrates reachability into motion planning by tracking planner trajectories with guaranteed error bounds. While game-theoretic methods provide strong safety guarantees, they are subject to the same dimensional limitations as classical HJ-based methods and are computationally expensive.




Figure 1.12: Backward reachable set (BRs) for a Dubins car over a 5 s horizon in (x, y, θ) . The green set encloses states that can reach the small goal region (blue cylinder) within the horizon under constant speed and bounded turn rate, yielding a characteristic helical geometry.




Figure 1.13: Forward reachable set (FRs) for a Dubins car over a 5 s horizon in (x, y, θ) . Starting from an initial set (blue cylinder), the green volume captures all states the vehicle can attain under constant speed and bounded curvature, illustrating the nonconvex, swirl-like structure induced by the kinematics.




Figure 1.14: Stochastic forward reachable set for a planar double integrator: estimated joint probability density $p(x_1, x_2)$ over a finite horizon. The surface illustrates how process uncertainty spreads probability mass across the state plane, emphasizing risk-aware reachability.


Figure 1.15: Decision Tree for Selecting a Reachability Analysis Method

Table 1.1: Comparison of Reachability Analysis Methods

Property	HJ (Hamilton–Jacobi)	Set-Based	Simulation-Guided	SMT / Logic-Based	Stochastic / Probabilistic	Neural Learning-Based	Game-Theoretic
Nonlinearity Support	Full	Partial (via linearisation)	Full	Symbolic	Partial (sample-based)	Full	Full
Model Requirement	Full dynamics	Symbolic or numeric model	Black-box / simulator	Symbolic model	Probabilistic model	Data / sim access	Full model + disturbance bounds
Scalability	Low ($\leq 5-6D$)	Medium ($\leq 5-10D$)	High ($< 10D$)	Low ($\leq 5D$)	Medium ($\leq 5-10D$)	High ($< 10D$)	Low ($\leq 5-6D$)
Guarantees	Yes	Yes	No	Yes	probabilistic	Approximate	Yes
Reachable Set	Yes	Yes	Yes	No	Yes	Yes	Yes
Example Tools	Level Set Toolbox, HelperOC, BEACLS	CORA, JuliaReach, Flow*, SpaceEx	DryVR, HyLaa, Verisig	dReal, HyST	SReachTools, ProbReach, PRISM	DeepReach	HelperOC, HJReachability.jl, FaSTrack

1.4. RESEARCH OBJECTIVES

The objective of this research is to identify and certify safe operating regions in the state space of dynamic systems. These regions, often referred to as safe sets, are critical for ensuring that autonomous or safety-critical systems operate within acceptable bounds—even in the presence of disturbances, uncertainties, or modelling errors. The focus lies on developing theoretical frameworks and practical algorithms for computing these sets under both deterministic and stochastic settings. Additionally, the research aims to provide formal certification for learned approximations of these sets, ensuring reliability even in high-dimensional systems where classical methods struggle.

Main Research Question

How can we formally define, compute, and certify safe operating regions for dynamic systems—under both deterministic and stochastic settings—while ensuring scalability, robustness to uncertainty, and formal correctness guarantees?

The dissertation is structured such that each chapter contributes to answering the main research question, while building on each other. The remainder of this section presents the research question and objective of each chapter.

Research question: Chapter 3 and 4

What is an appropriate definition of a safe set that ensures both viability and maneuverability, and how can such sets be computed using reachability analysis?

Existing definitions of safe sets neither guarantee viability nor ensure that a system can transition between any two states within the set in a given time horizon. The chapters identify this limitation and propose a new definition of safe sets (Chapter 4), and introduce the concept of manoeuvre sets (Chapter 3): control-invariant regions where two-way reachability is guaranteed within a time limit. They provide algorithms to compute such sets using Hamilton-Jacobi reachability tools, and show how they relate to forward and backward reachable sets.

Research question: Chapter 5

How can bounded uncertainties and unmodelled dynamics be systematically incorporated into reachability analysis?

Traditional Hamilton-Jacobi-Isaacs (HJI) formulations involve two-player games (controller vs disturbance) and assume perfect system models. To improve

robustness, this dissertation extends this to a three-player differential game, treating unmodelled dynamics as a third adversarial player. This framework enables the computation of robust avoid backward reachable tubes that remain valid under all realizations of disturbances and modelling errors.

Research question: Chapter 6

How can the correctness of learned reachable sets from neural approximators be formally certified?

Level-set methods for reachability analysis suffer from the curse of dimensionality, making them impractical for high-dimensional systems. Deep learning approaches like DeepReach overcome this limitation but lack theoretical guarantees on the safety of their outputs. This chapter introduces Certified Approximate Reachability (CARE), a framework that uses SMT solvers to certify satisfaction of a relaxed Hamilton-Jacobi PDE. Using Counterexample-Guided Inductive Synthesis (CEGIS), it refines the neural network iteratively to improve both accuracy and certifiability.

Research question: Chapter 7

What is the relationship between probabilistic backward reachable sets and viable sets in stochastic systems, and how can such sets be computed efficiently?

This chapter extends the concept of forced backward reachability to stochastic systems and analyses its relationship with probabilistic invariance. Algorithms for computing probabilistic forced backward reachable sets are developed, ensuring that the system is driven into the target set with high probability. This chapter establishes both the theoretical correctness and practical efficiency of the proposed methods, enabling applications such as air traffic control and obstacle avoidance.

Research question: Chapter 8

How can we compute stochastic backward reachable set that ensure safety with high probability?

In stochastic settings, it is often desirable to identify sets that the system can remain in with high probability while still reaching the target set or avoiding the unsafe region. This chapter introduces the Safe Probable Backward Reachable Set (SPBRS), a probabilistically control-invariant set guaranteeing that the system avoids unsafe regions until reaching the target set. It provides two algorithms for computing SPBRS in discrete and continuous state spaces, with convergence guarantees and efficiency trade-offs.

1.5. SCOPE OF THE RESEARCH

This research primarily focuses on the theoretical analysis and formal certification of safe sets for dynamical systems. The core contributions lie in defining new notions of safety, proving their properties mathematically, and developing or adapting computational methods to obtain these sets under both deterministic and stochastic settings.

The scope is deliberately limited to the computation and verification of the safe sets themselves. While the value functions obtained via Hamilton-Jacobi formulations inherently contain information about optimal controllers, this work does not explore the synthesis, analysis, or implementation of the associated control policies. Instead, the emphasis is placed on ensuring that the computed sets possess the desired safety, invariance, and reachability properties.

Furthermore, this dissertation remains largely theoretical and simulation-based. The proposed methods are validated on benchmark dynamical systems such as the double integrator and the Dubins car, which are widely used in the literature. However, the methods at the time of this writing not been applied to real-world systems or hardware platforms. As such, issues related to sensing, actuation, real-time computation, or implementation-specific constraints are beyond the current scope.

This scope allows for a deep and rigorous treatment of the core problem—identifying and certifying safe operating regions—while leaving the integration with control synthesis and real-world deployment as directions for future work.

1.6. METHODOLOGY

This dissertation follows a theoretical and algorithmic approach to identifying and certifying safe operating regions in dynamic systems. Across all research questions, the methodology consists of three main components: formal definition and mathematical proof, practical computation, and empirical validation. The selected tools—Hamilton-Jacobi reachability analysis, dynamic programming, viscosity solution theory, Satisfiability Modulo Theories (SMT), and learning-based approximations—are grounded in well-established mathematical principles and were chosen to address the central challenge of ensuring safety under both deterministic and stochastic settings.

Each contribution begins with a theoretical formulation wherein the proposed set—whether deterministic, robust, or stochastic—is shown to satisfy essential safety properties. These include viability, two-way reachability, and probabilistic safety guarantees. The properties are established through formal mathematical proofs, ensuring that the sets possess the claimed characteristics under the specified system dynamics and constraints. HJ reachability was foundational to this work due to its ability to characterise sets of states from which a system can satisfy a control objective under worst-case or constrained dynamics. These PDE-based methods naturally model safety, reachability, and invariance as properties of the value function—offering formal guarantees on system behaviour.

Viability theory, with its emphasis on control-invariance and safe evolution, aligns with the goal of constructing mathematically sound safe sets. The use of forward and backward reachable sets to construct both maneuverable and viable safe regions (as in Chapter 3 and 4) was especially useful because it allowed the articulation of time-aware, geometry-sensitive safety definitions. These sets encode both where a system can go and where it can stay—key for applications such as navigation, recovery, and motion planning.

Following the theoretical development, practical computation of these sets is carried out. In several chapters, new algorithms are designed to align with the proposed definitions—particularly for stochastic systems, where existing tools are insufficient. For deterministic settings, established toolboxes are extended through minimal but meaningful modifications. Specifically, the Level Set Toolbox [61] is utilised in the first three chapters, with adjustments such as altering the target set structure or switching between forward and backward reachability modes. In the case of neural network-based reachable set approximations, the DeepReach [35] toolbox is employed to address the curse of dimensionality, and is paired with the dReal [72] SMT solver to certify the correctness of the learned value function. For high-dimensional systems, deep learning approximators such as neural networks provide a scalable way to model reachable sets. However, scalability often comes at the cost of formal guarantees. The CARE framework innovatively employs SMT solvers and CEGIS to formally certify these learned approximations. This choice of methodology is essential for reintroducing trust and provability into learning-based reachability analysis—an area traditionally seen as black-box and unverifiable.

For the stochastic chapters, where no ready-made tools were available, custom algorithms were implemented in Matlab, designed to work with both discrete and continuous state spaces. These implementations include grid-based methods with convergence guarantees and optimisation-based formulations for probabilistic reachability. For Chapter 7 and 8, stochastic safety analysis is framed in discrete time. This modeling choice offers the dual benefit of computational tractability and compatibility with real-world digital control systems. Discrete-time dynamics also enable the formal definition of probabilistic safety sets (PFBRs, SPBRs) using measurable transitions and Markov policies.

To validate the theoretical results and demonstrate the effectiveness of the algorithms, numerical simulations are conducted using standard benchmark systems. Specifically, the double integrator and the Dubins car models—widely accepted in the reachability and control literature—are used. For each method, the computed safe sets are visualised and tested to confirm the theoretical properties established earlier. This consistent use of benchmark systems allows for comparative analysis across different methods and highlights the versatility and robustness of the proposed approaches.

In summary, the methodologies adopted in this work are well-suited to the research objectives and have proven essential for enabling new perspectives on safety certification under uncertainty. By combining control theory, reachability analysis, learning, and logic-based verification, the dissertation aims to ensure

that safe autonomy is mathematically grounded, computationally feasible, and practically impactful.

1.7. GUIDE TO THE READER

This dissertation is organised into eight chapters. [Chapter 1](#) detailed the motivation, surveyed the existing literature, set out the research objective and questions, and described its scope and methodology. [Chapter 2](#) provides the necessary background, covering reachability theory and various types of reachable sets, an overview of SMT theory for verification, the basics of neural networks for value function approximation, and key concepts from stochastic systems theory. [Chapter 3](#) to [Chapter 8](#) each address a specific research question and are written as standalone contributions. This format follows the article-based dissertation structure used at TU Delft. [Chapter 3](#) introduces the maneuverable set, a redefined safe set that ensures viability and two-way reachability, and presents methods to compute it using Hamilton-Jacobi tools. [Chapter 4](#) builds on the previous chapter to address the limitations of maneuverable sets introduced in it. [Chapter 5](#) extends reachability analysis to systems with unmodelled dynamics through a three-player differential game formulation. [Chapter 6](#) presents a method for certifying learned reachable sets using neural networks and SMT solvers, addressing the challenge of high dimensionality. [Chapter 7](#) focuses on computing probabilistic forced backward reachable sets and analysing their relationship with stochastic invariance, providing efficient algorithms with theoretical guarantees. [Chapter 8](#) introduces the Safe Probable Backward Reachable Set (SPBRS), a new notion of probabilistically control-invariant safety, along with algorithms for both discrete and continuous systems. Finally, [Chapter 9](#) concludes the dissertation by summarising the main findings, addressing the research questions, and outlining promising directions for future work, including controller synthesis and application to real-world systems.




Figure 1.16: Thesis flow diagram linking the main research question to deterministic (chapters 3 to 5) and stochastic (chapters 7 and 8) developments, with cross-cutting learning/certification in chapter 6. Solid arrows indicate conceptual/methodological progression; dashed arrows indicate future research areas.

REFERENCES

- [1] K. D. Kusano, K. Beatty, S. Schnelle, F. Favaro, C. Crary, and T. Victor. “Collision avoidance testing of the waymo automated driving system”. In: *arXiv preprint arXiv:2212.08148* (2022).
- [2] G. Ma, B. Wang, and S. S. Ge. “Robust optimal control of connected and automated vehicle platoons through improved particle swarm optimization”. In: *Transportation research part C: emerging technologies* 135 (2022), p. 103488.
- [3] S. Wen and G. Guo. “Distributed trajectory optimization and sliding mode control of heterogenous vehicular platoons”. In: *IEEE Transactions on Intelligent Transportation Systems* 23.7 (2021), pp. 7096–7111.
- [4] K. M. Kabore and S. Güler. “Distributed formation control of drones with onboard perception”. In: *IEEE/ASME Transactions on Mechatronics* 27.5 (2021), pp. 3121–3131.
- [5] X. Liu, K. Lam, B. Alkouz, B. Shahzaad, and A. Bouguettaya. “Constraint-based formation of drone swarms”. In: *2022 IEEE international conference on pervasive computing and communications workshops and other affiliated events (PerCom workshops)*. IEEE, 2022, pp. 73–75.
- [6] P. Zhang, G. Chen, Y. Li, and W. Dong. “Agile formation control of drone flocking enhanced with active vision-based relative localization”. In: *IEEE Robotics and Automation Letters* 7.3 (2022), pp. 6359–6366.
- [7] R. Ashima, A. Haleem, S. Bahl, M. Javaid, S. K. Mahla, and S. Singh. “Automation and manufacturing of smart materials in Additive Manufacturing technologies using Internet of Things towards the adoption of Industry 4.0”. In: *Materials Today: Proceedings* 45 (2021), pp. 5081–5088.
- [8] S. Li, P. Zheng, S. Liu, Z. Wang, X. V. Wang, L. Zheng, and L. Wang. “Proactive human–robot collaboration: Mutual-cognitive, predictable, and self-organising perspectives”. In: *Robotics and Computer-Integrated Manufacturing* 81 (2023), p. 102510.
- [9] E. Dinlersoz and Z. Wolf. “Automation, labor share, and productivity: Plant-level evidence from US Manufacturing”. In: *Economics of Innovation and New Technology* 33.4 (2024), pp. 604–626.
- [10] P. Aghion, C. Antonin, S. Bunel, and X. Jaravel. “The effects of automation on labor demand”. In: *Robots and AI* (2023), pp. 15–39.

- [11] T. R. Mohan, J. P. Roselyn, R. A. Uthra, D. Devaraj, and K. Umachandran. “Intelligent machine learning based total productive maintenance approach for achieving zero downtime in industrial machinery”. In: *Computers & Industrial Engineering* 157 (2021), p. 107267.
- [12] X. Shi, H. Yao, Z. Liang, and X. Li. “An empirical study on fuel consumption of commercial automated vehicles”. In: *Transportation Research Part D: Transport and Environment* 106 (2022), p. 103253.
- [13] Z. Yao, Y. Wang, B. Liu, B. Zhao, and Y. Jiang. “Fuel consumption and transportation emissions evaluation of mixed traffic flow with connected automated vehicles and human-driven vehicles on expressway”. In: *Energy* 230 (2021), p. 120766.
- [14] G. Wilk-Jakubowski, R. Harabin, and S. Ivanov. “Robotics in crisis management: A review”. In: *Technology in Society* 68 (2022), p. 101935.
- [15] G. Bathla, K. Bhadane, R. K. Singh, R. Kumar, R. Aluvalu, R. Krishnamurthi, A. Kumar, R. Thakur, and S. Basheer. “Autonomous vehicles and intelligent automation: Applications, challenges, and opportunities”. In: *Mobile Information Systems* 2022.1 (2022), p. 7632892.
- [16] M. A. Jamshed, K. Ali, Q. H. Abbasi, M. A. Imran, and M. Ur-Rehman. “Challenges, applications, and future of wireless sensors in Internet of Things: A review”. In: *IEEE Sensors Journal* 22.6 (2022), pp. 5482–5494.
- [17] M. Lungu. “Backstepping and dynamic inversion control techniques for automatic landing of fixed wing unmanned aerial vehicles”. In: *Aerospace Science and Technology* 120 (2022), p. 107261.
- [18] D. Hopkins and T. Schwanen. “Talking about automated vehicles: What do levels of automation do?” In: *Technology in Society* 64 (2021), p. 101488.
- [19] J. Herkert, J. Borenstein, and K. Miller. “The Boeing 737 MAX: Lessons for engineering ethics”. In: *Science and engineering ethics* 26 (2020), pp. 2957–2974.
- [20] T. Backström and M. Döös. “Absolute and relative frequencies of automation accidents at different kinds of equipment and for different occupational groups”. In: *Journal of safety research* 28.3 (1997), pp. 147–158.
- [21] B. Weng, G. A. Castillo, Y.-S. Kang, and A. Hereid. “Towards standardized disturbance rejection testing of legged robot locomotion with linear impactor: A preliminary study, observations, and implications”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2024, pp. 9946–9952.
- [22] N. Beckers, L. C. Siebert, M. Bruijnes, C. Jonker, and D. Abbink. “Drivers of partially automated vehicles are blamed for crashes that they cannot reasonably avoid”. In: *Scientific reports* 12.1 (2022), p. 16193.

- [23] National Transportation Safety Board. *Collision Between Vehicle Controlled by Developmental Automated Driving System and Pedestrian, Tempe, Arizona, March 18, 2018*. Tech. rep. NTSB/HAR-19/03. Corrected copy issued June 26, 2020. NTSB, 2019. URL: <https://www.nts.gov/investigations/accidentreports/reports/har1903.pdf>.
- [24] National Transportation Safety Board. *Collision Between a Car Operating With Automated Vehicle Control Systems and a Tractor-Semitrailer Truck Near Williston, Florida, May 7, 2016*. Tech. rep. NTSB/HAR-17/02. NTSB, 2017. URL: <https://www.nts.gov/Investigations/Accidentreports/Reports/Har1702.pdf>.
- [25] California Department of Motor Vehicles. *DMV Statement on Cruise LLC Suspension*. Oct. 2023. URL: <https://www.dmv.ca.gov/portal/news-and-media/dmv-statement-on-cruise-llc-suspension/>.
- [26] National Highway Traffic Safety Administration. *Part 573 Safety Recall Report 23E-086 (Cruise ADS Software)*. Tech. rep. NHTSA, Nov. 2023. URL: <https://static.nhtsa.gov/odi/rcl/2023/RCLRPT-23E086-7725.PDF>.
- [27] National Transportation Safety Board. *Collision of Two Washington Metropolitan Area Transit Authority Metrorail Trains Near Fort Totten Station, Washington, D.C., June 22, 2009*. Tech. rep. NTSB/RAR-10/02. NTSB, 2010. URL: <https://www.nts.gov/investigations/accidentreports/reports/rar1002.pdf>.
- [28] National Transportation Safety Board. *Descent Below Visual Glidepath and Impact With Seawall, Asiana Airlines Flight 214, Boeing 777-200ER, San Francisco, California, July 6, 2013*. Tech. rep. NTSB/AAR-14/01. NTSB, 2014. URL: <https://www.nts.gov/investigations/accidentreports/reports/aar1401.pdf>.
- [29] Australian Transport Safety Bureau. *In-flight upset, 154 km west of Learmonth, Western Australia, 7 October 2008, VH-QPA, Airbus A330-303*. Tech. rep. AO-2008-070. ATSB, 2011. URL: <https://www.atsb.gov.au/sites/default/files/media/3532398/ao2008070.pdf>.
- [30] Guinness World Records. *First human to be killed by a robot*. Entry recounts the death of Robert Williams on Jan 25, 1979, at Ford's Flat Rock plant. 1979. URL: <https://guinnessworldrecords.com/world-records/first-human-to-be-killed-by-a-robot>.
- [31] *Williams v. Litton Systems, Inc.*, 449 N.W.2d 669 (Mich. 1989). Michigan Supreme Court decision. 1989. URL: <https://law.justia.com/cases/michigan/supreme-court/1989/81951-5.html>.
- [32] C. Belta, B. Yordanov, and E. A. Gol. *Formal methods for discrete-time dynamical systems*. Vol. 89. Springer, 2017.
- [33] C. Baier and J.-P. Katoen. *Principles of model checking*. MIT press, 2008.

- [34] S. L. Herbert, S. Bansal, S. Ghosh, and C. J. Tomlin. “Reachability-based safety guarantees using efficient initializations”. In: *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE. 2019, pp. 4810–4816.
- [35] S. Bansal and C. J. Tomlin. “Deepreach: A deep learning approach to high-dimensional reachability”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 1817–1824.
- [36] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin. “Decomposition of reachable sets and tubes for a class of nonlinear systems”. In: *IEEE Transactions on Automatic Control* 63.11 (2018), pp. 3675–3688.
- [37] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-jacobi reachability: A brief overview and recent advances”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 2242–2253.
- [38] J. Darbon and S. Osher. “Algorithms for overcoming the curse of dimensionality for certain Hamilton–Jacobi equations arising in control theory and elsewhere”. In: *Research in the Mathematical Sciences* 3.1 (2016), p. 19.
- [39] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin. “Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking”. In: *IEEE Transactions on Automatic Control* 66.12 (2021), pp. 5861–5876.
- [40] T. Dreossi, T. Dang, and C. Piazza. “Parallelotope bundles for polynomial reachability”. In: *Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control*. 2016, pp. 297–306.
- [41] D. Henrion and M. Korda. “Convex computation of the region of attraction of polynomial control systems”. In: *IEEE Transactions on Automatic Control* 59.2 (2013), pp. 297–312.
- [42] A. Majumdar, A. A. Ahmadi, and R. Tedrake. “Control design along trajectories with sums of squares programming”. In: *2013 IEEE International Conference on Robotics and Automation*. IEEE. 2013, pp. 4054–4061.
- [43] G. Frehse, C. Le Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler. “SpaceEx: Scalable verification of hybrid systems”. In: *Computer Aided Verification: 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings* 23. Springer. 2011, pp. 379–395.
- [44] A. B. Kurzhanski and P. Varaiya. “On ellipsoidal techniques for reachability analysis. part ii: Internal approximations box-valued constraints”. In: *Optimization methods and software* 17.2 (2002), pp. 207–237.
- [45] J. N. Maidens, S. Kaynama, I. M. Mitchell, M. M. Oishi, and G. A. Dumont. “Lagrangian methods for approximating the viability kernel in high-dimensional systems”. In: *Automatica* 49.7 (2013), pp. 2017–2029.

- [46] P. Solanki, J. J. Van Beers, A. Jamshidnejad, and C. C. De Visser. “A novel algorithmic approach to obtaining maneuverable control-invariant sets”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2024, pp. 9916–9922.
- [47] T. Lombaerts, S. Schuet, K. Wheeler, D. M. Acosta, and J. Kaneshige. “Safe maneuvering envelope estimation based on a physical approach”. In: *Aiaa guidance, navigation, and control (gnc) conference*. 2013, p. 4618.
- [48] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin. “A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games”. In: *IEEE Transactions on automatic control* 50.7 (2005), pp. 947–957.
- [49] E. N. Barron. *Game theory: an introduction*. John Wiley & Sons, 2024.
- [50] O. Bokanowski and H. Zidani. “Minimal time problems with moving targets and obstacles”. In: *IFAC Proceedings Volumes* 44.1 (2011), pp. 2589–2593.
- [51] J. F. Fisac, M. Chen, C. J. Tomlin, and S. S. Sastry. “Reach-avoid problems with time-varying dynamics, targets and constraints”. In: *Proceedings of the 18th international conference on hybrid systems: computation and control*. 2015, pp. 11–20.
- [52] A. Abate. “A contractivity approach for probabilistic bisimulations of diffusion processes”. In: *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*. IEEE. 2009, pp. 2230–2235.
- [53] A. J. Summers and P. Müller. “Automating deductive verification for weak-memory programs”. In: *Tools and Algorithms for the Construction and Analysis of Systems: 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings, Part I* 24. Springer. 2018, pp. 190–209.
- [54] N. Malone, K. Lesser, M. Oishi, and L. Tapia. “Stochastic reachability based motion planning for multiple moving obstacle avoidance”. In: *Proceedings of the 17th international conference on Hybrid systems: computation and control*. 2014, pp. 51–60.
- [55] A. P. Vinod and M. M. Oishi. “Scalable underapproximation for the stochastic reach-avoid problem for high-dimensional LTI systems using Fourier transforms”. In: *IEEE control systems letters* 1.2 (2017), pp. 316–321.
- [56] A. P. Vinod and M. M. Oishi. “Scalable underapproximative verification of stochastic LTI systems using convexity and compactness”. In: *Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control (Part of CPS Week)*. 2018, pp. 1–10.
- [57] A. P. Vinod and M. M. Oishi. “Stochastic reachability of a target tube: Theory and computation”. In: *Automatica* 125 (2021), p. 109458.

- [58] M. Althoff, G. Frehse, and A. Girard. “Set propagation techniques for reachability analysis”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 4.1 (2021), pp. 369–395.
- [59] J. Lygeros. “On reachability and minimum cost optimal control”. In: *Automatica* 40.6 (2004), pp. 917–927.
- [60] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin. “FaSTrack: A modular framework for fast and guaranteed safe motion planning”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 1517–1522.
- [61] I. Mitchell, A. Bayen, and C. J. Tomlin. “Computing reachable sets for continuous dynamic games using level set methods”. In: *Submitted January* (2004).
- [62] C. Jongho. *helperOC*. URL: <https://github.com/HJReachability/helperOC>.
- [63] C. Jongho. *BEACLS*. URL: <https://hjreachability.github.io/beacsl/>.
- [64] M. Althoff. “An introduction to CORA 2015”. In: *Proc. of the workshop on applied verification for continuous and hybrid systems*. 2015, pp. 120–151.
- [65] S. Bogomolov, M. Forets, G. Frehse, K. Potomkin, and C. Schilling. “JuliaReach: a toolbox for set-based reachability”. In: *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*. 2019, pp. 39–44.
- [66] X. Chen. *User’s Manual of Flow v1. 2.0*. 2013.
- [67] W. Xiang, H.-D. Tran, X. Yang, and T. T. Johnson. “Reachable set estimation for neural network control systems: A simulation-guided approach”. In: *IEEE Transactions on Neural Networks and Learning Systems* 32.5 (2020), pp. 1821–1830.
- [68] C. Fan, B. Qi, S. Mitra, and M. Viswanathan. “DryVR: Data-driven verification and compositional reasoning for automotive systems”. In: *International Conference on Computer Aided Verification*. Springer. 2017, pp. 441–461.
- [69] S. Bak and P. S. Duggirala. “Hylaa: A tool for computing simulation-equivalent reachability for linear systems”. In: *Proceedings of the 20th International Conference on Hybrid Systems: Computation and Control*. 2017, pp. 173–178.
- [70] R. Ivanov, J. Weimer, R. Alur, G. J. Pappas, and I. Lee. “Verisig: verifying safety properties of hybrid systems with neural network controllers”. In: *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*. 2019, pp. 169–178.
- [71] C. Barrett and C. Tinelli. “Satisfiability modulo theories”. In: *Handbook of model checking* (2018), pp. 305–343.

- [72] S. Gao, S. Kong, and E. M. Clarke. “dReal: An SMT solver for nonlinear theories over the reals”. In: *International conference on automated deduction*. Springer. 2013, pp. 208–214.
- [73] S. Bak, S. Bogomolov, and T. T. Johnson. “HYST: a source transformation and translation tool for hybrid automaton models”. In: *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*. 2015, pp. 128–133.
- [74] A. Abate, A. Edwards, and M. Giacobbe. “Neural abstractions”. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 26432–26447.
- [75] A. Abate, C. David, P. Kesseli, D. Kroening, and E. Polgreen. “Counterexample guided inductive synthesis modulo theories”. In: *International Conference on Computer Aided Verification*. Springer. 2018, pp. 270–288.
- [76] J. D. Gleason, A. P. Vinod, and M. M. Oishi. “Lagrangian approximations for stochastic reachability of a target tube”. In: *Automatica* 128 (2021), p. 109546.
- [77] L. M. Bujorianu. *Stochastic reachability analysis of hybrid systems*. Springer Science & Business Media, 2012.
- [78] A. P. Vinod, B. HomChaudhuri, and M. M. Oishi. “Forward stochastic reachability analysis for uncontrolled linear systems using fourier transforms”. In: *Proceedings of the 20th International Conference on Hybrid Systems: Computation and Control*. 2017, pp. 35–44.
- [79] A. P. Vinod, J. D. Gleason, and M. M. Oishi. “SReachTools: a MATLAB stochastic reachability toolbox”. In: *Proceedings of the 22nd ACM international conference on hybrid systems: computation and control*. 2019, pp. 33–38.
- [80] F. Shmarov and P. Zuliani. “ProbReach: A Tool for Guaranteed Reachability Analysis of Stochastic Hybrid Systems.” In: *SNR@ CAV*. 2015, pp. 40–48.
- [81] M. Kwiatkowska, G. Norman, and D. Parker. “PRISM 4.0: Verification of probabilistic real-time systems”. In: *Computer Aided Verification: 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings* 23. Springer. 2011, pp. 585–591.
- [82] A. Lin and S. Bansal. “Verification of neural reachable tubes via scenario optimization and conformal prediction”. In: *6th Annual Learning for Dynamics & Control Conference*. PMLR. 2024, pp. 719–731.
- [83] A. Lin and S. Bansal. “Generating formal safety assurances for high-dimensional reachability”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2023, pp. 10525–10531.

2

BACKGROUND

This chapter presents the theoretical foundations necessary for the developments in subsequent chapters, each of which contributes to safe decision-making for autonomous systems under various forms of uncertainty. It begins with Hamilton-Jacobi (HJ) reachability analysis in [section 2.1](#), a robust optimal control framework that underlies our formulations of manoeuvrable and safe sets, and which we later extend to model uncertainty and stochastic settings. It then covers neural networks in [section 2.2](#), which serve as function approximators for high-dimensional PDE solutions, and Satisfiability Modulo Theories (SMT) in [section 2.3](#), which provide formal tools for certifying solution error bounds. Finally, it reviews essential concepts from measure theory and probability in [sections 2.4](#) and [2.5](#) respectively, which are foundational to our construction of stochastic reachable and avoid sets, and their associated properties. Together, these topics form the mathematical and computational basis for the novel theoretical contributions and algorithms introduced in the remainder of the thesis.

2.1. HAMILTON-JACOBI REACHABILITY: FOUNDATIONS AND APPLICATIONS

Hamilton-Jacobi (HJ) reachability analysis is a powerful control-theoretic framework for reasoning about safety and performance guarantees in dynamical systems operating under uncertainty [1–3]. It is particularly suited to safety-critical applications such as autonomous driving[4], aircraft collision avoidance[5], and human-robot interaction[6], where worst-case behaviours must be accounted for[7, 8].

2.1.1. Optimal Control and the Hamilton-Jacobi-Bellman Equation

Optimal control theory provides a mathematical framework for making optimal decisions for dynamical systems. Consider a control-affine system of the form:

$$\dot{x}(t) = f(x(t), u(t)), \quad u(t) \in \mathcal{U}, \quad (2.1)$$

where $x(t) \in \mathbb{R}^n$ is the system state, $u(t)$ is the control input, and f is the system dynamics. The goal is to find a control input $u(\cdot)$ that minimises a cost function of the form:

$$J_t(x; u(\cdot)) = \int_t^T L(x(s), u(s)) ds + M(x(T)), \quad (2.2)$$

where $L(x, u)$ is a running cost and $M(x)$ is a terminal cost.

The optimal cost-to-go from a state x at time t is captured by the value function:

$$V(t, x) = \inf_{u(\cdot)} J_t(x; u(\cdot)). \quad (2.3)$$

Using Bellman's principle of optimality, the value function satisfies the Hamilton-Jacobi-Bellman (HJB) partial differential equation [9]:

$$\frac{\partial V}{\partial t}(t, x) + \inf_{u \in \mathcal{U}} \{L(x, u) + \nabla_x V(t, x) \cdot f(x, u)\} = 0, \quad V(T, x) = M(x). \quad (2.4)$$

This is a final-time Partial Differential Equation (PDE) that propagates the value function backwards in time. The optimal control can be recovered via [9]:

$$u^*(t, x) = \arg \min_{u \in \mathcal{U}} \{L(x, u) + \nabla_x V(t, x) \cdot f(x, u)\}. \quad (2.5)$$

The quantity inside the minimisation ($L(x, u) + \nabla_x V(t, x) \cdot f(x, u)$), is known as the Hamiltonian of the optimal control problem. In classical mechanics, the Hamiltonian represents the total energy (kinetic plus potential) of a system. In optimal control, it generalises this notion by encoding both the dynamics and the cost in a single scalar quantity that guides decision-making over time. The structure of the HJB equation is deeply connected to the Pontryagin Maximum Principle (PMP), which provides necessary conditions for optimality. PMP states

that for an optimal trajectory, there exists a costate variable such that the control input must maximise (or minimise) the Hamiltonian pointwise in time. Thus, the Hamiltonian serves as the key object linking dynamic programming and variational approaches to optimal control.

Even though the value function may not be differentiable everywhere, it can be shown that the HJB equation admits a unique viscosity solution under standard Lipschitz continuity assumptions on f , L , and M [9].

2.1.2. Robust Optimal Control and Hamilton-Jacobi-Isaacs Equations

In many real-world systems, there exist uncertainties in the form of model mismatch or external disturbances. To handle such situations, we move from standard optimal control to *robust optimal control*, which seeks control strategies that perform well under worst-case disturbances. This gives rise to a two-player zero-sum dynamic game between the controller and the disturbance [7, 10].

Consider the disturbed dynamics:

$$\dot{x}(t) = f(x(t), u(t), d(t)), \quad u(t) \in \mathcal{U}, \quad d(t) \in \mathcal{D}, \quad (2.6)$$

where $d(t)$ denotes a disturbance or adversarial input bounded in a compact set $\mathcal{D}$. The goal is now to find a control policy that minimises the cost in the presence of worst-case disturbance:

$$J_t(x; u(\cdot), d(\cdot)) = \int_t^T L(x(s), u(s), d(s)) ds + M(x(T)), \quad (2.7)$$

with the robust value function given by:

$$V(t, x) = \sup_{d(\cdot)} \inf_{u(\cdot)} J_t(x; u(\cdot), d(\cdot)). \quad (2.8)$$

This formulation is known as a zero-sum differential game, where the controller (Player 1) aims to minimise the cost while the disturbance (Player 2) aims to maximise it. Importantly, the disturbance is assumed to use *non-anticipative strategies*—meaning it cannot anticipate future values of the control input and can only react causally. This assumption balances the information available to both players during gameplay [1, 10].

Using dynamic programming principles, the value function satisfies the Hamilton-Jacobi-Isaacs (HJI) partial differential equation [10]:

$$\frac{\partial V}{\partial t}(t, x) + \min_{u \in \mathcal{U}} \max_{d \in \mathcal{D}} \{L(x, u, d) + \nabla_x V(t, x) \cdot f(x, u, d)\} = 0, \quad V(T, x) = M(x). \quad (2.9)$$

The optimal feedback strategies for both players are obtained as:

$$u^*(t, x) = \arg \min_{u \in \mathcal{U}} \max_{d \in \mathcal{D}} \{L(x, u, d) + \nabla_x V(t, x) \cdot f(x, u, d)\}, \quad (2.10)$$

$$d^*(t, x) = \arg \max_{d \in \mathcal{D}} \{L(x, u^*(t, x), d) + \nabla_x V(t, x) \cdot f(x, u^*(t, x), d)\}. \quad (2.11)$$

This robust optimal control framework forms the foundation of Hamilton-Jacobi reachability, where the goal is to reason about sets of states that can (or cannot) be reached under adversarial dynamics. In the next section, we discuss how these problems are formulated within the HJI framework.

2.1.3. Hamilton-Jacobi Reachability as a Robust Optimal Control Problem

Hamilton-Jacobi (HJ) reachability analysis arises naturally from robust optimal control theory. Instead of minimising a cumulative cost, HJ reachability focuses on whether or not the system can reach a desired target set (or avoid an undesired one), regardless of worst-case disturbances. These problems are therefore formulated as games of kind, where the outcome is binary — success or failure — rather than numerical cost [2, 7].

To formalise this, consider a system with the following dynamics:

$$\dot{x}(t) = f(x(t), u(t), d(t)), \quad u(t) \in \mathcal{U}, \quad d(t) \in \mathcal{D}, \quad (2.12)$$

and a target set $\mathcal{L} \subset \mathbb{R}^n$, represented as the subzero level set of a Lipschitz continuous function $l : \mathbb{R}^n \rightarrow \mathbb{R}$:

$$\mathcal{L} = \{x \in \mathbb{R}^n \mid l(x) \leq 0\}. \quad (2.13)$$

We define the cost function as:

$$J_t(x; u(\cdot), d(\cdot)) = l(x(T)), \quad (2.14)$$

which encodes the signed distance to the target set at the terminal time T . The value function is then given by:

$$V(t, x) = \inf_{u(\cdot)} \sup_{d(\cdot)} l(\xi_{x,t}^{u,d}(T)), \quad (2.15)$$

where $\xi_{x,t}^{u,d}(\cdot)$ denotes the system trajectory starting from x at time t under control $u(\cdot)$ and disturbance $d(\cdot)$.

The corresponding Hamilton-Jacobi-Isaacs equation is:

$$\frac{\partial V}{\partial t}(t, x) + \min_{u \in \mathcal{U}} \max_{d \in \mathcal{D}} \nabla_x V(t, x) \cdot f(x, u, d) = 0, \quad V(T, x) = l(x). \quad (2.16)$$

The subzero level set of this value function at time t ,

$$\{x \in \mathbb{R}^n \mid V(t, x) \leq 0\}, \quad (2.17)$$

defines the **Backward Reachable Set (BRS)** — the set of states from which the system can be driven into the target set by time T , despite the worst-case disturbance [1].

HJ reachability problems are often games of kind, where the outcome is binary: either the system reaches a target set or it does not. These are reformulated as games of degree by defining a continuous value function whose zero sublevel set encodes the reachability property. This transforms the binary decision problem into a continuous optimisation problem over trajectories and controls. The resulting Hamilton-Jacobi-Isaacs PDE is then solved numerically using the **level set method** [11, 12], and the solution is interpreted in the **viscosity sense** [10].

The power of this approach lies in the viscosity solution framework, which allows for the existence and uniqueness of solutions even when the value function is not differentiable [7, 13]. The level set method provides a robust computational scheme for propagating the value function on a fixed grid, without requiring explicit parameterisation of the evolving set boundary [11, 14]. This flexibility enables generalisation to reach-avoid, viability, and time-varying formulations, which we explore in the following sections.

2.1.4. Reachability and the HJI PDE

Reachability analysis seeks to determine the set of states from which a system can be guaranteed to reach (or avoid) a specified set within a given time horizon, under constrained dynamics and adversarial or cooperative interactions [1, 3]. These problems are formulated as optimal control or differential game problems, and are characterised by the **value function** $V(t, x)$, which represents the minimal cost-to-go under worst-case scenarios.

This value function is obtained as the **viscosity solution** of a corresponding **Hamilton-Jacobi-Isaacs partial differential equation**. The desired reachable or avoidable set is given by the **zero sublevel set** of the value function [10]:

$$\{x \in \mathbb{R}^n \mid V(t, x) \leq 0\} \quad (2.18)$$

Given system dynamics

$$\dot{x} = f(x, u, d), \quad u \in \mathcal{U}, \quad d \in \mathcal{D}, \quad (2.19)$$

and terminal cost $V(0, x) = l(x)$, the HJI PDE takes the form:

$$\frac{\partial V}{\partial t}(x, t) + \min_{u \in \mathcal{U}} \max_{d \in \mathcal{D}} \nabla_x V(t, x) \cdot f(x, u, d) = 0 \quad (2.20)$$

2.1.5. Variations of Reachability

Hamilton-Jacobi reachability is a highly expressive framework that supports a variety of problem formulations depending on the desired safety or performance guarantee. In this section, we explore common variations and explain how different assumptions about control and disturbance roles lead to different optimisation structures in the Hamiltonian.

Roles of the Control and Disturbance

In robust control and differential games, the controller and the disturbance play adversarial roles. Depending on the question posed, the order of infimum and supremum in the value function — and thus in the Hamiltonian — changes.

For example, in a safety problem where we wish to compute the set of states from which the disturbance can drive the system into a target set despite any control:

$$V_{\text{ABRS}}(t) = \{x \mid \forall u(\cdot), \exists d(\cdot) : \xi_{x,t}^{u,d}(T) \in \mathcal{L}\}, \quad (2.21)$$

the corresponding Hamiltonian takes the form:

$$H(x, \nabla V) = \min_{u \in \mathcal{U}} \max_{d \in \mathcal{D}} \nabla V \cdot f(x, u, d). \quad (2.22)$$

On the other hand, in a goal-satisfaction problem where the controller is trying to steer the system into a target set despite all disturbances:

$$V_{\text{BRS}}(t) = \{x \mid \exists u(\cdot), \forall d(\cdot) : \xi_{x,t}^{u,d}(T) \in \mathcal{L}\}, \quad (2.23)$$

we use:

$$H(x, \nabla V) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \nabla V \cdot f(x, u, d). \quad (2.24)$$

In general, whenever we seek a guarantee that something holds for *all* controls or disturbances, we use the outer maximisation. When we are trying to *find an existence*, the inner minimisation applies.

Reachable Sets vs Reachable Tubes

The traditional reachable set considers states from which the system can enter the target at exactly time T . This is appropriate for planning or reach-avoid problems where timing matters[15].

However, for safety applications, we are often interested in whether the system can reach the target at any point during the time horizon. This leads to the concept of the **reachable tube**, [15] which tracks reachability over the full interval $[t, T]$. The cost function becomes:

$$J_t(x; u(\cdot), d(\cdot)) = \min_{s \in [t, T]} l(x(s)), \quad (2.25)$$

and the associated value function solves a Hamilton-Jacobi variational inequality (HJI VI) [16]:

$$\min \left\{ \frac{\partial V}{\partial t} + H(x, \nabla V), l(x) - V(t, x) \right\} = 0, \quad V(T, x) = l(x). \quad (2.26)$$

Viability Kernels

Closely related to reachable tubes is the **viability kernel**, the set of states from which the system can stay within a given safe set $\mathcal{L}$ for the entire horizon. It is defined as[15]:

$$K(t) = \{x \mid \exists u(\cdot), \forall d(\cdot), \forall s \in [t, T] : \xi_{x,t}^{u,d}(s) \in \mathcal{L}\}. \quad (2.27)$$

This can be computed by solving the Backward Reachable Tube (BRT) for the complement of the safe set $\mathcal{L}^c$. Let $l(x)$ represent the signed distance to the safe set. Then the HJI VI for the viability kernel becomes [7, 15, 17]:

$$\min \left\{ \frac{\partial V}{\partial t} + H(x, \nabla V), -l(x) - V(t, x) \right\} = 0, \quad V(T, x) = -l(x), \quad (2.28)$$

and the viability kernel is recovered as:

$$K(t) = \{x \in \mathbb{R}^n \mid V(t, x) > 0\}. \quad (2.29)$$

Forward vs Backward Reachability

In contrast to **backward** reachable sets, the **forward reachable set (FRS)** describes the set of states that the system can reach from a known initial set. The HJI formulation for forward reachability is an initial value PDE [15]:

$$\frac{\partial W}{\partial t} + \sup_{u \in \mathcal{U}} \inf_{d \in \mathcal{D}} \nabla W \cdot f(x, u, d) = 0, \quad W(0, x) = l(x), \quad (2.30)$$

where $l(x)$ represents the initial condition (e.g., an indicator or distance function for the initial set). The forward reachable set at time t is:

$$V_{\text{FRS}}(t) = \{x \in \mathbb{R}^n \mid W(t, x) \leq 0\}. \quad (2.31)$$

Similarly, the **forward reachable tube (FRT)** tracks the system's evolution over the interval and can be computed via an HJI VI:

$$\max \left\{ \frac{\partial W}{\partial t} + H(x, \nabla W), W(x, t) - l(x) \right\} = 0. \quad (2.32)$$

Reach-Avoid Sets with State Constraints

Finally, many real-world problems involve reaching a target set while avoiding an unsafe set (obstacle) for the entire time horizon. This leads to the **reach-avoid problem**, defined as[15]:

$$V_{\text{RAS}}(t) = \{x \mid \exists u(\cdot), \forall d(\cdot), \exists \tau \in [t, T], \xi_{x,t}^{u,d}(\tau) \in \mathcal{G}, \xi_{x,t}^{u,d}(s) \notin \mathcal{O}, \forall s \in [t, \tau]\}. \quad (2.33)$$

This corresponds to a **double-obstacle HJI VI**, with terminal condition:

$$V(T, x) = \max\{l(x), -g(T, x)\}, \quad (2.34)$$

and dynamics governed by:

$$\max \left\{ \min \left\{ \frac{\partial V}{\partial t} + H(x, \nabla V, l(x) - V(t, x)), -g(t, x) - V(t, x) \right\}, -g(t, x) - V(t, x) \right\} = 0. \quad (2.35)$$

The level set framework allows elegant encoding of both time-varying targets and obstacles, enabling general and robust analysis of reachability under complex constraints.

2.1.6. Formal Definitions of Reachable Sets and HJI Formulation

Let $\xi_{x,t}^{u,d}(\tau)$ denote the trajectory of the system starting from state x at time t under control $u(\cdot)$ and disturbance $d(\cdot)$.

Avoid Backward Reachable Set (Avoid-BRS): The Avoid-BRS is the set of all states from which the system can be forced into a target set $\mathcal{L} = \{x \mid l(x) \leq 0\}$ at the final time T , regardless of the control action [15].

$$V_{\text{BRS}}(t) = \{x \mid \exists d(\cdot), \forall u(\cdot), \xi_{x,t}^{u,d}(T) \in \mathcal{L}\} \quad (2.36)$$

Value function:

$$V(t, x) = \sup_{u(\cdot)} \inf_{d(\cdot)} l(\xi_{x,t}^{u,d}(T)) \quad (2.37)$$

HJI equation:

$$\frac{\partial V}{\partial t} + \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \nabla V \cdot f(x, u, d) = 0 \quad (2.38)$$

Avoid Backward Reachable Tube (Avoid-BRT): The Avoid-BRT includes all states from which the system can be forced into the target set at some time $\tau \in [t, T]$ [15].

$$V_{\text{BRT}}(t) = \{x \mid \exists d(\cdot), \forall u(\cdot), \exists \tau \in [t, T] : \xi_{x,t}^{u,d}(\tau) \in \mathcal{L}\} \quad (2.39)$$

Value function:

$$V(t, x) = \sup_{u(\cdot)} \inf_{d(\cdot)} \min_{\tau \in [t, T]} l(\xi_{x,t}^{u,d}(\tau)) \quad (2.40)$$

HJI variational inequality:

$$\min \left\{ \frac{\partial V}{\partial t} + H(x, \nabla V), V(t, x) - l(x) \right\} = 0 \quad (2.41)$$

BRS: The BRS is the set of states from which the system is guaranteed to enter a set $\mathcal{O} = \{x \mid o(x) \leq 0\}$ at time T , despite any disturbance [15].

$$V_{\text{ABRS}}(t) = \{x \mid \exists u(\cdot), \forall d(\cdot), \xi_{x,t}^{u,d}(T) \in \mathcal{O}\} \quad (2.42)$$

Value function:

$$V(t, x) = \inf_{u(\cdot)} \sup_{d(\cdot)} o(\xi_{x,t}^{u,d}(T)) \quad (2.43)$$

HJI equation:

$$\frac{\partial V}{\partial t} + \min_{u \in \mathcal{U}} \max_{d \in \mathcal{D}} \nabla V \cdot f(x, u, d) = 0 \quad (2.44)$$

BRT: This set includes all states from which the system is guaranteed to enter the set at some time in $[t, T]$ [15].

$$V_{\text{ABRT}}(t) = \{x \mid \forall u(\cdot), \exists d(\cdot), \exists \tau \in [t, T] : \xi_{x,t}^{u,d}(\tau) \in \mathcal{O}\} \quad (2.45)$$

Value function:

$$V(t, x) = \inf_{u(\cdot)} \sup_{d(\cdot)} \min_{\tau \in [t, T]} o(\xi_{x,t}^{u,d}(\tau)) \quad (2.46)$$

HJI variational inequality:

$$\min \left\{ \frac{\partial V}{\partial t} + H(x, \nabla V), V(t, x) - o(x) \right\} = 0 \quad (2.47)$$

Reach-Avoid BRT: This set consists of all states from which the system can be driven to a goal set $\mathcal{G} = \{x \mid g(x) \leq 0\}$ at any time $\tau \in [t, T]$, while avoiding an unsafe set $\mathcal{O}$ throughout the trajectory [15].

$$V_{\text{RAT}}(t) = \{x \mid \exists u(\cdot), \forall d(\cdot), \exists \tau \in [t, T], \xi_{x,t}^{u,d}(\tau) \in \mathcal{G}, \xi_{x,t}^{u,d}(s) \notin \mathcal{O}, \forall s \in [t, \tau]\} \quad (2.48)$$

Value function:

$$V(t, x) = \inf_{u(\cdot)} \sup_{d(\cdot)} \min_{\tau \in [t, T]} \max \left\{ g(\xi_{x,t}^{u,d}(\tau)), \max_{s \in [t, \tau]} o(\xi_{x,t}^{u,d}(s)) \right\} \quad (2.49)$$

HJI variational inequality:

$$\max \left\{ \min \left\{ \frac{\partial V}{\partial t} + H, V(t, x) - g(x) \right\}, V(t, x) - o(x) \right\} = 0 \quad (2.50)$$

Forward Reachable Set (FRS): The FRS contains all states that the system can reach at time T from a given initial condition x_0 , under control $u(\cdot)$ and against adversarial disturbance $d(\cdot)$ [15].

$$V_{\text{FRS}}(T) = \{x \mid \exists u(\cdot), \forall d(\cdot) : \xi_{x_0, t}^{u, d}(T) = x\} \quad (2.51)$$

Value function:

$$W(x, t) = \inf_{u(\cdot)} \sup_{d(\cdot)} \phi(\xi_{x_0, t}^{u, d}(T)) \quad (2.52)$$

HJI PDE (forward-time, with adversarial disturbance):

$$\frac{\partial W}{\partial t} + \min_{u \in \mathcal{U}} \max_{d \in \mathcal{D}} \nabla W \cdot f(x, u, d) = 0, \quad W(x, 0) = \phi(x) \quad (2.53)$$

where $\phi(x)$ is an initial condition encoding the initial state x_0 , often a signed distance or indicator function.

Forward Reachable Tube (FRT): The FRT contains all states that the system can reach at some time within the time horizon $[t, T]$, under admissible control and despite adversarial disturbance [15].

$$V_{\text{FRT}}(t) = \{x \mid \exists u(\cdot), \forall d(\cdot), \exists \tau \in [t, T] : \xi_{x_0, t}^{u, d}(\tau) = x\} \quad (2.54)$$

Value function:

$$W(x, t) = \inf_{u(\cdot)} \sup_{d(\cdot)} \min_{\tau \in [t, T]} \phi(\xi_{x_0, t}^{u, d}(\tau)) \quad (2.55)$$

HJI Variational Inequality:

$$\min \left\{ \frac{\partial W}{\partial t} + \min_u \max_d \nabla W \cdot f(x, u, d), W(x, t) - \phi(x) \right\} = 0 \quad (2.56)$$

Reachable Set (Reach Set): The reach set generally refers to all states that can be reached by the system from an initial condition, and is commonly associated with the forward reachable set. In safety analysis, it may be interpreted as the complement of an avoid set. It is a core object of interest in control synthesis, verification, and exploration [7, 18]. The example in the following subsection serves as an illustration of the section above.

2.1.7. Example: Double Integrator and Backward Reachable Tube (BRT)

Consider the one-dimensional double integrator system:

$$\dot{x} = v, \quad \dot{v} = u, \quad u \in [-1, 1] \quad (2.57)$$

where:

- $x \in \mathbb{R}$ is the position,
- $v \in \mathbb{R}$ is the velocity,
- $u \in [-1, 1]$ is the control input.

The state vector is represented as $s = (x, v) \in \mathbb{R}^2$, and the system dynamics can be written compactly as:

$$\dot{s} = \begin{bmatrix} v \\ u \end{bmatrix} \quad (2.58)$$

We define the target set $\mathcal{L} \subset \mathbb{R}^2$ as a ball of radius $R > 0$ centered at the origin:

$$\mathcal{L} = \{(x, v) \in \mathbb{R}^2 \mid x^2 + v^2 \leq R^2\} \quad (2.59)$$

The corresponding target function (or terminal cost) is defined as:

$$l(x, v) = \sqrt{x^2 + v^2} - R \quad (2.60)$$

This function is negative inside the target set, zero on its boundary, and positive outside.

We are interested in computing the BRT: the set of all initial states from which the system can be driven into the target set $\mathcal{L}$ at some time $\tau \in [t, T]$, using admissible control inputs.

The value function that encodes this reachability information is given by:

$$V(x, v, t) = \sup_{u(\cdot)} \inf_{d(\cdot)} \min_{\tau \in [t, T]} l(\xi_{(x, v), t}^{u, d}(\tau)) \quad (2.61)$$

where:

- $\xi_{(x, v), t}^{u, d}(\tau)$ is the state trajectory starting from (x, v) at time t ,
- $u(\cdot)$ is the control policy,
- $d(\cdot)$ is the disturbance (can be ignored or set to 0 in this example).

The associated HJI variational inequality that characterises this value function is:

$$\min \left\{ \frac{\partial V}{\partial t} + H(x, v, \nabla V), V(x, v, t) - l(x, v) \right\} = 0 \quad (2.62)$$

The Hamiltonian is defined as:

$$H(x, v, p) = \max_{u \in [-1, 1]} (p_x v + p_v u) \quad (2.63)$$

where $p = (p_x, p_v) = \nabla V$ is the gradient (costate vector) of the value function.

This running example demonstrates how the BRT captures the evolution of the reachable set over time. As the time horizon T increases, the reachable tube expands to include more initial states from which it is possible to steer

the system into the target region $\mathcal{L}$. This is especially useful in verifying the temporal flexibility of reaching goals in safety-critical systems.




Figure 2.1: HJB Reachability logic flow: Both optimisation and HJI PDE yield the value function, whose zero level set defines the reach/avoid set.

2.2. BACKGROUND ON NEURAL NETWORKS

Neural networks (NNs) are a class of machine learning models inspired by the structure and function of biological neural systems. Over the past decade, they have become central to modern artificial intelligence, with applications ranging from image recognition and natural language processing to robotics and control. Neural networks are particularly effective due to their ability to approximate highly nonlinear functions from data, making them ideal for systems with complex dynamics and high-dimensional inputs.

Note - all the background on Neural Networks presented in the thesis is based on the [19]

2.2.1. Feedforward Neural Networks

A *feedforward neural network* (FNN), also known as a multilayer perceptron (MLP), is one of the most commonly used NN architectures. It is composed of an input layer, one or more hidden layers, and an output layer. Information flows through the network in a single direction—from inputs to outputs—without forming cycles.

Formally, for an input vector $\mathbf{x} \in \mathbb{R}^n$, a feedforward neural network with N layers (including the output layer) computes the output as:

$$f(\mathbf{x}; \theta) = \phi_N \circ \phi_{N-1} \circ \cdots \circ \phi_1(\mathbf{x}), \quad (2.64)$$

where θ denotes the collection of all trainable parameters in the network, operator $\circ$ represents composition of functions, and each ϕ_i corresponds to the operation of layer i :

$$\phi_i(\mathbf{z}) = \sigma_i(W_i \mathbf{z} + \mathbf{b}_i), \quad i = 1, \dots, N-1. \quad (2.65)$$

Here, $W_i \in \mathbb{R}^{m_i \times m_{i-1}}$ is the weight matrix, $\mathbf{b}_i \in \mathbb{R}^{m_i}$ is the bias vector, and σ_i is a nonlinear activation function (e.g., ReLU, sigmoid, or tanh) applied elementwise. For the final layer (often a regression or classification output), we typically apply no activation:

$$\phi_N(\mathbf{z}) = W_N \mathbf{z} + \mathbf{b}_N. \quad (2.66)$$

The ReLU (Rectified Linear Unit) activation is a popular choice for hidden layers, defined as:

$$\text{ReLU}(z) = \max(0, z), \quad \text{applied elementwise.} \quad (2.67)$$

Example: A neural network with a single hidden layer (i.e., two layers in total) has the form:

$$f(\mathbf{x}; \theta) = W_2 \sigma(W_1 \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2, \quad (2.68)$$

where σ is the activation function (e.g., ReLU or tanh), and the output is a linear combination of the hidden units.

Geometrically, each hidden unit computes a hyperplane partitioning of the input space, and applies a nonlinear transformation. Together, the network partitions the input space into polytopes, and within each region the function is affine. For ReLU networks, the output is a piecewise linear function.

2.2.2. Universal Approximation Theorem

A foundation underpinning the use of neural networks is the *universal approximation theorem*, which states that a feedforward network with a single hidden layer containing a finite number of neurons can approximate any continuous function on a compact subset of $\mathbb{R}^n$, given sufficient width and appropriate weights.

Formally, let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be continuous, and let $\epsilon > 0$. Then there exists a neural network f_{NN} such that:

$$\sup_{\mathbf{x} \in K} |f(\mathbf{x}) - f_{\text{NN}}(\mathbf{x})| < \epsilon, \quad (2.69)$$

for any compact set $K \subset \mathbb{R}^n$.

2.2.3. Training Neural Networks

Training a neural network involves adjusting the parameters $\theta = \{W_i, \mathbf{b}_i\}_{i=1}^N$ to minimise a loss function over a dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^M$. This is typically

formulated as an optimisation problem:

$$\theta^* = \arg \min_{\theta} \frac{1}{M} \sum_{i=1}^M \mathcal{L}(f(\mathbf{x}_i; \theta), \mathbf{y}_i), \quad (2.70)$$

where $\mathcal{L}$ is a task-dependent loss function. For regression tasks, a common choice is the *mean squared error* (MSE):

$$\mathcal{L}_{\text{MSE}} = \frac{1}{M} \sum_{i=1}^M \|f(\mathbf{x}_i; \theta) - \mathbf{y}_i\|^2. \quad (2.71)$$

Optimisation is typically performed via stochastic gradient descent (SGD) or its variants (e.g., Adam, RMSProp), using backpropagation to compute gradients efficiently via the chain rule.

2.2.4. Relevance to Control and Robotics

In robotics and control, neural networks are increasingly used to approximate value functions, policies, and system dynamics. Their expressiveness allows them to capture complex, nonlinear mappings—such as from sensory input to control actions or from system states to safety margins—especially when analytical models are unavailable or intractable. Recent advances have further combined NNs with tools from control theory (e.g., Lyapunov analysis [20], reachability [21]) to ensure safe and stable operation.

2.3. MATHEMATICAL FOUNDATIONS OF SATISFIABILITY MODULO THEORIES (SMT)

The *Satisfiability Modulo Theories* (SMT) problem generalises the classical Boolean satisfiability problem (SAT) to richer domains by allowing reasoning over *background theories*, such as arithmetic, arrays, and data structures. Formally, the SMT problem asks whether a first-order logic (FOL) formula φ is satisfiable with respect to a background theory $\mathcal{T}$:

$$\text{SMT}(\varphi, \mathcal{T}) := \exists \sigma \in \text{Model}(\mathcal{T}) \text{ such that } \sigma \models \varphi \quad (2.72)$$

Here, φ is a logical formula, $\mathcal{T}$ defines the interpretation of function and predicate symbols, and σ is a variable assignment consistent with both φ and $\mathcal{T}$.

Note - all the background on the SMT presented in the thesis is based on [22].

2.3.1. Preliminaries and Formal Definitions

Let Σ be a signature consisting of function symbols $\mathcal{F}$, predicate symbols $\mathcal{P}$, and constant symbols $\mathcal{C}$. A first-order formula $\varphi \in \text{FOL}[\Sigma]$ is constructed via:

- Terms: $t ::= c \mid x \mid f(t_1, \dots, t_n)$,

- Formulas: $\varphi ::= t_1 = t_2 \mid \rho(t_1, \dots, t_n) \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid \forall x. \varphi \mid \exists x. \varphi$.

Given a theory $\mathcal{T}$ with a model class $\mathcal{M}_{\mathcal{T}}$, the SMT problem becomes:

$$\text{SMT}(\varphi, \mathcal{T}) := \text{Is there } \mathcal{M} \in \mathcal{M}_{\mathcal{T}} \text{ such that } \mathcal{M} \models \varphi? \quad (2.73)$$

Example: Let $\varphi = (x > 3) \wedge (x < 1)$, and $\mathcal{T} = \text{Th}(\mathbb{R})$, the theory of real arithmetic. Then $\text{SMT}(\varphi, \mathcal{T}) = \text{UNSAT}$, since no real x satisfies both constraints.

2

2.3.2. Common Background Theories

The power of SMT arises from its support for various decidable theories:

Theory $\mathcal{T}$	Domain	Description
$\mathcal{T}_{\text{EUF}}$	Uninterpreted functions	Equality logic with uninterpreted functions
$\mathcal{T}_{\text{LIA}}$	Integers	Linear Integer Arithmetic
$\mathcal{T}_{\text{LRA}}$	Reals	Linear Real Arithmetic
$\mathcal{T}_{\text{BV}}$	Bit-vectors	Fixed-size word-level arithmetic
$\mathcal{T}_{\text{Arrays}}$	Arrays	Memory indexing with <code>select/store</code>
$\mathcal{T}_{\text{UF} + \text{RA}}$	Combined	Uninterpreted + real arithmetic
$\mathcal{T}_{\text{NRA}}$	Nonlinear reals	e.g., $x \cdot y, \sin(x)$ (often undecidable)

Table 2.1: Typical SMT background theories

Some theories are decidable (e.g., linear arithmetic), while others are only δ -decidable—solvable up to numerical error δ (e.g., nonlinear real arithmetic with transcendental functions as supported by `dReal`).

2.3.3. The DPLL(T) Framework

Modern SMT solvers use a modular architecture known as **DPLL(T)**. It combines:

1. **Boolean abstraction:** Replace each theory atom with a propositional variable to obtain a Boolean skeleton:

$$\varphi^{\text{bool}} = A_1 \wedge (A_2 \rightarrow A_3) \wedge \neg A_4 \wedge \dots \quad (2.74)$$

2. **SAT solving:** A SAT solver searches for a satisfying assignment $\nu : \{A_i\} \rightarrow \{\text{true}, \text{false}\}$.
3. **Theory consistency check:** The theory solver checks the theory literals under the current assignment:

$$\mathcal{T} \models A_i \text{ if } \nu(A_i) = \text{true} \quad (2.75)$$

If inconsistent, a conflict clause is added to guide backtracking.

A theory solver is a decision procedure:

$$\mathcal{S}_{\mathcal{T}} : \text{Set of literals} \rightarrow \{\text{SAT}, \text{UNSAT}\} \quad (2.76)$$

using techniques such as:

- Simplex algorithm (for $\mathcal{T}_{\text{LRA}}$),
- Fourier-Motzkin elimination (for quantified constraints),
- Congruence closure (for $\mathcal{T}_{\text{EUF}}$),
- Interval constraint propagation (for $\mathcal{T}_{\text{NRA}}$).

2.3.4. Model-Theoretic Semantics

Let $\mathcal{L}$ be a first-order language and $\mathcal{T} = \langle \mathcal{L}, \mathcal{A} \rangle$ be a theory with axioms $\mathcal{A}$. The SMT decision problem becomes:

$$\mathcal{T} \models \exists \mathbf{x}. \varphi(\mathbf{x}) \iff \text{Satisfiability of } \mathcal{T} \cup \{\varphi\} \quad (2.77)$$

That is, SMT determines if φ is satisfiable in any model of $\mathcal{T}$.

2.3.5. Applications in Formal Methods and Control

SMT solving is foundational in:

- **Program verification** and static analysis,
- **Symbolic execution** and test generation,
- **Formal synthesis** from temporal logic specifications,
- **Verification of hybrid systems** and reachability analysis,
- **Motion planning** under logical and geometric constraints,
- **Control synthesis** and falsification in cyber-physical systems.

For example, in reachability analysis for hybrid systems with dynamics $\dot{\mathbf{x}} = f(\mathbf{x}, u)$ and unsafe region $\mathcal{X}_{\text{unsafe}}$, we encode:

$$\text{SMT} \left(\bigwedge_{t=0}^T \mathbf{x}_t \notin \mathcal{X}_{\text{unsafe}} \wedge \mathbf{x}_{t+1} = f(\mathbf{x}_t, u_t), \mathcal{T}_{\text{NRA}} \right) \quad (2.78)$$

Such encodings are solvable using dReal [23] or Z3 [24] with nonlinear arithmetic capabilities. In summary, SMT provides a unifying logic-based formalism for reasoning about real-world systems with arithmetic, functional, and structural constraints. It serves as a powerful tool for verification, synthesis, and analysis in embedded, robotic, and safety-critical systems.

2.4. MEASURE THEORY

Measure theory was developed to provide a rigorous and general framework for dealing with the concepts of length, area, volume, and integration, especially in cases where classical methods like Riemann integration fail. In classical calculus, integration is defined over intervals using sums of function values over subdivisions of the domain. However, this approach breaks down for functions that are highly discontinuous or for domains that are not simple intervals.

Lebesgue's insight was to define the integral by slicing the *range* of the function rather than the *domain*, and summing up the measures of the level sets where the function takes certain values. This new perspective allows us to integrate a much wider class of functions, including those that are discontinuous on dense subsets of their domain.

Measure theory is indispensable in modern mathematics, particularly in probability theory, where we interpret probability as a measure. It also plays a central role in real analysis, functional analysis, ergodic theory, and the study of partial differential equations. Without measure theory, many tools used in modern analysis, such as L^p spaces and weak convergence, would not have a rigorous foundation.

Note - all the measure theory background presented in the thesis is based on [25]

2.4.1. Main Concepts in Measure Theory

To understand measure theory, we begin by building a framework that allows us to assign a meaningful notion of “size” or “volume” to sets, even those that are irregular or infinite, and to integrate functions that may be too discontinuous for classical (Riemann) methods. This framework is constructed using four central ideas: σ -algebras, measures, measurable functions, and the Lebesgue integral.

A **σ -algebra** $\mathcal{F}$ over a set X is a collection of subsets of X that satisfies three properties:

1. $X \in \mathcal{F}$,
2. If $A \in \mathcal{F}$, then the complement $A^c \in \mathcal{F}$,
3. If $A_1, A_2, A_3, \dots \in \mathcal{F}$, then $\bigcup_{n=1}^{\infty} A_n \in \mathcal{F}$.

This structure ensures that we can perform standard set operations (like unions and complements) without leaving the space of measurable sets. For example, the collection of all open intervals on $\mathbb{R}$ generates the *Borel σ -algebra*, which includes most sets used in real analysis.

A **measure** μ is a function that assigns a non-negative extended real number to each set in a σ -algebra. It must satisfy:

1. $\mu(\emptyset) = 0$,
2. If $A_1, A_2, \dots$ are disjoint sets in $\mathcal{F}$, then

$$\mu\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} \mu(A_n). \quad (2.79)$$

This property is called *countable additivity*, and it ensures that the measure behaves predictably under infinite operations. A canonical example is the *Lebesgue measure* on $\mathbb{R}$, which assigns length to intervals: for instance, $\mu([a, b]) = b - a$.

A **measurable function** is a function $f : X \rightarrow \mathbb{R}$ such that for every $a \in \mathbb{R}$, the set $\{x \in X : f(x) < a\} \in \mathcal{F}$. This condition ensures we can integrate f with respect to the measure. It generalises the idea of continuity: all continuous functions are measurable, but measurable functions can be much more irregular.

The **Lebesgue integral** defines integration using the measure of the function's value-level sets. Instead of summing over partitions of the domain, as in Riemann integration, the Lebesgue integral slices the range of the function. Formally, if f is a non-negative measurable function, the integral is defined as:

$$\int_X f d\mu = \sup \left\{ \int_X s d\mu : 0 \leq s \leq f, s \text{ is simple} \right\}, \quad (2.80)$$

where a **simple function** is a finite linear combination of indicator functions:

$$s(x) = \sum_{i=1}^n a_i \chi_{A_i}(x), \quad a_i \geq 0, A_i \in \mathcal{F}. \quad (2.81)$$

This approach is much more powerful. For example, let

$$f(x) = \chi_{\mathbb{Q} \cap [0, 1]}(x), \quad (2.82)$$

i.e., $f(x) = 1$ if $x \in \mathbb{Q} \cap [0, 1]$, and $f(x) = 0$ otherwise. This function is not Riemann integrable because it is discontinuous everywhere. But since the rationals in $[0, 1]$ have *Lebesgue measure zero*, we have:

$$\int_0^1 f(x) dx = 0, \quad (2.83)$$

showing that $f(\cdot)$ is Lebesgue integrable.

Monotone Convergence Theorem (MCT)

Theorem. Let $\{f_n\}_{n=1}^\infty$ be a sequence of measurable functions such that

$$0 \leq f_1(x) \leq f_2(x) \leq \dots \quad \text{for all } x \in X, \quad (2.84)$$

and define $f(x) = \lim_{n \rightarrow \infty} f_n(x)$. Then

$$\lim_{n \rightarrow \infty} \int_X f_n d\mu = \int_X f d\mu. \quad (2.85)$$

This theorem tells us that when a sequence of non-negative measurable functions increases pointwise to a limit, then the integrals of those functions also converge to the integral of the limiting function. It allows us to approximate

complex functions using simpler ones, building them from below. MCT is fundamental to the construction of the Lebesgue integral. It ensures that integration is well-behaved under limits, especially when we use simple functions to approximate more general ones. This theorem is particularly useful when defining integrals or when working with sequences of truncated or bounded functions. It guarantees convergence of integrals under monotonic growth.

Example. Let

$$f_n(x) = \min(x, n), \quad x \in [0, \infty). \quad (2.86)$$

Then $f_n(x) \uparrow x$, and for a fixed $k > 0$,

$$\int_0^k f_n(x) dx = \int_0^k \min(x, n) dx \rightarrow \int_0^k x dx = \frac{1}{2}k^2. \quad (2.87)$$

By MCT, we have:

$$\lim_{n \rightarrow \infty} \int_0^k f_n(x) dx = \int_0^k x dx. \quad (2.88)$$

Fatou's Lemma

Theorem. Let $\{f_n\}_{n=1}^\infty$ be a sequence of non-negative measurable functions. Then:

$$\int_X \liminf_{n \rightarrow \infty} f_n d\mu \leq \liminf_{n \rightarrow \infty} \int_X f_n d\mu. \quad (2.89)$$

Fatou's Lemma provides a lower bound on the integral of the limit inferior of a sequence of non-negative functions. It tells us that even if the sequence does not converge nicely, the integral of the $\liminf$ cannot be larger than the $\liminf$ of the integrals. It is a foundational result that allows us to deal with convergence problems in integration, especially when we do not have monotonicity. It often serves as a tool in existence proofs or in proving inequalities involving integrals. Fatou's Lemma is commonly used when studying the convergence of integrals or expectations in real analysis, probability, or functional analysis. It helps establish inequality relations in weak convergence arguments.

Example. Let:

$$f_n(x) = \frac{1}{n} \chi_{[n, n+1]}(x), \quad (2.90)$$

defined on $\mathbb{R}$. Then for any fixed x , we have $\liminf f_n(x) = 0$. Hence,

$$\int_{\mathbb{R}} \liminf f_n(x) dx = \int_{\mathbb{R}} 0 dx = 0. \quad (2.91)$$

Each individual integral is:

$$\int_{\mathbb{R}} f_n(x) dx = \frac{1}{n} \rightarrow 0. \quad (2.92)$$

So we get:

$$\int \liminf f_n \leq \liminf \int f_n, \quad (2.93)$$

with both sides equal to zero, verifying Fatou's Lemma in this case.

2

Dominated Convergence Theorem (DCT)

Theorem. Let $\{f_n\}_{n=1}^{\infty}$ be a sequence of measurable functions such that:

- $f_n(x) \rightarrow f(x)$ pointwise almost everywhere on X ,
- and there exists an integrable function $g \in L^1(\mu)$ such that $|f_n(x)| \leq g(x)$ for all n and almost every x .

Then:

$$\lim_{n \rightarrow \infty} \int_X f_n d\mu = \int_X f d\mu. \quad (2.94)$$

The Dominated Convergence Theorem states that if a sequence of functions converges pointwise to a limit and all the functions are bounded in absolute value by a single integrable function, then we can safely interchange the limit and the integral. DCT is widely used when dealing with convergence of expectations or integrals, especially in contexts where we lack monotonicity. It offers a practical and general condition under which limits and integrals commute. It is essential in probability theory, PDEs, and functional analysis, especially when verifying convergence of expectations, solutions, or likelihood functions.

Example. Define:

$$f_n(x) = x^n \quad \text{on } [0, 1]. \quad (2.95)$$

Then $f_n(x) \rightarrow 0$ for $x \in [0, 1)$, and $f_n(1) = 1$, so:

$$f(x) = \begin{cases} 0, & \text{if } x \in [0, 1), \\ 1, & \text{if } x = 1. \end{cases} \quad (2.96)$$

We can dominate all f_n by the constant function $g(x) = 1$, which is integrable. Hence, by DCT:

$$\lim_{n \rightarrow \infty} \int_0^1 f_n(x) dx = \int_0^1 f(x) dx = 1. \quad (2.97)$$

Fubini's Theorem

Theorem. Let $(X, \mathcal{F}, \mu)$ and $(Y, \mathcal{G}, \nu)$ be σ -finite measure spaces, and let $f : X \times Y \rightarrow \mathbb{R}$ be a measurable function such that:

$$\int_{X \times Y} |f(x, y)| d(\mu \times \nu)(x, y) < \infty. \quad (2.98)$$

Then:

$$\int_{X \times Y} f(x, y) d(\mu \times \nu)(x, y) = \int_X \left(\int_Y f(x, y) d\nu(y) \right) d\mu(x) \quad (2.99)$$

$$= \int_Y \left(\int_X f(x, y) d\mu(x) \right) d\nu(y) \quad (2.100)$$

Fubini's Theorem lets us compute a double integral by integrating one variable at a time. Provided the function is absolutely integrable on the product space, we can change the order of integration freely. This theorem is foundational for multivariable integration, joint probability distributions, and product measures. It allows us to simplify complex integrals and gives rigorous justification for separating variables. Fubini's Theorem is essential in probability when computing expectations of jointly distributed random variables. It is also used in PDEs and multivariate calculus when transforming integrals.

