Safe Flight Envelope Uncertainty Quantification

Using probabilistic reachability analysis

R. van den Brandt

Faculty of Aerospace Engineering

TU Delft

Safe Flight Envelope Uncertainty Quantification

Using probabilistic reachability analysis

by

R. van den Brandt

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday December 19, 2017 at 14:00.

Student number:	4150562	
Project duration:	November 2, 2016 – December 19, 2017	
Thesis committee:	Dr. Q.P. Chu,	TU Delft, Chair
	Dr. ir. C.C. de Visser,	TU Delft, Supervisor
	Dr. ir. H.G. Visser,	TU Delft, Air Transport & Operations
	Ir. Y. Zhang,	TU Delft

TU Delft

Preface

This report contains the result of work that was performed over a one year period as part of the master thesis assignment at the faculty of Aerospace Engineering. The research into Safe Flight Envelope (SFE) prediction is part of a larger research scope performed by the department of Control & Simulation. This scope includes damaged aircraft modeling, flight envelope prediction and stall model identification, in a continuous effort to improve aviation safety.

This report is written from an engineering standpoint. No prior knowledge of the reachability analysis was assumed. Readers which are already familiar with the field of SFE prediction and reachability analysis may choose to pass the first two chapters. The author feels that it is more valuable to create a thorough understanding of the problem and limitations by means of examples and illustrations, instead of mathematical proof. Throughout this report, the reader is referred to articles and papers which provide a more rigorous explanation.

Over the duration of this assignment, I received a lot of help and support from various people. I would especially like to thank Coen de Visser for his guidance throughout this thesis project, and for letting me choose my own path. Furthermore, Ye Zhang and Sihao Sun for their help in understanding the problem at hand and the support and usage of the quadrotor model. Of course the upper-house which kept me company during the duration of this project. I would like to thank my friends and roommates for all the good times, and taking my mind of graduating once in a while. And of course my family, which went through some pretty rough times while I was busy graduating.

R. van den Brandt
Delft, December 2017

Summary

Aviation is widely considered to be the safest mode of transport. While accidents are scarce, when they do happen, they have a large impact. The biggest share of these accidents can be accounted to Loss of Control (LOC). Envelope protection has been proposed as a method to prevent LOC related accidents. For such a system to be effective, knowledge of the Safe Flight Envelope (SFE) is required. Reachability analysis is employed to determine the set of all states that can safely be attained by the aircraft. Although theoretically accurate, some states may be practically unattainable under the influence of turbulence. An extension to the reachability analysis is introduced as a way to quantity the effects of turbulence on the SFE. Furthermore, a new Path Integral Monte Carlo (PIMC) approach is proposed in order to validate the results from the probabilistic reachability analysis.

This method is applied to a simplified aircraft model to compare the differences between the deterministic SFE and the stochastic SFE. The deterministic SFE turns out to be significantly larger than the $3\,\sigma$ contour. When comparing the area of both envelopes at different roll angles it can be observed that the stochastic envelope ranges from 71.4 % till 50.8 % of the area. This means that at large roll angles the SFE is almost half that of the case without disturbances. Furthermore, for some states in the deterministic SFE the transition probability under influence of external turbulence is 0 %, meaning that there is no possibility of returning to the initial set.

A longitudinal quadrotor model is considered as a second example. This model features a full three Degrees-Of-Freedom (DOF) dynamic model combined with a non-linear parametrized aerodynamic model. The high number of states, combined with the unstable nature of the open-loop model, add to the computational complexity making it difficult to solve. To overcome these limitations, the time-horizon was kept low and the maximum pitch command was limited to 10 %. Since the aerodynamic model is non-linear, it is not possible to separate the drift and diffusion terms as required for the probabilistic reachability analysis. Therefore, it is recommended to determine the gust derivative by locally linearizing the aerodynamic model. Furthermore, the addition of rotor spin-up effects could be achieved by adding additional states for the rotor spin-rate and changing the input to a torque.

Finally, some methods were explored in an effort to reduce the computational complexity. Under the assumption that the process remain Gaussian over a small time step, the PIMC method can be simplified using a one step-ahead covariance prediction. The achievable speedup is roughly equal to the amount of samples of the PIMC method. Furthermore, in an attempt to break the ‘curse of dimensionality’ a quadrature point method is introduced. This method can provide an estimation of the transition probability for a single state, thus circumventing the scaling issues associated with Eulerian methods. Although effective, it only provides a rough estimate of the true reachability probability.

Contents

List of Abbreviations

DOF Degrees-Of-Freedom. v, 3, 43, 46, 57

EOM Equations of Motion. 43, 61

GBM Geometric Brownian Motion. xi, 23, 24, 26

HJB Hamilton-Jacobi-Bellman. 3, 10, 11, 27–29, 63

LOC Loss of Control. v, 1, 5, 6, 39, 43
LS Level-Set. xi, xii, 3, 9, 12, 15, 17–19, 27–34, 36, 38–41, 44, 46, 47, 49–54, 57–60

MC Monte-Carlo. xi, 17–19, 28, 30, 32, 46, 47, 57, 58, 60

PDE Partial Differential Equation. 10, 11, 26, 27, 63
PDF Probability Density Function. xi, 21, 24–28, 37, 52, 53, 57
PIMC Path Integral Monte Carlo. v, xii, 27–34, 40, 41, 49–52, 54, 55, 57–60

RCAM Research Civil Aircraft Model. 35

SDE Stochastic Differential Equation. xi, 1, 21–24, 26, 28, 30, 32, 35, 37, 49, 57, 59
SFE Safe Flight Envelope. iii, v, xi, xii, 1, 3–6, 8, 9, 28, 35–39, 43–46, 49, 57, 58, 61, 62

UAV Unmanned Aerial Vehicle. 43
UT Unscented Transform. xii, 51, 52

List of Symbols

α = Angle of Attack [°]
α_g = Gust Angle of Attack [°]
β = Angle of Sideslip [°]
γ = Flight path angle [°]
Γ = Input distribution matrix
ξ = Trajectory
ρ = Air density [kg m^{-3}]
σ = Noise standard deviation
τ = Time-step
Φ = System transition matrix
Ω = Stochastic reachable set

g = Gravitational Acceleration [m s^{-2}]
l = Length
L_g = Turbulence scale length [m]
m = Mass [kg]
p = Co-states
S = Wing area [m^2]
T = Thrust [N]
V = Velocity [m s^{-1}]
q = Pitch rate [°s^{-1}]
Q = System noise covariance

$\mathscr{H}$ = Hamiltonian
$\mathscr{I}$ = Invariance Set
$\mathscr{K}$ = Initial Set
$\mathscr{R}$ = Reachable Set
$\mathscr{V}$ = Viability Set

List of Figures

1

Introduction

A recent study by Boeing [1] into fatal airplane accidents has shown that the majority of aviation fatalities are the result of Loss of Control (LOC) related incidents. Although the LOC occurrence category represented only 4 % of all 2016 accidents, it remains of significant concern since it accounts for 42.9 % of all fatal accidents [2]. LOC can occur when there is an unintended departure from the controlled flight. Often, this happens when an aircraft leaves the normal flight envelope. This may result in a state from which recovery is hard or impossible (e.g. entering a stall or spin). Seventy-four LOC related accidents occurred in the last fifteen years among which were 27 stall accidents, 20 accidents with ice contaminated airfoils, and eight spatial disorientation accidents.

In search of methods to prevent LOC related accidents, flight envelope protection is often considered to be very effective [3]. Flight envelope protection consists of a human-machine interface that prevents the pilot from steering the aircraft into a state which exceeds the structural and aerodynamic operating limits. For such a system to be effective, knowledge of the flight envelope is required. The flight envelope, according to the standard definition, refers to a region of velocity and altitude or load-factor where the aircraft can operate safely. Since such an envelope only contains quasi-stationary states (e.g. coordinated turn or level flight) it is insufficient for the purpose of full envelope protection. A more comprehensive definition which also includes non-stationary states is known as the Safe Flight Envelope (SFE). The SFE refers to the set of all states that can be reached, for which a return trajectory to the steady-state exists. It is formally defined as the intersection between the forwards and backwards reachable set. Computation of the reachable set is facilitated by making a connection between reachable set theory and the SFE. Reachability set theory is widely used in safety analysis where the interest is in finding all states that can be reached, starting from a set of initial conditions over a certain time horizon.

Although the computed SFE is theoretically accurate, some states may be practically unattainable under the influence of (external) disturbances. The presence of a disturbance could, for example, push the aircraft into a state outside the SFE so that it is unable to return to a steady-state. To quantify the effect of (external) disturbances on the SFE, the differential equations representing the model are augmented with a stochastic process, which gives rise to the field of Stochastic Differential Equations (SDEs). The stochastic counterpart of the reachability set theory will then be used to compute the probabilistic SFE. By analyzing the differences between the deterministic and probabilistic SFEs, an assessment of the sensitivity of the SFE to (external) disturbances can be made.

The structure of the report is as follows. Chapter 3 will introduce the basics of flight envelope protection and the definition of the SFE, which is required for the purpose of flight envelope protection. The definition of the SFE, although being rigorous, does not support the actual computation of the SFE. Therefore, Chapter 4 expresses the reachable set theory in an optimal control framework. An introduction to SDEs and the stochastic counterpart of the reachability analysis is provided in Chapter 5. The concept of stochastic reachability analysis will be applied to a simplified aircraft model in Chapter 6. This will also introduce the stochastic counterpart of the SFE, called the probabilistic SFE. In an attempt to explore the (practical) limits of the reachability analysis, a quadrotor model that features a high number of states and a parametrized aerodynamic model is considered in Chapter 7. Some techniques to reduce the computational load are explored in Chapter 8. Finally, Chapter 9 holds the conclusions and recommendations.

2

Thesis Objective

Most of the current research into SFE prediction is geared towards the practical application of the reachable set computation. Since the Level-Set (LS) method suffers from the so-called 'Curse of Dimensionality' (computational effort increases exponentially with the dimension of the problem), it is impractical to apply the LS method to problems with a many Degrees-Of-Freedom (DOF), or for problems that required (near) real-time computation.

For example, in De Weerdt et al. [4] the LS method is combined with interval analysis to give a guaranteed upper and lower bound on the reachable set using lower computational complexity. Furthermore, Stapel et al. [5] evaluates improvements in the numerical methods for solving the Hamilton-Jacobi-Bellman (HJB) equations. These improvements include the usage of non-iterative schemes known as fast marching methods. Other methods try to evaluate the LS method only close to the interface, thus reducing the computational load. A different approach towards real-time implementation of envelope protection is the application of a database driven design. In Zhang et al. [6] the reachable set is computed offline and stored in a database. Finally, some methods try to circumvent the usage of the LS method all together. For example the work of Helsen et al. [7] tries to solve the reachability problem using a novel distance-fields-over-grids method.

Although a lot of research is focused on improving the computational efficiency of the SFE prediction, almost all methods implement a model without the presence of disturbances or uncertainties. Only a very limited number of cases with disturbances or uncertainties are treated in literature. For example, in the work of Lombaerts et al. [8, 9] uncertainties on the aerodynamic coefficients are modeled using a two-player differential game for varying uncertainty bands. In the two-player differential game, the uncertainty (second player) will try to minimize the value function, while the other player tries to maximize it. Although this gives an indication of the robustness of the reachable set, it will only compute the minimal reachable set. Other scenarios, where the disturbance drives the system further away from the initial set are not considered, thus providing an incomplete view of the system. Furthermore, the disturbances are only present on the aerodynamic coefficients. External disturbances such as turbulence are not considered.

2.1. Research Question

Based on previous research, a more sophisticated method of including disturbances in the reachability analysis is required. This research aims to extend the reachability analysis with disturbances acting on the dynamic model in order to shed light on the sensitivity of the SFE under the influence of these disturbances. The main goal has been formulated as the following research question.

"The objective of the proposed research is to investigate the effect of (stochastic) disturbances on safe flight envelope models."

The main research question gives rise to a number of sub-questions which, when answered, will together answer the main research question.

- How can the deterministic reachability analysis be extended to include stochastic processes?
- How to validate the results from the stochastic envelope prediction?

- How does the SFE change under the influence of (stochastic) disturbances?
- Is it possible to provide an approximate answer for real-time envelope prediction?

The main addition of this proposed research in comparison to the current state-of-the-art is the addition of (stochastic) disturbances in the reachability analysis which will give a more realistic insight into the true reachable set.

3

Flight Envelope Protection

Envelope protection has been identified as a key element to prevent Loss of Control (LOC) related accidents. It prevents the pilot from steering the aircraft into a state from which it is unable to recover. For the purpose of flight envelope protection, knowledge about the flight envelope is required. Under nominal conditions the aircraft should always operate within this envelope. Some safety systems exists that warn the pilot when a flight envelope excursion is imminent, however these systems do not provide full envelope protection. The standard definition of the flight envelope does not include non-stationary states and is therefore unsuitable for the purpose of full flight envelope protection. Other deficiencies include a lack of adaptivity (e.g. in case of structural failure) and absence of information on the safety levels. The shortcomings to the standard definition of the flight envelope are discussed in Section 3.1. A better suited definition, called the Safe Flight Envelope (SFE) is introduced in Section 3.2. Finally, Section 3.3 will express the SFE in the framework of reachable set theory.

3.1. Flight Envelope

The flight envelope commonly refers to a region that indicates the capabilities of an aircraft in terms of velocity and altitude or load-factor. Figure 3.1 shows an example of such an envelope. Under nominal conditions the aircraft should always be operated within these limits in order to guarantee safety of the vehicle. Leaving the flight envelope could, for example result in a stall condition or structural failure of the airframe. This flight envelope is usually determined during the development of an aircraft by means of simulation or flight testing.

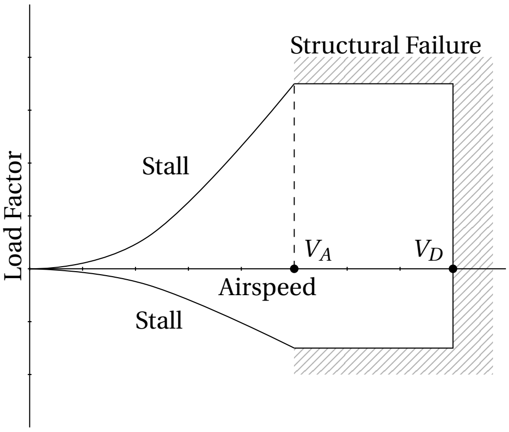

Figure 3.1: Example of a flight envelope.

When a flight envelope excursion happens, the aircraft can enter a state from which it is hard or impossible to recover, compromising safety. For example, when entering a stall in close proximity to the ground, upset recovery is almost impossible. Flight envelope protection restricts an aircraft from leaving the flight envelope which can help to prevent LOC, however it is not sufficient. Since the standard flight envelope only takes into account quasi-stationary states (e.g. coordinated turns or level flight) it does not provide a solution during dynamic maneuvers. In order to achieve full flight envelope protection the concept of a flight envelope has to be extended to include dynamic maneuvers.

3.2. Safe Flight Envelope

To overcome the shortcomings of the flight envelope, a more complete representation of the flight envelope is required. The concept of a SFE is introduced for the purpose of preventing LOC. It is formally stated according to Definition 3.1.

Definition 3.1. (Safe Flight Envelope): The part of the state-space for which safe operation of the aircraft and safety of its cargo can be guaranteed, and externally posed constraints will not be violated [10].

The externally posed constraints are composed of the following envelope's:

- Dynamic Envelope: Envelope that is constrained by the dynamic behavior of the aircraft, due to its aerodynamics and kinematics.
- Structural and Comfort envelope: Constraints posed by the airframe, pilot, passengers and cargo. These constraints are usually defined through maximum accelerations and loads.
- Environmental Envelope: Constraints due to the environment in which the aircraft operates.

The constraints posed in the last two envelopes are well-known and can be quantified quite easily, however this is not true for the dynamic envelope. The main focus of this study will be on the calculation of the dynamic envelope, where the complete envelope is given by the intersection of all three flight envelopes.

3.3. Reachable Set Theory

The formal definition of the SFE provides a high level description, however this definition is not sufficient to facilitate computation of the SFE. For this purpose the SFE will be expressed in the framework of reachable set theory. Reachability set theory is widely used in safety analysis where the interest is in finding all states that can be reached, starting from a set of initial conditions, over a certain time horizon. In this section the connection with the reachable set theory will be made [11].
Consider a continuous-time dynamic system

$$\dot{\mathbf{x}}(t) = f(t, \mathbf{x}(t), \mathbf{u}(t)) \quad \mathbf{x} \in \mathbb{R}^n, \mathbf{u}(\cdot) \in \mathscr{U} \subseteq \mathbb{R}^m \tag{3.1}$$

where $f(t, \mathbf{x}, \mathbf{u})$ is assumed to be Lipschitz continuous (i.e. $\dot{f}(t, \mathbf{x}, \mathbf{u}) < L$). With a time horizon $T \geq 0$ and a set of initial states $\mathscr{K} \subseteq \mathbb{R}^n$. A solution to the dynamic system is a trajectory through the state-space. The time horizon and the set of admissible control inputs ($\mathscr{U}$) provide a bound on the states that can be reached. A larger time horizon leads to a larger reachable set. Furthermore, when larger control inputs are permitted, a larger part of the state-space can be reached in the same amount of time.

Based on the definition of the dynamical system given by Equation (3.1), three different types of reachable sets can be defined.

Definition 3.2. (Invariance): $\mathscr{I}(t, \mathscr{K})$: Set of all states $\mathbf{x}(\cdot)$ for which there exists all the inputs $\mathbf{u}(\cdot) \in U_{[0,T]}$ such that $\mathbf{x}(t) \in \mathscr{K}$ for all $t \in [0, T]$.

Definition 3.3. (Viability): $\mathscr{V}(t, \mathscr{K})$: Set of all states $\mathbf{x}(\cdot)$ for which there exists at least one input $\mathbf{u}(\cdot) \in U_{[0,T]}$ such that $\mathbf{x}(t) \in \mathscr{K}$ for all $t \in [0, T]$.

Definition 3.4. (Reachability): $\mathscr{R}(t, \mathscr{K})$: Set of all the states $\mathbf{x}(\cdot)$ for which there exists at least one input $\mathbf{u}(\cdot) \in U_{[0,T]}$ and $t \in [0, T]$ such that $\mathbf{x}(t) \in \mathscr{K}$.

Figure 3.2 provides an illustration of the different reachable kernels. The *invariance* set is the smallest set. It contains all states for which all possible control inputs keep the trajectory within $\mathscr{K}$ over the entire time horizon. Set $\mathscr{V}$ is the set of all states for which there exists at least one control sequence which keeps the trajectory within $\mathscr{K}$ over the entire time horizon. This set, together with the *invariance* set are completely contained inside the initial set $\mathscr{K}$. Finally, $\mathscr{R}$ is the set of all states for which there exists at least one control sequence that can steer the trajectory into set $\mathscr{K}$ within the time horizon. The definition of the *invariance*, *viability* and *reachability* kernel are used extensively in safety analysis [10–12].

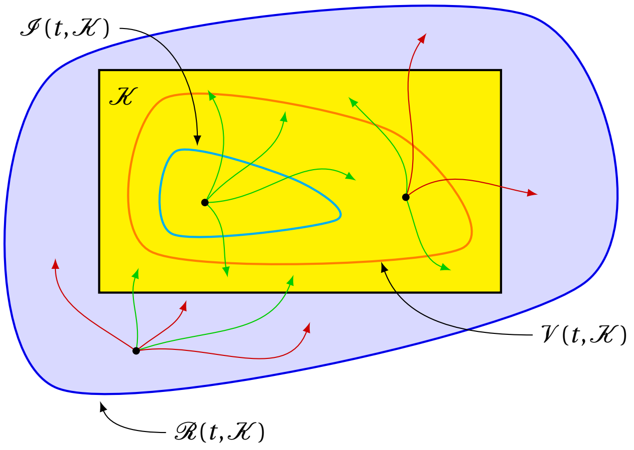

Figure 3.2: Illustration of the different reachable kernels. Viability (orange), Invariance (cyan) and (backwards) reachability Set (Blue), Green: Successful trajectories; Red: Unsuccessful trajectories [13].

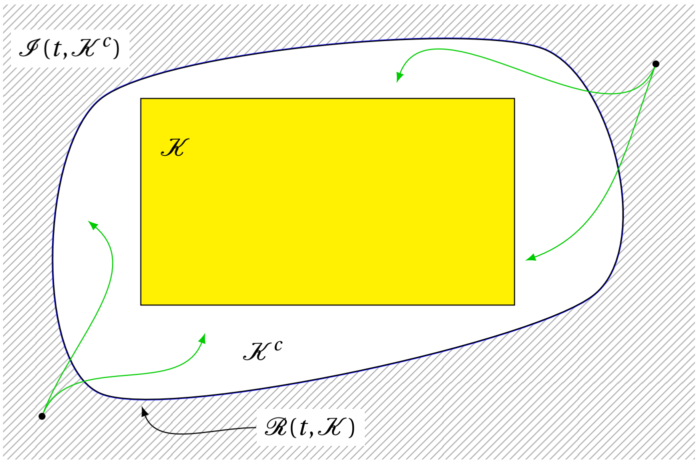

Figure 3.3: Duality principle shows that the reachability kernel is equal to the complement of all states for which no trajectory to set K exists.

A connection between the different reachable kernels can be made. The principle of duality

$$\mathscr{R}(t,K) = \left(\mathscr{I}(t,\mathscr{K})^c\right)^c \tag{3.2}$$

relates the *reachability* to the *invariance* set. Figure 3.3 illustrates the duality principle. It can be seen that all possible control sequences will result in a trajectory that stays within $\mathscr{K}^c$ over the entire time horizon. Alternatively, one can say that the complement of the reachability kernel consists of all the initial conditions for which there does not exist a single control sequence that can steer the trajectory into set $\mathscr{K}$ within the time horizon.

The dynamic system (3.1) can be solved forwards in time, but also backwards in time (terminal value problem instead of an initial value problem) which gives rise to two further types of reachability; *Forwards reachability* and *backwards reachability*. These reachable sets are defined according to Definition 3.5 and 3.6.

Definition 3.5. (Forwards Reachable Set): $\mathscr{R}_f(t,\mathscr{K})$: Set of all the states $\mathbf{x}(\cdot)$ for which there exists a control input $\mathbf{u}(\cdot) \in U_{0,T}$ such that this set can be approached starting from the initial set $\mathscr{K}$ for $t \in [0,T]$.

Definition 3.6. (Backwards Reachable Set): $\mathscr{R}_b(t,\mathscr{K})$: Set of all the states $\mathbf{x}(\cdot)$ for which a control input $\mathbf{u}(\cdot) \in U_{[0,T]}$ exists at time $t \in [0,T]$, such that at least one state in the initial set $\mathscr{K}$ can be reached.

Figure 3.4 shows both the forwards and backwards reachable sets. For any state in the forwards reachable set, a trajectory towards this state exists, starting from somewhere in the initial set. The backwards reachable

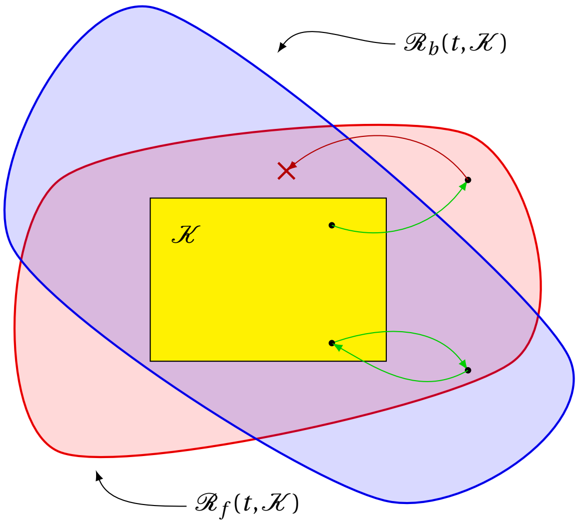

Figure 3.4: SFE defined as the intersection between the forwards (red) and backwards (blue) reachable set.

set consists of all the initial conditions for which a solution to the dynamic system exists that reaches set $\mathscr{K}$. The SFE is defined as the intersection between the two. For any point inside the SFE, there exists a trajectory from the initial set $\mathscr{K}$ to a state and back. If the state is in the forwards reachable set, but not in the backwards reachable set, it is possible to maneuverer to the state inside the forwards reachable set, but it is unable to return to the initial set within the time horizon.

4

Reachability Analysis

In the previous chapter a link between the SFE and the reachable set theory was established. Although this provides a thorough definition of the reachable sets, it does not support the actual computation of these sets. By expressing the reachable sets in an optimal control framework, computation of the flight envelopes is facilitated. With the exception of very simple cases, it is not possible to provide an analytical solution to this optimal control framework. Level-Set (LS) methods are introduced as a way to numerically approximate the solution. Section 4.1 will introduce the theory behind the optimal control framework. An intuitive way to describe the optimal control framework, called the Lagrangian method, is introduced in Section 4.2. Section 4.3 introduces the LS method which is used to numerically approximate the reachable set. Finally, verification and validation of these methods are presented in Section 4.4 and 4.5 respectively.

4.1. Optimal Control Formulation

The *viability* and *invariant* sets as introduced in Chapter 3 can be linked to SUPMIN and INFMIN optimal control problems. The relation between the reachable sets and optimal control problems has been stated and proven by Lygeros [11] and will be briefly summarized here.

Consider the continuous-time dynamic system

$$\dot{\mathbf{x}} = f(\mathbf{x}, t, \mathbf{a}, \mathbf{b}) \quad \mathbf{x} \in \mathbb{R}^n, \mathbf{a}(\cdot) \in \mathscr{A} \subseteq \mathbb{R}^{m_a}, \mathbf{b}(\cdot) \in \mathscr{B} \subseteq \mathbb{R}^{m_b} \tag{4.1}$$

where **a** and **b** are inputs to the system of differential equations. The inputs are considered to be two actors – pursuer and evader – with opposite goals. The pursuer tries to maximize the objective function while the evader tries to minimize it. This kind of problem is often referred to as a differential game. Solutions of the dynamic system (4.1) are trajectories of the system and will be denoted by [14]

$$\xi_f(\cdot; \mathbf{x}, t, \mathbf{a}(\cdot), \mathbf{b}(\cdot)) \tag{4.2}$$

which satisfies the initial condition $\xi_f(0; \mathbf{x}, t, \mathbf{a}(\cdot), \mathbf{b}(\cdot)) = x_0$ and is differentiable almost everywhere

$$\frac{\mathrm{d}}{\mathrm{d}\tau}\xi_f(\tau; \mathbf{x}, t, \mathbf{a}(\cdot), \mathbf{b}(\cdot)) = f\left(\xi_f(\tau; \mathbf{x}, t, \mathbf{a}(\cdot), \mathbf{b}(\cdot)), \tau, \mathbf{a}(\tau), \mathbf{b}(\tau)\right) \tag{4.3}$$

Note the semi-colon to distinguish between the argument τ of ξ_f and the trajectory parameters.

The objective of the differential game is defined by a value function

$$V_1(\mathbf{x}, t) = \sup_{\mathbf{a}(\cdot)\in\mathscr{A}} \inf_{\mathbf{b}(\cdot)\in\mathscr{B}} \min_{\tau\in[0,T]} l\left(\xi(\tau; \mathbf{x}, t, \mathbf{a}(\cdot), \mathbf{b}(\cdot))\right) \tag{4.4}$$

$$V_2(\mathbf{x}, t) = \inf_{\mathbf{a}(\cdot)\in\mathscr{A}} \sup_{\mathbf{b}(\cdot)\in\mathscr{B}} \min_{\tau\in[0,T]} l\left(\xi(\tau; \mathbf{x}, t, \mathbf{a}(\cdot), \mathbf{b}(\cdot))\right) \tag{4.5}$$

V is called the value function or the cost-to-go. It represents the value of the solution to the optimal control problem at time t and initial state $\mathbf{x}$. The input **a** tries to maximize the minimum value attained by the function l along the state trajectory over the time horizon $[0, T]$ while the evader **b** tries to minimize this.

The *viability* and *invariant* sets can now be linked to the value function by Equation (4.6,4.7).

$$\mathcal{V}(t,K)=\left\{\mathbf{x}\in\mathbb{R}^n | V_1(\mathbf{x},t)>0\right\} \tag{4.6}$$

$$\mathcal{I}(t,K)=\left\{\mathbf{x}\in\mathbb{R}^n | V_2(\mathbf{x},t))\geq 0\right\} \tag{4.7}$$

With $V(x,T)=l(x)$. The *viability* set is defined to be equal to all states where the V_1 is larger than zero. Furthermore, the *invariant* set is defined to be equal to all states where the V_2 is larger or equal than zero.

Bellman's Optimality Principle

'An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision'. [15]

Figure 4.1 illustrates the dynamic programming principle where the objective is to maximize the value attained over the trajectory. The value at each line represent the cost of moving from one state to another. An intuitive method would be to start at the beginning and calculate the value for all possible trajectories, then picking the maximum.

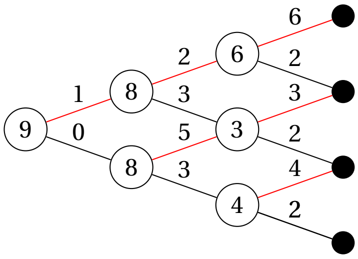

Figure 4.1: Dynamic programming principle showing the optimal path in red when moving backwards from right to left.

A more efficient way is to divide the problem into smaller sub problems, then solving them backwards in time. Starting at the right hand side, the path is chosen which yields the highest value. This optimal path is then marked in red and the cumulative value is stored in the node. Then, for the next step, the optimal path is one which maximizes the value attained over the path plus the value stored in the next node. By combining the optimal solutions to non-overlapping sub-problems, a solution can be found.

By solving the problem in this way, the optimal policy is found for every initial condition and at each node. Therefore, the backwards scheme yields the optimal control policy in the form of a state feedback law.

An intuitive derivation can be made if it is assumed that the value function is differentiable. The value function can be reformulated as a Hamilton-Jacobi-Bellman (HJB) Partial Differential Equation (PDE) using Bellman's optimality Principle and a Taylor expansion of the value function [10]

$$V_2(\mathbf{x},t)=\inf_{\mathbf{a}(\cdot)\in\mathscr{A}}\sup_{\mathbf{b}(\cdot)\in\mathscr{B}} V\left(\xi_f(t+\Delta t;x,t,\mathbf{a}(\cdot),\mathbf{b}(\cdot)),t+\Delta t\right) \tag{4.8}$$

First order Taylor expansion yields

$$V\left(\xi_f(t+\Delta t;x,t,\mathbf{a}(\cdot),\mathbf{b}(\cdot)),t+\Delta t\right)\approx V(x,t)+V_t(x,t)\Delta t+V_x(x,t)\Delta x \tag{4.9}$$

Re-arranging the terms of Equation (4.8) by the principle of optimality and dividing by Δt yields

$$\inf_{\mathbf{a}(\cdot)\in\mathscr{A}}\sup_{\mathbf{b}(\cdot)\in\mathscr{B}}\frac{V\left(\xi_f(t+\Delta t;x,t,\mathbf{a}(\cdot),\mathbf{b}(\cdot)),t+\Delta t\right)-V(x,t)}{\Delta t}=0 \tag{4.10}$$

Equation (4.10) can be thought of as one step/sub-problem of the dynamic programming principle as explained in Figure 4.1.

Substituting the Taylor expansion into Equation (4.10) yields

$$\inf_{\mathbf{a}(\cdot)\in\mathscr{A}}\sup_{\mathbf{b}(\cdot)\in\mathscr{B}}\frac{V_t(x,t)\Delta t+V_x(x,t)\cdot\Delta x}{\Delta t}=0 \tag{4.11}$$

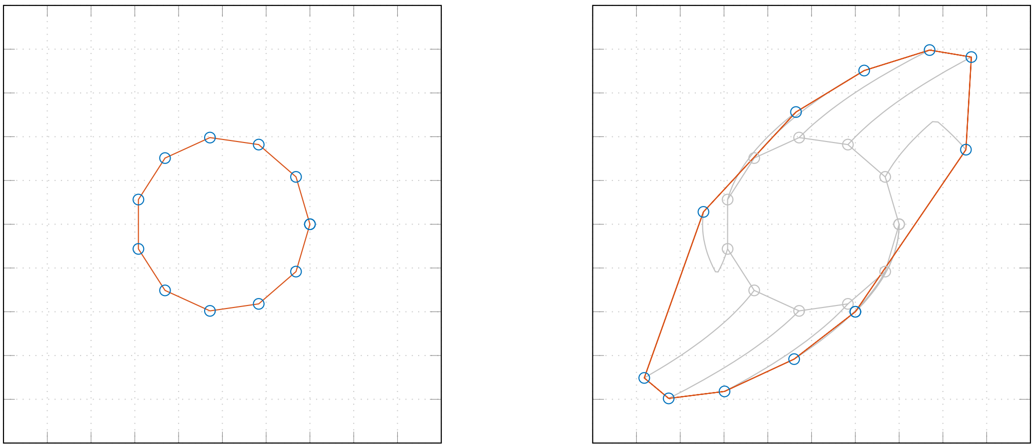

(a) Initial set (red) and discretization points (blue).

(b) Propagation of the discretization points and reconnection them according to the initial connectivity yields the reachable set.

Figure 4.2: Reachable set computation using the Lagrangian approach.