Example. Let:

$$f(x, y) = e^{-x-y} \quad \text{on } [0, \infty)^2. \quad (2.101)$$

We compute:

$$\int_0^\infty \int_0^\infty e^{-x-y} dy dx = \left(\int_0^\infty e^{-x} dx \right) \left(\int_0^\infty e^{-y} dy \right) = 1 \cdot 1 = 1. \quad (2.102)$$

Since $f \geq 0$ and integrable, Fubini's Theorem applies and confirms the result.

Carathéodory's Criterion for Measurability

Lemma. Let μ^* be an outer measure on $\mathbb{R}^n$. A set $E \subseteq \mathbb{R}^n$ is Lebesgue measurable if and only if for every set $A \subseteq \mathbb{R}^n$,

$$\mu^*(A) = \mu^*(A \cap E) + \mu^*(A \setminus E). \quad (2.103)$$

Explanation. This lemma provides a formal criterion to test whether a set behaves well with respect to the outer measure. If the measure of any set A can be broken into the measure of its parts inside and outside E , then E is measurable.

Why it's important. Carathéodory's criterion is the basis for defining Lebesgue measurable sets. Though rarely applied directly, it underpins most of the theory of measurable sets.

Use case. We use it implicitly every time we assert that unions, intersections, or complements of measurable sets are measurable.

Example. Any open interval $(a, b) \subset \mathbb{R}$ is measurable. Since the collection of open intervals generates the Borel σ -algebra, and the Lebesgue measure is complete, all Borel sets are measurable.

Radon–Nikodym Theorem

Theorem. Let ν and μ be σ -finite measures on a measurable space $(X, \mathcal{F})$, and suppose $\nu \ll \mu$, that is:

$$\mu(A) = 0 \Rightarrow \nu(A) = 0 \quad \text{for all } A \in \mathcal{F}. \quad (2.104)$$

Then there exists a measurable function $f : X \rightarrow [0, \infty)$ such that for every measurable set A ,

$$\nu(A) = \int_A f \, d\mu. \quad (2.105)$$

This function f is called the *Radon–Nikodym derivative* of ν with respect to μ , and is written:

$$f = \frac{d\nu}{d\mu}. \quad (2.106)$$

The theorem tells us that if one measure is absolutely continuous with respect to another, then it must arise from the other via integration of some density function. This derivative $\frac{d\nu}{d\mu}$ acts like a weight that locally relates ν to μ . The Radon–Nikodym Theorem is essential in probability theory for defining densities of distributions. It underlies concepts like likelihood functions, conditional probabilities, and change-of-measure techniques in stochastic processes. If a probability distribution $\mathbb{P}_X$ has a density with respect to Lebesgue measure λ , then:

$$\mathbb{P}(X \in A) = \int_A f(x) \, d\lambda(x), \quad (2.107)$$

where $f = \frac{d\mathbb{P}_X}{d\lambda}$.

Example. Let μ be Lebesgue measure on $[0, 1]$, and define:

$$\nu(A) = \int_A 3x^2 \, d\mu. \quad (2.108)$$

Then $\nu \ll \mu$, and the Radon–Nikodym derivative is:

$$\frac{d\nu}{d\mu}(x) = 3x^2. \quad (2.109)$$

Thus, for any measurable set $A \subseteq [0, 1]$,

$$\nu(A) = \int_A 3x^2 \, d\mu(x). \quad (2.110)$$

Why and How Do We Prove a Set is Measurable?

Measurable sets are essential in measure theory because integration is only defined over such sets. In probability theory, events must be measurable in

order for probabilities to be well-defined. Similarly, a function is considered measurable only if the pre-image of every Borel set is measurable. Without ensuring that sets and functions are measurable, the core theorems of measure theory — including the Monotone Convergence Theorem, Fatou's Lemma, and the Dominated Convergence Theorem — cannot be properly applied.

There are two standard methods to prove that a set is measurable. One is to use Carathéodory's criterion directly, which characterises measurable sets in terms of their compatibility with outer measure. The other, more common method, involves constructing measurable sets using known measurable sets (like open or closed intervals) and applying the closure properties of σ -algebras, such as countable unions, intersections, and complements. In practice, most constructions rely on the fact that Borel sets — generated from open intervals through these operations — are measurable with respect to the Lebesgue measure.

2.5. PROBABILITY THEORY: FUNDAMENTALS

Probability theory is deeply rooted in measure theory and provides a mathematical language for quantifying uncertainty. It is essential in systems where outcomes are inherently unpredictable, such as robotics, finance, and engineering systems.

Note - all the probability theory background presented in the thesis is based on [26]

2.5.1. Probability Space and Random Variables

A probability space provides a formal structure for mathematically modeling random experiments. It comprises three essential components: a sample space (Ω), which includes all possible outcomes; a sigma-algebra ($\mathcal{F}$), a structured collection of events (subsets of the sample space) that satisfies certain mathematical conditions such as closure under complementation and countable unions; and a probability measure (P), which assigns a numerical probability to each event, respecting axioms like non-negativity, normalization (total probability equals one), and countable additivity for disjoint events. This framework ensures that probabilistic analyses are rigorous and well-defined. A **probability space** is formally defined as a triple $(\Omega, \mathcal{F}, P)$ where:

- **Sample Space (Ω):** Set of all possible outcomes.
- **Sigma-Algebra ($\mathcal{F}$):** Collection of subsets of Ω , closed under complements and countable unions.
- **Probability Measure (P):** A measure satisfying:

$$P(\Omega) = 1, \quad P\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} P(A_i) \quad (2.111)$$

A random variable is a function mapping the outcomes of a random experiment to numerical values. By assigning numbers to outcomes, random variables allow for quantitative analyses of random phenomena, simplifying calculations and

interpretations. They form the cornerstone of probability theory, enabling us to succinctly describe events in terms of numbers. For example, when flipping a coin, a random variable could represent the number of heads obtained, converting qualitative outcomes (heads or tails) into quantitative values (0 or 1). A random variable X is a measurable function:

$$X : \Omega \rightarrow \mathbb{R}, \quad X^{-1}(B) \in \mathcal{F}, \quad \forall B \in \mathcal{B}(\mathbb{R}) \quad (2.112)$$

2.5.2. Distribution Functions and Probability Densities

The cumulative distribution function (CDF) describes the probability that a random variable takes a value less than or equal to a specific value. The CDF fully characterizes the distribution of a random variable and is used extensively in probability and statistics to summarize probabilities succinctly. For instance, if the random variable represents the outcome of rolling a die, then the CDF at 3 would indicate the probability of rolling a 3 or less. Formally **CDF** of random variable X is given as:

$$F_X(x) = P(X \leq x) \quad (2.113)$$

A probability density function (PDF) applies to continuous random variables and provides a detailed description of how probability is distributed over continuous intervals. Intuitively, the PDF gives the relative likelihood of the random variable being found at a particular point. Areas under the PDF curve between two points indicate the probability of the variable falling within that interval. Formally **PDF** for continuous random variables is given as:

$$F_X(x) = \int_{-\infty}^x f_X(t)dt, \quad f_X(x) = \frac{d}{dx}F_X(x) \quad (2.114)$$

2.5.3. Expectation and Variance

The expectation or mean of a random variable provides a measure of its central tendency or "average" value if the experiment is repeated infinitely many times. Mathematically, it is computed as the integral (or sum, in the discrete case) of the possible outcomes weighted by their probabilities. The expectation is a fundamental concept in probability and statistics, offering an intuitive measure of what one can typically expect to observe. Formally we find expectation as:

$$\mathbb{E}[X] = \int_{\Omega} X(\omega)dP(\omega) = \int_{-\infty}^{\infty} xf_X(x)dx \quad (2.115)$$

The variance of a random variable measures how spread out the values of that variable are around its expectation. Formally, it's the expectation of the squared deviation of the random variable from its mean. A large variance implies outcomes are widely dispersed from the mean, while a small variance means outcomes cluster closely around it. Variance is crucial for quantifying uncertainty

and variability in applications ranging from finance to robotics. Formally we find variance as:

$$\text{Var}(X) = \mathbb{E}[X^2] - \mathbb{E}[X]^2 \quad (2.116)$$

2.5.4. Important Inequalities

Markov's inequality provides a fundamental bound relating the probability that a non-negative random variable exceeds a certain value to its expectation. It is particularly useful when little is known about the detailed probability distribution beyond its mean. Markov's inequality states that the probability of the random variable exceeding a number a is at most the expectation of the variable divided by a . Formally Markov's inequality is given as:

$$P(X \geq a) \leq \frac{\mathbb{E}[X]}{a}, \quad a > 0 \quad (2.117)$$

Chebyshev's inequality offers a way to bound the probability that a random variable deviates from its expectation by more than a certain distance, using only its variance. It provides a conservative estimate, stating that the probability a variable deviates from its mean by at least a is no more than the variance divided by a^2 . This inequality is essential when exact distributions are unknown or hard to compute, ensuring robustness in probabilistic estimates. Formally Chebyshev's inequality is given as:

$$P(|X - \mathbb{E}[X]| \geq a) \leq \frac{\text{Var}(X)}{a^2} \quad (2.118)$$

2.5.5. Convergence of Random Variables

Convergence of random variables is concerned with how sequences of random variables behave as they approach some limit. Different types of convergence are used depending on context: Convergence in probability means that random variables get arbitrarily close to the limit with high probability:

$$X_n \xrightarrow{p} X \text{ if } \forall \epsilon > 0, P(|X_n - X| > \epsilon) \rightarrow 0 \quad (2.119)$$

While almost sure convergence is stronger, stating that variables approach the limit almost every time in repeated trials:

$$X_n \xrightarrow{a.s.} X \text{ if } P\left(\lim_{n \rightarrow \infty} X_n = X\right) = 1 \quad (2.120)$$

And convergence in distribution means the cumulative distribution functions approach the limit distribution.

$$X_n \xrightarrow{d} X \text{ if } F_{X_n}(x) \rightarrow F_X(x) \quad (2.121)$$

Understanding convergence is crucial, particularly for statistical inference and proving theorems such as the Law of Large Numbers or Central Limit Theorem.

2.5.6. Law of Large Numbers (LLN)

The Law of Large Numbers (LLN) describes the stability of long-term averages of random experiments. It states that as the number of trials increases, the average outcome will get closer to the expected value. There are two forms: the weak LLN states convergence in probability, while the strong LLN ensures almost sure convergence. The LLN is the reason empirical averages observed in experiments or surveys provide reliable estimates of underlying probabilities and expected values in real-world scenarios. They are formally defined as:

Weak LLN:

$$\frac{1}{n} \sum_{i=1}^n X_i \xrightarrow{p} \mu \quad (2.122)$$

Strong LLN:

$$\frac{1}{n} \sum_{i=1}^n X_i \xrightarrow{a.s.} \mu \quad (2.123)$$

2.5.7. Central Limit Theorem (CLT)

The Central Limit Theorem (CLT) explains why the normal (Gaussian) distribution frequently appears in nature. It states that the sum or average of many independent, identically distributed random variables with finite mean and variance tends toward a normal distribution, regardless of the original distribution's shape. The CLT underpins many statistical methods, including hypothesis testing and confidence intervals, allowing us to approximate complex distributions simply by knowing their mean and variance.

$$\frac{\sum_{i=1}^n (X_i - \mu)}{\sigma \sqrt{n}} \xrightarrow{d} \mathcal{N}(0, 1) \quad (2.124)$$

2.5.8. Conditional Probability and Independence

Conditional probability quantifies the probability of an event given that another event has already occurred. It updates our probability estimates based on new information, forming the foundation of Bayesian statistics and inference. Conditional Probability is formally given as:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}, \quad P(B) > 0 \quad (2.125)$$

Independence between two events means that the occurrence of one event does not influence the probability of the other. Formally, two events are independent if the probability of their intersection equals the product of their individual probabilities. Independence simplifies complex probability calculations and is a fundamental assumption in many probabilistic and statistical models. Independence is formally given as

$$P(A \cap B) = P(A)P(B) \quad (2.126)$$

2.5.9. Markov Chains and Markov Processes

The Markov property describes systems or processes in which the future state depends only on the current state, not on any past states. Such systems are called "memoryless." This property significantly simplifies analysis and computation for dynamic systems, enabling easier prediction of future states. For example, a robot's future position might depend solely on its current position and velocity, rather than its entire history of previous states. A Markov chain is a mathematical system evolving randomly from one state to another, possessing the Markov property. It describes systems where transitions depend only on the current state. Markov chains are widely applied in areas such as robotics, queuing theory, and biological modeling, providing a structured way to handle uncertainty over time.

$$P(X_t \in A | X_{t_1} = x_1, \dots, X_{t_n} = x_n) = P(X_t \in A | X_{t_n} = x_n) \quad (2.127)$$

A Markov Process is a generalization of a Markov chain, describing a system where transitions between states can occur continuously over time, and the state space can be continuous or discrete. It also satisfies the Markov property: the future state depends solely on the current state, not the entire past.

A transition probability matrix is a key concept in probability theory, particularly for describing Markov chains. It is a mathematical tool used to represent how a system moves from one state to another over discrete time steps. Formally transition probability matrix is defined as:

$$P_{ij} = P(X_{n+1} = j | X_n = i) \quad (2.128)$$

The Chapman-Kolmogorov equations provide a method for computing multi-step transition probabilities in Markov chains. These equations express multi-step transitions as sums of products of single-step transition probabilities. They are fundamental to analyzing the long-term behavior of Markov processes, enabling calculations of probabilities over extended sequences of events or states.

Chapman-Kolmogorov Equations:

$$P_{ij}^{(n+m)} = \sum_k P_{ik}^{(n)} P_{kj}^{(m)} \quad (2.129)$$

2.6. CONCLUSION

This chapter has outlined the theoretical foundations and computational techniques required to understand the reachability-based safety guarantees developed in the remainder of this dissertation. It began with optimal control and differential games, leading to Hamilton-Jacobi reachability as a robust verification framework. It then introduced neural networks and SMT-based formal methods to support data-driven and certifiable approximations of value

functions. Finally, it reviewed measure and probability theory as tools for reasoning about stochastic behaviour and probabilistic safety.

From this point onward, each chapter builds on these foundations to present a distinct research contribution. These include new definitions of manoeuvrability and safety grounded in HJ theory, extensions of the HJ framework to model uncertainty, certified neural approximations of PDE solutions, and the formulation of stochastic safe and control-invariant sets. Each contribution is self-contained but conceptually linked by the core themes introduced in this chapter.

REFERENCES

- [1] S. L. Herbert, S. Bansal, S. Ghosh, and C. J. Tomlin. “Reachability-based safety guarantees using efficient initializations”. In: *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE. 2019, pp. 4810–4816.
- [2] M. Chen and C. J. Tomlin. “Hamilton–jacobi reachability: Some recent theoretical advances and applications in unmanned airspace management”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 1.1 (2018), pp. 333–358.
- [3] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-jacobi reachability: A brief overview and recent advances”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 2242–2253.
- [4] Y. Zhang, X. Liang, D. Li, S. S. Ge, B. Gao, H. Chen, and T. H. Lee. “Barrier Lyapunov function-based safe reinforcement learning for autonomous vehicles with optimized backstepping”. In: *IEEE Transactions on Neural Networks and Learning Systems* 35.2 (2022), pp. 2066–2080.
- [5] C. Parzani and S. Puechmorel. “On a Hamilton-Jacobi-Bellman approach for coordinated optimal aircraft trajectories planning”. In: *Optimal Control Applications and Methods* 39.2 (2018), pp. 933–948.
- [6] S. Kousik, S. Vaskov, F. Bu, M. Johnson-Roberson, and R. Vasudevan. “Bridging the gap between safety and real-time performance in receding-horizon trajectory design for mobile robots”. In: *The International Journal of Robotics Research* 39.12 (2020), pp. 1419–1469.
- [7] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin. “A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games”. In: *IEEE Transactions on automatic control* 50.7 (2005), pp. 947–957.
- [8] P. Cardaliaguet, M. Quincampoix, and P. Saint-Pierre. “Differential games through viability theory: Old and recent results”. In: *Advances in dynamic game theory*. Springer, 2007, pp. 3–35.
- [9] M. Bardi, I. C. Dolcetta, et al. *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*. Vol. 12. Springer, 1997.
- [10] L. C. Evans and P. E. Souganidis. “Differential games and representation formulas for solutions of Hamilton-Jacobi-Isaacs equations”. In: *Indiana University mathematics journal* 33.5 (1984), pp. 773–797.
- [11] S. Osher and R. P. Fedkiw. *Level set methods and dynamic implicit surfaces*. Vol. 1. Springer New York, 2005.

- [12] I. M. Mitchell and J. A. Templeton. “A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems”. In: *Hybrid Systems: Computation and Control: 8th International Workshop, HSCC 2005, Zurich, Switzerland, March 9-11, 2005. Proceedings* 8. Springer. 2005, pp. 480–494.
- [13] J. Lygeros. “On reachability and minimum cost optimal control”. In: *Automatica* 40.6 (2004), pp. 917–927.
- [14] J. A. Sethian. *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*. Vol. 3. Cambridge university press, 1999.
- [15] S. L. Herbert. *Safe real-world autonomy in uncertain and unstructured environments*. University of California, Berkeley, 2020.
- [16] S. Bansal. *Safe and Data-Efficient Learning for Robotics: A Control Theoretic Approach*. University of California, Berkeley, 2020.
- [17] J.-P. Aubin and H. Frankowska. “Viability kernel of control systems”. In: *Nonlinear synthesis*. Springer, 1991, pp. 12–33.
- [18] I. M. Mitchell and C. J. Tomlin. “Overapproximating reachable sets by Hamilton-Jacobi projections”. In: *journal of Scientific Computing* 19 (2003), pp. 323–346.
- [19] C. C. Aggarwal *et al.* *Neural networks and deep learning*. Vol. 10. 978. Springer, 2018.
- [20] A. Abate, D. Ahmed, A. Edwards, M. Giacobbe, and A. Peruffo. “FOSSIL: a software tool for the formal synthesis of lyapunov functions and barrier certificates using neural networks”. In: *Proceedings of the 24th international conference on hybrid systems: computation and control*. 2021, pp. 1–11.
- [21] S. Bansal and C. J. Tomlin. “Deepreach: A deep learning approach to high-dimensional reachability”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 1817–1824.
- [22] D. Kroening and O. Strichman. *Decision procedures*. Vol. 5. Springer, 2008.
- [23] S. Gao, S. Kong, and E. M. Clarke. “dReal: An SMT solver for nonlinear theories over the reals”. In: *International conference on automated deduction*. Springer. 2013, pp. 208–214.
- [24] L. De Moura and N. Bjørner. “Z3: An efficient SMT solver”. In: *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer. 2008, pp. 337–340.
- [25] V. I. Bogachev and M. A. S. Ruas. *Measure theory*. Vol. 1. 1. Springer, 2007.
- [26] G. Grimmett and D. Stirzaker. *Probability and random processes*. Oxford university press, 2020.

3

A NOVEL ALGORITHMIC APPROACH TO OBTAINING MANEUVERABLE VIABLE SETS

Ensuring safety in autonomous systems is essential as they become more integrated with modern society. One way to accomplish this is to identify and maintain a safe operating space. To this end, much effort has been devoted in the field of reachability analysis to obtaining viable sets which ensure that a system inside of these sets can remain in these sets, and are thus essential for guaranteeing a system's safety. However, viability does not imply that a system can move from any state in the viable set to any other state in the viable set, within a given time horizon. In this paper, we develop an algorithm to obtain a viable set that allows a given system to move from any state in the set to any other state in the set within a given time horizon. We call this the 'maneuver set', $\mathcal{M}$. We substantiate the algorithm's efficacy through mathematical proof, affirming that the maneuver set obtained through the algorithm is indeed viable. Furthermore, we prove that the system is indeed able to move from any state within this set to any other state in the set. To illustrate the use of our algorithm, we provide the numerical example of a Dubins car, utilising Hamilton-Jacobi-Bellman reachability analysis along with the proposed algorithm in order to obtain $\mathcal{M}$.

3.1. INTRODUCTION

The proliferation of automated systems in our daily lives underscores the growing importance of ensuring their safety. Indeed, as shown in [1], the predominant factor contributing to robot failures is loss-of-control. Furthermore, recent research in [2] has flagged a noteworthy correlation between the escalation of automation and robotics, and an increased incidence of accidents. These findings emphasize the need for better safety measures and strategies within these domains. Various engineering disciplines explore the safety concerns associated with such systems. One effective approach for providing safety guarantees is reachability analysis, which revolves around the conceptual framework of reachable sets. Reachability problems focus on identifying the set of states for which an attainable control policy exists, either to guide the system towards a specified state or to avoid particular states. Much effort has been devoted in the field of reachability analysis to obtain viable sets. These sets ensure that, if the system is inside of these sets, there exists an attainable control input such that the system remains within the viable set and are thus essential for guaranteeing a system's safety. However, viability does not imply that the system can move from any state in the viable set to any other state in the viable set, within the given time horizon. Thus, this paper is motivated by the aim of finding a set that is both viable and allows the system to move from any state in the set to any other state in the set within the given time horizon. We call this set 'maneuver set', $\mathcal{M}$.

3.1.1. Related work

When addressing reachability problems, a variety of avenues have contributed useful approaches, often under different terminologies. Among these, the most prevalent term is 'capture basin' [3]. An alternative approach to obtain the capture basin is articulated within the framework of viability theory, as expounded in [4]. Another method for tackling the reachability problem involves formulating a value function incorporating optimal control principles, subsequently yielding a Hamilton-Jacobi-Bellman (HJB) equation. The solution to this equation is obtained through level set methods, wherein the zero level set of the HJB equation provides the solution to the reachability problem [5], [6], [7]. To address the challenges posed by the curse of dimensionality inherent to such grid-based approaches, innovative strategies have been explored, such as those based on Gaussian mixtures [8], convex optimization [9], particle filters [10], Lagrangian methods [11], and Monte Carlo simulations [12].

Apart from reachability analysis, another way for ensuring system safety involves the implicit design of a safety controller capable of modifying nominal control actions to preserve system integrity, as discussed in [13].

3.1.2. Contributions

The central contribution of this paper is the development of an algorithm for the computation of the maneuver set, $\mathcal{M}$. This set is a viable set that allows

the system to move from any state in the set to any other state in the set within the given time horizon. This set is constructed utilizing forward and backward reachable sets which are obtained through standard HJB reachability analysis. We prove that the maneuver set obtained through the application of the developed algorithm is indeed viable. Additionally, we further prove that the system is able to move from any state within $\mathcal{M}$ to any other state in $\mathcal{M}$, thereby affirming the utility of the algorithm. To illustrate our findings, we provide the numerical example of a Dubins car. In this example, HJB reachability analysis is used to obtain the forward and backward reachable sets of the car from which the maneuver set is derived using the developed algorithm.

3.1.3. Organisation

Section 3.2 establishes the foundational context required and describes the problem statement. Section 3.3 introduces the algorithm and presents the accompanying mathematical proofs. Section 3.4 delivers a concise introduction to HJB reachability analysis, while section 3.5 offers an introductory overview of the numerical example. Section 3.6 presents the results and Section 3.7 concludes the paper.

3.2. PROBLEM SETUP

Let the system be described by equation 3.1.

$$\frac{dx}{dt} = \dot{x} = f(x, u, d, \tau), \tau \in [t, T] \subseteq \mathbb{R}, t \geq 0, x \in \mathbb{X} \subseteq \mathbb{R}^n, u \in \mathbb{U} \subseteq \mathbb{R}^m, \\ d \in \mathbb{D} \subseteq \mathbb{R}^n \quad (3.1)$$

where x are the states of the system, $u(\cdot)$ is the input policy into the system, d is the disturbance to the system (input policy is also represented as u , due to space considerations). $\mathbb{X}, \mathbb{U}$ and $\mathbb{D}$ are continuous spaces (infinite sets), $\mathbb{R}$ is the set of real numbers and $\mathbb{F}$ is the set of all measurable functions. $\mathbb{U}_t^T$ is the set of all possible input policies in time horizon $[t, T]$ (whenever the time interval is obvious, $\mathbb{U}_t^T$ is replaced by $\mathbb{U}$ for the sake of readability). The system described by equation 3.1 is assumed to be locally Lipschitz continuous. $\mathbb{U}$ is compact and time invariant (for the given time interval), i.e., $\mathbb{U}_t^\tau = \mathbb{U}_{t+t_2}^{\tau+t_2}$. Furthermore, $\mathbb{U}_t^{T_1} \subseteq \mathbb{U}_t^{T_2}$ for $T_2 \geq T_1$. $\zeta_{x,t}^{u,d}(\tau)$ is the state of the system under control policy $u(\cdot)$, with initial state x and initial time t evaluated at time τ . It is further assumed that $\zeta_{x,t}^{u,d}(\tau) = \zeta_{x,t+t_2}^{u,d}(\tau + t_2)$ (this follows directly from the time-invariance of the system). The disturbance is assumed to be bounded, compact and time invariant (maximum and minimum disturbance is fixed).

In this paper, the disturbance, d , is a map, $d : \mathbb{U} \rightarrow \mathbb{D}$, i.e., player 2 (disturbance) is restricted to a strategy dependent upon player 1 (input), such that it always opposes the goals of player 1. Furthermore, it is assumed that

player 2 can only draw from non-anticipative strategies (equation (3.2)). This is a general assumption in robust control problems [14].

$$d \in \Lambda_t^T := \{\mathcal{N} : \mathbb{U}(t) \rightarrow \mathbb{D}(t) := u(\tau) = \hat{u}(\tau) \text{ a.e. } \tau \in [t, T]\} \quad (3.2)$$

$$\Rightarrow \mathcal{N}[u](\tau) = \mathcal{N}[\hat{u}](\tau) \text{ a.e. } \tau \in [t, T]$$

This means that player 2 cannot produce a different strategy in response to player 1's strategy until player 1's strategy changes. Thus, player 2 cannot set their strategy based on the anticipated future strategy of player 1. However, in this setting player 2 has the advantage of deciding their strategy based on player 1 at every time instance. Thus, the disturbance has an instantaneous informational advantage. The reason for this is that, in robust problems, the worst case scenario is expected. The above assumptions guarantee the existence and uniqueness of the solution of the system equation (3.1).

The forward reachable set is the set of states for which there exists at least one control policy, for all possible disturbances, that will drive the system into the set from its initial set, $\mathcal{L}$, within the given time horizon. It is given by equation 3.3. This set is also called the robust forward reachable set.

$$V_{\text{FRT}}([t, T], \mathcal{L}) := \{y : \exists u \in \mathbb{U}_t^T, \forall d \in \mathbb{D}_t^T, \exists \tau \in [t, T], x \in \mathcal{L}, \zeta_{x,t}^{u,d}(\tau) = y\} \quad (3.3)$$

The forced forward reachable set is the set of states for which there exists at least one possible disturbance, for all control policies, that will drive the system into the set from its initial set, $\mathcal{L}$, within the time horizon. It is given by equation 3.4.

$$V_{\text{FFRT}}([t, T], \mathcal{L}) := \{y : \forall u \in \mathbb{U}_t^T, \exists d \in \mathbb{D}_t^T, \exists \tau \in [t, T], x \in \mathcal{L}, \zeta_{x,t}^{u,d}(\tau) = y\} \quad (3.4)$$

The backward reachable set is the set of all initial states from which there exists at least one control policy, for all possible disturbances, that will lead the system into the required set, $\mathcal{L}$, within the time horizon. It is given by equation 3.5. This set is also known as the robust backward reachable set.

$$V_{\text{BRT}}([t, T], \mathcal{L}) := \{x : \exists u \in \mathbb{U}_t^T, \forall d \in \mathbb{D}_t^T, \exists \tau \in [t, T], \zeta_{x,t}^{u,d}(\tau) \in \mathcal{L}\} \quad (3.5)$$

The forced backward reachable set is the set of all the initial states from which there exists at least one possible disturbance, for all control policies, that will lead the system into the required set, $\mathcal{L}$, within the time horizon. It is given by equation 3.6. This set is also known as the avoid backward reachable set [7].

$$V_{\text{FBRT}}([t, T], \mathcal{L}) := \{x : \forall u \in \mathbb{U}_t^T, \exists d \in \mathbb{D}_t^T, \exists \tau \in [t, T], \zeta_{x,t}^{u,d}(\tau) \in \mathcal{L}\} \quad (3.6)$$

The viable set is defined as the set of states such that, if the system is inside of the viable set, there exists a control policy such that the system stays inside

of the viable set for the given time horizon. Equation 3.7 defines the viable set, $V_{\text{cinv}}([t, T])$. The union of all viable sets for the given time horizon is called the viable kernel and is denoted by $K_{\text{cinv}}([t, T])$, i.e., the viable kernel is the biggest viable set. It is also referred to as robust viable set and robust viable kernel, respectively.

$$V_{\text{cinv}}([t, T]) := \{x : \exists u \in \mathbb{U}_t^T, \forall d \in \mathbb{D}_t^T, \forall \tau \in [t, T], x \in V_{\text{cinv}}([t, \tau]), \zeta_{x,t}^{u,d}(\tau) \in V_{\text{cinv}}([t, \tau])\} \quad (3.7)$$

The intersection of forward and backward reachable set (often called the safe set in literature [15],[16],[17]) is shown in equation 3.8. In this paper we will refer to this set as intersection set.

$$V_{\text{int}}([t, T], \mathcal{L}) = V_{\text{BRT}}([t, T], \mathcal{L}) \cap V_{\text{FRT}}([t, T], \mathcal{L}) \quad (3.8)$$

In this paper we define a set called maneuver set. It is a viable set in which the system can travel from any state in the maneuver set to any other state in the maneuver set within the given time horizon. The set is denoted as $\mathcal{M}([t, T])$.

We also define a trajectory set. Given a control policy (u), a time horizon ($[t, T]$), and an initial state (x), this is the set of all the states that the system transitions to from the initial state in the time horizon under the control policy. We represent this set as $\{\zeta_{x,t}^{u,d}(T)\}$. One must note that since disturbance d is dependent on the control input (in a robust sense) the trajectory set is deterministic in nature.

The definitions of the sets provided above are valid under the assumption that the control policy set ($\mathbb{U}_t^T$) and the disturbance set ($\mathbb{D}_t^T$) are compact in nature, i.e., they are bounded and closed sets. The control set is usually assumed to be compact, which is true for the systems of interest (i.e., physical systems). However, the disturbance set can be bounded or unbounded based on the environment.

3.3. THE MANEUVER SET

The algorithm developed in this paper (algorithm 3.1) aims to find a viable set in which the system can travel from any state to any state within a given time horizon. Algorithm 1, in essence, utilises the forward and backward reachable sets of an initial set, $\mathcal{L}$ and obtains their intersection. Then it computes the complement of this intersection and finds the forced backward reachable set of this complement. It then subtracts this forced backward reachable set from the intersection, resulting in the set with the desired properties, i.e., the maneuver set (set $\mathcal{M}$).

The first part of the algorithm to obtain the maneuver set ($\mathcal{M}$) involves taking the intersection of the forward and backward reachable sets of an initial set, $\mathcal{L}$ and thus obtaining the intersection set, $S(V_{\text{int}}([t, T], \mathcal{L}))$. As proven in lemma

Algorithm 3.1: An algorithm to obtain maneuver set ($\mathcal{M}$)

Require: $V_{\text{BRT}}([t, T], \mathcal{L}), V_{\text{FRT}}([t, T], \mathcal{L})$
 $S \leftarrow V_{\text{BRT}}([t, T], \mathcal{L}) \cap V_{\text{FRT}}([t, T], \mathcal{L})$ Note: $V_{\text{int}}([t, T], \mathcal{L}) = S$
 Obtain $\bar{S} = \mathbb{X} \setminus S$
 $Q \leftarrow V_{\text{FBRT}}([t, T], \bar{S})$
 $\mathcal{M} = S \setminus Q$

3

1, within the set S , the system can transition from any state to any state. That is because any state inside of the intersection set is also part of set $\mathcal{L}$'s backward reachable set. Per definition this means that within the time horizon 0 to T the system can reach set $\mathcal{L}$. Furthermore, as per definition $\mathcal{L}$ is a subset of its own forward reachable set. Thus, the system can start from any state within $\mathcal{L}$ and can reach any state in the forward reachable set within the time horizon 0 to T .

Lemma 1: The system can transition from any state in the $V_{\text{int}}([0, T], \mathcal{L})$ to any other state in $V_{\text{int}}([0, T], \mathcal{L})$ within the time horizon of $[0, 2T]$.

Proof : Let $x, x_1 \in V_{\text{int}}([0, T], \mathcal{L})$

Since $x \in V_{\text{int}}([0, T], \mathcal{L}) \Rightarrow x \in V_{\text{BRT}}([0, T], \mathcal{L})$

$$\Rightarrow \exists u \in \mathbb{U}_0^T, \forall d \in \mathbb{D}_0^T, \exists \tau \in [0, T], \text{ s.t. } \zeta_{x,0}^{u,d}(\tau) = x_{\mathcal{L}} \in \mathcal{L}$$

Since the system is time invariant $V_{\text{FRT}}([0, T], \mathcal{L}) = V_{\text{FRT}}([\tau, T + \tau], \mathcal{L})$

$$\Rightarrow x' \in V_{\text{int}}([0, T], \mathcal{L}) \Rightarrow x' \in V_{\text{FRT}}([\tau, T + \tau], \mathcal{L})$$

$$\Rightarrow \exists u' \in \mathbb{U}_{\tau}^{T+\tau}, \forall d' \in \mathbb{D}_{\tau}^{T+\tau}, \exists \tau' \in [\tau, T + \tau], x_{\mathcal{L}} \in \mathcal{L}, \zeta_{x_{\mathcal{L}},\tau}^{u',d'}(\tau') = x'$$

$$\Rightarrow \exists u'' \in \mathbb{U}_0^{2T}, \forall d'' \in \mathbb{D}_0^{2T}, \exists \tau'' \in [0, 2T] \text{ s.t. } x, x' \in V_{\text{int}}([0, T], \mathcal{L}),$$

$$\zeta_{x,0}^{u'',d''}(\tau'') = x'$$

However, this does not imply that the intersection set is viable, as depicted in figure 3.1. It can be observed that the system at any state within the set S can transition to any state in $\mathcal{L}$ (as $\forall x \in S \Rightarrow x \in V_{\text{BRT}}([t, T], \mathcal{L})$) but the only way to do so might be by leaving set S , as illustrated by the red trajectory in figure 3.1. In contrast, for the blue dot there exists a control, and thus a trajectory, which leads the system into set $\mathcal{L}$ while ensuring that the complete trajectory stays inside the intersection set during the given time frame.

Lemma 2 proves that set $\mathcal{M}$ is viable. After finding the intersection of the forward and backward reachable set of set $\mathcal{L}$, we now want to find all the states in this intersection that force the system to leave the intersection in order to reach set $\mathcal{L}$, similar to the red trajectory in figure 3.1. These states are found

by using the complement of the intersection set (represented by $\bar{S}$) and finding the forced backward reachable set of this complement, $Q (V_{\text{FBRT}}([t, T], \bar{S}))$. Q consists of all the states that force the system into the complement of the intersection set $\bar{S}$ (which by definition is outside of the intersection set). Thus, we are provided with the set of states that we want to eliminate from the intersection set. These states are subsequently removed from the intersection set.

If we now take any state in this transformed intersection set (i.e. $\mathcal{M}$), we can be sure that the system will never leave the original intersection set since all states that would force it to leave have been eliminated from it. We also prove that the system will stay inside set $\mathcal{M}$ at all times. This is because any state inside set $\mathcal{M}$ is also a state in the backward reachable set of set $\mathcal{L}$. It is not possible that there exists a trajectory to $\mathcal{L}$ (e.g., the red trajectory in figure 2) that exits set $\mathcal{M}$ and crosses any state we have removed from the intersection set (illustrated by area Q in figure 3.2).

That is because if such a trajectory existed, it would imply that there is a direct trajectory from a state in Q to $\mathcal{L}$ which does not force the system outside of the intersection set. But if that was the case, this state would not have been part of Q in the first place. Therefore, the set $\mathcal{M}$ is viable.

Lemma 2: The set $\mathcal{M}$ is a viable.

Proof: Part I

Since $S \cup \bar{S} = \mathbb{X}$

$$\Rightarrow Q = V_{\text{FBRT}}([t, T], \bar{S}) \subseteq S$$




Figure 3.1: Illustration to provide a visual understanding of Lemma 1: (—), (—) and (—) represent the trajectory set.




Figure 3.2: Illustration to provide a visual understanding of Lemma 2: (—), (—) and (—) represent the trajectory set .

Part II

Let $x \in \mathcal{M} \Rightarrow x \in S \wedge x \notin Q$ where $\wedge$ is the ‘AND’ operation

if $x \in Q = V_{\text{FBRT}}([t, T], \bar{S}) \Rightarrow \forall u \in \mathbb{U}_t^T, \exists d \in \mathbb{D}_t^T, \exists \tau \in [t, T], \text{ s.t. } \zeta_{x,t}^{u,d}(\tau) \in \bar{S}$

Since $x \notin Q \Rightarrow \exists u \in \mathbb{U}_t^T, \forall d \in \mathbb{D}_t^T, \forall \tau \in [t, T], \text{ s.t. } \zeta_{x,t}^{u,d}(\tau) \in S$

$x \in \mathcal{M} = S \setminus Q \Rightarrow \exists u \in \mathbb{U}_t^T, \forall d \in \mathbb{D}_t^T, \forall \tau \in [t, T], \text{ s.t. } \zeta_{x,t}^{u,d}(\tau) \in S$

Part III

Let $x' \in \mathcal{M} \Rightarrow x' \in S \Rightarrow x' \in V_{\text{BRT}}([t, T], \mathcal{L})$

$\cap V_{\text{FRT}}([t, T], \mathcal{L}) \Rightarrow x' \in V_{\text{BRT}}([t, T], \mathcal{L})$

$\Rightarrow \exists u' \in \mathbb{U}_t^T, \forall d \in \mathbb{D}_t^T, \exists \tau' \in [t, T] \text{ s.t. } \zeta_{x',t}^{u',d}(\tau') \in \mathcal{L}$

It is evident from definition of the set S that $\mathcal{L} \subseteq S$

$$\Rightarrow \exists u' \in \mathbb{U}_t^T, \forall d \in \mathbb{D}_t^T, \exists \tau' \in [t, T] \text{ s.t. } \zeta_{x',t}^{u',d}(\tau') \in S$$

Let $\{\zeta_{x',t}^{u',d}(\tau')\}$ be the trajectory set then using part II $\{\zeta_{x',t}^{u',d}(\tau')\} \subseteq S$

Part IV

Claim: $Q \cap \{\zeta_{x',t}^{u',d}(\tau')\} = \emptyset$, Where $\emptyset$ is empty set

Let the claim be false $\Rightarrow Q \cap \{\zeta_{x',t}^{u',d}(\tau')\} \neq \emptyset \Rightarrow \exists \tau'' \in [t, \tau']$ s.t. $\zeta_{x',t}^{u',d}(\tau'')$
 $= x'' \in Q \cap \{\zeta_{x',t}^{u',d}(\tau')\}$

Since $x'' \in \{\zeta_{x',t}^{u',d}(\tau')\} \Rightarrow \exists u' \in \mathbb{U}_t^T, \forall d \in \mathbb{D}_t^T$,
 $\forall \tau''' \in [\tau'', \tau']$ s.t. $\zeta_{x'',\tau''}^{u',d}(\tau''') \in S$

Since $x'' \in Q \Rightarrow \forall u \in \mathbb{U}_t^T, \exists d \in \mathbb{D}_t^T, \exists \tau \in [t, T]$, s.t. $\zeta_{x,t}^{u,d}(\tau) \in \bar{S}$

This is a contradiction; thus the claim is true and by using part II, III and IV

$$x' \in \mathcal{M} \Rightarrow \exists u' \in \mathbb{U}_t^T, \forall d \in \mathbb{D}_t^T, \forall \tau' \in [t, T] \text{ s.t. } \zeta_{x',t}^{u',d}(\tau') \in S \setminus Q$$

Lemma 3 proves that within set $\mathcal{M}$ the system can travel from any state to any state. Here $\mathcal{M}$ is a subset of the intersection set, for which lemma 1 has proven that the system can travel from any state to any state. Using the same argument we can show that the system can travel from any state to any state in set $\mathcal{M}$.

As illustrated by figure 3, from any state in set $\mathcal{M}$ the system can take a trajectory to $\mathcal{L}$ (yellow trajectory), and from $\mathcal{L}$ the system can go to any state within $\mathcal{M}$ (blue trajectory).

Lemma 3: The system can transition from any state in $\mathcal{M}$ to any other state in $\mathcal{M}$, in time horizon $[0, 2T]$.

Proof: Let $x \in \mathcal{M}$

$$\text{Let } x \in \mathcal{M} \Rightarrow x \in S \Rightarrow x \in V_{\text{BRT}}([0, T], \mathcal{L}) \cap V_{\text{FRT}}([0, T], \mathcal{L})$$

$$\Rightarrow x \in V_{\text{BRT}}([0, T], \mathcal{L}) \Rightarrow \exists u \in \mathbb{U}, \forall d \in \mathbb{D}, \exists \tau \in [0, T] \text{ s.t. } \zeta_{x,0}^{u,d}(\tau) = x' \in \mathcal{L}$$

Since $\mathcal{L} \subseteq S \Rightarrow x' \in S \Rightarrow x' \in V_{\text{FRT}}([0, T], \mathcal{L})$

$$\Rightarrow \exists u' \in \mathbb{U}, \forall d \in \mathbb{D}, \exists \tau' \in [0, T] \text{ s.t. } \zeta_{x',0}^{u',d}(\tau') = x''$$

Since system is assumed to be time invariant

$$\Rightarrow \exists u' \in \mathbb{U}, \forall d \in \mathbb{D}, \exists \tau' \in [T, 2T] \text{ s.t. } \zeta_{x',T}^{u',d}(\tau') = x''$$

$$\Rightarrow \exists u'' \in \mathbb{U}, \forall d \in \mathbb{D}, \exists \tau'' \in [0, 2T], x, x'' \in \mathcal{M} \text{ s.t. } \zeta_{x,0}^{u'',d}(\tau'') = x''$$

3.4. HAMILTON JACOBI BELLMAN REACHABILITY ANALYSIS

In this subsection a short overview of Hamilton-Jacobi-Bellman (HJB) reachability analysis is provided. This technique is used to obtain the forward/backward reachable set pertaining to both reach and avoid scenarios [18–21]. A cost function is designed as shown in equation 3.9.

$$J(x, t, u) = \int_t^T c(x(\tau), u(\tau)) d\tau + l(x(T)) \quad (3.9)$$




Figure 3.3: Illustration to provide a visual understanding of Lemma 3: (—) and (—) represents the trajectory set and $S \setminus Q = \mathcal{M}$.

Then an optimal control problem is formulated such that the cost function is minimised with the system dynamics as one of the constraints as shown in equation (3.10).

$$V(x(t), t) = \inf_u [J(x, t, u)] \quad (3.10)$$

$$\text{subject to } \dot{x} = f(x(\tau), u(\tau)), \forall \tau \in [t, T]$$

$$u(\tau) \in \mathbb{U}_t^T$$

This optimal control problem is solved using dynamic programming, giving rise to a partial differential equation (PDE) called HJB PDE for the continuous case, or a recursive update equation called Bellman backup/equation in the case of discrete systems. The zero level set of the solution of this PDE provides us with the set of interest. This method further produces a control law or a controller that can keep the system in the set of interest. We obtain different value functions for different sets of interest, which correspond to different HJB PDE [22]. For backward reachable sets the value function is given by equation 3.11 and the corresponding PDE is given by equation 3.12

$$V_{\text{BRT}}([t, T], \mathcal{L}) = \inf_u \sup_d \min_{\tau \in t, T} l(\zeta_{x,t}^{u,d}(\tau)) \quad (3.11)$$

$$\min\{l(x, t) - V(x, t), D_t V + \inf_u \sup_d D_x V \cdot f(x, u)\} = 0 \quad (3.12)$$

The optimal control at any given time and state is given by equation 3.13

$$u^*(x, t) = \operatorname{arginf}_u \sup_d [D_x V \cdot f(x(t), u(t))] \quad (3.13)$$

For forced backward reachable sets the value function is given by equation 3.14 and the corresponding PDE is given by equation 3.15

$$V_{\text{FBRT}}([t, T], \mathcal{L}) = \sup_u \inf_d \min_{\tau \in t, T} l(\zeta_{x,t}^{u,d}(\tau)) \quad (3.14)$$

$$\min\{l(x, t) - V(x, t), D_t V + \sup_u \inf_d D_x V \cdot f(x, u)\} = 0 \quad (3.15)$$

The optimal control at any given time and state is given by equation 3.16

$$u^*(x, t) = \operatorname{argsup}_u \inf_d [D_x V \cdot f(x(t), u(t))] \quad (3.16)$$

For forward reachable sets the value function is given by equation 3.17 and the corresponding PDE by equation 3.18.

$$V_{\text{FRT}}([t, T], \mathcal{L}) = \inf_u \sup_d \max_{\tau \in t, T} l(\zeta_{x,t}^{u,d}(\tau)) \quad (3.17)$$

$$\max \{l(x, t) - V(x, t), D_t V - \inf_u \sup_d D_x V \cdot f(x, u)\} = 0 \quad (3.18)$$

The optimal control at any given time and state is given by equation 3.19

$$u^*(x, t) = \arg \inf_u \sup_d [-D_x V \cdot f(x(t), u(t))] \quad (3.19)$$

The obtained HJB PDE is solved using numerical methods. One of the most common methods is the level set method [23], [24]. There are also toolboxes available that are specifically created to solve HJB PDE for reachability analysis.

3

3.5. NUMERICAL EXAMPLE: DUBINS CAR

To demonstrate the theoretical results of the preceding sections, we employ the Dubins car with a constant speed v as the system model, shown in equation 3.20.

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} v \cos(\theta) \\ v \sin(\theta) \\ \omega \end{bmatrix} + \begin{bmatrix} d_x \\ d_y \\ d_\theta \end{bmatrix}, \omega \in [-1, 1] \quad (3.20)$$

where x is the position along the x-axis, y is the position along the y-axis, θ is the turn angle, ω is the input to the system and v is the constant velocity of the car which is assumed to be 1 m/s. It is assumed that ω is bounded between $[-1, 1]$. Furthermore, the disturbance is also assumed to be bounded with $d_x \in [-0.1, 0.1]$, $d_y \in [-0.1, 0.1]$ and $d_\theta \in [-0.1, 0.1]$. We follow the steps of the algorithm as explained in section III.

3.6. RESULTS AND DISCUSSION

To obtain the forward and backward reachable sets, a cylinder with a radius of 1 centred at origin was chosen as the initial set $\mathcal{L}$. Physically this initialisation means that the Dubins car could be situated at any state within a circle of radius 1 with an initial orientation of any angle between 0 and 360. Using the level set method and a customised version of the toolbox [25], equation 3.18 is solved to compute the forward reachable set of the initial set $\mathcal{L}$. Figure 3.4 shows the outer contour of the forward reachable set ($V_{\text{BRT}}([t, T], \mathcal{L})$) of the system represented by equation 3.20 for a time horizon of 5 seconds in green, while the outer contour of initial set ($\mathcal{L}$) is depicted in blue. The x-axis and y-axis of the figure represent distances in metres. The z-axis represents the turn angle or yaw angle of the system in radians.

Using the same toolbox, the backward reachable set of the initial set $\mathcal{L}$ is found through equation 3.12. Figure 3.5 shows the outer contour of the backward reachable set ($V_{\text{BRT}}([t, T], \mathcal{L})$) of the system represented by equation 3.20 for a time horizon of 5 seconds in green colour, while the outer contour of initial set ($\mathcal{L}$) is depicted using blue colour. The x-axis and y-axis of the figure represent




Figure 3.4: Outer contour of the forward reachable set denoted by colour (—) and initial set denoted by colour (—)




Figure 3.5: Outer contour of the backward reachable set denoted by colour (—) and initial set denoted by colour (—)

the distances travelled in metres. The z-axis represents the turn angle or yaw angle of the system in radians.

Using the forward and backward reachable sets as inputs into the algorithm developed in this paper and once again using the modified toolbox, we arrive at the maneuver set $\mathcal{M}$. Figure 3.6 shows the outer contour of the maneuver set ($\mathcal{M}$) of the system represented by equation 3.20 for a time horizon of 5 seconds in green colour. The x-axis and y-axis of the figure represent the distances travelled in metres. The z-axis represents the turn angle or yaw angle of the system in radians.

To illustrate that the system is able to stay within the maneuver set, we chose a random state within the maneuver set and show that the system is able to stay inside it for 5 seconds. This serves as an illustration for the set's viability. Figure 3.7 shows the outer contour of the maneuver set ($\mathcal{M}$) of the system represented by equation 3.20 for a time horizon of 5 seconds in green colour, an initial condition of the system in blue colour and the trajectory followed by the system




Figure 3.6: Outer contour of the maneuver set denoted by colour (—)

in red colour. The x-axis and y-axis of the figure represent the distances travelled in metres. The z-axis represents the turn angle or yaw angle of the system in radians.

3.7. CONCLUSIONS

To ensure the safety of automated systems it is essential that they are aware of the region in the state space where they have the control authority to move safely, and the capability to stay within that region. We have derived a definition of such a set, called the maneuver set, which can be computed from forward and backward reachable sets. We illustrated the computation of the set through an example. In future we would like to expand the concept of the maneuver set into the probabilistic domain to obtain probabilistic viable sets that guarantee the system to move from one state to another state with a certain probability.




Figure 3.7: Outer contour of the maneuver set denoted by colour (—), initial state denoted by colour (—) and the trajectory taken by the system denoted by colour (—)

REFERENCES

- [1] J. Carlson, R. R. Murphy, and A. Nelson. “Follow-up analysis of mobile robot failures”. In: *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA’04. 2004*. Vol. 5. IEEE. 2004, pp. 4987–4994.
- [2] S. Yang, Y. Zhong, D. Feng, R. Y. M. Li, X.-F. Shao, and W. Liu. “Robot application and occupational injuries: are robots necessarily safer?” In: *Safety science* 147 (2022), p. 105623.
- [3] J.-P. Aubin and H. Frankowska. “Viability kernel of control systems”. In: *Nonlinear synthesis*. Springer, 1991, pp. 12–33.
- [4] P. Cardaliaguet. “A differential game with two players and one target”. In: *SIAM Journal on Control and Optimization* 34.4 (1996), pp. 1441–1460.
- [5] J. Lygeros. “On reachability and minimum cost optimal control”. In: *Automatica* 40.6 (2004), pp. 917–927.
- [6] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-jacobi reachability: A brief overview and recent advances”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 2242–2253.
- [7] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin. “FaSTrack: A modular framework for fast and guaranteed safe motion planning”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 1517–1522.
- [8] N. Kariotoglou, K. Margellos, and J. Lygeros. “On the computational complexity and generalization properties of multi-stage and stage-wise coupled scenario programs”. In: *Systems & Control Letters* 94 (2016), pp. 63–69.
- [9] K. Lesser, M. Oishi, and R. S. Erwin. “Stochastic reachability for control of spacecraft relative motion”. In: *52nd IEEE Conference on Decision and Control*. IEEE. 2013, pp. 4705–4712.
- [10] G. Manganini, M. Pirodda, M. Restelli, L. Pirodda, and M. Prandini. “Policy search for the optimal control of Markov Decision Processes: A novel particle-based iterative scheme”. In: *IEEE transactions on cybernetics* 46.11 (2015), pp. 2643–2655.
- [11] J. N. Maidens, S. Kaynama, I. M. Mitchell, M. M. Oishi, and G. A. Dumont. “Lagrangian methods for approximating the viability kernel in high-dimensional systems”. In: *Automatica* 49.7 (2013), pp. 2017–2029.

- [12] S. Sun and C. C. de Visser. “Quadrotor safe flight envelope prediction in the high-speed regime: A Monte-Carlo approach”. In: *AIAA Scitech 2019 Forum*. 2019, p. 0948.
- [13] K. P. Wabersich and M. N. Zeilinger. “Linear model predictive safety certification for learning-based control”. In: *2018 IEEE Conference on Decision and Control (CDC)*. IEEE. 2018, pp. 7130–7135.
- [14] S. L. Herbert. *Safe real-world autonomy in uncertain and unstructured environments*. University of California, Berkeley, 2020.
- [15] E. Van Oort, Q. Chu, and J. Mulder. “Maneuver envelope determination through reachability analysis”. In: *Advances in Aerospace Guidance, Navigation and Control: Selected Papers of the 1st CEAS Specialist Conference on Guidance, Navigation and Control*. Springer. 2011, pp. 91–102.
- [16] T. Lombaerts, S. Schuet, K. Wheeler, D. M. Acosta, and J. Kaneshige. “Safe maneuvering envelope estimation based on a physical approach”. In: *Aiaa guidance, navigation, and control (gnc) conference*. 2013, p. 4618.
- [17] S. Schuet, T. Lombaerts, D. Acosta, K. Wheeler, and J. Kaneshige. “An adaptive nonlinear aircraft maneuvering envelope estimation approach for online applications”. In: *Aiaa guidance, navigation, and control conference*. 2014, p. 0268.
- [18] M. Chen, S. Herbert, and C. J. Tomlin. “Exact and efficient Hamilton-Jacobi guaranteed safety analysis via system decomposition”. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2017, pp. 87–92.
- [19] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin. “Decomposition of reachable sets and tubes for a class of nonlinear systems”. In: *IEEE Transactions on Automatic Control* 63.11 (2018), pp. 3675–3688.
- [20] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin. “Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking”. In: *IEEE Transactions on Automatic Control* 66.12 (2021), pp. 5861–5876.
- [21] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert. “Robust control barrier-value functions for safety-critical control”. In: *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE. 2021, pp. 6814–6821.
- [22] S. L. Herbert. *Safe real-world autonomy in uncertain and unstructured environments*. University of California, Berkeley, 2020.
- [23] J. A. Sethian. *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*. Vol. 3. Cambridge university press, 1999.
- [24] S. Osher and R. P. Fedkiw. *Level set methods and dynamic implicit surfaces*. Vol. 1. Springer New York, 2005.

- [25] I. M. Mitchell and J. A. Templeton. “A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems”. In: *Hybrid Systems: Computation and Control: 8th International Workshop, HSCC 2005, Zurich, Switzerland, March 9-11, 2005. Proceedings* 8. Springer. 2005, pp. 480–494.

4

A NOVEL APPROACH TO THE SAFE SET AND ITS ESTIMATION FOR SYSTEMS WITH BOUNDED DISTURBANCES

This paper introduces a new definition of the safe set for nonlinear control systems over finite horizons, motivated by the need for certifiable safe operating envelopes. The prevalent approach of defining the safe set as the intersection of forward and backward reachable tubes, does not, in general, guarantee control invariance or internal reachability. We instead define safety via round-trip reachability: a state is safe if it can be reached from, and returned to, a designated initial set within a finite horizon. We prove that the resulting set is control invariant over all sub-horizons. Moreover, if the initial set is internally controllable with a bounded transition time, then any two states in the safe set are mutually reachable without leaving the set. We further characterize the proposed safe set as a union of time-sliced intersections of forward and backward reachable tubes, enabling computation with reachability tools. In parallel, we formalize the limitation of the intersection-based definition, showing it is control invariant only under a restrictive conditions. All results extend to bounded disturbances under non-anticipative adversarial strategies. Together, these results provide a unified, theoretically grounded framework for finite-horizon safety verification, with implications for motion planning, control synthesis and envelope protection.

4.1. INTRODUCTION

4.1.1. Motivation

The increasing integration of autonomous systems into safety-critical applications, such as aircraft flight control, autonomous driving, and robotic manipulation, has intensified the need for rigorous safety verification methods. Empirical studies show that a large number of robotic failures arise due to loss of control [1], and recent research highlights a correlation between the rise in automation and increased accident rates [2]. These findings underscore the need for formal methods that guarantee system safety not only locally but also over long-term operation in uncertain environments.

In the context of flight safety the discussion is often through the notion of a safe operational/flight envelope. It is usually considered as the collection of aircraft states and operating conditions under which the vehicle remains controllable and is able to avoid loss-of-control scenarios despite actuator limits, aerodynamic nonlinearities, and external perturbations. A key point is that a practically meaningful flight envelope is not merely a static set of constraints, nor only a local invariance region around a single equilibrium. Rather, it must support maneuverability with recovery. The aircraft should be able to deviate from a trim or nominal operating set and still be able to return to that trim set within a relevant time horizon, while remaining within acceptable operating limits throughout the maneuver. This perspective is central to flight safety.

Reachability theory offers a principled way to reason about such questions over finite horizons under dynamics and control constraints. Reachability analysis provides tools to compute sets of states that a system can reach or avoid over a finite time horizon under given dynamics and control constraints. Classical constructs in this domain include the forward reachable tube (FRT)[3], which captures all states that can be reached from a given initial set, and the backward reachable tube (BRT) [4], which represents states that can reach a target set. Another widely studied concept is the control invariant set [5–7], which is the subset of the state space that can be rendered positively invariant by a suitable choice of control inputs.

However, in many practical applications, safety requires more than reachability or control invariance in isolation. For instance during level flight in aircraft, the goal is not just to remain within a control invariant operating region, but also to ensure that the system can recover back to a trim condition after desired deviation, and that it can transition freely between different states within a safe operational envelope. The control invariant set around a single equilibrium point does not guarantee such round-trip behaviour. The BRT alone does not indicate where one can go safely after reaching a target, and the FRT does not ensure whether return is possible. These are the issues that this paper aims to tackle.

4.1.2. Related work

The problem of computing safe operating regions in nonlinear systems has been studied extensively across control theory, robotics, and formal verification, under

a variety of names and formulations. One foundational perspective is offered by viability theory, where the goal is to compute the so-called viability kernel or capture basin: the set of states from which the system can remain within a constraint set under admissible control inputs [8]. Classical approaches to computing capture basins using differential inclusions have been developed in [9, 10], but such methods remain difficult to scale beyond low-dimensional systems.

Another powerful approach is based on Hamilton-Jacobi (HJ) reachability analysis, in which a value function encoding the worst-case cost-to-go is used to estimate the reachable sets and is governed by a HJ partial differential equation (PDE). The zero level set of this value function defines the reachable set under either optimal control or differential game formulations [4, 11, 12]. Grid-based HJ methods offer strong theoretical guarantees and are supported by open-source toolboxes [13, 14], but suffer from the curse of dimensionality. To address this, various approximation techniques have been proposed, including Gaussian mixture models [15], convex optimisation [16], particle filtering [17], and Monte Carlo-based reachability [18]. While these methods improve scalability, they typically sacrifice exactness or robustness guarantees.

From a system safety perspective, reachability-based methods can be grouped into four main paradigms. The first involves computing avoid backward reachable tubes from unsafe states and ensuring the system does not enter them [14, 19, 20]. The second uses forward reachable tubes to forecast system evolution and check for potential future violations [21–24]. The third approach maintains the system within a control invariant set [5, 25], while the fourth defines a “safe set” as the intersection of forward and backward reachable tubes from a given set [26–30].

This intersection-based safe set definition is particularly common in applications such as flight envelope protection and autonomous maneuvering. It is appealing in practice because it captures the idea that states in the set can be both reached from, and returned to, a designated operating set within a finite time horizon. However, it implicitly assumes that every point in the set is reachable from, and can return to, the initial set within a single finite time horizon. As demonstrated by [30] and reinforced in this paper, the control invariance of such sets is not ensured unless strict assumptions are met. While the forward/backward reachable tube-based approaches provide system safety for a certain time horizon, they do not guarantee that the system will not enter a state that will lead to an unavoidable unsafe scenario beyond the time horizon. In particular, the widely used definition fails to account for whether a full round-trip trajectory remains within the intersection set itself, or whether point-to-point transitions between internal states are feasible. These properties are crucial in practice for safety critical control synthesis, motion planning, and certification. Furthermore, this definition lacks a constructive relationship to control invariant sets or control invariance properties under disturbance, limiting its theoretical robustness.

4.1.3. Contributions

In this work, we adopt a control invariance theoretic perspective: a set is considered safe if it is control invariant, meaning that the system can remain within the set over a given time horizon using admissible control inputs. This control invariance based view leads to a new definition of safe sets that subsumes the forward-backward intersection approach, while ensuring stronger guarantees of control invariance and point-to-point maneuverability.

This paper introduces a new definition of the safe set for nonlinear control systems that guarantees both control invariance and internal reachability under significantly weaker assumptions than those required by intersection based formulation. The proposed safe set consists of all states that can be reached from, and returned to, a designated initial set within a finite time horizon under admissible control inputs. We prove that this round-trip condition is sufficient to ensure two essential properties: control invariance (the system can remain within the set for the entire horizon under some admissible control strategy); and internal reachability (any two states in the set can be connected by a trajectory that remains entirely within the set).

We rigorously analyse the standard definition of the safe set, which is based on the intersection of the forward and backward reachable tubes from an initial set. While this intersection is widely used in safety critical applications such as flight envelope protection and motion planning, we show that it is neither control invariant nor does it guarantee internal reachability, given time horizon. We prove that this intersection based set is only control invariant under a strict condition. This limitation is often overlooked in the literature and undermines key safety guarantees typically assumed.

A central contribution of this work is a constructive characterisation of the proposed safe set using classical reachability tools. Specifically, we show that the safe set can be expressed as the union of intersections of forward and backward reachable tubes over varying intermediate time splits. This not only generalises the standard definition but also offers a practical computational strategy for safe set construction using existing reachability algorithms.

Finally, we extend all theoretical results to systems subject to bounded disturbances by modelling disturbances as non-anticipative adversarial inputs. We prove that the proposed safe set remains control invariant and internally reachable in this robust setting, thus enabling safety guarantees even under worst-case disturbances. Together, these contributions offer a principled and computationally tractable framework for safe set computation, bridging control invariance, reachability analysis, and robust control.

4.1.4. Structure of the paper

The remainder of this paper is structured as follows. In [section 4.2](#) we define key concepts such as the forward and backward reachable tubes and the control invariant set and critically analyse the widely used definition of the safe set as the intersection of the forward and backward reachable tubes [28, 30]. In

section 4.3, we introduce the problem formulation and formalise the assumptions used throughout the paper. We then prove several theoretical properties of the proposed safe set, including its control invariance over any sub-horizon and its internal reachability: the ability to transition between any two points within the set without leaving it. In the same section we also show that the backward reachable tube is control invariant under a certain condition on the target set. Furthermore, we establish a constructive relationship between the proposed safe set and the reachability tubes, showing it can be expressed as a union over time-sliced intersections of forward and backward reachable tubes. In section 4.3.2, we extend all theoretical results to systems subject to bounded disturbances, using non-anticipative disturbance strategies. In section 4.4, we present several systems as a case study to demonstrate the validity of our results in practice. Finally, section 4.5 provides numerical results supporting the theoretical claims, and section 4.6 concludes the paper with a discussion of implications and future work.

4.2. CRITICAL ANALYSIS OF THE STANDARD SAFE SET DEFINITION

In order to critique the standard definition of the safe set, we begin by formally recalling the standard reachability constructs commonly used in safety-critical control literature. These definitions serve as the building blocks of our critical analysis and our proposed alternative.

4.2.1. Set Definitions and Safe Set Formulation

We begin by introducing standard definitions of reachable tubes and control invariant sets.

The forward reachable tube is defined in equation (4.1). It contains all states that can be reached from an initial set $\mathcal{L}$ within a time horizon $[t, T]$, under at least one admissible control policy:

$$V_{\text{FRT}}([t, T], \mathcal{L}) := \left\{ x : \exists u \in \mathbb{U}_t^T, \exists \tau \in [t, T], \exists x_o \in \mathcal{L}, \zeta_{x_o, t}^u(\tau) = x \right\} \quad (4.1)$$

The backward reachable tube, defined in equation (4.2), is the set of all states from which there exists a control policy that leads the system into the set $\mathcal{L}$ within the time horizon $[t, T]$:

$$V_{\text{BRT}}([t, T], \mathcal{L}) := \left\{ x : \exists u \in \mathbb{U}_t^T, \exists \tau \in [t, T], \zeta_{x, t}^u(\tau) \in \mathcal{L} \right\} \quad (4.2)$$

The (finite-horizon) control invariant set with respect to the constraint set $\mathcal{K}$ is the set of all states from which there exists a control policy that keeps the trajectory inside $\mathcal{K}$ for the entire time horizon $[t, T]$:

$$V_{\text{CI}}([t, T], \mathcal{K}) := \left\{ x \in \mathcal{K} : \exists u \in \mathbb{U}_t^T, \forall \tau \in [t, T], \zeta_{x, t}^u(\tau) \in \mathcal{K} \right\} \quad (4.3)$$

A common definition of the safe set, widely adopted in the literature on flight envelope protection and maneuvering safety (e.g., [26–28, 31]), is the intersection of the forward and backward reachable tubes from a designated set $\mathcal{L}$, typically representing a collection of trimmable or stable equilibrium states. Formally, this standard safe set is given by

$$V_{\text{int}}([t, T], \mathcal{L}) := V_{\text{FRT}}([t, T], \mathcal{L}) \cap V_{\text{BRT}}([t, T], \mathcal{L}). \quad (4.4)$$

The underlying assumption is that states in V_{int} are both forward and backward reachable from $\mathcal{L}$, and hence lie within a region of safe maneuverability [27]. In many applications, this set is treated as control invariant and further internal reachability is assumed between arbitrary states in V_{int} under admissible control policies.

However, these assumptions are not generally valid. In what follows, we provide a formal critique of this definition, showing that the intersection set V_{int} fails to guarantee either control invariance or internal reachability in general.

Non-control invariance of the intersection set

The fundamental flaw in (4.4) lies in the inability to guarantee that a trajectory remains in the set for the entire time horizon $[t, T]$. Specifically, while $x \in V_{\text{int}}([t, T], \mathcal{L})$ implies that the state x is reachable from $\mathcal{L}$ and can reach back to $\mathcal{L}$ within the horizon, it does not imply that the system can be kept within V_{int} during this round-trip evolution. This is because V_{FRT} and V_{BRT} are not control invariant sets in general, and their intersection need not be either.

Let us consider a trajectory $\zeta_{x_0, t}^u(\cdot)$ such that $x = \zeta_{x_0, t}^u(T') \in V_{\text{int}}$ and $\zeta_{x, T'}^u(\tau) \in \mathcal{L}$ for some $T' < \tau \leq T$. Even if x lies in the intersection, the trajectory may pass through states outside of V_{int} before returning to $\mathcal{L}$. This violates the requirement for control invariance and invalidates the assumption of guaranteed safety throughout the maneuver.

Control invariance does not survive intersection, even under strong assumptions

One might argue that if both the forward reachable tube and the backward reachable tube are themselves control invariant sets, then their intersection should also be control invariant. However, this assumption is flawed. Control invariance is not necessarily preserved under intersection. Formally, let

$$A := V_{\text{FRT}}([t, T], \mathcal{L}) \quad \text{and} \quad B := V_{\text{BRT}}([t, T], \mathcal{L}),$$

and suppose both A and B are control invariant sets over $[t, T]$. That is, for each $x \in A$ (resp. $x \in B$), there exists a control policy $u(\cdot)$ that ensures $\zeta_{x, t}^u(\tau) \in A$ (resp. in B) for all $\tau \in [t, T]$. Still, the intersection $A \cap B$ is not guaranteed to be control invariant.

The reason is that the control policy that keeps a trajectory in A may not keep it in B , and vice versa. Control invariance is a property that depends

jointly on both the set and the control strategy that maintains it. There is no reason to expect that the intersection inherits a common admissible control that respects the constraints of both A and B simultaneously. Hence, even under strong assumptions such as control invariance of the reachable tubes, their intersection cannot be assumed to be control invariant. This underscores the theoretical fragility of defining safety based solely on the intersection of forward and backward reachable tubes.

Intersection does not imply internal reachability

It is often implicitly assumed that the intersection of the forward and backward reachable tubes from a set $\mathcal{L}$ supports internal reachability. That is, for any two states $x_1, x_2 \in V_{\text{int}}([t, T], \mathcal{L})$, one might expect the existence of a control policy $u(\cdot)$ such that the system can be steered from x_1 to x_2 within the time horizon $[t, T]$, while staying in the set. However, this assumption is not valid in general. The key reason lies in the structure of the forward and backward reachable tubes. Specifically:

- $x_1 \in V_{\text{BRT}}([t, T], \mathcal{L})$ implies that the system can be driven from x_1 into the set $\mathcal{L}$ within time $[t, T]$ under some admissible control.
- $x_2 \in V_{\text{FRT}}([t, T], \mathcal{L})$ implies that the system can be driven to x_2 from some point in $\mathcal{L}$ within time $[t, T]$.

While this supports a two-hop argument, namely, $x_1 \rightarrow \mathcal{L} \rightarrow x_2$, it does not imply the existence of a direct trajectory from x_1 to x_2 . More importantly, there is no guarantee that the arrival state of the trajectory from x_1 to $\mathcal{L}$ is the same as the departure state needed to reach x_2 via the forward reachable tube.

To guarantee internal reachability, one would require that the forward reachable tube from $\mathcal{L}$ is reachable from every point in $\mathcal{L}$, that is,

$$x_2 \in \bigcap_{x' \in \mathcal{L}} V_{\text{FRT}}^{x'}([t, T]) \quad (4.5)$$

where $V_{\text{FRT}}^{x'}$ denotes the forward reachable set from the singleton $\{x'\}$. However, in general, $V_{\text{FRT}}([t, T], \mathcal{L})$ only guarantees that there exists *some* $x' \in \mathcal{L}$ such that x_2 is reachable from x' . Hence, without additional structure (e.g., full controllability or strong connectivity of $\mathcal{L}$ under system dynamics), we cannot assert that x_1 and x_2 are mutually reachable via a concatenated trajectory through $\mathcal{L}$.

Conclusion of the Critical Analysis

We conclude that the standard definition of the safe set as the intersection of forward and backward reachable tubes:

- **does not ensure control invariance.**
- **does not ensure internal reachability.**

These shortcomings highlight the need for a more robust definition of the safe set, one that guarantees both internal reachability and control invariance under admissible controls. In what follows, we formalise this intuition by presenting a new problem formulation grounded in control invariance theory. We introduce the mathematical framework and system assumptions that underpin our proposed definition, and demonstrate how it captures and extends the desirable properties of V_{int} , while avoiding its limitations.

4.3. PROBLEM FORMULATION AND PROPOSED APPROACH

This section presents the mathematical setting, formal assumptions, and problem scope considered in this work. We focus on systems governed by nonlinear control dynamics, operating either in the absence of external disturbances or under bounded disturbances. In both cases, our goal is to rigorously define and characterise safe sets, regions of the state space from which the system can be guaranteed to remain safe under admissible control inputs.

4.3.1. Absence of Disturbances

We begin by considering a deterministic nonlinear system of the form:

$$\dot{x} = f(x, u, t), \quad (4.6)$$

where $x \in \mathbb{X} \subseteq \mathbb{R}^m$ denotes the state vector, and $u(\cdot) \in \mathbb{U} \subseteq \mathbb{F}$ represents the control input. Here, $\mathbb{X}$ and $\mathbb{U}$ are continuous spaces, $\mathbb{R}$ is the set of real numbers, and $\mathbb{F}$ denotes the space of measurable functions. We define $\mathbb{U}_t^T$ as the set of all admissible control input functions defined on the time interval $[t, T]$. When the time interval is clear from context, we will drop the subscript and write $\mathbb{U}$ for readability.

Given a control policy $u(\cdot)$ and an initial state x at time t , we denote by $\zeta_{x,t}^u(\tau)$ the unique solution (trajectory) of (4.6) evaluated at time τ . The trajectory segment between t and τ is denoted by:

$$\{\zeta_{x,t}^u(\sigma) : \sigma \in [t, \tau]\} \quad (4.7)$$

Throughout the paper, we often abbreviate the control policy $u(\cdot)$ as u for notational simplicity.

We make the following standard assumptions on the system and input space:

Assumption 1 (Local Lipschitz Continuity) *The dynamics $f(x, u, t)$ are locally Lipschitz continuous in x and u for fixed t . This ensures the existence and uniqueness of solutions to (4.6) on finite time intervals.*

Assumption 2 (Compactness and Time Invariance of Input Space) *The admissible input space $\mathbb{U}_t^\tau$ is compact and time-invariant; that is,*

$$\mathbb{U}_t^\tau = \mathbb{U}_{t+t_2}^{\tau+t_2} \quad \text{for all } t_2 \in \mathbb{R}.$$

Assumption 3 (Temporal Nesting of Control Policies) For any $T_2 \geq T_1$, the admissible input sets satisfy

$$\mathbb{U}_t^{T_1} \subseteq \mathbb{U}_t^{T_2}.$$

Assumption 4 (Time-Invariance of Trajectories) The system is time-invariant in the sense that:

$$\zeta_{x,t}^u(\tau) = \zeta_{x,t+t_2}^u(\tau + t_2), \quad \text{for all } t_2 \in \mathbb{R}.$$

These assumptions are standard in the analysis of controlled dynamical systems. They hold for a wide class of models, including time-invariant nonlinear systems with compact input constraints. Together, they provide a foundation for developing the theoretical properties of the proposed safe set, particularly in relation to time-shifted reachability, concatenated trajectories, and control invariance.

4

Properties of the Proposed Safe Set

We now propose a new definition of the safe set, shown in equation (4.8). This set contains all states that can be both reached from the initial set $\mathcal{L}$ and returned to $\mathcal{L}$ within the time horizon $[t, T]$. Formally, the safe set is defined as:

$$V_s([t, T], \mathcal{L}) := \left\{ x \in \mathbb{R}^n : \exists u \in \mathbb{U}_t^T, \exists T_1, T_2 \in [t, T] \text{ with } T_1 \leq T_2, \exists x_0 \in \mathcal{L} : \right. \\ \left. \zeta_{x_0,t}^u(T_1) = x, \quad \zeta_{x,T_1}^u(T_2) \in \mathcal{L} \right\}. \quad (4.8)$$

In contrast, existing literature such as [26–28, 31] defines the safe set as the intersection of the forward and backward reachable tubes, shown in equation (4.9). We denote this as the "old safe set":

$$V_{\text{int}}([t, T], \mathcal{L}) := V_{\text{FRT}}([t, T], \mathcal{L}) \cap V_{\text{BRT}}([t, T], \mathcal{L}) \quad (4.9)$$

Having defined the proposed safe set $V_s([t, T], \mathcal{L})$ in equation (4.8), we now establish its key theoretical properties. Specifically, we show that it guarantees both reachability to and returnability from the designated set $\mathcal{L}$ under admissible control policies (internal reachability). Unlike the standard definition V_{int} , the set V_s satisfies the control invariance condition.

To support our new definition of the safe set, we define two properties that capture different levels of controllability of an initial set $\mathcal{L} \subseteq \mathbb{R}^n$.

Definition 1 (Initial Set Properties) Let $t < T$ be a finite time horizon. We define the following:

1. The set $\mathcal{L}$ is said to satisfy the strong property SP1(t, T) if it is **viable** over the interval $[t, T]$, i.e.,

$$\forall x \in \mathcal{L}, \exists u \in \mathbb{U}_t^T, \forall \tau \in [t, T] : \zeta_{x,t}^u(\tau) \in \mathcal{L}.$$

2. The set $\mathcal{L}$ is said to satisfy the weak property $\text{SP2}(t, T)$ if every initial state in $\mathcal{L}$ can return to $\mathcal{L}$ at some future time within the horizon:

$$\forall x \in \mathcal{L}, \exists u \in \mathbb{U}_t^T, \exists \tau \in (t, T] : \zeta_{x,t}^u(\tau) \in \mathcal{L}. \quad (4.10)$$

Remark 1 The SP2 property captures a minimal form of returnability: the system only needs to be able to re-enter the initial set once, not remain within it. SP1 implies a uniform stay-inside guarantee, and is thus strictly stronger.

Lemma 1 (SP1 Implies SP2) If $\mathcal{L}$ satisfies $\text{SP1}(t, T)$, then it also satisfies $\text{SP2}(t, T)$.

Proof : Let $x \in \mathcal{L}$. By SP1 , there exists a control policy $u \in \mathbb{U}_t^T$ such that:

$$\forall \tau \in [t, T], \quad \zeta_{x,t}^u(\tau) \in \mathcal{L}.$$

Choose any $\tau \in (t, T]$. Then:

$$\zeta_{x,t}^u(\tau) \in \mathcal{L},$$

which satisfies the SP2 condition $\square$.

Corollary 1 Any result proven under the weaker assumption $\text{SP2}(t, T)$ is also valid under $\text{SP1}(t, T)$.

This hierarchy of assumptions allows us to prove safe set properties under the minimal SP2 condition, ensuring the broadest applicability of our results. In what follows, we begin by studying how reach-return trajectories define membership in the safe set and then examine the set's control invariance properties over varying horizons.

To that end, we first formalise a basic monotonicity result for control invariant sets, which will be instrumental in establishing the control invariance of our proposed safe set under shorter horizons.

Theorem 1 Let $T_2 \geq T_1 \geq t$. Then the (finite-horizon) control invariant sets defined in equation (4.3) satisfy

$$V_{CI}([t, T_2], \mathcal{K}) \subseteq V_{CI}([t, T_1], \mathcal{K}).$$

Proof : Let $x \in V_{CI}([t, T_2], \mathcal{K})$. By definition of the (finite-horizon) control invariant set in equation (4.3), there exists a control policy $u \in \mathbb{U}_t^{T_2}$ such that

$$\forall \tau \in [t, T_2], \quad \zeta_{x,t}^u(\tau) \in \mathcal{K}.$$

In particular, this implies that

$$\forall \tau \in [t, T_1], \quad \zeta_{x,t}^u(\tau) \in \mathcal{K}.$$

Consider now the restriction of u to the shorter time interval $[t, T_1]$, which we denote by $u|_{[t, T_1]}$. This restricted control is admissible on $[t, T_1]$, i.e., $u|_{[t, T_1]} \in \mathbb{U}_t^{T_1}$, and it keeps the trajectory inside $\mathcal{K}$ for all $\tau \in [t, T_1]$.

Therefore, x satisfies the control invariance condition over the horizon $[t, T_1]$, and hence

$$x \in V_{\text{CI}}([t, T_1], \mathcal{K}).$$

Since x was arbitrary, we conclude that

$$V_{\text{CI}}([t, T_2], \mathcal{K}) \subseteq V_{\text{CI}}([t, T_1], \mathcal{K}).$$

□

To establish that the proposed safe set $V_s([t, T], \mathcal{L})$ is itself control invariant, we now introduce two supporting lemmas. The first shows that if the initial set $\mathcal{L}$ satisfies the returnability condition $\text{SP2}(t, \bar{T})$, then every point in $\mathcal{L}$ must also lie within the safe set. The second lemma demonstrates that the entire trajectory from any point in the safe set back to $\mathcal{L}$ remains contained within the safe set. Together, these results are key to proving the control invariance of $V_s([t, T], \mathcal{L})$ under suitable control policies.

Lemma 2 *Let $\mathcal{L}$ satisfy $\text{SP2}(t, \bar{T})$ and let $T \geq \bar{T}$. Then:*

$$\mathcal{L} \subseteq V_s([t, T], \mathcal{L}).$$

Proof : Let $x \in \mathcal{L}$. By $\text{SP2}(t, \bar{T})$, there exists a control $u \in \mathbb{U}_t^{\bar{T}}$ and a time $\tau \in (t, \bar{T}]$ such that

$$\zeta_{x,t}^u(\tau) \in \mathcal{L}.$$

Since $\bar{T} \leq T$, we can regard u as an element of $\mathbb{U}_t^T$ and we still have $\tau \in [t, T]$.

Now set

$$x_0 := x, \quad T_1 := t, \quad T_2 := \tau.$$

Then

$$\zeta_{x_0,t}^u(T_1) = x \in \mathcal{L} \quad \text{and} \quad \zeta_{x,T_1}^u(T_2) = \zeta_{x,t}^u(\tau) \in \mathcal{L}.$$

By definition of the safe set $V_s([t, T], \mathcal{L})$ (reach–return to $\mathcal{L}$ within $[t, T]$ under some admissible control), this implies

$$x \in V_s([t, T], \mathcal{L}).$$

Since $x \in \mathcal{L}$ was arbitrary, the inclusion $\mathcal{L} \subseteq V_s([t, T], \mathcal{L})$ follows.

Lemma 3 *Let $x \in V_s([t, T], \mathcal{L})$. Then there exists a control $u \in \mathbb{U}_t^T$ and a time $\tau \in [t, T]$ such that:*

$$\begin{aligned} \zeta_{x,t}^u(\tau) &\in \mathcal{L}, \quad \text{and} \\ \zeta_{x,t}^u(s) &\in V_s([t, T], \mathcal{L}) \quad \text{for all } s \in [t, \tau]. \end{aligned}$$

Proof : By definition of the safe set, there exists:

- a control input $u \in \mathbb{U}_t^T$,
- a point $x_0 \in \mathcal{L}$,
- times $T_1, T_2 \in [t, T]$ with $T_1 \leq T_2$,

such that:

$$\zeta_{x_0, t}^u(T_1) = x, \quad \text{and} \quad \zeta_{x, T_1}^u(T_2) \in \mathcal{L}.$$

Now define the composite trajectory:

$$\phi(s) := \begin{cases} \zeta_{x_0, t}^u(s) & \text{for } s \in [t, T_1], \\ \zeta_{x, T_1}^u(s) & \text{for } s \in [T_1, T_2]. \end{cases}$$

This defines a continuous trajectory from $x_0 \in \mathcal{L}$ to $\zeta_{x, T_1}^u(T_2) \in \mathcal{L}$, passing through x at $s = T_1$.

Let $\tau := T_2 - T_1 + t$. Then $\zeta_{x, t}^u(\tau) = \zeta_{x, T_1}^u(T_2) \in \mathcal{L}$ (by time invariance of the dynamics). Since each segment of the trajectory satisfies the safe set condition (reach and return to $\mathcal{L}$), and $x \in V_s([t, T], \mathcal{L})$, it follows that:

$$\zeta_{x, t}^u(s) \in V_s([t, T], \mathcal{L}) \quad \forall s \in [t, \tau].$$

□

Theorem 2 (Control invariance of the safe set) Let $\mathcal{L} \subseteq \mathbb{R}^n$ satisfy SP2($t, \bar{T}$), and let $T \geq \bar{T}$. Then, for any T_1 with $t \leq T_1 \leq T$, the safe set $V_s([t, T], \mathcal{L})$ is (finite-horizon) control invariant on $[t, T_1]$. That is,

$$\forall x \in V_s([t, T], \mathcal{L}), \exists u \in \mathbb{U}_t^{T_1}, \forall \tau \in [t, T_1] : \zeta_{x, t}^u(\tau) \in V_s([t, T], \mathcal{L}). \quad (4.11)$$

Equivalently,

$$V_s([t, T], \mathcal{L}) = V_{\text{CI}}([t, T_1], V_s([t, T], \mathcal{L})), \quad (4.12)$$

where V_{CI} is the (finite-horizon) control invariant set defined in equation (4.3).

Proof : Let

$$\mathcal{K} := V_s([t, T], \mathcal{L})$$

be the constraint set under consideration. We first show that

$$\mathcal{K} \subseteq V_{\text{CI}}([t, T_1], \mathcal{K}). \quad (4.13)$$

Step 1: Construction of a safe control for an arbitrary initial condition.

Fix an arbitrary $x_0 \in \mathcal{K}$. Following Lemma 3, there exists a control $u^{(0)} \in \mathbb{U}_t^T$ and a time $\tau_1 \in [t, T]$ such that

$$\begin{aligned} \zeta_{x_0, t}^{u^{(0)}}(\tau_1) &\in \mathcal{L} \subseteq \mathcal{K} \quad (\text{by Lemma 2}), \\ \zeta_{x_0, t}^{u^{(0)}}(s) &\in \mathcal{K} \quad \forall s \in [t, \tau_1]. \end{aligned}$$

Define

$$t_0 := t, \quad t_1 := \tau_1, \quad x_1 := \zeta_{x_0, t}^{u^{(0)}}(t_1) \in \mathcal{L} \subseteq \mathcal{K}.$$

Step 2: Iterative extension via concatenation.

Suppose we have constructed times

$$t = t_0 < t_1 < \dots < t_k \leq T,$$

states $x_j \in \mathcal{K}$ for $j = 0, \dots, k$, and a control $u^{(k)} \in \mathbb{U}_t^{t_k}$ such that

$$\zeta_{x_0, t}^{u^{(k)}}(s) \in \mathcal{K} \quad \forall s \in [t, t_k],$$

and $x_j = \zeta_{x_0, t}^{u^{(k)}}(t_j)$ for each j .

If $t_k \geq T_1$, we stop: the control $u := u^{(k)}$ restricted to $[t, T_1]$ already satisfies

$$\zeta_{x_0, t}^u(s) \in \mathcal{K} \quad \forall s \in [t, T_1],$$

and hence $x_0 \in V_{\text{CI}}([t, T_1], \mathcal{K})$.

Otherwise, if $t_k < T_1$, then in particular $x_k \in \mathcal{K}$. Apply Lemma 3 to x_k : there exists a control $\tilde{u}_k \in \mathbb{U}_{x_k}^T$ and a time $\tilde{t}_k \in [t, T]$ such that

$$\zeta_{x_k, t}^{\tilde{u}_k}(\tilde{t}_k) \in \mathcal{L} \subseteq \mathcal{K},$$

$$\zeta_{x_k, t}^{\tilde{u}_k}(s) \in \mathcal{K} \quad \forall s \in [t, \tilde{t}_k].$$

By time invariance of the dynamics (Assumption 4), we can shift $\tilde{u}_k$ to start at t_k : there exists a control segment u_k defined on an interval $[t_k, t_{k+1}]$ with

$$t_{k+1} := t_k + (\tilde{t}_k - t) \in (t_k, T],$$

such that the trajectory from (t_k, x_k) under u_k agrees with the shifted trajectory under $\tilde{u}_k$ and satisfies

$$\zeta_{x_k, t_k}^{u_k}(s) \in \mathcal{K} \quad \forall s \in [t_k, t_{k+1}], \quad \zeta_{x_k, t_k}^{u_k}(t_{k+1}) \in \mathcal{L} \subseteq \mathcal{K}.$$

We then define the extended control $u^{(k+1)}$ on $[t, t_{k+1}]$ by

$$u^{(k+1)}(s) := \begin{cases} u^{(k)}(s), & s \in [t, t_k], \\ u_k(s), & s \in (t_k, t_{k+1}]. \end{cases}$$

By construction, the trajectory under $u^{(k+1)}$ remains in $\mathcal{K}$ on $[t, t_{k+1}]$, and we set

$$x_{k+1} := \zeta_{x_0, t}^{u^{(k+1)}}(t_{k+1}) \in \mathcal{L} \subseteq \mathcal{K}.$$

Repeating this procedure as long as $t_k < T_1$ defines a control u on the entire interval $[t, T_1]$ for which

$$\zeta_{x_0, t}^u(s) \in \mathcal{K} \quad \forall s \in [t, T_1].$$

Hence $x_0 \in V_{CI}([t, T_1], \mathcal{K})$. Since $x_0 \in \mathcal{K}$ was arbitrary, we obtain the inclusion (4.13).

Step 3: Reverse inclusion and equivalence.

By the definition of the control invariant set in (4.3), the set $V_{CI}([t, T_1], \mathcal{K})$ is a subset of $\mathcal{K}$:

$$V_{CI}([t, T_1], \mathcal{K}) \subseteq \mathcal{K},$$

since membership requires $x \in \mathcal{K}$ as part of the definition.

Combining this with (4.13) gives

$$\mathcal{K} = V_{CI}([t, T_1], \mathcal{K}),$$

which is precisely (4.12) with $\mathcal{K} = V_s([t, T], \mathcal{L})$.

Finally, unwrapping the definition of V_{CI} in (4.3) yields the explicit control invariance property (4.11). This completes the proof $\square$.

Having established the control invariance of the safe set, we now turn to its structural characterisation. In particular, we show that the safe set can be expressed as a union over time-indexed intersections of forward and backward reachable tubes. To prove this result, we first establish time-invariance properties of the forward and backward reachable sets under the stated assumptions on the dynamics and control space.

Lemma 4 *Under Assumptions 1–4, the forward reachable set is time-invariant in the sense that:*

$$V_{FRT}([t, T], \mathcal{L}) = V_{FRT}([t + t_2, T + t_2], \mathcal{L}) \quad \forall t_2 \in \mathbb{R}.$$

Proof : By definition:

$$V_{FRT}([t, T], \mathcal{L}) = \left\{ x : \exists u \in \mathbb{U}_t^T, \exists \tau \in [t, T], \exists x_0 \in \mathcal{L}, \zeta_{x_0, t}^u(\tau) = x \right\}.$$

Using Assumptions 3 (time-shift invariance of control space) and 4 (time-invariance of trajectories), we get:

$$\zeta_{x_0, t}^u(\tau) = \zeta_{x_0, t+t_2}^u(\tau + t_2),$$

and

$$\mathbb{U}_t^T = \mathbb{U}_{t+t_2}^{T+t_2}.$$

Therefore, the same reachable states are obtained under the shifted time interval. Hence:

$$V_{FRT}([t, T], \mathcal{L}) = V_{FRT}([t + t_2, T + t_2], \mathcal{L}).$$

Lemma 5 *Under Assumptions 1–4, the backward reachable set is time-invariant in the sense that:*

$$V_{BRT}([t, T], \mathcal{L}) = V_{BRT}([t + t_2, T + t_2], \mathcal{L}) \quad \forall t_2 \in \mathbb{R}.$$

Proof : Follows directly from the same reasoning as Lemma 4, using time invariance of dynamics and control domain.

Theorem 3 (Safe Set as a Union of Intersected Reachable Tubes) *Let the forward reachable tube over $[t, \tau]$ be defined as:*

$$V_{\text{FRT}}([t, \tau], \mathcal{L}) := \left\{ x \in \mathbb{R}^n : \exists u \in \mathbb{U}_t^\tau, \exists s \in [t, \tau], \exists x_0 \in \mathcal{L}, \zeta_{x_0, t}^u(s) = x \right\},$$

and the backward reachable tube over $[\tau, T]$ as:

$$V_{\text{BRT}}([\tau, T], \mathcal{L}) := \left\{ x \in \mathbb{R}^n : \exists u \in \mathbb{U}_\tau^T, \exists s \in [\tau, T], \zeta_{x, \tau}^u(s) \in \mathcal{L} \right\}.$$

Then the safe set satisfies:

$$V_s([t, T], \mathcal{L}) = \bigcup_{\tau \in [t, T]} (V_{\text{FRT}}([t, \tau], \mathcal{L}) \cap V_{\text{BRT}}([\tau, T], \mathcal{L})).$$

Proof : We prove both inclusions.

(1) $\subseteq$ direction: Let $x \in V_s([t, T], \mathcal{L})$. Then, by definition of the safe set, there exist:

- a control input $u \in \mathbb{U}_t^T$,
- an initial point $x_0 \in \mathcal{L}$,
- times $T' \in [t, T]$ and $\tau \in [T', T]$,

such that:

$$\zeta_{x_0, t}^u(T') = x, \quad \zeta_{x, T'}^u(\tau) \in \mathcal{L}.$$

Then:

- $x \in V_{\text{FRT}}([t, T'], \mathcal{L})$, since it lies on a trajectory from $\mathcal{L}$ within $[t, T']$,
- $x \in V_{\text{BRT}}([T', T], \mathcal{L})$, since starting from x at time T' , the trajectory returns to $\mathcal{L}$ by time $\tau \in [T', T]$.

Letting $\tau := T'$, we conclude:

$$x \in V_{\text{FRT}}([t, \tau], \mathcal{L}) \cap V_{\text{BRT}}([\tau, T], \mathcal{L}),$$

so:

$$x \in \bigcup_{\tau \in [t, T]} (V_{\text{FRT}}([t, \tau], \mathcal{L}) \cap V_{\text{BRT}}([\tau, T], \mathcal{L})).$$

(2) $\supseteq$ direction: Let:

$$x \in V_{\text{FRT}}([t, \tau], \mathcal{L}) \cap V_{\text{BRT}}([\tau, T], \mathcal{L}) \quad \text{for some } \tau \in [t, T].$$

Then there exist: - $x_0 \in \mathcal{L}$, $u_1 \in \mathbb{U}_t^\tau$, and $s_1 \in [t, \tau]$ such that $\zeta_{x_0, t}^{u_1}(s_1) = x$; - $u_2 \in \mathbb{U}_\tau^T$, and $s_2 \in [\tau, T]$ such that $\zeta_{x, \tau}^{u_2}(s_2) \in \mathcal{L}$.

We now define a concatenated control input $u \in \mathbb{U}_t^T$ as:

$$u(s) := \begin{cases} u_1(s), & \text{for } s \in [t, \tau], \\ u_2(s), & \text{for } s \in (\tau, T]. \end{cases}$$

This is well-defined and admissible due to Assumption 2 (control space is time-invariant) and measurability.

Let $T' := s_1 \in [t, \tau]$. Then:

$$\zeta_{x_0, t}^u(T') = x.$$

To complete the safe set definition, we must show that starting from x at time T' , the system returns to $\mathcal{L}$ at some later time. We already know:

$$\zeta_{x, \tau}^u(s_2) \in \mathcal{L},$$

but this return is from time τ , not T' .

To fix this, we use Assumption 4 (time-invariance of the dynamics). Define:

$$\tau' := s_2 - (\tau - T') \in [T', T].$$

Then:

$$\zeta_{x, \tau}^u(s_2) = \zeta_{x, T'}^u(\tau') \in \mathcal{L}.$$

So we have: $\zeta_{x_0, t}^u(T') = x$, $\zeta_{x, T'}^u(\tau') \in \mathcal{L}$,
and hence $x \in V_s([t, T], \mathcal{L})$.

Conclusion: Both inclusions hold. Therefore,

$$V_s([t, T], \mathcal{L}) = \bigcup_{\tau \in [t, T]} (V_{\text{FRT}}([t, \tau], \mathcal{L}) \cap V_{\text{BRT}}([\tau, T], \mathcal{L})).$$

Assumption 5 (Internal controllability of $\mathcal{L}$ with bounded bridging time) *There exists $\delta_{\mathcal{L}} \geq 0$ such that for any start time τ and any $x_1, x_2 \in \mathcal{L}$, there exist a duration $\delta \leq \delta_{\mathcal{L}}$ and a control $u \in \mathbb{U}_\tau^{\tau+\delta}$ for which, with $\alpha := \tau + \delta$,*

$$\zeta_{x_1, \tau}^u(s) \in \mathcal{L} \quad \forall s \in [\tau, \alpha], \quad \zeta_{x_1, \tau}^u(\alpha) = x_2.$$

Remark 2 *Assumption 5, formalizes the ability to “re-dock” at any point of $\mathcal{L}$ within a uniformly bounded time $\delta_{\mathcal{L}}$. It is strictly stronger than SP2 and independent of SP1.*

Theorem 3 provides a structural decomposition of the safe set as a union of intersected reachable tubes, making explicit its dependency on both forward and backward dynamics relative to the set $\mathcal{L}$. Building on this characterisation,

we now show that the safe set supports $\mathcal{L}$ -mediated pairwise reachability with explicit horizon accounting: under the internal controllability assumption on $\mathcal{L}$, the system can transition between any two states in the safe set by first reaching $\mathcal{L}$, performing a bounded transfer within $\mathcal{L}$, and then departing to the target state, while staying entirely within the safe set.

Inbound/Outbound durations: For $x \in V_s([t, T], \mathcal{L})$ define the inbound duration

$$\tau_{\text{in}}(x) := \inf\{\Delta \in [0, T - t] : \exists u \in \mathbb{U}_t^{t+\Delta} \text{ s.t. } \zeta_{x,t}^u(t + \Delta) \in \mathcal{L}\}.$$

For $x' \in V_s([t, T], \mathcal{L})$ define the outbound duration

$$\tau_{\text{out}}(x') := \inf\{\Delta \in [0, T - t] : \exists u \in \mathbb{U}_t^{t+\Delta}, \exists x_L \in \mathcal{L} \text{ s.t. } \zeta_{x_L,t}^u(t + \Delta) = x'\}.$$

Theorem 4 ($\mathcal{L}$ -Mediated Pairwise Reachability) Assume $\mathcal{L}$ satisfies the internal controllability condition (Assumption 5). Suppose SP2($t, \tilde{T}$) holds and let $T - t \geq \tilde{T}$. Then, for any $x, x' \in V_s([t, T], \mathcal{L})$, there exists a control $u \in \mathbb{U}_t^\alpha$ and a time

$$\alpha := t + \tau_{\text{in}}(x) + \delta_{\mathcal{L}} + \tau_{\text{out}}(x') \in [t, t + 2(T - t) + \delta_{\mathcal{L}}]$$

such that

$$\zeta_{x,t}^u(\alpha) = x' \quad \text{and} \quad \zeta_{x,t}^u(s) \in V_s([t, T], \mathcal{L}) \quad \forall s \in [t, \alpha].$$

Moreover, if $\tau_{\text{in}}(x) + \delta_{\mathcal{L}} + \tau_{\text{out}}(x') \leq T - t$, then the transition is achieved within the original horizon $[t, T]$.

Proof. Let $x, x' \in V_s([t, T], \mathcal{L})$.

Step 1 (Inbound to $\mathcal{L}$). By Lemma 3, $\exists u_1 \in \mathbb{U}_t^T$ and $\tau_1 \in [t, T]$ with

$$\zeta_{x,t}^{u_1}(\tau_1) = x_{\text{in}} \in \mathcal{L}, \quad \zeta_{x,t}^{u_1}(s) \in V_s([t, T], \mathcal{L}) \quad \forall s \in [t, \tau_1].$$

Set $\tau_{\text{in}}(x) := \tau_1 - t$.

Step 2 (Bounded transfer in $\mathcal{L}$). By Assumption 5, $\exists u_L$ and $\delta \leq \delta_{\mathcal{L}}$ such that

$$\zeta_{x_{\text{in}},\tau_1}^{u_L}(s) \in \mathcal{L} \subseteq V_s([t, T], \mathcal{L}) \quad \forall s \in [\tau_1, \tau_1 + \delta], \quad \zeta_{x_{\text{in}},\tau_1}^{u_L}(\tau_1 + \delta) = x_{\text{out}} \in \mathcal{L}.$$

Step 3 (Outbound from $\mathcal{L}$ to x'). By Lemma 2, $\exists u_2 \in \mathbb{U}_t^T$ and $\tau_2 \in [t, T]$ with

$$\zeta_{x_{\text{out}},t}^{u_2}(\tau_2) = x', \quad \zeta_{x_{\text{out}},t}^{u_2}(s) \in V_s([t, T], \mathcal{L}) \quad \forall s \in [t, \tau_2].$$

By time invariance (Assumption 4), define the shift $\tilde{u}_2(s) := u_2(s - \tau_1 - \delta + t)$ for $s \in [\tau_1 + \delta, \tau_1 + \delta + \tau_2 - t]$. Set $\tau_{\text{out}}(x') := \tau_2 - t$.

Step 4 (Concatenation and budget). Define

$$u(s) = \begin{cases} u_1(s), & s \in [t, \tau_1], \\ u_L(s), & s \in (\tau_1, \tau_1 + \delta], \\ \tilde{u}_2(s), & s \in (\tau_1 + \delta, \tau_1 + \delta + \tau_2 - t], \end{cases}$$

and set $\alpha := \tau_1 + \delta + \tau_2 - t = t + \tau_{\text{in}}(x) + \delta_{\mathcal{L}} + \tau_{\text{out}}(x')$. Then $\alpha \leq t + 2(T - t) + \delta_{\mathcal{L}}$ and

$$\zeta_{x,t}^u(\alpha) = x', \quad \zeta_{x,t}^u(s) \in V_s([t, T], \mathcal{L}) \quad \forall s \in [t, \alpha].$$

If $\tau_{\text{in}}(x) + \delta_{\mathcal{L}} + \tau_{\text{out}}(x') \leq T - t$, then $\alpha \leq T$ and the transition fits within $[t, T]$.

While Theorem 4 establishes $\mathcal{L}$ -mediated maneuverability between elements of the safe set under Assumption 5, it is also important to relate this set to the standard definition based on intersecting reachable tubes. The following lemma shows that every state in the classical intersection-based set V_{int} also satisfies the round-trip reachability condition of our proposed safe set, albeit potentially over a longer time horizon.

Lemma 6 *The intersection set*

$$V_{\text{int}}([t, T], \mathcal{L}) := V_{\text{FRT}}([t, T], \mathcal{L}) \cap V_{\text{BRT}}([t, T], \mathcal{L})$$

is contained in the safe set over an extended horizon:

$$V_{\text{int}}([t, T], \mathcal{L}) \subseteq V_s([t, 2T - t], \mathcal{L}).$$

Proof: Let $x \in V_{\text{int}}([t, T], \mathcal{L})$. Then, by definition,

$$x \in V_{\text{FRT}}([t, T], \mathcal{L}) \quad \text{and} \quad x \in V_{\text{BRT}}([t, T], \mathcal{L}).$$

From the forward reachable tube membership, there exist

$$u_f \in \mathbb{U}_t^T, \quad T_1 \in [t, T], \quad x_0 \in \mathcal{L}$$

such that

$$\zeta_{x_0,t}^{u_f}(T_1) = x. \quad (4.14)$$

From the backward reachable tube membership, there exist

$$u_b \in \mathbb{U}_t^T, \quad \tau \in [t, T]$$

such that

$$\zeta_{x,t}^{u_b}(\tau) \in \mathcal{L}. \quad (4.15)$$

We now use time invariance of the dynamics (Assumption 4) to shift the backward trajectory so that it starts at time T_1 instead of t . Define the time shift

$$t_2 := T_1 - t,$$

and a shifted control

$$\tilde{u}_b(s) := u_b(s - t_2), \quad s \in [T_1, T_1 + (T - t)].$$

By Assumption 4, the corresponding trajectories satisfy

$$\zeta_{x,t}^{u_b}(\tau) = \zeta_{x,T_1}^{\tilde{u}_b}(T_1 + (\tau - t)).$$

Set

$$T_2 := T_1 + (\tau - t).$$

Since $\tau \in [t, T]$, we have

$$T_2 \in [T_1, T_1 + (T - t)] \subseteq [t, 2T - t].$$

Combining with (4.15), this gives

$$\zeta_{x, T_1}^{\tilde{u}_b}(T_2) \in \mathcal{L}, \quad (4.16)$$

with $T_1 \leq T_2 \leq 2T - t$.

We now construct a single control $u \in \mathbb{U}_t^{2T-t}$ by concatenating u_f and $\tilde{u}_b$:

$$u(s) := \begin{cases} u_f(s), & s \in [t, T_1], \\ \tilde{u}_b(s), & s \in (T_1, T_2], \\ \tilde{u}(s), & s \in (T_2, 2T - t], \end{cases}$$

where $\tilde{u}(\cdot)$ is any admissible control on $(T_2, 2T - t]$ (its specific choice is irrelevant for the round-trip construction). Since all pieces take values in the admissible input set and are measurable, the concatenated u is admissible on $[t, 2T - t]$, i.e. $u \in \mathbb{U}_t^{2T-t}$.

Under this control u , the trajectory from x_0 satisfies

$$\zeta_{x_0, t}^u(T_1) = \zeta_{x_0, t}^{u_f}(T_1) = x$$

by (4.14), and

$$\zeta_{x, T_1}^u(T_2) = \zeta_{x, T_1}^{\tilde{u}_b}(T_2) \in \mathcal{L}$$

by (4.16). Thus we have found

$$u \in \mathbb{U}_t^{2T-t}, \quad x_0 \in \mathcal{L}, \quad T_1, T_2 \in [t, 2T - t], \quad T_1 \leq T_2$$

such that

$$\zeta_{x_0, t}^u(T_1) = x, \quad \zeta_{x, T_1}^u(T_2) \in \mathcal{L}.$$

By the definition of the safe set $V_s([t, 2T - t], \mathcal{L})$ in (4.8), this implies

$$x \in V_s([t, 2T - t], \mathcal{L}).$$

Since $x \in V_{\text{int}}([t, T], \mathcal{L})$ was arbitrary, we conclude

$$V_{\text{int}}([t, T], \mathcal{L}) \subseteq V_s([t, 2T - t], \mathcal{L}).$$

□

Lemma 6 shows that the classical intersection-based set V_{int} is always contained within the proposed safe set V_s over an extended time horizon. However, this inclusion is generally strict. The next theorem provides a necessary and sufficient condition under which the two sets coincide, namely, when every state in the safe set can both be reached from and returned to $\mathcal{L}$ within the original time horizon $[t, T]$. This reveals that the classical definition only recovers the true safe set under strong symmetry assumptions on time-to-reach and time-to-return.

Theorem 5 *We have equality between the intersection set and the extended safe set:*

$$V_{\text{int}}([t, T], \mathcal{L}) = V_s([t, 2T - t], \mathcal{L})$$

if and only if for every $x \in V_s([t, 2T - t], \mathcal{L})$, both the reach time from $\mathcal{L}$ to x and the return time from x to $\mathcal{L}$ lie within the time interval $[t, T]$.

Proof : **If the condition holds, then equality follows:**

Assume that for every $x \in V_s([t, 2T - t], \mathcal{L})$, there exists:

$$T_1, T_2 \in [t, T], x_0 \in \mathcal{L}, u \in \mathbb{U}_t^T \text{ such that } \zeta_{x_0, t}^u(T_1) = x, \zeta_{x, t}^u(T_2) \in \mathcal{L}.$$

Then clearly: $x \in V_{\text{FRT}}([t, T], \mathcal{L}), x \in V_{\text{BRT}}([t, T], \mathcal{L})$

So $x \in V_{\text{int}}([t, T], \mathcal{L})$

Hence:

$$V_s([t, 2T - t], \mathcal{L}) \subseteq V_{\text{int}}([t, T], \mathcal{L}).$$

From Lemma 6, we already have the reverse inclusion.

Therefore:

$$V_s([t, 2T - t], \mathcal{L}) = V_{\text{int}}([t, T], \mathcal{L}).$$

If sets are equal, then reach and return must occur within $[t, T]$:

Assume $V_s([t, 2T - t], \mathcal{L}) = V_{\text{int}}([t, T], \mathcal{L})$.

Let $x \in V_s([t, 2T - t], \mathcal{L})$. Then:

$$x \in V_{\text{FRT}}([t, T], \mathcal{L}) \Rightarrow \exists u \in \mathbb{U}_t^T, \exists T' \in [t, T], \zeta_{x_0, t}^u(T') = x,$$

$$x \in V_{\text{BRT}}([t, T], \mathcal{L}) \Rightarrow \exists u \in \mathbb{U}_t^T, \exists \tau \in [t, T], \zeta_{x, t}^u(\tau) \in \mathcal{L}.$$

Therefore, both reach and return occur within $[t, T]$ for all $x \in V_s([t, 2T - t], \mathcal{L})$.

4.3.2. Bounded Disturbances

We now extend the results established for the disturbance-free case to the setting where the system is subject to bounded disturbances. This extension is particularly relevant for practical applications, where uncertainties or adversarial effects must be accounted for in safety analysis.

The system dynamics in the presence of bounded disturbances are given by:

$$\dot{x} = f(x, u, d, t), \quad (4.17)$$

where $x \in \mathbb{X} \subseteq \mathbb{R}^m$ is the system state, $u(\cdot) \in \mathbb{U} \subseteq \mathbb{F}$ is the control policy, and $d(\cdot) \in \mathbb{D} \subseteq \mathbb{F}$ is the disturbance policy. The spaces $\mathbb{X}, \mathbb{U}, \mathbb{D}$ are compact, and $\mathbb{F}$ denotes the space of measurable functions. For a time interval $[t, T]$, the admissible control and disturbance policy sets are denoted $\mathbb{U}_t^T$ and $\mathbb{D}_t^T$, respectively. The solution to (4.17) under input policy u and disturbance d , starting from initial state x at time t , evaluated at time τ , is denoted by $\zeta_{x, t}^{u, d}(\tau)$.

We adopt standard assumptions: The function f is Lipschitz continuous in x , uniformly continuous in (u, d, t) . Both control and disturbance sets $\mathbb{U}_t^T, \mathbb{D}_t^T$ are compact and time-invariant:

$$\mathbb{U}_t^T = \mathbb{U}_{t+t_2}^{T+t_2}, \quad \mathbb{U}_t^{T_1} \subseteq \mathbb{U}_t^{T_2} \text{ for } T_2 \geq T_1,$$

with analogous properties for $\mathbb{D}$. The system is time-invariant:

$$\zeta_{x,t}^{u,d}(\tau) = \zeta_{x,t+t_2}^{u,d}(\tau + t_2) \quad \forall t_2 \in \mathbb{R}.$$

In the robust control framework, the disturbance acts adversarially via a non-anticipative strategy [3]:

$$\Lambda_t^T := \{ \mathcal{N} : \mathbb{U}_t^T \rightarrow \mathbb{D}_t^T \mid \forall u, \hat{u} \in \mathbb{U}_t^T, \forall s \in [t, T], (u = \hat{u} \text{ a.e. on } [t, s]) \Rightarrow \mathcal{N}[u](s) = \mathcal{N}[\hat{u}](s) \}.$$

We write the disturbance signal induced by a control u as $d = \mathcal{N}[u]$, with $\mathcal{N} \in \Lambda_t^T$.

That is, player 2 (the disturbance) cannot respond to future values of the control input; they can only react based on its current and past values. However, player 2 has the advantage of observing the current control action before responding, reflecting the worst-case assumption typically made in robust formulations. We now define the robust counterparts of the forward and backward reachable tubes.

Robust forward reachable tube: the set of states for which there exists at least one control policy, for all the possible disturbances, that will drive the system into the set from its initial set $\mathcal{L}$ within the given time horizon $[t, T]$, as follows:

$$V_{\text{RFRT}}([t, T], \mathcal{L}) := \{x : \exists u \in \mathbb{U}_t^T, \forall \mathcal{N} \in \Lambda_t^T, \exists s \in [t, T], \exists x_0 \in \mathcal{L}, \zeta_{x_0,t}^{u,\mathcal{N}[u]}(s) = x\}. \quad (4.18)$$

Robust backward reachable tube: the set of all the initial states from which there exists at least one control policy, for all possible disturbances, that will lead the system into the required set $\mathcal{L}$ within the given time horizon $[t, T]$, as follows:

$$V_{\text{RBRT}}([t, T], \mathcal{L}) := \{x : \exists u \in \mathbb{U}_t^T, \forall \mathcal{N} \in \Lambda_t^T, \exists s \in [t, T], \zeta_{x,t}^{u,\mathcal{N}[u]}(s) \in \mathcal{L}\} \quad (4.19)$$

Robust control invariant set: Let $\mathcal{K} \subseteq \mathbb{R}^m$ be a constraint set and consider the disturbed system (4.17) with non-anticipative disturbance strategies Λ_t^T . The (finite-horizon) robust control invariant set with respect to $\mathcal{K}$ over the horizon $[t, T]$ is defined as

$$V_{\text{RCI}}([t, T], \mathcal{K}) := \{x \in \mathcal{K} : \exists u \in \mathbb{U}_t^T, \forall \mathcal{N} \in \Lambda_t^T, \forall \tau \in [t, T], \zeta_{x,t}^{u,\mathcal{N}[u]}(\tau) \in \mathcal{K}\}. \quad (4.20)$$

Robust Safe Set: We now define the robust safe set analogously to the disturbance-free case, with the added requirement that the reach-return condition holds uniformly over all disturbances.

$$V_{RS}([t, T], \mathcal{L}) := \{x : \exists u \in \mathbb{U}_t^T, \forall \mathcal{N} \in \Lambda_t^T, \exists T' \in [t, T], \exists \tau \in [T', T], \exists x_0 \in \mathcal{L}, \\ \zeta_{x_0, t}^{u, \mathcal{N}[u]}(T') = x, \quad \zeta_{x, T'}^{u, \mathcal{N}[u]}(\tau) \in \mathcal{L}\}. \quad (4.21)$$

That is, the system must be able to reach x from $\mathcal{L}$, and return to $\mathcal{L}$ from x , via a single control policy u , under all possible disturbance strategies.

Remark 3 *The robust safe set guarantees safety under the worst-case disturbance policy. In contrast to the disturbance-free case, the reach-return condition must be satisfied for all non-anticipative disturbance signals $d \in \Lambda_t^T$, which makes the problem inherently more conservative but better suited for safety-critical applications.*

Lemma 7 (Family of effective dynamics) *For each non-anticipative strategy $\mathcal{N} \in \Lambda_t^T$, define*

$$\bar{f}_{\mathcal{N}}(x, u, t) := f(x, u, \mathcal{N}[u](t), t).$$

Then trajectories under $(u, \mathcal{N})$ solve $\dot{x} = \bar{f}_{\mathcal{N}}(x, u, t)$. Moreover, if f is Lipschitz in x uniformly in (u, d, t) , then each $\bar{f}_{\mathcal{N}}$ is Lipschitz in x uniformly in (u, t) .

This is a notational device: all robust statements must still hold for every $\mathcal{N} \in \Lambda_t^T$. No single $\bar{f}_{\mathcal{N}}$ can replace the universal quantifier.

Assumption 6 (Robust internal controllability of $\mathcal{L}$) *There exists $\delta_{\mathcal{L}}^{\text{rob}} \geq 0$ such that for all start times σ and all $x_1, x_2 \in \mathcal{L}$ there exist $\delta \leq \delta_{\mathcal{L}}^{\text{rob}}$ and $u \in \mathbb{U}_{\sigma}^{\sigma+\delta}$ with*

$$\forall \mathcal{N} \in \Lambda_{\sigma}^{\sigma+\delta} : \zeta_{x_1, \sigma}^{u, \mathcal{N}[u]}(s) \in \mathcal{L} \quad \forall s \in [\sigma, \sigma + \delta], \quad \zeta_{x_1, \sigma}^{u, \mathcal{N}[u]}(\sigma + \delta) = x_2.$$

Theorem 6 (Robust pairwise reachability) *Assume SP2($t, \bar{T}$) and $T - t \geq \bar{T}$. Under Assumption 6, for any $x, x' \in V_{RS}([t, T], \mathcal{L})$ there exists a control u and a time*

$$\alpha = t + \tau_{\text{in}}^{\text{rob}}(x) + \delta_{\mathcal{L}}^{\text{rob}} + \tau_{\text{out}}^{\text{rob}}(x')$$

such that, for all $\mathcal{N} \in \Lambda_t^{\alpha}$,

$$\zeta_{x, t}^{u, \mathcal{N}[u]}(\alpha) = x', \quad \zeta_{x, t}^{u, \mathcal{N}[u]}(s) \in V_{RS}([t, T], \mathcal{L}) \quad \forall s \in [t, \alpha],$$

with $\tau_{\text{in}}^{\text{rob}}(x), \tau_{\text{out}}^{\text{rob}}(x') \leq T - t$. Hence $\alpha \leq t + 2(T - t) + \delta_{\mathcal{L}}^{\text{rob}}$. If $\tau_{\text{in}}^{\text{rob}}(x) + \delta_{\mathcal{L}}^{\text{rob}} + \tau_{\text{out}}^{\text{rob}}(x') \leq T - t$, the transition fits within $[t, T]$.

Conclusion Replacing each nominal tube by its robust counterpart (RFRT/RBRT and V_{RS}) preserves the structural results for every non-anticipative disturbance strategy. In particular, safe set results and under Assumption 6, the pairwise-reachability result holds.

4.4. NUMERICAL EXAMPLES

To illustrate and validate the theoretical results developed in this paper, we consider three benchmark systems commonly used in reachability-based safety analysis: the double integrator, a standard second-order model representing a 2D cart; the Dubins car, a nonlinear 3D system that captures curvature-constrained motion; and a point-mass aircraft model based on the RCAM (Research Civil Aircraft Model) slow dynamics, which has been extensively used for flight envelope estimation and loss-of-control prevention [11, 27]. These systems are selected for their simplicity, controllability, and relevance in practical scenarios such as mobile robotics, aerial vehicles, and transport aircraft.

Our objective is to compute the proposed safe set for each system and compare it against the classical intersection-based safe set. We adopt Hamilton–Jacobi (HJ) reachability analysis as our numerical framework, assuming full knowledge of the system dynamics. The HJ formalism enables the computation of forward and backward reachable tubes under optimal control policies, along with the corresponding value functions [32–35].

To solve the HJ partial differential equations (PDEs), we employ level set methods, which are widely recognised as a robust and accurate approach for reachable set estimation [36, 37]. These methods discretise the state space and evolve a value function over time to capture the evolving boundary of the reachable region. We implement this using the Level Set Toolbox and the HelperOC Toolbox [4, 38], which provide efficient and well-documented solvers tailored for safety-critical applications. These examples also demonstrate that the proposed safe set can be constructed directly using existing reachability tools, offering a practical pathway for implementation in real systems.

4.4.1. System Dynamics: Double Integrator

We consider a standard double integrator system, often used to model a simplified 2D cart or velocity-controlled robot. The system is described by the following dynamics:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ u \end{bmatrix}, \quad u \in \mathbb{U} := [-1, 1], \quad (4.22)$$

where x_1 denotes position, x_2 denotes velocity, and u is the control input constrained within a compact set $\mathbb{U}$.

To compute the safe set, we assume a circular target set (or operating region) defined as:

$$\mathcal{L} := \{x \in \mathbb{R}^2 : \|x\|_2 \leq R\}, \quad \text{with } R > 0, \quad (4.23)$$

where the origin corresponds to the nominal operating point, and R specifies the safe neighbourhood radius. To compute this set numerically, we adopt Hamilton Jacobi reachability analysis. Specifically, we first compute the forward reachable tube $V_{\text{FRT}}([t, \tau], \mathcal{L})$ from $\mathcal{L}$ using the system dynamics. We then compute the backward reachable tube $V_{\text{BRT}}([\tau, T], \mathcal{L})$ back into $\mathcal{L}$ and finally construct the safe set as shown in [Theorem 3](#). The resulting set is computed using HJ PDE solver : Level Set Toolbox [13] and the HelperOC Toolbox [33].

4.4.2. System Dynamics: Dubins Car

To further demonstrate the applicability of our proposed safe set construction, we consider the Dubins car model, a widely studied nonlinear system representing a unicycle with bounded turning rate. This system is relevant for applications such as mobile robot navigation, autonomous vehicles, and air traffic control where curvature constraints play a critical role.

The dynamics of the Dubins car are given by:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \\ u \end{bmatrix}, \quad u \in \mathbb{U} := [-1, 1], \quad (4.24)$$

4

where (x_1, x_2) denotes the position in the plane, θ the orientation, and u is the angular control input.

We define the target set as a small disc of radius $R > 0$ in the spatial dimensions centred at the origin with arbitrary heading:

$$\mathcal{L} := \{(x_1, x_2, \theta) \in \mathbb{R}^2 \times S^1 : \sqrt{x_1^2 + x_2^2} \leq R\}. \quad (4.25)$$

As with the double integrator case, we compute the safe set over a finite time horizon $[t, T]$ using the round-trip reachability formulation. Specifically, we calculate the forward reachable tube $V_{\text{FRT}}([t, \tau], \mathcal{L})$ and backward reachable tube $V_{\text{BRT}}([\tau, T], \mathcal{L})$ and then obtain the safe set.

4.4.3. System Dynamics: RCAM Slow Flight Envelope Model

Finally, we consider a more realistic aircraft example based on the slow dynamics of the Research Civil Aircraft Model (RCAM), following the physically motivated envelope-estimation framework of [27]. The goal is to characterise a safe manoeuvring envelope in the (V, γ) -plane, where V denotes airspeed and γ the flight-path angle.

The full RCAM model is nine-dimensional, but a time-scale separation argument allows the longitudinal motion in the vertical plane to be approximated by a two-dimensional point-mass model for (V, γ) , with roll angle ϕ , angle of attack α , sideslip β , and thrust T treated as (virtual) control inputs [27]. The resulting slow dynamics can be written as

$$x := \begin{bmatrix} V \\ \gamma \end{bmatrix}, \quad u := \begin{bmatrix} T \\ \alpha \\ \beta \\ \phi \end{bmatrix}, \quad (4.26)$$

and

$$\dot{x} = f(x, u) = \begin{bmatrix} \dot{V} \\ \dot{\gamma} \end{bmatrix}. \quad (4.27)$$

Under standard small-angle approximations for α and β , and neglecting high-frequency dynamics, the longitudinal point-mass equations reduce to (cf. [27])

$$\dot{V} = -\frac{\rho S}{2m} V^2 (C_{D0} + C_{D\alpha} \alpha + C_{D\alpha^2} \alpha^2) - g \sin \gamma + \frac{1}{m} T, \quad (4.28)$$

$$\dot{\gamma} = -\frac{g}{V} \cos \gamma + \frac{\rho S}{2m} V (C_{L0} + C_{L\alpha} \alpha) \cos \phi - \frac{\rho S}{2m} V C_{Y\beta} \beta \sin \phi, \quad (4.29)$$

where m is the aircraft mass, g the gravitational acceleration, ρ the air density, and S the wing reference area. The aerodynamic coefficients C_{D0} , $C_{D\alpha}$, $C_{D\alpha^2}$, C_{L0} , $C_{L\alpha}$, $C_{Y\beta}$ are taken from the RCAM database and represent drag, lift, and sideforce derivatives, respectively. For our numerical experiments we use the same parameter values and input bounds as in [27], corresponding to a large transport aircraft at near sea-level conditions. Typical limits are

$$T \in [T_{\min}, T_{\max}], \quad \alpha \in [\alpha_{\min}, \alpha_{\max}], \quad \beta \in [\beta_{\min}, \beta_{\max}], \quad \phi \in [\phi_{\min}, \phi_{\max}], \quad (4.30)$$

with $T_{\min} = 2.05 \times 10^4$ N, $T_{\max} = 4.11 \times 10^5$ N, $\alpha_{\min} = 0^\circ$, $\alpha_{\max} = 14.5^\circ$, $\beta_{\min} = -5^\circ$, $\beta_{\max} = 5^\circ$, and $\phi_{\min} = -60^\circ$, $\phi_{\max} = 60^\circ$. Following [27], we define a stable trim envelope as:

$$\mathcal{L} := \{(V, \gamma) \in \mathbb{R}^2 : V_{\min} \leq V \leq V_{\max}, \gamma_{\min} \leq \gamma \leq \gamma_{\max}\}, \quad (4.31)$$

with

$$V_{\min} = 60 \text{ m/s}, \quad V_{\max} = 100 \text{ m/s}, \quad \gamma_{\min} = -10^\circ, \quad \gamma_{\max} = 10^\circ. \quad (4.32)$$


Equivalently, $\mathcal{L}$ is the zero-superlevel set of the implicit function

$$\ell(x) := \min \{V - V_{\min}, V_{\max} - V, \gamma - \gamma_{\min}, \gamma_{\max} - \gamma\}, \quad (4.33)$$

4.5. RESULTS

This section presents numerical results that validate the theoretical framework developed in the preceding sections. We apply Hamilton Jacobi (HJ) reachability analysis to three representative systems: the double integrator, the Dubins car, and the RCAM based point mass aircraft model introduced in Section 4.4. For each case, we compute the forward reachable tube (FRT) and backward reachable tube (BRT) from a designated target set $\mathcal{L}$ over a finite time horizon. Using these, we construct two different notions of safe sets: the classical safe set, defined as the intersection of FRT and BRT over the same time interval, and the proposed safe set, defined as the union of their intersections across varying intermediate time splits, as formalised in Theorem 3. The HJ partial differential equations are solved numerically using level set methods, implemented via the

Level Set Toolbox and the HelperOC Toolbox [13, 33]. These tools allow efficient approximation of reachable sets by discretising the state space and evolving a value function backward or forward in time. For the double integrator and Dubins car, we fix the time horizon to $T = 5$ seconds, and perform computations over the domain $\tau \in [0, T]$. For the RCAM example, we follow the envelope estimation literature and use a shorter horizon of $T = 2$ seconds. For the double integrator system, the resulting sets are visualised in the (x_1, x_2) space. As shown in Figure 4.1, the classical intersection-based safe set $V_{\text{int}}([0, 5])$ is significantly smaller than the proposed safe set $V_s([0, 10])$, thereby confirming our theoretical result in Lemma 6. In the case of the Dubins car, the computed sets, shown in Figure 4.2a-4.3b, illustrate the progression of our analysis. The forward reachable tube (FRT) is depicted in Figure 4.2a. The backward reachable tube (BRT) is shown in Figure 4.2b. The classical safe set is presented in Figure 4.3a, and the proposed safe set obtained via round-trip reachability is shown in Figure 4.3b. For the RCAM slow-flight envelope example, where the state is given by the airspeed V and flight-path angle γ , and the target set $\mathcal{L}$ represents a rectangular trim envelope in the (V, γ) plane (see Section 4.4). The BRT and FRT for a horizon of $T = 2$ s are shown in Figures 4.4a and 4.4b. The classical safe set obtained as the intersection of FRT and BRT is plotted in Figure 4.5a. This set surrounds the trim region and includes pronounced bulges near the upper and lower left corners, corresponding to states with low speed and large pitch-up or pitch-down excursions that are both reachable and recoverable when considered separately. However, as shown in Figures 4.5b and 4.5c, many of these corner states do not admit a single round-trip trajectory that leaves $\mathcal{L}$ and returns within $T = 2$ s. The proposed safe set $V_s([0, 2], \mathcal{L})$ is therefore a strict subset of the classical intersection. Overall, the numerical results support the main theoretical claims of the paper.



4

Figure 4.1: Comparison of classical and proposed safe sets for the double integrator in the (x_1, x_2) -plane.







Figure 4.2: Forward and backward reachable tubes for the Dubins car.

(a) Classical safe set ($FRT \cap BRT$)

(b) Proposed safe set (round-trip)

Figure 4.3: Comparison of classical and proposed safe sets for the Dubins car.(a) Backward reachable tube (BRT, $T = 2$ s)(b) Forward reachable tube (FRT, $T = 2$ s)**Figure 4.4:** Forward and backward reachable tubes for the RCAM slow-flight envelope in the (V, γ) -plane.(a) Classical safe set $V_{\text{int}}([0, 2])$

(b) Classical vs. proposed safe set

(c) Proposed safe set $V_s([0, 2])$ **Figure 4.5:** Classical intersection-based safe set and proposed round-trip safe set for the RCAM slow-flight envelope.

4.6. CONCLUSIONS

Finite horizon safety cannot, in general, be captured by the intersection of forward and backward reachable tubes when safety is intended to mean control invariance, without strong symmetry assumptions, such an intersection may fail to guarantee control invariance and may not support safe connectivity between states.

A finite-horizon safe set defined via round-trip reachability, states that are reachable from a designated initial set and returnable to that set within the horizon does guarantee control invariance over all sub-horizons. Moreover, when the initial set is internally controllable, the resulting safe set supports internal reachability, any two states in the set can be connected by a trajectory that remains within the set throughout.

The round-trip safe set is equal to the union over intermediate times of intersections between forward and backward reachable tubes. This representation enables computation using standard Hamilton Jacobi reachability machinery without altering solver structure.

For systems with bounded disturbances, defining safety with non-anticipative disturbance strategies yields an analogous robust notion, the control invariance and internal reachability persist in a game theoretic (worst case) sense. Numerical studies demonstrate that the round trip definition can differ from tube intersections and leads to safety conclusions consistent with invariance and robust feasibility.

REFERENCES

- [1] J. Carlson, R. R. Murphy, and A. Nelson. “Follow-up analysis of mobile robot failures”. In: *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA’04. 2004*. Vol. 5. IEEE. 2004, pp. 4987–4994. DOI: [10.1109/robot.2004.1302508](https://doi.org/10.1109/robot.2004.1302508).
- [2] S. Yang, Y. Zhong, D. Feng, R. Y. M. Li, X.-F. Shao, and W. Liu. “Robot application and occupational injuries: are robots necessarily safer?” In: *Safety science* 147 (2022), p. 105623. DOI: [10.1016/j.ssci.2021.105623](https://doi.org/10.1016/j.ssci.2021.105623).
- [3] S. L. Herbert. *Safe real-world autonomy in uncertain and unstructured environments*. University of California, Berkeley, 2020.
- [4] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin. “A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games”. In: *IEEE Transactions on automatic control* 50.7 (2005), pp. 947–957. DOI: [10.1109/tac.2005.851439](https://doi.org/10.1109/tac.2005.851439).
- [5] F. Blanchini. “Set invariance in control”. In: *Automatica* 35.11 (1999), pp. 1747–1767. DOI: [10.1016/s0005-1098\(99\)00113-2](https://doi.org/10.1016/s0005-1098(99)00113-2).
- [6] M. Fiacchini, T. Alamo, and E. F. Camacho. “On the computation of convex robust control invariant sets for nonlinear systems”. In: *Automatica* 46.8 (2010), pp. 1334–1338. DOI: [10.1016/j.automatica.2010.05.007](https://doi.org/10.1016/j.automatica.2010.05.007).
- [7] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre. *Viability theory: new directions*. Springer Science & Business Media, 2011. DOI: [10.1007/978-3-642-16684-6](https://doi.org/10.1007/978-3-642-16684-6).
- [8] J.-P. Aubin and H. Frankowska. “Viability kernel of control systems”. In: *Nonlinear synthesis*. Springer, 1991, pp. 12–33. DOI: [10.1007/978-1-4757-2135-5_2](https://doi.org/10.1007/978-1-4757-2135-5_2).
- [9] P. Cardaliaguet. “A differential game with two players and one target”. In: *SIAM Journal on Control and Optimization* 34.4 (1996), pp. 1441–1460. DOI: [10.1137/s036301299427223x](https://doi.org/10.1137/s036301299427223x).
- [10] P. Cardaliaguet, M. Quincampoix, and P. Saint-Pierre. “Differential games through viability theory: Old and recent results”. In: *Advances in dynamic game theory*. Springer, 2007, pp. 3–35.
- [11] J. Lygeros. “On reachability and minimum cost optimal control”. In: *Automatica* 40.6 (2004), pp. 917–927.

- [12] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-jacobi reachability: A brief overview and recent advances”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 2242–2253. DOI: [10.1109/cdc.2017.8263977](https://doi.org/10.1109/cdc.2017.8263977).
- [13] I. M. Mitchell and J. A. Templeton. “A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems”. In: *Hybrid Systems: Computation and Control: 8th International Workshop, HSCC 2005, Zurich, Switzerland, March 9-11, 2005. Proceedings 8*. Springer. 2005, pp. 480–494. DOI: [10.1007/978-3-540-31954-2_31](https://doi.org/10.1007/978-3-540-31954-2_31).
- [14] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin. “FaSTrack: A modular framework for fast and guaranteed safe motion planning”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 1517–1522. DOI: [10.1109/cdc.2017.8263867](https://doi.org/10.1109/cdc.2017.8263867).
- [15] N. Kariotoglou, K. Margellos, and J. Lygeros. “On the computational complexity and generalization properties of multi-stage and stage-wise coupled scenario programs”. In: *Systems & Control Letters* 94 (2016), pp. 63–69. DOI: [10.1016/j.sysconle.2016.05.009](https://doi.org/10.1016/j.sysconle.2016.05.009).
- [16] K. Lesser, M. Oishi, and R. S. Erwin. “Stochastic reachability for control of spacecraft relative motion”. In: *52nd IEEE Conference on Decision and Control*. IEEE. 2013, pp. 4705–4712. DOI: [10.1109/cdc.2013.6760626](https://doi.org/10.1109/cdc.2013.6760626).
- [17] G. Manganini, M. Pirotta, M. Restelli, L. Piroddi, and M. Prandini. “Policy search for the optimal control of Markov Decision Processes: A novel particle-based iterative scheme”. In: *IEEE transactions on cybernetics* 46.11 (2015), pp. 2643–2655. DOI: [10.1109/tcyb.2015.2483780](https://doi.org/10.1109/tcyb.2015.2483780).
- [18] S. Sun and C. C. de Visser. “Quadrotor safe flight envelope prediction in the high-speed regime: A Monte-Carlo approach”. In: *AIAA Scitech 2019 Forum*. 2019, p. 0948. DOI: [10.2514/6.2019-0948](https://doi.org/10.2514/6.2019-0948).
- [19] J. D. Gleason, A. P. Vinod, and M. M. Oishi. “Underapproximation of reach-avoid sets for discrete-time stochastic systems via lagrangian methods”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 4283–4290. DOI: [10.1109/cdc.2017.8264291](https://doi.org/10.1109/cdc.2017.8264291).
- [20] S. Singh, M. Chen, S. L. Herbert, C. J. Tomlin, and M. Pavone. “Robust tracking with model mismatch for fast and safe planning: an SOS optimization approach”. In: *Algorithmic Foundations of Robotics XIII: Proceedings of the 13th Workshop on the Algorithmic Foundations of Robotics 13*. Springer. 2020, pp. 545–564. DOI: [10.1007/978-3-030-44051-0_32](https://doi.org/10.1007/978-3-030-44051-0_32).
- [21] M. Althoff. “Reachability analysis and its application to the safety assessment of autonomous cars”. PhD thesis. Technische Universität München, 2010.
- [22] C. Le Guernic. “Reachability analysis of hybrid systems with linear continuous dynamics”. PhD thesis. Université Joseph-Fourier-Grenoble I, 2009.

- [23] J. Michaux, P. Holmes, B. Zhang, C. Chen, B. Wang, S. Sahgal, T. Zhang, S. Dey, S. Kousik, and R. Vasudevan. “Can’t Touch This: Real-Time, Safe Motion Planning and Control for Manipulators Under Uncertainty”. In: *arXiv preprint arXiv:2301.13308* (2023). DOI: [10.1109/tro.2025.3584557](https://doi.org/10.1109/tro.2025.3584557).
- [24] S. Kousik. “Reachability-based trajectory design”. PhD thesis. University of Michigan, 2020.
- [25] J. N. Maidens, S. Kaynama, I. M. Mitchell, M. M. Oishi, and G. A. Dumont. “Lagrangian methods for approximating the viability kernel in high-dimensional systems”. In: *Automatica* 49.7 (2013), pp. 2017–2029. DOI: [10.1016/j.automatica.2013.03.020](https://doi.org/10.1016/j.automatica.2013.03.020).
- [26] E. Van Oort, Q. Chu, and J. Mulder. “Maneuver envelope determination through reachability analysis”. In: *Advances in Aerospace Guidance, Navigation and Control: Selected Papers of the 1st CEAS Specialist Conference on Guidance, Navigation and Control*. Springer. 2011, pp. 91–102. DOI: [10.1007/978-3-642-19817-5_8](https://doi.org/10.1007/978-3-642-19817-5_8).
- [27] T. Lombaerts, S. Schuet, K. Wheeler, D. M. Acosta, and J. Kaneshige. “Safe maneuvering envelope estimation based on a physical approach”. In: *Aiaa guidance, navigation, and control (gnc) conference*. 2013, p. 4618. DOI: [10.2514/6.2013-4618](https://doi.org/10.2514/6.2013-4618).
- [28] S. Schuet, T. Lombaerts, D. Acosta, K. Wheeler, and J. Kaneshige. “An adaptive nonlinear aircraft maneuvering envelope estimation approach for online applications”. In: *Aiaa guidance, navigation, and control conference*. 2014, p. 0268. DOI: [10.2514/6.2014-0268](https://doi.org/10.2514/6.2014-0268).
- [29] Y. Zhang, C. C. de Visser, and Q. P. Chu. “Online safe flight envelope prediction for damaged aircraft: a database-driven approach”. In: *AIAA Modeling and Simulation Technologies Conference*. 2016, p. 1189. DOI: [10.2514/6.2016-1189](https://doi.org/10.2514/6.2016-1189).
- [30] M. Yin, Q. P. Chu, Y. Zhang, M. A. Niestroy, and C. C. de Visser. “Probabilistic flight envelope estimation with application to unstable overactuated aircraft”. In: *Journal of Guidance, Control, and Dynamics* 42.12 (2019), pp. 2650–2663. DOI: [10.2514/6.2019-0951](https://doi.org/10.2514/6.2019-0951).
- [31] J. Stapel, C. C. de Visser, E.-J. Van Kampen, and Q. P. Chu. “Efficient methods for flight envelope estimation through reachability analysis”. In: *AIAA guidance, navigation, and control conference*. 2016, p. 0083. DOI: [10.2514/6.2016-0083](https://doi.org/10.2514/6.2016-0083).
- [32] M. Chen, S. Herbert, and C. J. Tomlin. “Exact and efficient Hamilton-Jacobi guaranteed safety analysis via system decomposition”. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2017, pp. 87–92. DOI: [10.1109/icra.2017.7989015](https://doi.org/10.1109/icra.2017.7989015).
- [33] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin. “Decomposition of reachable sets and tubes for a class of nonlinear systems”. In: *IEEE Transactions on Automatic Control* 63.11 (2018), pp. 3675–3688. DOI: [10.1109/tac.2018.2797194](https://doi.org/10.1109/tac.2018.2797194).

- [34] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin. “Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking”. In: *IEEE Transactions on Automatic Control* 66.12 (2021), pp. 5861–5876.
- [35] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert. “Robust control barrier–value functions for safety-critical control”. In: *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE. 2021, pp. 6814–6821. DOI: [10.1109/cdc45484.2021.9683085](https://doi.org/10.1109/cdc45484.2021.9683085).
- [36] J. A. Sethian. *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*. Vol. 3. Cambridge university press, 1999.
- [37] S. Osher and R. P. Fedkiw. *Level set methods and dynamic implicit surfaces*. Vol. 1. Springer New York, 2005.
- [38] I. M. Mitchell *et al.* “A toolbox of level set methods”. In: *UBC Department of Computer Science Technical Report TR-2007-11* 1 (2007), p. 6.

5

ROBUST AVOID BACKWARD REACHABLE TUBE

Ensuring safety in complex systems with uncertainty is a critical challenge in engineering. Reachability analysis, which computes sets of states satisfying safety criteria, is a key tool for verification. Traditional approaches use Hamilton-Jacobi-Bellman (HJB) or Hamilton-Jacobi-Isaacs (HJI) equations but assume deterministic models, limiting their applicability under unmodeled dynamics and uncertainties. In this paper we introduce a three-player game framework, extending the HJI equation to include unmodeled dynamics as a third adversarial player. We prove that the value function derived from this formulation is the viscosity solution to the extended HJI equation. Using this foundation, we define the robust avoid backward reachable tube, capturing initial states that will reach an avoid set under worst-case disturbances and all realisations of the system dynamics. We establish that the zero sub-level set of the value function characterises the robust avoid backward reachable tube. A numerical example demonstrates the framework's integration with existing level set-based toolboxes, validating its effectiveness in addressing safety verification in circumstances of uncertain system dynamics.

5.1. INTRODUCTION

As systems design grows increasingly complex, ensuring their accuracy and safety has become a significant challenge. Consequently, verification and validation have emerged as critical areas of study across various engineering disciplines. The simplest form of computational validation is simulation, which, although widely used, is inherently limited. Simulations are non-exhaustive by nature, particularly for continuous-state systems. To address these limitations, researchers have turned to the computation of reachable sets, which capture the behaviour of entire families of trajectories simultaneously. Due to their importance, questions of reachability have been extensively studied in the dynamics and control literature [1], [2], with applications spanning from engineering [3] to economics [4].

For instance, reachability analysis plays a pivotal role in addressing safety problems in domains such as ground transportation systems [5], [6], air traffic management systems [7] [8], and flight control [9–11]. Direct characterisation of reachability concepts has been a key topic in viability theory [12], [13], while hybrid systems literature proposes methods for exact and approximate computation of reachable sets for certain classes of continuous dynamics [14], [15]. These contributions have led to the development of computational tools for reachability analysis, enabling their application across diverse system classes [16–18].

Reachability problems can also be formulated as optimal control problems, where reachable, viable, or invariant sets are characterised as level sets of the value function. By leveraging dynamic programming, the value function is expressed as the solution to a Hamilton-Jacobi (HJ) partial differential equation (PDE) [1, 10, 19–21]. For instance, reachability problems can be encoded as optimal control problems where the cost function is the minimum value of a state-dependent function over a given time horizon. A notable example is the avoid-reachable set (or avoid-reachable tube), where the controller aims to maximise this minimum value [10, 22–24]. This framework is further extended to account for disturbances by formulating the problem as a two-player game, where the disturbance seeks to minimise the minimum value while the controller seeks to maximise it [18]. Recent advancements have also explored solving the HJ PDE for a specific class of systems [20], [25], and using deep learning methods to approximate the PDE solutions [2].

Despite these developments, most dynamic programming-based approaches assume that the system model is deterministic and fully known. Although there are extensions for multi-modal systems [26], these approaches are limited to discrete mode switches and do not address continuous uncertainties or the impact of disturbances over a given horizon. Furthermore, reachability analysis has historically given limited attention to uncertain system dynamics, and to the best of our knowledge, such uncertainty has not been previously considered together with bounded disturbances. This gap restricts the applicability of HJ reachability analysis in real-world scenarios, where unmodeled dynamics and uncertainties can significantly influence system behaviour.

To address these limitations, this work builds on the foundational contributions of [10, 27, 28] by formulating a three-player game framework to extend the classical Hamilton-Jacobi-Isaacs (HJI) PDE. The author of [27] rigorously established the connection between differential games and the HJI equation, providing a representation of the value function as a viscosity solution and enabling the formulation of optimal control problems under adversarial dynamics [27]. Building on this, we introduce a novel extension to the two player game by introducing a third adversarial player which represents the unmodeled which exacerbate the worst-case scenarios. That is, the adversarial players work 'collaboratively' against the controller. Thus, the proposed three-player game is where the controller seeks to maximise safety, the disturbance acts to minimise it, and a third player representing unmodeled dynamics seeks to further minimise safety by exacerbating worst-case scenarios.

In concrete terms, we first integrate this three-player game formulation into the value function and prove that this value function is the viscosity solution of a certain HJI equation. We then provide the definition for a robust avoid backward reachable tube, building upon the avoid backward reachable tube framework developed in [10]. Next, we establish the connection between the value function and the robust avoid backward reachable tube, demonstrating that the level set of the value function is the robust avoid backward reachable tube. By introducing unmodeled dynamics, we derive a robust characterisation of backward reachable tubes, enabling safety verification under both adversarial disturbances and dynamic uncertainties. This extension incorporates system uncertainties directly into the Hamiltonian, allowing robust solutions that account for real-world unpredictability.

Finally, we provide a numerical example showing that this extension can be easily incorporated with existing toolboxes such as [29], and those based on level set methods [30], [31]. Numerical simulations demonstrate the efficacy of this approach in ensuring system safety in uncertain systems.

The remainder of the paper is structured as follows: Section 5.2 develops the HJI PDE extension with the three-player game value function. Section 5.3 defines the robust avoid backward reachable tube and establishes its relationship with the value function and the corresponding HJI PDE. Section 5.4 provides a numerical example of the developed method using existing toolboxes, and section 5.5 concludes the paper.

5.2. ROBUST DYNAMIC PROGRAMMING FOR CONTROL SYSTEMS WITH UNCERTAINTY

In this section, we extend the dynamic programming framework to optimal control problems with uncertain dynamics by formulating a three-player game. This formulation builds on the classical Hamilton-Jacobi-Isaacs (HJI) equations by incorporating an additional player to account for unmodeled dynamics in the system. We are particularly interested in the robust optimisation of control systems in adversarial and uncertain systems. The primary focus is to characterise

the value function for this class of problems and demonstrate that it satisfies a Hamilton-Jacobi-Bellman-Isaacs (HJBI) PDE as a viscosity solution, establishing a sound mathematical foundation for robust control under uncertain dynamics.

Let us fix $T > t \geq 0$, $x \in \mathbb{R}^m$ and consider the ordinary differential equation (ODE) given by equation 5.1

$$\begin{aligned} \dot{x}(s) &= f(s, x(s), \alpha(s), u(s), d(s)) & t \leq s \leq T \\ x(t) &= x, & u : [t, T] \rightarrow \mathcal{U}, & d : [t, T] \rightarrow \mathcal{D}, \\ & & \alpha : [t, T] \rightarrow \mathcal{A} \end{aligned} \quad (5.1)$$

α, u and d are measurable functions, $\mathcal{U} \subset \mathbb{R}^k$, $\mathcal{D} \subset \mathbb{R}^l$ are compact sets, and $\mathcal{A}$ is a compact parametric set, where x is system state, t, s represents time, u is the input policy to the system, d is the disturbance policy, and α is the system parameters. We assume that:

$$f : [0, T] \times \mathbb{R}^m \times \mathcal{A} \times \mathcal{U} \times \mathcal{D} \rightarrow \mathbb{R}^m \quad (5.2)$$

We assume that f is uniformly continuous with the properties 5.3 and 5.4 for some constant C_1 and all $0 \leq t \leq T$

$$|f(t, x, \alpha, u, d)| \leq C_1 \quad (5.3)$$

$$|f(t, x, \alpha, u, d) - f(t, \hat{x}, \alpha, u, d)| \leq C_1 |x - \hat{x}| \quad (5.4)$$

and $x, \hat{x} \in \mathbb{R}^m$, $\alpha \in \mathcal{A}$, $u \in \mathcal{U}$, $d \in \mathcal{D}$. The (unique) solution for each $\alpha \in \mathcal{A}$ of the ODE is the response of the system to the control $u(\cdot)$, $d(\cdot)$, and is represented by $x(\cdot)$. We define an associated payoff/cost function to the ODE:

$$P(\alpha, u, d) \equiv P_{x,t}(\alpha(\cdot), u(\cdot), d(\cdot)) \equiv \int_t^T h(s, x(s), \alpha(s), u(s), d(s)) ds + g(x(T)) \quad (5.5)$$

where $g : \mathbb{R}^m \rightarrow \mathbb{R}$ is uniformly continuous and satisfies

$$|g(x)| \leq C_2 \quad (5.6)$$

$$|g(x) - g(\hat{x})| \leq C_2 |x - \hat{x}| \quad (5.7)$$

and $h : [0, T] \times \mathbb{R}^m \times \mathcal{A} \times \mathcal{U} \times \mathcal{D} \rightarrow \{0, \mathbb{R}^+\}$ is uniformly continuous with

$$|h(t, x, \alpha, u, d)| \leq C_3 \quad (5.8)$$

$$|h(t, x, \alpha, u, d) - h(t, \hat{x}, \alpha, u, d)| \leq C_3 |x - \hat{x}| \quad (5.9)$$

and we also assume that, $h(\cdot)$ is time invariant.

for constants C_2, C_3 and all $0 \leq t \leq T, x, \hat{x} \in \mathbb{R}^m, \alpha \in \mathcal{A}, u \in \mathcal{U}, d \in \mathcal{D}$ and $\alpha \in \mathcal{A}$. The goal of α and d is to minimise $P(\alpha, u, d)$, whereas the goal of u is to maximise the $P(\alpha, u, d)$.