For the limit of $\Delta t \to 0$

$$\frac{\partial V_1}{\partial t}(\mathbf{x}, t) + \min\left\{0, \sup_{\mathbf{a}(\cdot)\in\mathscr{A}} \inf_{\mathbf{b}(\cdot)\in\mathscr{B}} \frac{\partial V_1}{\partial \mathbf{x}}(\mathbf{x}, t) f(\mathbf{x}, t, \mathbf{a}, \mathbf{b})\right\} = 0 \tag{4.12}$$

$$\frac{\partial V_2}{\partial t}(\mathbf{x}, t) + \min\left\{0, \inf_{\mathbf{a}(\cdot)\in\mathscr{A}} \sup_{\mathbf{b}(\cdot)\in\mathscr{B}} \frac{\partial V_2}{\partial \mathbf{x}}(\mathbf{x}, t) f(\mathbf{x}, t, \mathbf{a}, \mathbf{b})\right\} = 0 \tag{4.13}$$

The result is a PDE which includes an optimization over the inputs **a** and **b**. Section 4.3.3 will elaborated on the optimization procedure. The minimization term with zero guarantees that the reachable set cannot shrink as time marches forwards. A state that already entered the target set is restricted from leaving before the final time by 'freezing' the evolution of the trajectory [6].

Unfortunately, the assumption made in Equation (4.8) that the value function is differentiable is often not valid, making it unable to interpret it as a solution in the 'classical' sense. Due to the switching of the optimal control and a discontinuity present in the right hand side of the HJB PDE the value function may not remain continuous. In order to obtain a solution to the HJB PDE a 'weak' formulation of solutions to these equations is necessary. Viscosity solutions form a general theory of 'weak' (i.e. non-differentiable) solutions. The interested reader is referred to Bressan [16] for more information on viscosity solutions. Note that viscosity solutions are not the same as vanishing viscosity solutions which are the limit of an additional viscosity term for $\epsilon \to 0$.

4.2. Lagrangian Method

The most intuitive way to explain the reachability analysis is using the Lagrangian method. This method works by propagating the dynamical system starting at a point on the boundary of the initial set. By keeping track of the connectivity between the different points, an approximation of the reachable set can be constructed. Consider the example of a simple linear system where the initial set is a circle with unit radius centered at the origin. This set can be approximated by placing a number of points, spaced evenly, on the boundary of the initial set. By storing the order of these points (connectivity) and interpolating in between them, the complete interface is reconstructed as shown in Figure 4.2a. Each of the points can be thought of as the initial condition of the dynamic system. The solution of the dynamic system under an optimal control law that maximizes the value function is a trajectory through the state-space. Figure 4.2b shows the reachable set that is constructed by connecting the endpoints of the individual trajectories under the initial connectivity.

Although being very intuitive, the Lagrangian method suffers from a number of important drawbacks. Keeping track of the connectivity (which points connect to which) might seem straightforward at first, but can quickly become challenging. When two trajectories cross path, the connectivity needs to change accordingly. Although this process is relatively simple in a two dimensional system, it is not trivial in higher dimensional systems.

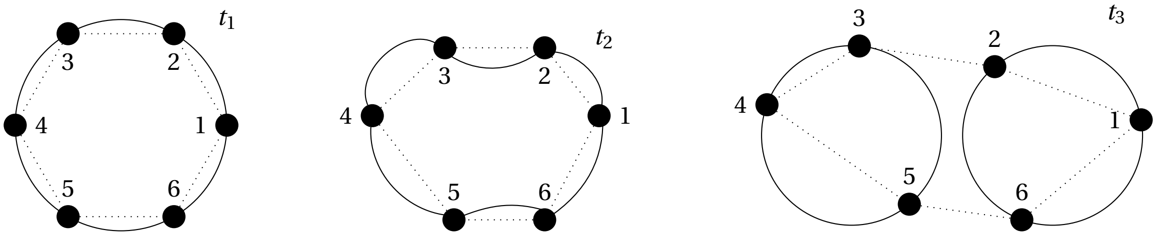

Figure 4.3: Effect of incorrect connectivity on the reachable set (polygon defined by the dotted line) [17].

A different problem arises when splitting or merging two or more separate sets. When not properly handling the merging and removal of discretization points, errors in the representation of the reachable set can arise. An example of a problem that can occur when splitting two sets is illustrated in Figure 4.3. The true reachable set (solid line) is approximate by connecting the discrete grid points (dotted line). After three time steps, the true reachable set has split into two sets. Since the connectivity has not been changed, the true reachable set is poorly approximated by the discrete grid points. Finally, although starting with an uniform spacing between the points in Figure 4.2a, after propagating them forwards in time and constructing the reachable set, it can be seen from Figure 4.2b that the spacing has changed. This effect could lead to discretization errors and loss of features. In some cases it may be needed to add new points or reinitialize the current points to adequately represent the interface, which is again very challenging in higher dimensions.

4.3. Level-Set Method

A more powerful and flexible alternative to the Lagrangian approach can be found in the so-called LS methods. LS methods, developed by Osher and Fedkiw [18] are a class of methods that add dynamics to implicit functions. These methods use a so called Eulerian approach which means that the computation is performed on a fixed grid without having to parametrize the interface first. This method has successfully been applied to (among others) reachability analysis of continuous-time dynamic systems. The implicit functions require $n+1$ dimensions to represent an n-dimensional set. While at first this may seem like waste of resources, it will quickly become clear that this method has significant advantages. Possibly the biggest benefit of the LS method, is that it is very straightforward to go from two to three or even higher dimensional systems. Other benefits include that the connectivity does not need to be determined for the discretization and that the numerical error is proportional to the mesh size, instead of the explicit discretization of the interface as is the case for the Lagrangian approach.

4.3.1. Implicit Functions

In the Lagrangian approach, the boundary is explicitly described by points that belong to the interface. Alternatively, an implicit interface representation defines the interface as the isocontour of some function. For example, the interface of the implicit function $\phi(x) = x^2 - 1$ is the set of points where $\phi(x) = 0$ (i.e. $\partial\Omega = \{-1, 1\}$). When representing more complex boundaries, the implicit function is often described by a signed distance function on a discretized grid. Each grid point then holds the value of the signed distance function (i.e. the distance to the boundary). In contrary to explicit contours as used in the Lagrangian method, the discretization of the implicit function is independent of the interface that is being described. Figure 4.4 shows an example of an implicit function. The grid on the x-y plane represents the discretized state-space. By interpolating the implicit function (ϕ) on the state-space, an implicit surface is formed. The interface is now described by the intersection of a level plane with the implicit function. For convenience this level plane is often defined at $\phi(x) = 0$ also known as the zero level-set. All the values belonging to the interior of the implicit contour are negative, while a positive value denotes the set outside of the implicit contour. Now that the implicit functions are described, the power of this method becomes clear. For example, operations such as merging between two sets, which was challenging for an explicit boundary, becomes very trivial in the case of an implicit function. It simply involves taking the minimum value of the two sets at each grid point.

4.3.2. Convection

LS methods were previously described as a way to add dynamics to implicit functions. Now that the implicit functions have been introduced, it is time to add the dynamics. The LS method describes how the value

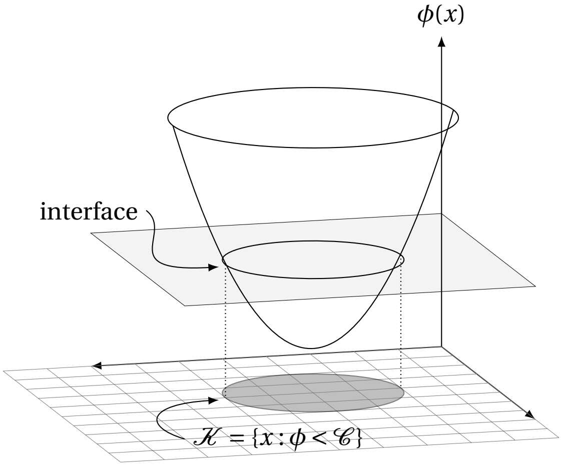

Figure 4.4: Interface of a set as represented by an implicit function.

function changes in time, which in term propagates the interface. In this case the boundary will be evolved in time by an external flow-field. This flow-field is the result of the dynamics of the system and the optimal control policy which will be elaborated in Section 4.3.3. For now it is assumed to be known and constant.

The motion of the implicit interface through a flow-field is described by

$$\phi_t + \mathbf{V} \cdot \nabla\phi = 0 \tag{4.14}$$

which is known as the convection equation. More specifically it describes that the time rate of change of ϕ is determined by the dot product of the velocity field $\mathbf{V}(\mathbf{x}, t)$ with the spatial gradient $\nabla\phi$. The numerical method of solving Equation (4.14) is to start with an initial implicit function $\phi_0(\mathbf{x})$. Then, based on the external velocity field, the change of the implicit function is calculated for each discrete point on the grid.

For a 1-dimensional problem Equation (4.14) reduces to

$$\Delta\phi = -V \cdot \phi_x \cdot \Delta t \tag{4.15}$$

The motion of an implicit boundary through an external flow-field is illustrated in Figure 4.5. The initial implicit function ϕ^n is shown in Figure 4.5a. It is a linear interpolation between the discrete grid points x^n. The boundary is located where the implicit function crosses the zero-level and is indicated as $\partial\Omega^n$. The implicit function will be propagated under the influence of an external flow-field. This flow-field is shown in Figure 4.5b. The value of the implicit function at each grid node will change according to the convection equation. It means that the time rate of change of the implicit function is equal to the velocity times the gradient of the implicit function times the time-step. This process is illustrated in Figure 4.5c. It can be seen that the value of the implicit function at time $n+1$ is decreased. The result is a new implicit function ϕ^{n+1} as shown in Figure 4.5d. By moving the implicit function down, the zero-level shifts to the right. The new location of interface ($\partial\Omega^{n+1}$) therefore also moves to the right in accordance with the flow-field.

4.3.3. Optimal Control

In the previous section, the flow-field used to propagate the implicit interface was assumed to be known and constant. In the case of reachability analysis this flow-field is the result of the underlying system dynamics, where the inputs are chosen to maximize the value function. This maximization is controlled by an optimal control policy which selects the best control sequence for each point in the state-space separately. Inserting the optimal input in the dynamical system yields the flow-field through which the boundary of the reachable set is evolved.

The second term in Equation (4.12) is often referred to as the Hamiltonian which plays a fundamental role in optimal control theory

$$\mathcal{H}(x, p) = \max_{a\in\mathcal{A}} \min_{b\in\mathcal{B}} p^T f(x, a, b) \tag{4.16}$$

with $p = \nabla\phi(x, t)$ the co-states, or Lagrange multipliers of the system. The nabla operator (∇) represents the gradient of the value function. Furthermore, the infimum and supremum are equal to concept the minimum and maximum since the admissible control set is bounded and compact.

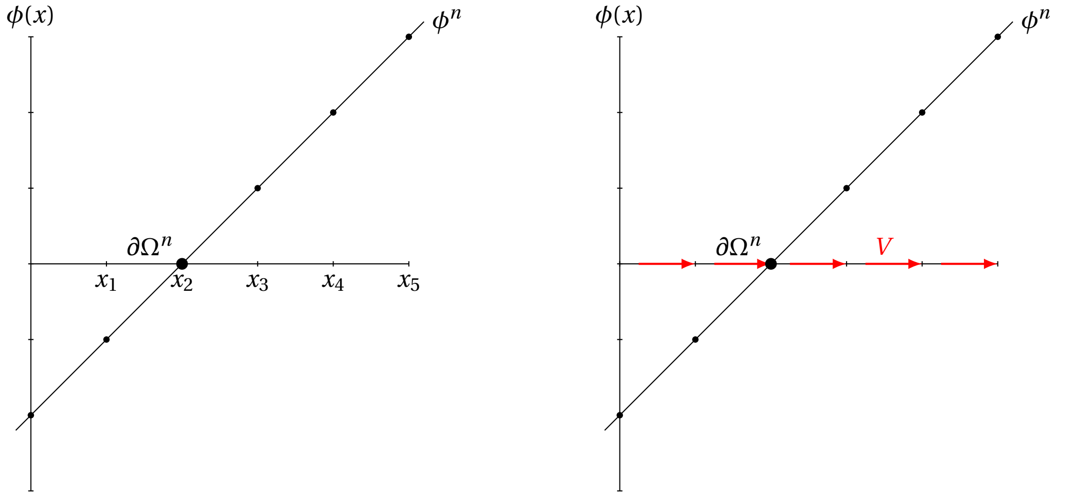

(a) Initial implicit function with boundary $\partial\Omega_n$.

(b) External flow-field.

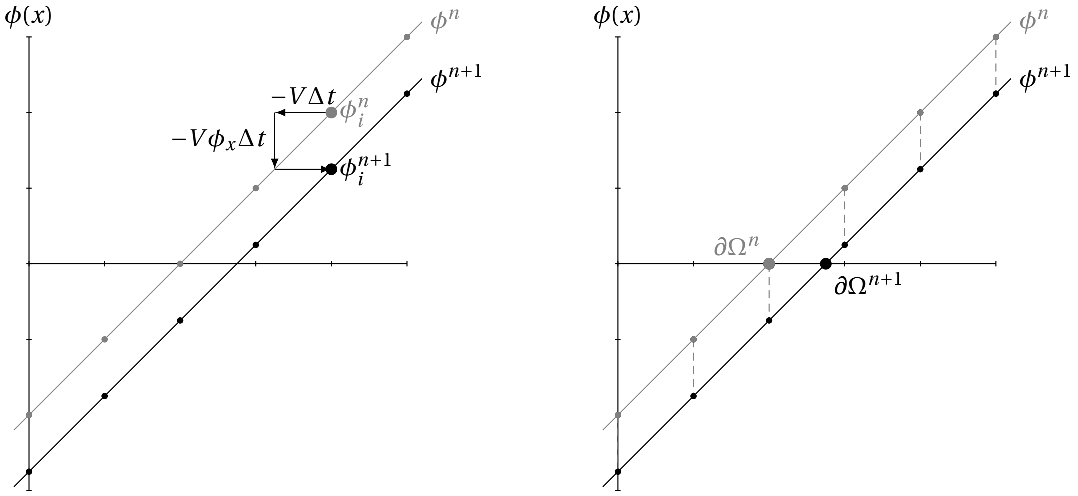

(c) Change in implicit function due to convection.

(d) Boundary at next time step.

Figure 4.5: Propagation of the implicit function by means of convection through an external flow-field [17].

Substituting Equation (4.16) into Equation (4.12) yields

$$\frac{\partial \phi(x,t)}{\partial t} + \min\left[0, \mathscr{H}(x, \nabla\phi(x,t))\right] = 0 \tag{4.17}$$

where the value function V is now represented by the value of the implicit function ϕ.

The method of Lagrange multiplier is used to solve an optimization problem under equality constraints. Essentially, the aim is to optimize the value function under equality constraints of the dynamics of the system. For each step the maximum rate of change of the value function is found when the gradient of the value function is tangential to the flow-vector of the dynamics. In mathematical terms this is equal to the dot-product of the state equations with the gradient of the value function. For each node in the grid, a control action is selected which maximizes (or minimizes) this product. Figure 4.6 shows an example for one grid node. $\vec{p}$ represents the gradient of the value function ($[p_1 p_2]^T$). The blue and red arrows show the motion of the dynamical system for the minimum and maximum input respectively. The dot product can be seen as the scalar projection of the blue and red vectors onto the gradient of the value function. In this case the minimum control input will give the largest value even though the magnitude of the red vector is larger.

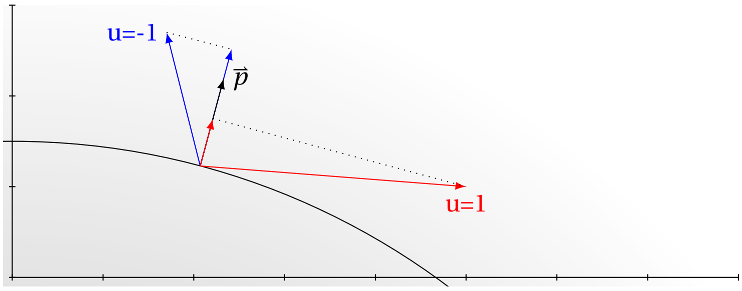

Figure 4.6: Optimal control procedure based on the gradient of the value function.

After the optimization is performed the evolution of the implicit boundary is performed in a similar way as explained in Section 4.3.2. Figure 4.7 shows the evolution of a simple linear system over a 1 s time horizon. The initial set is a circle with unit radius centered at the origin. Again, the blue and red arrows represent the flow-field of the dynamical system for the minimum and maximum control selected using the previously explained procedure.

4.4. Verification

To ensure that the LS method is implemented correctly, a verification will be performed. The solution of the LS method will be verified with an example for which an analytical solution is available. The example consists of a double integrator system. This system can be thought of as a cart on rails, with an input pushing it left or right. An illustration of the problem is shown in Figure 4.8.

The dynamics of the system are given by Equation (4.18, 4.19)

$$\dot{\boldsymbol{x}}(t) = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \boldsymbol{x}(t) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(t) \tag{4.18}$$

$$\boldsymbol{y}(t) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \boldsymbol{x}(t) + \begin{bmatrix} 0 \\ 0 \end{bmatrix} u(t) \tag{4.19}$$

where the input is bounded $u \in [-1; 1]$.

For the case where the initial condition is a point in the origin, an analytical solution can be found. This solution is given by [19]

$$\phi(x) = \begin{cases} x_2 + \sqrt{4x_1 + 2x_2^2} & \text{if } x_1 > \frac{1}{2}x_2|x_2| \\ -x_2 + \sqrt{-4x_1 + 2x_2^2} & \text{if } x_1 < \frac{1}{2}x_2|x_2| \\ |x_2| & \text{if } x_1 = -\frac{1}{2}x_2|x_2| \end{cases}$$

The numerical solution is found by solving Equation (4.17) using the LS Toolbox [20]. For the double

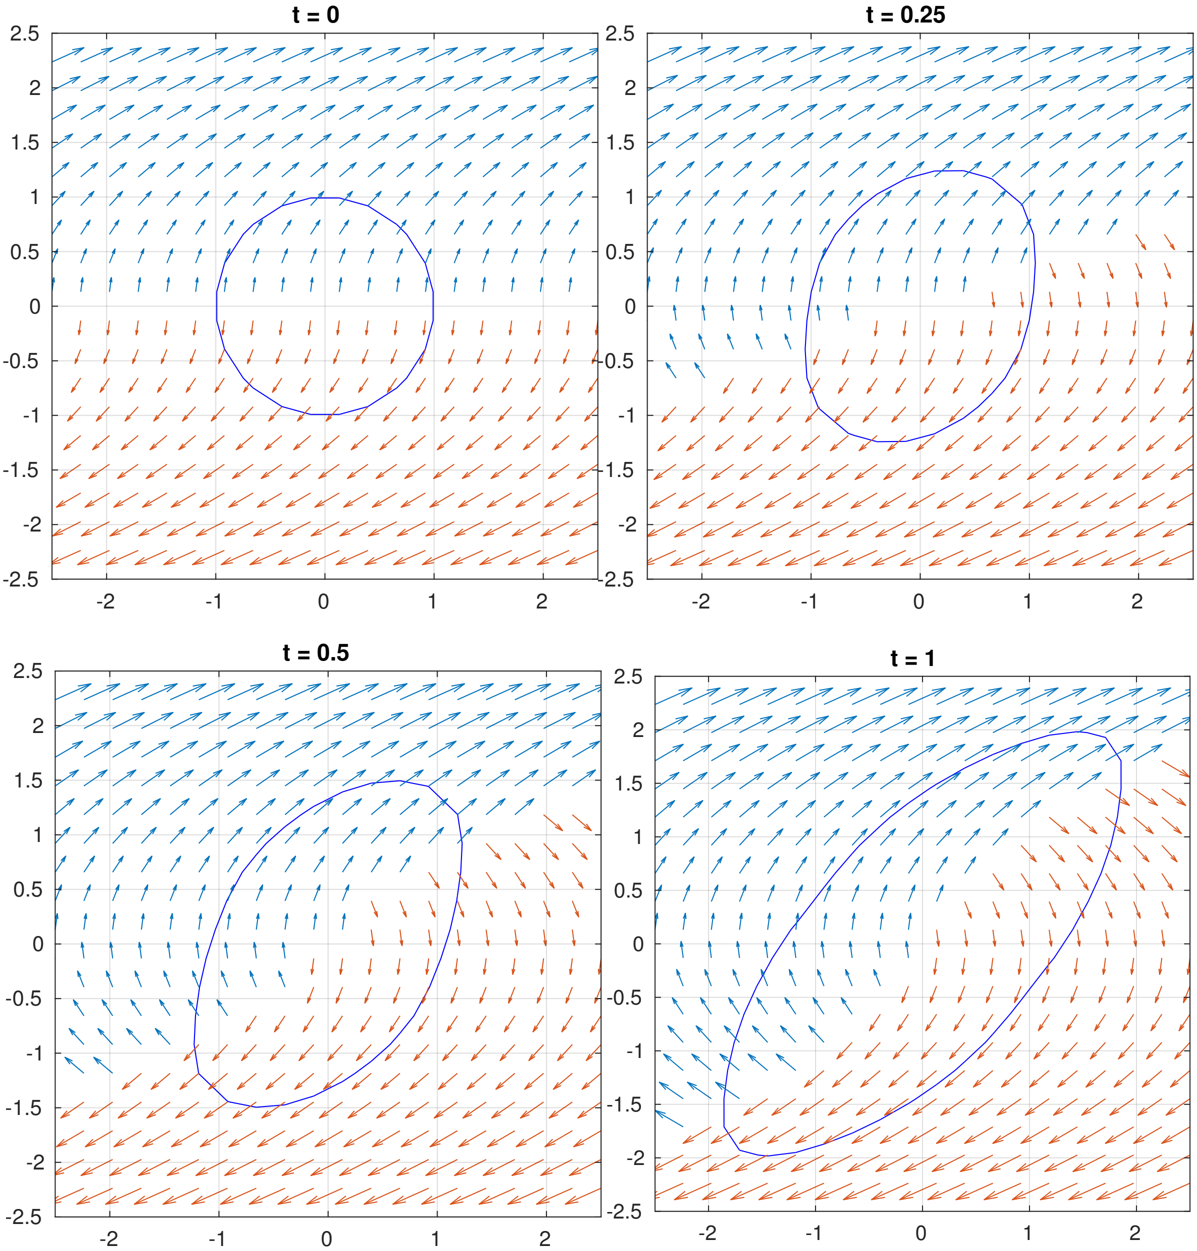

t = 0

t = 0.25

t = 0.5

t = 1

Figure 4.7: Evolution of a simple linear system; arrows show the flow-field of the dynamical system for both the maximum (blue) and minimum (red) optimal control inputs.

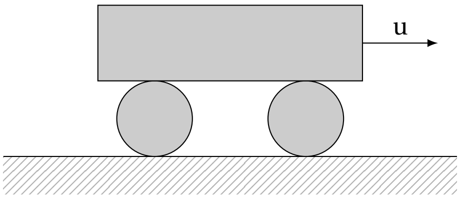

Figure 4.8: Schematic representation of the double integrator system.

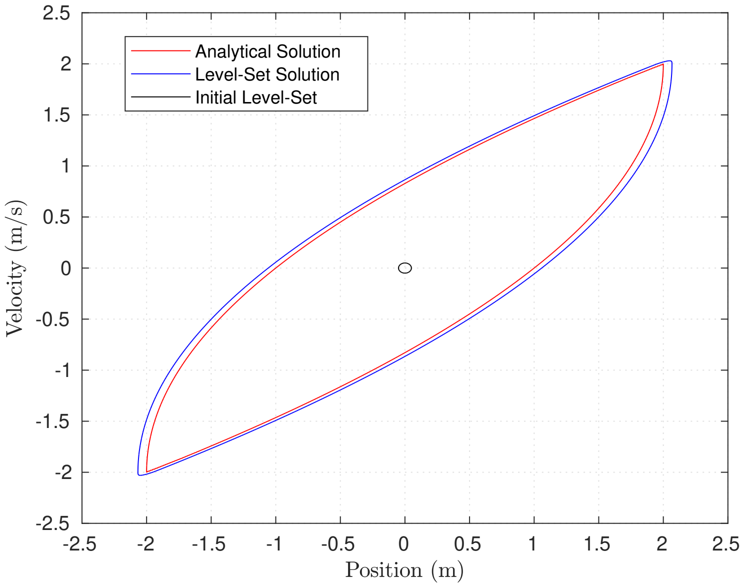

Figure 4.9: Verification of the double integrator forwards reachable set for $T = 2$ s.

integrator system the Hamiltonian becomes

$$\begin{aligned} \mathscr{H}(x,p) &= \max_{u\in\mathscr{U}} p^T f(x,u) \\ &= \max_{u\in\mathscr{U}}(p_1x_2 + p_2u) \\ &= p_1x_2 + |p_2u| \end{aligned} \tag{4.20}$$

with $p_1 = \frac{\partial\phi}{\partial x_1}$ and $p_2 = \frac{\partial\phi}{\partial x_2}$
The optimal control inputs for maximizing the Hamiltonian are given as

$$u = \begin{cases} 1 & \text{if } p_2 > 0 \\ -1 & \text{if } p_2 < 0 \\ [-1,1] & \text{if } p_2 = 0 \end{cases}$$

The initial condition will be approximated by a very small circle ($r = 0.05$) in the origin since the LS methods needs an implicit contour to work.

The results of the verification are shown in Figure 4.9. The red contour represents the true forwards reachable set. The blue contour represents the reachable set as calculated by the LS method starting from a small initial set (black). It can be seen that the two contours are almost identical. However, the LS method produces a reachable set that is slightly larger than the analytical solution. This is caused by the fact that the LS method is approximating the initial condition with a very small circle, essentially giving the LS method a head start. It is expected that as the state-space discretization is refined and the radius of the circle at the origin is reduced, the error will also reduce. Furthermore, the edges of the LS solution seem to be a bit more rounded off, which is an error introduced by the discretization.

4.5. Validation

For systems with a set as initial condition or non-linear systems, no analytical solution can be obtained. To validate the solution of the LS method, a Monte-Carlo (MC) method is used. The MC method works by simulating a large amount of individual trajectories that together will approximate the reachable set. For a single trajectory, the initial condition is chosen to be somewhere inside the initial set, then the trajectory is computed using a (semi-)random control sequence. The approximated reachable set is then computed as the boundary of all trajectories. The MC solution should approximate the reachable set as the number of samples

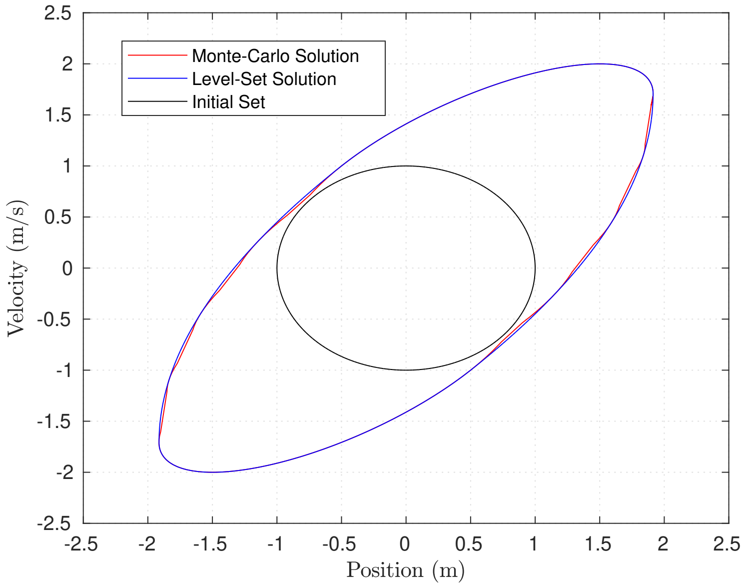

Figure 4.10: Validation of the double integrator forwards reachable set for $T = 1\,\mathrm{s}$.

increases. Since the MC method is an approximation of the reachable set, it is expected that the generated set is smaller than the LS solution. Furthermore, no MC samples should leave the LS reachable set. If this is the case, it could be a sign that the optimal control algorithm was not implemented correctly or indicate that numerical instabilities are present in the solver.

4.5.1. Double Integrator

The first example to be considered is the double integrator system as introduced in Section 4.4. However, instead of starting from a point in the origin, the initial condition is now represented by a set. This set is defined by a circle centered at the origin with unit radius. The initial set is shown in Figure 4.10 (black contour). The forwards reachable set is calculated using the LS method on a 200 x 200 grid with a time horizon of $T = 1\,\mathrm{s}$. The LS solution is plotted in Figure 4.10 as a blue contour. The MC solution was obtained by calculating the boundary of 16000 samples. These samples were started from 500 points inside the initial set with 32 different control sequences each. When comparing the LS solution with the results obtained from the MC analysis, it can be seen that the boundary of the MC solution is completely enclosed by the LS solution. Furthermore, the MC set is slightly smaller than the LS. This is expected since the MC solution only provides an approximation of the reachable set.

4.5.2. Pendulum on a Cart

The second example is a non-linear system that consists of a pendulum on a cart. A schematic representation of the system is shown in Figure 4.11. The dynamics of the system are

$$\frac{\mathrm{d}}{\mathrm{d}t}(\theta) = \dot{\theta} \tag{4.21}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}(\dot{\theta}) = \frac{u\cos\theta - (m_1 + m_2)g\sin\theta + ml(\cos\theta\sin\theta)\dot{\theta}^2}{ml\cos^2\theta - (m_1 + m_2)l} \tag{4.22}$$

with $m_1 = m_2 = 1\,\mathrm{kg}$, $l = 2\,\mathrm{m}$, $g = 9.81\,\mathrm{m\,s^{-2}}$ and $u = [-25;25]$ N. Note that the dynamics for the displacement of the cart have been left out to keep the dimension of the problem lower.

In order to solve this system using the LS toolbox it is necessary to define the Hamiltonian. Inserting the dynamic equation into Equation (4.17) yields

$$\begin{aligned} \mathscr{H}(x,p) &= \max_{u\in\mathscr{U}} p^T f(x,u) \\ &= \max_{u\in\mathscr{U}} \left(p_1 x_2 - p_2 \frac{u\cos\theta - (m_1 + m_2)g\sin\theta + ml(\cos\theta\sin\theta)\dot{\theta}^2}{ml\cos^2\theta - (m_1 + m_2)l} \right) \end{aligned} \tag{4.23}$$

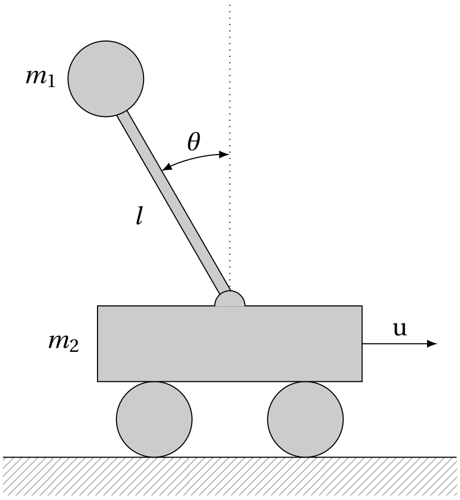

Figure 4.11: Schematic representation of the pendulum on a cart system.

The optimal control input for maximizing the Hamiltonian is

$$u = \begin{cases} 25 & \text{if } \frac{p_2 \cos\theta}{l(m_1+m_2)-lm\cos^2\theta} < 0 \\ -25 & \text{if } \frac{p_2 \cos\theta}{l(m_1+m_2)-lm\cos^2\theta} > 0 \\ [-25, 25] & \text{if } \frac{p_2 \cos\theta}{l(m_1+m_2)-lm\cos^2\theta} = 0 \end{cases}$$

The validation will consider the forwards reachable set of the pendulum on a cart model for both the LS and MC method. The initial condition for both methods will be a set of 1 by 2 centered at $\theta = \pm\pi$. Figure 4.12 shows the initial set in black. It should be noted that the boundary at $\theta = \pm\pi$ is periodic. The reachable set was computed using the LS method on a 200 x 200 grid for a time horizon of $T = 1$ s. The result of the LS method is shown in Figure 4.12 as a blue contour. The MC reachable set was computed as the boundary of 204800 samples started from 200 points inside the initial set with 1024 different control sequences. This boundary is shown as a red contour. When comparing the LS and MC solutions it can be seen that the MC result is contained within the LS solution, except for the far edges. This is caused by the rounding effect of the LS method. It is expected that refining the grid will decrease the round-off effect. Furthermore, the MC solution is slightly smaller which is also expected as the MC solution is an approximation to the reachable set. Increasing the number of samples should reduce this error.

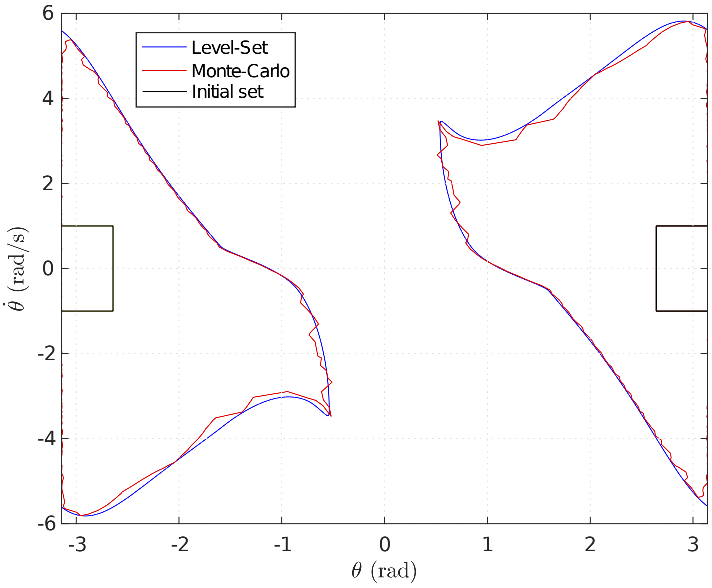

Figure 4.12: Validation of the pendulum on a cart forwards reachable set for $T = 1\,\text{s}$.