The sets $\mathcal{K}(t), \mathcal{M}(t)$, and $\mathcal{N}(t)$ are the sets of all parameters, control and disturbance policies respectively defined as

$$\begin{aligned} \mathcal{K}(t) &\equiv \{\alpha : [t, T] \rightarrow \mathcal{A} | \alpha \text{ measurable}\} \\ \mathcal{M}(t) &\equiv \{u : [t, T] \rightarrow \mathcal{U} | u \text{ measurable}\} \\ \mathcal{N}(t) &\equiv \{d : [t, T] \rightarrow \mathcal{U} | d \text{ measurable}\} \end{aligned}$$

Now we will identify the input policy (u), disturbance (d), and parameters (α) that agree a.e. (almost everywhere). Following [32], [27], and [33], a mapping is defined as

$$\beta : \mathcal{M}(t) \rightarrow \mathcal{N}(t)$$

which is the strategy of the disturbance (beginning at time t) provided for each $t \leq s \leq T, u, \hat{u} \in \mathcal{M}(t)$

$$u(\tau) = \hat{u}(\tau) \text{ for a.e. } t \leq \tau \leq s$$

$$\Rightarrow \beta[u](\tau) = \beta[\hat{u}](\tau) \text{ for a.e. } t \leq \tau \leq s$$

Let $\Delta(t)$ be the set of all such disturbance strategies beginning at time t . Similarly we define a mapping

$$\bar{\alpha} : \mathcal{M}(t) \times \mathcal{N}(t) \rightarrow \mathcal{K}(t)$$

which is the system parameter provided for each $t \leq s \leq T, u, \hat{u} \in \mathcal{M}(t), d, \hat{d} \in \mathcal{N}(t)$

$$u(\tau) = \hat{u}(\tau) \text{ for a.e. } t \leq \tau \leq s$$

$$d(\tau) = \hat{d}(\tau) \text{ for a.e. } t \leq \tau \leq s$$

$$\Rightarrow \bar{\alpha}[u, d](\tau) = \bar{\alpha}[\hat{u}, \hat{d}](\tau) \text{ for a.e. } t \leq \tau \leq s$$

We will for the sake of brevity use $\bar{\alpha}$ instead of $\bar{\alpha}[u, d]$. Let $\bar{\mathcal{A}}(t)$ be the set of all such system parameters beginning at time t . Now finally we define the value function as equation 5.10.

$$V(t, x) = \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \inf_{\tau \in [t, T]} \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \quad (5.10)$$

where $x(\cdot)$ solves the ODE in equation 5.1 for $u(\cdot)$, $\beta[u](\cdot)$ and $\bar{\alpha}$. Now in the following theorem we want to show that the value function defined in equation 5.10 has Bellman optimality principle [27]. This is necessary because we will later prove that the value function given in equation 5.10 is the viscosity solution of a certain HJI PDE.

Theorem 7 *The value function defined in equation 5.10 has Bellman optimality principle, i.e.,*

$$V(t, x) = \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \min \left\{ V((t + \sigma), x(t + \sigma)) + \int_t^{t+\sigma} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds, \right. \\ \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \right\}$$

Proof: The proof is based on a similar logic as the proof of the Bellman condition in [27].

$$\text{Let } W(t, x) = \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \min \left\{ V((t + \sigma), x(t + \sigma)) + \int_t^{t+\sigma} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds, \right. \\ \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \right\}$$

Then for a fixed $\epsilon > 0$, according to the definition of infimum, for all $u \in \mathcal{M}(t)$ there exists $\delta \in \Delta(t)$ and $\alpha^* \in \bar{\mathcal{A}}(t)$ such that:

$$W(t, x) \geq \min \{ V((t + \sigma), x(t + \sigma)) \\ + \int_t^{t+\sigma} h(s, x(s), \alpha^*(s), u(s), \delta[u](s)) ds, \\ \inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau h(s, x(s), \alpha^*(s), u(s), \delta[u](s)) ds + g(x(\tau)) \right] \} - 2\epsilon \quad (5.11)$$

From equation 5.10 we have:

$$V(t + \sigma, x(t + \sigma)) = \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t + \sigma)} \sup_{u \in \mathcal{M}(t + \sigma)} \inf_{\tau \in [t + \sigma, T]} \left[\int_{t + \sigma}^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right]$$

Based on the definition of infimum, for a fixed $\epsilon > 0$, for all $u \in \mathcal{M}(t + \sigma)$ there exists $\bar{\delta} \in \Delta(t + \sigma)$ and $\alpha^{**} \in \bar{\mathcal{A}}(t + \sigma)$, such that

$$V(t + \sigma, x(t + \sigma)) \geq \inf_{\tau \in [t + \sigma, T]} \left[\int_{t + \sigma}^{\tau} h(s, x(s), \alpha^{**}(t), u(s), \bar{\delta}[u](s)) ds + g(x(\tau)) \right] - 2\epsilon$$

$$\text{Let } \beta[u](s) \equiv \begin{cases} \delta(s) & \text{if } t \leq s < t + \sigma \\ \bar{\delta}(s) & \text{if } t + \sigma \leq s \leq T \end{cases}$$

$$\text{Let } \bar{\alpha}(s) = \bar{\alpha}[u, \beta](s) \equiv \begin{cases} \alpha^{*} & \text{if } t \leq s < t + \sigma \\ \alpha^{**} & \text{if } t + \sigma \leq s \leq T \end{cases}$$

then by substituting the above in equation 5.11 we get

$$\begin{aligned} W(t, x) &\geq \min \left\{ \inf_{\tau \in [t + \sigma, T]} \left[\int_{t + \sigma}^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \right. \\ &\quad \left. + \int_t^{t + \sigma} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds, \right. \\ &\quad \left. \inf_{\tau \in [t, t + \sigma]} \left[\int_t^{\tau} (h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(s))) \right] \right\} - 4\epsilon \end{aligned}$$

$$\begin{aligned} \Rightarrow W(t, x) &\geq \min \left\{ \inf_{\tau \in [t + \sigma, T]} \left[\int_t^{\tau} (h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau))) \right], \right. \\ &\quad \left. \inf_{\tau \in [t, t + \sigma]} \left[\int_t^{\tau} (h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \right. \right. \\ &\quad \left. \left. + g(x(\tau))) \right] \right\} - 4\epsilon \end{aligned}$$

$$\Rightarrow W(t, x) \geq \inf_{\tau \in [t, T]} \left[\int_t^\tau (h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s))) ds + g(x(\tau)) \right] - 4\epsilon$$

This holds true for all $u \in \mathcal{M}(t)$ thus:

$$W(t, x) \geq \sup_{u \in \mathcal{M}(t)} \inf_{\tau \in [t, T]} \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] - 4\epsilon$$

$$\Rightarrow W(t, x) \geq \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \inf_{\tau \in [t, T]} \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] - 4\epsilon$$

$$\Rightarrow W(t, x) \geq V(t, x) - 4\epsilon \quad (5.12)$$

Furthermore, for all $\bar{\alpha} \in \bar{\mathcal{A}}(t)$ and $\beta \in \Delta(t)$

$$W(t, x) \leq \sup_{u \in \mathcal{M}(t)} \min \{ V((t + \sigma), x(t + \sigma)) + \int_t^{t+\sigma} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds, \inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \}$$

Then for a fixed $\epsilon > 0$, according to the definition of the supremum, there exists a $u_1 \in \mathcal{M}(t)$, such that

$$W(t, x) \leq \min \{ V((t + \sigma), x(t + \sigma)) + \int_t^{t+\sigma} h(s, x(s), \bar{\alpha}(s), u_1(s), \beta[u_1](s)) ds, \inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u_1(s), \beta[u_1](s)) ds + g(x(\tau)) \right] \} + \epsilon$$

Now for each $u \in \mathcal{M}(t + \sigma)$ we define $\tilde{u} \in \mathcal{M}(t)$ by

$$\tilde{u}(s) \equiv \begin{cases} u_1(s) & \text{if } t \leq s < t + \sigma \\ u(s) & \text{if } t + \sigma \leq s \leq T \end{cases}$$

and we also define $\beta \in \Delta(t + \sigma)$ and $\alpha \in \bar{\mathcal{A}}(t + \sigma)$ by

$$\beta[u](s) \equiv \beta[\tilde{u}](s) \quad (t + \sigma \leq s \leq T)$$

$$\alpha(s) \equiv \bar{\alpha}[\tilde{u}, \beta](s) \quad (t + \sigma \leq s \leq T)$$

$$\begin{aligned} V(t + \sigma, x(t + \sigma)) &= \inf_{\alpha \in \bar{\mathcal{A}}(t + \sigma)} \inf_{\beta \in \Delta(t + \sigma)} \sup_{u \in \mathcal{M}(t + \sigma)} \inf_{\tau \in [t + \sigma, T]} \\ &\quad \left[\int_{t + \sigma}^{\tau} h(s, x(s), \alpha(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \end{aligned}$$

Then for all $\alpha \in \bar{\mathcal{A}}(t + \sigma)$ and $\beta \in \Delta(t + \sigma)$ we have

$$\begin{aligned} V(t + \sigma, x(t + \sigma)) &\leq \sup_{u \in \mathcal{M}(t + \sigma)} \inf_{\tau \in [t + \sigma, T]} \\ &\quad \left[\int_{t + \sigma}^{\tau} h(s, x(s), \alpha(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \end{aligned}$$

5

Then for a fixed $\epsilon > 0$, according to the definition of supremum, $\exists u_2 \in \mathcal{M}(t + \sigma)$, such that

$$\begin{aligned} V(t + \sigma, x(t + \sigma)) &\leq \inf_{\tau \in [t + \sigma, T]} \\ &\quad \left[\int_{t + \sigma}^{\tau} h(s, x(s), \alpha(s), u_2(s), \beta[u_2](s)) ds + g(x(\tau)) \right] + \epsilon \end{aligned}$$

$$u(s) \equiv \begin{cases} u_1(s) & \text{if } t \leq s < t + \sigma \\ u_2(s) & \text{if } t + \sigma \leq s \leq T \end{cases}$$

Thus we get

$$\begin{aligned} W(t, x) &\leq \min \left\{ \inf_{\tau \in [t + \sigma, T]} \right. \\ &\quad \left[\int_{t + \sigma}^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \\ &\quad + \int_t^{t + \sigma} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds, \\ &\quad \left. \inf_{\tau \in [t, t + \sigma]} \left[\int_t^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \right\} + 2\epsilon \end{aligned}$$

Since $\inf_x (f(x)) + g(y) = \inf_x (f(x) + g(y))$

$$\begin{aligned} \Rightarrow W(t, x) &\leq \min \left\{ \inf_{\tau \in [t+\sigma, T]} \right. \\ &\quad \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right], \\ &\quad \inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \right. \\ &\quad \left. \left. + g(x(\tau)) \right] \right\} + 2\epsilon \end{aligned}$$

$$\begin{aligned} \Rightarrow W(t, x) &\leq \inf_{\tau \in [t, T]} \\ &\quad \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] + 2\epsilon \quad (5.13) \end{aligned}$$

Now from equation 5.10

$$\begin{aligned} V(t, x) &= \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \inf_{\tau \in [t, T]} \\ &\quad \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \end{aligned}$$

Thus there exists $\bar{\alpha} \in \bar{\mathcal{A}}(t)$ and $\beta \in \Delta(t)$, such that for a given $\epsilon > 0$:

$$\begin{aligned} V(t, x) &\geq \sup_{u \in \mathcal{M}(t)} \inf_{\tau \in [t, T]} \\ &\quad \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] - 2\epsilon \end{aligned}$$

Thus for all $u \in \mathcal{M}(t)$ we get

$$\begin{aligned} V(t, x) + 2\epsilon &\geq \inf_{\tau \in [t, T]} \\ &\quad \left[\int_t^\tau h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \quad (5.14) \end{aligned}$$

Therefore, combining equations 5.13 and 5.14 yields:

$$\Rightarrow W(t, x) \leq V(t, x) + 4\epsilon \quad (5.15)$$

Finally, from equations 5.12 and 5.15

$$V(t, x) = \inf_{\bar{\alpha} \in \bar{A}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \min \left\{ V((t + \sigma), x(t + \sigma)) + \int_t^{t+\sigma} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds, \right. \\ \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \right\}$$

Now we will prove two Lemma that shows the monotonicity of the value function in both state and time. These Lemmas are required to prove the final result that the value function is the viscosity solution of a certain HJI PDE.

Lemma 8 For all $(t, x) \in [0, T] \times \mathbb{R}^m$

5

$$V(t, x) \leq V(t + \delta, x) + \inf_{\bar{\alpha} \in \bar{A}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \int_t^{t+\delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds$$

and $V(T, x) = g(x(T))$

Proof: The $V(T, x) = g(x(T))$ is trivial and can be directly observed from the definition of the value function (equation 5.10). Now

$$V(t, x) = \inf_{\bar{\alpha} \in \bar{A}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \inf_{\tau \in [t, T]} \left[\int_t^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right]$$

Thus let us assume for the sake of contradiction that

$$V(t, x) > V(t + \delta, x) + \inf_{\bar{\alpha} \in \bar{A}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \int_t^{t+\delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds$$

$$\begin{aligned}
&\Rightarrow \inf_{\bar{\alpha}_1 \in \bar{\mathcal{A}}(t)} \inf_{\beta_1 \in \Delta(t)} \sup_{u_1 \in \mathcal{M}(t)} \inf_{\tau \in [t, T]} \\
&\quad \left[\int_t^\tau h(s, x_1(s), \bar{\alpha}_1(s), u_1(s), \beta_1[u](s)) ds + g(x_1(\tau)) \right] \\
&\quad > \inf_{\bar{\alpha}_2 \in \bar{\mathcal{A}}(t+\delta)} \inf_{\beta_2 \in \Delta(t+\delta)} \sup_{u_2 \in \mathcal{M}(t+\delta)} \inf_{\tau \in [t+\delta, T]} \\
&\quad \left[\int_{t+\delta}^\tau h(s, x_2(s), \bar{\alpha}_2(s), u_2(s), \beta_2[u](s)) ds + g(x_2(\tau)) \right] \\
&\quad + \inf_{\bar{\alpha}_3 \in \bar{\mathcal{A}}(t)} \inf_{\beta_3 \in \Delta(t)} \sup_{u_3 \in \mathcal{M}(t)} \\
&\quad \left[\int_t^{t+\delta} h(s, x_3(s), \bar{\alpha}_3(s), u_3(s), \beta_3[u](s)) ds \right]
\end{aligned}$$

Then for all $\bar{\alpha}_1 \in \bar{\mathcal{A}}(t)$, there exists a pair $(\bar{\alpha}_2 \in \bar{\mathcal{A}}(t+\delta), \bar{\alpha}_3 \in \bar{\mathcal{A}}(t))$ and for all $\beta_1 \in \Delta(t)$ there exists a pair $(\beta_2 \in \Delta(t+\delta), \beta_3 \in \Delta(t))$, such that

$$\begin{aligned}
&\Rightarrow \sup_{u_1 \in \mathcal{M}(t)} \inf_{\tau \in [t, T]} \\
&\quad \left[\int_t^\tau h(s, x_1(s), \bar{\alpha}_1(s), u_1(s), \beta_1[u](s)) ds + g(x_1(\tau)) \right] \\
&\quad > \sup_{u_2 \in \mathcal{M}(t+\delta)} \inf_{\tau \in [t+\delta, T]} \\
&\quad \left[\int_{t+\delta}^\tau h(s, x_2(s), \bar{\alpha}_2(s), u_2(s), \beta_2[u](s)) ds + g(x_2(\tau)) \right] \\
&\quad + \sup_{u_3 \in \mathcal{M}(t)} \\
&\quad \left[\int_t^{t+\delta} h(s, x_3(s), \bar{\alpha}_3(s), u_3(s), \beta_3[u](s)) ds \right]
\end{aligned}$$

Furthermore, for all $(u_2 \in \mathcal{M}(t+\delta), u_3 \in \mathcal{M}(t))$ there exists $u_1 \in \mathcal{M}(t)$, such that

$$\begin{aligned}
&\inf_{\tau \in [t, T]} \\
&\quad \left[\int_t^\tau h(s, x_1(s), \bar{\alpha}_1(s), u_1(s), \beta_1[u](s)) ds + g(x_1(\tau)) \right] \\
&\quad > \inf_{\tau \in [t+\delta, T]} \\
&\quad \left[\int_{t+\delta}^\tau h(s, x_2(s), \bar{\alpha}_2(s), u_2(s), \beta_2[u](s)) ds + g(x_2(\tau)) \right] \\
&\quad + \int_t^{t+\delta} h(s, x_3(s), \bar{\alpha}_3(s), u_3(s), \beta_3[u](s)) ds
\end{aligned}$$

Now let us define:

$$u(s) \equiv \begin{cases} u_1(s) & \text{if } t \leq s < t + \delta \\ u_1(s) & \text{if } t + \delta \leq s \leq T \end{cases}$$

$$\beta[u](s) \equiv \begin{cases} \beta_3[u](s) & \text{if } t \leq s < t + \delta \\ \beta_2[u](s) & \text{if } t + \delta \leq s \leq T \end{cases}$$

$$\bar{\alpha}(s) \equiv \begin{cases} \bar{\alpha}_3(s) & \text{if } t \leq s < t + \delta \\ \bar{\alpha}_2(s) & \text{if } t + \delta \leq s \leq T \end{cases}$$

$$\begin{aligned} \Rightarrow \inf_{\tau \in [t, T]} & \left[\int_t^\tau h(s, x_1(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x_1(\tau)) \right] \\ & > \inf_{\tau \in [t+\delta, T]} \left[\int_{t+\delta}^\tau h(s, x_2(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x_2(\tau)) \right] \\ & \quad + \int_t^{t+\delta} h(s, x_3(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \end{aligned}$$

5

We also know that

$$\begin{cases} \frac{dx_1}{ds} = f(s, x_1(s), \alpha(s), u(s), \beta[u](s)) \\ x_1(t) = x \end{cases}$$

$$\begin{cases} \frac{dx_2}{ds} = f(s, x_2(s), \alpha(s), u(s), \beta[u](s)) \\ x_2(t + \delta) = x \end{cases}$$

Since $h(\cdot)$ is a non-negative function, and using the uniqueness of the solution of equation 5.1 we get

$$\begin{aligned} \Rightarrow \inf_{\tau \in [t, T]} & \left[\int_t^\tau h(s, x_1(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x_1(\tau)) \right] \\ & > \inf_{\tau \in [t+\delta, T]} \left[\int_t^\tau h(s, x_1(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x_1(\tau)) \right] \end{aligned}$$

This is a contradiction.

Lemma 9 For all $(t, x) \in [0, T] \times \mathbb{R}^m$

$$\begin{aligned}
 V(t, x) \leq & V(t + \delta, x(t + \delta)) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \\
 & \int_t^{t+\delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \leq \\
 & \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \\
 & \left(\int_t^T h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right)
 \end{aligned}$$

and $V(T, x) = g(x(T))$

Proof: The $V(T, x) = g(x(T))$ is trivial and can be directly observed from the definition of the value function (equation 5.10). Furthermore $V(t, x) \leq V(t + \delta, x(t + \delta)) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \int_t^{t+\delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds$ can be observed directly from the Theorem 7. Now we will prove the final part : Let us assume for the contradiction that

$$\begin{aligned}
 V(t + \delta, x(t + \delta)) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \\
 \int_t^{t+\delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds > \\
 \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \\
 \left(\int_t^T h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right)
 \end{aligned}$$

Then for all $\bar{\alpha}_1 \in \bar{\mathcal{A}}(t)$, there exists a pair $(\bar{\alpha}_2 \in \bar{\mathcal{A}}(t + \delta), \bar{\alpha}_3 \in \bar{\mathcal{A}}(t))$, for all $\beta_1 \in \Delta(t)$ there exists a pair $(\beta_2 \in \Delta(t + \delta), \beta_3 \in \Delta(t))$ and for all $(u_2 \in \mathcal{M}(t + \delta), u_3 \in \mathcal{M}(t))$ there exists $u_1 \in \mathcal{M}(t)$, such that:

$$\begin{aligned}
 & \int_t^T h(s, x(s), \bar{\alpha}_1(s), u_1(s), \beta_1[u](s)) ds + g(x(T)) \\
 & < \inf_{\tau \in [t+\delta, T]} \\
 & \left[\int_{t+\delta}^{\tau} h(s, x(s), \bar{\alpha}_2(s), u_2(s), \beta_2[u](s)) ds + g(x(\tau)) \right] \\
 & + \int_t^{t+\delta} h(s, x(s), \bar{\alpha}_3(s), u_3(s), \beta_3[u](s)) ds
 \end{aligned}$$

Now let us define:

$$u(s) \equiv \begin{cases} u_1(s) & \text{if } t \leq s < t + \delta \\ u_1(s) & \text{if } t + \delta \leq s \leq T \end{cases}$$

$$\beta[u](s) \equiv \begin{cases} \beta_3[u](s) & \text{if } t \leq s < t + \delta \\ \beta_2[u](s) & \text{if } t + \delta \leq s \leq T \end{cases}$$

$$\bar{\alpha}(s) \equiv \begin{cases} \bar{\alpha}_3(s) & \text{if } t \leq s < t + \delta \\ \bar{\alpha}_2(s) & \text{if } t + \delta \leq s \leq T \end{cases}$$

$$\begin{aligned} \Rightarrow \int_t^T h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(T)) \\ < \inf_{\tau \in [t+\delta, T]} \left[\int_t^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right] \end{aligned}$$

5

This is a contradiction as the infimum over a larger set is always less than or equal to any value in the set.

Now we will prove two lemmas that show that the value function given in equation 5.10 is bounded and Lipschitz continuous. These are standard results required to prove the existence and uniqueness of the solution [10] [28].

Lemma 10 *The value function given in equation 5.10 is bounded, i.e., $|V(t, x)| \leq C_4$, where C_4 is a constant.*

Proof:

$$P(\alpha, u, d) = \int_t^T h(s, x(s), \alpha(s), u(s), d(s)) ds + g(x(T))$$

Using equations 5.6 and 5.8 we can say that $|P(\alpha, u, d)| \leq TC_3 + C_2$. This holds for all $u \in \mathcal{U}$, $d \in \mathcal{D}$ and $\alpha \in \mathcal{A}$, thus implying $|V(t, x)| \leq C_4$.

Lemma 11 *The value function given in the equation 5.10 is Lipschitz continuous i.e. $|V(t, x) - V(\hat{t}, \hat{x})| \leq C_4(|t - \hat{t}| + |x - \hat{x}|)$ for all $0 \leq t \leq T$, $0 \leq \hat{t} \leq T$ and $x, \hat{x} \in \mathbb{R}^m$.*

Proof: The proof is based on the on the proof of theorem 3.2 of [27] and follows the same logic.

$$\text{Let } P(\tau, \bar{\alpha}, u, d) = \int_t^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), d(s)) ds + g(x(\tau))$$

Part I

We will show that $V(t_1, x_1) - V(t_2, x_2) \leq \bar{C}(|t_1 - t_2| + |x_1 - x_2|) + 4\epsilon$

$$V(t_1, x_1) = \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_1)} \inf_{\beta \in \Delta(t_1)} \sup_{u \in \mathcal{M}(t_1)} \inf_{\tau \in [t_1, T]} P(\tau, \bar{\alpha}, u, \beta[u])$$

Based on the definition of supremum we have for a fixed $\epsilon > 0$, $\exists u_0 \in \mathcal{M}(t_1)$, such that for all $\tau \in [t_1, T]$, for all $\bar{\alpha} \in \bar{\mathcal{A}}(t_1)$ and $\beta \in \Delta(t_1)$ we have

$$V(t_1, x_1) \leq P(\tau, \bar{\alpha}, u_0, \beta[u_0]) + \epsilon \quad (5.16)$$

We define for any $u \in \mathcal{M}(t_2)$, $\tilde{u} \in \mathcal{M}(t_1)$ as follows

$$\tilde{u}(s) \equiv \begin{cases} u_0(s) & \text{if } t_1 \leq t < t_2 \\ u(s) & \text{if } t_2 \leq t \leq T \end{cases}$$

Now we define for $\beta \in \Delta(t_2)$ by setting for each $u \in \mathcal{M}(t_2)$ as follows

$$\beta[u](s) \equiv \beta[\tilde{u}](s) \quad (t_2 \leq s \leq T)$$

We also define $\alpha \in \mathcal{A}(t_2)$ by setting for each $u \in \mathcal{M}(t_2)$ and $\beta \in \Delta(t_2)$ for $t_2 \leq s \leq T$

$$\alpha[u, \beta[u]](s) \equiv \bar{\alpha}[\tilde{u}, \beta[\tilde{u}]](s) \quad (t_2 \leq s \leq T)$$

Following the definition of infimum, $\exists \alpha \in \bar{\mathcal{A}}(t_2)$, $\beta \in \Delta(t_2)$ and $\tau \in [t_2, T]$ for all $u \in \mathcal{M}(t_2)$, such that

$$V(t_2, x_2) \geq P(\tau, \alpha, u, \beta[u]) - 3\epsilon \quad (5.17)$$

Then combining equations 5.16 and 5.17

$$V(t_1, x_1) - V(t_2, x_2) \leq P(\tau, \bar{\alpha}, \tilde{u}, \beta[\tilde{u}]) - P(\tau, \alpha, u, \beta[u]) + 4\epsilon \quad (5.18)$$

Now let $x_1(\cdot)$ solve the following system for the time horizon $(t_1 \leq s \leq T)$

$$\begin{cases} \frac{dx_1}{ds} = f(s, x_1(s), \bar{\alpha}(s), \tilde{u}(s), \beta[\tilde{u}](s)) \\ x_1(t_1) = x_1 \end{cases}$$

Let $x_2(\cdot)$ solve the following system for the time horizon $(t_2 \leq s \leq T)$

$$\begin{cases} \frac{dx_2}{ds} = f(s, x_2(s), \alpha(s), u(s), \beta[u](s)) \\ x_2(t_2) = x_2 \end{cases}$$

We can prove that

$$|x_1(t_2) - x_1| \leq C_1 |t_1 - t_2| \quad (5.19)$$

The proof of the statement above can be found in the appendix. Furthermore, since $u = \tilde{u}$, $\beta[u] = \beta[\tilde{u}]$ and $\underline{\alpha} = \bar{\alpha}$ for $t_2 \leq s \leq T$ we have

$$\begin{aligned} & |x_1(s) - x_2(s)| \\ &= \left| \int_{t_1}^s f(s, x_1(s), \bar{\alpha}(s), \tilde{u}(s), \beta[\tilde{u}](s)) ds - \int_{t_2}^s f(s, x_2(s), \underline{\alpha}(s), u(s), \beta[u](s)) ds \right| \\ &= \left| \int_{t_1}^{t_2} f(s, x_1(s), \bar{\alpha}(s), \tilde{u}(s), \beta[\tilde{u}](s)) ds \right. \\ &\quad \left. + \int_{t_2}^s (f(s, x_1(s), \bar{\alpha}(s), \tilde{u}(s), \beta[\tilde{u}](s)) \right. \\ &\quad \left. - f(s, x_2(s), \underline{\alpha}(s), u(s), \beta[u](s))) ds \right| \end{aligned}$$

Using equations 5.19, 5.4 we get

$$|x_1(s) - x_2(s)| \leq C(|t_1 - t_2| + |x_1 - x_2|) \text{ for } t_2 \leq s \leq T \quad (5.20)$$

where C is a constant. Now from equation 5.18 we get

$$V(t_1, x_1) - V(t_2, x_2) \leq P(\tau, \bar{\alpha}, \tilde{u}, \beta[\tilde{u}]) - P(\tau, \underline{\alpha}, u, \beta[u]) + 4\epsilon$$

This is valid for all $\tau \in (t, T)$, thus we get

$$\begin{aligned} & V(t_1, x_1) - V(t_2, x_2) \leq P(\tau, \bar{\alpha}, \tilde{u}, \beta[\tilde{u}]) - P(\tau, \underline{\alpha}, u, \beta[u]) + 4\epsilon \\ &= \int_{t_1}^{\tau} h(s, x_1(s), \bar{\alpha}(s), \tilde{u}(s), \beta[\tilde{u}](s)) ds + g(x_1(\tau)) \\ &\quad - \int_{t_2}^{\tau} h(s, x_2(s), \underline{\alpha}(s), u(s), \beta[u](s)) ds - g(x_2(\tau)) + 4\epsilon \\ &= \int_{t_1}^{t_2} h(s, x_1(s), \bar{\alpha}(s), \tilde{u}(s), \beta[\tilde{u}](s)) ds \\ &\quad + \int_{t_2}^{\tau} h(s, x_1(s), \bar{\alpha}(s), \tilde{u}(s), \beta[\tilde{u}](s)) ds + g(x_1(\tau)) \\ &\quad - \int_{t_2}^{\tau} h(s, x_2(s), \underline{\alpha}(s), u(s), \beta[u](s)) ds - g(x_2(\tau)) + 4\epsilon \end{aligned}$$

In the time interval of $t_2 \leq s \leq T$ we have that $\bar{\alpha} = \underline{\alpha}$, $\tilde{u} = u$ and $\beta[\tilde{u}] = \underline{\beta}[u]$, thus

$$\begin{aligned}
 &= \int_{t_1}^{t_2} h(s, x_1(s), \bar{\alpha}(s), \tilde{u}(s), \beta[\tilde{u}](s)) ds \\
 &\quad + \int_{t_2}^T (h(s, x_1(s), \bar{\alpha}(s), u(s), \beta[u](s)) \\
 &\quad - h(s, x_2(s), \bar{\alpha}(s), u(s), \beta[u](s))) ds \\
 &\quad + g(x_1(T)) - g(x_2(T)) + 4\epsilon
 \end{aligned}$$

Thus using equations 5.7, 5.8, 5.9, and 5.20 we get

$$V(t_1, x_1) - V(t_2, x_2) \leq \bar{C}(|t_1 - t_2| + |x_1 - x_2|) + 4\epsilon \quad (5.21)$$

where $\bar{C}$ is a constant.

5

Part II

Now we will show that $V(t_2, x_2) - V(t_1, x_1) \leq \bar{C}(|t_1 - t_2| + |x_1 - x_2|) + 4\epsilon$. Following the similar logic as part I we get that for all $\tau \in [t_2, T]$:

$$V(t_2, x_2) \leq P(\tau, \underline{\alpha}, \underline{u}, \beta[\underline{u}]) + \epsilon \quad (5.22)$$

Now for each $u \in \mathcal{M}(t_1)$ we defined $\underline{u} \in \mathcal{M}(t_2)$ by

$$\underline{u}(s) = u(s) \quad (t_2 \leq s \leq T)$$

Now for any fixed $\beta_0 \in \Delta(t_1)$ we define $\tilde{\beta} \in \Delta(t_1)$ by

$$\tilde{\beta} \equiv \begin{cases} \beta_0 & t_1 \leq s < t_2 \\ \beta[\tilde{u}] & t_2 \leq s \leq T \end{cases}$$

Also for any fixed $\alpha_0 \in \mathcal{A}(t_1)$ we define $\tilde{\alpha} \in \bar{\mathcal{A}}(t_1)$ by

$$\tilde{\alpha} \equiv \begin{cases} \alpha_0 & t_1 \leq s < t_2 \\ \underline{\alpha}[\underline{u}, \beta] & t_2 \leq s \leq T \end{cases}$$

Now we choose α_0 and β_0 such that all for $u \in \mathcal{M}(t_1)$

$$V(t_1, x_1) \geq \inf_{\tau \in [t_1, T]} P(\tau, \alpha_0, u, \beta_0[u]) - 2\epsilon$$

Then according to the definition of infimum there exists a $\tilde{\tau} \in [t_1, T]$ such that

$$V(t_1, x_1) \geq P(\tilde{\tau}, \tilde{\alpha}, u, \tilde{\beta}[u]) - 3\epsilon \quad (5.23)$$

Using equations 5.22 and 5.23 we get

$$V(t_2, x_2) - V(t_1, x_1) \leq P(\tau, \underline{\alpha}, \underline{u}, \beta[\underline{u}]) - P(\tilde{\tau}, \tilde{\alpha}, u, \tilde{\beta}[u]) + 4\epsilon \quad (5.24)$$

Now let $x_1(\cdot)$ solve the following from $t_1 \leq s < T$

$$\begin{cases} \frac{dx_1}{ds} = f(s, x_1(s), \tilde{\alpha}(s), u, \tilde{\beta}[u]) \\ x_1(t_1) = x_1 \end{cases}$$

And let $x_2(\cdot)$ solve the following from $t_2 \leq s < T$

$$\begin{cases} \frac{dx_2}{ds} = f(s, x_2(s), \underline{\alpha}(s), \underline{u}, \beta[\underline{u}]) \\ x_2(t_2) = x_2 \end{cases}$$

Using equation 5.19 And since $\tilde{\alpha} = \underline{\alpha}$ and $u = \tilde{u}$ and $\tilde{\beta}[u] = \beta[\underline{u}]$ for $t_2 \leq s \leq T$

$$|x_1(s) - x_2(s)| \leq C|x_1(t_2) - x_2| \leq C(|t_1 - t_2| + |x_1 - x_2|)$$

Now from equation 5.24

$$\begin{aligned} V(t_2, x_2) - V(t_1, x_1) &\leq \int_{t_2}^{\tau} h(s, x_2(s), \underline{\alpha}(s), \underline{u}(s), \beta[\underline{u}](s))ds + g(x_2(\tau)) \\ &\quad - \int_{t_1}^{\tilde{\tau}} h(s, x_1(s), \tilde{\alpha}(s), u(s), \tilde{\beta}[u](s))ds - g(x_1(\tilde{\tau})) + 4\epsilon \end{aligned}$$

Now we can have 2 cases:

- $\tilde{\tau} > t_2$
- $\tilde{\tau} \leq t_2$

For the case $\tilde{\tau} > t_2$, the proof is similar to **Part I**. Thus we get

$$V(t_2, x_2) - V(t_1, x_1) \leq \bar{C}(|t_1 - t_2| + |x_1 - x_2|) + 4\epsilon$$

For the case $\tilde{\tau} \leq t_2$

$$\begin{aligned} V(t_2, x_2) - V(t_1, x_1) &\leq \int_{t_2}^{\tau} h(s, x_2(s), \underline{\alpha}(s), \underline{u}(s), \beta[\underline{u}](s))ds + g(x_2(\tau)) - \\ &\quad \int_{t_1}^{\tilde{\tau}} h(s, x_1(s), \tilde{\alpha}(s), u(s), \tilde{\beta}[u](s))ds - g(x_1(\tau)) + 4\epsilon \quad (5.25) \end{aligned}$$

$$\begin{aligned}
&\leq \int_{t_2}^{\tau} h(s, x_2(s), \underline{\alpha}(s), \underline{u}(s), \beta[\underline{u}](s)) ds + g(x_2(\tau)) - \\
&\quad \int_{t_1}^{\tilde{\tau}} h(s, x_1(s), \tilde{\alpha}(s), u(s), \tilde{\beta}[u](s)) ds - g(x_1(\tau)) \\
&\quad + \int_{\tilde{\tau}}^{\tau} h(s, x_1(s), \tilde{\alpha}(s), u(s), \tilde{\beta}[u](s)) ds \\
&\quad - \int_{\tilde{\tau}}^{\tau} h(s, x_1(s), \tilde{\alpha}(s), u(s), \tilde{\beta}[u](s)) ds + 4\epsilon
\end{aligned}$$

Since $\tilde{\alpha} = \underline{\alpha}$ and $u = \tilde{u}$ and $\tilde{\beta}[u] = \beta[\underline{u}]$ for $t_2 \leq s \leq T$.

$$\begin{aligned}
&\leq \int_{t_2}^{\tau} (h(s, x_2(s), \underline{\alpha}(s), \underline{u}(s), \beta[\underline{u}](s)) \\
&\quad - h(s, x_1(s), \underline{\alpha}(s), \underline{u}(s), \beta[\underline{u}](s))) ds \\
&\quad + g(x_2(\tau)) - g(x_1(\tau)) \\
&\quad + \int_{\tilde{\tau}}^{\tau} h(s, x_1(s), \tilde{\alpha}(s), u(s), \tilde{\beta}[u](s)) ds \\
&\quad - \int_{t_1}^{t_2} h(s, x_1(s), \tilde{\alpha}(s), u(s), \tilde{\beta}[u](s)) ds + 4\epsilon
\end{aligned}$$

Thus using equations 5.7, 5.8, 5.9 and 5.25 we get

$$V(t_2, x_2) - V(t_1, x_1) \leq \bar{C}(|t_1 - t_2| + |x_1 - x_2|) + 4\epsilon \quad (5.26)$$

Thus using equations 5.21 and 5.26 we get $|V(t, x) - V(\hat{t}, \hat{x})| \leq C_4(|t - \hat{t}| + |x - \hat{x}|)$.

Now we will prove two more lemmas that enable us to prove one of the main results of the paper, namely that the value function given in equation 5.10 is the viscosity solution of a certain HJI PDE.

Lemma 12 *Let ϕ be C^1 . Then if ϕ satisfies*

$$\bullet D_t \phi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x \phi(t_0, x_0)) \leq -\theta$$

then for a small enough $\delta > 0$, there exists $\beta \in \Delta(t_0)$ and $\bar{\alpha} \in \bar{\mathcal{A}}(t_0)$ for all $u \in \mathcal{M}(t_0)$ such that

$$\begin{aligned}
&\phi((t_0 + \delta), x(t_0 + \delta)) - \phi(t_0, x_0) \\
&\quad + \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \leq -\frac{\theta \delta}{2}
\end{aligned}$$

and if

$$\bullet D_t \phi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x \phi(t_0, x_0)) \geq \theta$$

then for a small enough $\delta > 0$, there exists $u \in \mathcal{M}(t_0)$ for all $\beta \in \Delta(t_0)$ and $\bar{\alpha} \in \bar{\mathcal{A}}(t_0)$ such that

$$\begin{aligned} \phi((t_0 + \delta), x(t_0 + \delta)) - \phi(t_0, x_0) \\ + \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \geq \frac{\theta \delta}{2} \end{aligned}$$

$$\text{Where } \theta > 0 \text{ and } \mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} [f(t, x, \alpha, u, d) \cdot p + h(t, x, \alpha, u, d)]$$

$$\text{Note: } D_x \phi(t, x) = \frac{\partial \phi(t, x)}{\partial x} \text{ and } D_t \phi(t, x) = \frac{\partial \phi(t, x)}{\partial t}$$

Proof: Part I Let

$$\Lambda(t, x, \alpha, u, d) = D_t \phi(t, x) + f(t, x, \alpha, u, d) \cdot D_x \phi(t, x) + h(t, x, \alpha, u, d)$$

If $\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} \Lambda(t_0, x_0, \alpha, u, d) \leq -\theta < 0$, then for each $u \in \mathcal{U}$ there exists $d = d(u) \in \mathcal{D}$ and for each u, d there exists $\alpha = \alpha(u, d) \in \mathcal{A}$ such that $\Lambda(t_0, x_0, u, d, \alpha) \leq -\theta$. Since Λ is uniformly continuous we have in fact

$$\Lambda(t_0, x_0, \alpha, \sigma, d) \leq \frac{-3\theta}{4}$$

for all $\sigma \in \mathcal{B}(u, r) \cap \mathcal{U}$ and some $r = r(u) > 0$. Since $\mathcal{U}$ is compact, there exist finitely many distinct points $u_1, u_2, \dots, u_n \in \mathcal{U}, d_1, d_2, \dots, d_n \in \mathcal{D}, \alpha_1, \alpha_2, \dots, \alpha_n \in \mathcal{A}$ and $r_1, r_2, \dots, r_n > 0$ such that

$$\mathcal{U} \subset \bigcup_{i=1}^n \mathcal{B}(u_i, r_i)$$

and

$$\Lambda(t_0, x_0, \alpha_i, \sigma, d_i) \leq \frac{-3\theta}{4} \quad \text{for } \sigma \in \mathcal{B}(u_i, r_i)$$

We define

$$\begin{cases} \pi_1 : \mathcal{U} \rightarrow \mathcal{D} \\ \pi_2 : \mathcal{U} \times \mathcal{D} \rightarrow \mathcal{A} \end{cases} \quad \text{by setting} \quad \begin{cases} \pi_1(u) = d_k \text{ if } \pi_1(u) \in \mathcal{B}(d_k, r_k) \setminus \bigcup_{i=1}^{k-1} \mathcal{B}(d_i, r_i) \\ \pi_2(u, d) = \alpha_k \text{ if } \pi_2(u, d) \in \mathcal{B}(\alpha_k, r_k) \setminus \bigcup_{i=1}^{k-1} \mathcal{B}(\alpha_i, r_i) \end{cases}$$

Thus $\Lambda(t_0, x_0, \pi_2(u, \pi_1(u)), u, \pi_1(u)) \leq \frac{-3\theta}{4}$ for all $u \in \mathcal{U}$. Since Λ is uniformly continuous we therefore have for each sufficiently small $\delta > 0$

$$\Lambda(s, x(s), \pi_2(u, \pi_1(u)), u, \pi_1(u)) \leq \frac{-\theta}{2}$$

for all $u \in \mathcal{U}$, $t_0 \leq s \leq t_0 + \delta$ and any solution $x(\cdot)$ of ODE on $(t_0, t_0 + \delta)$ for any $d(\cdot)$, $u(\cdot)$, $\alpha(\cdot)$ with initial candidates $x(t_0) = x_0$. Finally we define $\beta \in \Delta(t_0)$ in the following way:

$$\beta[u](s) = \pi_1(u(s))$$

and define $\bar{\alpha} \in \bar{\mathcal{A}}(t_0)$ in the following way:

$$\bar{\alpha}(s) = \pi_2(u(s), \beta(s))$$

for each $u \in \mathcal{M}(t_0)$, $t_0 \leq s \leq T$,

$$\Lambda(s, x(s), \bar{\alpha}, u(s), \beta[u](s)) \leq \frac{-\theta}{2} \quad (t_0 \leq s \leq t)$$

for each $u \in \mathcal{M}(t_0)$. We integrate the above for t_0 to $t_0 + \delta$, adding and subtracting $\phi(t_0, x_0)$ and we get

$$\begin{aligned} \phi(t_0 + \delta, x(t_0 + \delta)) - \phi(t_0, x_0) \\ + \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \leq -\frac{\theta\delta}{2} \end{aligned}$$

Part II Let

$$\Lambda(t, x, \alpha, u, d) = D_t \phi(t, x) + f(t, x, \alpha, u, d) \cdot D_x \phi(t, x) + h(t, x, \alpha, u, d)$$

then if

$$\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} \Lambda(t_0, x_0, \alpha, u, d) \geq \theta > 0$$

Thus there exists a $u^* \in \mathcal{U}$ such that

$$\min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} \Lambda(t_0, x_0, \alpha, u^*, d) \geq \theta$$

Since Λ is uniformly continuous, we have that

$$\Lambda(s, x(s), \alpha, u^*, d) \geq \frac{\theta}{2}$$

provided $t_0 \leq s \leq t_0 + \delta$ (for any small $\delta > 0$) and $x(\cdot)$ which solves the ODE on $(t_0, t_0 + \delta)$ for any $u(\cdot)$, $d(\cdot)$, $\alpha(\cdot)$ with initial candidates $x(t_0) = x_0$. Hence, for $u(\cdot) \equiv u^*$ and any $\bar{\alpha} \in \bar{\mathcal{A}}(t_0)$, $\beta \in \Delta(t_0)$

$$D_t \phi(s, x(s)) + f(s, x(s), \bar{\alpha}, u(s), \beta[u](s)) \cdot D_x \phi(s, x(s)) + h(s, x(s), \bar{\alpha}, u(s), \beta[u](s)) \geq \frac{\theta}{2}$$

Integrating this expression from t_0 to $t_0 + \delta$, adding and subtracting $\phi(t_0, x_0)$ we get

$$\phi((t_0 + \delta), x(t_0 + \delta)) - \phi(t_0, x_0) + \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \geq \frac{\theta \delta}{2}$$

Lemma 13 Let ϕ be C^1 . Then if ϕ satisfies that

$$\bullet D_t \phi(t_0, x_0) + \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t_0, x_0, \alpha, u, d) \leq -\theta$$

then for a small enough $\delta > 0$, there exists $\beta \in \Delta(t_0)$ and $\bar{\alpha} \in \bar{\mathcal{A}}(t_0)$ for all $u \in \mathcal{M}(t_0)$, such that

$$\phi((t_0 + \delta), x(t_0)) - \phi(t_0, x_0) + \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \leq -\frac{\theta \delta}{2}$$

and if

$$\bullet D_t \phi(t_0, x_0) + \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t_0, x_0, \alpha, u, d) \geq \theta$$

then for a small enough $\delta > 0$, there exists $u \in \mathcal{M}(t_0)$ for all $\beta \in \Delta(t_0)$ and $\bar{\alpha} \in \bar{\mathcal{A}}(t_0)$, such that

$$\phi((t_0 + \delta), x(t_0)) - \phi(t_0, x_0) + \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \geq \frac{\theta \delta}{2}$$

Where $\theta > 0$

Proof: The proof of Lemma is similar to the proof of the Lemma 12 and follows the same logic.

Now we have all the intermediate results to prove that the value function is a viscosity solution of a certain HJI PDE. The proof follows a similar logic to the legacy proof of theorem 1 in [10], theorem 4.1 in [27].

Theorem 8 $V(t, x)$ is the viscosity solution of the terminal value HJI PDE:

$$D_t V(t, x) + \min\{\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t, x, \alpha, u, d), \mathcal{H}(x, t, D_x V)\} = 0$$

$$V(T, x) = g(x)$$

where

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} [f(t, x, \alpha, u, d) \cdot p + h(t, x, \alpha, u, d)]$$

Proof: To prove this, it is sufficient to prove the following claims [34], [27], [6], [28]:

1. For $\phi(x, t) \in C^1(\mathbb{R}^m \times (-\infty, 0])$, such that $V - \phi$ has a local maxima at $(x_0, t_0) \in \mathbb{R}^m \times (\infty, 0]$ then

$$D_t \phi(x_0, t_0) + \min\{\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t_0, x_0, \alpha, u, d), \mathcal{H}(t_0, x_0, D_x V)\} \geq 0$$

2. For $\phi(x, t) \in C^1(\mathbb{R}^m \times (-\infty, 0])$, such that $V - \phi$ has a local minimum at $(x_0, t_0) \in \mathbb{R}^m \times (-\infty, 0]$ then

$$D_t \phi(x_0, t_0) + \min\{\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t_0, x_0, \alpha, u, d), \mathcal{H}(t_0, x_0, D_x V)\} \leq 0$$

Claim 1

For $\phi(x, t) \in C^1(\mathbb{R}^m \times (-\infty, 0])$, such that $V - \phi$ has a local maxima at $(x_0, t_0) \in \mathbb{R}^m \times (\infty, 0]$ and let us assume for the sake of contradiction that $\exists \theta > 0$, such that

$$D_t \phi(x_0, t_0) + \min\{\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t_0, x_0, \alpha, u, d), \mathcal{H}(t_0, x_0, D_x V)\} \leq -\theta$$

Now there are two cases.

Case I

$$\mathcal{H}(t_0, x_0, D_x V) \leq \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t_0, x_0, \alpha, u, d)$$

then

$$D_t \phi(x_0, t_0) + \mathcal{H}(t_0, x_0, D_x V) \leq -\theta$$

⇒ According to Lemma 12

$$\begin{aligned} \phi((t_0 + \delta, x_0(t_0 + \delta))) - \phi(t_0, x_0) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \\ \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \leq \frac{-\theta \delta}{2} \end{aligned}$$

Since we know that $V - \phi$ has a local maxima at (x_0, t_0) thus

$$V(t_0, x_0) - \phi(t_0, x_0) \geq V(t_0 + \delta, x_0(t_0 + \delta)) - \phi(t_0 + \delta, x_0(t_0 + \delta))$$

$$\begin{aligned} \Rightarrow V(t_0 + \delta, x_0(t_0 + \delta)) - V(t_0, x_0) \leq \\ \phi(t_0 + \delta, x_0(t_0 + \delta)) - \phi(t_0, x_0) \end{aligned}$$

$$\begin{aligned} \Rightarrow V(t_0 + \delta, x_0(t_0 + \delta)) - V(t_0, x_0) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \\ \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \leq \frac{-\theta \delta}{2} \end{aligned}$$

But from Theorem 7 we know that

$$\begin{aligned} V(t, x) \leq V(t + \delta, x(t + \delta)) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t)} \inf_{\beta \in \Delta(t)} \sup_{u \in \mathcal{M}(t)} \\ \int_t^{t + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \end{aligned}$$

This leads to

$$0 \leq \frac{-\theta \delta}{2}$$

which is a contradiction.

Case II

$$\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t_0, x_0, \alpha, u, d) \leq \mathcal{H}(t_0, x_0, D_x V)$$

then

$$D_t \phi(x_0, t_0) + \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t_0, x_0, \alpha, u, d) \leq -\theta$$

according to Lemma 13

$$\phi((t_0 + \delta, x_0) - \phi(t_0, x_0) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \leq \frac{-\theta\delta}{2}$$

Since we know that $V - \phi$ has a local maxima at (x_0, t_0) thus

$$V(t_0, x_0) - \phi(t_0, x_0) \geq V(t_0 + \delta, x_0) - \phi(t_0 + \delta, x_0)$$

$$\Rightarrow V(t_0 + \delta, x_0) - V(t_0, x_0) \leq \phi(t_0 + \delta, x_0) - \phi(t_0, x_0)$$

$$\Rightarrow V(t_0 + \delta, x_0) - V(t_0, x_0) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \leq \frac{-\theta\delta}{2}$$

But from Lemma 8 we know that

$$V(t, x) \leq V(t + \delta, x) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds$$

This leads to

$$0 \leq \frac{-\theta\delta}{2}$$

which is a contradiction.

Thus claim 1 is true.

Claim 2

For $\phi(x, t) \in C^1(\mathbb{R}^m \times (-\infty, 0])$, such that $V(\cdot) - \phi(\cdot)$ has a local minimum at $(x_0, t_0) \in \mathbb{R}^m \times (-\infty, 0]$ and let us assume for the purpose of a contradiction that $\exists \theta \geq 0$, such that

$$D_t \phi(x_0, t_0) + \min_{u \in \mathcal{U}} \{ \max_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t, x, \alpha, u, d), \mathcal{H}(t_0, x_0, D_x V) \} \geq \theta$$

then

$$D_t \phi(x_0, t_0) + \max_{u \in \mathcal{U}} \{ \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} h(t_0, x_0, \alpha, u, d) \} \geq \theta$$

and furthermore

$$D_t \phi(x_0, t_0) + \mathcal{H}(t_0, x_0, D_x V) \geq \theta$$

and from Lemma 12 and Lemma 13 we get

$$\begin{aligned} \phi((t_0 + \delta), x(t_0 + \delta)) - \phi(t_0, x_0) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \\ \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \geq \frac{\theta \delta}{2} \end{aligned}$$

$$\begin{aligned} \phi((t_0 + \delta), x_0) - \phi(t_0, x_0) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \\ \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \geq \frac{\theta \delta}{2} \end{aligned}$$

Since $V(\cdot) - \phi(\cdot)$ has a local minimum at (t_0, x_0)

$$V(t_0, x_0) - \phi(t_0, x_0) \leq V(t_0 + \delta, x(t_0 + \delta)) - \phi(t_0 + \delta, x(t_0 + \delta))$$

$$\begin{aligned} \Rightarrow \phi(t_0 + \delta, x(t_0 + \delta)) - \phi(t_0, x_0) \\ \leq V(t_0 + \delta, x(t_0 + \delta)) - V(t_0, x_0) \end{aligned}$$

$$\begin{aligned} \Rightarrow V(t_0 + \delta, x(t_0 + \delta)) - V(t_0, x_0) + \\ \inf_{\bar{\alpha} \in \bar{\mathcal{A}}} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \\ \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \geq \frac{\theta \delta}{2} \quad (5.27) \end{aligned}$$

Using the same logic we can also get

$$\begin{aligned} \Rightarrow V(t_0 + \delta, x_0) - V(t_0, x_0) + \\ \inf_{\bar{\alpha} \in \bar{\mathcal{A}}} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \\ \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \geq \frac{\theta \delta}{2} \quad (5.28) \end{aligned}$$

Now there are two cases based on theorem 7.

Case I

$$\begin{aligned}
V(t_0 + \delta, x(t_0 + \delta)) + \int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \\
\leq \min_{\tau \in [t_0, t_0 + \delta]} \left(\int_{t_0}^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right)
\end{aligned}$$

then

$$\begin{aligned}
V(t_0, x) = V(t_0 + \delta, x(t_0 + \delta)) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \\
\int_{t_0}^{t_0 + \delta} h(s, x(s), u(s), \beta[u](s)) ds
\end{aligned}$$

substituting this in equation 5.27 leads to

$$\frac{\theta \delta}{2} \leq 0$$

This is a contradiction.

Case II

$$\begin{aligned}
\min_{\tau \in [t_0, t_0 + \delta]} \left(\int_{t_0}^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \right. \\
\left. + g(x(\tau)) \right) \leq V(t_0 + \delta, x(t_0 + \delta)) + \\
\int_{t_0}^{t_0 + \delta} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds
\end{aligned}$$

then

$$\begin{aligned}
V(t_0, x_0) = \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \min_{\tau \in [t_0, t_0 + \delta]} \\
\left(\int_{t_0}^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right)
\end{aligned}$$

Then there exists a $\tau \in [t_0, t_0 + \delta]$ that minimises $V(t_0, x_0)$ thus we have:

$$\begin{aligned}
V(t_0, x_0) = \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \\
\left(\int_{t_0}^{\tau} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds + g(x(\tau)) \right)
\end{aligned}$$

Then using lemma 9 we get that:

$$V(\tau', x(\tau')) + \inf_{\bar{\alpha} \in \bar{\mathcal{A}}(t_0)} \inf_{\beta \in \Delta(t_0)} \sup_{u \in \mathcal{M}(t_0)} \int_{t_0}^{\tau'} h(s, x(s), \bar{\alpha}(s), u(s), \beta[u](s)) ds \leq V(t_0, x_0)$$

$\tau' \in [t_0, \tau]$

This will lead to a contradiction with equation 5.27.

5.3. RELATIONSHIP WITH THE AVOID BACKWARD REACHABLE TUBE

After deriving the Hamilton-Jacobi-Isaacs (HJI) PDE for systems with uncertain dynamics, we now establish its connection to the concept of avoid backward reachable tubes. The avoid backward reachable tube for a given system is defined as the set of initial states for which there exists a disturbance that drives the system into a specified avoid set, irrespective of any control action taken, within a given time horizon [23], [24]. This concept builds on the concept of the backward reachable set, which was introduced in [11] as the set of initial states for which a control action exists to drive the system into a target set at an exact given time. In [10], this idea was generalised to include a time horizon, resulting in the backward reachable tube. This concept was further extended in [35] using a game-theoretic approach and showing that the avoid backward reachable tube could be computed using a Hamilton-Jacobi-Isaacs (HJI) equation. The level set of the value function in the HJI equation provides a representation of the avoid tube, as defined in [10] and other foundational works. This approach applies to systems with deterministic dynamics. However, if the dynamics are uncertain and the uncertainty is bounded, the problem becomes significantly more complex. To address this, we formulate a three-player game, where:

- Player 1 represents the control input,
- Player 2 represents the disturbance, and
- Player 3 represents the uncertain system dynamics.

In the robust framework, we require the avoid backward reachable tube to be valid for any realisation of the bounded system dynamics. This approach parallels the method proposed in [35], where disturbances are incorporated as adversarial players. However, here, the third player accounts for unmodeled dynamics, and robustness is achieved by ensuring that the avoid tube is valid for all possible realisations of the system.

We now define a complementary system to the ordinary differential equation described in equation (5.1):

$$\tilde{f} : [0, T] \times \mathbb{R}^m \times \mathcal{A} \times \mathcal{U} \times \mathcal{D} \rightarrow \mathbb{R}^m \quad (5.29)$$

such that:

$$\begin{aligned} \dot{\tilde{x}}(s) &= \tilde{f}(s, \tilde{x}(s), \tilde{\alpha}(s), \tilde{u}(s), \tilde{d}(s)) = \dot{x}(T-s) = \\ &f(T-s, x(T-s), \alpha(T-s), u(T-s), d(T-s)) \end{aligned}$$

We also define:

$$\kappa(T-t) \equiv \tilde{\kappa}(t)$$

$$\mathcal{M}(T-t) \equiv \tilde{\mathcal{M}}(t)$$

$$\mathcal{N}(T-t) \equiv \tilde{\mathcal{N}}(t)$$

and:

$$\Delta(T-t) \equiv \tilde{\Delta}(t) \quad \beta(T-t) = \tilde{\beta}(t)$$

$$\overline{\mathcal{A}}(T-t) \equiv \tilde{\mathcal{A}}(t) \quad \overline{\alpha}(T-t) = \tilde{\alpha}(t)$$

5

With this setup, the trajectory of the system given in equation (5.1) is identical to the trajectory of the complementary system described in equation (5.29), provided that $\tilde{x}(t) = x(T-t)$ for the respective parameters, control, and disturbance.

Now the robust avoid backward reachable tube is defined as the set of initial states for which there exists a disturbance that drives the system into a specified avoid set, irrespective of any control action taken, within a given time horizon and for any realisation of the system dynamics. The set is represented as shown in equation 5.30.

$$\begin{aligned} V_{RABRT}(T, \mathcal{L}) &= \{x \in \mathbb{R}^m : \exists \tilde{\alpha} \in \tilde{\mathcal{A}}(T), \forall \tilde{u} \in \tilde{\mathcal{M}}(T), \exists \tilde{\beta} \in \tilde{\Delta}(T), \\ &\quad \exists t \in [0, T], \tilde{x}(t) \in \mathcal{L}\} \quad (5.30) \end{aligned}$$

where $\mathcal{L}$ is the avoid set. In a differential game setting, the information available to each player significantly impacts the outcome. For this work, we adopt the following information structure:

- Player 1 (Input): Does not anticipate the actions of Players 2 and 3.
- Player 2 (Disturbance): Uses non-anticipative strategies, meaning it cannot react to input signals from Player 1 before they occur.
- Player 3 (System Dynamics): Has full knowledge of the strategies of both Player 1 and Player 2 and uses non-anticipative strategies for both.

This structure aligns with the assumptions made in [35] and [24]. The choice to give the disturbance and uncertain dynamics an informational advantage ensures that the derived robust sets account for worst-case scenarios, making them valid for all possible inputs and system realisations. We make the following assumption regarding the avoid set:

Assumption 7 *The avoid set of our reachability problem is closed and can be represented as the zero sub-level set of a bounded and Lipschitz-continuous function, as defined in Section 5.2.*

$$\mathcal{L} = \{x \in \mathbb{R}^m : g(x) \leq 0\}$$

Remark: If the avoid set is closed ($\mathcal{L}$), then the robust avoid backward reachable tube, $V_{RABRT}(\mathcal{L}, \tau)$ is also closed.

Lemma 14 *The value function given in equation 5.10 with $h(\cdot) = 0$ provides the avoid robust backward reachable set given in equation 5.30, i.e.,*

$$x \in V_{RABRT}(t, \mathcal{L}) \Leftrightarrow V(T - t, x) \leq 0$$

Note: The $V_{RABRT}(t, \mathcal{L})$ is for the complementary system (equation 5.29) while the value function is for the original system (equation 5.1).

Proof: To prove this we will show the following:

- $x \in V_{RABRT}(t, \mathcal{L}) \Rightarrow V(T - t, x) \leq 0$
- $V(T - t, x) \leq 0 \Rightarrow x \in V_{RABRT}(t, \mathcal{L})$

Part I Let us say for contradiction purposes that:

$$x \in V_{RABRT}(t, \mathcal{L}) \Rightarrow V(T - t, x) > 0$$

Since $V(T - t, x) > 0$, from the definition of the value function (equation 5.10), for all $\bar{\alpha} \in \bar{\mathcal{A}}(T - t)$ and $\beta \in \Delta(T - t)$ we have

$$\sup_{u \in \mathcal{M}(T-t)} \inf_{\tau \in [T-t, T]} g(x(\tau)) > 0$$

This implies that there exists a $u_1 \in \mathcal{M}(T - t)$, such that for any $\epsilon > 0$ we have, for all $\tau \in [T - t, T]$

$$g(x(\tau)) + \epsilon > 0$$

Now from the definition of $V_{RABRT}(t, \mathcal{L})$ (equation 5.30) we know that there exists $\tilde{\alpha} \in \tilde{\mathcal{A}}(t)$, $\tilde{\beta} \in \tilde{\Delta}(t)$, $\tilde{\tau} \in [0, t]$, such that for all $\tilde{u} \in \tilde{\mathcal{M}}(t)$ we have $\tilde{x}(\tilde{\tau}) \in \mathcal{L}$. From assumption 7 we know that

$$g(\tilde{x}(\tilde{\tau})) \leq 0$$

Let $\bar{\alpha} = \tilde{\alpha}$, $\beta = \tilde{\beta}$, $\tau = \tilde{\tau}$ and $\tilde{u} = u_1$. This implies that both the original system (5.1) and the complementary system (5.29) have the same trajectory. Thus $g(\tilde{x}(\tilde{\tau})) = g(x(\tau))$. This is a contradiction.

Part II For the second part we contradict the compliment of the statement. Since we want to prove $V(T-t, x) \leq 0 \Rightarrow x \in V_{RABRT}(t, \mathcal{L})$, this implies $x \notin V_{RABRT}(t, \mathcal{L}) \Rightarrow V(T-t, x) > 0$. Now we assume the contradiction:

$$x \notin V_{RABRT}(t, \mathcal{L}) \Rightarrow V(T-t, x) \leq 0$$

Since $V(T-t, x) \leq 0$, from the definition of the value function (equation 5.10) there exist $\bar{\alpha} \in \bar{\mathcal{A}}(T-t)$ and $\beta \in \Delta(T-t)$, such that for any $\epsilon > 0$ we have

$$\sup_{u \in \mathcal{M}(T-t)} \inf_{\tau \in [T-t, T]} g(x(\tau)) \leq 2\epsilon$$

This implies that for all $u \in \mathcal{M}(T-t)$ we have

$$\inf_{\tau \in [T-t, T]} g(x(\tau)) \leq 2\epsilon$$

Thus for any $\epsilon > 0$, there exists a $\tau \in [T-t, T]$

$$g(x(\tau)) \leq 3\epsilon$$

From the definition of $V_{RABRT}(t, \mathcal{L})$ (equation 5.30) we know that, for all $\tilde{\alpha} \in \tilde{\mathcal{A}}(t)$, $\tilde{\beta} \in \tilde{\Delta}(t)$, $\tilde{\tau} \in [0, t]$, there exists a $\tilde{u} \in \tilde{\mathcal{M}}(t)$, such that $\tilde{x}(\tilde{\tau}) \notin \mathcal{L}$. From assumption 7 we know that

$$g(\tilde{x}(\tilde{\tau})) > 0$$

Let $\tilde{\alpha} = \bar{\alpha}$, $\tilde{\beta} = \beta$, $\tilde{\tau} = \tau$ and $u = \tilde{u}$. This implies that both the original system (5.1) and the complementary system (5.29) have the same trajectory. Thus $g(\tilde{x}(\tilde{\tau})) = g(x(\tau))$. This is a contradiction.

Theorem 9 Let $V(t, x)$ be the viscosity solution of the terminal value HJI PDE given below

$$D_t V(t, x) + \min\{0, \mathcal{H}(x, t, D_x V)\} = 0$$

$$V(T, x) = g(x)$$

where

$$H(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \min_{\alpha \in \mathcal{A}} [f(t, x, \alpha, u, d) \cdot p]$$

Then the sub zero level set of the $V(t, x)$ provides the $V_{RABRT}(T-t, \mathcal{L})$ for the complementary system given in equation 5.29.

$$V_{RABRT}(T-t, \mathcal{L}) = \{x \in \mathbb{R}^m : V(t, x) \leq 0\}$$

Proof: The proof of the theorem follows directly from Lemma 14 and Theorem 8.

5.4. NUMERICAL EXAMPLE

In this section, we demonstrate the application of the proposed framework using the Dubins car system as an example. The Dubins car is a simple model often used to represent systems with bounded turning rates and forward velocity constraints. The system dynamics are given by:

$$\begin{aligned}\dot{x} &= v \cos(\theta) + d_x, \\ \dot{y} &= v \sin(\theta) + d_y, \\ \dot{\theta} &= u + d_\theta,\end{aligned}\tag{5.31}$$

where:

- $(x, y) \in \mathbb{R}^2$ represents the position of the car,
- $\theta \in \mathbb{R}$ represents the heading angle,
- $v \in [0.5, 0.7]$ is the forward velocity of the car,
- $u \in [-1, 1]$ is the control input (turning rate), and
- $d_x, d_y, d_\theta \in [-(0.1, 0.1, 0.7), (0.1, 0.1, 0.7)]$ is the disturbance acting on the system.

The objective is to compute the robust avoid backward reachable tube for an avoid set defined as:

$$\mathcal{L} = \{(x, y, \theta) \in \mathbb{R}^3 : g(x, y, \theta) \leq 0\},\tag{5.32}$$

where $g(x, y, \theta)$ is a Lipschitz-continuous function defining the unsafe states. The avoid set is a cylinder with height along θ and a radius of 1 centered at origin. Using the extended Hamilton-Jacobi-Isaacs (HJI) framework, we compute the value function $V(t, x, y, \theta)$, where the zero sub-level set $\{(x, y, \theta) : V(t, x, y, \theta) \leq 0\}$ characterises the robust avoid backward reachable tube.

Note: In this case, the original system and the complementary system are the same. This is typically true for time-invariant system dynamics.

The numerical computations were performed using level set methods implemented in the Level Set method based Toolbox [29] with a time horizon of $T = 2$. The results of the computation are visualised in Figures 5.1 and 5.2.

The results validate the proposed framework by demonstrating its ability to compute robust avoid backward reachable tubes for systems with bounded uncertainties. The inclusion of unmodeled dynamics (third player) ensures robustness to worst-case realisations of the system, as shown by the conservative nature of the computed tube. This example illustrates how the proposed three-player game framework can be integrated seamlessly with existing numerical tools, such as the Level Set Toolbox, to provide rigorous safety verification under uncertain dynamics.




Figure 5.1: Zero level contour of the V_{RABRT}

5.5. CONCLUSIONS

In this paper, we extended the Hamilton-Jacobi-Isaacs (HJI) framework to address optimal control problems for systems with bounded uncertain dynamics. By formulating a three-player game, we incorporated unmodeled dynamics alongside the controller and disturbance, enhancing the robustness of reachability analysis. We proved that the value function associated with this game is the viscosity solution to the extended HJI equation, providing a rigorous theoretical foundation for robust control. Using this framework, we defined the robust avoid backward reachable tube, which characterises the set of initial states that will reach an avoid set under worst-case disturbances and all realisations of system dynamics. We established a direct connection between the value function and the robust avoid backward reachable tube, demonstrating that the zero sub-level set of the value function precisely represents this robust set. Finally, a numerical example highlighted the practical applicability of our approach, showcasing its




Figure 5.2: Slices of the zero level contour of the V_{RABRT}

integration with existing level set-based toolboxes. This work provides a robust method for safety verification under uncertainty, offering a pathway for future developments in reachability analysis for complex and uncertain systems.

While this work extends the Hamilton-Jacobi-Isaacs (HJI) framework to incorporate bounded uncertain dynamics using a three-player game, there remain several avenues for further exploration. One promising direction is to extend this approach to stochastic HJI equations, where the system dynamics include stochastic perturbations or noise. Stochastic dynamics present additional challenges, as they require accounting for probabilistic uncertainties and formulating the value function in terms of expected costs or risks. Integrating stochasticity into the current framework would enable robust safety verification in environments with both adversarial and random uncertainties, expanding the applicability of this method to a broader class of real-world systems.

5.6. APPENDIX

5.7. PROOF OF EQUATION 5.19

The following is to prove the equation 5.19.

$$|x_1(t_2) - x_1| = |x_1(t_2) - x_1(t_1)|$$

$$= \left| \int_t^{t_2} f(s, x_2(s), \alpha(s), u(s), d(s)) - \int_t^{t_1} f(s, x(s), \alpha(s), u(s), d(s)) \right|$$

Using equation 5.3

$$\leq |C_1(t_2 - t) - C_1(t_1 - t)| \leq C_1|t_1 - t_2|$$

The above holds for any $u \in \mathcal{U}$, $d \in \mathcal{D}$ and $\alpha \in \mathcal{A}$.

REFERENCES

- [1] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-jacobi reachability: A brief overview and recent advances”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 2242–2253.
- [2] S. Bansal and C. J. Tomlin. “Deepreach: A deep learning approach to high-dimensional reachability”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 1817–1824.
- [3] S. Bansal, A. Bajcsy, E. Ratner, A. D. Dragan, and C. J. Tomlin. “A hamilton-jacobi reachability-based framework for predicting and analyzing human motion for safe planning”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2020, pp. 7149–7155.
- [4] P. A. Forsyth and G. Labahn. “Numerical methods for controlled Hamilton-Jacobi-Bellman PDEs in finance”. In: *Journal of Computational Finance* 11.2 (2007), p. 1.
- [5] C. Livadas and N. A. Lynch. “Formal verification of safety-critical hybrid systems”. In: *International Workshop on Hybrid Systems: Computation and Control*. Springer. 1998, pp. 253–272.
- [6] J. Lygeros, D. N. Godbole, and S. Sastry. “Verified hybrid controllers for automated vehicles”. In: *IEEE transactions on automatic control* 43.4 (1998), pp. 522–539.
- [7] C. Livadas, J. Lygeros, and N. A. Lynch. “High-level modeling and analysis of the traffic alert and collision avoidance system (TCAS)”. In: *Proceedings of the IEEE* 88.7 (2000), pp. 926–948.
- [8] C. Tomlin, I. Mitchell, and R. Ghosh. “Safety verification of conflict resolution manoeuvres”. In: *IEEE Transactions on Intelligent Transportation Systems* 2.2 (2001), pp. 110–120.
- [9] J. J. Choi, D. Lee, B. Li, J. P. How, K. Sreenath, S. L. Herbert, and C. J. Tomlin. “A forward reachability perspective on robust control invariance and discount factors in reachability analysis”. In: *arXiv preprint arXiv:2310.17180* (2023).
- [10] J. Lygeros. “On reachability and minimum cost optimal control”. In: *Automatica* 40.6 (2004), pp. 917–927.
- [11] J. Lygeros, C. Tomlin, and S. Sastry. “Controllers for reachability specifications for hybrid systems”. In: *Automatica* 35.3 (1999), pp. 349–370.

- [12] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre. *Viability theory: new directions*. Springer Science & Business Media, 2011.
- [13] J.-P. Aubin and H. Frankowska. “Viability kernel of control systems”. In: *Nonlinear synthesis*. Springer, 1991, pp. 12–33.
- [14] P. Cardaliaguet. “A differential game with two players and one target”. In: *SIAM Journal on Control and Optimization* 34.4 (1996), pp. 1441–1460.
- [15] P. Cardaliaguet, M. Quincampoix, and P. Saint-Pierre. “Differential games through viability theory: Old and recent results”. In: *Advances in dynamic game theory*. Springer, 2007, pp. 3–35.
- [16] E. Asarin, O. Bournez, T. Dang, O. Maler, and A. Pnueli. “Effective synthesis of switching controllers for linear systems”. In: *Proceedings of the IEEE* 88.7 (2000), pp. 1011–1025.
- [17] I. M. Mitchell *et al.* “A toolbox of level set methods”. In: *UBC Department of Computer Science Technical Report TR-2007-11 1* (2007), p. 6.
- [18] I. Mitchell, A. M. Bayen, and C. J. Tomlin. “Validating a Hamilton-Jacobi approximation to hybrid system reachable sets”. In: *International workshop on hybrid systems: Computation and control*. Springer, 2001, pp. 418–432.
- [19] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin. “Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking”. In: *IEEE Transactions on Automatic Control* 66.12 (2021), pp. 5861–5876.
- [20] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin. “Decomposition of reachable sets and tubes for a class of nonlinear systems”. In: *IEEE Transactions on Automatic Control* 63.11 (2018), pp. 3675–3688.
- [21] J. N. Maidens, S. Kaynama, I. M. Mitchell, M. M. Oishi, and G. A. Dumont. “Lagrangian methods for approximating the viability kernel in high-dimensional systems”. In: *Automatica* 49.7 (2013), pp. 2017–2029.
- [22] M. Chen and C. J. Tomlin. “Hamilton-jacobi reachability: Some recent theoretical advances and applications in unmanned airspace management”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 1.1 (2018), pp. 333–358.
- [23] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin. “FaSTrack: A modular framework for fast and guaranteed safe motion planning”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 1517–1522.
- [24] S. L. Herbert. *Safe real-world autonomy in uncertain and unstructured environments*. University of California, Berkeley, 2020.
- [25] S. Herbert, J. J. Choi, S. Sanjeev, M. Gibson, K. Sreenath, and C. J. Tomlin. “Scalable learning of safety guarantees for autonomous systems using Hamilton-Jacobi reachability”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 5914–5920.

- [26] V. G. Boltyanski and A. S. Poznyak. *The robust maximum principle: theory and applications*. Springer Science & Business Media, 2011.
- [27] L. C. Evans and P. E. Souganidis. “Differential games and representation formulas for solutions of Hamilton-Jacobi-Isaacs equations”. In: *Indiana University mathematics journal* 33.5 (1984), pp. 773–797.
- [28] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin. “A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games”. In: *IEEE Transactions on automatic control* 50.7 (2005), pp. 947–957.
- [29] I. M. Mitchell and J. A. Templeton. “A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems”. In: *Hybrid Systems: Computation and Control: 8th International Workshop, HSCC 2005, Zurich, Switzerland, March 9-11, 2005. Proceedings* 8. Springer, 2005, pp. 480–494.
- [30] J. A. Sethian. *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*. Vol. 3. Cambridge university press, 1999.
- [31] S. Osher and R. P. Fedkiw. *Level set methods and dynamic implicit surfaces*. Vol. 1. Springer New York, 2005.
- [32] R. J. Elliott and N. J. Kalton. *The existence of value in differential games*. Vol. 126. American Mathematical Soc., 1972.
- [33] E. Roxin. “Axiomatic approach in differential games”. In: *Journal of Optimization Theory and Applications* 3 (1969), pp. 153–163.
- [34] M. Bardi, I. C. Dolcetta, et al. *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*. Vol. 12. Springer, 1997.
- [35] C. J. Tomlin, J. Lygeros, and S. S. Sastry. “A game theoretic approach to controller design for hybrid systems”. In: *Proceedings of the IEEE* 88.7 (2000), pp. 949–970.