5

Probabilistic Reachability Analysis

The reachable set provides a bound on the states that can be reached within a given time horizon, starting from an initial set. However, while being theoretically accurate, they may not be reachable in reality. This is due to, for example, uncertainties in the model or external disturbances. Figure 5.1a shows the trajectories for a double integrator system to a target close to the initial set (purple circle). It can be seen that there are multiple trajectories that lead to this point. If a disturbance is present, there is most likely a different trajectory that still leads to the target. When the target is at the boundary of the reachable set, only one trajectory exists as shown in Figure 5.1b. Clearly there is a higher probability of reaching the first target than the second one. The purpose of this chapter is to develop a method for quantifying the effects of the disturbances on the reachable set. Section 5.1 will introduce the Brownian motion, which is used to model disturbances present in the system. Combining the Brownian motion with the dynamic system gives rise to an extension to ordinary differential equations. These so-called Stochastic Differential Equations (SDEs) are introduced in Section 5.2. In order to solve these SDEs, Section 5.3 explains the concept of Itô calculus. A link between SDEs and the Probability Density Functions (PDFs) is established in Section 5.4. Section 5.5 introduces the stochastic reachability analysis. An alternative method to approximate the probabilistic reachable set is explained in Section 5.6. Verification and validation of these methods is presented in Section 5.7 and 5.8 respectively.

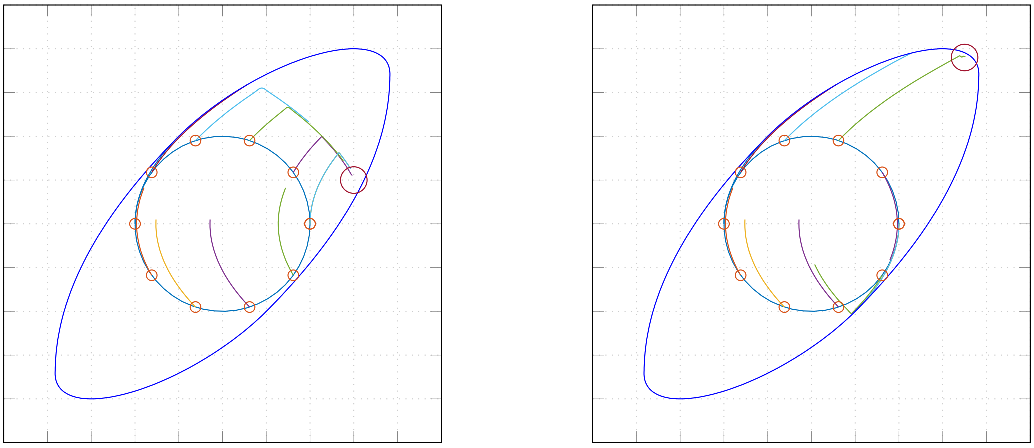
(a) Multiple trajectories to a target close to the initial set. (b) Trajectories to a target far from the initial set.

Figure 5.1: Difference between the reachability of states close and far from the initial set.

5.1. Brownian Motion

Brownian motion was discovered in the early nineteenth century by Robert Brown. It was used to describe the random motion of particles suspended in water. The Brownian motion is one of the simplest continuous-time stochastic processes. The Wiener process or Brownian motion can be seen as the continuous time variant of the simple random walk (the limit when Δt goes to zero). A Brownian motion is a random variable $B(t)$ that satisfies the following three conditions [21]

- $B(0) = 0$ (with probability 1).
- For $0 \leq s < t \leq T$ the random variable given by the increment $B(t) - B(s)$ is normally distributed with zero mean and variance $t - s$; equivalently, $B(t) - B(s) \sim \sqrt{t-s}N(0,1)$.
- For $0 \leq s < t < u < v \leq T$ the increments $B(t) - B(s)$ and $B(v) - B(u)$ are independent.

The most important properties of the Brownian motion are that the increments are normally distributed and that they are independent (i.e. the next increment is independent of all previous values). The Brownian motion is a continuous-time stochastic process which is used to represent the integral of a white noise process.

5.2. Stochastic Differential Equations

The first step in quantifying the effect of disturbances on the reachable set is to model the disturbances present in the dynamic system. When including stochastic processes, the ordinary differential equation becomes an SDE. An SDE is a differential equation in which one or more terms are stochastic processes. For example Equation (5.1) shows an SDE where b and σ are some given functions [22].

$$\frac{\mathrm{d}X_t}{\mathrm{d}t} = b(t, X_t) + \sigma(t, X_t) \cdot \text{“noise”} \tag{5.1}$$

It is reasonable to look for some stochastic process W_t to represent the noise term, so that

$$\frac{\mathrm{d}X_t}{\mathrm{d}t} = b(t, X_t) + \sigma(t, X_t) \cdot W_t \tag{5.2}$$

where W_t is a white noise process.

A suitable stochastic process with continuous paths can be found in the Brownian motion B_t which is obtained as the integral of a white noise signal. After substituting the Brownian motion as defined in Section 5.1, a solution to the SDE can be obtained by applying the usual integration notation

$$X(T) = X(0) + \int_0^T b(t, X_t)dt + \text{“}\int_0^T \sigma(t, X_t)dB_t\text{”} \tag{5.3}$$

where the first integral is the standard Riemann integral which can be seen as the drift term and the second is the Itô integral or the variance term.

A heuristic interpretation of the SDE is that over a small time interval $\mathrm{d}t$, the value of the stochastic process changes with an amount that is normally distributed with expectation $b(t, X_t)\mathrm{d}t$ and variance $\sigma(t, X_t)^2\mathrm{d}t$. Although being continuous in time, the Brownian motion is nowhere differentiable since it is defined as the limit of a simple random walk. When integrating white noise, the result is a Brownian motion. This stochastic process has a Gaussian distribution with quadratic variation. An extension of the methods of calculus for stochastic processes such as the Brownian motion is given by Itô's calculus. The concept of Itô's calculus can be seen as the stochastic counterpart of the Riemann-Stieltjes integral. Section 5.3 will provide a short introduction to Itô's lemma, which can be seen as the stochastic version of the chain-rule.

5.3. Itô Calculus

Itô calculus is an extension to the methods of calculus for stochastic processes such as the Wiener process. It will be the starting point for the analysis of stochastic differential equations and probabilistic reachability analysis.

Assume X_t is an Itô drift diffusion process [22]

$$\mathrm{d}X_t = \mu_t \mathrm{d}t + \sigma_t \mathrm{d}B_t \tag{5.4}$$

if $f(t, x)$ is a twice differentiable scalar function, its Taylor series expansion is given by

$$\mathrm{d}f = \frac{\partial f}{\partial t}\mathrm{d}t + \frac{\partial f}{\partial x}\mathrm{d}x + \frac{1}{2}\frac{\partial^2 f}{\partial x^2}\mathrm{d}x^2 + \ldots \tag{5.5}$$

Substituting X_t for x and therefore $\mu_t \mathrm{d}t + \sigma_t \mathrm{d}B_t$ for $\mathrm{d}x$ gives

$$\mathrm{d}f = \frac{\partial f}{\partial t}\mathrm{d}t + \frac{\partial f}{\partial x}(\mu_t \mathrm{d}t + \sigma_t \mathrm{d}B_t) + \frac{1}{2}\frac{\partial^2 f}{\partial x^2}(\mu_t^2 \mathrm{d}t^2 + 2\mu_t\sigma_t \mathrm{d}t\,\mathrm{d}B_t + \sigma_t^2 \mathrm{d}B_t^2) + \ldots \tag{5.6}$$

Because of the Wiener process, the term $\mathrm{d}B_t^2$ is of order $\mathcal{O}(\mathrm{d}t)$. For the limit of $\mathrm{d}t \to 0$, the terms $\mathrm{d}t^2$ and $\mathrm{d}t\ \mathrm{d}B_t$ will go to zero faster than $\mathrm{d}B_t^2$. Setting these terms to zero and substituting $\mathrm{d}t$ for $\mathrm{d}B_t^2$ yields

$$\mathrm{d}f = \left(\frac{\partial f}{\partial t} + \mu_t \frac{\partial f}{\partial x} + \frac{\sigma_t^2}{2}\frac{\partial^2 f}{\partial x^2}\right)\mathrm{d}t + \sigma_t \frac{\partial f}{\partial x}\mathrm{d}B_t \tag{5.7}$$

Itô's lemma also exists for higher dimensional systems. For a system of SDEs where $\mathbf{X}_t = \left(X_t^1, X_t^2, .., X_t^n\right)^T$ with stochastic process

$$\mathrm{d}\mathbf{X}_t = \boldsymbol{\mu}_t \mathrm{d}t + \mathbf{G}_t \mathrm{d}\mathbf{B}_t \tag{5.8}$$

Itô's lemma states that

$$\mathrm{d}f(t,\mathbf{X}_t) = \left\{\frac{\partial f}{\partial t} + (\nabla_X f)^T \boldsymbol{\mu}_t + \frac{1}{2}\mathrm{Tr}\left[\mathbf{G}_t^T (H_X f)\mathbf{G}_t\right]\right\}\mathrm{d}t + (\nabla_X f)^T \mathbf{G}_t \mathrm{d}\mathbf{B}_t \tag{5.9}$$

where $\nabla_x f$ is the gradient of f and H_x is the Hessian matrix of f with respect to X.

5.3.1. Example

In order to get a feel for the theory, a small example will be covered that applies Itô's lemma to the integration of an SDE. Consider the Geometric Brownian Motion (GBM) process defined by [22]

$$\mathrm{d}X_t = \mu X_t \mathrm{d}t + \sigma X_t \mathrm{d}B_t \tag{5.10}$$

or

$$\frac{\mathrm{d}X_t}{X_t} = \mu \mathrm{d}t + \sigma \mathrm{d}B_t \tag{5.11}$$

Hence

$$\int_0^t \frac{\mathrm{d}X_t}{X_t} = \mu t + \sigma B_t \quad (B_0 = 0) \tag{5.12}$$

Evaluate the integral on the left hand side by means of Itô's lemma for the function $f(t,x) = \ln(x)$

$$\begin{aligned} \mathrm{d}(\ln(X_t)) &= \frac{1}{X_t}\cdot \mathrm{d}X_t + \frac{1}{2}\left(-\frac{1}{X_t^2}\right)(\mathrm{d}X_t)^2 \\ &= \frac{\mathrm{d}X_t}{X_t} - \frac{1}{2X_t^2}\cdot\sigma^2 X_t^2 \mathrm{d}t \\ &= \frac{\mathrm{d}X_t}{X_t} - \frac{1}{2}\sigma^2 \mathrm{d}t \end{aligned} \tag{5.13}$$

Hence

$$\frac{\mathrm{d}X_t}{X_t} = \mathrm{d}(\ln(X_t)) + \frac{1}{2}\sigma^2 \tag{5.14}$$

Substitution in (5.12) yields

$$X_t = X_0 e^{(\mu - \frac{1}{2}\sigma^2)t + \sigma B_t} \tag{5.15}$$

Although the analytical solution is unique, this is only true for a single realization of the noise process. Every sample of the GBM process will therefore be different. Figure 5.2 shows 15 different samples of the GBM process. The analytical solution of an SDE is not always trivial, therefore several numerical methods are available in order to provide an iterative solution. An extension to the standard Euler integration scheme is given by [23]

$$x_{i+1} = x_i + b(x_i, t_i)\Delta t + \sigma(x_i, t_i)\Delta B_i \tag{5.16}$$

where

$$\Delta B_i = B(t_{i+1}) - B(t_i) \tag{5.17}$$

as defined in Section 5.1. An illustration of this iterative scheme is shown in Figure 5.3. Again, $B(t_n)$ are samples of the Brownian motion for a given realization of the stochastic process. The stochastic Euler scheme is very similar to the deterministic variant with the addition of the stochastic part that is dependent on the noise variance σ which scales with the square root of the time step.

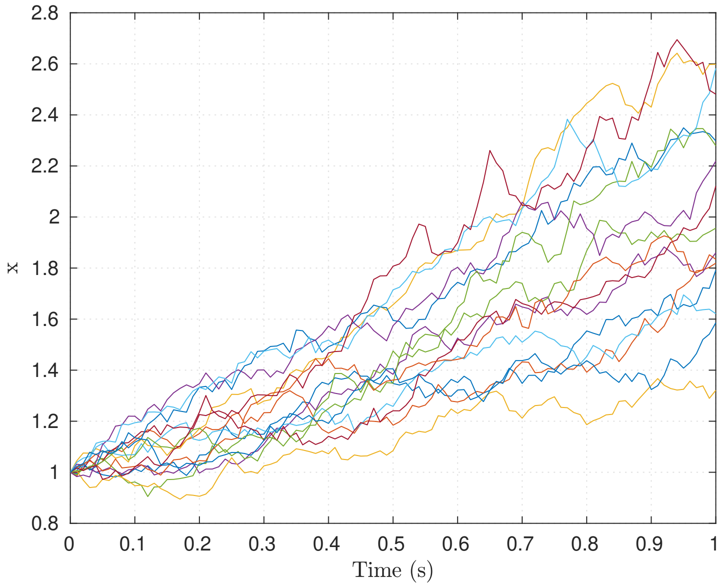

Figure 5.2: The 15 realizations of the GBM process all follow a slightly different trajectory.

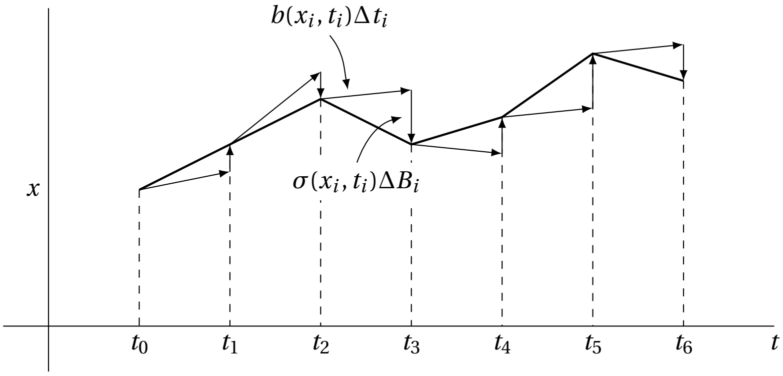

Figure 5.3: Euler integration scheme for constructing an approximate solution to the SDE.

5.4. Fokker-Planck-Kolmogorov Equation

In the previous section it was shown that each sample of a stochastic process follows a slightly different trajectory. However, when looking at the ensemble of all possible trajectories, trends can be detected. These trends are usually expressed in their moments such as mean (1^{st} moment) and variance (2^{nd} moment). A more complete picture can be obtained by looking at the distribution of the stochastic process. This distribution is often expressed as a PDF where the area under the PDF is related to the probability. It is possible to determine the statistics by sampling the process often, then determining the discrete PDF (i.e. histogram) from these samples. A more powerful approach is to formulate an expression for the PDF directly from the SDE. When starting from a known initial distribution, the distribution after a certain time can be described by the time rate of change of the probability. The forward Kolmogorov equation (or Fokker-Planck equation) describes the time evolution of the PDF. This section will try to provide an intuitive description of the Fokker-Planck equation. For the complete mathematical derivation the interested reader is referred to Risken [24].

Consider the two-dimensional lattice which can be thought of as a Cartesian grid with vertical spacing Dt and horizontal spacing Dx [25]. Points on the grid are labeled by a pair of integers (i, n) which specify their horizontal and vertical coordinate respectively. An example of such a grid is shown in Figure 5.4. The probability of each node at position i and time n is denoted $\tilde{p}(i, n)$. Furthermore, the transition probability from point j to i in one unit of time (Dt) is denoted by $T(i|j)$.

At each time step the particle must remain on the same place, or it can move to one of the neighbors. This

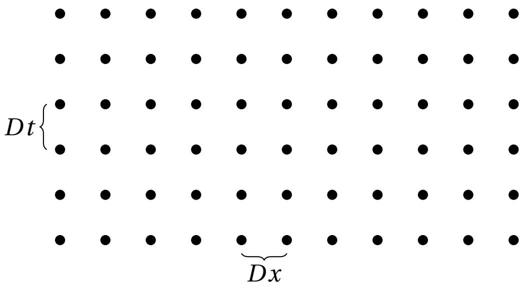

Figure 5.4: Two-dimensional lattice.

balance is shown in Figure 5.5.

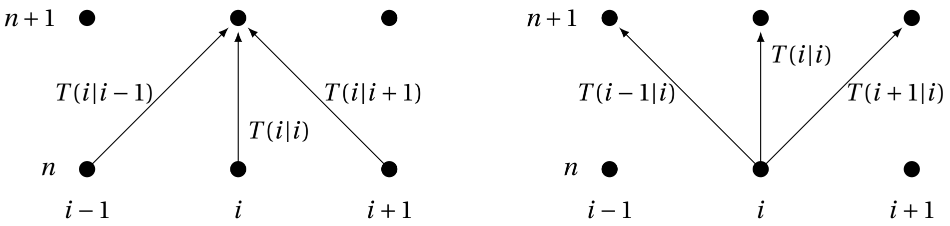

Figure 5.5: Probability balance.

The particle balance then becomes

$$\tilde{p}(i,n+1) = T(i|i-1)\tilde{p}(i-1|n) + T(i,i)\tilde{p}(i,n) + T(i|i+1)\tilde{p}(i+1,n) \tag{5.18}$$

Where the change of $\tilde{p}$ over one time step then is equal to

$$\tilde{p}(i,n+1) - \tilde{p}(i,n) = T(i|i-1)\tilde{p}(i-1,n) - \left[T(i+1|i) + T(i-1|i)\right]\tilde{p}(i,n) + T(i|i+1)\tilde{p}(i+1,n) \tag{5.19}$$

Now the spatial derivative over the lattice will be defined by the lattice derivative operation which is a first order Taylor expansion around n

$$D_{n,\pm 1}\tilde{f}(n) = \tilde{f}(n\pm 1) - \tilde{f}(n) \tag{5.20}$$

then the (first and second) derivative can be acquired by taking the limit for Dx to zero

$$\lim_{Dx\downarrow 0} \frac{D_{i,1}}{Dx}\tilde{f}(n) = \lim_{Dx\downarrow 0} \frac{f(x+Dx) - f(x)}{Dx} = \frac{\partial f}{\partial x}(x) \equiv \partial_x f \tag{5.21}$$

$$\lim_{Dx\downarrow 0} \frac{D_{i,1} + D_{i,-1}}{(Dx)^2}\tilde{f}(n) = \lim_{Dx\downarrow 0} \frac{f(x+Dx) - 2f(x) + f(x-Dx)}{(Dx)^2} = \frac{\partial^2 f}{\partial x^2} \equiv \partial_x^2 f \tag{5.22}$$

Suppose that $T(i,j)$ only distinguishes whether the particle is going to the right or to the left.

$$T(i|j) = \begin{cases} T_+ & i-j=1 \\ T_- & i-j=-1 \end{cases}$$

In such case the particle balance becomes

$$D_{n,1}\tilde{p}(i,n) = \frac{T_+ + T_-}{2}(D_{i,1} + D_{i,i-1})\tilde{p}(i,n) - (T_+ - T_-)\frac{D_{i,1} - D_{i,-1}}{2}\tilde{p}(i,n) \tag{5.23}$$

where $(T_+ + T_-)/2$ is equal to the diffusion and $T_+ - T_-$ is equal to the drift term.

Taking the limit of the lattice grid spacing Dx and Dt to zero, $D_{i,1} + D_{i,i-1}$ becomes the second derivative ∂_x^2 and $(D_{i,1} - D_{i,-1})/2$ is equal to the central difference derivative ∂_x. Finally, the Fokker-Planck equation is given as

$$\frac{\partial}{\partial t}p(x,t) = -\frac{\partial}{\partial x}\left(b(x,t)p(x,t)\right) + \frac{\partial^2}{\partial x^2}\left(\frac{\sigma(x,t)^2}{2}p(x,t)\right) \tag{5.24}$$

where the drift term is now represented by the function $b(x,t)$ and the diffusion by $\sigma(x,t)$. Note that $p(x,t)$ represents the PDF, not to be confused with the co-states of the Hamiltonian.

5.4.1. Example

Consider again the GBM process [26]

$$\mathrm{d}X_t = \mu X_t \mathrm{d}t + \sigma X_t \mathrm{d}B_t \tag{5.25}$$

Substitution of $b(x,t) = \mu X_t$ and $\sigma(x,t) = \sigma X_t$ into Equation (5.24) yields

$$\frac{\partial p}{\partial t} = -\frac{\partial}{\partial x}((\mu x)p(x,t)) + \frac{1}{2}\frac{\partial^2}{\partial x^2}((\sigma x)^2 p(x,t)) \tag{5.26}$$

The solution to this second order PDE can be expressed analytically for the case of a Dirac delta as initial condition ($p(x,0) = \delta(x)$) and is equal to

$$P(x;\mu,\sigma,t) = \frac{1}{\sqrt{2\pi}}\frac{1}{x\sigma\sqrt{t}} e^{-\frac{\left(\ln x - \ln x_0 - \left(\mu - \frac{1}{2}\sigma^2\right)t\right)^2}{2\sigma^2 t}} \tag{5.27}$$

Figure 5.6 shows individual samples of the GBM process on the $x-y$ plane. The analytical probability density is shown as a function of x and t. It can be seen that the mean of the sample paths follows the peak of the PDF. Furthermore, as the variance of the PDF increases, the sample paths seem to spread out more.

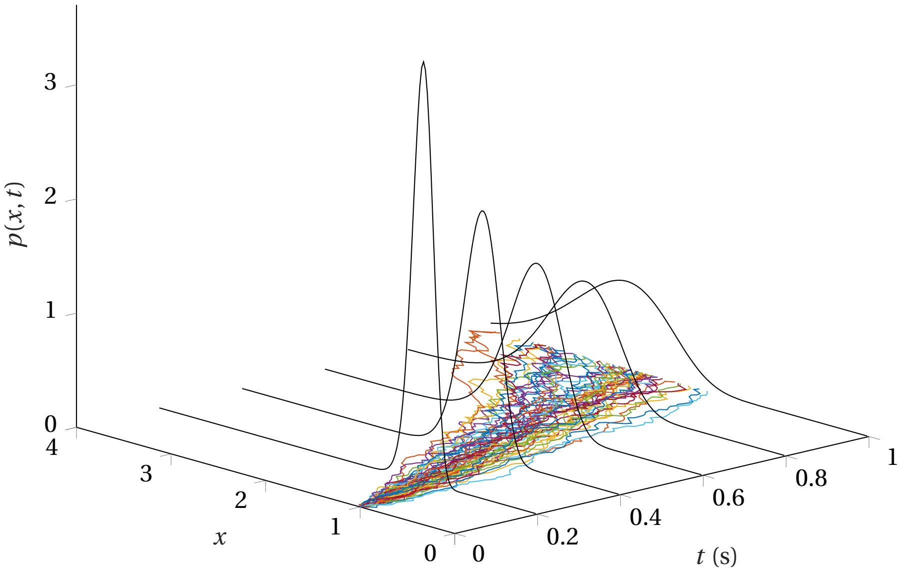

Figure 5.6: Individual samples of the GBM process and the corresponding PDF.

5.5. Probabilistic Reachability Analysis

Combining the SDEs with reachability analysis as introduced in Chapter 4 gives rise to the field of probabilistic reachability analysis. Instead of calculating the set of all states that can be reached starting from an initial set, the probabilistic reachability analysis will compute the transition probability to a certain state for a given initial probability. The transition probability is equal to the joint probability of being in the initial set $\mathscr{K}$ at time $t = 0$ and in a final state X at time $t = T$.

Consider a controlled stochastic process $X^u_{t,x}$ which is the solution to the following SDE

$$\mathrm{d}X(s) = f(X(s), s, u(s))\mathrm{d}t + \sigma(X(s), s)\mathrm{d}B_s, \quad \forall s \in [t, T] \tag{5.28}$$

Let $\mathscr{K}$ be the non-empty target set in $\mathbb{R}^n$, $\rho \in [0,1[$ and $t \leq T$. Consider the reachable set Ω_t under probability of success ρ, or the set of initial conditions x for which the probability that there exists a trajectory $X^u_{t,x}$, solution of (5.28), that reaches set $\mathscr{K}$ at time T, associated with the admissible control $u \in \mathscr{U}$ is at least ρ [27]

$$\Omega^\rho_t = \left\{x \in \mathbb{R}^n | \exists u \in \mathscr{U}, \mathbb{P}[X^u_{t,x}(T) \in \mathscr{K} > \rho]\right\} \tag{5.29}$$

Because the dynamics are stochastic, it is no longer possible to minimize the value function. Instead the expected pay-off is minimized over all possible further realizations of the Wiener process [28, 29].

$$V(x,t) = \min_{\mathbf{u}\in\mathcal{U}} \left\{ \mathbb{E}\left[V(x + \Delta f(x,u) + \xi, t + \Delta t) \right] \right\} \tag{5.30}$$

with $\xi \sim \mathcal{N}(0, \sigma(x,t))$ a process with normal distribution. The derivation is performed in a similar way as the deterministic case in Chapter 4. Taking the Taylor expansion of the value function by means of Itô's calculus. Since $\mathrm{d}x^2$ is of the order $\mathrm{d}t$ because of the Wiener process, the expansion must be performed up to the second order.

$$V(x+\Delta x, t+\Delta t) \approx V(x,t) + V_t(x,t)\Delta t + V_x(x,t)\Delta x + \frac{1}{2}\left(V_{xx}(x,t)\Delta x^2 + 2V_{xy}\Delta x \Delta t + V_{tt}(x,t)\Delta t^2\right) + \mathcal{O}(\Delta^3) \tag{5.31}$$

where $\Delta x = \Delta f(x,t,u) + \xi$

Keeping all terms of order $\mathcal{O}(\Delta t)$, similar to section 5.3

$$\mathbb{E}[V] = V(x,t) + \Delta f(x,u) V_x(x,t) + \frac{1}{2}\mathrm{Tr}(\sigma^2(x,t) V_{xx}(x,t)) \tag{5.32}$$

After rearranging by the principle of optimality and substituting into equation (5.30) then dividing by Δt

$$\frac{V(x,t) - V(x,t+\Delta t)}{\Delta t} = \min_{\mathbf{u}\in\mathcal{U}} \left\{ V_x f(x,t,u)^T + \frac{1}{2}\mathrm{Tr}(\sigma(x,t)\sigma(x,t)^T V_{xx}) \right\} \tag{5.33}$$

For the limit of $\Delta t \to 0$ the value function is characterized by the viscosity solution to the stochastic HJB equation [30]

$$\frac{\partial V}{\partial x}(x,t) + \inf_{u\in\mathcal{U}} \left\{ \frac{\partial V}{\partial x}(x,t) f(x,t,u) \right\} + \frac{1}{2}\mathrm{Tr}\left\{ \sigma(x,t)\sigma(x,t)^T \frac{\partial^2 V}{\partial x^2}(x,t) \right\} = 0 \tag{5.34}$$

The stochastic HJB equation can be solved using a similar method as for the deterministic HJB equation. Next to the convection term, an additional diffusion term is added to the LS method. In contrary to the deterministic reachable set, the probabilistic reachable set is not represented by a zero level contour of the value function. Instead, the signed distance function is replaced by an indicator function.

$$\mathbb{1}_{\mathcal{K}}(x) = \begin{cases} 1 & \text{if } x \in \mathcal{K} \\ 0 & \text{if } x \notin \mathcal{K} \end{cases}$$

The initial condition of the LS function is then equal to one for all states that are inside the initial set $\mathcal{K}$, and all other states are equal to zero. Figure 5.7a shows an illustration of the indicator function which represents the initial set. The regularization of the indicator function is shown in Figure 5.7b. Due to the discretization of the state-space, the LS function is slightly smaller which can introduce errors in the final solution. An error estimation is performed by Assellaou et al. [27]. The concavity of the LS function is represented by the arrows in Figure 5.7c. For a diffusion only problem, the second derivative, as found in the stochastic part of the HJB equation represents this concavity. Figure 5.7d shows how the concavity of the LS function is related to the diffusion.

5.6. Path Integral Monte Carlo

Since an analytical solution to the Fokker-Planck equation can only be found for several simple equations, numerical techniques have been developed to approximate the solution. In the previous section, the stochastic HJB PDE was introduced which could be solved using a LS method. An alternative to solving these PDEs is the Path Integral Monte Carlo (PIMC) method. The PIMC method allows a step-by-step capture of the entire evolution of the response process in terms of PDF, starting from a known initial condition (deterministic or stochasctic). The PIMC method has been used for a long time in physics to numerically solve the Fokker-Planck equation.

The starting point of the PIMC method is the Chapman-Kolmogorov equation [31]

$$p_x(x, t+\tau) = \int_D p_x(x, t+\tau | \bar{x}, t) p_x(\bar{x}, t) \mathrm{d}\bar{x} \tag{5.35}$$

When the PDF is known at the current time step t, it is possible to evaluate the PDF after a small time step τ which holds true, because of the Markovian property of the response.

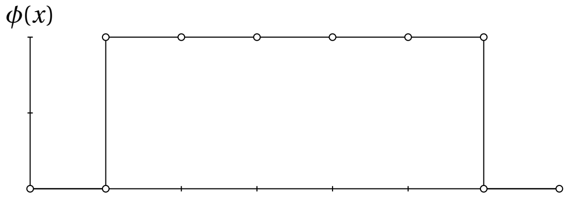

(a) Indicator function representing the initial set $\mathcal{K}$ (b) Regularization of the indicator function.

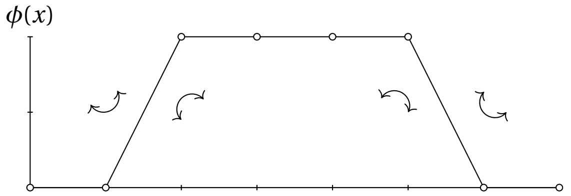

(c) Concavity of the LS function. (d) Diffusion of the LS function.

Figure 5.7: Illustration of the diffusion on the LS function.

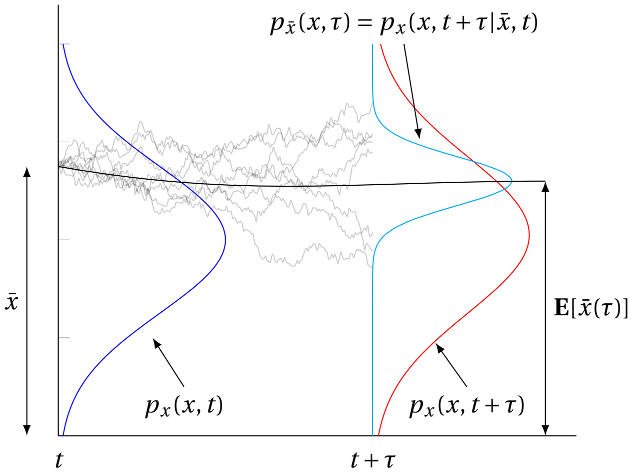

Figure 5.8: Path Integral Monte Carlo method.

Figure 5.8 illustrates the numerical solution to the Chapman-Kolmogorov equation for a single point. The blue line represents the (Gaussian) distribution of the SDE at the current time t. This distribution is denoted by $p_x(x,t)$. The PIMC method samples the SDE over the domain D. Figure 5.8 shows these samples for a single point denoted by $\bar{x}$ (not to be confused with the mean value). The gray samples are all unique realizations of the SFE starting from $\bar{x}$ over a time horizon τ. The PDF corresponding to these samples is shown in cyan and is denoted $p_x(x, t+\tau|\bar{x}, t)$. The PDF belonging to $\bar{x}$ (i.e. $p_x(x, t+\tau|\bar{x}, t)$) is then weighted by the previous probability at $p(\bar{x}, t)$. The PDF at the next time instant $(t+\tau)$ is shown in red. This PDF is the ensemble of all local (Gaussian) probability distributions over the domain D.

The size of the domain is identified by running a MC simulation with a low number of samples first. Then the domain is divided into a fine grid. The PIMC solution is then evaluated at each grid point using Equation (5.35).

5.7. Verification

For a few specific cases, an analytical solution to the stochastic HJB equation exists. The verification will compare the analytical solution with the numerical solution as calculated by the LS method. The system that will be considered is a stochastic system with constant diffusion and no drift. This system is described by

$$\begin{bmatrix} \mathrm{d}X_1 \\ \mathrm{d}X_2 \end{bmatrix} = \sigma \begin{bmatrix} c & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} \mathrm{d}W_t^1 \\ \mathrm{d}W_t^2 \end{bmatrix} \tag{5.36}$$

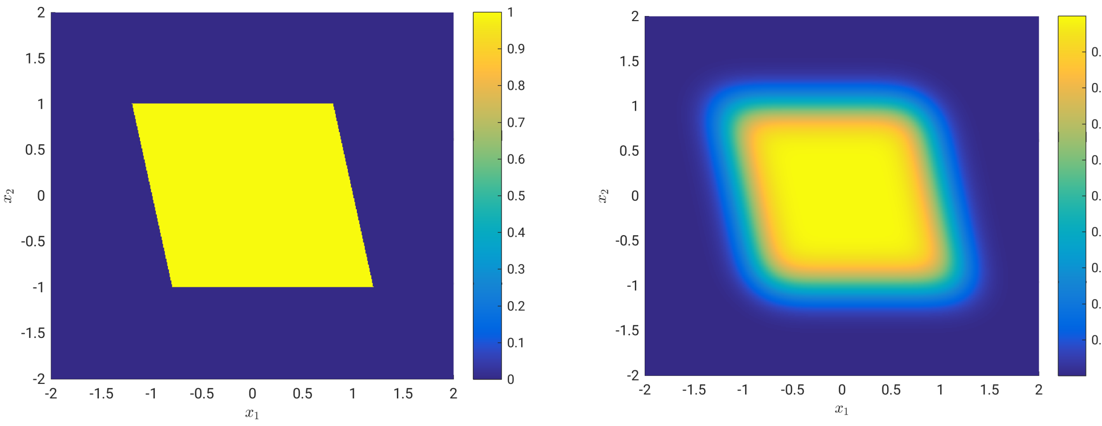

(a) Initial condition.

(b) LS Solution.

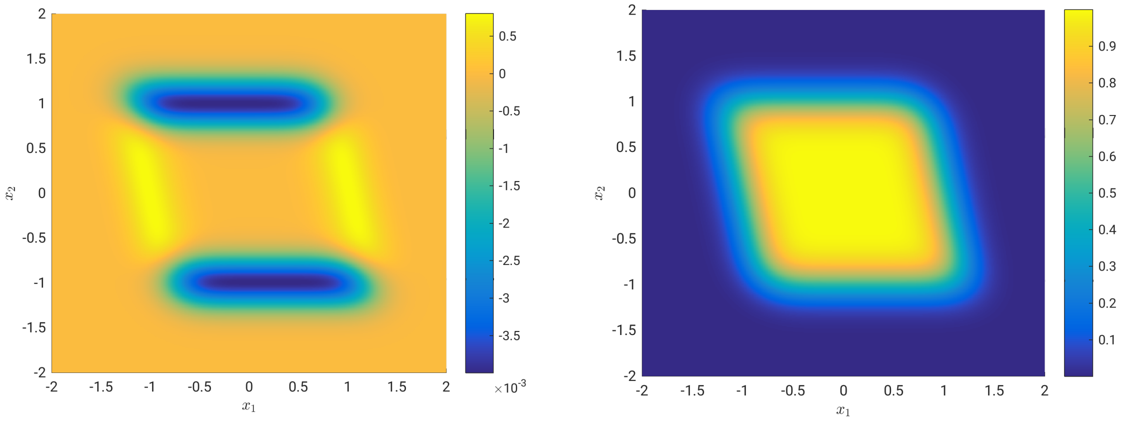

(c) Difference between the LS and analytical solution.

(d) Analytical Solution.

Figure 5.9: Verification of the diffusion only model probabilistic reachability for $T = 1$ s.

with $c = 0.2$ and $\sigma = 0.2$. The time horizon is $T = 1$ s and the initial set $\mathscr{K}$ is a tetragon with corners at $(-0.8,-1),(1.2,-1),(0.8,1),(-1.2,1)$ as shown in Figure 5.9a. The initial set is defined by an indicator function. For this special case an exact solution to the HJB equation exists and is given by [27]

$$v(t,x,y) = v_1(t,x+cy)v_1(t,y) \tag{5.37}$$

with

$$v_1(t,r) = \frac{1}{\sqrt{2\pi\sigma^2 t}} \int_{-1-r}^{1-r} e^{-\frac{s^2}{2\sigma^2 t}} \mathrm{d}s \tag{5.38}$$

The results of the verification are shown in Figure 5.9. Figure 5.9a shows the initial condition of the probabilistic reachability analysis. The analytical solution is shown in Figure 5.9d and the result from the LS method can be seen in Figure 5.9b. The difference between the analytical and the LS method is shown in Figure 5.9c. It can be seen that there is a small difference in the two solutions, however this is of order 10^{-3} which is considered to be negligible. Unfortunately, verification of non-linear or controlled systems is not possible since no analytical solution are know for these problems.

5.8. Validation

For more complex problems such as non-linear systems or controlled stochastic reachability, no analytical solution to the stochastic HJB equation exists. The PIMC method will be used to validate the results of the LS method. For a controlled stochastic process, this approach is used in combination with a state-feedback optimal control policy which steers the dynamic system towards the target set. By making use of Bellman's Optimality Principle as explained in Section 4.1, it is possible to extract the optimal control policy from the

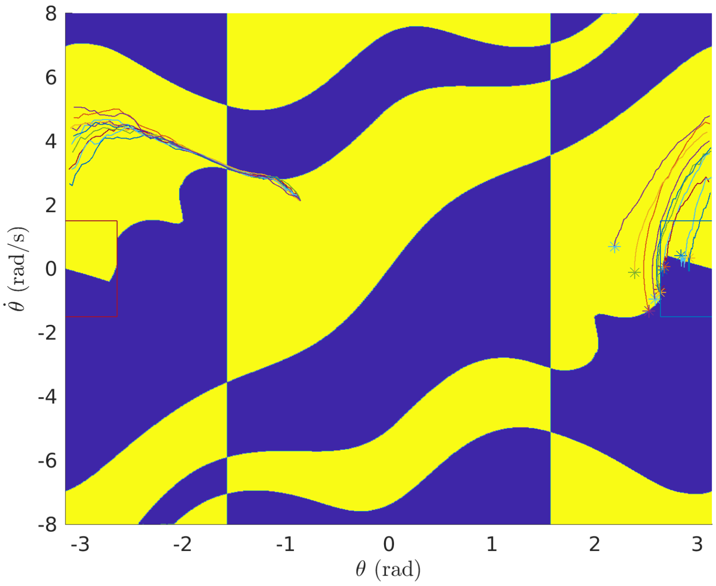

Figure 5.10: MC samples following the optimal control policy (Yellow maximum control input, Blue minimum control input)

deterministic LS solution in the form of a state feedback law. In order to extract the optimal control policy from the LS method, the value of the co-states is saved as the boundary of the LS method crosses a grid-point (i.e. when the sign of the value function changes).

With the co-states fixed, the optimal control that maximizes the Hamiltonian can then be computed. An example is shown in Figure 5.10. The yellow and blue patches represent the minimum and maximum controls for the bang-bang policy of the pendulum on a cart system. During the MC analysis this state-feedback is used to steer the trajectories towards the initial set. A number of samples starting from one point in the state-space are also plotted in Figure 5.10. It can be seen that not all trajectories reach the initial set under the influence of the external disturbance which is related to the reachability probability (it is essentially the same as the PIMC method with a very long time step).

5.8.1. Double Integrator

The first case for the validation is the stochastic counterpart of the double integrator system as introduced in Section 4.4, but with the addition of additive noise. This means that the input forcing the system is subjected to zero mean white Gaussian noise with a standard deviation of 0.5.
The SDEs describing the dynamic system now becomes

$$\dot{\boldsymbol{x}}(t) = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \boldsymbol{x}(t) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \boldsymbol{u}(t) + \begin{bmatrix} 0 \\ 0.5 \end{bmatrix} \boldsymbol{W} \tag{5.39}$$

$$\boldsymbol{y}(t) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \boldsymbol{x}(t) + \begin{bmatrix} 0 \\ 0 \end{bmatrix} \boldsymbol{u}(t) \tag{5.40}$$

The SDE is solved forwards in time over a 1 s time horizon. Figure 5.11a shows the initial set which is a circle located at the origin with unit radius. The probabilistic reachable set calculated using the LS method is shown in Figure 5.11b. Figure 5.11d shows the result generated by the PIMC method. The difference between the PIMC and the LS method is shown in Figure 5.11c. It can be seen that the difference between the two methods is generally small, except for two spots on the border of the initial set. Since the PIMC method is an approximation of the probabilistic reachable set, it is in the line of expectations that the transition probability is similar or lower than the LS method. This difference is thought to be caused by the underlying optimal control law. To check if the underlying optimal control policy as extracted from the deterministic LS solution is the cause of the errors, a separate MC analysis is performed which does not depend on the optimal control policy. Instead, a number of random bang-bang control sequences will be generated, similar to the deterministic MC analysis. For each control sequence, a number of samples will be simulated with different noise realizations. Then, the transition probability will be selected which corresponds to the control sequence

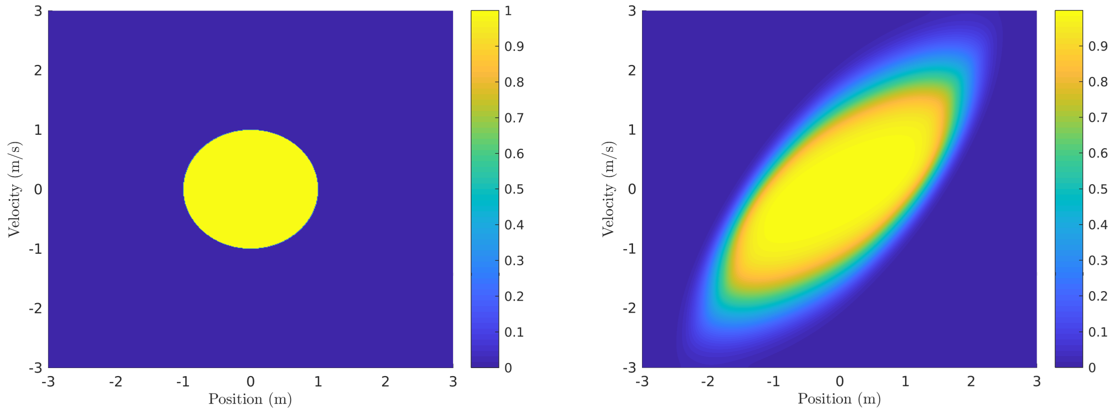

(a) Initial condition.

(b) LS solution.

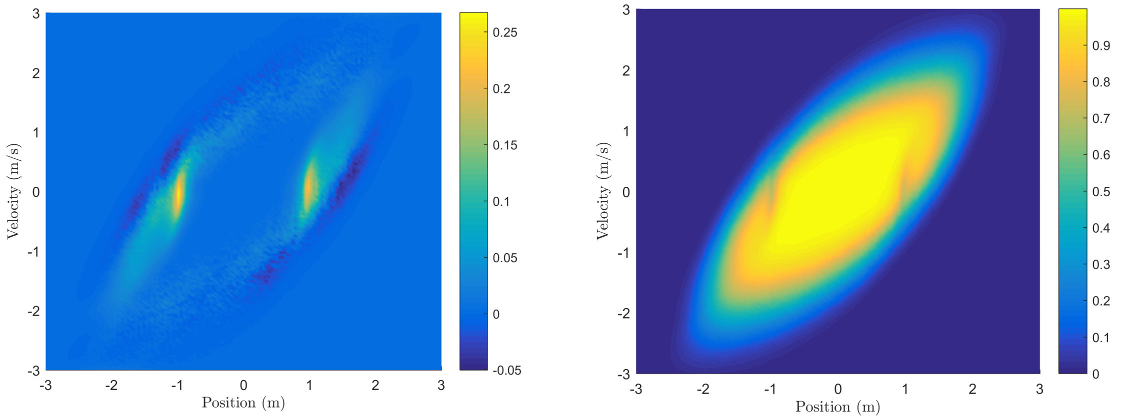

(c) Difference between the LS and PIMC method.

(d) PIMC solution.

Figure 5.11: Validation of the double integrator probabilistic reachability for $T = 1\,\mathrm{s}$.

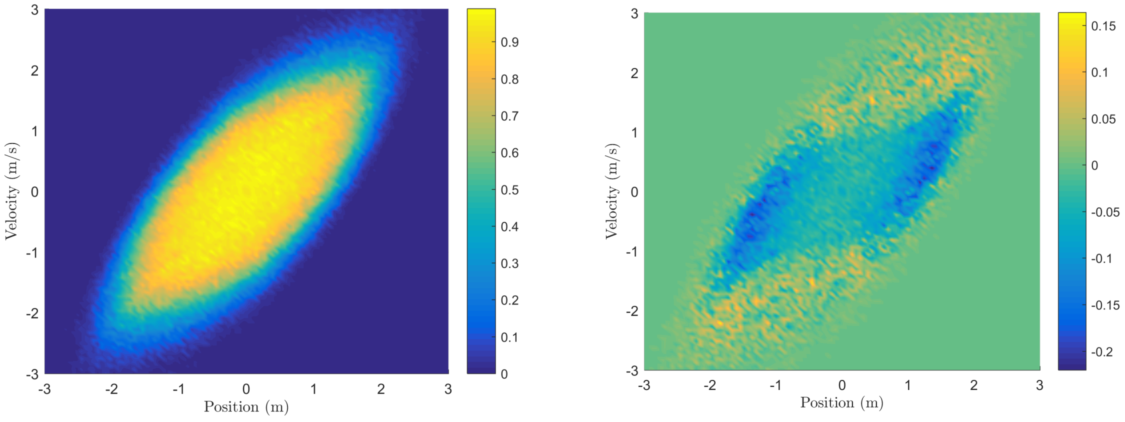

(a) MC Solution.

(b) Difference between the LS and bang-bang solution.

Figure 5.12: Validation of the double integrator without optimal control policy probabilistic reachability for $T = 1\,\mathrm{s}$.

that yields the highest value. It should be noted that these control sequences are essentially open loop control instead of the state feedback of the optimal control policy. The result of the MC analysis without the underlying optimal control policy is shown in Figure 5.12a. Finally, the difference between the two methods can be seen in Figure 5.12b. The result is very similar to the PIMC method with the optimal control policy. No control sequences were found that yield a significantly higher probability. Furthermore, the error on the boundary of the initial set is not present in this solution. This confirms that the discrepancy in the PIMC method is likely caused by the underlying optimal control policy as extracted from the deterministic LS method.

5.8.2. Pendulum on a Cart

As with the deterministic reachability analysis, the second validation case will be a non-linear pendulum on a cart model. The force driving the cart is now augmented with a disturbance force. In addition to the non-linear model, the noise is not additive anymore, but state dependent.

The resulting SDEs that describe the system dynamics are

$$\frac{\mathrm{d}}{\mathrm{d}t}(\theta) = \dot{\theta} \tag{5.41}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}(\dot{\theta}) = \frac{u\cos\theta - (m_1+m_2)g\sin\theta + ml(\cos\theta\sin\theta)\dot{\theta}^2}{ml\cos^2\theta - (m_1+m_2)l} + \frac{\sigma\cos\theta}{ml\cos^2\theta - (m_1+m_2)l}W \tag{5.42}$$

where $m_1 = m_2 = 1\,\mathrm{kg}$, $l = 2\,\mathrm{m}$, $g = 9.81\,\mathrm{m\,s^{-2}}$, $\sigma = 2\,\mathrm{N}$ and $u = [-25; 25]\,\mathrm{N}$. Since the noise is not dependent on the control, the Hamiltonian and the corresponding maximizers stay the same as presented in Section 4.5.2.

Figure 5.13 shows the reachability analysis for both the LS as well as the PIMC method. The initial set is again centered around $\theta = \pm\pi$ and is depicted in Figure 5.13a. Furthermore, the boundary at $\pm\pi$ is periodic. The LS solution was performed on a 1200 x 1200 grid (high resolution is a result of the convergence analysis performed in Appendix A) and is shown in Figure 5.13b. The PIMC solution as shown in Figure 5.13d was performed on a 150 x 150 grid with 100 samples per grid point. The difference between the PIMC and the LS method is plotted in Figure 5.13c. It can be seen that the two methods correspond quite well except for the area around $\theta = \pm 1$. At first, the error seemed to be the result of a numerical error. A convergence analysis was performed (Appendix A) to check this theory. However, increasing the resolution did not seem to decrease the error. A physical explanation can be found when looking at the phase plot of the pendulum model. Figure 5.13e shows the streamlines of the dynamical model for a constant input of $u = 25\,\mathrm{N}$. The unstable saddle point which is normally located at $\theta = 0$ (pendulum straight up) has shifted to the location of the discrepancy that is present in the PIMC method. This saddle point is especially sensitive to disturbances. When a disturbance pushes the pendulum in the wrong direction, it will not reach the target set within the time horizon. Figure 5.13f shows the quiver plot of the dynamic system when the optimal control policy is applied. It seems that the control policy as used in the probabilistic LS method provides a solution which increases the probability of reaching the target set.

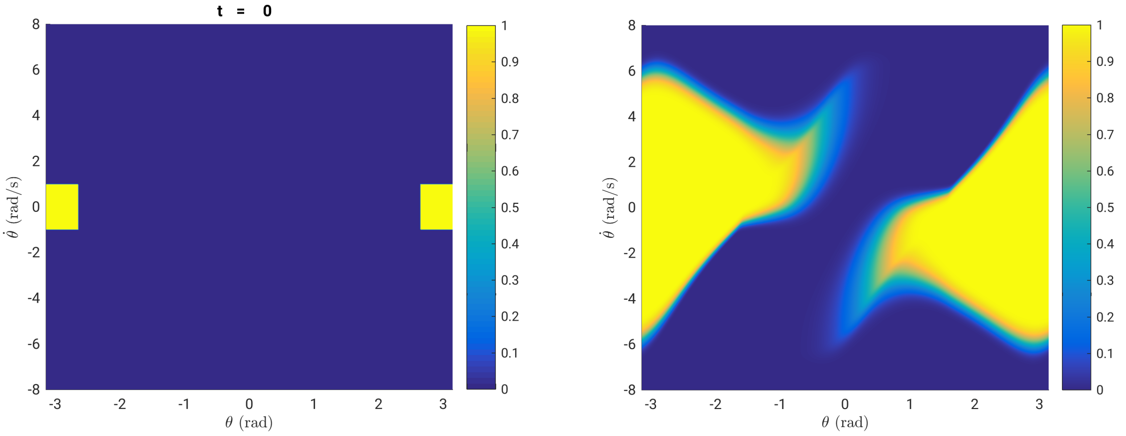

(a) Initial condition.

(b) LS Solution.

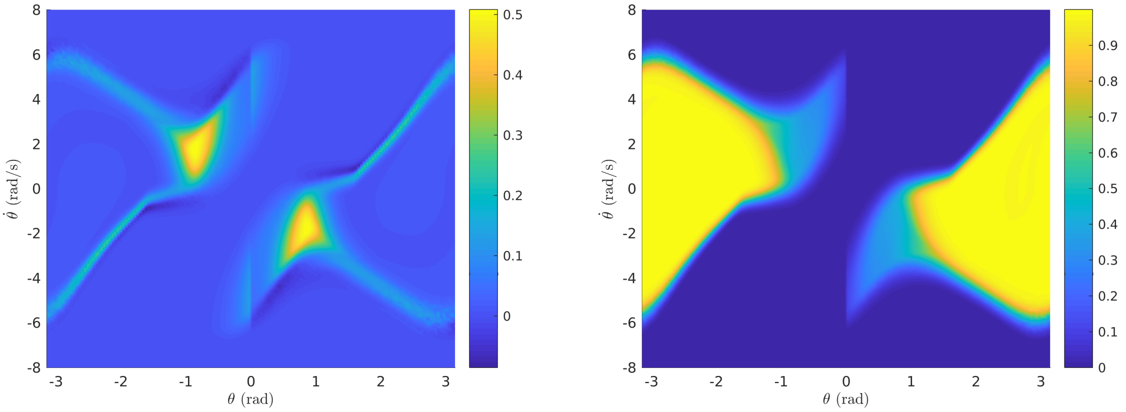

(c) Difference between the LS and PIMC result.

(d) PIMC solution.

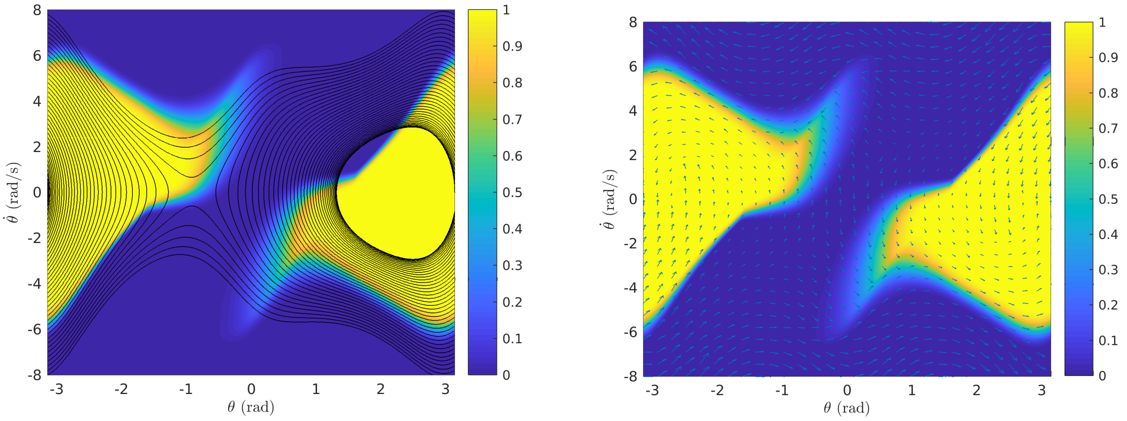

(e) Phase plot shows the streamlines of the pendulum model for a constant maximum control input.

(f) Quiver plot showing the flow-field when applying the optimal state feedback law.

Figure 5.13: Validation of the pendulum on a cart probabilistic reachability for $T = 1\,\mathrm{s}$.

To check that the optimal control policy is causing the difference between the two methods, a case is considered with only a disturbance acting on the model ($u = 0\,\text{N}$ and $\sigma = 10\,\text{N}$).

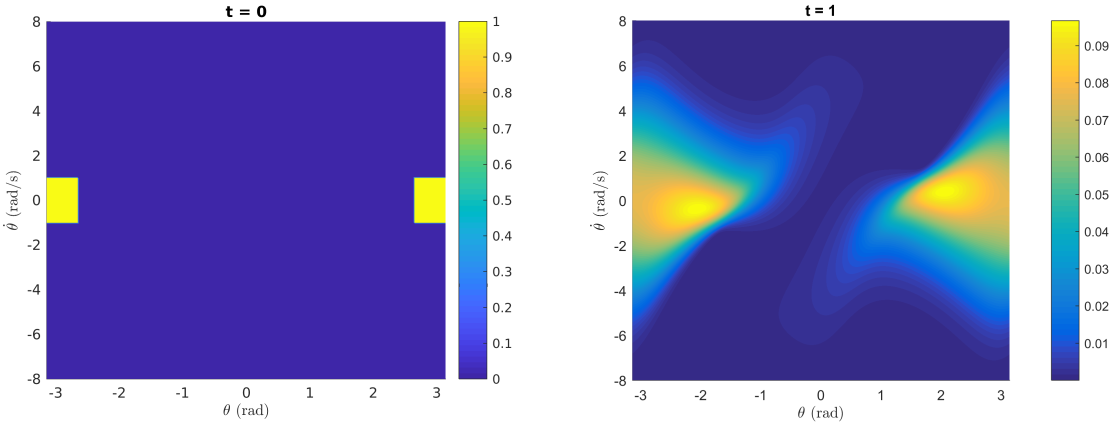

(a) Initial condition.

(b) LS solution.

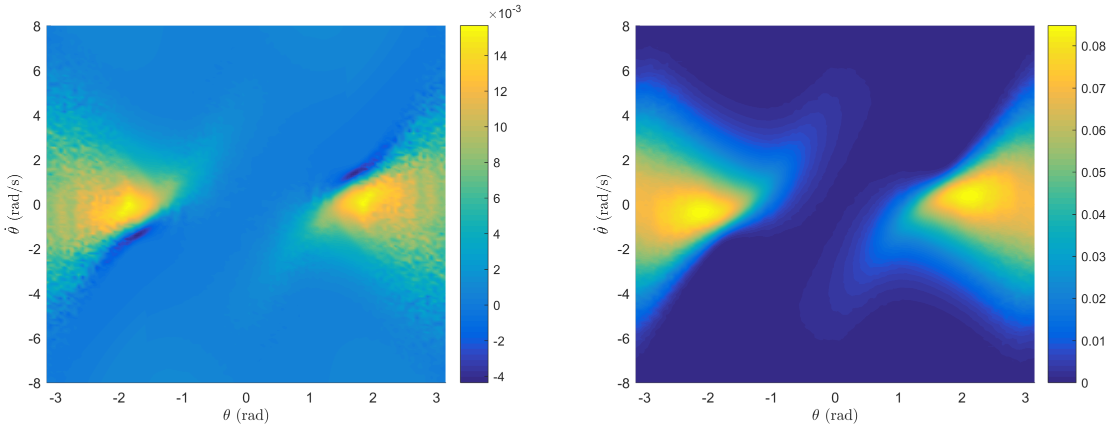

(c) Difference between the LS and PIMC result.

(d) PIMC solution.

Figure 5.14: Validation of the pendulum on a cart without control probabilistic reachability for $T = 1\,\text{s}$.

Figure 5.14 shows the solutions for this configuration. The initial set is centered at $\theta = \pm\pi$ as shown in Figure 5.14a. The LS solution on a 300 x 300 grid size is presented in Figure 5.14b. The PIMC solution was generated on a 150 x 150 grid. The result of this PIMC analysis can be seen in Figure 5.14d. Finally the difference between the LS and PIMC method is shown in Figure 5.14c. It is clear that without a control policy acting on the model, the LS and PIMC method produce the same solution.

6

Aircraft Model

In the previous chapters the theory of reachability analysis and probabilistic reachability analysis was introduced with the purpose of determining the SFE. This chapter will focus on the application of the theory on a simplified aircraft model. The dynamics of the aircraft model together with the optimal control strategy are introduced in Section 6.1. The deterministic SFE is presented in Section 6.2. A suitable turbulence model and the probabilistic SFE are introduced in Section 6.3 and 6.4 respectively. Validation of the probabilistic SFE is performed in Section 6.6. Finally, Section 6.5 will compare the difference between the deterministic and probabilistic SFE.

6.1. Dynamic Model

For the calculation of the SDE a nonlinear 3D simplified aircraft model will be considered [8]. The dynamics of the model are

$$\dot{V} = -\frac{\rho S}{2m}V^2 C_{D_0} - g\sin(\gamma) + \frac{T}{m} - \frac{\rho S}{2m}V^2(C_{D_\alpha}\alpha + C_{D_{\alpha^2}}\alpha^2) \tag{6.1}$$

$$\dot{\gamma} = -\frac{g}{V}\cos(\gamma) + \frac{\rho S}{2m}V(C_{L_0} + C_{L_\alpha}\alpha)\cos(\phi) - \frac{\rho S}{2m}VC_{Y_\beta}\beta\sin(\phi) \tag{6.2}$$

where $C_{L_\alpha} = 6.0723$, $C_{L_0} = 1.0656$, $C_{D_0} = 0.1599$, $C_{D_\alpha} = 0.5035$, $C_{D_{\alpha^2}} = 2.1175$, $m = 120\,000$ kg, $S = 260\,\text{m}^2$ and $\rho = 1.225\,\text{kg}\,\text{m}^{-3}$. The aerodynamic data is based on the Research Civil Aircraft Model (RCAM) model [32].

Both the thrust (T), the angle of attack (α) and the side-slip angle (β) are considered to be inputs to the system and are bounded. $\alpha \in [0;14.5]^\circ$, $\beta \in [-5,5]^\circ$ and $T \in [20\,546;410\,920]$ N. Furthermore, the roll angle is treated as a discretely gridded input.

The Hamiltonian of the system now becomes

$$\begin{aligned}\mathscr{H}(p,x,u) = p_1\left(-\frac{\rho S}{2m}V^2C_{D_0} - g\sin(\gamma) + \frac{T}{m} - \frac{\rho S}{2m}V^2(C_{D_\alpha}\alpha + C_{D_{\alpha^2}}\alpha^2)\right) + \\ p_2\left(-\frac{g}{V}\cos(\gamma) + \frac{\rho S}{2m}V(C_{L_0} + C_{L_\alpha}\alpha)\cos(\phi) - \frac{\rho S}{2m}VC_{Y_\beta}\beta\sin(\phi)\right)\end{aligned} \tag{6.3}$$

where $\mathbf{p}$ are the co-states of the value function: $p_1 = \frac{\partial\phi}{\partial V}$ and $p_2 = \frac{\partial\phi}{\partial\gamma}$. The inputs are linear for the angle of side-slip and thrust, and quadratic in angle of attack.

The maximizers for the Hamiltonian ($\hat{\alpha}$,$\hat{\beta}$ and $\hat{T}$) are

- If $p_1 > 0$ then $\hat{T} = T_{min}$ and
  - If $\hat{p} > \bar{\alpha}$ then $\hat{\alpha} = \alpha_{min}$
  - If $\hat{p} = \bar{\alpha}$ then $\hat{\alpha} = [\alpha_{min};\alpha_{max}]$
  - If $\hat{p} < \bar{\alpha}$ then $\hat{\alpha} = \alpha_{max}$
- If $p_1 = 0$ then $\hat{T} \in [T_{min};T_{max}]$ and

  - If $p_2 > 0$ then $\hat{\alpha} = \alpha_{min}$
  - If $p_2 = 0$ then $\hat{\alpha} \in [\alpha_{min}; \alpha_{max}]$
  - If $p_2 < 0$ then $\hat{\alpha} = \alpha_{max}$

- If $p_1 < 0$ then $\hat{T} = T_{max}$ and
  - If $\hat{p} \leq \alpha_{min}$ then $\hat{\alpha} = \alpha_{min}$
  - If $\alpha_{min} \leq \hat{p} \leq \alpha_{max}$ then $\hat{\alpha} = \hat{p}$
  - If $\hat{p} \geq \alpha_{max}$ then $\hat{\alpha} = \alpha_{max}$

besides, for the side-slip angle β holds

- If $p_2 \sin\phi > 0$ then $\hat{\beta} = \beta_{min}$
- If $p_2 \sin\phi = 0$ then $\hat{\beta} \in [\beta_{min}; \beta_{max}]$
- If $p_2 \sin\phi < 0$ then $\hat{\beta} = \beta_{max}$

where $\hat{p} = \frac{p_2 C_{L_\alpha} \cos\phi - p_1 V C_{D_\alpha}}{2 p_1 V C_{D_{\alpha^2}}}$ and $\bar{\alpha} = \frac{\alpha_{min} + \alpha_{max}}{2}$

6.2. Safe Flight Envelope

The SFE as introduced in Section 3.3 is defined by the intersection of the forwards and backwards reachable set. Figure 6.1 shows the result of the reachability analysis starting from an initial set $V \in [55; 85]$ m s^{-1} and $\gamma \in [-10; 10]°$ over a time horizon of 2 s. Figure 6.1a shows both the forwards (green) and backwards (red) reachable sets. The SFE is then computed by taking the intersection of both flight envelopes. Using the LS method, this is achieved by taking the maximum of both sets. The result is shown in Figure 6.1b. It can be seen that as the roll angle increases, the SFE becomes smaller. A similar trend can be seen for the flight path angle γ.

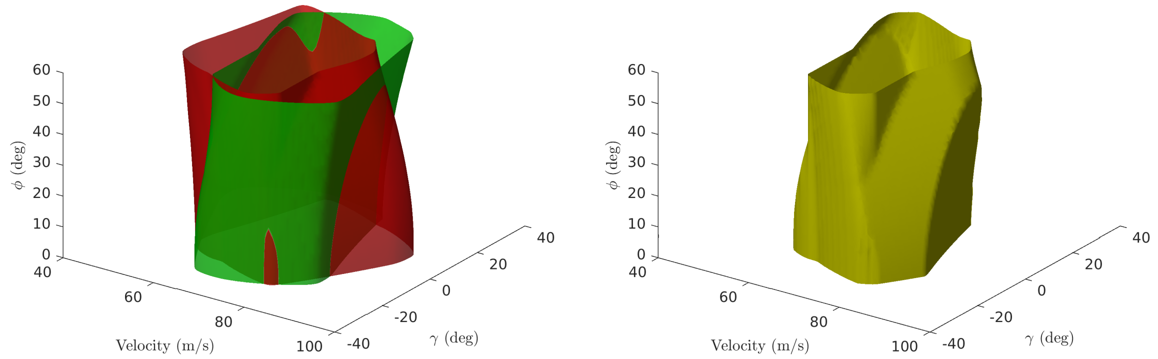

(a) Forwards (green) and Backwards (red) reachable set for the simplified aircraft model.

(b) Intersection between the forwards and backwards reachable set.

Figure 6.1: SFE of the simplified aircraft model for a 2 s time horizon.

6.3. Turbulence Model

In order to model the effect of disturbances on the SFE, a realistic turbulence model is required. A common way to describe atmospheric turbulence is given by the Dryden spectra [33]. This Dryden spectra is based on the assumptions that atmospheric turbulence is:

- Gaussian.
- Stationary and frozen.

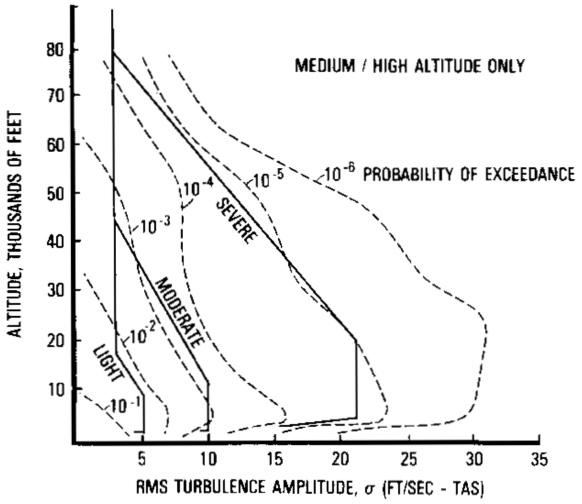

(a) Turbulence intensity as a function of altitude [34].

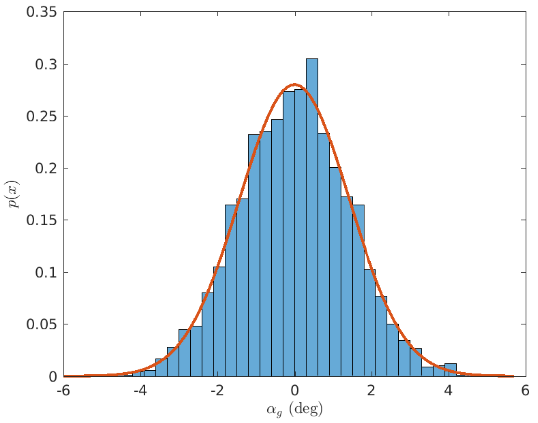

(b) Histogram of the gust angle of attack (blue); Fitted Gaussian distribution (red).

- Homogeneous along the flight path.
- Isotropic - Statistical properties are independent of the axis orientation.