6

CERTIFIED APPROXIMATE REACHABILITY (CARE)

Recent approaches to leveraging deep learning for computing reachable sets of continuous-time dynamical systems have gained popularity over traditional level-set methods, as they overcome the curse of dimensionality. However, as with level-set methods, considerable care needs to be taken in limiting approximation errors, particularly since no guarantees are provided during training on the accuracy of the learned reachable set. To address this limitation, we introduce an ϵ -approximate Hamilton-Jacobi Partial Differential Equation (HJ-PDE), which establishes a relationship between training loss and accuracy of the true reachable set. To formally certify this approximation, we leverage Satisfiability Modulo Theories (SMT) solvers to bound the residual error of the HJ-based loss function across the domain of interest. Leveraging Counter Example Guided Inductive Synthesis (CEGIS), we close the loop around learning and verification, by fine-tuning the neural network on counterexamples found by the SMT solver, thus improving the accuracy of the learned reachable set. To the best of our knowledge, Certified Approximate Reachability (CARE) is the first approach to provide soundness guarantees on learned reachable sets of continuous dynamical systems.

6.1. INTRODUCTION

Ensuring the safety of autonomous systems under uncertainty is a fundamental challenge in control and verification. Reachability analysis plays a key role in addressing safety concerns across various domains, including spacecraft trajectory design [1, 2], ground transportation systems [3, 4], air traffic management [5, 6], and flight control [7–9]. The theoretical foundations of reachability stem from viability theory [10, 11], while computational techniques have been developed for both exact and approximate reachable set computations in hybrid systems [12–16]. These advances have formulated reachability analysis as an optimal control problem, where reachable, viable, or invariant sets are represented as level sets of a value function satisfying a Hamilton-Jacobi (HJ) partial differential equation (PDE) [8, 17–20].

Traditionally, level set methods are used to compute the unique viscosity solution of the HJ equation [15]. However, even with high-order numerical schemes such as WENO methods, the non-differentiability of the value function introduces numerical inaccuracies [21]. While refining the grid used in finite element methods such as level set methods improves numerical accuracy, it is hard to predict what constitutes a sufficiently fine grid. Moreover, these grid-based approaches suffer from the curse of dimensionality, as the number of required grid points grows exponentially with system dimensionality.

Deepreach [22], a recent approach leveraging neural networks, has enabled solving the HJ equation without relying on finite difference methods. If the loss function enforcing the PDE conditions converges to zero in the region of interest, the learned value function is the unique viscosity solution of the HJ-PDE, ensuring that its zero-level set correctly represents the reachable set. This approach improves scalability to higher-dimensional systems by mitigating the curse of dimensionality.

In this work, we demonstrate an additional advantage of using neural networks to approximate the value function: the ability to provide a formally bounded ϵ -accurate reachable set. While both level-set methods and DeepReach yield empirically accurate solutions, their adherence to the HJ-PDE conditions—and thus the validity of the reachable set—has not been quantified. Follow-up work to DeepReach uses probabilistic methods to recover the reachable set [23–25]. In contrast, we provide a formal approach to verification that does not rely on recovery of the reachable set, but rather provides a sound over- and under-approximation of the reachable set. Our key contributions are:

1. We establish a formal bound ϵ on the residual error, ensuring that the absolute value of the HJ-based loss function remains within a specified threshold across the domain of interest.
2. We demonstrate that the neural network-based value function, constrained by ϵ , provides a certified approximation of the true value function, allowing for over- and under-approximation of the reachable set.

The remainder of the paper is structured as follows: Section 6.2 provides an overview of the reachability problems and introduces the Hamilton-Jacobi (HJ)

partial differential equation (PDE) whose solution we aim to learn. Section 6.3 presents the deep learning approach for computing the reachable set, while Section 6.4 discusses the formal verification method used to derive and certify an upper bound on the loss across the entire domain of interest. In Section 6.5, we introduce our main contribution, establishing a connection between the residual training loss and an under- and over-approximation of the reachable set. Finally, Section 6.6 provides implementation details and a case study.

6.2. PROBLEM SETUP

In reachability theory, a key objective is to compute the backward reachable set (BRS) of a dynamical system, which consists of all states from which trajectories can reach a given target set at the end of the time horizon. In contrast, the backward reachable tube (BRT) represents the set of states that can reach the target set within a specified time horizon [17]. If the target set represents unsafe states, the BRS/BRT identifies states that may lead to unsafe conditions and should be avoided. To account for adversarial disturbances, safety-critical scenarios are often modeled as a two-player game, where Player 1 represents the control input and Player 2 represents the disturbance input [26].

Mathematically, consider a dynamical system with state $x \in \mathbb{R}^m$ governed by the ordinary differential equation (ODE)

$$\dot{x}(s) \in f(s, x(s), u(s), d(s)), \quad t_0 \leq s \leq T, \quad (6.1)$$

where $T \geq t_0$ is the fixed time horizon. The initial state is given by $x(t_0) := x_0$, and the control and disturbance inputs are measurable functions

$$u : [t_0, T] \rightarrow \mathcal{U}, \quad d : [t_0, T] \rightarrow \mathcal{D},$$

with $\mathcal{U} \subset \mathbb{R}^k$ and $\mathcal{D} \subset \mathbb{R}^l$ being compact sets. The dynamics $f : [t_0, T] \times \mathbb{R}^m \times \mathcal{U} \times \mathcal{D} \rightarrow \mathbb{R}^m$ satisfy the following conditions for some constants C_1, C_2 :

$$|f(t, x, u, d)| \leq C_1, \quad (6.2)$$

$$|f(t, x, u, d) - f(t, \hat{x}, u, d)| \leq C_2 |x - \hat{x}|, \quad (6.3)$$

for all $t \in [t_0, T]$, $x, \hat{x} \in \mathbb{R}^m$, $u \in \mathcal{U}$, and $d \in \mathcal{D}$. These assumptions ensure that the ODE (6.1) admits a unique solution

$$x(t) = \phi(t, t_0, x_0, u(\cdot), d(\cdot)). \quad (6.4)$$

Next, we define the sets of control and disturbance policies:

$$\mathcal{M}_{[t, T]} \equiv \{u : [t, T] \rightarrow \mathcal{U} \mid u \text{ is measurable}\},$$

$$\mathcal{N}_{[t, T]} \equiv \{d : [t, T] \rightarrow \mathcal{D} \mid d \text{ is measurable}\}.$$

A nonanticipative strategy is a mapping $\beta : \mathcal{M}_{[t, T]} \rightarrow \mathcal{N}_{[t, T]}$ such that for all $s \in [t, T]$ and for all $u, \hat{u} \in \mathcal{M}_{[t, T]}$, if $u(\tau) = \hat{u}(\tau)$ for almost every $\tau \in [t, s]$,

then $\beta[u](\tau) = \beta[\hat{u}](\tau)$ for almost every $\tau \in [t, s]$. The class of such strategies is denoted as $\Delta_{[t, T]}$.

Given a target set $\mathcal{G}$ with a signed distance function $g(x)$ such that $g(x) \leq 0 \iff x \in \mathcal{G}$, the BRT is defined as

$$\text{BRT}_{\mathcal{G}}([t_0, T]) = \{x \mid \exists \beta \in \mathcal{N}_{[t_0, T]}, \forall u \in \mathcal{M}_{[t_0, T]}, \exists s \in [t_0, T] : \phi(s, t_0, x_0, u(\cdot), \beta(\cdot)) \in \mathcal{G}\}. \quad (6.5)$$

Equivalently, the BRT is given by the subzero level set of the value function $V^*(t, x)$, where

$$V^*(t, x) = \inf_{\beta \in \Delta_{[t, T]}} \sup_{u \in \mathcal{M}_{[t, T]}} \inf_{\tau \in [t, T]} g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))). \quad (6.6)$$

Let us introduce the shorthand notation $D_x V(t, x) = \frac{\partial V(t, x)}{\partial x}$ and $D_t V(t, x) = \frac{\partial V(t, x)}{\partial t}$. Then, a standard result of HJ reachability [17, 26, 27] allows us to reduce the optimization problem in Equation (6.6) from an infinite-dimensional optimization over policies to an optimization over control and disturbance inputs by formulating the value function as the viscosity solution of the HJ-PDE:

$$\begin{aligned} D_t V^*(t, x) + \min\{0, \mathcal{H}(t, x, D_x V^*)\} &= 0, \\ V^*(T, x) &= g(x), \end{aligned} \quad (6.7)$$

with the Hamiltonian defined as

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p].$$

The HJ-PDE can be solved using level-set methods [8, 16, 28] or through curriculum training of a neural network [22].

6.3. DEEP LEARNING APPROACH TO HJ REACHABILITY

Inspired by the self-supervised framework of Physics-Informed Neural Networks (PINNs) [29], the solution of the PDE (6.7) can be effectively learned, as demonstrated in [22].

Let V_θ denote the approximate value function learned by a neural network. To ensure that the zero level set of V_θ correctly represents the BRT, the network must achieve zero loss for the loss function defined as

$$\mathcal{L}(t_i, x_i; \theta) = 1(t_i = T) \mathcal{L}_1(x_i; \theta) + \lambda \mathcal{L}_2(t_i, x_i; \theta), \quad (6.8)$$

$$\mathcal{L}_1(x_i; \theta) = \|V_\theta(T, x_i) - g(x_i)\|, \quad (6.9)$$

$$\mathcal{L}_2(t_i, x_i; \theta) = \|D_t V_\theta(t_i, x_i) + \min[0, \mathcal{H}(t_i, x_i, D_x V_\theta)]\|, \quad (6.10)$$

where $1(t_i = T)$ is an indicator function ensuring that $\mathcal{L}_1$ is only applied at $t_i = T$, and λ controls the balance between the two loss terms. The non-linear nature of the neural network might result in areas of the state space with

significantly higher loss than empirically observed during training. This acute shortcoming necessitates a formal approach to verifying the approximation errors. To this end, in the subsequent sections, we will

1. establish a verification method to ensure $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$ for all x and t within the domain of interest; and
2. derive upper and lower bounds on V_θ in terms of viscosity solutions to the HJ-PDE, thus enabling under- and over-approximations of the BRT/BRS.

6.4. FORMAL SYNTHESIS OF NEURAL VALUE FUNCTIONS

Our approach to synthesizing the value function solution follows a two-phase process—a learning phase and a certification phase—that alternates in a Counter Example Guided Inductive Synthesis (CEGIS) loop [30]. The learning phase employs a curriculum training approach, discussed in detail in Section 6.6, terminating when we empirically satisfy $|\mathcal{L}_1(x_i; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t_i, x_i; \theta)| < \epsilon_2$, for N randomly sampled points $(t_1, x_1), \dots, (t_n, x_n)$.

To further ensure that $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$ for all x in the domain of interest $\mathcal{X}$ and for all $t \in [t_0, T]$, we design a sound certification phase: this certification is performed via SMT (Satisfiability Modulo Theories) solving, which symbolically reasons over the continuous domain $[t_0, T] \times \mathcal{X}$, hence generalizing across the sample-based loss error. While computationally intensive, formal verification provides soundness guarantees, allowing us to rigorously validate the learned value function as an approximate solution to the HJ-PDE.

Several approaches exist for certifying the accuracy of the learned value function. To remain general in our choice of dynamical systems, we employ an SMT solver capable of handling quantifier-free nonlinear real arithmetic formulae [31]. This allows us to incorporate arbitrary nonlinear activation functions into the neural network. Furthermore, it enables us to address systems with nonlinear dynamics, which induce nonlinearities in the Hamiltonian. Specifically, we employ dReal [32], which supports both polynomial and non-polynomial terms, such as transcendental functions (e.g., trigonometric and exponential functions).

After generating a symbolic representation of the neural network, the SMT solver searches for an assignment of variables (t, x) that satisfies the quantifier-free formula

$$(x \in \mathcal{X} \wedge \|V_\theta(T, x) - g(x)\| > \epsilon_1) \vee \quad (6.11)$$

$$(x \in \mathcal{X} \wedge t \in [t_0, T] \wedge \|D_t V_\theta(t, x) + \min[0, \mathcal{H}(t, x, D_x V_\theta)]\| > \epsilon_2) \quad (6.12)$$

The logical disjunction between (6.11) and (6.12) allows us to split the SMT call into separate parallel queries. If one query identifies a valid assignment,

i.e., a counterexample to either $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ or $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$, the remaining SMT calls terminate. A counterexample indicates the necessity to finetune the neural network, as it does not yet approximate a valid value function with sufficient accuracy. We leverage the counterexample for targeted training specifically at inputs where the network is insufficiently accurate. We sample additional training data points for finetuning in close proximity to the counterexample, thus reducing the need for extensive sampling in regions of the input space that are already well approximated. Finetuning and certification are alternately repeated within an inductive synthesis loop (CEGIS), progressively improving the neural value function in targeted regions of the input space, until the certifier determines the Formulae (6.11) and (6.12) to be unsatisfiable. This implies that there exists no counterexample to $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$, formally proving the validity of the learned value function over the entire domain of interest. In the following section, we derive results establishing how this formal certification can be leveraged to obtain sound over- and under-approximations of the BRT and BRS from the neural value function.

6

6.5. ϵ -ACCURATE VALUE FUNCTIONS

The certification of the neural value function in the previous section established that for all $(t, x) \in [t_0, T] \times \mathcal{X}$, there exists some $\hat{\epsilon} \in [0, \epsilon_2]$ such that

$$\|D_t V_\theta(t, x) + \min[0, \mathcal{H}(t, x, D_x V_\theta)]\| = \hat{\epsilon},$$

which can equivalently be rewritten as

$$\begin{aligned} \|D_t V_\theta(t, x) + \min[\pm \hat{\epsilon}, \tilde{\mathcal{H}}(t, x, D_x V_\theta)]\| &= \hat{\epsilon}, \\ \tilde{\mathcal{H}}(t, x, p) &= \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p \pm \hat{\epsilon}]. \end{aligned}$$

We consider the worst-case realisation of $\hat{\epsilon}$, i.e., $\hat{\epsilon} = \epsilon_2$. We will now show that, for $\epsilon = (\epsilon_1 + \epsilon_2(T - t))$,

$$V_\theta(t, x) - \epsilon \leq V^*(t, x) \leq V_\theta(t, x) + \epsilon, \quad (6.13)$$

which allows us to use the neural value function to reason about the true BRT and BRS.

6.5.1. Bounding the True Value Function

To derive these bounds, we introduce the modified value functions $\bar{V}$ and $\underline{V}$, which are the unique viscosity solutions under the worst-case realizations of $\hat{\epsilon}$:

$$\bar{V}(t, x) = \inf_{\beta(\cdot) \in \Delta[t, T]} \sup_{u(\cdot) \in \mathcal{M}[t, T]} \left[\int_t^\tau \epsilon_2 ds + g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) \right], \quad (6.14)$$

$$\underline{V}(t, x) = \inf_{\beta(\cdot) \in \Delta[t, T]} \sup_{u(\cdot) \in \mathcal{M}[t, T]} \left[\int_t^\tau -\epsilon_2 ds + g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) \right]. \quad (6.15)$$

Since ϵ_2 is a positive constant, it follows that

$$\underline{V}(t, x) \leq V^*(t, x) \leq \bar{V}(t, x). \quad (6.16)$$

Furthermore, using the inequality $-\epsilon_2(T-t) \leq \int_t^\tau -\epsilon_2 ds$, we obtain:

$$\begin{aligned} & \inf_{\tau \in [t, T]} g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) - \epsilon_2(T-t) \\ & \leq \inf_{\tau \in [t, T]} \left[g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) - \int_t^\tau \epsilon_2 ds \right]. \end{aligned}$$

Similarly, using $\epsilon_2(T-t) \geq \int_t^\tau \epsilon_2 ds$, we obtain:

$$\begin{aligned} & \inf_{\tau \in [t, T]} g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) + \epsilon_2(T-t) \\ & \geq \inf_{\tau \in [t, T]} \left[g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) + \int_t^\tau \epsilon_2 ds \right]. \end{aligned}$$

Thus, we derive the key bound:

$$\bar{V}(t, x) - \epsilon_2(T-t) \leq V^*(t, x) \leq \underline{V}(t, x) + \epsilon_2(T-t). \quad (6.17)$$

6.5.2. Relating Bounds to the Neural Value Function

To relate $\bar{V}$ and $\underline{V}$ to the neural value function, we need to establish that these auxiliary value functions are the unique viscosity solutions of the ϵ -modified HJ-PDE. This result is formalized in the following theorem.

Theorem 10 *The function $\bar{V}(t, x)$ is the unique viscosity solution of the ϵ -modified HJ-PDE:*

$$D_t \bar{V}(t, x) + \min \{ \epsilon_2, \mathcal{H}(t, x, D_x \bar{V}) \} = 0,$$

with the Hamiltonian defined as

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p + \epsilon_2],$$

and the terminal condition

$$\bar{V}(T, x) = g(x).$$

Similarly, $\underline{V}(t, x)$ is the unique viscosity solution of the ϵ -modified HJ-PDE:

$$D_t \underline{V}(t, x) + \min \{-\epsilon_2, \mathcal{H}(t, x, D_x \underline{V})\} = 0,$$

with the Hamiltonian

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p - \epsilon_2],$$

and the terminal condition

$$\underline{V}(T, x) = g(x).$$

The proof of Theorem 10 is given in the appendix.

As a consequence of Theorem 10, we obtain

$$\underline{V}(t, x) \leq V_\theta(t, x) \leq \bar{V}(t, x). \quad (6.18)$$

Finally, incorporating ϵ_1 , which perturbs the boundary condition, we have:

$$V_\theta(T, x) = g(x) \pm \epsilon_1. \quad (6.19)$$

Thus, by combining Equation (6.17) with Equations (6.18) and (6.19), we recover the bound in Equation (6.13).




Figure 6.1: Learner-verifier training framework with curriculum and certification loop.

6.6. CASE STUDIES

Given the page limitations, we prioritize detailing the key steps of the training and verification loop to ensure clarity, and thus omit a full comparison of

different reach/avoid examples. The provided toolbox¹ includes additional examples, such as an evader-pursuer scenario, forward and backward reachability problems, and implementations of both the reachable tube and the reachable set.

Let us consider the simple double integrator example, a canonical second-order control system, $\dot{x}_1 = x_2, \dot{x}_2 = u$, where x_1 represents position, x_2 represents velocity, and u is the control input. We consider the problem of learning the forward reachable set, with the initial set defined by

$$g(x) = x_1^2 + x_2^2 - R^2, \quad (6.20)$$

with $R^2 = 0.5$.

6.6.1. Training of value function

The training procedure, outlined in Algorithm 1 and illustrated in Fig. 6.1 follows a curriculum learning strategy with three phases:

- **Pretraining phase:** Focuses on boundary conditions at the final time $t = T$.
- **Curriculum phase:** Gradually extends the time horizon from $[T, T]$ to $[t_0, T]$.
- **Finetuning phase:** Further refines the model to minimize loss before certification.

Each training epoch involves generating a random batch of samples and computing the loss $\mathcal{L}(t_i, x_i, \theta)$ for each sample $i \in \{1, \dots, N\}$. The formal error bounds on the reachable set can be conservative as the results in Theorem 10 assume the worst-case realisation of ϵ_2 along any trajectory. Consequently, even when the mean training loss is low, the existence of regions of the state space with high network loss can lead to conservative over- or under-approximations of the reachable set. To mitigate this, we not only minimize the mean error over the batch but also the maximum error, leading to the training objective of minimizing:

$$\frac{1}{N} \sum_i \mathcal{L}(t_i, x_i, \theta) + \lambda_{\max} \max_i \mathcal{L}(t_i, x_i, \theta).$$

Here, $\lambda_{\max}$ is set to 10% during the curriculum phase and 30% during the finetuning phase. The finetuning phase uses a patience counter, terminating only when the loss consistently stays below a threshold $\lambda_{\epsilon} \epsilon_2$, with $\lambda_{\epsilon} = 95\%$. This modified approach to the curriculum training scheme presented in [22] reduces the total training time while ensuring sufficient confidence in the model prior to progressing to the certification phase.

To improve performance and enable smaller networks to approximate the value function effectively, we introduce a polynomial layer as the first layer of the network. This allows subsequent layers to operate not only on the original inputs (t, x) but also on polynomial transformations, such as (t^2, x^2) .

¹<https://github.com/nikovert/CARe>

Algorithm 6.1: Training of $V_\theta(t, x)$

```

1: Input: Thresholds  $\epsilon_1$  and  $\epsilon_2$ 
2: Initialize: Model  $V_\theta$ ,  $t_{\text{current}} \leftarrow T$ 
3: while Training do
4:   Generate samples  $x_i$  and  $t_i \in [t_{\text{current}}, T]$ 
5:   Compute model output  $\hat{y} = V_\theta(t_i, x_i)$ 
6:   Compute loss  $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{mean}} + \lambda_{\text{max}} \cdot \mathcal{L}_{\text{max}}$ 
7:   Backpropagate  $\mathcal{L}_{\text{total}}$  and update parameters  $\theta$ 
8:   if  $\mathcal{P}_{\text{pre-training}}$  and  $\max(\mathcal{L}_1) < \epsilon_1$  then
9:     Progress to  $\mathcal{P}_{\text{curriculum}}$ 
10:  else if  $\mathcal{P}_{\text{curriculum}}$  and  $\max(\mathcal{L}_2) < \epsilon_2$  then
11:    Expand time horizon:  $t_{\text{current}} \leftarrow t_{\text{current}} - \Delta_t$ 
12:    if  $t_{\text{current}} = 0$  then
13:      Progress to  $\mathcal{P}_{\text{finetune}}$ 
14:       $p \leftarrow 0$ 
15:    else if  $\mathcal{P}_{\text{finetune}}$  and  $\max(\mathcal{L}_2) < \lambda_{\epsilon} \epsilon_2$  then
16:       $p \leftarrow p + 1$ 
17:      if  $p > 1000$  then
18:        break ▷ Stop training
19: Output: Trained model  $V_\theta$ 

```

6

6.6.2. Formal Verification

To formally verify that the trained network represents an ϵ -accurate value function, we generate a symbolic representation $V_\theta(t, x)$ to instantiate the quantifier-free SMT query described in Section 6.4. To obtain the corresponding derivative terms $D_t V_\theta(t, x)$ and $\mathcal{H}(t, x, D_x V_\theta)$ we apply symbolic differentiation. To further parallelize the SMT call, we decompose the absolute value expressions of Equations (6.11)–(6.12) into separate calls, e.g., we seek x s.t.

$$\begin{aligned}
& x \in \mathcal{X} \wedge V_\theta(T, x) - g(x) > \epsilon_1, \\
& \vee x \in \mathcal{X} \wedge V_\theta(T, x) - g(x) < -\epsilon_1.
\end{aligned}$$

As SMT with nonlinear real arithmetic is undecidable in general, no exact decision procedure can exist for arbitrary HJ-PDEs [33]. However, `dreal` provides a δ -complete decision procedure that allows for user-specified δ deviations in the satisfying assignments [34, 35]. Consequently, while satisfying assignments may be spurious, unsatisfiability carries over to the exact problem. Hence, if the SMT query is unsatisfiable, it formally establishes a global upper bound on the loss function across the entire domain of interest.

Checking SMT queries with nonlinear real arithmetic can be computationally expensive, especially as unsatisfiability constitutes a full proof for validity. In practice, if a counterexample exists, the SMT query typically terminates within a reasonable timeframe.

For illustration, we consider training and verifying the double integrator




Figure 6.2: The over-approximated reachable set is shown in red, while the learned reachable set is depicted in blue. The initial set at time $t_0 = 0$ is a circle around the centre of radius 0.25, the time horizon is $T = 1$.

Table 6.1: Training and verification results of the CEGIS loop

Iter.	ϵ	Training (s)	Verif. (s)	Result
1	0.30	836.92	1367	Certified
2	0.27	-	22	Counterexample found
3	0.27	181.27	2735	Certified
4	0.243	-	5537	Certified
5	0.2187	-	237	Counterexample found

system using a single hidden layer with 16 neurons, employing sine activation functions. After training with $\epsilon := (\epsilon_1 + \epsilon_2(T - t_0))$, and $T = 1, t_0 = 0, \epsilon_2 = 0.95\epsilon = 0.285, \epsilon_1 = 0.05\epsilon = 0.015$ we progress to certification, which takes 1367 seconds resulting in no counterexample being found. Subsequently, we attempt to reduce ϵ by 10%. Within 22 seconds, one of the parallelized SMT queries finds a counterexample. Thus, at this stage, our trained neural network satisfies $\mathcal{L}(t, x) < 0.3$ but not $\mathcal{L}(t, x) < 0.27$. Following a Counterexample-Guided Inductive Synthesis (CEGIS) approach, we leverage the identified counterexample to refine our model (Fig. 6.1). Specifically, we re-enter the finetuning phase of training with the reduced ϵ this time ensuring that 10% of our samples are drawn near the counterexample. This refinement improves the model such that a subsequent verification attempt fails to find a new counterexample, thereby establishing a tighter bound on the loss. We continue the CEGIS loop for a total of 5 iterations with the results shown in Table 6.1. The fine-tuning phase of the final iteration is unable to obtain an empirical loss below 0.2187, thus, the best certifiable result with $\epsilon = 0.243$ is used to produce the final over-approximated reachable set, shown in Fig. 6.2.

Our approach is applicable to arbitrary continuous-time reachability problems. However, the use of `dreal` practically limits the network size, as increasing the depth and number of neurons adds complexity to the formulas (6.11)-(6.12). If an empirical result is satisfactory, Theorem 10 can provide an over- or under-approximation of the reachable set, even when only an empirical bound on the loss is available.

6.7. CONCLUSIONS

The primary goal of this paper is to establish a relation between the training loss and the induced level of approximation error on the final BRS/BRT. To this end, we introduce a modified HJ-PDE allowing us to relate a bound on the loss function used during training, to the true solutions of the reachability problem. To obtain the bound on the loss function, we could use an empirical data-driven approach to estimating the bound on the training loss, though this would fail to constitute a formal certificate. Thus we introduce an approach to verification using a Satisfiability Modulo Theory solver, enabling us to provide a formal certificate for approximate reachability.

This work opens several promising directions for future research. One avenue is the extension of the SMT-based verification framework to systems with stochastic dynamics, where uncertainty arises not just from bounded disturbances but also from probabilistic transitions. Another direction involves tightening the conservativeness of the computed certificates, either by improving the expressiveness of the neural network architecture or by integrating more refined numerical bounds on residual errors. Finally, exploring scalability of the proposed method to high-dimensional systems and longer time horizons remains a crucial step toward real-world deployment.

6.8. APPENDIX: CHAPTER 6

For the proofs in the remaining sections, we ease notation by omitting the subscript of ϵ , i.e. $\epsilon_2 = \epsilon$. We introduce the payoff function:

$$\begin{aligned} P(u, d) &= P_{x,t}(u(\cdot), \beta(\cdot)) \\ &= \int_t^T \epsilon ds + g(\phi(T, t, x, u(\cdot), \beta(\cdot))), \end{aligned} \quad (6.21)$$

where $g : \mathbb{R}^m \rightarrow \mathbb{R}$ satisfies the following conditions:

$$|g(x)| \leq C_2, \quad (6.22)$$

$$|g(x) - g(\hat{x})| \leq C_2 |x - \hat{x}|, \quad (6.23)$$

for some constant C_2 and for all $t_0 \leq t \leq T$, $x, \hat{x} \in \mathbb{R}^m$, $u \in \mathcal{U}$, and $d \in \mathcal{D}$. In the differential game setup considered in this paper, the control u is chosen to maximize $P(u, d)$, while the disturbance d is chosen to minimize $P(u, d)$. The term ϵ represents a positive perturbation ($\epsilon \in \mathbb{R}^+$). Since the proofs for $\underline{V}$ follow directly from those of $\bar{V}$, we focus only on the case of $\bar{V}$ and ease notation by using V to denote $\bar{V}$. We ease notation by introducing $\mathbf{x}_{u,\beta}(\tau) := \phi(\tau, t, x, u(\cdot), \beta(\cdot))$.

6

6.8.1. Bellman Optimality

Theorem 11 *The value function defined in Equation (6.14) satisfies the Bellman optimality principle:*

$$\begin{aligned} \bar{V}(t, x) &= \inf_{\beta(\cdot) \in \Delta_{[t, t+\sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t, t+\sigma]}} \min \left\{ \right. \\ &\quad \bar{V}(t + \sigma, \mathbf{x}_{u,\beta}(t + \sigma)) + \int_t^{t+\sigma} \epsilon ds, \\ &\quad \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau \epsilon ds + g(\mathbf{x}_{u,\beta}(\tau)) \right] \right\}. \end{aligned}$$

Proof:

We define

$$\begin{aligned} W(t, x) &= \inf_{\beta(\cdot) \in \Delta_{[t, t+\sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t, t+\sigma]}} \min \left\{ \right. \\ &\quad V(t + \sigma, \mathbf{x}_{u,\beta}(t + \sigma)) + \int_t^{t+\sigma} \epsilon ds, \\ &\quad \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau \epsilon ds + g(\mathbf{x}_{u,\beta}(\tau)) \right] \right\}. \end{aligned}$$

Part I: Lower bound of $W(t, x)$:

Fix $\gamma > 0$, and choose $\beta_1 \in \Delta_{[t, t+\sigma]}$ such that:

$$W(t, x) \geq \sup_{u_1(\cdot) \in \mathcal{M}_{[t, t+\sigma]}} \min \left\{ \begin{aligned} &V(t + \sigma, \mathbf{x}_{u_1, \beta_1}(t + \sigma)) + \int_t^{t+\sigma} \epsilon \, ds, \\ &\inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1, \beta_1}(\tau)) \right] \end{aligned} \right\} - \gamma.$$

Thus, for all $u_1(\cdot) \in \mathcal{M}_{[t, t+\sigma]}$, there exists a $\beta_1 \in \Delta_{[t, t+\sigma]}$ such that:

$$W(t, x) \geq \min \left\{ \begin{aligned} &V(t + \sigma, \mathbf{x}_{u_1, \beta_1}(t + \sigma)) + \int_t^{t+\sigma} \epsilon \, ds, \\ &\inf_{\tau \in [t, t+\sigma]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1, \beta_1}(\tau)) \right] \end{aligned} \right\} - \gamma. \quad (6.24)$$

6

Now, consider the term $V(t + \sigma, x)$:

$$V(t + \sigma, x) = \inf_{\beta(\cdot) \in \Delta_{[t+\sigma, T]}} \sup_{u(\cdot) \in \mathcal{M}_{[t+\sigma, T]}} \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^\tau \epsilon \, ds + g(\phi(\tau, t + \sigma, x, u(\cdot), \beta(\cdot))) \right].$$

Similarly, we can find a $\beta_2(\cdot) \in \Delta_{[t+\sigma, T]}$ and $u_2(\cdot) \in \mathcal{M}_{[t+\sigma, T]}$ such that:

$$V(t + \sigma, x) \geq \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^\tau \epsilon \, ds + g(\phi(\tau, t + \sigma, x, u_2(\cdot), \beta_2(\cdot))) \right] - \gamma.$$

Now, define the combined control policy $u(\cdot) \in \mathcal{M}_{[t, T]}$ and disturbance strategy $\beta(\cdot) \in \Delta_{[t, T]}$ as:

$$u(\tau) = \begin{cases} u_1(\tau) & \text{if } t \leq \tau < t + \sigma, \\ u_2(\tau) & \text{if } t + \sigma \leq \tau \leq T. \end{cases}$$

$$\beta[u](\tau) = \begin{cases} \beta_1[u_1](\tau) & \text{if } t \leq \tau < t + \sigma, \\ \beta_2[u_2](\tau) & \text{if } t + \sigma \leq \tau \leq T. \end{cases}$$

Subsequently, we obtain:

$$\begin{aligned}
 W(t, x) &\geq \min \left\{ \int_t^{t+\sigma} \epsilon \, ds + \right. \\
 &\quad \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^{\tau} \epsilon \, ds + g(\phi(\tau, t+\sigma, x, u_2(\cdot), \beta_2(\cdot))) \right], \\
 &\quad \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\phi(\tau, t, x, u_1(\cdot), \beta_1(\cdot))) \right] \right\} - 2\gamma, \\
 &\geq \min \left\{ \right. \\
 &\quad \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^{\tau} \epsilon \, ds + g(\phi(\tau, t+\sigma, x, u_2(\cdot), \beta_2(\cdot))) \right] \\
 &\quad \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\phi(\tau, t, x, u_1(\cdot), \beta_1(\cdot))) \right] \right\} - 2\gamma, \\
 &= \inf_{\tau \in [t, T]} \left[\int_t^{\tau} \epsilon \, ds + g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) \right] - 2\gamma.
 \end{aligned}$$

This holds true for all $u(\cdot) \in \mathcal{M}_{[t, T]}$, and hence we conclude:

$$W(t, x) \geq \sup_{u(\cdot) \in \mathcal{M}_{[t, T]}} \inf_{\tau \in [t, T]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u, \beta}(\tau)) \right] - 2\gamma.$$

Finally, this gives:

$$W(t, x) + 2\gamma \geq V(t, x).$$

Part II: Upper Bound $W(t, x)$

For all $\beta_1(\cdot) \in \Delta_{[t, t+\sigma]}$, we have:

$$\begin{aligned}
 W(t, x) &\leq \sup_{u(\cdot) \in \mathcal{M}_{[t, t+\sigma]}} \min \left\{ \right. \\
 &\quad V(t+\sigma, \mathbf{x}_{u, \beta_1}(t+\sigma)) + \int_t^{t+\sigma} \epsilon \, ds, \\
 &\quad \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u, \beta_1}(\tau)) \right] \right\}.
 \end{aligned}$$

Then for a fixed $\gamma > 0$, there exists a $u_1 \in \mathcal{M}_{[t, t+\sigma]}$ such that:

$$\begin{aligned}
 W(t, x) &\leq \min \left\{ V(t+\sigma, \mathbf{x}_{u_1, \beta_1}(t+\sigma)) + \int_t^{t+\sigma} \epsilon \, ds, \right. \\
 &\quad \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u_1, \beta_1}(\tau)) \right] \right\} + \gamma.
 \end{aligned}$$

Consider the term

$$V(t + \sigma, \mathbf{x}_{u,\beta}(t + \sigma)) = \inf_{\beta \in \Delta_{[t+\sigma, T]}} \sup_{u(\cdot) \in \mathcal{M}_{[t+\sigma, T]}} \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^{\tau} \epsilon ds + g(\phi(\tau, t + \sigma, \mathbf{x}, u(\cdot), \beta(\cdot))) \right].$$

Then for all $\beta_2 \in \Delta_{[t+\sigma, T]}$, we have:

$$V(t + \sigma, \mathbf{x}) \leq \sup_{u(\cdot) \in \mathcal{M}_{[t+\sigma, T]}} \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^{\tau} \epsilon ds + g(\phi(\tau, t + \sigma, \mathbf{x}, u(\cdot), \beta_2(\cdot))) \right].$$

For a fixed $\gamma > 0$, there exists a $u_2 \in \mathcal{M}_{[t+\sigma, T]}$ such that:

$$V(t + \sigma, \mathbf{x}) \leq \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^{\tau} \epsilon ds + g(\phi(\tau, t + \sigma, \mathbf{x}, u_2(\cdot), \beta_2(\cdot))) \right] + \gamma.$$

6

We define the combined control policy $u(\cdot) \in \mathcal{M}_{[t, T]}$ and disturbance strategy $\beta(\cdot) \in \Delta_{[t, T]}$ as before, such that:

$$\begin{aligned} W(t, \mathbf{x}) &\leq \min \left\{ \right. \\ &\quad \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^{\tau} \epsilon ds + g(\phi(\tau, t + \sigma, \mathbf{x}, u_2(\cdot), \beta_2(\cdot))) \right], \\ &\quad \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^{\tau} \epsilon ds + g(\mathbf{x}_{u_1, \beta_1}(\tau)) \right] \right\} + 2\gamma. \\ &\leq \min \left\{ \right. \\ &\quad \inf_{\tau \in [t+\sigma, T]} \left[\int_t^{\tau} \epsilon ds + g(\phi(\tau, t + \sigma, \mathbf{x}, u_2(\cdot), \beta_2(\cdot))) \right], \\ &\quad \left. \inf_{\tau \in [t, t+\sigma]} \left[\int_t^{\tau} \epsilon ds + g(\mathbf{x}_{u, \beta_1}(\tau)) \right] \right\} + 2\gamma. \\ &\leq \inf_{\tau \in [t, T]} \left[\int_t^{\tau} \epsilon ds + g(\mathbf{x}_{u, \beta}(\tau)) \right] + 2\gamma. \end{aligned} \tag{6.25}$$

It follows from Equation (6.14) that:

$$V(t, \mathbf{x}) = \inf_{\beta \in \Delta_{[t, T]}} \sup_{u \in \mathcal{M}_{[t, T]}} \inf_{\tau \in [t, T]} \left[\int_t^{\tau} \epsilon ds + g(\mathbf{x}_{u, \beta}(\tau)) \right].$$

Thus, there exists $\beta \in \Delta_{[t,T]}$ such that for a given $\gamma > 0$:

$$V(t, x) \geq \sup_{u \in \mathcal{M}_{[t,T]}} \inf_{\tau \in [t,T]} \left[\int_t^\tau \epsilon ds + g(\mathbf{x}_{u,\beta}(\tau)) \right] - \gamma.$$

Thus for all $u \in \mathcal{M}_{[t,T]}$, we get:

$$V(t, x) + \gamma \geq \inf_{\tau \in [t,T]} \left[\int_t^\tau \epsilon ds + g(\mathbf{x}_{u,\beta}(\tau)) \right]. \quad (6.26)$$

Combining this with Equation (6.25), we get:

$$W(t, x) \leq V(t, x) + 3\gamma. \quad (6.27)$$

Finally, since γ is arbitrary, we can let $\gamma \rightarrow 0$ and obtain:

$$\begin{aligned} V(t, x) = & \inf_{\beta(\cdot) \in \Delta_{[t,t+\sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t,t+\sigma]}} \min \left\{ \right. \\ & V(t + \sigma, \mathbf{x}_{u,\beta}(t + \sigma)) + \int_t^{t+\sigma} \epsilon ds, \\ & \left. \inf_{\tau \in [t,t+\sigma]} \left[\int_t^\tau \epsilon ds + g(\mathbf{x}_{u,\beta}(\tau)) \right] \right\}. \end{aligned}$$

6

The proof of Theorem 10 relies on several further properties of the value function, which will be provided in the form of Lemmas 15–18. As in the previous section, we will use V for ease of notation. We begin with the monotonicity of the value function.

6.8.2. Monotonicity of the value function

Lemma 15 For all $(t, x) \in [0, T] \times \mathbb{R}^m$

$$V(t, x) \leq V(t + \sigma, x) + \int_t^{t+\sigma} \epsilon ds$$

and $V(T, x) = g(x)$

Proof:

$V(T, x) = g(x)$ is trivial and can be directly observed from the definition of the value function, Equation (6.14). Consider

$$V(t, x) = \inf_{\beta \in \Delta_{[t,T]}} \sup_{u \in \mathcal{M}_{[t,T]}} \inf_{\tau \in [t,T]} \left[\int_t^\tau \epsilon ds + g(\mathbf{x}_{u,\beta}(\tau)) \right]$$

Let us assume for the sake of contradiction that

$$V(t, x) > V(t + \sigma, x) + \int_t^{t+\sigma} \epsilon \, ds,$$

which can be equivalently stated as

$$\begin{aligned} & \inf_{\beta_1 \in \Delta[t, T]} \sup_{u_1 \in \mathcal{M}[t, T]} \inf_{\tau \in [t, T]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1, \beta_1}(\tau)) \right] \\ & > \inf_{\beta_2 \in \Delta[t+\sigma, T]} \sup_{u_2 \in \mathcal{M}[t+\sigma, T]} \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^\tau \epsilon \, ds \right. \\ & \quad \left. + g(\phi(\tau, t + \sigma, x, u_2(\cdot), \beta_2(\cdot))) \right] + \int_t^{t+\sigma} \epsilon \, ds. \end{aligned}$$

Thus, for all $\beta_1 \in \Delta[t, T]$, there exists a $\beta_2 \in \Delta[t+\sigma, T]$ such that

$$\begin{aligned} & \sup_{u_1 \in \mathcal{M}[t, T]} \inf_{\tau \in [t, T]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1, \beta_1}(\tau)) \right] \\ & > \sup_{u_2 \in \mathcal{M}[t+\sigma, T]} \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^\tau \epsilon \, ds \right. \\ & \quad \left. + g(\phi(\tau, t + \sigma, x, u_2(\cdot), \beta_2(\cdot))) \right] + \int_t^{t+\sigma} \epsilon \, ds. \end{aligned}$$

Furthermore, for all $u_2 \in \mathcal{M}[t+\sigma, T]$, there exists $u_1 \in \mathcal{M}[t, T]$ such that

$$\begin{aligned} & \inf_{\tau \in [t, T]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1, \beta_1}(\tau)) \right] \\ & > \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^\tau \epsilon \, ds + g(\phi(\tau, t + \sigma, x, u_2(\cdot), \beta_2(\cdot))) \right] \\ & \quad + \int_t^{t+\sigma} \epsilon \, ds. \end{aligned}$$

Now let us define:

$$u(\tau) \equiv \begin{cases} u_1(\tau) & \text{if } t \leq \tau < t + \sigma, \\ u_2(\tau) & \text{if } t + \sigma \leq \tau \leq T. \end{cases}$$

and

$$\beta[u](\tau) \equiv \begin{cases} \beta_1[u](\tau) & \text{if } t \leq \tau < t + \sigma, \\ \beta_2[u](\tau) & \text{if } t + \sigma \leq \tau \leq T. \end{cases}$$

Using the uniqueness of the solution of the ODE (6.1), it follows that

$$\begin{aligned}
 & \inf_{\tau \in [t, T]} \left[\int_t^\tau \epsilon ds + g(\mathbf{x}_{u, \beta}(\tau)) \right] \\
 & > \inf_{\tau \in [t+\sigma, T]} \left[\int_{t+\sigma}^\tau \epsilon ds + g(\phi(\tau, t+\sigma, x, u(\cdot), \beta(\cdot))) \right] \\
 & > \inf_{\tau \in [t+\sigma, T]} \left[\int_t^\tau \epsilon ds + g(\phi(\tau, t+\sigma, x, u(\cdot), \beta(\cdot))) \right],
 \end{aligned}$$

which is a contradiction, as the infimum over a larger set is always less than or equal to the infimum over a smaller set.

As an immediate corollary of Lemma 15, we have

Corollary 2 For all $(t, x) \in [0, T] \times \mathbb{R}^m$

$$\begin{aligned}
 & V(\tau, x(\tau)) + \inf_{\beta \in \Delta[t, T]} \sup_{u \in \mathcal{M}[t, T]} \int_t^\tau \epsilon ds \leq \\
 & \inf_{\beta \in \Delta[t, T]} \sup_{u \in \mathcal{M}[t, T]} \left(\int_t^\tau \epsilon ds + g(\mathbf{x}_{u, \beta}(\tau)) \right)
 \end{aligned}$$

6

6.8.3. Existence and Uniqueness

We provide 2 lemmas that show that the value function is bounded and Lipschitz continuous. This is required for the existence and uniqueness guarantees.

Lemma 16 The value function, defined in the Equation (6.14), is bounded i.e. $|V(t, x)| \leq C_4$ Where C_4 is a constant

Proof:

$$P(u(\cdot), \beta(\cdot)) = \int_t^T \epsilon ds + g(\phi(T, t, x, u(\cdot), \beta(\cdot)))$$

Using Equation (6.22) and the fact that ϵ is a constant, it follows that $|P(u(\cdot), \beta(\cdot))| \leq (T-t)\epsilon + C_2$. This holds for all $u(\cdot) \in \mathcal{M}[t, T]$ and $\beta(\cdot) \in \Delta[t, T]$. Thus implying $|V(t, x)| \leq C_4$.

Lemma 17 The value function, defined in Equation (6.14), is Lipschitz continuous i.e. $|V(t, x) - V(\hat{t}, \hat{x})| \leq C_4(|t - \hat{t}| + |x - \hat{x}|)$ for all $0 \leq t \leq T$, $0 \leq \hat{t} \leq T$ and $x, \hat{x} \in \mathbb{R}^m$.

Proof:

We introduce the following notation

$$P_{t,x}(\tau, u, \beta[u]) := \int_t^\tau \epsilon ds + g(\mathbf{x}_{u,\beta}(\tau))$$

Part I

We will show that $V(t_1, x_1) - V(t_2, x_2) \leq \bar{C}(|t_1 - t_2| + |x_1 - x_2|) + 3\gamma$, where $\bar{C}$ is some constant. Without loss of generality we can assume that $t_1 \leq t_2$. Furthermore, for any β_1

$$\begin{aligned} V(t_1, x_1) &= \inf_{\beta \in \Delta_{[t_1, T]}} \sup_{u \in \mathcal{M}_{[t_1, T]}} \inf_{\tau \in [t_1, T]} P_{t_1, x_1}(\tau, u, \beta[u]) \\ &\leq \sup_{u \in \mathcal{M}_{[t_1, T]}} \inf_{\tau \in [t_1, T]} P_{t_1, x_1}(\tau, u, \beta_1[u]) \end{aligned}$$

Subsequently, for a fixed $\gamma > 0$, $\exists u_1 \in \mathcal{M}_{[t_1, T]}$ such that for all $\tau \in [t_1, T]$ we have

$$V(t_1, x_1) \leq P_{t_1, x_1}(\tau, u_1, \beta_1[u_1]) + \gamma \quad (6.28)$$

Similarly, we can show that for a fixed $\gamma > 0$, $\exists \beta_2 \in \Delta_{[t_2, T]}$ such that

$$V(t_2, x_2) \geq \sup_{u \in \mathcal{M}_{[t_2, T]}} \inf_{\tau \in [t_2, T]} P_{t_2, x_2}(\tau, u, \beta_2[u]) - \gamma \quad (6.29)$$

For an arbitrary $u_1 \in \mathcal{M}_{[t_2, T]}$ we have that

$$V(t_2, x_2) \geq \inf_{\tau \in [t_1, T]} P_{t_2, x_2}(\tau, u_1, \tilde{\beta}[u_1]) - \gamma$$

where $\tilde{\beta}[u_1]$ is defined as

$$\tilde{\beta}[u_1](\tau) \equiv \begin{cases} \beta_1[u_1](\tau) & \text{if } t_1 \leq \tau < t_2 \\ \beta_2[u_1](\tau) & \text{if } t_2 \leq \tau \leq T \end{cases}$$

Using the definition of infimum $\exists \tau \in [t_1, T]$ such that

$$\begin{aligned} V(t_2, x_2) &\geq P_{t_2, x_2}(\tau, u_1, \tilde{\beta}[u_1]) - 2\gamma \\ \Rightarrow -V(t_2, x_2) &\leq -P_{t_2, x_2}(\tau, u_1, \tilde{\beta}[u_1]) + 2\gamma \end{aligned} \quad (6.30)$$

Now let $x_1(\cdot)$ be the solution for the time horizon $(t_1 \leq s \leq T)$ of the following ODE

$$\begin{cases} \frac{dx_1}{ds} = f(s, x_1(s), u_1(s), \tilde{\beta}[u_1](s)) \\ x_1(t_1) = x_1 \end{cases}$$

Let $x_2(\cdot)$ be the solution for the time horizon $(t_2 \leq s \leq T)$ of the ODE

$$\begin{cases} \frac{dx_2}{ds} = f(s, x_2(s), u_1(s), \beta_2[u_1](s)) \\ x_2(t_2) = x_2 \end{cases}$$

Then it follows from equation (6.28) and (6.30), that

$$\begin{aligned} V(t_1, x_1) - V(t_2, x_2) &\leq \\ P_{t_1, x_1}(\tau, u_1, \beta_1[u_1]) - P_{t_2, x_2}(\tau, u_1, \tilde{\beta}[u_1]) + 3\gamma \\ &= \int_{t_1}^{\tau} \epsilon ds + g(x_1(\tau)) - \int_{t_2}^{\tau} \epsilon ds - g(x_2(\tau)) + 3\gamma \\ &= g(x_1(\tau)) - g(x_2(\tau)) + 3\gamma + (t_2 - t_1)\epsilon \end{aligned} \quad (6.31)$$

Using Equation (6.23), it follows that

$$|g(x_1(\tau)) - g(x_2(\tau))| \leq C_2 |x_1(\tau) - x_2(\tau)| \quad (6.32)$$

Furthermore, since $\tilde{\beta}[u_1](s) = \beta_2[u_1](s)$ for all $s \in [t_2, T]$, we have that

$$\begin{aligned} |x_1(\tau) - x_2(\tau)| &= \left| x_1 + \int_{t_1}^{\tau} f(s, x_1(s), u_1(s), \tilde{\beta}[u_1](s)) ds \right. \\ &\quad \left. - x_2 - \int_{t_2}^{\tau} f(s, x_2(s), u_1(s), \beta_2[u_1](s)) ds \right| \\ &= \left| x_1 + \int_{t_1}^{t_2} f(s, x_1(s), u_1(s), \tilde{\beta}[u_1](s)) ds - x_2 \right. \\ &\quad \left. + \int_{t_2}^{\tau} f(s, x_1(s), u_1(s), \beta_2[u_1](s)) \right. \\ &\quad \left. - f(s, x_2(s), u_1(s), \beta_2[u_1](s)) ds \right| \\ &\leq |x_1(t_2) - x_2(t_2)| \\ &\quad + \int_{t_2}^{\tau} |f(s, x_1(s), u_1(s), \beta_2[u_1](s)) \\ &\quad - f(s, x_2(s), u_1(s), \beta_2[u_1](s))| ds \\ &\leq |x_1(t_2) - x_2(t_2)| + C_1 \int_{t_2}^{\tau} |x_1(s) - x_2(s)| ds \\ &\leq |x_1(t_2) - x_2(t_2)| e^{(\tau - t_2)C_1} \end{aligned}$$

where the second inequality uses Equations (6.2) and (6.3) and the last inequality is due to the Bellman-Gronwall Lemma [36].

Using Equation (6.2), the following holds for any $u \in \mathcal{U}, d \in \mathcal{D}$

$$|x_1(t_2) - x_1| \leq C_3|t_1 - t_2|, \quad (6.33)$$

$$\begin{aligned} |x_1(t_2) - x_1(t_1)| &= \left| \int_t^{t_2} f(s, x_2(s), u(s), d(s)) ds \right. \\ &\quad \left. - \int_t^{t_1} f(s, x(s), u(s), d(s)) ds \right| \\ &\leq |C_1(t_2 - t) - C_1(t_1 - t)| \\ &\leq C_1|t_1 - t_2|. \end{aligned}$$

Thus we obtain

$$\begin{aligned} |g(x_1(\tau)) - g(x_2(\tau))| &\leq C_2|x_1(\tau) - x_2(\tau)| \\ &\leq C_2e^{(\tau-t_2)C_1}|x_1(t_2) - x_2(t_2)| \\ &\leq C_2e^{(\tau-t_2)C_1}|x_1(t_2) - x_1| \\ &\quad + C_2e^{(\tau-t_2)C_1}|x_1 - x_2| \\ &\leq C_3e^{(\tau-t_2)C_1}|t_1 - t_2| + e^{(\tau-t_2)C_1}|x_1 - x_2| \end{aligned}$$

Subsequently, there exists some constant $\bar{C}$, such that

$$V(t_1, x_1) - V(t_2, x_2) \leq \bar{C}(|t_1 - t_2| + |x_1 - x_2|) + 3\gamma \quad (6.34)$$

Part II

Now we will show that $V(t_2, x_2) - V(t_1, x_1) \leq \bar{C}(|t_1 - t_2| + |x_1 - x_2|) + 3\gamma$. We have that

$$V(t_2, x_2) = \inf_{\beta \in \Delta_{[t_2, T]}} \sup_{u \in \mathcal{M}_{[t_2, T]}} \inf_{\tau \in [t_2, T]} P_{t_2, x_2}(\tau, u, \beta[u])$$

Thus for all $\beta \in \Delta_{[t_2, T]}$.

$$V(t_2, x_2) \leq \sup_{u \in \mathcal{M}(t_2)} \inf_{\tau \in [t_2, T]} P_{t_2, x_2}(\tau, u, \beta[u])$$

Subsequently, for a fixed $\gamma > 0$, $\exists u_2 \in \mathcal{M}_{[t_2, T]}$ such that for all $\tau \in [t_2, T]$ we have

$$V(t_2, x_2) \leq P_{t_2, x_2}(\tau, u_2, \beta[u_1]) + \gamma \quad (6.35)$$

Similarly, we can show that for a fixed $\gamma > 0$, $\exists \beta_1 \in \Delta_{[t_1, T]}$ such that

$$V(t_1, x_1) \geq \sup_{u \in \mathcal{M}_{[t_1, T]}} \inf_{\tau \in [t_1, T]} P_{t_1, x_1}(\tau, u, \beta_1[u]) - \gamma \quad (6.36)$$

For an arbitrary $u_1 \in \mathcal{M}_{[t_2, T]}$ we have that

$$V(t_1, x_1) \geq \inf_{\tau \in [t_1, T]} P_{t_1, x_1}(\tau, u_1, \beta_1[u_1]) - \gamma$$

Using the definition of infimum, $\exists \tau \in [t_1, T]$ such that

$$\begin{aligned} V(t_1, x_1) &\geq P_{t_1, x_1}(\tau, u_1, \beta_1[u_1]) - 2\gamma \\ \Rightarrow -V(t_1, x_1) &\leq -P_{t_1, x_1}(\tau, u_1, \beta_1[u_1]) + 2\gamma \end{aligned} \quad (6.37)$$

Let us define the extended policy

$$\tilde{u}(\tau) \equiv \begin{cases} u_1(\tau) & \text{if } t_1 \leq \tau < t_2 \\ u_2(\tau) & \text{if } t_2 \leq \tau \leq T \end{cases}$$

Now by using Equation (6.35) and (6.37), it follows that

$$\begin{aligned} V(t_2, x_2) - V(t_1, x_1) &\leq \\ &P_{t_2, x_2}(\tau, \tilde{u}, \beta_1[\tilde{u}]) - P_{t_1, x_1}(\tau, \tilde{u}, \beta_1[\tilde{u}]) + 3\gamma \end{aligned} \quad (6.38)$$

Following similar arguments as in Part 1 and since we can choose γ to be arbitrarily small, we obtain $|V(t, x) - V(\hat{t}, \hat{x})| \leq C_4(|t - \hat{t}| + |x - \hat{x}|)$.

Next, let us recall the definition of the modified Hamiltonian

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p + \epsilon]. \quad (6.39)$$

Then the final Lemma required for Theorem 10 is

Lemma 18 *Let $\varphi \in C^1([t_0, T] \times \mathbb{R}^m)$ and $\theta > 0$. Then if φ satisfies*

$$D_t \varphi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0)) \leq -\theta \leq 0. \quad (6.40)$$

then for a small enough $\sigma > 0$, there exists $\beta \in \Delta_{[t_0, t_0 + \sigma]}$ such that for all $u \in \mathcal{M}_{[t_0, t_0 + \sigma]}$

$$\begin{aligned} &\varphi((t_0 + \sigma), \phi(t_0 + \sigma, t_0, x_0, u(\cdot), \beta(\cdot))) - \varphi(t_0, x_0) \\ &\quad + \int_{t_0}^{t_0 + \sigma} \epsilon ds \leq -\frac{\theta \sigma}{2} \end{aligned}$$

Conversely, if φ satisfies

$$D_t \varphi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0)) \geq \theta \geq 0 \quad (6.41)$$

then for a small enough $\sigma > 0$, there exists $u \in \mathcal{M}_{[t_0, t_0 + \sigma]}$ such that for all $\beta \in \Delta_{[t_0, t_0 + \sigma]}$

$$\begin{aligned} &\varphi((t_0 + \sigma), \phi(t_0 + \sigma, t_0, x_0, u(\cdot), \beta(\cdot))) - \varphi(t_0, x_0) \\ &\quad + \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \frac{\theta \sigma}{2}. \end{aligned}$$

Proof:

Part I we define

$$\Lambda(t, x, u, d) = D_t \phi(t, x) + f(t, x, u, d) \cdot D_x \phi(t, x) + \epsilon.$$

Then if $\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \Lambda(t_0, x_0, u, d) \leq -\theta < 0$ then for each $u \in \mathcal{U}$ there exists $d = d(u) \in \mathcal{D}$ such that $\Lambda(t_0, x_0, u, d) \leq -\theta$. Since Λ is uniformly continuous we have

$$\Lambda(t_0, x_0, \tilde{u}, d) \leq \frac{-3\theta}{4}$$

for all $\tilde{u} \in \mathcal{B}(u, r) \cap \mathcal{U}$ and some $r = r(u) > 0$. Since $\mathcal{U}$ is compact there exist finitely many distinct points $u_1, u_2, \dots, u_n \in \mathcal{U}$, $d_1, d_2, \dots, d_n \in \mathcal{D}$ and $r_1, r_2, \dots, r_n > 0$ such that

$$\mathcal{U} \subset \bigcup_{i=1}^n \mathcal{B}(u_i, r_i)$$

and

$$\Lambda(t_0, x_0, \tilde{u}, d_i) \leq \frac{-3\theta}{4} \quad \text{for } \tilde{u} \in \mathcal{B}(u_i, r_i)$$

We define $\beta_1 : \mathcal{U} \rightarrow \mathcal{D}$, setting

$$\beta_1(u) = d_k \text{ if } u \in \mathcal{B}(u_k, r_k) \setminus \bigcup_{i=1}^{k-1} \mathcal{B}(u_i, r_i) \quad (k = 1, \dots, n).$$

Thus $\Lambda(t_0, x_0, u, \beta_1(u)) \leq \frac{-3\theta}{4}$ for all $u \in \mathcal{U}$. Since Λ is uniformly continuous we therefore have for each sufficiently small $\delta > 0$

$$\Lambda(s, x(s), u, \beta_1(u)) \leq \frac{-\theta}{2}$$

for all $u \in \mathcal{U}$, $t_0 \leq s \leq t_0 + \delta$ and any solution $x(\cdot)$ of Equation (6.1) on $(t_0, t_0 + \delta)$ for any $d(\cdot)$, $u(\cdot)$ with initial condition $x(t_0) = x_0$.

Finally we define $\beta \in \Delta_{[t_0, T]}$ in the following way:

$$\beta[u](s) = \beta_1(u(s))$$

for each $u \in \mathcal{M}_{[t_0, T]}$. It then follows that

$$\Lambda(s, x(s), u(s), \beta[u](s)) \leq \frac{-\theta}{2} \quad (t_0 \leq s \leq t_0 + \sigma),$$

for each $u \in \mathcal{M}_{[t_0, T]}$. Notice that

$$\begin{aligned} \varphi(t_0 + \sigma, x(t_0 + \sigma)) &= \varphi(t_0, x_0) \\ &+ \int_{t_0}^{t_0 + \sigma} f(s, x(s), u(s), \beta[u](s)) \cdot D_x \varphi(s, x(s)) \\ &\quad + D_t \varphi(s, x(s)) ds, \end{aligned} \quad (6.42)$$

such that integrating over Λ yields

$$\varphi(t_0 + \sigma, x(t_0 + \sigma)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \leq -\frac{\theta \sigma}{2}$$

Part II Set

$$\Lambda(t, x, u, d) = D_t \varphi(t, x) + f(t, x, u, d) \cdot D_x \varphi(t, x) + \epsilon.$$

Then if

$$\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \Lambda(t_0, x_0, u, d) \geq \theta > 0,$$

there exists a $u^* \in \mathcal{U}$ such that

$$\min_{d \in \mathcal{D}} \Lambda(t_0, x_0, u^*, d) \geq \theta$$

Since Λ is uniformly continuous, we have that

$$\Lambda(s, x(s), u^*, d) \geq \frac{\theta}{2}$$

provided $t_0 \leq s \leq t_0 + \delta$ (for any small $\delta > 0$) and $x(\cdot)$ solves ODE on $(t_0, t_0 + \delta)$ for any $u(\cdot)$, $d(\cdot)$ with initial candidates $x(t_0) = x_0$. Hence for $u(\cdot) \equiv u^*$ and any $\beta \in \Delta_{[t_0, T]}$

$$\begin{aligned} D_t \varphi(s, x(s)) + f(s, x(s), u(s), \beta[u](s)) \cdot D_x \varphi(s, x(s)) \\ + \epsilon \geq \frac{\theta}{2} \end{aligned}$$

Integrating this expression from t_0 to $t_0 + \sigma$ and subtracting $\varphi(t_0, x_0)$ we obtain

$$\varphi((t_0 + \sigma), x(t_0 + \delta)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \frac{\theta \sigma}{2}$$

We are now in a position to provide the proof of Theorem 10

6.8.4. Proof of Theorem 10

Proof:

To prove this it is sufficient to prove the following ([37], [16], [6])

1. For $\varphi(t, x) \in C^1([0, T] \times \mathbb{R}^m)$ such that $V - \varphi$ attains a local maximum at $(t_0, x_0) \in [0, T] \times \mathbb{R}^m$ then

$$D_t \varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0))\} \geq 0$$

2. For $\varphi(t, x) \in C^1([0, T] \times \mathbb{R}^m)$ such that $V - \varphi$ has a local minimum at $(t_0, x_0) \in [0, T] \times \mathbb{R}^m$ then

$$D_t \varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0))\} \leq 0$$

Part 1

Let $\varphi(t, x) \in C^1([0, T] \times \mathbb{R}^m)$ and suppose that $V - \varphi$ attains a local maximum at $(t_0, x_0) \in [0, T] \times \mathbb{R}^m$. Let us assume, for the sake of contradiction, that $\exists \theta > 0$ such that

$$D_t \varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0))\} \leq -\theta \leq 0$$

Case I: $\mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0)) < \epsilon$.

$$D_t \varphi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0)) \leq -\theta$$

According to the Lemma 18, this implies

$$\varphi(t_0 + \sigma, x(t_0 + \sigma)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \leq \frac{-\theta \sigma}{2} \quad (6.43)$$

Since we know that $V - \varphi$ has a maximum at (t_0, x_0) , it follows that $V(t_0, x_0) - \varphi(t_0, x_0) \geq V(t_0 + \sigma, x_0(t_0 + \sigma)) - \varphi(t_0 + \sigma, x_0(t_0 + \sigma))$, such that, using the result of Equation (6.43),

$$V(t_0 + \sigma, x_0(t_0 + \sigma)) + \int_{t_0}^{t_0 + \sigma} \epsilon ds + \frac{\theta \sigma}{2} \leq V(t_0, x_0)$$

But this leads to a contradiction, since we know from Theorem 11 that

$$V(t, x) \leq V(t + \sigma, x(t + \sigma)) + \int_t^{t + \sigma} \epsilon ds.$$

Case II: $\epsilon \leq \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0))$

$$D_t \varphi(t_0, x_0) + \epsilon \leq -\theta$$

Integrating the above from t_0 to $t_0 + \sigma$, it follows that

$$\varphi(t_0 + \sigma, x_0) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \leq -\theta\sigma \quad (6.44)$$

Since we know that $V - \varphi$ has a maximum, following similar argumentation as before,

$$V(t_0 + \sigma, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds + \theta\sigma \leq V(t_0, x_0)$$

But from Lemma 15 we know that

$$V(t, x) \leq V(t + \sigma, x) + \int_{t_0}^{t_0 + \sigma} \epsilon ds.$$

This contradiction implies

$$D_t\varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x\varphi(t_0, x_0))\} \geq 0$$

Part 2

Let $\varphi(t, x) \in C^1([0, T] \times \mathbb{R}^m)$ and suppose that $V - \varphi$ attains a local minimum at $(t_0, x_0) \in [0, T] \times \mathbb{R}^m$. We assume, for the sake of contradiction, that $\exists \theta > 0$ such that

$$D_t\varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x V)\} \geq \theta. \quad (6.45)$$

This implies

$$D_t\varphi(t_0, x_0) + \epsilon \geq \theta$$

such that

$$\varphi(t_0 + \sigma, x(t_0 + \sigma)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \theta\sigma \geq \frac{\theta\sigma}{2}.$$

Furthermore, Equation (6.45) implies

$$D_t\varphi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x V) \geq \theta.$$

such that by Lemma 18, we obtain

$$\varphi((t_0 + \sigma), x(t_0 + \sigma)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \frac{\theta\sigma}{2}$$

Since $V - \varphi$ obtains a local minimum at (t_0, x_0) , i.e.

$$\begin{aligned} V(t_0, x_0) - \varphi(t_0, x_0) &\leq \\ V(t_0 + \sigma, x(t_0 + \sigma)) - \varphi(t_0 + \sigma, x(t_0 + \sigma)). \end{aligned}$$

it follows that

$$V(t_0 + \sigma, x(t_0 + \sigma)) - V(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \frac{\theta\sigma}{2}. \quad (6.46)$$

Following the same logic and we also obtain

$$V(t_0 + \sigma, x_0) - V(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \frac{\theta\sigma}{2}. \quad (6.47)$$

Recall the definition of $V(t, x)$ obtained from Theorem 11. We consider two cases

Case I

$$\begin{aligned} & V(t_0 + \sigma, x(t_0 + \sigma)) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \\ & \leq \inf_{\tau \in [t_0, t_0 + \sigma]} \left[\int_{t_0}^{\tau} \epsilon ds + g(x(\tau)) \right] \end{aligned}$$

then

$$V(t_0, x_0) = V(t_0 + \sigma, x(t_0 + \sigma)) + \int_{t_0}^{t_0 + \sigma} \epsilon ds$$

substituting this in equation (6.46) leads to a contradiction.

Case II

$$\begin{aligned} & \inf_{\tau \in [t_0, t_0 + \sigma]} \left[\int_{t_0}^{\tau} \epsilon ds + g(x(\tau)) \right] \leq \\ & V(t_0 + \sigma, x(t_0 + \sigma)) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \end{aligned}$$

Then

$$\begin{aligned} V(t_0, x_0) &= \inf_{\beta(\cdot) \in \Delta[t_0, t_0 + \sigma]} \sup_{v(\cdot) \in \mathcal{M}[t_0, t_0 + \sigma]} \\ & \inf_{\tau \in [t_0, t_0 + \sigma]} \left[\int_{t_0}^{\tau} \epsilon ds + g(x(\tau)) \right] \end{aligned}$$

The minima in τ must occur at $\tau = t_0$ and the minimizer is unique. If this was not true then there would exist some $\tau \in [t_0, t_0 + \sigma]$ such that

$$\begin{aligned} V(t_0, x_0) &= \inf_{\beta(\cdot) \in \Delta[t_0, t_0 + \sigma]} \sup_{u(\cdot) \in \mathcal{M}[t_0, t_0 + \sigma]} \\ & \left[\int_{t_0}^{\tau} \epsilon ds + g(x(\tau)) \right] \end{aligned}$$

Thus using corollary 2 we get that

$$V(\tau, x(\tau)) + \int_{t_0}^{\tau} \epsilon ds \leq V(t_0, x_0), \quad \tau \in [t_0, t_0 + \sigma]$$

This will lead to contradiction with equation (6.46). Furthermore, we know that for all $\tau \in [t_0, t_0 + \sigma]$ using Lemma 15

$$V(t_0, x_0) \leq V(t_0 + \sigma, x_0) + \inf_{\beta(\cdot) \in \Delta_{[t_0, t_0 + \sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t_0, t_0 + \sigma]}} \int_{t_0}^{t_0 + \sigma} \epsilon ds \leq g(x_0)$$

Therefore we have

$$V(t_0, x_0) = V(t_0 + \sigma, x_0) + \inf_{\beta(\cdot) \in \Delta_{[t_0, t_0 + \sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t_0, t_0 + \sigma]}} \int_{t_0}^{t_0 + \sigma} \epsilon ds = g(x_0)$$

This along with equation (6.47) leads to contradiction.

REFERENCES

- [1] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. “Multi-objective minimum time optimal control for low-thrust trajectory design”. In: *2021 European Control Conference (ECC)*. 2021, pp. 1975–1980. DOI: [10.23919/ECC54610.2021.9654919](https://doi.org/10.23919/ECC54610.2021.9654919).
- [2] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. “Multi-objective low-thrust spacecraft trajectory design using reachability analysis”. In: *European Journal of Control* 69 (2023), p. 100758. ISSN: 0947-3580. DOI: [10.1016/j.ejcon.2022.100758](https://doi.org/10.1016/j.ejcon.2022.100758).
- [3] C. Livadas and N. A. Lynch. “Formal verification of safety-critical hybrid systems”. In: *International Workshop on Hybrid Systems: Computation and Control*. Springer. 1998, pp. 253–272.
- [4] J. Lygeros, D. N. Godbole, and S. Sastry. “Verified hybrid controllers for automated vehicles”. In: *IEEE transactions on automatic control* 43.4 (1998), pp. 522–539.
- [5] C. Livadas, J. Lygeros, and N. A. Lynch. “High-level modeling and analysis of the traffic alert and collision avoidance system (TCAS)”. In: *Proceedings of the IEEE* 88.7 (2000), pp. 926–948.
- [6] C. Tomlin, I. Mitchell, and R. Ghosh. “Safety verification of conflict resolution manoeuvres”. In: *IEEE Transactions on Intelligent Transportation Systems* 2.2 (2001), pp. 110–120.
- [7] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. “Safety-Aware Hybrid Control of Airborne Wind Energy Systems”. In: *Journal of Guidance, Control, and Dynamics* 47.2 (2024), pp. 326–338. DOI: [10.2514/1.G007254](https://doi.org/10.2514/1.G007254).
- [8] J. Lygeros. “On reachability and minimum cost optimal control”. In: *Automatica* 40.6 (2004), pp. 917–927.
- [9] J. Lygeros, C. Tomlin, and S. Sastry. “Controllers for reachability specifications for hybrid systems”. In: *Automatica* 35.3 (1999), pp. 349–370.
- [10] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre. *Viability theory: new directions*. Springer Science & Business Media, 2011.
- [11] J.-P. Aubin and H. Frankowska. “Viability kernel of control systems”. In: *Nonlinear synthesis*. Springer, 1991, pp. 12–33.
- [12] P. Cardaliaguet. “A differential game with two players and one target”. In: *SIAM Journal on Control and Optimization* 34.4 (1996), pp. 1441–1460.

- [13] P. Cardaliaguet, M. Quincampoix, and P. Saint-Pierre. “Differential games through viability theory: Old and recent results”. In: *Advances in dynamic game theory*. Springer, 2007, pp. 3–35.
- [14] E. Asarin, O. Bournez, T. Dang, O. Maler, and A. Pnueli. “Effective synthesis of switching controllers for linear systems”. In: *Proceedings of the IEEE* 88.7 (2000), pp. 1011–1025.
- [15] I. M. Mitchell *et al.* “A toolbox of level set methods”. In: *UBC Department of Computer Science Technical Report TR-2007-11 1* (2007), p. 6.
- [16] I. Mitchell, A. M. Bayen, and C. J. Tomlin. “Validating a Hamilton-Jacobi approximation to hybrid system reachable sets”. In: *International workshop on hybrid systems: Computation and control*. Springer. 2001, pp. 418–432.
- [17] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-jacobi reachability: A brief overview and recent advances”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 2242–2253.
- [18] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin. “Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking”. In: *IEEE Transactions on Automatic Control* 66.12 (2021), pp. 5861–5876.
- [19] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin. “Decomposition of reachable sets and tubes for a class of nonlinear systems”. In: *IEEE Transactions on Automatic Control* 63.11 (2018), pp. 3675–3688.
- [20] J. N. Maidens, S. Kaynama, I. M. Mitchell, M. M. Oishi, and G. A. Dumont. “Lagrangian methods for approximating the viability kernel in high-dimensional systems”. In: *Automatica* 49.7 (2013), pp. 2017–2029.
- [21] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. “Verification of safety critical control policies using kernel methods”. In: *2022 European Control Conference (ECC)*. 2022, pp. 1870–1875. DOI: [10.23919/ECC55457.2022.9838224](https://doi.org/10.23919/ECC55457.2022.9838224).
- [22] S. Bansal and C. J. Tomlin. “Deepreach: A deep learning approach to high-dimensional reachability”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 1817–1824.
- [23] A. Lin and S. Bansal. “Generating Formal Safety Assurances for High-Dimensional Reachability”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2023, pp. 10525–10531. DOI: [10.1109/icra48891.2023.10160600](https://doi.org/10.1109/icra48891.2023.10160600). URL: <http://dx.doi.org/10.1109/icra48891.2023.10160600>.
- [24] A. Lin and S. Bansal. “Verification of neural reachable tubes via scenario optimization and conformal prediction”. In: *Proceedings of the 6th Annual Learning for Dynamics and Control Conference*. Vol. 242. Proceedings of Machine Learning Research. PMLR, 15–17 Jul 2024, pp. 719–731.

- [25] M. Tayal, A. Singh, S. Kolathaya, and S. Bansal. “A Physics-Informed Machine Learning Framework for Safe and Optimal Control of Autonomous Systems”. In: *arXiv preprint: 2502.11057* (2025). DOI: [10.48550/arXiv.2502.11057](https://doi.org/10.48550/arXiv.2502.11057). arXiv: [2502.11057](https://arxiv.org/abs/2502.11057) [cs.R0].
- [26] K. Margellos and J. Lygeros. “Hamilton–Jacobi formulation for reach–avoid differential games”. In: *IEEE Transactions on automatic control* 56.8 (2011), pp. 1849–1861.
- [27] L. C. Evans and P. E. Souganidis. “Differential games and representation formulas for solutions of Hamilton–Jacobi–Isaacs equations”. In: *Indiana University mathematics journal* 33.5 (1984), pp. 773–797.
- [28] C. J. Tomlin, J. Lygeros, and S. S. Sastry. “A game theoretic approach to controller design for hybrid systems”. In: *Proceedings of the IEEE* 88.7 (2000), pp. 949–970.
- [29] M. Raissi, P. Perdikaris, and G. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378 (Feb. 2019), pp. 686–707. ISSN: 0021-9991. DOI: [10.1016/j.jcp.2018.10.045](https://doi.org/10.1016/j.jcp.2018.10.045).
- [30] A. Abate, C. David, P. Kesseli, D. Kroening, and E. Polgreen. “Counterexample Guided Inductive Synthesis Modulo Theories”. In: *Computer Aided Verification*. Springer International Publishing, 2018, pp. 270–288. ISBN: 9783319961453. DOI: [10.1007/978-3-319-96145-3_15](https://doi.org/10.1007/978-3-319-96145-3_15). URL: http://dx.doi.org/10.1007/978-3-319-96145-3_15.
- [31] M. Fränzle, C. Herde, T. Teige, S. Ratschan, and T. Schubert. “Efficient Solving of Large Non-linear Arithmetic Constraint Systems with Complex Boolean Structure1”. In: *Journal on Satisfiability, Boolean Modeling and Computation* 1.3–4 (May 2007). Ed. by Y. Hamadi and L. Bordeaux, pp. 209–236. ISSN: 1574-0617. DOI: [10.3233/sat190012](https://doi.org/10.3233/sat190012). URL: <http://dx.doi.org/10.3233/sat190012>.
- [32] S. Gao, S. Kong, and E. M. Clarke. “dReal: An SMT Solver for Nonlinear Theories over the Reals”. In: *Automated Deduction – CADE-24*. Springer Berlin Heidelberg, 2013, pp. 208–214. ISBN: 9783642385742. DOI: [10.1007/978-3-642-38574-2_14](https://doi.org/10.1007/978-3-642-38574-2_14). URL: http://dx.doi.org/10.1007/978-3-642-38574-2_14.
- [33] D. Kroening and O. Strichman. *Decision Procedures - An Algorithmic Point of View*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2008.
- [34] S. Gao, J. Avigad, and E. M. Clarke. “ δ -Complete Decision Procedures for Satisfiability over the Reals”. In: *IJCAR*. Vol. 7364. Lecture Notes in Computer Science. Springer, 2012, pp. 286–300.
- [35] S. Gao, S. Kong, and E. M. Clarke. “dReal: An SMT solver for nonlinear theories over the reals”. In: *International conference on automated deduction*. Springer. 2013, pp. 208–214.