These assumptions are met for high altitude turbulence when the velocity of the aircraft is much higher than the air particles. The Dryden model is implemented as a white noise shaping filter. The white noise shaping filter in state-space representation is given by

$$\begin{bmatrix} \dot{w}_g(t) \\ \dot{w}_g^*(t) \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -\frac{V^2}{L_g^2} & -2\frac{V}{L_g} \end{bmatrix} \begin{bmatrix} w_g(t) \\ w_g^*(t) \end{bmatrix} + \begin{bmatrix} \sigma\sqrt{\frac{3V}{L_g}} \\ (1-2\sqrt{3})\sigma\sqrt{\left(\frac{V}{L_g}\right)^3} \end{bmatrix} w(t) \tag{6.4}$$

with L_g the turbulence scale length and σ the turbulence intensity. Then, the gust angle of attack is calculated as

$$\alpha_{total} = \alpha + \alpha_g \tag{6.5}$$

$$\alpha_g = \tan^{-1}\left(\frac{w_g}{V}\right) \approx \frac{w_g}{V} \tag{6.6}$$

For medium to high altitude turbulence, the turbulence scale length is chosen to be 500 m. Figure 6.2a shows the turbulence intensity as a function of altitude. For moderate high altitude turbulence a standard deviation of $2\,\mathrm{m\,s^{-1}}$ is selected.

Since the probabilistic reachability analysis uses Gaussian white noise inputs, the implementation of a white noise shaping filter would add an additional two dimensions to the problem, thus significantly increasing the computational cost. However, since the interest is in finding the PDF instead of the exact time response, it is possible the use the fact that the distribution of a Gaussian process remains Gaussian when integrating. As explained in Chapter 5, integrated white noise has a Gaussian distribution whose standard deviation increases with time (i.e. a random walk). As it was assumed that the atmospheric turbulence is stationary, the stationary solution for the gust angle of attack can be found. Furthermore, when applying the shaping filter to the Gaussian white noise input signal, the high frequencies are filtered out, thus changing the Power spectral density. The PDF however, remains Gaussian. Since it is complex to analytically integrate the SDE associated with the Dryden spectra, a numerical approximation is made. Figure 6.2b shows the histogram of the gust angle of attack as generated by the Dryden model. The red line represents a fitted Gaussian function ($\mu = 0.0°$ and $\sigma = 1.4245°$).

6.4. Probabilistic Safe Flight Envelope

For the probabilistic SFE the same model will be considered as for the deterministic SFE with the addition of a disturbance on the angle of attack (i.e. turbulence). The noise statistics as determined in Section 6.3 will be applied to the model.

The dynamical model for the simplified aircraft now becomes

$$\dot{V} = -\frac{\rho S}{2m}V^2 C_{D_0} - g\sin(\gamma) + \frac{T}{m} - \frac{\rho S}{2m}V^2(C_{D_\alpha}\alpha + C_{D_{\alpha^2}}\alpha^2) + W\frac{\rho S}{2m}V^2 C_{D_\alpha} \tag{6.7}$$

$$\dot{\gamma} = -\frac{g}{V}\cos(\gamma) + \frac{\rho S}{2m}V(C_{L_0} + C_{L_\alpha}\alpha)\cos(\phi) - \frac{\rho S}{2m}VC_{Y_\beta}\beta\sin(\phi) + W\frac{\rho S}{2m}VC_{L_\alpha} \tag{6.8}$$

Where W is the white noise process. It should be noted that the noise terms with α^2 are left out in order to keep the control independent of the noise.

The result of the probabilistic reachability analysis is shown in Figure 6.3. As with the deterministic reachability analysis, the analysis was performed from an initial set $V \in [55;85]\ \mathrm{m\,s^{-1}}$ and $\gamma \in [-10;10]°$ over a time horizon of 2 s. The initial LS was represented by an indicator function of the initial set. Figure 6.3a shows the $1\,\sigma$ isosurfaces for the forwards and backwards reachable set in green and red respectively. The $1\,\sigma$ confidence interval corresponds to all states with a 68.3 % reachability probability or higher. For the probabilistic reachability analysis, the SFE cannot be computed as the intersection of the two sets anymore. Instead, it is defined as the joint probability of the two sets (i.e. the probability to transition to a state and transition back to the initial set). For two independent processes this is equal to the product of the forwards transition probability with the backwards transition probability. The $1\,\sigma$ isosurface of the probabilistic SFE is shown in Figure 6.3b. Finally, Figure 6.3c shows slices of the probabilistic SFE that give a better insight in how the probability changes with the states.

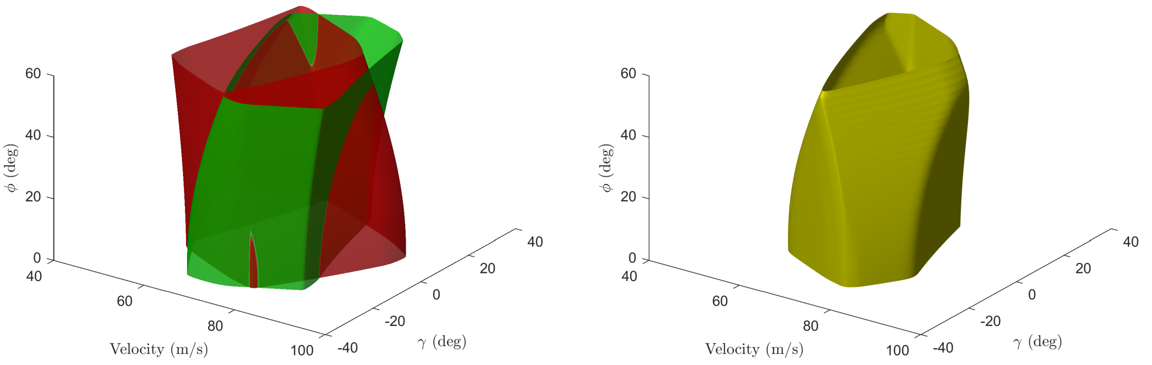

(a) Forwards (green) and Backwards (red) $1\,\sigma$ isosurface for the simplified aircraft model.

(b) Probabilistic SFE $1\,\sigma$ isosurface.

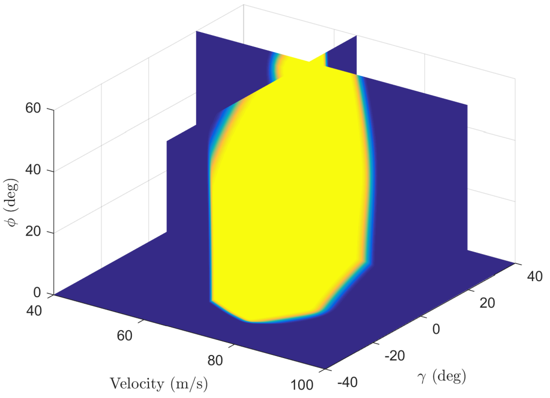

(c) Slice of the probabilistic SFE.

Figure 6.3: Probabilistic SFE of the simplified aircraft model for a 2 s time horizon.

Table 6.1: Size comparison between the deterministic and the probabilistic SFE.

Roll Angle/area	Deterministic	1σ	2σ	3σ
0°	100 %	94.3 %	81.9 %	71.4 %
30°	100 %	94.9 %	82.8 %	72.6 %
60°	100 %	78.0 %	62.8 %	50.8 %

6.5. Sensitivity

The objective of this research is to investigate the effect of disturbances on the SFE. By comparing the deterministic SFE with the confidence contours of the probabilistic SFE, it is possible to investigate the sensitivity. Since it is hard to compare the 3D reachable sets side-by-side, two slices have been made at roll angles of 0° and 60°.

Figure 6.4 shows both the deterministic SFE, as well as the probabilistic SFE. For the probabilistic SFE, the one, two and three σ isocontours are also plotted. The one, two and three σ isocontours represent the states in which a transition probability is achieved of at least 68.3 %, 95.5 % and 99.7 % respectively. Figure 6.4a shows the results for a roll angle of 0°. The outer contour represents the deterministic SFE. It can be seen that the 1σ contour, which has a probability of 68.3 % is roughly the same size (slightly smaller) than the deterministic contour. This means that when the deterministic SFE is used for the purpose of LOC prevention in a setting with turbulence, most states on the deterministic contour have a 31.7 % probability that the aircraft cannot return to the initial set. The probabilistic SFE can be defined for an arbitrary probability as chosen by the operator to ensure safety of the aircraft and occupants. The 3σ contour is more suitable for preventing LOC as it gives a high probability of returning to the initial set, even in the presence of atmospheric turbulence. The difference between the deterministic and probabilistic SFE becomes more pronounced at higher roll angles. Figure 6.4b shows the same information but now for a roll angle of 60°. It can be seen that the effect of disturbances on the SFE is greater. A fairly large part of the deterministic SFE has a transition probability of 0 %, which means that it is impossible for an aircraft which enters this part of the SFE to return to the initial set.

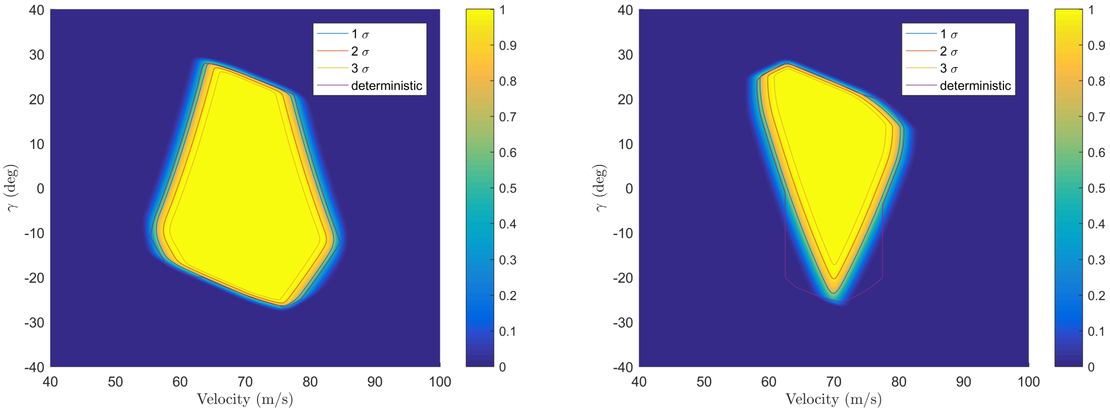

(a) SFE for a roll angle of 0°

(b) SFE for a roll angle of 60°

Figure 6.4: Sensitivity of the deterministic and stochastic SFE to the effect of turbulence.

To get some more insight in the size difference of the deterministic compared to the probabilistic SFE, the relative area between the deterministic set and the sigma contours is shown in Table 6.1. From this table it becomes clear that the disturbances have a greater effect on the SFE when the roll angle increases. For example, at a 60° roll angle, only 50 % of the deterministic SFE is enclosed by the 3σ contour.

6.6. Validation

To check that the aircraft model was correctly implemented in the LS toolbox, a validation is performed for the simplified aircraft model. The validation is performed in a similar fashion as explain in Section 5.8. In this validation the backwards probabilistic reachability for two cases is considered, a roll angle of 0° and 60°. The reachability is calculated over a time horizon of 2 s, starting from an initial set as represented by the black rectangle. Figure 6.5a, 6.5c and 6.5e show the LS solution, PIMC solution and the difference respectively for a roll angle of 0°. Furthermore, Figure 6.5b, 6.5d and 6.5f show the LS solution, PIMC solution and the difference respectively for a roll angle of 60°.

From Figure 6.5 it becomes clear that the LS solution is similar to the PIMC method. When looking at the magnitude of the transition probability, a slight difference can be seen in the corners of the PIMC method, However, the PIMC solution always produces a similar or lower estimation of the transition probability, as expected. This difference is again linked to the underlying optimal control policy. Furthermore, the direction of the uncertainty seems to be focused in the same direction for both methods.

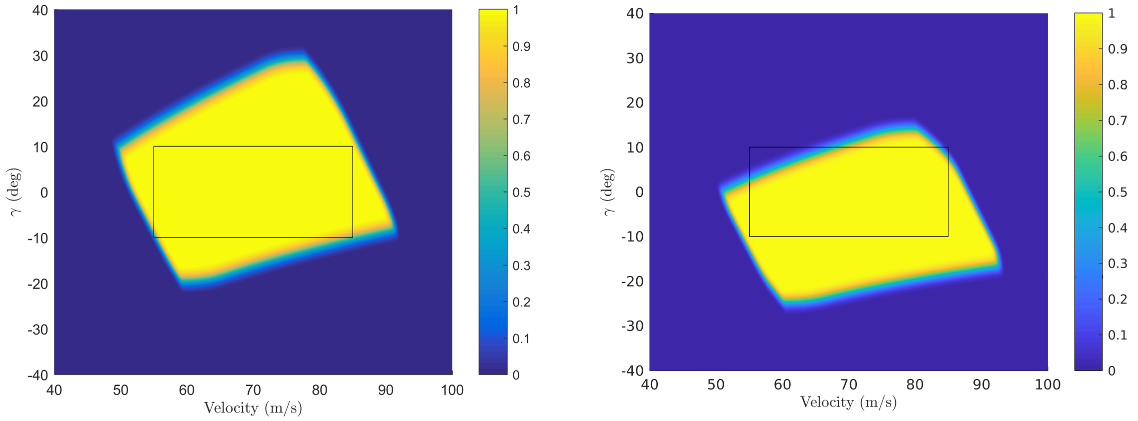

(a) LS solution for $\phi = 0°$.

(b) LS solution for $\phi = 60°$.

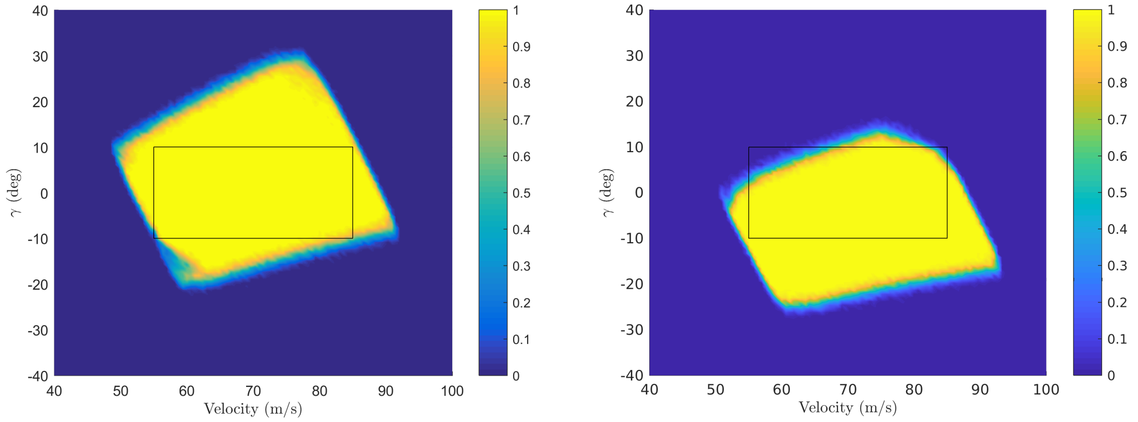

(c) PIMC solution for $\phi = 0°$.

(d) PIMC solution for $\phi = 60°$.

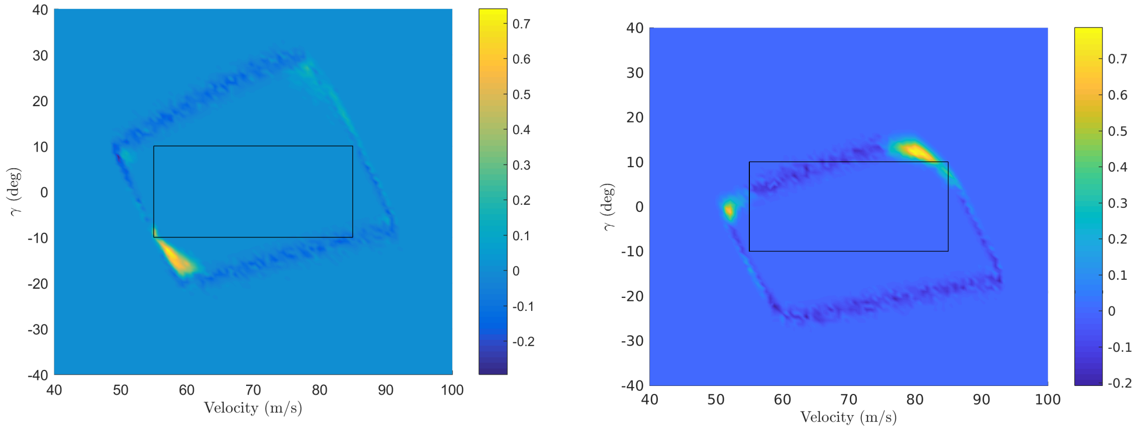

(e) Difference between the PIMC and the LS method for $\phi = 0°$.

(f) Difference between the PIMC and the LS method $\phi = 60°$.

Figure 6.5: Validation of the backwards reachable set for the simplified aircraft model for $T = 2\,\text{s}$.

7

Quadrotor Model

In resent years, Unmanned Aerial Vehicles (UAVs) are strongly gaining popularity. One of the most common variants is the quadrotor. These UAVs are often operated in and above densely populated areas which could be a safety concern. The application of envelope protection could contribute to safer operation of quadrotors by preventing envelope violation and LOC. The aim of this chapter is to investigate the SFE of a real life quadrotor. Section 7.1 introduces the Equations of Motions (EOMs) and the aerodynamic model for a Parrot Bebop quadrotor. The SFE of the quadrotor is presented in Section 7.2. The stochastic counterpart of the SFE is discussed in Section 7.3. Finally, a validation is performed in Section 7.4.

7.1. Dynamic Model

The Parrot Bebop is a small quadrotor that is commercially available. A schematic drawing including the axis system is shown in Figure 7.1. The longitudinal dynamics of the Parrot Bebop quadrotor are represented by a three Degrees-Of-Freedom (DOF) model. Equation (7.1-7.4) represent the EOM in a body fixed reference frame

$$\dot{u} = \frac{F_x(\delta_T, \delta_q, u, w)}{m} - qw - g\sin\theta \tag{7.1}$$

$$\dot{w} = \frac{F_z(\delta_T, \delta_q, u, w)}{m} + qu + g\cos\theta \tag{7.2}$$

$$\dot{q} = \frac{M_y(\delta_T, \delta_q, u, w)}{I_{yy}} \tag{7.3}$$

$$\dot{\theta} = q \tag{7.4}$$

where F_x, F_z and M_y are the aerodynamic forces and moments. The aerodynamics of the Parrot Bebop quadrotor are represented by a non-linear parametrized model that is identified using flight data and windtunnel testing [35]. The aerodynamic model is a function of the current body velocity and the engine settings. By changing the engine settings individually, it is possible to generate roll, pitch, yaw and thrust commands. For example, creating a thrust difference between the front and rear engines will result in a pitching moment. The

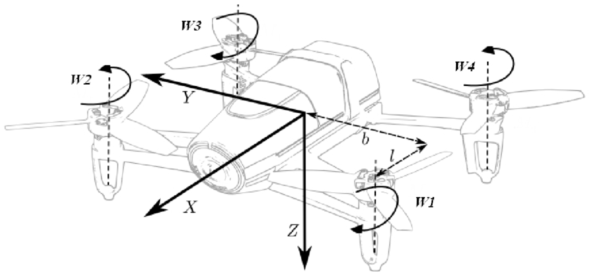

Figure 7.1: Schematic of the Parrot Bebop axis system [35].

mapping from individual engine settings to virtual control is given by the mixing logic

$$\begin{bmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & 1 & -1 & -1 \end{bmatrix} \begin{bmatrix} \delta_T \\ \delta_p \\ \delta_q \\ \delta_r \end{bmatrix} \tag{7.5}$$

where w_1 till w_4 are the engine settings and δ_T, δ_p, δ_q, δ_r are the thrust, roll, pitch and yaw commands respectively. The output of the mixing logic is unconstrained, meaning that the output could include engine commands greater than 1. Direct control allocation is applied to constrain the engine settings. Direct allocation methods scale the commands such that the resulting control allocation is projected on the boundary of the set of admissible control inputs. Since only the longitudinal motion is considered, the virtual control for roll and yaw will be set to zero.

The Hamiltonian of the system now becomes

$$\begin{aligned} \mathscr{H}(p,x,\delta_T,\delta_q) = \max_{\delta_T,\delta_q} p_4 q + p_2\left(\frac{F_z(\delta_T,\delta_q,u,w)}{m} + qu + g\cos\theta\right) - \\ p_1\left(qw - \frac{F_x(\delta_T,\delta_q,u,w)}{m} + g\sin\theta\right) + p_3\frac{M_y(\delta_T,\delta_q,u,w)}{I_{yy}} \end{aligned} \tag{7.6}$$

Unlike the previous examples, the policy that maximizes the Hamiltonian cannot be expressed analytically because of the non-linear aerodynamic model. To circumvent this issue, the Hamiltonian is calculated for a range of different virtual thrust and pitch combinations. Then, for each state the maximum value of the Hamiltonian is selected from the range of combinations.

7.2. Safe Flight Envelope

The reachability analysis is started from an initial set, represented by a 4D hypercube centered at the origin with $u \in [-1;1]$ m s^{-1}; $w \in [-1;1]$ m s^{-1}; $q \in [-5;5]$ °s^{-1} and $\theta \in [-1;1]$°. Ideally, the initial set would be represented by a point in the origin (i.e. when the quadrotor is hovering), however the LS method requires that the initial condition is represented by an implicit set. Therefore, the initial set is chosen to be a small hypercube around the origin. The size of the initial set can be decreased to approximate a point in the origin if required. The flight envelope was computed over a time horizon of 0.1 s. The time horizon is kept small for several reasons. First, a longer time horizon would increase the required computational time to impractical levels. Furthermore, since the dynamics of the quadrotor are inherently unstable, very large deflections can be achieved in a short amount of time. This is made worse by the fact that the control torques generated by the engines are very large in comparison to the inertia of the quadrotor, which allows for very large rotational accelerations. When a longer time horizon is selected, the reachability analysis would reach parts of the state-space that are not identified in the aerodynamic model, possibly leading to erroneous results. To further limit this effect, the maximum pitch command is constraint to only 10 % of the maximum. Finally, due to limitations in computational resources, the size of the state-space that can be discretized while keeping the resolution to reasonable levels is limited. A smaller time-horizon allows for a smaller domain, especially for the pitch rate. When simulating for larger time horizons, or using larger pitch commands, it is recommended to extend the pitch rate domain to $[0;2\pi]$rad s^{-1} using a periodic boundary condition. This would allow the quadrotor to make a full rotation around the pitch axis. Figure 7.2 shows both the forwards and backwards reachable set. The LS solution was generated on a 45 x 45 x 80 x 45 grid size. The computational domain was equal to $[-1.5;1.5]$ m s^{-1} x $[-3;3]$ m s^{-1} x $[-90;90]$ °s^{-1} x $[-6;6]$°. The current computational resources limits the use of a denser grid. Since the reachable set is four dimensional, a slice of the 4D implicit surface is created which allows it to be visualized in 3D. Figure 7.2 shows the SFE for a zero pitch rate slice ($q = 0$°s^{-1}). The forwards and backwards reachable sets are shown in Figure 7.2a. It is interesting to note that the forwards reachable set can achieve lower vertical velocities than the backwards reachable set. This is due to the maximum acceleration that can be generated by the quadrotor. The maximum upwards acceleration is almost twice that of the downwards acceleration (only gravity). The SFE is defined as the intersection between the forwards and backwards reachable set and is shown in Figure 7.2b. From the SFE it can be seen that the reachable vertical velocity decreases with an increasing pitch angle. This is caused by the engine control allocation. In order to generate a pitch moment, a thrust difference needs to be created between the front and rear engines. When applying full thrust, this difference can only be achieved by decreasing the thrust on the engines. Therefore, a pitch command lowers the maximum thrust. Similar, a pitch

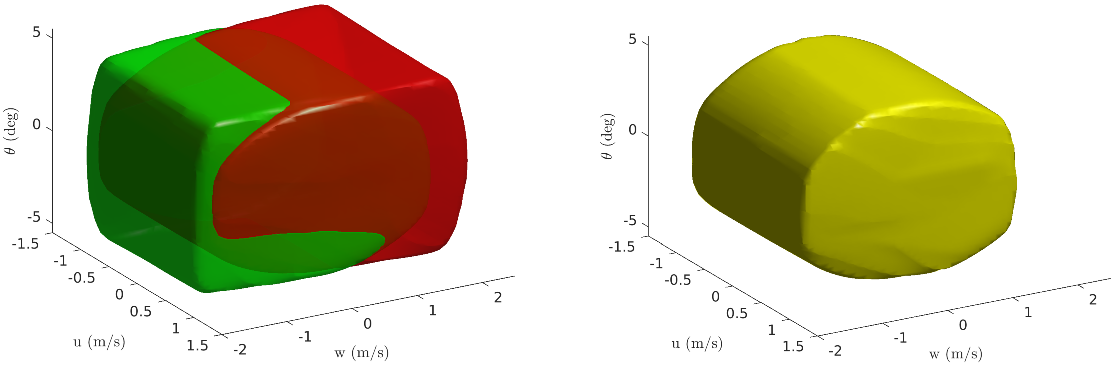

(a) Forwards (green) and Backwards (red) reachable set.

(b) Intersection between the forwards and backwards reachable set.

Figure 7.2: SFE of the quadrotor model for $q = 0°\mathrm{s}^{-1}$ and $T = 0.1\,\mathrm{s}$.

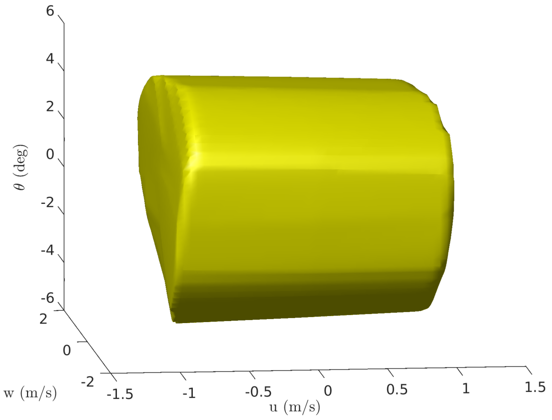

Figure 7.3: Relation of the SFE between the pitch angle and longitudinal velocity.

command would add some thrust to the minimum engine settings, thus decreasing the minimum reachable vertical velocity. Figure 7.3 shows the SFE from a different angle. It becomes clear that the longitudinal velocity is related to the pitch angle, which is expected as the engines need to be tilted forward to achieve thrust in the horizontal direction. A positive pitch angle relates to a lower velocity (flying backwards), while a negative pitch angle relates to a higher forward velocity.

A different slice is made to see how the SFE changes with pitch rate. Figure 7.4 shows a slice for $\theta = 0°$. The forwards and backwards reachable sets are shown in Figure 7.4a and the SFE is shown in Figure 7.4b. The slice for a constant pitch angle shows similar trends as seen in the slice for a constant pitch rate.

To circumvent the constraints posed by the limited state-space, it is recommended to perform the reachability analysis on the closed loop dynamics (i.e. the dynamical model that includes the controller logic) instead of the unstable open loop dynamics. Unfortunately, such a model is not available at the current time. Other techniques to improve the computational load include time-scale separation [36]. This technique separates the fast dynamics from the slower dynamics in order to circumvent the curse of dimensionality. For such an approach to be feasible, the slow dynamics (e.g. velocity) cannot be dependent on the fast dynamics (e.g. angular rates). In the case of the quadrotor model, both the vertical and the horizontal velocity are directly dependent upon the angular rates, making this approach infeasible for high angular rates. Other limitations are that the current model assumes that the engine commands are instantaneous which causes the flight envelope to be over estimated, as the rotors require some time to spin-up. Since the reachability analysis assumes that the process is Markov (i.e. not dependent on the history) is generally not possible to include these spin-up dynamics. One solution would be to add a state for the engine rotational rate. Then, instead of

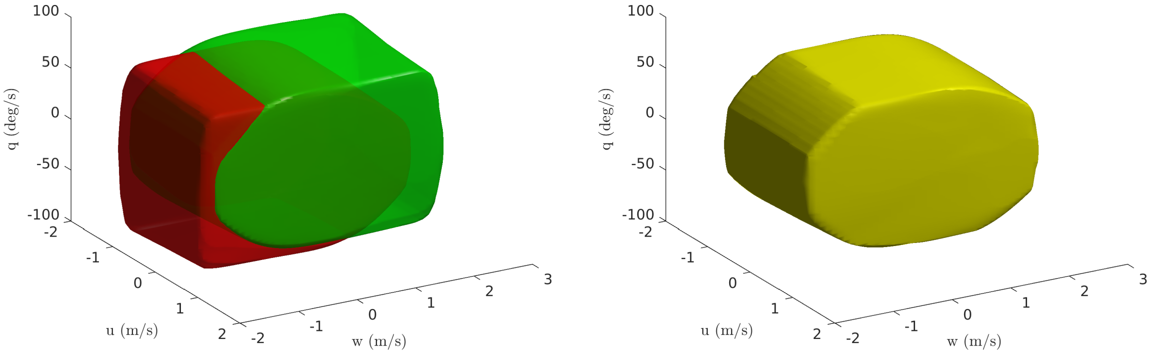

(a) Forwards (green) and Backwards (red) reachable set.

(b) Intersection between the forwards and backwards reachable set.

Figure 7.4: SFE of the quadrotor model for $\theta = 0°$ and $T = 0.1\,\mathrm{s}$.

an engine command, the input would be replaced by an engine torque setting. However, this approach would add more states to the model making it practically infeasible. For the 3 DOF model, two additional states are required (one for each set of engines). A full 6 DOF model would even require four additional states.

7.3. Probabilistic Safe Flight Envelope

In order to model the effect of turbulence on the SFE, the reachability analysis will have to be performed using a suitable turbulence model. In Section 6.3 the effect of the turbulence was modeled using a Dryden spectra that was added as a disturbance to the angle of attack of the simplified airplane model. The increased complexity of the quadrotor model poses some challenges if such an approach is taken. This section will discuss these challenges.

The Dryden model assumes a frozen wind field which is a valid assumption if the speed of the aircraft is much higher than that of the air particles. For the simplified aircraft model treated in Section 6.3, this assumption was not violated. However, the quadrotor operates in a region where the velocity is low, or even zero. Therefore, the Dryden spectra is not suitable for application to the quadrotor model. Furthermore, due to the non-linear nature of the aerodynamic model, it is not possible to separate the disturbance term from the drift term, as required for the stochastic reachability analysis. A proposed solution is to locally linearize the model in order to calculate the gust derivatives for the quadrotor model. However, the current aerodynamic model does not include these gust derivatives. Finally, the required computational cost to compute the probabilistic SFE is even higher than for the deterministic SFE, which was already on the limit of available computational resources. Since the indicator function used in the probabilistic reachability analysis is prone to round off errors, a higher resolution is required. Furthermore, the calculation of the second derivative requires additional computational resources and more memory. A more detailed analysis of the required computational resources is given in Section 8.3.

The limitations on the probabilistic reachability analysis posed by the quadrotor model are currently outside of the scope of this research. However, to get a rough idea of the probabilistic SFE for the quadrotor model, a low resolution analysis using an arbitrary disturbance is considered in Appendix B.

7.4. Validation

In this section the results of the reachability analysis are validated using a MC approach. The validation is performed for both the forwards and backwards reachable set separately. In both cases the initial set is the same as presented in Section 7.2. The MC analysis was performed on a 5 x 5 x 5 x 5 grid for the initial condition and 15625 different control sequences per initial condition for a total of 9765625 simulations. The MC analysis uses (semi-) random control sequences and is therefore independent of the optimal control formulation as used by the LS method. If the quadrotor model, or optimal controller implementation would contain errors, the LS method would yield a smaller reachable set. Figure 7.5 shows the validation result for both the forwards

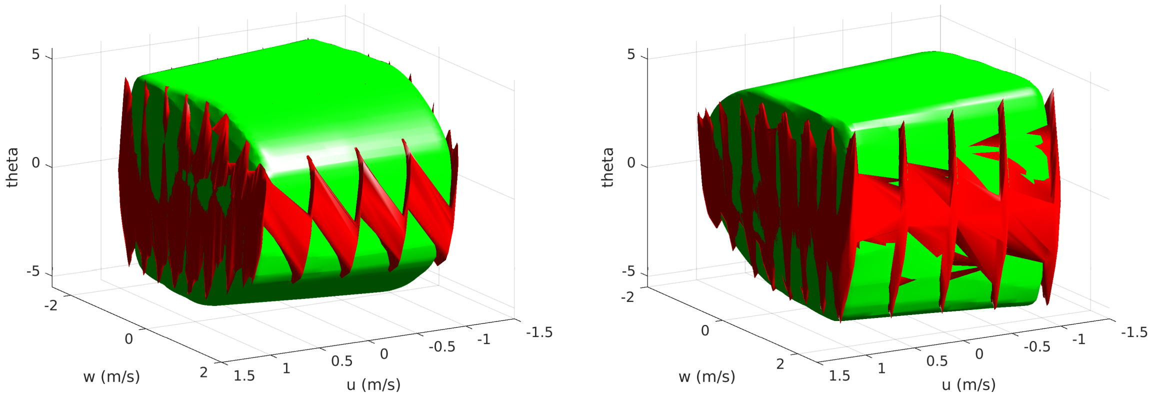

(a) Forwards reachable set.

(b) Backwards reachable set.

Figure 7.5: Validation of the quadrotor model (LS green; MC red) for $q = 0\,\mathrm{rad}$ and $T = 0.25\,\mathrm{s}$.

and backwards reachable set. From the MC validation it becomes clear that the LS and MC method produce roughly the same reachable set. The result obtained from the MC analysis might seem ’spiky’ which is caused by the low amount of initial conditions. In fact, five spikes can be seen for each dimension corresponding to the locations of the initial conditions. Increasing the number of initial conditions, or relaxing the criteria for the boundary calculation would resolve this discrepancy, however this is left out due to time constraints and computational limitations.

Under normal circumstances the MC set should not exceed the LS reachable set. From Figure 7.5 it can be seen that this is generally true. However, at the corners of the LS solution, the MC result extends slightly further. This is likely due to a rounding effect caused by the implicit surface. It is expected that as the resolution of the grid is increased, this effect will diminish.

8

Computational Efficiency

Previous chapters have introduced techniques for determining the deterministic and probabilistic SFE. The LS method which was used to solve this problem suffers from the 'curse of dimensionality'. In essence this means that the required computational effort increases exponentially with the dimension of the state-space. For simple problems this is generally not a big issue. However, when considering a more complex model such as the quadrotor model as introduced in Chapter 7, this quickly becomes the limiting factor. Although less computational intensive, the PIMC method also suffers from the same scaling issues. Next to limiting the analysis to simple models, the high computational resources also makes these methods unsuitable for real-time purposes (i.e. online computation of the SFE). Some solutions exist where the SFE is computed offline and stored in a database. Although this can facilitate real-time usage, it does not reduce the computational load, but merely shifts the problem. This chapter will focus on reducing the computational complexity of the probabilistic SFE computation. Section 8.1 introduces a method to reduce the computational complexity of the PIMC method. Furthermore, Section 8.2 makes an effort to break the 'curse of dimensionality'. Finally, an analysis of the required computational time is made in Section 8.3.

8.1. Covariance Prediction

The PIMC method was introduced in Section 5.6 for the purpose of validating the probabilistic reachable sets. This method uses an arbitrary number of samples to estimate the probability distribution after a small time step. This process is repeated for each node in the discretized state-space. As this method does not make any assumptions about the nature and distribution of the noise, it is very powerful. Depending on the non-linearity of the model and the size of the time-step, a large number of samples may be required to correctly capture the probability distribution. While some techniques like Sobol (also known as Quasi-Monte Carlo) sampling or kernel density estimation could reduce the number of samples required to capture the process variance, it is possible to further reduce the amount of samples. It can be observed that if the time step τ is small enough, the process seems to adhere to a Gaussian distribution. Effectively, this is equal to linearizing the model around the current state. By using this observation, the required number of samples can be reduced to only one.

Since the SDEs governing the probabilistic reachability analysis can be thought of having an expected value that is equal to the drift, the mean can be calculated using a single particle without noise. Then, under the assumption that the process behaves according to a Gaussian distribution, the probability distribution is calculated using the covariance estimation. The one-step ahead covariance estimate is equal to

$$p_{k+1} = \Phi_{k+1} P_k \Phi_{k+1}^T + \Gamma_{k+1} Q_k \Gamma_{k+1}^T \tag{8.1}$$

with Φ the system transition matrix, p the covariance matrix, Γ the input distribution matrix and Q the system noise covariance.

Figure 8.1 illustrates this process for a double integrator system. The red line represents the mean trajectory of the system and is represented by a single particle. The ellipses represent the 1σ confidence interval as calculated using the covariance prediction. It can be observed that the uncertainty increases as time marches forwards. The gray lines represent unique realizations of the stochastic process but are not used for the covariance prediction. The behavior of the samples is captured by the mean and covariance of the stochastic process. Because the double integrator system is linear, the assumption that the process behaves

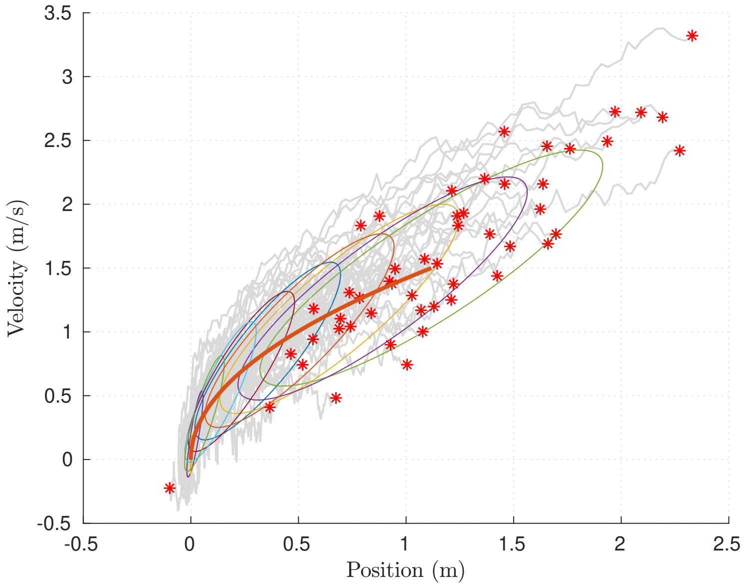

Figure 8.1: Covariance estimation for the double integrator system.

Gaussian is valid, even for a large time-step. Further computational effort can be reduced for systems with time independent dynamics and noise. When the time independence is present, the mean and covariance prediction can be reused for each time-step, only updating the transition probability on each step. Although the covariance prediction method reduces the computational load, it does not break the ‘curse of dimensionality’ since the state-space still has to be discretized.

8.1.1. Examples

To study how the assumptions of the covariance prediction influence the results of the probabilistic reachable set, two examples are considered. The first example is a double integrator system. It is similar to the validation performed in Section 5.8.1, but now the PIMC method with covariance prediction is used. Again, the initial set is a circle centered at the origin with unit radius. Since the double integrator model is linear, the covariance prediction should yield exactly the same result as the normal PIMC method.

Figure 8.2 shows the forwards probabilistic reachable set for the double integrator model. The result as obtained by the covariance prediction PIMC approach is shown in Figure 8.2a. The difference with the LS method is shown in Figure 8.2b. When comparing the PIMC method with the covariance prediction variant, it becomes clear that the result is very similar.

The second example that is considered is the pendulum on a cart model as introduced in Section 5.8.2. Since this model is not linear, the probabilistic reachable set will only be approximated correctly if the time step is small enough so that the process remains Gaussian. Figure 8.3 shows the forwards probabilistic reachable set for the pendulum on a cart model. The result from the covariance prediction PIMC method is shown in Figure 8.3a. The difference with the LS method is shown in Figure 8.3b. The biggest difference can be found at the $\theta = 0\,\mathrm{rad}$ line. This difference is caused by the fact that the computational domain was slightly too small. To prevent handling of the periodic boundary, the simulation was performed on a domain from 0 to $-2\pi\,\mathrm{rad}$, then split to make the comparison easier. This method places the edge of the domain on the $\theta = 0\,\mathrm{rad}$ line. Next to this, some slight differences can be observed around the edges of the probabilistic reachable set. Furthermore, the discrepancy that is caused by the underlying optimal control policy is still present, similar to the regular PIMC method.

8.2. Quadrature Point

The covariance prediction method, although being faster than the PIMC method, still requires discretization of the state-space. This means that both methods still suffer from the ‘curse of dimensionality’. In an attempt to circumvent this scaling issue, an approximate technique is introduced which is able to provide a rough

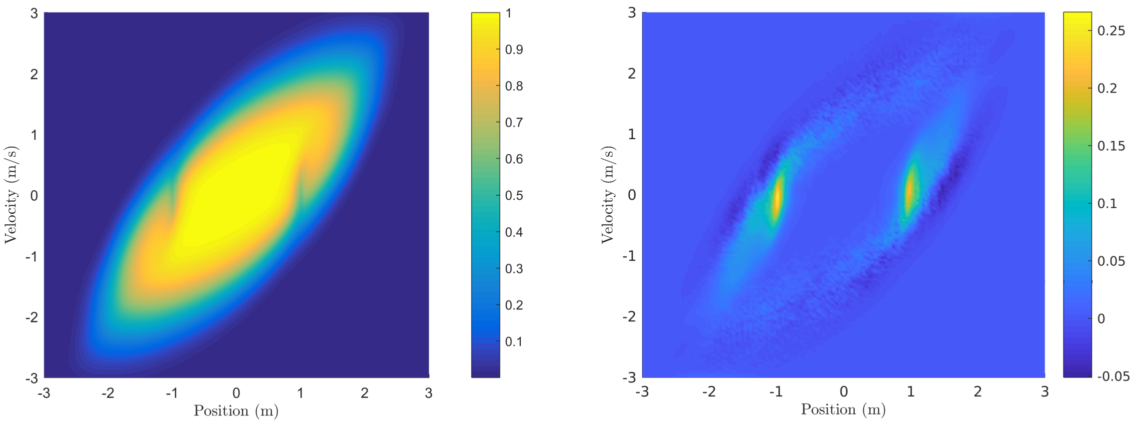

(a) Covariance prediction PIMC method.

(b) Difference between the covariance prediction PIMC method and the LS method.

Figure 8.2: Double integrator covariance prediction PIMC method.

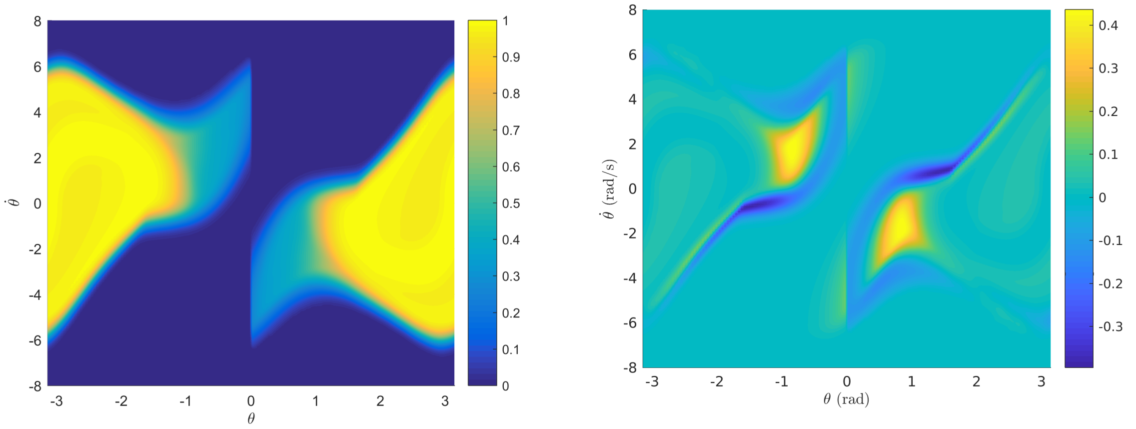

(a) Covariance prediction PIMC method.

(b) Difference between the covariance prediction PIMC method and the LS method.

Figure 8.3: Pendulum on a cart covariance prediction PIMC method.

estimate of the transition probability for a single point. The main idea behind the quadrature point method is to intelligently select a set of points which will be propagated through the (non-linear) dynamic model. The the statistics of the output (e.g. mean and variance) can be determined by a weighted average of the output. This process is similar to the one-step ahead prediction which is performed in a (unscented) Kalman filter. Instead of tracking the uncertainty over a small time step τ, it is tracked over the entire time horizon at once. The final transition probability is then equal to the area under the multivariate Gaussian that falls inside the target set.

8.2.1. Unscented Transform

The main idea of the quadrature point method is an intelligent selection of points which represent the statistics of the model. The Unscented Transform (UT) is a technique that relates these statistics to a minimum set of weighted points. This minimum set of points is often referred to as sigma points. The sigma points are then propagated through the system dynamics after which the statistics are recovered. This process is illustrated in Figure 8.4. From the initial distribution, $2N+1$ sigma points are derived. Each point is assigned a weight based on the initial covariance. Then, each point is transformed through the (non-linear) function $f(\theta)$. Finally the new distribution is calculated by combining the new sigma point location with the existing weights.

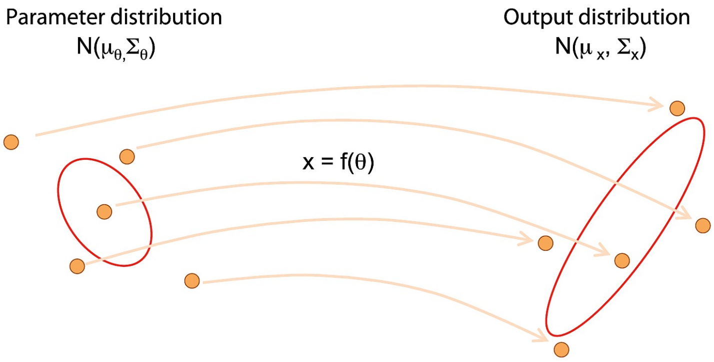

Figure 8.4: UT illustration.

The minimum set of $2L+1$ sigma points is derived, with L the dimension of the state [37].

$$\chi^0_{k-1|k-1} = \mathbf{x}^a_{k-1|k-1} \tag{8.2}$$

$$\chi^i_{k-1|k-1} = \mathbf{x}^a_{k-1|k-1} + \left(\sqrt{(L+\lambda)\mathbf{p}^a_{k-1|k-1}}\right)_i, \quad i = 1,\ldots,L \tag{8.3}$$

$$\chi^i_{k-1|k-1} = \mathbf{x}^a_{k-1|k-1} - \left(\sqrt{(L+\lambda)\mathbf{p}^a_{k-1|k-1}}\right)_{i-L}, \quad i = L+1,\ldots,2L \tag{8.4}$$

The sigma points are then propagated through the function describing the system dynamics

$$\chi^i_{k|k-1} = f\left(\chi^i_{k-1|k-1}\right) \tag{8.5}$$

Finally, the statistics can be recovered from the propagated sigma points.

$$\hat{\mathbf{X}}_{k|k-1} = \sum_{i=0}^{2L} W^i_s \chi^i_{k|k-1} \tag{8.6}$$

$$\mathbf{P}_{k|k-1} = \sum_{i=0}^{2L} W^i_c \left[\chi^i_{k|k-1} - \hat{\mathbf{x}}_{k|k-1}\right]\left[\chi^i_{k|k-1} - \hat{\mathbf{x}}_{k|k-1}\right]^T \tag{8.7}$$

The weights for each sigma point is given by

$$W^0_s = \frac{\lambda}{L+\lambda} \tag{8.8}$$

$$W^0_c = \frac{\lambda}{L+\lambda} + \left(1-\alpha^2+\beta\right) \tag{8.9}$$

$$W^i_s = W^i_c = \frac{1}{2(L+\lambda)} \tag{8.10}$$

The propagation of sigma points is shown Figure 8.5a for a double integrator system. Starting from a terminal condition at $x_T = (2,2)$, the sigma points are propagated through the model towards the target set (black circle). The ellipses represent the 1σ confidence interval as calculated using the UT. Since the mean does not lie within the target set, the point (2,2) would not be part of the deterministic reachable set. However, part of the 1σ confidence interval is located inside this set. In other words, some samples of the stochastic process would be able to reach the target set. Since the area under the PDF is related to the probability, the transition probability can be calculated by the area of the bivariate Gaussian that lies within the target set. This process is illustrated in Figure 8.5b.

8.2.2. Example

Since the quadrature point method makes a number of assumptions (e.g. Gaussian process), it is expected that the approximation of the probabilistic reachable set is of lower fidelity than the PIMC and LS methods. This section performs the quadrature point estimation over the entire state-space in order to compare the difference between these methods. Note that although the computation is performed on a grid, this method allows the computation of the transition probability for a single point. The current results are achieved using the optimal

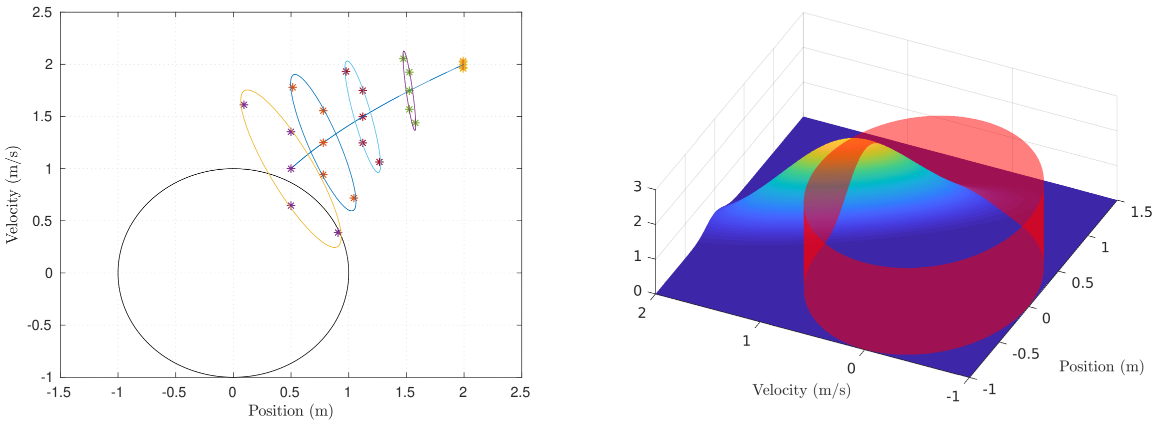

(a) Quadrature point method.

(b) Area of the PDF within the target set (red cylinder).

Figure 8.5: Illustration of the quadrature point method.

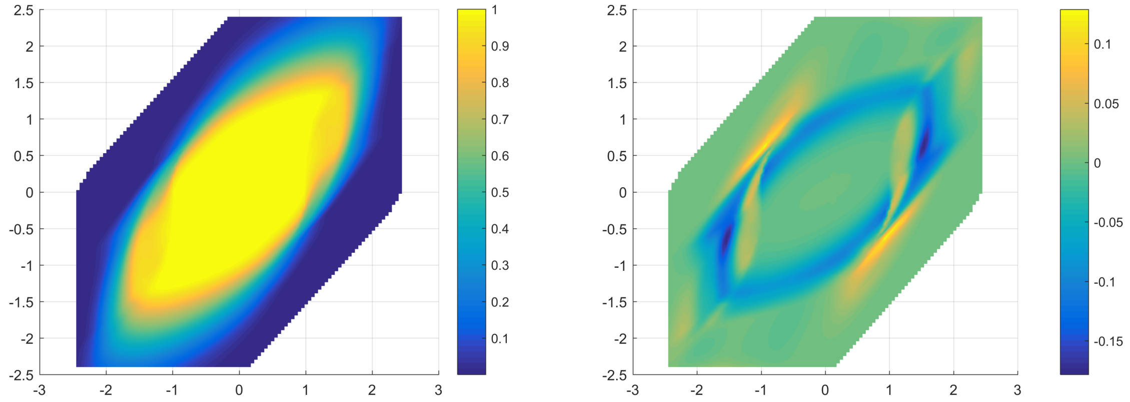

(a) Approximated transition probability.

(b) Difference with the LS method.

Figure 8.6: Quadrature point method approximation for the double integrator system.

control state-feedback law as derived from the LS toolbox, however when combined with other optimal control methods, such as shooting methods, the need for a state-feedback control law can be eliminated.

The first example is the double integrator system. The problem is again similar to the validation performed in Section 5.8.1. Figure 8.6a shows the approximated reachability probability as generated by the quadrature point method. The difference between this method and the LS method is shown in Figure 8.6b. At first glance, the quadrature point method seems to do a good job in approximating the transition probability as calculated using the LS method. The difference seems to be the biggest around the top right and bottom left part of the reachable set. This is most likely related to the switching point of the optimal control law. When the sigma points are located close to the switching point of the control law, some points can be subjected to a different input. Note that for some states, this method fails to determine a transition probability. This happens when the trajectory leaves the computational domain.

A more challenging example is the pendulum on a cart model. Since this model is non-linear, the assumption that the PDF remains Gaussian is no longer valid over a larger time horizon. Figure 8.7a shows the approximated transition probability using the quadrature point method. The difference of this method with the LS method is shown in Figure 8.7b. From these figures, it comes clear that the violation of the Gaussian assumption introduces errors into the transition probability approximation. Furthermore, the switching nature of the optimal control policy adds to this non-linearity. Changing the sigma point parameters (i.e. spacing between the points) could possibly improve the results, although a more detained analysis would be required.

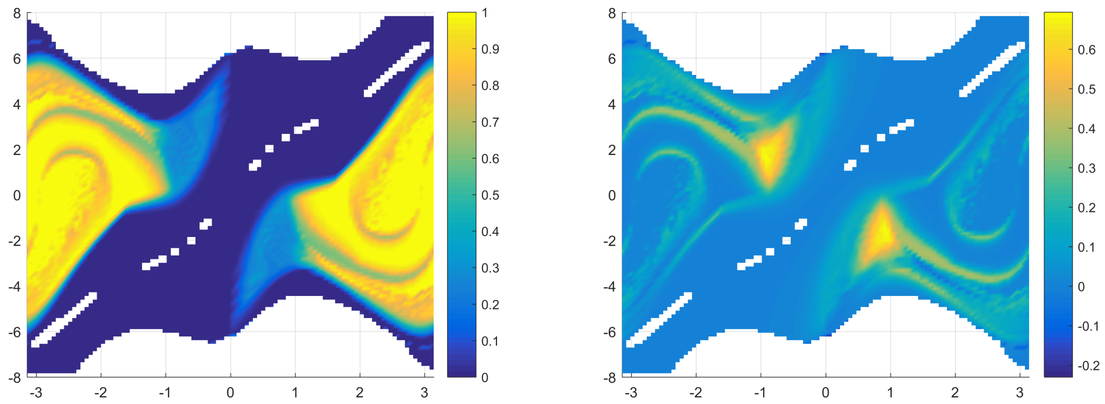

(a) Approximated transition probability.

(b) Difference with the LS method.

Figure 8.7: Quadrature point method approximation for the pendulum on a cart model.

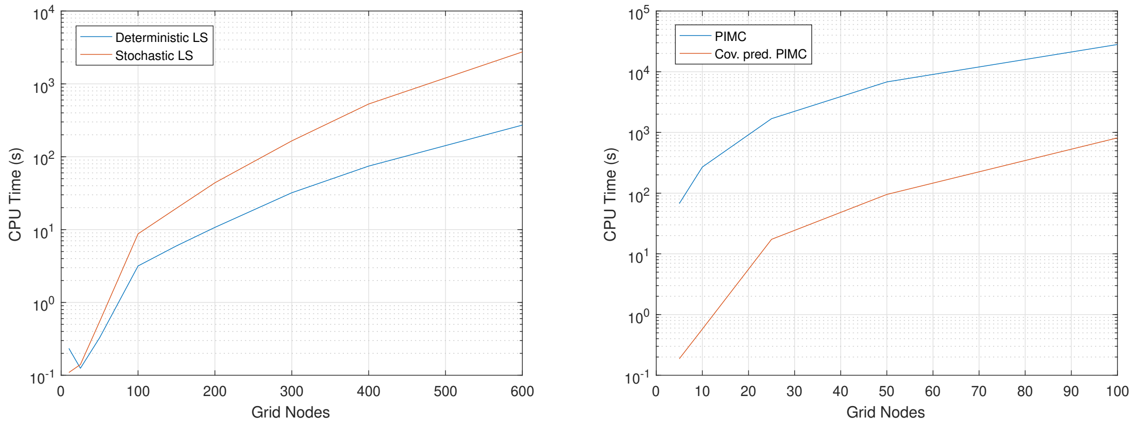

(a) CPU time for the deterministic and stochastic LS method.(b) CPU time for the (covariance prediction) PIMC method.

Figure 8.8: CPU time as a function of the number of grid nodes for the double integrator problem.

8.3. Computational Cost

Since LS methods are part of the Eulerian class of methods, they suffer from the 'curse of dimensionality'. The phenomena has been stated earlier and refers to the poor scaling with respect to the dimensionality of the problem. Consider an unit n-cube with evenly spaced samples so that the spacing between points is no more than 0.01. For a unit 1-cube (a line) 100 sample points are required. The required number of points for a unit n-cube is equal to 100^n which indicates exponential growth with the amount of dimensions.

This section will evaluate the computational costs of the methods that have been proposed in Section 8.1 and 8.2. First, a baseline is established for the LS and PIMC methods. This baseline will be based on a simple double integrator model as introduced in Section 4.4.

Figure 8.8 shows the CPU time for both the LS and PIMC solution on a y-log scale. It is expected that the LS method will have exponential scaling where the base is equal to the grid nodes and the exponent is equal to the dimension. From Figure 8.8a it is clear that the required CPU time increase linearly on a log scale. Furthermore, the probabilistic LS method requires roughly a order of magnitude more computational time due to the computation of the second derivative. However, the scaling is roughly the same as the deterministic LS method.

Since the LS method and PIMC method are implemented in a different way it is not possible to directly compare the computational complexity of the two methods. Therefore, the PIMC method is compared with the covariance prediction PIMC. Figure 8.8b shows the required CPU time for both methods as a function of the number of grid nodes. Again, both methods show exponential scaling with the amount of grid nodes.

Furthermore, the covariance prediction method is approximately one order of magnitude faster than the regular PIMC. Note that the exact difference is dependent on the amount of samples in the PIMC method. The more samples are used, the bigger the relative speedup of the other method will be.

9

Conclusions and Recommendations

Current methods for calculating the Safe Flight Envelope (SFE) often employ reachability analysis to determine the set of states that can safely be attained by an aircraft. Keeping the aircraft within this set guarantees safety of the vehicle and its occupants. Although the SFE is theoretically accurate, some states may be practically unattainable under the influence of (external) disturbances. To quantify the effect of (external) disturbances on the the SFE, a model is presented which combines the stochastic disturbances with the dynamical model of the aircraft. When including stochastic processes, the ordinary differential equation becomes an Stochastic Differential Equation (SDE). In order to perform calculations on these SDEs, Itô calculus was introduced. A combination of the reachability analysis with Itô's calculus yields the field of probabilistic reachability analysis. Instead of providing a set of reachable states, the probabilistic reachability analysis provides the probability that there exists a trajectory which reaches the target set, associated with the admissible control set, within a certain time horizon.

To validate the results of the reachability analysis, a Monte-Carlo (MC) analysis is often performed. The reachable set is approximated by the boundary of a large amount of individual trajectories. Unfortunately, such a method is not suitable for the purpose of validating the probabilistic reachable sets. A new Path Integral Monte Carlo (PIMC) method is introduced which samples the the stochastic process to capture the change of the Probability Density Function (PDF) over a small time step. Furthermore, an analytical verification was performed to ensure that the probabilistic reachability analysis was correctly implemented in the Level-Set (LS) toolbox.

The probabilistic reachability analysis was applied to a widely used simplified aircraft model to investigate how turbulence would change the SFE. The deterministic SFE is computed as the intersection between the forwards and backwards reachable set. Under the assumption that both sets are independent, the probabilistic SFE is determined as the product of the forwards and backwards probabilistic reachable sets which yields the joint probability, or the probability that there exists a trajectory from the initial set to a state and back to the initial set, under influence of some stochastic process. This stochastic process is modeled as a Dryden atmospheric turbulence model. The turbulence creates a gust angle of attack, thus disturbing the aircraft.

When comparing the deterministic SFE with the probabilistic counterpart, it becomes clear that the deterministic envelope roughly corresponds to the 1σ isocontour. This means that for most states in the deterministic SFE there is a 68.3 % or higher probability of returning to the initial set. States on the boundary of the deterministic reachable set would have a 31.7 % probability that the initial set cannot be reached anymore. In most scenarios, a 1σ reachability probability is not acceptable. A better suited definition would be the 3σ contour. This contour contains all states for which the transition probability is at least 99.7 %. The deterministic SFE is significantly larger than the 3σ contour, thus overestimating the SFE. When comparing the area of both envelopes of the simplified aircraft model at different roll angles, it can be observed that the probabilistic SFE ranges from 71.4 % till 50.8 % of the area. This means that at large roll angles the SFE is almost half that of the case without disturbances. Furthermore, for some states in the deterministic SFE, the transition probability under influence of external turbulence is 0 %, meaning that there is no possibility of returning to the initial set.

The second application that was considered is a longitudinal quadrotor model. This model is significantly more complex than the aircraft model. The dynamics of the quadrotor are represented by a full three Degrees-Of-Freedom (DOF) model. Furthermore, the aerodynamics are represented by a non-linear, parametrized model that is identified using flight tests and wind tunnel data. Both factors significantly add to the computational complexity required to compute the SFE. Furthermore, because of the open-loop dynamics, the quadrotor is inherently unstable. Combined with the fact that the control moments are very large in relation to the quadrotor inertia, requires a large part of the state-space to be discretized. This effect is limited by decreasing the time horizon of the reachability analysis and constraining the maximum allowable pitch command to 10 %. The resulting SFE is validated by means of a MC analysis. Comparison of the MC analysis with the results from the LS toolbox show no significant outliers. Finally, an attempt was made to calculate the probabilistic SFE which presented a number of challenges. The atmospheric turbulence model as used in the simplified aircraft example is no longer valid in the domain of the quadrotor. Furthermore, since the aerodynamic model is non-linear, it is not possible to separate the deterministic and stochastic part of the quadrotor model as required by the probabilistic reachability analysis. Finally, the spin-up time of the rotors was not included in the reachability analysis, which therefore overestimates the reachable set. These problems were not addressed in this study since they fall outside of the scope of this research. However, for future research it is recommended to perform the reachability analysis on the closed-loop dynamics in order to circumvent the unstable nature of the quadrotor model. Furthermore, a solution for the rotor spin-up is proposed by adding additional states for the rotor rotational rate and changing the input to a torque instead of an engine setting. Finally, by locally linearizing the aerodynamic model, the gust derivatives could be obtained. This technique would allow separation of the deterministic and stochastic parts of the quadrotor model.

A reoccurring issue is found in the ‘curse of dimensionality’ where the required computational effort increases exponentially with the dimension of the state-space. The probabilistic reachability analysis already adds an order of magnitude to the required computational time, with the same scaling issues. Several methods were considered in an effort to reduce this computational effort. First, the PIMC method which was used to validate the probabilistic reachability analysis was simplified under the assumption that the stochastic process remains Gaussian over a sufficiently small time step. This reduced the computational effort roughly proportional to the amount of amount of samples used in the PIMC method. However, the scaling remains the same. A different method that does not require discretization of the state-space uses a set of sigma points to capture the covariance over a trajectory. Although this method does not suffer from the ‘curse of dimensionality’ it fails to identify the whole state-space. Furthermore, it only provides an estimate of the transition probability. More research would be required to identify the feasibility of this method. Other methods such as Sobol sampling or kernel density estimation are also identified as possible ways to speedup the calculation.

The probabilistic reachability analysis shows great potential for safety analysis, and when applied to the field of envelope protection allows a quantification of the stochastic SFE under the influence of atmospheric turbulence. Although the PIMC method makes no assumptions on the noise, the probabilistic LS method as presented in this research is limited to Gaussian processes. Extending the reachability analysis to overcome these limitations would make this method suitable for more complex systems. The main limiting factor in the application of (probabilistic) reachability analysis is the poor scaling with increasing dimension. This makes any problem with more than four dimensions practically unusable. While it might not be feasible to break the ‘curse of dimensionality’, advances in modern computing, such as parallel and distributed computing, could put more complex problems within reach.

A

Convergence Analysis

One of the sources of the discrepancy between the Path Integral Monte Carlo (PIMC) and Level-Set (LS) method is thought to be the numerical convergence. Especially since numerical Stochastic Differential Equation (SDE) integrators have a low convergence order. A convergence analysis is performed for both methods. It is expected that as the resolution is increased, the solution for both methods will converge. If both methods converge, but not to the same value, the source of the discrepancy is most likely caused by a different phenomena. Figure A.1 shows the convergence of both the LS and the MC method for specific tests points on the pendulum on a cart model. The test points are located on a line as shown in Figure A.1d. The test points with the lowest value are located on the right of the line, while the test points with a high value are located on the left (e.g. the blue line corresponds to the most right test point, green to the most left). The value of the test points as function of the grid size for the LS method is shown in Figure A.1a. It is clear that the value of the test points converge to a steady solution at a grid size of 400. The value for the PIMC method as a function of the amount of time-steps is shown in Figure A.1b. When the amount of time-steps is increased, the value converges. Although both methods have converged, the result is still very different.

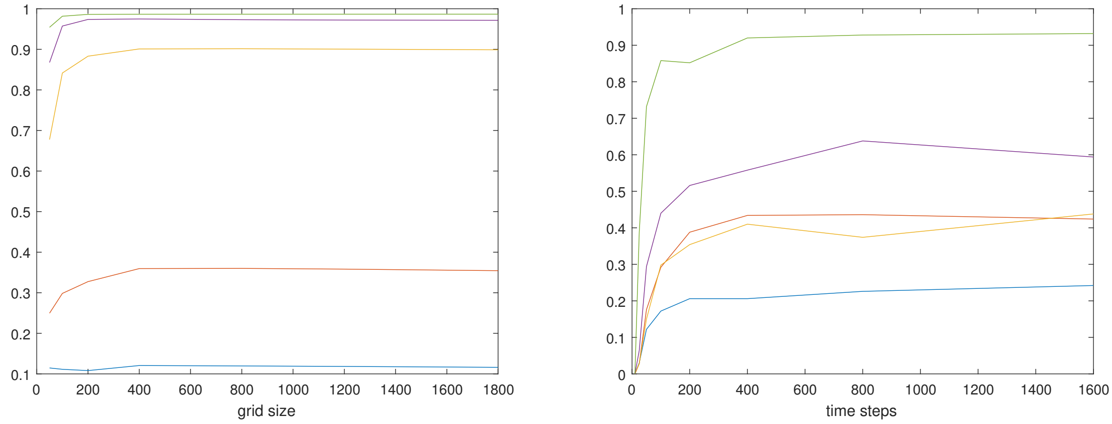

(a) Convergence of the LS method.

(b) Convergence of the PIMC method.