- [36] S. Sastry. *Nonlinear Systems*. Springer New York, 1999. ISBN: 9781475731088. DOI: 10.1007/978-1-4757-3108-8. URL: <http://dx.doi.org/10.1007/978-1-4757-3108-8>.
- [37] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert. “Robust control barrier–value functions for safety-critical control”. In: *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE. 2021, pp. 6814–6821.

7

THE STOCHASTIC AVOID PROBLEM WITH COMPUTATION

In safety-critical systems, concepts such as invariance and reachability play pivotal roles in ensuring safe operation regions. Deterministic or robust settings typically consider forced backward reachable sets, defining regions from which a system is inevitably driven into an avoid set regardless of control actions and invariant sets as the region in which system has the capability to stay for the whole time horizon. Extending this to stochastic domains, where disturbances are considered in a stochastic setting, introduces the concept of probabilistic forced backward reachable sets. This paper explores the relationship between the stochastic counterpart of forced backward reachable sets and invariant sets, addressing challenges in tractable algorithm development and convergence guarantees. By iteratively solving optimization problems, the paper provides algorithms for discrete and continuous systems to compute the probabilistic forced backward reachable set. In the paper we both the theoretical proofs that the developed algorithms indeed find the probabilistic forced backward reachable sets and the computational expense of the algorithms. The proposed framework enhances applicability in scenarios such as obstacle avoidance and air traffic control, showcasing its utility through illustrative examples. Through rigorous analysis and algorithmic development, this work contributes to advancing safety assurance methodologies in stochastic systems.

7.1. INTRODUCTION

Concepts such as invariance [1],[2],[3] and reachability [4],[5] play a central role in safety assurance and in the identification of safe operation regions [6],[7]. In a deterministic (no disturbance) or robust (for all possible realisations of disturbances) settings, a forced backward reachable set is a region in space from where the system will be driven to a particular set (usually called avoid set) irrespective of which control action is taken. The stochastic counterpart of such a set would be a region in space from where the probability of the system to be driven into a particular set is higher than a particular probability value.

Thorough research has been conducted on the forced backward reachable set [8],[9] and the control invariant set [10],[11] in deterministic and robust settings. These sets in robust settings are inherently more conservative, thus the stochastic counterpart of these sets reduces the conservatism by allowing some stochastic violations. This furthermore enhances the applicability of the forced backward reachable set for cases with unbounded disturbance. The paper is motivated by problems of reliable control [12],[13], stochastic MPC [14],[15], safety-critical control [16] and related applications such as path planning [17], obstacle avoidance [18] and air traffic control [19]. In this article we call the stochastic counterpart of the forced backward reachable set the 'probabilistic forced backward reachable set'.

The main problem this paper addresses is that, given a set of states $\mathcal{L}$, a time horizon T and a probability value $\mathcal{J} \in [0, 1]$, which state should the system start at such that despite all the available control actions the system enters $\mathcal{L}$ with a probability of at least $\mathcal{J}$. The problem of obtaining the forced backward reachable set is solved by converting the problem into an optimisation problem using dynamic programming and solving it backwards in time in an iterative manner for all possible disturbances [8]. Thus, for obtaining the probabilistic forced backward reachable set we follow a similar idea, but replace the forced backward reachable set with the probabilistic forced backward reachable set. The main challenges for achieving this are firstly to develop a tractable algorithm especially for continuous spaces to obtain the probabilistic forced backward reachable set, and secondly to guarantee convergence of the set in both space and time.

There are some stochastic extensions of sets similar to forced backward reachable sets such as the stochastic reach-avoid problem [20], in which the system has certain target states and avoid states. The solution of the stochastic reach-avoid problem is the set of initial states from which the system can reach the target states while avoiding the avoid states. This extension is not directly applicable to problems where the target set is unavailable or is constantly changing, and only the avoid set is known, i.e., the problem of obstacle avoidance. Another method developed in the context of a stochastic reach avoid problem includes methods based on stochastic contingent sets [21]. Geared towards a similar objective of almost-certain reachability, [22] introduced a new class of stochastic target problems and characterized their solution using dynamic programming. The authors of [22] also examined the differential properties

of the almost-certain reachable set based on a geometric partial differential equation analogous to the HJB equation [23] in that context. However, while almost-certain versions of reachability specifications are intriguing, they may be overly stringent in some applications, particularly when the goal is simply to control the probability of undesirable events occurring.

Other work related to safety assurance involves stochastic extension of viable sets. A viable set is a region in state space where the system has the ability to remain in the set using an available control action. The stochastic counterpart of such a set is a region in space where the system has an ability to remain in the set with a probability higher than a particular probability value. In [24] a target set is used to obtain a stabilisation value, and in [25] a definition of the stochastic control invariant set is provided. [26] is an application of [25] for the financial system. Similar to [25], [27] provides a definition of stochastic invariance for one-step model predictive control. [28] and [29] provide a definition of stochastic invariance for linear systems without input, the papers explore the polyhedral probabilistic invariant set approximations using Chebyshev's inequality. There exists a strong connection between the stochastic avoid problem and the stochastic control invariant set thus this paper is focused on the stochastic counterpart of the forced backward reachable set and its relationship with the stochastic counterpart of the control invariant set.

7.1.1. Contribution and organisation

The contributions of this paper are as follows: Firstly, we extend the definition of the forced backward reachable set from the robust domain to the stochastic domain: Given a designated set $\mathcal{L} \in \mathcal{B}(\mathbb{X})$, a horizon $T \in \mathbb{N}$, and a probability value $\mathcal{J} \in [0, 1]$, the set of all states ($x \in \mathbb{X}$) such that for all Markov control policies ($u \in \mathcal{M}$) there exists a step ($\tau \in [0, T]$) such that the value of the probability measure of the state to transition into the designated set at step τ is greater than or equal to the given probability value. *Note: We assume that the time step $\tau \in \mathbb{N}$.* This definition is appropriate if we do not have any target state and only obstacle avoidance is the main objective. The definition is applicable to discrete time stochastic systems. We then provide the fundamental properties such as measurability along with the proofs, and we also prove the relationship between probabilistic forced backward reachable sets and stochastic invariant sets. We furthermore provide and prove a sufficient condition for the convergence in time of a probabilistic forced backward reachable set of $\mathcal{J}$ probability. On the computational side, we provide an algorithm along with the cost of computation for discrete space systems that computes the exact probabilistic forced backward reachable set. We also provide an extension of this algorithm for continuous space systems and prove that the set obtained via this algorithm, while an estimation, converges to the probabilistic forced backward reachable set. The extension involves a discretisation of the state space. We furthermore provide the error bounds that are introduced due to the discretisation.

The paper is organised as follows. In Section 7.1 we introduce the problem statement and the need for the avoid set motivating the paper. In Section 7.2 we

set up the problem statement and some preliminaries. In Section 7.3 we develop an algorithm to obtain the probable forced backward reachable set for both discrete and continuous space systems. In 7.4 we prove the sufficient condition for the probable forced backward reachable set to exist and its relationship with the invariant set. In Section 7.5 we provide a toy example to demonstrate the algorithms and finally we conclude the paper and give future recommendations in Section 7.6.

7.2. SETTING OF THE PROBLEM STATEMENT

In our study, we consider a stochastic Markov control system denoted as $S = (\mathbb{X} = \mathbb{R}^n, \mathbb{U} = \mathbb{R}^m, \mathcal{P})$, where:

- $\mathbb{X}$ represents the state space, equipped with the corresponding Borel σ -algebra $\mathcal{B}(\mathbb{X})$.
- $\mathbb{U}$ is the control space, accompanied by the Borel σ -algebra $\mathcal{B}(\mathbb{U})$.
- $\mathcal{P}$ is a Borel-measurable stochastic kernel. Given $\mathbb{X} \times \mathbb{U}$, it assigns to each pair $(\mathbf{x}, u)$ a probability measure on the Borel space $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$ where $\mathbf{x} \in \mathbb{X}$ and $u \in \mathbb{U}$. Formally, probability measure on the Borel space $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$: $\mathcal{P}(\cdot | \mathbf{x}, u)$ can be defined as a function: $\mathcal{P} : \mathcal{B}(\mathbb{X}) \times \mathbb{X} \times \mathbb{U} \rightarrow [0, 1]$
- $n, m \in \mathbb{N}$ where $\mathbb{N}$ is the set of natural numbers.

We assume that the $\mathbb{U}_{\mathbf{x}}$ is a non-empty subset of $\mathbb{U}$ of all the control actions for each $\mathbf{x} \in \mathbb{X}$. Given a finite horizon $T \in \mathbb{N}$, a control policy is said to be a Markov control policy if it only depends on the current state, i.e., $u_k(\mathbf{x}(k)) = u_k(\mathbf{x}(k), \mathbf{x}(k-1), \dots), k \in \mathbb{N}$. For brevity $u_k(\mathbf{x}(k))$ will be represented as u .

Definition 1 (Markov control policy): A Markov control policy u for a given system S is a sequence $u = (u_1, u_2, \dots, u_{N-1})$ of universally measurable maps, i.e.:

$$u : \mathbb{N} \times \mathbb{X} \rightarrow \mathbb{U}, \quad \forall k \in \mathbb{N}$$

Note 1: A subset $\mathcal{A}$ of a space $\mathbb{Z}$ is universally measurable if the set $\mathcal{A}$ is measurable with respect to all the complete measures on the space $\mathbb{Z}$, that measure all Borel sets of space $\mathbb{Z}$ ($\mathcal{B}(\mathbb{Z})$). A function $f : \mathbb{Z} \rightarrow \mathbb{Y}$ is universally measurable if the inverse image ($f^{-1}(\mathcal{A})$) of all the sets $\mathcal{A} \in \mathcal{B}(\mathbb{Y})$ is universally measurable. For showing the existence of a solution to a stochastic optimal problem the condition of universal measurability is a weaker condition than the Borel measurability [30],[31],[32].

Note 2: In a substantial subset of stochastic optimal control problems, the utilization of Markov policies proves to be adequate in delineating the optimal policy [31]. Given that the adoption of a randomized Markov policy does not contribute to the increase in the maximum probability that the state stays within a designated set or reaches a particular set, our focus in this paper will be on deterministic Markov policies. We denote the set of all Markov policies by $\mathcal{M}$.

Given a control policy $u \in \mathcal{M}$, a horizon τ ($\tau \in \mathbb{N}$) and a state $x \in \mathbb{X}$, the probability represented as $\bar{\mathbb{P}}^\tau(x_\tau|x, u)$ is read as the probability to transition to state x_τ at step τ , given that the state at step 0 is x , under the control policy u , i.e. $\bar{\mathbb{P}}(X_\tau = x_\tau | X_0 = x, u)$, $\tau \in \mathbb{N}$. Given a control policy $u \in \mathcal{M}$, a time step $\tau \in \mathbb{N}$, a state $x \in \mathbb{X}$ and a set $A \in \mathcal{B}(\mathbb{X})$, the probability $\bar{\mathbb{P}}^\tau(x_\tau \in A|x, u)$ is read as the probability that the state x_τ is in set A , given that the state at step 0 is x , under the control policy u , i.e. $\bar{\mathbb{P}}(X_\tau = x_\tau \in A | X_0 = x, u)$, $\tau \in \mathbb{N}$.

Let us define an indicator function $f_A : \mathbb{X} \rightarrow \{0, 1\}$ of a Borel set $A \in \mathcal{B}(\mathbb{X})$, as shown in equation 7.1.

$$f_A(x) = \begin{cases} 1, & \text{if } x \in A \\ 0, & \text{if } x \notin A \end{cases} \quad (7.1)$$

Using the indicator function (equation 7.1) and following the dynamic programming principle as applied in [33], we can define a value function $V_{\tau, \mathcal{L}} : \mathbb{X} \rightarrow [0, 1]$, $\tau \in [0, T] \subseteq \mathbb{N}$ by the process of backward recursion as shown in equation 7.2, where $\bar{\mathcal{L}}$ is the compliment of set $\mathcal{L}$ in $\mathbb{X}$.

$$V_{\tau, \mathcal{L}}(x) = f_{\mathcal{L}}(x) + \inf_{u \in \mathcal{M}} f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{\tau+1, \mathcal{L}}(x') \mathcal{P}(dx'|x, u), x \in \mathbb{X} \quad (7.2)$$

This value function at $\tau = 0$ provides the probability distribution as $\inf_u \sup_t \bar{\mathbb{P}}^t(x' \in \mathcal{L}|x, u)$. The value function is initialised as $V_{T, \mathcal{L}}(x) = f_{\mathcal{L}}(x)$. The proof of above statement can be found in section 7.7.

Assumption 8 The set $\mathbb{U}_x(\tau, \mathcal{J})$ given by equation 7.3

$$\mathbb{U}_x(\tau, \mathcal{J}) = \{u \in \mathbb{U} \mid \int_{\mathbb{X}} V_{\tau+1, \mathcal{L}}(x') \mathcal{P}(dx'|x, u) \geq \mathcal{J}\} \quad (7.3)$$

is compact for all $x \in \mathbb{X}$, $\tau \in \mathbb{N}_{0, T}$, $\mathcal{J} \in [0, 1]$.

Thus if assumption 1 holds, then $\forall x \in \mathbb{X}$, the existence of a Markov policy that satisfies equation 7.2 is guaranteed and is given by equation 7.4.

$$u_k^*(x) = \arg \inf_{u \in \mathcal{M}} f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{\tau+1, \mathcal{L}}(x') \mathcal{P}(dx'|x, u), \quad x \in \mathbb{X} \setminus \mathcal{L}, k \in \mathbb{N}_{[0, T-1]} \quad (7.4)$$

7.3. AVOID SET PROPERTIES AND COMPUTATION

In this section we firstly prove that the probabilistic forced backward reachable set is universally measurable. After that we show that the value function used for obtaining the probabilistic forced backward reachable set converges. The convergence of the value function is used to obtain the termination time of the algorithm that computes the probabilistic forced backward reachable set. We

then introduce the algorithm for computing this set, along with its discrete space implementation and the associated computational cost. To extend these results to the continuous space, we propose a discretisation procedure for which we prove that the approximate value function used for obtaining the approximate probabilistic forced backward reachable set for the discretised space converges to the value function for the continuous space as the grid size approaches zero. This also proves that the approximate probabilistic avoid backward reachable set converges to the avoid backward reachable set for the continuous space as the grid size diminishes to zero.

Given a designated set $\mathcal{L} \in \mathcal{B}(\mathbb{X})$, a horizon $T \in \mathbb{N}$, and a probability value $\mathcal{J} \in [0, 1]$, the probable forced backward reachable set is the set of all states ($x \in \mathbb{X}$) such that, for all Markov control policies ($u \in \mathcal{M}$), there exists a step ($\tau \in [0, T]$) such that the value of the probability measure of the state to transition into the designated set at step τ is greater than or equal to the given probability value. We denote this set as $V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J})$. Mathematically the set is given by the equation 7.5

$$\begin{aligned} V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J}) &= \{x : \forall u \in \mathcal{M}, \exists \tau \in \mathbb{N}_{[0, T]}, \bar{\mathbb{P}}^\tau(x_\tau \in \mathcal{L} | x, u) \geq \mathcal{J}\} \\ &= \{x : \inf_u \sup_\tau \bar{\mathbb{P}}^\tau(x_\tau \in \mathcal{L} | x, u) \geq \mathcal{J}, u \in \mathcal{M}, \tau \in \mathbb{N}_{[0, T]}\} \end{aligned} \quad (7.5)$$

Let S be a set of all the singleton $x \in \mathbb{X}$ as shown in equation 7.6.

$$S = \{\{x\} : x \in \mathbb{X}\} \quad (7.6)$$

It is clear from the definition of the probable backward reachable set (equation 7.5) that it is contained in set $\sigma(S)$, where $\sigma(S)$ is the sigma algebra generated by S .

Lemma 19 *The $\sigma(S)$ is contained in the $\mathcal{B}(\mathbb{X})$ where $\sigma(S)$ is the sigma algebra generated by S (set of all singletons) and $\mathcal{B}(\mathbb{X})$ is the Borel sigma algebra generated by $\mathbb{X}$ ($\mathbb{X} = \mathbb{R}^n$).*

Proof : Let set $\mathcal{F}$ be defined as follows:

$$\mathcal{F} = \{F \subseteq \mathbb{X} : F \text{ is countable or } F^c \text{ is countable}\}$$

$$\text{Where } F^c \text{ is the compliment of } F \quad (7.7)$$

It can be trivially shown that $\mathcal{F}$ is a σ -algebra on $\mathbb{R}^n$. Furthermore, $\{x\} \in \mathcal{F}$ for every $x \in \mathbb{R}^n$, thus $\sigma(S) \subseteq \mathcal{F}$ as $\mathcal{F}$ is the sigma algebra that consists of singletons and $\sigma(S)$ is the smallest σ -algebra containing singletons. On the other hand if F is countable, then F can be expressed as the countable union of singletons implying $F \in \sigma(S)$, while if F^c is countable we can use the same logic to show $F^c \in \sigma(S)$. Since sigma algebra are closed under compliment, this implies $F \in \sigma(S)$ thus $\mathcal{F} \subseteq \sigma(S)$. Thus $\mathcal{F} = \sigma(S)$.

Intuitively it can be understood that since singletons are included in the Borel

set, the generated Borel σ -algebra must also contain the smallest sigma-algebra generated by the set of all singletons. Using Lemma 19 we can say that $V_{\text{PFBRs}}([T], \mathcal{L}, \mathcal{J}) \in \sigma(S) \subset \mathcal{B}(\mathbb{R}^n)$.

Note: The sigma algebra $\sigma(S)$ is a proper subset of $\mathcal{B}(\mathbb{R}^n)$. This can be easily shown as: $\sigma(S) \subset \mathcal{B}(\mathbb{R}^n)$ using Lemma 19 and $(0, 1)^n \in \mathbb{R}^n$ but $(0, 1)^n \notin \mathcal{F} = \sigma(S)$.

Lemma 20 For $S \in \mathcal{B}(\mathbb{X})$ the set $V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J}) \subseteq S$ is universally measurable.

Proof : If $V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J}) = S$, then it implies $V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J}) \in \mathcal{B}(\mathbb{X})$ and thus $V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J})$ is also measurable. If $V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J}) \subset S$ then the proof is as follows:

Let function $J_{\tau, \mathcal{L}} : \mathbb{X} \rightarrow \mathbb{R}$, $\tau \in \mathbb{N}_{[0, T]}$ be defined in the following manner:

$$J_{\tau, \mathcal{L}}(x) = V_{\tau, \mathcal{L}}(x), \forall x \in \mathbb{X} \quad (7.8)$$

Following the same logic as in [30] it can be shown that $J_{\tau, \mathcal{L}}(x)$ is lower semi-analytic on any set in $\mathcal{B}(\mathbb{X})$. Using definitions 7.20 and 7.21 [31], the function $J_{\tau, \mathcal{L}}(x)$ is also universally measurable for any set in $\mathcal{B}(\mathbb{X})$. Furthermore using Lemma 19, $J_{\tau, \mathcal{L}}(x)$ is also universally measurable for any set in $\sigma(S)$. Thus the function $J_{\tau, \mathcal{L}}(x)$ adheres to equation 7.9.

$$J_{\tau, \mathcal{L}}^{-1}(\mathcal{I}) = \{x \in \mathbb{X} \mid J_{\tau, \mathcal{L}}(x) \in \mathcal{I}\} \quad (7.9)$$

Where $\mathcal{I} \in \mathcal{B}(\mathbb{R})$ and the set $J_{\tau, \mathcal{L}}^{-1}(\mathcal{I})$ is universally measurable. From the definition of the probable forced backward reachable set we have

$$\begin{aligned} V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J}) &= \{x \in S : \inf_u \sup_{\tau} \tilde{\mathbb{P}}^{\tau}(x_{\tau} \in \mathcal{L} \mid x, u) \geq \mathcal{J}, u(\cdot) \in \mathcal{M}, \tau \in \mathbb{N}_{[0, T]}\} \\ &= \{x : V_{0, \mathcal{L}}(x) \geq \mathcal{J}\} \end{aligned} \quad (7.10)$$

$$V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J}) = \{x \in S : 1 \geq J_{T, \mathcal{L}}^{-1}(\mathcal{I}) \geq \mathcal{J}\} \quad (7.11)$$

$$V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J}) = J_{T, \mathcal{L}}^{-1}(\mathcal{I}) \quad (7.12)$$

Thus $V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J})$ is universally measurable.

The following preposition aims to show that the probable forced backward reachable set can only grow as we increase the time horizon.

Proposition 1 For $T_1, T_2 \in \mathbb{N}$ with $T_2 \geq T_1$ the set $V_{\text{PFBRs}}(T_1, \mathcal{L}, \mathcal{J}) \subseteq V_{\text{PFBRs}}(T_2, \mathcal{L}, \mathcal{J})$

Proof : The proof directly follows from the fact that $\inf_u \sup_{\tau \in [0, T_2]} \tilde{\mathbb{P}}(x_{\tau} \in \mathcal{L} \mid x, u) \geq \inf_u \sup_{\tau \in [0, T_1]} \tilde{\mathbb{P}}(x_{\tau} \in \mathcal{L} \mid x, u)$ for any $T_2 \geq T_1, T_1, T_2 \in \mathbb{N}$

Theorem 12 Given that $V_{\tau,\mathcal{L}}(x) = V_{\tau+1,\mathcal{L}}(x)$ implies $V_{\tau,\mathcal{L}}(x) = V_{0,\mathcal{L}}(x)$

Proof : From equation 7.2

$$\begin{aligned} V_{\tau,\mathcal{L}}(x) &= f_{\mathcal{L}}(x) + \inf_{u \in \mathcal{M}} f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{\tau+1,\mathcal{L}}(x') \mathcal{P}(dx'|x, u) = \\ &= f_{\mathcal{L}}(x) + \inf_{u \in \mathcal{M}} f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{\tau,\mathcal{L}}(x') \mathcal{P}(dx'|x, u) = V_{\tau-1,\mathcal{L}}(x) \end{aligned}$$

Thus using same logic we can say $V_{\tau,\mathcal{L}}(x) = V_{0,\mathcal{L}}(x)$

7.3.1. Probable forced backward reachable set computation

This subsection covers the problem of computing the probable forced backward reachable set given a time horizon (T), an avoid set ($\mathcal{L}$) and a probability value ($\mathcal{J} \in [0, 1]$). To tackle this problem, we start with an avoid set and find the probable forced backward reachable set either until the value function converges or until the time horizon ends, as shown in Algorithm 7.1.

Algorithm 7.1: Probable forced backward reachable set

Require: $\mathcal{L}, T, \mathcal{J}$

```

1: Initialize  $V_{T,\mathcal{L}}(x) = f_{\mathcal{L}}(x)$ 
2: Compute  $k_1 = 0, V_{T-k_1,\mathcal{L}}(x), \forall x \in \mathbb{X}$ 
3: Compute  $k_2 = k_1 + 1, V_{T-k_2,\mathcal{L}}(x), \forall x \in \mathbb{X}$ 
if  $V_{T-k_1,\mathcal{L}}(x) = V_{T-k_2,\mathcal{L}}(x), \forall x \in \mathbb{X}$  then
    Go to 4
else if  $V_{T-k_1,\mathcal{L}}(x) \neq V_{T-k_2,\mathcal{L}}(x), \forall x \in \mathbb{X}$  then
     $k_1 \leftarrow k_2, V_{T-k_1,\mathcal{L}}(x) \leftarrow V_{T-k_2,\mathcal{L}}(x)$ 
    Go to 3
4: Set  $V_{T-k_2,\mathcal{L}}(x) = V_{0,\mathcal{L}}(x)$  [Theorem 12]
 $V_{\text{PFBRs}}([T], \mathcal{L}, \mathcal{J}) \leftarrow V_{0,\mathcal{L}}(x) \geq \mathcal{L}$ 

```

To show that algorithm 7.1 is practically implementable, the algorithm must be computationally viable. In the following section we show that for discrete state and control space the computational tractability of the Algorithm 7.1 is guaranteed. For the continuous state and control space, we will provide a discretisation scheme that guarantees convergence and thus preserves computational tractability.

Discrete state space and control space

If the state space and control space are discrete, the stochastic kernel $\mathcal{P}$ is a transition matrix, denoted as $\mathcal{T}(x'|x, u)$, that provides the probability of the transition of state $x \in \mathbb{X}$ to state $x' \in \mathbb{X}$ under the control action $u \in \mathcal{M}$.

The probability transition matrix follows the total law of probability, i.e. $\sum_{x' \in \mathbb{X}} \mathcal{T}(x'|x, u) = 1, \forall x \in \mathbb{X}, u \in \mathcal{M}$. With state space and control space being discrete, the value function $V_{0,\mathcal{L}}(x)$ can be computed exactly, using linear programming [34]. Furthermore, the existence of the optimal Markov policy is always guaranteed.

Lemma 21 *Given an avoid set $\mathcal{L} \subset \mathbb{X}$ the value function $V_{\tau,\mathcal{L}}(x)$ as presented in equation 7.2 is obtained by solving the linear programming problem*

$$\begin{aligned} & \max \sum_{\tau \in [0, T]} \sum_{x \in \mathbb{X}} v_{\tau}(x) \\ & \text{s.t. } v_{\tau}(x) \leq v_{\tau}(x) + \sum_{x' \in \mathbb{X}} \mathcal{T}(x'|x, u) v_{\tau+1}(x) \\ & \quad \forall \tau \in \mathbb{N}_{[0, T-1]}, \forall x \in \mathbb{X}, \forall u \in \mathbb{U}_x \\ & \quad v_T(x) \leq f_{\mathcal{L}}(x), \forall x \in \mathbb{X} \\ & \quad v_{\tau}(x) \in [0, 1] \end{aligned} \tag{7.13}$$

Proof : The proof is exactly the same as the proof of Theorem 1 in [34].

The optimal solution obtained from the linear program 7.13 provides the value give function at respective time step i.e., $V_{\tau,\mathcal{L}} = v_{\tau}(x), \forall x \in \mathbb{X}$ and $\tau \in \mathbb{N}_{[0, T]}$. The linear program further provides the optimal Markov control policy u^* such that equation 7.14 is satisfied [34], where $v_{\tau}^*(x)$ is the optimal value.

$$v_{\tau}^*(x) = v_{\tau}(x) + \sum_{x' \in \mathbb{X}} \mathcal{T}(x'|x, u) v_{\tau+1}^*(x) \tag{7.14}$$

Note: If the cardinality of the state space is $|\mathbb{X}|$, the cardinality of the control space is $|\mathbb{U}|$, and the time horizon is T , then the maximum number of decision variables and constrains are $|\mathbb{X}|(T+1)$ and $|\mathbb{X}|(T|\mathbb{U}|+1)$ respectively. Thus solving the linear program implementation of algorithm 7.1 in the case of discrete state and control space has a maximum implementation time of $O(|\mathbb{X}|^2(T|\mathbb{U}|+1))$ [35].

Continuous state space and control space

For the computation tractability of the value function and the algorithm the continuous state space and the control space are discretised while providing a convergence guarantee. Furthermore, the Algorithm 7.1 is altered to provide an approximation of the probable forced backward reachable set. For any given set $\mathcal{L} \in \mathcal{B}(\mathbb{X})$, we denote the Lebesgue measure of the set as $\psi(\mathcal{L})$. The stochastic kernel $\mathcal{P}(\cdot|x, u)$ projects a probability density kernel denoted as $p(\cdot|x, u)$. The probability density kernel, $p(z|x, u)$ represents the probability density of z given the current state (x) and the control policy (u). We assume that there is a region of interest denoted $\mathcal{G}$ such that $\mathcal{L} \subset \mathcal{G} \subset \mathbb{X} \in \mathcal{B}(\mathbb{X})$. Set $\mathcal{G}$ is compact and $\mathcal{L}$ is a closed subset of $\mathcal{G}$ and thus compact, therefore $\psi(\mathcal{L})$ and $\psi(\mathcal{G})$ are bounded.

Furthermore we assume that the probability density function and the stochastic kernel $\mathcal{P}(\cdot|x, u)$ satisfy the following assumption:

Assumption 9 For any $x, x' \in \mathcal{G} \setminus \mathcal{L}, z, z' \in \mathcal{G} \setminus \mathcal{L}$ and $u, u' \in \mathbb{U}$, there exists a constant η such that equation 7.15 is valid, i.e., the probability density function is Lipschitz continuous.

$$|p(z|x, u) - p(z'|x', u')| \leq \eta(\|z - z'\| + \|x - x'\| + \|u - u'\|) \quad (7.15)$$

Assumption 10 For any $x, x' \in \mathcal{G} \setminus \mathcal{L}, z, z' \in \mathcal{G} \setminus \mathcal{L}$ and $u, u' \in \mathbb{U}$, there exists a constant κ such that equation 7.16 is valid, i.e., the stochastic kernel is Lipschitz continuous.

$$|\mathcal{P}(z|x, u) - \mathcal{P}(z'|x', u')| \leq \eta_2(\|z - z'\| + \|x - x'\| + \|u - u'\|) \quad (7.16)$$

Where η, η_2 are Lipschitz constant.

Assumption 11 Similar to [36], we assume that $p(z|x, u)$ is a probability density on region of interest i.e. $\mathcal{G}$. Thus $p(z|x, u)$ satisfies equation 7.17, where κ_p is a positive constant.

$$\int_{\mathcal{G}} p(z|x, u) dz = 1 \text{ while } p(z|x, u) \in [0, \kappa_p] \quad (7.17)$$

Lemma 22 The set $(\mathcal{G} \setminus \mathcal{L}) \in \mathcal{B}(\mathbb{X})$ if $\mathcal{G}, \mathcal{L} \in \mathcal{B}(\mathbb{X})$

Proof : The proof directly follows from the fact that $\mathcal{G} \setminus \mathcal{L} = \mathcal{G} \cap \mathcal{L}^c = (\mathcal{G}^c \cup \mathcal{L})^c$ and σ -algebra are closed under complements and countable union, thus $\mathcal{G} \setminus \mathcal{L} \in \mathcal{B}(\mathbb{X})$. Where A^c is compliment of set A

We discretise $\mathcal{G}$ into $k_x + 1$ pairwise disjoint non-empty Borel sets $\mathcal{G}_i$ and $\mathcal{L}$, where $i \in [1, k_x] \subseteq \mathbb{N}$, such that $\mathcal{G} = \bigcup_{i=1}^{k_x} \mathcal{G}_i \cup \mathcal{L}$, $\mathcal{G}_a \cap \mathcal{G}_b = \emptyset$, $\mathcal{G}_c \cap \mathcal{L} = \emptyset$ for any $a, b, c \in [1, k_x] \subseteq \mathbb{N}, a \neq b$. Each subset is then represented by a representative state denoted by $g_i \in \mathcal{G}_i$ and $\mathcal{L}$ is represented by a representative state $l \in \mathcal{L}$. Furthermore we define the set of all the representative states as $\hat{\mathcal{G}} = \{l, g_i : i \in [1, k_x] \subseteq \mathbb{N}\}$, maximum set member norm as $d_{x_i} = \sup_{x, x' \in \mathcal{G}_i} \|x - x'\|$ and $D_{x, \max} = \sup_{i \in [1, k_x]} (d_{x_i})$. In the same fashion we also discretise $\mathbb{U}$, into k_u pairwise disjoint non-empty Borel sets denoted by $\mathcal{F}_i, i \in [1, k_u] \subseteq \mathbb{N}$, thus satisfying $\mathbb{U} = \bigcup_{i \in [1, k_u]} \mathcal{F}_i$. Similar to the state space discretisation, each disjoint Borel set is represented by an element in the set, $\hat{u}_i \in \mathcal{F}_i$. Let the set of representative control policies

be denoted by $\hat{\mathcal{U}} = \{\hat{u}_i : i \in [1, k_u] \subseteq \mathbb{N}\}$, $d_{u_i} = \sup_{u, u' \in \mathcal{F}_i} \|u - u'\|$ and $D_{u, \max} = \sup_{i \in [1, k_u]} (d_{u_i})$.

We assume the grid size to be constant $\lambda \geq \max(D_{u, \max}, D_{x, \max})$. For $x \in \mathcal{G}$ we define the admissible control set as 7.18.

$$\hat{\mathcal{U}}_x = \{\hat{u} \in \hat{\mathcal{U}} : \|u - \hat{u}\| \leq \lambda \text{ for some } u \in \mathcal{U}_{g_i}, x \in \mathcal{G}_i\} \quad (7.18)$$

Where g_i is the representative state of $\mathcal{G}_i$ such that $x \in \mathcal{G}_i$. Similar to [36], in Lemma 23 we show that for each $x \in \mathcal{G}$ there exists a non-empty discrete control set.

Lemma 23 *For all $g_i \in \hat{\mathcal{G}}$ the set of available inputs is non-empty, i.e., $\hat{\mathcal{U}}_{g_i} \neq \emptyset$ and the available input at all state $x \in \mathcal{G}_i$ set is equal to $\hat{\mathcal{U}}_{g_i}$, i.e. $\hat{\mathcal{U}}_x = \hat{\mathcal{U}}_{g_i}, \forall x \in \mathcal{G}_i$*

Proof : Since $\mathcal{U}_x \neq \emptyset, \forall x \in \mathbb{X}$, $\mathcal{U}_{g_i}$ is also non-empty. Thus $\exists \hat{u} \in \hat{\mathcal{U}}$ such that $\|u - \hat{u}\| \leq \lambda$ for all $u \in \mathcal{U}_{g_i}$, thus $\hat{\mathcal{U}}_{g_i}$ is non-empty. Moreover, from the definition of $\hat{\mathcal{U}}_x$ it is trivial to obtain $\hat{\mathcal{U}}_{g_i} = \hat{\mathcal{U}}_x$.

Following the steps of [36], we define a function $\hat{p} : \mathcal{G} \setminus \mathcal{L} \times \mathcal{G} \times \hat{\mathcal{U}} \rightarrow \mathbb{R}^+$ as shown in equation 7.19 that provides a unified probability density function to all $x \in \mathcal{G}_i$, such that all $x \in \mathcal{G}_i$ have the same probability kernel. We also assume that $\int_{\mathcal{L}} p(z|x, u) dz = 1, \forall x \in \mathcal{L}$. Note that this assumption does not violate the assumption 4 as $\mathcal{L} \subset \mathcal{G}$

$$\hat{p}(x'|x, \hat{u}) = \begin{cases} \frac{p(g_{x'}|g_x, \hat{u})}{\int_{\mathcal{G}} p(g_{z'}|g_x, \hat{u}) dz} & \text{if } \int_{\mathcal{G}} p(g_{z'}|g_x, \hat{u}) dz \geq 1 \\ p(g_{x'}|g_x, \hat{u}) & \text{otherwise} \end{cases} \quad (7.19)$$

Where $g_{x'}$ and g_x are the respective representative states of the set of states that x' and x belong to. This gives rise to an approximate stochastic control system $\hat{\mathcal{S}}_{\mathcal{G}} = (\hat{\mathcal{G}}, \hat{\mathcal{U}}, \hat{\mathbb{P}})$, where $\hat{\mathbb{P}}$ is the transition probability defined as $\hat{\mathbb{P}}(g_j|g_i, \hat{u}) = \int_{\mathcal{G}_j} \hat{p}(x'|g_i, \hat{u}) dx'$, $(g_j, g_i) \in \hat{\mathcal{G}}$, $g_j \in \mathcal{G}_j, g_i \in \mathcal{G}_i, \hat{\mathbb{P}}(l|g_i, \hat{u}) = \int_{\mathcal{L}} \hat{p}(x'|g_i, \hat{u}) dx'$, $\hat{\mathbb{P}}(l|l, \hat{u}) = \int_{\mathcal{L}} p(x'|l, \hat{u}) dx' = 1$ and $\hat{u} \in \hat{\mathcal{U}}$

We define a function similar to equation 7.1 as $\hat{f}_l : \hat{\mathcal{G}} \rightarrow \{0, 1\}$ as shown in the following equation

$$\hat{f}_l(g_i) = \begin{cases} 1, & \text{if } g_i = l \\ 0, & \text{if } g_i \neq l \end{cases} \quad (7.20)$$

For the discrete system $\hat{\mathcal{S}}_{\mathcal{G}}$, the value function given by equation 7.2 boils down to equation 7.21, where l is a set of all states in $\hat{\mathcal{G}}$ except l , i.e., the complement of l in $\hat{\mathcal{G}}$.

$$\hat{V}_{\tau, l}(g_i) = \hat{f}_l(g_i) + \hat{f}_l(g_i) \min_{\hat{u} \in \hat{\mathcal{U}}} \left(\sum_{j=1}^{k_x} \hat{V}_{\tau+1, l}(g_j) \hat{\mathbb{P}}(g_j|g_i, \hat{u}) \right), \quad \forall \tau \in [0, T] \subseteq \mathbb{N} \quad (7.21)$$

For $x \in \mathcal{G}_i$, $\hat{V}_{\tau,l}(x) = \hat{V}_{\tau,l}(g_i)$ for all $\tau \in [0, T] \in \mathbb{N}$. The discrete optimal control policy is defined as shown in equation 7.22.

$$\hat{u}_{\tau, \text{PFBRs}}^*(g_i) = \arg \min_{\hat{u} \in \hat{\mathcal{U}}} \left(\sum_{j=1}^{k_x} \hat{V}_{\tau+1,l}(g_j) \hat{\mathbb{P}}(g_j | g_i, \hat{u}) \right) \quad (7.22)$$

Note: As observed above, the discrete system $\hat{\mathcal{S}}$ has a finite states and control policy space. Thus the value $\hat{V}_{\tau,l}(g_i)$ can be calculated using the linear program as shown in Lemma 21, while the optimal control policy can be computed using 7.14. The value of the value function $\hat{V}_{\tau,l}(g_i)$ represents the shared probability of all states in $\mathcal{G}_i$ to reach $\mathcal{L}$.

In the following, we guarantee the convergence of the value function as the grid size approaches zero. To prove the convergence of the approximate system's value function to the continuous state space value function we need two more Lemmas.

Lemma 24 *If assumptions 1 and 3 are satisfied, then, for any states $x, x' \in \mathcal{G} \setminus \mathcal{L}$, the value function is Lipschitz continuous, i.e., $V_{\tau,\mathcal{L}}(x)$ satisfies equation 7.23*

$$|V_{\tau,\mathcal{L}}(x) - V_{\tau,\mathcal{L}}(x')| \leq \kappa_1 \psi(\mathcal{G}) \|x - x'\|, \forall \tau \in [0, T] \subseteq \mathbb{N} \quad (7.23)$$

Proof : At $\tau = T$ the value function is equal to equation 7.1. Thus $|V_{T,\mathcal{L}}(x) - V_{T,\mathcal{L}}(x')| \leq \kappa_T \|x - x'\|$, $\forall x, x' \in \mathcal{G} \setminus \mathcal{L}$ holds trivially. For $\tau \in [0, T-1]$ using equation 7.2 we get:

$$\begin{aligned} |V_{\tau,\mathcal{L}}(x) - V_{\tau,\mathcal{L}}(x')| &= |f_{\mathcal{L}}(x) + \inf_u f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{\tau+1,\mathcal{L}}(\bar{x}) \mathcal{P}(d\bar{x} | x, u) \\ &\quad - f_{\mathcal{L}}(x') - \inf_u f_{\bar{\mathcal{L}}}(x') \int_{\mathbb{X}} V_{\tau+1,\mathcal{L}}(\bar{x}) \mathcal{P}(d\bar{x} | x', u)|, \quad \forall x, x' \in \mathcal{G} \\ &= |\inf_u \int_{\mathbb{X}} V_{\tau+1,\mathcal{L}}(\bar{x}) \mathcal{P}(d\bar{x} | x, u) - \inf_u \int_{\mathbb{X}} V_{\tau+1,\mathcal{L}}(\bar{x}) \mathcal{P}(d\bar{x} | x', u)|, \\ &\quad \forall x, x' \in \mathcal{G} \setminus \mathcal{L} \\ &\leq \inf_u \left| \int_{\mathbb{X}} V_{\tau+1,\mathcal{L}}(\bar{x}) \mathcal{P}(d\bar{x} | x, u) - \int_{\mathbb{X}} V_{\tau+1,\mathcal{L}}(\bar{x}) \mathcal{P}(d\bar{x} | x', u) \right|, \quad \forall x, x' \in \mathcal{G} \setminus \mathcal{L} \\ &\leq \inf_u \left| \int_{\mathbb{X}} V_{\tau+1,\mathcal{L}}(\bar{x}) (\mathcal{P}(d\bar{x} | x, u) - \mathcal{P}(d\bar{x} | x', u)) \right|, \quad \forall x, x' \in \mathcal{G} \setminus \mathcal{L} \end{aligned}$$

$$\leq \inf_u \int_{\mathbb{X}} |V_{\tau+1, \mathcal{L}}(\bar{x})| |(\mathcal{P}(d\bar{x}|x, u) - \mathcal{P}(d\bar{x}|x', u))|, \quad \forall x, x' \in \mathcal{G} \setminus \mathcal{L}$$

Since $V_{\tau+1, \mathcal{L}}(\bar{x}) \in [0, 1]$

$$\leq \inf_u \int_{\mathbb{X}} |(\mathcal{P}(d\bar{x}|x, u) - \mathcal{P}(d\bar{x}|x', u))|, \quad \forall x, x' \in \mathcal{G} \setminus \mathcal{L}$$

$$\leq \inf_u \psi(\mathcal{G}) \kappa_1 \|x - x'\| = \kappa_1 \psi(\mathcal{G}) \|x - x'\|$$

Lemma 25 *If assumptions 2 and 4 are satisfied, then $\forall x' \in \mathcal{G} \setminus \mathcal{L}$ and $g_i \in \hat{\mathcal{G}} \setminus \{l\}$ the integration of the modulus of difference of the approximate probability density with actual probability density over the space $\mathcal{G}$ is bounded for all $\hat{u} \in \hat{\mathcal{U}}$, i.e., equation 7.24 holds.*

$$\int_{\mathcal{G}} |\hat{p}(x'|g_i, \hat{u}) - p(x'|g_i, \hat{u})| dx' \leq 2\kappa_2 \psi(\mathcal{G})\lambda \quad (7.24)$$

Proof: If $\int_{\mathcal{G}} p(g_{x'}|g_i, \hat{u}) dx' \leq 1$ then using assumption 2, the above lemma holds trivially. If $\int_{\mathcal{G}} p(g_{x'}|g_i, \hat{u}) dx' \geq 1$ then:

$$\text{Firstly, } 0 \leq \int_{\mathcal{G}} p(g_{x'}|g_i, \hat{u}) dx' - 1$$

$$\begin{aligned} 0 &\leq \int_{\mathcal{G}} p(g_{x'}|g_i, \hat{u}) dx' - \int_{\mathcal{G}} p(x'|g_i, \hat{u}) dx' \\ &\leq \int_{\mathcal{G}} |p(g_{x'}|g_i, \hat{u}) - p(x'|g_i, \hat{u})| dx' \leq \kappa_2 \psi(\mathcal{G})\lambda \end{aligned}$$

Using assumption 2

$$\text{Secondly, we have } \int_{\mathcal{G}} |\hat{p}(x'|g_i, \hat{u}) - p(x'|g_i, \hat{u})| dx'$$

$$= \int_{\mathcal{G}} \frac{|p(g_{x'}|g_i, \hat{u}) - p(x'|g_i, \hat{u})| \int_{\mathcal{G}} p(g_z|g_i, \hat{u}) dz}{\int_{\mathcal{G}} p(g_z|g_i, \hat{u}) dz} dx'$$

$$\begin{aligned}
&\leq \int_{\mathcal{G}} |p(g_{x'}|g_i, \hat{u}) - p(x'|g_i, \hat{u})| \int_{\mathcal{G}} p(g_z|g_i, \hat{u}) dz dx' \\
&\leq \int_{\mathcal{G}} |p(g_{x'}|g_i, \hat{u}) - p(x'|g_i, \hat{u})| dx' + \\
&\quad \int_{\mathcal{G}} |p(x'|g_i, \hat{u}) - p(x'|g_i, \hat{u})| \int_{\mathcal{G}} p(g_z|g_i, \hat{u}) dz dx' \\
&\leq \int_{\mathcal{G}} |p(g_{x'}|g_i, \hat{u}) - p(x'|g_i, \hat{u})| dx' + \\
&\quad \left| \int_{\mathcal{G}} p(g_z|g_i, \hat{u}) dz - 1 \right| \int_{\mathcal{G}} |p(x'|g_i, \hat{u})| dx' \leq 2\kappa_2 \psi(\mathcal{G})\lambda
\end{aligned}$$

In the following lemma we prove the convergence of the discretised value function to the actual value function as the grid size tends to zero.

Lemma 26 *If assumption 2-4 are satisfied, then for all $x \in \mathcal{G}$ the difference between the value function and the discretised value function is bounded as shown in equation 7.25.*

7

$$|\hat{V}_{\tau, \mathcal{L}}(x) - V_{\tau, \mathcal{L}}(x)| \leq f_{\tau}(\mathcal{G})\lambda \quad (7.25)$$

Proof: For $x \in \mathcal{L}$ we have $\hat{V}_{\tau, \mathcal{L}}(x) = V_{\tau, \mathcal{L}}(x) = 1$ for all $\tau \in [0, T]$. Thus equation 7.25 holds trivially. For all $x \in \mathcal{G} \setminus \mathcal{L}$, and $\tau = T$, $\hat{V}_{T, \mathcal{L}}(x) = V_{T, \mathcal{L}}(x) = 0$. Proceeding using induction we assume that for some $\tau + 1$, $|\hat{V}_{\tau+1, \mathcal{L}}(x) - V_{\tau+1, \mathcal{L}}(x)| \leq f_{\tau+1}(\mathcal{G})\lambda$. For any $g_i \in \mathcal{G}_i$, $i \in [0, k_x] \subseteq \mathbb{N}$, we denote u_{τ}^* as the input policy that minimises the value function, i.e., $u_{\tau}^* = \arg \inf_{u \in \mathbb{U}} \int_{\mathcal{G}} V_{\tau+1, \mathcal{L}}(x') \mathcal{P}(dx'|g_i, u)$. Furthermore we denote $\hat{u}_{\tau}^*$ as the input policy that minimises the approximate value function, i.e $\hat{u}_{\tau}^* = \arg \inf_{\hat{u} \in \hat{\mathbb{U}}} \int_{\mathcal{G}} \hat{V}_{\tau+1, \mathcal{L}}(x') \hat{\mathcal{P}}(dx'|g_i, \hat{u})$. Thus we have

$$\begin{aligned}
V_{\tau, \mathcal{L}}(g_i) - \hat{V}_{\tau, \mathcal{L}}(g_i) &= \\
&\quad \int_{\mathcal{G}} V_{\tau+1, \mathcal{L}}(x') \mathcal{P}(dx'|g_i, u_{\tau}^*) - \int_{\mathcal{G}} \hat{V}_{\tau+1, \mathcal{L}}(x') \hat{\mathcal{P}}(dx'|g_i, \hat{u}_{\tau}^*) \\
&= \int_{\mathcal{G}} |V_{\tau+1, \mathcal{L}}(x') \mathcal{P}(dx'|g_i, u_{\tau}^*)| - \int_{\mathcal{G}} |\hat{V}_{\tau+1, \mathcal{L}}(x') \hat{\mathcal{P}}(dx'|g_i, \hat{u}_{\tau}^*)|
\end{aligned}$$

$$\begin{aligned}
&\leq \left| \int_{\mathcal{G}} |V_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, u_{\tau}^*)| - \int_{\mathcal{G}} |\hat{V}_{\tau+1,\mathcal{L}}(x')\hat{\mathcal{P}}(dx'|g_i, \hat{u}_{\tau}^*)| \right| \\
&\leq \int_{\mathcal{G}} |V_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, u_{\tau}^*) - \hat{V}_{\tau+1,\mathcal{L}}(x')\hat{\mathcal{P}}(dx'|g_i, \hat{u}_{\tau}^*)| \\
&\leq \int_{\mathcal{G}} |V_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, u_{\tau}^*) - \hat{V}_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, u_{\tau}^*)| + \\
&\quad \int_{\mathcal{G}} |\hat{V}_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, u_{\tau}^*) - \hat{V}_{\tau+1,\mathcal{L}}(x')\hat{\mathcal{P}}(dx'|g_i, \hat{u}_{\tau}^*)|
\end{aligned}$$

Since $\hat{V}_{\tau+1,\mathcal{L}}(x') \leq 1, \forall x' \in \mathcal{G}$

$$\begin{aligned}
&\leq \int_{\mathcal{G}} |\hat{\mathcal{P}}(dx'|g_i, \hat{u}_{\tau}^*) - \mathcal{P}(dx'|g_i, u_{\tau}^*)| + \int_{\mathcal{G}} |\hat{V}_{\tau+1,\mathcal{L}}(x) - V_{\tau+1,\mathcal{L}}(x)| |\mathcal{P}(dx'|g_i, u_{\tau}^*)| \\
&\leq 2\kappa_2\psi(\mathcal{G})\lambda + f_{\tau+1}(\mathcal{G})\lambda
\end{aligned}$$

and for $\hat{V}_{\tau,\mathcal{L}}(g_i) - V_{\tau,\mathcal{L}}(g_i) =$

$$\begin{aligned}
&\int_{\mathcal{G}} \hat{V}_{\tau+1,\mathcal{L}}(x')\hat{\mathcal{P}}(dx'|g_i, \hat{u}_{\tau}^*) - \int_{\mathcal{G}} V_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, u_{\tau}^*) \\
&\leq \int_{\mathcal{G}} |\hat{V}_{\tau+1,\mathcal{L}}(x')\hat{\mathcal{P}}(dx'|g_i, \hat{u}_{\tau}^*)| - \int_{\mathcal{G}} |V_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, u)| \\
&\leq \int_{\mathcal{G}} |\hat{V}_{\tau+1,\mathcal{L}}(x')\hat{\mathcal{P}}(dx'|g_i, \hat{u}_{\tau}^*) - \hat{V}_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, \hat{u}_{\tau}^*)| \\
&\quad + |\hat{V}_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, \hat{u}_{\tau}^*) - V_{\tau+1,\mathcal{L}}(x')\mathcal{P}(dx'|g_i, \hat{u}_{\tau}^*)| \\
&\leq 2\kappa_2\psi(\mathcal{G})\lambda + f_{\tau+1}(\mathcal{G})\lambda
\end{aligned}$$

This leads to $|V_{\tau,\mathcal{L}}(g_i) - \hat{V}_{\tau,\mathcal{L}}(g_i)| \leq 2\kappa_2\psi(\mathcal{G})\lambda + f_{\tau+1}(\mathcal{G})\lambda$

Since for any $x \in \mathcal{G}_i$, $i \in [0, T] \subseteq \mathbb{N}$, $\hat{V}_{\tau, \mathcal{L}}(x) = \hat{V}_{\tau, \mathcal{L}}(g_i)$

$$|V_{\tau, \mathcal{L}}(x) - \hat{V}_{\tau, \mathcal{L}}(x)| = |V_{\tau, \mathcal{L}}(x) - \hat{V}_{\tau, \mathcal{L}}(g_i)|$$

$$\leq |V_{\tau, \mathcal{L}}(x) - V_{\tau, \mathcal{L}}(g_i)| + |V_{\tau, \mathcal{L}}(g_i) - \hat{V}_{\tau, \mathcal{L}}(g_i)|$$

$$\leq (\kappa_1 + 2\kappa_2)\psi(\mathcal{G})\lambda + f_{\tau+1}(\mathcal{G})\lambda = f_{\tau}(\mathcal{G})\lambda$$

Thus the above lemma assures the convergence of the approximate value function to the true value function as the grid size tends to zero. For the discretised space, the probabilistic forced backward reachable set can be defined as shown in equation 7.26

$$\hat{V}_{\text{PFBRs}}(T, \mathcal{L}, \hat{\mathcal{J}}) = \{x \in \hat{\mathcal{S}} : \inf_{\hat{u}} \sup_{\tau} \hat{\mathbb{P}}^{\tau}(x' \in \mathcal{L} | x, \hat{u}) \geq \hat{\mathcal{J}}, \hat{u} \in \hat{\mathcal{U}}, \tau \in \mathbb{N}_{0, T}\} \quad (7.26)$$

Theorem 13 *If assumptions 1-4 hold, then the probabilistic forced backward reachable set's $V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J})$ approximate counterpart in the discretised space $\hat{\mathcal{S}}$ is given by $\hat{V}_{\text{PFBRs}}(T, \mathcal{L}, \hat{\mathcal{J}})$ where $\hat{\mathcal{J}} = \mathcal{J} + f_{\tau}(\mathcal{G})\lambda$*

Proof: As per the discretisation routine, for all $x \in \mathcal{G}_i$, for all $\tau \in [0, T] \in \mathbb{N}$ we have $\hat{V}_{\tau, \mathcal{L}}(x) = \hat{V}_{\tau, \mathcal{L}}(g_i)$. Thus if $\hat{V}_{0, \mathcal{L}}(x) \geq \hat{\mathcal{J}}$, then using Lemma 26 and triangle inequality $V_{0, \mathcal{L}}(x) \geq \hat{V}_{0, \mathcal{L}}(x) - f_{\tau}(\mathcal{G})$.

Thus with an appropriate grid size $0 < \lambda \leq \frac{1-\mathcal{J}}{f_{\tau}(\mathcal{G})}$, the problem of obtaining $V_{\text{PFBRs}}(T, \mathcal{L}, \mathcal{J})$ boils down to calculating the approximate probabilistic forced backward reachable set, $\hat{V}_{\text{PFBRs}}(T, \mathcal{L}, \hat{\mathcal{J}})$, with an appropriate $\hat{\mathcal{J}} = \mathcal{J} + f_{\tau}(\mathcal{G})\lambda$. Algorithm 7.2 is an adaptation of algorithm 7.1 for the approximate system space $\hat{\mathcal{S}}$ created by discretisation. After the discretisation process, algorithm 7.2 follows similar steps as algorithm 7.1 to obtain the approximate forced backward reachable sets.

Note: Following similar arguments as for the computation cost for algorithm 7.1 and [35] we can say that algorithm 7.2 has a computation time of $O(|k_x|^2(T|K_u| + 1))$.

7.4. CONVERGENCE IN TIME

In this section we focus on the convergence of the forced backward reachable set in time. Firstly we provide the sufficient condition required for the set to converge. Before we move to convergence in time, we need to define a few concepts.

Definition 2 (Stationary Markov control policy) : A stationary Markov policy is a Markov policy $u \in \mathcal{M}$ with $u = (\bar{u}, \bar{u}, \dots)$, $\bar{u} : \mathbb{X} \rightarrow \mathbb{U}$ that is not changing with

Algorithm 7.2: Approximate probable forced backward reachable set**Require:** $\mathcal{L}, T, \mathcal{J}, \hat{\mathcal{J}}, \mathcal{G}$ **1: Discretise** Select grid size $0 < \lambda \leq \frac{1-\mathcal{J}}{f_{\tau}(\mathcal{G})}$ and discretise $\mathcal{G} \setminus \mathcal{L}$ **2: Initialise** $\hat{V}_{T,\mathcal{L}}(x) = f_{\mathcal{L}}(x)$ **3: Compute** $k_1 = 0, \hat{V}_{T-k_1,\mathcal{L}}(x), \forall x \in \hat{\mathcal{G}}$ **4: Compute:** $k_2 = k_1 + 1, \hat{V}_{T-k_2,\mathcal{L}}(x), \forall x \in \hat{\mathcal{G}}$ **if** $\hat{V}_{T-k_1,\mathcal{L}}(x) = \hat{V}_{T-k_2,\mathcal{L}}(x), \forall x \in \hat{\mathcal{G}}$ **then**

Go to 5

else if $\hat{V}_{T-k_1,\mathcal{L}}(x) \neq \hat{V}_{T-k_2,\mathcal{L}}(x), \forall x \in \hat{\mathcal{G}}$ **then** $k_1 \leftarrow k_2, \hat{V}_{T-k_1,\mathcal{L}}(x) \leftarrow \hat{V}_{T-k_2,\mathcal{L}}(x)$

Go to 4

 $k_2 = T$ **5: Set** $V_{T-k_2,\mathcal{L}}(x) = V_{0,\mathcal{L}}(x)$ [Theorem 12] $V_{\text{PFBRs}}([T], \mathcal{L}, \mathcal{J}) \leftarrow V_{0,\mathcal{L}}(x) \geq \hat{\mathcal{J}}$ [Theorem 13]

time and is universally measurable.

Given a state x and a stationary Markov policy u , we define a $\bar{\mathbb{P}}^{\infty}(x' \in \mathcal{L}|x, u)$ as the maximum probability that the state x will transition into state $x' \in \mathcal{L}$ at any $\tau \in \mathbb{N}$. Thus we can define the infinite time variant of the probabilistic forced backward reachable set as shown in equation 7.27.

$$\begin{aligned}
 V_{\text{PFBRs}}^{\infty}(\mathcal{L}, \mathcal{J}) &= \{x : \forall u \in \mathcal{M}, \exists \tau \in \mathbb{N}, \bar{\mathbb{P}}^{\tau}(x_{\tau} \in \mathcal{L}|x, u) \geq \mathcal{J}\} \\
 &= \{x : \inf_u \sup_{\tau} \bar{\mathbb{P}}^{\tau}(x_{\tau} \in \mathcal{L}|x, u) \geq \mathcal{J}, u(\cdot) \in \mathcal{M}, \tau \in \mathbb{N}\} \\
 &= \{x : \inf_u \bar{\mathbb{P}}^{\infty}(x' \in \mathcal{L}|x, u) \geq \mathcal{J}, u \in \mathcal{M}\} \quad (7.27)
 \end{aligned}$$

For the set $V_{\text{PFBRs}}^{\infty}(\mathcal{L}, \mathcal{J})$, Lemma 20 holds, i.e., $V_{\text{PFBRs}}^{\infty}(\mathcal{L}, \mathcal{J})$ is universally measurable. Furthermore, following dynamic programming principles we define a value function $V_{\tau,\mathcal{L}}^* : \mathbb{X} \rightarrow [0, 1], \tau \in \mathbb{N}$ similar to [30] as shown in equation 7.28.

$$V_{\tau,\mathcal{L}}^*(x) = f_{\mathcal{L}}(x) + \inf_{u \in \mathcal{M}} f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{\tau+1,\mathcal{L}}^*(x') \mathcal{P}(dx'|x, u), x \in \mathbb{X} \quad (7.28)$$

Assumption 12 *There exists a $\tau \in \mathbb{N}$ such that the set $\mathbb{U}_x(\tau, \mathcal{J})$ as defined in equation 7.29 is compact for all $x \in \mathbb{X}, \mathcal{J} \in [0, 1]$ and $\tau \in \mathbb{N}$.*

$$\mathbb{U}_x(\tau, \mathcal{J}) = \{u \in \mathbb{U} \mid \int_{\mathbb{X}} V_{\tau,\mathcal{L}}^*(x') \mathcal{P}(dx'|x, u) \geq \mathcal{J}\} \quad (7.29)$$

If assumption 12 holds then as proven in [30], the limit $V_{\infty, \mathcal{L}}^*(x)$ exists for all $x \in \mathbb{X}$ and satisfies equation 7.30.

$$V_{\infty, \mathcal{L}}^*(x) = f_{\mathcal{L}}(x) + \inf_{u \in \mathcal{M}} f_{\mathcal{L}}(x) \int_{\mathbb{X}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx'|x, u), \quad x \in \mathbb{X} \quad (7.30)$$

Moreover, the optimal stationary Markov policy is given by equation 7.31.

$$\bar{u} = \arg \inf_{u \in \mathcal{U}_x(\tau, \mathcal{J})} f_{\mathcal{L}}(x) \int_{\mathbb{X}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx'|x, u), \quad x \in \mathbb{X} \quad (7.31)$$

This value function represents the probability measure, $\inf_u \bar{\mathbb{P}}^\infty(x' \in \mathcal{L}|x, u)$. The value function is initialised as $f_{\mathcal{L}}(x)$.

Given a set $V_{\text{SI}}(T, \mathcal{J})$ and a probability value $\mathcal{J}$ we call this set a stochastic invariant set if, for every state in the set, for all time horizons T , there exists a control input u such that the probability of a state to stay in the set is greater than $\mathcal{J}$. This is a refined definition of the maximum probabilistic safe set as defined in [30] as by requiring any state in the $V_{\text{SI}}(T, \mathcal{J})$ to have a probability to stay in $V_{\text{SI}}(T, \mathcal{J})$ of at least $\mathcal{J}$. The stochastic invariant set as defined above is called an absorbing set if the control input is given to the system and $\mathcal{J} = 1$ [37].

Note: Maximum probabilistic safe set as defined in [30] Given an initial set $\mathcal{A}$, a probability value $\mathcal{J}$, and a time horizon $T \in \mathbb{N}$, the set of all states $x \in \mathcal{A}$ such that, for all time horizons, there exists a control policy such that the probability of the system to stay in set $\mathcal{A}$ is greater than or equal to $\mathcal{J}$.

Let the maximum probabilistic safe set be denoted as $V_{\text{MPSS}}(T, \mathcal{A})$. It can be shown that the set $V_{\text{MPSS}}(T, \mathcal{A})$ is universally measurable under few assumptions. The proof of the previous statement along with the assumptions can be found in [30].

Lemma 27 $\inf_u \sup_\tau \mathcal{P}^\tau(x^\tau \in \mathcal{A}|x, u)$ is equal to $1 - \sup_u \inf_\tau \mathcal{P}^\tau(x^\tau \notin \mathcal{A}|x, u)$ given that $\mathcal{P}(\mathbb{X}|x, u) = 1$

Proof :

$$\begin{aligned} \inf_u \sup_\tau \mathcal{P}^\tau(x^\tau \in \mathcal{A}|x, u) &= \inf_u (-\inf_\tau (-\mathcal{P}^\tau(x^\tau \in \mathcal{A}|x, u))) \\ &= \inf_u (-\inf_\tau (1 - \mathcal{P}^\tau(x^\tau \notin \mathcal{A}|x, u))) = -\sup_u -(-\inf_\tau (1 - \mathcal{P}^\tau(x^\tau \notin \mathcal{A}|x, u))) \\ &= -\sup_u \inf_\tau (1 - \mathcal{P}^\tau(x^\tau \notin \mathcal{A}|x, u)) = 1 - \sup_u \inf_\tau \mathcal{P}^\tau(x^\tau \notin \mathcal{A}|x, u) \end{aligned}$$

Lemma 28 The set of states $\mathcal{Q}$ is a stochastic invariant set with a probability value $\mathcal{J} = 1$ and $T = \infty$ if and only if $V_{\infty, \mathcal{L}}^*(x) = 0, \forall x \in \mathcal{Q}$ given that $\mathcal{L} \cap \mathcal{Q} = \emptyset$ and assumption 1 holds.

Proof : Firstly for the if part, i.e., given that $V_{\infty, \mathcal{L}}^* = 0$ for all $x \in \mathcal{Q}$, $\mathcal{Q}$ is a stochastic invariant set.

$$V_{\infty, \mathcal{L}}^*(x) = f_{\mathcal{L}}(x) + \inf_{u \in \mathcal{M}} f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx'|x, u), \quad x \in \mathbb{X}$$

$$V_{\infty, \mathcal{L}}^*(x) = \inf_{u \in \mathcal{M}} \int_{\mathbb{X}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx'|x, u), \quad x \in \mathbb{X} - \mathcal{L}$$

Let $\bar{u}$ be the input policy that minimises the value function

$$\begin{aligned} V_{\infty, \mathcal{L}}^*(x) &= \int_{\mathbb{X} \setminus \mathcal{Q}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx'|x, \bar{u}) \\ &\quad + \int_{\mathcal{Q}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx'|x, \bar{u}), \quad x \in \mathbb{X} \setminus \mathcal{L} \end{aligned}$$

$$\int_{\mathbb{X} - \mathcal{Q}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx'|x, \bar{u}) = 0, \quad x \in \mathcal{Q}$$

Since we know $V_{\infty, \mathcal{L}}^*(x) \neq 0, \forall x \notin \mathcal{Q}$ thus

$$\int_{\mathbb{X} - \mathcal{Q}} \mathcal{P}(dx'|x, \bar{u}) = 0, \quad x \in \mathcal{Q}, \forall T \in \mathbb{N}$$

Thus $\mathcal{P}(\mathcal{Q}|x, \bar{u}) = 1, x \in \mathcal{Q}, \forall T \in \mathbb{N}$

Now for the only if part, i.e. if $\mathcal{Q}$ is a stochastic invariant set, then $V_{\infty, \mathcal{L}}^* = 0, \forall x \in \mathcal{Q}$. We are given that $\mathcal{Q}$ is an infinite time stochastic invariant set with $\mathcal{J} = 1$. Thus for all $x \in \mathcal{Q}$, $\sup_u \inf_{\tau} \mathcal{P}^{\tau}(x^{\tau} \in \mathcal{Q}|x, u) = 1$. From Lemma 27 we know that $\inf_u \sup_{\tau} \mathcal{P}^{\tau}(x^{\tau} \notin \mathcal{Q}|x, u) = 0$, thus $\inf_u \sup_{\tau} \mathcal{P}^{\tau}(x^{\tau} \in \mathbb{X} \setminus \mathcal{Q}|x, u) = 0$. Furthermore since $\mathcal{Q} \cap \mathcal{L} = \emptyset$, we receive $\inf_u \sup_{\tau} \mathcal{P}^{\tau}(x^{\tau} \in \mathcal{L}|x, u) = 0$. Hence $V_{\infty, \mathcal{L}}^*(x) = 0, \forall x \in \mathcal{Q}$.

Theorem 14 *If there exists an infinite time stochastic invariant set $\mathcal{Q}$ with $\mathcal{J} = 1$ and $\inf_u \inf_{\tau} \mathcal{P}^{\tau}(x^{\tau} \in \mathcal{Q}|x, u) \neq 0$ from some $x \in \mathbb{X}$, then there exists a probability value $0 < \mathcal{J}' < 1$ such that $V_{\text{PFBRs}}^{\infty}(\mathcal{L}, \mathcal{J}') \setminus \mathcal{L}$ is non-empty.*

Proof: We are given that there exists an infinite time stochastic invariant set $\mathcal{Q}$ with $\mathcal{J} = 1$ and $\inf_u \inf_{\tau} \mathcal{P}^{\tau}(x^{\tau} \in \mathcal{Q}|x, u) \neq 0$ from some $x \in \mathbb{X}$. We know that $V_{\infty, \mathcal{L}}^*(x) \neq 0$ for $x \in \mathbb{X} \setminus \mathcal{Q}$ using Lemma 28. Furthermore since

$\mathcal{P}(x' \in \mathcal{Q} | x, u) + \mathcal{P}(x' \in \mathbb{X} \setminus \mathcal{Q} | x, u) = \mathcal{P}(x' \in \mathbb{X} | x, u) = 1$ and $\inf_u \inf_{\tau} \mathcal{P}^{\tau}(x^{\tau} \in \mathcal{Q} | x, u) \leq \mathcal{P}(x' \in \mathcal{Q} | x, u)$ we get that $\mathcal{P}(x' \in \mathbb{X} \setminus \mathcal{Q} | x, u) < 1$.

$$V_{\infty, \mathcal{L}}^*(x) = f_{\mathcal{L}}(x) + \inf_{u \in \mathcal{M}} f_{\tilde{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx' | x, u), \quad x \in \mathbb{X}$$

$$V_{\infty, \mathcal{L}}^*(x) = \inf_{u \in \mathcal{M}} \int_{\mathbb{X}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx' | x, u), \quad x \in \mathbb{X} - \mathcal{L}$$

Let $\bar{u}$ be the input policy that minimises the value function

$$V_{\infty, \mathcal{L}}^*(x) = \int_{\mathbb{X} - \mathcal{Q}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx' | x, \bar{u}) + \int_{\mathcal{Q}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx' | x, \bar{u}),$$

$$x \in \mathbb{X} - \mathcal{L}$$

$$0 < V_{\infty, \mathcal{L}}^*(x) < \int_{\mathbb{X} \setminus \mathcal{Q}} V_{\infty, \mathcal{L}}^*(x') \mathcal{P}(dx' | x, \bar{u}) \leq$$

$$\int_{\mathbb{X} \setminus \mathcal{Q}} \mathcal{P}(dx' | x, \bar{u}) = \mathcal{P}(x' \in \mathbb{X} \setminus \mathcal{Q} | x, u) < 1, \quad x \in \mathbb{X} \setminus \mathcal{L} \cup \mathcal{Q}$$

7

7.5. EXAMPLE

We consider the system of a 2-D cart as used in [8] as an illustrative example. The system is given by equation 7.32.

$$x_{\tau+1} = Ax_{\tau} + Bu_{\tau} + Cw_{\tau}$$

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, C = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad (7.32)$$

The input policy is constrained by $|u| \leq 1$, w_{τ} is considered stochastic with zero mean Gaussian distribution with a variance of 2. Since the noise in the chosen example is supporting the whole space rather than only the region of interest, we decided to apply the boundary condition as value function at the or beyond the boundary is zero. This approximation introduces some errors which can be reduced by increasing the size on the grid and only checking the value function further away from the boundary. The value function of the computed probabilistic forced backward reachable set is shown in Figure 7.1. The initial set $\mathcal{L}$ for the example system was taken as a states in 1 meter radius circle with origin as center. The original grid was taken from $[-10$ to $10]$ with the region of interest being $[-5$ to $5]$ in both x, y dimensions. This was done to reduce the error generated by the boundary condition as stated above. Furthermore random points were selected in the .8 contour region of the value function of the example




Figure 7.1: Value function of computed probabilistic forced backward reachable set

system and was driven under the respective optimal control policy. It was observed that out of 1000 simulations, 782 times the system entered the avoid set. This of course does not serve as a validation but as a demonstrative example that if the system is present in the .8 contour region of the value function, then there is a 0.8 probability that the system will enter the avoid region even if the optimal control policy is applied to stop the system from entering the avoid set. The contour regions of the value function of the example system is shown in figure 7.2.

7.6. CONCLUSIONS

In this paper we focused on the forced backward reachable set in a stochastic setting. We proposed an algorithm for finding the set about certain sets of states that we want the system to avoid and in turn obtaining the region that the system needs to avoid entering the avoid states with a certain probability. The proposed algorithm is iterative in nature and computes the exact avoid region for a discrete space and time system. For a continuous space and discrete time system we append the previous algorithm and thus are able to estimate the probable forced backward reachable set with an error that is upper-bounded on the grid




Figure 7.2: contours of value function of computed probabilistic forced backward reachable set

7

size. We also proved the convergence of the estimated probable forced backward reachable set to the probable forced backward reachable set as the grid size tends to zero. The algorithms proposed in the paper are computationally tractable and we also provided the maximum computation cost of each algorithm. For the aspects of convergence of the set in time, we provided and proved a sufficient condition. We also provided the relationship of the probable forced backward reachable set with the invariant set.

7.7. APPENDIX

In the context of avoid problem the avoid set $\mathcal{L}$ is same as saying that state in the state will stay in the set for the whole time horizon with a probability of one i.e. if $x \in \mathcal{L} \Rightarrow \tilde{\mathbb{P}}^\tau(x_\tau \in \mathcal{L} | x, u) = 1, \forall \tau \in \mathbb{N}, \forall u \in \mathbb{U}$. This implies that the probability value of any state $x \in \mathbb{X}$ to transition in to the avoid set $\mathcal{L}$ is given by equation 7.33

$$\tilde{\mathbb{P}}^\tau(x_\tau \in \mathcal{L} | x, u) = \begin{cases} 1 & \text{if } x \in \mathcal{L} \\ \tilde{\mathbb{P}}^\tau(x_\tau \in \mathcal{L} | x, u) & \text{else} \end{cases} \quad (7.33)$$

For at $\tau = 0$ the we can safely assume that the $\bar{\mathbb{P}}^0(x_0 \in \mathcal{L}|x, u) = f_{\mathcal{L}}(x)$, $\forall u \in \mathbb{U}_x$ where $f_{\mathcal{L}}(x)$ is the indicator function. Furthermore using indicator function equation 7.33 can be given as 7.34.

$$\bar{\mathbb{P}}^\tau(x_\tau \in \mathcal{L}|x, u) = f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x)\bar{\mathbb{P}}^\tau(x_\tau \in \mathcal{L}|x, u) \quad (7.34)$$

In the lemma 29 we will prove the connection between the value function and the probability measure using induction.