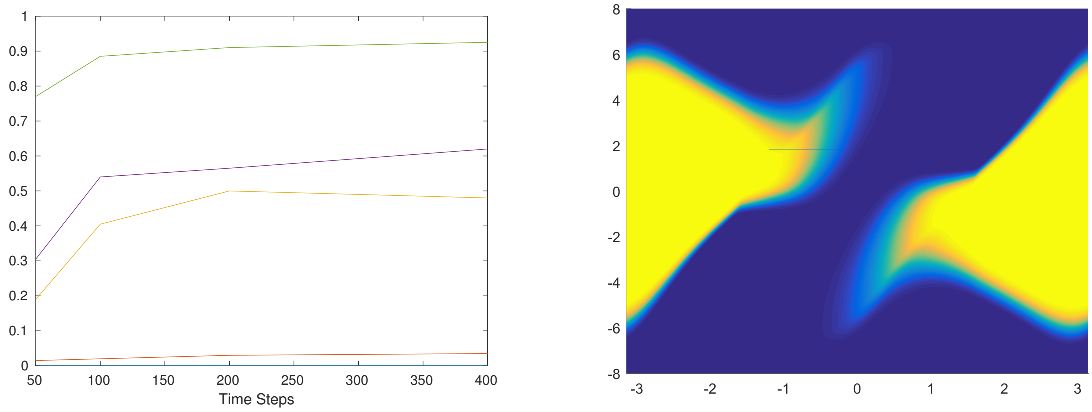

(c) Convergence of the brute force Monte-Carlo (MC).

(d) Location of the convergence sample points.

Figure A.1: Convergence of the LS and the PIMC solution.

B

Quadrotor Probabilistic Safe Flight Envelope

Due to the non-linear nature of the quadrotor aerodynamic model, it is not possible to separate the drift and diffusion parts of the model. This makes it unsuitable for application to the probabilistic reachability analysis. To get an impression of how the Safe Flight Envelope (SFE) changes under the influence of a stochastic process, an arbitrary disturbance is considered. Instead of a suitable turbulence model, the vertical acceleration will be disturbed by a white noise process with a standard deviation of $\sigma = 0.2\,\mathrm{m\,s^{-2}}$. The Equations of Motions (EOMs) will now be represented by

$$\dot{u} = \frac{F_x(\delta_T, \delta_q, u, w)}{m} - qw - g\sin\theta \tag{B.1}$$

$$\dot{w} = \frac{F_z(\delta_T, \delta_q, u, w)}{m} + qu + g\cos\theta + \sigma W \tag{B.2}$$

$$\dot{q} = \frac{M_y(\delta_T, \delta_q, u, w)}{I_{yy}} \tag{B.3}$$

$$\dot{\theta} = q \tag{B.4}$$

It should be noted that this disturbances does not have a physical meaning and does not accurately represent the SFE under the influence of atmospheric turbulence.

The result of the probabilistic reachability analysis is shown in Figure B.1 for a time horizon of $T = 0.1\,\mathrm{s}$. The forwards and backwards reachability analysis are shown in Figure B.1a in green and red respectively. The $1\,\sigma$ isosurface of the probabilistic SFE is shown in Figure B.1b. Finally, some slices of the probabilistic SFE are shown in Figure B.1c.

A comparison between the deterministic and probabilistic SFE for $q = 0°\mathrm{s^{-1}}$, $u = 0\,\mathrm{m\,s^{-1}}$ is shown in Figure B.2. Unlike the simplified aircraft model, the $1\,\sigma$ envelope is much smaller than the deterministic envelope. Furthermore, a substantial part of the deterministic SFE has a transition probability of 0 %.

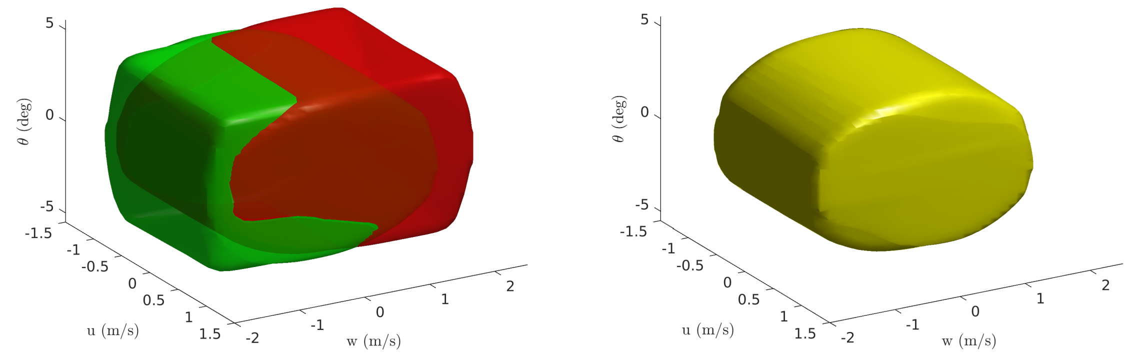

(a) Forwards (green) and Backwards (red) 1 σ isosurface.

(b) Probabilistic SFE 1 σ isosurface.

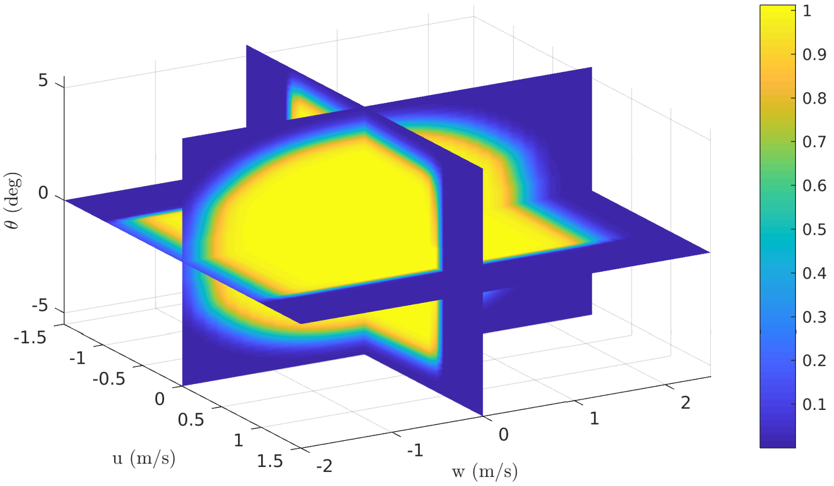

(c) Slice of the probabilistic SFE.

Figure B.1: Probabilistic SFE of the quadrotor model for $q = 0°\mathrm{s}^{-1}$ and $T = 0.1\,\mathrm{s}$.

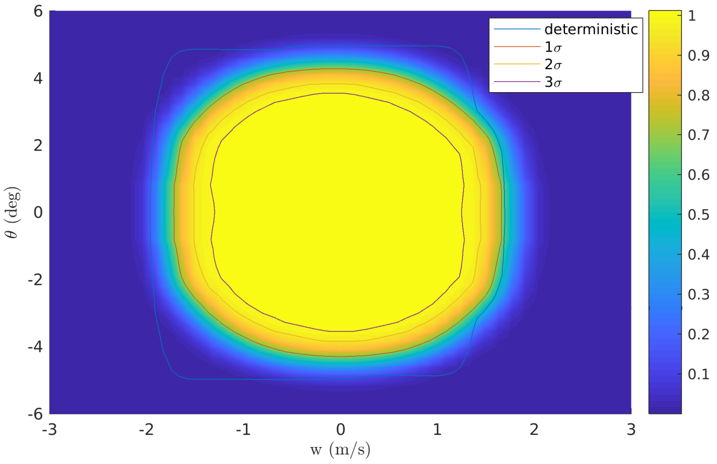

Figure B.2: Comparison between the deterministic and probabilistic SFE for $q = 0°\mathrm{s}^{-1}$, $u = 0\,\mathrm{m\,s}^{-1}$ and $T = 0.1\,\mathrm{s}$..

C

Stochastic Collision Avoidance

One of the most common examples in the topic of reachability analysis using Hamilton-Jacobi-Bellman (HJB) Partial Differential Equations (PDEs) is the game of identical vehicles, or the three dimensional aircraft avoidance example [14, 38, 39]. The problem consists of two vehicles (aircraft in this case) with a position and a heading. Each vehicle has a fixed velocity but can change its heading. In this differential game, the pursuer tries to cause a collision while the evader tries to prevent one. A collision occurs if both vehicles get within distance r of each other. Figure C.1 shows the coordinate system. In this examples, the game of identical vehicles is extended to include an uncertainty on the relative angle between the two aircraft.

The dynamics are now given by

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} -v_a + v_b \cos x_3 + a x_2 \\ v_b \sin x_3 - a x_1 \\ b - a \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \sigma_{x_3} \end{bmatrix} W(t) \tag{C.1}$$

where x_1 and x_2 are the relative distance between the two aircraft. Furthermore, x_3 is the relative angle between the two aircrafts and σ_{x_3} is the uncertainty on the heading angle.

Figure C.2a shows the deterministic reachable set for a time horizon of 2 s. Figure C.2b shows the 1σ isocontour for the probabilistic reachable set. This set includes all states with a 68.3 % or higher chance of a collision. A more detailed illustration of the probabilistic reachable set is shown in Figure C.2c.

From this analysis, it becomes clear that adding an uncertainty on the heading angle can have large effects on the presumed safe region. For some regions there is a significant chance of collision, where the deterministic reachability analysis predicts no collision at all.

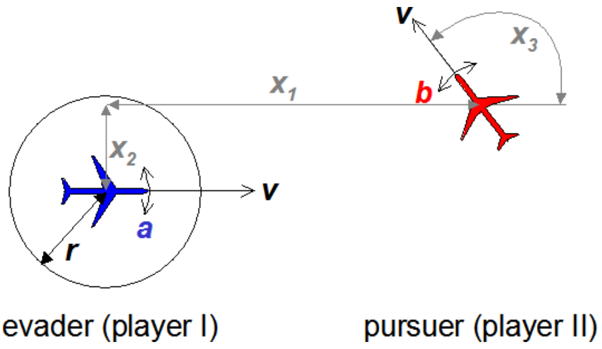

Figure C.1: Relative coordinate system for game of two identical vehicles [40].

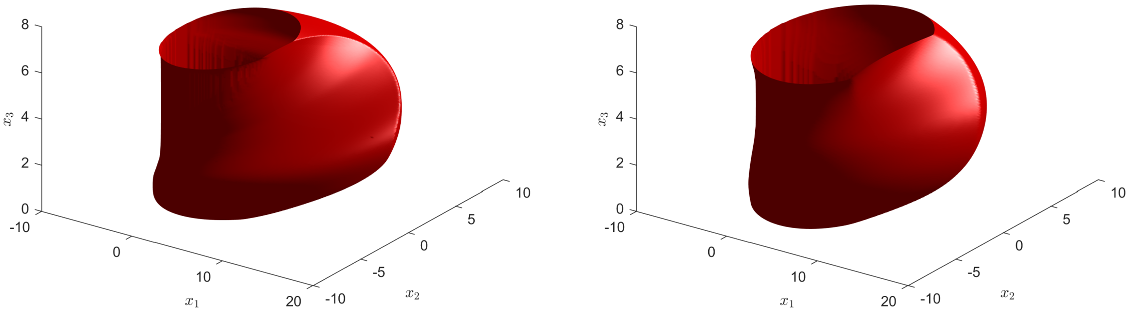

(a) Deterministic reachable set at $t = 2\,\mathrm{s}$.

(b) 1σ isosurface of the stochastic reachable set.

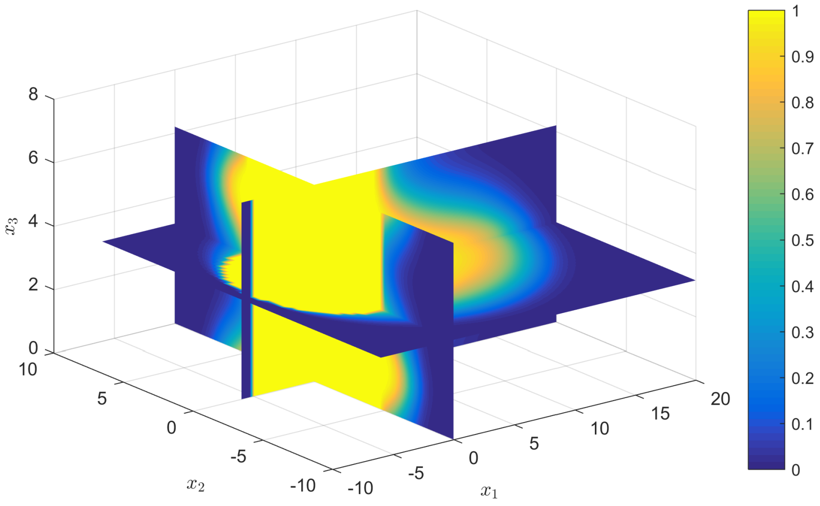

(c) Slices of the stochastic reachable set.

Figure C.2: Aircraft collision avoidance.

D

Draft Paper

Safe Flight Envelope Uncertainty Quantification using Probabilistic Reachability Analysis

R. van den Brandt * **C.C. de Visser** **

* *Graduate Student, Department of Control and Simulation, Delft University of Technology, Kluyverweg 1 2629 HS Delft, The Netherlands.*
** *Assistant Professor, Department of Control and Simulation, Delft University of Technology, Kluyverweg 1 2629 HS Delft, The Netherlands.*

Abstract: Loss of Control is the primary contributor to aviation fatalities. To prevent this type of accident, flight envelope protection is considered to be a necessary development. The calculation of the Safe Flight Envelope provides a bound on the states that can safely be approached by the aircraft. Although theoretically accurate, some states may not be reachable under the influence of disturbances (e.g. turbulence). In this paper a stochastic extension to the reachability analysis is applied to a simplified aircraft model. The probabilistic reachability analysis yields the transition probability from a state to the target set. By comparing the deterministic and probabilistic Safe Flight Envelope, it becomes clear that the Safe Flight Envelope can shrink considerably under the influence of turbulence. It is shown that for a $3\,\sigma$ (99.7 %) confidence interval, the envelope can shrink by as much as 50.8 % compared to the deterministic envelope. Furthermore, it is found that for high roll angles, some parts of the deterministic envelope have a 0 % transition probability under the influence of turbulence, further emphasizing the importance of probabilistic envelopes.

Keywords: Safe Flight Envelope, Reachability Analysis, Turbulence, Uncertainty Quantification, Probabilistic, Loss-of-Control, level-set, HJB

---

1. INTRODUCTION

A recent study by Boeing (2017) into fatal airplane accidents has shown that the majority of aviation fatalities are the result of Loss-of-Control (LOC) related incidents. Although the LOC occurrence category represented only 4 % of all 2016 accidents, it remains of significant concern since it accounts for 42.9 % of all fatal accidents (ICAO (2017)). LOC can occur when there is an unintended departure from controlled flight. Often, this happens when an aircraft leaves the normal flight envelope. This may result in a state from which recovery is hard or impossible (e.g. entering a stall or spin).

In search of methods to prevent LOC related accidents, flight envelope protection is considered to be very effective (Lambregts et al. (2008)). Flight envelope protection consists of a human-machine interface that prevents the pilot from steering the aircraft into a state which exceeds the structural and aerodynamic operating limits. For such a system to be effective, knowledge of the flight envelope is required. The flight envelope, according to the standard definition, refers to a region of velocity and altitude or load-factor where the aircraft can operate safely. Since such an envelope only contains quasi-stationary states it is insufficient for the purpose of full envelope protection. A more comprehensive definition which also includes non-stationary states is known as the Safe Flight Envelope (SFE) (van Oort et al. (2011); Zhang et al. (2016); Lombaerts et al. (2013)). The SFE refers to the set of all states that can be reached, for which a return trajectory to the steady-state exists. It is formally defined as the intersection between the forwards and backwards reachable set. Computation of the reachable set is facilitated by making a connection between reachable set theory and the SFE (Lygeros (2004)).

Although the computed SFE is theoretically accurate, some states may be practically unattainable under the influence of disturbances. The presence of a disturbance could, for example, push the aircraft into a state outside the SFE so that it is unable to return to a steady-state. The main contribution of this paper is a new approach for flight envelope prediction from a probabilistic point of view that can quantify practical attainability. The stochastic counterpart of the reachability set theory will be used to compute the probabilistic SFE. By analyzing the differences between the deterministic and probabilistic SFEs, an assessment of the sensitivity of the SFE to disturbances can be made. We demonstrate our new approach using a widely used simplified aircraft model and show that current deterministic envelope prediction methods may produce dangerously optimistic results when neglecting the influence of disturbances.

2. SAFE FLIGHT ENVELOPE

The flight envelope commonly refers to a region that indicates the capabilities of an aircraft in terms of velocity and altitude or load-factor. Under nominal conditions the aircraft should always be operated within these limits. When a flight envelope excursion happens, the aircraft can enter an upset condition from which it is hard or impossible to recover. Flight envelope protection is necessary (but not sufficient) for preventing the aircraft from leaving the flight envelope, thereby preventing LOC. Since the standard flight envelope only takes into account quasi-stationary states (e.g. coordinated turns or level flight) it does not provide bounds on dynamic maneuvers. To overcome the shortcomings of the standard flight envelope, a more complete representation of the flight envelope is required. The concept of a Safe flight Envelope (SFE) is introduced for the purpose of preventing LOC. It is formally stated as: *"The part of the state-space for which safe operation of the aircraft and safety of its cargo can be guaranteed, and externally posed constraints will not be violated"* (Van Oort (2011)).

The externally posed constraints are composed of the following envelope's:

- Dynamic Envelope: Envelope that is constrained by the dynamic behavior of the aircraft, due to its aerodynamics and kinematics.
- Structural and Comfort envelope: Constraints posed by the airframe, pilot, passengers and cargo. These constraints are usually defined through maximum accelerations and loads.
- Environmental Envelope: Constraints due to the environment in which the aircraft operates.

The constraints posed in the last two envelopes are well-known and can be quantified quite easily, however this is not true for the dynamic envelope. The main focus of this paper will be on the calculation of the dynamic envelope, where the complete envelope is given by the intersection of all three flight envelopes.

3. REACHABLE SET THEORY

The formal definition of the SFE provides a high level description, however this definition is not sufficient to facilitate computation of the SFE. For this purpose the SFE will be expressed in the framework of reachable set theory. Reachability set theory is widely used in safety analysis where the interest is in finding all states that can be reached, starting from a set of initial conditions, over a certain time horizon. In this section the connection with the reachable set theory will be made (Lygeros (2004)).

Consider a continuous-time dynamic system

$$\dot{\mathbf{x}}(t) = f(t, \mathbf{x}(t), \mathbf{u}(t)) \quad \mathbf{x} \in \mathbb{R}^n, \mathbf{u}(\cdot) \in \mathcal{U} \subseteq \mathbb{R}^m \qquad (1)$$

where $f(t, \mathbf{x}, \mathbf{u})$ is assumed to be Lipschitz continuous (i.e. $\dot{f}(t, \mathbf{x}, \mathbf{u}) < L$). With a time horizon $T \geq 0$ and a set of initial states $\mathcal{K} \subseteq \mathbb{R}^n$. A solution to the dynamic system is a trajectory through the state-space. The time horizon and the set of admissible control inputs ($\mathcal{U}$) provide a bound on the states that can be reached. A larger time horizon leads to a larger reachable set. Furthermore, when larger control inputs are permitted, a larger part of the state-space can be reached in the same amount of time. Based on the definition of the dynamical system given by (1), three different types of reachable sets can be defined.

- *Invariance:* $\mathcal{I}(t, \mathcal{K})$: Set of all states $\mathbf{x}(\cdot)$ for which there exists all the inputs $\mathbf{u}(\cdot) \in U_{[0,T]}$ such that $\mathbf{x}(t) \in \mathcal{K}$ for all $t \in [0, T]$.
- *Viability:* $\mathcal{V}(t, \mathcal{K})$: Set of all states $\mathbf{x}(\cdot)$ for which there exists at least one input $\mathbf{u}(\cdot) \in U_{[0,T]}$ such that $\mathbf{x}(t) \in \mathcal{K}$ for all $t \in [0, T]$.
- *Reachability:* $\mathcal{R}(t, \mathcal{K})$: Set of all the states $\mathbf{x}(\cdot)$ for which there exists at least one input $\mathbf{u}(\cdot) \in U_{[0,T]}$ and $t \in [0, T]$ such that $\mathbf{x}(t) \in \mathcal{K}$.

Fig. 1 provides an illustration of the different reachable kernels. The *invariance* set is the smallest set. It contains all states for which all possible control inputs keep the trajectory within $\mathcal{K}$ over the entire time horizon. Set $\mathcal{V}$ is the set of all states for which there exists at least one control sequence which keeps the trajectory within $\mathcal{K}$ over the entire time horizon. This set, together with the *invariance* set are completely contained inside the initial set $\mathcal{K}$. Finally, $\mathcal{R}$ is the set of all states for which there exists at least one control sequence that can steer the trajectory into set $\mathcal{K}$ within the time horizon. The definition of the *invariance*, *viability* and *reachability* kernel are used extensively in safety analysis (Lygeros (2004); Van Oort (2011); van Oort et al. (2011)).

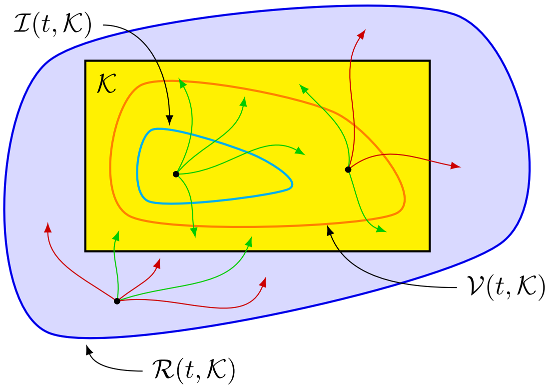

Fig. 1. Illustration of the different reachable kernels. Viability (orange), Invariance (cyan) and (backwards) reachability Set (Blue), Green: Successful trajectories; Red: Unsuccessful trajectories (Nabi (2016)).

The dynamic system (1) can be solved forwards in time, but also backwards in time which gives rise to two further types of reachability:

- *Forwards reachable Set:* $\mathcal{R}_f(t, \mathcal{K})$: Set of all the states $\mathbf{x}(\cdot)$ for which there exists a control input $\mathbf{u}(\cdot) \in U_{0,T}$ such that this set can be approached starting from the initial set $\mathcal{K}$ for $t \in [0, T]$.
- *Backwards Reachable Set:* $\mathcal{R}_b(t, \mathcal{K})$: Set of all the states $\mathbf{x}(\cdot)$ for which a control input $\mathbf{u}(\cdot) \in U_{[0,T]}$ exists at time $t \in [0, T]$, such that at least one state in the initial set $\mathcal{K}$ can be reached.

Fig. 2 shows both the forwards and backwards reachable set. For any state in the forwards reachable set, a trajectory towards this state exists, starting from somewhere in

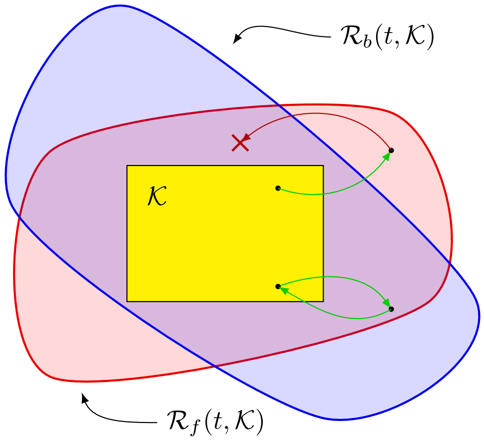

Fig. 2. SFE defined as the intersection between the forwards and backwards reachability set.

the initial set. The backwards reachable set consists of all the initial conditions for which a solution to the dynamic system exists that reaches the set $\mathcal{K}$. The SFE is defined as the intersection between the two. For any point inside the SFE, there exists a trajectory from the initial set $\mathcal{K}$ to a state and back. If the state is in the forwards reachable set, but not in the backwards reachable set, it is possible to maneuverer to the state inside the forwards reachable set, but it is unable to return to the initial set.

4. OPTIMAL CONTROL FORMULATION

The *viability* and *invariant* sets as introduced in Section 3 can be linked to SUPMIN and INFMIN optimal control problems. The relation between the reachable sets and optimal control problems has been stated and proven by Lygeros (2004) and will be briefly summarized here.

Consider the continuous-time dynamic system

$$\dot{\mathbf{x}} = f(\mathbf{x}, t, \mathbf{a}, \mathbf{b}) \quad \mathbf{x} \in \mathbb{R}^n, \mathbf{a} \in \mathcal{A} \subseteq \mathbb{R}^{m_a}, \mathbf{b} \in \mathcal{B} \subseteq \mathbb{R}^{m_b} \tag{2}$$

where $\mathbf{a}$ and $\mathbf{b}$ are inputs to the system of differential equations. The inputs are considered to be two actors – pursuer and evader – with opposite goals. The pursuer tries to maximize the objective function while the evader tries to minimize it. This kind of problem is often referred to as a differential game. Solutions of the dynamic system (2) are trajectories of the system and will be denoted by (Mitchell et al. (2005))

$$\xi_f(\cdot; \mathbf{x}, t, \mathbf{a}, \mathbf{b}) \tag{3}$$

which satisfies the initial condition $\xi_f(0; \mathbf{x}, t, \mathbf{a}, \mathbf{b}) = x_0$ and is differentiable almost everywhere

$$\frac{\mathrm{d}}{\mathrm{d}\tau}\xi_f\left(\tau; \mathbf{x}, t, \mathbf{a}, \mathbf{b}\right) = f\left(\xi_f\left(\tau; \mathbf{x}, t, \mathbf{a}, \mathbf{b}\right), \tau, \mathbf{a}(\tau), \mathbf{b}(\tau)\right) \tag{4}$$

Note the semi-colon to distinguish between the argument τ of ξ_f and the trajectory parameters.

The objective of the differential game is defined by a value function

$$V_1(\mathbf{x}, t) = \sup_{\mathbf{a}\in\mathcal{A}} \inf_{\mathbf{b}\in\mathcal{B}} \min_{\tau\in[0,T]} l\left(\xi\left(\tau; \mathbf{x}, t, \mathbf{a}, \mathbf{b}\right)\right) \tag{5}$$

$$V_2(\mathbf{x}, t) = \inf_{\mathbf{a}\in\mathcal{A}} \sup_{\mathbf{b}\in\mathcal{B}} \min_{\tau\in[0,T]} l\left(\xi\left(\tau; \mathbf{x}, t, \mathbf{a}, \mathbf{b}\right)\right) \tag{6}$$

V is called the value function or the cost-to-go. It represents the value of the solution to the optimal control problem at time t and initial state $\mathbf{x}$. The input $\mathbf{a}$ tries to maximize the minimum value attained by the function l along the state trajectory over the time horizon $[0, T]$ while the evader $\mathbf{b}$ tries to minimize this.

The *viability* and *invariant* sets can now be linked to the value function by (7,8)

$$\mathcal{V}(t, K) = \{\mathbf{x} \in \mathbb{R}^n | V_1(\mathbf{x}, t) > 0\} \tag{7}$$

$$\mathcal{I}(t, K) = \{\mathbf{x} \in \mathbb{R}^n | V_2(\mathbf{x}, t)) \geq 0\} \tag{8}$$

with $V(x, T) = l(x)$. An intuitive derivation can be made if it is assumed that the value function is differentiable. The value function can be reformulated as a Hamilton Jacobi Bellman (HJB) Partial Differential Equation (PDE) using Bellman's optimality Principle and a Taylor expansion of the value function (Van Oort (2011))

$$V_2(\mathbf{x}, t) = \inf_{\mathbf{a}\in\mathcal{A}} \sup_{\mathbf{b}\in\mathcal{B}} V\left(\xi_f\left(t + \Delta t; x, t, \mathbf{a}, \mathbf{b}\right), t + \Delta t\right) \tag{9}$$

First order Taylor expansion yields

$$V\left(\xi_f\left(t + \Delta t; x, t, \mathbf{a}, \mathbf{b}\right), t + \Delta t\right) \approx V(x, t) + V_t(x, t)\Delta t + V_x(x, t)\Delta x \tag{10}$$

Re-arranging the terms of (9) by the principle of optimality and dividing by Δt yields

$$\inf_{\mathbf{a}\in\mathcal{A}} \sup_{\mathbf{b}\in\mathcal{B}} \frac{V\left(\xi_f\left(t + \Delta t; x, t, \mathbf{a}, \mathbf{b}\right), t + \Delta t\right) - V(x, t)}{\Delta t} = 0 \tag{11}$$

Equation (11) can be thought of as one step/sub-problem of the dynamic programming principle.

Substituting the Taylor expansion into (11) yields

$$\inf_{\mathbf{a}\in\mathcal{A}} \sup_{\mathbf{b}\in\mathcal{B}} \frac{V_t(x, t)\Delta t + V_x(x, t) \cdot \Delta x}{\Delta t} = 0 \tag{12}$$

For the limit of $\Delta t \to 0$

$$\frac{\partial V_1}{\partial t}(\mathbf{x}, t) + \min\left\{0, \sup_{\mathbf{a}\in\mathcal{A}} \inf_{\mathbf{b}\in\mathcal{B}} \frac{\partial V_1}{\partial \mathbf{x}}(\mathbf{x}, t) f(\mathbf{x}, t, \mathbf{a}, \mathbf{b})\right\} = 0 \tag{13}$$

$$\frac{\partial V_2}{\partial t}(\mathbf{x}, t) + \min\left\{0, \inf_{\mathbf{a}\in\mathcal{A}} \sup_{\mathbf{b}\in\mathcal{B}} \frac{\partial V_2}{\partial \mathbf{x}}(\mathbf{x}, t) f(\mathbf{x}, t, \mathbf{a}, \mathbf{b})\right\} = 0 \tag{14}$$

The result is a PDE which includes an optimization over the inputs $\mathbf{a}$ and $\mathbf{b}$. The minimization term with zero guarantees that the reachable set cannot shrink as time marches forwards. A state that already entered the target set is restricted from leaving before the final time by 'freezing' the evolution of the trajectory (Zhang et al. (2016)).

Unfortunately, the assumption made in (9) that the value function is differentiable is often not valid, making it unable to interpret it as a solution in the 'classical' sense. Due to the switching of the optimal control and a discontinuity present in the right hand side of the HJB PDE the value function may not remain continuous. In order to obtain a solution to the HJB PDE a 'weak' formulation of solutions to these equations is necessary. Viscosity solutions form a general theory of 'weak' (i.e. non-differentiable) solutions. The interested reader is referred to Bressan (2011) for more information on viscosity solutions. Note that viscosity solutions are not the same as vanishing viscosity solutions which are the limit of an additional viscosity term for $\epsilon \to 0$.

5. PROBABILISTIC REACHABILITY ANALYSIS

Stochastic Differential Equations (SDEs) can be used to describe the behavior of a stochastic process. Instead of simulating each realization of the stochastic process, the forward Kolmogorov equation (or Fokker-Planck equation) can be used to describe the time rate of change of the Probability Density Function (PDF). Combining the optimal control framework with the time-evolution of the PDF will give rise to the field of probabilistic reachability analysis. The probabilistic reachability analysis will compute the transition probability to a certain state. Alternatively, the value function will be equal to the conditional probability of being in the initial set $\mathcal{K}$ at time $t = 0$ and in a final state X at time $t = T$.

Consider a controlled stochastic process $X^u_{t,x}$

$$\mathrm{d}X(s) = f(X(s), s, u(s))\mathrm{d}t + \sigma(X(s), s)\mathrm{d}W_s, \quad \forall s \in [t, T] \tag{15}$$

Let $\mathcal{K}$ be the non-empty target set in $\mathbb{R}^d$, $\rho \in [0, 1[$ and $t \leq T$. Consider the reachable set Ω_t under probability of success ρ, or the set of initial conditions x for which the probability that there exists a trajectory $X^u_{t,x}$ that reaches set $\mathcal{K}$ at time T, associated with the admissible control $u \in \mathcal{U}$ is at least ρ (Assellaou et al. (2015))

$$\Omega^\rho_t = \left\{x \in \mathbb{R}^d | \exists u \in \mathcal{U}, \mathbb{P}[X^u_{t,x}(T) \in \mathcal{K} > \rho]\right\} \tag{16}$$

Because the dynamics are stochastic, it is no longer possible to minimize the value function. Instead, the expected pay-off is minimized over all possible further realizations of the Wiener process (Kappen et al. (2011); Doya (2007)).

$$V(x,t) = \min_{\mathbf{u} \in \mathcal{U}} \left\{\mathbb{E}\left[V(x + \Delta f(x,u) + \xi, t + \Delta t)\right]\right\} \tag{17}$$

with $\xi \sim \mathcal{N}(0, \sigma(x,t))$ a process with normal distribution. The derivation is performed in a similar way as the deterministic case in Section 4. Taking the Taylor expansion of the value function by means of Itô's calculus. Since $\mathrm{d}x^2$ is of the order $\mathrm{d}t$ because of the Wiener process, the expansion must be performed up to the second order.

$$V(x + \Delta x, t + \Delta t) \approx V(x,t) + V_t(x,t)\Delta t + V_x(x,t)\Delta x + \frac{1}{2}\left(V_{xx}(x,t)\Delta x^2 + 2V_{xy}\Delta x \Delta t + V_{tt}(x,t)\Delta t^2\right) + \mathcal{O}(\delta^3) \tag{18}$$

where $\Delta x = \Delta f(x,t,u) + \xi$

Keeping all terms of order $\mathcal{O}(\Delta t)$

$$\mathbb{E}[V] = V(x,t) + \Delta f(x,u)V_x(x,t) + \frac{1}{2}\operatorname{Tr}(\sigma^2(x,t)V_{xx}(x,t)) \tag{19}$$

Substituting into (17) and dividing by Δt

$$\frac{V(x,t) - v(x, t + \Delta t)}{\Delta t} = \min_{\mathbf{u} \in \mathcal{U}} \left\{V_x f(x,t)^T + \frac{1}{2}\operatorname{Tr}(\sigma(x,t)\sigma(x,t)^T V_{xx})\right\} \tag{20}$$

For the limit of $\Delta t \to 0$ the value function is characterized by the viscosity solution to the stochastic HJB equation (Esfahani et al. (2016))

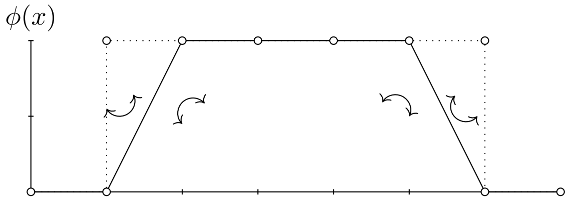

Fig. 3. Regularization of the indicator function.

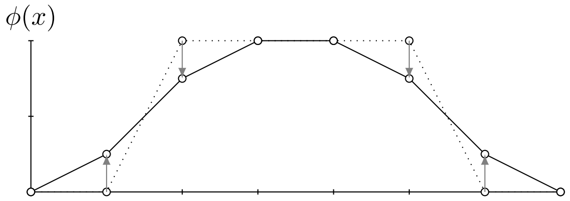

Fig. 4. Diffusion of the LS function.

$$\frac{\partial V}{\partial x}(x,t) + \inf_{u \in \mathcal{U}}\left\{\frac{\partial V}{\partial x}(x,t) f(x,t,u)\right\} + \frac{1}{2}\operatorname{Tr}\left\{\sigma(x,t)\sigma(x,t)^T \frac{\partial^2 V}{\partial x^2}(x,t)\right\} = 0 \tag{21}$$

The stochastic HJB equation can be solved using a similar method as for the deterministic HJB equation. Next to the convection term, an additional diffusion term is added to the Level-Set (LS) method. In contrary to the deterministic reachable set, the probabilistic reachable set is not represented by a zero level contour of the value function. Instead, the signed distance function is replaced by an indicator function.

$$\mathbb{1}_{\mathcal{K}}(x) = \begin{cases} 1 & \text{if } x \in \mathcal{K} \\ 0 & \text{if } x \notin \mathcal{K} \end{cases}$$

The initial condition of the LS function is then equal to one for all states that are inside the initial set $\mathcal{K}$, and all other states are equal to zero. Fig. 3 shows the indicator function as a dotted line. The regularization of the indicator function is shown as a solid line. Due to the discretization of the state-space, the LS function is slightly smaller which can introduce errors to the final solution. An error estimation is performed by Assellaou et al. (2015). The arrows in Fig. 3 represent the concavity of the implicit function. For a diffusion only problem, the second derivative, as found in the stochastic part of the HJB equation represents this concavity. Fig. 4 shows how the concavity of the LS function is related to the diffusion.

6. SIMPLIFIED AIRCRAFT MODEL

To analyze how external disturbances influence the SFE, a simplified aircraft model is considered. The deterministic SFE is calculated from the intersection of the forwards and backwards reachable set. Then, the probabilistic SFE is computed for the aircraft under influence of a suitable turbulence model.

6.1 Dynamic model

For the calculation of the SFE a nonlinear 3D simplified aircraft model will be considered (Lombaerts et al. (2013)). The dynamics of the model are

$$\dot{V} = -\frac{\rho S}{2m}V^2 C_{D_0} - g\sin(\gamma) + \frac{T}{m} - \frac{\rho S}{2m}V^2(C_{D_\alpha}\alpha + C_{D_{\alpha^2}}\alpha^2) \tag{22}$$

$$\dot{\gamma} = -\frac{g}{V}\cos(\gamma) + \frac{\rho S}{2m}V(C_{L_0} + C_{L_\alpha}\alpha)\cos(\phi) - \frac{\rho S}{2m}V C_{Y_\beta}\beta\sin(\phi) \tag{23}$$

where $C_{L_\alpha} = 6.0723$, $C_{L_0} = 1.0656$, $C_{D_0} = 0.1599$, $C_{D_\alpha} = 0.5035$, $C_{D_{\alpha^2}} = 2.1175$, $m = 120\,000$ kg, $S = 260\,\mathrm{m}^2$ and $\rho = 1.225\,\mathrm{kg\,m}^{-3}$. The aerodynamic coefficients are based on the RCAM model (Looye and Bennani (1997)).

Both the thrust (T), the angle of attack (α) and the side-slip angle (β) are considered to be inputs to the system and are bounded. $\alpha \in [0; 14.5]^\circ$, $\beta \in [-5, 5]^\circ$ and $T \in [20\,546; 410\,920]$ N. Furthermore, the roll angle is threated as a discretely gridded input.

The maximizers for the Hamiltonian ($\hat{\alpha}$,$\hat{\beta}$ and $\hat{T}$) are

- If $p_1 > 0$ then $\hat{T} = T_{min}$ and
  - · If $\hat{p} > \bar{\alpha}$ then $\hat{\alpha} = \alpha_{min}$
  - · If $\hat{p} = \bar{\alpha}$ then $\hat{\alpha} = [\alpha_{min}; \alpha_{max}]$
  - · If $\hat{p} < \bar{\alpha}$ then $\hat{\alpha} = \alpha_{max}$
- If $p_1 = 0$ then $\hat{T} \in [T_{min}; T_{max}]$ and
  - · If $p_2 > 0$ then $\hat{\alpha} = \alpha_{min}$
  - · If $p_2 = 0$ then $\hat{\alpha} \in [\alpha_{min}; \alpha_{max}]$
  - · If $p_2 < 0$ then $\hat{\alpha} = \alpha_{max}$
- If $p_1 < 0$ then $\hat{T} = T_{max}$ and
  - · If $\hat{p} \le \alpha_{min}$ then $\hat{\alpha} = \alpha_{min}$
  - · If $\alpha_{min} \le \hat{p} \le \alpha_{max}$ then $\hat{\alpha} = \hat{p}$
  - · If $\hat{p} \ge \alpha_{max}$ then $\hat{\alpha} = \alpha_{max}$

besides, for the side-slip angle β holds

- If $p_2 \sin\phi > 0$ then $\hat{\beta} = \beta_{min}$
- If $p_2 \sin\phi = 0$ then $\hat{\beta} \in [\beta_{min}; \beta_{max}]$
- If $p_2 \sin\phi < 0$ then $\hat{\beta} = \beta_{max}$

where $\hat{p} = \frac{p_2 C_{L_\alpha}\cos\phi - p_1 V C_{D_\alpha}}{2 p_1 V C_{D_{\alpha^2}}}$ and $\bar{\alpha} = \frac{\alpha_{min} + \alpha_{max}}{2}$.

6.2 Safe Flight Envelope

The SFE as introduced in Section 3 is defined by the intersection of the forwards and backwards reachable set. These reachable sets are calculated by implementing the simplified aircraft model in the Level-Set Toolbox by Mitchell (2008). Fig. 5 shows the result of the reachability analysis starting from an initial set $V \in [55; 85]\ \mathrm{m\,s}^{-1}$ and $\gamma \in [-10; 10]^\circ$ over a time horizon of 2 s. The forwards and backwards reachable set are shown in green and red respectively. The SFE is then computed by taking the intersection of both flight envelopes. Using the LS method, this is achieved by taking the maximum of both sets. The resulting intersection is marked in yellow. It can be seen that as the roll angle increases, the SFE becomes smaller. A similar trend can be seen for the flight path angle γ.

6.3 Probabilistic Safe Flight Envelope

For the probabilistic SFE the same model will be considered as for the deterministic SFE with the addition of a disturbance on the angle of attack (i.e. turbulence). The atmospheric turbulence for the simplified aircraft model is given by the Dryden spectra. For medium to high altitude turbulence, the turbulence scale length is chosen to be 500 m with a turbulence intensity of $2\,\mathrm{m\,s}^{-1}$ (MIL (1990)).

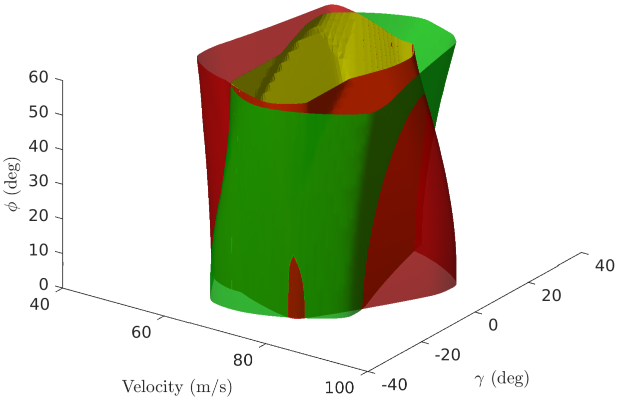

Fig. 5. Forward and Backward reachable set for the simplified aircraft model.

Then the gust angle of attack is calculated as

$$\alpha_{total} = \alpha + \alpha_g \tag{24}$$

$$\alpha_g = \tan^{-1}\left(\frac{w_g}{V}\right) \approx \frac{w_g}{V} \tag{25}$$

The dynamical model for the simplified aircraft now becomes

$$\dot{V} = -\frac{\rho S}{2m}V^2 C_{D_0} - g\sin(\gamma) + \frac{T}{m} - \frac{\rho S}{2m}V^2(C_{D_\alpha}\alpha + C_{D_{\alpha^2}}\alpha^2) + W\frac{\rho S}{2m}V^2 C_{D_\alpha} \tag{26}$$

$$\dot{\gamma} = -\frac{g}{V}\cos(\gamma) + \frac{\rho S}{2m}V(C_{L_0} + C_{L_\alpha}\alpha)\cos(\phi) - \frac{\rho S}{2m}V C_{Y_\beta}\beta\sin(\phi) + W\frac{\rho S}{2m}V C_{L_\alpha} \tag{27}$$

Where W is the Brownian motion process representing the gust angle of attack. It should be noted that the noise terms with α^2 are left out in order to keep the control independent of the noise.

Fig. 6 shows the results of the probabilistic reachability analysis, starting from the same initial set $V \in [55; 85]\ \mathrm{m\,s}^{-1}$ and $\gamma \in [-10; 10]^\circ$ over a time horizon of 2 s. The $1\,\sigma$ isosurfaces for the forwards and backwards reachable set are shown in green and red respectively. The $1\,\sigma$ confidence interval corresponds to all states with a 68.3 % reachability probability or higher. For the probabilistic reachability analysis, the SFE cannot be computed as the intersection of the two sets anymore. Instead, it is defined as the joint probability of the two sets (i.e. the probability to transition to a state and transition back to the initial set). For two independent processes this is equal to the product of the forwards transition probability with the backwards transition probability. Fig. 7 shows slices of the probabilistic SFE that give a better insight in how the probability changes with the states.

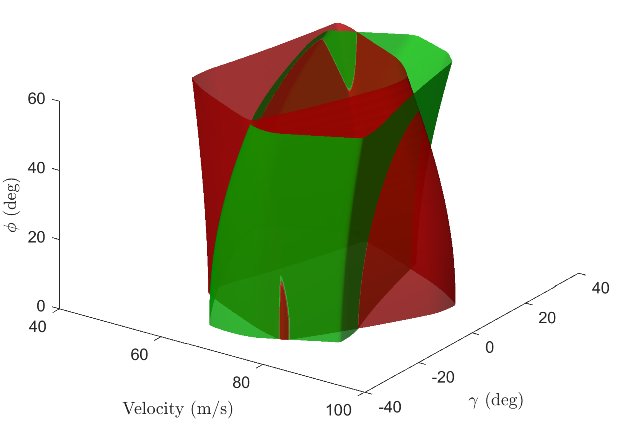

Fig. 6. Forwards and Backwards reachable set for the simplified aircraft model.

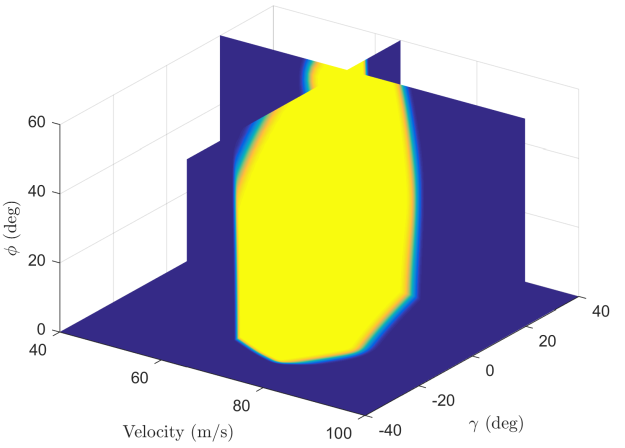

Fig. 7. Product of the forwards and backwards reachable set.

6.4 Sensitivity

The objective of this paper is to investigate the effect of disturbances on the SFE. By comparing the deterministic SFE with the confidence contours of the probabilistic SFE, it is possible to investigate the sensitivity. Since it is hard to compare the 3D reachable sets side-by-side, two slices have been made at roll angles of 0° and 60°.

Fig. 8 and Fig. 9 show both the deterministic SFE, as well as the probabilistic SFE. For the probabilistic SFE, the one, two and three σ isocontours are also plotted. The one, two and three σ isocontours represent the states in which a transition probability is achieved of at least 68.3 %, 95.5 % and 99.7 % respectively. Fig. 8 shows the results for a roll angle of 0°. The outer contour represents the deterministic SFE. It can be seen that the $1\,\sigma$ contour, which has a probability of 68.3 % is roughly the same size (slightly smaller) than the deterministic contour. This means that when the deterministic SFE is used for the purpose of LOC prevention in a setting with turbulence, there are some states that have a 31.7 % probability that the aircraft cannot return to the initial set. The probabilistic SFE can be defined at an arbitrary probability as chosen by the operator to ensure safety of the aircraft and occupants. The $3\,\sigma$ contour is more suitable for preventing LOC as it gives a high probability of returning to the initial set, even in the presence of atmospheric turbulence. The difference between the deterministic and probabilistic SFE becomes more pronounced at higher roll angles. Fig. 9 shows the same information but now for a roll angle of 60°. It can be seen that the effect of disturbances on the SFE is greater. A fairly large part of the deterministic SFE has a transition probability of 0 %, which means that it is impossible for an aircraft which enters this part of the SFE to return to the initial set.

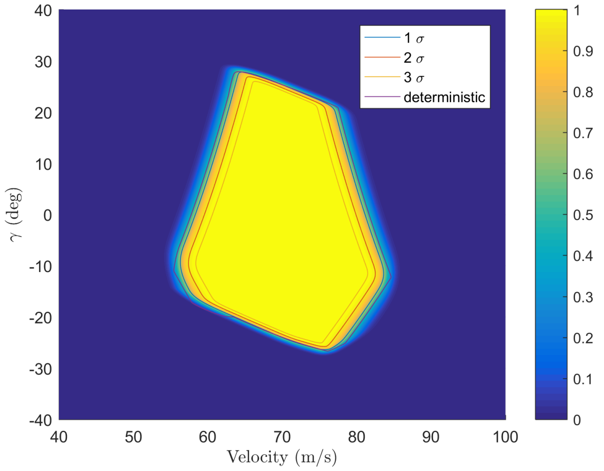

Fig. 8. SFE for a roll angle of 0°

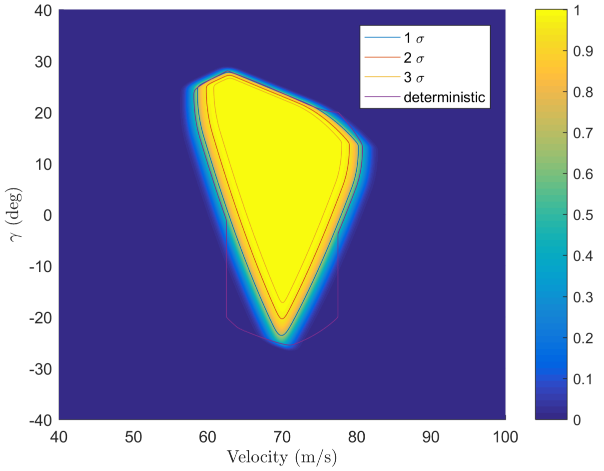

Fig. 9. SFE for a roll angle of 60°

Table 1. Size comparison between the deterministic and the probabilistic SFE.

Roll Angle/area	Deterministic	1σ	2σ	3σ
0°	100 %	94.3 %	81.9 %	71.4 %
30°	100 %	94.9 %	82.8 %	72.6 %
60°	100 %	78.0 %	62.8 %	50.8 %

To get some more insight in the size difference of the deterministic compared to the probabilistic SFE, the relative area between the deterministic set and the sigma contours is shown in Table 1. From this table it becomes clear that the disturbances have a greater effect on the SFE when the roll angle increases. For example, at a 60° roll angle, only 50 % of the deterministic SFE is enclosed by the $3\,\sigma$ contour.

7. VALIDATION

Generally, a Monte-Carlo method is employed to validate the results of the deterministic reachability analysis. This validation consists of a large number of individual trajectories which initial condition are located inside the initial set, for a number of (semi-) random control sequences. The reachable set is then approximated by the boundary of all trajectories. However, this approach is not suitable for application to the probabilistic reachability analysis. An alternative is presented as the Path Integral Monte Carlo (PIMC) method. The PIMC allows a step-by-step capture of the entire evolution of the response process in terms of PDF, starting from a known initial condition (deterministic or stochastic). The PIMC method has been used for a long time in physics to numerically solve the Fokker-Planck equation.

The starting point of the PIMC method is the Chapman-Kolmogorov equation (Pirrotta and Santoro (2011))

$$p_x(x, t+\tau) = \int_D p_x(x, t+\tau|\bar{x}, t)\, p_x(\bar{x}, t)\, d\bar{x} \qquad (28)$$

When the PDF is known at the current time step t, it is possible to evaluate the PDF after a small time step τ which holds true, because of the Markovian property of the response.

Fig. 10 illustrates the numerical solution to the Chapman-Kolmogorov equation for a single point. The blue line represents the (Gaussian) distribution of the SDE at the current time t. This distribution is denoted by $p_x(x, t)$. The PIMC method samples the SDE over the domain D. Fig. 10 shows these samples for a single point denoted by $\bar{x}$. The gray samples are all unique realizations of the SFE starting from $\bar{x}$ over a time horizon τ. The PDF corresponding to these samples is shown in cyan and is denoted $p_x(x, t+\tau|\bar{x}, t)$. The PDF belonging to $\bar{x}$ ($p_x(x, t+\tau|\bar{x}, t)$) is then weighted by the previous probability at $p(\bar{x}, t)$. The PDF at the next time instant $(t+\tau)$ is shown in red. This PDF is the ensemble of all local Gaussian probability distributions over the domain D.

To check that the aircraft model was correctly implemented in the LS toolbox, a validation is performed for the simplified aircraft model. In this validation the backwards reachability for two cases is considered, a roll angle of 0° and 60°. The difference between the LS and PIMC method is shown in Fig. 11 and Fig. 12 for a roll angle of 0° and 60° respectively. The reachability is calculated over a time horizon of 2 s, starting from an initial set as represented by the black rectangle.

From the validation it becomes clear that the LS solution is similar to the PIMC method. When looking at the magnitude of the transition probability, a slight difference can be seen in the corners of the PIMC method, However, the PIMC solution always produces a similar or lower estimation of the transition probability, as expected.

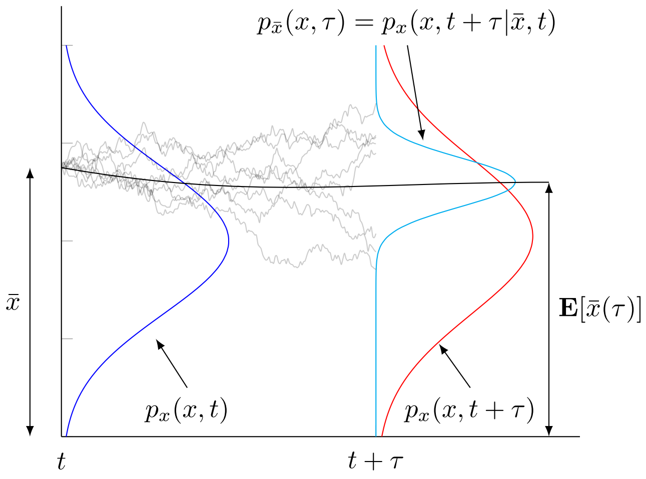

Fig. 10. Path Integral Monte Carlo method.

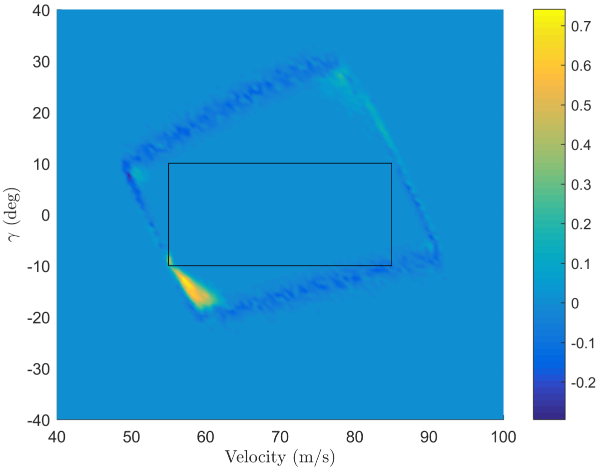

Fig. 11. Difference between the PIMC and the LS solution for $\phi = 0°$.

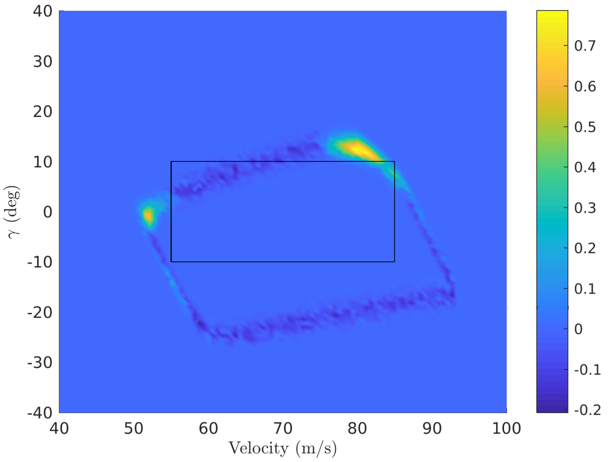

Fig. 12. Difference between the PIMC and the LS solution $\phi = 60°$.

8. CONCLUSION

In this paper, the reachability analysis which is commonly used to computed the Safe Flight Envelope (SFE), is extended to quantify the reachability probability of stochastic processes. This set is referred to as the probabilistic reachable set and describes the transition probability of a Stochastic Differential Equation (SDE) to a target set. The conditional probability between the forwards and backwards probabilistic reachable set then yields the probabilistic SFE. The difference between the deterministic and probabilistic SFE is analyzed for a simplified aircraft model.

When comparing the deterministic SFE with the probabilistic counterpart, it becomes clear that the deterministic envelope roughly corresponds to the $1\,\sigma$ isocontour. This means that for most states in the deterministic SFE there is a 68.3 % or higher probability of returning to the initial set. In most scenarios, a $1\,\sigma$ reachability probability is not acceptable. A better suited definition would be the $3\,\sigma$ contour. This contour contains all states for which the transition probability is at least 99.7 %. The deterministic SFE is significantly larger than the $3\,\sigma$ contour. When comparing the area of both envelopes at different roll angles it can be observed that the stochastic envelope ranges from 71.4 % till 50.8 % of the area. This means that at large roll angles the SFE is almost half that of the case without disturbances. Furthermore, for some states in the deterministic SFE, the transition probability under influence of external turbulence is 0 %, meaning that there is no possibility of returning to the initial set.

Finally, validation of the probabilistic reachable set is performed by means of a Path Integral Monte Carlo method. When looking at the magnitude of the transition probability, a slight difference can be seen in the corners of the Path Integral Solution. However, the Path Integral Monte Carlo method always produces a similar or lower estimation of the transition probability, as expected.

REFERENCES

(1990). *Flying Qualities of Piloted Vehicles.* Department of Defense. MIL-STD-1797A.

Assellaou, M., Bokanowski, O., and Zidani, H. (2015). Error estimates for second order hamilton-jacobi-bellman equations. approximation of probabilistic reachable sets. *Discrete and Continuous Dynamical Systems-Series A (DCDS-A)*, 35(9), 3933–3964.

Boeing (2017). Statistical summary of commercial jet airplane accidents worldwide operations 1959–2016. Technical report, Aviation Safety Boeing Commercial Airplanes, Seattle, Washington 98124-2207, USA.

Bressan, A. (2011). Viscosity solutions of hamilton-jacobi equations and optimal control problems. *Lecture notes.*

Doya, K. (2007). *Bayesian brain: Probabilistic approaches to neural coding.* MIT press.

Esfahani, P.M., Chatterjee, D., and Lygeros, J. (2016). The stochastic reach-avoid problem and set characterization for diffusions. *Automatica*, 70, 43–56.

ICAO (2017). 2017 safety report. Technical report, International Civil Aviation Organization, Quebec, Montral 999 Boulevard Robert-Bourassa, Canada.

Kappen, H.J. et al. (2011). Optimal control theory and the linear bellman equation. *Inference and Learning in Dynamic Models*, 363–387.

Lambregts, A., Nesemeier, G., Wilborn, J., and Newman, R. (2008). Airplane upsets: Old problem, new issues. In *AIAA Modeling and Simulation Technologies Conference and Exhibit*, 6867.

Lombaerts, T., Schuet, S., Wheeler, K., Acosta, D.M., and Kaneshige, J. (2013). Safe maneuvering envelope estimation based on a physical approach. In *AIAA Guidance, Navigation, and Control (GNC) Conference*, 4618.

Looye, G. and Bennani, S. (1997). Description and analysis of the research civil aircraft model (rcam). *Technical publication TP-088-27, Group for Aeronautical Research and Technology in EURope (GARTEUR).*

Lygeros, J. (2004). On reachability and minimum cost optimal control. *Automatica*, 40(6), 917–927.

Mitchell, I.M. (2008). The flexible, extensible and efficient toolbox of level set methods. *Journal of Scientific Computing*, 35(2), 300–329.

Mitchell, I.M., Bayen, A.M., and Tomlin, C.J. (2005). A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on automatic control*, 50(7), 947–957.

Nabi, H. (2016). *Effects of Structural Failures on the Safe Flight Envelope of Aircraft.* Master's thesis, TU Delft, Delft University of Technology.

Pirrotta, A. and Santoro, R. (2011). Probabilistic response of nonlinear systems under combined normal and poisson white noise via path integral method. *Probabilistic Engineering Mechanics*, 26(1), 26–32.

Van Oort, E.R. (2011). *Adaptive backstepping control and safety analysis for modern fighter aircraft.* TU Delft, Delft University of Technology.

van Oort, E., Chu, Q., and Mulder, J.A. (2011). Maneuver envelope determination through reachability analysis. In *Advances in Aerospace Guidance, Navigation and Control*, 91–102. Springer.

Zhang, Y., de Visser, C.C., and Chu, Q.P. (2016). Online safe flight envelope prediction for damaged aircraft: A database-driven approach. In *AIAA Modeling and Simulation Technologies Conference*, 1189.

Bibliography

[1] Boeing. Statistical summary of commercial jet airplane accidents worldwide operations 1959–2016. Technical report, Aviation Safety Boeing Commercial Airplanes, Seattle, Washington 98124-2207, USA, 2017.

[2] ICAO. 2017 safety report. Technical report, International Civil Aviation Organization, Quebec, Montréal 999 Boulevard Robert-Bourassa, Canada, 2017.

[3] AA Lambregts, Gregg Nesemeier, JE Wilborn, and RL Newman. Airplane upsets: Old problem, new issues. In *AIAA Modeling and Simulation Technologies Conference and Exhibit*, page 6867, 2008.

[4] Elwin De Weerdt, Eddy van Oort, Erik-Jan Van Kampen, Ping Chu, and Jan Mulder. *Level Set Method Based on Interval Analysis*. American Institute of Aeronautics and Astronautics (AIAA), 2011.

[5] Jork Stapel, Coen C de Visser, Erik-Jan Van Kampen, and Q Ping Chu. Efficient methods for flight envelope estimation through reachability analysis. In *AIAA Guidance, Navigation, and Control Conference*, page 0083, 2016.

[6] Ye Zhang, Coen C de Visser, and Q Ping Chu. Online safe flight envelope prediction for damaged aircraft: A database-driven approach. In *AIAA Modeling and Simulation Technologies Conference*, page 1189, 2016.

[7] Roel Helsen, Erik-Jan Van Kampen, Cornelis C de Visser, and Qiping Chu. Distance-fields-over-grids method for aircraft envelope determination. *Journal of Guidance, Control, and Dynamics*, pages 1–11, 2016.

[8] Thomas Lombaerts, Stefan Schuet, Kevin Wheeler, Diana M Acosta, and John Kaneshige. Safe maneuvering envelope estimation based on a physical approach. In *AIAA Guidance, Navigation, and Control (GNC) Conference*, page 4618, 2013.

[9] Thomas Lombaerts, Stefan Schuet, Kevin Wheeler, Diana Acosta, and John Kaneshige. Robust maneuvering envelope estimation based on reachability analysis in an optimal control formulation. In *2013 Conference on Control and Fault-Tolerant Systems (SysTol)*, pages 318–323. IEEE, 2013.

[10] Eduard Richard Van Oort. *Adaptive backstepping control and safety analysis for modern fighter aircraft.* TU Delft, Delft University of Technology, 2011.

[11] John Lygeros. On reachability and minimum cost optimal control. *Automatica*, 40(6):917–927, 2004.

[12] ER van Oort, QP Chu, and Jan Albert Mulder. Maneuver envelope determination through reachability analysis. In *Advances in Aerospace Guidance, Navigation and Control*, pages 91–102. Springer, 2011.

[13] H.N. Nabi. Effects of structural failures on the safe flight envelope of aircraft. Master's thesis, TU Delft, Delft University of Technology, 2016.

[14] Ian M Mitchell, Alexandre M Bayen, and Claire J Tomlin. A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on automatic control*, 50(7):947–957, 2005.

[15] Richard Bellman. A markovian decision process. Technical report, DTIC Document, 1957.

[16] Alberto Bressan. Viscosity solutions of hamilton-jacobi equations and optimal control problems. *Lecture notes*, 2011.

[17] A. Adardour. Nonlinear reachability analysis using multivariate b-splines. Master's thesis, TU Delft, Delft University of Technology, 2016.

[18] Stanley Osher and Ronald Fedkiw. *Level set methods and dynamic implicit surfaces*, volume 153. Springer Science & Business Media, 2006.

[19] M Athans and PL Falb. Optimal control. 1966. *McGraw- Iliil.*

[20] Ian M Mitchell. The flexible, extensible and efficient toolbox of level set methods. *Journal of Scientific Computing*, 35(2):300–329, 2008.

[21] Desmond J Higham. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM review*, 43(3):525–546, 2001.

[22] Bernt Oksendal. *Stochastic differential equations: an introduction with applications.* Springer Science & Business Media, 2013.

[23] CW Gardiner. Handbook of stochastic methods for physics, chemistry and the natural sciences. *Applied Optics*, 25:3145, 1986.

[24] Hannes Risken. Fokker-planck equation. In *The Fokker-Planck Equation*, pages 63–95. Springer, 1996.

[25] Lawrence C Evans. *An introduction to stochastic differential equations*, volume 82. American Mathematical Soc., 2012.

[26] Steven R Dunbar. Stochastic processes and advanced mathematical finance. *Department of Mathematics, University of Nebraska-Lincoln, USA*, 2016.

[27] Mohamed Assellaou, Olivier Bokanowski, and Hasnaa Zidani. Error estimates for second order hamilton-jacobi-bellman equations. approximation of probabilistic reachable sets. *Discrete and Continuous Dynamical Systems-Series A (DCDS-A)*, 35(9):3933–3964, 2015.

[28] Hilbert J Kappen et al. Optimal control theory and the linear bellman equation. *Inference and Learning in Dynamic Models*, pages 363–387, 2011.

[29] Kenji Doya. *Bayesian brain: Probabilistic approaches to neural coding.* MIT press, 2007.

[30] Peyman Mohajerin Esfahani, Debasish Chatterjee, and John Lygeros. The stochastic reach-avoid problem and set characterization for diffusions. *Automatica*, 70:43–56, 2016.

[31] A Pirrotta and R Santoro. Probabilistic response of nonlinear systems under combined normal and poisson white noise via path integral method. *Probabilistic Engineering Mechanics*, 26(1):26–32, 2011.

[32] Gertjan Looye and Samir Bennani. Description and analysis of the research civil aircraft model (rcam). *Technical publication TP-088-27, Group for Aeronautical Research and Technology in EURope (GARTEUR)*, 1997.

[33] TR Beal. Digital simulation of atmospheric turbulence for dryden and von karman models. *Journal of Guidance Control and Dynamics*, 16:132–132, 1993.

[34] *Flying Qualities of Piloted Vehicles.* Department of Defense, 1990. MIL-STD-1797A.

[35] S. Sun, R.J. Schilder, and C.C. de Visser. Identification of quadrotor aerodynamic model from high speed flight data. *AIAA Guidance, Navigation, and Control Conference*, in prep.

[36] I Kitsios and J Lygeros. Final glide-back envelope computation for reusable launch vehicle using reachability. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 4059–4064. IEEE, 2005.

[37] Simon J Julier. The scaled unscented transformation. In *American Control Conference, 2002. Proceedings of the 2002*, volume 6, pages 4555–4559. IEEE, 2002.

[38] Ian Mitchell and Claire Tomlin. Level set methods for computation in hybrid systems. In *HSCC*, volume 1790, pages 310–323. Springer, 2000.

[39] Ian Michael Mitchell. Application of level set methods to control and reachability problems in continuous and hybrid systems. 2003.

[40] Ian M Mitchell. A toolbox of level set methods. *Dept. Comput. Sci., Univ. British Columbia, Vancouver, BC, Canada, http://www. cs. ubc. ca/˜ mitchell/ToolboxLS/toolboxLS. pdf, Tech. Rep. TR-2004-09*, 2004.