Lemma 29 For a given $u \in \mathbb{U}$ the function $V_{T-k, \mathcal{L}}(x) = f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{T-k+1, \mathcal{L}}(x') \mathcal{P}(dx'|x, u)$ $\bar{\mathbb{P}}^k(x'' \in \mathcal{L}|x, u)$, $x \in \mathbb{X}$ given that $V_{T, \mathcal{L}}(x) = f_{\mathcal{L}}(x)$

Proof:

$$V_{T, \mathcal{L}}(x) = f_{\mathcal{L}}(x) = \bar{\mathbb{P}}^0(x'' \in \mathcal{L}|x, u)$$

$$\begin{aligned} \text{For } k = 1 \quad \bar{\mathbb{P}}^1(x'' \in \mathcal{L}|x, u) &= f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x)\bar{\mathbb{P}}^1(x'' \in \mathcal{L}|x, u) = \\ &= f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} \bar{\mathbb{P}}^0(x'' \in \mathcal{L}|x', u) \mathcal{P}(dx'|x, u) = \\ &= f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} V_{T, \mathcal{L}}(x') \mathcal{P}(dx'|x, u) = V_{T-1, \mathcal{L}}(x) \end{aligned}$$

$$\begin{aligned} \text{Assume for } k \quad \bar{\mathbb{P}}^k(x'' \in \mathcal{L}|x, u) &= f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x)\bar{\mathbb{P}}^k(x'' \in \mathcal{L}|x, u) = \\ &= f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} \bar{\mathbb{P}}^{k-1}(x'' \in \mathcal{L}|x', u) \mathcal{P}(dx'|x, u) = \\ &= f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} V_{T-k+1, \mathcal{L}}(x') \mathcal{P}(dx'|x, u) = V_{T-k, \mathcal{L}}(x) \end{aligned}$$

$$\begin{aligned} \text{For } k+1 \quad \bar{\mathbb{P}}^{k+1}(x'' \in \mathcal{L}|x, u) &= \\ &= f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x)\bar{\mathbb{P}}^{k+1}(x'' \in \mathcal{L}|x, u) = \\ &= f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} \bar{\mathbb{P}}^k(x'' \in \mathcal{L}|x', u) \mathcal{P}(dx'|x, u) \\ &= f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} V_{T-k, \mathcal{L}}(x') \mathcal{P}(dx'|x, u) = V_{T-k-1, \mathcal{L}}(x) \end{aligned}$$

We can further change the variable in the subscript of the value function as $V_{\tau, \mathcal{L}}(x)$ thus the value function is given as $V_{\tau, \mathcal{L}}(x) = f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{\tau+1, \mathcal{L}}(x') \mathcal{P}(dx'|x, u) = \bar{\mathbb{P}}^{T-\tau}(x'' \in \mathcal{L}|x, u)$, $x \in \mathbb{X}$.

Thus for a given input policy the probability that a state x transition to a given set $\mathcal{L}$ within a interval $[0, T]$ is given by $\sup_{\tau \in [0, T]} \bar{\mathbb{P}}^\tau(x' \in \mathcal{L}|x, u)$. We will now prove that for the probabilistic avoid problem $\sup_{\tau \in [0, T]} \bar{\mathbb{P}}^\tau(x' \in \mathcal{L}|x, u) = \bar{\mathbb{P}}^T(x' \in \mathcal{L}|x, u)$. For that we need to prove another result shown in Lemma 30. The proof of the following lemma is based on induction.

Lemma 30 *The value function as given in equation 7.2 is a non decreasing function, i.e. $V_{\tau+1, \mathcal{L}}(x) \leq V_{\tau, \mathcal{L}}(x)$, $\forall \tau \in [0, T]$.*

Proof:

$$\text{For } \tau = T, V_{T, \mathcal{L}}(x) - V_{T-1, \mathcal{L}}(x) = f_{\mathcal{L}}(x) - \left(f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{T, \mathcal{L}}(x') \mathcal{P}(dx'|x, u) \right) = -f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{T, \mathcal{L}}(x') \mathcal{P}(dx'|x, u)$$

$$\Rightarrow V_{T, \mathcal{L}}(x) \leq V_{T-1, \mathcal{L}}(x)$$

$$\text{Assume for } k \Rightarrow V_{k+1, \mathcal{L}}(x) \leq V_{k, \mathcal{L}}(x) \Rightarrow V_{k+1, \mathcal{L}}(x) - V_{k, \mathcal{L}}(x) \leq 0$$

$$\begin{aligned} \text{For } k+1 V_{k+2, \mathcal{L}}(x) - V_{k+1, \mathcal{L}}(x) &= \\ & f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{k+1, \mathcal{L}}(x') \mathcal{P}(dx'|x, u) \\ & - f_{\mathcal{L}}(x) - f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} V_{k, \mathcal{L}}(x') \mathcal{P}(dx'|x, u) = \\ & f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} (V_{k+1, \mathcal{L}}(x') - V_{k, \mathcal{L}}(x')) \mathcal{P}(dx'|x, u) \leq 0 \end{aligned}$$

Now using Lemmas 30 and 29 it can be trivially shown that the $\bar{\mathbb{P}}^\tau(x' \in \mathcal{L}|x, u) \leq \bar{\mathbb{P}}^{\tau+1}(x' \in \mathcal{L}|x, u)$. Thus implying that the $\sup_{\tau \in [0, T]} \bar{\mathbb{P}}^\tau(x' \in \mathcal{L}|x, u) = \bar{\mathbb{P}}^T(x' \in \mathcal{L}|x, u)$.

Finally forced backward reachable set is given by equation 7.5 can be obtained by taking the infimum over all the control policies. Thus we obtain equation 7.2 and the value function at $\tau = 0$ provides the probability distribution as $\inf_u \sup_t \bar{\mathbb{P}}^t(x' \in \mathcal{L}|x, u)$.

REFERENCES

- [1] D. Bertsekas. “Infinite time reachability of state-space regions by using feedback control”. In: *IEEE Transactions on Automatic Control* 17.5 (1972), pp. 604–613.
- [2] F. Blanchini. “Controlling systems via set-theoretic methods: some perspectives”. In: *Proceedings of the 45th IEEE Conference on Decision and Control*. IEEE. 2006, pp. 4550–4555.
- [3] F. Blanchini. “Set invariance in control”. In: *Automatica* 35.11 (1999), pp. 1747–1767.
- [4] M. Chen, S. Herbert, and C. J. Tomlin. “Exact and efficient Hamilton-Jacobi guaranteed safety analysis via system decomposition”. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2017, pp. 87–92.
- [5] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert. “Robust control barrier–value functions for safety-critical control”. In: *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE. 2021, pp. 6814–6821.
- [6] I. M. Mitchell and J. A. Templeton. “A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems”. In: *Hybrid Systems: Computation and Control: 8th International Workshop, HSCC 2005, Zurich, Switzerland, March 9-11, 2005. Proceedings 8*. Springer. 2005, pp. 480–494.
- [7] R. van den Brandt and C. de Visser. “Safe flight envelope uncertainty quantification using probabilistic reachability analysis”. In: *IFAC-PapersOnLine* 51.24 (2018), pp. 628–635.
- [8] S. L. Herbert. *Safe real-world autonomy in uncertain and unstructured environments*. University of California, Berkeley, 2020.
- [9] S. Herbert, J. J. Choi, S. Sanjeev, M. Gibson, K. Sreenath, and C. J. Tomlin. “Scalable learning of safety guarantees for autonomous systems using Hamilton-Jacobi reachability”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 5914–5920.
- [10] S. V. Rakovic and K. I. Kouramas. “Invariant approximations of the minimal robust positively invariant set via finite time Aumann integrals”. In: *2007 46th IEEE Conference on Decision and Control*. IEEE. 2007, pp. 194–199.

- [11] M. Rungger and P. Tabuada. “Computing robust controlled invariant sets of linear systems”. In: *IEEE Transactions on Automatic Control* 62.7 (2017), pp. 3665–3670.
- [12] M. A. Hernández-Mejías, A. Sala, C. Arino, and A. Querol. “Reliable controllable sets for constrained Markov-jump linear systems”. In: *International Journal of Robust and Nonlinear Control* 26.10 (2016), pp. 2075–2089.
- [13] M. A. Hernández-Mejías and A. Sala. “Reliability and time-to-failure bounds for discrete-time constrained Markov jump linear systems”. In: *International Journal of Robust and Nonlinear Control* 27.10 (2017), pp. 1773–1791.
- [14] M. Cannon, B. Kouvaritakis, S. V. Raković, and Q. Cheng. “Stochastic tubes in model predictive control with probabilistic constraints”. In: *IEEE Transactions on Automatic Control* 56.1 (2010), pp. 194–200.
- [15] A. Mesbah. “Stochastic model predictive control: An overview and perspectives for future research”. In: *IEEE Control Systems Magazine* 36.6 (2016), pp. 30–44.
- [16] I. M. Mitchell, S. Kaynama, M. Chen, and M. Oishi. “Safety preserving control synthesis for sampled data systems”. In: *Nonlinear Analysis: Hybrid Systems* 10 (2013), pp. 63–82.
- [17] J. Burlet, O. Aycard, and T. Fraichard. “Robust motion planning using markov decision processes and quadtree decomposition”. In: *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA’04. 2004*. Vol. 3. IEEE. 2004, pp. 2820–2825.
- [18] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin. “FaSTrack: A modular framework for fast and guaranteed safe motion planning”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 1517–1522.
- [19] J. Ding, M. Kamgarpour, S. Summers, A. Abate, J. Lygeros, and C. Tomlin. “A stochastic games framework for verification and control of discrete time stochastic hybrid systems”. In: *Automatica* 49.9 (2013), pp. 2665–2674.
- [20] P. M. Esfahani, D. Chatterjee, and J. Lygeros. “The stochastic reach-avoid problem and set characterization for diffusions”. In: *Automatica* 70 (2016), pp. 43–56.
- [21] J.-P. Aubin, G. Da Prato, and H. Frankowska. “Stochastic invariance for differential inclusions”. In: *Set-Valued Analysis* 8 (2000), pp. 181–201.
- [22] H. M. Soner and N. Touzi. “Dynamic programming for stochastic target problems and geometric flows”. In: *Journal of the European Mathematical Society* 4.3 (2002), pp. 201–236.
- [23] H. M. Soner and N. Touzi. “Stochastic target problems, dynamic programming, and viscosity solutions”. In: *SIAM Journal on Control and Optimization* 41.2 (2002), pp. 404–424.

- [24] S. Battilotti and A. De Santis. “Stabilization in probability of nonlinear stochastic systems with guaranteed region of attraction and target set”. In: *IEEE Transactions on Automatic Control* 48.9 (2003), pp. 1585–1599.
- [25] G. Pola, J. Lygeros, and M. D. Di Benedetto. “Invariance in stochastic dynamical control systems”. In: *International Symposium on Mathematical Theory of Networks and Systems*. 2006.
- [26] G. Pola and G. Pola. “A stochastic reachability approach to portfolio construction in finance industry”. In: *IEEE transactions on control systems technology* 20.1 (2011), pp. 189–195.
- [27] M. Cannon, B. Kouvaritakis, and X. Wu. “Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty”. In: *IEEE Transactions on Automatic Control* 54.7 (2009), pp. 1626–1632.
- [28] E. Kofman, J. A. De Doná, and M. M. Seron. “Probabilistic set invariance and ultimate boundedness”. In: *Automatica* 48.10 (2012), pp. 2670–2676.
- [29] E. Kofman, J. A. De Doná, M. M. Seron, and N. Pizzi. “Continuous-time probabilistic ultimate bounds and invariant sets: Computation and assignment”. In: *Automatica* 71 (2016), pp. 98–105.
- [30] A. Abate. *Probabilistic reachability for stochastic hybrid systems: theory, computations, and applications*. University of California, Berkeley, 2007.
- [31] D. Bertsekas and S. E. Shreve. *Stochastic optimal control: the discrete-time case*. Vol. 5. Athena Scientific, 1996.
- [32] M. B. Stinchcombe and H. White. “Some measurability results for extrema of random functions over random sets”. In: *The Review of Economic Studies* 59.3 (1992), pp. 495–514.
- [33] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. “Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems”. In: *Automatica* 44.11 (2008), pp. 2724–2734.
- [34] A. Bhattacharya and J. P. Kharoufeh. “Linear programming formulation for non-stationary, finite-horizon Markov decision process models”. In: *Operations Research Letters* 45.6 (2017), pp. 570–574.
- [35] M. B. Cohen, Y. T. Lee, and Z. Song. “Solving linear programs in the current matrix multiplication time”. In: *Journal of the ACM (JACM)* 68.1 (2021), pp. 1–39.
- [36] C.-S. Chow and J. N. Tsitsiklis. “An optimal one-way multigrid algorithm for discrete-time stochastic control”. In: *IEEE transactions on automatic control* 36.8 (1991), pp. 898–914.
- [37] I. Tkachev and A. Abate. “On infinite-horizon probabilistic properties and stochastic bisimulation functions”. In: *2011 50th IEEE Conference on Decision and Control and European Control Conference*. IEEE. 2011, pp. 526–531.

8

SAFE PROBABLE BACKWARD REACHABLE SETS: THEORY AND ALGORITHMS

In this paper, we extend the existing concept of the stochastic backward reachable set by introducing a novel set termed the 'Safe Probable Backward Reachable Set' ($V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$). Unlike the stochastic backward reachable set, the $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$ is stochastically viable, ensuring that the system remains within this set with a certain probability until it reaches the target set. The advantage of a set defined in this manner is that it can be directly used for safety assurances in stochastic settings. We present two algorithms for computing $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$ and rigorously prove that the sets generated by both algorithms are indeed the $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$. Additionally, we demonstrate that the set produced by Algorithm 8.2 is the largest possible $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$. We implement both algorithms for discrete state-space systems, where they compute the exact $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$. Moreover, we extend these algorithms to continuous state spaces by discretising the state space, and we prove that the sets obtained converge to the true $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$ as the grid size approaches zero. Finally, we analyse the computational costs for both algorithms for both discrete and continuous spaces, showing that Algorithm 8.1 is computationally more efficient than Algorithm 8.2. However, while algorithm 1 produces a set which satisfies the $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$ properties, it is not guaranteed to be the largest possible set.


8.1. INTRODUCTION

For many automated systems, the guarantee of their safe operation is integral to their accepted use in society. Concepts such as invariance, viability [1–3] and reachability [4–6] play a central role in safety assurance and in the identification of safe operating regions [7]. One such safe operating region is the so-called backward reachable set. This set is a region in state space from where the system has the control authority to reach a particular set, often referred to as the target set. Thorough research has been conducted on the backward reachable set [8], [9] and the viable set [10],[11] in both deterministic (i.e. no disturbance) and robust (i.e. worst-case disturbance) settings. These sets in robust settings are inherently more conservative. The conservatism of these sets can be reduced through the stochastic counterpart which allows stochastic violations. The stochastic counterpart of the backward reachable set is the region in state space from where the probability of the system to reach the target set is higher than a particular probability value [12], [13], [14]. Furthermore, in the robust context, the worst case realisation of the disturbance is assumed which forces the disturbance to be bounded. Similarly, in deterministic settings, the disturbance is neglected entirely. In contrast to these approaches, the stochastic counterpart of the backward reachable set permits cases with unbounded disturbances. Stochastic reachability provides an established mathematical framework to address the question of the stochastic counterpart of the backward reachable set [15], [12]. In this article we will call the stochastic counterpart of the backward reachable set the 'stochastic backward reachable set' (SBRS). This set has applications in guaranteeing the safety and performance of the stochastic dynamical system in the fields of robotics [12, 14–19], biomedical applications [20], [21], and spacecraft [22],[23].


8

The formal definition of the SBRS is as follows: Given a time horizon T , a target set $\mathcal{L}$, and a probability value $\mathcal{J}$, the SBRS is the set of all states for which there exists an input policy u such that the probability of the system to reach $\mathcal{L}$ within the time horizon T is equal to or greater than $\mathcal{J}$ [12],[14]. Although this definition looks like a natural counterpart of the deterministic or robust backward reachable set, the set suffers from a crucial drawback. Consequently, it cannot be directly utilised to guarantee the safety of systems. As illustrated by figure 8.1, in the case of the non-stochastic sets, the system can remain within the backward reachable set at all times during its trajectory towards the target set. However, in the case of the SBRS, it not possible to guarantee that the system remains within the SBRS at all times during its trajectory towards the target set. In the example shown in figure 8.1b, there is a probability that the system leaves the SBRS during the transition. The system's safety cannot be guaranteed while it is outside the SBRS.


Similar to the SBRS, there are some other stochastic extensions of sets similar to backward reachable sets such as the stochastic reach-avoid problem [24]. For such sets, the system has certain target states and avoid states with a certain probability.



(a) Backward Reachable Set (BRS)



(b) Stochastic Backward Reachable Set (SBRS)



(c) Safe Probable Backward Reachable Set (SPBRS)

Figure 8.1: In the BRS, trajectories leading to $\mathcal{L}$ stay inside the BRS for the whole time horizon. The SBRS provides us with the net probability of a state transitioning to $\mathcal{L}$. The SPBRS provides us with the probability of the set transitioning to $\mathcal{L}$ while staying in a certain region.

The solution of the stochastic reach-avoid problem is the set of initial states from which the system can reach the target states while avoiding the avoid states. Another method developed in the context of a stochastic reach avoid problem is based on stochastic contingent sets [25]. Geared towards a similar objective of almost-certain reachability, [26] introduced a new class of stochastic target problems and characterized their solution using dynamic programming. The authors of [26] also examined the differential properties of the almost-certain reachable set based on a geometric partial differential equation analogous to the HJB equation [27] in that context. However, none of these extensions are directly applicable to problems where the avoid set is unavailable or is constantly changing. We therefore introduce an alternative stochastic extension of the SBRS that guarantees system safety even when the avoid set is unknown.

The contributions of this paper are as follows: First, we extend the existing definition of the SBRS and introduce a new set which we call the 'Safe probable backward reachable set' ($V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$). Given a time horizon T , a target set $\mathcal{L}$, and a probability value $\mathcal{J}$, a set is called a $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ if for all states inside the set there exists an input policy u such that the probability of the system to reach $\mathcal{L}$ within the time horizon T while staying inside $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ is equal to or greater than $\mathcal{J}$. *Note: We assume that the time step $\tau \in \mathbb{N}$.* The definition is applicable to discrete time stochastic systems. The advantage of a set defined in this manner is that it can be directly used for safety assurances in stochastic settings. Secondly, we provide two algorithms to obtain $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ and formally prove that the sets produced by both the algorithms are indeed $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. We further prove that one of these algorithms in fact produces the largest possible $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. Then we prove that for the discrete time systems, the algorithms can be implemented using linear programming. For discrete space systems, both algorithms compute an exact $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. Furthermore, we provide an extension of both algorithms for continuous space systems. This extension involves a discretisation of the state space. Subsequently, we prove that the sets obtained via these algorithms for the discretised space, despite being estimates, converge to the actual $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ as the grid size approaches zero. Moreover, we provide the computational cost for both algorithms for both discrete and continuous spaces and show that Algorithm 8.1 is computationally less expensive than Algorithm 8.2.

The paper is motivated by problems of reliable control [28], [29], stochastic MPC [30], [31], safety-critical control [32] and related applications such as path planning [33], obstacle avoidance [34] and air traffic control [35].

The paper is structured as follows: In Section 8.2, we outline the problem statement and some preliminary information used throughout the paper. In Section 8.3, we provide Algorithm 8.1 and rigorously prove that the set obtained is indeed $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. Likewise, in Section 8.4 we provide Algorithm 8.2 and rigorously prove that the set obtained is indeed $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$, and that the set estimated using Algorithm 8.2 is indeed the biggest possible $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. Section 8.5 is divided one subsection per algorithm where we prove that for the discrete time systems, the algorithms can be implemented using linear

programming. In Section 8.6 we explain how we discretise the state space and provide the continuous space implementation of Algorithm 8.1 and Algorithm 8.2 respectively in two subsections. This is followed by a conclusion in Section 8.8.

8.2. SETTING AND PROBLEM STATEMENT

In this section, we provide the necessary definitions and preliminary information for setting up the proofs. We consider a stochastic Markov control system denoted as $S = (\mathbb{X} = \mathbb{R}^n, \mathbb{U} = \mathbb{R}^m, \mathcal{P})$, where:

- $\mathbb{X}$ represents the state space, equipped with the corresponding Borel σ -algebra $\mathcal{B}(\mathbb{X})$.
- $\mathbb{U}$ is the control space, accompanied by the Borel σ -algebra $\mathcal{B}(\mathbb{U})$.
- $\mathcal{P}$ is a Borel-measurable stochastic kernel. Given $(x, u) \in (\mathbb{X} \times \mathbb{U})$, it assigns to each pair (x, u) a probability measure on the Borel space $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$. Formally, the probability measure on the Borel space $(\mathbb{X}, \mathcal{B}(\mathbb{X})) : \mathcal{P}(\cdot|x, u)$ can be defined as a function: $\mathcal{P} : \mathcal{B}(\mathbb{X}) \times \mathbb{X} \times \mathbb{U} \rightarrow [0, 1]$
- $n, m \in \mathbb{N}$ where $\mathbb{N}$ is the set of natural numbers.

We assume that the $\mathbb{U}_x$ is a non-empty subset of $\mathbb{U}$ of all the control actions for each $x \in \mathbb{X}$. Given a finite time horizon $T \in \mathbb{N}$, a control policy is said to be a Markov control policy if it only depends on the current state, i.e., $u_k(x(k)) = u_k(x(k), x(k-1), \dots)$ for $k \in \mathbb{N}$.

Definition 1 (Markov control policy): A Markov control policy u for a given system S is a sequence $u = (u_1, u_2, \dots, u_{N-1})$ of universally measurable maps.

$$u : \mathbb{N} \times \mathbb{X} \rightarrow \mathbb{U}, \quad \forall k \in \mathbb{N}$$

For brevity $u_k(x(k))$ will be represented as u . *Note 1:* A subset $\mathcal{A}$ of a space $\mathbb{Z}$ is universally measurable if the set $\mathcal{A}$ is measurable with respect to all complete measures on space $\mathbb{Z}$ that measure all Borel sets of space $\mathbb{Z}$ ($\mathcal{B}(\mathbb{Z})$). A function $f : \mathbb{Z} \rightarrow \mathbb{Y}$ is universally measurable if the inverse image ($f^{-1}(\mathcal{A})$) of all the sets $\mathcal{A} \in \mathcal{B}(\mathbb{Y})$ is universally measurable. For showing the existence of a solution to a stochastic optimal problem the condition of universal measurability is a weaker condition than the Borel measurability [15],[36],[37].

Note 2: In a substantial subset of stochastic optimal control problems, the utilization of Markov policies proves to be adequate in delineating the optimal policy [36]. Given that the adoption of a randomized Markov policy does not contribute to the increase in the maximum probability of the state staying within a designated set or reaching a particular set, our focus in this paper will be on deterministic Markov policies. We denote the set of all Markov policies by $\mathcal{M}$.

Given a control policy $u \in \mathcal{M}$, a horizon $\tau \in \mathbb{N}$ and a state $x \in \mathbb{X}$, the probability represented as $\mathcal{P}(x^\tau|x, u)$ is read as the probability to transition to

state x^τ at step τ , given that the state at step 0 is x , under the control policy u , i.e. $\mathcal{P}(X^\tau = x^\tau | X^0 = x, u)$, $\tau \in \mathbb{N}$. Given a control policy $u \in \mathcal{M}$, a time step $\tau \in \mathbb{N}$, a state $x \in \mathbb{X}$ and a set $A \in \mathcal{B}(\mathbb{X})$, the probability $\mathcal{P}(x^\tau \in A | x, u)$ is read as the probability that the state x^τ is in set A , given that the state at step 0 is x , under the control policy u , i.e. $\mathcal{P}(X^\tau = x^\tau \in A | X^0 = x, u)$, $\tau \in \mathbb{N}$. For brevity the transition probability for one time step is represented by omitting the time step from x^τ to simply x , for example $\mathcal{P}(x | x', u)$.

Given a target set of states ($\mathcal{L}$), a time horizon ($[0, T]$) and a probability value ($\mathcal{J}$), the Safe Probable Backward Reachable Set ($V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J})$) is defined as the set of states ($x \in \mathbb{X}$) for which there exists a control action ($u \in \mathcal{M}$), a time step ($k \in [0, T] \subseteq \mathbb{N}$) such that the probability of the state to transition to the set $\mathcal{L}$, while staying in the $V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J})$ for all time steps is greater than or equal to $\mathcal{J}$. The set is defined in set builder notation in equation 8.1.

$$\begin{aligned} V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}) &= \{x : \exists u \in \mathcal{M}, k \in [0, T] \subseteq \mathbb{N}, \mathcal{P}(x^k \in \mathcal{L}, x^k \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), \\ &\quad x^{k-1} \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), \dots, |x^0 \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J})) \geq \mathcal{J} \\ &= \sup_{u \in \mathcal{M}} \sup_{k \in \mathbb{N}_{[0, T]}} \mathcal{P}(x^k \in \mathcal{L}, x^k \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), x^{k-1} \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), \dots, \\ &\quad x^1 \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}) | x^0 \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J})) \geq \mathcal{J} \} \quad (8.1) \end{aligned}$$

In the context of probable reach problems for the target set ($\mathcal{L}$), we assume that once the state is in the target set $\mathcal{L}$, it will remain in $\mathcal{L}$ for the whole time horizon with a probability of one, i.e., if $x^0 \in \mathcal{L} \Rightarrow \mathcal{P}(x^T \in \mathcal{L}, x^{T-1} \in \mathcal{L}, \dots, x^1 \in \mathcal{L} | x^0, u), \forall T \in \mathbb{N}, \forall u \in \mathcal{U}$. Thus, if the system has the Markov property then

$$\begin{aligned} &\mathcal{P}(x^T \in \mathcal{L}, x^{T-1} \in \mathcal{L}, \dots, x^{k+1} \in \mathcal{L}, x^k \in \mathcal{L}, x^k \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), \\ &\quad x^{k-1} \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), \dots, |x^0 \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J})) \\ &= \mathcal{P}(x^T \in \mathcal{L}, x^{T-1} \in \mathcal{L}, \dots, x^{k+1} \in \mathcal{L} | x^k \in \mathcal{L}, u) \cdot \mathcal{P}(x^k \in \mathcal{L}, x^k \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), \\ &\quad x^{k-1} \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), \dots, |x^0 \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J})) \\ &= \mathcal{P}(x^k \in \mathcal{L}, x^k \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), x^{k-1} \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J}), \dots, |x^0 \in V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J})) \end{aligned}$$

Furthermore, we define an indicator function of a Borel set $A \in \mathcal{B}(\mathbb{X})$, $f_A : \mathbb{X} \rightarrow \{0, 1\}$ as shown in equation 8.2.

$$f_A(x) = \begin{cases} 1, & \text{if } x \in A \\ 0, & \text{if } x \notin A \end{cases} \quad (8.2)$$

Similarly, we define an inclusion function $\Pi_a : \mathbb{X} \rightarrow \mathbb{R}$, $0 \leq a \in \mathbb{R}$ as shown in equation 8.3, where $V(x)$ is a function mapping $\mathbb{X}$ to $\mathbb{R}$

$$\Pi_a(V(x)) = \begin{cases} V(x), & \text{if } V(x) \geq a \\ 0, & \text{if } V(x) < a \end{cases} \quad (8.3)$$

Furthermore, we assume that $\mathcal{P}(x^T \in \mathcal{L} | x^T, u) = f_{\mathcal{L}}(x^T)$.

8.3. AN ALGORITHM TO COMPUTE A SAFE PROBABLE BACKWARD REACHABLE SET




Figure 8.2: Explanation of Algorithm 8.1 in four steps.

In this section we develop an algorithm that provides us with a $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. Among the two algorithms presented in the paper, this one is computationally less expensive. The algorithm finds the set using backward recursion, i.e., we start from the target set and move backwards in time to find the $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. Subsequently we prove that the set produced by the algorithm satisfies the properties of the $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. The algorithm works according to the following process: It sets the values of all the points in the target set as equal to 1 and the values of all other points equal to zero. Then it calculates the 1-step transition probability from every point in our domain to every other point. The value of each of these points in the next time step is equal to the product of their transition probabilities and the value assigned to them at the previous time step. Figure 8.2a shows the points with a probability of greater than $\mathcal{J}$ to transition into the target set in green and the others in red. Subsequently, all points with

a value of less than $\mathcal{J}$ are set as equal to 0 (Figure 8.2b). Thus, we receive a set from any point of which the system can transition into the target set with a probability of greater than $\mathcal{J}$. Then the process is repeated until the end of the time horizon (Figures 8.2c and 8.2d). In essence, with every time step we receive a set from which the system can transition to the previous set with a probability of greater than $\mathcal{J}$.

Using the indicator function (equation 8.2) and inclusion function (equation 8.3), we can define a value function $V_{T-k, \mathcal{L}} : \mathbb{X} \rightarrow [0, 1]$, $k \in [0, T] \subseteq \mathbb{N}$ by the process of backward recursion as shown in equation 8.4, where $\bar{\mathcal{L}}$ is the complement of set $\mathcal{L}$ in $\mathbb{X}$.

$$V_{T-k}(\mathbf{x}) = \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(\mathbf{x}) + f_{\bar{\mathcal{L}}}(\mathbf{x}) \int_{\mathbf{x}' \in \mathbb{X}} \Pi_{\mathcal{J}}(V_{T-k+1}(\mathbf{x}')) \mathcal{P}(d\mathbf{x}' | \mathbf{x}, u)),$$

$$\mathbf{x} \in \mathbb{X}, \mathcal{J} \in [0, 1] \quad (8.4)$$

Where $V_T(\mathbf{x}) = f_{\mathcal{L}}(\mathbf{x})$ and using the value function we define a set as shown in equation 8.5.

$$S_k = \{\mathbf{x} : V_{T-k}(\mathbf{x}) \geq \mathcal{J}\} \quad (8.5)$$

Assumption 13 *The set given by equation 8.6*

$$\mathbb{U}_{\mathbf{x}}(\tau, \mathcal{J}) = \{u \in \mathcal{M} \mid \int_{\mathbb{X}} \Pi_{\mathcal{J}}(V_{\tau+1, \mathcal{L}, \mathcal{J}}(\mathbf{x}')) \mathcal{P}(d\mathbf{x}' | \mathbf{x}, u) \geq \mathcal{J}\} \quad (8.6)$$

is compact for all $\mathbf{x} \in \mathbb{X}$, $\tau \in \mathbb{N}_{[0, T]}$, $\mathcal{J} \in [0, 1]$.

Thus, if Assumption 13 holds, then $\forall \mathbf{x} \in \mathbb{X}$, the existence of a Markov policy that satisfies equation 8.4 is guaranteed and is given by equation 8.7.

$$u_{\tau, \text{SPBRS}}^*(\mathbf{x}) = \arg \sup_{u \in \mathbb{U}} f_{\bar{\mathcal{L}}}(\mathbf{x}) \int_{\mathbb{X}} \Pi_{\mathcal{J}}(V_{\tau+1, \mathcal{L}, \mathcal{J}}(\mathbf{x}')) \mathcal{P}(d\mathbf{x}' | \mathbf{x}, u),$$

$$\mathbf{x} \in \mathbb{X} \setminus \mathcal{L}, \tau \in \mathbb{N}_{[0, T-1]} \quad (8.7)$$

Algorithm 8.1 is used to obtain the Safe Probable Backward Reachable Set.

Now we will prove that the S_k obtained using algorithm 8.1 is indeed a $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. For this we need the following two lemmas.

Lemma 31 *The value function in equation 8.4 is a non-decreasing function with k i.e., $V_{T-k}(\mathbf{x}) \leq V_{T-k-1}(\mathbf{x})$.*

Proof: We will use induction to prove the result. We know $V_T(\mathbf{x}) = f_{\mathcal{L}}(\mathbf{x})$, then based on the definition of the value function (equation 8.4) it is trivial

Algorithm 8.1: An algorithm to obtain a Safe Probable Backward Reachable Set ($V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$)

Require: $\mathcal{L}, \mathcal{J}, T$

$V_T(x) \leftarrow f_{\mathcal{L}}(x)$ equation 8.2

$k \leftarrow 1$

while $k \leq T$ **do**

 Obtain $V_{T-k}(x)$ using equation 8.4

if $V_{T-k}(x) = V_{T-k+1}(x)$ **then**

$V_0(x) \leftarrow V_{T-k}(x)$ using Theorem 15

$k \leftarrow T$

 Obtain S_k using equation 8.5

$k \leftarrow k + 1$

$V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J}) \leftarrow S_k$ using Theorem 16

to see that $V_T(x) \leq V_{T-1}(x)$. Furthermore, $V_T(x) \leq V_{T-1}(x) \Rightarrow \Pi_{\mathcal{J}}(V_T(x)) \leq \Pi_{\mathcal{J}}(V_{T-1}(x))$ (using equation 8.3). For $k = 1$ we have:

$$\begin{aligned} V_{T-1}(x) &= \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} \Pi_{\mathcal{J}}(V_T(x')) \mathcal{P}(dx'|x, u)) \\ &\leq \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} \Pi_{\mathcal{J}}(V_{T-1}(x')) \mathcal{P}(dx'|x, u)) = V_{T-2}(x) \\ &\Rightarrow V_{T-1}(x) \leq V_{T-2}(x) \text{ Hence true for } k = 1. \end{aligned}$$

Assume for $k = n < T \Rightarrow V_{T-n}(x) \leq V_{T-n-1}(x)$

By using equation 8.3 $\Rightarrow \Pi_{\mathcal{J}}(V_{T-n}(x)) \leq \Pi_{\mathcal{J}}(V_{T-n-1}(x))$

for $k = n + 1$

$$\begin{aligned} V_{T-n-1} &= \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} \Pi_{\mathcal{J}}(V_{T-n}(x')) \mathcal{P}(dx'|x, u)) \\ &\leq \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) \int_{x' \in \mathbb{X}} \Pi_{\mathcal{J}}(V_{T-n-1}(x')) \mathcal{P}(dx'|x, u)) = V_{T-n-2}(x) \end{aligned}$$

Lemma 31 states that the value function cannot decrease as we go backwards in time (i.e., as the time horizon increases). Since we start from set $\mathcal{L}$ and work our way backwards in time, the value function can only expand or stay the same as we proceed through the time step index k .

Theorem 15 $\sup_{k \in [0, T]} V_{T-k}(x) = V_0(x)$

Proof: The proof directly follows from Lemma 31.

Corollary 3 *The sets (S_k) produced by algorithm 8.1 are a non-decreasing sequence of sets, i.e., $S_k \subseteq S_{k+1} \forall k \in [0, T]$ and $S_0 = \mathcal{L}$.*

Proof: The proof is trivial and follows directly from the definition of set S_k (equation 8.5) and Lemma 31.

In the next lemma we prove that for any $k \in [0, T]$, the $V_{T-k}(x)$ is less than the joint probability of state x transitioning into the set $\mathcal{L}$ at time step k while staying in set S_k for all time steps until k .

Lemma 32 $V_{T-k}(x^0) = \sup_{u \in \mathcal{M}} \mathcal{P}(x^k \in \mathcal{L}, x^k \in S_0, x^{k-1} \in S_1, \dots, x^1 \in S_{k-1} | x^0, u)$ for any $k \in [1, T]$ and $x^0 \in S_k$

Proof: We use induction to prove the result for all $x \in S_k$.

For $k = 0$ we have: $V_{T-k}(x) = f_{\mathcal{L}}(x)$

and we know that

$$\mathcal{P}(x \in \mathcal{L} | x, u) = f_{\mathcal{L}}(x)$$

Thus, $V_T(x) = \mathcal{P}(x \in \mathcal{L} | x, u)$

For $k = 1$ we have

$$V_{T-1}(x) = \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) \int_{x_1 \in \mathbb{X}} \Pi_{\mathcal{J}}(V_T(x_1)) \mathcal{P}(dx_1 | x, u))$$

Since $\Pi_{\mathcal{J}}(V_T(x)) = 0$ for all $x \notin S_0$

$$\begin{aligned} &= \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) \int_{x_1 \in S_0} V_T(x_1) \mathcal{P}(dx_1 | x, u)) \\ &= \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) \int_{x_1 \in S_0} \mathcal{P}(x_1 \in \mathcal{L} | x_1, u) \mathcal{P}(dx_1 | x, u)) \\ &= \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) \mathcal{P}(x_1 \in \mathcal{L}, x_1 \in S_0 | x, u)) \end{aligned}$$

Since we know that $x \in S_1$ and $\mathcal{L} \subseteq S_1$

$$= \sup_{u \in \mathcal{M}} \mathcal{P}(x_1 \in \mathcal{L}, x_1 \in S_0 | x, u)$$

Assume for $k = n$:

$$V_{T-n}(x) = \sup_{u \in \mathcal{M}} (\mathcal{P}(x^n \in \mathcal{L}, x^n \in S_0, x^{n-1} \in S_1, \dots, x^1 \in S_{n-1} | x, u))$$

For $k = n + 1$

$$\begin{aligned} V_{T-n-1}(x) &= \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{x_1 \in \mathbb{X}} \Pi_{\mathcal{J}}(V_{T-n}(x_1)) \mathcal{P}(dx_1 | x, u)) \\ &= \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{x_1 \in S_n} V_{T-n}(x_1) \mathcal{P}(dx_1 | x, u)) \\ &= \sup_{u \in \mathcal{M}} (f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) \int_{x_1 \in S_n} \sup_{u \in \mathcal{M}} (\mathcal{P}(x^{n+1} \in \mathcal{L}, \\ &\quad x^{n+1} \in S_0, x^n \in S_1, \dots, x^2 \in S_{n-1} | x^1, u)) \mathcal{P}(dx_1 | x, u)) \end{aligned}$$

We know that the system has Markov property and, furthermore, that the probability is bounded and always non-negative. Thus, following the same logic as above, we get

$$= \sup_{u \in \mathcal{M}} (\mathcal{P}(x^{n+1} \in \mathcal{L}, x^{n+1} \in S_0, x^n \in S_1, \dots, x^1 \in S_n | x, u))$$

Corollary 4 $V_{T-k}(x) \leq \sup_{u \in \mathcal{M}} \mathcal{P}(x^k \in \mathcal{L}, x^k \in S_k, x^{k-1} \in S_k, \dots, x^1 \in S_k | x, u)$ for any given $x \in S_k$

Proof: The proof directly follows from the fact that, using Corollary 3, one can say $\mathcal{P}(x^k \in \mathcal{L}, x^k \in S_0, x^{k-1} \in S_1, \dots, x^1 \in S_{k-1} | x, u) \leq \mathcal{P}(x^k \in \mathcal{L}, x^k \in S_k, x^{k-1} \in S_k, \dots, x^1 \in S_k | x, u)$ for any $u \in \mathcal{M}$. Thus, using Lemma 4 implies $V_{T-k}(x) \leq \sup_{u \in \mathcal{M}} \mathcal{P}(x^k \in \mathcal{L}, x^k \in S_k, x^{k-1} \in S_k, \dots, x^1 \in S_k | x, u)$.

Finally, now we have all necessary intermediate results to prove that the set obtained using algorithm 8.1 is a $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$.

Theorem 16 The set S_T is a Safe Probable Backward Reachable Set $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$, i.e., $\forall x \in S_T$ the $\sup_{u \in \mathcal{M}} \sup_{k \in [0, T]} \mathcal{P}(x^k \in \mathcal{L}, x^k \in S_T, x^{k-1} \in S_T, \dots, x^1 \in S_T | x, u) \geq \mathcal{J}$

Proof: Using Corollary 4 we can say that for any $k \in [0, T]$

$$V_{T-k}(x) \leq \sup_{u \in \mathcal{M}} \mathcal{P}(x^k \in \mathcal{L}, x^k \in S_k, x^{k-1} \in S_k, \dots, x^1 \in S_k | x, u)$$

$$\begin{aligned} \Rightarrow \sup_{k \in [0, T]} V_{T-k}(x) &\leq \sup_{k \in [0, T]} \sup_{u \in \mathcal{M}} \\ &\quad \mathcal{P}(x^k \in \mathcal{L}, x^k \in S_k, x^{k-1} \in S_k, \dots, x^1 \in S_k | x, u) \end{aligned}$$

By using Theorem 15

$$\Rightarrow V_0(x) \leq \sup_{u \in \mathcal{M}} \sup_{k \in [0, T]} \mathcal{P}(x^k \in \mathcal{L}, x^k \in S_k, x^{k-1} \in S_k, \dots, x^1 \in S_k | x, u)$$

Since $\forall x \in S_k$, we have $\mathcal{J} \leq V_0(x)$

$$\Rightarrow \mathcal{J} \leq \sup_{u \in \mathcal{M}} \sup_{k \in [0, T]} \mathcal{P}(x^k \in \mathcal{L}, x^k \in S_k, x^{k-1} \in S_k, \dots, x^1 \in S_k | x, u)$$


Algorithm 8.1 produces a set S_T . Through theorem 16, we have shown that S_T is indeed a $V_{SPBRS}(\cdot)$. Thus, algorithm 8.1 produces a valid $V_{SPBRS}(\cdot)$. However, we do not yet have a guarantee that the obtained set is the largest possible set. To address this, we propose algorithm 8.2.

8.4. AN ALGORITHM TO COMPUTE THE LARGEST POSSIBLE SAFE PROBABLE BACKWARD REACHABLE SET


8

In this section, we develop an algorithm that provides the largest possible $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$. The algorithm finds this set using multiple loops of backward recursion. The process works in the following way: First, the algorithm obtains the stochastic backward reachable set (S_0^1 in figure 8.3) of the target set $\mathcal{L}$. This contains all points from which a system can reach the target set with a probability of greater than $\mathcal{J}$ (the green points in figure 8.3a). In the next iteration, the points outside this set (red in the figure) are given the value 0. Then the algorithm again finds the SBRS of the target set while keeping the points outside the earlier set fixed to 0. Thus, the next SBRS will be smaller (set S_0^2 in figure 8.3b) because it does not contain the points that require the system to leave the SBRS. We then repeat the iterations until the domain converges, implying that the obtained set is the $V_{SPBRS}(\cdot)$. In essence, the algorithm reduces the domain through which the system is allowed to transition, thus restricting it from moving out of S_0^1 . The set obtained in this manner is the largest possible $V_{SPBRS}(\cdot)$ because for the set to be equal to a $V_{SPBRS}(\cdot)$ it needs to converge.


In theorem 19 we prove that the set obtained is indeed a $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$. Subsequently, in theorem 20 we prove that the set obtained in this manner is indeed the largest possible $V_{SPBRS}(T, \mathcal{L}, \mathcal{J})$.



(a) Algorithm 8.2 explanation Step 1



(b) Algorithm 8.2 explanation Step 2



(c) Algorithm 8.2 explanation Step 3

Figure 8.3: In Step 8.3a, we obtain the probability of transition to $\mathcal{L}$ and set each state's value accordingly. If the value is less than $\mathcal{J}$, it is set to zero, yielding the set S_0^1 . Step 8.3b computes the value for transitioning to $\mathcal{L}$ while staying within S_0^1 , zeroing out values outside it, producing S_0^2 . This process repeats in Step 8.3c.

For the second algorithm we define a value function by using equations 8.8 and 8.9. Here, n is the number of times the whole iteration has been carried out while k is used for time-indexing.

$$V_{T-k}^n(x) = \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) f_{S^n}(x) \int_{x' \in \mathbb{X}} f_{S^n}(x') V_{T-k+1}^n(x') \mathcal{P}(dx'|x, u) \quad (8.8)$$

$$\text{Where } V_T^n = f_{\mathcal{L}}(x) \quad (8.9)$$

Furthermore, we define the set using equations 8.10 and 8.11.

$$S^n = \{x : V_0^{n-1}(x) \geq \mathcal{J}\} \quad (8.10)$$

$$\text{Where } S^0 = \mathbb{X} \quad (8.11)$$

Assumption 14 The set given by equation 8.12

$$\mathbb{U}_x(\tau, \mathcal{J}) = \{u \in \mathcal{M} | f_{S^n}(x) \int_{x' \in \mathbb{X}} f_{S^n}(x') V_{T-k+1}^n(x') \mathcal{P}(dx' | x, u) \geq \mathcal{J}\} \quad (8.12)$$

is compact for all $x \in \mathbb{X}$, $\tau \in \mathbb{N}_{0,T}$, $\mathcal{J} \in [0, 1]$.

Thus, if Assumption 14 holds, then $\forall x \in \mathbb{X}$, the existence of a Markov policy that satisfies equation 8.4 is guaranteed and is given by equation 8.13.

$$u_{k, \text{SPBRS}}^*(x) = \arg \sup_{u \in \mathbb{U}} f_{\hat{\mathcal{L}}}(x) f_{S^n}(x) \int_{x' \in \mathbb{X}} f_{S^n}(x') V_{T-k+1}^n(x') \mathcal{P}(dx' | x, u), x \in \mathbb{X} - \mathcal{L}, k \in \mathbb{N}_{[0, T-1]} \quad (8.13)$$

Algorithm 8.2 obtains the largest possible Safe Probable Backward Reachable Set.

Algorithm 8.2: An algorithm to obtain the largest Safe Probable Backward Reachable Set ($V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$)

Require: $\mathcal{L}, \mathcal{J}, T$

```

 $n \leftarrow 0$ 
 $S^0 \leftarrow \mathbb{X}$  As in equation 8.11
 $\alpha = 0$  Break condition for outer-loop
while  $\alpha \neq 1$  do
     $V_T^n(x) \leftarrow f_{\mathcal{L}}(x)$  As in equation 8.2
     $k \leftarrow 1$ 
    while  $k \leq T$  do
        Obtain  $V_{T-k}^n(x)$  Using equation 8.8
        if  $V_{T-k}^n(x) = V_{T-k+1}^n(x)$  then
             $V_0^n(x) \leftarrow V_{T-k}^n(x)$  Theorem 17
             $k \leftarrow T$ 
         $k \leftarrow k + 1$ 
     $n \leftarrow n + 1$ 
    Obtain  $S^n$  Using equation 8.10
    if  $S^n = S^{n-1}$  then
         $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J}) \leftarrow S^n$  Theorem 19
         $\alpha \leftarrow 1$ 

```

To prove that the set obtained using this algorithm is a $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$, as shown in Theorem 19, we firstly need Theorem 18, which requires Lemma 34, and Theorem 17, which requires Lemma 33. We furthermore need Lemma 35 and Lemma 36.

Lemma 33 *The value function in equation 8.8 is a non-decreasing function with k i.e., $V_{T-k}^n(x) \leq V_{T-k-1}^n(x)$.*

Proof: The proof is similar to the proof in Lemma 31.

Corollary 5 *The set $\mathcal{L}$ is always a subset of the set S^n i.e., $\mathcal{L} \subseteq S^n$*

Proof: The proof is trivial and directly follows from the fact that $\mathcal{J} \in [0, 1]$, $V_T^n(x) = f_{\mathcal{L}}(x)$, and from Lemma 33.

Theorem 17 $\sup_{k \in [0, T]} V_{T-k}^n(x) = V_0^n(x)$

Proof: The proof directly follows from Lemma 33.

Theorem 17 has proven that the value function is non-decreasing and thus a supremum over time is equal to the value function at the final time.

Lemma 34 *For a given $u \in \mathcal{M}$ and any $k \in [0, T]$, $\sup_{u \in \mathcal{M}} \mathcal{P}(x^k \in \mathcal{L}, x^k \in S^n, x^{k-1} \in S^n, \dots, x^1 \in S^n | x^0 \in S^n, u) = V_{T-k}^n(x^0)$ for all $x^0 \in S^n$.*

Proof: The proof uses induction. First of all, $k = 0$, $\mathcal{P}(x^0 \in \mathcal{L} | x^0 \in S^n, u) = f_{\mathcal{L}}(x^0) = V_T^n(x^0)$. This follows from the assumption that the system is Markov and from Corollary 5. For readability we will represent x^0 by x for this proof.

$$\begin{aligned}
 \text{For } k = 1 : V_{T-1}^n(x) &= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) f_{S^n}(x) \int_{x^1 \in \mathbb{X}} f_{S^n}(x^1) V_T^n(x^1) \mathcal{P}(dx^1 | x, u) \\
 &= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + \int_{x^1 \in S^n} \mathcal{P}(x^1 \in \mathcal{L} | x^1, u) \mathcal{P}(dx^1 | x \in S^n \setminus \mathcal{L}, u) \\
 &= \sup_{u \in \mathcal{M}} \mathcal{P}(x^1 \in \mathcal{L}, x^1 \in S^n | x \in S^n, u)
 \end{aligned}$$

Since the relation holds for $k = 1$, we assume for $k = m$

$$V_{T-m}^n(x) = \sup_{u \in \mathcal{M}} \mathcal{P}(x^m \in \mathcal{L}, x^m \in S^n, x^{m-1} \in S^n, \dots, x^1 \in S^n | x \in S^n, u)$$

Then for $k = m + 1$ we have that

$$\begin{aligned}
 V_{T-m-1}^n(x) &= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) f_{S^n}(x) \int_{x' \in \mathbb{X}} f_{S^n}(x') V_{T-m}^n(x') \mathcal{P}(dx' | x, u) \\
 &= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) f_{S^n}(x) \int_{x^1 \in S^n} \sup_{u \in \mathcal{M}} \mathcal{P}(x^{m+1} \in \mathcal{L}, x^{m+1} \in S^n, x^m \in S^n, \dots, \\
 &\quad x^2 \in S^n | x^1, u) \mathcal{P}(dx^1 | x, u)
 \end{aligned}$$

Since we know that the system has Markov property and that the probability is bounded and always non-negative, following a similar logic as above, we get

$$= \sup_{u \in \mathcal{M}} \mathcal{P}(x^{m+1} \in \mathcal{L}, x^{m+1} \in S^n, x^m \in S^n, \dots, x^2 \in S^n, x^1 \in S^n | x \in S^n, u)$$

Theorem 18 *The value function given by equation 8.8 provides the probability that there exists a control $u \in \mathcal{M}$ and there exists a time $k \in [0, T]$ such that a state inside set S^n transitions to set $\mathcal{L}$, while remaining in set S^n . i.e., $\sup_{u \in \mathcal{M}} \sup_{k \in [0, T]} \mathcal{P}(x^k \in \mathcal{L}, x^k \in S^n, x^{k-1} \in S^n, \dots, x^1 \in S^n | x^0 \in S^n, u) = V_0^n(x^0)$.*

Proof: The proof directly follows from Theorem 17 and Lemma 34.

Lemma 35 *If $V_0^{n-1}(x) = V_0^n(x) \Rightarrow S^n = S^{n+1}$*

Proof: The proof directly follows from the definition of set S^n given in equation 8.10. Since $S^n = \{x : V_0^{n-1}(x) \geq \mathcal{J}\} = \{x : V_0^n(x) \geq \mathcal{J}\} = S^{n+1}$

Lemma 36 *If $S^n = S^{n+1} \Rightarrow V_0^n(x) = V_0^{n+1}(x)$*

Proof: We will use induction for this proof. We are given that $S^n = S^{n+1}$. We know from equation 8.9 that $V_T^n(x) = V_T^{n+1}(x)$ and we know that the set $\mathcal{L}$ does not change in different time steps.

$$\begin{aligned} \text{For } k = 1, V_{T-k}^n &= V_{T-1}^n = \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + \\ &\quad f_{\bar{\mathcal{L}}}(x) f_{S^n}(x) \int_{x' \in \mathbb{X}} f_{S^n}(x') V_T^n(x') \mathcal{P}(dx' | x, u) \\ &= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) f_{S^n}(x) \int_{x' \in S^n} V_T^n(x') \mathcal{P}(dx' | x, u) \\ &= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\bar{\mathcal{L}}}(x) f_{S^{n+1}}(x) \int_{x' \in S^{n+1}} V_T^{n+1}(x') \mathcal{P}(dx' | x, u) = V_{T+1}^{n+1}(x) \end{aligned}$$

$$\text{Assume for } k = m : \quad V_{T-m}^n = V_{T-m}^{n+1}$$

$$\begin{aligned} \text{For } k = m + 1 : \quad V_{T-m-1}^n &= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + \\ &\quad f_{\bar{\mathcal{L}}}(x) f_{S^n}(x) \int_{x' \in \mathbb{X}} f_{S^n}(x') V_{T-m}^n(x') \mathcal{P}(dx' | x, u) \end{aligned}$$

$$\begin{aligned}
&= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) f_{S^n}(x) \int_{x' \in S^n} V_{T-m}^n(x') \mathcal{P}(dx'|x, u) \\
&= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) f_{S^{n+1}}(x) \int_{x' \in S^{n+1}} V_{T-m}^{n+1}(x') \mathcal{P}(dx'|x, u) = V_{T-m-1}^{n+1}(x)
\end{aligned}$$

Lemma 36 implies that if the set has converged, i.e., the set does not change over an iteration, then the corresponding value functions must also be equal to each other.

Theorem 19 *The set obtained from algorithm 8.2 is indeed a Safe Probable Backward Reachable Set. This means that S^n is a V_{SPBRS} if and only if S^n converges, i.e., $S^n = V_{SPBRS}$ if and only if $S^n = S^{n+1}$.*

Proof: If S^n is V_{SPBRS} , then for $x \in S^n$ the $\mathcal{J} \leq \sup_{u \in \mathcal{M}} \sup_{k \in [0, T]} \mathcal{P}(x^\tau \in \mathcal{L}, x^\tau \in S^n, x^{\tau-1} \in S^n, \dots, x^1 \in S^n | x^0 \in S^n, u) = V_0^n(x)$ (using Theorem 18). This implies $x \in S^{n+1}$ (using equation 8.10). Furthermore, if $S^n = S^{n+1}$, for $x \in S^{n+1}$ (using equation 8.10) we have $\mathcal{J} \leq V_0^n(x) = V_0^{n+1}(x)$ (using Lemma 36). This implies (using Theorem 18) that $\mathcal{J} \leq \sup_{u \in \mathcal{M}} \sup_{k \in [0, T]} \mathcal{P}(x^\tau \in \mathcal{L}, x^\tau \in S^{n+1}, x^{\tau-1} \in S^{n+1}, \dots, x^1 \in S^{n+1} | x^0 \in S^{n+1}, u)$ and S^n is a V_{SPBRS} . Now we will prove that the V_{SPBRS} obtained using algorithm 8.2 is the largest possible V_{SPBRS} if the $S^0 = \mathbb{X}$. To prove that, we first need to prove two other lemmas.

Lemma 37 *Given that $S^n \subseteq S^{n-1}$, the value function follows $V_{T-k}^n(x) \leq V_{T-k}^{n-1}(x)$.*

Proof: The proof is obtained using induction. We know that $V_T^n(x) = V_T^{n-1}(x) = f_{\mathcal{L}}(x)$ and using equation 8.8

$$\begin{aligned}
V_{T-1}^n(x) &= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) f_{S^n}(x) \int_{x' \in \mathbb{X}} f_{S^n}(x') V_T^n(x') \mathcal{P}(dx'|x, u) \\
&\leq \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) f_{S^{n-1}}(x) \int_{x' \in \mathbb{X}} f_{S^{n-1}}(x') V_T^n(x') \mathcal{P}(dx'|x, u) \\
&= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) f_{S^{n-1}}(x) \int_{x' \in \mathbb{X}} f_{S^{n-1}}(x') V_T^{n-1}(x') \mathcal{P}(dx'|x, u) = V_{T-1}^{n-1}(x)
\end{aligned}$$

Assume for $k = m$: $V_{T-m}^n(x) \leq V_{T-m}^{n-1}(x)$

For $k = m + 1$, we have:

$$\begin{aligned}
V_{T-m-1}^n(x) &= \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) f_{S^n}(x) \int_{x' \in \mathbb{X}} f_{S^n}(x') V_{T-m}^n(x') \mathcal{P}(dx'|x, u) \\
&\leq \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) f_{S^{n-1}}(x) \int_{x' \in \mathbb{X}} f_{S^{n-1}}(x') V_{T-m}^n(x') \mathcal{P}(dx'|x, u) \\
&\leq \sup_{u \in \mathcal{M}} f_{\mathcal{L}}(x) + f_{\tilde{\mathcal{L}}}(x) f_{S^{n-1}}(x) \int_{x' \in \mathbb{X}} f_{S^{n-1}}(x') V_{T-m}^{n-1}(x') \mathcal{P}(dx'|x, u) = V_{T-m-1}^{n-1}(x)
\end{aligned}$$

Corollary 6 Given that $S^n \subseteq S^{n-1}$, the value function follows $V_0^n(x) \leq V_0^{n-1}(x)$.

Proof: This corollary directly follows from Lemma 37.

Lemma 38 The set $S^n \subseteq S^{n-1}$ for any $n \in \mathbb{N}$.

Proof: The proof uses induction. We will use the fact that $V_{T-m}^n(x) \in [0, 1]$. The proof of this statement is trivial and can be directly obtained from equation 8.8.

For $n = 1$: $S^0 = \mathbb{X}$ and $S^1 = \{x : V_0^1(x) \geq \mathcal{J}\}$

We know that $V_0^1(x) \in [0, 1]$ and $\mathcal{J} \in [0, 1]$

$$\Rightarrow S^1 = \{x : V_0^0(x) \in [\mathcal{J}, 1]\} \subseteq \{x : V_0^0(x) \in [0, 1]\} = S^0$$

Assume for $n = k$ that $S^k \subseteq S^{k-1} \Rightarrow V_0^k(x) \leq V_0^{k-1}(x)$ (using Corollary 6)

$$\text{For } n = k + 1: S^{k+1} = \{x : V_0^k(x) \geq \mathcal{J}\} \subseteq \{x : V_0^{k-1}(x) \geq \mathcal{J}\} = S^k$$

Theorem 20 The V_{SPBRS} obtained via algorithm 8.2 is the largest possible V_{SPBRS} .

Proof: The theorem follows directly from the fact that the set obtained after each iteration of algorithm 8.2 is a subset of, or equal to, the previous set as shown in Lemma 38. By using theorem 19 it is shown that the produced set is a V_{SPBRS} once the set converges.

8.5. DISCRETE SPACE IMPLEMENTATION

In this section, we provide the discrete space implementation of both the algorithms. If the state space and control space are discrete, the stochastic kernel $\mathcal{P}$ is a transition matrix, denoted as $\mathcal{T}(x'|x, u)$, that provides the probability of the transition of state $x \in \mathbb{X}$ to state $x' \in \mathbb{X}$ under the control action $u \in \mathcal{M}$. The probability transition matrix follows the total law of probability, i.e., $\sum_{x' \in \mathbb{X}} \mathcal{T}(x'|x, u) = 1, \forall x \in \mathbb{X}, u \in \mathcal{M}$.

8.5.1. Discrete space implementation of Algorithm 8.1

With state space and control space being discrete, the value function $V_{T-k}(x)$ (equation 8.4) can be computed exactly, using linear programming [38]. Furthermore, the existence of an optimal Markov policy is always guaranteed.

Theorem 21 *Given a reach set $\mathcal{L} \subset \mathbb{X}$, the value function $V_0(x)$ as presented in equation 8.4 is obtained by solving the linear programming problem*

$$\max \sum_{\tau \in [0, T]} \sum_{x \in \mathcal{X}} V_{\tau}(x) \quad (8.14)$$

Such that

$$\begin{aligned} V_{\tau}(x) &\geq V_T(x) + \\ &f_{\mathcal{L}} \sum_{x' \in \mathbb{X}} \Pi_{\mathcal{L}}(V_{\tau+1}(x')) \mathcal{T}(x'|x, u) \\ &\quad \forall \tau \in \mathbb{N}_{[0, T-1]}, \forall x \in \mathbb{X}, \forall u \in \mathbb{U}_x \\ V_T(x) &\leq f_{\mathcal{L}}(x), \forall x \in \mathbb{X} \\ V_{\tau}(x) &\in [0, 1] \end{aligned}$$

Proof: The proof is exactly the same as the proof of Theorem 1 in [38].

The optimal solution obtained from the linear program 8.14 provides the value given in equation 8.4 at the final time step. That is, $V_0(x) = \max \sum_{\tau \in [0, T]} \sum_{x \in \mathcal{X}} V_{\tau}(x)$, for all $x \in \mathbb{X}$ and $\tau \in \mathbb{N}_{[0, T]}$. The linear program further provides the optimal Markov control policy u^* such that equation 8.15 is satisfied [38], where $V_{\tau}^*(x)$ is the optimal value.

$$V_{\tau}^*(x) = V_T(x) + \sum_{x' \in \mathbb{X}} \Pi_{\mathcal{L}}(V_{\tau+1}^*(x')) \mathcal{T}(x'|x, u^*) \quad (8.15)$$

Note: If the cardinality of the state space is $|N|$, the cardinality of the control space is $|U|$, and the time horizon is T steps, then the maximum number of decision variables and constraints are $|N|(T+1)$ and $|N|(T|U|+1)$ respectively. Thus, solving the linear program implementation of the algorithm in the case of discrete state and control space has a maximum implementation time of $O(|N|^2(T|U|+1))$ [39].

8.5.2. Discrete space implementation of Algorithm 8.2

With state space and control space being discrete, the value function $V_{T-k}^n(x)$ (equation 8.8) can be computed exactly through linear programming [38]. Furthermore, the existence of the optimal Markov policy is always guaranteed. For a single iteration the value function can be computed using Lemma 39. The V_{SPBRS} can be obtained following Algorithm 8.2.

Lemma 39 Given a reach set $\mathcal{L} \subset \mathbb{X}$, the value function $V_0^n(x)$ as presented in equation 8.8 is obtained by solving the linear programming problem

$$\max \sum_{\tau \in [0, T]} \sum_{x \in \mathcal{X}} V_\tau^n(x) \quad (8.16)$$

such that

$$\begin{aligned} V_\tau^n(x) &\geq V_T^n(x) + f_{\bar{\mathcal{L}}}(x) f_{S^n}(x) \\ &\quad \sum_{x' \in \mathbb{X}} f_{S^n}(x') V_{\tau+1}^n(x') \mathcal{T}(x'|x, u) \\ &\quad \forall \tau \in \mathbb{N}_{[0, T-1]}, \forall x \in \mathbb{X}, \forall u \in \mathbb{U}_x \\ V_T^n(x) &\leq f_{\mathcal{L}}(x), \forall x \in \mathbb{X} \\ V_\tau^n(x) &\in [0, 1] \end{aligned}$$

Proof: The proof is exactly the same as the proof of theorem 1 in [38].

The optimal solution obtained from the linear program 8.16 provides the value function given in equation 8.8 at the final time step for each iteration, i.e., $V_0^n = \max \sum_{\tau \in [0, T]} \sum_{x \in \mathcal{X}} V_\tau^n(x)$, $\forall x \in \mathbb{X}$ and $\tau \in \mathbb{N}_{[0, T]}$. The linear program further provides the optimal Markov control policy u^* such that equation 8.17 is satisfied [38], where $V_\tau^{n*}(x)$ is the optimal value.

$$V_\tau^{n*}(x) = V_T^n(x) + f_{\bar{\mathcal{L}}}(x) f_{S^n}(x) \sum_{x' \in \mathbb{X}} f_{S^n}(x') V_{\tau+1}^{n*}(x') \mathcal{T}(x'|x, u^*) \quad (8.17)$$

Note: If the cardinality of the state space is $|N|$, the cardinality of the control space is $|U|$, and the time horizon is T steps, then the maximum number of decision variables and constraints are $|N|(T+1)$ and $|N|(T|U|+1)$ respectively. Thus, solving the linear program implementation of algorithm 8.2 in the case of discrete state and control space has a maximum computational cost of $O(|N|^2(T|U|+1))$ per iteration [39]. It can be seen that for discrete systems, the computational cost of algorithm 8.1 is lower than the computational cost of algorithm 8.2.

8.6. CONTINUOUS SPACE IMPLEMENTATION

Many systems are continuous in nature. Thus, in this section, we provide a discretisation routine for continuous state and control spaces. We also provide a guarantee that the value function estimated via the discretised space converges to the true value function as the grid size tends to zero. For any given set $\mathcal{L} \in \mathcal{B}(\mathbb{X})$, we denote the Lebesgue measure of the set as $\psi(\mathcal{L})$. The stochastic kernel $\mathcal{P}(\cdot|x, u)$ projects a probability density kernel denoted as $p(\cdot|x, u)$. The probability density kernel $p(z|x, u)$ represents the probability density of z given the current state (x) and the control policy (u) . We assume that there is a region of interest denoted by $\mathcal{G}$ such that $\mathcal{L} \subset \mathcal{G} \subset \mathbb{X} \in \mathcal{B}(\mathbb{X})$. Set $\mathcal{G}$ is compact and $\mathcal{L}$

is a closed subset of $\mathcal{G}$ and is thus also compact. Therefore, $\psi(\mathcal{L})$ and $\psi(\mathcal{G})$ are bounded. Furthermore, we assume that the probability density function and the stochastic kernel $\mathcal{P}(\cdot|x, u)$ satisfy the following assumptions:

Assumption 15 *The probability density function is Lipschitz continuous, i.e., for any $x, x' \in \mathcal{G} \setminus \mathcal{L}, z, z' \in \mathcal{G} \setminus \mathcal{L}$ and $u, u' \in \mathbb{U}$, there exists a constant η such that equation 8.18 is valid.*

$$|p(z|x, u) - p(z'|x', u')| \leq \eta(\|z - z'\| + \|x - x'\| + \|u - u'\|) \quad (8.18)$$

Assumption 16 *The stochastic kernel is Lipschitz continuous, i.e., for any $x, x' \in \mathcal{G} \setminus \mathcal{L}, z, z' \in \mathcal{G} \setminus \mathcal{L}$ and $u, u' \in \mathbb{U}$, there exists a constant η_2 such that equation 8.19 is valid.*

$$|\mathcal{P}(z|x, u) - \mathcal{P}(z'|x', u')| \leq \eta_2(\|z - z'\| + \|x - x'\| + \|u - u'\|) \quad (8.19)$$

where η, η_2 are Lipschitz constants.

Assumption 17 *Similar to [40], we assume that $p(z|x, u)$ is a probability density on the region of interest $\mathcal{G}$. Thus, $p(z|x, u)$ satisfies equation 8.20, where κ_p is a positive constant.*

$$\int_{\mathcal{G}} p(z|x, u) dz = 1 \text{ while } p(z|x, u) \in [0, \kappa_p] \quad (8.20)$$

Lemma 40 *The set $(\mathcal{G} \setminus \mathcal{L}) \in \mathcal{B}(\mathbb{X})$ if $\mathcal{G}, \mathcal{L} \in \mathcal{B}(\mathbb{X})$*

Proof: The proof directly follows from the fact that $\mathcal{G} \setminus \mathcal{L} = \mathcal{G} \cap \mathcal{L}^c = (\mathcal{G}^c \cup \mathcal{L})^c$ and σ -algebra are closed under complements and countable union, thus $\mathcal{G} \setminus \mathcal{L} \in \mathcal{B}(\mathbb{X})$, where A^c is complement of set A

We discretise $\mathcal{G}$ into $k_x + 1$ pairwise disjoint non-empty Borel sets $\mathcal{G}_i$ and $\mathcal{L}$, where $i \in [1, k_x] \subseteq \mathbb{N}$, such that $\mathcal{G} = \bigcup_{i=1}^{k_x} \mathcal{G}_i \cup \mathcal{L}$ and $\mathcal{G}_a \cap \mathcal{G}_b = \emptyset$, $\mathcal{G}_c \cap \mathcal{L} = \emptyset$ for any $a, b, c \in [1, k_x] \subseteq \mathbb{N}, a \neq b$. Each subset is then represented by a representative state denoted by $g_i \in \mathcal{G}_i$ and $\mathcal{L}$ is represented by a representative state $l \in \mathcal{L}$. Furthermore, we define the set of all the representative states as $\hat{\mathcal{G}} = \{l, g_i : i \in [1, k_x] \subseteq \mathbb{N}\}$, the maximum set member norm as $d_{x_i} = \sup_{x, x' \in \mathcal{G}_i} \|x - x'\|$ and $D_{x, \max} = \sup_{i \in [1, k_x]} (d_{x_i})$.

In the same fashion we also discretise $\mathbb{U}$, into k_u pairwise disjoint non-empty Borel sets denoted by $\mathcal{F}_i, i \in [1, k_u] \subseteq \mathbb{N}$, thus satisfying $\mathbb{U} = \bigcup_{i \in [1, k_u]} \mathcal{F}_i$. Similar to the state space discretisation, each disjoint Borel set is represented by an element in the set, $\hat{u}_i \in \mathcal{F}_i$. Let the set of representative control policies be denoted by $\hat{\mathbb{U}} = \{\hat{u}_i : i \in [1, k_u] \subseteq \mathbb{N}\}$, the maximum set member norm as $d_{u_i} = \sup_{u, u' \in \mathcal{F}_i} \|u - u'\|$ and $D_{u, \max} = \sup_{i \in [1, k_u]} (d_{u_i})$.

We assume the grid size to be constant

$\lambda \geq \max(D_{u, \max}, D_{x, \max})$. For $x \in \mathcal{G}$ we define the admissible control 8.21.

$$\hat{\mathcal{U}}_x = \{\hat{u} \in \hat{\mathbb{U}} : \|u - \hat{u}\| \leq \lambda \text{ for some } u \in \mathbb{U}_{g_i}, x \in \mathcal{G}_i\} \quad (8.21)$$

where g_i is the representative state of $\mathcal{G}_i$ such that $x \in \mathcal{G}_i$. Similar to [40], in Lemma 41 we show that for each $x \in \mathcal{G}$ there exists a non-empty discrete control set.

Lemma 41 *For all $g_i \in \hat{\mathcal{G}}$ the set of available inputs is non-empty, i.e., $\hat{\mathcal{U}}_{g_i} \neq \emptyset$ and the available input at all states $x \in \mathcal{G}_i$ is equal to $\hat{\mathcal{U}}_{g_i}$, i.e. $\hat{\mathcal{U}}_x = \hat{\mathcal{U}}_{g_i}, \forall x \in \mathcal{G}_i$*

Proof: Since $\mathbb{U}_x \neq \emptyset, \forall x \in \mathbb{X}$, $\mathbb{U}_{g_i}$ is also non-empty, $\exists \hat{u} \in \hat{\mathbb{U}}$ such that $\|u - \hat{u}\| \leq \lambda$ for all $u \in \mathbb{U}_{g_i}$, thus $\hat{\mathcal{U}}_{g_i}$ is non-empty. Moreover, from the definition of $\hat{\mathcal{U}}_x$ it is trivial to obtain $\hat{\mathcal{U}}_{g_i} = \hat{\mathcal{U}}_x$.

Following the steps of [40], we define a function $\hat{p} : (\mathcal{G} \setminus \mathcal{L}) \times \mathcal{G} \times \hat{\mathcal{U}} \rightarrow \mathbb{R}^+$ as shown in equation 8.22 that provides a unified probability density function to all $x \in \mathcal{G}_i$, such that all $x \in \mathcal{G}_i$ have the same probability kernel. We also assume that $\int_{\mathcal{L}} p(z|x, u) dz = 1, \forall x \in \mathcal{L}$. Note that this assumption does not violate assumption 4 as $\mathcal{L} \subset \mathcal{G}$.

$$\hat{p}(x'|x, \hat{u}) = \begin{cases} \frac{p(g_{x'}|g_x, \hat{u})}{\int_{\mathcal{G}} p(g_{z'}|g_x, \hat{u}) dz} & \text{if } \int_{\mathcal{G}} p(g_{z'}|g_x, \hat{u}) dz \geq 1 \\ p(g_{x'}|g_x, \hat{u}) & \text{otherwise} \end{cases} \quad (8.22)$$

Here $g_{x'}$ and g_x are the respective representative states of the grid cell to which x' and x belong. This gives rise to an approximate stochastic control system $\hat{\mathcal{S}}_{\mathcal{G}} = (\hat{\mathcal{G}}, \hat{\mathcal{U}}, \hat{\mathbb{P}})$, where $\hat{\mathbb{P}}$ is the transition probability defined as $\hat{\mathbb{P}}(g_j|g_i, \hat{u}) = \int_{\mathcal{G}_j} \hat{p}(x'|g_i, \hat{u}) dx'$, $(g_j, g_i) \in \hat{\mathcal{G}}$, $g_j \in \mathcal{G}_j, g_i \in \mathcal{G}_i, \hat{\mathbb{P}}(l|g_i, \hat{u}) = \int_{\mathcal{L}} \hat{p}(x'|g_i, \hat{u}) dx'$, $\hat{\mathbb{P}}(l|l, \hat{u}) = \int_{\mathcal{L}} p(x'|l, \hat{u}) dx' = 1$ and $\hat{u} \in \hat{\mathbb{U}}$.

8

8.6.1. Continuous space implementation of Algorithm 8.1

We define a function similar to equation 8.2 as $\hat{f}_l : \hat{\mathcal{G}} \rightarrow \{0, 1\}$, shown in the following equation

$$\hat{f}_l(g_i) = \begin{cases} 1, & \text{if } g_i = l \\ 0, & \text{if } g_i \neq l \end{cases} \quad (8.23)$$

For the discrete system $\hat{\mathcal{S}}_{\mathcal{G}}$, the value function given by equation 8.4 boils down to equation 8.24, where $\bar{l}$ is the complement of l in $\hat{\mathcal{G}}$.

$$\hat{V}_{T-k}(g_i) = \hat{f}_l(g_i) + \hat{f}_{\bar{l}}(g_i) \max_{\hat{u} \in \hat{\mathbb{U}}} \left(\sum_{j=1}^{k_x} \Pi_{\mathcal{J}}(\hat{V}_{T-k+1}(g_j)) \hat{\mathbb{P}}(g_j|g_i, \hat{u}) \right), \quad \forall k \in [0, T] \subseteq \mathbb{N} \quad (8.24)$$

For $x \in \mathcal{G}_i$, $\hat{V}_{T-k}(x) = \hat{V}_{T-k}(g_i)$ for all $\tau \in [0, T] \in \mathbb{N}$. The discrete optimal control policy is defined by equation 8.25.

$$\hat{u}_{\tau, \text{SPBRS}}^*(g_i) = \arg \min_{\hat{u} \in \hat{\mathcal{U}}} \left(\sum_{j=1}^{k_x} \Pi_{\mathcal{J}}(\hat{V}_{T-k+1, l}(g_j)) \hat{P}(g_j | g_i, \hat{u}) \right) \quad (8.25)$$

Note: The discrete system $\hat{\mathcal{S}}$ has a finite state and control policy space. Thus, the value $\hat{V}_{T-k}(g_i)$ can be calculated using the linear program shown in Theorem 21. The corresponding optimal control policy can be computed using equation 8.14.

In the following, we guarantee the convergence of the value function as the grid size approaches zero. To prove the convergence of the approximate system's value function to its continuous counterpart we need two more Lemmas.

Lemma 42 *If assumptions 13 and 15 are satisfied and $0 < \mathcal{J}$, then, for any states $x, x' \in S_k \setminus \mathcal{L}$, the value function is Lipschitz continuous, i.e., $V_{T-k}(x)$ satisfies equation 8.26.*

$$|V_{T-k}(x) - V_{T-k}(x')| \leq \kappa_1 \psi(\mathcal{G}) \|x - x'\|, \forall k \in [0, T] \subseteq \mathbb{N} \quad (8.26)$$

Proof: At $k = 0$ the value function is equal to $f_{\mathcal{L}}(x)$. Thus, $|V_{T, \mathcal{L}}(x) - V_{T, \mathcal{L}}(x')| \leq \kappa_T \|x - x'\|$, $\forall x, x' \in S_k \setminus \mathcal{L}$ holds trivially. For $k \in [0, T]$ using equation 8.4 we get:

$$\begin{aligned} |V_{T-k}(x) - V_{T-k}(x')| &= |f_{\mathcal{L}}(x) + \inf_u f_{\bar{\mathcal{L}}}(x) \int_{\mathbb{X}} \Pi_{\mathcal{J}}(V_{T-k+1}(\bar{x})) \mathcal{P}(d\bar{x} | x, u) \\ &\quad - f_{\mathcal{L}}(x') - \inf_u f_{\bar{\mathcal{L}}}(x') \int_{\mathbb{X}} \Pi_{\mathcal{J}}(V_{T-k+1}(\bar{x})) \mathcal{P}(d\bar{x} | x', u)| \end{aligned}$$

8

Using equation 8.5 and since $0 < \mathcal{J}$

$$\begin{aligned} &= \left| \inf_u \int_{S_{k+1}} V_{T-k+1}(\bar{x}) \mathcal{P}(d\bar{x} | x, u) - \inf_u \int_{S_{k+1}} V_{T-k+1}(\bar{x}) \mathcal{P}(d\bar{x} | x', u) \right|, \\ &\quad \forall x, x' \in S_{T-k} \setminus \mathcal{L} \\ &\leq \inf_u \left| \int_{S_{k+1}} V_{T-k+1}(\bar{x}) (\mathcal{P}(d\bar{x} | x, u) - \mathcal{P}(d\bar{x} | x', u)) \right|, \forall x, x' \in S_{T-k} \setminus \mathcal{L} \end{aligned}$$

Since $V_{T-k+1}(\bar{x}) \in [0, 1]$

$$\leq \inf_u \int_{S_{k+1}} |(\mathcal{P}(d\tilde{x}|x, u) - \mathcal{P}(d\tilde{x}|x', u))|, \forall x, x' \in S_{T-k} \setminus \mathcal{L}$$

$$\leq \inf_u \psi(\mathcal{G}) \kappa_1 \|x - x'\| = \kappa_1 \psi(\mathcal{G}) \|x - x'\|$$

Lemma 43 *If assumptions 14 and 16 are satisfied, then $\forall x' \in \mathcal{G} \setminus \mathcal{L}$ and $g_i \in \hat{\mathcal{G}} \setminus \{l\}$ the integration of the modulus of the difference between the approximate probability density and the actual probability density over the space $\mathcal{G}$ is bounded for all $\hat{u} \in \hat{\mathcal{U}}$, i.e., equation 8.27 holds. Where κ_2 is a constant*

$$\int_{\mathcal{G}} |\hat{p}(x'|g_i, \hat{u}) - p(x'|g_i, \hat{u})| dx' \leq 2\kappa_2 \psi(\mathcal{G})\lambda \quad (8.27)$$

Proof: The proof of the following lemma is the same as the proof of Lemma 9 in [41],

In the following Theorem we prove the convergence of the discretised value function to the actual value function as the grid size tends to zero.

Theorem 22 *If assumptions 14 -16 are satisfied, then for all $x \in \mathcal{G}$, the difference between the continuous value function ($V_{T-k}(x)$) and the discretised value function $\hat{V}_{T-k}(x)$ is bounded through equation 8.28, where $h(k)$ is a constant depending on k .*

$$|\hat{V}_{T-k}(x) - V_{T-k}(x)| \leq h(k)\psi(\mathcal{G})\lambda \quad (8.28)$$

Proof: For $x \in \mathcal{L}$ we have $\hat{V}_{T-k}(x) = V_{T-k}(x) = 1$ for all $k \in [0, T]$. Thus, equation 8.28 holds trivially. Similarly for $x \in \mathcal{G} \setminus S_k$ we have $\hat{V}_{T-k}(x) = V_{T-k}(x) = 0$ for all $k \in [0, T]$. Thus, equation 8.28 holds trivially. For all $x \in S_k \setminus \mathcal{L}$, and $k = 0$, $\hat{V}_{T-k}(x) = V_{T-k}(x) = 0$. Proceeding using induction we assume that for some $k + 1$, $|\hat{V}_{T-k+1}(x) - V_{T-k+1}(x)| \leq h(k+1)\psi(\mathcal{G})\lambda$. For any $g_i \in \mathcal{G}_i$, $i \in [0, k_x] \subseteq \mathbb{N}$, we denote u_{T-k}^* as the input policy that minimises the value function. Furthermore, we denote $\hat{u}_{T-k}^*$ as the input policy that minimises the discretised value function. Thus, from equation 8.5 and since $0 < \mathcal{J}$, we have

$$\begin{aligned} V_{T-k}(g_i) - \hat{V}_{T-k}(g_i) &= \int_{S_{k+1}} V_{T-k+1}(x') \mathcal{P}(dx'|g_i, u_{T-k}^*) \\ &\quad - \int_{S_{k+1}} \hat{V}_{T-k+1}(x') \hat{\mathcal{P}}(dx'|g_i, \hat{u}_{T-k}^*) \\ &\leq \int_{S_{k+1}} |V_{T-k+1}(x') \mathcal{P}(dx'|g_i, u_{T-k}^*) - \hat{V}_{T-k+1}(x') \hat{\mathcal{P}}(dx'|g_i, \hat{u}_{T-k}^*)| \end{aligned}$$

$$\begin{aligned} &\leq \int_{S_{k+1}} |V_{T-k+1}(x') \mathcal{P}(dx' | g_i, u_{T-k}^*) - \hat{V}_{T-k+1}(x') \mathcal{P}(dx' | g_i, u_{T-k}^*)| \\ &\quad + \int_{S_{k+1}} |\hat{V}_{T-k+1}(x') \mathcal{P}(dx' | g_i, u_{T-k}^*) - \hat{V}_{T-k+1}(x') \hat{\mathcal{P}}(dx' | g_i, \hat{u}_{T-k}^*)| \end{aligned}$$

Since $\hat{V}_{T-k+1}(x') \leq 1, \forall x \in \mathcal{G}$

$$\begin{aligned} &\leq \int_{S_{k+1}} |\hat{\mathcal{P}}(dx' | g_i, \hat{u}_{T-k}^*) - \mathcal{P}(dx' | g_i, u_{T-k}^*)| + \\ &\quad \int_{S_{k+1}} |\hat{V}_{T-k+1}(x) - V_{T-k+1}(x)| \mathcal{P}(dx' | g_i, u_{T-k}^*)| \end{aligned}$$

Using Lemma 43 and the induction assumption, we get:

$$\leq 2\kappa_2 \psi(\mathcal{G})\lambda + h(k+1)\psi(\mathcal{G})\lambda$$

And similarly for

$$\begin{aligned} \hat{V}_{T-k}(g_i) - V_{T-k}(g_i) &= \int_{S_{k+1}} \hat{V}_{T-k+1}(x') \hat{\mathcal{P}}(dx' | g_i, \hat{u}_{T-k}^*) \\ &\quad - \int_{S_{k+1}} V_{T-k+1}(x') \mathcal{P}(dx' | g_i, u_{T-k}^*) \\ &\leq \int_{S_{k+1}} |\hat{V}_{T-k+1}(x') \hat{\mathcal{P}}(dx' | g_i, \hat{u}_{T-k}^*) - \hat{V}_{T-k+1}(x') \mathcal{P}(dx' | g_i, \hat{u}_{T-k}^*)| \\ &\quad + |\hat{V}_{T-k+1}(x') \mathcal{P}(dx' | g_i, \hat{u}_{T-k}^*) - V_{T-k+1}(x') \mathcal{P}(dx' | g_i, \hat{u}_{T-k}^*)| \\ &\leq 2\kappa_2 \psi(\mathcal{G})\lambda + h(k+1)\psi(\mathcal{G})\lambda \end{aligned}$$

This leads to

$$|V_{T-k}(g_i) - \hat{V}_{T-k}(g_i)| \leq 2\kappa_2 \psi(\mathcal{G})\lambda + h(k+1)\psi(\mathcal{G})\lambda$$

Since for any $x \in \mathcal{G}_i, i \in [0, k_x] \subseteq \mathbb{N}, \hat{V}_{T-k}(x) = \hat{V}_{T-k}(g_i)$

$$|V_{T-k}(x) - \hat{V}_{T-k}(x)| = |V_{T-k}(x) - \hat{V}_{T-k}(g_i)|$$

$$\leq |V_{T-k}(x) - V_{T-k}(g_i)| + |V_{T-k}(g_i) - \hat{V}_{T-k}(g_i)|$$

$$\leq h(k+1)\psi(\mathcal{G})\lambda$$

Thus, this theorem assures the convergence of the approximate value function to the true value function as the grid size tends to zero.

Note: Following similar arguments as for the computation cost for algorithm 8.1 and [39], we can say that for continuous space systems, algorithm 8.1 has a computation time of $O(|k_x|^2(T|k_u| + 1))$.

8.6.2. Continuous space implementation of Algorithm 8.2

In this subsection we prove that the safe probable set obtained via Algorithm 8.2 for the discretised space converges to the safe probable set of the continuous space as the grid size tends to zero. To prove this, we need two Lemmas.

Lemma 44 *If assumptions 1 and 3 are satisfied, then, for any states $x, x' \in \mathcal{G} \setminus \mathcal{L}$ and for any iteration n , the value function is Lipschitz continuous, i.e., $V_{T-k}(x)$ satisfies equation 8.29, where κ_3 is the Lipschitz constant.*

$$|V_{T-k}^n(x) - V_{T-k}^n(x')| \leq \kappa_3 \psi(\mathcal{G}) \|x - x'\|, \forall k \in [0, T] \subseteq \mathbb{N}, n \in \mathbb{N} \quad (8.29)$$

Proof: The proof of this Lemma is exactly same as the proof of Lemma 42.

In the following Lemma we prove the convergence of the discretised value function to the actual value function as the grid size tends to zero.

Theorem 23 *If assumptions 2-4 are satisfied, then for all $x \in \mathcal{G}$ the difference between the value function and the discretised value function is bounded for any iteration as shown in equation 8.30.*

$$|\hat{V}_{T-k}^n(x) - V_{T-k}^n(x)| \leq h(k)\psi(\mathcal{G})\lambda, \forall n \in \mathbb{N} \quad (8.30)$$

Proof: The proof of this Theorem is exactly the same as the proof of Theorem 22, but Lemma 42 is replaced by Lemma 8.29

Note: Following similar arguments as for the computation cost for algorithm 8.2 and [39], we can say that for continuous space systems, algorithm 8.2 has a computation time of $O(|k_x|^2(T|k_u| + 1))$ per iteration. Furthermore, similar to the case of discrete systems, here too the computational cost of algorithm 8.1 is lower than the computational cost of algorithm 8.2.

8.7. EXAMPLE

To demonstrate the above algorithms, we make use of a the discrete Dubin's car system. The model given in equation 8.31 describes the motion of a system in a

two-dimensional plane. The equations for the system's position and orientation at the next time step $k + 1$ are given by:

$$\begin{aligned} x_{k+1} &= x_k + v \cdot \cos(\theta_k) \cdot \Delta t, \\ y_{k+1} &= y_k + v \cdot \sin(\theta_k) \cdot \Delta t, \\ \theta_{k+1} &= \theta_k + \omega_k \cdot \Delta t + d_k \cdot \Delta t. \end{aligned} \quad (8.31)$$

Here, x_k and y_k represent the Cartesian coordinates of the system at time step k , and θ_k is its orientation in radians. The velocity v determines the speed of motion, while ω_k is the yaw rate, representing the rate of change of the orientation. d_k accounts for any disturbances affecting the orientation. The parameter Δt is the time step length.




Figure 8.4: The slices of $V_{\text{SPBRs}}(\cdot)$ for the system given in equation 8.31.

The disturbance was discretised in a grid of 4×1 and the probability mass

function of the disturbance for the range $[-0.8, 0.8]$ is given as follows:

$$\begin{aligned}\mu &= 0, \\ \sigma &= \frac{1.6}{6} \approx 0.2667, \\ \text{Probability mass}(d_i) &= \exp\left(-\frac{(d_i)^2}{2 \cdot (0.2667)^2}\right).\end{aligned}\tag{8.32}$$

The input was discretised in a grid of 3×1 in a range of $[-1, 1]$. The state space was discretised with a grid space of 0.2 in all the dimension of the system in a range of $[-5, 5]$ for x, y and $[0, 2\pi]$ for θ . Using this system, we obtained the largest possible $V_{\text{SPBRS}}(\cdot)$ with a time period of 4 seconds. $\mathcal{L}$ was taken to the cylinder with a radius of 1 about the origin, $\mathcal{J}$ is equal to 0.6, and the velocity of the system was taken to 1. The $V_{\text{SPBRS}}(\cdot)$ for the system given in equation 8.31 is shown in figure 8.4.

8.8. CONCLUSIONS

In this paper we provided a new extension of the stochastic reachable set, which we call the Safe Probable Backward Reachable Set $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. This new set is viable in the sense that if the system starts within $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$, it will remain within $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ until it reaches the target set with a certain probability. We then developed two algorithms to compute $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ and formally demonstrated that the sets generated by both algorithms are indeed $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$. Additionally, we established that the $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ produced by Algorithm 8.2 represents the largest possible set of this kind. We implemented these algorithms for discrete state-space systems, where both methods accurately compute the exact $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$, although we were only able to show one example due to a lack of space. Furthermore, we extended these algorithms to continuous state-space systems through a discretisation approach. We proved that the sets obtained in this discretised space, though estimates, converge to the true $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ as the grid size approaches zero. Finally, we analysed the computational costs associated with both algorithms for discrete and continuous spaces. Our findings show that Algorithm 8.1 is computationally more efficient than Algorithm 8.2.

In future works we would like to integrate the $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ with Control Barrier Functions (CBF) and Control Lyapunov Functions (CLF)-based controllers in order to overcome the problem of finding the CBF and CLF. We would also like to integrate the $V_{\text{SPBRS}}(T, \mathcal{L}, \mathcal{J})$ with Model Predictive Control (MPC) to produce a controller similar to the stochastic tube MPC. Furthermore, we also aim to explore the tractability of the developed algorithms for different systems and attempt to increase the efficiency of algorithms for certain classes of systems such as linear systems. Finally, we are also interested in exploring the use of function approximation techniques

to approximate the value function of the $V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J})$, thus enabling us to estimate the $V_{\text{SPBRs}}(T, \mathcal{L}, \mathcal{J})$ with very high computational efficiency.

REFERENCES

- [1] D. Bertsekas. “Infinite time reachability of state-space regions by using feedback control”. In: *IEEE Transactions on Automatic Control* 17.5 (1972), pp. 604–613.
- [2] F. Blanchini. “Controlling systems via set-theoretic methods: some perspectives”. In: *Proceedings of the 45th IEEE Conference on Decision and Control*. IEEE. 2006, pp. 4550–4555.
- [3] F. Blanchini. “Set invariance in control”. In: *Automatica* 35.11 (1999), pp. 1747–1767.
- [4] M. Chen, S. Herbert, and C. J. Tomlin. “Exact and efficient Hamilton-Jacobi guaranteed safety analysis via system decomposition”. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2017, pp. 87–92.
- [5] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert. “Robust control barrier–value functions for safety-critical control”. In: *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE. 2021, pp. 6814–6821.
- [6] I. M. Mitchell and J. A. Templeton. “A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems”. In: *Hybrid Systems: Computation and Control: 8th International Workshop, HSCC 2005, Zurich, Switzerland, March 9-11, 2005. Proceedings 8*. Springer. 2005, pp. 480–494.
- [7] R. van den Brandt and C. de Visser. “Safe flight envelope uncertainty quantification using probabilistic reachability analysis”. In: *IFAC-PapersOnLine* 51.24 (2018), pp. 628–635.
- [8] S. Herbert, J. J. Choi, S. Sanjeev, M. Gibson, K. Sreenath, and C. J. Tomlin. “Scalable learning of safety guarantees for autonomous systems using Hamilton-Jacobi reachability”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 5914–5920.
- [9] S. L. Herbert. *Safe real-world autonomy in uncertain and unstructured environments*. University of California, Berkeley, 2020.
- [10] S. V. Rakovic and K. I. Kouramas. “Invariant approximations of the minimal robust positively invariant set via finite time Aumann integrals”. In: *2007 46th IEEE Conference on Decision and Control*. IEEE. 2007, pp. 194–199.

- [11] M. Rungger and P. Tabuada. “Computing robust controlled invariant sets of linear systems”. In: *IEEE Transactions on Automatic Control* 62.7 (2017), pp. 3665–3670.
- [12] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. “Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems”. In: *Automatica* 44.11 (2008), pp. 2724–2734.
- [13] J. D. Gleason, A. P. Vinod, and M. M. Oishi. “Lagrangian approximations for stochastic reachability of a target tube”. In: *Automatica* 128 (2021), p. 109546.
- [14] A. P. Vinod and M. M. Oishi. “Stochastic reachability of a target tube: Theory and computation”. In: *Automatica* 125 (2021), p. 109458.
- [15] A. Abate. *Probabilistic reachability for stochastic hybrid systems: theory, computations, and applications*. University of California, Berkeley, 2007.
- [16] J. D. Gleason, A. P. Vinod, and M. M. Oishi. “Underapproximation of reach-avoid sets for discrete-time stochastic systems via lagrangian methods”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 4283–4290.
- [17] S. E. Z. Soudjani and A. Abate. “Probabilistic reach-avoid computation for partially degenerate stochastic processes”. In: *IEEE Transactions on Automatic Control* 59.2 (2013), pp. 528–534.
- [18] S. Summers and J. Lygeros. “Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem”. In: *Automatica* 46.12 (2010), pp. 1951–1961.
- [19] A. P. Vinod and M. M. Oishi. “Scalable underapproximative verification of stochastic LTI systems using convexity and compactness”. In: *Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control (Part of CPS Week)*. 2018, pp. 1–10.
- [20] S. Gao, S. Kong, W. Chen, and E. Clarke. “Delta-complete analysis for bounded reachability of hybrid systems”. In: *arXiv preprint arXiv:1404.7171* (2014).
- [21] S. Kong, S. Gao, W. Chen, and E. Clarke. “dReach: δ -reachability analysis for hybrid systems”. In: *Tools and Algorithms for the Construction and Analysis of Systems: 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015, Proceedings 21*. Springer. 2015, pp. 200–205.
- [22] K. Lesser, M. Oishi, and R. S. Erwin. “Stochastic reachability for control of spacecraft relative motion”. In: *52nd IEEE Conference on Decision and Control*. IEEE. 2013, pp. 4705–4712.

- [23] N. Malone, K. Lesser, M. Oishi, and L. Tapia. “Stochastic reachability based motion planning for multiple moving obstacle avoidance”. In: *Proceedings of the 17th international conference on Hybrid systems: computation and control*. 2014, pp. 51–60.
- [24] P. M. Esfahani, D. Chatterjee, and J. Lygeros. “The stochastic reach-avoid problem and set characterization for diffusions”. In: *Automatica* 70 (2016), pp. 43–56.
- [25] J.-P. Aubin, G. Da Prato, and H. Frankowska. “Stochastic invariance for differential inclusions”. In: *Set-Valued Analysis* 8 (2000), pp. 181–201.
- [26] H. M. Soner and N. Touzi. “Dynamic programming for stochastic target problems and geometric flows”. In: *Journal of the European Mathematical Society* 4.3 (2002), pp. 201–236.
- [27] H. M. Soner and N. Touzi. “Stochastic target problems, dynamic programming, and viscosity solutions”. In: *SIAM Journal on Control and Optimization* 41.2 (2002), pp. 404–424.
- [28] M. A. Hernández-Mejías and A. Sala. “Reliability and time-to-failure bounds for discrete-time constrained Markov jump linear systems”. In: *International Journal of Robust and Nonlinear Control* 27.10 (2017), pp. 1773–1791.
- [29] M. A. Hernández-Mejías, A. Sala, C. Arino, and A. Querol. “Reliable controllable sets for constrained Markov-jump linear systems”. In: *International Journal of Robust and Nonlinear Control* 26.10 (2016), pp. 2075–2089.
- [30] A. Mesbah. “Stochastic model predictive control: An overview and perspectives for future research”. In: *IEEE Control Systems Magazine* 36.6 (2016), pp. 30–44.
- [31] M. Cannon, B. Kouvaritakis, S. V. Raković, and Q. Cheng. “Stochastic tubes in model predictive control with probabilistic constraints”. In: *IEEE Transactions on Automatic Control* 56.1 (2010), pp. 194–200.
- [32] I. M. Mitchell, S. Kaynama, M. Chen, and M. Oishi. “Safety preserving control synthesis for sampled data systems”. In: *Nonlinear Analysis: Hybrid Systems* 10 (2013), pp. 63–82.
- [33] J. Burlet, O. Aycard, and T. Fraichard. “Robust motion planning using markov decision processes and quadtree decomposition”. In: *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA’04. 2004*. Vol. 3. IEEE. 2004, pp. 2820–2825.
- [34] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin. “FaSTrack: A modular framework for fast and guaranteed safe motion planning”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 1517–1522.
- [35] J. Ding, M. Kamgarpour, S. Summers, A. Abate, J. Lygeros, and C. Tomlin. “A stochastic games framework for verification and control of discrete time stochastic hybrid systems”. In: *Automatica* 49.9 (2013), pp. 2665–2674.

- [36] D. Bertsekas and S. E. Shreve. *Stochastic optimal control: the discrete-time case*. Vol. 5. Athena Scientific, 1996.
- [37] M. B. Stinchcombe and H. White. “Some measurability results for extrema of random functions over random sets”. In: *The Review of Economic Studies* 59.3 (1992), pp. 495–514.
- [38] A. Bhattacharya and J. P. Kharoufeh. “Linear programming formulation for non-stationary, finite-horizon Markov decision process models”. In: *Operations Research Letters* 45.6 (2017), pp. 570–574.
- [39] M. B. Cohen, Y. T. Lee, and Z. Song. “Solving linear programs in the current matrix multiplication time”. In: *Journal of the ACM (JACM)* 68.1 (2021), pp. 1–39.
- [40] C.-S. Chow and J. N. Tsitsiklis. “An optimal one-way multigrid algorithm for discrete-time stochastic control”. In: *IEEE transactions on automatic control* 36.8 (1991), pp. 898–914.
- [41] Y. Gao, K. H. Johansson, and L. Xie. “Computing probabilistic controlled invariant sets”. In: *IEEE Transactions on Automatic Control* 66.7 (2020), pp. 3138–3151.

9

CONCLUSIONS AND FUTURE DIRECTIONS

This chapter synthesises the core contributions of the dissertation, which centres on defining, computing, and certifying safe operating regions for autonomous dynamic systems under both deterministic and stochastic settings. It revisits the research objectives and outlines how each chapter addressed key questions—ranging from maneuverable safety definitions and robust reachable set formulations to neural network certification and probabilistic safety guarantees. The chapter also reflects on methodological limitations, discusses the broader impact of the work, and outlines promising future directions, establishing a foundation for scalable and formally sound safety assurance in real-world autonomous systems.

9.1. OVERVIEW

This work addressed the fundamental question of how to rigorously define, compute, and certify safe operating regions—commonly referred to as safe sets—in both deterministic and stochastic settings. The contributions include the introduction of maneuverable viable sets, which redefine what it means for a region of the state space to be not just safe but operationally viable; the refinement of classical safe set definitions through their relationship with forward and backward reachable sets; the formulation of robust avoid backward reachable sets under unmodelled dynamics; and the provision of formal certification methods for neural network-based reachable set approximations using SMT-based synthesis techniques. The work further extends into the stochastic domain, offering new constructs such as the Stochastic Avoid Backward Reachable Set and the Safe Probable Backward Reachable Set, which establish probabilistic safety guarantees under uncertainty.

At the technical core of the dissertation lies a unified and rigorous framework for certifying the safety of autonomous dynamical systems through reachability analysis. Central to this framework is the concept of maneuverable viable sets, a novel definition that ensures not only viability but also the ability to maneuver between states within a finite time horizon. The thesis further refines the classical notion of safe sets by characterizing their structure via Forward and Backward Reachable Sets, thereby enhancing both theoretical clarity and computational feasibility. To address robustness in uncertain environments, the concept of Robust Avoid Backward Reachable Sets is introduced and shown to satisfy a Hamilton-Jacobi-Isaacs (HJI) partial differential equation in the viscosity sense. Recognising the limitations of classical methods in high-dimensional spaces, the work provides formal certification of neural network-based reach and avoid sets using Satisfiability Modulo Theory (SMT) through Counterexample-Guided Inductive Synthesis (CEGIS). In the stochastic domain, the dissertation proposes both the Stochastic Avoid Backward Reachable Set and the Safe Probable Backward Reachable Set, the latter guaranteeing probabilistic viability until a target is reached. Together, these contributions advance the frontiers of formal safety analysis in autonomous systems, offering scalable, certifiable strategies applicable to a wide range of real-world scenarios.

The remainder of this chapter is organised as follows. Section 9.1.1 restates the central research objectives that shaped the trajectory of the dissertation. Section 9.2.1 synthesises the key contributions of each chapter, offering a detailed discussion of their theoretical and practical implications. Section 9.3 presents a critical assessment of the limitations associated with each contribution, both in terms of scope and applicability. Section 9.4 outlines several promising directions for future research, building upon the ideas introduced in this work. Section 9.5 discusses the broader impact of the research across various scientific and societal domains. Finally, Section 9.6 closes the dissertation with final remarks.

9.1.1. Restatement of Research Objectives

The primary objective of this research was to identify, compute, and certify safe operating regions—commonly referred to as safe sets—in the state space of dynamical systems. These regions are essential for ensuring that autonomous and safety-critical systems operate reliably, even in the presence of external disturbances, modelling uncertainties, or high-dimensional complexity. Traditional approaches to reachability and safety verification often fall short in handling these challenges, particularly in terms of robustness and scalability. This dissertation aimed to overcome these limitations through a combination of novel theoretical definitions, computational tools, and formal verification techniques.

In particular, the research sought to redefine safe sets in terms of both viability and maneuverability, integrate uncertainty into reachability formulations through robust and adversarial models, and extend safety analysis to probabilistic domains. Furthermore, it aimed to address the curse of dimensionality by leveraging neural networks for approximation—while simultaneously introducing SMT-based tools to provide formal certification of these approximations. Taken together, these objectives formed a cohesive and forward-looking agenda for the scalable and reliable certification of safety in autonomous dynamical systems.

9.2. CONCLUSIONS

9.2.1. Detailed Summary of Contributions

This dissertation presents a unified, rigorous, and scalable framework for the formal certification of safe operating regions for dynamic systems. It addresses the central research question: How can we formally define, compute, and certify safe operating regions for dynamic systems—under both deterministic and stochastic settings—while ensuring scalability, robustness to uncertainty, and formal correctness guarantees? Each chapter makes a distinct theoretical and practical contribution to this overarching goal, structured around the specific sub-questions posed in the introduction.

Research question: Chapter 3 and 4

What is an appropriate definition of a safe set that ensures both viability and maneuverability, and how can such sets be computed using reachability analysis?

The prevalence of control failures in robotics and automation highlights a pressing need to define safe operating regions that go beyond static containment. Traditional viable sets ensure that a system, once inside the set, can be kept there indefinitely through appropriate control inputs. However, such sets cannot guarantee the system’s ability to move within that set. For safety-critical autonomous systems, the ability to reconfigure and maneuver freely—without compromising safety—is indispensable.

To bridge this gap, this research introduced two related but conceptually distinct sets: the maneuverable viable set (or maneuver set) and an improved safe set. Both are defined using Hamilton-Jacobi reachability tools and constructed from specific combinations of Forward Reachable Sets (FRS) and Backward Reachable Sets (BRS). However, they differ fundamentally in what they guarantee.

The maneuver set is a region in the state space that satisfies two important properties: (i) it is viable, meaning the system can remain within the set over a given time horizon; and (ii) it allows the system to travel from any state in the set to any other state within the time horizon. These properties are verified using a constructive algorithm based on HJ reachability analysis, and theoretical results prove that such sets exist under well-defined conditions. However, as noted in the chapter, these two properties—viability and two-way maneuverability—cannot always be guaranteed simultaneously in practice. That is, the system may be able to remain within the maneuver set indefinitely or move between any two points, but not both without possibly leaving the set during the transition.

To resolve this limitation, an improved definition of the safe set was proposed. This set is more restrictive but more powerful: it contains only those states from which the system that can both remain inside the set (viability) and transition to any other point within the set without ever leaving it.

The contribution of this chapter, therefore, lies in both the conceptual innovation and the mathematical formalisation of these sets. The maneuver set provides a less restrictive, more flexible notion useful in scenarios where brief exits from a region are tolerable or can be accounted for. In contrast, the improved safe set offers a stricter—but operationally more robust—definition suitable for applications demanding continuous containment and transition capabilities.

These constructs were validated using numerical case studies, including the Dubins car and a double integrator system. These examples illustrate how the interplay between forward and backward reachability can be leveraged to define and compute safe sets that go beyond classical viability kernels. Together, the maneuver set and the improved safe set represent a foundational step forward in reachability-based safety analysis, offering richer guarantees and greater applicability in real-world autonomous systems.

Research question: Chapter 5

How can bounded uncertainties and unmodelled dynamics be systematically incorporated into reachability analysis?

As modern autonomous systems become more complex and safety-critical, classical validation techniques such as simulation fall short of offering guarantees. Simulations, while useful, are inherently incomplete for continuous-state systems, as they can only explore a limited subset of possible trajectories. To overcome this limitation, reachability analysis has emerged as a powerful framework for formal safety verification, as it captures the behaviour of all possible trajectories under worst-case disturbances. However, traditional Hamilton-Jacobi (HJ) reachability

formulations rely on fully known, deterministic system models and are primarily limited to two-player games—between the controller and a disturbance. This limits their robustness and applicability in real-world settings, where system models are often incomplete and exhibit unmodelled dynamics.

To address this challenge, this dissertation develops a significant extension of the classical HJ framework by formulating a three-player differential game, where the third player represents unmodelled dynamics. In this framework, the controller seeks to maintain safety, while two adversarial players—one representing known disturbances and the other representing unmodelled dynamics—act collaboratively to drive the system toward unsafe regions. This game-theoretic construction captures the compounded worst-case behaviour that arises when systems face both environmental disturbances and internal model uncertainties simultaneously.

The chapter defines a new construct: the Robust Avoid Backward Reachable Tube (RABRT). This tube characterises the set of initial states from which the system, under worst-case inputs from both adversaries, is guaranteed to enter an avoid set within a finite time horizon. Importantly, the value function corresponding to this game is shown to be the viscosity solution of an extended Hamilton-Jacobi-Isaacs (HJI) partial differential equation, thereby generalising the classical two-player framework in a mathematically rigorous way. This ensures that the RABRT is not just conceptually intuitive but also analytically well-defined and computationally tractable.

One of the major strengths of this formulation is its modularity and compatibility with existing tools. The robust Hamiltonian can be easily implemented using level-set methods and toolboxes such as those developed by [1], making it practical for integration into safety-critical applications. A numerical example is provided in the chapter, demonstrating the feasibility and effectiveness of computing robust reachable tubes in a system subject to uncertain dynamics.

This work provides a methodological advancement for safety analysis under bounded model uncertainty. It enables the direct incorporation of real-world unpredictability into the verification pipeline—something not previously addressed in HJ reachability analysis. By defining a principled way to account for adversarial model uncertainty, it paves the way for robust safety guarantees in applications such as aircraft collision avoidance, autonomous vehicle navigation, and robotic manipulation in unstructured environments.

Research question: Chapter 6

How can the correctness of learned reachable sets from neural approximators be formally certified?

As autonomous systems become more sophisticated and are deployed in increasingly high-dimensional and uncertain environments, the need for scalable yet reliable safety verification methods becomes paramount. To tackle the curse of dimensionality inherent to grid-based HJB reachability analysis, learning-based approaches such as DeepReach have emerged, using neural networks to approximate

the value function that characterises the reachable set. These methods bypass the need for spatial discretisation and exhibit impressive scalability. However, they trade off theoretical guarantees for empirical performance: although DeepReach solutions converge in practice, there exists no formal assurance that the neural network solution satisfies the HJ PDE or that the zero-level set of the approximation corresponds to the true reachable set.

This dissertation bridges that critical gap by introducing a formal framework—Certified Approximate Reachability (CARE), which allows for the certification of neural network approximations of reachable sets. The key idea is to treat the neural network approximation not merely as a solution surrogate but as a candidate for formal verification. The chapter first introduces a modified HJ-PDE formulation that connects the training loss directly to the approximation error on the final reachable set. It then establishes a formal residual bound, denoted ϵ , on the PDE loss, such that any deviation from the HJ-PDE remains within this bounded error across the domain of interest.

To go beyond empirical estimation, the framework employs Satisfiability Modulo Theory (SMT) solvers in combination with Counterexample-Guided Inductive Synthesis (CEGIS) to certify that the network approximation satisfies the modified PDE constraints within the ϵ -bound. This allows the construction of provably sound over- and under-approximations of the true reachable set, marking the first formal certification of deep learning-based reachability analysis.

This contribution has multiple implications. Theoretically, it provides the missing link between deep learning and formal verification in reachability settings, offering a path toward verifiable neural approximators. Practically, it enables scalable safety analysis in high-dimensional systems where classical methods fail, such as aerial swarms, robotic manipulators, and multi-agent systems, without compromising on correctness.

Research question: Chapter 7

What is the relationship between probabilistic backward reachable sets and viable sets in stochastic systems, and how can such sets be computed efficiently?

In stochastic systems, safety assurance requires fundamentally different tools than those used in deterministic or robust frameworks. While deterministic forced backward reachable sets define the states from which a system will inevitably reach an undesired region (an “avoid set”), the stochastic counterpart relaxes this notion by allowing for a probabilistic threshold. This is not only less conservative, but also essential in many practical domains where uncertainty is inherent and complete avoidance is infeasible. Applications such as air traffic control, stochastic model predictive control, and obstacle avoidance in uncertain environments demand such nuanced safety specifications.

To address this, the dissertation introduces and formalises the concept of the Probabilistic Forced Backward Reachable Set (PFBRs): the set of all states from

which the system, regardless of the applied control policy, will enter a designated avoid set with at least a specified probability within a finite time horizon. This probabilistic extension of the classical avoid problem makes the safety analysis framework significantly more applicable to systems influenced by random disturbances, such as sensor noise, environmental variability, or model uncertainties.

The work presents a rigorous mathematical formulation of PFBRs in discrete-time stochastic systems and establishes its fundamental properties, including measurability, convergence in time, and its relationship to stochastic viable sets. The chapter proves that under suitable conditions, the PFBRs converges in time and that it shares structural connections with the stochastic control viable set. The latter is a set in which the system can remain, with a high probability, for the whole time horizon—thereby linking safety in terms of avoidance and viable in a probabilistic setting.

From a computational standpoint, the dissertation provides an algorithmic framework for computing the PFBRs. For discrete space systems, an exact iterative algorithm is provided. For continuous state spaces, which are more typical in real-world applications, the chapter extends the method by discretising the state space, and proves that the resulting approximation converges to the true PFBRs as the grid resolution increases. The algorithm includes quantified error bounds due to discretisation, ensuring that the approximations can be interpreted within a well-defined tolerance. Furthermore, the computational complexity is analysed, confirming the tractability of the approach even for moderately large systems.

This contribution is especially valuable in scenarios where safety must be guaranteed not absolutely, but with high probability. It provides a theoretical and computational bridge between deterministic reachability and stochastic control, expanding the reach of formal safety analysis into domains previously beyond its scope. It also clarifies the nuanced relationship between reachability and viability in stochastic settings, offering a unified language for probabilistic safety verification.

Research question: Chapter 8

How can we compute stochastic backward reachable sets that ensure safety with high probability?

9

For autonomous systems operating under uncertainty, ensuring high-probability safety is critical to their acceptance in society and deployment in high-stakes environments. Classical tools such as deterministic and robust backward reachable sets provide rigorous guarantees, but often at the cost of conservatism. To overcome this, stochastic counterparts to these constructs—like the Stochastic Backward Reachable Set (SBRS)—have been introduced. These sets define the region of states from which a system can reach a target set within a time horizon with at least a certain probability. However, a critical limitation of the SBRS is that it does not ensure the system remains within the set during the transition. This lack of viability during execution means the SBRS alone is not sufficient for safety assurance

in many applications.

To address this fundamental limitation, this chapter introduces the Safe Probable Backward Reachable Set (SPBRS), a new set that guarantees not only that the system reaches a target set with high probability, but also that it remains within the set throughout the transition. Formally, the SPBRS is defined as the set of all initial states from which there exists a control policy such that the probability of reaching the target set within a given time horizon, while remaining within the SPBRS itself at all times, is at least a specified threshold. This definition endows the SPBRS with a probabilistic viability property—making it uniquely suitable for safety-critical applications where intermediate violations are unacceptable.

The dissertation provides two algorithms to compute the SPBRS in discrete-time stochastic systems. One algorithm guarantees the computation of the largest possible SPBRS, while the other is computationally more efficient. Both algorithms are proven to produce valid SPBRSs under the proposed definition. Furthermore, for discrete state spaces, these algorithms yield exact sets, while for continuous state spaces, the dissertation presents discretised approximations and formally proves convergence to the true SPBRS as the grid resolution is refined. The associated computational complexity is analysed for both algorithms, providing insight into their scalability and performance across different system types.

This contribution not only provides a theoretically grounded and computationally tractable framework for synthesising high-probability safe sets, but also addresses a longstanding gap in the literature: the lack of probabilistic reachability tools that ensure safety throughout the entire trajectory, not just at the endpoint. The SPBRS framework thus fills a critical void between reachability, viability, and practical safety assurances in uncertain environments.

In future work, the SPBRS can serve as a building block for probabilistic safety-aware control frameworks, including integration with Model Predictive Control (MPC), Control Barrier Functions (CBFs), and Control Lyapunov Functions (CLFs). Such integration would enable real-time, policy-driven safety enforcement in dynamic and uncertain environments. Additionally, the use of function approximation techniques to estimate the SPBRS could open the door to efficient computation even in high-dimensional systems, further extending the practical reach of this methodology.

9.3. LIMITATIONS

This section presents an assessment of the theoretical, computational, and practical limitations inherent to each contribution. These reflections are intended not only to contextualise the scope of the work but also to highlight promising directions for future research and refinement.

9.3.1. Critical Reflections: Limitations of Maneuverable viable Sets and Safe Sets Chapter 3 and 4

This chapter introduces two novel constructs that redefine how safety and viability are understood in autonomous systems: the maneuverable viable set (or maneuver set) and an improved definition of a set that ensures safety. It also proposes an algorithmic framework grounded in Hamilton-Jacobi reachability to compute these sets using the intersection of forward and backward reachable sets. The proposed framework presents a novel perspective on defining and computing safe regions for autonomous systems. Nonetheless, its reliance on deterministic, perfectly known models, high computational overhead, and absence of policy synthesis or robustness to uncertainty constrain its current impact.

High Computational Cost Due to Double Reachability Analysis

The algorithms require the computation of both forward and backward reachable sets—each of which is already computationally demanding in its own right. Their intersection to define the maneuver set introduces an additional layer of complexity, especially for high-dimensional systems. This makes real-time implementation or even offline computation for large-scale systems challenging, thus limiting practical scalability.

Conservative Approximation of the Safe Set

While the improved definition of the safe set provides simultaneous guarantees of viability and internal reachability, it does so at the cost of being more restrictive. As a result, the computed safe sets may be considerably smaller than traditional viable sets or maneuver sets. This conservatism, while ensuring stronger guarantees, may lead to underutilisation of the system's capabilities in applications where some degree of flexibility or risk is tolerable.

Applicability Limited to Perfectly Modelled Systems

The computation of both the maneuver set and the safe set relies on the assumption of accurate, fully known system dynamics. Any mismatch between the model used for reachability analysis and the true system dynamics can compromise the correctness of the computed sets. This limitation restricts the applicability of the framework in uncertain or learning-based environments where the system model is either unavailable or evolving.

Lack of Generalisation to Stochastic or Uncertain Domains

The current work is confined to deterministic systems. Neither the maneuver set nor the improved safe set is defined or proven to exist under stochastic dynamics. In practice, most systems are exposed to stochastic disturbances and noise, meaning that these deterministic guarantees, while strong, are not directly transferable to more realistic settings with stochastic disturbance without significant reformulation.

9.3.2. Critical Reflections: Limitations of Robust Avoid Backward Reachable Tube Chapter 5

The introduction of the Robust Avoid Backward Reachable Tube (RABRT) represents an advance in extending Hamilton-Jacobi (HJ) reachability analysis to systems subject to both bounded disturbances and unmodelled dynamics. By formulating the problem as a three-player differential game, this work offers a principled framework for synthesising robust safety certificates that remain valid even under adversarial and uncertain conditions. However, the RABRT comes with trade-offs in conservatism, computational complexity, and practical tractability.

Conservatism Due to Worst-Case Assumptions

The RABRT framework assumes that both the disturbance and the unmodelled dynamics act adversarially and jointly attempt to drive the system into unsafe regions. While this formulation ensures strong safety guarantees, it also induces conservatism. In practice, not all uncertainty behaves in a worst-case fashion. This could lead to overly pessimistic safe sets that limit the available state space unnecessarily, reducing system performance or feasibility in real-world deployments.

Applicability Limited to Bounded, Deterministic Uncertainty

Although the formulation captures unmodelled dynamics as an adversarial input, it assumes that the disturbances and unmodelled components are bounded and deterministic. This limits its applicability in settings involving stochastic or unbounded noise, such as those encountered in finance, epidemiology, or autonomous navigation in dynamic urban environments. In such cases, a stochastic formulation (e.g. SPBRS) may be more appropriate.

Computational Burden for High-Dimensional Systems

The RABRT, like all level-set-based reachability methods, suffers from the curse of dimensionality. The additional complexity introduced by the third player compounds the computational challenges already present in solving HJ-Isaacs PDEs. Computing RABRTs remains infeasible for systems of dimension higher than five to seven without substantial simplifications or problem-specific structure.

Static Disturbance Bounds and Lack of Online Adaptivity

The framework presumes that the bounds on disturbances and unmodelled dynamics are known a priori and remain constant throughout the analysis. In real systems, however, uncertainty can vary over time or across operating conditions. Without mechanisms for online adaptation or re-certification, the RABRT may either underestimate risk (if the bounds were underestimated) or be too conservative (if overestimated).

9.3.3. Critical Reflections: Limitations of Certified Approximate Reachability (CARE) Chapter 6

The Certified Approximate Reachability (CARE) framework presents a significant advancement in reconciling the scalability of neural approximators with the rigour of formal verification in reachability analysis. By leveraging Satisfiability Modulo Theory (SMT) solvers and Counterexample-Guided Inductive Synthesis (CEGIS), the framework is able to provide formal error bounds on neural network-based approximations of reachable sets. However, several theoretical and computational challenges remain. Most notably, the scalability of the verification stage and the lack of dynamic feedback during synthesis limit its immediate applicability to large-scale, real-world systems.

Scalability of the Formal Verification Pipeline

While the use of neural networks enables the approximation of reachable sets in high-dimensional spaces, the formal certification process using SMT solvers (e.g., dReal) is computationally expensive. In particular, the verification stage becomes a bottleneck as the size of the network or the complexity of the domain increases. As a result, CARE currently balances scalability in training with a verification procedure that does not yet scale accordingly.

Localised Error Guarantees and Coverage Limitations

The formal guarantees provided by CARE are typically local in nature—bounded within a specified region of interest or a discretised domain. While this is acceptable in many applications, it leaves open the possibility of unverified (and potentially unsafe) behaviour in regions not explicitly included in the verification set. The guarantees are thus conditional on the verification domain and do not necessarily imply global correctness across the entire state space.

Sensitivity to Network Architecture and Training Dynamics

The framework assumes that the neural network approximator is sufficiently expressive and well-trained to approximate the true value function. However, neural networks are known to be sensitive to hyperparameter choices, training loss landscapes, and data sampling strategies. A poorly trained model may converge to a function that appears accurate but exhibits subtle violations of the HJ PDE, thus increasing the burden on the verification part of the pipeline.

Limited Applicability to Non-HJ Reachability Problems

CARE is specifically designed around the structure of the Hamilton-Jacobi PDE and the notion of viscosity solutions. While this covers a broad class of reachability problems, it does not easily extend to other verification domains such as stochastic or probabilistic reachability (e.g., SPBRS), where the underlying mathematical formalism differs. Generalising the CARE methodology to other classes of dynamic programming equations remains a non-trivial challenge.

Non-stochastic Guarantees

The framework currently outputs a binary certification: the network either satisfies the error bound or it does not. This binary certification can be limiting in practical settings, where a “confidence gradient” would allow systems to operate more flexibly under partial or probabilistic guarantees. Developing probabilistic or risk-aware variants of CARE would enhance its utility for applications with varying safety requirements.

Dependence on Overapproximation of Error Bounds

To remain sound, the error bounds derived in CARE must conservatively enclose the true PDE residual. In doing so, they may be overly pessimistic, leading to tighter-than-necessary reachable set approximations. This conservatism, while necessary for formal guarantees, may limit the use of the certified sets in control synthesis, especially in tight or constrained operating environments.

9.3.4. Critical Reflections: Limitations of the Stochastic Avoid Problem with Computation Chapter 7

The stochastic avoid problem addresses one of the most foundational safety concerns in uncertain dynamical systems: how to ensure that a system avoids undesirable (unsafe) states with high probability over a given time horizon. This chapter introduces a formal framework for defining and computing the Probabilistic Forced Backward Reachable Set—the stochastic counterpart of the forced backward reachable set—and develops computationally tractable algorithms for its estimation. It makes a meaningful theoretical and computational contribution to safety analysis under uncertainty. Still, several limitations remain, particularly in its applicability to continuous-time systems, trajectory-wide safety guarantees, policy synthesis, model robustness, and computational scalability.

Applicability to Discrete-Time Systems Only

The theoretical formulation and algorithms developed are explicitly designed for discrete-time stochastic systems. While many control architectures operate in discrete time, this assumption restricts direct application to systems naturally modelled in continuous time—particularly in aerospace, finance, or certain biological processes. Extending the framework to continuous-time stochastic dynamics would require more sophisticated tools, such as stochastic differential equations or measure-theoretic formulations, which remain an open and non-trivial direction.

No Guarantees of Trajectory-Wide Safety

While the probabilistic forced backward reachable set provides guarantees that the system will eventually reach an avoid set with a certain probability under any control policy, it does not ensure that the system remains outside unsafe regions at all times during the trajectory. This limits its utility in scenarios where safety must be enforced continuously, not just asymptotically.

Heavy Dependence on Model Accuracy

The correctness of the computed probabilistic reachable sets hinges on an accurate stochastic model of the system's dynamics, including the transition probabilities. Any mismatch between the model used and the true system behaviour—such as unmodelled non-Gaussian noise or environment-dependent stochasticity—can undermine the reliability of the safety guarantees. This limits the method's robustness in data-driven, learned, or partially known systems.

Absence of Control Synthesis

The framework is focused on set computation, and while it formalises the probability that a system will violate safety under any policy, it does not provide or synthesise a control strategy that minimises or shapes this probability. In practice, it is often not enough to know that safety violations are probable—an actionable policy is needed to avoid such outcomes. This disconnect limits immediate control-theoretic application and leaves a gap between analysis and deployment.

Computational Tractability Limited by Discretisation

The algorithms proposed for continuous state spaces rely on discretising the space and simulating the stochastic evolution of the system. Although convergence guarantees are provided, the method becomes computationally intensive as the dimensionality increases or finer resolutions are required. In high-dimensional systems, the discretisation step can become infeasible without additional structure (e.g., sparsity, symmetry) or approximation techniques.

Rigid Dependence on a Predefined Avoid Set

The formulation presumes that the avoid set is well-defined and fixed over time. However, in many real-world applications—such as navigation through dynamic or partially known environments—the avoid set may be time-varying, probabilistic, or defined implicitly through sensor data or learned risk functions. The current framework does not directly support such generalisations, limiting its applicability in adaptive or perception-driven settings.

9.3.5. Critical Reflections: Limitations of Safe Probable Backward Reachability Chapter 8

While the introduction of the Safe Probable Backward Reachable Set (SPBRS) presents a novel and rigorously grounded method for certifying probabilistic safety in stochastic systems, its full real-world deployment will require addressing challenges of scalability, adaptability, and policy integration.

Discrete-Time Limitations and Generalisation Challenges

The current formulation and accompanying algorithms are designed specifically for discrete-time stochastic systems. While this makes the analysis mathemati-

cally tractable and aligns with many digital control systems, it may limit direct applicability in continuous-time domains. Extending the SPBRS to continuous-time systems would likely involve solving stochastic partial differential equations or employing time-discretisation approximations, both of which introduce additional complexity and potential approximation error.

Dependence on State-Space Discretisation

The computational algorithms for continuous-state systems rely on discretising the state space. Although convergence of the estimated SPBRS to the true SPBRS is formally proved as grid resolution increases, in practice this discretisation introduces significant computational burden, especially in higher-dimensional systems (more than four or five dimensions). This limitation is a manifestation of the well-known curse of dimensionality and poses challenges for scalability beyond low to moderately sized systems.

Assumption of Perfect Model Knowledge

The analysis assumes that the system dynamics (i.e. transition probabilities) and noise models are fully known and accurately specified. In many real-world applications, such complete model fidelity is not realistic. Model mismatch—especially in the stochastic transition kernel—can lead to mischaracterisation of the true safe region, thereby undermining the safety guarantees offered by the SPBRS.

No Explicit Handling of Time-Varying or Uncertain Targets

The current formulation of the SPBRS assumes a fixed target set $\mathcal{L}$ over a finite time horizon. However, many real-world systems, especially in robotics and air traffic control, operate in environments with dynamically evolving goals and constraints. Extending the SPBRS framework to handle such nonstationary or uncertain targets remains an open challenge.

Lack of Policy Synthesis Within the Framework

While the SPBRS guarantees the existence of a policy that satisfies the probabilistic safety requirement, the framework does not directly synthesise the corresponding control policy. This limits its immediate integration with feedback control architectures. A natural next step is to develop constructive methods for synthesising optimal or robust feedback policies that explicitly keep the system within the SPBRS while ensuring target reachability.

9.4. FUTURE WORK

The future directions outlined in this section are motivated by a desire to extend the applicability, scalability, and practical impact of the methodologies developed in this work. These include enhancing computational efficiency, improving policy

synthesis, adapting the frameworks to uncertain or time-varying models, and integrating the proposed methods into real-time control architectures. Some directions also aim to relax conservative assumptions, while others explore opportunities to bridge formal guarantees with data-driven learning approaches.

By pursuing these extensions, future work can continue to close the gap between formal verification and practical deployment—ensuring that autonomous systems can operate safely, flexibly, and reliably in increasingly uncertain and complex environments.

9.4.1. Future Work: Maneuverable viable Sets and Safe Sets Chapter 3 and 4

While the algorithmic framework developed for maneuverable viable and safe sets provides a strong theoretical foundation, its full potential lies in its adaptation to uncertainty, scalability, and integration with control and learning frameworks. These directions represent not only natural continuations of this work but also critical steps toward closing the gap between formal safety guarantees and real-world autonomy.

Policy Synthesis from Certified Sets

While the theoretical existence of control strategies is guaranteed for both the maneuver set and the improved safe set, no explicit method for synthesising such policies has been proposed. Developing constructive algorithms for control synthesis—especially those that provide feedback policies guaranteeing transitions within the set—would bridge the gap between offline analysis and online deployment. A natural direction is to integrate these sets with online control architectures such as Model Predictive Control (MPC) [2] or Control Barrier Function (CBF)–based safety filters [3]. Doing so would enable real-time enforcement of maneuverability or viability properties.

Generalisation to High-Dimensional and Hybrid Systems

The computational burden of computing maneuverable sets grows significantly with system dimensionality due to the need for both forward and backward reachability analysis. Future work could explore scalable approximations using decomposition techniques, sampling-based reachability, or neural approximators [4], [5],[6]. Some ideas in this direction include approaches similar to Deepreach [4], or using re-enforcement learning [7] to approximate the value function.

Data-Driven and Learning-Based Estimation

For systems where analytical models are unavailable, learning-based estimation of maneuverable or safe sets presents a compelling alternative. One promising direction involves using trajectory data or simulator rollouts to empirically estimate forward and backward reachability, from which the maneuver set can be inferred. Combining such methods with formal guarantees—e.g., via probabilistic bounds or

statistical learning theory—could open new frontiers for model-free safety analysis [8].

Time-Varying and Goal-Driven Extensions

In practical scenarios, safety requirements are often time-varying or goal-dependent. Extending the definitions of maneuverable and safe sets to incorporate time-dependent constraints, moving targets, or dynamic avoid regions would enhance their relevance in applications such as motion planning [9], pursuit-evasion games [10], and real-time navigation in dynamic environments [11].

9.4.2. Future Work: Robust Avoid Backward Reachable Tube Chapter 5

While the RABRT framework provides a robust foundation for safety verification under model uncertainty and adversarial disturbance, its full potential can be unlocked by addressing the challenges of scalability, stochasticity, and integration with data-driven modelling.

Extension to Stochastic and Probabilistic Uncertainty Models

While the current formulation accounts for bounded deterministic uncertainty and adversarial dynamics, real-world systems often experience unbounded stochastic disturbances. A natural extension would be to generalise the RABRT framework to stochastic dynamics, leading to a synthesis of robust and probabilistic reachability. This would require reformulating the Hamilton-Jacobi-Isaacs equation under a stochastic framework [12], possibly incorporating expectations or probabilistic risk metrics into the value function.

Scalable Numerical Methods and Decomposition Techniques

The computational burden of solving three-player HJI PDEs grows rapidly with system dimensionality. Future work should focus on approximate and scalable solvers, possibly using domain decomposition [13], tensor-train approximations [14], or deep learning-based PDE solvers [15],[16]. Exploring reduced-order models or structure-exploiting formulations could enable robust reachability analysis in high-dimensional applications such as robotic manipulators or UAV swarms.

Application to Hybrid and Switching Systems

The RABRT framework is currently limited to systems with continuous dynamics. Many modern autonomous systems exhibit hybrid or switching behaviour (e.g., drones switching between flight modes, or vehicles transitioning between highway and urban driving). Extending the robust reachability analysis to such hybrid systems—possibly via mode-dependent value functions or switched HJI formulations [17].

9.4.3. Future Work: Certified Approximate Reachability (CARE) Chapter 6

The Certified Approximate Reachability (CARE) framework presents a novel synthesis of deep learning and formal methods to certify neural approximations of reachable sets. It introduces a methodology to quantify the error between a neural value function and the true solution of the Hamilton-Jacobi PDE, offering formal guarantees in regions of interest. While the framework makes a significant contribution, it also opens several promising avenues for further research—particularly around improving expressivity, generalisation, and the tractability of verification.

Lipschitz-Based Certification for Scalable Guarantees

One promising direction is the incorporation of Lipschitz continuity bounds into the certification framework [18]. Neural networks with bounded Lipschitz constants can offer uniform guarantees on function deviation, enabling the derivation of global or local error margins without relying solely on expensive pointwise SMT checks. Such Lipschitz-based analysis may allow for over- and under-approximations of the reachable set that hold over continuous regions, offering a powerful alternative or complement to the current CEGIS-based verification.

Use of Kolmogorov–Arnold Neural Architectures

Standard feedforward neural networks (e.g. MLPs or ReLU-based DNNs) often require substantial capacity and careful training to accurately approximate non-smooth viscosity solutions of HJ PDEs. Kolmogorov–Arnold Networks (KANs), which are built on the Kolmogorov–Arnold representation theorem, offer stronger expressivity for approximating multivariate functions via compositions of univariate functions [19]. Future work could explore using KANs as an alternative architecture for value function learning, with the potential for better approximation accuracy and smoother loss landscapes. Additionally, their structure may allow tighter analytical bounds—potentially facilitating symbolic or hybrid verification methods in CARE.

Extension to Stochastic or Uncertain Systems

CARE is presently focused on certifying reachable sets for deterministic systems governed by HJ PDEs. Extending the framework to stochastic dynamics—such as those governed by Hamilton-Jacobi-Bellman (HJB) [12] or Fokker-Planck equations—would enable certification under probabilistic uncertainty [20]. This may require defining new loss functions, introducing expectation-based PDE relaxations, or integrating probabilistic SMT solvers to accommodate stochastic terms in the dynamics.

9.4.4. Future Work: Stochastic Avoidance and Safe Probable Backward Reachability Chapter 7 and 8

The probabilistic forced backward reachable set and the Safe Probable Backward Reachable Set (SPBRS) formalise stochastic safety guarantees by relaxing the conservatism of robust approaches, while still offering precise, probability-aware certification for systems subject to random disturbances. However, their full potential will be realised through extensions to control synthesis, high-dimensional scalability and continuous-time dynamics. These directions offer fertile ground for bridging theoretical safety guarantees with real-time, risk-aware decision-making in uncertain environments.

Extension to Continuous-Time Stochastic Systems

Both the stochastic avoid and SPBRS frameworks are currently formulated in discrete-time. However, many physical systems—such as those in aerospace, energy, and continuous control—are better modelled in continuous time [21]. Future work could adapt the reachability definitions and algorithms to stochastic differential equations (SDEs), possibly by leveraging stochastic viscosity solutions [22], [23]. This extension would close the gap between formal theory and high-fidelity physical modelling.

Policy Synthesis for Safe Execution

The SPBRS is a natural candidate for a terminal constraint or viable set in Stochastic MPC [24]. Integrating SPBRS into MPC pipelines would yield formal guarantees on recursive feasibility and safety under uncertainty. Moreover, SPBRS-based constraints could improve the interpretability and reliability of MPC in applications such as path planning, air traffic management, or automated vehicles.

Algorithmic Scalability and Learning-Based Approximations

While the current algorithms for computing the stochastic avoid and SPBRS sets are theoretically sound and converge with grid refinement, they face scalability issues in high-dimensional systems. One direction is to develop sampling-based or data-driven approximations of the value function—e.g., using Gaussian Processes [25], neural networks [4], [7], or randomized value iteration—with statistical confidence bounds.

Robustness to Model Uncertainty and Online Adaptivity

Both frameworks assume that the transition probabilities of the system are precisely known. In practice, these probabilities may be uncertain, time-varying, or learned from data. A promising direction is to develop robust stochastic reachability frameworks where the transition model lies within a known ambiguity set, such as a confidence region or a Wasserstein ball. This would ensure safety guarantees even in the presence of imperfect knowledge [26], [27].

9.5. RESEARCH IMPACT

This dissertation contributes fundamentally to the theory and computation of safety guarantees for autonomous dynamic systems, offering a set of general-purpose, mathematically rigorous tools that are both novel in formulation and adaptable to real-world contexts. While the work is grounded in control theory and formal methods, its broader implications extend beyond academic boundaries. The potential impact of this research spans several technological domains and addresses pressing societal challenges at the interface of autonomy, safety, and trust.

9.5.1. Autonomous Transportation Systems

In domains such as autonomous driving, drone navigation, and aerospace systems, ensuring the safety of motion under uncertainty is paramount. The definitions and algorithms developed in this work—particularly the maneuverable safe sets, robust avoid tubes, and the SPBRS—provide certifiable guarantees that can inform decision-making modules in autonomous vehicles. By supporting both deterministic and stochastic safety guarantees, this thesis contributes to the formal certification of autonomy stacks, potentially accelerating regulatory approval, public adoption, and integration into transportation infrastructure.

9.5.2. Robotics and Human–Robot Interaction

This dissertation contributes to the ongoing discourse on trustworthy AI by demonstrating how learning-based methods—such as neural networks—can be rigorously certified when used to approximate safety-critical functions. The CARE framework offers a hybrid of symbolic and data-driven methods, providing a path forward for verifiable AI in control systems. As concerns about the safety and reliability of black-box AI systems continue to grow, this work shows that formal guarantees can, in fact, be extended to neural approximators used in safety-critical loops.

9.5.3. Aerospace and Mission-Critical Systems

In aerospace applications, such as spacecraft trajectory planning, re-entry operations, and air traffic control, safety verification under uncertainty is a regulatory and operational necessity. The backward reachability tools, particularly those extended to the stochastic domain, can enable safer planning in these contexts. Moreover, the use of certified neural approximators could allow for faster onboard computation while maintaining formal guarantees—something traditional grid-based methods cannot achieve.

9.5.4. Healthcare and Biomedical Applications

In systems where safety is defined in probabilistic rather than absolute terms—such as in treatment planning, robotic surgery, or real-time diagnostic decision support—the stochastic reachability and viability frameworks developed here could play a crucial role. They provide a way to reason about patient-specific uncertainty and guarantee safe interventions within bounded risk thresholds.

9.5.5. Impact on the Foundations of Control Theory

From a theoretical standpoint, the dissertation contributes new mathematical structures—such as maneuverable viable sets and SPBRs—that refine the traditional notions of viability, reachability. These constructs enrich the vocabulary of control theory and open new directions in the design of safety certificates that account not only for where a system can remain, but also where it can go, and how reliably it can get there.

9.5.6. Societal Trust in Autonomous Technologies

Finally, this dissertation contributes to building the infrastructure of trust required for the societal deployment of autonomous systems. The proposed frameworks enable designers, regulators, and users to quantify and communicate risk in formal terms. By translating theoretical safety guarantees into computationally tractable and certifiable tools, this work supports the creation of systems that are not only autonomous, but verifiably safe.

9.6. FINAL REMARKS

This dissertation brings to a close several years of research that have been both challenging and rewarding. It presents a collection of ideas and tools aimed at improving how we understand, compute, and certify safety in autonomous dynamical systems—ranging from maneuverable viable sets to stochastic safety guarantees and certified learning-based approximations. While the contributions are specific to reachability and verification, I hope they offer useful building blocks for others working on similar problems in control, robotics, and formal methods.

REFERENCES

- [1] I. M. Mitchell *et al.* “A toolbox of level set methods”. In: *UBC Department of Computer Science Technical Report TR-2007-11* 1 (2007), p. 6.
- [2] K. P. Wabersich and M. N. Zeilinger. “Linear model predictive safety certification for learning-based control”. In: *2018 IEEE Conference on Decision and Control (CDC)*. IEEE. 2018, pp. 7130–7135.
- [3] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert. “Robust control barrier–value functions for safety-critical control”. In: *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE. 2021, pp. 6814–6821.
- [4] S. Bansal and C. J. Tomlin. “Deepreach: A deep learning approach to high-dimensional reachability”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2021, pp. 1817–1824.
- [5] Y. Zhang, X. Liang, D. Li, S. S. Ge, B. Gao, H. Chen, and T. H. Lee. “Barrier Lyapunov function-based safe reinforcement learning for autonomous vehicles with optimized backstepping”. In: *IEEE Transactions on Neural Networks and Learning Systems* 35.2 (2022), pp. 2066–2080.
- [6] M. Raissi, P. Perdikaris, and G. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378 (Feb. 2019), pp. 686–707. ISSN: 0021-9991. DOI: [10.1016/j.jcp.2018.10.045](https://doi.org/10.1016/j.jcp.2018.10.045).
- [7] M. Ganai, S. Gao, and S. Herbert. “Hamilton-jacobi reachability in reinforcement learning: A survey”. In: *IEEE Open Journal of Control Systems* (2024).
- [8] W. Zhao, T. He, and C. Liu. “Model-free safe control for zero-violation reinforcement learning”. In: *5th Annual Conference on Robot Learning*. 2021.
- [9] S. Kousik. “Reachability-based trajectory design”. PhD thesis. University of Michigan, 2020.
- [10] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-jacobi reachability: A brief overview and recent advances”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE. 2017, pp. 2242–2253.
- [11] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin. “Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking”. In: *IEEE Transactions on Automatic Control* 66.12 (2021), pp. 5861–5876.
- [12] J. Yong and X. Y. Zhou. *Stochastic controls: Hamiltonian systems and HJB equations*. Vol. 43. Springer Science & Business Media, 1999.

- [13] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin. “Decomposition of reachable sets and tubes for a class of nonlinear systems”. In: *IEEE Transactions on Automatic Control* 63.11 (2018), pp. 3675–3688.
- [14] A. Abate, A. Edwards, and M. Giacobbe. “Neural abstractions”. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 26432–26447.
- [15] C. C. Aggarwal *et al.* *Neural networks and deep learning*. Vol. 10. 978. Springer, 2018.
- [16] M. Tayal, A. Singh, S. Kolathaya, and S. Bansal. “A Physics-Informed Machine Learning Framework for Safe and Optimal Control of Autonomous Systems”. In: *arXiv preprint: 2502.11057* (2025). DOI: [10.48550/arXiv.2502.11057](https://doi.org/10.48550/arXiv.2502.11057). arXiv: [2502.11057](https://arxiv.org/abs/2502.11057) [cs.R0].
- [17] A. K. Shah and D. M. Adhyaru. “Clustering based multiple model control of hybrid dynamical systems using HJB solution”. In: *Applied Soft Computing* 31 (2015), pp. 103–117.
- [18] Y. Huang, H. Zhang, Y. Shi, J. Z. Kolter, and A. Anandkumar. “Training certifiably robust neural networks with efficient local lipschitz bounds”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 22745–22757.
- [19] Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljačić, T. Y. Hou, and M. Tegmark. “Kan: Kolmogorov-arnold networks”. In: *arXiv preprint arXiv:2404.19756* (2024).
- [20] N. M. Boffi and E. Vanden-Eijnden. “Probability flow solution of the fokker-planck equation”. In: *Machine Learning: Science and Technology* 4.3 (2023), p. 035012.
- [21] H. Pham. *Continuous-time stochastic control and optimization with financial applications*. Vol. 61. Springer Science & Business Media, 2009.
- [22] R. Buckdahn and J. Li. “Stochastic differential games and viscosity solutions of Hamilton–Jacobi–Bellman–Isaacs equations”. In: *SIAM Journal on Control and Optimization* 47.1 (2008), pp. 444–475.
- [23] R. Buckdahn and J. Ma. “Stochastic viscosity solutions for nonlinear stochastic partial differential equations. Part I”. In: *Stochastic processes and their applications* 93.2 (2001), pp. 181–204.
- [24] A. Mesbah. “Stochastic model predictive control: An overview and perspectives for future research”. In: *IEEE Control Systems Magazine* 36.6 (2016), pp. 30–44.
- [25] D. J. MacKay *et al.* “Introduction to Gaussian processes”. In: *NATO ASI series F computer and systems sciences* 168 (1998), pp. 133–166.
- [26] I. Yang. “Wasserstein distributionally robust stochastic control: A data-driven approach”. In: *IEEE Transactions on Automatic Control* 66.8 (2020), pp. 3863–3870.
- [27] A. Neufeld and J. Sester. “Robust Q-learning algorithm for Markov decision processes under Wasserstein uncertainty”. In: *Automatica* 168 (2024), p. 111825.

ACKNOWLEDGEMENTS

The completion of a PhD is never just about the technical work. The truth is that none of it happens in isolation. I have been fortunate to work in a supportive environment, surrounded by colleagues and friends whose feedback, encouragement, and companionship made the hard parts easier and the good parts more enjoyable. From conversations over equations and tea to spontaneous whiteboard breakthroughs to shared frustrations over failing code, this journey has been made meaningful by the people who shared it.

I would especially like to thank my supervisor, Prof. Coen de Visser, who has guided me with a rare balance of structure and freedom. His support enabled me to explore ambitious ideas, while his thoughtful feedback helped sharpen my thinking. His encouragement was consistent, and his trust in me gave me the space to grow into this work at my own pace.

I am also deeply grateful to Dr. Anahita Jamshidnejad, whose guidance and insight were invaluable throughout this journey. Her thoughtful mentorship shaped many aspects of my research.

My sincere thanks to Prof. Alessandro Abate, who hosted me during my research stay in Oxford and co-authored one of the central papers in this dissertation. Working with him and with Dr. Nikolaus Vertovec and Yannik Schnitzer, my co-authors on that same project was intellectually stimulating and immensely enjoyable.

A special thank you to Jasper van Beers, who has been part of this project since the beginning. We worked together closely on every chapter of this dissertation, and his input and collaboration have been integral to everything that follows. I would also like to thank Rowenna, for her kindness and support. She helped me greatly in many practical ways, especially by making sure I submitted things on time and by helping me manage my time effectively.

My deepest thanks go to my family. To my parents, Ravi and Jyotsna, thank you for your patience, faith, and unwavering support. To my wife, Josephine, who stood by me with compassion and constancy, listened, proofread, and anchored me throughout this process. And to my brother, Harshul, thank you for your insight, humour, and level-headed perspective, just what I needed when things felt overwhelming.

Looking back, it is clear to me that this dissertation is not the result of one person's effort, but it rests on the care, time, and encouragement of many others. I'm deeply aware that none of it would have been possible without that support, and for that, and for the community that made it possible, I am sincerely and profoundly grateful.

CURRICULUM VITÆ

Prashant Solanki

14-01-1993 Born in Delhi, India.

EDUCATION

2012–2017 Aerospace Engineering (BTech.) + Avionics (MTech.)
Amity University, Noida, Uttar Pradesh, India
Thesis: Flight controller design using dynamic inversion and neuro-adaptive techniques

2018–2020 MSc. Aerospace Engineering
TU Delft
Thesis: Attitude Estimation of a Quadcopter with one fully damaged rotor using on-board MARG Sensors

2021–2025 PhD. Aerospace Engineering
TU Delft
Thesis: Safe operation region estimation and certification
Promotor: Dr. ir. C.C. de Visser
Copromotor: Dr. A. Jamshidnejad

EXPERIENCE

2017 Research Assistant
Amity University, Noida, Uttar Pradesh, India

2021 Research Assistant
TU Delft, Netherlands

Since 2025 Postdoctoral researcher at TU Delft
Delft, NL

LIST OF PUBLICATIONS

1. P. Solanki and C. C. de Visser. “Attitude Estimation of a Quadcopter with one fully damaged rotor using on-board MARG Sensors”. In: *AIAA SCITECH 2022 Forum*. 2022, p. 0857
2. P. Solanki, J. J. Van Beers, A. Jamshidnejad, and C. C. De Visser. “A novel algorithmic approach to obtaining maneuverable control-invariant sets”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2024, pp. 9916–9922
3. J. J. Van Beers, P. Solanki, and C. C. De Visser. “A novel metric for detecting quadrotor loss-of-control”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2024, pp. 15570–15576
4. P. Solanki, N. Vertovec, Y. Schnitzer, J. Van Beers, C. de Visser, and A. Abate. “Certified approximate reachability (care): Formal error bounds on deep learning of reachable sets”. In: *2025 IEEE 64th Conference on Decision and Control (CDC)*. IEEE. 2025, pp. 3907–3912
5. I. El-Hajj, J. van Beers, and P. Solanki. “ZonoReach: A Reachability-Guided Controller Using Zonotopes and Local Hamilton–Jacobi Analysis”. In: *International Journal of Control, Automation and Systems* 23.11 (2025), pp. 3199–3208
6. P. Solanki, I. El-Hajj, J. van Beers, E.-J. van Kampen, and C. de Visser. “Formalizing the Relationship between Hamilton–Jacobi Reachability and Reinforcement Learning”. Manuscript under review at *Automatica*. 2026
7. P. Solanki, J. van Beers, A. Jamshidnejad, and C. de Visser. “Safe Probable Backward Reachable Sets: Theory and Algorithms”. Manuscript under review at *IEEE Transactions on Automatic Control*. 2026
8. P. Solanki, J. van Beers, A. Jamshidnejad, and C. de Visser. “Robust Avoid Backward Reachable Tube”. Manuscript under review at *IEEE Control Systems Letters*. 2026
9. P. Solanki, J. van Beers, A. Jamshidnejad, and C. de Visser. “The Stochastic Avoid Problem with Computation”. Manuscript under review at *IEEE Transactions on Automatic Control*. 2026
10. P. Solanki, J. van Beers, A. Jamshidnejad, and C. de Visser. “A Novel Approach to Safe Operational Envelope and Its Estimation”. Manuscript under review at *Journal of Guidance, Control, and Dynamics*. 2026
11. P. Solanki, I. El-Hajj, J. J. van Beers, E.-J. van Kampen, and C. C. de Visser. “Certifying Hamilton–Jacobi Reachability Learned via Reinforcement Learning”. Manuscript under review at *Automatica*. 2026

Propositions

accompanying the dissertation

Safe operation region estimation and certification

by

Prashant SOLANKI

1. Without knowledge of the safe operation region, executing any task amounts to gambling. Once the system enters an unsafe region, no controller can guarantee recovery or prevent failure. [this thesis]
2. Without mathematically sound guarantees, safety estimation reduces to guesswork. In the absence of rigorous proofs, any claim of safety remains speculative. [this thesis]
3. Neural networks can approximate the solutions to high-dimensional reachability problems, but they are admissible only when their error is controlled and certified. [this thesis]
4. In control synthesis, round-trip reachability is a necessary condition for point-to-point safety, and surpasses classical viability conditions when system reversibility or returnability is essential. [this thesis]
5. In stochastic safety, ambiguity in outcome is acceptable; ambiguity in definition is not.
6. Taking risks is a luxury reserved for those with infinite lives and nothing to lose. In safety-critical systems, we have neither. Where failure costs lives, guesswork is not courage, it is negligence.
7. An unsafe system is often tolerated as long as it maximises shareholder value.
8. Science is a tool. In the hands of the wise, it saves lives. In the hands of capital, it surveils, displaces, and destroys. Let this research be for protection, not policing; for movement, not militarisation.
9. Few experiences illustrate the limits of self-regulation better than missing a time management workshop because of bad time management.
10. Let it be known: no unsafe set, no failed proof, and no reviewer comment ever survived the healing powers of bubble water.

These propositions are regarded as opposable and defensible, and have been approved as such by the promotor Dr. ir. C.C. de Visser and the copromotor Dr.

A. Jamshidnejad.